


Adaptive Query Algorithms for Relational Structures Based on Homomorphism Counts


Balder ten Cate 

Institute for Logic, Language and Computation, University of Amsterdam, The Netherlands

Phokion G. Kolaitis 

University of California Santa Cruz, CA, USA

IBM Research - Almaden, CA, USA

Arnar Á. Kristjánsson 

Master of Logic, Institute for Logic, Language and Computation, University of Amsterdam, The Netherlands

Abstract

A query algorithm based on homomorphism counts is a procedure to decide membership for a class of finite relational structures using only homomorphism count queries. A left query algorithm can ask the number of homomorphisms from any structure *to* the input structure and a right query algorithm can ask the number of homomorphisms *from* the input structure to any other structure. We systematically compare the expressive power of different types of left or right query algorithms, including non-adaptive query algorithms, adaptive query algorithms that can ask a bounded number of queries, and adaptive query algorithms that can ask an unbounded number of queries. We also consider query algorithms where the homomorphism counting is done over the Boolean semiring \mathbb{B} , meaning that only the existence of a homomorphism is recorded, not the precise number of them.

2012 ACM Subject Classification Theory of computation \rightarrow Finite Model Theory

Keywords and phrases Query algorithms, homomorphisms, homomorphism counts, directed graphs, relational structures, Datalog, constraint satisfaction

Digital Object Identifier 10.4230/LIPIcs.MFCS.2025.34

Related Version *Full Version:* <https://arxiv.org/abs/2504.16567> [20]

1 Introduction

A classic result by Lovász [15] states that a finite relational structure A is determined up to isomorphism by the numbers $\text{hom}(B, A)$ of homomorphisms from B to A for every finite relational structure B with the same signature. This list of numbers is called the *left homomorphism profile* of A and is written as $\text{hom}(\mathcal{C}, A)$, where \mathcal{C} is the class of all finite relational structures with the same signature as A . Similarly, one can define the *right homomorphism profile* as the list of numbers of homomorphism from A to each structure. The right profile also determines the structure up to isomorphism, as shown by Chaudhuri and Vardi [6].

Recently, there has been considerable interest in exploring which structures can be distinguished by the homomorphism profiles when the class \mathcal{C} is restricted. As an example, in 2010, Dvořák [8] proved that if \mathcal{C} is the class of graphs of treewidth at most k , the left profile can distinguish precisely the same graphs as the k -dimensional Weisfeiler-Leman algorithm, a well-known polynomial time graph isomorphism test. Cai, Fürer, and Immerman [2] had already proved that the $k + 1$ variable fragment of first-order logic extended with counting quantifiers can distinguish precisely the same graphs as k -dimensional Weisfeiler-Leman algorithm. Also, in 2020, Grohe [12] proved that first-order logic with counting of quantifier rank at most k can distinguish exactly the same graphs as the left homomorphism profile



© Balder ten Cate, Phokion G. Kolaitis, and Arnar Á. Kristjánsson;
licensed under Creative Commons License CC-BY 4.0

50th International Symposium on Mathematical Foundations of Computer Science (MFCS 2025).

Editors: Paweł Gawrychowski, Filip Mazowiecki, and Michał Skrzypczak; Article No. 34; pp. 34:1–34:18

Leibniz International Proceedings in Informatics



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

restricted to graphs of treedepth k . Mančinska and Roberson [16] showed that the graphs that are quantum isomorphic are exactly the graphs that have the same left homomorphism profile restricted to planar graphs. In 2024, Seppelt [18] provides further insight into the logical equivalence relations for graphs that can be characterized by restricted left homomorphism profiles, using techniques from structural graph theory. These results highlight the importance of homomorphism profiles in finite model theory by describing a close connection between homomorphism profiles and logical languages.

Taking a slightly different approach, in 2022, Chen et al. [7] ask what properties of graphs can be decided based on *finitely many* homomorphism count queries. They did this by introducing the notion of a homomorphism-count based *query algorithm*. A *non-adaptive left k -query algorithm* (also referred to as a left k -query algorithm) consists of a finite tuple $\mathcal{F} = (F_1, \dots, F_k)$ of structures and a subset $X \subseteq \mathbb{N}^k$. The algorithm is said to decide the class of structures

$$\mathcal{A} := \{A : (\text{hom}(F_i, A))_{i=1}^k \in X\}.$$

Non-adaptive *right k -query algorithms* are defined similarly, but using $\text{hom}(A, F_i)$ instead of $\text{hom}(F_i, A)$. These notions allow for the study of what classes of structures can be defined using a finite restriction of the left (respectively, right) homomorphism profile. Chen et al. also define an *adaptive* version of left/right k -query algorithms. These are, roughly speaking, algorithms that can ask k homomorphism-count queries where the choice of the $(i + 1)$ -th query may depend on the answers to the first i queries. Chen et al. [7] studied the expressive power of non-adaptive and adaptive query algorithms for the specific case of *simple graphs* (i.e., undirected graphs without self-loops). Among other things, they showed that every first-order sentence ϕ that is a Boolean combination of universal first-order sentences can be decided by a non-adaptive k -left query algorithm for some k . On the other hand, they showed that there are first-order sentences ϕ that cannot be decided by a non-adaptive left k -query algorithm for any k . When it comes to adaptive query algorithms, they showed that every class of simple graphs can be decided by an adaptive left 3-query algorithm, but that there exists a class of simple graphs that is not decided by an adaptive right k -query algorithm for any k , when using only simple graphs for the queries.

In a 2024 paper, ten Cate et al. [19] extend the above analysis by exploring query algorithms for arbitrary relational structures (not only simple graphs) and by considering also query algorithms over the Boolean semiring \mathbb{B} , which means that instead of being able to query for the number of homomorphisms, the algorithm can only query for the existence of a homomorphism. The query algorithms as defined by Chen et al. [7] can be viewed as query algorithms over the semiring \mathbb{N} of natural numbers. It is shown in [19] that, for classes of structures that are closed under homomorphic equivalence, non-adaptive left k -query algorithms (over \mathbb{N}) are no more powerful than non-adaptive left k -query algorithms over \mathbb{B} . In other words, “counting does not help” for such classes. Moreover, for such classes, non-adaptive left query algorithms are equivalent in expressive power to first-order logic, meaning that such classes are first-order definable if and only if they are decided by a non-adaptive left query algorithm over \mathbb{N} (or \mathbb{B}).

Our contributions. We study the expressive power of different versions of these query algorithms for classes of relational structures. In doing so, we characterize their expressive power, relate it to that of logics and database query languages, and compare it across different types of algorithms. We focus mainly on adaptive query algorithms, including the previously unexplored notion of adaptive *unbounded* query algorithms – i.e., adaptive query algorithms for which there is no uniform bound on the number of queries they can ask on a given input, but which nevertheless terminate on every input.

$$\begin{array}{ccccccc}
\mathfrak{A}_1 & \subsetneq & \mathfrak{A}_2 & \subsetneq & \cdots & \mathfrak{A}_k & \cdots \subsetneq \bigcup_k \mathfrak{A}_k \subsetneq \mathfrak{A}_{unb} = \text{All classes of structures} \\
\parallel & & \cup & & & \cup & \\
\mathfrak{N}_1 & \subsetneq & \mathfrak{N}_2 & \subsetneq & \cdots & \mathfrak{N}_k & \cdots \subsetneq \bigcup_k \mathfrak{N}_k
\end{array}$$

■ **Figure 1** Summary of our results regarding the relative expressive power of different types of left query algorithms over \mathbb{N} . Here, \mathfrak{N}_k , \mathfrak{A}_k , and \mathfrak{A}_{unb} denote the collections of classes of structures that admit a non-adaptive left k -query algorithm, an adaptive left k -query algorithm, and an adaptive left unbounded query algorithm, respectively. No inclusions hold other than those depicted. In particular, $\mathfrak{A}_2 \not\subseteq \bigcup_k \mathfrak{N}_k$ and $\mathfrak{N}_{k+1} \not\subseteq \mathfrak{A}_k$.

Our main contributions are along two lines:

1. We exploit the symmetry of directed cycles and cycle-like structures to give a simple formula (Lemma 4.2) for the number of homomorphisms from a structure to a disjoint union of cycle-like structures of the same size. This formula serves as a powerful tool for deriving homomorphism indistinguishability results for finite profiles, and may be useful more broadly in related contexts. We use it to systematically investigate and compare the expressive power of adaptive and non-adaptive query algorithms over \mathbb{N} for relational structures. For instance, we show, in Theorem 4.5, that for every signature containing a non-unary relation, there exists a class of structures of that signature that is not decided by an adaptive left k -query algorithm over \mathbb{N} for any k . A concrete example of such a class is the class of directed graphs whose number of connected components is an even power of 2. This result is in sharp contrast to the result of Chen et al., which states that every class of simple graphs can be decided by an adaptive left 3-query algorithm over \mathbb{N} . We also show in Theorem 4.7 that, for each $k > 1$, there exists a class that is decided by a non-adaptive left k -query algorithm over \mathbb{N} but not by any adaptive left $(k - 1)$ -query algorithm over \mathbb{N} . However, in Theorem 4.8 we also note that there exists a class that is decided by an adaptive left 2-query algorithm but not by any non-adaptive left query algorithm. We thus establish all inclusion and non-inclusion relations between the expressive powers of adaptive and non-adaptive left k -query algorithms over \mathbb{N} (note: it follows from the proof of Lovász’s theorem that every class is decided by an adaptive left unbounded query algorithm). The results are summarized in Figure 1.

Interestingly, the situation is quite different for *right* query algorithms over \mathbb{N} . Indeed, it follows from results in [22] that every class of structures can be decided with only two adaptive right queries. This is again in contrast to a prior result of Chen et al. [7], namely that there is no k such that every class of simple graphs can be decided by an adaptive right k -query algorithm (when the queries themselves are also required to be simple graphs).

2. We investigate the expressive power of adaptive query algorithms over \mathbb{B} . We observe that adaptive *bounded* query algorithms over \mathbb{B} have the same expressive power as *non-adaptive* query algorithms over \mathbb{B} , whose expressive power was characterized by ten Cate et al. [19]. However, we show that adaptive unbounded query algorithms over \mathbb{B} are strictly more expressive. The classes decided by such algorithms exhibit an intricate pattern, which we show can be captured by a topological characterization. In the case of homomorphism-closed classes, we further show that the characterization can be simplified significantly, yielding an elegant condition for establishing whether such a class is decided by an adaptive unbounded query algorithm over \mathbb{B} or not.

We also apply the topological characterization to examine the expressive power in the special cases of classes that are homomorphism equivalence types, CSPs, or Datalog-definable. Specifically, we show that unboundedness does not increase the expressive power of adaptive left query algorithms over \mathbb{B} for deciding CSPs and homomorphic-equivalence types, thereby yielding, based on existing results, an effective criterion for determining whether such classes are decided by an adaptive left unbounded query algorithm over \mathbb{B} . Additionally, we characterize the Datalog definable classes that are decided by such an algorithm in terms of their upper envelopes. Finally, we show that analogous results hold for adaptive *right* unbounded query algorithms over \mathbb{B} .

Outline. Section 2 reviews preliminary definitions and background, including those of relational structures, digraphs, and walks in digraphs. Section 3 introduces the main concepts studied in this paper. We review the definition of non-adaptive query algorithms and introduce a new definition of adaptive query algorithms that applies in both the bounded and unbounded settings. Section 4 contains our main results regarding the expressive power of query algorithms over \mathbb{N} . Section 4.1 begins with the statement of a key lemma and then builds towards the result that not all classes can be expressed by an adaptive left bounded query algorithm over \mathbb{N} . In Section 4.2, we establish the results comparing adaptive and non-adaptive algorithms (as depicted in Figure 1). Section 4.3 concisely describes the simpler situation for adaptive right query algorithms over \mathbb{N} . Section 5 presents our results regarding the expressive power of query algorithms over \mathbb{B} . In Section 5.1, we show that unboundedness increases the expressive power of adaptive left query algorithms over \mathbb{B} . Section 5.2 provides the topological characterization of the classes expressed by adaptive left unbounded query algorithms over \mathbb{B} , and in Section 5.3 we apply this characterization to special cases. In Section 5.4, we show that similar results can be obtained for adaptive right query algorithms over \mathbb{B} . We conclude in Section 6 by listing directions for further research.

2 Preliminaries

Relational structures and homomorphism counts. A *signature* $\tau = \{R_1, \dots, R_n\}$ is a set of relational symbols where each symbol R_i has an associated arity r_i . A *relational structure* \mathbf{A} of signature τ consists of a non-empty set A and a relation $R_i^{\mathbf{A}} \subseteq A^{r_i}$ for each symbol $R_i \in \tau$. We use $\text{FIN}(\tau)$ to denote the class of all finite relational structures with signature τ . We will only study finite relational structures; we will thus write “structure” and mean “finite relational structure” throughout this text. If $\mathbf{A} = (A, (R_i^{\mathbf{A}})_{i \in \{1, \dots, n\}})$ and $\mathbf{B} = (B, (R_i^{\mathbf{B}})_{i \in \{1, \dots, n\}})$ are structures with signature τ , we say that $\varphi : A \rightarrow B$ is a *homomorphism* from \mathbf{A} to \mathbf{B} if it preserves all the relations, i.e. if

$$(a_1, \dots, a_{r_i}) \in R_i^{\mathbf{A}} \implies (\varphi(a_1), \dots, \varphi(a_{r_i})) \in R_i^{\mathbf{B}}$$

for each $R_i \in \tau$. We then write $\varphi : \mathbf{A} \rightarrow \mathbf{B}$. If there exists a homomorphism from \mathbf{A} to \mathbf{B} we write $\mathbf{A} \rightarrow \mathbf{B}$, otherwise, $\mathbf{A} \nrightarrow \mathbf{B}$. If $\mathbf{A} \rightarrow \mathbf{B}$ and $\mathbf{B} \rightarrow \mathbf{A}$, we then write $\mathbf{A} \leftrightarrow \mathbf{B}$ and say that \mathbf{A} and \mathbf{B} are homomorphically equivalent.

For structures A and B , we let $\text{hom}(A, B)$ denote the number of homomorphisms from A to B . We also use $\text{hom}_{\mathbb{N}}(A, B)$ to denote this same number. We then define

$$\text{hom}_{\mathbb{B}}(A, B) := \begin{cases} 0 & \text{if } \text{hom}(A, B) = 0 \\ 1 & \text{if } \text{hom}(A, B) > 0 \end{cases}$$

where \mathbb{B} denotes the Boolean semiring. For a class \mathcal{C} of structures, the *left homomorphism profile of A restricted to \mathcal{C}* is the tuple $(\text{hom}(C, A))_{C \in \mathcal{C}}$. Similarly, we can define the right homomorphism profile and the homomorphism profile over \mathbb{B} . When we speak of a class of structures, we will always assume it is closed under isomorphism.

Properties of structures. The incidence multigraph $\text{inc}(\mathbf{A})$ of a structure $\mathbf{A} = (A, (R^{\mathbf{A}})_{R \in \tau})$ is the bipartite multigraph whose parts are the sets A and $\text{facts}(\mathbf{A}) = \{(R, \mathbf{t}) : R \in \tau \text{ and } \mathbf{t} \in R^{\mathbf{A}}\}$, and there is an edge between a and (R, \mathbf{t}) for each entry in \mathbf{t} that is a . We say that \mathbf{A} is connected if $\text{inc}(\mathbf{A})$ is connected, and we say \mathbf{A} is acyclic if $\text{inc}(\mathbf{A})$ is acyclic. In particular, if an element appears multiple times in a fact, then the structure is *not* acyclic. This notion of acyclicity is also known as *Berge-acyclicity*. The number $c(\mathbf{A})$ of components of \mathbf{A} is the maximal n such that A can be written as the disjoint union of n structures.

For given structures A and B we let $A \oplus B$ denote their disjoint union and for a natural number m and a structure H we write

$$m \cdot H := \underbrace{H \oplus \dots \oplus H}_{m\text{-times}}.$$

Digraphs, walks, and homomorphisms to cycles. A *digraph* (or a *directed graph*) is a structure with exactly one binary relation. So it is a pair (V, R) where $R \subseteq V \times V$. A *simple graph* is an undirected graph without loops. We let C_n denote the directed cycle of length n :

$$C_n := (\{0, \dots, n-1\}, \{(0, 1), \dots, (n-2, n-1), (n-1, 0)\})$$

and P_n denote the directed path of length n :

$$P_n := (\{0, \dots, n\}, \{(0, 1), \dots, (n-1, n)\})$$

For our discussion of adaptive left query algorithms, walks in directed graphs play a large role. We will thus review the definition of an oriented walk and some related concepts.

► **Definition 2.1.** Let $A = (V, R)$ be a digraph.

(i) An oriented walk in A is a sequence $(a_0, r_0, a_1, r_1, \dots, a_n, r_n, a_{n+1})$ where $a_i \in V$ for each $i \in \{0, \dots, n\}$, $r_i \in \{-, +\}$ and if $r_i = +$ then $(a_i, a_{i+1}) \in R$ while if $r_i = -$ then $(a_{i+1}, a_i) \in R$.

(ii) The net length of an oriented walk $W = (a_0, r_0, a_1, r_1, \dots, a_n, r_n, a_{n+1})$ is defined as:

$$l(W) := |\{i : i \in \{0, \dots, n\} \text{ and } r_i = +\}| - |\{i : i \in \{0, \dots, n\} \text{ and } r_i = -\}|.$$

(iii) The oriented walk $(a_0, r_0, a_1, r_1, \dots, a_n, r_n, a_{n+1})$ is said to be closed if $a_0 = a_{n+1}$.

(iv) A closed oriented walk $(a_0, r_0, a_1, r_1, \dots, a_n, r_n, a_0)$ is called an oriented cycle if $a_i \neq a_j$ for all $i \neq j$ and $n \geq 0$.

(v) A directed walk is an oriented walk (a_0, r_0, \dots, a_n) such that $r_i = +$ for each i . Closed directed walks and directed cycles are defined analogously.

An important thing to note is that oriented walks and their net lengths are preserved under homomorphisms. So if $(a_0, r_0, a_1, r_1, \dots, a_n, r_n, a_{n+1})$ is an oriented walk in A and $\varphi : A \rightarrow B$ is a homomorphism then

$$(\varphi(a_0), r_0, \varphi(a_1), r_1, \dots, \varphi(a_n), r_n, \varphi(a_{n+1}))$$

is an oriented walk in B with the same net length.

For our proof of the existence of a class that is not decided by an adaptive left unbounded query algorithm, we take advantage of the fact that the existence of a homomorphism to a directed cycle can be described in an easy way.

► **Lemma 2.2** (Corollary 1.17 in [13]). *Let A be a digraph and n be a positive natural number. We have $A \rightarrow C_n$ if and only if n divides the net length of every closed oriented walk in A .*

We use \gcd to denote the greatest common divisor of a collection of natural numbers. For every finite digraph, we can define a parameter very related to the condition in Lemma 2.2.

► **Definition 2.3.** *For a finite digraph A we define*

$$\gamma(A) := \gcd(l_1, \dots, l_k)$$

where l_1, \dots, l_k is a listing of the net lengths of all cycles of positive net length in A .

Note, here we use the convention $\gcd(\emptyset) = 0$. For integers n, m we write $n \mid m$ if n divides m , and $n \nmid m$ otherwise. The next proposition relates this parameter to the condition in Lemma 2.2.

► **Proposition 2.4.** *Let A be a digraph and n be a positive integer. We have $A \rightarrow C_n$ if and only if $n \mid \gamma(A)$.*

Conjunctive queries. A Boolean conjunctive query (CQ) is a first-order logic (FO) sentence of the form $q = \exists \mathbf{x}(\bigwedge_i \alpha_i)$ such that each α_i is of the form $R(y_1, \dots, y_n)$ where each y_i is among the variables in \mathbf{x} . It is well known that every such CQ has an associated “canonical structure” A_q such that for every structure B we have $B \models q$ if and only if $A_q \rightarrow B$. Similarly, vice versa, for every (finite) structure A , there is a Boolean CQ q_A such that for every structure B we have $A \rightarrow B$ if and only if $B \models q_A$. A CQ is said to be Berge-acyclic if its canonical structure is acyclic. A (Boolean) *union of conjunctive queries* (UCQ) is a finite disjunction of (Boolean) CQs. Every Boolean UCQ is a positive existential FO sentence, and, conversely, every positive existential FO sentence is equivalent to a UCQ.

3 Definitions of Query Algorithms

First, we review the definition of non-adaptive query algorithms.

► **Definition 3.1.** *Let \mathcal{C} be a class of relational structures with the same signature, let $\mathcal{K} \in \{\mathbb{N}, \mathbb{B}\}$, and let k be a positive integer.*

- *A non-adaptive left k -query algorithm over \mathcal{K} for \mathcal{C} is a pair (\mathcal{F}, X) , where $\mathcal{F} = (F_1, \dots, F_k)$ is a tuple of relational structures with the same signature as the structures in \mathcal{C} , and X is a set of k -tuples over \mathcal{K} , such that for all structures D , we have $D \in \mathcal{C}$ if and only if*

$$(\text{hom}_{\mathcal{K}}(F_1, D), \text{hom}_{\mathcal{K}}(F_2, D), \dots, \text{hom}_{\mathcal{K}}(F_k, D)) \in X$$

- *A non-adaptive right k -query algorithm over \mathcal{K} for \mathcal{C} is a pair (\mathcal{F}, X) , where $\mathcal{F} = (F_1, \dots, F_k)$ is a tuple of relational structures with the same signature as the structures in \mathcal{C} , and X is a set of k -tuples over \mathcal{K} , such that for all structures D , we have that $D \in \mathcal{C}$ if and only if*

$$(\text{hom}_{\mathcal{K}}(D, F_1), \text{hom}_{\mathcal{K}}(D, F_2), \dots, \text{hom}_{\mathcal{K}}(D, F_k)) \in X.$$

► **Example 3.2.** Let $\mathcal{F} := \{F\}$ where $F = (\{0\}, \emptyset)$ is the singleton digraph with no edges, and let $X = \{3\}$. Then the non-adaptive left 1-query algorithm (\mathcal{F}, X) over \mathbb{N} decides the class \mathcal{C} of digraphs that have exactly 3 vertices. It is not difficult to see that \mathcal{C} does not admit a non-adaptive left query algorithm over \mathbb{B} . Indeed, this follows from the fact that the class is not closed under homomorphic equivalence. A query algorithm over \mathbb{B} cannot distinguish between structures that are homomorphically equivalent, since if $A \leftrightarrow B$ then $D \rightarrow A$ if and only if $D \rightarrow B$ for any D .

The example above shows that query algorithms over \mathbb{N} have more expressive power than query algorithms over \mathbb{B} in general. However, for classes closed under homomorphic equivalence, the same does not hold, as shown by ten Cate et al. in the following theorem.

► **Theorem 3.3** (Corollary 5.6 in [19]). *Let \mathcal{C} be a class of structures that is closed under homomorphic equivalence. Then the following are equivalent:*

- \mathcal{C} admits a non-adaptive left k -query algorithm over \mathbb{N} for some k .
- \mathcal{C} admits a non-adaptive left k -query algorithm over \mathbb{B} for some k .
- \mathcal{C} is FO-definable.

We now formulate the definition of an *adaptive unbounded query algorithm*. Adaptive query algorithms were defined in Chen et al. [7]. The definition there is limited in the sense that it only considers query algorithms whose number of queries is bounded by some constant. For a set Σ , we let $\Sigma^{<\omega}$ denote the set of finite strings with alphabet Σ .

- For $\sigma, \sigma' \in \Sigma^{<\omega}$, we write $\sigma' \sqsubseteq \sigma$ if σ' is an initial segment of σ .
- We let $\sigma \bullet \sigma'$ denote the concatenation of the strings σ, σ' (we also use this notation if σ' is an element of Σ).
- If $\sigma_0 \sqsubseteq \sigma_1 \sqsubseteq \sigma_2 \sqsubseteq \dots$ is a sequence, then we let $\bigsqcup_{n \geq 0} \sigma_n$ denote the (possibly infinite) string σ with length $m := \sup_{n \geq 0} |\sigma_n|$ such that for $i \leq m$, the i -th letter of σ is the i -th letter of σ_n for any n such that $|\sigma_n| \geq i$.

A set $\mathcal{T} \subseteq \Sigma^{<\omega}$ that is closed under initial segments (so if $\sigma \in \mathcal{T}$ and $\sigma' \sqsubseteq \sigma$ then $\sigma' \in \mathcal{T}$) is called a *subtree of $\Sigma^{<\omega}$* . An element σ of such a subtree \mathcal{T} is called a *leaf* if for every $\sigma' \in \mathcal{T}$ such that $\sigma \sqsubseteq \sigma'$ we have $\sigma' = \sigma$.

We are now ready to define adaptive left unbounded query algorithms.

► **Definition 3.4.** *Let $\mathcal{K} \in \{\mathbb{N}, \mathbb{B}\}$.*

- (i) *An adaptive left unbounded query algorithm over \mathcal{K} for structures with signature τ is a function*

$$G : \mathcal{T} \rightarrow \text{FIN}(\tau) \cup \{\text{YES}, \text{NO}\}$$

where \mathcal{T} is a subtree of $\mathcal{K}^{<\omega}$ and $G(\sigma) \in \{\text{YES}, \text{NO}\}$ if and only if σ is a leaf in \mathcal{T} .

- (ii) *For $A \in \text{FIN}(\tau)$, the computation path of an adaptive left unbounded query algorithm G on A is the string defined by $\sigma(A, G) := \bigsqcup_{n \geq 0} \sigma_n$ where*

$$\begin{aligned} \sigma_0 &= \varepsilon \\ \sigma_{n+1} &= \begin{cases} \sigma_n & \text{if } G(\sigma_n) \in \{\text{YES}, \text{NO}\} \\ \sigma_n \bullet \text{hom}_{\mathcal{K}}(G(\sigma_n), A) & \text{if } G(\sigma_n) \in \text{FIN}(\tau) \end{cases} \end{aligned}$$

The algorithm G halts on input A if $\sigma(A, G)$ is finite.

- (iii) *If G halts on all $A \in \text{FIN}(\tau)$, then G is said to be total.*
 (iv) *If G is total, we say that it decides the class*

$$\mathcal{C} := \{A \in \text{FIN}(\tau) : G(\sigma(A, G)) = \text{YES}\}.$$

- (v) We say that G is an adaptive left k -query algorithm if for every A , $\sigma(A, G)$ has length at most k . We say G is bounded if it is an adaptive left k -query algorithm for some k .

The notion of an adaptive right unbounded/bounded query algorithm G is defined analogously by appending $\text{hom}_K(A, G(\sigma_n))$ instead of $\text{hom}_K(G(\sigma_n), A)$ in the computation path.

We will only consider total adaptive query algorithms, so when we speak of adaptive query algorithms we always assume it is total.

► **Example 3.5.** Let G be an adaptive left query algorithm over \mathbb{N} for digraphs defined with the following:

$$G(\sigma) := \begin{cases} (\{0\}, \emptyset) & \text{if } \sigma = \varepsilon \\ P_n & \text{if } \sigma = 'n' \\ \text{YES} & \text{if } \sigma = 'a_0 a_1' \text{ and } a_1 \neq 0 \\ \text{NO} & \text{if } \sigma = 'a_0 a_1' \text{ and } a_1 = 0 \end{cases}.$$

The corresponding tree $\mathcal{T} \subseteq \mathbb{N}^{<\omega}$ would here be the set of all strings of length at most 2.

We can describe a run of the algorithm on a target structure A as follows: First, the algorithm queries for the number of homomorphisms from the singleton digraph with no edges to A . Let n denote the output of that query. It then queries for $\text{hom}_{\mathbb{N}}(P_n, A)$. If the result is positive, it accepts; otherwise, it rejects.

To see what class this algorithm decides, we first note that the number of homomorphisms from the singleton digraph with no edges to A is precisely the number of vertices of A . We also have that there exists a homomorphism from P_n to A if and only if A has a directed walk of length n . Since n is the size of A , we see that such a walk has to visit the same vertex twice. We thus get that such a walk exists if and only if A has a directed cycle. The algorithm, therefore, decides the class of digraphs that contain a directed cycle.

► **Example 3.6.** In this example, we show that every class \mathcal{C} of structures is decided by an adaptive left unbounded query algorithm over \mathbb{N} .

The algorithm first queries for the size of the structure. Then it queries the number of homomorphisms from every structure that is not larger than it. By the proof of Lovász's theorem (see, for example, Section III of [1] for a proof of this), the algorithm has now distinguished the structure (up to isomorphism) and can thus classify it correctly.

To formally define the algorithm, let τ be the signature of the structures in \mathcal{C} . We fix an enumeration A_1, A_2, A_3, \dots of all structures of signature τ such that for each $n \in \mathbb{N}$ there exists a number r_n such that A_1, \dots, A_{r_n} is an enumeration of all of the structures of size at most n . Then the algorithm can be defined with:

$$G(\sigma) := \begin{cases} (\{0\}, \emptyset) & \text{if } \sigma = \varepsilon \\ A_{k+1} & \text{if } \sigma = 'na_1 a_2 \dots a_k' \text{ and } k < r_n \\ \text{YES} & \text{if } \sigma = 'na_1 a_2 \dots a_{r_n}' \text{ and the unique } A \text{ s.t. } \text{hom}_{\mathbb{N}}(A_i, A) = a_i \\ & \text{for each } 1 \leq i \leq r_n \text{ satisfies } A \in \mathcal{C} \\ \text{NO} & \text{if } \sigma = 'na_1 a_2 \dots a_{r_n}' \text{ and the unique } A \text{ s.t. } \text{hom}_{\mathbb{N}}(A_i, A) = a_i \\ & \text{for each } 1 \leq i \leq r_n \text{ satisfies } A \notin \mathcal{C} \end{cases}$$

4 Query Algorithms over \mathbb{N}

4.1 Unboundedness helps for adaptive left query algorithms

A starting point for our work in this section is the following result by Chen et al:

► **Theorem 4.1** (Theorem 8.2 in [7]). *Every class of simple graphs can be decided by an adaptive left 3-query algorithm over \mathbb{N} .*

This result invites the question of what happens in the general case, for arbitrary relational structures. We prove that the same does not hold in this broader setting. We even show that there is a class of directed graphs that is not decided by an adaptive k -query algorithm over \mathbb{N} , for any k . To do this, we utilize the symmetry of directed cycles to prove that they are hard to distinguish from each other using left homomorphism queries.

We begin with a lemma that shows that the homomorphism count from a structure to a disjoint union of cycles can be described simply.

► **Lemma 4.2.** *For every digraph A and all positive integers n, m we have*

$$\text{hom}(A, m \cdot C_n) := \begin{cases} 0 & \text{if } n \nmid \gamma(A) \\ (m \cdot n)^{c(A)} & \text{if } n \mid \gamma(A) \end{cases}.$$

The proof of this theorem can be intuitively explained as follows: In disjoint unions of directed cycles, each vertex has in-degree 1 and out-degree 1. For a homomorphism from $A \rightarrow m \cdot C_n$, the whole map is therefore decided by the value of the map in one vertex from each component of A . Since there are $m \cdot n$ ways to pick the value of a homomorphism on that initial vertex (if such a homomorphism exists), the lemma follows.

We say that two classes \mathcal{A}, \mathcal{B} of structures are separated by a query algorithm if it accepts all structures in \mathcal{A} and rejects all structures in \mathcal{B} , or vice versa.

The importance of the above lemma is in the fact that it shows that for a given structure $m \cdot C_n$ there are only two possible outcomes for $\text{hom}(F, m \cdot C_n)$. This limits what can be done using few queries, as is shown in the following theorem:

► **Theorem 4.3.** *The class $\mathcal{D}_n := \{2^{n-m} \cdot C_{2^m} : 0 \leq m \leq n \text{ and } m \text{ is even}\}$ can be separated from the class $\mathcal{D}'_n := \{2^{n-m} \cdot C_{2^m} : 0 \leq m \leq n \text{ and } m \text{ is odd}\}$ by a non-adaptive left k -query algorithm over \mathbb{N} if and only if $k \geq n$.*

Proof. By Lemma 4.2 we have that for every digraph F and every $2^{n-m} \cdot C_{2^m} \in \mathcal{D}_n \cup \mathcal{D}'_n$:

$$\begin{aligned} \text{hom}(F, 2^{n-m} \cdot C_{2^m}) &= \begin{cases} 0 & \text{if } 2^m \nmid \gamma(F) \\ (2^n)^{c(F)} & \text{if } 2^m \mid \gamma(F) \end{cases} \\ &= \begin{cases} 0 & \text{if } \nu_2(\gamma(F)) < m \\ (2^n)^{c(F)} & \text{if } \nu_2(\gamma(F)) \geq m \end{cases} \end{aligned}$$

where $\nu_2(k)$ denotes the largest positive integer r such that $2^r \mid k$ if $k \neq 0$, if $k = 0$ then $\nu_2(k) = +\infty$. It follows from this that to have $\text{hom}(F, 2^{n-m} \cdot C_{2^m}) \neq \text{hom}(F, 2^{n-(m+1)} \cdot C_{2^{m+1}})$ we must have $\nu_2(\gamma(F)) = m + 1$. To distinguish between $2^{n-m} \cdot C_{2^m}$ and $2^{n-(m+1)} \cdot C_{2^{m+1}}$ for every $m \in \{0, \dots, n-1\}$ the algorithm then needs to query some F_m with $\nu_2(\gamma(F_m)) = m$ for each $m \in \{0, \dots, n-1\}$. The algorithm, therefore, needs at least n queries to separate the classes. It is clear that n non-adaptive queries suffice: the algorithm can query C_{2^m} for each $m \in \{0, \dots, n-1\}$. Every structure in $\mathcal{D}_n \cup \mathcal{D}'_n$ yields a unique outcome from the collection of these queries, and thus they can be classified correctly.

It is worth noting that homomorphism existence queries suffice here, and hence \mathcal{D}_n and \mathcal{D}'_n are already separated by a non-adaptive left n -query algorithm over \mathbb{B} . ◀

► **Corollary 4.4.** *The class \mathcal{D}_n can be separated from the class \mathcal{D}'_n by an adaptive left k -query algorithm over \mathbb{N} if and only if $k \geq \log(n+1)$.*

Proof. For any digraph $H \in \mathcal{D}_n \cup \mathcal{D}'_n$, there are only two possible outcomes for every query $\text{hom}(F, H)$ (as explained above). Therefore, an adaptive left k -query algorithm over \mathbb{N} that separates \mathcal{D}_n from \mathcal{D}'_n queries at most

$$2^0 + 2^1 + \dots + 2^{k-1} = 2^k - 1$$

different structures in all of its possible computation paths on inputs from $\mathcal{D}_n \cup \mathcal{D}'_n$. It can thus be translated into a non-adaptive left $(2^k - 1)$ -query algorithm over \mathbb{N} separating the two classes. By Theorem 4.3 it follows that such a non-adaptive left $(2^k - 1)$ -query algorithm over \mathbb{N} exists if and only if $2^k - 1 \geq n$, so $k \geq \log(n + 1)$. Thus, we have shown that the classes $\mathcal{D}_n, \mathcal{D}'_n$ can only be separated by an adaptive left k -query algorithm over \mathbb{N} if $k \geq \log(n + 1)$. It is also clear that $k \geq \log(n + 1)$ suffices. The algorithm can use queries of the form $\text{hom}(\mathbf{C}_{2^r}, 2^{n-m} \cdot \mathbf{C}_{2^m})$ to binary search for the m of an input structure of the form $2^{n-m} \cdot \mathbf{C}_{2^m}$, and use that to classify it correctly. ◀

Let \mathfrak{N}_k denote the set of classes that are decided by a non-adaptive left k -query algorithm over \mathbb{N} , and \mathfrak{A}_k denote the set of classes that are decided by an adaptive left k -query algorithm over \mathbb{N} . Theorem 4.3 and Corollary 4.4 show, among other things, that

$$\mathfrak{N}_1 \subsetneq \mathfrak{N}_2 \subsetneq \mathfrak{N}_3 \subsetneq \dots \quad \text{and} \quad \mathfrak{A}_1 \subsetneq \mathfrak{A}_2 \subsetneq \mathfrak{A}_3 \subsetneq \dots$$

Using Corollary 4.4, we can also prove an even stronger result:

► **Theorem 4.5.** *Let \mathcal{C} be the class of digraphs whose smallest directed cycle (if it exists) has length that is an even power of two. Then \mathcal{C} is not decided by an adaptive left k -query algorithm over \mathbb{N} for any k .*

Proof. Note that we have

$$\bigcup_{n>0} \mathcal{D}_n \subseteq \mathcal{C} \subseteq \left(\bigcup_{n>0} \mathcal{D}'_n \right)^c,$$

where X^c denotes the set theoretic complement of a set X . A given adaptive left k -query algorithm over \mathbb{N} cannot separate \mathcal{D}_{2^k} from \mathcal{D}'_{2^k} , by Corollary 4.4. To decide \mathcal{C} it must in particular separate these classes, therefore the algorithm can not decide \mathcal{C} . Thus it is clear that \mathcal{C} is not decided by any adaptive left k -query algorithm over \mathbb{N} for any k . ◀

Using the same method it can be shown that many other classes are not decided by an adaptive left k -query algorithm over \mathbb{N} for any k . In fact, to show that a class is not decided by such an algorithm it suffices to show that it separates \mathcal{D}_n and \mathcal{D}'_n for infinitely many n . Another example of such a class is the class of digraphs whose number of components is an even power of two.

► **Note.** As we saw in Example 3.6, every class is decided by an adaptive left unbounded query algorithm over \mathbb{N} . Theorem 4.5 therefore also shows that unboundedness increases the expressive power of adaptive left query algorithms over \mathbb{N} , i.e. $\bigcup_k \mathfrak{A}_k \subsetneq \mathfrak{A}_{unb}$ where \mathfrak{A}_{unb} denotes the class of structures decided by an adaptive left unbounded query algorithm over \mathbb{N} .

The preceding results show that there exist classes of digraphs that adaptive left bounded query algorithms over \mathbb{N} do not decide. It follows immediately that for every signature containing a binary relation, there exists a class of structures of that signature that left bounded query algorithms over \mathbb{N} cannot decide. For signatures containing only unary relations, it can be shown that every class can be decided with an adaptive left bounded query algorithm over \mathbb{N} . What about signatures that contain some n -ary relation with $n > 2$ and no binary relation? The following theorem answers this question.

► **Theorem 4.6.** *For every $n \geq 2$ and every signature τ containing an n -ary relation, there exists a class of structures with signature τ that is not decided by an adaptive left k -query algorithm over \mathbb{N} for any k .*

The class we construct in the proof of this theorem is similar to the class from Theorem 4.5, except we replace the directed cycles with n -ary “cycle-like” structures.

4.2 Adaptive versus non-adaptive left k -query algorithms

We now turn to a comparison of the expressive power of adaptive and non-adaptive left query algorithms over \mathbb{N} . It follows from Theorem 4.3 and Corollary 4.4 that $\mathfrak{N}_{2k} \not\subseteq \mathfrak{A}_k$. We also trivially have the relation $\mathfrak{N}_k \subseteq \mathfrak{A}_k$. This raises the question of where the bound lies, that is, for which l we have $\mathfrak{N}_l \subseteq \mathfrak{A}_k$ but $\mathfrak{N}_{l+1} \not\subseteq \mathfrak{A}_k$? This is answered by the following theorem:

► **Theorem 4.7.** *For each k , there exists a class of structures that is decided by a non-adaptive left k -query algorithm over \mathbb{N} but not by an adaptive left $(k-1)$ -query algorithm over \mathbb{N} .*

Proof. Let p_1, \dots, p_{2k} be distinct primes and write $P := \prod_{i=1}^{2k} p_i$ and $q_i := P/p_i$. Define a non-adaptive left k -query algorithm $((F_1, \dots, F_k), X)$ with

$$F_i := p_i \cdot C_{q_i} \text{ and } X := \{(P \cdot \delta_{1,j}, \dots, P \cdot \delta_{k,j}) : j \in \{1, \dots, k\}\}$$

where $\delta_{i,j} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases}$ is the Kronecker delta. Let \mathcal{C} be the class decided by the above query algorithm. It follows from Lemma 4.2 that

$$\text{hom}(F_i, p_j \cdot C_{q_j}) = P \cdot \delta_{i,j},$$

so for $j \in \{1, \dots, 2k\}$ we have $p_j \cdot C_{q_j} \in \mathcal{C}$ if and only if $j \leq k$. We will show that \mathcal{C} cannot be decided by an adaptive left $(k-1)$ -query algorithm over \mathbb{N} .

Let G be a given adaptive left $(k-1)$ -query algorithm over \mathbb{N} . Using an adversarial argument, we find two structures that are classified in the same way by G but differently by \mathcal{C} . We do this by defining sets \mathcal{A}_n such that G has the same computation path on all elements of \mathcal{A}_n up to stage n . We define $\mathcal{A}_0 := \{p_j \cdot C_{q_j} : j \in \{1, \dots, 2k\}\}$. Now let $\mathcal{A}_n \subseteq \mathcal{A}_0$ be given such that G has the same computation path σ on all elements of \mathcal{A}_n up to stage n . Write $F := G(\sigma)$. We have two cases:

- If $\text{hom}(F, A)$ is the same for all elements A of \mathcal{A}_n we can simply define $\mathcal{A}_{n+1} := \mathcal{A}_n$ and the computation path is the same up to stage $n+1$.
- Assume otherwise. Note that for any element $p_j \cdot C_{q_j}$ of \mathcal{A}_0 we have by Lemma 4.2 that

$$\text{hom}(F, p_j \cdot C_{q_j}) = \begin{cases} 0 & \text{if } q_j \nmid \gamma(F) \\ P^{c(F)} & \text{if } q_j \mid \gamma(F). \end{cases}$$

So there are only two possible outcomes. Moreover, if $q_i, q_j \mid \gamma(F)$ for $i \neq j$, then $P \mid \gamma(F)$ and thus $q_r \mid \gamma(F)$ for all $r \in \{1, \dots, 2k\}$ and we are in the former case. Thus, there is an i such that $q_j \mid \gamma(F)$ if and only if $j = i$. We can thus define $\mathcal{A}_{n+1} := \mathcal{A}_n \setminus \{p_i \cdot C_{q_i}\}$. It is clear that $\text{hom}(F, A) = 0$ for all elements A of \mathcal{A}_{n+1} , so they share the same computation path up to stage $n+1$.

The above construction gives us a set \mathcal{A}_{k-1} such that G has the same computation path up to stage $k-1$ on all elements of it. Clearly we also have that $|\mathcal{A}_{k-1}| \geq |\mathcal{A}_0| - (k-1) = k+1$. This means that there must exist $p_i \cdot C_{q_i}, p_j \cdot C_{q_j} \in \mathcal{A}_{k-1}$ such that $i \leq k$ and $j \geq k+1$.

Thus, we have found elements that G classifies in the same way but \mathcal{C} classifies differently. We have thus shown that G does not decide \mathcal{C} . ◀

The above theorem describes a class of structures where adaptiveness does not help at all when trying to decide the class; therefore, we have that $\mathfrak{N}_{k+1} \not\subseteq \mathfrak{A}_k$ for each k . This is interesting as in many cases, adaptiveness greatly reduces the number of queries needed.

In this regard, we now prove that there is an adaptive left 2-query algorithm over \mathbb{N} that decides a class that no non-adaptive left query algorithm over \mathbb{N} decides.

► **Theorem 4.8.** *The class of all digraphs that contain a directed cycle is decided by an adaptive left 2-query algorithm over \mathbb{N} , but not by any non-adaptive left k -query algorithm over \mathbb{N} , for any k .*

Proof. Note that directed cycles are preserved by homomorphisms. It then follows that the class is closed under homomorphic equivalence. It is also well known that this class is not first-order definable; this is a standard application of the Ehrenfeucht-Fraïssé method, which can be found in textbooks on finite model theory [9]. It then follows from Theorem 3.3 that the class is not decided by a non-adaptive left k -query algorithm over \mathbb{N} , for any k .

In example 3.5 we saw an adaptive left 2-query algorithm over \mathbb{N} that decides the class. This completes the proof. ◀

The preceding results now give the relations between \mathfrak{N}_k and \mathfrak{A}_k shown in Figure 1.

4.3 Adaptive right query algorithms over \mathbb{N}

In their paper, Chen et al. also prove a result about adaptive *right* query algorithms on simple graphs:

► **Proposition 4.9** (Proposition 9.3 in [7]). *There exists a class of simple graphs that is not decided by an adaptive right k -query algorithm over \mathbb{N} for any k , using only simple graphs in its queries.*

In the general case, we again get a different result. It turns out that using looped structures in the queries changes the picture completely. This result was essentially proved by Wu [22], but it is not stated in this form in his work.

► **Theorem 4.10** (From [22]). *Every class of structures can be decided by an adaptive right 2-query algorithm over \mathbb{N} .*

This result also is in contrast to our result for left query algorithms: the hierarchy of expressive power of adaptive *right* k -query algorithms collapses at $k = 2$, in contrast to the strict hierarchy shown in Figure 1 for adaptive *left* k -query algorithms. The cause of this is that for right query algorithms, we have more constructions that give useful information about the homomorphism counts. In particular, we have that if H is connected then

$$\text{hom}(H, A \oplus B) = \text{hom}(H, A) + \text{hom}(H, B).$$

No construction on the left side allows for addition of the hom-vectors.

5 Query Algorithms Over \mathbb{B}

We now move on to studying adaptive query algorithms over \mathbb{B} . We focus on left query algorithms, and we will comment on the case of right query algorithms afterwards.

5.1 Unboundedness helps for adaptive left query algorithms

Every class that can be decided by an adaptive k -query algorithm over \mathbb{B} can also be decided by a non-adaptive $(2^k - 1)$ -query algorithm over \mathbb{B} . This is done by querying all the structures that the algorithm can possibly query in the first k steps of its run. Since there are at most 2 branches from each node in the computation tree, this is at most $1 + 2 + 2^2 + \dots + 2^{k-1} \leq (2^k - 1)$ queries. This means that *adaptive* and *non-adaptive* left bounded query algorithms over \mathbb{B} have the same expressive power.

The question we answer below is whether every class decided by an adaptive unbounded query algorithm over \mathbb{B} can be decided by an adaptive bounded query algorithm (and thus a non-adaptive one)? The answer is no.

► **Theorem 5.1.** *The class $\mathcal{C} = \{A : A \text{ is a digraph that contains a directed cycle}\}$ can be decided by an adaptive left unbounded query algorithm over \mathbb{B} , but not by an adaptive left k -query algorithm over \mathbb{B} , for any k .*

Proof. In theorem 4.8 we showed that \mathcal{C} is not decided by a non-adaptive left k -query algorithm over \mathbb{N} , it is thus not decided by a non-adaptive left k -query algorithm over \mathbb{B} .

We now provide an adaptive left unbounded query algorithm over \mathbb{B} that decides \mathcal{C} . The algorithm runs as follows:

Let A be the given input. For $n = 1, 2, \dots$ query $\text{hom}_{\mathbb{B}}(P_n, A)$ and $\text{hom}_{\mathbb{B}}(C_n, A)$. If at any point $\text{hom}_{\mathbb{B}}(P_n, A) = 0$ then halt and output NO. If at any point $\text{hom}_{\mathbb{B}}(C_n, A) = 1$ then halt and output YES.

To see that this algorithm decides \mathcal{C} we note that the following equivalence holds:

$$\begin{aligned} A \text{ has no directed cycle} &\iff \text{the lengths of directed walks in } A \text{ are bounded} \\ &\quad \text{by a constant } k \\ &\iff \text{hom}_{\mathbb{B}}(P_n, A) = 0 \text{ for some } n. \end{aligned}$$

Also, directed cycles are preserved under homomorphism, so A has a directed cycle if and only if $\text{hom}_{\mathbb{B}}(C_n, A) = 1$ for some n . Thus, it is clear that the algorithm halts on all inputs and produces the correct answer. ◀

5.2 Characterization of the expressive power of adaptive left unbounded query algorithms over \mathbb{B}

We can use the language of topology to characterize the expressive power of adaptive unbounded query algorithms exactly. For background on topology, we refer the reader to [17]. The topology we define is generated by all “basic observations”, where a basic observation is a class defined by the answer to a single homomorphism query.

More formally, for a structure $A \in \text{FIN}(\tau)$ we define

$$\uparrow A := \{B : A \rightarrow B\},$$

which, phrased in terms of *conjunctive queries* (CQs) is simply the set of structures that satisfy the canonical CQ q_A of A . We then let

$$\mathcal{S} := \{\uparrow A : A \in \text{FIN}(\tau)\} \cup \{\text{FIN}(\tau) \setminus \uparrow A : A \in \text{FIN}(\tau)\}$$

be the subbasis for our topology \mathfrak{T} . In other words, phrased again in terms of conjunctive queries, the open sets of the topology are all classes of structures that can be defined by a CQ or the negation of a CQ, as well as all classes generated from these by closing under finite intersection and arbitrary union.

We have the following characterization of classes that are decided by adaptive left unbounded query algorithms over \mathbb{B} .

► **Theorem 5.2.** *A class $\mathcal{C} \subseteq \text{FIN}(\tau)$ is decided by an adaptive left unbounded query algorithm over \mathbb{B} if and only if \mathcal{C} is a clopen set in the topological space $(\text{FIN}(\tau), \mathfrak{T})$.*

Remember, a *clopen* set is a set that is both closed and open. We say that a class \mathcal{C} of structures is homomorphism-closed if whenever $A \in \mathcal{C}$ and $A \rightarrow B$ for some structure B , then $B \in \mathcal{C}$. We can use the topological characterization to get a simpler characterization for the existence of an adaptive left unbounded query algorithm over \mathbb{B} in the case of homomorphism-closed classes.

► **Theorem 5.3.** *Let \mathcal{C} be a homomorphism-closed class; then the following hold:*

- (i) $\mathcal{C} = \bigcup_{A \in \mathcal{C}} \uparrow A$, hence \mathcal{C} is open.
- (ii) Furthermore, \mathcal{C} is closed if and only if $\mathcal{C} = \bigcap_{i \in I} \bigcup_{j=0}^{n_i} \uparrow A_{i,j}$ for some structures $A_{i,j}$.

The preceding two theorems imply that a homomorphism-closed class \mathcal{C} is decided by an adaptive left unbounded query algorithm over \mathbb{B} if and only if \mathcal{C} is an intersection of UCQs.

► **Remark 5.4.** The first part of Theorem 5.1 can be derived from Theorem 5.3 since \mathcal{C} is homomorphism-closed and $\mathcal{C} = \bigcap_{n \geq 1} \mathcal{P}_n$.

5.3 Case studies

As our first case study, we consider the definability of *Constraint Satisfaction Problems* (CSPs) with adaptive left unbounded query algorithms over \mathbb{B} .

A CSP can be formulated as a problem of deciding whether there is a homomorphism from an input structure to a target structure. For a relational structure A we write $\text{CSP}(A) := \{B : B \rightarrow A\}$.

We can combine some known results with our observations from Section 5.2 to get the following result:

► **Theorem 5.5.** *Let \mathcal{C} be any class of the form $[A]_{\leftrightarrow}$ or $\text{CSP}(A)$ for a structure A . Then the following are equivalent:*

- (i) \mathcal{C} is decided by an adaptive left unbounded query algorithm over \mathbb{B} .
- (ii) \mathcal{C} is decided by a non-adaptive left k -query algorithm over \mathbb{B} for some k .

This theorem shows that unboundedness does not help when deciding CSPs or homomorphic equivalence classes, and it gives an effective characterization of the classes of this form that are decided by an adaptive left unbounded query algorithm over \mathbb{B} :

► **Corollary 5.6.** *$\text{CSP}(A)$ is decided by an adaptive left unbounded query algorithm over \mathbb{B} if and only if it is first-order definable. Moreover, there is an effective algorithm that decides for a given structure A whether this holds.*

The fact that $\text{CSP}(A)$ is decided by a non-adaptive left query algorithm over \mathbb{B} if and only if it is first-order definable follows from Theorem 3.3. The fact that there is an algorithm that decides whether $\text{CSP}(A)$ is first-order definable is a well-known result by Larose, Loten, and Tardif [14].

Consider now the database query language Datalog. To conserve space, we omit a detailed definition of the syntax and semantics of Datalog, which can be found in [9]. Here we only consider Boolean Datalog programs, which are Datalog programs with a 0-ary goal predicate. Such a program π defines a class of structures \mathcal{C}_π that is homomorphism-closed. A Datalog program is *monadic* if it only uses recursive predicates (a.k.a. IDBs) that are unary; it is *linear* if every rule body has at most one occurrence of a recursive predicate.

The class \mathcal{C} of digraphs that contain a directed cycle from Theorem 5.1 is Datalog definable, using the program:

$$\begin{aligned} X(x, y) &\leftarrow R(x, y) \\ X(x, y) &\leftarrow X(x, x_1), R(x_1, y) \\ \text{Ans}() &\leftarrow X(z, z) \end{aligned}$$

This shows that unboundedness helps even when we restrict ourselves to Datalog definable classes. Contrastingly, Theorem 5.5 shows that for every Datalog program that defines the complement of a CSP, unboundedness does *not* help. Corollary 5.6 furthermore shows that a class defined by such Datalog programs is decided by an adaptive left unbounded query algorithm if and only if it is first-order definable. Note that the Datalog programs that define the complement of a CSP form a non-trivial fragment of Datalog. Indeed, every (Boolean) monadic Datalog program whose rule bodies are “tree-shaped” defines a complement of a CSP (and moreover, the CSP in question can be constructed effectively from the Datalog program), cf. [10, 3].

We can use these observations to prove a negative result on the expressive power of adaptive left query algorithms over \mathbb{B} .

► **Theorem 5.7.** *There exists a class definable by a monadic linear Datalog program that is not decided by an adaptive left unbounded query algorithm over \mathbb{B} .*

We now give a characterization of the Datalog definable classes that are also decided by an adaptive left unbounded query algorithm in terms of their *upper envelopes*.

Chaudhuri and Kolaitis [4, 5], in their study of approximations of Datalog programs by non-recursive queries, they introduced the notion of an upper envelope. For a Boolean Datalog program deciding a class \mathcal{C} , an upper envelope is a union of conjunctive queries (UCQ) q that defines a class \mathcal{C}' such that $\mathcal{C} \subseteq \mathcal{C}'$. Such an upper envelope can be equivalently viewed as a class \mathcal{C}' of the form $\bigcup_{j=0}^n \uparrow A_j$ that contains \mathcal{C} . Since Datalog definable classes \mathcal{C} are homomorphism-closed, Theorem 5.3 therefore gives us the following:

► **Corollary 5.8.** *A Datalog definable class \mathcal{C} is decided by an adaptive left unbounded query algorithm over \mathbb{B} if and only if it is the intersection of its upper envelopes.*

5.4 Adaptive right query algorithms over \mathbb{B}

For adaptive *right* query algorithms over \mathbb{B} we have a result analogous to Theorem 5.1:

► **Theorem 5.9.** *The class \mathcal{C} of all digraphs that have an oriented cycle with non-zero net length can be decided by an adaptive right unbounded query algorithm over \mathbb{B} , and not by an adaptive right k -query algorithm over \mathbb{B} , for any k .*

It can be shown that this class is Datalog definable. So, as in the case for left query algorithms, unboundedness sometimes helps even when we restrict ourselves to Datalog definable classes.

We can also replicate the results from section 5.2 for right query algorithms. We define \mathfrak{T}' as the topology generated by the sub-basis

$$\mathcal{S}' := \{\downarrow A : A \in \text{FIN}(\tau)\} \cup \{\text{FIN}(\tau) \setminus \downarrow A : A \in \text{FIN}(\tau)\}.$$

where $\downarrow A$ denotes the set $\{B \in \text{FIN}(\tau) : B \rightarrow A\}$. In other words, \mathcal{S}' consists of all CSPs and their complements. Using the same argument as before, we can prove:

► **Theorem 5.10.** *A class $\mathcal{C} \subseteq \text{FIN}(\tau)$ is decided by an adaptive right unbounded query algorithm over \mathbb{B} if and only if \mathcal{C} is a clopen set in the topological space $(\text{FIN}(\tau), \mathfrak{T}')$.*

Then we also get

► **Theorem 5.11.** *Let \mathcal{C} be any class of the form $[A]_{\leftrightarrow}$ or $\uparrow A$ for a structure A . Then the following are equivalent:*

- (i) \mathcal{C} is decided by an adaptive right unbounded query algorithm over \mathbb{B} .
- (ii) \mathcal{C} is decided by a non-adaptive right k -query algorithm over \mathbb{B} for some k .

It then follows from a result by ten Cate et al. [19] that these conditions hold if and only if A is homomorphically equivalent to an acyclic structure. Using the language of conjunctive queries, this shows that a CQ q has an equivalent adaptive right unbounded query algorithm if and only if q is equivalent to a Berge-acyclic CQ. It is also good to note that since every CQ can be written as a Datalog program, this also shows that there are Datalog-definable classes that are not decided by any adaptive right unbounded query algorithm over \mathbb{B} .

6 Open Problems

Adaptive $f(n)$ -query algorithms. We can define an adaptive $f(n)$ -query algorithm as an adaptive unbounded query algorithm that uses at most $f(n)$ queries for inputs of size n . Then one could ask the question of what the growth rate of f needs to be for adaptive left $f(n)$ -query algorithms over \mathbb{N} to be able to decide any class of structures of a given signature. The algorithm described in Example 3.6 gives an upper bound of $2^{\sum_i (n^{r_i})}$ for such an f , where r_i are the arities of the relation symbols in the signature. On the other hand, Corollary 4.4 implies that such an f is not $o(\log \log n)$. This leaves quite a large gap between the known lower and upper bounds. Shrinking this gap is an interesting avenue for further research.

Decidability of the existence of an adaptive left unbounded query algorithm over \mathbb{B} . For classes of the form $\text{CSP}(A)$, Corollary 5.6 shows that determining whether they are decided by an adaptive left unbounded query algorithm is decidable. The criterion given in Corollary 5.8 for Datalog definable classes is, however, not effective. This raises the question of whether there exists such an effective criterion for Datalog definable classes, or whether that problem is undecidable. Note that the “Rice-style” general undecidability result for semantic properties of Datalog programs from [11] appears to be of no help here: although *admitting an adaptive left unbounded query algorithm over \mathbb{B}* is a semantic property of Datalog programs, it fails to satisfy their “strong non-triviality” requirement.

Adaptive hybrid k -query algorithms over \mathbb{B} . A query algorithm that can use both left and right queries can be called a hybrid query algorithm. Adaptive hybrid query algorithms over \mathbb{N} have been studied by Wu [21], but there has been no research on such algorithms over \mathbb{B} . Every class of structures that is closed under homomorphic equivalence can be decided by an adaptive hybrid unbounded query algorithm over \mathbb{B} . On input B , the algorithm simply queries for $\text{hom}_{\mathbb{B}}(A, B)$ and $\text{hom}_{\mathbb{B}}(B, A)$ for each structure A . At some point, both of these numbers are 1, then the homomorphic equivalence class of B is decided, which suffices to classify it correctly. This evokes the question of what the expressive power of hybrid k -query algorithms is: Do we get a strict hierarchy as in Figure 1? What is the relation between the expressive power of non-adaptive and adaptive hybrid query algorithms over \mathbb{B} ?

References

- 1 Albert Atserias, Phokion G Kolaitis, and Wei-Lin Wu. On the expressive power of homomorphism counts. In *2021 36th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–13. IEEE, 2021. doi:10.1109/LICS52264.2021.9470543.
- 2 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. doi:10.1007/BF01305232.
- 3 Balder ten Cate, Víctor Dalmau, and Jakub Opršal. Right-Adjoint for Datalog Programs. In Graham Cormode and Michael Shekelyan, editors, *27th International Conference on Database Theory (ICDT 2024)*, volume 290 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 10:1–10:20, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ICDT.2024.10.
- 4 Surajit Chaudhuri. Finding nonrecursive envelopes for Datalog predicates. In *Proceedings of the twelfth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems*, pages 135–146, 1993. doi:10.1145/153850.153862.
- 5 Surajit Chaudhuri and Phokion G Kolaitis. Can Datalog be approximated? In *Proceedings of the thirteenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems*, pages 86–96, 1994. doi:10.1145/182591.182602.
- 6 Surajit Chaudhuri and Moshe Y Vardi. Optimization of real conjunctive queries. In *Proceedings of the twelfth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems*, pages 59–70, 1993. doi:10.1145/153850.153856.
- 7 Yijia Chen, Jörg Flum, Mingjun Liu, and Zhiyang Xun. On algorithms based on finitely many homomorphism counts. In Stefan Szeider, Robert Ganian, and Alexandra Silva, editors, *47th International Symposium on Mathematical Foundations of Computer Science (MFCS 2022)*, volume 241 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 32:1–32:15, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.MFCS.2022.32.
- 8 Zdeněk Dvořák. On recognizing graphs by numbers of homomorphisms. *Journal of Graph Theory*, 64(4):330–342, 2010. doi:10.1002/jgt.20461.
- 9 Heinz-Dieter Ebbinghaus and Jörg Flum. *Finite model theory*. Perspectives in Mathematical Logic. Springer, 1995.
- 10 Péter L. Erdős, Dömötör Pálvölgyi, Claude Tardif, and Gábor Tardos. Regular families of forests, antichains and duality pairs of relational structures. *Combinatorica*, 37(4):651–672, 2017. doi:10.1007/s00493-015-3003-4.
- 11 Haim Gaifman, Harry G. Mairson, Yehoshua Sagiv, and Moshe Y. Vardi. Undecidable optimization problems for database logic programs. *J. ACM*, 40(3):683–713, July 1993. doi:10.1145/174130.174142.
- 12 Martin Grohe. Counting bounded tree depth homomorphisms. In *Proceedings of the 35th Annual ACM/IEEE Symposium on Logic in Computer Science*, pages 507–520, 2020. doi:10.1145/3373718.3394739.
- 13 Pavol Hell and Jaroslav Nešetřil. *Graphs and homomorphisms*, volume 28. OUP Oxford, 2004.
- 14 Benoît Larose, Cynthia Loten, and Claude Tardif. A characterisation of first-order constraint satisfaction problems. *Log. Methods Comput. Sci.*, 3(4), 2007. doi:10.2168/LMCS-3(4:6)2007.
- 15 László Lovász. Operations with structures. *Acta Mathematica Hungarica*, 18(3-4):321–328, 1967.
- 16 Laura Mančínska and David E. Roberson. Quantum isomorphism is equivalent to equality of homomorphism counts from planar graphs. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 661–672, 2020. doi:10.1109/FOCS46700.2020.00067.
- 17 James R. Munkres. *Topology*. Prentice Hall, Inc., 2 edition, January 2000. URL: <http://www.worldcat.org/isbn/0131816292>.

- 18 Tim Seppelt. Logical equivalences, homomorphism indistinguishability, and forbidden minors. *Information and Computation*, 301:105224, 2024. doi:10.1016/j.ic.2024.105224.
- 19 Balder ten Cate, Victor Dalmau, Phokion G. Kolaitis, and Wei-Lin Wu. When do homomorphism counts help in query algorithms? In Graham Cormode and Michael Shekelyan, editors, *27th International Conference on Database Theory, ICDT 2024, March 25-28, 2024, Paestum, Italy*, volume 290 of *LIPIcs*, pages 8:1–8:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPIcs.ICDT.2024.8.
- 20 Balder ten Cate, Phokion G. Kolaitis, and Arnar Á. Kristjánsson. Adaptive query algorithms for relational structures based on homomorphism counts, 2025. doi:10.48550/arXiv.2504.16567.
- 21 Wei-Lin Wu. Query algorithms based on homomorphism counts. In *Structure Meets Power Workshop (Contributed Talks)*, page 24, 2022.
- 22 Wei-Lin Wu. *A Study of the Expressive Power of Homomorphism Counts*. PhD thesis, University of California, Santa Cruz, USA, 2023. URL: <https://www.escholarship.org/uc/item/4647715d>.