Common Knowledge and Multi-Scale Locality Analysis in Cayley Structures

Felix Canavoi
TU Darmstadt

Martin Otto
TU Darmstadt

***Abstract*—We investigate multi-agent epistemic modal logic with common knowledge modalities for groups of agents and obtain van Benthem style model-theoretic characterisations, in terms of bisimulation invariance of classical first-order logic over the non-elementary classes of (finite or arbitrary) common knowledge Kripke frames. The fixpoint character of common knowledge modalities and the rôle that reachability and transitive closures play for the derived accessibility relations take our analysis beyond classical model-theoretic terrain and technically pose a novel challenge to the analysis of model-theoretic games. Over and above the more familiar locality-based techniques we exploit a specific structure theory for specially adapted Cayley groups: through the association of agents with sets of generators, all epistemic frames can be represented up to bisimilarity by suitable Cayley groups with specific acyclicity properties; these support a locality analysis at different levels of granularity as induced by distance measures w.r.t. various coalitions of agents.**

I. INTRODUCTION

Bisimulation invariance may be regarded as the crucial semantic feature of modal logics broadly conceived, with all their many and diverse applications in computer science that range from specification of process behaviours to reasoning about knowledge and agents in all kinds of distributed settings. Bisimulation equivalence is based on an intuitive back&forth probing of transitions between and/or passages between possible instantiations of data, possibly subject to observability by individual agents. As a core notion of procedural, behavioural or cognitive equivalence it underpins the very modelling of relevant phenomena in the state- and transition-based format of transition systems or Kripke structures. In this sense, bisimulation invariance is an essential 'sanity' requirement for any logical system that deals with relevant phenomena rather than artefacts of the encoding. Not surprisingly, modal logics in various formats share this preservation property. Moreover, modal logics can often be characterised in relation to classical logics of reference as precisely capturing the bisimulation invariant properties of relevant structures – which turns the required preservation property into a criterion of expressive completeness, and hence a model-theoretic characterisation of a natural level of expressiveness.

For classical basic modal logic, this characterisation is the content of van Benthem's classical theorem, which identifies basic modal logic ML as the bisimulation invariant fragment of first-order logic FO over the (elementary) class of all Kripke structures. In suggestive shorthand: $\mathsf{ML} \equiv \mathsf{FO}/{\sim}$, where $\mathsf{FO}/{\sim}$ stands for the set of those FO-formulae whose semantics is invariant under bisimulation equivalence $\sim$; a fragment that is syntactically undecidable, but equi-expressive with $\mathsf{ML} \subseteq \mathsf{FO}$ (identified with its standard translation to FO).

Theorem 1 (van Benthem [17]). $\mathsf{ML} \equiv \mathsf{FO}/{\sim}$.

Of the many extensions and variations on this theme that have been found, let us just mention two explicitly.

Firstly, by a result of Rosen [16], van Benthem's characterisation theorem $\mathsf{ML} \equiv \mathsf{FO}/{\sim}$ is also good as a theorem of finite model theory, where both, bisimulation-invariance and expressibility in modal logic are interpreted in restriction to the non-elementary class of all *finite* Kripke structures; this drastically changes the meaning and also requires a completely different proof technique. A transparent and constructive proof of expressive completeness that works in both the classical and the finite model theory settings is given in [13] and also in [8]; like many of the more challenging extensions and variations in [14], [5], [15], it relies on a model-theoretic upgrading argument that links finite approximation levels $\sim^\ell$ of full bisimulation equivalence $\sim$ to finite levels $\equiv_q$ of first-order equivalence. A combination of bisimulation respecting model transformations and an Ehrenfeucht–Fraïssé analysis establishes that every $\sim$-invariant first-order property is in fact invariant already under a finite level $\sim^\ell$ of bisimulation equivalence – a compactness phenomenon for $\sim$-invariant FO, despite the unavailability of compactness for FO in cases of interest.

Secondly, by a famous result of Janin and Walukiewicz, a similar characterisation is classically available for the modal μ-calculus L_μ in relation to monadic second-order logic MSO.

Theorem 2 (Janin–Walukiewicz [11]). $\mathsf{L}_\mu \equiv \mathsf{MSO}/{\sim}$.

In this case, the arguments are essentially automata-theoretic, and the status in finite model theory remains open – and a rather prominent open problem indeed.

Epistemic modal logics deal with information in a multi-agent setting, typically modelled by so-called S5 frames, in which accessibility relations for the individual agents are equivalence relations and reflect indistinguishability of possible worlds according to that agent's observations. A characterisation theorem for basic modal logic ML in this

Research of both authors partially supported by DFG grant OT 147/6-1: *Constructions and Analysis in Hypergraphs of Controlled Acyclicity.*

978-1-5090-3018-7/17/$31.00

epistemic setting was obtained in [5], both classically and in the sense of finite model theory. Like the van Benthem–Rosen characterisation, this deals with plain first-order logic (over the elementary class of S5 frames, or over its non-elementary finite counterpart) and can uniformly use Gaifman locality in the analysis of first-order expressiveness.

In contrast, the present paper explores the situation for epistemic modal logic $\mathsf{ML}[\mathsf{CK}]$ in a multi-agent setting with *common knowledge* operators that capture the essence of knowledge that is shared among a group of agents, not just as factual knowledge but also as knowledge of being shared to any iteration depth: everybody in the group also knows that everybody in the groups knows that ... ad libitum. Cf. [7] for a thorough discussion. This notion of common knowledge can be captured as a fixpoint construct, which is definable in MSO and in fact in L_μ. It can also be captured in plain ML in terms of augmented structures, with derived accessibility relations obtained as the transitive closures of combinations of the individual accessibility relations for the relevant agents: we here call these augmented structures *common knowledge structures* or *CK-structures* for short. But be it fixpoints, MSO, or the non-elementary and locality-averse class of CK-frames, all these variations rule out any straightforward use of simple locality-based techniques.

Here we use, as a template of highly intricate yet regular patterns of multi-scale transitive relations, the *coset structure* of Cayley groups w.r.t. various combinations of generators. We can show that (finite) *Cayley structures*, obtained as expansions of relational encodings of Cayley groups by propositional assignments, are universal representatives up to bisimulation of (finite) S5 structures. In this picture, generator combinations model coalitions of agents, cosets w.r.t. generated subgroups model islands of common knowledge or the induced accessibility relations of CK-frames. For the following cf. Definitions 9 and 10.

Lemma 3 (main lemma)**.** *Every (finite) CK-structure admits (finite) bisimilar coverings by Cayley structures (of various degrees of acyclicity w.r.t. their coset, i.e. epistemic structure).*

Cayley groups with suitable acyclicity properties for their coset structure are available from [15]; they are used here in a novel analysis of first-order expressiveness and Ehrenfeucht–Fraïssé games. This allows us to deal with the challenge of locality issues at different scales or levels of granularity, which are induced by reachability and transitivity phenomena for different groups of agents in CK-structures.

Our main theorem is the following.

Theorem 4. $\mathsf{ML}[\mathsf{CK}] \equiv \mathsf{FO}/{\sim}$ *over CK-structures, both classically and in the sense of finite model theory.*

An equivalent alternative formulation would characterise $\mathsf{ML}[\mathsf{CK}]$ as the $\sim$-invariant fragment of $\mathsf{FO}[\mathsf{CK}]$, the extension of FO that gives it access to the derived accessibility relations for common knowledge—now over all (finite) S5 structures.

A preliminary discussion of the technical challenges for the expressive completeness assertion in this theorem, also in comparison to those in related approaches to e.g. Theorem 1, can be found in Section II-D. The proof of the classical version over the class of all, finite or infinite CK-structures is given in Section IV. The techniques used there also pave the way to the more involved technical development in Section V that supports the finite model theory version.

II. Basic notions and terminology

A. S5 and CK Kripke structures and modal logic

For this paper we fix a finite set Γ of agents; individual agents are referred to by labels $a \in \Gamma$. In corresponding *S5 Kripke frames* $(W, (R_a)_{a\in\Gamma})$ the set W of possible worlds is split, for each $a \in \Gamma$, into equivalence classes $[w]_a$ w.r.t. the equivalence relations R_a that form the accessibility relations for the individual agents in this multi-modal Kripke frame. The epistemic reading is that agent a cannot directly distinguish worlds from the same class $[w]_a$; to simplify terminology we also speak of a-edges and a-equivalence classes. An *S5 Kripke structure* is an expansion of an S5 Kripke frame by a propositional assignment, for a given set of basic propositions $(P_i)_{i\in I}$. As individual formulae of all logics considered will only mention finitely many basic propositions, we may also think of the index set I for the basic propositions as a fixed finite set. The propositional assignment is encoded, in relational terms, by unary predicates P_i for $i \in I$, so a typical S5 Kripke structure is specified as

$$\mathfrak{M} = (W, (R_a)_{a\in\Gamma}, (P_i)_{i\in I}).^1$$

Basic modal logic ML for this setting has atomic formulae $\bot, \top$ and p_i for $i \in I$, and is closed under the usual boolean connectives, $\wedge, \vee, \neg$, as well as under the modal operators (modalities, modal quantifiers) $\Box_a$ and $\Diamond_a$ for $a \in \Gamma$. The semantics for ML is the standard one, with an intuitive epistemic reading of $\Box_a$ as "agent a knows that ..." and, dually, $\Diamond_a$ as "agent a regards it as possible that ...", inductively:

- $\mathfrak{M}, w \models p_i$ if $w \in P_i$;
  $\mathfrak{M}, w \models \top$ for all and $\mathfrak{M}, w \models \bot$ for no $w \in W$;
- boolean connectives are treated as usual;
- $\mathfrak{M}, w \models \Box_a\varphi$ if $\mathfrak{M}, w' \models \varphi$ for all $w' \in [w]_a$;
- $\mathfrak{M}, w \models \Diamond_a\varphi$ if $\mathfrak{M}, w' \models \varphi$ for some $w' \in [w]_a$.

The extension of ML to *common knowledge logic* $\mathsf{ML}[\mathsf{CK}]$ introduces further modalities $\Box_\alpha$ and $\Diamond_\alpha$ for every *group of agents* $\alpha \subseteq \Gamma$. The intuitive epistemic reading of $\Box_\alpha$ is that "it is common knowledge among agents in α that ...", and $\Diamond_\alpha$ is treated as the dual of $\Box_\alpha$. The semantics of $\Box_\alpha$ in an S5 Kripke structure $\mathfrak{M}$ as above is given by the condition that $\mathfrak{M}, w \models \Box_\alpha\varphi$ if φ is true in every world w' that is reachable from w by any path formed by edges in R_a for $a \in \alpha$, i.e., for any w' in the equivalence class $[w]_\alpha$ of the derived equivalence relation

$$R_\alpha := \mathsf{TC}(\textstyle\bigcup_{a\in\alpha} R_a),$$

[1] Where confusion is unlikely, we do not explicitly label the interpretations of the R_a and P_i by $\mathfrak{M}$.

where TC denotes (reflexive and symmetric) transitive closure.

- $\mathfrak{M}, w \models \Box_\alpha \varphi$ if $\mathfrak{M}, w' \models \varphi$ for all $w' \in [w]_\alpha$;
- $\mathfrak{M}, w \models \Diamond_\alpha \varphi$ if $\mathfrak{M}, w' \models \varphi$ for some $w' \in [w]_\alpha$.

Note that for singleton sets $\alpha = \{a\}$, $\Box_\alpha$ coincides with $\Box_a$ just as R_a coincides with $R_{\{a\}}$. The modal operators $\Box_\emptyset$ and $\Diamond_\emptyset$ are eliminable: they both refer to just truth in $[w]_\emptyset = \{w\}$. We also use $\tau := \mathcal{P}(\Gamma)$ for the labelling of the expanded list of modalities and the corresponding equivalence relations and classes, so α will range over τ.

Definition 5. With any S5 Kripke frame (or structure) we associate the CK-frame (or structure) obtained as the expansion by the family $(R_a)_{a\in\Gamma}$ to the family $(R_\alpha)_{\alpha\in\tau}$ for $\tau = \mathcal{P}(\Gamma)$, where $R_\alpha = \mathsf{TC}(\bigcup_{a\in\alpha} R_a)$.

We use notation $\mathfrak{M}^{\mathsf{CK}}$ to indicate the passage from the S5 Kripke structure $\mathfrak{M} = (W, (R_a)_{a\in\Gamma}, (P_i)_{i\in I})$ to its associated CK-structure,

$$\mathfrak{M}^{\mathsf{CK}} = (W, (R_a)_{\alpha\in\tau}, (P_i)_{i\in I}),$$

which is again an S5 Kripke structure. The resulting class of CK-structures is non-elementary since the defining conditions for the R_α are not first-order expressible.

Definition 6. The syntax of epistemic modal logic with common knowledge, $\mathsf{ML}[\mathsf{CK}]$, for the set of agents Γ is the same as the syntax of basic modal logic ML with modalities $\Box_\alpha$ and $\Diamond_\alpha$ for $\alpha \in \tau = \mathcal{P}(\Gamma)$. Its semantics, over S5 Kripke structures $\mathfrak{M}$ for the set of agents Γ, is the usual one, evaluated over the associated CK-structures $\mathfrak{M}^{\mathsf{CK}}$.

B. Bisimulation

We present the core ideas surrounding the notion of bisimulation equivalence in the language of model-theoretic back&forth games of the following format. Play is between two players, player **I** and **II**, and over two Kripke structures $\mathfrak{M} = (W, (R_a)_{a\in\Gamma}, (P_i)_{i\in I})$ and $\mathfrak{N} = (V, (R_a)_{a\in\Gamma}, (P_i)_{i\in I})$.

A position of the game consists of a pair of worlds $(w, v) \in W \times V$, marked by placement of single pebbles in $\mathfrak{M}$ and $\mathfrak{N}$.

In a round played from position (w, v), player **I** chooses one of the structures, $\mathfrak{M}$ or $\mathfrak{N}$, and one of the accessibility relations, i.e., one of the labels $a \in \Gamma$, and moves the pebble in the chosen structure along some edge of the chosen accessibility relation; player **II** has to move the pebble along an edge of the same accessibility relation in the opposite structure; the round results in a successor position (w', v').

Either player loses when stuck, **II** loses in any position (w, v) that violates propositional equivalence, i.e., whenever $\{i \in I : w \in P_i\} \neq \{i \in I : v \in P_i\}$; in this case the game terminates with a loss for **II**.

The unbounded game continues indefinitely, and any infinite play is won by **II**. The ℓ-round game is played for ℓ rounds, it is won by **II** if she can play through these ℓ rounds without violating propositional equivalence.

Definition 7. $\mathfrak{M}, w$ and $\mathfrak{N}, v$ are *bisimilar*, $\mathfrak{M}, w \sim \mathfrak{N}, v$, if **II** has a winning strategy in the unbounded bisimulation game on $\mathfrak{M}$ and $\mathfrak{N}$ starting from position (w, v). $\mathfrak{M}, w$ and $\mathfrak{N}, v$ are *ℓ-bisimilar*, $\mathfrak{M}, w \sim^\ell \mathfrak{N}, v$, if **II** has a winning strategy in the ℓ-round bisimulation game starting from position (w, v).

When a common background structure $\mathfrak{M}$ is clear from context we also write just $w \sim w'$ for $\mathfrak{M}, w \sim \mathfrak{M}, w'$, and similarly for $\sim^\ell$.

It is instructive to compare the bisimulation game on $\mathfrak{M}/\mathfrak{N}$ with the game on $\mathfrak{M}^{\mathsf{CK}}/\mathfrak{N}^{\mathsf{CK}}$. On one hand,

$$\mathfrak{M}, w \sim \mathfrak{N}, v \quad \text{iff} \quad \mathfrak{M}^{\mathsf{CK}}, w \sim \mathfrak{N}^{\mathsf{CK}}, v;$$

the non-trivial implication from left to right uses the fact that every move along an R_α-edge can be simulated by a finite number of moves along R_a-edges for $a \in \alpha$. This also means that, in the terminology of classical modal logic, passage from $\mathfrak{M}$ to $\mathfrak{M}^{\mathsf{CK}}$ is *safe for bisimulation*. On the other hand, there is no such correspondence at the level of finite approximations $\sim^\ell$, since the finite number of rounds needed to simulate a single round played on an R_α-edge cannot be uniformly bounded. This illustrates the infinitary character of passage from $\mathfrak{M}$ to $\mathfrak{M}^{\mathsf{CK}}$, and encapsulates central aspects of our concerns here:

- $\mathfrak{M} \mapsto \mathfrak{M}^{\mathsf{CK}}$ breaks standard notions of locality;
- $\mathfrak{M} \mapsto \mathfrak{M}^{\mathsf{CK}}$ is beyond first-order control.

Correspondingly, modal or first-order expressibility over $\mathfrak{M}^{\mathsf{CK}}$ transcends expressibility over $\mathfrak{M}$, and in particular $\mathsf{ML}[\mathsf{CK}]$ transcends ML while still being invariant under $\sim$.

The link between bisimulation and definability in modal logics, is the following well-known modal analogue of the classical Ehrenfeucht–Fraïssé theorem, cf. [3], [8]. Here and in the following we denote as

$$\mathfrak{M}, w \equiv_\ell^{\mathsf{ML}} \mathfrak{N}, v$$

indistinguishability by ML-formulae of modal nesting depth (quantifier rank) up to ℓ, just as $\equiv_q^{\mathsf{FO}}$ or just $\equiv_q$ will denote classical first-order (elementary) equivalence up to quantifier rank q. Over finite relational vocabularies all of these equivalences have finite index, which is crucial for the following.[2]

Theorem 8. *For any finite modal vocabularies (here: finite sets of agents and basic propositions), Kripke structures $\mathfrak{M}$ and $\mathfrak{N}$ with distinguished worlds w and v, and $\ell \in \mathbb{N}$:*

$$\mathfrak{M}, w \sim^\ell \mathfrak{N}, v \quad \textit{iff} \quad \mathfrak{M}, w \equiv_\ell^{\mathsf{ML}} \mathfrak{N}, v.$$

In particular, the semantics of any modal formula (in ML or in $\mathsf{ML}[\mathsf{CK}]$) is preserved under full bisimulation equivalence (of either the underlying plain S5 structures or their CK-expansions). Any formula of $\mathsf{ML}[\mathsf{CK}]$ is preserved under some level $\sim^\ell$ over CK-expansions (but not over the underlying plain S5 structures!).

The following notion will be of special interest for our constructions; it describes a particularly neat bisimulation relationship, mediated by a homomorphism (classical modal

[2] Finite index is crucial for the definability of the $\sim^\ell$-equivalence classes by so-called characteristic formulae $\chi^\ell_{\mathfrak{M},w}$ s.t. $\mathfrak{N}, v \models \chi^\ell_{\mathfrak{M},w}$ iff $\mathfrak{M}, w \sim^\ell \mathfrak{N}, v$.

terminology speaks of bounded morphisms). Bisimilar tree unfoldings are a well-known instance of (albeit, usually infinite) bisimilar coverings with many applications.

Definition 9. A surjective homomorphism $\pi\colon \hat{\mathfrak{M}} \to \mathfrak{M}$ between Kripke structures is called a *bisimilar covering* if $\hat{\mathfrak{M}}, \hat{w} \sim \mathfrak{M}, \pi(\hat{w})$ for all $\hat{w}$ from $\hat{\mathfrak{M}}$.

Control of multiplicities and cycles in Kripke structures plays an essential rôle towards the analysis of first-order expressiveness, simply because they are *not* controlled by bisimulation.

Core results from [5] deal with this at the level of plain S5 Kripke structures, where products with finite Cayley groups of sufficiently large girth suffice to avoid short cycles. These constructions would not avoid the kind of cycles we have to deal with in CK-structures: on one hand we will have to look to stronger acyclicity properties, viz. coset acyclicity of Cayley groups in Section III-A; on the other hand, we can naturally model any CK-scenario directly in a Cayley group.

That highly regular and homogeneous CK-structures can be obtained directly from Cayley groups, and that these structures are in fact generic up to bisimulation, forms a corner-stone of our approach to the analysis of the expressive power of first-order logic for $\sim$-invariant properties over CK-structures (cf. Main Lemma 3 in the introduction, and Lemma 11 below).

C. Common knowledge in Cayley structures

A *Cayley group* is a group $\mathbb{G} = (G, \cdot, 1)$ with a specified set of generators $E \subseteq G$, which in our case will always be distinct, non-trivial involutions: $e \neq 1$ and $e^2 = 1$ for all $e \in E$. $\mathbb{G}$ is generated by E in the sense that every $g \in G$ can be represented as a product of generators, i.e., as a word in E^*, which w.l.o.g. is reduced in the sense of not having any factors e^2. With the Cayley group $\mathbb{G} = (G, \cdot, 1)$ one associates its *Cayley graph*. Its vertex set is the set G of group elements; its edge relations are $R_e := \{(g, ge)\colon g \in G\}$, which in our case are symmetric and indeed complete matchings on G. That $\mathbb{G}$ is generated by E means that the edge-coloured graph $(G, (R_e)_{e\in E})$ is connected; it is also homogeneous in the sense that any two vertices g and h are related by a graph automorphism induced by right multiplication with $g^{-1}h$ in the group.

Partitioning the generator set E into subsets E_a associated with the agents $a \in \Gamma$, we consider subgroups $\mathbb{G}_a = \langle e\colon e \in E_a\rangle \subseteq \mathbb{G}$ generated by the $e \in E_a$, and regard cosets w.r.t. $\mathbb{G}_a$ as a-equivalence classes over G, turning G into the set of possible worlds of an S5 frame. Indeed, the associated equivalence relation

$$R_a := \{(g, gh)\colon h \in \mathbb{G}_a\} = \mathsf{TC}\big(\bigcup\{R_e\colon e \in E_a\}\big)$$

is the (reflexive, symmetric) transitive closure of the edge relation induced by corresponding generators in the Cayley graph. This pattern naturally extends to sets of agents $\alpha \in \tau = \mathcal{P}(\Gamma)$. Writing $\mathbb{G}_\alpha \subseteq \mathbb{G}$ for the subgroup generated by $E_\alpha := \bigcup\{E_a\colon a \in \alpha\}$, the equivalence relations

$$R_\alpha := \{(g, gh)\colon h \in \mathbb{G}_\alpha\} = \mathsf{TC}\big(\bigcup\{R_a\colon a \in \alpha\}\big)$$

are the accessibility relations in the CK-expansion: their equivalence classes *are* the cosets w.r.t. the subgroups $\mathbb{G}_\alpha$ generated by corresponding parts of the Γ-partitioned E.

Definition 10. With any Cayley group $\mathbb{G} = (G, \cdot, 1)$ with generator set E of involutions that is Γ-partitioned, $E = \dot{\bigcup}_{a\in\Gamma} E_a$, we associate the *Cayley CK-frame* (Cayley frame, for short) $\mathbb{G}^{\mathsf{CK}}$ over the set G of possible worlds with accessibility relations R_α for $\alpha \in \tau = \mathcal{P}(\Gamma)$. A *Cayley structure* consist of a Cayley frame together with a propositional assignment.

Note that any Cayley structure is a CK-structure, so that for Cayley structures $\mathfrak{M}$, always $\mathfrak{M}^{\mathsf{CK}} = \mathfrak{M}$. In the following we simply speak of α-edges, -classes, -cosets with reference to the R_α or $\mathbb{G}_\alpha$ in any Cayley structure.

Lemma 11. *Any connected (finite) CK-structure admits a bisimilar covering by a (finite) Cayley CK-structure.*

Proof. We may concentrate on the underlying plain S5 structures with accessibility relations R_a for $a \in \Gamma$ (bisimilar coverings are compatible with the bisimulation-safe passage to CK-structures). For given $\mathfrak{M} = (W, (R_a), (P_i))$ let $E := \dot{\bigcup}_{a\in\Gamma} R_a$ be the disjoint union of the edge sets R_a, partitioned into subsets E_a corresponding to the individual R_a. Let $\mathfrak{M} \oplus 2^E$ stand for the undirected E-edge-labelled graph formed by the disjoint union of $\mathfrak{M}$ with the $|E|$-dimensional hypercube 2^E. With $e \in E$ we associate the involutive permutation π_e of the vertex set V of $\mathfrak{M} \oplus 2^E$ that precisely swaps all pairs of vertices in e-labelled edges. For $\mathbb{G}$ we take the subgroup of the symmetric group on V that is generated by these π_e, which we identify with involutive generators e. Due to the 2^E-component, these $e = \pi_e \in E \subseteq \mathbb{G}$ are pairwise distinct and distinct from $1 \in \mathbb{G}$. We let $\mathbb{G}$ act on V in the natural fashion (from the right): for $g = e_1 \cdots e_n$,

$$g\colon v \mapsto we_1 \cdots e_n := (\pi_{e_n} \circ \cdots \circ \pi_{e_1})(v).$$

This operation is well-defined as a group action, since by definition $e_1 \cdots e_n = 1$ in $\mathbb{G}$ if, and only if, $\pi_{e_n} \circ \cdots \circ \pi_{e_1}$ fixes every $v \in W$. It also leaves $W \subseteq V$ invariant as a set, i.e. the action can be restricted to W. Then the map

$$\begin{array}{rccc} \hat{\pi}\colon & W \times \mathbb{G} & \longrightarrow & W \\ & (w, g) & \longmapsto & wg \end{array}$$

is a surjective homomorphism w.r.t. the edge relations R_a as interpreted in the direct product of $\mathfrak{M}$ and in the Cayley frame associated with $\mathbb{G}$. This homomorphism directly extends to the induced CK-frames with accessibility relations R_α for $\alpha \in \tau$. Moreover, since $\mathfrak{M}$ is connected, $\mathbb{G}$ acts transitively on W and we may restrict to a single connected sheet of the above multiple covering, corresponding to the identification of an (arbitrary) distinguished world $w_0 \in W$ as a base point. We obtain π as the restriction of $\hat{\pi}$ to the subset $\{(w_0 g, g)\colon g \in \mathbb{G}\}$, which is naturally isomorphic with $\mathbb{G}$ itself. We may expand the Cayley frame $(G, (R_\alpha))$ in a unique manner to a Cayley structure $(G, (R_\alpha), (P_i))$ for which π becomes a homomorphism onto $\mathfrak{M}^{\mathsf{CK}}$ – by pulling back $P_i \subseteq W$ to its pre-image $\pi^{-1}(P_i) \subseteq G$. The resulting

$$
\begin{array}{ccc}
\mathfrak{M}, w & \text{——— } \sim^\ell \text{ ———} & \mathfrak{N}, v \\
| & & | \\
\sim & & \sim \\
| & & | \\
\mathfrak{M}^*, w^* & \text{——— } \equiv_q \text{ ———} & \mathfrak{N}^*, v^*
\end{array}
$$

Fig. 1. Upgrading $\sim^\ell$ to $\equiv_q$ in bisimilar companions.

$\pi\colon (G, (R_\alpha), (P_i)) \to \mathfrak{M}^{\mathsf{CK}}$ is a bisimilar covering. Note that $\mathbb{G}$ and $(G, (R_\alpha), (P_i))$ are finite if W is. □

This crucial insight justifies the following, since – up to bisimulation – we may now transfer any model-theoretic question about (finite) CK-structures to (finite) Cayley structures.[3]

Proviso 12. *From now on consider Cayley structures as prototypical representatives of CK-structures.*

D. Upgrading for expressive completeness

The key to the expressive completeness results from [13] to [5], [15] lies in the establishment of the following finiteness or compactness phenomenon for $\sim$-invariant FO-formulae $\varphi(x)$ over the relevant classes $\mathcal{C}$ of structures:

$$
(\dagger) \qquad \begin{array}{c} \{\varphi \in \mathsf{FO}\colon \varphi \sim\text{-invariant over } \mathcal{C}\} \;= \\ \bigcup_{\ell\in\mathbb{N}}\{\varphi \in \mathsf{FO}\colon \varphi \sim^\ell\text{-invariant over } \mathcal{C}\}. \end{array}
$$

This finiteness property in turn follows if suitable levels $\sim^\ell$ can be upgraded in bisimilar companions within $\mathcal{C}$ so as to guarantee equivalence w.r.t. the given φ of quantifier rank q as follows. If, for suitable $\ell = \ell(q)$, any pair of pointed structures $\mathfrak{M}, w \sim^\ell \mathfrak{N}, v$ from $\mathcal{C}$ admits the construction of bisimilar companion structures $\mathfrak{M}^*, w^* \sim \mathfrak{M}, w$ and $\mathfrak{N}^*, v^* \sim \mathfrak{N}, v$ such that $\mathfrak{M}^*, w^* \equiv_q \mathfrak{N}^*, v^*$, as in Figure 1, then, over $\mathcal{C}$, any $\sim$-invariant FO-formula of quantifier rank q is indeed $\sim^\ell$-invariant, and hence expressible in ML at modal nesting depth ℓ over $\mathcal{C}$ by Theorem 8.

Obstructions to be overcome: Considering Figure 1, it is clear that $\mathfrak{M}^*$ and $\mathfrak{N}^*$ must avoid distinguishing features that are definable in FO_q (FO at quantifier rank up to q) but cannot be controlled by $\sim^\ell$ (for a level $\ell = \ell(q)$ to be determined). Potential features of this kind involve small multiplicities w.r.t. accessibility relations and basic propositions (i.e., small class sizes $[w]_\alpha$, possibly in restriction to some P_i) or short non-trivial cycles w.r.t. combinations of the accessibility relations R_α. In the setting of plain Kripke structures rather than our CK-structures, and thus for many of the more immediate variations on Theorem 1, it turns out that both these obstacles can be eliminated in bisimilar coverings by direct products: multiplicities can be boosted above critical thresholds in products with large enough cliques; short cycles can be eliminated in products with Cayley groups of large girth. While these aspects will play a rôle here, too, the greater challenge lies with the game arguments that typically allow us to establish $\equiv_q$. The classical q-round first-order Ehrenfeucht–Fraïssé or pebble game, which serves to establish $\equiv_q$-equivalence of two structures (cf., e.g. [10], [6]), has to be based on some useful structural analysis of the target structures $\mathfrak{M}^*$ and $\mathfrak{N}^*$. While many earlier upgrading results in this vein could rely on classical Gaifman locality arguments for this structural analysis, the situation here is different, because, naively, Gaifman locality is completely trivialised in connected CK-structures, which form a single Gaifman clique. In fact, it seems all but hopeless to use locality techniques in structures that are as dense in terms of their edge relations as CK-structures are. This is where Cayley structures, whose edge pattern is not only very dense but also highly regular and therefore amenable to structured analysis, provide a promising scenario. In this scenario we can perform a structural analysis that can deal with locality at different levels of granularity, depending on the combinations of R_α that are taken into account, and thus perform locality analysis at different scales of distance measures.

III. Cayley groups of controlled acyclicity

To win the classical q-round first-order Ehrenfeucht–Fraïssé or pebble game over two Cayley structures, player **II** must answer player **I**'s moves in such a way as to respect distances between pebbled worlds: short distances must be matched exactly, and long distances must be matched by distances that are also long. Since distances in Cayley structures can be measured on different scales, distance must be respected on all scales simultaneously. To control distances in this manner, we must in particular be able to avoid undesired short paths w.r.t. various combinations of R_α. Connected CK-structures have diameter 1 at the level of R_Γ, so non-trivial distances only arise when we zoom in to the level of α-steps for subsets $\alpha \subsetneq \Gamma$ that do not allow direct edges from source to target world. The notion of *non-trivial short coset paths* (formally defined in Definition 19) formed by overlapping α-classes, i.e. by α-cosets (cosets w.r.t. $\mathbb{G}_\alpha$), therefore becomes a key ingredient in achieving the desired control over distances at varying scales in Cayley groups that avoid (all or at least short) coset cycles.

A. Coset acyclicity vs. large girth

In the case of CK-frames and Cayley frames one cannot hope to avoid cycles outright.[4] Since any Cayley frame is connected, any two of its worlds w and w' are linked by a Γ-edge in any Cayley frame, this is of no concern for the upgrading (in fact, R_Γ is trivially FO-definable in Cayley frames). But crucial distinctions can occur w.r.t. the reducts of Cayley frames without Γ-edges: worlds w and w' may not be related by any single α-edge for $\alpha \subsetneq \Gamma$, but via a non-trivial short path that uses mixed edge relations. Assume we have Cayley structures $\mathfrak{M}$ and $\mathfrak{N}$, and pairs of worlds $(w, v), (w', v') \in W \times V$ such that $\mathfrak{M}, w \sim^\ell \mathfrak{N}, v$ and $\mathfrak{M}, w' \sim^\ell \mathfrak{N}, v'$. It is possible to have two different non-trivial short paths from w to w' but essentially only one such path

[3] Lemma 33 offers (finite) representations with specific acyclicity and richness properties, obtained as coverings by Cayley groups from [15].

[4] This is even true of S5 structures, but at least those cannot have short cycles w.r.t. long-range edge relations like our R_α.

from v to v'; and this difference could be expressible in FO_q. The solution is to find bisimilar companions to $\mathfrak{M}$ and $\mathfrak{N}$ that are locally acyclic w.r.t. non-trivial overlaps between α-classes, i.e. α-cosets $[w]_\alpha$ for various α. Simultaneously, every such coset $[w]_\alpha$ of the structures must be locally acyclic, in the same sense, w.r.t. β-classes for $\beta \subsetneq \alpha$. This is meant by *multi-scale acyclicity*, and it turns out that the following notions of *coset acyclicity* from [15] are what we can use.

Definition 13. Let $\mathfrak{M}$ be a Cayley frame. A *coset cycle of length m in* $\mathfrak{M}$ is a cyclic tuple $(w_i, \alpha_i)_{i \in \mathbb{Z}_m}$, where, for all $i \in \mathbb{Z}_m$, $(w_i, w_{i+1}) \in R_{\alpha_i}$ and

$$[w_i]_{\alpha_{i-1} \cap \alpha_i} \cap [w_{i+1}]_{\alpha_i \cap \alpha_{i+1}} = \emptyset.$$

Definition 14. A Cayley frame is *acyclic* if it does not contain a coset cycle, and *N-acyclic* if it does not contain a coset cycle of length up to N.

In Section IV and V we use bisimilar companions that are fully acyclic, or finite and N-acyclic, respectively.

Because of the transitive nature of the edge relations in Cayley frames, we cannot hope to construct bisimilar companions that have unique paths between any two vertices. However, paths in acyclic Cayley frames are unique in the sense that any two paths between worlds w and w' are just different variations of one and the same core path. In the case of N-acyclic Cayley frames arbitrary paths between worlds w and w' do not have to be related in any sense, but *short* paths behave as well as all paths would in fully acyclic frames. If the degree of acyclicity N is sufficiently high, this will be enough to show $\equiv_q$-equivalence between two ℓ-bisimilar structures.

The rest of this section is mostly concerned with the investigation of 2-acyclicity. It assumes a special rôle in the analysis of Cayley frames. We begin with a useful characterisation.

Lemma 15. *Let $\mathfrak{M}$ be a Cayley frame. $\mathfrak{M}$ is 2-acyclic if, and only if, for all $w \in W, \alpha, \beta \in \tau$ we have*

$$[w]_\alpha \cap [w]_\beta = [w]_{\alpha \cap \beta}.$$

Proof. "$\Leftarrow$" and $[w]_{\alpha\cap\beta} \subseteq [w]_\alpha \cap [w]_\beta$ obvious. For "$\Rightarrow$" consider a configuration $[w]_{\alpha\cap\beta} \subsetneq [w]_\alpha \cap [w]_\beta$. For $v \in [w]_\alpha \cap [w]_\beta \setminus [w]_{\alpha\cap\beta}$, $[w]_{\alpha\cap\beta} \cap [v]_{\alpha\cap\beta} = \emptyset$ so that (w, α, v, β, w) forms a 2-cycle. ☐

We use this characterisation to show that 2-acyclic Cayley frames display a high degree of regularity (and most of the notions that we will introduce in this and the following sections only make sense in 2-acyclic frames). In 2-acyclic Cayley frames, every family of α-classes (cosets) with non-trivial intersection intersects in a unique coset; and for any two vertices, there is a unique smallest set $\alpha \in \tau$ that connects these vertices. In contrast, arbitrary S5- or CK-frames impose very little structure on the overlap patterns between the equivalence classes formed by their accessibility relations.

We shall later draw on the structure of the dual hypergraph associated with a Cayley frame (cf. Definition 22).

Definition 16. In a Cayley frame $\mathfrak{M}$ define the *dual hyperedge* induced by a world w to be the set of cosets

$$[\![w]\!] := \{[w]_\alpha \colon \alpha \in \tau\}.$$

The following lemma collects some fundamental properties of 2-acyclic Cayley frames to be used throughout.

Lemma 17. *In a 2-acyclic Cayley frame $\mathfrak{M}$ with worlds $w, w_1, \ldots, w_k$ and groups of agents $\alpha, \alpha_1, \ldots, \alpha_k \in \tau$:*

(1) If $w \in \bigcap_{1\leqslant i\leqslant k} [w_i]_{\alpha_i}$, then $\bigcap_{1\leqslant i\leqslant k} [w_i]_{\alpha_i} = [w]_\beta$ for $\beta := \bigcap_{1\leqslant i\leqslant k} \alpha_i$.

(2) If $[w]_\alpha \in \bigcap_{1\leqslant i\leqslant k} [\![w_i]\!]$, then also $[w]_{\alpha'} \in \bigcap_{1\leqslant i\leqslant k} [\![w_i]\!]$ for any $\alpha' \in \tau$ s.t. $\alpha \subseteq \alpha'$.

(3) The set $\bigcap_{1\leqslant i\leqslant k} [\![w_i]\!]$ has a minimal element, i.e. there is an $\alpha_0 \in \tau$ s.t. $[w_1]_{\alpha_0} \in \bigcap_{1\leqslant i\leqslant k} [\![w_i]\!]$ and, for any $\alpha' \in \tau$: $[w_i]_{\alpha'} \in \bigcap_{1\leqslant i\leqslant k} [\![w_i]\!]$ iff $\alpha_0 \subseteq \alpha'$.

Proof. (1) and (2) are obvious from the definitions. For (3) we observe that 2-acyclicity implies that the collection

$$\{\alpha \in \tau \colon [w_1]_\alpha \in \textstyle\bigcap_{1\leqslant i\leqslant k} [\![w_i]\!]\}$$

is closed under intersections: otherwise there would be $[w_1]_\alpha, [w_1]_\beta \in \bigcap_{1\leqslant i\leqslant k} [\![w_i]\!]$ and $1 \leqslant i, j \leqslant k$ such that $[w_i]_{\beta\cap\alpha} \cap [w_j]_{\alpha\cap\beta} = \emptyset$. But this, together with $[w_1]_\alpha = [w_i]_\alpha = [w_j]_\alpha$ and $[w_1]_\beta = [w_i]_\beta = [w_j]_\beta$ would constitute a 2-cycle $(w_i, \alpha, w_j, \beta, w_i)$. ☐

In a 2-acyclic Cayley frame we denote the *minimal group of agents that connects the worlds in* $\mathbf{w}$ by $\operatorname{agt}(\mathbf{w}) \in \tau$ (as justified by (3) in the lemma). Intuitively, $\operatorname{agt}(\mathbf{w})$ sets the scale for zooming in to the link structure among $\mathbf{w}$ to $\alpha \subsetneq \operatorname{agt}(\mathbf{w})$. In 2-acyclic frames, it can be controlled in a regular manner.

Lemma 18. *In a 2-acyclic Cayley frame for worlds w, v:*

(i) For every agent $a \notin \operatorname{agt}(w, v)$ and every $v' \in [v]_a \setminus \{v\}$ we have $\operatorname{agt}(w, v') = \operatorname{agt}(w, v) \cup \{a\}$.

(ii) For every agent $a \in \operatorname{agt}(w, v)$ there is at most one $v' \in [v]_a$ such that $\operatorname{agt}(w, v') = \operatorname{agt}(w, v) \setminus \{a\}$.

As mentioned above, our upgrading has to deal with paths on all scales simultaneously. The real challenge in playing the pebble game lies in controlling all short paths. In sufficiently acyclic Cayley structures we only need to consider paths that link two worlds w, v via α-edges for $\alpha \subsetneq \operatorname{agt}(v, w)$; all other cases are trivial, as we shall show later. The quintessential kind of short paths to be controlled is isolated in the following definition of *non-trivial coset paths*.

Definition 19. Let $\mathfrak{M}$ be a 2-acyclic Cayley frame. A *coset path of length ℓ* from w_1 to $w_{\ell+1}$ is a labelled path $(w_1, \alpha_1, w_2, \alpha_2, \ldots, \alpha_\ell, w_{\ell+1})$ in $\mathfrak{M}$ such that, for $1 \leqslant i \leqslant \ell$: $[w_i]_{\alpha_{i-1}\cap\alpha_i} \cap [w_{i+1}]_{\alpha_i\cap\alpha_{i+1}} = \emptyset$, with $\alpha_0 = \alpha_{\ell+1} = \emptyset$.
A coset path $(w_1, \alpha_1, \ldots, \alpha_\ell, w_{\ell+1})$ in a 2-acyclic Cayley frame is *non-trivial* if $\alpha_i \subsetneq \operatorname{agt}(w_1, w_{\ell+1})$, for all $1 \leqslant i \leqslant \ell$.
A non-trivial coset path from v to $w \neq v$ is *minimal* if there is no shorter non-trivial coset path from v to w.

As we shall later show, the problem of avoiding short non-trivial coset paths does not pose great difficulties in Cayley frames that are fully acyclic in the sense of Definition 14; this will be dealt with in Section IV. Avoidance of non-trivial short paths in sufficiently acyclic *finite* Cayley frames, on the other hand, is the main focus of Sections V-A and V-B.

B. Coset acyclicity and hypergraph acyclicity

In this section we introduce the dual hypergraph of a Cayley frame or structure, and investigate the connections between acyclicity of Cayley frames and hypergraph acyclicity, and between coset paths in Cayley frames and chordless paths in hypergraphs. The dual hypergraph is an important part in describing the winning strategy for player **II** in the Ehrenfeucht–Fraïssé game. First, some basic notions.

Definition 20. A *hypergraph* is a structure $\mathcal{A} = (A, S)$, A its vertex set and $S \subseteq \mathcal{P}(A)$ its set of hyperedges.

With a hypergraph $\mathcal{A} = (A, S)$ we associate its *Gaifman graph* $G(\mathcal{A}) = (A, G(S))$; the undirected edge relation of $G(\mathcal{A})$ links $a \neq a'$ if $a, a' \in s$ for some $s \in S$.

An *n-cycle* in a hypergraph is cycle of length n in its Gaifman graph, and an *n-path* is a path of length n in its Gaifman graph. A *chord* of an n-cycle or n-path is an edge between vertices that are not next neighbours along the cycle or path. The following definition of hypergraph acyclicity is the classical one from [2], also known as α-acyclicity in [1].

Definition 21. A hypergraph $\mathcal{A} = (A, S)$ is *acyclic* if it is *conformal* and *chordal*:

(i) conformality requires that every clique in the Gaifman graph $G(\mathcal{A})$ is contained in some hyperedge $s \in S$;
(ii) chordality requires that every cycle in the Gaifman graph $G(\mathcal{A})$ of length greater than 3 has a chord.

For $N \geqslant 3$, $\mathcal{A} = (A, S)$ is *N-acyclic* if it is *N-conformal* and *N-chordal*:

(iii) N-conformality requires that every clique in $G(\mathcal{A})$ up to size N is contained in some hyperedge $s \in S$;
(iv) N-chordality requires that every cycle in $G(\mathcal{A})$ of length greater than 3 and up to N has a chord.

If a hypergraph is N-acyclic, then every induced substructure of size up to N is acyclic.

A hypergraph (A, S) is *tree decomposable* if it admits a tree decomposition $\mathcal{T} = (T, \delta)$: T is a tree and $\delta\colon T \to S$ is a map such that $\mathrm{image}(\delta) = S$ and, for every node $a \in A$, the set $\{v \in T : a \in \delta(v)\}$ is connected in T. A well-known result from classical hypergraph theory ([2],[1]) is that a hypergraph is tree decomposable if, and only if, it is acyclic.

Definition 22. Let $\mathfrak{M} = (W, (R_\alpha)_{\alpha\in\tau})$ be a Cayley frame. Its *dual hypergraph* is the vertex-coloured hypergraph

$$\begin{aligned} d(\mathfrak{M}) &:= (d(W), S, (Q_\alpha)_{\alpha\in\tau}) \text{ where} \\ d(W) &:= \dot{\bigcup}_{\alpha\in\tau} Q_\alpha \quad \text{for } Q_\alpha := W/R_\alpha, \\ S &:= \{[\![w]\!] \subseteq d(W) : w \in W\}. \end{aligned}$$

The notions of acyclicity for Cayley frames and hypergraph acyclicity are directly connected by the following.

Lemma 23. *[15] For $N \geqslant 3$, if $\mathfrak{M}$ is an N-acyclic Cayley frame, then $d(\mathfrak{M})$ is an N-acyclic hypergraph.*

Simple arguments also show that similarly the notions of non-trivial coset paths in a 2-acyclic Cayley frame and of chordless paths in its dual hypergraph are closely related as follows. A minimal non-trivial coset path

$$(w_1, \alpha_1, w_2, \ldots, w_\ell, \alpha_\ell, w_{\ell+1})$$

from w_1 to $w_{\ell+1}$ in a 2-acyclic Cayley frame $\mathfrak{M}$ induces, for $t = [\![w_1]\!] \cap [\![w_{\ell+1}]\!]$, a chordless path $[w_1]_{\alpha_1}, [w_2]_{\alpha_2}, \ldots, [w_\ell]_{\alpha_\ell}$ in $d(\mathfrak{M}) \setminus t$. A minimal (hence chordless) path

$$([w_1]_{\alpha_1}, [\![w_2]\!], [w_2]_{\alpha_2}, \ldots, [\![w_\ell]\!], [w_\ell]_{\alpha_\ell})$$

from $[w_1]_{\alpha_1}$ to $[w_\ell]_{\alpha_\ell}$ in $d(\mathfrak{M})$ that stays clear of $t := [\![w_1]\!] \cap [\![w_\ell]\!]$, induces a non-trivial coset path $(w_1, \alpha_1, w_2, \ldots, w_{\ell-1}, \alpha_{\ell-1}, w_\ell)$ in $\mathfrak{M}$.

IV. If finiteness were not an issue

In this section we prove the classical version of the characterisation theorem, Theorem 4, for $\mathsf{ML}[\mathsf{CK}]$ over the class of all (finite or infinite) CK-structures, which by Lemma 11 is represented up to $\sim$ by the class of all Cayley structures.

A. Free ω-unfoldings

For the upgrading argument in the spirit of Figure 1 we need to construct bisimilar companions to Cayley structures that avoid features that can be distinguished by FO-formula up to fixed quantifier rank q. As discussed in Section II-D, these concern small multiplicities and short non-trivial cycles. In the classical case we may use infinitely branching tree unfoldings, which may here directly be generated as bisimilar coverings by Cayley structures similar to Lemma 11. The differences are that here we want to put countably many generators for every edge of the original CK-frame (to boost multiplicities) and do not impose any non-trivial identities between generator words (to achieve full acyclicity). In other words, unlike for Lemma 11, we here use *free* Cayley groups, whose Cayley graphs are actual trees, i.e., acyclic. So, in order to boost all multiplicities uniformly to ω, we just replace the partition set of E_a-generators, which previously was the set of a-tagged individual irreflexive R_a-edges, by

$$E_a^\omega := R_a \times \omega \times \{a\}.$$

Picking any distinguished world w in the given CK-structure $\mathfrak{M}$ as a base point, we may proceed as in Lemma 11 to obtain a bisimilar covering of $\mathfrak{M}$ by a Cayley structure $\mathfrak{M}^\omega[w]$ that is acyclic in the sense of Definition 14 (and associates the identity element of the underlying Cayley group with w). Besides full acyclicity, this bisimilar covering by the infinite Cayley structure $\mathfrak{M}^\omega[w]$ has, for every α-neighbour u of any element w, $u \in [w]_\alpha$, countably infinitely many sibling worlds $u' \in [w]_\alpha$ that are bisimilar to u: we call this *ω-richness*. So we may, up to $\sim$ in the world of all (finite and infinite!) CK-structures, assume to be dealing with such *free*, *ω-rich* Cayley structures, as can be obtained as free ω-unfoldings of arbitrary CK-structures.

Proviso 24. *For the remainder of this section, we consider free and ω-rich Cayley structures only.*

B. Upgrading in free, ω-rich Cayley structures

We aim to show that for some suitable choice of $\ell = \ell(q)$, any free, ω-rich Cayley structures $\mathfrak{M}$ and $\mathfrak{N}$ satisfy:

$$\mathfrak{M}, w \sim^{\ell} \mathfrak{N}, v \;\Rightarrow\; \mathfrak{M}, w \equiv_q \mathfrak{N}, v.$$

We prove the statement by giving a winning strategy for player **II** in the standard q-round Ehrenfeucht–Fraïssé game on $\mathfrak{M}, w$ and $\mathfrak{N}, v$. In order to find suitable responses for **II** we use the dual hypergraphs $d(\mathfrak{M})$ and $d(\mathfrak{N})$ (cf. Definition 22) associated with $\mathfrak{M}$ and $\mathfrak{N}$ and employ techniques from [15] that were developed for playing Ehrenfeucht–Fraïssé games on α-acyclic hypergraphs. These techniques are critical in showing that the necessary level of $\sim^{\ell}$ can be bounded in terms of q.

To win the q-round Ehrenfeucht–Fraïssé game on $\mathfrak{M}$ and $\mathfrak{N}$ player **II** upholds the following invariant throughout her play. In round $i-1$ she finds substructures $\mathfrak{M}_{i-1} \subseteq \mathfrak{M}$ and $\mathfrak{N}_{i-1} \subseteq \mathfrak{N}$ that contain the pebbled worlds and an isomorphism $f\colon \mathfrak{M}_{i-1} \to \mathfrak{N}_{i-1}$ such that $\mathfrak{M}_{i-1}, u \sim^{\ell_{i-1}} \mathfrak{N}_{i-1}, f(u)$, for all $u \in \mathfrak{M}_{i-1}$ and a degree $\sim^{\ell_{i-1}}$ of bisimulation equivalence that depends on the current round i, and thus the number of rounds still to be played. Simultaneously, she keeps track of two subsets $Q_{i-1} \subseteq d(\mathfrak{M})$ and $Q'_{i-1} \subseteq d(\mathfrak{N})$ that induce isomorphic substructures $d(\mathfrak{M}) \restriction Q_{i-1}$ and $d(\mathfrak{N}) \restriction Q'_{i-1}$, as well as tree decompositions $\mathcal{T}_{i-1}$ and $\mathcal{T}'_{i-1}$ of these. The substructures $\mathfrak{M}_{i-1}$ and $\mathfrak{N}_{i-1}$ will be representations, in $\mathfrak{M}$ and $\mathfrak{N}$, of these two tree decompositons. Conversely, and roughly speaking, Q_{i-1} and Q'_{i-1} are the dual hypergraph images of $\mathfrak{M}_{i-1}$ and $\mathfrak{N}_{i-1}$.

If **I** places a pebble in $\mathfrak{M}$ say, then the structures $\mathfrak{M}_{i-1}$ and Q_{i-1} get expanded in a certain manner, and an analysis of these expansions tells **II** how to respond in $\mathfrak{N}$. Essentially, if **I** places a pebble on a world $w' \in \mathfrak{M}$ in round i, we have to add w' to the substructure $\mathfrak{M}_{i-1}$; for every equivalence class $[w'']_\alpha$ that is on a short coset path between $\mathfrak{M}_{i-1}$ and w', we have to add a representative w'' to $\mathfrak{M}_{i-1}$; and for every equivalence class on a short coset path between $\mathfrak{M}_{i-1}$ and the new vertices we have to add yet another representative and so forth. For the argument to work we need to figure out which representatives to choose and how to uniformly bound the number of vertices added to $\mathfrak{M}_{i-1}$.

We have seen above that short coset paths in a Cayley structure correspond to short chordless paths in its dual hypergraph. For this reason we investigate subsets of acyclic hypergraphs that are closed under short chordless paths.

Definition 25. Let $\mathcal{A} = (A, S)$ be a hypergraph.

(i) A subset $Q \subseteq A$ is *m-closed* if every chordless path up to length m between nodes $a, a' \in Q$ is contained in Q.

(ii) For $m \in \mathbb{N}$, the *convex m-closure* of a subset $P \subseteq A$ is the minimal m-closed subset that contains P:

$$\mathrm{cl}_m(P) := \bigcap \{ Q \supseteq P \colon Q \subseteq A \ m\text{-closed } \}.$$

In our Ehrenfeucht–Fraïssé game on $\mathfrak{M}$ and $\mathfrak{N}$, the auxiliary set $Q_{i-1} \subseteq d(\mathfrak{M})$ (as a record of $\mathfrak{M}_{i-1}$ in the dual) is chosen to be $2m+1$-closed, for some $m = m_i \in \mathbb{N}$ that is considered a short distance for round i. If **I** chooses a world w' for his move in round i, we take a closer look at the structure of the set $Q_i := \mathrm{cl}_m(Q_{i-1} \cup \{[w']_\emptyset\})$. Lemmas 26 and 27 show that in sufficiently acyclic hypergraphs, the addition of $[w']_\emptyset = \{w'\}$ to Q_{i-1} and closure of this set under short chordless paths updates Q_{i-1} to Q_i in a well-behaved manner. The following three lemmas are from [15]. Here $N^1(P) = \bigcup \{N^1(p) \colon p \in P\}$ refers to the 1-neighbourhood of the set P in the Gaifman graph; distance $d(P, q) = \min\{d(p, q) \colon p \in P\}$ between a set and a vertex similarly refers to distance in the Gaifman graph.

Lemma 26. *Let $\mathcal{A} = (A, S)$ be sufficiently acyclic, $m > 1$, $Q \subseteq A$ m-closed, $a \in A$ some vertex with $1 \leqslant d(Q, a) \leqslant m$. Let $\hat{Q} := \mathrm{cl}_m(Q \cup \{a\})$ and consider the region in which this extended closure attaches to Q:*

$$D := Q \cap N^1(\hat{Q} \setminus Q).$$

Then $\hat{Q} \setminus Q$ is connected, and D separates $\hat{Q} \setminus Q$ from $Q \setminus D$, hence $\hat{Q} = Q \cup \mathrm{cl}_m(D \cup \{a\})$.

Since we additionally assumed that Q_{i-1} is $2m+1$-closed, we can employ the following lemma, as well.

Lemma 27. *Let $\mathcal{A} = (A, S)$ be sufficiently acyclic, $Q \subseteq A$ m-closed, $a \in A$ some vertex with $1 \leqslant d(Q, a) \leqslant m$, and $\hat{Q} := \mathrm{cl}_m(Q \cup \{a\})$. If Q is even $(2m+1)$-closed, then $D = Q \cap N^1(\hat{Q} \setminus Q)$ is a clique.*

D being a clique implies that it is contained in a single bag of the tree decomposition of $d(\mathfrak{M}) \restriction Q_{i-1}$. This bag corresponds to some world u in $\mathfrak{M}_{i-1}$, in the sense that $D \subseteq [\![u]\!]$, which is mapped to $f(u) \in \mathfrak{N}_{i-1}$. This is our starting point for finding **II**'s response in $\mathfrak{N}$ to **I**'s move to w' in $\mathfrak{M}$.

Now that we know $Q_i = Q_{i-1} \cup \mathrm{cl}_m(D \cup \{[w']_\emptyset\})$, for some clique $D \subseteq Q_{i-1}$, we want to obtain a bound on the size of the extension $\mathrm{cl}_m(D \cup \{[w']_\emptyset\})$ that can occur in a single round; such a bound is critical in bounding the required level of $\sim^{\ell}$. Since our (dual) hypergraphs have a uniform width of $|\tau|$, which we regard as constant, we seek functions $f_m(k)$ that bound the size of m-closures of sets or tuples of size k in those hypergraphs, provided they are sufficiently acyclic.

Lemma 28. *For fixed width, there are functions $f_m(k)$ s.t. in hypergraphs $\mathcal{A}$ of that width that are sufficiently acyclic, $|\mathrm{cl}_m(P)| \leqslant f_m(k)$ for all $P \subseteq A$ of size $|P| \leqslant k$.*

The closures that we encounter in a single step of the Ehrenfeucht–Fraïssé game will be generated by a hyperedge and a single extra vertex, hence by at most $|\tau| + 1$ vertices. We can therefore bound the size of these closures, which also allows us to bound the depth of tree decompositions of the relevant sub-hypergraphs induced by the closures. Availability of suitable tree-like extensions of bounded depth, which can be used to cover the newly pebbled worlds, is controlled by bisimulation types of corresponding depth.

So we aim to describe the tree decomposition $\mathcal{T} = (T, \delta)$ of $d(\mathfrak{M}) \restriction Q$, for $Q := (Q_i \setminus Q_{i-1}) \cup D$ by an ML-formula. Since every bag $\delta(u)$ is a clique, it is covered a hyperedge $[\![w_u]\!]$ of $d(\mathfrak{M})$. Let $\hat{\delta}\colon V[T] \to \mathfrak{M}$ be a map such that $\delta(u) \subseteq [\![\hat{\delta}(u)]\!]$ for all $u \in V[T]$. In general, there is not a unique choice for $\hat{\delta}$, but we know $w' \in \mathrm{image}(\hat{\delta})$, since $[\![w']\!]$ is the only hyperedge that includes $[w']_\emptyset$. The set $\mathrm{image}(\hat{\delta})$ contains the vertices to be added to $\mathfrak{M}_{i-1}$ in round i, i.e. we define $\mathfrak{M}_i$ as the substructure of $\mathfrak{M}$ that is induced by $W_{i-1} \cup \mathrm{image}(\hat{\delta})$ – and we want to find a corresponding extension of $\mathfrak{N}_{i-1}$ through a description of $\mathcal{T} = (T, \delta)$ in ML. To prepare for that, we observe that in dual hypergraphs of 2-acyclic Cayley structures the set of equivalence classes $[\![u]\!] \cap [\![u']\!]$ is fully determined by the single set of agents $\mathrm{agt}(u, u')$ (cf. Lemma 17).

Let $w_u := \hat{\delta}(u)$, for $u \in V[T]$, and and $\lambda \in V[T]$ the vertex with $D \subseteq \delta(\lambda)$. We regard λ as the root of T. We describe the existential ML-type of $\mathfrak{M} \restriction \mathrm{image}(\hat{\delta}), w_\lambda$ by a formula $\varphi_\mathcal{T} := \varphi_{\mathcal{T},\lambda} \in \mathsf{ML}$. For every vertex $u \in V[T]$ we define a formula $\varphi_{\mathcal{T},u}$ by induction on the depth of u in T:

— For a leaf u let $\varphi_{\mathcal{T},u}$ be the formula of modal depth 0 that describes the atomic type of $w_u \in \mathfrak{M}$.

— For a non-leaf u with children $u_1, \ldots, u_k$ and associated formulae $\varphi_{\mathcal{T},u_j}$, let $\alpha_j = \mathrm{agt}(w_u, w_{u_j})$. Let $\chi \in \mathrm{ML}_0$ be the formula that describes the atomic type of w_u, then

$$\varphi_{\mathcal{T},u} := \chi \wedge \textstyle\bigwedge_j \Diamond_{\alpha_j} \varphi_{\mathcal{T},u_j}.$$

Note that the modal nesting depth of $\varphi_\mathcal{T}$ is uniformly bounded by the depth of T, which in turn was bounded by the size of the relevant m-closures. This means that values ℓ_{i-1} for the bisimulation level that needs to be respected by the isomorphism $f\colon \mathfrak{M}_{i-1} \simeq \mathfrak{N}_{i-1}$ can be chosen such that in round i of the game, $\varphi_\mathcal{T}$ is preserved by f. So, as clearly $\mathfrak{M}_i, w_\lambda \models \varphi_\mathcal{T}$ by construction, also $\mathfrak{N}_i, f(w_\lambda) \models \varphi_\mathcal{T}$. This will eventually allow us to expand $\mathfrak{N}_{i-1} \subseteq \mathfrak{N}$ to $\mathfrak{N}_i$ to keep our invariant alive. But there are two further obstacles:

Consider a subformula $\Diamond_\alpha \varphi_{\mathcal{T},u}$ of $\varphi_{\mathcal{T},\lambda}$. By construction, $\alpha = \mathrm{agt}(w_\lambda, w_u)$, and $\mathfrak{N}_i, f(w_\lambda) \models \Diamond_\alpha \varphi_{\mathcal{T},u}$ only implies that there is a suitable world $v_u \in [f(w_\lambda)]_\alpha$, but we might only have $\mathrm{agt}(f(w_\lambda), v_u) \subsetneq \alpha$. This is the first obstacle.

The second one comes in the form of undesirable non-trivial short paths. Assuming we had overcome the first obstacle, there might still be non-trivial short paths between $f(w_\lambda)$ and v_u. On the other hand, such paths do not exist between w_λ and w_u in $\mathfrak{M}$: as discussed in connection with Lemma 23, a short non-trivial coset path would induce a short chordless path between $[\![w_\lambda]\!]$ and $[\![w_u]\!]$ that is incompatible with the m-closed nature of Q.

Both problems can be solved in the same vein. We need to find a world x that is bisimilar to v_u and has the right reachability properties in relation to $f(w_\lambda)$. In acyclic Cayley frames this is easily done because coset paths between two vertices are in some sense unique.

Remark 29. *An acyclic Cayley frame $\mathfrak{M}$ has unique coset paths between any two vertices w and v in the single agent reduct $(W, (R_a)_{a \in \Gamma})$. Any coset path between w and v in $\mathfrak{M}$ is a contraction of this unique path.*

Recall that $\alpha = \mathrm{agt}(w_\lambda, w_u)$, so that non-trivial short paths from $f(w_\lambda)$ to v_u are based on sets $\beta \subsetneq \alpha$. Starting at v_u, we move to bisimilar copies of v_u via a-edges for $a \in \alpha$ that are not on the path between $f(w_\lambda)$ and v_u. If we move through all the agents $a \in \alpha$ sufficiently often, we arrive at a copy x of v_u that is not linked to $f(w_\lambda)$ by any non-trivial coset paths of length up to m. This is due to the fact that every set of agents $\beta \subsetneq \alpha$ that gets used in a non-trivial coset path can contain at most $|\alpha| - 1$ agents. Hence, after cycling through all the agents in α sufficiently many times, we need more than m edges to arrive at x on a non-trivial coset path. We thus obtain the following lemma.

Lemma 30. *Let $\mathfrak{M}$ be an acyclic, ω-rich Cayley structure, $\mathrm{agt}(w, v) \subseteq \alpha$ for some $w, v \in W$, and $m \in \mathbb{N}$. Then, there is a world $v' \sim v$ with $\mathrm{agt}(w, v') = \alpha$ that is not linked to w by non-trivial coset paths of length up to m.*

Lemma 30 allows us to expand $\mathfrak{N}_{i-1}$ to $\mathfrak{N}_i$ and Q'_{i-1} to Q'_i such that our invariant still holds after round i, because we can exactly replicate every intersection of hyperedges in Q_i and respect the propositional assignments. Furthermore, ω-richness allows us use a fresh branch of $\mathfrak{N}$ for the new part of $\mathfrak{N}_i$. The following sums up the results of this section.

Lemma 31. *For some suitable choice of $\ell = \ell(q)$, any free, ω-rich Cayley structures $\mathfrak{M}$ and $\mathfrak{N}$ satisfy:*

$$\mathfrak{M}, w \sim^\ell \mathfrak{N}, v \;\Rightarrow\; \mathfrak{M}, w \equiv_q \mathfrak{N}, v.$$

Proof. Let $(m_k)_{0 \leqslant k \leqslant q}$ and $(\ell_k)_{0 \leqslant k \leqslant q}$ be two sequences with $m_q \geqslant 2$, $m_{k-1} \geqslant 2m_k + 1$ and $\ell_q \geqslant 1$, $\ell_{k-1} \geqslant \ell_k + f_{m_k}(|\tau| + 1)$, where f_m is the bound on m-closures in sufficiently acyclic hypergraphs from Lemma 28. Put $\ell := \ell_0$.

We give a winning strategy for **II** in the q-round Ehrenfeucht–Fraïssé game on $\mathfrak{M}, w$ and $\mathfrak{N}, v$. She must uphold the following invariant: assume that after $0 \leqslant i \leqslant q$ rounds the pebbled vertices are $w_0, w_1, \ldots, w_i \in \mathfrak{M}$ and $v_0, v_1, \ldots, v_i \in \mathfrak{N}$, with $w_0 = w$ and $v_0 = v$. Then, there are substructures $\mathfrak{M}_i \subseteq \mathfrak{M}$ and $\mathfrak{N}_i \subseteq \mathfrak{N}$ that contain the pebbled vertices and an isomorphism $f_i\colon \mathfrak{M}_i \to \mathfrak{N}_i$ such that $\mathfrak{M}_i, u \sim^{\ell_i} \mathfrak{N}_i, f(u)$, for all $u \in \mathfrak{M}_i$, and $f_i(w_j) = v_j$, for all $0 \leqslant j \leqslant i$. Furthermore, there are m_i-closed subsets $Q_i \subseteq d(\mathfrak{M})$ and $Q'_i \subseteq d(\mathfrak{N})$ and tree decompositions $\mathcal{T}_i = (T_i, \delta_i)$ and $\mathcal{T}'_i = (T'_i, \delta'_i)$ of $d(\mathfrak{M}) \restriction Q_i$ and $d(\mathfrak{N}) \restriction Q'_i$, and mappings $\hat{\delta}_i\colon T_i \to \mathfrak{M}_i$ and $\hat{\delta}'_i\colon T'_i \to \mathfrak{N}_i$ with $\delta_i(u) \subseteq [\![\hat{\delta}_i(u)]\!]$, for all $u \in T_i$, and $\delta'_i(u) \subseteq [\![\hat{\delta}'_i(u)]\!]$, for all $u \in T'_i$.

Obviously, the invariant holds before the first round for $\mathfrak{M}_0 \subseteq \mathfrak{M}$ induced by $\{w_0\}$, $\mathfrak{N}_0 \subseteq \mathfrak{N}$ induced by $\{v_0\}$, $Q_0 = \{[\![w_0]\!]\}$ and $Q'_0 = \{[\![v_0]\!]\}$ because $\mathfrak{M}, w_0 \sim^{\ell_0} \mathfrak{N}, v_0$.

Assume the invariant after completion of the $(i-1)$-st round. In order to show that **II** can update the invariant for another round, we work in the expansions $\mathfrak{M}^{\ell_i}$ and $\mathfrak{N}^{\ell_i}$ of $\mathfrak{M}$ and $\mathfrak{N}$ by predicates that mark the ℓ_i-types of vertices. This means that f_{i-1} is compatible with $\sim^{\ell_{i-1} - \ell_i}$ over the expansions, and we need f_i to be just a local isomorphism w.r.t. these expansions.

W.l.o.g. we may assume that player **I** chooses a world $w_i \in \mathfrak{M}$. Put $Q_i := \mathrm{cl}_{m_i}(Q_{i-1} \cup \{[w_i]_\emptyset\})$. Since Q_{i-1} is $(2m_i+1)$-closed and $\mathfrak{M}$ is acyclic, Lemmas 26 and 27 imply that $Q_i = Q_{i-1} \cup \mathrm{cl}_{m_i}(D \cup \{[w_i]_\emptyset\})$, $Q_i \setminus Q_{i-1}$ is connected, and D separates $Q_i \setminus Q_{i-1}$ from $Q_{i-1} \setminus D$, for $D = Q_{i-1} \cap N^1(Q_i \setminus Q_{i-1})$. Lemma 28 implies $|Q_i \setminus Q_{i-1}| \leqslant |\mathrm{cl}_{m_i}(D \cup \{[w_i]_\emptyset\})| \leqslant f_{m_i}(|\tau|+1)$. We extend $\mathcal{T}_{i-1}$ to $\mathcal{T}_i$ in a straightforward manner and choose a suitable extension of $\hat{\delta}_{i-1}$ to $\hat{\delta}_i$ for all of $\mathcal{T}_i$. $\mathfrak{M}_i$ is the substructure that is induced by $\mathrm{image}(\hat{\delta}_i)$. Let $\lambda \in V[T_i]$ be the vertex with $D \subseteq \delta_i(\lambda)$, $w_\lambda = \hat{\delta}_{i-1}$, and $T_\lambda \subseteq T_i$ the subtree that represents $\mathrm{cl}_{m_i}(D \cup \{[w_i]_\emptyset\})$. We know that T_λ has at most depth $f_{m_i}(|\tau|+1) \leqslant \ell_{i-1} - \ell_i$, which means that the ML-formula φ_i that describes it has at most modal depth $\ell_{i-1} - \ell_i$. Therefore, $\mathfrak{N}_{i-1}^{\ell_i}, f_{i-1}(w_\lambda) \models \varphi_i$ and we may expand $\mathfrak{N}_{i-1}$ to $\mathfrak{N}_i$ according to φ_i in a fresh branch that starts at $f_{i-1}(w_\lambda)$ such that the distances on all scales are exactly the same as in $\mathfrak{M}_i$ (Lemma 30). This gives us an extension of f_{i-1} to $f_i \colon \mathfrak{M}_i \to \mathfrak{N}_i$ and **II**'s move goes to $v_i := f_i(w_i)$. If we correspondingly extend Q'_{i-1}, $\mathcal{T}'_{i-1}$ and $\hat{\delta}'_{i-1}$ to $\mathfrak{N}_i$, then Q'_i is m_i-closed and the substructures induced by Q_i and Q'_i are isomorphic, because we made sure that there are no non-trivial coset paths in $\mathfrak{N}_i$ that are not in $\mathfrak{M}_i$. Thus, **II** can keep the invariant alive for q rounds and wins the game. □

The previous lemma provides the upgrading needed for our main theorem over the class of all (finite or infinite) CK-structures (cf. Figure 1 and its discussion in Section II-D):

Theorem 32. $\mathsf{ML}[\mathsf{CK}] \equiv \mathsf{FO}/{\sim}$ *over the class of all Cayley structures, and hence over the class of all CK-structures.*

V. Where finiteness is essential

In this section we obtain the finite model theory version of our modal characterisation theorem. The argument follows the same pattern inspired by Figure 1: we upgrade ℓ-bisimulation equivalence to $\equiv_q$-equivalence for some ℓ that depends on q. But this time the upgrade has to be conducted over finite Cayley structures. This makes both the construction of the bisimilar companions and the analysis of the Ehrenfeucht–Fraïssé game considerably more involved.

The following main lemma is based on the same construction as Lemma 11 using compatible finite Cayley groups from [15] with the desired level of coset acyclicity after straightforward pre-processing to boost multiplicities.

Lemma 33. *Any connected finite CK-structure, and in particular any Cayley structure admits finite bisimilar coverings by Cayley structures that realise any given lower bounds for acyclicity and multiplicities.*

A. Structure theory for N-acyclic Cayley groups

Definition 34. Let $\mathfrak{M}$ be a Cayley frame that is $2N$-acyclic. We call a coset path *short* if its length is $\leqslant N$.

The main obstacle in the upgrading is to avoid short coset paths. To be more precise, for a pair of worlds w, v and a tuple of other worlds $\mathbf{z}$ without short non-trivial coset paths to v, we want to find a world $v' \sim v$, still without short non-trivial short links to $\mathbf{z}$, and such that there is also no short non-trivial coset path from w to v'.[5] In free ω-unfoldings this task could be accomplished because we had, at every world, an unbounded supply of copies of each branch and could thus move arbitrarily far away in an independent direction. Finite structures do not afford quite that luxury. But in a finite Cayley frame of sufficiently high finite acyclicity and richness, we can still repeatedly find suitable bisimilar variants v' that are good enough to win the q-round game. We call this property *freeness* and will define it formally in Definition 37. The main lemma of this section will be Lemma 38: it states that sufficiently acyclic Cayley structures with sufficiently high multiplicities have sufficiently free dual hypergraphs.

But first we treat a technical matter that also gives a glimpse of an intriguing structure theory for N-acyclic Cayley frames.

Very little can be said in general about two paths joining the same pair of vertices in arbitrary finite Cayley structures, or even in highly acyclic ones. Two *short* paths, however, will necessarily overlap in an interesting fashion if the structures are sufficiently acyclic. In fact, here short paths will be essentially unique, or just minor and local variations of one another; the following two lemmas make this precise. We first look at coset paths from one world to itself – a relaxation of the notion of a coset cycle in that it omits the condition on trivial intersection at the point of return.

Lemma 35. *Let w be a world in a Cayley frame $\mathfrak{M}$. If $\mathfrak{M}$ is N-acyclic, then there is no coset path of length up to N that starts at w and ends at w.*

Proof. The claim is shown by induction on the length m of the coset path, for $1 \leqslant m \leqslant N$. For $m = 1$, Definition 19 obviously rules out coset loops of the form (w, α, w).

Inductively, assume there are no coset paths in $\mathfrak{M}$ of length $1 \leqslant m < N$ from any world back to itself. Let $w \in W$ and consider a coset path

$$(w_1, \alpha_1, w_2, \dots, w_{m+1}, \alpha_{m+1}, w_{m+2})$$

of length $m+1$ with $w_1 = w_{m+2} = w$. N-acyclicity of $\mathfrak{M}$ implies $[w_1]_{\alpha_{m+1} \cap \alpha_1} \cap [w_2]_{\alpha_1 \cap \alpha_2} \neq \emptyset$ or $[w_1]_{\alpha_1 \cap \alpha_{m+1}} \cap [w_{m+1}]_{\alpha_{m+1} \cap \alpha_m} \neq \emptyset$. W.l.o.g. we assume that there is some $w' \in [w_1]_{\alpha_1 \cap \alpha_{m+1}} \cap [w_{m+1}]_{\alpha_{m+1} \cap \alpha_m}$. If $w' \notin [\alpha_2]_{\alpha_1 \cap \alpha_2}$, then $(w', \alpha_1, w_2, \dots, w_m, \alpha_m, w')$ is a coset path of length m that starts and ends at w'. Otherwise $(w', \alpha_2, w_3, \dots, w_m, \alpha_m, w')$ is a coset path of length $m-1$ that starts and ends at w'. In both cases, such a coset path cannot exist according to the induction hypothesis. □

From Lemma 35 we obtain the following lemma, a quasi-uniqueness property for short coset paths.

Lemma 36 (Zipper Lemma)**.** *Let $\mathfrak{M}$ be a $2N$-acyclic Cayley frame $w, v \in W$, and $(w, \alpha_1, u_1, \alpha_2, \dots, \alpha_\ell, v)$ and*

[5] In context, the tuple $\mathbf{z}$ will represent the previously pebbled worlds, and v' is meant to be made 'independent' from w.

$(w, \beta_1, r_1, \beta_2, \ldots, \beta_k, v)$ *be two coset paths from w to v of length up to N. Then*

(1) $[w]_{\beta_1\cap\alpha_1} \cap [u_1]_{\alpha_1\cap\alpha_2} \neq \emptyset$ *or* $[w]_{\beta_1\cap\alpha_1} \cap [r_1]_{\beta_2\cap\beta_1} \neq \emptyset$*;*

(2) $[v]_{\alpha_\ell\cap\beta_k} \cap [u_{\ell-1}]_{\alpha_{\ell-1}\cap\alpha_\ell} \neq \emptyset$ *or* $[v]_{\alpha_\ell\cap\beta_k} \cap [r_{k-1}]_{\beta_k\cap\beta_{k-1}} \neq \emptyset$*.*

Proof. Since $\mathfrak{M}$ is $2N$-acyclic we know that $[w]_{\beta_1\cap\alpha_1} \cap [u_1]_{\alpha_1\cap\alpha_2} \neq \emptyset$, or $[r_1]_{\beta_2\cap\beta_1} \cap [w]_{\beta_1\cap\alpha_1} \neq \emptyset$, or $[u_{\ell-1}]_{\alpha_{\ell-1}\cap\alpha_\ell} \cap [v]_{\alpha_\ell\cap\beta_k} \neq \emptyset$, or $[v]_{\alpha_\ell\cap\beta_k} \cap [r_{k-1}]_{\beta_k\cap\beta_{k-1}} \neq \emptyset$ occurs. Now assume we had, for instance,

$$[w]_{\beta_1\cap\alpha_1} \cap [u_1]_{\alpha_1\cap\alpha_2} = \emptyset \text{ and } [r_1]_{\beta_2\cap\beta_1} \cap [w]_{\beta_1\cap\alpha_1} = \emptyset.$$

This would imply a coset path of length $2N$ from v to itself, contradicting Lemma 35. □

Essentially, the lemma means that two short coset paths that share start and target vertex behave like a zipper that can be closed from both sides. This has some interesting consequences:

— if we have two short coset paths $(w, \alpha_1, \ldots, \alpha_\ell, v)$ and $(w, \beta_1, \ldots, \beta_k, v)$, then there is a short coset path from w to v that starts with an $(\alpha_1 \cap \beta_1)$-edge.

— if $(w_1, \alpha_1, w_2, \alpha_2, \ldots, \alpha_\ell, w_\ell)$ is a short coset path and $\alpha = \mathrm{agt}(w_1, w_\ell)$, then there are $w_i' \in [w_i]_{\alpha_{i-1}\cap\alpha_i}$, for $2 \leqslant i < \ell$, such that $(w_1, (\alpha_1 \cap \alpha), w_2', (\alpha_2 \cap \alpha), \ldots, (\alpha_\ell \cap \alpha), w_\ell)$ is a short coset path.

— if $(w_1, \alpha_1, w_2, \alpha_2, \ldots, \alpha_\ell, w_\ell)$ is a coset path in a $2N$-acyclic Cayley structure with $[w_1]_{\mathrm{agt}(w_1,w_\ell)} \cap [w_i]_{\alpha_{i-1}\cap\alpha_i} = \emptyset$ for some $2 \leqslant i < \ell$, then $\ell \geqslant N$. I.e. a coset path from w_1 to w_ℓ must be long if one of its links is disjoint from $[w_1]_{\mathrm{agt}(w_1,w_\ell)}$.

Hence, if two distinct worlds are not linked by a short non-trivial coset path, then any path between them must be long or of length 1.

B. Freeness and control of short paths

Recall the two challenges for player **II** in finding a suitable response to player **I**'s move from Section IV. For a world $v \in [w]_\alpha$, we must find a world $v' \sim v$ such that $\mathrm{agt}(w, v') = \alpha$ and such that there is no short non-trivial coset path between w and v'. Additionally, certain distance relations between v and worlds $\mathbf{z}$ must be preserved in the passage to v'. The property that allows her to find such a v' is here called *freeness*. Establishing that sufficiently acyclic Cayley structures with sufficiently high multiplicities are sufficiently free is the essential step in proving that we can maintain our invariant in the q-round pebble game on finite structures.

Freeness essentially is a property of the dual hypergraph $d(\mathfrak{M})$ of a Cayley structure $\mathfrak{M}$ (cf. Definition 16). The rôle of the worlds $\mathbf{z}$ is played by a set Z of hyperedges.

For $t, X, Y \subseteq A$ in a hypergraph $\mathcal{A} = (A, S)$, we denote as $d_t(X, Y)$ the distance between $X \setminus t$ and $Y \setminus t$ in the induced sub-hypergraph $\mathcal{A} \restriction (A \setminus t)$.

Definition 37. Let $\mathfrak{M}$ be a 2-acyclic Cayley structure and $d(\mathfrak{M}) = (d(W), S, (Q_\alpha)_{\alpha\in\tau})$ its dual hypergraph.

Let $[\![v]\!] \in S$ be a hyperedge, $Z \subseteq S$ a set of hyperedges, and $[\![w]\!] \in Z$. We say that $[\![v]\!]$ and Z are *n-free over $[\![w]\!]$* if $d_t([\![v]\!], (\bigcup Z)) > n$, where $t = [\![v]\!] \cap [\![w]\!]$.

We say that $\mathfrak{M}$ is *(n, K)-free*, for some $K \in \mathbb{N}$, if, for all $v \in W$, $[\![w]\!] \in Z \subseteq S$ with $|Z| \leqslant K$ and all sets of agents $\alpha \supseteq \mathrm{agt}(v, w)$, there is a $v' \sim v$ in $\mathfrak{M}$ such that $\mathrm{agt}(v', w) = \alpha$ and such that $[\![v']\!]$ and Z are n-free over $[\![w]\!]$.

The remainder of this section is devoted to a proof outline of the following main lemma.

Lemma 38. *Let $n, K \in \mathbb{N}$. If a Cayley structure $\mathfrak{M}$ is sufficiently acyclic and has sufficiently high multiplicities, then $d(\mathfrak{M})$ is (n, K)-free.*

Proofsketch. For the remainder of this section let $\mathfrak{M}$ be a Cayley structure with dual hypergraph $d(\mathfrak{M}) = (d(W), S, (Q_\alpha)_{\alpha\in\tau})$ that is sufficiently acyclic and has sufficiently high multiplicities, and let $[\![v]\!]$, Z, $[\![w]\!]$ and α be as above. We need to find a world $v' \sim v$ with $\mathrm{agt}(v', w) = \alpha$ such that $[\![v']\!]$ and Z are n-free over $[\![w]\!]$. Lemma 18 yields a world $v_0 \sim v$ with $\mathrm{agt}(v_0, w) = \alpha$; put $t := [\![v_0]\!] \cap [\![w]\!]$. Distances remain to be adjusted, which is achieved in two steps.

Step 1. We find $v_1 \sim v_0$ with $\mathrm{agt}(v_1, w) = \alpha$ such that $[\![v_1]\!] \cap [\![z]\!] \subseteq [\![v_1]\!] \cap [\![w]\!]$, which implies $d_t([\![v_1]\!], [\![z]\!]) \geqslant 1$, for all $[\![z]\!] \in Z$. The idea is to boost the sets $\mathrm{agt}(v_0, z)$ to $\mathrm{agt}(v_1, z)$ such that we have $\mathrm{agt}(v_1, z) \supseteq \mathrm{agt}(v_1, w)$. This can easily be done by using Lemma 18 multiple times on all worlds z with $[\![z]\!] \in Z$ in an inductive process. The second statement in Lemma 18 and sufficiently high multiplicities of $\mathfrak{M}$ ensure that it is possible to boost all the connecting sets of agents simultaneously. Then 2-acyclicity implies

$$\mathrm{agt}(v_1, z) \supseteq \mathrm{agt}(v_1, w) \Rightarrow [\![v_1]\!] \cap [\![z]\!] \subseteq [\![v_1]\!] \cap [\![w]\!].$$

Step 2. The goal of the second step is to find $v_2 \sim v_1$ with $\mathrm{agt}(v_2, w) = \alpha$ such that $d_t([\![v_2]\!], [\![z]\!]) > n$, for all $[\![z]\!] \in Z$ (while preserving the properties already achieved in Step 1). We crucially use that minimal coset paths in $\mathfrak{M}$ correspond to minimal paths in $d(\mathfrak{M})$. Assume the distance between $[\![v_1]\!] \setminus t$ and $[\![z]\!] \setminus t$ in $d(\mathfrak{M}) \setminus t$ is $\leqslant n$, for some $[\![z]\!] \in Z$. This implies a short path from a vertex of $[\![v_1]\!]$ to a vertex of $[\![z]\!]$ that runs outside $[\![v_1]\!] \cap [\![z]\!]$. Since $\mathfrak{M}$ is sufficiently acyclic this implies a non-trivial short coset path from v_1 to z. Hence, establishing $d_t([\![v_2]\!], [\![z]\!]) > n$ boils down to avoiding short non-trivial coset paths in $\mathfrak{M}$. The basic idea is the same as for Lemma 30. However, in the case of finite Cayley structures the proof is considerably more involved. In an inductive process we avoid short coset paths to the worlds z, one after the other, by moving multiple times to bisimilar copies of v_1 via a-edges for $a \in \mathrm{agt}(v_1, z)$. Cycling through all agents in $\mathrm{agt}(v_1, z)$ sufficiently often, we arrive at a suitably removed copy of v_1. The major complication in the finite case, compared to the treatment in Section IV, is that we do not have an unbounded supply of fresh unbounded branches to choose the next copy from. Hence, we must guarantee that we can move away from the current world z without getting closer to any one of the other worlds that have already been dealt with. Two key ingredients make this possible. Firstly, the Zipper lemma implies a kind of uniqueness for short coset paths in sufficiently acyclic structures. This essential uniqueness also gives rise to a direction in the sense of a set of agents for the first step on any short non-trivial path towards a target world. Secondly, up to inessential variations any pair of worlds can be connected by at most one short non-trivial coset path; so short links to Z can only rule out a few directions, and sufficiently high multiplicity gives us enough remaining possibilities to choose the next copy from.

Thus, we eventually find a world $v' \sim v$ with $\mathrm{agt}(v', w) = \alpha$ such that $[\![v']\!]$ and Z are n-free over $[\![w]\!]$. □

C. Upgrading in sufficiently acyclic Cayley structures

The upgrading argument follows the same pattern as in Section V, cf. Figure 1 in Section II-D. For two Cayley structures $\mathfrak{M}$ and $\mathfrak{N}$ that are sufficiently acyclic and have sufficiently high multiplicities we want to show

$$\mathfrak{M}, w \sim^{\ell} \mathfrak{N}, v \;\Rightarrow\; \mathfrak{M}, w \equiv_q \mathfrak{N}, v,$$

for some $\ell = \ell(q)$ by providing a winning strategy for **II** in the q-round Ehrenfeucht–Fraïssé game.

As before, we maintain isomorphic substructures $\mathfrak{M}_i \subseteq \mathfrak{M}$ and $\mathfrak{N}_i \subseteq \mathfrak{N}$ that contain all worlds pebbled so far. These substructures need to be consistently extended in each round. However, in the case of finite Cayley structures this extension is more involved. Free ω-unfoldings of Cayley structures are in some sense tree-like; this made it rather easy to expand $\mathfrak{N}_{i-1}$ to $\mathfrak{N}_i$ because there is always a fresh branch to choose which allows us to ensure that long distances are duplicated correctly. We do not have this advantage in the finite case, but a sufficient degree of acyclicity implies that *small* substructures are tree-like and that short coset paths are essentially unique by the Zipper Lemma. Together with sufficiently high multiplicities, this is enough for the upgrading to go through.

The possibility for player **II** to maintain the invariant that links the actual game to the auxiliary structure in the dual picture is the crucial step in the proof of the upgrading lemma that goes through due to a sufficient degree of freeness. Two things are essential in this approach. Firstly, because we essentially match tree decompositions in the dual picture, just two vertices need to be considered at a time, viz. parent and child in the tree decomposition. Secondly, the sufficiently free choice of the new vertex ensures that long distances in $\mathfrak{M}_i$ are matched by long distances in $\mathfrak{N}_i$.

We obtain the following upgrading lemma, and through it the crucial argument for our main theorem over the class of all finite CK-structures, again following Figure 1.

Lemma 39. *For some suitable choice of $\ell = \ell(q)$, any finite sufficiently acyclic Cayley structures with sufficiently high multiplicities $\mathfrak{M}$ and $\mathfrak{N}$ satisfy: $\mathfrak{M}, w \sim^{\ell} \mathfrak{N}, v \;\Rightarrow\; \mathfrak{M}, w \equiv_q \mathfrak{N}, v$.*

Theorem 40. $\mathsf{ML}[\mathsf{CK}] \equiv \mathsf{FO}/{\sim}$ *over the class of finite CK-structures.*

VI. Conclusion

Characterisations of natural levels of modal expressiveness beyond basic ML, and strictly between FO and full monadic second-order logic MSO, have primarily been obtained by automata-theoretic methods. The great success of those methods, most prominently the characterisation of L_μ in [11], but also of CTL^* in [9], [12] or of PDL in [4] in terms of natural fragments of MSO (over typically infinite trees), has so far not been matched for corresponding questions in finite model theory. Our approach in this paper applies rather more purely model-theoretic methods, of established significance in finite model theory and also in the modal setting [14], [5], [15], and takes them in a new direction beyond their customary first-order range. In particular, we use locality arguments in a seemingly locality-averse setting of structures with a complex, multi-scale connectivity pattern. This is facilitated by a direct algebraisation of the relevant frames and their special non-elementary frame properties in our main Lemma 3, which (unlike the classical use of tree unfoldings) does also work in the finite. Due to other inherent limitations, these new methods may not hold much promise for the great open problem concerning the status in finite model theory of the Janin–Walukiewicz result, but further extensions may treat other frame classes and/or extensions of ML of interest in a similar vein, by combinations of direct use of generic algebraic-combinatorial structures with suitable logical interpretations.

References

[1] C. Beeri, R. Fagin, D. Maier, and M. Yannakakis. On the desirability of acyclic database schemes. *Journal of the ACM*, 30:497–513, 1983.
[2] C. Berge. *Graphs and Hypergraphs*. North-Holland, 1973.
[3] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, 2001.
[4] F. Carreiro. PDL is the bisimulation-invariant fragment of weak chain logic. In *Proc. 30th ACM/IEEE Symposium on Logic in Computer Science LICS'15*, pages 341–352, 2015.
[5] A. Dawar and M. Otto. Modal characterisation theorems over special classes of frames. *Annals of Pure and Applied Logic*, 161:1–42, 2009.
[6] H.-D. Ebbinghaus and J. Flum. *Finite Model Theory*. Springer, 1999.
[7] R. Fagin, J. Halpern, Y. Moses, and M. Vardi. *Reasoning About Knowledge*. MIT Press, 1995.
[8] V. Goranko and M. Otto. Model theory of modal logic. In P. Blackburn, J. van Benthem, and F. Wolter, editors, *Handbook of Modal Logic*, pages 249–329. Elsevier, 2007.
[9] T. Hafer and W. Thomas. Computation tree logic CTL* and path quantifiers in the monadic theory of the binary tree. In *Proc. Automata, Languages and Programming*, pages 269–279. Springer, 1987.
[10] W. Hodges. *Model Theory*. Cambridge University Press, 1993.
[11] D. Janin and I. Walukiewicz. On the expressive completeness of the propositional μ-calculus with respect to monadic second-order logic. In *Proc. CONCUR*, pages 263–277, 1996.
[12] F. Moller and A. Rabinovich. Counting on CTL*: on the expressive power of monadic path logic. *Information and Computation*, 184:147–159, 2003.
[13] M. Otto. Elementary proof of the van Benthem–Rosen characterisation theorem. Technical Report 2342, FB Mathematik, TU Darmstadt, 2004.
[14] M. Otto. Modal and guarded characterisation theorems over finite transition systems. *Annals of Pure and Applied Logic*, 130:173–205, 2004.
[15] M. Otto. Highly acyclic groups, hypergraph covers and the guarded fragment. *Journal of the ACM*, 59 (1), 2012.
[16] E. Rosen. Modal logic over finite structures. *Journal of Logic, Language and Information*, 6:427–439, 1997.
[17] J. van Benthem. *Modal Logic and Classical Logic*. Bibliopolis, Napoli, 1983.