SIAM J. DISCRETE MATH.
Vol. 37, No. 2, pp. 405–486

PARAMETERIZED COUNTING AND CAYLEY GRAPH EXPANDERS[*]

NORBERT PEYERIMHOFF[†], MARC ROTH[‡], JOHANNES SCHMITT[§], JAKOB STIX[¶], ALINA VDOVINA[‖], AND PHILIP WELLNITZ[#]

Abstract. Given a graph property Φ, we consider the problem $\#$ EDGESUB(Φ), where the input is a pair of a graph G and a positive integer k, and the task is to compute the number of k-edge subgraphs in G that satisfy Φ. Specifically, we study the parameterized complexity of #EDGESUB(Φ) with respect to both approximate and exact counting, as well as its decision version EDGESUB(Φ). Among others, our main result fully resolves the case of minor-closed properties Φ: the decision problem EDGESUB(Φ) always admits a fixed-parameter tractable algorithm, and the counting problem #EDGESUB(Φ) always admits a fixed-parameter tractable randomized approximation scheme. For exact counting, we present an exhaustive and explicit criterion on the property Φ which, if satisfied, yields fixed-parameter tractability and otherwise #W[1]-hardness. Additionally, our hardness results come with an almost tight conditional lower bound under the exponential time hypothesis. Our main technical result concerns the exact counting problem: Building upon the breakthrough result of Curticapean, Dell, and Marx (Symposium on Theory of Computing 2017), we express the number of subgraphs satisfying Φ 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 analyze the coefficients in the aforementioned linear combinations over the field $\mathbb{F}_p$ which gives us significantly more control over the cancelation behavior of the coefficients.

Key words. counting complexity, fine-grained and parameterized complexity, graph homomorphisms, subgraphs, Caley graph expanders

MSC codes. 68Q27, 68Q17, 68R10, 05E40, 05C48, 11F06, 11R52, 20F65

DOI. 10.1137/22M1479804

1. Introduction.

Be it searching for cliques in social networks or understanding protein-protein interaction networks, many interesting real-life problems boil down to finding (or counting) small patterns in large graphs. Hence, to no surprise, finding (and counting) small patterns in large graphs is among the most well-studied

[*]Received by the editors February 22, 2022; accepted for publication (in revised form) September 13, 2022; published electronically April 26, 2023. Extended abstracts (without the proofs) presenting different cases of our main result appeared in the 2021 International Colloquium on Automata, Languages, and Programming [80] and the 2021 International Symposium on Mathematical Foundations of Computer Science [70].

https://doi.org/10.1137/22M1479804

Funding: The work of the third author was supported by the SNF early postdoc mobility grant 184245.

[†]Department of Mathematical Sciences, Durham University, Durham, United Kingdom (norbert.peyerimhoff@durham.ac.uk).

[‡]Merton College and Department of Computer Science, University of Oxford, Oxford, United Kingdom (marc.roth@cs.ox.ac.uk).

[§]Institute for Mathematics, University of Zürich, Zürich, Switzerland (johannes.schmitt@math.uzh.ch).

[¶]Mathematical Institute, Goethe-Universität Frankfurt, Frankfurt am Main, Germany (stix@math.uni-frankfurt.de).

[‖]Department of Mathematics, City College of New York, CUNY, New York, NY 10031 USA (avdovina@ccny.cuny.edu).

[#]Max Planck Institute for Informatics, Saarland Informatics Campus SIC, Saarbrücken, Germany (wellnitz@mpi-inf.mpg.de).

computational problems in the fields of database theory [14, 15, 31, 36, 46], molecular biology and bioinformatics [1, 43, 73, 82], and network science [66, 67, 83]. In fact, already in the 1970s, the relevance of *finding* patterns had become apparent in the context of finding cliques, finding Hamiltonian paths, or finding specific subgraphs in general [14, 20, 21, 88]. However, with the advent of motif counting for the frequency analysis of small structures in complex networks [66, 67], it became evident that detecting the existence of a pattern graph is not enough; we also need to *count* all of the occurrences of the pattern.

In this work, our patterns are (not necessarily induced) edge subgraphs that satisfy a certain graph property: for instance, given a graph, we want to count all occurrences of edge subgraphs that are are planar or connected. From a classical point of view, often the problem of *finding* patterns is already NP-hard: prime examples include the aforementioned problems of finding (maximum) cliques or Hamiltonian paths. However, for the task of network motif counting, the patterns are (almost) always much smaller than the network itself (see [1, 66, 67]). This motivates a *parameterized* view: can we obtain fast algorithms to compute the number of occurrences of "small" patterns? If we cannot, can we at least obtain fast (randomized) algorithms to compute an *estimate* of this number? And if we cannot even do this, can we at least obtain fast algorithms to *detect* an occurrence? In this work, we completely answer all of the above questions for patterns that are specified by *minor-closed* graph properties (such as planarity) or selected other graph properties (such as connectivity).

As it turns out, the techniques we develop for answering the above questions are quite powerful: they easily generalize to a parameterized version of the Tutte polynomial. Specifically, our techniques allow us to completely understand at which rational points we can evaluate the said parameterized Tutte polynomial in reasonable time and at which rational points this is not feasible. This dichotomy turns out to be similar, but not equal, to the complexity landscape of the classical Tutte polynomial due to Jaeger, Vertigan, and Welsh [49].

Parameterized counting and hardness. By now, *counting complexity theory* is a well-established subfield of theoretical computer science. Already in the 1970s, Valiant started a formal study of counting problems when investigating the complexity of the permanent [89, 90]: counting the number of perfect matchings in a graph is #P-complete and hence harder than any problem in the polynomial-time hierarchy PH by Toda's theorem [87]. In contrast, *detecting* a perfect matching in a graph is much easier and can be done in polynomial time [37]. Hence, counting problems can be much harder than their decision problem counterparts. As an attempt to overcome the hardness of counting problems in general, the focus shifted to a multivariate or *parameterized* view on these problems. Consider, for example, the following problem: given a query φ of size k and a database B of size n, we want to count the number of answers to φ in B. If we make the very reasonable assumption that k is much smaller than n, then we may consider an algorithm running in time $O(2^k \cdot n)$ as *tractable*. Note that in particular, such an algorithm may even outperform an algorithm running in time $O(n^2)$. Also consider [44] for a more detailed and formal discussion.

Formally, given a problem P and a *parameterization* κ that maps each instance I of P to a parameter $\kappa(I)$, we say that P is *fixed-parameter tractable* with respect to κ if there is an algorithm that solves each instance I of size n in time $f(\kappa(I)) \cdot n^{O(1)}$ for some computable function f. This notion was introduced by Downey and Fellows in the early 1990s [33, 34] and has itself spawned a rich body of literature (see [27, 35, 39]). In the context of the problems of detecting and counting small patterns in large networks, we parameterize by the size of the pattern: given a pattern of size k and a network of size n, we aim for algorithms that run in time $f(k) \cdot n^{O(1)}$ for some computable function f. However, for some patterns, even this goal is too ambitious: it is widely believed that even finding a clique of size k is not fixed-parameter tractable; in particular, a fixed-parameter tractable algorithm for finding a clique of size k would also imply a breakthrough result for the satisfiability problem and thereby refute the widely believed exponential time hypothesis (ETH) [17, 18]. If a problem P is at least as hard as finding a clique (or counting all cliques) of size k, we say that P is $\mathsf{W}[1]$-hard (or $\#\mathsf{W}[1]$-hard, respectively).

For such a (#)$\mathsf{W}[1]$-hard problem, the hope is to (significantly) improve upon the naive brute-force algorithm, which runs in time $n^{O(k)}$ for the problems considered in this work. However, in view of the aforementioned reduction from the satisfiability problem to the problem of finding cliques of size k [16, 17], we can see that for finding cliques this, too, would require a breakthrough for the satisfiability problem, which, again, is believed to be unlikely [48]. In our paper, via suitable reductions from the problem of finding cliques, we establish that exact algorithms significantly faster than the brute-force algorithms are unlikely for the problems we study.

Parameterized detection and counting of edge subgraphs. It is well known that *vertex-induced* subgraphs as patterns are notoriously hard to detect or to count. The long line of research on this problem [19, 25, 32, 50, 51, 52, 53, 65, 78, 79] showed that this holds even if the patterns are significantly smaller than the host graphs, as witnessed by $\mathsf{W}[1]$ and $\#\mathsf{W}[1]$-hardness results and almost tight conditional lower bounds. In the case of exact counting, it is in fact an open question whether there are nontrivial instances of induced subgraph counting that admit efficient algorithms; recent work [79] supports the conjecture that no such instances exist.

In search for fast algorithms, in this work, we hence consider a related but different version of network-motif counting: for a computable graph property Φ, in the problem #EdgeSub(Φ) we are given a graph G and a positive integer k, and the task is to compute the number of (not necessarily induced) edge subgraphs[1] with k edges in G that satisfy Φ. Similarly, we write EdgeSub(Φ) for the corresponding decision problem. Then, in contrast to the case of counting vertex-induced subgraphs, for (#)EdgeSub(Φ), we identify nontrivial properties Φ for which (#)EdgeSub(Φ) is fixed-parameter tractable; we discuss this in more detail later. First, however, let us take a detour to elaborate more on what is known already for (#)EdgeSub(Φ).

If the property Φ is satisfied by at most a single graph for each value of the parameter k, the decision problem EdgeSub(Φ) becomes the subgraph isomorphism problem. Hence, naturally there is a vast body of known techniques and results for special properties Φ: for fixed-parameter tractable algorithms, think of the color-coding technique by Alon, Yuster, and Zwick [3], the “divide and color” technique [18], narrow sieving [9], representative sets [40], and “extensor-coding” [11], to name but a few. For hardness results, apart from the aforementioned example of detecting a clique, Lin quite recently established that detecting a k-Biclique is also $\mathsf{W}[1]$-hard [58]. However, a complete understanding of the parameterized decision version of the subgraph isomorphism is one of the major open problems of parameterized complexity theory [35, Chapter 33.1] that is still to be solved.

In the setting of parameterized *counting*, the situation is much better understood: Flum and Grohe [38] proved #EdgeSub(Φ) to be $\#\mathsf{W}[1]$-hard when Φ is the property of being a cycle or the property of being a path. Curticapean [23] established the same result for the property of being a matching. In [26], Curticapean and Marx established a complete classification in the case that Φ does not hold on two different graphs with

[1]Recall that an edge subgraph G' of a graph G may have fewer edges than the subgraph of G that is induced by the vertices of G'.

the same number of edges, which is essentially the parameterized subgraph counting problem. In particular, they identified a bound on the matching number as the tractability criterion. In a later work, together with Dell [25], they presented what is now called the framework of complexity monotonicity, which can be considered to be one of the most powerful tools in the field of parameterized counting problems. Note that this does not classify the decision version.

In contrast to the parameterized subgraph detection/counting problems, the problem (#)EdgeSub(Φ) allows us to search for more general patterns. For example, while the (parameterized) complexity of counting all subgraphs of a graph G isomorphic to a *fixed* connected graph H with k edges is fully understood [26], the case of counting, for instance, all connected or all bipartite k-edge subgraphs of a graph G remains open so far. As consequences of our main technical result, we completely understand the problem #EdgeSub(Φ) for the properties $\Phi =$ connectivity and $\Phi =$ bipartiteness. In what follows, we present our results, followed by an exposition of the most important techniques.

1.1. Main results. In a first part, we present our results on (#)EdgeSub(Φ); we continue with a definition and our results for a parameterized Tutte polynomial in a second part.

Our main results on (#)EdgeSub(Φ) can be categorized in roughly three categories: (1) exact algorithms and hardness results for the counting problem, (2) approximation algorithms for the counting problem, and (3) algorithms for the decision problem. For minor-closed properties Φ, we obtain exhaustive results for all three categories; for other (classes of) properties that we study, we obtain partial criteria.

Complete classification for minor-closed properties. Let us start with the case where the graph property Φ is closed under taking minors, that is, if Φ holds for a graph, then Φ still holds after removing vertices or edges, or after contracting edges. For minor-closed properties Φ, we obtain a complete picture of the complexity of #EdgeSub(Φ) and EdgeSub(Φ). In what follows, we say that a property Φ has *bounded matching number* if there is a constant bound on the size of a largest matching in graphs satisfying Φ.

Main Theorem 1. *Let Φ denote a minor-closed graph property.*

1. Exact counting*: If Φ is either trivially true or of bounded matching number, then the (exact) counting version* #EdgeSub(Φ) *is fixed-parameter tractable. Otherwise, the problem* #EdgeSub(Φ) *is* #W[1]*-hard and, assuming the ETH, cannot be solved in time $f(k) \cdot |G|^{o(k/\log k)}$ for any functionf.*
2. Approximate counting*: The problem* #EdgeSub(Φ) *always has an FPTRAS.*[2]
3. Decision*: The problem* EdgeSub(Φ) *is always fixed-parameter tractable.*

Let us present an exemplary application of our main result; further discussions on those examples can be found in [70].

Corollary 1.1. *The following problems are* #W[1]*-hard and, assuming ETH, cannot be solved in time $f(k)\cdot|G|^{o(k/\log k)}$ for any function f: Given G and k, compute the number of*

[2]A fixed-parameter tractable randomized approximation scheme. The formal definition is given in section 2.2; intuitively an FPTRAS is the parameterized equivalent of a fully polynomial-time randomized approximation scheme (FPRAS).

- *planar subgraphs with k edges in G,*
- *linear k-forests (unions of paths) in G,*[3]
- *k-edge subgraphs in G with tree-depth at most t, for a constant $t \geq 2$,*
- *k-edge subgraphs in G with genus at most g, for a constant $g \geq 0$,*
- *k-edge subgraphs in G with Colin de Verdière's invariant at most c, for a constant $c \geq 1$.*

In contrast, all of the previous problems have an FPTRAS and a fixed-parameter tractable decision version.

Further results on exact counting. Let us return to the case of arbitrary graph properties Φ. Without any further assumptions on Φ, the naive algorithm for $\#\textsc{EdgeSub}(\Phi)$ on the input (k, G) proceeds by enumerating the k-edge subsets of G and counting the number of cases where the corresponding subgraph satisfies Φ. This leads to a running time of the form $f(k) \cdot |V(G)|^{2k+O(1)}$. However, at least the linear constant in the exponent can be substantially improved using the currently fastest known algorithm for counting subgraphs with k edges due to Curticapean, Dell, and Marx [25]. We will show that it easily extends to the case of $\#\textsc{EdgeSub}(\Phi)$:

PROPOSITION 1.2. *Let Φ denote a computable graph property. There exists a computable function f such that $\#\textsc{EdgeSub}(\Phi)$ can be solved in time*

$$f(k) \cdot |V(G)|^{0.174k+o(k)}.$$

On the other hand, it was shown by Curticapean and Marx [26] that for the property Φ of being a matching, the problem $\#\textsc{EdgeSub}(\Phi)$ cannot be solved in time $f(k) \cdot |V(G)|^{o(k/\log k)}$ for any function f unless ETH fails. In other words, asymptotically and up to a factor of $1/\log k$, the exponent of $|V(G)|$ in the running time of $\#\textsc{EdgeSub}(\Phi)$ cannot be improved without posing any restriction on Φ.

The goal is hence to identify properties Φ for which the algorithm in Proposition 1.2 *can* be (significantly) improved. In the best possible outcome, we hope to identify the properties for which the exponent of $|V(G)|$ does not depend on k; those cases are precisely the fixed-parameter tractable ones. An easy consequence of known results for subgraph counting (see, for instance, [26]) establishes the following tractability criterion; we will include the proof only for the sake of completeness.

PROPOSITION 1.3. *Let Φ denote a computable graph property satisfying that there is $M > 0$ such that for all k either the graphs with k edges satisfying Φ or the graphs with k edges satisfying $\neg\Phi$ have matching number bounded by M. Then $\#\textsc{EdgeSub}(\Phi)$ is fixed-parameter tractable.*

Examples of properties satisfying the tractability criterion of Proposition 1.3 include, among others, the property of being a star or the complement thereof. We conjecture that all remaining properties induce $\#\mathsf{W}[1]$-hardness and cannot be solved in time $f(k) \cdot |G|^{o(k/\log k)}$ for any function f unless ETH fails.[4] For the case of minor-closed graph properties, we have seen above that this conjecture holds.

Further, the techniques we develop to prove hardness of $\#\textsc{EdgeSub}(\Phi)$ for minor-closed properties Φ in Main Theorem 1 can also be applied directly to show hardness

[3] $\#\mathsf{W}[1]$-hardness and, implicitly, the same conditional lower bound for counting (not necessarily linear) k-forests was shown in [13] and is subsumed by our result as well.

[4] Note that it does not matter whether we choose $|G|$ or $|V(G)|$ for the size of the large graph since we care about the asymptotic behavior of the exponent.

for other specific properties Φ. Below, we record several natural examples of such properties which are covered by our methods.

MAIN THEOREM 2. *Consider the following graph properties:*

- $\Phi_C(H)=1$ *if and only if* H *is connected.*
- $\Phi_B(H)=1$ *if and only if* H *is bipartite.*
- $\Phi_H(H)=1$ *if and only if* H *is Hamiltonian.*
- $\Phi_E(H)=1$ *if and only if* H *is Eulerian.*
- $\Phi_{CF}(H)=1$ *if and only if* H *is claw-free.*

For each $\Phi \in \{\Phi_C, \Phi_B, \Phi_H, \Phi_E, \Phi_{CF}\}$*, the problem* $\#\text{EDGESUB}(\Phi)$ *is* $\#\mathsf{W}[1]$*-hard. Further, unless ETH fails,* $\#\text{EDGESUB}(\Phi)$ *cannot be solved in time* $f(k)\cdot|G|^{o(k/\log k)}$ *for any function* f*.*

Can you beat treewidth?. We conclude the presentation of our results on exact counting by commenting on the factor of $1/(\log \ldots)$ 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" [63]. 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 [64], is true.

Novel constructions of p-group Cayley graph expanders. Cayley graphs are recalled in section 3 and in the technical overview. 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 (section 1.2).

MAIN THEOREM 3. *Let* $p\geq 3$ *be a prime number, and let* $d\geq 2$ *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).*

For completeness, we emphasize that the case $p=2$ has already been resolved by two of the authors in [71].

Results for approximate counting and decision. Our results on exact counting indicate that we have to relax the problem if we aim for tractability results for a larger variety of properties. One approach is to only ask for an *approximate* count of the number of k-edge subgraphs satisfying Φ. Tractability of approximation in the parameterized setting is given by the notion of an FPTRAS as introduced by Arvind and Raman [6]. While we give the formal definition in section 2.2, it suffices for now to think of an FPTRAS as a fixed-parameter tractable algorithm that can compute an arbitrarily good approximation of the answer with high probability. Readers familiar with the classical notions of approximate counting algorithms should think of an FPTRAS as an FPRAS in which we additionally allow a factor of $f(k)$ in the running time for any computable function f.

For the statement of our results, we say that a property Φ satisfies the *matching criterion* if it is true for all but finitely many matchings, and we say that it satisfies the *star criterion* if it is true for all but finitely many stars. Furthermore, we say that Φ has bounded treewidth if there is a constant upper bound on the treewidth of graphs that satisfy Φ.

MAIN THEOREM 4. *Let Φ denote a computable graph property. If Φ satisfies the matching criterion* and *the star criterion, or if Φ has bounded treewidth, then* $\#\text{EDGESUB}(\Phi)$ *admits an FPTRAS.*

For example, the property of being planar satisfies both the star and the matching criteria. Moreover, we can show that every minor-closed graph property Φ either has bounded treewidth or satisfies the matching and star criteria and thus always admits an FPTRAS.

Additionally, if not only exact but also approximate counting is intractable, we ask whether we can at least obtain an efficient algorithm for the decision version $\text{EDGESUB}(\Phi)$. Again, we obtain a tractability criterion; observe the subtle difference in the tractability criterion compared to Main Theorem 4.

MAIN THEOREM 5. *Let Φ denote a computable graph property. If Φ satisfies the matching criterion* or *the star criterion, or if Φ has bounded treewidth, then* $\text{EDGESUB}(\Phi)$ *is fixed-parameter tractable.*

As an easy corollary, we can conclude that for monotone, that is, subgraph-closed properties Φ, the problem $\text{EDGESUB}(\Phi)$ is always fixed-parameter tractable.[5]

For many previously studied problems, the complexity analyses of approximate counting and decision were related: often an algorithm solving one setting can be used to solve the other setting [30, 65]. However, in our Main Theorems 4 and 5 we see an asymmetry between the two settings: it suffices for Φ to satisfy only one of the star and the matching criteria to induce tractability of the decision version, but we require satisfaction of both for approximate counting. One might expect that this reflects a shortcoming of our proof methods (and that in fact it suffices to check one of the criteria to have tractability of approximate counting). Interestingly, this is *not* the case: There exists a property Ψ that satisfies the matching criterion, but not the star criterion, such that $\text{EDGESUB}(\Psi)$ is fixed-parameter tractable but $\#\text{EDGESUB}(\Psi)$ does not admit an FPTRAS unless $\mathsf{W}[1]$ coincides with FPT (the class of all fixed-parameter tractable decision problems) under randomized parameterized reductions [80].

Dichotomy for evaluating a parameterized Tutte polynomial. As a final part of the presentation of our main results, let us discuss our results on a parameterized Tutte polynomial.

The classical Tutte polynomial (as well as its specializations like the chromatic, flow, or reliability polynomials) have received widespread attention from both a combinatorial and a complexity theoretic perspective [2, 8, 10, 12, 29, 42, 49, 91]. The classical Tutte polynomial is of special interest from a complexity theoretic perspective, as the Tutte polynomial encodes a plethora of properties of a graph: prominent examples include the chromatic number, the number of acyclic orientations, and the number of spanning trees; we refer the reader to the work of Jaeger, Vertigan, and Welsh [49] for a comprehensive overview. Formally, the Tutte polynomial is a bivariate graph polynomial defined as follows (see [49]):

$$T_G(x,y) := \sum_{A \subseteq E(G)} (x-1)^{k(A)-k(E(G))} \cdot (y-1)^{k(A)+\#A-\#V(G)},$$

where $k(S)$ is the number of connected components of the graph $(V(G), S)$. In the aforementioned work, Jaeger, Vertigan, and Welsh [49] also classified the complexity

[5] Every graph property has either bounded treewidth or unbounded matching number. In the latter case, if the property is additionally monotone, it must satisfy the matching criterion.

of evaluating the Tutte polynomial in every pair of (complex) coordinates; that is, for every pair (a,b), the complexity of computing the function $G \mapsto T_G(a,b)$ is fully understood.

In this work, we consider the following parameterized version of the Tutte polynomial by restricting to edge-subsets A in G of size k:

$$T^k_G(x,y) := \sum_{A \in \binom{E(G)}{k}} (x-1)^{k(A)-k(E(G))} \cdot (y-1)^{k(A)+k-\#V(G)}.$$

We observe that the parameterized Tutte polynomial can be seen as a weighted version of counting small k-edge subgraph patterns by assigning to each k-edge subset A of G the weight

$$(x-1)^{k(A)-k(E(G))} \cdot (y-1)^{k(A)+k-\#V(G)}.$$

Moreover, we point out that $T^k_G(x,y)$ is related to a generalization of the base-generating function for matroids [5]. By establishing a so-called deletion-contraction recurrence, we show that $T^k_G(x,y)$ has similar expressive power as its classical counterpart $T_G(x,y)$.

Main Theorem 6. *For any graph G and positive integer k, the following graph invariants are encoded in $T^k_G(x,y)$:*

1. *$T^k_G(2,1)$ is the number of k-forests in G. In other words $T^k_G(2,1)$ corresponds to the problem $\#$EdgeSub(Φ) for the property Φ of being a forest.*
2. *For each positive integer c, the values of $T^k_G(1-c,0)$ determine*[6] *the numbers of pairs (A,σ), where A is a k-edge subset of G, and σ is a proper c-coloring of $(V(G),A)$.*
3. *From $T^k_G(2,0)$ we can compute the numbers of pairs $(A,\vec{\eta})$, where A is a k-edge subset of G and $\vec{\eta}$ is an acyclic orientation of $(V(G),A)$.*
4. *$T^k_G(2,0)$ also determines the number of k-edge subsets A of G such that $(V(G),A)$ has even Betti number (we give a formal definition of the Betti number in section* 9.1*).*
5. *$T^k_G(0,2)$ determines the number of k-edge subsets A of G such that $(V(G),A)$ has an even number of components.*

Note that, while $\#$EdgeSub(Φ) only allows us to count the number of subgraphs with k edges that satisfy Φ, the parameterized Tutte polynomial allows us to count more intricate objects, such as tuples of an edge-subset and a coloring (or acyclic orientation) on the induced graph. From a complexity theoretic point of view, we obtain a similar result as [49], albeit only for rational coordinates: for each *fixed* pair (x,y) of coordinates, we consider the problem receiving as input a graph G and a positive integer k and computing $T^k_G(x,y)$. Following the paradigm of this work, we choose k as a parameter; that is, we consider inputs in which k is significantly smaller than $|G|$.

Main Theorem 7. *Let (x,y) denote a pair of rational numbers. The problem of computing $T^k_G(x,y)$ is solvable in polynomial time if $x=y=1$ or $(x-1)(y-1)=1$, fixed-parameter tractable, but $\#$P-hard if $x=1$ and $y \neq 1$ and $\#$W[1]-hard otherwise.*

The class $\#$P is the counting version of NP [89, 90], and, in particular, the $\#$P-hard cases in the above classification are not polynomial-time tractable unless the

[6]They are equal up to trivial modifications; in particular, their complexities coincide.

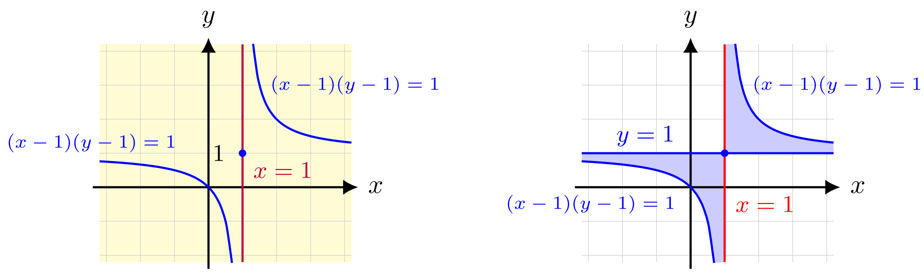

(a) Points of the parameterized Tutte polynomial that can be computed in polynomial-time (blue) and that are fixed-parameter tractable, but #P-hard (red). Exact computation at any other point (yellow) is #W[1]-hard.

(b) Points of the parameterized Tutte polynomial that allow for an FPRAS (blue) and for an FPTRAS (red); all points on the boundary of the blue area are included. The complexity of approximation is open for all points outside of the coloured region.

FIG. 1. *Points of the parameterized Tutte polynomial that can be computed in fixed-parameter tractable time* (a) *exactly or* (b) *approximately. We emphasize that a full classification for exact counting is established, while the complexity of approximation remains open outside of the colored area. (Color available online.)*

polynomial-time hierarchy collapses to P [87]. Consider Figure 1 for a depiction of the classification. Note that Main Theorem 7 yields #W[1]-hardness for each of the aforementioned problems from Main Theorem 6. Note further that the tractable cases are similar but not equal to the classical counterpart [49].

Moreover, our proof uses entirely different tools than [49] and illustrates the power and utility of the method presented in the subsequent discussion of our techniques.

Having fully classified the complexity of exact evaluation of the parameterized Tutte polynomial, we also consider the complexity of approximate evaluation. We identify two regions bounded by the hyperbola $(x-1)(y-1)=1$ and the lines $x=1$ and $y=1$ as efficiently approximable; consider Figure 1(b) for a depiction.

MAIN THEOREM 8. *Let (x,y) denote a pair of rational numbers satisfying the condition $0 \leq (x-1)(y-1) \leq 1$. Then $T_G^k(x,y)$ has an FPTRAS. If additionally $x \neq 1$ or $y=1$, then $T_G^k(x,y)$ even has an FPRAS.*

1.2. Technical overview. Our Main Theorems 4, 5, and 8 are obtained easily: the proof of Main Theorem 4 is a standard application (see, for instance, [65]) of the Monte Carlo approach, in combination with Ramsey's theorem, and of Arvind and Raman's algorithm for approximately counting subgraphs of bounded treewidth [6]. The proof of Main Theorem 5 uses a standard parameterized win-win approach for graphs of bounded treewidth or bounded degree. Finally, the proof of Main Theorem 8 is an easy consequence of the work of Anari et al. [5] on approximate counting via log-concave polynomials.

Hence, in this discussion, we want to focus on our main technical results on exact counting as well as on the construction of novel low-degree Cayley graph expanders, which, in combination, will enable us to prove the lower bounds for Main Theorems 1 and 2 and, perhaps surprisingly, also for Main Theorem 7.

As a main component, we use the complexity monotonicity framework of Curticapean, Dell, and Marx [23]. Given a property Φ and a positive integer k, we write $\#\mathsf{EdgeSub}(\Phi, k \to \star)$ for the function that maps a graph G to the number of k-edge subgraphs of G that satisfy Φ. Using a well-known transformation via Möbius inversion [60, Chapter 5.2], we can show that there exists a function of finite support $a_{\Phi,k}$ from graphs to rational numbers such that for every graph G we have

$$\#\mathsf{EdgeSub}(\Phi, k \to G) = \sum_H a_{\Phi,k}(H) \cdot \#\mathsf{Hom}(H \to G)\,, \tag{1.1}$$

where $\#\mathsf{Hom}(H \to G)$ is the number of graph homomorphisms from H to G. In other words, we can express $\#\mathsf{EdgeSub}(\Phi, k \to \star)$ as a finite linear combination of homomorphism counts. Here, we can then apply the complexity monotonicity framework [25], which asserts that computing a finite linear combination of homomorphism counts is *precisely as hard as* its hardest term (among the terms with a nonzero coefficient). However, the complexity of computing the number of homomorphisms from small pattern graphs to large host graphs is very well understood [28, 63]. Roughly speaking, the higher the treewidth of the pattern graph, the harder the problem becomes; we make this formal in section 2.2.

Instead of our original problem $\#\textsc{EdgeSub}(\Phi)$, we can thus consider the problem of computing linear combinations of graph homomorphism counts. In particular, to obtain hardness, it suffices to understand for which graphs H the coefficient $a_{\Phi,k}(H)$ in (1.1) is nonzero.

Relying on the well-known fact that the Möbius function of the partition lattice alternates in sign, Curticapean, Dell, and Marx [25] observed that nontrivial cancelations cannot occur in (1.1) if, for each k, every k-edge graph that satisfies Φ must have the same number of vertices. Consequently, if the matching number is unbounded, those properties yield $\#\mathsf{W}[1]$-hardness. An example for such a property is the case of $\Phi(H) = 1$ if and only if H is a tree. In contrast, the intractability result for the case of $\Phi =$ acyclicity (that is, being a forest) turned out to be much harder to show [13], indicated by connections to parameterized counting problems in matroid theory.

More generally, 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, unfortunately, to be a difficult problem in its own right. For example, in the case of the vertex-induced subgraph counting problem, the coefficient of the clique is the reduced Euler characteristic of a simplicial graph complex [78], the coefficient of the biclique is the so-called alternating enumerator [32], 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 [79]. In all of the previous works mentioned here, the complexity analysis of the respective pattern counting problems therefore amounted to understanding the cancelation behavior of those invariants. To do so, the papers used tools from combinatorial commutative algebra and, to some extent, topological fixed-point theorems.

In this work, we provide additional insights into said coefficients $a_{\Phi,k}(H)$. More precisely, we show that the coefficients of high-treewidth low-degree vertex transitive graphs can be analyzed much more easily than generic graphs of high treewidth such as the clique or the biclique. First, we prove that the coefficient of a graph H with k edges in (1.1) is equal to the *indicator* of Φ and H, defined as follows:[7]

$$a(\Phi, H) := \sum_{\sigma \in \mathcal{L}(\Phi,H)} \prod_{v \in V(H)} (-1)^{|\sigma_v|-1}(|\sigma_v| - 1)!\,, \tag{1.2}$$

[7] To be precise, the identity in (1.2) will be obtained for a colored version of $\#\textsc{EdgeSub}(\Phi)$, which we show to be interreducible with the uncolored version.

FIG. 2. *Illustration of the construction of a fractured graph. The left picture shows a vertex v of a graph H with incident edges* $E_H(v) = \{\bullet, \bullet, \bullet, \bullet, \bullet, \bullet\}$. *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* $B_1 = \{\bullet, \bullet, \bullet\}$ *and* $B_2 = \{\bullet, \bullet, \bullet\}$. *(Color available online.)*

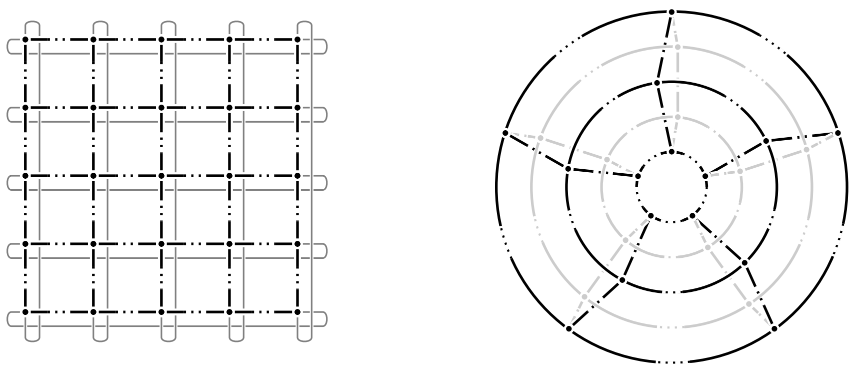

FIG. 3. *Two isomorphic representations of the toroidal grid $\circledcirc_\ell$: on the left-hand side as a grid with connected endpoints, on the right-hand side as a stylized torus.*

where $\mathcal{L}(\Phi, H)$ 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 by splitting each vertex $v \in V(H)$ according to σ_v; an illustration is provided in Figure 2.

As a consequence, the $\#\mathsf{W}[1]$-hardness for $\#\textsc{EdgeSub}(\Phi)$ 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 ETH will, additionally, require sparsity of the graphs.

Recall further 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 will use Cayley graphs as a natural class of vertex transitive graphs: The *Cayley graph* of a group Γ together with a symmetric generating set[8] $S \subseteq \Gamma$ 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. For example, setting $\Gamma = \mathbb{Z}/\ell\mathbb{Z} \times \mathbb{Z}/\ell\mathbb{Z}$ to be the twofold direct product of the group of integers modulo ℓ, and setting $S = \{(1,0), (-1,0), (0,1), (0,-1)\}$, the Cayley graph $\mathcal{C}(\Gamma, S)$ becomes the toroidal grid $\circledcirc_\ell$; a depiction is given in Figure 3.

[8]This means a subset S of the group Γ that generates this group and satisfies $S^{-1} = S$.

Cayley graphs $H=\mathcal{C}(\Gamma,S)$ have a vertex transitive symmetry group. The group Γ acts on the graph by letting $g\in\Gamma$ send the vertex $v\in V(H)=\Gamma$ to gv. This preserves edges because $g(xs)=(gx)s$ for all $g,x,s\in\Gamma$. 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 (1.2) are shown to be invariant under this action, the group Γ naturally permutes these summands.

Now let us assume that Γ is a p-group and that we are given a Cayley graph $H=\mathcal{C}(\Gamma,S)$ with $|E(H)|=k$. Since our Cayley graph H arises from a p-group Γ, it follows that when evaluating the indicator $a(\Phi,H)$ modulo p, only those contributions from fractures fixed under Γ 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. Consequently, the evaluation of the indicator $a(\Phi,H)$ modulo p requires us to only consider as many terms as there are partitions of the generating set S. In particular, if H has constant degree, then $|S|$ and hence the number of partitions of S must be bounded by a constant as well.

To illustrate the power of this strategy, we provide an explicit consequence when applied to the special case of toroidal grids (which satisfy $|S|=4$): Let us write M_k for the matching of size k, P_2 for the path consisting of 2 edges, C_k for the cycle of length k, and S_k for a sun (a cycle with dangling edges) of size k, and recall that $\circledcirc_k$ denotes the toroidal grid of size k.

THEOREM 1.4. *Let Φ denote a computable graph property, and assume that infinitely many primes ℓ satisfy the equation*[9]

$$\begin{aligned}(1.3)\qquad &-6\Phi(M_{2\ell^2})+4\Phi(M_{\ell^2}+\ell C_\ell)+8\Phi(\ell^2P_2)-\Phi(2\ell C_\ell)\\ &\qquad -2\Phi(\ell C_{2\ell})-4\Phi(\ell S_\ell)+\Phi(\circledcirc_\ell)\neq 0 \mod \ell.\end{aligned}$$

Then #EDGESUB(Φ) *is* #W[1]*-hard.*

As a toy example for an application of Theorem 1.4, let us consider the property Φ of being connected. Observe that among the graphs in (1.3), only $\circledcirc_\ell$ is connected, and thus the sum is always 1 for $\ell\geq 2$. Thus, indeed the left-hand side of (1.3) is nonzero, proving that #EDGESUB(Φ) is #W[1]-hard.

For our general results, especially for the fine-grained lower bounds, we need better classes of Cayley graphs. Let us summarize the necessary conditions: Our Cayley graphs must be

- sparse and of high treewidth (required for hardness) and
- generated by a p-group and of constant degree (required for the analysis of the indicators).

The natural choice of Cayley graphs satisfying both constraints are p-group Cayley graph expanders of low degree. An optimal construction (w.r.t. the degree) for the case $p=2$ is known and due to a subset of the authors [71]. However, we will see that 2-group Cayley graph expanders are not sufficient to prove our full classification. For this reason, we provide a novel construction of low-degree p-group Cayley graph expanders in this work. For the sake of presentation, we decided to encapsulate the treatment of our constructions in separate sections, both in the introduction 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 and

[9] We write $+$ for (disjoint) graph union and ℓH for the graph consisting of ℓ disjoint copies of H. Further, we set $\Phi(H)=1$ if H satisfies Φ and $\Phi(H)=0$ otherwise.

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 Main Theorem 3.

Construction of low-degree Cayley graph expanders. We prove Main Theorem 3 via an explicit construction of the groups Γ_i and the symmetric generating sets S_i in section 3 motivated by number theoretic objects.

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 [81] these groups were realized by mapping the generators a_k, b_j to explicit generalized quaternions, leading ultimately to an explicit injective group homomorphism

$$\Psi \colon \Gamma_{p;\alpha,\beta} \to \mathrm{GL}_3(\mathbb{F}_p[[t]])\,. \tag{1.4}$$

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 subsection 3.3 but could also be made explicit for any $p \geq 5$. Since the applications do not depend on concrete matrices, we merely state the existence of the injective group homomorphism Ψ.

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 with respect to $N_i = \ker(\pi_i \circ \Psi)$. Using results from [81], 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 [81] by Rungtanapirom and two of the authors, and it is based on a similar approach in the classical papers by Lubotzky, Phillips, and Sarnak [61] and by Morgenstern [69]. We note here that the results of [81] 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 program for GL_2.

At this point we have proven Main Theorem 3 for the particular valency $2d = 2(p+1)$. In order to obtain the more general valencies stated in the theorem, we recall in subsection 3.1 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 3.8 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 [85] on the local permutation structure of the underlying geometric square complex. To improve even further for $p=3$ we consider in subsection 3.3 a concrete presentation of $\Gamma_{3;0,1}$ which in Theorem 3.7 is shown to reduce to 2 generators. For $p=5$, the explicit Example 3.9 achieves a reduction to 2 generators for $\Gamma_{5;0,2}$. It follows again from the theory recalled in subsection 3.1 that by adding generators (as necessary) we obtain Main 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 Main Theorem 3 is computable. We also would like to emphasize again that the expanders constructed for the proof of Main 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.

2. Preliminaries.

Given a finite set S, we write $|S|$ and $\#S$ for the cardinality of S. Further, given a function $f : X \times Y \to Z$ and an element $x \in X$, we write $f(x,\star): B \to C$ for the function $y \mapsto f(x,y)$.

2.1. Graphs and homomorphisms.

Graphs in this work are simple and irreflexive; that is, we do not allow multiple edges or self-loops. Given a graph G and a subset A of $E(G)$, we write $G(A)$ for the graph $(V(G),A)$, and we write $G[A]$ for the graph obtained from $G(A)$ by deleting all isolated vertices. The degree of a graph is the maximum degree of its vertices, and we write $H+G$ for the union of H and G; formally, we assume that $V(G)\cap V(H)=\emptyset$ and set $H+G := (V(H)\cup V(G), E(H)\cup E(G))$.

Given graphs F and G, a *homomorphism* from F to G is an edge-preserving mapping $\varphi : V(F) \to V(G)$; that is, for each edge $\{u,v\} \in E(F)$ we have that $\{\varphi(u),\varphi(v)\} \in E(G)$. A homomorphism is called an *embedding* if it is injective (on the vertices). We write $\mathsf{Hom}(F \to G)$ and $\mathsf{Emb}(F \to G)$ for the set of all homomorphisms and embeddings, respectively, from F to G.

An *isomorphism* between two graphs F and G is a bijective embedding φ satisfying the stronger constraint $\{u,v\} \in E(F) \Leftrightarrow \{\varphi(u),\varphi(v)\} \in E(G)$. We say that F and G are isomorphic, denoted by $F \cong G$, if an isomorphism from F to G exists. An isomorphism from a graph F to itself is called an *automorphism,* and we write $\mathsf{Aut}(F)$ for the group formed by such automorphisms (where the group operation is the composition of automorphisms).

A *graph property* Φ is a function from graphs to $\{0,1\}$ which is invariant under isomorphisms, that is, $\Phi(H) = \Phi(F)$ whenever $H \cong F$. We say that a graph H satisfies Φ if $\Phi(H)=1$.

A graph F is a *minor* of a graph H if it can be obtained from H by a sequence of edge- and vertex-deletions and edge-contractions (with multiple edges and self-loops deleted). We write $F \prec H$ if F is a minor of H. A graph property Φ is *minor-closed* if $\Phi(H)=1$ implies $\Phi(F)=1$ whenever $F \prec H$. Every minor-closed graph property Φ is defined by a set $\mathcal{F}(\Phi)$ of minimal (w.r.t. $\prec$) forbidden minors, that is, $\Phi(H)=1$ if and only if $F \not\prec H$ for each $F \in \mathcal{F}(\Phi)$. The celebrated Robertson–Seymour theorem [75] implies that $\mathcal{F}(\Phi)$ is always finite. In particular, the latter implies that each minor-closed graph property is (polynomial-time) computable. For the purpose of this work, forbidden minors of degree at most two will be of particular importance; therefore, we define $\mathcal{F}_2(\Phi)$ to be the subset of $\mathcal{F}(\Phi)$ containing only the graphs of degree at most 2.

A graph G is called *k-edge-colored* if the edges of G are colored with (at most) k pairwise different colors. Given a homomorphism $\varphi \in \mathsf{Hom}(G \to H)$ for some graphs G and H, we also call φ an *H-coloring*. Moreover, an *H-colored graph* is a pair of a graph G and an H-coloring φ. For ease of readability, we say that a graph G is H-colored if the H-coloring is implicit or clear from the context. Observe that an H-coloring φ of a graph G induces a $\#E(H)$-edge-coloring by mapping an edge $\{u,v\} \in E(G)$ to the color $\{\varphi(u), \varphi(v)\}$. Throughout this work, we use the following notion of homomorphisms between H-colored graphs.

DEFINITION 2.1. *Let F and G denote two H-colored graphs, and let c_F and c_G denote their H-colorings. A homomorphism φ from F to G is called color-preserving if $c_G(\varphi(v)) = c_F(v)$ for every $v \in V(F)$. We write* $\mathsf{Hom}_{\mathsf{cp}}(F \to_H G)$ *for the set of all color-preserving homomorphisms from F to G. Color-preserving embeddings and the set* $\mathsf{Emb}_{\mathsf{cp}}(F \to_H G)$ *are defined similarly.*

Further, two H-colored graphs F and G are isomorphic as H-colored graphs, denoted by $F \cong_H G$, if there is a color-preserving isomorphism from F to G.

Note that, given two H-colored graphs F and G, we write $F \cong G$ (rather than $F \cong_H G$) if the underlying *uncolored* graphs are isomorphic.

For the treatment of decision and approximate counting, we introduce the following classification criteria on (computable) graph properties. To this end, we write $K_{\ell,r}$ for the biclique with ℓ vertices on the left and r vertices on the right side, respectively. In particular, $K_{1,r}$ denotes the star of size r.

DEFINITION 2.2. *Let Φ denote a computable graph property. We say that*

- *Φ satisfies the matching criterion if $\Phi(M_k) = 1$ for all but finitely many k,*
- *Φ satisfies the star criterion if $\Phi(K_{1,k}) = 1$ for all but finitely many k,*
- *Φ has bounded treewidth if there is a constant B such that Φ is false on all graphs of treewidth at least B.*

For example, the properties of being bipartite or planar satisfy both the matching and the star criteria. Furthermore, the property of being 2-regular has bounded treewidth, while the criterion of just being regular satisfies only the matching criterion. Further, the property of being a tree is an example that both satisfies the star criterion and is of bounded treewidth, while the property of being a forest satisfies all three criteria.

Expander graphs. All (almost-tight) conditional lower bounds in this work rely on the existence of certain (families of) expander graphs. Given a positive integer d, a rational $c > 0$, and a class of graphs $\mathcal{G} = \{G_1, G_2, \ldots\}$ with $\#V(G_i) = n_i$, we call $\mathcal{G}$ a family of *(n_i, d, c)-expanders* if, for all i, the graph G_i is d-regular and satisfies

$$\forall X \subseteq V(G_i) : |S(X)| \geq c\left(1 - \frac{|X|}{n_i}\right)|X|,$$

where $S(X)$ is the set of all vertices in $V(G_i) \setminus X$ that are adjacent to a vertex in X.

While being sparse due to d-regularity, expander graphs have treewidth[10] linear in the number of vertices (see, for instance, Proposition 1 in [45], and set $\alpha = 1/2$). Furthermore, they admit arbitrarily large clique minors [54]. Formally, we have the following fact.

FACT 2.3. *Fix a rational c and an integer d, and let $\mathcal{G}$ denote a family of (n_i, d, c)-expanders. Then, $\#E(G_i) \in \Theta(n_i)$ and $\mathsf{tw}(G_i) \in \Theta(n_i)$. Furthermore, for each*

[10] We use the graph parameter "treewidth" (tw) in a black-box manner only and refer the reader to, for instance, Chapter 7 of [27] for a detailed exposition.

positive integer k there is an index j such that for all $i \geq j$, the graph G_i contains the complete graph on k vertices as a minor.

In section 3.1 we recall a spectral reformulation of the expander property in terms of the nontrivial eigenvalues of the graph Laplace operator on $\mathbb{R}$-valued functions on the set of vertices.

2.2. Parameterized complexity theory. A *parameterized counting problem* is a pair of a counting problem $P : \Sigma^* \to \mathbb{N}$ and a *parameterization* $\kappa : \Sigma^* \to \mathbb{N}$. A *parameterized decision problem* is a pair (P, κ) of a decision problem $P : \Sigma^* \to \{0, 1\}$ and a parameterization κ. Consider, for example, the problems #Clique and Clique: on input of a graph G and a positive integer k, the task is either to compute the number of k-cliques in G or to detect the mere existence of a k-clique, respectively. The parameterization is given by $\kappa(G, k) := k$ for both problems.

A parameterized problem (P, κ) is called *fixed-parameter tractable* if there is a computable function f such that P can be computed in time $f(\kappa(x)) \cdot |x|^{O(1)}$. For historic reasons, the class of all fixed-parameter tractable *decision* problems is called FPT.[11] Furthermore, a *parameterized Turing reduction* from (P, κ) to $(\hat{P}, \hat{\kappa})$ is a Turing reduction from P to $\hat{P}$ that, on input x, runs in time $f(\kappa(x)) \cdot |x|^{O(1)}$ and additionally satisfies $\hat{\kappa}(y) \leq f(\kappa(x))$ for each oracle query y. Again, f only needs to be some fixed computable function. We write $(P, \kappa) \leq^{\mathrm{fpt}}_{\mathrm{T}} (\hat{P}, \hat{\kappa})$ if a parameterized Turing reduction exists.

A parameterized counting problem is #W[1]-*hard* if it can be reduced from #Clique, and, similarly, a parameterized decision problem is W[1]-*hard* if it can be reduced from Clique, both with respect to parameterized Turing reductions.

Under reasonable assumptions, such as the ETH [48] defined below, #W[1]- and W[1]-hard problems are not fixed-parameter tractable.[12]

Conjecture 2.4 (exponential time hypothesis). *The ETH asserts that* 3-SAT *cannot be solved in time* $\exp(o(n))$*, where n is the number of variables of the input formula.*

Our hardness results in this paper are obtained by reducing from the problem #Hom($\mathcal{H}$). Given a fixed class of graphs $\mathcal{H}$, in the problem #Hom($\mathcal{H}$) the input is a graph $H \in \mathcal{H}$ and an arbitrary graph G, and the task is to compute the number of homomorphisms from H to G; the parameter is $|H|$. Dalmau and Jonsson [28] established an exhaustive classification for this problem, stating that #Hom($\mathcal{H}$) is fixed-parameter tractable if the treewidth of graphs in $\mathcal{H}$ is bounded by a constant and #W[1]-hard otherwise.

Let Φ denote a graph property, that is, a function from (isomorphism classes) of graphs to $\{0, 1\}$. Setting

$$\mathsf{EdgeSub}(\Phi, k \to G) := \{A \subseteq E(G) | \#A = k \wedge \Phi(G[A]) = 1\},$$

we define #EdgeSub(Φ) as the parameterized counting problem in which on input of a graph G and a positive integer k, the task is to compute $\#\mathsf{EdgeSub}(\Phi, k \to G)$; the parameter is k.

[11] In some literature FPT is used for both parameterized decision and counting problems, while some authors write FFPT for the class of all fixed-parameter tractable parameterized counting problems.

[12] In fact, Chen et al. [16, 17] proved the much stronger statement that #Clique cannot be solved in time $f(k) \cdot |V(G)|^{o(k)}$ for any function f, unless ETH fails.

In this paper, we often rely on the following important but easy observation: write Φ_k for the set of graphs H with k edges and without isolated vertices that satisfy Φ. Then we have

$$\#\mathsf{EdgeSub}(\Phi, k \to G) = \sum_{H \in \Phi_k} \#\mathsf{Sub}(H \to G)\,, \tag{2.1}$$

where $\#\mathsf{Sub}(H \to G)$ is the number of subgraphs of G that are isomorphic to H.

Using the aforementioned transformation, both Proposition 1.2 and Proposition 1.3 are easy consequences of known results regarding the subgraph counting problem. We add their proofs only for the sake of completeness.

Proposition 2.5. *Let Φ denote a computable graph property. There exists a computable function f such that $\#\textsc{EdgeSub}(\Phi)$ can be solved in time*

$$f(k) \cdot |V(G)|^{0.174k+o(k)}\,.$$

Proof. The fastest known algorithm for computing $\#\mathsf{Sub}(H \to G)$ for a k-edge graph H runs in time $k^{O(k)} \cdot |V(G)|^{0.174k+o(k)}$ and is due to Curticapean, Dell, and Marx [25]. Now observe that the size of Φ_k is bounded by a function in k, since graphs in Φ_k have k edges and no isolated vertices and thus have at most $2k$ vertices. Consequently, their algorithm extends to $\#\textsc{EdgeSub}(\Phi)$ by computing the number $\#\mathsf{EdgeSub}(\Phi, k \to G)$ as given in (2.1). ∎

Note that the growth of f in the previous result depends, among other factors, on the complexity of verifying Φ.

For the following, recall that a property Φ has *bounded matching number* if there is a constant c such that Φ is false on all graphs containing a matching of size at least c. Furthermore, write $\neg\Phi$ for the complement of Φ.

Proposition 2.6. *Let Φ denote a computable graph property satisfying that there is $M > 0$ such that for all k either the graphs with k edges satisfying Φ or the graphs with k edges satisfying $\neg\Phi$ have matching number bounded by M. Then $\#\textsc{EdgeSub}(\Phi)$ is fixed-parameter tractable.*

Proof. Applying (2.1), we observe that counting subgraphs isomorphic to H is fixed-parameter tractable (even polynomial-time solvable) if there is a constant upper bound on the size of the largest matching of H [26]. This allows us to compute $\#\mathsf{EdgeSub}(\Phi, k \to G)$ in the desired running time if the graphs in Φ_k have matching number bounded by M. In the case that the latter is true for $\neg\Phi_k$ instead, we use the fact that $\#\mathsf{EdgeSub}(\Phi, k \to G) = \binom{\#E(G)}{k} - \#\mathsf{EdgeSub}(\neg\Phi, k \to G)$ and proceed similarly. ∎

2.3. Combinatorial commutative algebra. We assume familiarity with the notions of basic group theory and refer the reader to, for instance, [57] for a detailed introduction. Given a positive integer ℓ, we write $\mathbb{Z}_\ell$ for the group of integers modulo ℓ, and we write $\mathbb{Z}_\ell^k$ for the k-fold *direct product* of $\mathbb{Z}_\ell$; recall that the binary operation of the direct product is defined coordinatewise.

For a prime p, recall that a finite group G is called a p-group if the order $\#G$ of G is a power of p. Recall that by Lagrange's theorem, this implies that the order of any subgroup H of G is likewise a power of p.

Given a group G and a subgroup $H \subseteq G$, we write G/H for the set of *left cosets* of H. Formally, a left coset of H is an equivalence class of the following equivalence relation on G: two elements $g, g' \in G$ are equivalent if and only if $g' = gh$ for some

$h \in H$. We write gH for the equivalence class of $g \in G$. We define the *index* of H in G as the cardinality $[G : H] = \#G/H$ of the set of left cosets. Note that $[G : H]$ might be finite even though $\#G$ is infinite. The index satisfies the basic identity $\#G = [G : H] \cdot \#H$, and again, with a slight abuse of notation, we observe that the identity remains well defined in the infinite case: $\#G$ is infinite if and only if one of $[G : H]$ or $\#H$ is infinite. If the subgroup H is *normal* in G (that is, for each $g \in G$ we have that the subset $gHg^{-1} = \{ghg^{-1} | h \in H\}$ of G is equal to H), then the set G/H naturally carries the structure of a group, with the group operation defined by $(g_1 H) \cdot (g_2 H) = (g_1 \cdot g_2)H$. In this case, we call G/H a *quotient group*.

Given an element $g \in G$ we write

$$\langle g \rangle = \{g^a : a \in \mathbb{Z}\} \subseteq G$$

for the *subgroup generated by* g (recall that for a negative integer a we define g^a as the $|a|$th power of the inverse element of g). If there is a positive integer a such that g^a equals the neutral element of G, we define the *order* $\mathrm{ord}_G(g)$ of g as the smallest such positive integer a (and set $\mathrm{ord}_G(g) = \infty$ otherwise). If g has finite order, the subgroup $\langle g \rangle$ of G generated by g is isomorphic to the cyclic group $\mathbb{Z}_{\mathrm{ord}_G(g)}$.

Given a finite group Γ and a symmetric set S of generators of Γ, the *Cayley graph* of Γ and S, denoted by $\mathcal{C}(\Gamma, S)$, has as vertices the elements of Γ, and two vertices u and v are adjacent[13] if there is an $s \in S$ such that $v = us$. For example, the Cayley graph $\mathcal{C}(\mathbb{Z}/\ell\mathbb{Z}, \{1, -1\})$ is isomorphic to the cycle of length ℓ.

Möbius inversion and the partition lattice. We follow the notation of the standard textbook of Stanley [84]. Given a finite partially ordered set $(L, \leq)$ and a function $f : L \to \mathbb{Q}$, the *zeta transformation* $\zeta f : L \to \mathbb{Q}$ is defined as

$$\zeta f(\sigma) := \sum_{\rho \geq \sigma} f(\rho)\,.$$

The principle of Möbius inversion allows us to invert a zeta transformation; a proof of the following theorem can be found in [84, Chapter 3].

Theorem 2.7. *Given a partially ordered set* $(L, \leq)$*, there is a computable function* $\mu : L \times L \to \mathbb{Z}$*, called the Möbius function such that for all* $f : L \to \mathbb{Q}$ *and* $\sigma \in L$ *we have*

$$f(\sigma) = \sum_{\rho \geq \sigma} \mu(\sigma, \rho) \cdot \zeta f(\rho)\,.$$

We use Möbius inversion on the ordering of partitions of finite sets. Given two partitions σ and ρ of a finite set S, we say that σ *refines* ρ if every block of σ is a subset of a block of ρ, and in this case we write $\sigma \leq \rho$. This induces a partial order, called the *partition lattice*[14] of S. The explicit formula of the Möbius function over the partition lattice is of particular importance in this work.

Theorem 2.8 (see Chapter 3 in [84]). *Let* σ *denote a partition of a finite set* S*. We have*

$$\mu(\sigma, \top) = (-1)^{|\sigma|-1} \cdot (|\sigma| - 1)!\,,$$

where $\top = \{S\}$ *denotes the coarsest partition.*

[13]Note that, in some literature, Cayley graphs are colored and directed. However, we only consider the underlying uncolored and simple graph.

[14]A lattice is a partially ordered set in which every pair of elements has a least upper bound and a greatest lower bound. For a formal definition we refer to [84].

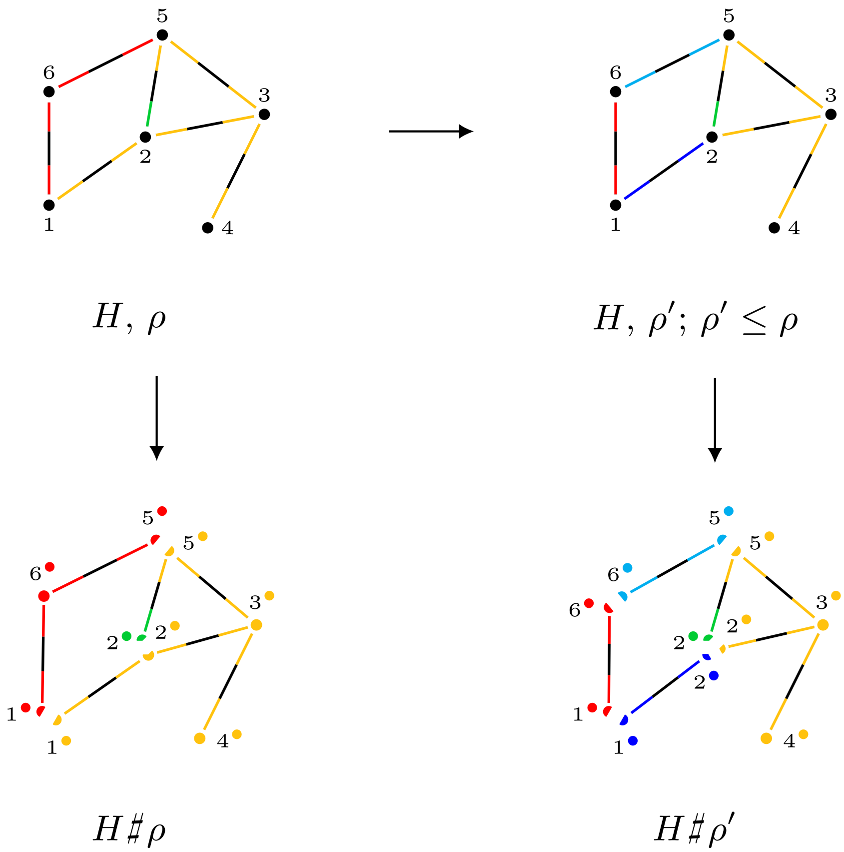

FIG. 4. *Two fractures ρ and $\rho' \leq \rho$ of a graph H (depicted by the colors of the outgoing edges of the vertices) and the corresponding fractured graphs $H\sharp\rho$ and $H\sharp\rho'$. (Color available online.)*

Fractured graphs.

DEFINITION 2.9 (fractures). *Let H denote a graph. A fracture of H is a tuple*

$$\rho = (\rho_v)_{v\in V(H)}\,,$$

where ρ_v is a partition of the set of edges $E_H(v)$ of H incident to v.

Note that the set of all fractures of H, denoted by $\mathcal{L}(H)$, is a lattice that is isomorphic to the (pointwise) product of the partition lattices of $E_H(v)$ for each $v \in V(H)$. In particular, we write $\sigma \leq \rho$ if, for each $v \in V(H)$, the partition σ_v refines the partition ρ_v. Consider Figure 4 for a visualization of a fracture and its refinement.

Note further that a fracture describes how to split (or *fracture*) each vertex of a given graph: for each vertex v, create a vertex v^B for each block B in the partition ρ_v; edges originally incident to v are made incident to v^B if and only if they are contained in B. We call the resulting graph the *fractured graph* $H\sharp\rho$; a formal definition is given in Definition 2.10, and a visualization is given in Figure 4.

DEFINITION 2.10 (fractured graph $H\sharp\rho$). *Given a graph H, we consider the matching M_H containing one edge for each edge of H; formally,*

$$V(M_H) := \bigcup_{e=\{u,v\}\in E(H)} \{u_e, v_e\} \quad \text{and} \quad E(M_H) := \{\{u_e, v_e\} \mid e = \{u,v\} \in E(H)\}.$$

For a fracture ρ of H, we define[15] *the graph $H\sharp\rho$ as the quotient graph of M_H under the equivalence relation on $V(M_H)$ which identifies two vertices v_e, w_f of M_H if and only if $v = w$ and e, f are in the same block B of the partition ρ_v of $E_H(v)$. We write v^B for the vertex of $H\sharp\rho$ given by the equivalence class of the vertices v_e (for which $e \in B$) of M_H.*

The fractured graph $H\sharp\rho$ comes with a natural H-coloring. Indeed, the homomorphism $M_H \to H$ which sends $v_e \in V(M_H)$ to $v \in V(H)$ descends to $H\sharp\rho$ so that we always have a canonical diagram $M_H \to H\sharp\rho \to H$ of graph homomorphisms. For example, for any graph H, the fracture $\perp$, with $\perp_v := \{\{e\} \mid e \in E_H(v)\}$, induces the fractured graph $H\sharp\perp = M_H$; the fracture $\top$, with $\top_v := \{E_H(v)\}$, induces the fractured graph $H\sharp\top = H$. The fractures $\perp, \top$ are the minimal and maximal elements of the lattice $\mathcal{L}(H)$, respectively.

Given a graph property Φ and a graph H, we write $\mathcal{L}(\Phi, H)$ for the set of all fractures ρ of H that satisfy $\Phi(H\sharp\rho) = 1$.

3. Construction of Cayley graph expanders.

As mentioned before, this section encapsulates our novel construction of p-group Cayley graph expanders with low degree. Readers only interested in the application of those Cayley graphs in our hardness proofs can safely skip this section. Recall that the *Cayley graph* of a group Γ is generated by a symmetric set $S \subseteq \Gamma$, i.e., such that $S^{-1} = S$, 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.

In this section we recall the construction of certain discrete groups from [81, 85] that are lattices in generalized quaternion algebras. Their group theory is controlled by representations with values in power series rings with coefficients in $\mathbb{F}_p$. In particular, these representations lead to well-chosen finite congruence quotients that are p-groups. With a natural set of generators as in [81, 85] these lead to sequences of Cayley graph expanders of valency $2(p+1)$ and vertex set of size a power of p; here p is odd. We explain the known fact that a change of generators for the lattice does not destroy the expander property. Moreover, we analyze the relations and find that we can reduce the number of generators to $(p+3)/2$ (here p is odd), so the valency of the corresponding Cayley graph is $p+3$. For $p \geq 5$ (with an extra argument for $p = 5$) we deduce the existence of a series of Cayley graph expanders of p-power order and now valency $2(p-2)$ that is necessary for the applications in this paper. Theorem 3.10 gives a more precise statement for the range of possible valencies of our construction.

3.1. Change of generators in Cayley graphs.

This section is concerned with the expansion property of Cayley graphs under change of generators. In general, the expansion property is not preserved (see, e.g., [47, section 11.4] for a brief discussion). However, if the two families of generators remain bounded and can be mutually written in words of bounded lengths of each other, then the expansion property is preserved (see, e.g., [55, Proposition 3.5.1]). The remainder of this section is devoted to an easy self-contained proof in the following special case: Let Γ be a finitely generated

[15] The notation $H\sharp\rho$ stems from the fact that the symbol "$\sharp$" is commonly used for medical fractures.

infinite group generated by two finite symmetric sets of generators S_1 and S_2, that is, $S_1^{-1} = S_1$ and $S_2^{-1} = S_2$. Let Γ_k be finite index subgroups of Γ with $[\Gamma : \Gamma_k] \to \infty$, and let $H_k = \Gamma/\Gamma_k$ be the corresponding finite quotients. The sets S_i of generators of Γ can also be viewed as sets of generators of the quotients H_k. Then the following holds.

PROPOSITION 3.1 (see, e.g., [55, Proposition 3.5.1]). *If the Cayley graphs $\mathcal{C}(H_k, S_1)$ represent a family of expander graphs, then $\mathcal{C}(H_k, S_2)$ is also a family of expander graphs.*

The proof is based on a well-known spectral description of expander graphs: Let $G = (V, E)$ be a connected graph with vertex set V and edge set E. The degree of a vertex $x \in V$ is denoted by d_x. The Laplacian L_G of a function $f : V \to \mathbb{R}$ is defined as follows:

$$L_G f(x) = \sum_{y \sim x} (f(x) - f(y)) = d_x f(x) - \sum_{y \sim x} f(y)\,.$$

Then L_G is a symmetric operator with nonnegative real eigenvalues. Let $\mu_1(G) > 0$ be the smallest nonzero eigenvalue of L_G. Since Cayley graphs are regular graphs, there is a close relationship between the eigenvalues of L_G and the adjacency matrix A_G (see, e.g., [47, Lemma 4.7]), namely, $\mu_1(G) = \lambda_1(A_G) - \lambda_2(A_G)$, where $\lambda_1(A_G), \lambda_2(A_G)$ are the two largest eigenvalues of A_G. The Cayley graphs $\mathcal{C}(H_k, S_1)$ from above are an expander family if and only if there exists a positive constant $C > 0$ such that, for all $k \in \mathbb{N}$ (see, e.g., [47, Theorem 4.11]),

$$\mu_1(\mathcal{C}(H_k, S_1)) \geq C\,.$$

Via the Rayleigh quotient

$$\mathcal{R}(f) := \frac{\sum_{x \in V} f(x) \cdot L_G f(x)}{\sum_{x \in V} f(x)^2} = \frac{\sum_{\{x,y\} \in E} (f(x) - f(y))^2}{\sum_{x \in V} f(x)^2}$$

the eigenvalue $\mu_1(G)$ has the following variational description (see, e.g., [56, Proposition 1.82] and [55, Proposition 3.4.3]):

$$\mu_1(G) = \inf \Big\{ \mathcal{R}(f) : \sum_{x \in V} f(x) = 0, f \neq 0 \Big\}.$$

We aim to show that there exists a positive constant $K > 0$ such that, for all $k \in \mathbb{N}$,

$$\mu_1(\mathcal{C}(H_k, S_2)) \geq K \mu_1(\mathcal{C}(H_k, S_1))\,.$$

Then the expanding property of the family $\mathcal{C}(H_k, S_1)$ implies a similar expanding property of the family $\mathcal{C}(H_k, S_2)$, albeit with a different spectral expansion constant.

Before starting the proof, we first consider the description of the Rayleigh quotient $\mathcal{R}(f)$ in the case of a Cayley graph $\mathcal{C}(H, S)$: The vertex set of this graph is given by $V = H$, and the set of edges is given by $E = \{\{x, xs\} : x \in V, s \in S\}$, where E is understood as a multiset. Then the Rayleigh quotient for a function $f : H \to \mathbb{R}$ can be written as

$$\mathcal{R}_S(f) = \frac{1}{2} \frac{\sum_{x \in H} \sum_{s \in S} (f(x) - f(xs))^2}{\sum_{x \in H} f(x)^2}.$$

The Rayleigh quotient is equipped with the index S, that is, the set of generators of the Cayley graph, since it depends on this choice of generators. Note that $\mathcal{C}(H,S)$ is vertex transitive since H acts on the vertex set $V=H$ by group left-multiplication. The proof is complete if there exists a constant $K>0$ such that, for all $k\in\mathbb{N}$ and all $f: H_k\to\mathbb{R}$,

$$\mathcal{R}_{S_2}(f)\geq K\mathcal{R}_{S_1}(f). \tag{3.1}$$

This is the case if there exists, for every generator $s\in S_1$, a constant $K'(s)>0$ such that

$$\sum_{x\in H_k}(f(x)-f(xs))^2\leq K'(s)\sum_{x\in H_k}\sum_{t\in S_2}(f(x)-f(xt))^2 \tag{3.2}$$

and that this constant $K'(s)$ does not depend on $k\in\mathbb{N}$. Since S_2 is a set of generators of Γ, any $s\in S_1\subset\Gamma$ can be written in the form

$$s=t_1t_2\cdots t_n$$

with $t_1,\ldots,t_n\in S_2$. The same relation between the generators holds in each of the quotients H_k, $k\in\mathbb{N}$. We abbreviate $s_j=t_1\cdots t_j$, so $s_0=1$ and $s_n=s$, and $s_j=s_{j-1}t_j$. Let us now show (3.2):

$$\begin{aligned}\sum_{x\in H_k}(f(x)-f(xs))^2&=\sum_{x\in H_k}\Big(\sum_{j=1}^{n}(f(xs_{j-1})-f(xs_j))\Big)^2\\&\leq n\sum_{x\in H_k}\sum_{j=1}^{n}\big(f(xs_{j-1})-f(xs_{j-1}t_j)\big)^2=n\sum_{y\in H_k}\sum_{j=1}^{n}(f(y)-f(yt_j))^2\\&\leq n^2\sum_{y\in H_k}\sum_{t\in S_2}(f(y)-f(yt))^2.\end{aligned}$$

This shows (3.2) with $K'(s)=n^2$, where n is the length of a word in S_2 expressing s.

Let us finally show (3.1):

$$\begin{aligned}\mathcal{R}_{S_1}(f)&=\frac{\sum_{s\in S_1}\sum_{x\in H_k}(f(x)-f(xs))^2}{2\sum_{x\in H_k}f(x)^2}\\&\leq\frac{\sum_{s\in S_1}K'(s)\sum_{x\in H_k}\sum_{t\in S_2}(f(x)-f(xt))^2}{2\sum_{x\in H_k}f(x)^2}\\&\leq\Big(\sum_{s\in S_1}K'(s)\Big)\frac{\sum_{x\in H_k}\sum_{t\in S_2}(f(x)-f(xt))^2}{2\sum_{x\in H_k}f(x)^2}=\Big(\sum_{s\in S_1}K'(s)\Big)\mathcal{R}_{S_2}(f).\end{aligned}$$

This implies that (3.1) holds with $K=(\sum_{s\in S_1}K'(s))^{-1}$.

3.2. Unbounded torsion order in congruence quotients. Let Γ be a group, and let R be a ring. Let $\mathrm{GL}_n(R)$ denote the *general linear group* of $n\times n$ matrices with entries in R and whose determinant is a unit in R. A *representation* of Γ with coefficients in R of dimension n is a group homomorphism

$$\rho\colon\Gamma\to\mathrm{GL}_n(R),$$

i.e., $\rho(g_1 \cdot g_2) = \rho(g_1) \cdot \rho(g_2)$ for all $g_1, g_2 \in \Gamma$. The representation ρ is said to be *faithful* if ρ is injective.

The ring $\mathbb{F}_p[[t]]$ is the ring of formal power series in the variable t and coefficients in the finite field $\mathbb{F}_p$ of p elements. Truncating a formal power series $f(t) = \sum_{k\geq 0} a_k t^k$ at order t^{i+1}, that is, mapping it to

$$\sum_{k=0}^{i} a_k t^k + (t^{i+1}) \in \mathbb{F}_p[t]/(t^{i+1}),$$

can be understood as a ring homomorphism

$$\pi_i \colon \mathbb{F}_p[[t]] \to \mathbb{F}_p[[t]]/(t^{i+1}) = \mathbb{F}_p[t]/(t^{i+1}).$$

A group is said to be *torsion-free* if all nontrivial elements have infinite order.

LEMMA 3.2. *Let Γ be an infinite, torsion-free group, and let*

$$\rho \colon \Gamma \to \mathrm{GL}_n(\mathbb{F}_p[[t]])$$

be a faithful representation. For any $i \geq 0$ let Γ_i be the finite group which is the image of the composition

$$\Gamma \xrightarrow{\rho} \mathrm{GL}_n(\mathbb{F}_p[[t]]) \xrightarrow{\pi_i} \mathrm{GL}_n(\mathbb{F}_p[t]/(t^{i+1}))$$

and we denote by $\psi_i = \pi_i \circ \rho : \Gamma \to \Gamma_i$ the corresponding surjective map from Γ to Γ_i.

Then for $g \in \Gamma$ a nontrivial element, the order of $\psi_i(g)$ in Γ_i tends to infinity as i increases.

Proof. We argue by contradiction. Note that the orders of $\psi_i(g)$ are monotone in i. If N is an upper bound for the orders of $\psi_i(g)$ for all i, then with $\ell = N!$ we have $\psi_i(g^\ell) = \psi_i(g)^\ell = 1$ for all i. Since $\rho(g^\ell)$ is described formally to arbitrary precision by $\psi_i(g^\ell) = 1$, we conclude that $\rho(g^\ell) = 1$. Because ρ is faithful, it follows that $g^\ell = 1$. Moreover, as Γ is torsion-free, we conclude $g = 1$, a contradiction. □

3.3. Quartic Cayley graph expanders of 3-groups. Let us begin with the least complex example underlying our construction of Cayley graph expanders, namely, a group Γ generated by a_1, a_2, a_3, a_4 subject to relations

$$a_1a_3a_1a_4, \quad a_1a_3^{-1}a_2a_3^{-1}, \quad a_1a_4^{-1}a_2^{-1}a_4^{-1}, \quad a_2a_3a_2a_4^{-1}.$$

The group Γ agrees with the group Γ_2 on page 457 of [85], where the correspondence is $a_1 = g_0$, $a_2 = g_2$, $a_3 = g_3$, $a_4 = g_1$. The group Γ can be obtained as the group $\Gamma_{3;0,1}$ described in section 3.4 (a particular case of the group $\Gamma_{M,\delta}$ of [81, section 2.8] for $q = 3$) by means of the identification $\Gamma \cong \Gamma_{3;0,1}$ via $a_1 \mapsto a_0$, $a_2 \mapsto a_2$, $a_3 \mapsto b_3^{-1}$, and $a_4 \mapsto b_1^{-1}$. It was shown by Stix and Vdovina in [85], that this group is a quaternionic lattice. The group Γ is torsion-free by [85, Theorem 30].

Consider the function field $\mathbb{F}_3(x)$ over the finite field $\mathbb{F}_3$ in one variable x. Using computer algebra, it is easy to verify that the following assignments give a well-defined representation $\Psi : \Gamma \to \mathrm{GL}_3(\mathbb{F}_3(x))$:

$$\Psi(a_1) = \begin{pmatrix} \frac{1}{x} & 1+\frac{1}{x} & -x+1-\frac{1}{x} \\ -\frac{1}{x} & -\frac{1}{x} & -x+\frac{1}{x} \\ -\frac{1}{x} & 1-\frac{1}{x} & x+1+\frac{1}{x} \end{pmatrix}, \tag{3.3}$$
$$\Psi(a_2) = \begin{pmatrix} \frac{1}{x} & -x+1 & -x^3-x^2-x \\ 0 & 1 & -x^2+x \\ 0 & 0 & x \end{pmatrix},$$
$$\Psi(a_3) = \begin{pmatrix} 0 & 0 & x^2+1 \\ 0 & -1 & -x-1 \\ \frac{1}{x^2+1} & \frac{x+1}{x^2+1} & -1+\frac{x}{x^2+1} \end{pmatrix},$$
$$\Psi(a_4) = \begin{pmatrix} \frac{2x}{x^2+1}-1 & -x-\frac{1}{x^2+1} & x^2-x+1-\frac{x}{x^2+1} \\ \frac{2x+1}{x^2+1} & \frac{2x-1}{x^2+1} & \frac{2x+1}{x^2+1}-x-1 \\ \frac{1}{x^2+1} & \frac{2x}{x^2+1} & \frac{1}{x^2+1}-1 \end{pmatrix}.$$

In fact, all matrices $\Psi(a_i)$ for $i=1,\dots,4$ have determinant 1.

The denominators of the matrix entries are nonzero at $x=1$. Therefore we can substitute $x=1+t$ and expand the matrix entries, which are now rational functions in t, as formal power series in $\mathbb{F}_3[[t]]$.

The determinant is still 1, so that the resulting matrices are invertible as matrices with values in $\mathbb{F}_3[[t]]$. We thus obtain a representation

$$\widetilde{\Psi} : \Gamma \to \mathrm{GL}_3(\mathbb{F}_3[[t]]).$$

Lemma 3.3. *The group homomorphism $\widetilde{\Psi} : \Gamma \to \mathrm{GL}_3(\mathbb{F}_3[[t]])$ given above is injective and defines a faithful representation of Γ.*

Proof. Because $\widetilde{\Psi}(a_2)$ agrees with $\Psi(a_2)$ under the necessary identifications, it has the same order, which is visibly infinite (due to the upper triangular shape and the entry x on the diagonal). Therefore $\widetilde{\Psi}$ has infinite image.

The group Γ is an arithmetic lattice in a group of rank 2 by construction in [85]. By [62], therefore, all homomorphic images of Γ are finite, or the kernel of the homomorphism is finite. As $\widetilde{\Psi}(\Gamma)$ is infinite, it follows that the kernel of $\widetilde{\Psi}$ is finite. But Γ is torsion-free; hence all finite subgroups are trivial, and a trivial kernel means that $\widetilde{\Psi}$ is faithful. ∎

As in Lemma 3.2 we consider the truncated representations up to orders t^{i+1} and higher,

$$\widetilde{\Psi}_i : \Gamma \xrightarrow{\widetilde{\Psi}} \mathrm{GL}_3(\mathbb{F}_3[[t]]) \xrightarrow{\pi_i} \mathrm{GL}_3(\mathbb{F}_3[t]/(t^{i+1})),$$

and denote the image by $\Gamma_i := \widetilde{\Psi}_i(\Gamma)$.

Remark 3.4. As an alternative to the proof of Lemma 3.3 one may observe that Γ is a lattice in a quaternion algebra D over a field K (or rather an arithmetic lattice in $D^\times/K^\times$) that is split by $\mathbb{F}_3(x)$ (as an extension of K). The homomorphism Ψ is nothing but the one induced from D acting on the 3-dimensional purely imaginary quaternions in D by conjugation. It follows that Γ acts faithfully since $D^\times/K^\times$ acts faithfully.

More importantly, it follows from this construction that the truncated homomorphisms $\Gamma \to \Gamma_i$ are finite congruence quotients of Γ.

Lemma 3.5. *The finite groups Γ_i are 3-groups: the order of Γ_i is a power of 3.*

Proof. We denote by $\mathbf{id}_3$ the 3×3 unit matrix. Forgetting the term of order t^i yields a group homomorphism $\Gamma_i \to \Gamma_{i-1}$. Its kernel is naturally a subgroup of the additive group $\mathrm{M}_3(\mathbb{F}_3)$ of 3×3 matrices with coefficients in $\mathbb{F}_3$, hence a 3-group. Indeed, the elements in $\ker(\Gamma_i \to \Gamma_{i-1})$ have the form $\mathbf{id}_3 + At^i$ for an $A \in \mathrm{M}_3(\mathbb{F}_3)$ and the identity matrix $\mathbf{id}_3$. The assignment $\mathbf{id}_3 + At^i \mapsto A$ is an injective group homomorphism $\ker(\Gamma_i \to \Gamma_{i-1}) \to \mathrm{M}_3(\mathbb{F}_3)$ because

$$(\mathbf{id}_3 + At^i)(\mathbf{id}_3 + Bt^i) = \mathbf{id}_3 + (A+B)t^i \quad \in \Gamma_i.$$

It remains to prove the claim for $i = 0$, the leading constant term, and to conclude by induction thanks to Lagrange's theorem on group orders.

For $i = 0$, we simply plug in $t = 0$ or, what is the same, $x = 1$, into (3.3) to get formulas for $\widetilde{\Psi}_0 : \Gamma \to \Gamma_0 \subseteq \mathrm{GL}_3(\mathbb{F}_3)$ as follows:

$$\widetilde{\Psi}_0(a_1) = \begin{pmatrix} 1 & -1 & -1 \\ -1 & -1 & 0 \\ -1 & 0 & 0 \end{pmatrix}, \quad \widetilde{\Psi}_0(a_2) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix},$$

$$\widetilde{\Psi}_0(a_3) = \widetilde{\Psi}_0(a_4) = \begin{pmatrix} 0 & 0 & -1 \\ 0 & -1 & 1 \\ -1 & 1 & 1 \end{pmatrix},$$

which generate a cyclic group of order 3, namely, Γ_0. ∎

Remark 3.6. In fact, the kernel $N_i = \ker(\widetilde{\Psi}_i : \Gamma \to \Gamma_i)$ is by construction described by congruences modulo t^{i+1}, and—unraveling the definition of [81, 85]—we see that N_i is a congruence subgroup of the lattice Γ. There is a corresponding infinite series of square complexes P_i, the quotient of the product of trees $T_4 \times T_4$ by N_i, with the number of vertices being a power of 3. Indeed, the 3-group $\Gamma_i = \Gamma/N_i$ acts simply transitively on the vertices of P_i. The 1-skeleton of P_i is the Cayley graph $\tilde{G}_i = \mathcal{C}(\Gamma_i, S_i)$ for S_i the image in Γ_i of the set of generators $\{a_1^{\pm 1}, a_2^{\pm 1}, a_3^{\pm 1}, a_4^{\pm 1}\}$ of Γ as considered in [81, 85].

THEOREM 3.7. *The quaternionic lattice* Γ_2 *of page* 457 *of* [85], *introduced above as* Γ, *is an infinite group with two generators* x_0, x_1 *and an infinite sequence* $\{N_i\}_{i \in \mathbb{N}}$ *of normal subgroups such that the following are true:*

1. *The indices* $[\Gamma : N_i] = n_i$ *are powers of* 3.
2. *Let* $v_0 = x_0 N_i$ *(resp.,* $v_1 = x_1 N_i$*) be the image of the generator* x_0 *(resp.,* x_1*) under the quotient map* $\Gamma \to \Gamma_i := \Gamma/N_i$. *The orders of the four subgroups* H_i^1, H_i^2, H_i^3, H_i^4 *of* Γ_i, *generated by* v_0, v_1, $v_1^{-1}v_0$, $v_1 v_0$, *respectively, converge to infinity as* i *increases.*
3. *There exists a positive constant* $c > 0$ *such that the set* $\mathcal{G} = \{G_1, G_2, \ldots\}$ *of Cayley graphs* $G_i = \mathcal{C}(\Gamma_i, \{v_0^{\pm 1}, v_1^{\pm 1}\})$ *is a family of* $(n_i, 4, c)$*-expanders.*

Proof. The generators are the elements $x_0 = a_2$ and $x_1 = a_3$ of the above. These two elements generate Γ because the other defining generators can be written (using the defining relations) as

$$a_1 = a_3 a_2^{-1} a_3, \quad a_4 = a_2 a_3 a_2.$$

The normal subgroups are the groups $N_i = \ker(\widetilde{\Psi}_i)$ of the above. The indices $[\Gamma : N_i] = \#\Gamma_i$ are powers of 3 due to Lemma 3.5. Lemma 3.2 shows the assertion on the asymptotic of the orders of the images of specific elements. In order to apply this to v_0, v_1, $v_1^{-1}v_0$, and v_1v_0 we must show that x_0, x_1, $x_1^{-1}x_0$, and x_1x_0 are all nontrivial. Consider the homomorphism $\Gamma \to \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2\mathbb{Z}$ given by counting modulo 2 the number of even and odd indexed a_i occurring in a word representing a group element of Γ. Then x_0, x_1, $x_1^{-1}x_0$, and x_1x_0 all have nontrivial image: $(1,0)$, $(0,1)$, $(1,1)$, and $(1,1)$, respectively; hence these elements are nontrivial.

Our next aim is to prove the expander property of the graphs G_i. In order to do so, we first consider the Cayley graphs $\tilde{G}_i = \mathcal{C}(\Gamma_i, S_i)$ for the larger set of generators S_i that is the image in Γ_i of the set of generators $\{a_1^{\pm 1}, a_2^{\pm 1}, a_3^{\pm 1}, a_4^{\pm 1}\}$ of Γ. As recalled in Remark 3.6, the graph $\tilde{G}_i$ is the 1-skeleton of a square complex $P_i = (T_4 \times T_4)/N_i$. As explained in [81, section 6], there are two Hecke operators on functions on the vertices $V(P_i)$ of P_i, averaging functions over neighbors that differ in the first component ($A_{v,i}$ with v: vertical neighbor) or, respectively, in the second component ($A_{h,i}$ with h: horizontal neighbor). Because the universal covering is a product of trees, these Hecke operators commmute: $A_{v,i}A_{h,i} = A_{h,i}A_{v,i}$. It follows that $A_{v,i}$ and $A_{h,i}$ are simultaneously diagonalizable. Therefore the spectrum of the adjacency operator $A_{\tilde{G}_i}$ of $\tilde{G}_i$, which is

$$A_{\tilde{G}_i} = A_{v,i} + A_{h,i},$$

can be related to the spectrum of both Hecke operators. Eigenvalues of $A_{\tilde{G}_i}$ are sums of eigenvalues, one of each Hecke operator. By [81, Theorem 6.14] the square complex P_i is Ramanujan (here it is crucial that N_i is a congruence subgroup), and thus the eigenvalues of both Hecke operators satisfy a spectral gap (even a Ramanujan gap). Using [81, Proposition 6.17] to bound the multiplicity of the extremal eigenvalues of the Hecke operators, also a spectral gap is inherited by $A_{\tilde{G}_i}$ (but now without the optimal Ramanujan property). Therefore the graphs $\tilde{G}_i$ form a family of Cayley graph expanders (of valency 8 for 3-groups).

The result of section 3.1 shows that the Cayley graphs $G_i = \mathcal{C}(\Gamma_i, \{v_0^{\pm 1}, v_1^{\pm 1}\})$ of Γ_i with respect to the generating set given by the images of $\{a_2^{\pm}, a_3^{\pm}\}$ in Γ_i are still a sequence of expander graphs (although with a different expansion constant). The valency is now 4, and the vertex set has cardinality n_i, a power of 3. ∎

3.4. Explicit Cayley graphs p-expanders with $p-2$ generators. We recall the explicit description of the lattices from [81, section 2.8]. For details we refer to loc. cit.

Let $p \geq 2$ be a prime number, and let $\mathbb{F}_{p^2}$ be the field with p^2 elements. Its multiplicative group $\mathbb{F}_{p^2}^{\times}$ is cyclic, and we fix a generator $\delta \in \mathbb{F}_{p^2}^{\times}$. We define for $k, j \in \mathbb{Z}/(p^2-1)\mathbb{Z}$ such that $k \not\equiv j \pmod{p-1}$ the elements $x_{k,j}, y_{k,j} \in \mathbb{Z}/(p^2-1)\mathbb{Z}$ uniquely by

$$\delta^{x_{k,j}} = 1 + \delta^{j-k}, \quad \delta^{y_{k,j}} = 1 + \delta^{k-j}.$$

(This is possible since $\delta^{j-k} \neq -1$.) We set further, in $\mathbb{Z}/(p^2-1)\mathbb{Z}$,

$$i(k,j) = j - y_{k,j}(p-1), \quad \ell(k,j) = k - x_{k,j}(p-1).$$

We now fix two elements $\alpha \neq \beta \in \mathbb{Z}/(p-1)\mathbb{Z}$, consider the reduction modulo $p-1$ given by

$$\mathrm{pr} : \mathbb{Z}/(p^2-1)\mathbb{Z} \to \mathbb{Z}/(p-1)\mathbb{Z},$$

and define $(p+1)$-element sets $K = \mathrm{pr}^{-1}(\alpha)$ and $J = \mathrm{pr}^{-1}(\beta)$. Since $p-1$ divides $\mu = (p^2-1)/2$, the sets K and J are preserved under translation by μ.

The group $\Gamma_{p;\alpha,\beta}$ (the dependence on δ is implicit) is defined by generators a_k for $k \in K$ and b_j for $j \in J$ subject to the relations: for all $k \in K$ and $j \in J$ we have $a_k a_{k+\mu} = 1$ and $b_j b_{j+\mu} = 1$ and

$$a_k b_j a_{\ell(k,j)}^{-1} b_{i(k,j)}^{-1} = 1. \tag{3.4}$$

It was proven in [81] that $\Gamma_{p;\alpha,\beta}$ is a quaternionic arithmetic lattice of rank 2 and residually pro-p by congruence p-group quotients. Moreover, congruence quotients yield Cayley graphs with respect to the given generators $A = \{a_k; k \in K\}$ together with $B = \{b_j; j \in J\}$ that form a sequence of expanders (as the 1-skeleton of 2-dimensional Ramanujan expander complexes). The valency of these graphs is $\#(K \cup J) = 2(p+1)$.

For the applications in this note we must reduce the number of essential generators (not counting inverses).

Proposition 3.8. *The group $\Gamma_{p;\alpha,\beta}$ can be generated by $(p+3)/2$ elements. More precisely, half of the elements of B (omitting inverses) together with one $a \in A$ generate $\Gamma_{p;\alpha,\beta}$.*

Proof. The relation (3.4) comes from the squares in the one vertex square complex with complete bipartite link whose fundamental group is $\Gamma_{p;\alpha,\beta}$. It follows that for all $j \in J$ the map

$$\sigma_j \colon K \to K, \quad \sigma_j(k) = \ell(k,j)$$

is bijective. The argument of [85, Proposition 35] works also for the groups $\Gamma_{p;\alpha,\beta}$ and shows that the group P_B generated by all σ_j for $j \in J$ in the symmetric group on the set K acts transitively on K. Since the relation (3.4) shows that

$$a_{\sigma_j(k)} = b_{i(k,j)}^{-1} a_k b_j,$$

a subgroup containing all $b \in B$ will contain with any a_k also the $a_{k'}$ for $k' \in K$ in the P_B-orbit of k. By transitivity of P_B on K this automatically involves all of A. This proves the theorem. ∎

Example 3.9. The following is an explicit example for the above construction for $p = 5$ and $\alpha = 0, \beta = 2$ in $\mathbb{Z}/4\mathbb{Z}$. The group $\Gamma_{5;0,2}$ is generated by elements (indices are to be considered modulo 24)

$$a_0, a_4, a_8, b_2, b_6, b_{10}$$

and 9 relations of length 4,

$$a_0 b_2 a_0 b_{10},\ a_0 b_6 a_0^{-1} b_6^{-1},\ a_0 b_2^{-1} a_4^{-1} b_2^{-1},\ a_0 b_{10}^{-1} a_8 b_{10}^{-1},\ a_4 b_6 a_4 b_2^{-1},$$

$$a_4 b_{10} a_4^{-1} b_{10}^{-1},\ a_4 b_6^{-1} a_8^{-1} b_6^{-1},\ a_8 b_2 a_8^{-1} b_2^{-1},\ a_8 b_{10} a_8 b_6^{-1}.$$

A direct inspection shows that $\Gamma_{5;0,2}$ can be generated by 2 elements, for example, a_0, b_2.

Theorem 3.10. *Let $p \geq 3$ be a prime number, and let $d \geq 2$ be an integer. We assume that $d \geq (p+3)/2$ if $p \geq 7$.*

There are infinitely many finite p-groups Γ_i with order tending to infinity and generating sets T_i of cardinality d such that the Cayley graphs $\mathcal{C}(\Gamma_i, T_i)$ form a family of expanders with number of vertices a power of p and valency $2d$.

In particular, for $p \geq 5$ we may take $d = p - 2$.

Proof. We consider the groups constructed in [81, section 2.8] as recalled in section 3.4 with notation $\Gamma_{p;\alpha,\beta}$. By Proposition 3.8, these groups can be generated by $(p+3)/2$ elements. Moreover, for $p = 5$ we take the group $\Gamma_{5;0,2}$ considered in Example 3.9, which can be generated by 2 elements. For $p = 3$, we take the 2-generated group $\Gamma_{3,0,1}$ of Theorem 3.7. For any d as in the statement of the theorem, we may therefore choose a set T of d generators of the respective infinity group $\Gamma_{p;\alpha,\beta}$ (adding arbitrary elements if necessary for values of d larger than the given minimal values).

If $p \geq 7$, then $p - 2 \geq (p+3)/2$, and $d = p - 2$ is a possible choice. For $p = 5$, we have $p - 2 \geq 2$, so $d = p - 2$ is a valid choice for all $p \geq 5$.

It follows from [81, Proposition 2.22] that these groups $\Gamma_{p;\alpha,\beta}$ are residually pro-p with respect to a suitable infinite sequence of congruence subgroup quotients of p-power order.

The standard generating sets $A \cup B$ of $\Gamma_{p;\alpha,\beta}$ yield for the sequence of congruence quotients of $\Gamma_{p;\alpha,\beta}$ that the corresponding Cayley graphs form a series of expanders (as the 1-skeleton of 2-dimensional Ramanujan expander complexes); see [81, section 6]. Details of the transition to the 1-skeleton are as in the proof of Theorem 3.7.

Indeed, by the results of section 3.1, all these Cayley graphs of the finite congruence p-group quotients of $\Gamma_{p;\alpha,\beta}$ with respect to the alternative set T of d generators are still expanders (but not necessarily Ramanujan). This proves the theorem. ∎

Let us conclude this section by emphasizing that, in combination, Theorems 3.7 and 3.10 yield Main Theorem 3.

4. The color-preserving homomorphism basis. It turns out that the analysis of the complexity of $\#\text{EdgeSub}(\Phi)$ is much easier if a *colorful* version of the problem is considered. For our hardness results, we then show that the colorful version reduces to the uncolored version. To this end, recall that an H-coloring of a graph G is a homomorphism from G to H and that a graph G is *H-colored* if G is equipped with an H-coloring c. Recall further the implicitly defined $\#E(H)$-edge-coloring of G. In the colorful version of $\#\text{EdgeSub}(\Phi)$, denoted by $\#\text{ColSubSub}(\Phi)$, the task is to compute the cardinality of the set

$$\mathsf{ColEdgeSub}(\Phi \to_H G) := \{A \subseteq E(G) | \#A = \#E(H) \wedge c(A) = E(H) \wedge \Phi(G[A]) = 1\},$$

and the parameter is $k := \#E(H)$. In particular, we write $\#\mathsf{ColEdgeSub}(\Phi \to_H \star)$ for the function that maps an H-colored graph G to $\#\mathsf{ColEdgeSub}(\Phi \to_H G)$. Note that $\#A = \#E(H) \wedge c(A) = E(H)$ implies that the A contains each of the $\#E(H)$ colors precisely once. Further, note that $\Phi(G[A]) = 1$ if and only if Φ holds on the (uncolored) graph $G[A]$.

Each element $A \in \mathsf{ColEdgeSub}(\Phi \to_H G)$ induces a fracture ρ of H, where for $v \in V(H)$ two edges $e, f \in E_H(v)$ are in the same block of ρ_v if and only if their (unique) preimages $\widehat{e}, \widehat{f} \in A$ under $c : A \to E(H)$ are connected to the *same* endpoint in $c^{-1}(v) \subseteq V(G)$. From the construction, it immediately follows that $G[A]$ and $H\sharp\rho$ are canonically isomorphic as H-colored graphs, that is, $G[A] \cong_H H\sharp\rho$.

Our goal is to express $\#\mathsf{ColEdgeSub}(\Phi \to_H \star)$ as a linear combination of (color-preserving) homomorphism counts from graphs only depending on Φ and H. In the case that H is a torus, we establish an explicit criterion sufficient for the term

$\#\mathsf{Hom}_{\mathsf{cp}}(\circledcirc_\ell \to_{\circledcirc_\ell} \star)$ to survive with a nonzero coefficient in this linear combination. The existence of the linear combination is given by the following lemma.

LEMMA 4.1. *Let H denote a graph. We have*

$$\#\mathsf{ColEdgeSub}(\Phi \to_H \star) = \sum_{\sigma \in \mathcal{L}(\Phi,H)} \sum_{\rho \geq \sigma} \mu(\sigma,\rho) \cdot \#\mathsf{Hom}_{\mathsf{cp}}(H\sharp\rho \to_H \star)\,,$$

where the relation $\leq$ and the Möbius function μ are over the lattice of fractures $\mathcal{L}(H)$.

Proof. Let G denote an H-colored graph. We first partition the elements A of the set $\mathsf{ColEdgeSub}(\Phi \to_H G)$ according to their induced fractures. Now, writing $\mathsf{ColEdgeSub}(\Phi \to_H G)[\sigma]$ for the set of A inducing the fracture $\sigma \in \mathcal{L}(H)$, we obtain

$$\#\mathsf{ColEdgeSub}(\Phi \to_H G) = \sum_{\sigma \in \mathcal{L}(\Phi,H)} \#\mathsf{ColEdgeSub}(\Phi \to_H G)[\sigma]\,,$$

since $\#\mathsf{ColEdgeSub}(\Phi \to_H G)[\sigma] = 0$ for all $\sigma \notin \mathcal{L}(\Phi,H)$. From the fact that $G[A]$ is canonically isomorphic to $H\sharp\sigma$ as an H-colored graph, for σ associated to $A \in \mathsf{ColEdgeSub}(\Phi \to_H G)$, it follows that

$$\#\mathsf{ColEdgeSub}(\Phi \to_H G)[\sigma] = \#\mathsf{Emb}_{\mathsf{cp}}(H\sharp\sigma \to_H G)\,.$$

Note that we are using that graphs of the form $H\sharp\sigma$ can have no nontrivial automorphisms as H-colored graphs (since all edges must be fixed). It remains to show that

$$\#\mathsf{Emb}_{\mathsf{cp}}(H\sharp\sigma \to_H G) = \sum_{\rho \geq \sigma} \mu(\sigma,\rho) \cdot \#\mathsf{Hom}_{\mathsf{cp}}(H\sharp\rho \to_H G). \tag{4.1}$$

To this end, we establish the following zeta transformation, which should be considered as a color-preserving version of the standard transformation in the uncolored case (see, e.g., [60, section 5.2.3]).

CLAIM 4.2. *For every fracture σ of H, we have*

$$\#\mathsf{Hom}_{\mathsf{cp}}(H\sharp\sigma \to_H G) = \sum_{\rho \geq \sigma} \#\mathsf{Emb}_{\mathsf{cp}}(H\sharp\rho \to_H G).$$

Proof. Every color-preserving homomorphism φ from $H\sharp\sigma$ to G induces a fracture $\rho \geq \sigma$, that is, ρ_v is a coarsening of σ_v for every $v \in V(H)$. Indeed, recall that the vertices of $H\sharp\sigma$ over $v \in V(H)$ correspond to the blocks B of the partition σ_v of the edges $E_H(v)$. Then the partition ρ_v of $E_H(v)$ is obtained from σ_v by joining those blocks B, B' whose associated vertices in $H\sharp\sigma$ map to the same vertex of G under φ. We have that the subgraph of G given by the image of $H\sharp\sigma$ under φ is canonically isomorphic to $H\sharp\rho$ as an H-colored graph.

Let us call two homomorphisms in $\mathsf{Hom}_{\mathsf{cp}}(H\sharp\sigma \to_H G)$ equivalent if they induce the same fracture and write $\mathsf{Hom}_{\mathsf{cp}}(H\sharp\sigma \to_H G)[\rho]$ for the equivalence class of all homomorphisms inducing ρ. The claim then follows by partitioning the set $\mathsf{Hom}_{\mathsf{cp}}(H\sharp\sigma \to_H G)$ into those equivalence classes and observing that

$$\#\mathsf{Hom}_{\mathsf{cp}}(H\sharp\sigma \to_H G)[\rho] = \#\mathsf{Emb}_{\mathsf{cp}}(H\sharp\rho \to_H G)\,. \qquad \square$$

Equation (4.1) is now obtained by using Möbius inversion (Theorem 2.7) on the zeta transformation given by the previous claim. This concludes the proof. $\square$

Let us now collect for the coefficient of the term $\#\mathsf{Hom}_{\mathsf{cp}}(H\sharp\top \to_H G)$, where $\top$ is the maximum fracture of H with respect to the ordering $\leq$. In particular, each partition of $\top$ only consists of a single block, and thus $H\sharp\top \cong H$, where the isomorphism is given by the H-coloring of $H\sharp\top$.

COROLLARY 4.3. *Let Φ denote a computable graph property, and let H denote a graph. There is a unique computable function $a_{\Phi,H} : \mathcal{L}(H) \to \mathbb{Z}$ such that*

$$\#\mathsf{ColEdgeSub}(\Phi \to_H \star) = \sum_{\rho \in \mathcal{L}(H)} a_{\Phi,H}(\rho) \cdot \#\mathsf{Hom}_{\mathsf{cp}}(H\sharp\rho \to_H \star)\,.$$

For $\rho = \top$ we have

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

Here, $|\sigma_v|$ denotes the number of blocks of partition σ_v.

Proof. The first claim follows immediately from Lemma 4.1 by collecting coefficients; note that Φ and μ are computable and that the image of μ is a subset of $\mathbb{Z}$. For the second claim, we collect the coefficients of $\#\mathsf{Hom}_{\mathsf{cp}}(H\sharp\top \to_H *)$ in Lemma 4.1 and obtain

$$a_{\Phi,H}(\top) = \sum_{\sigma \in \mathcal{L}(\Phi,H)} \mu(\sigma, \top)\,.$$

Recall that μ is the Möbius function of $\mathcal{L}(H)$ and that the latter is the product of the partition lattices of $N_H(v)$ for each $v \in V(H)$. Using that the Möbius function is multiplicative with respect to the product (see, for instance, [84, Proposition 3.8.2]), and applying the explicit formula for the partition lattice (Theorem 2.8), we obtain the second claim. ∎

In the remainder of the paper, given Φ and H, we refer to the function $a_{\Phi,H}$ from Corollary 4.3 as the *coefficient function* of Φ and H.[16]

4.1. A colored version of complexity monotonicity. Our next goal is to prove that computing a finite linear combination of color-preserving homomorphism counts, as given by Corollary 4.3, is precisely as hard as computing its hardest term. While the proof strategy follows the approach used in [25] and [32], we need to adapt to the setting of color-preserving homomorphisms between fractured graphs.

We rely on the tensor product of H-colored graphs in the following way: let H denote a fixed graph, and let G and F denote H-colored graphs with colorings c_G and c_F. The *color-preserving tensor product* $G \times_H F$ has vertices $\{(x,y) \in V(G) \times V(F) | c_G(x) = c_F(y)\}$, and two vertices (x,y) and (x',y') are made adjacent in $G \times_H F$ if (and only if) $\{x,x'\} \in E(G)$ and $\{y,y'\} \in E(F)$. Observe that the graph $G \times_H F$ is H-colored as well by the coloring $(x,y) \mapsto c_G(x) = c_F(y)$.

LEMMA 4.4. *Let H denote a graph, and let F, G_1 and G_2 denote H-colored graphs. We have*

$$\#\mathsf{Hom}_{\mathsf{cp}}(F \to_H G_1 \times_H G_2) = \#\mathsf{Hom}_{\mathsf{cp}}(F \to_H G_1) \cdot \#\mathsf{Hom}_{\mathsf{cp}}(F \to_H G_2)\,.$$

[16]We note that the value $a_{\Phi,H}(\top)$ of the coefficient function on the largest partition $\top$ is equal to the indicator $a(\Phi, H)$ used in the proof outline in the introduction.

Proof. The function

$$b : \mathsf{Hom}_{\mathsf{cp}}(F \to_H G_1) \times \mathsf{Hom}_{\mathsf{cp}}(F \to_H G_2) \to \mathsf{Hom}_{\mathsf{cp}}(F \to_H G_1 \times_H G_2),$$
$$b(\varphi, \psi)(u) := (\varphi(u), \psi(u)) \text{ for } u \in V(F)$$

is the canonical bijection. ∎

The reduction for isolating terms with nonzero coefficient requires us to solve a system of linear equations. For the definition of the corresponding matrix, we fix a linear extension $\preccurlyeq$ of the order $\leq$ of the H-fractures. Recall that $\leq$ is also the order of the product of the partition lattices of the set $E(v)$ for all $v \in V(H)$. In particular, $\sigma \leq \rho$ if and only if σ_v refines ρ_v for all $v \in V(H)$. As a consequence, we observe that $\rho \succ \sigma$, that is, $\neg(\rho \preccurlyeq \sigma)$, implies the existence of a vertex $v \in V(H)$ such that ρ_v does not refine σ_v. Now let $\mathcal{M}_H$ denote the matrix whose columns and rows are associated with the set of all H-fractures, ordered by $\preccurlyeq$, and whose entries are given by

$$\mathcal{M}_H[\rho, \sigma] := \#\mathsf{Hom}_{\mathsf{cp}}(H \sharp \rho \to H \sharp \sigma)\,.$$

LEMMA 4.5. *For each H, the matrix $\mathcal{M}_H$ is upper triangular with entries* 1 *on the diagonal.*

Proof. Let us first consider the diagonal. We claim that $\#\mathsf{Hom}_{\mathsf{cp}}(H\sharp\rho \to H\sharp\rho) = 1$. Due to the trivial (identity) homomorphism we have $\#\mathsf{Hom}_{\mathsf{cp}}(H\sharp\rho \to H\sharp\rho) \geq 1$. On the other hand, the canonical coloring $H\sharp\rho \to H$ induces a bijection from the edges of $H\sharp\rho$ to the edges of H that preserves the coloring. Since a color-preserving homomorphism $H\sharp\rho \to H\sharp\rho$ must commute with this map, it must act as the identity on all edges of $H\sharp\rho$ and is thus equal to the identity.

It remains to prove that $\mathcal{M}_H[\rho, \sigma] = 0$ for every $\rho \succ \sigma$. Recall that the latter implies the existence of a vertex $v \in V(H)$ such that ρ_v does not refine σ_v; that is, there is a block B of ρ_v which is not a subset of any block of σ_v. Thus, there are edges $e, f \in B \subseteq E_H(v)$ such that e, f are in different blocks of σ_v. Identifying $E(H\sharp\sigma) = E(H\sharp\rho) = E(H)$ using the coloring, we see that e, f are adjacent to the *same* vertex in $H\sharp\rho$ (corresponding to the block B) but to *different* vertices in $H\sharp\sigma$. This implies that there cannot be a color-preserving homomorphism $\varphi : H\sharp\rho \to H\sharp\sigma$ since e, f being incident at v^B in $H\sharp\rho$ would imply that $e = \varphi(e), f = \varphi(f)$ must be incident at $\varphi(v^B)$ in $H\sharp\sigma$. ∎

We are now able to prove a version of the complexity monotonicity principle which is sufficient for the purposes in this work. In what follows, given a graph property Φ, we write $\#\mathsf{ColEdgeSub}(\Phi \to_\star \star)$ for the function that expects as input a graph H and an H-colored graph G, and outputs $\#\mathsf{ColEdgeSub}(\Phi \to_H G)$.

LEMMA 4.6. *Let Φ denote a computable graph property. Then there exists a deterministic algorithm $\mathbb{A}$ which has oracle access to the function $\#\mathsf{ColEdgeSub}(\Phi \to_\star \star)$ and computes, given as input a graph H and an H-colored graph G, the numbers $\#\mathsf{Hom}_{\mathsf{cp}}(H\sharp\rho \to_H G)$ for every H-fracture ρ satisfying that $a_{\Phi,H}(\rho) \neq 0$, where $a_{\Phi,H}$ is the coefficient function of Φ and H.*

Furthermore, there is a computable function f such that $\mathbb{A}$ runs in time $f(|H|) \cdot |G|^{O(1)}$ and every posed oracle query $(\hat{H}, \hat{G})$ satisfies $\hat{H} = H$ and $|\hat{G}| \leq f(|H|) \cdot |G|$.

Proof. Given H and G, we can obtain $\#\mathsf{ColEdgeSub}(\Phi \to_H G \times_H (H\sharp\sigma))$ for all H-fractures σ via access to the oracle. By Corollary 4.3 and Lemma 4.4, we have

$$\begin{aligned}
&\#\mathsf{ColEdgeSub}(\Phi \to_H G \times_H (H\sharp\sigma)) \\
&\quad = \sum_{\rho\in\mathcal{L}(H)} a_{\Phi,H}(\rho)\cdot \#\mathsf{Hom}_{\mathsf{cp}}(H\sharp\rho \to_H G\times_H (H\sharp\sigma)) \\
&\quad = \sum_{\rho\in\mathcal{L}(H)} a_{\Phi,H}(\rho)\cdot \#\mathsf{Hom}_{\mathsf{cp}}(H\sharp\rho \to_H G)\cdot \#\mathsf{Hom}_{\mathsf{cp}}(H\sharp\rho \to_H H\sharp\sigma).
\end{aligned}$$

Observe that the latter yields a system of linear equations for the numbers

$$a_{\Phi,H}(\rho)\cdot \#\mathsf{Hom}_{\mathsf{cp}}(H\sharp\rho \to_H G)$$

with matrix $\mathcal{M}_H$ which has full rank according to Lemma 4.5. Consequently $\mathbb{A}$ can compute the number $a_{\Phi,H}(\rho)\cdot\#\mathsf{Hom}_{\mathsf{cp}}(H\sharp\rho \to_H G)$ for each H-fracture ρ. Therefore, $\#\mathsf{Hom}_{\mathsf{cp}}(H\sharp\rho \to_H G)$ can be computed whenever $a_{\Phi,H}(\rho)\neq 0$.

Now observe that $a_{\Phi,H}$, which is computable, only depends on Φ, which is fixed, and H. Furthermore $\mathcal{L}(H)$ and all $H\sharp\rho$ only depend on H. Thus the computation of $\#\mathsf{Hom}_{\mathsf{cp}}(H\sharp\rho \to_H H\sharp\sigma)$ takes time only depending on H as well. Consequently, the system of linear equations can be solved in time $f'(|H|)\cdot|G|^{O(1)}$ for some computable function f'. Furthermore, the size of $G\times_H(H\sharp\sigma)$ is bounded by $|H\sharp\sigma|\cdot|G|$. Setting $f(|H|) := \max\{f'(|H|), \max_{\sigma\in\mathcal{L}(H)}|H\sharp\sigma|\}$ concludes the proof since each fractured graph $H\sharp\sigma$ has only $\#E(H)$ many edges. ∎

4.2. Intractability of counting homomorphisms from tori and expanders. The final step of this section is to establish $\#\mathsf{W}[1]$-hardness of the (uncolored) problem $\#\textsc{EdgeSub}(\Phi)$ whenever $a_{\Phi,\circledcirc_\ell}(\top)\neq 0$ for infinitely many ℓ. Essentially, we rely on the fact that tori have high treewidth and that the problem of counting (color-preserving) homomorphisms from high-treewidth graphs is hard. We can proceed similarly in the case of expanders, and due to the fact that expanders have high treewidth and are sparse (see Fact 2.3), we even obtain an almost tight conditional lower bound.

In both cases, we use complexity monotonicity, which yields hardness of the (edge-)colorful version of $\#\textsc{EdgeSub}(\Phi)$. Consequently, we need to show that the colorful version reduces to the uncolored version. This can be achieved by a standard inclusion-exclusion argument.

Lemma 4.7. *Let Φ denote a computable graph property. There exists a deterministic algorithm $\mathbb{A}$, equipped with oracle access to the function*

$$(k,\hat{G}) \mapsto \#\mathsf{EdgeSub}(\Phi, k\to\hat{G}),$$

which expects as input a graph H and an H-colored graph G and computes in time $2^{|E(H)|}\cdot|G|^{O(1)}$ the cardinality $\#\mathsf{ColEdgeSub}(\Phi\to_H G)$. Furthermore, every oracle query $(k,\hat{G})$ posed by $\mathbb{A}$ satisfies $k=|E(H)|$ and $|\hat{G}|\leq|G|$.

Proof. Given H and an H-colored graph G, we write $c: E(G)\to E(H)$ for the induced edge-coloring of G. Given a set of edge-colors $J\subseteq E(H)$, we write $G\setminus J$ for the graph obtained from G by deleting all edges e with $c(e)\in J$. Now recall that

$$\mathsf{EdgeSub}(\Phi,|E(H)|\to G) = \{A\subseteq E(G) \mid \#A = |E(H)| \wedge \Phi(G[A])=1\} \text{ and}$$

$$\mathsf{ColEdgeSub}(\Phi\to_H G) = \{A\subseteq E(G) \mid \#A = |E(H)| \wedge c(A)=E(H) \wedge \Phi(G[A])=1\}.$$

Next set $k := |E(H)|$; then we have

$$
\begin{aligned}
&\#\mathsf{ColEdgeSub}(\Phi \to_H G)\\
&\quad = \#\mathsf{EdgeSub}(\Phi, k \to G)\\
&\qquad - \#\left(\bigcup_{e\in E(H)} \{A \in \mathsf{EdgeSub}(\Phi, k \to G) | e \notin c(A)\}\right)\\
&\quad = \#\mathsf{EdgeSub}(\Phi, k \to G)\\
&\qquad - \sum_{\emptyset \neq J \subseteq E(H)} (-1)^{|J|+1} \cdot \#\left(\bigcap_{e\in J} \{A \in \mathsf{EdgeSub}(\Phi, k \to G) | e \notin c(A)\}\right)\\
&\quad = \#\mathsf{EdgeSub}(\Phi, k \to G)\\
&\qquad - \sum_{\emptyset \neq J \subseteq E(H)} (-1)^{|J|+1} \cdot \#\{A \in \mathsf{EdgeSub}(\Phi, k \to G) | \forall e \in J : e \notin c(A)\}\\
&\quad = \#\mathsf{EdgeSub}(\Phi, k \to G)\\
&\qquad - \sum_{\emptyset \neq J \subseteq E(H)} (-1)^{|J|+1} \cdot \#\mathsf{EdgeSub}(\Phi, k \to G \setminus J)\\
&\quad = \sum_{J \subseteq E(H)} (-1)^{|J|} \cdot \#\mathsf{EdgeSub}(\Phi, k \to G \setminus J).
\end{aligned}
$$

Note that the second equation is due to the inclusion-exclusion principle. We conclude that the desired number $\#\mathsf{ColEdgeSub}(\Phi \to_H G)$ can be computed using $2^{|E(H)|}$ many oracle calls of the form $\#\mathsf{EdgeSub}(\Phi, |E(H)| \to G \setminus J)$. The claim of the lemma follows since $|G \setminus J| \leq |G|$. ∎

For the formal statement of this section's main lemma, we define $\mathcal{H}[\Phi, ⊚]$ as the set of all $⊚_\ell$ such that $a_{\Phi,⊚_\ell}(\top) \neq 0$. Furthermore, given a family $\mathcal{G} = \{G_1, G_2, \ldots\}$ of (n_i, d, c)-expanders, we write $\mathcal{H}[\Phi, \mathcal{G}]$ for the set of all G_i such that $a_{\Phi,G_i}(\top) \neq 0$.

LEMMA 4.8. *Let Φ denote a computable graph property, fix a rational c and an integer d, and let $\mathcal{G} = \{G_1, G_2, \ldots\}$ denote a family of (n_i, d, c)-expanders. If at least one of $\mathcal{H}[\Phi, ⊚]$ and $\mathcal{H}[\Phi, \mathcal{G}]$ is infinite, then* #COLSUBSUB(Φ) *and* #EDGESUB(Φ) *are* #W[1]*-hard. Moreover, if $\mathcal{H}[\Phi, \mathcal{G}]$ is infinite, then, assuming ETH, both problems cannot be solved in time*

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

for any function f.

Proof. We first consider the colored version #COLSUBSUB(Φ).

We start with the case of $\mathcal{H}[\Phi, ⊚]$ being infinite. If the latter is true, then $\mathcal{H}[\Phi, ⊚]$ has unbounded treewidth since it contains graphs with arbitrary large grid minors [74]. This allows us to reduce from the problem #HOM($\mathcal{H}[\Phi, ⊚]$) which is known to be #W[1]-hard since $\mathcal{H}[\Phi, ⊚]$ has unbounded treewidth [28]. It is convenient to consider the following intermediate problem: given a graph $H \in \mathcal{H}[\Phi, ⊚]$ and an H-colored graph G with coloring c, in the problem $\#CP - HOM(\mathcal{H}[\Phi, ⊚])$ the task is to compute the number $\#\mathsf{cp\text{-}Hom}(H \to G)$ of homomorphisms $\varphi \in \mathsf{Hom}(H \to G)$ such that $c(\varphi(v)) = v$ for each vertex v of H. It is well known that #HOM($\mathcal{H}$) reduces to $\#CP - HOM(\mathcal{H})$ for every class of graphs $\mathcal{H}$; see, for instance, [24, 31, 32, 77]—note that, in the latter, the problem is referred to as #PartitionedSub($\mathcal{H}$). Thus we have

$$\#\text{Hom}(\mathcal{H}[\Phi,\circledcirc]) \le_{\mathrm{T}}^{\mathrm{fpt}} \#CP-HOM(\mathcal{H}[\Phi,\circledcirc])\,. \tag{4.2}$$

Now observe that $\#\textsf{cp-Hom}(H \to G) = \#\textsf{Hom}_{\textsf{cp}}(H\sharp\top \to_H G)$ for every graph H and H-colored graph G, since $H\sharp\top = H$. By definition of $\mathcal{H}[\Phi,\circledcirc]$, we have that $a_{\Phi,\circledcirc_\ell}(\top) \neq 0$ whenever $\circledcirc_\ell \in \mathcal{H}[\Phi,\circledcirc]$. Thus we can use complexity monotonicity (Lemma 4.6) which yields the reduction

$$\#CP-HOM(\mathcal{H}[\Phi,\circledcirc]) \le_{\mathrm{T}}^{\mathrm{fpt}} \#\text{ColSubSub}(\Phi)\,. \tag{4.3}$$

Consequently, #ColSubSub(Φ) is #W[1]-hard by (4.2)–(4.4) in combination with #W[1]-hardness of #Hom($\mathcal{H}[\Phi,\circledcirc]$).

In the case of $\mathcal{H}[\Phi,\mathcal{G}]$, we reduce from the homomorphism counting problem as well and obtain #W[1]-hardness analogously. However, for the almost tight conditional lower bound, we rely on a result of Marx [63] implying that for any class $\mathcal{H}$ of unbounded treewidth, the problem #Hom($\mathcal{H}$) cannot be solved in time

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

for any function f, unless ETH fails.[17] Let us use the aforementioned lower bound for the case of $\mathcal{H}=\mathcal{H}[\Phi,\mathcal{G}]$. Observe that the reduction sequence from #Hom($\mathcal{H}[\Phi,\mathcal{G}]$) to #ColSubSub($\Phi$) as illustrated before only leads to a linear blow up of the parameter: given an input (G_i,G) for which we wish to compute $\#\textsf{Hom}(G_i \to G)$, we only query the oracle for #ColSubSub(Φ) on instances (k,G') where $k=\#E(G_i)$ and $|G'| \le f''(\#E(G_i))\cdot|G|$ for some function f''. Since both $\#E(G_i)$ and the treewidth of G_i are linear in $|V(G_i)|$ (see Fact 2.3), any algorithm that, for some function f', solves #ColSubSub(Φ) in time

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

yields an algorithm for #Hom($\mathcal{H}[\Phi,\mathcal{G}]$), running in time

$$f(|G_i|)\cdot|G|^{o(\mathsf{tw}(G_i)/\log\mathsf{tw}(G_i))}\,,$$

for some function f (depending only on f' and f''), contradicting ETH by Marx's lower bound.

Finally, we reduce the colored version to the uncolored version by Lemma 4.7 and obtain

$$\#\text{ColSubSub}(\Phi) \le_{\mathrm{T}}^{\mathrm{fpt}} \#\text{EdgeSub}(\Phi)\,. \tag{4.4}$$

Again, this reduction is tight in the sense that the parameter does not increase by the condition on the oracle queries in Lemma 4.7. Consequently, #W[1]-hardness and the conditional lower bound hold for #EdgeSub(Φ) as well. ∎

Regarding the previous proof, observe that we cannot obtain a similar conditional lower bound if only $\mathcal{H}[\Phi,\circledcirc]$ is infinite, since in that case the parameter grows quadratically: while $\circledcirc_\ell$ has treewidth $O(\ell)$, it has $2\ell^2$ edges.

[17] Observe that this result follows only implicitly from [63], but we made it explicit in [79].

In the next section, we will focus on analyzing the classes $\mathcal{H}[\Phi, \circledcirc]$ and $\mathcal{H}[\Phi, \mathcal{G}]$, the latter for a variety of families of p-group Cayley graph expanders constructed in section 3. More precisely, we will establish conditions on properties Φ under which those classes are infinite. The parameterized and fine-grained lower bounds on exact counting will then be derived in section 6 as an immediate consequence of Lemma 4.8.

5. Coefficients of tori and Cayley graph expanders. The previous section allows us to establish hardness of $\#\textsc{EdgeSub}(\Phi)$ by the purely combinatorial problem of determining whether one of the sets $\mathcal{H}[\Phi, \circledcirc]$ and $\mathcal{H}[\Phi, \mathcal{G}]$, for some family of expanders $\mathcal{G}$, is infinite. Still, this is a challenging combinatorial problem, and we consider the treatment of the coefficients of the tori and Cayley graph expanders to be our main technical contribution in this work.

Recall from Corollary 4.3 that the coefficient function of Φ and H satisfies

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

In the case that H satisfies certain symmetry properties, we prove that it suffices to consider only those fractures in the previous sum that are fixed points under suitable group actions. More formally, we obtain the desired symmetries from the structure of the groups underlying the Cayley graph constructions for tori and expanders as introduced in the subsequent subsections.

We start with general conditions on properties under which $\mathcal{H}[\Phi, \circledcirc]$ and $\mathcal{H}[\Phi, \mathcal{G}]$, the latter for a family of 4-regular 2-group Cayley graph expanders, are infinite.

5.1. Tori and 2-group Cayley graph expanders.

5.1.1. Symmetries of the torus. Even though most (but not all) of our hardness results rely on Cayley graph expanders, we start our presentation of the analysis of the coefficient function with the toroidal grid (or just the "torus"). Since the structure of the torus is much simpler than the Cayley graph expanders, the group action on the fractures of the torus will be much easier to visualize. In particular, we hope that the analysis of the torus provides the intuition for actions on our Cayley graph expanders needed in the subsequent sections which generalize the case of the torus in the natural way.

Formally, the torus is a simple Cayley graph given by the direct product of two cyclic groups.

Definition 5.1 (the torus). *Let $\ell \geq 3$ denote an integer. The torus, also called the toroidal grid, $\circledcirc_\ell$ of size ℓ is the Cayley graph of $\mathbb{Z}_\ell^2$ with generators* $(\pm 1, 0), (0, \pm 1)$, *that is,*

$$\circledcirc_\ell := \mathcal{C}\left(\mathbb{Z}_\ell^2, \{(1,0),(-1,0),(0,1),(0,-1)\}\right)\,.$$

Equivalently, the vertices of $\circledcirc_\ell$ are $\mathbb{Z}_\ell^2$ and two vertices (x,y) and (x',y') are adjacent if and only if

$$x = x' \textit{ and } y' = y \pm 1 \mod \ell, or y = y' and\ x' = x \pm 1 \mod \ell.$$

Consult Figure 3 *for a visualization.*

In the following, for simplicity, we write $+$ for (pointwise) addition modulo ℓ. Our goal is to understand the symmetries of $\circledcirc_\ell$. Consider the following action of $\mathbb{Z}_\ell^2$ on

$$\begin{tikzcd} M_H \arrow[r] \arrow[d, "\widehat{\varphi}"] & H\sharp\rho \arrow[r] \arrow[d] & H \arrow[d, "\varphi"] \\ M_H \arrow[r] & H\sharp\varphi(\rho) \arrow[r] & H \end{tikzcd}$$

FIG. 5. *Each automorphism φ of H lifts to an automorphism $\hat{\varphi}$ of M_H. The latter descends to an isomorphism from $H\sharp\rho$ to $H\sharp\varphi(\rho)$ for every fracture ρ of H.*

the vertices of $\circledcirc_\ell$. Let $(i,j) \in \mathbb{Z}_\ell^2$, and let $(x,y) \in V(\circledcirc_\ell)$. We set $(i,j) \vdash (x,y) := (i,j) + (x,y)$. The following is immediate.

FACT 5.2. *The action of $\mathbb{Z}_\ell^2$ on $V(\circledcirc_\ell)$ is transitive. In particular, for every $(i,j) \in \mathbb{Z}_\ell^2$, the function $(i,j) \vdash \star$ is an automorphism of $\circledcirc_\ell$.*

The fact above allows us to consider the set $\mathbb{Z}_\ell^2$ of all (i,j)-"shifts" as a subgroup of the automorphism group of $\circledcirc_\ell$. We remark that not all automorphisms are given by such shifts, but for our arguments we will not need to consider the full group of automorphisms.

Fractures of the torus. Recall that a fracture ρ of a graph H is a tuple $\rho = (\rho_v)_{v\in V(H)}$, where ρ_v is a partition of the set $E_H(v)$ of edges of H incident to v. Now given an automorphism $\varphi : H \to H$ of H, it gives a bijection from the edges $E_H(v)$ at v to the edges $E_H(\varphi(v))$ at $\varphi(v)$. Thus, given a fracture ρ of H, we obtain a fracture $\varphi(\rho)$ of H such that two edges $e_1, e_2 \in E_H(\varphi(v))$ are in the same block of $\varphi(\rho)_{\varphi(v)}$ if and only if their preimages $\varphi^{-1}(e_1), \varphi^{-1}(e_2) \in E_H(v)$ are in the same block of ρ_v.

We claim that that the two fractured graphs $H\sharp\rho \cong H\sharp\varphi(\rho)$ are isomorphic. To see this, note that the automorphism $\varphi : H \to H$ lifts to an automorphism $\widehat{\varphi} : M_H \to M_H$ of the matching M_H associated to H, where $\widehat{\varphi}$ sends the vertex v_e of M_H to $\varphi(v)_{\varphi(e)}$. The map $\widehat{\varphi}$ sends the equivalence relation on M_H associated to ρ (with quotient $H\sharp\rho$) to the equivalence relation associated to $\varphi(\rho)$ (with quotient $H\sharp\varphi(\rho)$). Thus $\widehat{\varphi} : M_H \to M_H$ descends to an isomorphism $H\sharp\rho \to H\sharp\varphi(\rho)$ fitting in a diagram of graph homomorphisms, depicted in Figure 5.

Given a finite group G acting on the graph H by graph isomorphisms $\varphi_g : H \to H$ (for $g \in G$), we obtain an action $\Vdash$ of G on the lattice $\mathcal{L}(H)$ of fractures on H, where $g \in G$ acts by $g \Vdash \rho = \varphi_g(\rho)$. Clearly, this action respects the order of the lattice ($\rho \le \rho'$ if and only if $g \Vdash \rho \le g \Vdash \rho'$), and as seen above, for any $\widetilde{\rho}$ in the G-orbit of ρ we have $H\sharp\widetilde{\rho} \cong H\sharp\rho$.

We now return to the special case when $H = \circledcirc_\ell$ is a torus. Here, given a vertex (i,j) of H it is convenient to identify the edges incident to the vertex (connecting it to $(i,j+1), (i,j-1), (i-1,j)$, and $(i+1,j)$) with the four "directions" $\vartriangle, \triangledown, \triangleleft$, and $\triangleright$, respectively, so that each $\rho_{(i,j)}$ is a partition of the set $\{\vartriangle, \triangledown, \triangleleft, \triangleright\}$.

We have seen that $\mathbb{Z}_\ell^2$ acts transitively on the vertices of $\circledcirc_\ell$ in such a way that every element of $\mathbb{Z}_\ell^2$ induces an automorphism of $\circledcirc_\ell$. Thus, by the discussion above, we obtain an action $\Vdash$ of $\mathbb{Z}_\ell^2$ on the set of fractures of $\circledcirc_\ell$. Let us make this action explicit: $(i,j) \Vdash \rho := \hat{\rho}$, where $\hat{\rho}_{((i,j)\vdash(i',j'))} = \rho_{(i',j')}$ for all $(i',j') \in \mathbb{Z}_\ell^2$.

Analysis of the fixed points. We proceed with the fixed points of the action $\Vdash$ of $\mathbb{Z}_\ell^2$ on the fractures $\circledcirc_\ell$. Since this action consists of (all possible) (i,j)-shifts, the fixed points are precisely those fractures ρ for which all partitions $\rho_{(i,j)}$ are equal—recall that we assumed every $\rho_{(i,j)}$ to be a partition of $\{\vartriangle, \triangledown, \triangleleft, \triangleright\}$. Fortunately, there are only 15 partitions of the four-element set, and thus we can analyze the fixed points by hand. Indeed, one special case of our main result, as well as the classification of the parameterized Tutte polynomial, relies on the understanding of all of those 15 fixed

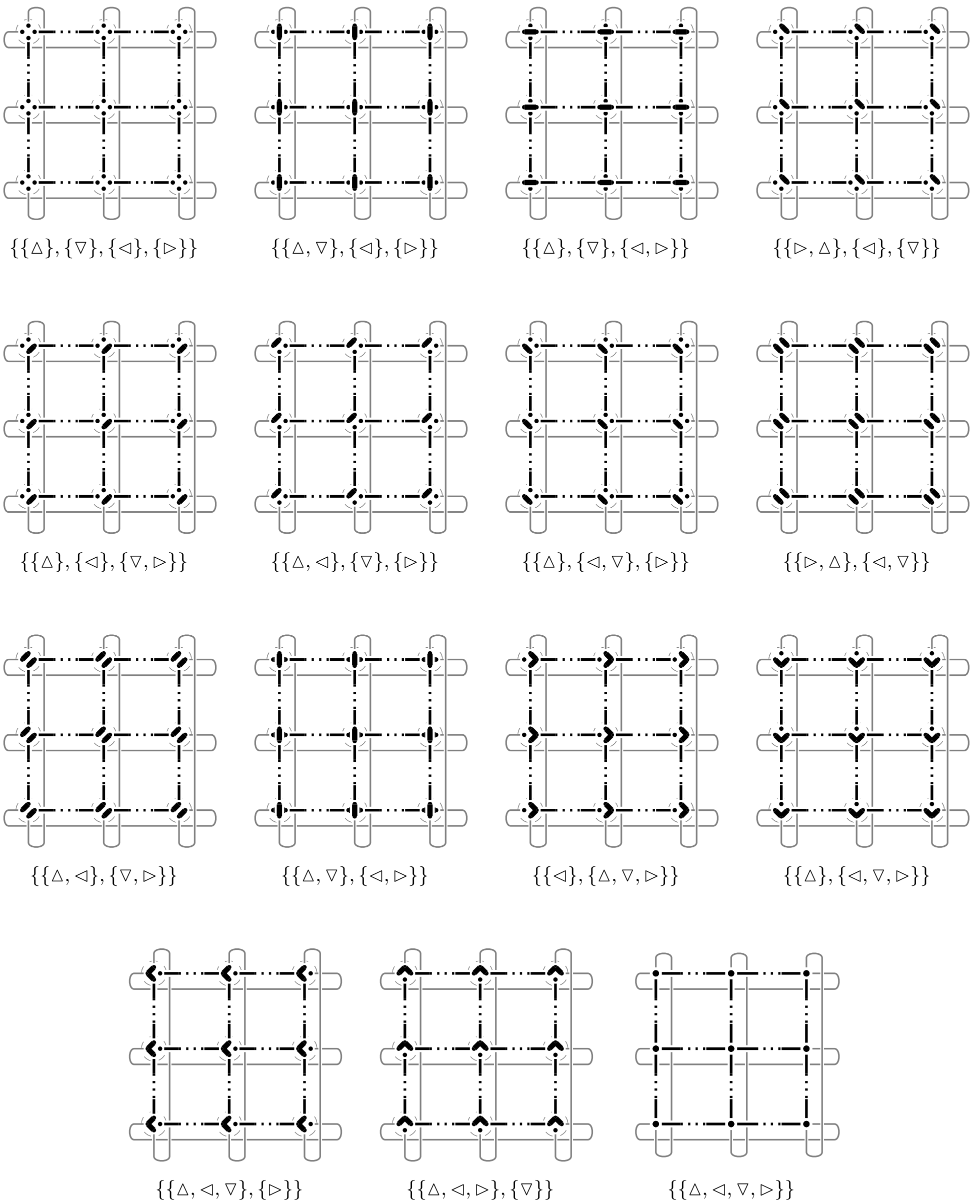

FIG. 6. *The 15 fixed points of the action $\Vdash$ of $\mathbb{Z}_\ell^2$ on the fractures of $\circledcirc_\ell$. Vertices in the fractured graphs that correspond to the same vertex in the original graph are encircled with a dashed line; the corresponding fixed point is denoted below its graphical representation.*

points. However, while there are 15 fixed points ρ, we can group those into 7 types according to the isomorphism class of $\circledcirc_\ell\sharp\rho$; an illustration of all fixed points is given in Figure 6.

Observation 5.3. The fixed points of the action of $\mathbb{Z}_\ell^2$ on the fractures of $\circledcirc_\ell$ are as follows.

Matching: $\circledcirc_\ell\sharp\rho \cong M_{2\ell^2}$, the matching of size $2\ell^2$.

1. $\rho_{(i,j)} = \{\{\triangle\},\{\triangledown\},\{\triangleleft\},\{\triangleright\}\}$ for all $(i,j)\in\mathbb{Z}_\ell^2$, that is, $\rho=\perp$.

Matching and cycles: $\circledcirc_\ell \sharp \rho \cong M_{\ell^2} + \ell C_\ell$, the union of a matching of size ℓ^2 and ℓ disjoint cycles of length ℓ.

2. $\rho_{(i,j)} = \{\{\vartriangle, \triangledown\}, \{\triangleleft\}, \{\triangleright\}\}$ for all $(i,j) \in \mathbb{Z}_\ell^2$.
3. $\rho_{(i,j)} = \{\{\vartriangle\}, \{\triangledown\}, \{\triangleleft, \triangleright\}\}$ for all $(i,j) \in \mathbb{Z}_\ell^2$.

Wedge packing: $\circledcirc_\ell \sharp \rho \cong \ell^2 P_2$, the union of ℓ^2 disjoint paths of length 2.

4. $\rho_{(i,j)} = \{\{\vartriangle, \triangleright\}, \{\triangledown\}, \{\triangleleft\}\}$ for all $(i,j) \in \mathbb{Z}_\ell^2$.
5. $\rho_{(i,j)} = \{\{\vartriangle, \triangleleft\}, \{\triangledown\}, \{\triangleright\}\}$ for all $(i,j) \in \mathbb{Z}_\ell^2$.
6. $\rho_{(i,j)} = \{\{\triangledown, \triangleleft\}, \{\vartriangle\}, \{\triangleright\}\}$ for all $(i,j) \in \mathbb{Z}_\ell^2$.
7. $\rho_{(i,j)} = \{\{\triangledown, \triangleright\}, \{\vartriangle\}, \{\triangleleft\}\}$ for all $(i,j) \in \mathbb{Z}_\ell^2$.

Cycle packing I: $\circledcirc_\ell \sharp \rho \cong 2\ell C_\ell$, the union of 2ℓ disjoint cycles of length ℓ.

8. $\rho_{(i,j)} = \{\{\vartriangle, \triangledown\}, \{\triangleleft, \triangleright\}\}$ for all $(i,j) \in \mathbb{Z}_\ell^2$.

Cycle packing II: $\circledcirc_\ell \sharp \rho \cong \ell C_{2\ell}$, the union of ℓ disjoint cycles of length 2ℓ.

9. $\rho_{(i,j)} = \{\{\vartriangle, \triangleright\}, \{\triangledown, \triangleleft\}\}$ for all $(i,j) \in \mathbb{Z}_\ell^2$.
10. $\rho_{(i,j)} = \{\{\vartriangle, \triangleleft\}, \{\triangledown, \triangleright\}\}$ for all $(i,j) \in \mathbb{Z}_\ell^2$.

Sun packing: $\circledcirc_\ell \sharp \rho \cong \ell S_\ell$, the union of ℓ suns of size ℓ. Here a a sun of size ℓ is obtained from a cycle of length ℓ by adding one "dangling" edge at every vertex of the cycle.

11. $\rho_{(i,j)} = \{\{\vartriangle\}, \{\triangledown, \triangleleft, \triangleright\}\}$ for all $(i,j) \in \mathbb{Z}_\ell^2$.
12. $\rho_{(i,j)} = \{\{\triangledown\}, \{\vartriangle, \triangleleft, \triangleright\}\}$ for all $(i,j) \in \mathbb{Z}_\ell^2$.
13. $\rho_{(i,j)} = \{\{\triangleleft\}, \{\vartriangle, \triangledown, \triangleright\}\}$ for all $(i,j) \in \mathbb{Z}_\ell^2$.
14. $\rho_{(i,j)} = \{\{\triangleright\}, \{\vartriangle, \triangleleft, \triangledown\}\}$ for all $(i,j) \in \mathbb{Z}_\ell^2$.

Torus: $\circledcirc_\ell \sharp \rho \cong \circledcirc_\ell$, the torus of size ℓ.

15. $\rho_{(i,j)} = \{\{\vartriangle, \triangledown, \triangleleft, \triangleright\}\}$ for all $(i,j) \in \mathbb{Z}_\ell^2$, that is, $\rho = \top$.

While it might be surprising at first glance, we observe that, for many properties Φ, our analysis of the complexity of $\#\text{EdgeSub}(\Phi)$ *only* depends on which of the previous 15 fixed points ρ satisfy that $\circledcirc_\ell \sharp \rho$ has the property Φ.

5.1.2. Symmetries of Cayley graph expanders of 2-groups. For the second family of Cayley graphs, we rely on an explicit construction of 4-regular Cayley graph expanders due to a subset of the authors [71]. They are constructed from an explicit infinite group Γ with generators x_0, x_1 and a sequence

$$\Gamma \supseteq N_0 \supseteq N_1 \supseteq N_2 \supseteq \cdots$$

of normal subgroups N_i of Γ such that the indices $[\Gamma : N_i] = 2^{t_i} = n_i$ are powers of 2 converging to infinity as i increases. Moreover, writing $\mathcal{K}_i$ for the quotient group Γ/N_i, the set of Cayley graphs $G_i := \mathcal{C}(\mathcal{K}_i, \{v_0^{\pm 1}, v_1^{\pm 1}\})$ is a family $\mathcal{G}$ of $(n_i, 4, c)$-expanders for some constant $c > 0$. Here, $v_0 = x_0 N_i$ and $v_1 = x_1 N_i$ are generators of $\mathcal{K}_i$.

Similar to the case of the toroidal grid, we obtain an action of the group $\mathcal{K}_i$ on the graph G_i, where an element $g \in \mathcal{K}_i$ acts on a vertex $v \in V(G_i) = \mathcal{K}_i$ sending it to $g \vdash v = gv$, where the latter is the product of g and v in the group $\mathcal{K}_i$. The action of g defines a graph automorphism of G_i since the four edges $\{v, vv_j^{\pm 1}\}$ (for $j = 0, 1$) at v are sent to the four edges $\{gv, gvv_j^{\pm 1}\}$ incident to gv.

Fractures of the Cayley graph expanders. For $v \in V(G_i) \cong \mathcal{K}_i$, the edges adjacent to v connect v to the vertices vs for $s \in S$ and thus can be uniquely labeled[18] by

[18]We use the same notation here as in the section about the torus grid $\circledcirc_\ell$. Many of the intuitions that we gained so far are still valid; for instance, the edge $\triangleright$ going out from the vertex v to vv_0 is equal to the edge $\triangleleft$ associated to the vertex vv_0. On the other hand, we also need to be more careful in our proofs, since, e.g., going along an edge $\triangleright$ followed by $\vartriangle$ does not necessarily go to the same vertex as the path $\vartriangle$ followed by $\triangleright$, since the group $\mathcal{K}_i$ is in general not abelian.

$$\rhd = \{v, vv_0\},\ \lhd = \{v, vv_0^{-1}\},\ \triangle = \{v, vv_1\},\ \triangledown = \{v, vv_1^{-1}\}.$$

Thus a fracture $\rho \in \mathcal{L}(G_i)$ is a collection $\rho = (\rho_v)_{v\in V(G_i)}$ of partitions of the set $\{\rhd, \lhd, \triangle, \triangledown\}$.

Analysis of the fixed points. As seen before, the action $\vdash$ of $\mathcal{K}_i$ on G_i induces an action $\Vdash$ of $\mathcal{K}_i$ on the lattice of partitions $\mathcal{L}(G_i)$. A fracture $\rho = (\rho_v)_{v\in V(G_i)}$ is invariant under the action of $\mathcal{K}_i$ if and only if ρ_v does not depend on v.

Later we want to compute the coefficient $a_{\Phi,G_i}(\top)$ modulo two. As before we observe that only fixed points of the action of $\mathcal{K}_i$ contribute, and additionally we observe that only such fixed points ρ can contribute where ρ_v has at most two blocks: The contribution of ρ_v to $a_{\Phi,G_i}(\top)$ contains a factor $(|\rho_v|-1)!$ which is even if ρ_v has at least 3 blocks (see Corollary 4.3). Thus in the following we consider fixed points in which each partition has at most 2 blocks.

LEMMA 5.4. *Fix $i \geq 2$, and let us denote $b_0 = \mathrm{ord}_{\mathcal{K}_i}(v_0), b_1 = \mathrm{ord}_{\mathcal{K}_i}(v_1)$, $b_2 = \mathrm{ord}_{\mathcal{K}_i}(v_1^{-1}v_0)$, $b_3 = \mathrm{ord}_{\mathcal{K}_i}(v_1v_0)$, and $a_j = \#\mathcal{K}_i/b_j$ for $j = 0,\ldots,3$. Then the fixed points $\rho = (\rho_v)_{v\in V(G_i)}$ of the action of $\mathcal{K}_i$ on the fractures of G_i satisfying that all ρ_v have at most two blocks are as follows:*

Cycle packing I:

1. $\rho_v = \{\{\triangle, \triangledown\}, \{\lhd, \rhd\}\}$ *for all* $v \in V(G_i)$ *and* $G_i\sharp\rho \cong a_0 \cdot C_{b_0} + a_1 \cdot C_{b_1}$.

Cycle packing II:

2. $\rho_v = \{\{\triangle, \rhd\}, \{\triangledown, \lhd\}\}$ *for all* $v \in V(G_i)$ *and* $G_i\sharp\rho \cong a_2 \cdot C_{2b_2}$.
3. $\rho_v = \{\{\triangle, \lhd\}, \{\triangledown, \rhd\}\}$ *for all* $v \in V(G_i)$ *and* $G_i\sharp\rho \cong a_3 \cdot C_{2b_3}$.

Sun packing:

4. $\rho_v = \{\{\triangle\}, \{\triangledown, \lhd, \rhd\}\}$ *for all* $v \in V(G_i)$ *and* $G_i\sharp\rho \cong a_0 \cdot S_{b_0}$.
5. $\rho_v = \{\{\triangledown\}, \{\triangle, \lhd, \rhd\}\}$ *for all* $v \in V(G_i)$ *and* $G_i\sharp\rho \cong a_0 \cdot S_{b_0}$.
6. $\rho_v = \{\{\lhd\}, \{\triangle, \triangledown, \rhd\}\}$ *for all* $v \in V(G_i)$ *and* $G_i\sharp\rho \cong a_1 \cdot S_{b_1}$.
7. $\rho_v = \{\{\rhd\}, \{\triangle, \lhd, \triangledown\}\}$ *for all* $v \in V(G_i)$ *and* $G_i\sharp\rho \cong a_1 \cdot S_{b_1}$.

Full graph:

8. $\rho_v = \{\{\triangle, \triangledown, \lhd, \rhd\}\}$ *for all* $v \in V(G_i)$*, that is,* $\rho = \top$ *and* $G_i\sharp\rho \cong G_i$.

Moreover, the numbers a_j, b_j are all powers of 2 and $a_j \geq 8$.

Proof. In cases 1,2,3 it follows from the definition of the fractured graph that $G_i\sharp\rho$ is 2-regular and thus a union of circles. In case 1 the first type of circles (associated to the directions $\lhd, \rhd$) is given by

$$w_0 \to w_0v_0 \to w_0v_0^2 \to \cdots \to w_0v_0^{b_0-1} \to w_0v_0^{b_0} = w_0 \tag{5.1}$$

and thus isomorphic to C_{b_0}, with one circle for each $w_0 \in K_i/\langle v_0\rangle$ giving a total number of $\#(\mathcal{K}_i/\langle v_0\rangle) = \#\mathcal{K}_i/b_i = a_i$. Analogously we obtain a_1 copies of C_{b_1} associated to the directions $\triangle, \triangledown$.

In case 2 the circles are of the form

$$\begin{aligned} w_0 &\to w_0v_1^{-1} \to w_0v_1^{-1}v_0 \to w_0v_1^{-1}v_0v_1^{-1} \to w_0(v_1^{-1}v_0)^2 \to \cdots \\ &\to w_0(v_1^{-1}v_0)^{b_2-1}v_1^{-1} \to w_0(v_1^{-1}v_0)^{b_2} = w_0. \end{aligned}$$

Thus they are isomorphic to C_{2b_2}, and since the total number of edges of $G_i\sharp\rho$ is equal to $\#E(G_i) = 4\#\mathcal{K}_i/2 = 2\#\mathcal{K}_i$, the number of copies of C_{2b_2} is given by $2\#\mathcal{K}_i/(2b_2) = a_2$. Case 3 is treated analogously.

In case 4, the connected component of a vertex $w_0 \in G_i\sharp\rho$ associated to the directions $\triangledown, \lhd, \rhd$ certainly contains the circle C_{b_0} given by (5.1), and in addition, each vertex $w_0 v_0^\ell$ is connected to $w_0 v_0^\ell v_1^{-1}$, which forms a leaf of $G_i \sharp \rho$. Thus, these are the only additional vertices connected to the circle, and thus the connected component of each vertex in $G_i \sharp \rho$ forms a sun S_{b_0}. The total number of suns is $\#E(G_i)/\#E(S_{b_0}) = (2\#\mathcal{K}_i)/(2b_0) = a_0$. The cases 5,6,7 are treated completely analogously.

Finally, case 8 follows from the general property $H \sharp \top \cong H$. The fact that a_j, b_j divide the order of $\#\mathcal{K}_i$, together with the property that $\mathcal{K}_i$ is a 2-group, implies that a_j, b_j are powers of 2. Finally, we show the inequality $a_j \geq 8$ by induction on $i \geq 2$. Note that in the case $i = 2$ this can be checked by hand. For this one uses the explicit description of the group law of $\mathcal{K}_2$ presented in [71, section 3] and verifies that the orders b_j of elements $v_0, v_1, v_1^{-1} v_0, v_1 v_0$ are precisely 4, so that $a_j = \#\mathcal{K}_2 / b_j = 8$.

To conclude the general case for $i \geq 2$, denote $V_0^i = \langle v_0 \rangle \subseteq \mathcal{K}_i$ the subgroup generated by v_0, so that $a_0 = [\mathcal{K}_i : V_0^i]$. Recalling the facts from the start of the section, we saw that $\mathcal{K}_i = G/N_i$ with $N_2 \supseteq N_i$ for $i \geq 3$. Thus we have a surjective group homomorphism

$$\varphi_i : \mathcal{K}_i = G/N_i \to G/N_2 = \mathcal{K}_2 \, , xN_i \mapsto xN_2$$

sending $V_0^i \subseteq \mathcal{K}_i$ to $V_0^2 \subseteq \mathcal{K}_2$ (this follows since the generator $v_0 = x_0 N_i$ of V_0^i maps to the generator $v_0 = x_0 N_2$ of V_0^2). As mentioned above, we checked by hand that V_0^2 has index 8 in $\mathcal{K}_2$. Then by Lemma 5.5 we have that $8 = [\mathcal{K}_2 : V_0^2] = [\varphi(\mathcal{K}_i) : \varphi(V_0^i)]$ divides $[\mathcal{K}_i : V_0^i] = a_i$ and thus $a_0 \geq 8$. The bounds for a_1, a_2, a_3 work exactly the same way. □

LEMMA 5.5. *Let Γ, Γ' denote finite groups and $\varphi : \Gamma \to \Gamma'$ a group homomorphism. Then, for any subgroup $H \subseteq \Gamma$ we have that $[\varphi(\Gamma) : \varphi(H)]$ divides $[\Gamma : H]$.*

Proof. Let $K = \ker(\varphi)$ and $K_H = K \cap H = \ker(\varphi|_H)$; then by the first isomorphism theorem we have $\varphi(\Gamma) \cong \Gamma/K$ and $\varphi(H) \cong H/K_H$. Using this, we have

$$\begin{aligned}[\Gamma : H] &= \frac{\#\Gamma}{\#H} = \frac{\#\Gamma/\#K}{\#H/\#K_H} \cdot \frac{\#K}{\#K_H} = \frac{\#(\Gamma/K)}{\#(H/K_H)} \cdot \frac{\#K}{\#K_H} = \frac{\#\varphi(\Gamma)}{\#\varphi(H)} \cdot \frac{\#K}{\#K_H} \\ &= [\varphi(\Gamma) : \varphi(H)] \cdot \frac{\#K}{\#K_H}.\end{aligned}$$

But $K_H \subseteq K$ is a subgroup, so by Lagrange's theorem, the order of K_H divides the order of K, so that $\#K/\#K_H$ is an integer. Thus the above equality shows that $[\varphi(\Gamma) : \varphi(H)]$ divides $[\Gamma : H]$. □

5.1.3. Analysis of the coefficient function via fixed points. While the value $a_{\Phi,H}(\top)$ of the coefficient function seems to be very difficult to handle for arbitrary graphs H, we now use our observations on the symmetries of the torus and the Cayley graph expanders to prove that the coefficient function does not vanish infinitely often under specific constraints on the behavior of Φ on the fixed points presented in the preceding section.

We start with the case of $a_{\Phi,\circledcirc_\ell}(\top)$, which, while being simple, turns out to be required for one of the special cases in our main classification for minor-closed graph properties.

LEMMA 5.6. *Let ℓ denote a prime, and let Φ denote a computable graph property. We have that*

$$\begin{aligned}a_{\Phi,\circledcirc_\ell}(\top) = &-6\Phi(M_{2\ell^2}) + 4\Phi(M_{\ell^2} + \ell C_\ell) + 8\Phi(\ell^2 P_2) \\ &- \Phi(2\ell C_\ell) - 2\Phi(\ell C_{2\ell}) - 4\Phi(\ell S_\ell) + \Phi(\circledcirc_\ell) \mod \ell\,.\end{aligned}$$

Proof. By Corollary 4.3 we have

$$a_{\Phi,\circledcirc_\ell}(\top) = \sum_{\sigma \in \mathcal{L}(\Phi,\circledcirc_\ell)} \prod_{v \in V(\circledcirc_\ell)} (-1)^{|\sigma_v|-1} \cdot (|\sigma_v| - 1)!.$$

Setting $f(\sigma) := \prod_{v \in V(\circledcirc_\ell)} (-1)^{|\sigma_v|-1} \cdot (|\sigma_v| - 1)!$, this rewrites to

$$a_{\Phi,\circledcirc_\ell}(\top) = \sum_{\sigma \in \mathcal{L}(\Phi,\circledcirc_\ell)} f(\sigma)\,.$$

We now use the action $\Vdash$ of $\mathbb{Z}_\ell^2$ on the subset $\mathcal{L}(\Phi,\circledcirc_\ell)$ of $\mathcal{L}(\circledcirc_\ell)$, given by permuting the elements of a fracture ρ according to the coordinate shift induced by an element $(i,j) \in \mathbb{Z}_\ell^2$. Restricting this action to $\mathcal{L}(\Phi,\circledcirc_\ell)$ is well defined since the action does not change the isomorphism class[19] of $\circledcirc_\ell \sharp \rho$. In particular, we have that $f(\sigma) = f(\rho)$ whenever σ and ρ are in the same orbit of the action. This allows us to rewrite as follows; the sum is taken over all orbits $[\sigma]$ of the group action:

$$a_{\Phi,\circledcirc_\ell}(\top) = \sum_{[\sigma]} \#[\sigma] \cdot f(\sigma).$$

Since ℓ is a prime, the group order of $\mathbb{Z}_\ell^2$ is a power of ℓ. As the size of every orbit must divide the group order, we can ignore all orbits which are not fixed points, that is, σ for which $\#[\sigma] \neq 1$, if we take the sum modulo ℓ. All 15 fixed points are explicitly given in Observation 5.3. Let us now compute the coefficients of each collection of fixed points that induce the same graph, up to isomorphism; we use Fermat's little theorem—recall that ℓ is a prime.

Matching: One fixed point ρ satisfies $\circledcirc_\ell \sharp \rho \cong M_{2\ell^2}$. The contribution to $a_{\Phi,\circledcirc_\ell}(\top)$ is thus

$$1 \cdot f(\rho) \cdot \Phi(M_{2\ell^2}) = ((-1)^{4-1} \cdot (4-1)!)^{\ell^2} \Phi(M_{2\ell^2}) = -6\Phi(M_{2\ell^2}) \mod \ell\,.$$

Matching and cycles: Two fixed points ρ satisfy $\circledcirc_\ell \sharp \rho \cong M_{\ell^2} + \ell C_\ell$. The contribution to $a_{\Phi,\circledcirc_\ell}(\top)$ is thus

$$2 \cdot f(\rho) \cdot \Phi(M_{\ell^2} + \ell C_\ell) = 2 \cdot ((-1)^{3-1} \cdot (3-1)!)^{\ell^2} \Phi(M_{\ell^2} + \ell C_\ell) = 4\Phi(M_{\ell^2} + \ell C_\ell) \mod \ell\,.$$

Wedge packing: Four fixed points ρ satisfy $\circledcirc_\ell \sharp \rho \cong \ell^2 P_2$. The contribution to $a_{\Phi,\circledcirc_\ell}(\top)$ is thus

$$4 \cdot f(\rho) \cdot \Phi(\ell^2 P_2) = 4 \cdot ((-1)^{3-1} \cdot (3-1)!)^{\ell^2} \Phi(\ell^2 P_2) = 8\Phi(\ell^2 P_2) \mod \ell\,.$$

Cycle packing I: One fixed point ρ satisfies $\circledcirc_\ell \sharp \rho \cong 2\ell C_\ell$. The contribution to $a_{\Phi,\circledcirc_\ell}(\top)$ is thus

$$1 \cdot f(\rho) \cdot \Phi(2\ell C_\ell) = ((-1)^{2-1} \cdot (2-1)!)^{\ell^2} \Phi(2\ell C_\ell) = -\Phi(2\ell C_\ell) \mod \ell\,.$$

Cycle packing II: Two fixed points ρ satisfy $\circledcirc_\ell \sharp \rho \cong \ell C_{2\ell}$. The contribution to $a_{\Phi,\circledcirc_\ell}(\top)$ is thus

$$2 \cdot f(\rho) \cdot \Phi(\ell C_{2\ell}) = 2 \cdot ((-1)^{2-1} \cdot (2-1)!)^{\ell^2} \Phi(\ell C_{2\ell}) = -2\Phi(\ell C_{2\ell}) \mod \ell\,.$$

[19] Note that while the action *can* change the isomorphism class *as a $\circledcirc_\ell$-colored graph*, the property Φ only depends on the underlying uncolored graph, which is unchanged, and thus $\mathcal{L}(\Phi,\circledcirc_\ell)$ is indeed invariant under the action.

Sun packing: Four fixed points ρ satisfy $⊚_\ell\sharp\rho \cong \ell S_\ell$. The contribution to $a_{\Phi,⊚_\ell}(\top)$ is thus

$$4 \cdot f(\rho) \cdot \Phi(\ell S_\ell) = 4 \cdot ((-1)^{2-1} \cdot (2-1)!)^{\ell^2} \Phi(\ell S_\ell) = -4\Phi(\ell S_\ell) \mod \ell\,.$$

Torus: One fixed point ρ satisfies $⊚_\ell\sharp\rho \cong ⊚_\ell$. The contribution to $a_{\Phi,⊚_\ell}(\top)$ is thus

$$1 \cdot f(\rho) \cdot \Phi(⊚_\ell) = ((-1)^{1-1} \cdot (1-1)!)^{\ell^2} \Phi(⊚_\ell) = \Phi(⊚_\ell) \mod \ell\,.$$

Taking the sum of the previous terms (modulo ℓ) concludes the proof. ∎

We proceed with a similar lemma for the Cayley graph expanders.

LEMMA 5.7. *Let $\mathcal{G} = \{G_1, G_2, \ldots\}$ denote the family of Cayley graph expanders given in subsection* 5.1.2*, and let Φ denote a computable graph property. For $i \geq 2$ we have*

$$a_{\Phi,G_i}(\top) = \Phi(a_0 \cdot C_{b_0} + a_1 \cdot C_{b_1}) + \Phi(a_2 \cdot C_{2b_2}) + \Phi(a_3 \cdot C_{2b_3}) + \Phi(G_i) \mod 2\,.$$

Moreover, the numbers a_j, b_j are all powers of 2 *and $a_j \geq 8$.*

Proof. By Corollary 4.3 we have

$$a_{\Phi,G_i}(\top) = \sum_{\rho \in \mathcal{L}(\Phi,G_i)} \prod_{v \in V(G_i)} (-1)^{|\rho_v|-1} \cdot (|\rho_v|-1)!.$$

Setting $f(\rho) := \prod_{v \in V(G_i)} (-1)^{|\rho_v|-1} \cdot (|\rho_v|-1)!$, this rewrites to

$$a_{\Phi,G_i}(\top) = \sum_{\rho \in \mathcal{L}(\Phi,G_i)} f(\rho)\,.$$

As before, the action of the 2-group $\mathcal{K}_i$ leaves the set $\mathcal{L}(\Phi, G_i)$ invariant, and modulo 2 the contribution of all elements ρ not fixed under $\mathcal{K}_i$ vanishes. Thus we only consider the fixed points $\rho = (\rho_v)_{v \in V(G_i)}$, for which ρ_v is independent of v.

From the formula of $f(\rho)$ it is easy to see that $f(\rho) = 1 \mod 2$ if ρ has at most two blocks and $f(\rho) = 0 \mod 2$ otherwise. Thus only the fractures ρ from cases 1 to 8 of Lemma 5.4 can give a nontrivial contribution to $a_{\Phi,G_i}(\top)$. The fixed point ρ contributes if and only if $\Phi(G_i\sharp\rho) = 1$. Finally, since the pairs of cases 4,5 and 6,7 lead to isomorphic graphs $G_i\sharp\rho$, any possible contributions from these cancel modulo 2, and we are left with the four summands above. ∎

5.2. 3-group Cayley graph expanders and suitable graph properties. The proof of our classification for minor-closed properties will require a separate treatment of properties that are false on cycles. Formally, we consider *suitable* graph properties; recall that C_k denotes the cycle of length k.

DEFINITION 5.8. *A graph property Φ is called suitable if there exists $k \geq 3$ such that $\Phi(C_k) = 0$.*

Our goal in this subsection is to show that certain suitable and minor-closed graph properties yield nonzero values of the coefficient function for the family of 4-regular 3-group Cayley graph expanders constructed in Theorem 3.7. This will later allow us to infer hardness by Lemma 4.8.

For the formal statement, we call a path with two edges a *wedge*, and we define a *wedge packing* of size k, denoted by kP_2, as the graph by a (disjoint) union of k

wedges. Furthermore, we say that a graph property Φ has *bounded wedge-number* if there is a constant d such that $\Phi(kP_2)=0$ for each $k\geq d$. Otherwise, Φ has *unbounded wedge-number*.

Now, given a family $\mathcal{G}=\{G_1,G_2,\dots\}$ of expander graphs, we recall that $\mathcal{H}[\Phi,\mathcal{G}]$ denotes the set of all G_i such that $a_{\Phi,G_i}(\top)\neq 0$.

LEMMA 5.9. *Let Φ be a minor-closed graph property. If Φ is suitable and of unbounded wedge-number, then there exists a family $\mathcal{G}=\{G_1,G_2,\dots\}$ of 4-regular 3-group Cayley graph expanders such that $\mathcal{H}[\Phi,\mathcal{G}]$ is infinite.*

Proof. We prove the lemma for the family $\mathcal{G}=\{G_1,G_2,\dots\}$ constructed in Theorem 3.7.

Since Φ is suitable, there exists t such that $\Phi(C_t)=0$. Now choose any i such that the orders of the groups H_i^1, H_i^2, H_i^3, and H_i^4 (from Theorem 3.7) are at least t. We will show that for each such choice of i, the coefficient function $a_{\Phi,G_i}(\top)$ is nonzero. Since, by Theorem 3.7, the orders of the four groups are unbounded, the coefficient function will then be nonzero infinitely often, proving that $\mathcal{G}$ is an obstruction.

Recall that

$$a_{\Phi,G_i}(\top)=\sum_{\sigma\in\mathcal{L}(\Phi,G_i)}\prod_{v\in V(G_i)}(-1)^{|\sigma_v|-1}\cdot(|\sigma_v|-1)!\,.$$

Recall further that $\mathcal{L}(\Phi,G_i)$ is the set of fractures ρ of G_i such that $G_i\sharp\rho$ satisfies Φ.

The graph G_i is 4-regular, and every vertex $v\in V(G_i)$ corresponds to a coset of the quotient group Γ/N_i. Moreover, every v is adjacent to vv_0, vv_0^{-1}, vv_1, and vv_1^{-1}, where v_0 and v_1 are the generators of Γ/N_i. It will be convenient to label the edges incident to v by

$$\rhd=\{v,vv_0\},\lhd=\{v,vv_0^{-1}\},\vartriangle=\{v,vv_1\},\triangledown=\{v,vv_1^{-1}\}.$$

Thus a fracture of G_i is a tuple $\rho=(\rho_v)_{v\in V(G_i)}$ of partitions of the set $\{\rhd,\lhd,\vartriangle,\triangledown\}$.

Similarly as in section 5.1.2, we observe that the quotient group Γ/N_i acts on the graph G_i by setting $g\vdash v:=gv$ for each $g\in\Gamma/N_i$ and $v\in V(G_i)=\Gamma/N_i$. Moreover, this action is transitive, and for each $g\in\Gamma/N_i$, the function $g\vdash\star$ is an automorphism of G_i.

The action $\vdash$ extends to an action $\Vdash$ of Γ/N_i on the set $\mathcal{L}(\Phi,H)$: Given $g\in\Gamma/N_i$ and $\rho\in\mathcal{L}(\Phi,H)$, the fracture $g\Vdash\rho$ is obtained from ρ by permuting its entries according to the automorphism $g\vdash\star$. Again, similarly as in section 5.1.2, this action is well defined since $H\sharp\rho$ is isomorphic to $H\sharp(g\Vdash\rho)$. Thus $H\sharp(g\Vdash\rho)$ satisfies Φ if and only if $H\sharp\rho$ does. In particular, setting

$$f(\sigma):=\prod_{v\in V(G_i)}(-1)^{|\sigma_v|-1}\cdot(|\sigma_v|-1)!\,,$$

we have that $f(\sigma)=f(\rho)$ whenever σ and ρ are in the same orbit.

This enables us to rewrite

$$a_{\Phi,G_i}(\top):=\sum_{[\rho]}\#[\rho]\cdot f(\rho)\,,$$

where the sum is over all orbits $[\rho]$ of the action. We will proceed by considering $a_{\Phi,G_i}(\top)\mod 3$. Recall that the size of any orbit must divide the order of the group. Since Γ/N_i is a 3-group, only orbits of size 1, i.e., fixed points, survive modulo 3. Due to transitivity of the group action on the vertices, the only fixed points are fractures ρ for which all ρ_v are equal. In what follows, we will thus abuse notation and write, e.g., $\rho = \{\{\triangleright, \triangleleft\}, \{\triangle\}, \{\triangledown\}\}$ for the fixed point ρ in which $\rho_v = \{\{\triangleright, \triangleleft\}, \{\triangle\}, \{\triangledown\}\}$ for all $v \in V(G_i)$.

We will analyze the contribution to $a_{\Phi,G_i}(\top)$ (modulo 3) of any possible fixed points in the subsequent series of claims.

CLAIM 5.10. *If $\triangleleft$ and $\triangleright$ are in the same block of ρ, then $\rho \notin \mathcal{L}(\Phi, G_i)$.*

Proof. If $\triangleleft$ and $\triangleright$ are in the same block, then $G_i \sharp \rho$ contains the (simple) cycle C_b given by

$$e \to v_0 \to v_0^2 \to \cdots \to v_0^{b-1} \to e\,,$$

where b is the order of v_0 (and thus equal to the order of H_i^1). By our choice of i, we have $b \geq t$. Consequently $C_t \prec G_i \sharp \rho$. Since $\Phi(C_t) = 0$ and Φ is minor-closed, we conclude that $\Phi(G_i \sharp \rho) = 0$ and thus $\rho \notin \mathcal{L}(\Phi, G_i)$. ∎

CLAIM 5.11. *If $\triangle$ and $\triangledown$ are in the same block of ρ, then $\rho \notin \mathcal{L}(\Phi, G_i)$.*

Proof. Analogously to the previous claim, substitute v_0 by v_1, and H_i^1 by H_i^2. ∎

CLAIM 5.12. $\{\{\triangle, \triangleright\}, \{\triangledown, \triangleleft\}\} \notin \mathcal{L}(\Phi, G_i)$.

Proof. The fractured graph $G_i \sharp \{\{\triangle, \triangleright\}, \{\triangledown, \triangleleft\}\}$ contains the (simple) cycle C_{2b} given by

$$e \to v_1^{-1} \to v_1^{-1}v_0 \to v_1^{-1}v_0v_1^{-1} \to (v_1^{-1}v_0)^2 \to \cdots \to (v_1^{-1}v_0)^{b-1} \to e\,,$$

where b is the order of $v_1^{-1}v_0$ (and thus equal to the order of H_i^3). By our choice of i, we have $2b \geq t$. Consequently $C_t \prec G_i \sharp \rho$. Since $\Phi(C_t) = 0$ and Φ is minor-closed, we conclude that $\Phi(G_i \sharp \rho) = 0$ and thus $\rho \notin \mathcal{L}(\Phi, G_i)$. ∎

CLAIM 5.13. $\{\{\triangle, \triangleleft\}, \{\triangledown, \triangleright\}\} \notin \mathcal{L}(\Phi, G_i)$.

Proof. Analogously to the previous claim, substitute v_1^{-1} by v_1, and H_i^3 by H_i^4. ∎

The only partitions of $\{\triangle, \triangledown, \triangleleft, \triangleright\}$ not covered by one of the previous four claims are the finest partition $\perp = \{\{\triangle\}, \{\triangledown\}, \{\triangleleft\}, \{\triangleright\}\}$, as well as the following four:

- $\rho_1 = \{\{\triangle, \triangleleft\}, \{\triangledown\}, \{\triangleright\}\}$,
- $\rho_2 = \{\{\triangle, \triangleright\}, \{\triangledown\}, \{\triangleleft\}\}$,
- $\rho_3 = \{\{\triangledown, \triangleleft\}, \{\triangle\}, \{\triangleright\}\}$,
- $\rho_4 = \{\{\triangledown, \triangleright\}, \{\triangle\}, \{\triangleleft\}\}$.

First note that $\perp \in \mathcal{L}(\Phi, G_i)$ since Φ has unbounded wedge-number, the fractured graph of $\perp$ is a matching, and every matching is a minor of a sufficiently large wedge packing. However, the contribution from $\perp$ still vanishes, since we have

$$f(\perp) = \prod_{v \in V(G_i)} (-1)^{4-1} \cdot (4-1)! = 0 \mod 3\,.$$

Finally, we observe that $G_i \sharp \rho_j$ is a wedge packing for $j = 1, \ldots, 4$. Since Φ is minor-closed and of unbounded wedge-number, we obtain that $\Phi(G_i \sharp \rho_j) = 1$ and thus $G_i \sharp \rho_j \in \mathcal{L}(\Phi, G_i)$ for $j = 1, \ldots, 4$. The latter implies

$$a_{\Phi,G_i}(\top) = 4 \cdot \prod_{v \in V(G_i)} (-1)^{3-1} \cdot (3-1)! = 2^{|V(G_i)|} \mod 3\,,$$

concluding the proof, since $2^{|V(G_i)|} \neq 0 \mod 3$. ∎

5.3. p-group Cayley graph expanders and forests. While the property of being a forest is minor-closed and will thus be covered by our classification for minor-closed graph properties, we nevertheless treat this case separately. The reason for the latter is the fact that the coefficient function of this property turns out to not vanish modulo p for any prime $p \geq 3$, which has strong implications for modular counting as discussed in section 6.3.

Since the case of 3-group Cayley graph expanders was covered in the previous subsection—note that the property of being a forest is suitable and of unbounded wedge-number—we focus on $p \geq 5$ in what follows.

The crucial ingredients of our proof are the p-group Cayley graph expanders with $p-2$ generators constructed in Theorem 3.10. In particular, the relatively low degree of $2(p-2)$ provides us with additional control over their coefficient in the color-prescribed homomorphism basis and ultimately allows us to establish that those Cayley graph expanders have nonzero coefficients modulo p.

Let Γ be a finite group, let $S_0 \subseteq \Gamma$ be a set of m generators, and let $S = \{g^{\pm 1} : g \in S_0\} \subseteq \Gamma$ be the associated symmetric set of $2m$ generators.[20] Let $G = \mathcal{C}(\Gamma, S)$ be the associated Cayley graph. Similarly as in the proof of Lemma 5.9, we have an action of Γ on G which extends to the fractures of G. We write $\mathcal{L}(\Phi, G)^{\Gamma}$ for the set of fixed points, i.e., fractures σ which are invariant under this action. As seen before, we can interpret σ as a partition of the set S into blocks.

We begin with a general construction that will appear prominently later: Let $m \geq 1$, and let S and S_0 be as above. Then for any set partition σ of S, we define a graph $\mathcal{H}(\sigma)$. It has a vertex w^B for each block B of σ, and its set of (multi-)edges $E(\mathcal{H}(\sigma))$ is given by

$$\left\{ \{w^B, w^{B'}\} : \text{ one multiedge for each } g \in S_0 \text{ such that } g \in B, g^{-1} \in B' \right\} . \tag{5.2}$$

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 m edges. An alternative construction of $\mathcal{H}(\sigma)$ is by taking the matching M_m on the vertex set S defined by the involution $s \mapsto s^{-1}$ on S and identifying all vertices in the same block of the partition σ.

The following two lemmas will be needed below.

LEMMA 5.14. *Let $m \geq 1$, and let S be a finite set with $2m$ elements. Given any (simple) graph T with m edges, the number of set partitions σ of S such that $\mathcal{H}(\sigma)$ is isomorphic to T is given by $2^m m!/|\mathsf{Aut}(T)|$.*

Proof. Given the data above, consider the two sets

$$M_0 = \{\sigma : \sigma \text{ partition of } S \text{ such that } \mathcal{H}(\sigma) \cong T\} ,$$
$$M = \left\{ (\sigma, \varphi) : \sigma \text{ partition of } S, \text{ and } \varphi : \mathcal{H}(\sigma) \xrightarrow{\sim} T \text{ isomorphism} \right\} .$$

In the lemma we want to count the number of elements of M_0, but as an auxiliary set we use M, which explicitly records the data of the isomorphism $\varphi : \mathcal{H}(\sigma) \xrightarrow{\sim} T$. First, we note that the automorphism group $\mathsf{Aut}(T)$ acts on M, where $\eta \in \mathsf{Aut}(T)$ sends the pair (σ, φ) to $(\sigma, \eta \circ \varphi)$. We claim that the action is free: since φ is an isomorphism, the equality $\eta \circ \varphi = \varphi$ implies that η is the identity. Thus the orbits of the action all have cardinality $|\mathsf{Aut}(T)|$. On the other hand, we observe that the map

$$M \to M_0, \quad (\sigma, \varphi) \mapsto \sigma$$

[20]In particular, we require that for $g \in S_0$ we have $g^{-1} \notin S_0$.

is surjective and the fibers of this map are precisely the orbits of the above action of $\mathsf{Aut}(T)$. Indeed, the surjectivity is clear from the definition, and given (σ,φ) and (σ,φ') in the fiber of σ, we have that the automorphism $\eta=\varphi'\circ\varphi^{-1}$ of T satisfies $\eta\cdot(\sigma,\varphi)=(\sigma,\varphi')$. Combining these two observations we see

$$|M|=|M_0|\cdot|\mathsf{Aut}(T)|.$$

Thus to conclude we need to show that $|M|=2^m m!$. To do this observe that we can identify the elements of S with the vertices of the matching M_m in such a way that there is an edge $\{g,g^{-1}\}$ for each $g\in S_0$. Given a partition σ of S, note that we can see $\mathcal{H}(\sigma)$ as the quotient of M_m obtained by identifying the vertices belonging to the blocks of σ. In particular there is a canonical, well-defined quotient map $q_\sigma : M_m\to\mathcal{H}(\sigma)$. Let $\mathsf{Sur}(M_m,T)$ be the set of surjections from M_m to the graph T (where we mean graph homomorphisms that are surjective, hence bijective, on the set of edges). Then we have a map

$$G:\mathsf{Sur}(M_m,T)\to M,\quad \psi\mapsto(\sigma=\{\psi^{-1}(w):w\in V(T)\},\overline{\psi}),$$

where $\overline{\psi}:\mathcal{H}(\sigma)\to T$ is the unique map such that $\psi=\overline{\psi}\circ q_\sigma$. A short computation shows that G is a bijection with inverse given by

$$G^{-1}:M\to\mathsf{Sur}(M_m,T),\quad (\sigma,\varphi)\mapsto\varphi\circ q_\sigma\,.$$

Thus the proof is finished once we show that $|\mathsf{Sur}(M_m,T)|=2^m m!$. But this is easy to see: to specify a surjection from M_m to T we exactly have to give a bijection from the set of edges of M_m to the m edges of T (for which we have $m!$ possibilities), and for each of these edges we have two choices of orientation in our map (giving the factor of 2^m), because T is a simple graph. ∎

LEMMA 5.15. *Given $n\geq 1$ we have*

$$\sum_{\substack{T \text{ tree on}\\ n \text{ vertices}}}\frac{1}{|\mathsf{Aut}(T)|}=\frac{n^{n-2}}{n!}\,,\tag{5.3}$$

where the sum goes over isomorphism classes of trees T.

Proof. For the proof we use Cayley's formula: the number of labeled trees $\widehat{T}$ on n vertices is given by n^{n-2}. The natural action of the symmetric group S_n on the n vertices induces an action on the set of labeled trees $\widehat{T}$, and the stabilizer of such a tree is equal to its automorphism group. Moreover, two labeled trees $\widehat{T_1},\widehat{T_2}$ are in the same orbit if and only if their underlying unlabeled graphs are isomorphic. Thus

$$n^{n-2}=\sum_{\substack{T \text{ tree on}\\ n\text{vertices}}}\frac{|S_n|}{|\mathsf{Aut}(T)|}=n!\cdot\sum_{\substack{T \text{ tree on}\\ n\text{vertices}}}\frac{1}{|\mathsf{Aut}(T)|}$$

by the orbit stabilizer theorem. ∎

Now recall that G is the Cayley graph of Γ and S. We show that the fractured graph $G\sharp\sigma$ is a forest if and only if $\mathcal{H}(\sigma)$ is.

LEMMA 5.16. *There exists a natural graph homomorphism*

$$\Psi:G\sharp\sigma\to\mathcal{H}(\sigma),\ v^B\mapsto w^B\tag{5.4}$$

which is surjective and a local isomorphism (i.e., the edges incident to v^B map bijectively to the edges at w^B). Moreover, the graph $G\sharp\sigma$ is a forest if and only if $\mathcal{H}(\sigma)$ is.

Proof. First we show that Ψ is a well-defined graph homomorphism. Given $v \in V(G) = \Gamma$, the edges of G incident to v are given by $\{v, vg\}$ for $g \in S$. Then, given a block B of σ, the edges incident to v^B inside $G\sharp\sigma$ are in bijection with B and given by

$$\{v^B, (vg)^{B'}\} \text{ for } g \in B, g^{-1} \in B'.$$

Compared to the edges (5.2) of $\mathcal{H}(\sigma)$ we see that Ψ is not only a well-defined graph homomorphism but in fact, as claimed above, a local isomorphism. Finally, the surjectivity (both on vertices and edges) is also clear.

To see the last claim, first assume that $G\sharp\sigma$ is *not* a forest, and let C be a a circular walk without backtracking[21] inside $G\sharp\sigma$. Under the homomorphism Ψ it maps to a circular walk $\Psi(C)$, and, since Ψ is a local isomorphism, there is again no backtracking in $\Psi(C)$. Thus, the graph $\mathcal{H}(\sigma)$ is not a forest.

Conversely let C' in $\mathcal{H}(\sigma)$ be a circular walk without backtracking starting at some vertex w^B. Choose a vertex v_0^B in the preimage of w^B under Ψ, and let C in $G\sharp\sigma$ be the unique lift of the walk C'. By this we mean that we start at v_0^B, and for the first edge taken by the path C' from w^B, we take the unique edge incident to v_0^B mapping to it. Iterating the process for the subsequent edges of C' we obtain the walk C, which terminates at some vertex v_1^B. The whole process can now itself be iterated: we continue the walk C by concatenating it with the unique lift of C' starting this time at vertex v_1^B, terminating at v_2^B, and we can continue from there. In this way we can obtain an arbitrarily long walk in the graph $G\sharp\sigma$. But note that this walk involves no backtracking (since the original walk C' in $\mathcal{H}(\sigma)$ had no backtracking and Ψ is a local isomorphism). Thus, since the graph $G\sharp\sigma$ is finite, the infinite walk must contain a circular subwalk which, as seen before, involves no backtracking. Thus $G\sharp\sigma$ is not a forest. ∎

PROPOSITION 5.17. *Let $p \geq 5$ be a prime, Γ a finite p-group, and $S_0 \subseteq \Gamma$ a set of $q = p - 2$ generators such that $S = \{g^{\pm 1} : g \in S_0\}$ has $2q$ elements. Then for the property Φ of being a forest, we have $a_{\Phi,G}(\top) \neq 0 \mod p$.*

Proof. Recall that the number $a_{\Phi,G}(\top)$ is given by

$$a_{\Phi,G}(\top) = \sum_{\sigma \in \mathcal{L}(\Phi,G)} \prod_{v \in V(G)} (-1)^{|\sigma_v|-1} \cdot (|\sigma_v| - 1)!.$$

As before, when evaluating modulo p, we can reduce to the fractures $\sigma \in \mathcal{L}(\Phi, G)^\Gamma$ invariant under the action of Γ. Such σ can be interpreted as partitions of the set S. Rewriting the above formula we have

$$a_{\Phi,G}(\top) \equiv \sum_{\sigma \in \mathcal{L}(\Phi,G)^\Gamma} \Big((-1)^{|\sigma|-1} \cdot (|\sigma| - 1)!\Big)^{|V(G)|} \mod p.$$

Since $|V(G)| = |\Gamma|$ is a power of p by the assumption that Γ is a p-group, by Fermat's little theorem we have $u^{|V(G)|} \equiv u \mod p$ for all integers u so that we can remove the exponent $|V(G)|$ in the formula above. Looking at the index set of the sum, we note that for $\sigma \in \mathcal{L}(\Phi, G)^\Gamma$ with $|\sigma| > p$ we have that p divides the term $(|\sigma| - 1)!$ inside

[21]A walk is a sequence of vertices $w_0, w_1, \ldots, w_\ell$ such that w_{i-1} and w_i are connected by an edge for $i = 1, \ldots, \ell$. We say that the walk is without backtracking if there does not exist an i such that $w_{i-1} = w_{i+1}$.

the sum, so that the corresponding summands vanish modulo p. On the other hand, for the property Φ of being a forest, we have by Lemma 5.16 that a partition σ of S is contained in $\mathcal{L}(\Phi,G)^\Gamma$ if and only if $\mathcal{H}(\sigma)$ is a forest.

The graph $\mathcal{H}(\sigma)$ has $|S_0|=p-2$ edges and $|\sigma|$ many vertices, where, as seen above, we can assume $|\sigma|\le p$. If $\mathcal{H}(\sigma)$ is a forest, then the number of trees it contains (the connected components) is its Euler characteristic $|V(\mathcal{H}(\sigma))|-|E(\mathcal{H}(\sigma))|=|\sigma|-(p-2)$. For $|\sigma|\le p-2$, the graph $\mathcal{H}(\sigma)$ has at least as many edges as it has vertices, and thus it can never be a forest. We are left with two cases:[22]

- Case I : $|V(\mathcal{H}(\sigma))|=|\sigma|=p-1$, which forces $\mathcal{H}(\sigma)\cong T$ to be a tree T.
- Case II : $|V(\mathcal{H}(\sigma))|=|\sigma|=p$, which forces $\mathcal{H}(\sigma)\cong T_1+T_2$ to be a union of two trees T_1,T_2. We remark for later that, since the total number p of vertices is odd, the two trees T_1,T_2 cannot be isomorphic (since one has odd and one has even number of vertices). Therefore we have

$$|\mathsf{Aut}(T_1+T_2)|=|\mathsf{Aut}(T_1)|\cdot|\mathsf{Aut}(T_2)|\,. \tag{5.5}$$

Denote by $\mathcal{L}_{\mathrm{I}},\mathcal{L}_{\mathrm{II}}$ the set of partitions of S corresponding to Cases I, II above. Then the current status of the calculation is that

$$\begin{aligned}a_{\Phi,G}(\top)&\equiv\sum_{\sigma\in\mathcal{L}_{\mathrm{I}}\sqcup\mathcal{L}_{\mathrm{II}}}(-1)^{|\sigma|-1}\cdot(|\sigma|-1)!\mod p\\&\equiv\sum_{\sigma\in\mathcal{L}_{\mathrm{I}}}(-1)^{p-2}\cdot(p-2)!+\sum_{\sigma\in\mathcal{L}_{\mathrm{II}}}(-1)^{p-1}\cdot(p-1)!\mod p\\&\equiv-|\mathcal{L}_{\mathrm{I}}|-|\mathcal{L}_{\mathrm{II}}|\mod p\,,\end{aligned}\tag{5.6}$$

where we used that p is odd and, due to Wilson's theorem, we have

$$-(p-2)!\equiv(p-1)\cdot(p-2)!\equiv(p-1)!\equiv-1\mod p\,.$$

To count the number of elements $\sigma\in\mathcal{L}_{\mathrm{I}}$, we can group them according to the isomorphism class T of the tree $\mathcal{H}(\sigma)$. Then, by Lemma 5.14 we have

$$|\mathcal{L}_{\mathrm{I}}|=2^{p-2}(p-2)!\sum_{\substack{T\text{ tree on}\\p-1\text{ vertices}}}\frac{1}{|\mathsf{Aut}(T)|}=2^{p-2}(p-2)!\frac{(p-1)^{p-3}}{(p-1)!}=2^{p-2}(p-1)^{p-4}\,.$$

Here in the third equality we used Lemma 5.15. Plugging into the formula above we compute

$$\begin{aligned}-|\mathcal{L}_{\mathrm{I}}|&=-2^{p-2}(p-1)^{p-4}\equiv-2^{p-2}(-1)^{p-4}\equiv2^{p-1}\cdot\frac{p+1}{2}\cdot(-1)^{p-3}\\&\equiv\frac{p+1}{2}\mod p\,,\end{aligned}$$

where in the fourth congruence we used Fermat's little theorem.

We now turn to the sum in Case II. We are counting every forest twice by choosing a numbering T_1, T_2 for the two trees in the forest. It is important that never $T_1\cong T_2$ as remarked above, so that for all forests we overcount with the factor 2. Then we have

[22] As a reassurance to the reader, a priori, the graph $\mathcal{H}(\sigma)$ was allowed to have loops or multiedges. However, any graph $\mathcal{H}(\sigma)$ having either of those is certainly not a forest (since it has a cycle of length 1 or 2, respectively), and thus from here on, only standard (simple) graphs appear.

(5.7)

$$
\begin{aligned}
|\mathcal{L}_{\mathrm{II}}| &= \frac{1}{2} \cdot 2^{p-2}(p-2)! \sum_{j=1}^{p-1} \sum_{\substack{T_1 \text{ tree, } j \text{ vertices,} \\ T_2 \text{ tree, } p-j \text{ vertices}}} \frac{1}{|\mathsf{Aut}(T_1+T_2)|} \\
&= 2^{p-3}(p-2)! \sum_{j=1}^{p-1} \left(\sum_{T_1 \text{ tree, } j \text{ vertices}} \frac{1}{|\mathsf{Aut}(T_1)|} \right) \cdot \left(\sum_{T_2 \text{ tree, } p-j \text{ vertices}} \frac{1}{|\mathsf{Aut}(T_2)|} \right) \\
&= 2^{p-3}(p-2)! \sum_{j=1}^{p-1} \frac{j^{j-2}}{j!} \cdot \frac{(p-j)^{p-j-2}}{(p-j)!} \\
&= 2^{p-3} \sum_{j=1}^{p-1} \big(j^{j-2}(p-j)^{p-j-3} \cdot \prod_{k=2}^{j} \frac{p-k}{k} \big) ,
\end{aligned}
$$

where the first equality uses Lemma 5.14, the second uses the observation (5.5), and the third uses Lemma 5.15. In the fourth equation we rearranged the factors of the factorials.

To continue, we observe that the final formula (5.7) for $|\mathcal{L}_{\mathrm{II}}|$ in fact *is well defined modulo* p. That is, we never divide by a number divisible by p. Thus we are allowed to evaluate and simplify this expression modulo p:

$$
\begin{aligned}
-|\mathcal{L}_{\mathrm{II}}| &= -2^{p-3} \sum_{j=1}^{p-1} \big(j^{j-2}(p-j)^{p-j-3} \cdot \prod_{k=2}^{j} \frac{p-k}{k} \big) \mod p \\
&\equiv -\frac{1}{4} 2^{p-1} \sum_{j=1}^{p-1} j^{j-2}(-j)^{p-j-3} \cdot (-1)^{j-1} \mod p \\
&\equiv \frac{1}{4} \sum_{j=1}^{p-1} j^{p-5}(-1)^{p-5} \equiv \frac{1}{4} \sum_{j=1}^{p-1} j^{p-5} \mod p \,.
\end{aligned}
$$

To simplify the sum from $j=1$ to $p-1$, we observe that j ranges over $\mathbb{F}_p^{\times}$. The map $\mathbb{F}_p^{\times} \to \mathbb{F}_p^{\times}, j \mapsto 2j$ is a bijection. Thus we have

(5.8)

$$
16 \sum_{j=1}^{p-1} j^{p-5} \equiv 16 \sum_{j=1}^{p-1} (2j)^{p-5} \equiv 16 \cdot 2^{p-5} \sum_{j=1}^{p-1} j^{p-5} \equiv 2^{p-1} \sum_{j=1}^{p-1} j^{p-5} \equiv \sum_{j=1}^{p-1} j^{p-5} \mod p \,,
$$

where the last equality is again Fermat's little theorem. Subtracting the right-hand side of (5.8) from the left-hand side, we see that that p divides $15 \cdot \sum_{j=1}^{p-1} j^{p-5}$. If $p > 5$, this implies that p divides $\sum_{j=1}^{p-1} j^{p-5}$, and so by the previous computation we conclude $-|\mathcal{L}_{\mathrm{II}}| \equiv 0 \mod p$.

We are now ready to plug in our computations in the formula (5.6) and to conclude the proof. In the case $p = 5$ we obtain

$$
a_{\Phi,G}(\top) \equiv -|\mathcal{L}_{\mathrm{I}}| - |\mathcal{L}_{\mathrm{II}}| \equiv \frac{5+1}{2} + \frac{1}{4} \cdot 4 \equiv 4 \neq 0 \mod 5 \,.
$$

On the other hand, for $p > 5$ we saw that the contribution from Case II vanishes, and thus we have $a_{\Phi,G}(\top) = (p+1)/2 \neq 0 \mod p$. In any case we can conclude that $a_{\Phi,G}(\top)$ does not vanish modulo p, finishing the proof. ∎

COROLLARY 5.18. *Let Φ be the property of being a forest. Then for every prime $p \geq 3$ there exists a family of $\max\{4, 2(p-2)\}$-regular p-group Cayley graph expanders $\mathcal{G} = \{G_1, G_2, \ldots\}$ such that*

$$a_{\Phi,G_i}(\top) \neq 0 \mod p$$

for infinitely many positive integers i.

Proof. The case $p = 3$ was treated in section 5.2 (the property of being a forest is suitable and of unbounded wedge-number). For the case $p \geq 5$, we apply the previous proposition to the Cayley graph expanders constructed in Theorem 3.10. ∎

5.4. p-group Cayley graph expanders and bipartite graphs. So far, we have only considered minor-closed graph properties. In the current section, we show that our technique also applies to more general properties, using bipartiteness as an example.

Similarly as in the previous section, we will use the graph $\mathcal{H}(\sigma)$ to prove that the property of being bipartite has nonzero coefficients on Cayley graph expanders. The central argument necessary for the latter is given by the following proposition.

PROPOSITION 5.19. *Let Γ be a finite group of odd order, let $S_0 \subseteq \Gamma$ be a set of m generators, and let $S = \{g^{\pm 1} : g \in S_0\} \subseteq \Gamma$ be the associated symmetric set of $2m$ generators. Let $G = \mathcal{C}(\Gamma, S)$ be the associated Cayley graph, and let σ be a partition of S. Then, the graph $G\sharp\sigma$ is bipartite if and only if $\mathcal{H}(\sigma)$ is.*[23]

Proof. By Lemma 5.16 there is a well-defined graph homomorphism $\Psi : G\sharp\sigma \to \mathcal{H}(\sigma)$. Thus if $\mathcal{H}(\sigma)$ is bipartite with a partition $V(\mathcal{H}(\sigma)) = L \sqcup R$ of the vertices, then the partition $V(G\sharp\sigma) = \Psi^{-1}(L) \sqcup \Psi^{-1}(R)$ shows that $G\sharp\sigma$ is bipartite.

For the converse direction, assume that $\mathcal{H}(\sigma)$ is not bipartite. This means that there is a cycle

$$w^{B_0}, w^{B_1}, \ldots, w^{B_\ell} = w^{B_0} \in V(\mathcal{H}(\sigma))$$

in $\mathcal{H}(\sigma)$ of odd length ℓ, which can be specified by a starting vertex $w^{B_0} \in V(\mathcal{H}(\sigma))$ together with a choice of ℓ elements $g_1, \ldots, g_\ell \in S$ of the generating set (corresponding to the oriented edges that our cycle follows). Choose any $v \in V(G\sharp\sigma) = \Gamma$; then we can lift the cycle above to a walk

$$v^{B_0}, (v \cdot g_1)^{B_1}, \ldots, (v \cdot g_1 \cdots g_\ell)^{B_0} \in V(G\sharp\sigma)$$

in the graph $G\sharp\sigma$. Let $o \geq 1$ be the order of the element $g_1 \cdots g_\ell$ in the group Γ, which must be an odd number since it divides the order of the group which is assumed to be odd. Then we can repeat the lifting procedure above o times (always starting the lift at the endpoint of the previous walk) to obtain a walk in $G\sharp\sigma$ of length $o \cdot \ell$ which is odd. The endpoint of this walk is the vertex

$$(v \cdot (g_1 \cdots g_\ell)^o)^{B_0} = v^{B_0}$$

that we started at, so indeed we found a walk of odd length (which must contain a cycle of odd length), and so the graph $G\sharp\sigma$ is not bipartite. ∎

[23]The property of being bipartite extends naturally to graphs with loops and multiedges: a graph H is bipartite if there is a partition $V(H) = L \sqcup R$ of the vertices into two disjoint sets such that for an edge connecting vertices v_1, v_2, one of them is in L and one of them is in R. This is easily seen to be equivalent to the property that the graph has no odd cycle.

PROPOSITION 5.20. *Let Φ be the property of being bipartite, and let $p \geq 3$ be a prime such that there exists a family $\mathcal{G} = \{G_1, G_2, \ldots\}$ of $(n_i, 6, c)$-expanders for some positive c such that the G_i are Cayley graphs for some p-groups Γ_i. Then*

$$a_{\Phi,G_i}(\top) \neq 0 \mod p$$

for all positive integers i.

Proof. Since the graph G_i is a 6-regular Cayley graph for Γ_i, there is a set $S_0 \subseteq \Gamma_i$ of three generators such that $G_i = \mathcal{C}(\Gamma, S)$ for $S = \{g^{\pm 1} : g \in S_0\}$.

As before, when evaluating $a_{\Phi,G_i}(\top)$ modulo p, we can reduce to the fractures $\sigma \in \mathcal{L}(\Phi, G_i)^{\Gamma}$ invariant under the action of Γ. Such a σ can be interpreted as a partition of the set S. Rewriting the above formula we have

$$a_{\Phi,G_i}(\top) \equiv \sum_{\sigma \in \mathcal{L}(\Phi,G_i)^{\Gamma_i}} \Big((-1)^{|\sigma|-1} \cdot (|\sigma|-1)!\Big)^{|V(G)|} \mod p\,.$$

Since $|V(G)| = |\Gamma|$ is a power of p by the assumption that Γ is a p-group, by Fermat's little theorem we have $u^{|V(G)|} \equiv u \mod p$ for all integers u so that we can remove the exponent $|V(G)|$ in the formula above:

$$a_{\Phi,G_i}(\top) \equiv \sum_{\sigma \in \mathcal{L}(\Phi,G_i)^{\Gamma_i}} (-1)^{|\sigma|-1} \cdot (|\sigma|-1)! \mod p\,. \tag{5.9}$$

By Proposition 5.19, the graph $G_i\sharp\sigma$ is bipartite if and only if $\mathcal{H}(\sigma)$ is. In Table 1 we list all possible isomorphism classes of bipartite graphs $\mathcal{H}(\sigma)$ as σ varies through the partitions of the set S of size 6—here we use that our graphs are 6-regular. Summing all the contributions to (5.9) we see $a_{\Phi,G_i}(\top) = -16 \neq 0 \mod p$. ∎

Note that, for $p = 5$, a family of 6-regular expanders required in the previous proposition exists by Theorem 3.10. In combination with Lemma 4.8, those expanders will thus ultimately allow us to prove hardness of counting bipartite k-edge subgraphs.

Remark 5.21. Instead of considering 6-regular expanders as above, we could consider more generally $(2m)$-regular expanders for some $m \geq 2$ and use the same method as in the proof of Proposition 5.20 to compute $a_{\Phi,G_i}(\top)$ modulo p. The results for the first few m, computed using the software SageMath [86], are as follows:

m	2	3	4	5	6
$a_{\Phi,G_i}(\top) \mod p$	0	-16	192	-16576	1109760

From this we see two things: Firstly, 4-regular expanders (such as used in section 5.1.2) cannot be used to show hardness since for $m = 2$ the value of $a_{\Phi,G_i}(\top)$ vanishes modulo p for all p. Secondly, for $p = 2$ the number $a_{\Phi,G_i}(\top)$ vanishes for all m that we checked. If it vanishes for all m, then the question arises whether the problem of counting bipartite k-edge subgraphs modulo 2 might actually be fixed-parameter tractable or at least allow for a significant improvement over the brute-force algorithm.

6. Hardness of exact counting. Building upon our analysis of the coefficient function of the torus and the Cayley graph expanders above, we are now able to present the proofs of our results on #EDGESUB(Φ) with respect to exact counting.

TABLE 1

List of bipartite graphs $\mathcal{H}(\sigma)$ for $m = 3$ 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 $a(\Phi, G_i) \mod p$. The number of possible σ can be computed by enumeration or via a variant ofLemma 5.14 *where the graph is allowed to be a multigraph with loops (using the correct notion of the automorphism group of such a graph).*

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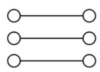	1	6	$-120 \cdot 1$
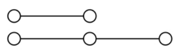	12	5	$24 \cdot 12$
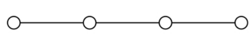	24	4	$-6 \cdot 24$
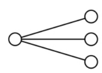	8	4	$-6 \cdot 8$
	6	4	$-6 \cdot 6$
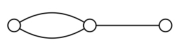	24	3	$2 \cdot 24$
	4	2	$-1 \cdot 4$
Total contribution			-16

6.1. Minor-closed properties. We begin with the classification for minor-closed graph properties. Our goal is to show that #EDGESUB(Φ) is hard whenever Φ is minor-closed, not trivially true, and of unbounded matching number. Our proof relies on Lemma 4.8 and the analysis of the coefficient function in the previous section. We need to treat the following cases separately; recall that $\mathcal{F}(\Phi)$ denotes the set of (minimal) forbidden minors of Φ.

I. Each graph in $\mathcal{F}(\Phi)$ has degree at least 3.
II. Φ has bounded wedge-number.
III. Φ is suitable and has unbounded wedge-number.
IV. Φ is unsuitable, and $\mathcal{F}(\Phi)$ contains a graph of degree at most 2.

Observe that the cases are exhaustive.

6.1.1. Case (I): Forbidden minors of degree at least 3.

LEMMA 6.1. *Let Φ denote a minor-closed graph property which is not trivially true, and assume that $\mathcal{F}(\Phi)$ does not contain a graph of degree at most* 2. *Then the problem* #EDGESUB(Φ) *is* #W[1]*-hard and, assuming ETH, cannot be solved in time*

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

for any function f.

Proof. Let $\mathcal{G} = \{G_1, G_2, \ldots\}$ denote the family of Cayley graph expanders given in subsection 5.1.2 We show that $\mathcal{H}[\Phi, \mathcal{G}]$ is infinite—this proved the claim by Lemma 4.8. Since Φ is not trivially true, the set $\mathcal{F}$ is nonempty. Thus let F denote an arbitrary graph in $\mathcal{F}$. By Fact 2.3, there is an index j such that for all $i \geq j$, the graph G_i contains the complete graph on $\#V(F)$ vertices (and thus also F) as a minor. In other words, $\Phi(G_i) = 0$ for all $i \geq j$. By Lemma 5.7, we have that for all $i \geq 2$

$$a_{\Phi,G_i}(\top) = \Phi(a_0 \cdot C_{b_0} + a_1 \cdot C_{b_1}) + \Phi(a_2 \cdot C_{2b_2}) + \Phi(a_3 \cdot C_{2b_3}) + \Phi(G_i) \mod 2\,.$$

Hence, for $i \geq \max\{2, j\}$, we have

$$a_{\Phi,G_i}(\top) = \Phi(a_0 \cdot C_{b_0} + a_1 \cdot C_{b_1}) + \Phi(a_2 \cdot C_{2b_2}) + \Phi(a_3 \cdot C_{2b_3}) \mod 2\,.$$

Finally, we rely on the premise of the lemma, implying that each graph in $\mathcal{F}$ has a vertex of degree at least 3. Consequently, no graph in $\mathcal{F}$ can be a minor of a cycle packing. Thus $\Phi(a_0 \cdot C_{b_0} + a_1 \cdot C_{b_1}) = \Phi(a_2 \cdot C_{2b_2}) = \Phi(a_3 \cdot C_{2b_3}) = 1$, and, consequently, $a_{\Phi,G_i}(\top) = 1 \mod 2$ for each $i \geq \max\{2, j\}$. ∎

6.1.2. Case (II): Properties of bounded wedge-number. For properties of bounded wedge-number (but unbounded matching-number), we do not rely on complexity monotonicity but instead on a more classical reduction from counting (edge-colorful) k-matchings. Furthermore, we prove hardness for the colored version, which will be required in an intermediate step of Case (IV).

Recall that a *wedge*, denoted by P_2, is a path with two edges, and a *k-wedge packing*, denoted by kP_2, is the disjoint union of k wedges. Recall further that a graph property Φ has *bounded wedge-number* if there exists a constant d such that $\Phi(kP_2) = 0$ for all $k \geq d$.

In what follows, we write #ColMatch for the problem of counting edge-colorful k-matchings; that is, on input of a graph G with k edge-colors, the goal is to compute the number of k-matchings in G that contain each color precisely once. The problem #ColMatch is known to be hard [26, Theorem III.1].[24] Observe that #ColMatch is equivalent to the problem #ColSubSub(Ψ) for the property Ψ of excluding P_2 as a minor.

Lemma 6.2. *Let Φ be a minor-closed graph property. If Φ has unbounded matching number but bounded wedge-number, then #ColSubSub(Φ) is #W[1]-hard and, assuming ETH, cannot be solved in time*

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

for any function f.

Proof. We construct a parameterized Turing reduction from #ColMatch to #ColSubSub(Φ) such that, additionally, on input G with k edge-colors, every oracle query (G', k') satisfies $|G'| \in O(|G|)$ and $k' \in O(k)$. Both #W[1]-hardness and the conditional lower bound then transfer from #ColMatch (see [24, section 5.2] and [26, Theorem III.1]).

Since Φ has bounded wedge-number, we have $\Phi(dP_2) = 0$ for some nonnegative integer d. Let s be the minimum nonnegative integer such that

$$\exists b \geq 0 : \Phi(sP_2 + M_b) = 0\,.$$

[24]To be precise, [26, Theorem III.1] establishes hardness if G has *at least* k edge-colors. The case of G having *precisely* k edge-colors is shown to be hard in [24, section 5.2].

We have that $s \leq d$ (for $b=0$), and observe further that $s>0$, as Φ has unbounded matching number and is minor-closed. Let furthermore $b(s)$ be the minimum over all b such that $\Phi(sP_2+M_b)=0$. This enables us to construct the reduction as follows:

Let G be the input of #ColMatch, that is, G has k edge-colors, and the goal is to compute the number of edge-colorful k-matchings in G.

We set $G' := G+(s-1)P_2+M_{b(s)}$ and color $(s-1)P_2+M_{b(s)}$ with $2(s-1)+b(s) = \#E((s-1)P_2+M_{b(s)})$ fresh colors. Setting $k' := k+2(s-1)+b(s)$ we observe that every k'-edge-colorful set A' of edges of G' decomposes into

$$A' = A\dot{\cup}E((s-1)P_2+M_{b(s)})\,,$$

where A is a k-edge-colorful set of edges in G. We claim that the number of edge-colorful k-matchings in G is equal to $\#\mathsf{ColEdgeSub}(\Phi, k' \to G')$.

To verify the latter claim, we prove that $\Phi(G'[A']) = 1$ if and only if A is a matching: If A is a matching, then $\Phi(G'[A']) = \Phi((s-1)P_2 + M_{k+b(s)}) = 1$, by our choice of s. If A is not a matching, then $G[A]$ contains a single wedge P_2 as a subgraph. Consequently $G'[A']$ contains $sP_2 + M_{b(s)}$ as a minor. Then $\Phi(G[A']) = 0$ by our choice of $b(s)$.

Our reduction thus computes G' and returns $\#\mathsf{ColEdgeSub}(\Phi, k' \to G')$ by querying the oracle. Since s, and thus $b(s)$ are fixed and independent of the input, we additionally obtain the desired conditions on the size of G' and k', which concludes the proof. ∎

6.1.3. Case (III): Suitable properties of unbounded wedge-number. For minor-closed graph properties that are suitable and of unbounded wedge-number, we will use our 4-regular 3-group Cayley graph expanders, the coefficients of which have been analyzed in Lemma 5.9. Similarly as in Case (II), we will establish hardness for the colored version.

Lemma 6.3. *Let Φ be a minor-closed graph property. If Φ is suitable and of unbounded wedge-number, then #ColSubSub(Φ) is #$\mathsf{W}[1]$-hard and, assuming ETH, cannot be solved in time*

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

for any function f.

Proof. The proof follows immediately by Lemma 5.9 and Lemma 4.8. ∎

6.1.4. Case (IV): Unsuitable properties. In what follows, we write $\mathcal{F}_2(\Phi)$ for the set of forbidden minors of Φ that have degree at most 2. We show that an unsuitable minor-closed property Φ can be reduced *from* a suitable property Ψ whenever $\mathcal{F}_2(\Phi)$ is nonempty.

Lemma 6.4. *Let Φ be a minor-closed graph property. If $\mathcal{F}_2(\Phi) \neq \emptyset$ but Φ is unsuitable, then there exists a graph property Ψ such that*

1. *Ψ is minor-closed and of unbounded matching number,*
2. *Ψ is suitable, and*
3. *#ColSubSub(Ψ)$\leq^{\mathrm{fpt}}_{\mathrm{T}}$#ColSubSub($\Phi$), and, on input G and k, every oracle query (G',k') of the reduction satisfies $|G'| \in O(|G|)$ and $k' \in O(k)$.*

Proof. Since Φ is not suitable, it is true on arbitrarily large cycles C_k. Since any finite union of paths is a minor of a sufficiently large cycle, any such union also satisfies Φ. Now any graph $F \in \mathcal{F}_2$ is a union of circles and paths,[25] and by the argument above, it must contain at least one circle. Using again that Φ is not suitable, we infer that each graph in $\mathcal{F}_2(\Phi)$ is the union of a cycle and a nonempty path-cycle packing. Since $\mathcal{F}_2(\Phi) \neq \emptyset$, we can choose $F \in \mathcal{F}_2(\Phi)$ such that the number of cycles in F is minimal among all graphs in $\mathcal{F}_2(\Phi)$. Thus $F = C + R$ for a cycle C and a (nonempty) path-cycle packing R. We define

$$\Psi(H) = 1 :\Leftrightarrow \Phi(H + R) = 1\,.$$

Let us now prove items 1–3:

1. First assume that $\Psi(H) = 1$ and $H' \prec H$; thus $H' + R \prec H + R$. Then

   $$\Psi(H) = 1 \Rightarrow \Phi(H + R) = 1 \Rightarrow \Phi(H' + R) = 1 \Rightarrow \Psi(H') = 1\,.$$

   The second implication holds as Φ is minor-closed. Consequently, Ψ is minor-closed as well.
   Next assume for contradiction that Ψ has bounded matching number. Then there exists a positive integer d such that $\Psi(M_d) = 0$ which implies that $\Phi(M_d + R) = 0$. Since $M_d + R$ is of degree 2, and every minor of a graph of degree 2 has degree at most 2, one of the graphs in $\mathcal{F}_2(\Phi)$ must be a minor of $M_d + R$. Observe that every minor of $M_d + R$ contains strictly fewer cycles than $F(= C + R)$. However, we choose F in such a way that the number of its cycles is minimal among all graphs in $\mathcal{F}_2(\Phi)$, which yields the desired contradiction.
2. Since $0 = \Phi(F) = \Phi(C + R)$ we immediately obtain that $\Psi(C) = 0$. Hence Ψ is suitable.
3. Let $r = \#E(R)$. The reduction is straightforward: Given G with k edge-colors for which we wish to compute $\#\mathsf{ColEdgeSub}(\Psi, k \to G)$, we construct G' as follows: We add a disjoint copy of R to G and color the edges of R arbitrarily with r fresh colors, yielding a $k + r$-edge-colored graph G' of size $|R| + |G| \in O(|G|)$—recall that $|R|$ is a constant. Then every $k + r$-edge-colorful subset of edges in G' consists precisely of all edges of R and a k-edge-colorful subset of edges in G. By definition of Ψ, we immediately obtain that

   $$\#\mathsf{ColEdgeSub}(\Psi, k \to G) = \#\mathsf{ColEdgeSub}(\Phi, k + r \to G')\,,$$

   which completes the reduction.

With all cases verified, the proof is concluded. ∎

6.1.5. Proof of the classification.

THEOREM 6.5. *Let Φ be a minor-closed graph property. If Φ is trivially true or of bounded matching number, then $\#\textsc{EdgeSub}(\Phi)$ is fixed-parameter tractable. Otherwise, $\#\textsc{EdgeSub}(\Phi)$ is $\#\mathsf{W}[1]$-hard and, assuming ETH, cannot be solved in time*

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

for any function f.

[25] We see an isolated vertex as a path of length 0.

Proof. The (fixed-parameter) tractability part of the classification was shown in Proposition 1.3.

Thus assume that Φ is not trivially true and of unbounded matching number. Then $\mathcal{F}(\Phi)$ is nonempty. If $\mathcal{F}_2(\Phi)$ is empty, then the result follows from Lemma 6.1. Hence consider the case $\mathcal{F}_2(\Phi) \neq \emptyset$.

We will show that the colorful version $\#\text{ColSubSub}(\Phi)$ satisfies the desired lower bound: a (tight) reduction to the uncolored version $\#\text{EdgeSub}(\Phi)$ given by Lemma 4.7.

If Φ has bounded wedge-number, then we apply Lemma 6.2. If Φ is of unbounded wedge-number and suitable, then we apply Lemma 6.3. If Φ is unsuitable, then we apply Lemma 6.4 which yields a (tight) reduction

$$\#\text{ColSubSub}(\Psi) \leq_{\mathrm{T}}^{\mathrm{fpt}} \#\text{ColSubSub}(\Phi)$$

for a minor-closed and suitable property Ψ of unbounded matching number. Now depending on whether Ψ has bounded wedge-number we, again, obtain hardness either by Lemma 6.2 or by Lemma 6.3. ∎

6.2. Selected natural properties and criteria for hardness. In this section, we will first present a concise criterion for graph properties Φ which, if satisfied, immediately yields $\#\mathsf{W}[1]$-hardness of $\#\text{EdgeSub}(\Phi)$. Afterwards, we present hardness results for a selected set of properties listed in Main Theorem 2. In what follows, we write ℓH for the graph consisting of ℓ disjoint copies of H.

We begin with a criterion that relies on the coefficient of the torus which was stated as Theorem 1.4 and which we restate for convenience.

THEOREM 1.4. *Let Φ denote a computable graph property, and assume that infinitely many primes ℓ satisfy the equation*[26]

$$\begin{aligned}(1.3)\quad & -6\Phi(M_{2\ell^2}) + 4\Phi(M_{\ell^2} + \ell C_\ell) + 8\Phi(\ell^2 P_2) - \Phi(2\ell C_\ell) - 2\Phi(\ell C_{2\ell}) \\ & \qquad - 4\Phi(\ell S_\ell) + \Phi(\circledcirc_\ell) \neq 0 \mod \ell.\end{aligned}$$

Then $\#\text{EdgeSub}(\Phi)$ is $\#\mathsf{W}[1]$-hard.

Proof. By Lemma 5.6, we have that $\mathcal{H}[\Phi, \circledcirc]$ is infinite for Φ that satisfy (1.3). $\#\mathsf{W}[1]$-hardness thus follows by Lemma 4.8. ∎

While it is easier to apply, recall that the torus cannot yield (almost) tight conditional lower bounds. If such fine-grained lower bounds are sought, it is necessary to choose a proper family of Cayley graph expanders as demonstrated by the examples given in Main Theorem 2, which we restate for convenience.

MAIN THEOREM 2. *Consider the following graph properties.*

- $\Phi_C(H) = 1$ *if and only if H is connected.*
- $\Phi_B(H) = 1$ *if and only if H is bipartite.*
- $\Phi_H(H) = 1$ *if and only if H is Hamiltonian.*
- $\Phi_E(H) = 1$ *if and only if H is Eulerian.*
- $\Phi_{CF}(H) = 1$ *if and only if H is claw-free.*

For each $\Phi \in \{\Phi_C, \Phi_B, \Phi_H, \Phi_E, \Phi_{CF}\}$, the problem $\#\text{EdgeSub}(\Phi)$ is $\#\mathsf{W}[1]$-hard. Further, unless ETH fails, $\#\text{EdgeSub}(\Phi)$ cannot be solved in time $f(k) \cdot |G|^{o(k/\log k)}$ for any function f.

[26]We write $+$ for (disjoint) graph union and ℓH for the graph consisting of ℓ disjoint copies of H. Further, we set $\Phi(H) = 1$ if H satisfies Φ and $\Phi(H) = 0$ otherwise.

Proof. We begin with the properties Φ_C, Φ_H, Φ_E, and Φ_{CF}; the case of bipartiteness is proved separately further below. Our proof proceeds by applying Lemma 5.7 to show $a_{\Phi,G_i}(\top) \neq 0$ for each of the properties Φ, allowing us to conclude using Lemma 4.8.

For $\Phi \in \{\Phi_C, \Phi_H, \Phi_E\}$, observe that the graphs $a_0 \cdot C_{b_0} + a_1 \cdot C_{b_1}$, $a_2 \cdot C_{2b_2}$, and $a_3 \cdot C_{2b_3}$ are each disconnected (and hence not Hamiltonian, nor Eulerian either) if $a_i \geq 8$ for $i = 1, 2, 3$. Further, the graphs G_i are connected since Cayley graphs are connected.[27] Thus, the graphs G_i are also Eulerian, since they are 4-regular. Moreover, Cayley graphs of p-groups are Hamiltonian [92]. Thus, by Lemma 5.7, we have that, for $i \geq 2$,

$$a_{\Phi,G_i}(\top) = \Phi(a_0 \cdot C_{b_0} + a_1 \cdot C_{b_1}) + \Phi(a_2 \cdot C_{2b_2}) + \Phi(a_3 \cdot C_{2b_3}) + \Phi(G_i) = 1 \mod 2\,.$$

Consequently, $\mathcal{H}[\Phi, \mathcal{G}]$ is infinite if $\Phi \in \{\Phi_C, \Phi_H, \Phi_E\}$. By Lemma 4.8, we obtain both $\#\mathsf{W}[1]$-hardness and the conditional lower bound.

For $\Phi = \Phi_{CF}$ we can perform a similar analysis: observe that cycle packings are always claw-free. On the other hand, for each $i > 2$, the graphs G_i do contain an (induced) claw. To see this, let $e_{\mathcal{K}_i}$ denote the neutral element of $\mathcal{K}_i$, and consider the vertices of G_i associated to $e_{\mathcal{K}_i}$, v_0, v_1, and v_1^{-1}. While $e_{\mathcal{K}_i}$ is adjacent to the remaining three cosets, it is easy to check by hand that v_0, v_1, and v_1^{-1} constitute an independent set in G_i.

Consequently, by Lemma 5.7, we have that, for $i > 2$,

$$\begin{aligned} a_{\Phi,G_i}(\top) = {} & \Phi(a_0 \cdot C_{b_0} + a_1 \cdot C_{b_1}) + \Phi(a_2 \cdot C_{2b_2}) \\ & + \Phi(a_3 \cdot C_{2b_3}) + \Phi(G_i) = 3 + 0 = 1 \mod 2\,. \end{aligned}$$

Thus, $\mathcal{H}[\Phi_{CF}, \mathcal{G}]$ is infinite. By Lemma 4.8, we hence obtain both $\#\mathsf{W}[1]$-hardness and the conditional lower bound.

Finally, for the property $\Phi = \Phi_B$ of being bipartite, we invoke Proposition 5.20 and use the 6-regular 5-group Cayley graph expanders given by Theorem 3.10. Again, in combination with Lemma 4.8, this yields the desired lower bounds. ∎

6.3. A comment on modular counting. It was shown in the conference version [70] that our technique of proving lower bounds for #EdgeSub(Φ) via the coefficients of Cayley graph expanders in the homomorphism basis also applies to modular counting. More precisely, let us write $\#_p$EdgeSub(Φ) for the problem of, given G and k, computing $\#\mathsf{EdgeSub}(\Phi, k \to G) \mod p$. Observing that the reduction sequence in section 4 applies to counting modulo p as well[28] and relying on known hardness results for counting color-prescribed homomorphisms modulo p, the following criterion was established for $\#_p$EdgeSub(Φ).

Lemma 6.6 (see [70]). *Let Φ be a computable graph property, and let furthermore $\mathcal{G} = \{G_1, G_2, \dots\}$ be a family of (n_i, d, c)-expanders for some positive integers c and d. If for infinitely many i we have*

$$a_{\Phi,G_i}(\top) \neq 0 \mod p\,,$$

[27]Recall that our definition of Cayley graphs enforces the set S to be a set of generators of the group.

[28]The matrix in Lemma 4.5 is triangular with 1s on the diagonal and thus nonsingular over $\mathsf{GF}[p]$ for any prime p.

then $\#_p$EdgeSub(Φ) *is* $\mathsf{Mod}_p\mathsf{W}[1]$*-hard*[29] *and, assuming the randomized ETH, cannot be solved in time* $f(k)\cdot|G|^{o(k/\log k)}$ *for any function* f*.*

Due to the length of this paper, we decided to not include an introduction to parameterized modular counting. However, we wish to emphasize one particular consequence of our treatment of forests in subsection 5.3 on the complexity of modular counting: Let us write $\#_p$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; the parameterization is given by k. Similarly, we write $\#_p$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 parameterization is given by the rank k of M.

Theorem 6.7. *Given a prime number* $p\geq 3$*, the problems* $\#_p$Forests *and* $\#_p$Bases *are* $\mathsf{Mod}_p\mathsf{W}[1]$*-hard and, assuming the randomized ETH, cannot be solved in time*

$$f(k)\cdot|G|^{o(k/\log k)}\quad (\textit{resp., } f(k)\cdot|M|^{o(k/\log k)})$$

for any function f*.*

Proof. We prove the lower bound for $\#_p$Forests; a parsimonious reduction to $\#_p$Bases follows easily by (deterministic) polynomial-time matroid truncation: Given a graph G, the number of k-forests in G is equal to the number of k-independent sets in the graphic matroid $M(G)$ (which is a linear matroid). A k-truncation of $M(G)$ is a matroid M of rank k whose bases are in one-to-one correspondence to the k-independent sets of $M(G)$ and, thus, to the k-forests in G. We refer the reader to [59] for a detailed exposition of k-truncations and, in particular, for a deterministic algorithm which, on input of a linear matroid and a positive integer k, computes a (linear) k-truncation of the matroid in polynomial time.

Now observe that $\#_p$Forests is equal to the problem $\#_p$EdgeSub(Φ) for Φ being the property of being a forest. Hence we can apply Lemma 6.6 to the Cayley graph expanders given by Corollary 5.18—the latter satisfy that the coefficients are $a_{\Phi,G_i}(\top)$ are nonzero modulo p infinitely often. ∎

We leave it as an open question whether or not the case of $p=2$, that is, counting k-forests modulo 2, is fixed-parameter tractable.

7. Approximate counting of small subgraph patterns. Recall that we identified #EdgeSub(Φ) as an inherently hard problem in the case when we aim for *exactly* counting the solutions. In particular, we established $\#\mathsf{W}[1]$-hardness for any nontrivial minor-closed property Φ of unbounded matching number. For this reason, the section below deals with the complexity of *approximating* the number of solutions. Tractability of approximating the solutions of parameterized counting problems is given by the notion of an FPTRAS.

Definition 7.1 (FPTRAS [6, 65]). *Let* (P,κ) *denote a parameterized counting problem. A fixed-parameter tractable randomized approximation scheme (FPTRAS) for* (P,κ) *is a randomized algorithm* $\mathbb{A}$ *that, given* $x\in\Sigma^*$ *and rational numbers* $\varepsilon>0$ *and* $0<\delta<1$*, computes an integer* z *such that*

$$\mathsf{Pr}[\,(1-\varepsilon)P(x)\leq z\leq(1+\varepsilon)P(x)\,]\geq 1-\delta\,.$$

[29] $\mathsf{Mod}_p\mathsf{W}[1]$-hard problems are at least as hard as counting k-cliques modulo p with respect to parameterized Turing-reductions.

The running time of $\mathbb{A}$ *must be bounded by* $f(\kappa(x)) \cdot \mathsf{poly}(|x|, 1/\varepsilon, \log(1/\delta))$ *for some computable function* f.

Indeed, we can show that #EdgeSub(Φ) allows an FPTRAS for every minor-closed property Φ. In fact, we prove the following general criterion, which implies the existence of an FPTRAS for minor-closed properties.

Main Theorem 4. *Let* Φ *denote a computable graph property. If* Φ *satisfies the matching criterion* and *the star criterion, or if* Φ *has bounded treewidth, then* #EdgeSub(Φ) *admits an FPTRAS.*

We start with the case of Φ satisfying both the matching and the star criteria. For readers familiar with the meta-theorem of Dell, Lapinskas, and Meeks [30], we point out that their method cannot be used to achieve the desired goal in the current setting: the results in [30, section 1.3] imply that #EdgeSub(Φ) admits an FPTRAS whenever the *edge-colorful decision version* of EdgeSub(Φ) is fixed-parameter tractable; in the latter, we expect as input a graph G with k different edge-colors, and the goal is to decide whether there is a subset A of edges containing each color exactly once such that $G[A]$ satisfies Φ (w.r.t. the underlying uncolored graph). Thus, if we could show that the edge-colorful decision version is fixed-parameter tractable for properties satisfying the matching and the star criteria, Main Theorem 4 would follow.

However, the latter cannot be true (unless $\mathsf{FPT} = \mathsf{W}[1]$) since the following property Φ induces a $\mathsf{W}[1]$-hard colorful decision version while satisfying both the matching and the star criteria: $\Phi(H) = 1$ if and only if H is either a star, a matching, or the union of a clique and a triangle. $\mathsf{W}[1]$-hardness follows from a reduction from finding edge-colorful k-cliques in a graph, which is known to be $\mathsf{W}[1]$-hard.[30] The reduction is straightforward: given a graph G with $\binom{k}{2}$ edge colors, we construct a graph G' by adding a triangle with three fresh colors to the graph. Then G' contains a colorful $\binom{k}{2} + 3$-edge-subset A that satisfies Φ if and only if G contains an edge-colorful k-clique. The latter is true since any colorful $\binom{k}{2} + 3$-edge-subset must contain the triangle with the three fresh colors and can thus induce neither a star nor a matching.

Being unable to rely on the colorful decision version, we thus use a different approach using Ramsey's theorem, similar to the one in [65]. More precisely, we use the following consequence.

Lemma 7.2. *Let* $k \geq 4$ *denote a positive integer, and let* G *denote a graph with at least* $R(k,k)$ *edges. Then* G *contains either* $K_{1,k}$ *or* M_k *as a subgraph.*

Proof. We apply Ramsey's theorem to the line graph $L(G)$ of G: The vertices of $L(G)$ are the edges of G, and two vertices e and e' of $L(G)$ are adjacent if and only if $e \cap e' \neq 0$. Since $L(G)$ contains at least $R(k,k)$ vertices, Ramsey's theorem implies that $L(G)$ either contains an independent set or a clique of size k. Note that a k-independent set of $L(G)$ corresponds to a k-matching in G and that a k-clique in $L(G)$ corresponds to a star $K_{1,k}$ in G; the latter requires that $k \geq 4$ since the line graph of a triangle is a triangle (and thus a clique) as well. ∎

The subsequent observation enables our Monte Carlo algorithm to only rely on "fixed-parameter tractable many" samples.

[30] The *vertex-colorful* clique problem is $\mathsf{W}[1]$-hard (see Chapter 13 in [27]) and reduces to the *edge-colorful* clique problem by assigning an edge $\{u, v\}$ the color $\{c(u), c(v)\}$, where $c(u)$ and $c(v)$ are the vertex-colors of u and v.

LEMMA 7.3. *Let $k \geq 4$ denote a positive integer, and let G denote a graph with at least $R(k,k)$ edges. Assume a subset A of k edges is sampled uniformly at random. We have*

$$\mathsf{Pr}[\, G[A] \cong M_k \vee G[A] \cong K_{1,k}\,] \geq \binom{R(k,k)}{k}^{-1}.$$

Proof. Set $m = |E(G)|$ and $r = R(k,k)$. It is convenient to assume that A is sampled as follows: we first choose r edges uniformly at random, denote this set by S, and afterwards obtain A by choosing k edges among S uniformly at random; of course, we need to show that this yields a uniform distribution. Let B denote any k-edge subset of G. By the law of total probability, we have that

$$\begin{aligned}\mathsf{Pr}[\, A = B\,] &= \sum_{T \in \binom{E(G)}{r}} \mathsf{Pr}[\, S = T\,] \cdot \mathsf{Pr}[\, A = B \mid S = T\,] \\ &= \sum_{T \in \binom{E(G)}{r}} \binom{m}{r}^{-1} \cdot \mathsf{Pr}[\, A = B \mid S = T\,].\end{aligned}$$

Note that $\mathsf{Pr}[\, A = B \mid S = T\,] = \binom{r}{k}^{-1}$ if $B \subseteq T$, and $\mathsf{Pr}[\, A = B \mid S = T\,] = 0$ otherwise. Consequently

$$\begin{aligned}\mathsf{Pr}[\, A = B\,] &= \#\{T \subseteq E(G) | B \subseteq T \wedge \#T = r\} \cdot \binom{m}{r}^{-1} \binom{r}{k}^{-1} \\ &= \binom{m-k}{r-k} \binom{m}{r}^{-1} \binom{r}{k}^{-1} = \binom{m}{k}^{-1}.\end{aligned}$$

Now let $\mathcal{E}$ denote the event $G[A] \cong M_k \vee G[A] \cong K_{1,k}$, and note that for every r-edge subset T of G we have that $\mathsf{Pr}[\, \mathcal{E} \mid S = T\,] \geq \binom{r}{k}^{-1}$ since, by the previous lemma, $G[T]$ contains either M_k or $K_{1,k}$ as a subgraph. We conclude that $\mathsf{Pr}[\, \mathcal{E}\,]$ is equal to

$$\sum_{T \in \binom{E(G)}{r}} \mathsf{Pr}[\, S = T\,] \cdot \mathsf{Pr}[\, \mathcal{E} \mid S = T\,] = \binom{m}{r} \binom{m}{r}^{-1} \cdot \mathsf{Pr}[\, \mathcal{E} \mid S = T\,] \geq \binom{r}{k}^{-1},$$

which concludes the proof. $\square$

For our FPTRAS, we use the following (consequence of a) Chernoff bound.

THEOREM 7.4 (see Theorem 11.1 in [68]). *Let $X_1, \ldots, X_t$ denote independent and identically distributed indicator random variables with expectation $\eta = E[X_i]$, and let $0 < \varepsilon, \delta < 1$ denote positive rationals. If $t \geq (3 \ln(2/\delta))/(\varepsilon^2 \eta)$, then*

$$\mathsf{Pr}\left[\left| \frac{1}{t} \cdot \sum_{i=1}^{t} X_i - \eta \right| < \varepsilon \eta \right] \geq 1 - \delta\,.$$

LEMMA 7.5. *Let Φ denote a computable graph property satisfying both the matching criterion and the star criterion. Then $\#\text{EdgeSub}(\Phi)$ has an FPTRAS.*

Proof. By assumption, there is a constant c' such that Φ is true for all matchings and stars of size at least c'; we set $c = \max(c', 4)$. Our FPTRAS $\mathbb{A}$ is constructed as follows: If $k < c$ or if $|E(G)| \leq R(k,k)$, then we solve the problem (exactly) by the naive brute-force algorithm. Otherwise, we take

Algorithm 1 An FPTRAS for $\#\textsc{EdgeSub}(\Phi)$, where Φ satisfies the matching and the star criteria.

```
MatchingsAndStarsFPTRAS(G, k, ε, δ)
    if k < c or |E(G)| ≤ R(k,k) then
        Solve the problem exactly by brute force;
    end
    else
        X ← 0; t ← binom(R(k,k), k) · 3ln(2/δ)/ε²;
        for i ← 1 to t do
            Sample a k-edge subset A of G uniformly at random;
            if Φ(G[A]) = 1 then X ← X + 1;
        end
        return X/t · binom(|E(G)|, k);
    end
end
```

$$\binom{R(k,k)}{k} \cdot \frac{3\ln(2/\delta)}{\varepsilon^2}$$

many independent samples of k-edge sets A of G, each taken uniformly at random. Finally, we output the fraction of those samples A such that $\Phi(G[A]) = 1$. Consult Algorithm 1 for a visualization as pseudocode.

Let us first argue about the running time: if $k < c$, then the brute-force algorithm takes time at most $|G|^c$,[31] and if $|E(G)| \leq R(k,k)$, then the brute-force algorithm takes time at most $|G| + R(k,k)^k$. Otherwise, we iterate through the loop t times, and each iteration can clearly be done in time $f'(k) \cdot \mathsf{poly}(|G|)$ for some computable function f'—note that the factor $f'(k)$ depends on the complexity of verifying whether $\Phi(G[A])$ holds, which might require superpolynomial time in $|G[A]| \in O(k)$. The overall running time is thus bounded by

$$\max\left\{|G|^c, |G| + R(k,k)^k, \binom{R(k,k)}{k} \cdot (3\ln(2/\delta))/\varepsilon^2 \cdot f'(k) \cdot \mathsf{poly}(|G|)\right\},$$

which is bounded by $f(k) \cdot \mathsf{poly}(|G|, 1/\varepsilon, \log(1/\delta))$ for some computable function f.

Next note that correctness is trivial in the case when the brute-force algorithm is executed. Hence assume that $k \geq c$ and $|E(G)| > R(k,k)$. To avoid notational clutter, we set $r := R(k,k)$ and $m := |E(G)|$. Now let X_i denote the indicator variable defined to be 1 if the ith sample, denoted A_i, satisfies $\Phi(G[A_i]) = 1$, and $X_i = 0$ otherwise. Observe that $E[X_i] = \#\mathsf{EdgeSub}(\Phi, k \to G) \cdot \binom{m}{k}^{-1}$ for all i. In what follows, we thus just set $\eta := E[X_i]$. Since Φ is true for M_k and $K_{1,k}$, by Lemma 7.3 we furthermore have

$$\eta = \mathsf{Pr}[\,\Phi(G[A]) = 1\,] \geq \mathsf{Pr}[\,G[A] \cong M_k \vee G[A] \cong K_{1,k}\,] \geq \binom{r}{k}^{-1}.$$

[31] $|G|$ rather than $|E(G)|$ since G might contain many isolated vertices.

Consequently, $t \geq (3\ln(2/\delta))/(\varepsilon^2\eta)$. By the previous Chernoff bound, we thus have

$$\mathsf{Pr}\left[\left|\frac{1}{t}\cdot\sum_{i=1}^{t} X_i - \eta\right| < \varepsilon\eta\right] \geq 1-\delta\,.$$

Finally, recall that $X = \sum_{i=1}^{t} X_i$, and observe that

$$\begin{aligned}
&\left|\frac{1}{t}\cdot\sum_{i=1}^{t} X_i - \eta\right| < \varepsilon\eta \\
\Rightarrow &\left|\frac{X}{t} - \frac{\#\mathsf{EdgeSub}(\Phi, k\to G)}{\binom{m}{k}}\right| < \varepsilon\cdot\frac{\#\mathsf{EdgeSub}(\Phi, k\to G)}{\binom{m}{k}} \\
\Rightarrow &\left|\frac{X}{t}\cdot\binom{m}{k} - \#\mathsf{EdgeSub}(\Phi, k\to G)\right| < \varepsilon\cdot\#\mathsf{EdgeSub}(\Phi, k\to G)\,.
\end{aligned}$$

We conclude the proof by pointing out that the latter implies

$$(1-\varepsilon)\cdot\#\mathsf{EdgeSub}(\Phi, k\to G) \leq \frac{X}{t}\cdot\binom{m}{k} \leq (1+\varepsilon)\cdot\#\mathsf{EdgeSub}(\Phi, k\to G)\,. \qquad \square$$

For the case of Φ having bounded treewidth, we rely on the following result of Arvind and Raman; to this end, given a fixed positive integer T, let $\#\textsc{Sub}(T)$ denote the problem that, on input of a graph H of treewidth at most T and an arbitrary graph G, requires us to compute $\#\mathsf{Sub}(H\to G)$.

Theorem 7.6 (see [6]). *For each positive integer T, there is an FPTRAS for $\#\textsc{Sub}(T)$ if it is parameterized by the size of the graph H.*

Lemma 7.7. *Let Φ denote a computable graph property of bounded treewidth. Then $\#\textsc{EdgeSub}(\Phi)$ admits an FPTRAS.*

Proof. By assumption, there is a constant T such that the treewidth of each graph H with $\Phi(H)=1$ is at most T. Define $g(k) := |\Phi_k|$, and observe that g is computable as Φ is.

Recall from (2.1) that for each G and k we have

$$\#\mathsf{EdgeSub}(\Phi, k\to G) = \sum_{H\in\Phi_k}\#\mathsf{Sub}(H\to G).$$

We thus just use the FPTRAS from Theorem 7.6 to approximate (with probability $1-\delta/g(k)$) each term $\#\mathsf{Sub}(H\to G)$ with $H\in\Phi_k$ and output the sum given by the previous equation.

Observe that approximating each term $\#\mathsf{Sub}(H\to G)$ takes time at most

$$f'(|H|)\cdot\mathsf{poly}(|G|, 1/\varepsilon, \log(g(k)/\delta))$$

for some computable function f'.

Since each $H\in\Phi_k$ has k edges, the overall running time is thus clearly bounded by

$$f(k)\cdot\mathsf{poly}(|G|, 1/\varepsilon, \log(\delta))$$

for some computable function f—note that f depends on Φ, f', and g, but the latter three are independent of the input. Now let r denote the output of our algorithm. It remains to show that

$$\mathsf{Pr}[\,(1-\varepsilon)\cdot\#\mathsf{EdgeSub}(\Phi, k\to G) \leq r \leq (1+\varepsilon)\cdot\#\mathsf{EdgeSub}(\Phi, k\to G)\,] \geq 1-\delta\,.$$

Write r_H for the output of the FPRAS from Theorem 7.6 on input G, H, ε, and $\delta/g(k)$. Then

$$r = \sum_{H\in\Phi_k} r_H\,,$$

and the following holds for each $H \in \Phi_k$:

$$\mathsf{Pr}[\,(1-\varepsilon)\cdot\#\mathsf{Sub}(H\to G)\le r_H \le (1+\varepsilon)\cdot\#\mathsf{Sub}(H\to G)\,]\ge 1-\delta/g(k)\,.$$

Since the outcomes r_H are independent and $g(k) = |\Phi_k|$, we have

$$\begin{aligned}&\mathsf{Pr}[\forall H\in\Phi_H : (1-\varepsilon)\cdot\#\mathsf{Sub}(H\to G)\le r_H\le(1+\varepsilon)\cdot\#\mathsf{Sub}(H\to G)\,]\\ &\ge (1-\delta/g(k))^{g(k)}\,,\end{aligned}$$

which is at most $(1 - \delta)$ by Bernoulli's inequality.

Consequently, with probability at least $(1 - \delta)$, we have that

$$\begin{aligned}(1-\varepsilon)\cdot\#\mathsf{EdgeSub}(\Phi,k\to G) &= (1-\varepsilon)\sum_{H\in\Phi_k}\#\mathsf{Sub}(H\to G)\\ &= \sum_{H\in\Phi_k}(1-\varepsilon)\cdot\#\mathsf{Sub}(H\to G)\\ &\le \sum_{H\in\Phi_k} r_H(=r)\\ &\le \sum_{H\in\Phi_k}(1+\varepsilon)\cdot\#\mathsf{Sub}(H\to G)\\ &= (1+\varepsilon)\sum_{H\in\Phi_k}\#\mathsf{Sub}(H\to G)\\ &= (1+\varepsilon)\cdot\#\mathsf{EdgeSub}(\Phi,k\to G)\,,\end{aligned}$$

which concludes the proof. ∎

Proof of Main Theorem 4. Main Theorem 4 holds by Lemmas 7.5 and 7.7. ∎

8. Detection of small subgraph patterns. In this section, we study the complexity of the decision problem EdgeSub(Φ). As a first observation we observe that EdgeSub(Φ) essentially subsumes the (parameterized) subgraph isomorphism problem: consider, for instance, the property Φ defined as $\Phi(H) = 1$ if and only if $H \cong K_{\ell,\ell}$ for some positive integer ℓ. Then EdgeSub(Φ) is equivalent to the problem k-Biclique which was only recently shown to be W[1]-hard by the seminal result of Lin [58] after being unresolved for at least a decade.

More generally, let $\mathcal{H}$ denote a class of graphs, and define Emb($\mathcal{H}$) as the problem that asks, given a graph $H \in \mathcal{H}$ and an arbitrary graph G, whether there is a subgraph embedding from H to G; the parameterization is given by $|H|$. Plehn and Voigt [72] proved Emb($\mathcal{H}$) to be fixed-parameter tractable whenever the treewidth of graphs in $\mathcal{H}$ is bounded by a constant. On the other hand, the question whether Emb($\mathcal{H}$) is W[1]-hard in all remaining cases is one of the "most infamous" [35, Chapter 33.1] open problems in parameterized complexity. Since EdgeSub(Φ) subsumes[32] Emb($\mathcal{H}$)

[32]To be precise, EDGESUB(Φ) subsumes EMB($\mathcal{H}$) whenever $\mathcal{H}$ does not contain two graphs with the same number of edges, which is, however, true for most of the natural instances of the subgraph isomorphism problem such as finding cliques, bicliques, cycles, paths, and matchings, to name only a few.

as we have seen in the case of k-BICLIQUE, a complete classification of EDGESUB(Φ) seems to be elusive at the moment.

However, we identify the following tractable instances of EDGESUB(Φ), which significantly extends the case of bounded treewidth.

MAIN THEOREM 5. *Let Φ denote a computable graph property. If Φ satisfies the matching criterion* or *the star criterion, or if Φ has bounded treewidth, then* EDGESUB(Φ) *is fixed-parameter tractable.*

In the case that Φ satisfies the matching or the star criterion, fixed-parameter tractability is obtained by a surprisingly simple win-win approach relying on the treewidth and the maximum degree of a graph. Assume, for example, that Φ is true for all matchings. Now, given a graph G and an integer k, we can easily verify whether G contains a maximum matching of size at least k. If the latter is true, G contains a subgraph with k edges that satisfies Φ. More interestingly, if the matching number of G is bounded by k, then its vertex-cover number (and thus its treewidth) is bounded by $2k$, and we can efficiently use dynamic programming over a tree-decomposition of small width of G to verify whether $\mathsf{EdgeSub}(\Phi, k \to G) \neq \emptyset$. Formally, the latter can be established by an easy application of Courcelle's theorem [22] as shown in the following lemma.

LEMMA 8.1. *Let Φ denote a computable graph property. There is a computable function g and an algorithm $\mathbb{A}$ that, given a graph G and a positive integer k, correctly decides whether* $\mathsf{EdgeSub}(\Phi, k \to G) \neq \emptyset$ *in time* $g(\mathsf{tw}(G), k) \cdot |G|$.

Proof. We use Courcelle's theorem as stated in [39, Theorem 11.37]. Thus it remains to provide a monadic second-order sentence[33] φ such that G satisfies φ if and only if $\mathsf{EdgeSub}(\Phi, k \to G) \neq \emptyset$. To this end, let $H \in \Phi_k$, and assume that $V(H) = \{1, \ldots, v_H\}$. Consider the following sentence:

$$\varphi_H := \exists x_1, \ldots, \exists x_{v_H} : \bigwedge_{i \neq j} x_i \neq x_j \wedge \bigwedge_{\{i,j\} \in E(H)} E(x_i, x_j)\,.$$

Observe that G satisfies φ_H if and only if H is a subgraph of G. Consequently, we set

$$\varphi := \bigvee_{H \in \Phi_k} \varphi_H\,.$$

Since the length of φ only depends on Φ and k, the lemma holds by Courcelle's theorem. ∎

We are now able to establish fixed-parameter tractability of EDGESUB(Φ) whenever Φ satisfies the matching criterion.

LEMMA 8.2. *Let Φ denote a computable graph property that satisfies the matching criterion. Then the problem* EDGESUB(Φ) *is fixed-parameter tractable.*

Proof. Since Φ satisfies the matching criterion, there is a constant c (only depending on Φ) such that $\Phi(M_k) = 1$ for all $k \geq c$. The fixed-parameter tractable algorithm is constructed as follows:

Given a graph G and a positive integer k, we can assume that $k \geq c$, solving the case $k < c$ by brute-force enumeration of all k-subsets of edges. In the case $k \geq c$,

[33]We refer the reader to, e.g., Chapter 4 in [39] for an introduction to monadic second-order logic.

we compute a maximum matching M of G in polynomial time by, e.g., the blossom algorithm [37]. If $|M| \geq k$, then we can output 1, since any k-subset A of M satisfies that $\Phi(G[A]) = 1$ by assumption.

In the remaining case, we can thus assume that the matching number of G is bounded by k. Consequently, the vertex cover number of G is bounded by $2k$. Since the treewidth of a graph is bounded by its vertex cover number, we conclude that $\mathsf{tw}(G) \leq 2k$. Invoking the algorithm from the previous lemma thus yields an overall running time bounded by

$$m^{O(1)} + g(2k, k) \cdot |G| ,$$

which proves fixed-parameter tractability. ∎

We continue with the case of Φ satisfying the star criterion. To this end, we require the following result, which is implicitly implied by the counting version of the Frick–Grohe theorem [41]; we provide a proof based on the bounded search-tree paradigm for completeness.

LEMMA 8.3. *Let Φ denote a computable graph property. There is a computable function g and an algorithm $\mathbb{A}$ that, given a graph G and a positive integer k, correctly decides whether* $\mathsf{EdgeSub}(\Phi, k \to G) \neq \emptyset$ *in time* $g(\mathsf{deg}(G), k) \cdot |G|$.

Proof. We check for each $H \in \Phi_k$ whether H is a subgraph of G and output 1 if (and only if) at least one of those checks is positive.

Assume for a moment that H is connected. In this case, the strategy is very simple: We guess a vertex v of G and search for a subgraph embedding of H in G that includes v. Since H is connected and has k edges, the image of the subgraph embedding can only contain vertices of distance at most k from v. This allows us to search for a copy of H in the $\leq k$ neighborhood of v by brute force, since the latter contains at most $\mathsf{deg}(G)^k$ vertices. The overall running time of finding a subgraph isomorphic to H in G is thus bounded by $|V(G)| \cdot \mathsf{deg}(G)^k$.

The situation becomes slightly more complicated if H is not connected. We would like to perform the previous strategy for each connected component of H, adding an additional factor of k in the worst case. However, since a subgraph embedding needs to be injective, we have to guarantee that we do not construct a solution that uses vertices of G twice. This issue is solved by a standard application of color-coding: We choose a function $\mathsf{col} : V(G) \to V(H)$ uniformly at random. If G contains a subgraph isomorphic to H, then with probability at least $p(k) > 0$ there is a subgraph embedding $\psi : V(H) \to V(G)$ such that additionally $\mathsf{col}(\psi(v)) = v$ for each vertex $v \in V(H)$, and such a subgraph embedding can be found in time $O(k \cdot |V(G)| \cdot \mathsf{deg}(G)^k)$ by adapting the above strategy for every connected component H accordingly. Finally, derandomization can be achieved by perfect hashing as shown in [3] (see also [39, Chapter 13.3]). ∎

Let us now establish fixed-parameter tractability of EDGESUB(Φ) whenever Φ satisfies the star criterion.

LEMMA 8.4. *Let Φ denote a computable graph property satisfying the star criterion. Then the problem* EDGESUB(Φ) *is fixed-parameter tractable.*

Proof. Since Φ satisfies the star criterion, there is a constant c (only depending on Φ) such that $\Phi(K_{1,k}) = 1$ for all $k \geq c$. The fixed-parameter tractable algorithm is constructed as follows:

Given a graph G and a positive integer k, we can again solve the case $k < c$ by brute force and thus assume $k \geq c$. Then, we check whether G contains a vertex v of degree at least k, in which case we can output 1, since any k-subset A of the incident edges of v satisfies that $\Phi(G[A]) = 1$ by assumption.

In the remaining case, we can thus assume that $\mathsf{deg}(G) \leq k$. Invoking the algorithm from the previous lemma thus yields an overall running time bounded by

$$m^{O(1)} + g(k,k) \cdot |G|\,,$$

which proves fixed-parameter tractability. ∎

Proof of Main Theorem 5. In the case that Φ satisfies the matching criterion or the star criterion, the claim holds by Lemma 8.2 and Lemma 8.4. If Φ has bounded treewidth, then, given G and k, we can use the algorithm of Plehn and Voigt [72] for each $H \in \Phi_k$. Since the size of Φ_k is bounded by a function in k, the overall running time still yields fixed-parameter tractability. ∎

Our main result regarding minor-closed properties is now obtained by the combination of our results in the realms of exact counting and approximate counting, as well as decision.

MAIN THEOREM 1. *Let Φ denote a minor-closed graph property.*

1. Exact counting*: If Φ is either trivially true or of bounded matching number, then the (exact) counting version* #EDGESUB(Φ) *is fixed-parameter tractable. Otherwise, the problem* #EDGESUB(Φ) *is* #W[1]*-hard and, assuming the ETH, cannot be solved in time $f(k) \cdot |G|^{o(k/\log k)}$ for any function f.*
2. Approximate counting*: The problem* #EDGESUB(Φ) *always has an FPTRAS.*
3. Decision: *The problem* EDGESUB(Φ) *is always fixed-parameter tractable.*

Proof. Note that each minor-closed property is computable (even in polynomial time) by the Robertson–Seymour theorem [75]. The classification of exact counting follows by Theorem 6.5. For approximate counting and decision, we claim that each minor-closed property Φ either has bounded treewidth or satisfies both the matching *and* the star criteria. If the latter holds, then the existence of an FPTRAS for approximate counting follows by Main Theorem 4, and the fixed-parameter tractable algorithm for decision follows by Main Theorem 5.

To prove the claim, we assume that Φ has unbounded treewidth; otherwise we are done. In that case, by the excluded grid theorem [74], Φ holds for a sequence of graphs containing arbitrarily large grids as minors. Since every planar graph (including matchings and stars) is a minor of a grid [76], and Φ is minor-closed, we conclude that Φ holds for all matchings and all stars and thus satisfies both the matching and the star criteria. ∎

9. A parameterized Tutte polynomial. In the last part of the paper, we take a step back and revisit exact counting: Recall that problem #EDGESUB(Φ) can be interpreted as the problem of evaluating a linear combination of subgraph counts, given by

$$\#\mathsf{EdgeSub}(\Phi, k \to *) = \sum_{H \in \Phi_k} \#\mathsf{Sub}(H \to G)\,,$$

where Φ_k is the set of all k-edge graphs that satisfy Φ. In particular, each coefficient in this linear combination is 0 or 1. We have seen that the values of Φ on the fixed points of certain group actions on (fractures of) Cayley graphs can be used to obtain explicit criteria for (#$\mathsf{W}[1]$-)hardness of #EdgeSub(Φ). In the current section, we show that the aforementioned method applies to the significantly more general problem of computing weighted linear combinations of k-edge subgraph counts. More precisely, we consider a natural parameterized variant of the Tutte polynomial and obtain an exhaustive classification for the complexity of evaluating it at any rational coordinates.

Recall that the (classical) Tutte polynomial is defined as follows:

$$T_G(x,y) := \sum_{A\subseteq E(G)} (x-1)^{k(A)-k(E(G))} \cdot (y-1)^{k(A)+\#A-\#V(G)}\,,$$

where $k(S)$ is the number of connected components of the graph $(V(G),S)$.
In this work, we consider the specialization of the Tutte polynomial to edge-subsets of size k, which we call the *parameterized Tutte polynomial*:

$$T_G^k(x,y) := \sum_{A\in\binom{E(G)}{k}} (x-1)^{k(A)-k(E(G))} \cdot (y-1)^{k(A)+k-\#V(G)}\,.$$

We emphasize that the parameterized Tutte polynomial is related to a generalization of the bases generating function for matroids investigated by Anari et al. in their work on approximate counting (and sampling) via log-concave generating polynomials [5, section 1.2].

Similarly to the classical counterpart due to Jaeger, Vertigan, and Welsh [49], our goal is to understand the parameterized complexity of evaluating $T_G^k(x,y)$ for any fixed pair of coordinates (x,y) when parameterized by k. Note that at points (x,y) with $x\neq 1, y\neq 1$ we can write the polynomial as

$$T_G^k(x,y) = (x-1)^{-k(E(G))}(y-1)^{k-\#V(G)} \sum_{A\in\binom{E(G)}{k}} ((x-1)\cdot(y-1))^{k(A)}\,.$$

So, up to the global factor $(x-1)^{-k(E(G))}(y-1)^{k-\#V(G)}$ (which can be computed in linear time in the input size) in this region the polynomial is really just a polynomial in the single variable $z=(x-1)(y-1)$. Still, we keep the variables x,y separate in the treatment below. On the one hand, this facilitates comparisons to the classical Tutte polynomial. On the other hand, we see some interesting behavior of $T_G^k(x,y)$ at points with $x=1$ or $y=1$. Indeed, let us start by investigating the expressibility of the parameterized Tutte polynomial in some individual points.

9.1. Interpretation in individual points. Recall that, given a graph G and a subset $A\subseteq E(G)$ of its edges, we write $G(A)=(V(G),A)$ for the graph induced by A. We emphasize the difference from the construction $G[A]$ we saw before: the graph $G[A]$ is obtained from $G(A)$ by removing all isolated vertices.

The most immediate information encoded in the parameterized Tutte polynomial is the number of k-forests in a graph.

Observation 9.1. The number of forests with k edges in a graph G is given by $T_G^k(2,1)$.

In particular, evaluating $T_G^k(2,1)$ is equivalent to evaluation $\#\mathsf{IndSub}(\Phi,k\to G)$ for the (minor-closed) property of being acyclic.

For further individual points, it is convenient to consider the following modification.

DEFINITION 9.2. *Define the modified Tutte polynomial of a graph G as*

$$\widetilde{T}_G(x,y) := \sum_{A \subseteq E(G)} (x-1)^{k(A)} \cdot (y-1)^{k(A)+\#A},$$

so that $\widetilde{T}_G(x,y) = (x-1)^{k(E(G))}(y-1)^{\#V(G)} T_G(x,y)$. *Similarly we define the parameterized version as*

$$\widetilde{T}^k_G(x,y) := \sum_{A \in \binom{E(G)}{k}} (x-1)^{k(A)} \cdot (y-1)^{k(A)+\#A}.$$

As for its classical counterpart, we observe a deletion-contraction recurrence, which enables us establish the properties at individual points. Setting $\widetilde{T}^{-1}_G(x,y) = 0$ we obtain the following lemma.

LEMMA 9.3. *Given a graph G and an edge $e \in E(G)$ we have*

$$\widetilde{T}^k_G(x,y) = \widetilde{T}^k_{G\setminus e}(x,y) + (y-1)\widetilde{T}^{k-1}_{G/e}(x,y)$$

for any $k \geq 0$ and similarly

$$\widetilde{T}_G(x,y) = \widetilde{T}_{G\setminus e}(x,y) + (y-1)\widetilde{T}_{G/e}(x,y).$$

Proof. In the definition of $\widetilde{T}^k_G$ we split the sum over $A \in \binom{E(G)}{k}$ as

$$\widetilde{T}^k_G(x,y) = \sum_{\substack{A \in \binom{E(G)}{k} \\ e \notin A}} (x-1)^{k(A)} \cdot (y-1)^{k(A)+\#A} + \sum_{\substack{A \in \binom{E(G)}{k} \\ e \in A}} (x-1)^{k(A)} \cdot (y-1)^{k(A)+\#A}. \tag{9.1}$$

The subsets $A \in \binom{E(G)}{k}$ with $e \notin A$ are naturally identified with the subsets $A \in \binom{E(G\setminus e)}{k}$, and we have $G(A) = (G \setminus e)(A)$. Thus the first sum in (9.1) is equal to $\widetilde{T}^k_{G\setminus e}(x,y)$. On the other hand, the subsets $A \in \binom{E(G)}{k}$ with $e \in A$ are naturally identified with the subsets $A' \in \binom{E(G/e)}{k-1}$ by $A \mapsto A' = A \setminus \{e\}$, and we have $k(A) = k(A')$ (in their respective graphs G and G/e). Thus the second summand in (9.1) equals $(y-1)\widetilde{T}^{k-1}_{G/e}(x,y)$, with the factor $(y-1)$ coming from the fact that $\#A = \#A' + 1$ in the above correspondence. The deletion-contraction formula for the (unparameterized) modified Tutte polynomial is obtained by summing over all k. ∎

Using the previous recurrence, the following transformation encapsulates the relation between the parameterized and the classical Tutte polynomial.

PROPOSITION 9.4. *Given a graph G and $k \geq 0$ we have*

$$\sum_{\ell=0}^{k} \binom{\#E(G)-\ell}{k-\ell} \cdot \widetilde{T}^\ell_G(x,y) = \sum_{A \in \binom{E(G)}{k}} \widetilde{T}_{G(A)}(x,y). \tag{9.2}$$

Proof. We prove the statement by induction on the number of edges. For $E(G) = \emptyset$ the two sides are zero for $k \neq 0$ and equal to $\widetilde{T}^0_G(x,y) = \widetilde{T}_G(x,y)$ for $k = 0$.

We show the induction step using the deletion-contraction relations above. Let G denote a graph with at least one edge e. Then we have

$$\sum_{\ell=0}^{k} \binom{\#E(G)-\ell}{k-\ell} \cdot \widetilde{T}^{\ell}_{G}(x,y) \tag{9.3}$$

$$= \sum_{\ell=0}^{k} \binom{\#E(G)-\ell}{k-\ell} \cdot \widetilde{T}^{\ell}_{G\setminus e}(x,y) + \binom{\#E(G)-\ell}{k-\ell} \cdot (y-1)\widetilde{T}^{\ell-1}_{G/e}(x,y)\,. \tag{9.4}$$

Furthermore, since $\#E(G) = \#E(G\setminus e)+1$, we can use the usual recursion of binomial coefficients to see

$$\begin{aligned}
&\sum_{\ell=0}^{k} \binom{\#E(G)-\ell}{k-\ell} \cdot \widetilde{T}^{\ell}_{G\setminus e}(x,y) \\
&= \sum_{\ell=0}^{k} \binom{\#E(G\setminus e)-\ell}{k-\ell} \cdot \widetilde{T}^{\ell}_{G\setminus e}(x,y) + \binom{\#E(G\setminus e)-\ell}{(k-1)-\ell} \cdot \widetilde{T}^{\ell}_{G\setminus e}(x,y) \\
&= \sum_{A\in\binom{E(G\setminus e)}{k}} \widetilde{T}_{(G\setminus e)(A)}(x,y) + \sum_{A'\in\binom{E(G\setminus e)}{k-1}} \widetilde{T}_{(G\setminus e)(A')}(x,y)\,,
\end{aligned}$$

where we have used the induction step. For the second summand in (9.4) we make the index shift $\ell' = \ell - 1$ and obtain

$$\begin{aligned}
&\sum_{\ell=0}^{k} \binom{\#E(G)-\ell}{k-\ell} \cdot (y-1)\widetilde{T}^{\ell-1}_{G/e}(x,y) \\
&= \sum_{\ell'=0}^{k-1} \binom{\#E(G/e)-\ell'}{(k-1)-\ell'} \cdot (y-1)\widetilde{T}^{\ell'}_{G/e}(x,y) \\
&= \sum_{A'\in\binom{E(G/e)}{k-1}} (y-1)\widetilde{T}_{(G/e)(A')}(x,y)\,.
\end{aligned}$$

Combining the last two equations we can conclude using suitable identifications, for instance, identifying the $A \in \binom{E(G)}{k}$ with $e \in A$ with $A' \in \binom{E(G\setminus e)}{k-1}$ via $A \mapsto A' = A\setminus\{e\}$ and using

$$(G\setminus e)(A') = G(A)\setminus e \text{ and } (G/e)(A') = G(A)/e\,.$$

Then we see that (9.4) equals

$$\begin{aligned}
&\sum_{A\in\binom{E(G\setminus e)}{k}} \widetilde{T}_{(G\setminus e)(A)}(x,y) + \sum_{A'\in\binom{E(G\setminus e)}{k-1}} \widetilde{T}_{(G\setminus e)(A')}(x,y) + (y-1)\widetilde{T}_{(G/e)(A')}(x,y) \\
&= \sum_{\substack{A\in\binom{E(G)}{k} \\ e\notin A}} \widetilde{T}_{G(A)}(x,y) + \sum_{\substack{A\in\binom{E(G)}{k} \\ e\in A}} \widetilde{T}_{G(A)}(x,y) \\
&= \sum_{A\in\binom{E(G)}{k}} \widetilde{T}_{G(A)}(x,y)\,.
\end{aligned}$$

∎

Using Proposition 9.4 we can now present combinatorial interpretations of the specialization of $T^k_G(x,y)$ to some individual points.

Chromatic polynomial. For $x = 1-c, y = 0$ the modified Tutte polynomial $\widetilde{T}_G(x,y)$ encodes the chromatic polynomial $\chi_G(c)$ (see, for instance, [49, Chapter 2]), so we see that the $\widetilde{T}_G^\ell(1-c,0)$ (for $0 \le \ell \le k$) contain the information of the number of pairs (A,σ) with $A \subseteq E(G)$ with $\#A = k$ and σ a c-coloring on $G(A)$:

$$\sum_{\ell=0}^{k} \binom{\#E(G)-\ell}{k-\ell} \cdot \widetilde{T}_G^\ell(1-c,0) = \#\left\{(A,\sigma) : A \in \binom{E(G)}{k}, \sigma \text{ } c\text{-coloring on } G(A)\right\}.$$

Acyclic orientations. For $x=2, y=0$ the Tutte polynomial $T_G(x,y)$ specializes to the number of acyclic orientations of G. We have $\widetilde{T}_G^k(2,0) = (-1)^{\#V(G)} T_G^k(2,0)$. Thus the $\widetilde{T}_G^\ell(2,0)$ (for $0 \le \ell \le k$) contain the information of the number of pairs $(A,\vec{\eta})$, where $A \subseteq E(G)$ with $\#A = k$ and $\vec{\eta}$ is an acyclic orientation on $G(A)$. Indeed, multiplying (9.2) with $(-1)^{\#V(G)}$ we obtain

$$\begin{aligned}&\sum_{\ell=0}^{k} \binom{\#E(G)-\ell}{k-\ell} \cdot T_G^\ell(2,0)\\ &= \#\left\{(A,\vec{\eta}) : A \in \binom{E(G)}{k}, \vec{\eta} \text{ acyclic orientation on } G(A)\right\}.\end{aligned}$$

k-edge sets inducing an even number of components.

PROPOSITION 9.5. *Given a graph G and a positive integer k, we have*

$$\begin{aligned}&\frac{1}{2}\left(\binom{\#E(G)}{k} + (-1)^{k(E(G))} T_G^k(0,2)\right)\\ &= \#\{A \subseteq E(G) : \#A = k \wedge k(A) = 0 \mod 2\}.\end{aligned}$$

Proof. Let $E = E(G)$. We have

$$\begin{aligned}\frac{1}{2}\left(\binom{\#E}{k} + (-1)^{k(E)} T_G^k(0,2)\right) &= \frac{1}{2}\left(\binom{\#E}{k} + (-1)^{k(E)} \sum_{A \in \binom{E}{k}} (-1)^{k(A)+k(E)}\right)\\ &= \frac{1}{2}\left(\sum_{A \in \binom{E}{k}} 1 + (-1)^{k(A)}\right).\end{aligned}$$

But observe that the summand above is 0 for $k(A)$ odd and 2 for $k(A)$ even. Thus after summing and dividing by 2 we count the subsets A with the graph $G(A)$ having an even number of components. ∎

k-edge sets of even Betti number. The (first) *Betti number*[34] of a graph is defined as $b_1(G) = k(E(G)) + \#E(G) - \#V(G)$ (cf. [7, Chapter 4]).

PROPOSITION 9.6. *Given a graph G and a positive integer k, we have*

$$\frac{1}{2}\left(\binom{\#E(G)}{k} + T_G^k(2,0)\right) = \#\{A \subseteq E(G) : \#A = k \wedge b_1(G(A)) = 0 \mod 2\}.$$

[34]The first Betti number is also called the circuit rank, cyclomatic number, cycle rank, or nullity.

Proof. We have

$$\frac{1}{2}\left(\binom{\#E(G)}{k}+T_G^k(2,0)\right)=\frac{1}{2}\left(\binom{\#E(G)}{k}+\sum_{A\in\binom{E(G)}{k}}(-1)^{k(A)+\#A-\#V(G)}\right)$$
$$=\frac{1}{2}\left(\sum_{A\in\binom{E(G)}{k}}1+(-1)^{b_1(G(A))}\right),$$

where we use $b_1(G(A))=k(A)+\#A-\#V(G)$. But observe that the summand above is 0 for $b_1(G(A))$ odd and 2 for $b_1(G(A))$ even. Thus after summing and dividing by 2 we count the subsets A with $G(A)$ having even Betti number. ∎

9.2. Classification for rational coordinates. We now classify the complexity of computing $T_G^k(x,y)$ for each pair of rational coordinates x and y. Formally, for each such pair, we consider the parameterized problem which expects as input G and k and outputs the value $T_G^k(x,y)$; the parameterization is given by k. Let us start with the following easy fact.

LEMMA 9.7. *For any $y\in\mathbb{Q}$, the problem of computing $T_G^k(1,y)$ is fixed-parameter tractable.*

Proof. Observe that $T_G^k(1,y)=0$ unless there is $A\subseteq E(G)$ of size k such that $k(A)=k(E(G))$. In other words, G has a spanning subgraph of k edges. Consequently, G can have at most $2k$ vertices, implying that G has at most $\binom{2k}{2}\le 4k^2$ many edges. Therefore an algorithm for computing $T_G^k(1,y)$ is obtained as follows: Given G and k, first check whether $|V(G)|>2k$, and output 0 in that case. Otherwise, obtain $T_G^k(x,y)$ by naively computing the sum, which takes time

$$O\left(\binom{4k^2}{k}\cdot|G|\right),$$

concluding the proof. ∎

Next, similarly to the classical counterpart [49], we obtain a trivial algorithm for coordinates x and y that lie on the hyperbola $(x-1)(y-1)=1$.

LEMMA 9.8. *Let x and y denote rational numbers such that $(x-1)(y-1)=1$. Then the problem of computing $T_G^k(x,y)$ is solvable in polynomial time (and thus fixed-parameter tractable as well).*

Proof. Observe that, given $(x-1)(y-1)=1$, and setting $V=V(G)$ and $E=E(G)$, we have

$$T_G^k(x,y)=\sum_{A\in\binom{E}{k}}(x-1)^{k(A)-k(E)}\cdot(y-1)^{k(A)+k-\#V}=(x-1)^{-k(E)}(y-1)^{k+\#V}\binom{\#E}{k},$$

which can be computed trivially. ∎

In what follows, we show that computing $T_G^k(x,y)$ is $\#\mathsf{W}[1]$-hard for *all remaining* rational coordinates x and y. First, it is convenient to rewrite the quantity $k(A)$ as follows: given an edge-subset A of a graph G, recall that $G[A]$ is the graph obtained from $(V(G),A)$ by deleting isolated vertices. Let us write $\mathsf{cc}(H)$ for the number of connected components of a graph H.

FACT 9.9. *Let G denote a graph, and let A denote a subset of edges of G. We have*

$$k(A) = \mathsf{cc}(G[A]) + \#V(G) - \#V(G[A]) .$$

Similarly as in the case of #EDGESUB(Φ), our goal is to reduce from a linear combination of (color-preserving) homomorphism counts. For this reason, we again consider an easy modification by excluding the term $(x-1)^{\#V(G)-\mathsf{cc}(G)}$; more precisely, consider

$$\widehat{T}^k_G(x,y) := \sum_{A \subseteq \binom{E(G)}{k}} (x-1)^{\mathsf{cc}(G[A])-\#V(G[A])} \cdot (y-1)^{\mathsf{cc}(G[A])-\#V(G[A])+k} ,$$

and observe that

$$T^k_G(x,y) = (x-1)^{\#V(G)-\mathsf{cc}(G)} \cdot \widehat{T}^k_G(x,y) .$$

In particular $\widehat{T}^k_G(x,y)$ is trivially interreducible with $T^k_G(x,y)$ if $x \neq 1$. Next we introduce an (H-)colored version of the parameterized Tutte polynomial; given an edge-subset A of a k-edge-colored graph, we write $\mathsf{cful}(A)$ if A contains each of the k colors precisely once.

DEFINITION 9.10 (colorful parameterized Tutte polynomial). *Let G denote a k-edge-colored graph. We define*

$$\mathsf{col}\text{-}\widehat{T}^k_G := \sum_{\substack{A \subseteq \binom{E(G)}{k} \\ \mathsf{cful}(A)}} (x-1)^{\mathsf{cc}(G[A])-\#V(G[A])} \cdot (y-1)^{\mathsf{cc}(G[A])-\#V(G[A])+k}$$

as the colorful parameterized Tutte polynomial.

The next lemma allows us to reduce the colorful version to the uncolored version.

LEMMA 9.11. *Let G denote a k-edge-colored graph, and assume that the set of colors is $[k]$. For each pair (x,y) we have*

$$\mathsf{col}\text{-}\widehat{T}^k_G(x,y) = \sum_{J \subseteq [k]} (-1)^{\#J} \cdot \tilde{T}^k_{G \setminus J}(x,y) ,$$

where $G \setminus J$ is the graph obtained from G by deleting all edges colored with an element of J.

Proof. The proof follows by the inclusion-exclusion principle (similarly as in Lemma 4.7) and the fact that, given a k-edge-subset A of G, deleting edges in $E(G)\setminus A$ does not change the quantity

$$(x-1)^{\mathsf{cc}(G[A])-\#V(G[A])} \cdot (y-1)^{\mathsf{cc}(G[A])-\#V(G[A])+k} . \qquad \square$$

Next, we express $\mathsf{col}\text{-}\widehat{T}^k_G$ as a linear combination of color-preserving homomorphisms counts. More precisely, given an H-colored graph G such that H has k edges, we implicitly assume the k-edge-coloring of G induced by its H-coloring. Further, given a fracture ρ of a graph H, we set $r(\sigma) := \mathsf{cc}(H \sharp \sigma) - \#V(H \sharp \sigma)$.

LEMMA 9.12. *Let H denote a graph with k edges. For every H-colored graph G, we have*

$$\#\mathsf{col}\text{-}\widehat{T}^k_G(x,y) = \sum_{\sigma\in\mathcal{L}(H)} (x-1)^{r(\sigma)} \cdot (y-1)^{r(\sigma)+k} \cdot \sum_{\rho\geq\sigma} \mu(\sigma,\rho)\cdot \#\mathsf{Hom}_{\mathsf{cp}}(H\sharp\rho \to_H G)\,,$$

where the relation $\leq$ and the Möbius function μ are over the lattice of fractures $\mathcal{L}(H)$.

Proof. Every colorful k-edge-subset A of G induces a fracture of H, similarly as we have seen in section 4. In particular, if A and A' induce the same fracture σ, then $G[A]\cong G[A']\cong H\sharp\sigma$. Writing $[\sigma]$ for the equivalence class of the induced fracture σ, we obtain

$$\#\mathsf{col}\text{-}\widehat{T}^k_G(x,y) = \sum_{\sigma\in\mathcal{L}(H)} (x-1)^{r(\sigma)} \cdot (y-1)^{r(\sigma)+k} \cdot \#[\sigma]\,.$$

Next observe that $\#[\sigma] = \#\mathsf{ColEdgeSub}(H\sharp\sigma \to_H G)$. Finally, we have already seen in (the proof of) Lemma 4.1 that

$$\#\mathsf{ColEdgeSub}(H\sharp\sigma \to_H G) = \sum_{\rho\geq\sigma} \mu(\sigma,\rho)\cdot \#\mathsf{Hom}_{\mathsf{cp}}(H\sharp\rho \to_H G)\,,$$

which concludes the proof. ∎

The following lemma establishes that the coefficient of the torus does not vanish apart from a few exceptions, which eventually allows us to prove $\#\mathsf{W}[1]$-hardness.

LEMMA 9.13. *Let $\ell > 2$ denote a prime, and let $(x,y)\in\mathbb{Q}^2$. There is a unique and computable function $a^\ell_{(x,y)}$ from fractures of $\circledcirc_\ell$ to rational numbers such that*

$$\mathsf{col}\text{-}\widehat{T}^{2\ell^2}_\star = \sum_{\rho\in\mathcal{L}(\circledcirc_\ell)} a^\ell_{(x,y)}(\rho)\cdot \#\mathsf{Hom}_{\mathsf{cp}}(\circledcirc_\ell\sharp\rho \to_{\circledcirc_\ell} \star)\,.$$

Moreover, if both the denominators of x,y and the numerators[35] *of $x-1$ and $(x-1)(y-1)-1$ are not divisible by ℓ, then $a^\ell_{(x,y)}(\top)\neq 0$.*

Proof. The first claim follows immediately from the previous lemma. For the second claim, we rely on the following fact from commutative algebra.

FACT 9.14. *Let $q\in\mathbb{Z}$ denote a nonzero integer; then the localization $\mathbb{Z}[1/q]$ of $\mathbb{Z}$ at q is the set*

$$\mathbb{Z}[1/q] = \left\{ u\in\mathbb{Q} : \exists v\in\mathbb{Z}, m\in\mathbb{N} \text{ with } u=\frac{v}{q^m}\right\}$$

of rational numbers which can be brought to a denominator which is a power of q. The subset $\mathbb{Z}[1/q]$ of $\mathbb{Q}$ is closed under addition and multiplication. Let furthermore ℓ denote a prime not dividing q, implying that q has an inverse q^{-1} mod ℓ. Then there is a well-defined map

$$\mathbb{Z}[1/q]\to\mathbb{Z}_\ell,\ \frac{v}{q^m}\mapsto v\cdot(q^{-1})^m,$$

and this map is compatible with addition and multiplication.

[35] In both cases, we refer to denominators and numerators of the corresponding shortened fractions.

Let us now collect the coefficients of $\#\mathsf{Hom}_{\mathsf{cp}}(\circledcirc_\ell \to_{\circledcirc_\ell} \star)$ in the sum appearing in Lemma 9.12. Completely similar to Corollary 4.3 we obtain

$$\text{(9.5)}\quad a^{\ell}_{(x,y)}(\top) = \sum_{\sigma\in\mathcal{L}(\circledcirc_\ell)} (x-1)^{r(\sigma)}(y-1)^{r(\sigma)+2\ell^2} \prod_{v\in V(\circledcirc_\ell)} (-1)^{|\sigma_v|-1}(|\sigma_v|-1)! \in \mathbb{Q}.$$

Note that in this expression we have for the exponents of $x-1$ and $y-1$ that $r(\sigma)\leq 0$ but $r(\sigma)+k\geq 0$. Let q denote the least common multiple of the denominators of x, y and the numerator of $x-1$; then we have that $(x-1)^{\pm 1}$ and $(y-1)$ are elements in $\mathbb{Z}[1/q]$. By the assumption that ℓ does not divide q together with Fact 9.14 above, we can see these expressions (and thus the entire sum (9.5)) as an element of $\mathbb{Z}_\ell$. Now recall the 15 fixed points of the action of $\mathbb{Z}_\ell^2$ on the fractures of $\circledcirc_\ell$ as given by Observation 5.3. Counting modulo ℓ allows us to rely on the same analysis as presented in the proof of Lemma 5.6, which yields that $a^{\ell}_{(x,y)}(\top)$ is, modulo ℓ, equal to

$$-6R(M_{2\ell^2}) + 4R(M_{\ell^2}+\ell C_\ell) + 8R(\ell^2 P_2) - R(2\ell C_\ell) - 2R(\ell C_{2\ell}) - 4R(\ell S_\ell) + R(\circledcirc_\ell)\,,$$

where $R(H) := (x-1)^{\mathsf{cc}(H)-\#V(H)}\cdot(y-1)^{2\ell^2+\mathsf{cc}(H)-\#V(H)}$.

Consequently, we have

$$\begin{aligned}
a^{\ell}_{(x,y)}(\top) = &-6(x-1)^{2\ell^2-4\ell^2}\cdot(y-1)^{2\ell^2+2\ell^2-4\ell^2}\\
&+4(x-1)^{\ell^2+\ell-3\ell^2}\cdot(y-1)^{2\ell^2+\ell^2+\ell-3\ell^2}\\
&+8(x-1)^{\ell^2-3\ell^2}\cdot(y-1)^{2\ell^2+\ell^2-3\ell^2}\\
&-1(x-1)^{2\ell-2\ell^2}\cdot(y-1)^{2\ell^2+2\ell-2\ell^2}\\
&-2(x-1)^{\ell-2\ell^2}\cdot(y-1)^{2\ell^2+\ell-2\ell^2}\\
&-4(x-1)^{\ell-2\ell^2}\cdot(y-1)^{2\ell^2+\ell-2\ell^2}\\
&+1(x-1)^{1-\ell^2}\cdot(y-1)^{2\ell^2+1-\ell^2}\quad \bmod \ell.
\end{aligned}$$

The first simplification is obtained by observing that the first and the third terms and the second, fifth, and sixth terms, respectively, contain the same monomial. Consequently, we have that (modulo ℓ)

$$\begin{aligned}
a^{\ell}_{(x,y)}(\top) = &+2(x-1)^{-2\ell^2}\\
&-2(x-1)^{-2\ell^2+\ell}\cdot(y-1)^{\ell}\\
&-(x-1)^{-2\ell^2+2\ell}\cdot(y-1)^{2\ell}\\
&+(x-1)^{1-\ell^2}\cdot(y-1)^{\ell^2+1}.
\end{aligned}$$

Using Fermat's little theorem, we obtain

$$\begin{aligned}
&a^{\ell}_{(x,y)}(\top)\\
=&2(x-1)^{-2} - 2(x-1)^{-2}\cdot(y-1) - (y-1)^2 + (y-1)^2 \quad \bmod \ell\\
=&2(x-1)^{-2}\cdot(1-(x-1)\cdot(y-1)) \quad \bmod \ell.
\end{aligned}$$

The assumption that the denominator of x (which is the numerator of $(x-1)^{-2}$) and the numerator of $(x-1)(y-1)-1$ are not divisible by ℓ implies that each factor in the product $2(x-1)^{-2}\cdot(1-(x-1)\cdot(y-1))$ gives a nonzero residue class mod ℓ.

Since ℓ is a prime, their product is still nonzero in $\mathbb{Z}_\ell$, and thus the original rational number $a^{\ell}_{(x,y)}(\top)$ is likewise nonzero, concluding the proof. ∎

We are thus able to rely on complexity monotonicity to establish hardness as promised.

LEMMA 9.15. *Let (x, y) denote a pair of rational numbers such that $(x-1)(y-1) \neq 1$ and $x \neq 1$. Then the problem of computing $T^k_G(x, y)$ is* #W[1]*-hard.*

Proof. Let $\mathcal{H}[x, y]$ denote the set of all $\circledcirc_\ell$ such that ℓ is prime and both the denominators of x and y as well as the numerators of $x-1$ and $(x-1)(y-1)-1$ are not divisible by ℓ. Since x and y are fixed, the latter is true for infinitely many primes ℓ, and thus $\mathcal{H}[x, y]$ contains tori of unbounded size. In particular, it contains graphs with arbitrary large grid minors and has thus unbounded treewidth [74], and hence, the problem #HOM($\mathcal{H}[x, y]$) is #W[1]-hard by the classification of Dalmau and Jonsson [28].

Completely analogously to the proof of Lemma 4.8, the problem #HOM($\mathcal{H}[x, y]$) reduces to computing col-$\widehat{T}^k_G(x, y)$ via complexity monotonicity (Lemma 4.6), since the coefficients of the tori do not vanish by the previous lemma.

Next, reducing to the uncolored version $\widehat{T}^k_G(x, y)$ can be done via Lemma 9.11, and, finally, $\widehat{T}^k_G(x, y)$ is trivially interreducible with $T^k_G(x, y)$ whenever $x \neq 1$. ∎

At last, we are able to prove this section's dichotomy theorem.

THEOREM 9.16. *Let (x, y) denote a pair of rational numbers. The problem of computing $T^k_G(x, y)$ is fixed-parameter tractable if $x = 1$ or $(x-1)(y-1) = 1$ and* #W[1]*-hard otherwise.*

Proof. The fixed-parameter tractable cases follow from Lemmas 9.7 and 9.8, and the #W[1]-hard cases follow from the previous lemma. ∎

As an immediate consequence, the computation of each individual point considered in subsection 9.1 is #W[1]-hard. Moreover, observe that the transformation

$$\sum_{\ell=0}^{k} \binom{\#E(G)-\ell}{k-\ell} \cdot \widetilde{T}^{\ell}_G(x, y) = \sum_{A \in \binom{E(G)}{k}} \widetilde{T}_{G(A)}(x, y)\,,$$

given by Proposition 9.4, is invertible in the sense that the numbers

$$\sum_{A \in \binom{E(G)}{\ell}} \widetilde{T}_{G(A)}(x, y)$$

for $\ell = 0, \ldots, k$ reveal $\widetilde{T}^k_G(x, y)$. Consequently, we obtain #W[1]-hardness of the information encoded in all considered individual points as well.

COROLLARY 9.17. *The following problems are* #W[1]*-hard when parameterized by k:*

- *Given G and k, compute the number of k-edge subsets A of G such that $G(A)$ has an even number of components.*
- *Given G and k, compute the number of pairs (A, σ) such that A is a k-edge subset of G and σ is a c-coloring of $G(A)$. Here $c \geq 2$ is a fixed integer.*
- *Given G and k, compute the number of pairs $(A, \vec{\eta})$ such that A is a k-edge subset of G and $\vec{\eta}$ is an acyclic orientation of $G(A)$.*

- *Given G and k, compute the number of k-edge subsets A of G such that $G(A)$ has even Betti number.*

9.2.1. Comparison to the classical dichotomy and real fixed-parameter tractable cases. In this section, we ask which of the fixed-parameter tractable cases allow for a polynomial-time algorithm. We can answer this question under the assumption $\mathsf{P} \neq \#\mathsf{P}$ by considering the classical dichotomy of Jaeger, Vertigan, and Welsh.[36]

THEOREM 9.18 (see [49]). *Let (x,y) be a pair of rational numbers. Then computing $T_G(x,y)$ is solvable in polynomial time if $(x,y) \in \{(1,1),(-1,-1),(0,-1),(-1,0)\}$ or if $(x-1)(y-1)=1$. In all other cases the problem is $\#\mathsf{P}$-hard.*

First, we observe that the parameterized dichotomy coincides with the classical dichotomy, except for the three points $(-1,-1)$, $(0,-1)$, and $(-1,0)$, in which the parameterized Tutte polynomial is ($\#\mathsf{W}[1]$-)hard to compute but the nonparameterized one is polynomial-time solvable. The latter indicates that taking the sum only over the k-edge subsets can, in fact, make the problem harder.

However, the nonparameterized Tutte polynomial always reduces to the parameterized Tutte polynomial *via polynomial-time Turing reductions*, since we can compute $T_G^0(x,y)+\cdots+T_G^{\#E(G)}(x,y)$ which is equal to $T_G(x,y)$. Thus any point (x,y) in which the nonparameterized Tutte polynomial is $\#\mathsf{P}$ hard and in which the parameterized Tutte polynomial is fixed-parameter tractable constitutes a "real" fixed-parameter tractable case. In particular, the latter shows that each point on the line $x=1$ yields a real fixed-parameter tractable case, except for the point $(1,1)$, which needs special treatment. More precisely, we have to determine whether computing $T_G^k(1,1)$ is not only fixed-parameter tractable (see Lemma 9.7) but also polynomial-time solvable. To this end, observe that

$$T_G^k(1,1) = \begin{cases} T_G(1,1), & k = \#V(G) - k(E(G)), \\ 0, & k \neq \#V(G) - k(E(G)), \end{cases}$$

since, for $x=y=1$, we have

$$(x-1)^{k(A)-k(E(G))}(y-1)^{k(A)+\#A-\#V(G)} = 0\,,$$

unless $\#A = \#V(G) - k(E(G))$. Thus, in point $(1,1)$, the parameterized Tutte polynomial can be computed in polynomial time by relying on the algorithm given by Theorem 9.18 in the case $k = \#V(G) - k(E(G))$ and outputting 0 otherwise.

Finally, recall that by Lemma 9.8, the case $(x-1)(y-1) = 1$ allows for a polynomial-time algorithm. The complete picture is hence given by the following refined classification; consider Figure 1 for a depiction of the tractable cases.

MAIN THEOREM 7. *Let (x,y) denote a pair of rational numbers. The problem of computing $T_G^k(x,y)$ is solvable in polynomial time if $x=y=1$ or $(x-1)(y-1)=1$, fixed-parameter tractable, but $\#\mathsf{P}$-hard if $x=1$ and $y\neq 1$ and $\#\mathsf{W}[1]$-hard otherwise.*

9.3. Approximating the parameterized Tutte polynomial. In the very last part of this paper, we identify rational points (x,y) for which $T_G^k(x,y)$ can be

[36]We state their classification only for rational numbers but point out that the full dichotomy includes all complex pairs.

approximated efficiently. Recall from Definition 7.1 that an FPTRAS for a parameterized counting problem (P,κ) is a (randomized) algorithm $\mathbb{A}$ which, on input I, ε, and δ, outputs a value z with probability

$$\mathsf{Pr}[\,(1-\varepsilon)P(I)\le z\le(1+\varepsilon)P(I)\,]\ge 1-\delta$$

in time $f(\kappa(I))\cdot\mathsf{poly}(|I|,\varepsilon^{-1},\log(1/\delta))$ for some computable function f. If f is a polynomial as well, then $\mathbb{A}$ is called an FPRAS (cf. [68, Definition 11.2]).

We have to be careful when speaking about approximating $T^k_G(x,y)$ since the latter can have negative values. One way of dealing with negative valued functions is to require that an FPTRAS/FPRAS outputs a pair (z,s) such that z is an approximation of the absolute value and s is the sign; that is, we require that with probability at least $(1-\delta)$, we have

$$(1-\varepsilon)|T^k_G(x,y)|\le z\le(1+\varepsilon)|T^k_G(x,y)| \text{ and } s=\mathsf{sign}(T^k_G(x,y))\,.$$

We are now able to establish a region of rational points for which the parameterized Tutte polynomial admits an FPTRAS or even an FPRAS; the proof is a simple consequence of the work of Anari et al. on approximate counting via log-concave polynomials [5].

Main Theorem 8. *Let (x,y) denote a pair of rational numbers satisfying the condition $0\le(x-1)(y-1)\le 1$. Then $T^k_G(x,y)$ has an FPTRAS. If additionally $x\neq 1$ or $y=1$, then $T^k_G(x,y)$ even has an FPRAS.*

Proof. The case $x=1$ is a trivial consequence of Main Theorem 7. If $x=1$, then exact counting is fixed-parameter tractable, and thus there is an FPTRAS. We consider two cases for the remaining points.

First, consider $y=1$. If $x=1$ as well, then we obtain by Main Theorem 7 a polynomial-time algorithm for exact counting and thus an FPRAS. Otherwise, we have that $T^k_G(x,1)$ is equal to

$$\sum_{A\in\binom{E(G)}{k}}(x-1)^{k(A)-k(E(G))}\cdot 0^{k(A)+k-\#V(G)}=\sum_{\substack{A\in\binom{E(G)}{k}\\ G(A)\text{ acyclic}}}(x-1)^{k(A)-k(E(G))}\,,$$

since $k(A)+k-\#V(G)=0$ if and only if $G(A)$ is acyclic. Recall that $k(A)$ is the number of connected components of $G(A)(=(V(G),A))$, and observe that for acyclic sets of edges A with $|A|=k$, we have $k(A)=\#V(G)-k$. Consequently,

$$T^k_G(x,1)=(x-1)^{\#V(G)-k-k(E(G))}\cdot\#\{A\subseteq E(G)\,|\,\#A=k\wedge G(A)\text{ acyclic}\}\,.$$

Since computing the number of acyclic edge-subsets of size k admits an FPRAS [4, 5], we can conclude this case.

In the remaining case, we have $x\neq 1$ and $y\neq 1$ (and thus $0<(x-1)(y-1)\le 1$). Let $q=(x-1)(y-1)$, and let G denote a graph with edges $E=[m]$, that is, G has m edges labeled with $1,\dots,m$. Consider the polynomial

$$f_{G,k,q}(x_1,\dots,x_m):=\sum_{A\in\binom{E}{k}}q^{-\mathsf{rk}(A)}\prod_{i\in A}x_i\,,$$

where $\mathsf{rk}(A):=\#V(G)-k(A)$ is the rank of A with respect to the graphic matroid of G. Anari et al. have established the existence of an FPRAS for evaluating $f_{G,k,q}(1_m)$ whenever $0<q\le 1$ [5, section 1.2]. Now observe that

$$T^k_G(x,y)=(x-1)^{\#V(G)-k(E(G))}(y-1)^k\cdot f_{G,k,q}(1_m)\,.$$

Since neither $x = 1$ nor $y = 1$, we conclude that an FPRAS for $f_{G,k,q}(1_m)$ yields the desired FPRAS for $T^k_G(x, y)$. ∎

Acknowledgments. We thank Johannes Lengler and Holger Dell for helpful discussions on early drafts of this work. The third author thanks the Max Planck Institute for Mathematics in Bonn for its hospitality.

REFERENCES

[1] N. Alon, P. Dao, I. Hajirasouliha, F. Hormozdiari, and S. C. Sahinalp, *Biomolecular network motif counting and discovery by color coding*, Bioinformatics, 24 (2008), pp. i241–i249, https://doi.org/10.1093/bioinformatics/btn163.

[2] N. Alon, A. M. Frieze, and D. Welsh, *Polynomial time randomized approximation schemes for Tutte-Gröthendieck invariants: The dense case*, Random Structures Algorithms, 6 (1995), pp. 459–478, https://doi.org/10.1002/rsa.3240060409.

[3] N. Alon, R. Yuster, and U. Zwick, *Color-Coding*, J. ACM, 42 (1995), pp. 844–856, https://doi.org/10.1145/210332.210337.

[4] N. Anari and M. Derezinski, *Isotropy and log-concave polynomials: Accelerated sampling and high-precision counting of matroid bases*, in Proceedings of the 61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, IEEE, 2020, pp. 1331–1344, https://doi.org/10.1109/FOCS46700.2020.00126.

[5] N. Anari, K. Liu, S. O. Gharan, and C. Vinzant, *Log-concave polynomials* II*: High-dimensional walks and an FPRAS for counting bases of a matroid*, in Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, 2019, pp. 1–12, https://doi.org/10.1145/3313276.3316385.

[6] V. Arvind and V. Raman, *Approximation Algorithms for Some Parameterized Counting Problems*, in Proceedings of the 13th International Symposium on Algorithms and Computation, ISAAC 2002, Vancouver, BC, Canada, 2002, pp. 453–464, https://doi.org/10.1007/3-540-36136-7.

[7] C. Berge, *The Theory of Graphs*, Dover Publications, Mineola, NY, 2001.

[8] A. Björklund, T. Husfeldt, P. Kaski, and M. Koivisto, *Computing the Tutte polynomial in vertex-exponential time*, in Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, IEEE Computer Society, Philadelphia, PA, 2008, pp. 677–686, https://doi.org/10.1109/FOCS.2008.40.

[9] A. Björklund, T. Husfeldt, P. Kaski, and M. Koivisto, *Narrow sieves for parameterized paths and packings*, J. Comput. Syst. Sci., 87 (2017), pp. 119–139, https://doi.org/10.1016/j.jcss.2017.03.003.

[10] A. Björklund and P. Kaski, *The fine-grained complexity of computing the Tutte polynomial of a linear matroid*, in Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Alexandria, VA, 2021, pp. 2333–2345.

[11] C. Brand, H. Dell, and T. Husfeldt, *Extensor-coding*, in Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, 2018, pp. 151–164, https://doi.org/10.1145/3188745.3188902.

[12] C. Brand, H. Dell, and M. Roth, *Fine-grained dichotomies for the Tutte plane and boolean #CSP*, Algorithmica, 81 (2019), pp. 541–556, https://doi.org/10.1007/s00453-018-0472-z.

[13] C. Brand and M. Roth, *Parameterized counting of trees, forests and matroid bases*, in Proceedings of the 12th International Computer Science Symposium in Russia, CSR 2017, Kazan, Russia, 2017, pp. 85–98, https://doi.org/10.1007/978-3-319-58747-9.

[14] A. K. Chandra and P. M. Merlin, *Optimal implementation of conjunctive queries in relational data bases*, in Proceedings of the 9th Annual ACM Symposium on Theory of Computing, Boulder, CO, 1977, pp. 77–90, https://doi.org/10.1145/800105.803397.

[15] H. Chen and S. 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, 2016, pp. 315–326, https://doi.org/10.1145/2902251.2902279.

[16] J. Chen, B. Chor, M. Fellows, X. Huang, D. W. Juedes, I. A. Kanj, and G. Xia, *Tight lower bounds for certain parameterized NP-hard problems*, Inform. and Comput., 201 (2005), pp. 216–231, https://doi.org/10.1016/j.ic.2005.05.001.

[17] J. Chen, X. Huang, I. A. Kanj, and G. Xia, *Strong computational lower bounds via parameterized complexity*, J. Comput. Syst. Sci., 72 (2006), pp. 1346–1367, https://doi.org/10.1016/j.jcss.2006.04.007.

[18] J. Chen, J. Kneis, S. Lu, D. Mölle, S. Richter, P. Rossmanith, S.-H. Sze, and F. Zhang, *Randomized divide-and-conquer: Improved path, matching, and packing algorithms*, SIAM J. Comput., 38 (2009), pp. 2526–2547, https://doi.org/10.1137/080716475.

[19] Y. Chen, M. Thurley, and M. Weyer, *Understanding the complexity of induced subgraph isomorphisms*, in Proceedings of the 35th International Colloquium on Automata, Languages, and Programming, ICALP 2008, Reykjavik, Iceland, 2008, pp. 587–596, https://doi.org/10.1007/978-3-540-70575-8.

[20] S. A. Cook, *The complexity of theorem-proving procedures*, in Proceedings of the 3rd Annual ACM Symposium on Theory of Computing, Shaker Heights, OH, 1971, pp. 151–158, 1971, https://doi.org/10.1145/800157.805047.

[21] D. G. Corneil and C. C. Gotlieb, *An efficient algorithm for graph isomorphism*, J. ACM, 17 (1970), pp. 51–64, https://doi.org/10.1145/321556.321562.

[22] B. Courcelle, *Graph rewriting: An algebraic and logic approach*, in Handbook of Theoretical Computer Science, Volume B: Formal Models and Sematics (B), Elsevier Science, New York, 1990, pp. 193–242.

[23] R. Curticapean, *Counting matchings of size k is* #W[1]*-hard*, in Proceedings of the 40th International Colloquium on Automatata, Languages, and Programming, ICALP 2013, F. V. Fomin, R. Freivalds, M. Z. Kwiatkowska, and D. Peleg, eds., Lecture Notes in Comput. Sci. 7965, Springer, Riga, Latvia, 2013, pp. 352–363, https://doi.org/10.1007/978-3-642-39206-1.

[24] R. Curticapean, *The Simple, Little and Slow things Count: On Parameterized Counting Complexity*, Ph.D. thesis, Saarland University, 2015, http://scidok.sulb.uni-saarland.de/volltexte/2015/6217/.

[25] R. Curticapean, H. Dell, and D. 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, 2017, pp. 210–223, https://doi.org/10.1145/3055399.3055502.

[26] R. Curticapean and D. Marx, *Complexity of counting subgraphs: Only the boundedness of the vertex-cover number counts*, in Proceedings of the 55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, 2014, pp. 130–139, https://doi.org/10.1109/FOCS.2014.22.

[27] M. Cygan, F. V. Fomin, L. Kowalik, D. Lokshtanov, D. Marx, M. Pilipczuk, M. Pilipczuk, and S. Saurabh, *Parameterized Algorithms*, Springer, Cham, 2015, https://doi.org/10.1007/978-3-319-21275-3.

[28] V. Dalmau and P. Jonsson, *The complexity of counting homomorphisms seen from the other side*, Theoret. Comput. Sci., 329 (2004), pp. 315–323, https://doi.org/10.1016/j.tcs.2004.08.008.

[29] H. Dell, T. Husfeldt, D. Marx, N. Taslaman, and M. Wahlen, *Exponential time complexity of the permanent and the tutte polynomial*, ACM Trans. Algorithms, 10 (2014), 21, https://doi.org/10.1145/2635812.

[30] H. Dell, J. Lapinskas, and K. Meeks, *Approximately counting and sampling small witnesses using a colourful decision oracle*, inProceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, 2020, pp. 2201–2211, https://doi.org/10.1137/1.9781611975994.135.

[31] H. Dell, M. Roth, and P. Wellnitz, *Counting answers to existential questions*, in Proceedings of the 46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, Patras, Greece, 2019, pp. 113:1–113:15, https://doi.org/10.4230/LIPIcs.ICALP.2019.113.

[32] J. Dörfler, M. Roth, J. Schmitt, and P. Wellnitz, *Counting induced subgraphs: An algebraic approach to #W[1]-hardness*, in Proceedings of the 44th International Symposium on Mathematical Foundations of Computer Science, MFCS 2019, P. Rossmanith, P. Heggernes, and J. Katoen, eds., LIPIcs. Leibniz Int. Proc. Inform. 138, Schloss Dagstuhl Leibniz-Zentrum für Informatik, Wadern, Germany, 2019, https://doi.org/10.4230/LIPIcs.MFCS.2019.26.

[33] R. G. Downey and M. R. Fellows, *Fixed-parameter tractability and completeness* I*: Basic results*, SIAM J. Comput., 24 (1995), pp. 873–921, https://doi.org/10.1137/S0097539792228228.

[34] R. G. Downey and M. R. Fellows, *Fixed-parameter tractability and completeness* II*: On completeness for W*[1], Theoret. Comput. Sci., 141 (1995), pp. 109–131, https://doi.org/10.1016/0304-3975(94)00097-3.

[35] R. G. Downey and M. R. Fellows, *Fundamentals of Parameterized Complexity,* Texts Comput. Sci., Springer, Cham, 2013, https://doi.org/10.1007/978-1-4471-5559-1

[36] A. Durand and S. Mengel, *Structural tractability of counting of solutions to conjunctive queries*, Theory Comput. Syst., 57 (2015), pp. 1202–1249, https://doi.org/10.1007/s00224-014-9543-y.

[37] J. Edmonds, *Paths, trees, and flowers*, Canad. J. Math., 17 (1965), pp. 449–467, https://doi.org/10.4153/CJM-1965-045-4.

[38] J. Flum and M. Grohe, *The parameterized complexity of counting problems*, SIAM J. Comput., 33 (2004), pp. 892–922, https://doi.org/10.1137/S0097539703427203.

[39] J. Flum and M. Grohe, Parameterized Complexity Theory, Texts Theoret. Comput. Sci. EATCS Ser., Springer, Cham, 2006, https://doi.org/10.1007/3-540-29953-X.

[40] F. V. Fomin, D. Lokshtanov, and S. Saurabh, *Efficient computation of representative sets with applications in parameterized and exact algorithms*, in Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, OR, 2014, pp. 142–151, https://doi.org/10.1137/1.9781611973402.10.

[41] M. Frick and M. Grohe, *Deciding first-order properties of locally tree-decomposable structures*, J. ACM, 48 (2001), pp. 1184–1206, https://doi.org/10.1145/504794.504798.

[42] L. A. Goldberg and M. Jerrum, *The complexity of computing the sign of the tutte polynomial*, SIAM J. Comput., 43 (2014), pp. 1921–1952, https://doi.org/10.1137/12088330X.

[43] J. A. Grochow and M. Kellis, *Network motif discovery using subgraph enumeration and symmetry-breaking*, in Research in Computational Molecular Biology, T. Speed and H. Huang, eds., Springer, Berlin, 2007, pp. 92–106.

[44] M. Grohe, *Parameterized complexity for the database theorist*, ACM SIGMOD Rec., 31 (2002), pp. 86–96, https://doi.org/10.1145/637411.637428.

[45] M. Grohe and D. Marx, *On tree width, bramble size, and expansion*, J. Combin. Theory Ser. B, 99 (2009), pp. 218–228, https://doi.org/10.1016/j.jctb.2008.06.004.

[46] M. Grohe, T. Schwentick, and L. Segoufin, *When is the evaluation of conjunctive queries tractable?*, in Proceedings of the 33rd Annual ACM Symposium on Theory of Computing, Heraklion, Crete, Greece, 2001, pp. 657–666, 2001, https://doi.org/10.1145/380752.380867.

[47] S. Hoory, N. Linial, and A. Wigderson, *Expander graphs and their applications*, Bull. Amer. Math. Soc. (N.S.), 43 (2006), pp. 439–561.

[48] R. Impagliazzo and R. Paturi, *On the complexity of k-SAT*, J. Comput. Syst. Sci., 62 (2001), pp. 367–375, https://doi.org/10.1006/jcss.2000.1727.

[49] F. Jaeger, D. L. Vertigan, and D. J. A. Welsh, *On the computational complexity of the Jones and Tutte polynomials*, Math. Proc. Cambridge Philos. Soc., 108 (1990), pp. 35–53, https://doi.org/10.1017/S0305004100068936.

[50] M. Jerrum and K. Meeks, *The parameterised complexity of counting connected subgraphs and graph motifs*, J. Comput. Syst. Sci., 81 (2015), pp. 702–716, https://doi.org/10.1016/j.jcss.2014.11.015.

[51] M. Jerrum and K. Meeks, *Some hard families of parameterized counting problems*, ACM Trans. Comput. Theory, 7 (2015), 11, https://doi.org/10.1145/2786017.

[52] M. Jerrum and K. Meeks, *The parameterised complexity of counting even and odd induced subgraphs*, Combinatorica, 37 (2017), pp. 965–990, https://doi.org/10.1007/s00493-016-3338-5.

[53] S. Khot and V. Raman, *Parameterized complexity of finding subgraphs with hereditary properties*, Theoret. Comput. Sci., 289 (2002), pp. 997–1008, https://doi.org/10.1016/S0304-3975(01)00414-5.

[54] J. M. Kleinberg and R. Rubinfeld, *Short paths in expander graphs*, in Proceedings of the 37th Annual Symposium on Foundations of Computer Science, FOCS '96, Burlington, VT, 1996, pp. 86–95, https://doi.org/10.1109/SFCS.1996.548467.

[55] E. Kowalski, *An Introduction to Expander Graphs*, Cours Spéc. 26, Société Mathématique de France, Paris, 2019.

[56] M. Krebs and A. Shaheen, *Expander Families and Cayley Graphs: A Beginner's Guide*, Oxford University Press, Oxford, 2011.

[57] S. Lang, *Algebra*, 3rd ed., Addison-Wesley, Reading, MA, 1993.

[58] B. Lin, *The parameterized complexity of the k-biclique problem*, J. ACM, 65 (2018), 34, https://doi.org/10.1145/3212622.

[59] D. Lokshtanov, P. Misra, F. Panolan, and S. Saurabh, *Deterministic truncation of linear matroids*, ACM Trans. Algorithms, 14 (2018), pp. 14:1–14:20, https://doi.org/10.1145/3170444.

[60] L. Lovász, *Large Networks and Graph Limits*, Amer. Math. Soc. Colloq. Publ. 60, American Mathematical Society, Providence, RI, 2012.

[61] A. Lubotzky, R. Phillips, and P. Sarnak, *Ramanujan graphs*, Combinatorica, 8 (1988), pp. 261–277, https://doi.org/10.1007/BF02126799.

[62] G. A. Margulis, *Discrete Subgroups of Semisimple Lie Groups*, Vol. 3, Springer-Verlag, Berlin, 1991.

[63] D. Marx, *Can you beat treewidth?*, Theory Comput., 6 (2010), pp. 85–112, https://doi.org/10.4086/toc.2010.v006a005.

[64] D. Marx, *Tractable hypergraph properties for constraint satisfaction and conjunctive queries*, J. ACM, 60 (2013), 42, https://doi.org/10.1145/2535926.

[65] K. Meeks, *The challenges of unbounded treewidth in parameterised subgraph counting problems*, Discrete Appl. Math., 198 (2016), pp. 170–194, https://doi.org/10.1016/j.dam.2015.06.019.

[66] R. Milo, S. Itzkovitz, N. Kashtan, R. Levitt, S. Shen-Orr, I. Ayzenshtat, M. Sheffer, and U. Alon, *Superfamilies of evolved and designed networks*, Science, 303 (2004), pp. 1538–1542, https://doi.org/10.1126/science.1089167.

[67] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon, *Network motifs: Simple building blocks of complex networks*, Science, 298 (2002), pp. 824–827, https://doi.org/10.1126/science.298.5594.824.

[68] M. Mitzenmacher and E. Upfal, *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*, 2nd ed., Cambridge University Press, Cambridge, UK, 2017.

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

[70] N. Peyerimhoff, M. Roth, J. Schmitt, J. Stix, and A. Vdovina, *Parameterized (modular) counting and Cayley graph expanders*, in Proceedings of the 46th International Symposium on Mathematical Foundations of Computer Science, LIPIcs. Leibniz Int. Proc. Inform. 202, Schloss Dagstuhl Leibniz-Zentrum für Informatik, Wadern, Germany, 2021, pp. 84:1–84:15.

[71] N. Peyerimhoff and A. Vdovina, *Cayley graph expanders and groups of finite width*, J. Pure Appl. Algebra, 215 (2011), pp. 2780–2788, https://doi.org/10.1016/j.jpaa.2011.03.018.

[72] J. Plehn and B. Voigt, *Finding minimally weighted subgraph*, in Proceedings of the 16th International Workshop on Graph-Theoretic Concepts in Computer Science, Lecture Notes in Comput. Sci. 484, Springer, Berlin, 1990, pp. 18–29, https://doi.org/10.1007/3-540-53832-1.

[73] O. Rahat, U. Alon, Y. Levy, and G. Schreiber, *Understanding hydrogen-bond patterns in proteins using network motifs*, Bioinformatics, 25 (2009), pp. 2921–2928, https://doi.org/10.1093/bioinformatics/btp541.

[74] N. Robertson and P. D. Seymour, *Graph minors.* V. *Excluding a planar graph*, J. Combin. Theory Ser. B, 41 (1986), pp. 92–114, https://doi.org/10.1016/0095-8956(86)90030-4.

[75] N. Robertson and P. D.Seymour, *Graph minors.* XX. *Wagner's conjecture*, J. Combin. Theory Ser. B, 92 (2004), pp. 325–357, https://doi.org/10.1016/j.jctb.2004.08.001.

[76] N. Robertson, P. D. Seymour, and R. Thomas, *Quickly excluding a planar graph*, J. Combin. Theory Ser. B, 62 (1994), pp. 323–348, https://doi.org/10.1006/jctb.1994.1073.

[77] M. Roth, *Counting Problems on Quantum Graphs: Parameterized and Exact Complexity Classifications*, Ph.D. thesis, Saarland University, 2019, https://scidok.sulb.uni-saarland.de/bitstream/20.500.11880/27575/1/thesis.pdf.

[78] M. Roth and J. Schmitt, *Counting induced subgraphs: A topological approach to* $\#W[1]$*-hardness*, Algorithmica, 82 (2020), pp. 2267–2291, https://doi.org/10.1007/s00453-020-00676-9.

[79] M. Roth, J. Schmitt, and P. Wellnitz, *Counting small induced subgraphs satisfying monotone properties*, in Proceedings of the 61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, 2020, pp. 1356–1367, https://doi.org/10.1109/FOCS46700.2020.00128.

[80] M. Roth, J. Schmitt, and P. Wellnitz, *Detecting and counting small subgraphs, and evaluating a parameterized Tutte polynomial: Lower bounds via toroidal grids and Cayley graph expanders*, in Proceedings of the 48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, LIPIcs. Leibniz Int. Proc. Inform. 198, Schloss Dagstuhl Leibniz-Zentrum für Informatik, Wadern, Germany, 2021, pp. 108:1–108:16, https://doi.org/10.4230/LIPIcs.ICALP.2021.108.

[81] N. Rungtanapirom, J. Stix, and A. Vdovina, *Infinite series of quaternionic* 1*-vertex cube complexes, the doubling construction, and explicit cubical Ramanujan complexes*, Internat. J. Algebra Comput., 29 (2019), pp. 951–1007.

[82] B. Schiller, S. Jager, K. Hamacher, T. Strufe, and Strea, M, *StreaM - a stream-based algorithm for counting motifs in dynamic graphs*, in Algorithms for Computational Biology, A.-H. Dediu, F. Hernández-Quiroz, C. Martín-Vide, and D. A. Rosenblueth, eds., Springer International Publishing, Cham, 2015, pp. 53–67.

[83] F. Schreiber and H. Schwöbbermeyer, *Frequency concepts and pattern detection for the analysis of motifs in networks*, in Transactions on Computational Systems Biology III, C. Priami, E. Merelli, P. Gonzalez, and A. Omicini, eds., Springer, Berlin, 2005, pp. 89–104.
[84] R. P. Stanley, *Enumerative Combinatorics*, Vol. 1, Cambridge University Press, Cambridge, UK, 2011.
[85] J. Stix and A. Vdovina, *Simply transitive quaternionic lattices of rank* 2 *over* $\mathbb{F}_{q(t)}$ *and a non-classical fake quadric*, Math. Proc. Cambridge Philos. Soc., 163 (2017), pp. 453–498.
[86] The Sage Developers, *SageMath, the Sage Mathematics Software System (Version* 9.1*)*, 2020, https://www.sagemath.org.
[87] S. Toda, *PP is as hard as the polynomial-time hierarchy*, SIAM J. Comput., 20 (1991), pp. 865–877, https://doi.org/10.1137/0220053.
[88] J. R. Ullmann, *An algorithm for subgraph isomorphism*, J. ACM, 23 (1976), pp. 31–42, https://doi.org/10.1145/321921.321925.
[89] L. G. Valiant, *The complexity of computing the permanent*, Theoret. Comput. Sci., 8 (1979), pp. 189–201, https://doi.org/10.1016/0304-3975(79)90044-6.
[90] L. G. Valiant, *The complexity of enumeration and reliability problems*, SIAM J. Comput., 8 (1979), pp. 410–421, https://doi.org/10.1137/0208032.
[91] D. Vertigan, *Bicycle dimension and special points of the Tutte polynomial*, J. Combin. Theory Ser. B, 74 (1998), pp. 378–396, https://doi.org/10.1006/jctb.1998.1860.
[92] D. Witte, *Cayley digraphs of prime-power order are Hamiltonian*, J. Combin. Theory Ser. B, 40 (1986), pp. 107–112, https://doi.org/10.1016/0095-8956(86)90068-7.