

On Modal Refinement and Consistency

Kim G. Larsen, Ulrik Nyman, and Andrzej Wąsowski

Department of Computer Science, Aalborg University, Denmark
{kgl,ulrik,wasowski}@cs.aau.dk

Abstract. Almost 20 years after the original conception, we revisit several fundamental questions about modal transition systems. First, we demonstrate the incompleteness of the standard modal refinement using a counterexample due to Hüttel. Deciding any refinement, complete with respect to the standard notions of implementation, is shown to be computationally hard (co-NP hard). Second, we consider four forms of consistency (existence of implementations) for modal specifications. We characterize each operationally, giving algorithms for deciding, and for synthesizing implementations, together with their complexities.

1 Background and Overview

Modal transition systems (MTSs) are a generalization of labeled transition systems (LTSs). Similarly to LTSs modal transition systems use labeled transitions between states to model behaviors. Unlike LTSs, they distinguish allowed and required behaviors (over- and under-approximations), which makes them a suitable semantic model for abstraction in program analysis and verification.

MTSs, originally introduced by Larsen and Thomsen almost 20 years ago [1], have since been applied in program analysis [2, 3], model checking [4, 5], verification [6, 7], equation solving [8], interface theories [9], software product lines [9, 10] and model merging [11, 12]. Foundational work on modal transition systems included extensions to modal hybrid systems [13], timed modal specifications [14–16] and variants of disjunctive MTSs [8, 17, 18]. Surprisingly though, several fundamental questions about the theory of MTSs have never been addressed.

Refinement relations for modal transition systems are defined contravariantly. If S refines T then all allowed behaviors of S need to be allowed in T , while all required behaviors of T need also be required by S . An *implementation* is an MTS that has been completely specified, i.e. all its allowed behavior is also required, leaving no further choice for refinement. One fundamental issue for a modal refinement is to see whether it characterizes the inclusion of implementation sets thoroughly: can one for an MTS S refining an MTS T imply that all implementations of S are also implementations of T ? And vice-versa?

Standard modal refinement is sound, but not complete in this sense. Meaning that here exist MTSs for which implementation inclusion holds, but which do not refine each other. We show that deciding any sound and complete refinement, preserving the set of implementations of standard modal refinement or weak

modal refinement is co-NP hard. We conjecture the same for may-weak modal refinement [9] and branching refinement [10].

Modal transition systems of [1] are *syntactically consistent*, meaning that any required transition must also be allowed. This effectively disallows reasoning about inconsistencies, which is necessary for proper treatment of logical connectives in the context of modal transition systems (for example one would like to be able to express a modal transition system expressing a conjunction of two other MTSs that represent contradictory specifications). On the other hand, in [9], we have observed that other, more behavioral, notions of consistency might be useful. We have shown that systems that are *observationally consistent* with respect to some set of hidden actions, can be decomposed using parallel decomposition. We used this observation to build a product line theory in which modal transition systems play the role of behavioral variability models.

We believe that consistency should be decoupled from the basic definition of a modal transition system. In our opinion understanding a notion of consistency requires relating it to a notion of satisfiability, as typically done in logics. For example: a propositional formula is consistent if there exists a truth assignment on which the formula evaluates to true. In our context, modal transition systems play the role of formulæ, truth assignments are concrete implementations, and a refinement preorder is our satisfaction relation. Consequently, instead of proposing ad hoc *criteria* for consistency, we define consistency of a specification *semantically* as existence of a concrete implementation refining it.

Altogether we discuss four modal refinements and their induced consistencies. For each of these we define consistency semantically and find a computable criterion (a consistency relation) for deciding it. Then we study the complexity of consistency and the criterion. The results are summarized in Table 1.

Our choice of refinements and consistencies for this study is driven by existing work. We choose one known consistency (syntactic consistency) that have not been characterized using a refinement, and three known refinements (strong, may-weak and weak modal refinement) for which the related notions of consistency had never been formulated. However, we believe that consistency is not only of theoretical interest. Inconsistencies in specifications typically indicate modeling errors and thus procedures for detecting them find use in tools.

The contents of this paper are: the definition of modal transition systems and their refinement (Section 2), complexity analysis of completeness of this refinement (Section 3), a discussion of consistency notions induced by four modal refinements (Sections 4–7), a summary and a list of open problems (Section 8).

Table 1. Summary of consistency-related results.

Modal refinement	Consistency	Lower bound	Upper bound	Section
syntactic	syntactic consistency [1]	linear time	linear time	4
strong [1]	strong consistency	NP-hard	exponential time	5
weak [19]	weak consistency	NP-hard	exponential time	6
may-weak [9]	may-weak consistency	NP-hard	exponential time	7

2 Modal Transition Systems

We introduce the basics following Larsen and Thomsen [1]. Assume a global set of actions act and write act^τ for $act \cup \{\tau\}$, where τ is a distinct internal action, such that $\tau \notin act$. A modal transition system is a triple $S = (states_S, \longrightarrow^S, \dashrightarrow^S)$, where $states_S$ is a set of states, also known as specifications [1] or processes. Then $\longrightarrow^S \subseteq states_S \times act^\tau \times states_S$ is a must-transition relation representing required transitions, and $\dashrightarrow^S \subseteq states_S \times act^\tau \times states_S$ is a may-transition relation representing allowed transitions.

In general the sets of states and transitions may be infinite, but we restrict ourselves to finite state systems with finite sets of actions in this paper. For simplicity we write $s \xrightarrow{a}^S s'$ iff $(s, a, s') \in \longrightarrow^S$, and $s \dashrightarrow^S s'$ iff $(s, a, s') \in \dashrightarrow^S$.

Larsen and Thomsen originally designed modal transition systems to be syntactically consistent meaning that all required transitions are also allowed: $\longrightarrow^S \subseteq \dashrightarrow^S$. Already in [14] Larsen lifts this restriction, with the argument that any sufficiently expressive specification language needs to be able to specify inconsistent specifications. This means that our transition systems are very much like mixed transition systems of Dams [20]. In Section 3 we follow the syntactic consistency requirement, while we relax it in later sections, generalizing the notion of consistency to strong and weak behavioral preorders. Regardless whether the consistency assumption is in place or not, we always separate the two transition relations explicitly to avoid confusion. A solid arrow represents just a must transition, without the possible related may transition. We draw both arrows when talking about a syntactically consistent must transition.

A modal transition system I is an *implementation* when the two transition relations coincide, $\longrightarrow^I = \dashrightarrow^I$. We use capital I to denote implementations and always state explicitly whenever a modal transition system is an implementation.

The following is the standard notion of strong refinement for modal transition systems introduced in [1] and generally accepted ever since:

Definition 1 (Modal Refinement). *For a pair of modal transition systems S and T a binary relation $\mathcal{R} \subseteq states_S \times states_T$ is a modal refinement between states of S and T iff for all $(s, t) \in \mathcal{R}$ and all actions a it holds that:*

- for all $t' \in states_T$ such that $t \xrightarrow{a}^T t'$*
- there exists an $s' \in states_S$ such that $s \xrightarrow{a}^S s'$ and $(s', t') \in \mathcal{R}$,*
- for all $s' \in states_S$ such that $s \dashrightarrow^S s'$*
- there exists a $t' \in states_T$ such that $t \dashrightarrow^T t'$ and $(s', t') \in \mathcal{R}$.*

We say that a state $s \in states_S$ refines a state $t \in states_T$, written $s \leq_m t$, iff there exists a modal refinement containing (s, t) .

If $\longrightarrow^T = \emptyset$ then this refinement collapses to regular simulation [21, 22], while it coincides with bisimulation equivalence [23, 24] if S and T are implementations.

3 Non-thoroughness of Modal Refinement

Already in the eighties there have been rumors of modal refinement being incomplete. However we were unable to find a published account of this fact, so we

decided to include it here. We shall now define what we mean by completeness, proceeding to a counterexample witnessing the incompleteness of modal refinement. After this brief introduction we move to the first contribution of the paper: a discussion of the complexity class of a hypothetical complete refinement.




Fig. 1. $\llbracket S, s \rrbracket \subseteq \llbracket T, t \rrbracket$ and $s \not\leq_m t$

For a state $s \in \text{states}_S$ let $\llbracket S, s \rrbracket$ denote the set of all its implementations such that $\llbracket S, s \rrbracket = \{(I, i) \mid i \leq_m s \text{ and } \xrightarrow{I} = \dashrightarrow^I\}$. Modal refinement is known to be sound, with respect to implementation inclusion: for $s \in \text{states}_S \wedge t \in \text{states}_T$, if $s \leq_m t$ then also $\llbracket S, s \rrbracket \subseteq \llbracket T, t \rrbracket$, which follows directly from transitivity of \leq_m . However \leq_m is not complete in this sense: there exist specifications S and T , with states s, t , such that $\llbracket S, s \rrbracket \subseteq \llbracket T, t \rrbracket$ but $s \not\leq_m t$. This property of modal refinement is sometimes known as non-thoroughness [25]. Figure 1 presents a counterexample originating in the thesis of Hüttel [26, p. 32], also found in the thesis of Xinxin [27, p. 87] and in [18], albeit disguised in the context of disjunctive modal transition systems [8]¹. It contains two specifications S, T . It is a simple exercise to see that $\llbracket S, s \rrbracket = \llbracket T, t \rrbracket$, while $s \not\leq_m t$.

3.1 A Thorough Refinement is Co-NP Hard

Despite the non-thoroughness (incompleteness) of modal refinement its usefulness has never been questioned. This is probably because modal refinement is a natural generalization of both simulation and bisimulation and because it can be established efficiently (in time polynomial in the size of the transition systems). By showing that any complete refinement preserving precisely the same set of implementations as \leq_m cannot be decided in polynomial time (unless $P=NP$), we give yet another argument in favor of \leq_m .

We show co-NP hardness by reducing 3-DNF-TAUTOLOGY to checking a sound and complete modal refinement in the above sense. Consider a propositional formula φ over n variables x_1, \dots, x_n . It is clear that φ is a tautology iff $true \Rightarrow \varphi$ is a tautology. We will show how to construct, in polynomial time, a modal transition systems T_φ (representing a tautology over $x_1 \dots x_n$) and S_φ (representing φ), so that $true \Rightarrow \varphi$ is a tautology iff $\llbracket T_\varphi, true \rrbracket \subseteq \llbracket S_\varphi, \varphi \rrbracket$, for selected initial states $true$ and φ of T_φ and S_φ respectively. For simplicity we will assume that all clauses of φ are satisfiable. Satisfiability of a clause consisting of three conjunctions can be decided in constant time. Unsatisfiable clauses can thus be removed from φ in polynomial time, before we construct T_φ and S_φ . We choose the following states and actions for S_φ :

$$\text{states}_{S_\varphi} = \{\varphi, c_1, \dots, c_m, \mathbf{0}\} \quad \{a, x_1, \dots, x_n\} \subseteq \text{act} \quad , \quad (1)$$

where c_i are clauses of φ , while $\mathbf{0}$ and a are fresh names.

First we explain how a single literal can be represented as a state with at most $n + 1$ outgoing transitions. For a positive literal x_i we introduce a state x_i with a required transition $x_i \xrightarrow{x_i} \mathbf{0}$ and allowed transitions $x_i \xrightarrow{x_k} \mathbf{0}$ for all $k = 1 \dots n$.

¹ We thank Michael Huth, Harald Fecher, Heiko Schmidt and one of the anonymous reviewers for helping to track down its origins.




Fig. 2. Representing (a) a positive literal, (b) a negative literal, (c) a 3-DNF formula $\varphi = c_1 \vee \dots \vee c_m$ and (d) a tautology over variables $x_1 \dots x_n$.

For a negative literal $\neg x_i$ we allow no outgoing must transitions and create may transitions $(\neg x_i) \xrightarrow{x_k} \mathbf{0}$ for all $k \neq i$. Positive assignments are represented by must transitions, and negative assignments are represented by lack of may transitions. Assignments with no effect on satisfaction of the formula are modeled by may transitions with no corresponding must transitions. See Figure 2ab.

Now generalize this to conjunctive clauses of a 3-DNF formula. A clause $l_1 \wedge l_2 \wedge l_3$ is translated into a state labeled $l_1 \wedge l_2 \wedge l_3$ with the following transitions:

- 1° $(l_1 \wedge l_2 \wedge l_3) \xrightarrow{x_i} S_\varphi \mathbf{0}$ iff $l_k = x_i$ for some $k = 1 \dots 3$.
- 2° $(l_1 \wedge l_2 \wedge l_3) \xrightarrow{x_i} S_\varphi \mathbf{0}$ iff $l_k \neq \neg x_i$ for all $k = 1 \dots 3$.

Since we only consider satisfiable clauses, modal transition systems created this way are syntactically consistent (all required transitions are allowed). A satisfying truth assignment to $l_1 \wedge l_2 \wedge l_3$ can be extracted from any implementation I refining the state with the same label—just set x_i to *true* iff $I \xrightarrow{x_i}$ and set x_i to *false* otherwise. Similarly we can construct an implementation refining $l_1 \wedge l_2 \wedge l_3$ given any satisfying assignment to this clause.




Fig. 3. Reduction for $\varphi = (x_1 \wedge \neg x_2 \wedge x_3) \vee (\neg x_1 \wedge x_2 \wedge x_3) \vee (x_1 \wedge \neg x_2 \wedge \neg x_3)$.

Also each implementation of S_φ embeds at most one satisfying assignment to φ extracted using the same rules as discussed for clauses (one per each non-deterministic choice in the initial state of the implementation). Clearly S_φ can be constructed in time polynomial in the size of φ .

We now consider construction of T_φ . First let $states_{T_\varphi} = \{true, t_\varphi, \mathbf{0}\}$. We also create the following transitions: $true \xrightarrow{a} T_\varphi t_\varphi$, $true \xrightarrow{\neg a} T_\varphi t_\varphi$, and $t_\varphi \xrightarrow{x_i} T_\varphi \mathbf{0}$ for all variables x_i of φ (See Fig. 2d). Clearly T_φ can be constructed in time at most polynomial in size of φ .

A 3-DNF formula $\varphi = c_1 \vee \dots \vee c_m$ is represented using a state labeled φ and may transitions to its clauses: $\varphi \xrightarrow{a} S_\varphi c_i$ for $i = 1 \dots m$. No must transitions are generated. See Figure 2c and 3. States labeled c_i represent processes resulting from translation of the individual clauses as presented above.

Observe that each satisfying assignment to formula φ has a corresponding deterministic implementation of S_φ .

The following lemma states the correctness of our reduction.

Lemma 2. *A 3-DNF formula φ with all satisfiable clauses is a tautology iff $\llbracket T_\varphi, true \rrbracket \subseteq \llbracket S_\varphi, \varphi \rrbracket$.*

Proof. We first consider the direction right to left, i.e. assume that $\llbracket T_\varphi, true \rrbracket \subseteq \llbracket S_\varphi, \varphi \rrbracket$ and take any truth assignment ϱ to variables x_i of φ . We construct a deterministic implementation I_ϱ in the following way: $states_{I_\varrho} = \{t, \varrho, \mathbf{0}\}$, where there are two transition from t to ϱ : $t \xrightarrow{a} \varrho \wedge t \xrightarrow{-a} \varrho$ and for all x_i such that $\varrho(x_i) = true$: $\varrho \xrightarrow{x_i} \mathbf{0} \wedge \varrho \xrightarrow{-x_i} \mathbf{0}$. Due to the construction of our reduction this means that ϱ satisfies φ . Since for any assignment ϱ we can conclude that φ holds, φ is a tautology.

Now consider the claim of the lemma from left to right. We address its contrapositive. Assume that there exists an implementation I and its state t such that $t \leq_m true$, but $t \not\leq_m \varphi$. We want to show that φ is not a tautology. Observe that since $t \not\leq_m \varphi$ there must exist a state $s \in states_I$ such that $t \xrightarrow{a} s$ and for all clause states c_i of S_φ it is the case that $s \not\leq_m c_i$. But this means that the assignment represented by s (present x_i -transitions give rise to $x_i = true$, absent to $x_i = false$) falsifies φ meaning that φ is not a tautology. \square

Theorem 3. *The problem of deciding $\llbracket T, t \rrbracket \subseteq \llbracket S, s \rrbracket$ for states t and s of arbitrary modal transition systems T and S respectively is co-NP hard.*

Co-NP hardness follows from the above reduction and co-NP hardness of 3-DNF-TAUTOLOGY. The same reduction can be used to show that the thorough refinement induced by weak modal refinement (Section 6) is also co-NP hard to decide. We omit that proof as the argument is rather similar to the above.

4 Syntactic Consistency and Syntactic Refinement

From now on we relax the syntactic consistency requirement presented in Section 2, and allow reasoning about systems for which $\longrightarrow \not\subseteq \dashrightarrow$. We will introduce a syntactic refinement \subseteq_m with its induced notion of consistency and prove that it is (almost) precisely characterized by the syntactic consistency. These results are very simple, but we include them for three reasons. First, we cannot avoid discussing the most well known notion of consistency for modal transition systems (a notion that had never been characterized using a refinement relation). Second, we can show a refinement inducing this consistency (a refinement that had never been explicitly linked to any consistency notion). Third, we want to present all ingredients of a consistency study using a simple example: a refinement, its induced consistency, operational characterization in form of a consistency relation, and a coincidence proof. Later sections will follow exactly the same pattern.

Definition 4 (Syntactic Refinement). *For two modal transition systems S and T a syntactic refinement \mathcal{R} is a partial injective function on states s into states T such that for all pairs (s, t) , $t = \mathcal{R}(s)$, and all actions a it holds that*

for all $t' \in \text{states}_T$ such that $t \xrightarrow{a}^T t'$
there exists an $s' \in \text{states}_S$ such that $s \xrightarrow{a}^S s'$ and $t' = \mathcal{R}(s')$,
for all $s' \in \text{states}_S$ such that $s \xrightarrow{a}^S s'$
there exists a $t' \in \text{states}_T$ such that $t \xrightarrow{a}^T t'$ and $t' = \mathcal{R}(s')$.

A state s is said to be a syntactic refinement of a state t , written $s \subseteq_m t$, if there exists a syntactic refinement function \mathcal{R} such that $t = \mathcal{R}(s)$.

Intuitively this refinement establishes that the may-transition graph of S is a subgraph of the may-transition graph of T and that the must-transition graph of T is a subgraph of the must-transition graph of S .

Definition 5 (Syntactic Consistency). A state $s \in \text{states}_S$ is syntactically consistent iff there exists an implementation I and its state s^I such that $s^I \subseteq_m s$.

We claim that this notion of semantic consistency (almost) coincides with the one presented in Section 2. For the sake of uniformity let us reformulate that definition using an explicit notion of consistency relation:

Definition 6 (Syntactic Consistency Relation). Given a modal transition system S , a binary relation $\mathcal{S} \subseteq \text{states}_S \times \text{states}_S$ is a syntactic consistency relation on states of S iff for each state s if $(s, s) \in \mathcal{S}$ and each action $a \in \text{act}$ it holds that whenever $s \xrightarrow{a} s'$ for some $s' \in \text{states}_S$ then also $s \xrightarrow{a} s'$ and $(s', s') \in \mathcal{S}$.

For a syntactic consistency relation \mathcal{S} and a state $s \in \text{states}_S$ such that $(s, s) \in \mathcal{S}$, we synthesize an implementation $I_{\mathcal{S}}$ with a state s^I such that $s^I \subseteq_m s$. Take states of $I_{\mathcal{S}}$ to be consistent states of S : $\text{states}_{I_{\mathcal{S}}} = \{p \in \text{states}_S \mid (p, p) \in \mathcal{S}\}$ and $s^I = s$. The transition relation of $I_{\mathcal{S}}$ is the must transition relation of S projected on states of $I_{\mathcal{S}}$: $\xrightarrow{\cdot}^{I_{\mathcal{S}}} = \xrightarrow{\cdot}^S \cap (\text{states}_{I_{\mathcal{S}}} \times \text{act}^{\tau} \times \text{states}_{I_{\mathcal{S}}})$.

Theorem 7 (Soundness). If there exists a syntactic consistency relation containing a state s of S then s is a syntactically consistent state in the sense of Definition 5. Moreover the implementation $I_{\mathcal{S}}$ constructed above is one of its refinements: $s^I \subseteq_m s$.

It turns out that syntactic consistency relations characterize syntactic consistency in the sense of Definition 5 in a complete manner. Given a syntactic implementation I of a modal transition system S ($I \subseteq_m S$) we can construct a syntactic consistency relation in the following way:

$$\mathcal{S}_I = \{(q, q) \in \text{states}_S \mid \text{exists } p \in \text{states}_I \wedge p \subseteq_m q\} \quad (2)$$

Theorem 8 (Completeness). Let s be a state of a modal transition system S and s^I be a state of an implementation I such that $s^I \subseteq_m s$. Then there exists a syntactic consistency relation for S containing (s, s) , and \mathcal{S}_I is one of such.

Since establishing consistency of models is a useful feature in modeling tools, we remark that the cost of deciding existence of syntactic implementations (via consistency relations) for a state $s \in \text{states}_S$ is at most (and at least) linear in

the size of S . The algorithm corresponds to a traversal of the must-transition graph starting in s , and checking the consistency requirement in each state.

Syntactic consistency relations characterize syntactic consistency in the sense of [1] *almost* precisely. In fact the two notions coincide if all states of S are reachable from s via must transitions. Otherwise Definitions 5 and 6 allow inconsistencies in unreachable parts, which has not been taken into account in [1].

5 Strong Modal Refinement and Strong Consistency

In Section 2 we have recalled the notion of (strong) modal refinement. Now we introduce its induced notion of consistency and characterize it operationally.

Definition 9 (Strong Consistency). *A state s of a modal transition system S is strongly consistent iff there exists an implementation I and its state s^I such that $s^I \leq_m s$.*

In order to give an operational characterization of strong consistency we need to lift the transition relations to sets of states. For sets $\sigma, \sigma' \subseteq \text{states}_S$ we write:

$$\sigma \xrightarrow{a} \sigma' \quad \text{iff} \quad \exists s \in \sigma. \exists s' \in \sigma'. s \xrightarrow{a} s' , \quad (3)$$

$$\sigma \xrightarrow{a} \sigma' \quad \text{iff} \quad \forall s \in \sigma. \exists s' \in \sigma'. s \xrightarrow{a} s' . \quad (4)$$

Definition 10 (Strong Consistency Relation). *Given a modal transition system S , a relation $\mathcal{B} \subseteq \mathcal{P}(\text{states}_S)$ is a strong consistency relation on states_S iff for all actions $a \in \text{act}$ and all $\sigma \in \mathcal{B}$ the following condition is satisfied:*

*whenever $s \xrightarrow{a} s'$ for some $s \in \sigma$ and some $s' \in \text{states}_S$
then also $\sigma \xrightarrow{a} \sigma'$ and $\sigma \xrightarrow{a} \sigma'$ for some $\sigma' \in \mathcal{B}$ containing s' .*

Elements of \mathcal{B} are called consistency classes. \mathcal{B} is a strong consistency relation for a state $s \in \text{states}_S$ iff it contains a consistency class σ_s such that $s \in \sigma_s$.

Given a consistency relation \mathcal{B} for a state $s \in \text{states}_S$ we can synthesize an implementation $I_{\mathcal{B}}$ with a state $s^I \in \text{states}_{I_{\mathcal{B}}}$, such that $s^I \leq_m s$. Take the consistency classes of \mathcal{B} , to be the states of $I_{\mathcal{B}}$: $\text{states}_{I_{\mathcal{B}}} = \mathcal{B}$ and s^I be the class σ_s containing s . Both transition relations of $I_{\mathcal{B}}$ equal the intersection of *must* and *may* transition relations of S lifted to consistency classes of \mathcal{B} :

$$\sigma \xrightarrow{a} \sigma' \text{ and } \sigma \xrightarrow{a} \sigma' \quad \text{iff} \quad \sigma \xrightarrow{a} \sigma' \text{ and } \sigma \xrightarrow{a} \sigma' . \quad (5)$$

Theorem 11 (Soundness). *If there exists a consistency relation \mathcal{B} for a modal transition system S then S is strongly consistent in the sense of Definition 9. Moreover $I_{\mathcal{B}}$ constructed as above is one of its refinements: $s^I \leq_m s$.*

Strong consistency relations characterize strong consistency in a sound and complete manner. Given a state s^I of an implementation I refining a state $s \in \text{states}_S$ ($s^I \leq_m s$) we can construct a consistency relation \mathcal{B}_I for S following (6):

$$\mathcal{B}_I = \{ \sigma_p \subseteq \text{states}_S \mid p \in \text{states}_I \text{ and } \sigma_p \neq \emptyset \text{ and } \forall q \in \sigma_p. p \leq_m q \} \quad (6)$$

Observe that the σ_p sets above are not necessarily maximal.




Fig. 4. Representing (a) a disjunctive clause and (b) a translation for φ .

Theorem 12 (Completeness). *Let $s \in \text{states}_S$ and let I be an implementation, let $s^I \in \text{states}_I$ and $s^I \leq_m s$. Then there exists a consistency relation for the state s . Also relation \mathcal{B}_I defined above is one of such relations.*

Definition 10 can be interpreted operationally giving a simple exponential fixpoint algorithm: start with a singleton class containing s and apply the rule generating classes until a fixpoint is reached.

We demonstrate that the problem of deciding strong consistency is in fact NP-hard using a reduction from 3-CNF-SAT. Let $\varphi = c_1 \wedge \dots \wedge c_m$ be a 3-CNF formula over variables x_1, \dots, x_n . Construct a modal transition system S_φ such that its state labeled c_m is consistent iff φ is satisfiable. The states of S_φ are literals of φ , a $\mathbf{0}$ state, a $\mathbf{1}$ state (a state allowing any behavior: $\mathbf{1} \xrightarrow{x_i} S_\varphi \mathbf{1}$ for all $i = 1 \dots n$ and $\mathbf{1} \xrightarrow{a} S_\varphi \mathbf{1}$), plus a polynomial number of auxiliary states. We shall use an action per each variable x_i and one auxiliary action a .

Literals in φ are translated to states using the principle shown in Figure 2ab. A disjunction of three literals $l_1 \vee l_2 \vee l_3$ is represented by a state labeled $(l_1 \vee l_2 \vee l_3)$ such that $(l_1 \vee l_2 \vee l_3) \xrightarrow{a} S_\varphi \mathbf{1}$ and $(l_1 \vee l_2 \vee l_3) \xrightarrow{a} S_\varphi l_k$ for all $k = 1 \dots 3$. Now each clause c_i is represented by a state labeled c_i followed by a sequence of exactly i may a -transitions leading to the state representing the disjunction. For regularity we assume that there is a special *true* clause c_0 , that we translate to $\mathbf{1}$. Figure 4a shows the result of translating a clause $c_3 = x_1 \vee \neg x_2 \vee \neg x_3$. Recall that states labeled with literals are actually results of translation of Figure 2ab.

Now the top-level conjunction is translated inductively. First representations of c_1, \dots, c_m are created as above, then they are conjoined using must transitions. The i th clause is conjoined by a must transition from c_i to c_{i-1} : $c_i \xrightarrow{a} S_\varphi c_{i-1}$. Note that we add at most a quadratic number of auxiliary states this way (and a similar number of transitions). After conjoining c_m we obtain a representation of the whole formula. Figure 4b presents a complete translation for a formula $\varphi = (x_1 \vee \neg x_2 \vee x_3) \wedge (\neg x_1 \vee x_2 \vee x_3) \wedge (x_1 \vee \neg x_2 \vee \neg x_3)$. All unlabeled transitions should actually be labeled by a (removed to decrease clutter).

It is not hard to see that if the c_m state has an implementation then it actually has a state that satisfies the requirements of all the states representing disjunctions, and thus it induces a satisfiable assignment to φ .

6 Weak Refinement and Weak Consistency

We shall now discuss what is considered a classic form of a weak modal refinement (obtained by transforming modal refinement in the same way as bisimulation is transformed in order to obtain its weak form; to the best of our knowledge first published by Hüttel and Larsen in [19]). The definition uses a notion of weak transition relations that we introduce first. We shall write:

$$s \xrightarrow{a}_*^S s' \quad \text{iff} \quad s (\xrightarrow{\tau}^S)^* \xrightarrow{a}^S (\xrightarrow{\tau}^S)^* s' \quad (7)$$

$$s \xrightarrow{-a}_*^S s' \quad \text{iff} \quad s (\xrightarrow{-\tau}^S)^* \xrightarrow{-a}^S (\xrightarrow{-\tau}^S)^* s' , \quad (8)$$

where \mathcal{R}^* denotes zero or more transitive applications of a binary relation \mathcal{R} . Finally we write $s \xrightarrow{\hat{a}}_*^S s'$ whenever $s \xrightarrow{a}_*^S s'$ and $a \neq \tau$, or whenever $s (\xrightarrow{-\tau}^S)^* s'$ and $a = \tau$. Similarly for the may transition relation.

Definition 13 (Weak Modal Refinement). *Let S, T be modal transition systems. A binary relation $\mathcal{R} \subseteq \text{states}_S \times \text{states}_T$ is a weak modal refinement iff for each pair $(s, t) \in \mathcal{R}$ and each action $a \in \text{act}^\tau$ it holds that:*

*for all $t' \in \text{states}_T$ such that $t \xrightarrow{a}^T t'$
there exists $s' \in \text{states}_S$ such that $s \xrightarrow{\hat{a}}_*^S s'$ and $(s', t') \in \mathcal{R}$,
for all $s' \in \text{states}_S$ such that $s \xrightarrow{-a}_*^S s'$
there exists $t' \in \text{states}_T$ such that $t \xrightarrow{-\hat{a}}_*^T t'$ and $(s', t') \in \mathcal{R}$.*

We say that a state $s \in \text{states}_S$ weakly refines a state $t \in \text{states}_T$, written $s \leq_m^* t$ iff there exists a weak modal refinement containing (s, t) .

Definition 14 (Weak Consistency). *A state s of a modal transition system S is weakly consistent iff there exists an implementation I and its state s^I such that $s^I \leq_m^* s$.*

We characterize weak consistency using consistency relations as before. In order to do this we need to lift weak transition relations $\xrightarrow{-a}_*$ and \xrightarrow{a}_* to sets of states. For two sets of states $\sigma, \sigma' \subseteq \text{states}_S$ write:

$$\sigma \xrightarrow{\hat{a}}_*^{[S]} \sigma' \quad \text{iff} \quad \exists s \in \sigma. \exists s' \in \sigma'. s \xrightarrow{\hat{a}}_*^S s' , \quad (9)$$

$$\sigma \xrightarrow{-\hat{a}}_*^{[S]} \sigma' \quad \text{iff} \quad \forall s \in \sigma. \exists s' \in \sigma'. s \xrightarrow{-\hat{a}}_*^S s' . \quad (10)$$

Definition 15 (Weak Consistency Relation). *Let S be a modal transition system. A relation $\mathcal{O} \subseteq \mathcal{P}(\text{states}_S)$ is a weak consistency relation on states_S iff for any set $\sigma \in \mathcal{O}$, for any state $s \in \sigma$, and for any action $a \in \text{act}^\tau$ it holds that:*

whenever $s \xrightarrow{a}_^S s'$ for some $s' \in \text{states}_S$
then also $\sigma \xrightarrow{\hat{a}}_*^{[S]} \sigma'$ and $\sigma \xrightarrow{-\hat{a}}_*^{[S]} \sigma'$ for some $\sigma' \in \mathcal{O}$ containing s' .*

\mathcal{O} is a weak consistency relation for a state $s \in \text{states}_S$ iff it contains a consistency class σ_s such that $\text{start}_s \in \sigma_s$.

As before, we claim that weak consistency relations (Definition 15) soundly characterize weak consistency (Definition 14): for a state $s \in \text{states}_S$ with a known weak consistency relation \mathcal{O} , one can construct a weak implementation $I_{\mathcal{O}}$ containing a state s^I such that $s^I \leq_m^* s$. Take states of $I_{\mathcal{O}}$ to be consistency classes of \mathcal{O} ($\text{states}_{I_{\mathcal{O}}} = \mathcal{O}$), and s^I be a class σ_s containing s . The transition relations of $I_{\mathcal{O}}$ are the intersection of the weak transition relations of S lifted to consistency classes of \mathcal{O} . For all actions $a \in \text{act}^\tau$:

$$\sigma \xrightarrow{a}^{I_{\mathcal{O}}} \sigma' \text{ and } \sigma \xrightarrow{-a}^{I_{\mathcal{O}}} \sigma' \quad \text{iff} \quad \sigma \xrightarrow{\hat{a}}_*^{[S]} \sigma' \text{ and } \sigma \xrightarrow{-\hat{a}}_*^{[S]} \sigma' . \quad (11)$$

Theorem 16 (Soundness). *Let S be a modal transition system, $s \in \text{states}_S$, and \mathcal{O} be a weak consistency relation for s . Then s is weakly consistent and $s^I \in \text{states}_{I_{\mathcal{O}}}$ is one of its implementations: $s^I \leq_m^* s$.*

Consistency relations characterize weak consistency precisely. Assume that a state $s \in \text{states}_S$ is refined by a state s^I of an implementation I ($I \leq_m^* S$). Then one can use this implementation to construct the consistency relation \mathcal{O}_I :

$$\mathcal{O}_I = \{\sigma_p \subseteq \text{states}_S \mid p \in \text{states}_I \text{ and } \sigma_p \neq \emptyset \text{ and } \forall q \in \sigma_p. p \leq_m^* q\} \quad (12)$$

Theorem 17 (Completeness). *Let S be a modal transition system, I be an implementation, and let $s^I \leq_m^* s$ for some $s^I \in \text{states}_I$ and $s \in \text{states}_S$. Then there exist weak consistency relations for s , and \mathcal{O}_I is one of them.*

Definition 15 can be interpreted operationally giving rise to an exponential algorithm for constructing a consistency relation and deciding weak consistency. Weak consistency collapses to strong consistency for systems without transitions labeled with τ . Consequently the problem of deciding it is at least NP-hard, by reduction from 3-CNF-SAT presented in Section 5.

We conclude this section with a comment on synthesis of a weak implementation $I_{\mathcal{O}}$ from a consistency relation \mathcal{O} . The implementation synthesized by the algorithm presented above will contain internal transitions, if the specification contained them. In fact this is not always necessary—there definitely exist specifications with internal transitions that can be realized without hidden behavior. However, hidden transitions are unavoidable for some specifications. Figure 5 shows such a specification (in fact even a syntactically consistent one).




Fig. 5. All implementations of T have τ -transitions.

7 May-weak Modal Refinement and Its Consistency

In [9] we have proposed another weakening of modal refinement, generalizing alternating simulation [28] for two players as used in interface automata [29]. We call it *may-weak* here, as it preserves strong behavior on must transitions, only allowing weak matching on may transitions. It has been demonstrated that

may-weak modal refinement is a sound basis for assume/guarantee reasoning: it preserves absence of deadlocks on guaranteed behaviors (details in [9]).

Before we can define the may-weak refinement, let us define the may-weak transition relation as used in this refinement. We shall write

$$s \xrightarrow{a}_{\triangleleft}^S s' \quad \text{iff} \quad s(-\tau \xrightarrow{S})^* s'' \xrightarrow{a}_{\triangleleft}^S s' \quad (13)$$

Similarly as before we write $s \xrightarrow{\hat{a}}_{\triangleleft}^S s'$ meaning $s \xrightarrow{a}_{\triangleleft}^S s'$ if $a \in \text{act}$ and $s(-\tau \xrightarrow{S})^* s'$ if $a = \tau$. We use the regular (strong) must-transition relation lifted to sets of states as in Section 5. We also lift our new may-weak transition relation:

$$\sigma \xrightarrow{\hat{a}}_{\triangleleft}^{\lfloor S \rfloor} \sigma' \quad \text{iff} \quad \forall s \in \sigma. \exists s' \in \sigma'. s \xrightarrow{\hat{a}}_{\triangleleft}^S s' . \quad (14)$$

Let us now define may-weak modal refinement [9] using may-weak transitions:

Definition 18 (May-weak Modal Refinement). *A binary relation $\mathcal{R} \subseteq \text{states}_S \times \text{states}_T$ is a may-weak refinement between states of two modal transition systems S and T iff for each pair of states $(s, t) \in \mathcal{R}$ it holds that:*

*for all $a \in \text{act}$ and for all $t' \in \text{states}_T$ such that $t \xrightarrow{a}_T t'$
there exists $s' \in \text{states}_S$ such that $s \xrightarrow{a}_S s'$ and $(s', t') \in \mathcal{R}$,
for all $a \in \text{act}^\tau$ and for all $s' \in \text{states}_S$ $s \xrightarrow{a}_S s'$
there exists $t' \in \text{states}_T$ such that $t \xrightarrow{\hat{a}}_T t'$ and $(s', t') \in \mathcal{R}$.*

A state $s \in \text{states}_S$ may-weakly refines a state $t \in \text{states}_T$, written $s \leq_m^{\triangleleft} t$ iff there exists a may-weak modal refinement containing (s, t) .

Definition 19 (May-weak Consistency). *A state s of a modal transition system S is may-weak consistent iff there exists an implementation I and its state s^I such that $s^I \leq_m^{\triangleleft} s$.*

Definition 20 (May-weak Consistency Relation). *Let S be a modal transition system. A relation $\mathcal{U} \subseteq \mathcal{P}(\text{states}_S)$ is a may-weak consistency relation on states_S iff for any set of states $\sigma \in \mathcal{U}$, for any state $s \in \sigma$, and for any action $a \in \text{act}$ the following holds:*

*whenever $s \xrightarrow{a}_S s'$ for some $s' \in \text{states}_S$
then also $\sigma \xrightarrow{a}_{\triangleleft}^{\lfloor S \rfloor} \sigma'$ and $\sigma \xrightarrow{a}_{\triangleleft}^{\lfloor S \rfloor} \sigma'$ for some $\sigma' \in \mathcal{U}$ containing s' .*

\mathcal{U} is a may-weak consistency relation for a state $s \in \text{states}_S$ iff it contains a consistency class $\sigma_s \in \mathcal{U}$ such that $s \in \sigma_s$.

Given a consistency relation \mathcal{U} for a state s of a modal transition system S , we can synthesize an implementation $I_{\mathcal{U}}$ with a state s^I refining s . The states of $I_{\mathcal{U}}$ are the consistency classes of \mathcal{U} : $\text{states}_{I_{\mathcal{U}}} = \mathcal{U}$ and s^I is the consistency class containing s . Transition relations of $I_{\mathcal{U}}$ equal intersection of *must* and may-weak transition relations of S lifted to consistency classes in \mathcal{U} (for $a \neq \tau$):

$$\sigma \xrightarrow{a}_{\rightarrow}^{I_{\mathcal{U}}} \sigma' \quad \text{and} \quad \sigma \xrightarrow{a}_{\rightarrow}^{I_{\mathcal{U}}} \sigma' \quad \text{iff} \quad \sigma \xrightarrow{a}_{\rightarrow}^{\lfloor S \rfloor} \sigma' \quad \text{and} \quad \sigma \xrightarrow{a}_{\rightarrow}^{\lfloor S \rfloor} \sigma' , \quad (15)$$

Theorem 21 (Soundness). *Let $s \in \text{states}_S$. If \mathcal{U} is a may-weak consistency relation for s then s is may-weakly consistent and $s^I \in \text{states}_{I_{\mathcal{U}}}$ constructed as above is one of its implementations: $s^I \leq_m^{\triangleleft} s$.*

For the completeness of characterization consider an implementation I , a state $s^I \in \text{states}_I$ such that $s^I \leq_m^{\triangleleft} s$, where $s \in \text{states}_S$. We construct a consistency relation \mathcal{U}_I for s in the following way:

$$\mathcal{U}_I = \{ \sigma_p \subseteq \text{states}_S \mid p \in \text{states}_I \text{ and } \sigma_p \neq \emptyset \text{ and } \forall q \in \sigma_p. p \leq_m^{\triangleleft} q \} . \quad (16)$$

Theorem 22 (Completeness). *Let S be a modal transition system, $s \in \text{states}_S$ and let I be an implementation such that $s^I \leq_m^{\triangleleft} s$ for some $s^I \in \text{states}_I$. Then there exist a may-weak consistency relation for s , and \mathcal{U}_I is one such relation.*

Existence of a may-weak consistency relation for a given state s can be decided in exponential time, using an algorithm that is easy to extract from Definition 20. As previously this problem is also NP-hard, as may-weak consistency collapses to strong consistency for specifications without τ transitions.

A remarkable property of may-weak modal refinement, which we have not realized when writing [9], is that a may-weak consistent system always has implementations that contain no hidden actions ($I_{\mathcal{U}}$ above is actually constructed without introducing internal transitions). This is because this refinement captures a kind of (observation) determinism of required behaviors in specifications. We find this property appealing for applications again: it describes a class of specifications which allow implementations that are predictable (provided that they are deterministic). As predictability is an important property of software systems, the above decision procedure is likely to prove useful in practice.

8 Conclusion and Open Problems

We have addressed several basic questions in the theory of modal transition systems. We have shown that deciding any refinement that captures, in a precise way, the same set of concrete implementations as the standard modal refinement (or weak modal refinement) is co-NP hard. This lower bound is not tight. An upper bound of EXPTIME is easily established by casting the problem as checking satisfiability of implication between two characteristic formulas, in the modal μ -calculus. Finding a tight bound remains an open problem that we shall address shortly. We also hope to study hardness of thorough refinements induced by may-weak modal refinement and branching modal refinement [10].

Furthermore we have contributed to the understanding of the relation between refinements and consistencies studying notions of consistency for modal transition systems induced by four different refinement relations: syntactic consistency [1] (induced by a graph inclusion refinement), strong consistency (induced by a regular modal refinement [1]), weak consistency (induced by weak modal refinement [19]) and may-weak consistency (induced by may-weak modal refinement [9]). For each of these we have given a sound and complete operational characterization. The upper bound on establishing the last three of these

consistencies is exponential, and they are NP-hard. Syntactic consistency can be established in linear time.

There is a range of open problems related to these results. First, it is an interesting question whether there exists a useful alternative to modal refinement that completely characterizes its own (as opposed to the currently accepted) set of implementations and that can be decided in polynomial time. The main challenge here is to argue that the set of implementations considered is interesting from a practical point of view. Alternatively, as suggested to us by Michael Huth, one can try to characterize broad classes of modal transition systems for which the currently used refinement is complete.

Finding a uniform formulation for four consistency studies as presented in this paper was a rather challenging but rewarding task. Given that they can be described so similarly one could try to take this analogy further and design a more abstract meta-consistency theory, parameterized only by a refinement.

Furthermore it is interesting to study the relation between consistency and parallel decomposition. We have done some preliminary work on that topic in [9], though in a rather restricted setting. We intend to generalize observational consistency of [9], and to understand its semantics building on the results of the present paper; ultimately employing it in a larger study of decomposition.

References

1. Larsen, K.G., Thomsen, B.: A modal process logic. In: LICS, IEEE Computer Society (1988)
2. Huth, M., Jagadeesan, R., Schmidt, D.: Modal transition systems: A foundation for three-valued program analysis. *Lecture Notes in Computer Science* **2028** (2001)
3. Schmidt, D.: From trace sets to modal-transition systems by stepwise abstract interpretation (2001)
4. Godefroid, P., Huth, M., Jagadeesan, R.: Abstraction-based model checking using modal transition systems. *Lecture Notes in Computer Science* **2154** (2001) 426+
5. Børjesson, A., Larsen, K.G., Skou, A.: Generality in design and compositional verification using tav. In: FORTE '92 Proceedings, Amsterdam, The Netherlands, The Netherlands, North-Holland Publishing Co. (1993) 449–464
6. Larsen, K.G., Steffen, B., Weise, C.: A constraint oriented proof methodology based on modal transition systems. In: *Tools and Algorithms for Construction and Analysis of Systems*. (1995) 17–40
7. Bruns, G.: An industrial application of modal process logic. *Sci. Comput. Program.* **29**(1-2) (1997) 3–22
8. Larsen, K.G., Xinxin, L.: Equation solving using modal transition systems. In: *Fifth Annual IEEE Symposium on Logics in Computer Science (LICS)*, 4–7 June 1990, Philadelphia, PA, USA. (1990) 108–117
9. Larsen, K.G., Nyman, U., Wąsowski, A.: Modal i/o automata for interface and product line theories. In Nicola, R.D., ed.: *Programming Languages and Systems, ESOP 2007*. Volume 4421 of LNCS., Springer (2007) 64–79
10. Fischbein, D., Uchitel, S., Braberman, V.: A foundation for behavioural conformance in software product line architectures. In: *ROSATEA '06 Proceedings*, New York, NY, USA, ACM Press (2006) 39–48

11. Uchitel, S., Chechik, M.: Merging partial behavioural models. In Taylor, R.N., Dwyer, M.B., eds.: SIGSOFT FSE, ACM (2004) 43–52
12. Brunet, G., Chechik, M., Uchitel, S.: Properties of behavioural model merging. In Misra, J., Nipkow, T., Sekerinski, E., eds.: FM. Volume 4085 of Lecture Notes in Computer Science., Springer (2006) 98–114
13. C. Weise, D. Lenzkes: Weak refinement for modal hybrid systems. In O. Maler, ed.: Hybrid and Real-Time Systems, Grenoble, France, Springer Verlag, LNCS 1201 (1997) 316–330
14. Larsen, K.G.: Modal specifications. In Sifakis, J., ed.: Automatic Verification Methods for Finite State Systems. Volume 407 of Lecture Notes in Computer Science., Springer (1989) 232–246
15. Cerans, K., Godskesen, J.C., Larsen, K.G.: Timed modal specification - theory and tools. In: CAV '93: Proceedings of the 5th International Conference on Computer Aided Verification, London, UK, Springer-Verlag (1993) 253–267
16. Larsen, K.G., Steffen, B., Weise, C.: Fischer's protocol revisited: a simple proof using modal constraints. Lecture Notes in Computer Science **1066** (1996) 604–615
17. Fecher, H., Huth, M.: Ranked predicate abstraction for branching time: Complete incremental, and precise. In: ATVA. Volume 4218 of Lecture Notes in Computer Science., Springer (2006) 322–336
18. Schmidt, H., Fecher, H.: Comparing disjunctive modal transition systems with a one-selecting variant. (2007) submitted for publication.
19. Hüttel, H., Larsen, K.G.: The use of static constructs in a modal process logic. In: LFCS: The 1st International Symposium on Logical Foundations of Computer Science. (1989)
20. Dams, D.: Abstract Interpretation and Partition Refinement for Model Checking. PhD thesis, Eindhoven University of Technology (July 1996)
21. Hennessy, M., Milner, R.: Algebraic laws for nondeterminism and concurrency. Journal of the ACM (1985) 137–161
22. Larsen, K.G.: A context dependent bisimulation between processes. Theoretical Computer Science **49** (1987)
23. Park, D.: Concurrency and automata on infinite sequences. In: Proceedings of 5th GI Conference. Volume 104. (1981)
24. Milner, R.: Calculi for synchrony and asynchrony. Theoretical Computer Science **25** (1983)
25. Godefroid, P., Jagadeesan, R.: Automatic abstraction using generalized model checking. In Brinksma, E., Larsen, K.G., eds.: CAV. Volume 2404 of Lecture Notes in Computer Science., Springer (2002) 137–150
26. Hüttel, H.: Operational and denotational properties of modal process logic. Master's thesis, Computer Science Department. Aalborg University (1988)
27. Xinxin, L.: Specification and Decomposition in Concurrency. PhD thesis, Department of Mathematics and Computer Science, Aalborg University (April 1992)
28. Alur, R., Henzinger, T.A., Kupferman, O., Vardi, M.: Alternating refinement relations. In Sangiorgi, D., de Simone, R., eds.: Proceedings of the Ninth International Conference on Concurrency Theory (CONCUR'98). Volume 1466. (1998) 163–178
29. Alfaro, L., Henzinger, T.A.: Interface automata. In: Proceedings of the Ninth Annual Symposium on Foundations of Software Engineering (FSE), Vienna, Austria (September 2001) 109–120