Taming First-Order Logic

SZABOLCS MIKULÁS, *Department of Computer Science, King's College, London, Strand, London WC2R 2LS, UK. E-mail: szabolcs@dcs.kcl.ac.uk*

Abstract

In this paper we define computationally well-behaved versions of classical first-order logic and prove that the validity problem is decidable[1].

Keywords: first-order logic, decidability, relativization, mosaic, polyadic and counting quantifiers.

1 Taming

In [5], we developed a strategy for taming logics. The idea of taming can be described as follows. Let us assume that we have a well-investigated logic with some undesirable metalogical properties. An example is the incompleteness and undecidability of the finite variable fragment of classical first-order logic, *FOL*, with at least three variables, cf. [4] 4.1.3 and 4.2.18 for the equivalent algebraic results. Taming a logic amounts to finding a version of the logic such that (i) this version has nicer properties than the original logic and (ii) its power is close to that of the original logic. Usually, one can achieve these two goals in two steps: (a) by weakening the logic (e.g., by widening the class of models) such that the weakened logic has desirable properties, and (b) by strengthening this weakened version (e.g., by (re-)introducing connectives that are not definable after weakening) without losing the nice properties.

In [5], we stated that if we relativize the square version of pair arrow logic with arbitrary, or with reflexive and/or symmetric relations, then these relativized versions have nicer properties, e.g., they are complete and decidable. In pair arrow logic relativization amounts to the following. In the square version, the frames are Cartesian spaces: $W = U \times U$. In the relativized versions we require that W be an arbitrary, or a reflexive and/or symmetric relation. Further, we could strengthen these relativized versions by adding the difference operator to the language without losing completeness and decidability.

In this paper, we will apply this strategy for taming the finite variable fragment of *FOL*. Our main concern will be decidability. We will define generalizations of reflexivity and symmetry to relations of higher rank, and define relativized versions of *FOL*. These relativized versions of *FOL* are decidable, cf. [8], Theorem 2.10 and Corollary 2.11 below. These versions remain decidable after strengthening by adding polyadic quantifiers and graded modalities, cf. Theorem 2.10 and Corollary 2.11.

[1]Supported by Hungarian National Foundation for Scientific Research grant Nos. F17452 and T16448 and by UK EPSRC grant No. GR/K54946.

2 First-order logic

We define the following versions of classical first-order logic with finitely many variables, cf. [4].

DEFINITION 2.1 (First-order logics with n variables: L_n and ${}^rL_n$)
Let n be a fixed natural number.[2]

Ordinary first-order logic with n variables is defined as the ordered tuple $L_n = \langle F, M, \models\rangle$ for which the following conditions hold.

1. Let $V = \{v_0, \ldots, v_{n-1}\}$ be the set of *variables*. Let P denote the set of *atomic formulas*, i.e., $P = \{r(v_{j_0}, \ldots, v_{j_{n-1}}) : j_0, \ldots, j_{n-1} \in n, r \in R\}$ for some set R; the set $R = \{r_i : i \in I\}$ is called the set of *relation* or *predicate symbols*. Then the set F of *formulas* is the smallest set H satisfying:
   - $P \subseteq H$
   - $v_i = v_j \in H$ for every $i, j \in n$
   - $\varphi, \psi \in H, i \in n \Rightarrow (\varphi \wedge \psi), \neg\varphi, \exists v_i\varphi \in H$.
2. The class M of *models* is defined by

$$M = \{\langle U, I\rangle : U \neq \emptyset, I : R \longrightarrow \mathcal{P}({}^nU)\},$$

   where $\mathcal{P}(X)$ denotes the power set of X. If $\mathcal{M} = \langle U, I\rangle \in M$, then U is called the *universe* of $\mathcal{M}$. Let $k \in {}^nU$. We define the *satisfaction* relation $\langle \mathcal{M}, k\rangle \Vdash \varphi$ by induction on the complexity of φ:
   - $\langle \mathcal{M}, k\rangle \Vdash r(v_{j_0}, \ldots, v_{j_{n-1}}) \Longleftrightarrow \langle k(j_0), \ldots, k(j_{n-1})\rangle \in I(r) \quad (r \in R)$
   - $\langle \mathcal{M}, k\rangle \Vdash v_i = v_j \Longleftrightarrow k(i) = k(j) \quad (i, j \in n)$
   - if $\psi_1, \psi_2 \in F$ and $i \in n$, then

$$\begin{aligned}
\langle \mathcal{M}, k\rangle \Vdash \neg\psi_1 &\iff \text{not } \langle \mathcal{M}, k\rangle \Vdash \psi_1 \\
\langle \mathcal{M}, k\rangle \Vdash \psi_1 \wedge \psi_2 &\iff \langle \mathcal{M}, k\rangle \Vdash \psi_1 \ \&\ \langle \mathcal{M}, k\rangle \Vdash \psi_2 \\
\langle \mathcal{M}, k\rangle \Vdash \exists v_i\psi_1 &\iff (\exists k' \in {}^nU)(\forall j \neq i)(k'(j) = k(j) \ \&\ \langle \mathcal{M}, k'\rangle \Vdash \psi_1).
\end{aligned}$$

   If $\langle \mathcal{M}, k\rangle \Vdash \varphi$, then we say that the evaluation k *satisfies* the formula φ in the model $\mathcal{M}$. We say that φ is *true* in $\mathcal{M}$, in symbols $\mathcal{M} \models \varphi$, iff for every $k \in {}^nU$, $\langle \mathcal{M}, k\rangle \Vdash \varphi$. The formula φ is the *semantical consequence* of the set Γ of formulas if $\mathcal{M} \models \Gamma$ implies $\mathcal{M} \models \varphi$, for every model $\mathcal{M}$. Finally, φ is *valid* if it is a semantical consequence of the empty set of formulas.

Restricted first-order logic with n variables, ${}^rL_n$, differs from the ordinary logic in the following: in restricted logic the order of the variables in atomic formulas $r(v_0, \ldots, v_{n-1})$ is fixed.

In the above definition we required that every relation symbol r should have arity n. This is not a real restriction as the following example shows. If the intended meaning of r is a binary relation, we may require that $(\forall u, v, w, w', \ldots \in U)\langle u, v, w, \ldots\rangle \in I(r) \Longleftrightarrow \langle u, v, w', \ldots\rangle \in I(r)$, i.e., the meaning of r depends on only two arguments.

We note that in the above definition we restricted *FOL* to a relational vocabulary. The reason for this is that including function symbols allows us to interpret the undecidable tiling problem into the logic: see [3] for the undecidability of first-order logic with two variables and function symbols, and [7] for the undecidability of relativized versions (see below) of logics.

[2] We use the convention that $n = \{0, \ldots, n-1\}$.

2.1 Modalizing

To apply our taming strategy it is convenient to consider first-order logic as a multi-modal logic, cf., e.g., [9] and [10].

Let W be an n-ary relation, i.e., $W \subseteq {}^nU$ for some set U. Let, for every $i \in n$, the binary relation T_i on W be defined as:

$$(\forall w, w' \in W) w T_i w' \iff w \lceil n \setminus \{i\} = w' \lceil n \setminus \{i\},$$

i.e., $(\forall j \neq i) w(j) = w'(j)$. Let $\tau \in {}^nn$, i.e., τ be a map from n into n. For every $\tau \in {}^nn$, let the binary relation S_τ on W be defined as:

$$(\forall w, w' \in W) w S_\tau w' \iff w' = w \circ \tau,$$

i.e., $w' = \langle w'(0), \ldots, w'(n-1)\rangle = \langle w(\tau(0)), \ldots, w(\tau(n-1))\rangle$. For every $i, j \in n$, the unary relation D_{ij} on W is defined as:

$$(\forall w \in W) D_{ij} w \iff w(i) = w(j).$$

We are ready to define the modal versions of L_n and ${}^rL_n$.

Definition 2.2 (Modal versions of L_n and ${}^rL_n$: $LIQS_n$ and LIQ_n)
The logic $LIQS_n$ is defined as the ordered tuple $\langle F, M, \models\rangle$ for which the following hold.

F is the set of formulas built up from a set R of propositional variables using the Boolean connectives, the unary connectives $\Diamond_i$ $(i \in n)$ and σ_τ $(\tau \in {}^nn)$, and the constants δ_{ij} $(i, j \in n)$. The connectives σ_τ and δ_{ij} are called *substitution* and *identity*, respectively.

A *frame* for $LIQS_n$ is an ordered tuple $\langle W, T_i, S_\tau, D_{ij} : i, j \in n, \tau \in {}^nn\rangle$ such that $W = {}^nU$ for some set U. A *model* is a frame–evaluation pair, where an evaluation $I : R \longrightarrow \mathcal{P}(W)$ is a map associating a subset of W to every propositional variable $r \in R$. M denotes the class of models. If no confusion is likely to arise, we will denote a model $\langle W, T_i, S_\tau, D_{ij}, I : i, j \in n, \tau \in {}^nn\rangle$ as $\langle W, T_i, S_\tau, D_{ij}, I\rangle$.

Let $\mathcal{M}$ be a model, and let $w \in W$ be an element of the universe of $\mathcal{M}$. We define *satisfaction* $\langle \mathcal{M}, w\rangle \Vdash \varphi$ by induction on the complexity of φ:

- $\langle \mathcal{M}, w\rangle \Vdash r \iff w \in I(r) \quad (r \in R)$
- $\langle \mathcal{M}, w\rangle \Vdash \delta_{ij} \iff D_{ij} w \quad (i, j \in n)$
- if $\psi_1, \psi_2 \in F$ and $i \in n$, then

$$\begin{array}{rcl}
\langle \mathcal{M}, w\rangle \Vdash \neg\psi_1 & \iff & \text{not } \langle \mathcal{M}, w\rangle \Vdash \psi_1 \\
\langle \mathcal{M}, w\rangle \Vdash \psi_1 \wedge \psi_2 & \iff & \langle \mathcal{M}, w\rangle \Vdash \psi_1 \ \& \ \langle \mathcal{M}, w\rangle \Vdash \psi_2 \\
\langle \mathcal{M}, w\rangle \Vdash \Diamond_i \psi_1 & \iff & (\exists w' \in W)(w T_i w' \ \& \ \langle \mathcal{M}, w'\rangle \Vdash \psi_1) \\
\langle \mathcal{M}, w\rangle \Vdash \sigma_\tau \psi_1 & \iff & (\exists w' \in W)(w S_\tau w' \ \& \ \langle \mathcal{M}, w'\rangle \Vdash \psi_1).
\end{array}$$

Truth in a model ($\models$) and the semantical consequence relation are defined in the usual way (cf. Definition 2.1).

The logic LIQ_n is defined as the substitution-free fragment of $LIQS_n$.

This formalism reflects the connection between L_n and its modal counterpart using the following translation ST:

$$\begin{aligned}&ST(r(v_0,\ldots,v_{n-1}))=r\\&ST(r(v_{\tau(0)},\ldots,v_{\tau(n-1)}))=\sigma_\tau r\\&ST(v_i=v_j)=\delta_{ij}\\&ST(\neg\varphi)=\neg ST\varphi\\&ST(\varphi\wedge\psi)=ST(\varphi)\wedge ST(\psi)\\&ST(\exists v_i\varphi)=\Diamond_i ST(\varphi).\end{aligned}$$

The following proposition ensures that the logic L_n (and ${}^rL_n$) can be interpreted into $LIQS_n$ (and LIQ_n) using ST.

PROPOSITION 2.3
Let φ be any formula of L_n. Let $\langle U,I\rangle$ be a model for L_n, and let k be an evaluation of the variables: $k\in{}^nU$. Then

$$\langle U,I,k\rangle\Vdash\varphi\Longleftrightarrow\langle{}^nU,T_i,S_\tau,D_{ij},I,k\rangle\Vdash ST(\varphi).$$

The same holds for ${}^rL_n$ and LIQ_n.

PROOF. It is a straightforward induction on the complexity of φ. ∎

By the above proposition, for $n>2$, the logics $LIQS_n$ and LIQ_n must be undecidable, otherwise L_n and ${}^rL_n$ were decidable.

2.2 Weakening

Now we define relativized versions $RLIQS_n$ and $RLIQ_n$ of $LIQS_n$ and LIQ_n. The reason for calling these versions *relativized* originates from algebraic logic. The algebraic counterparts of L_n and ${}^rL_n$ consist of algebras of n-ary relations with top element of the form ${}^nU$, cf. [4]. In [4], a method called relativization is introduced to define another class of algebras. If we intersect each element of an algebra with top element ${}^nU$ with a relation $W\subseteq{}^nU$, then we get another algebra whose elements are the elements of the original algebra intersected with W, and whose operations are the operations of the original algebra restricted to W.

We will apply the same technique below. Let $W\subseteq{}^nU$. We say that W is *locally cubic* if the following holds:

$$(\forall w\in W)(\forall\tau\in{}^nn)w\circ\tau\in W.$$

Let $ran(w)=\{w(0),\ldots,w(n-1)\}$. Then the above condition says that ${}^n ran(w)\subseteq W$.

REMARK 2.4
We note that the above condition can be considered as the generalization of being a symmetric and reflexive (binary) relation. The generalization of symmetry amounts to requiring the above condition on W only for permutations $\tau\in{}^nn$, i.e., for one-one maps. Another option is if we require that W satisfies the following:

$$(\forall w\in W)(\forall i,j\in n)w(i/j)\in W,$$

where $w(i/j) = \langle w(0), \ldots, w(i-1), w(j), w(i+1), \ldots, w(j-1), w(j), w(j+1), \ldots, w(n-1)\rangle$. That is, we substitute the jth value of w for the ith value. This condition may be considered as the generalization of reflexivity.

We will relativize the models for $LIQS_n$ with locally cubic relations. The reason for this is (i) that we want a logic as strong as possible[3] and (ii) that we want the alphabetical variants of formulas to be equivalent, cf. Definition 2.8 and Proposition 2.9 below. We note that problem (ii) does not arise in the case of the restricted logic ${}^rL_n$. Indeed, there are versions of ${}^rL_n$ such that the universes of the models are non-locally cubic relations. See [8] for the decidability of these weakened versions.

Definition 2.5 (Relativized versions of $LIQS_n$ and LIQ_n: $RLIQS_n$ and $RLIQ_n$) We define the logic $RLIQS_n$ as $LIQS_n$ with the following modification. In the definition of a frame for $RLIQS_n$, we require the universe $W \subseteq {}^nU$ be a locally cubic relation. The logic $RLIQ_n$ is defined analogously.

As we mentioned above we may relativize with arbitrary, with reflexive, etc. relations as well. These versions will be called *completely relativized*, *reflexive*, etc. versions.

The decidability of $RLIQ_n$ is proved in [8], and the decidability of $RLIQS_n$ follows from Theorem 2.10 below.

Remark 2.6
(Generalized first-order models) The logics $RLIQS_n$ and $RLIQ_n$ correspond to the following versions of L_n and ${}^rL_n$. In the truth definition of formulas, we require that every valuation k must be in some (fixed) local cube $W \subseteq {}^nU$. For instance, the truth of the formula $\exists v_i \varphi$ at a valuation $k \in W$ depends on whether there is a valuation k' *belonging to* W such that k' satisfies φ and kT_ik'. This kind of models are called generalized first-order models in [1] and [8].

Remark 2.7
(Syntactical approach) We weakened L_n and ${}^rL_n$ by widening the class of models, i.e., by allowing generalized models, while we left the syntax the same. A more syntactical approach is described in [1]. This amounts to defining fragments of FOL by restricting quantification: quantified formulas must have the form $\exists x(r(\ldots, x, \ldots) \wedge \varphi(\ldots, x, \ldots))$. In many cases the two approaches are equivalent, cf. [1]. For instance, relativizing with locally cubic relations is equivalent to considering the fragment with quantified formulas of the form $\exists v_{\tau(i)}(r(v_{\tau(0)}, \ldots, v_{\tau(n-1)}) \wedge \varphi(v_{\tau(0)}, \ldots, v_{\tau(n-1)}))$ where r is a fixed relation symbol satisfying the following axiom:

$$\forall v_0, \ldots, v_{n-1}(r(v_0, \ldots, v_{n-1}) \leftrightarrow r(v_{\tau(0)}, \ldots, v_{\tau(n-1)}))$$

for every τ.

We mentioned above that relativizing with non-locally cubic relations may cause strange properties. Consider the formulas $\forall v_0 r(v_0, v_0)$ and $\forall v_1 r(v_1, v_1)$. These formulas are equivalent in L_n. On the other hand, if we relativize with a non-locally cubic relation, these formulas are not equivalent. Let $W = \{\langle 0,0\rangle, \langle 0,1\rangle\}$, and let $I(r) = \{\langle 0,0\rangle\}$ in a model $\mathcal{M}$ with universe W. Then $\mathcal{M} \models ST(\forall v_0 r(v_0, v_0))$ but $\mathcal{M} \not\models ST(\forall v_1 r(v_1, v_1))$.

---

[3] For instance, if we relativize with arbitrary relation, then $ST(\forall v_0 \exists v_1(v_0 = v_1))$ is not valid.

Definition 2.8 (Alphabetical variants)
Let π be a permutation of n, and let φ be an L_n-formula. We define the *alphabetical variant (along π)* φ_π of φ as:

$$\begin{array}{l}(r(v_{\tau(0)},\ldots,v_{\tau(n-1)}))_\pi = r(v_{\pi\circ\tau(0)},\ldots,v_{\pi\circ\tau(n-1)})\\ (v_i = v_j)_\pi = (v_{\pi(i)} = v_{\pi(j)})\\ (\neg\varphi)_\pi = \neg(\varphi_\pi)\\ (\varphi\wedge\psi)_\pi = (\varphi_\pi \wedge \psi_\pi)\\ (\exists v_i\varphi)_\pi = \exists v_{\pi(i)}(\varphi_\pi).\end{array}$$

The following proposition states that relativizing with locally cubic relations preserves that property of L_n that alphabetical variants are equivalent.

Proposition 2.9
Let π be any permutation of n.

1. Let $\mathcal{M} = \langle W, T_i, S_\tau, , D_{ij}, I\rangle$ be a model for $RLIQS_n$, and let $w \in W$. Then

$$\langle\mathcal{M}, w\circ\pi^{-1}\rangle \Vdash ST(\varphi_\pi) \Longleftrightarrow \langle\mathcal{M}, w\rangle \Vdash ST(\varphi).$$

2. For every model $\mathcal{M}$ of $RLIQS_n$,

$$\mathcal{M} \models ST(\varphi_\pi) \Longleftrightarrow \mathcal{M} \models ST(\varphi).$$

Proof. 1: It is an easy induction on the complexity of φ. We show the most complicated case. Let $\langle\mathcal{M}, w\rangle \Vdash ST(\exists v_i\varphi)$. Then $(\exists w' \in W) wT_iw'$ & $\langle\mathcal{M}, w'\rangle \Vdash ST(\varphi)$. By the induction hypothesis, $\langle\mathcal{M}, w'\circ\pi^{-1}\rangle \Vdash ST(\varphi_\pi)$. It is easy to check that $w\circ\pi^{-1}T_{\pi(i)}w'\circ\pi^{-1}$, whence $\langle\mathcal{M}, w\circ\pi^{-1}\rangle \Vdash \Diamond_{\pi(i)}ST(\varphi_\pi)$. By the definition of ST, $\langle\mathcal{M}, w\circ\pi^{-1}\rangle \Vdash ST(\exists v_{\pi(i)}\varphi_\pi)$, i.e., $\langle\mathcal{M}, w\circ\pi^{-1}\rangle \Vdash ST((\exists v_i\varphi)_\pi)$. The other direction is completely analogous, since π^{-1} is a permutation as well.

2: This easily follows from 1, once we realized that every w is a π-image of some w'. ■

The problem with $RLIQS_n$ and $RLIQ_n$ is that their expressive power is remarkably weaker than that of the original versions $LIQS_n$ and LIQ_n.

2.3 Strengthening

It is a natural question whether we can strengthen relativized logics so that the nice properties are preserved, cf. [1] 6.3. Below we will add generalized diamonds (or polyadic quantifiers) and graded modalities (or counting quantifiers) to $RLIQS_n$ and $RLIQ_n$, and show that these logics are still decidable. We will show an example why the expressive power of the relativized logics $RLIQS_n$ and $RLIQ_n$ is much weaker than that of $LIQS_n$ and LIQ_n.

In $LIQS_n$ and LIQ_n the $\Diamond$'s commute, i.e., the following is a valid formula: $\Diamond_i\Diamond_j\varphi \leftrightarrow \Diamond_j\Diamond_i\varphi$, corresponding to the first-order validity $\exists v_i\exists v_j\varphi \leftrightarrow \exists v_j\exists v_i\varphi$. By this observation we can define generalized diamonds as follows. Let $\alpha \subseteq n$ and $w, w' \in W$. The binary relation T_α on W is defined as:

$$wT_\alpha w' \Longleftrightarrow w\lceil n\setminus\alpha = w'\lceil n\setminus\alpha.$$

For any model $\mathcal{M} = \langle W, T_\alpha, S_\tau, D_{ij}, I\rangle$, we let

$$(\forall w \in W)(\langle \mathcal{M}, w\rangle \Vdash \Diamond_\alpha \varphi \iff (\exists w' \in W) w T_\alpha w' \ \& \ \langle \mathcal{M}, w'\rangle \Vdash \varphi).$$

For $\alpha = \{a_0, \ldots, a_k\}$, $\Diamond_\alpha$ can be defined as $\Diamond_{a_0} \ldots \Diamond_{a_k}$ in $LIQS_n$ and LIQ_n.

On the other hand, $\Diamond_\alpha$ is not definable in $RLIQS_n$ and $RLIQ_n$, since the diamonds do not commute. To see this let us consider the following example. Let W be the following locally square binary relation: ${}^2\{a, b\} \cup {}^2\{b, c\}$. Let $\mathcal{M}$ be a model with universe W, and let r be interpreted as $\langle c, c\rangle$: $I(r) = \{\langle c, c\rangle\}$. Then $\langle \mathcal{M}, \langle b, a\rangle\rangle \Vdash \Diamond_1 \Diamond_0 r$, while $\langle \mathcal{M}, \langle b, a\rangle\rangle \not\Vdash \Diamond_0 \Diamond_1 r$.

We can add $\Diamond_\alpha$ with the above interpretation as a primitive connective to $RLIQS_n$ and $RLIQ_n$, by adding the accessibility relations T_α ($\alpha \subseteq n$) to the frames of $RLIQS_n$ and $RLIQ_n$. This way we can strengthen these logics, and they are decidable as an immediate consequence of Theorem 2.10, cf. Corollary 2.11.

Another way to look at $\Diamond_\alpha$ for $\alpha = n$ is to consider it as a graded modality $\langle 1\rangle$. For larger integers κ, $\langle \kappa\rangle$ is not definable in the above logic. We can define stronger logics without losing decidability in the following way. For every $\kappa \in \omega \setminus \{0\}$, let, in a model $\mathcal{M} = \langle W, T_\alpha, S_\tau, D_{ij}, I\rangle$,

$$\begin{aligned}\langle \mathcal{M}, w\rangle \Vdash \langle \kappa\rangle \varphi \quad \iff \quad & (\exists w_0, \ldots, w_{\kappa-1} \in W)|\{w_0, \ldots, w_{\kappa-1}\}| = \kappa \ \& \\ & (\forall \lambda < \kappa)\langle \mathcal{M}, w_\lambda\rangle \Vdash \varphi.\end{aligned}$$

Thus, $\langle 1\rangle$ coincides with $\Diamond_n$, since T_n is the universal relation on W.

For a logic L, let $L^{\alpha\kappa}$ denote that expansion where the connectives $\Diamond_\alpha$ and $\langle \kappa\rangle$ are defined. Thus, $RLIQS_n^{\alpha\kappa}$ and $RLIQ_n^{\alpha\kappa}$ denote the above strengthenings of $RLIQS_n$ and $RLIQ_n$. We will use the abbreviation $\langle \kappa!\rangle\varphi$ for $\neg\langle \kappa + 1\rangle\varphi \wedge \langle \kappa\rangle\varphi$.

We are going to prove the following theorem in Section 3.

Theorem 2.10
The logics $RLIQS_n^{\alpha\kappa}$ and $RLIQ_n^{\alpha\kappa}$ are decidable.

The obvious modifications in the proof of Theorem 2.10 yields the following, cf. Remark 2.4.

Corollary 2.11
The completely relativized, and the reflexive and/or symmetric versions of $LIQS_n^{\alpha\kappa}$ and $LIQ_n^{\alpha\kappa}$ are decidable.

In [8], it is shown that the decidability of $RLIQ_\omega$ follows from the decidability of every $RLIQ_n$ ($n \in \omega$). We define $RLIQS_\omega^{\alpha\kappa}$ in the above manner but we include $\Diamond_\alpha$ only for finite $\alpha \subset \omega$ and σ_τ only for finite transformations τ.[4] Straightforward modification of Németi's argument gives us the following.

Corollary 2.12
The logics $RLIQ_\omega^{\alpha\kappa}$ and $RLIQS_\omega^{\alpha\kappa}$ are decidable.

One might suggest to add the coordinatewise versions $\langle \kappa\rangle_\alpha$ of the graded modalities, too, with the following definition:

$$\begin{aligned}\langle \mathcal{M}, w\rangle \Vdash \langle \kappa\rangle_\alpha \varphi \quad \iff \quad & (\exists w_0, \ldots, w_{\kappa-1} \in W)|\{w_0, \ldots, w_{\kappa-1}\}| = \kappa \ \& \\ & (\forall \lambda < \kappa) w T_\alpha w_\lambda \ \& \ \langle \mathcal{M}, w_\lambda\rangle \Vdash \varphi.\end{aligned}$$

We showed in [7] that adding these connectives yields undecidable logics for dimensions ≥ 3.

[4] This is justified by the fact that in classical first-order logic formulas have finitely many variables.

3 Deciding by mosaics

In this section we will prove Theorem 2.10 by the so-called mosaic method, cf. [8], [6] and [10]. We will prove that a formula is satisfiable iff there is a finite set of finite mosaics satisfying some coherence conditions. These conditions will ensure that we can build a model using the mosaics as building blocks. The other main step in the proof is that it is decidable whether there is such a set of mosaics for a given formula.

We note that during the construction of the model, we may use a mosaic several (probably infinitely many) times. Thus, this decidability proof does not prove finite model property. Recently [2] showed that in fact our logic has the finite model property.

Let $n \in \omega$ and a set F of formulas be fixed.

Definition 3.1 (Mosaic)
A *mosaic* is a tuple $\mu = \langle U_\mu, E_\mu, l_\mu \rangle$ for which the following conditions hold.

The *universe* E_μ of μ is a locally cubic relation, and the *base* U_μ of μ is the smallest set such that $E_\mu \subseteq {}^nU_\mu$. The labelling function $l_\mu : {}^nU_\mu \longrightarrow \mathcal{P}(F)$ must satisfy the following conditions: for every formula in F, $(\forall e, e' \in E_\mu)(\forall s, s' \in {}^nU_\mu)(\forall i, j \in n)(\forall \tau \in {}^nn)(\forall \kappa \in \omega \setminus \{0\})(\forall \alpha \subseteq n)$

$$\begin{array}{lrcl}
(A1) & \varphi, \psi \in l_\mu(e) & \Leftrightarrow & \varphi \wedge \psi \in l_\mu(e) \\
(A2) & \varphi \in l_\mu(e) & \Leftrightarrow & \neg\varphi \notin l_\mu(e) \\
(A3) & \delta_{ij} \in l_\mu(e) & \Leftrightarrow & D_{ij}e \\
(A4) & \neg\langle\kappa\rangle\varphi \in l_\mu(e) & \Rightarrow & \neg[(\exists e_0, \ldots, e_{\kappa-1} \in E_\mu)(|\{e_0, \ldots, e_{\kappa-1}\}| = \kappa \ \& \\
 & & & (\forall \lambda < \kappa)\varphi \in l_\mu(e_\lambda))] \\
(A5) & \langle\kappa!\rangle\varphi \in l_\mu(e) & \Rightarrow & (\exists! e_0, \ldots, e_{\kappa-1} \in E_\mu)(|\{e_0, \ldots, e_{\kappa-1}\}| = \kappa \ \& \\
 & & & (\forall \lambda < \kappa)\varphi \in l_\mu(e_\lambda)) \\
(A6) & \langle\kappa\rangle\varphi \in l_\mu(e) & \Rightarrow & (\forall \lambda \leq \kappa)\langle\lambda\rangle\varphi \in l_\mu(e') \\
(A7) & \langle 1\rangle\varphi \in l(e) & \Leftrightarrow & \Diamond_n\varphi \in l(e) \\
(A8) & \Diamond_\alpha\varphi \in l_\mu(s) \ \& \ sT_\alpha s' & \Rightarrow & \Diamond_\alpha\varphi \in l_\mu(s') \\
(A9) & \Diamond_\alpha\varphi \in l_\mu(s) & \Rightarrow & (\forall \beta \supset \alpha)\Diamond_\beta\varphi \in l_\mu(s) \\
(A10) & \varphi \in l_\mu(e) \ \& \ e'S_\tau e & \Rightarrow & \sigma_\tau\varphi \in l_\mu(e') \\
(A11) & \sigma_\tau\varphi \in l_\mu(e) \ \& \ eS_\tau e' & \Rightarrow & \varphi \in l_\mu(e').
\end{array}$$

provided that the corresponding formulas are in F.

Remark 3.2
Since, in the above definition, we required E_μ be a locally cubic relation, it is not necessary to label the sequences $s \in {}^nU_\mu \setminus E_\mu$. We labeled these sequences because we want the method described below to work for the other relativized versions as well. Below we will define how to glue two mosaics together to form a bigger mosaic. If E_μ is not closed under substitutions, it may happen that, in the big mosaic, two sequences e and e' are in the T_α relation, while there is no sequence e'' in the common part of the two mosaics such that $eT_\alpha e''T_\alpha e'$. As we will see later, the lack of such a "witness" may cause difficulties in the proof. To prove Corollary 2.11, it is enough to modify the definition of a mosaic such that its universe E_μ is an appropriate (arbitrary, or reflexive and/or symmetric) relation.

We define the *distinguished* part $P_\mu \subseteq {}^nU_\mu$ of a mosaic μ as follows:

$$P_\mu = \{e \in E_\mu : \exists\varphi\exists\kappa \ \langle\kappa!\rangle\varphi \wedge \varphi \in l_\mu(e)\}.$$

Let $V_\mu = \bigcup\{ran(e) : e \in P_\mu\}$.

The set D_μ of *defects* of a mosaic is defined as

$$\begin{array}{rcl} D_\mu & = & \{\langle e, \langle\kappa\rangle\varphi\rangle : \neg[(\exists e_0, \ldots, e_{\kappa-1} \in E_\mu)(|\{e_0, \ldots, e_{\kappa-1}\}| = \kappa \;\& \\ & & (\forall\lambda < \kappa)\varphi \in l_\mu(e_\lambda))]\}\cup \\ & & \{\langle e, \Diamond_\alpha\varphi\rangle : \neg[(\exists e' \in E_\mu)\varphi \in l_\mu(e') \;\&\; eT_\alpha e')]\}. \end{array}$$

An *isomorphism*, in symbol $\cong$, between mosaics is a bijection between the universes preserving the relations T_α, S_τ, D_{ij}, and the labels of the sequences.

μ is a *submosaic* of μ', in symbols $\mu \subseteq \mu'$, if $P_\mu = P_{\mu'}$, $U_\mu \subseteq U_{\mu'}$, $E_\mu \subseteq E_{\mu'}$, and

$$(\forall s \in {}^nU_\mu)s \in E_{\mu'} \Rightarrow s \in E_\mu \;\&\; l_\mu(s) = l_{\mu'}(s).$$

The idea of the following definition is to define a finite set of "small" mosaics such that the defects of the mosaics disappear by gluing mosaics together to form a bigger mosaic.

Definition 3.3 (Good set of mosaics, GSM)
Let M be a finite set of mosaics. M is a *good set of mosaics*, $M \in GSM$, if the following conditions hold.

1. $(\forall\mu \in M)U_\mu \subseteq U \cup V_\mu$ for some set U such that $|U| \leq n$;
2. $(\forall\mu, \mu' \in M)P_\mu = P_{\mu'}$;
3. Condition for $\langle\kappa\rangle$:

$$\begin{array}{l} (\forall\mu \in M)(\forall e \in E_\mu)(\forall\langle\kappa\rangle\varphi \in l_\mu(e))(\exists\mu_0, \ldots, \mu_{\kappa-1} \in M) \\ (\exists\nu, \nu_0, \ldots, \nu_{\kappa-1}, \nu_\kappa \text{ mosaics})(\exists h, h_0, \ldots, h_{\kappa-1} \text{ isomorphisms})(\forall\lambda < \kappa) \\ h : \mu \longrightarrow \nu \;\&\; h_\lambda : \mu_\lambda \longrightarrow \nu_\lambda \;\&\; \nu \subseteq \nu_\kappa \;\&\; \nu_\lambda \subseteq \nu_\kappa \;\& \\ (\exists e_\lambda \in E_{\mu_\lambda})\varphi \in l_{\mu_\lambda}(e_\lambda) \;\&\; (\forall\iota \neq \lambda)h_\iota(e_\iota) \neq h_\lambda(e_\lambda). \end{array}$$

4. Condition for $\Diamond_\alpha$:

$$\begin{array}{l} (\forall\mu \in M)(\forall e \in E_\mu)(\forall\Diamond_\alpha\varphi \in l_\mu(e)) \\ (\exists\mu_0 \in M)(\exists\nu, \nu_0, \nu_1 \text{ mosaics})(\exists h, h_0 \text{ isomorphisms}) \\ h : \mu \longrightarrow \nu \;\&\; h_0 : \mu_0 \longrightarrow \nu_0 \;\&\; \nu \subseteq \nu_1 \;\&\; \nu_0 \subseteq \nu_1 \;\& \\ (\exists e_0 \in E_{\mu_0})\varphi \in l_{\mu_0}(e_0) \;\&\; h(e)T_\alpha h_0(e_0). \end{array}$$

Now we turn to proving Theorem 2.10.

Proof. (of Theorem 2.10) Our goal is to decide whether a formula is valid. A formula is valid iff its negation is not satisfiable. A formula χ is satisfiable iff so is $\chi \wedge \langle 1!\rangle p \wedge p$ where p is a propositional variable not occurring in χ. Thus, we can assume that χ has the above form, and decide whether χ is satisfiable.

The proof will consist of two steps: (i) a formula is satisfiable iff there is a good set of mosaics for this formula, and (ii) it is decidable whether there is a good set of mosaics for a given formula.

Let χ be given. We define an appropriate closure set F of formulas:

$$\begin{array}{rcl} F_0 & = & \text{subformulas of } \chi \\ F_1 & = & \{\langle\lambda!\rangle\varphi \wedge \varphi : (\exists\iota \geq \lambda)\langle\iota\rangle\varphi \in F_0\} \cup F_0 \\ F & = & \text{subformulas of the elements of } F_1. \end{array}$$

Note that F is a finite set of formulas, and that the size of F is computable from χ.

LEMMA 3.4
χ is satisfiable $\iff$ there is a GSM M for F such that $(\exists\mu \in M)(\exists e \in E_\mu)\chi \in l_\mu(e)$.

LEMMA 3.5
It is decidable whether there is a GSM M for F such that $(\exists\mu \in M)(\exists e \in E_\mu)\chi \in l_\mu(e)$.

PROOF. (OF LEMMA 3.5) If there is a GSM satisfying the conditions of Lemma 3.4, then there is such a GSM where the size of the mosaics is bounded. Indeed, let k be the size of F, and let j be the largest number such that $\langle j!\rangle\varphi \wedge \varphi$ occurs in F. Then there are at most $k \cdot j$ sequences labeled by formulas of the form $\langle\kappa!\rangle\psi \wedge \psi$, since a formula of the above kind occurs precisely on κ sequences by $(A5)$. Thus the size of P_μ is not greater than $k \cdot j$ for every mosaic μ. Hence $|V_\mu| \leq n \cdot k \cdot j$ and then $|U_\mu| \leq n \cdot k \cdot j + n$. There are only finitely many sets of sequences of this size (up to isomorphism), and they can be labeled by F only in finitely many ways, since F is finite. That is, we need to check only finitely many finite sets of mosaics whether at least one of them satisfies the conditions of the lemma, and this is a decidable procedure. ∎

Thus to prove Theorem 2.10 it suffices to prove Lemma 3.4.

PROOF. (OF LEMMA 3.4) $\Rightarrow$: It is not hard to see that we can "cut out" the appropriate set of mosaics from a model where χ is satisfied at some world.

$\Leftarrow$: First, using the elements of M as building blocks, we will construct a (probably infinite) mosaic without defects. Then we will define a model where there is a world satisfying χ.

0TH STEP: Take a mosaic $\mu \in M$ such that $(\exists e \in E_\mu)\chi \in l_\mu(e)$.

$k+1$ST STEP: By the induction hypothesis, the finite mosaic $G_k = \langle U_{G_k}, E_{G_k}, l_{G_k}\rangle$ constructed so far consists of isomorphic copies of members of M, i.e.,

$$(\forall e \in E_{G_k})(\exists\pi \subseteq G_k)(\exists\mu \in M)e \in E_\pi \;\&\; \pi \cong \mu.$$

Enumerate all the defects D_{G_k} of G_k.

CASE 1: the first defect has the form $\langle e, \langle\kappa\rangle\varphi\rangle$. Note that $\langle\lambda!\rangle\varphi \notin l(e)$ $(\lambda \in \omega)$, otherwise, by $(A5)$, the witnesses were in the distinguished part P_{G_k}. We make the following construction to make this defect disappear. By the induction hypothesis, there is a π such that $e \in E_\pi$, $\pi \subseteq G_k$, and there is a $\mu \in M$ such that $\pi \cong \mu$. Then, by the $\langle\kappa\rangle$-condition in Definition 3.3 of a GSM, there are $\mu_0, \ldots, \mu_\kappa \in M$ and mosaics $\nu, \nu_0, \ldots, \nu_{\kappa-1}, \nu_\kappa$ satisfying the $\langle\kappa\rangle$-condition. We will add an isomorphic copy of ν_κ to G_k.

Since $\pi \cong \mu \cong \nu$, there is an isomorphism $f : \nu \longrightarrow \pi$. This f induces a map $f' : U_\nu \longrightarrow U_\pi$. Let f'' be an extension of f' to U_{ν_κ} such that $(\forall u, u' \in U_{\nu_\kappa} \setminus U_\nu)f''(u) \notin U_{G_k} \;\&\; u \neq u' \Rightarrow f''(u) \neq f''(u')$. Let the f''-image of U_{ν_κ} be denoted by U. Let f^+ be the following extension of f: $(\forall s \in {}^nU_{\nu_\kappa})f^+(s) = \langle f''(s(i)) : i < n\rangle$. If we define the label of $f^+(s)$ as $l_{\nu_\kappa}(s)$, then the restriction of f^+ to ν_λ is an isomorphism to some π_λ for every $\lambda < \kappa$.

Case 2: the first defect has the form $\langle e, \Diamond_\alpha\varphi\rangle$. Then we make essentially the same construction as above. The only difference is that we need mosaics ν, ν_0, ν_1 satisfying the $\Diamond_\alpha$-condition of Definition 3.3, and add ν_1 to G_k.

Let $G' = \langle U_{G'}, E_{G'}, l_{G'}\rangle$ be defined as follows:

$$\begin{aligned} U_{G'} &= U_{G_k} \cup U \\ E_{G'} &= E_{G_k} \cup \{f^+(e) : e \in E_{\nu_\kappa}\} \\ l'_{G'} &= l_{G_k} \cup \{\langle f^+(s), l_{\nu_\kappa}(s)\rangle : s \in {}^nU_{\nu_\kappa}\}. \end{aligned}$$

Note that we did not define the labels of the sequences $s \in {}^nU_{G'} \setminus ({}^nU_{G_k} \cup {}^nU_{\nu_\kappa})$ yet. Let $l_{G'}$ be any extension of $l'_{G'}$ such that $(\forall s \in {}^nU_{G'} \setminus ({}^nU_{G_k} \cup {}^nU_{\nu_\kappa}))(\forall \Diamond_\beta\psi \in F)$

$$\Diamond_\beta\psi \in l_{G'}(s) \Leftrightarrow (\exists s' \in {}^nU_{G_k} \cup {}^nU_{\nu_\kappa})s'T_\beta s \ \&\ \Diamond_\beta\psi \in l'_{G'}.$$

Such an extension exists, since $l'_{G'}$ satisfies condition $(A8)$, cf. below.

It is easy to see that in G' there are witnesses for the label $\langle\kappa\rangle\varphi$ (or, in case 2, for $\Diamond_\alpha\varphi$) of e.

It remains to check that G' satisfies the induction hypothesis.

Since $\pi \subseteq G_k$, $P_\pi = P_{G_k}$. Since $\nu, \nu_\lambda \subseteq \nu_\kappa$ $(\lambda < \kappa)$, $P_\nu = P_{\nu_\lambda} = P_{\nu_\kappa}$. Then $P_{\pi_\lambda} = P_\pi = P_{G_k}$, whence $P_{G'} = P_{G_k}$. From this easily follows that $\pi_\lambda \subseteq G'$.

We have to check that G' is indeed a mosaic. Conditions $(A1) - (A3)$, $(A6)$, $(A7)$, and $(A9) - (A11)$ clearly hold. (For $(A6)$ use that $P_{G'}$ is not empty, for $(A10)$ and $(A11)$ use that the universes of mosaics are closed under substitution.)

To check the other conditions we use the fact that, in E_{G_k}, there is a distinguished sequence, say r, with a $\langle 1!\rangle p \wedge p$ label. (Actually, this is the reason why we decide a formula of the form $\chi \wedge \langle 1!\rangle p \wedge p$.) Clearly, $r \in P_{G'} = P_{G_k} = P_{\pi_\kappa}$.

Let us check $(A8)$. Assume that $sT_\beta s'$ and $\Diamond_\beta\psi \in l_{G'}(s)$. If s and s' are in the same mosaic G_k or π_κ, then $(A8)$ holds by the induction hypothesis. So, assume that $s \in {}^nU_{G_k} \setminus {}^nU_{\pi_\kappa}$ and $s' \in {}^nU_{\pi_\kappa}$. First assume $\beta \neq n$, and let $V = \{s(i) : i \in n\} \cap \{s'(i) : i \in n\}$. Then $V \neq \emptyset$. Let $s'' \in {}^nV$ such that $s''(j) = s(j) = s'(j)$ for every $j \in n \setminus \beta$. Then $sT_\beta s''T_\beta s'$ and $s'' \in {}^nU_{G'} \cap {}^nU_{\pi_\kappa}$. Hence, if $\Diamond_\beta\psi \in l_{G'}(s) = l_{G_k}(s)$, then by the induction hypothesis $\Diamond_\beta\psi \in l_{G_k}(s'') = l_{G'}(s'') = l_{\pi_\kappa}(s'')$. Since π_κ is a mosaic, $\Diamond_\beta\psi \in l_{\pi_\kappa}(s') = l_{G'}(s')$. The same argument works for s and s' interchanged. If $\beta = n$, then use the distinguished sequence r (note that in this case V may be empty). If s or s' is in ${}^nU_{G'} \setminus ({}^nU_{G_k} \cup {}^nU_{\nu_\kappa}))$, then $(A8)$ holds by the definition of $l_{G'}$.

$(A5)$ holds because of the following. If $\langle\lambda!\rangle\psi \in l_{G'}(e)$, then there are precisely λ sequences $e_0, \ldots, e_{\lambda-1}$ in the mosaic where e is such that they are labeled by $\langle\lambda!\rangle\psi\wedge\psi$. Moreover, these sequences are in the distinguished part $P_{G'}$, and every sequence with a label $\langle\lambda!\rangle\psi \wedge \psi$ is in $P_{G'} = P_{\pi_\kappa}$. Thus, E_{π_κ} and $E_{G'}$ contain precisely λ sequences of the above kind.

Let us check $(A4)$. Let $\neg\langle\lambda\rangle\psi \in l_{G'}(e)$ for some $e \in E_{G'}$. Let $\iota = \max\{k : \langle k!\rangle\psi \in l_{G'}(e)\}$. Then $\iota \leq \lambda$ because of $(A6)$. First assume that $\iota \neq 0$. This means that $\langle\iota!\rangle\psi \in l_{G'}(e)$. Then use $(A5)$. If $\iota = 0$, then by $(A6)$ $\langle 1\rangle\psi \notin l(e')$ for every e'. This implies, by $(A4)$ for G' and π_κ, that there is no sequence e'' labeled by ψ.

We make the same construction with this new mosaic G' and the second enumerated defect, etc. In finitely many steps this construction terminates, and we get a mosaic G_{k+1} satisfying the induction hypothesis such that $D_{G_k} \cap D_{G_{k+1}} = \emptyset$.

ωTH STEP: Take the union of the already constructed mosaics. Then we get a (probably infinite) mosaic $G_\omega = \langle U_{G_\omega}, E_{G_\omega}, l_{G_\omega} \rangle$ without defects: $D_{G_\omega} = \emptyset$.

Let the valuation I be defined as

$$I(q) = \{e \in E_{G_\omega} : q \in l_{G_\omega}(e)\}$$

for every propositional variable q.

PROPOSITION 3.6
For every $\varphi \in F$ and $e \in E_{G_\omega}$,

$$\varphi \in l_{G_\omega}(e) \Longleftrightarrow \langle E_{G_\omega}, T_\alpha, S_\tau, D_{ij}, I, e \rangle \Vdash \varphi.$$

PROOF. It is an easy induction on the complexity of φ. For every connective use the corresponding conditions in Definition 3.1, and that G_ω does not contain any defect. ∎

By the previous proposition, we have a model and a world where χ is satisfied, finishing the proof of Lemma 3.4. ∎

As we mentioned above this finishes the proof of Theorem 2.10. ∎

Thanks are due to Maarten Marx for valuable suggestions.

References

[1] H. Andréka, J. van Benthem & I. Németi, *Back and Forth between Modal Logic and Classical Logic,* Journal of the IGPL 3(5), 1996.

[2] H. Andréka, I. Hodkinson & I. Németi, *Finite Algebras of Relations are Representable on Finite Sets*, submitted.

[3] Y. Gurevich, *The Decision Problem for Standard Classes*, Journal of Symbolic Logic 41:460–464, 1976.

[4] L. Henkin, D. Monk & A. Tarski, *Cylindric Algebras, Part I,II,* North-Holland, Amsterdam, 1971, 1985.

[5] M. Marx, Sz. Mikulás & I. Németi, *Taming Logic*, Journal of Logic, Language and Information, 4:207–226, 1995.

[6] Sz. Mikulás, *Taming Logics*, ILLC Dissertation Series 1995-12, University of Amsterdam, 1995.

[7] Sz. Mikulás, M. Marx, *Undecidable Relativizations of Algebras of Relations*, Journal of Symbolic Logic, accepted for publication.

[8] I. Németi, *Decidability of Weakened Versions of First-order Logic,* L. CSIRMAZ, D.M. GABBAY & M. DE RIJKE (eds.), *Logic Colloquium'92*, CSLI Publications, 177–241, 1995.

[9] Y. Venema, *Modal Logic of Quantification and Substitution*, Bulletin of the IGPL 2(1), 1994.

[10] Y. Venema & M. Marx, *A Modal Logic of Relations*, to appear.

Received 1 April 1996