

Symmetries in Modal Logics

Carlos Areces

FaMAF
Universidad Nacional de Córdoba
Córdoba, Argentina
CONICET, Argentina

Guillaume Hoffmann

FaMAF
Universidad Nacional de Córdoba
Córdoba, Argentina

Ezequiel Orbe

FaMAF
Universidad Nacional de Córdoba
Córdoba, Argentina
CONICET, Argentina

We generalize the notion of symmetries of propositional formulas in conjunctive normal form to modal formulas. Our framework uses the coinductive models introduced in [4] and, hence, the results apply to a wide class of modal logics including, for example, hybrid logics. Our main result shows that the symmetries of a modal formula preserve entailment: if σ is a symmetry of φ then $\varphi \models \psi$ if and only if $\varphi \models \sigma(\psi)$.

1 Symmetries in Automated Theorem Proving

Many concrete, real life problems present symmetries. For instance, if we want to know whether trying to place three pigeons in two pigeonholes results in two occupying the same nest, it does not really matter which of all pigeons gets in each pigeonhole. Starting by putting the first pigeon to the first pigeonhole is the same as if we put the second one in it. In mathematical and common-sense reasoning these kinds of symmetries are often used to reduce the difficulty of reasoning — one can analyze in detail only one of the symmetric cases and generalize the result to the others. The exact same is done in propositional theorem proving. Many problem classes and, in particular, those arising from real world applications, display a large number of symmetries; and current SAT solvers take into account these symmetries to avoid exploring duplicate branches of the search space. In the last years there has been extensive research in this area, focusing on how to define symmetries, how to detect them efficiently, and how SAT solvers can better profit from them [27].

Informally, we can define a symmetry of a discrete object as a permutation of its components that leaves the object, or some aspect of it, intact (think of the rotations of a spatial solid). In the context of SAT solving we can formally define a symmetry as a permutation of the variables (or literals) of a problem that preserves its structure and, in particular, its set of solutions. Depending on which aspect of the problem is kept invariant, symmetries are classified in the literature into semantic or syntactic [9]. Semantic symmetries are intrinsic properties of a Boolean function that are independent of any particular representation, i.e., a permutation of variables that does not change the value of the function under any variable assignment. Syntactic symmetries, on the other hand, correspond to the specific algebraic representation of the function, i.e., a permutation of variables or literals that does not change the representation. A syntactic symmetry is also a semantic symmetry, but the converse does not always hold.

In [23], Krishnamurthy used symmetries in the context of SAT solving. In this article, the notions of *global* and *local* symmetries as inference rules are used to strengthen resolution-based proof systems for propositional logic, showing that they can shorten the proofs of certain difficult propositional problems like the pigeonhole principle. Since then, many articles discuss how to detect and exploit symmetries. Most of them can be grouped into two different approaches: static symmetry breaking and dynamic symmetry breaking. In the first approach [14, 15, 1], symmetries are detected and eliminated from the

problem statement before the SAT solver is used. They work as a preprocessing step. In contrast, dynamic symmetry breaking [13, 9, 10] detects and breaks symmetries during the search space exploration. The first approach can be used with any theorem prover; the second is prover dependent but it can take advantage of symmetries that emerge during search. Despite their differences they share the same goal: to identify symmetric branches of the search space and guide the SAT solver away from symmetric branches already explored. A third alternative was introduced in [8], which combines symmetry reasoning with clause learning [28] in Conflict-Driven Clause Learning SAT solvers [18]. The idea is to augment clause learning by using the symmetries of the problem to learn the symmetric equivalents of conflict-induced clauses. This approach is particularly appealing as it does not imply major modifications to the search procedure and the required modifications to the clause learning process are minor.

Symmetries have been extensively investigated and successfully exploited for propositional logic SAT and some results involve other logics, see [6, 7, 17]. To the best of our knowledge, symmetries remain largely unexplored in automated theorem proving for modal logics.

In this paper, we generalize the notion of symmetries to modal formulas in conjunctive normal form for different modal logics including the basic modal language over different model classes (e.g., reflexive, linear or transitive models), and logics with additional modal operators (e.g., universal and hybrid operators). The main result of the article shows that symmetries of a modal formula preserve entailment: if σ is a symmetry of φ then $\varphi \models \psi$ if and only if $\varphi \models \sigma(\psi)$. In cases where the modal language has a tree model property, we can actually use a more flexible notion of symmetry that enables different permutations to be applied at each modal depth. We also present a method to detect the symmetries of modal formulas in conjunctive normal form. This method reduces the symmetry detection problem to the graph automorphism problem. A general graph construction algorithm, suitable for many modal logics, is presented. In order to tackle a broad range of modal languages that may or may not enjoy the tree model property, we use in our work the semantics provided by coinductive modal models [4] instead of the more familiar Kripke relational semantics. Coinductive modal models provide a homogeneous framework to investigate different modal languages at a greater level of abstraction. A consequence of this is that results obtained in the coinductive framework can be easily extended to concrete modal languages by just giving the appropriate definition of the model classes and fixing some parameters.

In Section 2 we present modal logics and coinductive modal models. In Section 3 we define modal symmetries, together with the appropriate notion of simulation to show that symmetries preserve modal entailment. In Section 4 we introduce layered permutations and show that they can be used when the modal logic has the adequate notion of the tree model property. In Section 5 we present a graph construction algorithm to detect symmetries in modal formulas and prove its correctness. We draw our conclusions and discuss future research in Section 6.

2 Modal Logics and Coinductive Models

In what follows, we will assume basic knowledge of classical modal logics and refer the reader to [11, 12] for technical details. The coinductive framework for modal logics was introduced in [4] to investigate normal forms for a wide number of modal logics. Its main characteristic is that it allows the representation of different modal logics in a homogeneous form. For a start, the set of formulas is as for the basic (multi) modal logic.

Definition 1 (Modal formula). *A modal signature is a pair $\langle Atom, Mod \rangle$ where $Atom$ and Mod are two countable, disjoint sets. We usually assume that $Atom$ is infinite. The set of modal formulas over $\langle Atom,$*

$Mod\}$ is defined as

$$\varphi ::= a \mid \neg\varphi \mid \varphi \vee \varphi \mid [m]\varphi,$$

for $a \in Atom$, $m \in Mod$. \top and \perp stand for an arbitrary tautology and contradiction, respectively. Connectives such as \wedge, \rightarrow and $\langle m \rangle$, are defined as usual.

We will define a symmetry as a permutation of literals that preserve the structure of formulas in conjunctive normal form (CNF).

Definition 2 (Literals and modal CNF). A literal l is either an atom a or its negation $\neg a$. The set of literals over $Atom$ is $ALit = Atom \cup \{\neg a \mid a \in Atom\}$.

A modal formula is in modal conjunctive normal form (modal CNF) if it is a conjunction of modal CNF clauses. A modal CNF clause is a disjunction of atoms and modal literals. A modal literal is a formula of the form $[m]C$ or $\neg[m]C$ where C is a modal CNF clause. Every modal formula can be transformed into an equisatisfiable formula in modal CNF in polynomial time (see [4, 26] for details).

A formula in modal CNF can be represented as a set of modal CNF clauses (interpreted conjunctively), and each clause can be represented as a set of atom and modal literals (interpreted disjunctively). With the set representation we can disregard the order and multiplicity in which clauses and literals appear. This will be important when we define symmetries below. In the rest of the paper we will assume that modal formulas are in modal CNF, and we will refer to them as modal CNF formulas.

Example 1. The modal formula $\varphi = \langle m \rangle(p \wedge q \wedge p) \wedge [m]\neg r$ is equisatisfiable to the modal CNF formula $\varphi' = \{\{\neg[m]\{\neg p, \neg q\}\}, \{[m]\{\neg r\}\}\}$.

Up to now, we have not departed from the standard presentation of classical modal logic in important ways. The main change introduced by the coinductive approach is with the definition of model and semantic conditions.

Definition 3 (Models). Let $\mathcal{S} = \langle Atom, Mod \rangle$ be a modal signature and W be a fixed, non-empty set. \mathbf{Mods}_W , the class of all models with domain W , for the signature \mathcal{S} , is the class of all tuples $\langle w, W, V, R \rangle$ such that $w \in W$, $V(v) \subseteq Atom$ for all $v \in W$, and

$$R(m, v) \subseteq \mathbf{Mods}_W \text{ for } m \in Mod \text{ and } v \in W.$$

Given a model $\mathcal{M} = \langle w, W, V, R \rangle$ we will say that w is the point of evaluation and denote it as $w^{\mathcal{M}}$, W is the domain and denote it as $|\mathcal{M}|$, V is the modal valuation and denote it as $V^{\mathcal{M}}$, and R is the accessibility relation and denote it as $R^{\mathcal{M}}$. \mathbf{Mods} denotes the class of all models over all domains, $\mathbf{Mods} = \bigcup_W \mathbf{Mods}_W$.

Given $\mathcal{M} \in \mathbf{Mods}_W$, let $\text{Ext}(\mathcal{M})$, the extension of \mathcal{M} , be the smallest subset of \mathbf{Mods}_W that contains \mathcal{M} and is such that if $\mathcal{N} \in \text{Ext}(\mathcal{M})$, then $R^{\mathcal{N}}(m, v) \subseteq \text{Ext}(\mathcal{M})$ for all $m \in Mod$, $v \in W$.

The definition of a coinductive modal model is similar to the usual definition of a Kripke pointed model. The difference lies in the way the accessibility relation is defined. In particular, for each m and each state w , $R(m, w)$ is defined as the set of (potentially different) models accessible from w through the m modality. Observe that for each W , \mathbf{Mods}_W is well-defined (coinductively), and so does \mathbf{Mods} , the class of all models. Our results will apply not only to \mathbf{Mods} but to many of its subclasses. We will be interested in classes which are *closed under accessibility relations* (closed classes for short): $\mathcal{M} \in \mathcal{C}$ implies $\text{Ext}(\mathcal{M}) \subseteq \mathcal{C}$. In the rest of the paper we will only consider classes of models closed under accessibility relations.




Figure 1: a) A Kripke model. b) Equivalent coinductive model.

Example 2. Consider the pointed Kripke model in Figure 1a, and its equivalent coinductive modal model in Figure 1b. The point of evaluation in each model is circled. The main difference is that the relation of a coinductive model leads to another coinductive model, whereas in a Kripke model the relation leads to another point of the same model.

We are now ready to introduce the definition of the satisfiability relation \models .

Definition 4 (Semantics). Let φ be a formula in modal CNF and $\mathcal{M} = \langle w, W, V, R \rangle$ a model in **Mods**. We define \models for modal CNF formulas, clauses and literals as

$$\begin{aligned}
 \mathcal{M} \models \varphi & \quad \text{iff} \quad \text{for all clauses } C \in \varphi \text{ we have } \mathcal{M} \models C \\
 \mathcal{M} \models C & \quad \text{iff} \quad \text{there is some literal } l \in C \text{ such that } \mathcal{M} \models l \\
 \mathcal{M} \models a & \quad \text{iff} \quad a \in V(w) \text{ for } a \in \text{Atom} \\
 \mathcal{M} \models \neg a & \quad \text{iff} \quad a \notin V(w) \text{ for } a \in \text{Atom} \\
 \mathcal{M} \models [m]C & \quad \text{iff} \quad \mathcal{M}' \models C, \text{ for all } \mathcal{M}' \in R(m, w) \\
 \mathcal{M} \models \neg[m]C & \quad \text{iff} \quad \mathcal{M} \not\models [m]C.
 \end{aligned}$$

For \mathcal{C} a class of models, we write $\mathcal{C} \models \varphi$ whenever $\mathcal{M} \models \varphi$ for every \mathcal{M} in \mathcal{C} , and we say that $\Gamma_{\mathcal{C}} = \{\varphi \mid \mathcal{C} \models \varphi\}$ is the logic defined by \mathcal{C} .

The set of models in \mathcal{C} of a formula φ is the set $\text{Mod}_{\mathcal{C}}(\varphi) = \{\mathcal{M} \mid \mathcal{M} \in \mathcal{C} \text{ and } \mathcal{M} \models \varphi\}$ (when \mathcal{C} is clear from the context we will just write $\text{Mod}(\varphi)$). We say that ψ can be inferred from φ in \mathcal{C} and write $\varphi \models_{\mathcal{C}} \psi$ if $\text{Mod}_{\mathcal{C}}(\varphi) \subseteq \text{Mod}_{\mathcal{C}}(\psi)$.

As shown in [4], the logic Γ_{Mods} (generated by the class of all possible models) coincides with the basic multi-modal logic K. By properly restricting the model class we can capture different modal logics. Let us call a predicate P on models a *defining condition* for a class \mathcal{C} whenever \mathcal{C} is such that $\mathcal{M} \in \mathcal{C}$ if and only if $P(\mathcal{M})$ holds. Consider the signature $\mathcal{S} = \langle \text{Atom}, \text{Mod} \rangle$ where $\text{Atom} = \text{Prop} \cup \text{Nom}$, $\text{Mod} = \text{Rel} \cup \{A\} \cup \{\@_i \mid i \in \text{Nom}\}$; and $\text{Prop} = \{p_1, p_2, \dots\}$, $\text{Nom} = \{n_1, n_2, \dots\}$ and $\text{Rel} = \{r_1, r_2, \dots\}$ are mutually disjoint, countable infinite sets. In what follows, we will usually be interested in sub-languages of the language defined over \mathcal{S} by Definition 1.

Class	Defining condition
\mathcal{C}_m^K	$\mathcal{P}_m^K(\mathcal{M}) \iff R^{\mathcal{M}}(m, w) \subseteq \{\langle v, \mathcal{M} , V^{\mathcal{M}}, R^{\mathcal{M}} \rangle \mid v \in \mathcal{M} \}, m \in \text{Rel}$
\mathcal{C}_A	$\mathcal{P}_A(\mathcal{M}) \iff R^{\mathcal{M}}(A, w) = \{\langle v, \mathcal{M} , V^{\mathcal{M}}, R^{\mathcal{M}} \rangle \mid v \in \mathcal{M} \}$
$\mathcal{C}_{@_i}$	$\mathcal{P}_{@_i}(\mathcal{M}) \iff R^{\mathcal{M}}(@_i, w) = \{\langle v, \mathcal{M} , V^{\mathcal{M}}, R^{\mathcal{M}} \rangle \mid i \in V(v)\}, i \in \text{Nom}$
\mathcal{C}_{Nom}	$\mathcal{P}_{\text{Nom}}(\mathcal{M}) \iff \{w \mid i \in V^{\mathcal{M}}(w)\}$ is a singleton, $\forall i \in \text{Nom}$

Figure 2: Defining conditions for different modal logics

Figure 2 introduces a number of closed model classes by means of their defining conditions. Observe that \mathcal{P}_m^K is true for a model \mathcal{M} if every successor of $w^{\mathcal{M}}$ is identical to \mathcal{M} except perhaps on its point

of evaluation. We call m a *relational modality* when it is interpreted in \mathcal{C}_m^K because over this class they behave as classical relational modalities [4].

We can capture different modal operators, like the ones from hybrid logics [2], by choosing the proper class of models. Predicates \mathcal{P}_A and $\mathcal{P}_{@,i}$, for instance, impose conditions on the point of evaluation of the accessible models restricting the evaluation to the class of models where the relation is, respectively, the total relation ($\forall xy.R(x,y)$) and the ‘point to all i ’ relation ($\forall xy.R(x,y) \leftrightarrow i(y)$). Observe that whenever the atom i is interpreted as a singleton set, the ‘point to all i ’ relation becomes the usual ‘point to i ’ relation ($\forall xy.R(x,y) \leftrightarrow y = i$) of hybrid logics. Finally, predicate \mathcal{P}_{Nom} turns elements of Nom into nominals, i.e., true at a unique element of the domain of the model.

An interesting feature of this setting is that we can express the combination of modalities as the intersection of their respective classes. For example, $\mathcal{C}_{\mathcal{H}(@)}$, the class of models for the hybrid logic $\mathcal{H}(@)$, can be defined as follows:

$$\mathcal{C}_{\mathcal{H}(@)} = \mathcal{C}_{\text{Nom}} \cap \mathcal{C}_{@} \cap \mathcal{C}_{\text{Rel}}, \text{ where} \\ \mathcal{C}_{@} = \bigcap_{i \in \text{Nom}} \mathcal{C}_{@,i}, \text{ and } \mathcal{C}_{\text{Rel}} = \bigcap_{m \in \text{Rel}} \mathcal{C}_m^K.$$

The crucial characteristic of the coinductive approach is that all these different modal operators are captured using the same semantic condition introduced in Definition 4. All the details defining each particular operator are now introduced as properties of the accessibility relation. As a result, a unique notion of bisimulation is sufficient to cover all of them.

Definition 5 (Bisimulations). *Given two models \mathcal{M} and \mathcal{M}' we say that \mathcal{M} and \mathcal{M}' are bisimilar (notation $\mathcal{M} \underline{\leftrightarrow} \mathcal{M}'$) if $\mathcal{M} Z \mathcal{M}'$ for some relation $Z \subseteq \text{Ext}(\mathcal{M}) \times \text{Ext}(\mathcal{M}')$ such that whenever $\langle w, W, V, R \rangle Z \langle w', W', V', R' \rangle$ we have the following properties:*

- **Harmony:** $a \in V(w)$ iff $a \in V'(w')$, for all $a \in \text{Atom}$.
- **Zig:** $\mathcal{N} \in R(m, w)$ implies $\mathcal{N} Z \mathcal{N}'$ for some $\mathcal{N}' \in R'(m, w')$.
- **Zag:** $\mathcal{N}' \in R'(m, w')$ implies $\mathcal{N} Z \mathcal{N}'$ for some $\mathcal{N} \in R(m, w)$.

Such Z is called a bisimulation between \mathcal{M} and \mathcal{M}' .

The classic result of invariance of modal formulas under bisimulation [11] can easily be proved.

Theorem 1. *If $\mathcal{M} \underline{\leftrightarrow} \mathcal{M}'$, then $\mathcal{M} \models \varphi$ iff $\mathcal{M}' \models \varphi$, for all φ .*

As stated in [4], this general notion of bisimulation works for every modal logic definable as a closed subclass of **Mods**.

3 Modal Symmetries

We will show that consistent symmetries for modal formulas behave similarly as in the propositional case and, hence, could assist in modal theorem proving. Let us start by introducing the basic notions.

Definition 6 (Complete, Consistent and Generated sets of literals). *A set of literals L is complete if for each $a \in \text{Atom}$ either $a \in L$ or $\neg a \in L$. It is consistent if for each $a \in \text{Atom}$ either $a \notin L$ or $\neg a \notin L$. Any complete and consistent set of literals L defines a unique valuation $v \subseteq \text{Atom}$ as $a \in v$ if $a \in L$ and $a \notin v$ if $\neg a \in L$. For $S \subseteq \text{Atom}$, the consistent and complete set of literals generated by S (notation L_S) is $S \cup \{\neg a \mid a \in \text{Atom} \setminus S\}$.*

Definition 7 (Permutation). *A permutation is a bijective function $\sigma : \text{ALit} \mapsto \text{ALit}$. For L a set of literals, $\sigma(L) = \{\sigma(l) \mid l \in L\}$.*

In this work, we only deal with permutations defined over finite sets of literals, namely, those occurring in the formula φ under consideration. This restricts us to finite groups of symmetries [19]. This kind of permutations can be succinctly defined using cyclic notation, e.g., $\sigma = (p \neg q)(\neg p q)$ is the permutation that makes $\sigma(p) = \neg q$, $\sigma(\neg q) = p$, $\sigma(\neg p) = q$ and $\sigma(q) = \neg p$ and leaves unchanged all other literals; $\sigma = (p q r)(\neg p \neg q \neg r)$ is the permutation $\sigma(p) = q$, $\sigma(q) = r$ and $\sigma(r) = p$ and similarly for the negations. Finally, for $n \in \mathbb{Z}_{\geq 1}$ and σ a permutation, we denote the composition of σ with itself n times by σ^n . σ^0 denote the identity permutation, σ^{-1} the inverse of σ , and σ^{-n} , for $n \in \mathbb{Z}_{\leq -1}$ is the n -times composition of σ^{-1} with itself.

Because, in our language, atoms may occur in some modalities (like $@_i$) we should take some care when we apply permutations to modal formulas. We will say that a modality is *indexed by atoms* if its definition depends on the value of an atom. If m is indexed by an atom a we will sometimes write $m(a)$.

Definition 8. [Permutation of a formula] Let φ be a modal CNF formula and σ a permutation. We define $\sigma(\varphi)$ recursively:

$$\begin{aligned} \sigma(\varphi) &= \{\sigma(C) \mid C \in \varphi\} \quad \text{for } \varphi \text{ a modal CNF formula} \\ \sigma(C) &= \{\sigma(A) \mid A \in C\} \quad \text{for } C \text{ a modal CNF clause} \\ \sigma([m]C) &= [\sigma(m)]\sigma(C) \\ \sigma(\neg[m]C) &= \neg[\sigma(m)]\sigma(C) \end{aligned}$$

where $\sigma(m) = \sigma(m(a)) = m(\sigma(a))$ if m is indexed by a , and $\sigma(m) = m$ otherwise.

Definition 9. A permutation σ is consistent if for every literal l , $\sigma(\neg l) = \neg\sigma(l)$. A permutation σ is a symmetry for φ if $\varphi = \sigma(\varphi)$, when conjunctions and disjunctions in φ are represented as sets.

Example 3. Trivially, the identity permutation $\sigma(l) = l$ is a consistent symmetry of any formula φ . More interestingly, consider $\varphi = \{\{\neg p, r\}, \{q, r\}, \{r, [m]\{\neg p, q\}\}\}$, then the permutation $\sigma = (p \neg q)(\neg p q)$ is a consistent symmetry of φ .

Now, since a permutation over literals can be lifted to transform some formula φ into another formula $\sigma(\varphi)$, we also want to consider permutations applied to models. Indeed, if φ is true in some \mathcal{M} , we intuitively want $\sigma(\varphi)$ to be true in some model obtained from lifting σ to \mathcal{M} . Thus the next step is to define the notion of applying permutations to models.

Definition 10 (Permutation of a model). Let σ be a permutation and $\mathcal{M} = \langle w, W, V, R \rangle$ a model. Then $\sigma(\mathcal{M}) = \langle w, W, V', R' \rangle$, where,

$$\begin{aligned} V'(v) &= \sigma(L_{V(v)}) \cap \text{Atom} \quad \text{for all } v \in W, \text{ and,} \\ R'(m, v) &= \{\sigma(\mathcal{N}) \mid \mathcal{N} \in R(\sigma(m), v)\} \quad \text{for all } m \in \text{Mod and } v \in W. \end{aligned}$$

For M a set of models, $\sigma(M) = \{\sigma(\mathcal{M}) \mid \mathcal{M} \in M\}$.

The main ingredient to prove that symmetries preserve entailment is the relation between models that we call σ -simulation.

Definition 11 (σ -simulation). Let σ be a permutation. A σ -simulation between models $\mathcal{M} = \langle w, W, V, R \rangle$ and $\mathcal{M}' = \langle w', W', V', R' \rangle$ is a non-empty relation $Z \subseteq \text{Ext}(\mathcal{M}) \times \text{Ext}(\mathcal{M}')$ that satisfies the following conditions:

- **Root:** $\mathcal{M} Z \mathcal{M}'$.
- **Harmony:** $l \in L_{V(w)} \text{ iff } \sigma(l) \in L_{V'(w')}$.

- **Zig:** If $\mathcal{M}Z\mathcal{M}'$ and $\mathcal{N} \in R(m, w)$ then $\mathcal{N}Z\mathcal{N}'$ for some $\mathcal{N}' \in R'(\sigma(m), w')$.
- **Zag:** If $\mathcal{M}Z\mathcal{M}'$ and $\mathcal{N}' \in R'(m, w')$ then $\mathcal{N}Z\mathcal{N}'$ for some $\mathcal{N} \in R(\sigma^{-1}(m), w)$.

We say that two models \mathcal{M} and \mathcal{M}' are σ -similar (notation $\mathcal{M} \xrightarrow{\sigma} \mathcal{M}'$) if there is a σ -simulation Z between them.

Notice that while $\mathcal{M} \xrightarrow{\sigma} \mathcal{M}'$ implies $\mathcal{M}' \xrightarrow{\sigma^{-1}} \mathcal{M}$, the relation $\xrightarrow{\sigma}$ is not symmetric (in particular σ might differ from σ^{-1}). From the definition of σ -simulations it intuitively follows that while they do not preserve validity of modal formulas (as is the case with bisimulations) they do preserve validity of *permutations* of formulas.

Proposition 1. *Let σ be a consistent permutation, φ a modal CNF formula and $\mathcal{M} = \langle w, W, V, R \rangle$, $\mathcal{M}' = \langle w', W', V', R' \rangle$ models such that $\mathcal{M} \xrightarrow{\sigma} \mathcal{M}'$. Then $\mathcal{M} \models \varphi$ iff $\mathcal{M}' \models \sigma(\varphi)$.*

Proof. The proof is by induction on φ . Base Case: Suppose $\varphi = a$ then, $\mathcal{M} \models a$ iff $a \in V(w)$ iff $a \in L_{V(w)}$ iff, by definition of σ -simulation, $\sigma(a) \in L_{V'(w')}$ iff $\mathcal{M}' \models \sigma(a)$.

Suppose $\varphi = \neg a$ then, $\mathcal{M} \models \neg a$ iff $a \notin V(w)$ iff $\neg a \in L_{V(w)}$ iff, by definition of σ -simulation, $\sigma(\neg a) = \neg\sigma(a) \in L_{V'(w')}$ iff $\sigma(a) \notin V'(w')$ iff $\mathcal{M}' \models \neg\sigma(a)$.

When $\varphi = C$, with C a clause or a conjunction of clauses, the proof follows by induction directly.

Inductive Step: Suppose $\varphi = [m]\psi$. Then $\mathcal{M} \models [m]\psi$ iff $\mathcal{N} \models \psi$ for all $\mathcal{N} \in R(m, w)$. Given that $\mathcal{M} \xrightarrow{\sigma} \mathcal{M}'$, by Zig we know that for all \mathcal{N} exist \mathcal{N}' such that $\mathcal{N} \xrightarrow{\sigma} \mathcal{N}'$ and $\mathcal{N}' \in R'(\sigma(m), w')$. Then, by inductive hypothesis, $\mathcal{N}' \models \sigma(\psi)$ for all $\mathcal{N}' \in R'(\sigma(m), w')$ iff $\mathcal{M}' \models [\sigma(m)]\sigma(\psi)$. Then, by Definition 8, $\mathcal{M}' \models \sigma([m]\psi)$. The converse uses Zag and the inductive hypothesis.

Suppose $\varphi = \neg[m]\psi$. Then $\mathcal{M} \models \neg[m]\psi$ iff there exists $\mathcal{N} \in R(m, w)$ such that, $\mathcal{N} \models \neg\psi$. Given that $\mathcal{M} \xrightarrow{\sigma} \mathcal{M}'$, by Zig we know that for all \mathcal{N} exist \mathcal{N}' such that $\mathcal{N} \xrightarrow{\sigma} \mathcal{N}'$ and $\mathcal{N}' \in R'(\sigma(m), w')$. Then, by inductive hypothesis, $\mathcal{N}' \models \sigma(\neg\psi) = \neg\sigma(\psi)$ iff $\mathcal{M}' \models \neg[\sigma(m)]\sigma(\psi)$. Then, by Definition 8, $\mathcal{M}' \models \sigma(\neg[m]\psi)$. The converse follows using Zag and the inductive hypothesis. \square

An easily verifiable consequence of Definitions 10 and 11 is that \mathcal{M} and $\sigma(\mathcal{M})$ are always σ -similar.

Proposition 2. *Let σ be a consistent permutation and $\mathcal{M} = \langle w, W, V, R \rangle$ a model. Then $\mathcal{M} \xrightarrow{\sigma} \sigma(\mathcal{M})$.*

Proof. Let us define the relation $Z = \{(\mathcal{N}, \sigma(\mathcal{N})) \mid \mathcal{N} \in \text{Ext}(\mathcal{M})\}$ and show that it is a σ -simulation between \mathcal{M} and $\sigma(\mathcal{M})$. The Zig and Zag conditions are trivial by definition of $\sigma(\mathcal{M})$.

For *Harmony*, we have to check that $l \in L_{V(w)}$ iff $\sigma(l) \in L_{V'(w')}$. From the definition of $\sigma(\mathcal{M})$, $L_{V'(w')} = \sigma(L_{V(w)})$, hence if $l \in L_{V(w)}$ then $\sigma(l) \in \sigma(L_{V(w)})$. Moreover, $\sigma(L_{V(w)})$ is a complete set of literals because $L_{V(w)}$ is a complete set of literals and σ is a consistent permutation, and hence the converse also follows. \square

Interestingly, if σ is a symmetry of φ then for any model \mathcal{M} , \mathcal{M} is a model of φ if and only if $\sigma(\mathcal{M})$ is. This will be a direct corollary of the following proposition in the particular case when σ is a symmetry and hence $\sigma(\varphi) = \varphi$.

Proposition 3. *Let σ be a consistent permutation, \mathcal{M} a model and φ a modal CNF formula. Then $\mathcal{M} \models \varphi$ iff $\sigma(\mathcal{M}) \models \sigma(\varphi)$.*

Proof. From Proposition 2 ($\mathcal{M} \xrightarrow{\sigma} \sigma(\mathcal{M})$) and Proposition 1. \square

Corollary 1. *If σ is a symmetry of φ then $\mathcal{M} \in \text{Mod}(\varphi)$ iff $\sigma(\mathcal{M}) \in \text{Mod}(\varphi)$.*

To clarify the implications of the Corollary 1, consider the following example.

Example 4. Let $\varphi = (p \vee q \vee r) \wedge (s \vee q \vee r) \wedge (\neg p \vee \neg s) \wedge \langle m \rangle (p \vee s) \wedge [A](\neg r)$. From Figure 3a we can verify that $\mathcal{M}_1 \models \varphi$.

Now $\sigma = (p\ s)(\neg p\ \neg s)$ is a symmetry of φ . Then, by Corollary 1, we have $\sigma(\mathcal{M}_1) \models \varphi$, which can be verified in the model of Figure 3b.

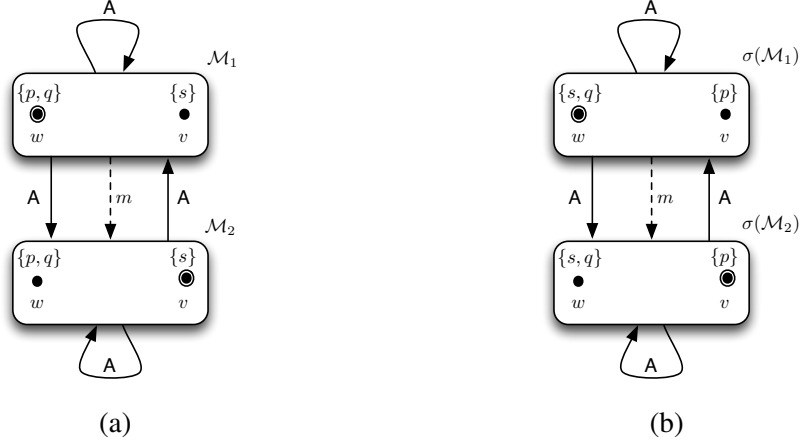


Figure 3: a) Model \mathcal{M}_1 . b) Model $\sigma(\mathcal{M}_1)$.

The notion of σ -simulation in coinductive modal models is general enough to be applicable to a wide range of modal logics. Notice though, that our definition of σ -simulation makes no assumption about the models being in the same class. Consider, for example, a model $\mathcal{M} \in \mathcal{C}_{\mathcal{H}(\@)}$ and a permutation $\sigma = (i\ p)(\neg i\ \neg p)$ for $i \in \text{Nom}$, $p \in \text{Prop}$. By the defining condition $\mathcal{C}_{\mathcal{H}(\@)}$, nominals in \mathcal{M} are true at a unique element in the domain, but this does not necessary hold for $\sigma(\mathcal{M})$, and hence $\sigma(\mathcal{M})$ might not be in $\mathcal{C}_{\mathcal{H}(\@)}$. Hence, when working with subclasses of **Mods** we will often have to require additional conditions to a permutation σ to ensure that for every \mathcal{M} , $\sigma(\mathcal{M})$ is in the intended class.

Definition 12. Let σ be a permutation and \mathcal{C} a closed class of models. We say that \mathcal{C} is closed under σ if for every $\mathcal{M} \in \mathcal{C}$, $\sigma(\mathcal{M}) \in \mathcal{C}$.

Example 5 (σ -simulation in hybrid logic). Consider the class $\mathcal{C}_{\mathcal{H}(\@)}$. $\mathcal{C}_{\mathcal{H}(\@)}$ is not closed under arbitrary permutations, but it is closed under permutations that send nominals to nominals.

Everything is now in place to show that modal entailment is preserved under symmetries.

Theorem 2. Let φ and ψ be modal formulas, let σ be a consistent symmetry of φ and \mathcal{C} a class of models closed under σ . Then $\varphi \models_{\mathcal{C}} \psi$ if and only if $\varphi \models_{\mathcal{C}} \sigma(\psi)$.

Proof. We first show that under the hypothesis of the theorem the following property holds

Claim: $\text{Mod}_{\mathcal{C}}(\varphi) = \sigma(\text{Mod}_{\mathcal{C}}(\varphi))$.

[\supseteq] Let $\mathcal{N} \in \sigma(\text{Mod}_{\mathcal{C}}(\varphi))$ and $\mathcal{M} \in \text{Mod}_{\mathcal{C}}(\varphi)$ be such that $\mathcal{N} = \sigma(\mathcal{M})$. Then $\mathcal{M} \models \varphi$ and by Corollary 1, $\sigma(\mathcal{M}) \models \varphi$. Given that \mathcal{C} is closed under σ , $\sigma(\mathcal{M}) \in \mathcal{C}$ and, hence, $\sigma(\mathcal{M}) \in \text{Mod}_{\mathcal{C}}(\varphi)$.

[\subseteq] Let $\mathcal{M} \in \text{Mod}_{\mathcal{C}}(\varphi)$, then $\mathcal{M} \models \varphi$. By Corollary 1, $\sigma(\mathcal{M}) \models \varphi$ and, given that \mathcal{C} is closed under σ , $\sigma(\mathcal{M}) \in \mathcal{C}$. Therefore, $\sigma(\mathcal{M}) \in \text{Mod}_{\mathcal{C}}(\varphi)$. Because σ is arbitrary, the results holds also for σ^k , $k \in \mathbb{Z}$.

Because σ is a permutation over a finite set, there exists n such that σ^n is the identity permutation. Now consider $\sigma^{n-1}(\mathcal{M})$, we know $\sigma^{n-1}(\mathcal{M}) \in \text{Mod}_{\mathcal{C}}(\varphi)$. Hence $\sigma^n(\mathcal{M}) = \mathcal{M} \in \sigma(\text{Mod}_{\mathcal{C}}(\varphi))$.

Now, we have to prove that $\varphi \models_{\mathcal{C}} \psi$ iff $\varphi \models_{\mathcal{C}} \sigma(\psi)$. By definition, $\varphi \models_{\mathcal{C}} \psi$ iff $\text{Mod}_{\mathcal{C}}(\varphi) \models_{\mathcal{C}} \psi$. By Proposition 3, this is the case if and only if $\sigma(\text{Mod}_{\mathcal{C}}(\varphi)) \models_{\mathcal{C}} \sigma(\psi)$.

Given that σ is a symmetry of φ , by the Claim above, $\sigma(\text{Mod}_{\mathcal{C}}(\varphi)) \models_{\mathcal{C}} \sigma(\psi)$ iff $\text{Mod}_{\mathcal{C}}(\varphi) \models_{\mathcal{C}} \sigma(\psi)$, which by definition means $\varphi \models_{\mathcal{C}} \sigma(\psi)$. \square

Theorem 2 provides an inexpensive inference mechanism that can be used in every situation where entailment is involved during modal automated reasoning. Indeed, applying a permutation on a formula is a calculation that is arguably computationally cheaper than a tableau expansion or a resolution step. Therefore, new formulas obtained by this mean may reduce the total running time of an inference algorithm. In the case of propositional logic, the strengthening of the learning mechanism has already shown its results in [8]. In the case of modal logic, it remains to see when cases of $\varphi \models \psi$ occur during a decision procedure, and how to better take advantage of them.

4 Layered Permutations

In this section we present the notion of layered permutations. First, we present a definition of the tree model property [11] for coinductive modal models that we will use.

Given a model \mathcal{M} , a (finite) *path rooted at \mathcal{M}* is a sequence $\pi = (\mathcal{M}_0, m_1, \mathcal{M}_1, \dots, m_k, \mathcal{M}_k)$, for $m_i \in \text{Mod}$ where $\mathcal{M}_0 = \mathcal{M}$, $k \geq 0$, and $\mathcal{M}_i \in R(m_i, w^{\mathcal{M}_i-1})$ for $i = 1, \dots, k$. For a path $\pi = (\mathcal{M}_0, m_1, \mathcal{M}_1, \dots, m_k, \mathcal{M}_k)$ we define $\text{first}(\pi) = \mathcal{M}_0$, $\text{last}(\pi) = \mathcal{M}_k$, and $\text{length}(\pi) = k$. We denote the set of all paths rooted at \mathcal{M} as $\Pi[\mathcal{M}]$. A coinductive tree model is a model that has a unique path to every reachable model (every model in $\text{Ext}(\mathcal{M})$). Formally we can define the class of all coinductive tree models, $\mathcal{C}_{\text{Tree}}$, with the following defining condition:

$$\mathcal{C}_{\text{Tree}} := P_{\text{Tree}}(\mathcal{M}) \iff \text{last} : \Pi[\mathcal{M}] \mapsto \text{Ext}(\mathcal{M}) \text{ is bijective.}$$

For example, the *unravelling* construction (in its version for coinductive modal models) shown below always defines a model in $\mathcal{C}_{\text{Tree}}$.

Definition 13 (Model Unravelling). *Given a model $\mathcal{M} = \langle w, W, V, R \rangle$, the unravelling of \mathcal{M} , (notation $\mathcal{T}(\mathcal{M})$), is the rooted coinductive model $\mathcal{T}(\mathcal{M}) = \langle (\mathcal{M}), \Pi[\mathcal{M}], V', R' \rangle$ where*

$$\begin{aligned} V'(\pi) &= V(w^{\text{last}(\pi)}), \text{ for all } \pi \in \Pi[\mathcal{M}], \\ R'(m, \pi) &= \{ \langle \pi', \Pi[\mathcal{M}], V', R' \rangle \mid \text{last}(\pi') \in R(m, \text{last}(\pi)) \}, \text{ for } m \in \text{Mod}, \pi \in \Pi[\mathcal{M}]. \end{aligned}$$

It is easy to verify that given a model \mathcal{M} , its unravelling $\mathcal{T}(\mathcal{M})$ is a tree ($\mathcal{T}(\mathcal{M}) \in \mathcal{C}_{\text{Tree}}$) and, as expected, \mathcal{M} and $\mathcal{T}(\mathcal{M})$ are bisimilar.

In what follows, we will use trees to define a more flexible family of symmetries that we call *layered symmetries*. The following will give a sufficient condition ensuring that layered symmetries also preserve entailment.

Definition 14 (Tree model closure property). *We say that a class \mathcal{C} of models is closed under trees if for every model $\mathcal{M} \in \mathcal{C}$ there is a tree model $\mathcal{T} \in \mathcal{C}$ such that $\mathcal{M} \leftrightarrow \mathcal{T}$.*

From this definition, it follows that a class of models \mathcal{C} closed under unravellings ($\mathcal{T}(\mathcal{M}) \in \mathcal{C}$ for all $\mathcal{M} \in \mathcal{C}$) is also closed under trees.

Example 6. *Trivially the class **Mods** (i.e., the basic modal logic) is closed under trees, and so does the class $\mathcal{C}_{\text{KAlt}_1}$ of models where the accessibility relation is a partial function. Many classes like \mathcal{C}_{A} , $\mathcal{C}_{\text{@}}$, and \mathcal{C}_{Nom} fail to be closed under trees.*

Logics defined over classes closed under trees have an interesting property: there is a direct correlation between the syntactical modal depth of the formula and the depth in a tree model satisfying it. In tree models, a notion of layer is induced by the depth (distance from the root) of the nodes in the model. Similarly, in modal formulas, a notion of layer is induced by the nesting of the modal operators. A consequence of this correspondence is that literals occurring at different formula layers are semantically independent of each other (see [3] for further discussion), i.e., at different layers the same literal can be assigned a different value.

Example 7. Consider the formula $\varphi = (p \vee q) \wedge (r \vee \neg \Box(\neg p \vee q \vee \Box \neg r))$ and a tree model \mathcal{M} of φ . Figure 4 shows the layers induced by the modal depth of the formula and the corresponding depth in \mathcal{M} .

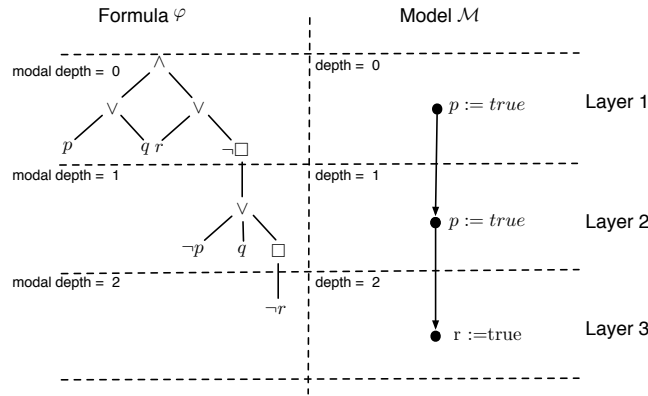


Figure 4: Induced layering on a model and a formula.

The independence between literals at different layers enables us to give a more flexible notion of a permutation that we will call *layered permutation*. Key to the notion of layered permutation is that of a *permutation sequence*.

Definition 15 (Permutation Sequence). We define a finite permutation sequence $\bar{\sigma}$ as either $\bar{\sigma} = \langle \rangle$ (i.e., $\bar{\sigma}$ is the empty sequence) or $\bar{\sigma} = \sigma : \bar{\sigma}_2$ with σ a permutation and $\bar{\sigma}_2$ a permutation sequence. Alternatively we can use the notation $\bar{\sigma} = \langle \sigma_1, \dots, \sigma_n \rangle$ instead of $\bar{\sigma} = \sigma_1 : \dots : \sigma_n : \langle \rangle$.

Let $|\bar{\sigma}| = n$ be the length of $\bar{\sigma}$ ($\langle \rangle$ has length 0). For $1 \leq i \leq n$, we write $\bar{\sigma}_i$ for the subsequence that starts from the i^{th} element of $\bar{\sigma}$. For $i \geq n$, we define $\bar{\sigma}_i = \langle \rangle$. In particular $\bar{\sigma} = \bar{\sigma}_1$. Given a permutation sequence $\sigma_1 : \bar{\sigma}_2$ we define $\text{head}(\sigma_1 : \bar{\sigma}_2) = \sigma_1$ and $\text{head}(\langle \rangle) = \sigma_{1d}$, where σ_{1d} is the identity permutation. We say that a permutation sequence is consistent if all of its permutations are consistent.

Applying a permutation sequence to a modal CNF formula can be defined as follows:

Definition 16 (Layered permutation of a formula). Let φ be a modal CNF formula and $\bar{\sigma}$ a permutation sequence. We define $\bar{\sigma}(\varphi)$ recursively:

$$\begin{aligned} \langle \rangle(\varphi) &= \varphi \\ (\sigma_1 : \bar{\sigma}_2)(l) &= \sigma_1(l) && \text{for } l \in \text{ALit} \\ (\sigma_1 : \bar{\sigma}_2)([m]C) &= [\sigma_1(m)]\bar{\sigma}_2(C) \\ \bar{\sigma}(C) &= \{\bar{\sigma}(A) \mid A \in C\} && \text{for } C \text{ a clause or a formula.} \end{aligned}$$

Notice that layered permutations are well defined even if the modal depth of the formula is greater than the size of the permutation sequence. Layered permutations let us use a different permutation at each modal depth. This enables symmetries (layered symmetries) to be found, that would not be found otherwise.

Example 8. Consider the formula $\varphi = (p \vee [m](p \vee \neg r)) \wedge (\neg q \vee [m](\neg p \vee r))$. If we only consider non-layered symmetries then φ has none. However, the permutation sequence $\langle \sigma_1, \sigma_2 \rangle$ generated by $\sigma_1 = (p \neg q)$ and $\sigma_2 = (p \neg r)$ is a layered symmetry of φ .

As we can see from the previous example, layered permutations let us map the same literal to different targets at each different modal depth. This additional degree of freedom can result in new symmetries for a given formula.

From now on we can mostly repeat the work we did in the previous section to arrive to a result similar to Theorem 2 but involving permutation sequences, with one caveat: the obvious extension of the notion of permuted model $\sigma(\mathcal{M})$ to layered permutations is ill defined if \mathcal{M} is not a tree. Hence, we need the additional requirement that the class \mathcal{C} of models is closed under trees for the result to go through.

Definition 17 (Layered Permutation of a model). Let $\bar{\sigma}$ be a permutation sequence and $\mathcal{M} = \langle w, W, V, R \rangle$ a tree model. Then $\bar{\sigma}(\mathcal{M}) = \langle w, W, V', R' \rangle$, where,

$$\begin{aligned} V'(v) &= \text{head}(\bar{\sigma})(L_{V(v)}) \cap \text{Atom} \quad \text{for all } v \in W, \text{ and,} \\ R'(m, v) &= \{ \bar{\sigma}_2(\mathcal{N}) \mid \mathcal{N} \in R(\text{head}(\bar{\sigma})(m), v) \} \quad \text{for all } m \in \text{Mod and } v \in W. \end{aligned}$$

For M a set of tree models, $\bar{\sigma}(M) = \{ \bar{\sigma}(\mathcal{M}) \mid \mathcal{M} \in M \}$.

We can now extend the notion of σ -simulation to permutation sequences.

Definition 18 ($\bar{\sigma}$ -simulation). Let $\bar{\sigma}$ be a permutation sequence. A $\bar{\sigma}$ -simulation between models $\mathcal{M} = \langle w, W, V, R \rangle$ and $\mathcal{M}' = \langle w', W', V', R' \rangle$ is a family of relations $Z_{\bar{\sigma}_i} \subseteq \text{Ext}(\mathcal{M}) \times \text{Ext}(\mathcal{M}')$, $1 \leq i$, that satisfies the following conditions:

- **Root:** $\mathcal{M} Z_{\bar{\sigma}_1} \mathcal{M}'$.
- **Harmony:** If $w Z_{\bar{\sigma}_i} w'$ then $l \in L_{V(w)}$ iff $\text{head}(\bar{\sigma}_i)(l) \in L_{V'(w')}$.
- **Zig:** If $\mathcal{M} Z_{\bar{\sigma}_i} \mathcal{M}'$ and $\mathcal{N} \in R(m, w)$ then $\mathcal{N} Z_{\bar{\sigma}_{i+1}} \mathcal{N}'$ for some $\mathcal{N}' \in R'(\text{head}(\bar{\sigma}_i)(m), w')$.
- **Zag:** If $\mathcal{M} Z_{\bar{\sigma}_i} \mathcal{M}'$ and $\mathcal{N}' \in R'(m, w')$ then $\mathcal{N} Z_{\bar{\sigma}_{i+1}} \mathcal{N}'$ for some $\mathcal{N} \in R(\text{head}(\bar{\sigma}_i)^{-1}(m), w)$.

We say that two models \mathcal{M} and \mathcal{M}' are $\bar{\sigma}$ -similar (notation $\mathcal{M} \rightrightarrows_{\bar{\sigma}} \mathcal{M}'$), if there is a $\bar{\sigma}$ -simulation between them.

An important remark about the previous definition is that it does not make any assumption about the size of the permutation sequence. In fact, it is well defined even if the permutation sequence at hand is the empty sequence. In that case, it just behave as the identity permutation at each layer, thus the relation defines a bisimulation between the models.

Given a closed class of tree models \mathcal{C} and $\bar{\sigma}$ a permutation sequence, we say that \mathcal{C} is closed under $\bar{\sigma}$ if for every $\mathcal{M} \in \mathcal{C}$, $\bar{\sigma}(\mathcal{M}) \in \mathcal{C}$.

Now we are ready to prove the main result concerning layered symmetries and entailment.

Theorem 3. Let φ and ψ be modal formulas and let $\bar{\sigma}$ be a consistent permutation sequence, and let \mathcal{C} be a class of models closed under trees and $\mathcal{C} \cap \mathcal{C}_{Tree}$ closed under $\bar{\sigma}$. If $\bar{\sigma}$ is a symmetry of φ then for any ψ we have that $\varphi \models_{\mathcal{C}} \psi$ if and only if $\varphi \models_{\mathcal{C}} \bar{\sigma}(\psi)$.

Proof. We first show that under the hypothesis of the theorem the following two properties hold.

Claim 1: $\text{Mod}_{\mathcal{L} \cap \mathcal{L}_{Tree}}(\varphi) = \bar{\sigma}(\text{Mod}_{\mathcal{L} \cap \mathcal{L}_{Tree}}(\varphi))$.

The argument is the same as for the Claim in Theorem 2 but using permutation sequences.

Claim 2: $\text{Mod}_{\mathcal{L}}(\varphi) \models_{\mathcal{L}} \varphi$ iff $\text{Mod}_{\mathcal{L} \cap \mathcal{L}_{Tree}}(\varphi) \models_{\mathcal{L}} \varphi$.

The left-to-right direction is trivial by the fact that $\text{Mod}_{\mathcal{L} \cap \mathcal{L}_{Tree}}(\varphi) \subseteq \text{Mod}_{\mathcal{L}}(\varphi)$. For the other direction, assume $\text{Mod}_{\mathcal{L} \cap \mathcal{L}_{Tree}}(\varphi) \models_{\mathcal{L}} \varphi$ and $\text{Mod}_{\mathcal{L}}(\varphi) \not\models_{\mathcal{L}} \varphi$. Then there is $\mathcal{M} \in \text{Mod}_{\mathcal{L}}(\varphi)$ such that $\mathcal{M} \not\models_{\mathcal{L}} \varphi$. But we know that $\mathcal{M} \xleftrightarrow{\tau} \mathcal{T}$, and $\mathcal{T} \in \text{Mod}_{\mathcal{L} \cap \mathcal{L}_{Tree}}(\varphi)$. Hence $\mathcal{T} \models_{\mathcal{L}} \varphi$ which contradicts our assumption.

It rests to prove that $\varphi \models_{\mathcal{L}} \psi$ if and only if $\varphi \models_{\mathcal{L}} \bar{\sigma}(\psi)$. By definition, $\varphi \models_{\mathcal{L}} \psi$ if and only if $\text{Mod}_{\mathcal{L}}(\varphi) \models_{\mathcal{L}} \psi$. By Claim 2, this is the case if and only if $\text{Mod}_{\mathcal{L} \cap \mathcal{L}_{Tree}}(\varphi) \models_{\mathcal{L}} \psi$. By the layered version of Proposition 3, this is the case if and only if $\bar{\sigma}(\text{Mod}_{\mathcal{L} \cap \mathcal{L}_{Tree}}(\varphi)) \models_{\mathcal{L}} \bar{\sigma}(\psi)$. Given that $\bar{\sigma}$ is a symmetry of φ , by Claim 1, $\bar{\sigma}(\text{Mod}_{\mathcal{L} \cap \mathcal{L}_{Tree}}(\varphi)) \models_{\mathcal{L}} \bar{\sigma}(\psi)$ if and only if $\text{Mod}_{\mathcal{L} \cap \mathcal{L}_{Tree}}(\varphi) \models_{\mathcal{L}} \bar{\sigma}(\psi)$, which by Claim 2 is the case if and only if $\text{Mod}_{\mathcal{L}}(\varphi) \models_{\mathcal{L}} \bar{\sigma}(\psi)$ which by definition means that $\varphi \models_{\mathcal{L}} \bar{\sigma}(\psi)$. \square

5 Symmetry Detection

Different techniques have been proposed for detecting symmetries of propositional formulas in clausal form. Some of them, deal directly with the formula [10], while others, reduce the problem to the problem of finding automorphisms in colored graphs constructed in such a way that the automorphism group of the graph is isomorphic to the symmetry group of the formula under consideration [14, 15, 1].

The availability of efficient tools to detect graph automorphisms (e.g., [24, 16, 21]) has made the later approach the most successful one because it is fast and easy to integrate.

In this section we present a technique for the detection of symmetries in modal formulas that extends the construction proposed for propositional formulas to modal CNF formulas. We present the graph construction algorithm and prove its correctness.

We now introduce some notation and definitions. In what follows, we consider modal CNF formulas as set of sets as defined in Section 2 and write $\psi \in \varphi$ to express that ψ is subformula of φ . Clauses occurring at modal depth 0 are named *top clauses* and clauses occurring in modal literals are named *modal clauses*. Let $s : \text{Mod} \times \{0, 1\} \mapsto \mathbb{N} \setminus \{0, 1\}$ be an injective function and let $t : \text{Sub}(\varphi) \mapsto \mathbb{N}$ be a partial function defined as:

$$t(\psi) = \begin{cases} 1 & \text{if } \psi \text{ is a top clause} \\ s(m, 0) & \text{if } \psi = [m]C \\ s(m, 1) & \text{if } \psi = \neg[m]C \end{cases}$$

The *typing* function t assigns a numeric type to every clause (top or modal). For modal clauses, the type is based on the modality and the polarity of the modal literal in which it occurs. Let us assign to each clause C occurring in φ a unique identifier $id(C) = \langle m, k, i \rangle$ where m is the modal depth at which the clause occurs, $k = t(\psi)$ is the type of the clause as returned by the *typing* function t and $i \in \mathbb{N}$ is different for each clause. To simplify notation, in what follows we will assume that each clause C is labeled by its unique identifier $id(C) = \langle m, k, i \rangle$ and write $C_{m,k,i}$.

By definition, a symmetry of a formula φ is a bijective function that maps literals to literals. It can naturally be extended to a function σ_{ext} that also maps each clause C to $\sigma(C)$. Notice that because σ is a symmetry of φ both C and $\sigma(C)$ are clauses in φ . Hence, both C and $\sigma(C)$ will be assigned some identifier by the *id* function.

The following are properties of σ_{ext} that are easy to verify.

Proposition 4. *Let φ be a modal CNF formula and σ a symmetry of φ . Then for the extension of σ , σ_{ext} , the following holds:*

- i) σ_{ext} is a bijective function.
- ii) If $\sigma_{ext}(C_{m,k,i}) = C_{m',k',i'}$ then $m = m'$.
- iii) If $\sigma_{ext}(C_{m,k,i}) = C_{m',k',i'}$ then $k = k'$.
- iv) If $l \in C_{m,k,i}$ then $\sigma_{ext}(l) \in \sigma_{ext}(C_{m,k,i})$.
- v) σ_{ext} is a symmetry of φ .

We can now introduce the construction of the colored graph corresponding to a given formula φ . We will construct an undirected colored graph with two types of edges. As coloring function we will use the typing function t introduced earlier in this section.

Definition 19. *Let φ be a modal CNF formula and let $At(\varphi)$ denote the set of atoms occurring in φ . The colored graph $G(\varphi) = (V, E_1, E_2)$ is constructed as follows:*

1. For each atom $a \in At(\varphi)$:
 - (a) Add two literal nodes of color 0: one labelled a and one labelled $\neg a$.
 - (b) Add an edge to E_1 between these two nodes to ensure Boolean consistency.
2. For each top clause C of φ add a clause node of color $t(C)$.
3. For each atom literal occurring in C , add an edge to E_1 from C to the corresponding literal node.
4. For each modal literal $[m]C'$ ($\neg[m]C'$) occurring in C :
 - (a) Add a clause node of color $t([m]C')$ ($t(\neg[m]C')$) to represent the modal clause C' .
 - (b) Add an edge to E_1 from the node of C to this node.
 - (c) If m is indexed by an atom literal l then add an edge to E_2 from the node C' to the indexing literal l .
 - (d) Repeat the process from point 3 for each literal (atom or modal) occurring in C' .


This construction creates a graph with $2 + 2|\text{Mod}|$ colours and at most $(2|V| + \#(\text{TopClauses}) + \#(\text{ModalClauses}))$ nodes.

Example 9. *Let us consider the following formula $\varphi = (a \vee [m](b \vee \neg[m]c)) \wedge (b \vee [m](a \vee \neg[m]c))$. This formula has six clauses (2 at modal depth 0, 2 at modal depth 1 and 2 at modal depth 2) and three atoms (six literals). The associated colored graph, $G(\varphi)$, is shown in Figure 5 (colors are represented by shapes in the figure).*

Note that the construction of Definition 19 induces a mapping g that associates to each literal and clause the corresponding node in the graph.

To prove that the proposed construction is correct, we first have to show that each symmetry of the formula is a colored automorphism of the graph.

Proposition 5. *Let φ be a modal CNF formula, σ a symmetry of φ , $G(\varphi) = (V, E_1, E_2)$ the colored graph of φ as defined by Definition 19, and g the mapping induced by the construction of $G(\varphi)$. Then $\pi = g \circ \sigma_{ext}$ is an automorphism of $G(\varphi)$.*



$$\begin{aligned}
 A &= (a \vee [m](b \vee \neg[m]c)) \\
 B &= (b \vee [m](a \vee \neg[m]c)) \\
 C &= [m](b \vee \neg[m]c) \\
 D &= [m](a \vee \neg[m]c) \\
 E &= [m]c \\
 F &= [m]c
 \end{aligned}$$

Figure 5: Graph representation of φ .

Proof. To simplify notation let's assume that g is the identity function (i.e., we do not differentiate between a clause (or literal) and its associated node in the graph) and as a consequence $\pi = \sigma_{ext}$. Then π is an automorphism of $G(\varphi)$ if the following holds:

1. $(l, \neg l) \in E_1$ iff $(\pi(l), \pi(\neg l)) \in E_1$ for all $l \in V$.

We have to consider the following cases:

- σ is a permutational symmetry:

(\rightarrow) : Assume $\pi = \sigma_{ext} = (a\ b)(\neg a\ \neg b)$. Then $(\pi(a), \pi(\neg a)) = (b, \neg b) \in E_1$ by construction of $G(\varphi)$.

(\leftarrow) : $(a, \neg a) \in E_1$ by construction of $G(\varphi)$.

- σ is a phase-shift symmetry:

(\rightarrow) : Assume $\pi = \sigma_{ext} = (a\ \neg a)$. Then $(\pi(a), \pi(\neg a)) = (\neg a, a)$, but given that $G(\varphi)$ is an undirected graph $(\neg a, a) = (a, \neg a) \in E_1$ by construction.

(\leftarrow) : $(a, \neg a) \in E_1$ by construction of $G(\varphi)$.

- σ is a compositional symmetry: It follows directly from the previous two cases.

2. $(l, C_{m,k,i}) \in E_1$ iff $(\pi(l), \pi(C_{m,k,i})) \in E_1$ for all $l, C_{m,k,i} \in V$.

(\rightarrow) : By construction, $(l, C_{m,k,i}) \in E_1$ only if $l \in C_{m,k,i}$. Then, given that σ_{ext} is a symmetry and by Proposition 4iv), we know that $\sigma_{ext}(l)$ and $\sigma_{ext}(C_{m,k,i})$ both occur in φ and $\sigma_{ext}(l) \in \sigma_{ext}(C_{m,k,i})$. Then, by construction, we have that $(\pi(l), \pi(C_{m,k,i})) \in E_1$.

(\leftarrow) : It follows directly by construction of $G(\varphi)$.

3. $(C_{m,k,i}, C_{m',k',i'}) \in E_1$ iff $(\pi(C_{m,k,i}), \pi(C_{m',k',i'})) \in E_1$ for all $C_{m,k,i}, C_{m',k',i'} \in V$.

(\rightarrow) : If $(C_{m,k,i}, C_{m',k',i'}) \in E_1$ we know that either $m < m'$ or $m > m'$. Assume $m < m'$ then $C_{m',k',i'}$ is a modal clause occurring in $C_{m,k,i}$. By Proposition 4iv), we have that $\sigma_{ext}(C_{m',k',i'})$ is a modal clause occurring in $\sigma_{ext}(C_{m,k,i})$, and given that σ_{ext} is a symmetry of φ , $\sigma_{ext}(C_{m,k,i})$ and $\sigma_{ext}(C_{m',k',i'})$ both occur in φ , therefore, by construction, $\pi(C_{m,k,i}) \in V$ and $\pi(C_{m',k',i'}) \in V$, therefore $(\pi(C_{m,k,i}), \pi(C_{m',k',i'})) \in E_1$.

(\leftarrow) : It follows directly by construction of $G(\varphi)$.

4. $(l, C_{m,k,i}) \in E_2$ iff $(\pi(l), \pi(C_{m,k,i})) \in E_2$ for all $l, C_{m,k,i} \in V$.

(\rightarrow) : $(l, C_{m,k,i}) \in E_2$ if the modality of the modal clause $C_{m,k,i}$ is indexed by l . Given that σ_{ext} is a symmetry of φ , we know that $\sigma_{ext}(l) \in \varphi$ and $\sigma_{ext}(C_{m,k,i}) \in \varphi$ and that $\sigma_{ext}(l)$ indexes the modality of the modal clause $\sigma_{ext}(C_{m,k,i})$, therefore, by construction, $\pi(l) \in V$ and $\pi(C_{m,k,i}) \in V$ and $(\pi(l), \pi(C_{m,k,i})) \in E_2$.

(\leftarrow) : It follows directly by construction of $G(\varphi)$.

5. For every cycle $(x y) \in \pi$, x and y have the same color.

Follows from Proposition 4iii) and the fact that by construction different types of clauses are assigned different colors in the graph.

□

We now prove that any colored automorphism of $G(\varphi)$ induces a symmetry of φ .

Proposition 6. *Let φ be a modal CNF formula, $G(\varphi) = (V, E_1, E_2)$ the colored graph of φ as defined by Definition 19, π an automorphism of $G(\varphi)$ and g the mapping induced by the construction of $G(\varphi)$. Then $\sigma_{ext} = g^{-1} \circ \pi$ is a symmetry of φ .*

Proof. Once more, assume g is the identity. To prove that σ_{ext} is a symmetry of φ , we have to prove the following properties:

1. σ_{ext} is a consistent permutation, i.e., $\sigma_{ext}(\neg l) = \neg \sigma_{ext}(l)$ for all $l \in \text{ALit}$.

By construction, Boolean consistency edges only connect literal nodes. Let $l_i \in V$ be a literal node. Then by construction we have that $(l_i, \neg l_i) \in E_1$. Now assume that $\pi(l_i) = l_j$ for $l_j \in V$. Given that π is an automorphism it must be the case that $(\pi(l_i), \pi(\neg l_i)) \in E_1$, and therefore that $\pi(\neg l_i) = \neg l_j = \neg \pi(l_i)$, which implies that $\sigma_{ext}(\neg l_i) = \neg \sigma_{ext}(l_i)$.

2. If $C_{m,k,i} \in \varphi$ then $\sigma_{ext}(C_{m,k,i}) \in \varphi$.

By construction $C_{m,k,i} \in V$ implies that $C_{m,k,i} \in \varphi$. As π is an automorphism, $\pi(C_{m,k,i}) \in V$, therefore, $\pi(C_{m,k,i}) \in \varphi$ which implies that $\sigma_{ext}(C_{m,k,i}) \in \varphi$.

3. If $l \in \varphi$ then $\sigma_{ext}(l) \in \varphi$.

It follows by the same argument as in the previous case.

4. If $\sigma_{ext}(C_{m,k,i}) = C_{m',k',i'}$ then $k = k'$.

It follows from the fact that π is a colored automorphism, mapping only nodes of the same color, and that by construction, clauses of the same type are assigned the same color in the graph.

5. If $\sigma_{ext}(C_{m,k,i}) = C_{m',k',i'}$ then $m = m'$.

We prove this by induction on m , the modal depth at which a clause occurs in φ .

Base Case: $m = 0$. We have to prove that if $\sigma_{ext}(C_{0,k,i}) = C_{m',k',i'}$ then $m' = 0$. Assume that $m' \neq 0$. Then, there is a clause $C_{n,s,j}$, with $n < m'$ such that, $C_{m',k',i'}$ is a modal clause occurring in it. By construction, we then have that $(C_{n,s,j}, C_{m',k',i'}) \in E_1$. As π is an automorphism of $G(\varphi)$, we should have $(\pi(C_{n,s,j}), \pi(C_{m',k',i'})) = (\pi(C_{n,s,j}), C_{0,k,i}) \in E_1$, but by construction there is no such edge.

Inductive Step: $n < m \implies m$. By construction of $G(\varphi)$ if $(C_{m,k,i}, C_{n,l,j}) \in E_1$ then $|m - n| = 1$. Now, assume $m \neq m'$. We know that there is a clause $C_{(m'-1),s,j}$ such that $(C_{(m'-1),s,j}, C_{m',k',i'}) \in E_1$. Then, as π is an automorphism of $G(\varphi)$, it must be the case that $(\pi(C_{(m'-1),s,j}), \pi(C_{m',k',i'})) = (\pi(C_{(m'-1),s,j}), C_{m,k,i}) \in E_1$. By the inductive hypothesis we know that $\pi(C_{(m'-1),s,j}) = C_{(m'-1),s,j'}$ and therefore, we have that $(C_{(m'-1),s,j'}, C_{m,k,i}) \in E_1$. But then we get that $|m - (m' - 1)| \geq 2$, which by construction cannot happen. Therefore $(C_{(m'-1),s,j'}, C_{m,k,i}) \notin E_1$, contradicting the fact that π is an automorphism of $G(\varphi)$.

6. If l index a clause $C_{m,k,i}$ then $\sigma_{ext}(l)$ index the clause $\sigma_{ext}(C_{m,k,i})$.

If l index a clause $C_{m,k,i}$, then by construction $(l, C_{m,k,i}) \in E_2$. Given that π is an automorphism, $(\pi(l), \pi(C_{m,k,i})) \in E_2$, which implies that $\sigma_{ext}(l)$ index the clause $\sigma_{ext}(C_{m,k,i})$.

We have proved that σ_{ext} , the extension of σ , obtained from an automorphism of the graph, is a symmetry of φ . To obtain the original symmetry σ we just take the restriction of σ_{ext} to atom literals. \square

Finally we can prove that our construction is correct.

Theorem 4. *Let φ be a modal CNF formula and $G(\varphi) = (V, E_1, E_2)$ the colored graph constructed following the construction of Definition 19. Then every symmetry σ of φ corresponds one-to-one to an automorphism π of $G(\varphi)$.*

Proof. Immediate from Proposition 5 and 6. \square

This construction enables the detection of symmetries as defined in Section 3, that is, symmetries defined over literals that can appear at various modal depths of a given formula. To detect layered symmetries (see Section 4) for logics with the tree model closure property, the construction needs to be modified to capture the notion of layers. This is easy to achieve by just changing the way literals occurring in the formula are handled (see [25] for details). Properties 5 and 6 (suitably generalized) also hold in this case.

6 Conclusions and Further Work

The notion of symmetry has been well studied in propositional logic, and various optimizations of decision procedures based on it are known. In this article, we extend the notion of syntactic and semantic symmetries to many different modal logics using the framework of coinductive models. The main contribution is that a symmetry σ preserves entailments whenever the class is closed by σ . For example, arbitrary symmetries preserve entailments in the basic modal logic, but for the hybrid logic $\mathcal{H}(@)$ we can only consider symmetries that map nominals to nominals. The second contribution of the paper is to show that if the class of models is closed under trees, then the more flexible notion of layered symmetry also preserve entailments.

To arrive at the previous results, we defined the concept of σ -simulation and showed that it preserves σ -permutation of formulas. We then presented permutation sequences $\bar{\sigma}$, and $\bar{\sigma}$ -simulations. Permutation sequences are relevant in those classes of models that are closed under trees. This property enables the use of layered symmetries, a notion that can capture more symmetries than the ordinary symmetry definition. Indeed, layered symmetries can be detected independently within atoms at each modal depth of a formula. $\bar{\sigma}$ -simulations extend the notion of σ -simulations to permutation sequences, and enabled us to prove that layered symmetries also preserve entailment.

Finally, we presented a method to detect symmetries in modal formulas that reduces the problem to the graph automorphism problem. Given a formula φ , the idea is to build a graph in such a way that the automorphism group of the graph is isomorphic to the symmetry group of the formula that generated it. A general construction algorithm, suitable for many modal logics, was presented and its correctness proved. The presented graph construction algorithm can be extended to detect layered symmetries. Preliminary results on modal symmetries concerning this last construction can be found in [25] where we developed an efficient algorithm to detect symmetries for the basic modal logic, and empirically verified that many modal problems (both randomly and hand generated) contain symmetries.

Our ongoing research focuses on the incorporation of symmetry information into a modal tableau calculi such as [22, 20] or modal resolution calculi such as [5]. One promising theme that we will investigate in the future is permutations involving also modal literals.

Acknowledgments. This work was partially supported by grants ANPCyT-PICT-2008-306, ANPCyT-PICT-2010-688, the FP7-PEOPLE-2011-IRSES Project “Mobility between Europe and Argentina applying Logics to Systems” (MEALS) and the Laboratoire International Associé “INFINIS”.

References

- [1] F. Aloul, A. Ramani, I. Markov & K. Sakallah (2003): *Solving difficult instances of Boolean satisfiability in the presence of symmetry*. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 22(9), pp. 1117–1137, doi:10.1109/TCAD.2003.816218.
- [2] C. Areces & B. ten Cate (2006): *Hybrid Logics*. In P. Blackburn, F. Wolter & J. van Benthem, editors: *Handbook of Modal Logics*, Elsevier, pp. 821–868, doi:10.1016/S1570-2464(07)80017-6.
- [3] C. Areces, R. Gennari, J. Heguiabehere & M. de Rijke (2000): *Tree-Based Heuristics in Modal Theorem Proving*. In: *Proceedings of ECAI’2000*, Berlin, Germany, pp. 199–203.
- [4] C. Areces & D. Gorín (2010): *Coinductive models and normal forms for modal logics (or how we learned to stop worrying and love coinduction)*. *Journal of Applied Logic* 8(4), pp. 305–318, doi:10.1016/j.jal.2010.08.010.
- [5] C. Areces & D. Gorín (2011): *Resolution with Order and Selection for Hybrid Logics*. *Journal of Automated Reasoning* 46(1), pp. 1–42, doi:10.1007/s10817-010-9167-0.
- [6] G. Audemard (2002): *Reasoning by symmetry and function ordering in finite model generation*. In: *Proceedings of CADE-18*, pp. 226–240, doi:10.1007/3-540-45620-1_19.
- [7] G. Audemard, B. Mazure & L. Sais (2004): *Dealing with Symmetries in Quantified Boolean Formulas*. In: *Proceedings of the 7th International Conference on Theory and Applications of Satisfiability Testing (SAT’04)*, pp. 257–262.
- [8] B. Benhamou, T. Nabhani, R. Ostrowski & M. Saidi (2010): *Enhancing Clause Learning by Symmetry in SAT Solvers*. In: *Proceedings of the 22nd IEEE International Conference on Tools with Artificial Intelligence (ICTAI)*, pp. 329–335, doi:10.1109/ICTAI.2010.55.
- [9] B. Benhamou & L. Sais (1992): *Theoretical Study of Symmetries in Propositional Calculus and Applications*. In: *Proceedings of CADE-11*, pp. 281–294, doi:10.1007/3-540-55602-8_172.
- [10] B. Benhamou & L. Sais (1994): *Tractability Through Symmetries in Propositional Calculus*. *Journal of Automated Reasoning* 12(1), pp. 89–102, doi:10.1007/BF00881844.
- [11] P. Blackburn, M. de Rijke & Y. Venema (2001): *Modal Logic*. Cambridge University Press.
- [12] P. Blackburn, J. van Benthem & F. Wolter (2006): *Handbook of Modal Logic*. *Studies in Logic and Practical Reasoning* 3, Elsevier Science Inc., New York, NY, USA, doi:10.1016/S1570-2464(07)80004-8.
- [13] C. Brown, L. Finkelstein & P. Purdom, Jr. (1996): *Backtrack searching in the presence of symmetry*. *Nordic Journal of Computing* 3(3), pp. 203–219, doi:10.1007/3-540-51083-4_51.
- [14] J. Crawford (1992): *A Theoretical Analysis of Reasoning By Symmetry in First-Order Logic*. In: *Proceedings of AAAI Workshop on Tractable Reasoning*, San Jose, CA, pp. 17–22.
- [15] J. Crawford, M. Ginsberg, E. Luks & A. Roy (1996): *Symmetry-Breaking Predicates for Search Problems*. In: *Proceedings of KR 1996*, pp. 148–159.
- [16] P. Darga, M. Liffiton, K. Sakallah & I. Markov (2004): *Exploiting structure in symmetry detection for CNF*. In: *Design Automation Conference, 2004. Proceedings. 41st*, pp. 530–534, doi:10.1145/996566.996712.
- [17] D. Déharbe, P. Fontaine, S. Merz & B. Woltzenlogel Paleo (2011): *Exploiting Symmetry in SMT Problems*. In: *Proceedings of CADE-23, Lecture Notes in Computer Science* 6803, Springer Berlin Heidelberg, pp. 222–236, doi:10.1007/978-3-642-22438-6_18.

- [18] N. Een & N. Sörensson (2003): *An Extensible SAT-solver*. In: *Proceedings of the 6th International Conference on Theory and Applications of Satisfiability Testing (SAT'03)*, pp. 502–518, doi:10.1007/978-3-540-24605-3_37.
- [19] J.B. Fraleigh & V.J. Katz (2003): *A first course in abstract algebra*. Addison-Wesley world student series, Addison-Wesley.
- [20] G. Hoffmann (2010): *Lightweight Hybrid Tableaux*. *Journal of Applied Logic* 8(4), pp. 397–408, doi:10.1016/j.jal.2010.08.003.
- [21] T. Junttila & P. Kaski (2007): *Engineering an Efficient Canonical Labeling Tool for Large and Sparse Graphs*. In: *Proceedings of the Workshop on Algorithm Engineering and Experiments, ALENEX 2007*, SIAM.
- [22] M. Kaminski & G. Smolka (2009): *Terminating Tableau Systems for Hybrid Logic with Difference and Converse*. *Journal of Logic, Language and Information* 18(4), pp. 437–464, doi:10.1007/s10849-009-9087-8.
- [23] B. Krishnamurthy (1985): *Short Proofs for Tricky Formulas*. *Acta Informatica* 22(3), pp. 253–275, doi:10.1007/BF00265682.
- [24] B. McKay (1990): *Nauty User's Guide*. Technical Report, Australian National University, Computer Science Department.
- [25] E. Orbe, C. Areces & G. Infante-López (2012): *A Note about Modal Symmetries*. Technical Report, FaMAF, UNC. Available at http://www.famaf.unc.edu.ar/publicaciones/documents/serie_a/AInf6.pdf.
- [26] P. Patel-Schneider & R. Sebastiani (2003): *A New General Method to Generate Random Modal Formulae for Testing Decision Procedures*. *Journal of Artificial Intelligence Research* 18, pp. 351–389.
- [27] M. Prasad, A. Biere & A. Gupta (2005): *A survey of recent advances in SAT-based formal verification*. *International Journal on Software Tools for Technology Transfer* 7, pp. 156–173, doi:10.1007/s10009-004-0183-4.
- [28] L. Ryan (2004): *Efficient Algorithms For Clause-Learning SAT Solvers*. Master's thesis, Simon Fraser University.