Relating Structure and Power: Comonadic Semantics for Computational Resources

Extended Abstract

Samson Abramsky and Nihil Shah

Department of Computer Science, University of Oxford

samson.abramsky@cs.ox.ac.uk nihil@berkeley.edu

1 Introduction

There is a remarkable divide in the field of logic in Computer Science, between two distinct strands: one focussing on semantics and compositionality ("Structure"), the other on expressiveness and efficiency ("Power"). It is remarkable because these two fundamental aspects of our field are studied using almost disjoint technical languages and methods, by almost disjoint research communiities. We believe that bridging this divide is a major issue in Computer Science, and may hold the key to fundamental advances in the field.

In this paper, we develop a novel approach to relating categorical semantics, which exemplifies the first strand, to finite model theory, which exemplifies the second. It builds on the ideas introduced in [1], but goes much further, showing clearly that there is a strong and robust connection, which can serve as a basis for many further developments.

The setting

Relational structures and the homomorphisms between them play a fundamental rôle in finite model theory, constraint satisfaction and database theory. The existence of a homomorphism $A \to B$ is an equivalent formulation of constraint satisfaction, and also equivalent to the preservation of existential positive sentences [3]. This setting also generalizes what has become a central perspective in graph theory [4].

Model theory and deception

In a sense, the purpose of model theory is "deception". It allows us to see structures not "as they really are", *i.e.* up to isomorphism, but only up to *definable properties*, where definability is relative to a logical language $\mathcal{L}$. The key notion is *logical equivalence* $\equiv^{\mathcal{L}}$. Given structures $\mathcal{A}$, $\mathcal{B}$ over the same vocabulary:

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

If a class of structures $\mathcal{K}$ is definable in $\mathcal{L}$, then it must be saturated under $\equiv^{\mathcal{L}}$. Moreover, for a wide class of cases of interest in finite model theory, the converse holds [5].

The idea of syntax-independent characterizations of logical equivalence is quite a classical one in model theory, exemplified by the Keisler-Shelah theorem [9]. It acquires additional significance in finite model theory, where model comparison games such as Ehrenfeucht-Fraïssé games, pebble games and bisimulation games play a central role [6].

We offer a new perspective on these ideas. We shall study these games, not as external artefacts, but as semantic constructions in their own right. Each model-theoretic comparison game encodes "deception" in terms of limited access to the structure. These limitations are indexed by a parameter which quantifies the resources which control this access. For Ehrenfeucht-Fraïssé games, this is the number of rounds; for pebble games, the number of pebbles; and for bisimulation games, the modal depth.

2 Main Results

We now give a conceptual overview of our main results. Technical details will be provided in a forthcoming paper.

We shall consider three forms of model comparison game: Ehrenfeucht-Fraïssé games, pebble games and bisimulation games [6]. For each of these notions of game G, and value of the resource parameter k, we shall define a corresponding *comonad* $\mathbb{C}_k$ on the category of relational structures and homomorphisms over some relational vocabulary. For each structure $\mathcal{A}$, $\mathbb{C}_k\mathcal{A}$ is another structure over the same vocabulary, which encodes the limited access to $\mathcal{A}$ afforded by playing the game on $\mathcal{A}$ with k resources. There is always an associated homomorphism $\varepsilon_{\mathcal{A}} : \mathbb{C}_k\mathcal{A} \to \mathcal{A}$ (the *counit* of the comonad), so that $\mathbb{C}_k\mathcal{A}$ "covers" $\mathcal{A}$. Moreover, given a homomorphism $h : \mathbb{C}_k\mathcal{A} \to \mathcal{B}$, there is a *Kleisli coextension* homomorphism $h^* : \mathbb{C}_k\mathcal{A} \to \mathbb{C}_k\mathcal{B}$. This allows us to form the *coKleisli category* $\mathsf{Kl}(\mathbb{C}_k)$ for the comonad. The objects are relational structures, while the morphisms from $\mathcal{A}$ to $\mathcal{B}$ in $\mathsf{Kl}(\mathbb{C}_k)$ are exactly the homomorphisms of the form $\mathbb{C}_k\mathcal{A} \to \mathcal{B}$. Composition of these morphisms uses the Kleisli coextension. The connection between this construction and the corresponding form of game G is expressed by the following result:

Theorem 1. *The following are equivalent:*

1. *There is a coKleisli morphism* $\mathbb{C}_k\mathcal{A} \to \mathcal{B}$
2. *Duplicator has a winning strategy for the existential* G*-game with* k *resources, played from* $\mathcal{A}$ *to* $\mathcal{B}$.

The existential form of the game has only a "forth" aspect, without the "back". This means that Spoiler can only play in $\mathcal{A}$, while Duplicator only plays in $\mathcal{B}$. This corresponds to the asymmetric form of the coKleisli morphisms $\mathbb{C}_k\mathcal{A} \to \mathcal{B}$. Intuitively, Spoiler plays in $\mathbb{C}_k\mathcal{A}$, which gives them limited access to $\mathcal{A}$, while Duplicator plays in $\mathcal{B}$. The Kleisli coextension guarantees that Duplicator's strategies can always be lifted to $\mathbb{C}_k\mathcal{B}$; while we can always compose a strategy $\mathbb{C}_k\mathcal{A} \to \mathbb{C}_k\mathcal{B}$ with the counit on $\mathcal{B}$ to obtain a coKleisli morphism.

This asymmetric form may seem to limit the scope of this approach, but in fact this is not the case. For each of these comonads $\mathbb{C}_k$, we have the following three equivalences:

- (E1) $\mathcal{A} \rightleftarrows_k \mathcal{B}$ iff there are coKleisli morphisms $\mathbb{C}_k\mathcal{A} \to \mathcal{B}$ and $\mathbb{C}_k\mathcal{B} \to \mathcal{A}$. Note that there need to be no relationship between these morphisms.
- (E2) $\mathcal{A} \leftrightarrow_k \mathcal{B}$ iff there is a "coupled" pair of coKleisli morphisms $\mathbb{C}_k\mathcal{A} \to \mathcal{B}$ and $\mathbb{C}_k\mathcal{B} \to \mathcal{A}$. The idea is that each coKleisli morphism $\mathbb{C}_k\mathcal{A} \to \mathcal{B}$ expresses a strategy which can be represented in a "positional form". A pair of morphisms is coupled if their positional forms are relational converses of each other. For details, see [1, 8].
- (E3) $\mathcal{A} \cong_{\mathsf{Kl}(\mathbb{C}_k)} \mathcal{B}$ iff $\mathcal{A}$ and $\mathcal{B}$ are isomorphic in the coKleisli category $\mathsf{Kl}(\mathbb{C}_k)$. This means that there are morphisms $\mathbb{C}_k\mathcal{A} \to \mathcal{B}$ and $\mathbb{C}_k\mathcal{B} \to \mathcal{A}$ which are inverses of each other in $\mathsf{Kl}(\mathbb{C}_k)$.

These three notions form a strict hierarchy: (E3) implies (E2), which implies (E1).

For each of our three types of game, there are corresponding fragments $\mathcal{L}_k$ of first-order logic:

- For Ehrenfeucht-Fraïssé games, $\mathcal{L}_k$ is the fragment of quantifier-rank $\leq k$.
- For pebble games, $\mathcal{L}_k$ is the k-variable fragment.
- For bismulation games over relational vocabularies with symbols of arity at most 2, $\mathcal{L}_k$ is the modal fragment [2] with modal depth $\leq k$.

In each case, we write $\exists\mathcal{L}_k$ for the existential positive fragment of $\mathcal{L}_k$, and $\mathcal{L}_k^{\#}$ for the extension of $\mathcal{L}_k$ with counting quantifiers [6].

We can now state our first main result, in a suitably generic form.

Theorem 2. *For finite structures* $\mathcal{A}$ *and* $\mathcal{B}$*:*

$$(1)\ \mathcal{A} \equiv^{\exists\mathcal{L}_k} \mathcal{B} \iff \mathcal{A} \rightleftarrows_k \mathcal{B}.$$

$$(2)\ \mathcal{A} \equiv^{\mathcal{L}_k} \mathcal{B} \iff \mathcal{A} \leftrightarrow_k \mathcal{B}.$$

$$(3)\ \mathcal{A} \equiv^{\mathcal{L}_k^{\#}} \mathcal{B} \iff \mathcal{A} \cong_{\mathsf{Kl}(\mathbb{C}_k)} \mathcal{B}.$$

Note that this is really a family of three theorems. Thus in each case, we capture the salient logical equivalences in syntax-free, categorical form.

We now turn to the significance of indexing by the resource parameter k. When $k \leq l$, we have a natural inclusion morphism $\mathbb{C}_k\mathcal{A} \to \mathbb{C}_l\mathcal{A}$, since playing with k resources is a special case of playing with $l \geq k$ resources. This tells us that the smaller k is, the easier it is to find a morphism $\mathbb{C}_k\mathcal{A} \to \mathcal{B}$. Intuitively, the more we restrict Spoiler's abilities to access the structure of $\mathcal{A}$, the easier it is for Duplicator to win the game.

The contrary analysis applies to morphisms $A \to \mathbb{C}_k\mathcal{B}$. The smaller k is, the *harder* it is find such a morphism. Note, however, that if $\mathcal{A}$ is a finite structure of cardinality k, then $\mathcal{A} \rightleftarrows_k \mathbb{C}_k\mathcal{A}$. In this case, with k resources we can access the whole of $\mathcal{A}$. What can we say when k is strictly smaller than the cardinality of $\mathcal{A}$?

It turns out that there is a beautiful connection between these indexed comonads and combinatorial invariants of structures. This is mediated by the notion of *coalgebra*, another fundamental (and completely general) aspect of comonads. A coalgebra for a comonad $\mathbb{C}_k$ on a structure $\mathcal{A}$ is a morphism $\mathcal{A} \to \mathbb{C}_k\mathcal{A}$ satisfying certain properties. We define the *coalgebra number* of. a structure $\mathcal{A}$, with respect to the indexed family of comonads $\mathbb{C}_k$, to be the least k such that there is a $\mathbb{C}_k$-coalgebra on $\mathcal{A}$.

We now come to our second main result.

Theorem 3.
- *For the pebbling comonad, the coalgebra number of $\mathcal{A}$ corresponds precisely to the* treewidth *of $\mathcal{A}$.*
- *For the Ehrenfeucht-Fraïssé comonad, the coalgebra number of $\mathcal{A}$ corresponds precisely to the* tree-depth *of $\mathcal{A}$ [7].*
- *For the modal comonad, the coalgebra number of $\mathcal{A}$ corresponds precisely to the* forest depth *of $\mathcal{A}$.*

The main idea behind these results is that coalgebras on $\mathcal{A}$ are in bijective correspondence with decompositions of $\mathcal{A}$ of the appropriate form. We thus obtain categorical characterizations of these key combinatorial invariants.

References

1. Samson Abramsky, Anuj Dawar, and Pengming Wang. The pebbling comonad in finite model theory. In *Logic in Computer Science (LICS), 2017 32nd Annual ACM/IEEE Symposium on*, pages 1–12. IEEE, 2017.
2. Hajnal Andréka, István Németi, and Johan van Benthem. Modal languages and bounded fragments of predicate logic. *Journal of Philosophical Logic*, 27(3):217–274, 1998.
3. Ashok K Chandra and Philip M Merlin. Optimal implementation of conjunctive queries in relational data bases. In *Proceedings of the Ninth Annual ACM Symposium on Theory of Computing*, pages 77–90. ACM, 1977.
4. Pavol Hell and Jaroslav Nesetril. *Graphs and homomorphisms.* Oxford University Press, 2004.
5. Phokion G Kolaitis and Moshe Y Vardi. Infinitary logics and 0–1 laws. *Information and Computation*, 98(2):258–294, 1992.
6. Leonid Libkin. *Elements of Finite Model Theory (Texts in Theoretical Computer Science. An EATCS Series).* Springer, 2004.
7. J. Něsetřil and P.O. De Mendez. A unified approach to structural limits, and limits of graphs with bounded tree-depth. *arXiv preprint*, page arXiv:1303.6471, 2013.
8. Nihil Shah. Game comonads in finite model theory. Master's thesis, University of Oxford, 2017.
9. Saharon Shelah. Every two elementarily equivalent models have isomorphic ultrapowers. *Israel Journal of Mathematics*, 10(2):224–233, 1971.