Counting Homomorphisms to Trees Modulo a Prime

ANDREAS GÖBEL, J. A. GREGOR LAGODZINSKI, and KAREN SEIDEL, Hasso Plattner
Institute, University of Potsdam, Germany

Many important graph-theoretic notions can be encoded as counting graph homomorphism problems, such as partition functions in statistical physics, in particular independent sets and colourings. In this article, we study the complexity of $\#_p\text{HomsToH}$, the problem of counting graph homomorphisms from an input graph to a graph H modulo a prime number p . Dyer and Greenhill proved a dichotomy stating that the tractability of non-modular counting graph homomorphisms depends on the structure of the target graph. Many intractable cases in non-modular counting become tractable in modular counting due to the common phenomenon of cancellation. In subsequent studies on counting modulo 2, however, the influence of the structure of H on the tractability was shown to persist, which yields similar dichotomies.

Our main result states that for every tree H and every prime p the problem $\#_p\text{HomsToH}$ is either polynomial time computable or $\#_p\text{P}$ -complete. This relates to the conjecture of Faben and Jerrum stating that this dichotomy holds for every graph H when counting modulo 2. In contrast to previous results on modular counting, the tractable cases of $\#_p\text{HomsToH}$ are essentially the same for all values of the modulo when H is a tree. To prove this result, we study the structural properties of a homomorphism. As an important interim result, our study yields a dichotomy for the problem of counting weighted independent sets in a bipartite graph modulo some prime p . These results are the first suggesting that such dichotomies hold not only for the modulo 2 case but also for the modular counting functions of all primes p .

CCS Concepts: • **Mathematics of computing** → **Graph algorithms; Enumeration; • Theory of computation** → *Randomness, geometry and discrete structures*;

Additional Key Words and Phrases: Graph homomorphisms, modular counting, complexity dichotomy

ACM Reference format:

Andreas Göbel, J. A. Gregor Lagodzinski, and Karen Seidel. 2021. Counting Homomorphisms to Trees Modulo a Prime. *ACM Trans. Comput. Theory* 13, 3, Article 19 (August 2021), 33 pages.
<https://doi.org/10.1145/3460958>

1 INTRODUCTION

Graph homomorphisms generate a powerful language expressing important notions; examples include constraint satisfaction problems and partition functions in statistical physics. As such, the computational complexity of graph homomorphism problems has been studied extensively from a wide range of views. Early results include that of Hell and Nešetřil [18], who study the complexity of HomsToH , the problem of deciding if there exists a homomorphism from an input graph G to a fixed graph H . They show the following dichotomy: If H is bipartite or has a loop, then the

Author's address: A. Göbel, J. A. G. Lagodzinski, and K. Seidel, Hasso Plattner Institute, University of Potsdam, Prof.-Dr.-Helmert-Str. 2–3, Potsdam, Germany, D-14482; emails: {andreas.gobel, gregor.lagodzinski, karen.seidel}@hpi.de.

This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2021 Copyright held by the owner/author(s).

1942-3454/2021/08-ART19 \$15.00

<https://doi.org/10.1145/3460958>

Fig. 1. The graph H will be our recurring example and the labelling of the vertices is explained later in the introduction.

problem is in P and in every other case HOMSToH is NP -complete. This is particularly interesting as a result of Ladner [20] shows that if $P \neq \text{NP}$, then there exist problems that are neither in P nor NP -hard.

Dyer and Greenhill [7] show a dichotomy for the problem $\# \text{HOMSToH}$, the problem of counting the homomorphisms from an input graph G to H . Their theorem states that $\# \text{HOMSToH}$ is tractable if H is a complete bipartite graph or a complete graph with loops on all vertices; otherwise, $\# \text{HOMSToH}$ is $\# P$ -complete. This dichotomy was progressively extended to weighted sums of homomorphisms with integer weights by Bulatov and Grohe [2], with real weights by Goldberg et al. [15], and, finally, with complex weights by Cai, Chen, and Lu [3].

We study the complexity of counting homomorphisms modulo a prime p . The set of homomorphisms from the input graph G to the target graph H is denoted by $\text{Hom}(G \rightarrow H)$. For each pair of fixed parameters p and H we study the computational problem $\#_p \text{HOMSToH}$, that is the problem of computing $|\text{Hom}(G \rightarrow H)|$ modulo p . The value of p and the structure of the target graph H influence the complexity of $\#_p \text{HOMSToH}$. Consider the graph H in Figure 1. Our results show that $\#_p \text{HOMSToH}$ is computable in polynomial time if $p = 2, 3$, while it is $\#_p P$ -hard for any other prime p , where $\#_p P$ is the canonical hardness class for modular counting problems, as we discuss in Section 1.1.

Our main goal is to fully characterise the complexity of $\#_p \text{HOMSToH}$ in a dichotomy theorem, when H is a forest. In this manner, we aim to determine for which pair of parameters (H, p) the problem is tractable and show that for every other pair of parameters the problem is hard. As the theorem of Ladner [20] extends to the modular counting problems, it is not obvious that there are no instances of $\#_p \text{HOMSToH}$ with an intermediate complexity.

The first study of graph homomorphisms under the setting of modular counting has been conducted by Faben and Jerrum [10]. Their work is briefly described in the following, and we assume the reader to be familiar with the notion of an automorphism and its order. We provide the formal introduction in Section 2. Given a graph H and an automorphism ϱ of H , H^ϱ denotes the subgraph of H induced by the fixpoints of ϱ . We write $H \Rightarrow_k H'$ if there is an automorphism ϱ of order k of H such that $H^\varrho = H'$, and we write $H \Rightarrow_k^* H'$ if either H is isomorphic to H' (written $H \cong H'$) or, for some positive integer t , there are graphs $H_1, \dots, H_t$ such that $H \cong H_1$, $H_1 \Rightarrow_k \dots \Rightarrow_k H_t$, and $H_t \cong H'$.

Faben and Jerrum showed [10, Lemma 3.3] that if the order of ϱ is a prime p , then we have that $|\text{Hom}(G \rightarrow H)| \equiv |\text{Hom}(G \rightarrow H^\varrho)| \pmod{p}$. Furthermore, they showed [10, Theorem 3.7] that there is (up to isomorphism) exactly one graph H^{*p} without automorphisms of order p , such that $H \Rightarrow_p^* H^{*p}$. This graph H^{*p} is called the *order p reduced form* of H . If H^{*p} falls into the polynomial computable cases of the theorem of Dyer and Greenhill, then $\#_p \text{HOMSToH}$ is computable in polynomial time as well. For $p = 2$, Faben and Jerrum conjectured that these are the only instances computable in polynomial time.

CONJECTURE 1.1 (FABEN AND JERRUM [10]). *Let H be a graph. If its order 2 reduced form H^{*2} has at most one vertex, then $\#_2 \text{HOMSToH}$ is in FP ; otherwise, $\#_2 \text{HOMSToH}$ is $\#_2 P$ -complete.*

Faben and Jerrum [10, Theorem 3.8] underlined their conjecture by proving it for the case in which H is a tree. In subsequent works, this proof was extended to cactus graphs [13] and to square-free graphs [14] by Göbel, Goldberg, and Richerby and to K_4 -minor free graphs by Focke et al. [11].

The present work follows a direction orthogonal to the aforementioned. Instead of proving the conjecture for richer classes of graphs, we show a dichotomy for all primes, starting again by studying trees.

THEOREM 1.2. *Let p be a prime and let H be a graph, such that its order p reduced form H^{*p} is a tree. If H^{*p} is a star, then $\#_p \text{HomsToH}$ is computable in polynomial time; otherwise, $\#_p \text{HomsToH}$ is $\#_p \text{P-complete}$.*

Our results are the first to suggest that the conjecture of Faben and Jerrum might apply to counting graph homomorphisms modulo every prime p instead of counting modulo 2. This suggestion, however, remains hypothetical. Borrowing the words of Dyer, Frieze, and Jerrum [6]: “One might even rashly conjecture” it “(though we shall not do so).”

Kazeminia and Bulatov [19] building upon our techniques have extended the dichotomy to include all square-free target graphs H (trees are by definition square-free). Recently, Lagodzinski et al. [21] have proved the dichotomy for an even broader class of graphs. To justify our title, we give the following corollary, stating a dichotomy for all trees H .

COROLLARY 1.3. *Let p be a prime and let H be a tree. If the order p reduced form H^{*p} of H is a star, then $\#_p \text{HomsToH}$ is computable in polynomial time; otherwise, $\#_p \text{HomsToH}$ is $\#_p \text{P-complete}$.*

Furthermore, for graphs H that consist of more than one components our results have the following implication.

COROLLARY 1.4. *Let H be a graph whose order p reduced form H^{*p} is a forest. If every component of H^{*p} is a star, $\#_p \text{HomsToH}$ is computable in polynomial time; otherwise, $\#_p \text{HomsToH}$ is $\#_p \text{P-complete}$.*

We illustrate Theorem 1.2 using the following discussion on Figure 1. The order 2 and the order 3 reduced form of H both are the graph with one vertex, whereas for any other prime the graph stays as such, that is $H^{*p} = H$.

The polynomial-time computable cases follow directly from the results of Faben and Jerrum. Thus, to prove Theorem 1.2 it suffices to show that $\#_p \text{HomsToH}$ is $\#_p \text{P-complete}$ for every tree H that is not a star and has no automorphism of order p . The reductions [10, 13, 14] show hard instances of $\#_2 \text{HomsToH}$ by starting from $\#_2 \text{IS}$, the problem of computing $|I(G)| \pmod{2}$, where $I(G)$ is the set of independent sets of G . $\#_2 \text{IS}$ was shown to be $\#_2 \text{P-complete}$ by Valiant [26]. Later, Faben [8] extended this result by proving $\#_k \text{IS}$ to be $\#_k \text{P-complete}$ for all integers k . For reasons to be explained in Section 1.3 we do not use this problem as a starting point for our reductions.

We turn our attention to $\#_p \text{BIS}$, the problem of counting the independent sets of a bipartite graph modulo p . In the same work Faben [8] includes a construction to show hardness for $\#_p \text{BIS}$. We employ the weighted version $\#_p \text{BIS}_{\lambda_\ell, \lambda_r}$ as a starting point for our reduction.

PROBLEM 1.5. *Name.* $\#_p \text{BIS}_{\lambda_\ell, \lambda_r}$.

Parameter. p prime and $\lambda_\ell, \lambda_r \in \mathbb{Z}_p$.

Input. Bipartite graph $G = (V_L, V_R, E)$.

Output. $Z_{\lambda_\ell, \lambda_r}(G) = \sum_{I \in I(G)} \lambda_\ell^{|V_L \cap I|} \lambda_r^{|V_R \cap I|} \pmod{p}$.

In fact, we obtain the following dichotomy.

THEOREM 1.6. *Let p be a prime and let $\lambda_\ell, \lambda_r \in \mathbb{Z}_p$. If $\lambda_\ell \equiv 0 \pmod{p}$ or $\lambda_r \equiv 0 \pmod{p}$, then $\#_p \text{BIS}_{\lambda_\ell, \lambda_r}$ is computable in polynomial time. Otherwise, $\#_p \text{BIS}_{\lambda_\ell, \lambda_r}$ is $\#_p \text{P}$ -complete.*

To prove hardness for $\#_p \text{HOMSTOH}$, we employ a reduction in three phases: (i) We reduce the “canonical” $\#_p \text{P}$ -complete problem $\#_p \text{SAT}$ to $\#_p \text{BIS}_{\lambda_\ell, \lambda_r}$; (ii) we reduce $\#_p \text{BIS}_{\lambda_\ell, \lambda_r}$ to $\#_p \text{PARTLABHOMSTOH}$, a restricted version of $\#_p \text{HOMSTOH}$, which we define in Section 1.3; and (iii) we reduce $\#_p \text{PARTLABHOMSTOH}$ to $\#_p \text{HOMSTOH}$.

Section 1.1 provides background knowledge on modular counting. In Section 1.2, we will discuss some related work. A high level proof of our three way reduction is provided in Section 1.3. There we also explain the technical obstacles arising from values of the modulo $p > 2$ and how we overcome them by generalising the techniques used for the case $p = 2$. First, we explain step (i), the reduction from $\#_p \text{SAT}$ to $\#_p \text{BIS}_{\lambda_\ell, \lambda_r}$. Afterwards, we describe step (iii), the reduction from $\#_p \text{PARTLABHOMSTOH}$ to $\#_p \text{HOMSTOH}$ establishing the required notation for the subsequent illustration of step (ii), the reduction from $\#_p \text{BIS}_{\lambda_\ell, \lambda_r}$ to $\#_p \text{PARTLABHOMSTOH}$. In Section 1.4, we discuss the limits of our techniques, which do not yield a dichotomy modulo any integer k . Finally, Section 1.5 outlines the rest of this article.

1.1 Modular Counting

Modular counting was originally studied from the decision problem’s point of view. Here, the objective is to determine whether the number of solutions is non-zero modulo k . The complexity class $\oplus \text{P}$ was first studied by Papadimitriou and Zachos [23] and by Goldschlager and Parberry [16]. $\oplus \text{P}$ consists of all problems of the form “is $f(x)$ odd or even?”, where $f(x)$ is a function in $\# \text{P}$. A result of Toda [25] states that every problem in the polynomial time hierarchy reduces in polynomial time to some problem in $\oplus \text{P}$. This result suggests that $\oplus \text{P}$ -completeness represents strong evidence for intractability.

For an integer k , the complexity class $\#_k \text{P}$ consists of all problems of the form “compute $f(x)$ modulo k ,” where $f(x)$ is a function in $\# \text{P}$. In the special case of $k = 2$, $\#_2 \text{P} = \oplus \text{P}$, as the instances of $\#_2 \text{P}$ require a one bit answer. Throughout this article, though, instead of the more traditional notation $\oplus \text{P}$, we will use $\#_2 \text{P}$ to emphasise our interest in computing functions.

If a counting problem can be solved in polynomial time, then the corresponding decision and modular counting problems can also be solved in polynomial time. The converse, though, does not necessarily hold. The reason is that efficient counting algorithms rely usually on an exponential number of cancellations that occur in the problem, e.g., compute the determinant of a non-negative matrix. The modulo operator introduces a natural setting for such cancellations to occur.

For instance, consider the $\# \text{P}$ -complete problem of counting proper 3-colourings of a graph G in the modulo 3 (or even modulo 6) setting. 3-colourings of a graph assigning all three colours can be grouped in sets of size 6, since there are $3! = 6$ permutations of the colours. Thus, the answer to these instances is always a multiple of 6 and therefore “cancels out.” It remains to compute the number of 3-colourings assigning less than 3 colours. For the case of using exactly 2 colours, we distinguish the following two cases: G is not bipartite and there are no such colourings; G is bipartite and the number of 3-colourings of G that use exactly 2 colours is $3(2^c)$, where c is the number of components of G . Finally, computing the number of proper 3-colourings of G that use exactly one colour is an easy task. Either G has an edge and there are no such colourings, or G has no edges and for every vertex there are three colours to choose from.

Valiant [26] observed a surprising phenomenon in the tractability of modular counting problems. He showed that for a restricted version of 3-SAT computing the number of solutions modulo 7 is in FP , but computing this number modulo 2 is $\#_2 \text{P}$ -complete. This mysterious number 7 was later explained by Cai and Lu [4], who showed that the k -SAT version of Valiant’s problem is tractable modulo $2^k - 1$.

1.2 Related Work

We have already mentioned earlier work on counting graph homomorphisms. In this section, we highlight the work of Faben [8] and the work of Guo et al. [17] on the complexity of the modular counting variant of the constraint satisfaction problem.

PROBLEM 1.7. *Name.* $\#_k\text{CSP}(\mathcal{F})$.

Parameter. $k \in \mathbb{Z}_{>0}$ and a set of functions $\mathcal{F} = \{f_1, \dots, f_m\}$, where for each $j \in [m]$, $f_j : \{0, 1\}^{r_j} \rightarrow \mathbb{Z}_k$ and $r_j \in \mathbb{Z}_{>0}$.

Input. Finite set of constraints over Boolean variables $x_1, \dots, x_n$ of the form

$$f_{j_\ell}(x_{i_{\ell,1}}, x_{i_{\ell,2}}, \dots, x_{i_{\ell,r_{j_\ell}}}).$$

Output. $\sum_{x_1, \dots, x_n \in \{0,1\}} \prod_\ell f_{j_\ell}(x_{i_{\ell,1}}, x_{i_{\ell,2}}, \dots, x_{i_{\ell,r_{j_\ell}}}) \pmod{k}$.

Faben showed a dichotomy theorem [8, Theorem 4.11] if the functions in $\mathcal{F}$ have Boolean domain and Boolean range, i.e., $f : \{0, 1\} \rightarrow \{0, 1\}$. Guo et al. extended this dichotomy [17, Theorem 4.1] to $\#_k\text{CSP}$ if the functions in $\mathcal{F}$ have Boolean domain $\{0, 1\}$ but range in $\mathbb{Z}_k$.

Constraint satisfaction problems generalise graph homomorphism problems if the domain of the constraint functions is arbitrarily large. To illustrate that $\#_k\text{CSP}$ is a generalisation of $\#_k\text{HOMSTO}H$ let G be an input for $\#_k\text{HOMSTO}H$, for which we describe an equivalent $\#_k\text{CSP}$ instance. The domain of the constraint satisfaction problem is $D = V(H)$ and $\mathcal{F}$ contains a single binary relation R_H , with $R_H(u, v) = 1$ if $(u, v) \in E(H)$ and $R_H(u, v) = 0$ otherwise. Thus, $\#_k\text{HOMSTO}H$ is an instance of $\#_k\text{CSP}(\{R_H\})$. The input of $\#_k\text{CSP}(\{R_H\})$ contains a variable x_v for every vertex $v \in V(G)$ and a constraint $R_H(x_u, x_v)$ for every edge $(u, v) \in E(G)$. As can be observed from the construction, every valid homomorphism $\sigma : V(G) \rightarrow V(H)$ corresponds to an assignment of the variables $\{x_v\}_{v \in V(G)}$ satisfying every constraint in the CSP.

These results are incomparable to ours. We consider prime values of the modulo and a single binary relation; however, the domain of our relations is arbitrarily large. Furthermore, the results of Faben [8, Theorem 4.11] show that the constraint language $\mathcal{F}$ for which $\#_2\text{CSP}$ is tractable is richer than the constraint language for which $\#_k\text{CSP}$ is tractable, where $k > 2$. In contrast, our results show that the dichotomy criterion of $\#_p\text{HOMSTO}H$ remains the same for all primes p if H is a tree.

1.3 Beyond One-bit Functions

Weighted bipartite independent sets. To explain how we prove Theorem 1.6, consider a bipartite graph $G = (V_L, V_R, E)$ and let $\lambda_\ell = 0$ (the case $\lambda_r = 0$ is symmetric). We observe that every independent set I that contributes a non-zero summand to $Z_{\lambda_\ell, \lambda_r}(G)$ can only contain vertices in V_R ($Z_{\lambda_\ell, \lambda_r}(G)$ is defined in Problem 1.5). This yields the closed form $Z_{\lambda_\ell, \lambda_r}(G) = (\lambda_r + 1)^{|V_R|}$, which is computable in polynomial time. Regarding the case $\lambda_\ell, \lambda_r \not\equiv 0 \pmod{p}$, we employ a generalisation of a reduction used by Faben. In [8, Theorem 3.7] Faben reduces $\#_p\text{SAT}$ to $\#_p\text{BIS}_{1,1}$, the problem of counting independent sets of a bipartite graph.

We have to generalise this reduction for the weighted setting, in particular allowing different vertex weights for the vertices of each partition. Furthermore, during the construction we have to keep track of the assignment of vertices to their corresponding part, V_L or V_R . For this purpose, we need to show the existence of bipartite graphs B , where $Z_{\lambda_\ell, \lambda_r}(B)$ takes specific values. These graphs are then used as gadgets in our reduction. In the unweighted setting $\#_p\text{BIS}_{1,1}$ the graphs B are complete bipartite graphs. However, in the weighted setting $\#_p\text{BIS}_{\lambda_\ell, \lambda_r}$ complete bipartite graphs are not sufficient. Therefore, we prove the existence of the necessary bipartite gadgets B constructively. The technical proofs appear in Section 3.

Pinning. Similarly to the existing hardness proofs on modular counting graph homomorphisms, we deploy a “pinning” technique. A *partial function* from a set X to a set Y is a function $f : X' \rightarrow Y$ for some $X' \subseteq X$. For any graph H a *partially H -labelled graph* $J = (G, \tau)$ consists of an *underlying graph* $G = G(J)$ and a *pinning function* $\tau = \tau(J)$, which is a partial function from $V(G)$ to $V(H)$. A homomorphism from a partially labelled graph $J = (G, \tau)$ to H is a homomorphism $\sigma : G \rightarrow H$, such that for all vertices $v \in \text{dom}(\tau)$ it holds $\sigma(v) = \tau(v)$. The resulting problem is denoted by $\#_p\text{PARTLABHOMSTOH}$, that is, given a prime p and graph H , compute $|\text{Hom}(J \rightarrow H)| \pmod{p}$. In Section 5, we show that $\#_p\text{PARTLABHOMSTOH}$ reduces to $\#_p\text{HOMSTOH}$ when H has no automorphisms of order p . This allows us to establish hardness for $\#_p\text{HOMSTOH}$ by proving hardness for $\#_p\text{PARTLABHOMSTOH}$. The reduction generalises the pinning reduction of Göbel, Goldberg and Richerby [14] from $\#_2\text{PARTLABHOMSTOH}$ to $\#_2\text{HOMSTOH}$.

To illustrate how we reduce $\#_2\text{PARTLABHOMSTOH}$ to $\#_2\text{HOMSTOH}$, we restrict the value of the modulo to 2 and the pinning function $\tau(J) = \{u \mapsto v\}$ to “pin” a single vertex. Given two graphs with distinguished vertices (G, u) and (H, v) , let $\text{Hom}((G, u) \rightarrow (H, v))$ be the set of homomorphisms from G to H mapping u to v . We define $\mathbf{w}_H(G)$ to be the $\{0, 1\}$ -vector containing the entries $|\text{Hom}((G, u) \rightarrow (H, v))| \pmod{2}$ for each vertex $v \in V(H)$. Observe that for two vertices $v_1, v_2 \in V(H)$, such that $(H, v_1) \cong (H, v_2)$, i.e., there is an automorphism of H mapping v_1 to v_2 , and any graph G the relevant entries in $\mathbf{w}_G(H)$ will always be equal. Therefore, we can contract all such entries to obtain the *orbit vectors* $\mathbf{v}_H(G)$. Suppose that there exists a graph with a distinguished vertex (Θ, u_Θ) , such that $\mathbf{v}_H(\Theta) = 0 \dots 010 \dots 0$, where the 1-entry corresponds to the vertex v of H . Given our input J for $\#_2\text{PARTLABHOMSTOH}$, we can now define an input G for $\#_2\text{HOMSTOH}$, such that $|\text{Hom}(J \rightarrow H)| \equiv |\text{Hom}((G(J), u) \rightarrow (H, v))| \equiv |\text{Hom}(G \rightarrow H)| \pmod{2}$. G contains a disjoint copy of $G(J)$ and Θ , where the vertices u and u_Θ are identified (recall that u is the vertex of J mapped by $\tau(J)$). Due to the value of $\mathbf{v}_H(\Theta)$ and the structure of G there is an even number of homomorphisms mapping u to any vertex $v' \neq v$, which establishes the claim.

Such a graph Θ , however, is not guaranteed to exist. Instead, we can define a set of operations on the vectors $\mathbf{v}_H$ corresponding to graph operations and show that for any vector in $\{0, 1\}^{|V(H)|}$ there exists a sequence of graphs with distinguished vertices $(\Theta_1, u_1), \dots, (\Theta_t, u_t)$ that “generate” this vector. Thus, there exists a set of graphs that “generate” $\mathbf{v} = 0 \dots 010 \dots 0$, which yields the desired reduction. This technique of [14] exploits the value of the modulo to be 2. Directly applying this technique to counting modulo any prime p , we can only establish pinning for asymmetric graphs, that is graphs whose automorphism group contains only the identity. A dichotomy for $\#_p\text{HOMSTOH}$ if H is an asymmetric tree appears in the first author’s doctoral thesis [12].

To go beyond asymmetric graphs, we observe that information might be lost from the contraction of the vectors $\mathbf{w}_H$ to the vectors $\mathbf{v}_H$. We note that in asymmetric graphs these two vectors are identical. For general graphs however, we have to restore pinning for counting homomorphisms modulo any prime p by utilising the non-contracted vectors $\mathbf{w}_H$.

THEOREM 1.8. *Let p be a prime and let H be a graph with no automorphism of order p . Then $\#_p\text{PARTLABHOMSTOH}$ reduces to $\#_p\text{HOMSTOH}$ via a polynomial-time Turing reduction.*

To obtain hardness for $\#_p\text{HOMSTOH}$, we only need to pin two vertices when H is a tree, i.e., the domain of the pinning function τ has size two. For a study of a more general class of target graphs H (see [14]) the size of the domain has to be larger. As our pinning theorem applies to all primes p , all graphs H and pinning functions of arbitrary domain size, it can potentially be used to show hardness for $\#_p\text{HOMSTOH}$ for all primes and any class of target graphs H . The formal proofs appear in Section 5.

Gadgets. Gadgets are structures appearing in the target graph H that allow to reduce $\#_2\text{IS}$ to $\#_2\text{PARTLABHOMSTOH}$ (the hardness of $\#_2\text{HOMSTOH}$ is then immediate from Theorem 1.8). For

illustrative purposes, we simplify the definitions appearing [14]. $\#_2\text{HomsToH}$ -gadgets consist of two partially H -labelled graphs with distinguished vertices (J_1, y) , (J_2, y, z) along with two “special” vertices $i, o \in V(H)$. Given the input G for $\#_2\text{IS}$, we construct an input G' for $\#_2\text{PARTLABHomsToH}$ as follows. We attach a copy of J_1 to every vertex u of G (identifying u with y) and replace every edge (u, v) of G with a copy of J_2 (identifying u with y and v with z). The properties of J_1 ensure that there is an odd number of homomorphisms from G' to H if the original vertices of G are mapped to i or o , while the number of the remaining homomorphisms cancels out. The properties of J_2 ensure that there is an even number of homomorphisms from G' to H if two adjacent vertices of G are both mapped to i , and an odd number of homomorphisms in every other case. Now we observe that $|I(G)| \equiv |\text{Hom}(G' \rightarrow H)| \pmod{2}$, because the set of homomorphisms that do not cancel out must map every vertex of G to i or o and no pair of adjacent vertices both to i . Every vertex of G that is in an independent set must be mapped to i , and every vertex that is out of the independent set must be mapped to o .

Generalising the described approach to any prime modulus $p > 2$ one would end up reducing from a restricted $\#_p\text{CSP}$ instance, containing a binary relation and a unary weight that must be applied to every variable of the instance (this is known as *external field* in statistical physics). Similar to the modulo 2 case the edge interaction is captured by the binary relation and size of the set of “special” vertices by the unary weights. Since for primes $p > 2$ there are more non-zero values than 1 (odd) a study of the external field is no longer trivial in this case. Instead, we choose a different approach and reduce from $\#_p\text{BIS}_{\lambda_\ell, \lambda_r}$. This seems to capture the structure that produces hardness in $\#_p\text{HomsToH}$ in a more natural way.

We formally present our reduction in Section 6. In the following, we sketch our proof method and focus our attention on the example graph H in Figure 1. Let $G = (V_L, V_R, E)$ be a bipartite graph. Homomorphisms from G to H must respect the partition of G , i.e., the vertices in V_L can only be mapped to the vertices in $\{x_L, u_1, u_2, u_3\}$ and the vertices in V_R can only be mapped to the vertices in $\{x_R, v_1, v_2\}$, or vice versa. Any homomorphism σ from G to H , which maps the vertex $w \in V(G)$ to any vertex in $\{u_1, u_2, u_3\}$, must map every neighbour of w to x_R . Similarly, any homomorphism σ from G to H , which maps the vertex $w \in V(G)$ to any vertex in $\{v_1, v_2\}$, must map every neighbour of w to x_L . Thus, homomorphisms from G to H express independent sets of G : $\{u_1, u_2, u_3\}$ represent the vertices of V_L in the independent set and $\{v_1, v_2\}$ represent the vertices of V_R in the independent set, or vice versa. We construct a partially H -labelled graph J from G to fix the choice of V_L and V_R in the set of homomorphisms from G to H . $G(J)$ contains a copy of G together with two new vertices $\hat{u}, \hat{v}$, where every vertex in V_L is attached to the new vertex $\hat{u}$ and every vertex in V_R is attached to the new vertex $\hat{v}$. In addition, $\tau(J) = \{\hat{u} \mapsto x_R, \hat{v} \mapsto x_L\}$ is the pinning function. We observe that the vertices in V_L can only be mapped to vertices in $\{x_L, u_1, u_2, u_3\}$ and vertices in V_R can only be mapped to vertices in $\{x_R, v_1, v_2\}$. This observation yields that the number of homomorphisms from J to H is equivalent to $\sum_{I \in I(G)} 3^{|V_L \cap I|} 2^{|V_R \cap I|}$ modulo p . Furthermore, the cardinality of the sets $\{u_1, u_2, u_3\}$ and $\{v_1, v_2\}$ introduces weights in a natural way.

For the reduction above, we need the following property easily observable in H : There exist two adjacent vertices of degree $a = \lambda_\ell + 1 \not\equiv 1 \pmod{p}$ and $b = \lambda_r + 1 \not\equiv 1 \pmod{p}$. Recall that to obtain hardness for $\#_p\text{BIS}_{\lambda_\ell, \lambda_r}$ Theorem 1.6 requires $\lambda_\ell, \lambda_r \not\equiv 0 \pmod{p}$. In fact, as we will show in Section 6, these vertices need not be adjacent. During the construction of J , we can replace the edges of G with paths of appropriate length. We call such a structure in H an (a, b, p) -path. In Lemma 6.7, we formally prove that if H has an (a, b, p) -path, then $\#_p\text{HomsToH}$ is $\#_p$ P-hard. In particular, observe that stars cannot contain (a, b, p) -paths. Finally, we show that every non-star tree H contains an (a, b, p) -path, which yields our main result on $\#_p\text{HomsToH}$ (Lemma 6.2).

1.4 Composites

We outline the obstacles occurring when extending the dichotomy for $\#_k\text{HomsToH}$ to any integer k . Let H be a graph and let k be an integer with prime factorisation $k = \prod_{i=1}^m k_i$, where $k_i = p_i^{r_i}$. Assuming $\#_k\text{HomsToH}$ can be solved in polynomial time, then for each $i \in [m]$, $\#_{k_i}\text{HomsToH}$ can also be solved in polynomial time. The reason is that k_i is a factor of k , and we can apply the modulo k_i operator to the answer for the $\#_k\text{HomsToH}$ instance. The Chinese remainder theorem shows that the converse is also true: If for each $i \in [m]$, then we can solve $\#_{k_i}\text{HomsToH}$ in polynomial time, then we can also solve $\#_k\text{HomsToH}$ in polynomial time. By the previous observations, we can now focus on powers of primes $k = p^r$. Assuming $\#_k\text{HomsToH}$ is computable in polynomial time yields again that $\#_p\text{HomsToH}$ is also computable in polynomial time. However, the converse is not always true.

Guo et al. [17] were able to obtain this reverse implication for the constraint satisfaction problem. They showed [17, Lemma 4.1 and Lemma 4.3] for p a prime that $\#_{p^r}\text{CSP}$ is computable in polynomial time if $\#_p\text{CSP}$ is computable in polynomial time. In Section 8, we show that their technique cannot be transferred to the $\#_k\text{HomsToH}$ setting. We show that there is a graph (P_4) such that $\#_2\text{HomsToP}_4$ is computable in polynomial time, while $\#_4\text{HomsToP}_4$ is $\#_2$ P-hard.

1.5 Organisation

Our notation is introduced in Section 2. In Section 3, we study the complexity of the weighted bipartite independent sets problem modulo any prime. Section 4 presents the connection to the polynomial time algorithm of Faben and Jerrum for $\#_p\text{HomsToH}$. Our pinning method is explained in Section 5. Section 6 contains the hardness reduction for $\#_p\text{HomsToH}$. Our results are collected into a dichotomy theorem in Section 7. Finally, in Section 8 we discuss the obstacles arising when counting modulo any integer.

2 PRELIMINARIES

We denote by $[n]$ the set $\{1, \dots, n\}$. Further, if v is an element of the set S , then we write $S - v$ for $S \setminus \{v\}$. We denote the composition of two functions f and g by $(f \circ g)(x) = f(g(x))$. Let k be a positive integer $k \in \mathbb{Z}_{>0}$, then for a function f its k -fold composition is denoted by $f^{(k)} = f \circ f \circ \dots \circ f$.

For a detailed introduction to graph theory and the used notation the reader is referred to Reference [27].

(Simple) graphs. Unless otherwise specified, *graphs* are undirected and simple, requiring them to contain neither parallel edges nor loops. For a graph G , we denote its vertex set by $V(G)$ and its edge set by $E(G)$. For all vertices $v \in V(G)$ of a graph G with a subgraph H , we denote by $\Gamma_H(v) = \{u \in V(H) \mid (u, v) \in E\}$ the *neighbourhood of v in H* containing all vertices in $V(H)$ adjacent to v , and we denote by $\deg_H(v)$ the size of $\Gamma_H(v)$. Paths in graphs do not repeat vertices; walks may repeat both vertices and edges. The *distance of two connected vertices u, v in G* , denoted by $d_G(u, v)$, is the length of a shortest path in G connecting u and v . An *independent set* of a graph G is a set of vertices $I \subseteq V(G)$, such that no pair of vertices in I is adjacent in G . We denote the *set of independent sets of G* by $\mathcal{I}(G)$. We write $G = (V_L, V_R, E)$ for the bipartite graph with fixed bipartition V_L and V_R .

Graph homomorphisms. Let G and H be graphs. A *homomorphism from G to H* is a function $\sigma : V(G) \rightarrow V(H)$, such that $(v_1, v_2) \in E(G)$ implies $(\sigma(v_1), \sigma(v_2)) \in E(H)$. Moreover, $\text{Hom}(G \rightarrow H)$ denotes the set of homomorphisms from G to H . An *isomorphism between G and H* is a bijective function $\varphi : V(G) \rightarrow V(H)$ preserving the edge relation in both directions, meaning $(v_1, v_2) \in E(G)$

if and only if $(\varrho(v_1), \varrho(v_2)) \in E(H)$. If such an isomorphism exists, then we say that G is *isomorphic to H* and denote it by $G \cong H$. An *automorphism of G* is an isomorphism from the graph G to itself. $\text{Aut}(G)$ denotes the *automorphism group of G* . An automorphism ϱ is an *automorphism of order k* in case it is not the identity and k is the smallest positive integer such that $\varrho^{(k)}$ is the identity.

Partially labelled graphs. Let H be a graph. A *partially H -labelled graph* $J = (G, \tau)$ consists of an *underlying graph* $G(J) = G$ and a (partial) *pinning function* $\tau(J) = \tau : V(G) \rightarrow V(H)$, mapping vertices in G to vertices in H . Every vertex v in the domain $\text{dom}(\tau)$ of τ is said to be *H -pinned to $\tau(v)$* . We omit H in case it is immediate from the context. We denote a partial function τ with finite domain $\{v_1, \dots, v_r\}$ also in the form $\tau = \{v_1 \mapsto \tau(v_1), \dots, v_r \mapsto \tau(v_r)\}$. A *homomorphism from a partially labelled graph J to a graph H* is a homomorphism from $G(J)$ to H that respects τ , i.e., for all $v \in \text{dom}(\tau)$ holds $\sigma(v) = \tau(v)$. By $\text{Hom}(J \rightarrow H)$, we denote the set of homomorphisms from J to H that respect the labelling.

Graphs with distinguished vertices. Let G and H be graphs. It is often convenient to regard a graph with a number of (not necessarily distinct) distinguished vertices $v_1, \dots, v_r$, which we denote by $(G, v_1, \dots, v_r)$. A *sequence of vertices* $v_1 \dots v_r$ may be abbreviated by $\bar{v}$ and $G[\bar{v}]$ stands for the subgraph of G induced by the set of vertices $\{v_1, \dots, v_r\}$. A *homomorphism from $(G, \bar{u})$ to $(H, \bar{v})$* with $r = |\bar{u}| = |\bar{v}|$ is a homomorphism σ from G to H with $\sigma(u_i) = v_i$ for each $i \in [r]$. Such a homomorphism immediately yields a homomorphism from the partially labelled graph $(G, \{u_1 \mapsto v_1, \dots, u_r \mapsto v_r\})$ to H and vice versa. For a partially labelled graph J and vertices $u_1, \dots, u_r \notin \text{dom}(\tau(J))$, we identify a homomorphism from $(J, \bar{u})$ to $(H, \bar{v})$ with the corresponding homomorphism from $(G(J), \tau(J) \cup \{u_1 \mapsto v_1, \dots, u_r \mapsto v_r\})$ to H . Similarly, $(G, \bar{u})$ and $(H, \bar{v})$ are *isomorphic* if $r = |\bar{u}| = |\bar{v}|$ and there is an isomorphism ϱ from G to H , such that $\varrho(u_i) = v_i$ for each $i \in [r]$. An *automorphism of $(G, \bar{u})$* is an automorphism ϱ of G with the property that $\varrho(u_i) = u_i$ for each $i \in [r]$ and $\text{Aut}(G, \bar{u})$ denotes the *automorphism group of $(G, \bar{u})$* .

Graph constructions. We often describe graph constructions by the operation of combining copies of two (or more) given graphs G_1, G_2 into a new graph by G by *identifying* a vertex $v_1 \in G_1$ with a vertex $v_2 \in G_2$ and naming this vertex with a new name, say v . This formally gives $V(G) = (\{v\} \cup V(G_1) \cup V(G_2)) \setminus \{v_1, v_2\}$ and $E(G) = \bigcup_{j=1}^2 \{(v, u) \mid (v_j, u) \in E(V_j)\} \cup (E(G_1) \setminus \{(v_1, u) \mid (v_1, u) \in E(V_1)\}) \cup (E(G_2) \setminus \{(v_2, u) \mid (v_2, u) \in E(V_2)\})$. If we use such a construction on two partially H -labelled graphs $J_1 = (G_1, \tau_1)$ and $J_2 = (G_2, \tau_2)$ (for the same graph H) with $v_j \notin \text{dom}(\tau_j)$ for $j = 1, 2$, then this creates a new partially H -labelled graph $J = (G, \tau)$, where G is the combination of G_1, G_2 as described above and $\tau(u) = \tau_j(u)$ for all $u \in \text{dom}(\tau_j)$; $j = 1, 2$ (hence for a vertex u in the domain of τ_j , we also copy its mapping into the new graph G). Similarly, we perform such combinations with graphs that have distinguished vertices. We choose and list whichever vertices we want to be distinguished in the new graph G .

Basic algebra. For an introduction to abstract algebra, we refer the reader to [5]. Finally, we assume familiarity with the notion of a *group*, an *action of a group on a set* and modular arithmetic in the *field* $\mathbb{Z}_p$, where p is a *prime* in $\mathbb{Z}$. We denote with $\mathbb{Z}_p^* = \mathbb{Z}_p \setminus \{0\}$. We are going to apply Fermat's little theorem (see [1, Theorem 11.6]) and Cauchy's group theorem (see, e.g., [1, Theorem 13.1]) frequently.

THEOREM 2.1 (FERMAT'S LITTLE THEOREM). *Let p be a prime. If $a \in \mathbb{Z}$ is not a multiple of p , then $a^{p-1} \equiv 1 \pmod{p}$.*

THEOREM 2.2 (CAUCHY'S GROUP THEOREM). *Let p be prime. If $\mathcal{G}$ is a finite group and p divides $|\mathcal{G}|$, then $\mathcal{G}$ contains an element of order p .*

3 WEIGHTED BIPARTITE INDEPENDENT SET

We study the complexity of computing the weighted sum over independent sets in a bipartite graph modulo a prime. Note that the set of independent sets of a graph does not change if the graph contains multiedges and that a bipartite graph cannot contain loops. For this reason, in this section we do not have to distinguish between a bipartite multigraph or a bipartite simple graph. For the unweighted version Faben [8, Theorem 3.7] showed that the problem $\#_k \text{BIS}$ of counting the independent sets of a graph modulo any integer k is hard, even when the input graph is restricted to be bipartite.

THEOREM 3.1 (FABEN). *For all positive integers k , $\#_k \text{BIS}$ is $\#_k \text{P-complete}$.*

To define the weighted version let p be a prime, $G = (V_L, V_R, E)$ be a bipartite graph with given bipartition and $\lambda_\ell, \lambda_r \in \mathbb{Z}_p$ be the weights contributed by vertices depending on the part they belong to. Analogously to Faben, we denote by $\#_p \text{BIS}_{\lambda_\ell, \lambda_r}$ the problem of computing the following weighed sum over independent sets of a bipartite graph $G = (V_L, V_R, E)$ modulo p

$$Z_{\lambda_\ell, \lambda_r}(G) = \sum_{I \in \mathcal{I}(G)} \lambda_\ell^{|V_L \cap I|} \lambda_r^{|V_R \cap I|}.$$

We note that every bipartite graph has a bipartition V_L, V_R and declaring a bipartite graph with $G = (V_L, V_R, E)$ is the same as having the graph G along with the bipartition as input. A fixed bipartition is necessary when studying weighted independent sets, since changing the bipartitioning changes the value of the weighted sum. In the unweighted sum of Theorem 3.1, there is no need to give a fixed partition as input, as it does not change the number of independent sets. Moreover, $\#_p \text{BIS}_{1,1}$ corresponds to the special case $\#_p \text{BIS}$ and Theorem 3.1 directly implies that $\#_p \text{BIS}_{1,1}$ is $\#_p \text{P-complete}$ for all primes p .

Formally, we study the complexity of the following problem.

PROBLEM 3.2. *Name.* $\#_p \text{BIS}_{\lambda_\ell, \lambda_r}$.

Parameter. p prime and $\lambda_\ell, \lambda_r \in \mathbb{Z}_p$.

Input. Bipartite graph $G = (V_L, V_R, E)$.

Output. $Z_{\lambda_\ell, \lambda_r}(G) \pmod{p}$.

We begin by identifying the tractable instances of $\#_p \text{BIS}_{\lambda_\ell, \lambda_r}$.

PROPOSITION 3.3. *If $\lambda_\ell \equiv 0 \pmod{p}$ or $\lambda_r \equiv 0 \pmod{p}$, then $\#_p \text{BIS}_{\lambda_\ell, \lambda_r}$ is computable in polynomial time.*

PROOF. Without loss of generality, we assume $\lambda_\ell \equiv 0 \pmod{p}$. Thus, any independent set that contains at least one vertex from V_L contributes zero to the sum in $Z_{\lambda_\ell, \lambda_r}(G)$. Therefore, we only need to consider the independent sets I with $I \subseteq V_R$. Since any subset of V_R yields an independent set, we obtain

$$\begin{aligned} Z_{\lambda_\ell, \lambda_r}(G) &\equiv 1 + \sum_{i=1}^{|V_R|} \binom{|V_R|}{i} (\lambda_r)^i \pmod{p} \\ &= \sum_{i=0}^{|V_R|} \binom{|V_R|}{i} (\lambda_r)^i = (\lambda_r + 1)^{|V_R|}, \end{aligned}$$

which can be computed in polynomial time. $\square$

The remainder of the section is dedicated to proving that $\#_p \text{BIS}_{\lambda_\ell, \lambda_r}$ is hard in all other cases. Our reduction is inspired by the reduction of Faben [8, Theorem 3.7].

To avoid double counting in the following proofs, we define a partition of the independent sets.

Definition 3.4. Let $G = (V_L, V_R, E)$ be a bipartite graph. We denote by $\mathcal{I}_L(G)$ the set $\{I \in \mathcal{I}(G) \setminus \{\emptyset\} \mid I \subseteq V_L\}$ of non-empty independent sets containing only vertices from V_L . Similarly, we write $\mathcal{I}_R(G)$ for the set of non-empty independent sets that contain only vertices from V_R . Finally, we denote by $\mathcal{I}_{LR}(G)$ the set $\mathcal{I}(G) \setminus (\mathcal{I}_L(G) \cup \mathcal{I}_R(G) \cup \{\emptyset\})$ of independent sets containing at least one vertex in V_L and at least one vertex in V_R .

The following lemma expresses $Z_{\lambda_\ell, \lambda_r}(G)$ in terms of the partitioning defined above.

LEMMA 3.5. *Let $G = (V_L, V_R, E)$ be a bipartite graph. Then,*

$$Z_{\lambda_\ell, \lambda_r}(G) = (\lambda_\ell + 1)^{|V_L|} + (\lambda_r + 1)^{|V_R|} - 1 + \sum_{I \in \mathcal{I}_{LR}(G)} \lambda_\ell^{|V_L \cap I|} \lambda_r^{|V_R \cap I|}.$$

PROOF. By Definition 3.4 the set $\mathcal{I}(G)$ partitions into $\{\mathcal{I}_L(G), \mathcal{I}_R(G), \mathcal{I}_{LR}(G), \{\emptyset\}\}$, which yields

$$\begin{aligned} Z_{\lambda_\ell, \lambda_r}(G) &= \sum_{I \in \mathcal{I}(G)} \lambda_\ell^{|V_L \cap I|} \lambda_r^{|V_R \cap I|} \\ &= \sum_{I \in \mathcal{I}_L(G)} \lambda_\ell^{|I|} + \sum_{I \in \mathcal{I}_R(G)} \lambda_r^{|I|} + \sum_{I \in \mathcal{I}_{LR}(G)} \lambda_\ell^{|V_L \cap I|} \lambda_r^{|V_R \cap I|} + 1. \end{aligned} \quad (1)$$

As in the proof of Proposition 3.3, we obtain

$$\sum_{I \in \mathcal{I}_L(G)} \lambda_\ell^{|I|} = \sum_{i=0}^{|V_L|} \binom{|V_L|}{i} \lambda_\ell^i - 1 = (\lambda_\ell + 1)^{|V_L|} - 1, \quad \text{and analogously} \quad (2)$$

$$\sum_{I \in \mathcal{I}_R(G)} \lambda_r^{|I|} = (\lambda_r + 1)^{|V_R|} - 1. \quad (3)$$

Inserting (3) and (2) into (1) yields the lemma. $\square$

For our reduction to work, we must design gadgets that are tailored to our general setting of weighted independent sets.

Definition 3.6. Let p be a prime and let $\lambda_\ell, \lambda_r \in \mathbb{Z}_p^*$.

For every $k \in [p]$, we denote by $B(k, p) = (V_L, V_R, E)$ the bipartite graph with $4(p-1)$ vertices in two disjoint vertex sets $V_L := \{u_1, \dots, u_{2(p-1)}\}$, $V_R := \{v_1, \dots, v_{2(p-1)}\}$ and the edge set

$$E := \{(u_i, v_j) \mid i, j \in [2(p-1)], \text{ where } i \neq j\} \cup \{(u_i, v_i) \mid i \notin [k]\},$$

consisting of all edges in the complete bipartite graph $K_{2(p-1), 2(p-1)}$ except (u_i, v_i) with $i \in [k]$.

See Figure 2 for the exemplary graph $B(1, 3)$.

$B(k, p)$ has two types of vertices in each part: the vertices in $\{u_i, v_i\}_{i \leq k}$ of degree $2(p-1) - 1$ and the vertices in $\{u_i, v_i\}_{i > k}$ of degree $2(p-1)$. Since the size of the vertex sets is a multiple of $(p-1)$, we are able to apply Fermat's little Theorem 2.1 in our reductions later on. Moreover, the size of the gadget is large enough to generate every necessary value of $k \in [p]$. This freedom of choice for k will entail the possibility, given $\lambda_\ell, \lambda_r \not\equiv 0 \pmod{p}$, to choose k such that $Z_{\lambda_\ell, \lambda_r}(B(k, p)) \equiv 0 \pmod{p}$. Given such a k , we will see that in each part there exists a vertex v such that removing this vertex from $B(k, p)$ will yield $Z_{\lambda_\ell, \lambda_r}(B(k, p) - v) \not\equiv 0 \pmod{p}$. This property will be crucial later on.

The following lemma establishes the key properties of the $B(k, p)$ defined above and will be later used to show the crucial properties of our reduction gadgets.

Fig. 2. Constructive route for $p = 3$ and $k = 1$. Starting with the complete bipartite graph $K_{4,4}$ the edge (u_1, v_1) is removed.

LEMMA 3.7. *Let p be a prime, $\lambda_\ell, \lambda_r \in \mathbb{Z}_p^*$, $k \in \mathbb{Z}_p$ and $B = B(k, p)$ as in Definition 3.6. Then,*

$$\sum_{I \in \mathcal{I}_{LR}(B)} \lambda_\ell^{|V_L \cap I|} \lambda_r^{|V_R \cap I|} \equiv k \lambda_\ell \lambda_r \pmod{p}.$$

PROOF. Let $I \in \mathcal{I}_{LR}(B)$ be a non-empty independent set containing a vertex $u_i \in V_R$ and a vertex $v_j \in V_L$. By the definition of B there is no independent set containing two vertices u_i and v_j with $i \neq j$. Thus $i = j$ and $V_L \cap I = \{u_i\}$ as well as $V_R \cap I = \{v_i\}$. We obtain $I = \{u_i, v_i\}$ yielding $\mathcal{I}_{LR} = \{\{u_i, v_i\} \mid i \in [k]\}$. $\square$

The following lemma states the properties of the graphs we will use as gadgets, namely a copy of a $B(k, p)$ for an appropriately chosen $k \in [p]$, together with two distinguished vertices.

LEMMA 3.8. *Let p be a prime and $\lambda_\ell, \lambda_r \in \mathbb{Z}_p^*$. There exists a bipartite graph $B = (V_L, V_R, E)$ with distinguished vertices $u_L \in V_L$ and $v_R \in V_R$, that satisfies*

- (1) $Z_{\lambda_\ell, \lambda_r}(B) \equiv 0 \pmod{p}$,
- (2) $Z_{\lambda_\ell, \lambda_r}(B - u_L) \not\equiv 0 \pmod{p}$,
- (3) $Z_{\lambda_\ell, \lambda_r}(B - v_R) \not\equiv 0 \pmod{p}$.

PROOF. For every graph $B = B(k, p)$, we apply Lemma 3.5 to obtain

$$\begin{aligned} Z_{\lambda_\ell, \lambda_r}(B) &= (\lambda_\ell + 1)^{|V_L|} + (\lambda_r + 1)^{|V_R|} - 1 + \sum_{I \in \mathcal{I}_{LR}} \lambda_\ell^{|V_L \cap I_{LR}|} \lambda_r^{|V_R \cap I_{LR}|} \\ &= (\lambda_\ell + 1)^{2(p-1)} + (\lambda_r + 1)^{2(p-1)} - 1 + \sum_{I \in \mathcal{I}_{LR}} \lambda_\ell^{|V_L \cap I_{LR}|} \lambda_r^{|V_R \cap I_{LR}|}. \end{aligned} \quad (4)$$

If one of the weights is equivalent to -1 in $\mathbb{Z}_p$, then the corresponding term in (4) vanishes. Otherwise, we are allowed to apply Fermat's little Theorem 2.1 and the corresponding term is equivalent to 1. Therefore, we have to distinguish cases.

i. $\lambda_\ell, \lambda_r \not\equiv -1 \pmod{p}$.

We choose $k \equiv -(\lambda_\ell \lambda_r)^{-1} \pmod{p}$, $u_L = u_{2(p-1)}$ and $v_L = v_{2(p-1)}$. Note that such a k uniquely exists, since p is a prime and $\mathbb{Z}_p$ a field. Applying Lemma 3.7 and Fermat's little Theorem 2.1 on (4) yields

$$Z_{\lambda_\ell, \lambda_r}(B) \equiv 1 + k \lambda_\ell \lambda_r \equiv 0 \pmod{p}.$$

Removing any of the chosen two vertices from $V(B)$ does not affect the independent sets in $\mathcal{I}_{LR}$ and thus

$$\begin{aligned} Z_{\lambda_\ell, \lambda_r}(B - u_L) &\equiv (\lambda_\ell + 1)^{2(p-1)-1} - 1 \equiv (\lambda_\ell + 1)^{-1} - 1 \pmod{p}; \\ Z_{\lambda_\ell, \lambda_r}(B - v_R) &\equiv (\lambda_r + 1)^{2(p-1)-1} - 1 \equiv (\lambda_r + 1)^{-1} - 1 \pmod{p}. \end{aligned}$$

Neither of these expressions is equivalent to 0 in $\mathbb{Z}_p$, since both weights are in $\mathbb{Z}_p^*$.

ii. $\lambda_\ell \equiv -1 \pmod{p}$, $\lambda_r \not\equiv -1 \pmod{p}$.

We choose $k = p$, $u_L = u_k$ and $v_R = v_{2(p-1)}$. Analogously to the first case, we obtain due to the choice of k

$$Z_{\lambda_\ell, \lambda_r}(B) \equiv k\lambda_\ell\lambda_r \equiv 0 \pmod{p}.$$

Similarly to the observation in the first case, we have due to the choice of v_R ,

$$Z_{\lambda_\ell, \lambda_r}(B - v_R) \equiv (\lambda_r + 1)^{2(p-1)-1} - 1 \equiv (\lambda_r + 1)^{-1} - 1 \not\equiv 0 \pmod{p}.$$

We note that the edge (u_k, v_k) is missing in B and therefore the set $\{u_k, v_k\}$ is in $\mathcal{I}_{LR}(B)$. Due to the choice of u_L , we deduce $\mathcal{I}_{LR}(B - u_L) = \mathcal{I}_{LR}(B) - \{u_k, v_k\}$ and thus

$$Z_{\lambda_\ell, \lambda_r}(B - u_L) \equiv \sum_{I \in \mathcal{I}_{LR}(B - u_L)} \lambda_\ell^{|V_L \cap I_{LR}|} \lambda_r^{|V_R \cap I_{LR}|} = (k-1)\lambda_\ell\lambda_r \equiv \lambda_r \not\equiv 0 \pmod{p}.$$

iii. $\lambda_\ell \not\equiv -1$, $\lambda_r \equiv -1 \pmod{p}$.

The proof of this case is similar to the second case. In particular, choosing $k = p$ as well as $u_L = u_{2(p-1)}$ and $v_R = v_k$ establishes this case.

iv. $\lambda_\ell, \lambda_r \equiv -1 \pmod{p}$.

We choose $k = 1$, $u_L = u_k$ and $v_R = v_k$ and obtain

$$Z_{\lambda_\ell, \lambda_r}(B) \equiv -1 + k\lambda_\ell\lambda_r \equiv 0 \pmod{p}.$$

The particular choice of u_L and v_R has the same effect on $\mathcal{I}_{LR}$ as in the previous two cases, and we deduce $\sum_{I \in \mathcal{I}_{LR}(B - u_L)} \lambda_\ell^{|V_L \cap I_{LR}|} \lambda_r^{|V_R \cap I_{LR}|} = (k-1)\lambda_\ell\lambda_r \equiv 0 \pmod{p}$. Hence,

$$\begin{aligned} Z_{\lambda_\ell, \lambda_r}(B - u_L) &= -1 + \sum_{I \in \mathcal{I}_{LR}(B - u_L)} \lambda_\ell^{|V_L \cap I_{LR}|} \lambda_r^{|V_R \cap I_{LR}|} = -1, \quad \text{and analogously} \\ Z_{\lambda_\ell, \lambda_r}(B - v_R) &= -1. \end{aligned}$$

This establishes the lemma. $\square$

Regarding the starting problem for the reduction, given a Boolean formula φ then $\text{sat}(\varphi)$ denotes the set of the satisfying assignments of φ .

PROBLEM 3.9. *Name.* $\#_k\text{SAT}$.

Parameter. k integer.

Input. Boolean formula φ in conjunctive normal form.

Output. $|\text{sat}(\varphi)| \pmod{k}$.

Simon in his thesis [24, Theorem 4.1] shows how the original reduction of Cook can be made parsimonious. As Faben observes [9, Theorem 3.1.17] any parsimonious reduction is parsimonious modulo k , for any integer k , hence $\#_k\text{SAT}$ is $\#_k$ P-complete.

THEOREM 3.10 (SIMON). $\#_k\text{SAT}$ is $\#_k$ P-complete under parsimonious reductions for all integers k .

Fig. 3. The graphs G'_φ and G_φ for $\varphi = (\mathbf{x}_1 \vee \mathbf{x}_2) \wedge (\bar{\mathbf{x}}_1 \vee \mathbf{x}_2 \vee \mathbf{x}_3)$.

Our reduction starts from a Boolean formula φ with n literals and m clauses, input for $\#_p\text{SAT}$, and constructs in two stages a graph G_φ , that is an input for $\#_p\text{BIS}_{\lambda_\ell, \lambda_r}$.

In the first stage, we define the graph G'_φ . For every variable x_i in φ , G'_φ contains three vertices u_i , $\bar{u}_i$, and w_i in the left vertex set $V_L(G'_\varphi)$ as well as three vertices v_i , $\bar{v}_i$, and z_i in the right vertex set $V_R(G'_\varphi)$. Furthermore, for every clause c_j of φ , G'_φ contains a vertex y_j in the right vertex set $V_R(G'_\varphi)$. Regarding the edges, for every variable x_i in φ , we introduce the edges forming the cycle $u_i v_i w_i \bar{v}_i \bar{u}_i z_i u_i$ to $E(G'_\varphi)$. Additionally for all $i \in [n]$, if x_i appears as a literal in clause c_j of φ , we introduce the edge (u_i, y_j) in G'_φ , and if $\bar{x}_i$ appears as a literal in clause c_j , then we introduce the edge $(\bar{u}_i, y_j)$ in G'_φ . The left part of Figure 3 illustrates an example of this construction. Formally, G'_φ is defined as follows.

Definition 3.11. Let φ be a Boolean formula in conjunctive normal form with variables $x_1, \dots, x_n$ and clauses $c_1, \dots, c_m$. The bipartite graph $G'_\varphi = (V_L(G'_\varphi), V_R(G'_\varphi), E(G'_\varphi))$ is defined by

$$\begin{aligned}
 V_L(G'_\varphi) &= \{ u_i, \bar{u}_i, w_i \mid i \in [n] \}, \\
 V_R(G'_\varphi) &= \{ v_i, \bar{v}_i, z_i \mid i \in [n] \} \cup \{ y_j \mid j \in [m] \} \text{ and} \\
 E(G'_\varphi) &= \{ (u_i, v_i), (w_i, v_i), (w_i, \bar{v}_i), (\bar{u}_i, \bar{v}_i), (\bar{u}_i, z_i), (u_i, z_i) \mid i \in [n] \} \\
 &\quad \cup \{ (u_i, y_j) \mid i \in [n], j \in [m] \text{ and } x_i \text{ occurs in } c_j \} \\
 &\quad \cup \{ (\bar{u}_i, y_j) \mid i \in [n], j \in [m] \text{ and } \bar{x}_i \text{ occurs in } c_j \}.
 \end{aligned}$$

Note that G'_φ is bipartite, since there are no adjacent vertices in either part. In the second and final stage, we construct the graph G_φ from G'_φ and $2n + m$ copies of the graph B that is provided by Lemma 3.8 once we have given a prime p and two weights $\lambda_\ell, \lambda_r \in \mathbb{Z}_p^*$. Then, for every $i \in [n]$ we adjoin a copy of B to the vertices $\{w_i, z_i\}$ in G'_φ associated to the literal i , where for positive variables w_i is identified with $u_L \in V(B)$ and for negative variables z_i is identified with $v_R \in V(B)$. Additionally, for every $j \in [m]$ we adjoin a copy of B to the vertex y_j in G'_φ associated to the clause c_j using $v_R \in V(B)$ again. For an example see the right part of Figure 3. Formally, we have the following definition.

Definition 3.12. Let φ be a Boolean formula in conjunctive normal form with variables $x_1, \dots, x_n$ and clauses $c_1, \dots, c_m$. Moreover, let G'_φ denote the bipartite graph from Definition 3.11 with $2n + m$ vertices. Further, let p be a prime, $\lambda_\ell, \lambda_r \in \mathbb{Z}_p^*$ and B be the bipartite graph with the distinguished vertices $u_L \in V_L(B)$ and $v_R \in V_R(B)$ as provided by Lemma 3.8.

For every $j \in [2n + m]$ denote by B^j a copy of B where every vertex $v \in V(B)$ is renamed v^j . The bipartite graph G_φ consists of the disjoint union of G'_φ and $\bigcup_{j \in [2n+m]} B^j$ with the following identifications: For all $i \in [n]$ identify w_i with u_L^i and z_i with v_R^{n+i} . For every $j \in [m]$ identify y_j with v_R^{2n+j} .

We observe that the graph G_φ is bipartite. Moreover, the identification of the vertices is such that the assignment of vertices to each part is preserved, i.e., $v \in V_L(G_\varphi)$ if and only if $v \in V_L(G'_\varphi)$ or $v \in V_L(B^j)$ for some $j \in [2n + m]$. This is justified, since vertices in $V_L(G'_\varphi)$ are identified exclusively with vertices in $V_L(B)$ and vertices in $V_R(G'_\varphi)$ are identified exclusively with vertices in $V_R(B)$ in the above construction.

We will employ the following partition.

Definition 3.13. Let φ be a Boolean formula in conjunctive normal form with n variables and m clauses and let G_φ be the associated bipartite gadget graph from Definition 3.12. We recursively define a partition $\{S_j\}_{j=0}^{2n+m}$ of $\mathcal{I}(G_\varphi)$ by

$$\begin{aligned} S_1 &:= \{I \in \mathcal{I}(G_\varphi) \mid v_1, \bar{v}_1 \notin I\} \\ S_j &:= \begin{cases} \{I \in \mathcal{I}(G_\varphi) \setminus \bigcup_{i=1}^{j-1} S_i \mid \Gamma_{G'_\varphi}(w_j) \cap I = \emptyset\} & \text{for } j \in \{2, \dots, n\}, \\ \{I \in \mathcal{I}(G_\varphi) \setminus \bigcup_{i=1}^{j-1} S_i \mid \Gamma_{G'_\varphi}(z_{j-n}) \cap I = \emptyset\} & \text{for } j \in \{n+1, \dots, 2n\}, \\ \{I \in \mathcal{I}(G_\varphi) \setminus \bigcup_{i=1}^{j-1} S_i \mid \Gamma_{G'_\varphi}(y_{j-2n}) \cap I = \emptyset\} & \text{for } j \in \{2n+1, \dots, 2n+m\}. \end{cases} \\ S_0 &:= \mathcal{I}(G_\varphi) \setminus \bigcup_{i=1}^{2n+m} S_i. \end{aligned}$$

We observe that for every $i \in [n]$ the neighbourhood in G'_φ of w_i contains only $v_i, \bar{v}_i$. Thus, for all $i \in [n]$ and for any independent set $I \in S_i$ we have $v_i, \bar{v}_i \notin I$. Similarly, we deduce that for every $i \in [n]$ and every $I \in S_{n+i}$ it holds $u_i, \bar{u}_i \notin I$. For every $j \in [m]$, S_{2n+j} does not contain independent sets of G_φ that intersect with the neighbourhood $\Gamma_{G'_\varphi}(y_j) = \{u_i \mid x_i \text{ is a literal in } c_j\} \cup \{\bar{u}_i \mid \bar{x}_i \text{ is a literal in } c_j\}$. Consequently, S_0 contains any independent set I in G_φ , such that, for all $i \in [n]$, at least one of $u_i, \bar{u}_i$ and at least one of $v_i, \bar{v}_i$ are in I and furthermore, for all $j \in [m]$, $\Gamma_{G'_\varphi}(y_j) \cap I \neq \emptyset$.

The following lemma shows that the independent sets of every set of the partition except S_0 cancel out when counting modulo p .

LEMMA 3.14. Let φ be a Boolean formula in conjunctive normal form with n variables and m clauses and let $G_\varphi = (V_L, V_R, E)$ be the associated bipartite gadget graph from Definition 3.12 as well as $\{S_j\}_{j=0}^{2n+m}$ the partition of $\mathcal{I}(G_\varphi)$ as defined in Definition 3.13. Then, for every $j \in [2n + m]$

$$\sum_{I \in S_j} \lambda_\ell^{|V_L \cap I|} \lambda_r^{|V_R \cap I|} \equiv 0 \pmod{p}.$$

PROOF. We fix $j \in [2n+m]$ and commence by defining the equivalence relation $\sim_j$ on S_j . For any two independent sets $I, I' \in S_j$, we have $I \sim_j I'$ if and only if $I \setminus V(B^j) = I' \setminus V(B^j)$. That is, I and I' are equivalent if and only if they differ solely in the vertices of B^j . We denote the $\sim_j$ -equivalence class of I by $\llbracket I \rrbracket_j$. Thus, $(\llbracket I \rrbracket_j)_{I \in S_j}$ is a partition of S_j .

Let $I_1, \dots, I_{t_j}$ be representatives from each $\sim_j$ -equivalence class. We obtain

$$\sum_{I \in S_j} \lambda_\ell^{|V_L \cap I|} \lambda_r^{|V_R \cap I|} = \sum_{s=1}^{t_j} \sum_{I \in \llbracket I_s \rrbracket_j} \lambda_\ell^{|V_L \cap I|} \lambda_r^{|V_R \cap I|}.$$

Therefore, it suffices to establish $\sum_{I \in \llbracket I_s \rrbracket_j} \lambda_\ell^{|V_L \cap I|} \lambda_r^{|V_R \cap I|} \equiv 0 \pmod{p}$ for every $s \in [t_j]$.

Let I_s be one of the representatives $I_1, \dots, I_{t_j}$ with its associated equivalence class $\llbracket I_s \rrbracket_j$. We continue by studying the set $I_B = I_s \setminus V(B^j)$ of common vertices among the independent sets of $\llbracket I_s \rrbracket_j$. Every independent set $I \in \llbracket I_s \rrbracket_j$ contains the vertices in I_B . However, let $I'_B = \{I \setminus I_B \mid I \in \llbracket I_s \rrbracket_j\}$. Since B^j is a bipartite graph and the assignment of vertices to their relative part is preserved in the construction of G_φ , we obtain

$$\sum_{I \in \llbracket I_s \rrbracket_j} \lambda_\ell^{|V_L \cap I|} \lambda_r^{|V_R \cap I|} = \lambda_\ell^{|V_L \cap I_B|} \lambda_r^{|V_R \cap I_B|} \sum_{I \in I'_B} \lambda_\ell^{|V_L(B^j) \cap I|} \lambda_r^{|V_R(B^j) \cap I|}. \quad (5)$$

Let v_B^j be the vertex of B^j that is identified with one of the vertices of G'_φ for the construction of G_φ , i.e., $v_B^j = u_L^i$ if $j \leq n$, and $v_B^j = v_R^i$ otherwise. By Definition 3.13, we observe that for any $I \in S_j$ no neighbour of x_j outside B^j is in I , whereas there is no restriction whether x_j is in I or not. Hence, any independent set $I' \in \mathcal{I}(B^j)$ yields an independent set in $\llbracket I_s \rrbracket_j$ and vice versa.

We deduce that $I'_B = \mathcal{I}(B^j)$. Therefore, the sum in the right-hand side of (5) is equal to $Z_{\lambda_\ell, \lambda_r}(B^j)$. For this, we recall that each B^j was chosen utilizing Lemma 3.8, whose Part 1 yields $Z_{\lambda_\ell, \lambda_r}(B^j) \equiv 0 \pmod{p}$. We obtain

$$\sum_{I \in \llbracket I_s \rrbracket_j} \lambda_\ell^{|V_L \cap I|} \lambda_r^{|V_R \cap I|} = \lambda_\ell^{|V_L \cap I_B|} \lambda_r^{|V_R \cap I_B|} Z_{\lambda_\ell, \lambda_r}(B^j) \equiv 0 \pmod{p},$$

which proves the lemma. $\square$

With these results at hand, we now show the main result of this section.

THEOREM 1.6. *Let p be a prime and let $\lambda_\ell, \lambda_r \in \mathbb{Z}_p$. If $\lambda_\ell \equiv 0 \pmod{p}$ or $\lambda_r \equiv 0 \pmod{p}$, then $\#_p \text{BIS}_{\lambda_\ell, \lambda_r}$ is computable in polynomial time. Otherwise, $\#_p \text{BIS}_{\lambda_\ell, \lambda_r}$ is $\#_p$ P-complete.*

PROOF. The first statement is a direct consequence of Proposition 3.3. Thus, let λ_ℓ, λ_r be in $\mathbb{Z}_p^*$. We are going to show hardness for $\#_p \text{BIS}_{\lambda_\ell, \lambda_r}$ by establishing a Turing reduction from $\#_p \text{SAT}$, which is known to be $\#_p$ P-complete by Simon's Theorem 3.10.

Let φ be a Boolean formula in conjunctive normal form with n variables and m clauses. We show that the constructed bipartite graph $G_\varphi = (V_L, V_R, E)$ from Definition 3.12 satisfies $Z_{\lambda_\ell, \lambda_r}(G_\varphi) \equiv K |\text{sat}(\varphi)| \pmod{p}$ for some constant $K \not\equiv 0 \pmod{p}$. The exact value of K depends on the values of the weights corresponding to the cases in the proof of Lemma 3.8.

Based on the partition $\{S_j\}_{j=0}^{2n+m}$ given by Definition 3.13, we obtain

$$Z_{\lambda_\ell, \lambda_r}(G_\varphi) = \sum_{j=0}^{2n+m} \sum_{I \in S_j} \lambda_\ell^{|V_L \cap I|} \lambda_r^{|V_R \cap I|}.$$

By Lemma 3.14 every term is equivalent to 0 in $\mathbb{Z}_p$ *except* the one regarding S_0 . This yields

$$Z_{\lambda_\ell, \lambda_r}(G_\varphi) \equiv \sum_{I \in S_0} \lambda_\ell^{|V_L \cap I|} \lambda_r^{|V_R \cap I|} \pmod{p}. \quad (6)$$

As in the proof of Lemma 3.14, we are going to use an equivalence relation $\sim_0$ along with the associated equivalence classes $\llbracket I_s \rrbracket_0$. We define $U := \{u_i, \bar{u}_i, v_i, \bar{v}_i \mid i \in [n]\}$ and the equivalence relation for two independent sets $I, I' \in S_0$ by $I \sim_0 I'$ if and only if $I \cap U = I' \cap U$. That is, I and I' have the same assignments of vertices in U . Let $I_1, \dots, I_t$ be representatives for the $\sim_0$ -equivalence classes. We obtain

$$\sum_{I \in S_0} \lambda_\ell^{|V_L \cap I|} \lambda_r^{|V_R \cap I|} = \sum_{s=1}^t \sum_{I \in \llbracket I_s \rrbracket_0} \lambda_\ell^{|V_L \cap I|} \lambda_r^{|V_R \cap I|}. \quad (7)$$

Let $s \in [t]$ and $I \in \llbracket I_s \rrbracket_0$. Since $I \in S_0$, at least one of $u_i, \bar{u}_i$ and at least one of $v_i, \bar{v}_i$ are in I . We recall that for each $i \in [n]$ both (u_i, v_i) and $(\bar{u}_i, \bar{v}_i)$ are edges in G_φ . Therefore, either the pair $\{u_i, \bar{v}_i\} \subseteq I$ or the pair $\{\bar{u}_i, v_i\} \subseteq I$ and, consequently, for each $i \in [n]$ neither $w_i (= u_L^i)$ nor $z_i (= v_R^{n+i})$ can be in I . Furthermore, for each $j \in [m]$ there exists at least one vertex in $\Gamma_{G'_\varphi}(y_j) \cap I$ by the definition of S_0 . Hence, for each $j \in [m]$ the vertex $y_j = v_R^{2n+j}$ cannot be in I . We deduce that I contains exactly n vertices from $V_L(G'_\varphi)$ and exactly n vertices from $V_R(G'_\varphi)$.

Each graph B^j is a copy of the graph B and the vertices u_L^j for $j \leq n$ and v_R^j for $j > n$, respectively, are cut vertices in G_φ . There are n copies of B with u_L identified with a vertex in G'_φ and $n + m$ copies of B with v_R identified with a vertex in G'_φ . Clearly, for arbitrary graphs G_1 and G_2 with disjoint vertex sets it holds $Z_{\lambda_\ell, \lambda_r}(G_1 \cup G_2) = Z_{\lambda_\ell, \lambda_r}(G_1) Z_{\lambda_\ell, \lambda_r}(G_2)$. This yields for every $s \in [t]$

$$\sum_{I \in \llbracket I_s \rrbracket_0} \lambda_\ell^{|V_L \cap I|} \lambda_r^{|V_R \cap I|} = (\lambda_\ell \lambda_r)^n \left(\sum_{I \in \mathcal{I}(B-u_L)} \lambda_\ell^{|V_L(B-u_L) \cap I|} \lambda_r^{|V_R(B-u_L) \cap I|} \right)^n \left(\sum_{I \in \mathcal{I}(B-v_R)} \lambda_\ell^{|V_L(B-v_R) \cap I|} \lambda_r^{|V_R(B-v_R) \cap I|} \right)^{n+m}.$$

Since B , $B - u_L$, and $B - v_R$ are bipartite graphs, we obtain

$$\sum_{I \in \llbracket I_s \rrbracket_0} \lambda_\ell^{|V_L \cap I|} \lambda_r^{|V_R \cap I|} = (\lambda_\ell \lambda_r)^n (Z_{\lambda_\ell, \lambda_r}(B - u_L))^n (Z_{\lambda_\ell, \lambda_r}(B - v_R))^{n+m}.$$

We recall that B was chosen due to Lemma 3.8, whose Property 2 and Property 3 assure

$$K := \sum_{I \in \llbracket I_s \rrbracket_0} \lambda_\ell^{|V_L \cap I|} \lambda_r^{|V_R \cap I|} \not\equiv 0 \pmod{p}. \quad (8)$$

Combining Inequivalence (8) and (7) in conjunction with (6), we derive

$$Z_{\lambda_\ell, \lambda_r}(G_\varphi) \equiv tK \pmod{p}. \quad (9)$$

We will conclude the proof by constructing a bijection between the $\sim_0$ -equivalence classes and the satisfying assignments of φ . In this manner, we will obtain $t = |\text{sat}(\varphi)|$.

For every equivalence class $\llbracket I_s \rrbracket_0$ with $s \in [t]$, we denote the set of common vertices in $\llbracket I_s \rrbracket_0$ by $U_s = \bigcap_{I \in \llbracket I_s \rrbracket_0} I$. Due to the definition of $\sim_0$ for every $i \in [n]$ either the pair $\{u_i, \bar{v}_i\}$ or the pair $\{\bar{u}_i, v_i\}$ is shared by all elements of $\llbracket I_s \rrbracket_0$. Hence, U_s contains exactly n such pairs of vertices.

Given an equivalence class $\llbracket I_s \rrbracket_0$ utilizing U_s , we obtain an assignment a_s for φ by assigning for all $i \in [n]$

$$x_i \mapsto \begin{cases} \text{true}, & \text{if } \{u_i, \bar{v}_i\} \subseteq U_s; \\ \text{false}, & \text{if } \{\bar{u}_i, v_i\} \subseteq U_s. \end{cases}$$

We observe that each $\llbracket I_s \rrbracket_0$ yields a unique assignment a_s . To show that it is a satisfying assignment it suffices to show that each clause of φ is satisfied when we apply a_s .

Since $I_s \in S_0$, for each clause c_j of φ there exists at least one vertex $u \in \Gamma_{G'_\varphi}(y_j)$ with $u \in I_s$. Due to the construction of G_φ this vertex u is either u_i , if x_i appears non-negated in the clause c_j , or $\bar{u}_i$, if x_i appears negated in the clause c_j . Hence, a_s satisfies c_j at least once.

Vice versa, we now argue that every satisfying assignment can be obtained from an equivalence class $\llbracket I_s \rrbracket_0$ for some $s \in [t]$. Let a be a satisfying assignment for φ , this assignment gives rise to the set

$$U_a = \bigcup_{i \in [n]} \{u_i, \bar{v}_i \mid \text{if } x_i \text{ is set "true" by } a\} \cup \{\bar{u}_i, v_i \mid \text{if } x_i \text{ is set "false" by } a\}.$$

From the structure of G_φ , we deduce that U_a is an independent set. Furthermore, from Definition 3.13 we have that $U_a \in S_0$. Thus for s such that $\llbracket U_a \rrbracket_0 = \llbracket I_s \rrbracket_0$ it holds $a_s = a$.

We deduce that there are t satisfying assignments of φ and by (9)

$$Z_{\lambda_\ell, \lambda_r}(G_\varphi) \equiv K |\text{sat}(\varphi)| \pmod{p},$$

which establishes the theorem. $\square$

4 TRACTABLE GRAPHS

We identify the classes of graphs H for which $\#_p \text{HomsTo} H$ can be solved in polynomial time. When counting graph homomorphisms modulo a prime p , the automorphisms of order p of a target graph H help us identify groups of homomorphisms that cancel out. More specifically, assume that the target graph H has an automorphism ϱ of order p . For any homomorphism σ from the input graph G to H , $\varrho \circ \sigma$ is also a homomorphism from G to H . This shows that the sets consisting of the homomorphisms $\varrho^{(j)} \circ \sigma$, for $j \in [p]$, have cardinality a multiple of p and cancel out. This intuition is captured by the theorem of Faben and Jerrum [10, Theorem 3.4]. Before we formally state their theorem, we need the following definition.

Definition 4.1. Let H be a graph and ϱ an automorphism of H . H^ϱ is the subgraph of H induced by the vertices fixed by ϱ .

THEOREM 4.2 (FABEN AND JERRUM). *Let G, H be graphs, p a prime and ϱ an automorphism of H of order p . Then $|\text{Hom}(G \rightarrow H)| \equiv |\text{Hom}(G \rightarrow H^\varrho)| \pmod{p}$.*

We can repeat the above reduction of H recursively in the following way.

Definition 4.3. Let H, H' be graphs and p a prime. We write $H \Rightarrow_p H'$ if there is an automorphism ϱ of H of order p such that $H^\varrho = H'$. We will also write $H \Rightarrow_p^* H'$ if either $H \cong H'$ or, for some positive integer k , there are graphs $H_1, \dots, H_k$ such that $H \cong H_1$, $H_1 \Rightarrow_p \dots \Rightarrow_p H_k$, and $H_k \cong H'$.

Faben and Jerrum [10, Theorem 3.7] show for any choice of intermediate homomorphisms of order p , the reduction $H \Rightarrow_p^* H'$ will end up in a unique graph up to isomorphism.

THEOREM 4.4 (FABEN AND JERRUM). *Given a graph H and a prime p , there is (up to isomorphism) exactly one graph H^{*p} that has no automorphism of order p and $H \Rightarrow_p^* H^{*p}$.*

The latter suggest the following definition.

Fig. 4. An example of the order 3 reduced form H^{*3} of the graph H . Here we indicate two different ways of $H \Rightarrow_3 H^{*3}$. The automorphism ϱ has order 3. It is indicated with red colour and $H^\varrho = H^{*3}$. σ , τ and ν each are automorphisms of order 3. These are indicated with blue colour and $((H^\sigma)^\tau)^\nu = H^{*3}$.

Definition 4.5. Let H be a graph and p a prime. We call the unique (up to isomorphism) graph H^{*p} , with $H \Rightarrow_p^* H^{*p}$, the *order p reduced form* of H .

Figure 4 illustrates Theorem 4.4 with an example of an order 3 reduced form of a graph. Note that if H has no loops, then the repeated application of the “ $\Rightarrow_p$ ” operation does not introduce any loops.

To compute the number of homomorphisms from G to H modulo p , denoted by $\#_p \text{HomSTo}H$, it suffices to compute the number of homomorphisms from G to H^{*p} modulo p . We refer to the dichotomy theorem by Dyer and Greenhill [7, Theorem 1.1] to obtain the graphs for which $\#_p \text{HomSTo}H$ is computed in polynomial time.

THEOREM 4.6 (DYER AND GREENHIL). *Let H be a graph that can contain loops. If every component of H is a complete bipartite graph with no loops or a complete graph with all loops present, then $\#HOMsToH$ can be solved in polynomial time. Otherwise $\#HOMsToH$ is $\#P$ -complete.*

We notice that a polynomial time algorithm for $\#HOMSTOH$, gives a polynomial time algorithm for $\#_pHOMSTOH$ by simply applying the modulo p operation. In our setting, H contains no loops, so we have the following characterisation for the polynomial time computable instances of $\#_pHOMSTOH$.

COROLLARY 4.7. *Let H be a graph and p a prime. If every component of $H^{\ast p}$ is a complete bipartite graph, then $\#_p \text{HOMsTo } H$ is computable in polynomial time.*

5 HOMOMORPHISMS OF PARTIALLY LABELLED GRAPHS

We prove that counting the number of homomorphisms from a partially labelled graph J to a fixed graph H modulo p reduces to the problem of counting homomorphisms from a graph G to H modulo p . This generalises the results of Göbel, Goldberg, and Richerby [14]. Many of the definitions and key lemmas we use in this sections are generalisation of the ones [14, Section 3], so our presentation follows the one of [14] closely.

We study the following problem.

PROBLEM 5.1. *Name.* $\#_p\text{PARTLABHOMSTO}H$.

Parameter. Graph H and prime p .

Input. Partially H -labelled graph $J = (G, \tau)$.

Output. $|\text{Hom}(J \rightarrow H)| \pmod{p}$.

As Lovász has shown [22], two graphs H and H' are isomorphic if and only if for every graph G holds $|\text{Hom}(G \rightarrow H)| = |\text{Hom}(G \rightarrow H')|$. In [10, Lemma 4.5] Faben and Jerrum used a slightly different terminology and showed that this result holds for partially labelled graphs J modulo all primes p if the pinning function is restricted to map exactly one vertex of $G(J)$ to a vertex of H . In [14, Lemma 3.6] the following version of this result was shown.

LEMMA 5.2 (GÖBEL, GOLDBERG AND RICHERBY). *Let $(H, \bar{v})$ and $(H', \bar{v}')$ be graphs that both have no automorphism of order 2, each with r distinguished vertices. Then $(H, \bar{v}) \cong (H', \bar{v}')$ if and only if, for all (not necessarily connected) graphs $(G, \bar{u})$ with r distinguished vertices,*

$$|\text{Hom}((G, \bar{u}) \rightarrow (H, \bar{v}))| \equiv |\text{Hom}((G, \bar{u}) \rightarrow (H', \bar{v}'))| \pmod{2}.$$

In a sense, this version is more general than the result by Faben and Jerrum as the pinning function can map any number of vertices, but it is only stated for modulo 2. A discussion about the subtle differences of the two results appears in [14, Section 3.4]. For our purposes, we observe that the proof of Lemma 5.2 holds modulo all primes p .

LEMMA 5.3. *Let p be a prime and let $(H, \bar{v})$ and $(H', \bar{v}')$ be graphs having no automorphism of order p , each with r distinguished vertices. Then $(H, \bar{v}) \cong (H', \bar{v}')$ if and only if, for all (not necessarily connected) graphs $(G, \bar{u})$ with r distinguished vertices,*

$$|\text{Hom}((G, \bar{u}) \rightarrow (H, \bar{v}))| \equiv |\text{Hom}((G, \bar{u}) \rightarrow (H', \bar{v}'))| \pmod{p}.$$

Explanation. Let $\text{InjHom}((G, \bar{u}) \rightarrow (H, \bar{v}))$ denote the set of injective homomorphisms from $(G, \bar{u})$ to $(H, \bar{v})$. In the proof of [14, Lemma 3.6], the following equivalence is shown as Equation (2):

$$|\text{InjHom}((G, \bar{u}) \rightarrow (H, \bar{v}))| \equiv |\text{InjHom}((G, \bar{u}) \rightarrow (H', \bar{v}'))| \pmod{2}.$$

By reviewing the proof, we observe that no modular equivalences are used and in fact the following equation holds:

$$|\text{InjHom}((G, \bar{u}) \rightarrow (H, \bar{v}))| = |\text{InjHom}((G, \bar{u}) \rightarrow (H', \bar{v}'))|. \quad (10)$$

Now we show that if (10) holds for all graphs $(G, \bar{u})$ with r distinguished vertices, then $(H, \bar{v}) \cong (H', \bar{v}')$. We consider first $(G, \bar{u}) = (H, \bar{v})$. An injective homomorphism from a finite graph to itself is an automorphism and, since $(H, \bar{v})$ has no automorphism of order p , $\text{Aut}(H, \bar{v})$ has no element of order p , so $|\text{Aut}(H, \bar{v})| \not\equiv 0 \pmod{p}$ by Cauchy's group theorem (Theorem 2.2). By (10), the number of injective homomorphisms from $(H, \bar{v})$ to $(H', \bar{v}')$ is not equivalent to 0 $\pmod{p}$, which means that there is at least one such homomorphism. Similarly, considering $(G, \bar{u}) = (H', \bar{v}')$ yields the existence of an injective homomorphism from $(H', \bar{v}')$ to $(H, \bar{v})$. Due to the existence of both injective homomorphisms, we conclude that the two graphs are isomorphic. $\square$

A complete, self-contained proof of Lemma 5.3 can also be found [12]. We now introduce orbit vectors as [14] but generalised to an arbitrary prime p .

Definition 5.4. Let H be a graph with no automorphism of order p and $r \in \mathbb{Z}_{>0}$. An enumeration $\bar{v}_1, \dots, \bar{v}_\mu$ of elements of $(V(H))^r$ such that, for every $\bar{v} \in (V(H))^r$, there is exactly one $i \in [\mu]$ such that $(H, \bar{v}) \cong (H, \bar{v}_i)$ is referred to as an *enumeration of $(V(H))^r$ up to isomorphism*.

The number μ of tuples in the enumeration depends on the structure of H and not only on $|V(H)|$.

Definition 5.5. Let H be a graph with no automorphism of order p , $r \in \mathbb{Z}_{>0}$ and let $\bar{v}_1, \dots, \bar{v}_\mu$ be an enumeration of $(V(H))^r$ up to isomorphism. Further, let $(G, \bar{u})$ be a graph with r distinguished vertices. We define the *orbit vector* $\mathbf{v}_H(G, \bar{u}) \in (\mathbb{Z}_p)^\mu$ where, for each $i \in [\mu]$, the i th component of $\mathbf{v}_H(G, \bar{u})$ is given by

$$(\mathbf{v}_H(G, \bar{u}))_i \equiv |\text{Hom}((G, \bar{u}) \rightarrow (H, \bar{v}_i))| \pmod{p}.$$

We say that $(G, \bar{u})$ *implements* this vector.

Due to Lemma 5.3, for every graph H and for all $\bar{v} \in (V(H))^r$ and $i \in [\mu]$ such that $(H, \bar{v}) \cong (H, \bar{v}_i)$, we have that $(\mathbf{v}_H(G, \bar{u}))_i \equiv |\text{Hom}((G, \bar{u}) \rightarrow (H, \bar{v}))| \pmod{p}$.

For a group $\mathcal{G}$ acting on a set X , the *orbit* of an element $x \in X$ is defined to be the set $\text{Orb}_{\mathcal{G}}(x) = \{\pi(x) \mid \pi \in \mathcal{G}\}$. For a graph H , we will abuse notation and write $\text{Orb}_H(\cdot)$ instead of $\text{Orb}_{\text{Aut}(H)}(\cdot)$. Using this notation, for a graph H and a tuple of vertices $\bar{v} = (v_1, v_2, \dots, v_\mu)$ in $V(H)^\mu$, $\text{Orb}_H(\bar{v})$ is the set of all tuples $\varrho(\bar{v}) = (\varrho(v_1), \varrho(v_2), \dots, \varrho(v_\mu))$, where $\varrho \in \text{Aut}(H)$. Thus, for $r \in \mathbb{Z}_{>0}$ and an enumeration $\bar{v}_1, \dots, \bar{v}_\mu$ of $(V(H))^r$ up to isomorphism, $|\{\bar{v} \in (V(H))^r \mid (H, \bar{v}) \cong (H, \bar{v}_i)\}| = |\text{Orb}_H(\bar{v}_i)|$ for every $i \in [\mu]$.

Defining the vectors using the enumeration up to isomorphism hides the size of the orbit of a tuple $\bar{v}_i \in (V(H))^r$, as each orbit gets contracted to a single entry. This information is not needed when counting modulo 2, because we can prove that, for the graphs we are interested in, for every tuple $\bar{v}_i$ the cardinality $|\text{Orb}_H(\bar{v}_i)|$ is odd. In contrast, this information is needed when counting modulo an odd prime. Since H is fixed, we can recover this information at any point. As it is more convenient to prove the technical lemmas using the contracted vectors of Definition 5.5, we will employ this recovery at a later point.

We denote by $\oplus^p$ and $\otimes^p$ componentwise addition and multiplication modulo p , of vectors in $(\mathbb{Z}_p)^\mu$, respectively.

LEMMA 5.6. Let $(G_1, \bar{u})$, $(G_2, \bar{u})$ be graphs, where $\bar{u} = u_1 \dots u_r$ with $r \in \mathbb{Z}_{>0}$, such that $V(G_1) \cap V(G_2) = \{u_1, \dots, u_r\}$. Further, let H be a graph with no automorphism of order p with an enumeration of $(V(H))^r$ up to isomorphism. Then

$$\mathbf{v}_H(G_1 \cup G_2, \bar{u}) = \mathbf{v}_H(G_1, \bar{u}) \otimes^p \mathbf{v}_H(G_2, \bar{u}).$$

PROOF. A function $\sigma: V(G_1) \cup V(G_2) \rightarrow V(H)$ is a homomorphism from $(G_1 \cup G_2, \bar{u})$ to $(H, \bar{v})$ if and only if for each $i \in \{1, 2\}$ the restriction of σ to $V(G_i)$ is a homomorphism from $(G_i, \bar{u})$ to $(H, \bar{v})$. $\square$

Componentwise multiplication of $\mathbf{v}_H(G_1, \bar{u})$ and $\mathbf{v}_H(G_2, \bar{u})$ for two given graphs $(G_1, \bar{u})$ and $(G_2, \bar{u})$ can be expressed as an orbit vector of a single graph. This is not the case for componentwise addition $\mathbf{v}_H(G_1, \bar{u}_1) \oplus^p \mathbf{v}_H(G_2, \bar{u}_2)$. For our purposes it is sufficient that a set of graphs exists, whose vectors sum componentwise to a desired vector.

For graphs with distinguished vertices $(G_1, \bar{u}_1), \dots, (G_t, \bar{u}_t)$, we define

$$\mathbf{v}_H((G_1, \bar{u}_1) + \dots + (G_t, \bar{u}_t)) = \mathbf{v}_H(G_1, \bar{u}_1) \oplus^p \dots \oplus^p \mathbf{v}_H(G_t, \bar{u}_t)$$

and say that a vector $\mathbf{v} \in (\mathbb{Z}_p)^\mu$ is *H-implementable* if it can be expressed as such a sum.

The modulo 2 version of the following lemma appears [10, Lemma 4.16] and is used for all pinning techniques in the literature so far. We reprove the lemma for the vectors in $(\mathbb{Z}_p)^\mu$ when p is an arbitrary prime.

LEMMA 5.7. *Let $\mu \in \mathbb{Z}_{>0}$ and $S \subseteq (\mathbb{Z}_p)^\mu$ be closed under $\oplus^p$ and $\otimes^p$. If $1^\mu \in S$ and, for every distinct $i, j \in [\mu]$, there is a tuple $s = s_1 \dots s_\mu \in S$ with $s_i \neq s_j$, then $S = (\mathbb{Z}_p)^\mu$.*

PROOF. It suffices to show that all of the basis vectors of the standard basis¹ in $(\mathbb{Z}_p)^\mu$ are in S . Since S is closed under $\oplus^p$ and $\otimes^p$ it follows that all of $(\mathbb{Z}_p)^\mu$ is in S .

We show that all the basis vectors are in S by induction on μ . If $\mu = 1$, then the lemma clearly holds as the all-ones vector is the only vector in the standard basis. Assume that $\mu > 1$ and that the induction hypothesis holds for $\mu - 1$. Then we can construct vectors that agree with the standard basis in the first $\mu - 1$ places without being able to control what happens in the μ th place. By the latter and $1^\mu \in S$, we obtain the following vectors:

$$\begin{array}{rcll} \mathbf{v}_0 & = & 1 & 1 & 1 & \dots & 1 & 1 \\ \mathbf{v}_1 & = & 1 & 0 & 0 & \dots & 0 & x_1 \\ \mathbf{v}_2 & = & 0 & 1 & 0 & \dots & 0 & x_2 \\ \vdots & & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{v}_{\mu-1} & = & 0 & 0 & 0 & \dots & 1 & x_{\mu-1} \end{array},$$

where the x_i can take any value in $\mathbb{Z}_p$.

Let r be an integer and let $\mathbf{v} \in (\mathbb{Z}_p)^\mu$. We use the notation $\mathbf{v}^r = \mathbf{v} \otimes^p \dots \otimes^p \mathbf{v}$ for the r -fold componentwise product and let $r\mathbf{v} = \mathbf{v} \oplus^p \dots \oplus^p \mathbf{v}$ denote the r -fold componentwise sum of $\mathbf{v}$. Consider the values of each x_i . If $x_i \neq 0$, then by Theorem 2.1 we have $x_i^{p-1} \equiv 1 \pmod{p}$. Hence $\mathbf{v}_i^{p-1} = 00 \dots 010 \dots 01$. So from now on, we can assume that for each $i \in [\mu]$, $x_i \in \{0, 1\}$. We have the following three cases.

Case 1. For all $i \in [\mu - 1]$, $x_i = 0$. Then the vector $\mathbf{v} = \mathbf{v}_0 \oplus^p \bigoplus_{i \in [\mu]}^p (p-1)\mathbf{v}_i = 0 \dots 01$ is the remaining vector that completes the standard basis.

Case 2. There are at least two j, ℓ such that $x_j, x_\ell = 1$. Then $\mathbf{v} = \mathbf{v}_j \otimes^p \mathbf{v}_\ell = 0 \dots 01$. To obtain the remaining vectors of the standard basis, for all $i \in [\mu - 1]$ with $x_i \neq 0$, we replace $\mathbf{v}_i$ with $\mathbf{v}_i \oplus^p (p-1)\mathbf{v}$.

Case 3. There is exactly one $i \in [\mu - 1]$ with $x_i = 1$. From the statement of the lemma there is a vector $\mathbf{u} \in S$ with $(\mathbf{u})_i = a$ and $(\mathbf{u})_\mu = b$, where $a \neq b$. Let $\mathbf{u}_i = \mathbf{u} \otimes^p \mathbf{v}_i = 0 \dots 0a0 \dots 0b$ and let $\mathbf{v}_a = (p-a)\mathbf{v}_i = 0 \dots 0(p-a)0 \dots 0(p-a)$. Then $\mathbf{u}_i \oplus^p \mathbf{v}_a = 0 \dots 0(p-a+b)$. Since $a \neq b$, $(p-a+b)$ is not a multiple of p , hence by Theorem 2.1 we have $(p-a+b)^{p-1} \equiv 1 \pmod{p}$. Therefore, we derive that $\mathbf{v} = (\mathbf{u}_i \oplus^p \mathbf{v}_a)^{p-1} = 0 \dots 01$ and $\mathbf{v}'_i = (p-1)\mathbf{v} \otimes^p \mathbf{v}_i = 0 \dots 010 \dots 0$ complete the standard basis. $\square$

COROLLARY 5.8. *Let H be a graph with no automorphism of order p with an enumeration $\bar{v}_1, \dots, \bar{v}_\mu$ of $(V(H))^r$ up to isomorphism. Then every $\mathbf{v} \in (\mathbb{Z}_p)^\mu$ is H -implementable.*

PROOF. Let S be the set of H -implementable vectors. S is clearly closed under $\oplus^p$ and by Lemma 5.6 is also closed under $\otimes^p$. Let G be the graph on vertices $\{u_1, \dots, u_r\}$ without edges. 1^μ is implemented by $(G, u_1, \dots, u_r)$, which has exactly one homomorphism to every $(H, \bar{v}_i)$. Finally by Lemma 5.3, for every pair $i, j \in [\mu]$ with $(H, \bar{v}_i)$ and $(H, \bar{v}_j)$ not isomorphic there is a graph $(G, \bar{u})$, such that

$$|\text{Hom}((G, \bar{u}) \rightarrow (H, \bar{v}_i))| \not\equiv |\text{Hom}((G, \bar{u}) \rightarrow (H, \bar{v}_j))| \pmod{p}.$$

¹The standard basis is the set $\{100 \dots 00, 010 \dots 00, \dots, 000 \dots 01\}$.

$(G, \bar{u})$ implements a vector $\mathbf{v}$ whose i th and j th components are different and the corollary follows from Lemma 5.7. $\square$

At this point, we have shown that all orbit vectors in $(\mathbb{Z}_p)^\mu$ are H -implementable. Now, we define the tuple vectors that have an entry for each r -tuple. From these tuple vectors, we can infer the sizes of the orbits $\text{Orb}_H(\bar{v})$ for all $v \in V(H)^r$. This information is vital for the proof of our main theorem.

Definition 5.9. Let H be a graph with no automorphism of order p , let $r \in \mathbb{Z}_{>0}$ and let $\bar{w}_1, \dots, \bar{w}_\nu$ be an enumeration of $(V(H))^r$, i.e., $\nu = |V(H)|^r$. Let $(G, \bar{u})$ be a graph with r distinguished vertices. We define the *tuple vector* $\mathbf{w}_H(G, \bar{u}) \in (\mathbb{Z}_p)^\nu$ where, for each $j \in [\nu]$, the j th component of $\mathbf{w}_H(G, \bar{u})$ is given by

$$(\mathbf{w}_H(G, \bar{u}))_j \equiv |\text{Hom}((G, \bar{u}) \rightarrow (H, \bar{w}_j))| \pmod{p}.$$

We say that $(G, \bar{u})$ *implements* this vector.

Definition 5.10. Let H be a graph with no automorphism of order p , $r \in \mathbb{Z}_{>0}$ and let $\bar{w}_1, \dots, \bar{w}_\nu$ be an enumeration of $(V(H))^r$, i.e., $\nu = |V(H)|^r$. Denote by $F(H, r) \subseteq (\mathbb{Z}_p)^\nu$ the set of vectors $\mathbf{w}$, such that, for all $i, j \in [\nu]$ with $(H, \bar{w}_i) \cong (H, \bar{w}_j)$, we have $(\mathbf{w})_i = (\mathbf{w})_j$.

The following lemma shows which tuple vectors are H -implementable. The proof uses the H -implementable orbit vectors and retracts the information that gets lost by using the enumeration up to isomorphism of the r -tuples.

LEMMA 5.11. *Let H be a graph with no automorphism of order p , $r \in \mathbb{Z}_{>0}$ and $\bar{w}_1, \dots, \bar{w}_\nu$ an enumeration of $(V(H))^r$. Then every $\mathbf{w} \in F(H, r)$ is H -implementable.*

PROOF. Let $\bar{v}_1, \dots, \bar{v}_\mu$ be an enumeration up to isomorphism of $(V(H))^r$. We denote by $f : [\mu] \rightarrow [\nu]$ the associated function with $\bar{v}_i = \bar{w}_{f(i)}$ for all $i \in [\mu]$, i.e., f tells us which coordinates of the tuple vector are representatives for the equivalence classes providing the coordinates of the orbit vector. Given $\mathbf{w} \in F(H, r)$, we compute the corresponding vector $\mathbf{v} \in (\mathbb{Z}_p)^\mu$ by letting $(\mathbf{v})_i = (\mathbf{w})_{f(i)}$ for all $i \in [\mu]$. The vector $\mathbf{v}$ is H -implementable by Corollary 5.8. Now, if $(G, \bar{u})$ is a graph with r distinguished vertices such that $(\mathbf{v})_i \equiv |\text{Hom}((G, \bar{u}) \rightarrow (H, \bar{v}_i))| \pmod{p}$ for all $i \in [\mu]$, then we also have $(\mathbf{w})_j \equiv |\text{Hom}((G, \bar{u}) \rightarrow (H, \bar{w}_j))| \pmod{p}$ for all $j \in [\nu]$. $\square$

Before we prove the main theorem of this section, we need the following lemma.

LEMMA 5.12. *Let H be a graph with no automorphism of order p , $r \in \mathbb{Z}_{>0}$ and $\bar{w}_1, \dots, \bar{w}_\nu$ an enumeration of $(V(H))^r$. Then for every graph $(G, \bar{u})$ with r distinguished vertices*

$$|\text{Hom}(G \rightarrow H)| \equiv \sum_{j \in [\nu]} (\mathbf{w}_H(G, \bar{u}))_j \pmod{p}.$$

PROOF. We have

$$\begin{aligned} \sum_{j \in [\nu]} (\mathbf{w}_H(G, \bar{u}))_j &\equiv \sum_{j \in [\nu]} |\text{Hom}((G, \bar{u}) \rightarrow (H, \bar{w}_j))| \pmod{p} \\ &= |\text{Hom}(G \rightarrow H)| \pmod{p}. \end{aligned}$$

The equivalence holds by the definition of $\mathbf{w}_H(G, \bar{u})$. The equality holds, because every homomorphism from G to H must map $\bar{u}$ to some r -tuple $\bar{w}$. Since $[\nu]$ indexes all r -tuples, we obtain all homomorphisms from G to H . $\square$

THEOREM 1.8. *Let p be a prime and let H be a graph with no automorphism of order p . Then $\#_p \text{PARTLABHOMSTO}H$ reduces to $\#_p \text{HOMSTO}H$ via a polynomial time Turing reduction.*

PROOF. Let $J = (G, \tau)$ be an instance of $\#_p\text{PARTLABHOMSTO}H$. Let $\bar{u} = u_1 \dots u_r$ be an enumeration of $\text{dom}(\tau)$ and let $\bar{w} = w_1 \dots w_r = \tau(u_1) \dots \tau(u_r)$. We translate the notion of partially H -labelled graphs to the equivalent notion of graphs with distinguished vertices and aim to compute $|\text{Hom}((G, \bar{u}) \rightarrow (H, \bar{w}))|$ modulo p . Let $\bar{w}_1, \dots, \bar{w}_v$ be an enumeration of $(V(H))^r$ and let $\mathbf{w} \in \{0, 1\}^v$ be the vector with $(\mathbf{w})_j = 1$ if $(H, \bar{w}_j) \cong (H, \bar{w})$, and 0 for all other $j \in [v]$; $\mathbf{w}$ has exactly $|\text{Orb}_H(\bar{w})|$ 1-entries. Since $\mathbf{w} \in F(H, r)$ by Lemma 5.11 $\mathbf{w}$ is H -implemented by some sequence $(\Theta_1, \bar{u}_1), \dots, (\Theta_t, \bar{u}_t)$ of graphs with r -tuples of distinguished vertices.

For each $s \in [t]$ let $(G_s, \bar{u})$ be the graph that results from taking the disjoint union of a copy of G and Θ_s and identifying the i th element of $\bar{u}$ with the i th element of $\bar{u}_s$ for each $i \in [r]$. Observe that Lemma 5.6 applies with $\mathbf{v}_H$ replaced by $\mathbf{w}_H$. We have

$$\mathbf{w}_H(G_s, \bar{u}) = \mathbf{w}_H(G, \bar{u}) \otimes^p \mathbf{w}_H(\Theta_s, \bar{u}_s).$$

With this, we obtain

$$\begin{aligned} \mathbf{w}_H(G, \bar{u}) \otimes^p \mathbf{w} &= \mathbf{w}_H(G, \bar{u}) \otimes^p \mathbf{w}_H((\Theta_1, \bar{u}_1) + \dots + (\Theta_t, \bar{u}_t)) \\ &= \mathbf{w}_H(G, \bar{u}) \otimes^p (\mathbf{w}_H(\Theta_1, \bar{u}_1) \oplus^p \dots \oplus^p \mathbf{w}_H(\Theta_t, \bar{u}_t)) \\ &= \bigoplus_{s \in [t]}^p (\mathbf{w}_H(G, \bar{u}) \otimes^p \mathbf{w}_H(\Theta_s, \bar{u}_s)) \\ &= \bigoplus_{s \in [t]}^p \mathbf{w}_H(G_s, \bar{u}). \end{aligned}$$

Since $\mathbf{w}$ contains a 1-entry for each $\bar{w}_k \in \text{Orb}_H(\bar{w})$ and a 0-entry everywhere else we have by summing the components of the vector $\mathbf{w}_H(G, \bar{u}) \otimes^p \mathbf{w}$

$$\sum_{j \in [v]} (\mathbf{w}_H(G, \bar{u}) \otimes^p \mathbf{w})_j \equiv |\text{Orb}_H(\bar{w})| \cdot |\text{Hom}((G, \bar{u}) \rightarrow (H, \bar{w}))| \pmod{p}. \quad (11)$$

Summing the components of the vector $\bigoplus_{s \in [t]}^p \mathbf{w}_H(G_s, \bar{u})$, we derive

$$\sum_{j \in [v]} \left(\bigoplus_{s \in [t]}^p \mathbf{w}_H(G_s, \bar{u}) \right)_j = \sum_{s \in [t]} \sum_{j \in [v]} (\mathbf{w}_H(G_s, \bar{u}))_j. \quad (12)$$

By applying Lemma 5.12, we have that the values of (12) are modulo p congruent to the sum $\sum_{s \in [t]} |\text{Hom}(G_s \rightarrow H)|$. Thus, by the equivalence of (11) and (12) we deduce

$$|\text{Orb}_H(\bar{w})| \cdot |\text{Hom}((G, \bar{u}) \rightarrow (H, \bar{w}))| \equiv \sum_{s \in [t]} |\text{Hom}(G_s \rightarrow H)| \pmod{p}.$$

The right side can be computed by making t calls to an oracle for $\#_p\text{HOMSTO}H$. Since H is fixed and r is finite, we can trivially compute $|\text{Orb}_H(\bar{w})|$, which allows us to recover $|\text{Hom}((G, \bar{u}) \rightarrow (H, \bar{w}))|$. $\square$

6 HARD CASES FOR TREES

We are going to identify the classes of trees H , for which $\#_p\text{HOMSTO}H$ is $\#_p$ P-hard. Due to Section 4, we focus on graphs that have no automorphism of order p . In particular, Corollary 4.7 yields that $\#_p\text{HOMSTO}H$ is tractable when H is a star. A tree that is not a star contains a path of length at least 3, and this path is the structure that gives us hardness. We formally define.

Definition 6.1. Let H be a graph, p be a prime and $a, b \in \mathbb{Z}_p \setminus \{1\}$. Assume H contains a path $x_0 \dots x_k$ for $k > 0$, such that the following hold

- (1) $x_0 \dots x_k$ is the unique path between x_0 and x_k in H .
- (2) $\deg_H(x_0) \equiv a \pmod{p}$ and $\deg_H(x_k) \equiv b \pmod{p}$.
- (3) For all $0 < i < k$, $\deg_H(x_i) \equiv 1 \pmod{p}$.

Then, we will call $x_0 \dots x_k$ an (a, b, p) -path in H and denote it by Q_H .

We proceed by showing that every non-star tree H without automorphisms of order p contains such a path.

LEMMA 6.2. Let H be a tree that has no automorphism of order p . Then, either H is a star or there are $a, b \in \mathbb{Z}_p \setminus \{1\}$ such that H contains an (a, b, p) -path.

PROOF. We assume that H is not a star and let $P = x_{-1}x_0 \dots x_\ell$ be a maximal path in H of length $\ell + 1$. We are going to prove that P contains an (a, b, p) -path.

Since H is not a star, P contains at least four vertices yielding $\ell > 1$. To prove that any vertex in $\Gamma_H(x_0) - x_1$ must be a leaf, we assume the contrary. Let $v \in \Gamma_H(x_0) - x_1$ be not a leaf and $v' \neq x_0$ be a neighbour of v . Since H is a tree, v' is distinct from all vertices in P . Then, $v'vx_0 \dots x_\ell$ is a path of length $\ell + 2$ contradicting the maximality of P . The very same argument yields that any vertex in $\Gamma_H(x_{\ell-1}) - x_{\ell-2}$ must be a leaf as well.

We assume toward a contradiction that $|\Gamma_H(x_0)| > p$. Let $Y = \{y_1, \dots, y_p\} \subseteq \Gamma_H(x_0) - x_1$ be a set of neighbours of x_0 , which are not equal to x_1 . Let ϱ be a mapping from H to itself defined as follows: for every vertex $y_i \in Y$, let $\varrho(y_i) = y_{i+1}$, where the indices are taken modulo p ; for any other vertex $v \in V(H) \setminus Y$, let $\varrho(v) = v$. As we have observed above, for all $i \in [p]$, y_i is a leaf only adjacent to x_0 . Therefore, ϱ is an automorphism of H of order p , which is a contradiction.

Hence, x_0 has at least two and at most p neighbours, which yields $\deg_H(x_0) \not\equiv 1 \pmod{p}$. Similarly, we obtain $\deg_H(x_{\ell-1}) \not\equiv 1 \pmod{p}$. Consequently, there exists the minimum

$$k = \min\{k' \in [\ell - 1] \mid \deg(x_{k'}) \not\equiv 1 \pmod{p}\},$$

which yields the subpath $P' = x_0 \dots x_k$ of P . Since H contains no cycles, P' is the unique path in H connecting x_0 and x_k . Finally, due to the choice of k we deduce $\deg_H(x_i) \equiv 1 \pmod{p}$ for all internal vertices x_i of P' with $i \in [k - 1]$. We conclude that P' is an (a, b, p) -path in H . $\square$

The following lemma helps us translate walks to homomorphisms and vice versa, which will be helpful in our arguments later.

LEMMA 6.3. Let H be a graph and let $x, y \in V(H)$. If P is the path $z_0z_1 \dots z_k$, then the number of k -walks in H from x to y is equal to $|\text{Hom}((P, z_0, z_k) \rightarrow (H, x, y))|$.

PROOF. Let $W(x, y, k)$ denote the number of k -walks in H between the vertices x and y . We prove the lemma by induction on k .

In the base case with $k = 1$ the path P consists only of the edge (z_0, z_1) . If x is adjacent to y in H , then there is only one homomorphism $\sigma : (P, z_0, z_k) \rightarrow (H, x, y)$ implying $W(x, y, 1) = 1 = |\text{Hom}((P, z_0, z_1) \rightarrow (H, x, y))|$. Otherwise, x and y are not adjacent. Hence, there cannot exist a homomorphism $\sigma : (P, z_0, z_1) \rightarrow (H, x, y)$ implying $W(x, y, 1) = 0 = |\text{Hom}((P, z_0, z_k) \rightarrow (H, x, y))|$.

Regarding the induction step, we assume $W(x, y, i) = |\text{Hom}((P, z_0, z_i) \rightarrow (H, x, y))|$ holds for all paths P of length $i < k$, and we will to show that $W(x, y, k) = |\text{Hom}((P, z_0, z_k) \rightarrow (H, x, y))|$. Let W be a k -walk in H from x to y . The first edge in W must be (x, u) for some neighbour u of x . Deleting the edge (x, u) from W yields a walk of length $k - 1$ from u to y . However, if for a neighbour

Fig. 5. Constructive route for J given G and the $(4, 2, 5)$ -path in H for $p = 5$.

u' of x there exists no k -walk from x to y with first edge (x, u') , then there is no $(k - 1)$ -walk from u' to y . This yields

$$W(x, y, k) = \sum_{u \in \Gamma_H(x)} W(u, y, k - 1). \quad (13)$$

Let $P' = z_1 \dots z_k$ be the path obtained from P by deleting the edge (z_0, z_1) . Since z_1 is adjacent to z_0 and every homomorphism $\sigma : (P, z_0, z_k) \rightarrow (H, x, y)$ maps z_0 to x , z_1 must be mapped to a neighbour of x . Hence, for every neighbour u of x a homomorphism $\sigma' : (P', z_1, z_k) \rightarrow (H, u, y)$ yields a homomorphism from (P, z_0, z_k) to (H, x, y) and vice versa. We deduce

$$|\text{Hom}((P, z_0, z_k) \rightarrow (H, x, y))| = \sum_{u \in \Gamma_H(x)} |\text{Hom}((P', z_1, z_k) \rightarrow (H, u, y))|. \quad (14)$$

Finally, by combining (14) with the induction hypothesis and (13) we obtain the desired

$$|\text{Hom}((P, z_0, z_k) \rightarrow (H, x, y))| = \sum_{u \in \Gamma_H(x)} W(u, y, k - 1) = W(x, y, k). \quad \square$$

COROLLARY 6.4. *Let G, H be graphs and let $u, v \in V(G)$. Then, for every homomorphism $\sigma : G \rightarrow H$ holds $d_H(\sigma(u), \sigma(v)) \leq d_G(u, v)$.*

PROOF. We assume toward a contradiction that there exists a homomorphism σ from G to H with $d_G(u, v) < d_H(\sigma(u), \sigma(v))$ and let $k = d_G(u, v)$. Since the distance between $\sigma(u)$ and $\sigma(v)$ in H is larger than k , there exists no k -walk in H between $\sigma(u)$ and $\sigma(v)$. Therefore, by Lemma 6.3 σ cannot exist. $\square$

To show that $\#_p \text{HOMSTO}H$ is $\#_p P$ -hard, we are going to establish a reduction from $\#_p \text{BIS}_{\lambda_\ell, \lambda_r}$ to $\#_p \text{PARTLABHOMSTO}H$. That is, given a graph G input for $\#_p \text{BIS}_{\lambda_\ell, \lambda_r}$, we construct a partially labelled graph J , input for $\#_p \text{PARTLABHOMSTO}H$, such that $Z_{\lambda_\ell, \lambda_r}(G) \equiv |\text{Hom}(J \rightarrow H)| \pmod{p}$. The construction of J as stated uses any path in H . When we define the actual reduction though, we will require that this path is an (a, b, p) -path in H .

Let p be a prime, $G = (V_L, V_R, E)$ be the bipartite input graph of $\#_p \text{BIS}_{\lambda_\ell, \lambda_r}$ and H be a tree, that is the target graph in $\#_p \text{PARTLABHOMSTO}H$. Assume H contains a path $Q = x_0 \dots x_k$. Then, J is constructed starting with G by adding two vertices $\hat{u}$ and $\hat{v}$ and connecting them to every vertex in V_L and V_R , respectively. Subsequently, every edge $e \in E$ is substituted with a copy P_k^e of the k -path by identifying the endpoints of the edge with the endpoints of the path. Finally, the pinning function of J maps $\hat{u}$ to x_0 as well as $\hat{v}$ to x_k . See Figure 5 for an example. Formally, we have the following definition.

Definition 6.5. Let p be a prime and H be a graph containing the path $Q = x_0 \dots x_k$. Given a bipartite graph $G = (V_L, V_R, E)$, J is the partially labelled graph with vertex set

$$V(G(J)) = \{ \hat{u}, \hat{v} \} \cup V_L \cup V_R \cup \{ z_i^e \mid i \in [k-1], e \in E \}$$

and edge set

$$\begin{aligned} E(G(J)) = & \{ (\hat{u}, u) \mid u \in V_L \} \cup \{ (z_j^e, z_{j+1}^e) \mid e \in E, j \in [k-2] \} \\ & \cup \{ (u, z_1^e), (z_{k-1}^e, v) \mid e = (u, v) \in E \} \cup \{ (v, \hat{v}) \mid v \in V_R \}. \end{aligned}$$

Finally, let $\tau(J) = \{ \hat{u} \mapsto x_0, \hat{v} \mapsto x_k \}$ be the partial labelling from $G(J)$ to H .

The following lemma requires the existence of an (a, b, p) -path in H and identifies the properties of J , which will help us to establish the reduction.

LEMMA 6.6. *Let p be a prime, let $G = (V_L, V_R, E)$ a bipartite graph and let H be a graph. Assume there are $a, b \in \mathbb{Z}_p \setminus \{1\}$ such that H contains an (a, b, p) -path $Q_H = x_0 \dots x_k$. We denote the diminished neighbourhoods of x_0 and x_k by $W_L = \Gamma_H(x_0) - x_1$ and $W_R = \Gamma_H(x_k) - x_{k-1}$, respectively. Additionally, let J be the partially labelled graph according to Definition 6.5. Then, for every homomorphism σ from J to H the following hold.*

- (1) *Let $u \in V_L$ and $v \in V_R$, then $\sigma(u) \in \Gamma_H(x_0)$ and $\sigma(v) \in \Gamma_H(x_k)$.*
- (2) *Let $\mathfrak{D}_\sigma = \{u \in V_L \mid \sigma(u) = x_1\} \cup \{v \in V_R \mid \sigma(v) = x_{k-1}\}$ and $\mathfrak{I}_\sigma = (V_L \cup V_R) \setminus \mathfrak{D}_\sigma$. Given another homomorphism σ' from J to H , the relation $\sigma \sim_{\mathfrak{I}} \sigma'$ if $\mathfrak{I}_\sigma = \mathfrak{I}_{\sigma'}$ is an equivalence relation with equivalence class denoted $\llbracket \sigma \rrbracket_{\mathfrak{I}}$.*
- (3) *Let $\sigma_1, \dots, \sigma_\mu$ be representatives from each $\sim_{\mathfrak{I}}$ -equivalence class. Then, the set $I(G)$ of independent sets of G is exactly the set $\{ \mathfrak{I}_{\sigma_i} \mid i \in [\mu] \}$.*
- (4) *For the diminished neighbourhoods holds $|\llbracket \sigma \rrbracket_{\mathfrak{I}}| \equiv |W_L|^{| \mathfrak{I}_\sigma \cap V_L |} |W_R|^{| \mathfrak{I}_\sigma \cap V_R |} \pmod{p}$.*

PROOF. We will prove each statement in order.

- (1) We observe that $\tau(J)(\hat{u}) = x_0$ and $\hat{u}$ is adjacent to every vertex in V_L . Therefore, σ has to map each vertex $u \in V_L$ to a vertex in the neighbourhood of x_0 . The analogous argument shows the second result regarding $\hat{v}$ and the neighbourhood of x_k .
- (2) The statement follows from the observation that each class $\llbracket \sigma \rrbracket_{\mathfrak{I}}$ is uniquely determined by the set $\mathfrak{I}_\sigma$.
- (3) We commence the proof with establishing that mapping σ to $\mathfrak{I}_\sigma$ defines a surjection from $\text{Hom}(J \rightarrow H)$ to $I(G)$. Then, we obtain a bijection from $\{ \llbracket \sigma_i \rrbracket_{\mathfrak{I}} \mid i \in [\mu] \}$ to $I(G)$, as with $\sim_{\mathfrak{I}}$ we identify exactly the σ and σ' , for which $\mathfrak{I}_\sigma = \mathfrak{I}_{\sigma'}$.

We first argue that, for every $\sigma \in \text{Hom}(J \rightarrow H)$, $\mathfrak{I}_\sigma$ is an independent set in G . Assume toward a contradiction that there exists $\sigma \in \text{Hom}(J \rightarrow H)$ and a pair of vertices $u, v \in \mathfrak{I}_\sigma$ with $(u, v) \in E$. Without loss of generality let $u \in V_L$ and $v \in V_R$. Due to Property 1 and $u, v \in \mathfrak{I}_\sigma$, we obtain that $\sigma(u) \in W_L$ and $\sigma(v) \in W_R$. The path $\sigma(u)x_0 \dots x_k \sigma(v)$ is the unique path of length $k+2$ connecting $\sigma(u)$ and $\sigma(v)$ in H , because Q_H is by the definition of an (a, b, p) -path the unique path connecting x_0 and x_k . Therefore, $\sigma(u)$ and $\sigma(v)$ have distance $k+2$ in H . However, by the construction of J we have $d_{G(J)}(u, v) = k$, which due to Corollary 6.4 contradicts the existence of σ .

Regarding surjectivity, let $I \in I(G)$. We are going to define a mapping σ_I that is a homomorphism from J to H with $\mathfrak{I}_{\sigma_I} = I$. To do so, let $x_{-1} \in W_L$ and $x_{k+1} \in W_R$. This is possible as Q_H is a (a, b, p) -path, and thus we have $W_L \neq \emptyset$ and $W_R \neq \emptyset$. Now, let σ_I be defined as follows.

First, σ_I maps $\hat{u}$ to x_0 and $\hat{v}$ to x_k , respecting the pinning $\tau(J)$. For every $u \in V_L \cap I$ and every $e \in E$ incident with u , σ_I maps the vertices $z_0^e, \dots, z_k^e$ to $x_{-1}, \dots, x_{k-1}$, respectively. For every $v \in V_R \cap I$ and every edge $e \in E$ incident with v , σ_I maps the vertices $z_0^e, \dots, z_k^e$ to $x_1, \dots, x_{k+1}$,

respectively. Finally for each edge $e \in E$ with neither of its endpoints in I , σ_I maps $z_0^e, \dots, z_k^e$ to $x_1, \dots, x_{k-1}, x_k, x_{k-1}$, respectively. From the construction of J it follows that $\sigma_I \in \text{Hom}(J \rightarrow H)$ and $\mathfrak{I}_{\sigma_I} = I$.

(4) Let $\sigma' : J \rightarrow H$ be a homomorphism in $\llbracket \sigma \rrbracket_{\mathfrak{I}}$. We commence with proving that, for every edge $e \in E$,

$$|\text{Hom}((P_k^e, z_0^e, z_k^e) \rightarrow (H, \sigma'(z_0^e), \sigma'(z_k^e)))| \equiv 1 \pmod{p}. \quad (15)$$

Let $r = |\text{Hom}((P_k^e, z_0^e, z_k^e) \rightarrow (H, \sigma'(z_0^e), \sigma'(z_k^e)))|$. Due to Lemma 6.3, r is equal to the number of k -walks in H from $\sigma'(z_0^e)$ to $\sigma'(z_k^e)$. We consider the following four cases for $\sigma'(z_0^e)$ and $\sigma'(z_k^e)$.

- (a) $\sigma'(z_0^e) \in W_L$ and $\sigma'(z_k^e) \in W_R$. This yields a contradiction as argued in the proof of Property 3.
- (b) $\sigma'(z_0^e) = x_1$ and $\sigma'(z_k^e) = x \in W_R$. Since Q_H is an (a, b, p) -path, $x_1 \dots x_k x$ is the only k -walk in H between x_1 and x . Hence, Lemma 6.3 yields $r = 1$.
- (c) $\sigma'(z_0^e) = x \in W_L$ and $\sigma'(z_k^e) = x_{k-1}$. Similarly to (b), this also yields $r = 1$.
- (d) $\sigma'(z_0^e) = x_1$ and $\sigma'(z_k^e) = x_{k-1}$. Consider the number of k -walks in H between x_1 and x_{k-1} denoted $W(x_1, x_{k-1}, k)$ and recall that $r = W(x_1, x_{k-1}, k)$. We denote by $Q' = x_1 \dots x_{k-1}$ the subpath of Q_H connecting x_1 and x_{k-1} , by which we derive $d_H(x_1, x_{k-1}) = k - 2$, because Q_H is an (a, b, p) -path and thus Q' is the unique path in H between x_1 and x_{k-1} . Furthermore, every k -walk in H between x_1 and x_{k-1} can be constructed from Q' by adding a walk of length 2 to any vertex x_i in Q' . Therefore, every vertex x_i yields one k -walk for every vertex in its neighbourhood. We note that by this construction the walk Q cannot contain a cycle as otherwise the uniqueness of Q_H would be violated. Thus, we only double-counted the walks entirely contained in Q' . That is, for every vertex x_i with $2 \leq i \leq k - 1$ in Q' the walk revisiting x_{i-1} after reaching x_i . Removing every such walk once from the calculation yields

$$W(x_1, x_{k-1}, k) = \left(\sum_{i=1}^{k-1} \deg_H(x_i) \right) - (k - 2).$$

Since Q_H is an (a, b, p) -path, we obtain, for all $2 \leq i \leq k - 1$, that $\deg_H(x_i) \equiv 1 \pmod{p}$ yielding $r \equiv 1 \pmod{p}$.

To show Property 4, we note that the set $\mathfrak{I}_{\sigma}$ uniquely determines $\llbracket \sigma \rrbracket_{\mathfrak{I}}$. Therefore, for any homomorphism $\sigma' \in \llbracket \sigma \rrbracket_{\mathfrak{I}}$ the labelling of vertices in $\mathfrak{I}_{\sigma}$ as well as $\hat{u}$ and $\hat{v}$ is fixed. Concerning the vertices in $\mathfrak{I}_{\sigma}$, due to Property 1, σ' maps $\mathfrak{I}_{\sigma} \cap V_L$ to W_L and $\mathfrak{I}_{\sigma} \cap V_R$ to W_R . Due to Definition 6.5 of $G(J)$ every vertex z_0^e and z_k^e is identified with a vertex in V_L and V_R , respectively. Finally, due to (15) once we have fixed a partial labelling σ' of every vertex z_0^e and z_k^e the number of homomorphisms respecting σ' from any path P^e to H is equivalent to 1 modulo p . This establishes the proof of

$$|\llbracket \sigma \rrbracket_{\mathfrak{I}}| \equiv |W_L|^{|I_{\sigma} \cap V_L}| |W_R|^{|I_{\sigma} \cap V_R}| \pmod{p}. \quad \square$$

Finally, with the above properties at hand we show that the existence of an (a, b, p) -path in H yields hardness for $\#_p \text{PARTLABHOMSTOH}$.

LEMMA 6.7. *Let p be a prime and let H be a tree with no automorphism of order p . If there are $a, b \in \mathbb{Z}_p \setminus \{1\}$ such that H has an (a, b, p) -path Q_H , then $\#_p \text{HOMSTOH}$ is $\#_p P$ -hard under Turing reductions.*

PROOF. We will show that $\#_p \text{BIS}_{a-1, b-1}$ reduces to $\#_p \text{PARTLABHOMSTOH}$ under polynomial time Turing reductions. Since $a, b \not\equiv 1 \pmod{p}$, the lemma then results from Theorems 1.6 and 1.8. Let $G = (V_L, V_R, E)$ be the bipartite graph, input for $\#_p \text{BIS}_{a-1, b-1}$. We construct the partially labelled graph J according to Definition 6.5, using the path Q_H . Note that $|V(G(J))|$ is polynomial in $|V(G)|$. Since Q_H is an (a, b, p) -path in H , the conditions of Lemma 6.6 are satisfied.

Let $\sigma_1, \dots, \sigma_\mu$ be representatives from each $\sim_3$ -equivalence class as given by Property 3 of Lemma 6.6. We obtain

$$|\text{Hom}(J \rightarrow H)| = \sum_{i=1}^{\mu} |\llbracket \sigma_i \rrbracket_3|.$$

By Property 4, for every $i \in [\mu]$, $|\llbracket \sigma_i \rrbracket_3| \equiv |W_L|^{|\mathcal{I}_{\sigma_i} \cap V_L|} |W_R|^{|\mathcal{I}_{\sigma_i} \cap V_R|} \pmod{p}$. Additionally, due to Definition 6.1 of an (a, b, p) -path $|W_L| \equiv a - 1 \pmod{p}$ and $|W_R| \equiv b - 1 \pmod{p}$. We deduce

$$|\text{Hom}(J \rightarrow H)| \equiv \sum_{i=1}^{\mu} (a - 1)^{|\mathcal{I}_{\sigma_i} \cap V_L|} (b - 1)^{|\mathcal{I}_{\sigma_i} \cap V_R|} \pmod{p}.$$

Finally, we recall Property 3 of Lemma 6.6, which yields the equality of the set $\{\mathcal{I}_{\sigma_i} \mid i \in [\mu]\}$ with the set $\mathcal{I}_G$ of independent sets of G . We obtain

$$\begin{aligned} |\text{Hom}(J \rightarrow H)| &\equiv \sum_{i=1}^{\mu} (a - 1)^{|\mathcal{I}_{\sigma_i} \cap V_L|} (b - 1)^{|\mathcal{I}_{\sigma_i} \cap V_R|} \pmod{p} \\ &\equiv \sum_{I \in \mathcal{I}(G)} (a - 1)^{|I \cap V_L|} (b - 1)^{|I \cap V_R|} \pmod{p}. \end{aligned}$$

The latter is exactly the definition of $Z_{a-1, b-1}(G)$, which concludes the proof. $\square$

7 DICHOTOMY THEOREMS

In this section, we gather our results into the following dichotomy theorem.

THEOREM 1.2. *Let p be a prime and let H be a graph, such that its order p reduced form H^{*p} is a tree. If H^{*p} is a star, then $\#_p \text{HOMSToH}$ is computable in polynomial time; otherwise, $\#_p \text{HOMSToH}$ is $\#_p$ P-complete.*

PROOF. If H^{*p} is a complete bipartite graph, then Corollary 4.7 yields that $\#_p \text{HOMSToH}^{*p}$ can be computed in polynomial time. We note that in this case H^{*p} has to be a star. Otherwise, H^{*p} is not a star and by Lemma 6.2, H^{*p} contains an (a, b, p) -path and thus Lemma 6.7 shows that $\#_p \text{HOMSToH}^{*p}$ is $\#_p$ P-hard. The theorem then follows from Theorem 4.2. $\square$

To justify our title, we use the following proposition showing that our dichotomy theorem holds for all trees. In [10, Section 5.3], this was stated as an obvious fact; however, for the sake of completeness we provide a formal proof.

PROPOSITION 7.2. *Let H be a tree and ϱ an automorphism of H . The subgraph H^ϱ of H induced by the fixed points of ϱ is also a tree.*

PROOF. Let H be a tree and ϱ an automorphism of H . H^ϱ is a subgraph of H , so it suffices to show that H^ϱ is connected. Toward a contradiction, we assume that H^ϱ is not connected. Thus, there exist two vertices $u, v \in V(H)$, whose images $\varrho(u), \varrho(v)$ belong to distinct components in H^ϱ . Since H^ϱ only contains the fixed points under ϱ , we obtain $\varrho(u) = u$ and $\varrho(v) = v$. Therefore, there has to exist a vertex w on a path P from u to v in H with $\varrho(w) \neq w$. Since ϱ is an automorphism, w must be mapped by ϱ to a vertex of some other path P' connecting u and v , where $P' \neq P$. The latter contradicts the assumption that H is a tree. $\square$

The claim implies that if H is a tree, then its order p reduced form H^{*p} is also a tree. This yields the following corollary.

COROLLARY 1.3. *Let p be a prime and let H be a tree. If the order p reduced form H^{*p} of H is a star, then $\#_p \text{HOMSToH}$ is computable in polynomial time; otherwise, $\#_p \text{HOMSToH}$ is $\#_p$ P-complete.*

To allow for disconnected graphs Faben and Jerrum [10, Theorem 6.1] show the following theorem.

THEOREM 7.4 (FABEN AND JERRUM). *Let H be a graph that has no automorphism of order 2. If H' is a connected component of H and $\#_2\text{HOMSToH}'$ is $\#_2$ P-hard, then $\#_2\text{HOMSToH}$ is $\#_2$ P-hard.*

The only part of the proof [10] requiring the value 2 of the modulo is the application of their pinning theorem [10, Theorem 4.7]. Since we have already shown the more general Theorem 1.8, we conclude that the theorem holds in the following form.

THEOREM 7.5. *Let p be a prime and let H be a graph that has no automorphism of order p . If H_1 is a connected component of H and $\#_p\text{HOMSToH}_1$ is $\#_p$ P-hard, then $\#_p\text{HOMSToH}$ is $\#_p$ P-hard.*

The latter strengthens Theorem 1.2 to the following version.

COROLLARY 1.4. *Let H be a graph whose order p reduced form H^{*p} is a forest. If every component of H^{*p} is a star, then $\#_p\text{HOMSToH}$ is computable in polynomial time; otherwise, $\#_p\text{HOMSToH}$ is $\#_p$ P-complete.*

8 COMPOSITE NUMBERS

We investigate counting homomorphisms modulo a composite integer k and observe that we may restrict our attention to powers of primes. With this the natural question arises of whether $\#_{p^r}\text{HOMSToH}$ being computable in polynomial time is equivalent to $\#_p\text{HOMSToH}$ being computable in polynomial time, where p is a prime and r a positive integer. We answer this question negatively, by presenting a graph H for which $\#_2\text{HOMSToH}$ is computable in polynomial time, while $\#_4\text{HOMSToH}$ is $\#_2$ P-hard. This contrasts results by Guo et al. [17] on counting constraint satisfaction problems modulo an integer.

To study the complexity of $\#_k\text{HOMSToH}$ for composite integers k , we will use the Chinese remainder theorem. Recall that integers k_1 and k_2 are said to be *relatively prime* if their only common divisor is 1.

THEOREM 8.1 (CHINESE REMAINDER THEOREM). *Let $\{k_i\}_{i=1}^m$ be a pairwise relatively prime family of positive integers, and let $a_1, \dots, a_m$ be arbitrary integers. Then there exists a solution $a \in \mathbb{N}$ to the system of congruences*

$$a \equiv a_i \pmod{k_i} \quad (i = 1, \dots, m).$$

Moreover, any $a' \in \mathbb{N}$ is a solution to this system of congruences if and only if $a \equiv a' \pmod{k}$, where $k = \prod_{i=1}^m k_i$.

For a proof see, e.g., [5, Theorem 17, Chapter 7].

LEMMA 8.2. *Let $k \in \mathbb{Z}_{>0}$ be an integer and $\prod_{i=1}^m k_i$ with $k_i = p_i^{r_i}$ its prime factorisation with primes $p_1, \dots, p_m$ and positive integers $r_1, \dots, r_m \in \mathbb{Z}_{>0}$. Then $\#_k\text{HOMSToH}$ can be solved in polynomial time if and only if, for each $i \in [m]$, $\#_{k_i}\text{HOMSToH}$ can also be solved in polynomial time.*

PROOF. Since k_i is a factor of k , we take the solution of $\#_k\text{HOMSToH}$ modulo k_i and obtain a solution for $\#_{k_i}\text{HOMSToH}$. From the Chinese remainder theorem (Theorem 8.1) the converse is also true: if for each $i \in [m]$ we can solve $\#_{k_i}\text{HOMSToH}$ in polynomial time, then we can also solve $\#_k\text{HOMSToH}$ in polynomial time. $\square$

With this lemma in mind the subsequent question is whether $\#_{p^r}\text{HOMSToH}$ is computable in polynomial time if and only if $\#_p\text{HOMSToH}$ is computable in polynomial time. Clearly, the first argument in the proof of Lemma 8.2 shows that if $\#_k\text{HOMSToH}$ is computable in polynomial time, then so is $\#_p\text{HOMSToH}$, as we can apply the modulo p operation to a solution of an instance of $\#_k\text{HOMSToH}$. We will show by a counterexample that the reverse implication does not hold.

Namely, we show that for the path with 4 vertices P_4 , $\#_2\text{HomsTo}P_4$ is computable in polynomial time whereas $\#_4\text{HomsTo}P_4$ is $\#_2$ P-hard.

LEMMA 8.3. *Let P_4 denote the path $w_1w_2w_3w_4$. Then $\#_2\text{HomsTo}P_4$ is computable in polynomial time.*

PROOF. The function $\varrho = \{w_1 \mapsto w_4, w_4 \mapsto w_1, w_2 \mapsto w_3, w_3 \mapsto w_2\}$ is an automorphism of order 2 for P_4 without fixed points, so P_4^{*2} is the empty graph. Trivially, for any non-empty input graph G $\#_2\text{HomsTo}P_4^{*2}$ is always zero. Thus, $\#_2\text{HomsTo}P_4$ is computable in polynomial time by Corollary 4.7. $\square$

Regarding the hardness of $\#_4\text{HomsTo}P_4$, we will use the following problem as an intermediate stop in our chain of reductions.

PROBLEM 8.4. *Name.* $\#_k\text{CONBIS}$.

Parameter. Positive integer k .

Input. Connected bipartite graph G .

Output. $|\mathcal{I}(G)| \pmod k$.

Recall Theorem 3.1 showing that $\#_k\text{BIS}$ is $\#_k$ P-complete for all integers k . The next lemma shows that $\#_k\text{CONBIS}$ is also hard for all positive integers.

LEMMA 8.5. *For all integers k , $\#_k\text{CONBIS}$ is $\#_k$ P-complete.*

PROOF. We will provide a Turing reduction from $\#_k\text{BIS}$ and then the lemma follows from Theorem 3.1. Let $G = (V_L, V_R, E)$ be a bipartite graph, input for $\#_k\text{BIS}$. Assume, without loss of generality, that all the isolated vertices of G are contained in V_L . We construct an instance G' for $\#_k\text{CONBIS}$ by adding an extra vertex v_0 to a copy of G and connecting v_0 with all the vertices in V_L . That is, $V(G') = V(G) \cup \{v_0\}$ and $E(G') = E \cup \{(v, v_0) \mid v \in V_L\}$.

We claim that $|\mathcal{I}(G)| + 2^{|V_R|} = |\mathcal{I}(G')|$. Let $\mathcal{I}_1(G') = \{I \in \mathcal{I}(G') \mid v_0 \in I\}$ and let $\mathcal{I}_2(G') = \{I \in \mathcal{I}(G') \mid v_0 \notin I\}$. $\mathcal{I}_1(G')$ and $\mathcal{I}_2(G')$ partition $\mathcal{I}(G')$. For every $I \in \mathcal{I}_1(G')$, we have that $I \cap V_L = \emptyset$, because in G' every vertex in V_L is adjacent to v_0 . Any subset of V_R can be an independent set in $\mathcal{I}_1(G')$, hence $|\mathcal{I}_1(G')| = 2^{|V_R|}$. It remains to show that $|\mathcal{I}_2(G')| = |\mathcal{I}(G)|$. Since v_0 is not in any independent set in $|\mathcal{I}_2(G')|$, every independent set of G is an independent set in $\mathcal{I}_2(G')$ and vice versa. $\square$

We now come to prove the $\#_2$ P-hardness of counting the homomorphisms to P_4 modulo 4.

PROPOSITION 8.6. *Let P_4 be the path $w_1w_2w_3w_4$. Then $\#_4\text{HomsTo}P_4$ is $\#_2$ P-hard.*

PROOF. We will show that $\#_2\text{CONBIS}$ reduces to $\#_4\text{HomsTo}P_4$. Let $G = (V_L, V_R, E)$ be an instance of $\#_2\text{CONBIS}$. We proceed by showing $2|\mathcal{I}(G)| = |\text{Hom}(G \rightarrow P_4)|$.

Given $I \in \mathcal{I}(G)$, we define $\sigma_I : V(G) \rightarrow V(P_4)$ to be the following mapping:

$$\sigma_I(v) = \begin{cases} w_1, & \text{if } v \in V_L \cap I \\ w_2, & \text{if } v \in V_R \setminus I \\ w_3, & \text{if } v \in V_L \setminus I \\ w_4, & \text{if } v \in V_R \cap I. \end{cases}$$

To show that σ_I is a homomorphism, we will show that for all $(v_1, v_2) \in E$, $(\sigma_I(v_1), \sigma_I(v_2)) \in E(P_4)$. Without loss of generality, we assume v_1 is in V_L , then $\sigma_I(v_1) \in \{w_1, w_3\}$, and since $v_2 \in V_R$, we have $\sigma_I(v_2) \in \{w_2, w_4\}$ by the definition of σ_I . Assume toward a contradiction $\sigma_I(v_1) = w_1$ and $\sigma_I(v_2) = w_4$. For the latter to hold, v_1 and v_2 must both lie in I , which is not possible, since I is

an independent set and $(v_1, v_2) \in E$. With this, we obtain $(\sigma_I(v_1), \sigma_I(v_2)) \in E(P_4)$, and therefore $\sigma_I \in \text{Hom}(G \rightarrow P_4)$.

Let $\varrho = \{w_1 \mapsto w_4, w_4 \mapsto w_1, w_2 \mapsto w_3, w_3 \mapsto w_2\}$ be the automorphism of order 2 of P_4 . Clearly, $\varrho \circ \sigma_I$ is a homomorphism different from σ_I , as they differ on all $v \in V(G)$. Thus, every I yields the two homomorphisms $\sigma_I, \varrho \circ \sigma_I \in \text{Hom}(G \rightarrow P_4)$.

Let $I, I' \in \mathcal{I}(G)$ with $I \neq I'$ be two different independent sets in G . Without loss of generality there exists $v \in I \setminus I'$. For this v all four values $\sigma_I(v)$, $(\varrho \circ \sigma_I)(v)$, $\sigma_{I'}(v)$ and $(\varrho \circ \sigma_{I'})(v)$ are different, thus $\sigma_I, \varrho \circ \sigma_I, \sigma_{I'}$ and $\varrho \circ \sigma_{I'}$ are four different elements of $\text{Hom}(G \rightarrow P_4)$.

It remains to argue that for every $\sigma \in \text{Hom}(G \rightarrow P_4)$ there is some $I \in \mathcal{I}(G)$, such that either $\sigma = \sigma_I$ or $\sigma = \varrho \circ \sigma_I$. To this end, we argue that

$$I_\sigma = \{v \in V(G) \mid \sigma(v) \in \{w_1, w_4\}\}$$

is an independent set of G . Let $v_1, v_2 \in I_\sigma$. The definition of I_σ yields $(\sigma(v_1), \sigma(v_2)) \notin E(P)$. As σ is a homomorphism, there can be no edge $(v_1, v_2) \in E$, so I_σ is an independent set of G . We conclude the proof by showing that either $\sigma = \sigma_{I_\sigma}$ or $\sigma = \varrho \circ \sigma_{I_\sigma}$. Let $v \in V_L \cap I_\sigma$. If $\sigma(v) = w_1$, then $\sigma = \sigma_{I_\sigma}$, because G is connected. However, $\sigma(v) = w_4$ implies $\sigma = \varrho \circ \sigma_{I_\sigma}$ and the proposition follows. $\square$

ACKNOWLEDGMENTS

The first author thanks Leslie Ann Goldberg and David Richerby for fruitful discussions during the early stages of this work.

REFERENCES

- [1] M. A. Armstrong. 1988. *Groups and Symmetry*. Springer-Verlag.
- [2] A. A. Bulatov and M. Grohe. 2005. The complexity of partition functions. *Theor. Comput. Sci.* 348, 2–3 (2005), 148–186. <https://doi.org/10.1016/j.tcs.2005.09.011>
- [3] J.-Y. Cai, X. Chen, and P. Lu. 2013. Graph homomorphisms with complex values: A dichotomy theorem. *SIAM J. Comput.* 42, 3 (2013), 924–1029. <https://doi.org/10.1137/110840194>
- [4] J.-Y. Cai and P. Lu. 2011. Holographic algorithms: From art to science. *J. Comput. Syst. Sci.* 77, 1 (2011), 41–61.
- [5] D. S. Dummit and R. M. Foote. 1991. *Abstract Algebra*. Prentice Hall.
- [6] M. E. Dyer, A. M. Frieze, and M. Jerrum. 2002. On counting independent sets in sparse graphs. *SIAM J. Comput.* 31, 5 (2002), 1527–1541. <https://doi.org/10.1137/S0097539701383844>
- [7] M. E. Dyer and C. S. Greenhill. 2000. The complexity of counting graph homomorphisms. *Rand. Struct. Algor.* 17, 3–4 (2000), 260–289.
- [8] J. Faben. 2008. The complexity of counting solutions to generalised satisfiability problems modulo k . arXiv:0809.1836. Retrieved from <https://arxiv.org/abs/0809.1836>.
- [9] J. Faben. 2012. The Complexity of Modular Counting in Constraint Satisfaction Problems. Ph.D. Dissertation. Queen Mary, University of London.
- [10] J. Faben and M. Jerrum. 2015. The complexity of parity graph homomorphism: An initial investigation. *Theory Comput.* 11 (2015), 35–57.
- [11] J. Focke, L. A. Goldberg, M. Roth, and S. Zivný. 2021. counting homomorphisms to K_4 -minor-free graphs, modulo 2. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA'21)*. 2303–2314.
- [12] A. Göbel (A. Gkoppel-Magkakakis). 2016. *Counting, Modular Counting and Graph Homomorphisms*. Ph.D. Dissertation. University of Oxford.
- [13] A. Göbel, L. A. Goldberg, and D. Richerby. 2014. The complexity of counting homomorphisms to cactus graphs modulo 2. *ACM Trans. Comput. Theory* 6, 4 (2014), 17:1–17:29.
- [14] A. Göbel, L. A. Goldberg, and D. Richerby. 2016. Counting homomorphisms to square-free graphs, modulo 2. *ACM Trans. Comput. Theory*, 8, 3 (2016), 12:1–12:29.
- [15] L. A. Goldberg, M. Grohe, M. Jerrum, and M. Thurley. 2010. A complexity dichotomy for partition functions with mixed signs. *SIAM J. Comput.* 39, 7 (2010), 3336–3402.
- [16] L. M. Goldschlager and I. Parberry. 1986. On the construction of parallel computers from various bases of Boolean functions. *Theor. Comput. Sci.* 43 (1986), 43–58.
- [17] H. Guo, S. Huang, P. Lu, and M. Xia. 2011. The complexity of weighted boolean #CSP modulo k . In *Proceedings of the Symposium on Theoretical Aspects of Computer Science (STACS'11)*. 249–260.

- [18] P. Hell and J. Nešetřil. 1990. On the complexity of H -coloring. *J. Combin. Theory Ser. B* 48, 1 (1990), 92–110.
- [19] A. Kazeminia and A. A. Bulatov. 2019. Counting homomorphisms modulo a prime number. In *Proceedings of the International Symposium on Mathematical Foundations of Computer Science (MFCS'19)*. 59:1–59:13. <https://doi.org/10.4230/LIPIcs.MFCS.2019.59>
- [20] R. E. Ladner. 1975. On the structure of polynomial time reducibility. *J. ACM* 22, 1 (1975), 155–171.
- [21] J. A. G. Lagodzinski, A. Göbel, K. Casel, and T. Friedrich. 2020. On counting (quantum-)graph homomorphisms in finite fields. In *Proceedings of the International Colloquium on Automata, Languages and Programming (ICALP'21)*. To appear.
- [22] L. Lovász. 1967. Operations with structures. *Acta Math. Acad. Sci. Hung.* 18, 3–4 (1967), 321–328.
- [23] C. H. Papadimitriou and S. Zachos. 1982. Two remarks on the power of counting. In *Proceedings of the GI-Conference on Theoretical Computer Science*. 269–276.
- [24] J. Simon. 1975. *On Some Central Problems in Computational Complexity*. Ph.D. Dissertation. Ithaca, NY.
- [25] S. Toda. 1991. PP is as hard as the polynomial-time hierarchy. *SIAM J. Comput.* 20, 5 (1991), 865–877.
- [26] L. G. Valiant. 2006. Accidental algorithms. In *Proceedings of the IEEE Symposium on Foundations of Computer Science (FOCS'06)*. 509–517.
- [27] D. B. West. 2000. *Introduction to Graph Theory* (2nd ed.). Prentice Hall.

Received December 2019; revised January 2021; accepted March 2021