Simulating Logspace-Recursion with Logarithmic Quantifier Depth

Steffen van Bergerem[*], Martin Grohe[†], Sandra Kiefer[‡], and Luca Oeljeklaus[†]
[*]Humboldt-Universität zu Berlin, Berlin, Germany
Email: steffen.van.bergerem@informatik.hu-berlin.de
[†]RWTH Aachen University, Aachen, Germany
Email: {grohe,oeljeklaus}@informatik.rwth-aachen.de
[‡]University of Oxford, Oxford, United Kingdom
Email: sandra.kiefer@cs.ox.ac.uk

Abstract

The fixed-point logic LREC$_=$ was developed by Grohe et al. (CSL 2011) in the quest for a logic to capture all problems decidable in logarithmic space. It extends FO+C, first-order logic with counting, by an operator that formalises a limited form of recursion. We show that for every LREC$_=$-definable property on relational structures, there is a constant k such that the k-variable fragment of first-order logic with counting quantifiers expresses the property via formulae of logarithmic quantifier depth. This yields that any pair of graphs separable by the property can be distinguished with the k-dimensional Weisfeiler–Leman algorithm in a logarithmic number of iterations. In particular, it implies that a constant dimension of the algorithm identifies every interval graph and every chordal claw-free graph in logarithmically many iterations, since every such graph admits LREC$_=$-definable canonisation.

Index Terms

counting logic, Weisfeiler–Leman algorithm, graph isomorphism, interval graphs

I. Introduction

By Fagin's celebrated theorem [12], over all finite structures, the complexity class NP is precisely the class of all computational problems that can be expressed via formulae in existential second-order logic. This means that the problem to decide whether a structure has a certain property is in NP if and only if there is a formula in existential second-order logic that defines the property. The theorem can be seen as the starting point of the field of descriptive complexity theory [31, 17], which aims at describing or *capturing* complexity classes via logics. Milestones include the results that, on ordered structures, fixed-point logic FP captures PTIME [29, 46] and deterministic transitive closure logic DTC captures LOGSPACE [30]. However, on general unordered structures, neither of these two results

1

holds; for both complexity classes, the quest for capturing logics still continues and has possibly become the most important question in the field.

Concerning LOGSPACE, even adding counting operators to DTC does not capture the class on trees yet [10]. Towards tackling this, Grohe et al. designed in [24] the logic LREC, which captures LOGSPACE on directed trees and is strictly contained in FP+C, the extension of FP by counting quantifiers. LREC extends first-order logic with counting by an operator which enables a limited version of recursion. The idea behind it is that, as in fixed-point logics, some power of fixed-point operators should be allowed, but the amount of possible recursion shall not lead to the expressive power exceeding logarithmic-space computation.

By extending LREC further to the logic LREC$_=$, a logic to capture LOGSPACE on all undirected trees and on all interval graphs was found [24]. LREC$_=$ is strictly contained in FP+C [7] and in the logic Choiceless Logarithmic Space [15].

More standard "first-order" logics such as transitive closure logic and its fragments quantify over vertices of the input graph. Addressing a single vertex in an n-element graph requires logarithmically many bits. Thus, to remain in LOGSPACE, such a logic can only store a bounded number of vertices at any time. This severely limits the expressiveness. The limited-recursion operator of LREC as an explicit resource management allows it to use the logarithmic space in a more effective way and look at more vertices at the same time. An easy example illustrating how this might be possible is the following: suppose we have already stored a vertex v of degree d. Then we only need $\log d$ bits to address any of its neighbours. As d may be much smaller than n, this may allow us to save space. In this way, the logic LREC is similar to choiceless polynomial time (CPT) [4, 14], which also has an explicit resource control, albeit to be able to more effectively exploit polynomial time. In the same way that CPT is strictly more powerful than FP+C and does not fit into the standard framework of finite-variable logics, LREC is more powerful than deterministic and symmetric transitive closure logic with counting, and it is not even contained in full transitive closure logic with counting, a logic capturing nondeterministic logarithmic space on ordered structures. Larger complexity classes such as AC^1, which describes parallel logarithmic-time complexity, are naturally described in terms of fixed-point logic with logarithmically many iterations. While it is known that LREC is contained in FP+C, it is not obvious that it can be simulated by logarithmically many fixed-point iterations, as the limited recursion of LREC may be polynomially deep. This is the question we address in this paper.

To understand the expressiveness of fixed-point logic with counting, it has turned out to be very fruitful to embed it into the finite-variable logics of first-order logic with counting [5, 44]. Then the number of fixed-point iterations corresponds naturally to the quantifier depth.

Over the decades of research, these counting-logic fragments have exhibited links to many other areas from practical and theoretical computer science [1, 2, 5, 8, 22, 18, 43]. A striking connection exists to the Weisfeiler–Leman algorithm, which is a procedure that computes and refines in an iterative and isomorphism-invariant way colours in the input graph. For every $k \in \mathbb{N}_{>0}$, its k-dimensional version runs in polynomial time and the computed colourings can often be used to detect non-isomorphism of graphs. As it turns out, the k-dimensional Weisfeiler–Leman algorithm (k-WL) is just as expressive as the logic C_{k+1}, the $(k+1)$-variable fragment of first-order logic enriched with counting quantifiers: it computes distinct colourings on two input graphs if and only if there is a distinguishing

formula in C_{k+1} for them. Moreover, the number of iterations of the algorithm needed to obtain distinct colourings corresponds to the quantifier depth of a distinguishing formula. By this correspondence, from a perspective of descriptive complexity theory, both the dimension of the algorithm and the number of iterations that it needs to produce an output are parameters worth being studied.

Concerning the dimension of the algorithm that is needed to distinguish two graphs, over the past years, many new insights have been obtained, also exploiting the link to counting logics. For example, forests can be identified with 1-WL, interval graphs with 2-WL [11], and planar graphs with 3-WL [36]. Also for many other natural graph classes, bounds on the necessary dimension to tackle the isomorphism problem in the class are known [19, 21, 16, 34].

Concerning the number of iterations, much less is known. Fürer proved a linear lower bound on the number of iterations of k-WL [13], which was improved to $n^{\Omega(k/\log k)}$ [3] and recently to $n^{\Omega(k)}$ [26] on k-ary relational structures. As to upper bounds, for $k = 2$, first progress over the trivial upper bound of $\Theta(n^k)$ has been made in [35], and the best known upper bound is $O(n \log n)$ on graphs of order n [41]. This has been generalised to a bound of $O(n^{k-1} \log n)$ for all $k \geq 2$ in [26]. The number of iterations is crucial for the parallelisability of the algorithm. For $\ell \geq \log n$, it holds that ℓ iterations of k-WL can be simulated in $\mathcal{O}(\ell)$ steps on a PRAM with $O(n^k)$ processors. This implies that if k-WL distinguishes all pairs of graphs of order n in a class $\mathcal{C}$ in $\mathcal{O}(\log n)$ iterations, then deciding isomorphism for $\mathcal{C}$ is in the complexity class TC^1. This is the case for all graph classes of bounded treewidth and all maps [23] as well as all planar graphs [20].

In this paper, we extend the result to all classes of interval graphs and, as a by-product, we obtain the same for chordal claw-free graphs.

Our results: We study the expressive power of the Weisfeiler–Leman algorithm when restricted to a logarithmic number of iterations. This restriction was first introduced as a means of showing that the graph isomorphism problem for graphs of bounded treewidth is in TC^1 [23]. In fact, we transcend to the logical perspective on the algorithm and prove that for every property on relational structures that is definable in the logic $\mathsf{LREC}_=$, there is a number $k \in \mathbb{N}$ such that the logic C_k expresses the property via a family of formulae of logarithmic quantifier depth, which is equivalent to $(k-1)$-WL detecting the property in a logarithmic number of iterations. That is, intuitively speaking, we can simulate logspace recursion with a logarithmic number of iterations of a suitable dimension of the Weisfeiler–Leman algorithm or, equivalently, with a logarithmic quantifier depth in C_k.

The formal statement of the result is as follows.

Theorem I.1 *For every vocabulary τ and every $\mathsf{LREC}_=[\tau]$-formula $\varphi(\bar{x}, \bar{\kappa})$, there is a constant $k \in \mathbb{N}$ such that for every $n \in \mathbb{N}$, there is a family of $C_k^{\mathcal{O}(\log n)}$-formulae $\left(\psi_{\bar{m}}(\bar{x})\right)_{\bar{m} \in [n]^{|\bar{\kappa}|}}$ such that for all τ-structures $\mathcal{A}$ of size $|\mathcal{A}| \leq n$, all $\bar{v} \in \left(V(\mathcal{A})\right)^{|\bar{x}|}$, and all $\bar{m} \in [|\mathcal{A}|]^{|\bar{\kappa}|}$, it holds that*

$$\mathcal{A} \models \varphi(\bar{v}, \bar{m}) \iff \mathcal{A} \models \psi_{\bar{m}}(\bar{v}).$$

In the proof, we restructure the recursive computation of the $\mathsf{LREC}_=$-operator to obtain a computation tree of logarithmic height and small bag overlap. The construction of the desired formulae is then similar to the approach presented in [20]: we build the formula from bottom to top along the tree decomposition, resulting in logarithmic quantifier depth. Here, we need to take care that the number of variables really stays constant.

As an example for the usefulness of the result, we then apply the result to the class of all interval graphs, which is the class of graphs which initially motivated us to study the power of the logarithmic Weisfeiler–Leman algorithm. The class of interval graphs is relevant in many application areas, for example in biology [47] and in operations research [6], and many usually computationally hard problems are known to be tractable on them [28, 32]. Köbler et al. [37] gave a LOGSPACE-algorithm for the isomorphism problem on interval graphs. Their result, however, is purely algorithmic and does not translate to results in terms of logics, making it incomparable to our theorem. Since by [24], for every interval graph, there is an LREC$_=$-formula that identifies it, we can deduce that a constant number of variables suffices to identify every interval graph with a C-sentence of logarithmic quantifier depth.

Theorem I.2 *There is a $k \in \mathbb{N}$ such that for every $n \in \mathbb{N}$ and interval graph G of order n, there is a formula $\Phi^G \in C_k^{\mathcal{O}(\log n)}$ that describes G up to isomorphism.*

As a by-product, using [27], we obtain an analogous statement for the class of chordal claw-free graphs.

Theorem I.3 *There is a $k \in \mathbb{N}$ such that for every chordal claw-free graph G of order n, there is a formula $\Phi^G \in C_k^{\mathcal{O}(\log n)}$ that describes G up to isomorphism.*

Afterwards, we analyse interval graphs in more detail and sketch a second, direct proof to show that isomorphism types of those graphs are definable in C with a constant number of variables and logarithmic quantifier depth. The proof avoids translating LREC$_=$-formulae and proceeds straight via a decomposition of the graph.

II. Preliminaries

We denote a tuple of elements $(x_1, \ldots, x_k)$ as $\bar{x}$. Two tuples $\bar{x} = (x_1, \ldots, x_k)$, $\bar{y} = (y_1, \ldots, y_\ell)$ are said to be *compatible* if $k = \ell$ and hence $|\bar{x}| = |\bar{y}|$. We refer to the *i*th position of a tuple $\bar{x}$ as $\bar{x}_i$. For all $k, \ell \in \mathbb{N}$, we define $[k, \ell] := \{i \in \mathbb{N} \mid k \leq i \leq \ell\}$ and $[k] := [1, k]$.

A. Structures

A *vocabulary* is a non-empty finite set of relation symbols. Each symbol $R \in \tau$ has a fixed *arity* $a_R \in \mathbb{N}$. A *τ-structure* $\mathcal{A}$ consists of a *domain* — a non-empty, finite set $V(A)$ — and, for each $R \in \tau$, a relation $R(\mathcal{A}) \subseteq V(\mathcal{A})^{a_R}$. By the *order* of a structure, we refer to its cardinality $|\mathcal{A}|$. We may write $a_1 \ldots a_k \in R(\mathcal{A})$ instead of $(a_1, \ldots, a_k) \in R(\mathcal{A})$. An *isomorphism* between τ-structures $\mathcal{A}$ and $\mathcal{B}$ is a relation-preserving bijection $\mu \colon V(\mathcal{A}) \to V(\mathcal{B})$, i.e., for all k-ary $R \in \tau$, and all $a_1, \ldots, a_k \in V(\mathcal{A})$, it must hold that $(a_1, \ldots, a_k) \in R(\mathcal{A}) \iff (\mu(a_1), \ldots, \mu(a_k)) \in R(\mathcal{B})$. We then call $\mathcal{A}$ and $\mathcal{B}$ *isomorphic*, denoted as $\mu \colon \mathcal{A} \cong \mathcal{B}$. We may omit μ if the particular mapping does not interest us.

For a τ-structure $\mathcal{A}$, we define the *two-sorted structure* $\mathcal{A}^+$ extending $\mathcal{A}$ as

$$\mathcal{A}^+ := \left(V(\mathcal{A}), \{R(\mathcal{A})\}_{R \in \tau}, N(\mathcal{A}), \leq, S, \min, \max \right),$$

where $N(\mathcal{A}) := \{0, \ldots, |\mathcal{A}|\}$, $\leq$ is the corresponding linear order, min, max are unary singleton relations defining the minimum and maximum of $\leq$, and S is the binary successor relation. Every *domain variable* then ranges over the universe $V(\mathcal{A})$ and is from the set $x_1, x_2, \ldots$, whereas every *number variable* ranges on $N(\mathcal{A})$ and is from the set $\iota_1, \iota_2 \ldots$. We may deviate from this convention and use the symbols $x, y, z \ldots$ for domain, resp.

4

$\iota, \kappa, \lambda \ldots$ for number variables. However, these should be understood to be placeholders for values from $x_1, \ldots$ resp. $\iota_1, \ldots$ used in an effort towards enhanced readability. To represent elements of $V(\mathcal{A})$ resp. $N(\mathcal{A})$, we employ symbols from $u, v, w, \ldots$ resp. $i, j, p, q, \ldots$.

Two two-sorted structures are isomorphic if their underlying τ-structures are isomorphic.

Although being limited to an initial segment of the natural numbers, we can represent larger numbers through tuples $\bar{i} \in N(\mathcal{A})^k$, which are then interpreted as base-$(|\mathcal{A}| + 1)$ numbers as

$$\langle \bar{i} \rangle_{\mathcal{A}} := \sum_{j=1}^{k} \bar{i}_j \cdot (|\mathcal{A}| + 1)^{j-1}.$$

We define an *interpretation* (or *assignment*) to be a mapping α assigning, to each domain variable x_i a value in $V(\mathcal{A})$ and to each number variable ι_i a value in $N(\mathcal{A})$. Since we consider interpretations only together with concrete formulae, it is sufficient if every variable occurring in the formula is assigned a value. Given domain resp. number variable tuples $\bar{x}$ and $\bar{\iota}$ as well as compatible domain resp. number tuples $\bar{v} \in V(\mathcal{A})^k$ and $\bar{p} \in N(\mathcal{A})^{\ell}$, we write $\alpha[\bar{v}/\bar{x}, \bar{p}/\bar{\iota}]$ to mean the assignment α modified to the effect that, for all $i \in [k]$ and $j \in [\ell]$, $\bar{x}_i$ is assigned the value $\bar{v}_i$ and $\bar{\iota}_j$ is assigned the value $\bar{p}_j$.

We call a variable *bound* if it occurs within the scope of a corresponding quantifier. Otherwise, we call it *free*. In particular, given a formula φ, we write $\varphi(x_1, \ldots, x_k)$ to express that free$(\varphi) \subseteq \{x_1, \ldots, x_k\}$ are distinct and that they are those variables that may occur free within φ. Those formulae in which all variables occur bound are *sentences*.

B. Graphs

A *(directed) graph* is an $\{E\}$-structure $G := (V(G), E(G))$ over a domain of *vertices* $V(G)$ and a binary *edge* relation $E(G)$. The *order* of G is $|G| := |V(G)|$.

A graph is *undirected* if $E(G)$ is symmetric and irreflexive, in which case we consider the elements $vw \in E(G)$ to be unordered. For such G and $v \in V(G)$, we denote by $N_G(v) := \{w \in V(G) \mid vw \in E(G)\}$ resp. $N_G[v] := N_G(v) \cup \{v\}$ the *open* resp. *closed* neighbourhood of v in G. Letting $W \subseteq V(G)$, we define the *subgraph induced by W in G* as $G[W] := (W, \{vw \in E(G) \mid v, w \in W\})$.

A directed graph is *acyclic* if there is no $k \in \mathbb{N}$ for which there exists a sequence of distinct edges $(v_1, v_2), (v_2, v_3) \ldots, (v_{k-1}, v_k)$ with $v_i \neq v_j$ for distinct $i, j \in [k-1]$ such that $v_1 = v_k$. We refer to *directed acyclic graphs* as *DAGs*. The *height* of a DAG is the length of a longest path in it.

Now, let G be a directed graph of order $n := |G|$. For every $v \in V(G)$, we let $N_G^+(v) := \{w \in V(G) \mid vw \in E(G)\}$ and $N_G^-(v) := \{u \in V(G) \mid uv \in E(G)\}$ be the sets of *out-neighbours* and *in-neighbours* of v, and $\deg_G^+(v) := |N_G^+(v)|$ and $\deg_G^-(v) := |N_G^-(v)|$ be the *out-degree* and the *in-degree* of v.

Let $\trianglelefteq_G$ denote the reflexive transitive closure of $E(G)$. In these and similar notations, we omit the subscript $_G$ if the graph G is clear from the context. We call a node of out-degree 0 a *leaf* of G. We call G *rooted* if $\trianglelefteq$ has a unique minimal element r that we call the *root* of G. Note that every $v \in V(G)$ is reachable from r.

C. Logics

This section presumes familiarity with first-order logic (FO) and standard model-theoretic notation, suitable overviews for which can be found in [9, 40]. Given a formula φ, a

structure $\mathcal{A}$ and an assignment α, we write $(\mathcal{A}, \alpha) \models \varphi$ to express that the formula φ is *true* given $\mathcal{A}$ and having assigned the variables as in α. In particular, we use the shorthands $\top := \forall x \, (x = x)$ and $\bot := \neg\top$.

Let us introduce an extension of FO, *first-order logic with counting* (FO+C), which operates on the two-sorted structures described above (see also [17]). It extends FO over a second domain and adds two quantifiers, which we now define recursively. In the remainder of the section, let τ be an arbitrary vocabulary. Further, let $\mathcal{A}^+$ be a two-sorted τ-structure, let φ be an FO+C$[\tau]$-formula, α an interpretation and ι be a number variable. Then $\exists\iota \, \varphi$ is an FO+C$[\tau]$-formula, and it is satisfied by $(\mathcal{A}^+, \alpha)$ iff

$$\{i \in N(\mathcal{A}) \mid (\mathcal{A}^+, \alpha[i/\iota]) \models \varphi\} \neq \varnothing$$

holds.

FO+C also adds the #-quantifier, evaluating to some natural number. Let x be a domain variable and κ an additional number variable. Then $\#x \, \varphi = \kappa$ and $\#\iota \, \varphi = \kappa$ are FO+C$[\tau]$-formulae and

$$\mathcal{A}^+ \models \#x \, \varphi = \kappa$$
$$\iff \big| \{v \in V(\mathcal{A}) \mid (\mathcal{A}^+, \alpha[v/x]) \models \varphi\} \big| = \alpha(\kappa)$$

resp.

$$\mathcal{A}^+ \models \#\iota \, \varphi = \kappa$$
$$\iff \big| \{i \in N(\mathcal{A}) \mid (\mathcal{A}^+, \alpha[i/\iota]) \models \varphi\} \big| = \alpha(\kappa).$$

Given a formula $\varphi(\bar{x}, \bar{\iota})$, we define the set of domain/number tuples satisfying φ as

$$\varphi[\mathcal{A}^+, \alpha; \bar{x}, \bar{\iota}] :=$$
$$\left\{ \bar{v}\bar{\iota} \in V(\mathcal{A})^{|\bar{x}|} \times N(\mathcal{A})^{|\bar{\iota}|} \mid (\mathcal{A}^+, \alpha[\bar{v}/\bar{x}, \bar{\iota}/\bar{\iota}]) \models \varphi(\bar{x}, \bar{\iota}) \right\}.$$

We now turn towards describing LREC and LREC$_=$, two extensions of FO+C which were first introduced in [24] in the quest for a logic capturing LOGSPACE. The set of all LREC$[\tau]$-formulae is obtained by extending the syntax of FO+C$[\tau]$ by the following rule. Let $\bar{x}, \bar{y}_1, \bar{y}_2$ be compatible domain variable k-tuples, and $\bar{\iota}, \bar{\kappa}$ be non-empty number variable tuples. Then, if $\varphi_{\mathrm{E}}, \varphi_{\mathrm{C}}$ are LREC$[\tau]$-formulae,

$$\varphi := [\mathsf{lrec}_{\bar{y}_1, \bar{y}_2, \bar{\iota}} \varphi_{\mathrm{E}}, \varphi_{\mathrm{C}}](\bar{x}, \bar{\kappa})$$

is an LREC$[\tau]$-formula with $\mathrm{free}(\varphi) := (\mathrm{free}(\varphi_{\mathrm{E}}) \setminus (\bar{y}_1 \cup \bar{y}_2)) \cup (\mathrm{free}(\varphi_{\mathrm{C}}) \setminus (\bar{y}_1 \cup \bar{\iota})) \cup \bar{x} \cup \bar{\kappa}$. Given a two-sorted τ-structure $\mathcal{A}^+$ and an assignment α, the formula φ recursively defines a relation $X \subseteq V(\mathcal{A})^k \times \mathbb{N}$ such that

$$(\mathcal{A}^+, \alpha) \models \varphi \iff (\alpha(\bar{x}), \langle \alpha(\bar{\kappa}) \rangle_{\mathcal{A}}) \in X.$$

We now describe how X is obtained. Initially, define a graph $\mathsf{G} := (\mathsf{V}, \mathsf{E})$ with $\mathsf{V} := V(\mathcal{A})^k$ and $\mathsf{E} := \varphi_{\mathrm{E}}[\mathcal{A}, \alpha; \bar{y}_1, \bar{y}_2])$. That is, G is a directed graph on the k-tuples of $V(\mathcal{A})$, and the edges are precisely those pairs $(\bar{v}_1, \bar{v}_2)$ with $\bar{v}_1, \bar{v}_2 \in V(\mathcal{A})^k$ such that $(\mathcal{A}, \alpha[\bar{v}_1/\bar{y}_1, \bar{v}_2/\bar{y}_2]) \models \varphi_{\mathrm{E}}(\bar{y}_1, \bar{y}_2)$. Over G, the formula φ_{C} then defines a vertex labelling

$$\mathsf{C}(\bar{v}_1) := \left\{ \langle \bar{\iota} \rangle_{\mathcal{A}} \mid \bar{\iota} \in \varphi_{\mathrm{C}}[\mathcal{A}^+, \alpha[\bar{v}_1/\bar{y}_1]; \bar{\iota}] \right\}.$$

A tuple $(\bar{v}, \bar{\imath})$ with $\bar{v} \in V(\mathcal{A})^k$ and a "resource term" $\bar{\imath} \in N(\mathcal{A})^{|\iota|}$ is contained in X if $\langle i \rangle_{\mathcal{A}} > 0$ (that is, there are still resources left) and

$$\left| \left\{ \bar{w} \in N_{\mathsf{G}}^+(\bar{v}) \,\middle|\, \left(\bar{w}, \left\lfloor \frac{\langle i \rangle_{\mathcal{A}} - 1}{\deg_{\mathsf{G}}^-(\bar{w})} \right\rfloor \right) \in X \right\} \right| \in \mathsf{C}(\bar{v}).$$

The logic $\mathsf{LREC}_=$ replaces the lrec-operator by the $\mathsf{lrec}_=$-operator, which allows the definition of an equivalence relation on the constructed graph. We recall the above defined structure, assignment, and variable tuples. Then, letting $\varphi_=, \varphi_{\mathrm{E}}, \varphi_{\mathrm{C}}$, be $\mathsf{LREC}_=[\tau]$-formulae, we obtain a new $\mathsf{LREC}_=[\tau]$-formula

$$\varphi := [\mathsf{lrec}_{\bar{y}_1, \bar{y}_2, \bar{\imath}} \varphi_=, \varphi_{\mathrm{E}}, \varphi_{\mathrm{C}}](\bar{x}, \bar{\kappa}).$$

The semantics are a bit more complex. First, we construct a graph G' as before. Then, letting $\sim$ be the equivalence relation defined by $\varphi_=[\mathcal{A}^+, \alpha; \bar{y}_1, \bar{y}_2]$ on V', we define a new graph

$$\mathsf{G} := \Big(\mathsf{V} := \mathsf{V}' / _\sim,$$
$$\mathsf{E} := \Big\{ (\bar{v}_{1/\sim}, \bar{v}_{2/\sim}) \in \mathsf{V}^2 \mid (\bar{v}_1, \bar{v}_2) \in \mathsf{E}' \Big\} \Big)$$

contracting the vertices from V' into their equivalence classes while maintaining the edges. For all $\bar{v}_{1/\sim} \in \mathsf{V}$, we define

$$\mathsf{C}(\bar{v}_{1/\sim}) := \big\{ \langle \bar{\imath} \rangle_{\mathcal{A}} \mid \exists \bar{v}' \in \bar{v}_{1/\sim} : \bar{\imath} \in \varphi_{\mathrm{C}}[\mathcal{A}^+, \alpha[\bar{v}'/\bar{y}_1]; \bar{\imath}] \big\}.$$

The relation X is then defined as previously.

See Fig. 1 for an example of how the relation X is computed. Readers wishing to develop a more in-depth understanding of LREC and $\mathsf{LREC}_=$ may want to look into [24], which also includes concrete examples of properties that can be expressed in these logics.

We now move towards defining the logic C, *first-order logic with counting quantifiers*, as the syntactical extension of FO by C-*quantifiers* of the form $\exists^{\geq n} x\, \varphi(x)$ (*there are at least n elements x satisfying $\varphi(x)$*) for all $n \in \mathbb{N}$, which immediately yields the related quantifiers $\exists^{\leq n} x\, \varphi(x) \equiv \exists x\, \varphi(x) \wedge \neg \exists^{\geq n+1} x\, \varphi(x)$ and $\exists^{=n} x\, \varphi(x) \equiv \exists^{\leq n} x\, \varphi(x) \wedge \exists^{\geq n} x\, \varphi(x)$.

Before continuing, it is worth noting that C and $\mathsf{FO+C}$ are two distinct, separate logics, which happen to be similarly named. We attempt to shortly clarify their differences to preempt any confusion. The logic C is only a syntactical extension of FO on relational structures, whereas $\mathsf{FO+C}$ is defined on two-sorted structures and has some access to quantification over the natural numbers. For example, whether a graph is regular or not can be expressed by the $\mathsf{FO+C}$-formula

$$\varphi_{\mathrm{regular}}^{\mathsf{FO+C}} := \exists \iota \left[\forall x \# y\, E(x, y) = \iota \right],$$

which can be understood as "there exists a number ι such that every vertex x has exactly ι neighbours". On the other hand, as a consequence of the hardcoded aspect of numbers in C-quantifiers, any C-formula characterising regularity can only do so for graphs of bounded size:

$$\varphi_{\mathrm{regular}(n)}^{\mathsf{C}} := \bigvee_{0 \leq i \leq n} \forall x \exists^{=i} y\, E(x, y).$$

However, on graphs of fixed size, every $\mathsf{FO+C}$-formula can be simulated by (a family of) C-formulae [9, Proposition 8.4.18], which is foundational for our proof of Theorem I.1.
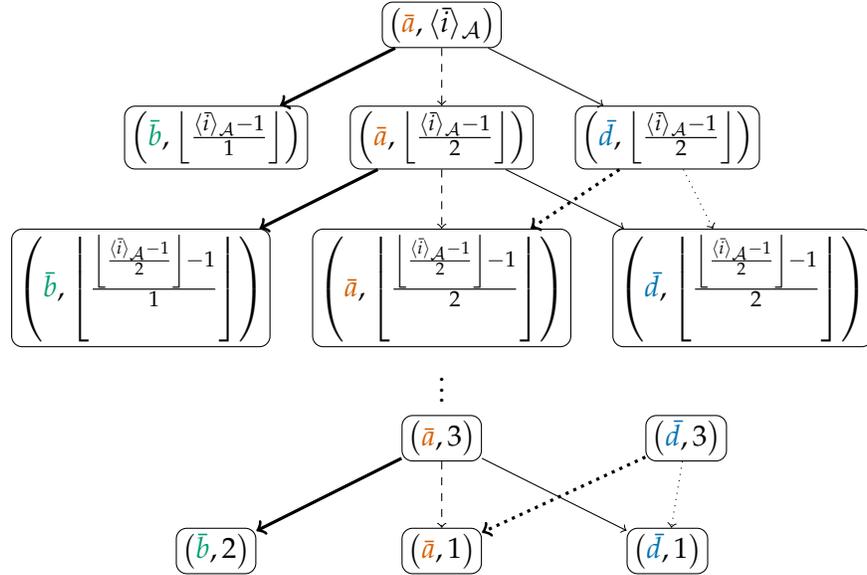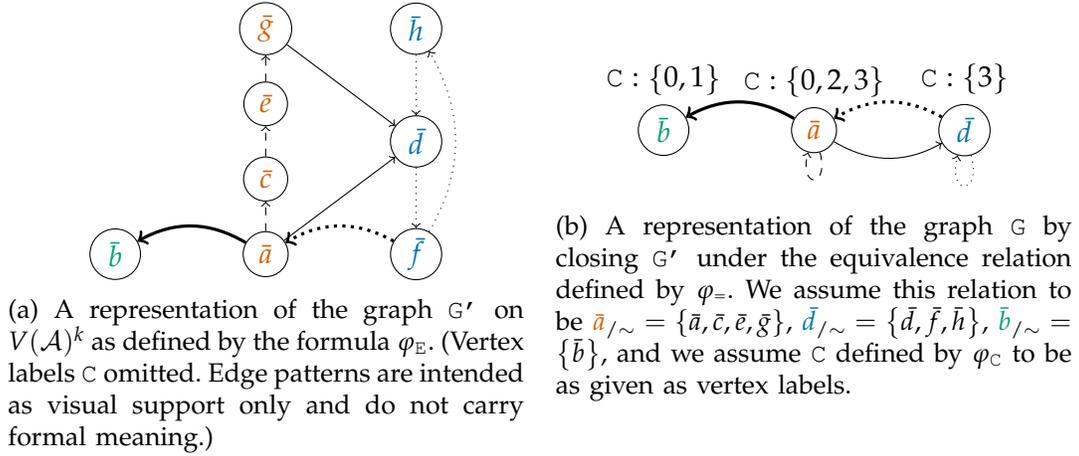
7

(a) A representation of the graph G′ on $V(\mathcal{A})^k$ as defined by the formula φ_{E}. (Vertex labels C omitted. Edge patterns are intended as visual support only and do not carry formal meaning.)

C : $\{0,1\}$ C : $\{0,2,3\}$ C : $\{3\}$

$\bar{b}$ $\bar{a}$ $\bar{d}$

(b) A representation of the graph G by closing G′ under the equivalence relation defined by $\varphi_{=}$. We assume this relation to be $\bar{a}_{/\sim} = \{\bar{a},\bar{c},\bar{e},\bar{g}\}$, $\bar{d}_{/\sim} = \{\bar{d},\bar{f},\bar{h}\}$, $\bar{b}_{/\sim} = \{\bar{b}\}$, and we assume C defined by φ_{C} to be as given as vertex labels.

$\left(\bar{a}, \langle \bar{\imath} \rangle_{\mathcal{A}}\right)$

$\left(\bar{b}, \left\lfloor \frac{\langle \bar{\imath} \rangle_{\mathcal{A}}-1}{1} \right\rfloor\right)$ $\left(\bar{a}, \left\lfloor \frac{\langle \bar{\imath} \rangle_{\mathcal{A}}-1}{2} \right\rfloor\right)$ $\left(\bar{d}, \left\lfloor \frac{\langle \bar{\imath} \rangle_{\mathcal{A}}-1}{2} \right\rfloor\right)$

$\left(\bar{b}, \left\lfloor \frac{\left\lfloor \frac{\langle \bar{\imath} \rangle_{\mathcal{A}}-1}{2} \right\rfloor -1}{1} \right\rfloor\right)$ $\left(\bar{a}, \left\lfloor \frac{\left\lfloor \frac{\langle \bar{\imath} \rangle_{\mathcal{A}}-1}{2} \right\rfloor -1}{2} \right\rfloor\right)$ $\left(\bar{d}, \left\lfloor \frac{\left\lfloor \frac{\langle \bar{\imath} \rangle_{\mathcal{A}}-1}{2} \right\rfloor -1}{2} \right\rfloor\right)$

$(\bar{a},3)$ $(\bar{d},3)$

$(\bar{b},2)$ $(\bar{a},1)$ $(\bar{d},1)$

(c) A visualisation of the DAG resulting from the recursive unfolding of the graph G with respect to the parameters $(\bar{a},\bar{\imath})$. For every recursive step, the resource term (i.e., $\langle \bar{\imath} \rangle_{\mathcal{A}}$, $\left\lfloor \frac{\langle \bar{\imath} \rangle_{\mathcal{A}}-1}{2} \right\rfloor$, etc.) can be understood to be "split" equitably among the in-neighbours of the vertex in G. Together with the requirement that the resource term must be positive, this ensures a logarithmic space bound.

Fig. 1: A visualisation of the computation of the relation X conducted when evaluating an $\mathsf{LREC}_{=}[\tau]$-formula $\varphi := [\mathsf{lrec}_{\bar{y}_1,\bar{y}_2,\bar{\imath}}\varphi_{=}, \varphi_{\mathrm{E}}, \varphi_{\mathrm{C}}](\bar{a},\bar{\imath})$. Let $\mathcal{A}^+$ be a two-sorted τ-structure. We assume the graph G′, defined via the formula φ_{E}, to be as shown in Fig. 1a. We then assume the formula $\varphi_{=}$ to yield the graph G and the formula φ_{C} to yield the vertex labelling C as shown in Fig. 1b. Using Fig. 1c, we now describe how the DAG is evaluated and thereby how the relation X is computed, in a bottom-to-top fashion. A leaf, say $(\bar{b},2)$, is part of the relation X if $0 \in \mathrm{C}(\bar{b})$, since the criterion is whether the number of $(\bar{b},2)$'s children which are in X (that is, 0) occurs in $\mathrm{C}(\bar{b})$. Therefore, here it holds that $(\bar{b},2), (\bar{a},1) \in X$, but $(\bar{d},1) \notin X$. Then their predecessor $(\bar{a},3)$ is in X since $\left|\{(\bar{b},2),(\bar{a},1)\}\right| = 2 \in \mathrm{C}(\bar{a})$, whereas $(\bar{d},3) \notin X$ because $\left|\{(\bar{a},1)\}\right| = 1 \notin \mathrm{C}(\bar{d})$. This then continues up to the root $(\bar{a}, \langle \bar{\imath} \rangle_{\mathcal{A}})$, which is contained in X if the number q of its children in X occurs in $\mathrm{C}(\bar{a})$. Finally, φ is satisfied iff the root is in X.

In itself, C is exactly as expressive as FO considering that

$$\exists^{\geq n}x\, \varphi(x) \equiv \exists x_1 \ldots \exists x_n \left[\bigvee_{\substack{i,j\in[n] \\ i\neq j}} x_i \neq x_j \wedge \bigwedge_{i\in[n]} \varphi(x_i) \right].$$

However, we are interested in the finite-variable fragments of C, denoted as C_k for $k \in \mathbb{N}$, where C_k contains exactly those formulae from C which only use variables from $\{x_1, \ldots, x_k\}$.

We inductively define the *quantifier depth* $\mathsf{qd}(\varphi)$ of a formula $\varphi \in \mathsf{C}$ as

$$\mathsf{qd}(\varphi) := \begin{cases} 0 & \text{if } \varphi \text{ is atomic,} \\ \mathsf{qd}(\psi) & \text{if } \varphi = \neg\psi, \\ \max(\mathsf{qd}(\psi_1), \mathsf{qd}(\psi_2)) & \text{if } \varphi = \psi_1 \vee \psi_2, \\ \mathsf{qd}(\varphi) + 1 & \text{if } \varphi = \exists^{(\geq n)}x\, \psi. \end{cases}$$

Now, we can further restrict C_k to C_k^r, which we define as the subset of formulae φ of C_k with $\mathsf{qd}(\varphi) \leq r$.

Example II.1 *The C_2^3-sentence*

$$\exists x\, \exists^{=3}y \left[E(x,y) \wedge \exists^{=4}x\, E(y,x) \right]$$

expresses that a graph satisfying it must admit a vertex with exactly three neighbours, each in turn admitting exactly 4 neighbours.

In the following, we will mostly be using the asymptotic notation $\mathsf{C}_k^{\mathcal{O}(\log n)}$, which should be understood as follows. Let τ be a vocabulary, and $\mathfrak{T}$ the class of all τ-structures. For any $m \in \mathbb{N}$ and τ-structure $\mathcal{B} \in \mathfrak{T}$ we denote, for all $\bar{b} \in V(\mathcal{B})^m$, the pair of a structure and an m-tuple of its elements as $(\mathcal{B}, \bar{b})$. Let $\mathfrak{B}$ be a class containing a subset of those pairs and suppose $\mathfrak{B}$ to be containing exactly those elements having some property P. We then say that P *can be expressed in* $\mathsf{C}_k^{\mathcal{O}(\log n)}$ if there exists a $k \in \mathbb{N}$ and a function $f(n) \in \mathcal{O}(\log n)$ such that for all $n \in \mathbb{N}$, there exists a formula $\varphi_{\mathfrak{B}}^{(n)}(\bar{v}) \in \mathsf{C}_k^{f(n)}[\tau]$ satisfying for all $(\mathcal{B}, \bar{b})$ with $|\mathcal{B}| = n$ that $\mathcal{B} \models \varphi_{\mathfrak{B}}^{(n)}(\bar{b}) \iff (\mathcal{B}, \bar{b}) \in \mathfrak{B}$.

D. The Weisfeiler–Leman Algorithm

The *(k-dimensional) Weisfeiler–Leman Algorithm* (*k-WL*, WL_k) is a combinatorial algorithm that iteratively computes a colouring $c_k^*: V(G)^k \to C$ on the k-vertex tuples of a graph. Applied to two graphs, it may be used to decide whether these are isomorphic or not. For our purposes, it suffices to know that the algorithm is initialised by colouring all k-tuples of vertices by their *atomic type*, which contains all information regarding connectivity and equality of the elements of such a tuple. The colouring is then iteratively *refined* by computing, for each k-tuple of vertices, a new colour based on the colours of the adjacent (that is, differing in one position) k-tuples. The final output is the first colouring c_k^i that partitions the k-tuples into the same colour classes as the previous iteration. More details can be found, for example, in [25, 33].

We denote by WL_k^r the restriction of WL_k that terminates after the first r refinement rounds, i.e., with c_k^r; WL_k^r then *distinguishes* two graphs G, H if there exists a colour c

such that after r rounds, G and H admit a different number of vertex k-tuples of that colour. WL_k^r then *identifies* G if it distinguishes it from all non-isomorphic graphs H. The connection between the Weisfeiler–Leman algorithm and the finite-variable fragment of counting logic is as follows.

Lemma II.2 ([23, 5]) *Let $k \in \mathbb{N}$. For graphs G and H of the same order and all $r \in \mathbb{N}$, the following statements are equivalent:*
- WL_k^r *distinguishes G and H.*
- *There is a C_{k+1}^r-sentence φ such that $G \models \varphi \iff H \not\models \varphi$.*

For all $k \in \mathbb{N}$, we define the *logarithmic Weisfeiler–Leman algorithm*, denoted as $\mathsf{WL}_k^{\mathcal{O}(\log n)}$, analogous to the asymptotic notation in $\mathsf{C}_k^{\mathcal{O}(\log n)}$. In particular, Lemma II.2 implies that if there is a $k \geq 2$ and a $\mathsf{C}_k^{\mathcal{O}(\log n)}$-sentence identifying a graph G, then $\mathsf{WL}_{k-1}^{\mathcal{O}(\log n)}$ identifies G.

III. The Treelike Decompositions

This section serves us to compute treelike decompositions of logarithmic height, along which we build our C-formulae of logarithmic quantifier depth in Section IV. We start off with a DAG since, as we will see in Section IV, the computation of the relation X from the definition of LREC results in such a graph. To account for X in C, we can decompose the graph as we are about to describe next. Crucially, the trees underlying the decomposition have logarithmic depth. We will then use the decomposition to construct $\mathsf{C}_k^{\mathcal{O}(\log n)}$-formulae.

Throughout this section, we assume that G is a rooted DAG and that r is the root of G.

The *tree unfolding* of G is the tree T_G whose vertices are paths $\bar{v} = (v_0, \ldots, v_k)$ in G with $v_0 = r$, and where $\bar{w} = (w_0, \ldots, w_\ell)$ is a child of $\bar{v}$ if $\ell = k + 1$ and $w_i = v_i$ for $i \in [0, k]$. For every $v \in V(G)$, let $\mathcal{P}_G(v)$ be the set of all paths $(v_0, \ldots, v_k)$ in G with $v_0 = r$ and $v_k = v$. Note that all $\bar{v} \in \mathcal{P}(v)$ are vertices of T_G. We call them the *copies* of v in T_G.

We define the *weight* of a vertex v in G to be

$$\mathrm{wt}_G(v) := |\mathcal{P}(v)|,$$

of which we omit the subscript if it is clear from context. We define the *aggregate weight* of G to be $\mathrm{awt}(G) := \sum_{v \in V(G)} \mathrm{wt}(v)$. Observe that

$$\mathrm{awt}(G) = |T_G|.$$

We define the *multiplicity* of a vertex v to be

$$\mathrm{mul}_G(v) := \max \left\{ \prod_{i=1}^{k} \deg^-(v_i) \;\middle|\; (v_0, \ldots, v_k) \in \mathcal{P}(v) \right\},$$

where the empty product is 1 (thus $\mathrm{mul}_G(r) = 1$). We may again omit the subscript. Further, we let the *aggregate multiplicity* of G be defined as $\mathrm{amul}(G) := \sum_{v \in V(G)} \mathrm{mul}(v)$.

Lemma III.1 *For all $v \in V(G)$, we have $\mathrm{wt}(v) \leq \mathrm{mul}(v)$.*

Proof: We prove the result by induction on the distance between v and the root r. For the root r, we have $\text{wt}(r) = \text{mul}(r) = 1$. So let v be a node with in-neighbours $u_1, \ldots, u_k$. Then $\deg^-(v) = k$ and

$$\text{wt}(v) = \sum_{i=1}^{k} \text{wt}(u_i) \leq \sum_{i=1}^{k} \text{mul}(u_i) \leq k \max_{i \in [k]} \text{mul}(u_i) = \text{mul}(v),$$

where the first inequality holds by the induction hypothesis. ∎

For $m \in \mathbb{N}$, we say that G has the *m-path property* if $\text{mul}(v) \leq m$ for all $v \in V(G)$. The m-path property allows us to control the size of a tree unfolding.

Corollary III.2 *If G has the m-path property, then $\text{awt}(G) \leq \text{amul}(G) \leq m \cdot |G|$.*

Let $v \in V(G)$ and $W \subseteq V(G)$ be such that $v \trianglelefteq_G w$ for all $w \in W$. Then, we define G_v^W to be the induced subgraph of G with vertex set

$$\big\{ u \in V(G) \mid \text{there is path } (v_0, \ldots, v_\ell) \text{ in } G \text{ with } v_0 = v \text{ and}$$
$$v_\ell = u \text{ and } \{v_0, v_1, \ldots, v_{\ell-1}\} \cap W = \varnothing \big\}.$$

In the definition, we stipulate that for $\ell = 0$, $\{v_0, v_1, \ldots, v_{\ell-1}\}$ is the empty set. Thus, if $v \in W$, then $G_v^W = G[\{v\}]$. Note that G_v^W is a rooted DAG with root v. If $W = \varnothing$, we write G_v instead of G_v^W, and if v is the root, we write G^W instead of G_r^W. If $W = \{w_1, \ldots, w_k\}$, we also write $G_v^{w_1, \ldots, w_k}$ instead of G_v^W.

Lemma III.3 *For all $v \in V(G)$, it holds that $\text{awt}(G_v) + \text{awt}(G^v) \leq \text{awt}(G) + 1$.*

Proof: Let $X := V(G_v) \setminus V(G^v)$, $Y := V(G^v) \setminus V(G_v)$, and $Z := \big(V(G_v) \cap V(G^v) \big) \setminus \{v\}$. We have

$\text{awt}(G)$	$=$	$\text{wt}_G(v)$	$+$	$\displaystyle\sum_{x \in X} \text{wt}_G(x)$
	$+$	$\displaystyle\sum_{y \in Y} \text{wt}_G(y)$	$+$	$\displaystyle\sum_{z \in Z} \text{wt}_G(z)$
$\text{awt}(G_v)$	$=$	$\text{wt}_{G_v}(v)$	$+$	$\displaystyle\sum_{x \in X} \text{wt}_{G_v}(x)$
			$+$	$\displaystyle\sum_{z \in Z} \text{wt}_{G_v}(z)$
$\text{awt}(G^v)$	$=$	$\text{wt}_{G^v}(v)$		
	$+$	$\displaystyle\sum_{y \in Y} \text{wt}_{G^v}(y)$	$+$	$\displaystyle\sum_{z \in Z} \text{wt}_{G^v}(z).$

We have $\text{wt}_{G_v}(v) = 1$ and $\text{wt}_{G^v}(v) = \text{wt}_G(v)$. Furthermore, for every $z \in Z$, we have $\text{wt}_{G_v}(z) + \text{wt}_{G^v}(z) \leq \text{wt}_G(z)$, because we can partition $\mathcal{P}_G(z)$ into $\mathcal{P}_{G^v}(z)$, consisting of all paths from r to z in G that avoid v, and the set $\mathcal{Q}$ consisting of all paths from r to z in G that contain v. We have $|\mathcal{P}_{G_v}(z)| \leq |\mathcal{Q}|$. With this, the assertion of the lemma follows. ∎

Lemma III.4 *Let $v \in V(G)$ such that v is not a leaf of G. Then there is an $a \in V(G)$ such that $v \trianglelefteq a$ and*

$$\text{awt}(G_v^a) \leq \frac{\text{awt}(G_v)}{2}, \tag{1}$$

$$\text{awt}(G_b) \leq \left\lceil \frac{\text{awt}(G_v)}{2} \right\rceil \qquad \text{for all } b \in N^+(a). \tag{2}$$

11

Proof: Let $m := \text{awt}(G_v)$ and note that $m \geq |G_v| \geq 2$.

Then $\text{awt}(G_v^v) = 1 \leq \frac{m}{2}$ and $\text{awt}(G_v^w) = \text{awt}(G_v) = m > \frac{m}{2}$ for every leaf w of G_v. Hence, there is an $a \in V(G_v)$ such that $\text{awt}(G_v^a) \leq \frac{m}{2}$ and $\text{awt}(G_v^b) > \frac{m}{2}$ for every $b \in N^+(a)$.

This a satisfies (1); to see that it satisfies (2), let $b \in N^+(a)$. Then $\text{awt}(G_v^b) > \frac{m}{2}$ and thus $\text{awt}(G_v^b) \geq \lfloor \frac{m}{2} \rfloor + 1$. By Lemma III.3, we have $\text{awt}(G_b) + \text{awt}(G_v^b) \leq m + 1$ and thus

$$\text{awt}(G_b) \leq m + 1 - \left(\left\lfloor \frac{m}{2} \right\rfloor + 1 \right) = \left\lceil \frac{m}{2} \right\rceil.$$

∎

Lemma III.5 *Let $v, w \in V(G)$ be such that $v \lhd w$. Then there is an $a \in V(G)$ such that $v \unlhd a \lhd w$ and*

$$\text{awt}(G_v^a) \leq \frac{\text{awt}(G_v^w)}{2}, \tag{3}$$

$$\text{awt}(G_b^w) \leq \left\lceil \frac{\text{awt}(G_v^w)}{2} \right\rceil \qquad \text{for all } b \in N^+(a) \text{ with } b \unlhd w. \tag{4}$$

Proof: Let $m := \text{awt}(G_v^w)$ and note that $m \geq |G_v^w| \geq 2$. Let $\bar{v} = (v_0, \ldots, v_k)$ be a path in G with $v_0 = v$ and $v_k = w$. We have $\text{awt}(G_v^{v_0}) = 1 \leq \frac{m}{2}$ and $\text{awt}(G_v^{v_k}) = \text{awt}(G_v^w) > \frac{m}{2}$. Thus there is a (unique) $i \in \{0, \ldots, k-1\}$ such that $\text{awt}(G_v^{v_i}) \leq \frac{m}{2}$ and $\text{awt}(G_v^{v_{i+1}}) > \frac{m}{2}$. Let $a(\bar{v}) := v_i$, and let a be $\unlhd$-maximal among all $a(\bar{v})$, where $\bar{v}$ ranges over all paths from v to w.

Then (3) is trivially satisfied by all $a(\bar{v})$ and in particular by a.

To prove (4), let $b \in N^+(a)$ such that $b \unlhd w$. As $v \unlhd a$ and $ab \in E(G)$ and $b \unlhd w$, there exists a path $(v_0, \ldots, v_k)$ from v to w such that $a = v_i$ and $b = v_{i+1}$. By the maximality of a, we have $\text{awt}(G_v^b) > \frac{m}{2}$. By Lemma III.3 applied to the graph G_v^w and b, we get $\text{awt}(G_b^w) \leq \lceil \frac{m}{2} \rceil$. ∎

We can now use Lemmas III.4 and III.5 to inductively construct a representation of G by a tree of logarithmic height.

Lemma III.6 *There are a rooted tree T and mappings $v \colon V(T) \to V(G)$, $W \colon V(T) \to 2^{V(G)}$ such that the following conditions are satisfied.*

1) $|W(t)| \leq 1$ for all $t \in V(T)$.
2) $t \in V(T)$ is a leaf of T if and only if $v(t)$ is a leaf of G or $W(t) = \{v(t)\}$.
3) *If $t \in V(T)$ is not a leaf of T and $W(t) = \{w\}$, then $v(t) \lhd w$.*
4) *If $t \in V(T)$ with children $u_1, \ldots, u_k$ for some $k \geq 1$, then*

$$V\left(G_{v(t)}^{W(t)} \right) \setminus \{v(t)\} \subseteq \bigcup_{i=1}^{k} V\left(G_{v(u_i)}^{W(u_i)} \right).$$

5) *The height of T is at most $2\log(\text{awt}(G))$.*

Proof: We define the tree T inductively. We start with a root r_T and let $v(r_T) := r$ and $W(r_T) := \varnothing$.

To extend the tree, let t be a node in T where the children are not yet defined. If $v(t)$ is a leaf of G or $v(t) \in W(t)$, then t is a leaf of T. Now suppose that $v(t)$ is not a leaf of G and $v(t) \notin W(t)$. Let $v := v(t)$ and $W := W(t)$. By induction, we assume $|W| \leq 1$ and $v \lhd_G w$ if $W = \{w\}$.

Case 1: $W = \varnothing$. (We say that t is a node of type 0.)

Then $G_v^W = G_v$. By Lemma III.4, there is an $a \in V(G_v)$ such that $\mathrm{awt}(G_v^a) \leq \frac{\mathrm{awt}(G_v)}{2}$ and $\mathrm{awt}(G_b) \leq \left\lceil \frac{\mathrm{awt}(G_v)}{2} \right\rceil$ for all $b \in N^+(a)$.

We add a child u_a of t with $v(u_a) := v$ and $W(u_a) := \{a\}$. For every $b \in N^+(a)$, we add a child u_b with $v(u_b) := b$ and $W(u_b) := \varnothing$.

Case 2: $W = \{w\}$ for some w. (We say that t is a node of type 1.)

Then $G_v^W = G_v^w$. By Lemma III.5, there is an $a \in V(G_v)$ such that $\mathrm{awt}(G_v^a) \leq \frac{\mathrm{awt}(G_b^w)}{2}$ and $\mathrm{awt}(G_b) \leq \left\lceil \frac{\mathrm{awt}(G_v^w)}{2} \right\rceil$ for all $b \in N^+(a)$ with $b \unlhd w$.

We add a child u_a of t with $v(u_a) := v$ and $W(u_a) := \{a\}$. For every $b \in N^+(a)$ with $b \unlhd w$, we add a child u_b with $v(u_b) := b$ and $W(u_b) := \{w\}$. For every $b \in N^+(a)$ with $b \ntrianglelefteq w$, we add a child u_b with $v(u_b) := b$ and $W(u_b) := \varnothing$.

It is immediate from the construction that T, v, W satisfy Items 1 to 4 of Lemma III.6. We need to prove that they satisfy Item 5. For every $t \in V(T)$, let $A(t) := \mathrm{awt}\left(G_{v(t)}^{W(t)} \right)$. Observe that for all nodes $t \in V(T)$ and all children u of t the following holds:

- $A(u) < A(t)$;
- if t is of type 0, then $A(u) \leq \left\lceil \frac{A(t)}{2} \right\rceil$;
- if $A(u) > \left\lceil \frac{A(t)}{2} \right\rceil$, then t is of type 1 and $W(u) = \varnothing$, so u is of type 0.

This implies that for all grandchildren v of t, it holds that $A(v) \leq \frac{A(t)}{2}$, and as $A(r_T) = \mathrm{awt}(G)$, Item 5 follows. ∎

IV. From LREC$_=$ to $\mathsf{C}_k^{\mathcal{O}(\log n)}$

Let G be a directed graph. A *cardinality condition* for G is a mapping C that associates to each $v \in V(G)$ a set $C(v) \subseteq [0, \deg^+(v)]$. Given a cardinality condition C, we define, analogously to the definition of LREC$_=$, $X = X(G, C) \subseteq V(G) \times \mathbb{N}_{>0}$ to be the inclusionwise smallest set such that for all $v \in V(G)$ and $i \in \mathbb{N}_{>0}$, it holds that $(v, i) \in X$ if and only if

$$\left| \left\{ w \in N^+(v) \;\middle|\; \left(w, \left\lfloor \frac{i-1}{\deg^-(w)} \right\rfloor \right) \in X \right\} \right| \in C(v).$$

For every $n \in \mathbb{N}_{>0}$, we define a vocabulary $\tau^{(n)} := \{E, P_0, \ldots, P_n\}$, where E is a binary relation symbol and the P_i are unary relation symbols. We can represent a tuple (G, C) consisting of a directed graph G of order $|G| \leq n$ and a cardinality condition C for G as a $\tau^{(n)}$-structure $\mathcal{A} = \mathcal{A}(G, C)$ with $V(\mathcal{A}) := V(G)$, $E(\mathcal{A}) := E(G)$, and, for all $i \in [0, n]$,

$$P_i(\mathcal{A}) := \{v \in V(G) \mid i \in C(v)\}.$$

The following theorem enables us to check X-membership via formulae in counting logics with logarithmic quantifier depth.

Theorem IV.1 *There is a $k \in \mathbb{N}_{>0}$ such that for all $n, r \in \mathbb{N}_{>0}$ and $i \in [(n+1)^r]$, there is a $\mathsf{C}_k^{\mathcal{O}(r \log n)}$-formula $\varphi_i^{(n)}(x)$ such that for all directed graphs G of order $|G| \leq n$, all cardinality conditions C for G, and all $v \in V(G)$, it holds that*

$$\mathcal{A}(G, C) \models \varphi_i^{(n)}(v) \iff (v, i) \in X(G, C).$$

Proof: Let $n, r \in \mathbb{N}_{>0}$, and let $i \in [(n+1)^r]$. First, for every directed graph G of order $|G| \leq n$ and every cardinality condition C on G, we are going to describe, for all $v \in V(G)$,
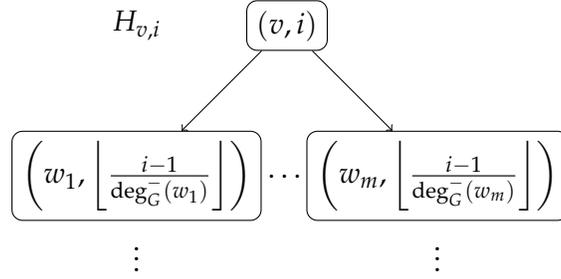
13

Fig. 2: The recursive construction of $H_{v,i}$: in this example, it holds that $N_G^+(v) = \{w_1, \ldots, w_m\}$.

rooted DAGs $H_{v,i}$ that may be used to decide whether (v,i) is contained in $X(G,C)$. Then, Lemma III.6 yields trees $T_{v,i}$ of logarithmic height based on the $H_{v,i}$. We describe how to use those trees to decide whether $(v,i) \in X(G,C)$ holds. At the end of this proof, we recursively construct formulae that check containment in $X(G,C)$ and have a structure that closely follows the structure of the described trees. Since the tree from Lemma III.6 has logarithmic height, the formulae will have a logarithmic quantifier depth.

Let G be a directed graph of order $|G| \leq n$, and let $v \in V(G)$. We inductively define a rooted DAG $H_{v,i}$, see also Fig. 2. We start with the root (v,i). Then, repeatedly, for every vertex $(v',i') \in V(H_{v,i})$ and every neighbour $w \in N_G^+(v')$ where for $j := \left\lfloor \frac{i'-1}{\deg_G^-(w)} \right\rfloor$, it holds that $j \geq 1$, we add a vertex (w,j) to $H_{v,i}$ (unless it already exists) and insert an edge from (v',i') to (w,j).

We could decide "$(v,i) \in X(G,C)$?" as follows. First, we go through all leaves (w,j) in $H_{v,i}$ and mark them as positive if $0 \in C(w)$ and as negative otherwise. Then, for every vertex (w,j) that has only marked children, we mark the vertex as positive if and only if the number of positively marked children is contained in $C(w)$, and we mark it as negative otherwise. Once all vertices have been marked, we have $(v,i) \in X(G,C)$ if and only if (v,i) is marked as positive. Since the height of $H_{v,i}$ may be linear in the size of G, this process might take a linear number of steps. Thus, we use a tree $T_{v,i}$ of logarithmic height instead, which we describe below.

Note that for a node $(w,j) \in V(H_{v,i})$, the graph $H_{w,j}$ is the induced subgraph of $H_{v,i}$ on all nodes below (or equal to) (w,j).

Claim 1 $H_{v,i}$ *has the* $(n+1)^r$-*path property.*

Proof: We prove the equivalent statement $\mathrm{mul}_{H_{v,i}}((w,j)) \leq (|H_{v,i}|+1)^r$ for all $(w,j) \in V(H_{v,i})$.

Let $(v'_0, \ldots, v'_p)$ be a path in $H_{v,i}$ with $v'_0 = (v,i)$ and $v'_p = (w,j)$. Moreover, let $(v_s, \ell_s) := v'_s$ for all $s \in [0,p]$. Then, we have

$$\ell_s = \left\lfloor \frac{\ell_{s-1}-1}{\deg_G^-(v_s)} \right\rfloor \leq \frac{\ell_{s-1}}{\deg_G^-(v_s)} = \frac{i}{\prod_{t=1}^s \deg_G^-(v_t)}$$

for all $s \in [p]$. With $\ell_p = j \geq 1$, it holds that $i \geq \prod_{s=1}^p \deg_G^-(v_s) \geq \prod_{s=1}^p \deg_{H_{v,i}}^-(v_s)$. Thus, $\mathrm{mul}_{H_{v,i}}((w,j)) \leq i \leq (n+1)^r$. Hence, $H_{v,i}$ has the $(n+1)^r$-path property. ⌟
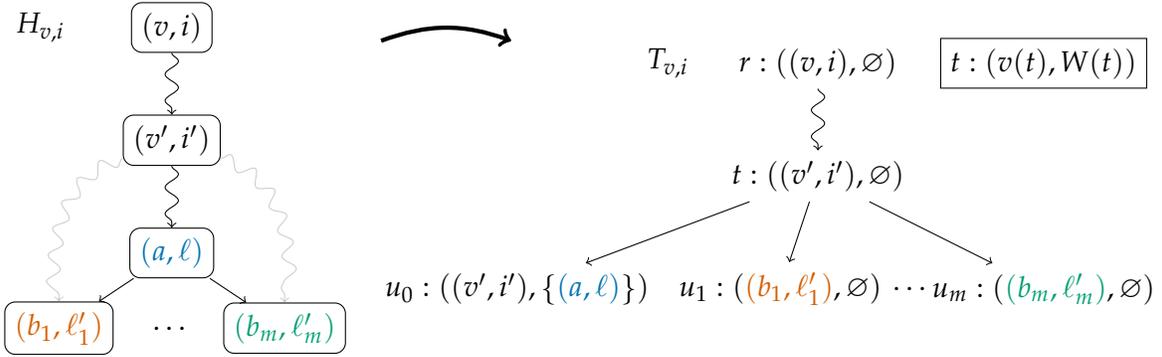
Fig. 3: Letting (v, i) be the vertex for which we want to know whether "$(v, i) \in X(G, C)$?", the figure shows how to recursively obtain $T_{v,i}$ at a vertex $t \in V(T_{v,i})$ of type 0 with $v(t) = (v', i'), W = \varnothing$.

Next, we apply Lemma III.6 to $H_{v,i}$ and obtain (the existence of) a rooted tree $T_{v,i}$ and mappings $v \colon V(T_{v,i}) \to V(H_{v,i})$ and $W \colon V(T_{v,i}) \to 2^{V(H_{v,i})}$. Let $T_{v,i}, v, W$ be as described in the proof of Lemma III.6.

Now, we describe how to decide "$(v, i) \in X(G, C)$?" using $T_{v,i}$. Let $H \coloneqq H_{v,i}$ and $T \coloneqq T_{v,i}$. We start with the root r of T, which, by the construction from Lemma III.6, is a node of type 0 (that is, at node t with $W(t) = \varnothing$) with $v(r) = (v, i)$ and $W(r) = \varnothing$. At every node $t \in V(T)$ of type 0 with $v(t) = (v', i')$ and $W(t) = \varnothing$, our goal is to decide whether $(v', i') \in X(G, C)$ by recursively checking the children of t. At every node $t \in V(T)$ of type 1 with $v(t) = (v', i')$ and $W(t) = \{(w, j)\}$, we are additionally given a number c and our goal is to decide whether $(v', i') \in X(G, C)$ by recursively checking the children of t under the assumption that exactly c of the children of (w, j) in H are contained in $X(G, C)$. In detail, the computations work as follows.

For the following, see also Fig. 3. Consider a node $t \in V(T)$ of type 0 with children u_0 of type 1 and $u_1, \ldots, u_m$ of type 0. Let $v(t) = (v', i')$, $W(t) = \varnothing$, $v(u_0) = (v', i')$, $W(u_0) = \{(a, \ell)\}$, and $v(u_s) = (b_s, \ell'_s)$, $W(u_s) = \varnothing$ for all $s \in [m]$. By the construction of T in Lemma III.6, we have $(v, i) \trianglelefteq_H (v', i') \triangleleft_H (a, \ell) \triangleleft_H (b_s, \ell'_s)$ for all $s \in [m]$, where the (b_s, ℓ'_s) are children of (a, ℓ) in H.

To decide whether "$(v', i') \in X(G, C)$?", we first decide $(b_s, \ell'_s) \in X(G, C)$ recursively for all $s \in [m]$. This is equivalent to running our procedure recursively on the children $u_1, \ldots, u_m$ of t of type 0. Let c be the number of the (b_s, ℓ'_s) contained in $X(G, C)$. We then run our procedure recursively on the child u_0 of type 1 with $v(u_0) = (v', i')$ and $W(u_0) = \{(a, \ell)\}$ to check whether (v', i') is contained in $X(G, C)$ under the assumption that exactly c of the children of (a, ℓ) are contained in $X(G, C)$.

Now consider a node $t \in V(T)$ of type 1 with children u_0 of type 1, $u_1, \ldots, u_m$ of type 1, and $u_{m+1}, \ldots, u_{m+p}$ of type 0 (see Fig. 4). Let $v(t) = (v', i')$, $W(t) = \{(w, j)\}$, $v(u_0) = (v', i')$, $W(u_0) = (a, \ell)$, $v(u_s) = (b_s, \ell'_s)$, $W(u_s) = \{(w, j)\}$ for all $s \in [m]$, and $v(u_s) = (b_s, \ell'_s)$, $W(u_s) = \varnothing$ for all $s \in [m+1, p]$. By the construction of T in Lemma III.6, we have $(v, i) \trianglelefteq_H (v', i') \triangleleft_H (a, \ell) \triangleleft_H (b_s, \ell'_s)$ for all $s \in [m+p]$, where the (b_s, ℓ'_s) are children of (a, ℓ) in H; we have $(b_s, \ell'_s) \trianglelefteq_H (w, j)$ for all $s \in [m]$ and $(b_s, \ell'_s) \ntrianglelefteq_H (w, j)$ for all $s \in [m+1, p]$.

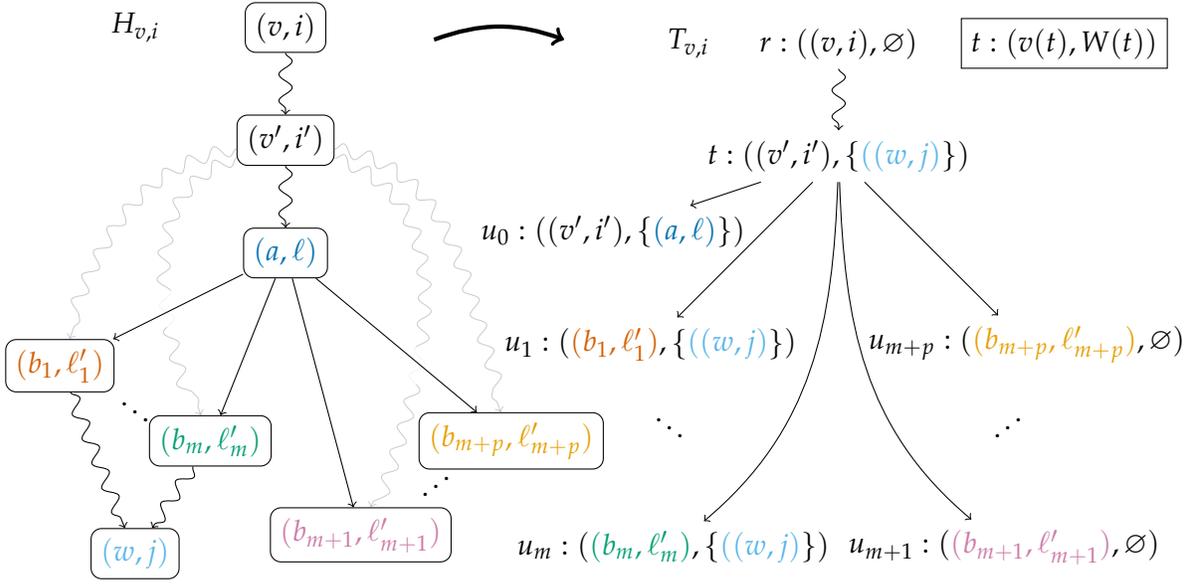To decide whether "$(v', i') \in X(G, C)$?" under the assumption that exactly c children of

15

Fig. 4: Letting (v, i) be the vertex for which we want to know whether "$(v, i) \in X(G, C)$?", the figure shows how to recursively obtain $T_{v,i}$ at a vertex $t \in V(T_{v,i})$ of type 1 with $v(t) = (v', i'), W = \{((w, j)\}.$

(w, j) in H are contained in $X(G, C)$, we recursively run our procedure on the children $u_1, \ldots, u_m$ of type 1 under the above-mentioned assumption and thereby decide $(b_s, \ell'_s) \in X(G, C)$ for $s \in [m]$. Next, we recursively run the procedure on the children $u_{m+1}, \ldots, u_{m+p}$ of type 0 and thereby decide $(b_s, \ell'_s) \in X(G, C)$ for $s \in [m+1, m+p]$. Let c' be the number of $(b_s, \ell'_s) \in X(G, C)$ for $s \in [m+p]$. Then, finally, we run our procedure recursively on the child u_0 of type 1 to check whether (v', i') is contained in $X(G, C)$ under the assumption that exactly c' of the children of (a, ℓ) are contained in $X(G, C)$.

For now, we assumed that the considered nodes are not leaves. For a leaf t of type 0 with $v(t) = (v', i')$ and $W(t) = \varnothing$, we have $(v', i') \in X(G, C)$ if and only if $0 \in C(v')$ since (v', i') is a leaf of $H_{v,i}$ by Lemma III.6 (Item 2). For a leaf t of type 1 with $v(t) = (v', i')$, by Lemma III.6 (Item 2), we have $W(t) = \{(v', i')\}$. Hence, we have to decide $(v', i') \in X(G, C)$ under the assumption that exactly c children of (v', i') in H are contained in $X(G, C)$ for some number c. This holds if and only if $c \in C(v')$.

Before we translate the structure of the trees $T_{v,i}$ into formulae, we first check that the trees are of logarithmic height. This is essential to obtain formulae of logarithmic quantifier depth.

Claim 2 $T_{v,i}$ *has height at most* $(4r + 2) \cdot \log(n + 1)$.

Proof: By Claim 1, $H_{v,\ell}$ has the $(n + 1)^r$-path property. Thus, with Corollary III.2, we obtain $\mathrm{awt}(H_{v,\ell}) \leq (n + 1)^r \cdot |H_{v,\ell}|$, and hence, by Lemma III.6, the height of $T_{v,\ell}$ is at most $2 \log((n + 1)^r \cdot |H_{v,\ell}|) \leq 2 \log((n + 1)^r \cdot (n + 1)^r \cdot n) \leq 2 \cdot (2r + 1) \cdot \log(n + 1)$. $\lrcorner$

In the following, we recursively construct formulae of the form $\psi_{t0,i'}^{h,n}(x)$ and $\psi_{t1,i',j,c}^{h,n}(x, y)$ with domain variables x, y. Before beginning, it is appropriate to shortly analyse the syntax. By n, we refer to the order of the structure, which is fixed. With $t0$ resp. $t1$, we keep track of whether the formula at hand corresponds to a node of type 0 or 1. The number

16

$h \in \mathbb{N}$ tracks our recursion depth and ensures that we do not produce formulae with non-logarithmic depth.

The ultimate goal is to check whether (v, i) is contained in $X(G, C)$. In the process, we check for all (v', i') from the DAG $H_{v,i}$ whether they are contained in X. We want that $G \models \psi_{t0,i'}^{h,n}(v')$ if and only if we can verify $(v', i') \in X(G, C)$ with recursion depth h. This corresponds to a node t of type 0 with height at most h in the tree $T_{v,i}$ with $v(t) = (v', i')$ and $W(t) = \varnothing$, and (v', i') is in $X(G, C)$. We want that $G \models \psi_{t1,i',j,c}^{h,n}(v', w)$ if and only if we can verify $(v', i') \in X(G, C)$ with recursion depth h while stopping the recursion whenever we reach (w, j). In these cases, we assume that (w, j) has exactly c children that are contained in $X(G, C)$. This corresponds to a node t of type 1 with height at most h in $T_{v,i}$ with $v(t) = (v', i')$, $W(t) = \{(w, j)\}$, and $(v', i') \in X(G, C)$ holds if (v', i') has exactly c children in $H_{v,\ell}$ that are contained in $X(G, C)$.

Since, for every $i' < 1$, (v', i') is not contained in $X(G, C)$, we can already set $\psi_{t0,i'}^{h,n}(x) := \bot$ and $\psi_{t1,i',j,c}^{h,n}(x, y) := \bot$ for all $h \in \mathbb{N}$, $i' \in \mathbb{Z}_{\leq 0}$, $j \in \mathbb{Z}$, and $c \in \mathbb{N}$. Moreover, for all $h \in \mathbb{N}$, $i' \in \mathbb{Z}$, $j \in \mathbb{Z}_{\leq 0}$, and $c \in \mathbb{N}$, we set $\psi_{t1,i',j,c}^{h,n}(x, y) := \bot$.

Preparation: Before proceeding, let us introduce a few formulae. For all $d \in \mathbb{N}$, let $\deg_d^-(x) := \exists^{=d} y \big(E(y, x) \big)$. Then, $G \models \deg_d^-(v)$ if and only if $\deg_G^-(v) = d$. Further, for all $\ell, \ell' \in \mathbb{N}$, we inductively define $\mathsf{path}_{\ell,\ell'}^{0,n}(x, y) := (x = y)$ if $\ell = \ell'$ and

$$\mathsf{path}_{\ell,\ell'}^{0,n}(x, y) = E(x, y) \wedge \bigvee_{d \in [n], \lfloor \frac{\ell-1}{d} \rfloor = \ell'} \deg_d^-(y)$$

else, and, for $h \geq 1$,

$$\mathsf{path}_{\ell,\ell'}^{h,n}(x, y) = \exists z \bigvee_{j=\ell'}^{\ell} \Big[\mathsf{path}_{\ell,j}^{h-1,n}(x, z) \wedge \mathsf{path}_{j,\ell'}^{h-1,n}(z, y) \Big].$$

We have $G \models \mathsf{path}_{\ell,\ell'}^{h,n}(v, w)$ if and only if there is a path in $H_{v,i}$ (and thus also in any other $H_{v',j}$ that includes (v, i)) from (v, ℓ) to (w, ℓ') that can be verified in h recursion steps.

Formulae of type 0: First, we construct formulae corresponding to nodes of type 0. Let T be a tree according to the proof of Lemma III.6 and let $t \in V(T)$ be a node of type 0 of the form $v(t) = (v', i')$ and $W(t) = \varnothing$.

In the base case $h = 0$, where we do not have any further recursion steps left, we check that t is a leaf in T. This is the case if v' does not have any successors w in G with $\lfloor \frac{i'-1}{\deg^-(w)} \rfloor \geq 1$, which is equivalent to $i' - 1 \geq \deg^-(w)$. For a leaf t as described above, we have $(v', i') \in X(G, C)$ if and only if $0 \in C(v')$. Thus, for all $i' \in \mathbb{N}_{>0}$, we set

$$\psi_{t0,i'}^{0,n}(x) := P_0(x) \wedge \forall y \left[E(x, y) \rightarrow \bigvee_{d=i}^{n} \deg_d^-(y) \right].$$

In the recursion step for $h \in \mathbb{N}_{>0}$, there is some vertex $a \in V(G)$ and a number $\ell < i'$ such that $(v', i') \lhd_{H_{v,i}} (a, \ell)$, i.e., (a, ℓ) is below (v', i') in $H_{v,i}$. Intuitively, the node (a, ℓ) should split the DAG $H_{v,i}$ into parts of almost equal size. We guess the number c of children of (a, ℓ) in $H_{v,i}$ that are in $X(G, C)$. Then, we verify that exactly this number of children is contained in $X(G, C)$ via formulae of type 0 and $h - 1$ remaining recursion steps. Using the number c, we can verify $(v', i') \in X(G, C)$ with a formula of type 1 and

17

$h - 1$ remaining recursion steps by passing the information that exactly c children of (a, ℓ) are contained in $X(G, C)$.

Hence, for $h, i' \in \mathbb{N}_{>0}$, we set

$$\psi_{t0,i'}^{h,n}(x) := \psi_{t0,i'}^{0,n}(x) \vee \exists y \bigvee_{\ell=1}^{i'-1} \bigvee_{c=0}^{n} \left[\mathsf{path}_{i',\ell}^{h,n}(x,y) \right.$$
$$\left. \wedge\ \mathsf{children}_{t0,\ell,c}^{h-1,n}(y) \wedge \psi_{t1,i',\ell,c}^{h-1,n}(x,y) \right]$$

with

$$\mathsf{children}_{t0,\ell,c}^{h,n}(y) :=$$
$$\exists^{=c} z \left[E(y,z) \wedge \bigvee_{d=1}^{n} \left(\deg_d^-(z) \wedge \psi_{t0,\lfloor \frac{\ell-1}{d} \rfloor}^{h,n}(z) \right) \right]$$

for all $h \in \mathbb{N}, \ell \in \mathbb{N}_{>0}$, and $c \in \mathbb{N}$, expressing that "(y, ℓ) admits c children of type 0 which lie in $X(G, C)$. This can be verified in h recursion steps."

Formulae of type 1: Now, we construct formulae corresponding to nodes of type 1. Let T be a tree according to the proof of Lemma III.6 and let $t \in V(T)$ be a node of type 1 of the form $v(t) = (v', i')$ and $W(t) = \{(w, j)\}$.

In the case $h = 0$, where we do not have any further recursion steps left, we check that t is a leaf in T. This happens if $v' = w$ and $i' = j$. For such a leaf, assuming that exactly c children of (v', i') are in $X(G, C)$, we have $(v', i') \in X(G, C)$ if and only if $c \in C(v')$. Thus, for all $i' \in \mathbb{N}_{>0}$ and $c \in \mathbb{N}$, we set $\psi_{t1,i',i',c}^{0,n}(x,y) := P_c(x) \wedge x = y$. Furthermore, for all $i', \ell \in \mathbb{N}_{>0}$ and $c \in \mathbb{N}$ with $i' \neq \ell$, we set $\psi_{t1,i',\ell,c}^{0,n}(x,y) := \bot$.

In the recursion step for $h \in \mathbb{N}_{>0}$, there is some vertex $a \in V(G)$ and an ℓ with $j < \ell < i'$ such that $(v', i') \lhd_{H_{v,i}} (a, \ell) \lhd_{H_{v,i}} (w, j)$. We guess the number c' of children of (a, ℓ) in $H_{v,i}$ that are in $X(G, C)$. Then, we verify that exactly this number is contained in $X(G, C)$. If the child is not above (w, j) in $H_{v,i}$, then it is a node of type 0, and we use a formula of type 0 with $h - 1$ remaining recursion steps. If the child is above (w, j) in $H_{v,i}$, then it is a node of type 1, and we use a formula of type 1 with $h - 1$ remaining recursion steps, passing the information that exactly c children of (w, j) are contained in $X(G, C)$. Then, using the guessed number c', we can verify $(v', i') \in X(G, C)$ with a formula of type 1 and $h - 1$ remaining recursion steps by passing the information that exactly c' children of (a, ℓ) are contained in $X(G, C)$.

Hence, for all $h, i', j \in \mathbb{N}_{>0}$ and $c \in \mathbb{N}$, we set

$$\psi_{t1,i',j,c}^{h,n}(x,y) := \psi_{t1,i',j,c}^{0,n}(x,y)$$
$$\vee \exists z \bigvee_{\ell=j+1}^{i'-1} \bigvee_{c'=0}^{n} \left(\mathsf{path}_{i',\ell}^{h,n}(x,z) \wedge \mathsf{path}_{\ell,j}^{h,n}(z,y) \right.$$
$$\left. \wedge\ \mathsf{children}_{t1,\ell,j,c,c'}^{h-1,n}(z,y) \wedge \psi_{t1,i',\ell,c'}^{h-1,n}(x,z) \right)$$

18

with

$$\mathsf{children}_{t1,\ell,j,c,c'}^{h,n}(z,y) :=$$

$$\exists^{=c'} z' \Bigg(E(z,z') \wedge \bigvee_{d=1}^{n} \bigg[\mathsf{deg}_d^-(z')$$

$$\wedge \ \bigg(\Big[\psi_{t0,\lfloor \frac{\ell-1}{d} \rfloor}^{h,n}(z') \wedge \neg\mathsf{path}_{\lfloor \frac{\ell-1}{d} \rfloor,j}^{h,n}(z',y) \Big]$$

$$\vee \ \Big[\psi_{t1,\lfloor \frac{\ell-1}{d} \rfloor,j,c}^{h,n}(z',y) \wedge \mathsf{path}_{\lfloor \frac{\ell-1}{d} \rfloor,j}^{h,n}(z',y) \Big] \bigg) \bigg] \Bigg)$$

for all $h \in \mathbb{N}$, $\ell, j \in \mathbb{N}_{>0}$, and $c, c' \in \mathbb{N}$, expressing that "(z,ℓ) admits c' children which lie in $X(G,C)$ if c children of (y,j) are contained in $X(G,C)$. The children of (z,ℓ) above (y,j) are of type 1 and those not above (y,j) are of type 0. All of this can be verified in h recursion steps"

Nesting depth of the formulae: Since, by Lemma III.6, there is a tree $T_{v,i}$ that, by Claim 2, has height $(4r+2) \cdot \log(n+1)$, it suffices to have formulae $\psi_{t0,i}^{h,n}(x)$ with logarithmic nesting depth. That is, we choose

$$\varphi_i^{(n)}(x) := \psi_{t0,i}^{((4r+2)\cdot\log(n+1)),n}(x).$$

Then, $\mathcal{A}(G,C) \models \varphi_i^{(n)}(v)$ if and only if $(v,i) \in X(G,C)$ for all graphs G of size $|G| \leq n$, all cardinality conditions C for G, and all $v \in V(G)$. ∎

We are ready to prove the main result of this section.

Theorem I.1 *For every vocabulary τ and every $\mathsf{LREC}_=[\tau]$-formula $\varphi(\bar{x}, \bar{\kappa})$, there is a constant $k \in \mathbb{N}$ such that for every $n \in \mathbb{N}$, there is a family of $\mathsf{C}_k^{\mathcal{O}(\log n)}$-formulae $\big(\psi_{\bar{m}}(\bar{x})\big)_{\bar{m}\in[n]^{|\bar{\kappa}|}}$ such that for all τ-structures $\mathcal{A}$ of size $|\mathcal{A}| \leq n$, all $\bar{v} \in \big(V(\mathcal{A})\big)^{|\bar{x}|}$, and all $\bar{m} \in [|\mathcal{A}|]^{|\bar{\kappa}|}$, it holds that*

$$\mathcal{A} \models \varphi(\bar{v}, \bar{m}) \iff \mathcal{A} \models \psi_{\bar{m}}(\bar{v}).$$

Proof: We proceed by induction on the structure of φ. For formulae $\varphi(\bar{x}, \bar{\kappa}) \in \mathsf{FO+C}[\tau]$, since we only need equivalence on structures of size at most n, we can apply the arguments from the proof of [9, Proposition 8.4.18], first replacing #-operators by counting quantifiers, then hard-coding families of formulae inductively, beginning with $\top$ resp. $\bot$ for atomic number sentences ($\leq, \min, \max, S$) and then replacing existential numeric quantification by disjunctions over $N(\mathcal{A})$ for every possible assignment of the previously quantified variable. Hence, there are constants $k, r \in \mathbb{N}$ and, for every $n \in \mathbb{N}$, a family of $\mathsf{C}_k^r[\tau]$-formulae $\big(\psi_{\bar{j}}(\bar{x})\big)_{\bar{j}\in[n]^{|\bar{\kappa}|}}$ such that for all τ-structures $\mathcal{A}$ of size $|\mathcal{A}| \leq n$, all $\bar{v} \in \big(V(\mathcal{A})\big)^{|\bar{x}|}$, and all $\bar{j} \in [|\mathcal{A}|]^{|\bar{\kappa}|}$, it holds that

$$\mathcal{A}^+ \models \varphi(\bar{v}, \bar{j}) \iff \mathcal{A} \models \psi_{\bar{j}}(\bar{v}).$$

For $\varphi = (\neg\varphi_1)$, $\varphi = (\varphi_1 \vee \varphi_2)$, $\varphi = (\exists y \, \varphi_1)$, $\varphi = (\#\iota \, \varphi_1 = \kappa)$, or $\varphi = (\#y \, \varphi_1 = \kappa)$, where
- y is a domain variable,
- ι, κ are number variables,
- $\bar{x}$ is a tuple of domain variables,
- $\bar{\kappa}$ is a tuple of number variables we assume to contain κ,

19

- free$(\varphi) \subseteq \bar{x} \cup \bar{\kappa}$, and
- φ_1, φ_2 are $\mathsf{LREC}_=$-formulae,

we first construct families of $\mathsf{C}_k^{\mathcal{O}(\log n)}$-formulae $(\psi_{1,\bar{j}}(\bar{x}))_{\bar{j} \in [n]^{|\bar{\kappa}|}}$, $(\psi_{2,\bar{j}}(\bar{x}))_{\bar{j} \in [n]^{|\bar{\kappa}|}}$ recursively and then again apply the arguments from the proof of [9, Proposition 8.4.18].

Now let $\varphi = \left[\mathsf{lrec}_{\bar{y}_1, \bar{y}_2, \bar{\iota}} \, \varphi_=, \varphi_{\mathrm{E}}, \varphi_{\mathrm{C}}\right](\bar{x}, \bar{\kappa})$ for some compatible tuples of domain variables $\bar{y}_1, \bar{y}_2, \bar{x}$, non-empty tuples of number variables $\bar{\iota}, \bar{\kappa}$, and $\mathsf{LREC}_=$-formulae $\varphi_=(\bar{y}_1, \bar{y}_2, \bar{x})$, $\varphi_{\mathrm{E}}(\bar{y}_1, \bar{y}_2, \bar{x})$, $\varphi_{\mathrm{C}}(\bar{y}_1, \bar{x}, \bar{\iota}, \bar{\kappa})$.

By the induction hypothesis, there is a constant $k' \in \mathbb{N}$ such that for every $n \in \mathbb{N}$, there are $\mathsf{C}_{k'}^{\mathcal{O}(\log n)}$-formulae $\psi_=(\bar{y}_1, \bar{y}_2, \bar{x}), \psi_{\mathrm{E}}(\bar{y}_1, \bar{y}_2, \bar{x})$ as well as a family of $\mathsf{C}_{k'}^{\mathcal{O}(\log n)}$-formulae $\left(\psi_{\mathrm{C},\bar{\imath}\bar{\jmath}}(\bar{y}_1, \bar{x})\right)_{\bar{\imath}\bar{\jmath} \in [n]^{|\bar{\iota}\bar{\kappa}|}}$ with

$$\begin{aligned}
\mathcal{A}^+ \models \varphi_=(\bar{u}_1, \bar{u}_2, \bar{v}) &\iff \mathcal{A} \models \psi_=(\bar{u}_1, \bar{u}_2, \bar{v}), \\
\mathcal{A}^+ \models \varphi_{\mathrm{E}}(\bar{u}_1, \bar{u}_2, \bar{v}) &\iff \mathcal{A} \models \psi_{\mathrm{E}}(\bar{u}_1, \bar{u}_2, \bar{v}), \text{ and} \\
\mathcal{A}^+ \models \varphi_{\mathrm{C}}(\bar{u}_1, \bar{v}, \bar{\imath}, \bar{\jmath}) &\iff \mathcal{A} \models \psi_{\mathrm{C},\bar{\imath}\bar{\jmath}}(\bar{u}_1, \bar{v})
\end{aligned}$$

for all structures $\mathcal{A}$ of size $|\mathcal{A}| \leq n$ and all $\bar{u}_1 \in \left(V(\mathcal{A})\right)^{|\bar{y}_1|}$, $\bar{u}_2 \in \left(V(\mathcal{A})\right)^{|\bar{y}_2|}$, $\bar{v} \in \left(V(\mathcal{A})\right)^{|\bar{x}|}$, $\bar{\imath} \in N(\mathcal{A})^{|\bar{\iota}|}$, and $\bar{\jmath} \in N(\mathcal{A})^{|\bar{\kappa}|}$. Moreover, by Theorem IV.1, there is a $k'' \in \mathbb{N}_{>0}$ such that for all $n \in \mathbb{N}_{>0}$ and $\ell \in \{1, \ldots, (n+1)^{|\bar{\kappa}|}\}$, there is a $\mathsf{C}_{k''}^{\mathcal{O}(\log n)}[\{E, P_0, \ldots, P_n\}]$-formula $\varphi_{X,\ell}(x)$ such that for all τ-structures $\mathcal{A}$ of size at most n, and for $G = (V, E)$ and C from the $\mathsf{LREC}_=$ definition in Section II for φ, it holds that $\mathcal{A}(G, C) \models \varphi_{X,\ell}(v) \iff (v, \ell) \in X(G, C)$ for all $v \in V(G)$. We turn these into $\mathsf{C}_k^{\mathcal{O}(\log n)}[\tau]$-formulae $\psi_{X,\ell}(\bar{x})$ by replacing every occurrence of

- $z_1 = z_2$ by
$$\exists z_{1,1} \cdots \exists z_{1,|\bar{x}|} \exists z_{2,1} \cdots \exists z_{2,|\bar{x}|} \left(\psi_=(\bar{z}_1, \bar{z}_2, \bar{x})\right)$$

  with $\bar{z}_i = (z_{i,1}, \ldots, z_{i,|\bar{x}|})$ for $i \in \{1, 2\}$,
- $\exists z_1$ by $\exists z_{1,1} \ldots \exists z_{1,|x|}$,
- $E(z_1, z_2)$ by

$$\begin{aligned}
\exists z_{1,1}' \cdots \exists z_{1,|\bar{x}|}' \exists z_{2,1}' \cdots \exists z_{2,|\bar{x}|}' \big(\\
\psi_=(\bar{z}_1', \bar{z}_1, \bar{x}) \wedge \psi_=(\bar{z}_2', \bar{z}_2, \bar{x}) \wedge \psi_{\mathrm{E}}(\bar{z}_1', \bar{z}_2', \bar{x})\big),
\end{aligned}$$

- and $P_i(z_1)$ by $\exists z_{1,1}' \cdots \exists z_{1,|\bar{x}|}' \left(\psi_=(\bar{z}', \bar{z}, \bar{x}) \wedge \psi_{\mathrm{C},\bar{\imath}\ell}(\bar{z}', \bar{x})\right)$

for any variables z_1, z_2 that occur in $\varphi_{X,\ell}$.

Let $\psi_{\bar{m}}(\bar{x}) := \psi_{X,\langle\bar{m}\rangle}(\bar{x})$ for all $\bar{m} \in [n]^{|\bar{\kappa}|}$. Then, for all τ-structures $\mathcal{A}$ of size $|\mathcal{A}| \leq n$, all $\bar{v} \in \left(V(\mathcal{A})\right)^{|\bar{x}|}$, and all $\bar{m} \in [|\mathcal{A}|]^{|\bar{\kappa}|}$, it holds that

$$\begin{aligned}
\mathcal{A}^+ \models \varphi(\bar{v}, \bar{m}) &\iff (\bar{v}, \langle\bar{m}\rangle) \in X(G, C) \\
&\iff \mathcal{A} \models \psi_{X,\langle\bar{m}\rangle}(\bar{v}) \\
&\iff \mathcal{A} \models \psi_{\bar{m}}(\bar{v}).
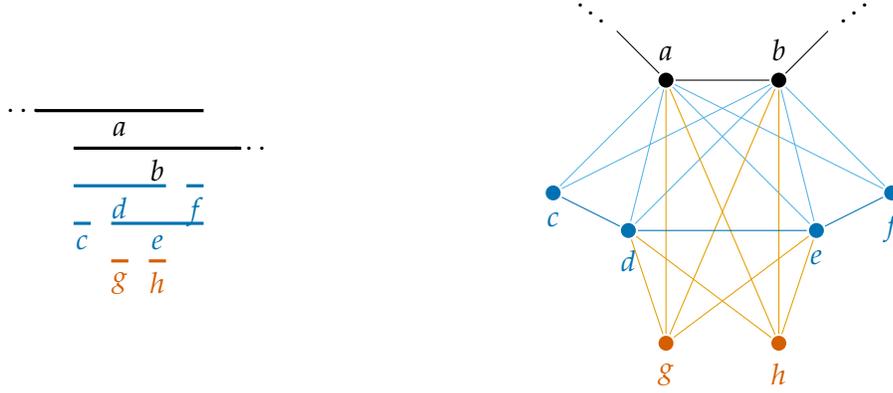\end{aligned}$$

$\blacksquare$

Fig. 5: Part of an interval representation (left) and its interval graph G admitting nested modules (right). The four (visible) maxcliques $C_{i,i\in[4]}$ of G are $\{a,b,c,d\}$, $\{a,b,d,e,g\}$, $\{a,b,d,e,h\}$, and $\{a,b,e,f\}$.

V. INTERVAL GRAPHS

In this section, we describe how the result from Theorem I.1 allows us to obtain, from an $\mathsf{LREC}_=$-definable canonisation of interval graphs, a $k \in \mathbb{N}$ such that, for every $n \in \mathbb{N}$, $\mathsf{C}_k^{\mathcal{O}(\log n)}$ identifies every interval graph of order n. We obtain a similar result for chordal claw-free graphs. Finally, we sketch how the result we obtained for interval graphs can be shown without the need for $\mathsf{LREC}_=$, using an $\mathsf{STC+C}$-definable canonisation for a subclass of interval graphs and the fact that every interval graph can be decomposed into interval graphs of that subclass.

An *interval* is a set of consecutive integers. An *interval representation* $\mathcal{I}$ is a set of intervals, from which we get its graph $G_{\mathcal{I}}$ with $V(G_{\mathcal{I}}) := \mathcal{I}$ and $E(G_{\mathcal{I}}) := \{\{I,J\} \subseteq \mathcal{I} \mid I \cap J \neq \varnothing\}$. An undirected graph G is an *interval graph* if there exists an interval representation $\mathcal{I}$ such that $G \cong G_{\mathcal{I}}$, see Fig. 5 for an example. An interval representation $\mathcal{I}$ is *(cardinalitywise) minimal* if $\bigcup \mathcal{I} \subset \mathbb{N}$ is minimal with respect to all interval representations $\mathcal{I}'$ with $G_{\mathcal{I}} \cong G_{\mathcal{I}'}$. An interval graph G is *proper* if there is an interval representation $\mathcal{I}$ with $G_{\mathcal{I}} \cong G$ and, for all $I, J \in \mathcal{I}$, $I \not\subseteq J$.

Lemma V.1 ([24]) *There exists an $\mathsf{LREC}_=$-definable canonisation of interval graphs $\psi(\iota,\kappa)$ such that, for all $n \in \mathbb{N}$ and interval graphs G of order n, it holds that $\mu \colon G^+ \cong ([n], \psi[G^+; \iota, \kappa])$.*

In particular, since ψ is the result of a canonisation, it holds for all $n \in \mathbb{N}$ and interval graphs G and H of order n that $G \cong H$ iff $([n], \psi[G^+; \iota, \kappa]) = ([n], \psi[H^+; \iota, \kappa])$.

Applying Theorem I.1, we obtain:

Corollary V.2 *There exists a $k \in \mathbb{N}$ such that for all $n \in \mathbb{N}$ there is a family of $\mathsf{C}_k^{\mathcal{O}(\log n)}$-sentences $(\psi_{ij})_{i,j\in[n]}$ such that for all interval graphs G of order n and $i,j \in [n]$*
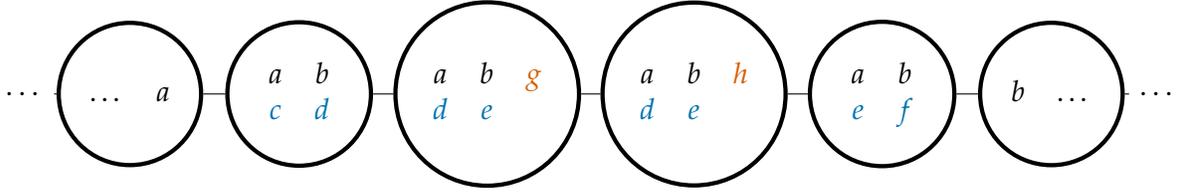
$$G \models \psi_{ij} \iff G^+ \models \psi(i,j).$$

Fig. 6: Part of a possible path decomposition for the graph G from Fig. 5. In particular, observe that other possible path decompositions can be obtained by totally reversing the contents of the four, resp. two centremost nodes.

Now, let, for every $n \in \mathbb{N}$, interval graph G of order n and all $i, j \in [n]$,

$$
\psi_{ij}^G := \begin{cases} \top & \text{if } G \models \psi_{ij}, \\ \bot & \text{if } G \not\models \psi_{ij}. \end{cases}
$$

Lemma V.3 *There exists a $k \in \mathbb{N}$ such that, for every $n \in \mathbb{N}$, every interval graph G of order n admits a $\mathsf{C}_k^{\mathcal{O}(\log n)}$-formula φ_G satisfying, for every interval graph H of order n,*

$$
H \models \varphi_G \iff H \cong G.
$$

Proof: Let G be an interval graph of order n. We claim that

$$
\varphi^G := \bigwedge_{i,j \in [n]} \psi_{ij} \leftrightarrow \psi_{ij}^G
$$

is that formula. To that end, let H be an interval graph of order n.

($\implies$) Suppose that $H \models \varphi_G$. Then, for all $i, j \in [n]$, $H \models \psi_{ij} \leftrightarrow \psi_{ij}^G$, implying that $H \models \psi_{ij}$ iff $G \models \psi_{ij}$. Thus, $([n], \psi[G^+; \iota, \kappa]) = ([n], \psi[H^+; \iota, \kappa])$ and hence, $G \cong H$.

($\impliedby$) Suppose that $G \cong H$. Then, $([n], \psi[G^+; \iota, \kappa]) = ([n], \psi[H^+; \iota, \kappa])$. Thus, for all $i, j \in [n]$, $H \models \psi_{ij}$ iff $G \models \psi_{ij}$ and thus $\psi_{ij} \leftrightarrow \psi_{ij}^G$. Hence, $H \models \varphi_G$. ∎

It thus remains to show that we can separate an interval graph from those that are not interval graphs or are of a different order. For this, we need the logics STC and STC+C.

Symmetric transitive closure logic STC (see [39]) extends FO by the stc-*operator*, which, for all vocabularies τ, τ-structures $\mathcal{A}$ and $k \in \mathbb{N}$, allows the definition of an undirected graph over vertex k-tuples. Syntactically, if ψ is an STC$[\tau]$-formula, $\bar{x}$ and $\bar{y}$ are k-tuples of variables, and $\bar{v}, \bar{w} \in V(\mathcal{A})^k$, then $\varphi := [\mathsf{stc}_{\bar{x},\bar{y}} \psi](\bar{v}, \bar{w})$ is also an STC-formula. Concerning the semantics, it suffices for us to know that in φ, ψ defines an undirected graph over $V(\mathcal{A})^k$. The stc-operator then tests whether $(\bar{v}, \bar{w})$ is an edge in the symmetric transitive closure of said graph. In that sense, it is very close to, though more restrictive than, the tc-operator from transitive closure logic [31]. The extension of STC to two-sorted structures then yields STC+C, where the stc-operator is extended over mixed domain/number tuples. STC+C has been found to be contained in LREC$_=$.

Lemma V.4 ([24]) STC+C $\leq$ LREC$_=$.

Thus, applying Theorem I.1, we immediately obtain the following.

22

Lemma V.5 *For every vocabulary τ and every* STC+C$[\tau]$*-formula $\varphi(\bar{x}, \bar{\iota})$, there is a constant $k \in \mathbb{N}$ such that for every $n \in \mathbb{N}$, there is a family of* C$_k^{\mathcal{O}(\log n)}$*-formulae $(\psi_{\bar{\iota}}(\bar{x}))_{\bar{\iota} \in [n]^{|\bar{\iota}|}}$ such that for all τ-structures $\mathcal{A}$ of size $|\mathcal{A}| \leq n$, all $\bar{v} \in (V(\mathcal{A}))^{|\bar{x}|}$, and all $\bar{\iota} \in (N(\mathcal{A}))^{|\bar{\iota}|}$, it holds that*

$$\mathcal{A}^+ \models \varphi(\bar{v}, \bar{\iota}) \iff \mathcal{A} \models \psi_{\bar{\iota}}(\bar{v}).$$

It is worth mentioning that we can also obtain the above result without passing through LREC$_=$. The approach is similar to the one of Theorem I.1, an induction over the structure of the formula, although easier, as the stc-operator can be modelled as a simple connectivity test on a graph defined over vertex/number k-tuples. This can be expressed through a formula of depth logarithmic in the size of the graph (see [20, Example 3]).

Lemma V.6 ([39]) *The class of interval graphs is* STC*-definable.*

Corollary V.7 *There is a $k \in \mathbb{N}$ such that, for every $n \in \mathbb{N}$, there exists a* C$_k^{\mathcal{O}(\log n)}$*-sentence $\varphi_{interval}^{(n)}$ such that for every graph G of order n, it holds that $G \models \varphi_{interval}^{(n)}$ if and only if G is an interval graph.*

By combining Lemma V.3 and Corollary V.7, we can now prove our main theorem of this section.

Theorem I.2 *There is a $k \in \mathbb{N}$ such that for every $n \in \mathbb{N}$ and interval graph G of order n, there is a formula $\Phi^G \in$ C$_k^{\mathcal{O}(\log n)}$ that describes G up to isomorphism.*

Proof: Let H be a graph with $G \not\cong H$. If $|V(H)| \neq |V(G)|$, then G and H are separated by the formula $\exists^{=n} x \, (x = x)$, where n is the order of G. Thus, let $|V(H)| = |V(G)|$. If $V(H)$ is not an interval graph, then, by Corollary V.7, G and H are separated by some formula $\varphi_{interval}^{(n)} \in$ C$_{k'}^{\mathcal{O}(\log n)}[\{E\}]$ for some fixed $k' \in \mathbb{N}$. Therefore, suppose that H is an interval graph of order n. Then, by Lemma V.3 $H \not\models \varphi^G$ with $\varphi^G \in$ C$_{k''}^{\mathcal{O}(\log n)}[\{E\}]$ for some fixed $k'' \in \mathbb{N}$. Hence, letting $k = \max(k', k'')$, the formula $\Phi^G \in$ C$_k^{\mathcal{O}(\log n)}[E]$ defined as

$$\Phi^G := \exists^{=n} x \, (x = x) \wedge \varphi_{interval}^{(n)} \wedge \varphi^G$$

describes G up to isomorphism. ∎

Corollary V.8 *There is a $k \in \mathbb{N}$ such that* WL$_k^{\mathcal{O}(\log n)}$ *identifies every n-vertex interval graph.*

We obtain a similar result for chordal claw-free graphs. A graph is *chordal* if every cycle of length at least 4 admits a chord. This can be expressed by an STC-sentence which, for every path of length 3 of a graph G, tests that it cannot be closed to an induced cycle of length at least 4. A graph is *claw-free* if it has no induced subgraph isomorphic to the complete bipartite graph $K_{1,3}$. It is clear that this can be tested by a C$_4^4$-formula. Thus, whether a graph is chordal and claw-free can be tested by an STC-sentence. In addition, chordal claw-free graphs admit LREC$_=$-definable canonisation.

Lemma V.9 ([27]) *There exists an* LREC$_=$*-definable canonisation of chordal claw-free graphs $\psi(\iota, \kappa)$ such that, for all interval graphs G,*

$$\mu \colon G^+ \cong ([n], \psi[G^+; \iota, \kappa]).$$

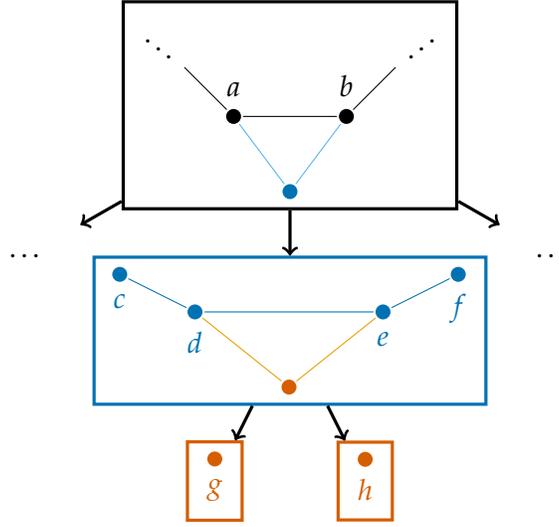An analogous argumentation yields Theorem I.3.

Fig. 7: A simplified sketch of the modular decomposition corresponding to G from Fig. 5. Letting $M \neq C_{i,i \in [4]}$ be a possible end of G, the vertex set $S := \bigcup_{i \in [4]} C_i \setminus \bigcup(\mathcal{M}(G) \setminus \{C_{i,i \in [4]}\}) = \{c,d,e,f,g,h\}$ is a module of G, whereas the vertex sets $\{g\}$ and $\{h\}$ are modules of $G[S]$, with $\{c,d\}$ and $\{e,f\}$ as possible ends of $G[S]$.

Theorem I.3 *There is a $k \in \mathbb{N}$ such that for every chordal claw-free graph G of order n, there is a formula $\Phi^G \in \mathsf{C}_k^{\mathcal{O}(\log n)}$ that describes G up to isomorphism.*

A. Circumventing $\mathsf{LREC}_=$

Finally, we sketch how Theorem I.2 can be proved directly without considering $\mathsf{LREC}_=$ and applying the results due to Laubner [38] and Grußien [27]. We begin with some insights into the properties of interval graphs. A *maxclique* of a graph G is a vertex subset $C \subseteq V(G)$ such that C forms a clique and, for all $v \in V(G) \setminus C$, $C \cup \{v\}$ does not form a clique. By $\mathcal{M}(G)$, we denote the set of all maxcliques of G.

Observe that interval graphs are exactly those graphs whose maxcliques can be brought into a (not necessarily unique) linear order such that every vertex is contained in consecutive maxcliques of that order [42]. Equivalently, they are those graphs admitting path decompositions (see [45]) of which every bag corresponds to a maxclique, as shown in Fig. 6.

A *possible end* of an interval graph G is a maxclique $M \in \mathcal{M}(G)$ such that G admits a path decomposition (P, β) satisfying, for some $p \in V(P)$ with $\deg_P(p) = 1$, that $\beta(p) = M$. Given such a possible end M of an interval graph G, we can obtain an STC-definable strict weak order $\prec_M$ which captures all the information about the order of the maxcliques [38]. Formally, $\prec_M$ is initialised as $M \prec_M C$ for all $C \in \mathcal{M}(G) \setminus \{M\}$ and is then recursively extended through

$$C \prec_M D \text{ if } \exists X \in \mathcal{M}(G) : \begin{cases} X \prec_M D, (X \cap C) \setminus D \neq \varnothing, \\ C \prec_M X, (X \cap D) \setminus C \neq \varnothing. \end{cases}$$

Notably, for some interval graphs G and possible ends M, the order $\prec_M$ becomes a linear order over $\mathcal{M}(G)$; it can be shown that extending it over $V(G)$ induces a strict weak order
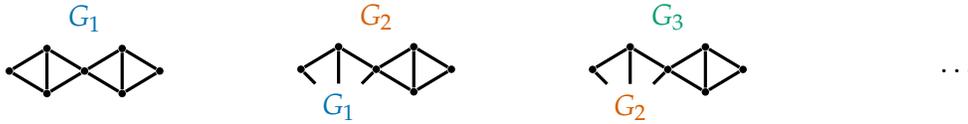
Fig. 8: First members of a family $(G_i)_{i \geq 1}$ of interval graphs such that, for all $i \in \mathbb{N}_{>0}$, the modular decomposition of G_i has a height of i. This is because, for $G_{i,i \geq 2}$, a large fraction $\left(> \frac{|V(G)|}{2} \right)$ of its vertices is contained in its (unique) module.

$<_G$ in which two vertices $v, w \in V(G)$ are $<_G$-incomparable iff $N_G[v] = N_G[w]$. This yields an STC+C-definable canonisation for those interval graphs [38]. In particular, thanks to Lemma V.5, we obtain a fixed k^* and a $\mathsf{C}_{k^*}^{\mathcal{O}(\log n)}$-formula for each such interval graph describing it up to isomorphism.

Now, we consider those interval graphs G for which $\prec_M$ is not a linear order. A set $\mathcal{C} \subseteq \mathcal{M}(G)$ of maxcliques is *incomparable* wrt. $\prec_M$ if for all pairwise distinct maxcliques $C_1, C_2 \in \mathcal{C}$, neither $C_1 \prec_M C_2$ or $C_2 \prec_M C_1$. A set $\mathcal{C} \subseteq \mathcal{M}(G)$ is *maximal* if, for all $C_1 \in \mathcal{C}$ and $C \in \mathcal{M}(G) \setminus \mathcal{C}$, $C_1 \prec_M C$ or $C \prec_M C_1$. For each such maximal incomparable set of maxcliques $\mathcal{C} \subseteq \mathcal{M}(G)$ and $\mathcal{D} := \mathcal{M}(G) \setminus \mathcal{C}$, the vertex set $S_{\mathcal{C}} := \bigcup \mathcal{C} \setminus \bigcup \mathcal{D}$ is a *module*, a vertex subset with a uniform connectivity behaviour towards the rest of the graph, that is, for every $v \in V(G) \setminus S_{\mathcal{C}}$, either $vw \in E(G)$ or $vw \notin E(G)$ for all $w \in S_{\mathcal{C}}$. In the following, we consider only those modules that are obtained through maximal incomparable sets of maxcliques. Note that for each module $S_{\mathcal{C}}$, it holds that $G[S_{\mathcal{C}}]$ is also an interval graph, see Fig. 7.

There, we see how inductively replacing modules by individual vertices (which maintain the same connectivity to the rest of the graph as the modules they replace) yields a *modular decomposition (tree) T*. Let G_S denote the copy of G in which all modules have been contracted to vertices. Conveniently, each such G_S admits a linear order $\prec_M$ over its maxcliques for some possible end $M \in \mathcal{M}(G_S)$ and can thus be described up to isomorphism in $\mathsf{C}_{k^*}^{\mathcal{O}(\log n)}$ by Lemma V.5.

Thus, for each $t \in V(T)$, the subgraph it corresponds to can be described in $\mathsf{C}_{k^*}^{\mathcal{O}(\log n)}$. Hence, developing a formula which, in an inductive, bottom-up fashion, characterises G up to isomorphism is not very difficult. The main issue is the height of T. This is due to the fact that an interval graph G may have up to and at most one module $S_{\mathcal{C}}$ with $|S_{\mathcal{C}}| > |G|/2$. Thus, we can construct families of interval graphs whose modular decomposition trees are linear in the size of the graph (see Fig. 8 for an example).

The idea is thus to build a treelike decomposition of the modular decomposition tree similarly as Lemma III.6. Since this decomposition has a logarithmic height in the size of the input graph, it then suffices to inductively describe our graph based on it, yielding formulae of logarithmic quantifier depth.

VI. Conclusion

We have shown that for every LREC$_=$-definable property, there is a constant k such that for every size bound n, the property can be expressed on structures of size at most n via a family of $\mathsf{C}_k^{\mathcal{O}(\log n)}$-formulae. This implies that the k-dimensional Weisfeiler–Leman algorithm distinguishes every pair of graphs separable by the property in a logarithmic

number of iterations. Moreover, building on results by Grohe et al. [24] and Grußien [27], this yields that the algorithm identifies every interval graph and every chordal claw-free graph in logarithmically many iterations.

It remains an interesting project to investigate the power of counting logics with logarithmic quantifier depth, or equivalently, a logarithmic number of iterations of the Weisfeiler–Leman algorithm, on other graph classes. A natural target class would be graphs defined by a finite set of excluded minors.

Also, since our results are non-uniform as we obtain formulae for each size bound n, a follow-up question could ask how to obtain similar uniform statements: is every LREC$_=$-formula equivalent to a formula of fixed-point logic with counting that only uses logarithmically many iterations?

References

[1] Albert Atserias and Elitza N. Maneva. Sherali–Adams relaxations and indistinguishability in counting logics. *SIAM J. Comput.*, 42(1):112–137, 2013. doi: 10.1137/120867834.

[2] Albert Atserias and Joanna Ochremiak. Definable ellipsoid method, sums-of-squares proofs, and the isomorphism problem. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018*, pages 66–75. ACM, 2018. doi: 10.1145/3209108.3209186.

[3] Christoph Berkholz and Jakob Nordström. Near-optimal lower bounds on quantifier depth and Weisfeiler–Leman refinement steps. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2016*, pages 267–276. ACM, 2016. doi: 10.1145/2933575.2934560.

[4] Andreas Blass, Yuri Gurevich, and Saharon Shelah. Choiceless polynomial time. *Ann. Pure Appl. Log.*, 100(1-3):141–187, 1999. doi: 10.1016/S0168-0072(99)00005-6.

[5] Jin-Yi Cai, Martin Fürer, and Neil Immerman. An optimal lower bound on the number of variables for graph identification. *Combinatorica*, 12(4):389–410, 1992.

[6] Andreas Darmann, Ulrich Pferschy, and Joachim Schauer. Resource allocation with time intervals. *Theor. Comput. Sci.*, 411(49):4217–4234, 2010. doi: 10.1016/j.tcs.2010.08.028.

[7] Anuj Dawar and Felipe Ferreira Santos. Separating LREC from LFP. In *LICS '22: 37th Annual ACM/IEEE Symposium on Logic in Computer Science*, pages 55:1–55:13. ACM, 2022.

[8] Holger Dell, Martin Grohe, and Gaurav Rattan. Lovász meets Weisfeiler and Leman. In *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018*, volume 107 of *LIPIcs*, pages 40:1–40:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. doi: 10.4230/LIPIcs.ICALP.2018.40.

[9] Heinz-Dieter Ebbinghaus and Jörg Flum. *Finite model theory*. Springer Science & Business Media, 2005.

[10] Kousha Etessami and Neil Immerman. Tree canonization and transitive closure. *Inf. Comput.*, 157(1-2):2–24, 2000. doi: 10.1006/inco.1999.2835.

[11] Sergei Evdokimov, Ilia Ponomarenko, and Gottfried Tinhofer. Forestal algebras and algebraic forests (on a new class of weakly compact graphs). *Discret. Math.*, 225(1-3): 149–172, 2000. doi: 10.1016/S0012-365X(00)00152-7.

[12] Ronald Fagin. Generalized first-order spectra and polynomial-time recognizable sets. In *Complexity of computation (Proc. SIAM-AMS Sympos., New York, 1973)*, SIAM-AMS Proc., Vol. VII, pages 43–73. Amer. Math. Soc., Providence, R.I., 1974.

[13] Martin Fürer. Weisfeiler–Lehman refinement requires at least a linear number of iterations. In *Automata, Languages and Programming, 28th International Colloquium, ICALP 2001*, volume 2076 of *Lecture Notes in Computer Science*, pages 322–333. Springer, 2001. doi: 10.1007/3-540-48224-5_27.

[14] Erich Grädel and Martin Grohe. Is polynomial time choiceless? In *Fields of Logic and Computation II - Essays Dedicated to Yuri Gurevich on the Occasion of His 75th Birthday*, volume 9300 of *Lecture Notes in Computer Science*, pages 193–209. Springer, 2015. doi: 10.1007/978-3-319-23534-9_11.

[15] Erich Grädel and Svenja Schalthöfer. Choiceless logarithmic space. In *44th International Symposium on Mathematical Foundations of Computer Science, MFCS 2019*, volume 138 of *LIPIcs*, pages 31:1–31:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

[16] Martin Grohe. Isomorphism testing for embeddable graphs through definability. In *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing, 2000*, pages 63–72. ACM, 2000. doi: 10.1145/335305.335313.

[17] Martin Grohe. *Descriptive Complexity, Canonisation, and Definable Graph Structure Theory*, volume 47 of *Lecture Notes in Logic*. Cambridge University Press, 2017. doi: 10.1017/9781139028868.

[18] Martin Grohe. The logic of graph neural networks. In *36th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2021, Rome, Italy, June 29 - July 2, 2021*, pages 1–17. IEEE, 2021. doi: 10.1109/LICS52264.2021.9470677.

[19] Martin Grohe and Sandra Kiefer. A linear upper bound on the Weisfeiler–Leman dimension of graphs of bounded genus. In *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019*, volume 132 of *LIPIcs*, pages 117:1–117:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019. doi: 10.4230/LIPIcs.ICALP.2019.117.

[20] Martin Grohe and Sandra Kiefer. Logarithmic Weisfeiler–Leman identifies all planar graphs. In *48th International Colloquium on Automata, Languages, and Programming (ICALP 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021.

[21] Martin Grohe and Daniel Neuen. Canonisation and definability for graphs of bounded rank width. In *34th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2019*, pages 1–13. IEEE, 2019. doi: 10.1109/LICS.2019.8785682.

[22] Martin Grohe and Martin Otto. Pebble games and linear equations. *J. Symb. Log.*, 80 (3):797–844, 2015. doi: 10.1017/jsl.2015.28.

[23] Martin Grohe and Oleg Verbitsky. Testing graph isomorphism in parallel by playing a game. In *Automata, Languages and Programming, 33rd International Colloquium, ICALP 2006*, volume 4051 of *Lecture Notes in Computer Science*, pages 3–14. Springer, 2006. doi: 10.1007/11786986_2.

[24] Martin Grohe, Berit Grußien, André Hernich, and Bastian Laubner. L-recursion and a new logic for logarithmic space. *Logical Methods in Computer Science*, 9, 2013.

[25] Martin Grohe, Kristian Kersting, Martin Mladenov, and Pascal Schweitzer. Color refinement and its applications. In *An Introduction to Lifted Probabilistic Inference*. The MIT Press, 08 2021.

[26] Martin Grohe, Moritz Lichter, and Daniel Neuen. The iteration number of the weisfeiler-leman algorithm. *ArXiv*, 2301.13317, 2023. Conference version in this conference.

[27] Berit Grußien. Capturing logarithmic space and polynomial time on chordal claw-free

graphs. *Log. Methods Comput. Sci.*, 15(3), 2019.

[28] Udaiprakash I. Gupta, Der-Tsai Lee, and JY-T Leung. Efficient algorithms for interval graphs and circular-arc graphs. *Networks*, 12(4):459–467, 1982.

[29] Neil Immerman. Relational queries computable in polynomial time. *Inf. Control.*, 68 (1-3):86–104, 1986. doi: 10.1016/S0019-9958(86)80029-8.

[30] Neil Immerman. Languages that capture complexity classes. *SIAM J. Comput.*, 16(4): 760–778, 1987. doi: 10.1137/0216051.

[31] Neil Immerman. *Descriptive complexity*. Graduate texts in computer science. Springer, 1999. doi: 10.1007/978-1-4612-0539-5.

[32] J. Mark Keil. Finding Hamiltonian circuits in interval graphs. *Information Processing Letters*, 20(4):201–206, 1985.

[33] Sandra Kiefer. *Power and Limits of the Weisfeiler–Leman Algorithm*. PhD thesis, RWTH Aachen University, Aachen, 2020.

[34] Sandra Kiefer and Daniel Neuen. The power of the Weisfeiler–Leman algorithm to decompose graphs. *SIAM J. Discret. Math.*, 36(1):252–298, 2022. doi: 10.1137/20m1314987.

[35] Sandra Kiefer and Pascal Schweitzer. Upper bounds on the quantifier depth for graph differentiation in first-order logic. *Log. Methods Comput. Sci.*, 15(2), 2019. doi: 10.23638/LMCS-15(2:19)2019.

[36] Sandra Kiefer, Ilia Ponomarenko, and Pascal Schweitzer. The Weisfeiler–Leman dimension of planar graphs is at most 3. *J. ACM*, 66(6):44:1–44:31, 2019. doi: 10.1145/3333003.

[37] Johannes Köbler, Sebastian Kuhnert, Bastian Laubner, and Oleg Verbitsky. Interval graphs: Canonical representations in logspace. *SIAM J. Comput.*, 40(5):1292–1315, 2011. doi: 10.1137/10080395X.

[38] Bastian Laubner. Capturing polynomial time on interval graphs. In *2010 25th Annual IEEE Symposium on Logic in Computer Science*, pages 199–208. IEEE, 2010.

[39] Bastian Laubner. *The structure of graphs and new logics for the characterization of Polynomial Time*. PhD thesis, Humboldt Universität zu Berlin, 2011.

[40] Leonid Libkin. *Elements of finite model theory*, volume 41. Springer, 2004.

[41] Moritz Lichter, Ilia Ponomarenko, and Pascal Schweitzer. Walk refinement, walk logic, and the iteration number of the Weisfeiler–Leman algorithm. In *34th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2019*, pages 1–13. IEEE, 2019. doi: 10.1109/LICS.2019.8785694.

[42] Rolf H. Möhring. Algorithmic aspects of comparability graphs and interval graphs. In *Graphs and Order*, pages 41–101. Springer, 1985.

[43] Christopher Morris, Martin Ritzert, Matthias Fey, William L. Hamilton, Jan Eric Lenssen, Gaurav Rattan, and Martin Grohe. Weisfeiler and Leman go neural: Higher-order graph neural networks. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019*, pages 4602–4609. AAAI Press, 2019. doi: 10.1609/aaai.v33i01.33014602.

[44] Martin Otto. *Bounded variable logics and counting – A study in finite models*, volume 9 of *Lecture Notes in Logic*. Springer Verlag, 1997.

[45] Neil Robertson and Paul D. Seymour. Graph minors. I. excluding a forest. *Journal of Combinatorial Theory, Series B*, 35(1):39–61, 1983.

[46] Moshe Y. Vardi. The complexity of relational query languages (extended abstract). In *Proceedings of the 14th Annual ACM Symposium on Theory of Computing, 1982*, pages

137–146. ACM, 1982. doi: 10.1145/800070.802186.

[47] Peisen Zhang, Eric A. Schon, Stuart G Fischer, Eftihia Cayanis, Janie Weiss, Susan Kistler, and Philip E. Bourne. An algorithm based on graph theory for the assembly of contigs in physical mapping of dna. *Bioinformatics*, 10(3):309–317, 1994.