Structure vs. Power: game comonads in finite model theory

Samson Abramsky

Department of Computer Science, University of Oxford

Structure vs. Power

Foundations of CS: the Great Divide

- **Structure**: structure and compositionality. E.g. semantics, type theory, ...
- **Power**: expressiveness and efficiency. Track A, but also "Track B/A"!

We understand very little about how **Structure** can talk to **Power**, and *vice versa*.

The Logical Structures in Computation program aimed to address this.

We shall talk about one direct off-shoot of the program:

- "The Pebbling Comonad in Finite Model Theory", SA, Anuj Dawar and Pengming Wang, LiCS 2017.
- "Game Comonads in Finite Model Theory", Nihil Shah, Oxford M.Sc. dissertation supervised by SA, 2017.

Model theory and deception

- In model theory, we see a structure, not "as it really is" (up to isomorphism) but only up to **definable properties**.
- The crucial notion is equivalence of structures up to the equivalence $\equiv^{\mathcal{L}}$ induced by the logic $\mathcal{L}$:

$$\mathcal{A} \equiv^{\mathcal{L}} \mathcal{B} \;:=\; \forall \phi \in \mathcal{L}.\, \mathcal{A} \models \phi \iff \mathcal{B} \models \phi$$

- It is always true that if a class of structures $\mathcal{K}$ is definable in $\mathcal{L}$, then $\mathcal{K}$ must be saturated under $\equiv^{\mathcal{L}}$.
- In most cases of interest in FMT, the converse is true too.
- In descriptive complexity, we seek to characterize a complexity class **C** (for decision problems) as those classes of structures $\mathcal{K}$ (e.g. graphs) definable in $\mathcal{L}$.

Homomorphisms and pebble games

Homomorphisms play a fundamental role in FMT, CSP, DB:

- existence of homomorphisms $\equiv$ CSP
- preservation of conjunctive queries, fundamental in DB

Existential k-pebble games (Kolaitis and Vardi 90): Spoiler moves pebbles in $\mathcal{A}$, Duplicator responds in $\mathcal{B}$.

Proposition (KV90)

The following are equivalent:

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

A novel perspective

We shall study k-pebble games, not as an external artefact, but as a semantic construction on relation structures.

Given a structure A over a relational signature σ, we shall introduce a new structure $\mathbb{T}_k A$ corresponding to Spoiler playing his part of an existential k-pebble game on A, with the potential codomain B left unspecified.

The idea is that we can exactly recover the content of a Duplicator strategy in B by giving a homomorphism from $\mathbb{T}_k A$ to B.

Thus the notion of **local approximation** built into the k-pebble game is internalised into the category of σ-structures and homomorphisms.

Formally, this construction will be shown to give a **comonad** on this category.

This leads to comonadic characterisations of a number of central concepts in Finite Model Theory.

Pebbling as a semantic construction

Given a structure A, the set of plays in A by the Spoiler is represented by the set of finite non-empty sequences of moves (p, a), where $p \in [k]$ is a pebble index, and $a \in A$.

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

This forms the universe of $\mathbb{T}_k A$.

How do we lift the relations on A to $\mathbb{T}_k A$?

Given e.g. a binary relation E, we define $E^{\mathbb{T}_k A}$ to be the set of pairs of plays $s, t \in \mathbb{T}_k A$ such that

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

The pebbling functor

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

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

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

We have already defined the map $\varepsilon_A : \mathbb{T}_k A \longrightarrow A$ for each structure A. It is easy to see that it is a homomorphism. We also note that this defines a natural transformation. That is, for each homomorphism $f : A \longrightarrow B$, the following diagram commutes.

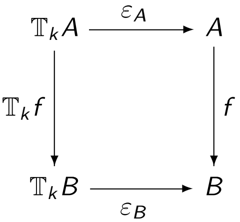

The Co-Kleisli Category

The objects are the same as those of $\mathcal{R}(\sigma)$, while a morphism from A to B in $\mathcal{K}(\mathbb{T}_k)$ is a homomorphism $f : \mathbb{T}_k A \longrightarrow B$.

To compose Kleisli maps, we use the **Kleisli coextension**, which takes $f : \mathbb{T}_k A \longrightarrow B$ to $f^* : \mathbb{T}_k A \longrightarrow \mathbb{T}_k B$.

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

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

Then the composition of $f : \mathbb{T}_k A \longrightarrow B$ and $g : \mathbb{T}_k B \longrightarrow C$ is given by $g \circ f^* : \mathbb{T}_k A \longrightarrow C$.

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

Theorem

The following are equivalent:

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

Grading

Conceptually, we can think of the morphisms $f : A \to_k B$ in the co-Kleisli category for $\mathbb{T}_k$ as those which only have to respect the k-local structure of A.

The lower the value of k, the less information available to Spoiler, and the easier it is for Duplicator to have a winning strategy.

Equivalently, the easier it is to have a morphism $A \to_k B$, *i.e.* a morphism from A to B in the co-Kleisli category.

This leads to a natural **weakening principle**: if we have a morphism from $\mathbb{T}_k A$ to B, then this should yield a morphism from $\mathbb{T}_l A$ to B when $l < k$.

Note that there is an inclusion $\mathbb{T}_l A \hookrightarrow \mathbb{T}_k A$ when $l < k$.

Proposition

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

Game comonads as syntax-free model theory

We can use the pebbling comonad to characterize several key logical equivalences:

- We have the logic C^k, the k-variable logic with counting quantifiers, which plays a central rôle in finite model theory.

Theorem

For all finite A, B: $A \cong_{\mathcal{K}} B \iff A \equiv^{C^k} B$.

- We can also capture equivalence in k-variable logic, and in existential k-variable logic.

Tree-width

A beautiful feature of these comonads is that they let us capture crucial combinatorial invariants of structures using the indexed comonadic structure.

Conceptually, we can think of the morphisms $f : A \to_k B$ in the co-Kleisli category for $\mathbb{T}_k$ as those which only have to respect the k-local structure of A.

The lower the value of k, the less information available to Spoiler, and the easier it is for Duplicator to have a winning strategy.

Equivalently, the easier it is to have a morphism $A \to_k B$, *i.e.* a morphism from A to B in the co-Kleisli category.

What about morphisms $A \to \mathbb{T}_k B$?
Pebbling B makes it **harder** for Duplicator to win the homomorphism game.

Coalgebras

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

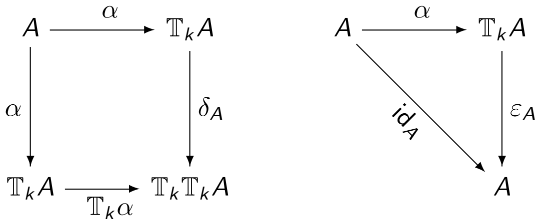

N.B. $\delta : \mathbb{T}_k A \longrightarrow \mathbb{T}_k \mathbb{T}_k A$ is the comultiplication of the comonad, which is the coextension of $\mathrm{id}_{\mathbb{T}_k A}$.

Note that a coalgebra structure α on A implies that a homomorphism exists from A to B whenever a homomorphism exists from $\mathbb{T}_k A$ to B. Given $h : \mathbb{T}_k A \longrightarrow B$, we can form $h \circ \alpha : A \longrightarrow B$.

Thus we should only expect a coalgebra structure to exist when the k-local information on A is sufficient to determine the structure of A.

Coalgebra number

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

Theorem

For all structures A, $\mathrm{tw}(\mathbb{T}_k A) < k$.

Thus although $\mathbb{T}_k A$ is always infinite, it has treewidth bounded by k.

We define the **coalgebra number** $\kappa(A)$ of a finite structure A to be the least k such that there is a coalgebra $\alpha : A \longrightarrow \mathbb{T}_k A$.

Theorem

For all finite structures A:

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

In fact, coalgebras on a structure correspond bijectively to certain "nice" tree decompositions.

Game Comonads

Further work has shown that the same ideas apply to other important logic comparison games:

- Ehrenfeucht-Fraissé games, comparison by quantifier rank
- Bisimulation games for the modal fragment

In each case, we can define corresponding comonads which characterize the logical equivalences.

Moreover, the coalgebra numbers pick out important invariants.

There is clearly a general paradigm here . . .

The Ehrenfeucht-Fraissé Comonad $\mathbb{E}_k$

- $\mathbb{E}_k A$ – sequences of elements of A of length $\leq k$.
- Relations defined in similar fashion to $\mathbb{T}_k A$ (without need for pebble condition).

This gives syntax-free characterizations of 3 important equivalences:

- quantifier-rank fragments of existential-positive logic
- quantifier-rank fragments of FO logic
- quantifier-rank fragments of FO logic with counting

The coalgebra number for these comonads characerizes **tree-depth** (Nešetřil and Ossona de Mendez) in exactly the same way as for tree-width in the case of the pebbling comonad.

This invariant plays an important rôle in Rossman's proof of the Homomorphism Preservation Theorem.

The Unravelling Comonad $\mathbb{U}_k$

Given a Kripke structure (labelled transition system with unary predicates):

- $\mathbb{U}_k A$ is the set of all paths $w_0 i_1 w_1 \cdots i_{k-1} w_n$ where $R^A_{i_j}(w_{j-1}, w_j)$, $1 \leq j \leq n$, and $0 \leq n \leq k$.
- $R_i^{\mathbb{U}_k A}(s, t)$ iff $t = siw$; $P^{\mathbb{U}_k A}(siw)$ iff $P^A(w)$ iff $P^{\mathbb{U}_k A}(w)$.

This yields syntax-free characterizations of equivalences modulo the modal fragment up to modal depth k.

The coalgebra number characterizes the following property:

Theorem

There is a k-coalgebra on A if and only if the multigraph $G(A)$ is a rooted forest of height $\leq k$.

It should be possible to generalize this comonad to the Guarded fragment.

No Mo' Comonads?

Many further directions to pursue:

- Other cases, e.g. games for branching quantifiers or Dependence Logic. May need richer forms of indexing or grading.
- What is the general pattern?
- There is a colimit structure

  $$\mathbb{T}_{\omega}A = \varinjlim \; \mathbb{T}_0 A \to \mathbb{T}_1 A \to \mathbb{T}_2 A \to \cdots$$

  Similarly for $\mathbb{E}_k$ and $\mathbb{U}_k$. What are their properties?
- No-go theorem for finite representation of $\mathbb{T}_k A$.
- Combining these comonads with **quantum monads**.
- Analyze, and give structural, and hopefully more general forms of major results such as Rossman's Homomorphism Preservation Theorem.
- Combine with categorical treatments of the Logic-Automata nexus (Bojanczyk, Gehrke-Petrisan-Colcombet, Adamek-Milius) to look at results such as decidability of the Guarded fragment.