Coalescing: Syntactic Abstraction for Reasoning in First-Order Modal Logics *

Damien Doligez[1], Jael Kriener[2], Leslie Lamport[3], Tomer Libal[2], and Stephan Merz[4]

[1] Inria, Paris, France
[2] MSR-Inria Joint Centre, Saclay, France
[3] Microsoft Research, Mountain View, CA, U.S.A.
[4] Inria, Villers-lès-Nancy, France

Abstract

We present a syntactic abstraction method to reason about first-order modal logics by using theorem provers for standard first-order logic and for propositional modal logic.

1 Introduction

Verification of distributed and concurrent systems requires reasoning about temporal behaviors. A common approach is to express the properties to be proved in a modal logic having one or more temporal modalities. For verifying real-world systems, a proof language must also include equality, quantification, and local definitions. It must therefore encompass FOML (*F*irst *O*rder *M*odal *L*ogic) and support operator definitions. One such language is TLA$^+$ [8], based on the logic TLA that has two temporal modalities: the usual $\Box$ (always) operator of linear-time temporal logic and a restricted next-state operator represented by priming. (The syntax does not permit priming of an expression containing a modal operator.)

A common way to prove an FOML theorem $\Gamma \models \varphi$ (φ holds in context Γ) is to translate it to a semantically equivalent FOL theorem $\Gamma^* \models_{FOL} \varphi^*$ and to prove this FOL theorem. For some FOMLs, this method is semantically complete—that is, $\Gamma \models \varphi$ is valid iff $\Gamma^* \models_{FOL} \varphi^*$ is [10]. This approach has been followed for embedding FOML in SPASS [7], Saturate [6], and other theorem provers.

Such a semantic translation may be appropriate for completely automatic provers. However, we are very far from being able to automatically prove a formula that expresses a correctness property of a non-trivial system. A person must break the proof into smaller steps that we call *proof obligations*, usually by interacting with the prover. Requiring the user to interactively prove the semantic translation of the FOML formula destroys the whole purpose of using modal logic, which is to allow her to think in terms of the simpler FOML abstraction of the theorem. The user should therefore decompose the FOML proof into FOML proof obligations.

In this paper we describe a method called *coalescing* that handles many FOML proof obligations by soundly abstracting them into formulas of either FOL or propositional modal logic (ML). The resulting formulas are dealt with by existing theorem provers for these logics. Although the basic idea of coalescing is simple, some care has to be taken in the presence of equality and bound variables. The translation becomes trickier in the presence of defined operators.

*This work has been partially funded by the Microsoft Research-Inria Joint Centre, France. It has also been supported by the European Union Seventh Framework Programme under grant agreement no. 295261 (MEALS) and by the French BGLE Project ADN4SE.

Outline of this Paper. Section 2 motivates our proposal by its application within the TLA$^+$ Proof System TLAPS where coalescing can be complete over a fragment of proofs involving temporal logic. Section 3 formally introduces FOML and its two fragments, FOL and ML. Sections 4 and 5 present coalescing for modal and first-order expressions respectively, proving their soundness. Section 6 extends the results to languages containing local definitions. In Section 7 we prove the completeness of coalescing for proving safety properties. Section 8 discusses semantic translation vs. coalescing and suggests some optimizations and future work.

2 A Motivating Example

Our motivation comes from designing the TLAPS proof system [4] for TLA$^+$, which can check correctness proofs of complex, real-world algorithms [9]. The essence of TLA proofs is to decompose proofs of temporal logic formulas so that most of the obligations contain no modal operator except *prime*. Figure 1 contains the outline of the proof of a simple safety property in TLAPS that illustrates this decomposition. The system specification is formula *Spec*, defined to equal $\mathit{Init} \wedge \Box[\mathit{Step}]_v$. In this formula, *Init* is a *state predicate* that describes the possible initial states, and *Step* is an *action predicate* that describes possible state transitions. Syntactically: *Init* is a FOL formula containing state (a.k.a. flexible) variables; *Step* is a formula containing state variables, FOL operators, and the *prime* operator; and v is a tuple of all state variables in the specification. The formula $[\mathit{Step}]_v$ is a shorthand for $\mathit{Step} \vee (v' = v)$, and $\Box$ is the usual "always" operator of temporal logic. We wish to prove that a state formula $\mathit{Safe}(p)$ is true throughout any behavior described by *Spec*, for every process $p \in \mathit{Proc}$. The definitions of these formulas, and the reason for writing $\Box[\mathit{Step}]_v$ instead of $\Box\mathit{Step}$, are irrelevant for understanding the proof.

The right-hand side of Figure 1 shows the assertion and proof of the theorem. The first step in the proof is purely first-order: it introduces a fresh constant p, assumes $p \in \mathit{Proc}$, and reduces the overall proof to showing the implication $\mathit{Spec} \Rightarrow \Box\mathit{Safe}(p)$. Step $\langle 1\rangle 1$ asserts that the initial condition implies $\mathit{Safe}(p)$. This formula does not contain any modal operators. Step $\langle 1\rangle 2$ shows that $\mathit{Safe}(p)$ is preserved by every transition (as specified by $[\mathit{Step}]_v$). The proof of this step is essentially first-order, although TLAPS must handle the *prime* modality. The basic idea is to distribute primes inward in expressions using rules such as $(x + y)' = x' + y'$, and then to replace the remaining primed expressions by new atoms. For this example, we are assuming that the specification is so simple that, after the definitions of *Init*, *Next*, v, and *Safe* have been expanded, the FOL proof obligations generated for these two steps can be discharged

$\mathit{Init} \triangleq \ldots$	THEOREM $\mathit{Spec} \Rightarrow \forall\, p \in \mathit{Proc} : \Box\mathit{Safe}(p)$
$\mathit{Step} \triangleq \ldots$	$\langle 1\rangle$. SUFFICES ASSUME NEW $p \in \mathit{Proc}$
$v \triangleq \ldots$	PROVE $\mathit{Spec} \Rightarrow \Box\mathit{Safe}(p)$
$\mathit{Spec} \triangleq \mathit{Init} \wedge \Box[\mathit{Step}]_v$	OBVIOUS
$\mathit{Safe}(p) \triangleq \ldots$	$\langle 1\rangle 1$. $\mathit{Init} \Rightarrow \mathit{Safe}(p)$
	BY DEF *Init*, *Safe*
	$\langle 1\rangle 2$. $\mathit{Safe}(p) \wedge [\mathit{Step}]_v \Rightarrow \mathit{Safe}(p)'$
	BY DEF *Safe*, *Step*, v
	$\langle 1\rangle 3$. QED
	BY $\langle 1\rangle 1$, $\langle 1\rangle 2$, PTL DEF *Spec*

Figure 1: Proof of a safety property in TLAPS.

by a theorem prover.

Step $\langle 1 \rangle 3$ concludes the proof. It is justified by propositional temporal reasoning, in particular the principle

$$\frac{P \wedge A \Rightarrow P'}{P \wedge \Box A \Rightarrow \Box P}$$

The *PTL* in the step's proof tells TLAPS to invoke a PTL decision procedure, which it does after replacing *Spec* by its definition and the formulas *Init*, *Safe*(p) and $[Next]_v$ by fresh atoms. This effectively hides all operators other than those of propositional logic, $\Box$, and *prime*.

We call the process of replacing expressions by atoms *coalescing*. It is similar to the introduction of names for subformulas that theorem provers apply during pre-processing steps such as CNF transformation. However, it has a different purpose: the fresh names hide complex formulas that are meaningless to a proof backend for a fragment of the original logic. As explained in the example above, TLAPS uses coalescing in its translations to invoke FOL and PTL backend provers, where the first do not support the modal operators $\Box$ and *prime*, and the second do not support first-order constructs such as quantification, equality or terms.

Coalescing cannot in itself be semantically complete because it cannot support proof steps that rely on the interplay of the sublogics. For example, separate FOL and PTL provers cannot prove rules that distribute quantifiers over temporal modalities. Similarly, proofs of liveness properties via well-founded orderings essentially mix quantification and temporal logic. However, we need very few such proof steps in actual proofs, and we can handle them using a more traditional backend that relies on a FOL translation of temporal modalities. Coalescing is complete for a subclass of temporal logic properties that includes safety properties, which can be established by propositional temporal logic from action-level hypotheses. For these applications, we have found coalescing to be more flexible and more powerful in practice than a more traditional FOL translation. In particular, proofs need not follow the simple schema of the proof shown in Figure 1 but can invoke auxiliary invariants or lemmas. The inductive reasoning underlying much of temporal logic is embedded in PTL decision procedures but would be difficult to automate in a FOL prover. On the other hand, the *prime* modality by itself is simple enough so that it can be handled by a pre-processing step applied before passing the proof obligation to a FOL prover.

We believe that translation by coalescing will be useful for proofs in modal logics other than TLA$^+$. We therefore present its fundamental principles here using a simpler FOML containing a single modal operator $\Box$. Corresponding to the translations we have implemented in TLAPS, we give two translations of FOML obligations, one into FOL and the other into ML, and we prove their soundness.

The idea underlying coalescing is very simple: abstract away a class of operators by introducing a fresh atom in place of a subformula whose principal operator is in that class. However, doing this in a sound way in the presence of equality is not trivial because of the *Leibniz principle*, which asserts $(d = e) \Rightarrow (P(d) = P(e))$ for any expressions d and e and operator P. The Leibniz principle is valid in FOL but not FOML, which makes translating from FOML obligations to FOL obligations tricky [5].

For example, the formula $(v = 0) \Rightarrow \Box(v = 0)$ is not valid in TLA$^+$ or more generally in FOML when v is flexible. A naive application of standard FOL provers could propagate the equality in the antecedent by substituting 0 for v throughout this formula, effectively applying the instance $((v = 0) = \text{TRUE}) \Rightarrow (\Box(v = 0) = \Box\text{TRUE})$ of the Leibniz principle, and consequently prove the formula using the axiom $\Box\text{TRUE}$. Such an approach is clearly unsound. The standard translation of FOML into predicate logic [10] avoids this problem by making explicit the states at which formulas are evaluated, but at the price of adding complexity to the formula. Moreover, one typically assumes specific properties about the accessibility relation(s) underlying modal logics. Incorporating these into first-order reasoning may not be easy. For example, in TLA$^+$ the $\Box$ modality corresponds to the transitive closure of the *prime* modality, and this is not first-order axiomatizable. Of course, whether this is an issue or not depends on the particular modal logic one is interested in: semantic translation works very well in applications such as [2] that are based on a modal logic whose frame conditions are first-order axiomatizable.

Our approach is to coalesce expressions and formulas that are outside the scope of a given theorem prover. For the example above, coalescing to FOL yields $(v = 0) \Rightarrow \boxed{\Box(v = 0)}$ where $\boxed{\Box(v = 0)}$ is a new 0-ary predicate symbol, and this formula is clearly not provable. Similarly, coalescing to ML yields $\boxed{v = 0} \Rightarrow \Box\boxed{v = 0}$ of propositional modal logic, and again, this formula is not provable. We give a detailed description of how to derive a new symbol $\boxed{exp}$ for an arbitrary expression *exp*. Care has to be taken when the coalesced expression contains bound variables. For example, a naive coalescing of the expression $\{a, a\}$ in the valid formula $\forall\, a : \{a, a\} = \{a\}$ would yield $\forall\, a : \boxed{\{a, a\}} = \{a\}$, from which we can deduce $\boxed{\{a, a\}} = \{1\}$ and $\boxed{\{a, a\}} = \{2\}$, proving $1 = 2$. A correct coalescing yields $\forall\, a : \boxed{\{a, a\}}(a) = \{a\}$.

Operator Definitions. Coalescing is trickier for a language with operator definitions like $P(x, y) \triangleq exp$, where *exp* does not contain free variables other than x and y. Definitions are necessary for structuring specifications and for managing the complexity of proofs through lemmas about the defined operators. We therefore do not want to systematically expand all defined operators in order to obtain formulas of basic FOML. The Leibniz principle may not hold for an expression $P(a, b)$ if the operator P is defined in terms of modal operators—that is, $(a = c) \land (b = d)$ need not imply $P(a, b) = P(c, d)$. It would therefore be unsound to encode P as an uninterpreted predicate symbol in first-order logic. We show how soundness is preserved by replacing an expression $P(a, b)$ with $\boxed{P, \epsilon_1, \epsilon_2}(a, b)$, for suitable expressions ϵ_1 and ϵ_2, where $\boxed{P, \epsilon_1, \epsilon_2}$ can be defined so it satisfies the Leibniz principle and also satisfies

$$\boxed{P, \epsilon_1, \epsilon_2}(a, b) = P(a, b)$$

in suitably extended models of FOML, ensuring equisatisfiability of the original and the coalesced formula. Since it satisfies the Leibniz principle, $\boxed{P, \epsilon_1, \epsilon_2}$ can be taken to be an uninterpreted predicate symbol by a first-order theorem prover. Our construction extends to the case of definitions of second-order operators, which are allowed in TLA$^+$.

3 First-Order Modal Logic

3.1 Syntax.

We introduce a language of first-order modal logic whose modal operator we denote by ∇ in order to avoid confusion with the $\Box$ of TLA$^+$. The language omits the customary distinction between function and predicate symbols, and hence between terms and formulas. This simplifies notation and allows our results to apply to TLA$^+$ as well as to a conventional language that does distinguish terms and formulas—the conventional language just having a smaller set of legal formulas.

We assume a first-order signature consisting of non-empty distinct denumerable sets $\mathcal{X}$ of rigid variables, $\mathcal{V}$ of flexible variables, and $\mathcal{O}$ of operator symbols. Operator symbols have arities in $\mathbb{N}$ and generalize both function and predicate symbols. Expressions e of FOML are then inductively defined by the following grammar:

$$e \quad ::= \quad x \mid v \mid op(e, \ldots, e) \mid e = e \mid \textsc{false} \mid e \Rightarrow e \mid \forall\, x : e \mid \nabla e$$

where $x \in \mathcal{X}$, $v \in \mathcal{V}$, $op \in \mathcal{O}$, and arities are respected (empty parentheses are omitted for 0-ary symbols). We do not allow quantification over flexible variables, so our flexible variables are really "flexible function symbols of arity 0". While TLA$^+$ allows quantification over flexible variables, it can be considered as another modal operator for the purposes of coalescing.

The notions of free and bound (rigid) variables are the usual ones. We say that an expression is *rigid* iff it contains neither flexible variables nor subexpressions of the form ∇e. The standard propositional ($\textsc{true}$, $\neg$, $\wedge$, $\vee$, $\equiv$) and first-order ($\exists$) connectives are defined in the usual way. The dual modality Δ is introduced by defining Δe as $\neg\nabla\neg e$. The extension to a multi-modal language is straightforward.

3.2 Semantics.

A *Kripke model* $\mathcal{M}$ for FOML is a 6-tuple $(\mathcal{I}, \xi, \mathcal{W}, R, \zeta, \nabla_{\mathcal{M}})$, where:

- $\mathcal{I}$ is a standard first-order interpretation consisting of a universe $|\mathcal{I}|$ and, for every operator symbol op, an interpretation $\mathcal{I}(op) : |\mathcal{I}|^n \to |\mathcal{I}|$ where n agrees with the arity of op. We assume that the universe $|\mathcal{I}|$ contains two distinguished, distinct values tt and ff.
- $\xi : \mathcal{X} \to |\mathcal{I}|$ is a valuation of the rigid variables.
- $\mathcal{W}$ is a non-empty set of states, and $R \subseteq \mathcal{W} \times \mathcal{W}$ is the accessibility relation.
- $\zeta : \mathcal{V} \times \mathcal{W} \to |\mathcal{I}|$ is a valuation of the flexible variables at the different states of the model.
- $\nabla_{\mathcal{M}} : 2^{|\mathcal{I}|} \to |\mathcal{I}|$ is a function such that $\nabla_{\mathcal{M}}(S) = \mathsf{tt}$ iff $S \subseteq \{\mathsf{tt}\}$.

Note that we assume a constant universe, independent of the states of the model, and we also assume that all operators in $\mathcal{O}$ are rigid—i.e., interpreted independently of the states.

We inductively define the interpretations of expressions $[\![e]\!]_w^{\mathcal{M}}$ at state w of model $\mathcal{M}$. When the model $\mathcal{M}$ is understood from the context, we drop it from the notation.

- $[\![x]\!]_w^{\mathcal{M}} \;=_{\mathsf{def}}\; \xi(x) \quad$ for $x \in \mathcal{X}$
- $[\![v]\!]_w^{\mathcal{M}} \;=_{\mathsf{def}}\; \zeta(v, w) \quad$ for $v \in \mathcal{V}$
- $[\![op(e_1, \ldots, e_n)]\!]_w^{\mathcal{M}} \;=_{\mathsf{def}}\; \mathcal{I}(op)([\![e_1]\!]_w^{\mathcal{M}}, \ldots, [\![e_n]\!]_w^{\mathcal{M}}) \quad$ for $op \in \mathcal{O}$
- $[\![e_1 = e_2]\!]_w^{\mathcal{M}} \;=_{\mathsf{def}}\; \begin{cases} \mathsf{tt} & \text{if } [\![e_1]\!]_w^{\mathcal{M}} = [\![e_2]\!]_w^{\mathcal{M}} \\ \mathsf{ff} & \text{otherwise} \end{cases}$
- $[\![\textsc{false}]\!]_w^{\mathcal{M}} \;=_{\mathsf{def}}\; \mathsf{ff}$
- $[\![\varphi \Rightarrow \psi]\!]_w^{\mathcal{M}} \;=_{\mathsf{def}}\; \begin{cases} \mathsf{tt} & \text{if } [\![\varphi]\!]_w^{\mathcal{M}} \neq \mathsf{tt} \text{ or } [\![\psi]\!]_w^{\mathcal{M}} = \mathsf{tt} \\ \mathsf{ff} & \text{otherwise} \end{cases}$
- $[\![\forall\, x : \varphi]\!]_w^{\mathcal{M}} \;=_{\mathsf{def}}\; \begin{cases} \mathsf{tt} & \text{if } [\![\varphi]\!]_w^{\mathcal{M}'} = \mathsf{tt} \text{ for all } \mathcal{M}' = (\mathcal{I}, \xi', \mathcal{W}, R, \zeta) \text{ such that} \\ & \quad \xi'(y) = \xi(y) \text{ for all } y \in \mathcal{X} \text{ different from } x \\ \mathsf{ff} & \text{otherwise} \end{cases}$
- $[\![\nabla\varphi]\!]_w^{\mathcal{M}} \;=_{\mathsf{def}}\; \nabla_{\mathcal{M}}(\{[\![\varphi]\!]_{w'}^{\mathcal{M}} : (w, w') \in R\})$

We write $\mathcal{M}, w \models \varphi$ instead of $[\![\varphi]\!]^{\mathcal{M}}_w = \mathsf{tt}$. We say that φ is *valid* iff $\mathcal{M}, w \models \varphi$ holds for all $\mathcal{M}$ and w, and that it is *satisfiable* iff $\mathcal{M}, w \models \varphi$ for some $\mathcal{M}$ and w. We define a consequence relation $\models$ as follows (where Γ is a set of formulas): $\Gamma \models \varphi$ iff for all $\mathcal{M}$, if $\mathcal{M}, w \models \psi$ for all $\psi \in \Gamma$ and $w \in \mathcal{W}$, then $\mathcal{M}, w \models \varphi$ for all $w \in \mathcal{W}$.

Our definition of the semantics is a straightforward extension of the standard Kripke semantics to our setting, where ∇e need not denote a truth value. The condition on the function $\nabla_{\mathcal{M}}$ used for interpreting the modality ensures that $\mathcal{M}, w \models \nabla\varphi$ iff $\mathcal{M}, w' \models \varphi$ for all w' such that $(w, w') \in R$ as in the standard Kripke semantics. Because we assume a constant domain of interpretation, both Barcan formulas are valid—that is, we have validity of

$$(\forall x : \nabla\varphi) \;\equiv\; \nabla(\forall x : \varphi). \tag{1}$$

Moreover, since all operator symbols have rigid interpretations, it is easy to prove by induction on the complexity of expressions that $[\![e]\!]_w = [\![e]\!]_{w'}$ holds for all states w, w' whenever e is a rigid expression. It follows that implications of the form $\varphi \Rightarrow \nabla\varphi$ are valid for rigid φ—for example:

$$\forall\, x, y : (x = y) \Rightarrow \nabla(x = y). \tag{2}$$

3.3 FOL and ML fragments of FOML

Two natural sublogics of FOML are first-order logic (FOL) and propositional modal logic (ML).

FOL does not have flexible variables $\mathcal{V}$ or expressions ∇e. A first-order structure $(\mathcal{I}, \xi)$ consists of an interpretation $\mathcal{I}$ as above and a valuation ξ of the (rigid) variables. The inductive definition of the semantics consists of the relevant clauses of the one given above for FOML, and the notions of first-order validity $\models_{FOL} \varphi$, satisfiability, and consequence carry over in the usual way.

ML does not have rigid variables, quantifiers, operator symbols or equality. A (propositional) Kripke model for ML is given as $\mathcal{K} = (\mathcal{W}, R, \zeta)$ where the set of states $\mathcal{W}$ and the accessibility relation R are as for FOML, and the valuation $\zeta : \mathcal{V} \times \mathcal{W} \to \{\mathsf{tt}, \mathsf{ff}\}$ assigns truth values to flexible variables at every state. The inductive definition of $[\![e]\!]^{\mathcal{K}}_w \in \{\mathsf{tt}, \mathsf{ff}\}$ specializes to the following clauses:

- $[\![v]\!]^{\mathcal{M}}_w \;=\; \zeta(v, w)$ for $v \in \mathcal{V}$
- $[\![\mathrm{FALSE}]\!]^{\mathcal{M}}_w \;=\; \mathsf{ff}$
- $[\![\varphi \Rightarrow \psi]\!]^{\mathcal{M}}_w \;=\; \mathsf{tt}$ iff $[\![\varphi]\!]^{\mathcal{M}}_w = \mathsf{ff}$ or $[\![\psi]\!]^{\mathcal{M}}_w = \mathsf{tt}$
- $[\![\nabla\varphi]\!]^{\mathcal{M}}_w \;=\; \mathsf{tt}$ iff $[\![\varphi]\!]^{\mathcal{M}}_{w'} = \mathsf{tt}$ for all $w' \in W$ such that $(w, w') \in R$

The notions of validity $\models_{ML} \varphi$, satisfiability, and consequence carry over as usual.

4 Coalescing Modal Expressions

4.1 Definition of the abstraction e_{FOL}

One of our objectives is to apply standard first-order theorem provers for proving theorems of FOML that are instances of first-order reasoning. Since the operator ∇ is not available in first-order logic, we must translate FOML formulas ψ to purely first-order formulas ψ_{FOL} such that the consequence $\Gamma_{FOL} \models_{FOL} \varphi_{FOL}$ entails $\Gamma \models \varphi$. A naive but unsound approach would be to replace the modal operator ∇ by a fresh monadic operator symbol *Nec*. As explained in Section 2, this approach would allow one to prove the invalid formula $(v = 0) \Rightarrow \nabla(v = 0)$.

(The formula is not valid because it is false at a state w of a model in which $\zeta(v, w) = \mathcal{I}(0)$, but $\zeta(v, w') \neq \mathcal{I}(0)$ for some state w' accessible from w.) As we observed, a sound approach is to define φ_{FOL} by using the well-known standard translation from modal logic to first-order logic [3, 10] that makes explicit the FOML semantics. However, that translation introduces additional complexity—complexity that is unnecessary for proof obligations that follow from ordinary first-order reasoning.

Instead, we define φ_{FOL} to be a syntactic first-order abstraction of φ in which modal subexpressions are coalesced—that is, replaced by fresh operators. If φ is $(v = 0) \Rightarrow \nabla(v = 0)$, then φ_{FOL} is $(v = 0) \Rightarrow \boxed{\nabla(v = 0)}$, where $\boxed{\nabla(v = 0)}$ is a new 0-ary operator symbol. The variable v is considered a free variable in φ_{FOL}.

We want to ensure that subexpressions appearing more than once are abstracted by the same operators, allowing for instances of first-order theorems to remain valid. This requires some care for expressions that contain bound variables. For example, we expect to prove

$$(\exists\, x, z : \nabla(v = x)) \equiv (\exists\, y : \nabla(v = y)) \tag{3}$$

The fresh operator symbols $\boxed{\nabla e}$ are therefore defined as λ-abstractions over the bound variables occurring in e, and these are identified modulo α-equivalence. Formally, we let $e_{FOL} = e^{\varepsilon}_{FOL}$ where, for a list $\vec{y}$ of rigid variables, the first-order expression $e^{\vec{y}}_{FOL}$ over the extended set of variables $\mathcal{X} \cup \mathcal{V}$ is defined inductively as follows.

- $x^{\vec{y}}_{FOL} \;=_{\mathsf{def}}\; x$ for $x \in \mathcal{X}$ a rigid variable,
- $v^{\vec{y}}_{FOL} \;=_{\mathsf{def}}\; v$ for $v \in \mathcal{V}$ a flexible variable,
- $(op(e_1, \ldots, e_n))^{\vec{y}}_{FOL} \;=_{\mathsf{def}}\; op((e_1)^{\vec{y}}_{FOL}, \ldots, (e_n)^{\vec{y}}_{FOL})$ for $op \in \mathcal{O}$,
- $(e_1 = e_2)^{\vec{y}}_{FOL} \;=_{\mathsf{def}}\; (e_1)^{\vec{y}}_{FOL} = (e_2)^{\vec{y}}_{FOL}$,
- $\text{FALSE}^{\vec{y}}_{FOL} \;=_{\mathsf{def}}\; \text{FALSE}$
- $(e_1 \Rightarrow e_2)^{\vec{y}}_{FOL} \;=_{\mathsf{def}}\; (e_1)^{\vec{y}}_{FOL} \Rightarrow (e_1)^{\vec{y}}_{FOL}$,
- $(\forall\, x : e)^{\vec{y}}_{FOL} \;=_{\mathsf{def}}\; \forall\, x : e^{x,\vec{y}}_{FOL}$,
- $(\nabla e)^{\vec{y}}_{FOL} \;=_{\mathsf{def}}\; \boxed{\lambda \vec{z} : \nabla e}(\vec{z})$ where $\vec{z}$ is the subsequence of rigid variables in $\vec{y}$ that appear free in e. (If z is the empty sequence, this is simply $\boxed{\nabla e}$.)

With these definitions, the formula (3) is coalesced as

$$(\exists\, x, z : \boxed{\lambda x : \nabla(v = x)}(x)) \;\equiv\; (\exists\, y : \boxed{\lambda y : \nabla(v = y)}(y)) \tag{4}$$

which is an instance of the valid first-order equivalence

$$(\exists\, x, z : P(x)) \;\equiv\; (\exists\, y : P(y))$$

In particular, the two operator symbols occurring in (4) are identified because the two λ-expressions are α-equivalent. Identification of coalesced formulas modulo α-equivalence ensures that the translation is insensitive to the names of bound (rigid) variables. Section 8 discusses techniques for abstracting from less superficial differences in first-order expressions, such as between $\lambda x, y$ and $\lambda y, x$ and between $a = b$ and $b = a$.

4.2 Soundness of coalescing to FOL

For a set Γ of FOML formulas, we denote by Γ_{FOL} the set of all formulas ψ_{FOL}, for $\psi \in \Gamma$. We now show the soundness of the abstraction.

Theorem 1. *For any set Γ of FOML formulas and any FOML formula φ, if $\Gamma_{FOL} \models_{FOL} \varphi_{FOL}$ then $\Gamma \models \varphi$.*

Proof (sketch). Assume that $\Gamma \not\models \varphi$, so $\mathcal{M} = (\mathcal{I}, \xi, \mathcal{W}, R, \zeta, \nabla_{\mathcal{M}})$ is a Kripke model such that $\mathcal{M}, w' \models \psi$ for all $\psi \in \Gamma$ and $w' \in \mathcal{W}$, but that $\mathcal{M}, w \not\models \varphi$ for some $w \in \mathcal{W}$.

For the extended set of variables $\mathcal{X} \cup \mathcal{V}$, define the first-order structure $\mathcal{S} = (\mathcal{I}', \xi')$ where $\mathcal{I}'$ agrees with $\mathcal{I}$ for all operator symbols that appear in Γ or φ, and where the valuation ξ' is defined by $\xi'(x) = \xi(x)$ for $x \in \mathcal{X}$ and $\xi'(v) = \zeta(w, v)$ for $v \in \mathcal{V}$. For the additional operator symbols introduced in Γ_{FOL} and φ_{FOL}, we define

$$\mathcal{I}'(\boxed{\lambda \vec{z} : \nabla e})(d_1, \ldots, d_n) \quad = \quad [\![\nabla e]\!]_w^{\mathcal{M}'}$$

where $\mathcal{M}'$ agrees with $\mathcal{M}$ except for the valuation ξ' that assigns the i^{th} variable of $\vec{z}$ to d_i. This interpretation is well-defined: if ∇e_1 and ∇e_2 are two expressions in Γ or φ that give rise to the same operator symbol, then $(\lambda \vec{z}_1 : \nabla e_1)$ and $(\lambda \vec{z}_2 : \nabla e_2)$ must be α-equivalent, and therefore $\mathcal{I}'(\boxed{\lambda \vec{z}_1 : \nabla e_1})(d_1, \ldots, d_n) = \mathcal{I}'(\boxed{\lambda \vec{z}_2 : \nabla e_2})(d_1, \ldots, d_n)$.

It is straightforward to prove that $[\![e_{FOL}]\!]^{\mathcal{S}} = [\![e]\!]_w^{\mathcal{M}}$ holds for all expressions e_{FOL} that appear in Γ_{FOL} or φ_{FOL}. In particular, it follows that $\mathcal{S} \models_{FOL} \psi_{FOL}$ for all $\psi \in \Gamma$ and $\mathcal{S} \not\models_{FOL} \varphi_{FOL}$. This shows that $\Gamma_{FOL} \not\models_{FOL} \varphi_{FOL}$ and concludes the proof. Q.E.D.

5 Coalescing First-Order Expressions

We now define an abstraction φ_{ML} of FOML formulas to formulas of propositional modal logic. Again, we require for soundness that $\Gamma \models \varphi$ whenever $\Gamma_{ML} \models_{ML} \varphi_{ML}$—that is, consequence between abstracted formulas implies consequence between the original ones. In this way, we can use theorem provers for propositional modal logic to carry out FOML proofs that are instances of propositional modal reasoning. The abstraction φ_{ML} replaces all first-order subexpressions e of φ by new (propositional) flexible variables $\boxed{e}$, where variables $\boxed{\forall x : e}$ are once again identified modulo α-equivalence. Formally, the translation is defined as follows.

- $x_{ML} \ =_{\mathsf{def}} \ \boxed{x}$ for $x \in \mathcal{X}$ a rigid variable,
- $v_{ML} \ =_{\mathsf{def}} \ v$ for $v \in \mathcal{V}$ a flexible variable,
- $(op(t_1, \ldots, t_n))_{ML} \ =_{\mathsf{def}} \ \boxed{op(t_1, \ldots, t_n)}$ for $op \in \mathcal{O}$,
- $(e_1 = e_2)_{ML} \ =_{\mathsf{def}} \ \boxed{e_1 = e_2}$,
- $\text{FALSE}_{ML} \ =_{\mathsf{def}} \ \text{FALSE}$,
- $(e_1 \Rightarrow e_2)_{ML} \ =_{\mathsf{def}} \ (e_1)_{ML} \Rightarrow (e_2)_{ML}$,
- $(\forall x : e)_{ML} \ =_{\mathsf{def}} \ \boxed{\forall x : e}$,
- $(\nabla e)_{ML} \ =_{\mathsf{def}} \ \nabla e_{ML}$.

As an example, coalescing the formula

$$(x = y) \land \nabla \Delta \text{TRUE} \Rightarrow \nabla \Delta (x = y)$$

yields the ML-formula

$$\boxed{x = y} \ \wedge\ \nabla\Delta\text{TRUE} \ \Rightarrow\ \nabla\Delta\boxed{x = y} \tag{5}$$

The implication (5) is not ML-valid. However, for rigid variables x and y, it follows from the hypothesis $\boxed{x = y} \Rightarrow \nabla\boxed{x = y}$, which is justified by the FOML law (2).

For a set Γ of FOML formulas, we denote by Γ_{ML} the set of modal abstractions ψ_{ML}, for all $\psi \in \Gamma$. Moreover, we define the set $\mathcal{H}(\Gamma)$ to consist of all formulas of the form $\boxed{e} \Rightarrow \nabla\boxed{e}$, for all flexible variables $\boxed{e}$ introduced in Γ_{ML} that correspond to rigid expressions e in Γ.

Theorem 2. *Assume that Γ is a set of FOML formulas and that φ is a FOML formula. If $\Gamma_{ML}, \mathcal{H}(\Gamma \cup \{\varphi\}) \models_{ML} \varphi_{ML}$ then $\Gamma \models \varphi$.*

Proof (sketch). As in Theorem 1, we prove the contra-positive. Assume that $\mathcal{M} = (\mathcal{I}, \xi, \mathcal{W}, R, \zeta, \nabla_{\mathcal{M}})$ is a Kripke model such that $\mathcal{M}, w' \models \psi$ for all $\psi \in \Gamma$ and $w' \in \mathcal{W}$, but $\mathcal{M}, w \not\models \varphi$ for a certain $w \in \mathcal{W}$.

Define the propositional Kripke model $\mathcal{K} = (\mathcal{W}, R, \zeta')$ where ζ' assigns truth values in $\{\mathsf{tt}, \mathsf{ff}\}$ to all states $w' \in \mathcal{W}$ and flexible variables in Γ_{ML} or φ_{ML}:

$$\begin{aligned}
&\zeta'(w', v) = \mathsf{tt} \ \text{ iff } \ \zeta(w', v) = \mathsf{tt} \ \text{ for } v \in \mathcal{V}\\
&\zeta'(w', \boxed{x}) = \mathsf{tt} \ \text{ iff } \ \xi(x) = \mathsf{tt} \ \text{ for } x \in \mathcal{X}\\
&\zeta'(w', \boxed{op(t_1, \ldots, t_n)}) = \mathsf{tt} \ \text{ iff } \ [\![op(t_1, \ldots, t_n)]\!]^{\mathcal{M}}_{w'} = \mathsf{tt}\\
&\zeta'(w', \boxed{e_1 = e_2}) = \mathsf{tt} \ \text{ iff } \ [\![e_1]\!]^{\mathcal{M}}_{w'} = [\![e_2]\!]^{\mathcal{M}}_{w'}\\
&\zeta'(w', \boxed{\forall x : e}) = \mathsf{tt} \ \text{ iff } \ \mathcal{M}, w' \models \forall x : e
\end{aligned}$$

Again, ζ' is well-defined. It is easy to prove, for all $w' \in \mathcal{W}$ and all e such that e_{ML} appears in Γ_{ML} or φ_{ML}, that $\mathcal{K}, w' \models e_{ML}$ iff $[\![e]\!]^{\mathcal{M}}_{w'} = \mathsf{tt}$. In particular, it follows that $\mathcal{K}, w' \models \psi_{ML}$ for all $\psi \in \Gamma$ and that $\mathcal{K}, w \not\models_{ML} \varphi_{ML}$.

Furthermore, the definition of $\mathcal{K}$ ensures that $\mathcal{K}, w' \models \psi$ holds for all $\psi \in \mathcal{H}(\Gamma \cup \{\varphi\})$ because $[\![e]\!]^{\mathcal{M}}_{w'} = [\![e]\!]^{\mathcal{M}}_{w''}$ holds for all rigid expressions e and all states $w', w'' \in \mathcal{W}$.

In summary, it follows that $\Gamma_{ML}, \mathcal{H}(\Gamma \cup \{\varphi\}) \not\models_{ML} \varphi_{ML}$, which concludes the proof. Q.E.D.

6 Coalescing in the presence of operator definitions

6.1 Operator definitions

We now extend our language to allow definitions of the form

$$d(x_1, \ldots, x_n) \ \triangleq\ e$$

where d is a fresh symbol, $x_1, \ldots, x_n$ are pairwise distinct rigid variables, and e is an expression whose free rigid variables are among $x_1, \ldots, x_n$.

For an operator d defined as above and expressions $e_1, \ldots, e_n$, the application $d(e_1, \ldots, e_n)$ is a well-formed expression whose semantics is given by:

$$[\![d(e_1, \ldots, e_n)]\!]^{\mathcal{M}}_{w} = [\![e[e_1/x_1, \ldots, e_n/x_n]]\!]^{\mathcal{M}}_{w}$$

In other words, the defining expression is evaluated when the arguments have been substituted for the variables. However, when reasoning about expressions containing defined operators, one does not wish to systematically expand definitions. If the precise definition is unimportant, it is better to leave the operator unexpanded in order to keep the formulas small. We now extend the coalescing techniques introduced in the preceding sections to handle expressions that may contain defined operators.

It is easy to see that the algorithm introduced in Section 5 for abstracting first-order subexpressions remains sound if we handle defined operators like operators in $\mathcal{O}$. In particular, two expressions $d(\vec{e}_1)$ and $d(\vec{e}_2)$ are abstracted by the same flexible variable only if they are syntactically equal up to α-equivalence. However, this simple approach does not work for the algorithm of Section 4 that abstracts modal subexpressions. As an example, consider the definition

$$cst(x) \;\triangleq\; \exists\, y : \nabla(x = y) \tag{6}$$

and the formula

$$(v = w) \;\Rightarrow\; (cst(v) \equiv cst(w)) \tag{7}$$

where v and w are flexible variables. An expression e satisfies $cst(e)$ at state w iff the value of e is the same at all reachable states w'. Hence, formula (7) is obviously not valid. If cst were treated like an operator in $\mathcal{O}$, the algorithm of Section 4 leaves (7) unchanged. However, v and w would be considered ordinary (rigid) variables and cst would be considered an uninterpreted operator symbol, so (7), seen as a FOL formula, would be provable. Thus, it would be unsound to simply treat defined operators like operators in $\mathcal{O}$ in our algorithm for coalescing modal subexpressions.

6.2 Rigid arguments and Leibniz positions

The example above shows that in the presence of definitions, FOML formulas without any visible modal operators may violate the Leibniz principle that substituting equals for equals should yield equal results. However, a first observation shows that the Leibniz principle still holds for rigid arguments.

Lemma 3. *For any defined n-ary operator d, expressions $e_1, \ldots, e_n$ with e_i rigid (for some $i \in 1..n$), Kripke model $\mathcal{M}$, state w, and rigid variable x that does not occur free in any e_j, we have*

$$[\![d(e_1, \ldots, e_n)]\!]^{\mathcal{M}}_w = [\![d(e_1, \ldots, e_{i-1}, x, e_{i+1}, \ldots, e_n)]\!]^{\mathcal{M}'}_w$$

where $\mathcal{M}'$ agrees with $\mathcal{M}$ except for the valuation ξ' of rigid variables, which is like ξ but assigns x to $[\![e_i]\!]^{\mathcal{M}}_w$.

Proof (sketch). Since e_i is rigid, the value of $[\![e_i]\!]^{\mathcal{M}}_{w'}$, for any $w' \in W$, is independent of the state w'. The assertion is then proved by induction on the defining expression for operator d.

Q.E.D.

For a non-rigid argument of a defined operator, the Leibniz principle is preserved when the argument does not appear in a modal context in the defining expression. We inductively define which argument positions of an FOML operator or connective are Leibniz (satisfy the Leibniz principle).

Definition 4 (Leibniz argument positions)**.**

- *All argument positions of the operators in $\mathcal{O}$ and of all FOML connectives except ∇ are Leibniz. The single argument position of ∇ is not Leibniz.*
- *For an operator defined by $d(x_1, \ldots, x_n) \triangleq e$, the i^{th} argument position of d is Leibniz iff x_i does not occur within a non-Leibniz argument position in e.*

In other words, the i^{th} argument position of a defined operator is Leibniz iff the i^{th} parameter does not appear in the scope of any occurrence of ∇ in the full expansion of the defining expression.

Lemma 5. *Assume that d is a defined n-ary operator whose i^{th} argument position is Leibniz. For any expressions $e_1, \ldots, e_n$, $i \in 1..n$, Kripke model $\mathcal{M}$, state w and rigid variable x that does not occur free in any e_i, we have*

$$[\![d(e_1, \ldots, e_n)]\!]_w^{\mathcal{M}} = [\![d(e_1, \ldots, e_{i-1}, x, e_{i+1}, \ldots, e_n)]\!]_w^{\mathcal{M}'}$$

where $\mathcal{M}'$ agrees with $\mathcal{M}$ except for the valuation ξ' of rigid variables, which is like ξ but assigns x to $[\![e_i]\!]_w^{\mathcal{M}}$.

Proof (sketch). Induction on the syntax of the defining expression for d. Q.E.D. It follows from Lemmas 3 and 5 that the implication

$$(e_i = f) \;\Rightarrow\; (d(e_1, \ldots, e_n) = d(e_1, \ldots, e_{i-1}, f, e_{i+1}, \ldots, e_n))$$

is valid when e_i and f are rigid expressions or when the i^{th} argument position of d is Leibniz.

6.3 Coalescing for defined operators

The definition of the syntactic abstraction e_{FOL} for the extended language is now completed by defining

- $(d(e_1, \ldots, e_n))^{\vec{y}}_{FOL} \;=_{\mathsf{def}}\; \boxed{d, \epsilon_1, \ldots, \epsilon_n}((e_1)^{\vec{y}}_{FOL}, \ldots, (e_n)^{\vec{y}}_{FOL})$ for a defined n-ary operator d where

$$\begin{array}{ll} \epsilon_i = * & \text{if the } i^{\text{th}} \text{ position of } d \text{ is Leibniz or } e_i \text{ is a rigid expression,} \\ \epsilon_i = e_i & \text{otherwise.} \end{array}$$

With these definitions, the single argument position of operator cst introduced by (6) is not Leibniz, and coalescing formula (7) yields

$$(v = w) \;\Rightarrow\; (\boxed{cst, v}(v) \equiv \boxed{cst, w}(w))$$

for two distinct fresh operators $\boxed{cst, v}$ and $\boxed{cst, w}$. As expected, this formula cannot be proved. However, the formula $\forall\, x, y : (x = y) \;\Rightarrow\; (cst(x) \equiv cst(y))$ is coalesced as $\forall\, x, y : (x = y) \;\Rightarrow$ $(\boxed{cst, *}(x) \equiv \boxed{cst, *}(y))$ and is valid.

Theorem 6. *Theorem 1 remains valid for FOML formulas in the presence of defined operator symbols.*

Proof (sketch). Extending the proof of Theorem 1, we define the interpretation of the fresh operator symbols as follows:

$$\begin{array}{l} \mathcal{I}'(\boxed{d, \epsilon_1, \ldots, \epsilon_n})(d_1, \ldots, d_n) \;=\; [\![d(\alpha_1, \ldots, \alpha_n)]\!]_w^{\mathcal{M}'} \\ \text{where } \alpha_i = \left\{ \begin{array}{ll} e_i & \text{if } \epsilon_i = e_i \\ x_i & \text{if } \epsilon_i = * \end{array} \right. \end{array}$$

In this definition, w is the state fixed in the proof and $\mathcal{M}'$ agrees with $\mathcal{M}$ except for the valuation ξ' that assigns the variables x_i to d_i.

Again, one proves that $[\![e_{FOL}]\!]^{\mathcal{S}} = [\![e]\!]^{\mathcal{M}}_w$ for all expressions e_{FOL} that appear in Γ_{FOL} or φ_{FOL}. For the expressions corresponding to applications of defined operators, the proof is obvious for those arguments where $\epsilon_i = e_i$, and it makes use of Lemmas 3 and 5 when $\epsilon_i = *$. Q.E.D.

7 Proving Safety Properties by Coalescing in TLA

In Section 2 we gave an example of using coalescing to prove a safety property in TLA and we claimed that it is always possible to do so. In this section we will give an informal argument to support that claim.

We start with the definition of *safety property*: a safety property is a property that holds for every prefix of a sequence of states if and only if it holds for the whole sequence.

The standard form of a TLA specification is $\mathtt{INIT} \wedge \Box[\mathtt{NEXT}]_v$. Given a specification $\mathtt{INIT}_0 \wedge \Box[\mathtt{NEXT}_0]_{v_0}$ and a safety property P_0, we want to prove the assertion $\mathtt{INIT}_0 \wedge \Box[\mathtt{NEXT}_0]_{v_0} \Rightarrow P_0$.

The first step is to reformulate it as an *invariant assertion*, i.e. an assertion of the form $\mathtt{INIT}_1 \wedge \Box[\mathtt{NEXT}_1]_{v_1} \Rightarrow \Box P_1$ equivalent to our initial assertion, where P_1 is a state predicate.

This is done by adding to $\mathtt{INIT}_0 \wedge \Box[\mathtt{NEXT}_0]_{v_0}$ a *history variable* that records all past states. By the definition of safety properties, P_0 holds for a sequence of states if and only if it holds for every prefix of it. We construct P_1 so it is true for the history variable of the last state of a prefix if and only if P_0 is true for the prefix.

The second step is to turn the invariant P_1 into an *inductive invariant*: a state predicate P_2 such that $\mathtt{INIT}_1 \wedge \Box[\mathtt{NEXT}_1]_{v_1} \Rightarrow \Box P_1$ is a theorem if and only if the following are theorems:

1. $\mathtt{INIT} \Rightarrow P_2$
2. $\mathtt{NEXT} \wedge P_2 \Rightarrow P_2'$
3. $P_2 \Rightarrow P_1$

This invariant exists under standard assumptions on the expressiveness of the language of state predicates. The statement of these standard assumptions (which are satisfied by TLA$^+$) and the proof that this transformation is always possible is essentially the same as in [1].

Once we have the inductive invariant P_2, we can easily prove the validity of the above equivalence: by coalescing first-order expressions we get a simple ML theorem that automatic tools handle without problems. This allows us not only to apply the above *induction rule*, but also to extend our toolset to include variations of this rule (for example by splitting the invariant into several mutually-inductive formulas), and in fact arbitrary ML theorems. This eases the proving of safety properties, and also enables us to prove some (but not all) liveness properties.

It is important to note that a ML prover will have no problem proving the above FOML induction theorem with the aid of coalescing because temporal induction is built into such provers. On the other hand, a standard FOL prover would have a very hard time with a FOL translation of the FOML formula because that involves induction over the naturals.

We can summarize the results obtained so far by stating that the validity of any safety property is equivalent to the validity of the three expressions $1, 2$ and 3 above, two of which are states predicates and one is an action predicate.

We then argue that coalescing is complete for action predicates and therefore that coalescing gives a sound and complete method for proving safety properties.

Given an action predicate A, we first eliminate all defined operators by expanding their definitions. This yields an action predicate B, equivalent to A, whose operators are all built-in TLA operators.

Then we use the fact that *prime* distributes over all built-in TLA operators to push the *prime*s downward as far as possible. This yields an action predicate C, equivalent to B and A, where *prime* is only applied to flexible variables.

We then show that coalescing is complete for such action predicates, i.e. that if C is valid in TLA, then C_{FOL} must be valid in the first-order fragment of TLA. This is done by contradiction: assuming a counter-model of C_{FOL} and building a counter-model of C.

This concludes our sequence of transformations, starting from any safety property, and ending with a ML formula and a few FOL formulas such that the safety property is true if and only if this handful of formulas are all true.

We have thus shown that the two kinds of coalescing presented in this paper are sufficient for proving safety properties, whose traditional FOL translations is usually beyond the capabilities of FOL provers.

8 Conclusion

We have found that our techniques for coalescing FOML formulas to FOL and ML are useful for verifying temporal logic properties of TLA$^+$ specifications. In particular, the overwhelming majority of proof obligations that arise during TLA$^+$ proofs contain only the *prime* modal operator. For this fragment, rewriting by the valid equality $op(e_1, \ldots, e_n)' = op(e'_1, \ldots, e'_n)$, for operators $op \in \mathcal{O}$, followed by coalescing to FOL is complete. Many of the proof obligations that involve the $\Box$ modality of TLA$^+$ are instances of propositional temporal reasoning, and these can be handled by coalescing to ML and invoking a decision procedure for propositional temporal logic.

Coalescing to FOL eschews semantic translation of FOML formulas [10] in favor of replacing a subformula whose principal operator is modal by a fresh operator symbol. The resulting formulas are simpler than those obtained by semantic translation, and they can readily be understood in terms of the original FOML formulation of the problem. The price to pay is a loss of completeness. For example, the valid Barcan formula (1) cannot be proved using only our two translations. TLA proofs contain only a small number of such proof obligations, and we expect TLAPS to be able to handle them easily with a semantic translation to FOL. For applications other than TLA$^+$ theorem proving that require first-order modal reasoning, the trade-off in choosing between semantic translation and coalescing will depend upon how effective one expects semantic translation and standard first-order theorem proving to work in practice. One recent experiment [2] found this technique entirely satisfactory, but it used a modal logic too weak to handle the applications that concern us. The validity problem for the first-order temporal logic we use is Π^1_1-complete, and semantic translation cannot be expected to work satisfactorily due to the need for inductive reasoning over natural numbers. An FOL prover applied to a semantic translation would probably not be able to prove obligations that a propositional temporal logic decision procedure easily handles with our ML translation.

The definition of coalescing to FOL presented in Section 4 identifies modal subformulas such as (3) that are identical up to the names of bound rigid variables that they contain. This definition can be refined to identify formulas that differ in less superficial ways. For example, it may be desirable to reorder bound variables according to their appearance in coalesced subformulas. This would allow us to coalesce the formula

$$(\exists\, y\, \forall\, x : \Box P(x, y)) \Rightarrow (\forall\, x\, \exists\, y : \Box P(x, y))$$

to the valid FOL formula

$$(\exists\, y\,\forall\, x : \boxed{\lambda x, y : \Box P(x,y)}(x,y)) \Rightarrow (\forall\, x\,\exists\, y : \boxed{\lambda x, y : \Box P(x,y)}(x,y))$$

rather than the formula

$$(\exists\, y\,\forall\, x : \boxed{\lambda y, x : \Box P(x,y)}(y,x)) \Rightarrow (\forall\, x\,\exists\, y : \boxed{\lambda x, y : \Box P(x,y)}(x,y))$$

obtained according to the definition given in Section 4, which results in the two fresh operators being distinct. In general, we would like coalesced versions of different expressions to use the same atomic symbol wherever that would be valid. For example, $\boxed{e_1 = e_2}$ and $\boxed{e_2 = e_1}$ could be the same symbol.

Rewriting a formula before coalescing can also make the translated obligation easier to prove. For example, the formula $\Box e$ for a rigid expression e can be replaced by $\boxed{\Box\text{FALSE}} \vee e$. In a modal logic whose $\Box$ modality is reflexive, the disjunct $\boxed{\Box\text{FALSE}}$ is not necessary. This allows the formula

$$\forall\, x, y : \Box(x = y) \Rightarrow \Box(f(x) = f(y))$$

for $f \in \mathcal{O}$ to be proved directly by translating with coalescing to FOL instead of requiring two steps, the first proving $(x = y) \texttt{=>} (f(x) = f(y))$ with FOL and the second being translated to ML. Another such rewriting is distributing TLA's modal *prime* operator over rigid operators used by TLAPS when translating to FOL.

We don't know yet if optimizations of the translations beyond those we have already implemented in TLAPS will be useful in practice. So far, we have proved only safety properties for realistic algorithms, which in TLA requires little temporal reasoning. We have begun writing formal liveness proofs, but TLAPS will not completely check them until we have a translation that can handle formulas which, like the Barcan formula, inextricably mix quantifiers and modal operators. We also have not yet implemented coalescing of non-Leibniz defined operators, but we expect to do that before we prepare the final version of this paper.

References

[1] Krzysztof R. Apt. Ten years of Hoare's logic: A survey—part I. *ACM Trans. Program. Lang. Syst.*, 3(4):431–483, October 1981.

[2] Christoph Benzmüller and Bruno Woltzenlogel Paleo. Gödel's God on the computer. In Stephan Schulz, Geoff Sutcliffe, and Boris Konev, editors, *10th Intl. Workshop Implementation of Logics*, EPiC Series. EasyChair, 2013.

[3] Torben Braüner and Silvio Ghilardi. First order modal logic. In Patrick Blackburn, Johan van Benthem, and Frank Wolter, editors, *Handbook of Modal Logic*, pages 549–620. Elsevier, 2007.

[4] Denis Cousineau, Damien Doligez, Leslie Lamport, Stephan Merz, Daniel Ricketts, and Hernán Vanzetto. TLA$^+$ proofs. In Dimitra Giannakopoulou and Dominique Méry, editors, *18th Intl. Symp. Formal Methods (FM 2012)*, volume 7436 of *LNCS*, pages 147–154, Paris, France, 2012. Springer.

[5] Melvin Fitting and Richard L. Mendelsohn. *First-Order Modal Logic.* Synthese Library. Springer, 1998.

[6] Harald Ganzinger, Robert Nieuwenhuis, and Pilar Nivela. The Saturate system, 1998. `http://www.mpi-inf.mpg.de/SATURATE/doc/Saturate/Saturate.html`.

[7] Ullrich Hustadt and Renate A. Schmidt. MSPASS: Modal reasoning by translation and first-order resolution. In Roy Dyckhoff, editor, *TABLEAUX*, volume 1847 of *LNCS*, pages 67–71. Springer, 2000.

[8] Leslie Lamport. *Specifying Systems, The TLA$^+$ Language and Tools for Hardware and Software Engineers.* Addison-Wesley, 2002.

[9] Leslie Lamport. Byzantizing Paxos by refinement. In David Peleg, editor, *Distributed Computing: 25th Intl. Symp. (DISC 2011)*, pages 211–224. Springer-Verlag, 2011.

[10] Hans Jürgen Ohlbach. Semantics-based translation methods for modal logics. *J. Log. Comput.*, 1(5):691–746, 1991.