

On Counting (Quantum-)Graph Homomorphisms in Finite Fields of Prime Order

J. A. Gregor Lagodzinski ✉ 

Hasso Plattner Institute, University of Potsdam, Potsdam, Germany

Andreas Göbel ✉ 

Hasso Plattner Institute, University of Potsdam, Potsdam, Germany

Katrin Casel ✉ 

Hasso Plattner Institute, University of Potsdam, Potsdam, Germany

Tobias Friedrich ✉ 

Hasso Plattner Institute, University of Potsdam, Potsdam, Germany

Abstract

We study the problem of counting the number of homomorphisms from an input graph G to a fixed (quantum) graph \bar{H} in any finite field of prime order \mathbb{Z}_p . The subproblem with graph H was introduced by Faben and Jerrum [ToC'15] and its complexity is still uncharacterised despite active research, e.g. the very recent work of Focke, Goldberg, Roth, and Zivný [SODA'21]. Our contribution is threefold.

First, we introduce the study of quantum graphs to the study of modular counting homomorphisms. We show that the complexity for a quantum graph \bar{H} collapses to the complexity criteria found at dimension 1: graphs. Second, in order to prove cases of intractability we establish a further reduction to the study of bipartite graphs. Lastly, we establish a dichotomy for all bipartite $(K_{3,3}\{e\}, \text{domino})$ -free graphs by a thorough structural study incorporating both local and global arguments. This result subsumes all results on bipartite graphs known for all prime moduli and extends them significantly. Even for the subproblem with $p = 2$ this establishes new results.

2012 ACM Subject Classification Theory of computation \rightarrow Problems, reductions and completeness; Mathematics of computing \rightarrow Discrete mathematics

Keywords and phrases Algorithms, Theory, Quantum Graphs, Bipartite Graphs, Graph Homomorphisms, Modular Counting, Complexity Dichotomy

Digital Object Identifier 10.4230/LIPIcs.ICALP.2021.86

Category Track A: Algorithms, Complexity and Games

Related Version full version hosted on arXiv

Full Version: <https://arxiv.org/abs/2011.04827>

Acknowledgements The authors would like to thank Holger Dell for bringing [12] to their attention. Our gratitude also goes to Jacob Focke and Marc Roth for their valuable insights on partially surjective homomorphisms and for pointing out a mistake in a previous version.

1 Introduction

The study of graph homomorphisms represents one of the classic bodies of work in both discrete mathematics and computer science but remains a very active research area. These homomorphisms play a crucial role in the study of graph limits and networks [4, 19, 20, 47], in the study of databases [11, 33, 39, 40], and in the study of spin-systems in statistical physics [2, 5]. Formally, a graph-homomorphism from G to H is a map from $V(G)$ to $V(H)$ that preserves edges. Many classic problems studied in computer science can be expressed with graph homomorphisms. Examples range from the *decision* problem of determining the



© J. A. Gregor Lagodzinski, Andreas Göbel, Katrin Casel, and Tobias Friedrich;
licensed under Creative Commons License CC-BY 4.0

48th International Colloquium on Automata, Languages, and Programming (ICALP 2021).

Editors: Nikhil Bansal, Emanuela Merelli, and James Worrell; Article No. 86; pp. 86:1–86:15

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

chromatic number of a graph, through the problem of *counting* the number of independent sets, to the problem of *counting* the number of k -colourings using all k colours. The latter can be expressed by a *linear combination* of the number of graph homomorphisms to a set of non-isomorphic graphs.

Graph homomorphisms are a prime example of a very general class of problems that frequently yields complexity dichotomies with structural characterisations, where the properties of a graph implying (in)tractability are easily computable. However, the dichotomy itself is hard to establish and by Ladner [37] not obvious to exist. Hell and Nešetřil studied the *decision* problem $\text{HOMSTO}H$ with fixed image graph H , that asks whether there exists a homomorphism from an input graph G to H . In [35] they showed that the problem $\text{HOMSTO}H$ can be solved in polynomial time if H contains a loop or is bipartite; otherwise it is NP-complete. Dyer and Greenhill introduced the *counting* problem $\#\text{HOMSTO}H$ with fixed image graph H , that asks for the number of homomorphisms from an input graph G to H . In their seminal work [17] they showed that $\#\text{HOMSTO}H$ can be solved in polynomial time if the connected components of H are complete bipartite graphs or reflexive complete graphs; otherwise it is $\#\text{P}$ -complete.

Lovász [38] observed that there are many graph parameters that can only be expressed by a linear combination of computational problems $\#\text{HOMSTO}H$ for a set of at least two graphs $H \in \mathcal{H}$. Examples are the class of *vertex surjective homomorphisms* and *compactness* studied in this context by Focke, Goldberg, and Zivný [24]. Lovász [38] introduced the notion of a *quantum graph* for a linear combination of finitely many graphs called its *constituents*. We refer by the dimension of a quantum graph to its number of constituents and find the set of graphs at dimension 1. With every increase of dimension, the set of graph parameters expressible by $\#\text{HOMSTO}H$ increases as well. For a quantum graph \bar{H} the counting problem $\#\text{HOMSTO}\bar{H}$ denotes then *linear combination* of problems $\#\text{HOMSTO}H$ for all constituents H of \bar{H} . Chen, Curticapean and Dell [12] studied the complexity of $\#\text{HOMSTO}\bar{H}$ and showed that the complexity is inherited from the complexity of $\#\text{HOMSTO}H$ for all constituents H of \bar{H} , which is given by the criterion of Dyer and Greenhill. Motivated by this strong connection, Chen et al. raised the question of whether techniques based on quantum graphs can advance the state of the art of open problems regarding modular counting homomorphisms.

We study the complexity of the problem $\#_p\text{HOMSTO}\bar{H}$ for any prime p and answer the question of Chen et al. in the affirmative, where the problem $\#_p\text{HOMSTO}\bar{H}$ asks for the value of $\#\text{HOMSTO}\bar{H}$ in the finite field \mathbb{Z}_p . Our contribution is threefold. First, we obtain results for the whole class of quantum graphs by showing that the complexity of $\#_p\text{HOMSTO}\bar{H}$ is inherited from the complexity $\#_p\text{HOMSTO}H$. Second, we reduce the study of $\#_p\text{HOMSTO}H$ to a study of bipartite graphs by establishing a reduction to a restricted homomorphism problem. Finally, we employ a structural analysis on the set of $(K_{3,3} \setminus \{e\}, \text{domino})$ -free graphs and establish a dichotomy for these.

The line of research on modular counting homomorphisms was initiated with the study of the parity of $\#\text{HOMSTO}H$ by Faben and Jerrum [21]. Despite the clear picture on the non-modular version $\#\text{HOMSTO}H$ the modulus implies additional cases of tractability as structures in H implying intractability for $\#\text{HOMSTO}H$ get ‘cancelled’ when counting in a finite field \mathbb{Z}_p . Faben and Jerrum [21] showed that automorphisms of order p capture a subset of these ‘cancellations’ and reduced the study to a structural analysis of parameter graphs H that do not enjoy such automorphisms. These graphs are called *order p reduced*. In particular, for $p = 2$ they conjectured that automorphisms of order 2 capture all cancellations and that $\#_2\text{HOMSTO}H$ for an order 2 reduced graph enjoys the same complexity criterion as the non-modular version $\#\text{HOMSTO}H$ given by Dyer and Greenhill. Despite a growing

	Mod	Trees	Cactus	Square-free	K_4 -minor-free	$(K_{3,3} \setminus \{e\}, \text{domino})$ -free
Faben and Jerrum [21]	2	×				
Göbel et al. [26]	2	×	×			
Göbel et al. [27]	2	×		×		
Focke et al. [23]	2	×	×	(×)	×	
Göbel et al. [28]	p	×				
Kazeminia and Bulatov [36]	p	×		×		
This paper	p	×	×	×		×

■ **Table 1** History of the study of $\#_p \text{HOMSTO}H$ on **bipartite** graphs H . (Note that the complexity study can be restricted to bipartite graphs by the bipartization result of this paper.) Crosses denote that the result incorporates the dichotomy for the graphclass, a p denotes that the result holds for all primes. Parenthesis denote that the result is not intrinsic but given by additional argumentation.

line of research by Göbel, Goldberg, and Richerby [26, 27] and the very recent work of Focke, Goldberg, Roth, and Zivný [23] the conjecture remains open. The body of work is dominated by a study of structures as the modulus commands incorporating not only the local but also global properties of the graph H .

The research on $\#_p \text{HOMSTO}H$ for arbitrary primes p was already suggested by Faben and Jerrum [21] as they showed that their results concerning automorphisms of order p apply for any prime p . However, Valiant [48] showed the existence of computational counting problems that enjoy a change of complexity with respect to different moduli. Therefore, a uniform complexity criterion for $\#_p \text{HOMSTO}H$ would emphasize the special role of graph homomorphisms even more. The study of $\#_p \text{HOMSTO}H$ was finally initiated by Göbel, Lagodzinski, and Seidel [28] and followed by Kazeminia and Bulatov [36]. In light of the richer structure due to the higher moduli far less is known about the complexity of $\#_p \text{HOMSTO}H$ compared to $\#_2 \text{HOMSTO}H$. Even though Faben and Jerrum [21] as well as Göbel et al. [28] considered an extension of the conjecture to all prime moduli and the results so far suggest it, no one has gone that far yet. We illustrate the individual contributions and the state of the art in Table 1.

1.1 Contribution

We establish a plethora of technical results, which we believe to be a major asset to future works on the complexity of $\#_p \text{HOMSTO}H$ and may be of independent interest to different lines of research. The main contributions are given in the following and discussed in more depth in the subsequent subsection.

Quantum Homomorphisms

We introduce the study of quantum graphs to the study of $\#_p \text{HOMSTO}H$. For any quantum graph \bar{H} we find that $\#_p \text{HOMSTO}\bar{H}$ is equivalent to $\#_p \text{HOMSTO}\bar{H}'$, where \bar{H}' is a quantum graph whose constituents are order p reduced with coefficients in $\mathbb{Z}_p^* = \mathbb{Z}_p \setminus \{0\}$. We call these constituents the p -constituents of \bar{H}' . Focusing on these ‘reduced’ quantum graphs we obtain the following inheritance theorem.

► **Theorem 1.1.** *Let p be a prime and $\bar{H} = \sum_{H \in \mathcal{H}} \alpha_H H$ be a quantum graph with p -constituents $\mathcal{H} = \{H_1, \dots, H_r\}$ that are order p reduced pairwise non-isomorphic graphs and a set of associated constants $\{\alpha_H\}_{H \in \mathcal{H}}$ that are in \mathbb{Z}_p^* . Then,*

- if there exists a graph $H \in \mathcal{H}$ such that $\#_p \text{HOMSTO}H$ is $\#_p \text{P-hard}$, then $\#_p \text{HOMSTO}\bar{H}$ is $\#_p \text{P-hard}$;
- if, for all $H \in \mathcal{H}$, $\#_p \text{HOMSTO}H$ is solvable in polynomial time, then $\#_p \text{HOMSTO}\bar{H}$ is also solvable in polynomial time.

This shows that the complexity of $\#_p \text{HOMSTO}\bar{H}$ collapses to the complexity of $\#_p \text{HOMSTO}H$. Even though the set of graph parameters expressible by $\#_p \text{HOMSTO}\bar{H}$ is arbitrarily larger compared to the parameters expressible by $\#_p \text{HOMSTO}H$, the complexity behaviour is captured at dimension $r = 1$, i.e. graphs.

In the same spirit, we show that the reduction technique applied to show Theorem 1.1 yields a universal technique that can be applied to obtain so-called *pinning* in classes of graph-homomorphisms closed under the composition. This technique is helpful for our study as we also obtain pinning for the restricted class of homomorphisms introduced in the following.

Bipartization

We restrict the study of $\#_p \text{HOMSTO}H$ to the study of bipartite graphs by a restricted class of homomorphisms. For two bipartite graphs $G = (L_G, R_G, E_G)$ and $H = (L_H, R_H, E_H)$ with fixed bipartition we say that a homomorphism from G to H *preserves the order of the bipartition* if the homomorphism maps L_G to L_H and R_G to R_H . The problem $\#_p \text{BIPHOMSTO}H$ with fixed bipartite graph H then asks for the number of these homomorphisms to H . It allows us to restrict the study of $\#_p \text{HOMSTO}H$ to the study of bipartite graphs by the following theorem.


- **Theorem 1.2.** *For any prime p and any graph H , there exists a bipartite graph H' such that*
- if $\#_p \text{BIPHOMSTO}H'$ is $\#_p \text{P-hard}$ then $\#_p \text{HOMSTO}H$ is $\#_p \text{P-hard}$;
 - if $\#_p \text{BIPHOMSTO}H'$ is solvable in polynomial time then $\#_p \text{HOMSTO}H$ is solvable in polynomial time.

This implies that a dichotomy for $\#_p \text{BIPHOMSTO}H'$ yields a dichotomy for $\#_p \text{HOMSTO}H$. As we will later show, the graph H' is a collection of complete bipartite graphs if and only if H satisfies the Dyer and Greenhill criterion. An additional feature of Theorem 1.2 is that it allows for the graph H to contain loops whereas the bipartite graph H' is always loop-less by definition. So far no study of $\#_p \text{HOMSTO}H$ allowed for loops. The structural implications of a bipartite graph H are also heavily exploited in the following analysis.


Hardness in Bipartite $(K_{3,3} \setminus \{e\}, \text{domino})$ -Free Graphs

In the longest and most technically involved part of the paper we study bipartite graphs H not satisfying the Dyer and Greenhill criterion with the goal of finding enough structural information to establish hardness of $\#_p \text{BIPHOMSTO}H$. We find that it suffices to study bipartite graphs in the class denoted $\mathcal{G}_{\text{bip}}^{*p}$ consisting of bipartite graphs without automorphisms of order p , that preserve the order of the bipartition. To this end, we conduct a rigorous structural analysis of the class of bipartite graphs that contain no induced subgraph isomorphic to $K_{3,3} \setminus \{e\}$ or *domino* (see Figure 1 for an illustration). Our insights on the structure of bipartite graphs allow us to establish the following theorem.

- **Theorem 1.3.** *Let p be a prime and $H \in \mathcal{G}_{\text{bip}}^{*p}$ be a $(K_{3,3} \setminus \{e\}, \text{domino})$ -free graph. If there exists a connected component of H that is not a complete bipartite graph, then $\#_p \text{BIPHOMSTO}H$ is $\#_p \text{P-hard}$.*



■ **Figure 1** From left to right: $K_{3,3} \setminus \{e\}$; domino; Example of a bipartite $(K_{3,3} \setminus \{e\}, \text{domino})$ -free and asymmetric graph containing locally and globally K_4 as a minor.



■ **Figure 2** Depiction of the seven minimal fertile graphs as given in Brightwell and Winkler [5, Fig. 6].

In many cases, a *domino* as induced subgraph yields a pair of vertices x, y where x dominates y . The class of bipartite domination-free $K_{3,3} \setminus \{e\}$ -free graphs is one of the focal points of the seminal work by Feder and Vardi [22]. They showed that the class of *graph retract* problems - a notion equivalent to a partially labelled graph homomorphism - to the class of bipartite domination-free $K_{3,3} \setminus \{e\}$ -free graphs contains as much computational power as the whole class of *constraint satisfaction problems* (CSP's), i.e. every CSP is polynomially equivalent to a partially labelled graph homomorphism problem, where the image is a bipartite domination-free $K_{3,3} \setminus \{e\}$ -free graph.

Consider the graphs studied in the work of Brightwell and Winkler [5] shown in Figure 2. The set of graph homomorphisms to these graphs played a key role in their study of spin systems in statistical physics. Prior results incorporate only two out of the seven minimal fertile graphs: ‘the stick’ and ‘the key’. Following the line of argumentation, our results incorporate the previous and three additional minimal fertile graphs. The only missing ones are ‘the hinge’ and ‘the gun’ as the construction used for bipartization yields graphs that are not *domino*-free. In fact, the class of bipartite $(K_{3,3} \setminus \{e\}, \text{domino})$ -free graphs captures all the classes of bipartite graphs studied in previous works on $\#_2\text{HOMSToH}$ and $\#_p\text{HOMSToH}$ except for the recent work by Focke et al. [23] on K_4 -minor-free graphs. Clearly, every biclique with at least 3 vertices in each part contains a K_4 as minor, as is the case with $K_{3,3} \setminus \{e\}$. A *domino* is K_4 -minor-free. Hence, our result given by a local property is orthogonal to the result of Focke et al. [23] given by a global property. An example is depicted in Figure 1.

1.2 Technical Overview

We are going to explain our results and argumentative routes in more detail. Due to their length, details are omitted but can be found in the full version.

In this work, hardness for modular counting problems is indicated by reducing from problems that are $\#_p\text{P}$ -hard. The class $\#_p\text{P}$ contains functions of the form “ $f \bmod p$ ”, where $f \in \#\text{P}$. Notably, for the case $p = 2$ the whole polynomial hierarchy reduces to problems in $\#_2\text{P}$ by Toda [46].

We briefly discuss the insights by Faben and Jerrum [21]. For a pair of graphs G, H we denote by $\text{Hom}(G \rightarrow H)$ the set of homomorphisms from G to H . By $\text{hom}(G \rightarrow H)$ we denote the cardinality $|\text{Hom}(G \rightarrow H)|$ and, for a modulus p , by $\text{hom}_p(G \rightarrow H)$ we denote $\text{hom}(G \rightarrow H) \pmod{p}$. The computational problem $\#\text{HOMSTO}H$ with parameter H then asks to compute $\text{hom}(G \rightarrow H)$ for an input G . Similarly, $\#_p\text{HOMSTO}H$ asks to compute $\text{hom}_p(G \rightarrow H)$. A central point in the study of $\#_p\text{HOMSTO}H$ is the (non)-existence of automorphisms of order p , where for $p = 2$ these automorphisms are called *involutions*. Given an automorphism ρ of order p acting as a derangement on the subset $V' \subseteq V(H)$, Faben and Jerrum [21] showed that the number of homomorphisms σ from any input graph G to H is equivalent to 0 in \mathbb{Z}_p if the image of σ intersects V' . They deduced that, for the subgraph H^ρ of H induced by the fixpoints $V(H) \setminus V'$, there exists a parsimonious reduction from $\#_p\text{HOMSTO}H^\rho$ to $\#_p\text{HOMSTO}H$. Iteratively applying this reduction, one ends up with a subgraph H^{*p} of H that admits no automorphism of order p called the *order p reduced* form of H . This subgraph is also unique up to isomorphism and thus well defined by [21, Theorem 3.7]. The study of $\#_p\text{HOMSTO}H$ focusses on graphs H that do not admit automorphisms of order p , which are called *order p reduced*.

1.2.1 Quantum Homomorphisms

It has been observed by Borgs, Chayes, Kahn, and Lovász [3] that the study of linear combinations of homomorphisms provides great insights especially on the comparability of pairs of graphs, for instance if one is a subgraph of the other. Lovász [38] introduced the term *quantum graph*, denoted \bar{H} , for a linear combination of finitely many graphs and calls the set of pairwise non-isomorphic graphs \mathcal{H} with coefficient $\alpha_H \neq 0$ in \bar{H} its *constituents*:

$$\bar{H} = \sum_{H \in \mathcal{H}} \alpha_H H.$$

A computational problem on \bar{H} translates into the linear combination of computational problems on entities H in \mathcal{H} with coefficient α_H . By Lovász [38] every graph parameter has - if any - a unique expression by a linear combination of finitely many graph homomorphisms up to isomorphisms.

We establish a polynomial time reduction from $\#_p\text{HOMSTO}H$ to $\#_p\text{HOMSTO}\bar{H}$, for any quantum graph \bar{H} and any constituent H in \bar{H} . Such a polynomial time reduction is commonly referred to as a *pinning*-reduction because it enables us to consider the subproblem where a partially mapping is already fixed. One of the main problems of reduction algorithms on modular counting problems is the loss of control of summations in a finite field because we cannot infer from a number of non-zero summands that the sum is non-zero. For instance, let $p = 2$ and \bar{H} be the quantum graph consisting of the two graphs H_1 and H_2 with coefficients $\alpha_{H_1} = \alpha_{H_2} = 1$, where H_1 is an asymmetric tree and H_2 is the disjoint union of a copy of H_1 and an isolated vertex. Let G be a connected graph and input for $\#\text{HOMSTO}\bar{H}$, then we obtain $\text{hom}(G \rightarrow H_2) = \text{hom}(G \rightarrow H_1) + \text{hom}(G \rightarrow K_1)$. Consequently, when

computing $\text{hom}(G \rightarrow H_1) + \text{hom}(G \rightarrow H_2)$ in \mathbb{Z}_2 the term referring to H_1 vanishes and this amounts to computing $\text{hom}(G \rightarrow K_1)$, which is polynomial time solvable. However, Theorem 1.1 yields that $\#_2\text{HOMSTO}\bar{H}$ is $\#_2\text{P-hard}$. The reason is that the split into $\text{hom}(G \rightarrow H_1) + \text{hom}(G \rightarrow K_1)$ only works if G is connected and by utilizing disconnected graphs the additional vertex in H_2 yields enough information to distinguish between H_1 and H_2 . Therefore, we can extract $\text{hom}_p(G \rightarrow H_1)$ from $\text{hom}_p(G \rightarrow \bar{H})$.

In finite fields, reduction algorithms usually rely heavily on multiplication. We find that the beautiful insight on specific matrices defined on families \mathcal{F} of simple graphs provided by Borgs, Chayes, Kahn, and Lovász [3, Lemma 4.2] is able to lift us above this hurdle. In order to adapt this result we first extend it to allow for graphs that contain loops. Then, we translate the result to counting in a finite field of prime order. A straightforward application of the modulo operator is not sufficient as the graphs in \mathcal{F} might contain a number of automorphisms that is a multiple of p . We restrict to order p reduced graphs and argue why this allows for an application of the modulo operator. In this way, we show the following.

► **Corollary 1.4.** *Let $k \geq 1$ and let $\mathcal{F} = \{F_1, \dots, F_k\}$ be a finite family of non-isomorphic order p reduced graphs closed under surjective homomorphic image, that contain no multi-edge. Then the matrix*

$$M_{\text{hom}} = [\text{hom}_p(F_i \rightarrow F_j)]_{i,j=1}^k$$

is nonsingular.

The strength of this result for our purposes is twofold. First, it allows us to show Theorem 1.1 in a concise manner. Given a quantum graph \bar{F} with set of p -constituents $\mathcal{F} = (F_1, \dots, F_r)$ closed under surjective homomorphic image, we obtain by Corollary 1.4 that any system of linear equations of the form $\bar{x} M_{\text{hom}} = \bar{v}$ has a unique solution in the field \mathbb{Z}_p . Therefore, for any vector \bar{v} with entries $(v_i)_{i \in [k]}$ there exists a unique linear combination of entities in \mathcal{F} with coefficients $\alpha_{\mathcal{F}}$ that yield the vector \bar{v} . In fact, we observe that this corresponds to a quantum graph \bar{F}' with $\text{hom}_p(\bar{F}' \rightarrow F_i) = v_i$ that *implements* the vector \bar{v} . In particular, there exists a quantum graph \bar{F}' implementing the i -th standard vector allowing us to ‘pick’ the i -th entry of \mathcal{F} , i.e. $\text{hom}_p(\bar{F}' \rightarrow F_j) = 1$ if $j = i$ and $\text{hom}_p(\bar{F}' \rightarrow F_j) = 0$ otherwise. Given an input graph G for $\text{hom}_p(G \rightarrow \bar{F})$, we then construct a quantum graph \bar{F}^* from G and \bar{F}' such that $\text{hom}_p(\bar{F}^* \rightarrow \bar{F}) = \text{hom}_p(G \rightarrow F_i)$. The main problem for this application is then that the set \mathcal{F} of p -constituents might not be closed under surjective homomorphic image. Given any quantum graph \bar{H} with set of p -constituents \mathcal{H} , we need to define a suitable family \mathcal{F} that contains all the image graphs needed. We find that the subgraphs of the maximal closure are sufficient for this purpose and obtain Theorem 1.1.

The second strength is the adaptability to subproblems of homomorphisms. The main property needed is that the subset of homomorphisms has to be closed under composition, i.e. the subset is actually a *subgroup* of the group of homomorphisms. Examples are *vertex surjective homomorphisms* and *compactness* as studied by Focke et al. [24]. A homomorphism $\sigma \in \text{Hom}(G \rightarrow H)$ is *vertex surjective* if the image-set of σ is the whole set $V(H)$. The homomorphism σ is a *compactness* if it is vertex surjective and every non-loop edge e is in the image of σ . A closely related example is the problem of counting *partially labelled* homomorphisms $\#_{\text{PARTLABHOMSTO}H}$, that are homomorphisms from an input graph G to H that have to respect a given mapping from a subset $V_G \subset V(G)$ to a subset $V_H \subset V(H)$ and are also referred to as *retractions* (see e.g. Focke et al. [24]). The reduction from $\#_{\text{PARTLABHOMSTO}H}$ to $\#_{\text{HOMSTO}H}$ is a building stone of every paper in the study of $\#_{\text{HOMSTO}H}$ and can be obtained in a swift manner due to the strength of Corollary 1.4. A third example will be discussed in the next subsection.

1.2.2 Bipartization

Chen et al. [12] employed the tensor product $H \otimes K_2 = H'$ to reduce to $\#HOMSTOH$ from $\#HOMSTOH'$, where H' is bipartite. The main problem when adapting this construction to modular counting $\#_pHOMSTOH$ is that for every graph G the number of homomorphisms $hom_2(G \rightarrow K_2)$ is 0 and thus the tensor product with K_2 annihilates seemingly any structure that might imply hardness. Instead of branching the study of $\#_pHOMSTOH$ into one studying the modulus 2 and one studying the modulus of odd primes, we solve this issue in a uniform way for all prime moduli.

The key insight towards this is that for an involution-free graph H the tensor product $H' = H \otimes K_2$ only yields involutions on H' that exchange the parts of the bipartition. A very important example is the graph H consisting of a single edge with one loop, for which it is known that $\#HOMSTOH$ is equivalent to counting the number of independent sets. The graph $H' = H \otimes K_2$ is then the path with 4 vertices (see Figure 2), that admits only the reflection across the middle edge. It is known that $\#HOMSTOH'$ is equivalent to counting the number of bipartite independent sets $\#BIS$ and also that $\#_4HOMSTOH'$ is $\#_2$ P-hard (see [28]) whereas $\#_2HOMSTOH'$ is polynomial time solvable. In order to evade the artificial involutions yielded by the tensor product with K_2 we introduce the study on the problem of counting homomorphisms between bipartite graphs that preserve the order of the bipartition denoted $\#BIPHOMSTOH$. For example, if H is the path with 4 vertices then $\#_2BIPHOMSTOH$ is equivalent to $\#_4HOMSTOH$.


► **Lemma 1.5.** *Let p be a prime, let H be a graph, and let $H' = H \otimes K_2$. Then, $\#_pBIPHOMSTOH'$ reduces to $\#_pHOMSTOH$ under parsimonious reduction.*

We note that the graph $H' = H \otimes K_2$ is a collection of complete bipartite graphs if and only if H satisfies the Dyer and Greenhill criterion, for these graphs $\#_pBIPHOMSTOH'$ is solvable in polynomial time.

The reduction from $\#_pBIPHOMSTOH$ has the downside that the machinery developed over the course of multiple papers on $\#_pHOMSTOH$ is not stated for the subgroup of homomorphisms counted by $\#_pBIPHOMSTOH$. We remedy this. First, by the strong adaptability of Corollary 1.4 and the subsequent reduction algorithm we obtain pinning for the problem $\#_pBIPHOMSTOH$. Second, using automorphisms of order p that preserve the order of the bipartition we reduce the bipartite graph H to a *part-wise order p reduced* bipartite graph $(H)^{*p}$. We deduce that the goal towards a dichotomy for $\#HOMSTOH$ is captured by Theorem 1.2. The chain of reductions is displayed below.

$$\#_pBIPPARTLABHOMSTOH^* \stackrel{(H^* = (H')^{*p})}{=} \#_pBIPHOMSTOH^* \stackrel{(H' = H \otimes K_2)}{\leq_{pars}} \#_pBIPHOMSTOH' \leq_{pars} \#_pHOMSTOH$$

We employ a gadgetry that establishes a reduction to $\#_pBIPHOMSTOH$ from a variant of $\#BIS$ with weights on both types of vertices. Such a gadgetry yielding hardness is called a *p-hardness gadget*. By an adaptation of the dichotomy for $\#BIS$ with weights on the vertices in the independent set given in [28] this reduction establishes hardness when counting in \mathbb{Z}_p if and only if none of the weights is equivalent to 0 in \mathbb{Z}_p . The problem of $\#BIS$ with weights on the vertices in the independent set is established as terminal problem yielding hardness in the study of $\#_pHOMSTOH$ [28, 36] as the bigger modulus implies a richer structure compared to the study of $\#_2HOMSTOH$ that traditionally focusses on counting the number of independent sets.



■ **Figure 3** For $p = 3$ the tree H contains no ball of radius 2 around any vertex with enough structure to yield hardness even though $\#_3\text{HOMSTO}H$ is $\#_3$ P-hard.

1.2.3 Hardness in Bipartite $(K_{3,3} \setminus \{e\}, \text{domino})$ -Free Graphs

A central argument in the work of Chen et al. [12] is that for a bipartite graph H and the problem $\# \text{HOMSTO}H$ there exists a simple reduction from $\# \text{HOMSTO}B$, where B is the ball of radius 2 around a vertex v in H denoted $B_2(v)$, i.e. vertices of distance at most 2 from v . By an iterative application of this argument they establish a reduction from $\# \text{HOMSTO}P$, where P is a generalization of the path with 4 vertices. Even though the reduction argument can be made valid for $\#_p \text{HOMSTO}H$ and $\#_p \text{BIPHOMSTO}H$ the restriction to a substructure might destroy the properties that yield hardness already for trees H , a class of graphs for which the dichotomy is proved (see [21, 28]). An example is depicted in Figure 3.

In a nutshell, the induced subgraphs of radius at most 2 can admit too many automorphisms of order p . In Figure 3 we observe that the problem originates from too many instances of complete bipartite graphs $K_{1,b}$, where $b \equiv 0 \pmod{3}$ or $b \equiv 1 \pmod{3}$. The way to overcome this is to also consider the global structure. In the case displayed in Figure 3 the number of walks of length 4 from v_4 to a vertex v in the neighbourhood of v_1 is $0 \pmod{3}$ if $v = v_2$ and 1 else. The goal is then to construct a reduction restricting the study to $B_2(v_0) \setminus \{v_2\}$, a graph that yields hardness. We do this in a general form by a second type of gadgetry called (B, p) -gadget that reduces $\#_p \text{BIPHOMSTO}H$ from $\#_p \text{BIPHOMSTO}B$.

As we have argued, one of the main obstacles for the study on $\#_p \text{HOMSTO}H$ are complete bipartite graphs. The graph $K_{3,3} \setminus \{e\}$ denotes the graph obtained from $K_{3,3}$ by deleting an edge, and the graph *domino* denotes the graph obtained from $K_{3,3}$ by deleting two edges without introducing a cut-vertex (see Figure 1). By the restriction to $(K_{3,3} \setminus \{e\}, \text{domino})$ -free bipartite graphs we study exactly the case of a great many of complete bipartite induced subgraphs. To this end, we observe that for every vertex $v \in H$ the induced subgraph $B_2(v)$ splits into connected components obtained from deleting v . The *split* of $B_2(v)$ at v corresponds to the set of these connected components, where every component contains a copy of v . By the absence of induced subgraphs isomorphic to $K_{3,3} \setminus \{e\}$ or *domino* we deduce that the blocks containing v in these components have to be complete bipartite.

The overall line of argumentations towards Theorem 1.3 is then the following. First, we establish the dichotomy for all $(K_{3,3} \setminus \{e\}, \text{domino})$ -free bipartite graphs H of radius at most 2 by a combination of (B, p) - and p -hardness gadgets. This is done by a careful structural study of the split of H at a central vertex v . An important first result is then that any bipartite graph H in $\mathcal{G}_{\text{bip}}^{*p}$ that contains a vertex v where $B_2(v)$ falls into the hard cases of the dichotomy, is itself such that $\#_p \text{BIPHOMSTO}H$ is $\#_p$ P-hard. Second, we study graphs H of radius larger than 2 in order to establish enough structural information of H allowing us to construct either a p -hardness gadget for H or a (B, p) -gadget such that $\#_p \text{BIPHOMSTO}B$ is $\#_p$ P-hard. This second step is very long and technically involved because the class of $(K_{3,3} \setminus \{e\}, \text{domino})$ -free graphs allows for many cases commanding us to explore the global structure of H . Before we shed more light on how we proceed towards the second step, we

display the chain of reduction arguments below, where the intermediate steps for H_i refer to H itself or an induced subgraph obtained by a (H_i, p) -gadget.

$$\#_p \text{BIS}_{\lambda_\ell, \lambda_r}^{\kappa_\ell, \kappa_r} \leq_P \#_p \text{BIPHOMSTO}H_k \leq_P \cdots \leq_P \#_p \text{BIPHOMSTO}H_1 \leq_P \#_p \text{BIPHOMSTO}H$$

Informally, we split the study of bipartite $(K_{3,3} \setminus \{e\}, \text{domino})$ -free graphs of radius larger than 2 into two broad cases: graphs with no pair of vertices that have a multiple of p common neighbours, and graphs with such a pair of vertices. Focusing on the first case, if the graph H contains a cycle of length at least 6 we argue that the cycle provides enough structure to show hardness. Otherwise, the more restricted structure of H renders the graph ‘tree-like’. The ‘leaves’ of H are vertices v , such that the split of $B_2(v)$ at v contains at most one component adjacent to a vertex not in $B_2(v)$. We call such a vertex a *dead end* and traverse the graph H along a path P starting at a dead end v . The path P is constructed such that it allows us to establish hardness depending only on the local properties of its endvertices, we call such a path P a *generalized hardness path*; an example is depicted in Figure 4. The first broad case is then established by showing that H contains a generalized hardness path whose endvertices are such that P yields hardness.

Turning towards the second broad case, we traverse the graph H again along a path P . Contrary to the first case, we can encounter a pair of vertices with a multiple of p common neighbours. For our purposes, it is important to evade such pairs of vertices. We argue that by the property of H being in $\mathcal{G}_{\text{bip}}^{*p}$ we are always able to do so. This leads to a structure we call *p-mosaic path* that, similar to a generalized hardness path, provides enough structural information towards establishing hardness depending only on the local properties of its endvertices. An example is shown in Figure 4. We find that if a p -mosaic path is a cycle then this cycle satisfies the same properties we found to be sufficient to yield hardness in the first broad case. Interestingly, we show that if the p -mosaic path does not yield hardness the only remaining option for an endvertex of a p -mosaic path is to also be the endvertex of a generalized hardness path. We conclude that H has to contain a concatenation of generalized hardness paths and p -mosaic paths. Arguing by the finiteness of H we show that this concatenation has to yield hardness.


1.2.4 Beyond $(K_{3,3} \setminus \{e\}, \text{domino})$ -Free Graphs

In light of our findings we conjecture that for a bipartite graph H in $\mathcal{G}_{\text{bip}}^{*p}$ the problem $\#_p \text{BIPHOMSTO}H$ is $\#_p \mathbf{P}$ -hard if H is not a collection of complete bipartite graphs. This conjecture then extends towards a conjecture on $\#_p \text{HOMSTO}H$ and also incorporates the conjecture of Faben and Jerrum.

► **Conjecture 1.6.** *Let p be a prime and H a graph with order p reduced form H^{*p} . Then, $\#_p \text{HOMSTO}H$ is solvable in polynomial time if the connected components of H^{*p} are complete bipartite or reflexive complete. Otherwise, $\#_p \text{HOMSTO}H$ is $\#_p \mathbf{P}$ -complete.*

We emphasize this conjecture by a study of the set of *partially surjective homomorphisms* from a graph G to a graph H denoted $\text{PartSurj}(G \rightarrow H)$. Partially surjective homomorphisms have to be surjective on a set of distinguished vertices $V^{\text{dist}} \subseteq V(H)$ and a set of distinguished edges $E^{\text{dist}} \subseteq E(H)$. We deduce that it suffices to study graphs H without automorphisms of order p acting bijectively on V^{dist} and E^{dist} . However, this reduction does not capture all cancellations because the graph H might still admit too many general automorphisms.

By our results on quantum graphs and an application of the inclusion-exclusion principle we find that the dichotomy presented in Conjecture 1.6 extends to a dichotomy on the



■ **Figure 4** For $p = 3$ the left figure depicts an example of a generalized hardness path and the right figure depicts an example of a p -mosaic path.

whole class of partially surjective homomorphisms. Contrary to the non-modular version established in [12], this dichotomy does not state a clear structural characterisation of the hard instances due to the mentioned possibility of additional cancellations. For the special cases in which the parameter graph H is order p reduced we amplify the dichotomy such that it states clear structural characterisations. Two examples for this case are the problems $\#_p \text{VERTSURJHOMSTO}H$ and $\#_p \text{COMPTO}H$ of counting in \mathbb{Z}_p the number of vertex surjective homomorphisms and compactions, respectively. We obtain the following criteria analogous to the criteria in the non-modular setting given by Focke et al. [24].

► **Corollary 1.7.** *Let p be a prime and H be a graph. The problem $\#_p \text{VERTSURJHOMSTO}H$ is solvable in polynomial time if either H admits an automorphism of order p , or every connected component of H is a complete bipartite graph or a reflexive complete graph.*

The problem $\#_p \text{COMPTO}H$ is solvable in polynomial time if either H admits an automorphism of order p , or every connected component of H is an irreflexive star or a reflexive complete graph of size at most two.


Assuming Conjecture 1.6 both problems are $\#_p \text{P}$ -hard in every other case.

In order to prove Conjecture 1.6, we need to study bipartite graphs that contain $K_{3,3} \setminus \{e\}$ or *domino* as an induced subgraph. The strong restrictions on the structure of the graphs under study are a double-edged sword. On one hand, it is more plausible to find enough structure that yields hardness. On the other hand, it is more difficult to pin the structural analysis down to a handful of cases. Furthermore, the higher moduli imply even more complexity of the structural analysis. We illustrate an especially difficult example in Figure 5.

We call such a graph as illustrated in Figure 5 a *p-Catherine wheel*. Even though these graphs are 2-connected and of radius 2 their highly symmetric global structure together with the lack of small structure in the sets R_i makes it difficult to identify sources for hardness. Here, the case displayed in Figure 5 where the sets R_i only yield a collection of trees and the sum of degrees of the vertices in the sets is $0 \pmod p$ is especially difficult. We note that such a graph cannot be part-wise order p reduced for $p = 2, 3$, which highlights the gain of complexity due to higher moduli.

1.3 Related Literature

Before we conclude the introduction, we mention related bodies of work. The study of homomorphisms under the point of view of *parameterized algorithms* has been long established (see Diaz, Serna, and Thilikos [15]) but enriched by the work of Amini, Fomin, and Saurabh [1]



■ **Figure 5** Illustration of a 5-Catherine wheel. Edges to encircled sets illustrate edges to every vertex in the set. The smaller substructure in the sets R_1, \dots, R_5 is illustrated to the right, where the vertices in the sets R_i are the more prominent ones at the bottom of the row.

and by Curticapean, Dell, and Marx [13], who also introduced linear combinations of graph homomorphisms to the study and motivated subsequent works, for instance Roth and Wellnitz [41].

The study of homomorphisms from the point of view of *extremal combinatorics* incorporates important conjectures like *Sidorenko's conjecture* [43, 44], which states a universal lower bound on the number of homomorphisms from a bipartite graph and, in a weaker version, can be found in the work of Simonovits [45]. Until today the conjecture remains open but still enjoys new contributions like the recent article by Shams, Ruoizzi, and Csikvári [42].

This leads to the body of work studying *approximation algorithms* including the work of Goldberg and Jerrum [30] on tree homomorphisms and the work of Galanis, Goldberg, and Jerrum [25], who showed that approximating the number of homomorphisms to a fixed graph H is $\#BIS$ -hard, a notorious complexity class in this body of work. These findings yield an interesting connection to ours in the form of the reduction from (versions of) $\#BIS$.

The body of studies concerning different versions of homomorphism problems is vast. It contains dichotomies for the affiliated problem, where the pre-image is from a fixed class of graphs, given by Dalmau and Jonsson [14] and Grohe [32]. Turning towards versions of the problem with fixed image, Focke, Goldberg, and Zivný [24] gave a dichotomy for surjective homomorphisms and compactions, and Dyer, Goldberg, and Paterson [16] gave a dichotomy for directed homomorphisms if the target is acyclic. The line of research towards the dichotomy for the generalization of $\#HOMSTO H$ allowing weights by Cai, Chen, and Lu [9] incorporates works by Bulatov and Grohe [7] and Goldberg, Grohe, Jerrum, and Thurley [29]. Recently, Govorov, Cai, and Dyer [31] extended this research body.

The connection of homomorphisms and CSP's was already shown by Feder and Vardi [22]. Bulatov [6] showed that the problem of counting satisfying assignments to a CSP enjoys a dichotomy theorem, a result on which Dyer and Richerby [18] shed more light. Furthermore, a complete dichotomy for directed homomorphism can be found in the dichotomy on counting weighted versions of CSP's by Cai and Chen [8]. Guo, Huang, Lu, and Xia [34] gave a dichotomy for the associated modular problem.

Finally, the recent work by Cai and Govorov [10] studied the power of expression of the class of homomorphisms. By studying algebras of quantum graphs they provide a general technique and showed, for instance, that the problem of counting perfect matchings cannot be expressed by counting homomorphisms to a fixed graph H regardless of possible weights.

References

- 1 Omid Amini, Fedor V. Fomin, and Saket Saurabh. Counting subgraphs via homomorphisms. *SIAM J. Discret. Math.*, 26(2):695–717, 2012. doi:10.1137/100789403.
- 2 Christian Borgs, Jennifer Chayes, László Lovász, Vera T. Sós, and Katalin Vesztegombi. Counting graph homomorphisms. In *Topics in Discrete Mathematics*, pages 315–371, 2006. doi:10.1007/3-540-33700-8_18.
- 3 Christian Borgs, Jennifer T. Chayes, Jeff Kahn, and László Lovász. Left and right convergence of graphs with bounded degree. *Random Structures and Algorithms*, 42(1):1–28, 2013. doi:10.1002/rsa.20414.
- 4 Christian Borgs, Jennifer T. Chayes, László Lovász, Vera T. Sós, Balázs Szegedy, and Katalin Vesztegombi. Graph limits and parameter testing. In *Proceedings of STOC 2006*, pages 261–270, 2006. doi:10.1145/1132516.1132556.
- 5 Graham R. Brightwell and Peter Winkler. Graph homomorphisms and phase transitions. *J. Comb. Theory, Ser. B*, 77(2):221–262, 1999. doi:10.1006/jctb.1999.1899.
- 6 Andrei A. Bulatov. The complexity of the counting constraint satisfaction problem. *J. ACM*, 60(5):34:1–34:41, 2013. doi:10.1145/2528400.
- 7 Andrei A. Bulatov and Martin Grohe. The complexity of partition functions. *Theor. Comput. Sci.*, 348(2-3):148–186, 2005. doi:10.1016/j.tcs.2005.09.011.
- 8 Jin-Yi Cai and Xi Chen. Complexity of counting CSP with complex weights. *J. ACM*, 64(3):19:1–19:39, 2017. doi:10.1145/2822891.
- 9 Jin-Yi Cai, Xi Chen, and Pinyan Lu. Graph homomorphisms with complex values: A dichotomy theorem. *SIAM J. Comput.*, 42(3):924–1029, 2013. doi:10.1137/110840194.
- 10 Jin-Yi Cai and Artem Govorov. Perfect matchings, rank of connection tensors and graph homomorphisms. In *Proceedings of SODA 2019*, pages 476–495, 2019. doi:10.1137/1.9781611975482.30.
- 11 Ashok K. Chandra and Philip M. Merlin. Optimal implementation of conjunctive queries in relational data bases. In *Proceedings of STOC 1977*, pages 77–90, 1977. doi:10.1145/800105.803397.
- 12 Hubie Chen, Radu Curticapean, and Holger Dell. The exponential-time complexity of counting (quantum) graph homomorphisms. In *Proceedings of WG 2019*, pages 364–378, 2019. doi:10.1007/978-3-030-30786-8_28.
- 13 Radu Curticapean, Holger Dell, and Dániel Marx. Homomorphisms are a good basis for counting small subgraphs. In *Proceedings of STOC 2017*, page 210–223, 2017. doi:10.1145/3055399.3055502.
- 14 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.
- 15 Josep Díaz, Maria Serna, and Dimitrios M. Thilikos. Counting H-colorings of partial k-trees. *Theor. Comput. Sci.*, 281(1):291 – 309, 2002. Selected Papers in honour of Maurice Nivat. doi:https://doi.org/10.1016/S0304-3975(02)00017-8.
- 16 Martin E. Dyer, Leslie Ann Goldberg, and Mike Paterson. On counting homomorphisms to directed acyclic graphs. *J. ACM*, 54(6):27, 2007. doi:10.1145/1314690.1314691.
- 17 Martin E. Dyer and Catherine Greenhill. The complexity of counting graph homomorphisms. *Random Structures and Algorithms*, 17(3-4):260–289, 2000. doi:10.1002/1098-2418(200010/12)17:3/4<260::AID-RSA5>3.0.CO;2-W.
- 18 Martin E. Dyer and David Richerby. An effective dichotomy for the counting constraint satisfaction problem. *SIAM J. Comput.*, 42(3):1245–1274, 2013. doi:10.1137/100811258.
- 19 Ethan R. Elenberg, Karthikeyan Shanmugam, Michael Borokhovich, and Alexandros G. Dimakis. Beyond triangles: A distributed framework for estimating 3-profiles of large graphs. In *Proceedings of KDD 2015*, pages 229–238, 2015. doi:10.1145/2783258.2783413.
- 20 Ethan R. Elenberg, Karthikeyan Shanmugam, Michael Borokhovich, and Alexandros G. Dimakis. Distributed estimation of graph 4-profiles. In *Proceedings of WWW 2016*, pages 483–493, 2016. doi:10.1145/2872427.2883082.

- 21 John Faben and Mark Jerrum. The complexity of parity graph homomorphism: An initial investigation. *Theory of Computing*, 11:35–57, 2015. doi:10.4086/toc.2015.v011a002.
- 22 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.
- 23 Jacob Focke, Leslie Ann Goldberg, Marc Roth, and Stanislav Zivný. Counting homomorphisms to K_4 -minor-free graphs, modulo 2. In *Proceedings of SODA 2021*, pages 2303–2314, 2021. doi:10.1137/1.9781611976465.137.
- 24 Jacob Focke, Leslie Ann Goldberg, and Stanislav Zivný. The complexity of counting surjective homomorphisms and compactions. *SIAM J. Discret. Math.*, 33(2):1006–1043, 2019. doi:10.1137/17M1153182.
- 25 Andreas Galanis, Leslie Ann Goldberg, and Mark Jerrum. Approximately counting H-colorings is #BIS-hard. *SIAM J. Comput.*, 45(3):680–711, 2016. doi:10.1137/15M1020551.
- 26 Andreas Göbel, Leslie Ann Goldberg, and David Richerby. The complexity of counting homomorphisms to cactus graphs modulo 2. *ACM Trans. Comput. Theory*, 6(4):17:1–17:29, 2014. doi:10.1145/2635825.
- 27 Andreas Göbel, Leslie Ann Goldberg, and David Richerby. Counting homomorphisms to square-free graphs, modulo 2. *ACM Trans. Comput. Theory*, 8(3):12:1–12:29, 2016. doi:10.1145/2898441.
- 28 Andreas Göbel, J. A. Gregor Lagodzinski, and Karen Seidel. Counting homomorphisms to trees modulo a prime. In *Proceedings of MFCS 2018*, pages 49:1–49:13, 2018. doi:10.4230/LIPIcs.MFCS.2018.49.
- 29 Leslie Ann Goldberg, Martin Grohe, Mark Jerrum, and Marc Thurley. A complexity dichotomy for partition functions with mixed signs. *SIAM J. Comput.*, 39(7):3336–3402, 2010. doi:10.1137/090757496.
- 30 Leslie Ann Goldberg and Mark Jerrum. The complexity of approximately counting tree homomorphisms. *ACM Trans. Comput. Theory*, 6(2):8:1–8:31, 2014. doi:10.1145/2600917.
- 31 Artem Govorov, Jin-Yi Cai, and Martin E. Dyer. A dichotomy for bounded degree graph homomorphisms with nonnegative weights. In *Proceedings of ICALP 2020*, pages 66:1–66:18, 2020. doi:10.4230/LIPIcs.ICALP.2020.66.
- 32 Martin Grohe. The complexity of homomorphism and constraint satisfaction problems seen from the other side. *J. ACM*, 54(1):1:1–1:24, 2007. doi:10.1145/1206035.1206036.
- 33 Martin Grohe, Thomas Schwentick, and Luc Segoufin. When is the evaluation of conjunctive queries tractable? In *Proceedings of STOC 2001*, pages 657–666, 2001. doi:10.1145/380752.380867.
- 34 Heng Guo, Sangxia Huang, Pinyan Lu, and Mingji Xia. The complexity of weighted boolean #csp modulo k . In *Proceedings of STACS 2011*, pages 249–260, 2011. doi:10.4230/LIPIcs.STACS.2011.249.
- 35 Pavol Hell and Jaroslav Nešetřil. On the complexity of H -coloring. *J. Comb. Theory, Ser. B*, 48(1):92–110, 1990. doi:10.1016/0095-8956(90)90132-J.
- 36 Amirhossein Kazemini and Andrei A. Bulatov. Counting homomorphisms modulo a prime number. In *Proceedings of MFCS 2019*, pages 59:1–59:13, 2019. doi:10.4230/LIPIcs.MFCS.2019.59.
- 37 Richard E. Ladner. On the structure of polynomial time reducibility. *J. ACM*, 22(1):155–171, 1975. doi:10.1145/321864.321877.
- 38 László Lovász. *Large Networks and Graph Limits*, volume 60 of *Colloquium Publications*. AMS, 2012.
- 39 Benjamin Rossman. Homomorphism preservation theorems. *J. ACM*, 55(3):15:1–15:53, 2008. doi:10.1145/1379759.1379763.
- 40 Benjamin Rossman. An improved homomorphism preservation theorem from lower bounds in circuit complexity. In *Proceedings of ITCS 2017*, pages 27:1–27:17, 2017. doi:10.4230/LIPIcs.ITCS.2017.27.

- 41 Marc Roth and Philip Wellnitz. Counting and finding homomorphisms is universal for parameterized complexity theory. In *Proceedings of SODA 2020*, pages 2161–2180, 2020. doi:10.1137/1.9781611975994.133.
- 42 Shahab Shams, Nicholas Ruozzi, and Peter Csikvári. Counting homomorphisms in bipartite graphs. In *Proceedings of ISIT 2019*, pages 1487–1491, 2019. doi:10.1109/ISIT.2019.8849389.
- 43 Alexander F. Sidorenko. Inequalities for functionals generated by bipartite graphs. (*Russian*) *Diskret. Mat.*, 3(3):50–65, 1991.
- 44 Alexander F. Sidorenko. A correlation inequality for bipartite graphs. *Graphs and Combin.*, 9(2-4):201–204, 1993. doi:10.1007/BF02988307.
- 45 Miklós Simonovits. Extremal graph problems, degenerate extremal problems, and supersaturated graphs. *Progress in Graph Theory*, page 419–437, 1984.
- 46 Seinosuke Toda. PP is as hard as the polynomial-time hierarchy. *SIAM J. Comput.*, 20(5):865–877, 1991. doi:10.1137/0220053.
- 47 Johan Ugander, Lars Backstrom, and Jon M. Kleinberg. Subgraph frequencies: mapping the empirical and extremal geography of large graph collections. In *Proceedings of WWW 2013*, pages 1307–1318, 2013. doi:10.1145/2488388.2488502.
- 48 Leslie G. Valiant. Accidental algorithms. In *Proceedings of FOCS 2006*, pages 509–517, 2006. doi:10.1109/FOCS.2006.7.