

Game Comonads & Generalised Quantifiers

Adam Ó Conghaile 

Department of Computer Science and Technology, University of Cambridge, UK
ac891@cl.cam.ac.uk

Anuj Dawar 

Department of Computer Science and Technology, University of Cambridge, UK
anuj.dawar@cl.cam.ac.uk

Abstract

Game comonads, introduced by Abramsky, Dawar and Wang and developed by Abramsky and Shah, give an interesting categorical semantics to some Spoiler-Duplicator games that are common in finite model theory. In particular they expose connections between one-sided and two-sided games, and parameters such as treewidth and treedepth and corresponding notions of decomposition. In the present paper, we expand the realm of game comonads to logics with generalised quantifiers. In particular, we introduce a comonad graded by two parameter $n \leq k$ such that isomorphisms in the resulting Kleisli category are exactly Duplicator winning strategies in Hella's n -bijection game with k pebbles. We define a one-sided version of this game which allows us to provide a categorical semantics for a number of logics with generalised quantifiers. We also give a novel notion of tree decomposition that emerges from the construction.

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

Keywords and phrases Logic, Finite Model Theory, Game Comonads, Generalised Quantifiers

Digital Object Identifier 10.4230/LIPIcs.CSL.2021.16

Related Version A full version of the paper is available at <https://www.arxiv.org/abs/2006.16039>.

Funding Research funded in part by EPSRC grant EP/T007257/1.

1 Introduction

Model-comparison games, such as Ehrenfeucht-Fraïssé games and pebble games, play a central role in finite model theory. Recent work by Abramsky et al. [3, 4] provides a category-theoretic view of such games which yields new insights. In particular, the *pebbling comonad* \mathbb{P}_k introduced in [3] reveals an interesting relationship between one-sided and two-sided pebble games. The morphisms in the Kleisli category associated with \mathbb{P}_k correspond exactly to winning strategies in the existential positive k -pebble game. This game was introduced by Kolaitis and Vardi [18] to study the expressive power of **Datalog**. A winning strategy for Duplicator in the game played on structures \mathcal{A} and \mathcal{B} implies that all formulas of existential positive k -variable logic true in \mathcal{A} are also true in \mathcal{B} . The game has found widespread application in the study of database query languages as well as constraint satisfaction problems. Indeed, the widely used k -local consistency algorithms for solving constraint satisfaction can be understood as computing the approximation to homomorphism given by such strategies [19]. At the same time, isomorphisms in the Kleisli category associated with \mathbb{P}_k correspond to winning strategies in the k -pebble *bijection* game. This game, introduced by Hella [16], characterises equivalence in the k -variable logic with counting. This gives a family of equivalence relations (parameterised by k) which has been widely studied as approximations of graph isomorphism. It is often called the Weisfeiler-Leman family of equivalences and has a number of characterisations in logic, algebra and combinatorics (see the discussion in [13]).



© Adam Ó Conghaile and Anuj Dawar;
licensed under Creative Commons License CC-BY

29th EACSL Annual Conference on Computer Science Logic (CSL 2021).

Editors: Christel Baier and Jean Goubault-Larrecq; Article No. 16; pp. 16:1–16:17

Leibniz International Proceedings in Informatics



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

The bijection game introduced by Hella is actually the initial level of a hierarchy of games that he defined to characterise equivalence in logics with generalised (i.e. Lindström) quantifiers. For each $n, k \in \mathbb{N}$ we have a k -pebble n -bijection game that characterises equivalence with respect to an infinitary k -variable logic with quantifiers of arity at most n . In the present paper, we introduce a graded comonad associated with this game. Our comonad $\mathbb{G}_{n,k}$ is obtained as a quotient of the comonad \mathbb{P}_k and we are able to show that isomorphisms in the associated Kleisli category correspond to winning strategies for Duplicator in the k -pebble n -bijection game. The morphisms then correspond to a new one-way game we define, which we call the k -pebble n -function game. We are able to show that this relates to a natural logic: a k -variable positive infinitary logic with n -ary *homomorphism-closed quantifiers*.

This leads us to a systematic eight-way classification of model-comparison games based on what kinds of functions Duplicator is permitted (arbitrary functions, injections, surjections or bijections) and what the partial maps in game positions are required to preserve: just atomic information or also negated atoms. We show that each of these variations correspond to preservation of formulas in a natural fragment of bounded-variable infinitary logic with n -ary Lindström quantifiers. Moreover, winning strategies in these games also correspond to natural restrictions of the morphisms in the Kleisli category of $\mathbb{G}_{n,k}$ that are well-motivated from the category-theoretic point of view.

Another key insight provided by the work of Abramsky et al. is that coalgebras in the pebbling comonad \mathbb{P}_k correspond exactly to tree decompositions of width k . Similarly, the coalgebras in the Ehrenfeucht-Fraïssé comonad introduced by Abramsky and Shah characterise the *treedepth* of structures. This motivates us to look at coalgebras in $\mathbb{G}_{n,k}$ and we show that they yield a new and natural notion of generalised tree decomposition.

In what follows, after a review of the necessary background in Section 2, we introduce the various games and logics in Section 3 and establish the relationships between them. Section 4 contains the definition of the comonad $\mathbb{G}_{n,k}$ and shows that interesting classes of morphisms in the associated Kleisli category correspond to winning strategies in the games. Section 5 defines a new class of extended tree decompositions and traversals and relates them to the coalgebras of the comonad $\mathbb{G}_{n,k}$. Proofs are omitted due to lack of space and may be found in the appendix.

2 Background

In this section we introduce notation that we use throughout the paper and give a brief overview of background we assume.

For a positive integer n , we write $[n]$ for the set $\{1, \dots, n\}$.

A tree T is a set with a partial order \leq such that for all $t \in T$, the set $\{x \mid x \leq t\}$ is linearly ordered by \leq and such that there is an element $r \in T$ called *the root* such that $r \leq t$ for all $t \in T$. If $t < t'$ in T and there is no x with $t < x < t'$, we call t' a *child* of t and t the *parent* of t' .

For X a set, we write X^* for the set of *lists* over elements of X and X^+ for the set of non-empty lists. We write the list with elements x_1, \dots, x_m in that order as $[x_1, \dots, x_m]$. For two lists $s_1, s_2 \in X^*$ we write $s_1 \cdot s_2$ for the list formed by concatenating s_1 and s_2 . For $x \in X$ and $s \in X^*$ we write $x; s$ for the list $[x] \cdot s$ and $s; x$ for the list $s \cdot [x]$. We occasionally underline the fact that $s_1 \cdot s_2, x; s$, and $s; x$ are lists by writing them enclosed in square brackets, as $[s_1 \cdot s_2], [x; s]$, and $[s; x]$.

2.1 Logics

We work with finite relational signatures and assume a fixed signature σ . Unless stated otherwise, the structures we consider are finite σ -structures. We write $\mathcal{A}, \mathcal{B}, \mathcal{C}$ etc. to denote such structures, and the corresponding roman letters A, B, C etc. to denote their universes.

We assume a standard syntax and semantics for first-order logic (as in [20]), which we denote **FO**. We write \mathcal{L}_∞ for the infinitary logic that is obtained from **FO** by allowing conjunctions and disjunctions over arbitrary sets of formulas. We write $\exists^+ \mathcal{L}_\infty$ and $\exists^+ \mathbf{FO}$ for the restriction of \mathcal{L}_∞ and **FO** to existential positive formulas, i.e. those without negations or universal quantifiers. We use natural number superscripts to denote restrictions of the logic to a fixed number of variables. We write \mathcal{C} to denote the extension of \mathcal{L}_∞ with counting quantifiers. We are mainly interested in the k -variable fragments of this logic \mathcal{C}^k .

2.2 Generalised quantifiers

We use the term *generalised quantifier* in the sense of Lindström [21]. These have been extensively studied in finite model theory (see [16, 8, 5]). In what follows, we give a brief account of the basic variant that is of interest to us here. For more on Lindström quantifiers, consult [11, Chap. 12]. We only consider quantifiers without relativisation, vectorisation or taking quotients in the interpretation.

Any isomorphism-closed class of structures K over a signature τ gives rise to a *generalised quantifier* Q_K . For a logic L , we write $L(Q_K)$ for its extension with the quantifier Q_K . We define the *arity* of the quantifier Q_K to be the maximum arity of any relation in τ . For Q_K with arity m , the formula $Q_K x_1, \dots, x_m. (\psi_R(x_1^R, \dots, x_m^R, \mathbf{z}^R))_{R \in \tau}$ where $x_i^R \in \mathbf{x}$ and $\mathbf{z}^R \subset \mathbf{z}$ is true on \mathcal{A}, \mathbf{a} if the τ -structure $\langle \mathcal{A}, (\psi_R(\cdot, \mathbf{a}^R))_{R \in \tau} \rangle$ is in K . We write $\mathcal{L}_\infty^k(Q_n)$ for the extension of \mathcal{L}_∞^k with *all* quantifiers of arity n . This is only of interest when $n \leq k$. Kolaitis and Väänänen [17] showed that $\mathcal{L}_\infty^k(Q_1)$ is equivalent to \mathcal{C}^k . However, allowing quantifiers of higher arity gives logics of considerably more expressive power. In particular, if σ is a signature with all relations of arity at most n , then *any* property of σ -structures is expressible in $\mathcal{L}_\infty^n(Q_n)$. Thus, all properties of graphs, for instance, are expressible in $\mathcal{L}_\infty^2(Q_2)$.

2.3 Games

For a pair of structures \mathcal{A} and \mathcal{B} and a logic L , we write $\mathcal{A} \Rightarrow_L \mathcal{B}$ to denote that every sentence of L that is true in \mathcal{A} is also true in \mathcal{B} . When the logic is closed under negation, as is the case with **FO** and \mathcal{L}_∞ , for instance, $\mathcal{A} \Rightarrow_L \mathcal{B}$ implies $\mathcal{B} \Rightarrow_L \mathcal{A}$. In this case, we have an equivalence relation between structures and we write $\mathcal{A} \equiv_L \mathcal{B}$. When \mathcal{A} and \mathcal{B} are finite structures, $\mathcal{A} \Rightarrow_{\mathbf{FO}} \mathcal{B}$ implies $\mathcal{A} \Rightarrow_{\mathcal{L}_\infty} \mathcal{B}$, and the same holds for the k -variable fragments of these logics (see [10]).

The relations \Rightarrow_L are often characterised in terms of games which we generically call *Spoiler-Duplicator* games. For instance, the existential-positive k -pebble game introduced by Kolaitis and Vardi [18], which we denote $\exists \mathbf{Peb}_k$, characterises the relation $\Rightarrow_{\exists^+ \mathcal{L}_\infty^k}$. In this game, Spoiler and Duplicator each has a collection of k pebbles indexed $1, \dots, k$. In each round Spoiler places one of its pebbles on an element of \mathcal{A} and Duplicator responds by placing its corresponding pebble (i.e. the one of the same index) on an element of \mathcal{B} . Note the game can go on for more than k rounds and pebbles can be repositioned throughout. If the partial map taking the element of \mathcal{A} on which Spoiler's pebble i sits to the element of \mathcal{B} on which Duplicator's pebble i is, fails to be a partial homomorphism, then Spoiler has won the game. Duplicator wins by playing forever without losing. We get a game characterising

$\equiv_{\mathcal{L}^\infty}^k$ if (i) Spoiler is allowed to choose, at each move, on which of the two structures it places a pebble and Duplicator is required to respond in the other structure; and (ii) Duplicator is required to ensure that the pebbled positions form a partial *isomorphism*.

Hella [16] introduced a *bijection game* which characterises the equivalence $\equiv_{\mathcal{L}^k}$. We write $\mathbf{Bij}_k(\mathcal{A}, \mathcal{B})$ for the bijection game played on \mathcal{A} and \mathcal{B} . At each move, Spoiler chooses an index $i \in [k]$ and Duplicator is required to respond with a bijection $f : A \rightarrow B$. Spoiler then chooses an element $a \in A$ and pebbles indexed i are placed on a and $f(a)$. If the partial map defined by the pebbled positions is not a partial isomorphism, then Spoiler has won. Duplicator wins by playing forever without losing.

In Hella's original work, the bijection games appear as a special case of the *n-bijjective k-pebble game*, which we denote $\mathbf{Bij}_n^k(\mathcal{A}, \mathcal{B})$ when played on structures \mathcal{A} and \mathcal{B} . This characterises the equivalence relation $\equiv_{\mathcal{L}^\infty(\mathcal{Q}_n)}^k$. Once again, we have a set of k pebbles associated with each of the structures \mathcal{A} and \mathcal{B} and indexed by $[k]$. At each move, Duplicator is required to give a bijection $f : A \rightarrow B$ and Spoiler chooses a set of up to n pebble indices $p_1, \dots, p_n \in [k]$ and moves the corresponding indices to elements $a_1, \dots, a_n \in A$ and $f(a_1), \dots, f(a_n)$ in B . If the partial map defined by the pebbled positions is not a partial isomorphism, then Spoiler has won. Duplicator wins by playing forever without losing. Note, in particular, that for Duplicator to have a winning strategy it is necessary that the reducts of \mathcal{A} and \mathcal{B} to relations of arity at most n are isomorphic. For example, on graphs Spoiler wins any game on non-isomorphic graphs with $n, k \geq 2$.

2.4 Comonads

We assume that the reader is familiar with basic definitions from category theory, in particular the notions of category, functor and natural transformation. For a finite signature σ , we are interested in the category $\mathcal{R}(\sigma)$ of relational structures over σ . The objects of the category are such structures and the maps are homomorphisms between structures.

Comonads on a category \mathcal{C} are triples (T, ϵ, δ) where T is an endofunctor on \mathcal{C} and ϵ and δ are natural transformations of type $T \rightarrow 1$ and $TT \rightarrow T$ respectively, satisfying certain comonad laws. The Kleisli category $\mathcal{K}(T)$ is the category with the same objects as \mathcal{C} where $\mathcal{K}(T)$ -morphisms are \mathcal{C} -morphisms of type $TA \rightarrow B$. Composition is defined with the help of δ . A T -coalgebra on an object $A \in \mathcal{C}$ is a \mathcal{C} -map $\alpha : A \rightarrow TA$ such that $\epsilon_A \circ \alpha = 1_A$.

Abramsky et al. [3] describe the construction of a comonad \mathbb{P}_k , graded by k , on the category $\mathcal{R}(\sigma)$ which exposes an interesting relationship between the games $\exists\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ and $\mathbf{Bij}_k(\mathcal{A}, \mathcal{B})$. Specifically, this construction shows that Duplicator winning strategies in the latter are exactly the isomorphisms in a category in which the morphisms are winning strategies in the former.

For any \mathcal{A} , $\mathbb{P}_k\mathcal{A}$ is an infinite structure (even when \mathcal{A} is finite) with universe $(A \times [k])^+$. The counit $\epsilon_{\mathcal{A}}$ takes a sequence $[(a_1, p_1), \dots, (a_m, p_m)]$ to a_m , i.e. the first component of the last element of the sequence. The comultiplication $\delta_{\mathcal{A}}$ takes a sequence $[(a_1, p_1), \dots, (a_m, p_m)]$ to the sequence $[(s_1, p_1), \dots, (s_m, p_m)]$ where $s_i = [(a_1, p_1), \dots, (a_i, p_i)]$. The relations are defined so that $(s_1, \dots, s_r) \in R^{\mathbb{P}_k\mathcal{A}}$ if, and only if, the s_i are all comparable in the prefix order of sequences, $R^{\mathcal{A}}(\epsilon_{\mathcal{A}}(s_1), \dots, \epsilon_{\mathcal{A}}(s_r))$ and whenever s_i is a prefix of s_j and ends with the pair (a, p) , there is no prefix of s_j properly extending s_i which ends with (a', p) for any $a' \in A$.

It is convenient to consider structures over a signature $\sigma \cup \{I\}$ where I is a new binary relation symbol. An I -structure is a structure over this signature which interprets I as the identity relation. Note that even when \mathcal{A} is an I -structure, $\mathbb{P}_k\mathcal{A}$ is not one. The key results from [3] relating the comonad with pebble games can now be stated as establishing a

precise translation between (i) $\mathcal{K}(\mathbb{P}_k)$ -morphisms $\mathcal{A} \longrightarrow \mathcal{B}$ for I -structures \mathcal{A} and \mathcal{B} ; and (ii) winning strategies for Duplicator in $\exists\text{Peb}_k(\mathcal{A}, \mathcal{B})$; and similarly a precise translation between (i) isomorphisms in $\mathcal{K}(\mathbb{P}_k)$ between \mathcal{A} and \mathcal{B} for I -structures \mathcal{A} and \mathcal{B} ; and (ii) winning strategies for Duplicator in $\mathbf{Bij}_k(\mathcal{A}, \mathcal{B})$.

A key result from the construction of the comonad \mathbb{P}_k is the relationship between the coalgebras of this comonad and tree decompositions. In particular, it is shown that a structure \mathcal{A} has a \mathbb{P}_k -coalgebra if, and only if, the treewidth of \mathcal{A} is at most $k - 1$. This relationship between coalgebras and tree decompositions is established through a definition of a *tree traversal* which we review in Section 5 below.

3 Games and Logic with Generalised Quantifiers

Hella's n -bijjective k -pebble game, \mathbf{Bij}_n^k is a model-comparison game which captures equivalence of structures over the logic $\mathcal{L}_\infty^k(\mathcal{Q}_n)$, i.e. k -variable infinitary logic where the allowed quantifiers are all generalised quantifiers with arity $\leq n$. This game generalises the bijection game \mathbf{Bij}_k which captures equivalence over \mathcal{C}^k , k -variable infinitary logic with counting quantifiers (which is equivalent to $\mathcal{L}_\infty^k(\mathcal{Q}_1)$ as shown by Kolaitis and Väänänen[17]). In this section, we introduce a family of games which relax the rules of \mathbf{Bij}_n^k and show their correspondence to different fragments of $\mathcal{L}_\infty^k(\mathcal{Q}_n)$. In particular, we introduce a “one-way” version of \mathbf{Bij}_n^k which is crucial to our construction of a generalised version of the \mathbb{P}_k comonad.

3.1 Relaxing \mathbf{Bij}_n^k

Recall that each round of $\mathbf{Bij}_n^k(\mathcal{A}, \mathcal{B})$ involves Duplicator selecting a bijection $f : A \rightarrow B$ and ends with a test of whether for the pebbled positions $(a_i, b_i)_{i \in [k]}$ it is the case that, for any $\{i_1, \dots, i_r\} \subset [k]$, $(a_{i_1}, \dots, a_{i_r}) \in R^{\mathcal{A}} \iff (b_{i_1}, \dots, b_{i_r}) \in R^{\mathcal{B}}$ where Duplicator loses if the test is failed. For the rest of the round, Spoiler rearranges up to n pebbles on \mathcal{A} with the corresponding pebbles on \mathcal{B} moved according to f .

So, to create from \mathbf{Bij}_n^k a “one-way” game from \mathcal{A} to \mathcal{B} we need to relax the condition that f be a bijection and the \iff in the final test. The following definition captures the most basic such relaxation:

► **Definition 1.** For two relational structures \mathcal{A}, \mathcal{B} , the positive k -pebble n -function game, $+\text{Fun}_n^k(\mathcal{A}, \mathcal{B})$ is played by Spoiler and Duplicator. Prior to the j th round the position consists of partial maps $\pi_{j-1}^a : [k] \rightarrow A$ and $\pi_{j-1}^b : [k] \rightarrow B$. In Round j

- Duplicator provides a function $h_j : A \rightarrow B$ such that for each $i \in [k]$, $h_j(\pi_{j-1}^a(i)) = \pi_{j-1}^b(i)$.
- Spoiler picks up to n distinct pebbles, i.e. elements $p_1, \dots, p_m \in [k]$ ($m \leq n$) and m elements $x_1, \dots, x_m \in A$.
- The updated position is given by $\pi_j^a(p_l) = x_l$ and $\pi_j^b(p_l) = h_j(x_l)$ for $l \in [m]$; and $\pi_j^a(i) = \pi_{j-1}^a(i)$ and $\pi_j^b(i) = \pi_{j-1}^b(i)$ for $i \notin \{p_1, \dots, p_m\}$.
- Spoiler has won the game if there is some $R \in \sigma$ and $(i_1, \dots, i_r) \in [k]^r$ with $(\pi_j^a(i_1), \dots, \pi_j^a(i_r)) \in R^{\mathcal{A}}$ but $(\pi_j^b(i_1), \dots, \pi_j^b(i_r)) \notin R^{\mathcal{B}}$.

Duplicator wins by preventing Spoiler from winning.

As this game is to serve as the appropriate one-way game for \mathbf{Bij}_n^k , it is worth asking why this game is a reasonable generalisation of $\exists\text{Peb}_k$ (the one-way game for \mathbf{Bij}_k). The answer comes in recalling Abramsky et al.'s presentation of a (deterministic) strategy for Duplicator in $\exists\text{Peb}_k(\mathcal{A}, \mathcal{B})$ as a collection of *branch maps* $\phi_{s,i} : A \rightarrow B$ for each $s \in (A \times [k])^*$, a history

of Spoiler moves and $i \in [k]$ a pebble index. These branch maps tell us how Duplicator would respond to Spoiler moving pebble i to any element in A given the moves s that Spoiler has played in preceding rounds. The functions h_j in Definition 1 serve as just such branch maps.


In addition to this game, we now define some other relaxations of \mathbf{Bij}_n^k which are important. In particular we define the following *positive* games by retaining that the pebbled position need only preserve positive atoms at the end of each round but varying the condition on f .

► **Definition 2.** For two relational structures \mathcal{A}, \mathcal{B} , the positive k -pebble n -injection (resp. surjection, bijection) game, $+\mathbf{Inj}_n^k(\mathcal{A}, \mathcal{B})$ (resp. $+\mathbf{Surj}_n^k(\mathcal{A}, \mathcal{B})$, $+\mathbf{Bij}_n^k(\mathcal{A}, \mathcal{B})$) is played by Spoiler and Duplicator. Prior to the j th round the position consists of partial maps $\pi_{j-1}^a : [k] \rightarrow A$ and $\pi_{j-1}^b : [k] \rightarrow B$. In Round j

- Duplicator provides an injection (resp. a surjection, bijection) $h_j : A \rightarrow B$ such that for each $i \in [k]$, $h_j(\pi_{j-1}^a(i)) = \pi_{j-1}^b(i)$.
- Spoiler picks up to n distinct pebbles, i.e. elements $p_1, \dots, p_m \in [k]$ ($m \leq n$) and m elements $x_1, \dots, x_m \in A$.
- The updated position is given by $\pi_j^a(p_l) = x_l$ and $\pi_j^b(p_l) = h_j(x_l)$ for $l \in [m]$; and $\pi_j^a(i) = \pi_{j-1}^a(i)$ and $\pi_j^b(i) = \pi_{j-1}^b(i)$ for $i \notin \{p_1, \dots, p_m\}$.
- Spoiler has won the game if there is some $R \in \sigma$ and $(i_1, \dots, i_r) \in [k]^r$ with $(\pi_j^a(i_1), \dots, \pi_j^a(i_r)) \in R^A$ but $(\pi_j^b(i_1), \dots, \pi_j^b(i_r)) \notin R^B$.

Duplicator wins by preventing Spoiler from winning.

Strengthening the test condition in each round so that Spoiler wins if there is some $R \in \sigma$ and $(i_1, \dots, i_r) \in [k]^r$ with $(\pi_j^a(i_1), \dots, \pi_j^a(i_r)) \in R^A$ if, and only if, $(\pi_j^b(i_1), \dots, \pi_j^b(i_r)) \notin R^B$, we get the definitions for the games \mathbf{Fun}_n^k , \mathbf{Inj}_n^k , \mathbf{Surj}_n^k and \mathbf{Bij}_n^k where the latter is precisely the n -bijjective k -pebble game of Hella. We recap the poset of the games we've just defined ordered by strengthening of the rules/restrictions on Duplicator in the leftmost Hasse diagram in Figure 1. Here a game \mathcal{G} is above \mathcal{G}' if a Duplicator winning strategy in \mathcal{G} is also one in \mathcal{G}' .



■ **Figure 1** Hasse Diagrams of Games and Logics with Labels For Reference.

3.2 Logics with generalised quantifiers

In Section 2, we introduce for each $n, k \in \mathbb{N}$ the logics, $\mathcal{L}_\infty^k(\mathcal{Q}_n)$ as the infinitary logic extended with all generalised quantifiers of arity n . For $n = 1$ this logic leads somewhat of a double life. Kolaitis and Väänänen [17] show that this logic is equivalent to \mathcal{C}^k , the infinitary logic with counting quantifiers and at most k variables.

In this section we explore fragments of $\mathcal{L}_\infty^k(Q_n)$ defined by restricted classes of generalised quantifiers, which we introduce next.

► **Definition 3.** *A class of σ -structures K is homomorphism-closed if for all homomorphisms $f : \mathcal{A} \rightarrow \mathcal{B}$ we have that $\mathcal{A} \in K \implies \mathcal{B} \in K$. Similarly, we say K is injection-closed (resp. surjection-closed, bijection-closed) if for all injective homomorphisms (resp. surjective, bijective homomorphisms) $f : \mathcal{A} \rightarrow \mathcal{B}$, we have $\mathcal{A} \in K \implies \mathcal{B} \in K$.*

We write \mathcal{Q}_n^h for the class of all generalised quantifiers Q_K of arity n where K is homomorphism-closed. Similarly, we write \mathcal{Q}_n^i , \mathcal{Q}_n^s and \mathcal{Q}_n^b for the collections of n -ary quantifiers based on injection-closed, surjection-closed and bijection-closed classes.

In order to define logics which incorporate these restricted classes of quantifiers, we first define a base logic without quantifiers or negation.

► **Definition 4.** *We denote by $+\mathcal{L}^k$ the class of positive infinitary k -variable quantifier-free formulas. That means the k -variable fragment of the class of formulas $+\mathcal{L}[\sigma]$ (for any signature σ), given by the grammar*

$$\phi ::= R(x_1, \dots, x_m) \mid \bigwedge_{\mathcal{I}} \phi \mid \bigvee_{\mathcal{J}} \phi$$

for $R \in \sigma$. We use \mathcal{L}^k to denote a similar class of formulas but with negation permitted on atoms.

This basic set of formulas can be extended into a logic by adding some set of quantifiers as described here:

► **Definition 5.** *For \mathcal{Q} some collection of generalised quantifiers, we denote by $+\mathcal{L}^k(\mathcal{Q})$ the smallest extension of $+\mathcal{L}^k$ to include the all quantifiers $Q_K \in \mathcal{Q}$, closed under quantification and ordinary boolean operations (excluding negation). $\mathcal{L}^k(\mathcal{Q})$ is the same logic but with negation on atoms. Note that $\exists^+ \mathcal{L}_\infty^k \equiv +\mathcal{L}^k(\exists)$ and, as we can always push negation down to the level of atoms in \mathcal{L}_∞^k , $\mathcal{L}_\infty^k \equiv \mathcal{L}^k(\exists, \forall)$.*

With this definition we are ready to introduce our logics. These are $\mathcal{L}^k(Q_n^h)$, $\mathcal{L}^k(Q_n^i)$, $\mathcal{L}^k(Q_n^s)$ and $\mathcal{L}^k(Q_n^b)$ and their positive counterparts $+\mathcal{L}^k(Q_n^h)$, $+\mathcal{L}^k(Q_n^i)$, $+\mathcal{L}^k(Q_n^s)$ and $+\mathcal{L}^k(Q_n^b)$. The obvious inclusion relationships between these logics are given by the middle Hasse diagram in Figure 1. As we shall see, these logics are governed exactly by the games pictured in the leftmost diagram in Figure 1.

Here we highlight two results relating this family of logics with more familiar infinitary logics. These results classify the two extreme logics in the diagram from Figure 1, namely $\mathcal{L}^k(Q_n^b)$ and $+\mathcal{L}^k(Q_n^h)$.

► **Lemma 6.** *For all $n, k \in \mathbb{N}$, $\mathcal{L}^k(Q_n^b) \equiv \mathcal{L}_\infty^k(Q_n)$.*

► **Lemma 7.** $+\mathcal{L}^k(Q_1^h) \equiv +\mathcal{L}^k(\exists)$

3.3 Games and logics correspond

So far we have introduced a series of games and logics which are all variations on Hella's n -bijection k -pebble game, \mathbf{Bij}_n^k , and the corresponding logic $\mathcal{L}_\infty^k(Q_n)$. Here we show that these games and logics match up in the way as one would expect looking at the respective refinement posets in Figures 1.

In order to present the proof of this in a uniform fashion, we label the corners of these cubes by three parameters $x_i, x_s, x_n \in \{0, 1\}$, standing for *injection*, *surjection* and *negated atoms* respectively. This can be seen in Figure 1.

Now we define the aliases of each of the games which modify \mathbf{Fun}_n^k as follows, with the games defined lining up with the games defined in Section 3.1.

► **Definition 8.** For two σ -structures \mathcal{A} and \mathcal{B} , the game $(x_i, x_s, x_n)\text{-Fun}_n^k(\mathcal{A}, \mathcal{B})$ is played by Spoiler and Duplicator in the same fashion as the game $\mathbf{Fun}_n^k(\mathcal{A}, \mathcal{B})$ with the following additional rules:

1. When Duplicator provides a function $f : A \rightarrow B$ at the beginning of a round, f is required to be injective if $x_i = 1$ and surjective if $x_s = 1$.
2. If $x_n = 1$, Spoiler wins at move j if the partial map taking $\pi_j^a(i)$ to $\pi_j^b(i)$ fails to preserve negated atoms as well as atoms.

Similarly, we define parameterised aliases for the logics introduced in Section 3.2. To lighten our notational burden, we use $\mathcal{H}^{n,k}$ to denote the logic $+\mathcal{L}^k(Q_n^h)$ throughout this section.

► **Definition 9.** We define $\mathcal{H}_x^{n,k}$ to be the logic $\mathcal{H}^{n,k}$ extended by

1. all n -ary generalised quantifiers closed by all homomorphisms which are injective, if $x_i = 1$; and surjective, if $x_s = 1$.
2. if $x_n = 1$, negation on atoms.

For example, $\mathcal{H}_{001}^{n,k}$ extends $\mathcal{H}^{n,k}$ with negation on atoms but contains no additional quantifiers as all n -ary quantifiers closed under homomorphisms are already in $\mathcal{H}^{n,k}$. On the other hand, $\mathcal{H}_{110}^{n,k}$ does not allow negation on atoms but allows all quantifiers that are closed under bijective homomorphisms.

Now to prove the desired correspondence between $\mathbf{x-Fun}_n^k$ and $\mathcal{H}_x^{n,k}$, we adapt a proof from Hella[16] to work for this parameterised set of games. For this we need the language of *back-and-forth* systems which Hella uses as an explicit representation of a Duplicator winning strategy. We provide the appropriate generalised definition here:

► **Definition 10.** Let $\mathbf{Part}_{x_n}^k(\mathcal{A}, \mathcal{B})$ be the set of all partial functions $A \rightarrow B$ which preserve atoms (i.e. are partial homomorphisms) and, if $x_n = 1$ additionally preserve negated atoms.

A set $\mathcal{S} \subset \mathbf{Part}_{x_n}^k(\mathcal{A}, \mathcal{B})$ is a *back-and-forth system* for the game $(x_i, x_s, x_n)\text{-Fun}_n^k(\mathcal{A}, \mathcal{B})$ if it satisfies the following properties:

- **Closure under subfunction:** If $f \in \mathcal{S}$ then $g \in \mathcal{S}$ for any $g \subset f$
- **(x_i, x_s) -forth property** For any f in \mathcal{S} s.t. $|f| < k$, there exists a function $\phi_f : A \rightarrow B$, which is injective if $x_i = 1$ and surjective if $x_s = 1$ s.t. for every $C \subset \text{dom}(f), D \subset A$ with $|D| \leq n$ we have $(f \upharpoonright C) \cup (\phi_f \upharpoonright D) \in \mathcal{S}$.

As this definition is essentially an unravelling of a Duplicator winning strategy for the game $(x_i, x_s, x_n)\text{-Fun}_n^k(\mathcal{A}, \mathcal{B})$ we see that

► **Lemma 11.** There is a back-and-forth system \mathcal{S} containing the empty partial homomorphism \emptyset if, and only if, Duplicator has a winning strategy for the game $(x_i, x_s, x_n)\text{-Fun}_n^k(\mathcal{A}, \mathcal{B})$

Following Hella, we define the *canonical* back-and-forth system for a game as follows:

► **Definition 12.** The canonical back-and-forth system for $(x_i, x_s, x_n)\text{-Fun}_n^k(\mathcal{A}, \mathcal{B})$ is denoted $I_x^{n,k}(\mathcal{A}, \mathcal{B})$ and is given by the intersection $\bigcap_m I_x^{n,k,m}(\mathcal{A}, \mathcal{B})$, whose conjuncts are defined inductively by setting $I_x^{n,k,0}(\mathcal{A}, \mathcal{B}) := \mathbf{Part}_{x_n}^k(\mathcal{A}, \mathcal{B})$ and letting $I_x^{n,k,m+1}(\mathcal{A}, \mathcal{B})$ be the set of $\rho \in I_x^{n,k,m}(\mathcal{A}, \mathcal{B})$ such that ρ satisfies the (x_i, x_s) -forth condition with respect to the set $I_x^{n,k,m}(\mathcal{A}, \mathcal{B})$

It is not difficult to see that for any back-and-forth system \mathcal{S} for $\mathbf{x}\text{-Fun}_n^k(\mathcal{A}, \mathcal{B})$ we have $\mathcal{S} \subset I_{\mathbf{x}}^{n,k}(\mathcal{A}, \mathcal{B})$. This means that there is a winning strategy for Duplicator in the game $\mathbf{x}\text{-Fun}_n^k(\mathcal{A}, \mathcal{B})$ if, and only if, $I_{\mathbf{x}}^{n,k}(\mathcal{A}, \mathcal{B})$ is not empty.

To complete the vocabulary needed to emulate Hella's proof in this setting we introduce the following generalisations of Hella's definitions.

► **Definition 13.** Denote by $J_{\mathbf{x}}^{n,k}(\mathcal{A}, \mathcal{B})$ the set of all $\rho \in \text{Part}_{x_n}^k(\mathcal{A}, \mathcal{B})$ which preserve the validity of all $\mathcal{H}_{\mathbf{x}}^{n,k}$ formulas in elements of $\text{dom}(\rho)$. Let $\exists^+ \mathbf{FO}_{\mathbf{x}}^{n,k}$ denote the fragment of $\mathcal{H}_{\mathbf{x}}^{n,k}$ with only finitary conjunctions and disjunctions. Denote by $K_{\mathbf{x}}^{n,k}(\mathcal{A}, \mathcal{B})$ the set of all $\rho \in \text{Part}_{x_n}^k(\mathcal{A}, \mathcal{B})$ which preserve the validity of all $\exists^+ \mathbf{FO}_{\mathbf{x}}^{n,k}$ formulas in elements of $\text{dom}(\rho)$.

An adaptation of Hella's argument yields the following Lemma:

► **Lemma 14.** For \mathcal{A}, \mathcal{B} finite relational structures, $I_{\mathbf{x}}^{n,k}(\mathcal{A}, \mathcal{B}) = J_{\mathbf{x}}^{n,k}(\mathcal{A}, \mathcal{B}) = K_{\mathbf{x}}^{n,k}(\mathcal{A}, \mathcal{B})$

We conclude this section by showing the desired correspondence for the whole family of games and logics we have introduced.

► **Theorem 15.** For $\mathbf{x} \in \{0, 1\}^3$ and all $n, k \in \mathbb{N}$ the following are equivalent:

- Duplicator has a winning strategy for $\mathbf{x}\text{-Fun}_n^k(\mathcal{A}, \mathcal{B})$
- $\mathcal{A} \rightleftharpoons_{\mathcal{H}_{\mathbf{x}}^{n,k}} \mathcal{B}$
- $\mathcal{A} \rightleftharpoons_{\exists^+ \mathbf{FO}_{\mathbf{x}}^{n,k}} \mathcal{B}$

The case of $n = 1$ for this correspondence is particularly interesting as we can show that unary injection-closed and surjection-closed quantifiers are generated by all counting quantifiers and the quantifiers $\{\exists, \forall\}$ respectively.

4 The Comonad and Kleisli Category

In this section, we show how to construct a game comonad $\mathbb{G}_{n,k}$ which captures the strategies of $+\mathbf{Fun}_n^k$ in the same way that \mathbb{P}_k captures the strategies of $\exists \mathbf{Peb}_k$. We do this using a new technique for constructing new game comonads from old based on strategy translation. We then show that different types of morphism in the Kleisli category of this new comonad correspond to Duplicator strategies for the games introduced in Section 3.

4.1 Translating between games

The pebbling comonad is obtained by defining a structure $\mathbb{P}_k \mathcal{A}$ for each \mathcal{A} whose universe consists of (non-empty) lists in $(A \times [k])^*$ which we think of as sequences of moves by Spoiler in a game $\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$, with \mathcal{B} unspecified. With this in mind, we call a sequence in $(A \times [k])^*$ a k -history (allowing the empty sequence). In contrast, a move in the game $+\mathbf{Fun}_n^k(\mathcal{A}, \mathcal{B})$ involves Spoiler moving up to n pebbles and therefore a history of Spoiler moves is a sequence in $((A \times [k])^{\leq n})^*$. We call such a sequence an n, k -history. With this set-up, (deterministic) strategies are given by functions $((A \times [k])^* \times [k]) \rightarrow (A \rightarrow B)$ for $\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ and $((A \times [k])^{\leq n})^* \rightarrow (A \rightarrow B)$ for $+\mathbf{Fun}_n^k(\mathcal{A}, \mathcal{B})$.

A winning strategy for Duplicator in $+\mathbf{Fun}_n^k(\mathcal{A}, \mathcal{B})$ can always be translated into one in $\mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$. We aim now to establish conditions when a translation can be made in the reverse direction. For this, it is useful to establish some machinery.

There is a natural *flattening* operation that takes n, k -histories to k -histories. We denote the operation by F , so $F([s_1, s_2, \dots, s_m]) = s_1 \cdot s_2 \cdots s_m$, where $s_1, \dots, s_m \in (A \times [k])^{\leq n}$. Of course, the function F is not injective and has no inverse. It is worth, however, considering

16:10 Game Comonads & Generalised Quantifiers

functions G from k -histories to n, k -histories that are inverse to F in the sense that $F(G(t)) = t$. One obvious such function takes a k -history s_1, \dots, s_m to the n, k -history $[[s_1], \dots, [s_m]]$, i.e. the sequence of one-element sequences. This is, in some sense, minimal in that imposes the minimal amount of structure on $G(t)$. We are interested in a *maximal* such function. For this, recall that the sequences in $(A \times [k])^{\leq n}$ that form the elements of an n, k -history have length at most n and do not have a repeated index from $[k]$. We aim to break a k -history t into maximal such blocks. This leads us to the following definition.

► **Definition 16.** A list $s \in (A \times [k])^*$ is called *basic* if it contains fewer than or equal to n pairs and the pebble indices are all distinct.

The n -structure function $S_n : (A \times [k])^* \rightarrow ((A \times [k])^{\leq n})^*$ is defined recursively as follows:

- $S_n(s) = [s]$ if s is basic
- otherwise, $S_n(s) = [a]; S_n(t)$ where $s = a \cdot t$ such that a is the largest basic prefix of s .

It is immediate from the definition that $F(S_n(t)) = t$. It is useful to characterise the range of the function S_n , which we do through the following definition.

► **Definition 17.** An n, k -history t is *structured* if whenever s and s' are successive elements of t , then either s has length exactly n or s' begins with a pair (a, p) such that p occurs in s .

It is immediate from the definitions that $S_n(s)$ is structured for all k -histories s and that an n, k -history is structured if, and only if, $S_n(F(s)) = s$.

We are now ready to characterise those Duplicator winning strategies for $\exists \mathbf{Peb}_k$ that can be lifted to $+\mathbf{Fun}_n^k$. First, we define a function that lifts a position in $\exists \mathbf{Peb}_k$ that Duplicator must respond to, i.e. a pair (s, p) where s is a k -history and p a pebble index, to a position in $+\mathbf{Fun}_n^k$, i.e. an n, k -history.

► **Definition 18.** Suppose s is a k -history and s' is the last basic list in $S_n(s)$, so $S_n(s) = t; [s']$. Let $p \in [k]$ be a pebble index.

Define the n -structuring $\alpha_n(s, p)$ of (s, p) by

$$\alpha_n(s, p) = \begin{cases} t; [s'] & \text{if } |s'| = n \text{ or } p \text{ occurs in } s' \\ t & \text{otherwise.} \end{cases}$$

► **Definition 19.** Say that a Duplicator strategy $\Psi : ((A \times [k])^* \times [k]) \rightarrow (A \rightarrow B)$ in $\exists \mathbf{Peb}_k$ is n -consistent if for all k -histories s and s' and all pebble indices p and p' :

$$\alpha_n(s, p) = \alpha_n(s', p') \quad \Rightarrow \quad \Psi(s, p) = \Psi(s', p').$$

Intuitively, an n -consistent Duplicator strategy in the game $\exists \mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ is one where Duplicator plays the same function in all moves that could be part of the same Spoiler move in the game $+\mathbf{Fun}_n^k(\mathcal{A}, \mathcal{B})$. We are then ready to prove the main result of this subsection.

► **Lemma 20.** Duplicator has an n -consistent winning strategy in $\exists \mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$ if, and only if, it has a winning strategy in $+\mathbf{Fun}_n^k(\mathcal{A}, \mathcal{B})$.

4.2 Lifting the comonad \mathbb{P}_k to $\mathbb{G}_{n,k}$

Central to Abramsky et al.'s construction of the pebbling comonad is the observation that for I -structures (defined in Section 2), maps in the Kleisli category $\mathcal{K}(\mathbb{P}_k)$ correspond to Duplicator winning strategies in $\exists \mathbf{Peb}_k(\mathcal{A}, \mathcal{B})$.

► **Lemma 21** ([3]). *For \mathcal{A} and \mathcal{B} I -structures over the signature σ , there is a homomorphism $\mathbb{P}_k\mathcal{A} \rightarrow \mathcal{B}$ if and only if there is a (deterministic) winning strategy for Duplicator in the game $\exists\text{Peb}_k(\mathcal{A}, \mathcal{B})$*

The relation to strategies is clear in the context of elements $s \in \mathbb{P}_k\mathcal{A}$ representing histories of Spoiler moves up to and including the current move in the game $\exists\text{Peb}_k(\mathcal{A}, \mathcal{B})$. The relational structure given to this set by Abramsky, Dawar and Wang ensures that pebbled positions preserve relations in σ , while the caveat here about I -structures is a technicality to ensure that the pebbled positions when “playing” according to a map f all define partial homomorphisms, in particular they give well defined partial maps from A to B .

As we saw in Lemma 20 a Duplicator winning strategy in $+\text{Fun}_n^k(\mathcal{A}, \mathcal{B})$ is given by an n -consistent strategy in $\exists\text{Peb}_k(\mathcal{A}, \mathcal{B})$. The n -consistency condition can be seen as saying that the corresponding map $f : \mathbb{P}_k\mathcal{A} \rightarrow \mathcal{B}$ must, on certain “equivalent” elements of $\mathbb{P}_k\mathcal{A}$ give the same value. We can formally define the equivalence relation as follows.

► **Definition 22.** *For $n \in \mathbb{N}$ and \mathcal{A} a relational structure. Define \approx_n on the universe of $\mathbb{P}_k\mathcal{A}$ as follows:*

$$[s; (a, i)] \approx_n [t; (b, j)] \iff a = b \text{ and } \alpha_n((s, i)) = \alpha_n((t, j))$$

In general, for any structured n, k -history t , we write $[t|a]$ to denote the \approx_n -equivalence class of an element $[s; (a, i)] \in \mathbb{P}_k\mathcal{A}$ with $\alpha_n(s, i) = t$.

This allows us to define the main construction of this section as a quotient of the relational structure $\mathbb{P}_k\mathcal{A}$. Note that the relation \approx_n is not a congruence of this structure, so there is not a canonical quotient. This is because don't have that $\mathbf{a} \approx_n \mathbf{b}$ does not imply $\mathbf{a} \in R^A \iff \mathbf{b} \in R^A$. Given an arbitrary equivalence relation \sim over a relational structure \mathcal{M} , there are two standard ways to define relations in a quotient \mathcal{M}/\sim . We could say that a tuple (c_1, \dots, c_r) of equivalence classes is in a relation $R^{\mathcal{M}/\sim}$ if, and only if, every choice of representatives is in $R^{\mathcal{M}}$ or if some choice of representatives is in $R^{\mathcal{M}}$. The latter definition has the advantage that the quotient map from \mathcal{M} to \mathcal{M}/\sim is a homomorphism.

► **Definition 23.** *For $n, k \in \mathbb{N}$, $k \geq n$ and σ a relational signature, we define the functor $\mathbb{G}_{n,k} : \mathcal{R}(\sigma) \rightarrow \mathcal{R}(\sigma)$ by:*

- **On objects** $\mathbb{G}_{n,k}\mathcal{A} := \mathbb{P}_k\mathcal{A}/\approx_n$.
- **On morphisms** $\mathbb{P}_k f / \approx_n$ is well-defined as $\mathbb{P}_k f$ only changes the elements not the pebble indices.

Writing $q_n : \mathbb{P}_k\mathcal{A} \rightarrow \mathbb{G}_{n,k}\mathcal{A}$ for the quotient map enables us to establish the following useful property.

► **Observation 24.** *$f : \mathbb{G}_{n,k}\mathcal{A} \rightarrow \mathcal{B}$ is a homomorphism if, and only if, $f \circ q_n : \mathbb{P}_k\mathcal{A} \rightarrow \mathcal{B}$ is a homomorphism.*

Combining this with Lemma 20, we have the appropriate generalisation of Lemma 21.

► **Lemma 25.** *For I -structures \mathcal{A} and \mathcal{B} , there is a homomorphism $f : \mathbb{G}_{n,k}\mathcal{A} \rightarrow \mathcal{B}$ if, and only if, there is a winning strategy for the Duplicator in the game $+\text{Fun}_n^k$*

Furthermore, we can see that

► **Lemma 26.** *The counit ϵ and comultiplication δ for \mathbb{P}_k lift to well-defined natural transformations for $\mathbb{G}_{n,k}$*

16:12 Game Comonads & Generalised Quantifiers

We will call these lifted natural transformations $\epsilon^{n,k} : \mathbb{G}_{n,k} \rightarrow 1$ and $\delta^{n,k} : \mathbb{G}_{n,k} \rightarrow \mathbb{G}_{n,k}\mathbb{G}_{n,k}$. As $q_n \circ \mathbb{P}_k q_n = \mathbb{G}_{n,k} q_n \circ q_n$, we have that for any $t \in (\mathbb{P}_k)^m \mathcal{A}$ the notion of “the” equivalence class of t , $\mathbf{q}_n(t) \in (\mathbb{G}_{n,k})^m \mathcal{A}$ is well-defined. So for any term T built from composing ϵ, δ and \mathbb{P}_k we have that the term \tilde{T} , obtained by replacing ϵ by $\epsilon^{n,k}$, δ with $\delta^{n,k}$ and \mathbb{P}_k with $\mathbb{G}_{n,k}$ satisfies $\mathbf{q}_n(T(t)) = \tilde{T}(\mathbf{q}_n(t))$ by the above proof. Now as the counit and coassociativity laws are equations in ϵ and δ which remain true on taking the quotient we have the following result.

► **Theorem 27.** $(\mathbb{G}_{n,k}, \epsilon^{n,k}, \delta^{n,k})$ is a comonad on $\mathcal{R}(\sigma)$

4.3 Classifying the morphisms of $\mathcal{K}(\mathbb{G}_{n,k})$

In Abramsky et al.’s treatment of the Kleisli category of \mathbb{P}_k [3] they classify the morphisms according to whether their *branch maps* are injective, surjective or bijective. We extend this definition to the comonad $\mathbb{G}_{n,k}$. This gives us a way of classifying the morphisms to match the classification of strategies given in Section 3.

► **Definition 28.** For $f : \mathbb{G}_{n,k}\mathcal{A} \rightarrow \mathcal{B}$ a Kleisli morphism of $\mathbb{G}_{n,k}$, the branch maps of f are defined as the following collection of functions $A \rightarrow B$, indexed by the structured n, k -histories $t \in ((A \times [k])^{\leq n})^*$ and defined as $\phi_t^f(x) := f([t|x])$. We say that such an f is branch-bijective (resp. branch-injective, -surjective) if for every t ϕ_t^f is bijective (resp. injective, surjective). We denote these maps by $\mathcal{A} \rightarrow_{n,k}^b \mathcal{B}$ (resp. $\mathcal{B} \mathcal{A} \rightarrow_k^i \mathcal{B}$ and $\mathcal{A} \rightarrow_k^s \mathcal{B}$)

Informally, the branch map ϕ_s^g is the response given by Duplicator in the $+\mathbf{Fun}_n^k(\mathcal{A}, \mathcal{B})$ when playing according to the strategy represented by g after Spoiler has made the series of plays in s . This gives us another way of classifying the Duplicator winning strategies for the games from Section 3.

► **Lemma 29.** There is a winning strategy for Duplicator in the game $+\mathbf{Bij}_n^k(\mathcal{A}, \mathcal{B})$ (resp. $+\mathbf{Inj}_n^k(\mathcal{A}, \mathcal{B})$, $+\mathbf{Surj}_n^k(\mathcal{A}, \mathcal{B})$) if and only if $\mathcal{A} \rightarrow_{n,k}^b \mathcal{B}$ (resp. $\mathcal{A} \rightarrow_{n,k}^i \mathcal{B}$, $\mathcal{A} \rightarrow_{n,k}^s \mathcal{B}$).

Expanding this connection between Kleisli maps and strategies, we define the following:

► **Definition 30.** We say a Kleisli map $f : \mathbb{G}_{n,k}\mathcal{A} \rightarrow \mathcal{B}$ is strongly branch-bijective (resp. strongly branch-injective, -surjective) if the strategy for $+\mathbf{Bij}_n^k(\mathcal{A}, \mathcal{B})$ (resp. $+\mathbf{Inj}_n^k(\mathcal{A}, \mathcal{B})$, $+\mathbf{Surj}_n^k(\mathcal{A}, \mathcal{B})$) is also a winning strategy for the game $\mathbf{Bij}_n^k(\mathcal{A}, \mathcal{B})$ (resp. $\mathbf{Inj}_n^k(\mathcal{A}, \mathcal{B})$, $\mathbf{Surj}_n^k(\mathcal{A}, \mathcal{B})$) and we denote these maps by $\mathcal{A} \rightarrow_{n,k}^b \mathcal{B}$ (resp. $\mathcal{B} \mathcal{A} \rightarrow_k^i \mathcal{B}$ and $\mathcal{A} \rightarrow_k^s \mathcal{B}$)

Now we generalise a result of Abramsky, Dawar and Wang to the Kleisli category $\mathcal{K}(\mathbb{G}_{n,k})$.

► **Lemma 31.** For \mathcal{A}, \mathcal{B} finite relational structures,

$$\mathcal{A} \xleftrightarrow[n,k]{i} \mathcal{B} \iff \mathcal{A} \xleftrightarrow[n,k]{s} \mathcal{B} \iff \mathcal{A} \rightarrow_{n,k}^b \mathcal{B} \iff \mathcal{A} \cong_{\mathcal{K}(\mathbb{G}_{n,k})} \mathcal{B}$$

This lemma allows us to conclude that the isomorphisms in the category $\mathcal{K}(\mathbb{G}_{n,k})$ correspond with equivalence of structures up to k variable infinitary logic extended by all generalised quantifiers of arity at most n and thus with winning strategies for Hella’s n -bijjective k -pebble game.

► **Theorem 32.** For two finite relational structures \mathcal{A} and \mathcal{B} the following are equivalent:

- $\mathcal{A} \cong_{\mathcal{K}(\mathbb{G}_{n,k})} \mathcal{B}$
- Duplicator has a winning strategy for $\mathbf{Bij}_n^k(\mathcal{A}, \mathcal{B})$
- $\mathcal{A} \equiv_{\mathcal{L}^k(\mathcal{Q}_n)} \mathcal{B}$

Proof. Immediate from Lemma 31 and Hella [15]. ◀

5 Coalgebras and Decompositions

Abramsky et al. [3] show that the coalgebras of the comonad \mathbb{P}_k are, in fact, objects of great interest to finite model theorists. That is, a coalgebra $\alpha : \mathcal{A} \rightarrow \mathbb{P}_k \mathcal{A}$ gives a tree decomposition of \mathcal{A} of width at most $k - 1$. Moreover, any such tree decomposition can be turned into a coalgebra. This result works, in essence, because $\mathbb{P}_k \mathcal{A}$ has a treelike structure where any pebble history, or branch, $s \in \mathbb{P}_k \mathcal{A}$ only witnesses the relations from the $\leq k$ elements of \mathcal{A} which make up the pebbled position on s . So a homomorphism $\mathcal{A} \rightarrow \mathbb{P}_k \mathcal{A}$ witnesses a sort of treelike k -locality of the relational structure of \mathcal{A} and the extra conditions of being a \mathbb{P}_k -coalgebra ensure this can be presented as a tree decomposition (of width $< k$).

In generalising this comonad to $\mathbb{G}_{n,k}$, we have given away some of the restrictive k -local nature of \mathbb{P}_k which makes this argument work. For example, we note that the substructure induced on elements of the form $\{[\epsilon|x] \mid x \in \mathcal{A}\}$ witnesses all relations in \mathcal{A} which have arity $\leq n$. So, in particular, if the maximum arity of σ the signature of \mathcal{A} is less than n , then it is not hard to see how to construct a homomorphism, indeed a coalgebra, $\mathcal{A} \rightarrow \mathbb{G}_{n,k} \mathcal{A}$. So our notion of n -generalised tree decomposition should clearly be more permissive than the notion of tree decomposition, collapsing for $n \geq \text{arity}(\sigma)$ and otherwise allowing a controlled amount of non-locality (parameterised by n). The correct definition, as we prove in this section, is the following.

► **Definition 33.** *An extended tree decomposition of a σ -structure \mathcal{A} is a triple (T, β, γ) with $\beta, \gamma : T \rightarrow 2^A$ such that:*

1. (T, B) with $B : T \rightarrow 2^A$ given by $B(t) = \beta(t) \cup \gamma(t)$ is a tree-decomposition of \mathcal{A} ; and
2. if $a \in \gamma(t)$ and $a \in B(t')$ then $t \leq t'$.

Thus, we can see an extended tree decomposition as a tree decomposition (T, B) where, additionally, at each node t we pick out a subset $\gamma(t)$ of $B(t)$ with the property that every element a of \mathcal{A} appears in at most one $\gamma(t)$ and when it does, this t is the root of the subtree of T in which a appears. We next define the width and arity of an extended tree decomposition.

► **Definition 34.** *The width of an extended tree decomposition (T, β, γ) is $\max_{t \in T} |\beta(t)|$.*

The arity of an extended tree decomposition (T, β, γ) of width k is the least $n \leq k$ such that:

1. if $t < t'$ then $|\beta(t') \cap \gamma(t)| \leq n$; and
2. for every tuple (a_1, \dots, a_m) in every relation R of \mathcal{A} , there is a $t \in T$ such that $\{a_1, \dots, a_m\} \subseteq B(t)$ and $|\{a_1, \dots, a_m\} \cap \gamma(t)| \leq n$.

Any extended tree decomposition $T(\beta, \gamma)$ of a structure \mathcal{A} can be transformed into one in which each $a \in A$ appears in exactly one $\gamma(t)$. Indeed, suppose there is some a for which this is not true and let t be the order minimal element such that $s \in B(t)$. We simply split t into two nodes adding a parent t_a (with no other children) with $\gamma(t_a) = \{a\}$ and $\beta(t_a) = \beta(t) \setminus \{a\}$. This is easily seen to be an extended tree decomposition with the same width and arity. We call such a tree decomposition one in *normal form*.

We are particularly interested in extended tree decompositions that are further well-structured, in the sense that is related to the definition of structured n, k -histories in Section 4.

► **Definition 35.** *An extended tree decomposition with width k and arity n is structured if for every $a \in A$ there is a $t \in T$ s.t. $a \in \gamma(t)$, for every node t , $\gamma(t) \neq \emptyset$, for any child t' of t $\beta(t') \cap \gamma(t) \neq \emptyset$ and for any t'' a child of t' we have that either:*

16:14 Game Comonads & Generalised Quantifiers

- $|\beta(t') \cap \gamma(t)| = n$; or
- $|\beta(t')| < k$; or
- $\gamma(t) \cap \beta(t') \setminus \beta(t'') \neq \emptyset$

For a node t in an extended tree decomposition, we call $\beta(t)$ the *fixed bag* at t and $\gamma(t)$ the *floating bag* at t .

In general, extended tree decompositions of width k and arity 1 correspond exactly with tree decompositions of width k .

► **Lemma 36.** *A relational structure \mathcal{A} has a tree decomposition of width k if, and only if, it has an extended tree decomposition of width k and arity 1*

Relating extended tree decompositions to our construction in Section 4, we note the following easy but important result.

► **Lemma 37.** *For any finite \mathcal{A} , there is a structured extended tree decomposition of $\mathbb{G}_{n,k}\mathcal{A}$ of width k and arity n for some $k, n \in \mathbb{N}$*

We now prove the main claim of this section, that the $\mathbb{G}_{n,k}$ -coalgebras are in correspondence with structured extended tree decompositions of width k and arity n . The correspondence between tree decompositions and coalgebras of \mathbb{P}_k was established in [3] through a partial order on a structure \mathcal{A} called a *tree traversal*. We now introduce an analogous traversal structure to link $\mathbb{G}_{n,k}$ -coalgebras and extended tree decompositions of width k and arity n . The following definitions provide precisely such a structure.

► **Definition 38.** *An n -tree order is a triple $(X, <, \sim)$ where $<$ is a partial order and \sim an equivalence relation, both on the set X , such that:*

1. for all $x, y, z \in X$, $x < y$ and $y \sim z$ implies $x < z$;
2. $(X/\sim, <)$ is a tree order; and
3. for each $x \in X$ and each \sim -equivalence class η , there are at most n elements $y \in \eta$ such that $y < x$.

An n -tree order provides the order structure allowing us to define the traversals we need.

► **Definition 39.** *For a σ -structure \mathcal{A} , let $(A, <, \sim)$ be an n -tree order and $\iota : O \rightarrow 2^{[k]}$ a function, where $O = \{(a, b) \mid a < b\}$ such that*

1. if $b < b'$ or $b \sim b'$, then $\iota(a, b) = \iota(a, b')$; and
2. if $a \neq a'$ and $a \sim a'$ then $\iota(a, b) \cap \iota(a', b) = \emptyset$.
3. if \mathcal{C} is a \sim -equivalence class then $|\bigcup_{a \in \mathcal{C}} \iota(a, b)| \leq n$

This is an n, k -traversal of \mathcal{A} if, for each tuple (a_1, \dots, a_m) in any relation R of \mathcal{A} , we have:

1. for each $i, j \in [m]$ either $a_i < a_j$, $a_j < a_i$ or $a_i \sim a_j$;
2. no more than n elements of $\{a_1, \dots, a_m\}$ belong to any one \sim -equivalence class; and
3. if $a_i < a_j$, there exists $p_i \in \iota(a_i, a_j)$ such that for all $c \in A$ with $a_i < c < a_j$ then $p_i \notin \iota(c, a_j)$.

An n, k -traversal is structured if for any $a < b < c$ such that there is no d with $a < d < b$, we have that either:

- $|\bigcup_{\{a' \mid a \sim a' \text{ and } a' < c\}} \iota(a', c)| = n$; or
- $\bigcup_{\{a' \mid a \sim a' \text{ and } a' < c\}} \iota(a', c) \cap \bigcup_{\{b' \mid b \sim b' \text{ and } b' < c\}} \iota(b', c) \neq \emptyset$

We establish the relationship between extended tree decompositions and n, k -traversals in the following lemma.

► **Lemma 40.** *For a finite structure \mathcal{A} , if \mathcal{A} has an extended tree decomposition of width k and arity n then it has an n, k -traversal. Furthermore, if the extended tree decomposition is structured then there is a structured n, k -traversal.*

We are ready to establish the relationship between n, k traversals and coalgebras of $\mathbb{G}_{n,k}$.

► **Lemma 41.** *There is a coalgebra $\alpha : \mathcal{A} \rightarrow \mathbb{G}_{n,k}\mathcal{A}$ if, and only if, there is a structured n, k -traversal of \mathcal{A}*

We finish this section by putting together these results into a single theorem.

► **Theorem 42.** *For \mathcal{A} a finite relational structure the following are equivalent:*

1. *there is a $\mathbb{G}_{n,k}$ -coalgebra $\alpha : \mathcal{A} \rightarrow \mathbb{G}_{n,k}\mathcal{A}$*
2. *there is a structured extended tree decomposition of \mathcal{A} with width k and arity n*
3. *there is a structured n, k -traversal of \mathcal{A}*

6 Concluding Remarks

The work of Abramsky et al., giving comonadic accounts of pebble games and their relationship to logic has opened up a number of avenues of research. It raises the possibility of studying logical resources through a categorical lens and introduces the notion of *coresources*. This view has been applied to pebble games [3], Ehrenfeucht-Fraïssé games, bisimulation games [4] and also to quantum resources [1, 2]. In this paper we have extended this approach to logics with generalised quantifiers.

The construction of the comonad $\mathbb{G}_{n,k}$ introduces interesting new techniques to this project. The pebbling comonad \mathbb{P}_k is graded by the value of k which we think of as a *coresource* increasing which constrains the morphisms. The new parameter n provides a second coresource, increasing which further constrains the moves of Duplicator. It is interesting that the resulting comonad can be obtained as a quotient of \mathbb{P}_k and the strategy lifting argument developed in Section 4 could prove useful in other contexts. The morphisms in the Kleisli category correspond to winning strategies in a new game we introduce which characterises a natural logic: the positive logic of homomorphism-closed quantifiers. The isomorphisms correspond to an already established game: Hella's n -bijective game with k pebbles. This relationship allows for a systematic exploration of variations characterising a number of natural fragments of the logic with n -ary quantifiers. One natural fragment that is not yet within this framework and worth investigating is the logic of embedding-closed quantifiers of Haigora and Luosto [14].

This work opens up a number of perspectives. Logics with generalised quantifiers have been widely studied in finite model theory. They are less of interest in themselves and more as tools for proving inexpressibility in specific extensions of first-order or fixed-point logic. For instance, the logics with rank operators [7, 12], of great interest in descriptive complexity, have been analysed as fragments of a more general logic with linear-algebraic quantifiers [6]. It would be interesting to explore whether the comonad $\mathbb{G}_{n,k}$ could be combined with a vector space construction to obtain a categorical account of this logic.

More generally, the methods illustrated by our work could provide a way to deconstruct pebble games into their component parts and find ways of constructing entirely new forms of games and corresponding logics. The games we consider and classify are based on Duplicator playing different kinds of functions (i.e. morphisms on finite sets) and maintaining different kinds of homomorphisms (i.e. morphisms in the category of σ -structures). Could we build reasonable pebble games and logics on other categories? In particular, can we bring the algebraic pebble games of [9] into this framework?

References

- 1 Samson Abramsky, Rui Soares Barbosa, Nadish de Silva, and Octavio Zapata. The Quantum Monad on Relational Structures. In Kim G. Larsen, Hans L. Bodlaender, and Jean-Francois Raskin, editors, *42nd International Symposium on Mathematical Foundations of Computer Science (MFCS 2017)*, volume 83 of *LIPICs*, pages 35:1–35:19, 2017. doi:10.4230/LIPICs.MFCS.2017.35.
- 2 Samson Abramsky, Rui Soares Barbosa, Martti Karvonen, and Shane Mansfield. A comonadic view of simulation and quantum resources. In *34th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS, 2019*. doi:10.1109/LICS.2019.8785677.
- 3 Samson Abramsky, Anuj Dawar, and Pengming Wang. The pebbling comonad in finite model theory. In *2017 32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–12, June 2017. doi:10.1109/LICS.2017.8005129.
- 4 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*, volume 119 of *LIPICs*, pages 2:1–2:17, 2018. doi:10.4230/LIPICs.CSL.2018.2.
- 5 Anuj Dawar. Generalized Quantifiers and Logical Reducibilities. *Journal of Logic and Computation*, 5(2):213–226, 1995. doi:10.1093/logcom/5.2.213.
- 6 Anuj Dawar, Erich Grädel, and Wied Pakusa. Approximations of isomorphism and logics with linear-algebraic operators. In *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9-12, 2019, Patras, Greece*, volume 132 of *LIPICs*, pages 112:1–112:14, 2019. doi:10.4230/LIPICs.ICALP.2019.112.
- 7 Anuj Dawar, Martin Grohe, Bjarki Holm, and Bastian Laubner. Logics with rank operators. In *Proceedings of the 2009 24th Annual IEEE Symposium on Logic In Computer Science, LICS '09*, pages 113–122, Washington, DC, USA, 2009. IEEE Computer Society. doi:10.1109/LICS.2009.24.
- 8 Anuj Dawar and Lauri Hella. The expressive power of finitely many generalized quantifiers. *Information and Computation*, 123(2):172–184, 1995.
- 9 Anuj Dawar and Bjarki Holm. Pebble games with algebraic rules. *Fundamenta Informaticae*, Vol. 150, nr 3/4:281–316, 2017.
- 10 Anuj Dawar, Steven Lindell, and Scott Weinstein. Infinitary logic and inductive definability over finite structures. *Information and Computation*, 119(2):160–175, 1995.
- 11 Heinz-Dieter Ebbinghaus and Jörg Flum. *Finite Model Theory*. Springer, 2nd edition, 1999.
- 12 Erich Grädel and Wied Pakusa. Rank logic is dead, long live rank logic! *The Journal of Symbolic Logic*, 84(1):54–87, 2019. doi:10.1017/jsl.2018.33.
- 13 Martin Grohe. *Descriptive Complexity, Canonisation, and Definable Graph Structure Theory*, volume 47 of *Lecture Notes in Logic*. Cambridge University Press, 2017.
- 14 Jevgeni Haigora and Kerkko Luosto. On logics extended with embedding-closed quantifiers, 2014. arXiv:1401.6682.
- 15 Lauri Hella. Definability hierarchies of generalized quantifiers. *Annals of Pure and Applied Logic*, 43(3):235–271, 1989. doi:10.1016/0168-0072(89)90070-5.
- 16 Lauri Hella. Logical hierarchies in PTIME. *Information and Computation*, 129(1):1–19, 1996. doi:10.1006/inco.1996.0070.
- 17 Phokion G. Kolaitis and Jouko A. Väänänen. Generalized quantifiers and pebble games on finite structures. *Annals of Pure and Applied Logic*, 74(1):23–75, 1995. doi:10.1016/0168-0072(94)00025-X.
- 18 Phokion G. Kolaitis and Moshe Y. Vardi. Infinitary logic for computer science. In *Automata, Languages and Programming*, pages 450–473, 1992.
- 19 Phokion G. Kolaitis and Moshe Y. Vardi. A game-theoretic approach to constraint satisfaction. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on Innovative Applications of Artificial Intelligence*, pages 175–181, 2000.

- 20 Leonid Libkin. *Elements Of Finite Model Theory (Texts in Theoretical Computer Science. An Eats Series)*. SpringerVerlag, 2004.
- 21 Per Lindström. First order predicate logic with generalized quantifiers. *Theoria*, 32(3):186–195, 1966. doi:10.1111/j.1755-2567.1966.tb00600.x.