

The Pebbling Comonad in Finite Model Theory

Samson Abramsky

Department of Computer Science, University of Oxford, Oxford, U.K.

Email: samson.abramsky@cs.ox.ac.uk

Anuj Dawar and Pengming Wang

Computer Laboratory, University of Cambridge

Email: anuj.dawar@cl.cam.ac.uk, pengming.wang@cl.cam.ac.uk

Abstract—Pebble games are a powerful tool in the study of finite model theory, constraint satisfaction and database theory. Monads and comonads are basic notions of category theory which are widely used in semantics of computation and in modern functional programming. We show that existential k -pebble games have a natural comonadic formulation. Winning strategies for Duplicator in the k -pebble game for structures A and B are equivalent to morphisms from A to B in the coKleisli category for this comonad. This leads on to comonadic characterisations of a number of central concepts in Finite Model Theory:

- Isomorphism in the co-Kleisli category characterises elementary equivalence in the k -variable logic with counting quantifiers.
- Symmetric games corresponding to equivalence in full k -variable logic are also characterized.
- The treewidth of a structure A is characterised in terms of its coalgebra number: the least k for which there is a coalgebra structure on A for the k -pebbling comonad.
- Co-Kleisli morphisms are used to characterize strong consistency, and to give an account of a Cai-Fürer-Immerman construction.
- The k -pebbling comonad is also used to give semantics to a novel modal operator.

These results lay the basis for some new and promising connections between two areas within logic in computer science which have largely been disjoint: (1) finite and algorithmic model theory, and (2) semantics and categorical structures of computation.

1. Introduction

Homomorphisms play a fundamental rôle in finite model theory, constraint satisfaction and database theory. The existence of a homomorphism $A \rightarrow B$ is an equivalent formulation of the basic CSP problem [1], [2], [3]. There is an equivalence between the existence of a homomorphism, and the property that every existential positive sentence satisfied by A is also satisfied by B [2]. Such sentences correspond to (disjunctions of) *conjunctive queries*, which are fundamental in database theory [4], [5].

One of the key tools in studying these notions is that of existential k -pebble games [6]. Such a game, for structures A, B , proceeds by Spoiler placing one of his k pebbles on an element of the universe of A . Duplicator then places one of her pebbles on an element of B . If Duplicator is always able to move so that the partial mapping from A to B defined by sending a_i , the element in A carrying the i 'th Spoiler pebble, to b_i , the corresponding element of B carrying the i 'th Duplicator pebble, is a homomorphism on the induced substructures, then Duplicator has a winning strategy.

Proposition 1 ([6]). *The following are equivalent:*

- Duplicator has a winning strategy in the existential k -pebble game.
- Every sentence of the existential positive k -variable fragment of first-order logic satisfied by A is also satisfied by B .

Our aim in this paper is to study these notions from a novel perspective, using the notion of *comonad* from category theory. Monads and comonads are basic notions of category theory which are widely used in semantics of computation and in modern functional programming [7], [8], [9]. We show that existential k -pebble games have a natural comonadic formulation. Given a structure A over a relational signature σ , we shall introduce a new structure $\mathbb{T}_k A$ corresponding to Spoiler playing his part of an existential k -pebble game on A , with the potential codomain B left unspecified. The idea is that we can exactly recover the content of a Duplicator strategy in B by giving a homomorphism from $\mathbb{T}_k A$ to B . Thus the notion of *local approximation* built into the k -pebble game is internalised into the category of σ -structures and homomorphisms. Formally, this construction will be shown to give a *comonad* on this category, which guarantees a wealth of further structural properties. This leads to comonadic characterisations of a number of central concepts in Finite Model Theory.

In Section 2, we introduce the pebbling comonads \mathbb{T}_k , which are graded by the number of pebbles k , and characterize their coalgebras. $\mathbb{T}_k A$ is always infinite. In Section 3, we prove a no-go theorem, to rule out any finite version. In Section 4, we show that the question of whether a morphism $\mathbb{T}_k A \rightarrow B$ exists is equivalent to the existence

of a winning strategy for the existential k -pebble game for A and B . We also show how this question can be finitized, and such morphisms can be realized by deterministic finite-state transducers. In Section 5 we study various notions of equivalence of structures, and show that isomorphism in the coKleisli category for \mathbb{T}_k coincides with elementary equivalence for C^k , the k -variable logic with counting quantifiers, which plays a central rôle in finite model theory. We also show that the equivalence of structures A and B for full k -variable logic is characterized by the existence of a pair of morphisms $\mathbb{T}_k A \rightarrow B$ and $\mathbb{T}_k B \rightarrow A$ satisfying a certain condition. This is important, as it shows that symmetric, back-and-forth game conditions also fall within the scope of the comonadic approach. In Section 6, we characterise treewidth, a combinatorial parameter which plays a pervasive rôle in algorithmic graph theory, in terms of the coalgebra number of a structure A : the least index k such that there is a coalgebra $A \rightarrow \mathbb{T}_k A$. Section 7 characterizes the weaker condition of existence of a homomorphism from A to $\mathbb{T}_k A$ in terms of the treewidth of the core of A . Section 8 discusses strong k -consistency, and gives a Cai-Fürer-Immerman construction in terms of coKleisli morphisms. In Section 9, we introduce the novel modality \Box_k , which can be read as “it is k -locally the case that”, and give its semantics using the pebbling comonad. Section 10 concludes.

Overall, these results lay the basis for some new and promising connections between two areas within logic in computer science which have largely been disjoint: (1) finite and algorithmic model theory, and (2) semantics and categorical structures of computation.

Notation and Background

We have attempted to make this paper accessible to non-specialists in category theory. We assume only very basic background in category theory (see e.g. [10], [11]), essentially the definitions of category and functor, spelling out other definitions as needed.

We also assume only basic background in (finite) model theory (see e.g. [1]). We use the letters A, B for σ -structures over some relational signature σ , not distinguishing notationally between the structures and their universes. The interpretation of a relation $R \in \sigma$ in a structure A is denoted by R^A . Our morphisms between structures will always be homomorphisms. We use $\mathcal{L}_{\infty, \omega}^k$ to denote the k -variable fragment of first-order logic extended with infinite disjunctions and conjunctions, $\exists^+ \mathcal{L}_{\infty, \omega}^k$ for its existential positive fragment (no universal quantifiers or negations), and $\exists^+ \mathcal{L}_{\omega, \omega}^k$ for the corresponding fragment of first-order logic.

We use the notation $[n] := \{1, \dots, n\}$, and $s \sqsubseteq t$ for the prefix ordering on sequences. If $s \sqsubseteq t$, there is a unique s' such that $ss' = t$, which we refer to as the suffix of s in t .

2. The Pebbling Comonad

For the remainder of the paper, we fix a finite relational signature σ . Our main setting is $\mathcal{R}(\sigma)$, the category

whose objects are σ -structures, and whose morphisms are homomorphisms of σ -structures. We shall work in $\mathcal{R}(\sigma)$ and various sub-categories thereof, in particular $\mathcal{R}_f(\sigma)$, the category of finite σ -structures. By “structure”, we always mean σ -structure.

Firstly, we can consider the set of all plays in A by the Spoiler. This can be represented by $([k] \times A)^+$, the set of finite non-empty sequences of moves (p, a) , where $p \in [k]$ is a pebble index, and $a \in A$. We shall use the notation $s = [(p_1, a_1), \dots, (p_n, a_n)]$ for these sequences. The set of plays forms the universe of $\mathbb{T}_k A$.

To complete the definition of $\mathbb{T}_k A$, we must define the relational structure on this universe. For simplicity, we first consider the case of a binary relation E . We define $E^{\mathbb{T}_k A}$ to be the set of pairs of plays $s, t \in \mathbb{T}_k A$ such that

- s and t are comparable in the prefix ordering, so $s \sqsubseteq t$ or $t \sqsubseteq s$.
- If $s \sqsubseteq t$, then the pebble index of the last move in s does not appear in the suffix of s in t ; and symmetrically if $t \sqsubseteq s$.
- $E^A(\varepsilon_A(s), \varepsilon_A(t))$, where $\varepsilon_A : \mathbb{T}_k A \rightarrow A$ sends a play $[(p_1, a_1), \dots, (p_n, a_n)]$ to a_n , the A -component of its last move.

To understand the second condition in this definition, note that for each pebble index p , only the *last* move with pebble index p is relevant to the current position. The idea is that, in placing pebble p on element a , Spoiler has to first remove it from its previous position. This is exactly the way in which the pebble game models bounded resources.

The extension to relations of arbitrary arity is straightforward. Given an m -ary relation $R \in \sigma$, $R^{\mathbb{T}_k A}(s_1, \dots, s_m)$ if the s_i are pairwise comparable in the prefix ordering, and hence form a chain with greatest element s ; the pebble index of the last move in each s_i does not appear in the suffix of s_i in s ; and $R^A(\varepsilon_A(s_1), \dots, \varepsilon_A(s_m))$.

The following is immediate from our definition.

Proposition 2. *The map $\varepsilon_A : \mathbb{T}_k A \rightarrow A$ is a homomorphism.*

We now extend \mathbb{T}_k to a functor $\mathbb{T}_k : \mathcal{R}(\sigma) \rightarrow \mathcal{R}(\sigma)$ by defining its action on morphisms. If $f : A \rightarrow B$ is a homomorphism, we define $\mathbb{T}_k f : \mathbb{T}_k A \rightarrow \mathbb{T}_k B$ to be the map

$$[(p_1, a_1), \dots, (p_n, a_n)] \mapsto [(p_1, f(a_1)), \dots, (p_n, f(a_n))].$$

It is clear that this is a homomorphism from $\mathbb{T}_k A$ to $\mathbb{T}_k B$. Moreover, it is easily verified that $\mathbb{T}_k(g \circ f) = \mathbb{T}_k(g) \circ \mathbb{T}_k(f)$, and $\mathbb{T}_k(\text{id}_A) = \text{id}_{\mathbb{T}_k A}$, so \mathbb{T}_k is a functor.

We have already defined the homomorphism $\varepsilon_A : \mathbb{T}_k A \rightarrow A$ for each structure A . We now note that this defines a natural transformation. That is, for each homomor-

phism $f : A \longrightarrow B$, the following diagram commutes.

$$\begin{array}{ccc} \mathbb{T}_k A & \xrightarrow{\varepsilon_A} & A \\ \mathbb{T}_k f \downarrow & & \downarrow f \\ \mathbb{T}_k B & \xrightarrow{\varepsilon_B} & B \end{array}$$

Now for each structure A , we define a map $\delta_A : \mathbb{T}_k A \longrightarrow \mathbb{T}_k \mathbb{T}_k A$, as follows. Given a play $s = [(p_1, a_1), \dots, (p_n, a_n)]$, define $s_i = [(p_1, a_1), \dots, (p_i, a_i)]$, $i = 1, \dots, n$. Then we define

$$\delta_A : s \mapsto [(p_1, s_1), \dots, (p_n, s_n)].$$

Proposition 3. *For each A , δ_A is a homomorphism. Moreover, δ is a natural transformation: for each homomorphism $f : A \longrightarrow B$, the following diagram commutes.*

$$\begin{array}{ccc} \mathbb{T}_k A & \xrightarrow{\delta_A} & \mathbb{T}_k \mathbb{T}_k A \\ \mathbb{T}_k f \downarrow & & \downarrow \mathbb{T}_k \mathbb{T}_k f \\ \mathbb{T}_k B & \xrightarrow{\delta_B} & \mathbb{T}_k \mathbb{T}_k B \end{array}$$

Proof. This is a straightforward unfolding of the definitions. Starting with a play $[(p_1, a_1), \dots, (p_n, a_n)] \in \mathbb{T}_k A$, going either way around the diagram results in $[(p_1, t_1), \dots, (p_n, t_n)]$, where $t_i = [(p_1, f(a_1)), \dots, (p_i, f(a_i))]$, $i = 1, \dots, n$. \square

The natural transformation ε is the *counit* of the comonad, while δ is the *comultiplication*.

We can now gather these elements together to complete our construction.

Theorem 4. *The triple $(\mathbb{T}_k, \varepsilon, \delta)$ forms a comonad on the category $\mathcal{R}(\sigma)$.*

Proof. The remaining points to be verified from the definition of a comonad are that the following diagrams commute, for all objects A of $\mathcal{R}(\sigma)$.

$$\begin{array}{ccc} \mathbb{T}_k A & \xrightarrow{\delta_A} & \mathbb{T}_k \mathbb{T}_k A \\ \delta_A \downarrow & & \downarrow \mathbb{T}_k \delta_A \\ \mathbb{T}_k \mathbb{T}_k A & \xrightarrow{\delta_{\mathbb{T}_k A}} & \mathbb{T}_k \mathbb{T}_k \mathbb{T}_k A \end{array} \quad \begin{array}{ccc} \mathbb{T}_k A & \xrightarrow{\delta_A} & \mathbb{T}_k \mathbb{T}_k A \\ \delta_A \downarrow & \searrow & \downarrow \mathbb{T}_k \varepsilon_A \\ \mathbb{T}_k \mathbb{T}_k A & \xrightarrow{\varepsilon_{\mathbb{T}_k A}} & \mathbb{T}_k A \end{array}$$

While these diagrams look somewhat formidable, verification that they commute again reduces to a straightforward diagram chase. For the first diagram, starting with a play $[(p_1, a_1), \dots, (p_n, a_n)] \in \mathbb{T}_k A$, either way around the diagram takes us to $[(p_1, t_1), \dots, (p_n, t_n)]$, where $s_i = [(p_1, a_1), \dots, (p_i, a_i)]$, $t_i = [(p_1, s_1), \dots, (p_i, s_i)]$, $i = 1, \dots, n$. \square

2.1. The Co-Kleisli Category

We now turn to one of the fundamental constructions associated with a comonad, the co-Kleisli category [12]. We use the notation $\mathcal{K}(\mathbb{T}_k)$ for this category. The objects are the same as those of $\mathcal{R}(\sigma)$, while a morphism from A to B in $\mathcal{K}(\mathbb{T}_k)$ is a $\mathcal{R}(\sigma)$ -morphism $f : \mathbb{T}_k A \longrightarrow B$. Given morphisms $f : \mathbb{T}_k A \longrightarrow B$ and $g : \mathbb{T}_k B \longrightarrow C$, we use the comonad structure to compose them:

$$\mathbb{T}_k A \xrightarrow{\delta_A} \mathbb{T}_k \mathbb{T}_k A \xrightarrow{\mathbb{T}_k f} \mathbb{T}_k B \xrightarrow{g} C$$

The identity morphisms are given by the counit of the comonad:

$$\varepsilon_A : \mathbb{T}_k A \longrightarrow A.$$

We write $A \rightarrow_k B$ if there exists a morphism from A to B in $\mathcal{K}(\mathbb{T}_k)$.

The Kleisli coextension. The operation $f \mapsto \mathbb{T}_k f \circ \delta_A$ which sends $f : \mathbb{T}_k A \longrightarrow B$ to $f^* : \mathbb{T}_k A \longrightarrow \mathbb{T}_k B$ is known as the *Kleisli coextension*. Comonads have an alternative presentation in terms of this operation and the counit maps [7]. For our purposes, it will be useful to have a concrete description of this operation. Given a co-Kleisli morphism $f : \mathbb{T}_k A \longrightarrow B$,

$$f^* : [(p_1, a_1), \dots, (p_n, a_n)] \mapsto [(p_1, f(s_1)), \dots, (p_n, f(s_n))]$$

where $s_i = [(p_1, a_1), \dots, (p_i, a_i)]$, $i = 1, \dots, n$.

2.2. Grading

We have defined a comonad \mathbb{T}_k for each positive integer k . We now consider how these are related.

We can think of the morphisms $f : A \longrightarrow B$ in the co-Kleisli category for \mathbb{T}_k as those which only have to respect the k -local structure of A . The lower the value of k , the less information available to Spoiler, and the easier it is for Duplicator to have a winning strategy. Equivalently by Theorem 13, the easier it is to have a homomorphism $\mathbb{T}_k A \longrightarrow B$, i.e. a morphism from A to B in the co-Kleisli category. This leads to a natural weakening principle: if we have a morphism from $\mathbb{T}_l A$ to B , then this should yield a morphism from $\mathbb{T}_k A$ to B when $l < k$.

This idea is directly captured in our construction. There is an inclusion $i_A^{l,k} : \mathbb{T}_l A \hookrightarrow \mathbb{T}_k A$ whenever $l \leq k$. The following is easily verified.

Proposition 5. *The inclusion maps form a natural transformation $i^{l,k} : \mathbb{T}_l \longrightarrow \mathbb{T}_k$ which is a morphism of comonads, i.e. it preserves the counit and comultiplication.*

This supports the weakening principle mentioned above. Given a morphism $f : \mathbb{T}_k A \longrightarrow B$, we have a morphism

$$\mathbb{T}_l A \xrightarrow{i_A^{l,k}} \mathbb{T}_k A \xrightarrow{f} B$$

2.3. Coalgebras

Another fundamental aspect of comonads is that they have an associated notion of *coalgebra*. A coalgebra for \mathbb{T}_k is a morphism $\alpha : A \longrightarrow \mathbb{T}_k A$ such that the following diagrams commute:

$$\begin{array}{ccc} A & \xrightarrow{\alpha} & \mathbb{T}_k A \\ \alpha \downarrow & & \downarrow \delta_A \\ \mathbb{T}_k A & \xrightarrow{\mathbb{T}_k \alpha} & \mathbb{T}_k \mathbb{T}_k A \end{array} \quad \begin{array}{ccc} A & \xrightarrow{\alpha} & \mathbb{T}_k A \\ \text{id}_A \searrow & & \downarrow \varepsilon_A \\ & & A \end{array}$$

Note in particular that a coalgebra structure on A makes it a retract of $\mathbb{T}_k A$ via the counit ε_A .

A morphism of coalgebras from (A, α) to (B, β) is a morphism $h : A \longrightarrow B$ such that the following diagram commutes:

$$\begin{array}{ccc} A & \xrightarrow{\alpha} & \mathbb{T}_k A \\ h \downarrow & & \downarrow \mathbb{T}_k h \\ B & \xrightarrow{\beta} & \mathbb{T}_k B \end{array}$$

Coalgebras and their morphisms form a category $\mathcal{R}(\sigma)^{\mathbb{T}_k}$, the *Eilenberg-Moore category* for the comonad \mathbb{T}_k . This provides another way of looking at the co-Kleisli category. We can think of the objects in it as the *cofree coalgebras* $\mathbb{T}_k A$, with structure maps given by comultiplication. Note that the diagrams for a comonad instantiate those for a coalgebra when we take the coalgebra to be $(\mathbb{T}_k A, \delta_A)$. The morphisms are then taken to be the coalgebra morphisms $h : \mathbb{T}_k A \longrightarrow \mathbb{T}_k B$. This is equivalent to the usual presentation, since we can pass from $f : \mathbb{T}_k A \longrightarrow B$ to its Kleisli coextension $f^* = Tf \circ \delta_A : \mathbb{T}_k A \longrightarrow \mathbb{T}_k B$, and from a coalgebra morphism $h : \mathbb{T}_k A \longrightarrow \mathbb{T}_k B$ to $\varepsilon_A \circ h : \mathbb{T}_k A \longrightarrow B$, and these two passages are mutually inverse. This representation displays the co-Kleisli category explicitly as a full subcategory of the Eilenberg-Moore category.

Note that a coalgebra structure α on A implies that a homomorphism exists from A to B whenever a homomorphism exists from $\mathbb{T}_k A$ to B . Given $h : \mathbb{T}_k A \longrightarrow B$, we can form $h \circ \alpha : A \longrightarrow B$. Thus we should only expect a coalgebra structure to exist when the k -local information on A is sufficient to determine the structure of A .

Given a structure A , we define a *k-traversal* of A to be a structure (A, \leq, i) , where \leq is a partial order on A which is a tree order, i.e. for each $a \in A$, the predecessors of a in the order form a chain; and $i : A \rightarrow [k]$ is a labelling map such that, whenever a is adjacent to b in the Gaifman graph of A , a is comparable to b , say $a \leq b$, and for all c such that $a < c \leq b$, $i(a) \neq i(c)$.

Theorem 6. *Let A be a finite structure. There is a bijective correspondence between*

- 1) *coalgebras* $\alpha : A \longrightarrow \mathbb{T}_k A$
- 2) *k-traversals* (A, \leq, i) .

Proof. Let $\alpha : A \longrightarrow \mathbb{T}_k A$ be a coalgebra. Let $T \subseteq \mathbb{T}_k A$ be the image of α . Unpacking the content of the two commutative diagrams for a coalgebra we see that if $\alpha(a) = [(p_1, a_1), \dots, (p_l, a_l)]$, then $a_l = a$, and $\alpha(a_i) = [(p_1, a_1), \dots, (p_i, a_i)]$, $i = 1, \dots, l$. Thus T is a prefix-closed subset of $\mathbb{T}_k A$, and moreover, for each $a \in A$, there is a unique $s \in T$ with last move of the form (p, a) . We can then define $a \leq b \iff \alpha(a) \sqsubseteq \alpha(b)$, and $i(a) = p$ where $s[p, a] \in T$. Since α is a homomorphism, it must be the case that when a and b are adjacent in the Gaifman graph, $\alpha(a)$ is comparable with $\alpha(b)$, say $\alpha(a) \sqsubseteq \alpha(b)$, and the pebble index $i(a)$ must not occur in the suffix of $\alpha(a)$ in $\alpha(b)$. Thus (A, \leq, i) is a k -traversal of A .

Conversely, let (A, \leq, i) be a k -traversal of A . We define $\alpha(a)$ by induction on the number of strict predecessors of a : $\alpha(a) = s[(i(a), a)]$, where $s = \alpha(a^-)$ if a^- is the immediate predecessor of a , and otherwise $s = []$. The k -traversal conditions imply that $\alpha : A \longrightarrow \mathbb{T}_k A$ is a coalgebra. It is easy to see that these passages between coalgebras and k -traversals are mutually inverse. \square

In the case that $k \geq n$, where n is the cardinality of A , there is a trivial k -traversal of A obtained by choosing a linear order $a_1 < \dots < a_n$ on A and defining $i(a_j) = j$.

We shall return to the issue of when coalgebras exist in our discussion of treewidth in Section 6.

3. Finite and infinite

Our primary focus is on finite structures. However, $\mathbb{T}_k A$ is always infinite. Is this necessary? As we shall see in the next section, given finite structures A and B , the question of whether there is a homomorphism $\mathbb{T}_k A \rightarrow B$ can be finitized, using a positional representation. However, to give a comonadic representation of k -locality, we have to define a structure on a given A which will allow us to characterize the situation for *all* choices of target B . Could this be done using a finite representation instead of $\mathbb{T}_k A$ with the same effect? The following no-go result says that this is not the case, and therefore an infinite representation cannot be avoided. We prove it for the special case $k = 2$.

Theorem 7. *There is no construction $A \mapsto QA$ on finite structures A such that QA is finite, and for all finite B :*

$$QA \rightarrow B \iff A \rightarrow_2 B.$$

Proof. We assume that σ contains a binary relation symbol E and let A be the structure with three elements a, b, c where E is interpreted as the three-element cycle, i.e. $E(a, b), E(b, c), E(c, a)$ hold and no other pairs are related by E . All other relations in σ are interpreted by the empty relation on A .

We claim that for any finite B , $A \rightarrow_2 B$ if and only if B contains an E -cycle. In one direction, we note that $\mathbb{T}_2 A$ contains an infinite E -path: $[(1, a)], [(1, a), (2, b)], [(1, a), (2, b), (1, c)], [(1, a), (2, b), (1, c), (2, a)], \dots$. The homomorphic image of an infinite path inside a finite structure must contain a cycle. In the other direction, assume that B contains an E -cycle C . We define a map $h : \mathbb{T}_2 A \rightarrow B$ by induction on the length of plays in \mathbb{T}_2 , such that the image of h is contained in the cycle C . For $s = [(p, x)]$ where $p \in \{1, 2\}$ and $x \in \{a, b, c\}$, choose $h(s)$ to be an arbitrary element of the cycle in B . Suppose now that s has length at least 2, the last move in s is (p, x) and by induction $h(s')$ has been defined for all proper prefixes s' of s . Let t be the longest prefix of s ending in a move (q, y) for $q \neq p$. We define $h(s)$ according to the three cases: if $x = y$ then $h(s) = h(t)$; if $E(x, y)$ then let $h(x)$ be the element of C with an E -edge to $h(t)$; and if $E(y, x)$, then let $h(x)$ be the element of C with an E -edge from $h(t)$. It is then easily checked that h is a homomorphism.

Now, suppose there was a finite QA as in the statement. Then, since $A \rightarrow_2 QA$, QA contains an E -cycle. Let m be the length of the shortest cycle in QA . Consider the structure C_{m+1} consisting of a single directed cycle of length $m+1$. Since it contains a cycle, $A \rightarrow_2 C_{m+1}$. However, it is not the case that $QA \rightarrow C_{m+1}$, since the homomorphic image of a cycle of length m must be a closed walk of length m/l for some l , and C_{m+1} contains no such walk. \square

We can apply this result to rule out a categorical formulation of our question.

Corollary 8. *There is no comonad defined on $\mathcal{R}_f(\sigma)$ whose co-Kleisli category has the same preorder collapse as $\mathcal{K}(\mathbb{T}_2)$.*

This version of the no-go theorem is stated in purely categorical terms, while the proof uses very characteristic finite model theory arguments — indeed, essentially this argument appeared in the proof of Proposition 7.9 in [13].

We conjecture that this result extends to all higher values of k . Indeed, a generalization of the construction in the proof above can be used to show that for each k there is a signature σ , containing a k -ary relation symbol, and a σ -structure A , for which there is no finite QA with $QA \rightarrow B$ iff $A \rightarrow_k B$. It remains to prove the result uniformly for all signatures.

4. Positions and strategies

We now turn to the situation where finite structures A and B are given. We define $\Gamma_k(A, B)$, the set of *configurations* or *positions* for the existential k -pebble game from A to B , to be the set of all partial functions $\gamma : [k] \rightarrow A \times B$. The idea is that γ represents the position left on the board after some rounds of the game. The domain of γ is the set of pebbles used so far; $\gamma(p) = (a, b)$ means that Spoiler currently has pebble p placed on $a \in A$, while Duplicator has her matching pebble on $b \in B$. We include the empty partial function, which represents the initial configuration.

We shall use the update operation $\gamma[p \mapsto (a, b)]$ on configurations. This yields the configuration with domain $\text{dom}(\gamma) \cup \{p\}$, and such that

$$\gamma[p \mapsto (a, b)](q) = \begin{cases} (a, b), & p = q \\ \gamma(q) & p \neq q \end{cases}$$

There are really two cases which are covered by this update operation, both of which have natural readings in terms of the pebble game. If $p \in \text{dom}(\gamma)$, then the update represents Spoiler moving pebble p from its current position to mark a , while Duplicator moves her matching pebble to b . If $p \notin \text{dom}(\gamma)$, then we are extending the domain of γ , which corresponds to Spoiler placing a previously unused pebble, and Duplicator her matching pebble.

It will be convenient to use the transition notation $\gamma \xrightarrow{(p,a):b} \gamma'$, where $\gamma' = \gamma[p \mapsto (a, b)]$; and $\gamma \longrightarrow \gamma'$ if $\gamma \xrightarrow{(p,a):b} \gamma'$ for some p, a, b .

We now relate positions to plays. A strategy for Duplicator in the existential k -pebble game from A to B can be represented by a function $f : \mathbb{T}_k A \rightarrow B$, which responds to each move of Spoiler, in the context of the previous history of the game, with a move in B . The coextension f^* , defined as in Section 2.1, makes explicit the sequence of responses by Duplicator, with the matching use of pebbles to those of Spoiler. Now consider a pair (s, t) in the graph of f^* , i.e. $s \in \mathbb{T}_k A$ and $t = f^*(s)$. Note that s and t will be the same length, and have the same sequence of pebble indices. Let p_1, \dots, p_l be the pebble indices occurring in s , with $1 \leq l \leq k$. Let (p_i, a_i) be the last occurrence of a move for pebble p_i in s , (p_i, b_i) the corresponding occurrence in t . We define position (s, t) to be the configuration γ with domain $\{p_1, \dots, p_l\}$ and $\gamma(p_i) = (a_i, b_i)$, $i = 1, \dots, l$. We write $\theta_f : \mathbb{T}_k A \rightarrow \Gamma_k(A, B)$ for the map

$$\theta_f(s) = \text{position}(s, f^*(s)).$$

Note that this is a map from the infinite set $\mathbb{T}_k A$ to the finite set $\Gamma_k(A, B)$.

A *strategy in positional form* is given by a set $S \subseteq \Gamma_k(A, B)$ satisfying the following conditions:

- (S0) $\emptyset \in S$
- (S1) For all $\gamma \in S$, $p \in [k]$ and $a \in A$, there is $b \in B$ such that $\gamma \xrightarrow{(p,a):b} \gamma' \in S$.
- (S1) S is reachable: for all $\gamma \in S$, there is a sequence

$$\gamma_0 \longrightarrow \dots \longrightarrow \gamma_n$$

with $\gamma_0 = \emptyset$, $\gamma_n = \gamma$, and $\gamma_i \in S$ for all i .

Given a Duplicator strategy $f : \mathbb{T}_k A \rightarrow B$, we define its representation in positional form as

$$S_f := \{\theta_f(s) \mid s \in \mathbb{T}_k A\} \cup \{\emptyset\},$$

the set of positions which can be reached following the strategy represented by f .

Proposition 9. *For any strategy, the set of positions S_f is a strategy in positional form. Conversely, for any strategy*

in positional form S , there is a function $f : \mathbb{T}_k A \longrightarrow B$ such that $S_f = S$.

Proof. Given a strategy f , (S0) holds by definition of S_f . For closure under (S1), suppose that

$$\gamma = \theta_f(s).$$

If Spoiler places pebble p on a , this leads to an extended play $s[(p, a)]$; for any such move, Duplicator has a response $b = f(s[(p, a)])$. Then

$$\gamma \xrightarrow{(p,a):b} \gamma' = \theta_f(s[(p, a)]) \in S_f.$$

Reachability of $\gamma = \theta_f(s)$ holds by induction on the length of s , since $\theta_f(s) \longrightarrow \theta_f(s[(p, a)])$.

For the converse, given S we can define $f(s)$ by induction on $|s|$. We choose a linear order on B . Consider a play $s[(p, a)]$. By induction, $f(s)$ has already been defined, with $\gamma = \theta_f(s) \in S$. We define $X := \{b \in B \mid \gamma \xrightarrow{(p,a):b} \gamma' \in S\}$,

$$Y := \{b \in X \mid \exists t = t'[(p, a)] \sqsubseteq s. \theta_f(t') = \gamma \wedge f(t) = b\}.$$

Then we define

$$f(s[(p, a)]) = \begin{cases} \min_B(X \setminus Y), & X \setminus Y \neq \emptyset \\ \min_B(X), & \text{otherwise.} \end{cases}$$

Clearly $\theta_f(s[(p, a)]) \in S$. Note that by (S0), this definition also covers the case when s is empty.

It remains to show that for all $\gamma \in S$, $\gamma \in S_f$. We argue by induction on the length of the shortest transition sequence from the empty configuration to γ . Consider $\gamma \xrightarrow{(p,a):b} \gamma'$ with $\gamma = \theta_f(s)$, where s is taken minimal in the prefix order in $\theta_f^{-1}(\gamma)$. Let X be the set specified above in the definition of f , and let b have i strict predecessors in X in the chosen linear order on B . Now consider the play $s' = s[\underbrace{(p, a), \dots, (p, a)}_{i+1}]$. By definition of f , $f(s') = b$, so

$$\theta_f(s') = \gamma'. \quad \square$$

The reason for the somewhat involved construction in the second part of the proof is that we have to construct a deterministic strategy at the level of plays which maps onto a possibly non-deterministic strategy at the level of positions. We shall return to this point in the sequel.

4.1. Winning conditions

A configuration γ is *winning for Duplicator* if the relation

$$R(\gamma) := \{\gamma(p) \mid p \in \text{dom}(\gamma)\} \subseteq A \times B$$

is a partial homomorphism from A to B .

A strategy in positional form S is winning if γ is winning for all $\gamma \in S$. A Duplicator strategy $f : \mathbb{T}_k A \longrightarrow B$ is winning if S_f is winning.

It will be convenient to make use of the following device. We consider the expansion of our relational signature σ with an additional binary relation I . We turn σ -structures into

$\sigma \cup \{I\}$ -structures by interpreting I as the identity relation. We refer to such structures as I -structures. Note that if we interpret $\mathbb{T}_k A$ over this expanded signature, it will not be an I -structure; however, the interpretation of the I relation on $\mathbb{T}_k A$ where A is an I -structure will hold when multiple pebbles have been placed by Spoiler on the same element of A . This will ensure that a homomorphism $\mathbb{T}_k A \rightarrow B$ will induce a single-valued mapping on the underlying positions.

Proposition 10. *Given I -structures A and B , and a function $f : \mathbb{T}_k A \longrightarrow B$, the following are equivalent:*

- 1) f is a winning strategy for Duplicator.
- 2) f is a homomorphism.

Proof. If f is a homomorphism, the relation $R(\gamma)$ arising from any position in S_f is single-valued, since A and B are I -structures, so multiple pebbles paced on the same element of A must be mapped to the same element of B . From the definition of the relational structure on $\mathbb{T}_k A$, preservation of the relation instances in a play s is easily seen to be equivalent to the partial homomorphism condition on $\gamma = \theta_f(s)$. \square

4.2. Determinization

We can view a strategy in positional form S as a *finite-state transducer*. The set of states is S , with initial state \emptyset . The input set is $[k] \times A$, while the output set is B . For each configuration $\gamma \in S$, there are transitions

$$\gamma \xrightarrow{(p,a):b} \gamma'.$$

Note that this transducer need not be deterministic. In general, functions $f : \mathbb{T}_k A \longrightarrow B$ will give rise to non-deterministic transducers S_f , since different plays mapping to the same position can give rise to different transitions for given Spoiler moves (p, a) . However, there is a simple determinization procedure. Given S , we choose a function $\theta_\gamma : [k] \times A \longrightarrow B$ for each $\gamma \in S$ such that for all $(p, a) \in [k] \times A$, $\gamma \xrightarrow{(p,a):\theta_\gamma(p,a)} \gamma' \in S$, and then define the reachable set of configurations S_d under this choice of transitions. Explicitly, S_d is defined as the least fixpoint of the following monotone function on sets of transitions:

$$\Phi(U) = \begin{cases} \{\emptyset\} \cup \\ \{\gamma' \mid \exists \gamma \in U, (p, a) \in [k] \times A. \gamma \xrightarrow{(p,a):\theta_\gamma(p,a)} \gamma'\} \end{cases}$$

The following is easily verified.

- Proposition 11.**
- 1) S_d is a strategy in positional form.
 - 2) It is deterministic: for each $\gamma \in S_d$ and $(p, a) \in [k] \times A$, there are unique $b \in B$, $\gamma' \in S_d$ such that $\gamma \xrightarrow{(p,a):b} \gamma'$.
 - 3) If S is winning, so is $S_d \subseteq S$.

For finite structures A, B , we write $A \rightarrow_k^d B$ if there is a co-Kleisli morphism realized by a deterministic finite-state strategy in positional form.

Proposition 12. For finite structures A, B , $A \rightarrow_k B \iff A \rightarrow_k^d B$.

Proof. Given $f : \mathbb{T}_k A \longrightarrow B$, by Proposition 9 we can pass to the positional representation S_f , then by Proposition 11 determinize S_f to obtain S_d , then, by Proposition 9 again, pass back to a function $f_d : \mathbb{T}_k A \longrightarrow B$ such that $S_{f_d} = S_d$. Moreover, by Proposition 10, if f is a homomorphism, so is f_d . \square

4.3. Co-Kleisli morphisms and winning strategies

The following result, which is a corollary to Propositions 9, 10, and 12, justifies our claim that the pebbling comonad captures the content of the existential k -pebble game.

Theorem 13. Given I -structures A and B , the following are equivalent:

- 1) There is a winning strategy for Duplicator in the existential k -pebble game from A to B .
- 2) $A \rightarrow_k B$.
- 3) $A \rightarrow_k^d B$.

5. Equivalences

Various notions of equivalence between relational structures play an important rôle in Finite Model Theory, in particular the elementary equivalences induced by various logics. If \mathcal{L} is a logic, the corresponding equivalence is denoted $\equiv^{\mathcal{L}}$, where

$$A \equiv^{\mathcal{L}} B \iff \forall \varphi \in \mathcal{L}. A \models \varphi \iff B \models \varphi.$$

These equivalences can be characterized by various combinatorial games.

Our aim in this section is to characterize three important such equivalences in terms of morphisms in the co-Kleisli category. We introduced the relation $A \rightarrow_k B$ on structures in Section 2.1. This relation is clearly reflexive and transitive. The corresponding equivalence relation is

$$A \rightleftharpoons_k B \iff A \rightarrow_k B \wedge B \rightarrow_k A.$$

We write \equiv^{\exists^+} for the elementary equivalence induced by the existential-positive fragment of $\mathcal{L}_{\infty, \omega}^k$.

Proposition 14. For all A, B : $A \rightleftharpoons_k B \iff A \equiv^{\exists^+} B$.

Proof. This is an immediate consequence of Proposition 1 and Theorem 13. \square

5.1. Counting logic equivalence

A more interesting question is posed by isomorphism in the co-Kleisli category, which we denote by \cong_K . What does the equivalence this induces correspond to in logical terms? Given a morphism $f : \mathbb{T}_k A \longrightarrow B$, for each $s \in \mathbb{T}_k A \cup \{\emptyset\}$ and $p \in [k]$, there is a function $\psi_{s,p} : A \longrightarrow B$ such that, for all $a \in A$, $f(s(p, a)) = \psi_{s,p}(a)$. If these functions

are all injective, respectively surjective, we write $A \rightarrow_k^i B$, respectively $A \rightarrow_k^s B$. The corresponding equivalences are denoted $A \rightleftharpoons_k^i B$, $A \rightleftharpoons_k^s B$. If there is a morphism $f : \mathbb{T}_k A \longrightarrow B$ such that the functions $\psi_{s,p}$ are all bijective, and moreover for all $s \in \mathbb{T}_k A$ with $\gamma = \text{position}(s, f^*(s))$, the relation $R(\gamma)$ is a partial isomorphism, we write $A \rightarrow_k^b B$.

It is standard that for finite structures A and B , if there are injective homomorphisms $A \rightarrow B$ and $B \rightarrow A$, then $A \cong B$; and similarly for surjective homomorphisms. Localizing these arguments to the functions in context $\psi_{s,p}$ yields the following result.

Theorem 15. For all finite A, B :

$$A \rightleftharpoons_k^i B \iff A \rightleftharpoons_k^s B \iff A \rightarrow_k^b B \iff A \cong_K B.$$

Proof. Since A and B are finite, if there exist injective maps from A to B and B to A , the two sets have the same number of elements. It follows that any injective map is in fact surjective, and the first equivalence is immediate. Now, fix morphisms $f : \mathbb{T}_k A \longrightarrow B$ and $g : \mathbb{T}_k B \longrightarrow A$ witnessing $A \rightarrow_k^i B$ and $B \rightarrow_k^i A$ respectively. Then f^* and g^* are both injective maps. Consider the finite substructures P_A and P_B of $\mathbb{T}_k A$ and $\mathbb{T}_k B$ respectively induced by sequences of a length at most n . It is easily seen that f^* and g^* restricted to these substructures are injective homomorphisms $P_A \rightarrow P_B$ and $P_B \rightarrow P_A$ respectively. It follows that they are isomorphisms of these finite structures, and so $R(\gamma)$ is a partial isomorphism when $\gamma = \text{position}(s, f^*(s))$ for any s . Thus, f witnesses $A \rightarrow_k^b B$.

For the final equivalence, suppose we are given a homomorphism $f : \mathbb{T}_k A \longrightarrow B$ witnessing $A \rightarrow_k^b B$. Then, a simple induction on the length of the sequences establishes that f^* is a bijection between $\mathbb{T}_k A$ and $\mathbb{T}_k B$ and indeed is an isomorphism. Therefore $A \cong_K B$. Conversely, if $h : \mathbb{T}_k A \longrightarrow \mathbb{T}_k B$ is an isomorphism, we let $f : \mathbb{T}_k A \longrightarrow B$ be given by $f = \varepsilon_B \circ h$ and note that this satisfies the condition that all functions $\psi_{s,p}$ are bijective. Moreover, if $t = h^*(s)$, then $s = k^*(t)$ where k is the inverse of h in $\mathcal{K}(\mathbb{T}_k)$. The fact that h and k are both morphisms implies that $R(\gamma)$ is a partial isomorphism, where $\gamma = \text{position}(s, t)$. \square

We now recall the *bijection game* from [14]. This is a variant of the pebble game, in which at each round with corresponding position γ , Spoiler selects a pebble p , and Duplicator then selects a bijection h between A and B such that, for all $q \neq p$ with $\gamma(q) = (a, b)$, $h(a) = b$. Duplicator wins the round if, for each $a \in A$, $R(\gamma[p \mapsto a])$ is a partial isomorphism.

The bijection game characterizes the equivalence induced by the counting logic C^k [15], the k -variable logic with counting quantifiers, which plays a central rôle in finite model theory.

Theorem 16 ([14]). For all finite A, B , the following are equivalent:

- 1) $A \equiv^{C^k} B$

- 2) *Duplicator has a winning strategy in the bijection game.*

The following is essentially immediate from the definitions:

Proposition 17. *For all A, B , Duplicator has a winning strategy in the bijection game if and only if $A \rightarrow_k^b B$.*

Combining these results with Theorem 15, we obtain the main result of this subsection:

Theorem 18. *For all finite A, B :*

$$A \cong_K B \iff A \equiv^{C^k} B.$$

Thus isomorphism in the co-Kleisli category for the pebbling comonad \mathbb{T}_k characterizes the elementary equivalence induced by k -variable counting logic.

5.2. Back-and-forth equivalence

Finally, we turn to the elementary equivalence induced by the full k -variable logic $\mathcal{L}_{\infty, \omega}^k$, which we write as \equiv^k . The standard game characterization of this uses back-and-forth pebble games, in which Spoiler can play in B as well as A , and Duplicator has to respond in the other structure [16]. This can be defined concisely using our present notation as follows. Given $\gamma \in \Gamma_k(A, B)$, define $\gamma^- \in \Gamma_k(B, A)$ with $\gamma^-(p) = (b, a)$ iff $\gamma(p) = (a, b)$; and given $S \subseteq \Gamma_k(A)$, define $S^- = \{\gamma^- \mid \gamma \in S\}$. Now a positional-form winning strategy for Duplicator in the k -pebble game between A and B is $S \subseteq \Gamma_k(A, B)$ such that S is a winning strategy in the existential k -pebble game from A to B , and S^- is a winning strategy in the existential k -pebble game from B to A . Spelling this out, we see that this requires that S satisfies a back-condition as well as the usual forth-condition as in the existential case; and that $R(\gamma)$ is a partial isomorphism for all $\gamma \in S$.

Theorem 19 ([16]). *For all finite structures A and B , there is a winning strategy for Duplicator in the k -pebble game between A and B iff $A \equiv^k B$.*

We now define the relation $A \rightleftarrows_k^- B$ if there are co-Kleisli arrows $f : \mathbb{T}_k A \longrightarrow B$ and $g : \mathbb{T}_k B \longrightarrow A$ such that $S_f^- = S_g$.

Theorem 20. *For all finite structures A and B :*

$$A \rightleftarrows_k^- B \iff A \equiv^k B.$$

Proof. The forward implication follows directly from Theorem 19. For the converse, given winning existential strategies S from A to B and S^- from B to A , we obtain corresponding f and g from Theorem 9. \square

An interesting point arising from this result is the necessity for non-deterministic positional strategies. While existential strategies witnessing homomorphisms can always be determinized by Proposition 11, in general a coupled pair of strategies (S, S^-) cannot both be deterministic. A simple example where this arises is given by the complete graphs K_k and K_{k+1} . Note that $K_k \equiv^k K_{k+1}$, while $K_k \not\equiv^{C^k} K_{k+1}$.

6. Coalgebra number and treewidth

The notion of treewidth of a graph [17], extended to general relational structures in [3], plays a pervasive rôle in identifying “islands of tractability” in algorithmic graph theory. We shall write $\text{tw}(A)$ for the treewidth of a structure A .

We now consider how the comonadic structure of k -pebbling can be used to characterize treewidth.

Theorem 21. *For all structures A , $\text{tw}(\mathbb{T}_k A) < k$.*

Proof. We shall define a tree decomposition (T, a) of $\mathbb{T}_k A$. The set of nodes of T is $\mathbb{T}_k A \cup \{[]\}$, the set of plays together with the empty sequence. We say that s is adjacent to t if $s \prec t$ or $t \prec s$, where \prec is the immediate predecessor relation induced by the prefix order; thus $s \prec t$ iff $t = s[(p, a)]$ for some p, a . The unique path between any s and t goes from s via \succ -instances of the adjacency relation to u , and then by \prec -instances from u to t , where u is the greatest common prefix of s and t . The labelling function $a(s)$ assigns the set of *active prefixes* of s to s , where t is an active prefix of s if $t \sqsubseteq s$, and the pebble index used in the last move in t does not occur in the suffix of t in s . Clearly, the maximum size of any $a(s)$ is k .

We now verify the conditions for (T, a) to be a tree decomposition of $\mathbb{T}_k A$. Firstly, $s \in a(s)$. Secondly, if $R^{\mathbb{T}_k A}(s_1, \dots, s_n)$, then the s_i must all be active prefixes of some u . Finally, suppose that s is in $a(t_1) \cap a(t_2)$. Let u be the largest common prefix of t_1 and t_2 . Then we must have $s \sqsubseteq u$, and moreover s active in t_1 implies that s is active in all v with $u \sqsubseteq v \sqsubseteq t_1$, and similarly for t_2 . Thus $s \in a(v)$ for all v in the unique path from t_1 to t_2 in T . \square

Thus, although $\mathbb{T}_k A$ is infinite, it has bounded treewidth. As an immediate consequence of this result, we have:

Proposition 22. *If there is an injective homomorphism $A \rightarrow \mathbb{T}_k B$, then $\text{tw}(A) < k$.*

Proof. If there is an injective homomorphism from A to $\mathbb{T}_k B$, then A is isomorphic to a substructure of $\mathbb{T}_k B$, hence $\text{tw}(A) \leq \text{tw}(\mathbb{T}_k B) < k$. \square

In particular, if there is a coalgebra $\alpha : A \longrightarrow \mathbb{T}_k A$, then $\text{tw}(A) < k$. We define the *coalgebra number* $\kappa(A)$ of a finite structure A to be the least k such that there is a coalgebra $\alpha : A \longrightarrow \mathbb{T}_k A$. Combining the previous Proposition and the remark after Theorem 6, we see that $\text{tw}(A) < \kappa(A) \leq \text{card}(A)$. In fact, $\kappa(A)$ yields an elegant, purely (co)algebraic characterization of treewidth.

Theorem 23. *For all finite structures A :*

$$\kappa(A) = \text{tw}(A) + 1.$$

Proof. We have already seen that $\text{tw}(A) < \kappa(A)$, so it suffices to show that $\kappa(A) \leq \text{tw}(A) + 1$. Suppose that $\text{tw}(A) = k - 1$. This means that there is a tree decomposition (T, a) where T is a tree, and for any node t of T , $a(t)$ is a set of at most k elements of A . By standard means, we can assume that T is a rooted, directed, tree and for any node

t of T , there is at most one element of A that appears in $a(t)$ that does not appear in $a(s)$ for any ancestor s of t (for instance, we could take (T, a) to be a *nice* decomposition as defined in [18]).

For each $x \in A$, we can then define $t(x)$ to be the least (in the tree order) element of T such that $x \in a(t(x))$. The properties of a tree decomposition guarantee that this is well defined, and our assumption that (T, a) is nice ensures that t is injective. The function t induces a partial order on A : $x \leq y$ if and only if $t(x)$ is an ancestor of $t(y)$. It is easily seen that this is a tree order. We now proceed to define a map $i : A \rightarrow [k]$ by induction on this order. Suppose $i(x)$ has been defined for all $x < y$. In particular, this means (since t is injective) that $i(x)$ is defined for all $x \in a(t(y))$ other than y . Since $a(t(y))$ has at most k elements, there is at least one value $p \in [k]$ such that $i(x) \neq p$ for all $x \in a(t(y))$ other than y . We set $i(y)$ to be the least such p .

It can now be verified that (A, \leq, i) is a k -traversal. By Theorem 6, this gives a coalgebra map $A \rightarrow \mathbb{T}_k A$, and establishes that $\kappa(A) \leq k$. \square

7. Cores and pebble number

We now turn to the question of when there exists a homomorphism (not necessarily a coalgebra) $A \rightarrow \mathbb{T}_k A$. We define the *pebble number* $\pi(A)$ of a finite structure A to be the least k such that $A \rightarrow \mathbb{T}_k A$.

The core of a structure is a key notion in graph theory and finite model theory [19]. We define the core of a finite structure A to be a substructure $A' \subseteq A$ which admits a homomorphism $A \rightarrow A'$, but for which no proper substructure does. We say that a structure is a core if it is the core of itself. The following summarizes some basic properties of cores.

- Proposition 24.**
- 1) Every finite structure has a core.
 - 2) The core of a structure is unique up to isomorphism.
 - 3) A substructure is a core of A if and only if it is a retract of A , and minimal among retracts of A with respect to inclusion.
 - 4) A structure is a core if and only if it admits no proper retracts.
 - 5) If A' is the core of A and B' the core of B , then $A \rightarrow B \iff A' \rightarrow B'$.

By virtue of these results, we can write $\text{core}(A)$ for the core of a finite structure A .

We shall make use of some results from the literature. We denote the *canonical conjunctive query* for a finite structure A by Q^A . This is a positive existential first-order sentence. We have the following classic result of Chandra and Merlin [2]:

Theorem 25. *The following are equivalent for finite structures A, B :*

- 1) $A \rightarrow B$.
- 2) $Q^B \models Q^A$.

- 3) $B \models Q^A$.

We also have the following results due to Dalmau, Kolaitis and Vardi [20].

Theorem 26. *The following are equivalent for finite structures A :*

- $\text{tw}(\text{core}(A)) < k$.
- Q^A is logically equivalent to a sentence of $\exists^+ \mathcal{L}_{\omega, \omega}^k$, the existential positive fragment of k -variable first-order logic.

We say that a sentence φ of the existential positive fragment of first-order logic can be transformed to another sentence ψ by *standard rewriting* if there are a sequence of steps, each of which is either (1) associative-commutative rewriting of conjunctions, (2) change of bound variable in an existential quantifier, or (3) replacing $\exists x. (\phi_1 \wedge \phi_2)$ by $\exists x. \phi_1 \wedge \phi_2$, where x does not occur free in ϕ_2 , which transform φ into ψ . The following result from [20] is a refinement of the previous theorem.

Theorem 27. *The following are equivalent for finite structures A :*

- $\text{tw}(A) < k$.
- Q^A can be transformed by standard rewriting to a sentence of $\exists^+ \mathcal{L}_{\omega, \omega}^k$.

The following result shows how $\mathbb{T}_k A$ captures the information about A which is expressible in $\exists^+ \mathcal{L}_{\omega, \omega}^k$.

Proposition 28. *Let A be a finite structure. Every sentence of $\exists^+ \mathcal{L}_{\omega, \omega}^k$ satisfied by A is also satisfied by a finite substructure of $\mathbb{T}_k A$.*

Proof. We list the variables of $\exists^+ \mathcal{L}_{\omega, \omega}^k$ as x_1, \dots, x_k . Given a formula $\varphi(\mathbf{x})$ with a list of distinct free variables $\mathbf{x} = x_{i_1}, \dots, x_{i_l}$, with $0 \leq l \leq k$, and a corresponding semantic tuple $\mathbf{a} = a_1, \dots, a_l$, we define a tuple $\mathbf{s}(\mathbf{x}, \mathbf{a})$ with $\mathbf{s}_0 = []$, and $\mathbf{s}_{i+1} = \mathbf{s}_i[(i+1, \mathbf{a}_i)]$, for $i = 0, \dots, l-1$.

We show by induction on formulas $\varphi(\mathbf{x}) \in \exists^+ \mathcal{L}_{\omega, \omega}^k$ that $A, \mathbf{a} \models \varphi(\mathbf{x})$ implies that $\mathbb{T}_k A, \mathbf{s} \models \varphi$ where $\mathbf{s} = \mathbf{s}(\mathbf{x}, \mathbf{a})$. The base case is for atomic formulas $R(\mathbf{x})$. This follows immediately from the definition of the relational structure on \mathbb{T}_k . For $\exists x. \varphi$, $A, \mathbf{a} \models \exists x. \varphi$ iff for some a : $A, \mathbf{a}a \models \varphi$, which by induction hypothesis implies that $\mathbb{T}_k A, \mathbf{s}s \models \varphi$, iff $\mathbb{T}_k A, \mathbf{s} \models \exists x. \varphi$. Finally, for a conjunction $\varphi \wedge \psi$, $A, \mathbf{a}bc \models \varphi \wedge \psi$ iff $A, \mathbf{a}b \models \varphi(\mathbf{x}, \mathbf{y})$, and $A, \mathbf{a}c \models \psi(\mathbf{x}, \mathbf{z})$, where the variables common to φ and ψ are listed as \mathbf{x} . By induction hypothesis, $\mathbb{T}_k A, \mathbf{s}t \models \varphi$ and $\mathbb{T}_k A, \mathbf{s}u \models \psi$. Then $\mathbb{T}_k A, \mathbf{s}' \models \varphi \wedge \psi$, where $\mathbf{s}' = \mathbf{s}(\mathbf{xyz}, \mathbf{abc})$. \square

Proposition 29. *For all finite structures A , $A \equiv^{\exists^+} \mathbb{T}_k A$.*

Proof. By Proposition 14, we just need to exhibit homomorphisms $\mathbb{T}_k T_k A \rightarrow A$ and $\mathbb{T}_k A \rightarrow \mathbb{T}_k A$. For the first, we can compose the counit maps for $\mathbb{T}_k A$ and for A , and for the second, we can just take the identity. \square

We can use these results to obtain our characterization of the relationship between treewidth and the pebble number.

Theorem 30. For all finite structures A :

$$\pi(A) = \text{tw}(\text{core}(A)) + 1.$$

Proof. It clearly suffices to prove that:

$$\text{tw}(\text{core}(A)) < k \iff \pi(a) \leq k.$$

Suppose firstly that $\text{tw}(\text{core}(A)) < k$. By Theorem 26, Q^A is logically equivalent to a sentence of $\exists^+ \mathcal{L}_{\omega, \omega}^k$. By Proposition 28, Q^A is satisfied by a finite substructure T of $\mathbb{T}_k A$. By Theorem 25, there is a homomorphism $A \rightarrow T \hookrightarrow \mathbb{T}_k A$.

Now suppose that there is a homomorphism $A \rightarrow \mathbb{T}_k A$. The image of this homomorphism is a finite substructure T of $\mathbb{T}_k A$. By Theorem 25, $Q^T \models Q^A$. By Proposition 22, $\text{tw}(T) < k$. Hence by Theorem 27, Q^T can be transformed by standard rewriting to a formula in $\exists^+ \mathcal{L}_{\omega, \omega}^k$. We can restrict the counit map $\varepsilon_A : \mathbb{T}_k A \rightarrow A$ to a homomorphism $T \rightarrow A$. Using Theorem 25 again, $Q^A \models Q^T$. Hence Q^A and Q^T are logically equivalent, which shows that Q^A is logically equivalent to a sentence in $\exists^+ \mathcal{L}_{\omega, \omega}^k$. By Theorem 26, this implies that $\text{tw}(\text{core}(A)) < k$. \square

This result shows that the existence of a homomorphism from A to $\mathbb{T}_k A$ characterizes the treewidth of the core of A . The stronger hypothesis of having a coalgebra map $A \rightarrow \mathbb{T}_k A$ characterizes the treewidth of A itself.

8. Strong k -consistency and contextuality

We now look at *strong k -consistency*, which is a fundamental notion in constraint satisfaction [21]. The non-uniform CSP problem $\text{CSP}(B)$ can be posed for a fixed finite structure B as the existence of a homomorphism $A \rightarrow B$ for a given finite structure A . An instance A is i -consistent if for every homomorphism h from an $(i-1)$ -element substructure of A to B and every $a \in A$, there is a homomorphism on a substructure of A including $\text{dom}(h) \cup \{a\}$ extending h . It is strongly k -consistent if it is i -consistent for all $i \leq k$.

This leads us to the following notion of *consistency number* $\text{sc}_B(A)$ of a finite structure A : the greatest k such that there is a homomorphism $\mathbb{T}_k A \rightarrow B$, i.e. $A \rightarrow_k B$. Note that, if we parameterize the pebble number to define $\pi_B(A)$ as the smallest k such that there is a homomorphism $A \rightarrow \mathbb{T}_k B$, then π_B and sc_B are formally dual.

If $k \geq \text{card}(A)$, then $A \rightarrow_k B$ if and only if $A \rightarrow B$, so we define $\text{sc}_B(A)$ as the largest $k \leq n$ such that $A \rightarrow_k B$. Also, because of the graded structure of the monads \mathbb{T}_k , note that it is strong k -consistency which we are capturing. Once we have $A \rightarrow_k B$, then also $A \rightarrow_l B$ for $l \leq k$.

The interesting case is where $\text{sc}_B(A) < \text{card}(A)$. This is the case where we can get partial solutions for up to k variables, but no global solutions. We note that this case covers a variety of examples of *contextuality* in quantum mechanics, as shown in [22]. In particular, many proofs of the Kochen-Specker theorem [23] provide examples of this phenomenon. As an illustration, we consider the Mermin Magic Square [24]:

A	B	C
D	E	F
G	H	I

The constraints are that each row and the first two columns have even parity, and the final column has odd parity. This translates into 6 linear equations over \mathbb{Z}_2 :

$$\begin{aligned} A \oplus B \oplus C &= 0 & A \oplus D \oplus G &= 0 \\ D \oplus E \oplus F &= 0 & B \oplus E \oplus H &= 0 \\ G \oplus H \oplus I &= 0 & C \oplus F \oplus I &= 1 \end{aligned}$$

Of course, these equations are not satisfiable in \mathbb{Z}_2 . The system is 8-consistent but not globally consistent. The significance of this construction in quantum mechanics is that we can interpret the variables with quantum observables in such a way that the specified constraints are exactly the predicted behaviour of measuring the observables. The global inconsistency corresponds to the impossibility of a non-contextual explanation for this behaviour. For more on this topic, see [25], [26].

We can also make a simple connection between pebble number and consistency number.

Proposition 31. If $\pi(A) \leq \text{sc}_B(A)$, then $A \rightarrow B$.

Proof. If $\pi(A) \leq \text{sc}_B(A)$, then for some k ,

$$A \rightarrow \mathbb{T}_k A \rightarrow B.$$

\square

8.1. Cai-Fürer-Immerman construction

Fix a relational signature with two ternary relations R_0 and R_1 . Consider the structure Z_2 in this signature with universe $\{0, 1\}$ where $R_0 = \{(i, j, k) \mid i \oplus j \oplus k = 0\}$ and $R_1 = \{(i, j, k) \mid i \oplus j \oplus k = 1\}$. Then, a system of equations over \mathbb{Z}_2 (as in the example above) can be seen as a structure A whose universe is the set of variables and $(v_1, v_2, v_3) \in R_0$ if $v_1 \oplus v_2 \oplus v_3 = 0$ is an equation in the system and $(v_1, v_2, v_3) \in R_1$ if $v_1 \oplus v_2 \oplus v_3 = 1$ is an equation in the system. It is then immediate that $A \rightarrow Z_2$ if and only if the system is solvable. This is the classical example of a constraint satisfaction problem that is not solvable by k -local consistency tests (though it is polynomial-time solvable by Gaussian elimination). Moreover, it is also known that the class of satisfiable systems of equations is not invariant under \equiv^{C^k} for any k . This is shown by Atserias et al. [27], based on the construction of Cai et al. [28] of a polynomial-time decidable class of graphs not definable in fixed-point logic with counting. Here, we give an account of this construction in the categorical framework we have developed.

Proposition 32. Suppose $A \rightarrow_k Z_2$. Then there is a pair of structures A_0 and A_1 with $A_0 \cong_K A_1$, $A \rightarrow A_0$, and $A_1 \rightarrow Z_2$.

Proof. We define A_0 and A_1 to both have universe $A \times \{0, 1\}$. The interpretation of the relations in A_0 is given by:

$$R_0 = \{(a, i), (b, j), (c, k) \mid (a, b, c) \in R_0^A \wedge i \oplus j \oplus k = 0\}$$

$$R_1 = \{(a, i), (b, j), (c, k) \mid (a, b, c) \in R_1^A \wedge i \oplus j \oplus k = 0\}.$$

The interpretation of the relations in A_1 is given by:

$$R_0 = \{(a, i), (b, j), (c, k) \mid (a, b, c) \in R_0^A \wedge i \oplus j \oplus k = 0\}$$

$$R_1 = \{(a, i), (b, j), (c, k) \mid (a, b, c) \in R_1^A \wedge i \oplus j \oplus k = 1\}.$$

The homomorphism $A \longrightarrow A_0$ is given by $a \mapsto (a, 0)$ and the homomorphism $A_1 \longrightarrow Z2$ by $(a, i) \mapsto i$.

Now, suppose $f : \mathbb{T}_k A \longrightarrow Z2$ is a homomorphism. We aim to use this to define an isomorphism $f' : \mathbb{T}_k A_0 \longrightarrow \mathbb{T}_k A_1$. For any $s \in \mathbb{T}_k A_0$, we write πs to denote the sequence in $\mathbb{T}_k A$ obtained by replacing each move $(p, (a, i))$ in s by (p, a) . Let s be the sequence $(p_1, (a_1, i_1)), \dots, (p_n, (a_n, i_n))$ and let s_1, \dots, s_n be the sequence of its non-empty prefixes. We then define $f'(s)$ to be the sequence in $\mathbb{T}_k A_1$ whose j th component is $(p_j, a_j, i_j \oplus f(\pi s_j))$ and verify that this is an isomorphism. \square

9. Modalities

Another facet of monads and comonads is their rôle in categorical logic, as categorified interpretations of S4 modalities [29]. The pebbling comonad \mathbb{T}_k provides a semantics for a modality \Box_k , which we can read as: “it is k -locally the case that ...”. We can take e.g. first-order logic as a base, and extend it with the formation rule that if φ is a formula and at most k free variables, then $\Box_k \varphi$ is a formula.

We extend the usual definition of satisfaction of a formula

$$A, \mathbf{a} \models \varphi(\mathbf{x})$$

where A is a structure, and \mathbf{a} a sequence of elements of A providing interpretations for the free variables \mathbf{x} of φ , with the following clause:

$$A, \mathbf{a} \models \Box_k \varphi \iff \mathbb{T}_k A, \alpha(\mathbf{a}) \models \varphi$$

where if $\mathbf{a} = \langle a_1, \dots, a_l \rangle$, $l \leq k$, $\alpha(a_i) = [(p_1, a_1), \dots, (p_i, a_i)]$, $i = 1, \dots, l$.

We can read off a number of properties of this modality directly from the comonadic structure. Firstly, the S4 axioms are valid in this semantics *with respect to* $\exists^+ \mathcal{L}_{\omega, \omega}^k$ formulas φ :

$$(T) \models \Box_k \varphi \supset \varphi$$

$$(4) \models \Box_k \varphi \supset \Box_k \Box_k \varphi$$

We also have the grading axiom:

$$\models \Box_k \varphi \supset \Box_l \varphi \quad (k \leq l)$$

Using Proposition 28, for formulas φ of $\exists^+ \mathcal{L}_{\omega, \omega}^k$ we also have the following:

$$\models \varphi \supset \Box_k \varphi.$$

If we want to reason about structures with bounded treewidth, we can use the results of the previous section to introduce suitable assumptions.

We can also consider a modality \Diamond , which can be read as “it is homomorphically true that”. This has the following semantics:

$$A, \mathbf{a} \models \Diamond \varphi \iff \exists h : A \rightarrow B : B, h(\mathbf{a}) \models \varphi.$$

This satisfies the dual S4 axioms with respect to $\exists^+ \mathcal{L}_{\omega, \omega}^k$ formulas:

$$(T) \models \varphi \supset \Diamond \varphi$$

$$(4) \models \Diamond \Diamond \varphi \supset \Diamond \varphi$$

Whereas the logics usually considered in finite model theory are interpreted in one structure at a time, these modalities allows us to navigate around the category of structures, treating them as different “possible worlds”.

10. Conclusions

While we are not aware of any closely related work, the work of Bojanczyk [30] and of Adamek *et al* [31] on recognizable languages over monads is in a broadly kindred spirit. The aim of these works is to use monads as a unifying notion for the many variations on the theme of recognizability. Also, we can mention the use of comonads in functional programming [32], [33]. In particular, the comonadic structure of lists is studied in [34]. The forgetful functor from structures to sets carries the pebbling comonad to this list comonad (although concretely the correspondence is up to list reversal).

The ideas developed in this paper suggest a number of further developments. Firstly, can we find similar characterizations of other forms of games, and connections with the corresponding logical notions? In fact, there is a natural family of comonads corresponding to Ehrenfeucht-Fraïssé games, which can be used to characterize equivalence up to given quantifier rank, using ideas analogous to those in Section 5.2. There is also an unravelling comonad, which arises in the analysis of bisimulation for modal and guarded logics [35]. Details of these constructions will appear in a sequel to the present paper.

There are many other natural lines for future investigation. Can we give categorical characterizations of key combinatorial parameters, in a similar fashion to treewidth? Can we use categorical limit and colimit constructions, possibly in an enriched form, to connect with the combinatorial developments in [36]? Can we leverage results in descriptive complexity to give categorical descriptions of complexity classes?

Another project is to analyze Rossman’s theorem on homomorphism preservation [37] from a categorical viewpoint. It seems that many of the technical notions developed in the proof of this result are susceptible of a more abstract formulation, which may lead to more conceptual proofs, and the possibility of finding wider applicability for the techniques developed there.

The ideas described in this paper provide a connection between two broad themes within the field of logic in computer science: the interaction of logic with the analysis of algorithms and computational complexity, and the study of the semantics of programs and processes. Examples of such connections are still fairly rare, and we hope that the ideas introduced here can be developed further in a fruitful fashion.

Acknowledgement: The work reported here was initiated at the programme on *Logical Structures in Computation* at the Simons Institute for the Theory of Computing, Berkeley.

References

- [1] L. Libkin, *Elements of finite model theory*. Springer Science & Business Media, 2013.
- [2] A. K. Chandra and P. M. Merlin, “Optimal implementation of conjunctive queries in relational data bases,” in *Proceedings of the ninth annual ACM symposium on Theory of computing*. ACM, 1977, pp. 77–90.
- [3] T. Feder and M. Y. Vardi, “The computational structure of monotone monadic SNP and constraint satisfaction: A study through Datalog and group theory,” *SIAM Journal on Computing*, vol. 28, no. 1, pp. 57–104, 1998.
- [4] S. Abiteboul, R. Hull, and V. Vianu, *Foundations of databases: the logical level*. Addison-Wesley Longman Publishing Co., Inc., 1995.
- [5] P. G. Kolaitis and M. Y. Vardi, “Conjunctive-query containment and constraint satisfaction,” in *Proceedings of the seventeenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems*. ACM, 1998, pp. 205–213.
- [6] —, “On the expressive power of Datalog: tools and a case study,” in *Proceedings of the ninth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems*. ACM, 1990, pp. 61–71.
- [7] E. Moggi, “Notions of computation and monads,” *Information and computation*, vol. 93, no. 1, pp. 55–92, 1991.
- [8] S. Brookes and S. Geva, “Computational comonads and intensional semantics,” in *Applications of Categories in Computer Science*, P. J. M.P. Fourman and A. Pitts, Eds. Cambridge University Press, 1992, pp. 1–40.
- [9] P. Wadler, “Monads for functional programming,” in *International School on Advanced Functional Programming*. Springer, 1995, pp. 24–52.
- [10] B. C. Pierce, “A taste of category theory for computer scientists,” CMU, Tech. Rep., 1988.
- [11] S. Abramsky and N. Tzevelekos, “Introduction to categories and categorical logic,” in *New structures for physics*. Springer, 2010, pp. 3–94.
- [12] H. Kleisli, “Every standard construction is induced by a pair of adjoint functors,” *Proceedings of the American Mathematical Society*, vol. 16, no. 3, pp. 544–546, 1965.
- [13] A. Atserias, A. Dawar, and P. G. Kolaitis, “On preservation under homomorphisms and unions of conjunctive queries,” *Journal of the ACM (JACM)*, vol. 53, no. 2, pp. 208–237, 2006.
- [14] L. Hella, “Logical hierarchies in PTIME,” *Information and Computation*, vol. 129, no. 1, pp. 1–19, 1996.
- [15] N. Immerman, “Relational queries computable in polynomial time,” *Information and control*, vol. 68, no. 1-3, pp. 86–104, 1986.
- [16] —, “Upper and lower bounds for first order expressibility,” *Journal of Computer and System Sciences*, vol. 25, no. 1, pp. 76–98, 1982.
- [17] N. Robertson and P. D. Seymour, “Graph minors. ii. Algorithmic aspects of tree-width,” *Journal of algorithms*, vol. 7, no. 3, pp. 309–322, 1986.
- [18] T. Kloks, *Treewidth, Computations and Approximations*, ser. Lecture Notes in Computer Science. Springer, 1994, vol. 842.
- [19] P. Hell and J. Nešetřil, “The core of a graph,” *Discrete Mathematics*, vol. 109, no. 1, pp. 117–126, 1992.
- [20] V. Dalmau, P. G. Kolaitis, and M. Y. Vardi, “Constraint satisfaction, bounded treewidth, and finite-variable logics,” in *International Conference on Principles and Practice of Constraint Programming*. Springer, 2002, pp. 310–326.
- [21] P. G. Kolaitis and M. Y. Vardi, “A game-theoretic approach to constraint satisfaction,” in *AAAI/IAAI*, 2000, pp. 175–181.
- [22] S. Abramsky and A. Brandenburger, “The sheaf-theoretic structure of non-locality and contextuality,” *New J. Phys.*, vol. 13, no. 11, p. 113036, 2011.
- [23] S. Kochen and E. P. Specker, “The problem of hidden variables in quantum mechanics,” *J. Math. Mech.*, vol. 17, no. 1, pp. 59–87, 1967.
- [24] N. D. Mermin, “Simple unified form for the major no-hidden-variables theorems,” *Physical Review Letters*, vol. 65, no. 27, p. 3373, 1990.
- [25] S. Abramsky, R. S. Barbosa, K. Kishida, R. Lal, and S. Mansfield, “Contextuality, cohomology and paradox,” in *24th EACSL Annual Conference on Computer Science Logic (CSL 2015)*, ser. Leibniz International Proceedings in Informatics (LIPIcs), S. Kreutzer, Ed., vol. 41. Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2015, pp. 211–228.
- [26] S. Abramsky, G. Gottlob, and P. G. Kolaitis, “Robust constraint satisfaction and local hidden variables in quantum mechanics,” in *IJCAI*, vol. 13, 2013, pp. 440–446.
- [27] A. Atserias, A. Bulatov, and A. Dawar, “Affine systems of equations and counting infinitary logic,” *Theoretical Computer Science*, vol. 410, no. 18, pp. 1666–1683, 2009.
- [28] J.-Y. Cai, M. Fürer, and N. Immerman, “An optimal lower bound on the number of variables for graph identification,” *Combinatorica*, vol. 12, no. 4, pp. 389–410, 1992.
- [29] G. M. Bierman and V. C. de Paiva, “On an intuitionistic modal logic,” *Studia Logica*, vol. 65, no. 3, pp. 383–416, 2000.
- [30] M. Bojańczyk, “Recognisable languages over monads,” in *International Conference on Developments in Language Theory*. Springer, 2015, pp. 1–13.
- [31] H. Urbat, J. Adámek, L.-T. Chen, and S. Milius, “One Eilenberg theorem to rule them all,” *arXiv preprint arXiv:1602.05831*, 2016.
- [32] T. Uustalu and V. Vene, “Comonadic notions of computation,” *Electronic Notes in Theoretical Computer Science*, vol. 203, no. 5, pp. 263–284, 2008.
- [33] —, “The essence of dataflow programming,” in *Central European Functional Programming School*. Springer, 2005, pp. 135–167.
- [34] D. Orchard, “Programming contextual computations,” *University of Cambridge, Computer Laboratory, Technical Report*, no. UCAM-CL-TR-854, 2014.
- [35] E. Grädel and M. Otto, “The freedoms of (guarded) bisimulation,” in *Johan van Benthem on Logic and Information Dynamics*. Springer, 2014, pp. 3–31.
- [36] J. Nešetřil and P. O. De Mendez, “A unified approach to structural limits, and limits of graphs with bounded tree-depth,” *arXiv preprint arXiv:1303.6471*, 2013.
- [37] B. Rossman, “Homomorphism preservation theorems,” *Journal of the ACM (JACM)*, vol. 55, no. 3, p. 15, 2008.