THE PEBBLE-RELATION COMONAD IN FINITE MODEL THEORY

YOÀV MONTACUTE ● [a] AND NIHIL SHAH ● [b]

[a] University of Cambridge, United Kingdom
   e-mail address: yoav.montacute@cl.cam.ac.uk

[b] University of Oxford, United Kingdom
   e-mail address: nihil.shah@cs.ox.ac.uk

ABSTRACT. The pebbling comonad, introduced by Abramsky, Dawar and Wang, provides a categorical interpretation for the k-pebble games from finite model theory. The coKleisli category of the pebbling comonad specifies equivalences under different fragments and extensions of infinitary k-variable logic. Moreover, the coalgebras over this pebbling comonad characterise treewidth and correspond to tree decompositions. In this paper we introduce the pebble-relation comonad, which characterises pathwidth and whose coalgebras correspond to path decompositions. We further show that the existence of a coKleisli morphism in this comonad is equivalent to truth preservation in the restricted conjunction fragment of k-variable infinitary logic. We do this using Dalmau's pebble-relation game and an equivalent all-in-one pebble game. We then provide a similar treatment to the corresponding coKleisli isomorphisms via a bijective version of the all-in-one pebble game. Finally, we show as a consequence a new Lovász-type theorem relating pathwidth to the restricted conjunction fragment of k-variable infinitary logic with counting quantifiers.

1. INTRODUCTION

Model theory is a field in which mathematical structures are not seen as they really are, i.e. up to isomorphism, but through the fuzzy reflection imposed by their definability in some logic $\mathscr{L}$. Namely, given two structures over the same signature $\mathcal{A}$ and $\mathcal{B}$, model theory is concerned with equivalence under the relation

$$\mathcal{A} \equiv^{\mathscr{L}} \mathcal{B} = \forall \phi \in \mathscr{L}, \ \mathcal{A} \models \phi \Leftrightarrow \mathcal{B} \models \phi.$$

Historically, a central theme in model theory has been to find syntax-free characterisations of these equivalences. This is exemplified by the Keisler-Shelah theorem [She71] for first-order logic. These equivalences are also characterised by model-comparison, or Spoiler-Duplicator games. Spoiler-Duplicator games are graded by a (typically finite) ordinal that corresponds to a grading of some syntactic resource. For example, the k-pebble game introduced by Immerman [Imm82] characterises equivalence in infinitary logic graded by the number of variables k.

Abramsky, Dawar and Wang [ADW17] provided this game, and two similar variants, with a categorical interpretation in terms of morphisms which involve the pebbling comonad $\mathbb{P}_k$.

Since then similar game comonads were discovered for the Ehrenfeucht-Fraïssé games [AS18], modal-bisimulation games [AS18], games for guarded logics [AM21] and games for finite-variable logics with generalised quantifiers [CD21]. In all of these cases the coalgebras over the game comonad correspond to decompositions of structures such as tree decompositions of width $< k$ and forest covers of height $\leq k$, in the cases of the pebbling comonad and the Ehrenfeucht-Fraïssé comonad respectively. As an immediate corollary of these correspondences, we obtain alternative and novel definitions for associated graph parameters such as treewidth in the case of the pebbling comonad and tree-depth in the case of the Ehrenfeucht-Fraïssé comonad.

Abramsky et al. [ADW17] proved that the coKleisli morphisms associated with $\mathbb{P}_k$ correspond to Duplicator's winning strategies in the one-sided k-pebble game. This one-sided game, introduced by Kolaitis and Vardi [KV90], was used to study expressivity in Datalog, preservation of existential positive sentences of k-variable logic $\exists^+ \mathscr{L}_\infty$ and **P**-tractable constraint satisfaction problems. Dalmau [Dal05] developed an analogous one-sided pebble-relation game that he used to study expressivity in linear Datalog, preservation of a restricted conjunction fragment of existential positive k-variable logic $\exists^+ \lambda \mathscr{L}^k$ and **NL**-tractable constraint satisfaction problems. Further, Abramsky et al. [ADW17] demonstrated that, for the pebbling comonad, the coKleisli isomorphism associated with $\mathbb{P}_k$ corresponds to Hella's bijective variant of the k-pebble game [Hel96]. However, the bijective version of the pebble-relation game has not been explored.

In this paper we widen the domain of these game comonads by introducing the pebble-relation comonad $\mathbb{PR}_k$, where the coalgebras correspond to path decompositions of width $< k$. This yields a new definition for pathwidth. Moreover, just as coKleisli morphisms of $\mathbb{P}_k$ correspond to Duplicator's winning strategies in the one-sided k-pebble game, we show that coKleisli morphisms of $\mathbb{PR}_k$ correspond to Duplicator's winning strategies in Dalmau's k-pebble relation game [Dal05]. We do this by introducing an equivalent game which, as will be detailed later, is much more suitable for our purpose. This game is called the *all-in-one k-pebble game* and the existence of a winning strategy for Duplicator in this game is equivalent to truth preservation in the restricted conjunction fragment of k-variable infinitary logic. Inspired by the isomorphism result for $\mathbb{P}_k$, we define a bijective variant of this game. Consequently, the coKleisli isomorphism allows us to obtain a new characterisation of equivalence in the, heretofore unexplored, restricted conjunction fragment of k-variable infinitary logic with counting quantifiers. It is worth mentioning that equivalence in this fragment may yield a new isomorphism approximation method [WL68], similarly to how equivalences in $k + 1$-variable counting logic can be algorithmically interpreted as graphs which are indistinguishable by the polynomial time k-Weisfeiler-Lehman isomorphism approximation method [Gro17].

As one immediate consequence of this comonadic interpretation, we will demonstrate that equivalence in the restricted conjunction fragment of k-variable logic with counting quantifiers is equivalent to counting homomorphisms from structures of pathwidth $< k$. This applies and expands the theoretical picture established in a recent paper by Dawar, Jakl and Reggio [DJR21] regarding a categorical Lovász-type theorem. It is also intertwined with the growing interest in the relationship between homomorphism counting and pathwidth as discussed for example in Dell, Grohe and Ratten [DGR18].

Outline. Section 2 introduces the necessary preliminaries and notation we use throughout the paper. It includes the necessary background on Spoiler-Duplicator game comonads and

the associated results that accompany them. Section 3 discusses the pebble-relation comonad $\mathbb{PR}_k$ and demonstrates the relationship between $\mathbb{PR}_k$ and $\mathbb{P}_k$. Section 4 proves that coalgebras over $\mathbb{PR}_k$ correspond to path decompositions of width $< k$, which provides a coalgebraic characterisation of pathwidth. Section 5 introduces Dalmau's one-sided pebble-relation game and the equivalent all-in-one pebble game in which Duplicator's winning strategies are captured through morphisms involving $\mathbb{PR}_k$. The same treatment is then given to the bijective version of the all-in-one pebble game in which Duplicator's winning strategies are captured by isomorphisms involving $\mathbb{PR}_k$. The section concludes with issuing a Lovász-type theorem for pathwidth. Section 6 concludes the paper with a summary of the results and a discussion on further research directions involving the pebble-relation comonad.

2. Preliminaries

In this section, we will establish some notational preliminaries and provide a short introduction to the relevant concepts in category theory and finite model theory used throughout the paper.

2.1. **Set notation.** Given a partially ordered set $(X, \leq)$ and $x \in X$, the *down-set* and *up-set* of x are $\downarrow x = \{y \in X \mid y \leq x\}$ and $\uparrow x = \{y \in X \mid y \geq x\}$ respectively. A partially ordered set $(X, \leq)$ is a *linear order*, or a *directed path*, if for every pair of elements $x, y \in X$, $x \leq y$ or $y \leq x$, i.e. every two elements are comparable. If $(T, \leq)$ is a partially ordered set such that for every $x \in T$, $\downarrow x$ is linearly ordered by $\leq$, then $\leq$ *forest orders* T and $(T, \leq)$ is a *forest*. The *height of an element* $x \in T$ of a forest $(T, \leq)$ is the cardinality of $\downarrow x \backslash \{x\}$. The *height of a forest* $(T, \leq)$ is the maximal height of an element $x \in T$. If $(T, \leq)$ is a forest and there exists a least element $\bot \in T$ such that, for all $x \in T$, $\bot \leq x$, then $\leq$ *tree orders* T and $(T, \leq)$ is a *tree*. For the purpose of forest orders, we will use interval notation, i.e. $(x, x'] = \{y \mid x < y \leq x'\}$ and $[x, x'] = \{y \mid x \leq y \leq x'\}$.

For a positive integer n, we denote by $[\mathsf{n}]$ the set $\{1, \ldots, n\}$. As a convention, we consider $[\mathsf{n}]$ as having the usual order on segments of natural numbers $\leq$. Given a set A, we denote the set of finite sequences of elements A as A^* and non-empty finite sequences as A^+. The set of sequences of length $\leq k$ is denoted by $A^{\leq k}$. We denote the sequence of elements $a_1, \ldots, a_n \in A$ as $[a_1, \ldots, a_n]$ and the empty sequence as ϵ. We write $|s| = n$ for the length of a sequence $s = [a_1, \ldots, a_n]$. Given two sequences $s, t \in A^*$, we denote the concatenation of s followed by t as st. If s is such that there exists a (possibly empty) sequence s' where $ss' = t$, then we write $s \sqsubseteq t$. Observe that $\sqsubseteq$ defines a relation on sequences and tree orders A^* and forest orders A^+. Hence, by the interval notation defined for forest orders, if $s \sqsubseteq t$, then $(s, t]$ denotes the suffix of s in t. For $s = [a_1, \ldots, a_n]$ and $i, j \in [\mathsf{n}]$, we define a range notation where $s(i, j] = [a_{i+1}, \ldots, a_j]$ if $i < j$ or ϵ if $i \geq j$, and $s[i, j] = [a_i, \ldots, a_j]$ if $i \leq j$ or ϵ if $i > j$.

2.2. **Category theory.** We assume familiarity with the standard category-theoretic notions of category, functor and natural transformation. Given a category $\mathfrak{C}$, we denote the class of its objects by $\mathfrak{C}_0$ and the class of its morphisms by $\mathfrak{C}_1$. We define the notion of a relative comonad which weakens the endofunctor requirement of the standard comonad. The relative comonad is the dual of the *relative monad* introduced in [ACU14]. Given a functor $J : \mathfrak{J} \to \mathfrak{C}$, a *relative comonad on* J is a triple $(\mathbb{T}, \varepsilon, (\cdot)^*)$, where $\mathbb{T} : \mathfrak{J}_0 \to \mathfrak{C}_0$ is an object mapping,

$\varepsilon_X : \mathbb{T}X \to JX \in \mathfrak{C}_1$ is defined for every object $X \in \mathfrak{J}_0$ and $f^* : \mathbb{T}X \to \mathbb{T}Y \in \mathfrak{C}_1$ is a coextension map, for each $f : \mathbb{T}X \to JY \in \mathfrak{C}_1$, satisfying the following equations:

$$\varepsilon_X^* = \mathsf{id}_{\mathbb{T}X}; \qquad \varepsilon \circ f^* = f; \qquad (g \circ f^*)^* = g^* \circ f^*.$$

These equations allow us to extend the object mapping $\mathbb{T}$ to a functor, where $\mathbb{T}f = (Jf \circ \varepsilon_X)^*$ for $f : X \to Y$. For every relative comonad $(\mathbb{T}, \varepsilon, (\cdot)^*)$ over $J : \mathfrak{J} \to \mathfrak{C}$, we can define an associated *coKleisli category*, denoted by $\mathcal{K}(\mathbb{T})$, where

- $\mathcal{K}(\mathbb{T})_0$ is the same as the class of objects $\mathfrak{J}_0$;
- $\mathcal{K}(\mathbb{T})_1$ are morphisms of type $f : \mathbb{T}X \to JY \in \mathfrak{C}_1$;
- The composition $g \circ_{\mathcal{K}} f : \mathbb{T}X \to JZ$ of two morphisms $f : \mathbb{T}X \to JY$ and $g : \mathbb{T}Y \to JZ$ is given by

$$\mathbb{T}X \xrightarrow{f^*} \mathbb{T}Y \xrightarrow{g} JZ;$$

- The identity morphisms are given by the counit

$$\varepsilon_X : \mathbb{T}X \to JX.$$

The ordinary notion of a *comonad in coKleisli form* [Man76] and the corresponding Kleisli category can be recovered when $\mathfrak{J} = \mathfrak{C}$ and $J = \mathsf{id}_{\mathfrak{C}}$. Observe that given a comonad $(\mathbb{T} : \mathfrak{C} \to \mathfrak{C}, \varepsilon, (\cdot)^*)$ and a functor $J : \mathfrak{J} \to \mathfrak{C}$, the functor $\mathbb{T}^J = \mathbb{T} \circ J$ can be made into a relative comonad $(\mathbb{T}^J, \varepsilon', (\cdot)')$ on J, where $\varepsilon'_X = \varepsilon_{JX}$ and the coextension mapping $()'$ is defined for $f : \mathbb{T}^J X \to JY$ to be the same as f^*, i.e. $f' = f^* : \mathbb{T}^J \to \mathbb{T}^J Y$.

Given an ordinary comonad in coKleisli form, we can define a comultiplication morphism $\delta_X : \mathbb{T}X \to \mathbb{T}\mathbb{T}X$, where $\delta_X = (\mathsf{id}_{\mathbb{T}X})^*$, which satisfies the following equations:

$$\mathbb{T}\delta_X \circ \delta_X = \delta_{\mathbb{T}X} \circ \delta_X; \qquad \mathbb{T}\varepsilon_X \circ \delta_X = \varepsilon_{\mathbb{T}X} \circ \delta_X = \mathsf{id}_{\mathbb{T}X}.$$

The triple $(\mathbb{T} : \mathfrak{C} \to \mathfrak{C}, \varepsilon, \delta)$, where $\mathbb{T}$ is a functor, is a *comonad in standard form*. Its coKleisli form can be recovered by defining the coextension mapping $(\cdot)^*$ as $f^* = \mathbb{T}f \circ \delta$.

Given a comonad $(\mathbb{T} : \mathfrak{C} \to \mathfrak{C}, \varepsilon, \delta)$, a *coalgebra* over $\mathbb{T}$ is defined as a pair $(A, \alpha : A \to \mathbb{T}A)$, where $A \in \mathfrak{C}_0$ and $\alpha \in \mathfrak{C}_1$, satisfying the following equations:

$$\delta_A \circ \alpha = \mathbb{T}\alpha \circ \alpha; \qquad \varepsilon_A \circ \alpha = \mathsf{id}_A.$$

The *category of coalgebras*, or *Eilenberg-Moore category*, associated with a comonad $\mathbb{T} : \mathfrak{C} \to \mathfrak{C}$, denoted by $\mathcal{EM}(\mathbb{T})$, is defined as

- $\mathcal{EM}(\mathbb{T})_0$ consists of coalgebras $(A, \alpha : A \to \mathbb{T}A)$;
- A morphism of type $h : (A, \alpha) \to (B, \beta) \in \mathcal{EM}(\mathbb{T})_1$ is a morphism $h : A \to B \in \mathfrak{C}_1$ such that $\mathbb{T}h \circ \alpha = \beta \circ h$;
- Identity and composition are defined as in $\mathfrak{C}$.

A *resolution* of $\mathbb{T}$ is an adjunction $L : \mathfrak{C} \to \mathfrak{D} \dashv R : \mathfrak{D} \to \mathfrak{C}$ that gives rise to $\mathbb{T}$ as a comonad. A resolution $L \dashv R$ is a *comonadic adjunction* if the comparison functor $K : \mathfrak{D} \to \mathcal{EM}(\mathbb{T})$ is an equivalence.

2.3. **Finite model theory.** We fix a vocabulary σ of relational symbols R, each with a positive arity $\rho(R)$. If R has arity $\rho(R)$, then R is called an $\rho(R)$-*ary relation*. A σ-*structure* $\mathcal{A}$ is specified by a universe of elements A and interpretations $R^{\mathcal{A}} \subseteq A^{\rho(R)}$ for each relation $R \in \sigma$. We use calligraphic letters $(\mathcal{A}, \mathcal{B}, \mathcal{C}, \text{etc.})$ to denote σ-structures and we use roman letters $(A, B, C, \text{etc.})$ to denote their underlying universes of elements.

Let $\mathcal{A}$ and $\mathcal{B}$ be σ-structures. If $B \subseteq A$ and $R^{\mathcal{B}} \subseteq R^{\mathcal{A}}$ for every relation $R \in \sigma$, then $\mathcal{B}$ is a σ-*substructure* of $\mathcal{A}$. If $B \subseteq A$, then we can form the B *induced* σ-*substructure* $\mathcal{B} = \mathcal{A}|_B$

with universe B and interpretations $R^{\mathcal{B}} = R^{\mathcal{A}} \cap B^{\rho(R)}$ for each relation $R \in \sigma$. The graph $\mathcal{G}(\mathcal{A}) = (A, \frown)$ is the Gaifman graph of $\mathcal{A}$, where $a \frown a'$ iff $a = a'$ or a, a' appear in some tuple of $R^{\mathcal{A}}$ for some $R \in \sigma$.

A σ-*morphism from $\mathcal{A}$ to $\mathcal{B}$*, or *homomorphism*, $h : \mathcal{A} \to \mathcal{B}$ is a set function from A to B such that $R^{\mathcal{A}}(a_1, \ldots, a_{\rho(R)})$ implies $R^{\mathcal{B}}(h(a_1), \ldots, h(a_{\rho(R)}))$ for every relation $R \in \sigma$. We will denote the category of σ-structures with σ-morphisms by $\mathfrak{R}(\sigma)$. The full subcategory of σ-structures with finite universes is denoted by $\mathfrak{R}_f(\sigma)$. If there exists a homomorphism $h : \mathcal{A} \to \mathcal{B}$, we write $\mathcal{A} \to \mathcal{B}$. In the category $\mathfrak{R}(\sigma)$, an isomorphism $f : \mathcal{A} \to \mathcal{B}$ is a bijective homomorphism that reflects relations, i.e. for each relation $R \in \sigma$,

$$R^{\mathcal{A}}(a_1, \ldots, a_m) \Leftrightarrow R^{\mathcal{B}}(f(a_1), \ldots, f(a_{\rho(R)})).$$

2.4. **Logical fragments.** We are mainly concerned with fragments of infinitary logic. The infinitary logic $\mathscr{L}_\infty$ has the standard syntax and semantics of first-order logic, but where disjunctions and conjunctions are allowed to be taken over arbitrary sets of formulas. We denote formulas with free variables among $\vec{x} = (x_1, \ldots, x_n)$ by $\phi(\vec{x})$. If $\mathcal{A}$ is a σ-structure, $\vec{a} \in A^n$, $\phi(\vec{x}) \in \mathscr{L}_\infty$ and $\mathcal{A}, \vec{a}$ satisfies $\phi(\vec{x})$, then we write $\mathcal{A}, \vec{a} \models \phi(\vec{x})$.

For vectors of variables and elements $\vec{x} = (x_1, \ldots, x_n)$, we write $[\vec{x}] = \{x_1, \ldots, x_n\}$ for the underlying *support* of the vector. The infinitary logic can be graded into k-variable fragments, denoted by $\mathscr{L}^k$, where formulas contain at most k-many variables.

For every logic $\mathscr{J}$ considered throughout the paper, we will also be interested in two variants. The first variant is the existential positive fragment $\exists^+ \mathscr{J}$, where we only consider formulas constructed using existential quantifiers, disjunctions, conjunctions and atomic formulas. The second variant involves a restricted conjunction. A *restricted conjunction* is a conjunction of the form $\bigwedge \Psi$, where Ψ is a set of formulas satisfying the following condition:

(R) At most one formula $\psi \in \Psi$ with quantifiers is not a sentence.

The main goal of this paper is to study the *restricted conjunction logics* $\exists^+ \curlywedge \mathscr{L}^k$ and $\# \curlywedge \mathscr{L}^k$ (Section 5.2).

Given two σ-structures $\mathcal{A}$ and $\mathcal{B}$, if for all sentences $\phi \in \mathscr{J}(\sigma)$, $\mathcal{A} \models \phi$ implies $\mathcal{B} \models \phi$, then we write $\mathcal{A} \Rightarrow^{\mathscr{J}} \mathcal{B}$. If $\mathcal{A} \Rightarrow^{\mathscr{J}} \mathcal{B}$ and $\mathcal{B} \Rightarrow^{\mathscr{J}} \mathcal{A}$, then we write $\mathcal{A} \equiv^{\mathscr{J}} \mathcal{B}$. For logics $\mathscr{J}$ closed under negation, we have that $\mathcal{A} \Rightarrow^{\mathscr{J}} \mathcal{B}$ implies $\mathcal{A} \equiv^{\mathscr{J}} \mathcal{B}$.

2.5. **Spoiler-Duplicator games.** The relations $\Rightarrow^{\mathscr{J}}$ and $\equiv^{\mathscr{J}}$ for specific choices of $\mathscr{J}$ are characterised, in a syntax-free fashion, by Spoiler-Duplicator games (also called model-comparison games or Ehrenfeucht-Fraïssé style games). In each game, we consider two σ-structures, $\mathcal{A}$ and $\mathcal{B}$, and two players: Spoiler, who tries to show that two structures are different under $\mathscr{J}$; and Duplicator, who tries to show the two structures are the same under $\mathscr{J}$ (i.e. $\mathcal{A} \equiv^{\mathscr{J}} \mathcal{B}$). The standard Ehrenfeucht-Fraïssé game appears in many introductory texts on model theory (see e.g. [Lib04]). The model-comparison games we are interested in are modified versions of the one-sided k-pebble game [Imm82] and the bijective k-pebble game [Hel96]. Each game is played in a number of rounds.

For the one-sided k-pebble game $\exists^+ \mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$, which characterises $\Rightarrow^{\mathscr{L}^k}$, both Spoiler and Duplicator have a set of k pebbles. At each n-th round such that $n \in \omega$:

- Spoiler places a pebble $p_n \in [\mathsf{k}]$ on an element $a_n \in A$. If the pebble $p_n \in [\mathsf{k}]$ is already placed on an element of A, Spoiler moves the pebble from that element to the chosen element;
- Duplicator places the pebble $p_n \in [\mathsf{k}]$ on an element $b_n \in B$.

At the end of the n-th round, we have a pair of sequences $s = [(p_1, a_1), \ldots, (p_n, a_n)]$ and $t = [(p_1, b_1), \ldots, (p_n, b_n)]$. For every p in the set $\mathsf{pebbles}(s) = \{p_1, \ldots, p_n\}$ of pebbles appearing in s, let $a^p = \mathsf{last}_p(s)$ and $b^p = \mathsf{last}_p(t)$ be the last elements pebbled with p in s and t respectively. If the relation $\gamma_n = \{(a^p, b^p) \mid p \in \mathsf{pebbles}(s)\}$ is a partial homomorphism from $\mathcal{A}$ to $\mathcal{B}$, then Duplicator wins the n-th round of the k-pebble game.[1] Duplicator has a winning strategy in $\exists^+ \mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ if for every round $n \in \omega$, and for every move by Spoiler in the n-th round, Duplicator has a winning move.

For the bijective k-pebble game $\#\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$, Spoiler wins automatically if there is no bijection between A and B. At each n-th round such that $n \in \omega$:

- Spoiler chooses a pebble $p \in [\mathsf{k}]$;
- Duplicator responds with a bijection $f_n : A \to B$ consistent with the previously placed pebbles, i.e. for every $q \neq p$ such that $(a^q, b^q) \in \gamma_{n-1}$, $f_n(a^q) = b^q$;
- Spoiler places the pebble p on an element $a \in A$. Duplicator places the pebble p on the element $f(a) \in B$.

Duplicator wins the n-th round of $\#\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ if the relation $\gamma_n = \{(a^p, f_n(a^p)) \mid p \in [\mathsf{k}]\}$ is a partial isomorphism. As before, Duplicator has a winning strategy if she can keep playing forever.

Capturing these games as constructions on the category of relational structures has been the underlying theme of the research program motivating this paper.

2.6. **Spoiler-Duplicator game comonads.** The larger research program motivating this paper is the discovery of Spoiler-Duplicator game comonads $\mathbb{C}_k$ associated with logics $\mathscr{J}$ graded by a resource k. Though these Spoiler-Duplicator game comonads have no general definition that applies to all cases, different indexed families of comonads $\mathbb{C}_k$ over $\mathfrak{R}(\sigma)$ exhibit the same pattern of results: the morphism power theorem, the isomorphism power theorem and the coalgebra characterisation theorem [ADW17, AS18, CD21, AM21]. In this section, we will review the general schema for each of these results.

Morphism and Isomorphism Power Theorems. Given a logic $\mathscr{J}_k$, graded by some syntactic resource k (e.g. number of variables, quantifier rank and modal depth) and corresponding model-comparison game $\mathbf{G}_k$, the Spoiler-Duplicator comonad associated with $\mathbf{G}_k$ is an indexed family of comonads $\mathbb{C}_k$ over $\mathfrak{R}(\sigma)$ [AS18]. We can then leverage the coKleisli category $\mathcal{K}(\mathbb{C}_k)$ associated with $\mathbb{C}_k$ to capture the relation $\Rrightarrow^{\exists^+ \mathscr{J}_k}$ of the sentences in the existential positive fragment of $\mathscr{J}_k$ limited by the resource k.

Theorem 2.1 (Morphism Power Theorem). *For all σ-structures $\mathcal{A}$ and $\mathcal{B}$, the following are equivalent:*

(1) *Duplicator has a winning strategy in $\exists^+ \mathbf{G}_k(\mathcal{A}, \mathcal{B})$.*
(2) $\mathcal{A} \Rrightarrow^{\exists^+ \mathscr{J}_k} \mathcal{B}$.
(3) *There exists a coKleisli morphism $f : \mathbb{C}_k \mathcal{A} \to \mathcal{B}$.*

---

[1] When Spoiler (resp. Duplicator) makes a choice that results in her winning a round, we often say that this choice is a *winning move* for her.

For example, in the case of the pebbling comonad $\mathbb{P}_k$, a coKleisli morphism $\mathbb{P}_k\mathcal{A} \to \mathcal{B}$ corresponds to a winning strategy for Duplicator in $\exists^+\mathbf{Peb}_k(\mathcal{A},\mathcal{B})$ and characterises the relation $\Rrightarrow^{\exists^+\mathscr{L}^k}$ [ADW17]. The indexing of $\mathbb{C}_k$ models resources in the corresponding game. Namely, for $k \leq l$, there is a comonad inclusion $\mathbb{C}_k \hookrightarrow \mathbb{C}_l$. Interpreting the inclusion in terms of the game corresponds to the fact that Spoiler playing with k resources is a special case of Spoiler playing with $l \geq k$ resources. This means that the smaller the k is, the easier it is to find morphisms of type $\mathbb{C}_k\mathcal{A} \to \mathcal{B}$.

Ostensibly, the asymmetry of a coKleisli morphism means that only the "forth" aspect of the game $\mathbf{G}_k$ (the game $\exists^+\mathbf{G}_k$) can be captured by this comonadic approach. A natural candidate for capturing the symmetric game would be to consider the symmetric relation of coKleisli isomorphism, i.e. there exist morphisms $f : \mathbb{C}_k\mathcal{A} \to \mathcal{B}$ and $g : \mathbb{C}_k\mathcal{B} \to \mathcal{A}$ such that $g \circ_{\mathcal{K}} f = \varepsilon_{\mathcal{A}}$ and $f \circ_{\mathcal{K}} g = \varepsilon_{\mathcal{B}}$. However, the isomorphisms in $\mathcal{K}(\mathbb{C}_k)$ characterise the equivalence relation $\equiv^{\#\mathscr{I}_k}$ for the logic $\mathscr{I}_k$ extended with counting quantifiers. This results in a logic stronger than $\mathscr{I}_k$.

Theorem 2.2 (Isomorphism Power Theorem). *For all finite σ-structures $\mathcal{A}$ and $\mathcal{B}$, the following are equivalent:*

(1) *Duplicator has a winning strategy in $\#\mathbf{G}_k(\mathcal{A},\mathcal{B})$.*
(2) $\mathcal{A} \equiv^{\#\mathscr{I}_k} \mathcal{B}$.
(3) *There exists a coKleisli isomorphism $f : \mathbb{C}_k\mathcal{A} \to \mathcal{B}$.*

For example, in the case of the pebbling comonad $\mathbb{P}_k$, a coKleisli isomorphism $\mathbb{P}_k\mathcal{A} \to \mathcal{B}$ corresponds to a winning strategy for Duplicator in $\#\mathbf{Peb}_k(\mathcal{A},\mathcal{B})$ and characterises the relation $\equiv^{\#\mathscr{L}^k}$ [ADW17].

Coalgebras and Adjunctions. A natural inquiry regarding the comonads $\mathbb{C}_k$ is to investigate the category of coalgebras $\mathcal{EM}(\mathbb{C}_k)$. It turns out that for all the cases of $\mathbb{C}_k$ constructed from some game $\mathbf{G}_k$, the coalgebras $\mathcal{A} \to \mathbb{C}_k\mathcal{A}$ correspond to forest covers of structures $\mathcal{A}$.

Define a *forest cover for $\mathcal{A}$* to be a forest $(S, \leq)$, where $A \subseteq S$, such that if $a \frown a' \in \mathcal{A}$, i.e. a and a' are related in some tuple of $\mathcal{A}$, then $a \leq a'$ or $a' \leq a$. In this paper, we assume that all forest covers are *tight* forest covers, i.e. where $S = A$.

Theorem 2.3 (coalgebra characterisation). *The following statements are equivalent:*

(1) $\mathcal{A}$ *has a forest cover of parameter $\leq k$.*
(2) *There exists a coalgebra $\alpha : \mathcal{A} \to \mathbb{C}_k\mathcal{A}$.*

For every game comonad $\mathbb{C}_k$, we define the $\mathbb{C}$-*coalgebra number of $\mathcal{A}$* to be the least k such that there exists a coalgebra $\mathcal{A} \to \mathbb{C}_k\mathcal{A}$. We denote it by $\kappa^{\mathbb{C}}(\mathcal{A})$. This definition together with the coalgebra characterisation theorems allow us to obtain new alternative definitions for various combinatorial invariants of relational structures.

For example, in the case of the Ehrenfeucht-Fraïssé comonad $\mathbb{E}_k$, a coalgebra over $\mathcal{A}$ corresponds to a forest cover of $\mathcal{A}$ with height $\leq k$. Consequently, $\kappa^{\mathbb{E}}(\mathcal{A})$ is equal to the tree-depth of $\mathcal{A}$ [AS18]. Moreover, in the case of the pebbling comonad, a coalgebra over $\mathcal{A}$ corresponds to a forest cover over $\mathcal{A}$ with an additional pebbling function $p : A \to [\mathsf{k}]$ that encodes the data of a tree decomposition for $\mathcal{A}$ of width $< k$. Consequently, $\kappa^{\mathbb{P}}(\mathcal{A})$ corresponds to the treewidth of $\mathcal{A}$ [ADW17].

The bijective correspondence between forest covers of a certain type and coalgebras can be extended to isomorphisms of the respective categories. In the language of adjunctions, for every $\mathbb{C}_k$, we can form the category of forest-ordered σ-structures $\mathfrak{T}^{\mathbb{C}_k}(\sigma)$ consisting of

- $\mathfrak{T}^{\mathbb{C}_k}(\sigma)_0$ are pairs $(\mathcal{A}, \leq)$, where $\mathcal{A} \in \mathfrak{R}(\sigma)_0$ and $\leq$ forest-orders the universe A of $\mathcal{A}$ such that the following condition holds:

$$\text{(E) if } a \frown a' \in \mathcal{A}, \text{ then } a \leq a' \text{ or } a' \leq a.$$

- $\mathfrak{T}^{\mathbb{C}_k}(\sigma)_1$ are σ-morphisms $f : \mathcal{A} \to \mathcal{B} \in \mathfrak{R}(\sigma)_1$ that preserve roots and the covering relation of $\leq$. That is, f preserves immediate successors in $\leq$.
- The identity and composition are inherited from $\mathfrak{R}(\sigma)$.

Evidently, there exists a forgetful functor $U_k : \mathfrak{T}^{\mathbb{C}_k}(\sigma) \to \mathfrak{R}(\sigma)$. We may view $\mathbb{C}_k$ as the comonad arising from constructing a functor $F_k : \mathfrak{R}(\sigma) \to \mathfrak{T}^{\mathbb{C}_k}(\sigma)$ that is right adjoint to U_k and where the adjunction $U_k \dashv F_k$ is comonadic. That is, $\mathfrak{T}^{\mathbb{C}_k}(\sigma)$ and $\mathcal{EM}(\mathbb{C}_k)$ are equivalent categories.

For example, the category of coalgebras for the Ehrenfeucht-Fraïssé comonad is equivalent to the category of forest covers of height $\leq k$ [AS21]. Similarly, the category of coalgebras $\mathcal{EM}(\mathbb{P}_k)$ for the pebbling comonad is equivalent to the category of k-pebble forest covers [AS21].

For each of the result schemata (morphism power theorem, isomorphism power theorem and coalgebra characterisation theorem), we mentioned the corresponding result for the pebbling comonad. However, we would like to stress that analogous results hold for all of the comonads previously discussed in the literature [ADW17, AS18, AM21, CD21]. Obtaining these results from a general definition is an active topic of research.

3. Pebble-Relation Comonad

The pebble-relation comonad, which captures pathwidth, is closely related to the pebbling comonad $(\mathbb{P}_k, \varepsilon, ()^*)$ that is used to define treewidth [ADW17, AS18]. Given a σ-structure $\mathcal{A}$, we define the universe of $\mathbb{P}_k\mathcal{A}$ as $([\mathsf{k}] \times A)^+$. The counit morphism $\varepsilon_\mathcal{A} : \mathbb{P}_k\mathcal{A} \to \mathcal{A}$ is defined as the second component of the last element of the sequence, i.e. $\varepsilon_\mathcal{A}[(p_1, a_1), \ldots, (p_n, a_n)] = a_n$. The coextension of a morphism $f : \mathbb{P}_k\mathcal{A} \to \mathcal{B}$ is defined as

$$f^*[(p_1, a_1), \ldots, (p_n, a_n)] = [(p_1, b_1), \ldots, (p_n, b_n)],$$

where $b_i = f[(p_1, a_1), \ldots, (p_i, a_i)]$, for all $i \in [\mathsf{n}]$.

To define the σ-structure on $\mathbb{P}_k\mathcal{A}$, we will need to introduce some new notation. The mapping $\pi_A : \mathbb{P}_k A \to [\mathsf{k}]$ is defined as the first component of the last element of the sequence, i.e. $\pi_A[(p_1, a_1), \ldots, (p_n, a_n)] = p_n$. Suppose $R \in \sigma$ is an m-ary relation. Then $R^{\mathbb{P}_k\mathcal{A}}(s_1, \ldots, s_m)$ iff

(1) $\forall i, j \in [\mathsf{m}], s_i \sqsubseteq s_j$ or $s_j \sqsubseteq s_i$;   *(pairwise comparability)*
(2) If $s_i \sqsubseteq s_j$, then $\pi_A(s_i)$ is not in $(s_i, s_j]$,
    and similarly for $s_j \sqsubseteq s_i$;   *(active pebble)*
(3) $R^\mathcal{A}(\varepsilon_\mathcal{A}(s_1), \ldots, \varepsilon_\mathcal{A}(s_m))$.   *(compatibility)*

The elements of $\mathbb{P}_k\mathcal{A}$ can be seen as the set of Spoiler plays in the k-pebble game on the structure $\mathcal{A}$. The additional active pebble condition models how Spoiler's choice of k pebble placements amounts to moving a k-sized window of variables that are assigned to elements in the structure $\mathcal{A}$. When working with $\mathbb{P}_k\mathcal{A}$, as before, we will use $\mathsf{pebbles}(s)$ to denote the set of pebbles appearing in s, $\mathsf{last}_p(s)$ to denote the last element $a \in \mathcal{A}$ in the sequence s pebbled by $p \in \mathsf{pebbles}(s)$. The set of *active elements of* s is defined as $\mathsf{Active}(s) = \{\mathsf{last}_p(s) \mid p \in \mathsf{pebbles}(s)\}$. Given two sequences $s \in \mathbb{P}_k A$ and $t \in \mathbb{P}_k B$, we define the relation of pebbled elements $\gamma_{s,t} = \{(\mathsf{last}_p(s), \mathsf{last}_p(t))\} \subseteq A \times B$.

We now introduce our main construction, which is the family of comonads $(\mathbb{PR}_k, \varepsilon, ()^*)$, for every $k \in \omega$ over $\mathfrak{R}(\sigma)$. Given a σ-structure $\mathcal{A}$, we define the universe of $\mathbb{PR}_k \mathcal{A}$ as

$$\mathbb{PR}_k A = \{([(p_1, a_1), \ldots, (p_n, a_n)], i)\},$$

where $(p_j, a_j) \in [\mathsf{k}] \times A$ and $i \in [\mathsf{n}]$. Intuitively, $\mathbb{PR}_k A$ is the set of Spoiler plays in the k-pebble game paired with an index denoting a move of the play. The counit morphism $\varepsilon_\mathcal{A} : \mathbb{PR}_k \mathcal{A} \to \mathcal{A}$ is defined as $\varepsilon_\mathcal{A}([(p_1, a_1), \ldots, (p_n, a_n)], i) = a_i$. The coextension of a morphism $f : \mathbb{PR}_k \mathcal{A} \to \mathcal{B}$ is defined as

$$f^*([(p_1, a_1), \ldots, (p_n, a_n)], i) = ([(p_1, b_1), \ldots, (p_n, b_n)], i),$$

where $b_j = f([(p_1, a_1), \ldots, (p_n, a_n)], j)$ for all $j \in [\mathsf{n}]$.

We can interpret the relations on $\mathbb{PR}_k \mathcal{A}$ in a similar manner to the interpretations given for $\mathbb{P}_k \mathcal{A}$: Suppose $R \in \sigma$ is an m-ary relation. Then $R^{\mathbb{PR}_k \mathcal{A}}((s_1, i_1), \ldots, (s_m, i_m))$ iff

(1) $\forall j \in [\mathsf{m}], s_j = s$;   (*equality*)
(2) $\pi_A(s, i_j)$ does not appear in $s(i_j, i]$;   (*active pebble*)
(3) $R^\mathcal{A}(\varepsilon_\mathcal{A}(s, i_1), \ldots, \varepsilon_\mathcal{A}(s, i_m))$;   (*compatibility*)

where $i = \max\{i_1, \ldots, i_m\}$ and $\pi_A : A \to [\mathsf{k}]$ is defined as

$$\pi_A([(p_1, a_1), \ldots, (p_n, a_n)], i) = p_i.$$

The active pebble and compatibility conditions play a similar role as in the definition of $\mathbb{P}_k \mathcal{A}$. As we will see in Section 5, the equality condition ensures that Duplicator must respond to a (full) Spoiler play in one round.

Proposition 3.1. $(\mathbb{PR}_k, \varepsilon, ()^*)$ *is a comonad in coKleisli form.*

Proof. It is easy to see that $\mathbb{PR}_k$ is a lifting of the pointed-list (see e.g. [Orc14]) comonad over **Set**. In order to show that $\mathbb{PR}_k$ is a comonad over $\mathfrak{R}(\sigma)$, we need to show that $\varepsilon_\mathcal{A}$ is a σ-morphism and that for every $f \in \mathfrak{R}(\sigma)_1$, $f^* \in \mathfrak{R}(\sigma)_1$. The fact that $\varepsilon_\mathcal{A}$ is a σ-morphism follows from the compatibility condition in the definition of $R^{\mathbb{PR}_k \mathcal{A}}$. To show that f^* is a σ-morphism, suppose that $R^{\mathbb{PR}_k \mathcal{A}}((s, i_1), \ldots, (s, i_m))$ and that $s = [(p_1, a_1), \ldots, (p_n, a_n)]$. Consider $t = [(p_1, b_1), \ldots, (p_n, b_n)]$, where $f(s, i) = b_i$. It follows that $f^*(s, i) = (t, i)$. By construction, $\pi_\mathcal{A}(s, i_j) = \pi_\mathcal{B}(t, i_j)$, so $((t, i_1), \ldots, (t, i_m))$ satisfies the active pebble condition. Since f is a σ-morphism, $R^{\mathbb{PR}_k \mathcal{A}}((s, i_1), \ldots, (s, i_m))$ and $\varepsilon_\mathcal{B}(t, i) = b_i$, we have that $R^\mathcal{B}(b_{i_1}, \ldots, b_{i_m})$. Therefore, the compatibility condition holds, $R^{\mathbb{PR}_k \mathcal{B}}((t, i_1), \ldots, (t, i_m))$ and f^* is a σ-morphism. $\square$

We also define for every $n \in \omega$, $\mathbb{PR}_{k,n} \mathcal{A} \subseteq \mathbb{PR}_k \mathcal{A}$ to be the substructure induced by pairs (s, i), where s is of length $\leq n$. The triple $(\mathbb{PR}_{k,n}, \varepsilon, ()^*)$, where ε and $()^*$ are restricted to the substructures of the form $\mathbb{PR}_{k,n} \mathcal{A}$, is also a comonad.

There is a comonad morphism $\nu : \mathbb{PR}_k \to \mathbb{P}_k$ with components $\nu_\mathcal{A} : \mathbb{PR}_k \mathcal{A} \to \mathbb{P}_k \mathcal{A}$, where $\nu_\mathcal{A}(s, i) = s[1, i]$, i.e. the length i prefix of s.

Proposition 3.2. $\nu : \mathbb{PR}_k \to \mathbb{P}_k$ *is a comonad morphism.*

Proof. We must confirm that ν is indeed a natural transformation, i.e. the following diagram commutes in $\mathfrak{R}(\sigma)$, for every $f : \mathcal{A} \to \mathcal{B}$:

$$
\begin{array}{ccc}
\mathbb{PR}_k \mathcal{A} & \xrightarrow{\nu_\mathcal{A}} & \mathbb{P}_k \mathcal{A} \\
{\scriptstyle \mathbb{PR}_k f} \downarrow & & \downarrow {\scriptstyle \mathbb{P}_k f} \\
\mathbb{PR}_k \mathcal{B} & \xrightarrow{\nu_\mathcal{B}} & \mathbb{P}_k \mathcal{B}
\end{array}
$$

It is clear that this diagram commutes by observing that $\mathbb{PR}_k f$ preserves the prefix relation on sequences. We present both comonads in the standard forms $(\mathbb{PR}_k, \varepsilon, \delta)$ and $(\mathbb{P}_k, \varepsilon', \delta')$. To confirm ν is a comonad morphism, we must show that the following diagrams commute in the category of endofunctors over $\mathfrak{R}(\sigma)$:

$$
\begin{array}{ccc}
\mathbb{PR}_k \xrightarrow{\ \nu\ } \mathbb{P}_k & \qquad & \mathbb{PR}_k^2 \xrightarrow{\ \mathbb{PR}_k\nu\ } \mathbb{PR}_k\mathbb{P}_k \xrightarrow{\ \nu\ } \mathbb{P}_k^2 \\
\ \ \searrow{\scriptstyle \varepsilon}\quad \downarrow{\scriptstyle \varepsilon'} & & {\scriptstyle \delta}\uparrow \qquad\qquad\qquad\qquad\qquad {\scriptstyle \delta'}\uparrow \\
\mathsf{Id}_{\mathfrak{R}(\sigma)} & & \mathbb{PR}_k \xrightarrow{\qquad\quad \nu \qquad\quad} \mathbb{P}_k
\end{array}
$$

The diagram on the left states that the last pebbled element of $\nu_{\mathcal{A}}(s, i) = s[1, i]$ is the i-th element pebbled in s which is clear by definition. To confirm the diagram on the right, recall that $\delta_{\mathcal{A}}$ and $\delta'_{\mathcal{A}}$ are the coextensions of $\mathsf{id}_{\mathbb{PR}_k\mathcal{A}}$ and $\mathsf{id}_{\mathbb{P}_k\mathcal{A}}$ respectively. Explicitly, for $s = [(p_1, a_1), \ldots, (p_n, a_n)]$ we have:

$$\delta_{\mathcal{A}}(s, i) = ([(p_1, (s, 1)), \ldots, (p_n, (s, n))], i);$$
$$\delta'_{\mathcal{A}}(s) = [(p_1, s[1, 1]), \ldots, (p_n, s[1, n])].$$

The confirmation is then straightforward:

$$
\begin{aligned}
& \nu_{\mathcal{A}} \circ \mathbb{PR}_k\nu_{\mathcal{A}} \circ \delta_{\mathcal{A}}(s, i) \\
&= \nu_{\mathcal{A}} \circ \mathbb{PR}_k\nu_{\mathcal{A}}([(p_1, (s, 1)), \ldots, (p_n, (s, n))], i) && (1) \\
&= \nu_{\mathcal{A}}([(p_1, \nu_{\mathcal{A}}(s, 1)), \ldots, (p_n, \nu_{\mathcal{A}}(s, n))], i) && (2) \\
&= \nu_{\mathcal{A}}([(p_1, s[1, 1]), \ldots, (p_n, s[1, n])], i) && (3) \\
&= [(p_1, s[1, 1]), \ldots, (p_i, s[1, i])] && (4) \\
&= \delta'_{\mathcal{A}}(s[1, i]) && (5) \\
&= \delta'_{\mathcal{A}} \circ \nu_{\mathcal{A}}(s, i); && (6)
\end{aligned}
$$

where (1) and (5) follow by $\delta'_{\mathcal{A}}$ as above, (3),(4) and (6) follow by the definition of $\nu_{\mathcal{A}}$ and (2) follows by the functoriality of $\mathbb{PR}_k$. $\qquad\square$

4. Coalgebras and Path Decompositions

We return to the initial motivation for constructing the pebble-relation comonad. That is, to give a categorical definition for the combinatorial parameter of pathwidth. There are many different characterisations of pathwidth. The original definition, introduced by Robertson and Seymour in [RS83], is in terms of path decompositions of a structure $\mathcal{A}$. We will show that coalgebras $\mathcal{A} \to \mathbb{PR}_k\mathcal{A}$ correspond to path decompositions of $\mathcal{A}$ of width $< k$. Given the relationship between $\mathbb{PR}_k$ and $\mathbb{P}_k$, our proof is very similar to the one given in [ADW17, AS21] demonstrating that coalgebras over the pebbling comonad correspond to tree decompositions. Intuitively, a sequence of pebbled elements with a common root forms a tree, so the mapping $f : A \to \mathbb{P}_k A$ associates with $a \in A$ a node in that tree. By contrast, a mapping $f : A \to \mathbb{PR}_k A$ associates with an element $a \in A$ a full path of pebbled elements and index into that path. It is this intuition that explains why coalgebras of $\mathbb{P}_k$ correspond to tree decompositions, while coalgebras of $\mathbb{PR}_k$ correspond to path decompositions.

Definition 4.1 (path decomposition). Given a finite σ-structure $\mathcal{A}$, a *path decomposition of $\mathcal{A}$* is a triple $(X, \leq_X, \lambda)$, where $(X, \leq_X)$ is a linearly ordered set and $\lambda : X \to \wp(A)$ is a function satisfying the following conditions:

(PD1) For every $a \in A$, there exists $x \in X$ such that $a \in \lambda(x)$.
(PD2) If $a \frown a' \in \mathcal{A}$, then $a, a' \in \lambda(x)$ for some $x \in X$.
(PD3) For all $y \in [x, x']$, $\lambda(x) \cap \lambda(x') \subseteq \lambda(y)$, where $[x, x']$ is an interval with respect to $\leq_X$.

The *width of a path decomposition* $(X, \leq_X, \lambda)$ for $\mathcal{A}$ is given by $k = \max_{x \in X} |\lambda(x)| - 1$.

Definition 4.2 (pathwidth). The pathwidth of a σ-structure $\mathcal{A}$, denoted by $\mathsf{pw}(\mathcal{A})$, is the least k such that $\mathcal{A}$ has a path decomposition of width k.

In order to show the correspondence between tree decompositions of width $< k$ and coalgebras over $\mathbb{P}_k$, Abramsky et al. made use of an intermediate structure called a k-pebble forest cover[2] [ADW17, AS18]. We will use the analogous notion in the linear-ordered case – the k-pebble linear forest cover.

Definition 4.3 (k-pebble linear forest cover). Given a σ-structure $\mathcal{A}$, a *k-pebble linear forest cover for $\mathcal{A}$* is a tuple $(\mathcal{F}, p)$, where $\mathcal{F} = \{(S_i, \leq_i)\}_{i \in I}$ is a partition of A into (possibly infinitely many) linearly ordered subsets and $p : A \to [\mathsf{k}]$ is a pebbling function such that the following conditions hold:

(FC1) If $a \frown a' \in \mathcal{A}$, then there exists $i \in I$ such that $a, a' \in S_i$.
(FC2) If $a \frown a' \in \mathcal{A}$ and $a \leq_i a'$, then for all $b \in (a, a']_i \subseteq S_i$, $p(b) \neq p(a)$.

We aim to show that the existence of a path decomposition $(X, \leq_X, \lambda)$ of width $< k$ for a finite $\mathcal{A}$ is equivalent to the existence of a k-pebble linear forest cover for $\mathcal{A}$. In order to do so, we need to define a pebbling function $p : A \to [\mathsf{k}]$. We accomplish this by defining functions $\tau_x : \lambda(x) \to [\mathsf{k}]$ on the subset associated with a node $x \in X$ in the path decomposition. If we define these functions in a consistent manner, then they can be 'glued' together to obtain p. We use the definition below to identify consistent families of τ_x.

Definition 4.4 (k-pebbling section family). Given a path decomposition $(X, \leq_X, \lambda)$ of width $< k$ for $\mathcal{A}$, we define a *k-pebbling section family for $(X, \leq_X, \lambda)$* as a family of functions $\{\tau_x : \lambda(x) \to [\mathsf{k}]\}$ indexed by $x \in X$, such that the following conditions hold:

(1) (Locally-injective) For every $x \in X$, τ_x is an injective function.
(2) (Glueability) For every $x, x' \in X$,

$$\tau_x|_{\lambda(x) \cap \lambda(x')} = \tau_{x'}|_{\lambda(x) \cap \lambda(x')}.$$

We show that every path decomposition has a pebbling section family.

Lemma 4.5. *If $(X, \leq_X, \lambda)$ is a path decomposition of width $< k$, then $(X, \leq_X, \lambda)$ has a k-pebbling section family $\{\tau_x\}_{x \in X}$.*

Proof. By induction on the linear order $\leq_X$ for every $x \in X$, we construct a k-pebbling section family $\{\tau_z\}_{z \in \downarrow x}$ for $(Y, \leq_Y, \lambda|_Y)$, where $Y = \downarrow x$ and $\leq_Y = \leq_X \cap (Y \times Y)$.

Base Case: Suppose r is the $\leq_X$-least element. By definition 4, the cardinality of $\lambda(r)$ is $\leq k$, thus we can enumerate the elements via an injective function $\tau_r : \lambda(r) \to [\mathsf{k}]$. By construction, τ_r is injective, so $\{\tau_r\}$ is locally-injective. Glueability follows trivially as every $x \in \downarrow r$ is equal to r.

---

[2]This is originally called k-traversal in [ADW17].

Inductive Step: Let x' be the immediate $\leq_X$-successor of x. By the induction hypothesis, there exists a k-pebbling section family $\{\tau_y\}_{y \in \downarrow x}$. Let $V_{x'}$ denote the subset of 'new' elements $a \in \lambda(x')$ such that $a \notin \lambda(y)$ for every $y <_X x'$.

Claim: For every $y <_X x'$, $\lambda(y) \cap \lambda(x') \subseteq \lambda(x) \cap \lambda(x')$. We show that this claim holds: By (PD3), for all $z \in [y, x']$, $\lambda(y) \cap \lambda(x') \subseteq \lambda(z)$. In particular, since $y \leq_X x <_X x'$, $\lambda(y) \cap \lambda(x') \subseteq \lambda(x)$. Therefore, $\lambda(y) \cap \lambda(x') \subseteq \lambda(x) \cap \lambda(x')$.

From the claim and the definition of $V_{x'}$, we have that $\lambda(x') = (\lambda(x) \cap \lambda(x')) \sqcup V_{x'}$, where $\sqcup$ denotes a disjoint union. This allows us to define $\tau_{x'} : \lambda(x') \to [\mathsf{k}]$ by cases on each of these parts. Fix an injective function $\upsilon_{x'} : V_{x'} \to [\mathsf{k}]$ enumerating $V_{x'}$. Let $\{i_1, \ldots, i_m\}$ enumerate the elements of $[\mathsf{k}]$ not in the image of $\tau_x|_{\lambda(x) \cap \lambda(x')}$. Define $\tau_{x'} : \lambda(x') \to [\mathsf{k}]$ as

$$
\tau_{x'}(a) = \begin{cases} \tau_x(a) & \text{if } a \in \lambda(x) \cap \lambda(x'); \\ i_j & \text{if } a \in V_{x'} \text{ and } \upsilon_{x'}(a) = j. \end{cases}
$$

Injectivity of $\tau_{x'}$ follows from the injectivity of $\tau_x|_{\lambda(x) \cap \lambda(x')}$ and $\upsilon_{x'}$. To verify glueability, it suffices to check that $\tau_{x'}|_{\lambda(y) \cap \lambda(x')} = \tau_y|_{\lambda(y) \cap \lambda(x')}$ for $y \in \downarrow x'$. Since $\{\tau_y\}_{y \in \downarrow x}$ is a k-pebbling section family, for all $y \in \downarrow x$, $\tau_y|_{\lambda(y) \cap \lambda(x)} = \tau_x|_{\lambda(y) \cap \lambda(x)}$. By construction, $\tau_x|_{\lambda(x) \cap \lambda(x')} = \tau_{x'}|_{\lambda(x) \cap \lambda(x')}$. By the claim, we have that $\lambda(y) \cap \lambda(x') \subseteq \lambda(x) \cap \lambda(x')$. Therefore, $\tau_y|_{\lambda(y) \cap \lambda(x')} = \tau_{x'}|_{\lambda(y) \cap \lambda(x')}$. $\qquad\square$

This leads to the following two theorems that are essential to the categorical characterisation of pathwidth.

Theorem 4.6. *The following are equivalent for all finite σ-structures $\mathcal{A}$:*

(1) *$\mathcal{A}$ has a path decomposition of width $< k$.*

(2) *$\mathcal{A}$ has a k-pebble linear forest cover.*

Proof. $(1) \Rightarrow (2)$ Suppose $(X, \leq_X, \lambda)$ is a path decomposition of $\mathcal{A}$ of width $< k$. We define a family of linearly ordered sets $\{(S_i, \leq_i)\}$, where each S_i is the vertex set of a connected component of $\mathcal{G}(\mathcal{A})$. To define the order $\leq_i$, we define an order on $\leq_A$ and realise $\leq_i$ as the restriction of $\leq_A$ to S_i. For every $a \in A$, let $x_a \in X$ denote the $\leq_X$-least element in X such that $a \in \lambda(x_a)$. Such an x_a always exists by (PD1). By lemma 4.5, there exists a k-pebbling section family $\{\tau_x\}_{x \in X}$. We then define $\leq_A$ as follows:

$$
a \leq_A a' \Leftrightarrow x_a <_X x_{a'} \text{ or } \tau_x(a) \leq \tau_x(a') \text{ if } x_a = x_{a'} = x.
$$

The glueability condition on k-pebble section family $\{\tau_x\}$ allows us to obtain a well-defined pebbling function $p : A \to [\mathsf{k}]$ from τ_x. Explicitly, thinking of functions as their sets of ordered pairs, $p = \bigcup_{x \in X} \tau_x$. The tuple $(\{(S_i, \leq_i)\}, p)$ is a k-pebble linear forest cover.

To verify that $\{(S_i, \leq_i)\}$ is a partition of A into linearly ordered subsets, we observe that by construction each S_i is a connected component of A, and so $\{S_i\}$ partitions A. Suppose $a, a' \in S_i$, then by $\leq_X$ being a linear order, either $x_a <_X x_{a'}$, $x_a >_X x_{a'}$, or $x_a = x_{a'}$. If $x_a <_X x_{a'}$ or $x_a >_X x_{a'}$, then $a <_i a'$ or $a >_i a'$ by the definition of $\leq_i$. If $x_a = x_{a'} = x$, then either $\tau_x(a) \leq \tau_x(a')$ or $\tau_x(a) \geq \tau_x(a')$ by the linear ordering $\leq$ on $[\mathsf{k}]$. Hence, in either case, $a \leq_i a'$ or $a \geq_i a'$, so $\leq_i$ is a linear ordering.

To verify (FC1), suppose $a \frown a' \in \mathcal{A}$. This means a, a' are connected in $\mathcal{G}(\mathcal{A})$, and so are in the same connected component S_i of $\mathcal{G}(\mathcal{A})$.

To verify (FC2), suppose $a \frown a' \in S_i$, $a \leq_i a'$, and $b \in (a, a']_i$. By definition of $\leq_i$, $x_a \leq_X x_b \leq_X x_{a'}$. Since $a \frown a'$, by (PD2) there exists $x \in X$ such that $a, a' \in \lambda(x)$. By the definition of $x_{a'}$ as the $\leq_X$-least element of X containing a', we have $x_{a'} \leq_X x$. By

transitivity of $\leq_X$ and $x_a \leq_X x_{a'} \leq_X x$, we have that $x_a \leq_X x$. By (PD3), for every $y \in [x_a, x]_X$, $\lambda(x_a) \cap \lambda(x) \subseteq \lambda(y)$. In particular, for $x_b \in [x_a, x_{a'}]_X \subseteq [x_a, x]_X$, we have $a \in \lambda(x_b)$. Hence, $a, b \in \lambda(x_b)$ and by the injectivity of τ_{x_b}, $\tau_{x_b}(a) \neq \tau_{x_b}(b)$. It follows that $p(a) \neq p(b)$.

$(2) \Rightarrow (1)$ Suppose $\mathcal{A}$ has k-pebble linear forest cover given by the partition $\{(S_i, \leq_i)\}_{i \in [n]}$ and pebbling function $p : A \to [k]$. We define a linearly ordered set $(A, \leq_A)$, where $\leq_A$ is the ordered sum of the family $\{(S_i, \leq_i)\}_{i \in [n]}$. Explicitly, $a \leq_A a'$ iff $a \in S_i$, $a' \in S_j$ for $i < j$ or $a \leq_i a'$ for $i = j$. We say that an element a is an *active predecessor* of a' if $a \leq_A a'$ and for all $b \in (a, a']_A$, $p(b) \neq p(a)$. Let $\lambda(a)$ be the set of active predecessors of a. The triple $(A, \leq_A, \lambda)$ is a path decomposition of $\mathcal{A}$ of width $< k$.

To verify (PD1), observe that, for every $a \in A$, a is an active predecessor of itself since $a \leq_A a$ and $(a, a]_A = \varnothing$. Hence, $a \in \lambda(a)$.

To verify (PD2), suppose $a \frown a' \in \mathcal{A}$. By (FC1), there exists an S_i where $a, a' \in S_i$. Without loss of generality, assume $a \leq_i a'$. By (FC2), for all $b \in (a, a']_i$, $p(b) \neq p(a)$. Therefore, a is an active predecessor of a', so $a, a' \in \lambda(a')$.

To verify (PD3), suppose $b \in [a, a']_A$ and that $c \in \lambda(a) \cap \lambda(a')$. By $c \in \lambda(a)$ and $b \in [a, a']_A$, we have that $c \leq_A a$ and $a \leq_A b$, so $c \leq_A b$. By $c \in \lambda(a')$, for all $d \in (c, a']_A$, $p(c) \neq p(d)$. In particular, for all $d \in (c, b]_A$, $p(c) \neq p(d)$. By definition, c is an active predecessor of b, so $c \in \lambda(b)$.

To verify the width of the decomposition $< k$, we need to show that for every $a' \in A$, $|\lambda(a')| \leq k$. Assume for contradiction that $|\lambda(a')| > k$ for some $a' \in A$. Consider the pebbling function restricted to $\lambda(a')$, $p|_{\lambda(a')} : \lambda(a') \to [k]$. By the Pigeonhole Principle, there must exist $a, c \in \lambda(a')$ with $a \neq c$, such that $p(a) = p(c)$. Without loss of generality assume that $a <_A c$. Since $a \in \lambda(a')$, a is an active predecessor of a', i.e. for all $b \in (a, a']_A$, $p(b) \neq p(a)$. In particular, since $c \in (a, a']_A$, as $a <_A c$ and $c \in \lambda(a')$, then $p(c) \neq p(a)$. This yields a contradiction. $\qquad\square$

In analogy to $\mathbb{P}_k$, where k-pebble forest covers correspond to $\mathbb{P}_k$-coalgebras, we also show that k-pebble linear forest covers correspond to $\mathbb{PR}_k$-coalgebras.

Theorem 4.7. *For all σ-structures $\mathcal{A}$, there is a bijective correspondence between the following:*

(1) *k-pebble linear forest covers of $\mathcal{A}$.*
(2) *coalgebras $\alpha : \mathcal{A} \to \mathbb{PR}_k\mathcal{A}$.*

Proof. $(1) \Rightarrow (2)$ Suppose $\mathcal{A}$ has k-pebble linear forest cover given by the partition into linear orders $\{(S_i, \leq_i)\}_{i \in [n]}$ and pebbling function $p : A \to [k]$. Since each $(S_i, \leq_i)$ is linearly ordered, we can present $(S_i, \leq_i)$ as a chain $a_1 \leq_i \cdots \leq_i a_{m_i}$. We define

$$t_i = [(p(a_1), a_1), \ldots, (p(a_m), a_{m_i})].$$

Intuitively, t_i is the enumeration induced by the linear order $\leq_i$ of S_i zipped with its image under p. For every $a_j \in S_i$ let $\alpha_i : S_i \to \mathbb{PR}_k\mathcal{A}$ be defined as $\alpha_i(a_j) = (t_i, j)$. Let $\alpha : A \to \mathbb{PR}_k A$ be $\alpha = \bigcup_{i \in [n]} \alpha_i$. Since the collection of S_i partitions A, α is well-defined. We must show that the function α is a coalgebra $\alpha : \mathcal{A} \to \mathbb{PR}_k\mathcal{A}$.

To verify that α is indeed a homomorphism, suppose $R \in \sigma$ is an m-ary relation and $R^\mathcal{A}(a_1, \ldots, a_m)$. By (FC1), $\{a_1, \ldots, a_m\} \subseteq S_i$ for some $i \in [n]$. Therefore, for all $j \in [m]$, $\alpha(a_j) = (t_i, z_j)$ for some $z_j \in \{1, \ldots, |t_i|\}$. Let z be the maximal index amongst the z_j. Assume $\alpha(a) = (t, z)$ for $a \in \{a_1, \ldots, a_m\}$. By (FC2), for every a_j and $b \in (a_j, a]_i$,

$p(a_j) \neq p(b)$. Therefore, $\pi_A(t_i, z_j)$ does not appear in $t_i(z_j, z]$. Hence, by the definition of $R^{\mathbb{PR}_k \mathcal{A}}$ and the supposition that $R^{\mathcal{A}}(a_1, \ldots, a_m)$, we obtain $R^{\mathbb{PR}_k \mathcal{A}}((t_i, z_1), \ldots, (t_i, z_m))$.

To verify that α satisfies the counit-coalgebra law, suppose $a \in \mathcal{A}$. Then, by $\{S_i\}$ partitioning A, $a \in S_i$ for some $i \in [\mathsf{n}]$. Suppose a is the j-th element in the $\leq_i$ linear ordering:

$$
\begin{aligned}
\varepsilon_{\mathcal{A}} \circ \alpha(a) &= \varepsilon_{\mathcal{A}} \circ \alpha_i(a) && \text{(by supposition } a \in S_i) \\
&= \varepsilon_{\mathcal{A}}(t_i, j) && \text{(def. } \alpha_i) \\
&= a. && \text{(def. } t_i, \varepsilon_A)
\end{aligned}
$$

To verify that α satisfies the comultiplication-coalgebra law observe that

$$
\begin{aligned}
&\delta_{\mathcal{A}} \circ \alpha(a) \\
&= \delta_{\mathcal{A}} \circ \alpha_i(a) && \text{(by supposition } a \in S_i) \\
&= \delta_{\mathcal{A}}(t_i, j) && \text{(def. } \alpha_i) \\
&= ([(p(a_1), (t_i, 1)), \ldots, (p(a_{m_i}), (t_i, m_i))], j) && \text{(def. } t_i, \delta_{\mathcal{A}}) \\
&= ([(p(a_1), \alpha(a_1)), \ldots, (p(a_{m_i}), \alpha(a_{m_i}))], j) && \text{(def. } \alpha) \\
&= \mathbb{PR}_k \alpha(t_i, j) && \text{(def. } t_i \text{ and functoriality of } \mathbb{PR}_k) \\
&= \mathbb{PR}_k \alpha \circ \alpha(a). && \text{(def. } \alpha)
\end{aligned}
$$

$(2) \Rightarrow (1)$. We define a family of linearly ordered subsets $\{(S_t, \leq_t)\}$ of A:

$$
S_t := \{a \mid \alpha(a) = (t, j) \text{ for some } j \in [|t|]\};
$$
$$
a \leq_t a' \Leftrightarrow \alpha(a) = (t, j), \ \alpha(a') = (t, j') \text{ and } j \leq j'.
$$

By the counit-coalgebra axiom, if $\alpha(a) = (t, j)$, then a is the j-th element in the chain S_t. By the comultiplication-coalgebra axiom, the relation $\leq_t$ is a linear order. Let $p : A \to [\mathsf{k}]$ be $p = \pi_A \circ \alpha$. The tuple $(\{(S_t, \leq_t)\}, p)$ is a k-pebble linear forest cover of $\mathcal{A}$.

To verify (FC1), suppose that $a \frown a'$ and that $\alpha(a) = (t, j)$ and $\alpha(a') = (t', j')$. By α being a homomorphism, $\alpha(a) \frown \alpha(a')$, so $(t, j) \frown (t', j')$. However, by the definition of $R^{\mathbb{PR}_k \mathcal{A}}$, for all $R \in \sigma$, elements of $\mathbb{PR}_k \mathcal{A}$ are only related if they are part of the same pebble play, so $t = t'$. By definition, $a, a' \in S_t$.

To verify (FC2), suppose that $a \frown a'$ with $a \leq_t a'$ and $b \in (a, a']_t$. We want to show that $p(b) \neq p(a)$. Since $a \frown a'$, then there exists some m-tuple $\vec{a} \in R^{\mathcal{A}}$ for some m-ary relation $R \in \sigma$. By α being a homomorphism, there exist $(t, j), (t, j') \in \alpha(\vec{a}) \in R^{\mathbb{PR}_k \mathcal{A}}$ for some $j \leq j' \in [|t|]$ such that $\alpha(a) = (t, j)$ and $\alpha(a') = (t, j')$. Moreover, since $b \in (a, a']_t$, then there exists $i \in (j, j']$ such that $\alpha(b) = (t, i)$. By construction, $p(b) = \pi_A \circ \alpha(b) = \pi_A(t, i)$. However, by the first condition in the definition of $R^{\mathbb{PR}_k \mathcal{A}}$, $\pi_A(t, j)$ does not appear in $t(j, j']$, so $p(b) = \pi_A(t, i) \neq \pi_A(t, j) = p(a)$. $\qquad \square$

We can now use the coalgebra number $\kappa^{\mathbb{PR}}$ to define pathwidth.

Corollary 4.8. *For all finite σ-structures $\mathcal{A}$, $\mathsf{pw}(\mathcal{A}) = \kappa^{\mathbb{PR}}(\mathcal{A}) - 1$.*

Proof. By Theorem 4.6 and Theorem 4.7, a structure $\mathcal{A}$ has a path decomposition of width $< k$ iff $\mathcal{A}$ has a coalgebra $\mathcal{A} \to \mathbb{PR}_k \mathcal{A}$. Hence, $\mathsf{pw}(\mathcal{A}) + 1 \leq \kappa^{\mathbb{PR}}(\mathcal{A})$ by the definition of $\mathsf{pw}(\mathcal{A})$ as the minimal width of a path decomposition for $\mathcal{A}$, and $\kappa^{\mathbb{PR}}(\mathcal{A}) \leq \mathsf{pw}(\mathcal{A}) + 1$ by $\kappa^{\mathbb{PR}}$ being the minimal index for a $\mathbb{PR}_k$-coalgebra of $\mathcal{A}$. $\qquad \square$

We can extend Theorem 4.7 to capture an equivalence between two categories. Consider the category $\mathfrak{T}^{\mathbb{PR}_k}(\sigma)$ of forest-ordered σ-structures $(\mathcal{A}, \leq_A, p)$ with $\mathcal{A} \in \mathfrak{R}(\sigma)_0$ and pebbling function $p : A \to [\mathsf{k}]$ satisfying the following conditions:

(E) If $a \frown a' \in \mathcal{A}$, then either $a \leq_A a'$ or $a' \leq_A a$.

(P) If $a \frown a' \in \mathcal{A}$ and $a \leq_A a'$, then for all $b \in (a, a']_A$, $p(b) \neq p(a)$.

(L) For every $a \in A$, $\uparrow a$ is linearly-ordered by $\leq_A$.

The condition (L), taken together with $\leq_A$ being a tree-order, i.e. for every $a \in A$, $\downarrow a$ is linearly ordered by $\leq_A$, means that $(A, \leq_A)$ is a disjoint union of linearly ordered sets. In other words, $(A, \leq_A)$ is a linear forest. Morphisms in $\mathfrak{T}^{\mathbb{PR}_k}(\sigma)$ are σ-morphisms that preserve paths of $\leq$ and pebbling functions. There is an evident forgetful functor $U_k : \mathfrak{T}^{\mathbb{PR}_k}(\sigma) \to \mathfrak{R}(\sigma)$, which yields that $(\mathcal{A}, \leq_A, p) \mapsto \mathcal{A}$.

Consider the functor $F_k : \mathfrak{R}(\sigma) \to \mathfrak{T}^{\mathbb{PR}_k}(\sigma)$ with object mapping $\mathcal{A} \mapsto (\mathbb{PR}_k\mathcal{A}, \leq^*, \pi_\mathcal{A})$, where $(t, i) \leq^* (t', j)$ iff $t = t'$ and $i \leq j$, and morphism mapping $f \mapsto \mathbb{PR}_k f$. The functor F_k is right adjoint to U_k. In fact, this adjunction yields an equivalence between $\mathfrak{T}^{\mathbb{PR}_k}(\sigma)$ and $\mathcal{EM}(\mathbb{PR}_k)$.

Theorem 4.9. *For each $k > 0$, $U_k \dashv F_k$ is a comonadic adjunction. Moreover, $\mathbb{PR}_k$ is the comonad arising from this adjunction.*

Proof. By theorem 4.7, the objects of the two categories are in bijective correspondence. Thus to complete the argument, we must show that a coalgebra morphism coincides with a $\mathfrak{T}^{\mathbb{PR}_k}(\sigma)$-morphism. Coalgebra morphisms preserve paths and pebble indices, and this is equivalent to the conditions for a $\mathfrak{T}^{\mathbb{PR}_k}(\sigma)$-morphism. $\qquad\square$

5. Morphism and Isomorphism Power Theorems

While the construction of $\mathbb{PR}_k$ is motivated by capturing pathwidth as coalgebras, the other Spoiler-Duplicator game comonads $\mathbb{C}_k$, explored in previous work [ADW17, AS18, CD21, AM21], were constructed by capturing Duplicator's winning strategies as coKleisli morphisms of $\mathbb{C}_k$. In this section we prove a morphism power theorem for $\mathbb{PR}_k$, showing that coKleisli morphisms $\mathbb{PR}_k\mathcal{A} \to \mathcal{B}$ correspond to Duplicator's winning strategies in a one-sided pebble-relation game. This game characterises preservation of sentences in the existential positive k-variable logic with restricted conjunctions $\exists^+ \mathop{\wedge} \mathscr{L}^k$. We then prove an isomorphism power theorem for $\mathbb{PR}_k$ which works analogously by showing that coKleisli isomorphisms $\mathbb{PR}_k\mathcal{A} \to \mathcal{B}$ correspond to Duplicator's winning strategies in a new type of pebble game which will be described in detail. We conclude the section with an ensuing corollary of the exhibited results.

5.1. **Morphism power theorem.** We begin by introducing Dalmau's k-pebble relation game on $\mathcal{A}$ and $\mathcal{B}$, denoted by $\exists^+\mathbf{PebR}_k(\mathcal{A}, \mathcal{B})$, which characterises the relation $\mathcal{A} \Rrightarrow^{\exists^+ \mathop{\wedge} \mathscr{L}^k} \mathcal{B}$.

Definition 5.1 (Dalmau's k-pebble relation game [Dal05])**.** Each round of the game ends with a pair (I, T) where $I \subseteq A$ is a domain such that $|I| \leq k$ and $T \subseteq \mathfrak{R}(\sigma)_1(\mathcal{A}|_I, \mathcal{B})$ is a set of σ-morphisms from $\mathcal{A}|_I$ to $\mathcal{B}$. At round 0, $I = \varnothing$ and $T = \{\lambda\}$, where λ is the unique function from $\varnothing$ to $\mathcal{B}$. At each subsequent round $n > 0$, let (I, T) denote the configuration of the previous round $n - 1$. Spoiler then chooses between two possible moves:

(1) A shrinking move, where Spoiler chooses a smaller domain $I' \subseteq I$.
  - Duplicator then chooses T' to be the set of restrictions of the morphisms in T to I', i.e. $T' = \{h|_{I'} \mid h \in T\}$.
(2) A blowing move when $|I| < k$, where Spoiler chooses a larger domain $I' \supseteq I$ with $|I'| \leq k$.
  - Duplicator then chooses a set of σ-morphisms T' from $\mathcal{A}|_{I'}$ to $\mathcal{B}$ such that $T'|_I = \{h|_I \mid h \in T'\} \subseteq T$.

At the end of the n-th round, the configuration is (I', T'). Duplicator wins the n-th round if T' is non-empty. Duplicator wins the game if she can play forever. Otherwise, Spoiler wins the game.

Intuitively, $\exists^+\mathbf{PebR}_k$ is a revised version of the k-pebble game in which Duplicator is given the advantage of non-determinism. That is, she can respond with a set of partial homomorphisms. At each round she is only obligated to extend some of the partial homomorphisms chosen in the previous round. In fact, we can recover the one-sided k-pebble game $\exists^+\mathbf{Peb}_k$ by insisting that Duplicator always responds with a singleton set T.

In the same paper, Dalmau proves the following result:

Theorem 5.2 [Dal05]. *The following are equivalent for all σ-structures $\mathcal{A}$ and $\mathcal{B}$:*

(1) *Duplicator has a winning strategy in $\exists^+\mathbf{PebR}_k(\mathcal{A}, \mathcal{B})$.*
(2) $\mathcal{A} \Rrightarrow^{\exists^+ \wedge \mathscr{L}^k} \mathcal{B}$.

In fact, Dalmau shows that if Duplicator has a winning strategy, then Duplicator has a default winning strategy, which he calls complete, by playing the set of all partial morphisms T on Spoiler's choice of domain I.

We will show that coKleisli morphisms $\mathbb{PR}_k\mathcal{A} \to \mathcal{B}$ capture Duplicator's winning strategies in $\exists^+\mathbf{PebR}_k(\mathcal{A}, \mathcal{B})$. Note that it is not clear in what way elements of $\mathbb{PR}_k\mathcal{A}$ correspond to Spoiler's plays in $\exists^+\mathbf{PebR}_k$ on $\mathcal{A}$. It is also not clear in what way functions $\mathbb{PR}_k\mathcal{A} \to \mathcal{B}$ correspond to Duplicator's responses in $\exists^+\mathbf{PebR}_k(\mathcal{A}, \mathcal{B})$. Fortunately, we can view elements of $\mathbb{PR}_k\mathcal{A}$ as Spoiler's plays in a different, yet equivalent (Theorem 5.5), game – the all-in-one k-pebble game $\exists^+\mathbf{PPeb}_k(\mathcal{A}, \mathcal{B})$.[3]

Definition 5.3 (all-in-one k-pebble game). The game is played in one round during which:

- Spoiler chooses a sequence $s = [(p_1, a_1), \ldots, (p_n, a_n)]$ of pebble placements on elements of the σ-structure $\mathcal{A}$.
- Duplicator responds with a compatible (same length and corresponding pebble at each index) sequence of pebble placements $t = [(p_1, b_1), \ldots, (p_n, b_n)]$ on elements of $\mathcal{B}$.

Duplicator wins the game if for every index $i \in [\mathsf{n}]$, the relation

$$\gamma_i = \{(\mathsf{last}_p(s[1, i]), \mathsf{last}_p(t[1, i])) \mid p \in \mathsf{pebbles}(s[1, i])\}$$

is a partial homomorphism from $\mathcal{A}$ to $\mathcal{B}$.

Example 5.4 (separating the pebble games). In the 2-pebble game on $\mathcal{A}$ and $\mathcal{B}$ below, Spoiler has a winning strategy. Let R and G denote the unary predicates 'red' and 'green'

---

[3]This game was discovered independently by the authors. The authors later discovered that a similar game appeared in a currently unavailable draft written by Iain A. Stewart.

respectively. The winning strategy of Spoiler is expressed by the $\exists^+\mathscr{L}^2$-formula

$$\exists x \exists y \Big(\exists y (Exy \wedge Ry) \wedge \exists y (Exy \wedge Gy) \wedge$$

$$\exists x (Eyx \wedge Rx) \wedge \exists x (Eyx \wedge Gx) \wedge Exy \Big),$$

which is true in $\mathcal{A}$ but not in $\mathcal{B}$. In contrast, Duplicator wins the all-in-one 2-pebble game on $\mathcal{A}$ and $\mathcal{B}$. This is straightforward, since the only difference between the two structures (that is expressible in $\exists^+\mathscr{L}^2$) is the property characterised by the formula above, which cannot be converted to a formula in $\exists^+ \curlywedge \mathscr{L}^2$.
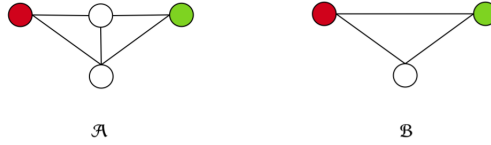

Figure 1: Separating example between the k-pebble game and the all-in-one k-pebble game for $k = 2$.

Besides its intuitive appeal and similarity to the original pebble game, as opposed to Dalmau's version, the all-in-one k-pebble game is a 'real' game in the sense that Duplicator has a choice and there is no default play that ensures a victory whenever one exists. Moreover, the elements of the comonad $\mathbb{PR}_k$ are plays in the all-in-one version of the game rather than the pebble-relation game, hence establishing this connection contributes to developing a more coherent relationship between the game and the logic. As we will see in the next subsection, this version of the game is also easier to generalise to the bijective case and work with in that setting.

In order to ensure that we capture k-variable logic with equality, we must impose an additional I-morphism condition on morphisms of type $\mathbb{PR}_k\mathcal{A} \to \mathcal{B}$. We can formulate the I-morphism condition by using the machinery of relative comonads. In order to do so, we expand the signature from σ to $\sigma^+ = \sigma \cup \{I\}$ which includes a new binary relation I, and we consider the functor $J : \mathfrak{R}(\sigma) \to \mathfrak{R}(\sigma^+)$, where a σ-structure $\mathcal{A}$ is mapped to the same σ-structure with $I^{J\mathcal{A}}$ interpreted as the identity relation, i.e. $I^{J\mathcal{A}} = \{(a, a) \mid a \in \mathcal{A}\}$. This leads us to the following theorem:

Theorem 5.5. *The following are equivalent for all σ-structures $\mathcal{A}$ and $\mathcal{B}$:*

(1) *Duplicator has a winning strategy in $\exists^+\mathbf{PPeb}_k(\mathcal{A}, \mathcal{B})$.*
(2) *There exists a coKleisli morphism $f : \mathbb{PR}_k\mathcal{A} \to \mathcal{B}$.*
(3) *There exists a coKleisli morphism $f^+ : \mathbb{PR}_k J\mathcal{A} \to J\mathcal{B}$.*
(4) *Duplicator has a winning strategy in $\exists^+\mathbf{PebR}_k(\mathcal{A}, \mathcal{B})$.*

Proof. $(1) \Rightarrow (2)$ We first define the coextension σ-morphism $f^* : \mathbb{PR}_k\mathcal{A} \to \mathbb{PR}_k\mathcal{B}$ and note that $f = \varepsilon_{\mathcal{B}} \circ f^*$. Suppose Duplicator has a winning strategy in $\exists^+\mathbf{PPeb}_k(\mathcal{A}, \mathcal{B})$. Then for every sequence of pebble placements $s = [(p_1, a_1), \ldots, (p_n, a_n)]$ on $\mathcal{A}$, Duplicator responds with a sequence of pebble placements $t = [(p_1, b_1), \ldots, (p_n, b_n)]$ with the same length and corresponding pebble placement at each index. We define $f^*(s, i) = (t, i)$ for every $i \in [\mathsf{n}]$.

To verify that f^* is indeed a homomorphism, suppose that $R^{\mathbb{PR}_k\mathcal{A}}((s, i_1), \ldots, (s, i_m))$ and let $i = \max\{i_1, \ldots, i_m\}$. By the active pebble condition, for all $i_j \in \{i_1, \ldots, i_m\}$, the pebble appearing at i_j does not appear in $s(i_j, i]$. Hence, the element $\varepsilon_{\mathcal{A}}(s, i_j)$ was the last

element pebbled by the pebble $\pi_A(s, i_j)$ in $s[1, i]$, and so $\varepsilon_{\mathcal{A}}(s, i_j)$ is in the domain of γ_i. By γ_i being a partial homomorphism, $R^{\mathcal{B}}(b_{i_1}, \ldots, b_{i_m})$, thus $R^{\mathcal{B}}(\varepsilon_{\mathcal{B}}(t, i_1), \ldots, \varepsilon_{\mathcal{B}}(t, i_m))$. The active pebble condition follows from t being compatible with s, so $R^{\mathbb{PR}_k \mathcal{B}}((t, i_1), \ldots, (t, i_m))$.

$(2) \Rightarrow (3)$ We say that a sequence $s = [(p_1, a_1), \ldots, (p_n, a_n)]$ is *duplicating* if there exists $i < j \in [\mathsf{n}]$ such that $\pi_A(s, i)$ does not appear in $s(i, j]$ (i.e. p_i is active at index j) with $a_i = a_j$. Let s' denote the longest subsequence of s such that s' is non-duplicating and for every $j \in [\mathsf{n}]$ there exists $j' \leq j$ with $\varepsilon_{\mathcal{A}}(s', j') = \varepsilon_{\mathcal{A}}(s, j)$. Such an s' can be shown to always exist for every sequence of pebble placements s. This can be shown inductively by removing moves (p_j, a_j) that are duplicating the placement of a different active pebble $p_i \neq p_j$ on the same element $a_i = a_j$. We define $f^+(s, j) = f(s', j')$.

To verify that $f^+ : \mathbb{PR}_k J\mathcal{A} \to J\mathcal{B}$ is a σ^+-morphism, suppose $((s, i), (t, j)) \in I^{\mathbb{PR}_k J\mathcal{A}}$. By the equality condition, $s = t$. If s is not duplicating, then $i = j$ and so $f^+(s, i) = f^+(t, j)$; therefore $(f^+(s, i), f^+(t, j)) \in I^{J\mathcal{B}}$ by $I^{J\mathcal{B}}$ being equality on $\mathcal{B}$. If s is duplicating, then by construction there exists some non-duplicating s' such that $f^+(s, i) = f^+(t, j) = f^+(s', j')$. As in the previous case, since $f^+(s, i) = f^+(t, j)$, we have that $(f^+(s, i), f^+(t, j)) \in I^{J\mathcal{B}}$. For the other relations $R \in \sigma \subseteq \sigma^+$, f^+ preserves R because f preserves R.

$(3) \Rightarrow (4)$ For each round n of the game with Spoiler choosing the domain I, we can construct a (possibly empty) sequence of pebbled elements s such that $I \subseteq \mathsf{Active}(s)$ inductively. Throughout the induction, we will keep track of the active pebbles $P = \mathsf{Active}(s)$ in s. For $n = 0$, $I = \varnothing$, so s is the empty sequence and $P = \varnothing$. For the inductive step, we assume for the inductive hypothesis, that s was constructed in round $n - 1$ from the move I and $P = \mathsf{Active}(s)$. There are two cases for Spoiler's move. If $I' \subseteq I$ is a shrinking move, then $s' = s$ and

$$P' = P \backslash \{p \mid \mathsf{last}_p(s) \in I \backslash I'\}.$$

If $I' = I \cup \{a_1, \ldots, a_j\}$ is growing move for $j \leq k$, then let

$$s' = s[(p_1, a_1), \ldots, (p_j, a_j)]$$

for $p_1, \ldots, p_j \notin P$. Clearly, $P' = \mathsf{Active}(s) = P \cup \{p_1, \ldots, p_j\}$.

We construct Duplicator's response in round n by considering the set

$$Z = \{t[1, m] \mid f^{\dagger}(s', m) = (t, m) \text{ and } s'[1, m] = s\},$$

where the map $f^{\dagger} : \mathbb{PR}_k J\mathcal{A} \to \mathbb{PR}_k J\mathcal{B}$ is the coextension of $f^+ : \mathbb{PR}_k J\mathcal{A} \to J\mathcal{B}$. For each $u = t[1, m] \in Z$, we can form the relation $\gamma_{s,u} = \{(\mathsf{last}_p(s), \mathsf{last}_p(u)) \mid p \in \mathsf{pebbles}(s)\}$ of pairs of active pebbled elements. Since $f^{\dagger}$ is a σ^+-morphism, each $\gamma_{s,u}$ is a partial homomorphism. Hence, Duplicator's response $T = \{\gamma_{s,u} \mid u \in Z\}$ is a winning move in $\exists^+ \mathbf{PebR}_k(\mathcal{A}, \mathcal{B})$.

$(4) \Rightarrow (1)$ Suppose Spoiler plays $s = [(p_1, a_1), \ldots, (p_n, a_n)]$ in the only round of the game $\exists^+ \mathbf{PPeb}_k(\mathcal{A}, \mathcal{B})$. We construct a list of Spoiler moves $(I_0^+, I_1^-, I_1^+, \ldots, I_n^-, I_n^+)$ in $\exists^+ \mathbf{PebR}_k(\mathcal{A}, \mathcal{B})$, where $I_0^+ = \varnothing$, I_i^- is a shrinking move removing the element previously pebbled with p_i from I_{i-1}^+ and I_i^+ is a blowing move adding a_i to I_i^-. Explicitly,

$$I_i^- = I_{i-1}^+ \backslash \{\mathsf{last}_{p_i}(s[1, i-1])\};$$
$$I_i^+ = I_i^- \cup \{a_i\}.$$

By assumption, Duplicator has a winning strategy in $\exists^+ \mathbf{PebR}_k(\mathcal{A}, \mathcal{B})$, so there exists a list of non-empty partial homomorphisms $(T_0^+, T_1^-, \ldots, T_n^-, T_n^+)$. Let $\gamma_1, \ldots, \gamma_n$ be a list of partial homomorphisms such that $\gamma_i \in T_i^+$ and $\gamma_{i+1}|_{I_{i+1}^-} = \gamma_i|_{I_{i+1}^-}$. Duplicator plays

$t = [(p_1, b_1), \ldots, (p_n, b_n)]$, where $b_i = \gamma_i(a_i)$ in $\exists^+\mathbf{PPeb}_k(\mathcal{A}, \mathcal{B})$. This is a winning move for Duplicator since the relation $\gamma_{s[1,i],t[1,i]} = \gamma_i$ by construction. $\qquad\square$

We combine Theorem 5.2 with Theorem 5.5 to derive the following result:

Corollary 5.6. *For all σ-structures $\mathcal{A}$ and $\mathcal{B}$,*

$$\mathcal{A} \Rrightarrow^{\exists^+ \curlywedge \mathscr{L}^k} \mathcal{B} \Leftrightarrow \mathbb{PR}_k \mathcal{A} \to \mathcal{B}.$$

5.2. **Isomorphism power theorem.** For the isomorphism power theorem, we first introduce the all-in-one bijective k-pebble game $\#\mathbf{PPeb}_k(\mathcal{A}, \mathcal{B})$ on $\mathcal{A}$ and $\mathcal{B}$, which is a revised bijective analogue of the all-in-one k-pebble game.

Definition 5.7 (all-in-one bijective k-pebble game). Suppose $\mathcal{A}, \mathcal{B}$ are σ-structures, we define the all-in-one bijective k-pebble game $\#\mathbf{PPeb}_k(\mathcal{A}, \mathcal{B})$. During the first and only round,

- Spoiler chooses a sequence of pebbles:

$$\vec{p} = [p_1, \ldots, p_n].$$

- Duplicator chooses a bijection $h_{\vec{p}} \colon A^n \to B^n$.
- Spoiler chooses a sequence of pebble placements respecting $\vec{p}$:

$$s = [(p_1, a_1), \ldots, (p_n, a_n)]$$

- Duplicator must respond with the list of pebble placements:

$$t = [(p_1, b_1), \ldots, (p_n, b_n)].$$

where $h_{\vec{p}}([a_1, \ldots, a_n]) = [b_1, \ldots, b_n]$.
Duplicator wins the round if for all $i \in [\mathsf{n}]$, the relation

$$\gamma_i = \{(a^p, b^p) \mid \forall p \in \{p_1, \ldots, p_n\}, a^p = \mathsf{last}_p(s[1, i]), b^p = \mathsf{last}_p(t[1, i])\}$$

is a partial isomorphism.

As previously mentioned, we work with the bijective variant $\#\mathbf{PPeb}_k(\mathcal{A}, \mathcal{B})$ of the all-in-one k-pebble game $\exists^+\mathbf{PPeb}_k(\mathcal{A}, \mathcal{B})$, since it is much easier to work with and generalise, in our context, than the k-pebble relation game.

The proof of this theorem follows similar lines to the analogous proof for the bijective Ehrenfeucht-Fraïssé game given by Libkin [Lib04].

We define the logic $\#\curlywedge\mathscr{L}^k$ where the equivalence relation $\equiv_{\#\curlywedge\mathscr{L}^k}$ is characterised by Duplicator winning strategies in $\#\mathbf{PPeb}_k(\mathcal{A}, \mathcal{B})$ and thus, isomorphism in $\mathcal{K}(\mathbb{PR}_k)$.

In order to define our logic $\#\curlywedge\mathscr{L}^k$, we first define a logic $\exists\curlywedge\mathscr{L}^k(Z)$. This logic is the same as $\exists^+\curlywedge\mathscr{L}^k$, but in addition to the k-many ordinary variables $X = \{x_1, \ldots, x_k\}$, we have an another sort of infinitely-many variables $Z = \{z_1, z_2, \ldots\}$ which we call 'walk pointer variables'. We also allow negations on atomic formulas in $\exists\curlywedge\mathscr{L}^k(Z)$ and require that every existential quantifier $\exists x_j$ is guarded with an equality $x_j = z_i$ equating a ordinary variable x_j with a walk pointer variable z_i. Explicitly, the formulas of $\exists\curlywedge\mathscr{L}^k(Z)$ can be defined recursively as follows:

$$\psi(\vec{x}, \vec{z}) ::= R(x_{i_1}, \ldots, x_{i_r}) \mid \neg p \mid \bigwedge \Psi \mid \bigvee \Psi \mid \exists x_j (x_j = z_i \wedge \psi(\vec{x'}, \vec{z'})),$$

where $R \in \sigma$ is an r-ary relation such that $\{x_{i_1}, \ldots, x_{i_r}\} \in [\vec{x}] \subseteq \{x_1, \ldots, x_k\}$, p is an atomic formula, $\bigwedge \Psi$ is a conjunction of formulas $\psi(\vec{x}, \vec{z})$ satisfying (R), $\bigvee \Psi$ is a disjunction of formulas $\psi(\vec{x}, \vec{z})$, $x_j \in [\vec{x'}]$ and $z_i \notin \vec{z'}$.

We can now define sentences of $\#\curlywedge\mathscr{L}^k$ as disjunctions of formulas that involve counting tuples of z-variables. Explicitly, the sentence of $\#\curlywedge\mathscr{L}^k$ can be defined recursively as follows:

$$\phi ::= \bigvee \Phi \mid \exists^j (z_{i_1}, \ldots, z_{i_m})\psi(z_{i_1}, \ldots, z_{i_m}),$$

where Φ is a collection of sentences ϕ, $j \in \mathbb{N}$ and $\psi(z_{i_1}, \ldots, z_{i_m}) \in \exists\curlywedge\mathscr{L}^k(Z)$.

Observe that in the syntax of the $\exists^j \vec{z}$ quantifier every free variable in the bound formula $\psi(\vec{z}) \in \exists\curlywedge\mathscr{L}^k(Z)$ is from the sort Z of walk pointer variables. We define the semantics of the operation above as

- $\mathcal{A} \vDash \exists^j (z_{i_1}, \ldots, z_{i_m})\psi(z_{i_1}, \ldots, z_{i_m})$ iff there exist exactly j-many tuples $\vec{a} = (a_1, \ldots, a_m)$ such that $\mathcal{A}, \vec{a} \vDash \psi(z_{i_1}, \ldots, z_{i_m})$.

The quantifier $\exists^j (z_{i_1}, \ldots, z_{i_m})$ in the logic $\#\curlywedge\mathscr{L}^k$ is inspired by a similar k-walk quantifier used in the logic defined in [LPS19].

A formula of $\varphi(\vec{x}) \in \exists^+\curlywedge\mathscr{L}^k$ is primitive if it contains no disjunctive subformula and every restricted conjunction subformula $\bigvee \Psi$ does not contain a sentence. From definition of primitive formula and the condition (R), it follows that every conjunctive subformula $\bigvee \Psi$ of a primitive formula $\varphi(x)$ contains only literals and at most one quantified formula with a free variable. We similarly can define primitive formulas of $\exists\curlywedge\mathscr{L}^k(Z)$. Thus, we extend this definition to $\#\curlywedge\mathscr{L}^k$, and say a sentence of $\phi \in \#\curlywedge\mathscr{L}^k$ is *primitive* if it is of the form $\phi = \exists^j (z_{i_1}, \ldots, z_{i_m})\psi(z_{i_1}, \ldots, z_{i_m})$ if $\psi(\vec{z}) \in \exists\curlywedge\mathscr{L}^k(Z)$ is primitive.

Proposition 5.8. *Every sentence $\phi \in \#\curlywedge\mathscr{L}^k$ is a disjunction of primitive subformulas.*

Proof. Given a sentence ϕ, we can inductively apply the standard rewrite rule

$$\exists y(\bigvee_{i \in I} \varphi(y, \vec{x})) \mapsto \bigvee_{i \in I} \exists y \varphi(y, \vec{x})$$

along with commutativity of $\bigwedge, \bigvee$ to obtain an equivalent formula. From the sematics it is clear that this standard rewrite rule also works for the walk counting quantifers $\exists^j (z_{i_1}, \ldots, z_{i_m})$. $\square$

We now prove the equivalence between truth preservation in $\#\curlywedge\mathscr{L}^k$ and the game $\#\mathbf{PPeb}_k(\mathcal{A}, \mathcal{B})$.

Theorem 5.9. *The following are equivalent for every σ-structures $\mathcal{A}$ and $\mathcal{B}$:*
(1) *Duplicator has a winning strategy in $\#\mathbf{PPeb}_k(\mathcal{A}, \mathcal{B})$.*
(2) $\mathcal{A} \equiv^{\#\curlywedge\mathscr{L}^k} \mathcal{B}$.

Proof. $\Rightarrow$ Suppose Duplicator has a winning strategy in $\#\mathbf{PPeb}_k(\mathcal{A}, \mathcal{B})$. Consider a primitive sentence ϕ of $\#\curlywedge\mathscr{L}^k$ such that $\mathcal{A} \vDash \phi$. As ϕ is primitive, it is of the form $\exists^j (z_{i_1}, \ldots, z_{i_m})\psi(z_{i_1}, \ldots, z_{i_m})$ for some $j \in \mathbb{N}$ and primitive $\psi(z_{i_1}, \ldots, z_{i_m}) \in \exists\curlywedge\mathscr{L}^k(Z)$. From ψ, we construct Spoiler's sequence of pebbles $\vec{p} = [p_1, \ldots, p_n]$ in the game $\#\mathbf{PPeb}_k(\mathcal{A}, \mathcal{B})$ such that for all $l \in [m]$, $x_{p_l} = z_{i_l}$ is subformula of ψ guarding a existential quantifier. Let $h_{\vec{p}} \colon A^m \to B^m$ denote Duplicator response in her winning strategy for $\#\mathbf{PPeb}_k(\mathcal{A}, \mathcal{B})$. By $\mathcal{A} \vDash \phi$, there exist exactly j-many tuples $\vec{a} = (a_1, \ldots, a_m)$ such that $\mathcal{A}, \vec{a} \vDash \psi(z_{i_1}, \ldots, z_{i_m})$. Thus, for each of these tuples, we examine the play where Spoiler plays the sequence $[(p_1, a_1), \ldots, (p_m, a_m)]$. Duplicator responds with the sequence $[(p_1, b_1), \ldots, (p_m, b_m)]$ such

that $h_{\vec{p}}(\vec{a}) = \vec{b} = [b_1, \ldots, b_m]$. From the winning condition it follows that $\mathcal{B}, \vec{b} \vDash \psi(z_{i_1}, \ldots, z_{i_m})$. Since $h_{\vec{p}}$ is a bijection there are j-many tuples $\vec{b}$ such that $\mathcal{B}, \vec{b} \vDash \psi(z_{i_1}, \ldots, z_{i_m})$. Thus, $\mathcal{B} \vDash \phi$. A similar argument shows that if $\mathcal{B} \vDash \phi$, then $\mathcal{A} \vDash \phi$. By Proposition 5.8, every sentence in $\#\curlywedge\mathscr{L}^k$ is a disjunction of primitive sentences ϕ. Therefore, $\mathcal{A} \equiv_{\#\curlywedge\mathscr{L}^k} \mathcal{B}$.

$\Leftarrow$ Before we proceed with this proof, we will need to define some formulas that are associated with plays in $\#\mathbf{PPeb}_k(\mathcal{A}, \mathcal{B})$.

First, we define, for every sequence $s = [(p_1, a_1), \ldots, (p_n, a_n)]$ of pebble placements on A, the literal diagram of s as

$$\mathbf{diag}_{\mathcal{A}}(s) = \{R(\mathbf{x}(p_{i_1}), \ldots, \mathbf{x}(p_{i_r})) \mid \forall R \in \sigma, \forall z \in [\mathsf{r}], \mathsf{last}_{p_{i_z}}(s) = a_{i_z} \wedge (a_{i_1}, \ldots, a_{i_r}) \in R^{\mathcal{A}}\}$$

$$\cup \{\neg R(\mathbf{x}(p_{i_1}), \ldots, \mathbf{x}(p_{i_r})) \mid \forall R \in \sigma, \forall z \in [\mathsf{r}], \mathsf{last}_{p_{i_z}}(s) = a_{i_z} \wedge (a_{i_1}, \ldots, a_{i_r}) \notin R^{\mathcal{A}}\}$$

where $\mathbf{x}(p) = x_p$, i.e. $\mathbf{x}$ returns the ordinary x variable corresponding to the pebble $p \in [\mathsf{k}]$. Next, for a sequence of pebble placements $s = [(p_1, a_1), \ldots, (p_n, a_n)]$ of length n, we define the formula $\mathbf{tp}_s(z_1, \ldots, z_n) \in \exists\curlywedge\mathscr{L}^k(Z)$ with walk pointer free variables $\{z_1, \ldots, z_n\} \subseteq Z$ such that $\mathcal{A}, (a_1 \ldots a_n) \vDash \mathbf{tp}_s(z_1, \ldots, z_n)$. To define $\mathbf{tp}_s$, for a sequence of pebble placements $s = [(p_1, a_1), \ldots, (p_n, a_n)]$ of length n, let $m \in \{0, \ldots, n-1\}$ be the largest index such that $p_m = p_n$. By induction on j up to m, we define the formulas $\varphi_j(\vec{z}, \vec{x})$, where $\vec{z} = (z_1, \ldots, z_j)$ are walk-pointer variables, and $\vec{x} = (\mathbf{x}(p_1), \ldots, \mathbf{x}(p_{j-1}))$ are ordinary variables corresponding to the pebbles $\{p_1, \ldots, p_{j-1}\}$ in $\mathsf{pebbles}(s[1,j])\backslash\{p_j\}$:

$$\varphi_0 = \top;$$

$$\varphi_j(z_j\vec{z}, \vec{x}) = \exists\mathbf{x}(p_j)(z_j = \mathbf{x}(p_j) \wedge \bigwedge \mathbf{diag}_{\mathcal{A}}(s[1,j]) \wedge \varphi_{j-1}(\vec{z}, \vec{x'}));$$

where $\mathbf{x}(p_j)$ is possibly in $\vec{x'}$. Finally, we define

$$\mathbf{tp}_s(z_1, \ldots, z_n) := \exists\vec{x} \left(\bigwedge_{j \in \{m+1, \ldots, n\}} z_j = \mathbf{x}(p_j) \wedge \bigwedge \mathbf{diag}_{\mathcal{A}}(s) \wedge \varphi_m(\vec{z'}, \vec{x'}) \right),$$

where $\vec{z'} = (z_1, \ldots, z_m)$, $\vec{x} = (\mathbf{x}(p_{m+1}), \ldots, \mathbf{x}(p_n))$ and $\exists\vec{x} := \exists\mathbf{x}(p_{m+1}) \ldots \exists\mathbf{x}(p_n)$.

We now proceed with the proof using these type formulas $\mathbf{tp}_s$. Suppose by hypothesis $\mathcal{A} \equiv_{\#\curlywedge\mathscr{L}^k} \mathcal{B}$. In particular, for every sentence of the form $\phi = \exists^j(z_1, \ldots, z_n)\mathbf{tp}_s(z_1, \ldots, z_n)$, $\mathcal{A} \vDash \phi \Leftrightarrow \mathcal{B} \vDash \phi$. This determines a bijection $h\colon A^n \to B^n$ such that:

$$\mathcal{A}, (a_1, \ldots, a_n) \vDash \mathbf{tp}_s(z_1, \ldots, z_n) \Leftrightarrow \mathcal{B}, (h(b_1), \ldots, h(b_n)) \vDash \mathbf{tp}_s(z_1, \ldots, z_n)$$

By construction of $\mathbf{tp}_s(z_1, \ldots, z_n)$, this yields a Duplicator winning strategy in $\#\mathbf{PPeb}_k(\mathcal{A}, \mathcal{B})$. Namely, for all indices $i \in [\mathsf{n}]$, either $i \leq m$ and $\mathbf{diag}_{\mathcal{A}}(s[1,i])$ is a subformula of $\mathbf{tp}_s(z_1, \ldots, z_n)$, or $i > m$ and $\mathbf{diag}_{\mathcal{A}}(s[1,i]) \subseteq \mathbf{diag}_{\mathcal{A}}(s)$. Therefore, $\mathcal{A}, (a_j, \ldots, a_i) \vDash \bigwedge \mathbf{diag}_{\mathcal{A}}(s[1,i])$ if and only if $\mathcal{B}, (h(a_j), \ldots, h(a_n)) \vDash \bigwedge \mathbf{diag}_{\mathcal{A}}(s[1,i])$, where $\mathsf{Active}(s[1,i]) = \{a_j, \ldots, a_i\}$. Thus, since $\bigwedge \mathbf{diag}_{\mathcal{A}}(s[1,i])$ expresses the satisfaction of all literals in $\{a_j, \ldots, a_i\} \subseteq A$ and $\{h(a_j), \ldots, h(a_i)\} \subseteq B$, γ_i is a partial isomorphism from $\mathcal{A}$ to $\mathcal{B}$. $\square$

We prove the correspondence Kleisli isomorphism in $\mathcal{K}(\mathbb{PR}_k)$ and $\#\mathbf{PPeb}_k(\mathcal{A}, \mathcal{B})$. Given a pebble sequence of some finite length $[p_1, \ldots, p_n] \in [\mathsf{k}]^{\leq\omega}$ and σ-structure $\mathcal{A}$, let $P_{\vec{p}}(\mathcal{A})$ be the induced substructure of $\mathbb{PR}_k(\mathcal{A})$ on the set

$$P_{\vec{p}}(A) = \{([(p_1, a_1), \ldots, (p_n, a_n)], i) \mid a_1, \ldots, a_n \in A \text{ and } i \in [\mathsf{n}]\}$$

Since the coextension $f^*\colon \mathbb{PR}_k\mathcal{A} \to \mathbb{PR}_k\mathcal{B}$ of a Kleisli morphism $f^*\colon \mathbb{PR}_k\mathcal{A} \to \mathcal{B}$ sends sequences of pebble placements on $\mathcal{A}$ paired with an index to compatible lists of pebble

placements on $\mathcal{B}$ paired with the same index, for every $\vec{p} \in [\mathsf{k}]^{\leq \omega}$, the restriction of f^* induces a σ-morphism $f_{\vec{p}} \colon P_{\vec{p}}(\mathcal{A}) \to P_{\vec{p}}(\mathcal{B})$. We also note that if $f \colon \mathbb{PR}_k \mathcal{A} \to \mathcal{B}$ is an isomorphism, then $f_{\vec{p}} \colon P_{\vec{p}}(\mathcal{A}) \to P_{\vec{p}}(\mathcal{B})$ is an isomorphism.

Theorem 5.10. *The following are equivalent:*

(1) *There exists a Kleisli isomorphism* $\mathbb{PR}_k J\mathcal{A} \to J\mathcal{B}$
(2) *Duplicator has a winning strategy in* $\#\mathbf{PPeb}_k(\mathcal{A}, \mathcal{B})$

Proof. $\Rightarrow$ Assume there exists a Kleisli isomorphism $f \colon \mathbb{PR}_k J\mathcal{A} \to \mathcal{B}$. We need construct a winning strategy for Duplicator in $\#\mathbf{PPeb}_k(\mathcal{A}, \mathcal{B})$, so suppose Spoiler plays the sequence $\vec{p} = [p_1, \ldots, p_n]$. The coextension f^* of the Kleisli isomorphism f when restricted to $P_{\vec{p}}(\mathcal{A})$ is the isomorphism $f_{\vec{p}} \colon P_{\vec{p}}(\mathcal{A}) \to P_{\vec{p}}(\mathcal{B})$. For every $[a_1, \ldots, a_n]$, there exists an $[b_1, \ldots, b_n]$, such that for all $i \in [\mathsf{n}]$, $f_{\vec{p}}(([(p_1, a_1), \ldots, (p_n, a_n)], i)) = ([(p_1, b_1), \ldots, (p_n, b_n)], i)$. We define Duplicator's response to be the function $\psi \colon A^n \to B^n$ which maps $[a_1, \ldots, a_n]$ to $[b_1, \ldots, b_n]$.

It follows from $f_{\vec{p}}$ being an isomorphism that ψ is a bijection. Namely. consider $g_{\vec{p}} \colon P_{\vec{p}}(\mathcal{B}) \to P_{\vec{p}}(\mathcal{A})$ to generated from the inverse $g^* \colon \mathbb{PR}_k J\mathcal{B} \to \mathbb{PR}_k J\mathcal{A}$ of f^*. There exists a $\phi \colon B^n \to A^n$ such that $\phi[b_1, \ldots, b_n] = [a_1, \ldots, a_n]$ if and only for all $i \in [\mathsf{n}]$ $g_{\vec{p}}([(p_1, b_1), \ldots, (p_n, b_n)], i) = ([(p_1, a_1), \ldots, (p_n, a_n)], i)$. Clearly, by construction $\psi^{-1} = \phi$.

Moreover, if Spoiler plays the sequence $s = [(p_1, a_1), \ldots, (p_n, a_n)]$, then Duplicator must respond with $t = [(p_1, b_1), \ldots, (p_n, b_n)]$ as $\psi[a_1, \ldots, a_n] = [b_1, \ldots, b_n]$. These sequences determine relation γ_i for all $i \in [\mathsf{n}]$. By $f_{\vec{p}}$ being an isomorphism, we can show that γ_i is a partial isomorphism. Suppose $R \in \sigma$ and consider the pairs $(a^{q_1}, b^{q_1}), \ldots, (a^{q_r}, b^{q_r}) \in \gamma_i$ such that $R^{\mathcal{A}}(a^{q_1}, \ldots, a^{q_r})$. By definition of γ_i, each $q_j = \pi_{\mathcal{A}}(s, z_j)$ and $a^{q_j} = \varepsilon_{\mathcal{A}}(s, j)$ for some $z_j \leq i \leq n$. Thus, by $R^{\mathcal{A}}(a^{q_1}, \ldots, a^{q_r})$, we have that $R^{\mathcal{A}}(\varepsilon(s, z_1), \ldots, \varepsilon(s, z_r))$ and $R^{\mathbb{PR}_k(\mathcal{A})}((s, z_1), \ldots, (s, z_r))$. By $f_{\vec{p}}$ being a morphism, $R^{\mathbb{PR}_k(\mathcal{B})}(f_{\vec{p}}(s, z_1), \ldots, f_{\vec{p}}(s, z_r))$ and $R^{\mathbb{PR}_k(\mathcal{B})}((t, z_1), \ldots, (t, z_r))$. Applying $\varepsilon_{\mathcal{B}}$, we obtain that $R^{\mathbb{PR}_k(\mathcal{B})}((t, z_1), \ldots, (t, z_r))$ and $R^{\mathcal{B}}(b^{q_1}, \ldots, b^{q_r})$. A similar proof, using the inverse of $f_{\vec{p}}$, show that if $R^{\mathcal{B}}(b^{q_1}, \ldots, b^{q_r})$, then $R^{\mathcal{A}}(a^{q_1}, \ldots, a^{q_r})$. Thus, γ_i is partial correspondence and using the fact that $J\mathcal{A}$ and $J\mathcal{B}$ interpret I as equality, we obtain that γ_i is a partial isomorphism for all $i \in [\mathsf{n}]$.

$\Leftarrow$ Conversely, suppose Duplicator has a winning strategy in $\#\mathbf{PPeb}_k$. Suppose Spoiler plays the pebble sequence $\vec{p} = [p_1, \ldots, p_n]$ and Duplicator responds with $\psi_{\vec{p}} \colon A^n \to B^n$. We define the σ-isomorphism $f_{\vec{p}} \colon P_{\vec{p}}(\mathcal{A}) \to P_{\vec{p}}(\mathcal{B})$ such that if $\psi_{\vec{p}}[a_1, \ldots, a_n] = [b_1, \ldots, b_n]$, then $f_{\vec{p}}([(p_1, a_1), \ldots, (p_n, a_n)], i) = ([(p_1, b_1), \ldots, (p_n, b_n)], i)$ for all $i \in [\mathsf{n}]$. Taking the coproduct of $f_{\vec{p}} \colon P_{\vec{p}}(\mathcal{A}) \to P_{\vec{p}}(\mathcal{B})$ over all $\vec{p} \in [\mathsf{k}]^{<\omega}$ we obtain the isomorphism $f^* \colon \mathbb{PR}_k J\mathcal{A} \to \mathbb{PR}_k J\mathcal{B}$. Composing with $\varepsilon_{J\mathcal{B}}$ yields the existence of the Kleisli isomorphism $f \colon \mathbb{PR}_k J\mathcal{A} \to \mathcal{B}$. $\quad\square$

We combine the above theorems to derive the following result:

Corollary 5.11. *For all finite σ-structures $\mathcal{A}$ and $\mathcal{B}$,*

$$\mathcal{A} \equiv^{\#\lambda\mathscr{L}^k} \mathcal{B} \Leftrightarrow J\mathcal{A} \cong^{\mathcal{K}(\mathbb{PR}_k)} J\mathcal{B}.$$

If we restrict the proofs of Theorems 5.9 and 5.10 to only sequences s of size up to n, we obtain a revised statement of Corollary 5.11 expressing that equivalence in the quantifier rank $\leq n$ fragment of $\#\lambda\mathscr{L}^k$ is characterised by a coKleisli isomorphism of $\mathbb{PR}_{k,n}$.

One of the contributions of the pebbling comonad was unifying two algorithms: the k-Weisfeiler-Lehman algorithm deciding the existence of a coKleisli isomorphism $J\mathcal{A} \cong^{\mathcal{K}(\mathbb{P}_k)} J\mathcal{B}$ and the k-consistency algorithm deciding the existence of a coKleisli morphism $\mathbb{P}_k \mathcal{A} \to \mathcal{B}$, where $\mathcal{A}$ and $\mathcal{B}$ represent the variables and values (respectively) of a given CSP [ADW17]. Corollary 5.11 suggests that understanding algorithms for deciding $\mathbb{PR}_k \mathcal{A} \to \mathcal{B}$ may provide

insights into how to construct algorithms for deciding $J\mathcal{A} \cong^{\mathcal{K}(\mathbb{PR}_k)} J\mathcal{B}$ (and vice-versa). We briefly discuss the complexity of such algorithms in the conclusion.

5.3. **Lovász-type theorem for pathwidth.** In one of his seminal papers, Lovász [Lov67] proved that two finite σ-structures $\mathcal{A}$ and $\mathcal{B}$ are isomorphic if and only if for all finite σ-structures $\mathcal{C}$, the number of homomorphisms from $\mathcal{C}$ to $\mathcal{A}$ is the same as the number of homomorphisms from $\mathcal{C}$ to $\mathcal{B}$. Following this result, weaker notions of equivalence have been found to have an analogous characterisation [Dvo09, Gro20, MR20]. These equivalences involve restricting the class of test structures $\mathcal{C}$. For instance, Dvořák [Dvo09] showed that counting homomorphisms from finite σ-structures of treewidth $< k$ yields equivalence in k-variable logic with counting quantifiers.

Counting homomorphisms from graphs of treewidth $< k$ also has a linear algebraic characterisation in terms of existence of a non-negative real solution to a certain system of linear equations [DGR18]. It was shown in a later work by Grohe, Rattan and Seppelt [GRS22] that removing this non-negativity constraint yields a homomorphism counting result from graphs of pathwidth $< k$. In their work they used techniques from representation theory. By contrast, we use techniques from category theory to obtain an analogue of Dvořák's result providing a logical characterisation of homomorphism counting from σ-structures with pathwidth $< k$.

In recent work, Dawar, Jakl and Reggio [DJR21] introduced a generalisation of Lovász-type theorems for comonads over $\mathfrak{R}(\sigma)$. In this section we utilise this generalisation in order to show that counting homomorphisms from finite σ-structures of pathwidth $< k$ yields equivalence in restricted conjunction k-variable logic with counting quantifiers.

We begin by introducing the categorical generalisation for Lovász-type theorems [DJR21].

Let $\mathbb{C}_k$ and $\mathbb{C}_k^+$ be comonads on $\mathfrak{R}(\sigma)$ and $\mathfrak{R}(\sigma^+)$, respectively. These comonads yield the forgetful-cofree adjunctions $U^{\mathbb{C}_k} \dashv F^{\mathbb{C}_k}$ and $U^{\mathbb{C}_k^+} \dashv F^{\mathbb{C}_k^+}$ where $U^{\mathbb{C}_k} : \mathcal{EM}(\mathbb{C}_k) \to \mathfrak{R}(\sigma)$ and $F^{\mathbb{C}_k} : \mathfrak{R}(\sigma^+) \to \mathcal{EM}(\mathbb{C}_k)$ (similarly for $\mathbb{C}_k^+$). Let $\mathbb{C}_k^+ J$ be the relative comonad on $J : \mathfrak{R}(\sigma) \to \mathfrak{R}(\sigma^+)$. The functor J has a left adjoint $H : \mathfrak{R}(\sigma^+) \to \mathfrak{R}(\sigma)$ sending $\mathcal{A} \in \mathfrak{R}(\sigma^+)$ to the quotient structure $\mathcal{A}^-/\sim \in \mathfrak{R}(\sigma)$, where $\mathcal{A}^-$ is the σ-reduct of $\mathcal{A}$ and $\sim$ is the least equivalence relation generated from $I^{\mathcal{A}}$. Below, recall that $\mathfrak{R}_f(\sigma)$ denotes the full subcategory of σ-structures with finite universes.

Theorem 5.12 [DJR21]. *Given a logic $\mathscr{L}_k$, suppose that the following three conditions are satisfied:*

(1) *For all finite σ-structures $\mathcal{A}$ and $\mathcal{B}$*

$$\mathcal{A} \equiv_{\mathscr{L}_k} \mathcal{B} \iff F^{\mathbb{C}_k^+}(\mathcal{A}) \cong F^{\mathbb{C}_k^+}(\mathcal{B});$$

(2) *$\mathbb{C}_k^+ J$ sends finite structures in $\mathfrak{R}(\sigma)$ to finite structures in $\mathfrak{R}(\sigma^+)$;*

(3) *$J : \mathfrak{R}(\sigma) \to \mathfrak{R}(\sigma^+)$ and $H : \mathfrak{R}(\sigma^+) \to \mathfrak{R}(\sigma)$ restrict to the subcategories of finite relational structures $\mathfrak{R}_f(\sigma) \cap \mathrm{im}(U^{\mathbb{C}_k})$ and $\mathfrak{R}_f(\sigma^+) \cap \mathrm{im}(U^{\mathbb{C}_k^+})$.*

Then for every finite σ-structures $\mathcal{A}$ and $\mathcal{B}$,

$$\mathcal{A} \equiv_{\mathscr{L}_k} \mathcal{B} \iff |\mathfrak{R}_f(\sigma)(\mathcal{C}, \mathcal{A})| = |\mathfrak{R}_f(\sigma)(\mathcal{C}, \mathcal{B})|,$$

for all finite σ-structure $\mathcal{C}$ in $\mathrm{im}(U^{\mathbb{C}_k})$.

The categorical interpretation of pathwidth and the isomorphism power theorem for $\mathbb{PR}_k$ allow us to derive the following consequence from Theorem 5.12:

Corollary 5.13 (Lovász-type theorem for pathwidth). *For all finite σ-structures $\mathcal{A}$ and $\mathcal{B}$,*

$$\mathcal{A} \equiv^{\#\wedge\mathscr{L}^k} \mathcal{B} \iff |\mathfrak{R}_f(\sigma)(\mathcal{C}, \mathcal{A})| = |\mathfrak{R}_f(\sigma)(\mathcal{C}, \mathcal{B})|,$$

for every finite σ-structure $\mathcal{C}$ with pathwidth $< k$.

Intuitively, since $\mathbb{PR}_k\mathcal{A}$ is infinite for every structure $\mathcal{A}$, in order to meet condition 2 we employ a similar technique to the one Dawar et al. used when applying Theorem 5.12 to $\mathbb{P}_k$. We consider the pebble-relation comonad $\mathbb{PR}_{k,n}$ that is also graded by sequence length. Recall that $\mathbb{PR}_{k,n}\mathcal{A}$ consists of pairs (s, i), where s is a sequence of length at most n. Each of the three conditions follows from the results in [DJR21] applied to $\mathbb{PR}_{k,n}$.

Condition 1 follows from Corollary 5.11 when applied to the comonad $\mathbb{PR}_{k,n}$, where $F^{\mathbb{C}_k^+}$ is instantiated as the cofree functor $F_{k,n} : \mathfrak{R}(\sigma^+) \to \mathcal{EM}(\mathbb{PR}_{k,n})$ from the adjunction $U_{k,n} \dashv F_{k,n}$.

Condition 2 follows from the fact that $\mathbb{PR}_{k,n}\mathcal{A}$ is finite whenever $\mathcal{A}$ is finite.

Condition 3 requires that if $\mathcal{A}$ has a k-pebble linear forest cover, then $J\mathcal{A}$ and $HJ\mathcal{A}$ have k-pebble linear forest covers. The paper by Dawar et al. proves this statement for general k-pebble forest covers (Lemma 22 of [DJR21]) as part of their application of Theorem 5.12 to $\mathbb{P}_k$. In particular, their argument holds for paths in the category of k-pebble forest covers. Therefore, the statement holds for coproducts of paths, i.e. k-pebble linear forest covers.

Finally, realising $\mathbb{PR}_k$ as the colimit of $\mathbb{PR}_{k,n}$ for all $n \in \omega$ allows us to obtain the desired result.

6. Conclusion

The results we exhibited extend the growing list of Spoiler-Duplicator game comonads that unify particular model-comparison games with combinatorial invariants of relational structures. This approach has applications in reformulations of Rossman's homomorphism preservation theorem [Pai20], and it provides a new perspective on finite model theory and descriptive complexity. In particular, $\mathbb{PR}_k$ provides a categorical definition for pathwidth as well as winning strategies in Dalmau's pebble-relation game and the all-in-one pebble game. This allowed us to obtain a syntax-free characterisation of equivalence in the restricted conjunction fragment of k-variable logic.

The comonad $\mathbb{PR}_k$ was obtained in a unique way by first trying to capture the combinatorial invariant, instead of internalising the corresponding game, as was done for previously studied game comonads [ADW17, AS18, CD21, AM21]. Moreover, some of the proofs in this paper use techniques distinct from those used for other Spoiler-Duplicator comonads (e.g. k-pebbling section families, fibres over $\mathbb{P}_k\mathcal{A}$ to model non-determinism). The coKleisli isomorphism of $\mathbb{PR}_k$ allowed us to define an original bijective game that characterises equivalence in the heretofore unexplored restricted conjunction fragment of k-variable logic with counting quantifiers. Finally, we proved a Lovász-type theorem linking equivalence in this logic with homomorphism counting from structures of pathwidth $< k$.

We share a few possible avenues for future work:

- Dalmau [Dal05] claims that constraint satisfaction problems (CSPs) with bounded pathwidth duality are in **NL**. All known CSPs in **NL** have bounded pathwidth duality. One possible inquiry would be to check if redefining bounded pathwidth duality in terms of $\mathbb{PR}_k$ can help in finding a proof for the converse or a construction of a counterexample.

Egri [Egr14] proved that CSPs that have bounded symmetric pathwidth duality are in **L**. Whereas bounded pathwidth duality can be seen as a local property of the obstruction set, bounded symmetric pathwidth duality is a global property of the obstruction set. By reformulating bounded symmetric pathwidth duality in terms of $\mathbb{PR}_k$, we can further reinforce the comonadic analysis to understand the complexity of CSPs.

- For the pebbling comonad $\mathbb{P}_k$, we can interpret a coKleisli morphism $\mathbb{P}_k\mathcal{A} \to \mathcal{B}$ as a simulation between Kripke models constructed from the sets of tuples $A^{\leq k}$ and $B^{\leq k}$. Analogously, a coKleisli morphism $\mathbb{PR}_k\mathcal{A} \to \mathcal{B}$ can be interpreted as trace inclusion on these very same Kripke models. Intuitively, Duplicator's response in the all-in-one k-pebble game maps a trace of the Kripke model built from $A^{\leq k}$ to a trace of the Kripke model built from $B^{\leq k}$. While simulation of Kripke models is P-complete, trace inclusion[4] is $PSPACE$-complete [HS96], leading to the following conjecture:

Conjecture 6.1. Deciding $\mathbb{PR}_k\mathcal{A} \to \mathcal{B}$ for inputs $\mathcal{A}, \mathcal{B}$ and fixed k is $PSPACE$-complete.

- The pebble-relation comonad $\mathbb{PR}_k$ can be seen as the 'linear', or path variant, of the 'tree shaped' pebbling comonad $\mathbb{P}_k$. In fact, as the notion of arboreal categories discussed in [AR21] shows, all of the previously discovered Spoiler-Duplicator game comonads $\mathbb{C}_k$ are 'tree shaped', i.e. $\mathcal{EM}(\mathbb{C}_k)$ is arboreal. One may investigate whether there exists a general method for obtaining a 'linear' variant of a 'tree shaped' Spoiler-Duplicator game comonad $\mathbb{C}_k$. Moreover, the framework of arboreal categories demonstrates that a certain categorical generalisation of bisimulation in $\mathcal{EM}(\mathbb{C}_k)$ captures equivalence in $\mathscr{J}_k$ associated with $\mathbb{C}_k$. Discovering which fragment of $\mathscr{L}^k$ is captured by a similar relation in $\mathcal{EM}(\mathbb{PR}_k)$ remains an open question.

REFERENCES

[ACU14]   Thosten Altenkirch, James Chapman, and Tarmo Uustalu. Monads need not be endofunctors. *Logical Methods in Computer Science*, 11(1), 2014.

[ADW17]   Samson Abramsky, Anuj Dawar, and Pengming Wang. The pebbling comonad in finite model theory. In *Logic in Computer Science (LICS), 2017 32nd Annual ACM/IEEE Symposium on*, pages 1–12. IEEE, 2017.

[AM21]   Samson Abramsky and Dan Marsden. Comonadic semantics for guarded fragments. In *2021 36th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*. IEEE, June 2021.

[AR21]   Samson Abramsky and Luca Reggio. Arboreal Categories and Resources. In Nikhil Bansal, Emanuela Merelli, and James Worrell, editors, *48th International Colloquium on Automata, Languages, and Programming (ICALP 2021)*, volume 198 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 115:1–115:20, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[AS18]   Samson Abramsky and Nihil Shah. Relating structure and power: Comonadic semantics for computational resources. In *27th EACSL Annual Conference on Computer Science Logic, CSL 2018, September 4-7, 2018, Birmingham, UK*, pages 2:1–2:17, 2018.

[AS21]   Samson Abramsky and Nihil Shah. Relating structure and power: Comonadic semantics for computational resources. *Journal of Logic and Computation*, 31(6):1390–1428, August 2021.

[CD21]   Adam Ó Conghaile and Anuj Dawar. Game comonads & generalised quantifiers. In *29th EACSL Annual Conference on Computer Science Logic (CSL 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021.

[Dal05]   Victor Dalmau. Linear Datalog and Bounded Path Duality of Relational Structures. *Logical Methods in Computer Science*, 1(1):5, April 2005.

---

[4]This is originally called trace preorder in [HS96].

[DGR18]   Holger Dell, Martin Grohe, and Gaurav Rattan. Lovász meets weisfeiler and leman. In Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Donald Sannella, editors, *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, July 9-13, 2018, Prague, Czech Republic*, volume 107 of *LIPIcs*, pages 40:1–40:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.

[DJR21]   Anuj Dawar, Tomas Jakl, and Luca Reggio. Lovász-type theorems and game comonads. In *2021 36th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*. IEEE, June 2021.

[Dvo09]   Zdeněk Dvořák. On recognizing graphs by numbers of homomorphisms. *Journal of Graph Theory*, 64(4):330–342, November 2009.

[Egr14]   László Egri. On constraint satisfaction problems below p. *Journal of Logic and Computation*, 26(3):893–922, 2014.

[Gro17]   Martin Grohe. *Descriptive Complexity, Canonisation, and Definable Graph Structure Theory*. Cambridge University Press, 2017.

[Gro20]   Martin Grohe. Counting bounded tree depth homomorphisms. In *Proceedings of the 35th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*. ACM, July 2020.

[GRS22]   Martin Grohe, Gaurav Rattan, and Tim Seppelt. Homomorphism tensors and linear equations, 2022. To appear in *49th EATCS International Colloquium on Automata, Languages and Programming (ICALP 2022)*.

[Hel96]   Lauri Hella. Logical hierarchies in PTIME. *Information and Computation*, 121:1–19, 1996.

[HS96]   Hans Hüttel and Sandeep Shukla. On the complexity of deciding behavioural equivalences and preorders. Technical report, USA, 1996.

[Imm82]   Neil Immerman. Upper and lower bounds for first order expressibility. *Journal of Computer and System Sciences*, 28:76–98, 1982.

[KV90]   Phokion G Kolaitis and Moshe Y Vardi. 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*, pages 61–71. ACM, 1990.

[Lib04]   Leonid Libkin. *Elements of Finite Model Theory (Texts in Theoretical Computer Science. An EATCS Series)*. Springer, 2004.

[Lov67]   László Lovász. Operations with structures. *Acta Mathematica Academiae Scientiarum Hungaricae*, 18(3-4):321–328, September 1967.

[LPS19]   Moritz Lichter, Ilia Ponomarenko, and Pascal Schweitzer. Walk refinement, walk logic, and the iteration number of the weisfeiler-leman algorithm. In *2019 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–13, 2019. doi:10.1109/LICS.2019.8785694.

[Man76]   Ernest G. Manes. *Algebraic Theories*. Springer New York, 1976.

[MR20]   Laura Mančinska 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. IEEE, 2020.

[Orc14]   Dominic Orchard. *Programming contextual computations*. PhD thesis, 2014.

[Pai20]   Tom Paine. A pebbling comonad for finite rank and variable logic, and an application to the equirank-variable homomorphism preservation theorem. In *Proceedings of MFPS*, 2020.

[RS83]   Neil Robertson and P.D. Seymour. Graph minors. i. excluding a forest. *Journal of Combinatorial Theory, Series B*, 35(1):39–61, August 1983.

[She71]   Saharon Shelah. Every two elementarily equivalent models have isomorphic ultrapowers. *Israel Journal of Mathematics*, 10(2):224–233, 1971.

[WL68]   Boris Weisfeiler and Andrei Leman. The reduction of a graph to canonical form and the algebra which appears therein. *NTI, Series*, 2(9):12–16, 1968.