

# Current Algorithms for Detecting Subgraphs of Bounded Treewidth are Probably Optimal

Karl Bringmann ✉

Saarland University, Saarland Informatics Campus, Germany

Max Planck Institute for Informatics, Saarland Informatics Campus, Germany

Jasper Slusallek ✉

Saarland University, Saarland Informatics Campus, Germany

---

## Abstract

---

The Subgraph Isomorphism problem is of considerable importance in computer science. We examine the problem when the pattern graph  $H$  is of bounded treewidth, as occurs in a variety of applications. This problem has a well-known algorithm via color-coding that runs in time  $O(n^{\text{tw}(H)+1})$  [Alon, Yuster, Zwick'95], where  $n$  is the number of vertices of the host graph  $G$ . While there are pattern graphs known for which Subgraph Isomorphism can be solved in an improved running time of  $O(n^{\text{tw}(H)+1-\varepsilon})$  or even faster (e.g. for  $k$ -cliques), it is not known whether such improvements are possible for all patterns. The only known lower bound rules out time  $n^{o(\text{tw}(H)/\log(\text{tw}(H)))}$  for any class of patterns of unbounded treewidth assuming the Exponential Time Hypothesis [Marx'07].

In this paper, we demonstrate the existence of maximally hard pattern graphs  $H$  that require time  $n^{\text{tw}(H)+1-o(1)}$ . Specifically, under the Strong Exponential Time Hypothesis (SETH), a standard assumption from fine-grained complexity theory, we prove the following asymptotic statement for large treewidth  $t$ :

For any  $\varepsilon > 0$  there exists  $t \geq 3$  and a pattern graph  $H$  of treewidth  $t$  such that Subgraph Isomorphism on pattern  $H$  has no algorithm running in time  $O(n^{t+1-\varepsilon})$ .

Under the more recent 3-uniform Hyperclique hypothesis, we even obtain tight lower bounds for each specific treewidth  $t \geq 3$ :

For any  $t \geq 3$  there exists a pattern graph  $H$  of treewidth  $t$  such that for any  $\varepsilon > 0$  Subgraph Isomorphism on pattern  $H$  has no algorithm running in time  $O(n^{t+1-\varepsilon})$ .

In addition to these main results, we explore **(1)** colored and uncolored problem variants (and why they are equivalent for most cases), **(2)** Subgraph Isomorphism for  $\text{tw} < 3$ , **(3)** Subgraph Isomorphism parameterized by pathwidth instead of treewidth, and **(4)** a weighted variant that we call Exact Weight Subgraph Isomorphism, for which we examine pseudo-polynomial time algorithms. For many of these settings we obtain similarly tight upper and lower bounds.

**2012 ACM Subject Classification** Theory of computation → Design and analysis of algorithms; Theory of computation → Computational complexity and cryptography

**Keywords and phrases** subgraph isomorphism, treewidth, fine-grained complexity, hyperclique

**Funding** *Karl Bringmann*: This work is part of the project TIPEA that has received funding from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (grant agreement No. 850979).

## 1 Introduction

The SUBGRAPH ISOMORPHISM problem is commonly defined as follows: Given a graph  $H$  on  $k$  vertices, and a graph  $G$  on  $n$  vertices, is there a (not necessarily induced) subgraph of  $G$  which is isomorphic to  $H$ ?

SUBGRAPH ISOMORPHISM generalizes many problems of independent interest, such as the  $k$ -PATH and  $k$ -CLIQUE problems. The problem is also of considerable interest when  $H$  is less structured, with applications to discovering patterns in graphs that, for example, arise from biological processes such as gene transcription or food networks, from social interaction, from electronic circuits, from neural networks [56], from chemical compounds [66] or from control flow in programs [26]. In some fields, the problem is sometimes referred to as the search for “network motifs”, i.e. subgraphs that appear more often than would normally be expected.

In its general form, the problem is NP-hard. We are interested in solving the problem when the pattern graph  $H$  is “tree-like” or “path-like”, i.e. when the treewidth  $\text{tw}(H)$  or the pathwidth  $\text{pw}(H)$  of  $H$  is bounded. Such pattern graphs of low treewidth or pathwidth often arise in practice when considering the structure of chemical compounds, the control flow of programs, syntactic relations in natural language, or many other graphs from practical applications (see e.g. [18, 21]). On the theoretical side, many restricted classes of graphs have bounded treewidth, see also [20]. Restricting NP-hard problems to graphs of bounded tree- and pathwidth often yields polynomial-time algorithms, and SUBGRAPH ISOMORPHISM is no exception. Most notably, the classic Color-Coding algorithm by Alon, Yuster and Zwick [12] solves the problem by a Las Vegas algorithm in expected time  $O(n^{\text{tw}(H)+1}g(k))$ , or by a deterministic algorithm in time  $\tilde{O}(n^{\text{tw}(H)+1}g(k))$ , where  $g$  is a computable function (and  $\tilde{O}(\cdot)$  is used to suppress factors that are polylogarithmic in the input size). In other words, if the pattern graph  $H$  has treewidth bounded by some constant, the problem is fixed-parameter tractable when parameterized by  $k$ . The Color-Coding algorithm is also relevant for practical purposes: Recently, it has received an efficient implementation, which tested well against state-of-the-art programs for SUBGRAPH ISOMORPHISM [53].

Many researchers wondered whether the Color-Coding algorithm can be improved. This question has been studied in many different directions, including the following:

- Marx [54] showed that no algorithm solves the SUBGRAPH ISOMORPHISM problem in time  $O(n^{o(\text{tw}(H)/\log(\text{tw}(H)))}g(k))$  unless the Exponential Time Hypothesis (ETH) fails, and this even holds when restricted to any class of pattern graphs of unbounded treewidth.
- A series of work has improved the computable function  $g$ , see e.g. [13, 40, 58].
- For many special pattern graphs faster algorithms have been found; the most famous example is the  $k$ -Clique problem, which can be solved in time  $O(n^{k\omega/3}g(k))$  [57]<sup>1</sup>.

In this paper, we use a different angle to approach the question whether Color-Coding can be improved. We ask whether there exist “hard” pattern graphs:

*Do there exist pattern graphs  $H$  for which SUBGRAPH ISOMORPHISM cannot be solved in time  $O(n^{\text{tw}(H)+1-\varepsilon})$  for any constant  $\varepsilon > 0$ ?*

To the best of our knowledge, this question has not been previously studied. As our main result, we (conditionally) give a positive answer to this question. More precisely, we show that for every  $t \geq 3$  there exists a pattern graph  $H$  with  $\text{tw}(H) = t$  for which SUBGRAPH ISOMORPHISM cannot be solved in time  $O(n^{\text{tw}(H)+1-\varepsilon})$  for any constant  $\varepsilon > 0$ , assuming

---

<sup>1</sup> This bound assumes that  $k$  is divisible by 3; there are similar results for general  $k$  [37].

the 3-uniform  $k$ -Hyperclique hypothesis; see Section 1.3 for details on this hypothesis. We also show a slightly weaker statement under the Strong Exponential Time Hypothesis. This conditionally shows that the Color-Coding algorithm by Alon, Yuster and Zwick cannot be significantly improved while still working for all pattern graphs.

For the case of  $\text{tw}(H) = 2$ , an algorithm of Curticapean, Dell and Marx [32] can be adapted such that it solves SUBGRAPH ISOMORPHISM in time  $\tilde{O}(n^\omega g(k))$ . We unify this with the algorithm of Alon, Yuster and Zwick by showing that both time bounds can be achieved within a simple framework. In particular, we use so-called  $k$ -wise matrix products, an operation which was introduced in its general form in [42] and studied further in [51].

We also study the SUBGRAPH ISOMORPHISM problem when the pathwidth of  $H$  is bounded, and specialize our framework to show slight improvements in running time compared to the case of bounded treewidth. Here, we use rectangular matrix products, for which faster-than-naive algorithms are known [41].

In further results, our focus is on the weighted variant EXACT WEIGHT SUBGRAPH ISOMORPHISM, where the subgraph must also have total weight equal to zero. In this work, we consider both the node-weighted and the edge-weighted variant of this problem, for both bounded treewidth and bounded pathwidth, allowing the maximum absolute weight  $W$  to appear in the running time (i.e. the pseudopolynomial-time setting). We show that our algorithms for the unweighted case can be adapted to the weighted case. We also speed up the weighted algorithms by using the fact that fast convolution (or rather, sumset computation), a folklore technique that lies at the core of many fast algorithms for problems with weights (e.g. [27, 24, 48, 46, 22, 15] and [31, exercise 30.1.7]), can easily be adapted to work with rectangular matrices and tensors. We furthermore show tight conditional lower bounds in many cases. Last but not least, we show that our algorithms can be slightly improved for the case of node-weighted instances for which either the pathwidth of  $H$  is bounded, or  $H$  is a tree. These algorithms also rely on fast rectangular matrix products.

## 1.1 Related Work

Additional to the conditional lower bound of  $O(n^{o(\text{tw}(H)/\log(\text{tw}(H)))}g(k))$  by Marx [54], there is an unconditional lower bound of  $O(n^{\kappa(H)})$  for the size of any  $\text{AC}^0$ -circuit, for some graph parameter  $\kappa(H) = \Omega(\text{tw}(H)/\log(\text{tw}(H)))$ , which holds even when considering the average case [50]. Interestingly, the factor of  $1/\log(\text{tw}(H))$  does not seem to be an artefact of the proof: There is an  $\text{AC}^0$ -circuit of size  $O(n^{o(\text{tw}(H))}g(k))$  that solves the problem on certain unbounded-treewidth classes in the average case [62].

In a different direction, Dalirrooyfard et al. [35] design various reductions from  $k$ -Clique to SUBGRAPH ISOMORPHISM, among other results. They also present results on the detection of induced subgraphs (we focus on non-induced subgraphs).

For the weighted variant of SUBGRAPH ISOMORPHISM, lower bounds under the  $k$ -SUM hypothesis for stars, paths, cycles and some other pattern graphs are presented in [6]. Edge-weighted triangle detection has a by-now classic  $O(n^{3-\epsilon})$  lower bound under both the 3SUM hypothesis and the APSP hypothesis [10]. On the other hand, in [7], it is proven that finding node-weighted  $k$ -cliques can be done almost as quickly as finding unweighted  $k$ -cliques. We are not aware of any results on the EXACT WEIGHT SUBGRAPH ISOMORPHISM problem when  $W$  may appear in the running time (i.e. a pseudopolynomial-time algorithm), which is what we focus on here.

In our work, we pose no restrictions on the host graph  $G$ . For an extensive classification of SUBGRAPH ISOMORPHISM with respect to various parameters of both  $G$  and  $H$ , see [55].

## 1.2 Hardness Assumptions

The most standard hypothesis from fine-grained complexity theory is the Strong Exponential Time Hypothesis (SETH) [43], which postulates that for any  $\varepsilon > 0$  there exists  $k \geq 3$  such that  $k$ -SAT on  $n$  variables cannot be solved in time  $O^*(2^{(1-\varepsilon)n})$ .

More recent is the Hyperclique hypothesis. In the  $h$ -uniform  $k$ -HYPERCLIQUE problem, for a given  $h$ -uniform hypergraph we want to decide whether there exist a set of  $k$  vertices such that every size- $h$  subset of these vertices forms a hyperedge. For any  $k > 3$ , the 3-uniform  $k$ -Hyperclique hypothesis postulates that this problem cannot be solved in time  $O(n^{k-\varepsilon})$  for any  $\varepsilon > 0$ . This hypothesis has also been formulated when replacing 3 with any  $h < k$ , getting progressively more believable with larger  $h$ . For a more in-depth discussion of the believability of this hypothesis we refer to [51, Section 7].

Note that we will also use the  $h$ -uniform Hyperclique hypothesis for various  $h$ , which is simply the conjecture that the  $h$ -uniform  $k$ -Hyperclique hypothesis is true for all  $k > h$ .

Related to this is the  $k$ -Clique conjecture, which postulates that the  $k$ -Clique problem (which is the 2-uniform  $k$ -Hyperclique problem) cannot be solved in time  $O(n^{\omega k/3-\varepsilon})$  for any constant  $\varepsilon > 0$ , where  $\omega < 2.373$  [49] is the exponent of matrix multiplication.

## 1.3 Our Results

**Unweighted Subgraph Isomorphism with Bounded Treewidth** First, consider the case of the unweighted SUBGRAPH ISOMORPHISM problem for bounded-treewidth pattern graphs  $H$ . As was said, and as we will re-prove with a unified algorithm later, this problem has an algorithm running in time  $\tilde{O}(n^{\text{tw}(H)+1})$  for  $\text{tw}(H) \geq 3$ . We show tight conditional lower bounds by proving the following obstacles to faster algorithms, which use the  $k$ -clique hypothesis and the  $h$ -uniform  $k$ -hyperclique hypothesis. Note that when we say, for some  $x$ , that an algorithm has running time  $O(n^{x-\varepsilon})$ , what we mean is that the algorithm runs in time  $O(n^{x-\varepsilon})$  for some constant  $\varepsilon > 0$ .

► **Theorem 1.** *The following statements are true.*

1. *For each  $t \geq 3$  and each  $3 \leq h \leq t$ , there exists a connected, bipartite pattern graph  $\mathcal{H}_{t,h}$  of treewidth  $t$  such that there cannot be an algorithm solving the SUBGRAPH ISOMORPHISM problem on pattern graph  $\mathcal{H}_{t,h}$  in time  $O(n^{t+1-\varepsilon})$  unless the  $h$ -uniform  $h(t+1)$ -hyperclique hypothesis fails.*
2. *For each  $t \geq 2$  and each  $h \geq 3$ , there exists a connected, bipartite pattern graph  $\mathcal{H}_{t,h}$  of treewidth  $t$  such that there cannot be an algorithm solving the SUBGRAPH ISOMORPHISM problem on pattern graph  $\mathcal{H}_{t,h}$  in time  $O(n^{t-\varepsilon})$  unless the  $h$ -uniform  $ht$ -hyperclique hypothesis fails.*
3. *For each  $t \geq 2$ , there exists a connected, bipartite pattern graph  $\mathcal{H}_t$  of treewidth  $t$  such that there cannot be an algorithm solving the SUBGRAPH ISOMORPHISM problem on pattern graph  $\mathcal{H}_t$  in time  $O(n^{(t+1)\omega/3-\varepsilon})$  unless the  $(t+1)$ -CLIQUE hypothesis fails.*

Indeed, with the very same reduction, we also get an obstacle from SETH. However, the lower bound it provides is not as tight as the above, and in the case of the second part does not work for each target treewidth  $t$ .

► **Theorem 2.** *Assuming SETH, the following two statements are true.*

1. *For any  $t \geq 3$  and any  $\varepsilon > 0$  there exists a pattern graph  $\mathcal{H}_{t,\varepsilon}$  of treewidth  $t$  such that there cannot be an algorithm solving all instances of SUBGRAPH ISOMORPHISM with pattern graph  $\mathcal{H}_{t,\varepsilon}$  in time  $O(n^{t-\varepsilon})$ .*

2. For any  $\varepsilon > 0$  there exists a  $t \geq 3$  and a pattern graph  $\mathcal{H}_\varepsilon$  of treewidth  $t$  such that there cannot be an algorithm solving all instances of SUBGRAPH ISOMORPHISM with pattern graph  $\mathcal{H}_\varepsilon$  in time  $O(n^{t+1-\varepsilon})$ .

On the algorithmic side, we present an algorithm that achieves matching running times (as listed in theorem 3 below). As was said, the results in the following theorem are not new. Part 1 was shown via Color-Coding in [12] and part 2 follows from techniques in [32]. We unify these two results by providing a single, relatively simple algorithmic technique achieving both, based on  $k$ -wise matrix products. These techniques are later expanded to also work for the weighted version, where they then achieve new results. In the following,  $\omega < 2.373$  [49] is the exponent of matrix multiplication.

► **Theorem 3.** *There are algorithms which, given an arbitrary instance  $\phi = (H, G)$  of SUBGRAPH ISOMORPHISM where  $H$  has treewidth  $\text{tw}(H)$ , solve  $\phi$  in*

1. time  $\tilde{O}(n^{\text{tw}(H)+1}g(k))$  when  $\text{tw}(H) \geq 3$ ,
2. time  $\tilde{O}(n^\omega g(k))$  when  $\text{tw}(H) = 2$ , and
3. time  $\tilde{O}(n^2 g(k))$  when  $\text{tw}(H) = 1$ ,

where  $k := |V(H)|$ ,  $n := |V(G)|$  and  $g$  is a computable function.

**Semi-Equivalence of Hyperclique and Subgraph Isomorphism** We also discuss how our results not only show a reduction from HYPERCLIQUE to SUBGRAPH ISOMORPHISM with bounded treewidth, but also in the other direction. For this, we show that calculating the boolean  $k$ -wise matrix products, which is the bottleneck in our algorithm for bounded-treewidth SUBGRAPH ISOMORPHISM, is actually equivalent to the  $k$ -uniform  $(k + 1)$ -HYPERGRAPH problem. Hence we have a reduction in the second direction. This gives an interesting intuition for why the Hyperclique hypothesis is the “correct” conjecture to prove conditional hardness of SUBGRAPH ISOMORPHISM for bounded treewidth.

We remark that this does not lead to a full equivalence of these problems because the uniformity (i.e. the size of hyperedges) of the HYPERCLIQUE problem we reduce from in the first reduction is much smaller than the size of the hypercliques we search for. Hence we only have a reduction from a HYPERCLIQUE instance with small edge uniformity to SUBGRAPH ISOMORPHISM, and a reduction from SUBGRAPH ISOMORPHISM to HYPERCLIQUE instances with large edge uniformity.

**Weighted Subgraph Isomorphism with Bounded Treewidth** Now consider the weighted version of SUBGRAPH ISOMORPHISM for bounded-treewidth graphs  $H$ . Recall that the weighted version can be either node- or edge-weighted and is defined such that the weights in the solution subgraph must have total weight zero. A trivial dynamic programming algorithm on the tree decomposition achieves a running time of  $\tilde{O}(n^{\text{tw}(H)+1} \cdot W \log W)$  for  $\text{tw}(H) \geq 3$ .

Note that these results show conditional lower bounds even when the maximum weight is restricted to  $W = \Theta(n^\gamma)$ , for any constant  $\gamma > 0$ .

► **Theorem 4.** *For both the node- and edge weighted variant of the problems, the following statements are true.*

1. For each  $t \geq 3$ , each  $\gamma \in \mathbb{R}^+$  and each  $3 \leq h \leq t$ , there exists a connected, bipartite graph  $\mathcal{H}_{t,h,\gamma}$  of treewidth  $t$  such that there cannot be an algorithm solving the EXACT WEIGHT SUBGRAPH ISOMORPHISM problem on pattern graph  $\mathcal{H}_{t,h,\gamma}$  for instances with maximum weight  $W = \Theta(n^\gamma)$  in time  $O(n^{t+1-\varepsilon}W)$ , unless the  $h$ -uniform Hyperclique hypothesis fails.

2. For each  $t \geq 1$ , each  $\gamma \in \mathbb{R}^+$  and each  $h \geq 3$ , there exists a connected, bipartite graph  $\mathcal{H}_{t,h,\gamma}$  of treewidth  $t$  such that there cannot be an algorithm solving the EXACT WEIGHT SUBGRAPH ISOMORPHISM problem on pattern graph  $\mathcal{H}_{t,h,\gamma}$  for instances with maximum weight  $W = \Theta(n^\gamma)$  in time  $O(n^{t-\varepsilon}W)$ , unless the  $h$ -uniform Hyperclique hypothesis fails.
3. For each  $t \geq 1$  and each  $\gamma \in \mathbb{R}^+$ , there exists a connected, bipartite graph  $\mathcal{H}_{t,\gamma}$  of treewidth  $t$  such that there cannot be an algorithm solving the EXACT WEIGHT SUBGRAPH ISOMORPHISM problem on pattern graph  $\mathcal{H}_{t,\gamma}$  for instances with maximum weight  $W = \Theta(n^\gamma)$  in time  $O(n^{(t+1)\omega/3-\varepsilon}W^{\omega/3})$ , unless the CLIQUE hypothesis fails.

Similar lower bounds also hold when trying to reduce the exponent of  $W$  instead of  $n$ . Meaning there is also no algorithm of running time  $O(n^{t+1}W^{1-\varepsilon})$  in part 1, etc.

On the algorithmic side, we present an algorithm that achieves matching running times for  $\text{tw}(H) \geq 3$ , and almost matching running times for  $\text{tw}(H) = 1, 2$ . Note that in terms of exponents, the first algorithm below is not better than the naive one with running time  $O(n^{\text{tw}+1}W \log W)$ . However, it avoids a factor of  $\log W$  in the largest term, and instead appends it to a smaller term, so in a way it presents an improvement of  $\log W$  in the running time. Specifically, we show

► **Theorem 5.** *There are algorithms which, given an arbitrary instance  $\phi = (H, G, w)$  of the EXACT WEIGHT SUBGRAPH ISOMORPHISM problem where  $H$  has treewidth  $\text{tw}(H)$ , solve  $\phi$  in*

1. time  $\tilde{O}((n^{\text{tw}(H)+1}W + n^{\text{tw}(H)}W \log W)g(k))$  when  $\text{tw}(H) \geq 3$ ,
2. time  $\tilde{O}((n^\omega W + n^2W \log W)g(k))$  when  $\text{tw}(H) = 2$ , or
3. time  $\tilde{O}((n^2W + nW \log W)g(k))$  when  $\text{tw}(H) = 1$ ,

where  $n := |V(G)|$ ,  $k := |V(H)|$ ,  $g$  is a computable function, and  $W$  is the maximum absolute weight in the image of  $w$ .

Comparing these upper bounds with the lower bounds from Theorem 5, we have a tight lower bound for the weighted case with  $\text{tw}(H) \geq 3$ . For weighted  $\text{tw}(H) = 2$ , we have a lower bound which is tight except for the exponent of  $\omega/3$  to  $W$ ; it is unclear whether this can be strengthened. The lower bound for weighted graphs with  $\text{tw}(H) = 1$  is obviously not tight: We have an upper bound of  $\tilde{O}(n^2W + nW \log W)$ , but our lower bounds only states that it requires time  $O(n^{1-o(1)}W^{1-o(1)})$  and  $O(n^{2\omega/3-o(1)}W^{\omega/3-o(1)})$ . Tighter lower bounds for this case remain an important open problem.

**Unweighted Subgraph Isomorphism with Bounded Pathwidth** So far we have only looked at the case of bounded treewidth. However, similar results hold for the case of bounded pathwidth. Let us start with the unweighted SUBGRAPH ISOMORPHISM problem.

Note that we do not get any lower bounds for the current setting. This is because we prove all our lower bounds by showing an equivalence of the standard Subgraph Isomorphism problem to a colored variant (see also Section 3.3), and then proving a lower bound for the colored version. We do not know how to prove such an equivalence for the current setting, therefore we do not get lower bounds in this case; we leave this as an open problem.

Since a path decomposition is always also a tree decomposition, we trivially get upper bounds as in theorem 3 (when replacing treewidth by pathwidth). However, we can do better by using rectangular matrix multiplication to speed up the computation. For  $z \in \mathbb{R}^+$ , let  $\omega(z)$  be the smallest real number such that multiplying a  $n \times n$  matrix with a  $n \times n^z$  matrix can be done in time  $O(n^{\omega(z)})$ , see Section 5.5 for discussion of this value. We prove the following upper bounds.

► **Theorem 6.** *There are algorithms which, given an arbitrary instance  $\phi = (H, G)$  of SUBGRAPH ISOMORPHISM where  $H$  has pathwidth  $p$ , solve  $\phi$  in*

1. *time  $\tilde{O}(n^{\omega(p-1)}g(k))$  when  $p \geq 2$ , and*
2. *time  $\tilde{O}(n^2g(k))$  when  $p = 1$*

*where  $k := |V(H)|$  and  $n := |V(G)|$ .*

We certainly have  $p \leq \omega(p-1) < p+1$ , so these results represent only a minor improvement, which is nonetheless important because it “beats” the lower bound for treewidth. Hence the lower bound for pathwidth cannot be the same as for treewidth.

**Weighted Subgraph Isomorphism with Bounded Pathwidth** We also analyze the bounded-pathwidth pattern graph version of WEIGHTED SUBGRAPH ISOMORPHISM. Specifically, we get the following lower bound.

► **Theorem 7** (Theorem 4 for pathwidth). *Parts 2 and 3 of Theorem 4 also hold when replacing the treewidth  $t$  by the pathwidth  $p$ . Part 1 does not hold.*

And on the algorithmic side, we can again use rectangular matrix multiplication to improve on the algorithms from the case of bounded treewidth. Specifically, we get:

► **Theorem 8.** *There are algorithms which, given an arbitrary instance  $\phi = (H, G, w)$  of the EXACT WEIGHT SUBGRAPH ISOMORPHISM problem, solve  $\phi$  in*

1. *time  $\tilde{O}((n^{\omega(\text{pw}(H)-1)}W + n^{\text{pw}(H)}W \log W)g(k))$  when  $\text{pw}(H) \geq 2$ ,*
2. *time  $\tilde{O}((n^2W + nW \log W)g(k))$  when  $\text{pw}(H) = 1$*

*where  $n := |V(G)|$ ,  $k := |V(H)|$ , and  $W$  is the maximum absolute weight in the image of  $w$ .*

For  $\text{pw}(H) \geq 3$ , the lower bounds are therefore obviously not tight (at least for current algorithms), unless significant advances in matrix multiplication techniques are made. For  $\text{pw}(H) = 1, 2$ , the situation is the same as with treewidth, see the discussion of Theorem 5.

**Improvements to Special Cases of Weighted Subgraph Isomorphism** It is natural to think that the exponents  $\frac{\omega}{3}$  to  $W$  in the lower bounds of Theorems 4 and 7 are only artefacts of the reduction, and that with more advanced methods, this exponent can be improved to 1. However, the following two theorems show that this notion is false for  $\text{tw}(H) = 1$  and  $\text{pw}(H) = 1, 2$ , at least when considering the node-weighted case. Indeed, for  $\text{tw}(H) = 1$  (or  $\text{pw}(H) = 1$ ) and  $W = n$ , these bounds are tight, so further general improvements on the exponent are impossible.

Specifically, Theorems 9 and 10 show the following improvements of the algorithms from Theorems 5 and 8 for small tree- or pathwidth. Let  $\text{MM}(n, n, x)$  be the time in which a  $n \times n$  matrix can be multiplied with with a  $n \times x$  matrix.

► **Theorem 9.** *There is an algorithm which, given an arbitrary instance  $\phi = (H, G, w)$  of the node-weighted EXACT WEIGHT SUBGRAPH ISOMORPHISM problem where  $H$  is a tree, solves  $\phi$  in time  $\tilde{O}((\text{MM}(n, n, W) + nW \log W)g(k))$ .*

► **Theorem 10.** *There is an algorithms which, given an arbitrary instance  $\phi = (H, G, w)$  of the node-weighted EXACT WEIGHT SUBGRAPH ISOMORPHISM problem, solves  $\phi$  in time  $\tilde{O}(\text{MM}(n, n, n^{\text{pw}(H)-1}W)g(k))$ .*

For  $W = O(n^\gamma)$ , the running time of Theorem 9 is  $\tilde{O}(n^{\omega(\gamma)} \text{poly}(k))$ . Using results from [41](discussed further in section 5.5), this implies several interesting facts. First, for  $\gamma < 0.31$ , the node-weighted problem on trees can be solved in time  $\tilde{O}(n^2 \text{poly}(k))$ , meaning

it can be solved in the same running time as the unweighted case. In particular, this applies to  $W = \text{polylog}(n)$  or  $W = O(\sqrt[4]{n})$ .

Second, for arbitrary  $\gamma$  we now have a running time of  $\tilde{O}(n^{\omega(\gamma)-\gamma}W \text{poly}(k))$ . Trivially, for any  $\gamma > 0$ ,  $\omega(\gamma) - \gamma < 2$ . Indeed, for  $\gamma \geq 5$  we have  $\omega(\gamma) - \gamma < 1.16$  by [41]. This shows that for node weighted trees, there cannot be a lower bound of  $n^{1.16}W$ , let alone  $n^2W$ . Indeed, it is known that  $\lim_{\gamma \rightarrow \infty} \omega(\gamma) - \gamma = 1$  [29], which implies that when restricting to instances where  $W = \Theta(n^\gamma)$ , there cannot be a lower bound of  $n^{1+\varepsilon}W$  for any  $\varepsilon > 0$  that holds for any constant  $\gamma > 0$ . Similar results hold for bounded-pathwidth graphs.

### Structure of the Paper

We begin by giving a simplified view of our two main proofs in Section 2. We then give formal definitions of key concepts in Section 3 before diving into the full proofs. In Section 4, we prove all lower bounds. In Section 5, we prove all upper bounds, including improved algorithms for node-weighted graphs. In Section 6, we prove the semi-equivalence between HYPERCLIQUE and SUBGRAPH ISOMORPHISM on bounded treewidth graphs. Finally, in Section 7, we prove the equivalences between the colored and uncolored problems which we use throughout the paper. We conclude by stating several important open problems in Section 8.

## 2 Technical Overview of Our Main Results

We now give proof sketches of our two main lower bound results to present the main ideas of the proofs without giving too much detail. The full details are available in Section 4.

### 2.1 Lower Bound for Subgraph Isomorphism

Our main result is the existence of the hard pattern graphs for bounded-treewidth SUBGRAPH ISOMORPHISM. We now prove their existence for treewidth at least 3 under the Hyperclique hypothesis, i.e. part 1 of Theorem 1.

The exact statement we prove is that for each  $t \geq 3$  and each  $3 \leq h \leq t$ , there exists a pattern graph of treewidth  $t$  such that SUBGRAPH ISOMORPHISM cannot be solved in time  $O(n^{t+1-\varepsilon})$  on that pattern graph unless the  $h$ -uniform  $h(t+1)$ -hyperclique hypothesis fails. Note that the proof actually shows this for the colored variant of SUBGRAPH ISOMORPHISM, after which we can use Lemma 14 to transfer the lower bound to the uncolored problem.


**Proof sketch.** See Figure 1 for a sketch of the reduction. Let  $t \geq 3$  and  $3 \leq h \leq t$  be given and assume that SUBGRAPH ISOMORPHISM can be solved in time  $O(n^{t+1-\varepsilon})$  on pattern graphs of treewidth  $t$ . We show that the  $h$ -uniform  $h(t+1)$ -hyperclique hypothesis fails.

**Construction of  $H$**  We construct a pattern graph  $H$  as a bipartite graph with vertex set  $A \cup B$  as follows. Writing  $[c] := \{1, \dots, c\}$ , we set  $A := [t+1]$  and  $B := \binom{A \times [h]}{h}$ . We connect a vertex  $b = ((a_1, j_1), \dots, (a_h, j_h))$  in  $B$  to a vertex  $a$  in  $A$  if  $a = a_\ell$  for some  $\ell$ . Set  $k := |A| + |B|$ .

We show that this pattern has a treewidth of  $t$ , via a well-known characterization of treewidth as a graph-theoretic game: A graph  $F$  has treewidth  $\leq t$  if and only if  $t+1$  cops can catch<sup>2</sup> a robber on  $F$  [63]. To show the bound on the treewidth of  $H$ , initially

<sup>2</sup> The game works as follows: The  $k+1$  cops select their starting vertices in the graph. Then the robber may choose their starting vertex. The cops can always see the robber and adapt their strategy





■ **Figure 1** A sketch of the reduction we use to prove part 1 of Theorem 1 where  $h = 3$  and  $t = 4$ . Note that this is only a partial sketch of the pattern graph. We use multiedges to signify that for the endpoints  $a \in A$  and  $b = ((a_1, j_1), \dots, (a_h, j_h)) \in B$  there exists more than one  $\ell$  such that  $a = a_\ell$ .

place a cop on each vertex of  $A$ . No matter on which vertex of  $B$  the robber starts, they are surrounded by cops. Since every vertex in  $B$  has  $h \leq t < t + 1$  neighbors in  $A$ , there must exist some cop which is not adjacent to the robber, so this cop can catch the robber in a single step. This concludes the proof that the pattern graph  $H$  has treewidth  $t$ .

**Construction of  $G$**  Now given a  $h(t + 1)$ -partite hypergraph  $H'$ , i.e. an instance of the  $h$ -uniform  $k'$ -HYPERCLIQUE problem for  $k' := h(t + 1)$ , we write the vertex set of  $H'$  as  $U_{1,1} \cup \dots \cup U_{1,h} \cup \dots \cup U_{t+1,1} \cup \dots \cup U_{t+1,h}$ . Let  $N_H$  be the number of vertices in each partition and  $n_H = O(N_H)$  the number of vertices overall.

We construct a  $k$ -partite graph  $G$  as follows. For  $a$  in  $A$  we set  $V_a := U_{a,1} \times \dots \times U_{a,h}$ . For  $b = ((a_1, j_1), \dots, (a_h, j_h))$  in  $B$ , we set  $V_b := E(H') \cap (U_{a_1, j_1} \times \dots \times U_{a_h, j_h})$ . This describes the  $k$  parts of the  $k$ -partite vertex set  $V(G)$ . Note that each part has size at most  $N_G := N_H^h$ . Now we construct the edges. For any  $a$  in  $A$  and  $b = ((a_1, j_1), \dots, (a_h, j_h))$  in  $B$  with  $(a, b) \in E(H)$ , consider an arbitrary  $u = (u_1, \dots, u_h)$  in  $V_a$  and  $u' = (u'_1, \dots, u'_h)$  in  $V_b$ . We say that  $u$  and  $u'$  are "compatible" if for every  $\ell$  with  $a_\ell = a$  we have  $u'_\ell = u_{j_\ell}$ ; in this case we connect  $u$  and  $u'$  by an edge. This finishes the construction of  $G$ .

**Correctness** Note that any colored subgraph isomorphism of  $H$  in  $G$  chooses vertices  $v_a$  in

---

accordingly. Similarly, the robber can see the cops. The game now proceeds in steps, where in each step, one of the cops chooses an arbitrary destination vertex and takes off via helicopter in the direction of that vertex. While the cop is travelling, the robber sees where they will land and may now move arbitrarily along edges of the graph, as long as they do not pass through stationary cops. When the robber has finished moving, the cop lands. The cops win if and only if they are guaranteed to catch the robber after a finite number of moves, and lose otherwise.

$V_a$  for all  $a$  in  $A$ . This corresponds to choosing vertices  $u_{i,j}$  in  $U_{i,j}$  for all  $i$  in  $[t+1]$ . Moreover, the edges of a  $h(t+1)$ -hyperclique are in one-to-one correspondence with the set  $B$ . Since for each  $b$  in  $B$  the colored subgraph isomorphism of  $H$  in  $G$  needs to choose a vertex  $v_b$  in  $V_b$ , which corresponds to an edge between certain vertices  $u_{i,j}$ , we indeed check that the chosen vertices  $u_{i,j}$  form an  $h$ -uniform  $h(t+1)$ -hyperclique.

**Running Time** Trivially, the construction time and output size are  $O(N_H^{2h})$  (actually, it is slightly better, but this is not important in this proof sketch), and  $G$  has  $n = O(N_G) = O(N_H^h)$  vertices. Now if we can solve SUBGRAPH ISOMORPHISM in time  $O(n^{t+1-\varepsilon})$ , we solve the  $h$ -uniform  $h(t+1)$ -HYPERCLIQUE instance in time  $O(N_H^{2h} + (N_H^h)^{t+1-\varepsilon}) = O(N_H^{h(t+1)-h\varepsilon}) = O(N_H^{h(t+1)-\varepsilon'}) = O(n^{h(t+1)-\varepsilon'})$ . ◀

This shows part 1 of Theorem 1. Part 2 can be shown by the almost the exact same proof, except that now we choose the size of  $A$  to be  $t$  instead of  $t+1$ , and we start with an  $ht$ -hyperclique instead of an  $h(t+1)$ -hyperclique. It can be seen that in this case, the pattern graph still has treewidth  $t$ . The third part of the theorem can be seen by simply taking an instance of  $(t+1)$ -CLIQUE and subdividing the edges in the obvious way to make the graph bipartite.

Let us also quickly mention how the proof of the slightly weaker bounds under SETH, i.e. Theorem 2, works. The split-and-list technique from [67] allows one to reduce the Satisfiability problem to HYPERCLIQUE. Using this technique, the following result was shown in [51, Lemma 9.1].

► **Lemma 11** ([51]). *Assuming SETH, for any  $\varepsilon > 0$  there exists  $h \geq 3$  such that for all  $k > h$ , the  $h$ -uniform  $k$ -HYPERCLIQUE problem is not in time  $O(n^{k-\varepsilon})$ .*

The SETH result now follows by using essentially the same reduction as above, but we prefix it by the reduction from SAT to HYPERCLIQUE.

## 2.2 Lower Bound for Exact Weight Subgraph Isomorphism

We also give lower bounds for the exact weight variant of the SUBGRAPH ISOMORPHISM problem. In particular, we prove the existence of hard pattern graphs for the bounded-treewidth EXACT WEIGHT SUBGRAPH ISOMORPHISM problem for any polynomial weight bound. We give this result for any treewidth which is at least 3, and under the Hyperclique hypothesis. This is part 1 of Theorem 4.

The exact statement we prove is that for each  $t \geq 3$ ,  $\gamma \in \mathbb{R}^+$  and  $3 \leq h \leq t$ , there exists a pattern graph of treewidth  $t$  such that EXACT WEIGHT SUBGRAPH ISOMORPHISM with maximum weight  $W = \Theta(n^\gamma)$  cannot be solved in time  $O(n^{t+1-\varepsilon}W)$  unless the  $h$ -uniform Hyperclique hypothesis fails. Again, we show this statement for the colored problem and transfer the lower bound via Lemma 14.

To do this, we will encode part of a hyperclique instance in the edges of the EXACT WEIGHT SUBGRAPH ISOMORPHISM problem, and the rest of the instance in the weights. To do the latter, we need to encode certain equality constraints only via weights. This can be done using so-called  $k$ -average free sets<sup>3</sup>, which we define below.

<sup>3</sup> These  $k$ -average-free sets are a tool which are very useful for weighted problems, especially when they have additive elements. Such problems include  $k$ -SUM, SUBSET SUM, BIN PACKING, various scheduling problems, TREE PARTITIONING, MAX-CUT, MAXIMUM/MINIMUM BISECTION, a DOMINATING SET variant with capacities, and similar [4, 5, 7, 14, 39, 44, 36]. Other uses of  $k$ -average-free sets in computer science include constructions in extremal graph theory, see e.g. [1, 2, 11].

► **Definition 12** (*k*-average free sets). A set  $S \subseteq \mathbb{Z}$  is called *k*-average-free if, for any  $s_1, \dots, s_{k'+1} \in S$  with  $k' \leq k$ , we have  $s_1 + \dots + s_{k'} = k' \cdot s_{k'+1}$  if and only if  $s_1 = \dots = s_{k'+1}$ . In other words, the average of  $s_1, \dots, s_{k'} \in S$  is in  $S$  if and only if all  $s_i$  are equal.

We use the following construction for *k*-average free sets, originally proven in [16], modified into a more useful version in [7] and formulated in this form in [4].

► **Lemma 13.** *There exists a universal constant  $c > 0$  such that, for all constants  $\varepsilon \in (0, 1)$  and  $k \geq 2$ , a *k*-average-free set  $S$  of size  $n$  with  $S \subseteq [0, k^{c/\varepsilon} n^{1+\varepsilon}]$  can be constructed in time  $\text{poly}(n)$ .*

Let us now prove the statement about EXACT WEIGHT SUBGRAPH ISOMORPHISM. We will construct an instance that is node-weighted, however this can easily be converted into an edge-weighted version by moving the weight of each vertex to all of its incident edges.

**Proof sketch.** Let  $t \geq 3$ ,  $3 \leq h \leq t$  and  $\gamma \in \mathbb{R}$  be given and assume that EXACT WEIGHT SUBGRAPH ISOMORPHISM can be solved in time  $O(n^{t+1-\varepsilon}W)$  on instances where the pattern graph has treewidth  $t$  and all weights are bounded by  $W = \Theta(n^\gamma)$ . We show that the *h*-uniform *k*-hyperclique hypothesis fails for some large enough *k*.

**Construction of H** We construct a pattern graph  $H$  as a graph with vertex set  $(A_1 \cup A_2) \cup B$  as follows. We set  $A_1 := [t+1]$ ,  $A_2 := [r]$  (for some  $r$  large enough) and  $B := \binom{(A_1 \cup A_2) \times [h]}{h}$ . We connect a vertex  $b = ((a_1, j_1), \dots, (a_h, j_h))$  in  $B$  to a vertex  $a$  in  $A_1$  (not in  $A_2$ ) if  $a = a_\ell$  for some  $\ell$ . Set  $k := |A_1| + |A_2| + |B|$ . By almost the same proof as in the unweighted version, it can be shown that this pattern  $H$  has treewidth  $t$ .

**Grouping partitions** Now let an instance of the *h*-uniform *k'*-HYPERCLIQUE problem be given, and write the vertex set of  $H'$  as  $U_1 \cup \dots \cup U_k$ . Let  $N_H$  be the number of vertices in each partition and  $n_H = O(N_H)$  the number of vertices overall.

We construct the *k*-partite graph  $G$  with at most some number  $N_G$  of vertices in each partition as follows. We will encode a  $\beta$ -fraction of the HYPERCLIQUE instance in the weights of the final EXACT WEIGHT SUBGRAPH ISOMORPHISM instance, and a  $(1 - \beta)$ -fraction in the edges, for some  $\beta$  chosen appropriately. To do this, we will choose  $\beta$  such that  $\frac{\beta k'}{hr}, \frac{(1-\beta)k'}{h(t+1)} \in \mathbb{N}$  and then group the sets  $U_1, \dots, U_{\beta k'}$  into  $hr$  groups and the sets  $U_{\beta k'+1}, \dots, U_{k'}$  into  $h(t+1)$  groups. Specifically, for each  $(x, y) \in [r] \times [h]$ , we create the set  $U_{x,y}^1 = U_{(xr+y-1)\frac{\beta k'}{hr}+1} \times \dots \times U_{(xr+y)\frac{\beta k'}{hr}}$ , and for each  $(x, y) \in [t+1] \times [h]$ , we create the set  $U_{x,y}^2 = U_{\beta k'+(x(t+1)+y-1)\frac{(1-\beta)k'}{h(t+1)}+1} \times \dots \times U_{\beta k'+(x(t+1)+y)\frac{(1-\beta)k'}{h(t+1)}}$ .

**Vertices of G** Now for each  $a$  in  $A_1$  we set  $V_a := U_{a,1}^1 \times \dots \times U_{a,h}^1$  and for each  $a$  in  $A_2$  we set  $V_a := U_{a,1}^2 \times \dots \times U_{a,h}^2$ . Finally, for each  $b = ((a_1, j_1), \dots, (a_h, j_h))$  in  $B$ , where for each  $\ell$  we have  $a_\ell \in A_{i_\ell}$ , we set  $V_b := E(H') \cap (U_{a_1, j_1}^{i_1} \times \dots \times U_{a_h, j_h}^{i_h})$ . This describes the  $k$  parts of the *k*-partite vertex set  $V(G)$ . We choose  $r$  large enough so that the maximum size of each part is  $N_G := N_H^{(1-\beta)k'/(t+1)}$ .

**Edges of G** Now we construct the edges and weights of the graph. Let us start with the edges. The construction here is basically the same as the construction of the edges in the unweighted proof in the last section. For each  $a$  in  $A_2$  and  $b = ((a_1, j_1), \dots, (a_h, j_h))$  in  $B$  with  $(a, b)$  in  $E(H)$ , consider an arbitrary  $u = (u_1, \dots, u_h)$  in  $V_a$  and  $u' = (u'_1, \dots, u'_h)$  in  $V_b$ . We say that  $u$  and  $u'$  are “compatible” if for every  $\ell$  with  $a_\ell = a$  we have  $u'_\ell = u_{j_\ell}$ ; in this case we connect  $u$  and  $u'$  by an edge. This finishes the construction of the edges of  $G$ .

**Weights of G** Now we construct the weights. We want to encode the same edge constraints as we just encoded for  $A_2$ , but now for  $A_1$ , and we have to use weights instead of edges.

To do this, we use  $|B|$ -average free sets via the construction of Lemma 28. We simplify the usage in this shortened proof to avoid dealing with too many variables. We use the lemma to obtain in polynomial time (which we will treat as negligible here) a  $|B|$ -average free set  $S$  of size  $N_H^{\beta k'/(hr)}$  such that  $S \subseteq [0, C]$ , where  $C \approx O(N_H^{\beta k'/(hr)})$  (up to a factor of  $(1 + \varepsilon)$  in the exponent, but we will ignore this here for simplicity). From this, we can construct an arbitrary bijection  $\varrho_S : [N_H^{\beta k'/(hr)}] \rightarrow S$ .

To simplify our construction, we specify a target weight  $T$  (instead of the default target zero). We can easily get rid of this again later by subtracting  $T$  from the weights of all vertices of some set of the partition. The binary representation of  $T$  consists of  $hr$  blocks of  $\lceil 2|B|C \rceil$  bits, indexed by pairs  $(i, j) \in [r] \times [h]$ , each containing the binary representation of  $|B|C$ . The block  $(i, j)$  represents the group  $U_{i,j}^1$ . The size of the blocks is large enough to prevent overflow between the blocks. Note that the maximum weight  $W$  now satisfies  $\log_2(W) = \Theta(hr \log_2(2|B|C))$  and hence  $W \approx O(N_H^{\beta k'})$ .

Let us now actually specify the weights of the vertices, beginning with the vertices in  $V_a$  for  $a \in A_1$ . For each  $(i, j) \in [r] \times [h]$ , we relabel the elements of each  $U_{i,j}^1$  as  $\{1, \dots, N_H^{\beta k'/(hr)}\}$ . Now we define the weight of the vertex  $V_a \ni u = (u_1, \dots, u_h)$  to have, for each  $i \in [h]$ , the value  $|B|C - |N(a)| \cdot \varrho_S(u_i)$  in the block  $(a, i)$  of its binary representation. Now we move on to the vertices in  $V_b$  for  $b = ((a_1, j_1), \dots, (a_h, j_h))$ . We define the weight of the vertex  $V_b \ni u' = (u'_1, \dots, u'_h)$  to have, for each  $i \in [h]$  such that  $a_i \in A_1$ , the value  $\varrho_S(u'_i)$  in the block  $(a_i, j_i)$  of its binary representation.

All blocks and vertices which have not been assigned a weight yet are assigned a value of zero. This concludes the construction of  $G$ .

**Correctness** Note that any colored subgraph isomorphism of  $H$  in  $G$  chooses vertices  $v_a$  in  $V_a$  for all  $a$  in  $A_1 \cup A_2$ . This corresponds to choosing vertices  $u_i$  in  $U_i$  for each  $i \in [k']$ . Moreover, the edges of a  $k'$ -hyperclique are in one-to-one correspondence with the set  $B$ . We simply need to show that the choice of hyperclique vertices induced by the choice of vertices in  $A_1 \cup A_2$  agrees with the choice of hyperclique edges induced by the choice of vertices in  $B$ . For the vertices in  $A_2$ , this is easily seen to be ensured by the edges. For the vertices in  $A_1$ , we need to prove that the weights encode the same constraint. This, however, is simply the definition of a  $|B|$ -average free set: Consider the block  $(i, j)$  (where  $(i, j) \in [r] \times [h]$ ) in the binary representation of the total weight of the subgraph. Suppose that for  $i \in A_1$  the vertex  $V_i \ni u = (u_1, \dots, u_h)$  was selected, and that for each  $B \ni b = ((a_1, j_1), \dots, (a_h, j_h))$  with  $\exists \ell : (a_\ell, j_\ell) = (i, j)$  the vertex  $V_b \ni u' = (u'_1, \dots, u'_h)$  was selected. Then by construction, the total value in the block  $(i, j)$  is the value  $|B|C - |N(i)|\varrho_S(u_j)$  (where  $N(i)$  is the neighbourhood of  $i \in A_1$ ), plus the value  $\varrho_S(u'_\ell)$  for all  $b$  as above. Note that the latter term has exactly  $|N(i)|$  summands, hence in order for the value in the block to be equal to  $|B|C$  as specified by the target weight, we must have that the value of  $\varrho_S(u_j)$  is equal to the value of each of the  $\varrho_S(u'_\ell)$  by the definition of  $|B|$ -average free sets. Since  $\varrho_S$  is a bijection, this ensures that the choice of hyperclique vertices in the sets appearing in the Cartesian product defining  $U_{i,j}^1$  – i.e.  $U_{(ir+j-1)\frac{be+ak'}{hr}+1}, \dots, U_{(ir+j)\frac{\beta k'}{hr}}$  – agree with the choice of hyperclique edges. This is true for all  $i, j$  and hence for each  $U_\ell$  for  $\ell \in [k']$ .

The other direction is easy to see via a similar, simpler argument. This concludes the correctness proof.

**Running Time** It can be seen that the running time of this reduction is  $O(N_H^{2h})$ , up to the running time of the algorithm for the construction of the  $|B|$ -average free set, which we will ignore here for sake of simplicity. Now suppose we can solve EXACT WEIGHT SUBGRAPH ISOMORPHISM in time  $O(n^{t+1-\varepsilon}W)$ . We use the reduction above to convert

a HYPERCLIQUE instance with  $n_H = O(N_H)$  nodes to an EXACT WEIGHT SUBGRAPH ISOMORPHISM instance where  $W \approx \Theta(N_H^{\beta k'})$  and  $n = O(N_H^{(1-\beta)k'/(t+1)})$ . Choosing  $\beta$  carefully, we get  $W = \Theta(N_H^\gamma)$ ; note that we are ignoring some intricacies in the choice of  $\beta$  that arise when you consider the running time of the algorithm that constructs the  $|B|$ -average free set – the details are available in Section 4. Now via the algorithm for EXACT WEIGHT SUBGRAPH ISOMORPHISM, we can solve this instance and hence the original HYPERCLIQUE problem in time  $O(N_H^{2h} + N_H^{((1-\beta)k'/(t+1))(t+1-\varepsilon)} \cdot N_H^{\beta k'}) = O(N_H^{k'-\varepsilon'}) = O(n_H^{k'-\varepsilon'})$ . ◀

It is easy to see that the same proof also rules out algorithms running in time  $O(n^{t+1}W^{1-\varepsilon})$ . Similar as with the proof for the unweighted problem, basically the same techniques can be used to prove the other parts of Theorem 4.

### 3 Preliminaries

#### 3.1 General Notation and Nomenclature

We denote by  $\mathbb{N}$  the set of positive integers. For  $p \in \mathbb{N}$ , we use  $[p]$  to denote the set  $\{1, \dots, p\}$ . For a statement or predicate  $P$ , we define the Iverson bracket  $[P]$  as 1 if  $[P]$  is true, and zero otherwise. To declutter notation that relies heavily on the Iverson bracket, we will often use truth values and 0/1 interchangeably, where true will be indicated by 1 and false by 0.

For a function  $f : \mathcal{A} \rightarrow \mathcal{B}$  and a set  $S \subseteq \mathcal{A}$ , we denote with  $f|_S : S \rightarrow \mathcal{B}$  the function  $f$  restricted to  $S$ . That is,  $\forall s \in S : f|_S(s) = f(s)$ . Furthermore, for  $u \notin \mathcal{A}$  and  $v \notin \mathcal{B}$  we define the function extension  $(f \cup \{u \mapsto v\}) : \mathcal{A} \cup \{u\} \rightarrow \mathcal{B} \cup \{v\}$  as

$$(f \cup \{u \mapsto v\})(c) := \begin{cases} v & \text{if } c = u \\ f(c) & \text{otherwise} \end{cases}$$

We use standard notation for graphs. In particular, for a graph  $G$ , we let  $V(G)$  be its set of vertices and  $E(G)$  its set of edges. For a set  $X \subseteq V(G)$ , we denote the induced subgraph by  $G[X]$ . For a vertex  $v \in V(G)$ , we denote its neighbourhood as  $N_G(v)$ , or as  $N(v)$  when  $G$  is clear from context. We denote the treewidth and pathwidth (see Section 3.4) of  $G$  as  $\text{tw}(G)$  and  $\text{pw}(G)$ , respectively. All graphs are, unless otherwise stated, simple, undirected and without self-loops.

We use  $\text{poly}(n)$  to denote functions which are upper-bounded by  $O(n^c)$  for some  $c \in \mathbb{N}$ , and  $\text{polylog}(n)$  to denote functions upper-bounded by  $O(\log^c(n))$  for some  $c \in \mathbb{N}$ . In running times, we use  $O^*(\cdot)$  to suppress factors that are polynomial in the input size, and  $\tilde{O}(\cdot)$  to suppress factors that are polylogarithmic in the input size.

In all weighted problems, we assume without further mention that the target weight or maximum absolute weight is at least 1. This is to avoid special cases with the running time.

#### 3.2 Notation and Nomenclature for Colored Subgraph Isomorphism

We now define some nomenclature for the (EXACT WEIGHT) COLORED SUBGRAPH ISOMORPHISM problem. Not that instead of talking about colors, we will talk about a “color homomorphism”  $f$ . Specifically, the COLORED SUBGRAPH ISOMORPHISM is defined as follows: Given a pattern graph  $H$  and a host graph  $G$  along with a graph homomorphism  $f : V(G) \rightarrow V(H)$ , is it possible to pick a set  $S$  with exactly one vertex from the preimage of each  $v \in V(H)$  such that the subgraph induced by  $S$  is isomorphic to  $H$ ? The homomorphism

$f$  simulates the colors, with all the vertices in a preimage of  $f$  being of equal color (which is unique over all preimages). The EXACT WEIGHT COLORED SUBGRAPH ISOMORPHISM is defined analogously<sup>4</sup>. The weight function is always be denoted by  $w$ .

Fix an instance  $(G, H, f)$  or  $(G, H, f, w)$ .  $H$  is the pattern graph which is to be found in the large graph  $G$  when given color homomorphism  $f : V(G) \rightarrow V(H)$  and weight function  $w$ . For a subset  $I \subseteq V(H)$ , we call a function  $R : I \rightarrow f^{-1}(I)$  a **configuration of  $I$**  if  $\forall v \in I : R(v) \in f^{-1}(v)$ . We define  $\text{Conf}(I) \subseteq (I \rightarrow f^{-1}(I))$  to be the set of configurations of  $I$ . A configuration of  $I$  is called **valid configuration** if  $\forall uv \in E(H[I]) : R(u)R(v) \in E(G)$ . Finally, we call a configuration  $R$  of  $I$  a **partial solution of  $I$  in  $J$** , for some  $I \subseteq J \subseteq V(H)$ , if there is a valid configuration  $S$  of  $J$  such that  $S|_I = R$ . We may shorten this to  $R$  being a **partial solution for  $J$**  if  $I$  is clear from context.

For a valid configuration  $R$  of  $I \subseteq V(H)$ , we call  $w(R)$  its **weight**. The exact definition of the weight  $w(R)$  depends on whether  $G$  is node-weighted or edge-weighted. If  $G$  is node-weighted with weight function  $w : V(G) \rightarrow \mathbb{Z}$ , we define  $w(R) := \sum_{u \in I} w(f^{-1}(u))$ . If it is edge-weighted with weight function  $w : E(G) \rightarrow \mathbb{Z}$ , we define  $w(R) := \sum_{uv \in E(H[I])} w(R(u)R(v))$ , where it is guaranteed that  $R(u)R(v) \in E(G)$  because  $R$  is a valid configuration. Furthermore, we say that a partial solution  $R$  of  $I$  in  $J$  **has an extension of weight  $W'$**  if there is a valid configuration  $S$  of  $J$  such that  $S|_I = R$  and the nodes (respectively: the edges) of  $S$  which are not in  $I$  have combined weight  $W'$ , i.e.  $w_{\text{ext}}(S, R) := w(S) - w(R) = W'$ .

For brevity, we further define the following predicates:

- $\text{ParSol}(R; I; J) = R$  is a partial solution of  $I$  in  $J$
- $\text{ParSolE}(R; I; J; W) = R$  is a partial solution of  $I$  in  $J$  with an extension of weight  $W$
- $\text{ValConf}(R; I) = R$  is a valid configuration of  $I$

### 3.3 Equivalence of the Colored and Uncolored Problems

All mentioned algorithms and conditional lower bounds are shown for the restricted problem of COLORED SUBGRAPH ISOMORPHISM, where the nodes of  $G$  and  $H$  are colored with  $|V(H)|$  colors and the isomorphism must preserve colors, as also studied in [54]. In section 7, we prove that the standard and the colored variant of SUBGRAPH ISOMORPHISM can be solved in essentially the same running time in almost all cases. Specifically, we show the following lemma.

► **Lemma 14.** *Let  $\rho$  be any graph parameter.*

1. *If there is a  $T(n, k, \rho(H))$  time algorithm for COLORED SUBGRAPH ISOMORPHISM, then there is a  $\tilde{O}(T(kn, k, \rho(H))g(k))$  time algorithm for SUBGRAPH ISOMORPHISM, where  $g$  is some computable function.*
2. *If there is a  $T(n, k, \rho(H), W)$  time algorithm for EXACT WEIGHT COLORED SUBGRAPH ISOMORPHISM, then there is a  $\tilde{O}(T(kn, k, \rho(H), W)g(k))$  time algorithm for EXACT WEIGHT SUBGRAPH ISOMORPHISM, where  $g$  is some computable function.*
3. *Let  $\text{tw}(H) \geq 2$ . If there is a  $T(n, k, \text{tw}(H))$  time algorithm for SUBGRAPH ISOMORPHISM, then there is a  $O(T(\text{poly}(k)n, \text{poly}(k), \text{tw}(H)) + \text{poly}(k)n^2)$  time algorithm for COLORED SUBGRAPH ISOMORPHISM.*
4. *If there is a  $T(n, k, \rho(H), W)$  time algorithm for EXACT WEIGHT SUBGRAPH ISOMORPHISM, then there is a  $O(T(2n, 2k, \rho(H), 2^k W) + \text{poly}(k)n^2)$  time algorithm for EXACT WEIGHT COLORED SUBGRAPH ISOMORPHISM.*

<sup>4</sup> Depending on whether the instance is node- or edge-weighted, we require of the solution subgraph that the sum of either its node or its edge weights is zero.

This lemma enables us to prove results for (EXACT WEIGHT) SUBGRAPH ISOMORPHISM while only talking about the more structured colored variants of the problem.

Regarding treewidth, the only case the above lemma does not cover is how to transform an algorithm for unweighted SUBGRAPH ISOMORPHISM to an algorithm for COLORED SUBGRAPH ISOMORPHISM for  $\text{tw}(H) = 1$ . For our purposes, this is not a problem, since SUBGRAPH ISOMORPHISM for trees already has a trivial unconditional lower bound of  $\Omega(n^2)$ , which is tight. Note that for this unconditional lower bound, we must assume that the graph is dense, i.e. has  $\Theta(n^2)$  edges.

The lemma also cannot transform algorithms for unweighted SUBGRAPH ISOMORPHISM to algorithms for COLORED SUBGRAPH ISOMORPHISM for bounded pathwidth. This means that we cannot show the same lower bounds for the unweighted SUBGRAPH ISOMORPHISM problem for bounded pathwidth as we can for the COLORED SUBGRAPH ISOMORPHISM problem. This is a shortcoming of the lemma that we could not fix, and hence we leave it as an open problem whether the lower bounds for the SUBGRAPH ISOMORPHISM problem for bounded pathwidth can even be improved.

### 3.4 Treewidth and Pathwidth

We give a very short introduction to treewidth and pathwidth, and state some auxiliary definitions and notation used throughout the paper. For an thorough introduction to treewidth, pathwidth and their many applications, we refer the reader to [34, Chapter 7].

► **Definition 15** (Tree Decomposition). *Let  $H$  be a graph. A **tree decomposition** of  $H$  is a pair  $\mathcal{T} = (T, \{X_t\}_{t \in V(T)})$  consisting of a tree  $T$  and along with a set of “bags”  $X_t \subseteq V(H)$ , one for each vertex of  $T$ . It must satisfy the following properties:*

(T1)  $\bigcup_{t \in V(T)} X_t = V(H)$

(T2)  $\forall uv \in E(H) : \exists t \in V(T) : \{u, v\} \subseteq X_t$

(T3)  $\forall u \in V(H) : \text{The subgraph induced by } \{t \in V(T) \mid u \in X_t\} \text{ is a connected subtree}$

► **Definition 16** (Treewidth). *Let  $\mathcal{T} = (T, \{X_t\}_{t \in V(T)})$  be a tree decomposition of  $H$ . We define its **width** to be  $\max_{t \in V(T)} |X_t| - 1$ . We define the **treewidth** of  $H$  to be the minimum width of all tree decompositions of  $H$  and denote it as  $\text{tw}(H)$ .*

► **Definition 17** (Path Decomposition). *Let  $H$  be a graph. A **path decomposition** of  $H$  is a tree decomposition where  $T$  is a path.*

► **Definition 18** (Pathwidth). *The **width** of path decompositions is defined as for tree decompositions. The **pathwidth** of  $H$  is defined to be the minimum width of all path decompositions of  $H$ , and is denoted as  $\text{pw}(H)$ .*

A classic algorithm by Bodlaender [19] computes an optimal tree decomposition or path decomposition for an input graph  $H$  in time  $O(f(|\text{tw}(H)|)|V(H)|)$  for a computable function  $f$ . For our purposes, this is almost excessive: For our results, we only need an algorithm which computes an optimal tree decomposition in time  $g(|V(H)|)$ , for some computable function  $g$ .

Clearly, the treewidth of a graph is always smaller than or equal to its pathwidth. It should also be noted that while graphs of treewidth one are exactly the class of tree graphs, graphs of pathwidth one encompass more than just paths. Rather, they are the class of graphs where each connected component is a caterpillar graph, i.e. consists of a single path with arbitrarily many degree-one nodes attached at any node of the path [59]. The latter is vital for our conditional lower bounds for pathwidth one.

For a tree or path decomposition with underlying tree  $T$  and for  $u \in V(T)$ , we define  $T_u$  to be the subtree rooted at  $u$ . The **cone**  $V_u$  is then defined to be  $V_u := \bigcup_{v \in V(T_u)} X_v$ .


**4 Hardness Results**

**4.1 Twin Water Lilies**

We will obtain our lower bounds by reducing HYPERCLIQUE instances to (un)weighted COLORED SUBGRAPH ISOMORPHISM instances, where the pattern graph  $H$  is of a special form defined below. In other words, the pattern graphs below are the “maximally hard” pattern graphs for the SUBGRAPH ISOMORPHISM problem. See Figure 2 for an illustration.

► **Definition 19** (Twin Water Lily). *For any  $h, s_1, s_2$ , we define the graph  $\text{TWL}(h, s_1, s_2)$  as follows and call it a  $h$ -wide **Twin Water Lily of order**  $(s_1, s_2)$ . The vertex set of  $\text{TWL}(h, s_1, s_2)$  consists of two independent sets  $S_1$  and  $S_2$  with  $r_1$  and  $r_2$  vertices, respectively. Additionally, for every size  $h$  subset  $\{(v_1, s_1), \dots, (v_h, s_h)\} \in \binom{(S_1 \cup S_2) \times [h]}{h}$ , it has a vertex  $v$  which is connected to all vertices in  $\{v_1, \dots, v_h\} \cap S_2$ . We define  $P$  to be the set of all such  $v$  created in this way.*

*Note that  $S_1$  consists only of isolated vertices, and that  $H$  is bipartite.*



■ **Figure 2** Partial sketch of the 3-wide Twin Water Lily of order  $(6, 4)$ . The dashed edges are not actual edges, they just represent which other vertices are in the set  $\{v_1, \dots, v_h\}$  of the vertex  $u \in P$  that they are connected to.

► **Proposition 20.** *If a graph  $H$  is a  $h$ -wide Twin Water Lily of order  $(s_1, s_2)$ , then its treewidth is bounded by*

$$\text{tw}(H) \leq \begin{cases} s_2 - 1 & \text{if } s_2 > h \\ s_2 & \text{otherwise} \end{cases}$$

*and its pathwidth is bounded by  $\text{pw}(H) \leq s_2$ .*

**Proof.** There is a very useful characterization of treewidth using a graph-theoretic game: A graph  $G$  has treewidth  $\leq k$  if and only if  $k + 1$  cops can catch<sup>5</sup> a visible robber on  $G$  [63].

<sup>5</sup> See the corresponding footnote in Section 2 for a description of the game.



To show the bound on the treewidth of  $H$ , simply place the cops on all vertices of  $S_2$ . No matter where the robber starts, it is surrounded by cops or is on an isolated vertex. If  $s_2 > h$ , then there must exist some cop which is not adjacent to the robber, whom we can use to catch him in a single step. If  $s_2 \leq h$ , it is not guaranteed that there is a non-adjacent cop. However, we have an additional cop which can start at any vertex. As soon as the robber is positioned, we use the additional cop to capture them.

A similar characterization exists for pathwidth: A graph  $G$  has pathwidth  $\leq k$  if and only if  $k + 1$  cops can catch<sup>6</sup> an invisible robber on  $G$  [38].

For the pathwidth of  $H$ , we again place all cops on the vertices of  $S_2$  and have one left over. We use this cop to go through all vertices not in  $S_2$ , one in each step. The robber is always surrounded by the cops in  $S_2$  and hence cannot move, so after going through all vertices with the additional cop, we must have caught him.  $\blacktriangleleft$

## 4.2 Unweighted Colored Subgraph Isomorphism

► **Lemma 21** (The Unweighted Lemma). *For any  $h \geq 3$  and any  $r \geq 2$ , if we can solve COLORED SUBGRAPH ISOMORPHISM with pattern graph  $\text{TWL}(h, 0, r)$  in time  $O(N^{r-\varepsilon})$ , then we can solve  $h$ -uniform  $hr$ -HYPERCLIQUE in time  $O(n^{hr-\varepsilon})$ .*

**Proof.** Let an instance  $\mathcal{G}_0 = G_0$  of  $h$ -uniform  $hr$ -HYPERCLIQUE be given. We convert this to a COLORED SUBGRAPH ISOMORPHISM instance with a Twin Water Lily as pattern graph in three steps, each of which we explain in detail below: First, we convert it to a COLORED  $hr$ -HYPERCLIQUE instance in a standard way. Second, we go from hypercliques to COLORED SUBGRAPH ISOMORPHISM by replacing each hyperedge by an intermediate vertex. Finally, we merge preimages to homogenize preimage sizes.

**1. Converting to Colored Hyperclique:** We convert  $\mathcal{G}_0$  to a  $h$ -uniform COLORED HYPERCLIQUE instance with  $hr$  colors. The converted instance should have a hyperclique where all vertices has different colors if and only if the old instance has a hyperclique. The new instance have the form  $\mathcal{G}_1 = (G_1, f_1)$ , where the color homomorphism  $f_1 : V(G_1) \rightarrow V(C_{hr})$  assigns each vertex of  $G_1$  a vertex in the  $h$ -uniform  $hr$ -hyperclique  $C_{hr}$ .

Let  $V(C_{hr}) = \{1, \dots, hr\}$ . For each  $i$ , the preimage  $f_1^{-1}(i)$  is a copy of  $V(G_0)$ . Let  $g : V(G_1) \rightarrow V(G_0)$  be a function between sets that indicates which vertex in  $G_0$  the vertex in  $G_1$  is a copy of. Now for each set  $\{v_1, \dots, v_h\} \in \binom{V(C_{hr})}{h}$ , we go through all tuples  $(w_1, \dots, w_h) \in f_1^{-1}(v_1) \times \dots \times f_1^{-1}(v_h)$  and create the edge  $\{w_1, \dots, w_h\} \in E(G_1)$  if and only if  $\{g(w_1), \dots, g(w_h)\} \in E(G_0)$ .

The correctness of this construction is easy to see.

**2. Representing Hyperedges by Intermediate Vertices:** We now go from the COLORED HYPERCLIQUE instance  $\mathcal{G}_1 = (G_1, f_1)$  to a (structured) COLORED SUBGRAPH ISOMORPHISM instance  $\mathcal{G}_2 = (H_2, G_2, f_2)$ . The reduction is done in a standard way: We replace each hyperedge with a vertex connected to all its endpoints.

Formally, we need to construct  $H_2$  and  $G_2$ .  $H_2$  has two sets of vertices  $S'_2$  and  $P$ .  $S'_2$  is a copy of  $V(C_{hr})$  from the last step, including its preimages. Accordingly, we write  $S'_2 = \{1, \dots, hr\}$ . In  $P$ , we have one vertex  $u$  for every subset  $\{w_1, \dots, w_h\} \in \binom{C_{hr}}{h}$ , and we have  $\forall \ell \in [h] : uw_\ell \in E(H_2)$ . Now for every hyperedge  $\{w'_1, \dots, w'_h\} \in E(G_1)$  with

<sup>6</sup> The game can be formulated such that it is the same as the one for treewidth, but the cops simply cannot see the robber and must therefore have a universal strategy for catching him on  $G$ . This is also sometimes referred to as the contamination cleansing or infection cleansing game.

$\forall \ell \in [h] : w'_\ell \in f_1^{-1}(G_1)$ , we add a vertex  $u' \in f_2^{-1}(u)$  which is connected to all vertices  $w'_1, \dots, w'_h$ . This concludes the construction of  $\mathcal{G}_2$ .

Correctness of this construction is again easy to see. As for the size, note that preimages of vertices in  $S'_2$  still have size  $n$ . The preimages of vertices in  $P$ , however, have at most  $n^h$  vertices.

**3. Merging Preimages in  $S'_2$ :** Lastly, we go from  $\mathcal{G}_2 = (H_2, G_2, f_2)$  to the final instance  $\mathcal{G}_3 = (H_3, G_3, f_3)$  where  $H_3$  is a Twin Water Lily of order  $(0, r)$ . Note that  $\mathcal{G}_2$  is “almost” the instance we want, save for the fact that the preimages  $P$  are much larger (size up to  $n^h$ ) than the preimages of  $S'_2$  (size  $n$ ). We rectify this by merging groups of vertices within  $S'_2$ . These groups have size  $h$ .

We split  $S'_2$  into  $r$  groups  $X_1, \dots, X_r$  of size  $h$ . In  $H_3$ , we have for every  $i \in [r]$  a vertex  $x_i$  representing  $X_i$ . Each vertex  $x'_i \in f_3^{-1}(x_i)$  corresponds to a configuration  $\text{conf}(x'_i) \in \text{Conf}(X_i)$ . In accordance with the definition of a Twin Water Lily, we define  $S_1 = \emptyset$  and  $S_2 = \{x_1, \dots, x_r\}$ .

The set  $P \subseteq V(H_3)$  remains the same as in the preceding step, including its preimages. For each  $u \in P$ , we connect  $u$  to all  $x_i$  such that  $X_i \cap N_{H_2}(u) \neq \emptyset$  (recall that  $N_{H_2}(u)$  is defined to be the neighbourhood of  $u$  in  $H_2$ ). Note that each  $u \in P$  is still connected to at most  $h$  other vertices, but that it can be less if multiple vertices of its neighborhood came from the same group. Finally, for an edge  $ux_i \in E(H_3)$ , we connect a vertex  $u' \in f_3^{-1}(u)$  to a vertex  $x'_i \in f_3^{-1}(x_i)$  if and only if  $\forall v \in N_{H_2}(u) \cap X_i : u'(\text{conf}(x'_i)(v)) \in E(G_2)$ .

Correctness is easy to see. Note that  $H_3$  is an  $h$ -wide Twin Water Lily of order  $(0, r)$  now, and all preimages are of size at most  $N := n^h$ .

This completes the construction of the reduction algorithm. Each of the steps runs in  $O(n^{2h-1})$  time: In the intermediate-vertex step, each of the vertices  $u'$  in a preimage  $f^{-1}(u)$  represents a hyperedge and is hence only connected to  $h$  vertices. After the next step, it is connected to all vertices which represent compatible configurations. Each vertex of  $S_2$  that  $u$  is now connected to must represent a non-empty intersection of the neighbourhood of  $u$ . Hence it has cardinality at least 1. Hence there are at most  $n^{h-1}$  compatible vertices in its preimage. Hence overall, there are at most  $O(n^{2h-1})$  edges, and the graph can also be constructed in this time.

Now suppose there is an algorithm solving the COLORED SUBGRAPH ISOMORPHISM problem with pattern  $\text{TWL}(h, 0, r)$  in time  $O(N^{r-\varepsilon})$ . Then for any  $h$ -uniform  $hr$ -HYPERCLIQUE instance, we run the reduction in time  $O(n^{2h-1})$  and then solve the new instance in time  $O(N^{r-\varepsilon}) = O(n^{hr-h\varepsilon})$ . Since  $r \geq 2$ , the algorithm hence takes total time  $O(n^{hr-h\varepsilon}) = O(n^{hr-\varepsilon'})$ .  $\blacktriangleleft$

We now use the Unweighted Lemma to prove the lower bounds for COLORED SUBGRAPH ISOMORPHISM. In particular, we prove the following theorem, which implies Theorem 1 from the Results section via the Equivalence Lemma (Lemma 14).

► **Theorem 22.** *The following statements are true.*

1. *For each  $t \geq 3$  and any  $3 \leq h \leq t$ , there exists a connected, bipartite pattern graph  $\mathcal{H}_{t,h}$  of treewidth  $t$  such that there cannot be an algorithm solving the COLORED SUBGRAPH ISOMORPHISM problem on pattern graph  $\mathcal{H}_{t,h}$  in time  $O(n^{t+1-\varepsilon})$  unless the  $h$ -uniform  $h(t+1)$ -HYPERCLIQUE hypothesis fails.*
2. *For each  $t \geq 2$  and any  $h \geq 3$ , there exists a connected, bipartite pattern graph  $\mathcal{H}_{t,h}$  of treewidth  $t$  such that there cannot be an algorithm solving the COLORED SUBGRAPH ISOMORPHISM problem on pattern graph  $\mathcal{H}_{t,h}$  in time  $O(n^{t-\varepsilon})$  unless the  $h$ -uniform  $ht$ -HYPERCLIQUE hypothesis fails.*

3. For each  $t \geq 2$ , there exists a connected, bipartite pattern graph  $\mathcal{H}_t$  of treewidth  $t$  such that there cannot be an algorithm solving the COLORED SUBGRAPH ISOMORPHISM problem on pattern graph  $\mathcal{H}_t$  in time  $O(n^{(t+1)\omega/3})$  unless the  $(t+1)$ -CLIQUE hypothesis fails.

► **Theorem 23** (Theorem 22 for pathwidth). *Part 2 of Theorem 22 also holds when replacing the treewidth  $t$  by the pathwidth  $p$ . Part 3 only holds when replacing  $t+1$  by  $p+1$ , and the pattern graph is not bipartite anymore. Part 1 does not hold.*

Unfortunately, Lemma 14 cannot be applied to Theorem 23, and hence we have no lower bounds for the uncolored, unweighted case of bounded pathwidth. We believe it is unlikely that the techniques used to prove 14 generalize to pathwidth.

We now prove the theorems above.

**Proof (of Theorem 22).** For the proof of this theorem, we use the Unweighted Lemma.

**Part 1:** Let  $t \geq 3$  and  $3 \leq h \leq t$  be given. It suffices to apply the Unweighted Lemma with  $h' := h$  and  $r := t+1$ , and set  $\mathcal{H}_{t,h} = \text{TWL}(h, 0, t+1)$ . As proven in Proposition 20,  $\text{TWL}(h, 0, t+1)$  has treewidth  $t$  since  $t+1 > h$ .

**Part 2:** Let  $t \geq 3$  and  $h \geq 3$  be given. It suffices to apply the Unweighted Lemma with  $h' := h$  and  $r := t$ , and set  $\mathcal{H}_{t,h} = \text{TWL}(h, 0, t+1)$ . The loss of the  $+1$  in the exponent is due to the weaker bound in Proposition 20 for  $t \leq h$ .

**Part 3:** Let  $t \geq 2$  be given. We know that by the  $t+1$ -CLIQUE hypothesis, COLORED SUBGRAPH ISOMORPHISM on pattern graph  $C_{t+1}$  cannot be solved in time  $O(n^{(t+1)\omega/3})$ . We can make  $C_{t+1}$  bipartite in the obvious way by subdividing the edges. This subdivided graph is  $\mathcal{H}_t$ . ◀

**Proof (of Theorem 23).** Completely analogous. The deviation in bounds with respect to Theorem 22 is due to the difference in bounds for treewidth and pathwidth in Proposition 20.

In part 3, we cannot subdivide the edges of  $\mathcal{H}_t$  without changing the pathwidth, hence the pattern graph stays a clique and is thus not bipartite. ◀

Indeed, the Unweighted Lemma can also be used to prove the results under SETH. SETH is beyond doubt the most widely used for conditional lower bounds for problems in P, which is why the following results are still interesting, even though they only give smaller lower bounds than the results under the HYPERCLIQUE hypothesis. For context on SETH and the many conditional lower bounds it enables, see e.g. [3, 9, 8, 4, 23, 25, 61].

The following theorem implies Theorem 2 via the Equivalence Lemma (Lemma 14).

► **Theorem 24.** *Assuming SETH, the following two statements are true.*

1. For any  $t \geq 3$  and any  $\varepsilon > 0$  there exists a pattern graph  $\mathcal{H}_{t,\varepsilon}$  of treewidth  $t$  such that there cannot be an algorithm solving all instances of COLORED SUBGRAPH ISOMORPHISM with pattern graph  $\mathcal{H}_{t,\varepsilon}$  in time  $O(n^{t-\varepsilon})$ .
2. For any  $\varepsilon > 0$  there exists a  $t \geq 3$  and a pattern graph  $\mathcal{H}_\varepsilon$  of treewidth  $t$  such that there cannot be an algorithm solving all instances of COLORED SUBGRAPH ISOMORPHISM with pattern graph  $\mathcal{H}_\varepsilon$  in time  $O(n^{t+1-\varepsilon})$ .

**Proof.** This result also follows via Unweighted Lemma and hence via the HYPERCLIQUE problem. Specifically, the split-and-list technique from [67] allows one to reduce the Satisfiability problem to HYPERCLIQUE. Along the same lines, the following result was shown in [51, Lemma 9.1].

► **Lemma 25** ([51]). *Assuming SETH, for any  $\varepsilon > 0$  there exists  $h \geq 3$  such that for all  $k > h$ , the  $h$ -uniform  $k$ -HYPERCLIQUE problem is not in time  $O(n^{k-\varepsilon})$ .*

Using this, we now prove parts 1 and 2 of the theorem.

**Part 1:** Let  $t \geq 3$  and  $\varepsilon > 0$  be given. We use Lemma 25 to obtain  $h \geq 3$  such that  $h$ -uniform  $ht$ -HYPERCLIQUE is not in time  $O(n^{ht-\varepsilon})$ . W.l.o.g. assume  $h \geq t$ . Now it suffices to apply the Unweighted Lemma with  $h' := h$  and  $r := t$ , and set  $\mathcal{H}_{t,\varepsilon} = \text{TWL}(h, 0, t)$ . Via Proposition 20 and the fact that  $h \geq t$ , we know that  $\text{TWL}(h, 0, t)$  has treewidth exactly  $t$ .

**Part 2:** Let  $\varepsilon > 0$  be given. Lemma 25 gives a  $h \geq 3$  such for all  $k > h$ ,  $h$ -uniform  $k$ -HYPERCLIQUE is not in time  $O(n^{k-\varepsilon})$ . We choose  $t := h$  and get that  $h$ -uniform  $ht$ -HYPERCLIQUE is not in time  $O(n^{ht-\varepsilon})$ . Now it suffices to apply the Unweighted Lemma with  $h' := h$  and  $r := t + 1 = h + 1$  and set  $\mathcal{H}_{t,\varepsilon} = \text{TWL}(h, 0, h + 1)$ . Via Proposition 20, we know that  $\text{TWL}(h, 0, t + 1)$  has treewidth exactly  $h = t$ .  $\blacktriangleleft$

### 4.3 Exact Weight Colored Subgraph Isomorphism

First, we state the Weighted Lemma. This result enables us to prove lower bounds for EXACT WEIGHT COLORED SUBGRAPH ISOMORPHISM.

- **Lemma 26 (The Weighted Lemma).** *For any  $\varepsilon \in (0, 1)$  and any constant parameters  $h \in \mathbb{N} \setminus \{1\}, r_1 \in \mathbb{N}, r_2 \in \mathbb{N}, \beta \in (0, 1) \cap \mathbb{Q}$ , there exists a  $k \in \mathbb{N}$  and an algorithm  $\mathcal{A}$  which*
- (a) *accepts as input an instance  $\mathcal{I} = G$  of  $h$ -uniform  $k$ -HYPERCLIQUE.*
  - (b) *produces an equivalent instance  $\mathcal{I}' = (H', G', f', w')$  of EXACT WEIGHT COLORED SUBGRAPH ISOMORPHISM, where  $H'$  is a  $h$ -wide Twin Water Lily of order  $(r_1, r_2)$ . The preimages of  $\mathcal{I}'$  have size at most  $\max\{n^{\beta k/r_1}, n^{(1-\beta)k/r_2}\}$ , and the maximum weight is  $W = \Theta(n^{(1+\varepsilon)\beta k})$ .*
  - (c) *runs in time  $O(n^{2h-1} + (n^{\beta k/(hr_1)})^{\hat{c}})$  for some universal constant  $\hat{c} \in \mathbb{N}$ .*

Intuitively, the parameter  $\beta$  indicates what percentage of the instance  $\mathcal{I}$  should be encoded in which part of the Twin Water Lily. A percentage of  $\beta$  is encoded in the weights, while the remaining percentage of  $(1 - \beta)$  is encoded in the edges.

We now prove the Weighted Lemma.

**Proof.** Let an instance  $\mathcal{I}_0 = G_0$  of  $h$ -uniform  $k$ -HYPERCLIQUE be given, where  $k$  is chosen later. We convert this to an EXACT WEIGHT COLORED SUBGRAPH ISOMORPHISM instance with a Twin Water Lily as pattern graph in five steps, each of which we explain in detail below: First, we convert it to a COLORED  $k$ -HYPERCLIQUE instance in a standard way. Second, we split the instance into the part that we want to encode in the weights and the part that we want to encode in the edges. In both of these parts, we merge large groups of preimages such that we are left with only  $hr_1$  in the weight part, and  $hr_2$  in the edges part. Third, we go from hypercliques to COLORED SUBGRAPH ISOMORPHISM in a standard way while preserving the preimages. Fourth, we convert the weight part of the instance into actually using weights by replacing edge constraints by weight constraints, using a construction known as  $k$ -average free sets. Finally, we merge preimages in both parts again to obtain the final Twin Water Lily instance.

- 1. Converting to Colored Hyperclique:** This step works exactly like step 1 in the proof of the Unweighted Lemma. As described there, We convert  $\mathcal{I}_0$  to a  $h$ -uniform COLORED HYPERCLIQUE instance. The latter has the form  $\mathcal{I}_1 = (G_1, f_1)$ , where the color homomorphism  $f_1 : V(G_1) \rightarrow V(C_k)$  assigns each vertex of  $G_1$  a vertex in the  $k$ -hyperclique  $C_k$ .
- 2. Merging Preimages:** We now convert  $\mathcal{I}_1$  to a  $h$ -uniform COLORED  $(hr_1 + hr_2)$ -HYPERCLIQUE instance, by condensing groups of (small) preimages into single (large) preimages. In the converted instance  $\mathcal{I}_2 = (G_2, f_2)$  with color homomorphism  $f_2 : V(G_2) \rightarrow V(C_{r_1+r_2})$ , we

ensure that  $V(C_{hr_1+hr_2})$  can be divided into two sets  $H_1, H_2$  such that  $|H_1| = hr_1, |H_2| = hr_2$  and  $\forall v \in H_1 : |f_2^{-1}(v)| = n^{\beta k/(hr_1)}, \forall v \in H_2 : |f_2^{-1}(v)| = n^{(1-\beta)k/(hr_2)}$ . These two sets correspond to the two water lilies constructed in later steps.

At this point, we must make our choice of  $k$ . We will need that  $k_1 := \beta k$  is an integer and divisible by  $hr_1$ . Furthermore,  $k_2 := (1 - \beta)k$  must also be an integer and divisible by  $hr_2$ . Letting  $\beta = \frac{p}{q} \in (0, 1) \cap \mathbb{Q}$  where  $p, q \in \mathbb{N}$ , it hence suffices to choose  $k = hr_1 r_2 q$ . Now, we split  $V(C_k) = \{1, \dots, k\}$  into two groups  $V_1 = \{1, \dots, k_1\}$  and  $V_2 = \{k_1 + 1, \dots, k_2\}$ . Going further, we split each of these sets again:  $V_1$  is split into  $hr_1$  disjoint groups  $V_{1,1}, \dots, V_{1,hr_1}$  of size  $\frac{k_1}{hr_1}$  each. Analogously, we split  $V_2$  into  $hr_2$  disjoint groups  $V_{2,1}, \dots, V_{2,hr_2}$  of size  $\frac{k_2}{hr_2}$  each.

Now define  $V(C_{hr_1+hr_2}) = \{v_1, \dots, v_{hr_1+hr_2}\}$  and furthermore let  $S'_1 = \{v_1, \dots, v_{hr_1}\}, S'_2 = \{v_{hr_1+1}, \dots, v_{hr_1+hr_2}\}$ . The vertex  $v_i$  corresponds to  $V_i$ , for every  $i$ . The vertices in their preimages represent all the configurations of the sets. In particular, we let every  $v'_i \in f_2^{-1}(v_i)$  represent a configuration  $\text{conf}(v'_i) \in \text{Conf}(V_i)$ . Hence, for  $v_i \in S'_1$  the preimage  $f_2^{-1}(v_i)$  has size  $n^{k_1/(hr_1)}$ , and for  $v_i \in S'_2$  the preimage  $f_2^{-1}(v_i)$  has size  $n^{k_2/(hr_2)}$ .

Now for the edges. For every subset  $\{v_{i_1}, \dots, v_{i_h}\} \in \binom{S'_1 \cup S'_2}{h}$ , we iterate over all  $(v'_{i_1}, \dots, v'_{i_h}) \in f_2^{-1}(v_{i_1}) \times \dots \times f_2^{-1}(v_{i_h})$ . We combine the configurations that they represent by defining  $R \in \text{Conf}(V_{i_1} \cup \dots \cup V_{i_h})$  as  $\forall \ell \in [h] : \forall v'_{i_\ell} \in f_2^{-1}(v_{i_\ell}) : R(v'_{i_\ell}) := \text{conf}(v'_{i_\ell})$ . Now we add  $\{w_1, \dots, w_h\}$  as a hyperedge to  $E(G_2)$  if and only if  $R$  is a valid configuration. That is, if the image of  $R$  induces a hyperclique in  $G_1$ .

Again, correctness is easy to see.

**3. Representing Hyperedges by Intermediate Vertices:** We now go from the COLORED HYPERCLIQUE instance  $\mathcal{G}_2 = (G_2, f_2)$  to a (structured) COLORED SUBGRAPH ISOMORPHISM instance  $\mathcal{G}_3 = (H_3, G_3, f_3)$ . The reduction is essentially the same as the one in step 2 in the proof of the Unweighted Lemma.

We construct  $H_3$  and  $G_3$ .  $H_3$  has three sets of vertices  $S'_1, S'_2$  and  $P$ .  $S'_1$  and  $S'_2$  copy  $S'_1$  and  $S'_2$  from the last step, including their preimages. Accordingly, we write  $S'_1 = \{v_1, \dots, v_{hr_1}\}$  and  $S'_2 = \{v_{hr_1+1}, \dots, v_{hr_1+hr_2}\}$ . In  $P$ , we have one vertex  $u$  for every subset  $\{w_1, \dots, w_h\} \in \binom{C_{hr_1+hr_2}}{h}$ , and we have  $\forall \ell \in [h] : uw_\ell \in E(H_3)$ . For every hyperedge  $\{w'_1, \dots, w'_h\} \in E(G_2)$  with  $\forall \ell \in [h] : w'_\ell \in f_2^{-1}(G_2)$ , we add a vertex  $u' \in f_3^{-1}(u)$  which is connected to all vertices  $w'_1, \dots, w'_h$ . Note here that  $u'$  is only connected to one vertex from each preimage, which is a property that will be needed in the fourth step.

Correctness of this construction is easy to see. Note that preimages of vertices in  $S'_1$  and  $S'_2$  still have size  $n^{k_1/(hr_1)}$  and  $n^{k_2/(hr_2)}$ , respectively. The preimages of vertices in  $P$ , have at most  $(\max\{n^{k_1/(hr_1)}, n^{k_2/(hr_2)}\})^h = \max\{n^{k_1/r_1}, n^{k_2/r_2}\}$  vertices.

**4. Replacing Some of the Edges with Weights** We now come to the crucial step of converting some of the edge constraints to weight constraints. We convert the COLORED SUBGRAPH ISOMORPHISM instance  $\mathcal{G}_3 = (H_3, G_3, f_3)$  of the preceding step into an EXACT WEIGHT COLORED SUBGRAPH ISOMORPHISM instance  $\mathcal{G}_4 = (H_4, G_4, f_4, w_4)$ . To do this, we will need so-called  $k$ -average free sets.

► **Definition 27** ( *$k$ -average free sets*). *A set  $S \subseteq \mathbb{Z}$  is called  $k$ -average-free if, for any  $s_1, \dots, s_{k'} \in S$  with  $k' \leq k$ , we have  $s_1 + \dots + s_{k'} = k' \cdot s_{k'+1}$  if and only if  $s_1 = \dots = s_{k'+1}$ . In other words, the average of  $s_1, \dots, s_{k'} \in S$  is in  $S$  if and only if all  $s_i$  are equal.*

We use the following construction for  $k$ -average free sets, originally proven in [16], modified into a more useful version in [7] and formulated in this form in [4].

► **Lemma 28.** *There exists a universal constant  $c > 0$  such that, for all constants  $\varepsilon \in (0, 1)$  and  $k \geq 2$ , a  $k$ -average-free set  $S$  of size  $n$  with  $S \subseteq [0, k^{c/\varepsilon} n^{1+\varepsilon}]$  can be constructed in time  $\text{poly}(n)$ .*

Specifically, we use Lemma 28 with  $\varepsilon' = \varepsilon, k' = \lambda := |P|$  and  $n' = n^{k_1/(hr_1)}$ . Letting  $B := \lambda^{c/\varepsilon} n^{(1+\varepsilon)k_1/(hr_1)}$ , this yields a  $\lambda$ -average free set  $S \subseteq [0, B]$  of size  $n^{k_1/(hr_1)}$ . From this, we can construct an arbitrary bijection  $\varrho_S : [n^{k_1/(hr_1)}] \rightarrow S$ .

Now, to construct  $\mathcal{G}_4$ , we first copy  $\mathcal{G}_3$ , giving each node a default weight of “infinity” (i.e. something otherwise unobtainable, e.g.  $k^2W + 1$ ). Now we delete all edges in  $H_3$  which are incident to a vertex in  $S'_1$ , along with the corresponding edges in  $G_3$ . These are the edges that we replace by weight constraints.

Hence we now describe the weights. To make our construction easier, we specify a target value  $T$  (instead of the default target zero). We can easily get rid of this again by picking some vertex  $\hat{w} \in V(H_3)$  and subtracting  $T$  from the weights of all of its preimages. The binary representation of  $T$  consists of  $h \cdot a$  blocks of  $\lceil \log(2\lambda B) \rceil$  bits, each containing the binary representation of  $\lambda B$ . The  $i$ -th block represents the vertex  $v_i \in S'_1$ .

We move to the weights of the vertices, starting with vertices in the preimages of  $S'_1$ . Somewhat abusing notation, we define  $\forall i \in [hr_1] : f_4^{-1}(v_i) = \{1, \dots, n^{k_1/(hr_1)}\}$ . The weight of vertex  $v'_i \in f_4^{-1}(v_i)$  has a value of  $\lambda B - |N(v_i)| \cdot \varrho_S(v'_i)$  in the  $i$ -th block, and a value of zero in all other blocks. Now for vertices in the preimages of  $P$ . Let  $u \in P$  correspond to the set  $\{w_1, \dots, w_h\} \in \binom{C_{hr_1+hr_2}}{h}$ . As observed in the preceding step, each  $u' \in f_3^{-1}(u)$  is connected to exactly one vertex  $w'_i$  from each preimage  $f_3^{-1}(w_i)$ . In the current step, for each  $i \in [h]$  with  $w_i \in S'_1$ , we have deleted the edges  $u'w'_i$ . To replace them, for each such  $i$ , we give  $u'$  a value of  $\varrho_S(w'_i)$  in the  $i$ -th block. We have just changed the weight of  $u'$  at  $|N(u) \cap S'_1|$  blocks of its binary representation. All other blocks have a value of zero.

All vertices with so far unspecified weight have weight zero. This concludes the construction of  $\mathcal{G}_4$ .

We show correctness of this construction. It suffices to show that any configuration  $R$  that is a solution for  $\mathcal{G}_3$  is also a solution for  $\mathcal{G}_4$  and vice versa. Hence, suppose  $R$  is a solution for  $\mathcal{G}_3$ . Then all edge constraints of  $H_4$  are trivially fulfilled and we need only show that the total weight is  $T$ . For ease of discussion, we denote by  $\alpha[i]$  the value of the  $i$ -th block of a weight  $\alpha$ . Consider the blocks of the binary representation of the sum of weights  $w(R) = \sum_{v \in \text{Im}(R)} w_4(v)$ , and let  $i$  be fixed. The large block size prevents overflow, so  $w(R)[i] = \sum_{v \in \text{Im}(R)} w_4(v)[i]$ . By construction, we have  $w(R)[i] = w(R(v_i))[i] + \sum_{u \in N(v_i)} w(R(u))[i]$ . However since  $R$  is a valid configuration in  $\mathcal{G}_3$ , we have that for each  $u, \forall v_j \in N(u) : R(u)R(v_j) \in E(G_3)$ . In particular,  $R(u)R(v_i) \in E(G_3)$  and hence by construction  $w(R(u))[i] = \varrho_S(v_i)$ . We conclude  $w(R)[i] = \lambda B - |N(v_i)| \cdot \varrho_S(R(v_i)) + \sum_{i \in N(v_i)} \varrho_S(R(v_i)) = \lambda B = T[i]$ . Hence  $w(R)$  is equal to  $T$  in each of its blocks, which was to be proven.

Conversely, suppose  $R$  is a solution for  $\mathcal{G}_4$ . Then all edge constraints in  $G_3[V(G_3) \setminus S'_1]$  are trivially satisfied and we need only show that  $\forall v_i \in S'_1 : \forall u \in N(v_i) : R(v_i)R(u) \in E(G_3)$ . Fix  $v_i \in S'_1$ . We have that  $\lambda B = T[i] = w(R)[i] = w(R(v_i))[i] + \sum_{u \in N(v_i)} w(R(u))[i]$ . Let  $N(v_i) = \{u_1, \dots, u_{|N(v_i)|}\}$ . For each  $\ell \in |N(v_i)|$ , we have that  $R(u_\ell)$  is connected to some vertex  $v_i^{(\ell)} \in f_3^{-1}(v_i)$ , and hence that  $w(R(u_\ell))[i] = \varrho_S(v_i^{(\ell)})$ . Hence we have that  $\lambda B = \lambda B - |N(v_i)| \cdot \varrho_S(R(v_i)) + \sum_{\ell \in |N(v_i)|} \varrho_S(v_i^{(\ell)})$ . Hence  $|N(v_i)| \cdot \varrho_S(R(v_i)) = \sum_{\ell \in |N(v_i)|} \varrho_S(v_i^{(\ell)})$ . But because the values in the image of the bijection  $\varrho_S$  are a  $\lambda$ -average free set and  $N(v_i) \subseteq P$  certainly has size less than  $|P| = \lambda$ , we have that  $\forall \ell : v_i^{(\ell)} = R(v_i)$ . Thus for all  $\ell$ ,  $R(u_\ell)$  is connected to  $R(v_i)$ , which was to be proven.

**5. Merging Preimages in  $S'_1$  and  $S'_2$ :** Lastly, we go from  $\mathcal{G}_4 = (H_4, G_4, f_4, w_4)$  to the final instance  $\mathcal{G}_5 = (H_5, G_5, f_5, w_5)$  where  $H_5$  is a Twin Water Lily of order  $(r_1, r_2)$ . This step is similar to step 3 of the proof of the Unweighted Lemma.

We merge groups of vertices within  $S'_1$  and  $S'_2$ . In both sets, these groups have size  $h$ .

We split  $S'_1$  into  $r_1$  groups  $X_1, \dots, X_{r_1}$  of size  $h$ , and we split  $S'_2$  into  $r_2$  groups  $X_{r_1+1}, \dots, X_{r_1+r_2}$  of size  $h$ . In  $H_5$ , we have for every  $i \in [r_1+r_2]$  a vertex  $x_i$  representing  $X_i$ . Each vertex  $x'_i \in f_5^{-1}(x_i)$  corresponds to a configuration  $\text{conf}(x'_i) \in \text{Conf}(X_i)$ . The weight of  $x'_i$  is  $w(\text{conf}(x'_i))$ , i.e. the sum of the weights of the vertices in the image of the configuration. In accordance with the definition of a Twin Water Lily, we define  $S_1 = \{x_1, \dots, x_{r_1}\}$  and  $S_2 = \{x_{r_1+1}, \dots, x_{r_1+r_2}\}$ .

The set  $P \subseteq V(H_5)$  remains the same as in the preceding step, including its preimages and the weights of the vertices in the preimages. For each  $u \in P$ , we go through the vertices  $v$  in the neighbourhood of  $u$  in  $H_4$ , and connect  $u$  to  $x_i$  such that  $v \in X_i$ . Note that each  $u \in P$  is still connected to at most  $h$  other vertices, but that it can be less if multiple vertices of its neighborhood came from the same group. Finally, we connect a vertex  $u' \in f^{-1}(u)$  to a vertex  $x'_i \in f_5^{-1}(x_i)$  if and only if  $\forall v \in N_{H_2}(u) \cap X_i : u'(\text{conf}(x'_i)(v)) \in E(G_4)$ .

Correctness is easy to see. Note that  $H_5$  is a Twin Water Lily of order  $(r_1, r_2)$  now. The set  $S_1$  has size  $r_1$  and  $S_2$  has size  $r_2$ , with respective preimages of size  $n^{k_1/r_1}$  and  $n^{k_2/r_2}$ . The preimages of  $P$  still have size  $\max\{n^{k_1/r_1}, n^{k_2/r_2}\}$ . Furthermore, the weights constructed in step 4 have  $hr_1$  blocks of  $\lceil \log(2\lambda B) \rceil$  bits, hence the maximum weight is  $\Theta(2^\zeta)$  where  $\zeta = hr_1 \cdot (\log(\lambda B) + O(1)) = hr_1 \cdot \log(|P|^{1+c/\varepsilon} n^{(1+\varepsilon)k_1/(hr_1)}) + O(1) = hr_1 \cdot \log(n^{(1+\varepsilon)k_1/(hr_1)}) + O(1) = (1+\varepsilon)k_1 \log(n) + O(1)$ , hence as promised in the statement of the lemma the maximum weight is  $\Theta(n^{(1+\varepsilon)k_1})$ .

We have shown a reduction that has properties (b) and (c) from the lemma. We still have to analyze the running time. It is easy to see that steps 1, 2 and 3 run in time  $O(n^{2h-1})$ , just as in the proof of the unweighted lower bound. In step 4, we need to construct the  $k$ -average free set  $S$ , which is done in time  $\text{poly}(n^{k_1/(hr_1)}) = O((n^{k_1/(hr_1)})^{\hat{c}})$  for some universal constant  $\hat{c} \in \mathbb{Z}$ . The rest of step 4 as well as step 5 can again be done in time  $O(n^{2h-1})$ . This concludes the proof.  $\blacktriangleleft$

We now use this lemma to prove the lower bound for EXACT WEIGHT COLORED SUBGRAPH ISOMORPHISM. Note that it only gives node-weighted instances. But as the following lemma shows, this is enough.. It shows that we can always convert node-weighted instances to edge-weighted instances, hence showing lower bounds for node-weighted instances immediately shows lower bounds for edge-weighted instances. It can also be reused later in the algorithms, as we can then always assume that the instances are edge-weighted.

**► Proposition 29.** *Given an instance of EXACT NODE WEIGHT COLORED SUBGRAPH ISOMORPHISM where each vertex  $v \in V(G)$  with  $w(v) \neq 0$  has degree at least one, we can transform it into an equivalent instance of EXACT EDGE WEIGHT COLORED SUBGRAPH ISOMORPHISM in such a way that only the weight function changes. Furthermore, this reduction runs in time linear in the input size.*

**Proof.** We construct a new weight function  $w' : E(G) \rightarrow \mathbb{Z}$ . Initially,  $w'(e) = 0$  for all  $e \in E(G)$ . The idea is to push the weight of each vertex of non-zero weight in  $G$  onto one of its edges. Hence let  $v \in V(G)$  with  $w(v) \neq 0$ . Then there must be  $u \in V(G)$  with  $uv \in E(G)$ . We add  $w(v)$  to  $w'(uv)$ . This completes the reduction.

It can easily be seen that if there is a solution in the node-weighted instance with  $w$ , then that same solution must work with  $w'$  and vice versa.  $\blacktriangleleft$

This enables us to prove the following theorem. Note that it implies Theorem 4 via Lemma 14.

► **Theorem 30.** *For both the node- and edge weighted variant of the problems, the following statements are true.*

1. *For each  $t \geq 3$ , each  $\gamma \in \mathbb{R}^+$  and any  $3 \leq h \leq t$ , there exists a connected, bipartite graph  $\mathcal{H}_{t,h,\gamma}$  of treewidth  $t$  such that there cannot be an algorithm solving the EXACT WEIGHT COLORED SUBGRAPH ISOMORPHISM problem on pattern graph  $\mathcal{H}_{t,h,\gamma}$  for instances with maximum weight  $W = \Theta(n^\gamma)$  in time  $O(n^{t+1-\varepsilon}W)$ , unless the  $h$ -uniform HYPERCLIQUE hypothesis fails.*
2. *For each  $t \geq 1$ , each  $\gamma \in \mathbb{R}^+$  and any  $h \geq 3$ , there exists a connected, bipartite graph  $\mathcal{H}_{t,h,\gamma}$  of treewidth  $t$  such that there cannot be an algorithm solving the EXACT WEIGHT COLORED SUBGRAPH ISOMORPHISM problem on pattern graph  $\mathcal{H}_{t,h,\gamma}$  for instances with maximum weight  $W = \Theta(n^\gamma)$  in time  $O(n^{t-\varepsilon}W)$ , unless the  $h$ -uniform HYPERCLIQUE hypothesis fails.*
3. *For each  $t \geq 1$  and each  $\gamma \in \mathbb{R}^+$ , there exists a connected, bipartite graph  $\mathcal{H}_{t,\gamma}$  of treewidth  $t$  such that there cannot be an algorithm solving the EXACT WEIGHT COLORED SUBGRAPH ISOMORPHISM problem on pattern graph  $\mathcal{H}_{t,\gamma}$  for instances with maximum weight  $W = \Theta(n^\gamma)$  in time  $O(n^{(t+1)\omega/3-\varepsilon}W^{\omega/3})$ , unless the CLIQUE hypothesis fails.*

The following implies Theorem 7 from the results section.

► **Theorem 31** (Theorem 30 for pathwidth). *Parts 2 and 3 of Theorem 30 also hold when replacing the treewidth  $t$  by the pathwidth  $p$ . Part 1 does not hold.*

We remark that by the algorithm presented in Theorem 38, we cannot hope to obtain a lower bound as in part 1 of Theorem 30 for the case of pathwidth.

**Proof (of theorem 30).** Note that by Proposition 29, it suffices to prove lower bounds for the node-weighted case.

We begin with **part 1** of the theorem. Let  $t \geq 3$  and  $\gamma \in \mathbb{R}^+$ , as well as  $3 \leq h \leq t$  be given. We apply the Weighted Lemma with

- some  $\varepsilon' > 0$  chosen later,
- some  $\beta' \in (0, 1) \cap \mathbb{Q}$  chosen later,
- $h' := h$ ,
- $r'_2 := t + 1$  and
- some arbitrary  $r'_1 \in \mathbb{N}$  with
  - $r'_1 > \frac{\hat{c}\beta'}{h}$  (this ensures that the running time  $O(n^{2h-1} + n^{\hat{c}\beta'k/(hr'_1)})$  of the reduction is equal to  $O(n^{k-\varepsilon})$  for some  $\varepsilon > 0$ , and can hence be ignored in the analysis) and
  - $r'_1 > \frac{\beta'(t+1)}{1-\beta'}$  (this ensures that  $\max\{n^{\beta'k/r'_1}, n^{(1-\beta')k/(t+1)}\} = n^{(1-\beta')k/(t+1)}$ ).

This produces a  $k \in \mathbb{N}$  and a reduction algorithm  $\mathcal{A}$  with the properties from the lemma. In particular, the reduction algorithm produces instances where the pattern graph  $H$  is a Twin Water Lily of order  $(r_1, r_2)$ , which we define to be our graph  $\mathcal{H}_{t,\gamma}$ .

Now suppose there is an algorithm for the EXACT WEIGHT COLORED SUBGRAPH ISOMORPHISM problem on pattern graph  $\mathcal{H}_{t,\gamma}$  running in time  $O(N^{t+1-\varepsilon}W)$  (the case  $O(N^{t+1}W^{1-\varepsilon})$  is analogous). We show that the  $h$ -uniform HYPERCLIQUE hypothesis fails by showing that there is an algorithm for  $h$ -uniform  $k$ -HYPERCLIQUE running in time  $O(n^{k-\varepsilon})$  for some  $\varepsilon > 0$ .

Given an  $h$ -uniform  $k$ -HYPERCLIQUE instance, we use algorithm  $\mathcal{A}$  to obtain an equivalent instance of EXACT WEIGHT COLORED SUBGRAPH ISOMORPHISM where the pattern graph is a



Twin Water Lily of order  $(r_1', t+1)$ , the preimages have size  $N = \max\{n^{\beta'k/r_1'}, n^{(1-\beta')k/(t+1)}\} = n^{(1-\beta')k/(t+1)}$ , and the maximum weight is  $W = \Theta(n^{(1+\varepsilon')\beta'k})$ .

First, we make sure that  $W = \Theta(N^\gamma)$  by choosing  $\beta'$  and  $\varepsilon'$  accordingly. Substituting, we get  $n^{(1+\varepsilon')\beta'k} = \Theta(n^{\gamma(1-\beta')k/(t+1)})$ , which is true if and only if

$$(1 + \varepsilon')\beta'k = \frac{\gamma(1 - \beta')k}{t + 1} \iff \frac{\beta'}{1 - \beta'} = \frac{\gamma}{(t + 1)(1 + \varepsilon')} \iff \varepsilon' = \frac{\gamma(1 - \beta')}{(t + 1)\beta'} - 1$$

Hence we choose  $\varepsilon'$  as such. However, to apply the Weighted Lemma, we must have  $\varepsilon' \in (0, 1)$ . Hence we get the following two constraints for  $\beta'$ :

$$\begin{aligned} \frac{\gamma(1 - \beta')}{(t + 1)\beta'} - 1 > 0 &\iff \beta' < \frac{\gamma}{(t + 1) + \gamma} \\ \frac{\gamma(1 - \beta')}{(t + 1)\beta'} - 1 < 1 &\iff \beta' > \frac{\gamma}{2(t + 1) + \gamma} \end{aligned}$$

We incorporate these constraints later.

Now we need to ensure that the new running time we get is also small. We solve the EXACT WEIGHT COLORED SUBGRAPH ISOMORPHISM instance in time  $O(N^{t+1-\varepsilon}W) = O(N^{t+1+\gamma-\varepsilon}) = O(n^{(t+1+\gamma-\varepsilon)(1-\beta')k/(t+1)})$ . Hence we get the following additional constraints on  $\beta'$ :

$$\frac{(t + 1 + \gamma - \varepsilon)(1 - \beta')}{t + 1} < 1 \iff 1 - \beta' < \frac{t + 1}{t + 1 + \gamma - \varepsilon} \iff \beta' > \frac{\gamma - \varepsilon}{(t + 1) + \gamma - \varepsilon}$$

Combining these three constraints on  $\beta'$ , we get

$$\max \left\{ \frac{\gamma - \varepsilon}{(t + 1) + \gamma - \varepsilon}, \frac{\gamma}{2(t + 1) + \gamma} \right\} < \beta' < \frac{\gamma}{(t + 1) + \gamma}$$

Clearly, it is always possible to choose a  $\beta' \in (0, 1) \cap \mathbb{Q}$  such that this is true.

**Part 2** of the theorem is very much analogous. Note that the loss of the 1 in the exponent is due to the weaker bound in Proposition 20.

**Part 3** is completely analogous for the case  $t \geq 2$ ; we simply always choose  $h = 2$ . The  $\omega/3$  in the bound comes from the CLIQUE hypothesis.

However, a small trick has to be used for the case  $t = 1$ , since a 2-wide Twin Water Lily of order  $(r_1, 2)$  has treewidth 2, not 1. To get the better lower bound, we have to slightly modify the proof of the Weighted Lemma for  $h = 2$  in step 3. Instead of replacing the edges between vertices of  $S_2$  by intermediate vertices, we simply leave them as-is. Now the resulting graph is not a Twin Water Lily anymore, but does always have treewidth  $r_2$ . The rest of the proof is analogous. ◀

Finally, we prove the same theorem for pathwidth.

**Proof (of Theorem 31).** Proving part 2 is exactly analogous to part 2 of the theorem for treewidth.

Now remember that we needed a slight modification of the proof of the Weighted Lemma for part 3 of the theorem for treewidth for  $t = 1$ . For part 3 of the theorem for pathwidth, we actually need that modification for all  $t$ , i.e. we always leave  $S_2$  as-is in step 3. The rest of the proof is analogous. ◀

### 4.3.1 Subset Sum

We remark that with basically the same technique as is used to prove the Weighted Lemma, we can also prove a lower bound on the SUBSET SUM problem<sup>7</sup>.

► **Theorem 32.** *For no  $\varepsilon > 0$  can there be an algorithm which solves SUBSET SUM in time  $O(T^{1-\varepsilon} \text{poly}(n))$  unless the  $h$ -uniform HYPERCLIQUE hypothesis fails for all  $h \geq 3$ .*

In [4], a slightly better lower bound for SUBSET SUM is proven under SETH: They prove that unless SETH fails, SUBSET SUM cannot have an algorithm running in time  $O(T^{1-\varepsilon} 2^{o(n)})$ .

We briefly discuss other current algorithms and lower bounds for SUBSET SUM. The SUBSET SUM problem has a well-known  $O(Tn)$  time algorithm using dynamic programming [60]. Very recently, surprising new algorithms with running time  $\tilde{O}(\sqrt{n}T)$  [46, 47] and  $\tilde{O}(T+n)$  [24] have been shown, the latter matching several conditional lower bounds from SETH [4], SET COVER [33],  $k$ -CLIQUE (observed in [24] via techniques from [7]), and now from HYPERCLIQUE. The algorithm in [24] is slightly simplified in [45], with improvements in the log factors of the running time. The SUBSET SUM problem is closely related to the  $k$ -SUM problem (see also Appendix A.1).

There are two ways to see why Theorem 32 is true, and they both more or less lead to the same reduction. Both involve first reducing a HYPERCLIQUE instance to a  $k$ -SUM instance, after which a well-known reduction from  $k$ -SUM to SUBSET SUM can be used. We describe the latter reduction formally in Appendix A.1.

The first way to see the result is a generalization of the reduction from  $k$ -CLIQUE to  $k$ -SUM described by [7]. Instead of encoding edges with only two endpoints in the weights, we encode hyperedges. The second way to see the result (as stated, they lead to the same reduction) is via a slight modification of the Weighted Lemma to encode everything in the weighted part. We give details for the second way in Appendix A.2.

## 5 Algorithmic Results

### 5.1 $k$ -Wise Matrix Products

In our algorithms, the following generalization of matrix multiplication to tensors is both a crucial building block and a bottleneck. It was defined in its general form in [42] and explored further algorithmically in [51].

Given  $k$  tensors  $A^1, \dots, A^k$  of order  $k$  with dimensions  $\overbrace{n \times \dots \times n}^{k \text{ times}}$ , we define the  **$k$ -wise matrix product**  $\text{MP}_k(A^1, \dots, A^k)$  to be the tensor given by

$$\text{MP}_k(A^1, \dots, A^k)[i_1, \dots, i_k] := \sum_{\ell \in [n]} A^1[\ell, i_2, \dots, i_k] \cdot A^2[i_1, \ell, i_3, \dots, i_k] \cdots A^k[i_1, \dots, i_{k-1}, \ell]$$

Clearly, for  $k = 2$  this product is exactly matrix multiplication. There is also a boolean version of this generalized matrix product, just as there is a boolean version of the standard matrix product. In this boolean version, the tensors contain truth values (or equivalently 0/1 values) and the sum is replaced by an OR, while the products are replaced by ANDs.

We briefly discuss the computational complexity of  $k$ -wise matrix products. They can trivially be computed in time  $O(n^{k+1})$  for all  $k$ , and in time  $O(n^\omega)$  for  $k = 2$  via techniques

<sup>7</sup> Defined as: Given a set  $A \subseteq \mathbb{N}$  of  $n$  numbers and a target  $T \in \mathbb{N}$ , determine whether  $\exists B \subseteq A : \sum_{b \in B} b = T$ .

originating from Strassen [65] (for a history and introduction, see Bläser [17]). Unfortunately, as Lincoln, Williams and Williams observe in [51], it is impossible that faster Strassen-like algorithms with running time  $O(n^{k+1-\varepsilon})$  with  $\varepsilon > 0$  exist for  $k \geq 3$ : Just as the operation of  $n \times n$  by  $n \times n$  matrix multiplication has a corresponding order 3 tensor of dimensions  $n \times \dots \times n$ , there is an order  $k+1$  tensor of dimension  $n \times \dots \times n$  corresponding to the  $k$ -wise matrix product. For  $k \geq 3$ , this tensor has border rank  $n^{k+1}$ , i.e. it cannot be expressed as the limit of a sequence of tensors of rank smaller than  $n^{k+1}$ .

## 5.2 $k$ -Wise Tree Decompositions

In our algorithms, we use a structured form of tree decompositions that allow us to apply  $k$ -wise matrix products very easily to solve the COLORED SUBGRAPH ISOMORPHISM problems (either weighted or unweighted) on them. Any tree decomposition can be converted to this structured form without changing its width. Its definition is very loosely based on the structured tree decompositions that [32] describes for treewidth 2 graphs. See figure 3 for a partial illustration of the structured form in the context of the algorithm for the unweighted problem.

► **Definition 33.** We call a tree decomposition  $\mathcal{T} = (T, \{X_t\}_{t \in V(T)})$  a  **$k$ -wise tree decomposition** if it satisfies the following requirements.  $T$  must be a rooted tree, and each of its nodes has one of three types: It is either an **intermediate-result node**, a  **$k$ -wise node**, or a **merge node**. Intermediate-result and merge nodes have bags of size  $k$ , while  $k$ -wise nodes have bag size  $k+1$ . We require the root to be an intermediate-result node, and all leaves to be merge nodes. Finally, the three types of nodes are defined as follows:

- (i) If  $t \in V(T)$  is an intermediate-result node, it has two children: a  $k$ -wise node and a merge node  $t'$  with  $X_t = X_{t'}$ .
- (ii) If  $t \in V(T)$  is a  $k$ -wise node, its parent  $\text{par}(t)$  is an intermediate-result node and its set of children  $\text{children}(t)$  consists of exactly  $k$  merge nodes. Furthermore, we can rename the nodes in its bag to  $X_t = \{v, u_1, \dots, u_k\}$  such that  $X_t = X_{\text{par}(t)} \cup \{v\}$  and  $\forall i \in [k] : \exists c(i) \in \text{children}(t) : X_t = X_{c(i)} \cup \{u_i\}$ .
- (iii) If  $t \in V(T)$  is a merge node, it has arbitrarily many children which must all be intermediate-result nodes. Furthermore, for each child  $t'$  we have  $X_t = X_{t'}$ .

We call intermediate-result and merge nodes **helper nodes**, and define  $\text{help}(T) \subseteq V(T)$  to be the set of all helper nodes.

Note that in particular, a  $k$ -wise tree decomposition has width  $k$ .

► **Lemma 34.** Let  $H$  be a graph and let  $\mathcal{T} = (T, \{X_t\}_{t \in V(T)})$  be a tree decomposition of  $H$  that has width  $\text{width}(\mathcal{T})$ . Then we can convert  $\mathcal{T}$  into a  $\text{width}(\mathcal{T})$ -wise tree decomposition  $\mathcal{T}' = (T', \{X_t\}_{t \in V(T')})$ . Furthermore,  $|V(T')| = \text{poly}(|V(T)|) \cdot \text{width}(\mathcal{T})$  and the conversion can be done in time  $O(\text{poly}(|V(T)|) \cdot \text{poly}(\text{width}(\mathcal{T})))$ .

**Proof.** First, we go from the input tree decomposition  $\mathcal{T}$  to one where all bags have size  $\text{width}(\mathcal{T}) + 1$ , and where for any adjacent nodes  $t, t' \in V(T)$ , we have  $X_t \cap X_{t'} = \text{width}(\mathcal{T})$ . To do this, we roughly follow the outline of an algorithm that [32] describes for treewidth 2 graphs. First, we merge adjacent nodes  $t, t' \in V(T)$  with  $X_t = X_{t'}$ . If there exists a node  $t$  with  $|X_t| \leq \text{width}(\mathcal{T}) + 1$  and a neighbor  $t'$  such that  $X_{t'} \not\subseteq X_t$ , then we simply add an element of  $X_{t'} \setminus X_t$  to  $X_t$ . Applying this rule exhaustively, we obtain a tree decomposition where all bags have exactly  $\text{width}(\mathcal{T}) + 1$ . Applying this rule exhaustively, the resulting tree decomposition has the desired properties. Now we take any adjacent nodes  $t, t' \in V(T)$  with

$|X_t \cap X_{t'}| < \text{width}(\mathcal{T})$  and, letting  $u \in X_t \setminus X_{t'}$  and  $v \in X_{t'} \setminus X_t$ , insert a vertex with bag  $(X_t \cup \{v\}) \setminus \{u\}$  between them.

Now we root  $T$  in an arbitrary node. For any  $t \in V(T)$  and any  $t' \in \text{children}(t)$ , we subdivide the edge  $tt'$ , and give the new node the bag  $X_t \cap X_{t'}$ . We call the newly inserted vertices small, and all other vertices large. The root node  $r'$  must be a large node. We give it a small parent by selecting an arbitrary subset  $X \subseteq X_{r'}$  of size  $\text{width}(\mathcal{T})$ , adding a new root  $r$  to  $T$  with bag  $X$  and making  $r'$  a child of  $r$ .

We now iterate over the large nodes in  $T$  (as it is now) in depth-first search order. For each large vertex  $t$ , we (temporarily) rename the vertices of its bag to  $X_t = \{v, u_1, \dots, u_{\text{width}(\mathcal{T})}\}$  such that its (small) parent has bag  $X_{\text{par}(t)} = \{u_1, \dots, u_{\text{width}(\mathcal{T})}\}$ . Furthermore, for each  $i$ , let  $C_i \subseteq \text{children}(t)$  be the (possibly empty) set of all children  $t'$  such that  $X_{t'} = X_t \setminus \{u_i\}$ . We add a node  $c(i)$  with bag  $X_t \setminus \{u_i\}$  to  $V(T)$  and connect all children  $t' \in C_i$  to  $c(i)$  instead of  $t$ . Finally, we make  $c(i)$  a child of  $t$ . We do something analogous for  $v$ , namely letting  $C_v \subseteq \text{children}(t)$  be the (possibly empty) set of all children  $t'$  such that  $X_{t'} = X_t \setminus \{v\}$ , we add a node  $c(v)$  with bag  $X_t \setminus \{v\}$  to  $V(T)$  and connect all children  $t' \in C_v$  to  $c(v)$  instead of  $t$ . However, we make  $c(v)$  a child of  $\text{par}(t)$ , not  $t$ . Now we are done: The newly added vertices  $c(i)$  (for all  $i$ ) and  $c(v)$  are merge nodes, their children are intermediate-result nodes, and the large nodes  $t$  are  $k$ -wise nodes. Note also that because we allowed the sets  $C_i$  (for all  $i$ ) and  $C_v$  to be empty, we now have that all leaves are merge nodes.

This concludes the construction. It is easily seen that the number of nodes of the new tree decomposition is  $\text{poly}(|V(T)|) \cdot \text{width}(\mathcal{T})$ , and that this conversion algorithm works in time  $O(\text{poly}(|V(T)|) \cdot \text{poly}(\text{width}(\mathcal{T})))$ .  $\blacktriangleleft$

### 5.3 Colored Subgraph Isomorphism for Bounded Treewidth

We begin by looking at the unweighted version of COLORED SUBGRAPH ISOMORPHISM. As was previously mentioned for Theorem 3, these results essentially follow from [12] and [32], but are now unified via a single technique.

► **Theorem 35.** *There is an algorithm which, given an arbitrary instance  $\phi = (H, G, f)$  of COLORED SUBGRAPH ISOMORPHISM, solves  $\phi$  in time*

1.  $O(n^{\text{tw}(H)+1} \text{poly}(k) + g(k))$  when  $\text{tw}(H) \geq 3$ ,
2.  $O(n^\omega \text{poly}(k) + g(k))$  when  $\text{tw}(H) = 2$ , where  $\omega$  is the exponent of matrix multiplication, and
3.  $O(n^2 \text{poly}(k) + g(k))$  when  $\text{tw}(H) = 1$ .

where  $k := |V(H)|$ ,  $n$  is the size of the preimages of  $f$ , and  $g$  is a computable function.

Obviously, a proof of this theorem suffices to prove Theorem 3, since we can simply plug the algorithm into Lemma 14.

**Proof (of part 1 of Theorem 35).** We describe an algorithm which, given  $G, H$  and  $f : V(G) \rightarrow V(H)$ , first calculates an optimal tree decomposition  $\mathcal{T}_{\text{initial}}$  of width  $\text{tw}(H)$  for  $H$  in time  $g(k)$  (via the algorithm by Bodlaender [19], see also Section 3.4), then finds a solution via dynamic programming over the tree decomposition. The algorithm that computes the optimal tree decomposition also ensures that its tree graph  $T_{\text{initial}}$  has size  $|T_{\text{initial}}| = \text{poly}(k)$ . By assumption,  $\text{tw}(H) \geq 3$ . We shorten  $\text{tw}(H)$  to  $\text{tw}$  in the following. We use a slightly more complicated framework than necessary, because it generalizes nicely to a proof of part 2 and to a proof of parts 1 and 2 of Theorem 36.

To make our algorithm as easy as possible, we begin by applying Lemma 34 to convert our tree decomposition into a  $\text{tw}$ -wise tree decomposition  $\mathcal{T} = (T, \{X_t\}_{t \in V(T)})$ . Since


$|T_{\text{initial}}| = \text{poly}(k)$ , the time this conversion takes is certainly negligible. Furthermore,  $|T| = \text{poly}(|T_{\text{initial}}|) \cdot \text{poly}(\text{width}(\mathcal{T}_{\text{initial}})) = \text{poly}(\text{poly}(k)) \cdot \text{tw} = \text{poly}(k)$ .

We now do dynamic programming over the tw-wise tree decomposition. Using notation and nomenclature from Section 3.2, we only store values for each configuration of the bags of helper nodes, not for  $k$ -wise nodes. In particular, for each helper node  $t \in \text{help}(T)$  we store the following function of finite domain from configurations of  $X_t$  to truth values. Remember that  $\text{ParSol}(S; I; J)$  is true if and only if  $S$  is a partial solution of  $I$  in  $J$ , for  $I \subseteq J \subseteq V(H)$ .

$$\begin{aligned} d_t &: \text{Conf}(X_t) \rightarrow \{\text{true}, \text{false}\} \\ d_t(R) &:= \text{ParSol}(R; X_t; V_t) \end{aligned} \tag{1}$$

i.e. we store for each configuration whether it is a partial solution of  $X_t$  in the cone  $V_t$ . We call these functions DP functions (where DP stands for dynamic programming). Since there are  $n^{\text{tw}}$  many configurations for  $X_t$ , each DP function  $d_t$  can be specified using  $n^{\text{tw}}$  many bits.

We calculate the DP functions  $d_t$  for all  $t \in \text{help}(T)$  in a bottom-up manner. The overall picture of the algorithm is very simple: At a merge node  $t$ , we take the DP functions of all children and do a pointwise AND. At an intermediate-vertex node  $t$ , we first calculate a tw-wise matrix product for its tw-wise node child, then AND the result with the DP function of its merge node child. See figure 3 for a conceptual illustration.



■ **Figure 3** Partial sketch of the tree  $T$  of the tw-wise tree decompositions. Colored arrows represent operations of the algorithms, with red indicating that the result of the tw-wise operations is calculated via a tw-wise matrix product, and green indicating that the result of that subtree is ANDed pointwise with the result of all other subtrees.

We now describe the algorithm in detail. We begin with the leaves of  $T$ , which must be merge nodes. Hence let  $t$  be a merge node with no children. Since  $X_t = V_t$ , any configuration of  $X_t$  is a configuration of  $V_t$ . Thus, to calculate  $d_t(R)$  as in equation 1, we simply have

to check whether  $R$  is a valid configuration. The latter can be done in time  $\text{poly}(k)$ , which leads to a total time of  $O(n^{\text{tw}} \text{poly}(k))$  per leaf.

Now let  $t \in \text{help}(T)$  be an inner node of  $T$ . There are two cases: either  $t$  is an intermediate-result node, or  $t$  is a merge node.

**$t$  is a merge node:** Let  $t$  be a merge node. We continue to denote the set of children of  $t$  by  $\text{children}(t)$ . For each configuration  $R$  we have

$$\begin{aligned} d_t(R) &= \text{ParSol}(R; X_t; V_t) = \text{ParSol}\left(R; X_t; \bigcup_{c \in \text{children}(t)} V_c\right) \\ &= \bigwedge_{c \in \text{children}(t)} \text{ParSol}(R; X_t; V_c) = \bigwedge_{c \in \text{children}(t)} d_c(R) \end{aligned}$$

Since  $t$  has at most  $\text{poly}(k)$  children, and since there are at most  $n^{\text{tw}}$  many possible configurations  $R$ , this can be done in time  $O(n^{\text{tw}} \text{poly}(k))$  per merge node.

**$t$  is an intermediate-result node:** Let  $t$  be an intermediate-result node with  $\text{tw}$ -wise child node  $t'$  and merge node  $t''$ . By definition we have  $X_t = X_{t''}$ . Furthermore  $|\text{children}(t')| = \text{tw}$ , and we can rename the nodes in  $X_{t'}$  to  $X_{t'} = \{v, u_1, \dots, u_{\text{tw}}\}$  such that  $X_{t'} = X_t \cup \{v\}$  and  $\forall i \in [k] : \exists c(i) \in \text{children}(t') : X_{t'} = X_{c(i)} \cup \{u_i\}$ . We have that

$$\begin{aligned} d_t(R) &= \text{ParSol}(R; X_t; V_t) \\ &= \text{ParSol}(R; X_{t''}; V_{t''}) \wedge \\ &\quad \exists v' \in f^{-1}(v) : \forall i \in [\text{tw}] : \text{ParSol}((R \cup \{v \mapsto v'\})|_{X_{c(i)}}; X_{c(i)}; V_{c(i)}) \\ &= d_{t''}(R) \wedge \exists v' \in f^{-1}(v) : \forall i \in [\text{tw}] : d_{c(i)}((R \cup \{v \mapsto v'\})|_{X_{c(i)}}) \end{aligned} \quad (*)$$

We show how to calculate the values of  $d_t$  via a  $\text{tw}$ -wise boolean matrix product with 0-1-tensors of dimension  $n \times \dots \times n$ . Note that we will use truth values and 0/1 interchangeably to declutter notation. True is interchangeable with 1, false with 0.

For convenience, all tensors from this point onward are indexed via configurations, where each dimension is indexed by a single vertex. Formally, we call a tensor  $A$  **indexed by configurations of  $X = \{v_1, \dots, v_h\} \subseteq V(H)$  with ordering  $(v_{i_1}, \dots, v_{i_h})$**  (where  $\{v_{i_1}, \dots, v_{i_h}\} = \{v_1, \dots, v_h\}$ ) when it is an order  $h$  tensor of dimension  $n \times \dots \times n$ . Abusing notation, we rename the vertices of  $f^{-1}(v_i)$  for each  $v_i \in X$  to  $\{1, \dots, n\}$  and use them as if they were numbers. Now for a configuration  $R$  of  $X$  we define  $A[R] = A[R(v_{i_1}), \dots, R(v_{i_h})]$ .

For our  $\text{tw}$ -wise matrix product, we have  $\text{tw}$  input tensors, specifically one tensor  $p_i$  for each  $i \in [\text{tw}]$ .  $p_i$  is indexed by configurations of  $X_{c(i)} = X_{t'} \setminus \{u_i\}$  with ordering  $(u_1, \dots, u_{i-1}, v, u_{i+1}, \dots, u_{\text{tw}})$ . Specifically, for a configuration  $R$  of  $X_{c(i)}$  we define  $p_i[R] := d_{c(i)}(R)$  (it is still a 0-1-tensor, remember that truth values and 0/1 are interchangeable).

Now we calculate the  $\text{tw}$ -wise matrix product  $p_{\text{res}} := \text{MP}_{\text{tw}}(p_1, \dots, p_{\text{tw}})$ , which we use as a tensor indexed by configurations of  $X_t$  with ordering  $(u_1, \dots, u_{\text{tw}})$ . Hence, for a configuration  $R$  of  $X_t$ , we have

$$\begin{aligned} p_{\text{res}}[R] &= \text{MP}_{\text{tw}}(p_1, \dots, p_{\text{tw}})[R(u_1), \dots, R(u_{\text{tw}})] \\ &= \bigvee_{\ell \in [n]} p_1[\ell, R(u_2), \dots, R(u_{\text{tw}})] \wedge \dots \wedge p_{\text{tw}}[R(u_1), \dots, R(u_{\text{tw}-1}), \ell] \\ &= \exists v' \in f^{-1}(v) : \forall i \in [\text{tw}] : p_i[R(u_1), \dots, R(u_{i-1}), v', R(u_{i+1}), \dots, R(u_{\text{tw}})] \\ &= \exists v' \in f^{-1}(v) : \forall i \in [\text{tw}] : d_{c(i)}((R \cup \{v \mapsto v'\})|_{X_{c(i)}}) \end{aligned}$$

From this, we can directly calculate  $d_t(R)$  via equation (\*): We have  $d_t(R) = d_{t''}(R) \wedge p_{\text{res}}[R]$ . Indeed, this gives a very simple algorithm for computing  $d_t$ . The tw-wise matrix product for  $p_{\text{res}}$  can be done in time  $O(n^{\text{tw}+1})$ , and after that calculating  $d_t$  is a simple pointwise AND with  $d_{t''}$ , which takes time  $O(n^{\text{tw}})$ . Hence overall we have running time  $O(n^{\text{tw}+1})$ .

Since the tree decomposition has  $\text{poly}(k)$  nodes, the overall time the algorithm takes is  $O(n^{\text{tw}+1} \text{poly}(k))$ . Now we simply have to extract the answer from DP function of the root node. Let  $r$  be the root node. We have that  $d_r(R)$  is true if and only if there is a partial solution of  $X_r$  in the cone  $V_r = V(H)$ . Hence it is true only if  $R$  can be expanded to a solution. Thus, the input is a YES-instance for COLORED SUBGRAPH ISOMORPHISM if and only if there is an  $R$  such that  $d_r(R)$  is true. ◀

So far, we have only looked at the case that  $\text{tw}(H) \geq 3$ . However, this exact algorithm also achieves the second result.

**Proof (of part 2 of Theorem 35).** Note that in the algorithm for part 1, all steps run in time  $O(n^{\text{tw}(H)} \text{poly}(k))$ , except for the  $\text{tw}(H)$ -wise matrix product, which can be done in time  $O(n^{\text{tw}(H)+1})$ . However, for  $\text{tw}(H) = 2$ ,  $\text{tw}(H)$ -wise matrix product is exactly matrix multiplication, which runs in time  $O(n^\omega)$ . Thus we obtain our second result. ◀

The third result with  $\text{tw}(H) = 1$  cannot be achieved by this algorithm directly. We shortly outline why. Rooting  $H$  in some arbitrary node, consider the tree decomposition that takes exactly the edges of  $H$  as bags, and constructs  $T$  such that two nodes are connected by an edge if and only if their bags have a non-empty intersection. Now consider a graph  $H$  which contains nodes  $u, v, w$  such that  $u$  is the parent of  $v$  and  $v$  is the parent of  $w$ . Let  $t, t' \in T$  be such that  $X_t = \{u, v\}$ ,  $X_{t'} = \{v, w\}$ . But now the edge  $vw$  is not covered by any of the subsets  $X_{t'} \setminus \{v\}$  or  $X_t \setminus \{w\}$  and is thus not considered in the algorithm at all. This leads to the algorithm failing. For  $\text{tw}(H) \geq 2$ , this does not happen, since any pair of nodes is contained in a bag  $X_t$  of size  $\geq 3$ , hence there is some  $u$  such that the edge is covered by the subset  $X_t \setminus \{u\}$ .

Thus, to obtain our third result, we must employ a different technique. However, this part of the theorem turns out to be easy.

**Proof (of part 3 of Theorem 35).** We only sketch the result, since it is easy to see. It suffices to employ the trivial dynamic programming solution on the tree  $G$ , which already has a running time of  $O(n^2 \text{poly}(k))$ . ◀

## 5.4 Exact Weight Colored Subgraph Isomorphism for Bounded Treewidth

We now move on to the weighted version of COLORED SUBGRAPH ISOMORPHISM. Specifically, we show how to solve the EXACT WEIGHT COLORED SUBGRAPH ISOMORPHISM problem for bounded-treewidth pattern graphs using dynamic programming. Remember that the instances we consider here may be either node- or edge-weighted.

We also remark that the restriction of the target weight  $T$  to zero in the EXACT WEIGHT SUBGRAPH ISOMORPHISM problems is simply for ease of discussion. We use this version to both simplify our algorithms and circumvent any problems that might arise from having  $T$  be part of the input. Trivially, all of our results also hold for the problem where a target  $T = O(W)$  is given.

► **Theorem 36.** *There is an algorithm which, given an arbitrary instance  $\phi = (H, G, f, w)$  of the EXACT WEIGHT COLORED SUBGRAPH ISOMORPHISM problem, solves  $\phi$  in time*

1.  $O((n^{\text{tw}(H)+1}W + n^{\text{tw}(H)}W \log W) \text{poly}(k) + g(k))$  when  $\text{tw}(H) \geq 3$ ,
2.  $O((n^\omega W + n^2W \log W) \text{poly}(k) + g(k))$  when  $\text{tw}(H) = 2$ , and
3.  $O((n^2W + nW \log W) \text{poly}(k) + g(k))$  when  $\text{tw}(H) = 1$ .

where  $k := |H|$ ,  $n$  is the size of the preimages of  $f$ , and  $W$  is the maximum absolute weight in the image of  $w$ .

Again, note that this implies Theorem 5 via a simple application of Lemma 14 to the resulting algorithm.

Before we prove this, we establish an important lemma, showing that you can basically do the  $k$ -wise matrix product of tensors of polynomials faster than naively by utilizing the Fast Fourier Transform. Recall that a Laurent polynomial  $p \in \mathbb{C}[X, X^{-1}]$  is simply a polynomial which may have negative powers of the  $X$ .

► **Lemma 37.** *Let  $q \in \mathbb{N}$  tensors  $A^1, \dots, A^q$  of order  $q$  be given, each of dimensions  $n \times \dots \times n$  and such that each of their entries  $A^i_{j_1, \dots, j_q} \in \mathbb{C}[Z, Z^{-1}]$  ( $i \in [q]$  and  $\forall \ell \in [q] : j_\ell \in [n]$ ) is a Laurent polynomial of degree bounded by  $W$  in both the positive and negative direction. Then their  $q$ -wise matrix product can be computed in time*

1.  $O(n^{q+1}W + qn^qW \log W)$  for  $q \geq 3$  and
2.  $O(n^\omega W + n^2W \log W)$  for  $q = 2$ .

**Proof.** It suffices to prove the result for standard polynomials of degree bounded by  $2W$ , since we can shift the exponents of the polynomials such that they only have positive exponents, do the  $q$ -wise matrix product, then shift back.

We assume for now that  $q \geq 3$ . Our algorithm is a generalization of the Fast Fourier Transform algorithm for standard polynomials (e.g. [28]). Specifically, we evaluate each  $A^i$  at the set  $S$  of the  $2W$ -th roots of unity, obtaining  $q \cdot |S|$  tensors with complex entries. Since each entry of  $A^i$  is a polynomial, this can be done separately for each entry. Then, for each  $s \in S$ , we compute  $\text{MP}_q(A^1(s), \dots, A^q(s))$ , obtaining  $|S|$  tensors with complex entries. Obviously, these are exactly the evaluations of  $\text{MP}_q(A^1, \dots, A^q)$  at  $S$ . At this point we can use interpolation via the inverse Fast Fourier Transform separately for each entry to recover the result.

Evaluating all entries of the tensors  $A^i$  at the roots of unity can be done using the classic Fast Fourier Transform algorithm. Since there are  $q \cdot n^q$  such entries, each with polynomials of degree bounded by  $W$ , this runs in time  $O(qn^qW \log W)$ . Similarly, the interpolation of the result can be done by the inverse Fast Fourier Transform algorithm. Since there are  $n^q$  entries to interpolate, each with a degree bound of  $O(W)$ , this runs in time  $O(qn^qW \log W)$ . Finally, computing the  $q$ -wise matrix products for each root of unity can be done in  $|S| \cdot n^{q+1} = O(n^{q+1}W)$ . Thus, our total running time is  $O(n^{q+1}W + qn^qW \log W)$ .

For the case that  $q = 2$ , note that in the above algorithm, all steps except for the  $q$ -wise matrix product run in time  $O(n^qW \log W)$ . For  $q = 2$ , the  $q$ -wise matrix product is exactly matrix multiplication, which can be done in  $O(n^\omega)$ . Thus the  $O(W)$  matrix multiplications can be done in  $O(n^\omega W)$ . ◀

We use this lemma as an important tool in our proof of Theorem 36. The framework of the proof is somewhat analogous to that of Theorem 35, but instead of storing values for every configuration of bags of nodes in  $\text{help}(T)$ , they we store the values for each configuration and each achievable weight. The computation of the dynamic programming table entries is slightly more complex, with some shifting of the entries being required.



**Proof (of part 1 of Theorem 36).** Due to Proposition 29, we only need to consider the case of edge weights. We describe an algorithm which takes as inputs  $G, H, f : V(G) \rightarrow V(H)$  and a weight function  $w$  describing the edge weights. It calculates an optimal tree decomposition  $\mathcal{T}_1 := (T_1, \{X_t\}_{t \in V(T_1)})$  for  $H$  in time  $g(k)$  and then computes a solution via dynamic programming over the tree decomposition. By assumption,  $\text{tw} := \text{tw}(H) \geq 3$ . We continue using the notation and terminology (e.g. “configuration”) as described in Section 3.2.

We apply Lemma 34 to convert the tree decomposition into a tw-wise tree decomposition  $\mathcal{T}_2 = (T_2, \{X_t\}_{t \in V(T_2)})$ . Since  $|T_1| = \text{poly}(k)$ , the time this conversion takes is certainly negligible. Furthermore,  $|T_2| = \text{poly}(k)$ . To simplify our algorithm further, we also modify  $\mathcal{T}_2$  further. Specifically, we introduce a new type of helper node, a **binary-merge node**, which replaces merge nodes. It is a node that has either zero or two children, both with the same bag as itself. Those children are either binary-merge nodes or intermediate-result nodes. We convert  $\mathcal{T}_2$  into a new tree decomposition  $\mathcal{T} := (T, \{X_t\}_{t \in V(T)})$  that has binary-merge nodes instead of merge nodes by repeatedly taking any remaining merge node  $t$  with more than two children, splitting its set of children into two non-empty sets  $A$  and  $B$ , creating two new children  $t'$  and  $t''$  of  $t$  and making the vertices from  $A$  children of  $t'$  and the vertices of  $B$  children of  $t''$ . For any merge node  $t$  with only a single child, we create a new leaf with the same bag and let it be a child of  $t$ . Correctness and negligibility of the conversion time is immediate, and the new tree decomposition still has size  $\text{poly}(k)$ . We now do dynamic programming over  $\mathcal{T}$ .

In deviation from the proof of Theorem 35, our dynamic programming table is structured differently. Instead of having only a single function of finite domain for each helper node  $t$ , we have one function for each  $t$  and each achievable weight. Since the instance is edge-weighted, the achievable weights must all lie in  $\mathcal{W} := \{-k^2W, \dots, k^2W\}$ .

Specifically, for each helper node  $t \in \text{help}(T)$  and each weight  $W' \in \mathcal{W}$  we store the following function of finite domain from configurations of  $X_t$  to truth values. Remember that  $\text{ParSolE}(R; I; J; W)$  is true if and only if  $R$  is a partial solution of  $I$  in  $J$  with an extension of weight  $W$ .

$$d_{t, W'} : \text{Conf}(X_t) \rightarrow \{\text{true}, \text{false}\}$$

$$d_{t, W'}(R) := \text{ParSolE}(R; X_t; V_t; W')$$

i.e. we store for each configuration whether it is a partial solution of  $X_t$  in the cone  $V_t$  that has an extension of weight  $W'$ . We call these functions DP functions (where DP stands for dynamic programming). Since there are  $n^{\text{tw}}$  many configurations for  $X_t$ , each DP function  $d_{t, W'}$  can be specified using  $n^{\text{tw}}$  many bits.

We calculate the DP functions  $d_t$  for all  $t \in \text{help}(T)$  in a bottom-up manner. The overall picture of the algorithm is as follows: At a non-leaf binary-merge node  $t$ , we take the DP functions of both children and do a boolean convolution. At an intermediate-vertex node  $t$ , we first convert the DP functions of the children of its tw-wise child to tensors with Laurent polynomials as entries, calculate their tw-wise matrix product, convert it back to a DP function and then do a boolean convolution with the DP function of its merge node child.

We now describe the algorithm in detail, beginning with the leaves of  $T$ . Hence  $t$  be a merge node with no children. Since  $X_t = V_t$ , we have that  $d_{t, W'}(R)$  is true if and only if  $R$  is a valid configuration and  $W' = 0$ . Clearly this takes total time at most  $O(n^{\text{tw}}W \text{poly}(k))$ .

Now let  $t \in \text{help}(T)$  be an inner node of  $T$ . There are two cases: either  $t$  is an intermediate-result node, or  $t$  is a merge node.

**$t$  is a binary-merge node:** Let  $t$  be a binary-merge node with children  $t'$  and  $t''$ . For each

configuration  $R$  we have

$$\begin{aligned} d_{t,W'}(R) &= \text{ParSolE}(R; X_t; V_t; W') = \text{ParSolE}(R; X_t; V_{t'} \cup V_{t''}; W') \\ &= \exists W_1 : \text{ParSolE}(R; X_{t'}; V_{t'}; W_1) \wedge \text{ParSolE}(R; X_{t''}; W' - W_1) \\ &= \exists W_1 : d_{t',W_1}(R) \wedge d_{t'',W'-W_1}(R) \end{aligned}$$

This can be calculated using a boolean convolution for each configuration. Specifically, for each configuration  $R$  we define the finitely supported functions  $f_R, g_R : \mathbb{Z} \rightarrow \{\text{true}, \text{false}\}$  as  $f_R(x) := d_{t',x}(R)$  if  $x \in \mathcal{W}$  and false otherwise, and analogously  $g_R(x) := d_{t'',x}(R)$  if  $x \in \mathcal{W}$  and false otherwise. We now use the boolean convolution of these functions, defined as  $(f_R * g_R)(x) := \bigvee_{z=-\infty}^{\infty} f_R(z) \wedge g_R(x - z)$ . We have that  $(f_R * g_R)(W') = (\exists W_1 : f_R(W_1) \wedge g_R(W' - W_1)) = d_{t,W'}(R)$ . Since  $f_R$  and  $g_R$  only have non-zero values in a range of size  $O(W \text{poly}(k))$ , the boolean convolution  $f_R * g_R$  can be calculated in time  $O(W \log W \text{poly}(k))^8$ .

Since there are at most  $n^{\text{tw}}$  many possible configurations  $R$ ,  $d_{t,W'}$  can be calculated for all  $W'$  in time  $O(n^{\text{tw}} W \log W \text{poly}(k))$ . Hence this is the maximum running time we need per binary-merge node.

**$t$  is an intermediate-result node:** Let  $t$  be an intermediate-result node with  $\text{tw}$ -wise child node  $t'$  and merge node  $t''$ . By definition we have  $X_t = X_{t''}$ . Furthermore  $|\text{children}(t')| = \text{tw}$ , and we can (temporarily) rename the nodes in  $X_{t'}$  to  $X_{t'} = \{v, u_1, \dots, u_{\text{tw}}\}$  such that  $X_{t'} = X_t \cup \{v\}$  and  $\forall i \in [k] : \exists c(i) \in \text{children}(t') : X_{t'} = X_{c(i)} \cup \{u_i\}$ .

First, we deal with calculating the weight of an extension for a configuration of the bag of the  $\text{tw}$ -wise node. Let  $R$  be a partial solution of  $X_{t'}$  in  $V_{t'}$  and let  $S$  be an extension of  $R$  (i.e. a compatible configuration of  $V_t$ ). Notice that we get the following formula for the weight of the extension  $S$ :

$$w_{\text{ext}}(S, R) \stackrel{\text{def.}}{=} w(S) - w(S|_{X_{t'}}) = \underbrace{\left( \sum_{i \in [\text{tw}]} w_{\text{ext}}(S|_{V_{c(i)}}, S|_{X_{c(i)}}) \right)}_{\text{Part A}} + \underbrace{w_{\text{ext}}(S|_{X_{t'}}, S|_{X_{t'} \setminus \{v\}})}_{\text{Part B}}$$

Part A is the combined weight of the edges which are not in  $X_{t'}$ . Compared to the left hand side, it is missing the weight of all edges going from  $v$  to  $X_{t'} \setminus \{v\}$ , which is exactly what is then added in Part B. Recall that  $w_{\text{ext}}(S|_{X_{t'}}, S|_{X_{t'} \setminus \{v\}}) = \sum_{u \in X_{t'} \setminus \{v\}} w(S(v)S(u))$ . Hence we can calculate the DP function for  $t$  as follows for each configuration  $R$ :

$$\begin{aligned} d_{t,W'}(R) &= \text{ParSolE}(R; X_t; V_t; W') \\ &= \exists W_v : \text{ParSolE}(R; X_{t''}; V_{t''}; W_v) \wedge \exists W_1, \dots, W_{\text{tw}} : \exists v' \in f^{-1}(v) : \\ &\quad w_{\text{ext}}((R \cup \{v \mapsto v'\}), R) + \sum_{i=1}^{\text{tw}} W_i = W' - W_v \wedge \\ &\quad \forall i \in [\text{tw}] : \text{ParSolE}((R \cup \{v \mapsto v'\})|_{X_{c(i)}}; X_{c(i)}; V_{c(i)}; W_i) \\ &= \exists W_v : d_{t'',W_v}(R) \wedge \exists W_1, \dots, W_{\text{tw}} : \exists v' \in f^{-1}(v) : \\ &\quad w_{\text{ext}}((R \cup \{v \mapsto v'\}), R) + \sum_{i=1}^{\text{tw}} W_i = W' - W_v \wedge \\ &\quad \forall i \in [\text{tw}] : d_{c(i),W_i}((R \cup \{v \mapsto v'\})|_{X_{c(i)}}) \end{aligned} \tag{†}$$

<sup>8</sup> This is a standard result which can be achieved e.g. by using the Fast Fourier Transform.

We show how to calculate the values of  $d_{t,W'}$  for all  $W'$  via a tw-wise matrix product of Laurent polynomials. We continue using the notation and nomenclature for indexing tensors via configurations, as described in the proof of part 1 of Theorem 35, and continue using truth values and 0/1 interchangeably.

We now define one 0-1-tensors  $p_{i,W'}$  for each  $i \in [\text{tw}]$  and  $W' \in \mathcal{W}$ , which is indexed by configurations of  $X_{c(i)} = X_{t'} \setminus \{u_i\}$  with ordering  $(u_1, \dots, u_{i-1}, v, u_{i+1}, \dots, u_{\text{tw}})$ . For a configuration  $R$  of  $X_{c(i)}$  we define

$$p_{i,W'}[R] := \begin{cases} d_{c(1),W' - \sum_{i=2}^{\text{tw}} w(R(v)R(u_i))}(R) & \text{if } i = 1 \\ d_{c(2),W' - w(R(v)R(u_1))}(R) & \text{if } i = 2 \\ d_{c(i),W'}(R) & \text{otherwise} \end{cases}$$

Note that the weight of the part B above is now encoded in these tensors, specifically in the ones for  $i = 1$  and  $i = 2$ . In particular, we can expand equation (†) as

$$\begin{aligned} &= \exists W_v : d_{t',W_v}(R) \wedge \exists W_1, \dots, W_{\text{tw}} : \sum_{i=1}^{\text{tw}} W_i = W' - W_v \wedge \\ &\quad \exists v' \in f^{-1}(v) : \forall i \in [\text{tw}] : p_{i,W_i}((R \cup \{v \mapsto v'\})|_{X_{c(i)}}) \end{aligned}$$

Analogously to the computation of  $p_{\text{res}}[R]$  in the proof of part 1 of Theorem 35, we have

$$= \exists W_v : d_{t',W_v}(R) \wedge \exists W_1, \dots, W_{\text{tw}} : \sum_{i=1}^{\text{tw}} W_i = W' - W_v \wedge MP_{\text{tw}}(p_{1,W_1}, \dots, p_{\text{tw},W_{\text{tw}}})[R]$$

We now show how to compute tensors  $D_{W'}$  for each  $W' \in \mathcal{W}$  such that  $D_{W'}(R) = \exists W_1, \dots, W_{\text{tw}} : \sum_{i=1}^{\text{tw}} W_i = W' - W_v \wedge MP_{\text{tw}}(p_{1,W_1}, \dots, p_{\text{tw},W_{\text{tw}}})[R]$ , which would simplify the above to

$$= \exists W_v : d_{t',W_v}(R) \wedge D_{W' - W_v}(R)$$

and hence make it computable via a boolean convolution in time  $O(n^{\text{tw}} W \log W \text{ poly}(k))$  (given the tensors  $D_{W'}$ ).

We want to compute  $D_{\widetilde{W}}$  for each  $\widetilde{W} \in \mathcal{W}$ . Note how, when defining the OR of tensors to be calculated entrywise, we have

$$D_{W'} = \bigvee_{\substack{W_1, \dots, W_{\text{tw}} \in \mathcal{W} \\ W_1 + \dots + W_{\text{tw}} = W'}} MP_{\text{tw}}(p_{1,W_1}, \dots, p_{\text{tw},W_{\text{tw}}(H)})$$

Indeed, the right-hand side can be calculated using a single tw-wise matrix product of tensors with polynomials as entries. Let  $T_{\text{tw}}$  be the group of order-tw( $H$ ) tensors of dimensions  $n \times \dots \times n$  with entries from  $\{\text{true}, \text{false}\}$ . Define  $T_{\text{tw}}[X, X^{-1}]$  as the group of Laurent polynomials with elements of  $T_{\text{tw}(H)}$  as coefficients. Note how elements of  $T_{\text{tw}}[X, X^{-1}]$  may also be viewed as tensors with Laurent polynomials as entries, or as functions  $f : \{0, 1\} \rightarrow T_{\text{tw}(H)}$  when using the usual definition of scalar-tensor AND.

We define  $p_i(X) := \sum_{j \in \mathcal{W}} p_{i,j} \cdot X^j \in T_{\text{tw}(H)}[X, X^{-1}]$  for all  $i$ . Viewing them as tensors of polynomials, we may calculate their tw( $H$ )-wise matrix product. Viewing  $MP_{\text{tw}}(p_1, \dots, p_{\text{tw}})$  as polynomial again, it can be easily seen that the coefficient tensor for  $X^{W'}$  is exactly

$$\sum_{\substack{W_1, \dots, W_{\text{tw}} \in \mathcal{W} \\ W_1 + \dots + W_{\text{tw}} = W'}} MP_{\text{tw}}(p_{1,W_1}, \dots, p_{\text{tw},W_{\text{tw}}})$$

Thus, to compute  $D_{W'}$  for all  $W'$ , it suffices to compute  $\text{MP}_{\text{tw}}(p_1, \dots, p_{\text{tw}})$ .

Hence, we have reduced the problem of calculating the tensors  $D_{\tilde{W}}$  to computing the  $\text{tw}(H)$ -wise matrix product of tensors whose entries are Laurent polynomials of degree bounded (in both directions) by  $O(W \text{poly}(k))$ . By Lemma 37, this can be done in time  $O((n^{\text{tw}(H)+1}W + n^{\text{tw}(H)}W \log W) \text{poly}(k))$ .

Hence we have an overall running time of  $O((n^{\text{tw}(H)+1}W + n^{\text{tw}(H)}W \log W) \text{poly}(k))$  per intermediate-result node.

Again, since the tree decomposition has  $\text{poly}(k)$  nodes, the overall time the algorithm takes is  $O((n^{\text{tw}(H)+1}W + n^{\text{tw}(H)}W \log W) \text{poly}(k))$ .

After calculating  $d_{t', W'}$  for each  $t' \in T, W' \in \mathcal{W}$ , outputting the result is simple. We simply extract the answer from the DP function of the root node. Let  $r$  be the root of  $T$ . We output YES if and only if  $\exists R : d_{r, -w(R)}(R) = \text{true}$ . By the definition of  $d_{t', -w(R)}$ , this is the case if and only if  $R$  is a partial solution of  $X_r$  in the cone  $V_r = V(H)$  with an extension of weight  $-w(R)$ . This is the case if and only if  $R$  can be expanded to a configuration for all of  $V(H)$  such that its total weight is  $-w(R) + w(R) = 0$ . Hence,  $\exists R : d_{t', -w(R)}(R) = \text{true}$  if and only if the instance has a solution. Checking whether such an  $R$  exists can obviously be done in time  $O(n^{\text{tw}(H)} \text{poly}(k))$ . ◀

We now consider the second part of the theorem. Similarly to the weighted case, it suffices to employ the algorithm from part 1.

**Proof (of part 2 of Theorem 36).** In the algorithm for part 1, all running times except for the  $\text{tw}(H)$ -wise matrix product are bounded by  $O(n^{\text{tw}}W \log W \text{poly}(k))$ . For  $\text{tw} = 2$ , the  $\text{tw}$ -wise matrix product is simply matrix multiplication, which can be done in  $O(n^\omega)$ . Thus, for  $\text{tw} = 2$  the running time is  $O(n^\omega W \text{poly}(k) + n^2 W \log W \text{poly}(k))$ . ◀

Finally, we come to the third part of the theorem. For reasons outlined in the proof of the previous theorem, the algorithm from part 1 does not work for  $\text{tw}(H) = 1$ . Again, however, the result turns out to be quite simple for this case.

**Proof (of part 3 of Theorem 36).** Again, we only sketch the result, since it is easy to see. A simple dynamic programming algorithm on trees can be applied, storing for each node which configurations together with which weights can be achieved. ◀

## 5.5 Colored Subgraph Isomorphism for Bounded Pathwidth

Surprisingly, COLORED SUBGRAPH ISOMORPHISM can be solved slightly faster on graphs of bounded pathwidth. This algorithm leverages rectangular matrix multiplication.

We briefly discuss (rectangular) matrix multiplication and current algorithms solving it. Multiplication of  $n \times n$  by  $n \times n$  matrices is perhaps the most ubiquitous open problem in computer science, with the central question being whether it can be done in  $O(n^2)$  time. The matrix multiplication exponent  $\omega$  has been slowly inching toward, but not quite reaching, a value of 2 over the last few decades. In our algorithms, however, we also multiply rectangular matrices. It turns out that the techniques used in these fast algorithms for square matrix multiplication can also be generalized to the rectangular case. In the following, let  $\text{MM}(s, r, t)$  denote the time needed to multiply a matrix of size  $r \times s$  with a matrix of size  $s \times t$ . We are mostly interested in the case that  $s = n, r = n$  and  $t = n^k$  for some  $k \in \mathbb{R}^+$ .

A simple, well-known result is that MM is both convex and symmetrical in its arguments (see e.g. [52, 64]). In particular, we have  $\forall x \in \mathbb{R}^+ : \text{MM}(n, n, n^{k+x}) \leq n^x \cdot \text{MM}(n, n, n^k)$  and  $\text{MM}(n, n, n^k) = \text{MM}(n, n^k, n)$ . Letting  $\omega(k) := \log_n \text{MM}(n, n, n^k)$ , we immediately get  $\forall k \geq 1 : \omega(k) < k + 1.373$  via the current bounds of  $\omega(1) = \omega < 2.373$  [49].

We can, however, do better: Le Gall [41] has shown that there are faster algorithms based on the Coppersmith-Winograd method [30, 49] used for square matrix multiplication. Among other values, he shows  $\omega(0.31) = 2$ ,  $\omega(2) < 3.26$ ,  $\omega(3) < 4.2$ ,  $\omega(4) < 5.18$  and  $\omega(5) < 6.16$  (see [41] for an extensive table of such values).

We now show how rectangular matrix multiplication can be used for bounded-pathwidth pattern graphs.

► **Theorem 38.** *There is an algorithm which, given an arbitrary instance  $\phi = (H, G, f)$  of COLORED SUBGRAPH ISOMORPHISM, solves  $\phi$  in time*

1.  $O(n^{\omega(\text{pw}(H)-1)} \text{poly}(k) + g(k))$  when  $\text{pw}(H) \geq 2$ , and
2.  $O(n^2 \text{poly}(k) + g(k))$  when  $\text{pw}(H) = 1$

where  $k := |H|$ ,  $n$  is the size of the preimages of  $f$ , and  $g$  is a computable function.

This theorem implies Theorem 6 from the results section.

**Proof. Part 2** of the theorem is trivial, since for any graph, its treewidth is smaller than its pathwidth. Hence, by application of Theorem 35, we achieve the desired running time.

For **part 1**, let  $G, H$  and  $f : V(G) \rightarrow V(H)$  be given. The algorithm first computes an optimal path decomposition  $\mathcal{P} = (P, \{X_t\}_{t \in V(T)})$  of  $H$  in time  $g(k)$  (see preliminaries), then does dynamic programming over  $\mathcal{P}$ .

To unify nomenclature and notation with the case of treewidth, we talk about a path as a tree rooted at one of its endpoints. As in the proof of Theorem 35, we modify the path decomposition  $\mathcal{P}$  to satisfy certain properties. Specifically, we wish to obtain the following properties:

1. The bags of the root and leaf of the path have size  $\text{pw}(H)$ , and every other bag has size  $\text{pw}(H) + 1$
2. For every  $t \in V(P)$  with child  $t'$ , we have  $|X_t \cap X_{t'}| = \text{pw}(H)$

These properties can be obtained with similar techniques as described in the proof of part 1 of Theorem 35.

We now do dynamic programming on this modified path decomposition. We only store values for each configuration of the separators  $X_t \cap X_{t'}$ . In particular, let  $t'$  be a node other than the root, and let  $t$  be its parent. As in Theorem 35, we only store

$$\begin{aligned} d_{t'} &: \text{Conf}(X_t \cap X_{t'}) \rightarrow \{\text{true}, \text{false}\} \\ d_{t'}(R) &:= \text{ParSol}(R; X_t \cap X_{t'}; V_{t'}) \end{aligned} \tag{1}$$

We calculate these functions bottom-up. The case that  $t'$  is the leaf is analogous to the corresponding case in Theorem 35, taking time  $O(n^{\text{pw}(H)} \text{poly}(k))$ .

Now let  $t'$  with parent  $t$  and child  $t''$  be an inner node of  $P$ . We define  $\hat{v}$  to be the unique element with  $\hat{v} \in X_{t'} \setminus X_t$  and similarly,  $\hat{w} \in X_{t'} \setminus X_{t''}$  (or, if we would have  $\hat{v} = \hat{w}$ , we take  $\hat{w}$  to be some vertex from  $X_{t'} \setminus \{\hat{v}\}$  instead), and finally  $E := X_{t'} \setminus \{\hat{v}, \hat{w}\}$ . Now if  $\hat{v}\hat{w} \notin E(H)$ , the calculation is easy. Hence assume  $\hat{v}\hat{w} \in E(H)$ . For a configuration  $R$  of  $X_t \cap X_{t'}$ , we get the following alternate characterization of  $d_{t'}(R)$ :

$$\begin{aligned} d_{t'}(R) &= \text{ValConf}(R; \{\hat{w}\} \cup E) \wedge \\ &\exists v' \in f^{-1}(\hat{v}) : \text{ParSol}((R \cup \{\hat{v} \mapsto v'\})|_{X_{t'} \setminus \{\hat{w}\}}; \{\hat{v}\} \cup E; V_{t''}) \wedge v'R(w) \in E(G) \end{aligned}$$

We describe how to calculate  $d_{t'}$  via rectangular matrix multiplication. Much like in Theorem 35, the matrices are indexed by configurations. In contrast to the former, however, one of the two dimensions of the matrix might correspond to the configuration of multiple vertices. Formally, for a vertex subset  $Y \subseteq V(H)$  and a vertex  $x \in V(H), x \notin Y$ , we call a matrix  $A$  **indexed by configurations of  $(x, Y)$**  when it is of dimensions  $n \times n^{|Y|}$ . We use two arbitrary bijections  $g_{\{x\}} : \text{Conf}(\{x\}) \rightarrow [n]$  and  $g_Y : \text{Conf}(Y) \rightarrow [n^{|Y|}]$  which will help us map configurations of  $x$  and  $Y$  to indices of the matrix. Hence, for a configuration  $R$  of  $\{x\} \cup Y$ , we define  $A[R] := A[g_x(R|_{\{x\}}), g_Y(R|_Y)]$ .

For our rectangular matrix product, we define a  $n \times n^{\text{pw}(H)-1}$  matrix  $B_{t'}$  indexed by configurations of  $(\hat{v}, E)$ . For a configuration  $R'$  of  $\{\hat{v}\} \cup E$ , we define

$$B_{t'}[R'] := d_{t''}(R') = \text{ParSol}(R'; \{\hat{v}\} \cup E; V_{t''})$$

Now consider the adjacency matrix  $\text{Adj}_{\hat{v}, \hat{w}}$  of  $f^{-1}(\hat{v})$  and  $f^{-1}(\hat{w})$ , indexed by configurations of  $(\hat{w}, \{\hat{v}\})$ . We make sure that the indexing bijection  $g_{\{\hat{v}\}}$  as defined above is the same for both  $\text{Adj}_{\hat{v}, \hat{w}}$  and  $B_{t'}$  and then calculate the matrix product  $\text{Adj}_{\hat{v}, \hat{w}} \cdot B_{t'}$ . Naturally, the product is indexed by configurations of  $(\hat{w}, E)$  and can be expressed as

$$\begin{aligned} & (\text{Adj}_{\hat{v}, \hat{w}} \cdot B_{t'})[R'] \\ &= \exists v' \in f^{-1}(\hat{v}) : \text{ParSol}((R' \cup \{\hat{v} \mapsto v'\})|_{X_{t'} \setminus \{\hat{w}\}}; \{\hat{v}\} \cup E; V_{t''}) \wedge v'R(w) \in E(G) \end{aligned}$$

Hence, we may write  $d_{t'}(R)$  as

$$d_{t'}(R) = \text{ValConf}(R; \{\hat{w}\} \cup E) \wedge (\text{Adj}_{\hat{v}, \hat{w}} \cdot B_{t'})[R] \quad (1)$$

The algorithm to calculate  $d_{t'}$  is immediate. First, we calculate the rectangular matrix product of  $\text{Adj}_{\hat{v}, \hat{w}}$  and  $B_{t'}$  in time  $O(n^{\omega(\text{pw}(H)-1)})$ , then calculate  $d_{t'}$  via formula 1. Checking whether  $R$  is a valid configuration of  $\{\hat{w}\} \cup E$  can be done in time  $\text{poly}(\text{pw}(H)) \leq \text{poly}(k)$ , giving us a total time of at most  $O(n^{\omega(\text{pw}(H)-1)} \text{poly}(k))$  per inner node.

To output the answer, consider the child  $r'$  of the root  $r$ . We have by definition that  $d_{r'}(R)$  is 1 if and only if  $R$  is a partial solution for  $X_r$  in  $V(H)$ . Thus, the input is a YES-instance for COLORED SUBGRAPH ISOMORPHISM if and only if there is an  $R$  such that  $d_{r'}(R)$  is 1.

Since there is only a single leaf and  $\text{poly}(k)$  inner nodes, total running time of the dynamic programming algorithm is  $O(n^{\omega(\text{pw}(H)-1)} \text{poly}(k))$ . ◀

## 5.6 Exact Weight Colored Subgraph Isomorphism for Bounded Pathwidth

The techniques of using rectangular matrix multiplication for the case of pathwidth can also be applied to the weighted case, and they lead to improvements in the expected way. Again, the instances may be either node- or edge-weighted.

► **Theorem 39.** *There is an algorithm which, given an arbitrary instance  $\phi = (H, G, f, w)$  of the EXACT WEIGHT COLORED SUBGRAPH ISOMORPHISM problem, solves  $\phi$  in time*

1.  $O((n^{\omega(\text{pw}(H)-1)}W + n^{\text{pw}(H)}W \log W) \text{poly}(k) + g(k))$  when  $\text{pw}(H) \geq 2$ , and
2.  $O((n^2W + nW \log W) \text{poly}(k) + g(k))$  when  $\text{pw}(H) = 1$

where  $k := |V(H)|$ ,  $n$  is the size of the preimages of  $f$ ,  $W$  is the maximum absolute weight in the image of  $w$ , and  $g$  is a computable function.

From this, Theorem 8 from the results section follows directly via Lemma 14.

► **Lemma 40.** *Let two matrices  $A, B$  of dimensions  $r \times s$  and  $s \times t$  be given, such that each of their entries  $A_{i,j}, B_{j,k} \in \mathbb{C}[X, X^{-1}]$  (for all  $i \in [r], j \in [s], k \in [t]$ ) is a Laurent polynomial of degree bounded by  $W$  in both the positive and the negative direction. Then their product can be computed in time  $O(MM(r, s, t)W + (rs + st)W \log W)$*

**Proof.** Analogous to Lemma 37. ◀

**Proof (of theorem 39).** Again, we only need to prove an algorithm for the edge-weighted case due to Proposition 29.

**Part 2** is a corollary of Theorem 36, since for any graph, its treewidth is smaller than its pathwidth.

For **part 1**, we only sketch the proof, since it is a straightforward combination of the techniques used in the proofs of Theorems 36 and 38. We use notation and phrasing from both of those proofs without further mention.

Given  $G, H$  and  $f$ , we compute an optimal path decomposition  $(P, \{X_t\}_{t \in V(P)})$  for  $H$  and do dynamic programming on the modified path decomposition as described in the proof of Theorem 38. For each non-root node  $t' \in P$  with parent  $t$ , for each weight  $W' \in \mathcal{W}$  and for each configuration  $R$  of  $X_t \cap X_{t'}$ , we store

$$d_{t', W'} := \text{ParSolE}(R; X_t \cap X_{t'}; V_{t'}; W')$$

The case that  $t'$  is a leaf is clear. For the case that  $t'$  is an inner node with parent  $t$  and child  $t''$ , let  $\hat{v}, \hat{w}$  be the unique elements with  $\hat{v} \in X_{t'} \setminus X_t, \hat{w} \in X_{t'} \setminus X_{t''}$  (or, if we would have  $\hat{v} = \hat{w}$ , we take  $\hat{w}$  to be some vertex from  $X_{t'} \setminus \{\hat{v}\}$  instead) and  $E = X_{t'} \setminus \{\hat{v}, \hat{w}\}$ . If  $\hat{v}\hat{w} \notin E(H)$ , the calculation is easy, hence assume  $\hat{v}\hat{w} \in E(H)$ . For each weight  $W'$ , we build a rectangular matrix  $A_{t'}^{W'}$  indexed by  $(\hat{v}, E)$ . The entry corresponding to a configuration  $R$  of  $\{\hat{v}\} \cup E$  tells us whether  $R$  has an extension of weight  $W' + x$ , where  $x$  is the weight contributed by  $\hat{v}$  in  $\{\hat{v}\} \cup E$ . In particular  $x := \sum_{\hat{u} \in E} w(R|_{\{\hat{u}, \hat{v}\}})$ .

We also use, for each weight  $W'$ , an adjacency matrix  $\text{Adj}_{\hat{v}, \hat{w}}^{W'}$  defined for node-weighted instances as

$$\text{Adj}_{\hat{v}, \hat{w}}^{W'}[v', w'] = v'w' \in E(G) \wedge w(v'w') = W'$$

We then create two Laurent polynomials of degree  $|\mathcal{W}|$ , one with the matrices  $\text{Adj}_{\hat{v}, \hat{w}}^{W'}$  as coefficients, one with the matrices  $A_{t'}^{W'}$ . These can also be seen as matrices with Laurent polynomials as entries. Using Lemma 40, we then calculate their matrix product, which tells us for each weight  $W'$  and each configuration  $R$  of  $\{\hat{w}\} \cup E$  if  $R$  is a potential solution of  $\{\hat{w}\} \cup E$  in  $V_{t'}$  with an extension of weight  $W'$ , potentially missing edges between the preimage of  $w$  and the preimages of  $E$ . The latter can be checked for each entry. This leads to the desired running time. ◀

## 5.7 Improvements for the Node-Weighted Case

In the case of node weights instead of edge weights, some of the algorithms can be slightly improved using rectangular matrix multiplication. However, these improvements only work for the case that the treewidth of  $H$  is 1 and for the case of bounded pathwidth.

Specifically, we show the following two results, which imply Theorems 9 and 10 from the results section.

► **Theorem 41.** *There is an algorithm which, given an arbitrary instance  $\phi = (H, G, f, w)$  of the node-weighted EXACT WEIGHT COLORED SUBGRAPH ISOMORPHISM problem where  $H$  is a tree, solves  $\phi$  in time  $O(\text{MM}(n, n, W) \text{poly}(k) + nW \log W \text{poly}(k))$ .*

► **Theorem 42.** *There is an algorithm which, given an arbitrary instance  $\phi = (H, G, f, w)$  of the node-weighted EXACT WEIGHT COLORED SUBGRAPH ISOMORPHISM problem, solves  $\phi$  in time  $O(\text{MM}(n, n, n^{\text{pw}(H)-1}W) \text{poly}(k) + g(k))$ .*

**Proof (of Theorem 41).** Let  $G, f : V(G) \rightarrow V(H)$  and  $w : V(G) \rightarrow \mathbb{Z}$  be given. All achievable total weights must lie in  $\mathcal{W} := \{-kW, \dots, kW\}$ .

We do dynamic programming on the tree  $H$ . Rooting  $H$  in an arbitrary vertex, we define  $T_u$  to be the subtree rooted at a node  $u \in H$ . For each weight  $W \in \mathcal{W}$ , each node  $v \in H$ , we store the following function of finite domain:

$$\begin{aligned} d_{v,W} &: \text{Conf}(\{v\}) \rightarrow \{\text{true}, \text{false}\} \\ d_{v,W}(R) &:= \text{ParSolE}(R; \{v\}; T_v; W) \end{aligned}$$

We calculate the entries of these functions bottom-up, starting at the leaves of  $H$ . If  $v$  is a leaf,  $d_{v,W}(R)$  is 1 if and only if  $W = 0$ . Now suppose  $v$  is a non-leaf node of  $H$  with set of children  $C_v$ . For each child  $u \in C_v$ , we construct a rectangular matrix  $p_v^u$ . This matrix is indexed by  $w$  and  $\mathcal{W}$ . Formally, we call a matrix  $A$  **indexed by  $w$  and  $\mathcal{W}$**  if it has dimensions  $n \times |\mathcal{W}|$ . We use two arbitrary bijections  $g_{\{v\}} : \text{Conf}(\{v\}) \rightarrow [n]$  and  $g_w : \mathcal{W} \rightarrow [|\mathcal{W}|]$  to help us map weights from  $\mathcal{W}$  and configurations of  $v$  to indices of  $A$ . Correspondingly, we define  $A[R, W] := A[g_v(R), g_w(W)]$ .

We define  $p_v^u$  as

$$p_v^u[R, W] := \text{ParSolE}(R; \{v\}; \{v\} \cup T_u; W)$$

We may calculate this as follows. Let  $\text{Adj}_{\{v,u\}}$  be the adjacency matrix of  $f^{-1}(v)$  and  $f^{-1}(u)$  indexed by  $(v, \{u\})$  (as defined in the proof of Theorem 38), and let  $\tilde{d}_u$  be the rectangular matrix indexed by  $u$  and  $\mathcal{W}$  and defined as  $\tilde{d}_u[R, W] := d_{u, W-w(R)}(R)$ . We make sure that the indexing bijection  $g_{\{u\}}$  as defined above is the same for both  $\text{Adj}_{v,u}$  and  $\tilde{d}_u$  and then calculate the matrix product  $\text{Adj}_{v,u} \cdot \tilde{d}_u$ . The product is indexed by configurations of  $v$  and  $\mathcal{W}$  and can be expressed as

$$\begin{aligned} (\text{Adj}_{v,u} \cdot \tilde{d}_u)[R, W] &= \exists u' \in f^{-1}(u) : R(v)u' \in E(G) \wedge \text{ParSolE}(R; \{u\}; T_u; W - w(u')) \\ &= \text{ParSolE}(R; \{v\}; \{v\} \cup T_u; W) \\ &= p_v^u[R, W] \end{aligned}$$

Writing  $C_v = \{u_1, \dots, u_c\}$  with  $c = |C_v|$ , the function  $d_{v,W}$  may then be expressed as

$$d_{v,W}(R) = \exists W_1, \dots, W_c : \sum_{i=1}^c W_i = W \wedge \forall i : p_v^{u_i}[R, W_i]$$

Similarly to the proof of Theorem 36, this may be calculated using a boolean convolution. Accordingly, we define for each configuration  $R$  of  $v$  the finitely supported functions  $f_{v,R}^1, \dots, f_{v,R}^c : \mathbb{Z} \rightarrow \{\text{true}, \text{false}\}$  as  $f_{v,R}^i(x) := p_v^{u_i}(R)$  if  $x \in \mathcal{W}$ , and 0 otherwise. By a simple calculation, we get  $d_{v,W}(R) = (f_{v,R}^1 * \dots * f_{v,R}^c)(W)$ .

Finally, after having calculated all values of  $d_{t',W}(R)$  for all  $t', W$  and  $R$ , we wish to output the result. Let  $r$  be the root of  $H$ . By definition of  $d_{r,W}$ , there is some configuration  $R$  of  $r$  such that  $d_{r,-w(R)}(R) = \text{true}$  if and only if the instance has a solution.



It remains to analyze the running time. For the leaves of  $H$ , the calculation takes time  $O(nW)$ . For inner nodes, the calculation of the matrix product  $Adj_{\{v,u\}} \cdot \tilde{d}_u$  takes time  $MM(n, n, W)$ . Finally, calculating the discrete convolutions takes time  $O(nW \log W)$ , since any vertex of  $v$  is involved as a child in at most one discrete convolution. Hence, we arrive at the running time from the theorem.  $\blacktriangleleft$

**Proof (of Theorem 42).** The proof uses a combination of the techniques from the proofs of Theorem 41 and Theorem 38.

Let  $G, H, f : V(G) \rightarrow V(H)$  and  $w : V(G) \rightarrow \mathbb{Z}$  be given. We compute an optimal path decomposition in time  $g(k)$ , modify it as described in Theorem 38, obtaining a modified path decomposition  $\mathcal{P} = (P, \{X_t\}_{t \in P})$ , and then do dynamic programming on  $\mathcal{P}$ . Note that all achievable weights must lie in  $\mathcal{W} := \{-kW, \dots, kW\}$ .

We store, for each non-root node  $t'$  with parent  $t$  of  $P$ , the following function of finite domain:

$$d_{t',W}(R) := \text{ParSolE}(R; X_t \cap X_{t'}; V_{t'}; W)$$

For  $t'$  a leaf, the calculations is clear. Let  $t'$  be an inner node with parent  $t$  and child  $t''$  and define  $\hat{v}$  to be the unique element with  $\hat{v} \in X_{t'} \setminus X_t$  and similarly  $\hat{w} \in X_{t'} \setminus X_{t''}$  (or, if we would have  $\hat{v} = \hat{w}$ , we take  $\hat{w}$  to be some arbitrary vertex from  $X_{t'} \setminus \{v\}$  instead) and finally  $E := X_{t'} \setminus \{\hat{v}, \hat{w}\}$ . If  $\hat{v}\hat{w} \notin E(H)$ , the calculation is easy, hence assume  $\hat{v}\hat{w} \in E(H)$ . For a configuration  $R$  of  $X_t \cap X_{t'}$ , we get the following alternate characterization of  $d_{t',W}(R)$ :

$$d_{t',W}(R) = \text{ValConf}(R; \{\hat{w}\} \cup E) \wedge \exists v' \in f^{-1}(\hat{v}) : v'R(\hat{w}) \in E(G) \wedge \\ \text{ParSolE}((R \cup \{\hat{v} \mapsto v'\})|_{X_{t'} \setminus \{\hat{w}\}}; \{\hat{v}\} \cup E; V_{t''}; W - w(v'))$$

We now set up our rectangular matrix product. For a vertex  $x \in V(H), x \notin E$ , we call a matrix  $A$  indexed by  $x, E$  and  $\mathcal{W}$  if it has dimensions  $n \times n^{|\mathcal{W}|}$ . We use two arbitrary bijections  $g_{\{x\}} : \text{Conf}(\{x\}) \rightarrow [n]$  and  $g_{E,\mathcal{W}} : \text{Conf}(E) \times \mathcal{W} \rightarrow [n^{|\mathcal{W}|}]$  to help index the matrix and define, for a configuration  $R$  of  $\{x\} \cup E$ ,  $A[R, W] := A[g_x(R|x), g_{E,\mathcal{W}}(R|E, W)]$ .

We define the  $n \times n^{pw(H)-1|\mathcal{W}|}$  matrix  $B_{t'}$ , which is to be indexed by  $\hat{v}, E$  and  $\mathcal{W}$ , as follows:

$$B_{t'}[R', W'] := d_{t',W'-w(R|\hat{v})}(R')$$

As in the proof of Theorem 38, we also use the adjacency matrix  $Adj_{\hat{v},\hat{w}}$  indexed by configurations of  $(\hat{w}, \{\hat{v}\})$ . Ensuring that the indexing bijections  $g_{\{\hat{v}\}}$  are the same for both  $B_{t'}$  and  $Adj_{\hat{v},\hat{w}}$ , we calculate  $Adj_{\hat{v},\hat{w}} \cdot B_{t'}$  and obtain a matrix indexed by  $\hat{w}, E$  and  $\mathcal{W}$ . Its entries can be expressed as

$$(Adj_{\hat{v},\hat{w}} \cdot B_{t'})[R', W'] = \exists v' \in f^{-1}(\hat{v}) : v'R'(w) \in E(G) \wedge \\ \text{ParSolE}((R' \cup \{\hat{v} \mapsto v'\})|_{X_{t'} \setminus \{\hat{w}\}}; \{\hat{v}\} \cup E; V_{t''}; W' - w(v'))$$

Thus,  $d_{t'}(R)$  can be expressed as

$$d_{t',W}(R) = \text{ValConf}(R; \{\hat{w}\} \cup E) \wedge (Adj_{\hat{v},\hat{w}} \cdot B_{t'})[R, W]$$

This concludes the description of the computation of  $d_{t',W}$ . For the computation of the answer, consider the child  $r'$  of the root  $r$ . By definition of  $d_{r',W}(R)$ , we have that there exists a configuration  $R$  such that  $d_{t',-w(R|x_r)}(R) = \text{true}$  if and only if the instance has a solution.

By a simple argument, this dynamic programming algorithm has a running time of  $O(MM(n, n, n^{pw(H)-1}W) \text{poly}(k))$ .  $\blacktriangleleft$

## 6 Interconnections Between Subgraph Isomorphism, Boolean $k$ -Wise Matrix Products and Hyperclique

We have seen in the proof of the conditional lower bound for the unweighted COLORED SUBGRAPH ISOMORPHISM problem that for any  $t \in \mathbb{N}$  there is a polynomial-time reduction from  $h$ -uniform  $h(t+1)$ -HYPERCLIQUE with  $n$  nodes to COLORED SUBGRAPH ISOMORPHISM on pattern graphs of treewidth  $t$  with  $n^h$  nodes. Hence we have the following corollary, which follows directly from the Unweighted Lemma.

► **Corollary 43.** *If there is a  $t \geq 3$  such that the COLORED SUBGRAPH ISOMORPHISM problem for pattern graphs of treewidth  $t$  can be solved in time  $O(n^{t+1-\varepsilon})$  (for some  $\varepsilon > 0$ ), then for any  $3 \leq h \leq t$  the  $h$ -uniform  $h(t+1)$ -HYPERCLIQUE problem can be solved in time  $O(n^{h(t+1)-\varepsilon'})$  (for some  $\varepsilon' > 0$ ).*

But what about the other direction? Can we also provide a conditional lower bound for HYPERCLIQUE under the hypothesis that COLORED SUBGRAPH ISOMORPHISM cannot be solved faster, hence proving an equivalence? Indeed we can! However, the “equivalence” we get is not as strong as one might hope.

Specifically, it turns out that the algorithm we described for COLORED SUBGRAPH ISOMORPHISM already gives a Turing reduction from COLORED SUBGRAPH ISOMORPHISM with treewidth  $t$  to the BOOLEAN  $t$ -WISE MATRIX PRODUCT problem<sup>9</sup>. Indeed, it can also be seen that the BOOLEAN  $t$ -WISE MATRIX PRODUCT problem is equivalent to the  $t$ -uniform  $(t+1)$ -HYPERCLIQUE problem. More formally, we have the following two lemmas.

► **Lemma 44.** *For any  $t \geq 3$ , if the BOOLEAN  $t$ -WISE MATRIX PRODUCT problem can be solved in time  $O(n^{t+1-\varepsilon})$ , then COLORED SUBGRAPH ISOMORPHISM problem on pattern graphs of treewidth  $t$  can be solved in time  $O(n^{t+1-\varepsilon})$ .*

**Proof.** As was already mentioned in the algorithms for COLORED SUBGRAPH ISOMORPHISM (see the proof of Theorem 35), the boolean  $\text{tw}(H)$ -wise matrix product is the bottleneck for the running time. All other operations run in time  $O(n^{\text{tw}(H)} \text{poly}(|V(H)|))$ . Hence a  $O(n^{t+1-\varepsilon})$  algorithm for BOOLEAN  $t$ -WISE MATRIX PRODUCT translates directly to a  $O(n^{t+1-\varepsilon})$  algorithm for COLORED SUBGRAPH ISOMORPHISM with  $\text{tw}(H) = t$ . ◀

► **Lemma 45.** *If the  $t$ -uniform  $(t+1)$ -HYPERCLIQUE problem can be solved in time  $O(n^{t+1-\varepsilon})$  (for some  $\varepsilon > 0$ ), then the BOOLEAN  $t$ -WISE MATRIX PRODUCT problem can be solved in time  $O(n^{t+1-\varepsilon'})$  (for some  $\varepsilon' > 0$ ).*

We defer the proof of Lemma 45 to Appendix B.2. Indeed, in appendix B, we also show the other direction, i.e. we show that the existence of fast algorithms for these two problems is actually equivalent. This is a natural generalization of methods from [68], where this result is proven for combinatorial algorithms for the case  $t = 2$ .

Composing these two lemmas, we get the following theorem.

► **Theorem 46.** *If the  $t$ -uniform  $(t+1)$ -HYPERCLIQUE problem can be solved in time  $O(n^{t+1-\varepsilon})$  (for some  $\varepsilon > 0$ ), then the COLORED SUBGRAPH ISOMORPHISM problem on pattern graphs of treewidth  $t$  can be solved in time  $O(n^{t+1-\varepsilon'})$  (for some  $\varepsilon' > 0$ ).*

<sup>9</sup> Defined as: Given  $t$  tensors  $A^1, \dots, A^t$  of order  $t$  with dimensions  $n \times \dots \times n$ , calculate their boolean  $t$ -wise matrix product  $\text{MP}_t(A^1, \dots, A^t)$ .

Looking at Corollary 43 and Theorem 46, we have reductions in both directions, but they do not give a full equivalence. This is because the reduction from Corollary 43 only gives an algorithm for  $h$ -uniform  $h(t+1)$ -HYPERCLIQUE (for any  $h \geq 3$ ), but Theorem 46 needs an algorithm for the much “denser”  $t$ -uniform  $(t+1)$ -HYPERCLIQUE problem.

## 7 The Colored Problems are Equivalent to the Uncolored Problems

We now show Lemma 14 from the preliminaries (restated in a modified form below), which shows that the (EXACT WEIGHT) SUBGRAPH ISOMORPHISM and (EXACT WEIGHT) COLORED SUBGRAPH ISOMORPHISM problems are essentially equivalent with respect to running times. Hence, for most of our purposes, we can treat them as equal, which simplifies both the proofs of the algorithms and the lower bounds since the colored version is much more structured.

The reductions from (EXACT WEIGHT) SUBGRAPH ISOMORPHISM to (EXACT WEIGHT) COLORED SUBGRAPH ISOMORPHISM and the reduction from EXACT WEIGHT COLORED SUBGRAPH ISOMORPHISM to EXACT WEIGHT SUBGRAPH ISOMORPHISM leaves  $H$  unmodified. The reduction from COLORED SUBGRAPH ISOMORPHISM to SUBGRAPH ISOMORPHISM, however, modifies  $H$  in such a way that preserves treewidth, but may modify pathwidth.

We say that SUBGRAPH ISOMORPHISM or COLORED SUBGRAPH ISOMORPHISM have a  $T(n, k, \rho(H))$  algorithm (for some graph parameter  $\rho$ ) if there is an algorithm  $\mathcal{A}$  which decides a given instance  $\phi = (H, G, f)$  of either problem in time  $T(n, k, \rho(H))$ . Analogously, we define the phrase that EXACT WEIGHT SUBGRAPH ISOMORPHISM or EXACT WEIGHT COLORED SUBGRAPH ISOMORPHISM has a  $T(n, k, \rho(H), W)$  algorithm, the only difference being that  $\phi = (H, G, f, w)$ .  $W$  denotes the maximum absolute value of the weight function  $w$ .

Parts 1 and 2 of the lemma follows directly from the Color Coding technique [12].

► **Lemma 47** (reformulation of Lemma 14). *Let  $\rho$  be any graph parameter.*

1. *If there is a  $T(n, k, \rho(H))$  time deterministic algorithm for COLORED SUBGRAPH ISOMORPHISM, then there is a  $O(T(kn, k, \rho(H))g(k))$  expected time algorithm and furthermore a  $\tilde{O}(T(kn, k, \rho(H))g(k))$  time deterministic algorithm for SUBGRAPH ISOMORPHISM, for some computable function  $g$ .*
2. *If there is a  $T(n, k, \rho(H), W)$  time deterministic algorithm for EXACT WEIGHT COLORED SUBGRAPH ISOMORPHISM, then there is a  $O(T(kn, k, \rho(H), W)g(k))$  expected time algorithm and furthermore a  $\tilde{O}(T(kn, k, \rho(H), W)g(k))$  time deterministic algorithm for EXACT WEIGHT SUBGRAPH ISOMORPHISM, for some computable function  $g$ .*
3. *Let  $\text{tw}(H) \geq 2$ . If there is a  $T(n, k, \text{tw}(H))$  time algorithm for SUBGRAPH ISOMORPHISM, then there is a  $O(T(\text{poly}(k)n, \text{poly}(k), \text{tw}(H)) + \text{poly}(k)n^2)$  time algorithm for COLORED SUBGRAPH ISOMORPHISM.*
4. *If there is a  $T(n, k, \rho(H), W)$  time algorithm for EXACT WEIGHT SUBGRAPH ISOMORPHISM, then there is a  $O(T(2n, 2k, \rho(H), 2^k W) + \text{poly}(k)n^2)$  time algorithm for EXACT WEIGHT COLORED SUBGRAPH ISOMORPHISM.*

**Proof (of Lemma 14).** We start by showing **part 1** and **part 2**. As mentioned, this follows directly from a standard application of the Color Coding technique [12]. Briefly speaking, they use random colorings of the vertices of  $G$  to make the potential solution subgraph multicolored with some probability depending only on  $k$ . In our case, we may then try all  $k!$  mappings from colors to vertices of  $H$  to obtain the randomized algorithm; we delete any monochromatic edges to make sure that  $f$  is a homomorphism. The authors of [12] also

explain how to derandomize the algorithm using  $k$ -perfect hash functions, which results in the deterministic algorithm with an additional factor of  $2^{O(k)} \log(n)$ .

The factor of  $k$  in front of  $n$  in  $T(kn, k, \text{tw}(H))$  comes from the fact that in the COLORED SUBGRAPH ISOMORPHISM problem, we consider  $n$  to be the size of the preimages of  $f$ , while in the SUBGRAPH ISOMORPHISM problem, we consider it to be the size of  $V(G)$ .

Now we show **part 3**. Given an instance  $\phi$  of COLORED SUBGRAPH ISOMORPHISM, where  $H$  has  $k$  vertices, with preimages in  $G$  of  $n$  vertices each, we construct an equivalent SUBGRAPH ISOMORPHISM instance  $\phi'$ . This is done in two steps. First, we modify  $\phi$  into an equivalent, but more structured COLORED SUBGRAPH ISOMORPHISM instance  $\tilde{\phi}$ , which we then reduce to  $\phi'$ . See Figure 4 for an example of this reduction.


**Subdividing all edges:** In the first step, we construct  $\tilde{H}$ , which consists of a subdivided copy of  $H$ , where each vertex has a unique “signature” structure attached to it. These signatures have a triangle as a key component. Abusing notation, we write  $V(H) = \{1, \dots, k\}$  and use the vertices as numbers. First, for each  $i \in V(H)$ , we add a vertex  $\tilde{i}$  to  $\tilde{H}$ . Then, for each edge  $ij \in E(H)$  with  $i < j$ , we add a vertex  $\tilde{x}_{ij}$  and create two edges  $\tilde{i}\tilde{x}_{ij}$  and  $\tilde{x}_{ij}\tilde{j}$ , hence subdividing the edge  $ij$ . This ensures that for now, the new graph has no triangles. In  $\tilde{G}$ , we populate the preimages as follows: For each  $i \in V(H)$ , let  $f^{-1}(i) = \{a_i^1, \dots, a_i^n\}$  and add  $n$  vertices  $\{\tilde{a}_i^1, \dots, \tilde{a}_i^n\}$  to  $f^{-1}(\tilde{i})$ . For each  $\ell$ , the vertex  $\tilde{a}_i^\ell$  corresponds to  $a_i^\ell$ . The preimages of  $\tilde{x}_{ij}$  are populated with  $n$  vertices  $\{b_{ij}^1, \dots, b_{ij}^n\}$  via  $\tilde{f}$ . For each edge  $ij \in E(H)$  with  $i < j$ , we add an edge  $a_i^\ell b_{ij}^\ell$  for every  $\ell \in [n]$ . We also go through each edge  $a_i^\ell a_j^m \in E(G)$  and add a corresponding edge  $b_{ij}^\ell \tilde{a}_j^m$  to  $E(\tilde{G})$ .

**Signatures:** We now add the signatures. For each  $i \in [k]$ , we add a new vertex  $\tilde{t}_i$  and a new triangle  $\tilde{u}_i \tilde{v}_i \tilde{w}_i$  to  $\tilde{H}$ , and connect  $\tilde{t}_i$  to both  $\tilde{u}_i$  and  $\tilde{v}_i$  from  $V(\tilde{H})$ . Furthermore, we connect  $\tilde{v}_i$  to  $i+1$  other newly created vertices  $\tilde{y}_i^1, \dots, \tilde{y}_i^{i+1}$ . Let the set of all newly created vertices  $\tilde{t}_i, \tilde{u}_i, \tilde{v}_i, \tilde{w}_i, \tilde{y}_i^\ell$  ( $i \in [k], \ell \in [i+1]$ ) be named  $X$ . In  $\tilde{G}$ , we populate the preimages of these new vertices by adding  $n$  vertices  $\{z_1, \dots, z_n\}$  to  $V(\tilde{G})$  for each vertex  $v \in X$ , with  $\forall i \in [n] : \tilde{f}(z_i) = v$ . Now, for each  $u \in X$ , we pick an arbitrary node from  $f^{-1}(u)$  and call it active. Furthermore, for all  $\tilde{i} \in V(\tilde{H})$ , we call all vertices of  $f^{-1}(\tilde{i})$  active. Now for each  $v \in X$ , we connect its active vertex in  $f^{-1}(v)$  to all active vertices from the neighbourhood  $f^{-1}(N(v))$ . Note that of the vertices in  $f^{-1}(X)$ , only the active ones have edges at all. Indeed, for each  $i \in [k]$ ,  $\tilde{G}$  contains exactly one triangle such that one of its vertices has degree  $i+3$ .

We set  $\tilde{\phi}$  to be the instance  $(\tilde{H}, \tilde{G}, \tilde{f})$ . This concludes the construction of  $\tilde{\phi}$ .

Clearly, the new instance  $\tilde{\phi}$  has a solution if and only if  $\phi$  has one, and the size of the new instance is only a poly( $k$ ) factor larger than the size of  $\phi$ . We must also show that the reduction preserves treewidth. Note that  $\tilde{H}$  is obtained from  $H$  via two operations: Subdividing edges and connecting a graph of smaller or equal treewidth via a single edge. It is easy to see that both operations do not change the treewidth.

Now, to construct  $\phi'$ , we simply get rid of the mapping  $\tilde{f}$ . In other words,  $\phi' = (\tilde{H}, \tilde{G})$ . Obviously, if  $\tilde{\phi}$  has a solution, then  $\phi'$  has one. For the other direction, suppose  $\phi'$  has a solution, i.e. a subgraph  $S'$  of  $\tilde{G}$  along with an isomorphism  $h : V(\tilde{H}) \rightarrow V(G[S'])$  of  $\tilde{G}$ . Since both  $\tilde{H}$  and  $\tilde{G}$  have, for each  $i$ , exactly one triangle with a vertex of degree  $i$ ,  $h$  must map these triangles to their respective counterparts in  $\tilde{G}$ . In particular, each of the vertices in  $X$  is mapped to the active vertex in its preimage. Since the active node in  $f^{-1}(\tilde{t}_i)$  is connected only to nodes of  $f^{-1}(\tilde{i})$  (apart from the active node of  $f^{-1}(\tilde{u}_i)$ , which is already in the image of  $h(\tilde{u}_i)$ ), we know that  $h(\tilde{i}) \in f^{-1}(\tilde{i})$ . Analogously,  $h(\tilde{x}_{ij}) \in f^{-1}(x_{ij})$ . Hence,  $S'$  takes exactly one vertex from each preimage of vertices from  $\tilde{H}$ . Thus,  $\tilde{\phi}$  also has a solution.



■ **Figure 4** An example of the construction of the instance  $\tilde{\phi}$  from  $\phi = (H, G, f)$  where  $H$  is a triangle

We thus obtain a way to reduce  $\phi$  to  $\phi'$  with a size factor of only  $\text{poly}(k)$ . The reduction obviously runs in  $O(\text{poly}(k)n^2)$  time. This shows part 3.

Finally, we show **part 4**. We begin with the node-weighted version. Given an instance  $\phi$  of EXACT WEIGHT COLORED SUBGRAPH ISOMORPHISM where the pattern graph  $H$  has  $k$  vertices, we create an instance  $\phi'$  of EXACT WEIGHT SUBGRAPH ISOMORPHISM by simply dropping  $f$  and modifying the weights. We have to ensure that a solution of  $\phi'$  takes exactly one node from each preimage of  $H$ . To do this, we encode a checklist in the weights of the nodes. Again, let  $V(H) = \{1, \dots, k\}$ . Let  $u \in f^{-1}(i)$  for  $i \in V(H)$ , and consider its weight  $w(u)$ . We modify it by multiplying it with  $2^k$  and adding  $2^i$ . We call the added weight its “signature”. Now, since any solution must pick exactly  $k$  vertices, the only way that the signatures of the solution vertices sum up to  $2^k - 1$  is to pick vertices which have a sum of weight 0 according to the original weight function and furthermore have exactly one vertex with added weight  $2^i$  for each  $i = 1, \dots, k$ . To complete our reduction, we pick an arbitrary vertex  $v \in V(G)$  and subtract  $2^k - 1$  from all vertices in  $f^{-1}(v)$ , making the new target zero. This reduction does not alter  $G$  or  $H$ , and instead only modifies the weight function, resulting in the stated time bounds.

For the edge-weighted version, we can use essentially the same construction as for the node-weighted version. Again, all weights are multiplied by  $2^k$ , and each vertex of  $H$  has a unique “signature”. However, this time, we have to add the signatures to the edge weights. Consequently, for each  $i \in V(H) = \{1, \dots, k\}$ , we pick an arbitrary incident edge  $e \in E(H)$ . Let  $e = \{i, j\}$ . We add the signature  $2^i$  to every edge in the preimage of  $e$ . That is, to every edge  $\{e' \in E(G) \mid e' = \{u, v\} \text{ and } f(u) = i \text{ and } f(v) = j\}$ . This way, we must still pick a node from each preimage to ensure that the signatures sum to  $T := 2^k - 1$ . Again, we pick an arbitrary edge  $e \in E(H)$  and subtract  $T$  from all edges in its preimage. This almost completes the proof. However, we still have to handle nodes of degree 0 in  $H$ , since we cannot pick an incident edge for them. However, an isolated vertex  $i$  in  $H$  may be mapped to any vertex in  $G$ . Hence, we may simply skip the signature of  $i$ . We also have to modify the target  $T$  to be  $T - 2^i$ . ◀

## 8 Open Problems

In this paper we discussed many different variants of the SUBGRAPH ISOMORPHISM problem. For some of these variants we leave gaps, which gives rise to several open problems:

1. Can the algorithms for weighted trees be improved? We have shown that some improvements can be made for node-weighted trees (see Theorem 9), but are these optimal? What about edge-weighted trees?
2. Are there fast algorithms for unweighted SUBGRAPH ISOMORPHISM on graphs of bounded pathwidth that do not use rectangular matrix multiplication? Can the gap between exponent  $\omega(p - 1)$  and exponent  $p$  be closed? Similar questions apply to the weighted case; see Theorems 6 and 8.
3. Relatedly, are there good lower bounds for unweighted SUBGRAPH ISOMORPHISM on graphs of bounded pathwidth? These could be attained via a modification of the proof of part 2 of Lemma 14 for pathwidth (though we see no way to do this), or with completely new techniques.

We conclude with some more general open problems:

1. Do our algorithms and lower bounds also work for other types of graph homomorphisms, and for counting the number of solutions? It seems like techniques from [32] should apply.

2. In this work we demonstrated the existence of maximally hard patterns for which Subgraph Isomorphism requires time  $n^{\text{tw}(H)+1-o(1)}$ . Can we classify which (classes of) patterns are maximally hard?
3. Changing our focus from hard patterns to easy patterns, we can ask: do classes of patterns of unbounded treewidth exist for which Subgraph Isomorphism can be solved in time  $n^{o(\text{tw}(H))}$ ? Recall that a conditional lower bound rules out  $n^{o(\text{tw}(H)/\log \text{tw}(H))}$  [54].

---

## References

- 1 Amir Abboud and Greg Bodwin. The  $4/3$  additive spanner exponent is tight. *Journal of the ACM (JACM)*, 64(4):1–20, 2017.
- 2 Amir Abboud, Greg Bodwin, and Seth Pettie. A hierarchy of lower bounds for sublinear additive spanners. *SIAM Journal on Computing*, 47(6):2203–2236, 2018.
- 3 Amir Abboud, Karl Bringmann, Holger Dell, and Jesper Nederlof. More consequences of falsifying SETH and the Orthogonal Vectors conjecture. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 253–266, 2018.
- 4 Amir Abboud, Karl Bringmann, Danny Hermelin, and Dvir Shabtay. SETH-based lower bounds for Subset Sum and Bicriteria Path. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 41–57. SIAM, 2019.
- 5 Amir Abboud, Karl Bringmann, Danny Hermelin, and Dvir Shabtay. Scheduling lower bounds via AND Subset Sum. *arXiv preprint arXiv:2003.07113*, 2020.
- 6 Amir Abboud and Kevin Lewi. Exact weight subgraphs and the k-SUM conjecture. In *International Colloquium on Automata, Languages, and Programming*, pages 1–12. Springer, 2013.
- 7 Amir Abboud, Kevin Lewi, and Ryan Williams. Losing weight by gaining edges. In *European Symposium on Algorithms*, pages 1–12. Springer, 2014.
- 8 Amir Abboud and Virginia Vassilevska Williams. Popular conjectures imply strong lower bounds for dynamic problems. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 434–443. IEEE, 2014.
- 9 Amir Abboud, Virginia Vassilevska Williams, and Oren Weimann. Consequences of faster alignment of sequences. In *International Colloquium on Automata, Languages, and Programming*, pages 39–51. Springer, 2014.
- 10 Amir Abboud, Virginia Vassilevska Williams, and Huacheng Yu. Matching triangles and basing hardness on an extremely popular conjecture. *SIAM Journal on Computing*, 47(3):1098–1122, 2018.
- 11 Noga Alon. Testing subgraphs in large graphs. *Random Structures & Algorithms*, 21(3-4):359–370, 2002.
- 12 Noga Alon, Raphael Yuster, and Uri Zwick. Color-coding. *Journal of the ACM (JACM)*, 42(4):844–856, 1995.
- 13 Omid Amini, Fedor V Fomin, and Saket Saurabh. Counting subgraphs via homomorphisms. In *International Colloquium on Automata, Languages, and Programming*, pages 71–82. Springer, 2009.
- 14 Zhao An, Qilong Feng, Iyad Kanj, and Ge Xia. The complexity of tree partitioning. In *Workshop on Algorithms and Data Structures*, pages 37–48. Springer, 2017.
- 15 MohammadHossein Bateni, MohammadTaghi Hajiaghayi, Saeed Seddighin, and Cliff Stein. Fast algorithms for knapsack via convolution and prediction. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1269–1282, 2018.
- 16 Felix A Behrend. On sets of integers which contain no three terms in arithmetical progression. *Proceedings of the National Academy of Sciences of the United States of America*, 32(12):331, 1946.
- 17 Markus Bläser. Fast matrix multiplication. *Theory of Computing*, pages 1–60, 2013.
- 18 Hans L Bodlaender. A tourist guide through treewidth. *Acta cybernetica*, 11(1-2):1, 1994.

- 19 Hans L Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM Journal on computing*, 25(6):1305–1317, 1996.
- 20 Hans L Bodlaender. A partial k-arboretum of graphs with bounded treewidth. *Theoretical computer science*, 209(1-2):1–45, 1998.
- 21 Hans L Bodlaender. Discovering treewidth. In *International Conference on Current Trends in Theory and Practice of Computer Science*, pages 1–16. Springer, 2005.
- 22 David Bremner, Timothy M Chan, Erik D Demaine, Jeff Erickson, Ferran Hurtado, John Iacono, Stefan Langerman, and Perouz Taslakian. Necklaces, convolutions, and  $x+y$ . In *European Symposium on Algorithms*, pages 160–171. Springer, 2006.
- 23 Karl Bringmann. Why walking the dog takes time: Frechet distance has no strongly subquadratic algorithms unless SETH fails. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 661–670. IEEE, 2014.
- 24 Karl Bringmann. A near-linear pseudopolynomial time algorithm for Subset Sum. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1073–1084. SIAM, 2017.
- 25 Karl Bringmann and Marvin Künnemann. Quadratic conditional lower bounds for string problems and dynamic time warping. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 79–97. IEEE, 2015.
- 26 Danilo Bruschi, Lorenzo Martignoni, and Mattia Monga. Detecting self-mutating malware using control-flow graph matching. In *International conference on detection of intrusions and malware, and vulnerability assessment*, pages 129–143. Springer, 2006.
- 27 Timothy M Chan and Moshe Lewenstein. Clustered integer 3SUM via additive combinatorics. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 31–40, 2015.
- 28 James W Cooley and John W Tukey. An algorithm for the machine calculation of complex Fourier series. *Mathematics of computation*, 19(90):297–301, 1965.
- 29 Don Coppersmith. Rapid multiplication of rectangular matrices. *SIAM Journal on Computing*, 11(3):467–471, 1982.
- 30 Don Coppersmith and Shmuel Winograd. Matrix multiplication via arithmetic progressions. In *Proceedings of the nineteenth annual ACM symposium on Theory of computing*, pages 1–6, 1987.
- 31 Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, and Clifford Stein. *Introduction to algorithms*. MIT Press, 2009.
- 32 Radu Curticapean, Holger Dell, and Dániel Marx. Homomorphisms are a good basis for counting small subgraphs. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 210–223, 2017.
- 33 Marek Cygan, Holger Dell, Daniel Lokshtanov, Dániel Marx, Jesper Nederlof, Yoshio Okamoto, Ramamohan Paturi, Saket Saurabh, and Magnus Wahlström. On problems as hard as CNF-SAT. *ACM Transactions on Algorithms (TALG)*, 12(3):1–24, 2016.
- 34 Marek Cygan, Fedor V Fomin, Łukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized algorithms*, volume 4. Springer, 2015.
- 35 Mina Dalirrooyfard, Thuy Duong Vuong, and Virginia Vassilevska Williams. Graph pattern detection: Hardness for all induced patterns and faster non-induced cycles. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 1167–1178, 2019.
- 36 Bartłomiej Dudek, Paweł Gawrychowski, and Tatiana Starikovskaya. All non-trivial variants of 3-LDT are equivalent. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 974–981, 2020.
- 37 Friedrich Eisenbrand and Fabrizio Grandoni. On the complexity of fixed parameter clique and dominating set. *Theoretical Computer Science*, 326(1-3):57–67, 2004.
- 38 John A Ellis, Ivan Hal Sudborough, and Jonathan S Turner. The vertex separation and search number of a graph. *Information and Computation*, 113(1):50–79, 1994.



- 39 Fedor V Fomin, Petr A Golovach, Daniel Lokshtanov, and Saket Saurabh. Almost optimal lower bounds for problems parameterized by clique-width. *SIAM Journal on Computing*, 43(5):1541–1563, 2014.
- 40 Fedor V Fomin, Daniel Lokshtanov, Venkatesh Raman, Saket Saurabh, and BV Raghavendra Rao. Faster algorithms for finding and counting subgraphs. *Journal of Computer and System Sciences*, 78(3):698–706, 2012.
- 41 François Le Gall and Florent Urrutia. Improved rectangular matrix multiplication using powers of the Coppersmith-Winograd tensor. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1029–1046. SIAM, 2018.
- 42 Edinah K Gngang, Ahmed Elgammal, and Vladimir Retakh. A spectral theory for tensors. In *Annales de la Faculté des sciences de Toulouse: Mathématiques*, volume 20, pages 801–841, 2011.
- 43 Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-SAT. *Journal of Computer and System Sciences*, 62(2):367–375, 2001.
- 44 Klaus Jansen, Stefan Kratsch, Dániel Marx, and Ildikó Schlotter. Bin packing with fixed number of bins revisited. *Journal of Computer and System Sciences*, 79(1):39–49, 2013.
- 45 Ce Jin and Hongxun Wu. A simple near-linear pseudopolynomial time randomized algorithm for subset sum. In *2nd Symposium on Simplicity in Algorithms (SOSA 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- 46 Konstantinos Koiliaris and Chao Xu. A faster pseudopolynomial time algorithm for Subset Sum. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1062–1072. SIAM, 2017.
- 47 Konstantinos Koiliaris and Chao Xu. Subset Sum made simple. *arXiv preprint arXiv:1807.08248*, 2018.
- 48 Marvin Künnemann, Ramamohan Paturi, and Stefan Schneider. On the fine-grained complexity of one-dimensional dynamic programming. In *44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.
- 49 François Le Gall. Powers of tensors and fast matrix multiplication. In *Proceedings of the 39th international symposium on symbolic and algebraic computation*, pages 296–303, 2014.
- 50 Yuan Li, Alexander Razborov, and Benjamin Rossman. On the  $AC^0$  complexity of Subgraph Isomorphism. *SIAM Journal on Computing*, 46(3):936–971, 2017.
- 51 Andrea Lincoln, Virginia Vassilevska Williams, and Ryan Williams. Tight hardness for shortest cycles and paths in sparse graphs. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1236–1252. SIAM, 2018.
- 52 Grazia Lotti and Francesco Romani. On the asymptotic complexity of rectangular matrix multiplication. *Theoretical Computer Science*, 23(2):171–185, 1983.
- 53 Josef Malík, Ondřej Suchý, and Tomáš Valla. Efficient implementation of Color Coding algorithm for Subgraph Isomorphism problem. In *International Symposium on Experimental Algorithms*, pages 283–299. Springer, 2019.
- 54 Dániel Marx. Can you beat treewidth? In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*, pages 169–179. IEEE, 2007.
- 55 Dániel Marx and Michal Pilipczuk. Everything you always wanted to know about the parameterized complexity of Subgraph Isomorphism (but were afraid to ask). In Ernst W. Mayr and Natacha Portier, editors, *31st International Symposium on Theoretical Aspects of Computer Science (STACS 2014)*, volume 25 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 542–553, Dagstuhl, Germany, 2014. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. URL: <http://drops.dagstuhl.de/opus/volltexte/2014/4486>, doi:10.4230/LIPIcs.STACS.2014.542.
- 56 Ron Milo, Shai Shen-Orr, Shalev Itzkovitz, Nadav Kashtan, Dmitri Chklovskii, and Uri Alon. Network motifs: simple building blocks of complex networks. *Science*, 298(5594):824–827, 2002.

- 57 Jaroslav Nešetřil and Svatopluk Poljak. On the complexity of the subgraph problem. *Commentationes Mathematicae Universitatis Carolinae*, 26(2):415–419, 1985.
- 58 Kevin Pratt. Waring rank, parameterized and exact algorithms. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 806–823. IEEE, 2019.
- 59 Andrzej Proskurowski and Jan Arne Telle. Classes of graphs with restricted interval models. *Discrete Mathematics and Theoretical Computer Science*, 3(4), 1999.
- 60 Bellman Richard. Dynamic programming. *Princeton University Press*, 89:92, 1957.
- 61 Liam Roditty and Virginia Vassilevska Williams. Fast approximation algorithms for the diameter and radius of sparse graphs. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 515–524, 2013.
- 62 Gregory Rosenthal. Beating treewidth for average-case subgraph isomorphism. In *14th International Symposium on Parameterized and Exact Computation (IPEC 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.
- 63 Paul D Seymour and Robin Thomas. Graph searching and a min-max theorem for tree-width. *Journal of Combinatorial Theory, Series B*, 58(1):22–33, 1993.
- 64 Andrew James Stothers. On the complexity of matrix multiplication. 2010.
- 65 Volker Strassen. Gaussian elimination is not optimal. *Numerische mathematik*, 13(4):354–356, 1969.
- 66 Edward H Sussenguth. A graph-theoretic algorithm for matching chemical structures. *Journal of Chemical Documentation*, 5(1):36–43, 1965.
- 67 Ryan Williams. A new algorithm for optimal 2-constraint satisfaction and its implications. *Theoretical Computer Science*, 348(2-3):357–365, 2005.
- 68 Virginia Vassilevska Williams and Ryan Williams. Subcubic equivalences between path, matrix and triangle problems. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 645–654. IEEE, 2010.

## A Proof of the Conditional Lower Bound for Subset Sum

In Section 4.3.1, we stated Theorem 32, which gives the following lower bound on SUBSET SUM:

*For no  $\varepsilon > 0$  can there be an algorithm which solves SUBSET SUM in time  $O(T^{1-\varepsilon} \text{poly}(n))$  unless the  $h$ -uniform HYPERCLIQUE hypothesis fails for all  $h \geq 3$ .*

We now prove this.

### A.1 Reduction from $k$ -SUM to Subset Sum

Before proving the result, we need a lemma that shows that we can reduce  $k$ -SUM to SUBSET SUM with minimal overhead. This theorem is already known, but we could not find a formal proof of it in the literature. Therefore we provide one here.

► **Lemma 48** (Reducing  $k$ -SUM to SUBSET SUM). *There is an algorithm  $\mathcal{B}$  which, given as input a  $k$ -SUM instance with  $N$  values from  $[0, D]$  per set, as well as a target  $T$ , constructs an equivalent SUBSET SUM instance with  $k \cdot N$  values in  $[0, Dg(k)]$  and a target  $T'$  bounded by  $Tg(k)$  for a computable function  $g$ . Furthermore, for constant  $k$ ,  $\mathcal{B}$  runs in time linear in the input size.*

**Proof.** The values of the SUBSET SUM instance are the union of the  $k$  sets. However, we modify the weights and target as follows. At the front of the binary representation of the weights and the target, we add a buffer of  $\lceil \log(k) \rceil$  zero bits to avoid overflow, then another  $k$  bits constituting a “checklist”, then in front of that another buffer of  $\lceil \log(k) \rceil$  zero bits and finally another  $\lceil \log(k) \rceil$  bits which contains a counter for the number of nodes of the solution.

The target  $T$  has the binary representation of  $k$  in the counter bits and only ones in the checklist bits. Each weight has the binary representation of 1 in the counter bits, ensuring that we take exactly  $k$  weights. Furthermore, if the weight comes from the  $i$ -th set of the  $k$ -SUM instances, its checklist bits is zero except for the  $i$ -th position.

Now if one picks more than  $2^{\lceil \log(k) \rceil}$  values, the counter at the front overflows the length of the target, so that selection cannot be a solution. Since  $2^{\lceil \log(k) \rceil} < 2k$  and since there is a buffer of  $\lceil \log(k) \rceil$  bits, the checklist cannot overflow into the counter. Hence any solution must pick exactly  $k$  weights. Hence, the only way to achieve all ones in the checklist bits of a sum of  $k$  weights is to pick exactly one weight from each of the  $k$  sets. This completes the reduction.

Since we add  $3\lceil \log(k) \rceil + k$  bits to the weights and the target, their value is multiplied by at most  $2^k \cdot 2^3 \cdot k^3$ . Choosing  $g(k) = 2^{k+3}k^3$ , we obtain the bounds from the lemma. ◀

### A.2 Modifying the Weighted Lemma to Prove the Lower Bound

To prove the theorem above, we modify the Weighted Lemma by adding a special case of parameters: All parameters are as before, but  $\beta = 1$  and  $r_2 = 0$ . In this case, the preimages in the resulting instance  $\mathcal{G}'$  instead have size  $n^{k/r_1}$ , and the maximum weight is  $W = \Theta(n^{(1+\varepsilon)k})$ .

To prove this special case of the Weighted Lemma, we can almost use the reduction from its proof. We must simply purge all parts of the construction that relate to the edges part. That is, no construction of  $V_2$  in step 2, no construction of  $S'_2$  in steps 3 and no construction of  $S_2$  in step 5. The reduction then yields an instance  $\mathcal{G}'$  as described above.

We now prove the theorem about SUBSET SUM. Essentially, the instances we get from the special case of the Weighted Lemma are  $k$ -SUM instances that we can then reduce to SUBSET SUM via the lemma in Appendix A.1 above.

**Proof (of Theorem 32).** Let  $h$  be given. Suppose there is an algorithm solving SUBSET SUM in time  $O(T^{1-\varepsilon}N^z)$  for some  $z \in \mathbb{N}$ . We use the special case of the Weighted Lemma with

- some  $\varepsilon'$  chosen later,
- $h' := h$ , and
- some arbitrary  $r_1 \in \mathbb{N}$  such that
  - $r_1 > \frac{\hat{c}}{h}$  (this ensures, again, that the running time  $O(n^{2h-1} + n^{\hat{c}k/(hr_1)})$  of the reduction is equal to  $O(n^{2h-1} + n^{k-\varepsilon})$  for some  $\varepsilon > 0$ , and can hence be ignored in the analysis), and
  - $r_1 > z\varepsilon'$  (this ensures that  $n^{zk/r_1} < n^{\varepsilon'k}$ ).

This yields a  $k \in \mathbb{N}$  and a reduction algorithm  $\mathcal{A}$  with the properties from the description of the special case above. In particular, the reduction algorithm produces an EXACT WEIGHT COLORED SUBGRAPH ISOMORPHISM where  $H$  consists only of isolated vertices with preimages of size  $O(n^{k/r_1})$  and maximum absolute weight  $W = \Theta(n^{(1+\varepsilon)k})$ . We make all weights positive by adding a large number to each, such that the target is  $T = \Theta(n^{(1+\varepsilon)k})$ . Note that this instance is also a  $(r_1 + \binom{hr_1}{h})$ -SUM instance, with each set of numbers being the set of weights in a preimage.

We now use the algorithm  $\mathcal{B}$  from Lemma 48 to convert this to a Subset Sum instance with  $N = O(n^{k/r_1})$  values in  $[0, \Theta(n^{(1+\varepsilon')k})]$  and target  $T = \Theta(n^{(1+\varepsilon')k})$ .

We can now solve the instance in time  $O(n^{(1-\varepsilon)(1+\varepsilon')k} n^{zk/r_1}) = O(n^{((1-\varepsilon)(1+\varepsilon')+\varepsilon')k})$ . Hence it suffices to choose  $\varepsilon'$  such that

$$(1 - \varepsilon)(1 + \varepsilon') + \varepsilon' < 1 \iff \varepsilon' < \frac{\varepsilon}{2 - \varepsilon}$$

Since  $\varepsilon \in (0, 1)$  and hence  $\frac{\varepsilon}{2-\varepsilon} \in (0, 1)$ , this is always possible. ◀

## B Hyperclique and Boolean $k$ -Wise Matrix Product are Equivalent

In this section, we show that  $k$ -uniform  $(k+1)$ -HYPERCLIQUE has a fast algorithm if and only if the  $k$ -WISE MATRIX PRODUCT problem has a fast algorithm. More formally:

► **Theorem 49.** *The  $k$ -uniform  $(k+1)$ -HYPERCLIQUE problem has an algorithm running in time  $O(n^{k+1-\varepsilon})$  (for some  $\varepsilon > 0$ ) if and only if the  $k$ -WISE MATRIX PRODUCT problem has an algorithm running in time  $O(n^{k+1-\varepsilon'})$  (for some  $\varepsilon' > 0$ ).*

In the proofs of both directions, we will reduce from and to the COLORED HYPERCLIQUE problem instead of the HYPERCLIQUE problem. This is because these problems are equivalent: In the one direction, we can use the reduction in step 1 in the proof of the Unweighted Lemma (Section 4.2). In the other direction, a trivial argument shows that simply throwing away the color homomorphism is enough.

Reducing from and to the COLORED HYPERCLIQUE problem simplifies the proofs of both directions. Both are a straightforward generalization of the known proof that the standard

boolean matrix product and the TRIANGLE problem<sup>10</sup> are subcubically equivalent<sup>11</sup> with respect to combinatorial algorithms. This was proven in [68] for the more general result of subcubic equivalence of the  $(\min, +)$  matrix product<sup>12</sup> and the NEGATIVE TRIANGLE problem<sup>13</sup> and is a fundamental result in Fine-Grained Complexity Theory.

## B.1 From Hyperclique to Boolean $k$ -Wise Matrix Product

► **Lemma 50.** *If the  $k$ -WISE MATRIX PRODUCT problem has an algorithm running in time  $O(n^{k+1-\varepsilon})$  (for some  $\varepsilon > 0$ ), then the  $k$ -uniform  $(k+1)$ -HYPERCLIQUE problem has an algorithm running in time  $O(n^{k+1-\varepsilon'})$  (for some  $\varepsilon' > 0$ ).*

**Proof.** First, recall that in the proof of theorem 35 we defined how to index tensors via configurations. We re-use this definition here.

Given a hyperclique instance  $(G, f)$  where  $G$  is a  $k$ -uniform hypergraph and  $f : V(G) \rightarrow V(C_{k+1})$  is a color homomorphism to the  $k$ -uniform  $(k+1)$ -hyperclique  $C_{k+1}$ , we solve it via a single application of the  $k$ -wise matrix product. To do this, we create  $k$  tensors  $A^1, \dots, A^k$  of order  $k$  with dimensions  $n \times \dots \times n$ . We initialize them with zeroes as entries. The tensor  $A^i$  will be indexed by configurations of  $V(C_{k+1}) \setminus \{i\}$  with ordering  $(1, \dots, i-1, k+1, i+1, \dots, k)$ . For each  $i$  and each configurations  $R$  of  $V(C_{k+1}) \setminus \{i\}$ , we set  $A^i[R]$  to one if and only if  $R(1) \dots R(i-1)R(i+1) \dots R(k+1) \in E(G)$ . We use truth values and 0/1 interchangeably to declutter notation.

Now let  $A(R) := \text{MP}_k(A^1, \dots, A^k)$ . Then for a configuration  $R$  of  $V(C_{k+1}) \setminus \{k+1\}$ , we have that

$$\begin{aligned} A[R] &= \bigvee_{\ell \in [n]} A^1[\ell, R(2), \dots, R(k)] \wedge \dots \wedge A^k[R(1), \dots, R(k), \ell] \\ &= \exists v' \in f^{-1}(k+1) : \forall i \in [k] : A^i[R(1), \dots, R(i-1), v', R(i+1), \dots, R(k)] \\ &= \exists v' \in f^{-1}(k+1) : \forall i \in [k] : R(1) \dots R(i-1)R(i+1) \dots R(v)v' \in E(G) \quad (\dagger) \end{aligned}$$

But now note that there is a hyperclique in  $G$  if and only if there exists a configuration  $R$  of  $V(C_{k+1}) \setminus \{k+1\}$  such that  $(\dagger)$  is true and  $R(1) \dots R(k) \in E(G)$ . Hence we simply calculate  $A[R]$  via the fast  $k$ -wise matrix product algorithm, then check if such a configuration exists by iterating through every possible configuration. This takes time  $O(n^{k+1-\varepsilon} + n^k) = O(n^{k+1-\varepsilon})$ . ◀

## B.2 From Boolean $k$ -Wise Matrix Product to Hyperclique

► **Proposition 51.** *Let  $T(n) = O(n^c)$  for some  $c \geq 1$ . If there is a  $T(n)$  time algorithm for the  $k$ -uniform  $(k+1)$ -HYPERCLIQUE problem running in time  $T(n)$ , then there is an  $O(T(n))$  time algorithm solving the  $k$ -uniform  $(k+1)$ -HYPERCLIQUE problem which also outputs the vertices of the  $(k+1)$ -hyperclique (if one exists).*

**Proof.** The proof is a straightforward binary-search-like algorithm. Let  $C_{k+1}$  be the  $k$ -uniform  $(k+1)$ -hyperclique and define  $V(C_{k+1}) = \{1, \dots, k+1\}$ . For each  $i \in [k+1]$ , we

<sup>10</sup> Given a 3-colored graph, check if it contains a triangle. Note that this is the colored  $k$ -uniform  $(k+1)$ -HYPERCLIQUE problem for  $k=2$ .

<sup>11</sup> Meaning one problem has a  $O(n^{3-\varepsilon})$  algorithm if and only if the other has one.

<sup>12</sup> This is the standard matrix product, but addition is replaced by min and multiplication is replaced by addition.

<sup>13</sup> Given a 3-colored graph with edge weights, check if it contains a triangle of negative total weight.

split the preimage  $f^{-1}(i)$  into two equal parts of size  $n/2$ . Now for each of the  $2^{k+1}$  possible  $(k+1)$ -tuples of halves from different preimages, we check if there exists a  $(k+1)$ -hyperclique between these halves in time  $T(n/2)$ . If no tuple has a hyperclique, we return no. If at least one tuple has a hyperclique, we recurse on one of them, with the new preimages being the halves selected in the tuple. The base case is reached if there is a single vertex in each preimage, which we then return.

The running time of this recursive algorithm is  $S(n) = 2^{k+1}T(n/2) + S(n/2) = 2^{k+1}(T(n/2) + T(n/4) + T(n/8) + \dots)$ . Since  $T(n) = O(n^c)$  for  $c \geq 1$ , we have  $S(n) = O(T(n))$ . ◀

Hence if we can check if a hyperclique exists, we can also find it in essentially the same time. We now use this to prove the second direction, which we already stated in lemma 45. Recall that the statement of lemma 45 is the following:

*If the  $k$ -uniform  $(k+1)$ -HYPERCLIQUE problem has an algorithm running in time  $O(n^{k+1-\varepsilon})$  (for some  $\varepsilon > 0$ ), then the  $k$ -WISE MATRIX PRODUCT problem has an algorithm running in time  $O(n^{k+1-\varepsilon'})$  (for some  $\varepsilon' > 0$ ).*

**Proof.** In this proof, we again use the definition of indexing tensors via configurations from the proof of theorem 35 (also used above).

Let the input tensors  $A^1, \dots, A^k$ , each of order  $k$  and dimensions  $n \times \dots \times n$  be given. Now consider the  $k$ -uniform  $(k+1)$ -hyperclique  $C_{k+1}$  and define  $V(C_{k+1}) = \{1, \dots, k+1\}$ . We let  $G$  be the graph with  $n$  vertices per preimage of  $C_{k+1}$ , with no hyperedges (yet). In the following, we construct the edges of  $G$  and the color homomorphism  $f : V(G) \rightarrow V(C_{k+1})$ .

For each  $i \in [k+1]$ , we let the tensor  $A^i$  be indexed by configurations of  $V(C_{k+1}) \setminus \{i\}$  with ordering  $(1, \dots, i-1, k+1, i+1, \dots, k)$ . Now for each configuration  $R$  of  $V(C_{k+1}) \setminus \{i\}$ , we add the hyperedge  $R(1) \dots R(i-1)R(i+1) \dots R(k+1)$  to  $G$  if and only if  $A^i[R] = 1$ . Hence we are encoding the  $i$ -th input tensor in the hyperedges between the preimages of  $V(C_{k+1}) \setminus \{i\}$ .

Finally, in  $V(C_{k+1}) \setminus \{k+1\}$ , we encode the all-ones tensor. That is, for each configuration  $R$  of  $V(C_{k+1}) \setminus \{k+1\}$ , we add the hyperedge  $R(1) \dots R(k)$  to  $G$ .

We now show how to take advantage of the fast algorithm for  $k$ -uniform  $(k+1)$ -HYPERCLIQUE to calculate the  $k$ -wise matrix product by using it repeatedly on subinstances of this new hyperclique instance. Let  $A_{\text{res}}$  be the output tensor, initialized to all-zeroes. Now take  $g \in \mathbb{N}$  (chosen later). For each  $i \in [k+1]$ , we split the preimage  $f^{-1}(i)$  into  $g$  parts (i.e. subsets of vertices)  $(f^{-1}(i))_1, \dots, (f^{-1}(i))_g$ , each of size  $n/g$  up to rounding.

For any  $(k+1)$ -tuple  $(i(1), \dots, i(k+1)) \in [g]^{k+1}$ , we do the following: (★) While there is a hyperclique between the parts  $(f^{-1}(1))_{i(1)}, \dots, (f^{-1}(k+1))_{i(k+1)}$  (which can be checked via the fast algorithm in time  $O((n/g)^{k+1-\varepsilon})$ ), we find this hyperclique in time  $O((n/g)^{k+1-\varepsilon})$  via the algorithm from Proposition 51. Let the vertices of the found hyperclique be  $j(1), \dots, j(k+1)$ . Then we set  $A_{\text{res}}[j(1), \dots, j(k)]$  to one and delete the edge  $j(1) \dots j(k)$  from  $G$ .

After having gone through all  $(k+1)$ -tuples, we output  $A_{\text{res}}$ .

We argue correctness. We use truth values and 0/1 interchangeably. Let the correct output be  $A_{\text{cor}} := \text{MP}_k(A^1, \dots, A^k)$ . We want to prove  $A_{\text{cor}} = A_{\text{res}}$ . Let  $A_{\text{cor}}$  and  $A_{\text{res}}$  be indexed by configurations of  $V(C_{k+1}) \setminus \{k+1\}$  with ordering  $(1, \dots, k)$ . Then for a configuration  $R$  of  $V(C_{k+1}) \setminus \{k+1\}$ , we have that  $A_{\text{cor}}[R]$  also satisfies the equation (†) from the proof of lemma 50.

Now suppose  $A_{\text{res}}[R] = 1$ . Then certainly (†) is true, so  $A_{\text{cor}}[R] = 1$ . Conversely, suppose that  $A_{\text{cor}}[R] = 1$  and hence that (†) is true. Then when we will find a hyperclique for the

tuple  $(i(1), \dots, i(k))$  such that  $\forall j : R(j) \in (f^{-1}(i))_{i(j)}$ . When we find it, the hyperedge  $R(1) \dots R(k)$  either still exists (in which case we set  $A_{\text{cor}}[R]$  to one) or it does not, in which case it must have been deleted in a previous iteration where we must have then already set  $A_{\text{cor}}[R]$  to one.

Let us analyze the running time. For each successful check in the while loop  $(\star)$ , we delete an edge in  $E(G)$ . Hence there are at most  $O(n^k)$  successful checks. In each successful check, we take time  $O((n/g)^{k+1-\varepsilon})$  for the check itself and another  $O((n/g)^{k+1-\varepsilon})$  for the execution of the while loop content. As for unsuccessful while loop checks, we have at most one per  $(k+1)$  tuple of parts, so  $g^{k+1}$ . Each of these also takes  $O((n/g)^{k+1-\varepsilon})$ .

Hence the overall running time is  $O(n^k(n/g)^{k+1-\varepsilon} + g^{k+1}(n/g)^{k+1-\varepsilon})$ . Choosing  $g^{k+1} = n^k$  and hence  $g = n^{k/(k+1)}$ , we have a running time of  $O(n^k n^{(1-k/(k+1))(k+1-\varepsilon)})$ , which works out to  $O(n^{k+1-\varepsilon/(k+1)})$ .  $\blacktriangleleft$