Comonadic semantics for hybrid logic

Samson Abramsky
University College London, London, UK

Dan Marsden
University of Oxford, Oxford, UK

Abstract

Hybrid logic is a widely-studied extension of basic modal logic, which corresponds to the bounded fragment of first-order logic. We study it from two novel perspectives: (1) We apply the recently introduced paradigm of comonadic semantics, which provides a new set of tools drawing on ideas from categorical semantics which can be applied to finite model theory, descriptive complexity and combinatorics. (2) We give a novel semantic characterization of hybrid logic in terms of *invariance under disjoint extensions*, a minimal form of locality. A notable feature of this result is that we give a uniform proof, valid for both the finite and infinite cases.

2012 ACM Subject Classification Theory of computation → Modal and temporal logics

Keywords and phrases comonads, model comparison games, semantic characterizations, hybrid logic, bounded fragment

Digital Object Identifier 10.4230/LIPIcs.MFCS.2022.7

1 Introduction

Hybrid logic (see e.g. [8, 7]) has been widely studied as an expressive extension of basic modal logic. It is semantically natural, e.g. in the analysis of temporal reasoning [10], and since it allows an internalisation of relational semantics, it has a very well-behaved proof theory [11], without needing to resort to explicit labelling of proofs or tableaux. The corresponding fragment of first-order logic under modal translation is the *bounded fragment*, in which quantification is relativized to atomic formulas from the relational vocabulary. This fragment is important in set theory [19], and has been studied in general proof- and model-theoretic terms in [16, 15].

In the present paper we study hybrid logic with inverse modalities, which we shall refer to as hybrid temporal logic, from two novel perspectives:

- Firstly, we apply the recently introduced paradigm of *comonadic semantics* [1, 5], which gives a uniform description of a wide range of logic fragments indexed by resource parameters. These fragments play a key role in finite model theory and descriptive complexity. Examples include the Ehrenfeucht-Fraïssé comonads $\mathbb{E}_k$, which capture the quantifier-rank fragments; the pebbling comonads $\mathbb{P}_k$, which capture the finite variable fragments; and the modal comonads $\mathbb{M}_k$, which capture the modal fragments of bounded modal depth. In each case, the comonads induce a number of resource-indexed equivalences on structures, which can be shown to capture the equivalences induced by the corresponding logic fragments. Moreover, the coalgebras for these comonads can be shown to characterise important combinatorial invariants of structures. For example, in the case of $\mathbb{P}_k$, the corresponding invariant is *tree-width* [5, 1].

  The common structure exhibited by this wide range of examples has been axiomatised in a very general setting in terms of arboreal categories [4]. This provides a new set of tools drawing on ideas from categorical semantics which can be applied to finite model theory, descriptive complexity and combinatorics. Early examples of the use of these ideas can be found in [25, 13, 2, 14], and further results are emerging rapidly, see e.g. [22, 12].

47th International Symposium on Mathematical Foundations of Computer Science (MFCS 2022).
Editors: Stefan Szeider, Robert Ganian, and Alexandra Silva; Article No. 7; pp. 7:1–7:14
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In the present paper, we extend the program of comonadic semantics to hybrid logic. The comonad which captures hybrid logic is a natural restriction of a pointed version of the Ehrenfeucht-Fraïssé comonad previously introduced in [5]. This comonadic analysis nicely reveals, in a clear and conceptual way, the way in which hybrid logic sits between basic modal logic and first-order logic. We characterise the coalgebras for this comonad as tree covers of a relational structure with additional locality constraints. This enables a uniform treatment of logical equivalences, bisimulation games, and combinatorial parameters, within the axiomatic framework recently given in [4].

- Secondly, we give a novel semantic characterization of the version of hybrid logic we study, in terms of *invariance under disjoint extensions* (and various equivalent formulations). This is a minimal form of locality relative to a given base-point, and shows that hybrid logic is the maximal fragment of first-order logic retaining a local character. A notable feature of this result is that we give a uniform proof, valid for both the finite and infinite cases. In particular, we make no use of the compactness theorem, and instead use game-theoretic constructions, specifically a result we call the Workspace Lemma.

  Apart from the interest of the results pertaining to hybrid logic in themselves, we see our work here as fitting into and refining a larger picture, of an emerging landscape in which the tractability of various logic fragments is mirrored in the structural properties of the corresponding comonads. In particular, hybrid logic is undecidable, but still retains a local character. A salient property which the modal and guarded fragments have, and hybrid logic lacks, is the bisimilar companion property. This property plays a key role in the uniform proofs of the van Benthem-Rosen Theorem for these fragments [24, 3]. We mitigate the failure of this property for hybrid logic by the use of game-theoretic arguments. All of this will be explained in detail in Section 6.

After some preliminaries in Section 2, we shall introduce the hybrid comonad in Section 3, and study the coalgebras for this comonad in Section 4. In Section 5, we characterize the equivalence on structures induced by hybrid logic in terms of *spans of open pathwise embeddings* for this comonad, following the pattern established in [5] and axiomatised in [4]. Then we develop the results on the semantic characterisation of hybrid logic in Section 6.

2 Preliminaries

We shall need a few notions on posets. Given $x, y \in P$ for a poset $(P, \leq)$, we write $x\uparrow y$ if x and y are comparable in the order, i.e. $x \leq y$ or $y \leq x$. We will use finite sequences extensively; these are partially ordered by prefix, with notation $s \sqsubseteq t$ indicating that list s is a prefix of list t.

A relational vocabulary σ is a set of relation symbols R, each with a specified positive integer arity. A σ-structure $\mathfrak{A}$ is given by a set A, the universe of the structure, and for each R in σ with arity n, a relation $R^{\mathfrak{A}} \subseteq A^n$. A homomorphism $h : \mathfrak{A} \to \mathfrak{B}$ is a function $h : A \to B$ such that, for each relation symbol R of arity n in σ, for all $a_1, \ldots, a_n$ in A: $R^{\mathfrak{A}}(a_1, \ldots, a_n) \Rightarrow R^{\mathfrak{B}}(h(a_1), \ldots, h(a_n))$. We write $\mathsf{Struct}(\sigma)$ for the category of σ-structures and homomorphisms.

Since evaluation in modal logics is relative to a given world, we shall also use the pointed category $\mathsf{Struct}_{\star}(\sigma)$. Objects are pairs $(\mathfrak{A}, a)$, where $\mathfrak{A}$ is a σ-structure, and $a \in A$. Morphisms $h : (\mathfrak{A}, a) \to (\mathfrak{B}, b)$ are homomorphisms $h : \mathfrak{A} \to \mathfrak{B}$ such that $h(a) = b$.

A *modal vocabulary* has only relation symbols of arity ≤ 2: a set of unary predicate symbols P, which will correspond to modal propositional atoms; and a binary relation symbol E, which we think of as a *transition relation* (more traditionally referred to as an *accessibility relation*).

2.1 Hybrid Temporal Logic

The main system we shall study is Hybrid Temporal Logic (HTL). HTL formulas are built from propositional atoms p and *world variables* x, with the following syntax:

$$\varphi ::= p \mid x \mid \neg\varphi \mid \varphi \wedge \varphi' \mid \varphi \vee \varphi' \mid \Box\varphi \mid \Diamond\varphi \mid \Box^{-}\varphi \mid \Diamond^{-}\varphi \mid {\downarrow}x.\,\varphi \mid @_x\varphi.$$

We use a redundant syntax to make it more convenient to discuss fragments. The new features compared with basic modal logic, augmented with backwards modalities as is standard in temporal logic, are the world variables, which can be bound with $\downarrow$, and used to force evaluation at a given world with @. Hybrid formulae are graded by their *hybrid modal depth*. This is the usual notion of modal depth, with the adjustment that sub-formulae of the form $\Diamond x$ or $\Diamond^{-}x$, for some world variable x, are deemed to have zero depth.

The semantics of hybrid temporal logic is given by translation into first-order logic with equality over a unary modal vocabulary, with a unary predicate P for each proposition atom p, and a single transition relation E. World variables are treated as ordinary first-order variables. The translation is parameterised on a variable, corresponding to the world at which the formula is to be evaluated. We write $\psi[x/y]$ for the result of substituting x for the free occurrences of y in ψ.

$$\begin{aligned}
\mathsf{ST}_x(p) &= P(x) & \mathsf{ST}_x(\Box\varphi) &= \forall y.[E(x,y) \to \mathsf{ST}_y(\varphi)]\\
\mathsf{ST}_x(x') &= x = x' & \mathsf{ST}_x(\Diamond\varphi) &= \exists y.[E(x,y) \wedge \mathsf{ST}_y(\varphi)]\\
\mathsf{ST}_x(\neg\varphi) &= \neg\,\mathsf{ST}_x(\varphi) & \mathsf{ST}_x(\Box^{-}\varphi) &= \forall y.[E(y,x) \to \mathsf{ST}_y(\varphi)]\\
\mathsf{ST}_x(\varphi\wedge\varphi') &= \mathsf{ST}_x(\varphi) \wedge \mathsf{ST}_x(\varphi') & \mathsf{ST}_x(\Diamond^{-}\varphi) &= \exists y.[E(y,x) \wedge \mathsf{ST}_y(\varphi)]\\
\mathsf{ST}_x(\varphi\vee\varphi') &= \mathsf{ST}_x(\varphi) \vee \mathsf{ST}_x(\varphi') & \mathsf{ST}_x({\downarrow}x'.\varphi) &= \mathsf{ST}_x(\varphi)[x/x']\\
& & \mathsf{ST}_x(@_{x'}\varphi) &= \mathsf{ST}_x(\varphi)[x'/x]
\end{aligned}$$

The obvious stipulations about renaming bound variables to avoid variable capture apply.

The target of this translation is the *bounded fragment* of first-order logic with equality, with quantifiers restricted to those of the form $\exists y.[E(x,y) \wedge \varphi]$, $\forall y.[E(x,y) \to \varphi]$, $\exists y.[E(y,x) \wedge \varphi]$, $\forall y.[E(y,x) \to \varphi]$, with $x \neq y$. Hybrid temporal logic is in fact equiexpressive with this fragment [7].

Note that $\mathsf{ST}_x(\Diamond y)$ is logically equivalent to $E(x,y)$, and similarly $\mathsf{ST}_x(\Diamond^{-}y)$ is logically equivalent to $E(y,x)$. Thus these formulas test for the presence of a transition between worlds which have already been reached, justifying our assignment of modal depth 0.

3 The hybrid comonad

We shall now introduce the hybrid comonad on $\mathsf{Struct}_\star(\sigma)$ for modal vocabularies σ, motivating it as combining features of the Ehrenfeucht-Fraïssé and modal comonads from [5].

1. We recall firstly the Ehrenfeucht-Fraïssé comonad $\mathbb{E}_k$ on $\mathsf{Struct}(\sigma)$ for an arbitrary vocabulary σ. Given a structure $\mathfrak{A}$, the universe of $\mathbb{E}_k\mathfrak{A}$ is the set of non-empty sequences of elements of A of length $\leq k$. We think of these sequences as *plays* in the Ehrenfeucht-Fraïssé game on $\mathfrak{A}$. We define the map $\varepsilon_\mathfrak{A} : \mathbb{E}_k A \to A$ which sends a sequence to its last element, which we think of as the current move or *focus* of the play. For a relation R of arity n, we define $R^{\mathbb{E}_k\mathfrak{A}}(s_1,\ldots,s_n)$ to hold iff $s_i{\uparrow}s_j$ for all $1 \leq i,j \leq n$, and $R^\mathfrak{A}(\varepsilon_\mathfrak{A}(s_1),\ldots,\varepsilon_\mathfrak{A}(s_n))$. Explicitly, for unary predicates P, $P^{\mathbb{E}_k\mathfrak{A}}(s)$ iff $P^\mathfrak{A}(\varepsilon_\mathfrak{A}(s))$, and for a binary relation R, $R^{\mathbb{E}_k\mathfrak{A}}(s,t)$ iff $s{\uparrow}t$ and $R^\mathfrak{A}(\varepsilon_\mathfrak{A}(s),\varepsilon_\mathfrak{A}(t))$. Thus the relations hold *along* plays as one extends another, but not between *different* (i.e. incomparable) plays.

2. This construction lifts to the pointed category $\mathsf{Struct}_\star(\sigma)$. We define the universe of $\mathbb{E}_k(\mathfrak{A}, a)$ to comprise the non-empty sequences of length $\leq k+1$ which start with a. The distinguished element is $\langle a \rangle$. The relations are lifted in exactly the same way as previously.
3. The modal comonad $\mathbb{M}_k$ over a modal vocabulary with unary predicates P corresponding to propositional atoms, and a single transition relation E, restricts the sequences in $\mathbb{E}_k(\mathfrak{A}, a)$ to those of the form $\langle a_0, \ldots a_j \rangle$, $a_0 = a$, such that for all i with $0 \leq i < j$, $E^{\mathfrak{A}}(a_i, a_{i+1})$. Thus we can only extend a sequence with an element which the previous element "sees". Moreover, the transition relation E is lifted in a correspondingly local fashion, so that a sequence is only related to its immediate extensions: $E^{\mathbb{M}_k(\mathfrak{A},a)}(s, t)$ iff $t = s\langle a' \rangle$ and $E^{\mathfrak{A}}(\varepsilon_{\mathfrak{A}}(s), \varepsilon_{\mathfrak{A}}(t))$. This is the familiar *unravelling* construction for modal structures [9].
4. The hybrid comonad $\mathbb{H}_k$ is again defined on the pointed category $\mathsf{Struct}_\star(\sigma)$. $\mathbb{H}_k(\mathfrak{A}, a)$ has as universe the subset of $\mathbb{E}_k(A, a)$ of those sequences $\langle a_0, a_1, \ldots, a_l \rangle$ such that $a_0 = a$, and for all j with $0 < j \leq l$, for some i, $0 \leq i < j$, $E^{\mathfrak{A}}(a_i, a_j)$ or $E^{\mathfrak{A}}(a_j, a_i)$. Thus we relax the locality condition of $\mathbb{M}_k$ to the condition that a sequence can only be extended with an element if it is related to *some* element which has been played previously. The σ-relations on $\mathbb{H}_k(\mathfrak{A}, a)$ are defined exactly as for $\mathbb{E}_k(\mathfrak{A}, a)$, and the distinguished element is $\langle a \rangle$, so $\mathbb{H}_k(\mathfrak{A}, a)$ is the induced substructure of $\mathbb{E}_k(\mathfrak{A}, a)$ given by this restriction of the universe. In this sense, $\mathbb{H}_k$ is closer to $\mathbb{E}_k$ than to $\mathbb{M}_k$.

To complete the specification of $\mathbb{H}_k$, we define the *coKleisli extension*: given a morphism $h : \mathbb{H}_k(\mathfrak{A}, a) \to (\mathfrak{B}, b)$, we define $h^* : \mathbb{H}_k(\mathfrak{A}, a) \to \mathbb{H}_k(\mathfrak{B}, b)$ by

$$h^*(\langle a, a_1, \ldots, a_i \rangle) = \langle h(\langle a \rangle), h(\langle a, a_1 \rangle), \ldots, h(\langle a, a_1, \ldots, a_i \rangle) \rangle.$$

We can verify that for each structure $\mathfrak{A}$, $\varepsilon_{\mathfrak{A}} : \mathbb{H}_k\mathfrak{A} \to \mathfrak{A}$ is a morphism; that for each morphism $h : \mathbb{H}_k(\mathfrak{A}, a) \to (\mathfrak{B}, b)$, $h^* : \mathbb{H}_k(\mathfrak{A}, a) \to \mathbb{H}_k(\mathfrak{B}, b)$ is a morphism; and that the following equations are satisfied, for all morphisms $h : \mathbb{H}_k(\mathfrak{A}, a) \to (\mathfrak{B}, b)$, $g : \mathbb{H}_k(\mathfrak{B}, b) \to (\mathfrak{C}, c)$:

$$\varepsilon_{\mathfrak{A}} \circ h^* = h, \qquad \varepsilon_{\mathfrak{A}}^* = \mathsf{id}_{\mathbb{H}_k\mathfrak{A}}, \qquad (g \circ h^*)^* = g^* \circ h^*,$$

This establishes the following result.

▶ **Proposition 1.** *The triple* $(\mathbb{H}_k, \varepsilon, (\cdot)^*)$ *is a comonad in Kleisli form [21].*

It is then standard [21] that $\mathbb{H}_k$ extends to a functor by $\mathbb{H}_k f = (f \circ \epsilon)^*$; that ε is a natural transformation; and that if we define the comultiplication $\delta : \mathbb{H}_k \Rightarrow \mathbb{H}_k^2$ by $\delta_{\mathfrak{A}} = \mathsf{id}^*_{\mathbb{H}_k\mathfrak{A}}$, then $(\mathbb{H}_k, \varepsilon, \delta)$ is a comonad.

3.1 I-morphisms and equality

Like the Ehrenfeucht-Fraïssé comonad $\mathbb{E}_k$, and unlike the modal comonad $\mathbb{M}_k$, equality is important for $\mathbb{H}_k$, as we might expect from its appearance in the translation of hybrid temporal logic into first-order logic. We shall follow the procedure introduced in [5, Section 4] to ensure that equality is properly handled in $\mathbb{E}_k$.

The issue is that elements of A may be repeated in the plays in $\mathbb{H}_k(\mathfrak{A}, a)$. In particular, this happens when there are cycles in the graph $(A, E^{\mathfrak{A}})$ which are reachable from a. We wish to view coKleisli morphisms $f : \mathbb{H}_k(\mathfrak{A}, a) \to (\mathfrak{B}, b)$ as winning strategies for Duplicator in the one-sided (or existential) Spoiler-Duplicator game from $(\mathfrak{A}, a)$ to $(\mathfrak{B}, b)$, in which Spoiler plays in $\mathfrak{A}$ and Duplicator in $\mathfrak{B}$ [18]. In order to fulfil the partial homomorphism winning condition, f must map repeated occurrences of an element $a' \in A$ in a play s in $\mathbb{H}_k(\mathfrak{A}, a)$ to the same element of B. The same issue will recur when we deal with back-and-forth games in section 5. We seek a systematic means of enforcing this requirement.

Given a relational vocabulary σ, we produce a new one $\sigma^+ = \sigma \cup \{I\}$, where I is a binary relation symbol not in σ. If we interpret $I^{(\mathfrak{A},a)}$ and $I^{(\mathfrak{B},b)}$ as the identity relations on A and B, then, following the general prescription for relation lifting in $\mathbb{E}_k(\mathfrak{A}, a)$, and hence also in $\mathbb{H}_k(\mathfrak{A}, a)$ as an induced substructure of $\mathbb{E}_k(\mathfrak{A}, a)$, we have $I^{(\mathbb{H}_k\mathfrak{A},a)}(s, t)$ iff $s{\uparrow}t$ and $\varepsilon_{\mathfrak{A}}(s) = \varepsilon_{\mathfrak{A}}(t)$. Thus a σ-morphism $f : (\mathbb{H}_k\mathfrak{A}, a) \to (\mathfrak{B}, b)$ satisfies the required condition iff it is a σ^+-morphism.

As it stands, this is an ad hoc condition: it relies on a special interpretation of the I-relation. We want our objects to live in $\mathsf{Struct}_\star(\sigma)$, but our morphisms to live in $\mathsf{Struct}_\star(\sigma^+)$. To accomplish this, we use a simple special case of the notion of *relative comonad* [6]. We can take advantage of the fact that $\mathbb{E}_k$, and hence $\mathbb{H}_k$ as a sub-comonad of $\mathbb{E}_k$, is defined uniformly in the vocabulary. Given a vocabulary σ, there is a full and faithful embedding $J : \mathsf{Struct}_\star(\sigma) \to \mathsf{Struct}_\star(\sigma^+)$ such that $I^{J(\mathfrak{A},a)}$ is the identity on A. Moreover, we have a comonad $\mathbb{E}^I_k$, which is the $\mathbb{E}_k$ construction applied to $\mathsf{Struct}_\star(\sigma^+)$. Note that this treats I like any other binary relation in the vocabulary.

We correspondingly obtain $\mathbb{H}^I_k(\mathfrak{A}, a)$ as the substructure of $\mathbb{E}^I_k(\mathfrak{A}, a)$ induced by restricting the universe to that of $\mathbb{H}_k(\mathfrak{A}, a)$. It is important to note that only the transition relation E is used to restrict the universe.

We use this data to obtain the J-relative comonad $\mathbb{H}^+_k = \mathbb{H}^I_k \circ J$ on $\mathsf{Struct}_\star(\sigma)$. The objects of the coKleisli category for this relative comonad are those of $\mathsf{Struct}_\star(\sigma)$. CoKleisli morphisms have the form $\mathbb{H}^I_k J(\mathfrak{A}, a) \to J(\mathfrak{B}, b)$. The counit and coextension are the restrictions of those for $\mathbb{H}^I_k$ to the image of J.

3.2 CoKleisli maps, existential games, and the existential positive fragment

The standard k-round existential Ehrenfeucht-Fraïssé game from $\mathfrak{A}$ to $\mathfrak{B}$ [18, 5] is defined as follows. In each round i, Spoiler moves by choosing an element a_i from A, and Duplicator responds by choosing an element b_i from B. The winning condition for Duplicator is that the correspondence $a_i \mapsto b_i$ is a partial homomorphism from $\mathfrak{A}$ to $\mathfrak{B}$.

The k-round existential hybrid game from $(\mathfrak{A}, a)$ to $(\mathfrak{B}, b)$ is defined in exactly the same way, with two additional provisos:

- At round 0, Spoiler must play $a_0 = a$, and Duplicator must respond with $b_0 = b$.
- At round $j > 0$, Spoiler must play a move a_j such that, for some $i < j$, $E^{\mathfrak{A}}(a_i, a_j)$ or $E^{\mathfrak{A}}(a_j, a_i)$.

▶ **Proposition 2.** *There is a bijective correspondence between*

- *Winning strategies for Duplicator in the k-round existential hybrid game from $(\mathfrak{A}, a)$ to $(\mathfrak{B}, b)$*
- *CoKleisli morphisms $h : \mathbb{H}^+_k(\mathfrak{A}, a) \to J(\mathfrak{B}, b)$.*

The existential positive fragment $\mathsf{HTL}^{\Diamond}$ of hybrid temporal logic is defined by omitting negation and both $\Box$ and $\Box^-$ from the syntax for hybrid logic given in section 2.1. $\mathsf{HTL}^{\Diamond}_k$ is the fragment of $\mathsf{HTL}^{\Diamond}$ comprising formulas of hybrid modal depth $\leq k$.

This fragment induces a preorder on pointed structures. Define $(\mathfrak{A}, a) \Rrightarrow^{\mathsf{HTL}}_k (\mathfrak{B}, b)$ as:

$$\forall \varphi \in \mathsf{HTL}^{\Diamond}_k.\, [(\mathfrak{A}, a) \models \varphi \;\Rightarrow\; (\mathfrak{B}, b) \models \varphi].$$

Here by $(\mathfrak{A}, a) \models \varphi$ we mean $(\mathfrak{A}, a) \models \psi(x)$, where $\psi(x) = \mathsf{ST}_x(\varphi)$.

The following is a variation on standard results (see e.g. [9, 5]).

▶ **Proposition 3.** *There is a winning strategy for Duplicator in the k-round existential hybrid game from $(\mathfrak{A}, a)$ to $(\mathfrak{B}, b)$ iff $(\mathfrak{A}, a) \Rrightarrow_k^{\mathsf{HTL}} (\mathfrak{B}, b)$.*

We define another preorder on pointed structures: $(\mathfrak{A}, a) \rightarrow_k^{\mathbb{H}} (\mathfrak{B}, b)$ iff there is a coKleisli morphism $h : \mathbb{H}_k^+(\mathfrak{A}, a) \rightarrow J(\mathfrak{B}, b)$. The following is then an immediate consequence of Propositions 2 and 3.

▶ **Theorem 4.** *Let σ be a finite modal vocabulary. For all $(\mathfrak{A}, a)$, $(\mathfrak{B}, b)$ in $\mathsf{Struct}_\star(\sigma)$:*

$$(\mathfrak{A}, a) \Rrightarrow_k^{\mathsf{HTL}} (\mathfrak{B}, b) \iff (\mathfrak{A}, a) \rightarrow_k^{\mathbb{H}} (\mathfrak{B}, b).$$

4 Coalgebras

We now study coalgebras for the hybrid comonad. These will yield a natural combinatorial invariant associated with hybrid logic, and also provide a basis for the semantic characterization of the equivalence on structures induced by hybrid logic, which will be given in the following section.

A coalgebra for a comonad (G, ε, δ) is a morphism $\alpha : A \rightarrow GA$ such that $\varepsilon_A \circ \alpha = \mathsf{id}_A$ and $\delta_A \circ \alpha = G(\alpha) \circ \alpha$. Given G-coalgebras $\alpha : A \rightarrow GA$ and $\beta : B \rightarrow GB$, a coalgebra morphism from α to β is a morphism $h : A \rightarrow B$ such that $\beta \circ h = G(h) \circ \alpha$. This gives a category of coalgebras and coalgebra morphisms, denoted by $\mathbf{EM}(G)$, the *Eilenberg-Moore category* of G.

We will now analyze $\mathbf{EM}(\mathbb{H}_k)$, the category of coalgebras for the hybrid comonad on a unimodal vocabulary σ. This will lead to a natural combinatorial parameter associated with hybrid temporal logic, which is a refinement of *tree-depth* [23]. It will also provide a basis for a comonadic characterisation of bisimulation and the equivalence on structures induced by the full hybrid temporal logic, as we will see in the next section.

We will need a few more notions on posets. A chain in a poset $(P, \leq)$ is a subset $C \subseteq P$ such that, for all $x, y \in C$, $x \uparrow y$. A *forest* is a poset $(F, \leq)$ such that, for all $x \in F$, the set of predecessors $\downarrow(x) := \{y \in F \mid y \leq x\}$ is a finite chain. The height $\mathsf{ht}(F)$ of a forest F is $\sup_C |C|$, where C ranges over chains in F. Note that the height is either finite or ω. A *tree* is a forest with a least element (the root). We write the covering relation for a poset as $\prec$; thus $x \prec y$ iff $x \leq y$, $x \neq y$, and for all z, $x \leq z \leq y$ implies $z = x$ or $z = y$. Morphisms of trees are monotone maps preserving the root and the covering relation.

Given a σ-structure $\mathfrak{A}$, the *Gaifman graph* $\mathcal{G}(\mathfrak{A})$ is $(A, \frown)$, where $a \frown a'$ (a is adjacent to a') if they are distinct elements of A which both occur in a tuple of some relation $R^{\mathfrak{A}}$, R in σ.

A *tree cover* of a pointed σ-structure $(\mathfrak{A}, a)$ is a tree order $(A, \leq)$ on A with least element a, and such that if $a \frown a'$, then $a \uparrow a'$. Thus adjacent elements in the Gaifman graph must appear in the same branch of the tree. The tree cover is *generated* if for all $a' \in A$ with $a' \neq a$, for some $a'' \in A$, $a'' < a'$ and $a' \frown a''$. Tree covers of a pointed structure are neither unique, nor guaranteed to exist.

▶ **Theorem 5.** *For any pointed σ-structure $(\mathfrak{A}, a)$, and $k > 0$, there is a bijective correspondence between:*

- *$\mathbb{H}_k$-coalgebras $\alpha : (\mathfrak{A}, a) \rightarrow \mathbb{H}_k(\mathfrak{A}, a)$.*
- *Generated tree covers of $(\mathfrak{A}, a)$ of height $\leq k + 1$.*

We define the *generated tree depth* of $(\mathfrak{A}, a)$ to be the minimum height of any generated tree cover of $(\mathfrak{A}, a)$. This can be seen as a refinement of the standard notion of tree depth [23].

We define the *hybrid coalgebra number* of $(\mathfrak{A}, a)$ to be the least k such that there is an $\mathbb{H}_k$-coalgebra $\alpha : (\mathfrak{A}, a) \to \mathbb{H}_k(\mathfrak{A}, a)$. If there is no coalgebra for any k, the hybrid coalgebra number is ω. The following is an immediate consequence of Theorem 5.

▶ **Theorem 6.** *The generated tree depth of a structure $(\mathfrak{A}, a)$ coincides with its hybrid coalgebra number.*

We define a category $\mathbf{Tree}(\sigma)$ with objects $(\mathfrak{A}, a, \leq)$, where $(\mathfrak{A}, a)$ is a pointed σ-structure, and $\leq$ is a generated tree cover of $(\mathfrak{A}, a)$. Morphisms $h : (\mathfrak{A}, a, \leq) \to (\mathfrak{B}, b, \leq')$ are morphisms of pointed σ-structures which are also tree morphisms. That is, they preserve the covering relation $\prec$ in the tree order, and the root element. For each $k > 0$, there is a full subcategory $\mathbf{Tree}(\sigma)_k$ determined by those objects whose covers have height $\leq k$.

▶ **Theorem 7.** *For each $k > 0$, $\mathbf{Tree}(\sigma)_k$ is isomorphic to $\mathbf{EM}(\mathbb{H}_k)$.*

There is an evident forgetful functor $U_k : \mathbf{Tree}(\sigma)_k \to \mathsf{Struct}_\star(\sigma)$ which sends $(\mathfrak{A}, a, \leq)$ to $(\mathfrak{A}, a)$. The following is now an immediate consequence of Theorem 7.

▶ **Theorem 8.** *For each $k > 0$, U_k has a right adjoint $R_k : \mathsf{Struct}_\star(\sigma) \to \mathbf{Tree}(\sigma)_k$ given by $R_k(\mathfrak{A}, a) = (\mathbb{H}_k(\mathfrak{A}, a), \sqsubseteq)$. The comonad induced by this adjunction is $\mathbb{H}_k$. The adjunction is comonadic.*

5 Paths, open maps, and back-and-forth equivalence

The coalgebra category $\mathbf{EM}(\mathbb{H}_k)$ has a richer structure than $\mathsf{Struct}_\star(\sigma)$, articulated as $\mathbf{Tree}(\sigma)_k$ by Theorem 7. In fact, $\mathbf{Tree}(\sigma)_k$ is an *arboreal category* as defined in [4]. The axiomatic structure of an arboreal category allows us to define notions of bisimulation and games on this category, which can then be transferred to $\mathsf{Struct}_\star(\sigma)$ via the adjunction $U_k \dashv R_k$, following the general pattern laid out in [5]. This leads to a semantic characterisation of the equivalence on structures induced by hybrid logic.

To accommodate I-morphisms, as discussed in section 3.1, we work with the J-relative version of this adjunction, using $R_k^+ = R_k^I J$, where R_k^I is the instance of the adjunction for $\mathsf{Struct}_\star(\sigma^+)$.

5.1 Embeddings, paths and pathwise embeddings

A morphism e in $\mathbf{Tree}(\sigma)_k$ is an *embedding* if $U_k(e)$ is an embedding of relational structures. We write $e : T \rightarrowtail U$ to indicate that e is an embedding.

A *path* in $\mathbf{Tree}(\sigma)_k$ is an object P such that the associated tree cover is a finite linear order, so it comprises a single branch; moreover, I^P is the identity relation. We say that $e : P \rightarrowtail T$ is a *path embedding* if P is a path. A morphism $f : T \to U$ in $\mathbf{Tree}(\sigma)_k$ is a *pathwise embedding* if for any path embedding $e : P \rightarrowtail T$, $f \circ e$ is a path embedding.

5.2 Open maps

A morphism $f : T \to U$ in $\mathbf{Tree}(\sigma)_k$ is *open* if, whenever we have a commuting diagram such as 1, where P and Q are paths, there is an embedding $Q \rightarrowtail T$ such that 2 commutes.

$$\begin{tikzcd} P \arrow[r, rightarrowtail] \arrow[d, rightarrowtail] & Q \arrow[d, rightarrowtail] \\ T \arrow[r, "f"'] & U \end{tikzcd} \tag{1}$$

$$\begin{tikzcd} P \arrow[r, rightarrowtail] \arrow[d, rightarrowtail] & Q \arrow[d, rightarrowtail] \arrow[dl, rightarrowtail] \\ T \arrow[r, "f"'] & U \end{tikzcd} \tag{2}$$

This is often referred to as the *path-lifting property.* If we think of f as witnessing a simulation of T by U, path-lifting means that if we extend a given behaviour in U (expressed by extending the path P to Q), then we can find a matching behaviour in T to “cover” this extension. Thus it expresses an abstract form of the notion of “p-morphism” from modal logic [9], or of functional bisimulation.

5.3 Bisimulation

We can now define the *back-and-forth equivalence* $(\mathfrak{A}, a) \leftrightarrow^{\mathbb{H}}_{k} (\mathfrak{B}, b)$ between structures in $\mathsf{Struct}_\star(\sigma)$. This holds if there is a span of open pathwise embeddings in $\mathbf{Tree}(\sigma)_k$ $R^+_k(\mathfrak{A}, a) \leftarrow T \rightarrow R^+_k(\mathfrak{B}, b)$. Note that we are using the arboreal category $\mathbf{Tree}(\sigma)_k$ to define an equivalence on the “extensional category” $\mathsf{Struct}_\star(\sigma)$.

5.4 Games

We shall now define a *back-and-forth game* $\mathcal{G}_k((\mathfrak{A}, a), (\mathfrak{B}, b))$ played between $(\mathfrak{A}, a)$ and $(\mathfrak{B}, b)$, using the comonad $\mathbb{H}_k$. Positions of the game are pairs $(s, t) \in \mathbb{H}_k(\mathfrak{A}, a) \times \mathbb{H}_k(\mathfrak{B}, b)$. The initial position is $(\langle a \rangle, \langle b \rangle)$.

We define a relation $\mathsf{W}((\mathfrak{A}, a), (\mathfrak{B}, b))$ on positions as follows. A pair (s, t) is in $\mathsf{W}((\mathfrak{A}, a), (\mathfrak{B}, b))$ iff for some path P, path embeddings $e_1 : P \rightarrowtail \mathbb{H}_k(\mathfrak{A}, a)$ and $e_2 : P \rightarrowtail \mathbb{H}_k(\mathfrak{B}, b)$, and $p \in P$, $s = e_1(p)$ and $t = e_2(p)$. The intention is that $\mathsf{W}((\mathfrak{A}, a), (\mathfrak{B}, b))$ picks out the winning positions for Duplicator.

At the start of each round of the game, the position is specified by $(s, t) \in \mathbb{H}_k(\mathfrak{A}, a) \times \mathbb{H}_k(\mathfrak{B}, b)$. The round proceeds as follows. Either Spoiler chooses some $s' \succ s$, and Duplicator must respond with $t' \succ t$, resulting in a new position (s', t'); or Spoiler chooses some $t'' \succ t$ and Duplicator must respond with $s'' \succ s$, resulting in (s'', t''). Duplicator wins the round if they are able to respond, and the new position is in $\mathsf{W}((\mathfrak{A}, a), (\mathfrak{B}, b))$.

5.5 Results

▶ **Theorem 9.** *Given $(\mathfrak{A}, a)$, $(\mathfrak{B}, b)$ in $\mathsf{Struct}_\star(\sigma)$, then $(\mathfrak{A}, a) \leftrightarrow^{\mathbb{H}}_{k} (\mathfrak{B}, b)$ iff Duplicator has a winning strategy for $\mathcal{G}_k((\mathfrak{A}, a), (\mathfrak{B}, b))$.*

Proof. The proof is a minor variation of that for [5, Theorem 10.1], the corresponding result for $\mathbb{E}_k$. Alternatively, this is an instance of the very general [4, Theorem 6.9]. ◀

The standard k-round Ehrenfeucht-Fraïssé game between $\mathfrak{A}$ and $\mathfrak{B}$ [20] is defined as follows. In each round i, Spoiler moves by either

- choosing an $a_i \in A$, to which Duplicator responds by choosing a $b_i \in B$; or
- choosing a $b_i \in B$, to which Duplicator responds by choosing an $a_i \in A$.

The winning condition for Duplicator is that the correspondence $a_i \mapsto b_i$ is a partial isomorphism from $\mathfrak{A}$ to $\mathfrak{B}$.

The k-round back-and-forth hybrid game between $(\mathfrak{A}, a)$ and $(\mathfrak{B}, b)$ is defined in exactly the same way, with two additional provisos:

- At round 0, Spoiler must either play $a_0 = a$, to which Duplicator must respond with $b_0 = b$; or $b_0 = b$, to which Duplicator must respond with $a_0 = a$
- At round $j > 0$, if Spoiler plays a move $a_j \in A$ then, for some $i < j$, $E^{\mathfrak{A}}(a_i, a_j)$ or $E^{\mathfrak{A}}(a_j, a_i)$; while if Spoiler plays a move $b_j \in B$ then, for some $i < j$, $E^{\mathfrak{B}}(b_i, b_j)$ or $E^{\mathfrak{B}}(b_j, b_i)$.

The partial isomorphism winning condition ensures that Duplicator is subject to the same constraints.

We write HTL_k for the set of hybrid formulas of modal depth k. We define an equivalence relation on pointed structures $(\mathfrak{A}, a) \equiv_k^{\mathsf{HTL}} (\mathfrak{B}, b)$ as $\forall \varphi \in \mathsf{HTL}_k. [(\mathfrak{A}, a) \models \varphi \iff (\mathfrak{B}, b) \models \varphi]$.

▶ **Theorem 10.** *Let σ be a finite unimodal vocabulary. For all $(\mathfrak{A}, a)$, $(\mathfrak{B}, b)$ in $\mathsf{Struct}_\star(\sigma)$, the following are equivalent:*

1. $(\mathfrak{A}, a) \leftrightarrow_k^{\mathbb{H}} (\mathfrak{B}, b)$.
2. *Duplicator has a winning strategy for the k-round back-and-forth hybrid game between $(\mathfrak{A}, a)$ and $(\mathfrak{B}, b)$.*
3. $(\mathfrak{A}, a) \equiv_k^{\mathsf{HTL}} (\mathfrak{B}, b)$.

6 Semantic characterization of hybrid temporal logic

We shall now prove a semantic characterisation of hybrid temporal logic in terms of invariance under disjoint extensions. A related result is already known [16, 7], however there are several novel features in our account:

- The previous results are for general (possibly infinite) structures, using tools from infinite model theory. We will give a uniform proof, which applies both to general structures, and to the finite case, which, as for the van Benthem-Rosen characterisation of basic modal logic in terms of bisimulation invariance [27, 26], is an independent result.
- Our proof follows similar lines to the uniform proof by Otto of the van Benthem-Rosen Theorem [24]. In particular, we use constructive arguments based on model comparison games, rather than model-theoretic constructions involving compactness. However, a key property used in his proof no longer holds for the hybrid fragment, so the argument has to take a different path.
- We also identify a key combinatorial lemma, implicit in [24], which we call the Workspace Lemma.
- One of the equivalent conditions in our characterization, invariance under disjoint extensions, appears to be new in this context. We can regard invariance under disjoint extensions as a minimal form of locality relative to a given basepoint. Thus this characterization shows that hybrid temporal logic defines the maximal fragment of first-order logic which retains a local character in this sense.

6.1 Comonadic aspects

Comonadic semantics have now been given for a number of important fragments of first-order logic: the quantifier rank fragments, the finite variable fragments, the modal fragment, and guarded fragments. In the landscape emerging from these constructions, some salient properties have come to the fore. These are properties which a comonad, arising from an *arboreal cover* in the sense of [4], may or may not have:

- The comonad may be *idempotent*, meaning that the comultiplication is a natural isomorphism. Idempotent comonads correspond to *coreflective subcategories*, which form the Eilenberg-Moore categories of these comonads. The modal comonads $\mathbb{M}_k$ are idempotent. The corresponding coreflective subcategories are of those modal structures which are tree-models to depth k [5].

- The comonad C may satisfy the following property: for each structure $\mathfrak{A}$, $C\mathfrak{A} \leftrightarrow^{C} \mathfrak{A}$, where $\leftrightarrow^{C}$ is the back-and-forth equivalence associated with C. We shall call this the *bisimilar companion* property. Note that an idempotent comonad, such as $\mathbb{M}_k$, will automatically have this property. The guarded comonads $\mathbb{G}_k$ from [2] are not idempotent, but have the bisimilar companion property, which is thus strictly weaker.
- Finally, the comonads $\mathbb{E}_k$ and $\mathbb{P}_k$ have neither of the above properties. Unlike the modal and guarded fragments, the quantifier rank and finite variable fragments cover the whole of first-order logic, so we call these comonads *expressive*.

Thus we have a strict hierarchy of comonads in the arboreal categories framework:

$$\text{idempotent} \Rightarrow \text{bisimilar companions} \Rightarrow \text{arboreal}.$$

This hierarchy is correlated with tractability: the modal and guarded fragments are decidable, and have the tree-model property [28, 17], while the expressive fragments do not. We can regard these observations as a small first step towards using structural properties of comonadic semantics to classify logic fragments and their expressive power. In [3], idempotence is used to give simple, general proofs of homomorphism preservation theorems for counting quantifier fragments, with an application to graded modal logic; while the bisimilar companion property is used to give a general, uniform Otto-style proof of van Benthem-Rosen theorems.

As we have already remarked, the hybrid comonads are closer to the Ehrenfeucht-Fraïssé comonads $\mathbb{E}_k$ than to the modal comonads $\mathbb{M}_k$. Indeed, the $\mathbb{H}_k$ comonads are neither idempotent, nor have the bisimilar companion property. On the tractability side, they are not decidable [7]. At the same time, they are not fully expressive for first-order logic, thus refining the above hierarchy.

Otto's proof of the van Benthem-Rosen theorem in [24] uses the bisimilar companion property. This is made explicit in the account given in [3]. Because $\mathbb{H}_k$ does not have this property, we shall use a different comonad in our invariance proof for the hybrid fragment.

Given a structure $\mathfrak{A}$, we can define a metric on A valued in the extended natural numbers $\mathbb{N} \cup \{\infty\}$, given by the path distance in the Gaifman graph $\mathcal{G}(\mathfrak{A})$ [20]. We set $d(a, b) = \infty$ if there is no path between a and b. We write $A[a; k]$ for the closed ball, centred on a, also referred to as the k-neighbourhood of a. Given $(\mathfrak{A}, a)$, we define $\mathbb{S}_k(\mathfrak{A}, a)$ to be $(\mathfrak{A}[a; k], a)$, where $\mathfrak{A}[a; k]$ is the substructure of $\mathfrak{A}$ induced by $A[a; k]$. This defines a comonad on $\mathsf{Struct}_\star(\sigma)$. The counit is the inclusion map, while coextension is the identity operation on morphisms, $h^* = h$. The fact that h is a σ-homomorphism implies that paths are preserved, so this is well defined. It is easily verified that $\mathbb{S}_k$ is an idempotent comonad. The corresponding coreflective subcategory of $\mathsf{Struct}_\star(\sigma)$ is the full subcategory of structures which are k-reachable from the initial elements. We can also define an idempotent comonad $\mathbb{S}$, where $\mathbb{S}(\mathfrak{A}, a) := \bigcup_{k\in\mathbb{N}} \mathbb{S}_k(\mathfrak{A}, a)$.

We can use this comonad to state the invariance property of interest. We say that a first-order formula $\varphi(x)$ is *invariant under generated substructures* if for all $(\mathfrak{A}, a)$ in $\mathsf{Struct}_\star(\sigma)$: $(\mathfrak{A}, a) \models \varphi \iff \mathbb{S}(\mathfrak{A}, a) \models \varphi$. It is *invariant under k-generated substructures* if for all $(\mathfrak{A}, a)$ in $\mathsf{Struct}_\star(\sigma)$: $(\mathfrak{A}, a) \models \varphi \iff \mathbb{S}_k(\mathfrak{A}, a) \models \varphi$. We use the standard disjoint union of structures, $\mathfrak{A} + \mathfrak{B}$. This is the coproduct in $\mathsf{Struct}(\sigma)$. We say that a sentence φ is *invariant under disjoint extensions* if for all $(\mathfrak{A}, a)$, $\mathfrak{B}$: $(\mathfrak{A}, a) \models \varphi \iff (\mathfrak{A} + \mathfrak{B}, a) \models \varphi$.

We can now state our main result.

▶ **Theorem 11** (Characterisation Theorem). *For any first-order formula $\varphi(x)$ with quantifier rank q, the following are equivalent:*

1. *φ is invariant under generated substructures.*
2. *φ is invariant under $q2^q$-generated substructures.*
3. *φ is invariant under disjoint extensions.*
4. *φ is equivalent to a sentence ψ of hybrid temporal logic with modal depth $\leq q2^q$.*

Note that this theorem has two versions, depending on the ambient category $\mathcal{C}$ relative to which equivalence is defined: $\forall(\mathfrak{A}, a) \in \mathcal{C}. (\mathfrak{A}, a) \models \varphi \iff (\mathfrak{A}, a) \models \psi$. The first version, for general models, takes $\mathcal{C} = \mathsf{Struct}_\star(\sigma)$. The second, for finite models, takes $\mathcal{C} = \mathsf{Struct}^f_\star(\sigma)$, the full subcategory of finite structures. Neither of these two versions implies the other. Following Otto [24], we aim to give a uniform proof, valid for both versions.

6.2 Proof of the Characterisation Theorem

Firstly, since any sentence can only use a finite vocabulary, we can assume without loss of generality in what follows that σ is finite. This implies that up to logical equivalence, the fragment HTL_k is finite.

Given a formula φ, we write $\mathsf{Mod}(\varphi) := \{(\mathfrak{A}, a) \mid (\mathfrak{A}, a) \models \varphi\}$. We shall use the following variation of a standard result.

▶ **Lemma 12** (Definability Lemma)**.** *For each $k > 0$ and structure $(\mathfrak{A}, a)$, there is a sentence $\theta^{(k)}_{(\mathfrak{A},\vec{a})} \in \mathsf{HTL}_k$ such that, for all $(\mathfrak{B}, b)$, $(\mathfrak{A}, a) \equiv^{\mathsf{HTL}}_k (\mathfrak{B}, \vec{b}) \iff (\mathfrak{B}, b) \models \theta^{(k)}_{(\mathfrak{A},\vec{a})}$.*

This says that $[(\mathfrak{A}, a)]_{\equiv^{\mathsf{HTL}}_k} = \mathsf{Mod}(\theta^{(k)}_{(\mathfrak{A},\vec{a})})$. Since $\equiv^{\mathsf{HTL}}_k$ has finite index, this implies that if $\mathsf{Mod}(\varphi)$ is saturated under $\equiv^{\mathsf{HTL}}_k$, φ is equivalent to a finite disjunction $\bigvee_{i=1}^n \theta^{(k)}_{(\mathfrak{A}_i,\vec{a}_i)}$, and hence to a formula in HTL_k.

The Workspace Lemma

A key step in the argument is a general result we call the Workspace Lemma. A special case of this is implicit in [24]. Note that $\equiv_q$ is elementary equivalence up to quantifier rank q.

▶ **Lemma 13** (Workspace Lemma)**.** *Given $(\mathfrak{A}, a)$ and $q > 0$, there is a structure $\mathfrak{B}$ such that $(\mathfrak{A} + \mathfrak{B}, a) \equiv_q (\mathfrak{A}[a; k] + \mathfrak{B}, a)$, where $k = 2^q$. Moreover, $|B| \leq 2q|A|$. Hence if $\mathfrak{A}$ is finite, so is $\mathfrak{B}$.*

The intuition for the workspace lemma is that the structure $\mathfrak{B}$, the *workspace*, contains enough disjoint copies of $\mathfrak{A}$ and $\mathfrak{A}[a; k]$ that Spoiler cannot tell the composite structures apart in q rounds of the Ehrenfeucht-Fraïssé game. The idea of Duplicator's strategy is that if Spoiler plays a move that is close to a previous position, in terms of distance in the Gaifman graph, Duplicator responds in the corresponding component of the other structure. If on the other hand Spoiler chooses a position that is "far away" from previously chosen positions, Duplicator responds in a fresh component of the appropriate type. The number of copies of both structures in the workspace, and the distances involved, are chosen so that everything is kept sufficiently far apart that Spoiler cannot see the difference between $\mathfrak{A} + \mathfrak{B}$ and $\mathfrak{A}[a; k] + \mathfrak{B}$. The formal argument is a delicate induction, which can be carried out at the level of generality of metric spaces.

We shall also require a few additional lemmas in our proof of the characterisation theorem. The following is immediate from the definitions.

▶ **Proposition 14.** *For each structure $(\mathfrak{A}, a)$ in $\mathsf{Struct}_\star(\sigma)$, we have $\mathbb{S}_k\mathbb{S}(\mathfrak{A}, a) = \mathbb{S}_k(\mathfrak{A}, a)$.*

The following lemma allows us to restrict our attention to generated substructures when considering HTL equivalence.

▶ **Lemma 15.** *For all $k, m > 0$, if $(\mathfrak{A}, a) \equiv^{\mathsf{HTL}}_m (\mathfrak{B}, b)$ then $\mathbb{S}_k(\mathfrak{A}, a) \equiv^{\mathsf{HTL}}_m \mathbb{S}_k(\mathfrak{B}, b)$.*

We also need the following result to strengthen HTL equivalence to first-order equivalence. We do so by lifting a Duplicator strategy for the hybrid game to one for the Ehrenfeucht-Fraïssé game, at the expense of some logical resources needed to traverse the structures step-by-step in the hybrid game.

▶ **Lemma 16.** *For all $k, q > 0$, if $\mathbb{S}_k(\mathfrak{A}, a) \equiv^{\mathsf{HTL}}_{kq} \mathbb{S}_k(\mathfrak{B}, b)$ then $\mathbb{S}_k(\mathfrak{A}, a) \equiv_q \mathbb{S}_k(\mathfrak{B}, b)$.*

Proof of the characterisation theorem

Proof of Theorem 11.

(2) ⇒ (1). Assume that φ is $\mathbb{S}_k$-invariant. Using Proposition 14, $(\mathfrak{A}, a) \models \varphi$ iff $\mathbb{S}_k(\mathfrak{A}, a) \models \varphi$ iff $\mathbb{S}_k\mathbb{S}(\mathfrak{A}, a) \models \varphi$ iff $\mathbb{S}(\mathfrak{A}, a) \models \varphi$.

(1) ⇒ (3). This follows immediately from the fact that $\mathbb{S}(\mathfrak{A} + \mathfrak{B}, a) = \mathbb{S}(\mathfrak{A}, a)$.

(3) ⇒ (4). Suppose that φ is invariant under disjoint extensions (abbreviated as IDE). Let $k = 2^q$. We shall use Lemma 12. Suppose that (i) $(\mathfrak{A}, a) \models \varphi$, and (ii) $(\mathfrak{A}, a) \equiv^{\mathsf{HTL}}_{kq} (\mathfrak{B}, b)$. We must show that $(\mathfrak{B}, b) \models \varphi$. Applying the Workspace Lemma twice, let $\mathfrak{C}$, $\mathfrak{D}$ be such that *(iii)* $(\mathfrak{A} + \mathfrak{C}, a) \equiv_q (\mathfrak{A}[a; k] + \mathfrak{C}, a)$ and *(iv)* $(\mathfrak{B} + \mathfrak{D}, b) \equiv_q (\mathfrak{B}[b; k] + \mathfrak{D}, b)$. From (ii), applying lemmas 15 and 16, we have *(v)* $\mathbb{S}_k(\mathfrak{A}, a) \equiv_q \mathbb{S}_k(\mathfrak{B}, b)$. Now

$$
\begin{array}{lll}
(\mathfrak{A}, a) \models \varphi & \Rightarrow \ (\mathfrak{A} + \mathfrak{C}, a) \models \varphi & \mathsf{IDE} \\
& \Rightarrow \ (\mathfrak{A}[a; k] + \mathfrak{C}, a) \models \varphi & (iii) \\
& \Rightarrow \ \mathbb{S}_k(\mathfrak{A}, a) \models \varphi & \mathsf{IDE} \\
& \Rightarrow \ \mathbb{S}_k(\mathfrak{B}, b) \models \varphi & (v) \\
& \Rightarrow \ (\mathfrak{B}[b; k] + \mathfrak{D}, b) \models \varphi & \mathsf{IDE} \\
& \Rightarrow \ (\mathfrak{B} + \mathfrak{D}, b) \models \varphi & (iv) \\
& \Rightarrow \ (\mathfrak{B}, b) \models \varphi & \mathsf{IDE}
\end{array}
$$

(4) ⇒ (2). We must show that if ψ is a formula in HTL_k, then it is invariant under k-generated substructures. This follows by a straightforward induction on syntax. ◀

▶ **Question 1.** *In his proof of the van Benthem-Rosen Theorem, Otto establishes an exponential succinctness gap between first-order logic and basic modal logic. A bisimulation-invariant first order formula of quantifier rank q has a modal equivalent of modal depth $\leq 2^q$. He shows that this is optimal. In our case, we have a gap of $q2^q$. Is this optimal for hybrid temporal logic?*

7 Further Directions

Everything which has been done for hybrid temporal logic in the present paper can be extended to the bounded fragment of first-order logic. This generalises the hybrid comonad, allowing both arbitrary relational vocabularies and constant symbols. Allowing for constants $c_1, \ldots, c_m$ involves working with the m-pointed category $\mathsf{Struct}_m(\sigma)$. This has objects $(\mathfrak{A}, \vec{a})$, where $\vec{a} = \langle a_1, \ldots, a_m \rangle \in A^m$. Morphisms $h : (\mathfrak{A}, \vec{a}) \to (\mathfrak{B}, \vec{b})$ must preserve these tuples. The intention is that $a_i = c_i^{\mathfrak{A}}$. Note that $\mathsf{Struct}_\star(\sigma) = \mathsf{Struct}_1(\sigma)$. The comonad constructions can be adapted smoothly to this setting, and the corresponding results go through without any problems.

Another variation is to consider hybrid logic without the backwards modalities $\Box^-$, $\Diamond^-$. The semantic significance of this is that directed rather than undirected reachability becomes the salient notion. The comonadic constructions can be adapted to this setting straightforwardly, but the semantic characterization results cannot be transferred directly. We leave the resolution of this issue to future work.

References

1. 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. IEEE, IEEE, 2017.
2. Samson Abramsky and Dan Marsden. Comonadic semantics for guarded fragments. In *2021 36th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–13. IEEE, IEEE, 2021.
3. Samson Abramsky and Luca Reggio. Aboreal categories: applications to preservation and invariance theorems, 2021. In preparation.
4. Samson Abramsky and Luca Reggio. Arboreal categories and resources. In *48th International Colloquium on Automata, Languages, and Programming (ICALP 2021)*, pages 115:1–115:20. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
5. Samson Abramsky and Nihil Shah. Relating structure and power: Comonadic semantics for computational resources. *Journal of Logic and Computation*, 31(6):1390–1428, 2021.
6. Thorsten Altenkirch, James Chapman, and Tarmo Uustalu. Monads need not be endofunctors. In *International Conference on Foundations of Software Science and Computational Structures*, pages 297–311. Springer, 2010.
7. Carlos Areces, Patrick Blackburn, and Maarten Marx. Hybrid logics: Characterization, interpolation and complexity. *J. Symb. Log.*, 66(3):977–1010, 2001. `doi:10.2307/2695090`.
8. Patrick Blackburn. Representation, reasoning, and relational structures: a hybrid logic manifesto. *Logic Journal of the IGPL*, 8(3):339–365, 2000.
9. Patrick Blackburn, Maarten De Rijke, and Yde Venema. *Modal Logic*, volume 53 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2002.
10. Patrick Blackburn and Klaus Frovin Jørgensen. Reichenbach, Prior and hybrid tense logic. *Synthese*, 193(11):3677–3689, 2016.
11. Torben Braüner. *Hybrid logic and its proof-theory*, volume 37. Springer Science & Business Media, 2010.
12. Adam Ó Conghaile. Cohomological k-consistency, 2021. Technical Report.
13. Adam Ó Conghaile and Anuj Dawar. Game comonads and generalised quantifiers. In Christel Baier and Jean Goubault-Larrecq, editors, *29th EACSL Annual Conference on Computer Science Logic (CSL 2021)*, pages 16:1–16:17, 2021. `arXiv:2006.16039`.
14. Anuj Dawar, Tomáš Jakl, and Luca Reggio. Lovász-type theorems and game comonads. In *2021 36th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–13, 2021. `doi:10.1109/LICS52264.2021.9470609`.
15. Solomon Feferman. Persistent and invariant formulas for outer extensions. *Compositio Mathematica*, 20:29–52, 1968.
16. Solomon Feferman and Georg Kreisel. Persistent and invariant formulas relative to theories of higher order. *Bulletin of the American Mathematical Society*, 72(3):480–485, 1966.
17. Erich Grädel. Why are modal logics so robustly decidable? *Bull. EATCS*, 68:90–103, 1999.
18. Phokion G. Kolaitis and Moshe Y. Vardi. On the expressive power of datalog: Tools and a case study. In Daniel J. Rosenkrantz and Yehoshua Sagiv, editors, *Proceedings of the Ninth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, April 2-4, 1990, Nashville, Tennessee, USA*, pages 61–71. ACM Press, 1990. `doi:10.1145/298514.298542`.

19 Azriel Lévy. *A hierarchy of formulas in set theory.* Number 57 in Memoirs of the American Mathematical Society. American Mathematical Soc., 1965.

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

21 Ernest G. Manes. *Algebraic Theories*, volume 26 of *Graduate Texts in Mathematics.* Springer Science & Business Media, 2012.

22 Yoàv Montacute and Nihil Shah. The pebble-relation comonad in finite model theory. *arXiv preprint*, 2021. `arXiv:2110.08196`.

23 Jaroslav Nešetřil and Patrice Ossona De Mendez. Tree-depth, subgraph coloring and homomorphism bounds. *European Journal of Combinatorics*, 27(6):1022–1041, 2006.

24 Martin Otto. Elementary proof of the van Benthem-Rosen characterisation theorem. Technical Report 2342, Department of Mathematics, Technische Universität Darmstadt, 2004.

25 Tom Paine. A pebbling comonad for finite rank and variable logic, and an application to the equirank-variable homomorphism preservation theorem. *Electronic Notes in Theoretical Computer Science*, 352:191–209, 2020.

26 Eric Rosen. Modal logic over finite structures. *Journal of Logic, Language and Information*, 6(4):427–439, 1997.

27 Johan van Benthem. *Modal logic and classical logic.* Bibliopolis, 1983.

28 Moshe Y Vardi. Why is modal logic so robustly decidable? Technical Report TR-97-274, Rice University, 1997.