

Isolation schemes for problems on decomposable graphs

Jesper Nederlof^{*} Michał Pilipczuk[†] Céline M. F. Swennenhuis[‡] Karol Węgrzycki[§]

Abstract

The Isolation Lemma of Mulmuley, Vazirani and Vazirani [Combinatorica'87] provides a self-reduction scheme that allows one to assume that a given instance of a problem has a unique solution, provided a solution exists at all. Since its introduction, much effort has been dedicated towards derandomization of the Isolation Lemma for specific classes of problems. So far, the focus was mainly on problems solvable in polynomial time.

In this paper, we study a setting that is more typical for NP-complete problems, and obtain partial derandomizations in the form of significantly decreasing the number of required random bits. In particular, motivated by the advances in parameterized algorithms, we focus on problems on decomposable graphs. For example, for the problem of detecting a Hamiltonian cycle, we build upon the rank-based approach from [Bodlaender et al., Inf. Comput.'15] and design isolation schemes that use

- $\mathcal{O}(t \log n + \log^2 n)$ random bits on graphs of treewidth at most t ;
- $\mathcal{O}(\sqrt{n})$ random bits on planar or H -minor free graphs; and
- $\mathcal{O}(n)$ -random bits on general graphs.

In all these schemes, the weights are bounded exponentially in the number of random bits used. As a corollary, for every fixed H we obtain an algorithm for detecting a Hamiltonian cycle in an H -minor-free graph that runs in deterministic time $2^{\mathcal{O}(\sqrt{n})}$ and uses polynomial space; this is the first algorithm to achieve such complexity guarantees. For problems of more local nature, such as finding an independent set of maximum size, we obtain isolation schemes on graphs of treedepth at most d that use $\mathcal{O}(d)$ random bits and assign polynomially-bounded weights.

We also complement our findings with several unconditional and conditional lower bounds, which show that many of the results cannot be significantly improved.

^{*}Utrecht University, The Netherlands, j.nederlof@uu.nl. 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 No 853234).

[†]Institute of Informatics, University of Warsaw, Poland, michal.pilipczuk@mimuw.edu.pl. This work is a part of the project TOTAL that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 677651).

[‡]Eindhoven University of Technology, The Netherlands, c.m.f.swennenhuis@tue.nl. Supported by the Netherlands Organization for Scientific Research under project no. 613.009.031b.

[§]Saarland University and Max Planck Institute for Informatics, Saarbrücken, Germany, wegrzycki@cs.uni-saarland.de. This 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 No 850979).



European Research Council
Established by the European Commission



1 Introduction

Isolation is a procedure that allows to single out a unique solution to a given problem within a possibly larger solution space, thus effectively reducing the original problem to a variant where one may assume that if a solution exists, then there is a unique one. The classic Isolation Lemma of Mulmuley, Vazirani and Vazirani [43] can be used to achieve this at the cost of allowing randomization. In complexity theory, isolation is used to show that hard problems are not easier to solve on instances with unique solutions [54]. This idea has found numerous applications ranging from structural results in complexity theory (e.g. $\text{NL/poly} \subseteq \oplus\text{L/poly}$ [57] or $\text{NL/poly} = \text{UL/poly}$ [49]) to the design of parallel algorithms [43, 32, 25, 51].

Since obtaining a general derandomization of the Isolation Lemma is impossible by counting arguments [6, 11, 1], it is natural to ask whether the isolation step can be derandomized for specific problems with explicit representation. In this context, there has recently been an exciting progress in isolation for perfect matchings [2, 9, 18, 31, 5, 32], which culminated in an isolation scheme that uses $\mathcal{O}(\log^3 n)$ random bits, implying a quasi-NC algorithm for detecting a perfect matching [51].

In contrast to this, derandomization of isolation procedures for NP-complete problems is relatively less studied, and not because of a lack of motivation: Many contemporary fixed-parameter algorithms rely on the Isolation Lemma [39, 44, 7, 34, 35, 16, 58]. Usually, the isolation procedure is the only subroutine requiring randomness. Many of the algorithms mentioned above apply the Isolation Lemma in combination with a decomposition-based method such as Divide&Conquer or dynamic programming. This motivates us to study the following:

Main Question. *How much randomness is required for isolating problems with decomposable structure?*

More concretely, we focus on graph problems where the underlying graph is *decomposable*, in the sense that it can be decomposed using small separators. Examples of such graphs are planar graphs or graphs of bounded treewidth. It is well-known that for many NP-complete problems, the nice structure of such graphs can be leveraged to solve these problems faster than in general graphs. We show that a similar phenomenon occurs when one considers the amount of randomness needed to isolate a single solution.

The model for isolation schemes. Suppose U is a finite set and $\omega: U \rightarrow \mathbb{N}$ is a weight function. For $X \subseteq U$ we write $\omega(X) := \sum_{e \in X} \omega(e)$. For a set family $\mathcal{F} \subseteq 2^U$ we say that ω *isolates* \mathcal{F} if there is exactly one set $S \in \mathcal{F}$ such that $\omega(S)$ is the minimum possible among the weights of the sets in \mathcal{F} . The classic Isolation Lemma of Mulmuley et al. [43] states that a weight function $\omega: U \rightarrow \{1, \dots, 2|U|\}$ chosen uniformly at random isolates any family $\mathcal{F} \subseteq 2^U$ with probability at least $\frac{1}{2}$. Note that sampling such ω requires $\mathcal{O}(|U| \log |U|)$ random bits.

Most of our isolation schemes work in a very restricted model inspired by the discussion above, which we explain now. Intuitively, the scheme is not aware of the graph or its decomposition, but is only aware of the vertex count of the graph and the relevant width parameter, such as the treewidth or treedepth.

Formally, a *vertex selection problem* is a function \mathcal{P} that maps every graph G to a family $\mathcal{P}(G) \subseteq 2^{V(G)}$ consisting of subsets of the vertex set of G . Edge selection problems are defined analogously: $\mathcal{P}(G)$ consists of subsets of $E(G)$. For example, we could define a vertex selection problem $\text{MIS}(\cdot)$ that maps every graph G to the family $\text{MIS}(G)$ comprising all maximum-size independent sets in G , or an edge selection problem $\text{HC}(\cdot)$ that maps every graph G to the family $\text{HC}(G)$ comprising all (edge sets of) Hamiltonian cycles in G . Further, let \mathcal{C} be a class of graphs, that is, a set of graphs that is invariant under isomorphism. For instance, \mathcal{C} could be the class of planar graphs, or the class of graphs of treewidth at most k , for any fixed k . Then our definition of an isolation scheme reads as follows (here, we write $[n] := \{1, \dots, n\}$):

Definition 1.1. *For a graph class \mathcal{C} , we say that a vertex selection problem \mathcal{P} admits an isolation scheme on \mathcal{C} if for every $n \in \mathbb{N}$ there exist weight functions $\omega_1, \dots, \omega_\ell: [n] \rightarrow \mathbb{N}$ such that for every $G \in \mathcal{C}$ with vertex set $[n]$, ω_i isolates $\mathcal{P}(G)$ for at least half of the indices $i \in [\ell]$.*

Isolation schemes for edge selection problems are defined analogously: the weight functions $\omega_1, \dots, \omega_\ell$ have domain $[m]$ and should assign weights to all the edges in m -edge graphs in \mathcal{C} , where the edges are assumed to be enumerated with numbers in $[m]$.

The two main parameters of interest for isolation schemes will be the number of *random bits*, which is defined as $\log \ell$, and the *maximum weight*, defined as the maximum value that any of the functions ω_i may take. Although Definition 1.1 only assumes the *existence* of suitable weight functions, all the isolation schemes proposed in this paper are extremely simple and can be used as an effective derandomization tool.

1.1 Our contribution

In the following discussion we restrict attention to Hamiltonian cycles and maximum-size independent sets for concreteness, that is, to the edge- and vertex-selection problems $\text{HC}(\cdot)$ and $\text{MIS}(\cdot)$ described above. However, our techniques have a wider applicability, which we comment on throughout the presentation. On a very high level, the natural idea that permeates all our arguments is to reduce the randomness using Divide&Conquer along small separators: If a separator X splits the given graph G in a balanced way, then the same random bits can be reused in each part of $G - X$.

Isolation schemes for Hamiltonian cycles. We first consider the problem of detecting a Hamiltonian cycle, since it represents an important class of connectivity problems such as STEINER TREE or k -PATH. For these problems, the Isolation Lemma has been particularly useful in the design of parameterized algorithms [39, 44, 7, 34, 35, 16, 58]. Our first result concerns general graphs.

Theorem 1.2. *There is an isolation scheme for Hamiltonian cycles in undirected graphs that uses $\mathcal{O}(n)$ random bits and assigns weights upper bounded by $2^{\mathcal{O}(n)}$.*

Observe that in an n -vertex graph there can be as many as $n!$ different Hamiltonian cycles. Hence, the application of the general-usage isolation scheme of Chari et al. [11] would give an isolation scheme for Hamiltonian cycles in general graphs that uses $\mathcal{O}(\log(n!)) = \mathcal{O}(n \log n)$ random bits. Note that as proved in [11], isolating a family \mathcal{F} over a universe of size n requires $\Omega(\log |\mathcal{F}| + \log n)$ random bits in general, hence the shaving of the $\log n$ factor reported in Theorem 1.2 required a problem-specific insight into the family of Hamiltonian cycles in a graph. This insight is provided by the *rank-based approach*, a technique introduced in the context of detecting Hamiltonian cycles in graphs of bounded treewidth [8]. The fact that this works is unexpected because all known methods for derandomizing Hamiltonian cycle require at least exponential space (see [8] for overview).

Let us note that isolation of Hamiltonian cycles was used by Björklund [7] in his $\mathcal{O}(1.657^n)$ -time algorithm for detecting a Hamiltonian cycle in an undirected graph. This algorithm is randomized due to the usage of the Isolation Lemma, and derandomizing it, even within time complexity $\mathcal{O}((2 - \varepsilon)^n)$ for any $\varepsilon > 0$, is a major open problem. While the constant hidden in the $\mathcal{O}(\cdot)$ notation used in Theorem 1.2 is too large to allow exploring the whole space of random bits within time $\mathcal{O}((2 - \varepsilon)^n)$, in principle we show that the amount of randomness needed is of the same magnitude as would be required for an efficient derandomization of the algorithm of Björklund.

Next, we show that in the setting of graphs of bounded treewidth the amount of randomness can be reduced dramatically, to a polylogarithm in n .

Theorem 1.3. *For every $t \in \mathbb{N}$, there is an isolation scheme for Hamiltonian cycles in graphs of treewidth at most t that uses $\mathcal{O}(t \log n + \log^2(n))$ random bits and assigns weights upper bounded by $2^{\mathcal{O}(t \log n + \log^2 n)}$.*

The proof of Theorem 1.3 fully exploits the idea of using small separators to save on randomness. It also uses the rank-based approach to shave off a $\log t$ factor in the number of random bits.

Finally, we use the separator properties of H -minor free graphs to prove the following.

Theorem 1.4. *For every fixed H , there is an isolation scheme for Hamiltonian cycles in H -minor-free graphs that uses $\mathcal{O}(\sqrt{n})$ random bits and assigns weights upper bounded by $2^{\mathcal{O}(\sqrt{n})}$.*

Recently, in [44] the authors presented a randomized algorithm for detecting a Hamiltonian cycle in a graph of treedepth at most d that works in time $2^{\mathcal{O}(d)} \cdot (W + n)^{\mathcal{O}(1)}$ time and uses polynomial space; here, W is the maximum weight assigned by isolation scheme¹. The only source of randomness in the algorithm of [44] is the Isolation Lemma. Since H -minor free graphs have treedepth $\mathcal{O}(\sqrt{n})$, we can use the isolation scheme of Theorem 1.4 to derandomize this algorithm, thus obtaining the following result.

Theorem 1.5. *For every fixed H , there is a deterministic algorithm for detecting a Hamiltonian cycle in an H -minor-free graph that runs in time $2^{\mathcal{O}(\sqrt{n})}$ and uses polynomial space.*

To the best of our knowledge, this is the first application of a randomness-efficient isolation scheme for a full derandomization of an exponential-time algorithm without a significant loss on complexity guarantees. Further, we are not aware of any previous algorithms that would simultaneously achieve determinism, running time $2^{\mathcal{O}(\sqrt{n})}$, and polynomial space complexity, even in the setting of planar graphs². Finally, let us note that the algorithm of Theorem 1.5 does not rely on any topological properties of H -minor-free graphs: the existence of balanced separators of size $\mathcal{O}(\sqrt{n})$ is the only property we use.

MSO-definable problems on graphs of bounded treewidth. We observe that the approach used in the proof of Theorem 1.3 relies only on finite-state properties of the HAMILTONIAN CYCLE problem on graphs of bounded treewidth. The range of problems enjoying such properties is much wider and encompasses all problems definable in CMSO_2 : the Monadic Second-Order logic with modular counting predicates. Consequently, we can lift the proof of Theorem 1.3 to a generic reasoning that yields an analogous result for every CMSO_2 -definable problem. This proves the following (see Section 6 for definitions).

Theorem 1.6. *Let \mathcal{P} be a CMSO_2 -definable edge selection problem. There exists a computable function f such that for every $k \in \mathbb{N}$, \mathcal{P} admits an isolation scheme on graphs of treewidth at most k that uses $R := f(k) \cdot \log n + \mathcal{O}(\log^2 n)$ random bits and assigns weights upper bounded by 2^R .*

Lower bounds. We show that a significant improvement of the parameters in the isolation schemes presented above is unlikely. First, a counting argument shows that the $\log n$ factor is necessary.

Theorem 1.7. *There does not exist an isolation scheme for Hamiltonian cycles on graphs of treewidth at most 4 that uses $o(\log n)$ random bits and polynomially bounded weights.*

Using similar constructions we also provide analogous $\Omega(\log n)$ lower bounds for isolating other families of combinatorial objects related to NP-hard problems, such as maximum independent sets, minimum Steiner trees, and minimum maximal matchings. These lower bounds hold even in graphs of bounded treedepth, which is a more restrictive setting than bounded treewidth.

We also show using existing reductions that a significant improvement over the scheme of Theorem 1.2 would imply a surprising partial derandomization of isolation schemes for SAT.

Theorem 1.8. *Suppose there is an isolation scheme for Hamiltonian cycles in undirected graphs that uses $o(n)$ random bits and polynomially bounded weights. Then there is a randomized polynomial-time reduction from SAT to UNIQUE SAT that uses $o(n)$ random bits, where n is the number of variables.*

¹They did not consider the weighted case, but the statement is implied by a standard extension, see Section 5 for details.

²Deterministic $2^{\mathcal{O}(\sqrt{n})}$ -time algorithms were previously known, but all of these use exponential space [8, 26].

Observe that since an n -vertex graph has treewidth at most $n - 1$, Theorem 1.8 also implies that in Theorem 1.3 one cannot expect reducing the number of random bits to $o(t)$. However, we stress that the lower bounds of Theorems 1.7 and 1.8 are not completely tight with respect to the upper bounds of Theorems 1.2 and 1.3, because the latter allow superpolynomial weights. It remains open whether the weights used by the schemes of Theorems 1.2, 1.3, and 1.4 can be reduced to polynomial.

In Section 7 we further discuss consequences of the hypothetical existence of a polynomial-time reduction from SAT to UNIQUE SAT that would use $o(n)$ random bits.

Level-aware isolation schemes for independent sets. In the light of the $\Omega(\log n)$ lower bound of Theorem 1.7, we consider a relaxation of the model from Definition 1.1, where the graph is provided together with an *elimination forest* (a decomposition notion suited for the graph parameter *treedepth*), and the weight of a vertex may depend both on the vertex' identifier and its level in the elimination forest. We demonstrate that in this relaxed model, the $\Omega(\log n)$ lower bound can be circumvented.

Definition 1.9. We say that vertex selection problem \mathcal{P} admits a level-aware isolation scheme if for all $n, d \in \mathbb{N}$ there exist functions $\omega_1, \dots, \omega_\ell: [n] \times [d] \rightarrow \mathbb{N}$ such that for every graph G on vertex set $[n]$ and elimination forest F of G of height at most d , at least half of the functions $\omega_1, \dots, \omega_\ell$ isolate $\mathcal{P}(G)$. Here, when evaluating ω_i on a vertex $u \in [n]$, we apply ω_i to u and the index of the level of u in F .

Theorem 1.10. For every $d \in \mathbb{N}$, there is a level-aware isolation scheme for maximum-size independent sets in graphs of treedepth at most d that uses $\mathcal{O}(d)$ random bits and assigns weights bounded by $\mathcal{O}(n^6)$.

In the proof of Theorem 1.10 we describe an abstract condition, dubbed the *exchange property*, which is sufficient for the argument to go through. This property is enjoyed also by other families of combinatorial objects defined through constraints of local nature, such as minimum dominating sets or minimum vertex covers. Therefore, we can prove analogous isolation results for those families as well.

Also, in Section 9 we discuss a similar reasoning for edge-selection problems on the example of maximum matchings, achieving a level-aware isolation scheme that uses $\mathcal{O}(d \log n)$ random bits and assigns weights bounded by $n^{\mathcal{O}(\log n)}$. This provides another natural class of graphs where isolation-based algorithms for finding a maximum matching can be derandomized (see [2, 9, 18, 31]).

We summarize our results with Table 1.

Table 1: Summary of our results based on Theorems 1.2-1.10.

Problem	Random Bits	Max Weight	Graph Class
HAMILTONIAN CYCLE	$\mathcal{O}(n)$	$2^{\mathcal{O}(n)}$	General Graphs
	$\Omega(n)$	$\text{poly}(n)$	
	$\mathcal{O}(\sqrt{n})$	$2^{\mathcal{O}(\sqrt{n})}$	H -minor free graphs
	$\Omega(\sqrt{n})$	$\text{poly}(n)$	
	$\mathcal{O}(t \log(n) + \log^2(n))$	$n^{\mathcal{O}(t + \log(n))}$	Treewidth t graphs
	$\Omega(t + \log(n))$	$\text{poly}(n)$	
CMSO ₂	$f(t) \log(n) + \mathcal{O}(\log^2(n))$	$n^{f(t) + \mathcal{O}(\log(n))}$	Treedepth t graphs
MAX INDEPENDENT SET	$\mathcal{O}(d)$	$\text{poly}(n)$	Treedepth d graphs
	$\Omega(d)$	$\text{poly}(n)$	




Figure 1: Top figure presents an example graph with bounded pathwidth. Figure below presents a definition of a path decomposition of width k , with bags S_a, S_b and interval $I_{a,b}$.

1.2 Organization

In Section 2 we introduce the main techniques behind our isolation schemes in an informal way. In Section 3 we provide preliminaries. Section 4 is dedicated to the formal proofs of Theorems 1.2, 1.3 and 1.4. The derandomized algorithm from Theorem 1.5 is subsequently proved in Section 5, and the general CMSO₂-result of Theorem 1.6 is formally supported in Section 6. The lower bounds from Theorem 1.7 and Theorem 1.8 are proved in 7. Finally, the level-aware isolation schemes for local vertex (respectively, edge) selection problems are given in Sections 8 (respectively, Section 9), and we finish the paper with possible directions for further research in Section 10.

2 An informal introduction to our techniques

In this section we present an isolation scheme for Hamiltonian cycle on graphs of bounded *pathwidth* at most k that uses $\mathcal{O}(k \log k \log n + \log^2 n)$ random bits. The arguments in this section are informal in order to convey the underlying intuition, and merely serve as a preliminary overview of the general framework that we formalize and further develop in the subsequent sections.

Throughout the paper we heavily build upon the approach proposed by Kallampally et al. [36], who showed an isolation scheme for shortest paths that uses $\mathcal{O}(\log^2(n))$ random bits assigns weights upper bounded by $n^{\mathcal{O}(\log n)}$ (see e.g., [32] for a more recent application). In fact, our isolation schemes are almost identical to Kallampally et al. [36] except for a different selection of prime numbers. Our contribution comes with the new insight for NP-complete problems. To achieve this we use a modern toolset from parameterized algorithms.

Let $G = (V, E)$ be the given graph. Informally, the pathwidth of G is parameter that measures how well G can be represented as a *thickened path*, which formalized through the notion of a *path decomposition* of *width* k . The reader is invited to think about a path decomposition of G of width k as a sequence of bags $S_1, \dots, S_q \subseteq V$, each of size at most k , that traverse the whole graph, i.e., $S_1 \cup \dots \cup S_q = V$ and subsequent bags differ by exactly one vertex $|S_i \triangle S_{i+1}| \leq 1$ (see Figure 1 for an example of a graph with bounded pathwidth and a schematic view of path decomposition). We may assume that $q \leq n$. Each bag is a separator in the sense that vertices present only in the bags to the left of it are pairwise non-adjacent to the vertices present only in the bags to the right of it. In this section we focus on pathwidth in order to avoid several technical difficulties that arise when dealing with treewidth.

We first describe our isolation scheme for Hamiltonian cycles. The crucial ingredient in our methods is a well-known hashing scheme due to Fredman, Komlós and Szemerédi [27] (the FKS hashing lemma, see Section 3 for a proof): For any set A of n -bits integers with $|A| = n^{\mathcal{O}(1)}$, for most of the primes p of order $|A|^{\mathcal{O}(1)}$ it holds that $x \not\equiv y \pmod p$ for all distinct $x, y \in A$. An important property that is guaranteed by this lemma is that after hashing modulo a prime p , every element of the set A is given a different value.

Isolation scheme. Assume without loss of generality that $\log n$ is an integer. Our isolation scheme for n -vertex graphs of pathwidth k reads as follows. Let $\text{id}: E(G) \rightarrow \{1, \dots, |E(G)|\}$ be any bijection that assigns to each edge $e \in E(G)$ its unique *identifier* $\text{id}(e)$. First we select the range $M := k^{\mathcal{O}(k)} \cdot n^{\mathcal{O}(1)}$ and $\log n$ random prime numbers $p_1, \dots, p_{\log n} \in \{1, \dots, M\}$. Note that we need $\mathcal{O}(k \log k \log n + \log^2 n)$ random bits to sample these prime numbers.

Next, we inductively define weights functions $\omega_1, \dots, \omega_{\log n}$ on $E(G)$ as follows:

- Set $\omega_1(e) := 2^{\text{id}(e)} \bmod p_1$ for all $e \in E(G)$.
- For each $e \in E(G)$ and $i = 1, \dots, \log n$, set

$$\omega_i(e) := Mn \cdot \omega_{i-1}(e) + \left(2^{\text{id}(e)} \bmod p_i\right).$$

Let $\omega := \omega_{\log n}$ and observe that ω assigns weights bounded by $2^{\mathcal{O}(k \log k + \log^2 n)}$. Note that the path decomposition (S_1, \dots, S_q) of G is not used at all in the isolation scheme. We will use it only in the analysis, that is, the proof that the sampled weight function ω isolates the family of Hamiltonian cycles in G with probability at least $\frac{1}{2}$.

Analysis. We first introduce the notion of an *interval* in a path decomposition. This is just a graph induced by all the bags present between two given ones. More precisely, for $1 \leq a \leq b \leq q$, we define

$$I_{a,b} := \bigcup_{i \in [a,b]} S_i \subseteq V$$

to be the interval between bags S_a and S_b (see Figure 1). The length of this interval is $|b - a|$. For an interval $I_{a,b}$ we say that $P_{a,b} \subseteq E(I_{a,b})$ is a *partial solution* if it is a collection of vertex-disjoint paths with endpoints in $S_a \cup S_b$ that together visit all vertices of $I_{a,b}$. Note that if we want to extend $P_{a,b}$ to a Hamiltonian cycle with another edge set P' , in order to check the feasibility of this extension we only need to know the pattern of connections induced by $P_{a,b}$ on S_a and on S_b . More precisely, we only need to know the *configuration* of $P_{a,b}$ on its boundary: such a configuration is represented by a matching M on $S_a \cup S_b$, which indicates which vertices of $S_a \cup S_b$ are corresponding endpoints of a path in $P_{a,b}$, and information on how many edges of $P_{a,b}$ are incident on every vertex of $S_a \cup S_b$. Then $P_{a,b} \cup P'$ is a Hamiltonian cycle if and only if $P' \cup M$ is a Hamiltonian cycle on the vertices of $G - (I_{a+1,b-1} \cup V_2)$, where V_2 are vertices of $S_a \cup S_b$ incident on two edges of $P_{a,b}$.

For an example of a realization of a configuration on $I_{a,b}$, see Figure 2.

Since a configuration is composed of an information about a matching and a partition of vertices in $S_a \cup S_b$, the number of possible configurations within each interval $I_{a,b}$ is at most $2^{\mathcal{O}(k \log k)}$.

To prove that the weight function ω isolates the family of Hamiltonian cycles in G with high probability, we prove the following claim by induction on ℓ .

Induction hypothesis. For every $\ell \in \{1, \dots, \log n\}$, the following event happens with a sufficiently high probability: for every interval $I_{a,b}$ of length at most 2^ℓ and a configuration σ on $I_{a,b}$, the weight function ω_ℓ isolates the family of all partial solutions in $I_{a,b}$ whose configuration is σ .

This induction hypothesis for $\ell = \log n$ immediately shows that $\omega = \omega_{\log n}$ isolates all Hamiltonian cycles with a sufficiently high probability.

To prove the base case ($\ell = 1$), we look at intervals of length at most 2, that is, we look at the subgraphs $G[S_i]$ and $G[S_i \cup S_{i+1}]$, for $i \in \{1, \dots, q\}$. Each of these subgraphs has at most $2k$ vertices, hence also at most $k^{\mathcal{O}(k)}$ different partial solutions. Also, there are at most $2n$ such subgraphs in total. Hence, the total number of different partial solution in intervals of length 2 is at most $n \cdot k^{\mathcal{O}(k)}$. We find from the FKS hashing




Figure 2: The top figure represents a realization of configuration $\sigma_{a,b}$ in the interval $I_{a,b}$ (green stroked) and a complementary configuration for the rest of the graph (dashed lines). The figure below gives intuition for the induction argument. For a fixed $\sigma_{a,b}$ we know that $\omega_{\ell-1}$ isolates all realizations in $I_{a,c}$ and $I_{c,b}$ and there are at most $2^{\mathcal{O}(k \log k)}$ possible configurations $\sigma_{a,c}, \sigma_{b,c}$ that combined give $\sigma_{a,b}$.

lemma (see Lemma 3.1) that if we choose the prime p_1 uniformly at random from the interval $\{1, \dots, M\}$, with a high probability all those partial solutions will be assigned pairwise different weights by the weight function ω_1 . Indeed, take $\psi(e) = 2^{\text{id}(e)}$ and $W = \{\psi(X) : X \text{ a partial solution in interval of length at most } 2\}$. Note that all the $\psi(X)$ in W for different X are unique. Then FKS says that $\psi(X) \not\equiv \psi(X') \pmod{p_1}$ for any two such solutions X and X' with a high probability. This implies that with a high probability, $\omega_1(X) \neq \omega_1(X')$. The base case follows.

Now assume the induction hypothesis to be true for all $i < \ell$. Fix an interval $I_{a,b}$ for some $1 \leq a \leq b \leq n$ of length $\leq 2^\ell$. Observe that if $I_{a,b}$ has length at most $2^{\ell-1}$, then it is already appropriately taken care of by function $\omega_{\ell-1}$, and hence also by ω_ℓ .

Therefore, we can assume that the length of $I_{a,b}$ belongs to $[2^{\ell-1} + 1, 2^\ell]$. Fix some configuration $\sigma_{a,b}$ on $S_a \cup S_b$, the boundary of $I_{a,b}$. Observe that there exists $c \in (a, b)$, such that $I_{a,c} \cup I_{c,b} = I_{a,b}$ and both $I_{a,c}$ and $I_{c,b}$ have length at most $2^{\ell-1}$. Further, there are at most $2^{\mathcal{O}(k \log k)}$ different pairs of configurations $\sigma_{a,c}$ and $\sigma_{c,b}$ that, when naturally combined, give a configuration $\sigma_{a,b}$. See Figure 2 for a visualization.

The crucial observation is that by the induction hypothesis, for every pair of configurations $\sigma_{a,c}, \sigma_{c,b}$ as above, the weight function $\omega_{\ell-1}$ already isolates the family of partial solutions in $I_{a,c}$ with configuration $\sigma_{a,c}$, as well as the family of partial solutions in $I_{c,b}$ with configuration $\sigma_{c,b}$. Therefore, for a fixed interval $I_{a,b}$ there can be at most $2^{\mathcal{O}(k \log k)}$ different partial solution with configuration $\sigma_{a,b}$ that have minimum weight w.r.t. $\omega_{\ell-1}$. This is because they must be composed from partial solutions in $I_{a,c}$ and $I_{c,b}$ that have minimum weights for their configurations. Moreover, there are at most $\mathcal{O}(n^2)$ different intervals of length in $[2^{\ell-1} + 1, 2^\ell]$. This means that in total, there can be at most $2^{\mathcal{O}(k \log k)} \cdot n^2$ different partial solutions in intervals $I_{a,b}$ of length at most 2^ℓ that are minimum-weight realizations (w.r.t. $\omega_{\ell-1}$) of their respective configurations. Now moving from $\omega_{\ell-1}$ to ω_ℓ , we can argue using the FKS Lemma that all of these partial solutions will receive pairwise different values in ω_ℓ , with high probability.

This concludes the intuitive sketch of the proof of the induction hypothesis. For a formal argument, see Section 4.

Extensions of the method. All our isolation schemes for Hamiltonian cycle follow the same blueprint sketched above. The main difference, however, is that we select our primes to be of the order $2^{\mathcal{O}(k)} \cdot n^{\mathcal{O}(1)}$. To argue that this is sufficient, we employ the rank-based approach to argue that the set of partial solutions that are “representative enough” is much smaller than $2^{\mathcal{O}(k \log k)}$: it is actually of size $2^{\mathcal{O}(k)}$. In the above sketch, this reduces the number of random bits from $\mathcal{O}(k \log k \log n + \log^2 n)$ to $\mathcal{O}(k \log n + \log^2 n)$.

To complete the proof of Theorem 1.3 it remains to lift the reasoning from graphs of bounded path-width to graphs of bounded treewidth. We do this by carefully generalizing the notion of an interval in a path decomposition to a notion of a *segment* in a tree decomposition. In particular, a segment of a tree decomposition can be always partitioned into at most five segments of twice smaller sizes, similarly we partitioned an interval into two intervals of at most half the length.

If we directly applied our analysis to the problem of isolating Hamiltonian cycles in H -minor-free graphs, then even with the rank-based approach employed we would only obtain an isolation scheme that uses $\mathcal{O}(\sqrt{n} \log n)$ random bits. To shave off the additional $\mathcal{O}(\log n)$ factor, we use certain properties of the decompositions of H -minor-free graphs that guarantee that size of separator decreases geometrically. In a nutshell, these properties will allow us to select $\mathcal{O}(\log n)$ primes, but each prime will be selected using a number of random bits that follows a geometric progression (see Section 4.5 for details).

In Section 6 we generalize the ideas to prove a meta-statement about all problems definable in Monadic Second-Order logic, CMSO₂. The idea is that in the sketch above, we almost did not use any particular combinatorial properties of Hamiltonian cycles. The only property we relied on is that the behavior of a partial solution within an interval can be subsumed in a configuration on the interval's boundary, and the number of configurations is bounded by a function of k only. Such a “finite-state” property is enjoyed by all problems definable in CMSO₂, which allows us to perform the whole reasoning on the meta-level.

Methods presented in Section 8 for isolating local problems follow a completely different framework that uses additional information about a graph. The analysis in this section is arguably simpler. There, we use a technical contribution of Chari et al. [11] and extend it with new observations regarding pivotal vertices in treedepth bounded graphs.

3 Preliminaries

Notation. For an integer k , we write $[k] := \{1, \dots, k\}$. We use standard graph notation: $V(G)$ and $E(G)$ respectively denote the vertex set and the edge set of a graph G , for $X \subseteq V(G)$ the *closed neighborhood* $N_G[X]$ is X plus all the neighbors of vertices of X , and the *open neighborhood* is $N_G(X) := N_G[X] \setminus X$.

Hashing modulo primes. The following standard hashing lemma that dates back to the work of Fredman, Komlós, and Szemerédi [27], will be the main source of randomness in our isolation schemes.

Lemma 3.1 (FKS hashing lemma [27]). *Let $S \subseteq \{0, 1, \dots, 2^n\}$ be a set of k integers, where $n, k \geq 1$. Suppose that p is a prime number chosen uniformly at random among prime numbers in the range $\{1, \dots, M\}$, where $M \geq 2$. Then*

$$\mathbb{P}[x \not\equiv y \pmod{p} \text{ for all } x, y \in S, x \neq y] \geq 1 - \frac{nk^2}{\sqrt{M}}.$$

Proof. Let

$$R := \prod_{x, y \in S, x \neq y} |x - y|.$$

Note that $R \leq 2^{n \cdot \binom{k}{2}}$. This implies that R may have at most $n \cdot \binom{k}{2}$ different prime divisors. On the other hand, from the prime number theorem it follows that $\pi(M) \in \Omega\left(\frac{M}{\log M}\right)$, where $\pi(M)$ denotes the number of primes in the range $\{1, \dots, M\}$. In fact, using a more precise estimate of Rosser [50], for $M \geq 17$ we have $\pi(M) \geq \frac{M}{\ln M}$. For $2 \leq M \leq 17$ a direct check shows that $\pi(M) \geq \sqrt{M}/2$. Since $\frac{M}{\ln M} \geq \sqrt{M}/2$ for all $M \geq 2$, we conclude that the probability that a random prime in the range $\{1, \dots, M\}$ is not among the at most $n \cdot \binom{k}{2}$ prime divisors of R is at least

$$1 - \frac{n \cdot \binom{k}{2}}{\sqrt{M}/2} \geq 1 - \frac{nk^2}{\sqrt{M}}. \quad \square$$

Graph decompositions. A *rooted forest* is directed acyclic graph F where every node x has at most one outneighbor, called the *parent* of x . A *root* is a node with no parent. If a node y is reachable from x by a directed path, then we write $y \preceq_F x$ and say that y is an *ancestor* of x and x is a *descendant* of y . Note that every vertex is considered its own ancestor and descendant. For $x \in V(F)$, we write

$$\begin{aligned} \text{tail}_F[x] &:= \{y: y \preceq_F x\}, & \text{subtree}_F[x] &:= \{z: z \succeq_F x\}, \\ \text{tail}_F(x) &:= \text{tail}_F[x] \setminus \{x\}, & \text{subtree}_F(x) &:= \text{subtree}_F[x] \setminus \{x\}. \end{aligned}$$

The *level* of a node x in F , denoted $\text{lvl}_F(x)$, is the number of its strict ancestors, that is, $|\text{tail}_F(x)|$. Note that roots have level 0. The *height* of a forest F is the maximum level among its nodes, plus 1. If the forest F is clear from the context, then we may omit it in the above notation.

An *elimination forest* of a graph G is a rooted forest F with $V(F) = V(G)$ such that for every edge uv of G , either u is an ancestor of v in F or vice versa. The *treedepth* of a graph G is the least possible height of an elimination forest of G . Treedepth as a graph parameter plays a central role in the structural theory of sparse graphs, see [45, Chapters 6 and 7]. It also has several applications in parameterized complexity and algorithm design [12, 22, 28, 44, 46, 47], as well as exhibits interesting combinatorial properties [12, 17, 21] and connections to descriptive complexity theory [23]. We refer to the introductory sections of the above works for a wider discussion.

A *tree decomposition* of a graph G is a pair $\mathbb{T} = (T, \beta)$, where T is an (unrooted) tree and $\beta: V(T) \rightarrow 2^{V(G)}$ is a function that assigns to each node $x \in V(T)$ its *bag* $\beta(x) \subseteq V(G)$ so that the following two conditions are satisfied:

- for each $u \in V(G)$, the set $\{x: u \in \beta(x)\}$ induces a nonempty and connected subtree of T ; and
- for each $uv \in E(G)$, there exists $x \in V(T)$ such that $\{u, v\} \subseteq \beta(x)$.

The *width* of \mathbb{T} is $\max_{x \in V(T)} |\beta(x)| - 1$ and the *treewidth* of G is the minimum possible width of a tree decomposition of G . It is easy to see that the treedepth of a graph is at most its treewidth plus one. Conversely, the treewidth is upper bounded by the treedepth times the logarithm of the vertex count [45].

For surgery on tree decompositions we will use the following definition and standard lemma.

Definition 3.2 (Segment of a tree). *For an unrooted tree T , a segment of T is a nonempty and connected subtree I of T such that there are at most two vertices of I that have a neighbor outside of I . The set of those at most two vertices is the boundary of I , and is denoted by ∂I . The size of I is equal to $|E(I)|$.*

Lemma 3.3. *Let T be an unrooted tree and let I be a segment of T of size $\ell \geq 2$. Then there are at most 5 segments I_1, \dots, I_t of T ($t \leq 5$), each of size at most $\ell/2$, such that segments I_1, \dots, I_t have pairwise disjoint edge sets and $E(I_1) \cup \dots \cup E(I_t) = E(I)$.*

Proof. For each edge $xy \in E(I)$, let $I_{y,x}$ and $I_{x,y}$ be the connected components of $I - xy$ that contain x and y , respectively. Let \vec{I} be the orientation of I where each edge xy is oriented towards x if $|E(I_{y,x})| > |E(I_{x,y})|$ and towards y if $|E(I_{y,x})| < |E(I_{x,y})|$; in case $|E(I_{y,x})| = |E(I_{x,y})|$, the edge xy is oriented in any way. Since I has ℓ edges and $\ell + 1$ nodes, there is a node z of I that has outdegree 0 in \vec{I} . This means that for every neighbor x of z , we have $|E(I_{z,x})| \leq |E(I_{x,z})|$, implying $|E(I_{z,x})| < \ell/2$. Denote $I_x := I_{z,x}$ and let \widehat{I}_x be I_x with the edge xz added.

We first argue that I can be edge-partitioned into at most 3 subtrees (not necessarily segments), each with at most $\ell/2$ edges. Consider first the corner case when there exists a neighbor x of z such that \widehat{I}_x has more than $\ell/2$ edges. Then both $I_x = I_{z,x}$ and $I_{x,z}$ have exactly $\frac{\ell-1}{2}$ edges each, so we can partition I into $I_{z,x}$, $I_{x,z}$, and a separate subtree consisting only of the edge xz . This case being resolved, we can assume that each tree \widehat{I}_x has at most $\ell/2$ edges. Starting with the set of trees $\mathcal{T} := \{\widehat{I}_x: x \text{ is a neighbor of } z\}$, iteratively apply the following procedure: take two trees from \mathcal{T} with the smallest edge counts, and replace them with their union, provided this union has at most $\ell/2$ edges. The procedure stops when this assertion

fails to be satisfied. Observe that the procedure can be carried out as long as $|\mathcal{T}| \geq 4$, for then the two trees from \mathcal{T} that have the smallest edge counts together include at most half of the edges of I . Therefore, at the end we obtain the desired edge-partition of I into at most three subtrees.

All in all, in both cases we edge-partitioned I into at most three subtrees, each having at most $\ell/2$ edges. Since $|\partial I| \leq 2$, it is easy to see that all of those subtrees are already segments (i.e. have boundaries of size at most 2) apart from at most one, say J , which may have a boundary of size 3. Supposing that J exists, let $\partial J = \{a, b, c\}$. Then there exists a node d of J such that every connected component of $J - d$ contains at most one of the vertices a, b, c . It is now straightforward to edge-partition J into three trees so that the boundary of each of them consists of d and one of the vertices a, b, c . Thus, replacing J with those three segments yields an edge-partition of I into at most 5 segments, each with at most $\ell/2$ edges. \square

4 Isolating Hamiltonian cycles

In this section we prove Theorems 1.3, 1.2, and 1.4. We begin by defining *configurations* for Hamiltonian cycles, which reflect the states of a natural dynamic programming algorithm for detection of a Hamiltonian cycle in a bounded-treewidth graph. Then we use the rank-based approach to bound the number of *minimum weight compliant edge sets* (see Theorem 4.6). This technical result captures the essence of the rank-based approach and will be used in all subsections that follow. Next, we prove Theorem 1.2 in Section 4.3. Then Theorem 1.3 is proved in Section 4.4. Finally, in Section 4.5 we first recall basic definitions and facts about separable graph classes, then we give a decomposition theorem (Theorem 4.16) for such classes that produces a low-depth elimination forest with several important technical properties, and finally we use this decomposition theorem to prove Theorem 1.4.

4.1 Configurations for Hamiltonian cycles

Let us fix a graph G . An edge set $S \subseteq E(G)$ is called a *partial solution* if every vertex of G is incident to at most two edges of S and S has no cycles. The following notion of a *configuration* describes the behavior of a partial solution with respect to a set of vertices.

Definition 4.1 (Configurations). *For $X \subseteq V(G)$, we define the set of configurations on X as:*

$$\text{conf}(X) := \{ (V_0, V_1, V_2, M) : (V_0, V_1, V_2) \text{ is a partition of } X \text{ and } M \text{ is a perfect matching on } V_1 \}.$$

Given a subgraph H of G , one can view the configurations on $X \subseteq V(H)$ as all possible different ways that a partial solution may behave on X . A vertex is then in the set V_i if it is incident to exactly i edges of the partial solution. The matching M on V_1 describes the endpoints of each path in the partial solution. This intuition is formalized in the following definition.

Definition 4.2. *Let $X \subseteq V(G)$ be a set of vertices of G and let $S \subseteq E(G)$ be a partial solution. Then define the configuration of S on X as $c_X(S) := (V_0, V_1, V_2, M) \in \text{conf}(X)$, where*

- $V_0 := \{v \in X : v \text{ is not incident to any edge of } S\}$,
- $V_1 := \{v \in X : v \text{ is incident to exactly one edge of } S\}$,
- $V_2 := \{v \in X : v \text{ is incident to exactly two edges of } S\}$,
- $M := \{\{u, v\} \in \binom{V_1}{2} : \text{there is a path with edges from } S \text{ connecting } u \text{ and } v\}$,

We omit X in the notation and write $c(S)$ when X is clear from context.

Note that in the above definition M is indeed a matching, because each $v \in V_1$ is connected to exactly one $u \in V_1$ through S , as any partial solution covers each vertex at most twice. For an example of deriving $c_X(S)$ from a partial solution S , see Figure 3.




Figure 3: Example partial solution S and its configuration $c_X(S) = (V_0, V_1, V_2, M)$ on a set X .




Figure 4: Example compliant partial solution S for a configuration $c = (V_0, V_1, V_2, M) \in \text{conf}(X)$.

We can use configurations to tell whether two partial solutions together form a Hamiltonian cycle. Let H be a subgraph of G and let $X \subseteq V(H)$. Assume that there exists a partial solution S that visits only vertices from $(V(G) \setminus V(H)) \cup X$, where every vertex of $V(G) \setminus V(H)$ is visited exactly twice. Then we only need to know $c_X(S)$ to determine which partial solutions $S' \subseteq E(H)$ would combine with S to a Hamiltonian cycle in G . We say that any such partial solution is *compliant* with $c_X(S)$, as expressed formally in the next definition.

Definition 4.3 (Compliant partial solution). *Let H be a subgraph of G and let $X \subseteq V(H)$. A configuration $c = (V_0, V_1, V_2, M) \in \text{conf}(X)$ and a partial solution $S \subseteq E(H)$ are compliant if $S \cap M = \emptyset$ and $S \cup M$ forms a Hamiltonian cycle on $V(H) \setminus V_2$.*

See Figure 4 for an example of a compliant partial solution.

In the sequel we will be trying to argue that some weight function ω is isolating the family of Hamiltonian cycles in the given graph G with high probability. In all cases this will be done by induction on larger and larger subgraphs of G , where at each point we argue that a suitable family of partial solutions is isolated with high probability. The following definition facilitates this discussion.

Definition 4.4 (Minimum weight compliant partial solution). *Let H be a subgraph of G , $X \subseteq V(H)$, $c \in \text{conf}(X)$, and let $\omega: E(G) \rightarrow \mathbb{N}$ be a weight function on the edges of G . Then we define the set $\text{Min}(\omega, H, c)$ of minimum weight partial solutions compliant with c as the set of those partial solutions $S \subseteq E(H)$ that*

- *are compliant with c , and*
- *subject to the above, have the smallest possible weight $\omega(S)$.*

4.2 Rank-based approach

We will use the *rank-based approach*, introduced by Cygan et al. in [15], as a tool in our analysis of isolation schemes. Let X be a set of vertices. Then define the *compatibility matrix* \mathcal{H}_X as the matrix with entries indexed by $\mathcal{H}_X[M_1, M_2]$ for M_1, M_2 perfect matchings on X , where

$$\mathcal{H}_X[M_1, M_2] = \begin{cases} 1 & \text{if } M_1 \cup M_2 \text{ is a simple cycle,} \\ 0 & \text{otherwise.} \end{cases}$$

Note that $\mathcal{H}_X[M_1, M_2]$ has $2^{\mathcal{O}(|X| \log |X|)}$ rows and columns. The crux of the rank-based approach is that in spite of that, this matrix has a small rank over the two-element field \mathbb{F}_2 .

Theorem 4.5 (Rank-based approach,[15]). *For any set X , the rank of \mathcal{H}_X over \mathbb{F}_2 is equal to $2^{|X|/2-1}$.*

We use Theorem 4.5 to prove that the total number of minimum weight compliant solutions is always relatively small, no matter what the weight function is. The following statement will be reused several times in the sequel. Note that a trivial cardinality argument would yield an upper bound of the form $2^{\mathcal{O}(|X|\log|X|)}$; the point of the rank-based approach is to reduce this to $2^{\mathcal{O}(|X|)}$.

Theorem 4.6. *Let G be a graph, $X \subseteq V(G)$, and $\omega: V(G) \rightarrow \mathbb{N}$ be a weight function such that for all $c \in \text{conf}(X)$, we have $|\text{Min}(\omega, G, c)| \leq 1$. Then*

$$\left| \bigcup_{c \in \text{conf}(X)} \text{Min}(\omega, G, c) \right| \leq 2^{\mathcal{O}(|X|)}.$$

Proof. Let $K := \bigcup_{c \in \text{conf}(X)} \text{Min}(\omega, G, c)$ and let $C := \{c(S) : S \in K\}$.

We first verify that $|C| = |K|$. By construction, we have $|C| \leq |K|$. Assume for contradiction that $|C| < |K|$. Then there are two different partial solutions $S_1, S_2 \in K$ such that $c(S_1) = c(S_2)$. By construction and the assumptions, there are two different configurations $d_1, d_2 \in \text{conf}(X)$ such that $\text{Min}(\omega, G, d_1) = \{S_1\}$ and $\text{Min}(\omega, G, d_2) = \{S_2\}$. However, since $c(S_1) = c(S_2)$, it follows that for any configuration $d \in \text{conf}(X)$, S_1 is compliant with d if and only if S_2 is compliant with d . In particular, S_1 is compliant with d_2 and S_2 is compliant with d_1 . This implies that $\omega(S_1) = \omega(S_2)$ and $S_2 \in \text{Min}(\omega, G, d_1)$ and $S_1 \in \text{Min}(\omega, G, d_2)$, a contradiction. Hence $|C| = |K|$.

Define a matrix $\widehat{\mathcal{H}}$ with both coordinates indexed by $\text{conf}(X)$ such that for $c, c' \in \text{conf}(X)$, where $c = (V_0, V_1, V_2, M)$ and $c' = (V'_0, V'_1, V'_2, M')$:

$$\widehat{\mathcal{H}}[c, c'] = \begin{cases} 1 & \text{if } V_0 = V'_0, V_2 = V'_2, \text{ and } M \cup M' \text{ is a simple cycle,} \\ 0 & \text{otherwise.} \end{cases}$$

Notice that if we sort the indices of $\widehat{\mathcal{H}}$ by the partitions (V_0, V_1, V_2) , then $\widehat{\mathcal{H}}$ can be seen as a block diagonal matrix with one block for each partition, and this block is a compatibility matrix on V_1 . That is,

$$\widehat{\mathcal{H}} = \bigoplus_{V_0 \uplus V_1 \uplus V_2 = X} \mathcal{H}_{V_1},$$

where \bigoplus denotes the operator of combining several matrices into a single block diagonal matrix. By Theorem 4.5, the rank over \mathbb{F}_2 of each of these blocks is bounded by $2^{|X|/2-1}$, hence the rank over \mathbb{F}_2 of $\widehat{\mathcal{H}}$ is bounded by $2^{|X|/2-1} \cdot 3^{|X|} \leq 2^{\mathcal{O}(|X|)}$.

Next, we claim that the set of rows of $\widehat{\mathcal{H}}$ corresponding to the configurations of C is linearly independent over \mathbb{F}_2 . Assume not, hence there is a nonempty set of configurations $D \subseteq C$ such that

$$\sum_{d \in D} \widehat{\mathcal{H}}[d, \cdot] = \mathbf{0},$$

where $\mathbf{0}$ is the all-zero vector (all computations are performed in \mathbb{F}_2). For each $d \in D$ there is some $S_d \in K$ such that $d = c(S_d)$. Let d_{\max} be a configuration of D for which $\omega(S_{d_{\max}})$ is the largest possible. Since $d_{\max} \in C$, we have that $\text{Min}(\omega, G, c) = \{S_{d_{\max}}\}$ for some $c \in \text{conf}(X)$ and hence $\widehat{\mathcal{H}}[d_{\max}, c] = 1$. However, as $\sum_{d \in D} \widehat{\mathcal{H}}[d, \cdot] = \mathbf{0}$, there must be another $d' \in D$, $d' \neq d_{\max}$, such that also $\widehat{\mathcal{H}}[d', c] = 1$. This means that d' is compliant with c , which implies that $\omega(S_{d'}) > \omega(S_{d_{\max}})$ by $\text{Min}(\omega, G, c) = \{S_{d_{\max}}\}$. This contradicts the maximality of $\omega(S_{d_{\max}})$.

We conclude that the set of rows of $\widehat{\mathcal{H}}$ corresponding to C are indeed linearly independent over \mathbb{F}_2 . Therefore, $|K| = |C|$ is upper bounded by the rank of $\widehat{\mathcal{H}}$ over \mathbb{F}_2 , which is at most $2^{\mathcal{O}(|X|)}$. \square

4.3 Hamiltonian cycles in general graphs using $\mathcal{O}(n)$ random bits

We now use the tools prepared so far to prove Theorem 1.2. The goal is to isolate all Hamiltonian cycles in an undirected graph $G = (V, E)$ using $\mathcal{O}(n)$ random bits, where n is the vertex count. First we give the isolation procedure. Then we analyze the probability of isolating all Hamiltonian cycles using configurations, compliant partial solutions, and the rank-based approach (through Theorem 4.6). Throughout the subsection we assume without loss of generality that $\log n$ is an integer.

As usual with isolation schemes, we assume that the vertex set of the considered graph G is $V = [n]$. We will apply induction on specific subgraphs of G called *intervals*.

Definition 4.7 (Interval of G). *For integers $1 \leq s \leq t \leq n$ and $1 \leq s' \leq t' \leq n$, the interval $G\langle s, t, s', t' \rangle$ is the graph (V', E') , where*

$$V' := \{s, \dots, t\} \cup \{s', \dots, t'\} \quad \text{and} \quad E' := \{uv : u \in \{s, \dots, t\}, v \in \{s', \dots, t'\}, uv \in E\}.$$

By $V\langle s, t, s', t' \rangle$ we denote the vertex set V' of the interval $G\langle s, t, s', t' \rangle$.

Note that $G\langle s, t, s, t \rangle$ is just the subgraph of G induced by $\{s, \dots, t\}$. On the other hand, if $\{s, \dots, t\} \cap \{s', \dots, t'\} = \emptyset$, then $G\langle s, t, s', t' \rangle$ is a bipartite graph, with $\{s, \dots, t\}$ and $\{s', \dots, t'\}$ being the sides of the bipartition.

Isolation scheme. We first present the isolation scheme. Let $\text{id}: E(G) \rightarrow \{1, \dots, |E(G)|\}$ be any bijection that assigns to each edge $e \in E(G)$ its unique *identifier* $\text{id}(e)$. Let C be some large enough constant, to be chosen later. Then independently at random sample $1 + \log n$ primes $p_0, p_1, \dots, p_{\log n}$ so that p_i is sampled uniformly among primes in the range $\{1, \dots, M_i\}$, where $M_i := 2^{C(\log n + 2^i)}$. Note that choosing each p_i requires $C(\log n + 2^i)$ random bits, hence we have used $\mathcal{O}(n)$ random bits in total.

Next, we inductively define weights functions $\omega_0, \dots, \omega_{\log n}$ on $E(G)$ as follows:

- Set $\omega_0(e) := 2^{\text{id}(e)} \bmod p_0$ for all $e \in E(G)$.
- For each $e \in E(G)$ and $i = 1, \dots, \log n$, set

$$\omega_i(e) := M_{i-1}n \cdot \omega_{i-1}(e) + \left(2^{\text{id}(e)} \bmod p_i\right).$$

Let $\omega := \omega_{\log n}$ and observe that ω assigns weights bounded by $2^{\mathcal{O}(n)}$, as required.

Analysis. We will prove the following statement for all $0 \leq i \leq \log n$ using induction on i .

Induction hypothesis. With probability at least $\left(1 - \frac{1}{n^2}\right)^{i+1}$, for all intervals $G\langle s, t, s', t' \rangle$ s.t. $t-s \leq 2^i$ and $t' - s' \leq 2^i$ and for each configuration $c \in \text{conf}(V\langle s, t, s', t' \rangle)$, there is at most one minimum weight (w.r.t. ω_i) compliant partial solution, i.e. $|\text{Min}(\omega_i, G\langle s, t, s', t' \rangle, c)| \leq 1$.

For $i = \log n$, the induction hypothesis gives us that for the complete interval $G = G\langle 1, 1, n, n \rangle$ and for the configuration $c = (\emptyset, \emptyset, V(G), \emptyset)$, there is at most one minimum weight compliant partial solution w.r.t. ω . In other words, w.r.t. ω there is at most one minimum weight Hamiltonian cycle in G . This happens with probability at least $\left(1 - \frac{1}{n^2}\right)^{\log n + 1} \geq 1 - \frac{1}{n}$. So it remains to perform the induction.

Base step. For $i = 0$, we have $t - s \leq 1$ and $t' - s' \leq 1$. Hence each such interval $G\langle s, t, s', t' \rangle$ has at most 4 edges. Let

$$Y := \bigcup_{\substack{t-s \leq 1 \\ t'-s' \leq 1}} 2^{E(G\langle s, t, s', t' \rangle)}$$

and for each $S \in Y$, let

$$x_S := \sum_{e \in S} 2^{\text{id}(e)}.$$

Observe that since the identifiers assigned to the edges are unique, the numbers x_S are also pairwise different. Also, note that $|Y| \leq 16n^2$ as there are at most n^2 intervals considered, and for each of them there are at most 16 possible subsets of the at most four edges. Recall that $M_0 = 2^{C(\log n + 1)}$ and p_0 is drawn uniformly at random among the primes in the range $\{1, \dots, M_0\}$. Therefore, from Lemma 3.1 we can conclude that with probability at least

$$\left(1 - \frac{(n^2 + 1)(16n^2)^2}{2^{(C/2)(\log n + 1)}}\right) \geq \left(1 - \frac{1}{n^2}\right)$$

all the numbers $\{x_S : S \in Y\}$ have pairwise different remainders modulo p_0 ; here the last inequality holds for a large enough constant C . Since $\omega_0(S) \equiv x_S \pmod{p_0}$, this means that with probability at least $(1 - \frac{1}{n^2})$, all $S \in Y$ receive pairwise different weights with respect to ω_0 . Therefore, the induction hypothesis is true for $i = 0$.

Induction step. Assume the induction hypothesis is true for all intervals $G\langle s, t, s', t' \rangle$ such that $t - s \leq 2^{i-1}$ and $t' - s' \leq 2^{i-1}$. Let

$$Y' := \bigcup_{\substack{t-s \leq 2^{i-1} \\ t'-s' \leq 2^{i-1}}} \bigcup_{c \in \text{conf}(V\langle s, t, s', t' \rangle)} \text{Min}(\omega_{i-1}, G\langle s, t, s', t' \rangle, c)$$

be the set of all the minimal partial solutions for those intervals. Further, let

$$Y := \{S_1 \cup S_2 \cup S_3 \cup S_4 : S_1, S_2, S_3, S_4 \in Y'\}$$

be the set containing all combinations of four such partial solutions. The strategy is as follows. We first prove in Claim 4.8 that any relevant minimum weight compliant partial solution should be in Y . Then Claim 4.9 says that with high probability, all partial solutions $S \in Y$ have pairwise different weights with respect to ω_i . Hence, proving these two claims will be sufficient to make the induction hypothesis go through.

Claim 4.8. *Let $1 \leq a \leq b \leq n$ and $1 \leq a' \leq b' \leq n$ be such that $b - a \leq 2^i$ and $b' - a' \leq 2^i$, and let $c \in \text{conf}(a, b, a', b')$. Then $\text{Min}(\omega_i, G\langle a, b, a', b' \rangle, c) \subseteq Y$.*

Proof. Take any $S \in \text{Min}(\omega_i, G\langle a, b, a', b' \rangle, c)$. Let

$$r = \lceil (a + b)/2 \rceil \quad \text{and} \quad r' = \lceil (a' + b')/2 \rceil$$

and let us select

$$\begin{aligned} S_1 &\subseteq E(G\langle a, r - 1, a', r' - 1 \rangle), & S_2 &\subseteq E(G\langle a, r - 1, r', b' \rangle), \\ S_3 &\subseteq E(G\langle r, b, a', r' - 1 \rangle), & S_4 &\subseteq E(G\langle r, b, r', b' \rangle) \end{aligned}$$

so that S_1, S_2, S_3, S_4 are disjoint and $S = S_1 \cup S_2 \cup S_3 \cup S_4$. See Figure 5 for an example.

We argue that $S_1 \in \text{Min}(\omega_{i-1}, G\langle a, r-1, a', r'-1 \rangle, c_1)$ for some $c_1 \in \text{conf}(V\langle a, r-1, a', r'-1 \rangle)$. Let $c = (V_0, V_1, V_2, M)$. Since $S \cup M$ is a simple cycle that visits all vertices of $V\langle a, b, a', b' \rangle$, we see that $R := S_2 \cup S_3 \cup S_4 \cup M$ is a partial solution in the graph $G\langle a, b, a', b' \rangle$ with the edges of M added. Letting $(V'_0, V'_1, V'_2, M') := c_{V\langle a, r-1, b, r-1 \rangle}(R)$, it follows that S_1 is compliant with the configuration

$$c_1 := (V'_0 \setminus (V_2 \cap V\langle a, r-1, b, r-1 \rangle), V'_1, V'_2 \cup (V_2 \cap V\langle a, r-1, b, r-1 \rangle), M').$$

Moreover, that $S \in \text{Min}(\omega_i, G\langle a, b, a', b' \rangle, c)$ implies that $S_1 \in \text{Min}(\omega_i, G\langle a, r-1, a', r'-1 \rangle, c_1)$, for otherwise S_1 could be replaced in S with a smaller-weight partial solution S'_1 that would be still compliant with c_1 , and this would turn S into a smaller-weight partial solution $S' = S'_1 \cup S_2 \cup S_3 \cup S_4$ that would be still compliant with c . Finally, by the construction of ω_i , $S_1 \in \text{Min}(\omega_i, G\langle a, r-1, a', r'-1 \rangle, c_1)$ entails $S_1 \in \text{Min}(\omega_{i-1}, G\langle a, r-1, a', r'-1 \rangle, c_1)$.

Therefore $S_1 \in Y'$. Analogously we argue that $S_2, S_3, S_4 \in Y'$, hence we conclude that $S \in Y$. \square




Figure 5: Example of splitting a partial solution $S \in E(G\langle a, b, a', b' \rangle)$ into four partial solutions S_1, S_2, S_3, S_4 , where $S_1 \subseteq E(G\langle a, r-1, a', r'-1 \rangle)$, $S_2 \subseteq E(G\langle a, r-1, r', b' \rangle)$, $S_3 \subseteq E(G\langle r, b, a', r'-1 \rangle)$ and $S_4 \subseteq E(G\langle r, b, r', b' \rangle)$ with $r = \lceil (a+b)/2 \rceil$ and $r' = \lceil (a'+b')/2 \rceil$.

Claim 4.9. *The following event happens with probability at least $(1 - \frac{1}{n^2})^{i+1}$: for all different $S, S' \in Y$, it holds that $\omega_i(S) \neq \omega_i(S')$.*

Proof. For each $S \in Y$, let

$$x_S := \sum_{e \in S} 2^{\text{id}(e)}.$$

Observe that since identifiers assigned to the edges are unique, the numbers x_S are pairwise different. The induction hypothesis gives us that the following event A_{i-1} happens with probability at least $(1 - \frac{1}{n^2})^i$: for all $1 \leq s \leq t \leq n$ and $1 \leq s' \leq t' \leq n'$ with $t-s \leq 2^{i-1}$ and $t'-s' \leq 2^{i-1}$, and all $c \in \text{conf}(V\langle s, t, s', t' \rangle)$, we have $|\text{Min}(\omega_{i-1}, G\langle s, t, s', t' \rangle, c)| \leq 1$. Assuming now that A_{i-1} indeed happens, by Theorem 4.6 we conclude that for every fixed choice of s, t, s', t' as above, we have

$$\left| \bigcup_{c \in \text{conf}(V\langle s, t, s', t' \rangle)} \text{Min}(\omega_{i-1}, G\langle s, t, s', t' \rangle, c) \right| \leq 2^{\mathcal{O}(2^{i-1})}.$$

Since there are at most n^4 choices of s, t, s', t' , this implies that

$$|Y| \leq |Y'|^4 \leq 2^{\mathcal{O}(2^{i-1})} \cdot n^{16}.$$

Since $M_i = 2^{C(\log n + 2^i)}$ and p_i is drawn uniformly at random among the primes in the range $\{1, \dots, M_i\}$, from Lemma 3.1 we can conclude that, for large enough C , with probability at least

$$\left(1 - \frac{(n^2 + 1) \left(n^{16} 2^{\mathcal{O}(2^{i-1})}\right)^2}{2^{(C/2)(\log n + 2^i)}}\right) \cdot \left(1 - \frac{1}{n^2}\right)^i \geq \left(1 - \frac{1}{n^2}\right)^{i+1},$$

all the numbers $\{x_S : S \in Y\}$ have pairwise different remainders modulo p_i ; here, the term $\left(1 - \frac{1}{n^2}\right)^i$ corresponds to the probability that A_i happens. As a consequence, with the same probability we have that $\omega_i(S) \neq \omega_i(S')$ for all different $S, S' \in Y$. \square

Now the induction step follows directly from combining Claim 4.8 with Claim 4.9.

4.4 Hamiltonian cycles in graphs of bounded treewidth

We will now use the same approach to give a proof of Theorem 1.3. More precisely, assume we are given a graph G of treewidth at most k . Our goal is to isolate the family of Hamiltonian cycles in G using $\mathcal{O}(k \log n + \log^2 n)$ random bits.

The proof follows the same structure as that of Theorem 1.2. We first describe the isolation scheme and then analyze the scheme using a tree decomposition $\mathbb{T} = (T, \beta)$ of G of width at most k . Note that the actual decomposition is not needed for the isolation procedure, and is only used as a tool in the analysis.

Isolation scheme. We first present the isolation scheme. As before, we assume that $V(G) = [n]$ and n is a power of 2. Let $\text{id} : E(G) \rightarrow \{1, \dots, |E(G)|\}$ be any bijection that assigns to each edge $e \in E(G)$ its unique *identifier* $\text{id}(e)$. Let C be some large enough constant, to be chosen later. Then we independently sample $3 \log n$ primes $p_1, \dots, p_{3 \log n}$ so that each p_i is sampled uniformly among all primes in the interval $\{1, \dots, M\}$, where $M = 2^{C(k \log n)}$. Note that choosing each p_i requires $C(k + \log n)$ random bits, hence we have used $\mathcal{O}(k \log n + \log^2 n)$ random bits in total, as required.

Next, we inductively define weights functions $\omega_0, \dots, \omega_{3 \log n}$ on $E(G)$ as follows:

- Set $\omega_0(e) := 0$ for all $e \in E(G)$.
- For each $e \in E(G)$ and $i = 1, \dots, 3 \log n$, set

$$\omega_i(e) := Mn \cdot \omega_{i-1}(e) + \left(2^{\text{id}(e)} \bmod p_i\right).$$

We let $\omega := \omega_{3 \log n}$ and we observe that ω assigns weights bounded by $2^{\mathcal{O}(k \log n + \log^2 n)}$.

Analysis. Let $\mathbb{T} = (T, \beta)$ be a tree decomposition of G of width at most k . It is well-known that T can be chosen so that it has at most n nodes. Further, let $\eta := E(G) \rightarrow V(T)$ be any function that assigns to each edge e of G any node x of T such that $e \subseteq \beta(x)$. In the sequel we will assume that η is injective. This can be achieved by adding, for each node $x \in V(T)$, $|\eta^{-1}(x)| - 1$ new nodes with the same bag and adjacent only to x , and appropriately distributing the images of edges of $\eta^{-1}(x)$ among the new nodes. Note that after this modification, the number of nodes of T is bounded by $\binom{k+1}{2} \cdot n \leq n^3$.

Compared to the proof of Theorem 1.2, instead of intervals we will use *segments* in the tree T underlying the tree decomposition \mathbb{T} . Recall that segments have been defined and discussed in Section 3. We first observe that there are only few segments.

Claim 4.10. *There are at most n^9 segments of T .*

Proof. Note that a segment I in T can be uniquely determined by specifying the at most two vertices of ∂I and any vertex of $V(I) \setminus \partial I$, provided there exists any. Since T has at most n^3 nodes, there are at most n^9 choices for such a specification. \square

For a set of nodes $Z \subseteq V(T)$, we write $\beta(Z) := \bigcup_{z \in Z} \beta(z)$. Further, for a segment I of T we consider the graph

$$G\langle I \rangle := (\beta(V(I)), \eta^{-1}(V(I))).$$

Usually when speaking about partial solutions in $G\langle I \rangle$, we consider their configurations on the vertex subset $\beta(\partial I)$. Note that $G\langle T \rangle = G$.

We proceed to the induction. We will prove the following statement for all $0 \leq i \leq \log n$.

Induction hypothesis. With probability at least $(1 - \frac{1}{n^2})^i$, for all segments I of T of size at most 2^i and for each configuration $c \in \text{conf}(\beta(\partial I))$, there is at most one minimum weight (w.r.t. ω_i) compliant partial solution in $G\langle I \rangle$, i.e. $|\text{Min}(\omega_i, G\langle I \rangle, c)| \leq 1$.

Note that since $|V(T)| \leq n^3$, for $i = 3 \log n$ the induction hypothesis gives that for $G\langle T \rangle = G$, there is at most one Hamiltonian cycle that has the minimum weight w.r.t. ω with probability at least $(1 - \frac{1}{n^2})^{3 \log n} \geq (1 - \frac{1}{n})$.

Base step. For $i = 0$, we take segments of size at most 1, i.e. we prove the induction hypothesis for every segment I of T that has either one or two nodes. More precisely, we have to prove that (with suitably large probability), for every such segment I and configuration $c \in \text{conf}(\beta(\partial I))$, we have $|\text{Min}(\omega_0, G\langle I \rangle, c)| \leq 1$. Note that since I has at most two nodes and η is injective, the edge set $E(G\langle I \rangle)$ consists of at most two edges. Moreover, it cannot be that two different edge subsets $E_1, E_2 \subseteq E(G\langle I \rangle)$ are simultaneously compliant with the same configuration $c \in \text{conf}(\beta(\partial I))$. It follows that sets $\text{Min}(\omega_0, G\langle I \rangle, c)$ have sizes at most 1 always, so the induction hypothesis for $i = 0$ is true.

Induction step. Assume the induction hypothesis is true for all segments of size at most 2^{i-1} . Let

$$Y' := \bigcup_{I: \text{segment of size } \leq 2^{i-1}} \bigcup_{c \in \text{conf}(\beta(\partial I))} \text{Min}(\omega_{i-1}, G\langle I \rangle, c).$$

be the set of all minimum weight partial solutions for segments of size at most 2^{i-1} . Further, let

$$Y := \{ S_1 \cup S_2 \cup S_3 \cup S_4 \cup S_5 : S_1, S_2, S_3, S_4, S_5 \in Y' \}$$

be the set comprising all combinations of five such partial solutions.

We first prove with Claim 4.11 that every relevant minimum weight compliant edge is contained in Y . Then Claim 4.12 says that with high probability, all $S \in Y$ receive pairwise different weights with respect to ω_i . The induction hypothesis will follow directly from combining these two claims.

Claim 4.11. *Let I be any segment of size at most 2^i and let $c \in \text{conf}(\beta(\partial I))$. Then $\text{Min}(\omega_i, G\langle I \rangle, c) \subseteq Y$.*

Proof. Consider any $S \in \text{Min}(\omega_i, G\langle I \rangle, c)$. By Lemma 3.3, there exist segments I_1, \dots, I_t ($t \leq 5$), each of size at most 2^{i-1} , such that $E(I)$ is the disjoint union of $E(I_1), \dots, E(I_t)$. For each $j \in \{1, \dots, t\}$ choose $S_j \in E(G\langle I_j \rangle)$ so that S is the disjoint union of S_1, \dots, S_t . The same argument as that was used in the proof of Claim 4.8 shows that there exists $c_j \in \text{conf}(\beta(\partial I_j))$ such that $S_j \in \text{Min}(\omega_{i-1}, G\langle I_j \rangle, c_j)$. Hence $S_j \in Y'$ for all $j \in \{1, \dots, t\}$, so it follows that $S \in Y$. \square

Claim 4.12. *The probability of the following event is at least $(1 - \frac{1}{n^2})^i$: for all different $S, S' \in Y$, it holds that $\omega_i(S) \neq \omega_i(S')$.*

Proof. For each $S \in Y$ let us define

$$x_S = \sum_{e \in S} 2^{\text{id}(e)}.$$

Observe that since the identifiers assigned to the edges are unique, the numbers x_S are pairwise different. By the induction hypothesis, the following event A_{i-1} happens with probability at least $(1 - \frac{1}{n^2})^{i-1}$: for every segment I of size at most 2^{i-1} and each configuration $c \in \text{conf}(\beta(\partial I))$, we have $|\text{Min}(\omega_{i-1}, G\langle I, c \rangle)| \leq 1$. By Theorem 4.6 it follows that provided A_{i-1} happens, for every fixed segment I of size at most 2^{i-1} we have

$$\left| \bigcup_{c \in \text{conf}(\beta(\partial I))} \text{Min}(\omega_{i-1}, G\langle I, c \rangle) \right| \leq 2^{\mathcal{O}(|\beta(\partial I)|)} \leq 2^{\mathcal{O}(k)}.$$

By Claim 4.10 there are at most n^9 different segments, hence this implies that

$$|Y| \leq |Y'|^5 \leq 2^{\mathcal{O}(k)} \cdot n^{45}.$$

Recall now that $M = 2^{C(k+\log n)}$ and p_i is drawn uniformly at random among the primes in the range $\{1, \dots, M\}$. Hence, from Lemma 3.1 we can conclude that, for large enough C , with probability at least

$$\left(1 - \frac{(n^2 + 1) (2^{\mathcal{O}(k)} \cdot n^{45})^2}{2^{(C/2)(k+\log n)}}\right) \cdot \left(1 - \frac{1}{n^2}\right)^{i-1} \geq \left(1 - \frac{1}{n^2}\right)^i,$$

all the numbers in $\{x_S : S \in Y\}$ have pairwise different remainders modulo p_i . Here, the factor $(1 - \frac{1}{n^2})^{i-1}$ corresponds to the probability that A_{i-1} happens. As a consequence, with the same probability for all different $S, S' \in Y$ we have $\omega_i(S) \neq \omega_i(S')$. \square

The induction step now follows directly from combining Claims 4.11 and 4.12.

4.5 Separable graph classes

In this section we use our understanding of isolation schemes for Hamiltonian cycles in decomposable graphs to design such isolation schemes for *separable* graph classes, that is, classes of graphs that admit small balanced separators. More precisely, we will prove a generalization of Theorem 1.4. First, we need to establish certain terminology and decomposition results.

4.5.1 Definitions and a decomposition theorem

A *graph class* is a (possibly infinite) set of graphs that is closed under taking isomorphisms. A graph class is *hereditary* if it is closed under taking induced subgraphs. The following notion of separability expresses the condition that graphs from a given class can be broken in a balanced way using small separators.

Definition 4.13. A graph class \mathcal{C} is separable with degree $\alpha \in (0, 1)$ if for every graph $G \in \mathcal{C}$, say on n vertices, and vertex subset $S \subseteq V(G)$, there exists a set $X \subseteq V(G)$ such that $|X| \leq \mathcal{O}(n^\alpha)$ and every connected component of $G - X$ contains at most $|S|/2$ vertices of S . Class \mathcal{C} is separable if it is separable with some degree $\alpha \in (0, 1)$.

It is well-known that planar graphs [40] and, more generally, H -minor-free graphs [4] for every fixed H are separable with degree $\frac{1}{2}$. However, this notion is more general. For instance, the class of *1-planar graphs* — graphs that admit a planar embedding where every edge has at most one crossing — is also separable with degree $\frac{1}{2}$ [20]. More generally, every graph class of *polynomial expansion* is separable [48]

(see also the discussion in [45, Sections 16.3 and 16.4]). Examples here would include intersection graphs of bounded-ply families of fat objects in Euclidean spaces of fixed dimension [33]. In fact, subject to technical details, the notions of polynomial expansion and of separability coincide [48, 24].

Our isolation schemes will work on any graph class that is hereditary and separable. That is, we will prove the following generalization of Theorem 1.4.

Theorem 4.14. *Let \mathcal{C} be a hereditary class of graphs that is separable with degree $\alpha \in (0, 1)$. Then there is an isolation scheme for Hamiltonian cycles in graphs from \mathcal{C} that uses $\mathcal{O}(n^\alpha)$ random bits and assigns weights upper bounded by $2^{\mathcal{O}(n^\alpha)}$.*

An important ingredient in the proof of Theorem 4.14 is a decomposition theorem for graphs from a fixed separable class. Intuitively, the decomposition is obtained by recursively breaking the graphs by extracting small balanced separators. The shape of the decomposition will be captured by the following generalization of the notion of an elimination forest.

Definition 4.15. *A generalized elimination forest of a graph G is a rooted forest F together with a mapping $\eta: V(G) \rightarrow V(F)$ satisfying the following property: for every edge $uv \in E(G)$, we have $\eta(u) \preceq \eta(v)$ or $\eta(u) \succeq \eta(v)$ (note that it is possible that $\eta(u) = \eta(v)$). The topological height of (F, η) is simply the height of F , while the height of (F, η) is equal to*

$$\max_{x \in V(F)} \sum_{y \preceq x} |\eta^{-1}(y)|.$$

It is easy to see that if a graph G admits a generalized elimination forest (F, η) of height d , then it also admits an elimination forest of height d : for every node x of F , replace x with a path consisting of vertices of $\eta^{-1}(x)$, in any order. Thus, the intuition is that a generalized elimination forest is a compressed representation of an elimination forest, where some sets of interchangeable vertices — the preimages $\eta^{-1}(x)$ for $x \in V(F)$ — are grouped together in single nodes. The quality of this compression is measured by the parameter topological height.

In the sequel, we will use the following decomposition theorem that ties together separable graph classes and generalized elimination forests. We are not aware of the existence of this particular formulation in the literature, however the proof relies on rather standard techniques.

Theorem 4.16. *Let \mathcal{C} be a graph class that is hereditary and separable with degree $\alpha \in (0, 1)$. Then every graph $G \in \mathcal{C}$, say on n vertices, admits a generalized elimination forest (F, η) satisfying the following conditions:*

(C1) F has one root and every node of F has at most seven children.

(C2) (F, η) has topological height at most $1 + \log_2 n$ and height $\mathcal{O}(n^\alpha)$.

(C3) For every node x of F , if i is the depth of x in F , then

- $|\eta^{-1}(x)| \leq \mathcal{O}((n/2^i)^\alpha)$,
- $|\eta^{-1}(\text{subtree}[x])| \leq n/2^i$, and
- $|N_G(\eta^{-1}(\text{subtree}[x]))| \leq \mathcal{O}((n/2^i)^\alpha)$.

Proof. Let K be the constant hidden in the $\mathcal{O}(\cdot)$ notation in the bound on the sizes of balanced separators in graphs from \mathcal{C} , as prescribed by the definition of the separability of \mathcal{C} . We will use the following simple claim.

Claim 4.17. *Suppose Ω is a finite set and there are two weight functions $\omega_1, \omega_2: \Omega \rightarrow \mathbb{R}_{\geq 0}$ satisfying the following conditions:*

- $\omega_1(\Omega) \leq 1$ and $\omega_2(\Omega) \leq 1$; and
- for each $e \in \Omega$, $\omega_1(e) \leq 1/2$ and $\omega_2(e) \leq 1/2$.

Then there exists a partition \mathcal{P} of Ω into at most seven parts such that for each $P \in \mathcal{P}$, we have $\omega_1(P) \leq 1/2$ and $\omega_2(P) \leq 1/2$.

Proof. Let \mathcal{P}_0 be the partition of Ω that puts every element of Ω into a separate part. By the second assumed condition, \mathcal{P}_0 respects the following assertion (\star): for each part P , we have $\omega_1(P) \leq 1/2$ and $\omega_2(P) \leq 1/2$. We will gradually transform \mathcal{P}_0 into a partition \mathcal{P} consisting of at most seven parts while preserving assertion (\star).

More precisely, starting with \mathcal{P}_0 we define a sequence of partitions $\mathcal{P}_0, \mathcal{P}_1, \mathcal{P}_2, \dots$ that finishes at the first partition \mathcal{P}_i that has at most seven parts; then we set $\mathcal{P} := \mathcal{P}_i$. Each partition \mathcal{P}_{i+1} is obtained from \mathcal{P}_i by merging two parts as follows. Note that \mathcal{P}_i has at least eight parts, for otherwise the construction should have already finished. Call a part P of \mathcal{P}_i *bad* if $\omega_1(P) > 1/4$ or $\omega_2(P) > 1/4$. Clearly, there can be at most six bad parts (at most three due to the first reason, and at most three due to the second reason), which leaves us with at least two parts that are not bad. Then construct \mathcal{P}_{i+1} from \mathcal{P}_i by merging any two not bad parts. It is clear that in this way, assertion (\star) is preserved during the construction and we are done. \square

We proceed to the construction of the generalized elimination forest (F, η) , which will be done by means of a recursive procedure. For a nonempty subset of vertices $A \subseteq V(G)$, the procedure constructs a generalized elimination forest (F_A, η_A) of $H := G[A]$ as follows.

- Let $H' := G[N_G[A]]$. Note that since \mathcal{C} is hereditary, we have $H' \in \mathcal{C}$.
- By the separability of \mathcal{C} , there are vertex subsets $X, Y \subseteq V(H')$, each of size at most $K \cdot |V(H')|^\alpha$, such that
 - every connected component of $H' - X$ contains at most $|A|/2$ vertices of A ; and
 - every connected component of $H' - Y$ contains at most $|S|/2$ vertices of S , where $S := N_G(A)$.
- Let $Z := X \cup Y$. For every connected component C of $H' - Z$, let

$$\omega_1(C) := \frac{|V(C) \cap A|}{|A|} \quad \text{and} \quad \omega_2(C) := \frac{|V(C) \cap S|}{|S|}.$$

In case $S = \emptyset$, we set $\omega_2(C) := 0$.

- Noting that the set of connected components of $H' - Z$ with weight functions ω_1 and ω_2 satisfy the prerequisites of Claim 4.17, we can group the connected components of $H' - Z$ into at most seven graphs, say with vertex sets B_1, \dots, B_t ($t \leq 7$), such that

$$|B_i \cap A| \leq |A|/2 \quad \text{and} \quad |B_i \cap S| \leq |S|/2 \quad \text{for each } i \in [t]. \quad (1)$$

- Recursively apply the procedure to the sets $B_1 \cap A, \dots, B_t \cap A$, thus obtaining generalized elimination forests $(F_1, \eta_1), \dots, (F_t, \eta_t)$ of graphs $G[B_1 \cap A], \dots, G[B_t \cap A]$, respectively.
- Construct the generalized elimination forest (F_A, η_A) of H by taking the union of (F_j, η_j) for $j \in [t]$, adding a single root node r with $\eta_A^{-1}(r) = Z \cap A$, and making the roots of forests F_j into children of r .

It is clear that (F_A, η_A) constructed in this manner is a generalized elimination forest of $G[A]$. We construct the generalized elimination forest (F, η) of G by applying the procedure to $A = V(G)$. Condition (C1) is clear from the construction, hence we need to verify conditions (C2) and (C3).

We start with condition (C3). Observe that it suffices to prove that for every recursive call of the construction procedure, say at recursion depth i , it holds that

$$|Z| \leq \mathcal{O}\left((n/2^i)^\alpha\right), \quad |A| \leq n/2^i \quad \text{and} \quad |N_G(A)| \leq \mathcal{O}\left((n/2^i)^\alpha\right),$$

where Z is as defined in the construction procedure. The second bound follows from a straightforward induction on i using the first part of (1). For the first and third bound, we shall prove by induction on i that

$$|N_G(A)| \leq L \cdot (n/2^i)^\alpha, \tag{2}$$

where

$$L := \max\left(1, \left(\frac{8K}{1 - \frac{1}{2^{1-\alpha}}}\right)^{\frac{1}{1-\alpha}}\right).$$

Note that in the base step, for $i = 0$, we have $N_G(A) = N_G(V(G)) = \emptyset$. During the proof of (2) we will argue that in the considered recursive call of the construction procedure, it holds that

$$|Z| \leq 4KL^\alpha \cdot (n/2^i)^\alpha. \tag{3}$$

Thus, (2) and (3) imply the first and the third bound of condition (C3).

Assume then that (2) holds for the call on a vertex subset A . We need to prove that, assuming the notation from the description of the procedure, for each subsequent call on a subset $B_j \cap A$, $j \in [t]$, we have $|N(B_j \cap A)| \leq L \cdot (n/2^{i+1})^\alpha$. Observe that

$$|V(H')| = |A| + |N_G(A)| \leq n/2^i + L \cdot (n/2^i)^\alpha \leq 2L \cdot (n/2^i),$$

hence

$$|X| \leq K \cdot |V(H')|^\alpha \leq 2KL^\alpha \cdot (n/2^i)^\alpha,$$

and similarly

$$|Y| \leq 2KL^\alpha \cdot (n/2^i)^\alpha.$$

Therefore,

$$|Z| = |X \cup Y| \leq 4KL^\alpha \cdot (n/2^i)^\alpha,$$

which in particular proves (3). Now observe that for each $j \in [t]$, we have

$$N_G(B_j \cap A) \subseteq Z \cup (B_j \cap S).$$

Hence, using the second part of (1), we conclude that

$$\begin{aligned} |N_G(B_j \cap A)| &\leq |Z| + |B_j \cap S| \\ &\leq |Z| + |S|/2 \\ &= |Z| + |N_G(A)|/2 \\ &\leq 4KL^\alpha \cdot (n/2^i)^\alpha + L/2 \cdot (n/2^i)^\alpha \\ &= (4KL^\alpha + L/2) \cdot (n/2^i)^\alpha \\ &\leq (8KL^\alpha + 2^{\alpha-1}L) \cdot (n/2^{i+1})^\alpha \\ &\leq L \cdot (n/2^{i+1})^\alpha. \end{aligned}$$

Here, the last inequality follows from the choice of L . This concludes the inductive proof of (2) and finishes the proof of condition (C3).

We are left with showing condition (C2). The first assertion — that the topological height of (F, η) is bounded by $1 + \log_2 n$ — follows immediately from the first assertion of condition (C3). For the second assertion, observe that in the proof of condition (C3) we argued that $|Z| \leq 4KL^\alpha \cdot (n/2^i)^\alpha$ for calls at recursion depth i . This implies that whenever x is a node of F at depth i , we have $|\eta^{-1}(x)| \leq 4KL^\alpha \cdot (n/2^i)^\alpha$. Therefore, the height of (F, η) is bounded by

$$\sum_{i=0}^{\lfloor \log_2 n \rfloor} 4KL^\alpha \cdot (n/2^i)^\alpha \leq 4KL^\alpha \cdot n^\alpha \cdot \sum_{i=0}^{\infty} \frac{1}{(2^\alpha)^i} \leq \mathcal{O}(n^\alpha).$$

Here, the last inequality is implied by the convergence of the geometric series in question. \square

4.5.2 Isolation scheme

With Theorem 4.16 established, we can proceed to the proof of Theorem 4.14. Let us then fix a hereditary graph class \mathcal{C} that is separable with degree $\alpha \in (0, 1)$, and an n -vertex graph $G \in \mathcal{C}$. We first present the isolation scheme. Let $\text{id}: E(G) \rightarrow \{1, \dots, |E(G)|\}$ be any bijection that assigns to each edge $e \in E(G)$ its unique *identifier* $\text{id}(e)$. Let C be some large enough constant, to be determined later. We independently select $1 + \log n^\alpha + C$ primes

$$p_0, p_1, \dots, p_{\log n^\alpha + C}$$

so that each p_i is sampled uniformly among primes in the range $\{1, \dots, M_i\}$, where $M_i := 2^{C(\log n^\alpha + 2^i)}$. Further, we independently sample $1 + \log n$ primes

$$q_0, q_1, \dots, q_{\log n}$$

so that each q_i is sampled uniformly among primes in the range $\{1, \dots, N_i\}$, where $N_i := 2^{C(\log n + (n/2^i)^\alpha)}$. Note that choosing each p_i requires $C(\log n^\alpha + 2^i)$ random bits and choosing each q_i requires $C(\log n + (n/2^i)^\alpha)$ random bits. Hence we have used $\mathcal{O}(n^\alpha)$ random bits in total, as required.

Next, we inductively define weight functions $\omega_0, \dots, \omega_{\log n^\alpha + C}$ and $\xi_0, \dots, \xi_{\log n}$ on $E(G)$ as follows:

- Set $\omega_0(e) = 2^{\text{id}(e)} \bmod p_0$ for all $e \in E(G)$.
- For $e \in E(G)$ and $i = 1, \dots, \log n^\alpha + C$, set

$$\omega_i(e) = M_{i-1}n \cdot \omega_{i-1}(e) + \left(2^{\text{id}(e)} \bmod p_i\right).$$

- Then, let for all $e \in E(G)$, set:

$$\xi_{\log n}(e) = M_{\log n^\alpha + C}n \cdot \omega_{\log n^\alpha + C}(e) + \left(2^{\text{id}(e)} \bmod q_{\log n}\right)$$

- Finally, for all $e \in E(G)$ and $i = \log n - 1, \dots, 0$, set

$$\xi_i(e) = N_{i+1}n \cdot \xi_{i+1}(e) + \left(2^{\text{id}(e)} \bmod q_i\right).$$

Note that the weight functions ξ_i depend on the weight functions ω_j . Furthermore, the order of defining the weight functions ξ_i might seem inverse to what one might expect. Intuitively, this corresponds to proving a suitable isolation property by a bottom-up induction on the generalized elimination forest provided by

Theorem 4.16. Finally, we define $\omega := \omega_{\log n^\alpha + C}$ and $\xi := \xi_0$. Then we return ξ as the output weight function. Note that ξ assigns weights upper bounded by $2^{\mathcal{O}(n^\alpha)}$, as promised. Hence, it remains to prove that ξ isolates the family of Hamiltonian cycles in G with high probability.

The proof of isolation is done in two steps. First we show that the weight function $\omega = \omega_{\log n^\alpha + C}$ isolates (with high probability) partial solutions on all graphs that intuitively correspond to single nodes of the generalized elimination forest F provided by Theorem 4.16. The second step is to use this knowledge to perform a bottom up induction on F , using weight functions $\xi_{\log n}, \dots, \xi_0$ for the consecutive steps.

4.5.3 Isolation of partial solutions in single nodes

Let (F, η) be a generalized elimination forest of G provided by Theorem 4.16. We may assume that G is connected (as otherwise there are no Hamiltonian cycles in G), hence F is a tree. We may assume that $\eta^{-1}(x) \neq \emptyset$ for every leaf x of F (otherwise x can be disposed of), hence F has at most $(1 + \log_2 n) \cdot n \leq n^2$ nodes.

We first show that the weight function ω , which uses the prime numbers $p_0, \dots, p_{\log n^\alpha + C}$, is enough to isolate all relevant partial in graphs H_x for $x \in V(F)$ defined as follows:

Definition 4.18 (Graph H_x). For each node x in F , define $H_x := (V_x, E_x)$, where

$$V_x := N_G[\eta^{-1}(x)] \cap \eta^{-1}(\text{tail}[x]) \quad \text{and} \quad E_x := \{uv : uv \in E(G), u \in \eta^{-1}(x), \text{ and } v \in V_x\}.$$

In other words, H_x is the subgraph of G whose vertex set consists of all the vertices of $\eta^{-1}(x)$ and their neighbors that are mapped to a node of $\text{tail}[x]$ by η . Among the edges with endpoints in this vertex set we keep only those whose at least one endpoint belongs to $\eta^{-1}(x)$.

The following statement is a generalization of Theorem 1.2, where the identifiers come from a larger codomain and we assert a stronger isolation property. The proof follows from a straightforward adjustment of the proof of Theorem 1.2, so we only sketch it.

Theorem 4.19. Let G be a graph with n vertices and let $\text{id} : E(G) \rightarrow \{1, \dots, Z\}$ be an injective function. Choose prime numbers $p_0, \dots, p_{\log n}$ independently at random so that p_i is chosen uniformly among the primes in the range $\{1, \dots, M_i\}$, where $M_i = 2^{C(\log Z + \log n + 2^i)}$ for some large constant C . Then with probability at least $1 - \frac{1}{Z \cdot \rho(n)}$, for all configurations $c \in \text{conf}(V(G))$ we have $|\text{Min}(\omega, G, c)| \leq 1$, where ω is defined as in Subsection 4.3 and $\rho(n)$ is any fixed polynomial.

Proof. The proof follows the exact same reasoning as the proof of Theorem 1.2. The induction hypothesis then becomes:

Induction hypothesis. With probability at least $\left(1 - \frac{1}{Z \cdot n \rho(n)}\right)^{i+1}$, for all intervals $G\langle s, t, s', t' \rangle$ such that $t - s \leq 2^i$ and $t' - s' \leq 2^i$ and for each configuration $c \in \text{conf}(V\langle s, t, s', t' \rangle)$, there is at most one minimum weight (with respect to ω_i) compliant partial solution, i.e. $|\text{Min}(\omega_i, G\langle s, t, s', t' \rangle, c)| \leq 1$.

The only major change is that after replacing the codomain of the identifier function with $\{1, \dots, Z\}$, we now have $\{x_S : S \in Y\} \subseteq \{1, \dots, 2^{Z+1}\}$ instead of $\{x_S : S \in Y\} \subseteq \{1, \dots, 2^{n^2+1}\}$, and therefore we need to choose each prime p_i among primes in the range $\{1, \dots, 2^{C(\log Z + \log n + 2^i)}\}$. Hence, the success probability accordingly also changes. Note that the constant C will need to be larger, but will remain a constant. Finally, similarly as argued in the proof of Theorem 1.2, the probability that for each configuration $c \in \text{conf}(V(G))$ we have $|\text{Min}(\omega, G, c)| \leq 1$ is at least $\left(1 - \frac{1}{Z \cdot n \rho(n)}\right)^{\log n} \geq \left(1 - \frac{1}{Z \cdot \rho(n)}\right)$. \square

We now use Theorem 4.19 to argue the following.

Lemma 4.20. *Assuming C is chosen large enough, the following event happens with probability at least $1 - \frac{1}{n^2}$: for all $x \in V(F)$ and all $c \in \text{conf}(V_x)$, we have $|\text{Min}(\omega, H_x, c)| \leq 1$.*

Proof. Apply Theorem 4.19 on each graph H_x for $x \in V(F)$, where each time we let the identifier function on E_x be the identifier function on $E(G)$, restricted to E_x . Note that $|V_x| \leq \mathcal{O}(n^\alpha)$ for each x , for large enough C hence we can use the primes $p_0, p_1, \dots, p_{\log n^\alpha + C}$ in each of these applications, and therefore obtain the weights function $\omega_0, \dots, \omega_{\log n^\alpha}$ defined in the same way on all the graphs H_x . By choosing C appropriately large we can guarantee that for every fixed $x \in V(F)$, with probability at least $1 - \frac{1}{n^4}$ we have $|\text{Min}(\omega, H_x, c)| \leq 1$ for all $c \in \text{conf}(V_x)$. Since $|V(F)| \leq n^2$, it now follows from union bound that this assertion holds for all $x \in F$ simultaneously with probability at least $1 - \frac{1}{n^2}$. \square

4.5.4 Isolation partial solutions in subtrees

Our goal now is to extend the conclusion of Lemma 4.20 from graphs H_x that are associated with single nodes x of F to graphs G_x that reflect the whole subtree of F comprising the descendants of x .

Definition 4.21 (Graph G_x). *For each node x in F , define $G_x = (V_x^\downarrow, E_x^\downarrow)$, where*

$$V_x^\downarrow := N_G[\eta^{-1}(\text{subtree}[x])] \quad \text{and} \quad E_x^\downarrow := \{uv : uv \in E(G), u \in \eta^{-1}(\text{subtree}[x]), \text{ and } v \in V_x^\downarrow\}.$$

In other words, G_x is a subgraph of G , but now its vertex set comprises all the vertices that are mapped to nodes of $\text{subtree}[x]$ by η and their neighbors. Among edges with both endpoints in this vertex set we keep only those with at least one endpoint in $\eta^{-1}(\text{subtree}[x])$. Observe that actually,

$$V_x^\downarrow = \bigcup_{y \in \text{subtree}[x]} V_y \quad \text{and} \quad E_x^\downarrow = \bigcup_{y \in \text{subtree}[x]} E_y.$$

Also, denote

$$W_x := N_G(\eta^{-1}(\text{subtree}[x])).$$

As explained before, we will perform a bottom-up induction on F to prove that for each node x , the relevant partial solutions in the graph G_x are appropriately isolated. This will be done under that condition that the event A described in Lemma 4.20 holds: for all $x \in V(F)$ and all $c \in \text{conf}(V_x)$, we have $|\text{Min}(\omega, H_x, c)| \leq 1$. Formally we will prove the following induction hypothesis for all $i = \log n, \dots, 0$, starting with $i = \log n$ and decreasing i at each step.

Induction hypothesis. Conditioned on A happening, the following event happens with probability at least $(1 - \mathcal{O}(\frac{1}{n^2}))^{\log n - i}$: for all nodes x at level i in F and for any configuration $c \in \text{conf}(W_x)$, we have $|\text{Min}(\xi_i, G_x, c)| \leq 1$.

Note that if the induction hypothesis is true for $i = 0$, that is, for the unique root node r of F , then ξ isolates the family of all Hamiltonian cycles in $G_r = G$ with probability at least $(1 - \frac{1}{n^2}) \cdot (1 - \frac{1}{n^2})^{\log n} \geq 1 - \frac{1}{n}$; here, the first factor is the lower bound on the probability of A provided by Lemma 4.20. Therefore, it remains to perform the induction.

Base step. For $i = \log n$, every node x of F at level $\log n$ is a leaf with $|\eta^{-1}(x)| = 1$, say $\eta^{-1}(x) = \{v\}$. Hence we have to prove that for any configuration $c \in \text{conf}(N_G(v))$, we have $|\text{Min}(\xi_{\log n}, G_x, c)| \leq 1$. Notice that G_x only contains edges between v and its neighbors, hence for every configuration $c \in \text{conf}(N_G(v))$ there is at most one partial solution in G_x that is compliant with c . So for $i = \log n$ the induction hypothesis is true.

Induction step. Assume the induction hypothesis is true for all nodes x at level $i + 1$. Let

$$Y' := \bigcup_{y: \text{ node at level } i+1} \bigcup_{c \in \text{conf}(W_y)} \text{Min}(\xi_{i+1}, G_y, c)$$

be the set of all relevant minimum weight partial solutions in graphs G_x for x at level $i + 1$. Furthermore, let

$$Z := \bigcup_{x: \text{ node at level } i} \bigcup_{c \in \text{conf}(V_x)} \text{Min}(\omega, H_x, c)$$

be the set of all minimum weight compliant partial solutions for configurations on graphs H_x for x at depth i . Finally, let

$$Y := \{ S_0 \cup S_1 \cup \dots \cup S_7 : S_0 \in Z \text{ and } S_1, \dots, S_7 \in Y' \}$$

be the set comprising all combinations of 7 partial solutions from Y' and a single partial solution from Z .

We first prove with Claim 4.22 that every relevant minimum weight compliant partial solution is included in Y . Then Claim 4.23 says that with high probability, all partial solutions in Y have pairwise different weights with respect to ξ_i . Hence, proving these two claims is sufficient to make the induction step go through.

Claim 4.22. *Let x be a node of depth i and let $c \in \text{conf}(W_x)$. Then $\text{Min}(\xi_i, G_x, c) \subseteq Y$.*

Proof. Take any $S \in \text{Min}(\xi_i, G_x, c)$. Let y_1, \dots, y_t ($t \leq 7$) be the (at most) seven child nodes of x at depth $i + 1$. Let $S_0 := S \cap E_x$ and $S_j := S \cap E_{y_j}^\downarrow$ for $j \in \{1, \dots, t\}$. Note that the partial solutions S_0, S_1, \dots, S_t are pairwise disjoint and their union is equal to S . Further, since $S \in \text{Min}(\xi, G_x, c)$, an argument analogous to the one used in the proof of Claim 4.8 shows that

- $S_0 \in \text{Min}(\omega, H_x, c_0)$ for some $c_0 \in \text{conf}(V_x)$; and
- for each $j \in \{1, \dots, t\}$, $S_j \in \text{Min}(\xi_{i+1}, G_{y_j}, c_j)$ for some $c_j \in \text{conf}(W_{y_j})$.

This means that $S_0 \in Z$ and $S_1, \dots, S_t \in Y'$, implying that $S \in Y$. □

Claim 4.23. *Conditioned on A , the probability of the following event is at least $(1 - \frac{1}{n^2})^{\log n - i}$: for all different $S, S' \in Y$, we have $\xi_i(S) \neq \xi_i(S')$.*

Proof. For each $S \in Y$ let us define

$$x_S := \sum_{e \in S} 2^{\text{id}(e)}.$$

Observe that since the identifiers assigned to the edges are unique, the numbers x_S are pairwise different. The induction hypothesis gives us that conditioned on A , the probability of the following event is at least $(1 - \frac{1}{n^2})^{\log n - (i+1)}$: for every node y of F at level $i + 1$ and every $c \in \text{conf}(W_y)$, we have $|\text{Min}(\xi_{i+1}, G_y, c)| \leq 1$. We may then use Theorem 4.6 to conclude that for each such y ,

$$\left| \bigcup_{c \in \text{conf}(W_y)} \text{Min}(\xi_{i+1}, G_y, c) \right| \leq 2^{\mathcal{O}(|W_y|)} \leq 2^{\mathcal{O}((n/2^{i+1})^\alpha)}.$$

Since we assume that A happens, we have $|\text{Min}(\omega, H_x, c)| \leq 1$ for each node x at level i . We can use Theorem 4.6 again to infer that for each such x ,

$$\left| \bigcup_{c \in \text{conf}(V_x)} \text{Min}(\omega, H_x, c) \right| \leq 2^{\mathcal{O}(|V_x|)} \leq 2^{\mathcal{O}((n/2^i)^\alpha)}.$$

Since F has at most n^2 nodes in total, the above bounds imply that

$$|Y| \leq |Z| \cdot |Y'|^7 \leq 2^{\mathcal{O}((n/2^i)^\alpha)} \cdot n^{16}.$$

Since $N_i = 2^{C(\log n + (n/2^i)^\alpha)}$, and prime q_i is sampled uniformly at random among the primes in the range $\{1, \dots, N_i\}$, from Lemma 3.1 we can conclude that, provided C is chosen large enough, with probability at least

$$\left(1 - \frac{(n^2 + 1) \left(2^{\mathcal{O}((n/2^i)^\alpha)} \cdot n^{16}\right)^2}{2^{(C/2)(\log n + (n/2^i)^\alpha)}}\right) \cdot \left(1 - \frac{1}{n^2}\right)^{\log n - (i+1)} \geq \left(1 - \frac{1}{n^2}\right)^{\log n - i},$$

all the numbers in $\{x_S : S \in Y\}$ have pairwise different remainders modulo q_i . As a consequence, with at least the same probability we have $\xi_i(S) \neq \xi_i(S')$ for all distinct $S, S' \in Y$. \square

As argued, the induction step follows directly from combining Claims 4.22 and 4.23.

5 Deterministic algorithm for Hamiltonian cycle in separable classes

A graph class \mathcal{C} shall be called *efficiently separable* with degree α if it is separable with degree α in the sense of Definition 4.13, and moreover given $G \in \mathcal{C}$ and a vertex subset $S \subseteq V(G)$, a suitable balanced separator X witnessing separability can be computed in polynomial time. In this section we prove the following result.

Theorem 5.1. *Let \mathcal{C} be a hereditary graph class that is efficiently separable with degree $\alpha \in (0, 1)$. Then there is an algorithm for HAMILTONIAN CYCLE on graphs from \mathcal{C} that runs in deterministic time $2^{\mathcal{O}(n^\alpha)}$ and uses polynomial space.*

It is well-known that for every fixed H , the class of H -minor-free graphs is efficiently separable with degree $\frac{1}{2}$ [4, 38]. Hence, Theorem 5.1 implies Theorem 1.5.

The first step towards the proof of Theorem 5.1 is to revisit the approach of [44] and extend it to obtain the following result.

Lemma 5.2. *There is a deterministic algorithm that takes as input an undirected graph $G = (V, E)$ along with an elimination forest of height at most d , a weight function $\omega : E \rightarrow \{1, \dots, W\}$, and a target integer t . The algorithm runs in time $5^d W (n \log W)^{\mathcal{O}(1)}$, uses space that is polynomial in n and $\log W$, and detects whether G has a Hamiltonian cycle C satisfying $\omega(C) = t$, provided there is at most one such C .*

The extension is similar to that performed by Björklund in [7], where he extended his $\mathcal{O}(1.66^n)$ -time algorithm for HAMILTONIAN CYCLE to an $1.66^n W (n \log W)^{\mathcal{O}(1)}$ time algorithm for the TRAVELING SALESMAN problem on n cities. Therefore, we sketch here the extension assuming (but recalling) the basic understanding of the approach of [44].

Definition 5.3. *Suppose \mathbb{F} is a finite field. An element $\rho \in \mathbb{F}$ is a primitive N -root of unity if $\rho^N = 1$ and for every $0 < N' < N$ it holds that $\rho^{N'} \neq 1$.*

It is well known that the multiplicative group of every finite field is cyclic, that is, there is a generator $g \in \mathbb{F}$ such that $\{g^0, g^1, \dots, g^{|\mathbb{F}|-1}\}$ are all the elements of field. Then we must have $g^{-1} = g^{|\mathbb{F}|-1}$. So $g^{|\mathbb{F}|} = 1$, and therefore g is a primitive $(|\mathbb{F}| - 1)$ -root of unity. We will work with the field \mathbb{F}_p , for some prime p . First we address the issue of finding a generator of \mathbb{Z}_p^* , the multiplicative group of \mathbb{F}_p :

Lemma 5.4. *A generator of \mathbb{Z}_p^* can be found in $\mathcal{O}(p \cdot \text{polylog}(p))$ deterministic time and $\mathcal{O}(\text{polylog}(p))$ space.*

Proof. First, find the prime factors p_1, \dots, p_ℓ using any deterministic $\mathcal{O}(p)$ -time algorithm. Note that we have $\ell \leq \log p$. Next we rely on the well-known fact that an element e is a generator of \mathbb{Z}_p^* if and only if for every $i = 1, \dots, \ell$ it holds that $e^{(p-1)/p_i} \not\equiv 1 \pmod{p}$. This fact follows from Lagrange's theorem (see for example the discussion preceding [42, Theorem 14.16], where this fact was used in a similar way to find generators probabilistically). By this fact, we can check whether e is a generator or not in $\mathcal{O}(\text{polylog}(p))$ time. Thus we can find a generator by simply iterating over all elements $e \in \mathbb{Z}_p$ until the check succeeds. \square

We will use the following well-known statement about discrete Fourier transform in finite fields (see e.g. [13, Equation 30.11]).

Theorem 5.5 (Discrete Fourier Inversion). *Let \mathbb{F} be a finite field, let $\rho \in \mathbb{F}$ be a primitive N -root of unity, and let $P(x) \in \mathbb{F}[x]$ be a polynomial of degree at most $N - 1$ in x with coefficients from \mathbb{F} . If $P(x) = \sum_{i=0}^{N-1} c_i x^i$, then for every $0 \leq t \leq N - 1$ it holds that*

$$c_t = \frac{1}{N} \sum_{i=0}^{N-1} \rho^{-it} P(\rho^i).$$

Let $N := W \cdot n \cdot \log n$ and consider the field $\mathbb{F} := \mathbb{F}_p$, where p is a prime satisfying $p \in \Theta(N)$. Such a prime can be deterministically found in time $\mathcal{O}(p)$ and using $\text{polylog}(p)$ space using brute-force and the polynomial-time deterministic primality testing algorithm [3]. By the above discussion, the field \mathbb{F} has a $(p - 1)$ -root of unity $\rho \in \mathbb{F}$, and such a root can be found in time $\mathcal{O}(p \cdot \text{polylog}(p))$ and space $(n \log W)^{\mathcal{O}(1)}$. Next, we continue with analysis of methods presented in [44].

Recall that we are given a graph G and an elimination forest F of G of height at most d . Since we are interested in Hamiltonian cycles in G , we may assume that G is connected, hence F is a tree. The central objects studied in [44] are polynomials $P[u, f], P(u, f) \in \mathbb{Z}[\alpha, \beta, \gamma]$, defined for each $u \in V$ and function $f: \text{tail}[u] \rightarrow \{0, 1_L, 1_R, 2_L, 2_R\}$ (or $f: \text{tail}(u) \rightarrow \{0, 1_L, 1_R, 2_L, 2_R\}$ in case of $P(u, f)$). Here, α, β, γ are formal variables. In [44] it is shown how to compute those polynomials in a bottom-up manner over the given elimination forest F . Further, the parity of the number of Hamiltonian cycles of total weight w can be inferred from the coefficient of the monomial $\alpha^w \beta^n \gamma^n$ in the polynomial $P(r, \emptyset)$, where r is the root of F . Therefore, the idea in [44] was to use Isolation Lemma to ensure that provided the graph is Hamiltonian, with high probability there exists $w \in [N]$ for which there is exactly one Hamiltonian cycle of weight w . Then one explicitly computes all the polynomials $P[u, f], P(u, f)$ in a bottom-up manner over the tree F in time $5^d W (n \log W)^{\mathcal{O}(1)}$. Finally, the existence of a Hamiltonian cycle can be inferred from the analysis of the coefficients of $P(r, \emptyset)$.

In the setting of Lemma 5.2 we can almost use the same strategy, however there is a caveat. Namely, the expansion of each polynomial $P(u, f)$ and $P[u, f]$ into a sum of monomials of the form $\alpha^a \beta^b \gamma^c$ may have length as large as $(N + 1)(n + 1)^2$, because the relevant values of a, b, c are $a \in \{0, \dots, N\}$ and $b, c \in \{0, \dots, n\}$. Therefore, storing the coefficients of $P[u, f]$ explicitly would take at least space $\mathcal{O}(Nn^2) = \mathcal{O}(Wn^3)$, which is more than $(n + \log W)^{\mathcal{O}(1)}$ promised in the statement of Lemma 5.2.

Therefore, the idea is not to compute the whole polynomial $P(r, \emptyset)$ explicitly, but evaluate the relevant coefficients of $P(r, \emptyset)$ one by one using Theorem 5.5. Precisely, let $Q = \sum_{w=0}^N c_{w,n,n} \cdot \alpha^w \in \mathbb{Z}[\alpha]$, where $c_{w,n,n}$ is the coefficient of $\alpha^w \beta^n \gamma^n$ in $P(r, \emptyset)$. After casting Q as a polynomial $Q' \in \mathbb{F}_p[\alpha]$, we can use the method presented in [44] to give an algorithm that evaluates $Q'(e)$ for a given $e \in \mathbb{F}_p$ in time $5^d W (n \log W)^{\mathcal{O}(1)}$ and using $(n \log W)^{\mathcal{O}(1)}$ space, because storing an element of \mathbb{F}_p requires $(n \log W)^{\mathcal{O}(1)}$ space. This is enough to compute the formula described in Theorem 5.5 within the promised

complexity guarantees. This concludes the description of how to compute the coefficient $c'_{t,n,n}$ of $Q'[\alpha]$ within the stated resource bounds.

Now we can compute the matching coefficient $c_{t,n,n}$ of $Q[\alpha]$ by observing that $c'_{t,n,n} = c_{t,n,n} \pmod p$, applying the above step $\Theta(n)$ times for different primes p , and reconstructing $c_{t,n,n}$ with the Chinese Remainder Theorem. Here, it is important to note that the coefficient $c_{t,n,n}$ is of the order $2^{\mathcal{O}(n)}$, hence the information about $c_{t,n,n}$ modulo $\Theta(n)$ different primes is sufficient to reconstruct $c_{t,n,n}$ completely. This concludes the sketch of the proof of Lemma 5.2.

We now use Lemma 5.2 to prove Theorem 5.1.

Proof of Theorem 5.1. Let $G = (V, E) \in \mathcal{C}$ be the given graph, where $n := |V|$. By iteratively extracting balanced separators (cf., [44, Theorem A.1]) we can compute an elimination forest of G of height $\mathcal{O}(n^\alpha)$ in polynomial time.

Next we use Theorem 4.14 that gives us a set of weight functions $\omega_1, \dots, \omega_\ell: E \rightarrow [2^{\mathcal{O}(n^\alpha)}]$ with $\ell = 2^{\mathcal{O}(n^\alpha)}$ such that at least half of the functions ω_i isolate the family of Hamiltonian cycles in G .

It can be easily seen by inspecting the construction of the isolation scheme of Theorem 4.14 that the functions $\omega_1, \dots, \omega_\ell$ can be enumerated one by one using polynomial working space and $2^{\mathcal{O}(n^\alpha)}$ time. Namely, we simply need to iterate over every tuple of $\log n + 1$ primes $p_0 \in [M_0], \dots, p_{\log n} \in [M_{\log n}]$. To achieve that, we can iterate over every prime number $p_i \in [M_i]$ in time $M_i^{\mathcal{O}(1)}$ (just iterate through $[M_i]$ and deterministically check for primality in $M_i^{\mathcal{O}(1)}$ time with a brute-force algorithm). Therefore, enumerating all weight functions $\omega_1, \dots, \omega_\ell$ can be done in $\text{poly}(M_1 \dots M_{\log n}) \leq 2^{\mathcal{O}(n^\alpha)}$ additional time and polynomial space. Theorem 4.14 guarantees that among the enumerated weight functions $\omega_1, \dots, \omega_\ell$, there is at least one (and even half of them) that isolates the family of Hamiltonian cycles in G .

Therefore, it remains to apply the algorithm of Lemma 5.2 to each consecutive function ω_i and each possible minimum weight $t \leq 2^{\mathcal{O}(n^\alpha)}$, and report a positive outcome if any of these applications finds a Hamiltonian cycle in G . The time complexity is bounded by $2^{\mathcal{O}(n^\alpha)}$ and the space complexity is polynomial in n and $\log 2^{\mathcal{O}(n^\alpha)} = \mathcal{O}(n)$. \square

6 MSO-definable problems

6.1 Definitions

CMSO₂ Logic. We work with problems definable in logic CMSO₂, which stands for Monadic Second-Order logic on graphs with modular counting predicates and quantification over edge subsets. Recall that in this logic we have variables for individual vertices, individual edges, sets of vertices, and sets of edges; the latter two kinds are called *monadic variables*. The basic constructs in CMSO₂ are *atomic formulas* of the following forms:

- *Equality:* $x = y$, checking equality of x and y ;
- *Membership:* $x \in X$, checking that x belongs to X ;
- *Incidence:* $\text{inc}(u, e)$, checking that vertex u is incident on the edge e ; and
- *Congruence:* $|X| \equiv a \pmod p$, where a, p are constants, with the expected semantics.

CMSO₂ formulas can be constructed from atomic formulas using standard boolean connectives, negation, and quantification over variables of each of the four kinds (both existential and universal). Note that a CMSO₂ formula can have *free variables* that are not bound by any quantification. A formula can be applied on a graph supplied with a valuation of the free variables. For example, the formula

$$\varphi(X) = [\forall_u \forall_v (u \in X \wedge v \in X \wedge u \neq v) \implies (\neg \exists_e \text{inc}(u, e) \wedge \text{inc}(v, e))], \quad (4)$$

when applied on a graph G and a vertex subset A , checks whether A is an independent set in G . If this is the case, we write $G \models \varphi(A)$.

Let $\varphi(X)$ be a CMSO_2 formula with one free vertex set variable X . For a graph G , we define

$$\text{select}_\varphi(G) := \{S \subseteq V(G) \mid G \models \varphi(S)\}.$$

For example, if $\varphi(X)$ is the formula presented in (4), then $\text{select}_\varphi(G)$ consists of all independent sets in G . If X is an edge set variable, then $\text{select}_\varphi(G)$ is defined analogously: it comprises all subsets S of edges of G for which $\varphi(S)$ is satisfied. Thus, if $\varphi(X)$ is a CMSO_2 formula with a free monadic variable X , then select_φ is a vertex or edge selection problem, depending on whether X is a vertex set or an edge set variable. A vertex/edge selection problem is CMSO_2 -*definable* if it is of the form select_φ for a formula $\varphi(X)$ as above.

Boundaried graphs. Throughout this section we assume that all considered graphs have vertices from a fixed countable set Ω . The reader may think that $\Omega = \mathbb{N}$.

A *boundaried graph* is a pair consisting of a graph G and a subset of its vertices B , called the *boundary*. We have two natural operations on boundaried graphs:

- Suppose $\mathbf{G}_1 = (G_1, B_1)$ and $\mathbf{G}_2 = (G_2, B_2)$ are boundaried graphs such that $V(G_1) \cap V(G_2) \subseteq B_1 \cup B_2$. Then the *sum* of \mathbf{G}_1 and \mathbf{G}_2 is the boundaried graph

$$\mathbf{G}_1 \oplus \mathbf{G}_2 := (G_1 \cup G_2, B_1 \cup B_2),$$

where $G_1 \cup G_2 = (V(G_1) \cup V(G_2), E(G_1) \cup E(G_2))$. Note that the sum is not defined if the condition $V(G_1) \cap V(G_2) \subseteq B_1 \cup B_2$ does not hold.

- Suppose $\mathbf{G} = (G, B)$ is a boundaried graph and $A \subseteq B$. Then the operation of *forgetting* A in \mathbf{G} yields the boundaried graph

$$\text{forget}_A(\mathbf{G}) := (G, A).$$

Note that for notational convenience, the set indicated in the subscript is the new boundary set, and not the set $B \setminus A$ of vertices that get forgotten, i.e., removed from the boundary.

The following standard lemma connects the operations on boundaried graphs with the concept of treewidth.

Lemma 6.1. *A graph G has treewidth at most k if and only if (G, \emptyset) can be obtained from graphs on at most $k + 1$ vertices by a repeated use of the sum and forget operations, where at each moment in the construction all boundaried graphs have boundaries of size at most $k + 1$.*

Configuration schemes. We now present the formalism of configuration schemes for selection problems. The notational layer is directly taken from the recent work of Chen et al. [12]. However, in general, the algebraic approach to graphs of bounded treewidth and recognizability of their properties dates back to the foundational work of Courcelle and others done in the 90s. See the book of Courcelle and Engelfriet for an introduction to the area [14].

For concreteness we focus on edge selection problems. Adjusting the definitions to vertex selection problems is straightforward.

A *configuration scheme* is a pair of functions (conf, c) with the following properties:

- conf assigns to each finite subset $B \subseteq \Omega$ a finite *configuration set* $\text{conf}(B)$. We require that the cardinality of the configuration set is uniformly and effectively bounded in the size of B , that is, there exists a computable function g such that $|\text{conf}(B)| \leq g(|B|)$ for each finite $B \subseteq \Omega$.
- For every boundaried graph $\mathbf{G} = (G, B)$ and a subset of edges $S \subseteq E(G)$, c maps the pair (\mathbf{G}, S) to a configuration $c(\mathbf{G}, S) \in \text{conf}(B)$.

We say that a configuration scheme (conf, c) is *compositional* if the following two conditions hold:

- For every pair of boundaried graphs $\mathbf{G}_1 = (G_1, B_1)$ and $\mathbf{G}_2 = (G_2, B_2)$ (with defined sum), and subsets of edges $S_1 \subseteq E(G_1)$ and $S_2 \subseteq E(G_2)$, the configuration

$$c(\mathbf{G}_1 \oplus \mathbf{G}_2, S_1 \cup S_2)$$

depends only on the pair of configurations

$$c(\mathbf{G}_1, S_1) \quad \text{and} \quad c(\mathbf{G}_2, S_2).$$

- For every boundaried graph $\mathbf{G} = (G, B)$, $A \subseteq B$, and a subset of edges $S \subseteq E(G)$, the configuration

$$c(\text{forget}_A(\mathbf{G}), S)$$

depends only on the configuration

$$c(\mathbf{G}, S).$$

In other words, the first condition means that we can endow the set $\text{conf}(B_1) \times \text{conf}(B_2)$ with a sum operation \oplus so that the operators \oplus and $c(\cdot, \cdot)$ commute. Similarly, the second condition means that $\text{conf}(B)$ can be endowed with an operation $\text{forget}_A(\cdot)$ so that the operators $\text{forget}_A(\cdot)$ and $c(\cdot, \cdot)$ commute. We will therefore use the operators \oplus and forget_A also as operators defined on configuration sets provided by conf . Note here that since \oplus is commutative on boundaried graphs, the corresponding operator \oplus on configurations can also be chosen to be commutative.

Suppose \mathcal{P} is an edge selection problem, that is, a function that with each graph G associates a family of edge subsets $\mathcal{P}(G) \subseteq 2^{E(G)}$. We say that a configuration scheme (conf, c) *recognizes* \mathcal{P} if there exists a set of *final configurations* $F \subseteq \text{conf}(\emptyset)$ such that for every graph G and a subset of edges $S \subseteq E(G)$,

$$S \in \mathcal{P}(G) \quad \text{if and only if} \quad c((G, \emptyset), S) \in F.$$

The next lemma follows from well-known compositionality properties of the CMSO₂ logic. We sketch the proof for completeness, but we remark that an essentially the same sketch was also provided in [12].

Lemma 6.2. *Every CMSO₂-definable edge selection problem is recognized by a compositional configuration scheme.*

Proof sketch. Let $\mathcal{P} = \text{select}_\varphi$ be the problem in question, where $\varphi(X)$ is a CMSO₂ formula with a free edge set variable X . Let q be the quantifier rank of φ , that is, the maximum number of nested quantifiers in φ .

Consider a finite set $B \subseteq \Omega$ and all CMSO₂ formulas $\psi(X)$ of quantifier rank at most q , where X is an edge set variable, which can also use the elements of B as constants (formally, the signature additionally contains every element of B as a constant). It is well-known that there are only finitely many pairwise non-equivalent such formulas, in the sense that ψ and ψ' are equivalent if for every boundaried graph $\mathbf{G} = (G, B)$ and $S \subseteq E(G)$, we have $\mathbf{G} \models \psi(S)$ if and only if $\mathbf{G} \models \psi'(S)$. Moreover, the number of equivalence classes is bounded by a computable function of $|B|$. Then, let $\text{Formulas}^q(B)$ be a set comprised of one arbitrarily selected representative from each equivalence class.

We now define the configuration scheme (conf, c) as follows:

- For each finite $B \subseteq \Omega$, we set

$$\text{conf}(B) := 2^{\text{Formulas}^q(B)},$$

that is, $\text{conf}(B)$ is the powerset of $\text{Formulas}^q(B)$.

- For each boundaried graph $\mathbf{G} = (G, B)$ and an edge subset $S \subseteq E(G)$, we set

$$c(\mathbf{G}, S) := \{\psi(X) \in \text{Formulas}^q(B) \mid \mathbf{G} \models \psi(S)\}.$$

A standard argument involving Ehrenfeucht-Fraïsse games shows that this configuration scheme is compositional. It remains to observe that (conf, c) recognizes select_φ by taking

$$F := \{\Delta \subseteq \text{Formulas}^q(\emptyset) \mid \varphi \in \Delta\}. \quad \square$$

6.2 Isolation

The remainder of this section is devoted to the proof of Theorem 1.6. We remark that the same technique can be used to prove the same result also for vertex selection problems. We leave the straightforward modifications to the reader.

By Lemma 6.2, \mathcal{P} is recognized by a compositional configuration scheme (conf, c) . Let $F \subseteq \text{conf}(\emptyset)$ be the set of final configurations, as in the definition of recognizing \mathcal{P} . Recall that there exists a computable function $g: \mathbb{N} \rightarrow \mathbb{N}$ such that

$$|\text{conf}(B)| \leq g(|B|) \quad \text{for each finite } B \subseteq \Omega.$$

We may assume that the configuration scheme (conf, c) satisfies the following assertion for every boundaried graph $\mathbf{G} = (G, B)$ and $S, S' \subseteq E(G)$:

$$\text{if } c(\mathbf{G}, S) = c(\mathbf{G}, S'), \quad \text{then } S \cap \binom{B}{2} = S' \cap \binom{B}{2}. \quad (5)$$

Indeed, if assertion (5) is not satisfied, then we can instead use the configuration scheme (conf', c') defined as

$$\text{conf}'(B) := \text{conf}(B) \times 2^{\binom{B}{2}} \quad \text{and} \quad c'(\mathbf{G}, S) = \left(c(\mathbf{G}, S), S \cap \binom{B}{2} \right),$$

which already satisfies assertion (5). Note here that (conf', c') still recognizes \mathcal{P} (by taking $F' := F \times \{\emptyset\}$) and $|\text{conf}'(B)|$ is still bounded by a computable function of $|B|$. We remark, however, that the configuration scheme provided by Lemma 6.2 actually already satisfies (5), provided the quantifier rank q is positive.

To prove Theorem 1.6, it suffices to give an isolation scheme for \mathcal{P} on graphs of treewidth at most k that uses at most $\mathcal{O}(\log \Gamma \log n + \log^2 n)$ random bits, where

$$\Gamma := \max \left(g(2k+2)^5, 2^{\binom{2k+2}{2}} \right),$$

and assigns weights that are at most exponential in the number of random bits. From now on let us fix k , the graph G given to the isolation scheme on input, and a tree decomposition $\mathbb{T} = (T, \beta)$ of G of width at most k . We may assume that T has at most n nodes, where $n := |V(G)|$ is the vertex count of G .

Isolation scheme. We first present the isolation scheme, which is just a general version of the scheme for HAMILTONIAN CYCLE presented in Section 4.4 (but without the refinement of using the rank-based approach). First, let $\text{id}: E(G) \rightarrow \{0, \dots, |E(G)| - 1\}$ be any bijection that assigns to each edge $e \in E(G)$ its unique *identifier* $\text{id}(e)$. Then choose $1 + \log n$ primes

$$p_0, p_1, \dots, p_{\log n} \in \{1, \dots, M\}$$

uniformly and independently at random among primes in $\{1, \dots, M\}$, where

$$M := \Gamma^4 \cdot n^{14}.$$

Note that selecting one prime p_i requires $\mathcal{O}(\log M) = \mathcal{O}(\log \Gamma + \log n)$ random bits, hence we have used $\mathcal{O}(\log \Gamma \log n + \log^2 n)$ random bits in total.

Next, we inductively define weight functions $\omega_0, \omega_1, \dots, \omega_{\log n}$ on $E(G)$ as follows:

- Set $\omega_0(e) = 2^{\text{id}(e)} \bmod p_0$ for all $e \in E(G)$.
- For $e \in E(G)$ and $i = 1, 2, \dots, \log n$, set

$$\omega_i(e) := Mn^2 \cdot \omega_{i-1}(e) + \left(2^{\text{id}(e)} \bmod p_i\right).$$

Let

$$\omega := \omega_{\log n}$$

and observe that ω assigns weights upper bounded by $2 \cdot (Mn^2)^{\log n} = 2^{\mathcal{O}(\Gamma \log n + \log^2 n)}$, as required.

It remains to verify that ω isolates $\mathcal{P}(G)$ with high probability. We do this by induction. Recall that a *segment* in the tree T is a connected subtree I of T such that the boundary ∂I – the set of nodes of I with neighbors outside of I – has size at most 2. For a segment I we define the *boundaried graph*

$$\mathbf{G}\langle I \rangle := \left(G \left[\bigcup_{x \in I} \beta(x) \right], \bigcup_{x \in \partial I} \beta(x) \right).$$

Note that this definition is slightly different than the one used in Section 4, as $\mathbf{G}\langle I \rangle$ contains all the edges of G with both endpoints in $\bigcup_{x \in I} \beta(x)$. We will also write $\partial \mathbf{G}\langle I \rangle := \bigcup_{x \in \partial I} \beta(x)$ for the boundary of $\mathbf{G}\langle I \rangle$. For $\gamma \in \text{conf}(\partial \mathbf{G}\langle I \rangle)$, let us define the set of *partial solutions* in $\mathbf{G}\langle I \rangle$ that yield configuration γ as:

$$\mathcal{S}(I, \gamma) := \{S \subseteq E(\mathbf{G}\langle I \rangle) \mid c(\mathbf{G}\langle I \rangle, S) = \gamma\}.$$

We will prove the following statement by induction on i .

Claim 6.3. *For each $i \in \{0, 1, 2, \dots, \log n\}$ and each segment I of T with at most 2^i edges, the probability that both of the following events happen is at least $1 - \frac{8^i}{n^5}$:*

- ω_i isolates $\mathcal{S}(I, \gamma)$ for each $\gamma \in \text{conf}(\partial \mathbf{G}\langle I \rangle)$ for which $\mathcal{S}(I, \gamma) \neq \emptyset$; and
- $\min \omega_i(\mathcal{S}(I, \gamma)) \neq \min \omega_i(\mathcal{S}(I, \gamma'))$ for all $\gamma, \gamma' \in \text{conf}(\partial \mathbf{G}\langle I \rangle)$ such that $\gamma \neq \gamma'$, $\mathcal{S}(I, \gamma) \neq \emptyset$, and $\mathcal{S}(I, \gamma') \neq \emptyset$.

Note that

$$\mathcal{P}(G) = \bigcup_{\gamma \in F} \mathcal{S}(T, \gamma).$$

Hence, since T has less than $n = 2^{\log n}$ edges, Claim 6.3 for $i = \log n$ and $I = T$ implies that ω isolates $\mathcal{P}(G)$ with probability at least $1 - \frac{8^{\log n}}{n^5} = 1 - \frac{1}{n^2}$. So it remains to prove Claim 6.3, which we do by induction on i .

Base step. For $i = 0$ we have that segment I has at most one edge, so it has at most two nodes. Then $\mathbf{G}\langle I \rangle$ is a bounded graph on at most $2k + 2$ vertices, hence it has at most $\binom{2k+2}{2}$ edges. For each $S \subseteq E(\mathbf{G}\langle I \rangle)$, let

$$x_S = \sum_{e \in S} 2^{\text{id}(e)}.$$

Observe that since the identifiers assigned to edges are pairwise different, the numbers x_S for $S \subseteq E(\mathbf{G}\langle I \rangle)$ are also pairwise different. Since those numbers are upper bounded by $2^{\binom{n}{2}}$ and there are at most $2^{\binom{2k+2}{2}}$ of them, $M \geq 2^{4\binom{2k+2}{2}} \cdot n^{14}$ and p_0 is drawn uniformly at random among primes in the range $\{1, \dots, M\}$, from Lemma 3.1 we can conclude that with probability at least $1 - \frac{1}{n^5}$, all the numbers $\{x_S : S \subseteq E(\mathbf{G}\langle I \rangle)\}$ have pairwise different remainders modulo p_0 . Since $\omega_0(S) \equiv x_S \pmod{p_0}$, this means that with probability at least $1 - \frac{1}{n^5}$, all subsets of $E(\mathbf{G}\langle I \rangle)$ receive pairwise different weights with respect to ω_0 . This implies the conclusion of Claim 6.3 for the segment I .

Induction step. Consider now any $i \geq 1$ and a segment I in T that has more than one but at most 2^i edges. By Lemma 3.3, segment I can be partitioned into at most five segments I_1, \dots, I_t ($t \leq 5$), each with at most 2^{i-1} edges, so that

$$\mathbf{G}\langle I \rangle = \text{forget}_{\partial \mathbf{G}\langle I \rangle} (\mathbf{G}\langle I_1 \rangle \oplus \dots \oplus \mathbf{G}\langle I_t \rangle). \quad (6)$$

By induction assumption, for each segment I_j , $j \in \{1, \dots, 5\}$, the property described in Claim 6.3 holds with probability at least $1 - \frac{8^{i-1}}{n^5}$. By union bound, this property holds for all I_j , $j \in \{1, \dots, 5\}$, simultaneously with probability at least $1 - \frac{5 \cdot 8^{i-1}}{n^5}$. We proceed under the assumption that this is the case; we call this supposition (\star).

We define

$$\Lambda := \text{conf}(\partial \mathbf{G}\langle I_1 \rangle) \times \dots \times \text{conf}(\partial \mathbf{G}\langle I_t \rangle).$$

Consider any $S \subseteq E(\mathbf{G}\langle I \rangle)$. We shall say that S is *compatible* with $(\gamma_1, \dots, \gamma_t) \in \Lambda$ if there exists a partition $\{S_1, \dots, S_t\}$ of S such that

$$S_j \subseteq E(\mathbf{G}\langle I_j \rangle) \quad \text{and} \quad S_j \in \mathcal{S}(I_j, \gamma_j) \quad \text{for each } j \in \{1, \dots, t\}.$$

(Note that S can be simultaneously compatible with several elements of Λ , as edges of S that belong to several of the graphs $\mathbf{G}\langle I_j \rangle$ may be placed in different parts S_j .) For $\bar{\gamma} \in \Lambda$, we denote

$$\mathcal{R}_{\bar{\gamma}} := \{S \subseteq E(\mathbf{G}\langle I \rangle) \mid S \text{ is compatible with } \bar{\gamma}\}.$$

We claim the following.

Claim 6.4. *For each $\bar{\gamma} \in \Lambda$ the weight function ω_{i-1} isolates the family $\mathcal{R}_{\bar{\gamma}}$, provided this family is nonempty.*

Proof. Let $\bar{\gamma} = (\gamma_1, \dots, \gamma_t)$. Suppose there exist two edge subsets S and S' that are both compatible with $\bar{\gamma}$ and have the same minimum weight with respect to ω_{i-1} . Let $\{S_1, \dots, S_t\}$ and $\{S'_1, \dots, S'_t\}$ be partitions of S and S' , respectively, that witness compatibility. Observe that for each $j \in \{1, \dots, t\}$, S_j must be an element of minimum weight with respect to ω_{i-1} in the family $\mathcal{S}(I_j, \gamma_j)$. Indeed, otherwise we could replace S_j with a partial solution $\widehat{S}_j \in \mathcal{S}(I_j, \gamma_j)$ with $\omega_{i-1}(\widehat{S}_j) < \omega_{i-1}(S_j)$, thus obtaining an edge subset $\widehat{S} := (S \setminus S_j) \cup \widehat{S}_j$ of weight smaller than that of S that is also compatible with $\bar{\gamma}$. Note here that assumption (5) implies that $\{S_1, \dots, S_{j-1}, \widehat{S}_j, S_{j+1}, \dots, S_t\}$ is still a partition of S . The same observation applies also to S'_j . Assumption (\star) implies that ω_{i-1} isolates $\mathcal{S}(I_j, \gamma_j)$, which means that $S_j = S'_j$. This applies to each $j \in \{1, \dots, t\}$, hence $S = S'$. \square

For $\gamma \in \text{conf}(\partial \mathbf{G}\langle I \rangle)$, let Λ_γ be the set of all those $\bar{\gamma} = (\gamma_1, \dots, \gamma_t) \in \Lambda$ for which

- $\text{forget}_{\partial \mathbf{G}\langle I \rangle}(\gamma_1 \oplus \dots \oplus \gamma_t) = \gamma$ and
- the family $\mathcal{R}_{\bar{\gamma}}$ is nonempty.

Note that sets $\{\Lambda_\gamma : \gamma \in \text{conf}(\partial \mathbf{G}\langle I \rangle)\}$ are pairwise disjoint. For each $\bar{\gamma} \in \Lambda_\gamma$, let $S_{\bar{\gamma}}$ be the element of $\mathcal{R}_{\bar{\gamma}}$ that has the minimum weight with respect to ω_{i-1} . By Claim 6.4, there is a unique such element. Let now

$$\mathcal{M}_\gamma := \{S_{\bar{\gamma}} \mid \bar{\gamma} \in \Lambda_\gamma \text{ and } \omega_{i-1}(S_{\bar{\gamma}}) = \min\{\omega_{i-1}(S_{\bar{\delta}}) : \bar{\delta} \in \Lambda_\gamma\}\}.$$

In other words, \mathcal{M}_γ comprises sets $S_{\bar{\gamma}}$ of minimum weight with respect to ω_{i-1} among $\{S_{\bar{\delta}} : \bar{\delta} \in \Lambda_\gamma\}$. Let us observe the following.

Claim 6.5. *Suppose $S \in \mathcal{S}(I, \gamma)$ is such that $\omega_i(S) = \min \omega_i(\mathcal{S}(I, \gamma))$. Then $S \in \mathcal{M}_\gamma$.*

Proof. We first observe that

$$\sum_{e \in S} \left(2^{\text{id}(e)} \bmod p_i\right) \leq p_i \cdot |S| < Mn^2. \quad (7)$$

Therefore, by the definition of ω_i , assertion $\omega_i(S) = \min \omega_i(\mathcal{S}(I, \gamma))$ implies also that

$$\omega_{i-1}(S) = \min \omega_{i-1}(\mathcal{S}(I, \gamma)). \quad (8)$$

Consider any partition $\{S_1, \dots, S_t\}$ of S such that $S_j \subseteq E(\mathbf{G}\langle I_j \rangle)$, for each $j \in \{1, \dots, t\}$. Let $\gamma_j := c(\mathbf{G}\langle I_j \rangle, S_j)$. Clearly, partition $\{S_1, \dots, S_t\}$ witnesses that S is compatible with $\bar{\gamma} := (\gamma_1, \dots, \gamma_t)$, hence $S \in \mathcal{R}_{\bar{\gamma}}$. Further, by (6) we infer that $\text{forget}_{\partial \mathbf{G}\langle I \rangle}(\gamma_1 \oplus \dots \oplus \gamma_t) = \gamma$. We conclude that $\bar{\gamma} \in \Lambda_\gamma$. As $\mathcal{R}_{\bar{\gamma}} \subseteq \mathcal{S}(I, \gamma)$, (8) implies that $S = S_{\bar{\gamma}}$. We can now use (8) again to conclude that $S \in \mathcal{M}_\gamma$. \square

Let now

$$\mathcal{M} := \bigcup_{\gamma \in \text{conf}(\partial \mathbf{G}\langle I \rangle)} \mathcal{M}_\gamma.$$

Note that

$$|\mathcal{M}| \leq |\Lambda| \leq g(2k+2)^5 \leq \Gamma.$$

Since p_i is chosen uniformly at random among primes in the range $\{1, \dots, M\}$, where $M = \Gamma^4 \cdot n^{14}$, from Lemma 3.1 we infer that the following event happens with probability at least $1 - \frac{1}{n^5}$ (conditioned on (\star) happening):

$$\left(\sum_{e \in S} 2^{\text{id}(e)}\right) \bmod p_i \neq \left(\sum_{e \in S'} 2^{\text{id}(e)}\right) \bmod p_i \quad \text{for all } S, S' \in \mathcal{M}, S \neq S'. \quad (9)$$

As argued in (7), we have $\sum_{e \in S} (2^{\text{id}(e)} \bmod p_i) < Mn^2$, hence event (9) happening entails that the mapping

$$S \mapsto (\omega_i(S) \bmod Mn^2) \bmod p_i$$

is injective on \mathcal{M} . In particular, the elements of \mathcal{M} receive pairwise different weights w.r.t. ω_i .

All in all, we conclude that the probability that (\star) happens and that the elements of \mathcal{M} receive pairwise different weights w.r.t. ω_i is at least

$$\left(1 - \frac{5 \cdot 8^{i-1}}{n^5}\right) \left(1 - \frac{1}{n^5}\right) \geq 1 - \frac{8^i}{n^5}.$$

It now suffices to observe that Claim 6.5 together with injectivity of ω_i on \mathcal{M} directly imply the conclusion of Claim 6.3 for segment I .

7 Lower Bounds

7.1 Unconditional Lower Bounds

In this subsection we present several lower bounds against the existence of (oblivious) isolation schemes for NP-hard problems on graphs with constant treedepth or pathwidth. We start with a standard information-theoretical lower bound on the number of random bits needed for construction of a weight function. A similar statement can be found in [11]

Lemma 7.1 (cf. proof of Theorem 1 in [11]). *For every $\beta > 0$ there exists $\alpha > 0$ such that the following holds. Suppose $n, W \in \mathbb{N}$ are large enough (depending on β) and $\omega_1, \dots, \omega_t: [n] \rightarrow [W]$ are weight assignments, where $t \leq \alpha \cdot \frac{n}{\log(nW)}$. Then there exist two different subsets $A, B \subseteq [n]$ such that*

- $A \cup B = [n]$,
- $|A \setminus B| = |B \setminus A| \leq \beta n$, and
- for every $i \in [t]$ it holds that $\omega_i(A) = \omega_i(B)$.

Proof. We may assume that $\beta \leq \frac{1}{2}$. Observe that then there exists $\alpha > 0$ such that

$$\binom{n}{\lfloor \beta n \rfloor} > 1 + 2^{\alpha n} \quad \text{for all } n \in \mathbb{N} \text{ large enough.} \quad (10)$$

We verify that the statement of the lemma holds for α chosen as above.

Consider a function $\Phi: 2^{[n]} \rightarrow \mathbb{N}^t$ defined as:

$$\Phi(X) := (\omega_1(X), \omega_2(X), \dots, \omega_t(X)).$$

Since for every nonempty $X \subseteq [n]$ and $i \in [t]$ we have $1 \leq \omega_i(X) \leq nW$, function Φ can take at most $1 + (nW)^t$ different values. As we assumed that $t \leq \alpha \cdot \frac{n}{\log(nW)}$, by (10) we have

$$1 + (nW)^t \leq 1 + 2^{\alpha n} < \binom{n}{\lfloor \beta n \rfloor}.$$

By pigeonhole principle it follows that there exist two different sets $A', B' \subseteq [n]$, each of size exactly $\lfloor \beta n \rfloor$, that have the same value assigned by Φ . This means that $\omega_i(A') = \omega_i(B')$ for all $i \in [t]$. Let now

$$A := A' \cup ([n] \setminus (A' \cup B')) \quad \text{and} \quad B := B' \cup ([n] \setminus (A' \cup B')).$$

Note that for each $i \in [t]$, we have

$$\omega_i(A) = \omega_i(A') + \omega_i([n] \setminus (A' \cup B')) = \omega_i(B') + \omega_i([n] \setminus (A' \cup B')) = \omega_i(B).$$

Further,

$$|A \setminus B| = |A' \setminus B'| \leq |A'| \leq \beta n,$$

and similarly $|B \setminus A| \leq \beta n$. Thus, sets A and B have the desired property. \square

We now use Lemma 7.1 to prove the following statement, which will be used for establishing a lower bound against isolation schemes for maximum-size independent sets.

Lemma 7.2. *There exists $\alpha > 0$ such that the following holds. Suppose $n, W \in \mathbb{N}$ are large enough and $\omega_1, \dots, \omega_t: [n] \rightarrow [W]$ are weight assignments, where $t \leq \alpha \cdot \frac{n}{\log(nW)}$. Then there exists a graph G on vertex set $[n]$ and such that*

- the treedepth of G is at most 4, and




Figure 6: The left panel presents the construction of the graph G . The vertices of the only two maximum-size independent sets of G are marked blue and red, respectively. The right panel presents an elimination forest of G of height 4.

- G has exactly two different maximum-size independent sets A and B , which moreover satisfy $\omega_i(A) = \omega_i(B)$ for all $i \in [t]$.

Proof. By Lemma 7.1 applied for $\beta = \frac{1}{2}$, we can choose α small enough so that the assertion $t \leq \alpha \cdot \frac{n}{\log(nW)}$ implies the existence of two different sets $A, B \subseteq [n]$ such that $|A \setminus B| = |B \setminus A|$, $A \cup B = [n]$, and $\omega_i(A) = \omega_i(B)$ for all $i \in [t]$. We are going to construct a graph G on vertex sets $[n]$ of treedepth at most 4 such that A and B are the only two maximum-size independent sets in G . Let $k := |A \setminus B|$.

Let us arbitrarily enumerate $A \setminus B$ as $\{a_1, \dots, a_k\}$ and $B \setminus A$ as $\{b_1, \dots, b_k\}$. The edge set of G consists of the following edges, see the left panel of Figure 6:

$$\{a_i b_i : i \in [k]\} \cup \{a_i b_1 : i \in [k] \setminus \{1\}\} \cup \{a_1 b_i : i \in [k] \setminus \{1\}\}.$$

This concludes the construction of G . Note that vertices of $A \cap B$ are isolated in G . It is easy to see that G admits an elimination forest of height 4, see the right panel of Figure 6.

It remains to prove that A and B are the only two maximum-size independent sets in G . Note first that A and B are indeed independent and there are no larger independent sets, because every independent set in G contains at most one endpoint from each edge of the matching $M := \{a_i b_i : i \in [k]\}$. Therefore, if I is a maximum-size independent set in G , then I needs to contain all the (isolated) vertices from $A \cap B$, and one endpoint of each edge of M . Now if I contains the endpoint a_1 of the edge $a_1 b_1$, then I cannot contain any of the vertices b_i for $i \in [k]$, because all these vertices are adjacent to a_1 . Hence I must contain all vertices a_i for $i \in [k]$, implying that $I = A$. Analogously, if I contains b_1 , then $I = B$. \square

Corollary 7.3 (Lower bound for maximum-size independent sets). *There does not exist an isolation scheme for maximum-size independent sets on graphs of treedepth at most 4 that would use $o(\log n)$ random bits and assign polynomially bounded weights.*

Proof. By Lemma 7.2, such an isolation scheme would need to produce $\Omega(n/\log n)$ different weight assignments on n -vertex graphs, for otherwise there would exist a graph of treedepth at most 4 where the only two maximum-size independent sets receive the same weights in all possible weight assignments. It follows that the isolation scheme in question needs to use $\Omega(\log n)$ random bits. \square

We now present similar constructions for three other types of objects: minimum Steiner trees, minimum maximal matchings, and Hamiltonian cycles. In each case we first give a lemma presenting the construction, which is followed by a corollary stating the lower bound. In each case, the corollary follows from the same argument as that used in the proof of Corollary 7.3.

For a graph G and a set of terminals $T \subseteq V(G)$, a *minimum Steiner tree* is a minimum-size set of vertices $S \subseteq V(G)$ such that $T \subseteq S$ and $G[S]$ is connected.




Figure 7: Construction used in the proofs of Lemmas 7.4, 7.6, and 7.8, respectively.

Lemma 7.4. *There exists $\alpha > 0$ such that the following holds. Suppose $n, W \in \mathbb{N}$ are large enough and $\omega_1, \dots, \omega_t: [n] \rightarrow [W]$ are weight assignments, where $t \leq \alpha \cdot \frac{n}{\log(nW)}$. Then there exists a graph G on vertex set $[n]$ and a set of terminals $T \subseteq [n]$ such that*

- the treedepth of G is at most 4, and
- there are exactly two different minimum Steiner trees for G and T , say A and B , and they satisfy $\omega_i(A) = \omega_i(B)$ for all $i \in [t]$.

Proof. As in the proof of Lemma 7.2, we can use Lemma 7.1 for $\beta = \frac{1}{3}$ to make sure that provided α is chosen small enough, there are two different sets $A, B \subseteq [n]$ such that $A \cup B = [n]$, $|A \setminus B| = |B \setminus A| \leq n/3$, and $\omega_i(A) = \omega_i(B)$ for all $i \in [t]$. It is easy to modify the proof of Lemma 7.1 so that the following property is also guaranteed: if $k := |A \setminus B|$, then $k \geq 2$. We are going to construct a graph G on vertex set $[n]$ together with a set of terminals $T \subseteq V(G)$ so that A and B are the only two minimum Steiner trees for G and T , and G has treedepth at most 4.

Let us arbitrarily enumerate $A \setminus B$ as $\{a_1, \dots, a_k\}$ and $B \setminus A$ as $\{b_1, \dots, b_k\}$. Note that $k \leq n/3$ and $|A \cap B| \geq n/3$, hence $k \leq |A \cap B|$. We set $T := A \cap B$ to be the terminals. Further, let $T' \subseteq T$ be any subset of $k - 1$ terminals and let us arbitrarily enumerate T' as $\{t_1, \dots, t_{k-1}\}$. First, we make every terminal in $T \setminus T'$ adjacent to the terminal t_1 . Next for every $i \in [k - 1]$, we add edges $t_i a_i, t_i b_i, a_i a_k$ and $b_i b_k$. This concludes the construction of G . See the left panel of Figure 7 for a visualization.

Observe that each connected component of the graph $G - \{a_k, b_k\}$ is a star, and hence has treedepth at most 2. It follows that G has treedepth at most 4. That A and B are the only two minimum Steiner trees for G and T is straightforward; we leave the verification to the reader. \square

Corollary 7.5 (Lower bound for minimum Steiner trees). *There does not exist an isolation scheme for minimum Steiner trees on graphs of treedepth at most 4 that would use $o(\log n)$ random bits and assign polynomially bounded weights.*

Next, recall that a matching in a graph is *maximal* if no strict superset of it is a matching, and it is moreover a *minimum maximal matching* if it has the smallest possible size among maximal matchings.

Lemma 7.6. *There exists $\alpha > 0$ such that the following holds. Suppose $m, W \in \mathbb{N}$ are large enough and $\omega_1, \dots, \omega_t: [m] \rightarrow [W]$ are weight assignments, where $t \leq \alpha \cdot \frac{m}{\log(mW)}$. Then there exists a graph G with edge set $[m]$ such that*

- the treedepth of G is at most 4, and
- there are exactly two different minimum maximal matchings in G , say A and B , and they satisfy $\omega_i(A) = \omega_i(B)$ for all $i \in [t]$.

Proof. By applying Lemma 7.1 for $\beta = \frac{1}{5}$, we can choose α small enough so that the assertion $t \leq \alpha \cdot \frac{m}{\log(mW)}$ implies that there exists a pair of different subsets $A, B \subseteq [m]$ such that $|A \setminus B| = |B \setminus A| \leq m/5$,

$A \cup B = [m]$, and $\omega_i(A) = \omega_i(B)$ for all $i \in [t]$. Note that then $|A \cap B| \geq 3m/5$, hence $|A \setminus B| \leq 3|A \cap B|$. Let then $\overline{K} \subseteq A \cap B$ be any subset of $A \cap B$ of size exactly $3k - 2$, where $k := |A \setminus B|$, and let $K := (A \cap B) \setminus \overline{K}$. As in Lemma 7.4, we may assume that $k \geq 2$.

Let $\tilde{A} := (A \setminus B) \cup K$ and $\tilde{B} := (B \setminus A) \cup K$. Note that $\tilde{A} \cap \tilde{B} = K$. Further, for every $i \in [t]$ we have

$$\omega_i(\tilde{A}) = \omega_i(A) - \omega_i(\overline{K}) = \omega_i(B) - \omega_i(\overline{K}) = \omega_i(\tilde{B}).$$

We now construct a graph G with edge set $[m]$ and treedepth at most 4 which has exactly two different minimum maximal matching: \tilde{A} and \tilde{B} . See the center panel of Figure 7 for the construction.

Let us arbitrarily enumerate $A \setminus B$ as $\{a_1, \dots, a_k\}$ and $B \setminus A$ as $\{b_1, \dots, b_k\}$. Further, recalling that $|\overline{K}| = 3k - 2$, we arbitrarily enumerate \overline{K} as $\{d_1, \dots, d_k, c_1, \dots, c_{k-1}, c'_1, \dots, c'_{k-1}\}$. Let us create $4k$ vertices: $\{v_1^c, \dots, v_k^c\}, \{v_1^a, \dots, v_k^a\}, \{v_1^b, \dots, v_k^b\}, \{v_1^d, \dots, v_k^d\}$. For every $i \in [k]$ we connect:

- vertices v_i^c and v_i^a using edge a_i ;
- vertices v_i^c and v_i^b using edge b_i ; and
- vertices v_i^c and v_i^d using edge d_i .

Next, for every $i \in [k - 1]$ we connect:

- vertices v_k^b and v_i^a using edge c_i ; and
- vertices v_k^a and v_i^b using edge c'_i .

In this way we have defined the endpoints of all the edges of $[m]$ apart from the edges of K . To finish the construction, for each edge $e \in K$ we add two extra vertices and connect them using e . Thus, K becomes a matching in G consisting of isolated edges.

This concludes the construction of G . Note that if from G we remove v_k^a and v_k^b , then each of the remaining connected components is a star, and hence has treedepth at most 2. This proves that G has treedepth at most 4.

We are left with proving that \tilde{A} and \tilde{B} are the only two minimum maximal matchings in G . Clearly \tilde{A} and \tilde{B} are maximal matchings. Let M be any maximal matching of G . Since for every $i \in [k]$, vertex v_i^d has degree 1 and v_i^c is its only neighbor, it follows that each vertex v_i^c needs to be incident to an edge of M . Clearly, M also needs to contain each edge of K , as these edges are isolated in G . Since vertices v_i^c are pairwise nonadjacent, we conclude that every maximal matching of G has at least $k + |K|$ edges. As $|\tilde{A}| = |\tilde{B}| = k + |K|$, this implies that both \tilde{A} and \tilde{B} are minimum maximal matchings in G .

It remains to argue that there are no minimum maximal matchings other than \tilde{A} and \tilde{B} . Let M be any minimum maximal matching in G . As we argued, M needs to contain K , one edge incident to v_i^c for each $i \in [k]$, and no other edges. In particular, no edge c_i or c'_i for any $i \in [k - 1]$, is contained in M . By the maximality of M this means that for each $i \in [k - 1]$, at least one of the edges $\{a_k, b_i\}$ is included in M , and at least one of the edges $\{b_k, a_i\}$ is included in M . If neither a_k nor b_k was included in M , then this would mean that both a_1 and b_1 necessarily belong to M , a contradiction. If now $a_k \in M$ and $b_k \notin M$, then $a_i \in M$ for all $i \in [k - 1]$ and $M = \tilde{A}$. Similarly, if $a_k \notin M$ and $b_k \in M$, then $M = \tilde{B}$. \square

Corollary 7.7 (Lower bound for minimum maximal matchings). *There does not exist an isolation scheme for minimum maximal matchings on graphs of treedepth at most 4 that would use $o(\log n)$ random bits and assign polynomially bounded weights.*

Lemma 7.8. *There exists $\alpha > 0$ such that the following holds. Suppose $m, W \in \mathbb{N}$ are large enough and $\omega_1, \dots, \omega_t: [m] \rightarrow [W]$ are weight assignments, where $t \leq \alpha \cdot \frac{m}{\log(mW)}$. Then there exists a graph G with edge set $[m]$ such that*

- the pathwidth of G is at most 4, and
- there are exactly two different Hamiltonian cycles in G , say with edge sets A and B , and they satisfy $\omega_i(A) = \omega_i(B)$ for all $i \in [t]$.

Proof. We apply Lemma 7.1 again, this time for $\beta = \frac{1}{3}$. Hence, by selecting α small enough we can ensure that there exists different sets $A, B \subseteq [m]$ such that $|A \setminus B| = |B \setminus A| \leq m/3$, $A \cup B = [m]$, and $\omega_i(A) = \omega_i(B)$ for all $i \in [t]$. Note that then $|A \cap B| \geq m/3$, hence $|A \cap B| \geq k$, where $k := |A \setminus B|$. As before, we may assume that $k \geq 3$.

Let us arbitrarily enumerate $A \setminus B$ and $B \setminus A$ as $\{a_1, \dots, a_k\}$ and $\{b_1, \dots, b_k\}$, respectively. Further, let $\{c_1, \dots, c_k\}$ be k arbitrary elements of $A \cap B$.

We now construct a graph G with edge set $[m]$ as follows; see the right panel of Figure 7. First, create $2k$ vertices $\{u_1, \dots, u_k\}$ and $\{v_1, \dots, v_k\}$. For every $i \in [k]$ we connect vertices u_i and u_{i+1} using edge a_i , where $u_{k+1} = u_1$. Similarly, for every $i \in [k]$ we connect vertices v_i with v_{i+1} with edge b_i , where $v_{k+1} = v_1$. Next, for every $i \in [k]$ we connect vertices u_i to v_i with the edge c_i . So far we have defined the endpoints of all the edges apart from the edges of $(A \cap B) \setminus \{c_1, \dots, c_k\}$. To accommodate the remaining edges, replace the edge c_1 with a path of length $|A \cap B| - k + 1$ connecting u_1 and v_1 , and let the edges of this path be $(A \cap B) \setminus \{c_2, \dots, c_k\}$.

Note that removing u_1 and v_1 turns G into the union of a path and a $2 \times (k - 1)$ grid, which is a graph of pathwidth at most 2. It follows that G itself has pathwidth at most 4. Finally, it is easy to see that G has exactly two Hamiltonian cycles, with edge sets A and B , respectively. \square

Corollary 7.9 (Lower bound for Hamiltonian cycles). *There does not exist an isolation scheme for Hamiltonian cycles on graphs of pathwidth at most 4 that would use $o(\log n)$ random bits and assign polynomially bounded weights.*

7.2 Conditional Lower Bounds

An important result in the complexity theory is the randomized reduction from languages in NP to UNIQUE SAT due to Valiant and Vazirani [54]. This reduction is the essential procedure in plenty fundamental results in the Computational Complexity, most notably Toda's Theorem [52].

We would like to note, that the question of derandomization of this red was already subject to a rigorous research. For example, [19] showed that probably we cannot hope to improve the probability of success of Valiant and Vazirani [54] reduction to be $2/3$ unless $\text{NP} \subseteq \text{P/poly}$. In the opposite success probability regime, Calabro et al. [10] gave a randomized polynomial time reduction from k -SAT to UNIQUE k -SAT which works with probability $2^{-\mathcal{O}(n \log^2 k/k)}$. Their bound was subsequently improved by [53] to $2^{-\mathcal{O}(n \log k/k)}$. Very recently [56] showed that if UNIQUE k -SAT on n variable admits an $2^{n(1-f(k)/k)}$ time algorithm for some unbounded f , then k -SAT is in $2^{n(1-f(k)(1-\varepsilon)/k)}$ time for every $\varepsilon > 0$. These reductions work in exponential time.

The question of whether you can derandomize the reduction of Valiant and Vazirani [54] was already subject to a rigorous research. Dell et al. [19] showed that improving the success probability of Valiant and Vazirani [54] reduction to $2/3$ is not possible unless $\text{NP} \subseteq \text{P/poly}$. In the opposite success probability regime, Calabro et al. [10] gave a randomized polynomial time reduction from k -SAT to UNIQUE k -SAT which works with probability $2^{-\mathcal{O}(n \log^2 k/k)}$. Their bound was subsequently improved by [53] to $2^{-\mathcal{O}(n \log k/k)}$. Very recently [56] showed that if UNIQUE k -SAT on n variable admits an $2^{n(1-f(k)/k)}$ time algorithm for some unbounded f , then k -SAT is in $2^{n(1-f(k)(1-\varepsilon)/k)}$ time for every $\varepsilon > 0$. These reductions work in exponential time.

Montoya and Müller [41] considered a parameterized version of result of Valiant and Vazirani [54] and showed that certain parameterized problems are also as hard as their unique variants.

The lower bounds in this section build on the following working assumption that these types of randomized reductions cannot be derandomized in the following strong sense:

Conjecture 7.10 (Linear-Random-Bits Conjecture). *There is no randomized polynomial time reduction from SAT to UNIQUE SAT that uses $o(n)$ random bits, where n is the number of variables of the original instance.*

Falsifying Conjecture 7.10 would on its own contribute to great progress in computational complexity and hopefully inspire novel ideas in derandomization.

The Conjecture 7.10 enables us to express the natural barrier in currently known techniques. We hope to spark an interest in more randomness efficient reduction to UNIQUE SAT, as it would hopefully lead to better Isolation Schemes. All our isolation schemes can be even implemented in the restricted NC setting. We could even focus on weaker and more believable Conjecture 7.10 to exclude NC-reductions from SAT to UNIQUE SAT with $o(n)$ random bits. This weaker conjecture would enable us to exclude isolation schemes that can be computed with NC circuits only.

Moreover, note that the reductions proposed in previous work blow up the size of an instance to be $n^{\mathcal{O}(1)}$. We do not restrict a size of the output UNIQUE SAT instance, outside the fact that reduction needs to be in polynomial time.

In this section we prove the following theorem.

Theorem 7.11 (Conditional Lower Bound for MAXIMUM INDEPENDENT SET). *No isolation scheme that can be computed in polynomial time and uses $o(td)$ random bits with polynomially bounded maximum weight exists for MAXIMUM INDEPENDENT SET unless Conjecture 7.10 is false.*

To prove Theorem 7.11 we use the fact that MAXIMUM INDEPENDENT SET is reducible to SAT with a *parsimonious* reduction (cf., [30, Exercise 2.30 and 2.28]). Parsimonious reduction is a reduction that preserves number of solutions [30]. In fact, all known natural reductions between NP-complete problems are parsimonious or can be easily modified to be parsimonious [30, Section 6.2.1].




Figure 8: The left scheme presents the example construction of the graph G in the reduction from SAT to MAXIMUM INDEPENDENT SET. Dotted vertices are part of the variable clauses. Note that after removing them, the connected components are $\mathcal{O}(n)$ therefore treedepth of G is $\mathcal{O}(n)$. The right scheme presents the schematic view of proof of Theorem 7.11. We use the fact that reductions behind black arrows are known. In the proof we present red reductions.

Proof. First, we reduce SAT on n variables to MAXIMUM INDEPENDENT SET problem on $N = n^{\mathcal{O}(1)}$ vertices by a standard reduction [37]. The reduction gives us graph G and number $k \in \mathbb{N}$, such that there is no independent set of size greater than k on G and G has maximum independent set of size exactly k if the original formula was satisfiable. Moreover graph G has treedepth $\mathcal{O}(n)$ because after removing vertices responsible for variable vertices graph G has connected components of size $\mathcal{O}(n)$ and there are only $\mathcal{O}(n)$ variable vertices. See Figure 8 for example construction of graph G in reduction. Note that the reduction [37] is not parsimonious and we do not need it to be at this point.

Now, we assume, that there exists an isolation scheme for MAXIMUM INDEPENDENT SET on N -vertex graphs with d bounded treedepth that use $o(d)$ random bits and maximum weight is $M = n^{\mathcal{O}(1)}$.

Now, we invoke our isolation scheme on N -vertices graph G and select a random integer $W \in_R [NM]$. We get a graph G with weights $\omega_1, \dots, \omega_N \in [n^{\mathcal{O}(1)}]$ of vertices and the property that with $1/n^{\mathcal{O}(1)}$ probability there exists exactly one independent set in G of size exactly k and weight exactly W (if the original formula was satisfiable).

We modify our graph as follows. For every vertex $v_i \in V(G)$ the graph \tilde{G} have $\ell_i := 2NM + w_i$ vertices $\{v_1^i, \dots, v_{\ell_i}^i\}$. For each edge $(v_i, v_j) \in E(G)$ we add an edge $(v_a^i, v_b^j) \in E(\tilde{G})$ for every $a \in [\ell_i]$ and $b \in [\ell_j]$. Set $\tilde{k} = 2kNM + W$. The transformed graph \tilde{G} has the property:

- (i) If every independent set of G is $< k$, then every independent set of \tilde{G} is $< \tilde{k}$, and
- (ii) If G has independent set of size k , then with $1/n^{\mathcal{O}(1)}$ probability \tilde{G} has a *unique* independent set of size \tilde{k} .

Graph \tilde{G} is an instance of unweighted MAXIMUM INDEPENDENT SET with the property that it has a unique independent set of size \tilde{k} if the original formula was satisfiable. Finally, we reduce our maximum independent set instance \tilde{G} to the SAT with parsimonious reduction [30]. This guarantees that with inversely polynomial probability the final instance of SAT has unique solution (if the original formula was a yes-instance).

Because graph G has treedepth $\mathcal{O}(n)$, our isolation scheme uses $o(n)$ random bits and therefore Conjecture 7.10 is false. See Figure 8 for a overall scheme of the reduction to UNIQUE SAT and sizes and parameters of the produced instances. Note, that the final instance of UNIQUE SAT may be polynomially larger than the original instance and it does not contradict Conjecture 7.10. \square

Observe that this lower-bounds framework works for many other NP-complete problems. Basically, all we need is (1) the reduction from SAT creates a graph with $\mathcal{O}(n)$ treedepth/treewidth and (2) there is a polynomial time reduction from weighted to unweighted problem. We can generalize our lower bound to work for HAMILTONIAN CYCLE in bounded treewidth graphs. First we reduce SAT to HAMILTONIAN CYCLE to the graph with $t = \mathcal{O}(n)$ and apply $o(t)$ -random bits isolation scheme. We observe that we can subdivide every edge of weight ω with a path of length ω . This way we construct an instance of SUBSET TSP on unweighted graphs (where terminals are vertices of the original graphs) and use a parsimonious reduction from SUBSET TSP to SAT. We sum up this observation with the following remark.

Remark 7.12. *Assuming Conjecture 7.10 no isolation scheme that can be computed in polynomial time and uses $o(t)$ random bits with polynomially bounded maximum weight exists for HAMILTONIAN CYCLE in graphs of treewidth bounded by t .*

For HAMILTONIAN CYCLE in planar graphs, we can use the fact that NP-hardness reduction [29] from SAT on m -variables produces a graph G with $\mathcal{O}(m^2)$ vertices. Similarly, we subdivide every edge of weight ω with a ω -length path. We arrive at the instance of SUBSET TSP in unweighted graphs and use a parsimonious reduction from SUBSET TSP to SAT. Therefore any isolation scheme that needs $o(\sqrt{n})$ random bits and uses polynomial weights would analogously contradict Conjecture 7.10.

Remark 7.13. *Assuming Conjecture 7.10 no isolation scheme that can be computed in polynomial time and uses $o(\sqrt{n})$ random bits with polynomially bounded maximum weight exists for HAMILTONIAN CYCLE in planar graphs (where n is the number of vertices of the graph).*

8 Isolation of local vertex selection problems

Recall that an independent set in a graph is a set of pairwise nonadjacent vertices, and an independent set is *maximum* if it has the largest possible cardinality. In this section we prove Theorem 1.10, which in plain words can be restated as follows. For a graph G , let $\text{MIS}(G)$ denote the set of all maximum independent sets in G . Suppose we consider a graph $G \in \mathcal{G}_d$, say on vertex set $[n]$, given together with an elimination forest F of height at most d . The isolation scheme of Theorem 1.10 is a family of $\ell = 2^{\mathcal{O}(d)}$ weight

functions $\omega_1, \dots, \omega_\ell: [n] \times [d] \rightarrow [W]$, where $W = \mathcal{O}(n^6)$. Function ω_i assigns to each $v \in [n]$ the weight $\omega_i(v, \text{lvl}_F(v))$. The requirement is that for all G and F as above, at least half of the functions $\omega_1, \dots, \omega_\ell$ isolates $\text{MIS}(G)$.

In the following, when discussing level-aware isolation schemes, we consider all weight functions as acting on a single argument: the vertex in question. The level of this vertex is supplied implicitly, as there is always some elimination forest of the graph fixed in the context.

In Section 9 we show how to extend the reasoning from this section to an example edge selection problem – selection of maximum matchings.

8.1 Exchange property

In fact, we will prove a more general result than Theorem 1.10. Namely, we isolate an abstract property of a vertex selection problem, which we call the *exchange property*, which is enjoyed by $\text{MIS}(\cdot)$ and which is sufficient for our technique to work. There are multiple other problems of local nature that also have this property, hence the approach is indeed more general.

We will rely on the following two definitions.

Definition 8.1 (Pivotal vertex). *Let G be a graph and F be an elimination forest of G . For a family $\mathcal{F} \subseteq 2^{V(G)}$ and two different sets of vertices $A, B \in \mathcal{F}$, we say that a vertex u is pivotal for A and B in F if $u \in A \Delta B$ and $\text{tail}_F(u) \cap A = \text{tail}_F(u) \cap B$.*

Definition 8.2 (Exchange property). *We say that a vertex selection problem \mathcal{P} has the exchange property if for every graph G , elimination forest F of G , and weight function $\omega: V(G) \rightarrow \mathbb{N}$ the following holds: if there exist two different $A, B \in \mathcal{P}(G)$ that are minimizers of ω on $\mathcal{P}(G)$, then there also exist $A', B' \in \mathcal{P}(G)$ that are minimizers of ω on $\mathcal{P}(G)$ such that there is only one pivotal vertex for A' and B' in F .*




Figure 9: Schematic figure of the exchange argument used in the proof of Lemma 8.4. Both figures present the elimination forest F of the given graph G . Black-filled vertices on the left are in A . Black-filled vertices on the right are in B . The pivotal vertices are green-stroked. Vertex $u \in A \setminus B$ is a chosen pivotal vertex. Because A and B are both minimizers of ω , we can exchange the red subtree with the yellow one and construct a valid minimizer with one less pivotal vertex.

In the next sections we will focus on proving the following result.

Theorem 8.3. *Let \mathcal{P} be a vertex selection problem that enjoys the exchange property. Then for every $d \in \mathbb{N}$, there is a level-aware isolation scheme for \mathcal{P} on graphs of treedepth at most d that uses $\mathcal{O}(d)$ random bits assigns weights bounded by $\mathcal{O}(n^6)$.*

Therefore, Theorem 1.10 follows by combining Theorem 8.3 with the following result.

Lemma 8.4. *The vertex selection problem $\text{MIS}(\cdot)$ has the exchange property.*

Proof. Let G be a graph, F be an elimination forest of G , and $\omega: V(G) \rightarrow \mathbb{N}$ be a weight function. Assume that there exist two different $A, B \in \text{MIS}(G)$ that are minimizers for ω and have more than one pivotal vertex. We will construct two different sets A' and B' that are also maximum independent sets in G and have the same weight as A and B , but have one less pivotal vertex. This is sufficient to prove the lemma, since we can repeat the construction to get a pair of minimizers with exactly one pivotal vertex, as required.

Let $u \in V(G)$ be a pivotal vertex for A and B . By symmetry, assume that $u \in B \setminus A$. We consider

$$A' := A \quad \text{and} \quad B' := (B \setminus \text{subtree}[u]) \cup (\text{subtree}[u] \cap A).$$

It is easy to see that A' and B' have exactly one less pivotal vertex than A and B : u is pivotal for A and B but not for A' and B' , and all other pivotal vertices are the same. Because $A' = A$, we only need to prove that $B' \in \text{MIS}(G)$ and that B' is a minimizer of ω on $\text{MIS}(G)$.

First, note that $\text{tail}(u) \cap A = \text{tail}(u) \cap B$ because u is pivotal of A and B . Since the vertices of $\text{subtree}[u]$ have only neighbors in $\text{tail}(u)$, this implies that B' is an independent set, and analogously the set $A'' := (A \setminus \text{subtree}[u]) \cup (\text{subtree}[u] \cap B)$ is an independent set as well. Since $|B'| + |A''| = |A| + |B|$, both B' and A'' are independent sets, while A and B are maximum independent sets, it follows that both B' and A'' must be maximum independent sets as well. Similarly, we have $\omega(B') + \omega(A'') = \omega(A) + \omega(B)$, so the assumption that A and B are both minimizers of ω on $\text{MIS}(G)$ implies that A'' and B' are also minimizers of ω on $\text{MIS}(G)$. \square

The same argument as the one used in the proof of Lemma 8.4 can be also applied to other combinatorial objects in graphs, where validity of an object depends on checking the neighborhood of every vertex. For instance, it is easy to prove in this way that minimum vertex covers and minimum dominating sets also have the exchange property. Thus, Theorem 8.3 also applies to the corresponding vertex selection problems.

In the next sections we will work towards the proof of Theorem 8.3. Therefore, let us fix a vertex selection problem \mathcal{P} that has the exchange property.

8.2 Warm-up: a deterministic isolation scheme

Before commencing to the proof of Theorem 1.10, we will show a simple deterministic level-aware isolation scheme. Observe that the Isolation Lemma in its most general form can be trivially derandomized provided we allow the maximum weight to be 2^n , where n is the size of the universe. Namely, it is enough to select the weight function $\omega(i) := 2^i$. In general, allowing exponential weights is prohibitively expensive and not algorithmically useful, however considering this idea explains some intuition behind our techniques. As a warm-up, we now present a level-aware isolation scheme for the considered problem \mathcal{P} on graphs of treedepth at most d that is deterministic, but may use weights as large as 2^d . The scheme is captured by the following lemma.

Lemma 8.5 (Exponential-weight deterministic isolation). *For every graph G and an elimination forest F of G , function $\omega_{\text{det}}(v) := 2^{\text{vl}_F(v)}$ isolates the family $\mathcal{P}(G)$.*

Note that as Lemma 8.5 involves only one weight function ω_{det} , it provides a deterministic isolation scheme. Also, provided the height of F is at most d , the assigned weights are upper bounded by 2^d . Further, this is a level-aware isolation scheme, because ω_{det} is also supplied with the level of the vertex in the given elimination forest. In Section 7 we showed that this additional information is really necessary, since without it any isolation scheme for maximum independent sets needs to use $\Omega(\log(n))$ random bits, even on graphs of treedepth at most 4.

Proof of Lemma 8.5. Assume for contradiction that there exist two different sets in $\mathcal{P}(G)$ that are both minimizers of $\omega_{\det}(v) := 2^{\text{lvl}_F(v)}$ on $\mathcal{P}(G)$. By Lemma 8.4 we know that there also exist $A, B \in \mathcal{P}(G)$ that are also minimizers of ω_{\det} and have exactly one pivotal vertex.

Let u be the only pivotal vertex of A and B in F . Without loss of generality we may assume that $u \in A \setminus B$. Let

$$R := A \setminus \text{subtree}[u] = B \setminus \text{subtree}[u], \quad S_A := A \cap \text{subtree}(u), \quad \text{and} \quad S_B := B \cap \text{subtree}(u).$$

We know that $\omega_{\det}(A) = \omega_{\det}(B)$. Therefore,

$$\omega_{\det}(R) + \omega_{\det}(u) + \omega_{\det}(S_A) = \omega_{\det}(R) + \omega_{\det}(S_B),$$

hence

$$\omega_{\det}(u) + \omega_{\det}(S_A) = \omega_{\det}(S_B). \tag{11}$$

Let ℓ be the level of vertex u . We know that $\omega_{\det}(u) = 2^\ell$. Moreover, the level of every vertex in S_A and S_B is greater than ℓ . Therefore $\omega_{\det}(S_A)$ and $\omega_{\det}(S_B)$ are divisible by $2^{\ell+1}$. We conclude that the right hand side of (11) is divisible by $2^{\ell+1}$, while the left hand side is not. This is a contradiction, hence ω_{\det} must have a unique minimizer on $\mathcal{P}(G)$. \square

8.3 Warm-up continued: a randomized isolation scheme

There is also a relatively simple level-aware isolation scheme for \mathcal{P} on graphs of treedepth at most d that uses $\mathcal{O}(d \log(n))$ random bits and assigns weights bounded by $\mathcal{O}(n)$. Namely, for every level $i \in [d]$ chose uniformly at random a number $r_i \in [Cn]$ for some constant C large enough, and let $\omega_{\text{rnd}}(v) := r_{\text{lvl}_F(v)}$, where F is the given elimination forest. Clearly, this scheme uses $\mathcal{O}(d \log n)$ random bits. It is not hard to prove that the function ω_{rnd} isolates $\mathcal{P}(G)$ with high probability. We leave the details to the reader, as the isolation scheme presented in the next section will supersede this one.

The idea behind our proof of Theorem 8.3 is to combine the deterministic isolation scheme ω_{\det} , presented in Lemma 8.5, the with randomized scheme sketched above. This approach is inspired by the shifting idea presented in [55, 11], but the technique is adapted to the setting of problems on graphs of bounded treedepth.

8.4 Proof of Theorem 8.3

Fix the considered graph $G = (V, E)$ and elimination forest F of height at most d , where $V = [n]$. We may assume that $d \geq 5 \lceil \log n \rceil$, for otherwise the deterministic isolation scheme from Lemma 8.5 isolates $\mathcal{P}(G)$ and assigns weights upper bounded by $2^d \leq \mathcal{O}(n^6)$.

Isolation Procedure. We start with the definition of the isolation scheme. Similarly to the scheme considered in Lemma 8.5, the weight assigned to a vertex will only depend on its level in F (and the random bits).

For every $i \in [d]$, we can uniquely encode i as a pair of integers $(e(i), f(i))$ so that

$$i = \lceil \log n \rceil \cdot e(i) + f(i), \quad \text{where } 0 \leq f < \lceil \log n \rceil.$$

Note that then $e(i) \leq \kappa$, where $\kappa := \left\lfloor \frac{d}{\lceil \log n \rceil} \right\rfloor$. As $d \geq 5 \lceil \log n \rceil$, we have $\kappa > 0$.

Next, for every $e \in \{0, 1, \dots, \kappa\}$ choose an integer $r_e \in [32n^5]$ independently and uniformly at random. Then $\bar{r} := (r_0, r_1, \dots, r_\kappa)$ is the vector of random integers used by our weight function. Note that choosing

\bar{r} requires $(\kappa + 1) \cdot \mathcal{O}(\log n) = \mathcal{O}(d)$ random bits, as promised. In the following, we treat r_0, \dots, r_κ as random variables, thus \bar{r} is a random vector that is uniformly distributed in $\Omega := [32n^5]^{\kappa+1}$.

Finally, for a vertex v and $\bar{\rho} \in \Omega$, we define

$$\omega_{\bar{\rho}}(v) := \rho_{e(i)} \cdot 2^{f(i)}, \text{ where } i = \text{lvl}(v).$$

With this definition in place, our isolation scheme simply samples \bar{r} as above and outputs the weight function $\omega_{\bar{r}}$. We are left with arguing that $\omega_{\bar{r}}$ isolates $\mathcal{P}(G)$ with probability at least $\frac{1}{2}$.

Analysis. For a vertex v , we write $e(v) := e(i)$ and $f(v) := f(i)$, where i is the level of v in F . For every set $X \subseteq V$, let us define a linear form $\phi^X: \Omega \rightarrow \mathbb{N}$ as follows:

$$\phi^X(\bar{\rho}) := \sum_{v \in X} \rho_{e(v)} \cdot 2^{f(v)}.$$

Intuitively, we can think of ϕ^X as a compressed version of X from which we can compute $\omega_{\bar{\rho}}(X)$ once it is fixed. Indeed, ϕ^X can be thought of as a compressed version of X since $|\Omega| \leq 2^{\mathcal{O}(d)}$. In particular, note that we actually have

$$\phi^X(\bar{\rho}) = \omega_{\bar{\rho}}(X). \quad (12)$$

Further, let

$$\Phi := \{\phi^X : X \in \mathcal{P}(G)\}.$$

The following lemma is the key step in the proof.

Lemma 8.6. *Suppose $\bar{\rho} \in \Omega$ is such that $\omega_{\bar{\rho}}$ does not isolate $\mathcal{P}(G)$. Then there are two different linear forms $\alpha, \beta \in \Phi$ such that*

$$\alpha(\bar{\rho}) = \beta(\bar{\rho}) = \min_{\gamma \in \Phi} \gamma(\bar{\rho}).$$

Proof. Since $\omega_{\bar{\rho}}$ does not isolate $\mathcal{P}(G)$, there are two different minimizers of $\omega_{\bar{\rho}}$ on $\mathcal{P}(G)$. By the exchange property, there are also two different minimizers A and B such that there exists exactly one pivotal vertex for A and B , say u . Without loss of generality suppose that $u \in A \setminus B$. Since A and B are both minimizers, by (12) we have

$$\phi^A(\bar{\rho}) = \phi^B(\bar{\rho}) = \min_{\gamma \in \Phi} \gamma(\bar{\rho}).$$

Hence, it suffices to prove that $\phi^A \neq \phi^B$.

Let $(e, f) := (e(u), f(u))$. We claim that the coefficients of ϕ^A and ϕ^B standing by the variable ρ_e are different. Letting $L := \{v \mid e(v) = e\}$, we see that these coefficients are respectively equal to

$$\sum_{v \in A \cap L} 2^{f(v)} \quad \text{and} \quad \sum_{v \in B \cap L} 2^{f(v)}.$$

Suppose for contradiction that these coefficients are actually equal, that is,

$$\sum_{v \in A \cap L} 2^{f(v)} = \sum_{v \in B \cap L} 2^{f(v)}. \quad (13)$$

Recall that u is the only pivotal vertex for A and B , hence $A \setminus \text{subtree}[u] = B \setminus \text{subtree}[u]$. Therefore, from (13) we infer that

$$\sum_{v \in A \cap L \cap \text{subtree}[u]} 2^{f(v)} = \sum_{v \in B \cap L \cap \text{subtree}[u]} 2^{f(v)}. \quad (14)$$

Now observe that $u \in A \cap L \cap \text{subtree}[u]$, $u \notin B \cap L \cap \text{subtree}[u]$, and $f(v) > f(u)$ for each $v \in L \cap \text{subtree}(u)$. From this it follows that the right hand side of (14) is divisible by $2^{f(u)+1}$, while the left hand side is not. This is a contradiction. \square

We now combine Lemma 8.6 with the following result of Chari et al. [11].

Lemma 8.7 (Proposition 2 of [11]). *Let Φ be a set of linear forms over $t < N$ variables with coefficients belonging to $\{0, \dots, N^2 - 1\}$. Let \bar{r} be chosen from $[N^5]^t$ uniformly at random. Then the probability that there exist two different $\alpha, \beta \in \Phi$ such that $\alpha(\bar{r}) = \beta(\bar{r}) = \min_{\gamma \in \Phi} \gamma(\bar{r})$ is at most $\frac{1}{2}$.*

Now observe that the coefficients in the forms appearing in Φ are bounded by $n \cdot 2^{\lceil \log n \rceil} < (2n)^2$, because each coefficient is a sum of at most n summands of the form $2^{f(v)}$, and each of these is bounded by $2^{\lceil \log n \rceil}$. Since \bar{r} is chosen uniformly at random from $[(2n)^5]^{\kappa+1}$, by combining Lemmas 8.6 and 8.7, where the latter is applied for $N := 2n$, we conclude that $\omega_{\bar{r}}$ isolates $\mathcal{P}(G)$ with probability at least $\frac{1}{2}$.

9 Isolation of local edge selection problems

In the Section 8, we designed level-aware isolation schemes exclusively for vertex selection problems. In this section we demonstrate that our techniques can be also applied to edge-selection problems on the example of maximum matchings: formally, we consider the edge selection problem $\text{MM}(\cdot)$ that maps every graph G to the family $\text{MM}(G)$ consisting of all maximum-size matchings in G .

Suppose we are given a graph $G = (V, E)$ and an elimination forest F of G . The *level* of an edge $e = uv \in E$ in F is defined as

$$\text{lvl}(e) := \min\{\text{lvl}(u), \text{lvl}(v)\}.$$

In our level-aware isolation scheme, the weight function will be supplied with two parameters: an edge and its level.

We now introduce the analogues of pivotal vertices and the exchange property.

Definition 9.1 (Edge-pivotal vertex). *For a family $\mathcal{F} \subseteq 2^E$ and two different sets of vertices $A, B \in \mathcal{F}$, we say that a vertex u is edge-pivotal for A and B if:*

- for every vertex $x \in \text{tail}(u)$ it holds that $ux \in A$ if and only if $ux \in B$;
- there exists $x \in \text{subtree}(u)$, such that $ux \in A \Delta B$;
- no vertex $u' \in \text{tail}(u)$ has the two properties above.

Definition 9.2 (Exchange property). *We say that an edge selection problem \mathcal{P} has the exchange property if for every graph $G = (V, E)$, elimination forest F of G , and weight function $\omega: E \rightarrow \mathbb{N}$, if there exist two different minimizers of ω on $\mathcal{P}(G)$, then there also exist two different minimizers of ω on $\mathcal{P}(G)$ for which there is exactly one edge-pivotal vertex.*

Using a reasoning similar to that from the proof of Lemma 8.4, we get the following.

Lemma 9.3. *The edge selection problem $\text{MM}(\cdot)$ has the exchange property.*

Proof. Let $G = (V, E)$ be a graph, F be an elimination forest of G , and $\omega: E \rightarrow \mathbb{N}$ be a weight function such that there are two different minimizers A, B of ω on $\text{MM}(G)$. Let e be any edge of $A \Delta B$ with the minimum possible level. Say that $e = uv$, where u is an ancestor of v . Let us define

$$B' := \left(B \cap \binom{\text{subtree}[u]}{2} \right) \cup \left(A \setminus \binom{\text{subtree}[u]}{2} \right).$$

As in the proof of Lemma 8.4, using the assumptions that A, B are maximum matching that are minimizers of ω on $\text{MM}(G)$, and that e is an edge of $A \Delta B$ of minimum possible level, we can easily see that B' is also a maximum matching that is a minimizer of ω on $\text{MM}(G)$. It now follows that u is the only pivotal vertex for A and B' . \square

The remainder of this section is devoted to the

Theorem 9.4. *For every $d \in \mathbb{N}$, there is a level-aware isolation scheme for maximum matchings on graphs of treedepth at most d that uses $\mathcal{O}(d \log n)$ random bits assigns weights bounded by $n^{\mathcal{O}(1)}$.*

Note that contrary to the situation in Section 8, we will conduct our reasoning only for the edge selection problem $\text{MM}(\cdot)$, as we will use some additional combinatorial properties of maximum matchings. While strong Isolation Lemma's already exist for the case of maximum matchings [51], our approach uses less random bits and seems extendable to other edge selection problems.

For the rest of this section, let us fix the graph $G = (V, E)$, and enumeration of its edges $\text{id}: E \rightarrow [m]$, and an elimination forest F of G . Mirroring the structure from Section 8, we first give a deterministic isolation scheme, which will be subsequently randomized in order to reduce the weights at the cost of random bits.

Recall that our weight functions take two parameters: an edge and its level. Similarly as in Section 8, we think of weight functions as acting only on edges, while the level of an edge is inferred implicitly from the forest F .

9.1 Deterministic isolation scheme

Let us introduce the weight function

$$\omega_{\text{det}}(e) := \text{id}(e) \cdot n^{2\text{lvl}(e)}.$$

We observe the following.

Lemma 9.5 (Exponential weight isolation). *Function ω_{det} isolates the family $\text{MM}(G)$.*

Proof. Assume for the contrary, that there exists two different maximum matchings $A, B \in \text{MM}(G)$ that are both minimizers of ω_{det} . Because $\text{MM}(\cdot)$ has the edge-exchange property, we can assume that A and B have exactly one edge-pivotal vertex. Let u be such a vertex and let ℓ be its level in the elimination forest F .

Let

$$R := A \setminus \binom{\text{subtree}[u]}{2}$$

be the set of edges from A that have at least one endpoint outside of $\text{subtree}[u]$. Note that since u is the only edge-pivotal vertex for A and B , it holds that

$$R = B \setminus \binom{\text{subtree}[u]}{2}.$$

For a vertex $v \in V$, let $E[v]$ be the set of edges with at least one endpoint in v . Finally, let

$$S_A := A \setminus (R \cup E[u]) = A \cap \binom{\text{subtree}(u)}{2} \quad \text{and} \quad S_B := B \setminus (R \cup E[u]) = B \cap \binom{\text{subtree}(u)}{2}.$$

We assumed that $\omega_{\text{det}}(A) = \omega_{\text{det}}(B)$. Therefore,

$$\omega_{\text{det}}(R) + \omega_{\text{det}}(S_A) + \omega_{\text{det}}(E[u] \cap A) = \omega_{\text{det}}(R) + \omega_{\text{det}}(S_B) + \omega_{\text{det}}(E[u] \cap B),$$

implying that

$$\omega_{\text{det}}(S_A) + \omega_{\text{det}}(E[u] \cap A) = \omega_{\text{det}}(S_B) + \omega_{\text{det}}(E[u] \cap B) \tag{15}$$

Every vertex in the maximum matching has degree 1. Since u is pivotal for A and B , u is incident in A to an edge e_A and incident in B to a different edge e_B such that $\text{lvl}(e_A) = \text{lvl}(e_B) = \ell$. Therefore,

$$\omega_{\det}(E[u] \cap A) = \text{id}(e_A)n^{2\ell} \quad \text{and} \quad \omega_{\det}(E[u] \cap B) = \text{id}(e_B)n^{2\ell}.$$

Note that $\text{lvl}(e) > \ell$ for each edge $e \in S_A \cup S_B$, hence both $\omega_{\det}(S_A)$ and $\omega_{\det}(S_B)$ is divisible by $n^{2\ell+2}$. Since $\text{id}(e_A) \neq \text{id}(e_B)$, we conclude that the two sides of (15) give different remainders modulo $n^{2\ell+2}$, a contradiction. \square

9.2 Randomized isolation scheme

We now proceed to the proof of Theorem 9.4. We begin with the construction of the weight function. First, for every $i \in [d]$ we choose a number $r_i \in [n^{10}]$ independently and uniformly at random. Thus $\bar{r} := (r_1, \dots, r_d)$ is a random vector, distributed uniformly in $\Omega := [n^{10}]^d$. For $\bar{\rho} \in \Omega$, we define the weight function $\omega_{\bar{\rho}}$ as follows:

$$\omega_{\bar{\rho}}(e) := \text{id}(e) \cdot \rho_{\text{lvl}(e)}.$$

Our isolation scheme simply samples \bar{r} as above and returns the weight function $\omega_{\bar{r}}$. Note that $\omega_{\bar{r}}$ assigns weights upper bounded by $\mathcal{O}(n^{10}\text{id}(e)) = \mathcal{O}(n^{12})$ and uses $\mathcal{O}(d \log n)$ random bits, as promised, so it remains to prove that $\omega_{\bar{r}}$ isolates $\text{MM}(G)$ with probability at least $\frac{1}{2}$.

Analysis. The argument is similar to that used in the proof of Theorem 8.3. For each $X \subseteq E$, we define a linear form $\phi^X: \Omega \rightarrow \mathbb{N}$ as

$$\phi^X(\bar{\rho}) := \sum_{e \in X} \text{id}(e) \cdot \rho_{\text{lvl}(e)}.$$

Thus, we have

$$\phi^X(\bar{\rho}) = \omega_{\bar{\rho}}(X). \tag{16}$$

Let

$$\Phi := \{\phi^X: X \in \text{MM}(G)\}.$$

Again, the key step is captured by the following lemma.

Lemma 9.6. *Suppose $\bar{\rho} \in \Omega$ is such that $\omega_{\bar{\rho}}$ does not isolate $\mathcal{P}(G)$. Then there are two different linear forms $\alpha, \beta \in \Phi$ such that*

$$\alpha(\bar{\rho}) = \beta(\bar{\rho}) = \min_{\gamma \in \Phi} \gamma(\bar{\rho}).$$

Proof. Since $\omega_{\bar{\rho}}$ does not isolate $\text{MM}(G)$, there are two different minimizers of $\omega_{\bar{\rho}}$ on $\text{MM}(G)$. By the exchange property, there are also two different minimizers A and B such that there exists exactly one pivotal vertex for A and B , say u . Since A and B are both minimizers, by (16) we have

$$\phi^A(\bar{\rho}) = \phi^B(\bar{\rho}) = \min_{\gamma \in \Phi} \gamma(\bar{\rho}).$$

Hence, it suffices to prove that $\phi^A \neq \phi^B$.

Let $i := \text{lvl}(u)$. We claim that the coefficients of ϕ^A and ϕ^B standing by the variable ρ_i are different. Letting $L := \{e \mid \text{lvl}(e) = i\}$, we see that these coefficients are respectively equal to

$$\sum_{e \in A \cap L} \text{id}(e) \quad \text{and} \quad \sum_{e \in B \cap L} \text{id}(e).$$

Now recall that u is the only pivotal vertex of A and B . Hence, as both A and B are matchings, we observe that $A \cap L$ and $B \cap L$ differ only in the edge that is incident to u (or lack thereof). Since the two edges incident to u in $A \cap L$ and $B \cap L$ have different identifiers (or one is non-existent), it follows that $\sum_{e \in A \cap L} \text{id}(e) \neq \sum_{e \in B \cap L} \text{id}(e)$. This concludes the proof. \square

Now observe that linear forms from Φ have coefficients upper bounded by $m \cdot n < n^4$. Hence, we can again combine Lemma 9.6 with Lemma 8.7 (applied for $N = n^2$) to infer that $\omega_{\bar{r}}$ isolates $\text{MM}(G)$ with probability at least $\frac{1}{2}$.

10 Directions for further research

In this paper we presented several isolation schemes for NP-complete problems, and we showed that analogues of decomposition-based methods such as Divide&Conquer can also be used to design more randomness-efficient isolation schemes. While we provide nearly matching lower bounds for all our results, at least as far as the number of random bits is concerned, we still leave open a number of interesting open questions:

1. Can we improve our isolation schemes to have weights that are only polynomial in n , while not increasing the number of used random bits? Note that in our approach, the use of large weights is crucial for the application of Lemma 3.1 that deals with interactions between different partial solutions in our isolation schemes.³
2. Can we shave off the log factors in the number of used random bits in our results? While some of the $\log n$ factors seem to be inherent in our ideas, there still might be a little room. For example, Melkebeek and Prakriya [55] presented an isolation scheme for reachability that uses $\mathcal{O}(\log^{1.5}(n))$ -random bits. Perhaps with their ideas one can get the same guarantees for isolating Hamiltonian cycles in constant treewidth graphs.
3. Does the (even more) natural isolation scheme work as well? Many of our isolation schemes draw several random prime numbers and assign a weight that is obtained by concatenating the congruence class of the vertex/edge identifier with respect to the different primes. A more natural, but possibly harder to analyse, scheme would be to sample a single (larger) prime number and define the weights to be the congruence classes of the identifiers with respect to that single prime.
4. Our methods allowed us to derandomize polynomial-space algorithms for H -minor free graphs without significantly increase the running time. Can our methods be used to derandomize other algorithms likewise?

References

- [1] M. Agrawal, R. Gurjar, and T. Thierauf. Impossibility of Derandomizing the Isolation Lemma for all Families. *Electron. Colloquium Comput. Complex.*, 27:98, 2020.
- [2] M. Agrawal, T. M. Hoang, and T. Thierauf. The Polynomially Bounded Perfect Matching Problem Is in NC². In *STACS 2007, 24th Annual Symposium on Theoretical Aspects of Computer Science*, pages 489–499, 2007.
- [3] M. Agrawal, N. Kayal, and N. Saxena. PRIMES is in P. *Annals of mathematics*, pages 781–793, 2004.
- [4] N. Alon, P. D. Seymour, and R. Thomas. A Separator Theorem for Graphs with an Excluded Minor and its Applications. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing, STOC 1990*, pages 293–299. ACM, 1990.

³In [11] a similar lemma was used to obtain isolation schemes with polynomial weights, but since the objects of the set family are not decomposed, the authors did not have this issue of interactions between different partial solutions.

- [5] R. Arora, A. Gupta, R. Gurjar, and R. Tewari. Derandomizing Isolation Lemma for $K_{3,3}$ -free and K_5 -free Bipartite Graphs. In *33rd Symposium on Theoretical Aspects of Computer Science, STACS 2016*, pages 10:1–10:15, 2016.
- [6] V. Arvind and P. Mukhopadhyay. Derandomizing the isolation lemma and lower bounds for circuit size. In *Approximation, Randomization and Combinatorial Optimization. Algorithms and Techniques*, pages 276–289. Springer, 2008.
- [7] A. Björklund. Determinant sums for undirected hamiltonicity. *SIAM J. Comput.*, 43(1):280–299, 2014.
- [8] H. L. Bodlaender, M. Cygan, S. Kratsch, and J. Nederlof. Deterministic single exponential time algorithms for connectivity problems parameterized by treewidth. *Inf. Comput.*, 243:86–111, 2015.
- [9] C. Bourke, R. Tewari, and N. V. Vinodchandran. Directed planar reachability is in unambiguous log-space. *ACM Trans. Comput. Theory*, 1(1):4:1–4:17, 2009.
- [10] C. Calabro, R. Impagliazzo, V. Kabanets, and R. Paturi. The complexity of Unique k -SAT: An Isolation Lemma for k -CNFs. *J. Comput. Syst. Sci.*, 74(3):386–393, 2008.
- [11] S. Chari, P. Rohatgi, and A. Srinivasan. Randomness-Optimal Unique Element Isolation with Applications to Perfect Matching and Related Problems. *SIAM J. Comput.*, 24(5):1036–1050, 1995.
- [12] J. Chen, W. Czerwiński, Y. Disser, A. E. Feldmann, D. Hermelin, W. Nadara, M. Pilipczuk, M. Pilipczuk, M. Sorge, B. Wróblewski, and A. Zych-Pawlewicz. Efficient fully dynamic elimination forests with applications to detecting long paths and cycles. *CoRR*, abs/2006.00571, 2020. To appear in the proceedings of SODA 2021.
- [13] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms, 3rd Edition*. MIT Press, 2009.
- [14] B. Courcelle and J. Engelfriet. *Graph Structure and Monadic Second-Order Logic — A Language-Theoretic Approach*, volume 138 of *Encyclopedia of mathematics and its applications*. Cambridge University Press, 2012.
- [15] M. Cygan, S. Kratsch, and J. Nederlof. Fast Hamiltonicity Checking Via Bases of Perfect Matchings. *J. ACM*, 65(3):12:1–12:46, 2018.
- [16] 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. In *52nd Annual Symposium on Foundations of Computer Science, FOCS 2011*, pages 150–159. IEEE, 2011.
- [17] W. Czerwiński, W. Nadara, and M. Pilipczuk. Improved Bounds for the Excluded-Minor Approximation of Treedepth. In *27th Annual European Symposium on Algorithms, ESA 2019*, volume 144 of *LIPICs*, pages 34:1–34:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.
- [18] S. Datta, R. Kulkarni, and S. Roy. Deterministically isolating a perfect matching in bipartite planar graphs. *Theory Comput. Syst.*, 47(3):737–757, 2010.
- [19] H. Dell, V. Kabanets, D. van Melkebeek, and O. Watanabe. Is Valiant-Vazirani’s isolation probability improvable? *Comput. Complex.*, 22(2):345–383, 2013.
- [20] V. Dujmović, D. Eppstein, and D. R. Wood. Structure of Graphs with Locally Restricted Crossings. *SIAM J. Discret. Math.*, 31(2):805–824, 2017.

- [21] Z. Dvořák, A. C. Giannopoulou, and D. M. Thilikos. Forbidden graphs for tree-depth. *Eur. J. Comb.*, 33(5):969–979, 2012.
- [22] F. Eisenbrand, C. Hunkenschroder, K. Klein, M. Koutecký, A. Levin, and S. Onn. An algorithmic theory of integer programming. *CoRR*, abs/1904.01361, 2019.
- [23] M. Elberfeld, M. Grohe, and T. Tantau. Where first-order and monadic second-order logic coincide. In *27th Annual IEEE Symposium on Logic in Computer Science, LICS 2012*, pages 265–274. IEEE Computer Society, 2012.
- [24] L. Esperet and J. Raymond. Polynomial expansion and sublinear separators. *Eur. J. Comb.*, 69:49–53, 2018.
- [25] S. Fenner, R. Gurjar, and T. Thierauf. Bipartite perfect matching is in quasi-NC. *SIAM Journal on Computing*, (0):STOC16–218, 2019.
- [26] 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(4):29:1–29:60, 2016.
- [27] M. L. Fredman, J. Komlós, and E. Szemerédi. Storing a Sparse Table with $\mathcal{O}(1)$ Worst Case Access Time. *J. ACM*, 31(3):538–544, 1984.
- [28] M. Fürer and H. Yu. Space saving by dynamic algebraization based on tree-depth. *Theory Comput. Syst.*, 61(2):283–304, 2017.
- [29] M. R. Garey, D. S. Johnson, and R. E. Tarjan. The planar Hamiltonian circuit problem is NP-complete. *SIAM Journal on Computing*, 5(4):704–714, 1976.
- [30] O. Goldreich. *Computational complexity - a conceptual perspective*. Cambridge University Press, 2008.
- [31] D. Grigoriev and M. Karpinski. The matching problem for bipartite graphs with polynomially bounded permanents is in NC (extended abstract). In *28th Annual Symposium on Foundations of Computer Science, Los Angeles, California, USA, 27-29 October 1987*, pages 166–172. IEEE Computer Society, 1987.
- [32] C. Gupta, V. R. Sharma, and R. Tewari. Efficient Isolation of Perfect Matching in $\mathcal{O}(\log n)$ Genus Bipartite Graphs. In *45th International Symposium on Mathematical Foundations of Computer Science (MFCS 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.
- [33] S. Har-Peled and K. Quanrud. Approximation algorithms for polynomial-expansion and low-density graphs. *SIAM J. Comput.*, 46(6):1712–1744, 2017.
- [34] F. Hegerfeld and S. Kratsch. Solving Connectivity Problems Parameterized by Treedepth in Single-Exponential Time and Polynomial Space. In *37th International Symposium on Theoretical Aspects of Computer Science, STACS 2020*, pages 29:1–29:16, 2020.
- [35] B. M. P. Jansen and J. Nederlof. Computing the chromatic number using graph decompositions via matrix rank. *Theor. Comput. Sci.*, 795:520–539, 2019.
- [36] V. A. T. Kallampally and R. Tewari. Trading determinism for time in space bounded computations. In P. Faliszewski, A. Muscholl, and R. Niedermeier, editors, *41st International Symposium on Mathematical Foundations of Computer Science, MFCS 2016, August 22-26, 2016 - Kraków, Poland*, volume 58 of *LIPICs*, pages 10:1–10:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.

- [37] R. M. Karp. Reducibility among combinatorial problems. In R. E. Miller and J. W. Thatcher, editors, *Proceedings of a symposium on the Complexity of Computer Computations, held March 20-22, 1972, at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, USA*, The IBM Research Symposia Series, pages 85–103. Plenum Press, New York, 1972.
- [38] K. Kawarabayashi and B. A. Reed. A separator theorem in minor-closed classes. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010*, pages 153–162. IEEE Computer Society, 2010.
- [39] J. Li and J. Nederlof. Detecting Feedback Vertex Sets of Size k in $\mathcal{O}^*(2.7^k)$ Time. In S. Chawla, editor, *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 971–989. SIAM, 2020.
- [40] R. J. Lipton and R. E. Tarjan. Applications of a planar separator theorem. *SIAM J. Comput.*, 9(3):615–627, 1980.
- [41] J. A. Montoya and M. Müller. Parameterized random complexity. *Theory Comput. Syst.*, 52(2):221–270, 2013.
- [42] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, 1995.
- [43] K. Mulmuley, U. V. Vazirani, and V. V. Vazirani. Matching is as easy as matrix inversion. *Comb.*, 7(1):105–113, 1987.
- [44] J. Nederlof, M. Pilipczuk, C. M. F. Swennenhuis, and K. Węgrzycki. Hamiltonian cycle parameterized by treedepth in single exponential time and polynomial space. In *46th International Workshop on Graph-Theoretic Concepts in Computer Science, WG 2020*, volume 12301 of *Lecture Notes in Computer Science*, pages 27–39. Springer, 2020.
- [45] J. Nešetřil and P. Ossona de Mendez. *Sparsity — Graphs, Structures, and Algorithms*, volume 28 of *Algorithms and combinatorics*. Springer, 2012.
- [46] M. Pilipczuk and S. Siebertz. Polynomial bounds for centered colorings on proper minor-closed graph classes. In *30th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019*, pages 1501–1520. SIAM, 2019.
- [47] M. Pilipczuk and M. Wrochna. On space efficiency of algorithms working on structural decompositions of graphs. *ACM Trans. Comp. Theory*, 9(4):18:1–18:36, 2018.
- [48] S. A. Plotkin, S. Rao, and W. D. Smith. Shallow excluded minors and improved graph decompositions. In *Proceedings of the Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 1994*, pages 462–470. ACM/SIAM, 1994.
- [49] K. Reinhardt and E. Allender. Making nondeterminism unambiguous. *SIAM J. Comput.*, 29(4):1118–1131, 2000.
- [50] B. Rosser. Explicit bounds for some functions of prime numbers. *American Journal of Mathematics*, 63(1):211–232, 1941.
- [51] O. Svensson and J. Tarnawski. The matching problem in general graphs is in Quasi-NC. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017*, pages 696–707, 2017.
- [52] S. Toda. On the computational power of PP and $\oplus P$. In *30th Annual Symposium on Foundations of Computer Science (FOCS 1989)*, pages 514–519, 1989.

- [53] P. Traxler. The time complexity of constraint satisfaction. In M. Grohe and R. Niedermeier, editors, *Parameterized and Exact Computation*, pages 190–201, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg.
- [54] L. G. Valiant and V. V. Vazirani. NP is as Easy as Detecting Unique Solutions. *Theor. Comput. Sci.*, 47(3):85–93, 1986.
- [55] D. van Melkebeek and G. Prakriya. Derandomizing Isolation in Space-Bounded Settings. *SIAM J. Comput.*, 48(3):979–1021, 2019.
- [56] N. Vyas and R. R. Williams. On super strong ETH. In *Theory and Applications of Satisfiability Testing - SAT 2019 - 22nd International Conference, SAT 2019*, pages 406–423, 2019.
- [57] A. Wigderson. $NL/poly \subseteq \oplus L/poly$ (preliminary version). In *Proceedings of the Ninth Annual Structure in Complexity Theory Conference*, pages 59–62, 1994.
- [58] R. Williams. Finding paths of length k in $\mathcal{O}^*(2^k)$ time. *Inf. Process. Lett.*, 109(6):315–318, 2009.