EXPTIME-completeness of thorough refinement on modal transition systems

Nikola Beneš [a,1], Jan Křetínský [a,*,2], Kim G. Larsen [b,3], Jiří Srba [b]

[a] *Faculty of Informatics, Masaryk University, Botanická 68a, 60200 Brno, Czech Republic*
[b] *Department of Computer Science, Aalborg University, Selma Lagerlöfs Vej 300, 9220 Aalborg Øst, Denmark*

ARTICLE INFO

ABSTRACT

Modal transition systems (MTS), a specification formalism introduced more than 20 years ago, has recently received a considerable attention in several different areas. Many of the fundamental questions related to MTSs have already been answered. However, the problem of the exact computational complexity of thorough refinement checking between two finite MTSs remained unsolved.

We settle down this question by showing EXPTIME-completeness of thorough refinement checking on finite MTSs. The upper-bound result relies on a novel algorithm running in single exponential time providing a direct goal-oriented way to decide thorough refinement. If the right-hand side MTS is moreover deterministic, or has a fixed size, the running time of the algorithm becomes polynomial. The lower-bound proof is achieved by reduction from the acceptance problem of alternating linear bounded automata and the problem remains EXPTIME-hard even if the left-hand side MTS is fixed and deterministic.

© 2012 Elsevier Inc. All rights reserved.

1. Introduction

Modal transition systems (MTS) is a specification formalism which extends the standard labeled transition systems with two types of transitions, the *may* transitions that are allowed to be present in an implementation of a given modal transition system and *must* transitions that must be necessarily present in any implementation. Modal transition systems hence allow to specify both safety and liveness properties. The MTS framework was suggested more than 20 years ago by Larsen and Thomsen [17] and has recently brought a considerable attention due to several applications to e.g. component-based software development [20,7], interface theories [25,21], modal abstractions and program analysis [23,11,14,18] and other areas [10,26], just to mention a few of them. A renewed interest in tool support for modal transition systems is recently also emerging [8,9,5]. A recent overview article on the theoretical foundations of MTSs and early tool development is available in [1].

Modal transition systems were designed to support *component-based* system development via a *stepwise refinement* process where abstract specifications are gradually refined into more concrete ones until an *implementation* of the system (where the may and must transitions coincide) is obtained. One of the fundamental questions is the decidability of a

---

* Corresponding author. Fax: +49 89 289 17207.
  E-mail addresses: xbenes3@fi.muni.cz (N. Beneš), jan.kretinsky@fi.muni.cz (J. Křetínský), kgl@cs.aau.dk (K.G. Larsen), srba@cs.aau.dk (J. Srba).

thorough refinement relation between two specifications S and T. We say that S thoroughly refines T iff every implementation of S is also an implementation of T. While for a number of other problems, like the common implementation problem, a matching complexity lower and upper bounds were given [2,16,3], the question of the exact complexity of thorough refinement checking between two finite MTSs remained unanswered.

In this article, we prove EXPTIME-completeness of thorough refinement checking between two finite MTSs. The hardness result is achieved by a reduction from the acceptance problem of alternating linear bounded automata, a well-known EXPTIME-complete problem, and it improves the previously established PSPACE-hardness [2]. The main reduction idea is based on the fact that the existence of a computation step between two configurations of a Turing machine can be locally verified (one needs to consider the relationships between three tape symbols in the first configuration and the corresponding three tape symbols in the second one, see e.g. [24, Theorem 7.37]), however, a nonstandard encoding of computations of Turing machines (which is crucial for our reduction) and the addition of the alternation required a nontrivial technical treatment. Moreover, we show that the problem remains EXPTIME-hard even if the left-hand side MTS is deterministic and of a constant size.

Initial proof ideas for the containment of the thorough refinement problem in EXPTIME were mentioned in [2] where the authors suggest a reduction of the refinement problem to validity checking of vectorized modal μ-calculus, which can be solved in EXPTIME—the authors in [2] admit that such a reduction relies on an unpublished popular wisdom, and they only sketch the main ideas hinting at the EXPTIME algorithm. In our article, we describe a novel technique for deciding thorough refinement in EXPTIME. The result is achieved by a direct goal-oriented algorithm performing a least fixed-point computation, and can be easily turned into a tableau-based algorithm. As a corollary, we also get that if the right-hand side MTS is deterministic or of a constant size, the algorithm for solving the problem runs in deterministic polynomial time.

2. Basic definitions

A *modal transition system* (MTS) over an action alphabet Σ is a triple $(P, \dashrightarrow, \longrightarrow)$, where P is a set of *processes* and $\longrightarrow \subseteq \dashrightarrow \subseteq P \times \Sigma \times P$ are *must* and *may* transition relations, respectively. Because in MTS whenever $S \xrightarrow{a} S'$ then necessarily also $S \dashrightarrow^{a} S'$, we adopt the convention of drawing only the must transitions $S \xrightarrow{a} S'$ in such cases. An MTS is *finite* if P and Σ are finite sets.

An MTS is an *implementation* if $\dashrightarrow = \longrightarrow$. As in implementations the must and may relations coincide, we can consider such systems as the standard *labeled transition systems*. The class of all implementations is denoted by $i\mathcal{MTS}$.

Definition 2.1. Let $M_1 = (P_1, \dashrightarrow_1, \longrightarrow_1)$, $M_2 = (P_2, \dashrightarrow_2, \longrightarrow_2)$ be MTSs over the same action alphabet Σ and $S \in P_1$, $T \in P_2$ be processes. We say that S *modally refines* T, written $S \leqslant_m T$, if there is a relation $R \subseteq P_1 \times P_2$ such that $(S, T) \in R$ and for every $(A, B) \in R$ and every $a \in \Sigma$:

1. if $A \dashrightarrow^{a}_1 A'$ then there is a transition $B \dashrightarrow^{a}_2 B'$ s.t. $(A', B') \in R$, and
2. if $B \xrightarrow{a}_2 B'$ then there is a transition $A \xrightarrow{a}_1 A'$ s.t. $(A', B') \in R$.

We often omit the indices in the transition relations and only use the symbols $\dashrightarrow$ and $\longrightarrow$ whenever it is clear from the context what transition system we have in mind. Note that on implementations modal refinement coincides with the classical notion of strong bisimilarity, and on modal transition systems without any must transitions it corresponds to the well-studied simulation preorder. It is also easy to argue that modal refinement is transitive.

Proposition 2.2. *The relation $\leqslant_m$ of modal refinement is transitive.*

We shall now observe that the modal refinement problem, i.e. the question whether a given process modally refines another given process, is tractable for finite MTSs.

Theorem 2.3. *The modal refinement problem for finite MTSs is P-complete.*

Proof. Modal refinement can be computed in polynomial time by the standard greatest fixed-point computation, similarly as in the case of strong bisimulation (for efficient algorithms implementing this strategy see e.g. [15,19]). P-hardness of modal refinement follows from P-hardness of bisimulation [4] (see also [22]). □

For convenient argumentation in some of the later proofs we extend the standard game-theoretic characterization of bisimilarity to the game characterization of modal refinement.

A *modal refinement game* (or simply a *modal game*) on a pair of processes (S, T) is a two-player game between *Attacker* and *Defender*. The game is played in *rounds*. In each round the players change the *current pair of processes* (A, B) (initially $A = S$ and $B = T$) according to the following rule:

Fig. 1. $S \leqslant_t T$ but $S \not\leqslant_m T$, and $S \not\leqslant_t U$ and $S \not\leqslant_m U$.

1. Attacker chooses an action $a \in \Sigma$ and one of the processes A or B. If he chose A then he performs a move $A \overset{a}{\dashrightarrow} A'$ for some A'; if he chose B then he performs a move $B \overset{a}{\longrightarrow} B'$ for some B'.
2. Defender responds by choosing a transition under a in the other process. If Attacker chose the move from A, Defender has to answer by a move $B \overset{a}{\dashrightarrow} B'$ for some B'; if Attacker chose the move from B, Defender has to answer by a move $A \overset{a}{\longrightarrow} A'$ for some A'.
3. The new current pair of processes becomes (A', B') and the game continues with a next round.

The game is similar to standard bisimulation games with the exception that Attacker is only allowed to attack on the left-hand side under may transitions (and Defender answers by may transitions on the other side), while on the right-hand side Attacker attacks under must transitions (and Defender answers by must transitions in the left-hand side process).

Any *play* (of the modal game) thus corresponds to a sequence of pairs of processes formed according to the above rule. A play (and the corresponding sequence) is finite iff one of the players gets stuck (cannot make a move). The player who got stuck lost the play and the other player is the winner. If the play is infinite then Defender is the winner.

The following fact is by a standard argument in analogy with strong bisimulation games: $S \leqslant_m T$ iff Defender has a winning strategy in the modal game starting with the pair (S, T); and $S \not\leqslant_m T$ iff Attacker has a winning strategy.

Example 2.4. Consider processes S and T in Fig. 1. We prove that S does not modally refine T. Indeed, Attacker has the following winning strategy in the modal game starting from (S, T). Attacker plays $S \overset{a}{\dashrightarrow} S_1$ to which Defender can answer by entering either T_1 or T_2. In the first case Attacker wins by playing $T_1 \overset{a}{\longrightarrow} T$ for which Defender has no answer from S_1 (no must transition under a is available) and Defender loses. In the second case Attacker wins by playing $S_1 \overset{a}{\dashrightarrow} S$ and Defender loses as well because no may transition under a is available from T_2. Similarly, one can argue that $S \not\leqslant_m U$.

We proceed with the definition of thorough refinement, a relation that holds for two modal specifications S and T iff any implementation of S is also an implementation of T.

Definition 2.5. For a process S let us denote by $[\![S]\!] = \{I \in i\mathcal{MTS} \mid I \leqslant_m S\}$ the set of all implementations of S. We say that S *thoroughly refines* T, written $S \leqslant_t T$, if $[\![S]\!] \subseteq [\![T]\!]$.

Clearly, if $S \leqslant_m T$ then also $S \leqslant_t T$ because the relation $\leqslant_m$ is transitive by Proposition 2.2. The opposite implication, however, does not hold as demonstrated by the processes S and T in Fig. 1 where one can easily argue that every implementation of S is also an implementation of T. On the other hand, $S \not\leqslant_t U$ because a process with just a single a-transition is an implementation of S but not of U.

3. Thorough refinement is EXPTIME-hard

In this section we prove that the thorough refinement relation $\leqslant_t$ on finite modal transition systems is EXPTIME-hard by reduction from the acceptance problem of alternating linear bounded automata. After recalling this acceptance problem, we show that it is enough to consider thorough refinement on tree implementations and we demonstrate how computation trees of linear bounded automata can be encoded as tree implementations of two modal transitions systems. The first system will have (the encoding of) any computation tree as its implementation while the second one will only have computation trees of incorrect or rejecting computations.

3.1. Alternating linear bounded automata

Definition 3.1. An *alternating linear bounded automaton* (*ALBA*) is a tuple $\mathcal{M} = (Q, Q_\forall, Q_\exists, \Sigma, \Gamma, q_0, q_{acc}, q_{rej}, \vdash, \dashv, \delta)$ where Q is a finite set of control states partitioned into $Q_\forall$ and $Q_\exists$, universal and existential states, respectively, Σ is a finite input alphabet, $\Gamma \supseteq \Sigma$ is a finite tape alphabet, $q_0 \in Q$ is the initial control state, $q_{acc} \in Q$ is the accepting state, $q_{rej} \in Q$ is the rejecting state, $\vdash, \dashv \in \Gamma$ are the left-end and the right-end markers that cannot be overwritten or moved, and $\delta : (Q \setminus \{q_{acc}, q_{rej}\}) \times \Gamma \longrightarrow 2^{Q \times \Gamma \times \{L, R\}}$ is a computation step function such that for all $q, p \in Q$ if $\delta(q, \vdash) \ni (p, a, D)$ then $a = \vdash$, $D = R$; if $\delta(q, \dashv) \ni (p, a, D)$ then $a = \dashv$, $D = L$; if $\delta(q, a) \ni (p, \vdash, D)$ then $a = \vdash$; and if $\delta(q, a) \ni (p, \dashv, D)$ then $a = \dashv$.

Remark 3.2. W.l.o.g. we assume that $\Sigma = \{a, b\}$, $\Gamma = \{a, b, \vdash, \dashv\}$, $Q \cap \Gamma = \emptyset$ and that for each $q \in Q_\forall$ and $a \in \Gamma$ it holds that $\delta(q, a)$ has exactly two elements (q_1, a_1, D_1), (q_2, a_2, D_2) where moreover $a_1 = a_2$ and $D_1 = D_2$. We fix this ordering

and the successor states q_1 and q_2 are referred to as the *first* and the *second successor*, respectively. The states q_{acc}, q_{rej} have no successors.

A *configuration* of $\mathcal{M}$ is given by its current state, the position of the head and the content of the tape. For technical reasons, we write it as a word over the alphabet $\varXi = Q \cup \Gamma \cup \{\vdash, \dashv, \exists, \forall, 1, 2, *\}$ (where $\exists$, $\forall$, 1, 2, $*$ are fresh symbols) in the following way. If the tape contains a word $\vdash w_1 a w_2 \dashv$, where $w_1, w_2 \in \Gamma^*$ and $a \in \Gamma$, and the head is scanning the symbol a in a state q, we write the configuration as $\vdash w_1 \alpha\beta q a w_2 \dashv$ where $\alpha\beta \in \{\exists *, \forall 1, \forall 2\}$.

The two symbols $\alpha\beta$ before the control state in every configuration are nonstandard, though important for the encoding of the computations into modal transition systems to be checked for thorough refinement. Intuitively, if a control state q is preceded by $\forall 1$ then it signals that the previous configuration (in a given computation) contained a universal control state and the first successor was chosen; similarly $\forall 2$ reflects that the second successor was chosen. Finally, if the control state is preceded by $\exists *$ then the previous control state was existential and in this case we do not keep track of which successor it was, hence the symbol $*$ is used instead. The *initial configuration* for an input word w is by definition $\vdash \exists * q_0 w \dashv$.

Depending on the present control state, every configuration is called either *universal*, *existential*, *accepting* or *rejecting*.

A *step of computation* is a relation $\longrightarrow$ between configurations defined as follows (where $w_1, w_2 \in \Gamma^*$, $\alpha\beta \in \{\forall 1, \forall 2, \exists *\}$, $a, a', b \in \Gamma$, $i \in \{1, 2\}$, and $w_1 a w_2$ and $w_1 b a w_2$ both begin with $\vdash$ and end with $\dashv$):

- $w_1 \alpha\beta q a w_2 \longrightarrow w_1 a' \forall i p w_2$
  if $\delta(q, a) \ni (p, a', R)$, $q \in Q_\forall$ and (p, a', R) is the ith successor,
- $w_1 \alpha\beta q a w_2 \longrightarrow w_1 a' \exists * p w_2$
  if $\delta(q, a) \ni (p, a', R)$ and $q \in Q_\exists$,
- $w_1 b \alpha\beta q a w_2 \longrightarrow w_1 \forall i p b a' w_2$
  if $\delta(q, a) \ni (p, a', L)$, $q \in Q_\forall$ and (p, a', L) is the ith successor, and
- $w_1 b \alpha\beta q a w_2 \longrightarrow w_1 \exists * p b a' w_2$
  if $\delta(q, a) \ni (p, a', L)$ and $q \in Q_\exists$.

Note that for an input w of length n all reachable configurations are of length $n + 5$. A standard result is that one can efficiently compute the set $Comp \subseteq \varXi^{10}$ of all compatible 10-tuples such that for each sequence $C = c_1 c_2 \cdots c_k$ of configurations $c_1, c_2, \ldots, c_k$, with the length of the first configuration being $\ell = |c_1| = n + 5$, we have $c_1 \longrightarrow c_2 \longrightarrow \cdots \longrightarrow c_k$ if and only if for all i, $0 \leqslant i \leqslant (k-1)\ell - 5$,

$$\big(C(i+1), C(i+2), C(i+3), C(i+4), C(i+5),$$
$$C(i+1+\ell), C(i+2+\ell), C(i+3+\ell), C(i+4+\ell), C(i+5+\ell) \big) \in Comp$$

where $C(i)$ is the ith element in the string C.

A *computation tree* for $\mathcal{M}$ on an input $w \in \Sigma^*$ is a tree $\mathcal{T}$ satisfying the following: the root of $\mathcal{T}$ is (labeled by) the initial configuration, and whenever N is a node of $\mathcal{T}$ labeled by a configuration c then the following hold:

- if c is accepting or rejecting then N is a leaf;
- if c is existential then N has one child labeled by some d such that $c \longrightarrow d$;
- if c is universal then N has two children labeled by the first and the second successor of c, respectively.

Without loss of generality, we shall assume from now on that any computation tree for $\mathcal{M}$ on an input w is finite (see e.g. [24, p. 198]) and that every accepting configuration contains at least four other symbols following after the state q_{acc}. We call a computation tree *accepting* if all its leaves are labeled with accepting configurations, otherwise it is a *rejecting* computation tree. A tree with nodes labeled by configurations that is neither accepting or rejecting computation tree is called *incorrect*.

We say that $\mathcal{M}$ *accepts* w iff there is a (finite) accepting computation tree for $\mathcal{M}$ on w. The following fact is well known (see e.g. [24]).

Proposition 3.3. *Given an ALBA M and a word w, the problem whether M accepts w is EXPTIME-complete.*

3.2. Encoding of configurations and computation trees

In this subsection we shall discuss the particular encoding techniques necessary for showing the lower bound. For technical convenience we will consider only tree encodings and so we first introduce the notion of tree-thorough refinement.

Definition 3.4. Let *Tree* denote the class of all MTSs with their graphs being trees. We say that a process S *tree-thoroughly refines* a process T, denoted by $S \leqslant_{tt} T$, if $[\![S]\!] \cap Tree \subseteq [\![T]\!] \cap Tree$.

Lemma 3.5. *For any two processes S and T, $S \leqslant_{tt} T$ iff $S \leqslant_t T$.*

Proof. The if case is trivial. For the only if case, we define an *unfold* $U(S)$ of a process S over an MTS $M = (P, \dashrightarrow, \longrightarrow)$ with an alphabet Σ to be a process S over an MTS $U(M) = (P^*, \dashrightarrow_U, \longrightarrow_U)$ over the same alphabet and where P^* is the set of all finite sequences over the symbols from P. The transition relations are defined as follows: for all $a \in \Sigma$, $T, R \in P$ and $\alpha \in P^*$, whenever $T \overset{a}{\dashrightarrow} R$ then $\alpha T \overset{a}{\dashrightarrow}_U \alpha T R$, and whenever $T \overset{a}{\longrightarrow} R$ then $\alpha T \overset{a}{\longrightarrow}_U \alpha T R$. Since the transitions in $U(S)$ depend only on the last symbol, we can easily see that $U(S) \leqslant_m S$ and $S \leqslant_m U(S)$ for every process S.

Let I be now an implementation of S. Its unfold $U(I)$ is also an implementation of S by $U(I) \leqslant_m I \leqslant_m S$ and the transitivity of $\leqslant_m$. By our assumption that $S \leqslant_{tt} T$ and the fact that $U(I)$ is a tree, we get that $U(I)$ is also an implementation of T. Finally, $I \leqslant_m U(I) \leqslant_m T$ and the transitivity of $\leqslant_m$ allow us to conclude that I is an implementation of T. $\square$

Let $\mathcal{M} = (Q, Q_\forall, Q_\exists, \Sigma, \Gamma, q_0, q_{acc}, q_{rej}, \vdash, \dashv, \delta)$ be an ALBA and $w \in \Sigma^*$ an input word of length n. We shall construct (in polynomial time) modal transition systems L and R such that $\mathcal{M}$ accepts w iff $L \not\leqslant_{tt} R$. The system L will encode (almost) all trees beginning with the initial configuration, while the implementations of R encode only the incorrect or rejecting computation trees.

Configurations, i.e. sequences of letters from $\varXi$, are not encoded straightforwardly as sequences of actions (the reason why this naive encoding does not work is explained later on in Remark 3.12). Instead we use two auxiliary actions π and σ. The intended implementations of L and R alternate between the actions π and σ on a linear path, while the symbols in the encoded configuration are present as side-branches on the path.

Formally, a sequence $a_1 a_2 a_3 \cdots a_n \in \varXi^*$ is encoded as

and denoted by $\mathrm{code}(a_1 a_2 \cdots a_n)$.

We now describe how to transform computation trees into their corresponding implementations. We simply concatenate the subsequent codes of configurations in the computation tree so that the end node of the previous configuration is merged with the begin node of the successor configuration. Whenever there is a (universal) branching in the tree, we do not branch in the corresponding implementation at its beginning but we wait until we reach the occurrence of the symbol $\forall$. The branching happens exactly before the symbols 1 or 2 that follow after $\forall$. This occurs in the same place on the tape in both of the configurations due to the assumption that the first and the second successor move simultaneously either to the left or to the right, and write the same symbol (see Remark 3.2). A formal definition of the encoding of computation trees into implementations follows.

Definition 3.6 (*Encoding computation trees into implementations*). Let $\mathcal{T}$ be a (finite) computation tree. We define its tree implementation $\mathrm{code}(\mathcal{T})$ inductively as follows:

- if $\mathcal{T}$ is a leaf labeled with a configuration c then $\mathrm{code}(\mathcal{T}) = \mathrm{code}(c)$;
- if the root of $\mathcal{T}$ is labeled by an existential configuration c with a tree $\mathcal{T}'$ being its child, then $\mathrm{code}(\mathcal{T})$ is rooted in the begin node of $\mathrm{code}(c)$, followed by $\mathrm{code}(\mathcal{T}')$ where the end node of $\mathrm{code}(c)$ and the begin node of $\mathrm{code}(\mathcal{T}')$ are identified;
- if the root of $\mathcal{T}$ is labeled by a universal configuration c with two children $d_1 \ldots \forall 1 \ldots d_n^1$ and $d_1 \ldots \forall 2 \ldots d_n^2$ that are roots of the subtrees $\mathcal{T}_1$ and $\mathcal{T}_2$, respectively, then $\mathrm{code}(\mathcal{T})$ is rooted in the begin node of $\mathrm{code}(c)$, followed by two subtrees $\mathrm{code}(\mathcal{T}_1)$ and $\mathrm{code}(\mathcal{T}_2)$ where the nodes in $\mathrm{code}(d_1 \ldots \forall)$ of the initial part of $\mathrm{code}(\mathcal{T}_1)$ are identified with the corresponding nodes in the initial part of $\mathrm{code}(\mathcal{T}_2)$ (note that by Remark 3.2 this prefix is common in both subtrees), and finally the end node of $\mathrm{code}(c)$ is identified with now the common begin node of both subtrees.

Fig. 2 illustrates this definition on a part of a computation tree, where the first configuration $c_1 \ldots c_n$ is universal and has two successor configurations $d_1 \ldots \forall 1 \ldots d_n^1$ and $d_1 \ldots \forall 2 \ldots d_n^2$.

3.3. The reduction—Part 1

We now proceed with the reduction. As mentioned earlier, our aim is to construct for a given ALBA $\mathcal{M}$ and a string w two modal transition systems L and R such that $L \not\leqslant_{tt} R$ iff $\mathcal{M}$ accepts w. Implementations of L will include all (also incorrect) possible computation trees. We only require that they start with the encoding of the initial configuration and do not "cheat" in the universal branching (i.e. after the encoding of every symbol $\forall$ there must follow a branching such that at least one of the branches encodes the symbol 1 and at least another one encodes the symbol 2).

As L should capture implementations corresponding to computations starting in the initial configuration, we set L to be the begin of $\mathrm{code}(\vdash \exists * q_0 w \dashv)$ and denote its end by M. After the initial configuration has been forced, we allow all possible continuations of the computation. This can be simply done by setting

Fig. 2. Computation tree encoding.

for all $a \in \Xi \setminus \{\forall\}$

Fig. 3. Fragment of the process L.

Fig. 4. Full specification of the process L.

$$M \overset{\pi}{\dashrightarrow} M_a$$
$$M_a \overset{\sigma}{\longrightarrow} M$$
$$M_a \overset{a}{\longrightarrow} X_a$$

for all letters $a \in \Xi \setminus \{\forall\}$ where the state X_a has no outgoing transitions as depicted in Fig. 3.

Finally, we add a fragment of MTS into the constructed process L which will guarantee the universal branching as mentioned above whenever the symbol $\forall$ occurs on a side-branch. The complete modal transition system L is now depicted in Fig. 4.

We now observe some simple facts regarding tree implementations of the process L.

Proposition 3.7. *Every tree implementation I of the process L satisfies that*

1. *every branch in I is labeled by an alternating sequence of π and σ actions, beginning with the action π, and if the branch is finite then it ends either with the action σ or with an action $a \in \Xi \setminus \{\forall\}$, and*

2. *every state in I with an incoming transition under the action π has at least one outgoing transition under the action σ and at least one outgoing transition under an action $a \in \Xi$, and*
3. *whenever from any state in I there are two outgoing transitions under some $a \in \Xi$ and $b \in \Xi$ then $a = b$, and moreover no further actions are possible after taking any transition under $a \in \Xi$, and*
4. *every branch in I longer than $2(n + 5)$ begins with the encoding of the initial configuration $\vdash \exists * q_0 w \dashv$ where $n = |w|$, and*
5. *every state in I with an incoming transition under σ from a state where the action $\forall$ is enabled satisfies that every outgoing transition under π leads to a state where either the action 1 or 2 is enabled (but not both at the same time), and moreover it has at least one such transition that enables the action 1 and at least one that enables the action 2.*

Of course, not every tree implementation of the process L represents a correct computation tree of the given ALBA. Implementations of L can widely (even uncountably) branch at any point and sequences of configurations they encode on some (or all) of their branches may not be correct computations of the given ALBA. Nevertheless, the encoding of any computation tree of the given ALBA is an implementation of the processes L, as stated by the following lemma.

Lemma 3.8. *Let $\mathcal{T}$ be a computation tree of an ALBA $\mathcal{M}$ on an input w. Then $\mathrm{code}(\mathcal{T}) \leqslant_m L$.*

Proof. We shall describe Defender's winning strategy in the modal refinement game starting from the begin node of $\mathrm{code}(\mathcal{T})$ and the process L. The tree $\mathrm{code}(\mathcal{T})$ clearly begins with the encoding of the initial configuration and an identical part is contained also in the beginning of the process L. Hence the game surely continues from the begin node of the code of the next configuration(s) in $\mathcal{T}$ and the process M (otherwise Attacker loses immediately, should he choose any of the side-branches). Now Attacker must attack under the action π in the left-hand side tree (no must transitions are enabled on the right-hand side from M) and Defender is matching this move from M in two different ways. Should the Attacker's next state contain a branch with a label a such that $a \in \Xi \setminus \{\forall\}$ then Defender plays π and enters the process M_a on the right-hand side. In order for Attacker to still have a chance to win, he must play the action σ in either of the processes and the players return to the situation where the right-hand side process is again in M. On the other hand if Attacker's next state after playing π in the left-hand side process contains a branch with the label $\forall$, then Defender enters under π the state $M_\forall$. As before, the only reasonable continuation for Attacker is to play σ in one of the processes and the players reach a pair of states where the left-hand side process branches under π into two different paths (due to Definition 3.6 of $\mathrm{code}(\mathcal{T})$) and the right-hand side process is in the state M'. Attacker can now choose one of the branches in either of the processes but Defender can safely match such an attack and the players after two rounds return to the situation where the right-hand side is again in the state M. To sum up, Defender has a winning strategy and thus $\mathrm{code}(\mathcal{T}) \leqslant_m L$. $\quad\square$

3.4. The reduction—Part 2

We now proceed with the construction of the right-hand side process R. Its implementations should be the codes of all incorrect or rejecting computation trees. To cover the notion of incorrect computation, we define the so-called bad path (see p. 57 for the definition of the relation *Comp*).

Definition 3.9. A sequence

$$c_1 c_2 c_3 c_4 c_5 \underbrace{a_1 a_2 \ldots a_{n-6} a_{n-5}}_{n-5 \text{ elements from } \Xi} d_1 d_2 d_3 d_4 d_5$$

is called a *bad path* if $(c_1, c_2, c_3, c_4, c_5, d_1, d_2, d_3, d_4, d_5) \in \Xi^{10} \setminus Comp$.

To cover the incorrect or rejecting computations, we loop in the process R under all actions, including the auxiliary ones, except for q_{acc}. For convenience we denote $\Xi' = \Xi \cup \{\pi, \sigma\}$. For any bad path, the process R can at any time nondeterministically guess the beginning of its first quintuple, realize it, then perform $n - 5$ times a sequence of π and σ, and finally realize the second quintuple. Moreover, we have to allow all possible detours of newly created branches to end in the state U where all available actions from Ξ' are always enabled and hence the continuation of any implementation modally refines U. Formally, for any $(c_1, c_2, c_3, c_4, c_5, d_1, d_2, d_3, d_4, d_5) \in \Xi^{10} \setminus Comp$ we add (disjointly) the following fragment into the process R (see also Fig. 5).

$$R \dashrightarrow^{\pi} V_1$$
$$V_j \xrightarrow{\sigma} W_j \xrightarrow{\pi} V_{j+1} \quad \text{for } 1 \leqslant j < n + 5$$
$$V_j \xrightarrow{c_j} X_j \quad \text{for } 1 \leqslant j \leqslant 5$$
$$V_{n+j} \xrightarrow{d_j} X_{5+j} \quad \text{for } 1 \leqslant j \leqslant 5$$
$$V_j \dashrightarrow^{x} U, \qquad W_j \dashrightarrow^{x} U, \qquad V_{n+5} \dashrightarrow^{x} U \quad \text{for } 1 \leqslant j < n + 5 \text{ and } x \in \Xi'$$

Fig. 5. A fragment of the system R for a bad path $c_1 c_2 c_3 c_4 c_5 \ldots d_1 d_2 d_3 d_4 d_5$.

$$U \xrightarrow{x} U \quad \text{for all } x \in \varXi'$$
$$R \xrightarrow{x} R \quad \text{for all } x \in \varXi' \setminus \{q_{acc}\}$$

We also add ten new states $N_1, \ldots, N_9, N_X$ and the following transitions: $R \xdashrightarrow{\pi} N_1 \xdashrightarrow{\varXi'} N_2 \xdashrightarrow{\varXi'} N_3 \xdashrightarrow{\varXi'} N_4 \xdashrightarrow{\varXi'} \cdots \xdashrightarrow{\varXi'} N_9$ and $N_1 \xRightarrow{q_{acc}} N_X$ where any transition labeled by $\varXi'$ is the abbreviation for a number of transitions under all actions from $\varXi'$.

Remark 3.10. We do not draw these newly added states $N_1, \ldots, N_9, N_X$ into Fig. 5 in order not to obstruct its readability because the reason why these states are added is purely technical. It is possible that there is an incorrect computation that ends with the last symbol q_{acc} but it cannot be detected by any bad path as defined in Definition 3.9 because that requires (in some situations) that there should be present at least four other subsequent symbols. By adding these new states into the process R, we guarantee that such situations where a branch in a computation tree ends in q_{acc} without at least four additional symbols will be easily matched in R by entering the state N_1.

Lemma 3.11. *Let I be a tree implementation of L such that every occurrence of q_{acc} in I is either preceded by a code of a bad path or does not continue with the encoding of at least four other symbols. Then $I \leqslant_m R$.*

Proof. We synthesize a winning strategy for Defender starting from the root of I and the process R in order to prove $I \leqslant_m R$.

Note that as long as the right-hand side process is in the state R then Attacker can attack only from the left-hand side processes. Defender's strategy is as follows:

1. If Attacker attacks under a π transition leading to a subtree from which at least one branch begins with the encoding of a bad path, Defender answers by a π transition from R leading to a state V_1 representing the fragment in the right-hand side process corresponding to this bad path. After that, Attacker is forced to switch sides and play under the must transitions the whole sequence of σ and π actions until reaching V_{n+5}; Defender will match this sequence by following the branch corresponding to this bad path in the left-hand side process and finally win. Should Attacker during this phase at any time decide to play again from the left-hand side, Defender will "escape" immediately to the state U from which Defender has a clear winning strategy.
2. On any other Attacker's move, Defender simply loops in R and Attacker cannot have played the action q_{acc} yet (the only action disabled in R) due to the assumption of the lemma. Should Attacker play a π transition enabling q_{acc} such that there is no continuation with at least 9 other transitions (i.e. 4 additional encoded symbols), Defender will enter the state N_1 (see Remark 3.10) and win.

The above defined strategy is winning for Defender and so $I \leqslant_m R$. $\quad\square$

Remark 3.12. Lemma 3.11 demonstrates the point where we need our special encoding of configurations using the alternation of π and σ actions together with side-branches to represent the symbols in the configurations. If the configurations were encoded directly as sequences of symbols on a linear path, the construction would not work. Indeed, the must path of alternating σ and π actions in the process R is necessary to ensure that the bad path entered in the left-hand side implementation I is indeed realizable. This path cannot be replaced by a linear path of must transitions containing directly the symbols of the configurations because the sequence of $n-5$ symbols in the middle of the bad sequence would require exponentially large MTS to capture all such possible sequences explicitly and the reduction would not be polynomial.

Fig. 6. Full specification of the process R.

Let us now finish the definition of the process R. Note that in ALBA even rejecting computation trees can still contain several correct computation paths ending in accepting configurations. We can only assume that during any universal branching in a rejecting tree, *at least one* of the two possible successors forms a rejecting branch. The process R must thus have the possibility to discard the possibly correct computation branch in universal branching and it suffices to make sure that the computation will continue with only one of the branches.

So in order to finish the construction of R we add an additional fragment to R as depicted in Fig. 6 (it is the part below R that starts with branching to U_1 and U_2).

The construction of the process R is now finished (recall that the part of the construction going from R to the right is repeated for any bad path of the machine $\mathcal{M}$). Because the newly added part of the construction does not use any must transitions, it does not restrict the set of implementations and hence Lemma 3.11 still holds. The following two lemmas show that the added part of the construction correctly handles the universal branching.

Lemma 3.13. *Let I be a tree implementation of L which is not, even after removing any of its branches, a code of any accepting computation tree of $\mathcal{M}$ on the input w. Then $I \leqslant_m R$.*

Proof. For any tree implementation $I \leqslant_m L$, which cannot be pruned to a code of any accepting computation tree, we extend the strategy for Defender from Lemma 3.11. The first rule remains unchanged.

Otherwise, if Attacker chose a π transition leading to a state in the left-hand side process with the action $\forall$ enabled, we know (from the definition of the process L, see Proposition 3.7 parts 2 and 5) that a σ action must follow and then there are at least two branches under π, one of them enabling the action 1 and the other one the action 2; because I is not an encoding of an accepting computation tree (not even after being pruned out) either (i) all subtrees beginning with the action 1 are either rejecting or incorrect or (ii) all subtrees beginning with the action 2 are either rejecting or incorrect. In case (i) Defender responds by entering U_1, in case (ii) by entering U_2. After Attacker plays the above mentioned action σ followed by one of the π actions, Defender responds by the σ action leaving (i) U_1 or (ii) U_2 and then by one of the two π actions so that the Attacker's move on the left-hand side is correctly matched. In case (i) if Attacker plays the π move with a following branch under 2, or in case (ii) if Attacker plays the π move with the following branch under 1, Defender will aim at entering the process U' and after the following σ move wins as any implementation is a refinement of U'. Hence Attacker is forced to choose, in case (i), some first branch and, in case (ii), some second branch and after the necessary action σ the game continues from a configuration where the right-hand side process is again in R.

Finally, if none of the previous cases applies, Defender simply mimics any Attacker's move by looping in R. Note that in this case Attacker cannot have played the action q_{acc} because we assume that the implementation I, even after removing any of its branches, does not encode any accepting computation tree.

As the above defined strategy is winning for Defender, we conclude that $I \leqslant_m R$. □

Lemma 3.14. *Let $\mathcal{T}$ be an accepting computation tree of an ALBA $\mathcal{M}$ on the input w. Then $\mathrm{code}(\mathcal{T}) \not\leqslant_m R$.*

Proof. Note that in $\mathcal{T}$ every branch ends with a configuration containing the accepting state q_{acc}. It is so clear that Attacker can easily win by playing repeatedly the transition π followed by the transition σ in the tree $\mathrm{code}(\mathcal{T})$. Defender is forced to stay in the state R because any branch in the tree is correct and hence Defender cannot "escape" by playing the π move to the state V_1 for any bad path (should Defender play like this, Attacker would switch the sides and play the must sequence of π and σ transitions until Defender is proven to be cheating and Attacker wins).

Fig. 7. Modified process L'.

The only situation when Defender can play a π move going to the state U_1 or U_2 is when Attacker (in the left-hand side process) is inside a code of a configuration following a universal configuration and after he played π the next label is $\forall$. In case that Defender entered U_1, Attacker simply continues on the left-hand side by taking the first successor configuration, and in case that Defender entered U_2, Attacker chooses the second successor configuration. After the sequence of one π and one σ move in the left-hand side process, Defender is forced to return to the state R (otherwise Attacker wins by playing the action 1 respectively 2 in the left-hand side process). Eventually, after reaching an accepting leaf configuration in $\mathcal{T}$, Attacker will play the action q_{acc} in the left-hand side process to which Defender has no answer from the process R. As we have described Attacker's winning strategy, we conclude that $\text{code}(\mathcal{T}) \not\leqslant_m R$. $\square$

3.5. Summary

We can now combine the facts about the constructed systems L and R in the following theorem.

Theorem 3.15. *An ALBA $\mathcal{M}$ accepts an input w iff $L \not\leqslant_t R$.*

Proof. If $\mathcal{M}$ accepts the input w then clearly it has an accepting computation tree $\mathcal{T}$. By Lemma 3.8 $\text{code}(\mathcal{T}) \leqslant_m L$ and by Lemma 3.14 $\text{code}(\mathcal{T}) \not\leqslant_m R$. This implies that $L \not\leqslant_t R$.

On the other hand, if $\mathcal{M}$ does not accept w then none of the tree implementations of L represents a code of an accepting computation tree of $\mathcal{M}$ on w. By Lemma 3.13 this means that any tree I such that $I \leqslant_m L$ satisfies that $I \leqslant_m R$ and hence $L \leqslant_{tt} R$ which is by Lemma 3.5 equivalent to $L \leqslant_t R$. $\square$

Corollary 3.16. *The problem of checking thorough refinement on finite modal transition systems is EXPTIME-hard.*

3.6. Deterministic and fixed process L

We can strengthen the above hardness result by adapting the described reduction to the situation where the left-hand side system is deterministic (meaning that for every state and every action there is at most one outgoing may transition) and of a fixed size.

Theorem 3.17. *The problem of checking thorough refinement on finite modal transition systems is EXPTIME-hard even if the left-hand side system is deterministic and of a fixed size.*

Proof. The MTS L of the previous construction is nondeterministic, as the process M reachable from L has several outgoing π transitions. Its size also depends on the size of the acceptance problem, as it contains the code of the initial configuration of the ALBA.

We now show a different construction with a deterministic fixed-size left-hand side process. This construction will be a modification of the original one. To make the explanation easier to follow, we split the modification into two steps. In the first step, we modify L and R into modal transition systems L' and R' such that L' is deterministic and $L' \not\leqslant_t R'$ if and only if $\mathcal{M}$ accepts w. In the second step, we modify L' and R' into L'' and R'' such that L'' is both deterministic and of a fixed size and still $L'' \not\leqslant_t R''$ if and only if $\mathcal{M}$ accepts w.

We introduce L' in Fig. 7. The modification is as follows. The process M is changed into M' which has only one outgoing may transition to N' labeled with π. The process N' then has outgoing may transitions for each symbol of $\varXi$ (including $\forall$), leading to a process with no transitions. It is easy to see that $L \leqslant_m L'$ and thus every implementation of the original L is also an implementation of L'. Therefore, if there exists an implementation of L that is an encoding of a correct accepting tree, L' also possesses such implementation. However, L' has more implementations representing incorrect computation trees than L. They are of the following types (we only consider tree implementations here):

(a) tree implementations that possess at least one branch with a sequence of the form $J \xrightarrow{\pi} J' \xrightarrow{\sigma} J''$ without any transitions labeled with symbols from $\varXi$ outgoing from J',

(b) tree implementations possessing at least one branch with a sequence of the form $J \xrightarrow{\pi} J' \xrightarrow{\sigma} J''$ with at least two transitions outgoing from J' labeled with different symbols of $\varXi$, i.e. $J' \xrightarrow{x}$ and $J' \xrightarrow{y}$ where $x \neq y$, $x, y \in \varXi$,

for all $a \in \varXi$

Fig. 8. Process L''.

(c) tree implementations in which a process K is reachable such that $K \xrightarrow{\forall}$ and neither $K \xrightarrow{\sigma} \xrightarrow{\pi} \xrightarrow{1}$ nor $K \xrightarrow{\sigma} \xrightarrow{\pi} \xrightarrow{2}$, and

(d) tree implementations in which a process K is reachable such that $K \xrightarrow{\forall}$ and either $K \xrightarrow{\sigma} \xrightarrow{\pi} \xrightarrow{1}$ or $K \xrightarrow{\sigma} \xrightarrow{\pi} \xrightarrow{2}$ but not both.

We first discuss the implementations of the forms (b), (c) and (d) and claim that they are already implementations of R of the original construction.

(b) If the implementation contains a branch with a sequence $J \xrightarrow{\pi} J' \xrightarrow{\sigma} J''$ such that J' has two different outgoing $\varXi$ transitions then surely one of these transitions is a part of a code of a bad path. Therefore, such implementation is clearly matched by R.

(c) Similarly, a sequence containing $\forall$ followed by something else than 1 or 2 is a part of a bad path.

(d) Suppose the implementation has a reachable process K such that $K \xrightarrow{\forall}$ and only one of $K \xrightarrow{\sigma} \xrightarrow{\pi} \xrightarrow{1}$ or $K \xrightarrow{\sigma} \xrightarrow{\pi} \xrightarrow{2}$ holds. Then clearly K is an implementation of R as it can be matched by one of the two paths from R to U'. Therefore also the implementation that has a path to K has to be an implementation of R.

However, the implementations of the form (a) may not be implementations of R. We thus modify R into R' by adding an extra outgoing transition that matches such implementations. The process R' is therefore equivalent to R with the following addition: $R' \xdashrightarrow{\pi} S'$, $S' \xdashrightarrow{\sigma} T'$, $T' \xdashrightarrow{\varXi} T'$, where S' and T' are new processes. Clearly this modification does not allow R' to possess as an implementation an encoding of a correct computation tree, and thus $L' \not\leqslant_t R'$ if and only if $\mathcal{M}$ accepts w.

We now proceed with the second step. Clearly, there are two places in the system for L' (and also L) that are dependent on the size of the original ALBA $\mathcal{M}$ and the word w to be accepted.

First, it is the encoding of the initial configuration (the path from L to M) and second, it is the number of outgoing transitions from N' which is dependent on the number of ALBA's states.

The interesting part is that of the encoding of the initial configuration. We change the left-hand side process L' into L'' by eliminating the path representing the initial configuration. (Thus, L'' is equivalent to M' of the previous construction.) This new process is depicted in Fig. 8. Clearly, L'' has more implementations than the process L', namely those that have a branch which starts with something different than the encoding of the initial configuration. Thus, we need to extend the right-hand side process so that it also admits this kind of implementations.

The new process R'' is built as follows. There are paths encoding the incorrect initial configurations, one for each position of the configuration, i.e. there is a path representing that an error (incorrect symbol in the initial configuration) happens at the first symbol of the configuration, a path representing an error at the second symbol, etc. Similarly to the previous construction we also include may transitions for escaping to universal process, to allow for arbitrary branching in the implementations. To this we furthermore add all transitions the process R' has, including transitions to R' itself. Note that this implies that $R' \leqslant_m R''$. Formally, the process R'' is defined as follows (see Fig. 9 for illustration):

$$R'' \xdashrightarrow{\pi} S_{k1} \quad \text{for } 1 \leqslant k \leqslant n+5$$

$$R'' \xdashrightarrow{x} X \quad \text{whenever } R' \xdashrightarrow{x} X$$

$$S_{k\ell} \xrightarrow{\sigma} T_{k\ell} \quad \text{for } 1 \leqslant \ell < k \leqslant n+5$$

$$S_{k\ell} \xrightarrow{i_\ell} X_{k\ell} \quad i_\ell \text{ is the } \ell\text{th symbol of the initial configuration}$$

$$S_{k\ell} \xdashrightarrow{x} U'' \quad \text{for all } x \in \varXi'$$

$$S_{kk} \xdashrightarrow{\sigma} U''$$

$$S_{kk} \xdashrightarrow{x} X_{kk} \quad \text{for all } x \in \varXi \setminus \{i_k\}$$

$$T_{k\ell} \xrightarrow{\pi} S_{k(\ell+1)} \quad \text{for } 1 \leqslant \ell < k \leqslant n+5$$

$$U'' \xdashrightarrow{x} U'' \quad \text{for all } x \in \varXi'$$

Fig. 9. Process R'' that accepts all branches not starting with the initial configuration $\vdash \exists * \ldots \dashv$ or entering R'.

We now need to show two things. First, that any implementation of L'' that is not an encoding of a correct accepting computation tree even after pruning is an implementation of R'', and second, that the encoding of a correct accepting computation tree is not an implementation of R''.

Regarding the first claim let us fix an implementation I of L'' that is not a correct accepting computation tree. Consider transitions from I placed on branches that do not start with an encoding of the initial configuration. These are matched by the newly created transitions from R'' to the paths encoding the branches that begin with an error. Let us further consider all transitions from I that are placed only on branches beginning with correct encodings of the initial configuration. The corresponding subtrees are then implementations of R'. But we know that $R' \leqslant_m R''$, so the whole I is also an implementation of R''.

The second claim then easily follows from the fact that the encoding of a correct accepting computation tree is neither an implementation of R', nor can be matched by any of the bad paths added to R'', as it clearly has to start with a correct initial configuration. Thus we have that the given ALBA accepts the input w if and only if $L'' \not\leqslant_t R''$.

What remains is to cope with the number of states of the original ALBA so that $\varXi$ is of constant size. That can be dealt with in a straightforward manner by encoding the states with binary strings of two symbols, say q_0 and q_1. This changes the compatible 10-tuples into compatible $(2 \cdot \lceil \log_2 |Q| \rceil + 8)$-tuples, but this change does not affect the construction heavily. □

4. Thorough refinement is in EXPTIME

In this section we provide a direct algorithm for deciding thorough refinement between MTSs in exponential time. Given two processes A and B over some finite-state MTSs, the algorithm will decide if there exists an implementation I that implements A but not B, i.e. $I \leqslant_m A$ and $I \not\leqslant_m B$.

For a modal transition systems B, we introduce the syntactical notation $\overline{B}$ to denote the semantical complement of B, i.e. $I \leqslant_m \overline{B}$ iff $I \not\leqslant_m B$. Our algorithm now essentially checks for consistency (existence of a common implementation) between A and $\overline{B}$ with the outcome that they are consistent if and only if $A \not\leqslant_t B$. Note that we never use the system $\overline{B}$ in any of the constructions, it is introduced purely for notational reasons so that $I \leqslant_m \overline{B}$ is an abbreviation for $I \not\leqslant_m B$.

In general, we shall check for consistency of sets of the form $\{A, \overline{B}_1, \ldots, \overline{B}_k\}$ in the sense of existence of an implementation I such that $I \leqslant_m A$ but $I \not\leqslant_m B_i$ for all $i \in \{1, \ldots, k\}$. Before the full definition is given, let us get some intuition by considering the case of consistency of a simple pair $A, \overline{B}$. During the arguments, we shall use CCS-like constructs (summation and action-prefixing) for defining implementations.

Clearly, if for some B' with $B \xrightarrow{a} B'$ and for all A_i with $A \xrightarrow{a} A_i$ we can find an implementation I_i implementing A_i but not B' (i.e. we demonstrate consistency between all the pairs $A_i, \overline{B'}$), we can claim consistency between A and $\overline{B}$: as

a common implementation I simply take $H + \sum_i a.I_i$, where H is some arbitrary implementation of A with all a-derivatives removed. Clearly, $I \leqslant_m A$ but $I \not\leqslant_m B$. In this case of a must transition in B that is unrealized in A, we will store this information by setting $B' \in$ unrealized$_a$.

We may also conclude consistency of A and $\overline{B}$, if for some A' with $A \dashrightarrow^a A'$, we can find an implementation I' of A', which is not an implementation of any B' where $B \dashrightarrow^a B'$. Here a common implementation would simply be $H + a.I'$ where H is an arbitrary implementation of A. However, in this case we will need to determine the consistency of the set $\{A'\} \cup \{\overline{B'} \mid B \dashrightarrow^a B'\}$ which is in general not a simple pair. In this case of a may transition in A that is not allowed in B, we will store this information by setting $B \in$ disallowed$_a$.

Thus consistency can be shown by either of the two ways and for any letter of the alphabet. We formalize this in the following definition.

Definition 4.1. Let $M = (P, \dashrightarrow, \longrightarrow)$ be an MTS over the action alphabet Σ. The set of *consistent* sets of the form $\{A, \overline{B}_1, \ldots, \overline{B}_k\}$, where $A, B_1, \ldots, B_k \in P$, is the smallest set Con such that $\{A, \overline{B}_1, \ldots, \overline{B}_k\} \in$ Con whenever $k = 0$ or there are sets of processes unrealized$_a$ and disallowed$_a$ for every $a \in \Sigma$, such that for every $B \in \{B_1, \ldots, B_k\}$ either $B \in$ disallowed$_a$ for some $a \in \Sigma$, or $B \longrightarrow^a B'$ for some $a \in \Sigma$ and $B' \in$ unrealized$_a$, so that for every $a \in \Sigma$

1. for all $A \longrightarrow^a A'$ we have $\{A'\} \cup \{\overline{B'} \mid B' \in$ unrealized$_a\} \in$ Con, and
2. for all $B \in$ disallowed$_a$ there is $A \dashrightarrow^a A'$ with $\{A'\} \cup \{\overline{B'} \mid B \dashrightarrow^a B'\} \cup \{\overline{B'} \mid B' \in$ unrealized$_a\} \in$ Con.

Lemma 4.2. *Given processes $A, B_1, \ldots, B_k$ of some finite MTS, there exists an implementation I such that $I \leqslant_m A$ and $I \not\leqslant_m B_i$ for all $i \in \{1, \ldots, k\}$ if and only if $\{A, \overline{B}_1, \ldots, \overline{B}_k\} \in$ Con.*

Proof. We prove both directions by induction.

'*If*' part (soundness of the construction). Because Con is defined as the smallest set, let by $\text{Con}_0, \text{Con}_1, \text{Con}_2, \ldots$ denote the nondecreasing sequence of sets according to in which round the elements (consistent sets) where added to Con. So Con_0 contains exactly all the consistent sets of the form $\{A\}$, Con_1 contains all the consistent sets that were added to Con_0 in one iteration of the definition, etc.

We prove by induction on n that whenever $\{A, \overline{B}_1, \ldots, \overline{B}_k\} \in \text{Con}_n$ then there exists an implementation I such that $I \leqslant_m A$ and $I \not\leqslant_m B_i$ for all $i \in \{1, \ldots, k\}$.

The base case $n = 0$ is trivial. For the induction step assume that $\{A, \overline{B}_1, \ldots, \overline{B}_k\} \in \text{Con}_{n+1}$. Then for every $a \in \Sigma$ there exist sets disallowed$_a$ and unrealized$_a$ satisfying the conditions of Definition 4.1.

From the first clause, it follows from the induction hypothesis that for every $A \longrightarrow^a A'$ there exists $I_{A'}$ such that $I_{A'} \leqslant_m A'$ and $I_{A'} \not\leqslant_m B'$ for all $B' \in$ unrealized$_a$. Similarly, from the second clause, it follows from the induction hypothesis that for every $B \in$ disallowed$_a$ there exists $A \dashrightarrow^a A_B$ and an implementation I_{A_B} such that $I_{A_B} \leqslant_m A_B$ and $I_{A_B} \not\leqslant_m B'$ for all B' with $B \dashrightarrow^a B'$ or $B' \in$ unrealized$_a$. Now we define an implementation I witnessing the correctness of $\{A, \overline{B}_1, \ldots, \overline{B}_k\} \in$ Con:

$$I \equiv \sum_{a \in \Sigma} \left(\sum_{A'.A \overset{a}{\longrightarrow} A'} a.I_{A'} + \sum_{B \in \text{disallowed}_a} a.I_{A_B} \right)$$

We shall prove that $I \leqslant_m A$ and $I \not\leqslant_m B_i$ for all $i \in \{1, \ldots, k\}$. Let us first establish $I \leqslant_m A$. Assume that $A \longrightarrow^a A'$, then $I \longrightarrow^a I_{A'}$ will provide the match. For the other direction assume that $I \dashrightarrow^a I_X$, then $A \dashrightarrow^a X$ provides the match. To see that $I \not\leqslant_m B_i$, we distinguish two cases. Either $B_i \in$ disallowed$_a$ for some $a \in \Sigma$ and then $I \dashrightarrow^a I_{A_{B_i}}$ cannot be matched by any may transition $B_i \dashrightarrow^a B'$; or there is $B_i \longrightarrow^a B' \in$ unrealized$_a$ for some $a \in \Sigma$ that can be matched by neither $I \longrightarrow^a I_{A'}$ nor $I \longrightarrow^a I_{A_{B_i}}$.

'*Only-if*' part (completeness of the construction). Let us define $I \leqslant_m^n S$ if either $n = 0$ or (i) whenever $I \dashrightarrow^a I'$ then $S \dashrightarrow^a S'$ with $I' \leqslant_m^{n-1} S'$ and (ii) whenever $S \longrightarrow^a S'$ then $I \longrightarrow^a I'$ with $I' \leqslant_m^{n-1} S'$. Hence the relation $\leqslant_m^n$ is a natural generalization of the classical bisimulation approximations to modal refinement, and clearly (on finite MTS) we have that $I \leqslant_m S$ iff $I \leqslant_m^n S$ for all n.

We prove by induction on n that whenever there exists an implementation I such that $I \leqslant_m A$ and $I \not\leqslant_m^n B_i$ for all $i \in \{1, \ldots, k\}$ then $\{A, \overline{B}_1, \ldots, \overline{B}_k\} \in$ Con.

The base case $n = 0$ is trivial as in this case $k = 0$ and hence $\{A, \overline{B}_1, \ldots, \overline{B}_k\} = \{A\} \in$ Con. For the induction step assume that $I \leqslant_m A$ and $I \not\leqslant_m^{n+1} B_i$ for some I. We define the following sets of processes for every $a \in \Sigma$:

$$\text{disallowed}_a = \left\{ B \in \{B_1, \ldots, B_k\} \,\middle|\, \exists I \dashrightarrow^a I'. \forall B \dashrightarrow^a B'. I' \not\leqslant_m^n B' \right\}$$

$$\text{unrealized}_a = \left\{ B' \,\middle|\, \exists B \in \{B_1, \ldots, B_k\}. \exists B \longrightarrow^a B'. \forall I \longrightarrow^a I'. I' \not\leqslant_m^n B' \right\}$$

Note that due to the definition of the modal refinement, these sets satisfy the conditions of Definition 4.1 once we establish its two clauses.

As for the first clause, for every $A \xrightarrow{a} A'$ there is $I \xrightarrow{a} I'$ where $I' \leqslant_m A'$ and by definition of unrealized$_a$ also $I' \not\leqslant_m B'$ for $B' \in$ unrealized$_a$. Thus by induction hypothesis $\{A'\} \cup \{\overline{B'} \mid B' \in$ unrealized$_a\} \in$ Con. As for the second clause, for every $B \in$ disallowed$_a$ we have $I \dashrightarrow^a I'$ (hence $I' \leqslant_m A'$ for some $A \dashrightarrow^a A'$) with $I' \not\leqslant_m^n B'$ for all $B \dashrightarrow^a B'$ by definition of disallowed$_a$. Since also $I' \not\leqslant_m^n B' \in$ unrealized$_a$ using the definition of unrealized$_a$, induction hypothesis guarantees $\{A'\} \cup \{\overline{B'} \mid B \dashrightarrow^a B'\} \cup \{\overline{B'} \mid B' \in$ unrealized$_a\} \in$ Con.

It follows that the sets disallowed$_a$ and unrealized$_a$ provide the evidence required by the definition to conclude that $\{A, \overline{B}_1, \ldots, \overline{B}_k\} \in$ Con. $\square$

Computing the collection of consistent sets $\{A, \overline{B}_1, \ldots, \overline{B}_k\}$ over an MTS $(P, \dashrightarrow, \longrightarrow)$ may be done as a simple (least) fixed-point computation. The running time is polynomial in the number of potential sets of the form $\{A, \overline{B}_1, \ldots, \overline{B}_k\}$ where $A, B_1, \ldots, B_k \in P$, hence it is exponential in the number of states of the underlying MTS. This gives an exponential time algorithm to check for thorough refinement.

Theorem 4.3. *The problem of checking thorough refinement on finite modal transition systems belongs to EXPTIME.*

Example 4.4. Consider S and T from Fig. 1. We have already mentioned in Section 2 that $S \leqslant_t T$. To see this, we will attempt (and fail) to demonstrate consistency of $\{S, \overline{T}\}$ according to Definition 4.1, which essentially asks for a finite tableau to be constructed. Now, in order for $\{S, \overline{T}\}$ to be concluded consistent, we have to establish consistency of $\{S_1, \overline{T}_1, \overline{T}_2\}$—as T has no must transitions the only choice is unrealized$_a = \emptyset$ and thus disallowed$_a = \{T\}$. Now, to establish consistency of $\{S_1, \overline{T}_1, \overline{T}_2\}$ both unrealized$_a = \emptyset$ and unrealized$_a = \{T\}$ are possibilities. In the former case disallowed$_a = \{T_1, T_2\}$ and in the latter case $T_2 \in$ disallowed$_a$. However, in both cases the requirement will be that $\{S, \overline{T}\}$ must be consistent. Given this cyclic dependency together with the minimal fixed-point definition of Con it follows that $\{S, \overline{T}\}$ is *not* consistent, and hence that $S \leqslant_t T$.

Example 4.5. Consider S and U from Fig. 1. Here $S \not\leqslant_t U$ clearly with $I = a.0$ as a witness implementation. Let us demonstrate consistency of $\{S, \overline{U}\}$. Choosing unrealized$_a = \emptyset$ and disallowed$_a = \{U\}$, this will follow from the consistency of $\{S_1, \overline{U}_1\}$. To conclude this, note that unrealized$_a = \{U_1\}$ and disallowed$_a = \emptyset$ will leave us with the empty collection of sets—as S_1 has no must transitions—all of which are obviously consistent.

Note that in the case of B being deterministic, we only need to consider pairs of the form $\{A, \overline{B}\}$ for determining consistency. This results in a polynomial time algorithm (see also [6] for an alternative proof of this fact). Similarly, if the process B is of a constant size, our algorithm runs in polynomial time as well.

Corollary 4.6. *The problem of checking thorough refinement on finite modal transition systems with the right-hand side system deterministic or of fixed size belongs to P.*

To conclude, by Theorem 4.3 and Corollary 3.16 we get our main result.

Theorem 4.7. *The problem of checking thorough refinement on finite modal transition systems is EXPTIME-complete.*

5. Conclusion

We proved that the problem of checking the thorough refinement relation between two finite-state modal transition systems is EXPTIME-complete. This result completes related complexity results achieved in [3] as the thorough refinement relations on both modal and mixed (where the must transition relation is not necessarily included in the may transition relation) specifications, the common implementation problems on modal and mixed specifications, as well as the consistency problem on mixed specifications are now all EXPTIME-complete. Our EXPTIME-hardness result is proved by reduction from the acceptance problem for alternating linear bounded automata because the problems of consistency and common implementation mentioned above did not seem to provide suitable starting problems for the reduction.

The fact that the thorough refinement relation is computationally hard means that the relation of modal refinement is more suitable for practical purposes, even though it describes a less desirable notion of syntactic refinement. On the other hand, much of the recent work in the area focuses to a large extent on deterministic specifications (see e.g. [12,13]) and here the two notions of refinement coincide. A detailed study of computational complexity of the problems on deterministic modal transition systems is provided in [6].

Acknowledgments

References

[1] A. Antonik, M. Huth, K.G. Larsen, U. Nyman, A. Wasowski, 20 years of modal and mixed specifications, Bull. Eur. Assoc. Theor. Comput. Sci. EATCS 95 (2008) 94–129.
[2] A. Antonik, M. Huth, K.G. Larsen, U. Nyman, A. Wasowski, Complexity of decision problems for mixed and modal specifications, in: Proc. of FOSSACS'08, in: Lecture Notes in Comput. Sci., vol. 4962, Springer, 2008, pp. 112–126.
[3] A. Antonik, M. Huth, K.G. Larsen, U. Nyman, A. Wasowski, Modal and mixed specifications: key decision problems and their complexities, Math. Structures Comput. Sci. 20 (1) (2010) 75–103.
[4] J.L. Balcazar, J. Gabarró, M. Santha, Deciding bisimilarity is P-complete, Form. Asp. Comput. 4 (6A) (1992) 638–648.
[5] S.S. Bauer, P. Mayer, A. Schroeder, R. Hennicker, On weak modal compatibility, refinement, and the MIO workbench, in: Proc. of TACAS'10, in: Lecture Notes in Comput. Sci., vol. 6015, Springer, 2010, pp. 175–189.
[6] N. Beneš, J. Křetínský, K.G. Larsen, J. Srba, On determinism in modal transition systems, Theoret. Comput. Sci. 410 (41) (2009) 4026–4043.
[7] N. Bertrand, S. Pinchinat, J.-B. Raclet, Refinement and consistency of timed modal specifications, in: Proc. of LATA'09, in: Lecture Notes in Comput. Sci., vol. 5457, Springer, 2009, pp. 152–163.
[8] N. D'Ippolito, D. Fischbein, M. Chechik, S. Uchitel, MTSA: The modal transition system analyser, in: Proc. of ASE'08, IEEE, 2008, pp. 475–476.
[9] N. D'Ippolito, D. Fischbein, H. Foster, S. Uchitel, MTSA: Eclipse support for modal transition systems construction, analysis and elaboration, in: Proc. of (ETX'07), ACM, 2007, pp. 6–10.
[10] H. Fecher, H. Schmidt, Comparing disjunctive modal transition systems with an one-selecting variant, J. Log. Algebr. Program. 77 (1–2) (2008) 20–39.
[11] P. Godefroid, M. Huth, R. Jagadeesan, Abstraction-based model checking using modal transition systems, in: Proc. CONCUR'01, in: Lecture Notes in Comput. Sci., vol. 2154, Springer, 2001, pp. 426–440.
[12] T.A. Henzinger, J. Sifakis, The embedded systems design challenge, in: Proceedings of the 14th International Symposium on Formal Methods (FM'06), in: Lecture Notes in Comput. Sci., vol. 4085, Springer-Verlag, 2006, pp. 1–15.
[13] T.A. Henzinger, J. Sifakis, The discipline of embedded systems design, IEEE Comput. 40 (10) (2007) 32–40.
[14] M. Huth, R. Jagadeesan, D.A. Schmidt, Modal transition systems: A foundation for three-valued program analysis, in: Proc. of ESOP'01, in: Lecture Notes in Comput. Sci., vol. 2028, Springer, 2001, pp. 155–169.
[15] P.C. Kanellakis, S.A. Smolka, CCS expressions, finite state processes, and three problems of equivalence, Inform. and Comput. 86 (1) (1990) 43–68.
[16] K.G. Larsen, U. Nyman, A. Wasowski, On modal refinement and consistency, in: Proc. of CONCUR'07, in: Lecture Notes in Comput. Sci., vol. 4703, Springer, 2007, pp. 105–119.
[17] K.G. Larsen, B. Thomsen, A modal process logic, in: Proc. of LICS'88, IEEE, 1988, pp. 203–210.
[18] S. Nanz, F. Nielson, H.R. Nielson, Modal abstractions of concurrent behaviour, in: Proc. of SAS'08, in: Lecture Notes in Comput. Sci., vol. 5079, Springer, 2008, pp. 159–173.
[19] R. Paige, R. Tarjan, Three partition refinement algorithms, SIAM J. Comput. 16 (6) (1987) 973–989.
[20] J.-B. Raclet, Residual for component specifications, in: Proc. of the 4th International Workshop on Formal Aspects of Component Software, 2007.
[21] J.-B. Raclet, E. Badouel, A. Benveniste, B. Caillaud, R. Passerone, Why are modalities good for interface theories?, in: Proc. of ACSD'09, IEEE Computer Society, 2009, pp. 119–127.
[22] Z. Sawa, P. Jančar, Behavioural equivalences on finite-state systems are PTIME-hard, Comput. Inform. 24 (5) (2005) 513–528.
[23] S. Shoham, O. Grumberg, Compositional verification and 3-valued abstractions join forces, Inform. and Comput. 208 (2) (2010) 178–202.
[24] M. Sipser, Introduction to the Theory of Computation, Course Technology, 2006.
[25] S. Uchitel, M. Chechik, Merging partial behavioural models, in: Proc. of FSE'04, ACM, 2004, pp. 43–52.
[26] O. Wei, A. Gurfinkel, M. Chechik, Mixed transition systems revisited, in: Proc. of VMCAI'09, in: Lecture Notes in Comput. Sci., vol. 5403, Springer, 2009, pp. 349–365.