Decision Problems for Interval Markov Chains *

Benoît Delahaye[1], Kim G. Larsen[2], Axel Legay[3],
Mikkel L. Pedersen[2], and Andrzej Wąsowski[4]

[1] Université de Rennes 1/IRISA, France
[2] Aalborg University, Denmark
[3] INRIA/IRISA, France
[4] IT University of Copenhagen, Denmark

Abstract. Interval Markov Chains (IMCs) are the base of a classic probabilistic specification theory by Larsen and Jonsson in 1991. They are also a popular abstraction for probabilistic systems.
In this paper we study complexity of several problems for this abstraction, that stem from compositional modeling methodologies. In particular we close the complexity gap for thorough refinement of two IMCs and for deciding the existence of a common implementation for an unbounded number of IMCs, showing that these problems are EXPTIME-complete. We also prove that deciding consistency of an IMC is polynomial and discuss suitable notions of determinism for such specifications.

1 Introduction

Interval Markov Chains (IMCs for short) extend Markov Chains, by allowing to specify intervals of possible probabilities on state transitions. IMCs have been introduced by Larsen and Jonsson [10] as a *specification* formalism—a basis for a stepwise-refinement-like modeling method, where initial designs are very abstract and underspecified, and are then made continuously more precise, until they are concrete. Unlike richer specification models such as Constraint Markov Chains [4], IMCs are difficult to use for compositional specification due to lack of basic modeling operators. To address this, we study complexity and algorithms for deciding consistency of conjunctive sets of IMC specifications.

In [10] Jonsson and Larsen have introduced refinement for IMCs, but have not determined its computational complexity. We complete their work on refinement by classifying its complexity and characterizing it using structural coinductive algorithms in the style of simulation.

Consider the issue of combining multiple specifications of the same system. It turns out that conjunction of IMCs cannot be expressed as an IMC itself, due to a lack of expressiveness of intervals. Let us demonstrate this using a simple specification of a user of a coffee machine. Let the model prescribe that a typical

* This work was supported by the European STREP-COMBEST project no. 215543, by VKR Centre of Excellence MT-LAB, and by an "Action de Recherche Collaborative" ARC (TP)I.

user orders coffee with milk with probability $x \in [0, 0.5]$ and black coffee with probability $y \in [0.2, 0.7]$ (customers also buy tea with probability $t \in [0, 0.5]$). The vendor of the machine delivers another specification, which prescribes that the machine is serviceable only if coffee (white or black) is ordered with some probability $z \in [0.4, 0.8]$ from among other beverages, otherwise it will run out of coffee powder too frequently, or the powder becomes too old. A conjunction of these two models would describe users who have use patterns compatible with this particular machine. Such a conjunction effectively requires that all the interval constraints are satisfied and that $z = x + y$ holds. However, the solution of this constraint is not described by an interval over x and y. This can be seen by pointing out an extremal point, which is not a solution, while all its coordinates take part in some solution. Say $x = 0$ and $y = 0.2$ violates the interval for z, while for each of these two values it is possible to select another one in such a way that z's constraint is also held (for example $(x = 0, y = 0.4)$ and $(x = 0.2, y = 0.2)$). Thus the solution space is not an interval over x and y.

This lack of closure properties for IMCs motivates us to address the problem of reasoning about conjunction without constructing it — the, so called, common implementation problem. In this paper we provide algorithms and complexity results for consistency, common implementation and refinement of IMCs, in order to enable compositional modeling. We contribute the following new results:

- In [10] a *thorough refinement* (TR) between IMCs is defined as an inclusion of implementation sets. We define suitable notions of determinism for IMCs, and show that for deterministic IMCs TR coincides with two simulation-like preorders (the *weak refinement* and *strong refinement*), for which there exist co-inductive algorithms terminating in a polynomial number of iterations.
- We show that the thorough refinement procedure given in [10] can be implemented in single exponential time. Furthermore we provide a lower bound, concluding that TR is EXPTIME-complete. While the reduction from TR of modal transition systems [3] used to provide this lower bound is conceptually simple, it requires a rather involved proof of correctness, namely that it preserves sets of implementations in a sound and complete manner.
- A polynomial procedure for checking whether an IMC is *consistent* (C), i.e. it admits a Markov Chain as an implementation.
- An exponential procedure for checking whether k IMCs are consistent in the sense that they share a Markov Chain satisfying all—a *common implementation* (CI). We show that this problem is EXPTIME-complete.
- As a special case we observe, that CI is PTIME for any constant value of k. In particular checking whether two specifications can be simultaneously satisfied, and synthesizing their shared implementation can be done in polynomial time.

For functional analysis of discrete-time non-probabilistic systems, the theory of Modal Transition Systems (MTSs) [15] provides a specification formalism supporting refinement, conjunction and parallel composition. Earlier we have obtained EXPTIME-completeness both for the corresponding notion of CI [2] and of TR [3] for MTSs. In [10] it is shown that IMCs properly contain MTSs,

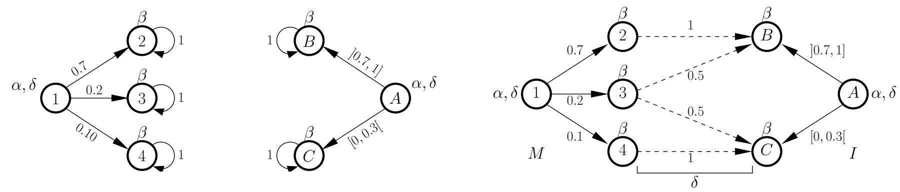

(a) A Markov Chain M (b) An IMC I (c) An example of satisfaction relation.

Fig. 1. Examples of Markov Chains, Interval Markov Chains and satisfaction relation.

which puts our new results in a somewhat surprising light: in the complexity theoretic sense, and as far as CI and TR are considered, the generalization of modalities by probabilities does come for free.

The paper proceeds as follows. In Section 2 we introduce the basic definitions. All results in subsequent sections are new and ours. In Section 3 we discuss deciding TR and other refinement procedures. We expand on the interplay of determinism and refinements in Section 4. The problems of C and CI are addressed in Section 5. We close by discussing the results and related work in Section 6. Due to space constraints, some algorithms and proofs are given in a long version of this paper [6].

2 Background

We shall now introduce the basic definitions used throughout the paper. In the following we will write $\mathsf{Intervals}_{[0,1]}$ for the set of all closed, half-open and open intervals included in $[0, 1]$.

We begin with settling notation for Markov Chains. A Markov Chain (sometimes MC in short) is a tuple $C = \langle P, p_0, \pi, A, V_C \rangle$, where P is a set of states containing the initial state p_0, A is a set of atomic propositions, $V_C : P \to 2^A$ is a state valuation labeling states with propositions, and $\pi : P \to \mathsf{Distr}(P)$ is a probability distribution assignment such that $\sum_{p' \in P} \pi(p)(p') = 1$ for all $p \in P$. The probability distribution assignment is the only component that is relaxed in IMCs:

Definition 1 (Interval Markov Chain). *An* Interval Markov Chain *is a tuple $I = \langle Q, q_0, \varphi, A, V_I \rangle$, where Q is a set of states containing the initial state q_0, A is a set of atomic propositions, $V_I : Q \to 2^A$ is a state valuation, and $\varphi : Q \to (Q \to \mathsf{Intervals}_{[0,1]})$, which for each $q \in Q$ and $q' \in Q$ gives an interval of probabilities.*

Instead of a distribution, as in MCs, in IMCs we have a function mapping elementary events (target states) to intervals of probabilities. We interpret this function as a constraint over distributions. This is expressed in our notation as follows. Given a state $q \in Q$ and a distribution $\sigma \in \mathsf{Distr}(Q)$, we say that

$\sigma \in \varphi(q)$ iff $\sigma(q') \in \varphi(q)(q')$ for all $q' \in Q$. Occasionally, it is convenient to think of a Markov Chain as an IMC, in which all probability intervals are closed point intervals.

We visualize IMCs as automata with intervals on transitions. As an example, consider the IMC in Figure 1b. It has two outgoing transitions from the initial state A. No arc is drawn between states if the probability is zero (or more precisely the interval is $[0, 0]$), so in the example there is zero probability of going from state A to A, or from B to C, etc. Otherwise the probability distribution over successors of A is constrained to fall into $]0.7, 1]$ and $[0, 0.3[$ for B and C respectively. States B and C have valuation β, whereas state A has valuation α, δ. Figure 1a presents a Markov Chain using the same convention, modulo the intervals. Notice that our formalism does not allow "sink states" with no outgoing transitions. In the figures, states with no outgoing transitions are meant to have a self-loop transition with probability 1 (a closed point interval).

There are three known ways of defining refinement for IMCs: strong refinement (introduced as *simulation* in [10]), weak refinement (introduced under the name of *probabilistic simulation* in [7]), and thorough refinement (introduced as *refinement* in [10]). We recall their formal definitions:

Definition 2 (Strong Refinement). *Let $I_1 = \langle Q, q_0, \varphi_1, A, V_1\rangle$ and $I_2 = \langle S, s_0, \varphi_2, A, V_2\rangle$ be IMCs. A relation $\mathcal{R} \subseteq Q \times S$ is a* strong refinement relation *if whenever $q \mathcal{R} s$ then*

1. *The valuation sets agree: $V_1(q) = V_2(s)$ and*
2. *There exists a correspondence function $\delta : Q \to (S \to [0, 1])$ such that, for all $\sigma \in \mathsf{Distr}(Q)$, if $\sigma \in \varphi_1(q)$, then*
   (a) *for all $q' \in Q$ such that $\sigma(q') > 0$, $\delta(q')$ is a distribution on S,*
   (b) *for all $s' \in S$, we have $\sum_{q' \in Q} \sigma(q') \cdot \delta(q')(s') \in \varphi_2(s)(s')$, and*
   (c) *for all $q' \in Q$ and $s' \in S$, if $\delta(q')(s') > 0$, then $q' \mathcal{R} s'$.*

I_1 strongly refines I_2, or $I_1 \leq_{\mathsf{S}} I_2$, iff there exists a strong refinement containing (q_0, s_0).

A strong refinement relation requires the existence of a single correspondence, which witnesses satisfaction for any resolution of probability constraint over successors of q and s. Figure 2a illustrates such a correspondence between states A and α of two IMCs. The correspondence function is given by labels on the dashed lines. It is easy to see that, regardless of how the probability constraints are resolved, the correspondence function distributes the probability mass in a fashion satisfying α.

A *weak refinement* relation requires that, for any resolution of probability constraint over successors in I_1, there exists a correspondence function, which witnesses satisfaction of I_2. The formal definition of weak refinement is identical to Def. 2, except that the condition opening Point (2) is replaced by a weaker one:

Definition 3 (Weak Refinement). *Let $I_1 = \langle Q, q_0, \varphi_1, A, V_1 \rangle$ and $I_2 = \langle S, s_0, \varphi_2, A, V_2 \rangle$ be IMCs. A relation $\mathcal{R} \subseteq Q \times S$ is a* weak refinement relation *if whenever $q \mathcal{R} s$, then*

1. *The valuation sets agree: $V_1(q) = V_2(s)$ and*
2. *For each $\sigma \in \mathsf{Distr}(Q)$ such that $\sigma \in \varphi_1(q)$, there exists a correspondence function $\delta : Q \to (S \to [0,1])$ such that*
   (a) *for all $q' \in Q$ such that $\sigma(q') > 0$, $\delta(q')$ is a distribution on S,*
   (b) *for all $s' \in S$, we have $\sum_{q' \in Q} \sigma(q') \cdot \delta(q')(s') \in \varphi_2(s)(s')$, and*
   (c) *for all $q' \in Q$ and $s' \in S$, if $\delta(q')(s') > 0$, then $q' \mathcal{R} s'$.*

I_1 weakly refines I_2, or $I_1 \leq_{\mathsf{W}} I_2$, iff there exists a weak refinement containing (q_0, s_0).

Figure 2b illustrates a weak refinement between states A and α of another two IMCs. Here, x stands for a value in $[0.2, 1]$ (arbitrary choice of probability of going to state C from A). Notably, for each choice of x, there exists $p \in [0,1]$ such that $p \cdot x \in [0, 0.6]$ and $(1-p) \cdot x \in [0.2, 0.4]$.

Satisfaction Relation. This relation establishes compatibility of Markov Chains (implementations) and IMCs (specifications). The original definition has been presented in [10, 11]. Consider a Markov chain $C = \langle P, p_0, \pi, A, V_C \rangle$ as an IMC with only closed point interval probabilities, and let $I = \langle Q, q_0, \varphi, A, V_I \rangle$ be an IMC. We say that C satisfies I, written $C \models I$, iff there exists a weak/strong refinement relation $\mathcal{R} \subseteq P \times Q$, called a *satisfaction relation*, containing (p_0, q_0). Remark that when C is a Markov Chain, the weak and strong notions of refinement coincide. Whenever $C \models I$, C is called an *implementation* of I. The set of implementations of I is written $[\![I]\!]$. Figure 1c presents an example of satisfaction on states 1 and A. The correspondence function is specified using labels on the dashed arrows i.e. the probability mass going from state 1 to 3 is distributed to state B and C with half going to each.

We say that a state q of an IMC is *consistent* if its interval constraint $\varphi(q)$ is satisfiable, i.e. there exists a distribution $\sigma \in \mathsf{Distr}(Q)$ satisfying $\varphi(q)$. Obviously, for a given IMC, it is sufficient that all its states are consistent in order to guarantee that the IMC is consistent itself—there exists a Markov Chain satisfying it. We discuss the problem of establishing consistency in a sound and complete manner in Section 5.

Finally, we introduce the thorough refinement as defined in [10]:

Definition 4 (Thorough Refinement). *IMC I_1* thoroughly refines *IMC I_2, written $I_1 \leq_{\mathsf{T}} I_2$, iff each implementation of I_1 implements I_2: $[\![I_1]\!] \subseteq [\![I_2]\!]$*

Thorough refinement is the ultimate refinement relation for any specification formalism, as it is based on the semantics of the models.

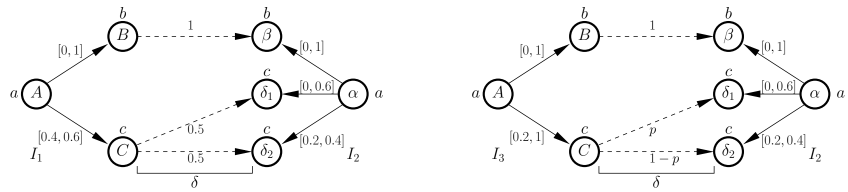

(a) Illustration of a strong refinement relation between an IMC I_1 and an IMC I_2.

(b) Illustration of a weak refinement relation between an IMC I_3 and an IMC I_2.

Fig. 2. Illustration of strong and weak refinement relations.

3 Refinement Relations

In this section, we compare the expressiveness of the refinement relations. It is not hard to see that both strong and weak refinements soundly approximate the thorough refinement (since they are transitive and degrade to satisfaction if the left argument is a Markov Chain). The converse does not hold. We will now discuss procedures to compute weak and strong refinements, and then compare the granularity of these relations, which will lead us to procedures for computing thorough refinement. Observe that both refinements are decidable, as they only rely on the first order theory of real numbers. In concrete cases below the calculations can be done more efficiently due to convexity of solution spaces for interval constraints.

Weak and Strong Refinement. Consider two IMCs $I_1 = \langle P, o_1, \varphi_1, A, V_1 \rangle$ and $I_2 = \langle Q, o_2, \varphi_2, A, V_2 \rangle$. Informally, checking whether a given relation $\mathcal{R} \subseteq P \times Q$ is a weak refinement relation reduces to checking, for each pair $(p, q) \in \mathcal{R}$, whether the following formula is true: $\forall \pi \in \varphi_1(p)$, $\exists \delta : P \to (Q \to [0, 1])$ such that $\pi \times \delta$ satisfies a system of linear equations / inequations. Since the set of distributions satisfying $\varphi_1(p)$ is convex, checking such a system is exponential in the number of variables, here $|P| \cdot |Q|$. As a consequence, checking whether a relation on $P \times Q$ is a weak refinement relation is exponential in $|P| \cdot |Q|$. For strong refinement relations, the only difference appears in the formula that must be checked: $\exists \delta : P \to (Q \to [0, 1])$ such that $\forall \pi \in \varphi_1(p)$, we have that $\pi \times \delta$ satisfies a system of linear equations / inequations. Therefore, checking whether a relation on $P \times Q$ is a strong refinement relation is also exponential in $|P| \cdot |Q|$.

Deciding whether weak (strong) refinement holds between I_1 and I_2 can be done in the usual coinductive fashion by considering the total relation $P \times Q$ and successively removing all the pairs that do not satisfy the above formulae. The refinement holds iff the relation we reach contains the pair (o_1, o_2). The algorithm will terminate after at most $|P| \cdot |Q|$ iterations. This gives an upper bound on the complexity to establish strong and weak refinements: a polynomial number of iterations over an exponential step. This upper bound may be loose. One could try to reuse techniques for nonstochastic systems [9] in order to reduce the number of iterations. This is left to future work.

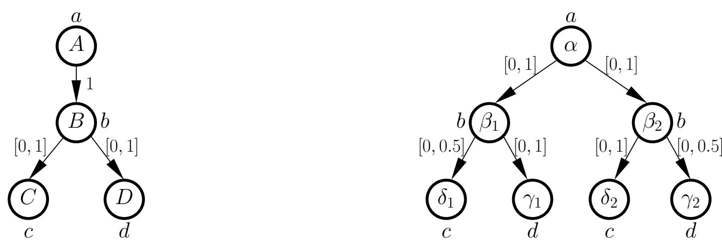

(a) IMC I_4 (b) IMC I_5

Fig. 3. IMCs I_4 and I_5 such that I_4 thoroughly but not weakly refines I_5

Granularity. In [10] an informal statement is made that the strong refinement is strictly stronger (finer) than the thorough refinement: $(\leq_{\mathsf{T}}) \supsetneq (\leq_{\mathsf{S}})$. In [7] the weak refinement is introduced, but without discussing its relations to neither the strong nor the thorough refinement. The following theorem resolves all open issues in relations between the three:

Theorem 1. *The thorough refinement is strictly weaker than the weak refinement, which is strictly weaker than the strong refinement:* $(\leq_{\mathsf{T}}) \supsetneq (\leq_{\mathsf{W}}) \supsetneq (\leq_{\mathsf{S}})$.

The first inequality is shown by exhibiting IMCs I_4 and I_5 such that I_4 thoroughly but not weakly refines I_5 (Figure 3). All implementations of I_4 satisfy I_5, but state B cannot refine any of β_1 or β_2: Let σ be a distribution admitted in B giving probability 1 to state C. Because of the interval $[0, 0.5]$ on the transition from β_1 to δ_1, at least 0.5 must be assigned to γ_1, but C and γ_1 cannot be related. A similar argument shows that B cannot refine β_2. The second inequality is shown by demonstrating two other IMCs, I_3 and I_2 such that I_3 weakly but not strongly refines I_2 (Figure 2b). State A weakly refines state α: Given a value x for the transition $A \to C$, we can split it in order to match both transitions $\alpha \xrightarrow{p \cdot x} \delta_1$ and $\alpha \xrightarrow{(1-p)\cdot x} \delta_2$. Define $\delta(C)(\delta_1) = p$ and $\delta(C)(\delta_2) = (1-p)$, with $p = 0$ if $0.2 \leq x \leq 0.4$, $p = \frac{x - 0.3}{x}$ if $0.4 < x < 0.8$, and $p = 0.6$ if $0.8 \leq x$. The correspondence function δ witnesses weak refinement between A and α. However, there is no such value of p that would work uniformly for all x, which is required by the strong refinement.

Deciding Thorough Refinement. As weak and strong refinements are strictly stronger than thorough refinement, it is interesting to investigate complexity of deciding TR. In [10] a procedure computing TR is given, albeit without a complexity class, which we establish now, closing the problem:

Theorem 2. *The decision problem TR of establishing whether there exists a thorough refinement between two given IMCs is EXPTIME-complete.*

The upper-bound in checking whether I_1 thoroughly refines I_2 is shown by observing that the complexity of the subset-simulation algorithm of [10] is $O(|Q| \cdot 2^{|P|})$, where Q and P are the set of states of I_1 and I_2, respectively (see [6]).

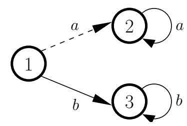

(a) A MTS M

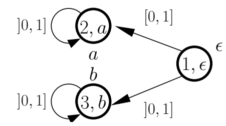

(b) The IMC $\widehat{M}$

Fig. 4. An example of the translation from Modal Transition Systems to IMCs

Summarizing, all three refinements are in EXPTIME. Still, weak refinement seems easier to check than thorough refinement. For TR, the number of iterations on the state-space of the relation is exponential while it is only polynomial for the weak refinement. Also, the constraint solved at each iteration involves a single quantifier alternation for the weak, and three alternations for the thorough refinement.

The lower bound of Theorem 2 is shown by a polynomial reduction of the thorough refinement problem for modal transition systems to TR of IMCs. The former problem is known to be EXPTIME-complete [3].

A modal transition system (an MTS in short) [15] is a tuple $M = (S, s_0, A, \rightarrow, \dashrightarrow)$, where S is the set of states, s_0 is the initial state, and $\rightarrow \subseteq S \times A \times S$ are the transitions that *must* be taken and $\dashrightarrow \subseteq S \times A \times S$ are the transitions that *may* be taken. In addition, it is assumed that $(\rightarrow) \subseteq (\dashrightarrow)$. An implementation of an MTS is a labelled transition system, i.e., an MTS where $(\rightarrow) = (\dashrightarrow)$. Formal definitions of refinement and satisfaction for MTSs are given in [6].

We describe here a translation of MTSs into IMCs which preserves implementations, while we delegate the technicalities of the proof to [6]. We assume we only work with modal transition systems that have no deadlock-states, in the sense that each state has at least one outgoing must transition. It is easy to transform two arbitrary MTSs into deadlock-free ones without affecting the thorough refinement between them [6].

The IMC $\widehat{M}$ corresponding to a MTS $M = (S, s_0, A, \rightarrow, \dashrightarrow)$ is defined by the tuple $\widehat{M} = \langle Q, q_0, A \cup \{\epsilon\}, \varphi, V\rangle$ where $Q = S \times (\{\epsilon\} \cup A)$, $q_0 = (s_0, \epsilon)$, for all $(s, x) \in Q$, $V((s, x)) = \{x\}$ and φ is defined as follows: for all $t, s \in S$ and $b, a \in (\{\epsilon\} \cup A)$, $\varphi((t, b))((s, a)) =]0, 1]$ if $t \xrightarrow{a} s$; $\varphi((t, b))((s, a)) = [0, 0]$ if $t \overset{a}{\not\dashrightarrow} s$; and $\varphi((t, b))((s, a)) = [0, 1]$ otherwise. The encoding is illustrated in Figure 4.

Now one can show that $I \models M$ iff $[\![\widehat{I}]\!] \subseteq [\![\widehat{M}]\!]$, and use this to show that the reduction preserves thorough refinement. This observation, which shows how deep is the link between IMCs and modal transition systems, is formalized in the following theorem lifting the syntactic reduction to the level of extensional semantics:

Theorem 3. *Let M and M' be two Modal Transition Systems and $\widehat{M}$ and $\widehat{M'}$ be the corresponding IMCs defined as above. We have*

$$M \leq_{\mathsf{T}} M' \iff \widehat{M} \leq_{\mathsf{T}} \widehat{M'}$$

Crucially the translation is polynomial. Thus if we had a subexponential algorithm for TR of IMCs, we could use it to obtain a subexponential algorithm for TR of MTSs, which is impossible [3].

4 Determinism

Although both are in EXPTIME, deciding weak refinement is easier than deciding thorough refinement. Nevertheless, since these two refinements do not coincide in general, a procedure to check weak refinement cannot be used to decide thorough refinement.

Observe that weak refinement has a syntactic definition very much like simulation for transition systems. On the other hand thorough refinement is a semantic concept, just as trace inclusion for transition systems. It is well known that simulation and trace inclusion coincide for deterministic automata. Similarly for MTSs it is known that TR coincides with modal refinement for deterministic objects. It is thus natural to define deterministic IMCs and check whether thorough and weak refinements coincide on these objects.

In our context, an IMC is deterministic if, from a given state, one cannot reach two states that share common atomic propositions.

Definition 5 (Determinism). *An IMC $I = \langle Q, q_0, \varphi, A, V \rangle$ is* deterministic *iff for all states $q, r, s \in Q$, if there exists a distribution $\sigma \in \varphi(q)$ such that $\sigma(r) > 0$ and $\sigma(s) > 0$, then $V(r) \neq V(s)$.*

Determinism ensures that two states reachable *with the same admissible distribution* always have different valuations. In a semantic interpretation this means that there exists no implementation of I, in which two states with the same valuation can be successors of the same source state. Another, slightly more syntactic but semantically equivalent notion of determinism is given in [6].

It is worth mentioning that deterministic IMCs are a strict subclass of IMCs. Figure 5 shows an IMC I whose set of implementations cannot be represented by a deterministic IMC.

We now state the main theorem of the section that shows that for deterministic IMCs, the weak refinement, and indeed also the strong refinement, correctly capture the thorough refinement:

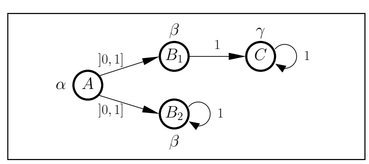

Fig. 5. An IMC I whose implementations cannot be captured by a deterministic IMC.

Theorem 4. *Given two deterministic IMCs I and I' with no inconsistent states, it holds that $I \leq_{\mathsf{T}} I'$ iff $I \leq_{\mathsf{W}} I'$ iff $I \leq_{\mathsf{S}} I'$.*

5 Common Implementation and Consistency

We now turn our attention to the problem of implementation of several IMC specifications by the same probabilistic system modeled as a Markov Chain. We start with a formal definition of the problem:

Definition 6 (Common Implementation (CI)). *Given $k > 1$ IMCs I_i, $i = 1 \dots k$, does there exist a Markov Chain C such that $C \models I_i$ for all i?*

Somewhat surprisingly we find out that, similarly to the case of TR, the CI problem is not harder for IMCs than for modal transition systems:

Theorem 5. *Deciding the existence of a CI between k IMCs is EXPTIME-complete.*

We sketch the line of argument below, delegating to [6] for details. To establish a lower bound for CI of IMCs, we reduce from CI of modal transition systems, which is known to be EXPTIME-complete [2]. For a set of modal transition systems M_i, $i = 1 \dots k$, translate each M_i, into an IMC $\widehat{M_i}$, using the same rules as in Section 3. It turns out that the set of created IMCs has a common implementation if and only if the original modal transition systems had. Since the translation is polynomial, the problem of CI for IMCs has to be at least EXPTIME-hard (otherwise it would give a sub-EXPTIME algorithm for CI of MTSs).

To address the upper bound we first propose a simple construction to check if there exists a CI for two IMCs. We start with the definition of *consistency relation* that witnesses a common implementation between two IMCs.

Definition 7. *Let $I_1 = \langle Q_1, q_0^1, \varphi_1, A, V_1 \rangle$ and $I_2 = \langle Q_2, q_0^2, \varphi_2, A, V_2 \rangle$ be IMCs. The relation $\mathcal{R} \subseteq Q_1 \times Q_2$ is a* consistency relation *on the states of I_1 and I_2 iff, whenever $(u, v) \in \mathcal{R}$, then*

- *$V_1(u) = V_2(v)$ and*
- *there exists a distribution $\rho \in \mathsf{Distr}(Q_1 \times Q_2)$ such that*
  1. *$\forall u' \in Q_1 : \sum_{v' \in Q_2} \rho(u', v') \in \varphi_1(u)(u') \ \wedge \ \forall v' \in Q_2 : \sum_{u' \in Q_1} \rho(u', v') \in \varphi_2(v)(v')$, and*
  2. *$\forall (u', v') \in Q_1 \times Q_2$, if $\rho(u', v') > 0$, then $(u', v') \in \mathcal{R}$.*

It can be shown that two IMCs indeed have a common implementation if and only if there exists a consistency relation containing their initial states. The consistency relation can be computed in polynomial time using a standard coinductive fixpoint iteration, where pairs violating Definition 7 are successively removed from $Q_1 \times Q_2$. Each iteration requires solving a polynomial number of linear systems, which can be done in polynomial time [14]. For the general problem of common implementation of k IMCs, we can extend the above definition of consistency relation to the k-ary relation in the obvious way, and the algorithm becomes exponential in the number of IMCs k, as the size of the state space $\prod_{i=1}^{k} |Q_i|$ is exponential in k.

As a side effect we observe that, exactly like MTSs, CI becomes polynomial for any constant value of k, i.e. when the number of components to be checked is bounded by a constant.

Consistency. A related problem is the one of checking consistency of a single IMC I, i.e. whether there exists a Markov chain M such that $M \models I$.

Definition 8 (Consistency (C)). *Given an IMC I, does it hold that $[\![I]\!] \neq \emptyset$?*

It turns out that, in the complexity theoretic sense, this problem is easy:

Theorem 6. *The problem C, to decide if a single IMC is consistent, is polynomial time solvable.*

Given an IMC $I = \langle Q, q_0, \varphi, A, V \rangle$, this problem can be solved by constructing a consistency relation over $Q \times Q$ (as if searching for a common implementation of I with itself). There exists an implementation of I iff there exists a consistency relation containing (q_0, q_0). Obviously, this can be checked in polynomial time.

The fact that C can be decided in polynomial time casts an interesting light on the ability of IMCs to express inconsistency. On one hand, one can clearly specify inconsistent states in IMCs (simply by giving intervals for successor probabilities that cannot be satisfied by any distribution). On the other hand, this inconsistency appears to be local. It does not induce any global constraints on implementations; it does not affect consistency of other states. In this sense IMCs resemble modal transition systems (which at all disallow expressing inconsistency), and are weaker than *mixed transition systems* [5]. Mixed transition systems relax the requirement of modal transition systems, not requiring that $(\longrightarrow) \subseteq (\dashrightarrow)$. It is known that C is trivial for modal transition systems, but EXPTIME-complete for mixed transition systems [2]. Clearly, with a polynomial time C, IMCs cannot possibly express global behaviour inconsistencies in the style of mixed transition systems, where the problem is much harder.

We conclude the section by observing that, given the IMC I and a consistency relation $\mathcal{R} \subseteq Q \times Q$, it is possible to derive a *pruned* IMC $I^* = \langle Q^*, q_0^*, \varphi^*, A, V^* \rangle$ that contains no inconsistent states and accepts the same set of implementations as I. The construction of I^* is as follows: $Q^* = \{q \in Q | (q, q) \in \mathcal{R}\}$, $q_0^* = q_0$, $V^*(q^*) = V(q^*)$ for all $q^* \in Q^*$, and for all $q_1^*, q_2^* \in Q^*$, $\varphi^*(q_1^*)(q_2^*) = \varphi(q_1^*)(q_2^*)$.

6 Related Work and Conclusion

This paper provides new results for IMCs [10] that is a specification formalism for probabilistic systems. We have studied the expressiveness and complexity of three refinement preorders for IMCs. The results are of interest as existing articles on IMCs often use one of these preorders to compare specifications (for abstraction) [10, 12, 7]. We have established complexity bounds and decision procedures for these relations, first introduced in [10]. Finally, we have studied the common implementation problem. Our solution is constructive in the sense that it can build such a common implementation.

There exist many other specification formalisms for describing and analyzing stochastic systems; the list includes process algebras [1, 16] or logical frameworks [8]. We believe that IMCs is a good unification model. A logical representation is suited for conjunction, but nor for refinement and vice-versa for process algebra. As an example, it is not clear how one can synthesize a MC (an implementation) that satisfies two Probabilistic Computation Tree Logic formulas.

In [12, 13], Katoen et al. have proposed an extension of IMCs to the continuous timed setting. It would be interesting to see our results extend to this new model.

References

[1] Andova, S.: Process algebra with probabilistic choice. In: ARTS, London, UK, Springer-Verlag (1999) 111–129

[2] Antonik, A., Huth, M., Larsen, K.G., Nyman, U., Wąsowski, A.: Modal and mixed specifications: key decision problems and their complexities. MSC **20**(01) (2010) 75–103

[3] Benes, N., Kretínský, J., Larsen, K.G., Srba, J.: Checking thorough refinement on modal transition systems is exptime-complete. In: ICTAC. (2009) 112–126

[4] Caillaud, B., Delahaye, B., Larsen, K.G., Legay, A., Pedersen, M.L., Wąsowski, A.: Compositional design methodology with constraint markov chains. In: QEST, IEEE Computer (2010)

[5] Dams, D.: Abstract Interpretation and Partition Refinement for Model Checking. PhD thesis, Eindhoven University of Technology (July 1996)

[6] Delahaye, B., Larsen, K.G., Legay, A., Pedersen, M.L., Wąsowski, A.: Decision problems for interval markov chains. `http://www.irisa.fr/s4/people/benoit.delahaye/rapports/LATA11-long.pdf` (2011)

[7] Fecher, H., Leucker, M., Wolf, V.: Don't Know in probabilistic systems. In: SPIN. Volume 3925 of LNCS, Springer (2006) 71–88

[8] Hansson, H., Jonsson, B.: A logic for reasoning about time and reliability. Formal Asp. Comput. **6**(5) (1994) 512–535

[9] Henzinger, M.R., Henzinger, T.A., Kopke, P.W.: Computing simulations on finite and infinite graphs. In: Proc. FOCS'95. (1995) 453–462

[10] Jonsson, B., Larsen, K.G.: Specification and refinement of probabilistic processes. In: LICS, IEEE Computer (1991) 266–277

[11] Jonsson, B., Larsen, K.G., Yi, W.: Probabilistic extensions of process algebras. In: Handbook of Process Algebra, Elsevier (2001) 685–710

[12] Katoen, J., Klink, D., Leucker, M., Wolf, V.: Three-valued abstraction for continuous-time Markov chains. In: CAV. Volume 4590 of LNCS, Springer (2007) 311–324

[13] Katoen, J., Klink, D., Neuhäußer, M.R.: Compositional abstraction for stochastic systems. In: FORMATS. Volume 5813 of LNCS, Springer (2009) 195–211

[14] Khachiyan, L.G.: A polynomial algorithm in linear programming. Dokl. Akad. Nauk SSSR **244**(5) (1979) 1093–1096

[15] Larsen, K.G.: Modal specifications. In: AVMS. Volume 407 of LNCS (1989) 232–246

[16] López, N., Núñez, M.: An overview of probabilistic process algebras and their equivalences. In: VSS. Volume 2925 of LNCS, Springer (2004) 89–123