Definable decompositions for graphs of bounded linear cliquewidth*

Mikołaj Bojańczyk† Martin Grohe‡ Michał Pilipczuk§

Abstract

We prove that for every positive integer k, there exists an $\textsc{mso}_1$-transduction that given a graph of linear cliquewidth at most k outputs, nondeterministically, some clique decomposition of the graph of width bounded by a function of k. A direct corollary of this result is the equivalence of the notions of $\textsc{cmso}_1$-definability and recognizability on graphs of bounded linear cliquewidth.

1 Introduction

Hierarchical decompositions of graphs have come to play an increasingly important role in logic, algorithms and many other areas of computer science. The treelike structure they impose often allows to process the data much more efficiently. The best-known and arguably most important graph decompositions are *tree decompositions*, which play a central role in a research direction at the boundary between logic, graph grammars, and generalizations of automata theory to graphs that was pioneered by Courcelle in the 1990s (see [5]). Recently, the first and third author of this paper answered a long standing open question in this area by showing that on graphs of bounded treewidth, the automata-theory-inspired notion of *recognizability* coincides with *definability* in monadic second order logic with modulo counting [2].

A drawback of tree decompositions is that they only yield meaningful results for sparse graphs. A suitable form of decomposition that also applies to dense graphs and that has a similarly nice, yet less developed Courcelle-style theory, is that of *clique decompositions*, introduced by Courcelle and Olariu [7]. A clique decomposition of a graph is a term in a suitable algebra consisting of (roughly) the following operations for constructing and manipulating colored graphs: (i) disjoint union; (ii) for a pair of colors i, j, simultaneously add an edge for every pair (i-colored vertex, j-colored vertex); and (iii) apply recolorings to entire colors. A natural notion of *width* for such a clique decomposition is the total number of colors used. The *cliquewidth* of a graph is the smallest width of a clique decomposition for it. An alternative notion of graph decomposition that also works well for dense graphs is that of *rank decompositions* introduced by Oum and Seymour [12, 13]. The corresponding notion of *rankwidth* turned out to be functionally equivalent to cliquewidth, that is, bounded cliquewidth is the same as bounded rankwidth. Bounded width clique or rank decompositions may be viewed as hierarchical decompositions that minimize "modular complexity" of cuts present in the decomposition, in the same way as treewidth corresponds to hierarchical decompositions using vertex cuts, where the complexity of a cut is its size.

*The work of M. Bojańczyk and Mi. Pilipczuk was supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (ERC consolidator grant LIPA, agreement no. 683080). Mi. Pilipczuk was supported by the Foundation for Polish Science via the START stipend programme.

†Institute of Informatics, University of Warsaw, Poland, bojan@mimuw.edu.pl

‡RWTH Aachen University, Germany, grohe@informatik.rwth-aachen.de

§Institute of Informatics, University of Warsaw, Poland, michal.pilipczuk@mimuw.edu.pl

Cliquewidth is tightly connected to $\textsc{mso}_1$ logic on graphs, in the same way as the $\textsc{mso}_2$ logic is connected to treewidth. Recall that in $\textsc{mso}_1$, one can quantify over vertices and sets of vertices, and check their adjacency, while $\textsc{mso}_2$ also allows quantification over sets of edges. Both these logics can be viewed as plain $\textsc{mso}$ logic on two different encodings of graphs as relational structures: for $\textsc{mso}_1$ the encoding uses only vertices as the universe and has a binary adjacency relation, while for $\textsc{mso}_2$ the encoding uses both vertices and edges as the universe and has an incidence relation binding every edge with its endpoints. These two logics are connected to cliquewidth and treewidth as follows. If a graph property Π is definable in $\textsc{mso}_2$, then tree decompositions of graphs in Π can be recognized by a finite state device (tree automaton). This leads, for instance, to a fixed-parameter model checking algorithm for $\textsc{mso}_2$-definable properties on graphs of bounded treewidth [4]. This notion of recognizability, where tree decompositions are processed, is called *HR-recognizability* [5]. Similarly, if Π is $\textsc{mso}_1$-definable, then clique decompositions of graphs in Π can be recognized by a finite state device. This notion of recognizability is called *VR-recognizability* [5], and it yields a fixed-parameter model checking algorithm for $\textsc{mso}_1$-definable properties on graphs of bounded cliquewidth [6].

It was conjectured by Courcelle [4] that $\textsc{mso}_2$-definability and recognizability for tree decompositions (i.e. HR-recognizability) are equivalent for every graph class of bounded treewidth, provided that $\textsc{mso}_2$ is extended by counting predicates of the form "the size of X is divisible by p", for every integer p (this logic is called $\textsc{cmso}_2$). This conjecture has been resolved by two of the current authors [2]. More precisely, in [2] it was shown that for every k there exists an $\textsc{mso}$ transduction which inputs a graph of treewidth at most k and (nondeterministically) outputs its tree decomposition of width bounded by a function of k. The graph is given via its incidence encoding. The conjecture of Courcelle then follows by composing this transduction with guessing the run of an automaton recognizing the property in question on the output decomposition.

The same question can be asked about cliquewidth: is it true that every property of graphs of bounded cliquewidth is $\textsc{mso}_1$-definable if and only if it is (VR-)recognizable? The present paper discusses this question, proving a special case of the equivalence.

Our contribution. Our main result (Theorem 2) is that for every $k \in \mathbb{N}$, there exists an $\textsc{mso}$-transduction which inputs a graph of *linear* cliquewidth at most k, and outputs a clique decomposition of it which has width bounded by a function of k. Here, we use the adjacency encoding of the graph. The *linear cliquewidth* of a graph is a linearized variant of cliquewidth, similarly as pathwidth is a linearized variant of treewidth; see Section 2 for definition and, e.g., [1, 8, 9, 10] for more background. An immediate consequence of this result (Theorem 3) is that every property of graphs of bounded linear cliquewidth is $\textsc{cmso}_1$-definable if and only if it is (VR-)recognizable. This gives a partial answer to the question above.

The proof of our main result shares one key idea with the proof of the $\textsc{mso}_2$-definability of tree decompositions, or more precisely, the pathwidth part of that proof [2, Lemma 2.5]. This is the use of Simon's Factorization Forest Theorem [14]. We view a linear clique decomposition of width k as a word over a finite alphabet and use the factorization theorem to construct a nested factorization of this word of depth bounded in terms of k. The overall $\textsc{mso}$ transduction computing a decomposition is then constructed by induction on the nesting depth of this factorization. The technical challenge in this paper is to analyze the composition of "subdecompositions", which is significantly more complicated in the cliquewidth case than in the treewidth/pathwidth case of [2]. In a path decomposition, each node of the path (over which we decompose) naturally corresponds to a separation of the graph, with the bag at the node being the separator. Thus, in the pathwidth case, each separation appearing in the decomposition essentially can be described by a tuple of vertices in the separator, with the left and the right side being essentially independent; this is a simple and easy to handle object. The difficulty in the cliquewidth case is that "separations" appearing in a linear clique decomposition are partitions of the vertex set into two sides with small "modular complexity": each side can be partitioned further into a bounded number of parts so that vertices from the same part have exactly the same neighbors on the other side. Such separations are much harder to control combinatorially, and hence capturing them using the resources of MSO requires a deep insight into the combinatorics of linear cliquewidth.

2 Preliminaries

Graphs and cliquewidth. All graphs considered in this paper are finite and simple. For the most part we use undirected graphs, if we use directed graphs than we remark this explicitly. We write $[k]$ for $\{1, 2, \ldots, k\}$ and $\binom{X}{1,2}$ for the family of nonempty subsets of a set X of size at most 2.

A *k-colored graph* is a graph with each vertex assigned a color from $[k]$. On k-colored graphs we define the following operations.

- **Recolor.** For every function $\phi\colon [k] \to [k]$ there is a unary operation which inputs one k-colored graph and outputs the same graph where each vertex is recolored to the image of its original color under ϕ.
- **Join.** For every family of subsets $S \subseteq \binom{[k]}{1,2}$ there is an operation that inputs a family of k-colored graphs, of arbitrary finite size, and outputs a single k-colored graph constructed as follows. Take the disjoint union of the input graphs and for each $\{i, j\} \in S$ (possibly $i = j$), add an edge between every pair of vertices that have colors i and j, respectively, and originate from different input graphs.
- **Constant.** For each color $i \in [k]$ there is a constant which represents a graph on a single vertex with color i.

Define a *width-k clique decomposition* to be a (rooted) tree where nodes are labelled by operation names in an arity preserving way, that is, all constants are leaves and all recolor operations have exactly one child. The tree does not have any order on siblings, because Join is a commutative operation. For a clique decomposition, we define its *result* to be the k-colored graph obtained by evaluating the operations in the decomposition. The *cliquewidth* of a graph is defined to be the minimum number k for which there is a width-k clique decomposition whose result is (some coloring of) the graph.

We remark that we somewhat diverge from the original definition of cliquewidth [7] in the following way. In [7], there is one binary disjoint union operation that just adds two input k-colored graphs, and for each pair of different colors i, j there is a unary operation that creates an edge between every pair of vertices of colors i and j, respectively. For our purposes, we need to have a union operation that takes an arbitrary number of input graphs. This is because an MSO transduction constructing a clique decomposition cannot break symmetries and join isomorphic parts of the graph in some arbitrarily chosen order, which would be necessary if we used binary disjoint union operations. Another difference is that our join does simultaneously two operations: it takes the disjoint union of several inputs, and adds edges between them. When using binary joins, the two operations can be separated by introducing temporary colors; however when the number of arguments is unbounded such a separation is not possible. It is easy to show that our definition of cliquewidth is at multiplicative factor at most 2 from the original definition.

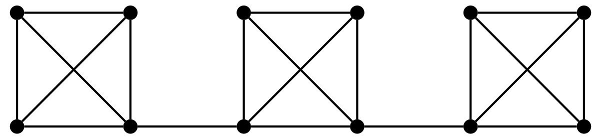

Figure 1: A graph of linear clique width 3

Linear cliquewidth is a linearized variant of cliquewidth, where we allow only restricted joins that add only a single vertex. More precisely, we replace the Join and Constant operations with one unary operation **Add Vertex**. This operation is parameterized by a color $i \in [k]$ and a color subset $X \subseteq [k]$, and it adds to the graph a new vertex of color i, adjacent exactly to vertices with colors belonging to X. A *width-k linear clique decomposition* is a word consisting of Add Vertex and Recolor operations, and the *result* of such a decomposition is the k-colored graph obtained by evaluating the operations over the empty graph. The *linear cliquewidth* of a graph is defined just like cliquewidth, but we consider only linear clique decompositions. Note that we can transform any linear clique decomposition of width k to a clique decomposition of width at most $(k+1)$ by replacing each subterm $\mathbf{Add\ Vertex}_{i,X}(\theta)$ by the term

$$\mathbf{Recolor}_{j\mapsto i}\Big(\mathbf{Join}_{\{\{j,x\}\colon x\in X\}}\big(\theta, \mathbf{Color}_j\big)\Big),$$

where j is a color not occurring in θ. Hence the linear cliquewidth of a graph is at least its cliquewidth minus one.

Example 1. Consider the graph G displayed in Figure 1. We will first argue that its cliquewidth is at most 3, and then show that also its linear clique width is at most 3.

We first construct a clique decomposition of G of width 3. The following three terms $\theta_1, \theta_2, \theta_3$ construct the three 4-cliques of G with appropriate colors:

$$\begin{aligned}
\theta_1 &= \mathbf{Join}_{\{\{1\},\{1,2\}\}}\Big(\big\{\mathbf{Color}_1, \mathbf{Color}_1, \mathbf{Color}_1, \mathbf{Color}_2\big\}\Big),\\
\theta_2 &= \mathbf{Join}_{\{\{1\},\{1,2\},\{1,3\},\{2,3\}\}}\Big(\big\{\mathbf{Color}_1, \mathbf{Color}_1, \mathbf{Color}_2, \mathbf{Color}_3\big\}\Big),\\
\theta_3 &= \mathbf{Join}_{\{\{1\},\{1,3\}\}}\Big(\big\{\mathbf{Color}_1, \mathbf{Color}_1, \mathbf{Color}_1, \mathbf{Color}_3\big\}\Big).
\end{aligned}$$

Then the term

$$\mathbf{Join}_{\{\{2\},\{3\}\}}\left(\{\theta_1, \theta_2, \theta_3\}\right)$$

is a width-3 clique decomposition of the graph G.

For a linear clique decomposition, it is convenient to denote the unary operation $\mathbf{Add\ Vertex}_{i,X}$ of adding a vertex of color i and connecting it to all vertices of color X by $\boldsymbol{a}_{i,X}$ and the recoloring operation by $\boldsymbol{r}_\phi$. Moreover, we apply both operations by multiplication from left to right, omitting parenthesis. This way, we may view a linear clique decomposition of width k as a word over the finite alphabet

$$\{\boldsymbol{a}_{i,X} \mid i \in [k], X \subseteq 2^{[k]}\} \cup \{\boldsymbol{r}_\phi \mid \phi\colon [k] \to [k]\}.$$

With this notation, our linear clique decomposition of the graph G looks as follows:

$$\underbrace{\boldsymbol{a}_{1,\emptyset}\boldsymbol{a}_{1,\{1\}}\boldsymbol{a}_{1,\{1\}}\boldsymbol{a}_{2,\{1\}}}_{\text{first clique}}\underbrace{\boldsymbol{a}_{3,\{2\}}\boldsymbol{r}_{\{2\mapsto 1,3\mapsto 2\}}\boldsymbol{a}_{2,\{2\}}\boldsymbol{a}_{2,\{2\}}\boldsymbol{a}_{3,\{2\}}\boldsymbol{r}_{\{2\mapsto 1,3\mapsto 2\}}}_{\text{second clique}}\underbrace{\boldsymbol{a}_{3,\{2\}}\boldsymbol{a}_{3,\{3\}}\boldsymbol{a}_{3,\{3\}}\boldsymbol{a}_{3,\{3\}}}_{\text{third clique}}.$$

Relational structures and logic. Define a *vocabulary* to be a set of *relation names*, each one with associated arity in $\mathbb{N}$. A *relational structure* over the vocabulary Σ consists of a set called the *universe*, and for each relation name in the vocabulary, an associated relation of the same arity over the universe. Note the possibility of relations of arity zero, such a relation stores a single bit of information about the structure. A graph is encoded as a relational structure as follows: the universe is the vertex set, and there is one symmetric binary relation that encodes adjacency.

A width-k clique decomposition of a graph is modeled as a relational structure whose universe is the set of nodes of the decomposition, there is a binary predicate "child", and for each operation from the definition of a clique decomposition there is a unary predicate (the set of these predicates depends on k) which selects nodes that use this operation. Note that the graph itself is not included in this structure, but it is straightforward to reconstruct it using an MSO transduction (see below).

To describe properties of relational structures, we use monadic second-order logic (MSO). This logic allows quantification both over single elements of the universe and also over subsets of the universe. For a precise definition of MSO, see [5]. We will also use counting MSO, denoted also by CMSO, which is the extension of MSO with predicates of the form "the size of X is divisible by p" for every $p \in \mathbb{N}$.

MSO transductions. We use the same notion of MSO transductions as in [2, 3]. For the sake of completeness, we now recall the definition of an MSO transduction, which is taken verbatim from [3]. We note that our MSO transductions differ syntactically from those used in the literature, see e.g. Courcelle and Engelfriet [5], but are essentially the same.

Suppose that Σ and Γ are finite vocabularies. Define a *transduction* with input vocabulary Σ and output vocabulary Γ to be a set of pairs

$$(\text{input structure over } \Sigma, \text{ output structure over } \Gamma)$$

which is invariant under isomorphism of relational structures. Note that a transduction is a relation and not necessarily a function, thus it can have many different possible outputs for the same input.

An MSO transduction is any transduction that can be obtained by composing a finite number of *atomic transductions* of the following kinds. Note that kind 1 is a partial function, kinds 2, 3, 4 are functions, and kind 5 is a relation.

1. **Filtering.** For every MSO sentence φ over the input vocabulary there is transduction that filters out structures where φ is satisfied. Formally, the transduction is the partial identity whose domain consists of the structures that satisfy the sentence. The input and output vocabularies are the same.

2. **Universe restriction.** For every MSO formula $\varphi(x)$ over the input vocabulary with one free first-order variable there is a transduction, which restricts the universe to those elements that satisfy φ. The input and output vocabularies are the same, the interpretation of each relation in the output structure is defined as the restriction of its interpretation in the input structure to tuples of elements that remain in the universe.

3. **Mso interpretation.** This kind of transduction changes the vocabulary of the structure while keeping the universe intact. For every relation name R of the output vocabulary, there is an MSO formula $\varphi_R(x_1, \ldots, x_k)$ over the input vocabulary which has as many free first-order variables as the arity of R. The output structure is obtained from the input structure by keeping the same universe, and interpreting each relation R of the output vocabulary as the set of those tuples $(x_1, \ldots, x_k)$ that satisfy φ_R.

4. **Copying.** For $k \in \{1, 2, \ldots\}$, define k-copying to be the transduction which inputs a structure and outputs a structure consisting of k disjoint copies of the input. Precisely, the output universe consists of k copies of the input universe. The output vocabulary is the input vocabulary enriched with a binary predicate copy that selects copies of the same element, and unary predicates $\mathsf{layer}_1, \mathsf{layer}_2, \ldots, \mathsf{layer}_k$ which select elements belonging to the first, second, etc. copies of the universe. In the output structure, a relation name R of the input vocabulary is interpreted as the set of all those tuples over the output structure, where the original elements of the copies were in relation R in the input structure.

5. **Coloring.** We add a new unary predicate to the input structure. Precisely, the universe as well as the interpretations of all relation names of the input vocabulary stay intact, but the output vocabulary has one more unary predicate. For every possible interpretation of this unary predicate, there is a different output with this interpretation implemented.

Note that each element v' of the output structure of an MSO-transduction is either identical to or a copy of an element v of the input structure. We call this element v the *origin* of v'. Thus we have a well-defined *origin mapping* from the output structure to the input structure. In general, this mapping is neither injective nor surjective.

Define the *size* of an atomic MSO transduction to be the size of its input and output vocabularies, plus the maximal quantifier rank of MSO formulas that appear in it (if the atomic type uses MSO formulas). Define the *size* of an MSO transduction to be the sum of sizes of atomic transductions that compose to the transduction. Note that there are finitely many MSO transductions of a given size, since there are finitely many MSO formulas (up to logical equivalence) once the vocabulary, the free variables, and the quantifier rank are fixed.

An MSO-transduction is *deterministic* if it uses no coloring. Note that a deterministic MSO-transduction is a partial function, that is, for each input structure there is at most one output structure.

Note that the composition of two MSO transductions is an MSO transduction by definition. Another well-known property that we will use, as expressed in the following lemma, is that the union of two MSO transductions is also an MSO transductions; recall that here we regard MSO transductions as relations between input and output structures. This property is Lemma 7.18 from [5]. Since our notion of an MSO transduction is a bit different from the one used in [5], we give a proof for completeness.

Lemma 1. *The union of two* MSO *transductions with the same input and output vocabularies is also an* MSO *transduction.*

Proof. First, using copying create two copies of the universe, called further the first and the second *layer*. Apply the first transduction only to the first layer, thus turning it into a (nondeterministically chosen) result of the first transduction. More precisely, all formulas used in the first transduction are relativized to the first layer, or any new elements originating from them. Also, each copying step is followed by an additional universe restriction step that removes unnecessary copies of the second layer. Then, analogously apply the second transduction only to the second layer. After this step, the structure is a disjoint union of some result of the first transduction applied to the initial structure, and some result of the second transduction applied to the initial structure. It remains to nondeterministically choose one of these results using coloring, and remove the other one using universe restriction. □

The key property of MSO transductions is that CMSO- and MSO-definable properties are closed under taking inverse images over MSO transductions. More precisely, we have the following.

Lemma 2 (Backwards Translation Theorem, [5])**.** *Let Σ, Γ be finite vocabularies and let $\mathcal{I}$ be an* MSO *transduction with input vocabulary Σ and output vocabulary Γ. Then for every* MSO *(resp.* CMSO*) sentence ψ over Γ there exists an* MSO *(resp.* CMSO*) sentence $\varphi = \mathcal{I}^{-1}(\psi)$ over Σ such that φ holds in exactly those Σ-structures on which $\mathcal{I}$ produces at least one output satisfying ψ.*

Simon Lemma. As we mentioned in Section 1, the main technical tool used in this work will be the Simon's Factorization Theorem [14]. We will use the following variant, which is an easy corollary of the original statement. Recall that a *semigroup* is an algebra with one associative binary operation, usually denoted as multiplication, and that an *idempotent* in a semigroup is an element e such that $e \cdot e = e$.

Lemma 3 (Simon Lemma)**.** *Suppose that S and T are semigroups, where S is finitely generated (but possibly infinite) and T is finite. Suppose further that $h\colon S \to T$ is a semigroup homomorphism and $f\colon \mathbb{N} \to \mathbb{N}$ and $\mu\colon S \to \mathbb{N}$ are functions such that*

$$\mu(s_1 \cdot \ldots \cdot s_n) \leqslant f(\max_{i \in [n]} \mu(s_i)) \tag{1}$$

holds whenever $n = 2$ or there is some idempotent $e \in T$ such that $e = h(s_1) = \ldots = h(s_n)$. Then μ has finite range, i.e. there exists $K \in \mathbb{N}$ such that $\mu(s) \leqslant K$ for all $s \in S$.

Proof. Before we proceed to the proof itself, we first recall the original statement of Simon's Factorization Theorem, as described by Kufleitner [11]. Suppose Σ is a finite alphabet and let Σ^+ is the semigroup of nonempty words over Σ with concatenation. Suppose further we are given a finite semigroups T and a homomorphism $h\colon \Sigma^+ \to T$.

For a word $u \in \Sigma^+$ of length more than 1, we define two types of factorizations:

- **Binary**: $u = u_1 u_2$ for some $u_1, u_2 \in \Sigma^+$, and
- **Idempotent**: $u = u_1 \cdots u_n$ for some $u_1, \ldots, u_n \in \Sigma^+$ such that all words u_i have the same image under h, which is moreover an idempotent in T.

Define the *h-rank* of a word $u \in \Sigma^+$ as follows. If u has length 1 then its h-rank is 1. Otherwise, we define the h-rank of u as

$$1 + \min_{u = u_1 \ldots u_n} \max_{i \in [n]} h\text{-rank of } u_i,$$

where the minimum is over all binary or idempotent factorizations of u. Simon's Factorization Theorem can be then stated as follows.

Theorem 1 (Simon's Factorization Theorem, [11, 14])**.** *If h is a homomorphism from Σ^+ to a finite semi-group T, then every word from Σ^+ has h-rank at most $3|T|$.*

The existence of an upper bound expressed only in terms of $|T|$ was first proved by Simon [14], while the improved upper bound of $3|T|$ is due to Kufleitner [11].

We proceed to the proof of our Simon Lemma. Let $\Sigma = \{g_1, \ldots, g_k\}$ be the generators of S, and let $\iota\colon \Sigma^+ \to S$ be the natural homomorphism that computes the product of sequences of generators in S. Consider the homomorphism $h'\colon \Sigma^+ \to T$ defined as the composition of ι and h.

We prove the following claim: for each $d \in \mathbb{N}$ there is a number K_d such that for every word $u \in \Sigma^+$ of h'-rank at most d, we have $\mu(\iota(u)) \leqslant K_d$. Observe that this will finish the proof for the following reason. By Theorem 1, every element $s \in S$ can be expressed as $s = \iota(u)$ for some $u \in \Sigma^+$ of h'-rank at most $3|T|$. Hence we can take $K = K_{3|T|}$.

We prove the claim by induction on d. For $d = 1$ we have that u has to consist of one symbol, so we can take

$$K_1 = \max_{i \in [k]} \mu(h(g_i)).$$

Suppose then that $d > 2$ and take any word u of h'-rank equal to d. By the definition of the h'-rank, u admits a factorization $u = u_1 \dots u_n$ into factors u_i of h'-rank smaller than d, such that either $n = 2$, or all words u_i have the same image under h', which is moreover an idempotent in T. By the supposition of the lemma and induction assumption, we have

$$\mu(\iota(u_1 \cdots u_n)) = \mu(\iota(u_1) \cdots \iota(u_n)) \leqslant f(\max_{i \in [n]} \mu(\iota(u_i))) \leqslant f(K_{d-1}).$$

Hence we can take $K_d = \max(K_{d-1}, f(K_{d-1}))$. □

3 Main results

Statement of the main result. Our main result is that for every k, there is an MSO transduction which maps every graph of linear cliquewidth k to some of its clique decompositions. The width of these decompositions is bounded by a function of k; we do not achieve the optimal value k. To state this result, we introduce a graph parameter, called *definable cliquewidth*, which measures the size of an MSO transduction necessary to transform the graph into its clique decomposition. Since the size of a transduction also includes its vocabulary and the width of the clique decompositions is visible in the output vocabulary, it follows that the definable cliquewidth is at least the cliquewidth.

Recall that we model a width-k clique decomposition of a graph as a (rooted) tree labelled by an alphabet of operations depending on k. Such a clique decomposition t constructs a graph G_t whose vertices are the leaves of the tree. More generally, we say that t is a clique decomposition of a graph G if there is an isomorphism from G_t to G.

Definition 1 (Decomposer)**.** A *width-k decomposer* is an MSO transduction $\mathcal{D}$ from the vocabulary of graphs to the vocabulary of width-k clique decompositions such that for every input-output pair (G, t) of $\mathcal{D}$ the following two conditions are satisfied.

(a) t is a width-k clique decomposition of G.

(b) The origin mapping from t to G restricted to the leaves of t is an isomorphism from G_t to G.

Condition (b) in the definition of decomposers may seem unnecessarily restrictive, but in fact will turn out to be very useful in the technical arguments (see Section 5.1). Furthermore, natural transductions satisfying (a) also tend to satisfy (b), because usually such transduction proceed by building the tree of a clique decomposition on top of the input graph.

Note that the *size* of a decomposer (as a particular MSO-transduction) is an upper bound for its width, because the size of a transduction is larger than the size of its output vocabulary.

Definition 2 (Definable cliquewidth)**.** The *definable cliquewidth* of a graph G, denoted by $\mathrm{dcw}(G)$, is the smallest size of a decomposer which produces at least one output on G.

Note that there are finitely many decomposers of a given size, and decomposers are closed under union by Lemma 1. Therefore, for every k there is a single decomposer (of width k and size $f(k)$) which produces at least one output on every graph with definable cliquewidth at most k, namely one can take the union of all decomposers of size at most k.

The main result of this paper is the following.

Theorem 2. *For every $k \in \mathbb{N}$ there exist a decomposer $\mathcal{D}$ that for every graph G of linear cliquewidth at most k produces at least one output. In other words, the definable cliquewidth of a graph is bounded by a function of its linear cliquewidth.*

The result above could be improved in two ways: first, we could make the transduction produce results for graphs of bounded cliquewidth (and not bounded linear cliquewidth), and second, we could produce clique decompositions of optimum width. We leave both of these improvements to future work. Note that it is impossible to find a decomposer which produces a linear clique decomposition for every graph of linear cliquewidth k; the reason is that such a decomposer would impose a total order on the vertices of the input graph, and this is impossible for some graphs, such as large independent sets.

We remark that, similarly to the case of treewidth [2], our proof is effective: the decomposer $\mathcal{D}$ can be computed from k. This essentially follows from a careful inspection of the proofs, so we usually omit the details in order not to obfuscate the main ideas with computability issues of secondary importance. There is, however, one step in the proof (Lemma 13) where computability of a bound is non-trivial, hence there we present an explicit discussion.

Recognizability. We now state an important corollary of the main theorem, namely that for graph classes with linear cliquewidth, being definable in (counting) MSO is the same as being recognizable. Let us first define the notion of recognizability that we use. For $k \in \mathbb{N}$, define a *k-context* to be a width-k clique decomposition with one distinguished leaf. If t is a k-context and G is a k-colored graph, then $t[G]$ is defined to be the k-colored graph obtained by replacing the distinguished leaf of t by G, and then applying all the operations in t.

Definition 3 (Recognizability, see [5], Def. 4.29)**.** Let L be a class of graphs. Two k-colored graphs G_1, G_2 are called L-equivalent if for every k-context t we have $t[G_1] \in L$ iff $t[G_2] \in L$, where membership in L is tested after ignoring the coloring. We say that L is *recognizable* if for every $k \in \mathbb{N}$ there are finitely many equivalence classes of L-equivalence.

Theorem 5.68(2) in [5] shows that if a class of graphs is definable in MSO (in the sense used here, i.e., MSO_1), then it is recognizable (in the sense of Definition 3, i.e. VR-recognizable). The converse implication is not true, e.g., there are uncountably many VR-recognizable graph classes. The following result, which is a corollary of our main theorem, says that the converse implication is true under the assumption of bounded linear cliquewidth.

Theorem 3. *If L is a class of graphs of bounded linear cliquewidth, then L is recognizable if and only if it is definable in* CMSO*.*

Proof. As mentioned above, the right-to-left implication is true even without assuming a bound on linear cliquewidth. For the converse, we use the following claim; since the proof is completely standard, we only sketch it.

Claim 1. If a class of graphs L is recognizable, then for every k the following language L_k of labelled trees is definable in CMSO.

$$L_k = \{t \colon t \text{ is a tree that is a width-}k\text{ clique decomposition whose resulting graph is in } L\}$$

Proof sketch. The language L_k is a set of (unranked) trees without sibling order. Define $\tilde{L}_k$ to be the language of sibling-ordered trees such that if the sibling order is ignored, then the resulting tree belongs L_k. Using the assumption that L is recognizable, one shows that $\tilde{L}_k$ is definable in

MSO; the idea is that using the sibling order an MSO formula can convert a tree into one which has binary branching, and then compute for each subtree its L-equivalence class. As shown in [4], if a language of sibling-ordered trees is definable in MSO and invariant under reordering siblings, then the language of sibling-unordered trees obtained from it by ignoring the sibling order is definable in CMSO without using the sibling order. Applying this to $\tilde{L}_k$ and L_k we obtain the claim. ⌟

Using Claim 1, we complete the left-to-right implication. Assume every graph from L has linear cliquewidth at most k. Apply Theorem 2, yielding a decomposer $\mathcal{D}$ from graphs to width-ℓ clique decompositions which produces at least one output for every graph in L. Apply Claim 1 to L and ℓ. Since $\mathcal{D}$ produces at least one output for every graph in L, we have that L is the inverse image under $\mathcal{D}$ of the language L_ℓ in the conclusion of the claim. It follows from the Backwards Translation Theorem that L is definable in CMSO. □

4 The proof strategy

In this section we present the proof strategy for our main contribution, Theorem 2.

A linear clique decomposition of width k, being a single path, can be viewed as a sequence of instructions. For such sequences of instructions (actually, for a similar but slightly more general object), we will use the name *k-derivations*. Intuitively speaking, a k-derivation corresponds to an infix of a linear clique decomposition of width k. We can concatenate k-derivations, which means that the set of k-derivations is endowed with a semigroup structure. The main idea is to use Simon's Factorization Theorem [14], in the flavor delivered by the Simon Lemma (Lemma 3), to factorize this product into a tree of bounded depth, so that definable clique decompositions of factors can be constructed via a bottom-up induction over the factorization.

More precisely, Simon Lemma is used to prove Theorem 2 as follows. As the semigroup S we use k-derivations. As the homomorphism h, we use a notion *abstraction*, which maps each k-derivation to a bounded-size combinatorial object consisting of all the information we need to remember about it. Composing k-derivations naturally corresponds to composing their abstractions, which formally means that the set of abstractions, whose size is bounded in terms of k, can be endowed with a semigroup structure so that taking an abstraction of a k-derivation is a semigroup homomorphism. By taking μ to be the definable cliquewidth of a graph, we use the Simon Lemma to show that μ has a finite range on the set of all k-derivations, i.e. there is a finite upper bound on the definable cliquewidth of all k-derivations. To this end, we need to prove that the assumptions of the Simon Lemma are satisfied, that is, condition (1) is satisfied when either $n = 2$ or all the abstractions of all k-derivations in the product are equal to some idempotent in the semigroup of abstractions.

We now set off to implement this plan formally. For the rest of the paper we fix $k \in \mathbb{N}$. Our goal is to show that graphs of linear cliquewidth at most k have bounded definable cliquewidth.

Derivations. We first introduce k-derivations and their semigroup.

Definition 4. A k-derivation σ is a triple (G, λ, ϕ), where

- G is a k-colored graph, called the *underlying graph of* σ;
- $\lambda\colon V(G) \to 2^{[k]}$ is a function that assigns to each vertex u its *profile* $\lambda(u) \subseteq [k]$; and
- $\phi\colon [k] \to [k]$ is a function called the *recoloring*.

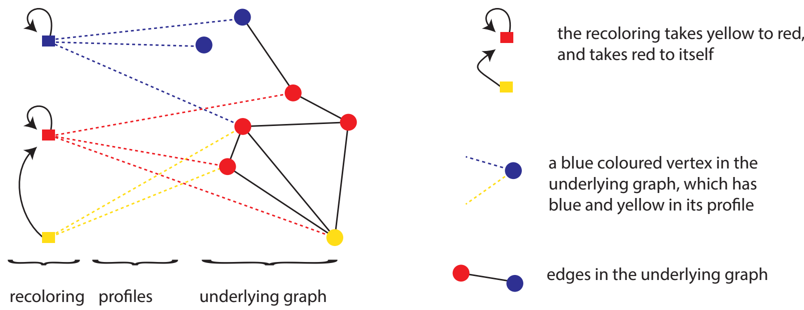

Figure 2: A k-derivation for $k = 3$, with the numbers $\{1, 2, 3\}$ being represented as colors {red, blue, yellow}. The red boxes indicate colors as used by the profiles and recoloring, and the circles indicate vertices of the underlying graph.

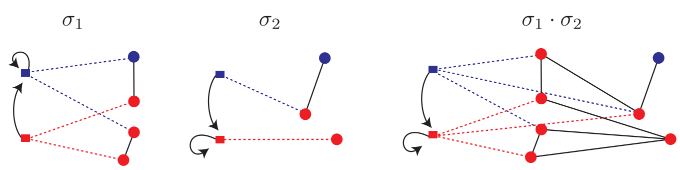

Figure 3: Composition of derivations.

Intuitively, if we treat a k-derivation $\sigma = (G, \lambda, \phi)$ as a subword of instructions in a linear clique decomposition, then G is the subgraph induced by vertices introduced by these instructions and ϕ is the composition of all recolorings applied. The profile λ has the following meaning: supposing there were some instructions preceding the k-derivation in question, it assigns each vertex u of G a subset $\lambda(u)$ of colors such that among vertices introduced by these preceding instructions, u is adjacent exactly to vertices with colors from $\lambda(u)$. See Figure 2 for an example.

By the definable cliquewidth of a k-derivation we mean the definable cliquewidth of its underlying graph, with the colors ignored. For a k-derivation $\sigma = (G, \lambda, \phi)$ and $c = (i, X) \in [k] \times 2^{[k]}$, the set of all vertices with color i and profile X is be called the *c-cell*, and denoted by $\sigma[c]$. For brevity, we denote $\mathcal{C}_k = [k] \times 2^{[k]}$ and interpret it as the index set of cells in k-derivations. By abuse of notation, we use the term *cell* also for the elements of $\mathcal{C}_k$.

We now describe the semigroup structure of k-derivations. We define the *composition* $\sigma_1 \cdot \sigma_2$ of two k-derivations $\sigma_1 = (G_1, \lambda_1, \phi_1)$ and $\sigma_2 = (G_2, \lambda_2, \phi_2)$ as follows; see Figure 3 for an illustration. The underlying graph of the composition is constructed by taking the disjoint union of $\phi_2(G_1)$ and G_2, where $\phi_2(G_1)$ denotes G_1 with the color of each vertex substituted with its image under ϕ_2, and adding an edge between a vertex $u \in G_1$ and a vertex $v \in G_2$ whenever the color of u in G_1 belongs to the profile $\lambda_2(v)$. The profile of a vertex u in the composition is equal to $\lambda_1(u)$ if u originates from G_1, and to $\phi_1^{-1}(\lambda_2(u))$ if u originates from G_2. Finally, the recoloring in the composition is the composition of recolorings, that is, $\phi_2 \circ \phi_1$. It is straightforward to see that composition is associative, and hence it turns the set of k-derivations into a semigroup.

Define an *atomic* k-derivation to be one where the underlying graph has at most one vertex. The number of different atomic k-derivations is finite and bounded only in terms of k, because the only freedom is the choice of the color and the profile of the unique vertex (if there is one), as well as the recoloring. Define S_k to be the subsemigroup of the semigroup of all k-derivations which is generated by the atomic k-derivations. By definition, S_k is finitely generated. The following claim is a straightforward reformulation of the definition of linear cliquewidth.

Proposition 4. *If a graph has linear cliquewidth at most k, then it is the underlying graph of some k-derivation $\sigma \in S_k$.*

Proof. Take any width-k linear clique decomposition of the graph G in question and turn it into a sequence of atomic k-derivations as follows. Every $\mathbf{Recolor}_\phi$ operation is replaced with an atomic k-derivation with empty underlying graph and recoloring ϕ, whereas every $\mathbf{AddVertex}_{i,X}$ operation is replaced by an atomic k-derivation with identity recoloring and underlying graph consisting of one vertex of color i and profile X. The composition of the obtained sequence yields a k-derivation $\sigma \in S_k$ whose underlying graph is G. □

Abstractions. Our goal is to apply the Simon Lemma to the finitely generated semigroup S_k, with μ being the definable cliquewidth of the underlying graph. To apply the Simon Lemma, we also need a homomorphism from S_k to some finite semigroup. This homomorphism is going to be abstraction, and we define it below.

To define the abstraction, we need one more auxiliary concept, namely the *flipping* a graph. For a graph G and vertex subsets $X, Y \subseteq V(G)$, the *flip between X and Y* is defined to be the following operation modifying G: for each $x \in X$, $y \in Y$, $x \neq y$, if there is an edge xy then remove it, and otherwise add it. In other words, flipping between X and Y means reversing the adjacency relation in all pairs of different elements from $X \times Y$. Note that in the flip operation, the sets X and Y need not be disjoint. Suppose that σ is a k-derivation. Recall that $\mathcal{C}_k$ represents the names of cells, i.e. each element of $\mathcal{C}_k$ is a pair (vertex color, profile). For a subset $Z \subseteq \binom{\mathcal{C}_k}{1,2}$ define the *Z-flip of σ* is a graph obtained from the underlying graph of G by performing the flip between $\sigma[c]$ and $\sigma[d]$ for each $\{c, d\} \in Z$. Note Z can contain singletons, i.e. we might have $c = d$.

Definition 5. For a k-derivation σ, its *abstraction*, denoted by $[\![\sigma]\!]$, is the triple (L, ρ, ϕ) consisting of the following information about σ:

- $L \subseteq \mathcal{C}_k$ is the set of cells that are non-empty in σ, called *essential*;
- $\rho \subseteq 2^{\binom{\mathcal{C}_k}{1,2}} \times \mathcal{C}_k \times \mathcal{C}_k \times 2^{\mathcal{C}_k}$ is the *connectivity registry*, which contains all tuples (Z, c, d, W) such that: in the Z-flip of σ there is a path that starts in a vertex of $\sigma[c]$, ends in a vertex of $\sigma[d]$, and whose all internal vertices belong to $\bigcup_{b \in W} \sigma[b]$;
- ϕ is the recoloring function of σ.

We briefly explain the idea behind the connectivity registry. In general, we would like to remember which pairs of cells can be connected by a path in the underlying graph of the derivation. However, in the proof we will sometimes be working not with a k-derivation, but with some Z-flip of it. Therefore, we want the abstraction to store the connectivity information after every possible flip. For technical reasons, we also remember the subset of cells that are traversed by the path.

Denote by T_k the set of all possible abstractions of k-derivations; note that T_k is a finite set whose size depends only on k, albeit it is doubly exponential in k. We leave it to the reader to prove that "having the same abstraction" is a congruence in the semigroup S_k, that is, an equivalence relation $\sim$ on S_k such that $s \sim s'$ and $t \sim t'$ imply $st \sim s't'$ for all $s, s', t, t' \in S_k$. It follows that we may endow T_k with a unique binary composition operation which makes it into a semigroup, and which makes the abstraction function a semigroup homomorphism from S_k to T_k.

Applying the Simon Lemma. We will apply the Simon Lemma for $S = S_k$, $T = T_k$, h being the abstraction, and μ being the definable cliquewidth of the underlying graph of a k-derivation (after forgetting the coloring). The conclusion of the Simon Lemma will say that μ has bounded range, i.e. there is a finite bound on the definable cliquewidth of the underlying graphs of derivations from S_k. Since these underlying graphs are the same as graphs of linear cliquewidth at most k by Proposition 4, this will mean that bounded linear cliquewidth implies bounded definable cliquewidth, thus proving Theorem 2.

To apply the Simon Lemma, we need to verify that assumption (1) is satisfied for some function $f\colon \mathbb{N} \to \mathbb{N}$. The treatment of cases when $n = 2$, and when all derivations have a common idempotent abstraction, is different, as encapsulated in the following two lemmas.

Lemma 5 (Binary Lemma)**.** *There is a function $f\colon \mathbb{N} \to \mathbb{N}$ such that*

$$\mathrm{dcw}(\sigma \cdot \tau) \leqslant f(\max(\mathrm{dcw}(\sigma), \mathrm{dcw}(\tau)))$$

for every $\sigma, \tau \in S_k$.

Lemma 6 (Idempotent Lemma)**.** *There is a function $f\colon \mathbb{N} \to \mathbb{N}$ such that*

$$\mathrm{dcw}(\sigma_1 \cdots \sigma_n) \leqslant f(\max_{i \in [n]} \mathrm{dcw}(\sigma_i))$$

for every $\sigma_1, \ldots, \sigma_n$ which have the same abstraction, and this abstraction is idempotent.

Condition (1) of Simon Lemma then follows by taking f to be the maximum of the functions given by the Binary and the Idempotent Lemma. Thus, we are left with proving these two results. The proof of the Binary Lemma is actually quite easy and we could present it right away, but it will be more convenient to use technical tools developed in the proof of the Idempotent Lemma, so we postpone it to Section 5.1.

5 Proof of the Idempotent Lemma

In this section we prove the Idempotent Lemma assuming a technical result called the Definable Order Lemma, which we will explain in a moment. Let us consider a sequence $\sigma_1, \ldots, \sigma_n$ of k-derivations such that for some abstraction e that is idempotent in T_k, we have $e = [\![\sigma_1]\!] = \ldots = [\![\sigma_n]\!]$. Let $\sigma = \sigma_1 \cdots \sigma_n$, and let G be the underlying graph of σ. Moreover, for $i \in [n]$ by G_i we denote the underlying graph of σ_i, and we call it also the *i-th block*.

Let $\preceq$ be the linear quasi-order (i.e. a total, transitive and reflexive relation) defined on the vertex set of G as follows: $u \preceq v$ holds if and only if u belongs to the i-th block and v belongs to the j-th block for some $i \leqslant j$. Similarly, let $\equiv$ be the equivalence relation on the vertex set of G defined as belonging to the same block; that is, $u \equiv v$ iff $u \preceq v$ and $v \preceq u$. The relations $\preceq$ and $\equiv$ will be called the *block order* and the *block equivalence*, respectively. Our general idea is to show that the block order, and hence also the block equivalence, can be interpreted using a bounded size (nondeterministic) MSO formula, i.e. that it has bounded (in terms of k) interpretation complexity as defined below.

Definition 6 (Interpretation complexity)**.** Suppose that $\mathfrak{A}$ is a relational structure, and let R be a relation on its universe, say of arity n. Define the *interpretation complexity* of R inside $\mathfrak{A}$ to be the smallest m such that there exist subsets $X_1, \ldots, X_m$ of the universe in $\mathfrak{A}$ and an MSO formula $\varphi(x_1, \ldots, x_n, X_1, \ldots, X_m)$ of quantifier rank at most m over the vocabulary of $\mathfrak{A}$ such that

$$(x_1, \ldots, x_n) \in R \quad \text{iff} \quad \varphi(x_1, \ldots, x_n, X_1, \ldots, X_m) \qquad \text{for all } x_1, \ldots, x_n \text{ in } \mathfrak{A}.$$

If the interpretation complexity of the block order was bounded by a function of k, then we would construct a clique decomposition of G as follows: first construct clique decompositions of all blocks, and then combine them sequentially along the block order. Unfortunately, in general we cannot hope for such a bound. To see this, consider the example where G consists of, say, two disjoint paths of length n each, plus an independent set of size n. In this example, each σ_i introduces the i-th vertex of each of the two paths and one vertex in the independent set. It is not difficult to see that in this example the interpretation complexity of the block order grows with the number of blocks. However, we can define the block order on each connected component (i.e. each of the two paths, and each vertex of the independent set) separately, and a clique decomposition of the whole graph can be obtained by putting a Join over decompositions of components. Thus, the obtained decomposition will have a different shape than the input linear decomposition corresponding to the product $\sigma_1 \cdots \sigma_n$. The following statement, which is our main technical result towards the proof of the Idempotent Lemma, explains how this plan can be implemented in general.

Lemma 7 (Definable Order Lemma)**.** *Let $\sigma_1, \ldots, \sigma_n$ be k-derivations as in the assumption of the Idempotent Lemma. There exists a set $Z \subseteq \binom{C_k}{1,2}$ such that if $\sim$ is the relation of being in the same connected component in the Z-flip of $\sigma_1 \cdots \sigma_n$, then the relation $\sim \cap \preceq$ has interpretation complexity over G bounded by a function of k.*

For now we postpone the proof of the Definable Order Lemma; it will be presented in Section 6. In the rest of this section we show how to use this result to prove the Idempotent Lemma. Along the way we will develop a relevant toolbox for handling decomposers and definable cliquewidth, and at some point the Binary Lemma will easily follow from the already gathered observations.

5.1 Toolbox for decomposers

Filtering and Transfering Structure. In this section, we establish two simple lemmas which crucially rely on decomposers being *origin-preserving*, that is, satisfying condition (b) of Definition 1. In the following, let $\{E\}$ be the vocabulary of graphs, where E is the binary adjacency relation, and let Δ_k be the vocabulary of width-k clique decompositions. We assume that $E \notin \Delta_k$.

The first of our lemmas allows us to make sure that a nondeterministic transduction that is supposed to be a decomposer is correct by filtering out outputs that are not clique decompositions of the input graph. Recall that for every input-output pair (G, t) of a decomposer $\mathcal{D}$ the structure t is a clique decomposition of G and the origin mapping of $\mathcal{D}$ induces an isomorphism from G_t to G. In general, for an MSO-transduction $\mathcal{I}$ from $\{E\}$ to Δ_k, we say that $\mathcal{I}$ *decomposes* a graph G if there is a some output t of $\mathcal{I}$ on input G such that t is a clique decomposition of G and the origin mapping of $\mathcal{D}$ induces an isomorphism from G_t to G.

Lemma 8 (Filter Lemma)**.** *Let $\mathcal{I}$ be an* MSO*-transduction with input vocabulary $\{E\}$ and output vocabulary Δ_k. Then there is a width-k decomposer $\mathcal{D}$ that decomposes the same graphs as $\mathcal{I}$.*

Proof. Let R be a binary relation symbol not contained in $\{E\} \cup \Delta_k$. By copying the input, we can modify $\mathcal{I}$ to obtain a transduction $\mathcal{I}'$ from $\{E\}$ to $\{E, R\} \cup \Delta_k$ that for every input-output pair (G, t) of $\mathcal{I}$ has an output pair (G, A), where A is the structure obtained from the disjoint union of G and t by adding a binary relation R that connects all elements with the same origin. Then, using filtering, we can discard those pairs (G, A) where the underlying t is not clique decomposition of G for which the origin mapping induces an isomorphism from G_t to G. (Note that we cannot check whether G_t is isomorphic to G, but we can check whether the origin mapping is an isomorphism. This is the main reason why we require decomposers to be origin preserving.) Finally, we can restrict the universe of A to retrieve the original t. $\square$

The Filter Lemma implies that to prove Theorem 2, it suffices to prove that for every k there is an ℓ and an MSO-transduction from $\{E\}$ to Δ_ℓ that decomposes all graphs of linear clique width k.

The second consequence of the decomposers being order-preserving is that we can transfer additional structure on the input graphs to the graphs constructed by the output decompositions. For vocabularies $\Sigma \subseteq \Sigma'$, the Σ-*restriction* of Σ'-structure A' is a Σ-structure A that has the same universe as A' and coincides with A' on all relations in Σ. Conversely, a Σ'-*expansion* of a Σ-structure A is a Σ'-structure A' such that A is the Σ-restriction of A'.

Lemma 9 (Transfer Lemma)**.** *Let $\mathcal{D}$ be a width-k decomposer and let Σ be a vocabulary disjoint from $\{E\} \cup \Delta_k$. Then there is an* MSO*-transduction $\mathcal{D}^\star$ from $\{E\} \cup \Sigma$ to $\Delta_k \cup \{E\} \cup \Sigma$ such that for every input-output pair (G, t) of $\mathcal{D}$ and every $\{E\} \cup \Sigma$-expansion $G^\star$ of G there is a unique $\Delta_k \cup \{E\} \cup \Sigma$-expansion $t^\star$ of t such that $(G^\star, t^\star)$ is an input output pair of $\mathcal{D}^\star$ and the origin mapping restricted to the leaves of t is an isomorphism from the induced substructure of $t^\star$ to $G^\star$.*

Proof. We apply exactly the same sequence of atomic transductions but we keep all the additional relations from Σ always intact. The claim follows by the assumption that $\mathcal{D}$ is origin-preserving (condition (b) of Definition 1). □

We will use this lemma to transfer colors of the input graph of a decomposer to the output.

Enforcing a fixed partition. Given a clique decomposition of a graph, say of width k, and a partition of the vertex set into p subsets, which may be non-related to the decomposition, one can adjust the decomposition at the cost of using $k \cdot p$ colors instead of k so that the final color partition of the decomposition matches the given one. Informally, this can be done by just enriching each original label with information to which subset of the final partition a vertex belongs. The following general-usage lemma formalizes this, and shows that the transformation may be performed by means of an MSO transduction

Lemma 10 (Color Enforcement Lemma)**.** *For every $k, p \in \mathbb{N}$, there exists a deterministic* MSO *transduction $\mathcal{E}_{k,p}$ with the following properties. The input vocabulary is the vocabulary of clique decompositions of width k with leaves colored using p unary predicates. The output vocabulary is the vocabulary of clique decomposition of width $k \cdot p$. Finally, on an input decomposition t with leaves partitioned into $(V_1, \ldots, V_p)$ using unary predicates, the output of $\mathcal{E}_{k,p}$ is a decomposition t' of the same graph, where in the result of t' the color of each vertex from V_i is equal to i, for all $i \in [p]$.*

Proof. The decomposition t is first adjusted to a decomposition t'' of width $k \cdot p$ with the following property: in the result of t'', the final color of every vertex is a pair consisting of its color in the result of t and the index i such that the leaf corresponding to the vertex belongs to V_i. This correction can be made by preserving the shape of t intact, and performing a straightforward modification to the labels of nodes. For instance, for a Join node, whenever the original label in t requested adding edges between colors c and d, the new label in t'' requests adding edges between colors (c, i) and (d, j) for all $i, j \in [p]$. Finally, we obtain t' by adding a recoloring step on top of t'' that removes the first coordinate of every color. □

Using the Color Enforcement Lemma, we can give a proof of the Binary Lemma.

Proof of the Binary Lemma. Let $m = \max(\mathrm{dcw}(\sigma), \mathrm{dcw}(\tau))$, and let $\mathcal{D}_\sigma$ and $\mathcal{D}_\tau$ be decomposers of size at most m such that $\mathcal{D}_\sigma$ produces at least one output on the underlying graph of σ, and similarly for $\mathcal{D}_\tau$. Using coloring, we first guess the partition of the vertex set into vertices that belong to the underlying graphs of σ and τ. Next, we guess the color partition in the underlying graph of σ. Finally, for the underlying graph of τ, we guess the partition of its vertices according to profiles in τ. Note that the validity of this guess, or more precisely the fact that the adjacency between the σ-part and the τ-part depends only on the (color,profile) pair of respective vertices, can be checked using a filtering step.

We now apply $\mathcal{D}_\sigma$ to the σ-part of the graph, yielding a clique decomposition of the underlying graph of σ of width at most m. By applying the transduction $\mathcal{E}_{m,k}$ given by the Color Enforcement Lemma to this decomposition, by the Transfer Lemma we can assume that the result of the obtained decomposition t_σ has the color partition equal to the color partition of σ. Similarly, by applying $\mathcal{D}_\tau$ followed by $\mathcal{E}_{m,2^k}$ for the profile partition, we turn the τ-part of the graph into its clique decomposition t_τ whose result has the color partition equal to the profile partition in τ. Since the adjacency between the σ-part and the τ-part depends only on the (color,profile) pair of respective vertices, it now suffices to add one binary Join node, with the roots of t_σ and t_τ as children, where we request adding edges between appropriate pairs of vertices, selected by color on the σ-side and profile on the τ-side. □

Combining many decomposers. In the setting of the Idempotent Lemma, the graph consists of multiple pieces, each having small definable cliquewidth. Thus, we may think that for each piece we have already constructed a decomposer (w.l.o.g. the same one, as we can take the union of the input decomposers), and now we need to put all these decomposers together. In particular, we will need to apply the decomposers "in parallel" to all the considered pieces. The following Parallel Application Lemma formalizes this idea.

For a vocabulary Σ and a sequence $\mathfrak{A}_1, \dots, \mathfrak{A}_n$ of Σ-structures, by their *disjoint union* $\biguplus_{i=1}^n \mathfrak{A}_i$ we denote a structure over vocabulary $\Sigma \cup \{\sim\}$, where $\sim$ is a binary symbol, defined as follows:

- the universe of $\biguplus_{i=1}^n \mathfrak{A}_i$ is the disjoint union of the universes of $\mathfrak{A}_i$ for $i \in [n]$;
- for each symbol $R \in \Sigma$, the interpretation of R in $\biguplus_{i=1}^n \mathfrak{A}_i$ is the union of its interpretations in structures $\mathfrak{A}_i$ for $i \in [n]$;
- $\sim$ is interpreted as an equivalence relation on the universe of $\biguplus_{i=1}^n \mathfrak{A}_i$ that selects pairs of elements originating in the same structure $\mathfrak{A}_i$.

Lemma 11 (Parallel Application Lemma)**.** *Let $\mathcal{I}$ be an* MSO *transduction with input vocabulary Σ and output vocabulary Γ. Then there is an* MSO *transduction $\widehat{\mathcal{I}}$ with input vocabulary $\Sigma \cup \{\sim\}$, output vocabulary $\Gamma \cup \{\sim\}$, and the following semantics: for every sequence $(\mathfrak{A}_1, \mathfrak{B}_1), \dots, (\mathfrak{A}_n, \mathfrak{B}_n)$ of pairs of Σ- and Γ-structures, we have $(\biguplus_{i=1}^n \mathfrak{A}_i, \biguplus_{i=1}^n \mathfrak{B}_i) \in \widehat{\mathcal{I}}$ if and only if $(\mathfrak{A}_i, \mathfrak{B}_i) \in \mathcal{I}$ for all $i \in [n]$.*

Proof. Observe that it suffices to verify the lemma on atomic transductions. For copying and coloring the claim is trivial: we can take the same operation. For universe restriction, say using MSO predicate $\psi(x)$, we use universe restriction using predicate $\psi'(x)$ that is constructed from $\psi(x)$ by relativizing it to the $\sim$-equivalence class of x, that is, adding a guard to every quantifier that restricts its range to (sets of) elements $\sim$-equivalent to x. For interpretation, we similarly modify each MSO formula $\varphi_R(x_1, \dots, x_r)$ by additionally requiring that the elements $x_1, \dots, x_r$ are pairwise $\sim$-equivalent, and relativizing the formula to the $\sim$-equivalence class of $x_1, \dots, x_r$. Finally, for filtering, say using an MSO sentence ψ, we use filtering using an MSO sentence saying that for every equivalence class of $\sim$, the formula ψ relativized to this equivalence class holds. □

We now proceed to the final tool for decomposers: the Combiner Lemma. In principle, it formalizes the idea that in the setting of the Idempotent Lemma, having defined the block order (roughly, using the Definable Order Lemma), we may construct clique decompositions for individual pieces (derivations), obtained by applying small decomposers in parallel, into a clique decomposition of the whole graph.

Define an *order-using decomposer* to be an MSO transduction which inputs a graph G together with a linear quasi-order on its vertices and which outputs clique decompositions of the input graph. On a given input, an order-using decomposer might produce several outputs, possibly zero.

Lemma 12 (Combiner Lemma). *For every $m \in \mathbb{N}$ there is an order-using decomposer $\mathcal{D}$ with the following property. Let $\tau_1, \ldots, \tau_n$ be k-derivations whose underlying graphs have definable cliquewidth at most m. Let G be the underlying graph of $\tau_1 \cdots \tau_n$ and $\preceq$ be the block order arising from decomposition $\tau_1 \cdots \tau_n$. Then $\mathcal{D}$ produces at least one output on $(G, \preceq)$.*

Proof. In the following, we describe the order-using decomposer $\mathcal{D}$. First, using coloring guess the partition of the vertex set into sets $\{U_c \colon c \in \mathcal{C}_k\}$ such that $U_c = \bigcup_{i \in [n]} \tau_i[c]$. Then the cell $\tau_i[c]$ may be recovered as the intersection of U_c with the underlying graph of τ_i, which in turn can be identified as a single equivalence class of the block equivalence.

By assumption, for each $i \in [n]$ there is a decomposer $\mathcal{D}_i$ of size at most m that applied to the underlying graph of τ_i produces at least one output. By the Color Enforcement Lemma applied to the cell partition $\{\tau_i[c] \colon c \in \mathcal{C}_k\}$ and the Transfer Lemma, we may assume that the result of each $\mathcal{D}_i$ has color partition coinciding with this cell partition. After this operation, the sizes of all decomposers $\mathcal{D}_i$ are still bounded by some m' depending only on m and k.

Let now $\mathcal{J}$ be the union of all decomposers of size at most m'. As we argued before, $\mathcal{J}$ is a decomposer of size bounded by a function of m' such that $\mathcal{J}$ applied to the underlying graph of any τ_i has at least one output. Moreover, this output is a clique decomposition t_i of the underlying graph of τ_i such that in the result of t_i, the color partition is equal to the cell partition in τ_i.

Let $\equiv$ be the block equivalence in $\tau_1 \cdots \tau_n$, interpreted from the block order $\preceq$. Apply Parallel Application Lemma to $\mathcal{J}$ and $\equiv$, yielding an MSO transduction $\widehat{\mathcal{J}}$ that, when applied to the whole structure, turns the underlying graph of each τ_i into its clique decomposition t_i as above. Since each t_i originates in vertices of the underlying graph of τ_i, on decompositions t_i we still have the order $\preceq$ present in the structure.

It now remains to combine decompositions t_i sequentially. We do it as follows. For every $i \in [n-1]$ we create two nodes a_i, b_i, for instance by copying the roots of decompositions t_i for $i \in [n-1]$ two times. Then we connect these nodes into a path, called the *spine*, so that each a_i is a child of b_i, and each b_i is a child of a_{i+1} (except $i = n-1$). It is easy to do it in a single interpretation step, as the order $\preceq$ is present in the structure, so for every decomposition t_i we can interpret the next decomposition t_{i+1}. Further, we make the root of t_{i+1} a child of a_i for each $i \in [n-1]$, and moreover the root of t_1 becomes a child of a_1; again, this can be done in one interpretation step. This establishes the shape of the final decomposition, where b_{n-1} is the root.

For the labels of nodes, each a_i is labeled by a Join operation, and each b_i is labeled by a Recolor operation. In the colored graphs computed along the spine, the consecutive colors assigned to every vertex, say originating from τ_i, are equal to the color that would be assigned to this vertex in derivations τ_i, $\tau_i \cdot \tau_{i+1}$, $\tau_i \cdot \tau_{i+1} \cdot \tau_{i+2}$, and so on.

The Join operation at node a_i requests adding edges between every vertex u coming from the child on the spine (b_{i-1}, or t_1 if $i = 1$), and every vertex v coming from decomposition t_{i+1}, whenever the color of u belongs to the profile of v. Recall here that the color partition in the result of t_{i+1} matches the cell partition in τ_{i+1}, so the profiles in τ_{i+1} are encoded in the colors in the result of t_{i+1}. Also, we may assume that the colors originating from the subtree below b_{i-1} are pairwise different than colors originating from t_{i+1}, so in the Join at a_i no two colors are merged.

The Recoloring operation at node b_i removes the information about profiles from the colors of vertices originating from t_{i+1}, and adjusts colors for vertices coming from below the spine (i.e., originating in decompositions $t_{i'}$ for $i' \leqslant i$) according to the recoloring applied in τ_{i+1}. Observe that for each node b_i we may guess this recoloring nondeterministically, by guessing, for every function $\phi\colon [k] \to [k]$, a unary predicate that selects nodes b_i where recoloring ϕ should be used. By appealing to the Filter Lemma, we can always check that in the end that we have indeed obtained a clique decomposition of the input graph. Hence, even though some of the nondeterministic guesses may lead to constructing a clique decomposition whose result is different from the input graph, these guesses will be filtered out at the end. □

5.2 Definable cliquewidth under restriction of the universe

In the Idempotent Lemma we assume that each individual k-derivation σ_i has bounded definable cliquewidth, say by K, which means that we have a decomposer of size at most K that constructs a clique decomposition of the underlying graph G_i of σ_i. However, recall that the Definable Order Lemma does not provide the full block order (that could be fed to the Combiner Lemma), only its restriction to the connected components of some Z-flip of the graph. Therefore, graphs G_i are not directly available to the transduction; we are able to construct only their restrictions to the connected components of the said Z-flip.

It would be now convenient to claim that definable cliquewidth is closed under taking induced subgraphs, similarly as the standard cliquewidth is, so that the Combiner Lemma could be applied to each connected component of the Z-flip separately. This, however, is not immediate, as the decomposer for the induced subgraph would need to work only on this induced subgraph. In fact, we do not know whether this statement is true at all, but we can prove a weaker variant that turns out to be sufficient for our needs.

Let G be an undirected graph and let (V_0, V_1) be a partition of the vertex set of G into two sets. We define the *rank* of the partition (V_0, V_1) as the number of equivalence classes in the following equivalence relation on vertices. Two vertices v, w are considered equivalent if they both belong to the same V_i for some $i \in \{0, 1\}$, and the sets of neighbors of v and w within V_{1-i} are the same. Define the *rank* of an induced subgraph H of G, denoted $\mathrm{rank}(G, H)$, to be the rank of the partition $(V(H), V(G) - V(H))$ in G. We prove that the definable cliquewidth of an induced subgraph is bounded by a function of the definable cliquewidth of the larger graph, provided the rank of the induced subgraph within the larger graph is bounded.

Lemma 13. *There is a function $f\colon \mathbb{N} \to \mathbb{N}$ such that for every graph G and its induced subgraph H,*

$$\mathrm{dcw}(H) \leqslant f(\max(\mathrm{dcw}(G), \mathrm{rank}(G, H))).$$

Before we proceed to the proof of Lemma 13, we give the following lemma about replacing a part of a graph subject to preserving the satisfaction of an MSO sentence. Its proof is completely standard, but we give it for completeness.

Lemma 14 (MSO Pumping Lemma)**.** *Let φ be an* MSO *sentence over graphs (i.e. the universe is the vertex set and the vocabulary consists of one binary edge predicate). Let G be a graph that satisfies φ. For every induced subgraph H of G, there is some G' such that:*

(1) G' satisfies φ;

(2) H is an induced subgraph of G'; and

(3) the number of vertices in $G' - H$ is bounded by a constant depending only on φ and $\mathrm{rank}(G, H)$.

Proof. Define $\mathcal{N}$ to be the family of neighborhoods in H of vertices from $G - H$, i.e.

$$\mathcal{N} \stackrel{\text{def}}{=} \{\{v : v \text{ is a vertex in } H \text{ adjacent to } w\} : w \text{ is a vertex in } G - H\}.$$

Observe that if the rank of H in G is r, then $|\mathcal{N}| \leqslant r$. This is because every neighborhood in $\mathcal{N}$ is the union of a collection of equivalence classes of the equivalence relation considered in the definition of $\operatorname{rank}(G, H)$.

Define F to be the $\mathcal{N}$-colored graph obtained from $G - H$ by coloring each vertex by its neighborhood in H. We treat F as a relational structure with one binary edge predicate and $|\mathcal{N}|$ unary predicates, one for each color. Let q be the quantifier rank of φ. Choose F' to be the smallest $\mathcal{N}$-colored graph which satisfies the same MSO sentences of quantifier rank q as F. The size of F' is bounded by a constant depending on $|\mathcal{N}|$ and q, which in turn depend only on $\operatorname{rank}(G, H)$ and φ. Define G' to be the following graph: we take the disjoint union of F' and H, forget the coloring in F', and then for each vertex v in F', we connect it to those vertices in H which were in its (now forgotten) color. Using an Ehrenfeucht-Fraisse argument, it is straightforward to show that G' and G satisfy the same MSO sentences of quantifier rank at most q. In particular, G' satisfies φ. □

We would like to remark that the number of vertices in $G' - H$ in the above lemma can be computed given: φ, the rank $\operatorname{rank}(G, h)$ and the cliquewidth of H. (Note that the lemma asserts a stronger property, namely that the number of vertices in $G' - H$ is bounded only by φ and the rank $\operatorname{rank}(G, h)$, and there is no dependency on the cliquewidth of H. Nevertheless, for computability we also use the cliquewidth of H, and this additional dependency is not an issue for our intended application in the proof of Lemma 13, where we have an upper bound on the cliquewidth of H anyway.) Indeed, given numbers $r, k, k' \in \{1, 2, \ldots\}$ and an MSO formula φ, consider the following statement:

> ($\star$) For every graph G of cliquewidth $\leqslant k$, and every induced subgraph H of G with
>
> $$\operatorname{rank}(G, H) < r$$
>
> there exists a graph G' which satisfies items (1) and (2) from Lemma 14 and such that the number of vertices in $G' - H$ is at most k'.

Lemma 14 implies that for every φ, r, k there exists some k' which makes ($\star$) true. By the following lemma, to compute (given φ, r, k) the smallest k' which makes ($\star$) we can go through each candidate for k' and check if ($\star$) is satisfied. Summing up, the bound on the size of $G' - H$ in item (3) of Lemma 14 is computable, assuming additionally that we know the cliquewidth of H (which is at most the cliquewidth of G).

Lemma 15. *For every φ, r, k, k' one can decide if ($\star$) holds.*

Proof. For an MSO formula φ and $k' \in \{1, 2, \ldots\}$, consider the following property of graphs:

$$\{H\colon \text{there is some } G' \models \varphi \text{ such that } \begin{cases} G' \text{ satisfies } \varphi, \\ H \text{ is an induced subgraph of } G', \text{ and} \\ \text{the number of vertices in } G' - H \text{ is at most } k' \end{cases} \} \qquad (2)$$

It is straightforward to see that the above property is also MSO definable. Indeed, one may existentially guess the isomorphism class of $G' - H$ and its adjacencies to H, because $G' - H$ has bounded size, and then rewrite φ by simulating quantification over the additional vertices in $G' - H$ within syntax. The property ($\star$) asks if the formula expressing (2) is true for all graphs of cliquewidth $\leqslant k$. Checking whether an MSO formula is true for all graphs of cliquewidth $\leqslant k$ is a decidable problem, see [5, Section 7.5]. □

Proof of Lemma 13. Let $\ell = \max(\mathrm{dcw}(G), \mathrm{rank}(G, H))$. Our goal is to show that the definable cliquewidth of H is bounded by a function of ℓ. By definition of definable cliquewidth, there is a decomposer $\mathcal{D}$ which produces at least one output on G. From the Backwards Translation Theorem applied to the sentence "true" it follows that the domain of the decomposer $\mathcal{D}$, i.e. those graphs where it produces at least one output, is MSO definable, say by a sentence φ. Apply the MSO Pumping Lemma to φ and the graph $H \subseteq G$, yielding a graph $G' \supseteq H$ in the domain of $\mathcal{D}$ such that H is an induced subgraph of G' and the number of vertices in $G' - H$ is bounded by a constant m depending only on ℓ. By the discussion after the MSO Pumping Lemma, the constant m can be effectively computed.

Since G' is in the domain of $\mathcal{D}$, we have that $\mathcal{D}$ applied on G' produces at least one output, say t. By the definition of a decomposer, t is a clique decomposition of G'. Then, a clique decomposition of H can be obtained from G' by removing all leaves of t corresponding to the vertices of $G' - H$, and performing straightforward cleaning operations.

We now construct a decomposer for H as follows. First, we copy any vertex of the graph m times, and using coloring and interpretation we (nondeterministically) turn H into G'. Then, we apply $\mathcal{D}$ to G', yielding some clique decomposition t of G'. By the Transfer Lemma, we can assume that the original relations are preserved on the leaves of t. Therefore, we can now remove the leaves of t corresponding to vertices of $G' - H$ and perform the clean-up operations; it is straightforward to see that this can be done by means of an MSO transduction. □

As argued, computability of the bound in item (3) of Lemma 14 entails computability of function f provided by Lemma 14.

5.3 Completing the proof

With all the tools prepared, we complete the proof of the Idempotent Lemma.

Proof of the Idempotent Lemma. Let $\sigma_1, \ldots, \sigma_n$ be k-derivations as in the Idempotent Lemma, i.e., with the same idempotent abstraction e. Let

$$K = \max_{i \in [n]} \mathrm{dcw}(\sigma_i).$$

Denote $\sigma = \sigma_1 \cdots \sigma_n$. Let G be the underlying graph of σ, let G_i be the underlying graph of σ_i for each $i \in [n]$, and let $\preceq$ be the block order in G compliant with the decomposition $\sigma_1 \cdots \sigma_n$. It suffices to describe a decomposer of size bounded in terms of k and K that constructs a clique decomposition of G.

First, using coloring we enrich the structure with unary predicates that encode the partition of the vertex set of G into cells $\sigma[c]$, for $c \in \mathcal{C}_k$. Then apply the Definable Order Lemma to $\sigma_1, \ldots, \sigma_n$, yielding Z and $\sim$. Note that Z is chosen among $2^{|\binom{\mathcal{C}_k}{1,2}|}$ options, so we can nondeterministically guess Z using, say, some coloring. Having Z fixed, the equivalence relation $\sim$ (being in the same connected component of the Z-flip of σ) can be added to the structure using interpretation. By the Definable Order Lemma, we can add also the relation $\preceq \cap \sim$ to the structure, as this increases the size of the transduction only by a function of k and K.

The next claim says that restricting blocks to equivalence classes of $\sim$ yields graphs of bounded definable cliquewidth. Here, we will crucially use the results from the last section on how definable cliquewidth behaves under restricting the vertex set.

Claim 2. There is $m \in \mathbb{N}$ depending only k and K such that for every equivalence class F of $\sim$ and every $i \in [n]$, the subgraph induced in G_i by vertices contained in F has definable cliquewidth at most m.

Proof. Take any $i \in [n]$ and any connected component F in the Z-flip of σ. Let $H_i = G_i[W \cap V(G_i)]$, that is, H_i is the subgraph induced in G_i by vertices contained in F. Observe that the rank of H_i inside G_i is bounded by $2k \cdot 2^k$, i.e., twice the number of cells. Indeed, since Z-flip changes only the adjacency between whole cells, and after performing the Z-flip there are no edges between the vertices of H_i and the rest of vertices of G_i, for every cell $c \in \mathcal{C}_k$ we have that both $(\sigma[c] \cap V(G_i)) \cap F$ and $(\sigma[c] \cap V(G_i)) - F$ are contained in the same equivalence class of the equivalence relation considered in the definition of the rank of H_i inside G_i. The claim now follows directly from Lemma 13. $\square$

Now apply the Combiner Lemma to the parameter m given by Claim 2, yielding an order-using decomposer $\mathcal{D}$ satisfying the following.

Claim 3. For each equivalence class F of $\sim$, the order-using decomposer $\mathcal{D}$ produces at least one output on $(G, \preceq)$ restricted to the vertices of F.

Recall that the cell partition $\{\sigma[c] \colon c \in \mathcal{C}_k\}$ has been guessed and added to the structure via unary predicates. By the Transfer Lemma, we may assume that this cell partition is preserved on the leaves of the output t_F of applying $\mathcal{D}$ to an equivalence class F of $\sim$. Hence, by appending an appropriate transduction given by the Color Enforcement Lemma for the cell partition, we may assume without loss of generality that the color partition in the result of t_F is equal to the restriction of the cell partition to F.

Let $\widehat{\mathcal{D}}$ be the MSO transduction given by the Parallel Application Lemma for $\mathcal{D}$. Using $\sim$ as the input equivalence relation for $\widehat{\mathcal{D}}$, we infer that $\widehat{\mathcal{D}}$ applied to the current structure turns every equivalence class F of $\sim$ into a clique decomposition t_F of the subgraph induced by F in G. Moreover, the color partition in the result of t_F is equal to the restriction of the cell partition to F.

It now suffices to add a new root node and attach all the root nodes of decompositions t_F as its children. This new root node is labeled by a Join operation, where we request that for all $c, d \in \mathcal{C}_k$ for which $\{c, d\} \in Z$, an edge should be added between every pair of vertices u, v such that u belongs to cell c, v belongs to cell d, and u, v originate from different decompositions t^F. Note that this is possible since in the result of each decomposition t^F, the color partition matches the appropriate restriction of the cell partition. Since the connected components of the Z-flip of G are pairwise non-adjacent in this Z-flip by definition, it is clear that the structure constructed in this manner is a clique decomposition of G. $\square$

6 Proof of the Definable Order Lemma

In this section we present the proof of the Definable Order Lemma. Let $\sigma_1, \ldots, \sigma_n$ be k-derivations which have the same idempotent abstraction. Recall that our goal is to prove that there is a set $Z \subseteq \binom{\mathcal{C}_k}{1,2}$ such that in the Z-flip of $\sigma_1 \cdots \sigma_n$ the relation $\sim \cap \preceq$ has interpretation complexity over G bounded by a function of k.

We first introduce some definitions. For a cell $c \in \mathcal{C}_k$, by U_c we denote the set of vertices of G comprising all vertices from c-cells in respective derivations σ_s. In other words,

$$U_c \stackrel{\text{def}}{=} \bigcup_{s \in [n]} \sigma_s[c].$$

Let us stress that U_c may be different from $\sigma[c]$, that is, the set of vertices from the c-cell in the overall derivation σ. This is because the profile and the color of a vertex in σ may differ from its profile and color in respective σ_s.

By $\preceq_c$ and $\equiv_c$ we denote the restriction of $\preceq$ and $\equiv$ to U_c, respectively. Moreover, for cells $c, d \in \mathcal{C}_k$, by $\preceq_{c,d}$ and $\equiv_{c,d}$ we denote the restriction of $\preceq$ and $\equiv$ to pairs from $U_c \times U_d$, respectively. Our first goal is to interpret the above relations for as large subset of cells (resp. pairs of cells) as possible.

For distinguishing neighboring blocks we introduce the following definitions. By *moduli* we mean the remainders modulo 7, that is, the elements of the set $\{0, 1, \ldots, 6\}$. All arithmetic on moduli is performed modulo 7. The *distance* between two moduli a, b, denoted $\mathrm{dist}(a, b)$, is the smaller of the numbers $(a - b) \bmod 7$ and $(b - a) \bmod 7$. Moduli a, b are *neighboring* if the distance between them is 1. Let the *modulus* of the i-th block be equal to $i \bmod 7$, for each $i \in [n]$. The modulus of a vertex u of G, denoted $\mathsf{mod}(u)$, is the modulus of the block to which it belongs. For a modulus $a \in \{0, 1, \ldots, 6\}$, by W_a we denote the set of all vertices of G that have modulus a.

We can now define a structure $\widehat{G}$ that is the enrichment of G with the following unary predicates:

- for each $c \in \mathcal{C}_k$, a unary predicate that selects the vertices of U_c; and
- for each modulus a, a unary predicate that selects the vertices of W_a.

The reader should think of these unary predicates as of some auxiliary information that is helpful in analyzing the graph for the purpose of interpreting the block order. These unary predicates will be (existentially guessed) monadic parameters X_i from the definition of interpretation complexity, hence we may simply assume that we work over structure $\widehat{G}$.

Idempotent recolorings. Observe that since the common abstraction $e = (L, \rho, \phi)$ is idempotent in T_k, all the input derivations σ_s have to have the same recoloring ϕ, and this recoloring has to be idempotent in the semigroup of functions from $[k]$ to $[k]$, endowed with the composition operation. This is because when composing two k-derivations, we compose their recolorings. It is easy to see that a function $\psi \colon [k] \to [k]$ is idempotent if and only if each i belonging to the image of ψ is a fixed point of ψ, that is, $\psi(i) = i$ for each $i \in \psi([k])$. Thus we have $\phi(i) = i$ for each $i \in \phi([k])$.

It is instructive to consider what this means in terms of recolorings, when each σ_j is treated as a sequence of instructions. Suppose that some vertex u belongs to the underlying graph of σ_s, for some $s \in [n]$, and its color in σ_s is i. This means that after applying all the operations in σ_s, the color of u is i, however this color may further change due to recolorings applied in each σ_t for $t > s$. For instance, the application of recoloring ϕ in σ_{s+1} changes the color of u from i to $\phi(i)$. However, since $\phi(\phi(i)) = \phi(i)$, the application of recoloring ϕ in every further σ_t, i.e. for $t \geqslant s + 2$, will not change the color of u, and this color will stay equal to $\phi(i)$ up to the end of the sequence.

Comparing pairs of cells. We now define *types* of pairs of cells as follows. A pair of (possibly equal) cells $(c, d) \in \mathcal{C}_k \times \mathcal{C}_k$, say $c = (i, X)$ and $d = (j, Y)$, is called:

- *negative* if $\phi(j) \notin X$ and $\phi(i) \notin Y$;
- *positive* if $\phi(j) \in X$ and $\phi(i) \in Y$; and
- *mixed* if $\phi(j) \in X$ and $\phi(i) \notin Y$, or $\phi(j) \notin X$ and $\phi(i) \in Y$.

Observe that for a cell c, the pair (c, c) is always either positive or negative, never mixed.

The following lemma shows that we can easily define the block order on each mixed pair of essential cells. Recall that a cell c is essential if it belongs to L, which means that the $\sigma_s[c]$ is nonempty for all $s \in [n]$. Note that for non-essential cells c, the set U_c is empty.

Lemma 16. *Suppose (c,d) is a mixed pair of essential cells. Then each of the following relations has interpretation complexity at most 2 over $\widehat{G}$: $\preceq_c$, $\preceq_d$, and $\preceq_{c,d}$.*

Proof. As (c,d) is mixed, we have $c \neq d$. Let $c = (i, X)$ and $d = (j, Y)$. By symmetry, without loss of generality assume that $\phi(j) \notin X$ and $\phi(i) \in Y$. We first claim that for pairs of vertices from $U_c \times U_d$ that originate in distant blocks, just the adjacency relation defines the block order.

Claim 4. Let $(u,v) \in U_c \times U_d$, and suppose $u \in G_s$ and $v \in G_t$ where $|s-t| > 1$. Then $u \preceq v$ if and only if u and v are adjacent in G.

Proof. Suppose first that $s < t-1$. Observe that u is colored with i in σ_s, then it is recolored to $\phi(i)$ when composing with σ_{s+1}, and it stays colored with $\phi(i)$ when composing with $\sigma_{s+2}, \ldots, \sigma_{t-1}$. As $\phi(i) \in Y$ and v resides in cell (j, Y), this means that when composing with σ_t, we add the edge between u and v.

Suppose next that $s > t+1$. Observe that v is colored with j in σ_t, then it is recolored to $\phi(j)$ when composing with σ_{t+1}, and it stays colored with $\phi(j)$ when composing with $\sigma_{t+2}, \ldots, \sigma_{s-1}$. As $\phi(j) \notin X$ and u resides in cell (i, X), this means that when composing with σ_s, we do not add the edge between u and v. ⌟

Next, elements of U_c from distant blocks can be compared using elements of U_d. For $u, v \in U_c$, a vertex $w \in U_d$ is called a *pivot for* (u,v) if the following conditions hold:

- $\mathrm{dist}(\mathsf{mod}(u), \mathsf{mod}(w)) > 1$, $\mathrm{dist}(\mathsf{mod}(v), \mathsf{mod}(w)) > 1$, and
- u and w are adjacent, whereas v and w are not adjacent.

First, we check that having a pivot implies the block order.

Claim 5. Let $u, v \in U_c$ and suppose there is a pivot for (u,v). Then $u \prec v$.

Proof. Let w be a pivot for (u,v), where $w \in G_p$ for some p. As $\mathrm{dist}(\mathsf{mod}(u), \mathsf{mod}(w)) > 1$ and $\mathrm{dist}(\mathsf{mod}(v), \mathsf{mod}(w)) > 1$, we in particular have that $|s-p| > 1$ and $|t-p| > 1$. By Claim 4, the adjacency between u and w implies that $u \preceq w$, and moreover $u \prec w$ since u and w have different moduli. Similarly, the non-adjacency between v and w implies that $w \prec v$. By transitivity we infer that $u \prec v$. ⌟

Next, we show that vertices of U_c that are in distant blocks always have a pivot.

Claim 6. Let $u, v \in U_c$, and suppose that $u \in G_s$ and $v \in G_t$ where $s < t-3$. Then there is a pivot for (u,v).

Proof. Note that among the 7 moduli, at most 3 are equal or neighboring $\mathsf{mod}(u)$, and at most 3 are equal or neighboring $\mathsf{mod}(v)$. Hence there exists a modulus r such that $\mathrm{dist}(\mathsf{mod}(u), r) > 1$ and $\mathrm{dist}(\mathsf{mod}(v), r) > 1$. It is easy to see that since $s < t-3$, there is a number p such that $s+1 < p < t-1$ and $p \bmod 7 = r$; in particular $|s-p| > 1$ and $|t-p| > 1$. Since d is an essential cell, there exists a vertex w in the block G_p such that $w \in U_d$. By Claim 4 we have that u and w are adjacent, while v and w are non-adjacent. As $\mathsf{mod}(w) = r$, indeed w is a pivot for (u,v). ⌟

We are ready to interpret $\preceq_c$.

Claim 7. The interpretation complexity of $\preceq_c$ over $\widehat{G}$ is at most 1.

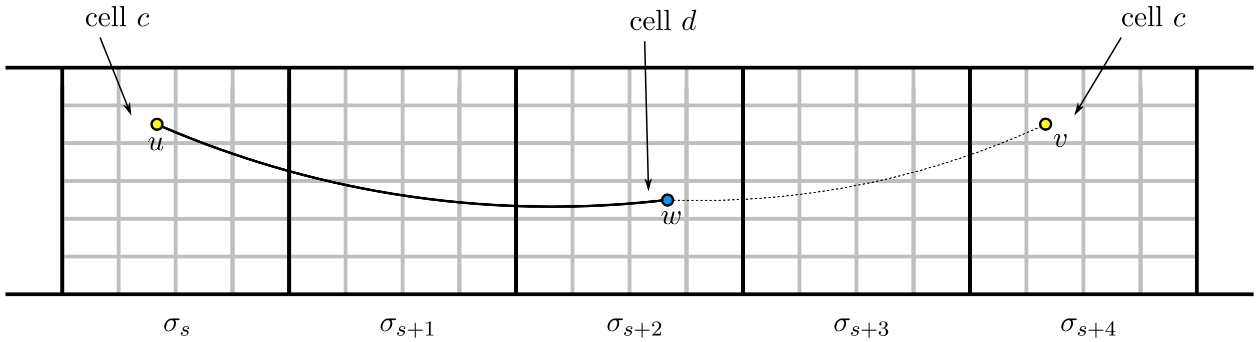

Figure 4: A pivot w for vertices $u \in G_s$ and $v \in G_{s+4}$. Solid line between vertices represents an edge, dotted represents a non-edge.

Proof. Given $u, v \in U_c$, we need to verify whether $u \preceq v$. We first check whether there is a pivot for (u, v) or whether there is a pivot for (v, u). If any of these checks holds, then by Claim 5 we can infer whether $u \preceq v$. On the other hand, if none of them holds, then by Claim 6 we have that $|s - t| \leqslant 3$, where $u \in G_s$ and $v \in G_t$. Then the block order between u and v can be inferred by comparing the moduli of u and v: $u \preceq v$ if and only if $\mathsf{mod}(u)$ is among $\{\mathsf{mod}(v) - x : x \in \{0, 1, 2, 3\}\}$. It is straightforward to implement this verification by a first-order formula with quantifier rank 1 working over $\widehat{G}$. ⌟

A symmetric reasoning yields the following.

Claim 8. The interpretation complexity of $\preceq_d$ over $\widehat{G}$ is at most 1.

Finally, we are left with showing that the interpretation complexity of $\preceq_{c,d}$ over $\widehat{G}$ is at most 2. Given $(u, v) \in U_c \times U_d$, we need to verify whether $u \preceq v$. Let $u \in G_s$, $v \in G_t$, $a = \mathsf{mod}(u) = s \bmod 7$, and $b = \mathsf{mod}(v) = t \bmod 7$. Observe first that if $\mathrm{dist}(a, b) > 1$, then also $|s - t| > 1$, and hence, by Claim 4, the block order between u and v is equivalent to adjacency between u and v. Hence, we are left with considering the case when $\mathrm{dist}(a, b) \leqslant 1$.

The reasoning will be similar as for interpreting $\preceq_c$. Call $w \in V_d$ a *pivot** for (u, v) if $\mathrm{dist}(a, \mathsf{mod}(w)) > 1$, u is adjacent to w, and $w \preceq_d v$. Observe that if there is a pivot* w for (u, v), then the adjacency between u and w together with $\mathrm{dist}(a, \mathsf{mod}(w)) > 1$ imply, by Claim 4, that $u \prec w$. Together with $w \preceq_d v$, this implies that $u \prec v$. Hence, the existence of a pivot* for (u, v) implies that $u \prec v$.

On the other hand, suppose for a moment that $s < t - 1$. Then, due to $\mathrm{dist}(a, b) \leqslant 1$, we actually have $s \leqslant t - 6$. Hence, there exists some modulus r and index p such that $s + 1 < p \leqslant t$, and $r = p \bmod 7$, and $\mathrm{dist}(r, a) > 1$. Since cell d is essential, there exists a vertex $w \in G_p$ such that also $w \in U_d$. Since $s + 1 < p$, by Claim 4 we have that u and w are adjacent, hence w is a pivot* for (u, v).

Therefore, in case $\mathrm{dist}(a, b) \leqslant 1$ we verify whether $u \preceq v$ as follows. First, check whether there is a pivot* for (u, v) or for (v, u). If any of these checks holds, then this forces the block order between u and v, and we can immediately infer whether $u \preceq v$. Otherwise there is neither a pivot* for (u, v) nor for (v, u), so we conclude that $|s - t| \leqslant 1$. Then the block order between u and v can be inferred by comparing the moduli: $u \preceq v$ if and only if $\mathsf{mod}(u)$ is equal either to $\mathsf{mod}(v)$ or to $\mathsf{mod}(v) - 1$.

It is straightforward to implement the above verification using a first-order formula of quantification rank 1 that uses the formula interpreting $\preceq_d$. Since the interpretation complexity of the latter relation is at most 1, it follows that the interpretation complexity of $\preceq_{c,d}$ is at most 2. □

Lemma 16 suggests the following classification of essential cells. An essential cell c is called *social* if there is another essential cell d such that (c, d) is a mixed pair of cells. Essential cells that are not social are called *solitary*, and a vertex u is *social* (resp. *solitary*) if $u \in U_c$ for some social (resp. solitary) cell c. Then Lemma 16 asserts that for any social cell c, the interpretation complexity of $\preceq_c$ over $\widehat{G}$ is at most 2. Intuitively, our next goal is to extend this order to solitary cells as much as possible, and to piece together the obtained orders $\preceq_c$ by interpreting the block order between elements from different cells.

First, we observe that we can in some sense compose the orders $\preceq_{c,d}$ that we have already interpreted. More precisely, let M be the *social graph* defined as follows: the vertex set of M comprises all social cells, and two cells are considered adjacent iff they form a mixed pair. For a component C of the social graph M, we denote $U_C = \bigcup_{c \in C} U_c$.

Lemma 17. *Suppose social cells c and d belong to the same connected component of the social graph M. Then the interpretation complexity of $\preceq_{c,d}$ over $\widehat{G}$ is at most $|\mathcal{C}_k|$.*

Proof. Let $Q = (c = c_1, c_2, \ldots, c_{p-1}, c_p = d)$ be any path in M between c and d. Obviously the length of Q is at most $|\mathcal{C}_k| - 1$, hence Q has at most $|\mathcal{C}_k| - 2$ internal vertices. Consider the following first-order formula with free variables $u \in U_c$ and $v \in U_d$: there exist vertices $w_2 \in V_{c_2}$, $w_3 \in V_{c_3}$, and so on up to $w_{p-1} \in V_{c_{p-1}}$, such that

$$u \preceq_{c,c_2} w_2 \preceq_{c_2,c_3} w_3 \preceq_{c_3,c_4} \cdots \preceq_{c_{p-3},c_{p-2}} w_{p-2} \preceq_{c_{p-2},c_{p-1}} w_{p-1} \preceq_{c_{p-1},d} v.$$

Since all cells $c_2, \ldots, c_{p-1}$ are essential, it is easy to prove that $u \preceq v$ if and only if this formula is satisfied. To see this, note that some or all of the w_i, w_j may belong to the same block, which guarantees $w_i \preceq_{c_i,c_j} w_j$. Moreover, by Lemma 16 and the definition of the social graph, each of the relations $\preceq_{c_i,c_{i+1}}$ has interpretation complexity at most 2. Since we quantify at most $|\mathcal{C}_k| - 2$ intermediate vertices, it follows that the constructed formula has quantifier rank at most $|\mathcal{C}_k|$. $\square$

For a connected component C of the social graph M, by $\preceq_C$ we denote the block order restricted to pairs of vertices from U_C. Lemma 17 immediately yields the following.

Corollary 18. *For any connected component C of M, the interpretation complexity of $\preceq_C$ over $\widehat{G}$ is at most $|\mathcal{C}_k|$.*

Finally, for the purpose of further reasoning we need to observe some basic properties of cells. More precisely, we analyze how the cell of a vertex in its block corresponds to its cell in the overall derivation σ. The following assertion follows immediately from the definition of composing derivations and the fact that ϕ is an idempotent function.

Observation 19. *Suppose $u \in G_s$ for some $s \in [n]$ and $u \in U_c$ for some cell $c = (i, X)$. Then in derivation σ, vertex u belongs to cell $\sigma[c']$ for $c' = (i', X')$ defined as follows*

- *$i' = i$ when $s = n$ and $i' = \phi(i)$ otherwise; and*
- *$X' = X$ when $s = 1$ and $X' = \phi^{-1}(X)$ otherwise.*

Next, we verify that cells c and c' as in Observation 19 behave in almost the same manner in the social graph M.

Lemma 20. *Suppose cells $c = (i, X)$ and $c' = (i', X')$ are such that $i' = i$ or $i' = \phi(i)$, and $X' = X$ or $X' = \phi^{-1}(X)$. Then for any cell d, the pair (c, d) is of the same type—negative, positive, or mixed—as the pair (c', d). Consequently, c is social if and only if c' is social, and provided c and c' are social, they belong to the same connected component of the social graph M.*

Proof. Since the second claim of the lemma statement follows immediately from the first claim, we focus on proving the latter. Let $d = (j, Y)$. Observe that $\phi(i') = \phi(i)$, because ϕ is idempotent and i' is equal to i or $\phi(i)$. Observe also that $\phi(j) \in X$ if and only if $\phi(j) \in \phi^{-1}(X)$, again because ϕ is idempotent. Hence, we have $\phi(i) \in Y$ if and only if $\phi(i') \in Y$, and $\phi(j) \in X$ if and only if $\phi(j) \in X'$, because X' is equal either to X or to $\phi^{-1}(X)$. It follows that pairs (c, d) and (c', d) are of the same type. $\square$

Note that Lemma 20 in particular applies to $d = c$ and $d = c'$. We can thus infer that pairs (c, c), (c, c'), and (c', c') are always of the same type.

Flipping the graph. In order to simplify the analysis of the instance, we perform an auxiliary flipping operation on G. Let $Z \subseteq \binom{\mathcal{C}_k}{1,2}$ be the subset of positive pairs of cells; that is, for each pair (c, d) of essential cells (possibly $c = d$), if (c, d) is positive then we put $\{c, d\}$ into Z. Then, construct a graph H from G by performing the flip between U_c and U_d for every $\{c, d\} \in Z$, and performing the flip between U_c and U_c for every $\{c\} \in Z$. Let $\widehat{H}$ be the structure obtained from $\widehat{G}$ by enriching it with the adjacency relation of H; note that $\widehat{H}$ still contains the adjacency relation of G as well as unary predicates for sets U_c and moduli. Observe that the adjacency relation in H can be interpreted by a first-order formula of quantifier rank 0 working on $\widehat{G}$, so its interpretation complexity over $\widehat{G}$ is 0. Hence, from now on we may assume that we work over $\widehat{H}$.

We first check that the construction of H is actually equivalent to taking the result of a Z-flip in σ. Note that this claim is non-trivial, since for each cell $c \in \mathcal{C}_k$, the set $\sigma[c]$ is not necessarily equal to U_c.

Lemma 21. *The Z-flip of σ is equal to H.*

Proof. Let H' be the Z-flip of σ. It suffices to show that a pair of vertices u, v is adjacent in H if and only if it is adjacent in H'. Let $u \in U_c$, $u \in \sigma[c']$, $v \in U_d$, and $v \in \sigma[d']$; here, pairs of cells (c, c') and (d, d') are as described in Observation 19. By applying Lemma 20 twice, we have that (c, d) is of the same type as (c', d), which in turn is of the same type as (c', d'). If (c, d) is negative, then (c', d') is also negative, and the adjacency between u and v from G is negated neither when constructing H nor when constructing H'. Consequently, u and v are adjacent in H iff they are adjacent in G iff they are adjacent in H'. The same holds also when (c, d) (equivalently, (c', d')) is mixed. However, when (c, d) is positive, then (c', d') is also positive, so the adjacency between u and v is negated both when constructing H and when constructing H'. Consequently, u and v are adjacent in H iff they are non-adjacent in G iff they are adjacent in H'. $\square$

The following lemma shows that the performed flipping operation simplifies the adjacency relation in the graph.

Lemma 22. *Suppose u and v are vertices adjacent in H, where $u \in G_s \cap U_c$ and $v \in G_t \cap U_d$ for some $s, t \in [n]$ and $c, d \in \mathcal{C}_k$; possibly $c = d$. Suppose further that (c, d) is not a mixed pair. Then $|s - t| \leqslant 1$.*

Proof. For the sake of contradiction, suppose that $|s - t| > 1$. By symmetry, we may further assume $s < t - 1$. Let $c = (i, X)$ and $d = (j, Y)$.

Observe that vertex u was colored with color i in σ_s, then it was recolored to $\phi(i)$ when composing with σ_{s+1}, and its color stayed equal to $\phi(i)$ when composing with $\sigma_{s+2}, \ldots, \sigma_{t-1}$. Hence, when composing with σ_t, the edge between u and v was added in G if and only if $\phi(i) \in Y$. As (c, d) is not a mixed pair, it is either positive or negative. If (c, d) is negative, then $\phi(i) \notin Y$, u and v were not adjacent in G and there was no flipping between U_c and U_d when constructing H. Consequently u

and v stay non-adjacent in H; a contradiction. Otherwise, if (c, d) is positive, then $\phi(i) \in Y$, u and v are adjacent in G, but we applied flipping between U_c and U_d when constructing H. Consequently u and v are non-adjacent in H; a contradiction again. $\square$

Lemma 22 tells us that apart from adjacencies in mixed pairs, the adjacency relation in H is local: vertices adjacent in H lie either in the same or in neighboring blocks, provided the relation between their cells is not mixed. Note that this prerequisite is satisfied always when at least one of the vertices is solitary.

Defining the block order. For a connected component F of H, by $\preceq_F$ we denote the block order restricted to the vertices of F. We now concentrate on proving the following result, after which we will be essentially done.

Lemma 23. *For each connected component F of H, the interpretation complexity of $\preceq_F$ over $\widehat{H}$ enriched with a unary predicate selecting vertices of F is at most $3|\mathcal{C}_k| + 5$.*

Before proving Lemma 23 let us see why the Definable Order Lemma follows from it.

Proof of the Definable Order Lemma assuming Lemma 23. Recall that by Lemma 21 we know that H is equal to the Z-flip of σ. As argued, we may assume that we work over $\widehat{H}$ instead of original G.

We now need to write an MSO formula that for given vertices u, v checks whether u and v are in the same connected component of H and moreover $u \preceq v$. The assertion that u and v are in the same connected component of H is clearly expressible in MSO, as the adjacency relation of H is present in the structure. For the assertion $u \preceq v$ we will use an appropriate interpretation provided by Lemma 23. More precisely, suppose F is the connected component of H that contains u and v. Then Lemma 23 asserts that the interpretation complexity of $\preceq$ restricted to F is at most $3|\mathcal{C}_k| + 5$; more precisely, there exists a formula φ_F of quantifier rank at most $3|\mathcal{C}_k| + 5$ and using at most $3|\mathcal{C}_k| + 5$ monadic variables that interprets $\preceq$ restricted to F. Note that the number of possible such formulas φ_F is bounded by a function of k. Hence, for every such possible formula ψ we may introduce an (existentially guessed) monadic parameter X_ψ that selects the union of vertex sets of those connected components F of H for which $\varphi_F = \psi$. Then to check whether $u \preceq v$ one may use the formula ψ for which $u, v \in X_\psi$ holds. $\square$

Hence, we are left with proving Lemma 23. The proof is divided into several steps. Intuitively, we try to connect components of the social graph M with each other, as well as vertices from solitary cells to these components, using paths in H whose internal vertices are all solitary. By Lemma 22 we know that such paths cannot jump between distant blocks, hence we will be able to control guessing them in MSO. We will also use the idempotence of the abstraction $e = [\![\sigma]\!]$, which is equal to $[\![\sigma_s]\!]$ for all $s \in [n]$, to reason about the existence of some “local” paths realizing sought connections.

We first formalize the kind of connections we are interested in. A path Q in H is called *solitary* if the following conditions hold:

- all internal vertices of Q are solitary; and
- if the endpoints of Q belong to U_c and U_d, respectively, then (c, d) is not a mixed pair.

Note that the endpoints of a solitary path may be social, but we explicitly exclude the case when they belong to cells forming a mixed pair. The following assertion about locality of solitary paths follows immediately from Lemma 22.

Observation 24. *If Q is a solitary path, then any two consecutive vertices on Q belong either to the same or to neighboring blocks.*

We first show that when two different components of the social graph M can be connected by a solitary path, then they can be connected by a local solitary path. For this we crucially use the idempotence of abstraction e.

Lemma 25. *Suppose C, D are two different connected components of the social graph M. Suppose further that there is a solitary path in H that starts in a vertex of U_C and ends in a vertex of U_D. Then, for each $s \in [n]$, there is a solitary path that starts in an vertex of U_C, ends in a vertex of U_D, and whose all vertices belong to the block G_s.*

Proof. Let Q be the path whose existence is asserted in the lemma statement. Consider Q as a path in the Z-flip of σ, which is equal to H by Lemma 21. Suppose Q starts in $\sigma[c']$, ends in $\sigma[d']$, and the set of cells of σ traversed by the internal vertices of Q is W'. Since Q is solitary, by Lemma 20 we infer that all the cells of W' are solitary. The endpoints of Q belong then to U_c and U_d for some $c, d \in \mathcal{C}_k$, respectively, such that $c \in C$, $d \in D$, and pairs c, c' and d, d' are as described in Observation 19. In particular, by Lemma 20 we have that $c' \in C$ and $d' \in D$.

Since the connectivity registries of $[\![\sigma]\!]$ and $[\![\sigma_s]\!]$ are equal for each $s \in [n]$, the existence of Q implies the existence of a path Q' in the Z-flip of σ_s such that Q' starts in $\sigma_s[c']$, ends in $\sigma_s[d']$, and all the internal vertices of Q' belong to cells $\sigma_s[b]$ for $b \in W'$. This means that Q' starts in $U_{c'} \subseteq U_C$, ends in $U_{d'} \subseteq U_D$, and every its internal vertex belongs to some U_b for $b \in W'$, which implies that it is solitary. Since (c, d) is not mixed due to Q being solitary, we infer that (c', d') is also not mixed by Lemma 20. Hence Q' is a solitary path and we are done. $\square$

Finally, we show that if a solitary vertex can be connected to a component of the social graph M via a solitary path, then there is also such a solitary path that is local.

Lemma 26. *Suppose $u \in G_s$ is a solitary vertex and there is a solitary path Q in H that leads from u to some social vertex v, say belonging to U_D for some connected component D of the social graph M. Then there is also a solitary path Q' in H that leads from u to some social vertex v' belonging to U_D such that all vertices traversed by Q' belong to blocks G_{s-1}, G_s, and G_{s+1}.*

Proof. By changing Q if necessary, we may assume that out of solitary paths that lead from u to any social vertex from U_D, Q is a path that uses the smallest number of vertices outside of the blocks G_{s-1}, G_s, and G_{s+1}. It suffices to show that Q chosen in this manner in fact traverses only vertices from these blocks, so assume otherwise. Let $d \in D$ be such that $v \in U_d$.

In the following, we regard Q as traversed in the direction from u to v. Let r be any vertex on Q that does not belong to any of the blocks G_{s-1}, G_s, and G_{s+1}. Let x be the earliest vertex before r on Q with the following property: all the vertices between x and r (inclusive) on Q do not belong to G_s. Similarly, let y be the latest vertex after r on Q with the following property: all the vertices between r and y (inclusive) on Q do not belong to G_s. Note that x has a predecessor on Q that belongs to G_s, while y either has a successor on Q that belongs to G_s, or $y = v$. Let R be the infix of Q between x and y (inclusive). Note that x is an internal vertex of Q, hence it is solitary.

By Lemma 22, we know that every pair of consecutive vertices on Q either belong to the same or to neighboring blocks. Consequently, by the definition of R we have that one of the following two cases holds:

(1) all vertices on R belong to $\bigcup_{t<s} G_t$, $x \in G_{s-1}$, and (either $y = v$ or $y \in G_{s-1}$); or

(2) all vertices on R belong to $\bigcup_{t>s} G_t$, $x \in G_{s+1}$, and (either $y = v$ or $y \in G_{s+1}$).

To prove the lemma, it suffices to show the following claim about the existence of a suitable replacement path for the infix R. See Figure 5 for reference.

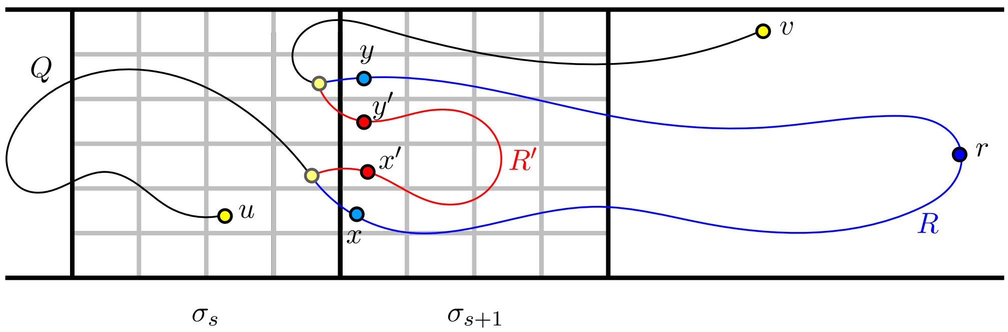

Figure 5: The situation in Claim 9 for the alternative (2). The replaced path R is depicted in blue, the replacement path R' is depicted in red.

Claim 9. There exists a solitary path R', say with endpoints x' and y', that satisfies the following properties:

(a) all vertices traversed by R' belong to the same block as x;

(b) x' is solitary and is adjacent in H to the predecessor of x on Q;

(c) if $y = v$, then y' belongs to U_D;

(d) if $y \neq v$, then y' is solitary and is adjacent in H to the successor of y on Q.

Observe that it suffices to prove Claim 9 for the following reason. Take Q' to be Q with the infix R replaced with R'. By (b), (c), and (d), Q' constructed in this manner is still a solitary path in H, and it ends in a vertex of U_D. However, by (a), Q' traverses strictly fewer vertices outside of blocks G_{s-1}, G_s, and G_{s+1}, because R' is entirely contained in these blocks, while R traversed vertex r which lies outside of these blocks. Thus, the existence of Q' contradicts the initial choice of Q.

Hence, from now on we focus on proving Claim 9. We consider two cases: either alternative (1) holds, or alternative (2) holds.

Suppose first that alternative (2) holds: R is contained in $\bigcup_{t>s} G_t$, $x \in G_{s+1}$, and (either $y = v$ or $y \in G_{s+1}$). Consider derivation

$$\sigma' = \sigma_{s+1} \cdot \sigma_{s+2} \cdot \ldots \cdot \sigma_n.$$

Since the abstraction e is idempotent, we have that $[\![\sigma']\!] = e$, in particular $[\![\sigma']\!]$ has the same connectivity registry as $[\![\sigma_t]\!]$ for each $t \in [n]$. Let $c_x, c_y, c'_x, c'_y \in \mathcal{C}_k$ be such that $x \in U_{c_x}$, $x \in \sigma'[c'_x]$, $y \in U_{c_y}$, and $y \in \sigma'[c'_y]$. By Observation 19, cells c_x and c'_x have the same profile, say X, and may differ in the color. Similarly, cells c_y and $c_{y'}$ have the same profile, say Y, and may differ in the color.

Consider now the path R as a path in the Z-flip of the derivation σ', which is equal to the subgraph induced in H by the vertex set of the underlying graph of σ', by Lemma 21. This path starts in $\sigma'[c'_x]$, ends in $\sigma'[c'_y]$, and let W be the set of cells in σ' traversed by the internal vertices of R. Note that since R is solitary, all the internal vertices of R are solitary, hence by Lemma 20 we have that all the cells of W are solitary. Observe now that the existence of R and the fact that the connectivity registries of σ' and σ_{s+1} are equal certify that there is a path R' in the Z-flip of σ_{s+1}

that starts in $\sigma_{s+1}[c'_x]$, ends in $\sigma_{s+1}[c'_y]$, and travels through cells of W. We claim that R' satisfies all the required properties.

First, observe that R' is a path in the Z-flip of σ_{s+1}, which is equal to the subgraph induced in H by the vertex set of G_{s+1}. Thus, R' is a path in H with all vertices belonging to G_{s+1}. It starts with some vertex x' belonging to $U_{c'_x}$. Note here that, by Lemma 20, c'_x is solitary because c_x is solitary, hence x' is solitary as claimed. Similarly, R' ends in some vertex y' belonging to $U_{c'_y}$. Note that y' is solitary if and only if y is solitary, which happened if and only if $y \neq v$. Observe further that R' is solitary because all the internal vertices traversed by R' belong to sets $\sigma_{s+1}[b] \subseteq U_b$ for $b \in W$, and all the cells of W are solitary.

Since the profiles of x and x' are equal, the predecessor of x on Q is adjacent to x' in G if and only if it is adjacent to x in G. Supposing this predecessor belongs to U_b for some $b \in \mathcal{C}_k$, by Lemma 20 we infer that type of the pair (b, c_x) is the same as the type of (b, c'_x). Consequently, the adjacency between the predecessor and x is negated when constructing H if and only if the adjacency between the predecessor and x' is negated. As the predecessor and x are adjacent in H, we infer that the predecessor is adjacent also to x' in H. A symmetric reasoning shows that provided $y \neq v$, the successor of y on Q is adjacent to y' in H. Finally, observe that if $y = v$, then since $y' \in U_{c'_y}$ and c_y, c'_y belong to the same connected component of the social graph M (by Lemma 20), we have that $y' \in U_D$. This concludes the verification that R' has all the claimed properties.

Next, consider the alternative (2): R is contained in $\bigcup_{t<s} G_t$, $x \in G_{s-1}$, and (either $y = v$ or $y \in G_{s-1}$). The proof is almost entirely symmetric; we simply consider the derivation $\sigma' = \sigma_1 \cdot \ldots \cdot \sigma_{s-1}$, and use the equality of connectivity registries of $[\![\sigma']\!]$ and $[\![\sigma_{s-1}]\!]$. The only different detail is the verification that x' is adjacent to the predecessor of x on Q, and that y' is adjacent to the successor of y on Q (provided $y \neq v$). Now, by Observation 19 we have that x has the same color in σ' and in σ_{s-1}, not the same profile as before. Consequently, path R' is chosen so that x and x' have the same color in σ', which implies that the predecessor of x on Q is adjacent to x' in G if and only if it is adjacent to x in G, because this assertion is equivalent to the common color of x and x' belonging to the profile of the predecessor. The rest of the reasoning, including the y-counterpart, is exactly symmetric; we leave the verification to the reader. $\square$

With all the tools gathered, we can extend further our interpretation of the block order. First, we say that two connected components C and D of the social graph M are *close* if there is a solitary path in H that starts in a vertex of U_C and ends in a vertex of U_D. Consider a graph with the connected components of M as the vertex set, where two components are considered adjacent whenever they are close. The connected components of this graph will be called *clusters*, and we identify each cluster A with the union of sets U_C for C belonging to A. For a cluster A, by $\preceq_A$ we denote the restriction of the block order $\preceq$ to the vertices of A.

Lemma 27. *For a cluster A, the interpretation complexity of $\preceq_A$ over $\widehat{H}$ is at most $3|\mathcal{C}_k|$.*

Proof. Take any vertices $u, v \in A$; we want to encode the check whether $u \preceq v$ using an MSO formula. Suppose then that $u \in G_s$ and $v \in G_t$.

Since $u, v \in A$, there is some $p \leqslant |\mathcal{C}_k|$ and components $C^1, C^2, \ldots, C^p$ of M such that $u \in U_{C^1}$, $v \in U_{C^p}$, and C^i is close to C^{i+1} for each $i \in [p-1]$. Let Q_i be a path that certifies this closeness; that is, Q_i is a solitary path that starts in U_{C^i} and ends in $U_{C^{i+1}}$. By Lemma 25, we can choose paths Q_i so that all their vertices are contained in G_s.

Consider an MSO formula with free variables u and v expressing the following property: for some $p \leqslant |\mathcal{C}_k|$, there are components $\{C^i : i \in [p]\}$ of M, vertices $\{x_i, y_i : i \in [p-1]\}$, and paths $\{Q_i : i \in [p]\}$ in H such that

- for each $i \in [p-1]$, we have that $x_i \in U_{C^i}$ and $y_i \in U_{C^{i+1}}$;

- denoting $y_0 = u$ and $x_p = v$, for each $i \in [p]$ we have $y_{i-1} \preceq_{C^i} x_i$; and
- each path Q_i is solitary, has endpoints x_i and y_i, and all vertices traversed by it have the same modulus.

Since the social graph M is fixed, and the quantifier rank of each $\preceq_{C^i}$ is at most $|\mathcal{C}_k|$ by Corollary 18, it is straightforward to construct such a formula of quantifier rank at most $3|\mathcal{C}_k|$. By the discussion of the previous paragraph, we see that provided $u \preceq v$, vertices u, v satisfy this formula. It remains to show that if the formula holds for some $u, v \in A$, then indeed $u \preceq v$.

To this end, fix components $\{C^i : i \in [p]\}$, vertices $\{x_i, y_i : i \in [p-1]\}$, and paths $\{Q_i : i \in [p]\}$ that witness the satisfaction of the formula for u and v. Each Q_i is a solitary path whose vertices have all the same modulus. From Observation 24 it follows that Q_i must be entirely contained in one block, hence $x_i \equiv y_i$ for all $i \in [p-1]$. We conclude that

$$u \preceq_{C^1} x_1 \equiv y_1 \preceq_{C^2} x_2 \equiv y_2 \preceq_{C^3} \ldots \preceq_{C^{p-2}} x_{p-1} \equiv y_{p-1} \preceq_{C^{p-1}} v,$$

so indeed $u \preceq v$. □

Before we proceed to the proof of Lemma 23, we need one more claim about the relation between clusters and connected components of H.

Lemma 28. *For each connected component F of H, there is at most one cluster that has a nonempty intersection with F.*

Proof. It suffices to prove that any two social vertices $u, v \in F$ belong to the same cluster. Let Q be a path in H with endpoints u and v, and let $u = w_1, w_2, \ldots, w_p = v$ be the social vertices traversed by Q, in the order of traversal. For each $i \in [p]$, let C^i be the component of the social graph M such that $w_i \in U_{C^i}$, and let $w_i \in U_{c^i}$ for some $c^i \in C^i$. If (c^i, c^{i+1}) is a mixed pair, then c^i and c^{i+1} are adjacent in M, and consequently $C^i = C^{i+1}$. Otherwise, the infix of Q between w_i and w_{i+1} is a solitary path that witnesses that C^i and C^{i+1} are close. Concluding, for each $i \in [p-1]$ we have that C^i and C^{i+1} are either close or equal, hence all the components C^i for $i \in [p]$ belong to the same cluster. □

We can now use the interpretation given by Lemma 27 together with local connections given by Lemma 26 to prove the conclusion of Lemma 23 for any component F of H that contains some social vertex.

Lemma 29. *Suppose F is a connected component of H that contains some social vertex. Then the interpretation complexity of $\preceq_F$ over $\widehat{H}$ enriched with a unary predicate selecting vertices of F is at most $3|\mathcal{C}_k| + 5$.*

Proof. Since F contains some social vertex, there is some cluster A that intersects F. By Lemma 28, this cluster A is unique. Observe that we can check in MSO that a subset of cells forms this unique cluster A, hence by making a disjunction over all subsets of cells, we can assume further that the constructed formula may use an additional unary predicate that selects the vertices of A.

Take any $u, v \in F$; we would like to express the verification whether $u \preceq v$ by an MSO formula. For simplicity of presentation, we assume that u and v are solitary. The construction of the formula in the other cases is even simpler and we discuss it at the end. Let $s, t \in [n]$ be such that $u \in G_s$ and $v \in G_t$.

Observe that there are solitary paths Q and R in H such that Q connects u with some social vertex $x \in A$, and R connects v with some social vertex $y \in A$. Indeed, it suffices to take shortest paths from u and v to the set of social vertices. By Lemma 26, we can choose Q so that it traverses only vertices from G_{s-1}, G_s, and G_{s+1}; in particular x belongs to one of these blocks. Similarly, we can choose R so that it traverses only vertices from G_{t-1}, G_t, and G_{t+1}; in particular y belongs to one of these blocks.

Consider an MSO formula, with free variables u, v and working over $\widehat{H}$ enriched with two unary predicates that select the vertices of A and F, respectively, which expresses the following properties:

- There exists a vertex $x \in A$ that can be connected with u by a solitary path traversing only vertices of the same or neighboring moduli as u.
- There exists a vertex $y \in A$ that can be connected with v by a solitary path traversing only vertices of the same or neighboring moduli as v.
- There exists a vertex $x' \in A$ such that x and x' belong to the same block if the moduli of u and x are equal, x' belongs to the block immediately before the block of x in case the modulus of x is one larger than the modulus of u, and x' belongs to the block immediately after the block of x in case the modulus of x is one smaller than the modulus of u.
- There exists a vertex $y' \in A$ such that y and y' belong to the same block if the moduli of u and y are equal, y' belongs to the block immediately before the block of y in case the modulus of y is one larger than the modulus of v, and y' belongs to the block immediately after the block of y in case the modulus of y is one smaller than the modulus of v.
- It holds that $x' \preceq_A y'$.

Having existentially quantified x, y, the first two properties can be easily checked using formulas of quantifier rank 2. Then, having existentially quantifier x', y', the next two properties can be checked using formulas of quantifier rank 1 that use the relation $\preceq_A$, so $1 + 3|\mathcal{C}_k|$ in total by Lemma 27. Indeed, for example to check whether x' is in the block immediately after the block x, it suffices to make sure that $x \prec_A x'$ and there is no vertex $x'' \in A$ such that $x \prec_A x'' \prec_A x'$. The last check requires quantifier rank $3|\mathcal{C}_k|$, by Lemma 27. From the discussion of the previous paragraph it is clear that the formula described above holds provided we have $u \preceq v$. Therefore, we are left with verifying that the satisfaction of the formula implies that $u \preceq v$.

Let x, x', y, y' be vertices whose existence is asserted by the satisfaction of the formula, and let Q, R be the solitary paths witnessing the satisfaction of the first two properties. Since Q is solitary, by Observation 24, every two consecutive vertices on Q belong to the same or to neighboring blocks. Since the vertices of Q have only one of three moduli: the modulus of u or the two neighboring ones, it follows that all vertices of Q in fact must be contained in the blocks G_{s-1}, G_s, or G_{s+1}. Then the quantification of x' ensures us that $x' \in G_s$. A symmetric reasoning shows that $y' \in G_t$, hence $u \preceq v$ is equivalent to $x' \preceq_A y'$, which we check in the last property.

It remains to argue what happens if one or two of the vertices u, v are social. In case both of them are social, we can simply use the already interpreted relation $\preceq_A$. In case one of them, say u, is social, we proceed exactly as above, except we put $u = x = x'$ instead of finding x, x' via an existential guess of the path Q. □

Finally, we resolve also connected components of H that have only solitary vertices.

Lemma 30. *Suppose F is a connected component of H that contains no social vertices. Then the interpretation complexity of $\preceq_F$ over $\widehat{H}$ enriched with a unary predicate selecting vertices of F is at most* 5.

Proof. Our first goal is to interpret $\equiv_F$, the block equivalence restricted to F. To this end, we prove locality of connections within F in a similar manner as in the proof of Lemma 26.

Claim 10. For any $u, v \in F$ such that $u, v \in G_s$ for some $s \in [n]$, there exists a path in H that connects u and v and whose all vertices belong to blocks G_{s-1}, G_s, and G_{s+1}.

Proof. Among paths in H that connect u and v, choose Q to be the one that minimizes the number of traversed vertices outside of blocks G_{s-1}, G_s, and G_{s+1}. Since F contains no social vertices, Q is solitary. In particular, by Observation 24 we have that every two consecutive vertices on Q belong to either to the same or to neighboring blocks.

For the sake of contradiction, suppose Q traverses some vertex r outside of blocks G_{s-1}, G_s, and G_{s+1}. Similarly as in the proof of Lemma 26, we choose x to the earliest vertex on Q before r such that all vertices between x (inclusive) and r are outside of the block G_s, and we choose y to be the latest vertex on Q after y such that all vertices between r and y (inclusive) are also outside of G_s. In particular, x and y belong to the same block, being either G_{s-1} or G_{s+1}. Let R be the infix of Q between x and y. By the same reasoning as in the proof of Lemma 26 (see Claim 9 therein), there exists a path R' in H such that:

- all vertices of R' are solitary and belong to the same block as x and y, which is either G_{s-1} or G_{s+1};
- R' starts in some vertex x' that is adjacent in H to the predecessor of x on Q;
- R' ends in some vertex y' that is adjacent in H to the successor of y on Q.

Then replacing R with R' in Q yields a path Q' in H that connects u and v, but has strictly less vertices outside of blocks G_{s-1}, G_s, and G_{s+1}. This is a contradiction with the choice of Q. ⌟

Given $u, v \in F$, consider an MSO formula expressing the following property: u and v have the same modulus, and there is a path Q in H connecting them which traverses only vertices of the same or neighboring moduli as u and v. Since all paths within F are solitary, by Observation 24 we have that the satisfaction of this formula implies that u and v are in the same block. On the other hand, by Claim 10 we have that the formula will be satisfied for all $u, v \in F$ that reside in the same block. Consequently, the presented formula, which has quantifier rank 2, interprets the block equivalence $\equiv_F$.

In order to interpret the block order $\preceq_F$, consider an MSO formula with free variables $u, v \in F$ expressing the following property. For every path Q in H that leads from u to v, if w is the last vertex on Q that is from the same block as u, then either $w = v$ or the successor of w on Q has modulus larger by one than the modulus of v. To quantify w we use the block equivalence $\equiv_F$ that we interpreted in the previous paragraph, hence it is easy to obtain such a formula with quantifier rank 5. To verify that this formula indeed defines the block order $\preceq_F$, observe that if $u \in G_s$ and w is the last vertex from G_s on Q, then either $w = v$, or the successor of w on Q belongs to G_{s-1} (which implies that $u \succ v$) or to G_{s+1} (which implies that $u \prec v$). □

Lemma 23 now follows directly from combining Lemmas 29 and 30: we check whether F contains a social vertex, and we apply the interpretation of Lemma 29 or 30 depending on the result. Hence, the proof of the Definable Order Lemma is complete.

7 Conclusions

We proved that for every k there is an MSO-transduction that defines for a given graph of linear cliquewidth k a width-$f(k)$ clique decomposition of this graph. A consequence of this result is that recognizability equals CMSO_1-definability on graphs of bounded linear cliquewidth.

The obvious open question is whether our result can be generalized from linear clique decompositions to general clique decompositions. The approach used in [2] for lifting the pathwidth case to the treewidth case heavily relies on combinatorial techniques specific to tree decompositions, and hence it seems hard to translate the ideas to the setting of clique decompositions.

References

[1] I. Adler and M. Kanté. Linear rank-width and linear clique-width of trees. *Theor. Comput. Sci.*, 589:87–98, 2015.

[2] M. Bojańczyk and M. Pilipczuk. Definability equals recognizability for graphs of bounded treewidth. In *LICS 2016*, pages 407–416. ACM, 2016.

[3] M. Bojańczyk and M. Pilipczuk. Optimizing tree decompositions in MSO. In *STACS 2017*, volume 66 of *LIPIcs*, pages 15:1–15:13. Schloss Dagstuhl — Leibniz-Zentrum für Informatik, 2017.

[4] B. Courcelle. The Monadic Second-Order logic of graphs. I. Recognizable sets of finite graphs. *Inf. Comput.*, 85(1):12–75, 1990.

[5] B. Courcelle and J. Engelfriet. *Graph Structure and Monadic Second-Order Logic — A Language-Theoretic Approach*, volume 138 of *Encyclopedia of mathematics and its applications*. Cambridge University Press, 2012.

[6] B. Courcelle, J. A. Makowsky, and U. Rotics. Linear time solvable optimization problems on graphs of bounded clique-width. *Theory Comput. Syst.*, 33(2):125–150, 2000.

[7] B. Courcelle and S. Olariu. Upper bounds to the clique width of graphs. *Discrete Applied Mathematics*, 101(1-3):77–114, 2000.

[8] F. Gurski and E. Wanke. On the relationship between NLC-width and linear NLC-width. *Theor. Comput. Sci.*, 347(1-2):76–89, 2005.

[9] P. Heggernes, D. Meister, and C. Papadopoulos. Graphs of linear clique-width at most 3. *Theor. Comput. Sci.*, 412(39):5466–5486, 2011.

[10] P. Heggernes, D. Meister, and C. Papadopoulos. Characterising the linear clique-width of a class of graphs by forbidden induced subgraphs. *Discrete Applied Mathematics*, 160(6):888–901, 2012.

[11] M. Kufleitner. The height of factorization forests. In *MFCS 2008*, volume 5162 of *LNCS*, pages 443–454. Springer, 2008.

[12] S. Oum. Rank-width and vertex-minors. *J. Comb. Theory, Ser. B*, 95(1):79–100, 2005.

[13] S. Oum and P. D. Seymour. Approximating clique-width and branch-width. *J. Comb. Theory, Ser. B*, 96(4):514–528, 2006.

[14] I. Simon. Factorization forests of finite height. *Theor. Comput. Sci.*, 72(1):65–94, 1990.