Counting list homomorphisms from graphs of bounded treewidth: tight complexity bounds

Jacob Focke[*]   Dániel Marx[†]   Paweł Rzążewski[‡]

Abstract

The goal of this work is to give precise bounds on the counting complexity of a family of generalized coloring problems (list homomorphisms) on bounded-treewidth graphs. Given graphs G, H, and lists $L(v) \subseteq V(H)$ for every $v \in V(G)$, a *list homomorphism* is a function $f : V(G) \to V(H)$ that preserves the edges (i.e., $uv \in E(G)$ implies $f(u)f(v) \in E(H)$) and respects the lists (i.e., $f(v) \in L(v)$). Standard techniques show that if G is given with a tree decomposition of width t, then the number of list homomorphisms can be counted in time $|V(H)|^t \cdot n^{\mathcal{O}(1)}$. Our main result is determining, for every fixed graph H, how much the base $|V(H)|$ in the running time can be improved. For a connected graph H we define $\mathrm{irr}(H)$ in the following way: if H has a loop or is nonbipartite, then $\mathrm{irr}(H)$ is the maximum size of a set $S \subseteq V(H)$ where any two vertices have different neighborhoods; if H is bipartite, then $\mathrm{irr}(H)$ is the maximum size of such a set that is fully in one of the bipartition classes. For disconnected H, we define $\mathrm{irr}(H)$ as the maximum of $\mathrm{irr}(C)$ over every connected component C of H. It follows from earlier results that if $\mathrm{irr}(H) = 1$, then the problem of counting list homomorphisms to H is polynomial-time solvable, and otherwise it is #P-hard. We show that, for every fixed graph H, the number of list homomorphisms from (G, L) to H

- can be counted in time $\mathrm{irr}(H)^t \cdot n^{\mathcal{O}(1)}$ if a tree decomposition of G having width at most t is given in the input, and

- given that $\mathrm{irr}(H) \geq 2$, cannot be counted in time $(\mathrm{irr}(H) - \varepsilon)^t \cdot n^{\mathcal{O}(1)}$ for any $\varepsilon > 0$, even if a tree decomposition of G having width at most t is given in the input, unless the Counting Strong Exponential-Time Hypothesis (#SETH) fails.

Thereby we give a precise and complete complexity classification featuring matching upper and lower bounds for all target graphs with or without loops.

1 Introduction

Many of the NP-hard problems studied in the literature are known to be polynomial-time solvable when restricted to graphs of bounded treewidth. In fact, the majority of these problems can be solved in time $f(t) \cdot n^{\mathcal{O}(1)}$ if a tree decomposition of width t is given in the input, that is, they are fixed-parameter tractable (FPT) parameterized by treewidth. As algorithms working on tree decompositions are useful building blocks in many types of FPT results and approximation schemes, determining the optimal dependence $f(t)$ on the width of the decomposition received significant attention.

On the upper bound side, new algorithmic techniques (such as Fast Subset Convolution, Cut & Count, representative sets) were developed to obtain improved algorithms. For lower bounds, conditional complexity results were given ruling out certain forms of running times. Lokshtanov, Marx, and Saurabh [22] considered problems that are known to be solvable in time $c^t \cdot n^{\mathcal{O}(1)}$ if a tree decomposition of width t is given in the input, and showed that these algorithms are essentially optimal, as no algorithm with running time $(c - \varepsilon)^t \cdot n^{\mathcal{O}(1)}$ can exist for any $\varepsilon > 0$, assuming the Strong Exponential-Time Hypothesis (SETH). In particular, for VERTEX COLORING with c colors, the textbook $c^t \cdot n^{\mathcal{O}(1)}$ algorithm based on dynamic programming cannot be improved to $(c - \varepsilon)^t \cdot n^{\mathcal{O}(1)}$. By now, there is a growing collection of tight lower bounds of this form in the literature [2, 6, 7, 11, 20, 23, 26, 27].

Vertex coloring with c colors can be seen as a homomorphism problem. Given graphs G and H, a *homomorphism* is a mapping $f : V(G) \rightarrow V(H)$ that preserves the edges of G, that is, $uv \in E(G)$ implies $f(u)f(v) \in E(H)$. Let us observe that a c-coloring of G can be seen equivalently as a homomorphism from G to K_c, the complete graph with c vertices. Thus generalized coloring problems defined by homomorphisms to a fixed graph H were intensively studied in the combinatorics literature and subsequently from the viewpoint of computational complexity [14, 15, 17–19, 24, 25]. The homomorphism problem can be generalized from graphs to arbitrary relational structures, giving a very direct connection to constraint satisfaction problems (CSPs), which are often described using the terminology of homomorphisms [13, 17–19].

Introducing lists of allowed images gives a more robust variant for homomorphism problems: formally, given graphs G and H, a *list assignment* is a function $L : V(G) \rightarrow 2^{V(H)}$. Then a list homomorphism from (G, L) to H is a homomorphism $f : V(G) \rightarrow V(H)$ that additionally respects the lists, that is, $f(v) \in L(v)$ for every $v \in V(G)$. Studying the list version of a homomorphism problem can be seen as analogous to studying the conservative version of CSP, where every unary constraint is allowed [1, 4, 5]. The study of conservative CSP often served as a starting point before more general investigations, motivating the exploration of the list version of homomorphism problems.

The polynomial-time solvable cases of (list) homomorphism is well understood. Nešetřil and Hell [15] showed that the non-list problem is polynomial-time solvable if H is bipartite or has a loop, and NP-hard when restricted to any other H. The complexity of the list version was characterized by Feder, Hell, and Huang [12]: it is polynomial-time solvable if H is a so-called bi-arc graph, and NP-hard for every other fixed H. But are there nontrivial algorithmic ideas that can help us obtain improved running times for the NP-hard cases? Similarly to coloring, the problem of finding a (list) homomorphism to H can be solved in time $|V(H)|^t \cdot n^{\mathcal{O}(1)}$ if G is given with a tree decomposition of width t. Can this straightforward dynamic programming algorithm be improved for a given fixed H, and if so, by how much?

This question was resolved first by Egri, Marx, and Rzążewski [11] for the list homomor-

phism problem in the special case when H is reflexive (that is, every vertex has a loop), which was extended by Okrasa, Piecyk, and Rzążewski [26] to every H (where each vertex may or may not have a loop). It turns out that there are a couple of algorithmic ideas that can be used to obtain $c^t \cdot n^{\mathcal{O}(1)}$ time algorithms with $c < |V(H)|$. In particular, there are delicate notions of decompositions of H, such that if they are present, then the problem can be reduced in a nontrivial way to multiple instances with some H' having fewer vertices than H. To formalize the optimality of these ideas, a combinatorial parameter $i^*(H)$ was introduced, and it was shown that the problem can be solved in time $i^*(H)^t \cdot n^{\mathcal{O}(1)}$, but there is no $(i^*(H) - \varepsilon)^t \cdot n^{\mathcal{O}(1)}$ time algorithm for any $\varepsilon > 0$, assuming the SETH. This means that the technical ideas behind the $i^*(H)^t \cdot n^{\mathcal{O}(1)}$ time algorithm already capture all the possible algorithmic insights that can be exploited when solving the problem on a given tree decomposition. A similar tight result was obtained by Okrasa and Rzążewski [27] for the non-list version of the problem. In that version, a different set of algorithmic ideas become relevant (reduction to a homomorphic core and factorizations), and the optimality of these ideas were proved assuming not only the SETH, but also two long-standing conjectures from algebraic graph theory.

A well-known phenomenon in computational complexity is that even if it is possible to find a solution efficiently in a combinatorial problem, counting the number of solutions can be hard. The most notable example is the perfect matching problem in bipartite graphs: finding a perfect matching is polynomial-time solvable, but counting the number of perfect matchings is #P-hard by the seminal result of Valiant [29]. There are algorithmic ideas that can be generalized from decision to counting (for example, simple forms of dynamic programming), but others may not be. In particular, arguments of the form "if there is a solution, then there is a solution with property P, hence we only need to look for solutions with property P" do not immediately generalize to counting, as they would fail to count the potential solutions that do not have property P. Unfortunately, the $i^*(H)^t \cdot n^{\mathcal{O}(1)}$ time algorithm of Okrasa, Piecyk, and Rzążewski [26] for list homomorphism heavily relies on such arguments. Careful observation shows that only a very limited set of algorithmic ideas remain relevant for the counting problem:

- **Connected components.** We may assume that G is connected: otherwise the number of solutions is the product of the number of solutions for each component of G. Furthermore, if G is connected, then every vertex of G is mapped into the same connected component of H. Thus we may assume that H is connected as well: the number of solutions is the sum of the number of solutions when restricted to each component of H.

- **Identical neighborhoods.** If two vertices u and v have identical neighborhoods, then we may pretend that only one of them, say u, is present in H. Then we need to solve a weighted problem that takes into account that whenever a vertex of G is mapped to u, then we can actually choose from two different copies of u. The extension to this weighted problem can be easily implemented in the dynamic programming algorithm working over a tree decomposition. Therefore, the number of possible images of a vertex can be restricted to the maximum size of a set $S \subseteq V(G)$ where the neighborhoods are pairwise different. We call such a set S an *irredundant* set.

- **Bipartite classes.** If H is bipartite (and has no loops), then G has to be bipartite as well. This means that each bipartite class of G is mapped to one of the two bipartite classes of H, resulting in two cases that we can treat separately. In each case, the possible images of a vertex v of G are restricted to a bipartite class of H. Now the

2

maximum number of possible images of v can be bounded by the maximum size of an irredundant set S that is contained in one bipartite class of H.

Our main negative result shows that these are all the algorithmic ideas that can be exploited in the list homomorphism problem for a fixed H. To state this formally, we define the following graph parameter.

Definition 1.1. Given a graph H, we say that $S \subseteq V(H)$ is *irredundant* if any two vertices have different neighborhoods. We say that graph H is irredundant if $V(H)$ is irredundant. For a connected bipartite graph H, set S is *one-sided* if it is fully contained in one of the bipartite classes.

For a connected graph H, we define $\mathrm{irr}(H)$ the following way. If H has a loop or is nonbipartite, then $\mathrm{irr}(H)$ is the maximum size of an irredundant set; if H is bipartite, then $\mathrm{irr}(H)$ is the maximum size of a one-sided irredundant set. For disconnected H, we define $\mathrm{irr}(H)$ as the maximum of $\mathrm{irr}(C)$ over every connected component C of H.

Note that if v has a loop, then the neighborhood of v includes v itself. This has to be carefully taken into account when interpreting different neighborhoods. For example, in a reflexive clique (where every vertex has a loop) every vertex has the same neighborhood.

For a fixed graph H (possibly with loops), we denote by $\#\mathrm{LHom}(H)$ the problem of counting the number of list homomorphisms from the given (G, L) to H. Our main result shows for $\#\mathrm{LHom}(H)$ that $\mathrm{irr}(H)$ is indeed the best possible base of the exponent that can appear in the running time. We obtain the lower bound under the #SETH, the counting version of the SETH, which is the natural variant of the assumption for lower bounds on counting problems (see [8]). Note that the SETH implies the #SETH, thus proving a result under the latter assumption make it somewhat stronger.

Theorem 1.2. *Let H be a graph that has a component which is not a biclique nor a reflexive clique. On n-vertex instances with treewidth* tw, *the $\#\mathrm{LHom}(H)$ problem*

1. *can be solved in time $\mathrm{irr}(H)^{\mathrm{tw}} \cdot n^{\mathcal{O}(1)}$, provided that a tree decomposition of width* tw *is given as part of the input.*
2. *cannot be solved in time $(\mathrm{irr}(H) - \varepsilon)^{\mathrm{tw}} \cdot n^{\mathcal{O}(1)}$, for any $\varepsilon > 0$, even if a tree decomposition of width* tw *is given as part of the input, unless the #SETH fails.*

While the counting complexity of homomorphism problems is well-studied, we may argue that our lower bound brings a new level of understanding on the real hardness of the problem. It is easy to deduce from earlier results that if H has a component that is neither a biclique nor a reflexive clique, then H contains an induced subgraph H' on at most 4 vertices such that $\#\mathrm{LHom}(H')$ is #P-hard. As we can use the lists to restrict the problem to $V(H')$, it follows that $\#\mathrm{LHom}(H)$ is #P-hard as well. Dyer and Greenhill [10] showed that for such an H the homomorphism counting problem is #P-hard even if no lists are allowed. Since we cannot use the lists to restrict the problem to $V(H')$, this proof is significantly more complicated. What all these proofs have in common is that they identify only one level of complexity: #P-hardness. Intuitively, counting 10-colorings feels harder than counting 3-colorings, as we need to consider a larger number of possible values at each vertex. Saying that both problems are #P-hard ignores this perceived difference of hardness. On the other hand, Theorem 1.2 establishes, in a formal way, a difference in complexity between $\#\mathrm{LHom}(K_{10})$ and $\#\mathrm{LHom}(K_3)$: the best possible base is 10 in the former and 3 in the latter.

Typical #P-hardness proofs use as the basis of the reduction a few known #P-hard problems, such as counting independent sets or counting 3-colorings. Once it is established that the problem can express one of these hard problems, then the proof can stop, as it is irrelevant whether some even harder problem can also be expressed. On the other hand, the lower bound proof of Theorem 1.2 has to delve deeper into the complexity of the problem, as it should extract as much hardness from the problem as possible. So instead of trying to show that the problem can express a simple relation corresponding to e.g., independent set or 3-coloring, our main goal is to show that, in a formal sense, #LHom(H) can express every relation over a domain of size irr(H). Going beyond the context of bounded-treewidth graphs, the main message of this paper is that the #LHom(H) problem is sufficiently rich to express all such relations.

1.1 Proof overview

As described above, the algorithmic part of the main result is fairly simple: we need to do only a few changes to the standard dynamic programming algorithm to take into account connected components, vertices with identical neighborhoods, and bipartiteness. Therefore, we focus on the main steps of the lower bound proof in this section.

From SAT to CSP. The #SETH states that for every $\varepsilon > 0$, there is a $d \geq 1$ such that n-variable #d-SAT (i.e., counting satisfying assignments of a Boolean formula where every clause has d literals) cannot be solved in time $\mathcal{O}((2 - \varepsilon)^n)$. Therefore, we need to show that the hypothetical fast algorithm for #LHom(H) would give an algorithm for #SAT violating this lower bound. It will be convenient to introduce an intermediate problem. By #CSP(q, r), we denote the problem of counting solutions in a CSP instance with domain $[q]$ and each constraint having arity at most r. It follows from the work of Lampis [21] that if for some $q \geq 2$ there is an $\varepsilon > 0$ such that #CSP(q, r) can be solved in time $(q - \varepsilon)^n \cdot (n + m)^{\mathcal{O}(1)}$ for every r, then the #SETH fails. Therefore, to prove our main lower bound for #LHom(H), we take a #CSP(q, r) instance with $q = $ irr(H) and show that it can be reduced to #LHom(H).

Focusing on bipartite H. Similarly to [26], first we prove the main result for bipartite H. Then it is a surprisingly simple task to raise the result to general H. Given a (not necessarily bipartite) graph H, we define the *associated* bipartite graph H^* the following way (see Figure 1): for every vertex v of H, there are are two vertices v', v'' in H^* and we introduce edges such that, if $uv \in E(H)$, then $u'v'', u''v' \in E(H^*)$ (if v has a loop in H, then $v'v'' \in E(H^*)$). It can be shown that irr(H^*) = irr(H) and the lower bound for #LHom(H) can be obtained by a simple reduction from #LHom(H^*). Thus in the following we can assume that H is bipartite. Furthermore, we may assume that H is irredundant, that is, any two vertices have distinct neighborhoods. It is easy to see that the lower bound for irredundant H can be extended to the general case.

New relations. Input to #LHom(H) can be seen as a CSP instance where the vertices are variables with domain $V(H)$ and each edge is a binary constraint restricting the combination of values appearing on two vertices. We can augment the problem by allowing a fixed finite set $\mathcal{R} = \{R_1, \ldots, R_c\}$ of other relations in the instance. Of course, introducing new relations may make the problem harder. If for every $1 \leq i \leq c$, we can provide an appropriate reduction from the problem augmented with $\{R_1, \ldots, R_i\}$ to the problem augmented only
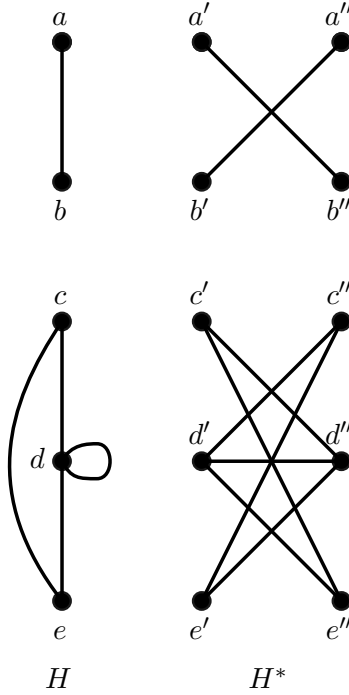
4

Figure 1: A graph H and its associated bipartite graph H^*.

with $\{R_1, \ldots, R_{i-1}\}$, then this chain of reductions shows that the problem augmented with $\mathcal{R}$ is not harder than the original #LHom(H) problem, hence any lower bound proved for the augmented problem holds for #LHom(H) as well. Therefore, we start with #LHom(H) and try to introduce new relations R_i one by one, in a way that does not make the problem any harder. We do this until we can introduce every r-ary relation over a domain S of size $q = \mathrm{irr}(H)$. At this point, the lower bound for #CSP(q, r) applies to our augmented problem, and hence to the original #LHom(H) problem as well.

Gadgets and interpolation. How can we show that introducing a new relation R_i does not make the problem harder? This can be done by showing that each occurrence of the R_i relation can be replaced by an appropriate gadget that, in a certain sense, simulates this relation. A *gadget* is a small instance where a subset $(x_1, \ldots, x_r)$ of the vertices are distinguished as the *interface*. Now if we fix the images $f(x_1) = v_1, \ldots, f(x_r) = v_r$, then, for each vector $(v_1, \ldots, v_r)$, there is some number of ways the mapping f can be extended to a list homomorphism from the gadget to H. If it so happens that for every $(v_1, \ldots, v_r) \in R_i$ there is exactly one extension, and for every $(v_1, \ldots, v_r) \notin R_i$ there are zero extensions, then this gadget effectively expresses a constraint with the relation R_i on the vertices $(x_1, \ldots, x_r)$. Now we can replace every occurrence of R_i with a copy of this gadget, showing that introducing R_i does not make the problem harder.

However, it is rare that a gadget can express the new relation R_i in such a clean way. Fortunately, the nature of counting problems allows us to use gadgets in a much more general setting. Suppose the number of possible extensions for every $(v_1, \ldots, v_r) \in R_i$ is always an integer from a set A (for example, $A = \{4, 5, 8, 11\}$), while the number of extensions for

$(v_1, \ldots, v_r) \notin R_i$ is always an integer from a set B (for example, $B = \{3, 7, 9\}$). Suppose that every $a \in A$ and $b \in B$ are coprime[1] (that is, there is no prime p dividing both an element of A and an element of B); this condition holds in our example. Then a polynomial interpolation technique of Dyer and Greenhill [10] allows us to count those assignments where $(v_1, \ldots, v_r) \in R_i$ is satisfied, effectively introducing a constraint with relation R_i.

A simple application of this technique allows us to express compositions of relations. Given two binary relations R_1 and R_2, their *composition* is defined as $R_1; R_2 = \{(a, b) \mid \exists c\colon (a, c) \in R_1, (c, b) \in R_2\}$. Given two interface vertices v_1 and v_2, we can construct a gadget by introducing a new vertex z and adding the constraint R_1 on v_1 and z, and adding the constraint R_2 on z and v_2. Now if the interface vertices are assigned a pair of values from $R_1; R_2$, then there is at least one extension to z; if they are assigned a pair of values not from $R_1; R_2$, then there is no extension. Thus by the technique mentioned in the previous paragraph, we can express the relation $R_1; R_2$.

Relations on a P_4. The path on four vertices (a P_4) is the smallest example of a connected bipartite graph H that is not a complete bipartite graph, and hence the #LHom(H) problem is #P-hard. Intuitively, #P-hardness should mean that whenever a P_4 appears in H, then we should be able to express arbitrarily complicated constraints when restricted to these vertices; however, this has not been stated explicitly in earlier work. We make this expectation formal by showing that if $(a - b - c - d)$ is a P_4 appearing in H, then any relation $R \subseteq \{a, c\}^r$ can be simulated with appropriate gadgets. For this statement it is essential that we consider the counting problem, and an analogous statement should not hold for the decision version. To see this, observe that in the path $(a - b - c - d)$, the neighborhood of c is a superset of the neighborhood of a. Thus if the list of a vertex contains both a and c, then a in the solution can always be replaced by c. This makes it impossible to simulate even very simple relations such as $R = \{aa, cc\}$ in the decision version since allowing aa and cc would automatically allow ac and ca as well.

Forcers. Forcers are crucial objects in our proof. Let $S \subseteq V(H)$, $x, y \in S$, and $\alpha, \beta \in V(H)$, where S, $\{x, y\}$, $\{\alpha, \beta\}$ are all one-sided sets. An (x, y, S)-forcer with respect to (α, β) is a binary relation $R \subseteq S \times \{\alpha, \beta\}$ that

- contains (x, α), but not (x, β),

- contains (y, β), but not (y, α), and

- for every $z \in S \setminus \{x, y\}$, contains at least one of (z, α) or (z, β).

Intuitively, the relation moves x and y to α and β, respectively, while other values of S can move to somewhere in $\{\alpha, \beta\}$. For example, both of the following two relations are $(x, y, \{x, y, z, q, w\})$-forcers with respect to (α, β):

$$R_1 = \{x\alpha, y\beta, z\alpha, z\beta, q\alpha, w\beta\},$$
$$R_2 = \{x\alpha, y\beta, z\alpha, q\alpha, w\alpha\}.$$

We will be particularly interested in forcers where α and β are vertices a and c of a P_4 $a - b - c - d$. The following claim will be crucial for our proof: if we can show that, for every

---

[1]More precisely, we allow 0 to appear only in B, and allow 1 to appear only in A.

$x, y \in S$, there are good (x, y, S)-forcers with respect to (a, c) of the same P_4, then we can realize any r-ary relation R over S. Therefore, if we choose an S with $|S| = \mathrm{irr}(H) = q$, then we can reduce $\#\mathrm{CSP}(q, r)$ to our problem, which was our main goal. Let us sketch how to prove the claim by realizing an r-ary relation R with the help of the forcers. Let $s = |S|$. Intuitively, by applying all the $s(s-1)$ possible (x, y, S)-forcers on a vertex v, we can extract $s(s-1)$ bits of information about the value of v. These bits are sufficient to identify the value of v, as for any two different values $x \neq y$, one of the bits will be different. Therefore, we can translate the values of r variables into $rs(s-1)$ bits and then we can implement R by an $rs(s-1)$-ary relation over $\{a, c\}$. As a and c are on a P_4, we have already seen that such a relation can be realized.

The connected structure of P_4's. As we have seen in the previous paragraph, a crucial step in the proof is to realize an (x, y, S)-forcer for every $x, y \in S$. An important detail is that we need all these (x, y, S)-forcers to be with respect to the same (a, c). But what happens if we have two P_4s $(a - b - c - d)$ and $(a' - b' - c' - d')$, can an (x, y, S)-forcer with respect to (a', c') be turned into one with respect to (a, c)? We can show that if the two P_4s are "adjacent" in the sense that they share two vertices in the same bipartite class of H, then this is indeed possible. Even if the two P_4s are not adjacent, but are in the "same connected component" under this definition of adjacency, then multiple applications of this argument show that a forcer with respect to (a', c') can be turned into a forcer with respect to (a, c). A graph-theoretic analysis shows that if H is connected and irredundant, then the P_4s have only a single connected component. Therefore, it does not matter in the definition of good forcers exactly where a and c are, as long as they are in a P_4. We can show that even this last condition is unnecessary: a forcer with respect to any two distinct (α, β) that are on the same side can be turned into a forcer with respect to (a, c).

The inductive proof. We prove by induction that there is a good (x, y, S)-forcer for every $S \subseteq V(H)$ and $x, y \in S$: when proving the statement for (x, y, S), we assume that it is true for every (x', y', S') where either

- $|S'| < |S|$ or

- $|S'| = |S|$ and the distance condition $\mathrm{dist}(\{x', y'\}, S' \setminus \{x', y'\}) < \mathrm{dist}(\{x, y\}, S \setminus \{x, y\})$ holds.

To explain the main ideas of the induction step, we need two new concepts. A weaker version of the forcer is the distinguisher, which additionally may allow $(y, \alpha) \in R$ as well. For example, the following relation is an $(x, y, \{x, y, z, q, w\})$-distinguisher with respect to (α, β):

$$R = \{x\alpha, y\alpha, y\beta, z\alpha, z\beta, q\alpha, w\beta\}.$$

As another application of the interpolation technique, we show that if we can realize a distinguisher, then we can realize a forcer as well, that is, we can eliminate the unwanted possibility (y, α).

For a partition (X, Y) of S, the (X, Y)-partitioner with respect to (α, β) is the relation

$$R = \{(x, \alpha) \mid x \in X\} \cup \{(y, \beta) \mid y \in Y\}.$$

It can be seen as a stricter version of the (x, y, S)-forcer: it is now specified for every element of S whether the relation should move it to α or β. But we can show that if we realize an
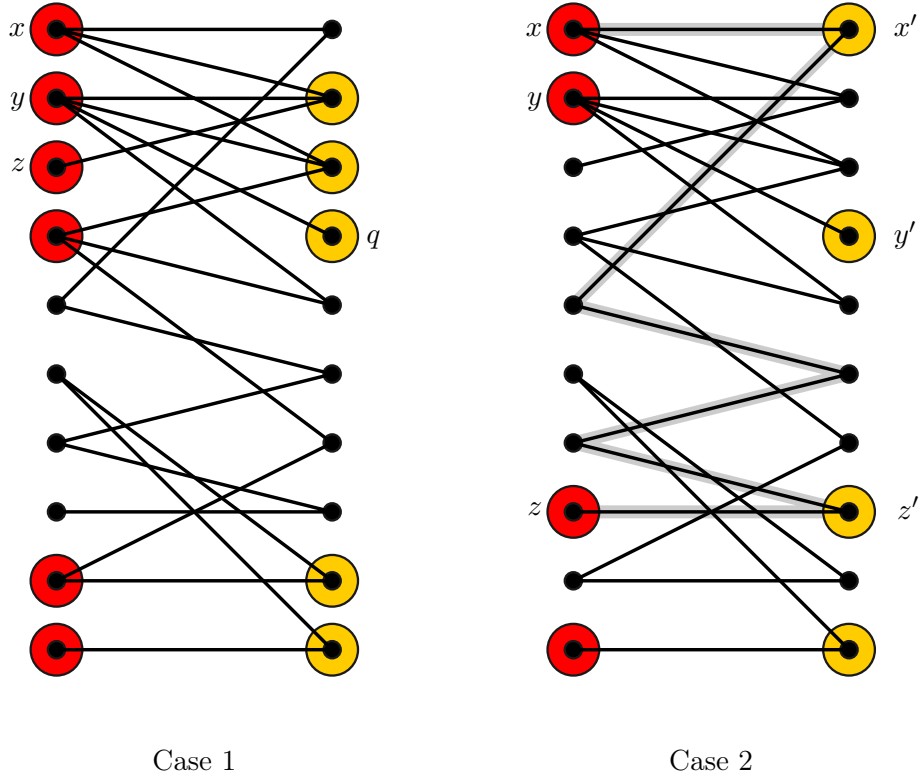
Figure 2: The two cases in the inductive proof. The set S is highlighted in red, the set S' is highlighted in yellow. Case 1: As z and y have a common neighbor, we can define S' with $|S'| < |S|$. Vertex $q \in S'$ is in $N(y) \setminus N(x)$. Case 2: In S, the distance of $\{x, y\}$ from the rest of the vertices is 6 (a shortest path P is highlighted in gray), while in (x', y', S') this distance is only 4.

(x, y, S)-forcer for every $x, y \in S$, then we can realize an (X, Y)-partitioner for every partition (X, Y) of S: the argument is similar to how forcers were used to realize arbitrary relations over S.

Now let us overview the main cases in the inductive proof.

- **Base case:** $S = \{x, y\}$. If x and y have a common neighbor, then x and y are part of a P_4 (as we assumed that H is irredundant) and we essentially have an (x, y, S)-distinguisher, which we can turn into an (x, y, S)-forcer. If x and y have distance at least 3, then we can use the edge relation E of H to move x and y closer to each other, until they have a common neighbor.

- **Case 1:** S contains a vertex z at distance two from $\{x, y\}$. Using that z has a neighbor that is also adjacent to one of x and y, we can define a set S' with $S \subseteq N(S')$ and $|S'| < |S|$ (see Figure 2). Furthermore, we can define S' in a way that, without loss of generality, the set $N(y) \setminus N(x)$ is nonempty. As $|S'| < |S|$, the induction hypothesis implies that there are (x', y', S')-forcers for every $x', y' \in S'$. Therefore, we may also assume that there are (X', Y')-partitioners for every partition (X', Y') of S'. Let R be an $(N(x) \cap S', S' \setminus N(x))$-partitioner with respect to some (α, β). If E is the edge relation

8

of H, then we claim that the composition $E;R$ is an (x, y, S)-distinguisher. Indeed, it moves x to α (as E has to move x to $N(x) \cap S'$ first, which is moved to α by R), y can move to β (as E can move y to the nonempty $(N(y) \setminus N(x)) \cap S'$, which is moved to β by R), and every element of S can move to at least one of α and β (as $S \subseteq N(S')$). Then this (x, y, S)-distinguisher can be turned into an (x, y, S)-forcer, as required.

- **Case 2:** Every vertex of $S \setminus \{x, y\}$ is at distance at least 4 from $\{x, y\}$. Let P be a shortest path between $\{x, y\}$ and $S \setminus \{x, y\}$. There are a few similar cases to consider, but let us assume, as a representative case, that P goes from x to some vertex z, and there is a vertex $y' \in N(y) \setminus N(x)$ (see Figure 2). Let x' be the neighbor of x on P, let z be the other endpoint of P, and let z' be the neighbor of z on P. We define S' starting from $\{x', y', z'\}$ and adding a neighbor of each $S \setminus \{x, y, z\}$; clearly, we have $|S'| \leq |S|$. Furthermore, $\mathrm{dist}(\{x', y'\}, S' \setminus \{x', y'\})$ is strictly less than $\mathrm{dist}(\{x, y\}, S \setminus \{x, y\})$: vertices x' and z' are closer to each other than x and z. Thus by the induction hypothesis, an (x', y', S')-forcer R exists with respect to some (α, β). Then we claim that the composition $E;R$ is an (x, y, S)-distinguisher. Indeed, x is moved to α (as x' is the only neighbor of x in S', which is moved to α by R), y can move to β (as y' is a neighbor of y and R can move y' to β), and every element of S can move to at least one of α and β (as $S \subseteq N(S')$). Then this (x, y, S)-distinguisher can be turned into an (x, y, S)-forcer, as required.

2 Preliminaries

For an integer n, we define $[n] := \{1, 2, \ldots, n\}$. Given a set S, by 2^S we denote the power set of S.

Graph theory By $(p_1 - p_2 - \ldots - p_q)$ we denote the path of length q whose consecutive vertices are $p_1, p_2, \ldots, p_q$. Let H be a graph. For vertices $u, v \in V(H)$, $\mathrm{dist}(u, v)$ is the length of a shortest path between u and v in H. $N(u) = \{v \in V(H) \mid uv \in E(H)\}$ is the (open) neighborhood of u in H. For a set $U \subseteq V(H)$, $N(U)$ is the set of vertices of H with at least one neighbor in U.

Structures A signature σ consists of a finite set of relation symbols with specified arities. For a relation (symbol) R, $\mathrm{ar}(R)$ denotes the arity of R. For $i \in [\mathrm{ar}(R)]$, $\Pi_i(R)$ is the projection onto the i-th entry of the tuples in R. A *structure* $\mathcal{H}$ with signature $\sigma(\mathcal{H})$ consists of a universe $V(\mathcal{H})$ together with a set of relations $\mathbf{R}(\mathcal{H}) = \{R^{\mathcal{H}} \mid R \in \sigma(\mathcal{H})\}$ over the universe $V(\mathcal{H})$. We write $\mathcal{H} = (V(\mathcal{H}), \mathbf{R}(\mathcal{H}))$. By $||\mathcal{H}||$ we denote $|\sigma(\mathcal{H})| + |V(\mathcal{H})| + \sum_{R \in \sigma(\mathcal{H})} |R^{\mathcal{H}}| \cdot \mathrm{ar}(R)$, which is the size of a "reasonable" encoding of $\mathcal{H}$. Note that if $\sigma(\mathcal{H})$ consists of a constant number of relations, each with constant arity, then $||\mathcal{H}|| = |V(\mathcal{H})|^{\mathcal{O}(1)}$. Throughout the paper we work in such a setting.

A structure $\mathcal{H}$ is a graph if $\mathbf{R}(\mathcal{H}) = \{E(\mathcal{H})\}$, where $E(\mathcal{H})$ is a symmetric binary relation over $V(\mathcal{H})$. In this case we write $\mathcal{H} = (V(\mathcal{H}), E(\mathcal{H}))$.

The *Gaifman graph* of a structure $\mathcal{H}$ has vertex set $V(\mathcal{H})$ and distinct $u, v \in V(\mathcal{H})$ are adjacent if and only if there is a relation $R \in \mathbf{R}(\mathcal{H})$ that involves u and v. We define the treewidth and a tree decomposition (resp., the pathwidth and a path decomposition) of $\mathcal{H}$ as the treewidth and a tree decomposition (resp., the pathwidth and a path decomposition) of the Gaifman graph of $\mathcal{H}$.

Given two structures $\mathcal{G}$ and $\mathcal{H}$ with the same signature σ, a function $f \colon V(\mathcal{G}) \to V(\mathcal{H})$ *respects* $R \in \sigma$ if, for each $\mathbf{x} \in R^{\mathcal{G}}$, $f(\mathbf{x}) \in R^{\mathcal{H}}$, where f is evaluated elementwise. A *homomorphism* from $\mathcal{G}$ to $\mathcal{H}$ is a function $h \colon V(\mathcal{G}) \to V(\mathcal{H})$ that respects every $R \in \sigma$. By $\hom(\mathcal{G} \to \mathcal{H})$ we denote the number of homomorphisms from $\mathcal{G}$ to $\mathcal{H}$. If $\mathrm{ar}(R) = 1$ then $R^{\mathcal{H}}$ is a subset of $V(\mathcal{H})$ and we also use the usual *list* notation: For an element $x \in V(\mathcal{G})$ we write $L(x) = R^{\mathcal{H}}$ if $x \in R^{\mathcal{G}}$. More generally, if L is a partial function from $V(\mathcal{G})$ to $2^{V(\mathcal{H})}$, then h is a homomorphism from $(\mathcal{G}, L)$ to $\mathcal{H}$ if h is a homomorphism from $\mathcal{G}$ to $\mathcal{H}$ with $h(x) \in L(x)$ for each x in the domain of L. Note that x being outside the domain of L is as restrictive as having $L(x) = V(\mathcal{H})$. By $\hom\big((\mathcal{G}, L) \to \mathcal{H}\big)$ we denote the number of homomorphisms from $(\mathcal{G}, L)$ to $\mathcal{H}$.

Let $\mathcal{H} = (U, \mathbf{S})$ be a fixed structure and let $\mathcal{J} = (Z, \mathbf{F})$ be a structure with the same signature as $\mathcal{H}$. For some integer p, let $\mathbf{x} = (x_1, x_2, \ldots, x_p) \in Z^p$ and $\mathbf{y} = (y_1, y_2, \ldots, y_p) \in U^p$. Then $\hom\big((\mathcal{J}, \mathbf{x}) \to (\mathcal{H}, \mathbf{y})\big)$ is the number of homomorphisms from $\mathcal{J}$ to $\mathcal{H}$ that map x_i to y_i for each $i \in [p]$. A p-tuple of distinguished (not necessarily pairwise distinct) elements of Z is also called an *interface* (of size p) of $\mathcal{J}$.

Counting Problems Let $\mathcal{H} = (V, \mathbf{R})$ be a structure. Then $\#\mathrm{Hom}(\mathcal{H})$ is the problem that takes as input a structure $\mathcal{G}$ with the same signature as $\mathcal{H}$ and asks to determine $\hom(\mathcal{G} \to \mathcal{H})$.

Note that the power set 2^V is the set of all unary relations over V. Then $\#\mathrm{LHom}(\mathcal{H})$ is the problem $\#\mathrm{Hom}(\mathcal{H}^+)$, where $\mathcal{H}^+ = (V, \mathbf{R} \cup (2^V \setminus \{\emptyset\}))$. Intuitively, $\mathcal{H}^+$ is the structure obtained from $\mathcal{H}$ by adding all non-empty unary relations over the universe of $\mathcal{H}$. In the special case, where $\mathcal{H}$ is a graph H we simplify the notation by writing instances of $\#\mathrm{LHom}(H)$ as (G, L), where G is a graph and L is a function from the vertices of G to subsets of $V(H)$ specifying the unary relations of the input (lists). For convenience we assume that for such instances every vertex of G has a list. Then $\#\mathrm{LHom}(H)$ on input (G, L) asks to determine $\hom\big((G, L) \to H\big)$.

Pathwidth and pathwidth-preserving reductions Let A and B be computational problems that take as input some structure given along with a path decomposition. A *pathwidth-preserving reduction* from A to B is a polynomial-time Turing reduction from A to B such that the corresponding algorithm, if executed on an instance $\mathcal{G}$ of A given with a path decomposition with width at most t, makes B-oracle calls only on structures of size polynomial in $\|\mathcal{G}\|$, given with a path decomposition with width at most $t + \mathcal{O}(1)$.

Let $\mathcal{H} = (V, \mathbf{R})$ be a structure. A relation $R \subseteq V^{\mathrm{ar}(R)}$ is *realizable* (by $\mathcal{H}$) if there is a pathwidth-preserving reduction from $\#\mathrm{LHom}((V, \mathbf{R} \cup \{R\}))$ to $\#\mathrm{LHom}(\mathcal{H})$.

Note that pathwidth-preserving reductions are transitive in the sense that a pathwidth-preserving reduction from A to B together with a pathwidth-preserving reduction from B to C implies that there is a pathwidth-preserving reduction from A to C. The concept of realizable relations is transitive as well: If a relation R is realizable by a structure $\mathcal{H} = (V, \mathbf{R})$, and a relation R' is realizable by $\mathcal{H}' = (V, \mathbf{R} \cup \{R\})$, then R' is realizable by $\mathcal{H}$.

However, there is a caveat here: we are only allowed to combine a constant number of pathwidth-preserving reductions to make sure that the total increase of the pathwidth bound is constant as well.

We say that $\mathcal{H} = (V, \mathbf{R})$ is an *induced substructure* of $\mathcal{H}' = (V', \mathbf{R}')$ if $V \subseteq V'$ and for every R we have $\mathbf{x} \in R^{\mathcal{H}}$ if and only if $\mathbf{x} \in R^{\mathcal{H}'} \cap V^{\mathrm{ar}(R)}$. Note that, since we can use arbitrary lists, we can see each instance of $\#\mathrm{LHom}(\mathcal{H})$ as an instance of $\#\mathrm{LHom}(\mathcal{H}')$, where each element of

the input universe has a list that is a subset of V. This gives us a simple pathwidth-preserving reduction from #LHom($\mathcal{H}$) to #LHom($\mathcal{H}'$). Therefore, if a relation R is realizable by $\mathcal{H}$, it is also realizable by $\mathcal{H}'$. We will use this fact implicitly throughout this work.

Relations and Operators Let us now define some notation concerning relations.

- Let $R_1 \subseteq X_1 \times X_2$ and $R_2 \subseteq X_2 \times X_3$. The *composition* of R_1 and R_2 is the relation $R_1; R_2 \subseteq X_1 \times X_3$ defined as $\{(u, v) \mid \exists\, w \colon (u, w) \in R_1, (w, v) \in R_2\}$.

- Given a graph H and two sets $U, W \subseteq V(H)$, we define $R_{U \to W} = \{(u, v) \mid u \in U, v \in W \cap N(u)\}$.

- For $q \geq 2$, a relation $R \subseteq X_1 \times \cdots \times X_q$, and $v_1 \in X_1$, we define $R(v_1) = \{(v_2, \ldots, v_q) \in X_2 \times \cdots \times X_q \mid (v_1, v_2, \ldots, v_q) \in R\}$.

Note that for a graph H, sets $U, W \subseteq V(H)$, and $u \in U$, we have $R_{U \to W}(u) = \{v \in W \mid (u, v) \in R_{U \to W}\} = V \cap N(u)$.

The following lemma shows three simple ways to obtain realizable relations.

Lemma 2.1. *Let $\mathcal{H} = (V, \mathbf{R})$. The following relations R are realizable by $\mathcal{H}$:*

1. *$R_1 \cap R_2$, where $R_1, R_2 \in \mathbf{R}$ are of the same arity,*
2. *$R_{U \to U'}$, where $\mathbf{R}$ contains $E(H)$ for some graph H with vertex set V, and $U, U' \subseteq V$,*
3. *$R_1; R_2$, where $R_1, R_2 \in \mathbf{R}$ are of arity 2.*

The proof of Lemma 2.1 is postponed to the next subsection, where we introduce necessary tools to prove the third item.

2.1 Pathwidth-Preserving Reductions and Interpolation

One well-known tool for proving hardness for exact counting problems is interpolation.

Lemma 2.2 ([10, Lemma 3.2]). *For distinct non-zero integers $\mathbf{a} = (a_1, \ldots, a_k)$ and integers $\mathbf{b} = (b_1, \ldots b_k)$, suppose that, for each $j \in [k]$, we have $\sum_{i=1}^k a_i^j \cdot x_i = b_j$ for some unknown $\mathbf{x} = (x_1, \ldots x_k)$. Then the values of $\mathbf{x}$ can be computed in time polynomial in k.*

The following lemma is a modification of [10, Corollary 3.3] that generalizes to structures and makes the reduction pathwidth-preserving. Note that the conditions on the numbers are slightly different than the conditions on the weights in [10, Corollary 3.3].

Lemma 2.3. *Let $\mathcal{H} = (V, \mathbf{R})$ be a structure with signature σ and, for some positive integer p, let $R \subseteq V^p$ be a relation that is not in $\mathbf{R}$. Suppose that there is a structure $\mathcal{J}$ with signature σ and an interface $\mathbf{x}$ of size p such that, for each $\mathbf{f} \in R$ and $\mathbf{g} \in V^p \setminus R$, we have the following:*

1. $\mathrm{hom}\big((\mathcal{J}, \mathbf{x}) \to (\mathcal{H}, \mathbf{f})\big) \neq 0$.
2. $\mathrm{hom}\big((\mathcal{J}, \mathbf{x}) \to (\mathcal{H}, \mathbf{g})\big) \neq 1$.
3. *If $\mathrm{hom}\big((\mathcal{J}, \mathbf{x}) \to (\mathcal{H}, \mathbf{g})\big) \neq 0$, then there is no prime that divides both $\mathrm{hom}\big((\mathcal{J}, \mathbf{x}) \to (\mathcal{H}, \mathbf{f})\big)$ and $\mathrm{hom}\big((\mathcal{J}, \mathbf{x}) \to (\mathcal{H}, \mathbf{g})\big)$.*

Then R is realizable by $\mathcal{H}$.

Proof. In order to show that R is realizable by $\mathcal{H}$, we need to show, for $\mathcal{H}' = (V, \mathbf{R} \cup \{R\})$, that there is a pathwidth-preserving reduction from $\#\text{LHom}(\mathcal{H}')$ to $\#\text{LHom}(\mathcal{H})$.

Consider an input $\mathcal{G} = (U, \mathbf{S})$ of the $\#\text{LHom}(\mathcal{H}')$ problem, given along with a path decomposition with width t. Slightly overloading notation, we will use R to denote both the relation and its symbol so the corresponding relation in $\mathcal{G}$ is $R^{\mathcal{G}}$. We set $n = ||\mathcal{G}||$ and $m = |R^{\mathcal{G}}|$. Note that if $m = 0$, then we are done, as $\mathcal{G}$ can be cast as a structure with signature σ and it therefore can be interpreted as an instance of the $\#\text{LHom}(\mathcal{H})$ problem.

Let k be a positive integer. For each $(v_1, v_2, \ldots, v_p) \in R^{\mathcal{G}}$, we introduce k copies of the structure $\mathcal{J}$ with interface $\mathbf{x} = (x_1, \ldots, x_p)$. For each copy we identify $x_1, \ldots, x_p$ with $v_1, v_2, \ldots, v_p$, respectively. Then we remove the relation $R^{\mathcal{G}}$ from $\mathcal{G}$. Let us call the obtained structure $\mathcal{G}_{\mathcal{J}^k}$. Note that the signature of $\mathcal{G}_{\mathcal{J}^k}$ is σ.

Let Φ be the class of functions from U to V that respect all relations from σ (but not necessarily R). Let Φ^+ be the class of functions in Φ that respect R.

For each function $f \in \Phi$, define

$$w(f) = \prod_{\mathbf{v} \in R^{\mathcal{G}}} \text{hom}\big((\mathcal{J}, \mathbf{x}) \to (\mathcal{H}, f(\mathbf{v}))\big).$$

We also say that f is a $w(f)$-function. For each integer w, by $\Phi(w)$ we denote the set of w-functions in Φ.

Define

$$\begin{aligned}
\mathcal{W} &:= \{w(f) \mid f \in \Phi \text{ and } w(f) \neq 0\} \\
\mathcal{W}^+ &:= \{w(f) \mid f \in \Phi^+ \text{ and } w(f) \neq 0\} \\
\mathcal{W}^- &:= \{w(f) \mid f \in \Phi \setminus \Phi^+ \text{ and } w(f) \neq 0\}.
\end{aligned}$$

First note that for $f \in \Phi^+$, we have $w(f) \neq 0$ according to property 1 as given in the statement of the lemma. Thus, $\mathcal{W}^+ = \{w(f) \mid f \in \Phi^+\}$.

We now argue that $\mathcal{W}^+$ and $\mathcal{W}^-$ are disjoint. Let $w^+ \in \mathcal{W}^+$ and $w^- \in \mathcal{W}^-$. We have $w^- \neq 0$, and by properties 2 and 3 as given in the statement of the lemma, we observe that w^- has a prime factor that does not divide w^+. Therefore, $w^+ \neq w^-$. We can conclude that

$$\text{hom}\big(\mathcal{G} \to \mathcal{H}'\big) = \sum_{w \in \mathcal{W}^+} |\Phi(w)|. \tag{1}$$

Furthermore, we have

$$\begin{aligned}
\text{hom}\big(\mathcal{G}_{\mathcal{J}^k} \to \mathcal{H}\big) &= \sum_{f \in \Phi} w(f)^k \\
&= \sum_{w \in \mathcal{W}} \sum_{f \in \Phi(w)} w^k \\
&= \sum_{w \in \mathcal{W}} w^k \cdot |\Phi(w)|. \tag{2}
\end{aligned}$$

Let $\mathbf{a} = (w)_{w \in \mathcal{W}}$ and $\mathbf{b} = (\text{hom}\big(\mathcal{G}_{\mathcal{J}^k} \to \mathcal{H}\big))_{k \in [|\mathcal{W}|]}$. Note that a factor of the form $\text{hom}\big((\mathcal{J}, \mathbf{x}) \to (\mathcal{H}, f(\mathbf{v}))\big)$ can have at most $|V|^p$ different values, each of which can have multiplicity between 0 and m in the product $w(f)$. Thus, there are at most $(m+1)^{(|V|^p)}$ distinct values for $w(f)$ and the set $\mathcal{W}$ can be computed in time polynomial in n.

So, the tuple $\mathbf{a}$ can be computed in time polynomial in n. The tuple $\mathbf{b}$ can be computed using the algorithm that solves $\#\mathrm{LHom}(\mathcal{H})$ for each $k \in [|\mathcal{W}|]$ on input $\mathcal{G}_{\mathcal{J}^k}$ We have $|V(\mathcal{G}_{\mathcal{J}^k})| \leq |U| + |R^{\mathcal{G}}| \cdot |V(\mathcal{J})| \cdot k$.

Consider a tuple $(v_1, v_2, \ldots, v_p) \in R^{\mathcal{G}}$ and note that it corresponds to a clique in the Gaifman graph of $\mathcal{G}$, so in a path decomposition $\mathcal{X}$ of this graph there is a bag that contains $\{v_1, v_2, \ldots, v_p\}$. Let X' be the first such a bag. We modify $\mathcal{X}$ by inserting right after X' bags $X'_1, X'_2, \ldots, X'_k$, where X'_i is the union of X' and the vertex set of the i-th copy of $\mathcal{J}$ with interface $(v_1, v_2, \ldots, v_p)$. Clearly the obtained sequence $\mathcal{X}'$ is a path decomposition of $\mathcal{G}_{\mathcal{J}^k}$. As the size of $\mathcal{J}$ depends only on R and $\mathcal{H}$, and thus is a constant, we obtain that the width of $\mathcal{X}'$ is $t + \mathcal{O}(1)$.

Thus, $\mathbf{b}$ can be computed using $|\mathcal{W}|$ calls to the $\#\mathrm{LHom}(\mathcal{H})$ oracle, and we have shown that $|\mathcal{W}| \in n^{\mathcal{O}(1)}$ and each of the oracle calls is on an instance given with a path decomposition of width at most $t + \mathcal{O}(1)$.

Note that the values of $\mathbf{a}$ are non-zero by definition of $\mathcal{W}$. Using (2) and the fact that $\mathbf{a}$ contains non-zero pairwise distinct values, we can apply Lemma 2.2. Thus, we obtain $\mathbf{x} = (|\Phi(w)|)_{w \in \mathcal{W}}$ from $\mathbf{a}$ and $\mathbf{b}$ in time polynomial in n. By computing, for each $\mathbf{f} \in R$, the value of $\hom((\mathcal{J}, \mathbf{x}) \to (\mathcal{H}, \mathbf{f}))$, one can decide in polynomial time which values of $\mathcal{W}$ belong to $\mathcal{W}^+$ (recall that each $w^- \in \mathcal{W}^-$ has a prime factor that does not appear in any value of $\hom((\mathcal{J}, \mathbf{x}) \to (\mathcal{H}, \mathbf{f}))$ with $\mathbf{f} \in R$). Thus, using (1), the sought-for number of (list) homomorphisms can be computed from $\mathbf{x} = (|\Phi(w)|)_{w \in \mathcal{W}}$ in time polynomial in n. $\qquad\square$

The following simple corollary formalizes the modeling of relations by gadgets.

Corollary 2.4. *Let $\mathcal{H} = (V, \mathbf{R})$ be a structure with signature σ and, for some positive integer p, let $R \in V^p$ be a relation whose relation symbol is not in σ. Suppose that there is a structure $\mathcal{J}$ with signature σ and an interface $\mathbf{x}$ of size p such that, for each $\mathbf{f} \in R$ and $\mathbf{g} \in V^p \setminus R$, we have $\hom((\mathcal{J}, \mathbf{x}) \to (\mathcal{H}, \mathbf{f})) \neq 0$ and $\hom((\mathcal{J}, \mathbf{x}) \to (\mathcal{H}, \mathbf{g})) = 0$, then R is realizable by $\mathcal{H}$.*

Now we are ready to prove Lemma 2.1 that we recall here.

Lemma 2.1. *Let $\mathcal{H} = (V, \mathbf{R})$. The following relations R are realizable by $\mathcal{H}$:*

1. *$R_1 \cap R_2$, where $R_1, R_2 \in \mathbf{R}$ are of the same arity,*
2. *$R_{U \to U'}$, where $\mathbf{R}$ contains $E(H)$ for some graph H with vertex set V, and $U, U' \subseteq V$,*
3. *$R_1; R_2$, where $R_1, R_2 \in \mathbf{R}$ are of arity 2.*

Proof. For each case we will build a structure $\mathcal{J} = (U, \mathbf{S})$ with the same signature as $\mathcal{H}$ and interface $\mathbf{x}$, which satisfies the assumptions of Corollary 2.4.

1. Let $R = R_1 \cap R_2$ for some $R_1, R_2 \in \mathbf{R}$ with $p = \mathrm{ar}(R_1) = \mathrm{ar}(R_2)$. We set $U = \{x_1, \ldots, x_p\}$, $\mathbf{x} = (x_1, \ldots, x_p)$, and $R_1^{\mathcal{J}} = R_2^{\mathcal{J}} = \{(x_1, \ldots, x_p)\}$.

2. Suppose $\mathbf{R}$ contains $E(H)$ for some graph H with vertex set V, and $R = R_{U \to U'}$, where $U, U' \subseteq V$. We set $U = \{x, y\}$, $\mathbf{x} = (x, y)$, and $E(H)^{\mathcal{J}} = \{(x, y), (y, x)\}$. We also introduce lists $L(x) = U$, $L(y) = U'$.

3. Let $R = R_1; R_2$ for some $R_1, R_2 \in \mathbf{R}$ or arity 2. We set $U = \{x, y, z\}$, $\mathbf{x} = (x, z)$, $R_1^{\mathcal{J}} = \{(x, y)\}$, and $R_2^{\mathcal{J}} = \{(y, z)\}$.

Note that in the first two cases all vertices of $\mathcal{J}$ are in $\mathbf{x}$, so it it straightforward to observe that $\hom((\mathcal{J}, \mathbf{x}) \to (\mathcal{H}, \mathbf{f})) = 1$ if $\mathbf{x} \in R$ and $\hom((\mathcal{J}, \mathbf{x}) \to (\mathcal{H}, \mathbf{f})) = 0$ if $\mathbf{f} \notin R$. In the last case, we observe that $\hom((\mathcal{J}, \mathbf{x}) \to (\mathcal{H}, \mathbf{f})) > 0$ if $\mathbf{f} \in R$ and $\hom((\mathcal{J}, \mathbf{x}) \to (\mathcal{H}, \mathbf{f})) = 0$ otherwise. Thus the lemma follows from Corollary 2.4. $\qquad\square$

2.2 Hardness of counting satisfying assignments to $\mathrm{CSP}(q, r)$

Recall that the #SETH states that for every $\varepsilon > 0$, there is a $d \geq 1$ such that n-variable #d-SAT cannot be solved in time $\mathcal{O}((2 - \varepsilon)^n)$. Note that since d is a constant, this is equivalent to saying that n-variable m-clause #d-SAT cannot be solved in time $(2 - \varepsilon)^n \cdot (n + m)^{\mathcal{O}(1)}$.

For integers $r, q \geq 2$, by $\mathrm{CSP}(q, r)$ we denote the CSP problem with domain $[q]$ and each constraint of arity at most r. By #$\mathrm{CSP}(q, r)$ we denote the problem of counting satisfying valuations of a given instance of $\mathrm{CSP}(q, r)$. In this section we show the hardness of computing #$\mathrm{CSP}(q, r)$, conditioned on the #SETH.

The decision version of this result, conditioned on the SETH, was proven by Lampis [21]. Our proof is just an adaptation of the original one to the counting world, so we will not elaborate on the details and refer the reader to the paper of Lampis [21].

Theorem 2.5. *For every integer $q \geq 2$ and $\varepsilon > 0$ there is an integer r, such that the following holds. Unless the #SETH fails, #$\mathrm{CSP}(q, r)$ with n variables and m constraints cannot be solved in time $(q - \varepsilon)^n \cdot (n + m)^{\mathcal{O}(1)}$.*

Proof. Fix $q \geq 2$ and $\varepsilon > 0$ and suppose that for every r there is an algorithm solving n-variable m-constraint #$\mathrm{CSP}(q, r)$ in time $(q - \varepsilon)^n \cdot (n + m)^{\mathcal{O}(1)}$. Suppose we are given a #d-SAT instance Φ with N variables and M clauses, where d is a constant. Without loss of generality we may assume that each variable is involved in some clause.

In the reduction we will carefully choose $r = r(d, q, \varepsilon)$ and build a #$\mathrm{CSP}(q, r)$ instance Ψ. To avoid confusion, we will refer to an assignment of values to the variables of Φ as *truth assignment*, while an assignment of values to the variables of Ψ will be called *valuation*.

First, select an integer p and a real $\delta > 0$, such that there exists an integer t satisfying the following (see [21] for the argument that such a choice is possible):

$$(q - \varepsilon)^p \leq (2 - \delta)^t \leq 2^t \leq q^p.$$

Note that δ does not depend on d. Now group the variables of Φ into $\gamma := \lceil N/t \rceil$ groups, each with at most t variables. Call these groups $V_1, V_2 \ldots, V_\gamma$. For each $i \in [\gamma]$, we create a group X_i of p variables of Ψ. Thus the total number of variables in Ψ is $n := \gamma \cdot p \leq \frac{pN}{t} + p$.

Note that the total number of truth assignments of variables in V_i is at most $2^t \leq q^p$, so it does not exceed the number of possible valuations of X_i. Let us fix some injective function that maps each truth assignment on V_i to a distinct valuation of X_i.

Now, let us define the constraints of Ψ. Consider any clause c of Φ and let $V_{i_1}, V_{i_2}, \ldots, V_{i_{s'}}$ for $d' \leq d$ be the groups that contain the variables of c. Note that each truth assignment of variables in $V_{i_1}, V_{i_2}, \ldots, V_{i_{d'}}$ that satisfies c corresponds to some valuation of variables in $X_{i_1}, X_{i_2}, \ldots, X_{i_{d'}}$. We introduce a new constraint $C(c)$ on all variables from $X_{i_1} \cup X_{i_2} \cup \ldots \cup X_{i_{d'}}$ that accepts only the valuations that correspond to the truth assignments that satisfy c. The arity of $C(c)$ is $|X_{i_1} \cup X_{i_2} \cup \ldots \cup X_{i_{d'}}| \leq d \cdot p$.

Summing up, Ψ has $n = \gamma \cdot p \leq \frac{pN}{t} + p$ variables, $m = M$ constraints, its domain is $[q]$, and each constraint has arity at most $r := dp$.

It is straightforward to observe that Φ has a satisfying truth assignment if and only if Ψ has a satisfying valuation. However, a stronger property holds as well: there is a bijection between truth assignments that satisfy Φ and valuations that satisfy Ψ. Indeed, recall that every variable of Φ appears in some clause, and thus, for each $i \in [\gamma]$, there is a constraint of Ψ involving all variables of X_i. Since the constraints of Ψ were defined in a way that the

only accepted valuations are in one-to-one correspondence to the truth assignments of Φ, we observe that the claimed bijection between truth assignments and valuations exists. Hence, the solution to $\#\mathrm{CSP}(q, r)$ on the instance Ψ is precisely the number of satisfying assignments of Φ, i.e., the solution of our instance of $\#d\text{-SAT}$. Let us call our hypothetical algorithm for Ψ. Its running time is:

$$
\begin{aligned}
(q - \varepsilon)^n \cdot (n + m)^{\mathcal{O}(1)} &\leq ((q - \varepsilon)^p)^{N/t+1} \cdot (N + M)^{\mathcal{O}(1)} \\
&\leq \left((2 - \delta)^t\right)^{N/t+1} \cdot (N + M)^{\mathcal{O}(1)} \\
&\leq 2^t \cdot (2 - \delta)^N \cdot (N + M)^{\mathcal{O}(1)}.
\end{aligned}
$$

As t depends only on q and ε, this running time is $(2 - \delta)^N \cdot (N + M)^{\mathcal{O}(1)}$, which violates the $\#\mathrm{SETH}$. $\qquad\square$

3 Counting list homomorphisms to bipartite graphs H

Let H be a bipartite graph. We say that a set $S \subseteq V(H)$ is *irredundant* if for all distinct $u, v \in S$ it holds that $N(u) \neq N(v)$. We say that a graph H is *irredundant* if $V(H)$ is irredundant. If H is connected, then $S \subseteq V(H)$ is *one-sided* if it is contained in one bipartition class. Then recall from Definition 1.1 that $\mathrm{irr}(H)$ is the maximum size of a one-sided irredundant subset of $V(H)$. We extend this to disconnected bipartite graphs by setting $\mathrm{irr}(H)$ to be the maximum value of $\mathrm{irr}(H')$ over all connected components H' of H. Observe that the following conditions are equivalent for every bipartite graph H:

  (i) $\mathrm{irr}(H) \geq 2$,

 (ii) H has a connected component that is not a biclique,

(iii) H contains an induced P_4.

For a connected bipartite graph H, an instance (G, L) of $\#\mathrm{LHom}(H)$ is *consistent*, if:

 • G is connected and bipartite, let X and Y be its bipartition classes,

 • $\bigcup_{x \in X} L(x)$ is contained in one bipartition class of H, and $\bigcup_{y \in Y} L(y)$ is contained in the other bipartition class of H.

3.1 Algorithm

Theorem 3.1. *For each bipartite graph H, the $\#\mathrm{LHom}(H)$ problem on n-vertex instances given along with a tree decomposition of width at most t can be solved in time $\mathrm{irr}(H)^t \cdot n^{\mathcal{O}(1)}$.*

Proof. Let (G, L) be an instance of $\#\mathrm{LHom}(H)$, where G has n vertices and is given along with a tree decomposition of width at most t. First observe that if G is not bipartite, then there is no homomorphism from G to H, thus we return 0. So from now on assume that G is bipartite.

First, assume that G and H are both connected. Let X, Y be the bipartition of G and A, B be the bipartition of H. We observe that in every homomorphism from G to H, either each vertex of X is mapped to a vertex of A and each vertex of Y is mapped to a vertex of B,

or each vertex of X is mapped to a vertex of B and each vertex of Y is mapped to a vertex of A.

Thus we can reduce the problem to solving two consistent instances of #LHom(H) as follows. Let L_1 be the lists obtained from L by setting $L_1(x) := L(x) \cap A$ for every $x \in X$ and $L_1(y) := L(y) \cap B$ for every $y \in Y$. Similarly, define L_2 as follows: $L_2(x) := L(x) \cap B$ for every $x \in X$ and $L_2(y) := L(y) \cap A$ for every $y \in Y$. By the above reasoning, we obtain that $\text{hom}\big((G, L) \to H\big) = \text{hom}\big((G, L_1) \to H\big) + \text{hom}\big((G, L_2) \to H\big)$.

So let us focus on computing $\text{hom}\big((G, L_1) \to H\big)$ as computing $\text{hom}\big((G, L_2) \to H\big)$ is analogous. We define lists L_1' and an auxiliary weight function $\mathfrak{w} \colon V(G) \times V(H) \to \mathbb{N}$ as follows. For each $v \in V(G)$, we partition the vertices of $L_1(v)$ into subsets consisting of vertices with exactly the same neighborhood. From each such subset W we include in $L_1'(v)$ exactly one vertex, say u, and set $\mathfrak{w}(v, u) = |W|$. On the other hand, for every $u \notin L_1'(v)$ we set $\mathfrak{w}(v, u) = 0$.

Observe that for each $v \in V(G)$ the set $L_1'(v)$ is irredundant and contained in one bipartition class of H. Thus, for each $v \in V(G)$, we have $|L_1'(v)| \leq \text{irr}(H)$.

Furthermore, denoting by Ω the set of all homomorphisms from (G, L_1') to H, we observe that:

$$\text{hom}\big((G, L_1) \to H\big) = \sum_{f \in \Omega} \prod_{v \in V(G)} \mathfrak{w}(v, f(v)). \tag{3}$$

Since every list in L_1' has size at most $\text{irr}(H)$, a standard bottom-up dynamic programming approach can be used to compute $\text{hom}\big((G, L_1') \to H\big)$ in time $\text{irr}(H)^t \cdot n^{\mathcal{O}(1)}$ [?]. Furthermore one can easily modify the algorithm to actually determine the sum in (3): whenever we decide to assign color u to some vertex $v \in V(G)$, we multiply the number of solutions by $\mathfrak{w}(w, u)$. This modification does not increase the running time.

Now consider the general case that graphs G and H are possibly disconnected. Let $G_1, G_2, \ldots, G_p$ be the connected components of G and let $H_1, H_2, \ldots, H_q$ be the connected components of H. It is straightforward to observe that

$$\text{hom}\big((G, L) \to H\big) = \prod_{i=1}^{p} \left(\sum_{j=1}^{q} \text{hom}\big((G_i, L) \to H_j\big) \right).$$

Thus, the total running time of such an algorithm is $\text{irr}(H)^t \cdot n^{\mathcal{O}(1)}$. $\qquad \square$

3.2 Hardness for bipartite target graphs

The following lemma is the main building block of our hardness reduction.

Lemma 3.2. *Let $H = (V, E)$ be a connected bipartite graph with $\text{irr}(H) \geq 2$, and let $S \subseteq V$ be a one-sided irredundant set. For every fixed $p \geq 1$, every relation $R \subseteq S^p$ is realizable by H.*

We postpone the proof of Lemma 3.2 to Section 3.6, and now we show how it implies the lower bounds for bipartite graphs H.

Theorem 3.3. *Let H be a bipartite graph with $\text{irr}(H) \geq 2$. Assuming the #SETH, there is no $\varepsilon > 0$, such that #LHom(H) on consistent n-vertex instances given with a path decomposition of width t can be solved in time $(\text{irr}(H) - \varepsilon)^t \cdot n^{\mathcal{O}(1)}$.*

Proof. Let H be as in the assumptions and let S be a maximum-size irredundant set contained in a bipartition class of some connected component H' of H. Let $q := |S| = \text{irr}(H)$. Note that H' and S satisfy the assumptions of Lemma 3.2. Suppose that there is $\varepsilon > 0$ and an algorithm that solves every consistent n-vertex m-edge instance (G, L) of $\#\text{LHom}(H)$ that is given along with a path decomposition of G with width t in time $(q - \varepsilon)^t \cdot (n + m)^{\mathcal{O}(1)}$.

We reduce from $\#\text{CSP}(q, r)$, where r is the constant given for q and ε by Theorem 2.5. Consider an instance Ψ with variables U, domain $D = [q]$, and let $\mathcal{R}$ be the set of relations used in the constraints of Ψ. Note that $|\mathcal{R}|$ depends only on q and r, i.e., $|\mathcal{R}|$ is a constant. Furthermore, the number of constraints in Ψ is polynomial in $|U|$.

Recall that $|S| = q$, so by fixing an arbitrary bijection between S and $[q]$, we can equivalently see Ψ as an instance with domain S. In other words, the instance Ψ can be equivalently seen as a structure, which is an instance of $\#\text{LHom}((S, \mathcal{R}))$. Note that the pathwidth of Ψ is clearly at most $|U|$.

As the arity of each relation in $\mathcal{R}$ is at most r, by Lemma 3.2, all relations in $\mathcal{R}$ are realizable by H' and thus by H. This means that there is a pathwidth-preserving reduction from $\#\text{LHom}((S, \mathcal{R}))$ to $\#\text{LHom}(H)$.

Let us point out that our reduction is in fact a chain of pathwidth-preserving reductions. However, the total number of steps in this chain is $\mathcal{O}(|\mathcal{R}|)$, which is a constant. Thus the total increase of the pathwidth is $\mathcal{O}(1)$.

So, using our hypothetical algorithm for $\#\text{LHom}(H)$, we can count satisfying assignments for Ψ in time

$$(q - \varepsilon)^{|U| + \mathcal{O}(1)} \cdot |U|^{\mathcal{O}(1)} = (q - \varepsilon)^{|U|} \cdot |U|^{\mathcal{O}(1)},$$

where we use the fact that q is a constant. By Theorem 2.5 the existence of such an algorithm for $\#\text{CSP}(q, r)$ contradicts the $\#\text{SETH}$. $\square$

3.3 Special case: $H = P_4$

As a warm-up, let us start with the special case that $H = P_4$. The following lemma can be seen as a restriction of Lemma 3.2 to this case. The proof will illustrate our approach.

Lemma 3.4. *Consider $P_4 = (a - b - c - d)$ and fix a positive integer q. Then any $R \subseteq \{a, c\}^q$ is realizable by P_4.*

Proof. We first show that the relations $\text{NEQ} := \{(a, c), (c, a)\}$ and $\text{OR}_q := \{a, c\}^q \setminus \{c^q\}$ are realizable by P_4. We will then use these relations to show the statement of the lemma.

First, let us focus on NEQ. We aim to use Lemma 2.3. Let $\mathcal{J}^{\text{NEQ}}$ be a five-vertex path $(u - w_1 - w_2 - w_3 - v)$. The interface of $\mathcal{J}^{\text{NEQ}}$ is $\mathbf{x} = (u, v)$. The lists are $L(u) = L(v) = L(w_2) = \{a, c\}$, and $L(w_1) = L(w_3) = \{b, d\}$.

Let us analyze the values of $\hom\big((\mathcal{J}^{\text{NEQ}}, \mathbf{x}) \to (P_4, \mathbf{f})\big)$ for distinct $\mathbf{f} \in V(P_4)^2$:

$$\hom\big((\mathcal{J}^{\text{NEQ}}, \mathbf{x}) \to (P_4, (a, a))\big) = 2,$$
$$\hom\big((\mathcal{J}^{\text{NEQ}}, \mathbf{x}) \to (P_4, (a, c))\big) = 3,$$
$$\hom\big((\mathcal{J}^{\text{NEQ}}, \mathbf{x}) \to (P_4, (c, a))\big) = 3,$$
$$\hom\big((\mathcal{J}^{\text{NEQ}}, \mathbf{x}) \to (P_4, (c, c))\big) = 5.$$
$$\hom\big((\mathcal{J}^{\text{NEQ}}, \mathbf{x}) \to (P_4, \mathbf{f})\big) = 0, \text{ otherwise.}$$

17

So, for each $\mathbf{f} \in \text{NEQ}$ and $\mathbf{g} \in V^2 \setminus \text{NEQ}$, we have $\hom\big((\mathcal{J}, (s,t)) \to (\mathcal{H}', \mathbf{f})\big) = 3$ and $\hom\big((\mathcal{J}, (s,t)) \to (\mathcal{H}', \mathbf{g})\big) \in \{0, 2, 5\}$. Thus, $\mathcal{J}^{\text{NEQ}}$ satisfies the assumptions of Lemma 2.3 and we obtain that NEQ is realizable by P_4.

Now let us consider OR_q. Define $\mathcal{J}^{\text{OR}_q}$ as follows. We introduce q vertices $v_1, v_2, \ldots, v_q$ and one additional vertex w, adjacent to all v_i's. We set $L(v_i) = \{a, c\}$ for all $i \in [q]$, and $L(w) = \{b, d\}$. The interface of $\mathcal{J}^{\text{OR}_q}$ is $\mathbf{x} = (v_1, v_2, \ldots, v_q)$.

For $\mathbf{f} \in V(P_4)^q$ we have

$$\hom\big((\mathcal{J}^{\text{OR}_q}, \mathbf{x}) \to (P_4, \mathbf{f})\big) = 0 \text{ if } \mathbf{f} \notin \{a, c\}^q,$$
$$\hom\big((\mathcal{J}^{\text{OR}_q}, \mathbf{x}) \to (P_4, \mathbf{f})\big) = 1 \text{ if } \mathbf{f} \in \{a, c\}^q \setminus c^q,$$
$$\hom\big((\mathcal{J}^{\text{OR}_q}, \mathbf{x}) \to (P_4, \mathbf{f})\big) = 2 \text{ if } \mathbf{f} = c^q.$$

Again, $\mathcal{J}^{\text{OR}_q}$ satisfies the assumptions of Lemma 2.3 and we obtain that OR_q is realizable by P_4.

Finally, consider an arbitrary relation $R \subseteq \{a, c\}^q$. Let $\{\mathbf{f}_1, \mathbf{f}_2, \ldots, \mathbf{f}_p\} = \{a, c\}^q \setminus R$, and for each $i \in [p]$ let $R_i := \{a, c\}^q \setminus \{\mathbf{f}_i\}$. Clearly $R = \bigcap_{i=1}^p R_i$.

Thus, by Lemma 2.1, it is sufficient to show that each R_i is realizable. Fix some i and let J be the set of the indices $j \in [q]$, such that the j-th coordinate of $\mathbf{f}_i$ is a (and the other ones are c). If $|J| = \emptyset$, then $R_i = \text{OR}_q$ and we are done. So suppose that $|J| \geq 1$ and we can realize all relations $R' = \{a, c\}^q \setminus \{\mathbf{f}\}$, where the number of coordinates of $\mathbf{f}$ that are equal to a is smaller than $|J|$.

So let us choose some $j \in J$. For each tuple $\mathbf{f} \in \{a, c\}^q$ let $\mathbf{f}'$ be the tuple in $\{a, c\}^q$ obtained from $\mathbf{f}$ by changing the j-th coordinate from a to c or from c to a, whichever applies. Consider $R_i' = \{a, c\}^q \setminus \{\mathbf{f}_i'\}$. Since $j \in J$, $\mathbf{f}_i'$ is obtained from $\mathbf{f}_i$ by swapping the j-th coordinate from a to c. By the inductive assumption, R_i' is realizable.

Note that $R_i = \{\mathbf{f} \mid \mathbf{f}' \in R_i'\}$. We will use Corollary 2.4 to show that R_i is realizable by the structure $\mathcal{H} = (V(P_4), \{E(P_4), R_i', \text{NEQ}\})$, which implies that R_i is realizable by $P_4 = (V(P_4), E(P_4))$. Slightly abusing notation, we use E, R_i', and NEQ also as corresponding relation symbols in the signature of $\mathcal{H}$. We define a gadget $\mathcal{J}$ on $q+1$ vertices $\{v_1, \ldots, v_q, u\}$ with interface $\mathbf{x} = (v_1, \ldots, v_q)$. We apply R_i' to the tuple $(v_1, \ldots v_{j-1}, u, v_{j+1}, \ldots v_q)$ and we apply NEQ to (u, v_j), i.e. $R_i'^{\mathcal{J}} = \{(v_1, \ldots v_{j-1}, u, v_{j+1}, \ldots v_q)\}$, $\text{NEQ}^{\mathcal{J}} = \{(u, v_j)\}$, and $E^{\mathcal{J}} = \emptyset$. Clearly, $\mathcal{J}$ has the same signature as $\mathcal{H}$. Moreover, $\hom\big((\mathcal{J}, \mathbf{x}) \to (\mathcal{H}, \mathbf{f})\big) \neq 0$ iff $\mathbf{f} \in R_i$. $\qquad\square$

Note that Lemma 3.4 together with the proof of Theorem 3.3 already yield the following result.

Corollary 3.5. *Assuming the #SETH, there is no $\varepsilon > 0$, such that #LHom(P_4) on n-vertex instances given with a path decomposition of width t can be solved in time $(2 - \varepsilon)^t \cdot n^{\mathcal{O}(1)}$.*

3.4 Structure of P_4s in H

Let H be a bipartite graph with bipartition (X, Y). We define

$$\mathcal{P}_3(H) := \{S \subseteq V(H) \mid H[S] \simeq P_3\},$$
$$\mathcal{P}_4(H) := \{S \subseteq V(H) \mid H[S] \simeq P_4\}.$$

By $\mathsf{Adj}_{P_4}(H)$ we denote the following graph:

$$V(\mathsf{Adj}_{P_4}(H)) := \mathcal{P}_4(H),$$
$$E(\mathsf{Adj}_{P_4}(H)) := \left\{ SS' \in \mathcal{P}_4(H)^2 \mid (|(S \cap S') \cap X| = 2 \vee |(S \cap S') \cap Y| = 2) \right\}.$$

Informally speaking, we think of two induced four-vertex paths as adjacent if they share two vertices from one bipartition class. Let us point out that in the definition of $E(\mathsf{Adj}_{P_4}(H))$ we do not insist that $S \neq S'$. Therefore the graph $\mathsf{Adj}_{P_4}(H)$ is reflexive, i.e., every vertex has a loop. This is a technical detail that allows us to simplify some arguments.

Lemma 3.6. *Let H be a connected irredundant bipartite graph. Then $\mathsf{Adj}_{P_4}(H)$ is connected.*

Proof. Note that if $|V(H)| \leq 3$, then $\mathsf{Adj}_{P_4}(H) = \emptyset$ and we are done. Thus, suppose that H has at least four vertices. Let $C_1, C_2, \ldots, C_p$ be the connected components of $\mathsf{Adj}_{P_4}(H)$. Let $c \colon \mathcal{P}_4(H) \to [p]$ be a function such that $c(S) = i$ if and only if S belongs to C_i.

First, let us observe the following.

Claim 3.6.1. *Every edge and every induced P_3 in H is contained in some induced P_4.*

Proof of Claim. First, let us argue that every induced P_3 is contained in some induced P_4. Let $(x - y - z)$ be an induced P_3. Since the vertices x and z do not have the same neighborhood, one of them, say x, has a neighbor w that is not adjacent to z. Then $(w - x - y - z)$ is an induced P_4. Now consider an edge xy. Since H is connected and has at least three vertices, one of the endvertices of xy, say y, has a neighbor z. Since H is bipartite, $(x - y - z)$ is an induced P_3. As we have already shown, it is contained in some induced P_4. ◁

Note that if $A \in \mathcal{P}_3(H)$, then the set $\mathcal{S}_A := \{S \in \mathcal{P}_4(H) \mid A \subseteq S\}$ is non-empty (by Claim 3.6.1) and induces a clique in $\mathsf{Adj}_{P_4}(H)$. In particular, for any $S, S' \in \mathcal{S}_A$ it holds that $c(S) = c(S')$. We introduce a labeling $\ell \colon \mathcal{P}_3(H) \to [p]$, where $\ell(A) = i$ if and only if $c(\mathcal{S}_A) = \{i\}$.

Claim 3.6.2. *For any edge $xy \in E(H)$ and any $A, B \in \mathcal{P}_3(H)$ such that $\{x, y\} \subseteq A \cap B$, it holds that $\ell(A) = \ell(B)$. Consequently, all sets $S \in \mathcal{P}_4(H)$ that contain xy belong to the same connected component of $\mathsf{Adj}_{P_4}(H)$.*

Proof of Claim. Let $A = \{x, y, a\}$ and $B = \{x, y, b\}$, where $a \neq b$. If $\{x, y, a, b\} \in \mathcal{P}_4(H)$, then $\{x, y, a, b\} \in \mathcal{S}_A \cap \mathcal{S}_B$ and thus $\ell(A) = \ell(B) = c(\{x, y, a, b\})$.

So suppose that $\{x, y, a, b\} \notin \mathcal{P}_4(H)$. This is possible in two cases: (1) if a and b are adjacent to the same vertex from $\{x, y\}$ (and thus $H[\{x, y, a, b\}]$ is a star with 3 leaves), or (2) if $ab \in E(H)$ (and thus $H[\{x, y, a, b\}]$ is a 4-cycle).

By Claim 3.6.1, each of the sets A, B belongs to some induced P_4, i.e., there are vertices c, d, such that $\{x, y, a, c\}, \{x, y, b, d\} \in \mathcal{P}_4(H)$. Note that it is possible that $c = d$.

We will show that there is a walk from $\{x, y, a, c\}$ to $\{x, y, b, d\}$ in $\mathsf{Adj}_{P_4}(H)$, which will prove that $c(\{x, y, a, c\}) = c(\{x, y, b, d\})$ and consequently $\ell(A) = \ell(B)$. We consider some cases.

Case 1. Suppose that $H[\{x, y, a, b\}]$ is a star with 3 leaves. By symmetry we assume that $a, b, x \in N(y)$. We consider possible positions of c and d. One of the following cases occurs, see also Figure 3.
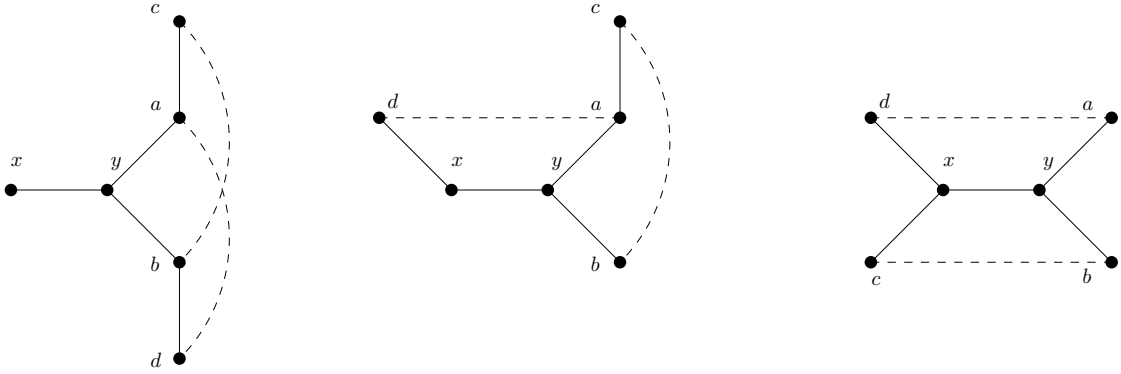
Figure 3: Possible configurations in the proof of Case 1 in Claim 3.6.2. Dashed edges may, but do not have to exist. Note that the first and last subcase vertices c and d may coincide.

Subcase 1.1. $N(x) \cap \{c, d\} = 0$. If none of the edges cb, ad exists in H, then we have a walk $\big(\{x, y, a, c\} - \{c, a, y, b\} - \{a, y, b, d\} - \{x, y, b, d\}\big)$ in $\mathsf{Adj}_{P_4}(H)$. If at least one of these edges, say cb, exists in H, then we have a walk $\big(\{x, y, a, c\} - \{x, y, b, c\} - \{x, y, b, d\}\big)$ in $\mathsf{Adj}_{P_4}(H)$.

Subcase 1.2. $N(x) \cap \{c, d\} = 1$. By symmetry, assume that $dx \in E(H)$. If $ad \notin E(H)$, then we have a walk $\big(\{x, y, a, c\} - \{d, x, y, a\} - \{d, x, y, b\}\big)$ in $\mathsf{Adj}_{P_4}(H)$. If $cb \in E(H)$, then we have a walk $\big(\{x, y, a, c\} - \{x, y, b, c\} - \{d, x, y, b\}\big)$ in $\mathsf{Adj}_{P_4}(H)$. If $cb \notin E(H)$ and $ad \in E(H)$, then we have a walk $\big(\{x, y, a, c\} - \{c, a, y, b\} - \{d, a, y, b\} - \{d, x, y, b\}\big)$ in $\mathsf{Adj}_{P_4}(H)$.

Subcase 1.3. $N(x) \cap \{c, d\} = 2$. If one of the edges cb, ad, say ad, does not exist in H, then we have a walk $\big(\{c, x, y, a\} - \{d, x, y, a\} - \{d, x, y, b\}\big)$ in $\mathsf{Adj}_{P_4}(H)$. If both $cb, ad \in E(H)$, then we have a walk $\big(\{c, x, y, a\} - \{c, x, d, a\} - \{d, x, c, b\} - \{d, x, y, b\}\big)$ in $\mathsf{Adj}_{P_4}(H)$.

Case 2. Suppose that $H[\{x, y, a, b\}]$ is a 4-cycle. By symmetry we assume that consecutive vertices of the cycle are a, x, y, b. Note that c is adjacent to exactly one of a, y, and d is adjacent to exactly one of x, b, see also Figure 4.

Subcase 2.1. $dx, cy \in E(H)$. If $cd \notin E(H)$, then we have a walk $\big(\{a, x, y, c\} - \{d, x, y, c\} - \{d, x, y, b\}\big)$ in $\mathsf{Adj}_{P_4}(H)$. If $cd \in E(H)$, then we have a walk $\big(\{a, x, y, c\} - \{a, x, d, c\} - \{b, a, x, d\} - \{b, y, x, d\}\big)$ in $\mathsf{Adj}_{P_4}(H)$.

Subcase 2.2. Exactly one of the edges xd, yc, say xd exists in H.
If $cd \notin E(H)$, then we have a walk $\big(\{c, a, x, y\} - \{c, a, x, d\} - \{b, a, x, d\}\big)$ in $\mathsf{Adj}_{P_4}(H)$. If $cd \in E(H)$, then we have a walk $\big(\{c, a, x, y\} - \{c, d, x, y\} - \{d, x, y, b\}\big)$ in $\mathsf{Adj}_{P_4}(H)$.

Subcase 2.3. $ac, bd \in E(H)$. If $cd \notin E(H)$, then we have a walk $\big(\{c, a, x, y\} - \{c, a, b, y\} - \{c, a, b, d\} - \{x, a, b, d\} - \{x, y, b, d\}\big)$ in $\mathsf{Adj}_{P_4}(H)$. If $cd \in E(H)$, then we have a walk $\big(\{c, a, x, y\} - \{x, a, c, d\} - \{x, a, b, d\} - \{x, y, b, d\}\big)$ in $\mathsf{Adj}_{P_4}(H)$.

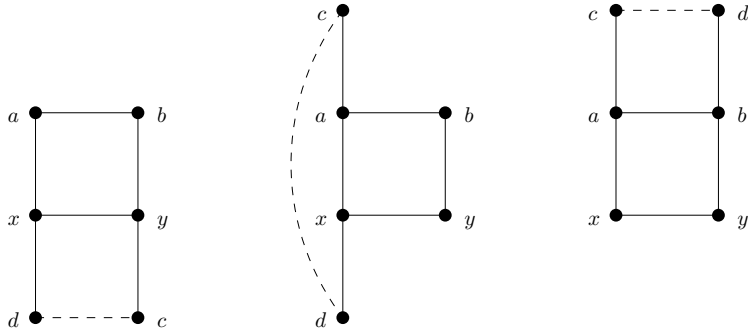As these are all possible cases, the proof of the claim is complete. ◁

20

Figure 4: Possible configurations in the proof of Case 2 in Claim 3.6.2. Dashed edges may, but do not have to exist.

Claim 3.6.2 allows us to define a labeling ℓ' of edges of H, analogous to ℓ: for $xy \in E(H)$, we set $\ell'(xy) = c(S)$, where S is any element of $\mathcal{P}_4(H)$ with $\{x, y\} \in S$. Note that by Claim 3.6.1, such a set S always exists and by Claim 3.6.2, the value of $\ell'(xy)$ does not depend on the choice of S. Note that for every induced path $(x-y-z)$ we have $\ell(\{x, y, z\}) = \ell'(xy) = \ell'(yz)$.

Now we claim that for any two edges $e, f \in E(H)$ we have $\ell'(e) = \ell'(f)$. For contradiction, suppose this is not the case. Since H is connected, this means that there are two edges xy and xz, where $y \neq z$, such that $\ell'(xy) \neq \ell'(xz)$. However, this cannot happen as $\ell'(xy) = \ell(\{x, y, z\}) = \ell'(yz)$, since $\{x, y, z\} \in \mathcal{P}_3(H)$. $\qquad\square$

3.5 Constructing (x, y)-Forcers

Definition 3.7 $((x, y, S)$-*distinguisher*, (x, y, S)-*forcer*)**.** Let $H = (V, E)$ be a bipartite graph. Let S be a one-sided subset of V with distinct vertices $x, y \in S$. Let $\{\alpha, \beta\}$ be a one-sided subset of V, and let R be a relation in $S \times \{\alpha, \beta\}$.

1. R is an (x, y, S)-*distinguisher with respect to* (α, β) if it has the following properties:

   - $(x, \alpha) \in R$, $(x, \beta) \notin R$.
   - $(y, \beta) \in R$.
   - $\forall v \in S : R \cap \{(v, \alpha), (v, \beta)\} \neq \emptyset$.

2. R is an (x, y, S)-*forcer with respect to* (α, β) if it is an (x, y, S)-distinguisher with respect to (α, β) with the additional property:

   - $(y, \alpha) \notin R$.

We say that a relation R is an (x, y, S)-*distinguisher* (resp. (x, y, S)-*forcer*) if there are is a one-sided set $\{\alpha, \beta\}$, such that R is an (x, y, S)-distinguisher (resp. (x, y, S)-forcer) with respect to (α, β).

The following lemma is a crucial building block that we will use repeatedly in the results leading up to the proof of Lemma 3.2. Here, given a realizable (x, y, S)-distinguisher with respect to some pair of vertices on a four-vertex path, we use Lemma 2.3 to turn this distinguisher into a forcer.

Lemma 3.8. *Let $H = (V, E)$ be an irredundant connected bipartite graph. Let S be a one-sided subset of V with distinct vertices $x, y \in S$. Let R be a relation that is realizable by H such that, for some induced 4-vertex path $(a - b - c - d)$ in H, R is an (x, y, S)-distinguisher with respect to (a, c) or with respect to (c, a). Then the following relations are realizable by H.*

1. *an (x, y, S)-forcer with respect to (a, c),*
2. *an (x, y, S)-forcer with respect to (c, a),*
3. *an (x, y, S)-forcer with respect to (b, d).*
4. *an (x, y, S)-forcer with respect to (d, b),*

Proof. First note that, since R is realizable by H, there is a pathwidth-preserving reduction from $\#\mathrm{LHom}((V, E \cup \{R\}))$ to $\#\mathrm{LHom}(H)$. So it suffices to show that the relations in items 1-4 are realizable by $\mathcal{H}' = (V, E \cup \{R\})$.

Recall that by Lemma 3.4, the relations $\mathrm{NEQ}(\{a, c\}) = \{(a, c), (c, a)\}$ and $\mathrm{NEQ}(\{b, d\}) = \{(b, d), (d, b)\}$ are realizable by P_4 and thus by H. Note that if R is an (x, y, S)-distinguisher with respect to (c, a), then $R;\mathrm{NEQ}(\{a, c\})$ is an (x, y, S)-distinguisher with respect to (a, c). Furthermore, by Lemma 2.1, this relation is realizable by $\mathcal{H}'$. So from now on we can assume that R is an (x, y, S)-distinguisher with respect to (a, c), realizable by $\mathcal{H}'$.

Let us first focus on proving item 1.

Claim 3.8.1. *Let R is an (x, y, S)-distinguisher with respect to (a, c), realizable by $\mathcal{H}'$. Then there is an (x, y, S)-forcer R' with respect to (a, c), realizable by $\mathcal{H}'$.*

Proof of Claim. Note that if $(y, a) \notin R$ then we can choose $R' = R$ and are done. So we can assume that $(y, a) \in R$.

Slightly overloading notation, we use E and R also as the relation symbols in $\sigma(\mathcal{H}')$ corresponding to the relations E and R, respectively. We now define a gadget $\mathcal{J} = (U, \mathbf{S})$ with two interface vertices s, t that has the same signature as $\mathcal{H}'$. The vertices of $\mathcal{J}$ are $U = \{s, t, t', u_1, u_2, u_3\}$, the edge relation is $E^{\mathcal{J}} = E((t - u_1 - t' - u_2 - u_3))$, and $R^{\mathcal{J}} = \{(s, t), (s, t')\}$. As lists we set $L(u_1) = L(u_2) = \{b, d\}$ and $L(u_2) = \{a, c\}$. Intuitively, $\mathcal{J}$ is a path $(t - u_1 - t' - u_2 - u_3)$, where in addition s, t and s, t' are subject to the relation R. The gadget $\mathcal{J}$ is illustrated in the following figure:
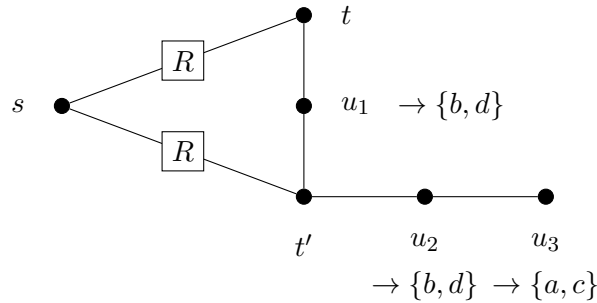
Consider the following subsets of S.

- $S_a = \{v \in S \mid (v, a) \in R, (v, c) \notin R\}$,
- $S_c = \{v \in S \mid (v, a) \notin R, (v, c) \in R\}$,
- $S_{ac} = \{v \in S \mid (v, a), (v, c) \in R\}$.

Note that $R = (S_a \times \{a\}) \cup (S_c \times \{c\}) \cup (S_{ac} \times \{a, c\})$, $x \in S_a$, and $y \in S_{ac}$. We choose $R' = R \setminus \{(v, a) \mid v \in S_{ac}\}$, i.e., $R' = (S_a \times \{a\}) \cup (S_c \times \{c\}) \cup (S_{ac} \times \{c\})$. Note that R' is an (x, y, S)-forcer with respect to (a, c). Let us now show that R' is realizable by $\mathcal{H}'$. For each tuple $(v, w) \in V^2$, we determine $\mathrm{hom}\big((\mathcal{J}, (s, t)) \to (\mathcal{H}', (v, w))\big)$:

- If $v \in S_a$, $\mathrm{hom}\big((\mathcal{J}, (s, t)) \to (\mathcal{H}', (v, a))\big) = 2$.

- For $v \in S_c$, $\mathrm{hom}\big((\mathcal{J}, (s, t)) \to (\mathcal{H}', (v, c))\big) = 6$.

- For $v \in S_{ac}$, $\mathrm{hom}\big((\mathcal{J}, (s, t)) \to (\mathcal{H}', (v, a))\big) = 5$.

- For $v \in S_{ac}$, $\mathrm{hom}\big((\mathcal{J}, (s, t)) \to (\mathcal{H}', (v, c))\big) = 8$.

- Otherwise, $\mathrm{hom}\big((\mathcal{J}, (s, t)) \to (\mathcal{H}', (v, w))\big) = 0$.

So, for each $\mathbf{f} \in R'$ and $\mathbf{g} \in V^2 \setminus R'$,

1. $\mathrm{hom}\big((\mathcal{J}, (s, t)) \to (\mathcal{H}', \mathbf{f})\big) \in \{2, 6, 8\}$.

2. $\mathrm{hom}\big((\mathcal{J}, (s, t)) \to (\mathcal{H}', \mathbf{g})\big) \in \{0, 5\}$.

Therefore, we can apply Lemma 2.3 to complete the proof. $\triangleleft$

Let R' be the (x, y, S)-forcer with respect to (a, c) given by Claim 3.8.1. To obtain item 2, notice that R';NEQ$(\{a, c\})$ is an (x, y, S)-forcer with respect to (c, a), and by Lemma 2.1 it is realizable by $\mathcal{H}'$.

Now let us focus on (3). Note that R';$R_{\{a,c\} \to \{b,d\}}$ is an (x, y, S)-distinguisher with respect to (b, d), which is realizable by $\mathcal{H}'$ by Lemma 2.1. Thus, applying Claim 3.8.1 with a switched to b and c switched to d, we obtain that an (x, y, S)-forcer R'' with respect to (b, d) is realizable by $\mathcal{H}'$.

Finally, to show (4), note that R'';NEQ$(\{b, d\}$ is an (x, y, S)-forcer with respect to (d, b), realizable by $\mathcal{H}'$ (by Lemma 2.1). This completes the proof. $\square$

We show a strengthening of Lemma 3.8: Given a realizable (x, y, S)-distinguisher with respect to some pair of vertices that are potentially far apart in H, we can obtain a realizable (x, y, S)-forcer with respect to some pair of vertices on an induced four-vertex path. Moreover, since P_4s form a connected structure in H (recall Lemma 3.6), we can even choose a P_4 in H and obtain a forcer with respect to any one-sided pair from this very P_4.

Lemma 3.9. *Let $H = (V, E)$ be an irredundant connected bipartite graph. Fix an induced 4-vertex path $(a - b - c - d)$ in H. Let S be a one-sided subset of V with distinct $x, y \in S$. Suppose there is an (x, y, S)-distinguisher R' that is realizable by H. Then there is an (x, y, S)-forcer R with respect to (a, c) such that R is realizable by H.*

Proof. The proof is divided into two parts. First, we show that there is a realizable forcer with respect to a one-sided pair of vertices on *some* induced P_4 in H.

Claim 3.9.1. *There exist an induced 4-vertex path $(a' - b' - c' - d')$ in H and an (x, y, S)-forcer with respect to (a', c') that is realizable by H.*

Proof of Claim. Let $\{\alpha, \beta\}$ be a one-sided set such that R' is an (x, y, S)-distinguisher with respect to (α, β). If α and β have a common neighbor then, since H is irredundant and has at least four vertices, there is an induced 4-vertex path in H that is of the form $(\alpha - b' - \beta - d')$ or of the form $(d' - \alpha - b' - \beta)$. In either case, the statement follows from Lemma 3.8.

Otherwise, there is a shortest path $P = (p_1 - \ldots - p_k)$ from α to β ($p_1 = \alpha$, $p_k = \beta$) with k vertices. Since α and β are in the same bipartition class we have $k \geq 5$. For each $i \in \{1, \ldots, k-4\}$, the relation

$$R_i = R_{\{p_i, p_k\} \to \{p_{i+1}, p_{k-1}\}}; R_{\{p_{i+1}, p_{k-1}\} \to \{p_{i+2}, p_k\}}$$

is a $(p_i, p_k, \{p_i, p_k\})$-forcer with respect to (p_{i+2}, p_k), and R_i is realizable by H according to Lemma 2.1. So, by Lemma 2.1, the relation $R'; R_1; \ldots; R_{k-4}$ is a realizable (x, y, S)-distinguisher with respect to (p_{k-2}, p_k), where $(p_{k-3} - p_{k-2} - p_{k-1} - p_k)$ is an induced path in H. By Lemma 3.8, there is also a realizable (x, y, S)-forcer with respect to (p_{k-2}, p_k). ◁

Let $(a' - b' - c' - d')$ be as in Claim 3.9.1. Note that H satisfies the assumptions of Lemma 3.6, so there is a sequence $P^{(1)}, P^{(2)}, \ldots, P^{(s)}$ of subsets of H, such that (i) each $P^{(i)}$ induces a P_4 in H, (ii) $P^{(1)} = \{a', b', c', d'\}$ and $P^{(s)} = \{a, b, c, d\}$, and (iii) for each $i \leq s-1$, sets $P^{(i)}$ and $P^{(i+1)}$ share two vertices from the same bipartition class.

We prove the statement by induction on s. If $s = 1$, then either $(a - b - c - d) = (a' - b' - c' - d')$ or $(a - b - c - d) = (d' - c' - b' - a')$. By Lemma 3.8, in either case there is a realizable (x, y, S)-forcer with respect to (a, c) and we are done.

So assume that $s \geq 2$ and there is an (x, y, S)-forcer R'' with respect to (a'', c''), such that R'' is realizable by H, where $H[P^{(s-1)}] = (a'' - b'' - c'' - d'')$. Let $(s, t) \in \{(a'', c''), (b'', d'')\}$ be such that $s, t \in P^{s-1} \cap P^s$. Applying Lemma 3.8 for the path $H[P^{s-1}]$, we obtain that an (x, y, S)-forcer with respect to (s, t) is realizable by H. Since $s, t \in P^s$, where s and t are in the same bipartition class, we have $(s, t) \in \{(a, c), (c, a), (b, d), (d, b)\}$. Now using Lemma 3.8 for the path $H[P^s] = (a - b - c - d)$, we conclude that an (x, y, S)-forcer with respect to (a, c) is realizable by H. □

Recall from Lemma 3.4 that the structure of a 4-vertex path is rich enough to encode basic binary relations. In Lemma 3.9 we showed how to obtain forcers with respect to some specified pair (a, c) on a 4-vertex path. In the next lemma we show that such a collection of forcers lets us realize more expressive relations.

Lemma 3.10. *Let $H = (V, E)$ be an irredundant connected bipartite graph. Let $(a - b - c - d)$ be an induced 4-vertex path in H, and let S be a one-sided subset of V. If for every $x, y \in S$ there is an (x, y, S)-distinguisher realizable by H, then, for every fixed $p \geq 1$ and $q \geq 0$, every relation $R \subseteq S^p \times \{a, c\}^q$ is realizable by H.*

Proof. Let $|S| = s$ and enumerate S as $\{x_1, x_2, \ldots, x_s\}$. Let $(a - b - c - d)$ be a fixed 4-vertex path in H. By the assumption of the lemma, for all distinct $x_i, x_j \in S$ there is some (x_i, x_j, S)-distinguisher realizable by H. Thus, by Lemma 3.9, there is also an (x_i, x_j, S)-forcer $R_{i,j}$ with respect to (a, c), realizable by H.

First, let us define an auxiliary relation $I \subseteq S \times \{a, c\}^{s(s-1)}$ and show that it is realizable by H. Recall from Section 2 that $I(x) = \{(v_1, \ldots, v_{s(s-1)}) \in \{a, c\}^{s(s-1)} \mid (x, v_1, \ldots, v_{s(s-1)}) \in I\}$. We say that a relation I is an *indicator* if

- $I(x) \neq \emptyset$ for all $x \in S$, and

24

- $I(x) \cap I(x') = \emptyset$ for all distinct $x, x' \in S$.

Intuitively, each element of $I(x)$ uniquely represents the value x, and every $x \in S$ admits such a representation.

Claim 3.10.1. *There is an indicator $I \subseteq S \times \{a, c\}^{s(s-1)}$ realizable by H.*

Proof of Claim. For brevity, denote $\mathsf{Ind} := \{(i, j) \mid i, j \in [s], i \neq j\}$. As for each $(i, j) \in \mathsf{Ind}$, the forcer $R_{i,j}$ is realizable by H, it is sufficient to show an indicator realizable by the structure $\mathcal{H}' = (V, \bigcup_{(i,j) \in \mathsf{Ind}} \{R_{i,j}\})$.

We define a structure $\mathcal{I} = (U, \mathbf{S})$ with the same signature as $\mathcal{H}'$. The vertex set of $\mathcal{I}$ is $U = \{u\} \cup \bigcup_{(i,j) \in \mathsf{Ind}} \{u_{i,j}\}$. We also define an interface $\mathbf{u}$ of $\mathcal{I}$ by setting

$$\mathbf{u} = (u, u_{1,2}, \dots, u_{1,s}, u_{2,1}, u_{2,3}, \dots, u_{2,s}, \dots, u_{s,1}, \dots, u_{s,s-1}).$$

For all $(i, j) \in \mathsf{Ind}$, we set $R_{i,j}^{\mathcal{I}} = \{(u, u_{i,j})\}$, i.e., the (x_i, x_j, S)-forcer $R_{i,j}$ is applied to the tuple $(u, u_{i,j})$. Consider the set

$$I := \{\mathbf{f} \in V^{s(s-1)+1} \mid \mathrm{hom}\big((\mathcal{I}, \mathbf{u}) \to (\mathcal{H}', \mathbf{f})\big) \neq 0\}$$

and let $(f, f_{1,2}, \dots, f_{1,s}, f_{2,1}, f_{2,3}, \dots, f_{2,s}, \dots, f_{s,1}, \dots, f_{s,s-1}) \in I$.

First, clearly $f \in S$ and for all $(i, j) \in \mathsf{Ind}$ we have $f_{i,j} \in \{a, c\}$. Next, if $f = x_i$ for $i \in [s]$, then the properties of forcers imply that

- $f_{i,j} = a$ for all j, such that $(i, j) \in \mathsf{Ind}$, and

- for all $i' \in [s] \setminus \{i\}$ there is j, such that $(i', j) \in \mathsf{Ind}$ and $f_{i',j} = c$.

Therefore $I(x_i) \cap I(x_{i'}) = \emptyset$ for all distinct $i, i' \in [s]$. Finally, by the definition of a forcer, we note that for all $i \in [s]$ and all $(i', j') \in \mathsf{Ind}$ we have that at least one of (x_i, a) or (x_i, c) is in $R_{i',j'}$, and thus $I(x_i) \neq \emptyset$. Summing up, I is an indicator and by Corollary 2.4 it is realizable by H. Note that the exact definition of I depends on the exact definitions of $R_{i,j}$ for $(i, j) \in \mathsf{Ind}$. ◁

Let I be the indicator given by Claim 3.10.1; it is realizable by H. For each $i \in [s]$, we fix some element $\mathsf{id}(x_i)$ from $I(x_i)$. We think of $\mathsf{id}(x_i)$ as a unique identifier of x_i. The properties of I ensure these identifiers are pairwise distinct. Note that $\mathsf{id}(x_i) \in \{a, c\}^{s(s-1)}$.

Let $p \geq 1$ and $q \geq 0$ be integers and consider a relation $R \subseteq S^p \times \{a, c\}^q$. Let $R_I \subseteq \{a, c\}^{p(s-1)s+q}$ be the relation defined as follows:

$$R_I = \{(\mathsf{id}(y_1), \mathsf{id}(y_2), \dots, \mathsf{id}(y_p), b_1, b_2, \dots, b_q) \mid (y_1, y_2, \dots, y_p, b_1, b_2, \dots, b_q) \in R\}.$$

Intuitively, we may see R_I as R "translated" to the the ground set $\{a, c\}$, where the translation of each element of S to a sequence over $\{a, c\}$ is provided by the choice of $\mathsf{id}(\cdot)$. By Lemma 3.4, the relation R_I is realizable by H. Thus, in order to prove that R is realizable by H, it is sufficient to show that R is realizable by a structure $\mathcal{H}' = (V, \{I, R_I\})$. Slightly abusing notation we use $\{I, R_I\}$ also to denote the signature of $\mathcal{H}'$. Consider a structure $\mathcal{R} = (U, \mathbf{S})$ with signature $\{I, R_I\}$, defined as follows:

$$U = \bigcup_{\ell \in [p]} \{y_\ell, u_1^\ell, \dots, u_{s(s-1)}^\ell\} \cup \bigcup_{\ell \in [q]} \{b_\ell\},$$

$$I^{\mathcal{R}} = \bigcup_{\ell \in [q]} \{(y_\ell, u_1^\ell, \dots, u_{s(s-1)}^\ell)\},$$

$$R_I^{\mathcal{R}} = \{(u_1^1, \dots, u_{s(s-1)}^1, u_1^2, \dots, u_{s(s-1)}^2, \dots, u_1^p, \dots, u_{s(s-1)}^p, b_1, \dots, b_q)\}.$$

25

Define $\mathbf{y} = (y_1, y_2, \ldots, y_p, b_1, \ldots, b_q)$. It is straightforward to verify that $\mathrm{hom}\big((\mathcal{R}, \mathbf{y}) \to (\mathcal{H}', \mathbf{f})\big) \neq 0$ if and only if $(\mathsf{id}(y_1), \mathsf{id}(y_2), \ldots, \mathsf{id}(y_p), b_1, \ldots, b_q) \in R_I$, which in turn happens if and only if $(y_1, y_2, \ldots, y_p, b_1, \ldots, b_q) \in R$. Thus by Corollary 2.4 the relation R is realizable by H. $\qquad\square$

Let us now introduce one more special type of relation. Consider a bipartite graph $H = (V, E)$ and a subset $S \subseteq V$. Let (X, Y) be a partition of S. Moreover, let a, c be two distinct vertices in V. A relation $R \subseteq S \times \{a, c\}$ is an (X, Y)-*partitioner with respect to* (a, c) if:

- for every $x \in X$ it holds that $R(x) = \{a\}$,

- for every $y \in Y$ it holds that $R(y) = \{c\}$.

As each partitioner is a relation in $S \times \{a, c\}$, Lemma 3.10 immediately yields the following.

Corollary 3.11. *Let $H = (V, E)$ be an irredundant connected bipartite graph. Let $(a-b-c-d)$ be an induced 4-vertex path in H, and let S be a one-sided subset of V. If for every $x, y \in S$ there is an (x, y, S)-distinguisher realizable by H, then for every partition (X, Y) of S, the (X, Y)-partitioner with respect to (a, c) is realizable by H.*

3.6   Proof of Lemma 3.2

In this section we finally prove Lemma 3.2. Let us first discuss the plan. Lemma 3.2 will easily follow from Lemma 3.10 (for $q = 0$), provided that for all distinct $x, y \in S$, there is an (x, y, S)-distinguisher realizable by H. We prove this statement inductively, essentially deriving new forcers from forcers that are "smaller" with respect to some measure on (x, y, S). On the way, we use Lemma 3.9 to turn distinguishers into forcers, and we use Corollary 3.11 to turn forcers into partitioners.

Lemma 3.12. *Let $H = (V, E)$ be an irredundant connected bipartite graph with $\mathrm{irr}(H) \geq 2$. Let S be a one-sided subset of V, and let x, y be distinct vertices in S. Fix any induced 4-vertex path $(a - b - c - d)$ in H. Then there is an (x, y, S)-forcer R with respect to (a, c) such that R is realizable by H.*

Proof. Let S and S' be one-sided irredundant sets (not necessarily in the same part of the bipartition) such that S contains two distinct vertices $x, y \in S$ and S' contains two distinct vertices $x', y' \in S'$. We define an order by setting $(x', y', S') < (x, y, S)$ if one of the following holds:

- $|S'| < |S|$, or
- $|S'| = |S|$ and $\mathrm{dist}(\{x', y'\}, S' \setminus \{x', y'\}) < \mathrm{dist}(\{x, y\}, S \setminus \{x, y\})$.

We prove the statement of the lemma by induction with respect to this order. Note that (x, y, S) is a minimal element if $S = \{x, y\}$.

Suppose that the following holds:

$$\text{There is an } (x, y, S)\text{-distinguisher or a } (y, x, S)\text{-distinguisher realizable by } H. \qquad (4)$$

In the first case we can apply Lemma 3.9 to obtain an (x, y, S)-forcer with respect to (a, c) that is realizable by H. In the second case, note that a (y, x, S)-forcer with respect to (a, c), obtained by the application of Lemma 3.9, is an (x, y, S)-forcer with respect to (c, a). Thus,

by Lemma 3.8, we again obtain an (x, y, S)-forcer with respect to (a, c) that is realizable by H. Therefore, in order to show the statement of the lemma, it suffices to show (4).

We first consider the base case $S = \{x, y\}$. Since H is irredundant we have $N(x) \neq N(y)$. Suppose there is a vertex $q \in N(y) \setminus N(x)$ (the other case is analogous). Since H is connected there is some vertex $p \in N(x)$ (possibly $p \in N(x) \cap N(y)$). Since S is one-sided $R_{\{x,y\} \to \{p,q\}}$ is an (x, y, S)-distinguisher with respect to (p, q) (if $p \notin N(y)$ it is even a forcer). This is what we need according to (4) and this completes the base case.

For the inductive step, let x, y, S be as given in the statement of the lemma. We can assume $S \neq \{x, y\}$. To shorten notation let $S_0 = S \setminus \{x, y\}$. Let P be a shortest path from $\{x, y\}$ to S_0. So P is of the form $(p_1 - \ldots - p_k)$, where $p_1 \in \{x, y\}$, $p_k \in S_0$, and $k \geq 3$ (since S_0 is assumed to be non-empty). Let $p = p_2$, then p is a neighbor of at least one of x or y. Since $N(x) \neq N(y)$ and H is connected we can choose a vertex q, such that both x and y have a neighbor in $\{p, q\}$, and at least one of p or q has exactly one neighbor in $\{x, y\}$. Note that, according to (4), it does not cause any issues to rename the vertices x and y because a realizable (y, x, S)-distinguisher with respect to (a, c) also gives a realizable (x, y, S)-forcer with respect to (a, c). So, without loss of generality (by renaming x and y), we can assume that $p \in N(x)$ and $q \in N(y) \setminus N(x)$.

Let S_0' be a minimal-size set with the properties $p_{k-1} \in S_0'$ and $S_0 \subseteq N(S_0')$. Thus, $|S_0'| \leq |S_0|$. Let $S' = \{p, q\} \cup S_0'$. Clearly, S' is a one-sided set with $|S'| \leq |S|$.

Case 1 Suppose that $k = 3$ and consequently $p = p_2 = p_{k-1}$. Then $|S'| < |S|$. So, for any distinct $x', y' \in S'$, we have $(x', y', S') < (x, y, S)$, and consequently, by the induction hypothesis, there is an (x', y', S')-forcer with respect to (a, c) that is realizable by H. From Corollary 3.11 we obtain that for every partition (X, Y) of S', the (X, Y)-partitioner with respect to (a, c) is realizable by H. In particular, let R' be the $(S' \cap N(x), S' \setminus N(x))$-partitioner. Then R' is realizable by H, and $R = R_{S \to S'}; R'$ is also realizable by H (Lemma 2.1). We now show that R is an (x, y, S)-distinguisher with respect to (a, c). We have $p \in R_{S \to S'}(x) = S' \cap N(x)$ and, for each $v \in S' \cap N(x)$, we have $R'(v) = \{a\}$. It follows that $R(x) = \{a\}$. Since $q \in R_{S \to S'}(y)$ but $q \notin N(x)$, we have $R'(q) = \{c\}$ and consequently $c \in R(y)$. Since $S_0 \subseteq N(S_0')$ every vertex in S has a neighbor in S'. Therefore, for $u \in S$, at least one of (u, a) or (u, c) is in R. This shows that R is an (x, y, S)-distinguisher with respect to (a, c). By (4), we are done.

Case 2 Suppose $k > 3$, i.e., $k \geq 5$ (since $\{x, y\}$ is one-sided). We know that $|S_0'| \leq |S_0|$ and consequently $|S'| \leq |S|$ — but now these sets could have the same cardinality. However, by the choice of P, note that $\text{dist}(\{p, q\}, S' \setminus \{p, q\}) \leq k - 2 < k = \text{dist}(\{x, y\}, S \setminus \{x, y\})$. So, we have $(p, q, S') < (x, y, S)$ and by the induction hypothesis there is a (p, q, S')-forcer R' with respect to (a, c). By the choice of P it also follows that no vertex in S_0 has a common neighbor with x or y, and therefore x has no neighbor in S_0'. Thus, p is the only neighbor of x in S'. Consequently, $R = R_{S \to S'}; R'$ is an (x, y, S)-distinguisher with respect to (a, c), and this is what we need according to (4).

This completes the proof. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

We can now prove Lemma 3.2, which we restate for convenience.

Lemma 3.2. *Let $H = (V, E)$ be a connected bipartite graph with $\text{irr}(H) \geq 2$, and let $S \subseteq V$ be a one-sided irredundant set. For every fixed $p \geq 1$, every relation $R \subseteq S^p$ is realizable by H.*

Proof. Let S' be a maximal-size irredundant superset of S in V. Note that every maximal-size irredundant set in H contains a vertex from each class of vertices with identical neighborhood. Thus, $H' = H[S']$ is connected as H is connected. Furthermore, we have $\mathrm{irr}(H') = \mathrm{irr}(H) \geq 2$. Since H' is an induced subgraph of H it suffices to show that every relation in S^p is realizable by H'. Since $\mathrm{irr}(H') \geq 2$, the graph H' is not a complete bipartite graph and therefore it contains an induced four-vertex path $(a-b-c-d)$. By Lemma 3.12, for every pair of distinct $x, y \in S$, there is an (x, y, S)-forcer R with respect to (a, c) such that R is realizable by H'. The statement of the lemma follows from Lemma 3.10 for $q = 0$. $\qquad\square$

4 Counting list homomorphisms to general graphs H

In this section we discuss how to lift the results from Section 3, where we assumed H to be bipartite, to the general case.

4.1 Associated bipartite graphs

For a graph $H = (V, E)$, by H^* we denote its *associated bipartite graph*, i.e., the graph with vertex set $\{v', v'' \mid v \in V\}$ and edge set $\{u'v'', u''v' \mid uv \in E\}$. We set $V' := \{v' \mid v \in V\}$ and $V'' := \{v'' \mid v \in V\}$. So (V', V'') is a bipartition of H^*. Note that if $H = H_1 + H_2$, where $+$ denotes the disjoint sum, then $H^* = H_1^* + H_2^*$.

Recall that if H is connected and nonbipartite, then $\mathrm{irr}(H)$ is the cardinality of the largest irredundant set in H. Let us point out that associated bipartite graphs allow us to provide a uniform definition of $\mathrm{irr}(H)$, which does not need to distinguish bipartite and nonbipartite graphs.

Observation 4.1. *For every graph H it holds that* $\mathrm{irr}(H) = \mathrm{irr}(H^*)$.

Proof. First, observe that $S \subseteq V(H)$ is irredundant if and only if $S' = \{v' \mid v \in S\}$ is irredundant in H^* if and only if $S'' = \{v'' \mid v \in S\}$ is irredundant in H^*.

Furthermore, we observe that there is a correspondence between connected components of H and connected components of H^*. Let H' be a connected component of H. If H' is bipartite, then H'^* consists of two disjoint copies of H' and so $\mathrm{irr}(H') = \mathrm{irr}(H'^*)$. If H' is nonbipartite, then H'^* is connected (and bipartite). Indeed, this follows from the fact that for every two vertices x, y of H' there is an even x-y-walk and an odd x-y-walk in H'. Thus all vertices x', x'', y', y'' are in the same connected component of H'.

Consequently, each bipartite component H' of H corresponds to two components of H^*, both isomorphic to H', and each nonbipartite component H' of H corresponds to one connected component of H^*, i.e., $(H')^*$.

Now the claim easily follows from the previous observations about S, S' and S''. $\qquad\square$

Note that H^* is a biclique if and only if H is a reflexive clique. Thus, we observe that $\mathrm{irr}(H) \geq 2$ if and only if H has a connected component that is not a biclique nor a reflexive clique. This allows us to restate the complexity dichotomy for #LHom(H), which was originally observed as a simple consequence of the non-list result of Dyer and Greenhill [10] by Díaz, Serna, and Thilikos [9], and also by Hell and Nešetřil [16].

Theorem 4.2 ([9, 10, 16])**.** *Let H be a fixed graph. If* $\mathrm{irr}(H) = 1$*, then* #LHom(H) *is polynomial-time solvable, and otherwise it is* #P*-complete.*

4.2 Algorithm for general graphs H

For an instance (G, L) of #LHom(H), we define its *associated instance* (G^*, L^*) of #LHom(H^*), where for all $v \in V(G)$

$$L^*(v') = \{x' \mid x \in L(v)\} \text{ and } L^*(v'') = \{x'' \mid x \in L(v)\}.$$

We say that a homomorphism $f \colon G^* \to H^*$ is *clean* if, for every $v \in V(G)$ and $x \in V(H)$ it holds that $f(v') = x'$ if and only if $f(v'') = x''$.

Lemma 4.3. *There is a bijection between homomorphisms from (G, L) to H and clean homomorphisms from (G^*, L^*) to H^*.*

Proof. Consider a homomorphism $f \colon (G, L) \to H$. Let $\sigma(f)$ be a mapping $f^* \colon V(G^*) \to V(H^*)$ defined as follows. Let $v \in V(G)$ and suppose $f(v) = x$. We set $f^*(v') = x'$ and $f^*(v'') = x''$. It is straightforward to verify that f^* is a clean homomorphism from (G^*, L^*) to H^* and that σ is a bijection. $\qquad\square$

Using Lemma 4.3 we can show the algorithmic statement of Theorem 1.2. Note that it follows from the subsequent slightly stronger result.

Theorem 4.4. *For each graph H, the #LHom(H) problem on n-vertex instances given along with a tree decomposition of width at most t can be solved in time $\mathrm{irr}(H)^t \cdot n^{\mathcal{O}(1)}$.*

Proof. Consider an instance (G, L) of #LHom(H). By Lemma 4.3, it suffices to count the number of clean homomorphisms from (G^*, L^*) to H^*.

Let $\mathcal{T}$ be a tree decomposition of G with width at most t. We modify it as follows: in every bag of $\mathcal{T}$ we replace every vertex $v \in V(G)$ with two vertices $v', v'' \in V(G^*)$. It is straightforward to verify that this way we obtain a tree decomposition of G^* with width at most $2t$; let us call this decomposition $\mathcal{T}^*$.

First, just like in the proof of Theorem 3.1, we reduce the problem to its equivalent weighted version with each list of size at most $\mathrm{irr}(H^*) = \mathrm{irr}(H)$. Consider one bag of $\mathcal{T}^*$ and recall that we are only interested in counting the number of clean homomorphisms from G^* to H^*. Therefore, even though the size of the bag is at most $2t$, there are at most $\mathrm{irr}(H)^t$ colorings that could possibly be extended to a clean homomorphism from G^* to H^*.

Thus, by an argument analogous to the one in the proof of Theorem 3.1, we obtain the desired running time. $\qquad\square$

4.3 Hardness for general graphs H

Recall the definition of consistent instances from the beginning of Section 3. The following lemma is a crucial tool used in our hardness reduction.

Lemma 4.5. *For a consistent instance (G, L') of #LHom(H^*), define $L \colon V(G) \to 2^{V(H)}$ as follows: $L(v) := \{x \mid \{x', x''\} \cap L'(v) \neq \emptyset\}$. There is a bijection between homomorphisms from (G, L') to H^* and homomorphisms from (G, L) to H.*

Proof. For a homomorphism $f \colon (G, L') \to H^*$, define $\sigma(f) \colon V(G) \to V(H)$ as follows. If $f(v) \in \{x', x''\}$, then $\sigma(f)(v) = x$. It is straightforward to verify that $\sigma(f)$ is a homomorphism from (G, L) to H. Furthermore, σ is a bijection – as the instance (G, L') is consistent, for no $v \in V(G)$ and $x \in V(H)$ it holds that $x', x'' \in L'(v)$. $\qquad\square$

Now we are ready to prove the complexity statement in Theorem 1.2. Again, we will actually show a slightly stronger result, where we consider *pathwidth* as the parameter.

Theorem 4.6. *Let H be a graph with $\mathrm{irr}(H) \geq 2$. Assuming the #SETH, there is no $\varepsilon > 0$, such that #LHom(H) on connected bipartite n-vertex instances G can be solved in time $(\mathrm{irr}(H) - \varepsilon)^{\mathrm{pw}(G)} \cdot n^{\mathcal{O}(1)}$, even if $\mathrm{pw}(G) = \mathrm{tw}(G)$ and G is given along with an optimal path decomposition.*

Proof. Note that the associated graph H^* contains a connected component other than a biclique and therefore $\mathrm{irr}(H^*) \geq 2$. We reduce from #LHom(H^*). Let (G', L') be an instance of #LHom(H^*), where G' has n vertices and is given along with a path decomposition $\mathcal{P}$ with width t. Recall from the proof of Theorem 3.1 that we can assume that the instance (G', L') is consistent.

Let (G', L) be an instance of #LHom(H) constructed as in Lemma 4.5, note that the lists L are computed in polynomial time. By Lemma 4.5, we know that $\mathrm{hom}\big((G', L') \to H^*\big) = \mathrm{hom}\big((G', L) \to H\big)$. Now we modify G' to make sure that the treewidth and the pathwidth of our instance is *exactly* t.

Let v be an arbitrary vertex of G' that appears in the last bag of $\mathcal{P}$. Let G be the graph obtained from G' by adding $2t - 1$ new vertices, which, together with v, induce a biclique $K_{t,t}$ and are nonadjacent to any other vertices of G'. Let A and B be the bipartition classes of this biclique, where $v \in A$.

The number of vertices of G is $n + 2t - 1 \leq 3n$. Furthermore, as $\mathrm{tw}(G') \leq \mathrm{pw}(G') \leq t$ and $\mathrm{tw}(K_{t,t}) = \mathrm{pw}(K_{t,t}) = t$, we conclude that $\mathrm{tw}(G) = \mathrm{pw}(G) = t$. Finally, it is easy to obtain an optimal path decomposition of G by appending to $\mathcal{P}$ an optimal path decomposition of $K_{t,t}$, where v appears in the first bag.

For each $a \in L(v)$, let a' be an arbitrary neighbor of a in H. Let $L_a : V(G) \to 2^{V(H)}$ be the list function defined as follows:

$$L_a(u) = \begin{cases} \{a\} & \text{if } u \in A, \\ \{a'\} & \text{if } u \in B, \\ L(u) & \text{otherwise.} \end{cases}$$

Since lists L_a force a unique coloring of vertices in $A \cup B$, we observe that $\mathrm{hom}\big((G, L_a) \to H\big)$ is equal the number of homomorphisms from (G', L) to H that map v to a. Consequently, we obtain

$$\mathrm{hom}\big((G', L') \to H^*\big) = \mathrm{hom}\big((G', L) \to H\big) = \sum_{a \in L(v)} \mathrm{hom}\big((G, L_a) \to H\big).$$

Thus, if for each $a \in L(v)$ we could compute $\mathrm{hom}\big((G, L_a) \to H\big)$ in time $(\mathrm{irr}(H) - \varepsilon)^{\mathrm{pw}(G)} \cdot (n + 2t - 1)^{\mathcal{O}(1)}$, then we could also compute $\mathrm{hom}\big((G', L') \to H^*\big)$ in time

$$(\mathrm{irr}(H) - \varepsilon)^{\mathrm{pw}(G)} \cdot (n + 2t - 1)^{\mathcal{O}(1)} = (\mathrm{irr}(H^*) - \varepsilon)^t \cdot n^{\mathcal{O}(1)}.$$

Recall that $\mathrm{irr}(H^*) \geq 2$. By Theorem 3.3, the existence of such an algorithm for #LHom(H^*) contradicts the #SETH. $\qquad\square$

4.4 From #LHom(H) to #Hom(H): special cases

In this section we show two corollaries of our main result that concern *non-list* problems. First, let us show that Corollary 3.5 implies the following lower bound.

Corollary 4.7. *There is no algorithm that counts all independents sets in n-vertex bipartite graphs given with a path decomposition of width t in time $(2 - \varepsilon)^t \cdot n^{\mathcal{O}(1)}$ for any $\varepsilon > 0$, unless the #SETH fails.*

Proof. We give a pathwidth-preserving reduction from #LHom(P_4). Let (G, L) be an instance of #LHom(P_4), where $P_4 = (a - b - c - d)$. Note that we can safely assume that G is bipartite, as otherwise the answer is clearly 0. Furthermore, similarly as in the proof of Theorem 3.1, we can assume that G is connected. Let (X, Y) be a bipartition of G.

Let L' (resp. L'') be the list function defined by setting $L'(x) = L(x) \cap \{a, c\}$ and $L'(y) = L(y) \cap \{b, d\}$ (resp. $L''(x) = L(x) \cap \{b, d\}$ and $L''(y) = L(y) \cap \{a, c\}$) for every $x \in X$ and $y \in Y$. Observe that in any homomorphism f from G to P_4, either $f(X) \subseteq \{a, c\}$ and $f(Y) \subseteq \{b, d\}$, or $f(X) \subseteq \{b, d\}$ and $f(Y) \subseteq \{a, c\}$. Thus we have

$$\hom\big((G, L) \to P_4\big) = \hom\big((G, L') \to P_4\big) + \hom\big((G, L'') \to P_4\big).$$

Let us focus on computing $\hom\big((G, L') \to P_4\big)$, as computing $\hom\big((G, L'') \to P_4\big)$ is symmetric. We observe that if for some vertex $v \in V(G)$ we have $L'(v) = \{a\}$ (resp. $L'(v) = \{d\}$), then we can safely remove d (resp. a) from the lists of all neighbors of v, and then delete the vertex v. On the other hand, if $L'(v) = \{b\}$ or $L'(v) = \{c\}$, then we can safely remove v from the graph. Let $(\widetilde{G}, \widetilde{L})$ the the instance obtained by exhaustive application of the rules above and note that

$$\hom\big((\widetilde{G}, \widetilde{L}) \to P_4\big) = \hom\big((G, L') \to P_4\big).$$

If the list of some vertex is empty, then $\hom\big((\widetilde{G}, \widetilde{L}) \to P_4\big) = 0$ and we are done. Otherwise, every list in $\widetilde{L}$ is either $\{a, c\}$ (for vertices in X) or $\{b, d\}$ (for vertices in Y). Consequently, $\hom\big((\widetilde{G}, \widetilde{L}) \to P_4\big)$ is precisely the number of independent sets in $\widetilde{G}$ (the independent set is formed by the vertices mapped to a and d).

Recall that $\widetilde{G}$ is a subgraph of G. Thus, given a path decomposition of G with width at most t, we can easily obtain a path decomposition of $\widetilde{G}$ with width at most t. Now the lower bound follows directly from Corollary 3.5. $\square$

Next, let us focus on the case that $H = K_q$ for $q \geq 3$, i.e., counting proper q-colorings of a given graph. As this problem is a special case of #LHom(K_q), the algorithmic statement follows immediately from Theorem 4.4. However, Theorem 4.6 proves hardness only for counting *list* q-colorings. Let us show a reduction from the problem of counting list q-colorings to the problem of counting q-colorings. For simplicity, we will only prove hardness parameterized by the treewidth of the instance graph, but using the same approach as in Theorem 4.6 we can obtain an analogous result parameterized by the pathwidth. Let us point out that the lower bound for counting proper q-colorings, conditioned on the SETH, follows from the result for the decision variant of the problem by Lokshtanov, Marx, Saurabh [22].

Corollary 4.8. *Let $q \geq 3$. On n-vertex instances with treewidth tw, the number of proper q-colorings cannot be counted in time $(q - \varepsilon)^{\mathrm{tw}} \cdot n^{\mathcal{O}(1)}$, for any $\varepsilon > 0$, even if a tree decomposition of width tw is given as part of the input, unless the #SETH fails.*

Proof. Let $G = (V, E)$ be an n-vertex graph and let $L \colon V \to 2^{[q]}$ be a list function. Suppose that the treewidth of G is tw and G is given along with a tree decomposition $\mathcal{T}$ of width tw.

Let G' be a graph obtained from G as follows. First, we introduce a q-vertex clique K with vertices $\{x_1, x_2, \ldots, x_q\}$. Next, for each $v \in V$ and each $i \in [q]$, we make v adjacent to x_i if and only if $i \notin L(v)$.

It is straightforward to observe that (G, L) is a yes-instance of list q-coloring if and only if G' is a yes-instance of q-coloring. Furthermore, each proper list coloring of (G, L) corresponds to exactly $q!$ proper colorings of G', one for each proper coloring of vertices of K. Thus $\hom\big((G, L) \to K_q\big) = \frac{1}{q!} \cdot \hom\big(G' \to K_q\big)$.

The number of vertices of G' is $n + q$. Now let us deal with the treewidth. We can easily modify the tree decomposition $\mathcal{T}$ of G into a tree decomposition $\mathcal{T}'$ of G' by including all vertices of K in every bag. This proves that the treewidth of G' is at most tw $+ q$. Let us further modify the instance, so that the treewidth is exactly tw $+ q$. We use a trick similar to the one in the proof of Theorem 4.6.

Let v be an arbitrary vertex of G' and let G'' be obtained from G' by introducing $2(\mathrm{tw} + q) - 1$ new vertices, which, together with v, form a biclique $K_{\mathrm{tw}+q, \mathrm{tw}+q}$. Recall that the treewidth of $K_{\mathrm{tw}+q, \mathrm{tw}+q}$ is exactly tw $+ q$, and the treewidth of G' is also tw $+ q$. On the other hand we can turn the tree decomposition $\mathcal{T}'$ of G' into a tree decomposition $\mathcal{T}''$ of G'' as follows. Let $\widetilde{\mathcal{T}}$ be an optimal tree decomposition of the biclique $K_{\mathrm{tw}+q, \mathrm{tw}+q}$. We choose any bag of $\mathcal{T}'$ containing v and make it adjacent to any bag of $\widetilde{\mathcal{T}}$ containing v. This way we obtain a tree decomposition of G'' with width tw $+ q$. The number of vertices of G'' is $n + q + 2(\mathrm{tw} + q) - 1 = \mathcal{O}(n)$.

Now, let $f(q, \mathrm{tw})$ be the number of proper q-colorings of $K_{\mathrm{tw}+q, \mathrm{tw}+q}$ with the color of one vertex fixed. We observe that $\hom\big(G' \to K_q\big) = \frac{1}{f(q, \mathrm{tw})} \cdot \hom\big(G'' \to K_q\big)$.

Thus, if we could count the number of proper q-colorings of G'' in time $(q - \varepsilon)^{\mathrm{tw}+q} n^{\mathcal{O}(1)}$, we could count the number of list q-colorings of (G, L) in time $(q - \varepsilon)^{\mathrm{tw}+q} n^{\mathcal{O}(1)} = (q - \varepsilon)^{\mathrm{tw}} n^{\mathcal{O}(1)}$. By Theorem 4.4, this would contradict the #SETH. □

5 Conclusion

Let us conclude the paper with the discussion of a potential extension of our results to the non-list variant of the #LHom(H) problem, i.e., the problem of counting homomorphisms to a fixed graph H. Denote this problem by #Hom(H). The complexity dichotomy for #Hom(H) was provided by Dyer and Greenhill [10] and it is exactly the same as for the list variant: the problem is polynomial-time solvable if every component of H is either a reflexive clique or an irreflexive biclique, and otherwise it is #P-complete.

While the algorithmic statement of Theorem 1.2 clearly carries over to #Hom(H) (as #Hom(H) is a restriction of #LHom(H) where all lists are equal to $V(H)$), our hardness proof heavily exploits non-trivial lists. The simple tricks we used in Corollaries 4.7 and 4.8 to reduce the list variant of coloring to the non-list variant cannot be easily generalized to arbitrary graphs H.

Let us point out that the fine-grained complexity of the decision variant of #Hom(H), parameterized by the treewidth of the instance graph is not fully understood [27].

Typically, the hardness proofs concerning the complexity of (non-list) graph homomorphism problems involve some tools from universal algebra and algebraic graph theory [3,15,28].

For graphs $H_1, H_2, \ldots, H_p$, we define their *direct product* $H_1 \times \cdots \times H_p$ as follows:

$$
\begin{aligned}
V(H_1 \times \cdots \times H_p) =& V(H_1) \times V(H_2) \times \cdots \times V(H_p), \\
E(H_1 \times \cdots \times H_p) =& \{(x_1, \ldots, x_p)(y_1, \ldots, y_p) \mid x_i y_i \in E(H_i) \text{ for all } i \in [p]\}.
\end{aligned}
$$

The following observation is straightforward.

Observation 5.1. *Let G and $H = H_1 \times \cdots \times H_p$ be two graphs. A function $f \colon V(G) \to V(H)$ is a homomorphism if and only if for every $i \in [p]$, the function $\Pi_i(f)$ is a homomorphism from G to H_i.*

As an immediate consequence, we obtain

$$
\hom\bigl(G \to H_1 \times \cdots \times H_p\bigr) = \prod_{i=1}^{p} \hom\bigl(G \to H_i\bigr).
$$

In other words, if H can be obtained as a direct products of some factors $H_1, H_2, \ldots, H_p$, we can reduce solving #Hom(H) to solving #Hom(H_i) for all i (with the same instance graph). As irr(H_i) can be much smaller than irr(H), we observe that for some graphs H, the #Hom(H) problem can be solved faster than the #LHom(H) problem. In particular, while it was the truth for #LHom(H), the parameter irr(H) is not always the correct base of the exponential factor appearing in the complexity of an optimal algorithm solving #Hom(H).

References

[1] Libor Barto. The dichotomy for conservative constraint satisfaction problems revisited. In *Proceedings of the 26th Annual IEEE Symposium on Logic in Computer Science, LICS 2011, June 21-24, 2011, Toronto, Ontario, Canada*, pages 301–310. IEEE Computer Society, 2011. doi:10.1109/LICS.2011.25.

[2] Glencora Borradaile and Hung Le. Optimal dynamic program for r-domination problems over tree decompositions. In Jiong Guo and Danny Hermelin, editors, *11th International Symposium on Parameterized and Exact Computation, IPEC 2016, August 24-26, 2016, Aarhus, Denmark*, volume 63 of *LIPIcs*, pages 8:1–8:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPIcs.IPEC.2016.8.

[3] Andrei A. Bulatov. H-coloring dichotomy revisited. *Theor. Comput. Sci.*, 349(1):31–39, 2005. [doi:10.1016/j.tcs.2005.09.028](doi:10.1016/j.tcs.2005.09.028).

[4] Andrei A. Bulatov. Complexity of conservative constraint satisfaction problems. *ACM Trans. Comput. Log.*, 12(4):24:1–24:66, 2011. [doi:10.1145/1970398.1970400](doi:10.1145/1970398.1970400).

[5] Andrei A. Bulatov. Conservative constraint satisfaction re-revisited. *J. Comput. Syst. Sci.*, 82(2):347–356, 2016. doi:10.1016/j.jcss.2015.07.004.

[6] Radu Curticapean, Nathan Lindzey, and Jesper Nederlof. A tight lower bound for counting Hamiltonian cycles via matrix rank. In Artur Czumaj, editor, *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 1080–1099. SIAM, 2018. [doi:10.1137/1.9781611975031.70](doi:10.1137/1.9781611975031.70).

[7] Radu Curticapean and Dániel Marx. Tight conditional lower bounds for counting perfect matchings on graphs of bounded treewidth, cliquewidth, and genus. In Robert Krauthgamer, editor, *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 1650–1669. SIAM, 2016. `doi:10.1137/1.9781611974331.ch113`.

[8] Holger Dell, Thore Husfeldt, Dániel Marx, Nina Taslaman, and Martin Wahlen. Exponential time complexity of the permanent and the Tutte polynomial. *ACM Trans. Algorithms*, 10(4):21:1–21:32, 2014. `doi:10.1145/2635812`.

[9] Josep Díaz, Maria J. Serna, and Dimitrios M. Thilikos. Recent results on parameterized *H*-colorings. In Jaroslav Nešetřil and Peter Winkler, editors, *Graphs, Morphisms and Statistical Physics, Proceedings of a DIMACS Workshop, New Brunswick, New Jersey, USA, March 19-21, 2001*, volume 63 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 65–85. DIMACS/AMS, 2001. `doi:10.1090/dimacs/063/05`.

[10] Martin Dyer and Catherine Greenhill. The complexity of counting graph homomorphisms. *Random Structures & Algorithms*, 17(3-4):260–289, 2000.

[11] László Egri, Dániel Marx, and Paweł Rzążewski. Finding list homomorphisms from bounded-treewidth graphs to reflexive graphs: a complete complexity characterization. In Rolf Niedermeier and Brigitte Vallée, editors, *35th Symposium on Theoretical Aspects of Computer Science, STACS 2018, February 28 to March 3, 2018, Caen, France*, volume 96 of *LIPIcs*, pages 27:1–27:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. `doi:10.4230/LIPIcs.STACS.2018.27`.

[12] Tomás Feder, Pavol Hell, and Jing Huang. Bi-arc graphs and the complexity of list homomorphisms. *J. Graph Theory*, 42(1):61–80, 2003. `doi:10.1002/jgt.10073`.

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

[14] Geňa Hahn and Claude Tardif. *Graph homomorphisms: structure and symmetry*, pages 107–166. Springer Netherlands, Dordrecht, 1997. `doi:10.1007/978-94-015-8937-6_4`.

[15] Pavol Hell and Jaroslav Nešetřil. On the complexity of *H*-coloring. *J. Comb. Theory, Ser. B*, 48(1):92–110, 1990. `doi:10.1016/0095-8956(90)90132-J`.

[16] Pavol Hell and Jaroslav Nešetřil. Counting list homomorphisms and graphs with bounded degrees. In Jaroslav Nešetřil and Peter Winkler, editors, *Graphs, Morphisms and Statistical Physics, Proceedings of a DIMACS Workshop, New Brunswick, New Jersey, USA, March 19-21, 2001*, volume 63 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 105–112. DIMACS/AMS, 2001. `doi:10.1090/dimacs/063/08`.

[17] Pavol Hell and Jaroslav Nešetřil. *Graphs and homomorphisms*, volume 28 of *Oxford lecture series in mathematics and its applications*. Oxford University Press, 2004.

[18] Pavol Hell and Jaroslav Nešetřil. Colouring, constraint satisfaction, and complexity. *Comput. Sci. Rev.*, 2(3):143–163, 2008. `doi:10.1016/j.cosrev.2008.10.003`.

[19] Pavol Hell and Jaroslav Nešetřil. In praise of homomorphisms. *Comput. Sci. Rev.*, 40:100352, 2021. `doi:10.1016/j.cosrev.2020.100352`.

[20] Ioannis Katsikarelis, Michael Lampis, and Vangelis Th. Paschos. Structural parameters, tight bounds, and approximation for (k, r)-center. *Discret. Appl. Math.*, 264:90–117, 2019. `doi:10.1016/j.dam.2018.11.002`.

[21] Michael Lampis. Finer tight bounds for coloring on clique-width. *SIAM J. Discret. Math.*, 34(3):1538–1558, 2020. `doi:10.1137/19M1280326`.

[22] Daniel Lokshtanov, Dániel Marx, and Saket Saurabh. Known algorithms on graphs of bounded treewidth are probably optimal. *ACM Trans. Algorithms*, 14(2):13:1–13:30, 2018. `doi:10.1145/3170442`.

[23] Dániel Marx, Govind S. Sankar, and Philipp Schepper. Degrees and gaps: Tight complexity results of general factor problems parameterized by treewidth and cutwidth. In Nikhil Bansal, Emanuela Merelli, and James Worrell, editors, *48th International Colloquium on Automata, Languages, and Programming, ICALP 2021, July 12-16, 2021, Glasgow, Scotland (Virtual Conference)*, volume 198 of *LIPIcs*, pages 95:1–95:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. `doi:10.4230/LIPIcs.ICALP.2021.95`.

[24] Hermann A. Maurer, Arto Salomaa, and Derick Wood. Colorings and interpretations: a connection between graphs and grammar forms. *Discret. Appl. Math.*, 3(2):119–135, 1981. `doi:10.1016/0166-218X(81)90037-8`.

[25] Hermann A. Maurer, Ivan Hal Sudborough, and Emo Welzl. On the complexity of the general coloring problem. *Inf. Control.*, 51(2):128–145, 1981. `doi:10.1016/S0019-9958(81)90226-6`.

[26] Karolina Okrasa, Marta Piecyk, and Paweł Rzążewski. Full complexity classification of the list homomorphism problem for bounded-treewidth graphs. In Fabrizio Grandoni, Grzegorz Herman, and Peter Sanders, editors, *28th Annual European Symposium on Algorithms, ESA 2020, September 7-9, 2020, Pisa, Italy (Virtual Conference)*, volume 173 of *LIPIcs*, pages 74:1–74:24. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. `doi:10.4230/LIPIcs.ESA.2020.74`.

[27] Karolina Okrasa and Paweł Rzążewski. Fine-grained complexity of the graph homomorphism problem for bounded-treewidth graphs. *SIAM J. Comput.*, 50(2):487–508, 2021. `doi:10.1137/20M1320146`.

[28] Mark H. Siggers. A new proof of the H-coloring dichotomy. *SIAM Journal on Discrete Mathematics*, 23(4):2204–2210, 2010. `arXiv:https://doi.org/10.1137/080736697`, `doi:10.1137/080736697`.

[29] Leslie G. Valiant. The complexity of enumeration and reliability problems. *SIAM J. Comput.*, 8(3):410–421, 1979. `doi:10.1137/0208032`.