

From Width-Based Model Checking to Width-Based Automated Theorem Proving

Mateus de Oliveira Oliveira and Farhad Vadi e

Department of Informatics - University of Bergen - Norway

`{mateus.oliveira,farhad.vadiee}@uib.no`

Abstract

In the field of parameterized complexity theory, the study of graph width measures has been intimately connected with the development of width-based model checking algorithms for combinatorial properties on graphs. In this work, we introduce a general framework to convert a large class of width-based model-checking algorithms into algorithms that can be used to test the validity of graph-theoretic conjectures on classes of graphs of bounded width. Our framework is modular and can be applied with respect to several well-studied width measures for graphs, including treewidth and cliquewidth.

As a quantitative application of our framework, we show that for several long-standing graph-theoretic conjectures, there exists an algorithm that takes a number k as input and correctly determines in time double-exponential in $k^{O(1)}$ whether the conjecture is valid on all graphs of treewidth at most k . This improves significantly on upper bounds obtained using previously available techniques.

Keywords: Dynamic Programming Cores, Width Measures, Automated Theorem Proving

1 Introduction

1.1 Motivation

When mathematicians are not able to solve a conjecture about a given class of mathematical objects, it is natural to try to test the validity of the conjecture on a smaller, or better behaved class of objects. In the realm of graph theory, a common approach is to try analyze the conjecture on restricted classes of graphs, defined by fixing some structural parameter. In this work, we push forward this approach from a computational perspective by focusing on parameters derived from graph width measures. Prominent examples of such parameters are the *treewidth* of a graph, which intuitively quantifies how much a graph is similar to a tree [70, 10, 42] and the cliquewidth of a graph, which intuitively quantifies how much a graph is similar to a clique [32]. More specifically, we are concerned with the following problem.

Problem 1 (Width-Based ATP). *Given a graph property \mathbb{P} and a positive integer k , is it the case that every graph of width at most k belongs to \mathbb{P} ?*

Problem 1 provides a width-based approach to the field of automated theorem proving (ATP). For instance, consider Tutte’s celebrated 5-flow conjecture [77], which states that every bridgeless graph has a nowhere-zero 5-flow. Let `HasBridge` be the set of all graphs that have a bridge, and `NZFlow(5)` be the set of all graphs that admit a nowhere-zero 5-flow. Then, proving Tutte’s 5-flow conjecture is equivalent to showing that every graph belongs to the graph property `HasBridge` \vee `NZFlow(5)`. Since Tutte’s conjecture has been unresolved for many decades, one possible approach for gaining understanding about the conjecture is to try to determine, for gradually increasing values of k , whether every graph of width at most k , with respect to some fixed width-measure,

belongs to $\text{HasBridge} \vee \text{NZFlow}(5)$. What makes this kind of question interesting from a proof theoretic point of view is that several important classes of graphs have small width with respect to some width measure. For instance, trees and forests have treewidth at most 1, series-parallel graphs and outerplanar graphs have treewidth at most 2, k -outerplanar graphs have treewidth at most $3k - 1$, co-graphs have cliquewidth at most 2, any distance hereditary graph has cliquewidth 3, etc [11, 14, 21, 17, 16, 47]. Therefore, proving the validity of a given conjecture on classes of graphs of small width corresponds to proving the conjecture on interesting classes of graphs.

1.2 Our Results

In this work, we introduce a general and modular framework that allows one to convert width-based dynamic programming algorithms for the model checking of graph properties into algorithms that can be used to address Problem 1. More specifically, our main contributions are threefold.

1. We start by defining the notions of a *treelike decomposition class* (Definition 2) and of a *treelike width-measure* (Definition 3). These two notions can be used to express several classic, well studied width measures for graphs, such as treewidth [13], pathwidth [54], carving width [76], cutwidth [24, 75], bandwidth [24], cliquewidth [32], etc, and some more recent measures such as ODD width [4].
2. Subsequently, we introduce the notion of a *treelike dynamic programming core* (Definition 8), a formalism for the specification of dynamic programming algorithms operating on treelike decompositions. As stand-alone objects, DP-cores are essentially a formalism for the specification of sets of terms, much like tree automata, but with the exception that transitions are specified implicitly, using functions. Nevertheless, when associated with the notion of a treelike decomposition class, DP-cores can be used to define classes of graphs. Furthermore, when satisfying a property called *coherency* (Definition 11), treelike DP-cores can be safely used for the purpose of model checking. Intuitively, coherency is a condition that requires that a DP-core gives the same answer when processing any two treelike decompositions of the same graph. Finally, our formalism is symbolic, in the sense that graphs are encoded as terms over a finite alphabet. This makes our approach suitable for the consideration of questions pertaining to the realm of automated theorem proving, as described next.
3. Intuitively, our main result (Theorem 33) states that if a graph property \mathbb{P} is a dynamic combination (see Definition 30) of graph properties $\mathbb{P}_1, \dots, \mathbb{P}_\ell$ that can be decided by coherent DP-cores D_1, D_2, \dots, D_ℓ satisfying certain finiteness conditions, then the process of determining whether every graph of width at most k belongs to \mathbb{P} can be decided roughly¹ in time

$$2^{O(\beta(k) \cdot \mu(k))} \leq 2^{2^{O(\beta(k))}},$$

where $\mu(k)$ and $\beta(k)$ are respectively the maximum multiplicity and the maximum bitlength of a DP-core from the list D_1, \dots, D_ℓ (see Definition 14). Additionally, if a counterexample of width at most k exists, then a term of height at most $2^{O(\beta(k) \cdot \mu(k))}$ representing such a counterexample can be constructed (Corollary 28).

The modularity of our approach makes it highly suitable to be applied in the context of automated theorem proving. For instance, when specialized to the context of treewidth, our approach can be used to infer that several long-standing conjectures in graph theory can be tested on the class of graphs of treewidth at most k in time double exponential in $k^{O(1)}$. Examples of such conjectures include Hadwiger conjecture [41], Tutte's flow conjectures [77], Barnette's conjecture [78], and many others (Section 6).

¹The precise statement of Theorem 33 involves other parameters that are negligible in typical applications.

1.3 Related Work

General automata-theoretic frameworks for the development of dynamic programming algorithms have been introduced under a wide variety of contexts [40, 64, 65, 56, 66, 60, 57, 6, 7, 8]. In most of these contexts, automata are used to encode the space of solutions of combinatorial problems when a graph G is given at the input. For instance, given a tree decomposition of width k of a graph G , one can construct a tree automaton representing the set of proper 3-colorings of G [7].

In our framework, treelike DP-cores are used to represent families of graphs satisfying a given property. For instance, one can define a treelike DP-core D , where for each $k \in \mathbb{N}$, $D[k]$ is a finite representation of the set of all graphs of treewidth at most k that are 3-colorable. In our context it is essential that graphs of width k are encoded as terms over a finite alphabet whose size only depends on k . We note that the idea of representing families of graphs as tree languages over a finite alphabets has been used in a wide variety of contexts [18, 2, 31, 37, 39, 30]. Nevertheless, the formalisms arising in these contexts are usually designed to be compatible with logical algorithmic meta-theorems, and for this reason, tree automata are meant to be compiled from logical specifications, rather than to be programmed. In contrast, we provide a framework that allows one to easily program state-of-the-art dynamic programming algorithms operating on treelike decompositions (see Section 5.3), and to safely combine these algorithms (just like plugins) for the purpose of width-based automated theorem proving.

The monadic second-order logic of graphs (MSO_2 logic) extends first-order logic by introducing quantification over sets of vertices and over sets of edges. This logic is powerful enough to express several well studied graph-theoretic properties such as connectivity, Hamiltonicity, 3-colorability, and many others. Given that for each $k \in \mathbb{N}$, the MSO_2 theory of graphs of treewidth at most k is decidable [73, 28], we have that if a graph property \mathbb{P} is definable in MSO_2 logic, then there is an algorithm that takes an integer k as input, and correctly determines whether every graph of treewidth at most k belongs to \mathbb{P} . A similar result can be proved with respect to graphs of constant cliquewidth using MSO_1 logic [73], and for graphs of bounded ODD width using FO logic [4].

One issue with addressing Problem 1 using the logic approach mentioned above is that algorithms obtained in this way are usually based quantifier-elimination. As a consequence, the function upper-bounding the running time of these algorithms in terms of the width parameter grows as a tower of exponentials whose height depends on the number of quantifier alternations of the logical sentence given as input to the algorithm. For instance, the algorithm obtained using this approach to test the validity of Hadwiger’s conjecture restricted to c colors on graphs of treewidth at most k has a very large dependency on the width parameter. In [51], the time necessary to perform this task was estimated in $f(c, k) \leq p^{p^p}$, where $p = (k + 1)^{(c-1)}$ [51]. Our approach yields a much smaller upper bound of the form $2^{2^{O(k \cdot \log k + c^2)}}$. Significant reductions in complexity can also be observed for other conjectures.

Courcelle’s model checking theorem and its subsequent variants [29, 5, 20] have had a significant impact in the development of width-based model checking algorithms. Indeed, once a graph property has been shown to be decidable in FPT time using the machinery surrounding Courcelle’s theorem, the next relevant question is how small can the dependency in the width parameter be. Algorithms with optimal dependency on the width parameter have been obtained for a large number of graph properties [58, 67] under standard complexity theoretic assumptions, such as the exponential time hypothesis (ETH) [46, 45] and related conjectures [22, 23]. It is worth noting that in many cases, the development of such optimal algorithms requires the use of advanced techniques borrowed from diverse subfields of mathematics, such as structural graph theory [71, 9], algebra [79, 15, 61], combinatorics [68, 59], etc. Our framework allows one to incorporate several of these techniques in the development of faster algorithms for width-based automated theorem proving.

2 Preliminaries

We let \mathbb{N} denote the set of natural numbers and \mathbb{N}_+ denote the set of positive natural numbers. We let $[0] \doteq \emptyset$, and for each $n \in \mathbb{N}_+$, we define $[n] \doteq \{1, \dots, n\}$. Given a set S , the set of finite subsets of S is denoted by $\mathcal{P}_{\text{fin}}(S)$.

In this work, a *graph* is a triple $G = (V, E, \rho)$ where $V \subseteq \mathbb{N}$ is a *finite* set of *vertices*, $E \subseteq \mathbb{N}$ is a finite set *edges*, and $\rho \subseteq E \times V$ is an *incidence relation*. For each edge $e \in E$, we let $\text{endpts}(e) = \{v \in V : (e, v) \in \rho\}$ be the set of vertices incident with e . In what follows, we may write V_G , E_G and ρ_G to denote the sets V , E and ρ respectively. We let $|G| = |V_G| + |E_G|$ be the *size* of G . We let GRAPHS denote the set of all graphs. For us, the *empty graph* is the graph $(\emptyset, \emptyset, \emptyset)$ with no vertices, no edges, and no incidence pairs.

An *isomorphism* from a graph G to a graph H is a pair $\phi = (\phi_1, \phi_2)$ where $\phi_1 : V_G \rightarrow V_H$ is a bijection from the vertices of G to the vertices of H and $\phi_2 : E_G \rightarrow E_H$ is a bijection from the edges of G to the edges of H with the property that for each vertex $v \in V_G$ and each edge $e \in E_G$, $(v, e) \in \rho_G$ if and only if $(\phi_1(v), \phi_2(e)) \in \rho_H$. If such a bijection exists, we say that G and H are *isomorphic*, and denote this fact by $G \sim H$.

A *graph property* is any subset $\mathbb{P} \subseteq \text{GRAPHS}$ closed under isomorphisms. That is to say, for each two isomorphic graphs G and H in GRAPHS , $G \in \mathbb{P}$ if and only if $H \in \mathbb{P}$. Note that the sets \emptyset and GRAPHS are graph properties. Given a set S of graphs, the *isomorphism closure* of S is defined as the set $\text{ISO}(S) = \{G \in \text{GRAPHS} : \exists H \in S, G \sim H\}$.

Given a graph property \mathbb{P} , a \mathbb{P} -*invariant* is a function $\mathcal{I} : \mathbb{P} \rightarrow S$, for some set S , that is invariant under graph isomorphisms. More precisely, $\mathcal{I}(G) = \mathcal{I}(H)$ for each two isomorphic graphs G and H in \mathbb{P} . If $\mathbb{P} = \text{GRAPHS}$, we may say that \mathcal{I} is simply a *graph invariant*. For instance, chromatic number, clique number, dominating number, etc., as well as width measures such as treewidth and cliquewidth, are all graph invariants. In this work the set S will be typically \mathbb{N} , when considering width measures, or $\{0, 1\}^*$ when considering other invariants encoded in binary. In order to avoid confusion we may use the letter \mathcal{M} to denote invariants corresponding to width measures, and the letter \mathcal{I} to denote general invariants.

A *ranked alphabet* is a finite non-empty set Σ together with function $\mathfrak{r} : \Sigma \rightarrow \mathbb{N}$, which intuitively specifies the arity of each symbol in Σ . The arity of r is the maximum arity of a symbol in Σ . A term over Σ is a pair $\tau = (T, \lambda)$ where T is a rooted tree, where the children of each node are ordered from left to right, and $\lambda : \text{Nodes}(T) \rightarrow \Sigma$ is a function that labels each node p in $\text{Nodes}(T)$ with a symbol from Σ of arity $|\text{Children}(p)|$, i.e., the number of children of p . In particular, leaf nodes are labeled with symbols of arity 0. We may write $\text{Nodes}(\tau)$ to refer to $\text{Nodes}(T)$. We write $|\tau|$ to denote $|\text{Nodes}(T)|$. The height of τ is defined as the height of T . We denote by $\text{Terms}(\Sigma)$ the set of all terms over Σ . If $\tau_1 = (T_1, \lambda_1), \dots, \tau_r = (T_r, \lambda_r)$ are terms in $\text{Terms}(\Sigma)$, and $a \in \Sigma$ is a symbol of arity r , then we let $a(\tau_1, \dots, \tau_r)$ denote the term $\tau = (T, \lambda)$ where $\text{Nodes}(T) = \{u\} \cup \text{Nodes}(T_1) \cup \dots \cup \text{Nodes}(T_r)$ for some fresh node u , $\text{root}(T) = u$, $\lambda(u) = a$, and $\lambda|_{\text{Nodes}(T_j)} = \lambda_j$ for each $j \in [r]$. A tree automaton is a tuple $\mathcal{A} = (\Sigma, Q, F, \Delta)$ where Σ is a ranked alphabet, Q is a finite set of states, F is a final set of states, and Δ is a set of transitions (i.e. rewriting rules) of the form $a(q_1, \dots, q_r) \rightarrow q$, where a is a symbol of arity r , and q_1, \dots, q_r, q are states in Q . A term τ is accepted by \mathcal{A} if it can be rewritten into a final state in F by transitions in Δ . The language of \mathcal{A} , denoted by $\mathcal{L}(\mathcal{A})$, is the set of all terms accepted by \mathcal{A} . A tree language $L \subseteq \text{Terms}(\Sigma)$ is said to be *regular* if there is a tree automaton \mathcal{A} over Σ such that $L = \mathcal{L}(\mathcal{A})$. A tree automaton $\mathcal{A} = (\Sigma, Q, F, \Delta)$ is said to be *deterministic* if for each symbol $a \in \Sigma$ of arity r , and each r -tuple of states (q_1, \dots, q_r) , there is at most one state q such that $a(q_1, \dots, q_r) \rightarrow q$ is a transition of \mathcal{A} . We refer to [27] for basic concepts on tree automata theory.

Let \mathbb{P} be a graph property, and G be a graph. We let $\mathbb{P}(G)$ be the Boolean value 1 if $G \in \mathbb{P}$ and the Boolean value 0 otherwise. Given graph properties $\mathbb{P}_1, \dots, \mathbb{P}_\ell$, and a Boolean function $\mathcal{C} : \{0, 1\}^\ell \rightarrow \{0, 1\}$, we let $\hat{\mathcal{C}}(\mathbb{P}_1, \dots, \mathbb{P}_\ell) = \{G : \mathcal{C}(\mathbb{P}_1(G), \dots, \mathbb{P}_\ell(G)) = 1\}$ be the Boolean \mathcal{C} -combination of $\mathbb{P}_1, \dots, \mathbb{P}_\ell$. For properties \mathbb{P}_1 and \mathbb{P}_2 we may write simply $\neg \mathbb{P}_1$ for the complement of \mathbb{P}_1 ; $\mathbb{P}_1 \wedge \mathbb{P}_2$ for the intersection of \mathbb{P}_1 and \mathbb{P}_2 ; $\mathbb{P}_1 \vee \mathbb{P}_2$ for the union of \mathbb{P}_1 and \mathbb{P}_2 ; and $\mathbb{P}_1 \rightarrow \mathbb{P}_2$

for the graph property $\neg\mathbb{P}_1 \vee \mathbb{P}_2$. We say that \mathbb{P} is a Boolean combination of graph properties $\mathbb{P}_1, \dots, \mathbb{P}_r$ if there is a function $\mathcal{C} : \{0, 1\}^r \rightarrow \{0, 1\}$ such that $\mathbb{P} = \hat{\mathcal{C}}(\mathbb{P}_1, \dots, \mathbb{P}_r)$. In Section 5.6 we will define a more general notion of combination of graph properties and graph invariants.

Let $k \in \mathbb{N}$. A k -boundaried graph is a pair (G, θ) where G is a graph and $\theta : B \rightarrow V_G$ is an injective map from some subset $B \subseteq [k + 1]$ to the vertex set of G . Given k -boundaried graphs (G_1, θ_1) and (G_2, θ_2) with $\text{Dom}(\theta_1) = \text{Dom}(\theta_2)$, we let $(G_1, \theta_1) \oplus (G_2, \theta_2)$ be the k -boundaried graph (G, θ) where $\theta = \theta_1$ and G is the graph obtained from G_1 and G_2 by identifying, for each $u \in B$, the vertex $\theta_1(u)$ of G_1 with the vertex $\theta_2(u)$ of G_2 . More precisely, let $a = \max\{x : x \in V_{G_1}\}$ and $b = \max\{z : z \in E_{G_1}\}$.

1. $V_G = V_{G_1} \cup \{x + a : x \in V_{G_2} \setminus \theta_2(B)\}$,
2. $E_G = E_{G_1} \cup \{e + b : e \in E_{G_2}\}$,
3. $\rho_G = \rho_{G_1} \cup \{(e + b, \theta_1(u)) : (e, \theta_2(u)) \in \rho_2\} \cup \{(e + b, x + a) : (e, x) \in \rho_1, x \in V_{G_1} \setminus \text{Im}(\theta_2)\}$.

3 Treelike Width Measures

In this section, we introduce the notion of a *treelike width measure*. Subsequently, we show that prominent width measures such as treewidth and cliquewidth fulfil the conditions of our definition. We start by introducing the notion of a *treelike decomposition class*.

Definition 2. Let $r \in \mathbb{N}$. A *treelike decomposition-class of arity r* is a sequence

$$\mathbf{C} = \{(\Sigma_k, \mathbf{L}_k, \mathcal{G}_k)\}_{k \in \mathbb{N}},$$

where for each $k \in \mathbb{N}$, Σ_k is a ranked alphabet of arity at most r , \mathbf{L}_k is a regular tree language over Σ_k , and $\mathcal{G}_k : \mathbf{L}_k \rightarrow \text{GRAPHS}$ is a function that assigns a graph $\mathcal{G}_k(\tau)$ to each $\tau \in \mathbf{L}_k$. Additionally, we require that for each $k \in \mathbb{N}$, $\Sigma_k \subseteq \Sigma_{k+1}$, $\mathbf{L}_k \subseteq \mathbf{L}_{k+1}$, and $\mathcal{G}_{k+1}|_{\mathbf{L}_k} = \mathcal{G}_k$.

Terms in the set $\mathbf{L}(\mathbf{C}) = \bigcup_{k \in \mathbb{N}} \mathbf{L}_k$ are called *\mathbf{C} -decompositions*. For each such a term τ , we may write simply $\mathcal{G}(\tau)$ to denote $\mathcal{G}_k(\tau)$. The *\mathbf{C} -width* of \mathbf{C} -decomposition τ , denoted by $w_{\mathbf{C}}(\tau)$, is the minimum k such that $\tau \in \mathbf{L}_k$. The *\mathbf{C} -width* of a graph G , denoted by $w_{\mathbf{C}}(G)$, is the minimum \mathbf{C} -width of a \mathbf{C} -decomposition τ with $\mathcal{G}(\tau) \simeq G$. We let $w_{\mathbf{C}}(G) = \infty$ if no such minimum k exists.

For each $k \in \mathbb{N}$, we may write $\mathbf{C}_k = (\Sigma_k, \mathbf{L}_k, \mathcal{G}_k)$ to denote the k -th triple in \mathbf{C} . The *graph property defined by \mathbf{C}_k* is the set $\mathbb{G}[\mathbf{C}_k] = \text{ISO}(\{\mathcal{G}(\tau) : \tau \in \mathbf{L}_k\})$. Note that every graph in $\mathbb{G}[\mathbf{C}_k]$ has \mathbf{C} -width at most k , and that $\mathbb{G}[\mathbf{C}_k] \subseteq \mathbb{G}[\mathbf{C}_{k+1}]$. We let $\mathbb{G}[\mathbf{C}] = \bigcup_{k \in \mathbb{N}} \mathbb{G}[\mathbf{C}_k]$ be the graph property defined by \mathbf{C} . We note that the \mathbf{C} -width of any graph in $\mathbb{G}[\mathbf{C}]$ is finite.

Definition 3 (Treelike Width Measure). Let \mathbb{P} be a graph property and $\mathcal{M} : \mathbb{P} \rightarrow \mathbb{N}$ be a \mathbb{P} -invariant. We say that \mathcal{M} is a *treelike width measure* if there is a treelike decomposition-class \mathbf{C} such that $\mathbb{P} = \mathbb{G}[\mathbf{C}]$, and for each graph $G \in \mathbb{P}$, $w_{\mathbf{C}}(G) = \mathcal{M}(G)$. In this case, we say that \mathbf{C} is a realization of \mathcal{M} .

It is worth noting that for any \mathbb{P} -invariant $\mathcal{M} : \mathbb{P} \rightarrow \mathbb{N}$, and any realization \mathbf{C} of \mathcal{M} , we have that $\mathbb{G}[\mathbf{C}_k]$ is the class of all graphs where the value of the invariant is at most k . In the context of our work, treelike width measures are meant to realize invariants corresponding to width measures such as treewidth, cliquewidth, etc. In this case, $\mathbb{P} = \text{GRAPHS}$, since the value of such a measure is defined on each graph. However, our notion of treelike width measure also allows one to capture width measures defined for more restricted classes of graphs, although we will not investigate these measures in this work.

The next theorem states that several well studied width measures for graphs are treelike. The proof of this theorem can be found in Appendix B.

Theorem 4. The width measures treewidth, pathwidth, carving width, cutwidth, and cliquewidth are automatic treelike width measures.

In our results related to width-based automated theorem proving, we will need to take into consideration the time necessary to construct a description of the languages associated with a treelike decomposition class. Let $\mathbf{C} = \{(\Sigma_k, \mathbf{L}_k, \mathcal{G}_k)\}_{k \in \mathbb{N}}$ be a treelike decomposition class of arity r . An automation for \mathbf{C} is a sequence $\mathcal{A} = \{\mathcal{A}_k\}_{k \in \mathbb{N}}$ of tree automata where for each $k \in \mathbb{N}$, $\mathcal{L}(\mathcal{A}_k) = \mathbf{L}_k$. We say that \mathcal{A} has complexity $f : \mathbb{N} \rightarrow \mathbb{N}$ if for each $k \in \mathbb{N}$, \mathcal{A}_k has at most $f(k)$ states, and there is an algorithm \mathfrak{A} that takes a number $k \in \mathbb{N}$ as input, and constructs \mathcal{A}_k in time $k^{O(1)} \cdot f(k)^{O(r)}$.

4 A DP-Friendly Realization of Treewidth

As stated in the proof of Theorem 4, a construction from [36] shows that treewidth fulfills our definition of a treelike width measure. Several more logically-oriented constructions have been considered in the literature [18, 2, 31, 37, 39]. In this section we introduce an alternative realization of treewidth as a treelike width measure. The reason for us to consider this realization is that, at the same time that it allows one to specify graphs of bounded treewidth as graphs over a finite alphabet, our realization is quite compatible with modern techniques for the development of algorithms on graphs of bounded treewidth using traditional tree decompositions. This implies in particular, such algorithms can be converted in our framework without much difficulty.

Definition 5. For each $k \in \mathbb{N}$, we let

$$\Sigma_k = \{\text{Leaf}, \text{IntroVertex}\{u\}, \text{ForgetVertex}\{u\}, \\ \text{IntroEdge}\{u, v\}, \text{Join} : u, v \in [k+1], u \neq v\}.$$

where **Leaf** is a symbol of arity 0, **IntroVertex** $\{u\}$, **ForgetVertex** $\{u\}$ and **IntroEdge** $\{u, v\}$ are symbols of arity 1, and **Join** is a symbol of arity 2. We call Σ_k the k -instructive alphabet.

Intuitively, the elements of Σ_k should be regarded as instructions that can be used to construct graphs inductively. Each such a graph has an associated set $\mathbf{b} \subseteq [k+1]$ of *active labels*. In the base case, the instruction **Leaf** creates an empty graph with an empty set of active labels. Now, let G be a graph with set of active labels \mathbf{b} . For each $u \in [k+1] \setminus \mathbf{b}$, the instruction **IntroVertex** $\{u\}$ adds a new vertex to G , labels this vertex with u , and adds u to \mathbf{b} . For each $u \in \mathbf{b}$, the instruction **ForgetVertex** $\{u\}$ erases the label from the current vertex labeled with u , and removes u from \mathbf{b} . The intuition is that the label u is now free and may be used later in the creation of another vertex. For each $u, v \in \mathbf{b}$, the instruction **IntroEdge** $\{u, v\}$ introduces a new edge between the current vertex labeled with u and the current vertex labeled with v . We note that multiedges are allowed in our graphs. Finally, if G and G' are two graphs, each having \mathbf{b} as the set of active labels, then the instruction **Join** creates a new graph by identifying, for each $u \in \mathbf{b}$, the vertex of G labeled with u with the vertex of G' labeled with u .

A graph constructed according to the process described above can be encoded by a term over the alphabet Σ_k . Not all such terms represent the construction of a graph though. For instance, if at a given step during the construction of a graph, a label u is active, then the next instruction cannot be **IntroVertex** $\{u\}$. We define the set of valid terms ITD_k as the language of a suitable tree automaton \mathcal{A}_k over the alphabet Σ_k . More specifically, we let $\text{ITD}_k = \mathcal{L}(\mathcal{A}_k)$ where $\mathcal{A}_k = (\Sigma_k, Q_k, F_k, \Delta_k)$ is a tree automaton with $Q_k = F_k = \mathcal{P}([k+1])$, and

$$\begin{aligned} \Delta_k = & \{\text{Leaf} \rightarrow \emptyset\} \cup \\ & \{\text{IntroVertex}\{u\}(\mathbf{b}) \rightarrow \mathbf{b} \cup \{u\} \mid \mathbf{b} \subseteq [k+1], u \in [k+1] \setminus \mathbf{b}\} \cup \\ & \{\text{ForgetVertex}\{u\}(\mathbf{b}) \rightarrow \mathbf{b} \setminus \{u\} \mid \mathbf{b} \subseteq [k+1], u \in \mathbf{b}\} \cup \\ & \{\text{IntroEdge}\{u, v\}(\mathbf{b}) \rightarrow \mathbf{b} \mid \mathbf{b} \subseteq [k+1], u, v \in \mathbf{b}, u \neq v\} \cup \\ & \{\text{Join}(\mathbf{b}, \mathbf{b}) \rightarrow \mathbf{b} \mid \mathbf{b} \subseteq [k+1]\}. \end{aligned}$$

Intuitively, states of \mathcal{A}_k are subsets of $[k+1]$ corresponding to subsets of active labels. The set of transitions specify both which instructions can be applied from a given set of active labels \mathbf{b} , and which labels are active after the application of a given instruction.




Figure 1: Left: a 2-instructive tree decomposition τ , and the graph $\mathcal{G}(\tau)$ associated with τ . Note that the graph has four vertices even though only elements from $\{1, 2, 3\}$ are used to label the nodes of the tree. Intuitively, once a label has been forgotten, it can be reused to define a new vertex.

Definition 6. The terms in ITD_k are called k -instructive tree decompositions. Terms in ITD_k that do not use the symbol **Join** are called k -instructive path decompositions. We let IPD_k denote the set of all k -instructive path decompositions.

For each $k \in \mathbb{N}$, we let $\mathcal{F}_k = \{f : \mathfrak{b} \rightarrow \mathbb{N} \mid \mathfrak{b} \subseteq [k+1], f \text{ is injective}\}$ be the set of injective functions from some subset $\mathfrak{b} \subseteq [k+1]$ to \mathbb{N} . As a last step, we define a function $\mathcal{G} : \bigcup_{k \in \mathbb{N}} \text{ITD}_k \rightarrow \text{GRAPHS}$ that assigns a graph $\mathcal{G}(\tau)$ to each $\tau \in \bigcup_k \text{ITD}_k$. This function is defined inductively below, together with an auxiliary function $\theta : \bigcup_k \text{ITD}_k \rightarrow \bigcup_k \mathcal{F}_k$ that assigns, for each $k \in \mathbb{N}$, and each k -instructive tree decomposition τ , an injective map $\theta[\tau] : \mathfrak{b} \rightarrow \mathbb{N}$ in the set \mathcal{F}_k . In this way the pair $(\mathcal{G}(\tau), \theta[\tau])$ forms a k -boundaried graph. Each element $u \in \mathfrak{b}$ is said to be an *active label* for $\mathcal{G}(\tau)$, and the vertex $\theta[\tau](u)$ is the active vertex labeled with u . The functions \mathcal{G} and θ are inductively defined as follows. We note that for each $\tau \in \bigcup_k \text{ITD}_k$, we specify the injective map $\theta[\tau]$ as a subset of pairs from $[k+1] \times \mathbb{N}$.

1. If $\tau = \text{Leaf}$, then $\mathcal{G}(\tau) = (\emptyset, \emptyset, \emptyset)$ and $\theta[\tau] = \emptyset$.

2. If $\tau = \text{IntroVertex}\{u\}(\sigma)$ then

$$\mathcal{G}(\tau) = (V_{\mathcal{G}(\sigma)} \cup \{|V_{\mathcal{G}(\sigma)}| + 1\}, E_{\mathcal{G}(\sigma)}, \rho_{\mathcal{G}(\sigma)}), \text{ and } \theta[\tau] = \theta[\sigma] \cup \{(u, |V_{\mathcal{G}(\sigma)}| + 1)\}.$$

3. If $\tau = \text{ForgetVertex}\{u\}(\sigma)$, then $\mathcal{G}(\tau) = \mathcal{G}(\sigma)$, and $\theta[\tau] = \theta[\sigma] \setminus \{(u, \theta[\sigma](u))\}$.

4. If $\tau = \text{IntroEdge}\{u, v\}(\sigma)$, then

$$\mathcal{G}(\tau) = (V_{\mathcal{G}(\sigma)}, E_{\mathcal{G}(\sigma)} \cup \{|E_{\mathcal{G}(\sigma)}| + 1\}, \rho_{\mathcal{G}(\sigma)} \cup \{(|E_{\mathcal{G}(\sigma)}| + 1, \theta[\sigma](u)), (|E_{\mathcal{G}(\sigma)}| + 1, \theta[\sigma](v))\}), \text{ and } \theta[\tau] = \theta[\sigma].$$

5. If $\tau = \text{Join}(\sigma_1, \sigma_2) \in \text{ITD}_k$, then $(\mathcal{G}(\tau), \theta[\tau]) = (\mathcal{G}(\sigma_1), \theta[\sigma_1]) \oplus (\mathcal{G}(\sigma_2), \theta[\sigma_2])$.

In Item 5, the operation \oplus is the join of two boundaried graphs (see Section 2).

By letting, for each $k \in \mathbb{N}$, \mathcal{G}_k be the restriction of \mathcal{G} to the set ITD_k , we have that the sequence $\text{ITD} = \{(\Sigma_k, \text{ITD}_k, \mathcal{G}_k)\}_{k \in \mathbb{N}}$ is a treelike decomposition class. We call this class the *instructive tree decomposition class*. Note that ITD has complexity 2^k , since as discussed above, for each $k \in \mathbb{N}$, ITD_k is accepted by a tree automaton \mathcal{A}_k with 2^k states. The following lemma implies that ITD realizes treewidth.

Lemma 7. Let $G \in \text{GRAPHS}$ and $k \in \mathbb{N}$. Then G has treewidth at most k if and only if there exists a k -instructive tree decomposition τ such that $\mathcal{G}(\tau) \simeq G$.

5 Treelike Dynamic Programming Cores

In this section, we introduce the notion of a *treelike dynamic-programming core* (treelike DP-core), a formalism intended to capture the behavior of dynamic programming algorithms operating on treelike decompositions. Our formalism generalizes and refines the notion of dynamic programming core introduced in [8]. There are two crucial differences. First, our framework can be used to define DP-cores for classes of dense graphs, such as graphs of constant cliquewidth, whereas the DP-cores devised in [8] are specialized to work on tree decompositions. Second, and most importantly, in our framework, graphs of width k can be represented as terms over ranked alphabets whose size depend only on k . This property makes our framework modular and particularly suitable for applications in the realm of automated theorem proving.

Definition 8 (Treelike DP-Cores). *A treelike dynamic programming core is a sequence of 6-tuples $D = \{(\Sigma_k, \mathcal{W}_k, \text{Final}_k, \Delta_k, \text{Clean}_k, \text{Inv}_k)\}_{k \in \mathbb{N}}$ where for each $k \in \mathbb{N}$,*

1. Σ_k is a ranked alphabet;
2. \mathcal{W}_k is a decidable subset of $\{0, 1\}^*$;
3. $\text{Final}_k : \mathcal{W}_k \rightarrow \{0, 1\}$ a function;
4. Δ_k is a set containing
 - a finite subset $\hat{a} \in \mathcal{P}_{\text{fin}}(\mathcal{W}_k)$ for each symbol a of arity 0,
 - a function $\hat{a} : \mathcal{W}_k^{\times \mathfrak{r}(a)} \rightarrow \mathcal{P}_{\text{fin}}(\mathcal{W}_k)$ for each symbol a of arity $\mathfrak{r}(a) \geq 1$;
5. $\text{Clean}_k : \mathcal{P}_{\text{fin}}(\mathcal{W}_k) \rightarrow \mathcal{P}_{\text{fin}}(\mathcal{W}_k)$ is a function;
6. $\text{Inv}_k : \mathcal{P}_{\text{fin}}(\mathcal{W}_k) \rightarrow \{0, 1\}^*$ is a function.

We let $D[k] = (\Sigma_k, \mathcal{W}_k, \text{Final}_k, \Delta_k, \text{Clean}_k, \text{Inv}_k)$ denote the k -th tuple of D . We may write $D[k].\Sigma$ to denote the set Σ_k , $D[k].\mathcal{W}$ to denote the set \mathcal{W}_k , and so on. Intuitively, for each k , $D[k]$ is a description of a dynamic programming algorithm that operates on terms from $\text{Terms}(\Sigma_k)$. This algorithm processes such a term τ from the leaves towards the root, and assigns a set of local witnesses to each node of τ . The algorithm starts by assigning the set $D[k].\hat{a}$ to each leaf node labeled with symbol a . Subsequently, the set of local witnesses to be assigned to each internal node p is computed by taking into consideration the label of the node, and the set (sets) of local witnesses assigned to the child (children) of p . The algorithm accepts τ if at the end of the process, the set of local witnesses associated with the root node $\text{root}(\tau)$ has some final local witness, i.e., some local witness $w \in \mathcal{W}$ such that $D[k].\text{Final}(w) = 1$.

Dynamic programming algorithms often make use of a function that removes redundant elements from a given set of local witnesses. In our framework, this is formalized by the function $D[k].\text{Clean}$, which is applied to each non-leaf node as soon as the set of local witnesses associated with this node has been computed. The function $D[k].\text{Inv}$ is useful in the context of optimization problems. For instance, given a set S of local witnesses encoding weighted partial solutions to a given problem, $D[k].\text{Inv}(S)$ may return (a binary encoding of) the minimum/maximum weight of a partial solution in the set.

We note that for the moment, terms in $\text{Terms}(\Sigma_k)$ have no semantic meaning. Nevertheless, later in this section, we will consider terms that correspond to \mathbf{C} -decompositions of width at most k , for a fixed decomposition class \mathbf{C} . In this context, the intuition is that if S is the set associated with the root of τ , then $D[k].\text{Inv}(S)$ corresponds to some invariant of the graph $\mathcal{G}(\tau)$, such as the minimum size of a vertex cover, the maximum size of an independent set, etc.

The dynamic programming process described above can be formalized in our framework using the notion of k -th *dynamization* of a dynamic core D , which is a function $\Gamma[D, k]$ that assigns a

set $\Gamma[D, k](\tau)$ of local witnesses to each term $\tau \in \text{Terms}(k)$. Given a symbol a of arity r in the set $D[k].\Sigma$, and subsets $S_1, \dots, S_r \subseteq D[k].\mathcal{W}$, we define the following set:

$$D[k].\hat{a}(S_1, \dots, S_r) \doteq D[k].\text{Clean} \left(\bigcup_{i \in [r], w_i \in S_i} D[k].\hat{a}(w_1, \dots, w_r) \right). \quad (1)$$

Using this notation, for each $k \in \mathbb{N}$, the function $\Gamma[D, k]$ is defined by induction on the structure of τ as follows.

Definition 9 (Dynamization). *Let D be a treelike DP-core. For each $k \in \mathbb{N}$, the k -th dynamization of D is the function $\Gamma[D, k] : \text{Terms}(D[k].\Sigma) \rightarrow \mathcal{P}_{\text{fin}}(D[k].\mathcal{W})$ inductively defined as follows.*

1. If $\tau = a$ for some symbol $a \in D[k].\Sigma$ of arity 0, then $\Gamma[D, k](\tau) = D[k].\hat{a}$.
2. If $\tau = a(\tau_1, \dots, \tau_r)$ for some $a \in D[k].\Sigma$ of arity r , and some terms τ_1, \dots, τ_r in $\text{Terms}(D[k].\Sigma)$, then $\Gamma[D, k](\tau) = D[k].\hat{a}(\Gamma[D, k](\tau_1), \dots, \Gamma[D, k](\tau_r))$.

For each $k \in \mathbb{N}$, we say that a term $\tau \in \text{Terms}(D[k].\Sigma)$ is *accepted* by $D[k]$ if $\Gamma[D, k](\tau)$ contains a final local witness, i.e., a local witness w with $D[k].\text{Final}(w) = 1$. We let $\text{Acc}(D[k])$ denote the set of all terms accepted by $D[k]$. We let $\text{Acc}(D) = \bigcup_{k \in \mathbb{N}} \text{Acc}(D[k])$.

So far, our notion of a treelike DP-core is just a symbolic formalism for the specification sequences of tree languages (one tree language $\text{Acc}(D[k])$ for each $k \in \mathbb{N}$). Our formalism is very close in spirit from tree automata, except for the fact that transitions and states are specified implicitly. Next, we show that when combined with the notion of a treelike decomposition class, DP-cores can be used to define graph properties.

Definition 10 (Graph Property of a DP-Core). *Let C be a treelike decomposition class, and D be a treelike DP-core. For each $k \in \mathbb{N}$, the graph property of $D[k]$ is the set*

$$\mathbb{G}[D[k], C] = \text{ISO}(\{\mathcal{G}(\tau) : \tau \in L_k \cap \text{Acc}(D[k])\}).$$

The graph property defined by D is the set $\mathbb{G}[D, C] = \bigcup_k \mathbb{G}[D[k], C]$.

We note that for each $k \in \mathbb{N}$, $\mathbb{G}[D[k], C] \subseteq \mathbb{G}[C_k]$, and hence, $\mathbb{G}[D, C] \subseteq \mathbb{G}[C]$.

5.1 Coherency

In order to be useful in the context of model-checking and automated theorem proving, DP-cores need to behave coherently with respect to distinct treelike decompositions of the same graph. This intuition is formalized by the following definition.

Definition 11 (Coherency). *Let $C = \{(\Sigma_k, L_k, \mathcal{G}_k)\}_{k \in \mathbb{N}}$ be a treelike decomposition class, and D be a treelike DP-core. We say that D is C -coherent if for each $k \in \mathbb{N}$, $\Sigma_k = D[k].\Sigma$, and for each $k, k' \in \mathbb{N}$, and each $\tau \in L_k$ and $\tau' \in L_{k'}$ with $\mathcal{G}(\tau) \simeq \mathcal{G}(\tau')$,*

1. $\tau \in \text{Acc}(D[k])$ if and only if $\tau' \in \text{Acc}(D[k'])$, and
2. $D.\text{Inv}(\Gamma[D, k](\tau)) = D.\text{Inv}(\Gamma[D, k'](\tau'))$.

Let D be a C -coherent treelike DP-core. Condition 1 of Definition 11 guarantees that if a graph G belongs to $\mathbb{G}[D, C]$, then for each $k \in \mathbb{N}$ and each C -decomposition τ of width at most k such that $\mathcal{G}(\tau) \simeq G$, we have that $\tau \in \text{Acc}(D[k])$. On the other hand, if G does not belong to $\mathbb{G}[D, C]$, then no C -decomposition τ with $\mathcal{G}(\tau) \simeq G$ belongs to $\text{Acc}(D)$. This discussion is formalized in the following proposition.

Proposition 12. *Let $C = \{(\Sigma_k, L_k, \mathcal{G}_k)\}_{k \in \mathbb{N}}$ be a treelike decomposition class, and D be a C -coherent treelike DP-core. Then for each $k \in \mathbb{N}$, and each $\tau \in L_k$, we have that $\mathcal{G}(\tau) \in \mathbb{G}[D, C]$ if and only if $\tau \in \text{Acc}(D[k])$.*

Proof. Let $k \in \mathbb{N}$ and $\tau \in \mathbb{L}_k$. Suppose that $\tau \in \text{Acc}(\mathbb{D}[k])$. Then, by Definition 10, $\mathcal{G}(\tau) \in \mathbb{G}[\mathbb{D}[k], \mathbb{C}]$, and therefore, we have that $\mathcal{G}(\tau) \in \mathbb{G}[\mathbb{D}, \mathbb{C}]$. Note that this direction holds even if \mathbb{D} is not coherent.

For the converse, we do need coherency. Suppose $\mathcal{G}(\tau) \in \mathbb{G}[\mathbb{D}, \mathbb{C}]$. Then, there is some $k' \in \mathbb{N}$ and some $\tau' \in \text{Acc}(\mathbb{D}[k'])$ such that $\mathcal{G}(\tau') \simeq \mathcal{G}(\tau)$. Since \mathbb{D} is \mathbb{C} -coherent, we can infer from Definition 11.1 that $\tau \in \text{Acc}(\mathbb{D}[k])$. \square

A nice consequence of Proposition 12 is that if \mathbb{D} is a \mathbb{C} -coherent treelike DP-core, then, in order to determine whether a given graph G belongs to $\mathbb{G}[\mathbb{D}, \mathbb{C}]$, it is enough to select an arbitrary \mathbb{C} -decomposition of G and then to determine whether τ belongs to $\text{Acc}(\mathbb{D}[k])$, where k is the \mathbb{C} -width of τ . In this way, the analysis of the complexity of testing whether $G \in \mathbb{G}[\mathbb{D}, \mathbb{C}]$ can be split into two parts. First, the analysis of the complexity of computing a \mathbb{C} -decomposition of minimum width k (or approximately minimum width). Second, the analysis of the complexity of verifying whether τ belongs to $\text{Acc}(\mathbb{D}[k])$. This second step is carried on in details in Theorem 22 using the complexity measures introduced in Definition 14. The construction of \mathbb{C} -decompositions of (approximately) minimum width is not a focus of this work. Nevertheless, it is worth noting that for several well studied width measures that can be formalized in our framework, decompositions of minimum width (or approximately minimum width) can be constructed in time FPT on the width parameter [53, 55, 19, 44, 63, 12]. Finally, it is worth noting that coherent DP-cores have also applications to the context of width-based automated theorem proving (Theorem 27 and Theorem 33). In this context, one does not need to consider the problem of computing a \mathbb{C} -decompositions of a given input graph.

Coherent DP-cores may be used to define not only graph properties but also graph invariants, as specified in Definition 13.

Definition 13 (Invariant of a DP-Core). *Let \mathbb{C} be a decomposition class and \mathbb{D} be a \mathbb{C} -coherent treelike DP-core. The $\mathbb{G}[\mathbb{D}, \mathbb{C}]$ -invariant defined by \mathbb{D} is the function $\mathcal{I}[\mathbb{D}, \mathbb{C}] : \mathbb{G}[\mathbb{D}, \mathbb{C}] \rightarrow \{0, 1\}^*$ that assigns to each graph $G \in \mathbb{G}[\mathbb{D}, \mathbb{C}]$, the string $\mathbb{D}[w_{\mathbb{C}}(\tau)].\text{Inv}(\Gamma[\mathbb{D}, k](\tau))$ where τ is an arbitrary \mathbb{C} -decomposition with $\mathcal{G}(\tau) \simeq G$.*

We note that Condition 2 of Definition 11 guarantees that

$$\mathbb{D}[w_{\mathbb{C}}(\tau)].\text{Inv}(\Gamma[\mathbb{D}, k](\tau)) = \mathbb{D}[w_{\mathbb{C}}(\tau')].\text{Inv}(\Gamma[\mathbb{D}, k](\tau'))$$

for any two \mathbb{C} -decompositions τ and τ' with $\mathcal{G}(\tau) \simeq \mathcal{G}(\tau')$. Therefore, for each graph $G \in \mathbb{G}[\mathbb{D}, \mathbb{C}]$, the value $\mathcal{I}[\mathbb{D}, \mathbb{C}](G)$ is well defined, and invariant under graph isomorphism.

5.2 Complexity Measures

In order to analyze the behavior of treelike DP-cores from a quantitative point of view we consider four complexity measures. We say that a set S of local witnesses is (\mathbb{D}, k, n) -useful if there is some $\tau \in \text{Terms}(\mathbb{D}[k].\Sigma)$ of size $|\tau|$ at most n such that $\Gamma[\mathbb{D}, k](\tau) = S$. We say that a local witness w is (\mathbb{D}, k, n) -useful if it belongs to some (\mathbb{D}, k, n) -useful set.

Definition 14 (Complexity Measures). *Let \mathbb{D} be a treelike DP-core, and $n \in \mathbb{N}$.*

1. *Bitlength:* we let $\beta_{\mathbb{D}}(k, n)$ denote the maximum number of bits in an (\mathbb{D}, k, n) -useful witness. We call $\beta_{\mathbb{D}}$ the bitlength of \mathbb{D} .
2. *Multiplicity:* we let $\mu_{\mathbb{D}}(k, n)$ denote the maximum number of elements in a (\mathbb{D}, k, n) -useful set. We call $\mu_{\mathbb{D}}$ the multiplicity of \mathbb{D} .
3. *State Complexity:* we let $\nu_{\mathbb{D}}(k, n)$ be the number of (\mathbb{D}, k, n) -useful witnesses. We call $\nu_{\mathbb{D}}$ the state complexity of \mathbb{D} .
4. *Deterministic State Complexity:* we let $\delta_{\mathbb{D}}(k, n)$ denote the number of (\mathbb{D}, k, n) -useful sets. We call $\delta_{\mathbb{D}}$ the deterministic state complexity of \mathbb{D} .

The next observation establishes some straightforward relations between these complexity measures.

Observation 15. *Let D be a k -abstract DP-core. Then, for each $n \in \mathbb{N}$, the following inequalities are verified.*

$$\begin{aligned}\mu_D(k, n) &\leq \nu_D(k, n) \leq 2^{\beta_D(k, n)} \\ \delta_D(k, n) &\leq \min\{2 \cdot \nu_D(k, n)^{\mu_D(k, n)}, 2^{\nu_D(k, n)}\} \leq 2^{2^{\beta_D(k, n)}}.\end{aligned}$$

Proof. The maximum size of a (D, k, n) -useful set of witnesses is clearly upper bounded by the total number of (D, k, n) -useful witnesses. Therefore, $\mu_D(k, n) \leq \nu_D(k, n)$. Since the maximum number of bits needed to represent a (D, k, n) -useful witness is $\beta_D(k, n)$, we have $\nu_D(k, n) \leq 2^{\beta_D(k, n)}$. Now, the number of (D, k, n) -useful sets of witnesses is upper bounded by $\sum_{i=0}^{\mu_D(k, n)} \binom{\nu_D(k, n)}{i}$ which is always smaller than both $2^{\nu_D(k, n)}$ and $2 \cdot \nu_D(k, n)^{\mu_D(k, n)}$. Finally, the last inequality is obtained by using the fact that $\nu_D(k, n) \leq 2^{\beta_D(k, n)}$. \square

An important class of DP-cores is the class of cores where maximum number of bits in a useful local witness is independent of the size of a term τ . In other words, the size may depend on k but not on $|\tau|$.

Definition 16 (Finite DP-cores). *We say that a treelike DP-core D is finite if there is a function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that for each $n \in \mathbb{N}$, $\beta_D(k, n) \leq f(k)$.*

If D is a finite DP-core then we may write simply $\beta_D(k)$, $\mu_D(k)$, $\nu_D(k)$, and $\delta_D(k)$ to denote the functions $\beta_D(k, n)$, $\mu_D(k, n)$, $\nu_D(k, n)$, and $\delta_D(k, n)$ respectively.

In this work, we will be concerned with DP-cores that are *internally polynomial*, as defined next. Typical dynamic programming algorithms operating on tree-like decompositions give rise to internally polynomial DP-cores.

Definition 17 (Internally Polynomial DP-Cores). *Let D be a treelike DP-core. We say that D is internally polynomial if the following conditions are satisfied.*

1. *For each $k \in \mathbb{N}$, and each (D, k, n) -useful set S , $|D[k].\text{Inv}(S)| = \beta_D(k, n)^{O(1)}$.*
2. *There is a deterministic algorithm \mathfrak{A} such that the following conditions are satisfied.*
 - (a) *Given $k \in \mathbb{N}$, and a string $w \in \{0, 1\}^*$, \mathfrak{A} decides whether $w \in D[k].\mathcal{W}$ in time $(k + |w|)^{O(1)}$,*
 - (b) *Given $k \in \mathbb{N}$, and a symbol $a \in \Sigma_k$ of arity 0, the algorithm constructs the set $D[k].\hat{a}$ in time $\beta_D(k, 0) \cdot \mu_D(k, 0)$.*
 - (c) *Given $k \in \mathbb{N}$, a symbol $a \in \Sigma_k$, and an input X for the function $D[k].\hat{a}$, the algorithm constructs the set $D[k].\hat{a}(X)$ in time $(k + |X|)^{O(1)}$.*
 - (d) *Given $k \in \mathbb{N}$, an element $\text{Function} \in \{\text{Final}, \text{Clean}, \text{Inv}\}$, and an input X for the function $D[k].\text{Function}$, the algorithm computes the value $D[k].\text{Function}(X)$ in time $(k + |X|)^{O(1)}$.*

Intuitively, a DP-core D is internally polynomial if there is an algorithm \mathfrak{A} that when given $k \in \mathbb{N}$ as input simulates the behavior of $D[k]$ in such a way that the output of the invariant function has polynomially many bits in the bitlength of $D[k]$; \mathfrak{A} decides membership in the set $D[k].\mathcal{W}$ in time polynomial in k plus the size of the queried string; for each symbol a of arity 0, \mathfrak{A} constructs the set $D[k].\hat{a}$ in time polynomial in k plus the maximum number of bits needed to describe such a set; and \mathfrak{A} computes each function in $D[k]$ in time polynomial in k plus the size of the input of the function. Note that the fact that D is internally polynomial does not imply that one can determine whether a given term τ is accepted by D in time polynomial in $|\tau|$. The complexity of this test is governed by the complexity measures of Definition 14 (see Theorem 22).

5.3 Example: A DP-Core for VertexCover_r

Let $G = (V, E, \rho)$ be a graph. A subset X of V is a vertex cover of G if every edge of G has at least one endpoint in X . We let VertexCover_r be the graph property consisting of all graphs that have a vertex cover of size at most r .

Let $\text{ITD} = \{(\Sigma_k, \text{ITD}_k, \mathcal{G}_k)\}_{k \in \mathbb{N}}$ be the decomposition class defined in Section 4, which realizes treewidth. Next, we give the specification of an ITD-coherent treelike DP-core $\mathbf{C}\text{-VertexCover}_r$, with graph property $\mathbb{G}[\mathbf{C}\text{-VertexCover}_r, \text{ITD}] = \text{VertexCover}_r$. It is enough to specify, for each $k \in \mathbb{N}$, the components of $\mathbf{C}\text{-VertexCover}_r[k]$. A local witness for $\mathbf{C}\text{-VertexCover}_r[k]$ is a pair $\mathbf{w} = (R, s)$ where $R \subseteq [k+1]$ and $s \in \mathbb{N}$. Intuitively, R denotes the set of active labels associated with vertices of a partial vertex cover, and s denotes the size of the partial vertex cover. Therefore, we set

$$\mathbf{C}\text{-VertexCover}_r[k].\mathcal{W} = \{(R, s) : R \subseteq [k+1], s \in \{0, 1, \dots, r\}\}.$$

In this particular DP-core, each local witness is final. In other words, for each local witness \mathbf{w} , we have

$$\mathbf{C}\text{-VertexCover}_r[k].\text{Final}(\mathbf{w}) = 1.$$

If S is a set of local witnesses, and (R, s) and (R, s') are local witnesses in S with $s < s'$, then (R, s') is redundant. The clean function of the DP-core takes a set of local witnesses as input and removes redundancies. More precisely,

$$\mathbf{C}\text{-VertexCover}_r[k].\text{Clean}(S) = \{(R, s) \in S : \nexists s' < s, (R, s') \in S\}.$$

The invariant function of the core takes a set of witnesses as input and returns the smallest value s with the property that there is some subset $R \subseteq [k+1]$ with $(R, s) \in S$.

$$\mathbf{C}\text{-VertexCover}_r[k].\text{Inv}(S) = \min\{s : \exists R \text{ s.t. } (R, s) \in S\}.$$

Next, we define the transition functions of the DP-core.

Definition 18. Let $\mathbf{w} = (R, s)$ and $\mathbf{w}' = (R', s')$ be local witnesses, and $u, v \in [k+1]$ be such that $u \neq v$.

$$1. \mathbf{C}\text{-VertexCover}_r[k].\text{LeafSet} = \{(\emptyset, 0)\}.$$

$$2. \mathbf{C}\text{-VertexCover}_r[k].\text{IntroVertex}\{u\}(\mathbf{w}) = \{\mathbf{w}\}.$$

$$3. \mathbf{C}\text{-VertexCover}_r[k].\text{ForgetVertex}\{u\}(\mathbf{w}) = \{(R \setminus \{u\}, s)\}.$$

$$4. \mathbf{C}\text{-VertexCover}_r[k].\text{IntroEdge}\{u, v\}(\mathbf{w}) = \begin{cases} \{\mathbf{w}\} & \text{if } u \text{ or } v \in R, \\ \emptyset & \text{if } u, v \notin R \text{ and } s = r, \\ \{(R \cup \{u\}, s+1), (R \cup \{v\}, s+1)\} & \text{otherwise.} \end{cases}$$

$$5. \mathbf{C}\text{-VertexCover}_r[k].\text{Join}(\mathbf{w}, \mathbf{w}') = \begin{cases} \{(R \cup R', s + s' - |R \cap R'|)\} & \text{If } s + s' - |R \cap R'| \leq r, \\ \{\} & \text{otherwise.} \end{cases}$$

It should be clear that $\mathbf{C}\text{-VertexCover}_r[k]$ is both finite and internally polynomial. The next lemma characterizes, for each k -instructive tree decomposition τ , the local witnesses \mathbf{w} that are present in the set $\Gamma[\mathbf{C}\text{-VertexCover}_r, k](\tau)$.

Lemma 19. Let $r \in \mathbb{N}$. For each $k \in \mathbb{N}$, each k -instructive tree decomposition τ , and each local witness $\mathbf{w} = (R, s)$ in $\mathbf{C}\text{-VertexCover}_r[k].\mathcal{W}$, $\mathbf{w} \in \Gamma[\mathbf{C}\text{-VertexCover}_r, k](\tau)$ if and only if the following predicate is satisfied:

- $\mathbf{P}\text{-VertexCover}_r[k](\tau, w) \equiv$

1. $s \leq r$,
2. s is the minimum size of a vertex cover X in $\mathcal{G}(\tau)$ with $\theta[\tau](R) = X \cap \text{Im}(\theta[\tau](B(\tau)))$.

The proof of Lemma 19 follows straightforwardly by induction on the structure of τ . For completeness, and also for illustration purposes, a detailed proof can be found in Appendix D. Lemma 19 implies that $\mathbf{C}\text{-VertexCover}_r$ is coherent and that for each $k \in \mathbb{N}$, the graph property $\mathbb{G}[\mathbf{C}\text{-VertexCover}_r[k], \text{ITD}]$ is the set of all graphs of treewidth at most k with a vertex cover of size at most r .

Theorem 20. *For each $r \in \mathbb{N}$, the DP-core $\mathbf{C}\text{-VertexCover}_r$ is coherent. Additionally, for each $k \in \mathbb{N}$, $\mathbb{G}[\mathbf{C}\text{-VertexCover}_r[k], \text{ITD}] = \text{VertexCover}_r \cap \text{GRAPHSTW}[k]$.*

Proof. Let $k \in \mathbb{N}$, G be a graph of treewidth at most k , and $\tau \in \text{ITD}_k$ be such that $\mathcal{G}(\tau) \simeq G$. Then, by Lemma 19, $\Gamma[\mathbf{C}\text{-VertexCover}_r, k](\tau)$ is nonempty (i.e. has some local witness) if and only if G has a vertex cover of size at most r . Since, for this particular DP-core, every local witness is final, we have that $\tau \in \text{Acc}(\mathbf{C}\text{-VertexCover}_r[k])$ if and only if G has a vertex cover of size at most r . Therefore, $G \in \mathbb{G}[\mathbf{C}\text{-VertexCover}_r[k], \text{ITD}]$ if and only if G has a vertex cover of size at most r . \square

Now, consider the DP-core $\mathbf{C}\text{-VertexCover}$ (i.e. without the subscript r) where for each $k \in \mathbb{N}$, all components are identical to the components of $\mathbf{C}\text{-VertexCover}_r$, except for the local witnesses (R, s) , where now s is allowed to be any number in \mathbb{N} , and for the edge introduction function $\mathbf{C}\text{-VertexCover}[k].\text{IntroEdge}\{u, v\}$ which is defined as follows on each local witness $w = (R, s)$.

$$\mathbf{C}\text{-VertexCover}[k].\text{IntroEdge}\{u, v\}(w) = \begin{cases} \{w\} & \text{if } u \text{ or } v \in R, \\ \{(R \cup \{u\}, s+1), (R \cup \{v\}, s+1)\} & \text{otherwise.} \end{cases}$$

Then, in this case, the core is not anymore finite because one can impose no bound on the value s of a local witness (R, s) . Still, we have that the multiplicity of $\mathbf{C}\text{-VertexCover}[k]$ is bounded by 2^{k+1} (i.e. a function of k only) because the clean function eliminates redundancies. Additionally, one can show that this variant actually computes the minimum size of a vertex cover in the graph represented by a k -instructive tree decomposition. Below, we let $D = \mathbf{C}\text{-VertexCover}$ and $\mathcal{I}[D, C] : \mathbb{G}[D, C] \rightarrow \{0, 1\}^*$ be the graph invariant computed by D .

Theorem 21. *Let τ be a k -instructive tree decomposition. Then, $\mathcal{I}[D, C](\mathcal{G}(\tau))$ is the (binary encoding of) the minimum size of a vertex cover in $\mathcal{G}(\tau)$.*

We omit the proof of Theorem 21 given that the proof is very similar to the proof of Theorem 20. In particular, this theorem is a direct consequence of an analog of Lemma 19 where $\mathbf{C}\text{-VertexCover}_r$ is replaced by $\mathbf{C}\text{-VertexCover}$ and the predicate $\mathbf{P}\text{-VertexCover}_r$ is replaced by the predicate $\mathbf{P}\text{-VertexCover}$ obtained by omitting the first condition ($s \leq r$).

5.4 Model Checking and Invariant Computation

Let $C = \{(\Sigma_k, L_k, \mathcal{G}_k)\}_{k \in \mathbb{N}}$ be a treelike decomposition class, and D be a C -coherent treelike DP-core. Given a C -decomposition of width at most k , we can use the notion of dynamization (Definition 9), to check whether the graph $\mathcal{G}(\tau)$ encoded by τ belongs to the graph property $\mathbb{G}[D, C]$ represented by D . The next theorem states that the complexity of this model-checking task is essentially governed by the bitlength and by the multiplicity of D . We note that in typical applications the arity r of a decomposition class is a constant (most often 1 or 2), and the width k is smaller than $\beta_D(k, n)$ for each $n \in N$. Nevertheless, for completeness, we explicitly include the dependence on $k^{O(1)}$ and $r^{O(1)}$ in the calculation of the running time.

Theorem 22 (Model Checking). *Let $\mathbf{C} = \{(\Sigma_k, \mathbf{L}_k, \mathcal{G}_k)\}_{k \in \mathbb{N}}$ be a treelike decomposition class of arity r , \mathbf{D} be an internally polynomial \mathbf{C} -coherent treelike DP-Core, and let τ be a \mathbf{C} -decomposition of \mathbf{C} -width at most k and size $|\tau| = n$.*

1. *One can determine whether $\mathcal{G}(\tau) \in \mathbb{G}[\mathbf{D}, \mathbf{C}]$ in time*

$$T(k, n) = n \cdot k^{O(1)} \cdot r^{O(1)} \cdot \beta_{\mathbf{D}}(k, n)^{O(1)} \cdot \mu_{\mathbf{D}}(k, n)^{r+O(1)}.$$

2. *One can compute the invariant $\mathcal{I}[\mathbf{D}, \mathbf{C}](\mathcal{G}(\tau))$ in time*

$$T(k, n) + k^{O(1)} \cdot \beta_{\mathbf{D}}(k, n)^{O(1)} \cdot \mu_{\mathbf{D}}(k, n)^{O(1)}.$$

Proof. Since \mathbf{D} is \mathbf{C} -coherent, and since τ has \mathbf{C} -width at most k , by Proposition 12, we have that $\mathcal{G}(\tau)$ belongs to $\mathbb{G}[\mathbf{D}, \mathbf{C}]$ if and only if $\tau \in \text{Acc}(\mathbf{D}[k])$. In other words, if and only if the set $\Gamma[\mathbf{D}, k](\tau)$ has some final local witness. Therefore, the model-checking algorithm consists of two steps: we first compute the set $\Gamma[\mathbf{D}, k](\tau)$ inductively using Definition 9, and subsequently, we test whether this set has some final local witness. Since $|\tau| = n$, we have that τ has at most n sub-terms. Let τ be such a subterm. If $\sigma = a$ for some symbol of arity 0, then the construction of the set $\Gamma[\mathbf{D}, k](\sigma)$ takes time at most $\beta_{\mathbf{D}}(k, n)^{O(1)} \cdot \mu_{\mathbf{D}}(k, n)^{O(1)}$. Now, suppose that $\sigma = a(\sigma_1, \dots, \sigma_{\tau(a)})$ for some symbol $a \in \Sigma_k$, and some terms $\sigma_1, \dots, \sigma_{\tau(a)}$ in $\text{Terms}(\Sigma_k)$, and assume that the sets $\Gamma[\mathbf{D}, k](\sigma_1), \dots, \Gamma[\mathbf{D}, k](\sigma_{\tau(a)})$ have been computed. We claim that using these precomputed sets, together with Equation 1, the set $\Gamma[\mathbf{D}, k](\sigma)$ can be constructed in time $k^{O(1)} \cdot r^{O(1)} \cdot \beta_{\mathbf{D}}(k, n)^{O(1)} \cdot \mu_{\mathbf{D}}(k, n)^{r+O(1)}$. To see this, we note that in order to construct $\Gamma[\mathbf{D}, k](\sigma)$, we need to construct, for each tuple $(\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_{\tau(a)})$ of local witnesses in the Cartesian product $\Gamma[\mathbf{D}, k](\sigma_1) \times \dots \times \Gamma[\mathbf{D}, k](\sigma_{\tau(a)})$, the set $\mathbf{D}[k].\hat{a}(\mathbf{w}_1, \dots, \mathbf{w}_{\tau(a)})$. Since, by assumption \mathbf{D} is internally polynomial, this set can be constructed in time at most $k^{O(1)} \cdot r^{O(1)} \cdot \beta_{\mathbf{D}}(k, n)^{O(1)}$. That is to say, polynomial in k plus the size of the input to this function, which is upper bounded by $r \cdot \beta_{\mathbf{D}}(k, n)$. In particular, this set has at most $k^{O(1)} \cdot r^{O(1)} \cdot \beta_{\mathbf{D}}(k, n)^{O(1)}$ local witnesses. Since we need to consider at most $\mu_{\mathbf{D}}(k, n)^r$ tuples, taking the union of all such sets $\mathbf{D}[k].\hat{a}(\mathbf{w}_1, \dots, \mathbf{w}_{\tau(a)})$ takes time $k^{O(1)} \cdot r^{O(1)} \cdot \beta_{\mathbf{D}}(k, n)^{O(1)} \cdot \mu_{\mathbf{D}}(k, n)^{r+O(1)}$. Finally, once the union has been computed, since \mathbf{D} is internally polynomial, the application of the function $\mathbf{D}[k].\text{Clean}$ takes an additional (additive) factor of $k^{O(1)} \cdot \beta_{\mathbf{D}}(k, n)^{O(1)} \cdot \mu_{\mathbf{D}}(k, n)^{O(1)}$. Therefore, the overall computation of the set $\Gamma[\mathbf{D}, k](\tau)$, and subsequent determination of whether this set contains a final local witness takes time

$$T(k, n) = n \cdot k^{O(1)} \cdot r^{O(1)} \cdot \beta_{\mathbf{D}}(k, n)^{O(1)} \cdot \mu_{\mathbf{D}}(k, n)^{r+O(1)}.$$

For the second item, after having computed $\Gamma[\mathbf{D}, k](\tau)$ in time $T(k, n)$, we need to compute the value of $\mathbf{D}[k].\text{Inv}$ on this set. This takes time $k^{O(1)} \cdot \beta_{\mathbf{D}}(k, n)^{O(1)} \cdot \mu_{\mathbf{D}}(k, n)^{O(1)}$. Therefore, in overall, we need time $T(k, n) + k^{O(1)} \cdot \beta_{\mathbf{D}}(k, n)^{O(1)} \cdot \mu_{\mathbf{D}}(k, n)^{O(1)}$ to compute the value $\mathcal{I}[\mathbf{D}, \mathbf{C}](\mathcal{G}(\tau))$. \square

It is worth noting that for finite cores, where the number of bits in a local witness depends only on k , but not on the size of a decomposition, the running times in Theorem 22 are of the form $f(k) \cdot n$, in other words, fixed-parameter linear with respect to k . On the other hand, even when a DP-core \mathbf{D} is not finite, the model-checking algorithm of Theorem 22 may still have a running time of the form $f(k) \cdot n^{O(1)}$. The reason is that for these algorithms, even though the corresponding cores are not finite, in the sense that the number bits in a witness may depend on the size of the processed decomposition, the multiplicity $\mu_{\mathbf{D}}(k, n)$ of \mathbf{D} may still be bounded by a function of k . And indeed, this is the case in typical FPT dynamic-programming algorithms operating on treelike tree decompositions.

Consider for instance the problem of computing the minimum vertex cover on a graph of treewidth at most k as described in the previous section. In our framework, a local witness (R, s) has size $k + \log n$ where k bits are used to represent a partial cover, and $\log n$ bits are used to represent the weight. Then, although the number of (\mathbf{D}, k, n) -useful local witnesses is $2^{k+\log n} = 2^k \cdot n$, the size of a (\mathbf{D}, k, n) -useful set can be bounded by 2^{k+1} because if (R, s) and (R, s') are two local witnesses, with the same partial cover R , but with distinct weights, then we only need to store the one with the smallest weight.

5.5 Inclusion Test

Let \mathbf{C} be a treelike decomposition class, and \mathbf{D} be a treelike DP-core. As discussed in the introduction, the problem of determining whether $\mathbb{G}[\mathbf{C}] \subseteq \mathbb{G}[\mathbf{D}, \mathbf{C}]$ can be regarded as a task in the realm of automated theorem proving. A width-based approach to testing whether this inclusion holds is to test for increasing values of k , whether the inclusion $\mathbb{G}[\mathbf{C}_k] \subseteq \mathbb{G}[\mathbf{D}, \mathbf{C}]$ holds. It turns out that if \mathbf{D} is \mathbf{C} -coherent, then testing whether $\mathbb{G}[\mathbf{C}_k] \subseteq \mathbb{G}[\mathbf{D}, \mathbf{C}]$ reduces to testing whether all \mathbf{C} -decompositions of width at most k are accepted by $\mathbf{D}[k]$, as stated in Lemma 23 below. We note that this is not necessarily true if \mathbf{D} is not \mathbf{C} -coherent.

Lemma 23. *Let $\mathbf{C} = \{(\Sigma_k, \mathbf{L}_k, \mathcal{G}_k)\}_{k \in \mathbb{N}}$ be a treelike decomposition class and \mathbf{D} be a \mathbf{C} -coherent treelike DP-core. Then, for each $k \in \mathbb{N}$, $\mathbb{G}[\mathbf{C}_k] \subseteq \mathbb{G}[\mathbf{D}, \mathbf{C}]$ if and only if $\mathbf{L}_k \subseteq \text{Acc}(\mathbf{D}[k])$.*

Proof. Suppose that $\mathbf{L}_k \subseteq \text{Acc}(\mathbf{D}[k])$. Let $G \in \mathbb{G}[\mathbf{C}_k]$. Then, there is some $k \in \mathbb{N}$, and some k -instructive tree decomposition $\tau \in \mathbf{L}_k$, such that $\mathcal{G}(\tau)$ is isomorphic to G . Since, by assumption, τ also belongs to $\text{Acc}(\mathbf{D}[k])$, we have that both $\mathcal{G}(\tau)$ and G belong to $\mathbb{G}[\mathbf{D}[k], \mathbf{C}]$. Therefore, G also belongs to $\mathbb{G}[\mathbf{D}, \mathbf{C}]$. Since G was chosen to be an arbitrary graph in $\mathbb{G}[\mathbf{C}_k]$, we have that $\mathbb{G}[\mathbf{C}_k] \subseteq \mathbb{G}[\mathbf{D}, \mathbf{C}]$. We note that for this direction we did not need the assumption that \mathbf{D} is \mathbf{C} -coherent.

For the converse, we do need the assumption that \mathbf{D} is coherent. Suppose that $\mathbb{G}[\mathbf{C}_k] \subseteq \mathbb{G}[\mathbf{D}, \mathbf{C}]$. Let $\tau \in \mathbf{L}_k$. Then, by the definition of graph property associated with a treelike decomposition class, we have that $\mathcal{G}(\tau)$ belongs to $\mathbb{G}[\mathbf{C}_k]$. Since, by assumption, $\mathcal{G}(\tau)$ also belongs to $\mathbb{G}[\mathbf{D}, \mathbf{C}]$, we have that there is some k' , and some k' -instructive tree decomposition τ' in $\text{Acc}(\mathbf{D}[k'])$ such that $\mathcal{G}(\tau')$ is isomorphic to $\mathcal{G}(\tau)$. But since \mathbf{D} is coherent, this implies that τ also belongs to $\text{Acc}(\mathbf{D}[k])$. Since τ was chosen to be an arbitrary treelike decomposition in \mathbf{L}_k , we have that $\mathbf{L}_k \subseteq \text{Acc}(\mathbf{D}[k])$. \square

Lemma 23 implies that if \mathbf{D} is coherent, then in order to show that $\mathbb{G}[\mathbf{C}_k] \not\subseteq \mathbb{G}[\mathbf{D}, \mathbf{C}]$ it is enough to show that there is some \mathbf{C} -decomposition τ of width at most k that belongs to \mathbf{L}_k but not to $\text{Acc}(\mathbf{D}[k])$. We will reduce this later task to the task of constructing a dynamic programming refutation (Definition 24).

Let \mathbf{C} be a decomposition class with automation \mathcal{A} , and let \mathbf{D} be a \mathbf{C} -coherent treelike DP-core. An $(\mathcal{A}, \mathbf{D}, k)$ -pair is a pair of the form (q, S) where q is a state of \mathcal{A}_k and $S \subseteq \mathbf{D}[k].\mathcal{W}$. We say that such pair (q, S) is $(\mathcal{A}, \mathbf{D}, k)$ -inconsistent if q is a final state of \mathcal{A}_k , but S has no final local witness for \mathbf{D} .

Definition 24 (DP-Refutation). *Let $\mathbf{C} = \{(\Sigma_k, \mathbf{L}_k, \mathcal{G}_k)\}_{k \in \mathbb{N}}$ be a decomposition class, \mathcal{A} be an automation for \mathbf{C} , \mathbf{D} be a \mathbf{C} -coherent treelike DP-core, and $k \in \mathbb{N}$. An $(\mathcal{A}, \mathbf{D}, k)$ -refutation is a sequence of $(\mathcal{A}, \mathbf{D}, k)$ -pairs*

$$R \equiv (q_1, S_1)(q_2, S_2) \dots (q_m, S_m)$$

satisfying the following conditions:

1. (q_m, S_m) is $(\mathcal{A}, \mathbf{D}, k)$ -inconsistent.

2. For each $i \in [m]$,

- (a) either $(q_i, S_i) = (q, \mathbf{D}[k].\hat{a})$ for some symbol a of arity 0 in Σ_k , and some state q such that $a \rightarrow q$ is a transition of \mathcal{A}_k , or
- (b) $(q_i, S_i) = (q, \mathbf{D}[k].\hat{a}(S_{j_1}, \dots, S_{j_{\mathfrak{r}(a)}}))$, for some $j_1, \dots, j_{\mathfrak{r}(a)} < i$, some symbol $a \in \Sigma_k$ of arity $\mathfrak{r}(a) > 0$, and some state q such that $a(q_{j_1}, \dots, q_{j_{\mathfrak{r}(a)}}) \rightarrow q$ is a transition of \mathcal{A}_k .

The following theorem shows that if \mathbf{C} is a decomposition class with automation \mathcal{A} and \mathbf{D} is a \mathbf{C} -coherent treelike DP-core, then showing that $\mathbb{G}[\mathbf{C}] \not\subseteq \mathbb{G}[\mathbf{D}, \mathbf{C}]$, is equivalent to showing the existence of some $(\mathcal{A}, \mathbf{D}, k)$ -refutation.

Theorem 25. Let $\mathcal{C} = \{(\Sigma_k, \mathcal{L}_k, \mathcal{G}_k)\}_{k \in \mathbb{N}}$ be a decomposition class with automation \mathcal{A} , and \mathcal{D} be a \mathcal{C} -coherent treelike DP-core. For each $k \in \mathbb{N}$, we have that $\mathbb{G}[\mathcal{C}_k] \not\subseteq \mathbb{G}[\mathcal{D}, \mathcal{C}]$ if and only if some $(\mathcal{A}, \mathcal{D}, k)$ -refutation exists.

Proof. Suppose that $\mathbb{G}[\mathcal{C}_k] \not\subseteq \mathbb{G}[\mathcal{D}, \mathcal{C}]$. Then, there is a graph G that belongs to $\mathbb{G}[\mathcal{C}_k]$, but not to $\mathbb{G}[\mathcal{D}, \mathcal{C}]$. Since $G \in \mathbb{G}[\mathcal{C}_k]$, we have that for some $\tau \in \mathcal{L}_k$, G is isomorphic to $\mathcal{G}(\tau)$. Since $G \notin \mathbb{G}[\mathcal{D}, \mathcal{C}]$ we have that $\tau \notin \text{Acc}(\mathcal{D})$, and therefore, $\tau \notin \text{Acc}(\mathcal{D}[k])$.

Let $\text{Sub}(\tau) = \{\sigma : \sigma \text{ is a subterm of } \tau\}$ be the set of subterms² of τ , and $\sigma_1, \sigma_2, \dots, \sigma_m$ be a topological ordering of the elements in $\text{Sub}(\tau)$. Since we ordered the subterms topologically, for each $i, j \in [m]$, if σ_i is a subterm of σ_j , then $i \leq j$. Additionally, $\sigma_m = \tau$. Now, consider the sequence

$$R \equiv (q_1, S_1)(q_2, S_2) \dots (q_m, S_m).$$

Since $\sigma_1, \sigma_2, \dots, \sigma_m$ are subterms of τ and ordered topologically, we have that for each $i \in [m]$, σ_i is either a symbol of arity zero or there is a symbol a of arity $\tau(a) > 0$, and $j_1, \dots, j_{\tau(a)} < i$ such that $\sigma_i = a(\sigma_{j_1}, \dots, \sigma_{j_{\tau(a)}})$.

- If $\sigma_i = a$ has arity zero, and q is the unique state of \mathcal{A}_k such that $a \rightarrow q$ is a transition of \mathcal{A}_k , then we set $(q_i, S_i) = (q, \mathcal{D}[k].\hat{a})$.
- If $\sigma_i = a(\sigma_{j_1}, \dots, \sigma_{j_{\tau(a)}})$ for some symbol a of arity $\tau(a) > 0$, and q is the unique state of \mathcal{A}_k such that $a(q_{j_1}, \dots, q_{j_{\tau(a)}}) \rightarrow q$ is a transition of \mathcal{A}_k , we set $(q_i, S_i) = (q, \mathcal{D}[k].\hat{a}(S_{j_1}, \dots, S_{j_{\tau(a)}}))$.

By construction, R satisfies Condition 2 of Definition 24. Now, we know that $\tau \in \mathcal{L}_k$ and that q_m is the state reached by τ in \mathcal{A}_k , and therefore, q_m is a final state. On the other hand, $S_m = \Gamma[\mathcal{D}, k](\tau)$ and $\tau \notin \text{Acc}(\mathcal{D}[k])$, and therefore, S_m has no final local witness. Therefore, the pair (q_m, S_m) is an $(\mathcal{A}, \mathcal{D}, k)$ -inconsistent pair, so Condition 1 of Definition 24 is satisfied. Consequently, the first direction of Theorem 25 is proved, i.e., R is an $(\mathcal{A}, \mathcal{D}, k)$ -refutation.

For the converse, assume that $R \equiv (q_1, S_1)(q_2, S_2) \dots (q_m, S_m)$ is an $(\mathcal{A}, \mathcal{D}, k)$ -refutation. Using this refutation, we will construct a sequence of terms $\sigma_1, \sigma_2, \dots, \sigma_m$ with the following property: for each $i \in [m]$, $\sigma_i \in \text{Terms}(\Sigma_k)$ and $S_i = \Gamma[\mathcal{D}, k](\sigma_i)$. Since q_m is a final state for \mathcal{A}_k but S_m has no final local witness for $\mathcal{D}[k]$, we have that σ_m is in \mathcal{L}_k but not in $\text{Acc}(\mathcal{D}[k])$. In other words, $\mathcal{G}(\sigma_m)$ is in $\mathbb{G}[\mathcal{C}_k]$ but not in $\mathbb{G}[\mathcal{D}[k], \mathcal{C}]$. Since \mathcal{D} is \mathcal{C} -coherent, we have that for each $k' \in \mathbb{N}$, there is no term $\sigma' \in \text{Acc}(\mathcal{D}[k'])$ with $\mathcal{G}(\sigma_m) \simeq \mathcal{G}(\sigma')$ (otherwise, σ_m would belong to $\text{Acc}(\mathcal{D}[k])$). Therefore, $\mathcal{G}(\sigma_m)$ is not in $\mathbb{G}[\mathcal{D}, \mathcal{C}]$ either. We infer that $\mathbb{G}[\mathcal{C}_k] \not\subseteq \mathbb{G}[\mathcal{D}, \mathcal{C}]$.

Now, for each $i \in \mathbb{N}$, the construction of σ_i proceeds as follows. If there is a symbol $a \in \Sigma_k$ of arity 0 such that $a \rightarrow q_i$ is a transition of \mathcal{A}_k , and $S_i = \mathcal{D}[k].\hat{a}$, then we let $\sigma_i = a$. On the other hand, if there is a symbol $a \in \Sigma_k$ of arity $\tau(a) > 0$ and some $j_1, \dots, j_r < i$ such that $a(q_{j_1}, \dots, q_{j_{\tau(a)}}) \rightarrow q_i$ is a transition of \mathcal{A}_k and $S_i = \mathcal{D}[k].\hat{a}(S_{j_1}, \dots, S_{j_{\tau(a)}})$, then we let $\sigma_i = a(\sigma_{j_1}, \dots, \sigma_{j_{\tau(a)}})$. It should be clear that for each $i \in [m]$, σ_i is a term in $\text{Terms}(\Sigma_k)$. Furthermore, using Definition 9, it follows by induction on i that for each $i \in [m]$, $S_i = \Gamma[\mathcal{D}, k](\sigma_i)$. This concludes the proof of the theorem. \square

The proof of Theorem 25 provides us with an algorithm to extract, from a given $(\mathcal{A}, \mathcal{D}, k)$ -refutation R , a \mathcal{C} -decomposition τ of width at most k such that $\mathcal{G}(\tau) \notin \mathbb{G}[\mathcal{D}, \mathcal{C}]$. The graph $\mathcal{G}(\tau)$ corresponding to τ may be regarded as a counter-example for the conjecture $\mathbb{G}[\mathcal{C}] \subseteq \mathbb{G}[\mathcal{D}, \mathcal{C}]$. Note that the minimum height of such a \mathcal{C} -decomposition τ is upper-bounded by the minimum length of a $(\mathcal{A}, \mathcal{D}, k)$ -refutation. Consequently, if \mathcal{C} is a treelike decomposition class of arity r , then τ has at most m nodes, if $r = 1$, and at most $\frac{r^m - 1}{r - 1}$ nodes, if $r > 1$.

Corollary 26. Let $\mathcal{C} = \{(\Sigma_k, \mathcal{L}_k, \mathcal{G}_k)\}_{k \in \mathbb{N}}$ be a treelike decomposition class of arity r with automation \mathcal{A} , \mathcal{D} be a \mathcal{C} -coherent treelike DP-core, and $R \equiv (q_1, S_1)(q_2, S_2) \dots (q_m, S_m)$ be a $(\mathcal{A}, \mathcal{D}, k)$ -refutation. Then, there is a \mathcal{C} -decomposition $\tau \in \mathcal{L}_k$, such that $\mathcal{G}(\tau) \in \mathbb{G}[\mathcal{C}_k] \setminus \mathbb{G}[\mathcal{D}, \mathcal{C}]$, τ has height at most $m - 1$, and size $|\tau|$ at most m , if $r = 1$ and at most $\frac{r^m - 1}{r - 1}$, if $r > 1$.

²Note that $|\text{Sub}(\tau)|$ may be smaller than $|\tau|$, since a given subterm may occur in several positions of τ .

Theorem 25 also implies the existence of a simple forward-chaining style algorithm for determining whether $\mathbb{G}[\mathbf{C}_k] \subseteq \mathbb{G}[\mathbf{D}, \mathbf{C}]$ when \mathbf{D} is a *finite* and \mathbf{C} -coherent treelike DP-core.

Theorem 27 (Inclusion Test). *Let \mathbf{C} be a treelike decomposition class of complexity $f(k)$ and arity r , and let \mathbf{D} be a finite, internally polynomial \mathbf{C} -coherent treelike DP core. One can determine whether $\mathbb{G}[\mathbf{C}_k] \subseteq \mathbb{G}[\mathbf{D}, \mathbf{C}]$ in time*

$$f(k)^{O(r)} \cdot 2^{O(r \cdot \beta_{\mathbf{D}}(k) \cdot \mu_{\mathbf{D}}(k))} \leq f(k)^{O(r)} \cdot 2^{r \cdot 2^{O(\beta_{\mathbf{D}}(k))}}.$$

Proof. Using breadth-first search, we successively enumerate the $(\mathcal{A}, \mathbf{D}, k)$ -pairs that can be derived using rule 2 of Definition 24. We may assume that the first traversed pairs are those in the set

$$\{(q, \mathbf{D}[k].\hat{a}) \mid \tau(a) = 0, a \rightarrow q \text{ is a transition of } \mathcal{A}_k\},$$

which correspond to symbols of arity 0. We run this search until either an inconsistent $(\mathcal{A}, \mathbf{D}, k)$ -pair has been reached, or until there is no $(\mathcal{A}, \mathbf{D}, k)$ -pair left to be enumerated. In the first case, the obtained list of pairs

$$R = (q_1, S_1)(q_2, S_2) \dots (q_m, S_m) \quad (2)$$

is, by construction, a $(\mathcal{A}, \mathbf{D}, k)$ -refutation since (q_m, S_m) is inconsistent and for each $i \in [m]$, (q_i, S_i) has been obtained by applying the rule 2 of Definition 24. In the second case, no such refutation exists. More specifically, in the beginning of the process, R is the empty list, and the pairs $\{(\mathcal{A}_k.a, \mathbf{D}[k].\hat{a}) : \tau(a) = 0\}$ are added to a *buffer set* Y . While Y is non-empty, we delete an arbitrary pair (q, S) from Y and append this pair to R . If the pair is inconsistent, we have constructed a refutation R , and therefore, we return R . Otherwise, for each (q', S') that can be obtained from (q, S) using the rule 2 of Definition 24 (together with another pair from R), we insert (q', S') to the buffer Y provided this pair is not already in R . We repeat this process until either R has been returned or until Y is empty. In this case, we conclude that $\mathbb{G}[\mathbf{C}_k] \subseteq \mathbb{G}[\mathbf{D}[k], \mathbf{C}]$, and return *Inclusion Holds*. This construction is detailed in Algorithm 1.

Now an upper bound on the running time of the algorithm can be established as follows. Suppose that \mathbf{C} has complexity $f(k)$ for some $f : \mathbb{N} \rightarrow \mathbb{N}$. First, we note that there are at most $f(k) \cdot \delta_{\mathbf{D}}(k)$ pairs of the form (q, S) where q is a state of \mathcal{A}_k , and $S \subseteq \mathbf{D}[k].\mathcal{W}$. Furthermore, since \mathbf{C} has arity at most r , the creation of a new pair may require the analysis of at most $f(k)^r \cdot \delta_{\mathbf{D}}(k)^r$ tuples of previously created pairs. For each such a tuple $[(q_1, S_1), (q_2, S_2), \dots, (q_r, S_r)]$ with $r' \leq r$, the computation of the state q' from the tuple (q_1, \dots, q_r) takes time $f(k)^{O(1)}$ and, as argued in the proof of Theorem 22, the computation of the set $S' = \mathbf{D}[k].\hat{a}(S_{j_1}, \dots, S_{j_{\tau(a)}})$ from the tuple $(S_{j_1}, \dots, S_{j_{\tau(a)}})$ takes time $k^{O(1)} \cdot r^{O(1)} \cdot \beta_{\mathbf{D}}(k)^{O(1)} \cdot \mu_{\mathbf{D}}(k)^{r+O(1)}$. Therefore, the whole process takes time at most

$$k^{O(1)} \cdot r^{O(1)} \cdot \beta_{\mathbf{D}}(k)^{O(1)} \cdot \mu_{\mathbf{D}}(k)^{r+O(1)} \cdot f(k)^{r+O(1)} \cdot \delta_{\mathbf{D}}(k)^{r+O(1)}. \quad (3)$$

Since, by assumption, $f(k) \geq k$ and by Observation 15, $\mu_{\mathbf{D}}(k) \leq 2^{\beta_{\mathbf{D}}(k)}$, and $\delta_{\mathbf{D}}(k) \leq 2^{\beta_{\mathbf{D}}(k) \cdot \mu_{\mathbf{D}}(k)+1}$, we have that Expression 3 can be simplified to

$$f(k)^{O(r)} \cdot 2^{O(r \cdot \beta_{\mathbf{D}}(k) \cdot \mu_{\mathbf{D}}(k))} \leq f(k)^{O(r)} \cdot 2^{r \cdot 2^{O(\beta_{\mathbf{D}}(k))}}.$$

□

Since the search space in the proof of Theorem 27 has at most $f(k) \cdot \delta_{\mathbf{D}}(k)$ distinct $(\mathcal{A}, \mathbf{D}, k)$ -pairs, a minimum-length $(\mathcal{A}, \mathbf{D}, k)$ -refutation has length at most $f(k) \cdot \delta_{\mathbf{D}}(k)$. Therefore, this fact together with Corollary 26 implies the following result.

Corollary 28. *Let $\mathbf{C} = \{(\Sigma_k, \mathbf{L}_k, \mathcal{G}_k)\}_{k \in \mathbb{N}}$ be a treelike decomposition class of complexity $f(k)$, and let \mathbf{D} be a finite, \mathbf{C} -coherent treelike DP core. If $\mathbb{G}[\mathbf{C}_k] \not\subseteq \mathbb{G}[\mathbf{D}, \mathbf{C}]$, then there is a \mathbf{C} -decomposition $\tau \in \mathbf{L}_k$, such that $\mathcal{G}(\tau) \in \mathbb{G}[\mathbf{C}_k] \setminus \mathbb{G}[\mathbf{D}, \mathbf{C}]$, τ has height at most $f(k) \cdot \delta_{\mathbf{D}}(k) - 1$, and size $|\tau|$ at most $f(k) \cdot \delta_{\mathbf{D}}(k)$, if $r = 1$ and at most $r^{f(k) \cdot \delta_{\mathbf{D}}(k)}$ if $r > 1$.*

The requirement that the DP-core D in Theorem 27 is finite can be relaxed if instead of asking whether $\mathbb{G}[C_k] \subseteq \mathbb{G}[D, C]$, we ask whether all graphs in $\mathbb{G}[C_k]$ that can be represented by a C -decomposition of size at most n belong to $\mathbb{G}[D, C]$.

Corollary 29 (Bounded-Size Inclusion Test). *Let C be a treelike decomposition class of complexity $f(k)$ and arity r , and let D be a (not necessarily finite) internally polynomial C -coherent treelike DP core. One can determine in time*

$$f(k)^{O(r)} \cdot 2^{O(r \cdot \beta_D(k, n) \cdot \mu_D(k, n))} \cdot n^{O(1)}$$

whether every graph corresponding to a C -decomposition of width at most k and size at most n belongs to $\mathbb{G}[D, C]$.

Proof. The proof is identical to the proof of Theorem 27, except that instead of performing a BFS over the space of pairs of the form (q, S) , we perform a BFS over the space of triples of the form (q, S, i) where $i \in \{0, 1, \dots, n\}$. More specifically, the BFS will enumerate all such triples with the property that there is a term $\tau \in \text{Terms}(\Sigma_k)$ such that q is the state reached in \mathcal{A}_k after reading τ , $S = \Gamma[D, k](\tau)$, and $i = |\tau|$. If during the enumeration, one finds a triple (q, S, i) where (q, S) is an inconsistent (\mathcal{A}, D, k) -pair, then we know that a counter-example of width at most k and size at most n exists, and this counter-example can be constructed by backtracking. Otherwise, no such a counter-example exists. \square

We note that whenever $\mu_D(k, n) = h_1(k)$ for some function $h_1 : \mathbb{N} \rightarrow \mathbb{N}$, and $\beta_D(k, n) = h_2(k) \cdot \log n$ for some function $h_2 : \mathbb{N} \rightarrow \mathbb{N}$, then the running time stated in Corollary 29 is of the form $n^{h_3(k)}$ for some function $h_3 : \mathbb{N} \rightarrow \mathbb{N}$. This is significantly faster than the naive brute-force approach of enumerating all terms of width at most k , and size at most n , and subsequently testing whether these terms belong to $\mathbb{G}[D, C]$.

Algorithm 1: Inclusion Test

Input : An automation for \mathcal{A} of C , a finite, C -coherent treelike DP-core D , and an integer $k \in \mathbb{N}$.

Output: An (\mathcal{A}, D, k) -refutation R if $\mathbb{G}[C_k] \not\subseteq \mathbb{G}[D[k], C]$, and "Inclusion Holds", otherwise.

```

1  $R \leftarrow []$  ; /*  $[]$  is the empty list. */
2  $\hat{R} \leftarrow \{ \}$  ; /*  $\{ \}$  is the empty set. */
3  $Y \leftarrow \{(q, D.\hat{a}) : \tau(a) = 0, a \rightarrow q \text{ is a transition of } \mathcal{A}_k.\}$ ;
4 InconsistentPair  $\leftarrow$  false;
5 while  $Y \neq \emptyset$  and InconsistentPair = false do
6   Remove some pair  $(q, S)$  from  $Y$ , append it to  $R$ , and insert it in  $\hat{R}$ ;
7   if  $(q, S)$  is an inconsistent  $(\mathcal{A}, D, k)$ -pair, then
8     return  $R$  ; /* In this case,  $R$  is a  $(\mathcal{A}, D, k)$ -refutation. */
9   else
10     foreach  $a \in \Sigma_k$  do
11       foreach sequence  $(q_1, S_1) \dots (q_{\tau(a)}, S_{\tau(a)})$  of pairs  $\in \hat{R}^{\tau(a)}$  having an occurrence
          of  $(q, S)$  do
12          $(q', S') \leftarrow (a(q_1, q_2, \dots, q_{\tau(a)}), D[k].\hat{a}(S_1, S_2, \dots, S_{\tau(a)}))$ ;
13         if  $(q', S')$  is not in  $\hat{R}$  then  $Y \leftarrow Y \cup \{(q', S')\}$  ;
14 return "Inclusion Holds."

```

5.6 Combinators and Combinations

Given a graph property \mathbb{P} , and a graph $G \in \text{GRAPHS}$, we let $\mathbb{P}(G)$ denote the Boolean value *true* if $G \in \mathbb{P}$ and the value *false*, if $G \notin \mathbb{P}$.

Definition 30 (Combinators). *Let $\ell \in \mathbb{N}$. An ℓ -combinator is a function*

$$\mathcal{C} : \{0, 1\}^\ell \times (\{0, 1\}^*)^\ell \rightarrow \{0, 1\}.$$

Given graph properties $\mathbb{P}_1, \dots, \mathbb{P}_\ell$ and graph invariants $\mathcal{I}_1, \dots, \mathcal{I}_\ell$, we define the following graph property:

$$\hat{\mathcal{C}}(\mathbb{P}_1, \dots, \mathbb{P}_\ell, \mathcal{I}_1, \dots, \mathcal{I}_\ell) = \{G : \mathcal{C}(\mathbb{P}_1(G), \dots, \mathbb{P}_\ell(G), \mathcal{I}_1(G), \dots, \mathcal{I}_\ell(G)) = 1\}.$$

We say that \mathcal{C} is polynomial if it can be computed in time $O_\ell(|X|^c)$ for some constant c on any given input X .

Intuitively, a combinator is a tool to define graph classes in terms of previously defined graph classes and previously defined graph invariants. It is worth noting that Boolean combinations of graph classes can be straightforwardly defined using combinators. Nevertheless, one can do more than that, since combinators can also be used to establish relations between graph invariants. For instance, using combinators one can define the class of graphs whose *covering number* (the smallest size of a vertex-cover) is equal to the *dominating number* (the smallest size of a dominating set). This is just a illustrative example. Other examples of invariants that can be related using combinators are: *clique number*, *independence number*, *chromatic number*, *diameter*, and many others. Next, we will use combinators as a tool to combine graph properties and graph invariants defined using DP-cores. Given a DP-core D , and a finite subset $S \subseteq D.\mathcal{W}$, we let $F(D, S)$ be the Boolean value 1 if S contains some final witness for D , and the value 0, otherwise.

Theorem 31. *Let \mathcal{C} be an ℓ -combinator, \mathcal{C} be a treelike decomposition class, and D_1, \dots, D_ℓ be \mathcal{C} -coherent treelike DP-cores. Then, there exists a \mathcal{C} -coherent treelike DP-core $D = D(\mathcal{C}, D_1, \dots, D_\ell)$ satisfying the following properties:*

1. $\mathbb{G}[D, \mathcal{C}] = \mathcal{C}(\mathbb{G}[D_1, \mathcal{C}], \dots, \mathbb{G}[D_\ell, \mathcal{C}], \mathcal{I}[D_1, \mathcal{C}], \dots, \mathcal{I}[D_\ell, \mathcal{C}])$.
2. D has bitlength $\beta_D(k, n) = \sum_{i=1}^\ell \beta_{D_i}(k, n) \cdot \mu_{D_i}(k, n)$.
3. D has multiplicity $\mu_D(k, n) = 1$.
4. D has deterministic state complexity $\delta_D(k, n) = \nu_D(k, n) \leq \prod_{i=1}^\ell \delta_{D_i}(k, n)$.

Proof. We let the \mathcal{C} -combination of (D_1, \dots, D_ℓ) be the DP-core D whose components are specified below. Here, we let $u, v \in [k+1]$, and $\mathbf{S} = (S_1, \dots, S_\ell)$ and $\mathbf{S}' = (S'_1, \dots, S'_\ell)$ be tuples in $\mathcal{P}_{\text{fin}}(D_1.\mathcal{W}) \times \dots \times \mathcal{P}_{\text{fin}}(D_\ell.\mathcal{W})$.

1. $D.\mathcal{W} = \mathcal{P}_{\text{fin}}(D_1.\mathcal{W}) \times \dots \times \mathcal{P}_{\text{fin}}(D_\ell.\mathcal{W})$.
2. $D.\text{Leaf} = \{(D_1.\text{Leaf}, \dots, D_\ell.\text{Leaf})\}$.
3. $D.\text{IntroVertex}\{u\}(\mathbf{S}) = \{(D_1.\text{IntroVertex}\{u\}(S_1), \dots, D_\ell.\text{IntroVertex}\{u\}(S_\ell))\}$.
4. $D.\text{ForgetVertex}\{u\}(\mathbf{S}) = \{(D_1.\text{ForgetVertex}\{u\}(S_1), \dots, D_\ell.\text{ForgetVertex}\{u\}(S_\ell))\}$.
5. $D.\text{IntroEdge}\{u, v\}(\mathbf{S}) = \{(D_1.\text{IntroEdge}\{u, v\}(S_1), \dots, D_\ell.\text{IntroEdge}\{u, v\}(S_\ell))\}$.
6. $D.\text{Join}(\mathbf{S}, \mathbf{S}') = \{(D_1.\text{Join}(S_1, S'_1), \dots, D_\ell.\text{Join}(S_\ell, S'_\ell))\}$.
7. $D.\text{Clean}(\{\mathbf{S}\}) = \{(D_1.\text{Clean}(S_1), \dots, D_\ell.\text{Clean}(S_\ell))\}$.
8. $D.\text{Final}(\mathbf{S}) = \mathcal{C}(F(D_1, S_1), \dots, F(D_\ell, S_\ell), D_1.\text{Inv}(S_1), \dots, D_\ell.\text{Inv}(S_\ell))$.
9. $D.\text{Inv}(\{\mathbf{S}\}) = (D_1.\text{Inv}(S_1), \dots, D_\ell.\text{Inv}(S_\ell))$.

The upper bounds for the functions $\beta_D(k, n)$, $\mu_D(k, n)$, $\nu_D(k, n)$ and $\delta_D(k, n)$, can be inferred directly from this construction. The fact that D is \mathcal{C} -coherent follows also immediately from the fact that D_1, \dots, D_ℓ are \mathcal{C} -coherent. Note that $\delta_D(k, n) = \nu_D(k, n)$, since $\mu_D(k, n) = 1$. \square

We call the DP-core $D = D(\mathcal{C}, D_1, \dots, D_\ell)$ the \mathcal{C} -combination of D_1, \dots, D_ℓ . As a corollary of Theorem 31 and Theorem 22, we have the following theorem relating the complexity of model-checking the graph property defined by D to the complexity of model-checking the properties defined by D_1, \dots, D_ℓ .

Theorem 32 (Model Checking for Combinations). *Let \mathcal{C} be a treelike decomposition class of arity r ; D_1, \dots, D_ℓ be internally polynomial, \mathcal{C} -coherent treelike DP-cores; and \mathcal{C} be a polynomial ℓ -combinator. Let $D = D(\mathcal{C}, D_1, \dots, D_\ell)$ be the \mathcal{C} -combination of D_1, \dots, D_ℓ , $\beta(k, n) = \max_i \beta_{D_i}(k, n)$ and $\mu(k, n) = \max_i \mu_{D_i}(k, n)$. Then, given a \mathcal{C} -decomposition τ of width at most k , and size $|\tau| = n$, one can determine whether $\mathcal{G}(\tau) \in \mathbb{G}[D, \mathcal{C}]$ in time*

$$n \cdot \ell \cdot k^{O(1)} \cdot r^{O(1)} \cdot \beta(k, n)^{O(1)} \cdot \mu(k, n)^{r+O(1)} + O_\ell(\beta(k, n)^{O(1)}).$$

If the DP-cores D_1, \dots, D_ℓ are also *finite*, besides being \mathcal{C} -coherent, and internally polynomial, then Theorem 31 together with Theorem 27 directly imply the following theorem.

Theorem 33 (Inclusion Test for Combinations). *Let \mathcal{C} be a treelike decomposition class of arity r ; D_1, \dots, D_ℓ be finite, internally polynomial, \mathcal{C} -coherent treelike DP-cores; and \mathcal{C} be a polynomial ℓ -combinator. Let $D = D(\mathcal{C}, D_1, \dots, D_\ell)$ be the \mathcal{C} -combination of D_1, \dots, D_ℓ , $\beta(k) = \max_i \beta_{D_i}(k)$ and $\mu(k) = \max_i \mu_{D_i}(k)$. Then, for each $k \in \mathbb{N}$, one can determine whether $\mathbb{G}[C_k] \subseteq \mathbb{G}[D, \mathcal{C}]$ in time*

$$f(k)^{O(r)} \cdot 2^{O(\ell \cdot r \cdot \beta(k) \cdot \mu(k))} \leq f(k)^{O(r)} \cdot 2^{\ell \cdot r \cdot 2^{O(\beta(k))}}.$$

We note that in typical applications, the parameters r and ℓ are constant, while the growth of the function $f(k)$ is negligible when compared with $2^{O(\beta(k) \cdot \mu(k))}$. Therefore, in these applications, the running time of our algorithm is of the form $2^{O(\beta(k) \cdot \mu(k))} \leq 2^{2^{O(\beta(k))}}$. It is also worth noting that if $\mathbb{G}[C_k] \not\subseteq \mathbb{G}[D, \mathcal{C}]$, then a term τ of height at most $2^{O(\beta(k) \cdot \mu(k))}$ encoding a graph in $\mathbb{G}[C_k] \setminus \mathbb{G}[D, \mathcal{C}]$ can be constructed (see Corollary 28).

6 Applications of Theorem 33

In this section, we illustrate the applicability of Theorem 33 to the realm of automated theorem proving. In our examples, we focus on the width measure *treewidth*, given that this is by far the most well studied width measure for graphs. Below, we list several graph properties that can be decided by finite, internally polynomial, ITD-coherent treelike DP-cores. Here, ITD is the class of instructive tree decompositions introduced in Section 4. This class has complexity 2^k .

1. **Simple**: the set of all simple graphs (i.e. without multiedges).
2. **MaxDeg $_{\geq}(c)$** : the set of graphs containing at least one vertex of degree at least c .
3. **MinDeg $_{\leq}(c)$** : the set of graphs containing at least one vertex of degree at most c .
4. **Colorable(c)**: the set of chromatic number at most c .
5. **Conn**: set of connected graphs.
6. **VConn $_{\leq}(c)$** : the set of graphs with vertex-connectivity at most c . A graph is c -vertex-connected if it has at least c vertices, and if it remains connected whenever fewer than c vertices are deleted.

7. **EConn $_{\leq}(c)$** : the set of graphs with edge-connectivity at most c . A graph is c -edge-connected if it remains connected whenever fewer than c edges are deleted.
8. **Hamiltonian**: the set of Hamiltonian graphs. A graph is Hamiltonian if it contains a cycle that spans all its vertices.
9. **NZFlow(\mathbb{Z}_m)**: set of graphs that admit a \mathbb{Z}_m -flow. Here, $\mathbb{Z}_m = \{0, \dots, m-1\}$ is the set of integers modulo m . A graph G admits a nowhere-zero \mathbb{Z}_m -flow if one can assign to each edge an orientation and a non-zero element of \mathbb{Z}_m in such a way that for each vertex, the sum of values associated with edges entering the vertex is equal to the sum of values associated with edges leaving the vertex.
10. **Minor(H)**: the set of graphs containing H as a minor. A graph H is a *minor* of a graph G if H can be obtained from G by a sequence of vertex/edge deletions and edge contractions.

Theorem 34. *Let ITD be the instructive tree decomposition class defined in Section 4. The properties specified above have ITD-coherent DP-cores with complexity parameters (bitlength β , multiplicity μ , state complexity ν , deterministic state complexity δ) as specified in Table 1.*

Property	$\beta(k)$	$\mu(k)$	$\nu(k)$	$\delta(k)$
Simple	$O(k^2)$	1	$2^{O(\beta(k))}$	$2^{O(\beta(k))}$
MaxDeg$_{\geq}(c)$	$O(k \cdot \log c)$	1	$2^{O(\beta(k))}$	$2^{O(\beta(k))}$
MinDeg$_{\leq}(c)$	$O(k \cdot \log c)$	1	$2^{O(\beta(k))}$	$2^{O(\beta(k))}$
Colorable(c)	$O(k \log c)$	$2^{O(\beta(k))}$	$2^{O(\beta(k))}$	$2^{2^{O(\beta(k))}}$
Conn	$O(k \log k)$	$2^{O(\beta(k))}$	$2^{O(\beta(k))}$	$2^{2^{O(\beta(k))}}$
VConn(c)	$O(\log c + k \log k)$	$2^{O(\beta(k))}$	$2^{O(\beta(k))}$	$2^{2^{O(\beta(k))}}$
EConn(c)	$O(\log c + k \log k)$	$2^{O(\beta(k))}$	$2^{O(\beta(k))}$	$2^{2^{O(\beta(k))}}$
Hamiltonian	$O(k \log k)$	$2^{O(k)}$	$2^{O(\beta(k))}$	$2^{2^{O(k)}}$
NZFlow(\mathbb{Z}_m)	$O(k \log m)$	$2^{O(\beta(k))}$	$2^{O(\beta(k))}$	$2^{2^{O(\beta(k))}}$
Minor(H)	$O(k \log k + V_H + E_H)$	$2^{\beta(k)}$	$2^{O(\beta(k))}$	$2^{2^{O(\beta(k))}}$

Table 1: Complexity measures for DP-cores deciding several graph properties.

The proof of Theorem 34 can be found in Appendix E. Note that in the case of the DP-core **C-Hamiltonian** the multiplicity $2^{O(k)}$ is smaller than the trivial upper bound of $2^{O(k \cdot \log k)}$ and consequently, the deterministic state complexity $2^{2^{O(k)}}$ is smaller than the trivial upper bound of $2^{2^{O(k \cdot \log k)}}$. We note that the proof of this fact is a consequence of the rank-based approach developed in [15]. Next, we will show how Theorem 33 together with Theorem 34 can be used to provide double-exponential upper bounds on the time necessary to verify long-standing graph-theoretic conjectures on graphs of treewidth at most k . If such a conjecture is false, then Corollary 28 can be used to establish an upper bound on minimum height of a term representing a counterexample for the conjecture.

Hadwiger's Conjecture. This conjecture states that for each $c \geq 1$, every graph with no K_{c+1} -minor has a c -coloring [41]. This conjecture, which suggests a far reaching generalization of the 4-colors theorem, is considered to be one of the most important open problems in graph theory. The conjecture has been resolved in the positive for the cases $c < 6$ [69], but remains open for each value of $c \geq 6$. By Theorem 34, **Colorable(c)** has DP-cores of deterministic state complexity $2^{2^{O(k \log c)}}$, while **Minor(K_{c+1})** has DP-cores of deterministic state complexity $2^{2^{O(k \log k + c^2)}}$. Therefore, by using Theorem 33, we have that the case c of Hadwiger's conjecture can be tested in time $2^{2^{O(k \log k + c^2)}}$ on graphs of treewidth at most k .

Using the fact for each fixed $c \in \mathbb{N}$, both the existence of K_{c+1} -minors and the existence of c -colorings are MSO-definable, together with the fact that the MSO theory of graphs of bounded

treewidth is decidable, one can infer that for each $c, k \in \mathbb{N}$, one can determine whether the case c of Hadwiger’s conjecture is true on graphs of treewidth at most k in time $f(c, k)$ for some computable function $f : \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$. Using Courcelle’s approach, one can estimate the growth of f as a tower of exponentials of height at most 10. In [51] Karawabayshi have estimated that $f(c, k) \leq p^{p^p}$, where $p = (k + 1)^{(c-1)}$. It is worth noting that our estimate of $2^{2^{O(k \log k + c^2)}}$ obtained by a combination Theorem 33 and Theorem 34 improves significantly on both the estimate obtained using the MSO approach and the estimate provided in [51].

Tutte’s Flow Conjectures. Tutte’s 5-flow, 4-flow, and 3-flow conjectures are some of the most well studied and important open problems in graph theory. The 5-flow conjecture states that every bridgeless graph G has a \mathbb{Z}_5 -flow. This conjecture is true if and only if every 2-edge-connected graph has a \mathbb{Z}_5 -flow [77]. By Theorem 34, both $\text{ECon}(2)$ and $\text{NZFlow}(\mathbb{Z}_5)$ have coherent DP-cores of deterministic state complexity $2^{2^{O(k \log k)}}$. Since Tutte’s 5-flow conjecture can be expressed in terms of a Boolean combination of these properties, we have that this conjecture can be tested on graphs of treewidth at most k in time $2^{2^{O(k \log k)}}$ on graphs of treewidth at most k . The 4-flow conjecture states that every bridgeless graph with no Petersen minor has a nowhere-zero 4-flow [80]. Since this conjecture can be formulated using a Boolean combination of the properties $\text{ECon}(2)$, $\text{Minor}(P)$ (where P is the Petersen graph), and $\text{NZFlow}(\mathbb{Z}_4)$, we have that this conjecture can be tested in time $2^{2^{O(k \log k)}}$ on graphs of treewidth at most k . Finally, Tutte’s 3-Flow conjecture states that every 4-edge connected graph has a nowhere-zero 3-flow [38]. Similarly to the other cases it can be expressed as a Boolean combination of $\text{ECon}(4)$ and $\text{NZFlow}(\mathbb{Z}_3)$. Therefore, it can be tested in time $2^{2^{O(k \log k)}}$ on graphs of treewidth at most k .

Barnette’s Conjecture. This conjecture states that every 3-connected, 3-regular, bipartite, planar graph is Hamiltonian. Since a graph is bipartite if and only if it is 2-colorable, and since a graph is planar if and only if it does not contain K_5 or $K_{3,3}$ as minors, Barnette’s conjecture can be stated as a combination of the cores $\text{VCon}(3)$, $\text{MaxDeg}_{\geq}(3)$, $\text{MinDeg}_{\leq}(3)$, $\text{Colorable}(2)$, $\text{Minor}(K_5)$ and $\text{Minor}(K_{3,3})$. Therefore, by Theorem 33, it can be tested in time $2^{2^{O(k \log k)}}$ on graphs of treewidth at most k .

7 Conclusion and Future Directions

In this work, we have introduced a general and modular framework that allows one to combine width-based dynamic programming algorithms for model-checking graph-theoretic properties into algorithms that can be used to provide a width-based attack to long-standing conjectures in graph theory. By generality, we mean that our framework can be applied with respect to any treelike width measure (Definition 3), including treewidth [13], cliquewidth [32], and many others [54, 76, 24, 75, 4]. By modularity, we mean that dynamic programming cores may be developed completely independently from each other as if they were plugins, and then combined together either with the purpose of model-checking more complicated graph-theoretic properties, or with the purpose of attacking a given graph theoretic conjecture.

As a concrete example, we have shown that the validity of several longstanding graph theoretic conjectures can be tested on graphs of treewidth at most k in time double exponential in $k^{O(1)}$. This upper bound follow from Theorem 33 together with upper bounds established on the bitlength/multiplicity of DP-cores deciding several well studied graph properties. Although still high, this upper bound improves significantly on approaches based on quantifier elimination. This is an indication that the expertise accumulated by parameterized complexity theorists in the development of more efficient width-based DP algorithms for model checking graph-theoretic properties have also relevance in the context of automated theorem proving. It is worth noting that this is the case even for graph properties that are computationally easy, such as connectivity and bounded degree.

For simplicity, have defined the notion of a treelike width measure (Definition 3) with respect to graphs. Nevertheless, this notion can be directly lifted to more general classes of relational structures. More specifically, given a class \mathfrak{R} of relational structures over some signature \mathfrak{s} , we define the notion of a *treelike \mathfrak{R} -decomposition class* as a sequence of triples $\mathcal{C} = \{(\Sigma_k, \mathsf{L}_k, \mathcal{G}_k)\}_{k \in \mathbb{N}}$ precisely as in Definition 2, with the only exception that now, \mathcal{G}_k is a function from L_k to \mathfrak{R} . With this adaptation, and by letting the relation \simeq in Definition 11 denote isomorphism between \mathfrak{s} -structures, all results in Section 5 generalize smoothly to relational structures from \mathfrak{R} . This generalization is relevant because it shows that our notion of width-based automated theorem proving can be extended to a much larger context than graph theory.

In the field of parameterized complexity theory, the irrelevant vertex technique is a set of theoretical tools [3, 71] that can be used to show that for certain graph properties \mathbb{P} there is a constant $K_{\mathbb{P}}$, such that if G is a graph of treewidth at least $K_{\mathbb{P}}$ then it contains an *irrelevant vertex* for \mathbb{P} . More specifically, there is a vertex x such that G belongs to \mathbb{P} if and only if the graph $G \setminus x$ obtained by deleting x from G belongs to \mathbb{P} . This tool, that builds on Robertson and Seymour’s celebrated excluded grid theorem [71, 25] and on the flat wall theorem [26, 52, 72], has found several applications in structural graph theory and in the development of parameterized algorithms [2, 71, 34, 50, 48]. Interestingly, the irrelevant vertex technique has also theoretical relevance in the framework of width-based automated theorem proving. More specifically, the existence of vertices that are irrelevant for \mathbb{P} on graphs of treewidth at least $K_{\mathbb{P}}$ implies that if there is some graph G that *does not* belong to \mathbb{P} , then there is some graph of treewidth at most $K_{\mathbb{P}}$ that *also does not* belong to \mathbb{P} . As a consequence, \mathbb{P} is equal to the class of *all graphs* (see Section 2 for a precise definition of this class) if and only if all graphs of treewidth at most $K_{\mathbb{P}}$ belong to \mathbb{P} . In other words, the irrelevant vertex technique allows one to show that certain conjectures are true in the class of all graphs if and only if they are true in the class of graphs of treewidth at most K for some *constant* K . This approach has been considered for instance in the study of Hadwiger’s conjecture (for each fixed number of colors c) [51, 74, 49]. Identifying further conjectures that can be studied under the framework of the irrelevant vertex technique would be very relevant to the framework of width-based automated theorem proving.

Acknowledgements

We acknowledge support from the Research Council of Norway in the context of the project *Automated Theorem Proving from the Mindset of Parameterized Complexity Theory* (proj. no. 288761).

References

- [1] Isolde Adler, Frederic Dorn, Fedor V Fomin, Ignasi Sau, and Dimitrios M Thilikos. Faster parameterized algorithms for minor containment. *Theoretical Computer Science*, 412(50):7018–7028, 2011.
- [2] Isolde Adler, Martin Grohe, and Stephan Kreutzer. Computing excluded minors. In *Proc. of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2008)*, pages 641–650. SIAM, 2008.
- [3] Isolde Adler, Stavros G. Kolliopoulos, Philipp Klaus Krause, Daniel Lokshantov, Saket Saurabh, and Dimitrios M. Thilikos. Irrelevant vertices for the planar disjoint paths problem. *J. Comb. Theory, Ser. B*, 122:815–843, 2017.
- [4] Alexsander Andrade de Melo and Mateus de Oliveira Oliveira. On the width of regular classes of finite structures. In *Proc. of the 27th International Conference on Automated Deduction (CADE 2019)*, pages 18–34. Springer, 2019.

- [5] Stefan Arnborg, Jens Lagergren, and Detlef Seese. Easy problems for tree-decomposable graphs. *Journal of Algorithms*, 12(2):308–340, 1991.
- [6] Max Bannach and Sebastian Berndt. Positive-instance driven dynamic programming for graph searching. In *Workshop on Algorithms and Data Structures*, pages 43–56. Springer, 2019.
- [7] Max Bannach and Sebastian Berndt. Practical access to dynamic programming on tree decompositions. *Algorithms*, 12(8):172, 2019.
- [8] Julien Baste, Michael R. Fellows, Lars Jaffke, Tomáš Masařík, Mateus de Oliveira Oliveira, Geevarghese Philip, and Frances A. Rosamond. Diversity of solutions: An exploration through the lens of fixed-parameter tractability theory. *Artificial Intelligence*, 303:103644, 2022.
- [9] Julien Baste, Ignasi Sau, and Dimitrios M. Thilikos. A complexity dichotomy for hitting connected minors on bounded treewidth graphs: the chair and the banner draw the boundary. In *SODA*, pages 951–970. SIAM, 2020.
- [10] Umberto Bertele and Francesco Brioschi. On non-serial dynamic programming. *J. Comb. Theory, Ser. A*, 14(2):137–148, 1973.
- [11] Therese C. Biedl. On triangulating k-outerplanar graphs. *Discret. Appl. Math.*, 181(1):275–279, 2015.
- [12] Hans L Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM Journal on computing*, 25(6):1305–1317, 1996.
- [13] Hans L Bodlaender. Treewidth: Algorithmic techniques and results. In *Proc. of the 22nd International Symposium on Mathematical Foundations of Computer Science*, pages 19–36. Springer, 1997.
- [14] Hans L Bodlaender. A partial k-arboretum of graphs with bounded treewidth. *Theoretical computer science*, 209(1-2):1–45, 1998.
- [15] Hans L Bodlaender, Marek Cygan, Stefan Kratsch, and Jesper Nederlof. Deterministic single exponential time algorithms for connectivity problems parameterized by treewidth. *Information and Computation*, 243:86–111, 2015.
- [16] Hans L Bodlaender and Arie MCA Koster. Combinatorial optimization on graphs of bounded treewidth. *The Computer Journal*, 51(3):255–269, 2008.
- [17] H.L. Bodlaender. Classes of graphs with bounded tree-width. Technical Report RUU-CS-86-22, Department of Information and Computing Sciences, Utrecht University, 1986.
- [18] Mikołaj Bojańczyk and Michał Pilipczuk. Definability equals recognizability for graphs of bounded treewidth. In *Proc. of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science (LICS 2016)*, pages 407–416. ACM, 2016.
- [19] HD Booth, Rajeev Govindan, Michael A Langston, and Siddharthan Ramachandramurthi. Cutwidth approximation in linear time. In *Proc. of the 2nd Great Lakes Symposium on VLSI*, pages 70–73. IEEE, 1992.
- [20] Richard B. Borie, R. Gary Parker, and Craig A. Tovey. Algorithms for recognition of regular properties and decomposition of recursive graph families. *Ann. Oper. Res.*, 33(3):125–149, 1991.

- [21] Andreas Brandstädt, Van Bang Le, and Jeremy P Spinrad. *Graph classes: a survey*. SIAM, 1999.
- [22] Chris Calabro, Russell Impagliazzo, and Ramamohan Paturi. The complexity of satisfiability of small depth circuits. In *Proc. of the 4th International Workshop on Parameterized and Exact Computation*, pages 75–85. Springer, 2009.
- [23] Marco L. Carmosino, Jiawei Gao, Russell Impagliazzo, Ivan Mihajlin, Ramamohan Paturi, and Stefan Schneider. Nondeterministic extensions of the strong exponential time hypothesis and consequences for non-reducibility. In *Proc. of the 7th ACM Conference on Innovations in Theoretical Computer Science (ITCS 2016)*, pages 261–270. ACM, 2016.
- [24] Fan RK Chung and Paul D Seymour. Graphs with small bandwidth and cutwidth. *Discrete Mathematics*, 75(1-3):113–119, 1989.
- [25] Julia Chuzhoy. Excluded grid theorem: Improved and simplified. In *Proc. of the 47th annual ACM symposium on Theory of Computing*, pages 645–654, 2015.
- [26] Julia Chuzhoy. Improved bounds for the flat wall theorem. In *Proc. of the 2015 ACM-SIAM Symposium on Discrete Algorithms*, pages 256–275. SIAM, 2015.
- [27] Hubert Comon, Max Dauchet, Rémi Gilleron, Florent Jacquemard, Denis Lugiez, Christof Löding, Sophie Tison, and Marc Tommasi. *Tree automata techniques and applications*, 2008.
- [28] Bruno Courcelle. The monadic second-order logic of graphs. I. Recognizable sets of finite graphs. *Information and Computation*, 85(1):12 – 75, 1990.
- [29] Bruno Courcelle. The monadic second-order logic of graphs XV: On a conjecture by d. Seese. *Journal of Applied Logic*, 4(1):79–114, 2006.
- [30] Bruno Courcelle and Irène Durand. Computations by fly-automata beyond monadic second-order logic. *Theor. Comput. Sci.*, 619:32–67, 2016.
- [31] Bruno Courcelle and Joost Engelfriet. *Graph structure and monadic second-order logic: A language-theoretic approach*, volume 138. Cambridge University Press, 2012.
- [32] Bruno Courcelle and Stephan Olariu. Upper bounds to the clique width of graphs. *Discrete Applied Mathematics*, 101(1-3):77–114, 2000.
- [33] Marek Cygan, Fedor V Fomin, Łukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized algorithms*, volume 5. Springer, 2015.
- [34] Anuj Dawar, Martin Grohe, and Stephan Kreutzer. Locally excluding a minor. In *Proc. of the 22nd ACM/IEEE Annual Symposium on Logic in Computer Science (LICS 2007)*, pages 270–279. IEEE Computer Society, 2007.
- [35] Mateus de Oliveira Oliveira. An algorithmic metatheorem for directed treewidth. *Discrete Applied Mathematics*, 204:49–76, 2016.
- [36] Rodney G Downey and Michael R. Fellows. *Parameterized complexity*. Springer Science & Business Media, 2012.
- [37] Michael Elberfeld. Context-free graph properties via definable decompositions. In *Proc. of the 25th Conference on Computer Science Logic (CSL 2016)*, volume 62 of *LIPIcs*, pages 17:1–17:16, 2016.

- [38] GH Fan. Tutte’s 3-flow conjecture and short cycle covers. *Journal of Combinatorial Theory, Series B*, 57(1):36–43, 1993.
- [39] Jörg Flum, Markus Frick, and Martin Grohe. Query evaluation via tree-decompositions. *Journal of the ACM (JACM)*, 49(6):716–752, 2002.
- [40] Stefania Gnesi, Ugo Montanari, and Alberto Martelli. Dynamic programming as graph searching: An algebraic approach. *Journal of the ACM (JACM)*, 28(4):737–751, 1981.
- [41] Hugo Hadwiger. Über eine klassifikation der streckenkomplexe. *Vierteljschr. Naturforsch. Ges. Zürich*, 88(2):133–142, 1943.
- [42] Rudolf Halin. S-functions for graphs. *Journal of geometry*, 8(1-2):171–186, 1976.
- [43] Illya V Hicks. Branch decompositions and minor containment. *Networks: An International Journal*, 43(1):1–9, 2004.
- [44] Illya V Hicks, Sang-il Oum, James J Cochran, Louis A Cox, Pinar Keskinocak, Jeffrey P Kharoufeh, and J Cole Smith. Branch-width and tangles. *Wiley Encyclopedia of Operations Research and Management Science*, 2010.
- [45] Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-SAT. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001.
- [46] Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, 2001.
- [47] Frank Kammer. Determining the smallest k such that g is k-outerplanar. In *European Symposium on Algorithms*, pages 359–370. Springer, 2007.
- [48] Ken-ichi Kawarabayashi and Yusuke Kobayashi. Algorithms for finding an induced cycle in planar graphs. *Comb.*, 30(6):715–734, 2010.
- [49] Ken-ichi Kawarabayashi and Bojan Mohar. Some recent progress and applications in graph minor theory. *Graphs Comb.*, 23(1):1–46, 2007.
- [50] Ken-ichi Kawarabayashi, Bojan Mohar, and Bruce A. Reed. A simpler linear time algorithm for embedding graphs into an arbitrary surface and the genus of graphs of bounded tree-width. In *FOCS*, pages 771–780. IEEE Computer Society, 2008.
- [51] Ken-ichi Kawarabayashi and Bruce Reed. Hadwiger’s conjecture is decidable. In *Proc. of the 41st Annual ACM Symposium on Theory of Computing*, pages 445–454, 2009.
- [52] Ken-ichi Kawarabayashi, Robin Thomas, and Paul Wollan. A new proof of the flat wall theorem. *Journal of Combinatorial Theory, Series B*, 129:204–238, 2018.
- [53] Eun Jung Kim, Sang-il Oum, Christophe Paul, Ignasi Sau, and Dimitrios M. Thilikos. An FPT 2-approximation for tree-cut decomposition. *Algorithmica*, 80(1):116–135, 2018.
- [54] Ephraim Korach and Nir Solel. Tree-width, path-width, and cutwidth. *Discrete Applied Mathematics*, 43(1):97–101, 1993.
- [55] Tuukka Korhonen. A single-exponential time 2-approximation algorithm for treewidth. In *Proc. of the 62nd Annual Symposium on Foundations of Computer Science (FOCS 2022)*, pages 184–192. IEEE, 2022.
- [56] Vipin Kumar and Laveen N Kanal. The cdp: A unifying formulation for heuristic search, dynamic programming, and branch-and-bound. In *Search in Artificial Intelligence*, pages 1–27. Springer, 1988.

- [57] C León, G Miranda, and C Rodríguez. Tools for tree searches: Dynamic programming. *Optimization Techniques for Solving Complex Problems*, pages 209–230, 2009.
- [58] Daniel Lokshtanov, Dániel Marx, and Saket Saurabh. Lower bounds based on the exponential time hypothesis. *Bull. EATCS*, 105:41–72, 2011.
- [59] Daniel Lokshtanov, Amer E. Mouawad, Saket Saurabh, and Meirav Zehavi. Packing cycles faster than erdos-posa. *SIAM J. Discret. Math.*, 33(3):1194–1215, 2019.
- [60] D Gonzalez Morales, Francisco Almeida, C Rodriguez, Jose L. Roda, I Coloma, and A Delgado. Parallel dynamic programming and automata theory. *Parallel computing*, 26(1):113–134, 2000.
- [61] Jesper Nederlof, Michal Pilipczuk, Céline M. F. Swennenhuis, and Karol Wegrzycki. Isolation schemes for problems on decomposable graphs. In *STACS*, volume 219 of *LIPICs*, pages 50:1–50:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.
- [62] Mateus de Oliveira Oliveira. Subgraphs satisfying MSO properties on z-topologically orderable digraphs. In *International Symposium on Parameterized and Exact Computation*, pages 123–136. Springer, 2013.
- [63] Sang-il Oum and Paul Seymour. Approximating clique-width and branch-width. *Journal of Combinatorial Theory, Series B*, 96(4):514–528, 2006.
- [64] Ivan Papusha, Jie Fu, Ufuk Topcu, and Richard M Murray. Automata theory meets approximate dynamic programming: Optimal control with temporal logic constraints. In *Proc. of the 55th IEEE Conference on Decision and Control (CDC 2016)*, pages 434–440. IEEE, 2016.
- [65] Miguel Angel Alonso Pardo, Eric De La Clergerie, and Manuel Vilares Ferro. Automata-based parsing in dynamic programming for linear indexed grammars. *Proc. of DIALOGUE*, 97:22–27, 1997.
- [66] D Stott Parker. Partial order programming. In *Proceedings of the 16th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 260–266, 1989.
- [67] Michał Pilipczuk. Problems parameterized by treewidth tractable in single exponential time: A logical approach. In *Proc. of the 36th International Symposium on Mathematical Foundations of Computer Science*, pages 520–531. Springer, 2011.
- [68] Jean-Florent Raymond and Dimitrios M Thilikos. Recent techniques and results on the erdős-pósa property. *Discrete Applied Mathematics*, 231:25–43, 2017.
- [69] Neil Robertson, Paul Seymour, and Robin Thomas. Hadwiger’s conjecture fork 6-free graphs. *Combinatorica*, 13(3):279–361, 1993.
- [70] Neil Robertson and Paul D Seymour. Graph minors. III. Planar tree-width. *Journal of Combinatorial Theory, Series B*, 36(1):49–64, 1984.
- [71] Neil Robertson and Paul D Seymour. Graph minors. V. Excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 41(1):92–114, 1986.
- [72] Ignasi Sau, Giannos Stamoulis, and Dimitrios M Thilikos. A more accurate view of the flat wall theorem. *arXiv preprint arXiv:2102.06463*, 2021.
- [73] Detlef Seese. The structure of the models of decidable monadic theories of graphs. *Annals of pure and applied logic*, 53(2):169–195, 1991.

- [74] Paul Seymour. *Hadwiger's conjecture*. Springer, 2016.
- [75] Dimitrios M Thilikos, Maria Serna, and Hans L Bodlaender. Cutwidth I: A linear time fixed parameter algorithm. *Journal of Algorithms*, 56(1):1–24, 2005.
- [76] Dimitrios M Thilikos, Maria J Serna, and Hans L Bodlaender. Constructive linear time algorithms for small cutwidth and carving-width. In *Proc. of the 11th International Symposium on Algorithms and Computation*, pages 192–203. Springer, 2000.
- [77] William Thomas Tutte. A contribution to the theory of chromatic polynomials. *Canadian journal of mathematics*, 6:80–91, 1954.
- [78] William Thomas Tutte. Recent progress in combinatorics. In *Proceedings of the 3rd Waterloo Conference on Combinatorics*. Academic Press, 1969.
- [79] Johan M. M. van Rooij, Hans L. Bodlaender, and Peter Rossmanith. Dynamic programming on tree decompositions using generalised fast subset convolution. In *ESA*, volume 5757 of *Lecture Notes in Computer Science*, pages 566–577. Springer, 2009.
- [80] Xiaofeng Wang, Cun-Quan Zhang, and Taoye Zhang. Nowhere-zero 4-flow in almost petersen-minor free graphs. *Discrete mathematics*, 309(5):1025–1032, 2009.
- [81] Michał Ziobro and Marcin Pilipczuk. Finding hamiltonian cycle in graphs of bounded treewidth: Experimental evaluation. *Journal of Experimental Algorithmics (JEA)*, 24:1–18, 2019.

A A more detailed version of Figure 1

In Figure 2 we depict a step-by-step construction of the graph of Figure 1.

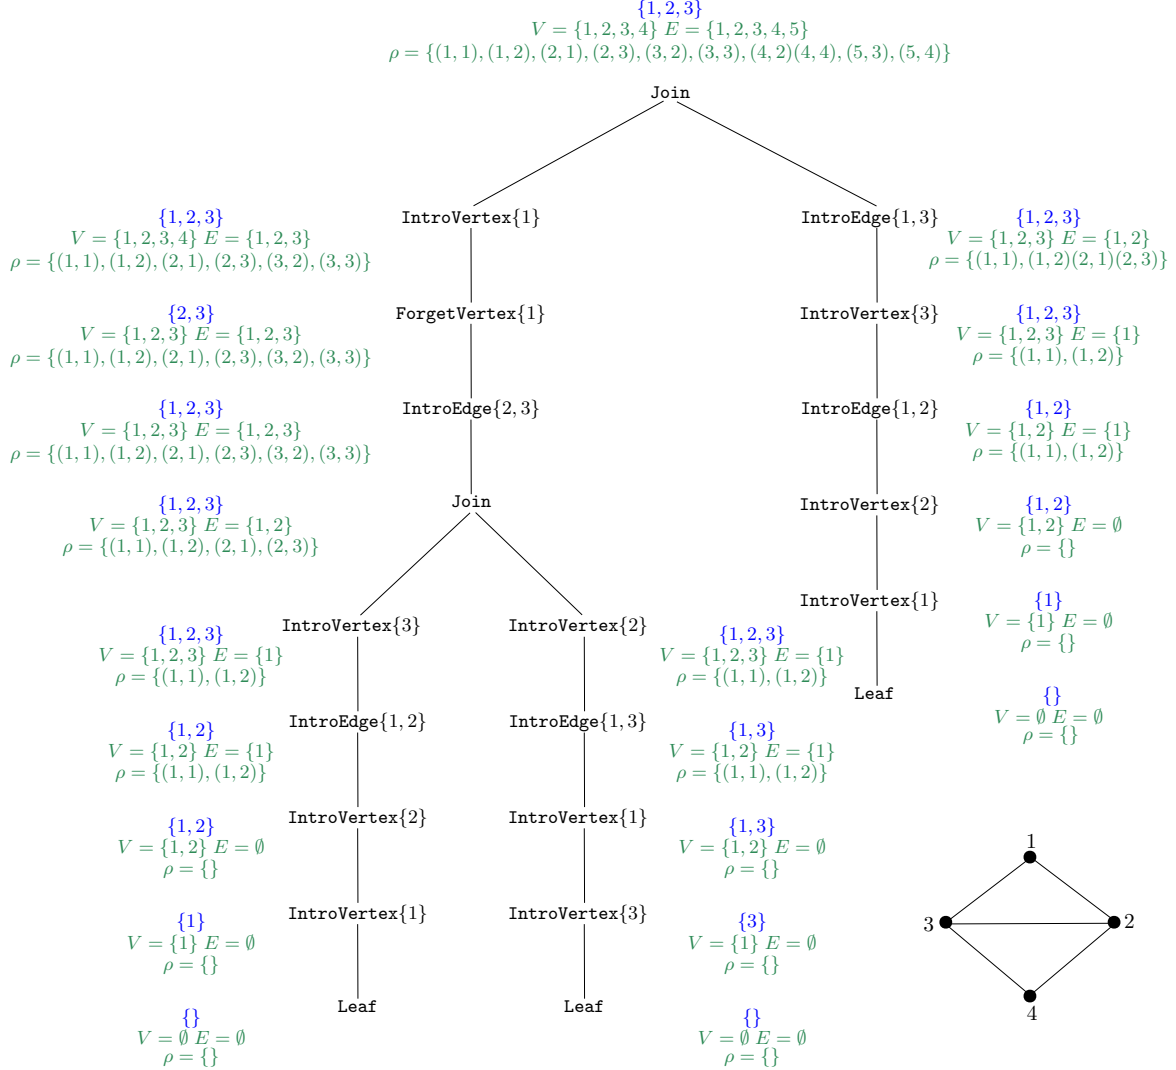


Figure 2: Construction of the graph associated to the 2-instructive tree decomposition of Fig. 1.

B Proof of Theorem 4

The proof of Theorem 4 readily follows from results available in the literature. A graph has cliquewidth at most k if and only if it can be defined as the graph associated with a k -expression as introduced in [32]. The fact that the set of all k -expressions is regular follows directly by the definition of k -expression [32, 30]. This shows that cliquewidth is a treelike width measure. In Chapter 12 of [36] it is shown how to define graphs of treewidth at most k using suitable parse trees, and how to associate a graph of treewidth at most k to each such parse tree. The fact that the set of all parse trees corresponding to graphs of treewidth at most k is regular is a direct consequence of the definition. This shows that treewidth is a treelike width measure. Graphs of pathwidth at most k can be obtained by considering a restricted version of the parse-trees considered in [36]. This restriction preserves regularity. Therefore pathwidth is also a treelike width measure. Cutwidth can be shown to be treelike using the notion of slice decompositions considered for instance in [62], while the fact that carving-width is treelike follows from the generalization of slice-decompositions to the context of trees defined in [35]. \square

C Proof of Lemma 7

Next, in Definition 35, we provide a slightly different, but equivalent definition of the notion of tree decomposition of a graph. The only difference is that we replace the condition that requires that for each edge of the graph there is a bag containing the endpoints of that edge, with a choice function β , which specifies, for each edge e , the bag $X_{\beta(e)}$ that contains the endpoints of e . We also assume that the tree is rooted, since choosing a root can be done without loss of generality.

Definition 35 (Tree Decomposition). *Let G be a graph. A tree decomposition of G is a triple (T, X, β) where $T = (N, F, r)$ is a rooted tree with set of nodes N , set of arcs F , and root r ; $\beta : E_G \rightarrow N$ is a function mapping each edge $e \in E_G$ to some node $p \in N$; and $X = \{X_u\}_{u \in N}$ is a collection of subsets of V_G satisfying the following properties.*

1. *For each vertex $x \in V_G$, there is some $p \in N$ such that $x \in X_p$.*
2. *For each edge $e \in E_G$, $\text{endpts}_G(e) \subseteq X_{\beta(e)}$.*
3. *For each vertex $x \in V_G$, the set $\{p \in N : x \in X_p\}$ induces a connected subtree of T .*

The subsets in X are called *bags*. The width of (T, X, β) is defined as the maximum size of a bag in X minus one: $w(T, X, \beta) = \max_{p \in N} |X_p| - 1$. The treewidth of a graph G is the minimum width of a tree decomposition of G .

We say that (T, X, β) is a *nice edge-introducing tree decomposition* of G if the function β is injective, and additionally, each node $p \in N$ is of one of the following types.

1. **Leaf** node: p has no child, and $X_p = \emptyset$.
2. **IntroVertex** node: p has a unique child p' and $X_p = X_{p'} \cup \{x\}$ for some vertex $x \in V_G$. We say that x is *introduced* at p .
3. **ForgetVertex** node: p has a unique child p' and $X_{p'} = X_p \cup \{x\}$ for some vertex $x \in V_G$. We say that x is *forgotten* at p .
4. **IntroEdge** node: p has a unique child p' , $X_p = X_{p'}$, and $p = \beta(e)$ for some $e \in E_G$. We say that e is *introduced* at p .
5. **Join** node: p has two children p' and p'' , and $X_p = X_{p'} = X_{p''}$.

We note that our notion of nice, edge-introducing tree decompositions of a graph is essentially identical to the notion of nice tree decompositions with introduce-edge nodes used in the literature (see [33], pages 161 and 168), except that we do not require the root bag to be empty. The fact that arbitrary tree decompositions can be efficiently transformed into nice, edge-introducing tree decompositions is an easy exercise (see for instance Lemma 7.4 of [33]).

Observation 36. *Let G be a graph, and (T, X, β) be a tree decomposition G of width at most k , where $T = (N, F, r)$. Then, one can construct in time $O(k^2 \cdot \max\{|V_G|, |N|\} + |E_G|)$ a nice, edge-introducing tree decomposition (T', X', β') of G of width at most k that has $O(k \cdot |V_G| + |E_G|)$ nodes.*

Now, we are in a position to prove that a graph G is isomorphic to the graph $\mathcal{G}(\tau)$ associated with some k -instructive tree decomposition τ if and only if G has treewidth at most k .

First, let τ be a k -instructive tree decomposition, and $\mathcal{G}(\tau)$ be its associated graph. Then this graph can be obtained from smaller graphs of size at most $k + 1$ by successively identifying sub-graphs of size at most $k + 1$. This implies that $\mathcal{G}(\tau)$ is a subgraph of a k -tree, and therefore that $\mathcal{G}(\tau)$ has treewidth at most k .

Claim 37. *Let τ be a k -instructive tree decomposition. Then, the graph $\mathcal{G}(\tau)$ has a nice edge-introducing tree decomposition (T, X, β) of width at most k where the bag associated with the root node r is the set $X_r = \theta[\tau](B(\tau))$.*

Proof. The proof is by induction on the height h of τ . In the base case, $h = 0$. In this case, $\tau = \text{Leaf}$. Therefore, since $\mathcal{G}(\tau)$ is the empty graph, the claim is vacuously satisfied. Now, suppose the claim is valid for every k -instructive tree decomposition τ of height at most h , and let τ be a k -instructive tree decomposition of height $h + 1$. There are four cases to be analyzed.

- Let σ be a k -instructive tree decompositions of height at most h and let (T, X, β) be a nice edge-introducing tree decompositions with roots r root-bag X_r satisfying the conditions of the claim. In this case we create a new root node r' with root bag $X_{r'}$, and set r as the child of r' .
 1. If $\tau = \text{IntroVertex}\{u\}(\sigma)$, we let $X_{r'} = X_r \cup \{\theta[\tau](u)\}$.
 2. If $\tau = \text{ForgetVertex}\{u\}(\sigma)$, we let $X_{r'} = X_r \setminus \{\theta[\tau](u)\}$.
 3. If $\tau = \text{IntroEdge}\{u, v\}(\sigma)$, we let $X_{r'} = X_r$.
- Let σ_1 and σ_2 be k -instructive tree decompositions of height at most h and let (T_1, X_1, β_1) and (T_2, X_2, β_2) be their respective nice edge-introducing tree decompositions with roots r_1 and r_2 and root-bags X_{r_1} and X_{r_2} satisfying the conditions of the claim.
 4. Let $\tau = \text{Join}(\sigma_1, \sigma_2)$. In this case, we create a new root bag r , labeled with the bag $X_r = X_{r_1}$ and set r_1 and r_2 as the children of r . Additionally, in the new tree decomposition, for each node p of T_2 , the bag X_p is replaced by the bag

$$X'_p = \{\theta[\tau_1](u) : u \in B(\tau), \theta[\tau_2](u) \in X_p\} \cup \{x + a : x \in X_p \setminus \theta[\tau](B(\tau))\}.$$

where $a = \max V_{\mathcal{G}(\tau)}$. Note that this renumbering is necessary because the vertices of $\mathcal{G}(\tau_2)$ are renumbered by the join operation $(\mathcal{G}(\tau_1), \theta[\tau]) \oplus (\mathcal{G}(\tau_2), \theta[\tau_2])$.

It is straightforward to check that in each of the four cases, the obtained tree decomposition satisfies the conditions of the claim. \square

Let G be a graph and $\alpha : V_G \rightarrow [k + 1]$ be a proper $(k + 1)$ -coloring of G . More specifically, for each edge $e \in E_G$ the endpoints of e are colored differently by α . We say that α is bag-injective with respect to a nice, edge introducing tree decomposition (T, X, β) of G if for each bag $X_p \in X$, the restriction $\alpha|_{X_p}$ is injective.

Observation 38. *Let G be a graph and (T, X, β) be a tree decomposition of width k of G , where $T = (N, F, r)$. Then, one can construct in time $O(k \cdot |N|)$ a proper $(k + 1)$ -coloring α of G that is bag-injective for (T, X, β) .*

Proof. We will construct a coloring $\alpha : V_G \rightarrow [k + 1]$ of G that is bag-injective for (T, X, β) by traversing the bags in X from the root towards the leaves. We start by choosing an arbitrary injective coloring of the vertices in the root bag X_r . Since X_r has at most $k + 1$ vertices, such an injective coloring with at most $(k + 1)$ colors exists. Now assume that the vertices of a given bag X_p have been injectively colored with at most $k + 1$ colors, and let p' be a child of p . Then we have three possibilities: (i) $X_{p'} = X_p$; (ii) $X_{p'} = X_p \setminus \{x\}$, for some vertex x ; (iii) $X_{p'} = X_p \cup \{x\}$ for some vertex x . In the first two cases, an injective coloring of the vertices in $X_{p'}$ has already been chosen. In the last case, we just choose some arbitrary color for x in the set $[k + 1] \setminus \alpha(X_p)$. We proceed in this way until all bags have been visited. Since each vertex occurs in some bag, each vertex has received some color from the set $[k + 1]$. Since the endpoints of each edge e are contained in some bag, and since by construction, the coloring in each bag is injective, we have that the endpoints of e receive distinct colors. Therefore, α is a proper $(k + 1)$ -coloring of G . \square

Now, let G be a graph of treewidth at most k . Let (T, X, β) be a nice, edge-introducing tree decomposition of G of width at most k . Let $\alpha : V_G \rightarrow [k+1]$ be a proper $(k+1)$ -coloring of G that is bag-injective for (T, X, β) . By Observation 38 such a coloring exists and can be constructed in time $O(k \cdot |N|)$ where N is the set of nodes of T . Then, a k -instructive tree decomposition τ of G can be constructed as follows. We let τ have the same structure as T , except that instead being labeled with bags, the nodes are now labeled with instructions from the alphabet Σ_k . More specifically, each leaf node is labeled with the instruction **Leaf**. Now, let p be a node of T . If a vertex x is introduced at p , then we label the node p of τ with the instruction **IntroVertex** $\{\alpha(x)\}$. If an edge $\{x, y\}$ is introduced at p , then the node p of τ is labeled with the instruction **IntroEdge** $\{\alpha(x), \alpha(y)\}$. If a vertex x is forgotten at p , then the node p of τ is labeled with the instruction **ForgetVertex** $\{\alpha(x)\}$. Finally, p is a join node, then the node p of τ is labeled with the instruction **Join**. This concludes the proof of Lemma 7. \square

D Proof of Lemma 19

In this section, we prove Lemma 19 by induction on the height of a k -instructive tree decomposition τ . Although the proof is straightforward, it serves as an illustration of how the objects $B(\tau)$, $\mathcal{G}(\tau)$, $\theta[\tau]$ and $\Gamma[\mathbf{D}, k](\tau)$ interact in an inductive proof of correctness of a DP-core \mathbf{D} . In particular, the structure of this proof can be adapted to provide inductive proofs for the correctness of DP-cores deciding other graph properties.

Base Case: In the base case, the height of τ is 0, and therefore $\tau = \mathbf{Leaf}$. By Definition 9 and Definition 18.1, $\Gamma[\mathbf{C}\text{-VertexCover}_r, k](\tau) = \mathbf{C}\text{-VertexCover}_r[k].\mathbf{LeafSet} = \{(\emptyset, 0)\}$. Since $\mathcal{G}(\tau) = (\emptyset, \emptyset, \emptyset)$, we have that \emptyset is the only vertex cover of $\mathcal{G}(\tau)$, and since this vertex cover has size 0, $\mathbf{w} = (\emptyset, 0)$ is the only local witness with $\mathbf{P}\text{-VertexCover}_r(\tau, \mathbf{w}) = 1$. Therefore, the lemma holds in the base case.

Inductive Step: Now, assume that the lemma holds for every k -instructive tree decomposition of height at most h . Let τ be a k -instructive tree decomposition of height $h+1$, and $\mathbf{w} = (R, s)$ be a local witness in $\mathbf{C}\text{-VertexCover}_r[k].\mathcal{W}$. There are four cases to be considered.

1. Let $\tau = \mathbf{IntroVertex}\{u\}(\sigma)$. Suppose $\mathbf{w} \in \Gamma[\mathbf{C}\text{-VertexCover}_r, k](\tau)$. By Definition 9 and Definition 18.2, we have that \mathbf{w} also belongs to $\Gamma[\mathbf{C}\text{-VertexCover}_r, k](\sigma)$. By the induction hypothesis, $\mathbf{P}\text{-VertexCover}_r[k](\sigma, \mathbf{w}) = 1$. Then, $s \leq r$, and s is the minimum size of a vertex cover X in $\mathcal{G}(\sigma)$ with $\theta[\sigma](R) = X \cap \text{Im}(\theta[\sigma](B(\sigma)))$. This implies that X is also a vertex cover in $\mathcal{G}(\tau)$ with $\theta[\tau](R) = X \cap \text{Im}(\theta[\tau](B(\tau)))$. Therefore, $\mathbf{P}\text{-VertexCover}_r[k](\tau, \mathbf{w}) = 1$.

For the converse, let $\mathbf{P}\text{-VertexCover}_r[k](\tau, \mathbf{w}) = 1$. Then, there is a vertex cover X of $\mathcal{G}(\tau)$ of size s where $\theta[\tau](R) = X \cap \text{Im}(\theta[\tau](B(\tau)))$ and s is the minimum size of such vertex cover. Since $E_{\mathcal{G}(\tau)} = E_{\mathcal{G}(\sigma)}$, the vertex cover X is a vertex cover of $\mathcal{G}(\sigma)$ such that $\theta[\sigma](R) = X \cap \text{Im}(\theta[\sigma](B(\sigma)))$. Therefore, $\mathbf{P}\text{-VertexCover}_r[k](\sigma, \mathbf{w}) = 1$. By the induction hypothesis, $\mathbf{w} \in \Gamma[\mathbf{C}\text{-VertexCover}_r, k](\sigma)$. By Definition 18.2, we have that $\mathbf{C}\text{-VertexCover}_r[k].\mathbf{IntroVertex}\{u\}(\mathbf{w}) = \{\mathbf{w}\}$. Therefore, $\mathbf{w} \in \Gamma[\mathbf{C}\text{-VertexCover}_r, k](\tau)$.

2. Let $\tau = \mathbf{ForgetVertex}\{u\}(\sigma)$. Suppose $\mathbf{w} \in \Gamma[\mathbf{C}\text{-VertexCover}_r, k](\tau)$. By Definition 18.3, there exists $\mathbf{w}' = (R', s') \in \Gamma[\mathbf{C}\text{-VertexCover}_r, k](\sigma)$ such that $\mathbf{C}\text{-VertexCover}_r[k].\mathbf{ForgetVertex}\{u\}(\mathbf{w}') = \{\mathbf{w}\}$. Note that $R' \setminus \{u\} = R$ and $s' = s$. By the induction hypothesis, $\mathbf{P}\text{-VertexCover}_r[k](\sigma, \mathbf{w}') = 1$, and therefore, there exists a vertex cover X of minimum size s' in $\mathcal{G}(\sigma)$ where $\theta[\sigma](R') = X \cap \text{Im}(\theta[\sigma](B(\sigma)))$. Since $\mathcal{G}(\sigma) = \mathcal{G}(\tau)$, the set X is a vertex cover of $\mathcal{G}(\tau)$. We can infer the following by the fact that $B(\tau) = B(\sigma) \setminus \{u\}$, $R = R' \setminus \{u\}$, and $\theta[\tau] = \theta[\sigma]|_{B(\tau)}$:

$$\theta[\tau](R) = X \cap \text{Im}(\theta[\tau](B(\tau))) \quad (4)$$

Since X is a vertex cover of minimum size in $\mathcal{G}(\sigma)$ such that $\theta[\sigma](R') = X \cap \text{Im}(\theta[\sigma](B(\sigma)))$ and since $\mathcal{G}(\tau) = \mathcal{G}(\sigma)$, we have that X is also a vertex cover of minimum size in $\mathcal{G}(\tau)$ that satisfies Equation 4. Consequently, $\text{P-VertexCover}_r[k](\tau, \mathbf{w}) = 1$.

For the converse, suppose $\text{P-VertexCover}_r[k](\tau, \mathbf{w}) = 1$. Then, there is a vertex cover X of minimum size such that $\theta[\tau](R) = X \cap \text{Im}(\theta[\tau](B(\tau)))$ and $|X| = s \leq r$. Since $\mathcal{G}(\tau) = \mathcal{G}(\sigma)$, the set X is a vertex cover of $\mathcal{G}(\sigma)$. There are two cases to be considered:

- (a) $\theta[\sigma](u) \in X$. In this case, let $\mathbf{w}' = (R', s)$ where $R' = R \cup \{u\}$.
- (b) $\theta[\sigma](u) \notin X$. In this case, let $\mathbf{w}' = (R', s)$ where $R' = R$.

In both cases we will show that $\mathbf{w}' \in \Gamma[\text{C-VertexCover}_r, k](\sigma)$. Recall that $B(\tau) = B(\sigma) \setminus \{u\}$ and $\theta[\tau] = \theta[\sigma]|_{B(\tau)}$, and therefore we have that

$$\theta[\sigma](R') = X \cap \text{Im}(\theta[\sigma](B(\sigma))). \quad (5)$$

Equation 5 holds for the two cases and X is a vertex cover of minimum size satisfying Equation 5. Then, $\text{P-VertexCover}_r[k](\sigma, \mathbf{w}') = 1$, and by the induction hypothesis, $\mathbf{w}' \in \Gamma[\text{C-VertexCover}_r, k](\sigma)$. By Definition 18.3, $\text{C-VertexCover}_r[k].\text{ForgetVertex}\{u\}(\mathbf{w}') = \{(R' \setminus \{u\}, s)\}$ and we have that $(R' \setminus \{u\}, s) = \mathbf{w}$, and therefore, $\mathbf{w} \in \Gamma[\text{C-VertexCover}_r, k](\tau)$.

3. Let $\tau = \text{IntroEdge}\{u, v\}(\sigma)$. Suppose $\mathbf{w} \in \Gamma[\text{C-VertexCover}_r, k](\sigma)$. By Definition 18.4, there exists $\mathbf{w}' = (R', s') \in \Gamma[\text{C-VertexCover}_r, k](\sigma)$ where $\text{C-VertexCover}_r[k].\text{IntroEdge}\{u, v\}(\mathbf{w}') = \{\mathbf{w}\}$ and either $(R, s) = (R', s')$ or $(R, s) = (R' \cup \{x\}, s' + 1)$ where $x \in \{u, v\}$. By the induction hypothesis, $\text{P-VertexCover}_r[k](\sigma, \mathbf{w}') = 1$, and therefore, there is a vertex cover X of minimum size where $\theta[\sigma](R') = X \cap \text{Im}(\theta[\sigma](B(\sigma)))$. We consider two cases and define X' with regards to these cases.

- (a) $R = R' \cup \{x\}$ where $x \in \{u, v\}$. In this case let $X' = X \cup \{\theta[\tau](x)\}$.
- (b) $R = R'$. In this case let $X' = X$.

Note that $E_{\mathcal{G}(\tau)} = E_{\mathcal{G}(\sigma)} \cup \{e\}$ where $(e, \theta[\tau](u)), (e, \theta[\tau](v)) \in \rho_{\mathcal{G}(\tau)}$ so that X' is a vertex cover of $\mathcal{G}(\tau)$. Recall that $B(\tau) = B(\sigma)$ and $\theta[\tau] = \theta[\sigma]$. We can infer the following equation by the fact that we have mentioned.

$$\theta[\tau](R) = X' \cap \text{Im}(\theta[\tau](B(\tau))) \quad (6)$$

If we show X' is a vertex cover of minimum size satisfying 6, then we have proved the statement. Next, we will prove the vertex cover X' is minimum size. Suppose $X' = X$ (the second case), then obviously X' is minimum size. Otherwise, suppose $X' = X \cup \{\theta[\tau](x)\}$ (the first case) and suppose X' is not a vertex cover of minimum size and there is a vertex cover $|X''| < |X'|$ where $\theta[\tau](R) = X'' \cap \text{Im}(\theta[\tau](B(\tau)))$. Therefore, we have $\theta[\tau](R') = (X'' \setminus \{\theta[\tau](x)\}) \cap \text{Im}(\theta[\tau](B(\sigma)))$ and we know $|X'' \setminus \{\theta[\tau](x)\}| < |X|$, and therefore, this contradicts the assumption that X is the minimum size vertex cover which satisfies the second condition of the definition of P-VertexCover_r .

For the converse, suppose $\text{P-VertexCover}_r[k](\tau, \mathbf{w}) = 1$. Then, there exists a vertex cover X where $\theta[\tau](R) = X \cap \text{Im}(\theta[\tau](B(\tau)))$ where X is the minimum size of such vertex cover. In the following let e be the edge that $\theta[\tau](u)$ and $\theta[\tau](v)$ are the endpoints of it.

- (a) $X \setminus \{\theta[\tau](x)\}$ where $x \in \{u, v\}$ is a vertex cover of $\mathcal{G}(\tau) \setminus \{e\}$. In this case let $\mathbf{w}' = (R', s')$ where $R' = R \setminus \{u, v\}$ and $s' = s - 1$, and let $X' = X \setminus \{\theta[\tau](x)\}$.
- (b) $X \setminus \{\theta[\tau](x)\}$ where $x \in \{u, v\}$ is not a vertex cover of $\mathcal{G}(\tau) \setminus \{e\}$. In this case, we set $\mathbf{w}' = (R', s')$ where $R' = R$ and $s' = s$, and let $X' = X$.

In both cases, we show that $\text{P-VertexCover}_r[k](\sigma, \mathbf{w}') = 1$. The condition $s' \leq r$, is satisfied. Since $\mathcal{G}(\tau) = \mathcal{G}(\sigma) \setminus \{e\}$, X' in both cases is a vertex cover of $\mathcal{G}(\sigma)$. We have $B(\tau) = B(\sigma)$ and $\theta[\tau] = \theta[\sigma]$, and therefore in the both cases the following equation is satisfied.

$$\theta[\sigma](R') = X' \cap \text{Im}(\theta[\sigma](B(\sigma))) \quad (7)$$

Now it remains that we show X' is the minimum size vertex cover where satisfying Equation 7. If $X' = X$ and since X is minimum size, X' is minimum size. If $X' = X \setminus \{\theta[\tau](x)\}$ and suppose X' is not minimum size, therefore, there exists a vertex cover X'' where $\theta[\sigma](R') = X'' \cap \text{Im}(\theta[\sigma](B(\sigma)))$ such that $|X''| < |X'|$, and consequently, we have $|X'' \cup \{\theta[\tau](x)\}| < |X|$ and $\theta[\tau](R) = (X'' \cup \{\theta[\tau](x)\}) \cap \text{Im}(\theta[\tau](B(\tau)))$ and this contradicts the assumption that X is minimum size. Therefore, X' is minimum size, and consequently, $\text{P-VertexCover}_r[k](\sigma, \mathbf{w}') = 1$. By the induction hypothesis, $\mathbf{w}' \in \Gamma[\text{C-VertexCover}_r, k](\sigma)$. By Definition 18.4, $\text{C-VertexCover}_r[k](\mathbf{w}') = \{\mathbf{w}\}$, and therefore, $\mathbf{w} \in \Gamma[\text{C-VertexCover}_r, k](\tau)$.

4. Let $\tau = \text{Join}(\sigma_1, \sigma_2)$. Suppose $\mathbf{w} \in \Gamma[\text{C-VertexCover}_r, k](\tau)$. Let $\mathbf{w}_1 = (R_1, s_1) \in \Gamma[\text{C-VertexCover}_r, k](\sigma_1)$ and $\mathbf{w}_2 = (R_2, s_2) \in \Gamma[\text{C-VertexCover}_r, k](\sigma_2)$ be two local witnesses such that $\text{C-VertexCover}_r[k].\text{Join}(\mathbf{w}_1, \mathbf{w}_2) = \{\mathbf{w}\}$, and by Definition 18.5, $R = R_1 \cup R_2$ and $s = s_1 + s_2 - |R_1 \cap R_2|$ where $s \leq r$. By the induction hypothesis, since $\mathbf{w}_i \in \Gamma[\text{C-VertexCover}_r, k](\sigma_i)$, we have $\text{P-VertexCover}_r[k](\sigma_i, \mathbf{w}_i) = 1$ for $i \in \{1, 2\}$. Therefore, there exists a vertex cover X_i of minimum size of $\mathcal{G}(\sigma_i)$ where $\theta[\sigma_i](R_i) = X_i \cap \text{Im}(\theta[\sigma_i](B(\sigma_i)))$. Let $X = X_1 \cup \{v + a : v \in X_2 \setminus \theta[\sigma_2](B(\sigma_2))\}$ where $a = \max\{v : v \in V_{\mathcal{G}(\sigma_1)}\}$. Let $\zeta : V_{\mathcal{G}(\sigma_2)} \rightarrow V_{\mathcal{G}(\tau)}$ be the map between corresponding labels of vertices in $\mathcal{G}(\sigma_2)$ and $\mathcal{G}(\tau)$. More specifically, we let $a = \max\{v : v \in V_{G_1}\}$, and set for each $x \in V_{\mathcal{G}(\sigma_2)}$

$$\zeta(x) = \begin{cases} \theta[\tau](\theta[\sigma_2]^{-1}(x)) & \text{if } x \in \text{Range}(\theta[\sigma_2]) \\ x + a & \text{if } x \notin \text{Range}(\theta[\sigma_2]) \end{cases}$$

By the fact that $B(\tau) = B(\sigma_2)$ and that $R_2 \subseteq B(\sigma_2)$, we can infer that $\zeta(\theta[\sigma_2](R_2)) = \theta[\tau](R_2)$, $\zeta(\text{Im}(\theta[\sigma_2](B(\sigma_2)))) = \text{Im}(\theta[\tau](B(\tau)))$, $\zeta(X_2) = (X \setminus X_1) \cup \theta[\tau](R_2)$, and $X = X_1 \cup \zeta(X_2)$. Therefore, we have the following.

$$\begin{aligned} \zeta(\theta[\sigma_2](R_2)) &= \zeta(X_2) \cap \zeta(\text{Im}(\theta[\sigma_2](B(\sigma_2)))) \\ \implies \theta[\tau](R_2) &= X'_2 \cap \text{Im}(\theta[\tau](B(\tau))) \end{aligned} \quad (8)$$

where $X'_2 = \zeta(X_2)$. The set X is a vertex cover of $\mathcal{G}(\tau)$ since each edge is either in $\mathcal{G}(\sigma_1)$ or $\mathcal{G}(\sigma_2)$. If an edge is in $\mathcal{G}(\sigma_1)$, then at least one of the endpoints is in X_1 and if an edge is in $\mathcal{G}(\sigma_2)$, then at least one endpoints is in X'_2 . By the fact $\theta[\sigma_1](R_1) = X_1 \cap \text{Im}(\theta[\sigma_1](B(\sigma_1)))$, $B(\tau) = B(\sigma_1) = B(\sigma_2)$, $\theta[\tau] = \theta[\sigma_1]$, and Equation 8, we can infer the following.

$$\begin{aligned} \theta[\tau](R_1) \cup \theta[\tau](R_2) &= (X_1 \cap \text{Im}(\theta[\tau](B(\tau)))) \cup (X'_2 \cap \text{Im}(\theta[\tau](B(\tau)))) \\ \implies \theta[\tau](R_1) \cup \theta[\tau](R_2) &= (X_1 \cup X'_2) \cap \text{Im}(\theta[\tau](B(\tau))) \\ \implies \theta[\tau](R) &= X \cap \text{Im}(\theta[\tau](B(\tau))) \end{aligned} \quad (9)$$

The vertex cover X satisfies Equation 9 and if we show that s is the minimum size of such vertex cover satisfying Equation 9, then we have proved the statement. Suppose X is not a vertex cover of minimum size and there exists a vertex cover \mathcal{M} such that $|\mathcal{M}| < |X|$ where $\theta[\tau](R) = \mathcal{M} \cap \text{Im}(\theta[\tau](B(\tau)))$. Let $\mathcal{M}_1 = \mathcal{M} \cap V_{\mathcal{G}(\sigma_1)}$ and $\mathcal{M}_2 = \zeta^{-1}(\mathcal{M} \cap \zeta(V_{\mathcal{G}(\sigma_2)}))$ be vertex covers of $\mathcal{G}(\sigma_1)$ and $\mathcal{G}(\sigma_2)$ satisfying the following equation.

$$\theta[\sigma_i](R_i) = \mathcal{M}_i \cap \text{Im}(\theta[\sigma_i](B(\sigma_i))) \quad (10)$$

for $i \in \{1, 2\}$. We have $\mathcal{M}_1 \cup \zeta(\mathcal{M}_2) = \mathcal{M}$, and consequently, the following equation holds.

$$|\mathcal{M}| = |\mathcal{M}_1 \cup \zeta(\mathcal{M}_2)| = |\mathcal{M}_1| + |\zeta(\mathcal{M}_2)| - |\mathcal{M}_1 \cap \zeta(\mathcal{M}_2)| \quad (11)$$

Also we have the following,

$$|X| = |X_1 \cup \zeta(X_2)| = |X_1| + |\zeta(X_2)| - |X_1 \cap \zeta(X_2)| \quad (12)$$

We know X_i is a vertex cover of minimum size of $\mathcal{G}(\sigma_i)$ satisfying Equation 10, and therefore, $|X_i| \leq |\mathcal{M}_i|$, and also by $\theta[\tau](R_1 \cap R_2) = \mathcal{M}_1 \cap \zeta(\mathcal{M}_2) = X_1 \cap \zeta(X_2)$, Equation 11, and Equation 12 we can infer $|X| \leq |\mathcal{M}|$. This contradicts the assumption that $|\mathcal{M}| < |X|$, and therefore, X is minimum size. As a consequence, $\text{P-VertexCover}_r[k](\tau, \mathbf{w}) = 1$.

For the converse, suppose $\text{P-VertexCover}_r[k](\tau, \mathbf{w}) = 1$. There is a vertex cover X where satisfies the second Condition of the definition of P-VertexCover_r . Let $R_1 = \theta[\tau]^{-1}(X \cap V_{\mathcal{G}(\sigma_1)} \cap \theta[\sigma_1](B(\sigma_1)))$, $s_1 = |X \cap V_{\mathcal{G}(\sigma_1)}|$, and $X_1 = X \cap V_{\mathcal{G}(\sigma_1)}$. Also, let $R_2 = \theta[\tau]^{-1}(X \cap \zeta(V_{\mathcal{G}(\sigma_2)}) \cap \zeta(\theta[\sigma_2](B(\sigma_2))))$, $s_2 = |X \cap \zeta(V_{\mathcal{G}(\sigma_2)})|$, and $X_2 = \zeta^{-1}(X \cap \zeta(V_{\mathcal{G}(\sigma_2)}))$. The following equation holds for the defined sets.

$$\theta[\sigma_i](R_i) = X_i \cap \text{Im}(\theta[\sigma_i](B(\sigma_i))) \quad (13)$$

Let $\mathbf{w}_i = (R_i, s_i)$ and we will show that $\text{P-VertexCover}_r[k](\sigma_i, \mathbf{w}_i) = 1$ for $i \in \{1, 2\}$. To prove it, we need to show X_i is a vertex cover of minimum size satisfying Equation 13. Suppose it is not minimum, and there is a vertex cover $|X'_1| < |X_1|$ (same argument for X_2) that satisfies $\theta[\sigma_1](R_1) = X'_1 \cap \text{Im}(\theta[\sigma_1](B(\sigma_1)))$. The set $X' = X'_1 \cup \zeta(X_2)$ is also a vertex cover of $\mathcal{G}(\tau)$ where $\theta[\tau](R) = X' \cap \text{Im}(\theta[\tau](B(\tau)))$ and $|X'| < |X|$. This contradicts the assumption that X is a vertex cover of minimum size.

We have proved $\text{P-VertexCover}_r[k](\sigma_i, \mathbf{w}_i) = 1$, and therefore, by the induction hypothesis, $\mathbf{w}_i \in \Gamma[\text{C-VertexCover}_r, k](\sigma_i)$ for $i \in \{1, 2\}$. By Definition 18.5, we have that $\text{C-VertexCover}_r[k].\text{Join}(\mathbf{w}_1, \mathbf{w}_2) = \{\mathbf{w}\}$, and therefore, $\mathbf{w} \in \Gamma[\text{C-VertexCover}_r, k](\tau)$. \square

E Proof of Theorem 34

In this section, we prove Theorem 34. More specifically, for each graph property \mathbb{P} listed in Theorem 34, we provide an upper bound for the bit-length $\beta_{\mathbb{D}}(k)$ of a suitable ITD-coherent DP-core \mathbb{D} with $\mathbb{G}[\mathbb{D}, \text{ITD}] = \mathbb{P}$. The multiplicity of such a DP-core \mathbb{D} is trivially upper bounded by $2^{O(\beta_{\mathbb{D}}(k))}$, the state complexity of \mathbb{D} is trivially upper bounded by $2^{O(\beta_{\mathbb{D}}(k))}$, and the deterministic state complexity of \mathbb{D} is trivially upper bounded by $2^{2^{O(\beta_{\mathbb{D}}(k))}}$. In some cases, we show that that these two last upper bounds can be improved using additional arguments.

The process of specifying a coherent DP-core \mathbb{D} deciding a given graph property \mathbb{P} can be split into four main steps, which should hold for each $k \in \mathbb{N}$.

1. The specification of a suitable set $\mathbb{D}[k].\mathcal{W}$ of local witnesses, and of a Boolean function $\mathbb{D}[k].\text{Final}$ determining which are the final local witnesses in $\mathbb{D}[k].\mathcal{W}$.
2. The specification of a predicate $\mathbb{P}[k] \subseteq \text{ITD}_k \times \mathbb{D}[k].\mathcal{W}$ such that for each $\tau \in \text{ITD}_k$, and each *final local witness* \mathbf{w} , $(\tau, \mathbf{w}) \in \mathbb{P}[k]$ if and only if the graph $\mathcal{G}(\tau)$ belongs to \mathbb{P} .
3. The specification of the functions belonging to $\mathbb{D}[k]$.
4. An inductive proof that for each pair $(\tau, \mathbf{w}) \in \text{ITD}_k \times \mathbb{D}[k].\mathcal{W}$, $(\tau, \mathbf{w}) \in \Gamma[\mathbb{D}, k](\tau)$ if and only if $(\tau, \mathbf{w}) \in \mathbb{P}[k]$. This last step together with Step 2 guarantees that the DP-core is coherent and that $\mathbb{G}[\mathbb{D}[k]] = \mathbb{P}$.

The process described above was exemplified in full detail with respect to the DP-core **C-VertexCover**, defined in Section 5.3. In this case, Steps 1, 2 and 3 were carried out in Section 5.3, and Step 4 was carried out in Appendix D. In this section, we are only interested in providing an upper bound on the bitlength of the DP-cores listed in Theorem 34. Therefore, we will only carry out in details Steps 1 and 2, which are the ones that require ingenuity. Once this is done, it is straightforward to specify the functions of $D[k]$ in such a way that the requirement

$$(\tau, \mathbf{w}) \in \Gamma[D, k](\tau) \Leftrightarrow (\tau, \mathbf{w}) \in P[k]$$

is satisfied by construction. A full inductive proof of correctness of the specification, as exemplified in Section D is also straightforward.

Our choices of local witnesses is based on simplicity, instead of efficiency, and our goal primarily to justify the asymptotic upper bounds listed in Theorem 34.

E.1 C-Simple

The graph property of the DP-core **C-Simple** is the set **Simple** of all simple graphs. For each $k \in \mathbb{N}$, a local witness for **C-Simple** $[k]$ is a subset $\mathbf{w} \subseteq \binom{[k+1]}{2}$. All local witnesses are final. The transitions of **C-Simple** $[k]$ are defined in such a way that for each k -instructive tree decomposition τ and each local witness \mathbf{w} , $\mathbf{w} \in \Gamma[\mathbf{C-Simple}, k](\tau)$ if and only if the following predicate is satisfied:

- **P-Simple** $[k](\tau, \mathbf{w}) \equiv \mathcal{G}(\tau)$ is a simple graph, and for each two distinct elements u and v in $[k+1]$, $\{u, v\} \in \mathbf{w}$ if and only if there is a unique edge $e \in E_{\mathcal{G}(\tau)}$ such that $\text{endpts}_{\mathcal{G}(\tau)}(e) = \{\theta[\tau](u), \theta[\tau](v)\}$.

Each local witness \mathbf{w} may be represented as a Boolean vector in $\{0, 1\}^{\binom{k+1}{2}}$ which has one coordinate for each pair $\{u, v\} \in \mathbf{w}$. The DP-core **C-Simple** $[k]$ can be defined in such a way that it is deterministic, in the sense that the set obtained from each local witness \mathbf{w} upon the application of each transition is a singleton. Therefore, the multiplicity of **C-Simple** $[k]$ is 1. This implies that the deterministic state complexity of **C-Simple** $[k]$ is at most $2^{\binom{k+1}{2}}$.

E.2 C-MaxDeg $_{\geq}(d)$

The graph property of the DP-core **C-MaxDeg $_{\geq}(d)$** is the set **MaxDeg $_{\geq}(d)$** of all graphs with maximum degree at least d . For each $k \in \mathbb{N}$, a local witness for **C-MaxDeg $_{\geq}(d)$** $[k]$ is a pair $(x, y) \in \{0, 1\} \times \{0, 1, \dots, d+1\}^{k+1}$. Such local witness is final if and only if $x = 1$ or there is some $u \in [k+1]$ such that $y_u \geq d$. The transitions of **C-MaxDeg $_{\geq}(d)$** $[k]$ are defined in such a way that for each k -instructive tree decomposition τ and each local witness (x, y) , $(x, y) \in \Gamma[\mathbf{C-MaxDeg}_{\geq}(d), k](\tau)$ if and only if the following predicate is satisfied:

- **P-MaxDeg $_{\geq}(d)$** $[k](\tau, (x, y)) \equiv$ for each $u \in [k+1]$, the vertex $\theta[\tau](u)$ in the graph $\mathcal{G}(\tau)$ has degree y_u if $y_u \leq d$ and degree at least $d+1$ if $y_u = d+1$; $x = 1$ if and only if there is some vertex of degree at least d in $V_{\mathcal{G}(\tau)} \setminus \theta[\tau](B(\tau))$.

Each local witness (x, y) can be represented as a binary string containing $1 + (k+1) \cdot \lceil \log(d+1) \rceil$ bits. The DP-core **C-MaxDeg $_{\geq}(d)$** $[k]$ can be defined in such a way that it is deterministic. Therefore, the multiplicity of **C-MaxDeg $_{\geq}(d)$** $[k]$ is 1. This implies that the deterministic state complexity of **C-MaxDeg $_{\geq}(d)$** $[k]$ is upper bounded by $2^{O(k \log d)}$.

E.3 C-MinDeg $_{\leq}(d)$

The graph property of the DP-core **C-MinDeg $_{\leq}(d)$** is the set **MinDeg $_{\leq}(d)$** of all graphs with minimum degree at most d . For each $k \in \mathbb{N}$, a local witness for **C-MinDeg $_{\leq}(d)$** $[k]$ is a pair $(x, y) \in \{0, 1\} \times \{0, 1, \dots, d+1\}^{k+1}$. Such local witness is final if and only if $x = 1$ or $y_u \leq d$

for some $u \in [k+1]$. The transitions of $\mathbf{C}\text{-MinDeg}_{\leq}(d)[k]$ are defined in such a way that for each k -instructive tree decomposition τ and each local witness y , $(x, y) \in \Gamma[\mathbf{C}\text{-MinDeg}_{\leq}(d), k](\tau)$ if and only if the following predicate is satisfied:

- $\mathbf{P}\text{-MinDeg}_{\leq}(d)[k](\tau, (x, y)) \equiv$ for each $u \in [k+1]$, the vertex $\theta[\tau](u)$ in the graph $\mathcal{G}(\tau)$ has degree y_u if $y_u \leq d$ and degree at least $d+1$ if $y_u = d+1$; $x = 1$ if and only if there is some vertex of degree at most d in $V_{\mathcal{G}(\tau)} \setminus \theta[\tau](B(\tau))$.

Each local witness (x, y) can be represented as a binary string containing $1 + (k+1) \cdot \lceil \log(d+1) \rceil$ bits. The DP-core $\mathbf{C}\text{-MinDeg}_{\leq}(d)[k]$ can be defined in such a way that it is deterministic. Therefore, the multiplicity of $\mathbf{C}\text{-MinDeg}_{\leq}(d)[k]$ is 1. This implies that the deterministic state complexity of $\mathbf{C}\text{-MinDeg}_{\leq}(d)[k]$ is upper bounded by $2^{O(k \log d)}$.

E.4 $\mathbf{C}\text{-Colorable}(c)$

The graph property of the DP-core $\mathbf{C}\text{-Colorable}(c)$ is the set $\mathbf{Colorable}(c)$ of all graphs that are c -colorable. For each $k \in \mathbb{N}$, a local witness for $\mathbf{C}\text{-Colorable}(c)[k]$ is a vector in $\mathbf{w} \in \{0, 1, \dots, c\}^{k+1}$. All local witnesses are final. The transitions of $\mathbf{C}\text{-Colorable}(c)[k]$ are defined in such a way that for each k -instructive tree decomposition τ , and each local witness \mathbf{w} , $\mathbf{w} \in \Gamma[\mathbf{C}\text{-Colorable}(c), k](\tau)$ if and only if the following predicate is satisfied:

- $\mathbf{P}\text{-Colorable}(c)[k](\tau, \mathbf{w}) \equiv$ there is a proper c -coloring $\alpha : V_{\mathcal{G}(\tau)} \rightarrow [c]$ where for each $u \in B(\tau)$, $\mathbf{w}_u = \alpha(\theta[\tau](u))$ and for each $u \in [k+1] \setminus B(\tau)$, $\mathbf{w}_u = 0$.

Each local witness can be represented using $\lceil \log(c+1) \rceil \cdot (k+1)$ bits. It is worth noting that every graph of treewidth at most k is $(k+1)$ -colorable, and therefore, for $c \geq k+1$, we can define $\mathbf{C}\text{-Colorable}(c)[k]$ as the trivial DP-core which accepts all k -instructive tree decompositions. This DP-core has a unique local witness, and this unique local witness is final.

E.5 $\mathbf{C}\text{-Conn}$

The graph property of the DP-core $\mathbf{C}\text{-Conn}$ is the set \mathbf{Conn} of all connected graphs. For each $k \in \mathbb{N}$, a local witness for $\mathbf{C}\text{-Conn}[k]$ is a pair (γ, P) where $\gamma \in \{0, 1, 2, 3\}$, and P is a partition of some subset of $[k+1]$. Such a local witness (γ, P) is final if $\gamma \neq 3$ and P has at most one cell (the empty partition with no cell is a legal partition of the empty set). The transitions of $\mathbf{C}\text{-Conn}[k]$ are defined in such a way that for each k -instructive tree decomposition τ , and each local witness (γ, P) , $(\gamma, P) \in \Gamma[\mathbf{C}\text{-Conn}, k](\tau)$ if and only if the following predicate is satisfied:

- $\mathbf{P}\text{-Conn}[k](\tau, (\gamma, P)) \equiv U(P) = B(\tau)$; for each two labels $u, v \in B(\tau)$, u and v are in the same cell of P if and only if $\theta[\tau](u)$ and $\theta[\tau](v)$ belong to the same connected component in the graph $\mathcal{G}(\tau)$; furthermore,

$$\gamma = \begin{cases} 0 & \text{if } \mathcal{G}(\tau) \text{ is the empty graph;} \\ 1 & \text{if } P \neq \emptyset \text{ and every vertex in } \mathcal{G}(\tau) \text{ is reachable from } \theta[\tau](U(P)); \\ 2 & \text{if } P = \emptyset, \mathcal{G}(\tau) \text{ is connected and not the empty graph;} \\ 3 & \text{if } P \neq \emptyset \text{ and some vertex in } \mathcal{G}(\tau) \text{ is not reachable from } \theta[\tau](U(P)), \\ & \text{or } P = \emptyset \text{ and } \mathcal{G}(\tau) \text{ is disconnected.} \end{cases}$$

Each local witness (γ, P) can be represented using $2 + (k+1) \cdot \lceil \log(k+2) \rceil = O(k \log k)$ bits. Additionally, $\mathbf{C}\text{-Conn}$ can be defined in such a way that it has multiplicity 1. Therefore, its deterministic state complexity is upper bounded by $2^{O(k \log k)}$.

E.6 C-VConn_≤(c)

The graph property of the DP-core **C-VConn_≤(c)** is the set of all graphs with vertex-connectivity at most c . A local witness is a triple (r, γ, P) where $r \in \{0, 1, \dots, c\}$, $\gamma \in \{0, 1, 2, 3\}$, and P is a partition of some subset of $[k+1] \setminus R$. A witness is final if $\gamma = 3$ or P has more than one cell (this means that after removing $r \leq c$ vertices the graph gets disconnected). The transitions of **C-VConn_≤(c)[k]** are defined in such a way that for each k -instructive tree decomposition τ , and each local witness (r, γ, P) , $(r, \gamma, P) \in \Gamma[\mathbf{C-VConn}_{\leq}(c), k](\tau)$ if and only if the following predicate is satisfied:

- **P-VConn_≤(c)[k](τ, \mathbf{w})** \equiv there is a subset of vertices $X \subseteq V_{\mathcal{G}(\tau)}$ of size r such that: $U(P) \subseteq B(\tau)$ and for each $u \in B(\tau)$, $u \in U(P)$ if and only if $\theta[\tau](u) \notin X$; for each two labels $u, v \in B(\tau)$, u and v are in the same cell of P if and only if $\theta[\tau](u)$ and $\theta[\tau](v)$ belong to the same connected component in the graph $\mathcal{G}(\tau) \setminus X$; furthermore,

$$\gamma = \begin{cases} 0 & \text{if } \mathcal{G}(\tau) \setminus X \text{ is the empty graph;} \\ 1 & \text{if } P \neq \emptyset \text{ and every vertex in } \mathcal{G}(\tau) \setminus X \text{ is reachable from } \theta[\tau](U(P)); \\ 2 & \text{if } P = \emptyset, \mathcal{G}(\tau) \setminus X \text{ is connected and not the empty graph;} \\ 3 & \text{if } P \neq \emptyset \text{ and some vertex in } \mathcal{G}(\tau) \setminus X \text{ is not reachable from } \theta[\tau](U(P)), \\ & \text{or } P = \emptyset \text{ and } \mathcal{G}(\tau) \setminus X \text{ is disconnected.} \end{cases}$$

Each local witness can be represented using $2 + \lceil \log c \rceil + (k+1) \cdot (1 + \lceil \log(k+2) \rceil) = O(\log c + k \log k)$ bits. It is worth noting that every graph of treewidth at most k has connectivity at most $k+1$, and therefore, for $c \geq k+1$, we can define **C-VConn_≤(c)[k]** as the trivial DP-core which accepts all k -instructive tree decompositions. This DP-core has a unique local witness, and this unique local witness is final.

E.7 C-EConn_≤(c)

The graph property of the DP-core **C-EConn_≤(c)** is the set of all graphs with edge-connectivity at most c . A local witness is a tuple (r, γ, P) where $r \in \{0, 1, \dots, c\}$, $\gamma \in \{0, 1, 2, 3\}$, and P is a partition of some subset of $[k+1]$. A witness is final if $\gamma = 3$ or P has more than one cell (this means that after removing $r \leq c$ edges, the graph gets disconnected). The transitions of **C-EConn_≤(c)[k]** are defined in such a way that for each k -instructive tree decomposition τ , and each local witness (r, γ, P) , $(r, \gamma, P) \in \Gamma[\mathbf{C-EConn}_{\leq}(c), k](\tau)$ if and only if the following predicate is satisfied:

- **P-EConn_≤(c)[k](τ, \mathbf{w})** \equiv there is a subset of edges $Y \subseteq E_{\mathcal{G}(\tau)}$ of size r such that: $U(P) = B(\tau)$; for each two labels $u, v \in B(\tau)$, u and v are in the same cell of P if and only if $\theta[\tau](u)$ and $\theta[\tau](v)$ belong to the same connected component in the graph $\mathcal{G}(\tau) \setminus Y$; furthermore,

$$\gamma = \begin{cases} 0 & \text{if } \mathcal{G}(\tau) \setminus Y \text{ is the empty graph;} \\ 2 & \text{if } P \neq \emptyset \text{ and every vertex in } \mathcal{G}(\tau) \setminus Y \text{ is reachable from } \theta[\tau](U(P)); \\ 1 & \text{if } P = \emptyset, \mathcal{G}(\tau) \setminus Y \text{ is connected and not the empty graph;} \\ 3 & \text{if } P \neq \emptyset \text{ and some vertex in } \mathcal{G}(\tau) \setminus Y \text{ is not reachable from } \theta[\tau](U(P)), \\ & \text{or } P = \emptyset \text{ and } \mathcal{G}(\tau) \setminus Y \text{ is disconnected.} \end{cases}$$

Each local witness can be represented using $2 + \lceil \log c \rceil + (k+1) \cdot \lceil \log(k+2) \rceil = O(\log c + k \log k)$ bits.

E.8 C-Hamiltonian

The graph property of the DP-core **C-Hamiltonian** is the set **Hamiltonian** of all graphs that are Hamiltonian. This DP-core is defined by a straightforward adaptation of the standard algorithm for testing Hamiltonicity parameterized by treewidth. See for instance [81].

For each $k \in \mathbb{N}$, a local witness for $\mathbf{C}\text{-Hamiltonian}(c)[k]$ is a pair (β, M) where $\beta : S \rightarrow \{0, 1, 2\}$ is a function whose domain S is a subset of $[k+1]$ and $M \subseteq \mathcal{P}(\beta^{-1}(1), 2)$ is a matching that relates pairs of labels in S that are sent to the value 1. The transitions of $\mathbf{C}\text{-Hamiltonian}(c)[k]$ are defined in such a way that for each k -instructive tree decomposition τ , and each local witness \mathbf{w} , $\mathbf{w} \in \Gamma[\mathbf{C}\text{-Hamiltonian}(c), k](\tau)$ if and only if the following predicate is satisfied:

- $\mathbf{P}\text{-Hamiltonian}(c)[k](\tau, \mathbf{w}) \equiv$ either $\mathcal{G}(\tau)$ is Hamiltonian and $\beta^{-1}(0) = \beta^{-1}(1) = \emptyset$, or there is a partition \mathcal{P} of $V_{\mathcal{G}(\tau)}$ into vertex-disjoint paths such that for each $u \in B(\tau)$, $\theta[\tau](u)$ has degree $\beta(u)$ in some path of \mathcal{P} .

Each local witness can be represented using $\lceil \log k \rceil \cdot (k+1) = O(k \log k)$ bits. This yields a multiplicity of at most $2^{O(k \log k)}$, and consequently a deterministic state complexity of $2^{2^{O(k \log k)}}$. It turns out that using a clever $\mathbf{D}[k].\mathbf{Clean}$ function, that applies the rank-based approach introduced in [15] to eliminate redundancies, one can guarantee that the number of local witnesses in a useful witness set (i.e. the multiplicity of the DP-core) is always bounded by $2^{|\beta^{-1}(1)|-1} \leq 2^k$. Algorithms with a better multiplicity have been devised in [33]. A nice discussion about the rank based approach and other approaches to solve the Hamiltonian cycle problem on graphs of bounded treewidth is also present in [81]. In other words, this clean function decreases the multiplicity of the DP-core without affecting the existence of a solution. As a consequence, the deterministic state complexity of the DP-core is upper bounded by $\binom{2^{O(k \log k)}}{2^{O(k)}} = 2^{2^{O(k)}}$.

E.9 $\mathbf{C}\text{-NZFlow}(\mathbb{Z}_m)$

We let $\mathbb{Z}_m = \{0, \dots, m-1\}$ be the set of integers modulo m . Let G be a graph and (t, h) be an orientation of G . In other words, $t : E_G \rightarrow V_G$ and $h : E_G \rightarrow V_G$ are maps that specify the tail $t(e)$ and the head $h(e)$ of each edge $e \in E_G$ ($t(e) \neq h(e)$). For each $v \in V_G$, we let $\delta^-(v) = \{e \mid t(v) = e\}$ be the set of edges whose tail is v , and $\delta^+(v) = \{e \mid h(v) = e\}$ be the set of edges whose head is v . We say that v satisfies the flow equation if the following condition is satisfied.

$$\sum_{e \in \delta^+(v)} \phi(e) = \sum_{e \in \delta^-(v)} \phi(e) \quad (14)$$

Definition 39 (Nowhere-Zero \mathbb{Z}_m -Flow). *Let G be a graph and (t, h) be an orientation of G . A nowhere-zero \mathbb{Z}_m -flow in (G, t, h) is a function $\phi : E_G \rightarrow \mathbb{Z}_m$ satisfying the following conditions:*

1. *for each $v \in V$, v satisfies the flow equation, and*
2. *for each $e \in E$, $\phi(e) \neq 0$.*

We say that $\phi : E_G \rightarrow \mathbb{Z}_m$ is a nowhere-zero \mathbb{Z}_m -flow in G if there is an orientation (t, h) such that ϕ is a \mathbb{Z}_m -flow in (G, t, h) .

For each $m \in \mathbb{N}$, we let $\mathbf{NZFlow}(\mathbb{Z}_m)$ be the graph property consisting of all graphs that admit a nowhere-zero \mathbb{Z}_m -flow. For each $k \in \mathbb{N}$, a local witness for $\mathbf{C}\text{-NZFlow}(\mathbb{Z}_m)[k]$ is a set of triples of the form (u, v, f) where u and v belong to $[k+1]$ and $f \in \mathbb{Z}_m$. Such a local witness is final, meaning that $\mathbf{C}\text{-NZFlow}(\mathbb{Z}_m)[k].\mathbf{Final}(u, v, f) = 1$ if and only if $\mathbf{Flow}(\mathbf{w}, u)$ is true for each $u \in \mathbf{Labels}(\mathbf{w})$. The transitions of $\mathbf{C}\text{-MaxDeg}_{\geq}(d)[k]$ are defined in such a way that for each k -instructive tree decomposition τ and each local witness \mathbf{w} , $\mathbf{w} \in \Gamma[\mathbf{C}\text{-NZFlow}(\mathbb{Z}_m), k](\tau)$ if and only if the following predicate is satisfied:

- $\mathbf{P}\text{-NZFlow}(\mathbb{Z}_m)[k](\tau, \mathbf{w}) \equiv$ there is an orientation (t, h) of the graph $\mathcal{G}(\tau)$, and a function $\phi : E_{\mathcal{G}(\tau)} \rightarrow \mathbb{Z}_m$ such that the following conditions are satisfied.
1. For each u and v in $B(\tau)$, and each $f \in \mathbb{Z}_m$, $(u, v, f) \in \mathbf{w}$ if and only if $\sum_e \phi(e) = f$, where e ranges over all edges $e \in E_{\mathcal{G}(\tau)}$ with $t(e) = \theta[\tau](u)$, and $h(e) = \theta[\tau](v)$.

2. For each $u \in B(\tau)$, $\text{Flow}(\mathbf{w}, u) = \text{true}$ if and only if $\theta[\tau](u)$ satisfies the flow equation (Equation 14). In particular, if \mathbf{w} is a final local witness, then for each $u \in B(\tau)$, the vertex $\theta[\tau](u)$ satisfies the flow equation.
3. Each vertex $x \in V_{\mathcal{G}(\tau)} \setminus \text{Im}(\theta[\tau])$, satisfies the flow equation (Equation 14).

Since a local witness for $\mathbf{C}\text{-NZFlow}(\mathbb{Z}_m)[k]$ is a set with $O(k^2)$ triples, such a witness can be represented by a binary vector of length $O(k^2 \cdot \log m)$, where $\log m$ bits are used to represent each flow value.

E.10 C-Minor(H)

Parameterized algorithms for determining whether a pattern graph H is a minor of a host graph G parameterized by the branchwidth of the host graph have been devised in [43], and subsequently improved in [1]. Both algorithms operate with branch decompositions instead of tree decompositions. The later algorithm works in time $O(2^{(2k+1) \cdot \log k} \cdot |V_H|^{2k} \cdot 2^{2|V_H|^2} \cdot |V_G|)$, where k is the width of the input branch decomposition. Partial solutions in both algorithms are combinatorial objects called *rooted packings*.

In this section, we describe the structure of local witnesses of a DP-core $\mathbf{C}\text{-Minor}(\mathbf{H})$ that solves the minor containment problem by analyzing k -instructive tree decompositions, instead of branching decompositions. In our setting, partial solutions are objects called *quasimodels*. A local witness may be regarded as an encoding of the restriction of a quasimodel to the active vertices of a given k -instructive tree decomposition. We note that our local witnesses have bitlength $O(k \cdot \log k + |V_H| + |E_H|)$, and as a consequence, our DP-core decides whether a graph G of treewidth at most k has an H -minor in time $2^{O(k \cdot \log k + |V_H| + |E_H|)} \cdot (|V_G| + |E_G|)$. This improves the algorithm of [1] for sparse pattern graphs.

We will need the following variant of the predicate $\mathbf{P}\text{-Conn}[k](\tau, (\gamma, P))$ defined in Subsection E.5. This variant is parameterized by a subset $X \subseteq \mathbb{N}$, where X is meant to be a subset of $V_{\mathcal{G}(\tau)}$. Note that this variant is obtained essentially, by replacing $\mathcal{G}(\tau)$ with $\mathcal{G}(\tau)[X]$ and by requiring that P is a partition of some subset of $B(\tau)$, instead of a partition of $B(\tau)$.

- $\mathbf{P}\text{-Conn}^X[k](\tau, (\gamma, P)) \equiv X \subseteq V_{\mathcal{G}(\tau)}$; $U(P) \subseteq B(\tau)$ and for each $u \in B(\tau)$, $u \in U(P)$ if and only if $\theta[\tau](u) \in X$; for each two labels $u, v \in B(\tau)$, u and v are in the same cell of P if and only if $\theta[\tau](u)$ and $\theta[\tau](v)$ belong to the same connected component in the graph $\mathcal{G}(\tau)[X]$; furthermore,

$$\gamma = \begin{cases} 0 & \text{if } \mathcal{G}(\tau)[X] \text{ is the empty graph;} \\ 1 & \text{if } P \neq \emptyset \text{ and every vertex in } \mathcal{G}(\tau)[X] \text{ is reachable from } \theta[\tau](U(P)); \\ 2 & \text{if } P = \emptyset, \mathcal{G}(\tau)[X] \text{ is connected and not the empty graph;} \\ 3 & \text{if } P \neq \emptyset \text{ and some vertex in } \mathcal{G}(\tau)[X] \text{ is not reachable from } \theta[\tau](U(P)), \\ & \text{or } P = \emptyset \text{ and } \mathcal{G}(\tau)[X] \text{ is disconnected.} \end{cases}$$

Below, we define the notions of quasimodel and model of a graph H in a graph G . Intuitively, a model should be regarded as a certificate that H is a minor of G , while a quasi-model is a substructure that can be potentially extended to a model.

Definition 40 (Quasimodels and Models). *Let H and G be graphs. A quasimodel of H in G is a pair of sequences $M = ([X_x]_{x \in V_H}, [y_e]_{e \in E_H})$ satisfying the following conditions:*

1. for each $x \in V_H$, X_x is a subset of V_G ,
2. for each two distinct $x, x' \in V_H$, X_x is disjoint from $X_{x'}$,
3. for each $e \in E_H$ with endpoints $\{x, x'\}$, either y_e is an edge in E_G with one endpoint in X_x and another endpoint in $X_{x'}$, or $y_e = 0$,

4. for each two distinct $e, e' \in E_H$, $y_e \neq y_{e'}$.

We say that M is a model of H in G if additionally, for each $x \in V_H$, $G[X_x]$ is a connected subgraph of G , and for each $e \in E_H$, $y_e > 0$.

We say that a graph H is a minor of a graph G if there is a model of H in G . The graph property of the DP-core $\mathbf{C-Minor}(H)$ is the set of all graphs that have a model of H . A local witness is a pair of tuples

$$\mathbf{w} = ([(\gamma_x, P_x)]_{x \in V_H}, [b_e]_{e \in E_H}) \quad (15)$$

where for each $x \in V_H$, $\gamma_x \in \{0, 1\}$, P_x is a partition of some subset of $[k+1]$, and for each $e \in E_H$, $b_e \in \{0, 1\}$. Such a local witness is final if and only if $b_e = 1$ for every $e \in E_H$, and for each $x \in V_H$, $\gamma_x \neq 3$ and P_x has at most one cell. Note that this last condition, imposed on the pair (γ_x, P_x) , is just the condition for a local witness to be final with respect to the DP-core $\mathbf{C-Conn}$ defined in Subsection E.5. Intuitively, this will be used to certify that for some $X_x \subseteq V_{\mathcal{G}(\tau)}$, the induced subgraph $\mathcal{G}(\tau)[X_x]$ is connected.

The transitions of $\mathbf{C-Minor}(H)[k]$ are defined in such a way that for each k -instructive tree decomposition τ , and each local witness $\mathbf{w} = ([(\gamma_x, P_x)]_{x \in V_H}, [b_e]_{e \in E_H})$, $\mathbf{w} \in \Gamma[\mathbf{C-Minor}(H), k](\tau)$ if and only if the following predicate is satisfied:

- $\mathbf{P-Minor}(H)[k](\tau, \mathbf{w}) \equiv$ there exists a quasi-model $([X_x]_{x \in V_H}, [y_e]_{e \in E_H})$ of H in G such that
 1. for each $x \in V_H$, $\mathbf{P-Conn}^{X_x}[k](\tau, (\gamma_x, P_x)) = 1$, and
 2. for each $e \in E_H$, $b_e = 1$ if and only if $y_e > 0$.

Each local witness $([(\gamma_x, P_x)]_{x \in V_H}, [b_e]_{e \in E_H})$ can be straightforwardly represented using

$$|V_H| \cdot (2 + k \cdot \lceil \log(k+2) \rceil) + |E_H| = O(|V_H| \cdot k \cdot \log k + |E_H|)$$

bits, given that each partition P_x can be represented by at most $k \log k$ bits. Nevertheless, this upper bound can be improved to $O(k \log k + |V_H| + |E_H|)$ by noting that one can assume that the vertices of V_H are ordered (any fixed ordering) and that for each $x, x' \in V_H$ with $x \neq x'$, each cell of P_x is disjoint from each cell of $P_{x'}$. More specifically, we may represent the sequence $[P_x]_{x \in V_H}$ as a sequence of symbols from the alphabet $\{1, \dots, k+1\} \cup \{ |, : \}$, where two consecutive occurrences of the symbol $|$ enclose labels occurring in some P_x , while the symbol $:$ is used to separate cells inside the partition. Numbers occurring after the x -th occurrence symbol $|$ and before the $(x+1)$ -th occurrence of this symbol correspond to partition P_x . The symbol $:$ is used to separate cells within two consecutive occurrences of the symbols $|$. Since no number in the set $\{1, \dots, k+1\}$ occurs twice in the sequence, the symbol $:$ occurs at most k times, and the symbol $|$ occurs at most $|V_H|$ times, the length of this sequence is at most $(k+1) + k + |V_H|$. Note that the numbers in $\{1, \dots, k+1\}$ require each $O(\log k)$ bits to be represented, while the symbols $|$ and $:$ can be represented with 2 bits. Therefore, the overall representation of the witness, including the vector $[b_e]_{e \in E_H}$ has $O(k \log k + |V_H| + |E_H|)$ bits.