The Complexity of One-agent Refinement Modal Logic*

Laura Bozzelli[†], Hans van Ditmarsch[‡], Sophie Pinchinat[§]

Abstract

We investigate the complexity of satisfiability for one-agent *refinement modal logic* (RML), an extension of basic modal logic (ML) obtained by adding *refinement quantifiers* on structures. RML is known to have the same expressiveness as ML, but the translation of RML into ML is of non-elementary complexity, and RML is at least *doubly* exponentially more succinct than ML. In this paper we show that RML-satisfiability is 'only' *singly* exponentially harder than ML-satisfiability, the latter being a well-known PSPACE-complete problem.

1 Introduction

Modal logics with propositional quantifiers have been investigated since Fine's seminal paper [Fine, 1970]. Fine distinguishes three different propositional quantifications, which allow different kinds of model transformations: quantifying over propositionally definable subsets (over booleans), quantifying over subsets definable in the logical language (of basic modalities and quantifiers), and quantifying over all subsets. Only the first two are, in our modern terms, bisimulation preserving. Propositional quantification can easily lead to undecidable logics [Fine, 1970; French, 2006]. Undecidability relies on the ability of propositional quantification to dictate the structural properties of the underlying model [French, 2006]. This has motivated, more recently, the introduction of bisimulation quantified logics [Visser, 1996; Hollenberg, 1998; French, 2006; Pinchinat, 2007]. In this framework, the quantification is over the models which are bisimilar to the current model except for a propositional variable p. This operation is bisimulation preserving, and these logics are decidable.

In [Balbiani *et al.*, 2008] the authors propose a novel way of quantifying, namely over modally definable submodels. Unlike the above proposals, this not merely involves changing the valuation of a proposition in a subdomain, but *restricting the model* to that subdomain. The setting for these logics is how to quantify over information change. In the logic APAL of [Balbiani *et al.*, 2008], an expression that we might write as $\exists\varphi$ for our purposes stands for 'there is a modal formula ψ such that in the submodel restriction to the states satisfying ψ it holds that φ'. This logic is undecidable [French and van Ditmarsch, 2008]. *Refinement modal logic* (RML) [van Ditmarsch and French, 2009; van Ditmarsch *et al.*, 2010; Bozzelli *et al.*, 2012a] is a generalization of this perspective to more complex model transformations than submodel restrictions. This is achieved by existential and universal quantifiers which range over the *refinements* of the current model. From the *atoms/forth/back* requirements of bisimulation, a refinement of a modal structure need only satisfy *atoms* and *back*. It is therefore the dual of a simulation that need only satisfy *atoms* and *forth*. Refinement is more general than model restriction, since it is equivalent to bisimulation followed by model restriction. From a syntactic point of view it is bisimulation quantification followed by a relativization [Bozzelli *et al.*, 2012a]. Just as in bisimulation quantified logics we have *explicit* quantification over propositional variables, refinement quantification as it is realized in refinement modal logic is *implicit* quantification over propositional variables, i.e., quantification over variables not occurring in the formula bound by the quantifier. In RML, an expression $\exists_r\varphi$ stands for 'there is a refinement wherein it holds that φ.'

As an example of a refinement consider the following four rooted (underlined) structures.

$$\begin{array}{rl} \bullet \longrightarrow \bullet \longrightarrow \bullet \longrightarrow \bullet & M \\ \bullet \longrightarrow \bullet \longrightarrow \bullet & M' \\ \bullet \longleftarrow \bullet \longrightarrow \bullet \longrightarrow \bullet & M'' \\ \bullet \longleftarrow \bullet \longleftarrow \bullet \longleftarrow \bullet \longrightarrow \bullet \longrightarrow \bullet \longrightarrow \bullet & M''' \end{array}$$

With respect to the first model, M, the second one, M', is a model restriction. Model M'' is a refinement of M. It is not a model restriction. However, it is a model restriction of M''', a bisimilar copy of M. Refinements have really different properties, e.g., a formula like $\Diamond\Box\bot \wedge \Diamond\Diamond\Box\bot$ is clearly false in any model restriction of M, but it true in its refinement M''. In the root of the original model M is satisfied the formula $\exists_r(\Diamond\Box\bot \wedge \Diamond\Diamond\Box\bot)$, where $\exists_r$ is the refinement quantifier.

As amply illustrated in [Bozzelli *et al.*, 2012a], refinement

*The paper on which this extended abstract is based was the recipient of the best paper award of the 2012 JELIA conference in Toulouse [Bozzelli *et al.*, 2012b]. Hans van Ditmarsch is also affiliated to IMSc, Chennai, as research associate.

[†]Technical University of Madrid (UPM), laura.bozzelli@fi.upm.es
[‡]LORIA, CNRS / Univ. of Lorraine, hans.van-ditmarsch@loria.fr
[§]IRISA/INRIA, University of Rennes, Sophie.Pinchinat@irisa.fr

quantification has applications in many settings: in logics for games [Alur *et al.*, 2002; Pinchinat, 2007], it may correspond to a player discarding some moves; for program logics [Harel *et al.*, 2000], it may correspond to operational refinement; and for logics for spatial reasoning, it may correspond to subspace projections [Parikh *et al.*, 2007].

We now get to the content of this paper and its novel contributions. We focus on complexity issues for (one-agent) refinement modal logic [van Ditmarsch and French, 2009; van Ditmarsch *et al.*, 2010; Bozzelli *et al.*, 2012a], the extension of (one-agent) basic modal logic (ML) obtained by adding the existential and universal refinement quantifiers $\exists_r$ and $\forall_r$. It is known [van Ditmarsch *et al.*, 2010; Bozzelli *et al.*, 2012a] that RML has the same expressivity as ML, but the translation of RML into ML is of non-elementary complexity and no elementary upper bound is known for its satisfiability problem [Bozzelli *et al.*, 2012a]. In fact, an upper bound in 2EXPTIME has been claimed in [van Ditmarsch *et al.*, 2010] by a tableaux-based procedure: the authors later concluded that the procedure is sound but not complete [Bozzelli *et al.*, 2012a]. In this paper, we close that gap. We also investigate the complexity of satisfiability for some equi-expressive fragments of RML. We associate with each RML formula φ a parameter $\Upsilon_w(\varphi)$ corresponding to a slight variant of the classical quantifier alternation depth (measured w.r.t. $\exists_r$ and $\forall_r$), and for each $k \geq 1$ we consider the fragment RML^k consisting of the RML formulas φ such that $\Upsilon_w(\varphi) \leq k$. Moreover, we consider the existential (resp., universal) fragment $\mathsf{RML}^\exists$ (resp. $\mathsf{RML}^\forall$) obtained by disallowing the universal (resp. exist.) refinement quantifier.

In order to present our results, first, we recall some computational complexity classes. We assume familiarity with the standard notions of complexity theory [Johnson, 1990; Papadimitriou, 1994]. We will make use of the levels Σ_k^{EXP} ($k \geq 1$) of the exponential-time hierarchy EH, which are defined similarly to the levels Σ_k^{P} of the polynomial-time hierarchy PH, but with NP replaced with NEXPTIME. In particular, Σ_k^{EXP} corresponds to the class of problems decided by single exponential-time bounded Alternating Turing Machines (ATM, for short) with at most $k-1$ alternations and where the initial state is existential [Johnson, 1990]. Note that Σ_1^{EXP} = NEXPTIME. Recall that EH $\subseteq$ EXPSPACE and EXPSPACE corresponds to the class of problems decided by single exponential-time bounded ATM (with no constraint on the number of alternations) [Chandra *et al.*, 1981]. We are also interested in an intermediate class between EH and EXPSPACE, here denoted by $\mathsf{AEXP_{pol}}$, that captures the precise complexity of some relevant problems [Ferrante and Rackoff, 1975; Johnson, 1990; Rybina and Voronkov, 2003] such as the first-order theory of real addition with order [Ferrante and Rackoff, 1975; Johnson, 1990]. Formally, $\mathsf{AEXP_{pol}}$ is the class of problems solvable by single exponential-time bounded ATM with a polynomial-bounded number of alternations.

Our complexity results are summarized in Figure 1 where we also recall the well-known complexity of ML-satisfiability. For the upper bounds, the (technically non-trivial) main step in the proposed approach exploits a "small" size model property: we establish that like basic modal logic ML, RML enjoys a single exponential size model property.

ML	$\mathsf{RML}^\exists = \mathsf{RML}^1$	$\mathsf{RML}^\forall \subseteq \mathsf{RML}^2$
PSPACE-complete	$\in$ NEXPTIME PSPACE-hard	$\in \Sigma_2^{\mathsf{EXP}}$ NEXPTIME-hard

RML^{k+1} $(k \geq 1)$	RML
$\in \Sigma_{k+1}^{\mathsf{EXP}}$ Σ_k^{EXP}-hard	$\mathsf{AEXP_{pol}}$-complete

Figure 1: Complexity results for satisfiability of RML and RML-fragments

Our results are surprising for the following reason. While our results essentially indicate that satisfiability of RML is "only" *singly* exponentially harder than satisfiability of ML, it is known [Bozzelli *et al.*, 2012a] that RML is *doubly* exponentially more succinct than ML.

2 Preliminaries

Fix a finite set P of atomic propositions.

A (*one-agent Kripke*) *structure* (*over* P) is a tuple $M = \langle S, E, V \rangle$, where S is a set of states (or worlds), $E \subseteq S \times S$ is a transition (or accessibility) relation, and $V : S \mapsto 2^P$ is a P-*valuation* assigning to each state s the set of propositions in P which hold at s. For states s and t of M such that $(s,t) \in E$, we say that t is a *successor* of s. A *pointed* structure is a pair (M, s) consisting of a structure M and a designated initial state s of M.

A tree T is a prefix-closed subset of $\mathbb{N}^*$, where $\mathbb{N}$ is the set of natural numbers. The elements of T are called *nodes* and the empty word ε is the *root* of T. For $x \in T$, the set of children (or successors) of x is $\{x \cdot i \in T \mid i \in \mathbb{N}\}$. The *size* $|T|$ of T is the number of T-nodes. A (*rooted*) *tree structure* (over P) is a pair $\langle T, V \rangle$ such that T is a tree and $V : T \mapsto 2^P$ is a P-valuation over T. For $x \in T$, the *tree substructure of* $\langle T, V \rangle$ *rooted at* x is the tree structure $\langle T_x, V_x \rangle$, also denoted by $\langle T, V \rangle_x$, where $T_x = \{y \in \mathbb{N}^* \mid x \cdot y \in T\}$ and $V_x(y) = V(x \cdot y)$ for all $y \in T_x$. Note that a tree structure $\langle T, V \rangle$ corresponds to the pointed structure $(\langle T, E, V \rangle, \varepsilon)$, where $(x,y) \in E$ iff y is a child of x. Moreover, we can associate with any pointed structure (M,s) a tree structure, denoted by $Unw(M,s)$, obtained by unwinding M from s.

For two structures $M = \langle S, E, V \rangle$ and $M' = \langle S', E', V' \rangle$, a *refinement from* M *to* M' is a relation $\mathfrak{R} \subseteq S \times S$ such that for all $(s, s') \in \mathfrak{R}$: (i) $V(s) = V'(s')$, and (ii) if $(s', t') \in E'$ for some $t' \in S'$, then there is some state $t \in S$ such that $(s,t) \in E$ and $(t, t') \in \mathfrak{R}$. If, additionally, the inverse of $\mathfrak{R}$ is a refinement from M' to M, then $\mathfrak{R}$ is a *bisimulation* from M to M'. For states $s \in S$ and $s' \in S'$, (M', s') is a *refinement* of (M, s) (resp., (M, s) and (M', s') are *bisimilar*), written $(M,s) \succcurlyeq (M', s')$ (resp.,

$(M,s) \approx (M',s')$), if there is a refinement (resp., bisimulation) $\mathfrak{R}$ from M to M' s.t. $(s,s') \in \mathfrak{R}$. Note that $\succcurlyeq$ is a preorder (i.e., reflexive and transitive) and $\approx$ is an equivalence relation over pointed structures. For each pointed structure (M,s), $(M,s) \approx Unw(M,s)$.

We recall the syntax and semantics of one-agent *refinement modal logic* (RML) [van Ditmarsch *et al.*, 2010; Bozzelli *et al.*, 2012a], an equally expressive extension of basic modal logic [Blackburn *et al.*, 2001] obtained by adding the *existential and universal refinement quantifiers*. For technical convenience, the syntax of RML formulas φ over P is given in *positive form* as:

$$\varphi ::= p \mid \neg p \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \Diamond\varphi \mid \Box\varphi \mid \exists_r\varphi \mid \forall_r\varphi$$

where $p \in P$, $\Diamond\varphi$ reads as "possibly φ", $\Box\varphi$ reads as "necessarily φ", and $\exists_r$ and $\forall_r$ are the existential and universal refinement quantifiers. The *dual* $\widetilde{\varphi}$ of a RML formula φ is inductively defined as: $\widetilde{p} = \neg p$, $\widetilde{\neg p} = p$, $\widetilde{\varphi \vee \psi} = \widetilde{\varphi} \wedge \widetilde{\psi}$, $\widetilde{\Diamond\varphi} = \Box\widetilde{\varphi}$, $\widetilde{\Box\varphi} = \Diamond\widetilde{\varphi}$, $\widetilde{\exists_r\varphi} = \forall_r\widetilde{\varphi}$, and $\widetilde{\forall_r\varphi} = \exists_r\widetilde{\varphi}$. The size $|\varphi|$ of a formula φ is the number of distinct subformulas of φ. RML is interpreted over pointed structures (M,s). The satisfaction relation $(M,s) \models \varphi$ is inductively defined as usual, we only give the clause for the existential refinement quantifier (the one for the universal quantifier is its dual).

$(M,s) \models \exists_r\varphi$ iff for some refinement (M',s') of (M,s), $(M',s') \models \varphi$.

Note that $(M,s) \models \varphi$ iff $(M,s) \not\models \widetilde{\varphi}$. If $(M,s) \models \varphi$, we say that (M,s) *satisfies* φ, or also that (M,s) is a *model* of φ. A RML formula φ is *satisfiable* if φ admits some model.

Let ML be the fragment of RML obtained by disallowing the refinement quantifiers, which corresponds to basic modal logic [Blackburn *et al.*, 2001], and $\mathsf{RML}^{\forall}$ and $\mathsf{RML}^{\exists}$ be the fragments of RML obtained by disallowing the existential refinement quantifier and the universal refinement quantifier, respectively. Moreover, we introduce a family $\{\mathsf{RML}^k\}_{k\geq 1}$ of RML-fragments, where RML^k consists of the RML formulas whose *weak refinement quantifier alternation depth* (see below) is at most k.

Definition 1 (Weak Refinement Quant. Alternation Depth) We first define the *weak alternation length* $\ell(\chi)$ of finite sequences $\chi \in \{\exists_r, \forall_r\}^*$ of refinement quantifiers: $\ell(\varepsilon) = 0$, $\ell(Q) = 1$ for every $Q \in \{\exists_r, \forall_r\}$, and $\ell(QQ'\chi)$ is $\ell(Q'\chi)$ if $Q = Q'$, and $\ell(Q'\chi) + 1$ otherwise. For a RML formula φ, let $T(\varphi)$ be the standard tree encoding of φ, where each node is labeled by either a modality, or a boolean connective, or an atomic proposition. The *weak refinement quantifier alternation depth* $\Upsilon_w(\varphi)$ of a RML formula φ is the maximum of the alternation lengths $\ell(\chi)$ where χ is the sequence of refinement quantifiers along a path of $T(\exists_r\varphi)$ (note that we consider $T(\exists_r\varphi)$ and not $T(\varphi)$).

As an example, for $\varphi = \forall_r\exists_r p \vee \Box\exists_r(p \wedge \forall_r q)$, $\Upsilon_w(\varphi) = 3$. Note that $\mathsf{RML}^{\exists} = \mathsf{RML}^1$ and $\mathsf{RML}^{\forall} \subseteq \mathsf{RML}^2$. Moreover, for each RML formula φ, $\Upsilon_w(\forall_r\varphi) = \Upsilon_w(\widetilde{\forall_r\varphi}) + 1$. The following illustrates the succinctness of $\mathsf{RML}^{\exists}$ w.r.t. ML.

Example 1 *For $n \geq 1$, a n-block is a sequence $b_1, \ldots, b_{n+1}$ of $n+1$ bits. The following $\mathsf{RML}^{\exists}$ formula φ_n is satisfied by a tree structure iff there are two paths from the root encoding two n-blocks of the form $b_1, \ldots, b_n, b_{n+1}$ and $b_1, \ldots, b_n, b'_{n+1}$ s.t. $b_{n+1} \neq b'_{n+1}$: $\varphi_n := \exists_r\big(\Diamond^{n+1}(0 \wedge \neg 1) \wedge \Diamond^{n+1}(1 \wedge \neg 0) \wedge \bigwedge_{i=1}^{n} \bigvee_{b\in\{0,1\}} \Box^i(b \wedge \neg(1-b))\big)$. One can easily show that any ML formula which is equivalent to φ_n has size singly exponential in n.*

For each RML-fragment $\mathfrak{F}$, let $\mathsf{SAT}(\mathfrak{F})$ be the set of satisfiable $\mathfrak{F}$ formulas. We investigate the complexity of $\mathsf{SAT}(\mathfrak{F})$ for any $\mathfrak{F} \in \{\mathsf{RML}, \mathsf{RML}^{\exists}, \mathsf{RML}^{\forall}, \mathsf{RML}^2, \ldots\}$. Figure 1 depicts our complexity results.

Since RML is bisimulation invariant [van Ditmarsch *et al.*, 2010; Bozzelli *et al.*, 2012a], and because each pointed structure is bisimilar to its tree unwinding, w.l.o.g. we can assume that the semantics of RML is restricted to tree structures. Further, since RML and ML have the same expressivity [van Ditmarsch *et al.*, 2010; Bozzelli *et al.*, 2012a], we easily obtain:

Proposition 1 (Finite Model Property) *Let φ be a RML formula and $\langle T, V\rangle$ be a tree structure satisfying φ. Then, there is a* finite refinement *$\langle T_r, V_r\rangle$ of $\langle T, V\rangle$ satisfying φ.*

3 Upper Bounds

We first provide the upper bounds illustrated in Figure 1. Our approach consists of two steps. First we show that RML enjoys a singly exponential size model property. Using this result, we then show that $\mathsf{SAT}(\mathsf{RML})$ can be decided by a singly exponential-time bounded ATM whose number of alternations on an input φ is at most $\Upsilon_w(\varphi) - 1$ and whose initial state is existential.

Theorem 2 (Exponential Size Model Property) *For all satisfiable RML formulas φ and tree structures $\langle T, V\rangle$ such that $\langle T, V\rangle$ satisfies φ: there exists a* finite refinement *$\langle T', V'\rangle$ of $\langle T, V\rangle$ such that $\langle T', V'\rangle$ satisfies φ and $|T'| \leq |\varphi|^{3|\varphi|^2}$.*

The main steps in the proof of Theorem 2 are as follows. Fix a finite set P of atomic propositions and consider RML formulas and tree structures over P. Given a RML formula φ, we associate with φ tableaux-based *finite* objects called *constraints systems for* φ. Essentially, a constraint system $\mathcal{S}$ for φ is a tuple of hierarchically ordered finite tree structures which intuitively represents an *extended model* of φ: (1) each node x in a tree structure of $\mathcal{S}$ is additionally labeled by a set of subformulas of φ which hold at the tree substructure rooted at node x, and, in particular, the first tree structure, called *main structure*, represents a model of φ, and (2) the other tree structures of $\mathcal{S}$ are used to manage the $\exists_r$-subformulas of φ. In fact, in order to be an *extended model* of φ, $\mathcal{S}$ has to satisfy additional structural requirements which capture the semantics of the boolean connectives and all the modalities except the universal refinement quantifier $\forall_r$, the latter being only semantically captured. Let $\mathcal{C}(\varphi)$ be the set of these constraints systems for φ, which are said to be *well-formed, saturated, and semantically $\forall_r$-consistent*. We individuate a subclass $\mathcal{C}_{min}(\varphi)$ of $\mathcal{C}(\varphi)$ consisting of "*minimal*" constraints systems for φ whose sizes are *singly exponential* in the size of φ, and which can be obtained from φ by applying structural *completion rules*. Furthermore, we introduce a notion of "refinement" between constraint systems for φ which preserves the semantic $\forall_r$-consistency requirement. Then, given a *finite* tree structure $\langle T, V\rangle$ satisfying φ, we show that: (1) there is a constraint system $\mathcal{S} \in \mathcal{C}(\varphi)$ whose main structure is $\langle T, V\rangle$, and (2) starting from $\mathcal{S}$, it is possible to construct a minimal constraint system $\mathcal{S}_{min} \in \mathcal{C}_{min}(\varphi)$ which is a *refinement* of $\mathcal{S}$. This entails that the main structure of $\mathcal{S}_{min}$ is a refinement of $\langle T, V\rangle$ satisfying φ and having a single exponential size. Hence, by Proposition 1, Theorem 2 follows.

Using Theorem 2, we then show Theorem 3 that says that SAT(RML) can be decided by a singly exponential-time bounded ATM whose number of alternations on an input φ is at most $\Upsilon_w(\varphi) - 1$ and whose initial state is existential.

Theorem 3 *SAT(RML) $\in$ AEXP$_{pol}$ and SAT(RML^k) $\in$ Σ_k^{EXP} for each $k \geq 1$.*

4 Lower Bounds

We now provide the lower bounds illustrated in Figure 1. The main contribution is AEXP$_{\mathsf{pol}}$-hardness of SAT(RML), which is proved by a polynomial-time reduction from a suitable AEXP$_{\mathsf{pol}}$-complete problem.

For $k \geq 1$, a *k-ary deterministic Turing Machine* is a deterministic Turing machine $\mathcal{M} = \langle k, I, A, Q, \{q_{acc}, q_{rej}\}, q_0, \delta\rangle$ operating on k ordered semi-infinite tapes and having just one read/write head, where: I (resp., $A \supset I$) is the input (resp., work) alphabet, A contains the blank symbol $\#$, Q is the set of states, q_{acc} (resp., q_{rej}) is the terminal accepting (resp., rejecting) state, q_0 is the initial state, and $\delta : (Q \setminus \{q_{acc}, q_{rej}\}) \times A \to (Q \times A \times \{-1, +1\}) \cup \{1, \ldots, k\}$ is the transition function. In each non-terminal step, if the read/write head scans a cell of the ℓth tape ($1 \leq \ell \leq k$) and $(q, a) \in (Q \setminus \{q_{acc}, q_{rej}\}) \times A$ is the current pair state/scanned cell content, the following occurs:

- $\delta(q, a) \in Q \times A \times \{-1, +1\}$ (*ordinary moves*): $\mathcal{M}$ overwrites the tape cell being scanned, there is a change of state, and the read/write head moves one position to the left (-1) or right (+1) in accordance with $\delta(q, a)$.
- $\delta(q, a) = h \in \{1, \ldots, k\}$ (*jump moves*): the read/write head jumps to the left-most cell of the hth tape and the state remains unchanged.

$\mathcal{M}$ *accepts* a k-ary input $(w_1, \ldots, w_k) \in (I^*)^k$, written $\mathcal{M}(w_1, \ldots, w_k)$, if the computation of $\mathcal{M}$ from $(w_1, \ldots, w_k)$ (initially, the ℓth tape contains the word w_ℓ, and the head points to the left-most cell of the first tape) is accepting.

An instance of the *Alternation Problem* is a triple $(k, n, \mathcal{M})$, where $k \geq 1$, $n > 1$, and a $\mathcal{M}$ is a *polynomial-time bounded* k-ary deterministic Turing Machine with input alphabet I. The instance $(k, n, \mathcal{M})$ is *positive* iff the following holds, where $\mathsf{Q}_\ell = \exists$ if ℓ is odd, and $\mathsf{Q}_\ell = \forall$ otherwise (for all $1 \leq \ell \leq k$), $\mathsf{Q}_1 x_1 \in I^{2^n}. \mathsf{Q}_2 x_2 \in I^{2^n}. \ldots \mathsf{Q}_k x_k \in I^{2^n}. \mathcal{M}(x_1, \ldots, x_k)$. Note that the quantifications Q_i are restricted to words over I of length 2^n.

For $k \geq 1$, the *k-Alternation Problem* is the Alternation Problem restricted to instances of the form $(k, n, \mathcal{M})$ (i.e., the first input parameter is fixed to k), and the *Linear Alternation Problem* is the Alternation Problem restricted to instances of the form $(n, n, \mathcal{M})$. The following Prop. is then standard.

Proposition 4 *The Linear Alternation Problem is AEXP$_{pol}$-complete and for all $k \geq 1$, the k-Alternation Problem is Σ_k^{EXP}-complete.*

We then proceed as follows. Fix an instance $(k, n, \mathcal{M})$ of the Alternation Problem with $\mathcal{M} = \langle k, I, A, Q, \{q_{acc}, q_{rej}\}, q_0, \delta\rangle$. Since $\mathcal{M}$ is polynomial-time bounded, there is an integer constant $c \geq 1$ such that when started on a k-ary input $(w_1, \ldots, w_k)$, $\mathcal{M}$ reaches a terminal configuration in at most $(|w_1| + \ldots + |w_k|)^c$ steps. A (k, n)-input is a k-ary input $(w_1, \ldots, w_k)$ such that $w_i \in I^{2^n}$ for all $1 \leq i \leq k$. Let $\mathsf{c}(k, n) := c \cdot (n + \lceil \log k \rceil)$, where $\lceil \log k \rceil$ denotes the smallest $i \in \mathbb{N}$ such that $i \geq \log k$. Note that a configuration of $\mathcal{M}$ reachable from a (k, n)-input, called (k, n)*-configuration*, can be described as a tuple $\overrightarrow{C} = (C_1, \ldots, C_k)$ of k words $C_1, \ldots, C_k$ over $A \cup (Q \times A)$ of length exactly $2^{\mathsf{c}(k,n)}$ such that for some $1 \leq \ell \leq k$, C_ℓ is of the form $w \cdot (q, a) \cdot w' \in A^* \times (Q \times A) \times A^*$, and for $i \neq \ell$, $C_i \in A^{2^{\mathsf{c}(k,n)}}$. For a (k, n)-input $(a \cdot w_1, \ldots, w_k)$, the associated *initial* (k, n)*-configuration* is $((q_0, a) \cdot w_1 \cdot \#^{2^{\mathsf{c}(k,n)} - 2^n}, \ldots, w_k \cdot \#^{2^{\mathsf{c}(k,n)} - 2^n})$. Thus, the computations of $\mathcal{M}$ from (k, n)-inputs, called (k, n)*-computations*, can be described by sequences π of at most $2^{\mathsf{c}(k,n)}$ (k, n)-configurations. In fact, w.l.o.g., we can assume that π has length *exactly* $2^{\mathsf{c}(k,n)}$. This makes it possible to prove Th. 5. In the omitted proof we need an encoding of (k, n)-computations by suitable tree structures, called (k, n)*-computation tree code*.

Theorem 5 *One can construct a RML^{k+1} formula φ in time polynomial in n, k, and the size of the TM $\mathcal{M}$ such that (i) φ is a RML$^\forall$ formula if $k = 1$, and (ii) φ is satisfiable iff the instance $(k, n, \mathcal{M})$ of the Alternation Problem is positive.*

From Proposition 4 and Theorem 5 we also obtain:

Corollary 6 *SAT(RML) is AEXP$_{pol}$-hard, SAT(RML$^\forall$) is NEXPTIME-hard, and for all $k \geq 1$, SAT(RML^{k+1}) is Σ_k^{EXP}-hard.*

5 Concluding Remarks

An intriguing question left open is the complexity of satisfiability for *multi-agent* RML [van Ditmarsch *et al.*, 2010; Bozzelli *et al.*, 2012a]. Our approach does not seem to scale to the multi-agent case. Another interesting direction is to investigate the exact complexity of the fragments RML$^\exists$, RML$^\forall$, and RML^k, and the succinctness gap between RML^k and RML^{k+1} for each $k \geq 1$. Furthermore, since the modal μ-calculus extended with refinement quantifiers (RML$^\mu$, for short) is non-elementarily decidable [Bozzelli *et al.*, 2012a], it would be interesting to individuate weak forms of interactions between fixed-points and refinement quantifiers, which may lead to elementarily decidable and interesting RML$^\mu$-fragments.

References

[Alur *et al.*, 2002] R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, 2002.

[Balbiani *et al.*, 2008] P. Balbiani, A. Baltag, H. van Ditmarsch, A. Herzig, T. Hoshi, and T. De Lima. 'Knowable' as 'known after an announcement'. *Review of Symbolic Logic*, 1(3):305–334, 2008.

[Blackburn *et al.*, 2001] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, Cambridge, 2001. Cambridge Tracts in Theoretical Computer Science 53.

[Bozzelli *et al.*, 2012a] L. Bozzelli, H. van Ditmarsch, T. French, J. Hales, and S. Pinchinat. Refinement modal logic. Available at http://arxiv.org/abs/1202.3538, 2012.

[Bozzelli *et al.*, 2012b] L. Bozzelli, H. van Ditmarsch, and S. Pinchinat. The complexity of one-agent refinement modal logic. In *Proc. of 13th JELIA*, LNCS 7519, pages 120–133. Springer, 2012.

[Chandra *et al.*, 1981] A.K. Chandra, D.C. Kozen, and L.J. Stockmeyer. Alternation. *Journal of the ACM*, 28(1):114–133, 1981.

[Ferrante and Rackoff, 1975] J. Ferrante and C. Rackoff. A decision procedure for the first order theory of real addition with order. *SIAM Journal on Computing*, 4(1):69–76, 1975.

[Fine, 1970] K. Fine. Propositional quantifiers in modal logic. *Theoria*, 36(3):336–346, 1970.

[French and van Ditmarsch, 2008] T. French and H. van Ditmarsch. Undecidability for arbitrary public announcement logic. In *Advances in Modal Logic 7*, pages 23–42. College Publications, 2008.

[French, 2006] T. French. *Bisimulation quantifiers for modal logic*. PhD thesis, University of Western Australia, 2006.

[Harel *et al.*, 2000] D. Harel, D. Kozen, and J. Tiuryn. *Dynamic Logic*. MIT Press, 2000.

[Hollenberg, 1998] M. Hollenberg. *Logic and bisimulation*. PhD thesis, University of Utrecht, 1998.

[Johnson, 1990] D.S. Johnson. A catalog of complexity classes. In *Handbook of Theoretical Computer Science*, pages A:67–161. MIT Press, 1990.

[Papadimitriou, 1994] C.H. Papadimitriou. *Computational Complexity*. Addison Wesley, 1994.

[Parikh *et al.*, 2007] R. Parikh, L. Moss, and C. Steinsvold. Topology and epistemic logic. In *Handbook of Spatial Logics*, pages 299–341. Springer Verlag, 2007.

[Pinchinat, 2007] Sophie Pinchinat. A generic constructive solution for concurrent games with expressive constraints on strategies. In *Proc. 5th ATVA*, LNCS 4762, pages 253–267. Springer, 2007.

[Rybina and Voronkov, 2003] T. Rybina and A. Voronkov. Upper bounds for a theory of queues. In *Proc. 30th ICALP*, LNCS 2719, pages 714–724. Springer, 2003.

[van Ditmarsch and French, 2009] H. van Ditmarsch and T. French. Simulation and information. In *Proc. Knowledge Representation for Agents and Multi-Agent Systems*, LNAI 5605, pages 51–65. Springer, 2009.

[van Ditmarsch *et al.*, 2010] H. van Ditmarsch, T. French, and S. Pinchinat. Future event logic - axioms and complexity. In *Proc. 7th Advances in Modal Logic*, volume 8, pages 77–99. College Publications, 2010.

[Visser, 1996] A. Visser. Bisimulations, model descriptions and propositional quantifiers, 1996. Logic Group Preprint Series 161, Department of Philosophy, Utrecht University.