Parameterized (Modular) Counting and Cayley Graph Expanders

Norbert Peyerimhoff
Department of Mathematical Sciences, Durham University, UK

Marc Roth
Merton College, University of Oxford, UK

Johannes Schmitt
Mathematical Institute, University of Bonn, Germany

Jakob Stix
Mathematical Institute, Goethe-Universität Frankfurt, Germany

Alina Vdovina
School of Mathematics and Statistics, Newcastle University, UK

Abstract

We study the problem #EdgeSub(Φ) of counting k-edge subgraphs satisfying a given graph property Φ in a large host graph G. Building upon the breakthrough result of Curticapean, Dell and Marx (STOC 17), we express the number of such subgraphs as a finite linear combination of graph homomorphism counts and derive the complexity of computing this number by studying its coefficients.

Our approach relies on novel constructions of low-degree Cayley graph expanders of p-groups, which might be of independent interest. The properties of those expanders allow us to analyse the coefficients in the aforementioned linear combinations over the field $\mathbb{F}_p$ which gives us significantly more control over the cancellation behaviour of the coefficients. Our main result is an exhaustive and fine-grained complexity classification of #EdgeSub(Φ) for minor-closed properties Φ, closing the missing gap in previous work by Roth, Schmitt and Wellnitz (ICALP 21).

Additionally, we observe that our methods also apply to modular counting. Among others, we obtain novel intractability results for the problems of counting k-forests and matroid bases modulo a prime p. Furthermore, from an algorithmic point of view, we construct algorithms for the problems of counting k-paths and k-cycles modulo 2 that outperform the best known algorithms for their non-modular counterparts.

In the course of our investigations we also provide an exhaustive parameterized complexity classification for the problem of counting graph homomorphisms modulo a prime p.

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 Cayley graphs, counting complexity, expander graphs, fine-grained complexity, parameterized complexity

Digital Object Identifier 10.4230/LIPIcs.MFCS.2021.84

Related Version *Full Version*: https://arxiv.org/abs/2104.14596

Funding *Johannes Schmitt*: The third author was supported by the SNF Early Postdoc.Mobility grant 184245 and thanks the Max Planck Institute for Mathematics in Bonn for its hospitality.

Acknowledgements The second author is grateful to Holger Dell for fruitful discussions on the connections between [10] and our work.

46th International Symposium on Mathematical Foundations of Computer Science (MFCS 2021).
Editors: Filippo Bonchi and Simon J. Puglisi; Article No. 84; pp. 84:1–84:15
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Extended Abstract

In this work we study the problem of counting small patterns in large host graphs. With applications in a diverse set of disciplines such as constraint satisfaction problems [13, 7], database theory [15, 8] and network science [34, 1, 41], it is unsurprising that this problem has received significant attention from the viewpoint of parameterized and fine-grained complexity theory in recent years [2, 17, 33, 9, 12, 11, 5, 37, 38, 28, 39].

We continue this line of work and study the problem of counting k-edge subgraphs that satisfy a graph property Φ: For any fixed Φ, the problem #EdgeSub(Φ) asks, on input a graph G and a positive integer k, to compute the number of (not necessarily induced) subgraphs with k edges in G that satisfy Φ. In particular, we focus on instances in which k is significantly smaller than G. Formally, we choose k to be the *parameter* of the problem and ask for which Φ there is a function f such that #EdgeSub(Φ) can be solved in time $f(k) \cdot |V(G)|^{O(1)}$; in this case we call the problem *fixed-parameter tractable* with respect to the parameter k.

If #EdgeSub(Φ) is not fixed-parameter tractable, it is desirable to improve the exponent of $|V(G)|$ in the running time as far as possible. For example, the best known algorithm for counting k-edge subgraphs [11] can be used to solve #EdgeSub(Φ) in time $f(k) \cdot |V(G)|^{0.174k+o(k)}$ [39]. Additionally, it was shown in recent work that #EdgeSub(Φ) is fixed-parameter tractable whenever Φ has *bounded matching number*, that is, whenever there is a constant upper bound on the size of the largest matching of any graph satisfying Φ [39]. If, for each k, the property Φ is true for only one graph on k edges, then the previous fixed-parameter tractability result is best possible: In this case, #EdgeSub(Φ) becomes an instance of the counting version of the parameterized subgraph isomorphism problem which has been fully classified by Curticapean and Marx [12].

However, for arbitrary Φ, much less is known about the complexity of #EdgeSub(Φ). In [39], two of the authors, together with Wellnitz, presented first results for more general properties such as connectivity, Eulerianity and, in particular, an *almost* exhaustive classification for minor-closed properties Φ, leaving (partially) open the case of forbidden minors of degree at most 2. In this work, we close this gap and provide a full dichotomy result:

▶ **Theorem 1.** *Let Φ be a minor-closed graph property. If Φ is trivially true or of bounded matching number, then #EdgeSub(Φ) is fixed-parameter tractable. Otherwise, #EdgeSub(Φ) is #W[1]-hard and, assuming the Exponential Time Hypothesis, it cannot be solved in time*

$$f(k) \cdot |G|^{o(k/\log k)}$$

for any function f.

Here, #W[1] is the parameterized counting analogue of NP; a formal definition is provided in the full version. Particular cases for which we obtain novel intractability results are given by the following (minor-closed) properties; the formal intractability results are stated and proved in the full version as well.

- $\Phi(H) = 1$ if H is a forest.
- $\Phi(H) = 1$ if H is a linear forest.
- $\Phi(H) = 1$ if the tree-depth of H is bounded by a constant.
- $\Phi(H) = 1$ if the Colin de Verdière Invariant of H is bounded by a constant.

Additionally, we investigate the property of being bipartite. For this case, we present not only a novel fine-grained lower bound, but also a #W[1]-hardness result, which was not known before.

▶ **Theorem 2.** *Let Φ be the property of being bipartite. Then $\#\text{EdgeSub}(\Phi)$ is $\#\mathsf{W}[1]$-hard and, assuming the Exponential Time Hypothesis, it cannot be solved in time*

$$f(k) \cdot |G|^{o(k/\log k)}$$

for any function f.

Our hardness results crucially rely on a novel construction of families of low-degree Cayley graph expanders of p-groups, which might be of independent interest. We will present the new Cayley graph expanders in the following theorem; their construction, as well as their role in the hardness proofs for $\#\text{EdgeSub}(\Phi)$ will be elaborated on in the technical discussion.

▶ **Theorem 3.** *Let $p \geq 3$ be a prime number, and $d \geq 2$ be an integer. We assume that $d \geq (p+3)/2$ if $p \geq 7$.*

Then there is an explicit construction of a sequence of finite p-groups Γ_i of orders that tend to infinity, with symmetric generating sets S_i of cardinality $2d$ such that the Cayley graphs $\mathcal{C}(\Gamma_i, S_i)$ form a family of expanders (of fixed valency $2d$ on a set of vertices of p-power orders and with vertex transitive automorphism groups).

Our methods do not only apply to exact counting, but also to modular pattern counting problems: Here the goal is to compute the number of occurrences of the pattern *modulo a fixed prime p*. In classical complexity theory, the study of modular counting problems has a rich history, such as the algorithm for computing the permanent modulo 2^ℓ [44], the so-called accidental algorithms [45], Toda's Theorem [43], classifications for modular #CSPs and Holants [24, 25] and the line of research on the modular homomorphism counting problem [16, 20, 21, 22, 26, 18, 27], only to name a few.

While results are scarcer, the *parameterized* complexity of modular (pattern) counting problems has also been studied in recent years [4, 14, 10], and we contribute to this line of research as follows: First, we provide a novel intractability result for modular counting of forests and matroid bases. We write $\#_p\text{Forests}$ for the problem of, given a graph G and a positive integer k, computing the number of forests with k edges in G, modulo p. Similarly, we write $\#_p\text{Bases}$ for the problem of, given a linear matroid M of rank k in matrix representation, computing the number of bases of M, modulo p; the parameter of both problems is given by k.

▶ **Theorem 4.** *For each prime $p \geq 3$, the problems $\#_p\text{Forests}$ and $\#_p\text{Bases}$ are $\mathsf{Mod}_p\mathsf{W}[1]$-hard and, assuming the randomised Exponential Time Hypothesis, cannot be solved in time $f(k) \cdot |G|^{o(k/\log k)}$ (resp. $f(k) \cdot |M|^{o(k/\log k)}$), for any function f.*

Here, $\mathsf{Mod}_p\mathsf{W}[1]$ is the parameterized modular counting version of NP. Roughly speaking, a problem is $\mathsf{Mod}_p\mathsf{W}[1]$ if it is at least as hard as counting k-cliques modulo p; we give a formal definition in the full version. Additionally, we provide an algorithmic result for counting k-paths and k-cycles modulo 2:

▶ **Theorem 5.** *The problems of counting k-paths and k-cycles in a graph G modulo 2 can be solved in time $k^{O(k)} \cdot |V(G)|^{k/6+O(1)}$.*

We emphasize that the algorithm in the previous theorem is faster than the best known algorithms for (non-modular) counting of k-cycles/k-paths, which run in time $k^{O(k)} \cdot |V(G)|^{13k/75+o(k)}$ [11]. Furthermore, it follows from a result by Curticapean, Dell and Husfeldt [10] that counting k-paths modulo 2 is $\mathsf{Mod_2W[1]}$-hard, implying that we cannot hope for an algorithm for k-paths running in time $f(k) \cdot |V(G)|^{O(1)}$. Finally, we study the parameterized complexity of counting homomorphisms modulo p. In the classical setting, the related problem of modular counting homomorphisms with *right-hand side* restrictions received much attention: For any fixed graph H, the problem $\#_p\textsc{HomsTo}(H)$[1] asks, on input a graph G, to compute the number of homomorphisms from G to H modulo p. Despite significant effort [16, 20, 21, 22, 26, 18, 27], the problem has not been fully classified for each graph H.[2]

In this work, we consider the related *left-hand side* version of the problem. Adapting the definitions of Grohe, Dalmau and Jonsson [23, 13] for detecting and exact counting of homomorphisms, we define a problem $\#_p\textsc{Hom}(\mathcal{H})$ for each class of graphs $\mathcal{H}$ and for each prime p: This problem expects as input a graph $H \in \mathcal{H}$ and an arbitrary graph G, and the goal is to compute the number of homomorphisms from H to G, modulo p. The problem is parameterized by the size of H, that is, we assume H to be significantly smaller than G.

It is known that the decision version $\textsc{Hom}(\mathcal{H})$ is fixed-parameter tractable (even solvable in polynomial time) if the treewidth of the cores of $\mathcal{H}$ is bounded by a constant, and $\mathsf{W[1]}$-hard otherwise [23]. Similarly, the (exact) counting version $\#\textsc{Hom}(\mathcal{H})$ is known to be fixed-parameter tractable (even solvable in polynomial time) if the treewidth of the graphs of $\mathcal{H}$ is bounded by a constant, and $\#\mathsf{W[1]}$-hard otherwise [13].

In case of counting modulo p, we establish an exhaustive classification for $\#_p\textsc{Hom}(\mathcal{H})$ along what we call the *p-reduced quotients* of the graphs in $\mathcal{H}$. Let H be a graph and let α be an automorphism of H of order p. Then we define the quotient graph H/α to have a vertex for each orbit of the action of α on $V(H)$, and two vertices corresponding to orbits $\mathcal{O}_1$ and $\mathcal{O}_2$ are made adjacent if and only if there are vertices $v_1 \in \mathcal{O}_1$ and $v_2 \in \mathcal{O}_2$ such that $\{v_1, v_2\} \in E(H)$. This induces a finite (possibly trivial) sequence $H = H_1, \dots, H_\ell$ where for $i = 1, \dots, \ell - 1$ we set $H_{i+1} = H_i/\alpha_i$ for some automorphism α_i of order p of H_i and where the last graph H_ℓ does not have an automorphism of order p. Then $H_p^* := H_\ell$ is called the p-reduced quotient of H.[3] We will see that H_p^* is well-defined by proving that each of the aforementioned sequences yields the same graph, up to isomorphism.

Let us now state our classification for $\#_p\textsc{Hom}(\mathcal{H})$. In what follows, given a class $\mathcal{H}$, we write $\mathcal{H}_p^*$ for the p-reduced quotients without self-loops of graphs in $\mathcal{H}$. We first present the algorithmic part:

▶ **Theorem 6.** *Let $p \geq 2$ be a prime and let $\mathcal{H}$ be a class of graphs. The problem $\#_p\textsc{Hom}(\mathcal{H})$ can be solved in time*

$$\exp(\mathsf{poly}(|V(H)|)) \cdot |V(G)|^{\mathsf{tw}(H_p^*)+O(1)}\,.$$

In particular, $\#_p\textsc{Hom}(\mathcal{H})$ is fixed-parameter tractable if the treewidth of $\mathcal{H}_p^*$ is bounded.

Here tw denotes treewidth.

---

[1] For $p = 2$, the problem is usually denoted by $\oplus\textsc{HomsTo}(H)$.
[2] However, very recently a full classification was announced by Bulatov and Kazeminia [6].
[3] We remark that H_p^* is related to the notion of involution-free reductions used in the analysis of the right-hand side version of the problem [16, 22]. However, the difference is that the p-reduced quotient identifies non-fixed points of an order-p automorphism by including a vertex for each orbit, while the involution-free reduction just deletes all non-fixed points.

▶ Remark 7. Using the quasi-polynomial time algorithm for GI due to Babai [3], we will also show how the algorithm in the previous theorem can be improved to run in quasi-polynomial time in the input length $|V(H)| + |V(G)|$. Additionally, proving that the construction of the p-reduced quotient is at least as hard as the graph automorphism problem, we observe that a polynomial-time algorithm is unlikely, unless the construction of the p-reduced quotients can be avoided.

For the intractability part of our classification, we show that unbounded treewidth of $\mathcal{H}_p^*$ yields hardness:

▶ **Theorem 8.** *Let $p \geq 2$ be a prime and let $\mathcal{H}$ be a computable class of graphs. If the treewidth of $\mathcal{H}_p^*$ is unbounded, then $\#_p\textsc{Hom}(\mathcal{H})$ is $\mathsf{Mod}_p\mathsf{W}[1]$-hard and, assuming the randomised Exponential Time Hypothesis, cannot be solved in time*

$$f(|H|) \cdot |G|^{o(\mathsf{tw}(H_p^*)/\log \mathsf{tw}(H_p^*))}$$

for any function f.

Can You Beat Treewidth?

We conclude the presentation of our results by commenting on the factor of $1/(\log \dots)$ in the exponents of all of our fine-grained lower bounds. This factor is related to the conjecture of whether it is possible to "beat treewidth" [31]. In particular, we point out that the factor can be dropped in *all* of our lower bounds if this conjecture, formally stated as Conjecture 1.3 in [32], is true.

Technical Overview

Our central approach follows the so-called Complexity Monotonicity framework due to Curticapean, Dell and Marx [11]. We express the counting problems considered in this work as formal linear combination of homomorphism counts, which allows us to derive the complexity of the problem at hand by analyzing the coefficients.

More precisely, let us fix a graph property Φ and a positive integer k. Given a graph G, we furthermore write $\#\mathsf{EdgeSub}(\Phi, k \to G)$ for the number of k-edge subgraphs of G that satisfy Φ. It was shown in [39] that there exists a function of finite support $a_{\Phi,k}$ from graphs to rationals such that for every graph G we have

$$\#\mathsf{EdgeSub}(\Phi, k \to G) = \sum_H a_{\Phi,k}(H) \cdot \#\mathsf{Hom}(H \to G)\,, \tag{1}$$

where $\#\mathsf{Hom}(H \to G)$ is the number of graph homomorphisms from H to G. Curticapean, Dell and Marx [11] have shown that computing a linear combination as in (1) is *precisely as hard as* computing its hardest term. Fortunately, the complexity of counting and detecting homomorphisms from H to G is thoroughly classified [13, 31]: Roughly speaking, the higher the treewidth of H, the harder it is to compute the number of homomorphisms from H to G. Therefore, proving hardness of computing $\#\mathsf{EdgeSub}(\Phi, k \to G)$ reduces to the purely combinatorial problem of determining which of the coefficients $a_{\Phi,k}(H)$ in (1) for high-treewidth graphs H are non-zero.

Unfortunately, it has turned out that the coefficients of such linear combinations for related pattern counting problems are often determined by (or even equal to) a variety of algebraic and topological invariants, whose analysis is known to be a difficult problem in its own right. For example, in case of the vertex-induced subgraph counting problem, the

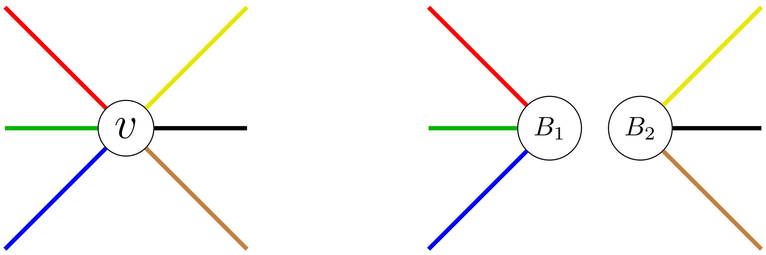

Figure 1 Illustration of the construction of a fractured graph. The left picture shows a vertex v of a graph H with incident edges $E_H(v) = \{$●, ●, ●, ●, ●, ●$\}$. The right picture shows the splitting of v in the construction of the fractured graph $H\sharp\sigma$ for a fracture σ satisfying that the partition σ_v contains two blocks $B_1 = \{$●, ●, ●$\}$, and $B_2 = \{$●, ●, ●$\}$.

coefficient of the clique is the reduced Euler characteristic of a simplicial graph complex [37], the coefficient of the biclique is the so-called alternating enumerator [14], and, more generally, the coefficients of dense graphs are related to the h-and f-vectors associated with the property of the patterns that are to be counted [38]. In all of the previous works mentioned here, the complexity analysis of the respective pattern counting problems therefore amounted to understanding the cancellation behaviour of those invariants. To do so, the papers used tools from combinatorial commutative algebra and, to some extent, topological fixed-point theorems.

In case of #EdgeSub(Φ), two of the authors, together with Wellnitz, observed that the coefficients of high-treewidth low-degree vertex-transitive graphs can be analysed much easier than generic graphs of high treewidth such as the clique or the biclique [39]. First, it was shown that the coefficient of a graph H with k edges in (1) is equal to the *indicator* of Φ and H, defined as follows:[4]

$$a(\Phi, H) := \sum_{\sigma \in \mathcal{L}(\Phi,H)} \prod_{v \in V(H)} (-1)^{|\sigma_v|-1}(|\sigma_v| - 1)!\,. \tag{2}$$

Here, $\mathcal{L}(\Phi, H)$ is the set of *fractures* σ of H such that the associated *fractured graph* $H\sharp\sigma$ satisfies Φ. Here, a fracture of a graph H is a tuple $\sigma = (\sigma_v)_{v\in V(H)}$, where σ_v is a partition of the set of edges $E_H(v)$ of H incident to v. Given a fracture ρ of H, the fractured graph $H\sharp\sigma$ is obtained from H be splitting each vertex $v \in V(H)$ according to σ_v; an illustration is provided in Figure 1.

As a consequence, the #W[1]-hardness results of Theorems 1 and 2 can be obtained if we find a family of graphs H of unbounded treewidth, such that $a(\Phi, H) \neq 0$ for infinitely many graphs H in this family. The almost tight conditional lower bound under the Exponential Time Hypothesis will, additionally, require sparsity of the graphs. In combination with the main observation in [39], stating that the indicator $a(\Phi, H)$ can be analysed much easier for vertex-transitive graphs, we propose that regular Cayley graph expanders are the right choice for the family of graphs to be considered. Indeed, those graphs are sparse, have high treewidth and are always vertex transitive. A particular family of Cayley graph expanders was already used in [39], but it turned out to be impossible to prove Theorems 1 and 2 relying only on this family of Cayley graph expanders; we discuss this in detail in the full version.

[4] To be precise, the identity in (2) was obtained for a coloured version of #EdgeSub(Φ). However, we will mostly rely on this result in a blackbox manner; all details of the coloured version necessary for the treatment in this paper will be carefully introduced when needed.

In this work, we therefore present novel constructions of families of low-degree Cayley graph expanders. Those will not only allow us to prove most of our main theorems by analysing their indicators, but might be of independent interest. For the sake of presentation, we decided to encapsulate the treatment of our constructions in separate sections, both in the extended abstract and the main part of the paper. We hope that this makes the paper accessible both for readers primarily interested in the novel construction of Cayley graph expanders, as well as for readers mainly interested in the analysis of the pattern counting problems. In particular, this last group may safely skip the next subsection and rely only on Theorem 3.

Construction of Low-Degree Cayley Graph Expanders

We prove Theorem 3 via an explicit construction of the groups Γ_i and the symmetric generating sets S_i in the full version, motivated by number theoretic objects. In what follows we present an overview of our construction.

Let us fix a prime $p \geq 3$. The starting point is an explicit arithmetic lattice (a discrete subgroup) in a group of generalized quaternions over a function field in characteristic p. The quaternion algebra is at the heart of the mathematical properties of extracting the finite p-groups and the expansion property of the resulting Cayley graphs, but it is not crucial for understanding the construction. Concretely, for any choice of elements $\alpha \neq \beta \in \mathbb{Z}/(p-1)\mathbb{Z}$ we construct an infinite group $\Gamma_{p;\alpha,\beta}$ defined in terms of $2(p+1)$ generators a_k, b_j (where the indices k, j run through sets $K, J \subseteq \mathbb{Z}/(p^2-1)\mathbb{Z}$ defined depending on α, β) and relations of length 4. The set of relations is described by explicit algebraic equations in the field $\mathbb{F}_{p^2}$. In [40] these groups were realized by mapping the generators a_k, b_j to explicit generalized quaternions, leading ultimately to an explicit injective group homomorphism

$$\Psi \colon \Gamma_{p;\alpha,\beta} \to \mathrm{GL}_3(\mathbb{F}_p[[t]])\,. \tag{3}$$

In other words, every element of $\Gamma_{p;\alpha,\beta}$ is sent to an invertible 3×3-matrix whose entries are power series in some formal variable t, whose coefficients live in the finite field $\mathbb{F}_p$ with p elements. This is made explicit for $p = 3$ in the full version, but could also be made explicit for any $p \geq 5$. Since the applications do not depend on concrete matrices, we merely state its existence.

To construct the finite p-groups Γ_i, consider the group homomorphism

$$\pi_i : \mathrm{GL}_3(\mathbb{F}_p[[t]]) \to \mathrm{GL}_3(\mathbb{F}_p[t]/(t^{i+1}))$$

taking a matrix with power series entries and truncating the power series after the term of order t^i. Then the group $\mathrm{GL}_3(\mathbb{F}_p[t]/(t^{i+1}))$ is finite, and we define Γ_i to be the image of the group $\Gamma_{p;\alpha,\beta}$ under the composition $\pi_i \circ \Psi$. These groups Γ_i are easily shown to be p-groups and they are what is called congruence quotients (by construction). The generators a_k, b_j from the construction of $\Gamma_{p;\alpha,\beta}$ map to symmetric generating sets T_i of Γ_i, i.e., to the set of cosets $a_k N_i, b_j N_i$ when $\Gamma_i = \Gamma_{p;\alpha,\beta}/N_i$ is considered as a factor group. Using results from [40], we know that the Cayley graphs $G_i = \mathcal{C}(\Gamma_i, T_i)$ associated to the congruence quotient groups Γ_i with respect to the generating sets T_i are expanders. This argument is worked out in [40] by Rungtanapirom and two of the authors, and it is based on a similar approach in the classical papers by Lubotzky, Phillips and Sarnak [30] and by Morgenstern [35]. We note here that the results of [40] ultimately rely on deep number theoretic results, namely a translation of the spectrum of the adjacency operator into Satake parameters of an associated automorphic representation and most crucially on work of Drinfeld on the geometric Langlands programme for GL_2.

At this point we have proven Theorem 3 for the particular valency $2d = 2(p+1)$. In order to obtain the more general valencies stated in the theorem (which will be crucial for some of our reductions), we recall in Section 3.1 of the full version that a uniformly controlled change of the generating sets T_i of the groups Γ_i (the generators must be mutually expressible in words of uniformly bounded length) preserves the expander property. This change of generating set is best performed by finding a smaller generating set for the underlying infinite group $\Gamma_{p;\alpha,\beta}$. This is done in Proposition 3.7 in the full version, reducing to $d = (p+3)/2$ for all $p \geq 3$. The reduction is based on the explicit form of the relations and a combinatorial group theoretic result from [42] on the local permutation structure of the underlying geometric square complex. To improve even further for $p = 3$ we consider in Section 3.3 of the full version a concrete presentation of $\Gamma_{3;0,1}$ which is shown to reduce to 2 generators. For $p = 5$, an explicit example given in the full version achieves a reduction to 2 generators for $\Gamma_{5;0,2}$. In the last step, we will then show that by adding generators (as necessary) we obtain Theorem 3 for all d's in the range that the theorem promises.

While this is not needed for the purposes of our hardness results, all of the constructions above are explicit, certainly in the weak sense that for a fixed p, the sequence of graphs G_i from Theorem 3 is computable. We also would like to emphasize again, that the expanders constructed for the proof of Theorem 3 consist of vertex transitive graphs, of prime power number of vertices, with a fairly low bound on the degree. All of this is made possible by working with very specific generalized quaternion groups in positive characteristic.

Analysis of the Indicators

Having established the existence of the low-degree Cayley graph expanders, we turn back to the analysis of the indicator

$$a(\Phi, H) = \sum_{\sigma \in \mathcal{L}(\Phi,H)} \prod_{v \in V(H)} (-1)^{|\sigma_v|-1}(|\sigma_v| - 1)!\,. \tag{4}$$

Recall that we claimed the analysis of $a(\Phi, H)$ to be easier for vertex-transitive graphs. Let us now elaborate on this claim. First of all, we restate the formal definition of Cayley graphs for readers who skipped the explicit construction of our expanders: the *Cayley graph* of a group Γ together with a symmetric generating set[5] $S \subseteq \Gamma$ is the graph $G = \mathcal{C}(\Gamma, S)$ with vertex set $V(G) = \Gamma$ and edge set

$$E(G) = \{(x, xs) \in V(G) \times V(G); x \in \Gamma,\ s \in S\}.$$

Since S is symmetric, with any edge (x, xs) the Cayley graph also contains the edge with opposite orientation $(xs, x) = (xs, (xs)s^{-1})$. Hence we consider Cayley graphs as the underlying unoriented graph.

Given a Cayley graph G as above, the group Γ acts on the graph by letting $g \in \Gamma$ send the vertex $v \in V(G) = \Gamma$ to gv. This action extends to the set of fractures $\mathcal{L}(\Phi, H)$ and since the terms $\prod_{v \in V(H)}(-1)^{|\sigma_v|-1}(|\sigma_v| - 1)!$ in the formula (4) are shown to be invariant under this action, the group Γ naturally permutes these summands. Since our Cayley graph expanders G_i arise from p-groups Γ_i, it follows that when evaluating the indicator $a(\Phi, G_i)$ *modulo* p, only those contributions from fractures fixed under Γ_i survive. Now recall that σ_v is a partition of the edges incident to v. The fixed-point fractures σ will satisfy that all σ_v are equal if we identify the edges incident to v with the elements of the generating set. Since, for

[5] This means a subset $S \subseteq \Gamma$ of the group that generates this group and satisfies $S^{-1} = S$.

fixed p, our Cayley graph expanders have constant degree, we can thus prove the indicator to be non-zero modulo p by considering just a constant number of fractures. This approach was first used in [39], and we will show that it becomes significantly more powerful if applied to our novel Cayley graph expanders.

Now let us illustrate and sketch this approach for the property Φ of being bipartite, that is, for proving Theorem 2. By Theorem 3, there is a family $\mathcal{G}$ of 5-group Cayley graph expanders of degree 6. For graphs $G \in \mathcal{G}$, we can show that the indicator $a(\Phi, G)$ does not vanish, given that Φ is the property of being bipartite. Theorem 2 will then follow by the argument outlined above; the detailed and formal proof is presented in the full version.

In the first step, given a graph $G = \mathcal{C}(\Gamma_i, S_i) \in \mathcal{G}$ we need to establish which fixed-point fractures σ are contained in $\mathcal{L}(\Phi, G)$, that is, for which σ the fractured graph $G\#\sigma$ is bipartite. Since fixed-point fractures $\sigma = (\sigma_v)_{v\in V(G)}$ of G satisfy that all σ_v correspond to one particular partition of S_i, we will ease notation and identify σ with this partition. Using that S_i is a symmetric set of generators of cardinality 6, that is, $S_i = \{g_1, g_2, g_3, g_1^{-1}, g_2^{-1}, g_3^{-1}\}$, we define a graph $\mathcal{H}(\sigma)$ as follows: It has a vertex w^B for each block B of σ, and its set of (multi)edges is given by

$$E(\mathcal{H}(\sigma)) = \left\{ \{w^B, w^{B'}\} : \text{one multiedge for each } g \in \{g_1, g_2, g_3\} \text{ s.t. } g \in B, g^{-1} \in B' \right\} . \quad (5)$$

Note that we see $\mathcal{H}(\sigma)$ as a graph with possible loops and possible multiedges. In particular, the graph $\mathcal{H}(\sigma)$ has precisely 3 edges.

The important property of $\mathcal{H}(\sigma)$, which we prove in the full version, is that $\mathcal{H}(\sigma)$ is bipartite if and only if $G\#\sigma$ is bipartite. Consequently, for Φ being the property of being bipartite, the indicator $a(\Phi, G)$ is given by the following drastically simplified[6] expression, if considered modulo 5:

$$a(\Phi, G) \equiv \sum_{\sigma:\mathcal{H}(\sigma) \text{ is bipartite}} (-1)^{|\sigma|-1} \cdot (|\sigma| - 1)! \mod 5 \,, \quad (6)$$

where the sum is over partitions σ of S_i. We provide the evaluation of the above expression in Table 1 and observe that the result is $-16 \neq 0 \mod 5$. Since this argument applies to all members of the family of 5-group Cayley graph expanders, we conclude that the indicator is non-zero infinitely often, which ultimately proves Theorem 2.

The proof of Theorem 1 will follow a comparable technique, but it will require multiple, more involved cases.

Extension to Modular Counting

Our understanding of the cancellation behaviour of the indicators $a(\Phi, H)$ modulo p does not only allow us to analyse the complexity of exact counting of k-edge subgraphs satisfying Φ, but also extends to counting k-edge subgraphs modulo p.

To this end, we provide the necessary set-up for parameterized modular counting. In particular, the basis for our intractability results on the modular counting versions is given by the *decision* problem of detecting so-called colour-prescribed homomorphisms, which is known to be hard for graphs of high treewidth due to Marx [31]. Using a version of the Schwartz-Zippel-Lemma due to Williams et al. [46], we are able to reduce an instance I of

[6] A priori, the formula (4) leads to the version of formula (6) with all summands taken to the power $|V(G)|$. However, since the number of vertices is a power of $p = 5$, by Fermat's little theorem it can be omitted.

Table 1 List of bipartite graphs $\mathcal{H}(\sigma)$ for 3 generators; here we give the isomorphism class of $\mathcal{H}(\sigma)$, the number of partitions σ with the corresponding isomorphism class, the number of blocks of sigma and the total contribution to $a(\Phi, G) \mod 5$ for the property Φ of being bipartite.

$\mathcal{H}(\sigma)$	No. of σ	$\lvert\sigma\rvert$	$(-1)^{\lvert\sigma\rvert-1} \cdot (\lvert\sigma\rvert - 1)!$
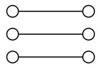	1	6	$-120 \cdot 1$
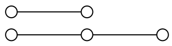	12	5	$24 \cdot 12$
	24	4	$-6 \cdot 24$
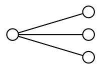	8	4	$-6 \cdot 8$
	6	4	$-6 \cdot 6$
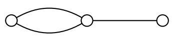	24	3	$2 \cdot 24$
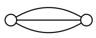	4	2	$-1 \cdot 4$
Total contribution			$-16 \equiv 4 \mod 5$

this problem to an instance I', such that, with high probability, I' has precisely one solution if I has at least one solution, and I' has no solutions if I has no solutions. In a second step, the obtained instance I' can then easily be reduced to the modular counting version for each prime p. Since the first step of this reduction is randomised, we need to assume the randomised Exponential Time Hypothesis for our fine-grained lower bounds.

Afterwards, we prove a variant of the Complexity Monotonicity principle for modular counting in the case of colour-prescribed homomorphisms. As a consequence, our hardness results for modular subgraph counting problems, including Theorem 4, can be proved following the same strategy as outlined in the previous subsection.

Moreover, instead of only presenting intractability results, we investigate whether the expression as a linear combination of homomorphism counts can also be used to achieve improved algorithms for modular subgraph counting problems. And indeed, considering the linear combination (1) modulo 2 for the property Φ of being a path or a cycle, allows us to prove that each graph H with degree at least 5 vanishes in the linear combination, that is, $a_{\Phi,k}(H) = 0$. More precisely, we will prove this for a version of the problem in which two vertices of the k-paths or the k-cycles are already fixed. This must be done to avoid automorphisms of even order, which turns out to be necessary for (1) to be well-defined modulo 2, since some of the coefficients $a_{\Phi,k}(H) = 0$ are of the form $\#\mathsf{Aut}(H)^{-1}$.

The algorithms for counting k-paths and k-cycles modulo 2 turn then out to be very simple: Essentially, we will see that it suffices to guess the two fixed vertices, and thereafter the algorithm evaluates Equation (1) modulo 2, by computing each non-vanishing term using a standard treewidth-based dynamic programming algorithm for counting homomorphisms. Since each graph H whose coefficient survives modulo 2 has degree at most 4, we can rely on known results on the treewidth of bounded degree graphs [19]. Ultimately, this allows us to prove Theorem 5.

Finally, our classification for counting homomorphisms modulo p builds upon the well-established algorithms and reduction sequences used both in the classification for the decision problem [23], as well as in the classification for the exact counting problem [13]. However, the difficulty in proving our classification for counting modulo p is due to graphs H which have high treewidth but admit automorphisms of order p. For those graphs, we can neither rely on an algorithm for exact counting, nor does the known hardness proof transfer.

We solve this problem by considering the p-reduced quotients. Let us denote the function that maps a graph G to the number of homomorphisms from H to G, modulo p, by $\#_p\mathsf{Hom}(H \to \star)$. We show that for each graph H we have

$$\#_p\mathsf{Hom}(H \to \star) = \#_p\mathsf{Hom}(H_p^* \to \star)\,.$$

As a consequence, it suffices to consider the p-reduced quotients for our classification. Since, by definition, those graphs to not admit an automorphism of order p, we are able to show that the known methods for proving classifications for homomorphism problems apply.

Let us conclude by pointing out that, while proving that the p-reduced quotient is uniquely defined up to isomorphism, we also establish a modular variant of Lovász' criterion for graph isomorphism via homomorphism counts (see Chapter 5 in [29]):

▶ **Lemma 9.** *Let H and H' be graphs, neither of which has an automorphism of order p. Suppose that for all graphs G we have that*

$$\#_p\mathsf{Hom}(H \to G) = \#_p\mathsf{Hom}(H' \to G)\,.$$

Then H and H' are isomorphic.

We point out that the previous result can be considered a dual version of [16, Lemma 3.10].

Conclusion and Open Questions

All of our hardness results for modular subgraph counting problems only apply to primes $p \geq 3$ and have, using the randomised Exponential Time Hypothesis as a slightly stronger assumption, the same complexity as their counterparts from exact counting. However, for $p = 2$ the complexity landscape seems different: We obtained an improvement for counting k-cycles and k-paths modulo 2. Moreover, there are known instances of the counting version of the parameterized subgraph isomorphism problem, such as counting k-matchings, where exact counting, as well as counting modulo p for each prime $p \geq 3$ is fixed-parameter **in**tractable, while the computation becomes fixed-parameter tractable if done modulo 2 [9, 10].

Since, additionally, many of our hardness proofs do not apply to the case of counting modulo 2, we propose a thorough investigation of the complexity of the parameterized subgraph counting problem modulo 2 as the next step in this line of research. As a starting point, we suggest the problem of counting bipartite k-edge subgraphs modulo 2: While our proofs extend to counting such subgraphs modulo p for some primes $p \geq 3$, a computer-aided search revealed that, for $p = 2$, our approach *cannot* work for any family of Cayley graph expanders of degree at most 12; details are provided in the full version. Indeed, we conjecture that our methods for proving intractability can be used to show that the problem is intractable for each prime $p > 2$, but not for $p = 2$, which leads to the question of whether this problem might be fixed-parameter tractable.

There are also interesting open questions concerning p-group Cayley graph expanders with low degree. To describe them, fix some prime p and consider the set $D(p) \subseteq \mathbb{Z}_{\geq 0}$ of integers d such that there exists a sequence of finite p-groups Γ_i of orders that tend to infinity, with symmetric generating sets S_i of cardinality $2d$ such that the Cayley graphs $\mathcal{C}(\Gamma_i, S_i)$ form a family of expanders. With any $d \in D(p)$ actually any $d' \geq d$ also lies in $D(p)$, because we can find a uniform bound on the length of a word in S_i to produce a new additional generator for Γ_i, showing $d + 1 \in D(p)$. So the ultimate question is the following:

▶ **Question 10.** What is the behaviour of the function $p \mapsto d(p) = \min D(p)$?

Since 2-regular graphs are never expanders, we know that $d(p) \geq 2$ for all primes p. Moreover, combining the construction of [36] (for $p = 2$) with our constructions and some further examples that we computed, we obtain the following values and bounds for the function d:

p	$p \in \{2, 3, 5, 7, 11, 13\}$	$17 \leq p \leq 83$	$89 \leq p$
d(p)	2	$d(p) \in \{2, 3\}$	$2 \leq d(p) \leq (p + 3)/2$

Based on this experimental evidence we make the following conjecture:

▶ **Conjecture 11.** *For every prime $p \geq 3$ there is a group among the $\Gamma_{p,\alpha,\beta}$ that is 3-generated. In particular, there are p-group Cayley graph expanders of fixed valency $2d$ for all $p \geq 3$ and all $d \geq 3$.*

If the conjecture is satisfied, the function d above would be uniformly bounded from above by 3.

References

1. Noga Alon, Phuong Dao, Iman Hajirasouliha, Fereydoun Hormozdiari, and S. Cenk Sahinalp. Biomolecular network motif counting and discovery by color coding. *Bioinformatics*, 24(13):i241–i249, July 2008. doi:10.1093/bioinformatics/btn163.
2. Vikraman Arvind and Venkatesh Raman. Approximation Algorithms for Some Parameterized Counting Problems. In *Algorithms and Computation, 13th International Symposium, ISAAC 2002 Vancouver, BC, Canada, November 21-23, 2002, Proceedings*, pages 453–464, 2002. doi:10.1007/3-540-36136-7_40.
3. László Babai. Graph isomorphism in quasipolynomial time [extended abstract]. In Daniel Wichs and Yishay Mansour, editors, *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 684–697. ACM, 2016. doi:10.1145/2897518.2897542.
4. Andreas Björklund, Holger Dell, and Thore Husfeldt. The Parity of Set Systems Under Random Restrictions with Applications to Exponential Time Problems. In *Automata, Languages, and Programming – 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I*, pages 231–242, 2015. doi:10.1007/978-3-662-47672-7_19.
5. Cornelius Brand, Holger Dell, and Thore Husfeldt. Extensor-coding. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 151–164, 2018. doi:10.1145/3188745.3188902.

6 Andrei A. Bulatov and Amirhossein Kazeminia. Complexity classification of counting graph homomorphisms modulo a prime number. *CoRR*, abs/2106.04086, 2021. arXiv:2106.04086.

7 Andrei A. Bulatov and Stanislav Zivný. Approximate counting CSP seen from the other side. *ACM Trans. Comput. Theory*, 12(2):11:1–11:19, 2020. doi:10.1145/3389390.

8 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.

9 Radu Curticapean. Counting Matchings of Size k Is W[1]-Hard. In Fedor V. Fomin, Rusins Freivalds, Marta Z. Kwiatkowska, and David Peleg, editors, *Automata, Languages, and Programming – 40th International Colloquium, ICALP 2013, Riga, Latvia, July 8-12, 2013, Proceedings, Part I*, volume 7965 of *Lecture Notes in Computer Science*, pages 352–363. Springer, 2013. doi:10.1007/978-3-642-39206-1_30.

10 Radu Curticapean, Holger Dell, and Thore Husfeldt. Modular counting of subgraphs: Matchings, matching-splittable graphs, and paths. In *29th Annual European Symposium on Algorithms, ESA 2021, September 6-8, 2021, Lisbon, Portugal*, 2021. to appear; preprint at arXiv:2107.00629.

11 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.

12 Radu Curticapean and Dániel Marx. Complexity of Counting Subgraphs: Only the Boundedness of the Vertex-Cover Number Counts. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 130–139, 2014. doi:10.1109/FOCS.2014.22.

13 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.

14 Julian Dörfler, Marc Roth, Johannes Schmitt, and Philip Wellnitz. Counting Induced Subgraphs: An Algebraic Approach to #W[1]-hardness. In Peter Rossmanith, Pinar Heggernes, and Joost-Pieter Katoen, editors, *44th International Symposium on Mathematical Foundations of Computer Science, MFCS 2019, August 26-30, 2019, Aachen, Germany*, volume 138 of *LIPIcs*, pages 26:1–26:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.MFCS.2019.26.

15 Arnaud Durand and Stefan Mengel. Structural Tractability of Counting of Solutions to Conjunctive Queries. *Theory Comput. Syst.*, 57(4):1202–1249, 2015. doi:10.1007/s00224-014-9543-y.

16 John D. Faben and Mark Jerrum. The Complexity of Parity Graph Homomorphism: An Initial Investigation. *Theory Comput.*, 11:35–57, 2015. doi:10.4086/toc.2015.v011a002.

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

18 Jacob Focke, Leslie Ann Goldberg, Marc Roth, and Stanislav Zivný. Counting Homomorphisms to K_4-minor-free Graphs, modulo 2. In Dániel Marx, editor, *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10–13, 2021*, pages 2303–2314. SIAM, 2021. doi:10.1137/1.9781611976465.137.

19 Fedor V. Fomin, Serge Gaspers, Saket Saurabh, and Alexey A. Stepanov. On Two Techniques of Combining Branching and Treewidth. *Algorithmica*, 54(2):181–207, 2009. doi:10.1007/s00453-007-9133-3.

20 Andreas Göbel, Leslie Ann Goldberg, and David Richerby. The complexity of counting homomorphisms to cactus graphs modulo 2. *ACM Trans. Comput. Theory*, 6(4):17:1–17:29, 2014. doi:10.1145/2635825.

21 Andreas Göbel, Leslie Ann Goldberg, and David Richerby. Counting Homomorphisms to Square-Free Graphs, Modulo 2. *ACM Trans. Comput. Theory*, 8(3):12:1–12:29, 2016. doi:10.1145/2898441.

22 Andreas Göbel, J. A. Gregor Lagodzinski, and Karen Seidel. Counting Homomorphisms to Trees Modulo a Prime. In Igor Potapov, Paul G. Spirakis, and James Worrell, editors, *43rd International Symposium on Mathematical Foundations of Computer Science, MFCS 2018, August 27-31, 2018, Liverpool, UK*, volume 117 of *LIPIcs*, pages 49:1–49:13. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.MFCS.2018.49.

23 Martin Grohe. The complexity of homomorphism and constraint satisfaction problems seen from the other side. *J. ACM*, 54(1):1:1–1:24, 2007. doi:10.1145/1206035.1206036.

24 Heng Guo, Sangxia Huang, Pinyan Lu, and Mingji Xia. The Complexity of Weighted Boolean #CSP Modulo k. In Thomas Schwentick and Christoph Dürr, editors, *28th International Symposium on Theoretical Aspects of Computer Science, STACS 2011, March 10-12, 2011, Dortmund, Germany*, volume 9 of *LIPIcs*, pages 249–260. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2011. doi:10.4230/LIPIcs.STACS.2011.249.

25 Heng Guo, Pinyan Lu, and Leslie G. Valiant. The Complexity of Symmetric Boolean Parity Holant Problems. *SIAM J. Comput.*, 42(1):324–356, 2013. doi:10.1137/100815530.

26 Amirhossein Kazeminia and Andrei A. Bulatov. Counting Homomorphisms Modulo a Prime Number. In Peter Rossmanith, Pinar Heggernes, and Joost-Pieter Katoen, editors, *44th International Symposium on Mathematical Foundations of Computer Science, MFCS 2019, August 26-30, 2019, Aachen, Germany*, volume 138 of *LIPIcs*, pages 59:1–59:13. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.MFCS.2019.59.

27 J. A. Gregor Lagodzinski, Andreas Göbel, Katrin Casel, and Tobias Friedrich. On Counting (Quantum-)Graph Homomorphisms in Finite Fields. In *48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, to appear*, 2021. Preprint on arXiv:2011.04827.

28 Daniel Lokshtanov, Saket Saurabh, and Meirav Zehavi. Efficient Computation of Representative Weight Functions with Applications to Parameterized Counting. Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Alexandria, VA, USA, January 10-13, 2021, to appear.

29 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.

30 A. Lubotzky, R. Phillips, and P. Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988. doi:10.1007/BF02126799.

31 Dániel Marx. Can You Beat Treewidth? *Theory of Computing*, 6(1):85–112, 2010. doi:10.4086/toc.2010.v006a005.

32 Dániel Marx. Tractable Hypergraph Properties for Constraint Satisfaction and Conjunctive Queries. *J. ACM*, 60(6):42:1–42:51, 2013. doi:10.1145/2535926.

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

34 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.

35 Moshe Morgenstern. Existence and explicit constructions of $q+1$ regular Ramanujan graphs for every prime power q. *J. Combin. Theory Ser. B*, 62(1):44–62, 1994. doi:10.1006/jctb.1994.1054.

36 Norbert Peyerimhoff and Alina Vdovina. Cayley graph expanders and groups of finite width. *J. Pure Appl. Algebra*, 215(11):2780–2788, 2011. doi:10.1016/j.jpaa.2011.03.018.

37 Marc Roth and Johannes Schmitt. Counting Induced Subgraphs: A Topological Approach to #W[1]-hardness. *Algorithmica*, 82(8):2267–2291, 2020. doi:10.1007/s00453-020-00676-9.

38 Marc Roth, Johannes Schmitt, and Philip Wellnitz. Counting Small Induced Subgraphs Satisfying Monotone Properties. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 1356–1367. IEEE, 2020. doi:10.1109/FOCS46700.2020.00128.

39 Marc Roth, Johannes Schmitt, and Philip Wellnitz. Detecting and Counting Small Subgraphs, and Evaluating a Parameterized Tutte Polynomial: Lower Bounds via Toroidal Grids and Cayley Graph Expanders. In *48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, to appear*, 2021. Preprint on arXiv:2011.03433.

40 Nithi Rungtanapirom, Jakob Stix, and Alina Vdovina. Infinite series of quaternionic 1-vertex cube complexes, the doubling construction, and explicit cubical ramanujan complexes. *International Journal of Algebra and Computation*, 29(6):951–1007, 2019.

41 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.

42 Jakob Stix and Alina Vdovina. Simply transitive quaternionic lattices of rank 2 over $\mathbb{F}_{q(t)}$ and a non-classical fake quadric. *Mathematical Proceedings of the Cambridge Philosophical Society*, 163(3):453–498, 2017.

43 Seinosuke Toda. PP is as Hard as the Polynomial-Time Hierarchy. *SIAM J. Comput.*, 20(5):865–877, 1991. doi:10.1137/0220053.

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

45 Leslie G. Valiant. Holographic Algorithms. *SIAM J. Comput.*, 37(5):1565–1594, 2008. doi:10.1137/070682575.

46 Virginia Vassilevska Williams, Joshua R. Wang, Richard Ryan Williams, and Huacheng Yu. Finding Four-Node Subgraphs in Triangle Time. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 1671–1680, 2015. doi:10.1137/1.9781611973730.111.