Restricting Power: The Pebble-Relation Comonad in Finite Model Theory

Nihil Shah
Department of Computer Science
University of Oxford
nihil.shah@cs.ox.ac.uk

The pebbling comonad, introduced by Abramsky, Dawar, and Wang, gave a categorical semantics to the k-pebble games used in finite model theory. They showed that the coKleisli category of the pebbling comonad can be used to give characterizations of equivalence under fragments and extensions of infinitary k-variable logic. Moreover, the coalgebras over this pebbling comonad correspond to tree decompositions and characterize treewidth of a structure. In this paper, we investigate the pebble-relation comonad whose coalgebras correspond to path decompositions and characterize pathwidth. We also demonstrate how coKleisli morphisms of the pebble-relation comonad give a categorical semantics to Duplicator's winning strategies in Dalmau's pebble-relation game. Consequently, the coKleisli morphisms characterize preservation under a restricted conjunction fragment of existential positive infinitary k-variable logic.

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 glasses imposed by definability in a logic $\mathscr{J}$. Namely, given two structures over the same signature $\mathscr{A}$ and $\mathscr{B}$, model theory is concerned with equivalence under the relation:

$$\mathscr{A} \equiv^{\mathscr{J}} \mathscr{B} := \forall \varphi \in \mathscr{J}, \mathscr{A} \vDash \varphi \Leftrightarrow \mathscr{B} \vDash \varphi$$

Historically, a theme in model theory has been to find syntax-free characterizations of these equivalences. This is exemplified by the Keisler-Shelah theorem [15] for first-order logic. These equivalences are also characterized 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 [9] characterizes equivalence in infinitary logic graded by the number of variables k. In [1], this game, and two similar variants, were given a semantics in terms of morphisms involving the pebbling comonad, $\mathbb{P}_k$. Since then, similar game comonads have been found for Ehrenfeucht-Fraïssé games, modal-bisimulation games [4], games for guarded logics [2], and games for finite-variable logics with generalised quantifiers [6].

Remarkably, in all of these cases, the coalgebras over the game comonad correspond to decompositions of structures, such as tree decompositions of width $< k$ (in the case of $\mathbb{P}_k$) and forest covers of height $\leq k$ (in the case of the Ehrenfeucht-Fraïssé comonad, $\mathbb{E}_k$). An immediate corollary of these correspondences are novel definitions for an associated graph parameter, such as treewidth (for $\mathbb{P}_k$) and tree-depth (for $\mathbb{E}_k$). In the present paper, we widen the domain of these game comonads by investigating a pebble-relation comonad $\mathbb{PR}_k$ where the coalgebras correspond to path decompositions of width $< k$; yielding a new definition of pathwidth.

Submitted to:
ACT 2021

In [1], Abramsky et. al. showed 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 in [11], was used to study expressivity in Datalog, preservation of existential positive formulas of k-variable logic $\exists^+\mathscr{L}_\infty$, and **PTIME**-tractable constraint satisfaction problems. Dalmau in [7] 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 $\exists^+\mathscr{N}_k$, and **NLOGSPACE**-tractable constraint satisfaction problems. Utilizing Dalmau's results and a general fact about comonads, we are able to show that coKleisli morphisms of our comonad $\mathbb{PR}_k$ correspond to winning strategies in this pebble-relation game. We then obtain this result directly. Curiously, there are features of the pebble-relation game, that ensure that this direct proof utilizes ideas distinct from the arguments used for the other game comonads explored in [1, 4, 6].

In section 2, we introduce our notation and necessary background. In section 3, we explain the emerging notion of a Spoiler-Duplicator game comonad and the associated family of results that have accompanied them. In section 4, we introduce our central contribution, the pebble relation comonad $\mathbb{PR}_k$, and demonstrate the relationship between $\mathbb{PR}_k$ and $\mathbb{P}_k$. In section 5, we prove that coalgebras over $\mathbb{PR}_k$ correspond to path decompositions of width $< k$, providing a coalgebraic characterization of pathwidth. In section 6, we introduce Dalmau's one-sided pebble-relation game, the novel variants of this game, and how winning strategies for Duplicator in these games are captured through morphisms involving $\mathbb{PR}_k$. In section 7, we conclude the paper with a summary of results and point out further directions of this comonadic view to descriptive complexity.

2 Background

In this section, we will establish some notational preliminaries and give a short introduction to the relevant concepts in category theory and finite model theory that we will need.

2.1 Set notations

Given a poset, i.e. a set with a partial order, $(X, \leq)$ and $x \in X$, we denote the down and up sets as $\downarrow x = \{y \in X \mid y \leq x\}$ and $\uparrow x = \{y \in X \mid y \geq x\}$. A poset $(X, \leq)$ 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 element is comparable. If $(T, \leq)$ is a poset, 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*. If $(T, \leq)$ is a forest and the 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*. In the presence of a forest order, we will also 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 will write $[\mathsf{n}]$ for the set $\{1, \ldots, n\}$. When convenient, we will consider $[\mathsf{n}]$ as having the usual order on segments of natural numbers $\leq$. Given a set A, we will 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 $A^{\leq k}$. We will denote a 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 sequence $s = [a_1, \ldots, a_n]$. Given two sequences $s, t \in A^*$, we will denote the concatenation of s followed by t by juxtaposition, i.e. 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^+. For $s = [a_1, \ldots, a_n]$ and $i, j \in [\mathsf{n}]$, let $s(i, j] = [a_{i+1}, \ldots, a_j]$ if $i < j$ or ε if $i \geq j$; and let $s[i, j] = [a_i, \ldots, a_j]$ if $i \leq j$ or ε if $i > j$.

2.2 Category Theory

We will assume the reader is familiar with standard the category-theoretic notions of category, functor, and natural transformation. Given a category $\mathfrak{C}$, the class of objects will be denoted as $\mathfrak{C}_0$ and the class of morphisms $\mathfrak{C}_1$. If $X, Y \in \mathfrak{C}_0$, then the class of morphisms from X to Y is denoted $\mathfrak{C}(X,Y)$. Some readers may be familiar with the notion of a comonad; an important class of endofunctor. We will need the more general notion of a relative comonad introduced in [5] which weakens this endofunctor requirement. 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.

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 *Kleisli category*, denoted $\mathscr{K}(\mathbb{T})$, where

- $\mathscr{K}(\mathbb{T})_0$ is the same as the class of objects $\mathfrak{J}_0$
- $\mathscr{K}(\mathbb{T})_1$ are morphisms of the type $f : \mathbb{T}X \to JY \in \mathfrak{C}_1$
- The composition $g \circ_{\mathscr{K}} f : \mathbb{T}X \to JZ$ of 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* [12], and the corresponding Klesli category, can be recovered when $\mathfrak{J} = \mathfrak{C}$ and $J = \mathsf{id}_{\mathfrak{C}}$. Given a 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*. The coKleisli form can be recovered by defining the coextension mapping $(\cdot)^*$ as $f^* = \mathbb{T}f \circ \delta$.

The *category of coalgebras* $\mathscr{EM}(\mathbb{T})$, or *Eilenberg-Moore category* associated with a comonad $(\mathbb{T} : \mathfrak{C} \to \mathfrak{C}, \varepsilon, \delta)$ is specified by:

- $\mathscr{EM}(\mathbb{T})_0$ are pairs $(A, \alpha : A \to \mathbb{T}A)$ with $A \in \mathfrak{C}_0$ and $\alpha \in \mathfrak{C}_1$ such that following diagrams commute in $\mathfrak{C}$:

$$\begin{array}{ccc} A & \xrightarrow{\alpha} & \mathbb{T}A \\ {\scriptstyle\alpha}\downarrow & & \downarrow{\scriptstyle\delta_A} \\ \mathbb{T}A & \xrightarrow{\mathbb{T}\alpha} & \mathbb{T}^2A \end{array} \qquad \begin{array}{ccc} A & \xrightarrow{\alpha} & \mathbb{T}A \\ & {\scriptstyle\mathsf{id}_A}\searrow & \downarrow{\scriptstyle\varepsilon_A} \\ & & A \end{array}$$

- $\mathscr{EM}(\mathbb{T})$ are morphisms $f : A \to B \in \mathfrak{C}_1$ from $(A, \alpha : A \to \mathbb{T}A)$ to $(B, \beta : B \to \mathbb{T}B)$ such that the following diagram commutes in $\mathfrak{C}$:

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

- The identity morphisms and composition are inherited from $\mathfrak{C}$

Monads and comonads are closely linked to adjunctions. Recall that an *adjunction*, denoted $L \dashv R$, is a pair of functors $L : \mathfrak{C} \to \mathfrak{D}$ and $R : \mathfrak{D} \to \mathfrak{C}$ together with natural transformations $\eta : \mathrm{id}_{\mathfrak{C}} \to RL$, and $\varepsilon : LR \to \mathrm{id}_{\mathfrak{D}}$ such that the maps $\theta_{A,B} : \mathfrak{C}(A, RB) \to \mathfrak{D}(LA, B)$ and $\theta'_{A,B} : \mathfrak{D}(LA, B) \to \mathfrak{C}(A, RB)$ given by are mutually inverse.

Conversely, every comonad $(\mathbb{T}, \varepsilon, \delta)$ arises from an adjunction, or *resolution* of $\mathbb{T}$. The class of all resolutions of a comonad $\mathbb{T}$ form a category. The inital object of this category is an adjunction associated with $\mathscr{K}(\mathbb{T})$. The terminal (final) object of this category is an adjunction associated with $\mathscr{EM}(\mathbb{T})$. Hence, for every

Observe that given a comonad $(\mathbb{T} : \mathfrak{C} \to \mathfrak{C}, \varepsilon, ()^*)$ and functor $J : \mathfrak{J} \to \mathfrak{C}$, the functor $\mathbb{T}^J = \mathbb{T} \circ J$ can be made into a relative comonad $(\mathbb{T}^J, \varepsilon', ()')$ on J where $\varepsilon'_X = \varepsilon_{JX}$ and coextension mapping $()'$ is defined for $f : \mathbb{T}^J X \to JY$ to be $f' = f^* : \mathbb{T}^J \to \mathbb{T}^J Y$.

2.3 Finite Model Theory

We will assume a fixed vocabulary σ of relational symbols R each with a positive integer arity $m = \rho(R)$. If $\rho(R) = m$, we will say R is an m-ary relation. A *σ-structure* $\mathscr{A}$ is specified by a universe of elements A, and interpretations $R^{\mathscr{A}} \subseteq A^m$ of each m-ary relation $R \in \sigma$. We will use calligraphic and boldface letters (e.g. $\mathscr{A}, \mathscr{B}, \mathscr{C}, \mathbf{P}, \mathbf{Q}, \mathbf{S}$ etc.) to denote σ-structures and the corresponding roman letters (e.g. A, B, C, P, Q, S etc.) to denote the underlying universe of elements.

Let $\mathscr{A}$ and $\mathscr{B}$ be σ-structures. If $B \subseteq A$ and $R^{\mathscr{B}} \subseteq R^{\mathscr{A}}$ for every relation $R \in \sigma$, then $\mathscr{B}$ is a *σ-substructure* of $\mathscr{A}$. If $B \subseteq A$, then we can form the *B induced σ-substructure*, $\mathscr{B} = \mathscr{A}|_B$ with universe B and interpretations $R^{\mathscr{B}} = R^{\mathscr{A}} \cap B^m$ for m-ary relation $R \in \sigma$. The graph $\mathscr{G}(\mathscr{A}) = (A, \frown)$ is the Gaifman graph of $\mathscr{A}$ where $a \frown a'$ iff $a = a'$ or a, a' appear in some tuple of $R^{\mathscr{A}}$ for some $R \in \sigma$.

A *σ-morphism from $\mathscr{A}$ to $\mathscr{B}$*, or *homomorphism*, denoted $h : \mathscr{A} \to \mathscr{B}$, is a set function $h : A \to B$ such that $R^{\mathscr{A}}(a_1, \ldots, a_m) \Rightarrow R^{\mathscr{B}}(h(a_1), \ldots, h(a_m))$ for every m-ary relation $R \in \sigma$. If $h : A \to B$ additionally satisfies $R^{\mathscr{B}}(h(a_1), \ldots, h(a_m)) \Rightarrow R^{\mathscr{A}}(a_1, \ldots, a_m)$ for every m-ary relation $R \in \sigma$, then $h : \mathscr{A} \to \mathscr{B}$ is called *strong*. We will denote the category of σ-structures with σ-morphisms as $\mathfrak{R}(\sigma)$. If there exists an $h : \mathscr{A} \to \mathscr{B}$, we write $\mathscr{A} \to \mathscr{B}$, and say that $\mathscr{A}$ *homomorphically maps to* $\mathscr{B}$. If $h : \mathscr{A} \to \mathscr{B}$ is a strong injective σ-morphism, then we write $h : \mathscr{A} \rightarrowtail \mathscr{B}$. If there exists an $h : \mathscr{A} \rightarrowtail \mathscr{B}$, we write $\mathscr{A} \rightarrowtail \mathscr{B}$, and say that $\mathscr{A}$ *embeds into* $\mathscr{B}$.

2.4 Logical Fragments

We are mainly concerned with fragments of infinitary logic. 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 $\mathbf{x} = (x_1, \ldots, x_n)$ as $\phi(\mathbf{x})$. The set of formulas in $\mathscr{L}_\infty$ with symbols in signature σ is $\mathscr{L}_\infty(\sigma)$. If $\mathscr{A}$ is a σ-structure $\mathbf{a} \in A^n$ and $\phi(\mathbf{x}) \in \mathscr{L}_\infty(\sigma)$ and $\mathscr{A}, \mathbf{a}$ satisfies $\phi(\mathbf{x})$, then we write $\mathscr{A}, \mathbf{a} \vDash \phi(\mathbf{x})$. The infinitary logic can be graded into k-variable fragments, denoted $\mathscr{L}_k$, where formulas can only contain at most k-many variables. We will also consider infinitary logic $\mathscr{L}_\infty$ graded by quantifier rank $\leq k$, denoted $\mathscr{L}^k$. Observe that for finite signatures σ (i.e. only finitely many relational symbols), $\mathscr{L}^k$ is equivalent to ordinary first-order logic graded by quantifier rank $\leq k$.

For every logic $\mathscr{J}$ we consider, we will also be interested in three variants. The first variant is the primitive positive fragment $\exists^p \mathscr{J}$ where we will only consider formulas built using existential quantifiers, conjunction, and atomic formulas. The second variant is the existential positive fragment $\exists^+ \mathscr{J}$ where we will only consider formulas built using existential quantifiers, disjunctions, conjunctions, and atomic formulas. The last variant is the extension of $\mathscr{J}$ with counting quantifiers $\#\mathscr{J}$. These are quantifiers of the form $\exists_{\leq m}$ and $\exists_{\geq m}$ where the semantics of $\mathscr{A} \vDash \exists_{\leq m} x \psi(x)$ is that there exist at most m distinct elements of A satisfying $\psi(x)$.

A *restricted conjunction (disjunction)* is a conjunction (disjunction) of the form $\bigwedge \Psi$ ($\bigvee \Psi$) where Ψ is a set of formulas satisfying the condition:

(R) *At most one formula in* Ψ *having quantifiers is not a sentence.*

The motivation of this paper is to study the *restricted conjunction fragment* $\exists^+ \mathscr{N}_k$ of $\exists^+ \mathscr{L}_k$ where all formulas are built using existential quantifiers, disjunctions, restricted conjunctions, and atomic formulas.

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

2.5 Spoiler-Duplicator games

The relations $\Rrightarrow^{\mathscr{J}}$ and $\equiv^{\mathscr{J}}$ for specific choices of $\mathscr{J}$ are characterized, in a syntax-free fashion, by Spoiler-Duplicator games (also called model-comparison games or Ehrenfeucht-Fraïssé style games). We will consider two structures $\mathscr{A}$ and $\mathscr{B}$. Each game has two players. Spoiler, who is trying to show that two structures are different under $\mathscr{J}$ and Duplicator, who is trying to show the two structures are the same under $\mathscr{J}$ (i.e $\mathscr{A} \equiv^{\mathscr{J}} \mathscr{B}$). Each game is played in a number of rounds.

- For the Ehrenfeucht-Fraïssé game $\mathbf{EF}_k(\mathscr{A}, \mathscr{B})$ characterizing $\equiv^{\mathscr{L}^k}$, at each n-th round, $n \in [\mathsf{k}]$:
  - Spoiler chooses an element in either structure $a_n \in A$ or $b_n \in B$.
  - Duplicator chooses an element in the other structure $b_n \in B$ or $a_n \in A$.

  At the end of the n-th round, a pair of sequences $[a_1, \ldots, a_n]$ and $[b_1, \ldots, b_n]$ were chosen. If the relation $\gamma_n = \{(a_j, b_j) \mid j \in [\mathsf{n}]\}$ is a partial isomorphism from $\mathscr{A}$ to $\mathscr{B}$, then Duplicator wins the n-th round. Otherwise, Spoiler wins. Duplicator has a winning strategy in $\mathbf{EF}_k(\mathscr{A}, \mathscr{B})$ if for every move Spoiler makes in the k-round game, Duplicator has a corresponding winning move. For this game, the resource is the number of rounds k corresponding to the syntactic resource of quantifier rank in the definition of $\mathscr{L}^k$.

- For the pebbling game $\mathbf{Peb}_k(\mathscr{A}, \mathscr{B})$ characterizing $\equiv^{\mathscr{L}_k}$, Spoiler and Duplicator both have a set of k pebbles. At each n-th round, $n \in \omega$:
  - Spoiler places some pebble $p_n \in [\mathsf{k}]$ on an element in either structure $a_n \in A$ or $b_n \in B$. If the pebble $p_n \in [\mathsf{k}]$ was already placed on a previous element, the player moves the pebble from that element to the newly chosen element.
  - Duplicator places their pebble $p_n \in [\mathsf{k}]$ on an element in the other structure $b_n \in B$ or $a_n \in A$.

  At the end of the n-th round, a pair of sequences $s = [(p_1, a_1), \ldots, (p_n, a_n)]$ and $t = [(p_1, b_1), \ldots, (p_n, b_n)]$ where chosen. For every $p \in [\mathsf{k}]$, let $a^p = \mathsf{last}_p(s)$ and $b^p = \mathsf{last}_p(t)$ be the last element pebbled with p in s and t (respectively). If the relation $\gamma_n = \{(a^p, b^p) \mid p \in [\mathsf{k}]\}$ is a partial isomorphism from $\mathscr{A}$ to $\mathscr{B}$, then Duplicator wins n-th round of the k-pebble game. Duplicator has a winning strategy in $\mathbf{Peb}_k(\mathscr{A}, \mathscr{B})$ if for every round $n \in \omega$ and move by Spoiler in the n-th round, Duplicator has a winning move. For this game, the resource is the number of pebbles k corresponding to the syntactic resource of variables in the definition of $\mathscr{L}_k$.

For both of the games $\mathbf{G}_k = \mathbf{EF}_k(\mathscr{A}, \mathscr{B})$ or $\mathbf{G}_k = \mathbf{Peb}_k(\mathscr{A}, \mathscr{B})$, we can consider one-sided variants $\exists \mathbf{G}_k = \exists \mathbf{EF}_k(\mathscr{A}, \mathscr{B})$ or $\exists \mathbf{G}_k = \exists \mathbf{Peb}_k(\mathscr{A}, \mathscr{B})$ (respectively). These games characterize $\mathscr{A} \Rrightarrow^{\exists^+ \mathscr{L}^k} \mathscr{B}$ and

$\mathscr{A} \Rrightarrow^{\exists^+\mathscr{L}_k} \mathscr{B}$ (respectively). The one-sided variants are obtained by restricting Spoiler to only choose (or place a pebble) on elements of $\mathscr{A}$ and consequently, Duplicator only plays in $\mathscr{B}$. We also relax the winning condition so that the relation γ is only required to be a partial homomorphism from $\mathscr{A}$ to $\mathscr{B}$.

Additionally, we can also consider bijection variants of these games $\#\mathbf{G}_k = \#\mathbf{EF}_k(\mathscr{A},\mathscr{B})$ or $\#\mathbf{G}_k = \#\mathbf{Peb}_k(\mathscr{A},\mathscr{B})$. For both games, Spoiler automatically wins if there does not exist a bijection between $A \to B$. At each n-th round in both games, Duplicator chooses a bijection $f_n : A \to B$. In the case of $\#\mathbf{EF}_k(\mathscr{A},\mathscr{B})$, Spoiler chooses an element $a \in A$. Duplicator wins the n-th round of $\#\mathbf{EF}_k(\mathscr{A},\mathscr{B})$ if the relation $\gamma = \{(a_j, f_j(a_j)) \mid j \in [\mathsf{n}]\}$ is a partial isomorphism. In the case of $\#\mathbf{Peb}_k(\mathscr{A},\mathscr{B})$, Spoiler plays first by placing a pebble p on an element $a \in A$ and Duplicator responds with a bijection $f_n : A \to B$ that is consistent with the previously placed pebbles, i.e. for every $q \neq p$ such that $(a^q, b^q) \in \gamma_{n-1}$, $f_n(a^q) = b^q$. Duplicator wins the n-th round of $\#\mathbf{Peb}_k(\mathscr{A},\mathscr{B})$ if the relation $\gamma = \{(a^p, f_n(b^p)) \mid p \in [\mathsf{k}]\}$ is a partial isomorphism.

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

3 Spoiler-Duplicator game comonads

The larger research program motivating this paper is the emerging notion of a Spoiler-Duplicator game comonad $\mathbb{C}_k$ associated with logic $\mathscr{J}$ graded by a resource k. Though this notion of a Spoiler-Duplicator game comonad has no formal definition, each of the papers [1, 4, 6, 2] define specific indexed families of comonads $\mathbb{C}_k$ over $\mathfrak{R}(\sigma)$ that all exhibit the same pattern of results: a family of morphism-power theorems, a coalgebra characterization theorem, and a parameterized Chandra-Merlin correspondence. In this section, we will go over the general scheme for each of these of results. Throughout, we will also state the result for the specific cases of the Ehrenfeucht-Fraïssé comonad $\mathbb{E}_k$ and pebbling comonad $\mathbb{P}_k$.

We first define the Ehrenfeucht-Fraïssé comonad $(\mathbb{E}_k, \varepsilon, ()^*)$. Given a σ-structure $\mathscr{A}$, we define the universe of $\mathbb{E}_k\mathscr{A}$ as $A^{\leq k} - \{\varepsilon\}$. The counit morphism $\varepsilon_{\mathscr{A}} : \mathbb{E}_k\mathscr{A} \to \mathscr{A}$ is defined as the last element of the sequence, i.e $\varepsilon_{\mathscr{A}}([a_1,\ldots,a_n]) = a_n$. The coextension of a morphism of type $f : \mathbb{E}_k\mathscr{A} \to \mathscr{B}$ is defined as $f^*[a_1,\ldots,a_j] = [b_1,\ldots,b_j]$ where $f[a_1,\ldots,a_i] = b_i$ for all $i \in [\mathsf{j}]$. To define the σ-structure on $\mathbb{E}_k\mathscr{A}$, suppose $R \in \sigma$ is an m-ary relation:

$$\begin{aligned} R^{\mathbb{E}_k\mathscr{A}}(s_1,\ldots,s_m) \Leftrightarrow\ & i,j \in [\mathsf{m}], s_i \sqsubseteq s_j \text{ or } s_j \sqsubseteq s_i && \textit{pairwise comparable} \\ & \text{and } R^{\mathscr{A}}(\varepsilon_{\mathscr{A}}(s_1),\ldots,\varepsilon_{\mathscr{A}}(s_m)) && \textit{compatibility} \end{aligned}$$

Intuitively, $\mathbb{E}_k\mathscr{A}$ is the set of Spoiler plays in k-round Ehrenfeucht-Fraïssé game on the structure $\mathscr{A}$. The pairwise comparability condition is used in the interpretation of relations in order to ensure that plays are only related when they are fragments of the same Spoiler play.

We define the pebbling comonad $(\mathbb{P}_k, \varepsilon, ()^*)$. Given a σ-structure $\mathscr{A}$, we define the universe of $\mathbb{P}_k\mathscr{A}$ as $([\mathsf{k}] \times A)^+$. The counit morphism $\varepsilon_{\mathscr{A}} : \mathbb{P}_k\mathscr{A} \to \mathscr{A}$ is defined as the second component of the last element of the sequence, i.e $\varepsilon_{\mathscr{A}}[(p_1,a_1),\ldots,(p_n,a_n)] = a_n$. The coextension of a morphism $f : \mathbb{P}_k\mathscr{A} \to \mathscr{B}$ is defined as $f^*[(p_1,a_1),\ldots,(p_n,a_n)] = [(p_1,b_1),\ldots,(p_n,b_n)]$ where $f[(p_1,a_1),\ldots,(p_i,a_i)] = b_i$ for all $i \in [\mathsf{n}]$. To define the σ-structure on $\mathbb{P}_k\mathscr{A}$, we will need some notation. The mapping $\pi_A : \mathbb{P}_kA \to [\mathsf{k}]$ 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:

$$R^{\mathbb{P}_k\mathscr{A}}(s_1,\ldots,s_m) \Leftrightarrow i,j \in [\mathsf{m}], s_i \sqsubseteq s_j \text{ or } s_j \sqsubseteq s_i \qquad \textit{pairwise comparable}$$
$$\text{and } \pi_A(s_i) \text{ does not appear in the suffix of } s_i \text{ in } s_j \text{ for any } s_i \sqsubseteq s_j \qquad \textit{active pebble}$$
$$\text{and } R^{\mathscr{A}}(\varepsilon_{\mathscr{A}}(s_1),\ldots,\varepsilon_{\mathscr{A}}(s_m)) \qquad \textit{compatibility}$$

Just as with the $\mathbb{E}_k$ case, $\mathbb{P}_k\mathscr{A}$ can be seen as the set of Spoiler plays in the k-pebble game on the structure $\mathscr{A}$. The additional active pebble condition models how Spoiler's choice of k pebble placements is moving a k- sized window of variables being assigned to elements in the structure $\mathscr{A}$. When working with $\mathbb{P}_k\mathscr{A}$, we will use the notation $\mathsf{last}_p(s)$ to denote the last element $a \in \mathscr{A}$ in the sequence s pebbled by $p \in [\mathsf{k}]$. The set of active elements of s is then $\mathsf{Active}(s) = \{\mathsf{last}_p(s) \mid p \in [\mathsf{k}]\}$. Given an $s \in \mathbb{P}_k A$ and $t \in \mathbb{P}_k B$ we can form the relation of pebbled elements $\gamma_{s,t} = \{(\mathsf{last}_p(s), \mathsf{last}_p(t))\} \subseteq A \times B$.

3.1 Morphism Power Theorems

Given a logic $\mathscr{J}_k$ graded by some syntactic resource k and corresponding modal comparison game $\mathbf{G}_k$, the Spoiler-Duplicator comonad associated with $\mathbf{G}_k$ is an indexed family of comonads $\mathbb{C}_k$ over $\mathfrak{R}(\sigma)$. We can then leverage the Kleisli category $\mathscr{K}(\mathbb{C}_k)$ associated with $\mathbb{C}_k$ to capture the preservation relation $\Rrightarrow^{\exists^+\mathscr{J}_k}$ of the sentences in the existential-positive fragment of $\mathscr{J}_k$ limited by the resource k.

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

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

For each $\mathbb{C}_k$, we can define the relation $\to_k^{\mathbb{C}}$ where $\mathscr{A} \to_k^{\mathbb{C}} \mathscr{B}$ if there exists a coKleisli morphism $f : \mathbb{C}_k\mathscr{A} \to \mathscr{B}$.

Theorem. *In the case of $\mathbb{E}_k$ and $\mathbb{P}_k$, we have the following morphism-power results:*

- $\mathscr{A} \to_k^{\mathbb{E}} \mathscr{B} \Leftrightarrow \mathscr{A} \Rrightarrow^{\exists^+\mathscr{L}^k} \mathscr{B} \Leftrightarrow$ *Duplicator has a winning strategy in $\exists\boldsymbol{EF}_k(\mathscr{A},\mathscr{B})$ [4]*
- $\mathscr{A} \to_k^{\mathbb{P}} \mathscr{B} \Leftrightarrow \mathscr{A} \Rrightarrow^{\exists^+\mathscr{L}_k} \mathscr{B} \Leftrightarrow$ *Duplicator has a winning strategy in $\exists\boldsymbol{Peb}_k(\mathscr{A},\mathscr{B})$ [1, 4]*

Ostensibly, the asymmetry of a coKleisli morphism means that only the "forth" aspect of the game $\mathbf{G}$ (i.e. the game $\exists\mathbf{G}$) 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 exists morphisms $f : \mathbb{C}_k\mathscr{A} \to \mathscr{B}$ and $g : \mathbb{C}_k\mathscr{B} \to \mathscr{A}$ such that $g \circ_{\mathscr{K}} f = \varepsilon_{\mathscr{A}}$ and $f \circ_{\mathscr{K}} g = \varepsilon_{\mathscr{B}}$. However, this relation turns out be too strong and characterizes equivalence in a logic stronger than $\mathscr{J}_k$. The isomorphisms in $\mathscr{K}(\mathbb{C}_k)$ characterize the equivalence relation $\equiv^{\#\mathscr{J}_k}$ for the logic $\mathscr{J}_k$ extended with counting quantifiers.

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

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

For each $\mathbb{C}_k$, we can define the relation $\cong_k^{\mathbb{C}}$ where $\mathscr{A} \cong_k^{\mathbb{C}} \mathscr{B}$ if there exists a coKleisli isomorphism between $\mathscr{A}$ and $\mathscr{B}$.

Theorem. *In the case of $\mathbb{E}_k$ and $\mathbb{P}_k$, we have the following isomorphism-power results:*

- $\mathscr{A} \cong_k^{\mathbb{E}} \mathscr{B} \Leftrightarrow \mathscr{A} \equiv^{\#\mathscr{L}^k} \mathscr{B} \Leftrightarrow$ *Duplicator has a winning strategy in* $\#\mathbf{EF}_k(\mathscr{A},\mathscr{B})$ *[4]*
- $\mathscr{A} \cong_k^{\mathbb{P}} \mathscr{B} \Leftrightarrow \mathscr{A} \equiv^{\#\mathscr{L}_k} \mathscr{B} \Leftrightarrow$ *Duplicator has a winning strategy in* $\#\mathbf{Peb}_k(\mathscr{A},\mathscr{B})$ *[1, 4]*

Since isomorphism is too strong, we need to define a different symmetric relation on the objects of $\mathfrak{R} = \mathfrak{R}(\sigma)$. In order to do this, we utilize a variant of the established notion of an *open map bisimulation* of Joyal, Nielsen, and Winksel [10], in the category of coalgebras $\mathfrak{R}^{\mathbb{C}_k}$. Since this new notion of bisimulation is constructed in the category of coalgebras, we discuss the bisimulation power theorem in section 3.2.

3.2 Coalgebras and adjunctions

A natural question for the comonads $\mathbb{C}_k$ is to understand the category of coalgebras $\mathscr{E}\mathscr{M}(\mathbb{C}_k)$. Beautifully, for all the cases of $\mathbb{C}_k$ constructed from some game $\mathbf{G}_k$, the coalgebras $\mathscr{A} \to \mathbb{C}_k\mathscr{A}$ correspond to forest covers of structures $\mathscr{A}$.

Theorem (Coalgebra Characterization)**.** *There is a bijective correspondence:*

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

In the case of $\mathbb{E}_k$, the coalgebras of $\mathscr{A} \to \mathbb{E}_k\mathscr{A}$ correspond to forest covers of $\mathscr{G}(\mathscr{A})$ of height k. Given a undirected graph $G = (V, \frown)$, a *forest cover of G* is tree-order $(F, \leq)$ where $F = V$ and $v \frown v' \in G$ implies that $v \leq v'$ or $v' \leq v$. The *height of the forest cover* $(F, \leq)$ is the length of the longest path in the forest F. The *tree-depth* of a structure $\mathsf{td}(\mathscr{A})$ is the minimum height of a forest covers $(F, \leq)$ of $\mathscr{G}(\mathscr{A})$.

Theorem. *There is a bijective correspondence between:*

1. *Forest covers of* $\mathscr{G}(\mathscr{A})$ *of height* $\leq k$
2. $\mathbb{E}_k$*-coalgebras* $\alpha : \mathscr{A} \to \mathbb{E}_k\mathscr{A}$

In the case of $\mathbb{P}_k$, the coalgebras of $\mathscr{A} \to \mathbb{P}_k\mathscr{A}$ correspond to k-pebble forest covers over $\mathscr{G}(\mathscr{A})$. A *k-pebble forest cover of a graph G* is a triple $(F, \leq, p : F \to [\mathsf{k}])$ where $(F, \leq)$ is a forest cover of G and p is pebble assignment function such that if $v \frown v'$ and $v \leq v'$, then for all $w \in (v, v']$, $p(v) \neq p(w)$. This notion was introduced in [1], under the name k-traversal, to provide a bridge between $\mathbb{P}_k$-coalgebras and the graph-theoretic notion of a tree decomposition of G of width $< k$.

Theorem. *The following are equivalent:*

1. $\mathscr{G}(\mathscr{A})$ *has a tree decompositions of width* $< k$
2. $\mathscr{G}(\mathscr{A})$ *has a k-pebble forest cover (i.e. there exists a* $\mathbb{P}_k$*-coalgebra* $\alpha : \mathscr{A} \to \mathbb{P}_k\mathscr{A}$*)*

We will introduce tree decompositions, and the linear variant–path decompositions, in section 5. The associated combinatorial invariant is the *treewidth* of a structure $\mathsf{tw}(\mathscr{A})$ which is the minimum width of a tree decomposition of $\mathscr{G}(\mathscr{A})$.

These results allows us to obtain coalgebraic characterizations of the combinatorial invariants that correspond to the tree-like decomposition. Namely, for every game comonad $\mathbb{C}_k$, we define the $\mathbb{C}$-*coalgebra number of* $\mathscr{A}$ to be the least k, denoted $\kappa^{\mathbb{C}}(\mathscr{A})$, such that there exists a coalgebra $\mathscr{A} \to \mathbb{C}_k\mathscr{A}$.

- For all finite σ-structures $\mathscr{A}$, $\kappa^{\mathbb{E}}(\mathscr{A}) = \mathsf{td}(\mathscr{A})$
- For all finite σ-structures $\mathscr{A}$, $\kappa^{\mathbb{P}}(\mathscr{A}) = \mathsf{tw}(\mathscr{A}) + 1$

The bijective correspondence between forest covers of a certain type, k-height or k-pebble, and coalgebras over $\mathbb{E}_k$ and $\mathbb{P}_k$ (respectively) can actually be extended to isomorphisms of the respective categories. That is, the category of k-height forest covers is isomorphic to $\mathscr{EM}(\mathbb{E}_k)$. Similarly, the category of k-pebble forest covers is isomorphic to $\mathscr{EM}(\mathbb{P}_k)$. In the language of adjunctions, for every $\mathbb{C}_k$, we can form a category of forest-ordered σ-structures $\mathfrak{T}^{\mathbb{C}}_k(\sigma)$ consisting of:

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

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

- $\mathfrak{T}^{\mathbb{C}}_k(\sigma)_1$ are σ-morphisms $f : \mathscr{A} \to \mathscr{B} \in \mathfrak{R}(\sigma)_1$ that preserve the ordering relation.
- 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 can realize $\mathbb{C}_k$ 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 where the adjunction $U_k \dashv F_k$ is comonadic. That is, $\mathfrak{T}^{\mathbb{C}}_k(\sigma)$ and $\mathscr{EM}(\mathbb{C}_k)$ are equivalent categories.

Theorem. *In the case of $\mathbb{E}_k$ and $\mathbb{P}_k$ we have the following constructions:*

- *For the case of $\mathbb{E}_k$, the right adjoint F_k is given by $F_k\mathscr{A} = (\mathbb{E}_k\mathscr{A}, \sqsubseteq)$ where $\mathfrak{T}^{\mathbb{E}}_k(\sigma)$ is the category of tree-ordered structures of depth $\leq k$.*
- *For the case of $\mathbb{P}_k$, the right adjoint F_k is given by $F_k\mathscr{A} = (\mathbb{P}_k\mathscr{A}, \sqsubseteq, \pi_{\mathscr{A}})$ where $\mathfrak{T}^{\mathbb{P}}_k(\sigma)$ is the category of tree-ordered structures with objects $(\mathscr{A}, \leq_{\mathscr{A}}, p)$ equipped with an additional pebbling function $p : A \to [k]$ satisfying the condition:*

$$\textit{(P) if } a \frown a' \textit{ and } a \leq_{\mathscr{A}} a', \textit{ then for all } b \in (a, a']_{\mathscr{A}}, p(a) \neq p(b).$$

  Morphisms in $\mathfrak{T}^{\mathbb{P}}_k(\sigma)$ also preserve the pebbling the function.

3.2.1 Bisimulation power theorem

Rendering the category of coalgebras of $\mathbb{C}_k$ as forest covers allows us to define a categorical notion of bisimulation between structures $\leftrightarrow^{\mathbb{C}}_k$. We use this notion of bisimulation in order. Alongside the morphism and isomorphism power theorems, such a bisimulation power theorem completes a picture that demonstrates $\mathbb{C}_k$ gives a syntax-free definition to three variants of the associated logic $\mathscr{J}_k$.:

Given a subcategory of "path" objects $\mathfrak{R}^P \hookrightarrow \mathfrak{R}^{\mathbb{C}_k}$, an $\mathfrak{R}^P$-*open morphism* is a morphism $h : (\mathscr{A}, \alpha) \to (\mathscr{B}, \beta)$ such that for every embedding $(\mathbf{P}, p) \rightarrowtail (\mathbf{Q}, q) \in \mathfrak{R}^P$ satisfying the diagram in $\mathfrak{R}^{\mathbb{C}_k}$:

$$\begin{array}{ccc} (\mathbf{P}, p) & \rightarrowtail & (\mathbf{Q}, q) \\ \downarrow & & \downarrow \\ (\mathscr{A}, \alpha) & \xrightarrow[h]{} & (\mathscr{B}, \beta) \end{array}$$

there exists an embedding $(\mathbf{Q}, q) \rightarrowtail (\mathscr{A}, \alpha)$ making the following diagram commute in $\mathfrak{R}^{\mathbb{C}_k}$:

$$\begin{array}{ccc} (\mathbf{P}, p) & \rightarrowtail & (\mathbf{Q}, q) \\ \downarrow & \swarrow & \downarrow \\ (\mathscr{A}, \alpha) & \xrightarrow[h]{} & (\mathscr{B}, \beta) \end{array}$$

This allows us to establish an intermediate power theorem demonstrating that $\mathbb{C}_k$ can be used to capture equivalence in the logic $\mathscr{J}_k$.

Theorem (Bisimulation Power Theorem)**.** *For all σ-structures $\mathscr{A}, \mathscr{B}$, the following are equivalent:*

1. *Duplicator has a winning strategy $\boldsymbol{G}_k(\mathscr{A}, \mathscr{B})$*
2. $\mathscr{A} \equiv^{\mathscr{J}_k} \mathscr{B}$
3. *There exists a span of $\mathfrak{R}^P$-open coalgebra morphisms:*

$$\begin{array}{ccc} & (\mathscr{R}, \rho) & \\ \swarrow^{f} & & \searrow^{g} \\ (\mathbb{C}_k\mathscr{A}, \delta_{\mathscr{A}}) & & (\mathbb{C}_k\mathscr{B}, \delta_{\mathscr{B}}) \end{array}$$

For each $\mathbb{C}_k$ we associate a subcategory of path objects $\mathfrak{R}^P \hookrightarrow \mathfrak{R}^{\mathbb{C}_k}$, this allows us define the relation $\leftrightarrow_k^{\mathbb{C}}$ where $\mathscr{A} \leftrightarrow_k^{\mathbb{C}} \mathscr{B}$ if there exists a span of $\mathfrak{R}^P$-open coalgebra morphisms $(\mathbb{C}_k\mathscr{A}, \delta_{\mathscr{A}}) \leftarrow (\mathscr{R}, \rho) \rightarrow (\mathbb{C}_k\mathscr{B}, \delta_{\mathscr{B}})$.

Theorem. *In the case of $\mathbb{E}_k$ and $\mathbb{P}_k$, we have the following bisimulation-power results:*

- $\mathscr{A} \leftrightarrow_k^{\mathbb{E}} \mathscr{B} \Leftrightarrow \mathscr{A} \equiv^{\mathscr{L}^k} \mathscr{B} \Leftrightarrow$ *Duplicator has a winning strategy in $\boldsymbol{EF}_k(\mathscr{A}, \mathscr{B})$ [4]*
- $\mathscr{A} \leftrightarrow_k^{\mathbb{P}} \mathscr{B} \Leftrightarrow \mathscr{A} \equiv^{\mathscr{L}_k} \mathscr{B} \Leftrightarrow$ *Duplicator has a winning strategy in $\boldsymbol{Peb}_k(\mathscr{A}, \mathscr{B})$ [1, 4]*

We now turn to the grading structure of the family of comonads $\mathbb{C}_k$. Namely, for $k \leq l$, there is comonad inclusion $\mathbb{C}_k \hookrightarrow \mathbb{C}_l$. Interpreting this in terms of the game, this corresponds to the fact that Spoiler playing with k resources is a special case of Spoiler playing with $l \geq k$ resources. This means the smaller the k, the easier it is to find morphisms of the type $\mathbb{C}_k\mathscr{A} \rightarrow \mathscr{B}$.

4 Pebble-relation comonad

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

$$\mathbb{PR}_k A := \{([(p_1, a_1), \ldots, (p_n, a_n)], i) \mid (p_j, a_j) \in [\mathsf{k}] \times A \text{ and } i \in [\mathsf{n}]\}$$

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 $\varepsilon_{\mathscr{A}} : \mathbb{PR}_k\mathscr{A} \rightarrow \mathscr{A}$ is defined as $\varepsilon_{\mathscr{A}}([(p_1, a_1), \ldots, (p_n, a_n)], i) = a_i$. The coextension of a morphism $f : \mathbb{PR}_k\mathscr{A} \rightarrow \mathscr{B}$ 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 $f([(p_1, a_1), \ldots, (p_n, a_n)], j) = b_j$ for all $j \in [\mathsf{n}]$. We define the σ-structure on $\mathbb{PR}_k\mathscr{A}$ similarly to $\mathbb{P}_k\mathscr{A}$. The mapping $\pi_A : A \rightarrow [\mathsf{k}]$ is defined as $\pi_A([(p_1, a_1), \ldots, (p_n, a_n)], i) = p_i$. Suppose $R \in \sigma$ is a m-ary relation, then:

$$\begin{array}{lr} R^{\mathbb{PR}_k\mathscr{A}}((s, i_1), \ldots, (s, i_m)) \Leftrightarrow \text{let } i = \max\{i_1, \ldots, i_m\}, \pi_A(s, i_j) \text{ does not appear in } s(i_j, i] & \textit{active pebble} \\ \text{and } R^{\mathscr{A}}(\varepsilon_{\mathscr{A}}(s, i_1), \ldots, \varepsilon_{\mathscr{A}}(s, i_m)) & \textit{compatibility} \end{array}$$

Proposition 1. $(\mathbb{PR}_k, \varepsilon, ()^*)$ *is a comonad in coKleisli form.*

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

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

Proposition 2. *$\nu : \mathbb{PR}_k \to \mathbb{P}_k$ is an 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 $f : \mathscr{A} \to \mathscr{B}$:

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

It is clear that this diagram commutes, by observing that the $\mathbb{PR}_k f$ preserves the prefix relation on sequences. We present both comonads in standard form $(\mathbb{PR}_k, \varepsilon, \delta)$ 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 $\mathfrak{R}(\sigma)$:

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

□

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

$$\delta_{\mathscr{A}}(s,i) = ([(p_1,(s,1)),\dots,(p_n,(s,n))],i)$$
$$\delta'_{\mathscr{A}}(s) = [(p_1,s[1,1]),\dots,(p_n,s[1,n])]$$

The confirmation is then straightforward:

$$\begin{aligned} \nu_{\mathscr{A}} \circ \mathbb{PR}_k\nu_{\mathscr{A}} \circ \delta_{\mathscr{A}}(s,i) &= \nu_{\mathscr{A}} \circ \mathbb{PR}_k\nu_{\mathscr{A}}([(p_1,(s,1)),\dots,(p_n,(s,n))],i) && \delta_{\mathscr{A}} \text{ as above} \\ &= \nu_{\mathscr{A}}([(p_1,\nu_{\mathscr{A}}(s,1)),\dots,(p_n,\nu_{\mathscr{A}}(s,n))],i) && \text{functoriality of } \mathbb{PR}_k \\ &= \nu_{\mathscr{A}}([(p_1,s[1,1]),\dots,(p_n,s[1,n])],i) && \text{defn of } \nu_{\mathscr{A}} \\ &= [(p_1,s[1,1]),\dots,(p_i,s[1,i])] && \text{defn of } \nu_{\mathscr{A}} \\ &= \delta'_{\mathscr{A}}(s[1,i]) && \delta'_{\mathscr{A}} \text{ as above} \\ &= \delta'_{\mathscr{A}} \circ \nu_{\mathscr{A}}(s,i) && \text{defn of } \nu_{\mathscr{A}} \end{aligned}$$

5 Coalgebras and path decompositions

We now turn the initial motivation for the construction of a pebble-relation comonad: to give a categorical definition for the combinatorial parameter of pathwidth. There are many different characterizations of pathwidth. The original definition, introduced by Robertson and Seymour in [14], is in terms of path decompositions of a structure $\mathscr{A}$. This is a specialization of the notion of tree decompositions.

Definition 1. Give a σ-structure $\mathscr{A}$, a *path decomposition for* $\mathscr{A}$ is a triple $(X, \leq_X, \lambda)$ where $(X, \leq_X)$ is a linear ordered set and $\lambda : X \to \mathscr{P}A$ is a function satisfying the following conditions:

(PD1) For every $a \in A$, there exists an $x \in X$ such that $a \in \lambda(x)$, i.e. $A = \bigcup_{x \in X} \lambda(x)$.

(PD2) If $a \frown a' \in \mathscr{A}$, then $a, a' \in \lambda(x)$ for some $x \in X$.

(PD3) For all $y \in [x, x']$ (i.e. the interval between x, x' w.r.t $\leq$), $\lambda(x) \cap \lambda(x') \subseteq \lambda(y)$.

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

Definition 2. The pathwidth of a σ-structure $\mathscr{A}$, denoted $\mathsf{pw}(\mathscr{A})$ is the least k such that $\mathscr{A}$ has a path decomposition of width k.

The proof of [1] showing the correspondence between tree decompositions of width $< k$ and coalgebras over $\mathbb{P}_k$ made use of an intermediate structure, k-pebble forest covers. We will make use of the analogous notion in the linearly-ordered case: k-pebble linear forest covers.

Definition 3. Given a σ-structure $\mathscr{A}$, a *k-pebble linear forest cover for* $\mathscr{A}$ is a tuple $(\mathscr{F}, p)$ where $\mathscr{F} = \{(S_i, \leq_i)\}_{i \in [\mathsf{n}]}$ is a partition of A into linear ordered subsets and $p : A \to [\mathsf{k}]$ is a pebbling function such that the following hold:

(FC1) If $a \frown a' \in \mathscr{A}$, then there exists an i such that $a, a' \in S_i$.

(FC2) If $a \frown a'$ such that $a \leq_i a'$, then for all $b \in (a, a']_i \subseteq S_i$, $p(b) \neq p(a)$.

We aim to show that the existence of path decomposition $(X, \leq_X, \lambda)$ of width $< k$ for $\mathscr{A}$ is equivalent to the existence of a k-pebble linear forest cover for $\mathscr{A}$. In order to do this, we will 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 $x \in X$ of the path decomposition. If we define these functions in a consistent way, then they can be "glued" together to obtain p. We use the following definition to pick out consistent families of τ_x.

Definition 4. Given a path decomposition $(X, \leq_X, \lambda)$ of width $< k$ for $\mathscr{A}$, 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 hold:

- (Locally-injective) For every $x \in X$, τ_x is an injective function.
- (Glueability) For every $x, x' \in X$, $\tau_x|_{\lambda(x) \cap \lambda(x')} = \tau_{x'}|_{\lambda(x) \cap \lambda(x')}$.

To make this machinery useful, we show that every path decomposition has a pebbling section family.

Lemma 3. *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. We prove the hypothesis:

$$\{\tau_z\}_{z \in \downarrow x} \text{ is a } k\text{-pebbling section family for } (Y, \leq_Y, \lambda|_Y) \text{ where } Y = \downarrow x \text{ and } \leq_Y = \leq_X \cap (Y \times Y)$$

by induction on the linear order $\leq_X$. *Base Case:* Suppose r is the $\leq_X$-least element. By 5, the cardinality of $\lambda(r)$ is $\leq k$, therefore 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 any $x \in \downarrow r$ is equal to r. *Inductive Step:* Let x' be the immediate $\leq_X$-successor of x. By the inductive 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 any $y <_X x'$.

Claim: For every $y <_X x'$, $\lambda(y) \cap \lambda(x') \subseteq \lambda(x) \cap \lambda(x')$. 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 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 $\{i_1, \dots, i_m\}$ 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, $\tau_y|_{\lambda(y)\cap\lambda(x')} = \tau_{x'}|_{\lambda(y)\cap\lambda(x')}$. □

Theorem 4. *The following are equivalent:*

1. *$\mathscr{A}$ has a path decomposition of width $< k$*
2. *$\mathscr{A}$ has k-pebble linear forest cover*

Proof. (1) $\Rightarrow$ (2) Suppose $(X, \leq_X, \lambda)$ is a path decomposition of $\mathscr{A}$ of width $< k$. We define a family of linear ordered sets $\{(S_i, \leq_i)\}$, where each S_i is the vertex set of a connected component of $\mathscr{G}(\mathscr{A})$. To define the order $\leq_i$, we define an order on $\leq_A$ and realize $\leq_i$ 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 3, there exists a k-pebbling section family$\{\tau_x\}_{x\in X}$. We then define $\leq_A$:

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

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 the 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 linear 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 $a \frown a' \in \mathscr{A}$. This means a, a' are connected in $\mathscr{G}(\mathscr{A})$, and so are in the same connected component S_i of $\mathscr{G}(\mathscr{A})$.

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 an $x \in X$ 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 $x_b \in [x_a, x_{a'}]_X \subseteq [x_a, x]_X$, $a \in \lambda(x_b)$. Hence, $a, b \in \lambda(x_b)$ and by the injectivity of τ_{x_b}, $\tau_{x_b}(a) \neq \tau_{x_b}(b)$. Therefore, $p(a) \neq p(b)$.

(2) $\Rightarrow$ (1) Suppose $\mathscr{A}$ has k-pebble linear forest cover given by the partition $\{(S_i, \leq_i)\}_{i\in[\mathsf{n}]}$ and pebbling function $p : A \to [\mathsf{k}]$. We define a linear ordered set $(A, \leq_A)$ where $\leq_A$ is the ordered sum of the family $\{(S_i, \leq_i)\}_{i\in[\mathsf{n}]}$. 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 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 $\lambda(a)$ to be the set of active predecessors of a. The triple $(A, \leq_A, \lambda)$ is a path decomposition of $\mathscr{A}$ 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 $a \frown a' \in \mathscr{A}$. 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, $|\lambda(a')| > k$ for some $a' \in A$. Consider the pebbling function restricted to $\lambda(a')$, $p|_{\lambda(a')} : \lambda(a') \to [\mathsf{k}]$. By the pidegeonhole principle, there must exist $a, c \in \lambda(a')$ with $a \neq c$, such that $p(a) = p(c)$. Without loss of generality, assume $a <_A c$. 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 $c \in \lambda(a')$, $p(c) \neq p(a)$. Contradiction.

Theorem 5. *There is a bijective correspondence between:*

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

(1) $\Rightarrow$ (2) Suppose $\mathscr{A}$ has k-pebble linear forest cover given by the partition $\{(S_i, \leq_i)\}_{i \in [\mathsf{n}]}$ and pebbling function $p : A \to [\mathsf{k}]$. Since each $(S_i, \leq_i)$ is a linear ordered, we can present $(S_, \leq_i)$ as a chain:

$$a_1 \leq_i \cdots \leq_i a_{m_i}$$

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 $\alpha_i : S_i \to \mathbb{PR}_k\mathscr{A}$ be defined as $\alpha_i(a_j) = (t_i, j)$. Let $\alpha : A \to \mathbb{PR}_k A$ be $\alpha = \bigcup_{i \in [\mathsf{n}]} \alpha_i$. Since the collection of S_i partition A, α is well-defined. We must show that the function α is a coalgebra $\alpha : \mathscr{A} \to \mathbb{PR}_k\mathscr{A}$.

To verify that α is indeed a homomorphism, suppose $R \in \sigma$ is an m-ary relation and $R^{\mathscr{A}}(a_1, \ldots, a_m)$. By (FC1), $\{a_1, \ldots, a_m\} \subseteq S_i$ for some $i \in [\mathsf{n}]$. Therefore, for all $j \in [\mathsf{m}]$, $\alpha(a_j) = (t_i, z_j)$ 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 $R^{\mathbb{PR}_k\mathscr{A}}$ and the supposition that $R^{\mathscr{A}}(a_1, \ldots, a_m)$, $R^{\mathbb{PR}_k\mathscr{A}}((t_i, z_1), \ldots, (t_i, z_m))$.

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

$$\begin{aligned}
\varepsilon_{\mathscr{A}} \circ \alpha(a) &= \varepsilon_{\mathscr{A}} \circ \alpha_i(a) && \text{by supposition } a \in S_i \\
&= \varepsilon_{\mathscr{A}}(t_i, j) && \text{defn of } \alpha_i \\
&= a && \text{defn of } t_i \text{ and } \varepsilon_A
\end{aligned}$$

To verify that α satisfies the comultiplication-coalgebra law,

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

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

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

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 $\mathscr{A}$.

To verify that $\{(S_t, \leq_t)\}$ is a partition of linear ordered subsets, suppose there exists an $a \in S_t \cap S_{t'}$, then we want to show that $t = t'$. Since $a \in S_t$, then $\alpha(a) = (t, j)$ for some $j \in [|t|]$. Similarily, since $a \in S_{t'}$, $\alpha(a) = (t', j')$ for some $j' \in [|t'|]$. By α being a well-defined function, $\alpha(a) = (t, j) = (t', j')$, so $t = t'$. Finally, by construction and $\{1,\dots,|t|\} = [|t|]$ being linear ordered by $\leq$, the order $\leq_t$ is a linear ordering.

To verify (FC1), suppose $a \frown a'$ 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 $R^{\mathbb{PR}_k\mathscr{A}}$ for all $R \in \sigma$, elements of $\mathbb{PR}_k\mathscr{A}$ 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 $a \frown a'$ with $a \leq_t a'$ and $b \in (a, a']_t$. We want to show that $p(b) \neq p(a)$. Since $a \frown a'$, there is some m-tuple $\vec{a} \in R^{\mathscr{A}}$ for some m-ary relation $R \in \sigma$. By α being a homomorphism, there exists $(t, j), (t, j') \in \alpha(\vec{a}) \in R^{\mathbb{PR}_k\mathscr{A}}$ for some $j \leq j' \in [|t|]$ such that $\alpha(a) = (t, j)$ and $\alpha(a') = (t, j')$. Moreover, since $b \in (a, a']_t$ there exists an $i \in (j, j']$ 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 $R^{\mathbb{PR}_k\mathscr{A}}$, $\pi_A(t, j)$ does not appear in $t(j, j']$, so $p(b) = \pi_A(t, i) \neq \pi_A(t, j) = p(a)$. ◻

Corollary 6. *For all σ-structures $\mathscr{A}$, $\mathsf{pw}(\mathscr{A}) = \kappa^{\mathbb{PR}}(\mathscr{A}) - 1$*

Proof. By theorem 4 and 5, a structure $\mathscr{A}$ has a path decomposition of width $< k$ iff $\mathscr{A}$ has a coalgebra $\mathscr{A} \to \mathbb{PR}_k\mathscr{A}$. Hence, $\mathsf{pw}(\mathscr{A}) + 1 \leq \kappa^{\mathbb{PR}}(\mathscr{A})$ and $\kappa^{\mathbb{PR}}(\mathscr{A}) \leq \mathsf{pw}(\mathscr{A}) + 1$, by the definition of $\mathsf{pw}(\mathscr{A})$ as the minimal width of a path decomposition for $\mathscr{A}$ and $\kappa^{\mathbb{PR}}$ as the minimal index for a $\mathbb{PR}_k$-coalgebra of $\mathscr{A}$. ◻

We can extend theorem 5 to an equivalence of categories. Consider the category $\mathfrak{T}_k^{\mathbb{PR}}(\sigma)$ of tree-ordered σ-structures $(\mathscr{A}, \leq_A, p)$ with $\mathscr{A} \in \mathfrak{R}(\sigma)_0$ and pebbling function $p : A \to [\mathsf{k}]$ satisfying the conditions:

- (E) If $a \frown a' \in \mathscr{A}$, then either $a \leq_A a'$ and $a' \leq_A a$.
- (P) If $a \frown a'$ 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 last condition (L) taken together with $\leq_A$ being a tree-order (i.e. for every $a \in A$, $\downarrow a$ is linear ordered by $\leq_A$) means $(A, \leq_A)$ is a disjoint union of linear ordered sets. That is, $(A, \leq_A)$ is a linear forest. Morphisms in $\mathfrak{T}_k^{\mathbb{PR}}(\sigma)$ are σ-morphisms that preserve the ordering relation and pebbling function. There is an evident forgetful functor $U_k : \mathfrak{T}_k^{\mathbb{PR}}(\sigma) \to \mathfrak{R}(\sigma)$, i.e. $(\mathscr{A}, \leq_A, p) \mapsto \mathscr{A}$. Consider the functor $F_k : \mathfrak{R}(\sigma) \to \mathfrak{T}_k^{\mathbb{PR}}(\sigma)$ with

- Object mapping $\mathscr{A} \mapsto (\mathbb{PR}_k\mathscr{A}, \leq^*, \pi_{\mathscr{A}})$ where $(t,i) \leq^* (t',j)$ iff $t = t'$ and $i \leq j$, and
- Morphism mapping $f \mapsto \mathbb{PR}_k f$

This functor F_k is right adjoint to U_k. In fact, we show that this adjunction yields an equivalence between $\mathfrak{T}_k^{\mathbb{PR}}(\sigma)$ and $\mathscr{E}\mathscr{M}(\mathbb{PR}_k)$.

Theorem 7. *For each $k > 0$, $U_k \dashv F_k$ is a comonadic adjunction. Moreover, $\mathbb{PR}_k$ is the comonad arising from this adjunction.*

6 Logical equivalences

Though the construction of $\mathbb{PR}_k$ was motivated by capturing pathwidth as coalgebras, the other Spoiler-Duplicator game comonads $\mathbb{C}_k$ explored in [1, 4, 6, 2] 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 $\mathbb{PR}_k\mathscr{A} \to \mathscr{B}$ correspond to Duplicator's winning strategy in a one-sided pebble-relation game. This game characterizes preservation of sentences in the existential positive k-variable logic with restricted conjunctions $\exists^+\mathscr{N}_k$. First, we will prove this result, indirectly, by utilizing Dalmau's results linking his pebble-relation game to pathwidth [7].

We began by introducing Dalmau's k-pebble relation game $\exists\mathbf{PebR}_k(\mathscr{A}, \mathscr{B})$ from $\mathscr{A}$ to $\mathscr{B}$ characterizing $\mathscr{A} \Rrightarrow^{\exists^+\mathscr{N}_k} \mathscr{B}$. Each round of the game ends with a pair (I, T) where $I \subseteq A$ is a domain such that $|I| \leq k$ and $T \subseteq \mathfrak{R}(\sigma)_1(\mathscr{A}|_I, \mathscr{B})$ is as set of σ-morphisms from $\mathscr{A}|_I$ to $\mathscr{B}$. At round 0, $I = \varnothing$ and $T = \{\lambda\}$ where λ is the unique function from $\varnothing$ to $\mathscr{B}$. At each subsequent round $n > 0$, suppose (I, T) is the configuration of the previous round $n-1$, then Spoiler moves first with two possible moves:

- A shrinking move, where Spoiler chooses a smaller domain $I' \subseteq I$
  - Duplicator chooses T' to be restrictions of the morphisms in T to I', i.e. $T' = \{h|_{I'} \mid h \in T\}$.
- A blowing move when $|I| < k$, where Spoiler chooses a larger domain $I \subseteq I'$ with $|I'| \leq k$.
  - Duplicator responds by choosing a set of σ-morphisms T' from $\mathscr{A}|_{I'}$ to $\mathscr{B}$ extending a subset S of the morphisms in 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. Spoiler wins otherwise. That is, Spoiler wins if Duplicator can not successfully extend any of the morphisms in T to I'.

Intuitively, in the $\exists\mathbf{PebR}_k$ game, Duplicator is given the advantage of non-determinism by being able to respond with a set of partial homomorphisms where she is only obligated to extend some of the partial homomorphisms chosen in the previous round. In fact, we can recover the ordinary one-sided k-pebble game $\exists\mathbf{Peb}_k$ by insisting that Duplicator always responds with a singleton set T. In order to show that coKleisli morphisms $\mathbb{PR}_k\mathscr{A} \to \mathscr{B}$ capture Duplicator's winning strategies in $\exists\mathbf{PebR}_k(\mathscr{A}, \mathscr{B})$, we need to introduce some notation.

Definition 5. For all comonads $\mathbb{T} : \mathfrak{C} \to \mathfrak{C}$ and objects $A, B \in \mathfrak{C}$, *A maps to B up to $\mathbb{T}$-coalgebras*, denoted $A \xrightarrow{\mathbb{T}} B$, if for all coalgebras $\alpha : C \to \mathbb{T}C$, $C \to A \Rightarrow C \to B$

Similarly, we say that $\mathscr{A}$ *maps to* $\mathscr{B}$ *up to pathwidth* $< k$, denoted $\mathscr{A} \xrightarrow{\mathsf{pw}<k} \mathscr{B}$, if for all σ-structures $\mathscr{C}$ such that $\mathsf{pw}(\mathscr{C}) < k$, $\mathscr{C} \to \mathscr{A} \Rightarrow \mathscr{C} \to \mathscr{B}$. A general fact about comonads is that the relation $A \xrightarrow{\mathbb{T}} B$ is exactly existence of a coKleisli morphism from A to B.

Proposition 8. *Let* $(\mathbb{T} : \mathfrak{C} \to \mathfrak{C}, \varepsilon, \delta)$ *be a comonad, then* $f : \mathbb{T}A \to B \Leftrightarrow A \xrightarrow{\mathbb{T}} B$

Proof. $\Rightarrow$ Suppose $h : C \to A \in \mathfrak{C}_1$, then given $f : \mathbb{T}A \to B$ and $\alpha : C \to \mathbb{T}C$, we can form the composition $f \circ \mathbb{T}h \circ \alpha : C \to B$.

$\Leftarrow$ Suppose that for all coalgebras, $\alpha : C \to \mathbb{T}C$, $C \to A \Rightarrow C \to B$. In particular, if we let $C = \mathbb{T}A$ and let $\alpha = \delta_A : \mathbb{T}A \to \mathbb{T}\mathbb{T}A$, i.e. the cofree coalgebra over A, then by $\varepsilon_A : \mathbb{T}A \to A$, there exists an $f : \mathbb{T}A \to B$. □

The seminal theorem of Dalmau's paper [7] links pathwidth, pebble-relation games, and preservation of sentences in the logic $\exists^+\mathscr{N}_k$. We will utilize this result, along with our coalgebra characterization theorem, i.e. theorems 4,5, in order to obtain indirectly the morphism power theorem.

Theorem 9 (Theorem 1 of [7])**.** *The following are equivalent for all* σ*-structures* $\mathscr{A}, \mathscr{B}$*:*

1. *Duplicator has a winning strategy in* $\exists\textbf{\textit{PebR}}_k(\mathscr{A}, \mathscr{B})$.
2. $\mathscr{A} \Rrightarrow^{\exists^+\mathscr{N}_k} \mathscr{B}$
3. $\mathscr{A} \xrightarrow{\mathsf{pw}<k} \mathscr{B}$

The comonad $\mathbb{PR}_k$ allows us to add coKleisli morphisms to this equivalence, obtaining the morphism power theorem for $\mathbb{PR}_k$:

Theorem 10. *The following are equivalent for all* σ*-structures* $\mathscr{A}, \mathscr{B}$

1. *There exists a coKleisli morphism* $f : \mathbb{PR}_k\mathscr{A} \to \mathscr{B}$
2. *Duplicator has a winning strategy in* $\exists\textbf{\textit{PebR}}_k(\mathscr{A}, \mathscr{B})$

Proof. We have the following chain of equivalences:

$$\begin{aligned} f : \mathbb{PR}_k\mathscr{A} \to \mathscr{B} &\Leftrightarrow \mathscr{A} \xrightarrow{\mathbb{PR}_k} \mathscr{B} && \text{proposition 8, } \mathbb{T} = \mathbb{PR}_k \\ &\Leftrightarrow \mathscr{A} \xrightarrow{\mathsf{pw}<k} \mathscr{B} && \text{theorems 4, 5} \\ &\Leftrightarrow \text{Duplicator has a winning strategy } \exists\mathbf{PebR}_k(\mathscr{A}, \mathscr{B}) && \text{theorem 9} \end{aligned}$$

□

Though we are able to obtain this morphism power theorem for $\mathbb{PR}_k$ indirectly, a direct proof would provide hints towards proofs for the bisimulation and isomorphism power theorems for $\mathbb{PR}_k$. The relationship between $\mathbb{PR}_k$ and the game $\exists\mathbf{PebR}_k$ is bit different than the other Spoiler-Duplicator game comonads. This ensures that a direct proof will be different in character than the analogous proofs for $\mathbb{E}_k$ and $\mathbb{P}_k$. For the other Spoiler-Duplicator game comonads, the action of the comonad on a structure has a nice intuitive description in terms of the corresponding game. Namely, elements of $\mathbb{E}_k\mathscr{A}$ could be seen as Spoiler's play in a one-sided k-round Ehrenfeucht-Fraïssé game $\exists\mathbf{EF}_k$ from $\mathscr{A}$ to another structure. Set functions from $\mathbb{E}_k\mathscr{A} \to \mathscr{B}$ could be seen as Duplicator's responses in $\exists\mathbf{EF}_k(\mathscr{A}, \mathscr{B})$. The interpretation $R^{\mathbb{E}_k\mathscr{A}}$ of the relations $R \in \sigma$ were chosen to force σ-morphisms $\mathbb{E}_k\mathscr{A} \to \mathscr{B}$ to correspond to **winning** strategies for Duplicator in $\exists\mathbf{EF}_k(\mathscr{A}, \mathscr{B})$. A similar story could be told for the pebbling comonad $\mathbb{P}_k$ and the other Spoiler-Duplicator game comonads appearing in [6, 2]. By contrast, it is not clear how elements of $\mathbb{PR}_k\mathscr{A}$ correspond to Spoiler's play in the $\exists\mathbf{PebR}_k$ game from $\mathscr{A}$. It is also not clear how functions $\mathbb{PR}_k\mathscr{A} \to \mathscr{B}$ correspond to Duplicator's responses in the game $\exists\mathbf{PebR}_k(\mathscr{A},\mathscr{B})$. However, we can view elements of $\mathbb{PR}_k\mathscr{A}$ as Spoiler's plays in a different, but equivalent, game: the one-sided all-in-one or pre-announced k-pebble game $\exists\mathbf{PPeb}_k(\mathscr{A},\mathscr{B})$ introduced in [16]. This game is played in one round. In the first and only round,

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

Duplicator wins if for every index $i \in [\mathsf{n}]$, the relation $\gamma_i = \{(\mathsf{last}_p(s[1,i]), \mathsf{last}_p(t[1,i])) \mid p \in [\mathsf{k}]\}$ is a partial homomorphism from $\mathscr{A}$ to $\mathscr{B}$.

Just as with the proofs with the morphism power theorems for $\mathbb{E}_k$ and $\mathbb{P}_k$, we must impose an additional I-morphism condition on morphisms of the type $\mathbb{PR}_k\mathscr{A} \to \mathscr{B}$. In the proof for $\mathbb{E}_k$, this condition was used to ensure that coKleisli morphisms $g : \mathbb{E}_k\mathscr{A} \to \mathscr{B}$ would map a sequence $s = [a_1,\dots,a_n]$ with a repetition, say $a_i = a_j$, to the same Duplicator response. That is, $b_i = g(s_i) = g(s_j) = b_j$ with s_i, s_j being the i-th subsequence and j-th subsequence of s (respectively). This ensured that Duplicator would respond with elements of $\mathscr{B}$ that generated a **well-defined** partial homomorphism. We can reformulate this use of I-morphism, by using the machinery of relative comonads. Namely, we expand the signature $\sigma^+ = \sigma \cup \{I\}$ to include a new binary relation I, and we consider the functor $J : \mathfrak{R}(\sigma) \to \mathfrak{R}(\sigma^+)$ where a σ-structure $\mathscr{A}$ is mapped to the same σ-structures with $I^{J\mathscr{A}}$ interpreted as the identity relation, i.e. $I^{J\mathscr{A}} = \{(a,a) \mid a \in \mathscr{A}\}$. Let $\mathbb{PR}_k^{\sigma^+}, \mathbb{P}_k^{\sigma^+}$ denote the comonads $\mathbb{PR}_k, \mathbb{P}_k$ on the category $\mathfrak{R}(\sigma^+)$ of this expanded signature σ^+. It is easy to show that $\mathbb{PR}_k^+ = \mathbb{PR}_k^{\sigma^+} \circ J$ and $\mathbb{P}_k^+ = \mathbb{P}_k^{\sigma^+} \circ J$ are relative comonads over J.

We can now give, using this $\exists\mathbf{PPeb}_k(\mathscr{A},\mathscr{B})$ game, a more direct proof of theorem 10.

Theorem 11. *The following are equivalent for all σ-structures $\mathscr{A},\mathscr{B}$:*

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

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

To verify f^* is indeed a homomorphism, suppose $R^{\mathbb{PR}_k\mathscr{A}}((s,i_1),\dots,(s,i_m))$ and let $i = \max\{i_1,\dots,i_m\}$. By the active pebble condition, for all $i_j \in \{i_1,\dots,i_m\}$, the pebble appearing at i_j does not appear in $s(i_j,i]$. Hence, the element $\varepsilon_{\mathscr{A}}(s,i_j)$ was the last element pebbled by the pebble $\pi_A(s,i_j)$ in $s[1,i]$, so $\varepsilon_{\mathscr{A}}(s,i_j)$ is in the domain of γ_i. By γ_i being a partial homomorphism, $R^{\mathscr{B}}(b_{i_1},\dots,b_{i_m})$, so $R^{\mathscr{B}}(\varepsilon_{\mathscr{B}}(t,i_1),\dots,\varepsilon_{\mathscr{B}}(t,i_m))$. The active pebble condition follows by t being compatible with s, so $R^{\mathbb{PR}_k\mathscr{B}}((t,i_1),\dots,(t,i_m))$.

(2) $\Rightarrow$ (3) We say a sequence $s = [(p_1,a_1),\dots,(p_n,a_n)]$ is *duplicating* if there exists a $i < j \in [\mathsf{n}]$ such that $p_i \neq p_j$, $\pi_A(s,i)$ does not appear in $s(i,j]$ (i.e. p_i is active) 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 a $j' \leq j$ with

$\varepsilon_{\mathscr{A}}(s',j') = \varepsilon_{\mathscr{A}}(s,j)$. Such a 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 $a_i = a_j$. Define $f^+(s,j) = f(s',j')$.

(3) $\Rightarrow$ (4) Recall, there exists a natural transformation $\nu_{J\mathscr{A}} : \mathbb{PR}_k^+ \to \mathbb{P}_k^+$. The pre-image of k-pebble plays in $\mathbb{P}_k\mathscr{A}$ under $\nu_{J\mathscr{A}}$ gives us a way to formulate Duplicator's response in $\exists\mathbf{PebR}_k(\mathscr{A},\mathscr{B})$. Namely, consider the coextension $f^\dagger : \mathbb{PR}_k^+\mathscr{A} \to \mathbb{PR}_k^+\mathscr{B}$ obtained from $f^+\mathbb{PR}_k\mathscr{A} \to JB$ (which exists by (3)). From $f^\dagger$, $\nu_{J\mathscr{A}}$, and $\nu_{J\mathscr{B}}$ we can define a set of partial homomorphisms for every $s \in \mathbb{P}_k\mathscr{A} \cup \{\varepsilon\}$:

$$T_s = \{\ \gamma_{s,t} \mid (u,i) \in \nu_{J\mathscr{A}}^{-1}(s) \text{ and } t = \nu_{J\mathscr{B}}(f^\dagger(u,i))\}$$

where $\nu_{J\mathscr{A}}^{-1}(s)$ is the fiber of s under $\nu_{J\mathscr{A}}$. Unpacking the definition of $\nu_{J\mathscr{A}}^{-1}(s')$, this set consists of pairs (u',i) such that $s' = u'[1,i]$ where i is the length of s'. Intuitively, each u' is a full Spoiler play of the one-sided k-pebble game which mirrors s' for the first i moves. By induction on the number rounds n of the game, we prove the following inductive hypothesis:

(IH1) There exists a pebble play $s \in \mathbb{P}_kA \cup \{\varepsilon\}$, I_n and T_n such that $I_n \subseteq \mathsf{Active}(s)$ and $T_n = T_s|_{I_n}$

(IH2) Duplicator has won the n-th round of the game, i.e. the configuration is (I_n,T_n) where T_n is a non-empty set of partial homomorphism from domain $I_n \subseteq \mathscr{A}$ to $\mathscr{B}$ and $|I_n| \leq k$.

For the base case $n = 0$, by definition, Spoiler chooses $I_0 = \varnothing$ and $T_0 = \{\lambda\}$ where λ is the unique homomorphism of type $\varnothing \to \mathscr{B}$. We let $s_0 = \varepsilon$. For the inductive step, round $n' = n+1$, assume (I,T) was the configuration and s was the pebble play of the previous round. If Spoiler chooses a shrinking move $I' \subseteq I$, then T' is just the restriction of functions $f \in T$ to the smaller domain I'. We let $s' = s$ in this case. By the construction of s and Duplicator's default response to shrinking moves, (IH1) and (IH2) hold for round n'. If Spoiler chooses a blowing move $I' \supseteq I$, then $|I'| \leq k$ and $|I| < k$. By (IH1), $I \subseteq \mathsf{Active}(s)$. Let $P \subseteq [\mathrm{k}]$ be the set of pebbles **not** pebbling elements of I in $\mathsf{Active}(s)$. Let $N = I' \backslash I$ be the set of new elements of I'. By cardinality considerations, there exists an injective function $p : N \to P$. Choose an enumeration of $N = \{a_1,\ldots,a_m\}$ and let $s' = s[(p(a_1),a_1),\ldots,(p(a_m),a_m)]$.

For $a \in I'$, either $a \in N$ or $a \in I \cap I'$. If $a \in N$, then a is pebbled by $p(a)$ and since P contains pebbles not in $\mathsf{Active}(s)$, $a \in \mathsf{Active}(s')$. If $a \in I \cap I'$, then $a \in \mathsf{Active}(s)$ and therefore pebbled by a pebble not in P, it follows that $a \in \mathsf{Active}(s')$. Let $T_n = T_{s'}$. Therefore, (IH1) holds for n'.

To verify (IH2), we need to show that $T_{s'}$ is non-empty set of partial homomorphisms. Let $\gamma_{s',t} \in T_{s'}$, then there exists a u' such that $(u',i) \in \mathbb{PR}_k\mathscr{A}$ where i is the length of s' and $t = f^\dagger(u',i)[1,i]$. Suppose $(a_1,\ldots,a_r) \in R^{\mathscr{A}}$ for r-ary relation symbol $R \in \sigma$ and $a_1,\ldots,a_r \in I'$, Since u' has s' as prefix, $a_1,\ldots,a_r$ appear in u' at some indices $j_1,\ldots,j_r \leq i$. Let $b_z = f^\dagger(u',j_z)$, then $(a_z,b_z) \in \gamma_{s',t}$ for $z = 1,\ldots,r$. By $f^\dagger$ being homomorphism, it follows that $R^{\mathscr{B}}(b_1,\ldots,b_r)$. Hence, $\gamma_{s',t}$ is a partial homomorphism. The set $T_{s'}$ is non-empty as $\gamma_{s,t} \in T_{s'}$ for $t = \nu_{J\mathscr{B}}(f^\dagger(s',i))$

(4) $\Rightarrow$ (1) This is theorem 14 of [16]. □

7 Conclusion

We have added to the growing list of Spoiler-Duplicator game comonads that unify particular model-comparison games with combinatorial invariants of relational structures. This perspective has found applications in reformulations of Rossman's homomorphism preservation theorem [13] and promises to provide new perspectives in finite model theory and descriptive complexity.

In particular, $\mathbb{PR}_k$ provides a categorical definition for pathwidth and winning strategies in Dalmau's pebble relation game. This allowed us to obtain a syntax-free characterization of equivalence in a restricted conjunction fragment of k-variable logic. In addition to this new comonad, we also obtained this comonad in novel way: by first trying to capture the combinatorial invariant, instead of internalizing the corresponding game. Moroever, some of the proofs, in this paper use techniques distinct from those used for other Spoiler-Duplicator comonads (e.g. k-pebbling section families, fibers over $\mathbb{P}_k\mathscr{A}$ to model non-determinism).

There are a few avenues for future work:

- The pebble-relation comonad $\mathbb{PR}_k$ can be seen as the "linear" or path variant of the "tree shaped" pebbling comonad. In fact, as the notion of arboreal categories discussed in [3] shows, all of the Spoiler-Duplicator game comonads discovered so far have been "tree shaped". Is there a general method for obtaining a "linear" variant of a "tree shaped" Spoiler-Duplicator game comonad?
- It was stated in [7] that constraint satisfaction problems (CSPs) with bounded pathwidth duality are in **NLOGSPACE**. All known CSPs in **NLOGSPACE** have bounded pathwidth duality. Could redefining bounded pathwidth duality in terms of $\mathbb{PR}_k$ aid in a proof of this converse or construction of a counterexample?
- In [8], CSPs in **LOGSPACE** were shown to have bounded symmetric pathwidth duality. Whereas, bounded pathwidth duality can be seen as a local property of the obstruction set, bounded symmetric pathwidth is a global property of the obstruction set. How do we formulate bounded symmetric pathwidth in terms of $\mathbb{PR}_k$?

References

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

[2] Samson Abramsky & Dan Marsden (2020): *Comonadic semantics for guarded fragments*.

[3] Samson Abramsky & Luca Reggio (2021): *Arboreal Categories: An Axiomatic Theory of Resources*.

[4] Samson Abramsky & Nihil Shah (2018): *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*, pp. 2:1–2:17, doi:10.4230/LIPIcs.CSL.2018.2.

[5] Thosten Altenkirch, James Chapman & Tarmo Uustalu (2014): *Monads need not be endofunctors*. *Logical Methods in Computer Science* 11(1), doi:10.2168/LMCS-11(1:3)2015. ArXiv: 1412.7148.

[6] Adam Ó Conghaile & Anuj Dawar (2020): *Game Comonads & Generalised Quantifiers*.

[7] Victor Dalmau (2005): *Linear Datalog and Bounded Path Duality of Relational Structures*. *Logical Methods in Computer Science* 1(1), p. 5, doi:10.2168/LMCS-1(1:5)2005. Available at `http://arxiv.org/abs/cs/0504027`. ArXiv: cs/0504027.

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

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

[10] André Joyal, Mogens Nielsen & Glynn Winskel (1996): *Bisimulation from Open Maps*. *Information and Computation* 127(2), pp. 164–185, doi:10.1006/inco.1996.0057. Available at `https://doi.org/10.1006/inco.1996.0057`.

[11] Phokion G Kolaitis & Moshe Y Vardi (1990): *On the expressive power of Datalog: tools and a case study*. In: *Proceedings of the ninth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems*, ACM, pp. 61–71.

[12] Ernest G. Manes (1976): *Algebraic Theories*. Springer New York, doi:10.1007/978-1-4612-9860-1. Available at `https://doi.org/10.1007/978-1-4612-9860-1`.

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

[14] Neil Robertson & P.D. Seymour (1983): *Graph minors. I. Excluding a forest*. *Journal of Combinatorial Theory, Series B* 35(1), pp. 39–61, doi:10.1016/0095-8956(83)90079-5. Available at `https://doi.org/10.1016/0095-8956(83)90079-5`.

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

[16] Iain A Stewart (2007): *Bounded pathwidth duality and uniform constraint satisfaction problems*. Available at `http://community.dur.ac.uk/i.a.stewart/Papers/Bddpathwidthduality.pdf`.