Counting Restricted Homomorphisms via Möbius Inversion over Matroid Lattices[*][†]

Marc Roth

Saarland University and Cluster of Excellence (MMCI), Saarbrücken, Germany
`mroth@mmci.uni-saarland.de`

⸻ Abstract ⸻

We present a framework for the complexity classification of parameterized counting problems that can be formulated as the summation over the numbers of homomorphisms from small pattern graphs $H_1, \ldots, H_\ell$ to a big host graph G with the restriction that the coefficients correspond to evaluations of the Möbius function over the lattice of a graphic matroid. This generalizes the idea of Curticapean, Dell and Marx [STOC 17] who used a result of Lovász stating that the number of subgraph embeddings from a graph H to a graph G can be expressed as such a sum over the lattice of partitions of H.

In the first step we introduce what we call graphically restricted homomorphisms that, inter alia, generalize subgraph embeddings as well as locally injective homomorphisms. We provide a complete parameterized complexity dichotomy for counting such homomorphisms, that is, we identify classes of patterns for which the problem is fixed-parameter tractable (FPT), including an algorithm, and prove that all other pattern classes lead to #W[1]-hard problems. The main ingredients of the proof are the complexity classification of linear combinations of homomorphisms due to Curticapean, Dell and Marx [STOC 17] as well as a corollary of Rota's NBC Theorem which states that the sign of the Möbius function over a geometric lattice only depends on the rank of its arguments.

We apply the general theorem to the problem of counting locally injective homomorphisms from small pattern graphs to big host graphs yielding a concrete dichotomy criterion. It turns out that – in contrast to subgraph embeddings – counting locally injective homomorphisms has "real" FPT cases, that is, cases that are fixed-parameter tractable but not polynomial time solvable under standard complexity assumptions. To prove this we show in an intermediate step that the subgraph counting problem remains #P-hard when both the pattern and the host graphs are restricted to be trees. We then investigate the more general problem of counting homomorphisms that are injective in the r-neighborhood of every vertex. As those are graphically restricted as well, they can also easily be classified via the general theorem.

Finally we show that the dichotomy for counting graphically restricted homomorphisms readily extends to so-called linear combinations.

---

[*] A full version of the paper is available at `https://arxiv.org/abs/1706.08414`.
[†] Part of this work was done while the author was visiting the Simons Institute for the Theory of Computing

1 Introduction

In his seminal work about the complexity of computing the permanent Valiant [29] introduced counting complexity which has since then evolved into a well-studied subfield of computational complexity. Despite some surprising positive results like polynomial time algorithms for counting perfect matchings in planar graphs by the FKT method [27, 17], counting spanning trees by Kirchhoff's Matrix Tree Theorem or counting Eulerian cycles in directed graphs using the "BEST"-Theorem (see e.g. [2]), most of the interesting problems turned out to be intractable. Therefore, several relaxations such as restrictions of input classes [33] and approximate counting [16, 11] were introduced. Another possible relaxation, the one this work deals with, is to consider parameterized counting problems as introduced by Flum and Grohe [14]. Here, problems come with an additional parameter k and a problem is fixed-parameter tractable (FPT) if it can be solved in time $g(k) \cdot \mathrm{poly}(n)$ where n is the input size and g is a computable function, which yields fast algorithms for large instances with small parameters. On the other hand, a problem is considered intractable if it is #W[1]-hard. This stems from the fact that #W[1]-hard problems do not allow an FPT algorithm unless standard assumptions such as the exponential time hypothesis (ETH) are wrong.

When investigating a family of related (counting) problems one could aim to simultaneously solve the complexity of as many problems as possible, rather than tackling a (possibly infinite) number of problems by hand. For example, instead of proving that counting paths in a graph is hard, then proving that counting cycles is hard and then proving that counting stars is easy, one should, if possible, find a criterion that allows a classification of those problems in hard and easy cases. Unfortunately, there are results like Ladner's Theorem [18], stating that there are problems neither in P nor NP-hard (assuming P ≠ NP), which give a negative answer to that goal in general. However, there are families of problems that have enough structure to allow so-called dichotomy results. One famous example, and to the best of the authors knowledge this was the first such result, is Schaefer's dichotomy [25], stating that every instance of the generalized satisfiability problem is either polynomial time solvable or NP-complete. Since then much work has been done to generalize this result, culminating in recent announcements ([3],[34],[23]) of a proof of the Feder-Vardi-Conjecture [12]. This question was open for almost twenty years and indicates the difficulty of proving such dichotomy results, at least for decision problems. In counting complexity, however, it seems that obtaining such results is less cumbersome. One reason for this is the existence of some powerful techniques like polynomial interpolation [28], the Holant framework [30, 31, 4] as well as the principle of inclusion-exclusion which all have been used to establish very revealing dichotomy results such as [5, 9].

Examples of dichotomies in parameterized counting complexity are the complete classifications of the homomorphism counting problem due to Dalmau and Jonsson [10][1] and the subgraph counting problem due to Curticapean and Marx [9]. For the latter, one is given graphs H and G and wants to count the number of subgraphs of G isomorphic to H, parameterized by the size of H. It is known that this problem is polynomial time solvable if there is a constant upper bound on the size of the largest matching of H and #W[1]-hard otherwise[2]. The first step in this proof was the hardness result of counting matchings of size k of Curticapean [6], which turned out to be the "bottleneck" problem and was then reduced to the general problem.

---

[1] Ultimately, the results of [7] and this work rely on the dichotomy for counting homomorphisms.

[2] On the other hand the complexity of the decision version of this problem, that is, finding a subgraph of G isomorphic to H, is still unresolved. Only recently it was shown in a major breakthrough that finding bicliques is hard [19].

This approach, first finding the hard obstructions and then reducing to the general case, seemed to be the canonical way to tackle such problems. However, recently Curticapean, Dell and Marx [7] discovered that a result of Lovász [20] implies the existence of parameterized reductions that, inter alia, allow a far easier proof of the general subgraph counting problem. Lovász result states that, given simple graphs H and G, it holds that

$$\#\mathsf{Emb}(H,G) = \sum_{\rho \geq \emptyset} \mu(\emptyset, \rho) \cdot \#\mathsf{Hom}(H/\rho, G) \,, \tag{1}$$

where the sum is over the elements of the partition lattice of $V(H)$, $\mathsf{Emb}(H,G)$ is the set of embeddings[3] from H to G and $\mathsf{Hom}(H/\rho, G)$ is the set of homomorphisms from the graph H/ρ obtained from H by identifying vertices along ρ to G. Furthermore μ is the Möbius function. In their work Curticapean, Dell and Marx showed in a general theorem that a summation $\sum_{i=1}^{\ell} c_i \cdot \#\mathsf{Hom}(H_i, G)$ for *pairwise non-isomorphic graphs H_i* is #W[1]-hard if there is no upper bound on the treewidth of the pattern graphs H_i and fixed-parameter tractable otherwise, using a dichotomy for counting homomorphisms due to Dalmau and Jonsson [10]. Having this, one only has to show two properties of (1) to obtain the dichotomy for #Emb. First, one has to show that a high matching number of H implies that one of the graphs H/ρ has high treewidth and second, that two (or more) terms with high treewidth and isomorphic graphs H/ρ and H/σ do not cancel out (note that the Möbius function can be negative). As there is a closed form for the Möbius function over the partition lattice it was possible to show that whenever H/ρ and H/σ are isomorphic the sign of the Möbius function is equal.

1.1 Our results

The motivation of this work is the question whether the result of Curticapean, Dell and Marx can be generalized to construct a framework for the complexity classification of counting problems that can be expressed as the summation over homomorphisms and it turns out that this is possible whenever the summation is over a the lattice of a graphic matroid and the coefficients are evaluations of the Möbius function over the lattice, capturing not only embeddings but also locally injective homomorphisms.

In Section 3 we introduce what we call *graphically restricted homomorphisms*: Intuitively, a graphical restriction $\tau(H)$ of a graph H is a set of forbidden binary vertex identifications of H, modeled as a graph with vertex set $V(H)$ and edges along the binary constraints. We write $\tau\text{-}\mathcal{M}(H)$ as the set of all graphs obtained from H by contracting vertices along edges in $\tau(H)$ and deleting multiedges, excluding those that contain selfloops. Now a graphically restricted homomorphism from H to G with respect to τ is a homomorphism from H to G that maps every pair of vertices $u, v \in V(H)$ that are adjacent in $\tau(H)$ to different vertices in G. We write $\mathsf{Hom}_\tau(H, G)$ for the set of all graphically restricted homomorphisms w.r.t. τ from H to G and provide a complete complexity classification for counting graphically restricted homomorphisms:

▶ **Theorem 1** (Intuitive version). *Computing $\#\mathsf{Hom}_\tau(H, G)$ is fixed-parameter tractable when parameterized by $|V(H)|$ if the treewidth of every graph in $\tau\text{-}\mathcal{M}(H)$ is small. Otherwise the problem is #W[1]-hard.*

---

[3] Note that embeddings and subgraphs are equal up to automorphisms, that is, counting embeddings and counting subgraphs are essentially the same problem.

In particular, we obtain the following algorithmic result:

▶ **Theorem 2.** *There exists a deterministic algorithm that computes* $\#\mathsf{Hom}_\tau(H, G)$ *in time* $g(|V(H)|) \cdot |V(G)|^{\mathsf{tw}(\tau\text{-}\mathcal{M}(H))+1}$, *where g is a computable function and* $\mathsf{tw}(\tau\text{-}\mathcal{M}(H))$ *is the maximum treewidth of every graph in* $\tau\text{-}\mathcal{M}(H)$.

Having established the general dichotomy we observe that there exist graphical restrictions τ_{clique} and τ_{Li} such that $\mathsf{Hom}_{\tau_{\mathsf{clique}}}(H, G)$ is the set of all subgraph embeddings from H to G and $\mathsf{Hom}_{\tau_{\mathsf{Li}}}(H, G)$ is the set of all *locally injective homomorphisms* from H to G.

As a consequence we obtain a full complexity dichotomy for counting locally injective homomorphisms from small pattern graphs H to big host graphs G. To the best of the author's knowledge, this is the first result about the complexity of counting locally injective homomorphisms.

▶ **Corollary 3** (Intuitive version). *Computing the number of locally injective homomorphisms from H to G is fixed-parameter tractable when parameterized by $|V(H)|$ if the treewidth of every graph in $\tau_{\mathsf{Li}}\text{-}\mathcal{M}(H)$ is small. Otherwise the problem is* $\#\mathrm{W}[1]$*-hard.*

Furthermore, there exists a deterministic algorithm that computes this number in time $g(|V(H)|) \cdot |V(G)|^{\mathsf{tw}(\tau_{\mathsf{Li}}\text{-}\mathcal{M}(H))+1}$, *where g is a computable function and* $\mathsf{tw}(\tau_{\mathsf{Li}}\text{-}\mathcal{M}(H))$ *is the maximum treewidth of every graph in* $\tau_{\mathsf{Li}}\text{-}\mathcal{M}(H)$.

We then observe that – in contrast to subgraph embeddings – counting locally injective homomorphisms has "real" FPT cases, that is, cases that are fixed-parameter tractable but not polynomial time solvable under standard assumptions. We show this by restricting the pattern graph to be a tree:

▶ **Corollary 4.** *Computing the number of locally injective homomorphisms from a tree T to a graph G can be done in deterministic time $g(|V(T)|) \cdot |V(G)|^2$, that is, the problem is fixed-parameter tractable when parameterized by $|V(T)|$. On the other hand, the problem is* $\#\mathrm{P}$*-hard.*

To prove $\#\mathrm{P}$-hardness, we prove in an intermediate step that the subgraph counting problem remains hard when both graphs are restricted to be trees, which may be of independent interest:

▶ **Lemma 5.** *The problem of, given trees T_1 and T_2, computing the number of subtrees of T_2 that are isomorphic to T_1 is* $\#\mathrm{P}$*-hard.*

After that we generalize locally injective homomorphisms to homomorphisms that are injective in the r-neighborhood of every vertex and observe that those are also graphically restricted and consequently obtain a counting dichotomy as well. Due to space constraints, the corresponding section is deferred to the full version of the paper.

Finally, it turns out that all results can easily be extended to so-called *linear combinations* of graphically restricted homomorphisms. Here one gets as input graphs $H_1, \ldots, H_\ell$ together with positive coefficients $c_1, \ldots, c_\ell$ and a graph G and the goal is to compute

$$\sum_{i=1}^{\ell} c_i \cdot \#\mathsf{Hom}_{\tau_i}(H_i, G),$$

for graphical restrictions $\tau_1, \ldots, \tau_\ell$. This generalizes for example problems like counting all trees of size k in G or counting all locally injective homomorphisms from all graphs of size k to G or a combination thereof. We find out that, under some conditions, the dichotomy criteria transfer immediately to linear combinations:

▶ **Theorem 6** (Intuitive version). *Computing $\sum_{i=1}^{\ell} c_i \cdot \#\mathsf{Hom}_{\tau_i}(H_i, G)$ is fixed-parameter tractable when parameterized by $\max_i\{|V(H_i)|\}$ if the maximum treewidth of every graph in $\bigcup_i^{\ell} \tau_i\text{-}\mathcal{M}(H_i)$ is small. Otherwise, if additionally $|V(H_i)|$ has the same parity for every $i \in [\ell]$, the problem is $\#\mathrm{W}[1]$-hard.*

Furthermore we observe that this theorem is not true on the $\#\mathrm{W}[1]$-hardness side if we omit the parity condition. Due to space constraints, the section dealing with linear combinations is deferred to the full version as well.

1.2 Techniques

The main ingredients of the proofs of Theorem 1 and Theorem 2 are the complexity classification of linear combinations of homomorphisms due to Curticapean, Dell and Marx (see Lemma 3.5 and Lemma 3.8 in [7]) as well as a corollary of Rota's NBC Theorem (see e.g. Theorem 4 in [24]). In the first step we prove the following identity for the number of graphically restricted homomorphisms via Möbius inversion:

$$\#\mathsf{Hom}_\tau(H, G) = \sum_{\rho \geq \emptyset} \mu(\emptyset, \rho) \cdot \#\mathsf{Hom}(H/\rho, G),$$

where the sum is over elements of the lattice of flats of the graphical matroid given by $\tau(H)$ and H/ρ is the graph obtained by contracting the vertices of H along the flat ρ. After that we use Rota's Theorem to prove that none of the terms cancel out[4], despite the fact that the Möbius function can be negative. More precisely we show that whenever $H/\rho \cong H/\sigma$, we have that $\mathsf{rk}(\rho) = \mathsf{rk}(\sigma)$ and therefore, by Rota's Theorem, $\mathsf{sgn}(\mu(\emptyset, \rho)) = \mathsf{sgn}(\mu(\emptyset, \sigma))$.

The dichotomies for locally injective homomorphisms and homomorphisms that are injective in the r-neighborhood of every vertex are mere applications of the general theorem. For $\#\mathrm{P}$-hardness of the subgraph counting problem restricted to trees, we adapt the idea of the "skeleton graph" by Goldberg and Jerrum [15] and reduce directly from computing the permanent. To transfer this result to locally injective homomorphisms we use the well-known observation that locally injective homomorphisms from a tree to a tree are embeddings.

Finally, we prove the dichotomy for linear combinations of graphically restricted homomorphisms by taking a closer look at the proof of Theorem 1. Here, the parity constraint of the vertices of the graphs in the linear combination assures that there are no graphs H_i and H_j and elements ρ_i and ρ_j of the matroid lattices of $\tau_i(H_i)$ and $\tau_j(H_j)$ such that H_i/ρ_i and H_j/ρ_j are isomorphic but ρ_i and ρ_j have ranks of different parities. Using this observation, Theorem 6 can be proven in the same spirit as Theorem 1.

2 Preliminaries

First we will introduce some basic notions: Given a finite set S, we write $|S|$ or $\#S$ for the cardinality of S. Given a natural number ℓ we let $[\ell]$ be the set $\{1, \ldots, \ell\}$. Given a real number r we define the *sign* $\mathsf{sgn}(r)$ of r to be 1 if $r > 0$, 0 if $r = 0$ and -1 if $r < 0$.

A *poset* is a pair $(P, \leq)$ where P is a set and $\leq$ is a binary relation on P that is reflexive, transitive and anti-symmetric. Throughout this paper we will write $y \geq x$ if $x \leq y$. A *lattice* is a poset $(L, \leq)$ such that every pair of elements $x, y \in L$ has a *least upper bound* $x \vee y$ and a *greatest lower bound* $x \wedge y$ that satisfy:

---

[4] Here "cancel out" means that it could be possible that H/ρ and H/σ are isomorphic, but $\mu(\emptyset, \rho) = -\mu(\emptyset, \sigma)$ and all other H/ρ' are not isomorphic to H/ρ. In this case, the term $\#\mathsf{Hom}(H/\rho, G)$ would vanish in the above identity.

- $x \vee y \geq x$, $x \vee y \geq y$ and for all z such that $z \geq x$ and $z \geq y$ it holds that $z \geq x \vee y$.
- $x \wedge y \leq x$, $x \wedge y \leq y$ and for all z such that $z \leq x$ and $z \leq y$ it holds that $z \leq x \wedge y$.

Given a finite set S, a *partition* of S is a set ρ of pairwise disjoint subsets of S such that $\dot{\bigcup}_{s \in \rho} s = S$. We call the elements of ρ *blocks*. For two partitions ρ and σ we write $\rho \leq \sigma$ if every element of ρ is a subset of some element of σ. This binary relation is a lattice and called the *partition lattice* of S. We will in particular encounter lattices of graphic matroids in our proofs.

2.1 Matroids

We will follow the definitions of Chapt. 1 of the textbook of Oxley [22].

▶ **Definition 7.** A *matroid M* is a pair $(E, \mathcal{I})$ where E is a finite set and $\mathcal{I} \subseteq \mathcal{P}(E)$ such that

(1) $\emptyset \in \mathcal{I}$,
(2) if $A \in \mathcal{I}$ and $B \subseteq A$ then $B \in \mathcal{I}$, and
(3) if $A, B \in \mathcal{I}$ and $|B| < |A|$ then there exists $a \in A \setminus B$ such that $B \cup \{a\} \in \mathcal{I}$.

We call E the *ground set* and an element $A \in \mathcal{I}$ an *independent set*. A maximal independent set is called a *basis*. The *rank* $\mathsf{rk}(M)$ of M is the size of its bases[5].

Given a subset $X \subseteq E$ we define $\mathcal{I}|X := \{A \subseteq X \mid A \in \mathcal{I}\}$. Then $M|X := (X, \mathcal{I}|X)$ is also a matroid and called the restriction of M to X. Now the *rank* $\mathsf{rk}(X)$ of X is the rank of $M|X$. Equivalently, the rank of X is the size of the largest independent set $A \subseteq X$.

Furthermore we define the *closure* of X as follows:

$$\mathsf{cl}(X) := \{e \in E \mid \mathsf{rk}(X \cup \{e\}) = \mathsf{rk}(X)\}.$$

Note that by definition $\mathsf{rk}(X) = \mathsf{rk}(\mathsf{cl}(X))$. We say that X is a *flat* if $\mathsf{cl}(X) = X$. We denote $L(M)$ as the set of flats of M. It holds that $L(M)$ together with the relation of inclusion is a lattice, called the *lattice of flats* of M. The least upper bound of two flats X and Y is $\mathsf{cl}(X \cup Y)$ and the greatest lower bound is $X \cap Y$. It is known that the lattices of flats of matroids are exactly the geometric lattices[6] and we denote the set of those lattices as $\mathcal{L}$.

In Section 3 we take a closer look at (lattices of flats of) graphic matroids:

▶ **Definition 8.** Given a graph $H = (V, E) \in \mathcal{G}$, the *graphic matroid $M(H)$* has ground set E and a set of edges is independent if and only if it does not contain a cycle.

If H is connected then a basis of H is a spanning tree of H. If H consists of several connected components then a basis of $M(H)$ induces spanning trees for each of those. Every subset X of E induces a partition of the vertices of H where the blocks are the vertices of the connected components of $H|_X$ and it holds that

$$\mathsf{rk}(X) = |V(H)| - c(H|_X). \tag{2}$$

In particular, the flats of $M(H)$ correspond bijectively to the partitions of vertices of H into connected components as adding an element to X such that the rank does not change will not change the connected components, too. For convenience we will therefore abuse notation

---

[5] This is well-defined as every maximal independent set has the same size due to (3).
[6] For the purpose of this paper we do not need the definition of geometric lattices but rather the equivalent one in terms of lattices of flats and therefore omit it. We recommend e.g. Chapt. 3 of [32] and Chapt. 1.7 of [22] to the interested reader.

and say, given an element ρ of the lattice of flats of $M(H)$, that ρ partitions the vertices of H where the blocks are the vertices of the connected components of $H|_\rho$. The following observation will be useful in Section 3:

▶ **Lemma 9.** *Let $\rho, \sigma \in L(M(H))$ for a graph $H \in \mathcal{G}$. If the number of blocks of ρ and σ are equal then $\mathsf{rk}(\rho) = \mathsf{rk}(\sigma)$.*

Proof. Immediately follows from Equation (2). ◀

We denote H/ρ as the graph obtained from H by contracting the vertices of H that are in the same component of ρ and deleting multiedges (but keeping selfloops). As the vertices of H/ρ partition the vertices of H, we think of the vertices of H/ρ as subsets of vertices of H and call them *blocks*. Furthermore we write $[v]$ for the block containing v.

2.2 Graphs and homomorphisms

In this work all graphs are considered unlabeled and simple but may allow selfloops unless stated otherwise. We denote the set of all those graphs as $\mathcal{G}^\circ$. Furthermore we denote $\mathcal{G}$ as the set of all unlabeled and simple graphs without selfloops.

For a graph G we write n for the number of vertices $V(G)$ of G and m for the number of edges $E(G)$ of G. We denote $c(G)$ as the number of connected components of G. Furthermore, given a subset X of edges, we denote $G|_X$ as the graph with vertices $V(G)$ and edges X. Given a partition of vertices ρ of a graph H, we write H/ρ as the graph obtained from H by contracting the vertices of H that are in the same component of ρ and deleting multiedges (but keeping selfloops). As the vertices of H/ρ partition the vertices of H, we think of the vertices as subsets of vertices of H and call them *blocks*. Furthermore we write $[v]$ for the block containing v.

Given graphs H and G, a *homomorphism* from H to G is a mapping $\varphi : V(H) \to V(G)$ such that $\{u, v\} \in E(H)$ implies that $\{\varphi(u), \varphi(v)\} \in E(G)$. We denote $\mathsf{Hom}(H, G)$ as the set of all homomorphisms from H to G. A homomorphism is called *embedding* if it is injective and we denote $\mathsf{Emb}(H, G)$ as the set of all embeddings from H to G. An embedding from H to H is called an *automorphism* of H. We denote $\mathsf{Aut}(H)$ as the set of all automorphisms of H. Furthermore we let $\mathsf{Sub}(H, G)$ be the set of all subgraphs of G that are isomorphic to H. Then it holds that $\#\mathsf{Aut}(H) \cdot \#\mathsf{Sub}(H, G) = \#\mathsf{Emb}(H, G)$ (see e.g. [20]).

Given a set S and a function $\alpha : S \to \mathbb{Q}$, we define the *support* of α as follows:

$$\mathsf{supp}(\alpha) := \{s \in S \mid \alpha(s) \neq 0\}.$$

A graph parameter that will be of quite some importance to define the dichotomy criteria is the *treewidth* of a graph, capturing how "tree-like" a graph is. We do not need the explicit definition of treewidth which is therefore, together with some examples, deferred to the full version. However, throughout this paper we will often say that a set C of graphs has *bounded treewidth* meaning that there is a constant B such that the treewidth of every graph $H \in C$ is bounded by B.

2.3 Parameterized counting

We will mainly follow the definitions of Chapt. 14 of the textbook of Flum and Grohe [14]. A *parameterized counting problem* is a function $F : \{0, 1\}^* \to \mathbb{N}$ together with a polynomial-time computable *parameterization* $k : \{0, 1\}^* \to \mathbb{N}$. A parameterized counting problem is *fixed-parameter tractable* if there exists a computable function g such that it can be solved

in time $g(k(x)) \cdot |x|^{O(1)}$ for any input x. A *parameterized Turing reduction* from (F, k) to (F', k') is an FPT algorithm w.r.t. parameterization k with oracle (F', k') that on input x computes $F(x)$ and additionally satisfies that there exists a function g' such that for every oracle query y it holds that $k'(y) \leq g(k(x))$. A parameterized counting problem (F, k) is #W[1]-hard if there exists a parameterized Turing reduction from #k-clique to (F, k), where #k-clique is the problem of, given a graph G and a parameter k, computing the number of cliques of size k in G[7]. Under standard assumptions (e.g. under the exponential time hypothesis) #W[1]-hard problems are not fixed-parameter tractable.

The following two parameterized counting problems will be of particular importance in this work: Given a class of graphs $C \subseteq \mathcal{G}$, #Hom(C) (#Emb(C)) is the problem of, given a graph $H \in C$ and a graph $G \in \mathcal{G}$, computing #Hom(H, G) (#Emb(H, G)). Both problems are parameterized by #$V(H)$. Their complexity has already been classified:

▶ **Theorem 10** ([10]). *Let C be a recursively enumerable class of graphs. If C has bounded treewdith then* #Hom(C) *can be solved in polynomial time. Otherwise* #Hom(C) *is* #W[1]-*hard.*

▶ **Theorem 11** ([9]). *Let C be a recursively enumerable class of graphs. If C has bounded matching number then* #Emb(C) *can be solved in polynomial time. Otherwise* #Emb(C) *is* #W[1]-*hard.*

Recall that "bounded treewidth (matching number)" means that there is a constant B such that the treewidth (size of the largest matching) of any graph in C is bounded by B.

2.4   Linear combinations of homomorphisms and Möbius inversion

Curticapean, Dell and Marx [7] introduced the following parameterized counting problem:

▶ **Definition 12** (Linear combinations of homomorphisms). Let $\mathcal{A}$ be a set of functions $a : \mathcal{G} \to \mathbb{Q}$ with finite support[8]. We define the parameterized counting problem #Hom($\mathcal{A}$) as follows:

Given $a \in \mathcal{A}$ and $G \in \mathcal{G}$, compute

$$\sum_{H \in \text{supp}(a)} a(H) \cdot \#\text{Hom}(H, G), \quad \text{parameterized by} \max_{H \in \text{supp}(a)} \#V(H).$$

Note that this problem generalizes #Hom(C). The following theorem will be the foundation of all complexity results in this paper:

▶ **Theorem 13** ([7], Lemma 3.5 and Lemma 3.8). *If $\mathcal{A}$ has bounded treewidth then* #Hom($\mathcal{A}$) *can be solved in time $g(|\text{supp}(\alpha)|) \cdot n^{O(1)}$ on input (α, G) where $n = |V(G)|$ and g is a computable function. Otherwise the problem is* #W[1]-*hard.*

In their paper, the authors show how this result can be used to give a much simpler proof of Theorem 11. The idea is that every problem #Emb(C) is equivalent to a problem #Hom($\mathcal{A}$). As all proofs in this work are in the same flavour, we will outline the technique here, using #Emb(C) as an example. Therefore, we first need to introduce the so called Möbius inversion (we recommend reading [26] for a more detailed introduction):

---

[7] For a more detailed introduction to #W[1] we recommend [14] to the interested reader.
[8] We can also think of $\mathcal{A}$ being a set of lists.

▶ **Definition 14.** Let $(P, \leq)$ be a poset and $h : P \to \mathbb{C}$ be a function. Then the *zeta transformation* ζh is defined as follows:

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

▶ **Theorem 15** (Möbius inversion, see [26] or [24]). *Let $(P, \leq)$ and h as in Definition 14. Then there is a function $\mu_P : P \times P \to \mathbb{Z}$ such that for all $\sigma \in P$ it holds that*

$$h(\sigma) = \sum_{\rho \geq \sigma} \mu_P(\sigma, \rho) \cdot \zeta h(\rho) \,.$$

μ_P *is called the* Möbius function.

The following identity is due to Lovász [20]:

$$\#\mathsf{Hom}(H/\sigma, G) = \sum_{\rho \geq \sigma} \#\mathsf{Emb}(H/\rho, G) \,,$$

where σ and ρ are partitions of vertices of H and $\geq$ is the partition lattice of H. Now Möbius inversion yields the following identity [20]:

$$\#\mathsf{Emb}(H, G) = \sum_{\rho \geq \emptyset} \mu(\emptyset, \rho) \cdot \#\mathsf{Hom}(H/\rho, G) \,,$$

where μ is the Möbius function over the partition lattice. Therefore, for every class of graphs C, there is a family of functions with finite support $\mathcal{A}$ such that $\#\mathsf{Emb}(C)$ and $\#\mathsf{Hom}(\mathcal{A})$ are the same problems. Now Curticapean, Dell and Marx show that C has unbounded matching number if and only if $\mathcal{A}$ has unbounded treewidth. The critical point in this proof was to show that the sign of $\mu(\emptyset, \rho)$ only depends on the number of blocks of ρ, which implies that for two isomorphic graphs H_1 and H_2, the terms $\#\mathsf{Hom}(H_1, G)$ and $\#\mathsf{Hom}(H_2, G)$ have the same sign in the above identity and therefore do not cancel out in the homomorphism basis. As there is a closed form for $\mu(\emptyset, \rho)$[9], the information about the sign could easily be extracted.

The motivation of this work is the question whether this can be made more general and it turns out that a corollary of Rota's NBC Theorem [24] (see also [1]) captures exactly what we need:

▶ **Theorem 16** (See e.g. Theorem 4 in [24]). *Let L be a geometric lattice with unique minimal element $\perp$ and let ρ be an element of L. Then it holds that*

$$\mathsf{sgn}(\mu_L(\perp, \rho)) = (-1)^{\mathsf{rk}(\rho)} \,.$$

In the following we will show that combining Rota's Theorem and the dichotomy for counting linear combinations of homomorphisms yields complete complexity classifications for the problems of counting those restricted homomorphisms that induce a Möbius inversion over the lattice of a graphic matroid, which are known to be geometric, when transformed into the homomorphism basis. Those include embeddings as well as locally injective homomorphisms.

---

[9] Here it is crucial that μ is the Möbius function over the (complete) partition lattice.

3 Graphically restricted homomorphisms

In the following we write $\emptyset$ for the minimal element of a matroid lattice.

▶ **Definition 17.** A *graphical restriction* is a computable mapping τ that maps a graph $H \in \mathcal{G}$ to a graph $H' \in \mathcal{G}$ such that $V(H) = V(H')$, that is, τ only modifies edges of H. We denote the set of all graphical restrictions as T. Given graphs H and G and a graphical restriction τ, we define the set of *graphically restricted homomorphisms* w.r.t. τ from H to G as follows:

$$\mathsf{Hom}_\tau(H, G) := \{\varphi \in \mathsf{Hom}(H, G) \mid \forall u, v \in V(H) : \{u, v\} \in E(\tau(H)) \Rightarrow \varphi(u) \neq \varphi(v)\}.$$

Given a recursively enumerable class of graphs $C \subseteq \mathcal{G}$, we define the parameterized counting problem $\#\mathsf{Hom}_\tau(C)$ as follows: Given a graph $H \in C$ and a graph $G \in \mathcal{G}$, we parameterize by $|V(H)|$ and wish to compute $\#\mathsf{Hom}_\tau(H, G)$.

Assume for example that τ_{clique} maps a graph H to the complete graph with vertices $V(H)$. Then one can easily verify that $\mathsf{Hom}_{\tau_{\mathsf{clique}}}(H, G) = \mathsf{Emb}(H, G)$.

The following lemma is an application of Möbius inversion (and slightly generalizes [20]). Due to space constraints, the proof is deferred to the full version.

▶ **Lemma 18.** *Let τ be a graphical restriction. Then for all graphs $H \in \mathcal{G}^\circ$ and $G \in \mathcal{G}$ it holds that*

$$\#\mathsf{Hom}_\tau(H, G) = \sum_{\rho \geq \emptyset} \mu(\emptyset, \rho) \cdot \#\mathsf{Hom}(H/\rho, G), \tag{3}$$

where $\leq$ and μ are the relation and the Möbius function of the lattice $L(M(\tau(H)))$.

Intuitively, we will now show that counting graphically restricted homomorphisms from H to G is hard if we can "glue" vertices of H together along edges of $\tau(H)$ such that the resulting graph has no selfloops and high treewidth. We will capture this intuition formally:

▶ **Definition 19.** Let $H \in \mathcal{G}$ be a graph and let τ be a graphical restriction. A graph $H' \in \mathcal{G}^\circ$ obtained from H by contracting pairs of vertices u and v such that $\{u, v\} \in E(\tau(H))$ and deleting multiedges (but keeping selfloops) is called a τ-*contraction* of H. If additionally $H' \in \mathcal{G}$, that is, the contraction did not yield selfloops, we call H' a τ-*minor* of H. We denote the set of all τ-minors of H as $\tau\text{-}\mathcal{M}(H)$ and given a class of graphs $C \subseteq \mathcal{G}$ we denote the set of all τ-minors of all graphs in C as $\tau\text{-}\mathcal{M}(C)$.

Finally, we can classify the complexity of counting graphically restricted homomorphisms along the treewidth of their τ-minors:

▶ **Theorem 20** (Theorem 1 and Theorem 2, restated). *Let τ be a graphical restriction and let $C \subseteq \mathcal{G}$ be a recursively enumerable class of graphs. Then $\#\mathsf{Hom}_\tau(C)$ is FPT if $\tau\text{-}\mathcal{M}(C)$ has bounded treewidth and $\#\mathrm{W}[1]$-hard otherwise. Furthermore, given $H, G \in \mathcal{G}$, there exists a deterministic algorithm that computes $\#\mathsf{Hom}_\tau(H, G)$ in time*

$$g(|V(H)|) \cdot |V(G)|^{\mathsf{tw}(\tau\text{-}\mathcal{M}(H))+1},$$

where g is a computable function.

Proof. By Lemma 18 we have that

$$\#\mathsf{Hom}_\tau(H, G) = \sum_{\rho \geq \emptyset} \mu(\emptyset, \rho) \cdot \#\mathsf{Hom}(H/\rho, G) .$$

Now, as G has no selfloops, a term $\#\mathsf{Hom}(H/\rho, G)$ is zero whenever H/ρ has a selfloop. Consequently, for every non-zero term $\#\mathsf{Hom}(H/\rho, G)$, it holds that $H/\rho \in \tau\text{-}\mathcal{M}(H)$. Therefore, by Lemma 3.5 in [7], we obtain an algorithm computing $\#\mathsf{Hom}_\tau(H, G)$ in time $g(|V(H)|) \cdot |V(G)|^{\mathsf{tw}(\tau\text{-}\mathcal{M}(H))+1}$, for a computable function g. This immediately implies that the problem $\#\mathsf{Hom}_\tau(C)$ is fixed-parameter tractable if $\tau\text{-}\mathcal{M}(C)$ has bounded treewidth. It remains to show that $\#\mathsf{Hom}_\tau(C)$ is $\#\mathrm{W}[1]$-hard otherwise. By condensing all terms $\#\mathsf{Hom}(H/\rho, G)$ and $\#\mathsf{Hom}(H/\sigma, G)$ where H/ρ and H/σ are isomorphic, it follows that there exist coefficients $c_H[H']$ for every $H' \in \tau\text{-}\mathcal{M}(H)$ such that

$$\#\mathsf{Hom}_\tau(H, G) = \sum_{H' \in \tau\text{-}\mathcal{M}(H)} c_H[H'] \cdot \#\mathsf{Hom}(H', G) .$$

We will now show that none of the $c_H[H']$ is zero: It holds that

$$c_H[H'] = \sum_{\substack{\rho \geq \emptyset \\ H' \cong H/\rho}} \mu(\emptyset, \rho) . \tag{4}$$

Consider ρ and ρ' such that $H/\rho \cong H/\rho' \cong H'$. It follows that $\mathsf{rk}(\rho) = |V(H)| - c(H/\rho) = |V(H)| - c(H') = |V(H)| - c(H/\rho') = \mathsf{rk}(\rho')$. Now, as the lattice of $M(\tau(H))$ is geometric, we can apply the corollary of Rota's NBC Theorem (Theorem 16) and obtain that $\mathsf{sgn}(\mu(\emptyset, \rho)) = (-1)^{\mathsf{rk}(\rho)} = (-1)^{\mathsf{rk}(\rho')} = \mathsf{sgn}(\mu(\emptyset, \rho'))$. Consequently every term in Equation (4) has the same sign and therefore $c_H[H'] \neq 0$. Now we define a function $a_H : \mathcal{G} \to \mathbb{Q}$ as follows

$$a_H(F) := \begin{cases} c_H[F] & \text{if } F \in \tau\text{-}\mathcal{M}(H) \\ 0 & \text{otherwise} \end{cases}$$

and we set $\mathcal{A}_C = \{a_H \mid H \in C\}$. Then the problems $\#\mathsf{Hom}(\mathcal{A}_C)$ and $\#\mathsf{Hom}_\tau(C)$ are equivalent w.r.t. parameterized turing reductions. As $c_H[H'] \neq 0$ for every $H' \in \tau\text{-}\mathcal{M}(H)$ it follows that $\mathcal{A}_C$ has unbounded treewidth if and only if $\tau\text{-}\mathcal{M}(C)$ has unbounded treewidth. We conclude by Theorem 13 that $\#\mathsf{Hom}_\tau(C)$ is $\#\mathrm{W}[1]$-hard in this case. ◀

4 Locally injective homomorphisms

In this section we are going to apply the general dichotomy theorem to the concrete case of counting locally injective homomorphisms. A homomorphism φ from H to G is *locally injective* if for every $v \in V(H)$ it holds that $\varphi|_{N(v)}$ is injective. We denote $\mathsf{Li\text{-}Hom}(H, G)$ as the set of all locally injective homomorphisms from H to G and we define the corresponding counting problem $\#\mathsf{Li\text{-}Hom}(C)$ for a class of graphs $C \subseteq \mathcal{G}$ as follows: Given graphs $H \in C$ and $G \in \mathcal{G}$, compute $\#\mathsf{Li\text{-}Hom}(H, G)$. The parameter is $|V(H)|$. Locally injective homomorphisms have already been studied by Nešetřil in 1971 [21] and were applied in the context of distance constrained labelings of graphs (see [13] for an overview). As well as subgraph embeddings, locally injective homomorphisms are graphically restricted homomorphisms. Due to space constraints, the proofs of this section are deferred to the full version.

▶ **Lemma 21.** *Let $H \in \mathcal{G}$ be a graph and let $\tau_{\mathsf{Li}}(H) = (V(H), E_{\mathsf{Li}}(H))$ be a graphical restriction defined as follows: $E_{\mathsf{Li}}(H) = \{\{u, w\} \mid u \neq w \wedge \exists v : \{u, v\}, \{w, v\} \in E(H)\}$. Then for all $G \in \mathcal{G}$ it holds that $\mathsf{Hom}_{\tau_{\mathsf{Li}}}(H, G) = \mathsf{Li\text{-}Hom}(H, G)$.*

We continue by stating the dichotomy for counting locally injective homomorphisms.

▶ **Corollary 22** (Corollary 3, restated). *Let $C \subseteq \mathcal{G}$ be a recursively enumerable class of graphs. Then #Li-Hom(C) is FPT if τ_{Li}-$\mathcal{M}(C)$ has bounded treewidth and #W[1]-hard otherwise. Furthermore, there exists a computable function g and a deterministic algorithm that computes #Li-Hom(H, G) in time*

$$g(|V(H)|) \cdot |V(G)|^{\mathsf{tw}(\tau_{\mathsf{Li}}\text{-}\mathcal{M}(H))+1} \,.$$

We give an example for a hard instance of the problem: Let W_k be the "windmill" graph of size k, i.e., the graph with vertices $a, v_1, \ldots, v_k, w_1, \ldots, w_k$ and edges $\{a, v_i\}, \{v_i, w_i\}$ and $\{w_i, a\}$ for each $i \in [k]$. Furthermore we let $\mathcal{W}$ be the set of all W_k for $k \in \mathbb{N}$.

▶ **Corollary 23.** *#Li-Hom($\mathcal{W}$) is #W[1]-hard.*

In contrast to embeddings where every FPT case is also polynomial time solvable, there are "real" FPT cases when it comes to locally injective homomorphisms. Let $\mathcal{T} \subseteq \mathcal{G}$ be the class of all trees. Counting locally injective homomorphisms from those graphs is fixed-parameter tractable:

▶ **Corollary 24.** *#Li-Hom($\mathcal{T}$) is FPT. In particular, there is a deterministic algorithm that computes #Li-Hom(T, G) for a tree T in time*

$$g(|V(T)|) \cdot |V(G)|^2 \,,$$

where g is a computable function.

On the other hand #Li-Hom($\mathcal{T}$) is unlikely to have a polynomial time algorithm.

▶ **Lemma 25.** *#Li-Hom($\mathcal{T}$) is #P-hard.*

The proof of this theorem as well as a brief introduction to classical counting complexity can be found in the full version. Corollary 4 follows then from Corollary 24 and Lemma 25.

5 Conclusion and further work

We have shown that various parameterized counting problems can be expressed as a linear combination of homomorphisms over the lattice of graphic matroids, implying immediate complexity classifications along with fixed-parameter tractable algorithms for the positive cases. These results can be obtained without using often cumbersome tools like "gadgeting" or interpolation and relies only on the knowledge of the problem of counting homomorphisms and the comprehension of the cancellation behaviour when transforming a problem into this "homomorphism basis". The latter, in turn, was nothing more than a question about the sign of the Möbius function, which was answered by Rota's Theorem.

This framework, however, still has limits: It seems that, e.g., neither induced subgraphs nor edge-injective homomorphisms [8] are graphically restricted. Indeed, both can be expressed as a sum of homomorphisms over (non-geometric) lattices but the problem is that there are isomorphic terms with different signs in both cases. This suggests that a better understanding of the Möbius function over those lattices could yield even more general complexity classifications of parameterized counting problems.

References

1 Andreas Blass and Bruce E. Sagan. Möbius functions of lattices. *Advances in mathematics*, 127(1):94–123, 1997.

2 Graham R. Brightwell and Peter Winkler. Note on counting eulerian circuits. *CoRR*, cs.CC/0405067, 2004. URL: `http://arxiv.org/abs/cs.CC/0405067`.

3 Andrei A. Bulatov. A dichotomy theorem for nonuniform CSPs. *CoRR*, abs/1703.03021, 2017. URL: `http://arxiv.org/abs/1703.03021`.

4 Jin-yi Cai, Pinyan Lu, and Mingji Xia. Holographic algorithms by fibonacci gates and holographic reductions for hardness. In *49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, October 25-28, 2008, Philadelphia, PA, USA*, pages 644–653, 2008. `doi:10.1109/FOCS.2008.34`.

5 Jin-yi Cai, Pinyan Lu, and Mingji Xia. Computational complexity of holant problems. *SIAM J. Comput.*, 40(4):1101–1132, 2011. `doi:10.1137/100814585`.

6 Radu Curticapean. Counting matchings of size k is #W[1]-hard. In *Proceedings of the 40th International Colloquium on Automata, Languages, and Programming, ICALP, Part I*, pages 352–363, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg. `doi:10.1007/978-3-642-39206-1_30`.

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

8 Radu Curticapean, Holger Dell, and Marc Roth. Counting edge-injective homomorphisms and matchings on restricted graph classes. In *34th Symposium on Theoretical Aspects of Computer Science, STACS 2017, March 8-11, 2017, Hannover, Germany*, pages 25:1–25:15, 2017. `doi:10.4230/LIPIcs.STACS.2017.25`.

9 Radu Curticapean and Dániel Marx. Complexity of counting subgraphs: Only the boundedness of the vertex-cover number counts. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 130–139, 2014. `doi:10.1109/FOCS.2014.22`.

10 Víctor Dalmau and Peter Jonsson. The complexity of counting homomorphisms seen from the other side. *Theor. Comput. Sci.*, 329(1-3):315–323, 2004. `doi:10.1016/j.tcs.2004.08.008`.

11 Martin E. Dyer, Leslie Ann Goldberg, and Mark Jerrum. An approximation trichotomy for Boolean #CSP. *J. Comput. Syst. Sci.*, 76(3-4):267–277, 2010. `doi:10.1016/j.jcss.2009.08.003`.

12 Tomás Feder and Moshe Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through datalog and group theory. *SIAM J. Comput.*, 28(1):57–104, 1998. `doi:10.1137/S0097539794266766`.

13 Jirí Fiala and Jan Kratochvíl. Locally constrained graph homomorphisms – structure, complexity, and applications. *Computer Science Review*, 2(2):97–111, 2008. `doi:10.1016/j.cosrev.2008.06.001`.

14 Jörg Flum and Martin Grohe. The parameterized complexity of counting problems. *SIAM J. Comput.*, 33(4):892–922, 2004. `doi:10.1137/S0097539703427203`.

15 Leslie Ann Goldberg and Mark Jerrum. Counting unlabelled subtrees of a tree is #P-complete. 1999.

16 Mark Jerrum and Alistair Sinclair. Approximating the permanent. *SIAM J. Comput.*, 18(6):1149–1178, 1989. `doi:10.1137/0218077`.

17 Pieter W. Kasteleyn. Graph theory and crystal physics. In *Graph Theory and Theoretical Physics*, pages 43–110. Academic Press, 1967.

18   Richard E. Ladner. On the structure of polynomial time reducibility. *J. ACM*, 22(1):155–171, 1975. `doi:10.1145/321864.321877`.

19   Bingkai Lin. The parameterized complexity of *k*-biclique. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 605–615, 2015. `doi:10.1137/1.9781611973730.41`.

20   László Lovász. Operations with structures. *Acta Mathematica Hungarica*, 18(3-4):321–328, 1967.

21   Jaroslav Nesetril. Homomorphisms of derivative graphs. *Discrete Mathematics*, 1(3):257–268, 1971. `doi:10.1016/0012-365X(71)90014-8`.

22   James G. Oxley. *Matroid theory*. Oxford University Press, 1992.

23   Arash Rafiey, Jeff Kinne, and Tomás Feder. Dichotomy for digraph homomorphism problems. *CoRR*, abs/1701.02409, 2017. URL: `http://arxiv.org/abs/1701.02409`.

24   Gian-Carlo Rota. On the foundations of combinatorial theory I. Theory of möbius functions. *Probability theory and related fields*, 2(4):340–368, 1964.

25   Thomas J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the 10th Annual ACM Symposium on Theory of Computing, May 1-3, 1978, San Diego, California, USA*, pages 216–226, 1978. `doi:10.1145/800133.804350`.

26   Richard P Stanley. Enumerative combinatorics: Volume 1. 2011.

27   Harold N. V. Temperley and Michael E. Fisher. Dimer problem in statistical mechanics – an exact result. *Philosophical Magazine*, 6(68):1478–6435, 1961.

28   Salil P. Vadhan. The complexity of counting in sparse, regular, and planar graphs. *SIAM J. Comput.*, 31(2):398–427, 2001. `doi:10.1137/S0097539797321602`.

29   Leslie G. Valiant. The complexity of computing the permanent. *Theor. Comput. Sci.*, 8:189–201, 1979. `doi:10.1016/0304-3975(79)90044-6`.

30   Leslie G. Valiant. Accidental algorithms. In *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006), 21-24 October 2006, Berkeley, California, USA, Proceedings*, pages 509–517, 2006. `doi:10.1109/FOCS.2006.7`.

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

32   Dominic J.A. Welsh. *Matroid theory*. Courier Corporation, 2010.

33   Mingji Xia, Peng Zhang, and Wenbo Zhao. Computational complexity of counting problems on 3-regular planar graphs. *Theor. Comput. Sci.*, 384(1):111–125, 2007. `doi:10.1016/j.tcs.2007.05.023`.

34   Dmitriy Zhuk. The proof of CSP dichotomy conjecture. *CoRR*, abs/1704.01914, 2017. URL: `http://arxiv.org/abs/1704.01914`.