

Cayley Structures and Common Knowledge

Felix Canavoi and Martin Otto *

revised, July 2020

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 technical challenges posed by the reachability and transitive closure features of the derived accessibility relations are dealt with through passage to (finite) bisimilar coverings of epistemic frames by Cayley graphs of permutation groups whose generators are associated with the agents. Epistemic frame structure is here induced by an algebraic coset structure. Cayley structures with specific acyclicity properties support a locality analysis at different levels of granularity as induced by distance measures w.r.t. various coalitions of agents.

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

1	Introduction	3
2	Basics	5
2.1	S5 and CK Kripke structures and modal logic	5
2.2	Common knowledge in Cayley structures	7
2.3	Bisimulation	8
2.4	Main lemmas	9
2.5	Upgrading for expressive completeness	13
3	Coset acyclicity and its structure theory	14
3.1	Coset acyclic Cayley structures	14
3.1.1	Coset cycles	15
3.1.2	2-acyclicity	16
3.1.3	Coset acyclicity and hypergraph acyclicity	18
3.2	Structure theory for n -acyclic Cayley structures	20
3.2.1	Freeness	20
3.2.2	Coset paths	23
3.2.3	The freeness theorem	24
4	Characterisation theorem	28
4.1	The invariant	30
4.2	Upgrading and characterisation	36

1 Introduction

Modal logics have diverse applications that range from specification of process behaviours in computer science to the reasoning about knowledge and the interaction of agents in all kinds of distributed settings. Across this broad conception of modal logics bisimulation invariance stands out as the crucial semantic feature uniting an extremely diverse family of logics. Bisimulation equivalence is based on an intuitive back&forth probing of transitions 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 is meant to deal 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. This results in a model-theoretic characterisation that casts a natural level of expressiveness in a new perspective.

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: $ML \equiv FO/\sim$, where 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 $ML \subseteq FO$ (identified with its standard translation into FO).

Theorem 1.1 (van Benthem [19]). $ML \equiv FO/\sim$.

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 [18], van Benthem’s characterisation theorem $ML \equiv FO/\sim$ 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 [15] and also in [12]; like many of the more challenging extensions and variations in [16, 9, 17], it relies on a model-theoretic upgrading argument that links finite approximation levels \sim^ℓ 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 must in fact be invariant under some finite level \sim^ℓ of bisimulation equivalence. This may be seen as a crucial compactness phenomenon for \sim -invariant FO , despite the unavailability of compactness for FO

in some 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 1.2 (Janin–Walukiewicz [14]). $L_\mu \equiv \text{MSO}/\sim$.

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 epistemic setting was obtained in [9], 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 the epistemic modal logic ML[CK] in a multi-agent setting with *common knowledge* operators. Common knowledge operators 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. [11] 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 for highly intricate yet regular patterns of multi-scale transitive relations, the *coset structure* of Cayley groups w.r.t. combinations of generators. We can show that *Cayley structures*, obtained as expansions of relational encodings of Cayley groups by propositional assignments, are universal representatives up to bisimulation of S5 structures – both in the general and in the finite setting. 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 2.3 and 2.6.

Lemma 1.3 (main lemma). *Every (finite) CK-structure admits (finite) bisimilar coverings by Cayley structures (of various degrees of acyclicity w.r.t. their epistemic or coset structure).*

Cayley groups with suitable acyclicity properties for their coset structure are available from [17]; 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 as induced by reachability and transitivity phenomena for different groups of agents in CK-structures. Our main theorem is the following.

Theorem 1.4. $\text{ML}[\text{CK}] \equiv \text{FO}/\sim$ over CK-structures, both classically and in the sense of finite model theory.

An equivalent alternative formulation would characterise $\text{ML}[\text{CK}]$ as the \sim -invariant fragment of $\text{FO}[\text{CK}]$, 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 1.1, can be found in Section 2.5.

2 Basics

2.1 S5 and CK Kripke structures and modal logic

For this paper we fix a finite non-empty 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 $(P_i)_{i \in I}$. Individual formulae of the logics considered will only mention finitely many basic propositions, and 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 $i \in I$, and 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 \perp, \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 $\mathfrak{M}, w \models \perp$ for no $w \in W$;

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

- boolean connectives are treated as usual;
- $\mathfrak{M}, w \models \Box_\alpha \varphi$ 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 $w' \in [w]_\alpha$.

The extension of ML to *common knowledge logic* ML[CK] introduces further modalities \Box_α and \Diamond_α for every *group of agents* $\alpha \subseteq \Gamma$. The intuitive epistemic reading of \Box_α is that “it is common knowledge among agents in α that ...”, and \Diamond_α is treated as the dual of \Box_α . The semantics of \Box_α 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 on any path formed edges from the R_a for $a \in \alpha$. The relevant set of worlds w' is the equivalence class $[w]_\alpha$ w.r.t. the derived equivalence relation

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

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_α coincides with \Box_a just as R_α 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 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 2.1. With any S5 Kripke frame (or structure) we associate the *CK-frame* (or structure) obtained as the expansion of the family $(R_a)_{a \in \Gamma}$ to the family $(R_\alpha)_{\alpha \in \tau}$ for $\tau = \mathcal{P}(\Gamma)$, where $R_\alpha = \text{TC}(\bigcup_{a \in \alpha} R_a)$.

We use notation \mathfrak{M}^{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,

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

which is again an S5 Kripke structure. The resulting class of CK-structures is non-elementary. Indeed, a simple compactness argument shows that the defining conditions for the R_α cannot be first-order expressible.

Definition 2.2. The syntax of *epistemic modal logic with common knowledge*, ML[CK], for the set of agents Γ is the same as the syntax of basic modal logic ML with modalities \Box_α and \Diamond_α 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}^{CK} .

We next look at a seemingly very special class of CK-structures. In these, the equivalence relations R_α are induced by the coset structure of an underlying

group w.r.t. designated (sets of) generators. We use the name *Cayley structures* for these special CK-structures whose epistemic structure is induced by the Cayley graph of a group, which relates its combinatorics to basic algebraic concepts as explored by Cayley in [7, 8]. As we shall see in Lemma 2.7, which is a cornerstone for the approach taken in this paper, the class of these Cayley structures is rich enough to represent any CK-structure up to bisimulation.

2.2 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 $R_e := \{(g, ge) : g \in G\}$, 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 left multiplication with hg^{-1} in the group.

We partition the generator set E into non-empty subsets E_a associated with the agents $a \in \Gamma$, and consider subgroups $\mathbb{G}_a = \langle e : e \in E_a \rangle \subseteq \mathbb{G}$ generated by the $e \in E_a$. This allows us to regard left 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

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

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 $E_\alpha := \bigcup\{E_a : a \in \alpha\}$, the equivalence relations

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

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

Definition 2.3. With any Cayley group $\mathbb{G} = (G, \cdot, 1)$ with generator set E of involutions that is Γ -partitioned according to $E = \bigcup_{a \in \Gamma} E_a$, we associate the *Cayley CK-frame* (Cayley frame, for short) \mathbb{G}^{CK} over the set G of possible worlds with accessibility relations R_α for $\alpha \in \tau = \mathcal{P}(\Gamma)$. A *Cayley structure* consists 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 $\mathfrak{M}^{\text{CK}} = \mathfrak{M}$. In the following we simply speak of α -edges, -classes, -cosets with reference to the R_α or \mathbb{G}_α in any Cayley structure.

2.3 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 $(w, v) \in W \times V$, which denotes a placement of a single pair of pebbles on w in \mathfrak{M} and on v in \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 2.4. \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^ℓ .

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

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

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}^{CK} is *safe for bisimulation*. On the other hand, there is no such correspondence at the level of the finite approximations \sim^ℓ , 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}^{CK} , and encapsulates central aspects of our concerns here:

the passage $\mathfrak{M} \mapsto \mathfrak{M}^{\text{CK}}$ is beyond first-order control
and breaks standard notions of locality.

Correspondingly, modal or first-order expressibility over \mathfrak{M}^{CK} transcends expressibility over \mathfrak{M} , and in particular $\text{ML}[\text{CK}]$ 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. [4, 12]. Here and in the following we denote as

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

indistinguishability by ML-formulae of modal nesting depth (quantifier rank) up to ℓ , just as \equiv_q^{FO} or just \equiv_q will denote classical first-order equivalence (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 2.5. *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 $\ell \in \mathbb{N}$:*

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

In particular, the semantics of any modal formula (in ML or in ML[CK]) is preserved under full bisimulation equivalence (of either the underlying plain S5 structures or their CK-expansions). Any formula of ML[CK] is preserved under some level \sim^{ℓ} 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 terminology speaks of bounded morphisms). Bisimilar tree unfoldings are a well-known instance of (albeit, usually infinite) bisimilar coverings with many applications.

Definition 2.6. A surjective homomorphism $\pi: \hat{\mathfrak{M}} \rightarrow \mathfrak{M}$ 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 $\hat{\mathfrak{M}}$.

2.4 Main lemmas

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 [9] 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. Instead we will have to look to stronger acyclicity properties, viz. coset acyclicity of Cayley groups in Section 3.1.1. On the other hand, we can naturally model any CK-scenario up to bisimulation, indeed up to a bisimilar covering, in a Cayley group directly. The following Lemma 2.7, which was already stated as Lemma 1.3 in the introduction, forms

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

a cornerstone of our approach to the analysis of the expressive power of first-order logic for \sim -invariant properties over CK-structures.

More fundamentally it says that, as far as bisimulation invariant phenomena are concerned, Cayley structures can serve as representatives of arbitrary CK-structures. And this is even true not just within the class of all CK-structures but also in the more restricted setting of just finite CK-structures.

Lemma 2.7. *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). Indeed, for the construction of the covering, we even go below that level and decompose the given accessibility relations R_a further into constituents induced by individual R_a -edges.

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 , where we identify an edge $e = (w, w')$ with its converse (w', w) (or think of the edge relations as sets of unordered pairs, or sets of size 1 for reflexive and size 2 for irreflexive edges). Formally we may represent this disjoint union by tagged copies of the individual edge pairs from each R_a for $a \in \Gamma$ as $E = \{(\{w, w'\}, a) : a \in \Gamma, (w, w') \in R_a\}$, which may be partitioned into subsets $E_a = \{(\{w, w'\}, a) : (w, w') \in R_a\}$ 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 . The vertices of this hypercube are the $\{0, 1\}$ -valued sequences indexed by the set E , with a symmetric e -edge between any pair of such sequences whose entries differ precisely in the e -component. 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. We note that W is closed under the action of π_e . For $e = (\{w, w'\}, a)$, the permutation π_e fixes all worlds in W other than w, w' ; and if $e = (\{w, w'\}, a)$ is a reflexive a -edge, then $\pi_e \upharpoonright W = \text{id}_W$. In restriction to 2^E on the other hand, π_e has no fixed points, and $\pi_e \neq \pi_{e'}$ whenever $e \neq e'$ (even if $\pi_e \upharpoonright W = \pi_{e'} \upharpoonright W$, which can occur for $a \neq a'$ if $e = (\{w, w'\}, a)$ and $e' = (\{w, w'\}, a')$).

For \mathbb{G} we take the subgroup of the symmetric group on V that is generated by these π_e , which we regard as involutive generators of \mathbb{G} . This is justified since, as just observed, the $(\pi_e)_{e \in E}$ and $1 \in \mathbb{G}$ are pairwise distinct due to the 2^E -component. We may thus identify π_e with e and regard the edge set E as the subset $E = \{\pi_e : e \in E\} \subseteq \mathbb{G}$, which generates \mathbb{G} as a group. We let \mathbb{G} act on V in the natural fashion (from the right): for $g = e_1 \cdots e_n$,

$$g: v \longmapsto ve_1 \cdots e_n := (\pi_{e_n} \circ \cdots \circ \pi_{e_1})(v).$$

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 $v \in V$. It also leaves $W \subseteq V$ invariant as a set, i.e. the action can be restricted to W . Then the map

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

is a bisimilar covering w.r.t. the following natural S5 interpretations of edge relations R_α as $R_\alpha := \text{TC}(\{(w, g), (w, ge)\}: w \in W, g \in \mathbb{G}, e \in E_\alpha\}$ over $W \times \mathbb{G}$. This bisimilar covering directly extends to the induced S5 frames with accessibility relations R_α for $\alpha \in \tau$ (again obtained as transitive closures of corresponding unions). Moreover, since \mathfrak{M} is connected, \mathbb{G} acts transitively on W and we may restrict to a single orbit, i.e. to a single connected sheet $\{w_0\} \times \mathbb{G}$ of the above multiple covering. This restriction corresponds 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 $\{(w_0, g): g \in \mathbb{G}\}$, which is naturally isomorphic with the Cayley frame of \mathbb{G} . 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}^{CK} . This is achieved by pulling back $P_i \subseteq W$ to its pre-image $\pi^{-1}(P_i) \subseteq G$, which becomes the assignment to proposition P_i on \mathbb{G} . The resulting

$$\pi: (G, (R_\alpha), (P_i)) \longrightarrow \mathfrak{M}^{\text{CK}}$$

provides the desired bisimilar covering of the CK-structure \mathfrak{M}^{CK} by a Cayley structure. Note that \mathbb{G} and $(G, (R_\alpha), (P_i))$ are finite if W is. \square

We used the hypercube structure 2^E in the above as an auxiliary component to adapt the group to its purposes in the covering, viz. in this case, to turn the set of individual accessibility edges of \mathfrak{M} into a set of non-trivial and mutually independent generators in the Cayley structure that covers \mathfrak{M} . Different variants of this idea are available. These allow us to adapt the group structure in order to make the bisimilar covering more amenable for specific purposes. We discuss some immediate such variants here; an even more important one will then be discussed in much greater detail in Section 3.

Firstly, the well-known tree-like bisimilar unfolding of S5 Kripke structures can be presented in a very similar fashion based on free groups and their Cayley graphs. We define the *free* or *acyclic* group with involutive generator set E over the set of reduced words over the alphabet E . An E -word $w = e_1 \dots e_n \in E^*$ is *reduced* if $e_{i+1} \neq e_i$, for all $1 \leq i < n$.

Definition 2.8. The *free group* $\mathbb{F}(E)$ with involutive generator set E is the group that consists of all reduced words over the alphabet E without any non-trivial equalities, together with the (reduced) concatenation of words as its operation and the empty word as its neutral element.

Using $\mathbb{F}(E)$ and its Cayley graph in place of the group \mathbb{G} (as was abstracted from permutation group action on $\mathfrak{M} \oplus 2^E$ above), we obtain the following.

Lemma 2.9. *Any connected CK-structure admits a bisimilar covering by a Cayley CK-structure based on the Cayley graph of a free or acyclic group with involutive generators, which itself is a tree structure.*

Note that the resulting bisimilar coverings are infinite in all but the most trivial cases. Also observe that non-trivial S5- and CK-frames cannot be trees. Rather, the above bisimilar coverings result in S5- or CK-structures that are

generated from actual tree structures through transitive closure operations; in a sense they are as close to trees as possible, and coset acyclic in the sense to be discussed in Section 3 (cf. Definition 3.1).

Cayley graphs of *large girth* have been obtained from permutation group actions in [3] and used in the construction of *finite* bisimilar coverings of multimodal Kripke structures [16] and of S5 structures [9]. We could here similarly obtain finite bisimilar coverings of CK-structures that are generated through transitive closures from Cayley graphs of large girth (i.e. without short generator cycles). It turns out, however, that much stronger acyclicity properties for Cayley structures are needed for our present purposes. The cyclic configurations that matter in Cayley frames are induced by R_α -edges (which includes R_a -edges as a special case). Arising from transitive closures, these edges stem from paths of a priori unbounded lengths in terms of the underlying generator edges; and equivalence classes for accessibility relations R_α are *cosets* w.r.t. generated subgroups. This is why levels of *coset acyclicity*, rather than just lower bounds on girth, will be extensively discussed in Section 3.

As another immediate variation of the main lemma, we consider *richness* criteria. Simple variants of the above covering construction allow us to locally boost multiplicities. The *multiplicity* of a bisimulation type in an α -class is the cardinality of the set of its realisations in this class, and k -richness requires that this multiplicity is at least k (if not 0).

Definition 2.10. A CK-structure \mathfrak{M} is k -rich, for some $k \in \mathbb{N}$, if for every $\alpha \in \tau$, the multiplicity of every bisimulation type that is realised in an α -class is at least k in that class. \mathfrak{M} is ω -rich if all these multiplicities are infinite.

By augmenting the number of generators in the group \mathbb{G} that are associated with $e \in E$ in the basic construction or its variants we can achieve k -richness in (finite) bisimilar coverings as in Lemma 2.7 as well as ω -richness in tree-based coverings as in Lemma 2.9. Technically it suffices to replace E by $E \times \{0, \dots, k\}$ or by $E \times \omega$, and to let the group operation $\pi_{(e,i)}$ of the copies (e, i) be the same as π_e on W while separating them in the hypercube component for the new E . This trick boosts multiplicities by a factor of 2^k or ω . (In fact the basic covering construction in the proof of Lemmas 2.7 and 2.9 typically already introduces quite some boost in multiplicities compared to \mathfrak{M} since the operation of π_e on W is rich in fixpoints.)

Lemma 2.11. *For all $k \in \mathbb{N}$, every connected (finite) CK-structure admits a (finite) bisimilar covering by a Cayley structure that is k -rich. Every connected CK-structure admits a bisimilar covering by an infinite ω -rich Cayley structure based on a free or acyclic group with involutive generators whose Cayley graph is a tree.*

The crucial insight of Lemma 2.7 justifies the following, since – up to bisimulation – we may now transfer any model-theoretic question about (finite) CK-structures to (finite) Cayley structures. Lemma 3.4 will actually offer (finite)




Figure 1: Upgrading \sim^ℓ to \equiv_q in bisimilar companions.

representations by Cayley structures with additional acyclicity and richness properties. Those are again obtained as coverings by Cayley groups with corresponding properties from [17].

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

2.5 Upgrading for expressive completeness

The key to the expressive completeness results from [15] to [9, 17] lies in establishing the following finiteness or compactness phenomenon for \sim -invariant FO-formulae $\varphi(x)$ over the relevant classes \mathcal{C} of structures:

$$(\dagger) \quad \{\varphi \in \text{FO}: \varphi \sim\text{-invariant over } \mathcal{C}\} = \bigcup_{\ell \in \mathbb{N}} \{\varphi \in \text{FO}: \varphi \sim^\ell\text{-invariant over } \mathcal{C}\}.$$

This finiteness property in turn follows if suitable levels \sim^ℓ can be upgraded in bisimilar companions within \mathcal{C} so as to guarantee equivalence w.r.t. the given φ of quantifier rank q . The argument is as follows. Assume that 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 $\mathfrak{N}^*, v^* \sim \mathfrak{N}, v$ in \mathcal{C} such that $\mathfrak{M}^*, w^* \equiv_q \mathfrak{N}^*, v^*$, as in Figure 1. Then the detour through the lower rung of Figure 1 shows that over \mathcal{C} any \sim -invariant FO-formula of quantifier rank q is indeed \sim^ℓ -invariant, and hence expressible in ML at modal nesting depth ℓ over \mathcal{C} by Theorem 2.5.

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^ℓ (for a level $\ell = \ell(q)$ to be determined). Features of this kind that would beat any level ℓ involve

- small multiplicities w.r.t. accessibility relations, like fixed but differing small cardinalities for definable properties of worlds in α -classes, and
- short cycles of fixed lengths 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 1.1, 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, and short cycles can be eliminated in products with Cayley groups of large girth.

We have also seen first indications above how to eliminate differences involving small multiplicities in (finite) bisimilar coverings by Cayley structures that are sufficiently rich as in Lemma 2.11; and Lemma 2.9 at least allows us to focus on Cayley structures that avoid cycles as far as possible at the level of the underlying Cayley graph if finiteness does not matter.

The great challenge, however, lies with the game arguments that are typically used 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. [13, 10]), 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. Indeed Gaifman locality is completely trivialised in connected CK-structures, which must form a single Gaifman clique w.r.t. R_Γ . Naively it thus seems all but hopeless to use locality techniques in structures that are as dense in terms of their edge relations as CK-structures are. But despite its denseness, the highly regular edge pattern of Cayley structures will allow us to invoke notions of locality at different levels of granularity, which are based on the various intermediate R_α between the extremes of the individual R_a and the trivial R_Γ .

3 Coset acyclicity and its structure theory

This section is the technical core of this paper. Here we introduce all tools required to overcome the immediate obstructions for upgrading \sim^ℓ to \equiv_q for suitable $\ell = \ell(q)$, over the class of (finite) Cayley structures. To overcome problems with avoidable short cycles, we introduce notions of *coset acyclicity* from [17] for our purposes, and show that every (finite) Cayley structure admits a bisimilar covering by a (finite) Cayley structure that is coset n -acyclic. We also introduce the *dual hypergraph* associated with a Cayley structure. This dual picture will be an important tool for our upgrading arguments later and also allows us to relate coset acyclicity to classical hypergraph acyclicity. The second part of this section, Section 3.2, then introduces *freeness* as the most important notion in the structure theory of suitable Cayley structures for our upgrading arguments, with core results from [6].

3.1 Coset acyclic Cayley structures

In the case of CK-frames and Cayley frames one cannot hope to avoid cycles outright.³ 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

³This 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 .

(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 from v to v' ; and this difference could be expressible in FO_q . The solution is to find bisimilar companions for \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 $\beta \subsetneq \alpha$. It turns out that the notion of *coset acyclicity* from [17] is what we can use.

3.1.1 Coset cycles

Recall from the definition of Cayley structures that the accessibility relations R_a (for individual agents $a \in \Gamma$) or R_α (for groups of agents $\alpha \in \tau = \mathcal{P}(\Gamma)$, of which the R_a are just a special case) arise from transitive closures of *sets* of edge relations induced by individual generators of the underlying group structure.

Cyclic configurations w.r.t. combinations of different R_α are cyclic configurations formed by cosets rather than by generators of the underlying group. Correspondingly we are interested in Cayley frames that avoid short coset cycles rather than just short generator cycles (i.e. large girth).

Definition 3.1. Let \mathfrak{M} be a Cayley frame.

- (i) A *coset cycle of length $m \geq 2$* 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

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

- (ii) \mathfrak{M} is *acyclic* if it does not have any coset cycles, and *n -acyclic* if it does not contain coset cycles of lengths up to n .

In Section 2.4 we showed that every Cayley structure has a bisimilar covering by an infinite Cayley structure that is based on a free group with involutive generators (cf. Lemma 2.9). It is easy to see that this covering is coset acyclic: any non-trivial coset cycle would stem from a non-trivial generator cycle. The two kinds of acyclicity coincide at the level of full acyclicity because the blowup in length in the passage from coset-steps to generator-steps is not taken into account. Together with the ω -richness property from Lemma 2.11 these coverings would suffice for proving the classical version of our characterisation theorem. For the finite model theory version, we need bisimilar coverings that are finite and, at the same time, *sufficiently coset acyclic* and *sufficiently rich*. The richness part is already covered by Lemma 2.11. Suitable levels of coset acyclicity in finite Cayley groups were introduced in [17].

Definition 3.2. Let \mathbb{G} be a Cayley group with generator set E .

- (i) A *coset cycle of length m* in \mathbb{G} is a cyclic tuple $(g_i, \alpha_i)_{i \in \mathbb{Z}_m}$ with $g_i \in \mathbb{G}$ and $\alpha_i \subseteq E$, for all $i \in \mathbb{Z}_m$, where $g_i^{-1}g_{i+1} \in \mathbb{G}_{\alpha_i}$ and

$$g_i \mathbb{G}_{\alpha_{i-1} \cap \alpha_i} \cap g_{i+1} \mathbb{G}_{\alpha_i \cap \alpha_{i+1}} = \emptyset.$$

- (ii) \mathbb{G} is *acyclic* if it does not contain any coset cycles, and *n -acyclic* if it does not contain coset cycles of lengths up to n .

As shown in [17], every finite Cayley group can be covered by a finite, n -acyclic Cayley group, for arbitrary $n \in \mathbb{N}$.

Lemma 3.3. *For every finite Cayley group \mathbb{G} with finite generator set E and every $n \in \mathbb{N}$, there is a finite, n -acyclic Cayley group $\hat{\mathbb{G}}$ with generator set E such that there is a surjective homomorphism $\pi: \hat{\mathbb{G}} \rightarrow \mathbb{G}$.*

Combining Lemma 3.3 with the main lemmas from Section 2.4, we obtain the desired coverings for finite CK-structures. We apply Lemma 3.3 to the Cayley group that underlies the Cayley structure that we get from Lemma 2.11 and define a bisimilar covering as in the proof of Lemma 2.7.

Lemma 3.4. *For all $k, n \in \mathbb{N}$, every connected (finite) CK-structure admits a (finite) bisimilar covering by a Cayley structure that is k -rich and coset n -acyclic.*

Here and in what follows it is important to keep in mind that α -classes in Cayley frames and structures are nothing but left cosets w.r.t. subgroups \mathbb{G}_α in the underlying Cayley group. This is clearly reflected in the parallelism of Definition 3.1 and 3.2. Absence of coset 2-cycles in a Cayley frame based on \mathbb{G} , for instance, just says that $\mathbb{G}_\alpha \cap \mathbb{G}_\beta = \mathbb{G}_{\alpha \cap \beta}$ for all $\alpha, \beta \subseteq \Gamma$.

Observation 3.5. *Let \mathfrak{M} be a Cayley frame based on the Cayley group \mathbb{G} . \mathfrak{M} is 2-acyclic if, and only if, for all $w \in W, \alpha, \beta \in \tau$, $[w]_\alpha \cap [w]_\beta = [w]_{\alpha \cap \beta}$, if, and only if, $\mathbb{G}_\alpha \cap \mathbb{G}_\beta = \mathbb{G}_{\alpha \cap \beta}$ for all $\alpha, \beta \in \tau$.*

3.1.2 2-acyclicity

We use Observation 3.5 to show that 2-acyclic Cayley frames display a high degree of regularity that will be essential for many of the notions to be introduced in this and the following sections. While arbitrary S5- or CK-frames impose very little structure on the overlap patterns between the equivalence classes w.r.t. various R_α , we shall see that in 2-acyclic Cayley frames, e.g. any pair of vertices is connected by R_α for a unique minimal set $\alpha \in \tau$.

In the next section, we shall investigate the structure of the dual hypergraph associated with a Cayley frame (cf. Definition 3.12). We anticipate the definition of the *dual hyperedge* which is convenient for notational purposes here as well.

Definition 3.6. In a Cayley frame \mathfrak{M} define the *dual hyperedge* induced by a world w to be the set of cosets

$$\llbracket w \rrbracket := \{[w]_\alpha : \alpha \in \tau\}.$$

The following lemma from [6] is straightforward from the definitions. It shows in particular that two worlds w, v in a 2-acyclic structure are connected by a unique minimal set of agents α , i.e. a set α for which $[w]_\beta = [v]_\beta$ if and only if $\beta \supseteq \alpha$. This then justifies Definition 3.8.

Lemma 3.7. *In a 2-acyclic Cayley frame \mathfrak{M} with worlds w, w_1, \dots, w_k and sets of agents $\alpha_1, \dots, \alpha_k \in \tau$:*

1. For $\beta := \bigcap_{1 \leq i \leq k} \alpha_i$: $w \in \bigcap_{1 \leq i \leq k} [w_i]_{\alpha_i} \Leftrightarrow \bigcap_{1 \leq i \leq k} [w_i]_{\alpha_i} = [w]_\beta$.
2. The set $\bigcap_{1 \leq i \leq k} \llbracket w_i \rrbracket$ has a least element in the sense that there is an $\alpha_0 \in \tau$ such that $[w_1]_{\alpha_0} \in \bigcap_{1 \leq i \leq k} \llbracket w_i \rrbracket$ and, for any $\alpha \in \tau$ and $1 \leq i \leq k$:

$$[w_i]_\alpha \in \bigcap_{1 \leq i \leq k} \llbracket w_i \rrbracket \Leftrightarrow \alpha_0 \subseteq \alpha.$$

We shall often blur the distinction between a finite set (of worlds) and its enumeration as a tuple, using notation like \mathbf{w} for a finite collection of worlds w .

Definition 3.8. In a 2-acyclic Cayley frame we denote the unique *minimal set of agents that connects the worlds in \mathbf{w}* by $\text{agt}(\mathbf{w}) \in \tau$.

So, for each one of the worlds $w \in \mathbf{w}$, $\text{agt}(\mathbf{w}) = \bigcap \{\alpha \in \tau : \mathbf{w} \subseteq [w]_\alpha\}$. Intuitively, $\text{agt}(\mathbf{w})$ sets the scale for zooming-in on the minimal substructure that connects the worlds \mathbf{w} . We shall see later that, regarding distances between the worlds \mathbf{w} , we only need to control cycles with β -steps for $\beta \subsetneq \text{agt}(\mathbf{w})$ within the cluster $[w]_{\text{agt}(\mathbf{w})}$ for $w \in \mathbf{w}$. For intersections between dual hyperedges, Lemma 3.7 implies that every intersection can be described by the unique set of agents $\text{agt}(\mathbf{w})$. This means, for every $w \in \mathbf{w}$:

$$[w]_\alpha \in \bigcap_{w \in \mathbf{w}} \llbracket w \rrbracket \Leftrightarrow \alpha \supseteq \text{agt}(\mathbf{w}).$$

The following lemma will be vital for many of the constructions to come, as it allows us to control $\text{agt}(\mathbf{w})$ in 2-acyclic frames.

Lemma 3.9. *In a 2-acyclic Cayley frame for worlds w, v :*

1. For every agent $a \notin \text{agt}(w, v)$ and every $v' \in [v]_a \setminus \{v\}$:

$$\text{agt}(w, v') = \text{agt}(w, v) \cup \{a\}.$$

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

$$\text{agt}(w, v') = \text{agt}(w, v) \setminus \{a\}.$$

Proof. For (1), the inclusion $\text{agt}(w, v') \subseteq \text{agt}(w, v) \cup \{a\}$ is obvious; for the converse, observe that $v, a, v', (\text{agt}(v', w) \cup \text{agt}(w, v))$ would be a 2-cycle if $a \notin \text{agt}(w, v')$, and that $w, \text{agt}(w, v), v, \text{agt}(w, v')$ would be a 2-cycle if $\text{agt}(w, v) \not\subseteq \text{agt}(w, v') = \text{agt}(w, v) \cup \{a\}$.

Similarly for (2), let $a \in \alpha := \text{agt}(w, v)$, $\beta := \alpha \setminus \{a\}$ and assume that $\text{agt}(w, v') = \text{agt}(w, v'') = \beta$ for two different worlds $v', v'' \in [v]_a \setminus \{v\}$. One checks that v', β, v'', a, v' would form a 2-cycle. \square

3.1.3 Coset acyclicity and hypergraph acyclicity

The dual hypergraph of a Cayley frame or structure will play a crucial rôle in the Ehrenfeucht–Fraïssé arguments in Section 4. We here investigate the connections between acyclicity of Cayley frames and hypergraph acyclicity, and between coset paths in Cayley frames and chordless paths in hypergraphs.

Definition 3.10. A *hypergraph* is a structure $\mathfrak{A} = (A, S)$, A its vertex set and $S \subseteq \mathcal{P}(A) \setminus \{\emptyset\}$ its set of hyperedges.

All hypergraphs to be considered here have finite width, where *width* refers to the maximal size of the hyperedges, $\max\{|s| : s \in S\}$. A hypergraph $\mathfrak{A} = (A, S)$ has as its *induced substructures* the hypergraphs $\mathfrak{A} \upharpoonright Q$ for $\emptyset \neq Q \subseteq A$ with vertex set B and hyperedges $s \cap Q$ for those $s \in S$ that intersect Q .

With a hypergraph $\mathfrak{A} = (A, S)$ we associate its *Gaifman graph* $G(\mathfrak{A}) = (A, G(S))$; the undirected edge relation $G(S)$ of $G(\mathfrak{A})$ links $a \neq a'$ if $a, a' \in s$ for some $s \in S$. An *n-cycle* in a hypergraph is a 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]; *n-acyclicity* was introduced in [17].

Definition 3.11. A hypergraph $\mathfrak{A} = (A, S)$ is *acyclic* if it is *conformal* and *chordal*:

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

For $n \geq 3$, $\mathfrak{A} = (A, S)$ is *n-acyclic* if it is *n-conformal* and *n-chordal*:

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

Remark ([17]). If a hypergraph $\mathfrak{A} = (A, S)$ is *n-acyclic*, then every induced substructure $\mathfrak{A} \upharpoonright Q$ for non-empty subsets $Q \subseteq A$ of up to n vertices is acyclic.

A hypergraph $\mathfrak{A} = (A, S)$ is *tree decomposable* if it admits a tree decomposition $\mathcal{T} = (T, \delta)$: T is a tree (an acyclic connected graph) and $\delta: T \rightarrow S$ is a surjective map that associates a hyperedge with every node of the tree such that, for every vertex $a \in A$, the set $\{v \in T : a \in \delta(v)\}$ is connected (induces a subtree) in T . So $\mathcal{T} = (T, \delta)$ organises the hyperedges of \mathfrak{A} , as *bags* in the terminology of tree decompositions, so as to reflect the tree-like nature of their overlap pattern in \mathfrak{A} . 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 3.12. Let $\mathfrak{M} = (W, (R_\alpha)_{\alpha \in \tau})$ be a Cayley frame. Its *dual hypergraph* is the vertex-coloured hypergraph $d(\mathfrak{M}) := (d(W), S, (Q_\alpha)_{\alpha \in \tau})$ where

$$d(W) := \bigcup_{\alpha \in \tau} Q_\alpha \quad \text{for } Q_\alpha := W/R_\alpha,$$

$$S := \{[w] \subseteq d(W) : w \in W\}.$$

Note that $d(\mathfrak{M})$ has width $|\tau| = 2^{|\Gamma|}$ and diameter exactly 2. The notions of acyclicity for Cayley frames and hypergraph acyclicity are directly connected by the following.

Lemma 3.13 ([17]). *For $n \geq 3$, if \mathfrak{M} is an n -acyclic Cayley frame, then $d(\mathfrak{M})$ is an n -acyclic hypergraph.*

When playing the Ehrenfeucht-Fraïssé game in Section 4 to prove the upgrading theorem over Cayley structures we use their dual hypergraphs as auxiliary structures to describe a winning strategy. For that we keep track of an invariant involving a substructure that contains the pebbled worlds. This invariant has an image in the dual hypergraph that we use to maintain and expand the invariant properly in every round. The key notion to describe this dual image is the *convex m -closure*, which was defined in [17] for a similar purpose.

Definition 3.14. Let $\mathfrak{A} = (A, S)$ be a hypergraph.

- (i) A subset $Q \subseteq A$ is *m -closed* if every chordless path of length up to m between vertices $a, a' \in Q$ 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 : $\text{cl}_m(P) := \bigcap \{Q \supseteq P : Q \subseteq A \text{ } m\text{-closed}\}$.

In the analysis of the Ehrenfeucht-Fraïssé game it will be important to maintain, as part of an invariant, convex closures of the representations of the pebbled configurations, which means that such convex closures need to be updated to incorporate newly added elements. The following lemma shows that, in sufficiently acyclic hypergraphs, such an extensions of convex closures are well-behaved and can be controlled. In the statement, distance $d(P, q) = \min\{d(p, q) : p \in P\}$ between a set and a vertex refers to distance in the Gaifman graph, and $N^1(P) = \bigcup \{N^1(p) : p \in P\}$ is the 1-neighbourhood of the set P in the Gaifman graph. Both the following lemmas are from [17].

Lemma 3.15 ([17]). *Let $m > 1$, $\mathfrak{A} = (A, S)$ be a hypergraph that is sufficiently acyclic, $Q \subseteq A$ m -closed and $a \in A$ some vertex with $1 \leq d(Q, a) \leq m$. Let $\hat{Q} := \text{cl}_m(Q \cup \{a\})$ and consider the region $D := Q \cap N^1(\hat{Q} \setminus Q)$ in which this extended closure attaches to Q . Then $\hat{Q} \setminus Q$ is connected, and D separates $\hat{Q} \setminus Q$ from $Q \setminus D$ (in the graph-theoretic sense in $G(\mathfrak{A})$), whence*

$$\hat{Q} = Q \cup \text{cl}_m(D \cup \{a\}).$$

If Q is even $(2m + 1)$ -closed, then $D = Q \cap N^1(\hat{Q} \setminus Q)$ is a clique.

As $\hat{Q} = Q \cup \text{cl}_m(D \cup \{a\})$ for some clique $D \subseteq Q$, it will be important to bound the size of the extension $\text{cl}_m(D \cup \{a\})$. This is the extension that occurs, in the dual image as part of the invariant, in a single round, and a size bound will be critical for bounding the required level of ℓ -bisimulation that is necessary to win the game. As pointed out above, the dual hypergraphs are of uniform width $|\tau|$, which we regard as constant; we therefore seek functions $f_m(k)$ that bound the size of m -closures of up to k elements in hypergraphs of fixed width $|\tau|$, provided they are sufficiently acyclic.

Lemma 3.16 ([17]). *For fixed width, there are functions $f_m(k)$ such that in hypergraphs \mathfrak{A} of that width that are sufficiently acyclic, $|\text{cl}_m(P)| \leq f_m(k)$, for all $P \subseteq A$ of size $|P| \leq k$.*

3.2 Structure theory for n -acyclic Cayley structures

Section 3.1 covered the first part of the upgrading argument and especially the availability of suitable coverings. This section provides the tools for the second part, viz. for showing that \sim^ℓ -equivalence of two such suitable Cayley structures implies \equiv_q -equivalence, by playing the first-order Ehrenfeucht-Fraïssé game. The central notion of this subsection is *freeness*, a special property of sufficiently rich and acyclic Cayley structures. Essentially, freeness governs a single round in the Ehrenfeucht-Fraïssé game, in the sense that it allows **II** to find a suitable response to any move of **I**. As with richness and acyclicity, there are different degrees of freeness. The main result of this section is the *freeness theorem*, which states that sufficiently rich and acyclic Cayley structures are sufficiently free. Beside a formal introduction and discussion of freeness, this section introduces *coset paths*. They generalise graph-theoretic paths in the same way that coset cycles generalise graph-theoretic cycles and play an important part in proving the freeness theorem in Section 3.2.3. Most of the auxiliary theoretical results on coset paths come from [6].

Many of the definitions and notions that we will introduce from now on only make sense in 2-acyclic Cayley frames, because they are based on the notion of the unique minimal connecting set of agents $\text{agt}(\mathbf{w})$ defined in the previous section. As every Cayley structure has a 2-acyclic bisimilar covering, the following is justified.

Proviso. *We assume for the remainder of this section that all Cayley frames are at least 2-acyclic.*

3.2.1 Freeness

Consider playing the i -th round of an Ehrenfeucht-Fraïssé game on Cayley structures \mathfrak{M} and \mathfrak{N} with worlds $w_1, \dots, w_{i-1} \in W$ and $v_1, \dots, v_{i-1} \in V$ already pebbled. If player **I** chooses the world $w_i \in W$ in his next move, how does player **II** respond to this? As usual, **II** has to maintain a partial isomorphism between the pebbled worlds. In order to cope with player **I**'s challenges in future rounds, she also needs to match short distances between worlds exactly and to

match long distances with long distances. Since we play on Cayley structures, she has to respect these distances on multiple scales.

Freeness, or to be more precise (m, k) -freeness for $m, k \in \mathbb{N}$, is the special property of Cayley structures that allows for making suitable choices. Freeness roughly means that for a world v and a set \mathbf{z} of k worlds there is some world $v^* \sim v$ such that $d(v^*, z) > m$, for all $z \in \mathbf{z}$. In the scenario of the Ehrenfeucht-Fraïssé game, the set of worlds \mathbf{z} comprises not just the worlds pebbled so far, but a certain small substructure of \mathfrak{M} spanned by the pebbled worlds. In Cayley structures that are sufficiently free, this substructure can be extended properly from round to round. The main result of this section states that a Cayley structure is (m, k) -free if it is sufficiently acyclic and rich.

In the context of freeness, distance is defined in terms of the dual hypergraph of a Cayley structure. We shall use the dual hypergraph in Section 4 as an auxiliary structure to describe a winning strategy for player **II**.

Recall from Section 3.1.3, Definition 3.12, the dual hypergraph associated with a Cayley frame and the notion of Gaifman distance in hypergraphs discussed there. The following notation is useful towards the formal definition of freeness: 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 $\mathcal{A} \setminus t := \mathcal{A} \upharpoonright (A \setminus t)$. For a set of worlds $\mathbf{z} \subseteq W$, we write $\llbracket \mathbf{z} \rrbracket$ for the set $\{\llbracket z \rrbracket : z \in \mathbf{z}\}$ of associated hyperedges. A *pointed set (of worlds)* is a pair (\mathbf{z}, z_0) , where \mathbf{z} is a set of worlds and $z_0 \in \mathbf{z}$.

Definition 3.17. Let \mathfrak{M} be a 2-acyclic Cayley structure and $m, k \in \mathbb{N}$. For $v \in W$ and a pointed set of worlds (\mathbf{z}, z_0) we say that (\mathbf{z}, z_0) and v are *m -free*, denoted as $(\mathbf{z}, z_0) \perp_m v$, if $d_t(\bigcup \llbracket \mathbf{z} \rrbracket, \llbracket v \rrbracket) > m$ in $d(\mathfrak{M})$, where $t = \llbracket v \rrbracket \cap \llbracket z_0 \rrbracket$. We say that \mathfrak{M} is *(m, k) -free* if for all $v \in W$, all pointed sets (\mathbf{z}, z_0) with $|\mathbf{z}| \leq k$, and all sets of agents $\gamma \supseteq \text{agt}(v, z_0)$, there is some $v^* \sim v$ such that $\text{agt}(v^*, z_0) = \gamma$ and $(\mathbf{z}, z_0) \perp_m v^*$.

The finite set \mathbf{z} represents a set in \mathfrak{M} that is spanned by the worlds that have already been pebbled. The world v is a possible next move for **II** that may not be entirely suitable because it is too close to \mathbf{z} , in the sense that v and (\mathbf{z}, z_0) are not m -free. In this context z_0 plays the rôle of the world in \mathbf{z} that is, in some sense, closest to v .

Freeness, as presented here, is from [17]. It was used there to define a winning strategy for an Ehrenfeucht-Fraïssé game played on n -acyclic hypergraphs, in order to show a characterisation theorem for guarded logic. We adapt the idea for our purposes to use it over Cayley structures and their dual hypergraphs. Essentially, freeness is applied in the same way as in [17], but the proof that sufficiently acyclic and rich Cayley structures are (m, k) -free is new here. Definition 3.17 speaks about worlds in the Cayley structure and about distances in the Gaifman graph of the dual hypergraph. Our proof of the freeness theorem finds the desired world v^* , which is far enough away from \mathbf{z} in terms of the dual hypergraph, through constructions on the original Cayley structure.

A world v and a pointed set (\mathbf{z}, z_0) are m -free if the distance between $\llbracket v \rrbracket \setminus t$ and $(\bigcup \llbracket \mathbf{z} \rrbracket) \setminus t$ in $d(\mathfrak{M}) \setminus t$ is strictly larger than m , for $t = \llbracket v \rrbracket \cap \llbracket z_0 \rrbracket$. In

other words, a minimal path between $\llbracket v \rrbracket \setminus t$ and $(\bigcup \llbracket \mathbf{z} \rrbracket) \setminus t$ in $d(\mathfrak{M}) \setminus t$ must be strictly longer than m . In fact we are only interested in those paths between $\llbracket v \rrbracket$ and $\bigcup \llbracket \mathbf{z} \rrbracket$ that do *not* go through t : the paths that go through t are all the trivial paths between $\llbracket v \rrbracket$ and $\bigcup \llbracket \mathbf{z} \rrbracket$, and the goal is to find some $v^* \sim v$ such that all the non-trivial paths are long. The set t is a set of equivalence classes in \mathfrak{M} . By definition, it contains exactly those classes that contain both v and z_0 , so that some class $[v]_\beta$ is an element of t if and only if $(z_0, v) \in R_\beta^{\mathfrak{M}}$, which by 2-acyclicity is equivalently expressed as

$$t = \{[v]_\beta : \beta \supseteq \text{agt}(z_0, v)\} = \{[v]_\beta : [v]_\beta \supseteq [v]_{\text{agt}(z_0, v)}\}.$$

So the classes in t represent the coset paths of length 1 from z_0 to v . These are the trivial paths, which we cannot and need not avoid. But in order to win the Ehrenfeucht-Fraïssé game we need to be able to ensure that a response in a given round of the game can match long paths with long paths.

In order to find a suitable v^* , we will deal with each world $z \in \mathbf{z}$ in turn. First, we find a copy v_0 of v such that $d_t(\llbracket v_0 \rrbracket, \llbracket z_0 \rrbracket) > m$, then we find a copy v_1 such that $d_t(\llbracket v_1 \rrbracket, \llbracket z \rrbracket) > m$, for another world $z \in \mathbf{z}$, while maintaining $d_t(\llbracket v_1 \rrbracket, \llbracket z_0 \rrbracket) > m$, and so forth. The last of these copies will be v^* . Take note of the fact that we always need to avoid the same set $t = \llbracket v \rrbracket \cap \llbracket z_0 \rrbracket$ (rather than $\llbracket v \rrbracket \cap \llbracket z \rrbracket$) when we want to increase the distance between $\llbracket v \rrbracket$ and $\llbracket z \rrbracket$. This complicates things on a technical level. Note that $d_t(\llbracket v \rrbracket, \llbracket z \rrbracket) > 1$ implies $\llbracket v \rrbracket \cap \llbracket z \rrbracket \subseteq t$, which means that all the classes that directly connect v and z will also be avoided.

To find some suitable world v^* in the Cayley structure \mathfrak{M} , we consider paths in $d(\mathfrak{M}) \setminus t$ that need to be avoided as paths in \mathfrak{M} , as certain *coset paths* to be introduced below. We close this section with a useful alternative description of the set t that needs to be avoided. Motivated by freeness, t was defined as $\llbracket v \rrbracket \cap \llbracket z_0 \rrbracket$, i.e. in terms of v and z_0 . Since we assume \mathfrak{M} to be 2-acyclic, t can also be described in terms of v and the set $\gamma := \text{agt}(z_0, v)$ as $t = \{[v]_\beta : \beta \supseteq \gamma\}$, as we saw above. This motivates the following definition.

Definition 3.18. For a 2-acyclic Cayley frame \mathfrak{M} with dual hypergraph $d(\mathfrak{M})$, we define the following mapping:

$$\begin{aligned} \rho^{\mathfrak{M}}: W \times \tau &\longrightarrow \mathcal{P}(d(W)) \\ (v, \gamma) &\longmapsto \{[v]_\beta : \beta \supseteq \gamma\} \end{aligned}$$

If it is clear from the context, we drop the superscript \mathfrak{M} and just write ρ instead of $\rho^{\mathfrak{M}}$. Note that the set t to be avoided will typically be $t = \rho(v, \text{agt}(z_0, v))$.

The following lemma characterises the relationship of the sets $\llbracket v \rrbracket \cap \llbracket z_0 \rrbracket$ and $\llbracket v \rrbracket \cap \llbracket z \rrbracket$ in $d(\mathfrak{M})$ in terms of $\text{agt}(z_0, v)$ and $\text{agt}(v, z)$. We can observe the usual duality in the transition from Cayley structures to their dual hypergraphs.

Lemma 3.19. *Let \mathfrak{M} be a 2-acyclic Cayley frame, v, z two worlds and γ a set of agents, then*

$$\llbracket z \rrbracket \cap \llbracket v \rrbracket \subseteq \rho(v, \gamma) \quad \Leftrightarrow \quad \gamma \subseteq \text{agt}(z, v).$$

3.2.2 Coset paths

A special case of the coset paths to be considered here are the non- t coset paths for some $t = \rho(v, \gamma)$. Those are the coset paths that correspond to the chordless paths that avoid t in the dual hypergraph. Based on non- t coset paths we present a multi-scale measure of distance in Cayley graphs and results from [6] that tie it to the measure of distance that we use in hypergraphs.

Definition 3.20. Let \mathfrak{M} be a Cayley frame. A *coset path of length $\ell \geq 1$* is a labelled path $w_1, \alpha_1, w_2, \alpha_2, \dots, \alpha_\ell, w_{\ell+1}$ such that, for $1 \leq i \leq \ell$,

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

where we trivially supplement the path with $\alpha_0 = \alpha_{\ell+1} := \emptyset$.

A coset path $w_1, \alpha_1, \dots, \alpha_\ell, w_{\ell+1}$ of length $\ell \geq 2$ with $\alpha := \text{agt}(w_1, w_{\ell+1})$ is

- *non-trivial* if $[w_i]_{\alpha_i} \not\supseteq [w_1]_\alpha$
- an *inner path* if $[w_i]_{\alpha_i} \subsetneq [w_1]_\alpha$

for all $1 \leq i \leq \ell$. It is

- *non- t* for some $t = \rho(w_{\ell+1}, \gamma)$ with $\gamma \in \tau$, if $[w_i]_{\alpha_i} \not\supseteq [w_{\ell+1}]_\gamma$

for all $1 \leq i \leq \ell$. A non- t (or non-trivial) coset path from w to $v \neq w$ is *minimal* if there is no shorter non- t (or non-trivial) coset path from w to v .

Remark. A non-trivial coset path from w to v is the same as a non- t coset path for $t = \rho(v, \text{agt}(w, v))$. A coset path $w_1, \alpha_1, \dots, \alpha_\ell, w_{\ell+1}$ of length $\ell \geq 2$ is an inner path if $\text{agt}(w_1, w_{\ell+1}) = \bigcup_i \alpha_i$, and any such inner coset path is non-trivial.

Definition 3.21. Let \mathfrak{M} be a Cayley frame that is $2n$ -acyclic. We call a coset path *short* if its length is at most n .

Defining a measure of distance in \mathfrak{M} is a non-trivial matter because of its highly intricate edge pattern. Every Cayley frame is a single clique with respect to R_Γ , the accessibility relation induced by the set Γ of all agents. So the diameter of a Cayley frame is at most 1, which trivialises the usual notion of distance and renders locality techniques seemingly useless.

However, in 2-acyclic structures we can define a sensible notion of distance that is based on non- t coset paths. Essentially, a non- t coset path between w and v excludes all trivial connections between w and v and only looks at the scale that we are interested in, which is set by t .

Definition 3.22. Let \mathfrak{M} be a 2-acyclic Cayley frame, $w \neq v$ two worlds, $\gamma \subseteq \Gamma$ and $t = \rho(v, \gamma)$. The *t -distance $d_t(w, v)$* between w and v is defined as the length of a minimal coset path from w to v that avoids t (∞ if there is no such path). For a set of worlds \mathbf{z} , the *t -distance $d_t(\mathbf{z}, v)$* between $\mathbf{z} \neq \emptyset$ and v is defined as

$$d_t(\mathbf{z}, v) := \min_{z \in \mathbf{z}} d_t(z, v).$$

Remark. Depending on t , t -distance allows for distance 1: $d_t(w, v) = 1$ if, and only if, $[v]_{\text{agt}(w, v)} \notin t$. However, we are usually interested in cases where $\gamma \subseteq \text{agt}(w, v)$, which implies $[v]_{\text{agt}(w, v)} \in t$, for $t = \rho(v, \gamma)$ and thus $d_t(w, v) \geq 2$.

As suggested by our notation, $d_t(w, v)$ and $d_t(\llbracket w \rrbracket, \llbracket v \rrbracket)$ are closely related. It is shown in [6] that the two measures correspond for sufficiently acyclic Cayley structures because a path in one structure is the dual image of a path in the dual structure. Furthermore, it is shown that in sufficiently acyclic Cayley structures the t -distance between two worlds is not short if there are no short *inner* non- t coset paths. This result is of crucial importance because it reduces the global distance between w and v in \mathfrak{M} to a local phenomenon within $\mathfrak{M} \upharpoonright \text{agt}(w, v)$.

Lemma 3.23 ([6]). *Let $\ell \geq 1$, \mathfrak{M} a sufficiently acyclic Cayley frame, w, v two worlds, $\gamma \subseteq \Gamma$ and $t = \rho(v, \gamma)$. If there is no inner non- t coset path of length ℓ from w to v , then $d_t(w, v) > \ell$ and $d_t(\llbracket w \rrbracket, \llbracket v \rrbracket) > \ell - 1$.*

Hence, given w, v and t , finding some $v^* \sim v$ such that $d_t(\llbracket w \rrbracket, \llbracket v^* \rrbracket) > \ell - 1$ reduces to finding some $v^* \sim v$ such that $d_t(w, v^*) > \ell$, which reduces to the local matter of eliminating, in some sense, all the short inner non- t coset paths. In the following section, we prove the freeness theorem. Lemma 3.23 from [6] is the cornerstone for this undertaking.

3.2.3 The freeness theorem

Let $m \geq 1$, \mathfrak{M} be a Cayley structure, v a world, (\mathbf{z}, z_0) a pointed set with $v \notin \mathbf{z}$ and $\gamma = \text{agt}(z_0, v)$. The challenge is to find a world $v^* \sim v$ with $\text{agt}(z_0, v^*) = \gamma$ such that v^* and (\mathbf{z}, z_0) are m -free, assuming \mathfrak{M} is sufficiently acyclic and sufficiently rich. The necessary levels of richness and acyclicity are determined by m and $|\mathbf{z}|$. Hence, we need a suitable v^* such that $d_t(\bigcup \llbracket \mathbf{z} \rrbracket, \llbracket v^* \rrbracket) > m$, for $t = \llbracket z_0 \rrbracket \cap \llbracket v^* \rrbracket$; by Lemma 3.23 it suffices to have a v^* such that $d_t(\mathbf{z}, v^*) > m + 1$. Since we need such v^* for arbitrary m , we will show how to obtain a v^* such that $d_t(\mathbf{z}, v^*) > m$ in order to make the following more readable. Proving the freeness theorem involves two steps.

The first step

The first step is to find some $v_1 \sim v$ with $\text{agt}(z_0, v_1) = \gamma$ such that $d_t(\mathbf{z}, v_1) > 1$, for $t = \llbracket z_0 \rrbracket \cap \llbracket v \rrbracket = \rho(v, \gamma) = \rho(v_1, \gamma)$. The choice of t immediately implies $d_t(z_0, v) > 1$, but we need to look for an appropriate bisimilar copy of v in $[v]_\gamma$ to increase the t -distance to the other worlds of \mathbf{z} . The condition $d_t(\mathbf{z}, v_1) > 1$ can be equivalently rephrased as $\text{agt}(z_0, v_1) \subseteq \text{agt}(z, v_1)$, for all $z \in \mathbf{z}$. Lemma 3.7 guarantees the uniqueness of the minimal connecting sets of agents $\text{agt}(\cdot, \cdot)$ in 2-acyclic \mathfrak{M} and thereby implies a triangle inequality with respect to these:

$$\text{agt}(v, z) \subseteq \text{agt}(v, z_0) \cup \text{agt}(z_0, z).$$

If we find a bisimilar copy v_1 of v with $\text{agt}(z_0, v_1) = \gamma$ such that

$$\text{agt}(v_1, z) = \text{agt}(v_1, z_0) \cup \text{agt}(z_0, z),$$

then $\text{agt}(z_0, v_1) \subseteq \text{agt}(z, v_1)$. In other words, in the passage from v to v_1 we need to increase the distance, with regard to connecting agents, from z without changing the distance from z_0 . Lemma 3.25 shows that this can be done in 2-acyclic, rich structures for multiple $z \in \mathbf{z}$ simultaneously.

We use the following argument in the proofs of the Lemmas 3.25 and 3.27 to find suitable bisimilar copies in rich structures.

Lemma 3.24. *Let v be a world and (\mathbf{z}, z_0) a finite pointed set in a Cayley structure \mathfrak{M} , with $\text{agt}(v, z_0) \subseteq \text{agt}(v, z)$ for all $z \in \mathbf{z}$. Let $a \in \text{agt}(v, z_0)$; if \mathfrak{M} is 2-acyclic and sufficiently rich, then there is some $v' \in [v]_a \setminus \{v\}$ with $v' \sim v$ such that, for all $z \in \mathbf{z}$,*

$$\text{agt}(v', z) = \text{agt}(v, z).$$

Proof. Let $B_z = \{u \in [v]_a : a \notin \text{agt}(z, u)\}$ for $z \in \mathbf{z}$. By Lemma 3.9(2), $|B_z| \leq 1$ for all $z \in \mathbf{z}$. Let $B = \bigcup_{z \in \mathbf{z}} B_z$; then $|B| \leq |\mathbf{z}|$. Since \mathfrak{M} is sufficiently rich, there is some $v' \in [v]_a \setminus (B \cup \{v\})$ with $v' \sim v$. It follows $\text{agt}(v', z) = \text{agt}(v, z)$, for all $z \in \mathbf{z}$, from the definition of B . \square

In the statement of Lemma 3.25, the worlds of \mathbf{z} are the ones that have already been taken care of and u is the world to be processed next.

Lemma 3.25. *Let v, u be worlds and (\mathbf{z}, z_0) a finite pointed set of worlds in a Cayley structure \mathfrak{M} , with $\text{agt}(v, z) = \text{agt}(v, z_0) \cup \text{agt}(z_0, z)$ for all $z \in \mathbf{z}$. If \mathfrak{M} is 2-acyclic and sufficiently rich, then there is a world $v^* \sim v$ with $\text{agt}(v^*, z) = \text{agt}(v, z)$ for all $z \in \mathbf{z}$, such that also*

$$\text{agt}(v^*, u) = \text{agt}(v^*, z_0) \cup \text{agt}(z_0, u).$$

Proof. Put $\alpha_1 := \text{agt}(v, z_0)$, $\alpha_2 := \text{agt}(z_0, u)$ and $\alpha_3 := \text{agt}(u, v)$. By Lemma 3.7, 2-acyclicity implies $\alpha_i \subseteq \alpha_j \cup \alpha_k$ whenever $\{i, j, k\} = \{1, 2, 3\}$. We show that if $\alpha_3 \subsetneq \alpha_1 \cup \alpha_2$, then for every agent $a \in (\alpha_1 \cup \alpha_2) \setminus \alpha_3$ there is a world $v' \in [v]_a$ with $v' \sim v$ such that

- $\text{agt}(u, v') = \alpha_3 \cup \{a\}$;
- $\text{agt}(v', z_0) = \alpha_1$, and $\text{agt}(v', z) = \text{agt}(v, z)$, for all $z \in \mathbf{z}$;
- $\alpha_1 \subseteq \alpha_2 \cup \text{agt}(u, v')$, $\alpha_2 \subseteq \alpha_1 \cup \text{agt}(u, v')$, $\text{agt}(u, v') \subseteq \alpha_1 \cup \alpha_2$.

Since $(\alpha_1 \cup \alpha_2) \setminus \alpha_3$ is finite, applying this argument a finite number of times leads to a suitable world v^* with, in particular, $(\alpha_1 \cup \alpha_2) \setminus \text{agt}(u, v^*) = \emptyset$. Let $a \in (\alpha_1 \cup \alpha_2) \setminus \alpha_3$. Then $a \in \alpha_1$ because if we assume $a \notin \alpha_1$, it follows

$$a \notin \alpha_1 \xrightarrow{a \in \alpha_1 \cup \alpha_2} a \notin \alpha_1, a \in \alpha_2 \xrightarrow{\alpha_2 \subseteq \alpha_1 \cup \alpha_3} a \in \alpha_3.$$

Since \mathfrak{M} is sufficiently rich and 2-acyclic and $a \in \alpha_1 \subseteq \text{agt}(v, z)$, for all $z \in \mathbf{z}$, Lemma 3.24 gives us a world $v' \in [v]_a \setminus \{v\}$ with $v' \sim v$ and $\text{agt}(v', z) = \text{agt}(v, z)$, for all $z \in \mathbf{z}$. Set $\alpha'_1 := \text{agt}(v', z_0)$, $\alpha'_2 := \text{agt}(z_0, u)$ and $\alpha'_3 := \text{agt}(u, v')$. Lemma 3.9 implies $\alpha'_3 = \alpha_3 \cup \{a\}$ because $a \notin \alpha_3$ and \mathfrak{M} is 2-acyclic, $a \in \alpha_1$ implies $\alpha'_1 = \alpha_1$, and 2-acyclicity gives us $\alpha'_i \subseteq \alpha'_j \cup \alpha'_k$ whenever $\{i, j, k\} = \{1, 2, 3\}$. \square

The second step

The second step is the more difficult one. We have to establish $d_t(\mathbf{z}, v^*) > m$, while maintaining $\text{agt}(v^*, z_0) = \gamma$. By Lemma 3.23 this means that we need to eliminate short inner non- t coset paths between v and the worlds in \mathbf{z} by moving to bisimilar copies of v within $[v]_\gamma$.

There are many possible bisimilar copies of v to choose from. The key is to find a suitable $a \in \gamma$ such that an a -step to a bisimilar copy of v in $[v]_a$ brings us closer to v^* . We define the set of “right” agents in γ by describing the “wrong” agents, i.e. the direction one has to take if one wants to move on a short path from v towards \mathbf{z} . If we can do that, we just move in any other direction.

Again, we resort to a result from [6]. There it is shown that the direction one has to take if one wants to move from v to z on a short non- t coset path in a sufficiently acyclic structure is unique in the following sense: there is a non-empty set α_0 such that if v, α, \dots, z is a short non- t coset path, then $\alpha_0 \subseteq \alpha$. We denote this set as

$$\text{short}_t(v, z).$$

It is shown in [6] that this set exists and is unique if there is a short non- t coset path from v to z . It is our goal to use $\text{short}_t(v, z)$ to find a suitable bisimilar copy v^* of v such that the t -distance between v^* and z increases. This means that we must take a different direction, i.e. some $a \notin \text{short}_t(v, z)$, and move to a bisimilar copy in $[v]_a$. The idea is to repeat this procedure with different suitable agents until we reach a copy of v that has a sufficiently large t -distance to z .

The agent $a \notin \text{short}_t(v, z)$ can be chosen to be in γ : if $v, \text{short}_t(v, z), \dots, z$ is a short coset path (cf. Definition 3.21) that avoids t (recall that $t = \rho(v, \gamma) = \{[v]_\beta : \beta \supseteq \gamma\}$), then $[v]_\gamma \not\subseteq [v]_{\text{short}_t(v, z)}$ implying $\gamma \not\subseteq \text{short}_t(v, z)$. The case $\gamma \subseteq \text{agt}(v, z)$ is of particular interest in the proof of the freeness theorem.

Remark ([6]). Let \mathfrak{M} be a 2-acyclic Cayley structure, $v, z \in \mathfrak{M}$ and $\gamma \subseteq \text{agt}(v, z)$ a set of agents. Then, for $t = \rho(v, \gamma)$, $\gamma \not\subseteq \text{short}_t(v, z)$.

Similar to the set $\text{agt}(v, z)$ in 2-acyclic structures, $\text{short}_t(v, z)$ behaves in a controlled manner in sufficiently acyclic structures.

Lemma 3.26 ([6]). *Let $m \in \mathbb{N}$, \mathfrak{M} a Cayley frame, z, v two worlds, $\gamma \subseteq \text{agt}(v, z)$ and $t = \rho(v, \gamma)$. Assume \mathfrak{M} is $(2m + 1)$ -acyclic, $d_t(z, v) \leq m$, and that there are $a \notin \text{short}_t(v, z)$ and $v' \in [v]_a \setminus \{v\}$ such that $d_t(v', z) \leq m$. Then $a \in \text{short}_t(v', z)$.*

The agents in $\text{short}_t(v, z)$ are the ones that represent the direction one needs to take if one wants to move from v to z on a short non- t coset path. Lemma 3.26 makes this notion precise and tells us how to use $\text{short}_t(v, z)$. We choose an agent $a \notin \text{short}_t(v, z)$ and move to a world $v' \in [v]_a \setminus \{v\}$. If the structure is sufficiently acyclic, every short non- t coset path from v' to z must start with a set that includes agent a .

Lemma 3.26 is the cornerstone for the second step in the proof of the freeness theorem, which establishes $d_t(\mathbf{z}, v^*) > m$. It will be utilised as follows. Let

$w_1, \alpha_1, \dots, \alpha_\ell, w_{\ell+1}$ be a short inner non- t coset path from z to v , for $t = \rho(v, \gamma)$. Then

- (i) every set α_i , $1 \leq i \leq \ell$, is a proper subset of $\text{agt}(z, v)$;
- (ii) no class $[w_i]_{\alpha_i}$ contains $[v]_\gamma$; in other words, if $[w_i]_{\alpha_i} \cap [v]_\gamma \neq \emptyset$, then $\gamma \not\subseteq \alpha_i$.

In particular, the relevant sets of agents α_i are bounded in terms of $\text{agt}(z, v)$ and γ .

Assume that we move along an a_1 -edge from v to v_1 , then along an a_2 -edge from v_1 to v_2 and so forth, for suitable agents $a_i \in \gamma$ until none remain. Then the set $\text{short}_t(v_1, z)$ must contain a_1 , the set $\text{short}_t(v_2, z)$ must contain a_2 , etc. Let $v' \sim v$ be the final world in this sequence. If we assume that the distance $d_t(z, v')$ is still $d_t(z, v)$, then there must be a non- t inner coset path $w_1, \alpha_1, \dots, \alpha_\ell, w_{\ell+1}$ from z to v' of length $\ell = d_t(z, v')$. In this case it is possible to show that the set α_ℓ contains all a_i and the rest of γ , which means $\gamma \subseteq \alpha_\ell$. This contradicts the assumption that $w_1, \alpha_1, \dots, \alpha_\ell, w_{\ell+1}$ is a non- t coset path for $t = \rho(v, \gamma)$. The precise statement necessary for this argument is encapsulated in Lemma 3.27 below; its rather technical proof can be found in the appendix.

Lemma 3.27. *Let $m \in \mathbb{N}$, \mathfrak{M} be a Cayley structure, v a world, (\mathbf{z}, z_0) a finite pointed set, $\mathbf{y} \subseteq \mathbf{z}$ a possibly empty subset and $w \in \mathbf{z} \setminus \mathbf{y}$; let $\gamma = \text{agt}(z_0, v)$ and $t = \rho(v, \gamma)$. Assume that \mathfrak{M} is sufficiently acyclic and sufficiently rich, and*

- $\gamma \subseteq \text{agt}(z, v)$, for all $z \in \mathbf{z}$;
- $d_t(\mathbf{y}, v) > m$.

Then there is a world $v^ \in [v]_\gamma$ with $\mathfrak{M}, v^* \sim \mathfrak{M}, v$ and $\text{agt}(z, v^*) = \text{agt}(z, v)$, for all $z \in \mathbf{z}$, such that*

$$d_t(\mathbf{y} \cup \{w\}, v^*) > m.$$

The freeness theorem

This section is devoted to the proof of the freeness theorem, which is the crucial tool for choosing suitable responses in the Ehrenfeucht-Fraïssé game on Cayley structures that are sufficiently rich and acyclic. The main ingredients are Lemma 3.25 for the first step and Lemma 3.27 for the second step.

Theorem 3.28 (freeness theorem). *Let $m, k \in \mathbb{N}$. If a Cayley structure \mathfrak{M} is sufficiently acyclic and sufficiently rich, then \mathfrak{M} is (m, k) -free.*

Proof. Let v be a world, (\mathbf{z}, z_0) a pointed set of size $|\mathbf{z}| = k$ enumerated as $(z_i)_{0 \leq i < k}$, and $\gamma \supseteq \text{agt}(v, z_0)$. We show that there is a world $v^* \sim v$ with $\text{agt}(v^*, z_0) = \gamma$ such that v^* and (\mathbf{z}, z_0) are m -free.

Preparation: 2-acyclicity and Lemma 3.9, together with sufficient richness, imply the existence of some $v' \sim v$ with $\text{agt}(v', z_0) = \gamma$. Replace v by this world v' so that $\text{agt}(v, z_0) = \gamma$.

We must now find some world $v^* \sim v$ with $\text{agt}(v^*, z_0) = \gamma$ such that $d_t(\cup \llbracket \mathbf{z} \rrbracket, \llbracket v^* \rrbracket) > m$, for $t := \rho(v, \gamma) = \llbracket v \rrbracket \cap \llbracket z_0 \rrbracket$. By Lemma 3.23 it suffices to show $d_t(\mathbf{z}, v^*) > m + 1$. We do this in two steps. Step 1 ensures

$\text{agt}(z_0, v^*) \subseteq \text{agt}(z, v^*)$ (which is equivalent to $d_t(\mathbf{z}, v^*) > 1$) for all $z \in \mathbf{z}$, and step 2 ensures $d_t(\mathbf{z}, v^*) > m + 1$.

Step 1: we show by induction on $0 \leq j < k$ that there are worlds $v_j \sim v$ such that $d_t(z_i, v_j) > 1$, for all $0 \leq i \leq j$. The base case works for $j = 0$, $v_0 := v$. For the induction step let $j \geq 1$ and assume there is a world $v_{j-1} \sim v$ with $\text{agt}(z_0, v_{j-1}) = \gamma$ such that $\text{agt}(z_0, v_{j-1}) \subseteq \text{agt}(z_i, v_{j-1})$ for all $0 \leq i < j$. Together with 2-acyclicity we get

$$\begin{aligned} \text{agt}(z_0, z_i) &\subseteq \text{agt}(z_0, v_{j-1}) \cup \text{agt}(v_{j-1}, z_i) = \text{agt}(v_{j-1}, z_i), \text{ and then} \\ \text{agt}(v_{j-1}, z_i) &\subseteq \text{agt}(v_{j-1}, z_0) \cup \text{agt}(z_0, z_i) \subseteq \text{agt}(v_{j-1}, z_i). \end{aligned}$$

Hence, $\text{agt}(v_{j-1}, z_i) = \text{agt}(v_{j-1}, z_0) \cup \text{agt}(z_0, z_i)$, for all $0 \leq i < j$. Now Lemma 3.25 yields a world v_j with $\text{agt}(v_j, z_j) = \text{agt}(v_j, z_0) \cup \text{agt}(z_0, z_j)$ so that $\text{agt}(v_j, z_j) \supseteq \text{agt}(v_j, z_0)$.

In particular, we obtained a world $v_{k-1} \sim v$ with $\text{agt}(v_{k-1}, z_0) = \gamma$ such that $\text{agt}(z_0, v_{k-1}) \subseteq \text{agt}(z_i, v_{k-1})$, for all $0 \leq i \leq k-1$, or equivalently, $d_t(\mathbf{z}, v_{k-1}) > 1$, for $t = \rho(v_{k-1}, \gamma)$. We set the new v to be v_{k-1} .

Step 2: we show by induction on $0 \leq i \leq k$ that there are worlds $v_i \in [v]_\gamma$ such that

- $v_i \sim v$,
- $\text{agt}(v_i, z) = \text{agt}(v, z)$, for all $z \in \mathbf{z}$, and
- $d_t(\mathbf{y}_i, v_i) > m + 1$, for $\mathbf{y}_i := \{z_j \in \mathbf{z} : 0 \leq j < i\}$.

The base case works for $v_0 := v$ with $\mathbf{y}_0 = \emptyset$. For the induction step let $0 \leq i < k$ and assume there is a world $v_i \in [v]_\gamma$ with $v_i \sim v$, $\text{agt}(v_i, z) = \text{agt}(v, z)$, for all $z \in \mathbf{z}$, and $d_t(\mathbf{y}_i, v_i) > m + 1$. Since $\gamma \subseteq \text{agt}(z, v)$, for all $z \in \mathbf{z}$, and $z_i \in \mathbf{z} \setminus \mathbf{y}_i$, Lemma 3.27 implies a world $v_{i+1} \in [v_i]_\gamma = [v]_\gamma$ with

- $v_{i+1} \sim v_i$,
- $\text{agt}(z, v_{i+1}) = \text{agt}(z, v_i) = \text{agt}(z, v)$, for all $z \in \mathbf{z}$, and
- $d_t(\mathbf{y}_i \cup \{z_i\}, v_{i+1}) > m + 1$.

We obtain the desired world $v^* = v_k$ by induction. \square

4 Characterisation theorem

Our main result is a modal characterisation theorem for common knowledge logic $\text{ML}[\text{CK}]$ over (finite) S5 structures. This section contains the final step of its proof. We described the strategy for the proof at the end of Section 2: if we can show that an FO -formula φ that is \sim -invariant over (finite) CK structures is \sim^ℓ -invariant over (finite) CK structures, for some $\ell \in \mathbb{N}$, then φ must be equivalent to an ML -formula over (finite) CK structures by the modal Ehrenfeucht-Fraïssé theorem. This is done by upgrading ℓ -bisimilarity to FO_q -equivalence over (finite) Cayley structures, i.e. we show for suitable pointed Cayley structures \mathfrak{M}, w and \mathfrak{N}, v that

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

where q is the quantifier rank of φ and ℓ depends on q . Upgrading over Cayley structures suffices because by Lemma 2.7 (Main Lemma 1.3) Cayley structures are, up to bisimulation, the universal representatives of CK structures. For the upgrading, we regard a Cayley structure as suitable if it is n -acyclic and k -rich, for sufficiently large $n, k \in \mathbb{N}$ that depend on q . The construction of sufficiently acyclic and rich (finite) coverings for (finite) CK structures, for the first part of the upgrading argument, was presented in Sections 2.4 and 3.1. It remains to show that sufficiently acyclic and rich ℓ -bisimilar Cayley structures are FO_q -equivalent. The necessary structure theory for playing first-order Ehrenfeucht-Fraïssé games on the non-elementary class of Cayley structures was developed in Section 3 partly based on results from [17, 6]. Its central notion of freeness will now play a crucial rôle in the analysis of the Ehrenfeucht-Fraïssé game to prove \equiv_q -equivalence of suitable Cayley structures on the basis of \sim^ℓ -equivalence for sufficiently large ℓ .

Sketch of the core idea. In order to win the q -round Ehrenfeucht-Fraïssé game on \mathfrak{M}, w and \mathfrak{N}, v player **II** needs to keep track of several features, which are built into an invariant to be maintained through the successive rounds. First and foremost we incorporate an increasing chain of partial isomorphisms $(\sigma_i)_{i \leq q}$ between \mathfrak{M} and \mathfrak{N} , where $\sigma_0 = \{w \mapsto v\}$ and, for $i < q$, $\sigma_{i+1} \supseteq \sigma_i$ covers the elements newly pebbled in the $(i+1)$ -st round. As the invariant needs to be good for the remaining rounds, $\sigma_i: \mathfrak{M} \upharpoonright \text{dom}(\sigma_i) =: \mathfrak{M}_i \simeq \mathfrak{N}_i := \mathfrak{N} \upharpoonright \text{im}(\sigma_i)$ has to include worlds that lie on short paths between pebbled worlds. What ‘short’ means in relation to the number of rounds still to be played, is specified by a suitably chosen decreasing sequence of critical distances $(m_i)_{i \leq q}$. To guarantee extendability throughout the remaining rounds, we also need σ_i to preserve the ℓ_i -bisimulation type of elements for a suitably chosen decreasing sequence $(\ell_i)_{i \leq q}$ starting from $\ell_0 = \ell$, the degree of initial bisimilarity in $\mathfrak{M}, w \sim^\ell \mathfrak{N}, v$. The extension steps $\sigma_{i+1} \supseteq \sigma_i$ reflect **II**’s response to **I**’s challenge in round $i+1$, which also updates the auxiliary information in the invariant. Part of this auxiliary information resides in the dual hypergraphs $d(\mathfrak{M})$ and $d(\mathfrak{N})$. The invariant includes substructures of the dual hypergraphs, which essentially are the dual images of \mathfrak{M}_i and \mathfrak{N}_i (cf. Definition 3.12 for the dual hypergraph). In fact, the choice of the decreasing sequence of critical distances $(m_i)_{i \leq q}$ (for short distances in round i) is rooted in the dual hypergraphs, where results from [17] can be used to bound the sizes of convex m_i -closures (cf. Definition 3.14 and Lemma 3.16); and these bounds in turn determine the decreasing sequence $(\ell_i)_{i \leq q}$ of required bisimulation levels between elements of \mathfrak{M} and \mathfrak{N} that are linked by σ_i . The rather complex overall structure of the invariant, which is schematically presented in Figure 2, is formally presented, and shown to be maintainable through q rounds, in Section 4.1. The proof of the upgrading and the characterisation theorem is then completed in Section 4.2.




Figure 2: A snapshot of the invariant, after round i of the game on $\mathfrak{M}, w; \mathfrak{N}, v$, with pebbles on $\mathbf{w} = (w, w_1, \dots, w_i)$ and $\mathbf{v} = (v, v_1, \dots, v_i)$, based on isomorphic underlying tree decompositions $\mathcal{T}_i = (T_i, \delta_i) \simeq (T_i, \delta'_i) = \mathcal{T}'_i$ of acyclic subhypergraphs $d(\mathfrak{M}) \upharpoonright Q_i \simeq d(\mathfrak{N}) \upharpoonright Q'_i$.

4.1 The invariant

Definition of the invariant. Player **II** wins a play in the q -round Ehrenfeucht-Fraïssé game on the pointed Kripke structures \mathfrak{M}, w_0 and \mathfrak{N}, v_0 if $(w_i \mapsto v_i)_{0 \leq i \leq q}$ induces a partial isomorphism σ_q , for the pebbled worlds $w_0, w_1, \dots, w_q \in W$ and $v_0, v_1, \dots, v_q \in V$. Starting with the mapping $\sigma_0 = \{w_0 \mapsto v_0\}$ before the first round is played, **II** extends σ_i , the partial isomorphism after the i -th round in response to player **I**'s challenges, from round to round. In order to be able to do that for q rounds in a foresighted manner, she needs to keep track of more information than just the current σ_i . This auxiliary information is built into the invariant, which depicts the current game position and maintains a measure of the similarity of between the \mathfrak{M} - and the \mathfrak{N} -parts. The required degree of similarity depends on the number of rounds still to be played (decreasing as $q - i$ with the index i of rounds). It is governed by two decreasing sequences of natural numbers, a locality parameter and a bisimilarity parameter. The locality parameter in the sequence $(m_i)_{0 \leq i \leq q}$ indicates that distances up to m_i are considered short in the i -th round; the bisimilarity parameter in the sequence $(\ell_i)_{0 \leq i \leq q}$ specifies the degree of bisimilarity that worlds $w' \in \mathfrak{M}$ and $v' \in \mathfrak{N}$ need to display if they are matched (by σ_i) in round i . As usual in Ehrenfeucht-Fraïssé games, the locality parameter m_i decreases by about one half in each round, as reflected in the recursive definition

$$m_q := 2 \quad \text{and} \quad m_{i-1} := 2m_i + 1 \quad \text{for } 0 < i \leq q;$$

and the recursive definition of the sequence $(\ell_i)_{0 \leq i \leq q}$ refers to the functions f_m from Lemma 3.16, which bound the size of m -closed sets, for the levels $m = m_i$:

$$\ell_q := 1 \quad \text{and} \quad \ell_{i-1} := \ell_i + f_{m_i}(|\tau| + 1) \quad \text{for } 0 < i \leq q.$$

The structural backbone of the invariant is a tree decomposition of matching

representations of relevant substructures $\mathfrak{M}_i \subseteq \mathfrak{M}$ and $\mathfrak{N}_i \subseteq \mathfrak{N}$ in their dual hypergraphs $d(\mathfrak{M})$ and $d(\mathfrak{N})$. These are formalised as $\mathcal{T}_i = (T_i, \delta_i)$ and $\mathcal{T}'_i = (T_i, \delta'_i)$ based on the same tree T_i but representing, as tree decompositions, acyclic induced sub-hypergraphs $d(\mathfrak{M}) \upharpoonright Q_i \subseteq d(\mathfrak{M})$ and $d(\mathfrak{M}) \upharpoonright Q_i \subseteq d(\mathfrak{N})$. These tree decompositions serve as the scaffolding for the updates and extensions that need to be performed from round to round. The common tree structure T_i in their tree decompositions, in particular, governs the structural similarity between the \mathfrak{M} - and \mathfrak{N} -parts of the current game position.

With respect to the dual hypergraphs $d(\mathfrak{M})$ and $d(\mathfrak{N})$ compare Definition 3.6 and discussion there. Note that \mathfrak{M} is FO-interpretable in $d(\mathfrak{M})$ if we identify $w \in M$ with the vertex $d(w) := [w]_\emptyset \in Q_\emptyset \subseteq d(W)$ and note that $(w, w') \in R_\alpha$ iff $[w']_\alpha \in [w]$ iff there is some vertex $a \in Q_\alpha \subseteq d(W)$ that is joined by hyperedges to both $[w]_\emptyset$ and to $[w']_\emptyset$. We also note in this connection that $\llbracket w \rrbracket$ is the unique hyperedge of (\mathfrak{M}) that is incident on $[w]_\emptyset$; so also $(w, w') \in R_\alpha$ iff $\llbracket w \rrbracket$ and $\llbracket w' \rrbracket$ overlap in a vertex of colour Q_α . In particular, $d(\mathfrak{M}), [w]_\emptyset$ and $d(\mathfrak{N}), [v]_\emptyset$ determine whether $\mathfrak{M}, w \equiv_q \mathfrak{N}, v$.

The full invariant can be described as follows, assuming that the worlds $w_0, w_1, \dots, w_i \in W$ (where $w_0 = w$) and $v_0, v_1, \dots, v_i \in V$ (where $v_0 = v$) have been pebbled after the i -th round (cf. Figure 2):

- (I1) two isomorphic induced substructures $\mathfrak{M}_i \subseteq \mathfrak{M}$ and $\mathfrak{N}_i \subseteq \mathfrak{N}$ that comprise the pebbled worlds in each structure, with an isomorphism

$$\sigma_i: \mathfrak{M}_i, w_0, \dots, w_i \simeq \mathfrak{N}_i, v_0, \dots, v_i$$

that preserves \sim^{ℓ_i} -types: $\mathfrak{M}, w \sim^{\ell_i} \mathfrak{N}, \sigma_i(w)$ for all $w \in \mathfrak{M}_i$;

- (I2) m_i -closed subsets $Q_i \subseteq d(W)$ and $Q'_i \subseteq d(V)$, that comprise the dual images $[w_j]_\emptyset = \{w_j\} \in d(W)$ and $[v_j]_\emptyset = \{v_j\} \in d(V)$ for $j \leq i$;
- (I3) isomorphic tree decompositions $\mathcal{T}_i = (T_i, \delta_i)$ of $d(\mathfrak{M}) \upharpoonright Q_i$ and $\mathcal{T}'_i = (T_i, \delta'_i)$ of $d(\mathfrak{N}) \upharpoonright Q'_i$ of the sub-hypergraphs $d(\mathfrak{M}) \upharpoonright Q_i \subseteq d(\mathfrak{M})$ and $d(\mathfrak{N}) \upharpoonright Q'_i \subseteq d(\mathfrak{N})$;
- (I4) associated maps $\hat{\delta}_i: T_i \rightarrow W$ and $\hat{\delta}'_i: T_i \rightarrow V$ that pick representatives $w_u := \hat{\delta}_i(u) \in \bigcap \delta(u)$ in \mathfrak{M}_i and $v_u := \hat{\delta}'_i(u) \in \bigcap \delta'(u)$ in \mathfrak{N}_i for $u \in T_i$ such that $\sigma_i(w_u) = v_u$.⁴

The induced substructures \mathfrak{M}_i and \mathfrak{N}_i comprise all worlds pebbled during the first i rounds together with, essentially, their closure under short coset paths. These closures are induced by the m_i -closures Q_i and Q'_i . The worlds of \mathfrak{M}_i and \mathfrak{N}_i arise as appropriate choices of the representatives $\hat{\delta}_i(u)$ and $\hat{\delta}'_i(u)$ in \mathfrak{M} and \mathfrak{N} for the bags of their isomorphic tree decompositions of the acyclic induced sub-hypergraphs $d(\mathfrak{M}) \upharpoonright Q_i$ and $d(\mathfrak{N}) \upharpoonright Q'_i$. The matching tree decompositions of $d(\mathfrak{M}) \upharpoonright Q_i$ and $d(\mathfrak{M}) \upharpoonright Q_i$ form the structural backbone of the invariant; and the careful choice of matching representatives in \mathfrak{M} and \mathfrak{N} govern player \mathbf{II} 's strategy

⁴Note that, as subsets of the universes $d(W)$ and $d(V)$ of the dual hypergraphs, $\delta_i(u)$ and $\delta'_i(u)$ are sets of equivalence classes, i.e. sets of subsets in M and N ; and, e.g. $w_u \in \bigcap \delta_i(u)$ means that $\delta_i(u) \subseteq \llbracket w_u \rrbracket$.




Figure 3: Analysis of pebble placement $w_i \in \mathfrak{M}$ in round i , in terms of extensions Q_i , \mathcal{T}_i , δ_i and $\hat{\delta}_i$ towards finding matching extensions Q'_i , \mathcal{T}'_i , δ'_i , $\hat{\delta}'_i$ and $\mathfrak{N}_i, \mathbf{v}v_i$.

to maintain the invariant in response to player **I**'s move in each round, over and above the actual placement of a single pebble. If player **II** manages to maintain the invariant throughout the q -round game, she wins after round q since σ_q is a partial isomorphism that matches pebble positions: $\sigma_q: \mathfrak{M}_q, w_0, \dots, w_q \simeq \mathfrak{N}_q, v_0, \dots, v_q$ for the induced substructures $\mathfrak{M}_q \subseteq \mathfrak{M}$ and $\mathfrak{N}_q \subseteq \mathfrak{N}$.

The invariant is initialised, for $i = 0$, as $\mathfrak{M}_0 := \mathfrak{M} \upharpoonright \{w_0\}$, $\mathfrak{N}_0 := \mathfrak{N} \upharpoonright \{v_0\}$ with $\sigma_0: w_0 \mapsto v_0$ such that $\mathfrak{M}, w = \mathfrak{M}, w_0 \sim^{\ell_0} \mathfrak{N}, v_0 = \mathfrak{N}, v$ is given; we let $Q_0 := \{[w_0]_\emptyset\}$ and $Q'_0 := \{[v_0]_\emptyset\}$, which are m_0 -closed by the acyclicity requirements on \mathfrak{M} and \mathfrak{N} ; and $d(\mathfrak{M}) \upharpoonright Q_0$ and $d(\mathfrak{N}) \upharpoonright Q'_0$ are trivially acyclic with tree decompositions $\mathcal{T}_0, \mathcal{T}'_0$, based on the trivial tree $T_0 = \{\lambda\}$ consisting of just its root.

The idea behind the strategy to maintain the invariant through round i , and if player **I** puts a pebble on $w_i \in \mathfrak{M}$ say, can be roughly described as follows (cf. Figure 3). We set $Q_i := \text{cl}_{m_i}(Q_{i-1} \cup \{[w_i]_\emptyset\})$. By our acyclicity assumptions, the substructure $d(\mathfrak{M}) \upharpoonright Q_i$ is acyclic (cf. Lemma 3.16) and admits a tree decomposition $\mathcal{T}_i = (T_i, \delta_i)$ extending the tree decomposition \mathcal{T}_{i-1} by a single new subtree (cf. Lemmas 3.15 and 3.16). Elements $\hat{\delta}_i(u)$ representing the new bags $\delta_i(u)$ for $u \in T_i \setminus T_{i-1}$ are chosen under suitable constraints that avoid unwanted short paths, crucially based on freeness criteria. As $\llbracket w_i \rrbracket$ is the unique hyperedge of $d(\mathfrak{M})$ that is incident with $d(w_i) = [w_i]_\emptyset = \{w_i\} \in Q_i$, the bag $\delta_i(u)$ containing this element of $d(\mathfrak{M}) \upharpoonright Q_i$ is necessarily represented by $\hat{\delta}_i(u) = w_i$. Putting $\mathfrak{M}_i := \mathfrak{M} \upharpoonright \text{image}(\hat{\delta}_i)$, w_i is part of the extension $\mathfrak{M}_i \supseteq \mathfrak{M}_{i-1}$ which we view as a representation in \mathfrak{M} of the tree decompositions $\mathcal{T}_i \supseteq \mathcal{T}_{i-1}$ of $d(\mathfrak{M}) \upharpoonright Q_i \supseteq d(\mathfrak{M}) \upharpoonright Q_{i-1}$.

The challenge for **II** lies in matching these extensions $Q_i \supseteq Q_{i-1}$, $\mathcal{T}_i \supseteq \mathcal{T}_{i-1}$ with $\delta_i \supseteq \delta_{i-1}$, and $\hat{\delta}_i \supseteq \hat{\delta}_{i-1}$ on the side of \mathfrak{N} and $d(\mathfrak{N})$ in order to maintain the invariant with all its constraints (I1)–(I4). The following section presents a detailed discussion.

Maintaining the invariant. We show how to maintain the invariant through round i , w.l.o.g. in response to a placement of the i th pebble on $w_i \in \mathfrak{M} \setminus \mathfrak{M}_i$. We assume the invariant after round $i - 1$ provides sets $Q_{i-1} \subseteq d(W)$ and $Q'_{i-1} \subseteq d(V)$ inducing acyclic sub-hypergraphs of $d(\mathfrak{M})$ and $d(\mathfrak{N})$ with tree

decompositions $\mathcal{T}_{i-1} = (T_{i-1}, \delta_{i-1})$ of $d(\mathfrak{M}) \upharpoonright Q_{i-1}$ and $\mathcal{T}'_{i-1} = (T_{i-1}, \delta'_{i-1})$ of $d(\mathfrak{N}) \upharpoonright Q'_{i-1}$ and surjective maps $\hat{\delta}_{i-1}: T_{i-1} \mapsto \mathfrak{M}_{i-1}$ and $\hat{\delta}'_{i-1}: T_{i-1} \mapsto \mathfrak{N}_{i-1}$ such that $\mathfrak{M}_{i-1} := \mathfrak{M} \upharpoonright \text{image}(\hat{\delta}_{i-1})$ and $\mathfrak{N}_{i-1} := \mathfrak{N} \upharpoonright \text{image}(\hat{\delta}'_{i-1})$ are isomorphic via

$$\sigma_{i-1}: \mathfrak{M}_{i-1}, w_0, \dots, w_{i-1} \simeq \mathfrak{N}_{i-1}, v_0, \dots, v_{i-1},$$

which preserves $\sim^{\ell_{i-1}}$ and is compatible with $\hat{\delta}_{i-1}$ and $\hat{\delta}'_{i-1}$ in the sense that the following diagram commutes:

$$\begin{array}{ccc} \mathfrak{M}_{i-1} & \xrightarrow{\sigma_{i-1}} & \mathfrak{N}_{i-1} \\ & \nwarrow \hat{\delta}_{i-1} & \nearrow \hat{\delta}'_{i-1} \\ & T_{i-1} & \end{array}$$

Player **I** having pebbled $w_i \in \mathfrak{M} \setminus \mathfrak{M}_i$ as in round i , we analyse the new configuration on the side of \mathfrak{M} in terms of extensions Q_i , \mathcal{T}_i , $\hat{\delta}_i$, which then need to be matched on the side of \mathfrak{N} . For $Q_i := \text{cl}_{m_i}(Q_{i-1} \cup \{[w_i]_{\emptyset}\})$ we have that

- $[w_i]_{\emptyset} \notin Q_{i-1}$ as $w_i \notin \mathfrak{M}_{i-1}$;
- $1 \leq d(Q_{i-1}, [w_i]_{\emptyset}) \leq 2 \leq m_i$ as $\text{diam}(d(\mathfrak{M})) = 2$.

As $d(\mathfrak{M})$ is sufficiently acyclic, Lemma 3.15 can be applied. Denoting as $D := Q_{i-1} \cap N^1(Q_i \setminus Q_{i-1})$ the region in which the extended closure attaches to Q_{i-1} , we know that $Q_i \setminus Q_{i-1}$ is connected, and that D separates $Q_i \setminus Q_{i-1}$ from $Q_{i-1} \setminus D$ so that $Q_i = Q_{i-1} \cup \text{cl}_{m_i}(D \cup \{[w_i]_{\emptyset}\})$. In fact D is a separator in the graph-theoretic sense so that every path linking $Q_{i-1} \setminus D$ to $Q_i \setminus Q_{i-1}$ in $d(\mathfrak{M}) \upharpoonright Q_i$ must go through D . By Lemma 3.15, D is a clique since Q_{i-1} is $(2m_i+1)$ -closed ($m_{i-1} = 2m_i+1$). By Lemma 3.16, the size of $\text{cl}_{m_i}(D \cup \{[w_i]_{\emptyset}\})$ is bounded by $f_{m_i}(|\tau| + 1)$, which implies that $d(\mathfrak{M}) \upharpoonright Q_i$ is tree decomposable. Let

$$Q := (Q_i \setminus Q_{i-1}) \cup D,$$

and let $u_0 \in T_{i-1}$ be a node of \mathcal{T}_{i-1} representing D , i.e. with $\delta_{i-1}(u_0) \supseteq D$. The rôle of D as separator in $d(\mathfrak{M}) \upharpoonright Q_i$ implies that $d(\mathfrak{M}) \upharpoonright Q_i$ admits a tree decomposition obtained as the fusion of the tree decomposition \mathcal{T}_{i-1} of $d(\mathfrak{M}) \upharpoonright Q_{i-1}$ and a tree decomposition $\mathcal{T} = (T, \delta)$ of $d(\mathfrak{M}) \upharpoonright Q$ whose root node λ represents D as $D = \delta(\lambda)$. The resulting tree decomposition $\mathcal{T}_i = (T_i, \delta_i)$ extends $\mathcal{T}_{i-1} = (T_{i-1}, \delta_{i-1})$ and is obtained by attaching the root λ of \mathcal{T} to the node u_0 of T_{i-1} that represents D . So T becomes the subtree of new nodes in T_i and $\delta_i = \delta_{i-1} \dot{\cup} \delta$. We correspondingly extend $\hat{\delta}_{i-1}: T_{i-1} \rightarrow \mathfrak{M}$ to $\hat{\delta}_i: T_i \rightarrow \mathfrak{M}$ such that $\hat{\delta}_i(u) \subseteq \llbracket \hat{\delta}_i(u) \rrbracket$, for all $u \in T_i$. Here again $\hat{\delta}_i = \hat{\delta}_{i-1} \dot{\cup} \hat{\delta}$ for $\hat{\delta}: T \rightarrow \mathfrak{M}$ stemming from the tree decomposition \mathcal{T} of $d(\mathfrak{M}) \upharpoonright Q$. In particular we choose $\hat{\delta}(\lambda) = \hat{\delta}_i(\lambda) = \hat{\delta}_{i-1}(u_0) \in \mathfrak{M}_{i-1}$.

The second stage of the construction is about finding mirror images of these augmentations on the side of \mathfrak{N} . These then yield the desired update of the invariant after round i , and in particular a response move for player **II**.

The given invariant from round $i - 1$ provides an isomorphic image $D' \subseteq \delta'_{i-1}(u_0)$ in $d(\mathfrak{N})$ of the set $D \subseteq \delta_{i-1}(u_0)$, and $\sigma_{i-1}(\hat{\delta}_{i-1}(u_0)) = \hat{\delta}'_{i-1}(u_0) \in \mathfrak{N}_{i-1} \subseteq \mathfrak{N}$ as an isomorphic copy of $\hat{\delta}_{i-1}(u_0) \in \mathfrak{M}$. Starting from this image, we need to find suitable isomorphic images for the remainder of Q . We do this through an induction on the structure of the tree T . Starting with a child of the root, we inductively add nodes in a breadth-first manner. In other words, we construct a suitable extension of \mathcal{T}'_{i-1} by growing a subtree isomorphic to \mathcal{T} at $\delta'_{i-1}(u_0)$ in \mathcal{T}'_{i-1} .

In order to construct this extension we first describe the structure of \mathcal{T} in its manifestations in $d(\mathfrak{M}) \upharpoonright Q \subseteq d(\mathfrak{M})$ and $\mathfrak{M} \upharpoonright \text{image}(\hat{\delta}) \subseteq \mathfrak{M}$ by an ML-formula. We make use of the fact that in dual hypergraphs of 2-acyclic Cayley structures the set of equivalence classes $\llbracket u \rrbracket \cap \llbracket u' \rrbracket$ is fully determined by the set of agents $\text{agt}(u, v)$ (cf. Lemma 3.7), i.e.

$$\llbracket u \rrbracket \cap \llbracket v \rrbracket = \{[u]_\beta : \beta \supseteq \text{agt}(u, v)\} = \{[v]_\beta : \beta \supseteq \text{agt}(u, v)\}.$$

This allows us to describe the overlap between the two bags at next neighbours u and v of the tree decomposition by a modality induced by the single set $\text{agt}(u, v)$ of agents. We then follow this description to find a suitable bisimilar image of $\text{image}(\hat{\delta})$ in \mathfrak{N} in order to extend \mathfrak{N}_{i-1} to \mathfrak{N}_i .

Let $w_u := \hat{\delta}(u)$, for $u \in T$, and $\lambda \in T$ the root node with $D = \delta(\lambda)$, $w_\lambda = \hat{\delta}_{i-1}(u_0) \in \mathfrak{M}_{i-1}$ as above. We describe the finite substructure $\mathfrak{M} \upharpoonright \text{image}(\hat{\delta})$, w_λ by a formula $\varphi_{\mathcal{T}} := \varphi_{\mathcal{T}, \lambda} \in \text{ML}$ of modal depth ℓ_{i-1} . Inductively, for $u \in T$ the formula $\varphi_{\mathcal{T}, u}$ is defined as follows. For $u \in T$ let χ_u be the formula of modal depth ℓ_i that describes the ℓ_i -bisimulation type of w_u . If $u \in T$ has children u_1, \dots, u_k with associated formulae $\varphi_{\mathcal{T}, u_j}$ already defined, let $\alpha_j := \text{agt}(w_u, w_{u_j})$ and put

$$\varphi_{\mathcal{T}, u} := \chi_u \wedge \bigwedge_{1 \leq j \leq k} \Diamond_{\alpha_j} \varphi_{\mathcal{T}, u_j}.$$

The modal nesting depth of $\varphi_{\mathcal{T}} := \varphi_{\mathcal{T}, \lambda}$ is uniformly bounded by $\ell_i + \text{depth}(\mathcal{T}) \leq \ell_i + f_{m_i}(|\tau| + 1)$, as the depth can be bounded by the size of the relevant m_i -closure. This guarantees that $\varphi_{\mathcal{T}}$ is preserved under $\sim^{\ell_{i-1}}$ and hence by σ_{i-1} . The subformulae χ_u , which characterise the ℓ_i -bisimulation type of every world w_u for $u \in T$, make sure that matching worlds in \mathfrak{N}_i can support an isomorphism $\sigma_i: \mathfrak{M}_i \simeq \mathfrak{N}_i$ that respects ℓ_i -bisimilarity.

The desired matches in $\mathfrak{N}_i \supseteq \mathfrak{N}_{i-1}$ and $Q'_i \supseteq Q'_{i-1}$ are obtained as follows.

Starting from $v_\lambda := \sigma_{i-1}(w_\lambda)$, we use that $\mathfrak{M}, w_\lambda \models \varphi_{\mathcal{T}}$ translates into $\mathfrak{N}, v_\lambda \models \varphi_{\mathcal{T}}$, for $v_\lambda := v_{u_0} = \sigma_{i-1}(w_{u_0}) = \sigma_{i-1}(w_\lambda)$, because $\text{md}(\varphi_{\mathcal{T}}) \leq \ell_{i-1}$. We proceed by induction on T , starting from the root λ and progressing in a depth-first manner. At the root, we simply have to replicate the stipulations from level $i - 1$: $v_\lambda = \sigma_{i-1}(w_\lambda) =: \sigma_i(w_\lambda)$ and $\delta'_i(\lambda) := \{[v_\lambda]_\beta : [w_\lambda]_\beta \in D\}$. For a typical extension step, consider nodes at depth 1 from the root. We need to find, for each child u of λ in T , a suitable world $v_u \in V$ with $\mathfrak{N}, v_u \sim^{\ell_{i-1}-1} \mathfrak{M}, w_u$ and to extend \mathfrak{N}_{i-1} , Q'_{i-1} , \mathcal{T}' and $\hat{\delta}'$ accordingly:

- for every equivalence class $[w_u]_\beta \in Q_i$, we add the class $[v_u]_\beta$ to Q'_{i-1} ,

- extend δ'_{i-1} by a bag $\delta'_i(u)$ that contains these added vertices,
- extend $\hat{\delta}'_{i-1}$ by $\hat{\delta}'_i: u \mapsto v_u$.

However, v_u cannot be chosen arbitrarily (of the right bisimulation type) in the α -class of v_λ for $\alpha = \text{agt}(w_\lambda, w_u)$. Bad choices could violate the m_i -closure condition on Q'_i and lead to a substructure \mathfrak{N}_i that is not isomorphic to \mathfrak{M}_i . Instead the following need to be guaranteed:

- $\text{agt}(v_\lambda, v_u) = \text{agt}(w_\lambda, w_u)$ and $\text{agt}(v_s, v_u) = \text{agt}(w_s, w_u)$ for all $s \in T_{i-1}$,
- Q'_i must be m_i -closed just as Q_i is.

Any violation of the first condition would immediately spoil the isomorphism condition on σ_i for $\text{agt}(v_s, v_u)$ -edges. Violation of the m_i -closure condition, a requirement the invariant by (I2), could be exploited by player **I** in the continuation of the game, if short paths in one configuration cannot be matched in the other.

To make this problem more explicit, consider a choice of v_u such that $\text{agt}(v_s, v_u) = \text{agt}(w_s, w_u)$ for all $s \in T_{i-1}$. Since \mathcal{T}_i is a tree decomposition and u is a child of λ , the bag $\delta_i(u)$ intersects bags of \mathcal{T}_{i-1} only within $\delta_i(\lambda)$, i.e. $\delta_i(s) \cap \delta_i(u) \subseteq \delta_i(\lambda) \cap \delta_i(u)$, for all $s \in T_{i-1}$. Together with Q_i being 2-closed this implies $\llbracket w_s \rrbracket \cap \llbracket w_u \rrbracket \subseteq \llbracket w_\lambda \rrbracket \cap \llbracket w_u \rrbracket$ which in turn implies $\llbracket v_s \rrbracket \cap \llbracket v_u \rrbracket \subseteq \llbracket v_\lambda \rrbracket \cap \llbracket v_u \rrbracket$, for all $s \in T_{i-1}$. We next need to add a vertex $[v_u]_\beta$ to Q'_i if and only if $\llbracket w_u \rrbracket_\beta \in Q_i$. This, however, might result in a set that is *not* m_i -closed. Since Q_i is m_i -closed, there are no short paths of length up to m_i from Q_{i-1} to $\delta_i(u) \setminus \delta_i(\lambda)$ that leave Q_i . And all such paths need to pass through $\delta_i(u) \cap \delta_i(\lambda)$ since \mathcal{T}_i is a tree decomposition. Hence, for $t = \delta_i(u) \cap \delta_i(\lambda)$, we have $d_t(\delta_i(u), Q_{i-1}) > m_i$ on the side of \mathfrak{M} and $d(\mathfrak{M})$, which must be matched on the side of \mathfrak{N} and $d(\mathfrak{N})$.

The key to overcoming these problems lies in freeness (cf. Definition 3.17). Since we assumed \mathfrak{M} and \mathfrak{N} to be sufficiently acyclic and rich, both structures are sufficiently free by Theorem 3.28. Let v' be some world in $[v_\lambda]_\alpha$ that is ℓ_i -bisimilar to w_u , and define $\mathbf{z} := \text{image}(\hat{\delta}'_{i-1}) = \{v_s : s \in T_{i-1}\}$. Then freeness of \mathfrak{N} implies that there is some $v_u \sim v'$ such that

- $\text{agt}(v_\lambda, v_u) = \alpha = \text{agt}(w_\lambda, w_u)$, and
- $(\mathbf{z}, v_\lambda) \perp_{m_i} v_u$, i.e. (\mathbf{z}, v_λ) and v_u are m_i -free.

This world v_u is a suitable choice for the extension of \mathfrak{N}_{i-1} towards \mathfrak{N}_i , with the corresponding extension of Q'_{i-1} towards Q'_i . Putting

$$\begin{aligned} \sigma_i(w_u) &:= \hat{\delta}'_i(u) := v_u, \\ \delta'_i(u) &:= \{[v_u]_\beta : [w_u]_\beta \in \delta_i(u)\}, \end{aligned}$$

we have, for $t = \llbracket v_\lambda \rrbracket \cap \llbracket v_u \rrbracket$, and since $(\mathbf{z}, v_\lambda) \perp_{m_i} v_u$, that

$$d_t(\llbracket v_u \rrbracket, \bigcup \llbracket \mathbf{z} \rrbracket) > m_i \text{ and } d_t(\delta'_i(u), \bigcup \text{image}(\delta'_{i-1})) > m_i,$$

because $\delta'_i(u) \subseteq \llbracket v_u \rrbracket$ and $\bigcup \text{image}(\delta'_{i-1}) \subseteq \bigcup \llbracket \mathbf{z} \rrbracket$. This implies that

$$\bigcup \text{image}(\delta'_{i-1}) \cup \delta'_i(u)$$

is m_i -closed. Furthermore,

$$\begin{aligned}
& d_i(\llbracket v_u \rrbracket, \bigcup \llbracket \mathbf{z} \rrbracket) > 1 \\
\Rightarrow & \llbracket z \rrbracket \cap \llbracket v_u \rrbracket \subseteq \llbracket v_\lambda \rrbracket \cap \llbracket v_u \rrbracket, \text{ for all } z \in \mathbf{z}, \\
\Rightarrow & \text{agt}(z, v_u) \supseteq \text{agt}(v_\lambda, v_u), \text{ for all } z \in \mathbf{z}.
\end{aligned}$$

Together with $\text{agt}(v_\lambda, v_s) = \text{agt}(w_\lambda, w_s)$, for all $s \in T_{i-1}$, and $\text{agt}(v_\lambda, v_u) = \text{agt}(w_\lambda, w_u)$ we obtain $\text{agt}(v_u, v_s) = \text{agt}(w_u, w_s)$, for all $s \in T_{i-1}$. This means that adding $\delta'(u) = \{[v_u]_\beta : [w_u]_\beta \in \delta_i(u)\}$ to Q'_{i-1} as a first step in defining Q'_i and defining $\sigma_i(w_u) := v_u = \hat{\delta}'_i(u)$ conforms to the isomorphism conditions between $d(\mathfrak{M}) \upharpoonright Q_i$ and $d(\mathfrak{N}) \upharpoonright Q'_i$ and for σ_i as defined so far. Thus, we have shown that we can find a counterpart for w_u in \mathfrak{N} such that the invariant is maintained.

The remainder of the tree T is treated in the same way in a breadth-first manner. As new nodes $\delta'(u) \in d(\mathfrak{N})$ for $u \in T$ already processed are added to Q'_i , the associated new vertices $\hat{\delta}'_i(u) := v_u$ are added into \mathfrak{N}_i and incorporated into \mathbf{z} for the next extension treating a new node $u \in T$. The distinguished world z_0 of the pointed set (\mathbf{z}, z_0) is the world that is associated with the father of the node in T that is to be processed. The freeness argument works for the whole tree T , node by node, because the size of Q and the depth of \mathcal{T} are a priori bounded, and \mathfrak{M} and \mathfrak{N} could be guaranteed to be sufficiently free for this argument through all q rounds. In particular, the bound $f_{m_i}(|\tau| + 1)$ on the size of Q translates into a trivial bound on the depth of T so that $\mathfrak{M}, w_\lambda \sim^{\ell_{i-1}} \mathfrak{N}, v_\lambda$ and $\ell_{i-1} = \ell_i + f_{m_i}(|\tau| + 1)$, guarantee that $\mathfrak{M}, w \sim^{\ell_i} \mathfrak{N}, \sigma_i(w)$ holds for all $w \in \mathfrak{M}_i$.

The completion of this construction for round i in particular yields the actual response for player **II**, placement of the pebble on $v_i \in \mathfrak{N}_i \subseteq \mathfrak{N}$ in response to player **I**'s pebble placement on $w_i \in \mathfrak{M}_i \subseteq \mathfrak{M}$. Thus, player **II** is able to maintain the invariant in the i -th round. The following lemma summarises this.

Lemma 4.1. *Let $q \in \mathbb{N}$, and \mathfrak{M}, w_0 and \mathfrak{N}, v_0 be pointed Cayley structures that are sufficiently acyclic and sufficiently rich. If the invariant described in Section 4.1 is true in the $(i - 1)$ -th round of the q -round Ehrenfeucht-Fraïssé game on \mathfrak{M}, w_0 and \mathfrak{N}, v_0 , then player **II** has a strategy to update and maintain this invariant in the i -th round.*

4.2 Upgrading and characterisation

This section can be regarded as the culmination of the work so far: the upgrading theorem and the characterisation of basic modal logic over (finite) Cayley structures. The elements of structural analysis developed so far contain all the building blocks for proving those two theorems. We speak of *two* theorems because the restriction to finite models and the unconstrained classical reading are a priori independent. Proving either one does not entail the other even though our specific proof method allows us to treat the two versions in parallel. In Sections 2.4 and 3.1 we showed that every (finite) CK structure can be covered by a

bisimilar (finite) Cayley structure that is arbitrarily acyclic and arbitrarily rich. Recall that the (finite) bisimilar coverings by Cayley structures from Lemma 2.7 were boosted to (finite) bisimilar coverings by Cayley structures satisfying additional acyclicity and richness requirements in Lemma 3.4. The main result of Section 3.2, the freeness theorem, then further showed that sufficient degrees of acyclicity and richness imply (m, k) -freeness, a special property of suitable Cayley structures that is essential for the upgrading. In particular we see that not just Cayley structures but even Cayley structures of any given finite degree of acyclicity, richness and freeness can, up to bisimulation, be taken as universal representatives of all (finite) CK structures.

Finally, the previous development in the current section has provided an invariant that affords player **II** a winning strategy in the q -round Ehrenfeucht-Fraïssé game on sufficiently free pointed Cayley structures that are ℓ -bisimilar for some sufficiently large ℓ . The upgrading theorem follows easily from that.

Theorem 4.2 (upgrading theorem). *Let $q \in \mathbb{N}$. For some suitable choice of $\ell = \ell(q)$, any sufficiently acyclic and sufficiently rich Cayley structures \mathfrak{M} and \mathfrak{N} satisfy*

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

Proof. Let $(\ell_k)_{0 \leq k \leq q}$ be the sequence of the same name from Section 4.1. Set $\ell := \ell_0$ and let $\mathfrak{M}, w, \mathfrak{N}, v$ be two sufficiently acyclic and rich pointed Cayley structures such that $\mathfrak{M}, w \sim^\ell \mathfrak{N}, v$. In order to prove that these structures are FO_q -equivalent we provide a winning strategy for player **II** in the q -round Ehrenfeucht-Fraïssé game on \mathfrak{M}, w and \mathfrak{N}, v . Her strategy is to preserve the invariant from Section 4.1 according to Lemma 4.1.

We need to check that the invariant can be set up before the first round. That $\mathfrak{M}, w \sim^\ell \mathfrak{N}, v$ implies that the substructures $\mathfrak{M}_0 := \mathfrak{M} \upharpoonright \{w\}$ and $\mathfrak{N}_0 := \mathfrak{N} \upharpoonright \{v\}$ are isomorphic via $\sigma_0 = \{w \mapsto v\}$ (w and v are atomically equivalent and all accessibility relations are reflexive), and that σ_0 respects \sim^{ℓ_0} . The singleton subsets $Q_0 := \{[w]_\emptyset\}$ in $d(\mathfrak{M})$ and $Q'_0 := \{[v]_\emptyset\}$ in $d(\mathfrak{N})$ are m_0 -closed because the Cayley structures, and with them their dual hypergraphs by Lemma 3.13, are sufficiently acyclic. The induced sub-hypergraphs $d(\mathfrak{M}) \upharpoonright Q_0$ and $d(\mathfrak{N}) \upharpoonright Q'_0$ are trivially isomorphic and tree decomposable. By Lemma 4.1 player **II** can therefore maintain the invariant through all q rounds, which implies that she wins the game. In the end, no matter what the moves of player **I**, the pebble placements are related by the isomorphism σ_q between induced substructures \mathfrak{M}_q and \mathfrak{N}_q . \square

The upgrading theorem, together with the existence of suitable bisimilar coverings, implies our main result, the characterisation of ML as the bisimulation-invariant fragment of FO over the non-elementary classes of all CK-structures and of all finite CK-structures, respectively.

Theorem 4.3 (main theorem). *Over the class of (finite) Cayley structures, and hence over the class of (finite) CK-structures:*

$$\text{ML}[\text{CK}] \equiv \text{ML} \equiv \text{FO}/\sim$$

$$\begin{array}{ccc}
\mathfrak{M}, w & \xrightarrow{\sim^\ell} & \mathfrak{N}, v \\
\downarrow \sim & & \downarrow \sim \\
\hat{\mathfrak{M}}, \hat{w} & \xrightarrow{\equiv_q} & \hat{\mathfrak{N}}, \hat{v}
\end{array}$$

Figure 4: Upgrading \sim^ℓ to \equiv_q .

Proof. The standard translation (cf. Section 1) implies $\text{ML} \subseteq \text{FO}/\sim$. For the crucial converse direction, establishing *expressive completeness*, let φ be an FO-formula with $\text{qr}(\varphi) = q$ that is bisimulation-invariant over (finite) Cayley structures. If we can show that φ is \sim^ℓ -invariant for some $\ell \in \mathbb{N}$ over (finite) Cayley structures, then there is an ML formula of modal depth ℓ that is logically equivalent to φ over (finite) Cayley structures (cf. Theorem 2.5).

We choose $\ell = \ell(q)$ from Theorem 4.2 above, and let \mathfrak{M}, w and \mathfrak{N}, v be pointed Cayley structures that are ℓ -bisimilar (compare Figure 4). By Lemma 2.11 (for the unrestricted reading) and Lemma 3.4 (to cover the restriction to finite structures) there are bisimilar coverings $\hat{\mathfrak{M}}, \hat{w} \sim \mathfrak{M}, w$ and $\hat{\mathfrak{N}}, \hat{v} \sim \mathfrak{N}, v$ that are sufficiently acyclic and rich such that Theorem 4.2 applies. In particular, Lemma 3.4 gives us such coverings that are finite if \mathfrak{M} and \mathfrak{N} are finite. Since in particular $\hat{\mathfrak{M}}, \hat{w} \sim^\ell \hat{\mathfrak{N}}, \hat{v}$, Theorem 4.2 implies $\hat{\mathfrak{M}}, \hat{w} \equiv_q \hat{\mathfrak{N}}, \hat{v}$, hence

$$\begin{aligned}
\mathfrak{M}, w \models \varphi & \Leftrightarrow \hat{\mathfrak{M}}, \hat{w} \models \varphi && (\varphi \sim\text{-inv.}) \\
& \Leftrightarrow \hat{\mathfrak{N}}, \hat{v} \models \varphi && (\text{qr}(\varphi) \leq q) \\
& \Leftrightarrow \mathfrak{N}, v \models \varphi && (\varphi \sim\text{-inv.})
\end{aligned}$$

which implies \sim^ℓ -invariance of φ over (finite) Cayley structures, as desired. \square

Acknowledgement. We are particularly indebted to an anonymous referee whose insightful and constructive report contained several corrections and suggestions that have led to essential improvements in the present account.

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] N. Biggs. Cubic graphs with large girth. In G. Blum et al., editor, *Annals of the New York Academy of Sciences*, volume 555, pages 56–62. 1989.
- [4] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, 2001.

- [5] F. Canavoi. *Cayley structures and the expressiveness of common knowledge logic*. PhD thesis, Department of Mathematics, TU Darmstadt, 108 pages, 2018.
- [6] F. Canavoi. Cayley structures and coset acyclicity. *arXiv e-prints*, arXiv:1907.07025, July 2019.
- [7] A. Cayley. Desiderata and suggestions: no. 1. the theory of groups. *American Journal of Mathematics*, 1, 1878. DOI: 10.2307/2369433.
- [8] A. Cayley. Desiderata and suggestions: no. 2. the theory of groups. *American Journal of Mathematics*, 1, 1878. DOI: 10.2307/2369306.
- [9] A. Dawar and M. Otto. Modal characterisation theorems over special classes of frames. *Annals of Pure and Applied Logic*, 161:1–42, 2009.
- [10] H.-D. Ebbinghaus and J. Flum. *Finite Model Theory*. Springer, 2nd edition, 1999.
- [11] R. Fagin, J. Halpern, Y. Moses, and M. Vardi. *Reasoning About Knowledge*. MIT Press, 1995.
- [12] 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.
- [13] W. Hodges. *Model Theory*. Cambridge University Press, 1993.
- [14] D. Janin and I. Walukiewicz. On the expressive completeness of the propositional μ -calculus with respect to monadic second-order logic. In *Proceedings of CONCUR*, pages 263–277, 1996.
- [15] M. Otto. Elementary proof of the van Benthem–Rosen characterisation theorem. Technical Report 2342, Fachbereich Mathematik, Technische Universität Darmstadt, 2004.
- [16] M. Otto. Modal and guarded characterisation theorems over finite transition systems. *Annals of Pure and Applied Logic*, 130:173–205, 2004.
- [17] M. Otto. Highly acyclic groups, hypergraph covers and the guarded fragment. *Journal of the ACM*, 59 (1), 2012.
- [18] E. Rosen. Modal logic over finite structures. *Journal of Logic, Language and Information*, 6:427–439, 1997.
- [19] J. van Benthem. *Modal Logic and Classical Logic*. Bibliopolis, Napoli, 1983.

Appendix

Lemma 3.27. *Let $m \in \mathbb{N}$, \mathfrak{M} be a Cayley structure, v a world, (\mathbf{z}, z_0) a finite pointed set, $\mathbf{y} \subseteq \mathbf{z}$ a possibly empty subset and $w \in \mathbf{z} \setminus \mathbf{y}$; set $\gamma = \text{agt}(z_0, v)$ and $t = \rho(v, \gamma)$. Assume that \mathfrak{M} is sufficiently acyclic and sufficiently rich, and*

- $\gamma \subseteq \text{agt}(z, v)$, for all $z \in \mathbf{z}$;
- $d_t(\mathbf{y}, v) > m$.

Then there is a world $v^ \in [v]_\gamma$ with $\mathfrak{M}, v^* \sim \mathfrak{M}, v$ and $\text{agt}(z, v^*) = \text{agt}(z, v)$, for all $z \in \mathbf{z}$, such that*

$$d_t(\mathbf{y} \cup \{w\}, v^*) > m.$$

Proof. If $d_t(w, v) > m$, simply set $v^* = v$. Otherwise there is some $1 < \ell \leq m$ such that $d_t(w, v) \leq \ell$. Since \mathfrak{M} is sufficiently acyclic there is an inner coset path of length ℓ from w to v that avoids t but no such path of length $< \ell$. We need to show that there is a suitable world v^* such that there is no inner coset path of length up to ℓ from w to v^* that avoids t . Then the statement follows from repeated application of the same argument.

Proof outline. We inductively find a sequence of worlds $(v_n)_{n \geq 1}$ in $[v]_\gamma$ that are bisimilar to v , along with three auxiliary sequences: two sequences of sets of agents $(\beta_n)_{n \geq 1}$, $(\gamma_n)_{n \geq 1}$ and a sequence of agents $(a_n)_{n \geq 1}$ in $\text{agt}(z_0, v)$.

We show that these sequences terminate after finitely many steps and that the last one of the v_n can serve as the desired world v^* . Intuitively, every v_n will be, in some sense, further away from v than its predecessor v_{n-1} ; β_n describes the direction back to w on short paths that avoid t ; γ_n the steps that still have to be taken to get far enough away from w ; and a_n is the direction we take to go from v_{n-1} to v_n .

To construct the sequences we need one more auxiliary statement that says that, as long as there is a short path from v to w , we can move in a suitable direction to a copy $v' \sim v$ without *decreasing* the distance to \mathbf{y} , i.e. that we can move away from several worlds simultaneously. It is similar in spirit to Lemma 3.24.

Claim 1. Let \mathfrak{M} be a Cayley structure, v a world, (\mathbf{z}, z_0) a finite pointed set, $\mathbf{y} \subseteq \mathbf{z}$ a possibly empty subset and $w \in \mathbf{z} \setminus \mathbf{y}$; set $\gamma = \text{agt}(z_0, v)$, $t = \rho(v, \gamma)$ and $m \geq 2$. Assume that \mathfrak{M} is sufficiently acyclic and sufficiently rich, and

- $\gamma \subseteq \text{agt}(z, v)$, for all $z \in \mathbf{z}$;
- $d_t(\mathbf{y}, v) > m$ and $d_t(w, v) = \ell \leq m$.

Then $\gamma \setminus \text{short}_t(v, w) \neq \emptyset$, and for every $a \in \gamma \setminus \text{short}_t(v, w)$ there is some $v' \in [v]_a \setminus \{v\}$ such that $\mathfrak{M}, v' \sim \mathfrak{M}, v$, and

$$d_t(\mathbf{y}, v') > m.$$

Proof of claim 1. Let $w, \alpha_1, \dots, \alpha_\ell, v$ be an inner non- t coset path. In particular, this means $\gamma \not\subseteq \alpha_\ell$ and also $\gamma \not\subseteq \text{short}_t(v, w)$ since $\emptyset \neq \text{short}_t(v, w) \subseteq \alpha_\ell$. Thus, we obtain the first statement: $\gamma \setminus \text{short}_t(v, w) \neq \emptyset$.

For the second statement, let $a \in \gamma \setminus \text{short}_t(v, w)$, $z \in \mathbf{y}$ and assume there is some $u \in [v]_a$ and $k \leq m$ such that there is a non- t coset path $z, \beta_1, \dots, \beta_k, u$. We claim that $d_t(z, u') > m$ for $u' \in [v]_a \setminus \{u\}$. Firstly, we show $a \notin \text{short}_t(u, z)$. Assume $a \in \text{short}_t(u, z)$, then $a \in \beta_k$ because $u, \beta_k, \dots, \beta_1, z$ is a short non- t coset path. It follows $v \in [u]_a \subseteq [u]_{\beta_k} = [v]_{\beta_k}$ which means that $z, \beta_1, \dots, \beta_k, v$ or $z, \beta_1, \dots, \beta_{k-1}, v$ is a short non- t coset path. This implies $d_t(z, v) \leq m$ contrary to $d_t(\mathbf{y}, v) > m$, so $a \notin \text{short}_t(u, z)$. Secondly, $d_t(z, u') \leq m$ implies $a \in \text{short}_t(u', z)$ (Lemma 3.26), which again implies $d_t(z, v) \leq m$, contrary to assumption.

Thus, for any $z \in \mathbf{y}$ there is at most one $u_z \in [v]_a$ such that $d_t(z, u_z) \leq m$. Since \mathfrak{M} is sufficiently rich, there remains a world $v' \in [v]_a \setminus \{v\}$ such that $v' \sim v$ and $d_t(\mathbf{y}, v') > m$.

End of proof of claim 1.

The construction. For $n = 1$, Claim 1 implies $\gamma \setminus \text{short}_t(v, w) \neq \emptyset$; let $a_1 \in \gamma \setminus \text{short}_t(v, w)$. As \mathfrak{M} is sufficiently rich, there is a world $v_1 \in [v]_{a_1} \setminus \{v\}$ that is bisimilar to v such that $\text{agt}(v_1, z) = \text{agt}(v, z)$, for all $z \in \mathbf{z}$, and $d_t(\mathbf{y}, v_1) > m$ (cf. Lemma 3.24 and Claim 1). If $d_t(w, v_1) \leq \ell$, set $\beta_1 := \text{short}_t(v_1, w)$ and $\gamma_1 := (\gamma \setminus \text{short}_t(v, w)) \setminus \beta_1$. If $d_t(w, v_1) > \ell$ or $\gamma_1 = \emptyset$, then the sequence terminates in v_1 .

For $n > 1$, assume that the worlds v_1, \dots, v_{n-1} and the sets $\beta_1, \dots, \beta_{n-1}$, $\gamma_1, \dots, \gamma_{n-1}$ have been defined and that the sets are non-empty. Let $a_n \in \gamma_{n-1}$. Since \mathfrak{M} is sufficiently rich, by Lemma 3.24 and Claim 1 there is again a $v_n \in [v_{n-1}]_{a_n} \setminus \{v_{n-1}\}$, bisimilar to v_{n-1} such that $\text{agt}(v_n, z) = \text{agt}(v_{n-1}, z)$, for all $z \in \mathbf{z}$, and $d_t(\mathbf{y}, v_n) > m$. If $d_t(w, v_n) \leq \ell$, set $\beta_n := \text{short}_t(v_n, w)$, and $\gamma_n := \gamma_{n-1} \setminus \beta_n$. If $d_t(w, v_n) > \ell$ or $\gamma_n = \emptyset$, then v_n is the last world and the sequence terminates in v_n .

We have constructed these four finite sequences:

$$\begin{aligned} (v_n)_{n \geq 1} &\in [v]_\gamma \subseteq [v]_{\text{agt}(w, v)} \quad \text{all bisimilar to } v; \\ (a_n)_{n \geq 1} &\in \Gamma; \\ (\beta_n)_{n \geq 1} &\in \tau; \\ (\gamma_n)_{n \geq 1} &\in \tau. \end{aligned}$$

Additionally set $v_0 := v$, $\beta_0 := \text{short}_t(v, w)$ and $\gamma_0 := \gamma \setminus \beta_0$.

Correctness. We show the following properties of the sequences by induction on $n \geq 1$.

- (1) $\beta_n = \{a_j, a_{j+1}, \dots, a_n\}$, for some $1 \leq j \leq n$, or $\beta_n \supseteq \beta_0 \cup \{a_1, \dots, a_n\}$.
- (2) The worlds v_0, \dots, v_n occur on every short inner coset path that avoids t from w to v_n in the order of their indices: let $w_1, \alpha_1, w_2, \dots, w_k, \alpha_k, w_{k+1}$ be such a path, and $0 \leq i < j \leq n$. If $1 \leq k_i, k_j \leq m$ are minimal such that $v_i \in [w_{k_i}]_{\alpha_{k_i}}$ and $v_j \in [w_{k_j}]_{\alpha_{k_j}}$, then $k_i \leq k_j$.
- (3) $\gamma_n \subsetneq \gamma_{n-1}$.

For $n = 1$, ad (1) and (2). Together with $a_1 \in \gamma \setminus \text{short}_t(v, w)$ and $v_1 \neq v_0$, Lemma 3.26 implies $a_1 \in \beta_1 = \text{short}_t(v_1, w)$. For every inner short coset path

$$w = w_1, \alpha_1, w_2, \dots, w_k, \alpha_k, w_{k+1} = v_1$$

from w to $v_1 \in [v_0]_{\alpha_1}$ that avoids t we have $v_0 \in [v_1]_{\alpha_k}$ because $\alpha_k \supseteq \beta_1 \ni a_1$. Furthermore, since k is the minimal index such that $v_1 \in [w_k]_{\alpha_k}$, the minimal index for v_0 can only be smaller or equal. If there is one such path with $v_0 \in [w_k]_{\alpha_{k-1} \cap \alpha_k}$, we have $\beta_1 = \{a_1\}$, because

$$w = w_1, \alpha_1, w_2, \dots, v_0, \{a_1\}, w_{k+1} = v_1$$

would be a short inner coset path from w to v_1 . If $v_0 \in [w_k]_{\alpha_k} \setminus [w_{k-1}]_{\alpha_{k-1}}$, for all short inner coset paths from w to v_1 , then $\beta_0 = \text{short}_t(v_0, w) \subseteq \alpha_k$ since every such path is a short inner coset path from w to v_0 that avoids t . Thus, $\beta_1 = \{a_1\}$ or $\beta_0 \cup \{a_1\} \subseteq \beta_1$.

For $n = 1$, ad (3), note that $\gamma_1 = (\gamma_0 \setminus \beta_0) \setminus \beta_1 = \gamma_0 \setminus (\beta_0 \cup \beta_1)$ implies $\gamma_1 \subseteq \gamma_0$, which together with $a_1 \in \gamma_0 \cap \beta_1$ implies $\gamma_1 \subsetneq \gamma_0$.

For $n > 1$ inductively assume that properties (1)–(3) hold for $1, \dots, n-1$.

For n , ad (1): we chose

$$a_n \in \gamma_{n-1} = \gamma_{n-2} \setminus \beta_{n-1} = \gamma_{n-2} \setminus \text{short}(v_{n-1}, w) \quad \text{and} \quad v_n \in [v_{n-1}]_{a_n} \setminus \{v_{n-1}\}.$$

Lemma 3.26 implies that $a_n \in \beta_n = \text{short}_t(v_n, w)$. If $\{a_1, \dots, a_n\} \subseteq \beta_n$, then $\beta_n = \{a_1, \dots, a_n\}$ or $\beta_0 \cup \{a_1, \dots, a_n\} \subseteq \beta_n$, similar to the base case. If there is a $1 \leq j < n$ such that $a_j \notin \beta_n$, let j be the largest such index. Thus, there is a short inner coset path

$$w = w_1, \alpha_1, w_2, \dots, v_j, \{a_{j+1}, \dots, a_n\}, v_n$$

from w to v_n that avoids t , which implies $\beta_n = \{a_{j+1}, \dots, a_n\}$.

For n , ad (2): let $w_1, \alpha_1, w_2, \dots, w_k, \alpha_k, w_{k+1}$ be a short inner coset path from w to v_n that avoids t . We showed $a_n \in \beta_n \subseteq \alpha_k$ which implies $v_{n-1} \in [w_{k+1}]_{\alpha_k}$. So

$$w = w_1, \alpha_1, w_2, \dots, w_k, \alpha_k, v_{n-1} \quad \text{or} \quad w = w_1, \alpha_1, w_2, \dots, w_{k-1}, \alpha_{k-1}, v_{n-1}$$

is a short inner coset path from w to v_{n-1} that avoids t . By induction hypothesis the worlds v_0, \dots, v_{n-1} must occur on such a path in order of their indices. The smallest index i such that $v_n \in [w_i]_{\alpha_i}$ is k . Thus, all worlds v_0, \dots, v_{n-1} occur in equivalence classes $[w_i]_{\alpha_i}$ with $i \leq k$.

For n , property (3) follows from $\gamma_n = \gamma_{n-1} \setminus \beta_n$ and the fact that $a_n \in \gamma_{n-1} \cap \beta_n$.

First of all, property (3) implies that the four sequences as constructed terminate after finitely many steps, since there are only finitely many agents. If v_k is the terminal world in the first sequence, we claim that $d_t(w, v_k) > \ell$:

There cannot be an inner coset path that avoids t of length $< \ell$ from w to v_k because that would imply an inner coset path from w to v that avoids t of

length $< \ell$ by property (2), which cannot exist by assumption. Hence, for the sake of contradiction, we assume that there is an inner coset path of length ℓ

$$w = w_1, \alpha_1, w_2, \dots, w_\ell, \alpha_\ell, w_{\ell+1} = v_k$$

from $w = w_1$ to $v_k = w_{\ell+1}$ that avoids t . Again, property (2) implies that v occurs somewhere on this path. Furthermore, the smallest index i such that $v \in [w_i]_{\alpha_i}$ must be ℓ , otherwise there would be an inner coset path from w to v that avoids t of length $< \ell$. In particular, $v = v_0 \in [w_\ell]_{\alpha_\ell} \setminus [w_{\ell-1}]_{\alpha_{\ell-1}}$. Property (2) states that all worlds v_1, \dots, v_k must occur after v_0 on all short inner coset paths from w to v_k that avoid t , hence $v_i \in [w_\ell]_{\alpha_\ell} \setminus [w_{\ell-1}]_{\alpha_{\ell-1}}$, for all $1 \leq i \leq k$. This implies $\bigcup_{i=0}^k \beta_i \subseteq \alpha_\ell$ because $\beta_i = \text{short}_t(v_i, w)$, for all $0 \leq i \leq k$. Furthermore,

$$\emptyset = \gamma_k = \gamma \setminus \bigcup_{i=0}^k \beta_i \quad \Rightarrow \quad \gamma \subseteq \bigcup_{i=0}^k \beta_i \subseteq \alpha_\ell.$$

But we also have $\gamma \not\subseteq \alpha_\ell$ because we assumed that the coset path

$$w = w_1, \alpha_1, \dots, \alpha_\ell, w_{\ell+1} = v_k$$

avoids $t = \rho(v, \gamma) = \rho(v_k, \gamma)$, contradicting the assumption $d_t(w, v_k) \leq \ell$.

Thus, since each agent a_i , $1 \leq i \leq k$, is an element of γ and each v_i , $1 \leq i \leq k$, was chosen such that $\text{agt}(z, v_i) = \text{agt}(z, v)$, for all $z \in \mathbf{z}$, $d_t(\mathbf{y}, v_i) > m$ and $\mathfrak{M}, v \sim \mathfrak{M}, v_i$, the world $v_k =: v^*$ is as desired. \square