Counting Induced Subgraphs: An Algebraic Approach to #W[1]-hardness*

Julian Dörfler

Saarbrücken Graduate School of Computer Science, Saarland Informatics Campus (SIC), Germany
s8judoer@stud.uni-saarland.de

Marc Roth

Cluster of Excellence (MMCI), Saarland Informatics Campus (SIC), Saarbrücken, Germany
https://www.roth-marc.com
mroth@mmci.uni-saarland.de

Johannes Schmitt

ETH Zürich, Switzerland
https://people.math.ethz.ch/~schmittj/
johannes.schmitt@math.ethz.ch

Philip Wellnitz

Max Planck Institute for Informatics, Saarland Informatics Campus (SIC), Saarbrücken, Germany
https://people.mpi-inf.mpg.de/~wellnitz/
wellnitz@mpi-inf.mpg.de

Abstract

We study the problem #IndSub(Φ) of counting all induced subgraphs of size k in a graph G that satisfy the property Φ. This problem was introduced by Jerrum and Meeks and shown to be #W[1]-hard when parameterized by k for some families of properties Φ including, among others, connectivity [JCSS 15] and even- or oddness of the number of edges [Combinatorica 17]. Very recently [IPEC 18], two of the authors introduced a novel technique for the complexity analysis of #IndSub(Φ), inspired by the "topological approach to evasiveness" of Kahn, Saks and Sturtevant [FOCS 83] and the framework of graph motif parameters due to Curticapean, Dell and Marx [STOC 17], allowing them to prove hardness of a wide range of properties Φ. In this work, we refine this technique for graph properties that are non-trivial on edge-transitive graphs with a prime power number of edges. In particular, we fully classify the case of monotone bipartite graph properties: It is shown that, given *any* graph property Φ that is closed under the removal of vertices and edges, and that is non-trivial for bipartite graphs, the problem #IndSub(Φ) is #W[1]-hard and cannot be solved in time $f(k) \cdot n^{o(k)}$ for any computable function f, unless the Exponential Time Hypothesis fails. This holds true even if the input graph is restricted to be bipartite and counting is done modulo a fixed prime. A similar result is shown for properties that are closed under the removal of edges only.

2012 ACM Subject Classification Theory of computation → Parameterized complexity and exact algorithms; Theory of computation → Problems, reductions and completeness; Mathematics of computing → Combinatorics; Mathematics of computing → Graph theory

Keywords and phrases counting complexity, edge-transitive graphs, graph homomorphisms, induced subgraphs, parameterized complexity

Digital Object Identifier 10.4230/LIPIcs.MFCS.2019.26

Related Version A full version of the paper is available at https://arxiv.org/abs/1904.10479.

Funding *Johannes Schmitt*: The third author has received funding from the European Research Council (ERC) under the EU's Horizon 2020 research and innovation programme (No 786580).

Acknowledgements We are very grateful to Radu Curticapean and Holger Dell for fruitful discussions and valuable feedback on early drafts of this work.

* All authors of this paper are students.

44th International Symposium on Mathematical Foundations of Computer Science (MFCS 2019).
Editors: Peter Rossmanith, Pinar Heggernes, and Joost-Pieter Katoen; Article No. 26; pp. 26:1–26:14
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

The study of the computational complexity of counting problems was initiated by Valiant's seminal work about the complexity of computing the permanent [22]. In contrast to a decision problem which requires to *verify* the existence of a solution, a counting problem asks to compute the *number* of solutions. Counting complexity theory is particularly interesting for problems whose decision versions are solvable efficiently but whose counting versions are intractable. One such example is the problem of finding/counting perfect matchings, whose decision version is solvable in polynomial time [7] and whose counting version is as least as hard as every problem in the Polynomial Hierarchy PH with respect to polynomial-time Turing reductions [22, 21]. In this work, we consider the following problem which was first introduced by Jerrum and Meeks [11]: Fix a graph property Φ, given a graph G and a positive integer k, compute the number of all induced subgraphs of G with k vertices that satisfy Φ. We denote this problem by #IndSub(Φ) and remark that, strictly speaking, #IndSub(Φ) is the *unlabeled* version of *p*-#InducedSubgraphWithProperty(Φ) as defined in [12, Section 1.3.1]. In particular, our properties only depend on the isomorphism type of a graph and not on any labeling of the vertices.

We study the *parameterized complexity* of #IndSub(Φ) depending on the property Φ. The underlying framework, known as *parameterized counting complexity theory*, was introduced independently by Flum and Grohe [8] and McCartin [16], and constitutes a hybrid of (classical) computational counting and parameterized complexity theory. Here, the method of parameterization allows us to perform a multivariate analysis of the complexity of #IndSub(Φ): Instead of the distinction between polynomial-time solvable and NP-hard cases, we search for properties Φ for which the problem is solvable in time $f(k) \cdot n^{O(1)}$, where n is the number of vertices of the graph and f can be any computable function. If this is the case, the problem is called *fixed-parameter tractable*. Unfortunately, the only known cases of Φ for which #IndSub(Φ) is fixed-parameter tractable are trivial in the sense that there are only finitely many k such that Φ is neither true nor false on the set of all graphs with k vertices. On the contrary, it is easy to see that #IndSub(Φ) is most likely not fixed-parameter tractable if Φ encodes a problem whose decision version is already known to be hard. An example of the latter is the property of being a complete graph. In this case, the problem #IndSub(Φ) is identical to the problem of counting cliques of size k, for which even the decision version, that is, *finding* a clique of size k in a graph with n vertices, cannot be done in time $f(k) \cdot n^{o(k)}$, unless the Exponential Time Hypothesis fails [3, 4].

The first non-trivial hardness result of #IndSub(Φ) was given by Jerrum and Meeks for Φ the property of being connected [11]. Note that, in this case, the decision version of the problem can be solved efficiently as, on input G and k, one only has to decide whether there exists a connected component of G of size at least k. This result initiated a line of research in which Jerrum and Meeks proved fixed-parameter tractability of #IndSub(Φ) to be unlikely for the property of having an even (or odd) number of edges [12], for properties that induce low edge densities [10] and for properties that are closed under the addition of edges and whose (edge-)minimal elements have large treewidth [17]. More precisely, all of those results established hardness for the parameterized complexity class #W[1], which can be seen as the parameterized counting equivalent of NP. In a recent breakthrough result [5], Curticapean, Dell and Marx have shown, that for every graph property Φ, the problem #IndSub(Φ) is either fixed-parameter tractable or hard for #W[1], that is, there are no cases of intermediate difficulty. On the downside, they did not provide an explicit criterion for #W[1]-hardness that allows to pin down the complexity of #IndSub(Φ), given a concrete property Φ.

However, combining the framework of [5] with tools from the “topological approach to evasiveness” by Kahn, Saks and Sturtevant [13], two of the authors of the current paper established #W[1]-hardness for a wide range of properties, including, for example, all non-trivial properties that are closed under the removal of edges and false on odd cycles [20]. Taken together, the above results suggest the following conjecture.

▶ **Conjecture 1.** *Let Φ be a computable graph property satisfying that there are infinitely many positive integers k such that Φ is neither true nor false on all graphs with k vertices. Then $\#\textsc{IndSub}(\Phi)$ is $\#\mathsf{W}[1]$-hard.*

Unfortunately, a proof of this conjecture seems to be a long way off. In this work however, building up on [5, 20], we introduce an algebraic approach that allows us to resolve the above conjecture in case of *all* non-trivial monotone properties on bipartite graphs. In particular, we obtain a matching lower bound under the Exponential Time Hypothesis.

Results and techniques

We call a graph property *monotone* if it is closed under the removal of vertices and edges and *edge-monotone* if it is closed under the removal of edges only. Furthermore, we write IS_k for the graph consisting of k isolated vertices and $K_{t,t}$ for the complete bipartite graph with t vertices on each side. Our main theorems read as follows.

▶ **Theorem 2.** *Let Φ be a computable graph property and let $\mathcal{K}$ be the set of all prime powers t such that $\Phi(\mathsf{IS}_{2t}) \neq \Phi(K_{t,t})$. If $\mathcal{K}$ is infinite then $\#\textsc{IndSub}(\Phi)$ is $\#\mathsf{W}[1]$ hard. If additionally $\mathcal{K}$ is dense then it cannot be solved in time $f(k) \cdot n^{o(k)}$ for any computable function f unless ETH fails. This holds true even if the input graphs to $\#\textsc{IndSub}(\Phi)$ are restricted to be bipartite.*

In the previous theorem, a set $\mathcal{K}$ is *dense* if there exists a constant c such that for every $m \in \mathbb{N}$, there exists a $k \in \mathcal{K}$ such that $m \leq k \leq cm$. While the hypotheses of Theorem 2 sound technical, the theorem applies in many situations. In particular, it is applicable to properties that are neither (edge-) monotone nor the complement thereof: Let Φ be the property of being Eulerian. The graph $K_{t,t}$ contains an Eulerian cycle if $t = 2^s$ for $s \geq 1$. Hence we can apply Theorem 2 with $\mathcal{K} = \{2^s \mid s \geq 1\}$, which is infinite and dense.

▶ **Corollary 3.** *Let Φ be the property of being Eulerian. Then $\#\textsc{IndSub}(\Phi)$ is $\#\mathsf{W}[1]$-hard and cannot be solved in time $f(k) \cdot n^{o(k)}$ for any computable function f unless the ETH fails. This holds true even if the input graphs to $\#\textsc{IndSub}(\Phi)$ are restricted to be bipartite.*

In case Φ is edge-monotone, the condition $\Phi(\mathsf{IS}_{2t}) \neq \Phi(K_{t,t})$ is equivalent to non-triviality and if Φ is monotone, we obtain the following, more concise statement of the hardness result.

▶ **Theorem 4.** *Let Φ be a computable monotone graph property such that Φ and $\neg\Phi$ hold on infinitely many bipartite graphs. Then $\#\textsc{IndSub}(\Phi)$ is $\#\mathsf{W}[1]$-hard and cannot be solved in time $f(k) \cdot n^{o(k)}$ for any computable function f unless the Exponential Time Hypothesis fails. This holds true even if the input graphs to $\#\textsc{IndSub}(\Phi)$ are restricted to be bipartite.*

Let us illustrate further consequences of the previous theorems with respect to (edge-) monotone properties. First of all, most of the prior hardness results ([11, 10, 17, 12, 20]) are shown to hold in the restricted case of bipartite graphs. We provide three examples:

▶ **Corollary 5.** *The problem $\#\textsc{IndSub}(\Phi)$, restricted to* bipartite input graphs*, is $\#\mathsf{W}[1]$-hard and cannot be solved in time $f(k) \cdot |V(G)|^{o(k)}$ for any computable function f unless ETH fails, if Φ is one of the properties of being disconnected, planar or non-hamiltonian.*

One example of a monotone property Φ for which the complexity of #IndSub(Φ) was unknown, even for general graphs, is given by the following corollary of Theorem 4.

▶ **Corollary 6.** *Let F be a fixed bipartite graph with at least one edge and define $\Phi(G) = 1$ if G does not contain a subgraph isomorphic to F. Then #IndSub(Φ) is #W[1]-hard and cannot be solved in time $f(k) \cdot |V(G)|^{o(k)}$ for any computable function f unless ETH fails. This holds true even if the input graphs of #IndSub(Φ) are restricted to be bipartite.*

As the number of induced subgraphs of size k that satisfy Φ equals $\binom{|V(G)|}{k}$ minus the number of induced subgraphs of size k that satisfy $\neg\Phi$, *all of the previous result remain true for the complementary properties* $\neg\Phi$.

In proving the previous theorems we build up on the approach in [5, 20], where it was shown that, given a graph property Φ and a positive integer k, the number of induced subgraphs of size k in a graph G that satisfy Φ can equivalently be expressed as the following sum over all (isomorphism types of) graphs H:

$$\sum_{H} a_\Phi(H) \cdot \#\mathsf{Hom}(H \to G)\,, \tag{1}$$

where a_Φ is a function from graphs to integers with finite support and $\#\mathsf{Hom}(H \to G)$ is the number of graph homomorphisms from H to G. It is known that computing a linear combination of homomorphism numbers, as in the above expression, is *precisely as hard as* computing its hardest term with a non-zero coefficient ([5], also implicitly proved in [2]). We refer to this property as *complexity monotonicity*. In [20] two of the authors of the current paper used a topological approach to analyze the coefficient $a_\Phi(K_k)$ of the complete graph on k vertices. If this coefficient is non-zero then complexity monotonicity implies that computing the number of induced subgraphs of size k in a graph G that satisfy Φ is at least as hard as computing the number $\#\mathsf{Hom}(K_k \to G)$. This, in turn, is equivalent to computing the number of cliques of size k in G, the canonical #W[1]-complete problem [8]. While this approach led to hardness proofs for a wide range of properties Φ, it seems that resolving Conjecture 1, even restricted to monotone properties, requires a significant amount of new ideas. Without going too much into the details[1] of [20], our analysis of $a_\Phi(K_k)$ is complicated by the fact that the number of edges of the complete graph on $k \geq 4$ vertices is not a prime power. In this work, we hence focus on the coefficient of $a_\Phi(H)$ for graphs H that have a prime power number of edges and for which computing $\#\mathsf{Hom}(H \to G)$ is hard. One example of such graphs is the biclique $K_{t,t}$ for some prime power t. Here a biclique $K_{t,t}$, also called a complete bipartite graph, has t vertices on each side and contains every edge from a vertex on the left side to a vertex to the right side. Hence the number of edges is t^2 which is a prime power if t is.

In analyzing the coefficient $a_\Phi(K_{t,t})$ of the complete bipartite graph, we invoke the results of Rivest and Vuillemin [19] who considered transitive boolean functions over a domain of prime power cardinality to resolve the asymptotic version of what is known as *Karp's evasiveness conjecture* (we recommend Miller's survey [18] for an excellent overview).

[1] Readers familiar with [20] might recall that fixed points of group actions have been used to derive a simpler formula to compute the number $a_\Phi(K_t)$ modulo a prime p for positive powers t of p. This formula would simplify greatly if the group had a p-power number of elements *and* acted transitively on the edges of K_t. Unfortunately, this can never happen for $t \geq 4$, since the number of edges of K_t is not itself a p-power.

Given a property Φ and a graph H, the *alternating enumerator* of Φ and H is defined to be

$$\hat{\chi}(\Phi, H) := \sum_{S \subseteq E(H)} \Phi(H[S]) \cdot (-1)^{\#S},$$

where $H[S]$ is the graph with vertices $V(H)$ and edges S. Roughly speaking, it will turn out that the value of $a_\Phi(H)$ is closely related to $\hat{\chi}(\Phi, H)$. We furthermore point out that, in case Φ is closed under the removal of edges, the alternating enumerator $\hat{\chi}(\Phi, H)$ equals what is called the reduced Euler characteristic of the simplicial complex on $E(H)$ associated to Φ [18, 20]. In Section 3 we study the alternating enumerator in case of edge-transitive graphs, that is, graphs whose automorphism groups act transitively on the set of edges. We give a self-contained proof of the following fact, which implicitly follows from [19].

▶ **Lemma 7.** *Let Φ be a graph property and let H be an edge-transitive graph with p^k edges such that p is a prime and $\Phi(H[\emptyset]) \neq \Phi(H)$. Then it holds that $\hat{\chi}(\Phi, H) = (\pm 1) \mod p$.*

Now, intuitively, Lemma 7 induces a strategy towards proving hardness of $\#\textsc{IndSub}(\Phi)$: Assume a family of edge-transitive graphs $\mathcal{H}$ can be found such that $\#E(H)$ is a prime power and $\Phi(H[\emptyset]) \neq \Phi(H)$ for every $H \in \mathcal{H}$. Then $\#\textsc{IndSub}(\Phi)$ is at least as hard as counting homomorphisms from graphs in $\mathcal{H}$, the latter of which is fully understood [6]. This observation gives a strong motivation for the study of edge-transitive graphs with a prime power number of edges. In the second part of Section 3, we fully classify those graphs as subgraphs of bipartite graphs or vertex-transitive subgraphs of wreath graphs; consult Section 3 for the formal definitions. The proof of the following theorem, which might be of independent interest, relies on a non-trivial application of Sylow's theorems.

▶ **Theorem 8.** *Let G be a connected edge-transitive graph with p^t edges for some prime p and positive integer t. Then either G is bipartite or G is vertex-transitive and can be obtained from the wreath graph W_{p^k} for $k \geq 1$ by removing edges (or both).*

With the analysis of $\hat{\chi}$ and edge-transitive graphs completed, we turn to the reduction from counting homomorphisms in Section 4. More precisely, given a class $\mathcal{H}$ of edge-transitive graphs with a prime power number of edges and a graph property Φ such that for every $H \in \mathcal{H}$ we have that $\Phi(H[\emptyset]) \neq \Phi(H)$, we construct a parameterized Turing reduction from $\#\textsc{Hom}(\mathcal{H})$ to $\#\textsc{IndSub}(\Phi)$. Here, the problem $\#\textsc{Hom}(\mathcal{H})$ is defined as follows: Given as input a graph $H \in \mathcal{H}$ and a graph G, compute the number of homomorphisms from H to G. For technical reasons, we cannot immediately transform the number of induced subgraphs that satisfy Φ to a linear combination of homomorphism numbers as in Equation (1). We solve this technical issue by introducing color-prescribed variants of those problems in an intermediate step. In this context we consider H-colored graphs. Recall that a graph G is H-colored if it comes with a homomorphism c from G to H. A homomorphism from H to G is then called color-prescribed if it maps every vertex v of H to a vertex u of G satisfying that $c(u) = v$. We demonstrate that, given an H-colored graph G and oracle access to $\#\textsc{IndSub}(\Phi)$, the following linear combination can be computed in time $f(|V(H)|) \cdot |V(G)|^{O(1)}$.

$$\sum_{S \subseteq E(H)} \hat{a}_\Phi(S) \cdot \#\textsf{cp-Hom}(H[S] \to G). \tag{2}$$

Here $\textsf{cp-Hom}(H[S] \to G)$ denotes the set of color-prescribed homomorphisms from $H[S]$ to G and $\hat{a}_\Phi$ is a function of finite support only depending in Φ. In particular, $\hat{a}_\Phi(E(H))$ and $\hat{\chi}(\Phi, H)$ are proved to agree up to a factor of -1. Finally, we establish complexity monotonicity for linear combinations of color-prescribed homomorphisms as in Equation (2), which in combination with Lemma 7 yields the desired reduction.

Combining the previous results, we invoke the reduction on graph properties that are non-trivial on bipartite graphs and prove Theorem 2 and Theorem 4, in Section 5. Furthermore, we illustrate that our algebraic approach readily extends to modular counting by proving that both, Theorem 2 and Theorem 4 remain true in case counting is done modulo a fixed prime. Due to space constraints, the formal statement and proof of the modular counting version, as well as some proofs of Sections 3 and 4, are deferred to the full version.

2 Preliminaries

Given a positive integer k, we write $[k]$ for the set $\{1,\dots,k\}$ and given a set A we write $\binom{A}{k}$ for the set of all subsets of size k of A. Furthermore, assuming that A is finite, we write $\#A$ or $|A|$ for its cardinality. Given a function $g : A \times B \to C$ and an element $a \in A$, we write $g(a,\star)$ for the function which maps $b \in B$ to $g(a,b)$. Some of our proofs rely on (elementary) group theory; due to the space constraints we refrain from an introduction and refer the reader to e.g. Chapter 1 in the standard textbook of Lang [14].

2.1 Graph theory

Graphs in this work are considered simple, undirected and without self-loops. More precisely, a graph G is a pair of a finite set $V(G)$ of vertices and a symmetric and irreflexive relation $E(G) \subseteq V(G)^2$. If a graph H is obtained from G by deleting a set of edges and a set of vertices of G, including incident edges, then H is called a *subgraph* of G. Given a subset $\hat{V}$ of $V(G)$ we write $G[\hat{V}]$ for the graph with vertices $\hat{V}$ and edges $E \cap \hat{V}^2$. The resulting graph is called an *induced subgraph* of G. An *edge-subgraph* of a graph H is a graph obtained from H by deleting edges. Given a set $S \subseteq E(H)$ we write $H[S]$ for the edge-subgraph $(V(H),S)$ of H.

Homomorphisms and embeddings

A *homomorphism* from a graph H to a graph G is a mapping $h : V(H) \to V(G)$ that preserves adjacencies. In other words, for every edge $\{u,v\} \in E(H)$ it holds that $\{h(u),h(v)\} \in E(G)$. We write $\mathsf{Hom}(H \to G)$ for the set of all homomorphisms from H to G. A homomorphism inducing a bijection of vertices and satisfying $\{u,v\} \in E(H)$ if and only if $\{f(u),f(v)\} \in E(G)$ is called an *isomorphism* and we say that two graphs H and $\hat{H}$ are *isomorphic* if there exists an isomorphism from H to $\hat{H}$. We write $\mathsf{Sub}(H \to G)$ and $\mathsf{IndSub}(H \to G)$ for the sets of all subgraphs and induced subgraphs of G, respectively, that are isomorphic to H.

An isomorphism from a graph to itself is called an *automorphism*. The set of automorphisms of a graph, together with the operation of functional composition constitutes a group, called the *automorphism group* of a graph. Slightly abusing notation, we will write $\mathsf{Aut}(H)$ for both the set of automorphisms of a graph H as well as for the automorphism group of H.

An *embedding* is an injective homomorphism and we write $\mathsf{Emb}(H \to G)$ for the set of embeddings from H to G. If an embedding h from H to G additionally satisfies that $\{h(u),h(v)\} \in E(G)$ implies $\{u,v\} \in E(H)$, we call it a *strong embedding*. We write $\mathsf{StrEmb}(H \to G)$ for the set of strong embeddings from H to G. Observe that the images of embeddings and strong embeddings from H to G are precisely the subgraphs and induced subgraphs of G that are isomorphic to H.

Colored variants

Given graphs G and H, we say that G is *H-colored* if G comes with a homomorphism c from G to H, called an *H-coloring*. Note that, in particular, every edge-subgraph of H can be H-colored by the identity function on $V(H)$, which is assumed to be the given coloring whenever we consider H-colored edge-subgraphs of H in this paper. Given an edge-subgraph F of H and a homomorphism h from F to a H-colored graph G, we say that h is *color-prescribed* if for all $v \in V(F) = V(H)$ it holds that $c(h(v)) = v$. We write $\mathsf{cp\text{-}Hom}(F \to G)$ for the set of all color-prescribed homomorphisms from F to G. $\mathsf{cp\text{-}StrEmb}(F \to G)$ is defined similarly for color-prescribed strong embeddings. We point out that a definition of $\mathsf{cp\text{-}Emb}$ is obsolete as every color-prescribed homomorphism is injective by definition and hence an embedding. Furthermore, we write $\mathsf{cp\text{-}Sub}(F \to G)$ and $\mathsf{cp\text{-}IndSub}(F \to G)$ for the sets of images of color-prescribed embeddings and strong embeddings from F to G, respectively. Elements of $\mathsf{cp\text{-}Sub}(F \to G)$ and $\mathsf{cp\text{-}IndSub}(F \to G)$ are referred to as color-prescribed subgraphs and induced subgraphs.[2]

Graph properties and the alternating enumerator

A *graph property* is a function Φ from graphs to $\{0, 1\}$ such that for any pair of isomorphic graphs H and $\hat{H}$ we have that $\Phi(H) = \Phi(\hat{H})$. Adapting the notation of Rivest and Vuillemin [19], we define the *alternating enumerator* of a property Φ and a graph H to be the function

$$\hat{\chi}(\Phi, H) := \sum_{S \subseteq E(H)} \Phi(H[S]) \cdot (-1)^{\#S} .$$

A graph property Φ is called *edge-monotone* if it is closed under the removal of edges. It is called *monotone* if it is closed under the removal of edges *and* vertices.[3] Given a graph property Φ, a positive integer k and a graph G, we write $\mathsf{IndSub}(\Phi, k \to G)$ for the set of all induced subgraphs of size k of G that satisfy Φ. Furthermore, given a graph property Φ and an H-colored graph G, we write $\mathsf{cp\text{-}IndSub}(\Phi \to G)$ for the set of all color-prescribed induced subgraphs of size $|V(H)|$ in G that satisfy Φ.

2.2 Parameterized counting complexity

The field of parameterized counting was introduced independently by McCartin [16] and Flum and Grohe [8] and constitutes a hybrid of classical computational counting and parameterized complexity theory. A *parameterized counting problem* is a pair of a function $P : \Sigma^* \to \mathbb{N}$ and a computable parameterization $\kappa : \Sigma^* \to \mathbb{N}$. It is called *fixed-parameter tractable* (FPT) if there exists a computable function f and a deterministic algorithm that computes $P(x)$ in time $f(\kappa(x)) \cdot |x|^{O(1)}$ for every $x \in \Sigma^*$. A *parameterized Turing reduction* from (P, κ) to $(\hat{P}, \hat{\kappa})$ is a deterministic FPT algorithm with respect to κ that is given oracle access to $\hat{P}$ and that on input x computes $P(x)$ with the additional restriction that there exists a computable function g such that for any oracle query y it holds that $\hat{\kappa}(y) \le g(\kappa(x))$. We write $(P, \kappa) \le^{\mathrm{fpt}}_{\mathrm{T}} (\hat{P}, \hat{\kappa})$ if a parameterized Turing reduction exists.

[2] The observant reader might have noticed that the sets $\mathsf{cp\text{-}Sub}(F \to G)$ and $\mathsf{cp\text{-}Hom}(F \to G)$ as well as $\mathsf{cp\text{-}IndSub}(F \to G)$ and $\mathsf{cp\text{-}StrEmb}(F \to G)$ are essentially the same as a color-prescribed homomorphism is uniquely identified by its image. However, we decided to distinguish those notions in order to make the combinatorial arguments in Section 4 more accessible.

[3] To avoid confusion, we remark that in some literature, e.g. in [17] a property is called monotone if it is closed under *addition* of vertices and edges.

Given a graph G and a positive integer k, the parameterized counting problem #Clique asks to compute the number of complete subgraphs of size k in G and is parameterized by k, that is $\kappa(G, k) := k$. It is complete for the class #W[1], which can be seen as a parameterized counting equivalent of NP [8]. Evidence for the fixed-parameter intractability of #W[1]-hard problems is given by the *Exponential Time Hypothesis* (ETH), which asserts that 3-SAT cannot be solved[4] in time $\exp(o(m))$ where m is the number of clauses of the input formula. Assuming ETH, #Clique cannot be solved in time $f(k) \cdot n^{o(k)}$ for any function f [3, 4] and hence #W[1]-hard problems are not fixed-parameter tractable.

Given a recursively enumerable class of graphs $\mathcal{H}$, the problem #Hom($\mathcal{H}$) asks, given a graph $H \in \mathcal{H}$ and an arbitrary graph G, to compute $\#\mathsf{Hom}(H \to G)$. Its parameterization is given by $\kappa(H, G) := |V(H)|$. The problems #cp-Hom($\mathcal{H}$) and #cp-IndSub($\mathcal{H}$) are defined similarly. Note that the inputs of the latter two problems are of the form (H, G) where $H \in \mathcal{H}$ and G comes with an explicitly given H-coloring.

Given a computable graph property Φ, the problem #IndSub(Φ) asks, given a graph G and a positive integer k, to compute $\#\mathsf{IndSub}(\Phi, k \to G)$ and the parameterization is given by $\kappa(G, k) := k$. Furthermore, we define #cp-IndSub(Φ) to be the problem of, given a graph G that is H-colored for some graph H, computing $\#\mathsf{cp\text{-}IndSub}(\Phi \to G)$ and parameterize it by $\kappa(G) := |V(H)|$. We emphasize that, similarly to #cp-Hom($\mathcal{H}$), the input graph G comes with an explicitly given H-coloring, from which H can be constructed and thus the parameterization is well-defined.

3 Alternating enumerators and p-edge-transitive graphs

In this part of the paper we will provide a rough exposition of the work of Rivest and Vuillemin [19] who studied transitive boolean functions to resolve the asymptotic version of Karp's evasiveness conjecture. We will then apply their result to graphs H that are both edge-transitive and have p^ℓ many edges for some prime p. This will enable us to conclude that the alternating enumerator of Φ and H is (± 1) modulo p whenever $\Phi(H[\emptyset]) \neq \Phi(H)$. We start by introducing some required notions from algebraic graph theory.
The automorphism group of a graph H induces a group action on the edges of H, given by $h\{u, v\} := \{h(u), h(v)\}$. A group action is *transitive* if there exists only one orbit and a graph H is called *edge-transitive* if the group action on the edges is transitive, that is, if for every pair of edges $\{u, v\}$ and $\{\hat{u}, \hat{v}\}$ there exists an automorphism $h \in \mathsf{Aut}(H)$ such that $h\{u, v\} = \{\hat{u}, \hat{v}\}$. If additionally the number of edges of an edge-transitive graph is a prime power p^ℓ we call the graph *p-edge-transitive*.

▶ **Lemma 9** (Lemma 7 restated). *Let Φ be a graph property and let H be a p-edge-transitive graph such that $\Phi(H[\emptyset]) \neq \Phi(H)$. Then it holds that $\hat{\chi}(\Phi, H) = (\pm 1) \mod p$.*

Lemma 9 is implicitly proven in [19, Theorem 4.3], but for completeness we will include a short and self-contained proof, demonstrating a first application of the machinery of Sylow subgroups that we will need later.

For the proofs in this section, let us recall some key results from group theory. Given a prime number p, a finite group Γ' is called a *p-group* if the order $\#\Gamma'$ is a power of p. The following is a well-known and central result from the theory of finite groups.

[4] We point out that this includes deterministic and randomized algorithms.

▶ **Theorem 10** (Sylow theorems)**.** *Let Γ be a finite group of order $\#\Gamma = p^k m$ for a prime p and an integer $m \geq 1$ coprime to p. Then Γ contains a subgroup Γ' of order p^k. Moreover, every other subgroup Γ'' of Γ of order p^k is conjugate to Γ', that is there exists $g \in \Gamma$ with $\Gamma'' = g\Gamma' g^{-1}$. In particular, the groups Γ', Γ'' are isomorphic (via the conjugation by g).*

Finally, every subgroup $\tilde{\Gamma} \subseteq \Gamma$ which is a p-group is actually contained in some conjugate $g\Gamma' g^{-1}$ of the group Γ'.

A subgroup $\Gamma' \subseteq \Gamma$ as above is called a *p-Sylow subgroup* of Γ.

The following result is a first important application of the Sylow theorems. It can be found as Exercise (E28) in [1]; for completeness we include a proof in the full version.

▶ **Lemma 11.** *Let Γ be a finite group acting transitively on a set T such that $\#T = p^l$ for some $l \geq 0$. Then the induced action of any p-Sylow subgroup $\Gamma' \subseteq \Gamma$ on T is still transitive.*

This result allows us to give a short proof of Lemma 9 above. We sketch the proof here and provide the details in the full version of the paper.

Proof sketch of Lemma 9. Let Γ' be a p-Sylow subgroup of $\mathsf{Aut}(H)$, then by Lemma 11 it acts transitively on $E(H)$. This action on the edges of H induces an action on the set of subsets $S \subseteq E(H)$ and by the Orbit-Stabilizer Theorem, for any S which is not invariant under Γ', the size of its orbit by Γ' is a positive power of p. Then in the sum

$$\hat{\chi}(\Phi, H) = \sum_{S \subseteq E(H)} \Phi(H[S]) \cdot (-1)^{\#S},$$

we group together summands belonging to S in the same Γ'-orbit. The contribution of any orbit of positive size is divisible by p and can be left out modulo p. Since Γ' acts transitively on $E(H)$, the only invariant sets S are $S = \emptyset$ and $S = E(H)$, so we have

$$\hat{\chi}(\Phi, H) = \Phi(H[\emptyset]) + \Phi(H[E(H)]) \cdot (-1)^{\#E(H)} = \Phi(H[\emptyset]) - \Phi(H) \mod p.$$

Note that we use the fact that for $p > 2$ we have that $\#E(H)$ is odd since it is a prime power and for $p = 2$ we have $-1 = 1$ modulo p. Now, the condition $\Phi(H[\emptyset]) \neq \Phi(H)$ exactly gives us $\Phi(H[\emptyset]) - \Phi(H) = \pm 1 \mod p$. ◀

There are two main examples for p-edge-transitive graphs. The first example is the class of the complete, bipartite graphs K_{p^l,p^m} with $l, m \geq 0$. The graph K_{p^l,p^m} has p^{l+m} edges and the automorphism group clearly acts transitively on the edges of that graph. The second example is the class of wreath graphs W_{p^k} for $k \geq 1$. The graph W_{p^k} has p^k vertices that can be decomposed in disjoint sets $V_0, \ldots, V_{p-1}$ of order p^{k-1} each, and edges $\{v_i, v_{i+1}\}$ for each $i = 0, \ldots, p-1$ and vertices $v_i \in V_i, v_{i+1} \in V_{i+1}$ (where it is understood that $V_p = V_0$). Thus in total, W_{p^k} has p^{2k-1} edges, except for $p = 2$ where it has 2^{2k-2} edges. The graph W_{p^k} can be seen as the lexicographical product of a p-cycle with a graph consisting of p^{k-1} disjoint vertices. For $k = 1$ we exactly obtain the p-cycle. To see that W_{p^k} is edge-transitive, we observe that on the one hand, for fixed i we can apply an arbitrary permutation on V_i leaving the graph invariant. On the other hand, there exists a "rotational action" sending V_j to V_{j+1} for $j = 0, \ldots, p-1$, which also leaves the graph invariant. Using these two types of automorphisms, we can map every edge to every other edge.

The following result tells us that in a certain sense the graphs K_{p^l,p^m} and W_{p^k} are the maximal p-edge-transitive graphs. A graph G is called *vertex-transitive* if its automorphism group $\mathsf{Aut}(G)$ acts transitively on its set of vertices $V(G)$.

▶ **Theorem 12** (Theorem 8 restated). *Let G be a connected p-edge-transitive graph. Then either G is bipartite (and thus a subgraph of a graph of the form K_{p^l,p^m} for some $l,m \geq 0$) or G is vertex-transitive and an edge-subgraph of W_{p^k} for $k \geq 1$ (or both).*

Due to the space constraints the proof is deferred to the full version of the paper.

4 The main reduction: From homomorphisms to induced subgraphs

In what follows we will construct a sequence of reductions, starting from $\#\textsc{Hom}(\mathcal{H})$ and ending in $\#\textsc{IndSub}(\Phi)$. Here, $\mathcal{H}$ is a recursively enumerable set of p-edge-transitive graphs and Φ is a graph property such that for every graph $H \in \mathcal{H}$ we have that $\Phi(H[\emptyset]) \neq \Phi(H)$. More precisely, we will prove that

$$\#\textsc{Hom}(\mathcal{H}) \leq^{\mathrm{fpt}}_{\mathrm{T}} \#\textsc{cp-Hom}(\mathcal{H}) \overset{\text{Lemma 17}}{\leq^{\mathrm{fpt}}_{\mathrm{T}}} \#\textsc{cp-IndSub}(\Phi) \leq^{\mathrm{fpt}}_{\mathrm{T}} \#\textsc{IndSub}(\Phi) \quad (3)$$

In particular, all of those reductions will be tight in the sense that conditional lower bounds on the fine-grained complexity of $\#\textsc{Hom}(\mathcal{H})$ immediately transfer to $\#\textsc{IndSub}(\Phi)$. For the hardness results we rely on a result of Dalmau and Jonsson [6] stating that the problem $\#\textsc{Hom}(\mathcal{H})$ is known to be #W[1]-hard whenever $\mathcal{H}$ is recursively enumerable and of unbounded treewidth.[5] Here a class of graphs is said to have unbounded treewidth if for every $b \in \mathbb{N}$ there exists a graph in the class with treewidth at least b. Due to space constraints, the remainder of this section is concerned with proving Lemma 17, that is, the second step of the reduction sequence; the first and the third step are deferred to the full version.

Reducing color-prescribed homomorphisms to color-prescribed induced subgraphs

The reduction from color-prescribed homomorphisms to color-prescribed induced subgraphs requires the introduction of an H-colored variant of the framework of graph motif parameters, which was explicitly introduced in [5] and implicitly used in [2]. More precisely, given an H-colored graph G and a property Φ, we will express $\#\mathsf{cp\text{-}IndSub}(\Phi \to G)$ as a linear combination of color-prescribed homomorphisms, that is, terms of the form $\#\mathsf{cp\text{-}Hom}(H[S] \to G)$. In a first step, we show complexity monotonicity for linear combinations of color-prescribed homomorphisms. While this property allows a quite simple proof, a second step, in which we study the coefficient of $\#\mathsf{cp\text{-}Hom}(H \to G)$ requires a thorough understanding of the alternating enumerator of Φ and H. In case of p-edge-transitive graphs, the latter is provided by Lemma 9.

We start by introducing a colored variant of the tensor product of graphs (see e.g. Chapter 5.4.2 in [15]). Given two H-colored graphs G and $\hat{G}$ with colorings c and $\hat{c}$ we define their *color-prescribed* tensor product $G \times_H \hat{G}$ as the graph with vertices $V = \{(v,\hat{v}) \in V(G) \times V(\hat{G}) \mid c(v) = \hat{c}(\hat{v})\}$ and edges between two vertices $(v,\hat{v})$ and $(u,\hat{u})$ if and only if $\{v,u\} \in E(G)$ and $\{\hat{v},\hat{u}\} \in E(\hat{G})$. The next lemma states that $\#\mathsf{cp\text{-}Hom}$ is linear with respect to $\times_H$, a short proof of which can be found in the full version of the paper.

▶ **Lemma 13.** *Let H be a graph, let F be an edge-subgraph of H and let G and $\hat{G}$ be H-colored. Then we have that*

$$\#\mathsf{cp\text{-}Hom}(F \to G \times_H \hat{G}) = \#\mathsf{cp\text{-}Hom}(F \to G) \cdot \#\mathsf{cp\text{-}Hom}(F \to \hat{G})\,.$$

[5] We remark that the graph parameter of treewidth is not used explicitly in this work. Hence we omit the definition and refer the interested reader e.g. to Chapter 11 in [9].

We are now prepared to prove the color-prescribed variant of complexity monotonicity.

▶ **Lemma 14** (Complexity monotonicity). *Let H be a graph and let a be a function from edge-subgraphs of H to rationals. There exists an algorithm $\mathbb{A}$ that is given an H-colored graph G as input and has oracle access to the function*

$$\sum_{S\subseteq E(H)} a(H[S])\cdot \#\mathsf{cp\text{-}Hom}(H[S]\to\star)\,,$$

and computes $\#\mathsf{cp\text{-}Hom}(H[S]\to G)$ for all S such that $a(H[S])\neq 0$ in time $f(|H|)\cdot|V(G)|$ where f is a computable function. Furthermore, every oracle query $\hat{G}$ satisfies $|V(\hat{G})|\leq f(|H|)\cdot|V(G)|$.

Proof. Using Lemma 13 we have that for every H-colored graph F it holds that

$$\sum_{S\subseteq E(H)} a(H[S])\cdot \#\mathsf{cp\text{-}Hom}(H[S]\to(G\times_H F)) \quad (4)$$

$$=\sum_{S\subseteq E(H)} a(H[S])\cdot \#\mathsf{cp\text{-}Hom}(H[S]\to G)\cdot\#\mathsf{cp\text{-}Hom}(H[S]\to F)\,, \quad (5)$$

which we can evaluate for $F=H[\emptyset],\dots,H[E(H)]$. This induces a system of linear equations which can easily be shown to have a unique solution; the proof that the corresponding matrix is non-singular can be found in the full version of the paper. Consequently, the numbers $a(H[S])\cdot\#\mathsf{cp\text{-}Hom}(H[S]\to G)$ are uniquely determined and can be computed by solving the system using Gaussian elimination. Finally, we obtain the numbers $\#\mathsf{cp\text{-}Hom}(H[S]\to G)$ by multiplying with $a(H[S])^{-1}$ whenever $a(H[S])\neq 0$. ◀

It remains to express the number of color-prescribed induced subgraphs that satisfy a property Φ as a linear combination of color-prescribed homomorphisms. We only sketch the proof of the following lemma and defer the details to the full version of the paper.

▶ **Lemma 15.** *Let H be a graph, let Φ be a graph property and let G be an H-colored graph. Then it holds that*

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

Moreover, the absolute values of the coefficient of $\#\mathsf{cp\text{-}Hom}(H\to G)$ and $\hat{\chi}(\Phi,H)$ are equal.

Proof sketch. We rely on the following claim which follows by inclusion-exclusion.

▷ Claim 16. Let H be graph, let $S\subseteq E(H)$ and let G be an H-colored graph. Then we have that

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

Now summing up over all S for which $\Phi(H[S])=1$ and applying Claim 16 yields

$$\#\mathsf{cp\text{-}IndSub}(\Phi\to G)=\sum_{S\subseteq E(H)}\Phi(H[S])\sum_{J\subseteq E(H)\setminus S}(-1)^{\#J}\cdot\#\mathsf{cp\text{-}Hom}(H[S\cup J]\to G)\,. \quad (6)$$

Finally, we collect for the coefficient of $\#\mathsf{cp\text{-}Hom}(H\to G)$ and obtain

$$\sum_{S\subseteq E(H)}\Phi(H[S])\cdot(-1)^{\#E(H)-\#S}=(-1)^{\#E(H)}\cdot\hat{\chi}(\Phi,H)\,. \quad (7)$$

◀

The application of the complexity monotonicity property for color-prescribed homomorphisms (Lemma 14) requires non-zero coefficients. However, this can be guaranteed for the coefficient of interest in case of p-edge-transitive graphs as shown in Section 3. Formally, the reduction is constructed as follows.

▶ **Lemma 17.** *Let Φ be a graph property and let H be a p-edge-transitive graph such that $\Phi(H[\emptyset]) \neq \Phi(H)$. There exists an algorithm $\mathbb{A}$ that is given an H-colored graph G as input and has oracle access to the function $\#\mathsf{cp\text{-}IndSub}(\Phi \to \star)$ and computes $\#\mathsf{cp\text{-}Hom}(H \to G)$ in time $f(|H|) \cdot |V(G)|$ where f is a computable function. Furthermore, every oracle query $\hat{G}$ is H-colored as well and satisfies $|V(\hat{G})| \leq f(|H|) \cdot |V(G)|$.*

Proof. Using Lemma 15 we can express $\#\mathsf{cp\text{-}IndSub}(\Phi \to \star)$ as a linear combination of color-prescribed homomorphisms. In particular, the coefficient of $\#\mathsf{cp\text{-}Hom}(H \to \star)$ is $(\pm 1) \cdot \hat{\chi}(\Phi, H)$ and by Lemma 9 we have that this number is non-zero whenever H is p-edge-transitive and $\Phi(H[\emptyset]) \neq \Phi(H)$. Hence we can use the algorithm from Lemma 14 to compute $\#\mathsf{cp\text{-}Hom}(H \to G)$ in the desired running time. ◀

5 Non-trivial monotone properties on bipartite graphs

In the last part of the paper, we apply the algebraic approach which was laid out in the preceding sections to bipartite graph properties. This will allow us to prove our main result. To this end, we say that a set $\mathcal{K} \subseteq \mathbb{N}$ is *dense* if there exists a constant c such that for every $k' \in \mathbb{N}$ there exists $k \in \mathcal{K}$ such that $k' \leq k \leq ck'$. Furthermore, we write IS_k for the graph with k isolated vertices. The following theorem is obtained by invoking the reduction sequence (3) to complete bipartite graphs $K_{t,t}$ for prime powers $t = p^k$, which are p-edge-transitive (see Section 3). Due to the space constraints, the details, as well as the case of modular counting, are deferred to the full version of the paper.

▶ **Theorem 18** (Theorem 2 restated). *Let Φ be a computable graph property and let $\mathcal{K}$ be the set of all prime powers t such that $\Phi(\mathsf{IS}_{2t}) \neq \Phi(K_{t,t})$. If $\mathcal{K}$ is infinite then $\#\textsc{IndSub}(\Phi)$ is $\#\mathsf{W}[1]$ hard. If additionally $\mathcal{K}$ is dense then it cannot be solved in time $f(k) \cdot n^{o(k)}$ for any computable function f unless ETH fails. This holds true even if the input graphs to $\#\textsc{IndSub}(\Phi)$ are restricted to be bipartite.*

Note that, in case Φ or its complement is edge-monotone, we only have to find infinitely many prime powers t for which Φ is neither true nor false on the set of all edge-subgraphs of $K_{t,t}$, which is the case for all sensible, non-trivial properties that do not rely on the number of vertices in some way. If Φ (or its complement) is monotone, that is, not only closed under the removal of edges, but also under the removal of vertices, then such artificial properties do not exist and we can state the result more clearly as follows.

▶ **Corollary 19** (Theorem 4 restated). *Let Φ be a computable monotone graph property such that Φ and $\neg\Phi$ hold on infinitely many bipartite graphs. Then $\#\textsc{IndSub}(\Phi)$ is $\#\mathsf{W}[1]$-hard and cannot be solved in time $f(k) \cdot n^{o(k)}$ for any computable function f unless ETH fails. This holds true even if the input graphs to $\#\textsc{IndSub}(\Phi)$ are restricted to be bipartite.*

Proof. If Φ is monotone and Φ and $\neg\Phi$ hold on infinitely many bipartite graphs, then $\Phi(\mathsf{IS}_k) = 1$ for all positive integers k and $\Phi(K_{t,t}) = 0$ for all but finitely many t. Hence we can apply Theorem 18 and, in particular, the set $\mathcal{K}$ will contain all but finitely many prime powers and is therefore dense. ◀

Conclusion

We have established hardness for #IndSub(Φ) for *any* (edge-)monotone property Φ that is non-trivial on bipartite graphs. In particular, this holds true even if we count modulo a prime and restrict the input graphs to be bipartite as well. Hence, we did not only significantly extend the set of graph properties Φ for which the (parameterized) complexity of #IndSub(Φ) is understood, but we also generalized many of the prior results, such as [11], [17] and parts of [20] to the cases of bipartite input graphs and modular counting.

As a next step towards a proof of Conjecture 1, we suggest the study of properties that are defined by forbidden induced subgraphs, for which the complexity of #IndSub(Φ) is only partially resolved at this point.

References

1. Shreeram S. Abhyankar. *Lectures on algebra. Vol. I.* World Scientific Publishing Co. Pte. Ltd., Hackensack, NJ, 2006. doi:10.1142/9789812773449.
2. Hubie Chen and Stefan Mengel. Counting Answers to Existential Positive Queries: A Complexity Classification. In *Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, PODS 2016, San Francisco, CA, USA, June 26 - July 01, 2016*, pages 315–326, 2016. doi:10.1145/2902251.2902279.
3. Jianer Chen, Benny Chor, Mike Fellows, Xiuzhen Huang, David W. Juedes, Iyad A. Kanj, and Ge Xia. Tight lower bounds for certain parameterized NP-hard problems. *Inf. Comput.*, 201(2):216–231, 2005. doi:10.1016/j.ic.2005.05.001.
4. Jianer Chen, Xiuzhen Huang, Iyad A. Kanj, and Ge Xia. Strong computational lower bounds via parameterized complexity. *J. Comput. Syst. Sci.*, 72(8):1346–1367, 2006. doi:10.1016/j.jcss.2006.04.007.
5. 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, 2017. doi:10.1145/3055399.3055502.
6. Víctor Dalmau and Peter Jonsson. The complexity of counting homomorphisms seen from the other side. *Theor. Comput. Sci.*, 329(1-3):315–323, 2004. doi:10.1016/j.tcs.2004.08.008.
7. Jack Edmonds. Paths, Trees, and Flowers. *Canadian Journal of Mathematics*, 17:449–467, 1965. doi:10.4153/CJM-1965-045-4.
8. Jörg Flum and Martin Grohe. The Parameterized Complexity of Counting Problems. *SIAM J. Comput.*, 33(4):892–922, 2004. doi:10.1137/S0097539703427203.
9. Jörg Flum and Martin Grohe. *Parameterized Complexity Theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2006. doi:10.1007/3-540-29953-X.
10. Mark Jerrum and Kitty Meeks. Some Hard Families of Parameterized Counting Problems. *TOCT*, 7(3):11:1–11:18, 2015. doi:10.1145/2786017.
11. Mark Jerrum and Kitty Meeks. The parameterised complexity of counting connected subgraphs and graph motifs. *J. Comput. Syst. Sci.*, 81(4):702–716, 2015. doi:10.1016/j.jcss.2014.11.015.
12. Mark Jerrum and Kitty Meeks. The parameterised complexity of counting even and odd induced subgraphs. *Combinatorica*, 37(5):965–990, 2017. doi:10.1007/s00493-016-3338-5.
13. Jeff Kahn, Michael E. Saks, and Dean Sturtevant. A topological approach to evasiveness. *Combinatorica*, 4(4):297–306, 1984. doi:10.1007/BF02579140.
14. Serge Lang. *Algebra*. Graduate Texts in Mathematics. Springer New York, 2005. doi:10.1007/978-1-4613-0041-0.
15. László Lovász. *Large Networks and Graph Limits*, volume 60 of *Colloquium Publications*. American Mathematical Society, 2012. URL: http://www.ams.org/bookstore-getitem/item=COLL-60.

16 Catherine McCartin. Parameterized counting problems. *Ann. Pure Appl. Logic*, 138(1-3):147–182, 2006. doi:10.1016/j.apal.2005.06.010.

17 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.

18 Carl A. Miller. Evasiveness of Graph Properties and Topological Fixed-Point Theorems. *Foundations and Trends in Theoretical Computer Science*, 7(4):337–415, 2013. doi:10.1561/0400000055.

19 Ronald L. Rivest and Jean Vuillemin. On recognizing graph properties from adjacency matrices. *Theoret. Comput. Sci.*, 3(3):371–384, 1976/77. doi:10.1016/0304-3975(76)90053-0.

20 Marc Roth and Johannes Schmitt. Counting Induced Subgraphs: A Topological Approach to #W[1]-hardness. In *13th International Symposium on Parameterized and Exact Computation, IPEC 2018, August 20-24, 2018, Helsinki, Finland*, pages 24:1–24:14, 2018. doi:10.4230/LIPIcs.IPEC.2018.24.

21 Seinosuke Toda. PP is as hard as the polynomial-time hierarchy. *SIAM J. Comput.*, 20(5):865–877, 1991. doi:10.1137/0220053.

22 Leslie G. Valiant. The Complexity of Computing the Permanent. *Theor. Comput. Sci.*, 8:189–201, 1979. doi:10.1016/0304-3975(79)90044-6.