Counting Small Induced Subgraphs with Edge-monotone Properties

Simon Döring
Max Planck Institute for Informatics and
Saarbrücken Graduate School of Computer Science, SIC,
Saarbrücken, Germany

Dániel Marx
CISPA Helmholtz Center for Information Security,
Saarbrücken, Germany

Philip Wellnitz
Max Planck Institute for Informatics, SIC,
Saarbrücken, Germany

Abstract

We study the parameterized complexity of #IndSub(Φ), where given a graph G and an integer k, the task is to count the number of induced subgraphs on k vertices that satisfy the graph property Φ. Focke and Roth [STOC 2022] completely characterized the complexity for each Φ that is a *hereditary property* (that is, closed under vertex deletions): #IndSub(Φ) is #W[1]-hard except in the degenerate cases when every graph satisfies Φ or only finitely many graphs satisfy Φ. We complement this result with a classification for each Φ that is *edge monotone* (that is, closed under edge deletions): #IndSub(Φ) is #W[1]-hard except in the degenerate case when there are only finitely many integers k such that Φ is nontrivial on k-vertex graphs. Our result generalizes earlier results for specific properties Φ that are related to the connectivity or density of the graph.

Further, we extend the #W[1]-hardness result by a lower bound which shows that #IndSub(Φ) cannot be solved in time $f(k) \cdot |V(G)|^{o(\sqrt{\log k}/\log\log k)}$ for any function f, unless the Exponential-Time Hypothesis (ETH) fails. For many natural properties, we obtain even a tight bound $f(k) \cdot |V(G)|^{o(k)}$; for example, this is the case for every property Φ that is nontrivial on k-vertex graphs for each k greater than some k_0.

Acknowledgements We thank Jacob Focke and Marc Roth for insightful discussions. We also thank Rebekka Burkholz for providing us with an Overleaf account.

Contents

1 Introduction

Searching and counting patterns is one of the oldest algorithmic tasks in computer science and has many applications in other scientific fields. The theoretical notion of graphs give a useful and widely used way to model various types of data. In this work, we focus on the parameterized complexity of counting small-size patterns in graphs, which is motivated by related applications in the study of database systems [GSS01], neural and social networks [MSOI+02, SJHS15], biology [SS05], and many other fields.

When counting patterns (such as cycles or connected graphs) of some small size k in an n-vertex graph, one can use an obvious brute-force approach that enumerates all $O(n^k)$ subsets of k vertices in the graph. While this is polynomial for fixed k, it would be desirable to have a running time less sensitive to the size k of the pattern. Flum and Grohe [FG04] initiated the study of the parameterized complexity of counting problems, with the goal of determining which counting problems are *fixed-parameter tractable (FPT)*, that is, can be solved in time $f(k)n^{O(1)}$ for some computable function f. It is widely believed that many basic counting problems, such as #Clique are *not* FPT. The notion of #W[1]-hardness can be used to give evidence that a counting problem is unlikely to be FPT by showing that it is at least as hard as #Clique. For example, it is known that counting paths of length k, cycles of length k, matchings of size k, and many other types of subgraphs are #W[1]-hard [FG04, Cur13, CM14, CDM17].

While the parameterized complexity of counting different types of subgraphs is well understood, the complexity of more general properties of patterns is far from clear. Formally, a *graph property* is a computable function Φ from the set of graphs to $\{0, 1\}$ that is invariant under relabeling. Common examples are *is bipartite*, *is clique*, *is independent set*, *is connected*, or *is planar*, to name but a few. Given a property Φ, we would like to count the number of subsets of vertices of a specified size that induce a graph with this property. That is, for a fixed graph property Φ, Jerrum and Meeks [JM15a, JM15b] introduced the #IndSub(Φ) problem, where given a graph G and a nonnegative integer k, the task is to compute the number of induced subgraphs of G of size k that satisfy Φ. We denote this number by $\#\mathrm{IndSub}((\Phi, k) \to G)$ and write $\#\mathrm{IndSub}((\Phi, k) \to \star)$ for the function that maps G to $\#\mathrm{IndSub}((\Phi, k) \to G)$. This problem for specific types of properties Φ has been in the focus of a large amount of research in recent years [JM17, CDM17, RS20, RSW23, DRSW22, FR22].

It is obvious that #IndSub(Φ) is #W[1]-hard if Φ is the graph property *is a complete graph* since $\#\mathrm{IndSub}((\Phi, k) \to G)$ is the number of k-cliques. However, it turns out that #IndSub(Φ) is also #W[1]-hard for many other graph properties. Early works in that area focused on showing #W[1]-hardness for specific graph properties [JM15a, JM17], and it looks like #IndSub(Φ) is #W[1]-hard for all possible graph properties Φ except trivial ones. In this setting, we say that Φ *is trivial on k* if it is constant on k-vertex graphs, hence $\#\mathrm{IndSub}((\Phi, k) \to G)$ is either 0 or $\binom{|V(G)|}{k}$. We say that Φ is *trivial* if there is an N such that Φ is trivial on all $k \geq N$. It is easy to verify that #IndSub(Φ) is FPT whenever Φ is trivial.

Conjecture 1.1 ([FR22, RSW23]). *For all nontrivial, computable graph properties Φ the #IndSub(Φ) problem is #W[1]-hard. Otherwise, #IndSub(Φ) is FPT.*

Curticapean, Dell, and Marx showed in [CDM17] that #IndSub(Φ) is always FPT or #W[1]-hard. However, the proof does not give an easy way to determine which case holds for a given property Φ. As we review it in Section 1.1, Conjecture 1.1 is known to hold for many classes of graph properties. In particular, Focke and Roth [FR22] showed that Conjecture 1.1 holds if Φ is a *hereditary* property, that is, closed under deletion of vertices. In this paper, we show a complementary result for graph properties that are *edge monotone* (closed under deletion of edges). We also provide quantitative lower bounds on the exponent of n under the *Exponential-Time Hypothesis (ETH)* [IPZ01].

Main Theorem 1. *Let Φ denote a nontrivial edge-monotone graph property.*

- *The problem #IndSub(Φ) is #W[1]-hard.*
- *Further, assuming ETH, there is a universal constant $\gamma > 0$ (independent of Φ) such that for any integer $k \geq 3$ on which Φ is nontrivial, no algorithm (that reads the whole input) computes for every graph G the number $\#\mathrm{IndSub}((\Phi, k) \to G)$ in time $O(|V(G)|^{\gamma\sqrt{\log k}/\log\log k})$.*

As an example, the following nontrivial graph properties are all edge-monotone and not covered by any method in Section 1.1, but Main Theorem 1 shows their #W[1]-hardness.

Example 1.2.

- $\Phi_1^c(G) = 1$ if and only if G is disconnected or $\mathrm{diam}(G) \geq c|V(G)|$ for a fixed constant $c \in (0, 1)$.
- $\Phi_2^c(G) = 1$ if and only if G is bipartite or contains an independent set of size at least c for a fixed constant $c \geq 3$.
- $\Phi_3^c(G) = 1$ if and only if the maximum degree of G is at most $c|V(G)|$ for a fixed constant $c \in (1/2, 1)$.

A notable difference between hereditary and edge-monotone properties is that an edge-monotone property might be nontrivial only on a sparse set of integers: for example, one can define $\Phi(G)$ to be 1 if and only if G is a clique with $|V(G)|$ being a power of two (or a prime, or the product of the first i primes, or ...). On the other hand, if Φ is a nontrivial hereditary property, then it is an easy exercise to show that it has to be nontrivial for every k larger than some N. Our algebraic proof techniques for Main Theorem 1 are very sensitive to the number-theoretic properties of the size k, hence it is a significant challenge to make it work when Φ can be nontrivial only on certain integers.

In [RS20, RSW23, DRSW22, FR22], it was proven that for specific classes of nontrivial graph properties Φ, there is a function g such that the problem #IndSub(Φ) cannot be solved in time $f(k) \cdot |V(G)|^{o(g(k))}$ for any computable function f, unless ETH fails. The second part of Main Theorem 1 shows a similar lower bound on the exponent of the running time for edge-monotone properties. However, readers familiar with the way ETH-based lower bounds are stated in the parameterized complexity literature should notice that the second part of Main Theorem 1 uses a very different formulation (a similar formulation was given by Cohen-Addad et al. [CCMdM21] for various problems). We use a stronger approach by showing that, assuming ETH, there is a constant $\gamma > 0$ such that for all edge-monotone graph properties Φ and *any* fixed nontrivial $k \geq 3$ the function $\#\mathrm{IndSub}((\Phi, k) \to G)$ cannot be computed for all graphs G in time $O(|V(G)|^{\gamma g(k)})$ unless ETH fails. That is, our stronger statement applies not only to algorithms solving the problem in general, but also gives a meaningful statement for algorithms solving the problem for a fixed k. It is easy to observe that this approach implies that no algorithm solves #IndSub(Φ) in time $f(k) \cdot |V(G)|^{o(g(k))}$ for any computable function f, unless ETH fails.

Corollary 1.3. *Let Φ denote a nontrivial edge-monotone graph property.*

- *The problem #IndSub(Φ) is #W[1]-hard.*
- *Further, assuming ETH, no algorithm computes for every graph G and every positive integer k the number $\#\mathrm{IndSub}((\Phi, k) \to G)$ in time $f(k) \cdot |V(G)|^{o(\sqrt{\log k}/\log\log k)}$ for any computable function f.*

Proof. Suppose that #IndSub(Φ) can be solved in $O(f(k)|V(G)|^{g(k)})$ for $g(k) \in o(\sqrt{\log k}/\log\log k)$. Then, there is an N with $g(k') \leq \gamma\sqrt{\log k'}/\log\log k'$ for all $k' \geq N \geq 3$. Since Φ is nontrivial, there is a $k \geq N$ such that Φ is nontrivial on k. Main Theorem 1 shows that no algorithm solves #IndSub(Φ) in $O(f(k)|V(G)|^{g(k)})$. ∎

Thus our formulation of the second part of Main Theorem 1 implies the usual formulation. Further, even though the statement is stronger, it turns out that it is somewhat easier to work with this formulation: technicalities involving the little-o notation and the function $f(k)$ disappear, which results in more streamlined proofs.

Observe that quantitative lower bounds for #IndSub(Φ) in Main Theorem 1 are fairly weak: we say that the exponent cannot be much better than $\sqrt{\log k}/\log\log k$. Naturally, we would like to show lower bounds for #IndSub(Φ) of the form $O(|V(G)|^{\gamma k})$, that is, we would like to show bounds that are tight in the sense that we can indeed solve #IndSub(Φ) in time $O(|V(G)|^{k})$ using a brute force approach. While we cannot prove such tight lower bounds in full generality, we are able to obtain tight lower bounds for specific cases that cover most properties of interest.

Main Theorem 2. *For each prime p, there is a constant $\gamma_p > 0$ such that for each integer m with $p^m \geq 3$ and each edge-monotone graph property Φ that is nontrivial on p^m, no algorithm (that reads the whole input) computes for every graph G the number* $\#\mathrm{IndSub}((\Phi, k) \to G)$ *in time $O(|V(G)|^{\gamma_p p^m})$, unless ETH fails.*

As before, we also restate Main Theorem 2 in terms of ruling out $o(k)$ in the exponent.

Corollary 1.4. *For all edge-monotone Φ that for fixed prime number p are nontrivial on infinitely many numbers of the form p^m, no algorithm computes for every graph G end every positive integer k the number* $\#\mathrm{IndSub}((\Phi, k) \to G)$ *in time $f(k) \cdot |V(G)|^{o(k)}$ for any computable function f, unless ETH fails.*

Proof. Suppose that #IndSub(Φ) can be solved in $O(f(k)|V(G)|^{g(k)})$ for $g(k) \in o(k)$. Then, there is an N with $g(k') \leq \gamma k$ for all $k' \geq N \geq 3$. Since Φ is nontrivial, there is a $k \geq N$ with Φ is nontrivial on k. Main Theorem 2 shows that no algorithm solves #IndSub(Φ) in $O(f(k)|V(G)|^{g(k)})$.

Observe that Corollary 1.4 holds whenever there is a constant N such that Φ is nontrivial on all $k \geq N$. It is easy to check that this condition holds for all Φ_i in Example 1.2. Thus, we obtain tight lower bounds for all #IndSub(Φ_i) using Corollary 1.4.

1.1 Prior Work

In the following, we summarize recent results for the parameterized complexity of the #IndSub(Φ) problem. The results are ordered by their date of publication.

(a) #IndSub(Φ) is #W[1]-hard for
  - **(i)** $\Phi(G) = 1$ if and only if G is connected [JM15a]
  - **(ii)** $\Phi(G) = 1$ if and only if $|E(G)|$ is even, $\Phi(G) = 1$ if and only if $|E(G)|$ is odd (both in [JM17])

(b) In [JM15b], Jerrum and Meeks proved #W[1]-hardness if the number of distinct edge densities of graphs that satisfy Φ is low. This also shows that #IndSub(Φ) is #W[1]-hard whenever Φ is minor-closed.

(c) In [Mee16], Meeks proved that #IndSub(Φ) is #W[1]-hard if Φ is closed under the addition of edges and the edge-minimal graphs of Φ have unbounded treewidth. A Graph is edge-minimal if $\Phi(G) = 1$ and $\Phi(G') = 0$ for all proper edge-subgraphs of G.

(d) In [RS20], Roth and Schmitt proved that #IndSub(Φ) is #W[1]-hard if Φ is nontrivial, edge-monotone and satisfies at least one of the following conditions.
  - **(i)** Φ is false for odd cycles. A cycle has vertex set $\{0, \dots, n-1\}$ and $\{a, b\}$ is an edge if and only if $a - b \equiv_n 1$.

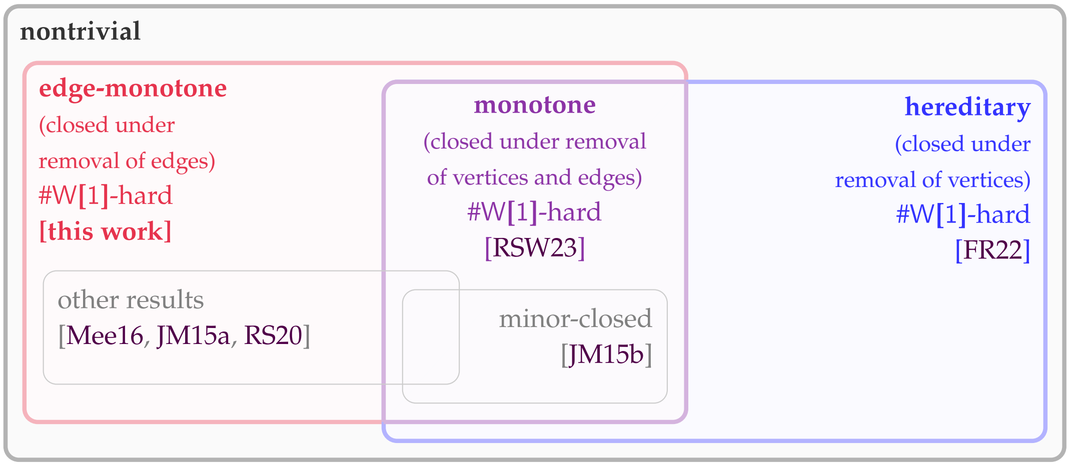

Figure 1 Hierarchy of classes of graph properties Φ, together with the results that show #W[1]-hardness for the corresponding problem #IndSub(Φ) (if such results exist).

(ii) Φ is true for odd anti-holes. An anti-hole is the complement graph of a cycle.
(iii) There is a $c \in \mathbb{N}$ such that $\Phi(H) = 1$ if and only if H is not c-edge-connected
(iv) There is a graph F such that $\Phi(H) = 1$ if and only if there is no homomorphism from F to H.

(e) In [DRSW22], Dörfler, Roth, Schmitt, and Wellnitz proved that #IndSub(Φ) is #W[1]-hard if there are infinitely many prime powers t such that $\Phi(K_{t,t}) \neq \Phi(\mathrm{IS}_{2t})$.

(f) In [RSW23], Roth, Schmitt, and Wellnitz proved the following criteria for checking #W[1]-hardness. Let $f_i^{\Phi,k} := \#\{A \subseteq E(K_k) : \#A = i \wedge \Phi(K_k[A]) = 1\}$ denote a vector and $\mathrm{hw}(f^{\Phi,k}) := \#\{i : f_i^{\Phi,k} \neq 0\}$ the hamming weight of $f^{\Phi,k}$. We define the function $\beta\colon \mathcal{K}(\Phi) \to \mathbb{Z}_{\geq 0}; k \mapsto \binom{k}{2} - \mathrm{hw}(f^{\Phi,k})$, where $\mathcal{K}(\Phi)$ is the set of $n \in \mathbb{N}$ with Φ is nontrivial on n. The problem #IndSub(Φ) is #W[1]-hard if $\beta(k) \in \omega(k)$. They also proved that #IndSub(Φ) is #W[1]-hard if Φ is *monotone*, meaning closed under taking subgraphs.

(g) In [FR22], Focke and Roth proved that #IndSub(Φ) is #W[1]-hard if Φ is nontrivial and hereditary. A property is called *hereditary* if it is closed under vertex-deletion, meaning that if G satisfies Φ, then each induced subgraph of G also satisfies Φ.

(h) Lastly, we observe that the counting problem #IndSub(Φ) is #W[1]-hard if and only if #IndSub(¬Φ) is #W[1]-hard; and #IndSub(Φ) is #W[1]-hard if and only if #IndSub($\overline{\Phi}$) is #W[1]-hard. Here, $\neg\Phi(G) := 1 - \Phi(G)$; and $\overline{\Phi}(G) := \Phi(\overline{G})$, where $\overline{G}$ is the complement of G (see [RSW23, see Fact 2.3]). This means that we can prove #W[1]-hardness of #IndSub(Φ) by analyzing #IndSub(¬Φ) or #IndSub($\overline{\Phi}$).

1.2 High-level Ideas

Our results build on concepts introduced in earlier work, but we need to develop substantial new technical ideas to be able to deploy them in our setting. We briefly review these concepts here and highlight our main new technical contributions; a more detailed technical overview follows in Section 2.

Given a graph property Φ, we define the *alternating enumerator* $\widehat{\Phi}$ as

$$\widehat{\Phi}(H) := \sum_{S \subseteq E(H)} \Phi(H\{S\})(-1)^{\#S}.$$

Here, $H\{S\}$ denotes the subgraph of H that contains the same set of vertices but only the edge set S. Dörfler et al. [DRSW22] proved that #IndSub(Φ) is #W[1]-hard if there are graphs with arbitrary large treewidth and nonzero alternating enumerator. Thus, #W[1]-hardness can be established by showing that such graphs exist. However, the complicated definition of $\widehat{\Phi}(H)$ does not make this task easy, even for a specific property Φ. Dörfler et al. [DRSW22] made the following observation that can help in arguing that $\widehat{\Phi}(H)$ is nonzero. Write Γ for a group that consists in automorphisms of $V(H)$ and suppose that the order of Γ is a power of p. We say that a subset $S \subseteq E(H)$ is a fixed point with respect to Γ if every automorphism in Γ moves every edge in S to an edge in S. Clearly, we can show that $\widehat{\Phi}(H)$ is nonzero by showing that it is nonzero modulo p. The main observation is that if we want to compute the sum in $\widehat{\Phi}(H)$ modulo p, then we do not need to sum over all sets $S \subseteq E(H)$: it is sufficient to sum over the fixed points with respect to Γ, as the other subgraphs somehow cancel out. Thus, #W[1]-hardness can be established by finding graphs H with large treewidth and appropriate groups Γ to show that $\widehat{\Phi}(H)$ is nonzero modulo some prime p.

It is not difficult to show that every fixed point with respect to Γ is the disjoint union of orbits of edges. This means that the fixed points have a natural lattice structure: we can imagine the fixed points that are the disjoint union of ℓ orbits as the ℓ-th level of the lattice. In broad terms, our approach is to define some group Γ, consider the fixed points of the complete graph on k vertices with respect to Γ, and then try to find a fixed point of sufficiently high level whose alternating enumerator is nonzero. Fixed points on higher levels have more edges and hence larger treewidth. Now, we invoke (adaptations of) the earlier results of Dörfler et al. [DRSW22] to show hardness with such graphs. In more detail, our proofs are based on the following four main technical ideas.

(1) **Duality of the highest nonzero level.** We use linear algebra arguments to show that if Φ is 0 on every fixed point on the topmost c levels, then there is a fixed point on level $\geq c$ whose alternating enumerator is nonzero, and the level implies that treewidth is at least c. Thus, we may assume that there is a fixed point of fairly high level that is nonzero in Φ. Note that this statement is true for every property Φ, even if Φ is not edge-monotone.

(2) **Avalanche effect for difference graphs on** $\mathbb{F}_{p^m}$. Write A for a subset of $\mathbb{F}_{p^m}$. Then we can define the difference graph on the vertex set $\mathbb{F}_{p^m}$, where there is an edge between x and y if and only if $x - y \in A$. Write Γ for the additive group of $\mathbb{F}_{p^m}$. One can observe that the fixed points with respect to Γ are exactly the difference graphs. If Φ is edge-monotone and Φ is nonzero on a fixed point S, then this implies that Φ should be nonzero on all the other fixed points that are subsets of S'. In particular, a single nonzero fixed point on one of the top c levels starts an "avalanche" that forces every fixed point on a level of at most roughly p^m/c to be nonzero in Φ. The proof is based on the fact that $\mathbb{F}_{p^m}$ multiplication is an isomorphism of the difference graph. Thus, if we pick the fixed point with the lowest level that does not satisfy Φ, then it has a level of at least p^m/c and hence fairly large treewidth. A simple calculation using the binomial theorem shows that the any fixed point of the lowest level that do not satisfy Φ always has a nonzero alternating enumerator.

We combine the previous two ideas to obtain #W[1]-hardness if Φ is nontrivial on infinitely many prime powers. If $k = p^m$, then we use the following win/win approach: if Φ is 0 on the topmost $c = \sqrt{p^m}$ levels, then (1) gives a fixed point of treewidth at least $\sqrt{p^m}$ with nonzero alternating enumerator; otherwise, (2) gives such a fixed point. To obtain the #W[1]-hardness result in Main Theorem 1, we extend our proof to the case when $k = d \cdot p^m$ with the following idea. We also prove the ETH-based quantitative part of Main Theorem 1 by observing that the largest prime power divisor of k is always at least logarithmic in k.

(3) Product construction and reduction. We extend the lower bound for properties that are nontrivial on some prime powers to the general case in the following way. Write Γ for a group of permutations on p^m vertices. There is a natural way to raise Γ to a group Γ^d of permutations on $d \cdot p^m$ vertices. We observe that the fixed points of Γ^d can be described as the disjoint unions of the fixed points of Γ, plus additionally fully connecting some pairs of these fixed points. Let us take a look at the fixed point F with the lowest level that is zero in Φ (which we know to have nonzero alternating enumerator). If F contains one of the aforementioned full connections, then F has high treewidth, which is what we wanted. Otherwise, F is the disjoint union of fixed points of Γ; write F' for one of them and write H for the union of the remaining $d-1$ fixed points. Then by "pinning" H we can define a nontrivial property Φ' on p^m: set $\Phi'(G) := \Phi(H \uplus G)$, which is zero on the p^m-vertex graph F'. Now a standard reduction based on the Inclusion-Exclusion principle shows how to reduce #IndSub(Φ') to #IndSub(Φ), hence the lower bounds for prime powers can be used.

Unfortunately, a square root loss between k and the treewidth of the identified fixed point is inevitable when using (1) and (2), hence they cannot lead to tight bounds. The last idea allows us to obtain tight bounds in at least some cases.

(4) Avalanche effect for lexicographic product of graphs. For a prime p and integer $m \geq 2$, we show another way of defining fixed points on p^m vertices that have better avalanche properties. We define a group Γ (which is in fact the Sylow p-group of the automorphism group of K_{p^m}) such that the fixed points with respect to Γ are exactly the so-called m-dimensional lexicographic products of difference graphs on $\mathbb{F}_p$. Let us consider the lowest level ℓ that contains a fixed point where Φ is 0. We observe that every level ℓ has fixed points that contain the complete bipartite graph $K_{p^{m-1},p^{m-1}}$. If Φ is 0 on one such fixed point F on level ℓ, then it is easy to show that the alternating enumerator of F is nonzero. Now, earlier work shows how to reduce the counting of p^{m-1}-cliques to #IndSub(Φ) with $k = p^m$. Otherwise, if Φ is nonzero on every such fixed point on level ℓ, then the avalanche effect shows that Φ is nonzero also on every other fixed point on the same level ℓ, a contradiction.

Specifically, if $k = p^m$ for some constant prime p, then we obtain the tight bound that #IndSub(Φ) for k is at least as hard as counting p^{m-1}-cliques, which proves Main Theorem 2.

2 Technical Overview

In this section, we present an overview of the most important techniques and ideas that we use to show #W[1]-hardness of #IndSub(Φ) for each nontrivial edge-monotone graph property Φ, as well as the ETH-based quantitative lower bounds. We start with a review of the techniques that we use from previous work and then elaborate on our novel technical ideas.

Alternating Enumerator, #W[1]-hardness, and Lower Bounds

A problem instance of #Hom($\mathcal{H}$) is a pair of a graph $H \in \mathcal{H}$ and a graph $G \in \mathcal{G}$ and the output is the number of homomorphisms from H to G (that is, #Hom($H \to G$)); we parameterize by $\kappa(H, G) := |V(H)|$. Dalmau and Jonsson [DJ04] proved that #Hom($\mathcal{H}$) is #W[1]-hard if and only if the treewidth of the set $\mathcal{H}$ is unbounded (that is, there is no constant c such that the treewidth of all elements in $\mathcal{H}$ is below c).

Further results follow from the fact that a certain colored version of Hom($\mathcal{H}$) can be reduced to #Hom($\mathcal{H}$) [CCMdM21, Mar10]. We prove our #W[1]-hardness results by using a parameterized Turing reduction from #Hom($\mathcal{H}$) to #IndSub(Φ), which was first developed by Dörfler et al. [DRSW22].

For this reduction to work, the methods of [DRSW22] require that the *alternating enumerator* $\widehat{\Phi}(H)$ is nonvanishing for each graph $H \in \mathcal{H}$.[1]

Definition 2.1. *For a graph property* Φ *and a graph* H*, we define the* alternating enumerator $\widehat{\Phi}(H)$ *as*

$$\widehat{\Phi}(H) := \sum_{S \subseteq E(H)} \Phi(H\{S\})(-1)^{\#S}.$$

We say that a graph H *is* nonvanishing for a graph property Φ *if the alternating enumerator of* Φ *and* H *is nonzero. We say that a sequence of graphs* H_k *is nonvanishing for a graph property* Φ *if every* H_k *is nonvanishing for* Φ*, that is, if for all* k*, we have* $\widehat{\Phi}(H_k) \neq 0$*.*

This means that we can show #W[1]-hardness by finding a nonvanishing sequence H_k that has unbounded treewidth. Further, assuming ETH, the reduction of [DRSW22] yields a lower bound for #IndSub$((\Phi, k) \to G)$ for a fixed k. To be more precise, we show that if H is nonvanishing for Φ, we can use the reduction of [DRSW22] to solve Hom$(\{H\})$ using an oracle for #IndSub$((\Phi, |V(H)|) \to \star)$.

From the work of Cohen-Addad et al. [CCMdM21], we obtain that #Hom$(\{H\})$ cannot be solved in time $O(n^{\alpha_{\text{Hom}} \cdot \text{tw}(H)/\log \text{tw}(H)})$, where n is the number of vertices of the input graph. Hence, if we could compute #IndSub$((\Phi, |V(H)|) \to \star)$ fast enough, then our reduction shows that we can solve Hom$(\{H\})$ in time $O(n^{\alpha_{\text{Hom}} \cdot \text{tw}(H)/\log \text{tw}(H)})$, which contradicts ETH. Summarizing the previous discussion, we obtain the following lemma, which essentially follows from previous work. For completeness, we include a proof in Appendix A.

Lemma A.8 ([DRSW22, CCMdM21])**.** *Let* Φ *denote a nontrivial graph property.*

- *If there is a sequence of graphs with unbounded treewidth where each graph has an alternating enumerator that is nonvanishing for* Φ*, then* #IndSub(Φ) *is* #W[1]*-hard.*
- *Assuming ETH, there is a universal constant* $\alpha_{IndSub} > 0$ *(that is independent of* Φ*) such that for any positive integer* k *for which there is a graph* H_k *with* k *vertices,* $\widehat{\Phi}(H_k) \neq 0$*, and* $\text{tw}(H_k) \geq 2$*, no algorithm (that reads the whole input) computes for every graph* G *the number* #IndSub$((\Phi, k) \to G)$ *in time* $O(|V(G)|^{\alpha_{IndSub}\, \text{tw}(H_k)/\log \text{tw}(H_k)})$*.*

For edge-monotone graph properties Φ, the alternating enumerator of the k-clique $\widehat{\Phi}(K_k)$ is equal to the reduced Euler characteristic of the simplicial graph complex $\hat{\chi}(\Delta(\Phi_k))$, where Φ_k is Φ restricted on k-vertex graphs (see [DRSW22, Lemma 14]). The reduced Euler characteristic is in turn closely related to Karp's famous evasiveness conjecture (see [KSS84]), which conjectures that each nontrivial edge-monotone Φ_k is evasive. This conjecture holds if the reduced Euler characteristic $\hat{\chi}(\Delta(\Phi_k))$ is nonvanishing (see [DRSW22, Theorem 4] and [KSS84]).

However, this means that the computation of $\widehat{\Phi}(H)$ is highly nontrivial, which makes it hard to apply Lemma A.8. Fortunately for us, it suffices to show that $\widehat{\Phi}(H) \not\equiv_p 0$ for a prime number p—which turns out to be easier. In particular, as observed by Dörfler et al. [DRSW22], for a prime p, we can compute $\widehat{\Phi}(H)$ mod p in an elegant way using the fixed points of a p-subgroup Γ of Aut(H) when acting on edge-subgraphs of H. Thus, we heavily rely on the following lemma (whose proof is presented in Appendix A for completeness).

Lemma A.1 ([DRSW22])**.** *Let* H *denote a graph and let* $\Gamma \subseteq \text{Aut}(H)$ *denote a* p*-group, then*

$$\widehat{\Phi}(H) \equiv_p \sum_{A \in \text{FP}(\Gamma, H)} \Phi(A)(-1)^{\#E(A)}.$$

[1] In [RS20, DRSW22], the authors use $\hat{\chi}(\Phi, H)$ to denote the alternating enumerator.

Combining Lemmas A.1 and A.8, we immediately obtain the following tool to show hardness.

Corollary 2.2. *Let Φ denote a graph property and let (H_k) denote a sequence of graphs such that*
- *(H_k) has unbounded treewidth and*
- *for each graph H_k, there is a prime p_k and a p_k-group $\Gamma_k \subseteq \mathrm{Aut}(H_k)$ such that*

$$\widehat{\Phi}(H_k) \equiv_{p_k} \sum_{A \in \mathrm{FP}(\Gamma_k, H_k)} \Phi(A)(-1)^{\#E(A)} \quad \text{is nonzero modulo } p_k.$$

Then, $\#\textsc{IndSub}(\Phi)$ is $\#\mathrm{W}[1]$-hard.

A fixed point A of Γ in H is an edge-subgraph of H such that $gA = A$ for all $g \in \Gamma$. We use $\mathrm{FP}(\Gamma, H)$ to denote the set of all fixed points. The advantage of this approach is that the set of fixed points $\mathrm{FP}(\Gamma, H)$ is usually much smaller than the set of all edge-subgraphs of H. As it turns out, the set $\mathrm{FP}(\Gamma, H)$ itself has a natural lattice structure, which we exploit to find nonvanishing graphs with large treewidth.

Fixed Points as a Union of Orbits

Our goal is to find, for a given graph property Φ and value k, a nonvanishing graph H with k vertices that has *large* treewidth. To find these graphs, we analyze the fixed point structure of a certain graph H under a certain p-group $\Gamma \subseteq \mathrm{Aut}(H)$.

In Section 4, we introduce a systematic way to analyze and describe the fixed points of a group Γ and a graph H. Write $E(H)/\Gamma := \{O_1, \dots, O_s\}$ for the orbits of the group action $\cdot\colon \Gamma \times E(H)$ that maps $g \in \Gamma$ and $\{u, v\} \in E(H)$ to $\{g(u), g(v)\}$.

Our first observation is that every fixed point $F \in \mathrm{FP}(\Gamma, H)$ decomposes into a set of orbits of $E(H)/\Gamma$; that is, we have $V(F) = V(H)$ and $E(F) = \cup_{i \in A} O_i$ for some $A \subseteq [s]$. This means that the orbits $E(H)/\Gamma$ are the basic building blocks of the fixed points $\mathrm{FP}(\Gamma, H)$.

Lemma 4.1. *Let H denote a graph and let $\Gamma \subseteq \mathrm{Aut}(H)$ denote a group. Further, let Γ act on $E(H)$ and write $E(H)/\Gamma$ for the set of all resulting orbits. Finally, let Γ act on edge-subgraphs of H and write $\mathrm{FP}(\Gamma, H)$ for the set of all resulting fixed points of Γ in H.*

Then, the edge set of each fixed point $F \in \mathrm{FP}(\Gamma, H)$ is the (possible empty) disjoint union of orbits $O_1, \dots, O_s$, where $O_i \in E(H)/\Gamma$ and each disjoint union of orbits yields a fixed point. The partition into orbits is unique.

The *level* of a fixed point F, denoted by $\ell(F)$, is the number of orbits that F is made up of. The level allows us to classify the fixed points into *low* fixed points (consisting of a few orbits) and *high* fixed points (consisting of many orbits). Usually, fixed points with a high level also have a high treewidth since they contain more edges. Additionally, if Γ is a p-group, we can use the Orbit-stabilizer Theorem (which implies that the size of the orbit divides the size of the group) to show that $(-1)^{\#E(F)} \equiv_p (-1)^{\ell(F)}$ for all $F \in \mathrm{FP}(\Gamma, H)$.

Our second observation is that fixed points of fixed points lie within each other, that is, for all $A \in \mathrm{FP}(\Gamma, H)$, we have

$$\mathrm{FP}(\Gamma, A) = \{B \in \mathrm{FP}(\Gamma, H) : E(B) \subseteq E(A)\}.$$

We use $B \subseteq A$ to denote that B is a fixed point that lies in A, and say that A is a sub-point B.

Combining our two observations with Lemma A.1, we obtain

$$\widehat{\Phi}(A) \equiv_p \sum_{\substack{B \in \mathrm{FP}(\Gamma, H) \\ B \subseteq A}} \Phi(B)(-1)^{\ell(A)}. \tag{1}$$

Observe that we are summing only over fixed points $B \in \mathrm{FP}(\Gamma, H)$ that lie in A.

Hence, for computing the alternating enumerator of all $A \in \mathrm{FP}(\Gamma, H)$ it suffices to analyze the fixed point structure $\mathrm{FP}(\Gamma, H)$ for a single H. Equation (1) helps significantly in understanding the alternating enumerator. For instance, Equation (1) plays a central role in proving the following key lemma.

Lemma 4.8. *Let H denote a graph, let $\Gamma \subseteq \mathrm{Aut}(H)$ denote a p-group, and let Φ denote a graph property. Further, let $A \in \mathrm{FP}(\Gamma, H)$ denote a fixed point without property Φ, such that all of the proper sub-points of A do have property Φ; that is, we have $\Phi(A) = 0$ and $\Phi(B) = 1$ for every $B \subsetneq A$. Then A is nonvanishing.*

We employ the following strategy. Instead of directly finding one specific nonvanishing graph H with k vertices and *large* treewidth, we use Equation (1) to find some nonvanishing fixed point $A \in \mathrm{FP}(\Gamma, H)$. The advantage of this approach is that we can use Equation (1) to analyze the alternating enumerator of many different fixed points simultaneously to find a nonvanishing fixed point. Naturally, we still have to ensure that our fixed point has large treewidth. To that end, we prove that for certain graphs H and groups Γ, the level of A is lower bound for the treewidth.

The Prime Power Case: Difference Graphs

For a prime p and a positive integer m, we write $\mathbb{F}_{p^m}$ for the finite field with p^m elements. For $m = 1$, we write $\mathbb{F}_p = [0 \,..\, p)$. The elements that are invertible are denoted with $\mathbb{F}^*_{p^m} = \mathbb{F}_{p^m} \setminus \{0\}$.

For a prime p, we set $\mathbb{F}^+_p := \{1, \dots, \lceil (p-1)/2 \rceil\}$, which is a subset of $\mathbb{F}^*_p$ that contains exactly one of x and $-x$ for every $x \in \mathbb{F}^*_p$. We generalize this notion to finite fields with a prime power number of elements and write $\mathbb{F}^+_{p^m}$ for a set of elements that we obtain by including into it exactly one of x and $-x$ for every $x \in \mathbb{F}^*_{p^m}$ (observe that if $p = 2$, then $x = -x$ and hence $\mathbb{F}^+_{p^m} = \mathbb{F}^*_{p^m}$). We use $\mathbb{F}^+_{p^m}$ only in situations where the specific choice of elements does not matter. It is instructive to make explicit the following easy observation.

Lemma 2.3. *Let p denote a prime and let $m > 0$ denote an integer. Then, $|\mathbb{F}^*_{p^m}| = p^m - 1$ and $|\mathbb{F}^+_{p^m}| \geq (p^m - 1)/2$.*

Proof. For the bound $|\mathbb{F}^+_{p^m}| \geq (p^m - 1)/2$, we observe that in the special case $p = 2$, we have $x = -x$ and thus $|\mathbb{F}^+_{2^m}| = |\mathbb{F}^*_{2^m}| = (2^m - 1)$.

Suppose that Φ is nontrivial on $k = p^m$. Let us consider the graph K_{p^m} whose vertex set is the finite field $\mathbb{F}_{p^m}$ with p^m elements. The *rotation subgroup* $\circlearrowleft_{p^m} \subseteq \mathrm{Aut}(K_{p^m})$ contains those permutations of K_{p^m} that are described by addition in F_{p^m}, that is,

$$\circlearrowleft_{p^m} := \{\varphi_c \in \mathrm{Aut}(K_{p^m}) : c \in \mathbb{F}_{p^m} \text{ and } \varphi_c(v) = v + c \text{ for all } v \in K_{p^m}\}.$$

We observe that fixed points of the group $\circlearrowleft_{p^m}$ acting on K_{p^m} are exactly the difference graphs, as defined below.

Definition 2.4. *For a prime p, an integer $m > 0$, and a set $A \subseteq \mathbb{F}^+_{p^m}$, we define the* difference graph *$C^A_{p^m}$ via*

$$V(C^A_{p^m}) := \mathbb{F}_{p^m} \quad \textit{and} \quad E(C^A_{p^m}) := \{\{u, v\} : u, v \in \mathbb{F}_{p^m}, (u - v) \in A \cup (-A)\},$$

where $-A = \{-x : x \in A\}$. Observe that $C^A_{p^m} = K_{p^m}$ whenever $A = \mathbb{F}^+_{p^m}$.

The level *of a difference graph $C^A_{p^m}$ is the cardinality of A; we write $\ell(C^A_{p^m})$ for the level of $C^A_{p^m}$.*

As $C^A_{p^m}$ is $2|A|$-regular if $p \neq 2$ and $|A|$-regular if $p = 2$, it has treewidth at least $|A|$. This means that it is sufficient to find a nonvanishing difference graph (fixed point) with a high level. To that end, we consider two different cases.

First, let us consider the case when $\Phi(C^A_{p^m}) = 0$ for all fixed points $C^A_{p^m}$ with $\ell(A) \geq p^m - \sqrt{p^m}$. We introduce two vectors: the vector showing the total value of Φ on the fixed points on each level, and the analogous vector for $\widehat{\Phi}$. The two vectors are related by an invertible linear transform. A careful inspection of the matrix of the transformation shows that if Φ is 0 on fixed points of a level of at least $p^m - \sqrt{p^m}$, then fixed points of a level of at least $\sqrt{p^m}$ cannot all have zero alternating enumerator.

Lemma 4.13. *Let H denote a graph, let $\Gamma \subseteq \mathrm{Aut}(H)$ denote a p-group, and let Φ denote a graph property. Let $c < \ell(H)$ denote a nonnegative integer. Suppose that we have $\Phi(\varnothing) = 1$ and $\Phi(F) = 0$ for every $F \in \mathrm{FP}(\Gamma, H)$ with level $\ell(F) > \ell(H) - c$. Then, there is a fixed point $S \in \mathrm{FP}(\Gamma, H)$ with $\widehat{\Phi}(S) \not\equiv_p 0$ and $\ell(S) \geq c$.*

Second, let us consider the case when $\Phi(C^A_{p^m}) \neq 0$ for some fixed point $C^A_{p^m}$ with $\ell(A) \geq p^m - \sqrt{p^m}$. Since $C^A_{p^m}$ is true in Φ, all edge-subgraphs of $C^A_{p^m}$ are also true (here is the point where we use that Φ is edge-monotonicity). Further, each fixed point $C^B_{p^m}$ that is isomorphic to an edge-subgraph of $C^A_{p^m}$ (that is, $C^B_{p^m}$ is isomorphic to a graph $C^{A^*}_{p^m}$ with $A^* \subseteq A$) satisfies Φ as well. Since $C^A_{p^m}$ has a high level, this "starts an avalanche" and at some point, all fixed points below a certain level satisfy Φ. We show that this happens at level roughly $\sqrt{p^m}$. Now let us look at the fixed point $C^A_{p^m}$ with $\Phi(C^A_{p^m}) = 0$ such that all proper edge-subgraphs satisfy Φ. We have that $\ell(C^A_{p^m})$ is at least roughly $\sqrt{p^m}$ and Lemma 4.8 yields that $C^A_{p^m}$ is nonvanishing.

Lemma 5.14. *Let p denote a prime, let $m > 0$ denote an integer, and let Φ denote an edge-monotone graph property that is nontrivial on p^m. Further, write c and d for positive integers with $cd \leq |\mathbb{F}^+_{p^m}|$. Suppose that there is a fixed point $C^A_{p^m} \in \mathrm{FP}(\circlearrowleft_{p^m}, K_{p^m})$ with $\ell(C^A_{p^m}) \geq |\mathbb{F}^+_{p^m}| - d$ and $\Phi(C^A_{p^m}) = 1$.*

Then, there is a fixed point $C^B_{p^m} \subseteq C^A_{p^m}$ with $\ell(C^B_{p^m}) \geq c$ and $\widehat{\Phi}(C^B_{p^m}) \not\equiv_p 0$.

Both cases together allow us to find a nonvanishing fixed point $C^A_{p^m}$ with a level of roughly $\sqrt{p^m}$, thus $C^A_{p^m}$ has treewidth roughly $\sqrt{p^m}$. Thus, if Φ is nontrivial on infinitely many prime powers, then we can use this insight to find a sequence of nonvanishing graphs with unbounded treewidth.

Theorem 5.16. *Let Φ denote an edge-monotone graph property.*

- *If Φ is nontrivial on infinitely many prime powers, then #IndSub(Φ) is #W[1]-hard.*
- *Further, assuming ETH, there is a universal constant $\alpha > 0$ (independent of Φ) such that for any prime power $k \geq 3$ on which Φ is nontrivial, no algorithm (that reads the whole input) computes for every graph G the number* #IndSub$((\Phi, k) \to G)$ *in time $O(|V(G)|^{\alpha\sqrt{k}/\log k})$.*

General Case: Reduction to Prime Powers

After showing #W[1]-hardness for edge-monotone graph properties Φ that are nontrivial on infinitely many prime powers, our next goal is to find a reduction from the general case (that is Φ is nontrivial on infinitely many numbers) to this prime power case. If Φ is nontrivial on a k, then we analyze the largest prime power factor $q(k)$ of k. Set $d := k/q(k)$ and $p^m := q(k)$. We *join* d copies of the complete graph $K_{q(k)}$ into a graph $K^d_{q(k)} := K_{q(k)} \triangledown \cdots \triangledown K_{q(k)}$ which is isomorphic to $K_{d \cdot q(k)}$. Further, we take the group-theoretical product of d copies of $\circlearrowleft_{q(k)}$ to obtain the p-group $\circlearrowleft^d_{q(k)}$ that acts on the vertices of $K^d_{q(k)}$.

The advantage of this construction is that we can understand the fixed points $\mathrm{FP}(\circlearrowleft^d_{q(k)}, K^d_{q(k)})$ in terms of the fixed points $\mathrm{FP}(\circlearrowleft_{q(k)}, K_{q(k)})$. Specifically, each fixed point $F \in \mathrm{FP}(\circlearrowleft^d_{q(k)}, K^d_{q(k)})$ is made up of

- a graph C with $E(C) \subseteq [d]$ and d fixed points $C^{A_1}_{q(k)}, \dots, C^{A_d}_{q(k)} \in \mathrm{FP}(\circlearrowright_{q(k)}, K_{q(k)})$,
- where each pair of fixed points $C^{A_i}_{q(k)}$ and $C^{A_j}_{q(k)}$ is either fully connected or not connected at all, depending on whether the edge $\{i, j\}$ is present in $E(C)$.

Consult Figure 4 for visualizations of examples.

We use $C\langle C^{A_1}_{q(k)}, \dots, C^{A_d}_{q(k)} \rangle$ to denote said fixed points; we have

$$V(C\langle C^{A_1}_{q(k)}, \dots, C^{A_d}_{q(k)} \rangle) := [d] \times \mathbb{F}_{q(k)}$$
$$E(C\langle C^{A_1}_{q(k)}, \dots, C^{A_d}_{q(k)} \rangle) := \{\{(i, v_i), (j, u_j)\} : \{i, j\} \in E(C) \text{ or } (i = j \text{ and } v_i - u_i \in A_i)\}.$$

We show that $\mathrm{FP}(\circlearrowright^d_{q(k)}, K^d_{q(k)}) = \{C\langle C^{A^1}_{q(k)}, \dots, C^{A^d}_{q(k)} \rangle : C \text{ is a } d\text{-vertex graph}, A^i \subseteq \mathbb{F}^+_{q(k)}\}$ (see Corollary 6.8). A very important observation is that if C is not the empty graph, then the fixed point $C\langle \cdots \rangle$ contains $K_{q(k),q(k)}$ as a subgraph, and has thus a treewidth of at lest $q(k)$. Our discussion motivations the following notation.

Definition 6.17. *Let Φ denote an edge-monotone graph property and write M_Φ for the set of numbers on which Φ is nontrivial. We say that Φ is* concentrated *on an integer $k \in M_\Phi$ if there is a graph H on k vertices with $\hat{\Phi}(H) \neq 0$ and H contains $K_{q(k),q(k)}$ as a subgraph.*

We say that Φ is scattered *on an integer $k \in M_\Phi$ if it is not concentrated for Φ.*

If Φ is nontrivial, then we consider the following case distinction. First, we assume that Φ is (nontrivial and) concentrated on infinitely many values k. For each concentrated k, by definition, we have a nonvanishing graph with a treewidth of at least $q(k)$. We observe that $q(k) \geq c\log(k)$ for some constant $c > 0$ (see Lemma 6.16). This means that if Φ is (nontrivial and) concentrated on infinitely many k, then we can use these values to construct a nonvanishing sequence with unbounded treewidth. Now, Lemma A.8 shows that #IndSub(Φ) is #W[1]-hard.

Otherwise, Φ is (nontrivial and) scattered on infinitely many values k. For any such k, let us consider a fixed point $C\langle C^{A_1}_{q(k)}, \dots, C^{A_d}_{q(k)} \rangle$ of minimum level that is zero in Φ; as we have seen (Lemma 4.8), the alternating enumerator is nonzero for this graph. As k is scattered, we have that C is the empty graph, hence the fixed point is the disjoint union of $C^{A_1}_{p^m}, \dots, C^{A_d}_{p^m}$. Assume without loss of generality that $A_d \neq \emptyset$ and let H denote the disjoint union of $C^{A_1}_{p^m}, \dots, C^{A_{d-1}}_{p^m}$. Let us define a graph property whose value is $\Phi(G \uplus H)$ on G. Then, this property is nontrivial on $q(k)$-vertex graphs: it is zero on $C^{A_d}_{p^m}$ and nonzero on IS_{p^m}. Thus for each k on which Φ is scattered, we can construct a graph property that is nontrivial on the prime power $q(k)$.

Lemma 6.19. *Let Φ denote an edge-monotone graph property and write M_Φ for the set of numbers on which Φ is nontrivial. For any number $k \in M_\Phi$ on which Φ is scattered, there is a graph H on $k - q(k)$ vertices such that the property $(\Phi - H) := \{G : G \uplus H \in \Phi\}$ is edge-monotone and nontrivial on $q(k)$.*

Now, the idea is to combine the infinitely many graph properties from Lemma 6.19 into a single graph property Φ' that is nontrivial on infinitely many prime powers. The problem #IndSub(Φ') is now #W[1]-hard due to Theorem 5.16. Further, we show how to compute #IndSub($(\Phi', p^m) \to \star$) with an oracle for #IndSub($(\Phi, k) \to \star$) using the Inclusion-Exclusion principle (see Lemma 6.12). Thereby we obtain a parameterized Turing reduction from #IndSub(Φ') to #IndSub(Φ). Since the problem #IndSub(Φ') is #W[1]-hard, we thus obtain that #IndSub(Φ) is #W[1]-hard. Combining both cases leads to Main Theorem 1.

Main Theorem 1. *Let Φ denote a nontrivial edge-monotone graph property.*

- *The problem #IndSub(Φ) is #W[1]-hard.*

- *Further, assuming ETH, there is a universal constant $\gamma > 0$ (independent of Φ) such that for any integer $k \geq 3$ on which Φ is nontrivial, no algorithm (that reads the whole input) computes for every graph G the number* $\#\text{IndSub}((\Phi, k) \to G)$ *in time* $O(|V(G)|^{\gamma\sqrt{\log k}/\log\log k})$.

Tight Lower Bounds via Large Bicliques

Lastly, in Section 7, we show stronger lower bounds assuming ETH. We write n for the number of vertices of the input graph. So far, we proved our lower bounds by showing that we can solve $\textsc{Hom}(\{H\})$ using an oracle for $\#\text{IndSub}((\Phi, |V(H)|) \to \star)$ whenever $\widehat{\Phi}(H) \neq 0$. Further, we used that $\textsc{Hom}(\{H\})$ cannot be solved in time $O(n^{\alpha_{\textsc{Hom}} \cdot \text{tw}(H)/\log \text{tw}(H)})$ which yields an $O(n^{\alpha_{\textsc{IndSub}} \cdot \text{tw}(H)/\log \text{tw}(H)})$ lower bound for $\#\text{IndSub}((\Phi, |V(H)|) \to \star)$. However, we would like to prove that $\#\text{IndSub}((\Phi, k) \to \star)$ cannot be solved in time $O(n^{\gamma k})$ for a global constant $\gamma > 0$ that does not depend on k. This is not possible with our current method since we cannot get rid of the log factor in the denominator of the exponent. A lower bound of $O(n^{\gamma k})$ for $\#\text{IndSub}((\Phi, k) \to \star)$ would also prove that we cannot solve $\#\textsc{IndSub}(\Phi)$ in time $f(k)n^{o(k)}$ for any computable function unless ETH fails. This is tight in the sense that there is a brute-force algorithm that solves $\#\textsc{IndSub}(\Phi)$ in time $O(f(k)n^k)$, which is achieved by simply iterating through induced subgraphs of size k.

To achieve this goal, we use a reduction from [DRSW22, Theorem 1] that uses nonvanishing graphs that contain large bicliques $K_{k,k}$. Moreover, instead of starting the reduction from $\#\textsc{Hom}(\{H\})$, we start from the k-$\textsc{Clique}$ problem (which is the problem of deciding whether an input graph G contains a k-clique). It is known that $\textsc{Clique}$ has no algorithm with running time $f(k)n^{o(k)}$ for any computable function f, unless ETH fails [CFK+15]. However, we need a stronger form of this statement saying that there exists a constant α such that k-$\textsc{Clique}$ cannot be solved in time $O(n^{\alpha k})$ unless ETH fails; we prove this statement in Appendix B. The idea is to find a reduction from k-$\textsc{Clique}$ to $\#\text{IndSub}((\Phi, k) \to \star)$ such that an algorithm computing $\#\text{IndSub}((\Phi, k) \to \star)$ in time $O(n^{\gamma k})$ could be used to compute k-$\textsc{Clique}$ in time $O(n^{\alpha k})$ which is not possible unless ETH fails. We reprove the reduction of [DRSW22, Theorem 1] in Appendix C in the form that we need.

Theorem C.2 (Modification of [DRSW22]). *There is a global constant $\beta > 0$ and a positive integer N such that for all graph properties Φ, functions h, numbers k with*

- $h(k) \geq N$
- *there is a graph F with k vertices and $\widehat{\Phi}(F) \neq 0$,*
- *and F contains $K_{h(k),h(k)}$ as a subgraph*

there is no algorithm (that reads the whole input) that for every G computes $\#\text{IndSub}((\Phi, k) \to G)$ *in time* $O(|V(G)|^{\beta h(k)})$ *unless ETH fails.*

This means that we can prove stronger lower bounds for $\#\text{IndSub}((\Phi, k) \to \star)$ by finding nonvanishing graphs that contain large bicliques. As discussed below, for a prime p and $m \geq 1$, we can construct a p-group Syl_{p^m} such that we can find a fixed point of Syl_{p^m} in K_{p^m} that is nonvanishing and contains $K_{p^{m-1},p^{m-1}}$ as a subgraph. If p is small (constant) compared to p^m, then the size of this biclique is approximately the same as the number of vertices (however, this means that we cannot use these fixed points if Φ is nontrivial only on prime numbers). Therefore, Theorem C.2 allow us to prove tight lower bounds whenever Φ is nontrivial on a prime power $k = p^m$.

Main Theorem 2. *For each prime p, there is a constant $\gamma_p > 0$ such that for each integer m with $p^m \geq 3$ and each edge-monotone graph property Φ that is nontrivial on p^m, no algorithm (that reads the whole input) computes for every graph G the number* $\#\text{IndSub}((\Phi, k) \to G)$ *in time* $O(|V(G)|^{\gamma_p p^m})$*, unless ETH fails.*

Finding Nonvanishing Graphs with Large Bicliques

To find nonvanishing graphs with large bicliques, we use a new p-group Syl_{p^m}, defined as follows. We still use the complete graph K_{p^m}, but with vertex set $\{0, 1, \ldots, p-1\}^m$. For all m-tuples $\overline{\varphi} = (\varphi_0, \ldots, \varphi_{m-1})$ of functions with $\varphi_j \colon [0\,..\,p)^j \to [0\,..\,p)$, we define the following function on $V(K_{p^m})$.[2]

$$\overline{\varphi}(x_1, \ldots, x_m) = (x_1 + \varphi_0, x_2 + \varphi_1(x_1), x_3 + \varphi_2(x_1, x_2), \ldots, x_m + \varphi_{m-1}(x_1, \ldots, x_{m-1}))$$

where all computations are done modulo p. It is easy to see that each $\overline{\varphi}$ is a bijection on the vertex set and therefore in $\mathrm{Aut}(K_{p^m}) \cong \mathfrak{S}_{p^m}$, where $\mathfrak{S}_{p^m}$ is the symmetric group on p^m elements. We denote by Syl_{p^m} the set of all m-tuples of functions $\overline{\varphi} = (\varphi_0, \ldots, \varphi_{m-1})$ with $\varphi_j \colon [0\,..\,p)^j \to [0\,..\,p)$.

It is easy to check that $|\mathrm{Syl}_{p^m}|$ is a p-power. To describe the fixed points of Syl_{p^m} in K_{p^m} we need a concept that is known as the *lexicographic product* of the graphs. We use the following standard definition [Har94, Page 22].

Definition 7.8. *For graphs $G_1, \ldots, G_m$, we define their* lexicographic product $G_1 \circ \cdots \circ G_m$ *via*

$$\begin{aligned} V(G_1 \circ \cdots \circ G_m) &\coloneqq V(G_1) \times \cdots \times V(G_m) \quad \text{and} \\ E(G_1 \circ \cdots \circ G_m) &\coloneqq \{\{(u_1, \ldots, u_m), (v_1, \ldots, v_m)\} \\ &\qquad : \text{there is an } i \in [m] \text{ with } u_j = v_j \text{ for all } j < i \text{ and } \{u_i, v_i\} \in E(G_i)\}. \end{aligned}$$

Lemma 7.11. *For any prime p and any positive integer m, we have*

$$\mathrm{FP}(\mathrm{Syl}_{p^m}, K_{p^m}) = \{C_p^{A_1} \circ \cdots \circ C_p^{A_m} : A_i \subseteq \mathbb{F}_p^+\}.$$

One important observation is that a fixed point $C_p^{A_1} \circ \ldots \circ C_p^{A_m}$ has a large biclique if there is a small number i with $A_i \neq \varnothing$. We capture this observation by introducing the *empty-prefix* of a graph. The empty-prefix of $H = C_p^{A_1} \circ \ldots \circ C_p^{A_m}$ is the smallest index i with $A_i \neq \varnothing$, minus one; that is, $\varepsilon(A_1, \ldots, A_m) \coloneqq i - 1$, where i is the smallest index with $A_i \neq \varnothing$. We observe that a graph with a low empty-prefix contains a large biclique as a subgraph.

Lemma 7.15. *Let p denote a prime number and let m denote a positive integer. For each $i \in [m]$, let $A_i \subseteq \mathbb{F}_p^+$ denote a subset and set $A \coloneqq (A_1, \ldots, A_m)$. Then, $C_p^{A_1} \circ \cdots \circ C_p^{A_m}$ contains $K_{p^{m-1-\varepsilon(A)}, p^{m-1-\varepsilon(A)}}$ as a subgraph.*

Suppose that our graph property is nontrivial on p^m for $m \geq 2$. If we find a fixed point $H = C_p^{A_1} \circ \ldots \circ C_p^{A_m}$ with a minimal empty-prefix of $\varepsilon(A_1, \ldots, A_m) = 0$, then we know that the treewidth of H is at least p^{m-1}. Thus, our goal is to find a fixed point H with $\varepsilon(A_1, \ldots, A_m) = 0$ and $\widehat{\Phi}(H) \not\equiv_p 0$.

To that end, we prove that a fixed point with a high empty-prefix is always isomorphic to an edge-subgraph of a fixed point with a low empty-prefix, which allows us to always consider fixed points with the low empty-prefix. To be more precise, we show that each fixed point $C_p^{\varnothing} \circ \ldots C_p^{\varnothing} \circ C_p^{A_1} \circ \ldots \circ C_p^{A_{m-j}}$ is isomorphic to an edge-subgraph of $C_p^{A_1} \circ \ldots \circ C_p^{A_m}$. Intuitively, we achieve thus by defining an isomorphism that pushes the edges of each $C_p^{A_i}$ *one level down*.

Lemma 7.17. *Let p denote a prime number and let m denote a positive integer. For each $i \in [m]$, let $A_i \subseteq \mathbb{F}_p^+$ denote a subset. Then, for all $j \in [m]$, the graph $C_p^{\varnothing} \circ \cdots C_p^{\varnothing} \circ C_p^{A_1} \circ \cdots \circ C_p^{A_{m-j}}$ is isomorphic to an edge-subgraph of $C_p^{A_1} \circ \cdots \circ C_p^{A_m}$.*

---

[2] We write φ_0 for $\varphi_0(\varnothing)$ since φ_0 is a function that is defined on a single element

Now, we are ready to show that there is always a nonvanishing fixed point of Syl_{p^m} in K_{p^m} that contains $K_{p^{m-1},p^{m-1}}$ as a subgraph. We show this by finding a fixed point H with $\Phi(H) = 0$, an empty-prefix of 0, and $\Phi(\tilde{H}) = 1$ for all $\tilde{H} \subseteq H$ (that is, fixed points $\tilde{H}$ that lie in H). To do this, we consider the first level i such that there is a fixed point of level i that does not satisfy Φ. Observe that the only fixed point of level 0 is the independent set and $\Phi(\mathrm{IS}_{p^m}) = 1$ thus $i \geq 1$. Further, if we assume that all fixed points $C_p^{A_1} \circ \ldots C_p^{A_m}$ of level i with $\varepsilon(A_1, \ldots, A_m) = 0$ satisfy Φ, then we use Lemma 7.17 to show that each fixed point of level i satisfy Φ, a contradiction. Thus, there is a fixed point H with $\Phi(H) = 0$, an empty-prefix of 0, and $\Phi(\tilde{H}) = 1$ for all $\tilde{H} \subseteq H$. Lastly, we use Lemmas 4.8 and 7.15 and to show that $\widehat{\Phi}(H) \neq 0$ and that H contains $K_{p^{m-1},p^{m-1}}$ as a subgraph.

Theorem 7.18. *Let Φ denote an edge-monotone graph property that is nontrivial on a prime power p^m, then there is a nonvanishing fixed point of* Syl_{p^m} *in K_{p^m} that contains $K_{p^{m-1},p^{m-1}}$ as a subgraph.*

3 Additional Preliminaries

Numbers and Sets

For a natural number n, we write $[n]$ for the set $\{1, \ldots, n\}$. For natural numbers a and b, we write $[a\,..\,b)$ for the set $\{a, \ldots, b-1\}$.

Next, we write $\binom{n}{k}$ for the binomial coefficient and we set $\binom{n}{k} := 0$ whenever $k < 0$ or $n < k$. If A is a set then $\binom{A}{k}$ is the set of all subsets $B \subseteq A$ with size k.

For a subset $B \subseteq R$ of a field R, we set $-B := \{-b : b \in B\}$; for a set $B \subseteq R$ and an element $\lambda \in R$ we set $\lambda B := \{\lambda b : b \in B\}$.

We use $a \equiv_m b$ as a shorthand for $a \equiv b \mod m$.

Graphs

We consider only simple graphs, that is undirected graphs that have neither weights, loops, nor parallel edges. We write $\mathcal{G}$ for the set of all (simple) graphs and $\mathcal{G}_n$ for the set of all (simple) graphs with the vertex set $[n]$.

For a graph G, we write $V(G)$ for its vertices and we write $E(G)$ for its edges. Given a subset of edges $A \subseteq E(G)$, we write $G\{A\}$ for the graph with vertex set $V(G\{A\}) := V(G)$ and edge set $E(G\{A\}) := A$. We say that $G\{A\}$ is an *edge-subgraph* of G and we write $\mathcal{E}(G)$ for the set of all edge-subgraphs of G. For a set $X \subseteq V(G)$, we write $G \setminus X$ for the graph that is obtained from G by removing the vertices in X and all edges with at least one endpoint in X.

For two graphs G and H with a common vertex set V, we write $G \cup H$ for the graph on V with edges $E(G) \cup E(H)$.

We write IS_k for the independent set with k vertices. Further, we write K_n for the complete graph on n vertices and we write $K_{n,m}$ for the complete bipartite graph on $n + m$ vertices.

We write $\mathrm{tw}(G)$ for the *treewidth* of a graph G. As we use the treewidth of a graph in a black-box manner, we refer an interested reader to [CFK+15, Chapter 7.2], for a formal definition. Intuitively, the treewidth of a graph measures how far away a given graph is from being a tree. For example, the treewidth of a tree is 1; the treewidth of a complete graph on n vertices is $n - 1$.

Group Theory and Morphisms between Graphs

For a finite set X of size n, the set of all bijections $X \to X$ and the function composition together form a group which we call the symmetric group $\mathfrak{S}_n$. We also write $\mathfrak{S}_X$ if we wish to emphasize the set X.

A group Γ is a *permutation group* if Γ is a subgroup of the symmetric group $\mathfrak{S}_X$ for some finite set X. A permutation group $\Gamma \subseteq \mathfrak{S}_X$ is *transitive* if every $x \in X$ can be mapped to any other element $y \in X$ via some $g \in \Gamma$. Finally, a *p-group* is a finite group Γ that has an order that is a power of p (that is, the number of elements of Γ is a power of p).

A graph *homomorphism* $h : V(H) \to V(G)$ from H to G is a function between the vertex sets of two graphs that preserves adjacencies (but not necessarily non-adjacencies), that is, h maps the vertices of every edge $\{u, v\} \in E(H)$ to vertices $\{h(u), h(v)\} \in E(G)$. We write $\mathrm{Hom}(H \to G)$ for the set of all homomorphisms from H to G.

We say a homomorphism $h : V(H) \to V(G)$ is an *isomorphism* if h is a bijection on the vertex sets and $h^{-1} : V(H) \to V(G)$ also defines a homomorphism (that is, $\{u, v\} \in E(H)$ if and only if $\{h(u), h(v)\} \in E(G)$). We say that two graphs G and H are *isomorphic*, denoted by $G \cong H$ if there is an isomorphism between them. An *automorphism* of G is an isomorphism from G to G. We write $\mathrm{Aut}(G)$ for the set of all *automorphism* of G.

For each graph G, the set of automorphisms $\mathrm{Aut}(G)$ forms a group with composition as the group operation. Observe that the automorphism group of a clique $\mathrm{Aut}(K_n)$ is just the symmetric group $\mathfrak{S}_n$.

For primes p and prime powers p^m, the group $\mathrm{Aut}(K_{p^m})$ contains a useful subgroup with order p^m: in particular, $\mathrm{Aut}(K_{p^m})$ contains automorphisms φ_c of G that "rotate" the vertices of G, that is, automorphisms that send every vertex i to vertex $i + c$ for some $c \in \mathbb{F}_{p^m}$ (where we identify the vertices of K_{p^m} with the elements of $\mathbb{F}_{p^m}$); we say that φ_c is a *rotation (by c)* of G. In particular, we write $\circlearrowleft_{p^m}$ for the subgroup of all rotations of $\mathrm{Aut}(K_{p^m})$, that is

$$\circlearrowleft_{p^m} := \{\varphi_c \in \mathrm{Aut}(K_{p^m}) : c \in \mathbb{F}_{p^m} \text{ and } \varphi_c(v) = v + c \text{ for all } v \in K_{p^m}\}.$$

Observe that $\circlearrowleft_{p^m}$ is a group with p^m elements.

Let G denote a graph and consider a subgroup $\Gamma \subseteq \mathrm{Aut}(G)$. Any $g \in G$ is a bijection on $V(G)$, which we may interpret as g permuting the vertices of G. Thus, we may also interpret g as an operation on the edges of G. However, as we wish to use results from the literature (such as the Orbit-stabilizer Theorem), we need to describe this operation on the edges of G using the language of group actions.

Formally, we can turn the operation on vertices into a *group action* $\cdot : \Gamma \times E(G) \to E(G)$ that tells us how each member of Γ moves the edges of G. Specifically, we define $g \cdot \{u, v\} := \{g(u), g(v)\}$. We interpret said group action $\cdot$ multiplicatively and typically write just $g\{u, v\}$.

Extending the previous group action, Γ also acts on edge-subgraphs of G via

$$\begin{aligned} V(g \cdot A) &:= V(A), \text{ and} \\ E(g \cdot A) &:= \{\{g(u), g(v)\} : \{u, v\} \in E(A)\}, \end{aligned}$$

for each $g \in \Gamma$ and each $A \in \mathcal{E}(G)$. Again, we interpret said group actions $\cdot$ multiplicatively and typically write just gA.

For each $\{u, v\} \in E(G)$ the set $\Gamma \cdot \{u, v\} := \{g \cdot \{u, v\} : g \in \Gamma\}$ is the *orbit* of $\{u, v\}$. Two different edges either have disjoint orbits or equal orbits. We write $E(G)/\Gamma$ to denote the set of all orbits of $\cdot$. Recall that $E(G)/\Gamma$ forms a partition of $E(G)$. In a slight abuse of notation, for an orbit O, we also write O for the edge-subgraph $G\{O\}$; similarly, we use $E(H)/\Gamma$ to denote the set of all such edge-subgraphs.

We say that an edge-subgraph $A \in \mathcal{E}(G)$ is a *fixed point* of Γ in G if $gA = A$ for all $g \in \Gamma$. We write $\mathrm{FP}(\Gamma, G)$ for the set of all fixed points of Γ in G.

As it turns out, a graph H being a fixed point of Γ in G is a strong property that is very useful to us. For now and as a first useful observation, we see that H inherits key properties from G.

Lemma 3.1. *Let G denote a graph and let $\Gamma \subseteq \mathrm{Aut}(G)$ denote a group. For any $H \in \mathrm{FP}(\Gamma, G)$ all of the following hold.*
(1) *We have $\Gamma \subseteq \mathrm{Aut}(H)$.*
(2) *Any fixed point of Γ in H is also a fixed point of Γ in G and we have $\mathrm{FP}(\Gamma, H) = \{A \in \mathrm{FP}(\Gamma, G) : E(A) \subseteq E(H)\}$.*

Proof. For (1), we first recall that as H is a fixed point of Γ in G, we have $gH = H$ for every $g \in \Gamma$. In particular, this means that every G-automorphism $g \in \Gamma$ is also an automorphism of H. This in turn yields the claim.

For (2), we first see that by (1), $\mathrm{FP}(\Gamma, H)$ is indeed well-defined. Now, consider a fixed point $A \in \mathrm{FP}(\Gamma, H)$. By definition, we have $E(A) \subseteq E(H) \subseteq E(G)$ and $gA = A$ for every $g \in \Gamma$, so A is indeed also a fixed point of Γ in G.

For the other direction, observe that a fixed point $A \in \mathrm{FP}(\Gamma, G)$ with $E(A) \subseteq E(H)$ is just a fixed point of Γ in H, which completes the proof. ∎

Next, we observe that the notion of fixed points is compatible with set operations (on the edges of the underlying graph).

Lemma 3.2. *Let G denote a graph and let $\Gamma \subseteq \mathrm{Aut}(G)$ denote a group. For any $H_1, H_2 \in \mathrm{FP}(\Gamma, G)$ all of the following hold.*
(1) *We have $H_1 \cup H_2 \in \mathrm{FP}(\Gamma, G)$.*[3]
(2) *If $H_2 \subseteq H_1$, then we have $H_1 \setminus H_2 \in \mathrm{FP}(\Gamma, G)$.*

Proof. For (1), we wish to show that for any $g \in \Gamma$, we have $g(H_1 \cup H_2) = H_1 \cup H_2$. To that end, observe that we have $gH_1 = H_1$ and $gH_2 = H_2$; that is, g is an automorphism when restricted to either H_1 or H_2. As unions of automorphisms are automorphisms, we obtain the claim.

For (2), as in the proof of (1), observe that any automorphism $g \in \Gamma$ (which is also an automorphism on H_1 as H_1 is a fixed point) stays an automorphism when restricted to H_2. In particular, this means that any such automorphism decomposes into an automorphism on H_2 and an automorphism on $H_1 \setminus H_2$. ∎

Counting Problems and Parameterized Complexity A *parameterized (counting) problem* consists of a function $P \colon \Sigma^* \to \mathbb{N}$ and a computable parameterization $\kappa \colon \Sigma^* \to \mathbb{N}$. A parameterized problem (P, κ) is called *fixed-parameter tractable* (FPT) if there is a computable function f and a deterministic algorithm $\mathbb{A}$ such that $\mathbb{A}$ computes $P(x)$ in time $f(\kappa(x))|x|^{O(1)}$ for all $x \in \Sigma^*$.

A parameterized Turing reduction from (P, κ) to (P', κ') is a deterministic FPT algorithm with oracle access to P' that computes $P(x)$ such that there is a computable function g with the property that $\kappa'(y) \le g(\kappa(y))$ holds for each oracle access to P'. We write $A \le_{\mathrm{T}}^{\mathrm{fpt}} B$ to denote that there is a parameterized Turing reduction from A to B.

In the counting version $\#P$ of a decision problem P, the task is to compute the number of valid solutions for a given input x. Of special importance is the counting problem #Clique, which gets as input a graph G and a natural number k. The output is the number of induced subgraphs of G of size k which form a k-clique. We parameterize #Clique by $\kappa(G, k) := k$. We are interested in barriers to obtain fast algorithms for #Clique and for parameterized counting problems in general.

---

[3] Slightly abusing notation, we use set operation for fixed points to mean the same operation on the edge sets of the corresponding edge-subgraphs.

A problem (P', κ') is #W[1]-hard if there is parameterized Turing reduction from (#Clique, κ) to (P', κ'). It is widely believed that #Clique has no FPT algorithm [FG04, CCF+05], ; hence #W[1]-hardness rules out FPT algorithms based on said belief.

For more fine-grained lower bounds, we also rely on the *Exponential Time Hypothesis* (ETH).

Exponential Time Hypothesis (ETH) ([CFK+15, Conjecture 14.1] [IPZ01]). *There is a positive real value $\varepsilon > 0$ such that the problem 3-SAT cannot be solved in time $O^*(2^{\varepsilon n})$, where n is the number of variables used in the formula.*

The Exponential Time Hypothesis is a stronger hypothesis and in fact implies there are no FPT algorithms for #W[1]-hard problems. For example, it can be used to show that no algorithm solves Clique in time $f(k)n^{o(k)}$, where n is the number of vertices. Observe, that this implies that there is no FPT algorithm for the #W[1]-hard Clique (see Lemma B.2).

4 Fixed Points and the Alternating Enumerator

In this section, we introduce the framework of how we can use fixed points of groups that act on graphs to show that the alternating enumerator is nonvanishing on certain graphs. In particular, the key technical result of this section is Lemma 4.13, which states a duality between the highest nonzero levels in Φ and $\widehat{\Phi}$. That is, it is not possible that both levels are low: their sum must be at least n.

4.1 Fixed Points as Unions of Orbits and Sub-points of Fixed Points

As a first step, we discuss how to construct any fixed point from a small set of basic building blocks.

Lemma 4.1. *Let H denote a graph and let $\Gamma \subseteq \mathrm{Aut}(H)$ denote a group. Further, let Γ act on $E(H)$ and write $E(H)/\Gamma$ for the set of all resulting orbits. Finally, let Γ act on edge-subgraphs of H and write $\mathrm{FP}(\Gamma, H)$ for the set of all resulting fixed points of Γ in H.*

Then, the edge set of each fixed point $F \in \mathrm{FP}(\Gamma, H)$ is the (possible empty) disjoint union of orbits $O_1, \dots, O_s$, where $O_i \in E(H)/\Gamma$ and each disjoint union of orbits yields a fixed point. The partition into orbits is unique.

Proof. We readily confirm that the empty set is indeed a fixed point; we may obtain the empty set as the empty union of orbits.

Next, we turn to single orbits and verify that, indeed, they are fixed points as well.

Claim 4.2. *We have $E(H)/\Gamma \subseteq \mathrm{FP}(\Gamma, H)$.*[4]

Proof. Consider an orbit $O \in E(H)/\Gamma$ and an automorphism $g \in \Gamma$. We intend to show that $gO = O$, which suffices to prove the claim.

To that end, by definition of an orbit, we have $ge \in O$ for every edge $e \in O$. As g is an automorphism of G, no two different edges from O are mapped to the same edge; this yields the claim.

Observe that Claim 4.2 and Lemma 3.2 together yield that the union of orbits is indeed a fixed point.

Finally, we show that we can always split off some orbit from a (non-empty) fixed point.

Claim 4.3. *(The edge set of) any fixed point $F \in \mathrm{FP}(\Gamma, H)$ can be obtained as the (disjoint) union of and orbit $O \in E(H)/\Gamma$ and a fixed point from $\mathrm{FP}(\Gamma, H)$.*

[4] Recall that we use $O \in E(H)/\Gamma$ to denote both the edge set as well as the edge-subgraph.

Proof. Consider an arbitrary edge $e \in E(F)$ and let O denote the corresponding orbit in $E(H)/\Gamma$. By definition of O and F, we have $O \subseteq E(F)$. Combining Claim 4.2 and Lemma 3.2 yields the claim. ◿

Iterating Claim 4.3 and recalling that the orbits $E(H)/\Gamma$ partition $E(H)$ completes the proof. ◾

Remark 4.4. Observe that Lemma 4.1 allows us to strengthen Lemma 3.2 to all set operations: the orbits of different edges are either equal or disjoint, which for any $H_1, H_2 \in \mathrm{FP}(\Gamma, G)$ directly yields all of the following.

(1) We have $H_1 \cup H_2 \in \mathrm{FP}(\Gamma, G)$.
(2) We have $H_1 \cap H_2 \in \mathrm{FP}(\Gamma, G)$.
(3) We have $H_1 \setminus H_2 \in \mathrm{FP}(\Gamma, G)$. ◾

Lemma 4.1 and in particular Claim 4.3 induce an order of the fixed points in $\mathrm{FP}(\Gamma, H)$: we say a fixed point F is a *sub-point* of another fixed point G if we can obtain G from the union of F and (potentially multiple) orbits in $E(H)/\Gamma$.

Definition 4.5. *Let H denote a graph and let $\Gamma \subseteq \mathrm{Aut}(H)$ denote a group. Further, let Γ act on $E(H)$ and write $E(H)/\Gamma$ for the set of all resulting orbits. Finally, let Γ act on edge-subgraphs of H and write $\mathrm{FP}(\Gamma, H)$ for the set of all resulting fixed points of Γ in H.*

For a fixed point $F \in \mathrm{FP}(\Gamma, H)$, its orbit factorization *$\mathbb{O}(F)$ is the unique subset of $E(H)/\Gamma$ whose union is F:*

$$F = \bigcup \mathbb{O}(F).$$

The level *of F, denoted by $\ell(F)$, is the size of the orbit factorization of F*

$$\ell(F) := |\mathbb{O}(F)|.$$

Finally, for two fixed points $F_1, F_2 \in \mathrm{FP}(\Gamma, H)$, we say that F_2 is a sub-point *of F_1, denoted by $F_2 \subseteq F_1$, if the orbit factorization of F_2 is a subset of the orbit factorization of F_1. If the inclusion is strict, we say that F_2 is a* proper sub-point *of F_1.* ◾

Observe that for two fixed points F_1 and F_2 of some group Γ with $F_2 \subseteq F_1$, we may equivalently write $F_2 \in \mathrm{FP}(\Gamma, F_1)$.

The next lemma makes it possible to group the sub-points of a fixed point according to their level.

Lemma 4.6. *Let H denote a graph and let $\Gamma \subseteq \mathrm{Aut}(H)$ denote a p-group.*
For any fixed point $F \in \mathrm{FP}(\Gamma, H)$, we have

$$(-1)^{\#E(F)} \equiv_p (-1)^{\ell(F)}.$$

Proof. Write $\mathbb{O}(F) = \{O_1, \ldots, O_{\ell(F)}\}$ for the orbit factorization of F. By the Orbit-stabilizer Theorem, the size of each O_i is a divider of the group order of Γ. Thus, $|O_i|$ is always odd if p is an odd prime number, which implies the claim.

For $p = 2$, we have $(-1) \equiv_2 1$, which immediately yields the claim. ◾

Using Lemma 4.6, we can obtain yet another expression for the alternating enumerator.

Corollary 4.7. *Let H denote a graph, let $\Gamma \subseteq \mathrm{Aut}(H)$ denote a p-group, and let Φ denote a graph property. We have*

$$\hat{\Phi}(H) \equiv_p \sum_{A \in \mathrm{FP}(\Gamma, H)} \Phi(A)(-1)^{\ell(A)}.$$

Further, for any fixed point $A \in \mathrm{FP}(\Gamma, H)$, we have

$$\hat{\Phi}(A) \equiv_p \sum_{B \subseteq A} \Phi(B)(-1)^{\ell(B)}.$$ ◾

4.2 Nonvanishing of the Alternating Enumerator via Sub-points

In the next step, we use sub-points to analyze the alternating enumerator. As a first result, we use Corollary 4.7 to obtain that the alternating enumerator is nonvanishing on fixed points that are "minimally vanishing" for the corresponding graph property.

Lemma 4.8. *Let H denote a graph, let $\Gamma \subseteq \mathrm{Aut}(H)$ denote a p-group, and let Φ denote a graph property. Further, let $A \in \mathrm{FP}(\Gamma, H)$ denote a fixed point without property Φ, such that all of the proper sub-points of A do have property Φ; that is, we have $\Phi(A) = 0$ and $\Phi(B) = 1$ for every $B \subsetneq A$. Then A is nonvanishing.*

Proof. By definition, the fixed point A has exactly $\binom{\ell(A)}{i}$ sub-points that have a level of exactly i. Starting from Corollary 4.7, we rewrite the alternating enumerator and obtain

$$\begin{aligned}
\widehat{\Phi}(A) &\equiv_p \sum_{B \subseteq A} \Phi(B)(-1)^{\ell(B)} \\
&\equiv_p -(-1)^{\ell(A)} + \sum_{B \subseteq A} (-1)^{\ell(B)} && (\Phi(A) = 0 \text{ and } \Phi(B) = 1 \text{ for all } B \subsetneq A) \\
&\equiv_p -(-1)^{\ell(A)} + \sum_{i=0}^{\ell(A)} \binom{\ell(A)}{i} (-1)^i && \text{(group by level)} \\
&\equiv_p -(-1)^{\ell(A)} + (-1+1)^{\ell(A)} && \text{(Binomial Theorem).}
\end{aligned}$$

In total, this yields the claim. ∎

For our next steps toward the proof Lemma 4.13, it is instructive to group sub-points by their level. In particular, we are interested in how many sub-points have a given graph property and we are interested in the sum of the alternating enumerators of said sub-points.

Definition 4.9. *Let H denote a graph, let $\Gamma \subseteq \mathrm{Aut}(H)$ denote a p-group, and let Φ denote a graph property.*

For every level $i \in [0\,..\,\ell(H)]$, we write w_i for the number (modulo p) of sub-points of H with level i that satisfy Φ; that is,

$$w_i := \sum_{\substack{A \in \mathrm{FP}(\Gamma, H) \\ \ell(A) = i}} \Phi(A) \bmod p.$$

We write $\mathbf{w} := (w_i)$ for the $(\ell(H)+1)$-dimensional vector that consists of the values w_i.

Further, for every level $i \in [0\,..\,\ell(H)]$, we write $\widehat{w}_i$ for the sum (modulo p) of the alternating enumerators of the sub-points of H with level i; that is,

$$\widehat{w}_i := \sum_{\substack{A \in \mathrm{FP}(\Gamma, H) \\ \ell(A) = i}} \widehat{\Phi}(A) \bmod p.$$

We write $\widehat{\mathbf{w}} := (\widehat{w}_i)$ for the $(\ell(H)+1)$-dimensional vector that consists of the values $\widehat{w}_i$. ∎

Let us take a step back and discuss Definition 4.9 for a bit. First, observe that $\widehat{w}_i \not\equiv_p 0$ implies $\widehat{\Phi}(A) \not\equiv_p 0$ for some sub-point A with level i. Hence, for our purposes, it suffices to understand when $\widehat{w}_i$ is nonzero.

Second, observe that understanding $\widehat{w}_i$ directly is thus not much easier compared to understanding a single alternating enumerator. However, this is where w_i—which is very easy to understand and compute—turns out to be useful. As we show next, we can express $\widehat{w}_i$-values as a linear combination of w_i-values and—more importantly—vice versa. In particular, this then allows us to show that some $\widehat{w}_i$-value has to be nonzero when enough w_j-values are nonzero.

Let us start by defining the transformation matrix that we use in the following.

Definition 4.10. *For a positive integer n, we define the $\widehat{\Phi}$-transformation matrix $C_n \in \mathbb{Z}^{(n+1)\times(n+1)}$ as*

$$(C_n)_{i,j} := (-1)^j \binom{n-j}{i-j}, \quad \text{where } i, j \in [0\,..\,n+1].$$

Further, for each $0 \le c \le n$, we define the c-restricted Φ-$\widehat{\Phi}$-transformation matrix $C_{n;c} \in \mathbb{Z}^{(n-c+1)\times(n-c+1)}$ as the square submatrix of C_n that consists in the first $n-c+1$ columns and the last $n-c+1$ rows of C_n; that is,

$$(C_{n;c})_{i,j} := (C_n)_{i+c,j} = (-1)^j \binom{n-j}{i+c-j},$$

where $i, j \in [0\,..\,n-c]$.

Recall that we set $\binom{a}{b} = 0$ whenever $b < 0$ or $b > a$. In particular, this means that we have $(C_n)_{i,j} = 0$ for all $j > i$, that is, C_n is a lower triangular matrix.

We proceed to show that Definition 4.10 indeed relates the $\widehat{\mathbf{w}}$-vector with the $\mathbf{w}$-vector.

Lemma 4.11. *Let H denote a graph, let $\Gamma \subseteq \mathrm{Aut}(H)$ denote a p-group, and let Φ denote a graph property. Define $n := \ell(H)$.*

Then, we have $\widehat{\mathbf{w}} \equiv_p C\ \mathbf{w}$; that is, for every level $i \in [0\,..\,n]$, we have

$$\widehat{w}_i \equiv_p \sum_{j=0}^{n} (-1)^j \binom{n-j}{i-j} w_j.$$

Proof. Fix a level $i \in [0\,..\,n]$. First, we use Corollary 4.7 to rewrite $\widehat{w}_i$. We obtain

$$\widehat{w}_i \equiv_p \sum_{\substack{A \in \mathrm{FP}(\Gamma,H)\\ \ell(A)=i}} \widehat{\Phi}(A) \equiv_p \sum_{\substack{A \in \mathrm{FP}(\Gamma,H)\\ \ell(A)=i}} \sum_{B \subseteq A} \Phi(B)(-1)^{\ell(B)}.$$

Next, we group the terms of $\sum_{B\subseteq A} \Phi(B)$ according to their level. We obtain

$$\widehat{w}_i \equiv_p \sum_{\substack{A \in \mathrm{FP}(\Gamma,H)\\ \ell(A)=i}} \sum_{j=0}^{i} (-1)^j \sum_{\substack{B \subseteq A\\ \ell(B)=j}} \Phi(B) \equiv_p \sum_{j=0}^{i} (-1)^j \sum_{\substack{A \in \mathrm{FP}(\Gamma,H)\\ \ell(A)=i}} \sum_{\substack{B \subseteq A\\ \ell(B)=j}} \Phi(B).$$

Next, we reorder the summation and group single fixed points B to obtain

$$\widehat{w}_i \equiv_p \sum_{j=0}^{i} (-1)^j \sum_{\substack{B \in \mathrm{FP}(\Gamma,H)\\ \ell(B)=j}} \sum_{\substack{A \in \mathrm{FP}(\Gamma,H)\\ A \supseteq B\\ \ell(A)=i}} \Phi(B).$$

Finally, we compute the cardinality of the set $\{A \in \mathrm{FP}(\Gamma, H) : A \supseteq B \text{ and } \ell(A) = i\}$ as the number of possible ways to choose $i - j$ elements from $n - j$ elements when disregarding the order of the chosen elements and allowing each element to be selected at most once. We obtain the claimed equality

$$\widehat{w}_i \equiv_p \sum_{j=0}^{i} (-1)^j \binom{n-j}{i-j} \sum_{\substack{B \in \mathrm{FP}(\Gamma,H)\\ \ell(B)=j}} \Phi(B) \equiv_p \sum_{j=0}^{n} (-1)^j \binom{n-j}{i-j} w_j,$$

where for the last step, we use the definition of w_j and we exploit that $\binom{a}{b} = 0$ whenever $b < 0$.

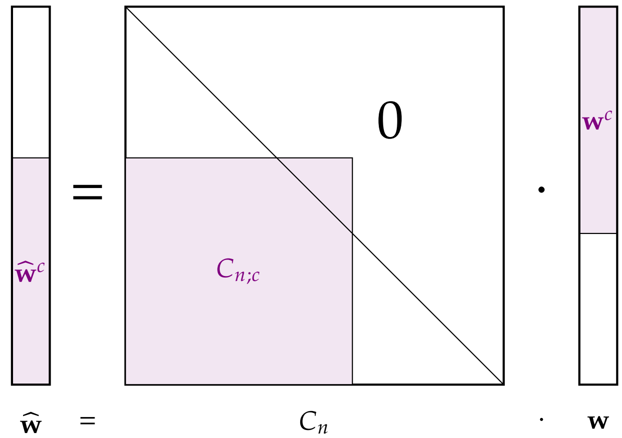

Figure 2 Lemma 4.11 visualized. Assume that the last c levels of the $\mathbf{w}$-vector are 0 (that is, $w_i = 0$ for $i > n - c$). Then the last $n - c + 1$ levels of the $\widehat{\mathbf{w}}$-vector can be obtained by multiplying the first $n - c + 1$ levels of the $\mathbf{w}$-vector by the lower left $(n - c + 1) \times (n - c + 1)$ submatrix $C_{n;c}$ of C. As $C_{n;c}$ is invertible, this transformation is a bijection (Lemma 4.12).

The key insight in the proof of Lemma 4.13 is the following. Assuming that the last c levels of $\mathbf{w}$-vector are 0 (that is, $w_i = 0$ for $i > n - c$), then the last $n - c + 1$ levels of the $\widehat{\mathbf{w}}$-vector are determined from $\mathbf{w}$ by the lower left $(n - c + 1) \times (n - c + 1)$ submatrix of C_n (consult Figure 2). We show next that this submatrix $C_{n;c}$ is invertible, hence we can revert the operation. This allows us to prove that the last $n - c + 1$ levels of $\widehat{\mathbf{w}}$ do not vanish.

Lemma 4.12. *For a prime p and integers $0 \le c \le n$, the matrix $C_{n;c}$ is invertible in $\mathbb{F}_p^{(n-c+1)\times(n-c+1)}$.*

Proof. We compute the determinant of $C_{n;c}$ and in particular show that $\det(C_{n;c}) \not\equiv_p 0$, which suffices to prove the claim.

To that end, first consider the matrix $D_{n;c}$ with

$$(D_{n;c})_{i,j} \coloneqq |(C_{n;c})_{i,j}| = \binom{n-j}{c+i-j}, \quad \text{for } i, j \in [0\,..\,n-c].$$

As we have $C_{n;c} = D_{n;c} \cdot \operatorname{diag}(1, -1, 1, -1, \dots, (-1)^{n-c+1})$, we also have

$$|\det(C_{n;c})| = |\det(D_{n;c}) \cdot \det(\operatorname{diag}(1, -1, \dots, (-1)^{n-c+1}))| = |\det(D_{n;c})|.$$

Next, write $E_{n;c}$ for the matrix that is obtained from $D_{n;c}$ by reversing its rows and columns, that is,

$$\begin{aligned}(E_{n;c})_{i,j} &\coloneqq (D_{n;c})_{(n-c-i),(n-c-j)} \\ &= \binom{n-(n-c-j)}{c+(n-c-i)-(n-c-j)} = \binom{c+j}{c+j-i} = \binom{c+j}{i}, \quad \text{for } i, j \in [0\,..\,n-c];\end{aligned}$$

where the last equality uses $\binom{a}{a-b} = \binom{a}{b}$. As permuting rows and columns of a matrix does not change (up to the sign) the determinant of a matrix, we have $|\det(D_{n;c})| = |\det(E_{n;c})|$.

Now, we use Vandermonde's Identity to obtain

$$(E_{n;c})_{i,j} = \binom{c+j}{i} = \sum_{a=0}^{i} \binom{c}{i-a}\binom{j}{a}. \tag{2}$$

Equation (2) in turn shows that we have $(E_{n;c}) = (L_{n;c})(U_{n;c})$ for the matrices

$$(L_{n;c})_{i,j} \coloneqq \binom{c}{i-j} \quad \text{and} \quad (U_{n;c})_{i,j} \coloneqq \binom{j}{i}, \quad \text{for } i,j \in [0\,..\,n-c].$$

Hence, we have $\det(E_{n;c}) = \det(L_{n;c})\det(U_{n;c})$.

Finally, observe that $L_{n;c}$ is a lower-triangular matrix (as $\binom{c}{i-j} = 0$ whenever $i < j$) and observe that $U_{n;c}$ is an upper-triangular matrix (as $\binom{j}{i} = 0$ whenever $i > j$). Hence, we obtain

$$\det(L_{n;c}) = \prod_{a=0}^{n-c}(L_{n;c})_{a,a} = \prod_{a=0}^{n-c}\binom{c}{0} = 1 \quad \text{and} \quad \det(U_{n;c}) = \prod_{a=0}^{n-c}(U_{n;c})_{a,a} = \prod_{a=0}^{n-c}\binom{a}{a} = 1;$$

which yields $\det(E_{n;c}) = 1$ and thus $\det(C_{n;c}) \in \{-1, 1\}$. Hence, we have $\det(C_{n;c}) \not\equiv_p 0$, which implies that $C_{n;c}$ is regular in $\mathbb{F}_p^{(n-c+1)\times(n-c+1)}$; completing the proof. ◾

We combine Lemmas 4.11 and 4.12 to obtain the main result of the section.

Lemma 4.13. *Let H denote a graph, let $\Gamma \subseteq \mathrm{Aut}(H)$ denote a p-group, and let Φ denote a graph property. Let $c < \ell(H)$ denote a nonnegative integer. Suppose that we have $\Phi(\varnothing) = 1$ and $\Phi(F) = 0$ for every $F \in \mathrm{FP}(\Gamma, H)$ with level $\ell(F) > \ell(H) - c$. Then, there is a fixed point $S \in \mathrm{FP}(\Gamma, H)$ with $\widehat{\Phi}(S) \not\equiv_p 0$ and $\ell(S) \geq c$.*

Proof. Let us define $n = \ell(H)$. Write $\mathbf{w}^c$ for the vector that consists in the $n - c + 1$ first entries of the vector $\mathbf{w}$ and write $\widehat{\mathbf{w}}^c$ for the vector that consists in the $n - c + 1$ last entries of the vector $\widehat{\mathbf{w}}$. First, we show $\widehat{\mathbf{w}}^c \equiv_p C_{n;c}\ \mathbf{w}^c$. Observe that we have $w_i = 0$ for all $i > n - c$. Hence, Lemma 4.11 yields

$$\widehat{w}_i \equiv_p \sum_{j=0}^{n}(-1)^j\binom{n-j}{i-j}\, w_j \equiv_p \sum_{j=0}^{n-c}(-1)^j\binom{n-j}{i-j}\, w_j;$$

This can be rewritten into $\widehat{\mathbf{w}}^c \equiv_p C_{n;c}\ \mathbf{w}^c$ by shifting the indices.

From $\Phi(\varnothing) = 1$, we conclude $w_0 = 1 \not\equiv_p 0$; hence, $\mathbf{w}^c$ is not the zero vector. Now, by Lemma 4.12, the matrix $C_{n,c}$ is regular; hence we obtain that $\widehat{\mathbf{w}}^c \equiv_p C_{n;c}\ \mathbf{w}^c$ cannot be the zero vector, either. As $\widehat{\mathbf{w}}^c$ contains the $n - c + 1$ last entries of $\widehat{\mathbf{w}}$, there is thus a fixed point $S \in \mathrm{FP}(\Gamma, H)$ with $\widehat{\Phi}(S) \not\equiv_p 0$ and $n \geq c$; completing the proof. ◾

5 Prime Powers and Difference Graphs

In this section, we obtain a first application of the techniques from Section 4. In particular, our goal for this section is to prove the following theorem.

Theorem 5.16. *Let Φ denote an edge-monotone graph property.*

- *If Φ is nontrivial on infinitely many prime powers, then #IndSub(Φ) is #W[1]-hard.*
- *Further, assuming ETH, there is a universal constant $\alpha > 0$ (independent of Φ) such that for any prime power $k \geq 3$ on which Φ is nontrivial, no algorithm (that reads the whole input) computes for every graph G the number #IndSub$((\Phi, k) \to G)$ in time $O(|V(G)|^{\alpha\sqrt{k}/\log k})$.* ◾

For our proof of Theorem 5.16, we consider the complete graph K_{p^m} for a prime p and a nonnegative integer m. In particular, we identify the vertices of K_{p^m} with elements of the finite field $\mathbb{F}_{p^m}$ and consider the rotation subgroup $\circlearrowleft_{p^m} \subseteq \mathrm{Aut}(K_{p^m})$—recall that the rotation group $\circlearrowleft_{p^m}$ consists in all automorphisms of K_{p^m} that map every vertex v to the vertex $v + b$ (for some $b \in \mathbb{F}_{p^m}$, where addition is in $\mathbb{F}_{p^m}$).

In a first step, we use Lemma 4.1 to understand the fixed points $\mathrm{FP}(\circlearrowleft_{p^m}, K_{p^m})$ by computing the orbits $E(K_{p^m})/\circlearrowleft_{p^m}$. In a second step, we then use Lemma 4.13 to obtain useful fixed points that have both nonvanishing alternating enumerator and high level. Finally, we show how to construct a suitable sequence of useful fixed points to obtain the desired hardness via Lemma A.8.

5.1 The Fixed Points of Rotations of K_{p^m}

As described, we wish to understand the orbits $E(K_{p^m})/\circlearrowleft_{p^m}$. To that end, it is instructive to recall the definition of difference graphs and of the level of a difference graph.

Definition 2.4. *For a prime p, an integer $m > 0$, and a set $A \subseteq \mathbb{F}^+_{p^m}$, we define the* difference graph $C^A_{p^m}$ *via*

$$V(C^A_{p^m}) := \mathbb{F}_{p^m} \quad \text{and} \quad E(C^A_{p^m}) := \{\{u, v\} : u, v \in \mathbb{F}_{p^m}, (u - v) \in A \cup (-A)\},$$

where $-A = \{-x : x \in A\}$. Observe that $C^A_{p^m} = K_{p^m}$ whenever $A = \mathbb{F}^+_{p^m}$.

The level *of a difference graph $C^A_{p^m}$ is the cardinality of A; we write $\ell(C^A_{p^m})$ for the level of $C^A_{p^m}$.*

For $m = 1$, the graphs C^A_p are also called *circulant graphs* in the literature [ABD⁺18, AP79].

Next, let us quickly confirm the intuitive statement that different sets $A, B \subseteq \mathbb{F}^+_{p^m}$ give rise to different difference graphs $C^A_{p^m}$ and $C^B_{p^m}$.

Lemma 5.1. *For any sets $A \neq B \subseteq \mathbb{F}^+_{p^m}$, we have $C^A_{p^m} \neq C^B_{p^m}$.*

Proof. Assume without loss of generality that $B \setminus A \neq \emptyset$; otherwise swap A and B. Now, consider an element $b \in B \setminus A$ and an $x \in \mathbb{F}_{p^m}$. Clearly, we have $\{x, x + b\} \in E(C^B_{p^m})$ and $\{x, x + b\} \notin E(C^A_{p^m})$.

Now, we obtain that the orbits $E(K_{p^m})/\circlearrowleft_{p^m}$ are the difference graphs of the singleton subsets of $\mathbb{F}^+_{p^m}$.

Lemma 5.2. *Let p denote a prime and let $m > 0$ denote an integer. Further, let $\circlearrowleft_{p^m}$ act on $E(K_{p^m})$. Then, we have $E(K_{p^m})/\circlearrowleft_{p^m} = \{C^{\{x\}}_{p^m} : x \in \mathbb{F}^+_{p^m}\}$.*[5]

Proof. First, we show that each edge set $E(C^{\{x\}}_{p^m})$ defines a unique orbit.

Claim 5.3. *For every $x \in \mathbb{F}^+_{p^m}$, we have $\circlearrowleft_{p^m} \cdot \{0, x\} = E(C^{\{x\}}_{p^m})$.*

Proof. For a $b \in \mathbb{F}_{p^m}$, write $\varphi_b \in \circlearrowleft_{p^m}$ for the rotation $x \mapsto (x + b)$. Fix an $x \in \mathbb{F}^+_{p^m}$. We compute the orbit of the edge $\{0, x\} \in E(K_{p^m})$ under $\circlearrowleft_{p^m}$ as

$$\begin{aligned} \circlearrowleft_{p^m} \cdot \{0, x\} &= \{\{\varphi_b(0), \varphi_b(x)\} : \varphi_b \in \circlearrowleft_{p^m}\} \\ &= \{\{b, x + b\} : b \in \mathbb{F}_{p^m}\} \\ &= \{\{u, v\} : u, v \in \mathbb{F}_{p^m}, u - v = x\} \\ &= E(C^{\{x\}}_{p^m}). \end{aligned}$$

Now, Claim 5.3 and Lemma 5.1 indeed yield that the edge sets $E(C^{\{x\}}_{p^m})$ form disjoint orbits.

Lastly, we check that $\{C^{\{x\}}_{p^m} : x \in \mathbb{F}^+_{p^m}\}$ are all orbits.

Claim 5.4. *We have $E(K_{p^m}) = \bigcup_{x \in \mathbb{F}^+_{p^m}} \circlearrowleft_{p^m} \cdot \{0, x\}$.*

[5] Again, we abuse notation and identify subsets of edges with the corresponding edge-subgraphs.

Proof. Consider an arbitrary edge $\{u, v\} \in E(K_{p^m})$. Via the rotation φ_u, we have $\{u, v\} \in \circlearrowleft_{p^m} \cdot \{0, v-u\}$; via the rotation φ_v, we have $\{u, v\} \in \circlearrowleft_{p^m} \cdot \{0, u-v\}$. Finally, we observe that $\mathbb{F}^{+}_{p^m}$ contains exactly one out of $u - v$ and $v - u$; completing the proof.

Taken together, Claims 5.3 and 5.4 yield the desired $E(K_{p^m})/\circlearrowleft_{p^m} = \{C^{\{x\}}_{p^m} : x \in \mathbb{F}^{+}_{p^m}\}$.

Combining Lemmas 4.1 and 5.2, we obtain a classification of the fixed points $\mathrm{FP}(\circlearrowleft_{p^m}, K_{p^m})$.

Lemma 5.5. *Let p denote a prime and let $m > 0$ denote an integer. Further, let $\circlearrowleft_{p^m}$ act on $\mathcal{E}(K_{p^m})$. Then, we have*

$$\mathrm{FP}(\circlearrowleft_{p^m}, K_{p^m}) = \{C^{A}_{p^m} : A \subseteq \mathbb{F}^{+}_{p^m}\}.$$

Proof. From Lemma 5.2, we obtain $E(K_{p^m})/\circlearrowleft_{p^m} = \{C^{\{x\}}_{p^m} : x \in \mathbb{F}^{+}_{p^m}\}$.

From Lemma 4.1, we obtain that each fixed point in $\mathrm{FP}(\circlearrowleft_{p^m}, K_{p^m})$ is the (disjoint) union of orbits from $E(K_{p^m})/\circlearrowleft_{p^m}$.

Finally, we readily convince ourselves that for sets $A, B \subseteq \mathbb{F}^{+}_{p^m}$ we have $C^{A\cup B}_{p^m} = C^{A}_{p^m} \cup C^{B}_{p^m}$.

Remark 5.6. Recall that the level of $C^{A}_{p^m}$ is the cardinality of A. From Lemma 5.5 we see that, indeed, $C^{A}_{p^m}$ (as a fixed point) consists in $|A|$ orbits and thus also has a level of $|A|$ as a fixed point.

Also consult Figure 3 for a visualization of an example.

Finally, we use that $C^{A}_{p^m}$ is regular to show a lower bound for the treewidth.

Corollary 5.7. *Let p denote a prime and let $m > 0$ denote an integer. Then, every fixed point $C^{A}_{p^m} \in \mathrm{FP}(\circlearrowleft_{p^m}, K_{p^m})$ has treewidth of at least $\ell(A)$.*

Proof. First, for all $x \in V(C^{A}_{p^m}) = \mathbb{F}_{p^m}$, we obtain that x is adjacent to all vertices of the form $x + c$ for $c \in A \cup (-A)$.

If $p \neq 2$, then A and $-A$ are disjoint by construction of $\mathbb{F}^{+}_{p^m}$, hence the graph is $2\ell(A)$-regular. Otherwise, $p = 2$ and $A = -A$, thus the graph is $\ell(A)$-regular. According to [BK11, Lemma 4] a d-regular graph has treewidth at least d, which proves the claim.

5.2 Nonvanishing Alternating Enumerators via Avalanches

Next, we wish to exploit our understanding of the fixed points $\mathrm{FP}(\circlearrowleft_{p^m}, K_{p^m})$. In particular, we intend to show that for an edge-monotone property Φ, a single high-level fixed point $A \in \mathrm{FP}(\circlearrowleft_{p^m}, K_{p^m})$ that satisfies Φ already implies that all fixed points of a small level also satisfy Φ. We can exploit this "avalanche" effect to obtain another criterion for #W[1]-hardness of #IndSub(Φ). To that end, we first need to understand when fixed points of $\mathrm{FP}(\circlearrowleft_{p^m}, K_{p^m})$ are isomorphic.

Definition 5.8. *Let p denote a prime and let $m > 0$ denote an integer. Two sets $A, B \subseteq \mathbb{F}^{+}_{p^m}$ are* isomorphic, *denoted by $A \sim B$, if there is a $\lambda \in \mathbb{F}^{*}_{p^m}$ such that $\lambda \cdot (A \cup (-A)) = B \cup (-B)$.*

Lemma 5.9. *Let p denote a prime and let $m > 0$ denote an integer. For any two sets $A, B \subseteq \mathbb{F}^{+}_{p^m}$ with $A \sim B$, we have $C^{A}_{p^m} \cong C^{B}_{p^m}$.*

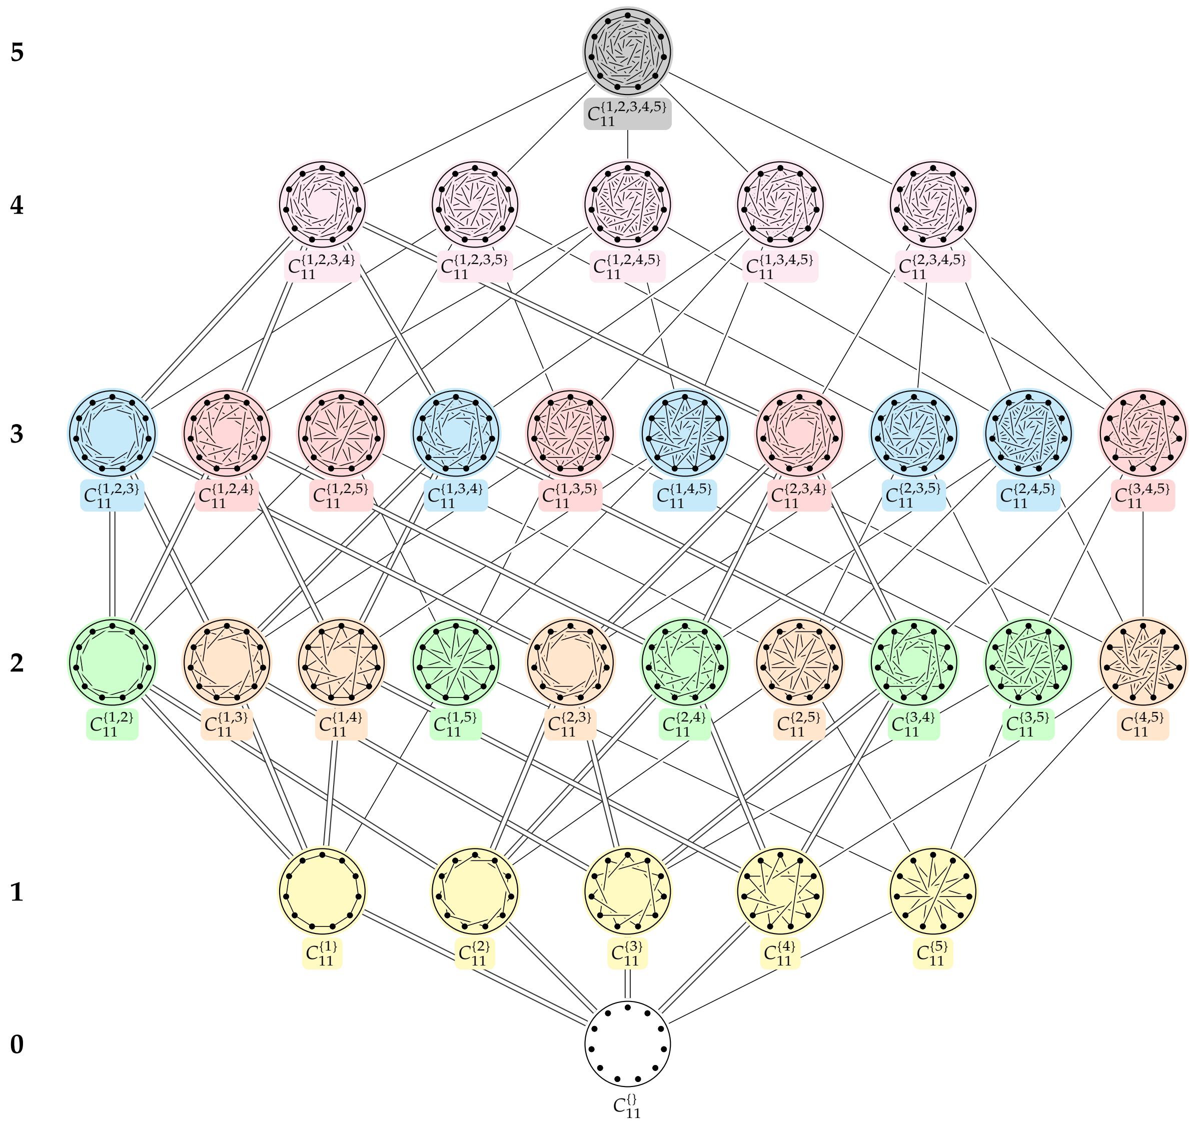

Figure 3 The fixed points $\mathrm{FP}(\circlearrowleft_{11}, K_{11})$ form the same lattice as subsets of a 5-element universe. We group the fixed points according to their level (which is denoted on the left side). Isomorphic graphs are colored with the same color. Two fixed points of adjacent levels are connected with an edge if one of them is a sub-point of the other. Double edges highlight the sub-points of $C_{11}^{\{1,2,3,4\}}$. Observe that if this fixed point satisfies an edge-monotone property Φ, then every fixed point on the third level (every fixed point in red or blue) satisfies Φ as well.

Proof. By definition, there is a $\lambda \in \mathbb{F}^+_{p^m}$ such that $\lambda \cdot (A \cup (-A)) = B \cup (-B)$.

Now, consider the function $\varphi_\lambda \colon V(K_{p^m}) \to V(K_{p^m})$ with $x \mapsto \lambda x$. As λ is invertible, φ_λ is an automorphism of K_{p^m} and thus bijective.

Next, for any $u, v \in \mathbb{F}_{p^m}$ with $u - v \in A \cup (-A)$, we have $\lambda u - \lambda v = \lambda(u - v) \in \lambda(A \cup (-A)) = B \cup (-B)$. Hence, φ_λ maps each edge $\{u, v\}$ of $C^A_{p^m}$ to the edge $\{\lambda u, \lambda v\}$ of $C^B_{p^m}$. Thus, φ_λ is also a homomorphism from $C^A_{p^m}$ to $C^B_{p^m}$; which is also surjective as $|\lambda \cdot (A \cup (-A))| = |B \cup (-B)|$.

In total, φ_λ is a homomorphism that is both surjective and injective; thus φ_λ is an isomorphism. ▟

Next, we show the aforementioned "avalanche" effect. Our key statement shows that every sufficiently small B set is isomorphic to a subset of A.

Lemma 5.10. *Let p denote a prime, let $m > 0$ denote an integer, and let $A \subsetneq \mathbb{F}^+_{p^m}$ denote a set.*
For every $B \subseteq \mathbb{F}^+_{p^m}$ with $|B| < |\mathbb{F}^+_{p^m}|/(|\mathbb{F}^+_{p^m}| - |A|)$, there is an isomorphic set $B' \sim B$ with $B' \subseteq A$.

Proof. Let $B \subseteq \mathbb{F}^+_{p^m}$ denote a set with $|B| < |\mathbb{F}^+_{p^m}|/(|\mathbb{F}^+_{p^m}| - |A|)$. We order B arbitrarily and define a vector $\vec{b} \in (\mathbb{F}^*_{p^m})^{|B|}$ via

$$\vec{b} := (B[0], \ldots, B[|B| - 1]).$$

It is easy but instructive to confirm that multiplying $\vec{b}$ with different nonzero elements yield vectors that differ at every position.

Claim 5.11. *Let $\lambda, \mu \in \mathbb{F}^*_{p^m}$ and let $i \in [0\,..\,|B|)$ denote a position.*
If we have $(\lambda\vec{b})_i = (\mu\vec{b})_i$, then $\lambda = \mu$.

Proof. By construction, each element $\vec{b}_i \in \mathbb{F}^*_{p^m}$ is invertible; which yields the claim. ◿

Next, write $(\star)^+$ for the function that maps each element from $\mathbb{F}^*_{p^m}$ to its corresponding representative in $\mathbb{F}^+_{p^m}$. Now, write $\mathbb{F}^*_{p^m}\vec{b} := \{\lambda\vec{b} : \lambda \in \mathbb{F}^*_{p^m}\}$ for the set of multiples of $\vec{b}$ and write $(\mathbb{F}^*_{p^m}\vec{b})^+ := \{(\lambda\vec{b})^+ : \lambda \in \mathbb{F}^*_{p^m}\}$ for the set of all multiples of $\vec{b}$ when "ignoring the signs" of the entries of the multiples. In particular, each vector in $(\mathbb{F}^*_{p^m}\vec{b})^+$ corresponds to some $B' \subseteq \mathbb{F}^+_{p^m}$ with $B \sim B'$ (which need not be different for each vector).

Claim 5.12. *We have $|\mathbb{F}^*_{p^m}\vec{b}| = |\mathbb{F}^*_{p^m}|$ and $|(\mathbb{F}^*_{p^m}\vec{b})^+| \geq |\mathbb{F}^+_{p^m}|$.*

Proof. The first equality is immediate from Claim 5.11 (applied to the first elements of the vectors in $\mathbb{F}^*_{p^m}\vec{b}$).

For the second equality, first consider the case $p = 2$. Now, we have $\mathbb{F}^*_{2^m} = \mathbb{F}^+_{2^m}$ and $(\star)^+$ is the identity; which yields the claim.

Now, for $p \neq 2$, we have $|\mathbb{F}^*_{p^m}| = 2|\mathbb{F}^+_{p^m}|$. Further, $(\star)^+$ identifies at most 2 different elements from $\mathbb{F}^*_{p^m}$ to the same element from $\mathbb{F}^+_{p^m}$; which in particular holds for the first elements of the vectors in $\mathbb{F}^*_{p^m}\vec{b}$. Taken together, we obtain the claim. ◿

Next, we say that an element $s \in \mathbb{F}^+_{p^m} \setminus A$ *sullies* a vector $\lambda\vec{b} \in (\mathbb{F}^*_{p^m}\vec{b})^+$ if there is a position $i \in [0\,..\,|B|)$ with $(\lambda\vec{b})_i = s$. Similarly, we say that an element $s \in \mathbb{F}^+_{p^m} \setminus A$ *sullies* a position $i \in [0\,..\,|B|)$ if there is a vector $\lambda\vec{b} \in (\mathbb{F}^*_{p^m}\vec{b})^+$ with $(\lambda\vec{b})_i = s$.

We observe that we can prove the lemma by showing that there is a vector in $(\mathbb{F}^*_{p^m}\vec{b})^+$ that is not sullied by any $s \in \mathbb{F}^+_{p^m} \setminus A$.

To that end, we observe that each element $s \in \mathbb{F}_{p^m}^+$ sullies at most one position (by construction, all entries of $\vec{b}$ are pairwise different) and at most one vector (by Claim 5.11). In particular, this means that each element $s \in \mathbb{F}_{p^m}^+$ sullies at most $|B|$ vectors of $(\mathbb{F}_{p^m}^* \vec{b})^+$. Hence, in total, at most $|B| \cdot |\mathbb{F}_{p^m}^+ \setminus A|$ vectors of $(\mathbb{F}_{p^m}^* \vec{b})^+$ are sullied. Finally, we plug in $|B| < |\mathbb{F}_{p^m}^+|/(|\mathbb{F}_{p^m}^+| - |A|)$ to obtain

$$|B| \cdot |\mathbb{F}_{p^m}^+ \setminus A| < |\mathbb{F}_{p^m}^+| \leq |(\mathbb{F}_{p^m}^* \vec{b})^+|;$$

from which we conclude that, indeed, there is a vector in $(\mathbb{F}_{p^m}^* \vec{b})^+$ that is not sullied. This in turn completes the proof. ◀

We combine Lemmas 5.9 and 5.10, to readily obtain that any (nontrivial) edge-monotone graph property has to be true for all difference graphs with a small level.

▶ **Corollary 5.13.** *Let p denote a prime, let $m > 0$ denote an integer, and let $A \subsetneq \mathbb{F}_{p^m}^+$ denote a set. Further, let Φ denote an edge-monotone graph property with $\Phi(C_{p^m}^A) = 1$.*

For every $B \subseteq \mathbb{F}_{p^m}^+$ with $|B| < |\mathbb{F}_{p^m}^+|/(|\mathbb{F}_{p^m}^+| - |A|)$, we have $\Phi(C_{p^m}^B) = 1$.

Proof. Let $B \subseteq \mathbb{F}_{p^m}^+$ denote a set with $|B| < |\mathbb{F}_{p^m}^+|/(|\mathbb{F}_{p^m}^+| - |A|)$.

From Lemma 5.10, we obtain an isomorphic subset $B' \sim B$ with $B' \subseteq A$. From Lemma 5.9, we obtain $C_{p^m}^{B'} \cong C_{p^m}^B$. Combined, we thus obtain that $C_{p^m}^B$ is isomorphic to an edge-subgraph of $C_{p^m}^A$. Finally, as Φ is edge-monotone, we obtain $\Phi(C_{p^m}^B) = 1$; which completes the proof. ◀

Next, we combine Corollary 5.13 and Lemma 4.8 to obtain another criterion for #W[1]-hardness of #IndSub(Φ).

▶ **Lemma 5.14.** *Let p denote a prime, let $m > 0$ denote an integer, and let Φ denote an edge-monotone graph property that is nontrivial on p^m. Further, write c and d for positive integers with $cd \leq |\mathbb{F}_{p^m}^+|$. Suppose that there is a fixed point $C_{p^m}^A \in \mathrm{FP}(\circlearrowleft_{p^m}, K_{p^m})$ with $\ell(C_{p^m}^A) \geq |\mathbb{F}_{p^m}^+| - d$ and $\Phi(C_{p^m}^A) = 1$.*

Then, there is a fixed point $C_{p^m}^B \subseteq C_{p^m}^A$ with $\ell(C_{p^m}^B) \geq c$ and $\widehat{\Phi}(C_{p^m}^B) \not\equiv_p 0$.

Proof. First, we recall that we have $\ell(C_{p^m}^A) = \ell(A) = |A|$. Further, we have

$$c = \frac{cd}{d} \leq \frac{|\mathbb{F}_{p^m}^+|}{|\mathbb{F}_{p^m}^+| - (|\mathbb{F}_{p^m}^+| - d)} \leq \frac{|\mathbb{F}_{p^m}^+|}{|\mathbb{F}_{p^m}^+| - |A|}.$$

Now, from Corollary 5.13, we obtain that $\Phi(C_{p^m}^S) = 1$ for all S with $|S| < c \leq |\mathbb{F}_{p^m}^+|/(|\mathbb{F}_{p^m}^+| - |A|)$.

As Φ is nontrivial on p^m, we have $\Phi(C_{p^m}^{\mathbb{F}_{p^m}^+}) = \Phi(K_{p^m}) = 0$. Thus, there is a minimal level $b \geq c$ on which not every fixed point satisfies Φ. Consider such a fixed point $C_{p^m}^B$ with level $\ell(B) = b$ that does not satisfy Φ. As b is minimal, all proper sub-points of $C_{p^m}^B$ do satisfy Φ. Hence, we may use Lemma 4.8 to obtain $\widehat{\Phi}(C_{p^m}^B) \not\equiv_p 0$; thus completing the proof. ◀

On the one hand, we may use Lemma 5.14 whenever we have a fixed point that satisfies Φ and has a high level. On the other hand, we may use Lemma 4.13 if there is no fixed point with a high level that satisfies Φ. Hence, if we combine Lemmas 4.13 and 5.14, then we can show that there is a fixed point with a nonvanishing alternating enumerator and a treewidth of roughly $p^{m/2}$ if Φ is nontrivial on p^m.

▶ **Corollary 5.15.** *Let p denote a prime, let $m > 0$ denote an integer, and let Φ denote an edge-monotone graph property that is nontrivial on p^m.*

Then, there is a fixed point $C_{p^m}^A \in \mathrm{FP}(\circlearrowleft_{p^m}, K_{p^m})$ that has a nonvanishing alternating enumerator and a treewidth of at least $p^{m/2}/2 - 2$.

Proof. We show that there is a fixed point $C^A_{p^m}$ with $\widehat{\Phi}(C^A_{p^m}) \not\equiv_p 0$ and $\ell(A) \geq p^{m/2}/2 - 2$. Then, the claim follows from Corollary 5.7.

First, suppose that a fixed point $C^S_{p^m}$ with $\ell(S) \geq |\mathbb{F}^+_{p^m}| - d$ satisfies Φ. As the property Φ is nontrivial on p^m, this means that we can use Lemma 5.14 to obtain a fixed point $C^A_{p^m}$ with $\widehat{\Phi}(C^A_{p^m}) \not\equiv_p 0$ and $\ell(A) \geq c$ for any c with $cd \leq |\mathbb{F}^+_{p^m}|$.

Next, suppose that no fixed point $C^S_{p^m}$ with $\ell(S) \geq n - d$ satisfies Φ. As the property Φ is nontrivial on p^m, this means that we can use Lemma 4.13 to obtain a fixed point $C^A_{p^m}$ with $\widehat{\Phi}(C^A_{p^m}) \not\equiv_p 0$ and $\ell(A) \geq d$.

In both cases we obtain a nonvanishing fixed point $C^A_{p^m}$ with level at least c. Choosing $c := d$, in both cases, we obtain a nonvanishing fixed point $C^A_{p^m}$ with level at least d.

Finally, to ensure that $cd = d^2$ is at most $|\mathbb{F}^+_{p^m}|$, set $d := \lfloor |\mathbb{F}^+_{p^m}|^{1/2} \rfloor$. Recalling Lemma 2.3, we observe

$$d = \lfloor |\mathbb{F}^+_{p^m}|^{1/2} \rfloor \geq |\mathbb{F}^+_{p^m}|^{1/2} - 1 \geq (p^m/2 - 1)^{1/2} - 1 \geq p^{m/2}/2 - 2.$$

Now, the claim follows from Corollary 5.7. ∎

5.3 #W[1]-hardness and Quantitative Lower Bounds for Prime Powers

Corollary 5.15 directly implies our first result.

Theorem 5.16. *Let Φ denote an edge-monotone graph property.*

- *If Φ is nontrivial on infinitely many prime powers, then #IndSub(Φ) is #W[1]-hard.*
- *Further, assuming ETH, there is a universal constant $\alpha > 0$ (independent of Φ) such that for any prime power $k \geq 3$ on which Φ is nontrivial, no algorithm (that reads the whole input) computes for every graph G the number $\#\mathrm{IndSub}((\Phi, k) \to G)$ in time $O(|V(G)|^{\alpha\sqrt{k}/\log k})$.*

Proof. We start with #W[1]-hardness. By assumption on Φ, for every $k \in \mathbb{N}$, there is a prime number p_k and a positive integer m_k such that Φ is nontrivial on $p_k^{m_k} \geq k$.

From Corollary 5.15, we obtain a graph $H_k := C^A_{p_k^{m_k}}$ with $p_k^{m_k}$ vertices, nonvanishing alternating enumerator, and $\mathrm{tw}(H_k) \geq p_k^{m_k/2}/2 - 2 \geq \sqrt{k}/2 - 2$. Now, using the constructed sequence of graphs, Lemma A.8 yields #W[1]-hardness of #IndSub(Φ).

We proceed to the ETH-based lower bound. To that end, write $\alpha_{\textsc{IndSub}} > 0$ for the constant from Lemma A.8 and set $\alpha' := \alpha_{\textsc{IndSub}}/3$.

Now, fix a prime power $k \geq 12^2$ such that Φ is nontrivial on k. From Corollary 5.15, we obtain a graph $H_k := C^A_k$ with k vertices, nonvanishing alternating enumerator, and

$$\mathrm{tw}(H_k) \geq \sqrt{k}/2 - 2 \geq \sqrt{12^2}/2 - 2 \geq 2.$$

Thus, by Lemma A.8 and assuming ETH, there is no algorithm that for each graph G computes the number $\#\mathrm{IndSub}((\Phi, k) \to G)$ in time $O(|V(G)|^{\alpha_{\textsc{IndSub}}\,\mathrm{tw}(H_k)/\log \mathrm{tw}(H_k)})$. We complete the proof by showing the following inequality.

Claim 5.17. *For $12^2 \leq k$, we have*

$$\alpha' \frac{\sqrt{k}}{\log k} = \alpha_{\textsc{IndSub}} \frac{\sqrt{k}/3}{\log k} \leq \alpha_{\textsc{IndSub}} \frac{\mathrm{tw}(H_k)}{\log \mathrm{tw}(H_k)}.$$

Proof. For $k \geq 12^2$, we have $\sqrt{k}/3 \leq \sqrt{k}/2 - 2$. Further, we define $h_1\colon \mathbb{R}_{>1} \to \mathbb{R}, x \mapsto x/\log(x)$. The derivative of this function is $(\log(x) - 1)/(\log^2(x))$, thus the function is monotonically increasing for $x \geq \mathrm{e}$. Since $\mathrm{e} < \sqrt{k}/3 \leq \sqrt{k}/2 - 2 \leq \mathrm{tw}(H_k)$, we obtain

$$\alpha_{\textsc{IndSub}} \frac{\sqrt{k}/3}{\log(k)} \leq \alpha_{\textsc{IndSub}} \frac{\sqrt{k}/3}{\log(\sqrt{k}/3)} \leq \alpha_{\textsc{IndSub}} \frac{\sqrt{k}/2 - 2}{\log(\sqrt{k}/2 - 2)} \leq \alpha_{\textsc{IndSub}} \frac{\mathrm{tw}(H_k)}{\log \mathrm{tw}(H_k)};$$

which completes the proof. ⌟

Now, from Claim 5.17, we obtain

$$O(|V(G)|^{\alpha' \sqrt{k}/\log k}) \subseteq O(|V(G)|^{\alpha_{\textsc{IndSub}} \mathrm{tw}(H_k)/\log \mathrm{tw}(H_k)}).$$

Finally, to obtain the claim also for all $3 \leq k < 12^2$, we chose $\alpha := \min(\alpha', 1/12)$. Observe that for $k < 12^2$, we obtain

$$O(|V(G)|^{\alpha \sqrt{k}/\log k}) = o(|V(G)|).$$

Now, such a running time is unconditionally unachievable for any algorithm that reads the whole input. This completes the proof. ∎

6 Main Result 1: #W[1]-hardness for Edge-monotone Properties

In this section, we prove Main Theorem 1.

Main Theorem 1. *Let Φ denote a nontrivial edge-monotone graph property.*
- *The problem $\#\textsc{IndSub}(\Phi)$ is #W[1]-hard.*
- *Further, assuming ETH, there is a universal constant $\gamma > 0$ (independent of Φ) such that for any integer $k \geq 3$ on which Φ is nontrivial, no algorithm (that reads the whole input) computes for every graph G the number $\#\mathrm{IndSub}((\Phi, k) \to G)$ in time $O(|V(G)|^{\gamma \sqrt{\log k}/\log\log k})$.* ∎

We prove Main Theorem 1 by raising the techniques of Section 5 from prime powers p^m to multiples $d \cdot p^m$ of prime powers. In Sections 6.1 and 6.2, we introduce group-theoretic and graph-theoretic concepts that allow us to do this generalization. While the notions and the arguments are fairly natural, the statements and the proofs come with some notational overhead. One important point to keep in mind is that the resulting fixed points on $d \cdot p^m$ vertices are not just the disjoint unions of d fixed points on p^m vertices, but may contain complete bipartite graphs between some of the copies.

In Sections 6.4 and 6.5, we exploit our understanding of these fixed points to prove Main Theorem 1. Compared to the hardness proofs in Section 5, there is an extra Inclusion-Exclusion step to obtain a reduction from p^m vertices to $d \cdot p^m$ vertices.

6.1 Product Groups, Graphs Unions, and Graph Joins

We start by defining the disjoint union of sets and product groups.

Definition 6.1. *For (not necessarily disjoint) sets $X_1, \dots, X_m$, the disjoint union $X_1 \uplus \cdots \uplus X_m$ is the set $\{(i, x) : i \in [m], x \in X_i\}$.*

For graphs $G_1, \dots, G_m$ and a graph $C \in \mathcal{G}_m$, we define the inhabited graph $C\langle G_1, \dots, G_m \rangle$ *via*

$$V(C\langle G_1, \dots, G_m \rangle) := V(G_1) \uplus \ldots \uplus V(G_m)$$
$$E(C\langle G_1, \dots, G_m \rangle) := \{\{(i, v_i), (j, u_j)\} : \{i, j\} \in E(C) \text{ or } (i = j \text{ and } \{v_i, u_i\} \in E(G_i))\}.$$

If C consists in a single edge $\{i,j\}$, then we also write

$$\{i,j\}\big\langle G_1,\dots,G_m\big\rangle := ([m],\{i,j\})\big\langle G_1,\dots,G_m\big\rangle.$$

If $C = IS_m$, then we call the graph $IS_m\big\langle G_1,\dots,G_m\big\rangle$ the disjoint union *of $G_1,\dots,G_m$ and also write*

$$G_1 \uplus \dots \uplus G_m := IS_m\big\langle G_1,\dots,G_m\big\rangle.$$

If $C = K_m$, then we call the graph $K_m\big\langle G_1,\dots,G_m\big\rangle$ the join *of $G_1,\dots,G_m$ and also write*

$$G_1 \triangledown \dots \triangledown G_m := K_m\big\langle G_1,\dots,G_m\big\rangle.$$

Consult Figure 4 for a visualization of an example for Definition 6.1. It is instructive to briefly discuss unions of inhabited graphs.

Lemma 6.2. *For two inhabited graphs $F_1 = C_1\big\langle G_1,\dots,G_m\big\rangle$ and $F_2 = C_2\big\langle G_1',\dots,G_m'\big\rangle$ with a common vertex set, we have*[6]

$$F_1 \cup F_2 = (C_1 \cup C_2)\big\langle G_1 \cup G_1',\dots,G_m \cup G_m'\big\rangle.$$

Proof. The union of (edge) sets is an associative operation. Unfolding the definitions of inhabited graphs and their union yields the claim.

Next, we define a standard operation on groups that mirrors the disjoint union of graphs. Namely, the products of groups.

Definition 6.3. *For permutation groups $\Gamma_1 = (G_1,\circ),\dots,\Gamma_m = (G_m,\circ)$ with $\Gamma_i \subseteq \mathfrak{S}_{X_i}$, their product group $\Gamma := \Gamma_1 \times \dots \times \Gamma_m$ is the set $G := G_1 \times \dots \times G_m$ together with the component-wise function composition, that is, for $(g_1,\dots,g_m),(g_1',\dots,g_m') \in \Gamma$, we set*

$$(g_1,\dots,g_m)\circ(g_1',\dots,g_m') := (g_1 \circ g_1',\dots,g_m \circ g_m').$$

We let Γ act on $X := X_1 \uplus \dots \uplus X_m$ via

$$g(j,x_j) := (j,g_j(x_j)),\ \text{for all } g = (g_1,\dots,g_m) \in G \text{ and } x = (j,x_j) \in X.$$

Fact 6.4. *The order of G is $\prod_{i=1}^m |G_i|$. In particular, the product group of p-groups is still a p-group. Further, observe that G is a permutation group of $X_1 \uplus \dots \uplus X_m$. Hence, we have $G \subseteq \mathfrak{S}_{X_1 \uplus \dots \uplus X_m}$.*

Finally, we let product groups act on joins of graphs; in particular, we are interested in the resulting fixed points. To that end, we first study the orbits that appear.

Lemma 6.5. *For $i \in [m]$, let G_i denote a graph with $n(i) := |V(G_i)|$ vertices and let $\Gamma_i \subseteq \mathrm{Aut}(G_i)$ denote a transitive permutation group. Further, let Γ_i act on $E(G_i)$ and write $E(G_i)/\Gamma_i = \{O_1^i,\dots,O_{s_i}^i\}$ for the resulting orbits. Finally, let $\bigtimes_{i=1}^m \Gamma_i$ act on $E(\triangledown_{i=1}^m G_i)$.*

Then, we have

$$E\left(\bigtriangledown_{i=1}^m G_i\right) \Big/ \bigtimes_{i=1}^m \Gamma_i = \{\{i,j\}\big\langle \mathrm{IS}_{n(1)},\dots,\mathrm{IS}_{n(m)}\big\rangle : i \neq j \in [m]\} \\ \cup \{\mathrm{IS}_m\big\langle \mathrm{IS}_{n(1)},\dots,\mathrm{IS}_{n(i-1)},O_j^i,\mathrm{IS}_{n(i+1)},\dots,\mathrm{IS}_{n(m)}\big\rangle : i \in [m], j \in [s_i]\}.$$

[6] Recall that we use $G \cup H$ to denote the graph on $V = V(G) = V(H)$ with edges $E(G) \cup E(H)$.

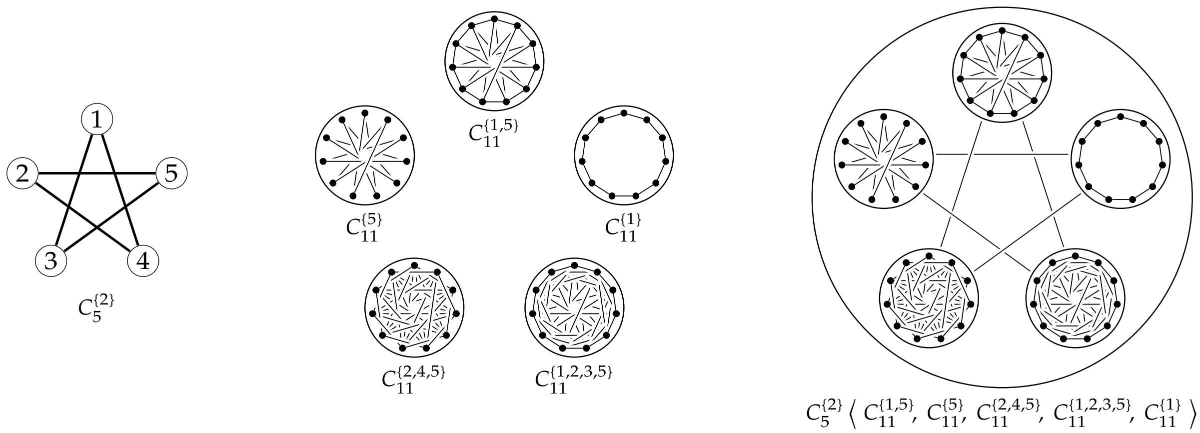

(a) A 5-cycle that is inhabited by 5 difference graphs.

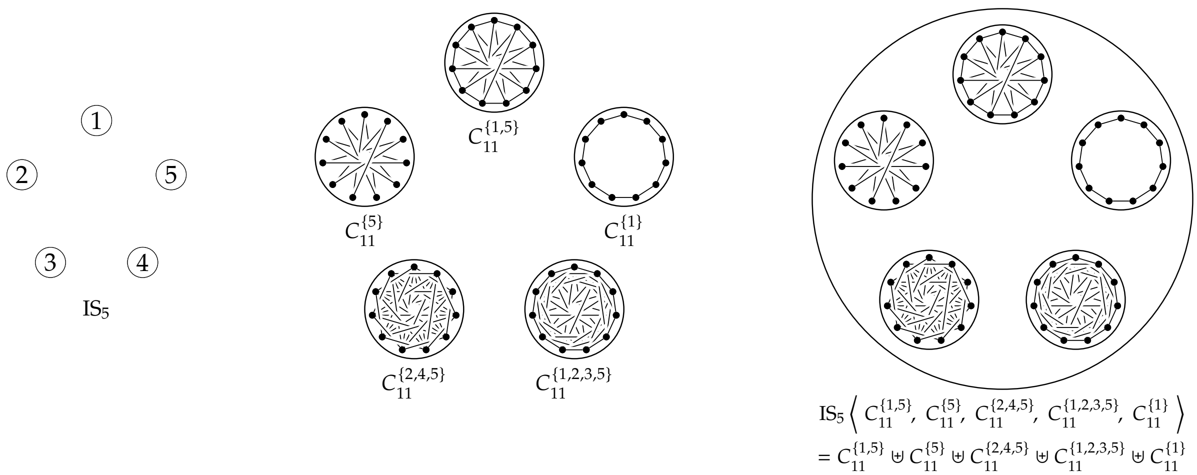

(b) An independent set that is inhabited by 5 difference graphs. The resulting graph is also the disjoint union of the 5 graphs.

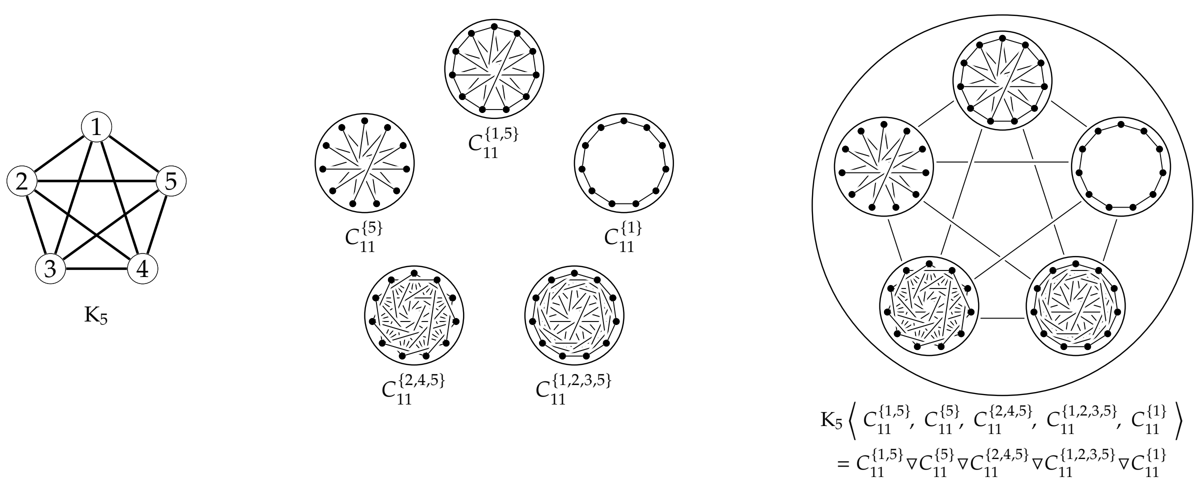

(c) An clique that is inhabited by 5 difference graphs. The resulting graph is also the join of the 5 graphs.

Figure 4 Examples for fixed points in $\mathrm{FP}(\cup_5^d, K_5^d)$, which are also inhabited graphs.

Proof. Recall that $\Gamma := \bigtimes_{i=1}^m \Gamma_i$ acts on $E(\nabla_{i=1}^m G_i)$ by mapping the group element $(\alpha_1, \dots, \alpha_m)$ and the edge $\{(i,u),(j,v)\}$ to $\{(i,\alpha_i(u)),(j,\alpha_j(v))\}$.

We start by analyzing the orbits of the edges inside of the graphs G_i.

Claim 6.6. *For every $i \in [m]$ and $O \in E(G_i)/\Gamma_i$, we have*

$$\mathrm{IS}_m \left\langle \mathrm{IS}_{n(1)}, \dots, \mathrm{IS}_{n(i-1)}, O, \mathrm{IS}_{n(i+1)}, \dots, \mathrm{IS}_{n(m)} \right\rangle \in E\left(\bigtriangledown_{i=1}^{m} G_i\right) / \bigtimes_{i=1}^{m} \Gamma_i.$$

Proof. Fix an orbit $O \in E(G_i)/\Gamma_i$ and write $\{u,v\} \in O$ for an edge of O. We compute the orbit of the edge $\{(i,u),(i,v)\}$ under Γ.

Observe that whenever Γ acts on an edge $\{(i,u),(i,v)\}$ of $\nabla_{i=1}^m G_i$, we recover the action of Γ_i on said edge: writing $\star|_2$ for the restriction of pairs to their second component, we obtain

$$(g_1, \dots, g_i, \dots, g_m)\,\{(i,u),(i,v)\}\,|_2 = \{(i,g_i(u)),(i,g_i(v))\}\,|_2 = \{g_i(u), g_i(v)\} = g_i\{u,v\}.$$

In particular, we obtain $\Gamma\{(i,u),(i,v)\}|_2 = \Gamma_i\{u,v\} = O$.

Finally, we observe that adding isolated vertices to a graph has no effect on the orbits of the edges of the graph, which yields

$$\Gamma\{(i,u),(i,v)\} = \mathrm{IS}_m \left\langle \mathrm{IS}_{n(1)}, \dots, \mathrm{IS}_{n(i-1)}, O, \mathrm{IS}_{n(i+1)}, \dots, \mathrm{IS}_{n(m)} \right\rangle;$$

and thus the claim. ◿

Next, we analyze the orbits of edges between graphs G_i.

Claim 6.7. *For every $i \neq j \in [m]$, we have*

$$\{i,j\}\left\langle \mathrm{IS}_{n(1)}, \dots, \mathrm{IS}_{n(m)} \right\rangle \in E\left(\bigtriangledown_{i=1}^{m} G_i\right) / \bigtimes_{i=1}^{m} \Gamma_i.$$

Proof. Fix $F := \{i,j\}\left\langle \mathrm{IS}_{n(1)}, \dots, \mathrm{IS}_{n(m)} \right\rangle$ and observe that F consists in isolated vertices and a complete bipartite graph on $n(i) + n(j)$ vertices.

We consider vertices $u \in V(G_i)$ and $v \in V(G_j)$ and compute the orbit of the edge $\{(i,u),(j,v)\}$ under Γ. To that end, we observe that, as Γ_i and Γ_j are transitive, we have $\Gamma_i u = V(G_i)$ and $\Gamma_j v = V(G_j)$. In particular, we have $\Gamma\{(i,u),(j,v)\} = \{(i,a),(j,b) : a \in V(G_i), b \in V(G_j)\}$—which are precisely the edges of F, which in turn yields the claim. ◿

Finally, we readily see that Claims 6.6 and 6.7 cover all orbits in $E\left(\nabla_{i=1}^m G_i\right)/\bigtimes_{i=1}^m \Gamma_i$. First, any edge $\{(i,u),(i,v)\}$ is covered by an orbit that corresponds to an orbit in $E(G_i)/\Gamma_i$. Second, any edge $\{(i,u),(j,v)\}$ is covered by the orbit $\{i,j\}\left\langle \mathrm{IS}_{n(1)}, \dots, \mathrm{IS}_{n(m)} \right\rangle$. In total, this completes the proof. ∎

As before, we use orbits to build fixed points.

Corollary 6.8. *For $i \in [m]$, let G_i denote a graph and let $\Gamma_i \subseteq \mathrm{Aut}(G_i)$ denote a transitive permutation group. Further, let Γ_i act on $E(G_i)$ and write $\mathrm{FP}(\Gamma_i, G_i)$ for the resulting fixed points. Finally, let $\bigtimes_{i=1}^m \Gamma_i$ act on $\mathcal{E}(\nabla_{i=1}^m G_i)$.*

Then, we have

$$\mathrm{FP}(\bigtimes_{i=1}^{m} \Gamma_i, \bigtriangledown_{i=1}^{m} G_i) = \left\{ C\left\langle A^1, \dots, A^m \right\rangle : C \in \mathcal{G}_m, A^i \in \mathrm{FP}(\Gamma_i, G_i) \right\},$$

Corollary 6.8 allows us to interpret the fixed points $\mathrm{FP}(\bigtimes_{i=1}^{m}\Gamma_i, \nabla_{i=1}^{m} G_i)$ as combinations of the fixed points of $\mathrm{FP}(\Gamma_i, G_i)$. In particular, each fixed point $C\langle A^1, \ldots, A^m\rangle$ of consists in m blocks, where the i-th block is a fixed point $A^i \in \mathrm{FP}(\Gamma_i, G_i)$ and we fully connect two blocks i and j with each other if $\{i, j\} \in E(C)$.

Proof. Recall that by Lemma 4.1 (the edge set of) each fixed point $F \in \mathrm{FP}(\bigtimes_{i=1}^{m}\Gamma_i, \nabla_{i=1}^{m} G_i)$ is a union of orbits from $E\left(\nabla_{i=1}^{m} G_i\right)/\bigtimes_{i=1}^{m}\Gamma_i$.

From Lemma 6.5, we understand the orbits $E\left(\nabla_{i=1}^{m} G_i\right)/\bigtimes_{i=1}^{m}\Gamma_i$ as inhabited graphs that are pairwise compatible. Further, from Lemma 6.2, we obtain that we may compute the union of compatible inhabited graphs in a block-wise fashion.

Next, we apply Lemma 4.1 to each of the blocks, that is, each fixed point $F_i \in \mathrm{FP}(\Gamma_i, G_i)$ is a union of orbits from $E(G_i)/\Gamma_i$ (which are precisely the inner blocks of the orbits $E\left(\nabla_{i=1}^{m} G_i\right)/\bigtimes_{i=1}^{m}\Gamma_i$). Finally, we observe that we may obtain any m-vertex graph as the (edge-)union of m-vertex graphs that have a single edge.

In total, this completes the proof. ∎

6.2 The Fixed Points of Rotations of $K_{p^m}^d$

For a positive integer d and a prime power p^m, write $K_{p^m}^d := K_{p^m} \nabla \cdots \nabla K_{p^m}$ for the join of d copies of the graph K_{p^m}—observe that $K_{p^m}^d$ is the complete graph on the vertex set $[d] \times \mathbb{F}_{p^m}$. Further, write $\circlearrowleft_{p^m}^d := \circlearrowleft_{p^m} \times \ldots \times \circlearrowleft_{p^m}$ for the product of d copies of the group $\circlearrowleft_{p^m}$.

Next, we use Lemma 5.5 and Corollary 6.8 to understand $\mathrm{FP}(\circlearrowleft_{p^m}^d, K_{p^m}^d)$.

Lemma 6.9. *Let p denote a prime and let m and d denote positive integers. Further, let $\circlearrowleft_{p^m}^d = K_{p^m} \nabla \cdots \nabla K_{p^m}$ act on $\mathcal{E}(K_{p^m}^d) = \circlearrowleft_{p^m} \times \ldots \times \circlearrowleft_{p^m}$. Then, we have*

$$\mathrm{FP}(\circlearrowleft_{p^m}^d, K_{p^m}^d) = \{G\langle C_{p^m}^{A^1}, \ldots, C_{p^m}^{A^d}\rangle : G \in \mathcal{G}_d, A^i \subseteq \mathbb{F}_{p^m}^{+}\}.$$

Proof. Recall that by Lemma 5.5, we have

$$\mathrm{FP}(\circlearrowleft_{p^m}, K_{p^m}) = \{C_{p^m}^{A} : A \subseteq \mathbb{F}_{p^m}^{+}\}.$$

Observe that $\circlearrowleft_{p^m}$ is transitive, since for all $x, y \in \mathbb{F}_{p^m}$ we can find a element $\varphi_{y-x} \in \circlearrowleft_{p^m}$ with $\varphi_{y-x}(x) = x + (y - x) = y$. Hence, we use Corollary 6.8 to obtain the claim. ∎

Consult again Figure 4 for visualizations of examples for fixed points $\mathrm{FP}(\circlearrowleft_{p^m}^d, K_{p^m}^d)$.

Next, for our ETH-based lower bounds, it is instructive to observe that every fixed point $G\langle \cdots \rangle \in \mathrm{FP}(\circlearrowleft_{p^m}^d, K_{p^m}^d)$ contains a large biclique as a subgraph as long as G contains at least one edge.

Lemma 6.10. *Let p denote a prime and let m and d denote positive integers.*

Every fixed point $G\langle C_{p^m}^{A^1}, \ldots, C_{p^m}^{A^d}\rangle \in \mathrm{FP}(\circlearrowleft_{p^m}^d, K_{p^m}^d)$ with $G \neq \mathrm{IS}_m$ contains the graph K_{p^m,p^m} as a subgraph and has treewidth of at least p^m.

Proof. Consider an edge $\{i, j\} \in E(G)$. Both graphs $C_{p^m}^{A^i}$ and $C_{p^m}^{A^j}$ have p^m vertices; by definition of an inhabited graph, said vertices are connected with a complete bipartite graph.

Finally, a graph with K_{p^m,p^m} as a subgraph has a treewidth of at least p^m. ∎

Finally, let us compute the level of a fixed point $F \in \mathrm{FP}(\circlearrowleft_{p^m}^d, K_{p^m}^d)$.

Lemma 6.11. *Let p denote a prime and let m and d denote positive integers. For every fixed point $F = G\langle C_{p^m}^{A^1}, \dots, C_{p^m}^{A^d}\rangle \in \mathrm{FP}(\mho_{p^m}^d, K_{p^m}^d))$, we have*

$$\ell(F) = \ell(G\langle C_{p^m}^{A^1}, \dots, C_{p^m}^{A^d}\rangle) = |E(G)| + \sum_{i=1}^{d} |A^i|.$$

Proof. Corollary 6.8 and Lemma 6.5 allow us to compute the level of F as the sum of the level of G and the levels of the fixed points $C_{p^m}^{A}$. To that end, we observe that every edge of G contributes one orbit to the level of G; each fixed point $C_{p^m}^{A}$ has a level of $|A|$. In total, this yields the claim. ∎

6.3 The Property $(\Phi - H)$

We often use the following definition: given a property Φ and a graph H, we define $(\Phi - H)$ to be the property that contains a graph only if it satisfies Φ when we extend the graph with H as a disjoint union. Using standard techniques, we show that #IndSub$(\Phi - H)$ is not harder than #IndSub(Φ): we simply add a copy of H to the input graph and use the Inclusion-Exclusion Principle to count only those induced subgraphs that fully contain this copy.

Lemma 6.12. *Let Φ denote a graph property and suppose that there is an algorithm $\mathbb{A}$ that computes for each graph G and integer k the value* #IndSub$((\Phi, k) \to G)$ *in time $g(k, |V(G)|)$ for some computable function g that is monotonically increasing. Finally, for a graph H, write $(\Phi - H) := \{G : G \uplus H \in \Phi\}$ for the graph property of all graphs that is extended by H to a graph in Φ.*

Then, there is an algorithm $\mathbb{B}$ with oracle access to $\mathbb{A}$ that computes for each graph G and positive integer k the value #IndSub$(((\Phi - H), k) \to G)$ *in time*

$$O(2^{|V(H)|} \cdot (|V(G)| + |V(H)|)^2 \cdot g(|V(H)| + k, |V(H)| + |V(G)|)).$$

The algorithm $\mathbb{B}$ queries $\mathbb{A}$ on instances with a parameter of $|V(H)| + k$.

Proof. Fix a graph G and a positive integer k. Observe that for any induced subgraph $G[X] \in \Phi - H$, we have $G[X] \uplus H \cong (G \uplus H)[X] \in \Phi$. In particular, once we extend G with H, we may use the algorithm $\mathbb{A}$ to compute the number #IndSub$((\Phi, k + |V(H)|) \to G \uplus H)$. We then use the Inclusion-Exclusion Principle to recover the number #IndSub$(((\Phi - H), k) \to G)$.

For a formal proof, write $(\mathcal{G} \uplus H)_{k+|V(H)|}$ for all size-$(k + |V(H)|)$ induced subgraphs of $G \uplus H$ that satisfy Φ, that is,

$$(\mathcal{G} \uplus H)_{k+|V(H)|} := \{(G \uplus H)[X] : |X| = |V(H)| + k\} \cap \Phi.$$

Next, for each vertex $x \in V(H)$, write $\mathcal{G}_x$ for the set of all graphs in $(\mathcal{G} \uplus H)_{k+|V(H)|}$ that contain the vertex x,[7] that is,

$$\mathcal{G}_x := \{F \in (\mathcal{G} \uplus H)_{k+|V(H)|} : x \in V(F)\}.$$

Now, we first show that we can rewrite #IndSub$(((\Phi - H)_k, k) \to G)$ as the number of graphs of $(\mathcal{G} \uplus H)_{k+|V(H)|}$ that contain all vertices of H.

[7] Technically, the operation $\triangledown$ renames the vertex x to (i, x) for some i. We may safely ignore this detail, as it is not relevant for our proof.

Claim 6.13. *We have*

$$\#\mathrm{IndSub}(((\Phi - H), k) \to G) = \Big| \bigcap_{x \in V(H)} \mathcal{G}_x \Big|.$$

Proof. We prove both inclusions separately.

First, fix a graph $G[X] \in \mathrm{IndSub}(((\Phi - H), k) \to G)$. By definition, the graph $G[X] \uplus H$ has $(k + |V(H)|)$ vertices and the property Φ. Further, as X and $V(H)$ are in disjoint subgraphs of $G[X] \uplus H$, the graph $G[X] \uplus H$ is isomorphic to the graph $(G \uplus H)[X]$; in particular, we have $G[X] \uplus H \in (\mathcal{G} \uplus H)_{k+|V(H)|}$. Finally, we observe that the graph $G[X] \uplus H$ contains every vertex $v \in V(H)$. Hence, we have

$$G[X] \uplus H \in \bigcap_{x \in V(H)} \mathcal{G}_x.$$

For the other direction, fix a graph $F := (G \uplus H)[X]$ with $F \in \bigcap_{x \in V(H)} \mathcal{G}_x$. In particular, the graph F has $k + |V(H)|$ vertices and we have $F \in \Phi$. As F contains every vertex of H, we also have that F contains H as an induced subgraph. This in turn means that we have $F \cong G[X] \uplus H$. Hence, we also have $G[X] \in \mathrm{IndSub}(((\Phi - H), k) \to G)$, which completes the proof. ◿

Ultimately, we wish to remove the graph H from our oracle calls. Toward an application of the Inclusion-Exclusion Principle, we need to understand how to partially remove H from our oracle calls. To that end, write $\overline{\mathcal{G}_x}$ for the complement of $\mathcal{G}_x$, that is, set $\overline{\mathcal{G}_x} := (\mathcal{G} \uplus H)_{k+|V(H)|} \setminus \mathcal{G}_x$. Now, we show that for any $X \subseteq V(H)$, we can rewrite $\#\mathrm{IndSub}((\Phi, |V(H)| + k) \to G \uplus (H \setminus X))$ as the number of graphs of $(\mathcal{G} \uplus H)_{k+|V(H)|}$ that contain no vertex in X.

Claim 6.14. *For every $X \subseteq V(H)$, we have*

$$\Big| \bigcap_{x \in X} \overline{\mathcal{G}_x} \Big| = \#\mathrm{IndSub}((\Phi, |V(H)| + k) \to G \uplus (H \setminus X)).$$

Proof. Unfolding the definition of $\overline{\mathcal{G}_x}$, we obtain

$$\bigcap_{x \in X} \overline{\mathcal{G}_x} = (\mathcal{G} \uplus H)_{k+|V(H)|} \setminus \bigcup_{x \in X} \mathcal{G}_x.$$

Now, $(\mathcal{G} \uplus H)_{k+|V(H)|} \setminus \bigcup_{x \in X} \mathcal{G}_x$ is the set of all induced subgraphs of $G \uplus H$ of size $k + |V(H)|$ that satisfy Φ and do not contain any vertex in X; which is precisely $\#\mathrm{IndSub}((\Phi, k + |V(H)|) \to G \uplus (H \setminus X))$. ◿

Finally we express $|\bigcap_{x \in V(H)} \mathcal{G}_x|$ in terms of $|\bigcap_{x \in X} \overline{\mathcal{G}_x}|$.

Claim 6.15. *We have*

$$\Big| \bigcap_{x \in V(H)} G_x \Big| = \sum_{X \subseteq V(H)} (-1)^{|X|+1} \Big| \bigcap_{x \in X} \overline{G_x} \Big|.$$

Proof. After unfolding the definition of $\overline{\mathcal{G}_x}$, we see that the claim is equivalent to the classical Inclusion-Exclusion Principle for set intersection; which yields the claim. ◿

Now, the combination of Claims 6.13 to 6.15 yields

$$\#\mathrm{IndSub}(((\Phi - H), k) \to G) = \sum_{X \subseteq V(H)} (-1)^{|X|+1} \#\mathrm{IndSub}((\Phi, k + |V(H)|) \to G \uplus (H \setminus X)).$$

In particular, we can compute $\#\mathrm{IndSub}(((\Phi - H), k) \to G)$ by calling the algorithm $\mathbb{A}$ with parameter $k + |V(H)|$ on $2^{|V(H)|}$ graphs G' with $|V(G')| \leq |V(G)| + |V(H)|$. Further, we can construct each graph G' in time $O((|V(G)| + |V(H)|)^2)$. In total, we obtain a running time of $O(2^{|V(H)|} \cdot (|V(G)| + |V(H)|)^2 \cdot g(k + |V(H)|, |V(G)| + |V(H)|))$; which completes the proof. ◼

6.4 Scattered Properties and Reducing to the Prime-Power Case

For a positive integer n, let us write $q(n)$ for the largest divisor of n that is a prime power; we set $q(1) := 1$. It is instructive to discuss some useful properties of the function $q(n)$.

Lemma 6.16. **(1)** *For every positive integer n, we have $n \leq q(n)^{q(n)}$.*
(2) *Write $n := p_1^{a_1} \cdots p_c^{a_c}$ for a positive integer and its corresponding prime factorization. Then, we have $\sqrt[c]{n} \leq q(n)$.*
(3) *There is a universal constant $c > 0$ such that for every positive integer n, we have $c \log(n) \leq q(n)$.*

Proof. For (1), write $n = p_1^{a_1} \cdots p_c^{a_c}$ for the prime factorization of n. We bound each prime power factor of n with $q(n)$, which yields $n \leq q(n)^c$. Finally, we use $c \leq \max(p_1, \dots, p_c) \leq q(n)$ to obtain the claim.

For (2), observe that we have $q(n) = p_i^{a_i}$ for some $i \in [c]$. Since $p_i^{a_i}$ is the largest prime power factor of n, we obtain $q(n)^c = (p_i^{a_i})^c \geq n$; which yields the claim.

For (3), write $\omega(n)$ for the number of prime divisors of n. We show $q(n) \geq c \log(n)$ for all $n \geq 26$; sufficiently decreasing c then yields the claim.

We use the following result due to Robin [Rob83, Theorem 13]. For every $n \geq 26$, we have

$$\omega(n) \leq \frac{\log n}{\log(\log(n)) - 1.1714}. \tag{3}$$

Combining Equation (3) with Claim (2), for every $n \geq 26$, we obtain

$$q(n) \geq n^{\omega(n)^{-1}} \geq n^{\frac{\log(\log(n))-1.1714}{\log(n)}} = \log(n)\, \mathrm{e}^{-1.1714};$$

which yields the claim. ∎

Let us also recall the definition of concentrated and scattered integers for Φ.

Definition 6.17. *Let Φ denote an edge-monotone graph property and write M_Φ for the set of numbers on which Φ is nontrivial. We say that Φ is* concentrated *on an integer $k \in M_\Phi$ if there is a graph H on k vertices with $\hat{\Phi}(H) \neq 0$ and H contains $K_{q(k),q(k)}$ as a subgraph.*

We say that Φ is scattered *on an integer $k \in M_\Phi$ if it is not concentrated for Φ.* ∎

Let us observe that if Φ is computable, then we can decide which case holds for a given k.

Lemma 6.18. *Let Φ denote a computable edge-monotone property. For every integer k, we can decide if Φ is trivial, concentrated, or scattered on k.*

Proof. By evaluating Φ on every k-vertex graph, we can decide if Φ is trivial. If Φ is not trivial on k, then computing the alternating enumerator of every k-vertex graph that contains $K_{q(k),q(k)}$ as a subgraph tells us if k is concentrated. ∎

In the next step, we show the key result toward our win-win approach: for an edge-monotone property Φ and any k on which Φ is nontrivial, we either win by obtaining a high-treewidth graph with nonvanishing alternating enumerator (if Φ is concentrated on k), or we win by reducing to the prime power case (if Φ is scattered on k). Formally, we prove Lemma 6.19.

Lemma 6.19. *Let Φ denote an edge-monotone graph property and write M_Φ for the set of numbers on which Φ is nontrivial. For any number $k \in M_\Phi$ on which Φ is scattered, there is a graph H on $k - q(k)$ vertices such that the property $(\Phi - H) := \{G : G \uplus H \in \Phi\}$ is edge-monotone and nontrivial on $q(k)$.*

Proof. Fix an integer $k \in M_\Phi$. We prove that either k is concentrated or the claim holds.

To that end, we let $\circlearrowleft_{q(k)}^{k/q(k)}$ act on $K_{q(k)}^{k/q(k)}$ and consider the resulting fixed points $\mathrm{FP}(\circlearrowleft_{q(k)}^{k/q(k)}, K_{q(k)}^{k/q(k)})$. As Φ is nontrivial on $k = q(k) \cdot k/q(k)$, there is a minimal integer i such that

- there is a fixed point with a level of i that does not satisfy Φ but
- all fixed points with a level of less than i satisfy Φ.

As Φ is nontrivial on k, we have $\Phi(\mathrm{IS}_{q(k)}^{q(k)/d}) = 1$ and $\Phi(K_{q(k)}^{q(k)/d}) = 0$. In particular, as $\mathrm{IS}_{q(k)}^{q(k)/d}$ is the (only) fixed point with the minimum level of 0 and as $K_{q(k)}^{q(k)/d}$ is the fixed point with the maximum level, we obtain that i satisfies

$$0 = \ell(\mathrm{IS}_{q(k)}^{q(k)/d}) < i \leq \ell(K_{q(k)}^{q(k)/d}) = \binom{k/q(k)}{2} + \#\mathbb{F}_{q(k)}^{+} \cdot k/q(k).$$

Write $F := C\langle C_{q(k)}^{A^1}, \dots, C_{q(k)}^{A^{k/q(k)}} \rangle$ for a fixed point with $\ell(F) = i$ and $\Phi(F) = 0$. Next, we distinguish two cases based on whether C contains an edge.

First, consider the case that C contains some edge. Now, Lemma 4.8 yields $\widehat{\Phi}(F) \neq 0$. Further, Lemma 6.10 yields that F has $K_{q(k),q(k)}$ as a subgraph. Thus, Φ is concentrated on k.

Second, consider the remaining case that $C = \mathrm{IS}_{k/q(k)}$. This in turn means that F is the disjoint union of $k/q(k)$ fixed points of $\mathrm{FP}(\circlearrowleft_{q(k)}, K_{q(k)})$; we write

$$F = C_{q(k)}^{A^1} \uplus \cdots \uplus C_{q(k)}^{A^{k/q(k)}}.$$

As $\ell(F) = i \geq 1$, there is at least one block $x \in [k/q(k)]$ with $A^x \neq \varnothing$. Without loss of generality, we may assume $x = 1$ (otherwise, pick the isomorphic graph with renamed blocks); and in particular $C_{q(k)}^{A^1} \neq \mathrm{IS}_{q(k)}$. Now, we set

$$H := C_{q(k)}^{A^2} \uplus \cdots \uplus C_{q(k)}^{A^{k/q(k)}}.$$

Further, we define the graph property $(\Phi - H)$ via

$$(\Phi - H) := \{G : G \uplus H \in \Phi\}.$$

As Φ is edge-monotone and computable, so is $(\Phi - H)$. Finally, we show that $(\Phi - H)$ is nontrivial on $q(k)$.

Claim 6.20. *We have* $\mathrm{IS}_{q(k)} \in (\Phi - H)$ *and* $C_{q(k)}^{A^1} \notin (\Phi - H)$.

Proof. First, the graph $\mathrm{IS}_{q(k)} \uplus H$ is isomorphic to the fixed point $F' = \mathrm{IS}_{q(k)} \uplus C_{q(k)}^{A^2} \uplus \cdots \uplus C_{q(k)}^{A^{k/q(k)}}$. Further, we have $\ell(F') < \ell(F) = i$ as $C_{q(k)}^{A^1}$ is not the empty graph. Thus, F' satisfies Φ as, by construction, all fixed points with a level less than i satisfy Φ.

Second, observe that we have that $C_{q(k)}^{A^1} \uplus H$ is isomorphic to F, hence

$$(\Phi - H)(C_{q(k)}^{A^1}) = \Phi(C_{q(k)}^{A^1} \uplus H) = \Phi(F) = 0.$$

Claim 6.20 shows that, indeed, Φ is scattered on k. This completes the proof.

Next, we show that scattered integers indeed yield a reduction to the prime power-power case. To that end, we first use the scattered part of a graph property define another graph property on prime powers.

Definition 6.21. *Let Φ denote a computable, edge-monotone graph property and write Sc_Φ for the set of all integers on which Φ is scattered. Further, write $q(Sc_\Phi) := \{q(k) : k \in Sc_\Phi\}$ for the set of all maximal prime powers corresponding to Sc_Φ. Finally, for each $m \in q(Sc_\Phi)$ write $q^{-1}(m)$ for the minimum $k \in Sc_\Phi$ with $q(k) = m$.*

Now, for each $m \in q(Sc_\Phi)$, let H_m denote the lexicographically first graph on $q^{-1}(m) - m$ vertices such that the graph property $(\Phi - H_m) := \{G : G \uplus H_m \in \Phi\}$ is nontrivial on m.

We define the scattered property Φ_{Sc} corresponding to Φ as

$$\Phi_{Sc}(G) = 1 \iff |V(G)| \in q(Sc_\Phi) \text{ and } \Phi(G \uplus H_{|V(G)|}) = 1$$

Observe that Lemma 6.19 implies that H_m is well-defined in Definition 6.21, as there is at least one such graph. Let us verify the main properties of the defined function Φ_{Sc}.

Lemma 6.22. *Let Φ denote a computable, edge-monotone graph property and write Sc_Φ for the set of all integers on which Φ is scattered. If Sc_Φ is infinite, then the scattered property Φ_{Sc} corresponding to Φ is computable, edge-monotone, and nontrivial on infinitely many prime powers.*

Proof. As in Definition 6.21, write $q(\mathrm{Sc}_\Phi) \coloneqq \{q(k) : k \in \mathrm{Sc}_\Phi\}$ for the set of all maximal prime powers corresponding to Sc_Φ. Next, for each $m \in q(\mathrm{Sc}_\Phi)$ write $q^{-1}(m)$ for the minimum $k \in \mathrm{Sc}_\Phi$ with $q(k) = m$.

Claim 6.23. *The sets Sc_Φ and $q(Sc_\Phi)$, the function q^{-1}, and the graphs H_m for every $m \in q(Sc_\Phi)$ are computable.*

Proof. Fix an integer k. As Φ is computable and by using Lemma 6.18, we can compute whether $k \in \mathrm{Sc}_\Phi$.

Next, we wish to decide $k \in q(\mathrm{Sc}_\Phi)$. To that end, we iterate through the integers i starting from k and ending with k^k. For each such j, we first check if $j \in \mathrm{Sc}_\Phi$. If this is the case, we compute $q(j)$ and check if $k = q(j)$. If indeed $k = q(j)$, we return $k \in q(\mathrm{Sc}_\Phi)$. Otherwise, if we find that for all integers $k \le j \le k^k$, we have $j \notin \mathrm{Sc}_\Phi$, we return $k \notin q(\mathrm{Sc}_\Phi)$.

This algorithm is correct, as by Lemma 6.16(1), any integer k may appear as the largest prime power only for integers that are at most k^k. Finally, said algorithm also yields the smallest value j with $q(j) = k$; which proves that the function q^{-1} is computable as well.

Recall that, for $m \in q(\mathrm{Sc}_\Phi)$, the graph H_m has $q^{-1}(m) - m$ vertices (which is a number we can compute). Let us enumerate in lexicographic order the graphs H_m on $q^{-1}(m) - m$ vertices and check if $(\Phi - H_m)$ is a nontrivial property on m vertices. As Φ is computable, this can be done by enumerating every graph on m vertices.

Finally, we observe that as Sc_Φ is infinite, so is $q(\mathrm{Sc}_\Phi)$ (as by Lemma 6.16(3), we have $q(k) \ge c\log(k)$ for all integer k). Thus Φ_{Sc} is nontrivial on infinitely many prime powers. The edge-monotonicity of Φ_{Sc} follows from its definition and from the edge-monotonicity of Φ.

Next, we show that if the scattered property Φ_{Sc} corresponding to Φ is computable, edge-monotone, and nontrivial on infinitely many prime powers, then we can reduce #IndSub(Φ) to #IndSub(Φ_{Sc}).

Corollary 6.24. *Let Φ denote a computable, edge-monotone graph property and write Φ_{Sc} for the scattered property corresponding to Φ. Then, there is a parameterized Turing reduction from #IndSub(Φ_{Sc}) to #IndSub(Φ).*

Proof. Write Sc_Φ for the set of all integers on which Φ is scattered. As in Definition 6.21, write $q(\mathrm{Sc}_\Phi) \coloneqq \{q(k) : k \in \mathrm{Sc}_\Phi\}$ for the set of all maximal prime powers corresponding to Sc_Φ. Next, for each $m \in q(\mathrm{Sc}_\Phi)$ write $q^{-1}(m)$ for the minimum $k \in \mathrm{Sc}_\Phi$ with $q(k) = m$. For $m \in q(\mathrm{Sc}_\Phi)$, define H_m as in Definition 6.21.

Fix a graph G and an integer k. We wish to compute #IndSub$((\Phi_{\mathrm{Sc}}, k) \to G)$. First, we check if k is a prime power that is contained in $q(\mathrm{Sc}_\Phi)$. If we observe $k \notin q(\mathrm{Sc}_\Phi)$, we return 0, as no graph with k vertices is in $\Phi_{\mathrm{Sc}}^{\mathcal{H}}$. Otherwise, a k-vertex graph is in Φ_{Sc} if and only if it is in $(\Phi - H_k)$. Thus, we can return #IndSub$(((\Phi - H_k), k) \to G)$ using Lemma 6.12 (which we supply with our oracle for #IndSub(Φ)).

Finally, we analyze the running time of the reduction. To that end, we assume that we have access to a #IndSub(Φ) oracle, that is, we assume that the running time $g(k, |V(G)|)$ of the oracle is constant. Now, first observe that all computations in the reduction (other than the call to Lemma 6.12) have a running time that depends only on the parameter k. Next, we observe that $|V(H_k)| = q^{-1}(k) - k$. Hence, the call to Lemma 6.12 takes time

$$O(2^{q^{-1}(k)-k} \cdot (q^{-1}(k) - k + |V(G)|)^2 \cdot g(q^{-1}(k), q^{-1}(k) - k + |V(G)|)) = O(g'(k) \cdot |V(G)|^2),$$

for some computable function g'. Lastly, we observe that the reduction of Lemma 6.12 uses only oracle calls where the parameter has size $(q^{-1}(k) - k) + k = q^{-1}(k)$.

Hence, there is a parameterized Turing reduction from #IndSub(Φ_{Sc}) to #IndSub(Φ). ◾

Remark 6.25. If the graph property Φ_{Sc} is edge-monotone and nontrivial infinitely often (that is, nontrivial on infinitely many prime powers), then Theorem 5.16 applies and the problem #IndSub(Φ_{Sc}) is #W[1]-hard (and we also obtain an ETH-based lower bound). ◾

6.5 #W[1]-hardness and Quantitative Lower Bounds for Edge-monotone Properties

We are finally ready to prove Main Theorem 1.

Main Theorem 1. *Let Φ denote a nontrivial edge-monotone graph property.*

- *The problem #IndSub(Φ) is #W[1]-hard.*
- *Further, assuming ETH, there is a universal constant $\gamma > 0$ (independent of Φ) such that for any integer $k \geq 3$ on which Φ is nontrivial, no algorithm (that reads the whole input) computes for every graph G the number #IndSub((Φ, k) $\to G$) in time $O(|V(G)|^{\gamma\sqrt{\log k}/\log\log k})$.*

Proof. Write Sc_Φ for the set of all integers on which Φ is scattered and write Co_Φ for the set of all integers on which Φ is concentrated.

For the #W[1]-hardness part, let us consider two cases: either Sc_Φ or Co_Φ is infinite. Suppose first that Sc_Φ is infinite. Then, the scattered property Φ_{Sc} corresponding to Φ is computable and nontrivial on infinitely many prime powers (Lemma 6.22). Hence, #IndSub(Φ_{Sc} is #W[1]-hard (Theorem 5.16) and can be reduced to #IndSub(Φ) (Lemma 6.12), showing that #IndSub(Φ) is also #W[1]-hard.

Assume now that Co_Φ is infinite. Let $\mathcal{H}$ contain every graph H whose alternating enumerator is nonvanishing for Φ and contains $K_{q(|V(H)|),q(|V(H)|)}$ as subgraph. The set $\mathcal{H}$ is clearly infinite and computable. Observe that every $H \in \mathcal{H}$ has treewidth at least $q(k)$.[8] In particular, by Lemma 6.16(3), as the family $\mathcal{H}$ is infinite, it has unbounded treewidth. Thus Lemma A.8 yields #W[1]-hardness for #IndSub(Φ).

We turn to the ETH-based lower bound next. To that end, fix a k for which Φ is nontrivial. Now, Φ is either scattered or concentrated on k. We provide a proof for both cases.

First, we consider the case that Φ is scattered on k. Define $m := q(k) \in q(\mathrm{Sc}_\Phi)$. Let H_m denote the graph defined in Definition 6.21. Then, $(\Phi - H_m) := \{G : G \uplus H_m \in \Phi\}$ is edge-monotone and nontrivial on $m = q(k) \geq c\log(k) \geq 3$ vertices (for some constant c and sufficiently large k).

Now, by Theorem 5.16 and assuming ETH, there is an α such that there is no algorithm that for each graph G computes the number #IndSub((($\Phi - H_m$), m) $\to G$) in time $O(|V(G)|^{\alpha\sqrt{q(k)}/\log(q(k))})$. Using the reduction of Lemma 6.12 (which has a quadratic overhead), we obtain that there is no algorithm that computes for every graph G the number #IndSub((Φ, k) $\to G$) in time

$$O(|V(G)|^{(\alpha\sqrt{q(k)}/\log(q(k)))-2}),$$

We conclude this case by proving the following inequalities.

[8] Consult for instance [BK11, Corollary 9 and Lemma 4] for a proof of this folklore fact.

Claim 6.26. *Choosing $\gamma' := \alpha\sqrt{c}/2$, for any sufficiently large k, we have*

$$\gamma' \frac{\sqrt{\log k}}{\log\log k} \leq \gamma' \frac{\sqrt{q(k)/c}}{\log(q(k)/c)} \leq \gamma' \frac{\sqrt{q(k)/c}}{\log(q(k))} \leq \alpha \frac{\sqrt{q(k)}}{\log(q(k))} - 2.$$

Proof. **(1)** For the first inequality we consider the function $h_1 \colon \mathbb{R}_{>1} \to \mathbb{R}; x \mapsto \sqrt{x}/\log(x)$. Observe that h_1 is monotonically increasing for $x \geq \mathrm{e}^2$.[9] Finally, we choose k large enough to satisfy $\mathrm{e}^2 \leq \log k \leq q(k)/c$.

(2) The second inequality is immediate.

(3) The last inequality is equivalent to $4 \leq \alpha\sqrt{q(k)}/\log(q(k))$, which holds for sufficiently large $q(k)$ since h_1 is unbounded and monotonically increasing.

Now, by Claim 6.26, for a $k \geq N_1$ (where N_1 is a sufficiently large constant) and assuming ETH, there is no algorithm that for each graph G computes the number #IndSub$((\Phi, k) \to G)$ in time

$$O(|V(G)|^{(\alpha\sqrt{q(k)}/\log(q(k)))-2}) \supseteq O(|V(G)|^{(\gamma'\sqrt{q(k)}/\log(q(k)))}).$$

Next, we consider the case that Φ is concentrated on k. By definition, there exists a graph H_k on k vertices with a nonvanishing alternating enumerator that contains $K_{q(k),q(k)}$ as a subgraph. As before, this implies that H_k has a treewidth of at least $q(k)$.

Now, from Lemma 6.16(3), for each integer N_2, we obtain that for some constant c and for $k \geq \beta$ (for some constant $\beta := \beta(c, N_2)$ that depends on c and N_2), we have

$$\mathrm{tw}(H_k) \geq q(k) \geq c\log(k) \geq N_2.$$

Now, by Lemma A.8, assuming ETH, and choosing $N_2 := 3$, there is no algorithm that for each graph G computes the number #IndSub$((\Phi, k) \to G)$ in time

$$O(|V(G)|^{\alpha_{\textsc{IndSub}}\,\mathrm{tw}(H_k)/\log \mathrm{tw}(H_k)}) \supseteq O(|V(G)|^{\alpha_{\textsc{IndSub}} c\log(k)/\log(c\log(k))}) \supseteq O(|V(G)|^{\alpha_{\textsc{IndSub}} c\log k/\log\log k}).$$

For the first step, we use that the function $h_2 \colon \mathbb{R}_{>1} \to \mathbb{R}, x \mapsto x/\log(x)$ is monotonically increasing for $x > \mathrm{e}$. For the second step, we use $0 < c < 1$ (which we can assume without loss of generality). Finally, to obtain the claim also for all values k that are less than $N_0 := \max(N_1, \beta(c, N_2))$, we choose $\gamma := \min(1/\sqrt{\log(N_0)}, \gamma', c\alpha_{\textsc{IndSub}})$. Observe that now, for $k < N_0$, we obtain that

$$O(|V(G)|^{\gamma\sqrt{\log k}/\log\log k}) = o(|V(G)|).$$

Now, such a running time is unconditionally unachievable for any algorithm that reads the whole input. This completes the proof.

7 Main Result 2: Tight Bounds for Edge-monotone Properties on Prime Powers

In this section, we prove Main Theorem 2. While we build on our structural understanding of fixed points from Section 4, the proof is independent of Sections 5 and 6.

Main Theorem 2. *For each prime p, there is a constant $\gamma_p > 0$ such that for each integer m with $p^m \geq 3$ and each edge-monotone graph property Φ that is nontrivial on p^m, no algorithm (that reads the whole input) computes for every graph G the number #IndSub$((\Phi, k) \to G)$ in time $O(|V(G)|^{\gamma_p p^m})$, unless ETH fails.*

[9] First, one may verify that the derivative of h_1 is $h_1'(x) = (\log(x) - 2)/(2\sqrt{x}\log^2(x))$. Next, one verifies that $h_1'(x) \geq 0$ for all $x \geq \mathrm{e}^2$. Thus, h_1 is monotonically increasing after e^2.

Recall that by Section 5, the rotation subgroup $\circlearrowleft_{p^m} \subseteq \mathrm{Aut}(K_{p^m})$ gives rise to the difference graphs over $\mathbb{F}_{p^m}$ as fixed points. In this section, we consider a much larger subgroup of $\mathrm{Aut}(K_{p^m})$, namely the Sylow p-subgroup. Thereby, we obtain fixed points of a different type, the lexicographic product of difference graphs over $\mathbb{F}_p$. To analyze such fixed points, for the rest of the section, we view the vertex set of K_{p^m} as $[0\,..\,p)^m$.

7.1 The Fixed Points of Sylow Groups on K_{p^m}

We intend to construct special subgroups of $\mathrm{Aut}(K_{p^m})$. To that end, we first define special types of bijections on $[0\,..\,p)^m$.

Definition 7.1. *Consider a prime p and a positive integer m. For each $j \in [0\,..\,m)$, write φ_i for a function $[0\,..\,p)^j \to [0\,..\,p)$ and set $\varphi := (\varphi_0, \dots, \varphi_{m-1})$. We define the function $\overline{\varphi} : [0\,..\,p)^m \to [0\,..\,p)^m$ via*[10]

$$\overline{\varphi}(x_1, \dots, x_m) := (x_1 + \varphi_0, x_2 + \varphi_1(x_1), x_3 + \varphi_2(x_1, x_2), \dots, x_m + \varphi_{m-1}(x_1, \dots, x_{m-1})),$$

where all computations are done modulo p.

We write $\overline{p}^m$ for the set of all functions $\overline{\varphi} : [0\,..\,p)^m \to [0\,..\,p)^m$ that are obtained in this fashion, that is,

$$\overline{p}^m := \{\overline{\varphi} : \varphi_i \in [0\,..\,p)^j \to [0\,..\,p)\}.$$

Lemma 7.2. *For each $j \in [0\,..\,m)$, write φ_i for a function $[0\,..\,p)^j \to [0\,..\,p)$ and set $\overline{\varphi} := (\varphi_0, \dots, \varphi_{m-1})$. Then, the function $\overline{\varphi}$ is a bijection on $[0\,..\,p)^m$.*

Proof. Consider two different tuples $x := (x_1, \dots, x_m)$ and $x' := (x'_1, \dots, x'_m)$ and write i for the minimal position with $x_i \neq x'_i$. We claim that $\overline{\varphi}(x)_i \neq \overline{\varphi}(x')_i$. To that end, observe that φ_{i-1} depends only on the values $x_1 = x'_1, \dots, x_{i-1} = x'_{i-1}$. Thus, we have $\varphi_i(x_1, \dots, x_{i-1}) = \varphi_i(x'_1, \dots, x'_{i-1})$. Hence, we have

$$\overline{\varphi}(x)_i = x_i + \varphi_i(x_1, \dots, x_{i-1}) \neq x'_i + \varphi_i(x_1, \dots, x_{i-1}) = x'_i + \varphi_i(x'_1, \dots, x'_{i-1}) = \overline{\varphi}(x')_i.$$

Hence, $\overline{\varphi}$ is injective, and as a mapping between equal-sized sets thus also bijective.

Recall that $[0\,..\,p)^m$ is also the vertex set of K_{p^m}. Thus, Lemma 7.2 implies $\overline{\varphi} \in \mathrm{Aut}(K_{p^m}) \cong \mathfrak{S}_{p^m}$.

Remark 7.3. Let us discuss and interpret bijections $\overline{\varphi} \in \overline{p}^m$. To that end, consider an ordered, complete p-tree T_{p^m} of height m. Label the root with the empty tuple (). Further, for each vertex with label $(v_1, \dots, v_j)$, label its i-th child with $(v_1, \dots, v_j, i)$ (where we number children 0-indexed). Observe that T_{p^m} has p^m leaves that correspond to the elements of $[0\,..\,p)^m$. Consult Figure 5a for a visualization of an example.

Now, $\overline{\varphi}$ acts on T_{p^m} in the following way. Starting with the root, each level i of T_{p^m} is rotated by $\varphi_i(v_1, \dots, v_i)$ nodes, that is, the j-th child of a vertex $(v_1, \dots, v_i)$ becomes the $((i+\varphi_i(v_1, \dots, v_i)) \bmod p)$-th child of $(v_1, \dots, v_i)$. Consult Figure 5b for a visualization of an example.

Next, we show that $\overline{p}^m$ forms a group with the function composition as the group operation.

Lemma 7.4. *The pair $\mathrm{Syl}_{p^m} := (\overline{p}^m, \circ)$ forms a p-group.*

Proof. We first show that Syl_{p^m} defines a group.

[10] We write φ_0 for $\varphi_0(())$ since φ_0 is a function that is defined on a single element

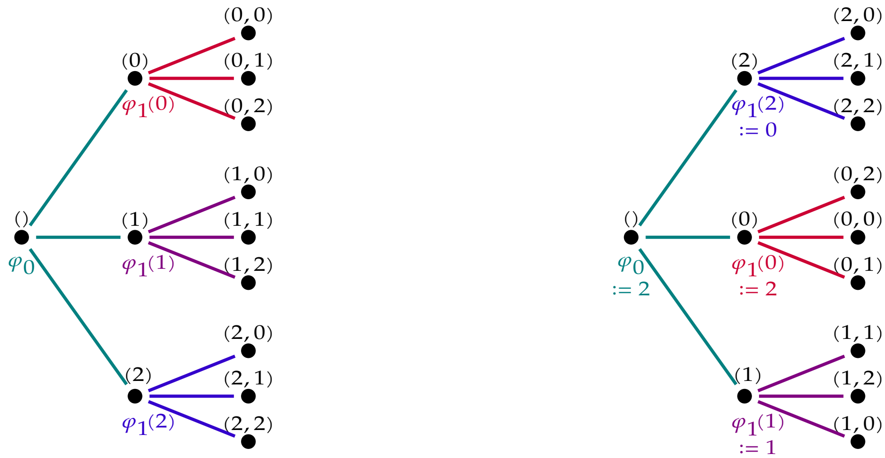

(a) The tree T_{3^2} from Remark 7.3. Each vertex is label with a label that is depicted above the corresponding vertex. Further, we depict below each vertex v the function (value) in $\overline{\varphi}$ that is responsible for rotating the subtree rooted at v.

(b) The image $\overline{\varphi}T_{3^2}$, where $\varphi_0 := 2$ and $\varphi_1 := \{0 \mapsto 2, 1 \mapsto 1, 2 \mapsto 0\}$. Reading of the labels of the leaves from top to bottom, we obtain the corresponding permutation of $[0\,..\,p)^m$.

Figure 5 Definition 7.1 and Remark 7.3 visualized.

Claim 7.5. *The composition of elements of $\overline{p}^m$ is an element of $\overline{p}^m$.*

Proof. Fix $\overline{\varphi}, \overline{\psi} \in \overline{p}^m$. Expanding the definition yields

$$\begin{aligned}\overline{\psi}(\overline{\varphi}(x_1, \dots, x_m)) = (&x_1 + \varphi_0 + \psi_0, x_2 + \varphi_1(x_1) + \psi_1((\varphi_0)(x_1)), \dots, \\ &x_m + \varphi_{m-1}(x_1, \dots, x_{m-1}) + \psi_m((\varphi_0, \dots, \varphi_{m-1})(x_1, \dots, x_{m-1}))) \\ = \overline{\lambda}&(x_1, \dots, x_m),\end{aligned}$$

where $\lambda_0 := \varphi_0 + \psi_0$ and $\lambda_j(x_1, \dots x_j) := \varphi_j(x_1, \dots, x_j) + \psi_j((\varphi_0, \dots, \varphi_{j-1})(x_1, \dots, x_{j-1}))$. Thus, indeed $\overline{\lambda} \in \mathrm{Syl}_{p^m}$, which yields the claim. ⌟

Claim 7.6. *Every element of $\overline{\varphi} \in \overline{p}^m$ has an inverse $\overline{\varphi}^{-1} \in \overline{p}^m$.*

Proof. We define $\overline{\varphi}^{-1}$ via $\varphi_0^{-1} := -\varphi_0$ and

$$\varphi_j^{-1}(x_1, \dots, x_j) := -\varphi_j((\varphi_0^{-1}, \dots, \varphi_{j-1}^{-1})(x_1, \dots, x_{j-1})) \text{ for } j \in [1\,..\,m).$$

An induction on $j \in [0\,..\,m)$ readily yields $\overline{\varphi}^{-1} \circ \overline{\varphi} = \overline{\varphi} \circ \overline{\varphi}^{-1} = \mathrm{id}$; which completes the proof. ⌟

By Claims 7.5 and 7.6, Syl_{p^m} is indeed a group.

Finally, we show that Syl_{p^m} is a p-group. To that end, observe that each m-tuple $(\varphi_0, \dots, \varphi_{m-1})$ with $\varphi_j \colon [0\,..\,p)^j \to [0\,..\,p)$ defines a different group element $\overline{\varphi}$. The number of functions φ_j from $[0\,..\,p)^j$ to $[0\,..\,p)$ is equal to p^{p^j}, thus we obtain $|\mathrm{Syl}_{p^m}| = |\overline{p}^m| = p^1 \cdot p^{p^1} \cdot \dots \cdot p^{p^{m-1}}$; which completes the proof. ∎

Remark 7.7. Let us briefly discuss how Syl_{p^m} relates to Sylow p-subgroups that appear in the literature. Classically, a *Sylow p-subgroup* of Γ is a p-subgroup that is not a proper subgroup of any other p-subgroup of Γ. Due to the Sylow Theorems (consult for instance [Rot94, Theorem 4.12]), there is only one Sylow p-subgroup in $\mathfrak{S}_{p^m}$ (up to isomorphism). Usually, the Sylow p-subgroup is constructed by taking m

copies of the cyclic group $\mathbb{Z}_p$ and using the wreath product $\mathbb{Z}_p \wr \cdots \wr \mathbb{Z}_p$ [Rot94, Theorem 7.27]. Indeed, one may prove that the group Syl_{p^m} defined via Lemma 7.4 is isomorphic to the Sylow p-subgroup of $\mathfrak{S}_{p^m}$, that is, we have $\mathrm{Syl}_{p^m} \cong \mathbb{Z}_p \wr \cdots \wr \mathbb{Z}_p$. ◾

Next, we analyze the fixed point structure of $\mathrm{FP}(\mathrm{Syl}_{p^m}, K_{p^m})$. To that end, we introduce the lexicographic product of graphs.

Definition 7.8. *For graphs $G_1, \dots, G_m$, we define their* lexicographic product $G_1 \circ \cdots \circ G_m$ *via*

$$\begin{aligned} V(G_1 \circ \cdots \circ G_m) &:= V(G_1) \times \cdots \times V(G_m) \quad \text{and} \\ E(G_1 \circ \cdots \circ G_m) &:= \{\{(u_1, \dots, u_m), (v_1, \dots, v_m)\} \\ &\qquad : \text{there is an } i \in [m] \text{ with } u_j = v_j \text{ for all } j < i \text{ and } \{u_i, v_i\} \in E(G_i)\}. \end{aligned}$$ ◾

As an easy example, observe that we have $G_1 \circ G_2 \cong G_1\langle G_2, \dots, G_2\rangle$ and $G_1 \circ G_2 \circ G_3 \cong G_1\langle G_2 \circ G_3, \dots, G_2 \circ G_3\rangle$.

Remark 7.9. Another name for the lexicographic product is *wreath product of graphs* [ABD+18, DM09]. The name wreath product emphasizes the close connection between the group theoretical wreath products of automorphism subgroups $A \subseteq \mathrm{Aut}(G)$ and $B \subseteq \mathrm{Aut}(H)$, and the wreath products of G and H. Further, there is a close connection between the group-theoretical wreath product and the wreath product of groups when considering fixed points. However, as our proofs do not use the terminology of wreath product of groups, we chose the more common name *lexicographic product*. ◾

Let us take a closer look at the lexicographic product of difference graphs.

Lemma 7.10. *Write p for a prime and m for a positive integer. Further, for each $i \in [m]$, write $A_i \subseteq \mathbb{F}_p^+$ for a subset. Then, $E(C_p^{A_1} \circ \cdots \circ C_p^{A_m})$ is equal to*

$$\{\{(a_1, \dots a_m), (b_1, \dots b_m)\} : \exists i \in [m] \colon (\forall j < i \colon a_j = b_j) \text{ and } a_i - b_i \in A_i \cup (-A_i)\}$$

Proof. Unfolding the definition of the lexicographic product, we observe that $\{(a_1, \dots, a_m), (b_1, \dots, b_m)\}$ is an edge of $C_p^{A_1} \circ \dots \circ C_p^{A_m}$ if we can find a i such that $a_j = b_j$ for $j < i$ and $\{a_i, b_i\} \in C_p^{A_i}$. Finally, we recall that $\{a_i, b_i\} \in C_p^{A_i}$ is equivalent to $a_i - b_i \in A_i \cup (-A_i)$. ◾

Next, we use the lexicographic product of difference graphs to understand the fixed points $\mathrm{FP}(\mathrm{Syl}_{p^m}, K_{p^m})$.

Lemma 7.11. *For any prime p and any positive integer m, we have*

$$\mathrm{FP}(\mathrm{Syl}_{p^m}, K_{p^m}) = \{C_p^{A_1} \circ \cdots \circ C_p^{A_m} : A_i \subseteq \mathbb{F}_p^+\}.$$

Proof. Write $C := \{C_p^{A_1} \circ \cdots \circ C_p^{A_m} : A_i \subseteq \mathbb{F}_p^+\}$.

$C \subseteq \mathrm{FP}(\mathrm{Syl}_{p^m}, K_{p^m})$. Fix sets A_i and the corresponding graph $H := C_p^{A_1} \circ \cdots \circ C_p^{A_m}$. We prove that H is a fixed point of Syl_{p^m} by verifying that $\overline{\varphi}(H) = H$ holds for every function $\overline{\varphi} \in \mathrm{Syl}_{p^m}$.

To that end, fix a $\overline{\varphi} \in \mathrm{Syl}_{p^m}$. Further, fix an edge $\{(a_1, \dots, a_m), (b_1, \dots, b_m)\} \in E(H)$. We show that $\{\overline{\varphi}(a_1, \dots, a_m), \overline{\varphi}(b_1, \dots, b_m)\} \in E(H)$. From Lemma 7.10 we obtain an $i \in [m]$ such that $a_j = b_j$ for all $j < i$ and $a_i - b_i \in A_i \cup (-A_i)$. Without loss of generality, we assume that $a_i - b_i \in A_i$ (otherwise switch the roles of a and b). Observe that we need only the values $a_1, \dots, a_j$ to compute the j-th coordinate of

$\overline{\varphi}(a)$ (the same is true for the j-th coordinate of $\overline{\varphi}(b)$). This implies $\overline{\varphi}(a)_j = \overline{\varphi}(b)_j$ for all $j < i$. Further, we obtain

$$\begin{aligned}\overline{\varphi}(a)_i - \overline{\varphi}(b)_i &= a_i + (\varphi_0, \dots, \varphi_{i-1})(a_1, \dots, a_{i-1}) - (b_i + (\varphi_0, \dots, \varphi_{i-1})(b_1, \dots, b_{i-1}))\\ &= a_i + (\varphi_0, \dots, \varphi_{i-1})(a_1, \dots, a_{i-1}) - (b_i + (\varphi_0, \dots, \varphi_{i-1})(a_1, \dots, a_{i-1}))\\ &= a_i - b_i \in A_i,\end{aligned}$$

which implies that $\{\overline{\varphi}(a_1, \dots, a_m), \overline{\varphi}(b_1, \dots, b_m)\}$ is an edge in H. Thus, $\overline{\varphi}$ maps edges of H to edges of H and is therefore an automorphism.

$\mathrm{FP}(\mathrm{Syl}_{p^m}, K_{p^m}) \subseteq C$. We prove the contrapositive, that is, we prove that each graph H with vertex set $V(K_{p^m})$ that is not in C is also not in $\mathrm{FP}(\mathrm{Syl}_{p^m}, K_{p^m})$.

To that end, we construct a $\overline{\varphi} \in \mathrm{Syl}_{p^m}$ with $\overline{\varphi}(H) \neq H$. If $H \notin C$, then, due to Lemma 7.10, there are an $i \in [m]$ and an $x \neq 0$ such that

- $u := \{(a_1, \dots, a_{i-1}, a_i, b_{i+1}, \dots, b_m), (a_1, \dots, a_{i-1}, a_i + x, c_{i+1}, \dots, c_m)\}$ is an edge in H but
- $v := \{(a_1, \dots, a_{i-1}, d_i, \beta_{i+1}, \dots, \beta_m), (a_1, \dots, a_{i-1}, d_i + x, \gamma_{i+1}, \dots, \gamma_m)\}$ is not an edge in H.

Next, we consider the group element $\overline{\varphi} := (\varphi_0, \dots \varphi_{m-1})$ with

$$\varphi_{m-1}(x_1, \dots x_{m-1}) := \begin{cases} \beta_m - b_m & \text{if } (x_1, \dots x_{m-1}) = (a_1, \dots, a_i, b_{i+1}, \dots, b_{m-1}) \\ \gamma_m - c_m & \text{if } (x_1, \dots x_{m-1}) = (a_1, \dots, a_i + x, c_{i+1}, \dots, c_{m-1}) \\ 0 & \text{otherwise} \end{cases}$$

and $\varphi_j := \mathrm{id}$ for $j \neq m-1$. We obtain

$$\overline{\varphi}(u) = \{(a_1, \dots, a_i, b_{i+1}, \dots, b_{m-1}, \beta_m), (a_1, \dots, a_i + x, c_{i+1}, \dots, c_{m-1}, \gamma_m).$$

In particular, the last coordinate of the edge $\overline{\varphi}(u)$ is equal to the last coordinate of the edge v. By iterating this construction, we obtain an $\overline{\psi} \in \mathrm{Syl}_{p^m}$ with

$$\overline{\psi}(u) = \{(a_1, \dots, a_i, \beta_{i+1}, \dots, \beta_m), (a_1, \dots, a_i + x, \gamma_{i+1}, \dots, \gamma_m)\}.$$

Lastly, we define $\tilde{\psi} := (\tilde{\psi}_0, \dots, \tilde{\psi}_{m-1})$ with

$$\tilde{\psi}_i(x_1, \dots x_{i-1}) := \begin{cases} d_i - a_i & \text{if } (x_1, \dots x_{i-1}) = (a_1, \dots, a_{i-1}) \\ 0 & \text{otherwise} \end{cases}$$

and $\tilde{\psi}_j := \mathrm{id}$ for $j \neq i$. Observe that $\tilde{\psi}(\overline{\psi}(u)) = v$. Thus there is an element in Syl_{p^m} that maps the edge u to the non-edge v which shows that H is not in $\mathrm{FP}(\mathrm{Syl}_{p^m}, K_{p^m})$. ◼

Lastly, we compute the level of fixed points in $\mathrm{FP}(\mathrm{Syl}_{p^m}, K_{p^m})$.

Lemma 7.12. *For any prime p and any positive integer m, the level of $C_p^{A_1} \circ \dots \circ C_p^{A_m} \in \mathrm{FP}(\mathrm{Syl}_{p^m}, K_{p^m})$ is*

$$\ell(C_p^{A_1} \circ \dots \circ C_p^{A_m}) = \sum_{i=1}^{m} |A_i|.$$

Proof. For an $x \in \mathbb{F}_p^+$, consider the fixed point $F := C_p^{A_1} \circ \dots \circ C_p^{A_j} \circ \dots \circ C_p^{A_m}$, where $A_j := \{x\}$ and $A_i := \emptyset$ otherwise.

Claim 7.13. *We have $\ell(F) = 1$.*

Proof. Recall from Definition 4.5 that the level of F is defined as the size of the orbit factorization of F. As F has edges, the orbit factorization of F is not empty. Thus, the level of F is at least 1.

Next, suppose that $\ell(F) > 1$. Then, the orbit factorization of F contains at least two orbits. This implies that there is a proper sub-point $F' \subsetneq F$ that contains at least one orbit and is therefore not the empty graph. However, if F' is a proper sub-point of F then

$$F' = C_p^{B_1} \circ \cdots \circ C_p^{B_j} \circ \cdots \circ C_p^{B_m},$$

with $\varnothing \subsetneq B_j \subsetneq A_j = \{x\}$ and $B_i \subseteq A_i = \varnothing$; which is a contradiction. Hence, we have $\ell(F) \leq 1$ which yields the claim. ⊿

Next, observe that for all $x \neq y \in \mathbb{F}_p^+$, the conditions

$$a_j - b_j \in \{x, -x\} \qquad a_j - b_j \in \{y, -y\}$$

are mutually exclusive. Hence, for a fixed point $F' \coloneqq C_p^\varnothing \circ \cdots \circ C_p^\varnothing \circ C_p^{A_j} \circ C_p^\varnothing \circ \cdots \circ C_p^\varnothing$ with $A_j \subseteq \mathbb{F}_{p^m}^+$, we obtain the following disjoint union

$$\begin{aligned} E(F) &= \{\{(a_1, \dots a_m), (b_1, \dots b_m)\} : (\forall i < j \colon a_i = b_i) \text{ and } a_j - b_j \in A_j \cup (-A_j)\} \\ &= \bigcup_{x \in A_j} \{\{(a_1, \dots a_m), (b_1, \dots b_m)\} : \forall i < j \colon a_i = b_i) \text{ and } a_j - b_j \in \{x, -x\}\} \\ &= \bigcup_{x \in A_j} E(C_p^\varnothing \circ \cdots \circ C_p^\varnothing \circ C_p^{\{x\}} \circ C_p^\varnothing \circ \cdots \circ C_p^\varnothing). \end{aligned}$$

Hence, we have $\ell(F') = |A_j|$.

Lastly, for a general fixed point we obtain

$$E(C_p^{A_1} \circ \cdots \circ C_p^{A_j} \circ \cdots \circ C_p^{A_m}) = \bigcup_{j \in [m]} E(C_p^\varnothing \circ \cdots \circ C_p^\varnothing \circ C_p^{A_j} \circ C_p^\varnothing \circ \cdots \circ C_p^\varnothing).$$

Thus, we have $\ell(C_p^{A_1} \circ \cdots \circ C_p^{A_j} \circ \cdots \circ C_p^{A_m}) = \sum_{i=1}^m |A_i|$; which completes the proof. ∎

Let us briefly discuss why the fixed points $\mathrm{FP}(\mathrm{Syl}_{p^m}, K_{p^m})$ are useful to us. Recall that our goal is the find nonvanishing fixed points with large treewidth. Achieving this goal is easier if many fixed points have a large treewidth. Fixed points have a large treewidth if they consist of orbits that have a large treewidth. Typically, orbits (and graphs in general) have a large treewidth if they have many edges. Since each orbit has the form $\mathrm{Syl}_{p^m} \cdot \{i, j\} \coloneqq \{\{\overline{\varphi}(i), \overline{\varphi}(j)\} : \overline{\varphi} \in \mathrm{Syl}_{p^m}\}$, for some edge $\{i, j\} \in E(K_{p^m})$, we observe that larger groups lead to larger orbits. Now, we additionally need that our group is a p-group. As the p-Sylow group Syl_{p^m} is by definition the largest p group in $\mathfrak{S}_{p^m}$, the choice of the group Syl_{p^m} is in some sense optimal.

More concretely, we may compare the number of orbits $E(K_{p^m})/\circlearrowright_{p^m}$ and the number of orbits $E(K_{p^m})/\mathrm{Syl}_{p^m}$. As orbits form a partition of the edge set of K_{p^m}, having fewer orbits leads to orbits that have more edges on average. Now, recall that by Lemma 5.2, $E(K_{p^m})/\circlearrowright_{p^m}$ contains $|\mathbb{F}_{p^m}^+|$ orbits, which is equal to $(p^m - 1)/2$ if p is odd. Next, Lemma 7.12 shows that $E(K_{p^m})/\mathrm{Syl}_{p^m}$ contains $m \cdot |\mathbb{F}_p^+|$ orbits, which is equal to $m \cdot (p-1)/2$ if p is odd. This is much smaller than $(p^m - 1)/2$.

7.2 Prime Powers Contain Large Bicliques

Next, we prove Theorem 7.18 by analyzing the fixed points $\mathrm{FP}(\mathrm{Syl}_{p^m}, K_{p^m})$. To that end, one important observation is that a fixed point $C_p^{A_1} \circ \ldots \circ C_p^{A_m}$ has a large biclique if there is a small position i with $A_i \neq \emptyset$. This implies that there is an edge between all vertices $(a_1, \ldots, a_{i-1}, a_i, a_{i+1}, \ldots a_m)$ and $(a_1, \ldots, a_{i-1}, b_i, b_{i+1}, \ldots, b_m)$ as long as $a_i - b_i \in A_i \cup (-A_i)$. As we are free in our choice of $a_{i+1}, \ldots, a_m, b_{i+1}, \ldots, b_m \in [0\,..\,p)$, this leads to a large biclique. Thus, it is useful to keep track of this value i.

Definition 7.14. *Write $H := C_p^{A_1} \circ \ldots \circ C_p^{A_m}$ for a fixed point of $\mathrm{FP}(\mathrm{Syl}_{p^m}, K_{p^m})$. The* empty-prefix *of H is the smallest index i with $A_i \neq \emptyset$, minus one; we write $\varepsilon(A_1, \ldots, A_m) := i - 1$.*

As discussed earlier, we use the empty-prefix of a fixed point to find large a biclique as a subgraph.

Lemma 7.15. *Let p denote a prime number and let m denote a positive integer. For each $i \in [m]$, let $A_i \subseteq \mathbb{F}_p^+$ denote a subset and set $A := (A_1, \ldots, A_m)$. Then, $C_p^{A_1} \circ \cdots \circ C_p^{A_m}$ contains $K_{p^{m-1-\varepsilon(A)}, p^{m-1-\varepsilon(A)}}$ as a subgraph.*

Proof. Set $t := \varepsilon(A) + 1$. We have $A_t \neq \emptyset$. Hence, there are $x \in A_t$ and $\alpha, \beta \in [0\,..\,p)$ with $\alpha - \beta = x$. Now, we obtain that all vertices of the form $(a_1, \ldots a_{t-1}, \alpha, b_{t+1}, \ldots, b_m)$ are connected to the vertices $(a_1, \ldots a_{t-1}, \beta, c_{t+1}, \ldots, c_m)$ since they coincide on the first $t-1$ elements and the t-th element differs by x. In particular, as we can freely choose $b_{t+1}, \ldots, b_m$ and $c_{t+1}, \ldots, c_m$, we obtain a complete bipartite subgraph where one side contains all vertices of the from $(a_1, \ldots a_{t-1}, \alpha, b_{t+1}, \ldots, b_m)$ and the other side contains all vertices of the from $(a_1, \ldots a_{t-1}, \beta, c_{t+1}, \ldots, c_m)$. Both sides contain $p^{m-(t+1)+1} = p^{m-t} = p^{m-1-\varepsilon(A)}$ vertices. Thus, $K_{p^{m-1-\varepsilon(A)}, p^{m-1-\varepsilon(A)}}$ is a subgraph of $C_p^{A_1} \circ \cdots \circ C_p^{A_m}$.

Imagine that our graph property is nontrivial on p^m for $m \geq 2$. If we find a fixed point $H = C_p^{A_1} \circ \cdots \circ C_p^{A_m}$ with a minimal empty-prefix of $\varepsilon(A_1, \ldots, A_m) = 0$, then we know that the treewidth of H is at least p^{m-1}. Thus, our goal is to find a fixed point H with $\varepsilon(A_1, \ldots, A_m) = 0$ and $\widehat{\Phi}(H) \not\equiv_p 0$.

To this end, we prove that a fixed point with a high empty-prefix is always isomorphic to an edge-subgraph of a fixed point with a low empty-prefix. This in turn allows us to consider only fixed points with low empty-prefix. To be more precise, we show that each fixed point $C_p^{\emptyset} \circ \cdots C_p^{\emptyset} \circ C_p^{A_1} \circ \cdots \circ C_p^{A_{m-j}}$ is isomorphic to an edge-subgraph of $C_p^{A_1} \circ \cdots \circ C_p^{A_m}$.

For a formal proof, we define the forward revolution.

Definition 7.16. *For a prime p and a positive integer m, we define the* forward revolution *of K_{p^m} as*

$$\overrightarrow{\varphi}_{p^m} : V(K_{p^m}) \to V(K_{p^m}); \quad (a_1 \ldots, a_m) \mapsto (a_m, a_1, \ldots, a_{m-1}).$$

Using Definition 7.16 allows us to prove the following lemma.

Lemma 7.17. *Let p denote a prime number and let m denote a positive integer. For each $i \in [m]$, let $A_i \subseteq \mathbb{F}_p^+$ denote a subset. Then, for all $j \in [m]$, the graph $C_p^{\emptyset} \circ \cdots C_p^{\emptyset} \circ C_p^{A_1} \circ \cdots \circ C_p^{A_{m-j}}$ is isomorphic to an edge-subgraph of $C_p^{A_1} \circ \cdots \circ C_p^{A_m}$.*

Proof. We only show that $\widetilde{H} := C_p^{\emptyset} \circ C_p^{A_1} \circ \cdots \circ C_p^{A_{m-1}}$ is isomorphic to an edge-subgraph of $H := C_p^{A_1} \circ \cdots \circ C_p^{A_m}$. Iterating then yields the general result.

We show that $\widetilde{H}$ is an edge-subgraph of $\overrightarrow{\varphi}_{p^m}(H)$ by showing that $\{a, b\} \in E(\widetilde{H})$ implies $\{a, b\} \in E(\overrightarrow{\varphi}_{p^m}(H))$. This proves the claim as $\overrightarrow{\varphi}_{p^m}$ is a bijective function and hence, $\overrightarrow{\varphi}_{p^m}(H)$ is isomorphic to H.

For all $\{a, b\} \in E(\widetilde{H})$, we obtain that there is an $i \in \{2, \ldots, m\}$ with $a_j = b_j$ for all $j < i$ and $a_i - b_i \in A_{i-1} \cup (-A_{i-1})$. Without loss of generality, we assume that $a_i - b_i \in A_{i-1}$ (otherwise we switch the roles of a and b).

Next, we show that $\{\widetilde{a}, \widetilde{b}\} \in E(H)$, where $\widetilde{a} = (a_2, \ldots, a_m, a_1)$ and $\widetilde{b} = (b_2, \ldots, b_m, b_1)$. Observe that for $s := i - 1 \in [m]$, we obtain $\widetilde{a}_j = a_{j+1} = b_{j+1} = \widetilde{b}_j$ for all $j + 1 < i$ which is equivalent to $j < s$. Further, we obtain $\widetilde{a}_s - \widetilde{b}_s \equiv a_i - b_i \in A_s$ which proves that $\{\widetilde{a}, \widetilde{b}\} \in E(H)$. This implies $\{\overrightarrow{\varphi}_{p^m}(\widetilde{a}), \overrightarrow{\varphi}_{p^m}(\widetilde{b})\} \in E(\overrightarrow{\varphi}_{p^m}(H))$. Finally, we observe that $\overrightarrow{\varphi}_{p^m}(\widetilde{a}) = a$ and $\overrightarrow{\varphi}_{p^m}(\widetilde{b}) = b$ which shows that $\{a, b\}$ is an edge of $\overrightarrow{\varphi}_{p^m}(H)$. This completes the proof. ∎

We are now ready to prove Theorem 7.18.

Theorem 7.18. *Let Φ denote an edge-monotone graph property that is nontrivial on a prime power p^m, then there is a nonvanishing fixed point of* Syl_{p^m} *in K_{p^m} that contains $K_{p^{m-1}, p^{m-1}}$ as a subgraph.*

Proof. Our goal is to show that there is a fixed point H with the following properties

- $H = C_p^{A_1} \circ \cdots \circ C_p^{A_m}$ for $A_i \subseteq \mathbb{F}_p^+$,
- $\Phi(H) = 0$,
- $\varepsilon(A_1, \ldots, A_m) = 0$, and
- $\Phi(\tilde{H}) = 1$ for all proper sub-points $\tilde{H}$ of H.

As Φ is nontrivial on p^m, we have $\Phi(K_{p^m}) = 0$ and $\Phi(\mathrm{IS}_{p^m}) = \Phi(C_p^{\varnothing} \circ \cdots \circ C_p^{\varnothing}) = 1$. Now, let i denote the smallest value such that there is a fixed point H of level i that does not satisfy Φ, but all fixed points of level smaller than i satisfy Φ. As $\Phi(\mathrm{IS}_{p^m}) = 1$, we have $i > 0$.

Claim 7.19. *There is a fixed point H of level i with $\Phi(H) = 0$ and whose empty-prefix is zero.*

Proof. Toward an indirect proof, assume that all fixed points $C_p^{A_1} \circ \cdots \circ C_p^{A_m}$ of level i with $\varepsilon(A_1, \ldots, A_m) = 0$ satisfy Φ. We show that now, all fixed points of level i satisfy Φ, which is a contradiction to our choice of i.

To that end, fix an $F := C_p^{A_1} \circ \ldots \circ C_p^{A_m}$. If $\varepsilon(A_1, \ldots, A_m) = 0$, then $\Phi(F) = 1$ due to our assumption. Otherwise, $\varepsilon(A_1, \ldots, A_m) > 0$, which means that F has the form $F = C_p^{\varnothing} \circ \cdots \circ C_p^{\varnothing} \circ C_p^{A_1} \circ \ldots \circ C_p^{A_{m-j}}$ and is thus, by Lemma 7.17, isomorphic to an edge-subgraph of

$$F' := C_p^{A_1} \circ \cdots \circ F_p^{A_{m-j}} \circ C_p^{\varnothing} \circ \cdots \circ C_p^{\varnothing}.$$

By Lemma 7.12, the level of F is equal to the level of F'. Thus by assumption, F' satisfies Φ and hence F also satisfies Φ as Φ is edge-monotone. ◿

By Claim 7.19, we may assume that there is a fixed point $H := C_p^{A_1} \circ \cdots \circ C_p^{A_m}$ of level i with $\varepsilon(A_1, \ldots, A_m) = 0$ that does not satisfy Φ. Now, as the empty-prefix of H is zero, Lemma 7.15 yields that H contains $K_{p^{m-1}, p^{m-1}}$ as a subgraph. Finally, we apply Lemma 4.8 to show that $\widehat{\Phi}(H)$ does not vanish; which completes the proof. ∎

7.3 Quantitative Lower Bounds

Next, we prove Main Theorem 2. Observe that Main Theorem 2 differs from Theorem 5.16 since the constant α of Theorem 5.16 does not depend on a specific prime.

Main Theorem 2. *For each prime p, there is a constant $\gamma_p > 0$ such that for each integer m with $p^m \geq 3$ and each edge-monotone graph property Φ that is nontrivial on p^m, no algorithm (that reads the whole input) computes for every graph G the number $\#\mathrm{IndSub}((\Phi, k) \to G)$ in time $O(|V(G)|^{\gamma_p p^m})$, unless ETH fails.*

Proof. First, we show that assuming ETH, there are a constant $\gamma'_p > 0$ and a constant N_p such that for all fixed $k = p^m \geq N_p$ and each edge-monotone graph property Φ that is nontrivial on k, no algorithm computes for each graph G the number $\#\mathrm{IndSub}((\Phi, k) \to G)$ in time $O(|V(G)|^{\gamma_p k})$.

We intend to use Theorem C.2, which shows that nonvanishing graphs with large bicliques are sufficient to prove the result. Write $\beta > 0$ and N' for the constants from Theorem C.2. We define $h(k) := k/p$ and $\gamma'_p := \beta/p$ and $N_p := N' \cdot p$.

Consider a $k = p^m \geq N_p$. Clearly, $h(k) \geq N'$. Further, write Φ for an edge-monotone graph property that is nontrivial on k. Theorem 7.18 yields a graph F with k vertices, $\widehat{\Phi}(F) \neq 0$, and that contains $K_{h(k),h(k)}$ as a subgraph. Now, if for every graph G, we could compute the number $\#\mathrm{IndSub}((\Phi, k) \to \star)$ in time $O(|V(G)|^{\gamma_p k}) = O(|V(G)|^{\beta h(k)})$, then Theorem C.2 would show that ETH fails.

To obtain our lower bound also for $k < N_p$, we set $\gamma_p := \min(\gamma'_p, 1/(N_p + 1))$. Observe that for $k = p^m < N_p$, we obtain

$$O(|V(G)|^{\gamma_p k}) = o(|V(G)|).$$

Now, such a running time is unconditionally unachievable for any algorithm that reads the whole input. This completes the proof. ∎

Finally, we extend Main Theorem 2 from prime powers to products of c prime powers (times a constant d). However, this comes at the cost of a weaker lower bound in the exponent.

Theorem 7.20. *Let $p_1, \dots, p_c$ denote primes and let d denote a positive integer. Then, there is a constant $\gamma > 0$ (that depends on $p_1, \dots, p_c$ and d) such that for all fixed $k = d \cdot p_1^{m_1} \cdots p_c^{m_c} \geq 3$ and all edge-monotone graph properties Φ that are nontrivial on k, no algorithm (that reads the whole input) computes for every graph G the number $\#\mathrm{IndSub}((\Phi, k) \to G)$ in time $O(|V(G)|^{\gamma \sqrt[c]{k}})$, unless ETH fails.*

Proof. First, we show that assuming ETH, there are a constant $\gamma' > 0$ and a constant N such that for all fixed $k = d \cdot p_1^{m_1} \cdot \dots \cdot p_c^{m_c} \geq N$ and each edge-monotone graph property Φ that is nontrivial on k, no algorithm computes for each graph G the number $\#\mathrm{IndSub}((\Phi, k) \to G)$ in time $|V(G)|^{\gamma \sqrt[c]{k}}$.

To that end, we set $h(k) := \lceil \sqrt[c]{k/d} \rceil$. Observe that from Lemma 6.16, we obtain $q(k) \geq h(k)$ for all k with $k = d \cdot p_1^{m_1} \cdot \dots \cdot p_c^{m_c}$. As $q(k)$ is a positive integer, we may safely round up $\sqrt[c]{k/d}$. Further, write $\hat{c} > 0$ for the constant of Lemma 6.16 (3).

Next, write $\gamma_0 > 0$ and N' for the universal constants from Theorem C.2. Further, for each prime number p_i we use Main Theorem 2 to obtain a constant γ_i. Now, set $\gamma' := \min(\gamma_0, \dots, \gamma_p)/(2\sqrt[c]{d})$ and $N := \max(d \cdot N'^c, 4^c/(d\gamma^c), \exp(d/\hat{c}) + 1)$.

Fix a $k = d \cdot p_1^{m_1} \cdot \dots \cdot p_c^{m_c} \geq N$ and an edge-monotone graph property Φ that is nontrivial on k. Now, Φ is either scattered or concentrated on k; we provide a proof for both cases.

First, we consider the case that Φ is scattered on k. Define $m = q(k)$. Let H_m be the graph defined in Definition 6.21 such that $(\Phi - H) := \{G : G \uplus H \in \Phi\}$ is edge-monotone and nontrivial on $q(k)$. Now, we use Lemma 6.16 to obtain $q(k) \geq \hat{c}\log(k) > \hat{c}\log(\exp(d/\hat{c})) = d$, which implies $q(k) = p_i^{\hat{m}_i}$ for some

$i \in [c]$ and $\hat{m}_i \geq m_i$. Further, we obtain $q(k) \geq \sqrt[c]{k/d} \geq N'$ due to Lemma 6.16. Similarly to the proof of Main Theorem 1, assuming ETH, Main Theorem 2 shows that there is no algorithm that for each graph G computes the number #IndSub$(((\Phi - H), q(k)) \to G)$ in time $O(|V(G)|^{\gamma_i q(k)})$.

Now, assume that for each graph G, we could compute #IndSub$((\Phi, k) \to G)$ in time $O(|V(G)|^{\gamma' \sqrt[c]{k}})$. Then, using Lemma 6.12 to compute #IndSub$(((\Phi - H), q(k)) \to G)$ in time

$$O(2^{k-q(k)}(|V(G)| + k - q(k))^2 \cdot (|V(G)| + k - q(k))^{\gamma' \sqrt[c]{k}}) = O(|V(G)|^2 \cdot |V(G)|^{\gamma' \sqrt[c]{k}})$$
$$\overset{*}{\subseteq} O(|V(G)|^{\gamma_i \sqrt[c]{k/d}}) \subseteq O(|V(G)|^{\gamma_i q(k)})$$

would show that ETH fails. The last step is justified by $\sqrt[c]{k/d} \leq q(k)$. The step $(*)$ is justified by

$$2 + \gamma \sqrt[c]{k} \leq 2 + \frac{\gamma_i}{2\sqrt[c]{d}} \sqrt[c]{k} \leq \frac{\gamma_i}{\sqrt[c]{d}} \sqrt[c]{k} = \gamma_i \sqrt[c]{k/d}.$$

The second inequality is equivalent to $4^c/(d\gamma_i^c) \leq k$, which is true since $N \leq k$.

Next, we consider the case that Φ is concentrated on k. Then by definition there is a k-vertex graph H with a nonvanishing alternating enumerator that contains $K_{q(k),q(k)}$ as a subgraph. Now, since $k \geq d \cdot N'^c$ and since h is monotonically increasing, we obtain $h(k) \geq h(d \cdot N'^c) \geq N'$. Further, we have $q(k) \geq h(k) \geq \sqrt[c]{k/d}$, thus Theorem C.2 yields that there is no algorithm that computes for each graph G the number #IndSub$((\Phi, k) \to G)$ in time $O(|V(G)|^{\gamma_0 \sqrt[c]{k/d}}) \supseteq O(|V(G)|^{\gamma' \sqrt[c]{k}})$, where this is justified by $\gamma' \leq \gamma_0/\sqrt[c]{d}$.

To obtain our lower bound also for $k < N$, we set $\gamma := \min(\gamma', 1/(N+1))$. Now, for $k < N$ we obtain

$$O(|V(G)|^{\gamma \sqrt[c]{k}}) = o(|V(G)|).$$

Now, such a running time is unconditionally unachievable for any algorithm that reads the whole input. This completes the proof. ∎

Bibliography

ABD+18 João Araújo, Wolfram Bentz, Edward Dobson, Janusz Konieczny, and Joy Morris. Automorphism Groups of Circulant Digraphs With Applications to Semigroup Theory. *Combinatorica*, 38:1–28, 2018. doi:10.1007/s00493-016-3403-0. 23, 43

AP79 B. Alspach and T. D. Parsons. Isomorphism of circulant graphs and digraphs. *Discrete Mathematics*, 25:97–108, 1979. 23

BK11 Hans L. Bodlaender and Arie M. C. A. Koster. Treewidth computations II. Lower bounds. *Information and Computation*, 209(7):1103–1119, 2011. doi:10.1016/j.ic.2011.04.003. 24, 39

CCF+05 Jianer Chen, Benny Chor, Mike Fellows, Xiuzhen Huang, David Juedes, Iyad A. Kanj, and Ge Xia. Tight lower bounds for certain parameterized NP-hard problems. *Information and Computation*, 201(2):216–231, 2005. doi:10.1016/j.ic.2005.05.001. 17

CCMdM21 Vincent Cohen-Addad, Éric Colin de Verdière, Dániel Marx, and Arnaud de Mesmay. Almost Tight Lower Bounds for Hard Cutting Problems in Embedded Graphs. *J. ACM*, 68(4), jul 2021. doi:10.1145/3450704. 2, 6, 7, 53, 54, 55

CDM17 Radu Curticapean, Holger Dell, and Dániel Marx. Homomorphisms are a good basis for counting small subgraphs. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 210–223. ACM, 2017. doi:10.1145/3055399.3055502. 1, 54

CFK+15 Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Micha l Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer Cham, 1. edition, 2015. 12, 14, 17, 56, 57

CM14 Radu Curticapean and Dániel Marx. Complexity of Counting Subgraphs: Only the Boundedness of the Vertex-Cover Number Counts. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 130–139, 2014. doi:10.1109/FOCS.2014.22. 1

Cur13 Radu Curticapean. Counting Matchings of Size k Is $\#W[1]$-Hard. In Fedor V. Fomin, Rūsiņš Freivalds, Marta Kwiatkowska, and David Peleg, editors, *Automata, Languages, and Programming*, pages 352–363, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg. 1

DJ04 Víctor Dalmau and Peter Jonsson. The complexity of counting homomorphisms seen from the other side. *Theoretical Computer Science*, 329(1):315–323, 2004. doi:10.1016/j.tcs.2004.08.008. 6

DM09 Edward Dobson and Joy Morris. Automorphism Groups of Wreath Product Digraphs. *Electr. J. Comb.*, 16, January 2009. doi:10.37236/106. 43

DRSW22 Julian Dörfler, Marc Roth, Johannes Schmitt, and Philip Wellnitz. Counting Induced Subgraphs: An Algebraic Approach to #W[1]-hardness. *Algorithmica*, 84:379–404, 2022. doi:doi.org/10.1007/s00453-021-00894-9. 1, 2, 4, 5, 6, 7, 12, 52, 53, 54, 55, 57, 58

FG04 Jörg Flum and Martin Grohe. The Parameterized Complexity of Counting Problems. *SIAM Journal on Computing*, 33(4):892–922, 2004. doi:10.1137/S0097539703427203. 1, 17

FR22 Jacob Focke and Marc Roth. Counting Small Induced Subgraphs with Hereditary Properties. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2022, pages 1543–1551, New York, NY, USA, 2022. Association for Computing Machinery. doi:10.1145/3519935.3520008. 1, 2, 4

GJS76 M.R. Garey, D.S. Johnson, and L. Stockmeyer. Some simplified NP-complete graph problems. *Theoretical Computer Science*, 1(3):237–267, 1976. doi:10.1016/0304-3975(76)90059-1. 57

GSS01 Martin Grohe, Thomas Schwentick, and Luc Segoufin. When is the evaluation of conjunctive queries tractable? In *Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing*, STOC '01, page 657–666, New York, NY, USA, 2001. Association for Computing Machinery. doi:10.1145/380752.380867. 1

Har94 Frank Harary. *Graph theory*. Addison-Wesley, 3. print. edition, 1994. 13

IPZ01 Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *Journal of Computer and System Sciences*, 63(4):512–530, 2001. doi:10.1006/jcss.2001.1774. 1, 17

JM15a Mark Jerrum and Kitty Meeks. The parameterised complexity of counting connected subgraphs and graph motifs. *Journal of Computer and System Sciences*, 81(4):702–716, 2015. doi:10.1016/j.jcss.2014.11.015. 1, 3, 4

JM15b Mark Jerrum and Kitty Meeks. Some hard families of parameterized counting problems. *ACM Trans. Comput. Theory*, 7(3), jul 2015. doi:10.1145/2786017. 1, 3, 4

JM17 Mark Jerrum and Kitty Meeks. The parameterised complexity of counting even and odd induced subgraphs. *Combinatorica*, 37:965–990, 2017. doi:10.1007/s00493-016-3338-5. 1, 3

KSS84 Jeff Kahn, Michael Saks, and Dean Sturtevant. A topological approach to evasiveness. *Combinatorica*, 4:297–306, 1984. doi:10.1007/BF02579140. 7

Mar10 Dániel Marx. Can you beat treewidth? *Theory Comput.*, 6(1):85–112, 2010. doi:10.4086/TOC.2010.V006A005. 6

Mee16 Kitty Meeks. The challenges of unbounded treewidth in parameterised subgraph counting problems. *Discrete Applied Mathematics*, 198:170–194, 2016. doi:10.1016/j.dam.2015.06.019. 3, 4

MSOI+02 R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon. Network Motifs: Simple Building Blocks of Complex Networks. *Science*, 298(5594):824–827, 2002. doi:10.1126/science.298.5594.824. 1

Rob83 Guy Robin. Estimation de la fonction de tchebychef θ sur le k-ième nombre premier et grandes valeurs de la fonction $\omega(n)$ nombre de diviseurs premiers de n. *Acta Arithmetica*, 42(4):367–389, 1983. URL: http://eudml.org/doc/205883. 36

Rot94 Joseph J. Rotman. *An Introduction to the Theory of Groups*. Springer New York, 4 edition, 1994. doi:10.1007/978-1-4612-4176-8. 42, 43

RS20 Marc Roth and Johannes Schmitt. Counting Induced Subgraphs: A Topological Approach To #W[1]-Hardness. *Algorithmica*, 82:2267–2291, 2020. doi:10.1007/s00453-020-00676-9. 1, 2, 3, 4, 7

RSW23 Marc Roth, Johannes Schmitt, and Philip Wellnitz. Counting Small Induced Subgraphs Satisfying Monotone Properties. *SIAM Journal on Computing*, pages FOCS20–139–FOCS20–174, 2023+. accepted. doi:10.1137/20M1365624. 1, 2, 4

SJHS15 Benjamin Schiller, Sven Jager, Kay Hamacher, and Thorsten Strufe. Stream - a stream-based algorithm for counting motifs in dynamic graphs. In Adrian-Horia Dediu, Francisco Hernández-Quiroz, Carlos Martín-Vide, and David A. Rosenblueth, editors, *Algorithms for Computational Biology*, pages 53–67, Cham, 2015. Springer International Publishing. 1

SS05 Falk Schreiber and Henning Schwöbbermeyer. Frequency concepts and pattern detection for the analysis of motifs in networks. In Corrado Priami, Emanuela Merelli, Pablo Gonzalez, and Andrea Omicini, editors, *Transactions on Computational Systems Biology III*, pages 89–104, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg. 1

A Graph Properties and the Alternating Enumerator

As described in the Technical Overview, using a p-subgroup of the automorphism group, we can simplify the computation of the alternating enumerator (modulo p).

Lemma A.1 ([DRSW22]). *Let H denote a graph and let $\Gamma \subseteq \mathrm{Aut}(H)$ denote a p-group, then*

$$\widehat{\Phi}(H) \equiv_p \sum_{A \in \mathrm{FP}(\Gamma, H)} \Phi(A)(-1)^{\#E(A)}.$$

Proof. We follow the proof of [DRSW22, Lemma 1].

First, we rewrite the definition of the alternating enumerator to use edge-subgraphs of H instead of subsets of the edges of H. To that end, we readily see that for each edge-subgraph $A \in \mathcal{E}(H)$, we can find a subset $S \subseteq E(G)$ with $A = H\{S\}$ and vice versa. Hence, we obtain

$$\widehat{\Phi}(H) = \sum_{S \subseteq E(H)} \Phi(H\{S\})(-1)^{\#E(S)} = \sum_{A \in \mathcal{E}(H)} \Phi(A)(-1)^{\#E(A)}.$$

Now, recall that Γ acts on $\mathcal{E}(G)$ and consider the orbits of this group action $\cdot : \Gamma \times \mathcal{E}(G) \to \mathcal{E}(G)$. We choose a representative for each orbit, which means that each orbit has the form $\Gamma A_0 \coloneqq \{gA_0 : g \in \Gamma\}$ for an edge-subgraph $A_0 \in \mathcal{E}(G)$. We write $\mathcal{A}$ to denote the set of all representatives.

The orbits of $\cdot$ partition the set $\mathcal{E}(G)$, which allows us the rewrite the alternating enumerator as

$$\widehat{\Phi}(H) = \sum_{A \in \mathcal{E}(H)} \Phi(A)(-1)^{\#E(A)} = \sum_{A_0 \in \mathcal{A}} \sum_{A \in \Gamma A_0} \Phi(A)(-1)^{\#E(A)}.$$

Now, fix an $A \in \Gamma A_0$. By construction, there is a graph automorphism g with $g(A) = A_0$. Thus, we have $\Phi(A) = \Phi(A_0)$ and $\#E(A) = \#E(A_0)$ and hence

$$\widehat{\Phi}(H) = \sum_{A_0 \in \mathcal{A}} \sum_{A \in \Gamma A_0} \Phi(A_0)(-1)^{\#E(A_0)} = \sum_{A_0 \in \mathcal{A}} (\#\Gamma A_0)\Phi(A_0)(-1)^{\#E(A_0)}.$$

Now, we use the Orbit Stabilizer Theorem to see that the size $\#\Gamma A_0$ of any orbit of $\cdot$ is a divider of the group order of Γ. As Γ is a p-group, its order is equal to p^k for some $k \in \mathbb{N}$. Hence, $\#\Gamma A_0 \mod p$ is either equal to 0 (if $\#\Gamma A_0 > 1$); or equal to 1 (if $\#\Gamma A_0 = 1$). Finally, observe that $\#\Gamma A_0 = 1$ if and only if A_0 is a fixed point of Γ. This in turn means that only the fixed points remain when computing $\widehat{\Phi}(H) \mod p$. Hence, we obtain the claimed equation

$$\widehat{\Phi}(H) \equiv_p \sum_{A \in \mathrm{FP}(\Gamma, H)} \Phi(A)(-1)^{\#E(A)}.$$

∎

Using Lemma 3.1, we obtain a useful strengthening of Lemma A.1.

Corollary A.2. *For a graph G, a p-group $\Gamma \subseteq \mathrm{Aut}(G)$, and a fixed point $H \in \mathrm{FP}(\Gamma, G)$, we have*

$$\widehat{\Phi}(H) \equiv_p \sum_{\substack{A \in \mathrm{FP}(\Gamma, G) \\ E(A) \subseteq E(H)}} \Phi(A)(-1)^{\#E(A)}.$$

Proof. We use Lemma 3.1(2), and in particular the characterization of FP H in terms of FP G. Now, Lemma A.1 yields the claim. ∎

A.1 Lower Bounds for Counting Induced Subgraphs via the Alternating Enumerator

In this section, we prove Lemma A.8; which was implicitly proved in [DRSW22] (for the #W[1]-hardness). As advertised, we augment Lemma A.8 with a slightly stronger quantitative lower bound by lifting a similar result for #Hom($\mathcal{H}$) due to [CCMdM21].

Lemma A.8 ([DRSW22, CCMdM21])**.** *Let Φ denote a nontrivial graph property.*

- *If there is a sequence of graphs with unbounded treewidth where each graph has an alternating enumerator that is nonvanishing for Φ, then #IndSub(Φ) is #W[1]-hard.*
- *Assuming ETH, there is a universal constant $\alpha_{\textsc{IndSub}} > 0$ (that is independent of Φ) such that for any positive integer k for which there is a graph H_k with k vertices, $\widehat{\Phi}(H_k) \neq 0$, and $\operatorname{tw}(H_k) \geq 2$, no algorithm (that reads the whole input) computes for every graph G the number $\#\mathrm{IndSub}((\Phi, k) \to G)$ in time $O(|V(G)|^{\alpha_{\textsc{IndSub}} \operatorname{tw}(H_k)/\log \operatorname{tw}(H_k)})$.*

To that end, we need to work with colored versions of #Hom($\mathcal{H}$) and #IndSub(Φ), which we define next. Informally, in a colored problem, the vertices of the input graph G are partitioned into classes and we have to select exactly one vertex from each class. In the colored version of the homomorphism problem, we are given two graphs G and H, with a coloring $c : V(G) \to V(H)$ (that is, a partitioning $V(G)$ into $|V(H)|$ classes). A color-prescribed homomorphism h from H to G is a homomorphism from H to G such that $c(h(v)) = v$ for all $v \in V(H)$. Observe that if u and v are not adjacent in H, then the existence of the edges in G between $c^{-1}(u)$ and $c^{-1}(v)$ does not play any role whatsoever in the problem. Hence we might as well assume that there are no such edges in G at all, which formally means that c is a homomorphism from G to H. Therefore, we assume that c is indeed such a homomorphism, or in other words, G is *H-colored via c*.

We write cp-Hom($H \to G$) for the set of all color-prescribed homomorphism from H to a graph G that is H-colored via c. For a recursively enumerable class of graphs $\mathcal{H}$, in the problem #cp-Hom($\mathcal{H}$) we are given a graph $H \in \mathcal{H}$ and a graph G that is H-colored via c, the task is to compute the value #cp-Hom($H \to G$). We parameterize #cp-Hom($\mathcal{H}$) by $\kappa(H, G) := |V(H)|$.

In the colored variant of #IndSub(Φ), the vertices of the input graph G are partitioned into k classes and we are counting the number of k-vertex induced subgraphs satisfying Φ that contains exactly one vertex from each class. However, we need to define the problem in a way that allows a closer connection to #cp-Hom($\mathcal{H}$). For a recursively enumerable class of graphs $\mathcal{H}$, the input of #cp-IndSub($\Phi, \mathcal{H}$) consists of a graph G, a graph $H \in \mathcal{H}$, and a H-coloring c of G. The task is to compute the number of $|V(H)|$-vertex induced subgraphs of G that satisfies Φ and has exactly one vertex with each of the $|V(H)|$ colors. We denote this number by #cp-IndSub((Φ, H) $\to G$). We parameterize #cp-IndSub($\Phi, \mathcal{H}$) by $\kappa(H, G) := |V(H)|$.

Write Φ for a property and let $\mathcal{H}$ contain every graph with nonvanishing alternating enumerator. The proof of Lemma A.8 use the hardness of #Hom($\mathcal{H}$) to obtain hardness for #IndSub(Φ). It relies on the following chain of reductions of [DRSW22].

$$\#\textsc{Hom}(\mathcal{H}) \overset{\text{[DRSW22, Lemma 4]}}{\leq^{\text{fpt}}_{\text{T}}} \#\textsc{cp-Hom}(\mathcal{H}) \overset{\text{[DRSW22, Lemmas 7 and 8]}}{\leq^{\text{fpt}}_{\text{T}}} \#\textsc{cp-IndSub}(\Phi, \mathcal{H}) \overset{\text{[DRSW22, Lemma 10]}}{\leq^{\text{fpt}}_{\text{T}}} \#\textsc{IndSub}(\Phi); \tag{4}$$

The first reduction is very simple (essentially, making $|V(H)|$ copies of the vertex set of G) and the last reduction is a standard application of the Inclusion-Exclusion principle. Thus, let us focus on the reduction #cp-Hom($\mathcal{H}$) $\leq^{\mathrm{fpt}}_{\mathrm{T}}$ #cp-IndSub(Φ) and the corresponding key lemma from [DRSW22].

Lemma A.3 ([DRSW22, Lemma 8])**.** *Let H denote a graph, let Φ denote a graph property, and let G denote an H-colored graph. Then, we have*

$$\#\text{cp-IndSub}((\Phi, H) \to G) = \sum_{S \subseteq E(H)} \Phi(H\{S\}) \sum_{J \subseteq E(H) \setminus S} (-1)^{\#J} \cdot \#\text{cp-Hom}(H\{S \cup J\} \to G).$$

Further, the absolute values of $\widehat{\Phi}(H)$ and of the coefficient of #cp-Hom($H \to G$) *are equal.*

In particular, the second part of Lemma A.3 yields that a term #cp-Hom($H \to G$) appears in the sum if and only if $\widehat{\Phi}(H) \neq 0$, which is part of our assumption in Lemma A.8.

Observe that Lemma A.3 in itself does not suffice to obtain the claimed reduction #cp-Hom($\mathcal{H}$) $\leq^{\mathrm{fpt}}_{\mathrm{T}}$ #cp-IndSub(Φ): the oracle for #cp-IndSub(Φ) computes only a sum in which #cp-Hom($H \to G$) occurs as some term—we still need to extract the value #cp-Hom($H \to G$) out of the result of the oracle. Fortunately for us, [DRSW22, Lemma 7] does exactly that by showing a generalization of the *Complexity Monotonicity* of [CDM17]. The other reductions of (4) can be used without modifications. In total, we obtain the #W**[**1**]**-hardness part of Lemma A.8.

Next, we turn to ETH-based lower bounds. We start from the binary constraint satisfaction problem (CSP).

Definition A.4. *A (binary) CSP instance is a triple (V, D, C) where*
- *V is a set of variables*
- *D is a domain of values,*
- *C is a set of constraints. Each constraint is a triple (u, v, R) where $(u, v) \in V^2$, and $R \subseteq D^2$.*

A solution *to (V, D, C) is a function $f : V \to D$ such that for all constraints (u, v, R), the pair $(f(u), f(v))$ is in R. The* primal graph *of a CSP insurance (V, D, C) is a graph H with vertex set V such that $u, v \in V(H)$ are adjacent if and only if there is a constraint in C of the form (u, v, R).*

In particular, we use the following result of Cohen-Addad, Colin de Verdière, Marx, and de Mesmay, which we can easily modify for our purposes.

Theorem A.5 ([CCMdM21, Theorem 2.7])**.** *Assuming ETH, there is a universal constant $\alpha_{CSP} > 0$ such that for any fixed graph H with* $\mathrm{tw}(H) \geq 2$*, there is no algorithm that decides the binary CSP instances whose primal graph is H in time $O(|D|^{\alpha_{CSP} \cdot \mathrm{tw}(H)/\log \mathrm{tw}(H)})$.*

Corollary A.6. *Assuming ETH, there is a universal constant $\alpha_{Hom} > 0$ such that for any fixed graph H with* $\mathrm{tw}(H) \geq 2$*, there is no algorithm that computes* #Hom($H \to \star$) *on input G in time $O(|V(G)|^{\alpha_{Hom} \cdot \mathrm{tw}(H)/\log \mathrm{tw}(H)})$.*

Proof. We model the decision problem Hom($\{H\}$) as a CSP. For a given graph G, we define the CSP instance $I = (V(H), V(G), C)$ with $C := \{(u, v, E(G)) : \{u, v\} \in E(H)\}$. Each solution $h : V(H) \to V(G)$ to I is also a homomorphism from H to G and each homomorphism from H to G is also a solution to I (observe that whenever the sole relation in I is the edge relation, the definition of a solution to I coincides with the definition of a graph homomorphism). Thus, Hom($H \to G$) is non-empty if and only if I has a solution.

As Hom($\{H\}$) can be solved via its counting version #Hom($\{H\}$), we obtain that #Hom($H \to G$) cannot be computed in time $O(|V(G)|^{\alpha_{\mathrm{CSP}} \cdot \mathrm{tw}(H)/\log \mathrm{tw}(H)})$. Choosing $\alpha_{\mathrm{Hom}} := \alpha_{\mathrm{CSP}}$ yields the claim.

Next, we again use the reductions of (4). In particular, we observe that said reductions preserve the exponent in the running time (up to some constant additive term). It is useful to have a separate (sub-)claim for a reduction that starts from color-prescribed homomorphisms.

Lemma A.7. *Let Φ denote a graph property and let H denote a graph such that $\widehat{\Phi}(H) \neq 0$.*

Any algorithm that computes for each graph G' the number $\#\mathrm{IndSub}((\Phi, |V(H)|) \to G')$ in time $O(|V(G')|^{\beta})$ implies

- *an algorithm that for each graph G computes the number $\#\text{cp-Hom}(H \to G)$ in time $O(|V(G)|^{\beta+2})$.*
- *an algorithm that for each graph G computes the number $\#\mathrm{Hom}(H \to G)$ in time $O(|V(G)|^{\beta+3})$.*

Proof. Write $k := |V(H)|$ and observe that for our purposes, k is a constant. We first construct an algorithm that solves $\#\textsc{cp-Hom}(\{H\})$ using an oracle for $\#\mathrm{IndSub}((\Phi, k) \to \star)$. To that end, suppose that we are given a graph G that is H-colored via c. We wish to compute the number $\#\mathrm{Hom}(H \to G)$.

(1) First, we use the reduction from $\#\textsc{cp-Hom}(\mathcal{H})$ to $\#\textsc{cp-IndSub}(\Phi)$ [DRSW22, Lemmas 7 and 8] to compute for the graph G the number $\#\text{cp-Hom}(H \to G)$ in time $O(f(k) \cdot |V(G)|)$ using an oracle for $\#\text{cp-IndSub}((\Phi, H) \to \star)$ for a computable function f. In said reduction, each graph G' that is used inside an oracle call satisfies $|V(G')| \leq f(k) \cdot |V(G)|$.

(2) Next, we use the reduction from $\#\textsc{cp-IndSub}(\Phi)$ to $\#\textsc{IndSub}(\Phi)$ [DRSW22, Lemma 10] to compute for each G' the number $\#\text{cp-IndSub}((\Phi, H) \to G')$ in time $O(g(k) \cdot |V(G')|)$ using an oracle for $\#\mathrm{IndSub}((\Phi, |V(H)|) \to \star)$ for a computable function g. In said reduction, each graph G'' that is used in oracle calls satisfy $|V(G'')| \leq |V(G')|$.

(3) Lastly, we compute for each G'' the number $\#\mathrm{IndSub}((\Phi, |V(H)|) \to G'')$ in time $O(|V(G'')|^{\beta})$ using the algorithm which we assumed to exist.

We combine the above steps to obtain an algorithm that computes $\#\text{cp-Hom}(H \to G)$ in time

$$O\Big(\underbrace{f(k) \cdot |V(G)|}_{(1)} \quad \cdot \quad \underbrace{g(k) \cdot (f(k) \cdot |V(G)|)}_{(2)} \quad \cdot \quad \underbrace{(f(k) \cdot |V(G)|)^{\beta}}_{(3)}\Big),$$

which can be rewritten into $O(|V(G)|^2 \cdot |V(G)|^{\beta})$, as $k = |V(H)|$ is a constant.

Next, we construct an algorithm that solves $\#\textsc{Hom}(\{H\})$ using an oracle for $\#\mathrm{IndSub}((\Phi, k) \to \star)$. To that end, suppose that we are given a graph G. We wish to compute the number $\#\mathrm{Hom}(H \to G)$.

First, we use the reduction from $\#\textsc{Hom}(\mathcal{H})$ to $\#\textsc{cp-Hom}(\mathcal{H})$ [DRSW22, Lemma 4] to compute the value $\#\mathrm{Hom}(H \to G)$ in time $O(h(k) \cdot |V(G)|)$ using an oracle for $\#\text{cp-Hom}(H \to \star)$ for a computable function h. In said reduction, each graph G' that is used in oracle calls satisfy $|V(G')| \leq h(k) \cdot |V(G)|$.

Next, our algorithm from the first part of the proof allows us to compute each value $\#\text{cp-Hom}(H \to G')$ in time $O((h(k)|V(G')|)^{\beta+2})$ using a oracle for $\#\mathrm{IndSub}((\Phi, |V(H)|) \to \star)$. Hence, we obtain a total running time of $O(h(k) \cdot |V(G)| \cdot (h(k) \cdot |V(G)|)^{\beta+2})$ which can be rewritten into $O(|V(G)|^3 \cdot |V(G)|^{\beta})$ as k is a constant. ◀

Putting everything together, we obtain Lemma A.8.

Lemma A.8 ([DRSW22, CCMdM21]). *Let Φ denote a nontrivial graph property.*

- *If there is a sequence of graphs with unbounded treewidth where each graph has an alternating enumerator that is nonvanishing for Φ, then $\#\textsc{IndSub}(\Phi)$ is $\#\mathrm{W}[1]$-hard.*
- *Assuming ETH, there is a universal constant $\alpha_{\textsc{IndSub}} > 0$ (that is independent of Φ) such that for any positive integer k for which there is a graph H_k with k vertices, $\widehat{\Phi}(H_k) \neq 0$, and $\mathrm{tw}(H_k) \geq 2$, no algorithm (that reads the whole input) computes for every graph G the number $\#\mathrm{IndSub}((\Phi, k) \to G)$ in time $O(|V(G)|^{\alpha_{\textsc{IndSub}} \mathrm{tw}(H_k)/\log \mathrm{tw}(H_k)})$.*

Proof. We write $\mathcal{H}$ for the set $\{H \in \mathcal{G} : \widehat{\Phi}(H) \neq 0\}$. The set $\mathcal{H}$ is recursively enumerable since the set of all graphs is recursively enumerable and Φ is computable. This in turn implies that the alternating enumerator is computable.

Observe that the treewidth of the elements of $\mathcal{H}$ is unbounded by the assumption of the lemma. This means that $\#\textsc{Hom}(\mathcal{H})$ is #W[1]-hard when parameterized by the pattern size $|V(H)|$. Finally, we use the parameterized reductions (4) from $\#\textsc{Hom}(\mathcal{H})$ to $\#\textsc{IndSub}(\Phi)$, which proves that $\#\textsc{IndSub}(\Phi)$ is also #W[1]-hard.

Next, we turn to the ETH-based lower bounds. Set $\alpha'_{\textsc{IndSub}} := \alpha_{\textsc{Hom}}/2$, $N := \max(2, (6/\alpha_{\textsc{Hom}})^2)$. Let k denote a fixed integer such that there is a nonvanishing k-vertex graph H_k with $\mathrm{tw}(H_k) \geq N$. We show that any algorithm $\mathbb{A}$ that computes for every graph G the value $\#\mathrm{IndSub}((\Phi, k) \to G)$ in time $O(|V(G)|^{\alpha'_{\textsc{IndSub}} \mathrm{tw}(H_k)/\log \mathrm{tw}(H_k)})$ implies that we can compute for any graph G the value $\#\mathrm{Hom}(H_k \to G)$ in time $O(|V(G)|^{\alpha_{\textsc{Hom}} \mathrm{tw}(H_k)/\log \mathrm{tw}(H_k)})$ (which would violate ETH due to Corollary A.6).

To that end, first observe that by Lemma A.7, the algorithm $\mathbb{A}$ yields an algorithm $\mathbb{B}$ to compute for every graph G the value $\#\mathrm{Hom}(H_k \to G)$ in time $O(|V(G)|^{\alpha'_{\textsc{IndSub}} \mathrm{tw}(H_k)/\log \mathrm{tw}(H_k)+3})$.

Claim A.9. *For* $\mathrm{tw}(H_k) \geq \max(2, (6/\alpha_{\textsc{Hom}})^2)$, *we have*

$$\alpha'_{\textsc{IndSub}} \mathrm{tw}(H)/\log \mathrm{tw}(H_k) + 3 \leq \alpha_{\textsc{Hom}} \mathrm{tw}(H_k)/\log \mathrm{tw}(H_k).$$

Proof. For $\mathrm{tw}(H_k) > 1$, we have $\mathrm{tw}(H_k)/\log \mathrm{tw}(H_k) > \sqrt{\mathrm{tw}(H_k)}$. Hence, for $\mathrm{tw}(H_k) \geq (6/\alpha_{\textsc{Hom}})^2$, we have

$$\mathrm{tw}(H)/\log \mathrm{tw}(H_k) \geq 6/\alpha_{\textsc{Hom}} = 3/(\alpha_{\textsc{Hom}} - \alpha_{\textsc{Hom}}/2) = 3/(\alpha_{\textsc{Hom}} - \alpha'_{\textsc{IndSub}}).$$

Now, rearranging yields the claim. ◻

From Claim A.9, we conclude that $\mathbb{B}$ computes $\#\mathrm{Hom}(H_k \to G)$ in time $O(|V(G)|^{\alpha_{\textsc{Hom}} \mathrm{tw}(H_k)/\log \mathrm{tw}(H_k)})$. From Corollary A.6 (and assuming $\mathrm{tw}(H_k) \geq 2$), we conclude that the algorithm $\mathbb{B}$ violates ETH.

Define $\alpha_{\textsc{IndSub}} := \min(\alpha'_{\textsc{IndSub}}, 1/N)$. For $2 \leq \mathrm{tw}(H_k) < N$, we obtain an algorithm that computes $\#\mathrm{IndSub}((\Phi, |V(H_k)|) \to G)$ in time $|V(G)|^{\alpha_{\textsc{IndSub}} \mathrm{tw}(H_k)/\log \mathrm{tw}(H_k)} = o(|V(G)|)$. Now, such a running time is unconditionally unachievable for any algorithm that reads the whole input. This completes the proof. ◼

B ETH-based Lower Bounds for k-Clique

We discuss the following useful result on finding k cliques.

Theorem B.1 (Theorem 14.21 in [CFK⁺15])**.** *Assuming ETH, there is no $f(k)|V(G)|^{o(k)}$-time algorithm for* Clique *or* Independent Set *for any computable function f.* ◼

Next, we modify Theorem B.1. The decision problem k-Clique gets as input a graph G and checks if G contains a k-clique as a subgraph.

Lemma B.2 (Modification of 14.21 in [CFK⁺15])**.** *Assuming ETH, there is a constant $\alpha > 0$ such that for $k \geq 3$, no algorithm (that reads the whole input) solves k-*Clique *on graph G in time $O(|V(G)|^{\alpha k})$.*

Proof. Write $n := |V(G)|$ for the number of vertices of the input graph G. Assuming ETH, there is a $\delta > 0$ such that no algorithm solves 3-SAT in time $O(2^{\delta v})$, where v is the number of variables. We show that a similar statement also holds for the 3-Coloring problem.

Claim B.3. *Assuming ETH, there is a $\delta' > 0$ such that no algorithm solves 3-*Coloring *in time $O(2^{\delta' n})$, where n is the number of vertices.*

Proof. For a 3-SAT formula ϕ with v variables and m clauses, the standard reduction [GJS76, Theorem 2.1] from 3-SAT to 3-Coloring constructs a graph with $3 + 2v + 6m$ vertices that contains a 3-coloring if and only if ϕ is satisfiable. An algorithm that solves 3-Coloring in time $O(2^{\beta n})$ could be used to solve 3-SAT in time $O(2^{6\beta(v+m)})$. The Sparsification Lemma [CFK+15, Theorem 14.4] yields the existence of a $\delta' > 0$ such that there is no algorithm that solves 3-Coloring in time $O(2^{\delta' n})$. ◿

Next, we show that we can solve 3-Coloring by using an algorithm for k-Clique that runs in $O(n^{\alpha k})$.

Claim B.4. *If we can solve k-Clique in time $O(n^{\alpha k})$, then we can solve 3-Coloring in time $O(3^{\alpha n} + n^2 \cdot 3^{2n/k})$.*

Proof. First, we split the n vertices of G into k blocks $V_1, \dots, V_k$ of size at most $\lceil n/k \rceil$ each. Next, we construct a graph H in the following way. For each proper 3-coloring of a block V_i, we create a vertex that represents this coloring and add said vertex to H. The number of vertices in H is at most

$$|V(H)| \leq k \cdot 3^{\lceil n/k \rceil} \leq k \cdot 3^{n/k+1}.$$

Write u for a vertex in H that represents a coloring of $G[V_i]$ and write v for a vertex in H that represents a coloring of $G[V_j]$. We add an edge between u and v if and only if $i \neq j$ and the coloring of u and v is a proper 3-coloring of $G[V_i \cup V_j]$. Observe that we can construct this graph in time $O((k \cdot 3^{n/k+1})^2 \cdot n^2) = O(n^2 \cdot 3^{2n/k})$.

It is easy to verify that H contains a k-clique if and only if there is a proper 3-coloring of G. Thus, we can use our $O(n^{\alpha k})$ time algorithm for k-Clique to solve 3-Coloring in time $O((k \cdot 3^{n/k+1})^{\alpha k} + n^2 \cdot 3^{2n/k}) = O(3^{\alpha n} + n^2 \cdot 3^{2n/k})$. ◿

Set $\alpha := \min(\log_3(2)\delta', \alpha/3)$. If there is a $k > 2/\alpha$ such that we can solve k-Clique in time $O(n^{\alpha k})$, then we can use Claim B.4 to solve 3-Coloring in time $O(3^{\alpha n} + n^2 \cdot 3^{2n/k}) \subseteq O(2^{\delta' n})$. According to Claim B.3, this is only possible if ETH fails. Otherwise, $\alpha k < 1$, and there is no sublinear algorithm that reads the whole input and solves k-Clique in sublinear time. ∎

C Tight Lower Bounds for Counting Induced Subgraphs

In this section, we show how to obtain tight lower bounds under ETH for #IndSub$((\Phi, k) \to \star)$ by modifying a reduction from [DRSW22]. We start from the non-parameterized decision problem k-Clique and solve it with an oracle for #IndSub$((\Phi, k) \to \star)$. Observe that k-Clique cannot be solved in time $O(n^{\alpha k})$ for a fixed $\alpha > 0$ unless ETH fails (see Lemma B.2), which yields a lower bound for computing #IndSub$((\Phi, k) \to \star)$. The reduction is similar to the reduction shown in Appendix A with the difference that we start from k-Clique instead of Hom$(\{H\})$. The main idea is that we can solve k-Clique using an oracle for #cp-Hom$(F \to \star)$ as long as F contains $K_{k,k}$ as a subgraph.

Lemma C.1 (Modification of [DRSW22, Lemma 11]). *There is an algorithm that given a positive integer $\ell > 1$, a graph F (that contains $K_{\ell,\ell}$ as a subgraph), and a graph G; computes a F-colored graph G' with $2\ell|V(G)| + (|V(F)| - 2\ell)$ vertices. Further, the number of cliques of size ℓ in G equals #cp-Hom$(F \to G')$. The running time of the algorithm is $O(|V(F)|^{|V(F)|+2} + |V(F)|^2|V(G)|^2)$*

Proof. Write $\tilde{\ell}$ for $|V(F)| - 2\ell$. For the proof, we modify the construction of [DRSW22, Lemma 11]. First, observe that we can locate the vertices of the subgraph $K_{\ell,\ell}$ in F by looping over all vertex sets A and B of size ℓ and checking if the induced graph has the complete bipartite graph $K_{\ell,\ell}$ as a subgraph. This can be done in $O(|V(F)|^2 \cdot |V(F)|^{2\ell}) \subseteq O(|V(F)|^{|V(F)|+2})$ by using a brute force implementation.

Next, we relabel the vertices of the graph F by splitting them up into a left side, a right side and the remaining vertices. Formally, we use $V(F) = \{a_i, b_i : i \in [\ell]\} \cup \{x_i : i \in [\tilde{\ell}]\}$, and ensure that the induced subgraph of $\{a_i, b_i : i \in [\ell]\}$ contains $K_{\ell,\ell}$, where $\{a_i : i \in [\ell]\}$ is the left side and $\{b_i : i \in [\ell]\}$ is the right side of the complete bipartite graph.[11]

Now, let G denote a graph with vertex set $\{v_i : i \in [n]\}$. We construct the graph G' on the vertex set $\{u_{i,j}, w_{i,j}, y_k : i \in [\ell], j \in [n], k \in [\tilde{\ell}]\}$ with the F-coloring given by $c(u_{i,j}) = a_i$, $c(w_{i,j}) = b_i$ and $c(y_k) = x_k$. For each k, we add an edge between y_k and all other vertices in G'. Additionally, we add an edge between $u_{i,j}$ and $w_{i',j'}$ if and only if

- either $(i, j) = (i', j')$,
- or $i < i'$, $j < j'$ and the vertices v_j and $v_{j'}$ are adjacent,
- or $i > i'$, $j > j'$ and the vertices v_j and $v_{j'}$ are adjacent.

Further, we add the edges $\{u_{i,j}, u_{i',j'}\}$ and $\{w_{i,j}, w_{i',j'}\}$ to G'.

Let $C = (v_{j_1}, \ldots, v_{j_\ell})$ denote an ordered tuple (that is $j_k < j_{k'}$ for $k < k'$) such that $\{v_{j_1}, \ldots, v_{j_\ell}\}$ is a ℓ-clique in G. We construct the homomorphism $h_C : V(F) \to V(G')$ with $h_C(a_i) = u_{i,j_i}$, $h_C(b_i) = v_{i,j_i}$ and $h_C(x_k) = y_k$. Observe that this defines a color-prescribed homomorphism $h_C \in$ cp-Hom$(F \to G')$.

Next, consider an $h' \in$ cp-Hom$(F \to G')$. Then, $h'(x_k) = y_k$ since y_k is the only vertex in G' with $c(y_k) = x_k$. Further, we obtain $h'(a_i) = u_{i,\alpha_i}$ and $h'(b_j) = w_{j,\beta_j}$. Observe that for all $i \in [\ell]$ the edge $\{a_i, b_i\}$ is in F. Thus $\{h'(a_i), h'(b_i)\} = \{u_{i,\alpha_i}, w_{i,\beta_i}\}$ is an edge in G' which is only possible if $\alpha_i = \beta_i$.

Next, we define the tuple $C := (v_{\alpha_1}, \ldots, v_{\alpha_\ell})$. For all $i < i'$, we know that the edge $\{a_i, b_{i'}\}$ is in F, hence the edge $\{u_{i,\alpha_i}, w_{i',\beta_{i'}}\}$ is also in G' which implies that $\alpha_i < \beta_{i'} = \alpha_{i'}$. Observe that this also implies that there is an edge between v_{α_i} and $v_{\alpha_{i'}}$. So, the indices of the tuple C are ordered and $\{v_{\alpha_1}, \ldots, v_{\alpha_\ell}\}$ is a ℓ-clique in G. Also, observe $h' = h_C$, where h_C is the color-prescribed homomorphism from above. Thus each ordered ℓ-cliques C yields a color-prescribed homomorphism h_C and each color-prescribed homomorphism h' yields an ordered ℓ-clique C. This shows an one-to-one correspondence between color-prescribed homomorphisms in cp-Hom$(F \to G')$ and ℓ-cliques in G. ∎

This means that we can solve the decision problem k-Clique for a graph G and a parameter k by computing #cp-Hom$(F \to \star)$. Further, if we assume that F is non-vanishing then we can use the reduction shown in Lemma A.7 to compute #cp-Hom$(F \to \star)$ using #IndSub$((\Phi, |V(F)|) \to \star)$. Thus we can use #IndSub$((\Phi, |V(F)|) \to \star)$ to solve k-Clique.

Theorem C.2 (Modification of [DRSW22]). *There is a global constant $\beta > 0$ and a positive integer N such that for all graph properties Φ, functions h, numbers k with*

- *$h(k) \geq N$*
- *there is a graph F with k vertices and $\hat{\Phi}(F) \neq 0$,*
- *and F contains $K_{h(k),h(k)}$ as a subgraph*

there is no algorithm (that reads the whole input) that for every G computes #IndSub$((\Phi, k) \to G)$ *in time $O(|V(G)|^{\beta h(k)})$ unless ETH fails.*

Proof. We show how to use #IndSub$((\Phi, k) \to \star)$ to solve $h(k)$-Clique.

Claim C.3. *For a fixed k such that there is a graph F with k vertices, $\hat{\Phi}(F) \neq 0$, and F contains $K_{h(k),h(k)}$ as a subgraph; if there is an algorithm that computes for each graph G' the value* #IndSub$((\Phi, k) \to G')$ *in time $O(|V(G')|^{\gamma})$, then $h(k)$-Clique can be computed for each graph G in time $O(|V(G)|^{\gamma+2})$.*

[11] Note that they may be edges of the form a_i to a_j or form b_i to b_j.

Proof. We construct an algorithm that solves $h(k)$-Clique using an oracle for #IndSub$((\Phi, j) \to \star)$. Since k and F are fixed, we can assume that our algorithm knows these elements. Fix a graph G. Then we use the algorithm from Lemma C.1 to construct a graph G' such that G contains an $h(k)$-clique if and only if #cp-Hom$(F \to G')$ is not zero. The running time of this construction is in $O(|V(G)|^2)$ since $|V(F)| = k$ is constant. Further, we obtain $|V(G')| \leq 2h(k)|V(G)| + k = O(|V(G)|)$.

If we can compute #IndSub$((\Phi, k) \to |V(G')|)$ in time $O(|V(G')|^{\gamma})$, then we can use Lemma A.7 to compute #cp-Hom$(F \to G')$ in time $O((2h(k)|V(G)| + k)^{\gamma+2})$. Thus, we can solve $h(k)$-Clique in time $O(|V(G)|^2 + (2h(k)|V(G)| + k)^{\gamma+2})$ which is in $O(|V(G)|^{\gamma+2})$ since k is fixed. ◿

According to Lemma B.2, there is a constant $\alpha > 0$ such that no algorithm solves $h(k)$-Clique in time $O(|V(G)|^{\alpha h(k)})$ for a fixed $h(k) \geq 3$ unless ETH fails. Set $\beta \coloneqq \alpha/2$ and $N \coloneqq \max(3, 4/\alpha)$. If there are a k with $h(k) \geq N \geq 3$, a graph F with k vertices such that $\widehat{\Phi}(F) \neq 0$, and F contains $K_{h(k),h(k)}$ as a subgraph, and an algorithm that solves #IndSub$((\Phi, k) \to G')$ in time $O(|V(G')|^{\beta h(k)})$, then we can use Claim C.3 to solve $h(k)$-Clique in time $O(|V(G)|^{\beta h(k)+2})$. Observe that $\beta h(k) + 2 \leq \alpha h(k)$ for $h(k) \geq 4/\alpha$. Thus, we can use solve $h(k)$-Clique in time $O(n^{\alpha h(k)})$. Hence, ETH fails. ◾