

On the Parameterized Complexity of Computing Tree-Partitions

Hans L. Bodlaender  

Department of Information and Computing Sciences, Utrecht University, The Netherlands

Carla Groenland  

Department of Information and Computing Sciences, Utrecht University, The Netherlands

Hugo Jacob  

ENS Paris-Saclay, France

Abstract

We study the parameterized complexity of computing the tree-partition-width, a graph parameter equivalent to treewidth on graphs of bounded maximum degree.

On one hand, we can obtain approximations of the tree-partition-width efficiently: we show that there is an algorithm that, given an n -vertex graph G and an integer k , constructs a tree-partition of width $O(k^7)$ for G or reports that G has tree-partition width more than k , in time $k^{O(1)}n^2$. We can improve slightly on the approximation factor by sacrificing the dependence on k , or on n .

On the other hand, we show the problem of computing tree-partition-width exactly is XALP-complete, which implies that it is $W[t]$ -hard for all t . We deduce XALP-completeness of the problem of computing the domino treewidth. Finally, we adapt some known results on the parameter tree-partition-width and the topological minor relation, and use them to compare tree-partition-width to tree-cut width.

2012 ACM Subject Classification Theory of computation \rightarrow Parameterized complexity and exact algorithms; Theory of computation \rightarrow Graph algorithms analysis; Theory of computation \rightarrow Approximation algorithms analysis

Keywords and phrases Parameterized algorithms, Tree partitions, tree-partition-width, Treewidth, Domino Treewidth, Approximation Algorithms, Parameterized Complexity

Digital Object Identifier 10.4230/LIPIcs.IPEC.2022.7

Related Version *Full Version*: <https://arxiv.org/abs/2206.11832> [8]

Funding *Carla Groenland*: Supported by the European Union's Horizon 2020 research and innovation programme under the ERC grant CRACKNP (number 853234) and the Marie Skłodowska-Curie grant GRAPHCOSY (number 101063180).

1 Introduction

Graph decompositions have been a very useful tool to draw the line between tractability and intractability of computational problems. There are many meta-theorems showing that a collection of problems can be solved efficiently if a decomposition of some form is given. By finding efficient algorithms to compute a decomposition if it exists, we deduce the existence of efficient algorithms even if the decomposition is not given. In particular, this proves useful when designing win-win arguments: for some problems, the existence of a solution and the existence of a decomposition are not independent, so that we can either use the decomposition for an efficient computation of the solution, or conclude that a solution must (or cannot) exist when there is no decomposition of small enough width.



© Hans L. Bodlaender, Carla Groenland, and Hugo Jacob;

licensed under Creative Commons License CC-BY 4.0

17th International Symposium on Parameterized and Exact Computation (IPEC 2022).

Editors: Holger Dell and Jesper Nederlof; Article No. 7; pp. 7:1–7:20

Leibniz International Proceedings in Informatics



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

The most successful notion of graph decomposition to date is certainly tree decompositions, and its corresponding parameter treewidth. Any problem expressible in MSO_2 ¹ can be solved in linear time in graphs of bounded treewidth due to a meta-theorem of Courcelle [13] and the algorithm of Bodlaender for computing an optimal tree decomposition [3]. Treewidth is a central tool in the study of minor-closed graph classes. A minor-closed graph class has bounded treewidth if and only if it contains no large grid minor.

In this paper, we focus on the parameter tree-partition-width (also called strong treewidth) which was independently introduced by Seese [27] and Halin [23]. It is known to have simple relations to treewidth [14, 29]: $\mathbf{tw} = O(\mathbf{tpw})$, and $\mathbf{tpw} = O(\Delta \mathbf{tw})$, where \mathbf{tw} , \mathbf{tpw} , Δ denote the treewidth, the tree-partition-width, and the maximum degree respectively. Applications of tree-partition-width include graph drawing and graph colouring [11, 22, 16, 17, 30, 2, 1]. Recently, Bodlaender, Cornelissen and Van der Wegen [5] showed for a number of problems (in particular, problems related to network flow) that these are intractable (XNLP-complete) when the pathwidth is used as parameter, but become fixed parameter tractable when parameterized by the width of a given tree-partition. This raises the question of the complexity of finding tree-partitions. We show that computing tree-partitions of approximate width is tractable.

► **Theorem 1.** *There is an algorithm that given an n -vertex graph G and an integer k , constructs a tree-partition of width $O(k^7)$ for G or reports that G has tree-partition width more than k , in time $k^{O(1)}n^2$.*

Thus, this removes the requirement from the results from [5] that a tree partition of small width is part of the input. Our technique is modular and allows us to also give alternatives running in FPT time or polynomial time with an improved approximation factor (see Theorem 10). Although not formulated as an algorithm, a construction of Ding and Oporowski [15] implies an FPT algorithm to compute tree-partitions of width $f(k)$ for graphs of tree-partition-width k , for some fixed computable function f . We adapt their construction and give some new arguments designed for our purposes. This significantly improves on the upper bounds to the width, and the running time.

The results from [5] are stated in terms of the notions of stable gonality, stable tree-partition-width and treebreadth. The notion of stability comes from the origin of the notion of gonality (from algebraic geometry); in graph terms, this implies that we look here at the minimum over all possible subdivision of edges. Tree-partition-width and stable tree-partition-width are bounded by polynomial functions of each other (see the appendix).

Related to tree-partition-width is the notion of domino treewidth, first studied by Bodlaender and Engelfriet [7]. A domino tree decomposition is a tree decomposition where each vertex is in at most two bags. Where graphs of small tree partition-width can have large degree, a graph of domino treewidth k has maximum degree at most $2k$. Bodlaender and Engelfriet show that DOMINO TREEWIDTH is hard for each class $W[t]$, $t \in \mathbb{N}$; we improve this result and show XALP-completeness .

► **Theorem 2.** *DOMINO TREEWIDTH and $\text{TREE PARTITION WIDTH}$ are XALP-complete .*

In [4], Bodlaender gave an algorithm to compute a domino tree decomposition of width $O(\mathbf{tw} \Delta^2)$ in $f(\mathbf{tw})n^2$ time for n -vertex graphs of treewidth \mathbf{tw} and maximum degree Δ , where f is a fixed computable function. This implies an approximation algorithm for domino treewidth.


¹ Formulae with quantification over sets of edges or vertices, quantification over vertices and edges, and with the incidence predicate.

We also consider the parameter *tree-cut width* introduced by Wollan in [28]. As the tractability results of Bodlaender et al. [5] use techniques similar to a previous work on algorithmic applications of tree-cut width [21], one may wonder whether there is a relationship between tree-cut width and tree-partition-width. We answer this question in two steps. We obtain a parameter that is polynomially tied to tree-partition-width and is topological minor monotone. In particular, this shows that tree-partition-width is relatively stable with respect to subdivisions. Then, we show how to relate tree-cut width to the tree-partition-width of a subdivision. We show that a bound on tree-partition-width does not imply a bound on tree-cut width and that tree-partition-width is polynomially bounded by tree-cut width.

Paper overview

In Section 3, we provide our results on approximating the tree-partition-width. In Section 4, we show that computing the tree-partition-width is XALP-complete. We then derive XALP-completeness of computing the domino treewidth. Our results relating tree-cut width to tree-partition-width are in the appendix.

2 Preliminaries



■ **Figure 1** An example of tree-partition.

A *tree-partition* of a graph $G = (V, E)$ is a tuple $(T, (B_i)_{i \in V(T)})$, where $B_i \subseteq V(G)$, with the following properties.

- T is a tree.
- For each $v \in V$ there is a unique $i(v) \in V(T)$ such that $v \in B_{i(v)}$.
- For any edge $uv \in E$, either $i(v) = i(u)$ or $i(u)i(v)$ is an edge of T .

The *size* of a bag B_i is $|B_i|$, the number of vertices it contains. The *width* of the decomposition is given by $\max_{i \in V(T)} |B_i|$. The *tree-partition-width* (**tpw**) of a graph G is the minimum width of a tree-partition of G .

A *tree decomposition* of a graph $G = (V, E)$ is a pair $(T = (I, F), \{X_i \mid i \in I\})$ with $T = (I, F)$ a tree and $\{X_i \mid i \in I\}$ a family of (not necessarily disjoint) subsets of V (called *bags*) such that $\bigcup_{i \in I} X_i = V$, for all edges $vw \in E$, there is an i with $v, w \in X_i$, and for all v , the nodes $\{i \in I \mid v \in X_i\}$ form a connected subtree of T . The *width* of a tree decomposition $(T, \{X_i \mid i \in I\})$ is $\max_{i \in I} |X_i| - 1$, and the *treewidth* (**tw**) of a graph G is the minimum width over all tree decompositions of G . The *domino treewidth* is the minimum width over all tree decompositions of G such that each vertex appears in at most two bags.

We say that two parameters α, β are (*polynomially*) *tied* if there exist (polynomial) functions f, g such that $\alpha \leq f(\beta)$ and $\beta \leq g(\alpha)$.

3 Approximation algorithm for tree-partition-width

We first describe our algorithm, then prove correctness and finally discuss the trade-offs between running time and solution quality.

3.1 Description of the algorithm

Let G be a graph, and k any positive integer. We describe a scheme that produces a tree-partition of G of width $O(wbk^3) = k^{O(1)}$, or reports that $\mathbf{tpw}(G) > k$. We will use various different functions of k for b and w , depending on the quality/time trade-offs of the black-box algorithms inserted into our algorithm (e.g. for approximating treewidth).

Step 1. *Compute a tree decomposition for G of width $w(k)$ or conclude that $\mathbf{tpw}(G) > k$.* As mentioned above, we do not directly specify the function $w = w(k)$, since different algorithms for step 1 give different solution qualities (bounds for $w(k)$) and running times. Since $\mathbf{tw} + 1 \leq 2 \cdot \mathbf{tpw}$, if $\mathbf{tw}(G) > 2k - 1$ it follows that $\mathbf{tpw}(G) > k$. If we obtained a decomposition of width w , we also know that $\mathbf{tw}(G) \leq 2k - 1$, and hence there are at most $(2k - 1)n$ edges in G .

We set a threshold $b \geq \max\{2k - 1, w + 1\}$. We define an auxiliary graph G^b as follows. The vertex set of G^b is $V(G)$. The edges of G^b are given by the pairs of vertices $u, v \in V(G)$ with minimum u - v separator of size at least b .

Step 2. *Construct the auxiliary graph G^b with connected components of size at most k or report that $\mathbf{tpw}(G) > k$.*

We later describe several ways of computing the edges of G^b . We will show in Claim 3 that vertices in the same connected component of G^b must be in the same bag for any tree-partition of width at most k . For this reason, we conclude that $\mathbf{tpw}(G) > k$ if a component of G^b has more than k vertices.

We define H , the b -reduction of G , which is the graph obtained from G by identifying the connected components of G^b .

Step 3. *We compute a tree decomposition of width w for each 2-connected component of H .* Given the components of G^b , we can compute H , and its 2-connected components in time $O(k^{O(1)}n)$. Using Claim 4, we obtain a tree decomposition of H by replacing vertices of G , in the tree decomposition of G , by their component in G^b .

By Claim 5, the maximum degree Δ_H within the 2-connected components of H is at most Cbk^2 for some constant C when $\mathbf{tpw}(G) \leq k$.

Step 4. *If $\Delta_H > Cbk^2$, report $\mathbf{tpw}(G) > k$. Else, compute a tree partition of width $O(w\Delta_H) = O(wbk^2)$ for H .*

By rooting the decomposition of H in 2-connected components, we can define a *parent* cutvertex for each 2-connected component except the root. We separately compute tree partitions for each 2-connected component of H with the constraint that their parent cutvertex should be the single vertex of its bag. A construction of Wood [29] enables us to compute a tree-partition of width $O(\Delta w)$ for any graph of maximum degree Δ and treewidth w ; this can be adjusted to allow for this isolation constraint without increasing the upper bound on the width. We give the details of this in Corollary 7. After doing this, the partitions of each component can be combined without increasing the width. Indeed, although cutvertices are shared, only one 2-connected component will consider putting other vertices in its bag. We obtain a tree-partition of H of width $O(wbk^2)$.

Step 5. Deduce a tree partition of width $O(wbk^3)$ for G .

We “expand” the vertices of H . In the tree-partition of H , each vertex of H is replaced by the vertices of the corresponding connected component of G^b . This gives a tree-partition of G of width $O(wbk^3)$.

3.2 Correctness

For $s, t \in V(G)$ we denote by $\mu(s, t)$ the size of a minimum s - t separator in $G - st$.

▷ **Claim 3.** Let G be a graph and $s, t \in V(G)$.

- If $\mu(s, t) \geq k + 1$, then in any tree-partition of width at most k , s and t must be in adjacent bags or the same bag;
- If $\mu(s, t) \geq 2k - 1$, then in any tree-partition of width at most k , s and t must be in the same bag.

Proof. Assume that s and t are not in adjacent bags nor in the same bag of a tree-partition of width at most k , then any internal bag on the path between their respective bags is an s - t separator. In particular, $\mu(s, t) \leq k$. This proves the first point by contraposition.

Assume that s and t are in adjacent bags but not in the same bag for some tree-partition of width at most k . We denote their respective bags by B_s and B_t . Then, $(B_s \cup B_t) \setminus \{s, t\}$ is an s - t separator of $G - st$. Consequently, $\mu(s, t) \leq 2k - 2$. This proves the second point. ◁

▷ **Claim 4.** Consider $(T, (X_i)_{i \in V(T)})$ a tree decomposition of width w of G , $b \geq w + 1$, and let Y_i be the set of connected components of G^b that intersect with X_i . Then $(T, (Y_i)_{i \in V(T)})$ is a tree decomposition of the b -reduction H of G .

Proof. Every component of G^b appears in at least one Y_i , because it contains a vertex which must appear in at least one X_i . Furthermore, for each edge UV of H , there must be vertices $u \in U, v \in V$ such that uv is an edge of G . Hence, there is a bag X_i containing u and v so Y_i contains u and v . Finally, suppose that there is a bag Y_i not containing a component C of G^b , and several components of $T - i$ have bags containing C . There must be an edge of G^b connecting vertices u and v of C such that u is in bags of X and v is in X' , where X and X' are in different components of $T - i$. By definition of G^b there are at least b vertex disjoint u, v -paths in G , so the minimal size of a separator of u and v is at least $b \geq w + 1$. However, since the tree decomposition $(T, (X_i))$ has width w and the bags containing u are disjoint from the bags containing v (in particular X_i separates them), there is a separator of u and v of size at most w , a contradiction. This concludes the proof that $(T, (Y_i))$ is a tree decomposition of H . ◁

▷ **Claim 5.** If H is the b -reduction of G , $\text{tpw}(G) \leq k$, and B is one of its 2-connected components, then the maximum degree in B is $O(bk^2)$.

Proof. Consider u a vertex achieving maximum degree in B . By definition of B , $B - u$ is connected. We denote by N the neighbourhood of u in B . Let T be a spanning tree of $B - u$. We iteratively remove leaves that are not in N , and contract edges with an endpoint of degree 2 that is not in N . This produces the reduced tree T' . The maximum degree in this tree is $b - 1$ as the set of edges incident to a given vertex can be extended to disjoint paths leading to vertices in N .

We call the number of vertices in the component of G^b associated to a vertex v of H the weight of v and denote it $|v|$.

Clearly the neighbours of u must be either in the same bag as u or in a neighbouring bag. Since the bag of u will be a separator of vertices that are in distinct neighbouring bags, in particular, it splits the graph into several components each containing neighbours of u of total weight at most k .

There must exist a subset of vertices of T' of size at most $k - |u|$ whose removal splits T' in components containing vertices of N of total weight at most k . Since the degree of a vertex of T' is at most $b - 1$, removing one of its vertices adds at most $b - 2$ new components. Hence, after removing $k - |U| \leq k - 1$ vertices, there are at most $1 + (k - 1)(b - 2)$ components. We conclude that $|N| \leq k(1 + (k - 1)(b - 2))$. Since u had maximum degree in B , we conclude. \triangleleft

In [29], Wood shows the following lemma.

► **Lemma 6.** *Let $\alpha = 1 + 1/\sqrt{2}$ and $\gamma = 1 + \sqrt{2}$. Let G be a graph with treewidth at most $k \geq 1$ and maximum degree at most $\Delta \geq 1$. Then G has tree-partition-width $\text{tpw}(G) \leq \gamma(k + 1)(3\gamma\Delta - 1)$.*

Moreover, for each set $S \subseteq V(G)$ such that $(\gamma + 1)(k + 1) \leq |S| \leq 3(\gamma + 1)(k + 1)\Delta$, there is a tree-partition of G with width at most $\gamma(k + 1)(3\gamma\Delta - 1)$ such that S is contained in a single bag containing at most $\alpha|S| - \gamma(k + 1)$ vertices.

We deduce this slightly stronger version of [29, Theorem 1]

► **Corollary 7.** *From a tree decomposition of width w in a graph G of maximum degree Δ , for any vertex v of G , we can produce a tree-partition of G of width $O(\Delta w)$ in which v is the only vertex of its bag.*

Proof. We wish to apply Lemma 6 to $S \supseteq N(v)$. Let $\gamma = 1 + \sqrt{2}$. We have $|N(v)| \leq \Delta$, so in particular, $|N(v)| \leq 3(\gamma + 1)(w + 1)\Delta$. In case $|N(v)| < (\gamma + 1)(w + 1)$, we can add arbitrary vertices to $N(v)$ to form S satisfying $|S| \geq (\gamma + 1)(w + 1)$. Otherwise, we simply set $S = N(v)$. We then apply the lemma to S in $G - v$. There is a single bag that contains $N(v)$, and so we may add the bag $\{v\}$ adjacent to this in order to deduce a tree partition of G of width $O(\Delta w)$ in which v is the only vertex of its bag. \blacktriangleleft

3.3 Time/quality trade-offs

For Step 1, we consider the following algorithms to compute tree decompositions:

- An algorithm of Korhonen [25] computes a tree decomposition of width at most $2k + 1$ or reports that $\text{tw}(G) > k$ in time $2^{O(k)}n$.
- An algorithm of Fomin et al. [19] computes a tree decomposition of width $O(k^2)$ or reports that $\text{tw}(G) > k$ in time $O(k^7 n \log n)$.
- An algorithm of Feige et al. [18] computes a tree decomposition of width $O(k\sqrt{\log k})$ or reports that $\text{tw}(G) > k$ in time $O(n^{O(1)})$.

Recall that we denote by w the width of the computed tree decomposition of G .

We give two methods to compute G^b in step 2 of the algorithm.

- We can use a maximum-flow algorithm (e.g. Ford-Fulkerson [20]) to compute for each pair $\{s, t\}$ of vertices of G whether there are at least b vertex disjoint paths from s to t , in time $O(bkn)$. To compute a minimum vertex cut, replace each vertex v by two vertices $v_{\text{in}}, v_{\text{out}}$ with an arc from v_{in} to v_{out} . All arcs going to v should go to v_{in} , and all arcs leaving v should leave v_{out} . All arcs are given capacity 1. We may stop the maximum flow algorithm as soon as a flow of at least b was found. Furthermore, we can reduce the number of pairs $\{s, t\}$ of vertices to check to $O(wn)$, as each pair must be contained in a bag due to $b \geq w + 1$. This results in a total time of $O(wbkn^2)$.

- We can also use dynamic programming to enumerate all possible ways of connecting pairs of vertices that are in the same bag in time $2^{O(w^2)}n$, which is sufficient to compute G^b . A state of the dynamic programming consists of the subset of vertices of the bag that are used by the partial solution, a matching on some of these vertices, up to two vertices that were decided as endpoints of the constructed paths, the number of already constructed paths between the endpoints, and two disjoint subsets of the used vertices that are not endpoints, nor in the matching such that we found a disjoint path from the first or second endpoint to them. We first tabulate answers for each subtree of the decomposition by starting from the leaves, and then tabulate answers for each complement of a subtree by starting from the root and, when branching to some child, combining with the partial solutions of the subtree of the other child.

The b -reduction H of G and its 2-connected components can be computed in $O(k^{O(1)}n)$ time (see e.g. [24]), since the size of the graph is $O(k^{O(1)}n)$ here.

We will now make use of the following result due to Bodlaender and Hagerup [10]:

- **Lemma 8.** *There is an algorithm that given a tree decomposition of width k with $O(n)$ nodes of a graph G , finds a rooted binary tree of G of width at most $3k + 2$ with depth $O(\log n)$ in $O(kn)$ time.*

When implementing the construction of Wood for 2-connected components of H , the running time is dominated by $O(n)$ queries to find a balanced separator with respect to a set W of size $k^{O(1)}$. After a preprocessing in time $O(kn)$, we can do this in time $k^{O(1)}d$ where d is the diameter of our tree decomposition. We first obtain a binary balanced decomposition using Lemma 8, then reindex the vertices in such a way that we can check if a vertex is in some bag of a given subtree of tree decomposition in constant time. Using this, we can in time $k^{O(1)}$ determine whether a bag is a balanced separator of W , and if not move to the subtree containing the most vertices of W . This procedure will consider at most d bags, hence the total running time of $k^{O(1)}d$. Since the decomposition has depth $O(\log n)$ it also has diameter $d = O(\log n)$. Hence the construction of Wood can be executed in time $k^{O(1)}n \log n$.

- **Lemma 9.** *We can compute a tree partition of width $O(\Delta \mathbf{tw})$ in time $O(k^{O(1)}n \log n)$ when given a tree decomposition of width $k^{O(1)}$.*

By combining the previous algorithms we obtain the following theorem.

- **Theorem 10.** *There is a polynomial time algorithm that constructs a tree-partition of width $O(k^5 \log k)$ or reports that the tree-partition-width is more than k .*

There is an algorithm running in time $2^{O(k^2)}n + k^{O(1)}n \log n$ that computes a tree-partition of width $O(k^5)$ or reports that the tree-partition-width is more than k .

There is an algorithm running in time $k^{O(1)}n^2$ that computes a tree-partition of width $O(k^7)$ or reports that the tree-partition-width is more than k .

Proof. The first algorithm uses the algorithm of Feige et al. to compute the tree decomposition, then naively computes G^b , and then finds balanced separators for Wood's construction using the tree decomposition in polynomial time (no need to balance the decomposition).

The second algorithm uses Korhonen's algorithm to compute the tree decomposition, then computes G^b using the dynamic programming approach, and then finds balanced separators for Wood's construction as described.

The third algorithm uses the algorithm of Fomin et al. to compute the tree decomposition, then computes G^b via a maximum-flow algorithm in time $O(wbkn^2) = O(k^5n^2)$, and then finds balanced separators for Wood's construction as described.

The guarantees on the width follow from the analysis of our scheme. ◀

4 XALP-completeness of Tree Partition Width

In this section, we show that the TREE PARTITION WIDTH problem is XALP-complete, even when we use the width target and the degree as combined parameter. As a relatively simple consequence, we obtain that DOMINO TREEWIDTH is XALP-complete.

XALP is the class of all parameterized problems that can be solved in $f(k)n^{O(1)}$ time and $f(k)\log n$ space on a nondeterministic Turing Machine with access to a push-down automaton, or equivalently way of the class of problems that can be solved by an alternating Turing Machine in $f(k)n^{O(1)}$ treesize and $f(k)\log(n)$ space. An alternating Turing Machine (ATM) is nondeterministic Turing Machine with some extra states where we ask for all of the transitions to lead to acceptance. This creates independent configurations that must all lead to acceptance, and we call “co-nondeterministic step” the process of obtaining these independent configurations.

XALP is closed by reductions using at most $f(k)\log n$ space and running in FPT time. These two conditions are implied by using at most $f(k) + \log n$ space. We call reductions respecting the latter condition *parameterized logspace* reductions (or pl-reductions).

This class is relevant here because the problems we consider are complete for it. Completeness for XALP has the following consequences: $W[t]$ -hardness for all positive integers t , membership in XP, and there is a conjecture that XP space is required for algorithms running in XP time. If the conjecture holds, this roughly means that the dynamic programming algorithm used for membership is optimal.

The following problem is shown to be XALP-complete in [9] and starts as the starting point of our reduction.

TREE-CHAINED MULTICOLOR INDEPENDENT SET

Input: A tree $T = (V_T, E_T)$, an integer k , and for each $i \in V_T$, a collection of k pairwise disjoint sets of vertices $V_{i,1}, \dots, V_{i,k}$ and a graph G with vertex set $V = \bigcup_{i \in V_T, j \in [1,k]} V_{i,j}$

Parameter: k

Question: Is there a set of vertices $W \subset V$, such that W contains exactly one vertex from each $V_{i,j}$ ($i \in V_T, j \in [1,k]$), and for each pair $V_{i,j}, V_{i',j'}$ with $i = i'$ or $ii' \in E_T$, $j, j' \in [1,k]$, $(i,j) \neq (i',j')$, the vertex in $W \cap V_{i,j}$ is non-adjacent to the vertex in $W \cap V_{i',j'}$?

We further use that we can assume the tree T to be binary without loss of generality (see [9] for more details).

► **Lemma 11.** TREE PARTITION WIDTH is in XALP.

Proof. To keep things simple, we will use as a black box the fact that reachability in undirected graphs can be decided in logspace [26]. We assume that the vertices have some arbitrary ordering σ .

For now, assume that the given graph is connected.

We begin by guessing at most k vertices to form an initial bag B_0 , and have an empty parent bag P_0 . We will recursively extend a partial tree-partition in the following manner. Suppose that we have a bag B with parent bag P , we must find a child bag for B in each connected component of $G - B$ that does not contain a vertex of P . We use the fact that a connected component can be identified by its vertex appearing first in σ , that the restriction of σ to these representatives gives an ordering on σ , and that we can compute such representatives in logspace. Let us denote by c the current vertex representative of a connected component of $G - B$. c is initially the first vertex in σ that is not in B and

cannot reach P in $G - B$. We do a co-nondeterministic step so that in one branch of the computation we find a tree-partition for the connected component with representative c , and in the other branch we find the representative of the next connected component. The representative c' of the next component is the first vertex in σ such that it cannot reach a vertex appearing before c (inclusive) in σ , nor a vertex of P in $G - B$. When found, c is replaced by c' and we repeat this computation. If we don't find such a vertex c' , then c must have represented the last connected component, so we simply accept.

Let us now describe what happens in the computation branch where we compute a new bag. We can iterate on vertices in the component of c , by iterating on vertices of $G - B$ and then skipping if they are not reachable from c in $G - B$. In particular, we can guess a subset B' of size at most k of vertices from this component. We then check that the neighbourhood of B in this component is contained in B' . If it is the case, we can set $P := B$ and $B := B'$ and recurse. If not, we reject.

If the graph is not connected, we can iterate on its connected components by using the same technique of remembering a vertex representative. For each of these components, we apply the above algorithm, with the modification that in each enumeration of the vertices we skip the vertex if it is not contained in the current component.

During these computations, we store at most $3k + O(1)$ vertices and use logspace subroutines. Furthermore, the described computation tree is of polynomial size. ◀

We first give a brief sketch of the structure of the hardness proof. We have a *trunk* gadget to enforce the shape of the tree from the TREE-CHAINED MULTICOLOR INDEPENDENT SET. On the trunk are attached *clique chains* which are longer than the part of trunk between their endpoints, and have some wider parts at some specific positions. The length of the chain gives us some slack which will be used to encode the choice of a vertex for some subset $V_{i,j}$. Based on the edges of G , we adjust the width along the trunk so that only one clique chain may place its wider part on each position of the trunk. In other words, part of the trunk is a collection of gadgets representing edges of G that allow for only one incident vertex to be chosen.

► **Lemma 12.** TREE PARTITION WIDTH *with target width and maximum degree as combined parameter is XALP-hard.*

Proof. We reduce from TREE-CHAINED MULTICOLOR INDEPENDENT SET.

Suppose that we are given a binary tree $T = (V_T, E_T)$, and for each node $i \in V_T$, a k -colored vertex set V_i . We denote the colors by integers in $[1, k]$, and write $V_{i,j}$ for the set of vertices in V_i with color j . We are also given a set of edges E of size m . Each edge in E is a pair of vertices in $V_i \times V_{i'}$ with $i = i'$ or ii' an edge in E_T . We can assume the edges are numbered: $E = \{e_1, \dots, e_m\}$.

In the TREE-CHAINED MULTICOLOR INDEPENDENT SET problem, we want to choose one vertex from each set $V_{i,j}$, $i \in V_T$, $j \in [1, k]$, such that for each edge $ii' \in E_T$, the chosen vertices in $V_i \cup V_{i'}$ form an independent set (which thus will be of size $2k$).

We assume that each set $V_{i,j}$ is of size r . (If not, we can add vertices adjacent to all other vertices in $V_{i,j}$, $j \in [1, k]$.) Write $V_{i,j} = \{v_{i,j,1}, v_{i,j,2}, \dots, v_{i,j,r}\}$.

Let $N = (m + 1)r$. Let $L = 36k + 5$.

Cluster Gadgets. In the construction, we use a *cluster gadget*. Suppose Z is a clique. Adding a cluster gadget for Z is the following operation on the graph that is constructed. Add a clique with vertex set $C_Z = \{c_{Z,1}, c_{Z,2}, \dots, c_{Z,2L}\}$ of size $2L$ to the graph, and add an edge between each vertex in Z and each vertex $c_{Z,j}$, $1 \leq j \leq L$, i.e. Z with the first L vertices in C_Z forms a clique.

7:10 On the Parameterized Complexity of Computing Tree-Partitions

In a tree partition of a graph, the vertices of a clique can belong to at most two different bags. The cluster gadget ensures that the vertices of clique Z belong to exactly one bag. This cluster gadget will be used in two different steps in the construction of the reduction.

► **Lemma 13.** *Suppose a graph H contains a clique Z with the cluster gadget for Z . In each tree partition of H of width at most L , there is a bag that contains all vertices from Z .*

Proof. There must be two adjacent bags that contain the vertices of C_Z and no other vertices. Similarly, there must be two adjacent bags containing all vertices in $Z \cup \{c_{Z,1}, \dots, c_{Z,L}\}$. This forces all vertices in $\{c_{Z,1}, \dots, c_{Z,L}\}$ to be in a single bag, and all vertices in Z to be in a single adjacent bag. ◀

A subdivision of T . The first step in the construction is to build a tree $T' = (V_{T'}, E_{T'})$, as follows. Choose an arbitrary node i from V_T . Add a new neighbor i' to i . Add a new neighbor r_0 to i' . Now subdivide each edge N times. The resulting tree is $T' = (V_{T'}, E_{T'})$. We view T' as a rooted tree, with root r_0 . We will use the word *grandparent* to refer to the parent of the parent of a vertex. The nodes that do not result from the subdivisions are referred to as *original nodes*. Nodes $i \in V_T$ and their copies in T' will be denoted with the same name.

The graph H consists of two main parts: the *trunk* and the *clique chains*. To several cliques in these parts, we add cluster gadgets.

The trunk. The trunk is obtained by taking for each node $i \in V_T$ a clique A_i . We specify below the size of these cliques. For each edge ii' in T' , we add an edge between each vertex in A_i to each vertex in $A_{i'}$. We add for each A_i a cluster gadget.

To specify the sizes of sets A_i , we first need to give some definitions:

- For each node $i' \in V_{T'}$, we let $p(i')$ be the number of nodes $i \in V_T$ (i.e., “original nodes”), such that i' is on the path (including endpoints) in T' from i to the vertex that is the grandparent of i in T . I.e., for each original node i , we look to the grandparent of i (if it exists), and then add 1 to the count of each node i' on the path between them in T' .
- For each edge $e_j \in E$, let $g(e_j) = 2jr$.
- For each edge $e_j = \{v_{i,c,s}, v_{i',c',s'}\}$, we have that $i = i'$ or i' is a child of i . Let i_{e_j} be the node in T' , obtained by making $g(e_j)$ steps up in T' from i : i.e., i_{e_j} is the ancestor of i with distance $g(e_j)$ in T' .

Now, for all nodes $i \in V_{T'}$,

- $|A_i|$ equals $L - 6k \cdot p(i) - 1$, if $i = i_{e_j}$ for some $e_j \in E$. At this node, we will verify that a choice (encoded by the clique chains, explained below), indeed gives an independent set: we check that we did not choose both endpoints of e_j .
- $|A_i|$ equals $L - 6k \cdot p(i) - 2$, otherwise.

The clique chains. For each $i \in V_T$, and each color class $c \in [1, k]$, we have a clique chain with $2N + r + 5$ cliques, denoted $CC'_{i,c,j}$, $j \in [1, 2N + r + 5]$. All vertices in the first clique $CC'_{i,c,1}$ are made incident to all vertices in A_i . All vertices in the last clique $CC'_{i,c,2N+r+5}$ are made incident to all vertices in $A_{i'}$ with i' the parent of the parent (i.e., the grandparent) of i in T . (Notice that the distance from i to i' in T' equals $2(N+1)$.) All vertices in $CC'_{i,c,j}$ are made incident to all vertices in $CC'_{i,c,j+1}$, i.e., all vertices in a clique are adjacent to all vertices in the next clique in the chain.

To each clique $CC'_{i,c,j}$ we add a cluster gadget.

The cliques have different sizes, which we now specify. Consider $i \in V_T$, $c \in [1, k]$, $\gamma \in [1, 2N + r + 5]$. The size of $CC_{i,c,\gamma}$ equals:

- $L - 7$, if $\gamma = 1$ or $\gamma = 2N + r + 5$ (i.e., for the first and last clique in the chain.)
- 7, if there is an edge e_j with one endpoint in $V_{i,c}$ for which one of the following cases holds:
 - $e_j = \{v_{i,c,\alpha}, v_{i',c',\alpha'}\}$, $c \neq c'$, i.e., one endpoint is in $V_{i,c}$, and the other endpoint is in another color class in V_i , and $\gamma = g(e_j) + 1 + \alpha$.
 - $e_j = \{v_{i,c,\alpha}, v_{i',c',\alpha'}\}$, i' is a child of i , and $\gamma = g(e_j) + 1 + \alpha$.
 - $e_j = \{v_{i,c,\alpha}, v_{i',c',\alpha'}\}$, i is a child of i' , and $\gamma = g(e_j) + N + 2 + \alpha$.
- 6, otherwise

Let H be the resulting graph.

► **Lemma 14.** *H has tree partition width at most L , if and only if the given instance of TREE-CHAINED MULTICOLOR INDEPENDENT SET has a solution.*

Proof. Suppose we have a solution of the TREE-CHAINED MULTICOLOR INDEPENDENT SET. Suppose for each class $V_{i,c}$, we choose the vertex $v_{i,c,h(i,c)}$. Now, we can construct the tree partition as follows. First, we take the tree T' , and for each node i in T' , we take a bag initially containing the vertices in A_i ; we later add more vertices to these bags in the construction.

Now, we add the chains, one by one. For a chain for $V_{i,c}$, take a new bag that contains $CC_{i,j,1}$, and make this bag incident to the bag of i . We add the vertices of $CC_{i,c,h(i,c)+1}$ to the bag of i . If $h(i,c) > 1$, then we place the vertices of cliques $CC_{i,c,\alpha+1}$ with $1 < \alpha < h(i,c)$ in bags outside the trunk: $CC_{i,c,h(i,c)}$ goes to the bag with $CC_{i,c,1}$; to this bag, we add an adjacent bag with $CC_{i,c,2} \cup CC_{i,j,h(i,c)-1}$; to this, we add an adjacent bag with $CC_{i,j,3} \cup CC_{i,j,h(i,c)-2}$, etc.

Now, add the vertices of $CC_{i,c,h(i,c)+2}$ to the bag of the parent of i , and continue this: each next clique is added to the next parent bag, until we add a clique to the bag of the parent of the parent of i in T ; name this node here i'' . Add a new bag incident to i'' and put $CC_{i,c,2N+r+5}$ in this bag (i.e., the last clique of the chain). Similar as at the start of the chain, fold the end of the chain (with possibly some additional new bags) such that a bag containing $CC_{i,c,2N+r+4}$ is adjacent to the bag with $CC_{i,c,2N+r+5}$.

Finally, for each cluster gadget, add two new bags, with the first incident to the bag containing the respective clique.

One easily verifies that this gives a tree partition of H . For bags outside the trunk, one easily observes that the size is at most L . Bags i in the trunk contain a set A_i , and precisely $p(i) \cdot k$ cliques of the clique chains: for each path that counts for the bag, and each color class in $[1, k]$, we have one chain with one clique. Each of these cliques has size six or seven. Now, we can notice that a clique of size 7 corresponds to an edge $e_{j'}$ with endpoint in the class. This clique will be mapped to a node in the trunk that equals $i_{e_{j'}}$, if and only if this endpoint is chosen; otherwise, the clique will be mapped to a trunk node with distance less than r to $i_{e_{j'}}$. Thus, there are two cases for a trunk node i :

- There is no edge e_j with $i = i_{e_j}$. Then, a close observation of the clique chains shows that there are at most two clique chains with size 7 mapped to i . Indeed, the construction is such that each edge has its private interval, and affects the trunk both between i and its parent i' , and between i' and its parent i'' .
- $i = i_{e_j}$. Now, at most one endpoint of e_j is in the solution. The clique chain of the color class of that endpoint can have a clique of size 7 mapped to i . The “offset” of the clique chain for the color class of the other endpoint is such that there is a clique of size 6 for that chain at i .

7:12 On the Parameterized Complexity of Computing Tree-Partitions

In both cases, the total size of the bag at i is at most L . Thus, the width of the tree partition is at most L .

Suppose we have a tree partition of H of width L . First, by the use of the cluster gadgets, each clique A_i is in one bag. A bag cannot contain two cliques A_i as each has size larger than $L/2$. Now, the bags containing A_i form a subtree of the partition tree that is isomorphic to T' . For each clique chain of a class $V_{i,c}$, we have that the first clique $CC_{i,c,1}$ is in a bag incident to i , and the last clique $CC_{i,c,2N+r+5}$ is in a bag incident to the trunk bag that corresponds to the grandparent of i in T , say i'' . Each trunk bag from i to i'' thus must contain a clique (of size 6 or 7) from the clique chain of $V_{i,c}$. It follows that each trunk bag i' contains at least $p(i) \cdot k$ cliques of size at least 6 each of the clique chains. Now, $|A_{i'}| + 6p(i) \cdot k \in \{L-1, L-2\}$, and thus we cannot add another clique of a clique chain to a trunk bag.

For a clique chain of $V_{i,c}$, there is a clique mapped to the trunk bag of i . Suppose $CC_{i,c,h(i,c)}$ is mapped to i . We claim that choosing from each $V_{i,c}$ the vertex $v_{i,c,h(i,c)}$ gives an independent set, and thus, we have a solution of the TREE CHAINED MULTICOLOR INDEPENDENT SET problem.

The vertices of $CC_{i,c,h(i,c)+2}$ must be mapped to the bag of the parent of i , as otherwise, i will contain an additional clique of size at least 6, and the size of the bag of i will become larger than L . By induction, we have that the α th parent of i , $\alpha \in [1, 2N+2]$ contains the vertices of $CC_{i,c,h(i,c)+\alpha+1}$. (Note that the $(2N+2)$ nd parent equals the node corresponding the grandparent of i in T .)

We now consider the node i_{e_j} for edge $e_j \in E$. Suppose $e_j = v_{i,c,\alpha}v_{i',c',\alpha'}$. Without loss of generality, suppose $i = i'$ or i' is a child of i ; otherwise, switch roles of i and i' . For each endpoint of this edge, if the endpoint is chosen (i.e., $\alpha = h(i,c)$ or $\alpha' = h(i',c')$), then the corresponding clique chain has a clique of size 7 in the bag i_{e_j} . This can be seen by the following case analysis:

- By assumption, $CC_{i,c,h(i,c)+1}$ is placed in the bag of i . As each successive clique in the chain is placed in one higher bag along the path from i to the grandparent of i (in T), we have that $CC_{i,c,h(i,c)+g(e_j)+1}$ is placed in the bag of i_{e_j} , as this node is the $g(e_j)$ th parent of i in T' . This clique has size 7.
- If $i' = i$, the same argument shows that $CC_{i',c',h(i',c')+1}$ is a clique of size 7 placed in the bag of i_{e_j} .
- If i' is a child of i in T , then $CC_{i',c',h(i',c')+1}$ is placed in the bag of i' . Again, each successive clique in the chain of $V_{i',c'}$ is placed in the next parent bag, for all nodes on the path from i' to the grandparent of i' in T (which is the parent of i in T .) This implies that $CC_{i',c',h(i',c')+N+2}$ is placed in the bag of i and $CC_{i',c',h(i',c')+N+2+g(e_j)}$ is placed in the bag of i_{e_j} ; this bag has size 7.

Now, if both endpoints would be chosen, then the size of the bag of i_{e_j} is larger than L : it contains $A_{i_{e_j}}$ (which has size $L - 6kp(i_{e_j}) - 1$), $6p(i_{e_j})$ bags of clique chains, of which all have size at least 6 and two have size 7; contradiction. So, at most one endpoint is chosen, so choosing vertices $v_{i,c,h(i,c)}$ gives an independent set. ◀

The maximum degree of a vertex in H is less than $5kL + 5L = O(k^2)$:

- Vertices in cluster gadgets have maximum degree less than $2L$.
- A vertex in a trunk clique A_i of a node i that resulted from a subdivision have maximum degree less than $4L$ as i has two incident nodes, each with a trunk clique of size less than L , and there is a cluster gadget attached to A_i .

- A vertex in a trunk clique A_i of a node i that is an original node (i.e., also in T) have less than L neighbors in A_i , less than $3L$ neighbors in $A_{i'}$ with i' incident to i in T' , less than kL neighbors of cliques $CC_{i,c,1}$ (one clique of size $L - 7$ for each class $c \in [1, k]$), less than $4kL$ neighbors of cliques $CC_{i',c,2N+r+5}$ (one clique of size $L - 7$ for each node of which i is the grandparent in T for each class $c \in [1, k]$), and less than L neighbors in the cluster gadget attached to A_i .

Finally, we conclude that the transformation can be carried out in $f(k) \log n$ space, thus the result follows. ◀

From Lemmas 11 and Lemma 12, the following result directly follows.

► **Theorem 15.** TREE PARTITION WIDTH is XALP-complete, both when the target value, and when the target value plus the maximum degree is used as parameter.

► **Theorem 16.** DOMINO TREewidth is XALP-complete.

Proof. Membership: We use the fact that the maximum degree of the graph is bounded by $2k$ where k is the domino treewidth. We can discard an instance where this condition on the maximum degree is not satisfied in logspace. We first assume that the given graph is connected.

The “certificate” used for this computation will be of size $O(k^2 + k \log(n))$ and consists of:

- The current bag and for each of its vertices whether it was contained in a previous bag or not. This requires at most $k + k \log(n)$ bits.
- For each neighbour of the bag, whether it was already covered by a bag. This requires $O(k^2)$ bits.

The algorithm works as follows. Given the current certificate, if all neighbours have been covered we accept. Otherwise, we guess a new child bag by picking a non-empty subset of $k + 1$ vertices among the vertices of the current bag that were contained only in this bag, and the neighbours that were not already covered. We then check that each vertex that is in both the current bag and child bag has all of its neighbours in these two bags. We then guess for each not already covered neighbour of the current bag if it should be covered by the subtree of this child. These vertices are then considered as covered in the current bag certificate. In the child bag certificate, the non covered neighbours are these vertices and the neighbours of the child bag that are not neighbours of the parent bag. We then recurse with both certificates, and accept if both recursions accept.

This computation uses $O(k^2 + k \log n)$ space and the computation tree has polynomial size.

We can handle disconnected graphs by iterating on component representatives and discarding vertices that are not reachable using the fact that reachability in undirected graphs can be computed in logspace (see membership for TREE PARTITION WIDTH for more details).

Hardness follows from a reduction from TREE PARTITION WIDTH when we use the target value and maximum degree as parameter.

Suppose we are given a graph $G = (V, E)$ of maximum degree d and an integer k .

Let $L = kd + 1$, and $M = (k + 1)L - 1$. Now, build a graph H as follows. For each vertex $v \in V$, we take a clique C_v with L vertices. For each vertex $w \in C_v$, we add a set S_w with $2M - 2$ vertices, and make one of the vertices in S_w incident to w and to all other vertices in S_w ; call this vertex y_w .

For each edge $e = vw \in E$, we add a vertex z_e , and make z_e incident to all vertices in C_v and all vertices in C_w . Let H be the resulting graph.

► **Lemma 17.** *G has tree partition width at most k , if and only if H has domino treewidth at most $M - 1$.*

Proof. Suppose H has domino treewidth at most $M - 1$. Suppose $(\{X_i | i \in I\}, T = (I, F))$ is a domino tree decomposition of H of width at most $M - 1$, i.e., each bag has size at most M .

First, consider a vertex w in some C_v . The vertex y_w has degree $2M - 2$, which implies that there are two adjacent bags that each contain y_w , and $M - 1$ neighbors of y_w . One of these bags contains w .

For each $v \in V$, there must be at least one bag that contains all vertices of C_v , by a well known property of tree decompositions. There can be also at most one such bag, because each vertex $w \in C_v$ is in another bag that is filled by w , y_w , and $M - 2$ other neighbors of y_w .

Let for $i \in I$, $Y_i \subseteq V$ be the set of vertices $v \in V$ with $C_v \subseteq X_i$. We claim that $(\{Y_i | i \in I\}, T = (I, F))$ is a tree partition of G of width at most k (some bags are empty). First, by the discussion above, each vertex $v \in V$ belongs to exactly one bag Y_i . Second, as $M < (k + 1)L$, each Y_i has size at most k . Third, if we have an edge $e = vw \in E$, then z_e is in the bag that contains C_v , and z_e is in the bag that contains C_w . As z_e is in at most two bags, these two bags must be the same, or adjacent, so in $(\{Y_i | i \in I\}, T = (I, F))$, v and w are in the same set Y_i or in sets Y_i and $Y_{i'}$ with i and i' adjacent in the decomposition tree T .

Now, suppose G has tree partition width at most k , say with tree partition $(\{Y_i | i \in I\}, T = (I, F))$. For each $i \in I$, let $X_i = \bigcup_{v \in Y_i} C_v \cup \{z_e \mid \exists v \in Y_i, w \in V : e = vw\}$. For each $v \in V$, $w \in C_v$, add two bags, one containing w , y_w , and $M - 2$ other neighbors of y_w , and the other containing y_w and the remaining $M - 1$ neighbors of y_w , and make these bags adjacent, and the first adjacent to the bag in T that also contains w . One easily verifies that this results in a domino tree decomposition of H with maximum bags size at most M , hence H has domino treewidth at most $M - 1$. ◀

It is easy to see that H can be constructed from G with $f(k) \log |V|$ memory. So, the hardness of DOMINO TREewidth follows from the previous lemma. ◀

5 Conclusion

We settle the question of the exact computation of tree-partition-width, and show that its approximation is tractable. However, many questions remain regarding approximation algorithms:

- Is a constant factor approximation tractable?
- Can we improve the approximation ratios with similar running times?
- Can we improve the dependence on n to something linear in n ?

Regarding the running time of the approximation algorithm, the following results could improve the final running times. Firstly, although the techniques of [6] are nontrivial, it seems reasonable to hope that they can be adapted to implement Wood's construction in time $f(k)n$. This would give an algorithm running in time $f(k)n$ to compute a tree-partition of width $k^{O(1)}$. Secondly, is there some value of b polynomial in w and k such that we can compute G^b in time $k^{O(1)}n \log n$ or in time $2^{o(k^2)}n$? This would directly give faster running times for the approximation algorithm, possibly at the cost of a worse approximation ratio.

The parameter treebreadth studied in [5] is polynomially tied to tree-partition-width, but might be easier to approximate.

Another interesting direction is to study the complexity of computing (approximate) tree decompositions on graphs of bounded tree-partition-width.

References

- 1 Noga Alon, Guoli Ding, Bogdan Oporowski, and Dirk Vertigan. Partitioning into graphs with only small components. *J. Comb. Theory, Ser. B*, 87(2):231–243, 2003. doi:10.1016/S0095-8956(02)00006-0.
- 2 János Barát and David R. Wood. Notes on nonrepetitive graph colouring. *Electron. J. Comb.*, 15(1), 2008. URL: http://www.combinatorics.org/Volume_15/Abstracts/v15i1r99.html.
- 3 Hans L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM J. Comput.*, 25(6):1305–1317, 1996. doi:10.1137/S0097539793251219.
- 4 Hans L. Bodlaender. A note on domino treewidth. *Discret. Math. Theor. Comput. Sci.*, 3(4):141–150, 1999. URL: <http://dmtcs.episciences.org/256>.
- 5 Hans L. Bodlaender, Gunther Cornelissen, and Marieke van der Wegen. Problems hard for treewidth but easy for stable gonality. *arXiv*, abs/2202.06838, 2022. Extended abstract to appear in Proceedings WG 2022. arXiv:2202.06838.
- 6 Hans L. Bodlaender, Pål Grønås Drange, Markus S. Dregi, Fedor V. Fomin, Daniel Lokshtanov, and Michał Pilipczuk. A $c^k n$ 5-approximation algorithm for treewidth. *SIAM J. Comput.*, 45(2):317–378, 2016. doi:10.1137/130947374.
- 7 Hans L. Bodlaender and Joost Engelfriet. Domino treewidth. *J. Algorithms*, 24(1):94–123, 1997. doi:10.1006/jagm.1996.0854.
- 8 Hans L. Bodlaender, Carla Groenland, and Hugo Jacob. On the parameterized complexity of computing tree-partitions. *arXiv*, abs/2206.11832, 2022. URL: <https://arxiv.org/abs/2206.11832>.
- 9 Hans L. Bodlaender, Carla Groenland, Hugo Jacob, Marcin Pilipczuk, and Michał Pilipczuk. On the complexity of problems on tree-structured graphs. *CoRR*, 2022. doi:10.48550/ARXIV.2206.11828.
- 10 Hans L. Bodlaender and Torben Hagerup. Parallel algorithms with optimal speedup for bounded treewidth. *SIAM J. Comput.*, 27(6):1725–1746, 1998. doi:10.1137/S0097539795289859.
- 11 Paz Carmi, Vida Dujmovic, Pat Morin, and David R. Wood. Distinct distances in graph drawings. *Electron. J. Comb.*, 15(1), 2008. URL: http://www.combinatorics.org/Volume_15/Abstracts/v15i1r107.html.
- 12 Julia Chuzhoy and Zihan Tan. Towards tight(er) bounds for the excluded grid theorem. *J. Comb. Theory, Ser. B*, 146:219–265, 2021. doi:10.1016/j.jctb.2020.09.010.
- 13 Bruno Courcelle. The monadic second-order logic of graphs. I. Recognizable sets of finite graphs. *Inf. Comput.*, 85(1):12–75, 1990. doi:10.1016/0890-5401(90)90043-H.
- 14 Guoli Ding and Bogdan Oporowski. Some results on tree decomposition of graphs. *J. Graph Theory*, 20(4):481–499, 1995. doi:10.1002/jgt.3190200412.
- 15 Guoli Ding and Bogdan Oporowski. On tree-partitions of graphs. *Discret. Math.*, 149(1-3):45–58, 1996. doi:10.1016/0012-365X(94)00337-I.
- 16 Vida Dujmovic, Pat Morin, and David R. Wood. Layout of graphs with bounded tree-width. *SIAM J. Comput.*, 34(3):553–579, 2005. doi:10.1137/S0097539702416141.
- 17 Vida Dujmovic, Matthew Suderman, and David R. Wood. Graph drawings with few slopes. *Comput. Geom.*, 38(3):181–193, 2007. doi:10.1016/j.comgeo.2006.08.002.
- 18 Uriel Feige, MohammadTaghi Hajiaghayi, and James R. Lee. Improved approximation algorithms for minimum weight vertex separators. *SIAM J. Comput.*, 38(2):629–657, 2008. doi:10.1137/05064299X.
- 19 Fedor V. Fomin, Daniel Lokshtanov, Saket Saurabh, Michał Pilipczuk, and Marcin Wrochna. Fully polynomial-time parameterized computations for graphs and matrices of low treewidth. *ACM Trans. Algorithms*, 14(3):34:1–34:45, 2018. doi:10.1145/3186898.
- 20 L. R. Ford and D. R. Fulkerson. Maximal flow through a network. *Canadian Journal of Mathematics*, 8:399–404, 1956. doi:10.4153/CJM-1956-045-5.

- 21 Robert Ganian, Eun Jung Kim, and Stefan Szeider. Algorithmic applications of tree-cut width. In Giuseppe F. Italiano, Giovanni Pighizzini, and Donald Sannella, editors, *Mathematical Foundations of Computer Science 2015 - 40th International Symposium, MFCS 2015, Milan, Italy, August 24-28, 2015, Proceedings, Part II*, volume 9235 of *Lecture Notes in Computer Science*, pages 348–360. Springer, 2015. doi:10.1007/978-3-662-48054-0_29.
- 22 Emilio Di Giacomo, Giuseppe Liotta, and Henk Meijer. Computing straight-line 3d grid drawings of graphs in linear volume. *Comput. Geom.*, 32(1):26–58, 2005. doi:10.1016/j.comgeo.2004.11.003.
- 23 Rudolf Halin. Tree-partitions of infinite graphs. *Discret. Math.*, 97(1-3):203–217, 1991. doi:10.1016/0012-365X(91)90436-6.
- 24 John Hopcroft and Robert Tarjan. Algorithm 447: Efficient algorithms for graph manipulation. *Commun. ACM*, 16(6):372–378, June 1973. doi:10.1145/362248.362272.
- 25 Tuukka Korhonen. A single-exponential time 2-approximation algorithm for treewidth. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*, pages 184–192. IEEE, 2021. doi:10.1109/FOCS52979.2021.00026.
- 26 Omer Reingold. Undirected connectivity in log-space. *J. ACM*, 55(4), 2008. doi:10.1145/1391289.1391291.
- 27 Detlef Seese. Tree-partite graphs and the complexity of algorithms. In Lothar Budach, editor, *5th International Conference on Fundamentals of Computation Theory, FCT 1985*, volume 199 of *Lecture Notes in Computer Science*, pages 412–421. Springer, 1985. doi:10.1007/BFb0028825.
- 28 Paul Wollan. The structure of graphs not admitting a fixed immersion. *J. Comb. Theory, Ser. B*, 110:47–66, 2015. doi:10.1016/j.jctb.2014.07.003.
- 29 David R. Wood. On tree-partition-width. *Eur. J. Comb.*, 30(5):1245–1253, 2009. doi:10.1016/j.ejc.2008.11.010.
- 30 David R. Wood and Jan Arne Telle. Planar decompositions and the crossing number of graphs with an excluded minor. In Michael Kaufmann and Dorothea Wagner, editors, *Graph Drawing, 14th International Symposium, GD 2006, Karlsruhe, Germany, September 18-20, 2006. Revised Papers*, volume 4372 of *Lecture Notes in Computer Science*, pages 150–161. Springer, 2006. doi:10.1007/978-3-540-70904-6_16.

A Tree-cut width and the stability of tree-partition-width

In this section, we consider the relation of the notion of tree-cut width with (stable) tree-partition-width. Tree-cut width was introduced by Wollan [28]. Ganian et al. [21] showed that several problems that are $W[1]$ -hard with treewidth as parameter are fixed parameter tractable with tree-cut width as parameter.

We begin by defining the *tree-cut width* of a graph $G = (V, E)$. A tree-cut decomposition (T, \mathcal{X}) consists of a rooted tree T and a family \mathcal{X} of *bags* $(X_i)_{i \in V(T)}$ which form a near partition of $V(G)$ (i.e. some bags may be empty, but nonempty bags form a partition of $V(G)$). For $t \in V(T)$, we denote by $e(t)$ the edge of T incident to t and its parent. For $e \in E(T)$, let T_1, T_2 denote the two connected components of $T - e$. We denote by $\text{cut}(e)$ the set of edges with an endpoint in both of $\bigcup_{i \in V(T_1)} X_i$ and $\bigcup_{i \in V(T_2)} X_i$. The adhesion of $t \in V(T)$ is $\text{adh}(t) = |\text{cut}(e(t))|$, and its torso-size is $\text{tor}(t) = |X_t| + b_t$ where b_t is the number of edges $e \in E(T)$ incident to t such that $|\text{cut}(e)| \geq 3$. The width of the decomposition is then $\max_{t \in V(T)} \{\text{adh}(t), \text{tor}(t)\}$. Note that edges are allowed to go between vertices that are not in the same bag. The tree-cut width of a graph is the minimal width of tree-cut decomposition. When $|\text{cut}(e)| \geq 3$, the edge e is called *bold*, and otherwise, e is called *thin*. When $\text{adh}(t) \leq 2$, node t is called *thin*, otherwise it is called *bold*. In [21], it is shown that a tree-cut decomposition can be assumed to be *nice*, meaning that if $t \in V(T)$ is thin then $N(Y_t) \cap \left(\bigcup_{b \text{ sibling of } t} Y_b \right) = \emptyset$, where Y_i is the union of X_j for j in the subtree of i .

Wollan shows that having bounded tree-cut width is equivalent to only having wall immersions of bounded size.

► **Observation 18.** $K_{3,n-3}$ has tree-partition-width at most 3 but unbounded tree-cut width.

Indeed, note that any graph on $n - 3$ vertices with maximum degree 3 can be immersed in $K_{3,n-3}$. In particular, this works for any wall. The lower bound given by Wollan shows that the tree-cut width of $K_{3,n-3}$ is $\Omega(n^{\frac{1}{4}})$.

We denote by $\underline{\text{tpw}}(G)$ the minimum tree-partition-width over subdivisions of G , and by $\overline{\text{tpw}}(G)$ the maximum tree-partition-width of subdivisions of G . We will show that both are polynomially tied to the tree-partition-width of G , which proves useful in polynomially bounding tree-partition-width by tree-cut width due to the following lemma.

► **Lemma 19.** $\underline{\text{tpw}} = O(\text{tcw}^2)$.

Proof. Consider a nice tree-cut decomposition (T, \mathcal{X}) of a graph G of width k . We will construct a tree-partition for a subdivision of G . Note that the bags are already disjoint, but some edges are not between neighbouring bags of T .

Each edge uv of G is subdivided $d_T(u, v)$ times, which is the distance between the nodes containing u and v respectively in their bag (recall that the bags form a near partition). We then add the vertices of the subdivided edge in the bags on the path in the decomposition between the bags containing their endpoints. This is sufficient to make the decomposition a tree-partition T' of a subdivision of G .

We now argue that T' has a width of $O(k^2)$. A bag Y_t of T' contains at most:

- k initial vertices
- $\max(2, k)$ vertices from subdivisions of edges in $\text{cut}(e(t))$ accounting for edges going from a child of t to an ancestor of t
- $k^2/2$ vertices from edges that are between bold children of t . For u, v children of T , there are only edges between T_u and T_v if both are bold. There are also at most k such edges incident to T_u for any child u of T , and we may divide by 2 since each edge will be counted twice this way. We stress that thin children do not contribute because the tree-cut decomposition is nice.

Hence, $\underline{\text{tpw}}(G) \leq 2 + k(k + 2)/2 + k = O(\text{tcw}(G)^2)$ ◀

Next, we consider the parameters tpw and $\overline{\text{tpw}}$.

► **Lemma 20.** $\text{tpw} \leq \overline{\text{tpw}} \leq \text{tpw}(\text{tpw} + 1)$

Proof. The lower bound is immediate. We prove the upper bound.

Consider a graph G with a tree-partition (T, \mathcal{X}) of width k , and a subdivision G' of G . We construct a tree-partition (T', \mathcal{X}') of G' of width at most $k(k + 1)$.

We root the decomposition T arbitrarily.

Suppose that u, v are in the same bag of T and the edge uv was subdivided to form the path u, a_1, \dots, a_ℓ, v . We add the vertices a_i in new bags containing, $\{a_1, a_\ell\}$, $\{a_2, a_{\ell-1}, \dots$ which corresponds to a new branch of the decomposition of width at most 2.

Consider next the vertices obtained by subdividing an edge uv for u in the child bag of the bag of v . If a subdivided edge was between two vertices of adjacent bags, we order the vertices of the path obtained by subdividing the edge from the vertex in the child bag to the vertex in the parent bag. We add the penultimate vertex to the child bag, and fold the remaining vertices of the path in a fresh branch of the decomposition of width at most 2 similarly to the previous case.

7:18 On the Parameterized Complexity of Computing Tree-Partitions

This gives a tree partition T' . Bags of T' that are not in T have size at most 2, and, to bags of T' that are also in T , we added at most k^2 vertices (at most one per edge between the bag and its parent). We conclude that T' has width at most $k(k+1)$. ◀

A result of Ding and Oporowski [15] shows that tree-partition-width is tied to a parameter γ that is (by design) monotonic with respect to the topological minor relation. We adapt their proof to derive the following stronger result.

► **Theorem 21.** *There exists a parameter γ which is polynomially tied to the tree-partition-width, and is monotonic with respect to the topological minor relation. More precisely, $\mathbf{tpw} = \Omega(\gamma)$ and $\mathbf{tpw} = O(\gamma^{24})$.*

We deduce the following statement.

► **Corollary 22.** *\mathbf{tpw} , $\overline{\mathbf{tpw}}$, and $\underline{\mathbf{tpw}}$ are polynomially tied.*

Proof. Lemma 20 shows that \mathbf{tpw} and $\overline{\mathbf{tpw}}$ are polynomially tied. Note that, by definition, $\underline{\mathbf{tpw}} \leq \mathbf{tpw}$. Then, for a fixed graph G , consider H a subdivision of G achieving $\mathbf{tpw}(H) = \underline{\mathbf{tpw}}(G)$. $\mathbf{tpw}(G)^{O(1)} \leq \gamma(G) \leq \gamma(H) \leq \mathbf{tpw}(H)^{O(1)} = \underline{\mathbf{tpw}}(G)^{O(1)}$. The first and last inequalities come from the fact that γ and \mathbf{tpw} are polynomially tied. The middle inequality is because γ is monotonic with respect to the topological minor relation. ◀

From Lemma 19 and Corollary 22, we deduce the following theorem.

► **Theorem 23.** *The parameter tree-partition-width is polynomially bounded by the parameter tree-cut width. In other words, we show that there exist constants $C, c > 0$ such that for any graph G , $\mathbf{tpw}(G) \leq C \mathbf{tcw}(G)^c$.*

We now turn our focus to the technical proof of Theorem 21. We define the m -grid as the graph on the vertex set $[m] \times [m]$ with edges $(i, j)(i', j')$ when $|i - i'| + |j - j'| \leq 1$. We then define the m -wall as the graph obtained from the m -grid by removing edges $(i, j)(i + 1, j)$ for $i + j$ even. The *wall number* of a graph G is then defined as the largest m such that G contains the m -wall as a (topological)² minor, and the *grid number* of G is the largest m such that G contains the m -grid as a minor. We denote them by $\mathbf{wn}(G)$ and $\mathbf{gn}(G)$ respectively.

► **Observation 24.** *The wall number and the grid number are linearly tied: $\mathbf{wn}(G) = \Theta(\mathbf{gn}(G))$.*

We use the following result of Chuzhoy and Tan [12] (the bound is weakened to have a lighter formula).

► **Lemma 25** (Chuzhoy and Tan [12]). *The treewidth is polynomially tied to the grid number: $\mathbf{tw} = \Omega(\mathbf{gn})$ and $\mathbf{tw} = O(\mathbf{gn}^{10})$.*

Hence, the treewidth is polynomially tied to the wall number: $\mathbf{tw} = \Omega(\mathbf{wn})$ and $\mathbf{tw} = O(\mathbf{wn}^{10})$.

We call m -fan the graph that consists of a path of order m with an additional vertex adjacent to all of the vertices of the path. We call m -branching-fans the graphs that consist of a tree T and a vertex v adjacent to a subset N of the vertices of T containing at least the leaves, such that m is the minimum size of a subset of vertices X of T such that each component of $T - X$ contains at most m vertices of N . In particular, the $(m+1)^2$ -fan is an

² The notions of minor and topological minor coincide for graphs of maximum degree at most 3.

$(m + 1)$ -branching-fan. We call m -multiple of a tree of order m a graph obtained from a tree of order m after replacing its edges by m parallel edges and then subdividing each edge once to keep the graph simple.

Let $\gamma_1(G)$ be the largest m such that G contains an m -branching-fan as a topological minor. Let $\gamma_2(G)$ be the largest m such that G contains an m -multiple of a tree of order m as a topological minor.

Let $\gamma(G)$ be the maximum of $\mathbf{wn}(G)$, $\gamma_1(G)$, and $\gamma_2(G)$.

▷ **Claim 26.** The parameter γ is monotonic with respect to the topological minor relation.

Proof. Let G be a graph and H be a topological minor of G . Any topological minor of H is also a topological minor of G , hence $\mathbf{wn}(G) \geq \mathbf{wn}(H)$, $\gamma_1(G) \geq \gamma_1(H)$, $\gamma_2(G) \geq \gamma_2(H)$. We conclude that $\gamma(G) \geq \gamma(H)$. ◁

► **Observation 27.** *The m -branching-fans, the m -multiples of trees of order m and the m -wall have tree-partition-width $\Omega(m)$. Hence, we have $\mathbf{tpw}(G) = \Omega(\gamma(G))$.*

We fix a graph G and let $m = \gamma(G)$. Note that $m \geq \gamma_2(G) > 0$.

We denote by G^b the graph on the vertex set of G , where xy is an edge if and only if there are at least b vertex disjoint paths from x to y . We now consider G^b for $b = \Omega(m^{10})$.

▷ **Claim 28.** The connected components of G^b have size at most m .

Proof. We proceed by contradiction, and assume there is a connected component C of size at least $m + 1$.

Since C is connected, it contains a spanning tree T . We number its edges e_1, \dots, e_ℓ such that every prefix induces a connected subtree of T . We construct a subgraph H of G that should be an $(m + 1)$ -multiple of a tree of order $m + 1$, contradicting the definition of m . For each edge uv , in order, we try to add to H $m + 1$ vertex disjoint paths from u to v that avoid vertices of C and the vertices already in H . If we manage to do this for at most m edges, then we have placed at most $m(m + 1)$ paths. Let uv be the first edge for which we could not find $m + 1$ vertex disjoint paths that do not intersect previous vertices (except for u or v). By definition of G^b , there are b vertex disjoint u, v -paths in G , we denote the set of such paths by π . At most m of the paths of π hit vertices of C already in H . Then, since at least $b - m$ are hit by previous paths and there are at most $m(m + 1)$ previous paths. By the pigeon hole principle, one of the previous paths P_0 must hit $\frac{b-m}{m(m+1)} \geq (m + 1)^2$ paths in π . By considering P_0 and the paths it hits in π , we easily obtain a subdivision of an $(m + 1)^2$ -fan. This is a contradiction with the definition of m . Hence, we must have been able to process m edges. Which means we obtained a subdivision of an $(m + 1)$ -multiple of a tree of order $m + 1$. This is a contradiction to the definition of m . We conclude that the connected component must have size at most m . ◁

Let H be the quotient of G by the connected components of G^b . We call it the b -reduction of G .

▷ **Claim 29.** The blocks of H have maximum degree at most bm^3 .

Proof. Assume by contradiction that the maximum degree is more than bm^3 . Let B be a block of H , and X be one of its vertices of maximum degree. X contains at most m vertices of G by Claim 28. The vertices of G in $B - X$ must be in the same connected component C of G since $b > m$. There are at least $bm^3 + 1$ edges between X and C . By the pigeon hole principle, one vertex v of X must have at least $bm^2 + 1$ neighbours in C . Consider a

7:20 On the Parameterized Complexity of Computing Tree-Partitions

spanning tree T of C . We iteratively remove leaves that are not neighbours of v , and then replace any vertex of degree 2 that is not a neighbour of v by an edge between its neighbours. We denote this reduced tree by T' .

First, note that the degree in T' is bounded by $b - 1$ because incident edges can be extended to vertex disjoint paths to leaves of T' which are neighbours of v by construction. We now use the fact that G contains no $(m + 1)$ -branching-fans as topological minors. In particular, there must be a set U of vertices of T' of size at most m such that components of $T' - U$ contain at most m neighbours of v . By removing at most m vertices of degree at most $(b - 1)$, we have at most $1 + (b - 2)m$ components in $T' - U$ meaning v has degree bounded by $(b - 1)m^2$. We found our contradiction. \triangleleft

\triangleright **Claim 30.** The treewidth of H is at most $O(b)$.

Proof. We first apply Lemma 25 to bound the treewidth of G by $O(b)$. Consider a tree decomposition of G of adequate width $\Theta(b)$, and replace each bag by the components of G^b that intersect it. By Claim 4, this is a decomposition of H . \triangleleft

Using Claim 29 and Claim 30 and the construction of Wood as we did in the approximation algorithm, we obtain a tree-partition of H of width $O(b^2m^3)$. We then replace components of G^b by their vertices, obtaining a tree-partition of G of width $O(b^2m^4)$ due to Claim 28.

We have obtained a tree-partition of width $O(b^2m^4) = O(m^{24})$. This concludes the proof of Theorem 21.