Parameter Synthesis for Parametric Interval Markov Chains

Benoit Delahaye, Didier Lime, Laure Petrucci

► To cite this version:

Benoit Delahaye, Didier Lime, Laure Petrucci. Parameter Synthesis for Parametric Interval Markov Chains. 17th International Conference on Verification, Model Checking, and Abstract Interpretation, Jan 2016, Saint Petersburg, Florida, United States. pp.372 - 390, 10.1007/978-3-662-49122-5_18 . hal-01590884

HAL Id: hal-01590884

<https://hal.science/hal-01590884>

Submitted on 20 Sep 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Parameter Synthesis for Parametric Interval Markov Chains

Benoît Delahaye¹, Didier Lime², and Laure Petrucci³

¹ Université de Nantes, LINA – Nantes, France

² École Centrale de Nantes, IRCCyN – Nantes, France

³ LIPN, UMR CNRS 7030

Université Paris 13, Sorbonne Paris Cité – Paris, France

Abstract. Interval Markov Chains (IMCs) are the base of a classic probabilistic specification theory introduced by Larsen and Jonsson in 1991. They are also a popular abstraction for probabilistic systems. In this paper we study parameter synthesis for a parametric extension of Interval Markov Chains in which the endpoints of intervals may be replaced with parameters. In particular, we propose constructions for the synthesis of *all* parameter values ensuring several properties such as consistency and consistent reachability in both the existential and universal settings with respect to implementations. We also discuss how our constructions can be modified in order to synthesise *all* parameter values ensuring other typical properties.

1 Introduction

Interval Markov Chains (IMCs for short) extend Markov Chains by allowing to specify intervals of possible probabilities on transitions instead of precise probabilities. When modelling real-life systems, the exact value of transition probabilities may not be known precisely. Indeed, in most cases, these values are measured from observations or experiments which are subject to imprecision. In this case, using intervals of probabilities that take into account the imprecision of the measures makes more sense than using an arbitrary but precise value.

IMCs have been introduced by Larsen and Jonsson [22] as a *specification* formalism—a basis for a stepwise-refinement-like modelling method, where initial designs are very abstract and underspecified, and then they are made continuously more precise, until they are concrete. Unlike richer specification models such as Constraint Markov Chains [7] or Abstract Probabilistic Automata [13], IMCs are difficult to use for compositional specification due to the lack of basic modelling operators. Nevertheless, IMCs have been intensively used in order to model real-life systems in domains such as systems biology, security or communication protocols [2, 6, 25, 17]. Going further in the abstraction hierarchy, one could then assume that the endpoints of probability intervals are also imprecise.

This work has been partially supported by project PACS ANR-14-CE28-0002 and Pays de la Loire research project AFSEC.

As an example, consider that a given component can be built with arbitrary quality by using different, more or less costly, materials. This quality can be related in practice to the maximal error rate of the component, which is reflected in our design by the upper endpoint of the interval associated with a transition leading to an error state. Since this value can be chosen arbitrarily, it can be represented as a parameter. Obviously, if several instances of this component are embedded in our design, the same parameter will be used in several places. In this setting, the designer will be interested in computing the set of acceptable values for this parameter – i.e. ensuring that the overall design satisfies some given properties; or synthesising the best acceptable value for this parameter – i.e. giving the best compromise between some given (quantitative?) property and the production cost.

This new setting thus calls for methods and tools for modelling and analysing IMCs where interval endpoints are not fixed in advance.

Parametric Interval Markov Chains (pIMCs for short) have been introduced in [15] as an extension of IMCs that allows for using parameters instead of numeric values as the lower or upper endpoint of intervals. The goal of using such a model is then to synthesise parameter values ensuring correctness w.r.t. given properties. In this paper, we focus on the first basic property of such models: consistency. Consistency of a parameter valuation in a given pIMC boils down to verifying that the chosen parameter values are not incoherent, i.e. that the resulting IMC can be implemented. While [15] focuses on deciding whether a consistent parameter valuation exists in a given pIMC, we propose in this paper constructions for *synthesising all consistent parameter valuations of a given pIMC*. In addition, we also consider other objectives such as reachability or avoidability while always guaranteeing consistency. Reachability can be formulated in two flavours: either universal reachability, i.e. ensuring that all implementations reach a given set of states, or existential reachability, i.e. ensuring that there exists at least one implementation that satisfies the property. We therefore propose constructions for solving both problems while still ensuring consistency of the model.

Related work. Our work is a follow-up on [15], which is to the best of our knowledge the only existing work addressing parametric probabilistic specification theories where parameters range over probability values. In [15], we only study the consistency problem in the existential setting and propose an algorithm for deciding whether there exists at least one parameter valuation ensuring consistency for a subclass of pIMCs. In contrast, the results we provide here are fully general, and, more importantly, we attack a slightly different problem that consists in *synthesising all parameter values* ensuring consistency and reachability.

Other classes of systems where parameters give some latitude on probability distributions, such as parametric Markov models [23], have been studied in the literature [24, 19]. The activity in this domain has yielded decidability results [21], parametric probabilistic model-checking algorithms [11] and even tools [20, 12]. Continuous-time parametric and probabilistic models have also

been considered in some very restricted settings [9]. Networks of probabilistic processes where the number of processes is a parameter have also been studied in [4, 5], and probabilistic timed automata with parameters in clock constraints and invariants have been studied in [1].

In another setting, the model checking problem for Interval Markov Chains has been addressed in [10, 3, 8]. In [10, 3], the authors propose algorithms and complexity bounds for checking respectively ω -regular and LTL properties on Interval Markov Chains with closed intervals. [3] assumes that parameters can be present in the models and formulae, but these parameters do not range on the probability endpoints of the intervals, as in our work. On the other hand, [8] focuses on Interval Markov Chains with open intervals and proposes algorithms for verifying PCTL properties but does not consider parameters.

Outline. First, Section 2 recalls the basic definitions and notations of Interval Markov chains and their parametric extension. Then, Section 3 explores the consistency of Parametric Interval Markov Chains and proposes a construction for synthesising *all* the parameter valuations that guarantee consistency. The problem of existential consistent reachability is addressed in Section 4, and we show how our constructions can be adapted to solve other problems such as consistent avoidability and universal consistent reachability. Finally, Section 5 summarises the paper contributions and gives hints for future work. For space reasons, our proofs are presented in an extended version of this paper [16].

2 Background

Throughout the paper, we use the notion of parameters. A parameter $p \in P$ is a variable ranging through the interval $[0, 1]$. A valuation for P is a function $\psi : P \rightarrow [0, 1]$ that associates values with each parameter in P . We write $\text{Int}_{[0,1]}(P)$ for the set of all closed parametric intervals of the form $[x, y]$ where x, y can be either reals in the interval $[0, 1]$ or parameters from P . When $P = \emptyset$, we write $\text{Int}_{[0,1]} = \text{Int}_{[0,1]}(\emptyset)$ to denote closed intervals with real-valued endpoints. Given an interval I of the form $I = [a, b]$, $\text{Low}(I)$ and $\text{Up}(I)$ respectively denote the lower and upper endpoints of I , i.e. a and b . Given an interval $I = [a, b] \in \text{Int}_{[0,1]}$, we say that I is well-formed whenever $a \leq b$. It is worth noting that, for readability reasons, we limit ourselves to closed intervals. Nevertheless, all the results we propose can be extended with minor modifications to open/semi-open intervals whose endpoints contain linear combinations of parameters and constants.

Given a parametric interval $I \in \text{Int}_{[0,1]}(P)$ and a parameter valuation $\psi : P \rightarrow [0, 1]$, we write $\psi(I)$ for the interval of $\text{Int}_{[0,1]}$ obtained by substituting in the endpoints of I each parameter p by the value $\psi(p)$. Constraints on parameter valuations are expressions on parameter variables that restrict their potential values. Given a constraint C over P and a parameter valuation $\psi : P \rightarrow [0, 1]$, we write $\psi \models C$ when the parameter valuation satisfies constraint C . In the following, we abuse notations and identify constraints on parameter valuations with the set of parameter valuations that satisfy them. Therefore, given a constraint

C over P , we sometimes write $\psi \in C$ instead of $\psi \models C$. We also use intersections (resp. unions) of constraints to represent the set of parameter valuations satisfying their conjunction (resp. disjunction).

Given a finite set S , we denote by $\text{Dist}(S)$ the set of distributions over S , i.e. the set of functions $\rho : S \rightarrow [0, 1]$ such that $\sum_{s \in S} \rho(s) = 1$. In the rest of the paper, we assume that all the states in our structures are equipped with labels taken from a fixed set of atomic propositions A . A state-labelling function over S is thus a function $V : S \rightarrow 2^A$ that assigns to each state a set of labels in A .

2.1 Markov Chains definitions

We recall the notion of Markov Chains (MCs), that will act as models for (parametric) IMCs. An example of a Markov Chain is given in Figure 1a.

Definition 1 (Markov Chain). *A Markov Chain is a tuple $\mathcal{M} = (S, s_0, M, A, V)$, where S is a finite set of states containing the initial state s_0 , A is a set of atomic propositions, $V : S \rightarrow 2^A$ is a labeling function, and $M : S \times S \rightarrow [0, 1]$ is a probabilistic transition function such that $\forall s \in S, \sum_{t \in S} M(s, t) = 1$.*

We now recall the notion of Interval Markov Chains (IMCs), adapted from [14]. IMCs are a specification formalism that allows one to represent an infinite set of MCs. Roughly, IMCs extend MCs by replacing exact probability values on transitions with intervals of allowed probability values. An example of an IMC is given in Figure 1b.

Definition 2 (Interval Markov Chain [14]). *An Interval Markov Chain (IMC) is a tuple $\mathcal{I} = (S, s_0, \varphi, A, V)$, where S , s_0 , A and V are as for MCs, and $\varphi : S \times S \rightarrow \text{Int}_{[0,1]}$ is a transition constraint that associates with each potential transition an interval of probabilities.*

The following definition recalls the notion of satisfaction introduced in [14]. Satisfaction (also called implementation in some cases) allows to characterise the set of MCs represented by a given IMC specification. Crucially, satisfaction abstracts from the syntactic structure of transitions in IMCs: a single transition in the implementation MC can contribute to satisfaction of more than one transition in the specification IMC, by distributing its probability mass against several transitions. Similarly many MC transitions can contribute to the satisfaction of just one specification transition. This crucial notion is embedded in the so-called *correspondence function* δ introduced below. Informally, such a function is given for all pairs of states (t, s) in the satisfaction relation, and associates with each successor state t' of t – in the implementation MC – a distribution over potential successor states s' of s – in the specification IMC – specifying how the transition $t \rightarrow t'$ contributes to the transition $s \rightarrow s'$.

Definition 3 (Satisfaction Relation [14]). *Let $\mathcal{I} = (S, s_0, \varphi, A, V^I)$ be an IMC and $\mathcal{M} = (T, t_0, M, A, V^M)$ be a MC. A relation $\mathcal{R} \subseteq T \times S$ is a satisfaction relation if whenever $t \mathcal{R} s$,*

1. the labels of s and t agree: $V^M(t) = V^I(s)$,
2. there exists a correspondence function $\delta : T \rightarrow (S \rightarrow [0, 1])$ such that
 - (a) for all $t' \in T$ such that $M(t, t') > 0$, $\delta(t')$ is a distribution on S ,
 - (b) for all $s' \in S$, we have $(\sum_{t' \in T} M(t, t') \cdot \delta(t')(s')) \in \varphi(s, s')$, and
 - (c) for all $t' \in T$ and $s' \in S$, if $\delta(t')(s') > 0$, then $(t', s') \in \mathcal{R}$.

We say that state $t \in T$ satisfies state $s \in S$ (written $t \models s$) iff there exists a (minimal) satisfaction relation containing (t, s) and that $\mathcal{M}$ satisfies $\mathcal{I}$ (written $\mathcal{M} \models \mathcal{I}$) iff $t_0 \models s_0$.

The notion of satisfaction between the MC $\mathcal{M}$ from Figure 1a and the IMC $\mathcal{I}$ from Figure 1b is illustrated in Figure 1c. In this figure, we remark that the transition $1 \rightarrow 3$ in the MC $\mathcal{M}$ partly contributes to the satisfaction of transitions $A \rightarrow B$ and $A \rightarrow C$ in the IMC $\mathcal{I}$. Similarly, transitions $1 \rightarrow 2$ and $1 \rightarrow 3$ in the MC $\mathcal{M}$ both contribute to the satisfaction of transition $A \rightarrow B$ in the IMC $\mathcal{I}$.

Fig. 1: Markov Chain, Interval Markov Chain and satisfaction relation [14]

The set of MCs satisfying a given IMC $\mathcal{I}$ is written $\llbracket \mathcal{I} \rrbracket$. Formally, $\llbracket \mathcal{I} \rrbracket = \{\mathcal{M} \mid \mathcal{M} \models \mathcal{I}\}$. We say that an IMC $\mathcal{I}$ is *consistent* iff $\llbracket \mathcal{I} \rrbracket \neq \emptyset$. Although the satisfaction relation abstracts from the syntactic structure of transitions, we recall the following result from [15], that states that whenever a given IMC is consistent, it admits at least one implementation that strictly respects its structure.

Theorem 1 ([15]). *An IMC $\mathcal{I} = (S, s_0, \varphi, A, V)$ is consistent iff it admits an implementation of the form $\mathcal{M} = (S, s_0, M, A, V)$ where, for all reachable states s in $\mathcal{M}$, it holds that $M(s, s') \in \varphi(s, s')$ for all s' .*

In the following, we say that state s is consistent in the IMC $\mathcal{I} = (S, s_0, \varphi, A, V)$ if there exists an implementation $\mathcal{M} = (S, s_0, M, A, V)$ of $\mathcal{I}$ in which state s is reachable with a non-zero probability.

2.2 pIMCs and their relations to IMCs/MCs

We now recall the notion of Parametric Interval Markov Chain (pIMC), previously introduced in [15]. Intuitively, pIMCs extend IMCs by allowing parameters to be used as interval endpoints.

Definition 4 (Parametric Interval Markov Chain). A parametric Interval Markov Chain (pIMC) is a tuple $\mathcal{I}^P = (S, s_0, \varphi_P, A, V, P)$, where S , s_0 , A and V are as for IMCs, P is a set of variables (parameters) ranging over $[0, 1]$ and $\varphi_P : S \times S \rightarrow \text{Int}_{[0,1]}(P)$ associates with each potential transition a (parametric) interval.

Given a pIMC $\mathcal{I}^P = (S, s_0, \varphi_P, A, V, P)$ and a parameter valuation $\psi : P \rightarrow [0, 1]$, we write $\psi(\mathcal{I}^P)$ for the IMC obtained by replacing φ_P by the function $\varphi : S \times S \rightarrow \text{Int}_{[0,1]}$ defined by $\forall s, s' \in S, \varphi(s, s') = \psi(\varphi_P(s, s'))$. The IMC $\psi(\mathcal{I}^P)$ is called an *instance* of pIMC $\mathcal{I}^P$.

Finally, we say that a MC $\mathcal{M} = (T, t_0, M, A, V^M)$ implements pIMC $\mathcal{I}^P$, written $\mathcal{M} \models \mathcal{I}^P$, iff there exists an instance $\mathcal{I}$ of $\mathcal{I}^P$ such that $\mathcal{M} \models \mathcal{I}$. We write $\llbracket \mathcal{I}^P \rrbracket$ for the set of MCs implementing $\mathcal{I}^P$ and say that a pIMC is *consistent* iff its set of implementations is not empty.

In the rest of the paper, and in particular in examples, we sometimes omit atomic propositions in our figures and reasonings as they do not impact any of the problems we solve.

3 Consistency

When considering IMCs, one question of interest is to decide whether it is consistent without computing its set of implementations. This problem has already been addressed in the literature [14, 15], yielding polynomial decision algorithms and procedures that produce one implementation when the IMC is consistent. The same question holds for pIMCs, although in a slightly different setting. In [15], we have proposed a polynomial algorithm for deciding whether a given pIMC is consistent, in the sense that it admits at least one parameter valuation for which the resulting IMC is consistent.

Fig. 2: Consistent pIMC $\mathcal{I}^P$

Fig. 3: MC $\mathcal{M}$ implementing $\mathcal{I}^P$

Example 1. Consider pIMC $\mathcal{I}^P$ given in Figure 2. In this pIMC, parameters p and q appear in the outgoing transitions of several states, therefore the algorithm presented in [15] cannot be used in order to decide if $\mathcal{I}^P$ is consistent. From the outgoing transitions of state 4, we can extract constraints stating that the value of parameter p must be at the same time greater than 0.5 and lower than 0.3. Although state 4 is thus clearly inconsistent, $\mathcal{I}^P$ can still be consistent if there exists implementations avoiding state 4. Hence, the probability to move from state 2 to state 4 must be 0. Such an implementation is given in Figure 3 for the parameter valuation $p = q = 0.5$.

In this section, we move one step further and introduce a construction that synthesises all parameter valuations ensuring that a given pIMC is consistent. Observe that consistency is a recursive notion: a state is consistent iff there exists a distribution matching its outgoing intervals and such that all states reached through this distribution are themselves consistent. Based on this observation, we propose an inductive notion of n -consistency that follows this reasoning to a given depth n . We then build on this notion to synthesise the set of parameter valuations ensuring that a given pIMC is consistent. The section is organised as follows.

We start by introducing notions and notations that will be used throughout the rest of the paper. We then introduce the notion of n -consistency in the IMC setting, adapt it to the pIMC setting and finally present our main contribution: a construction that synthesises all parameter valuations ensuring that a given pIMC is consistent.

3.1 Notations

Let $\mathcal{I}^P = (S, s_0, \varphi_P, A, V, P)$ be a pIMC and let $s \in S$ be a state of $\mathcal{I}^P$. We say that state s is *consistent* in pIMC $\mathcal{I}^P$ if there exists an implementation $\mathcal{M} = (S, s_0, M, A, V)$ of $\mathcal{I}^P$ in which s is reachable from s_0 .

In order to decide whether a given IMC is consistent, we need to address the set of potential successors of a given state s . Obviously, this set of potential successors will depend on the values given to the parameters in $\mathcal{I}^P$. Nevertheless, we can rule out all states s' for which the interval of probabilities going from s to s' in $\mathcal{I}^P$ is $[0, 0]$. We thus write $\text{Succ}(s)$ for the set of states that can be reached from s with a probability interval not reduced to $[0, 0]$. Formally, $\text{Succ}(s) = \{s' \in S \mid \varphi_P(s, s') \neq [0, 0]\}$.

Other states of interest are the states s' for which $\varphi_P(s, s')$ is not reduced to $[0, 0]$, but that can still be avoided as successors by setting the actual probability of going from s to s' to 0 in an implementation. In order to be able to set this probability to 0, the subsequent interval must contain the value 0. As a consequence, s' must be such that $\text{Low}(\varphi_P(s, s')) = 0$ or such that the lower endpoint of the interval of probability is a parameter, i.e. $\text{Low}(\varphi_P(s, s')) \in P$. Indeed, in this case, we can force this interval to contain 0 by setting the value of its lower endpoint to 0. We thus define $LP(s) = \{s' \in \text{Succ}(s) \mid \text{Low}(\varphi_P(s, s')) \in P\}$ and $Z(s) = LP(s) \cup \{s' \in \text{Succ}(s) \mid \text{Low}(\varphi_P(s, s')) = 0\}$. Therefore, states in

$Z(s)$ can be avoided as successors of s in some implementations. We now propose a constraint on parameter valuations that ensures that a probability distribution exists that matches the outgoing intervals of s while reaching only states from a given set S' .

$$LC(s, S') = \left[\sum_{s' \in S'} \text{Up}(\varphi_P(s, s')) \geq 1 \right] \cap \left[\sum_{s' \in S'} \text{Low}(\varphi_P(s, s')) \leq 1 \right] \\ \cap \left[\bigcap_{s' \in S'} \text{Low}(\varphi_P(s, s')) \leq \text{Up}(\varphi_P(s, s')) \right]$$

Informally, $LC(s, S')$ represents all parameter valuations ensuring that all outgoing intervals of s are well-formed and that the sum of their lower endpoints is lower or equal to 1 while the sum of their upper endpoints is greater or equal to 1.

Example 2. Consider pIMC $\mathcal{I}^P$ from Figure 2. We illustrate the construction of LC for state 2 of $\mathcal{I}^P$. Let $S' = \{0, 1, 2, 3\}$. From the definition of LC , we obtain $LC(2, \{0, 1, 2, 3\}) = (p + q \geq 1) \cap (0 \leq 1) \cap (p \geq 0) \cap (q \geq 0)$. As a consequence, $\psi \in LC(2, \{0, 1, 2, 3\})$ iff $\psi(p) + \psi(q) \geq 1$.

As a clear consequence of the definition of LC , any parameter valuation ψ is in $LC(s, S')$ iff there exists a distribution in the IMC $\psi(\mathcal{I}^P)$ that avoids all states not in S' and satisfies all the intervals of probability going from s to S' .

Proposition 1. *Given a pIMC $\mathcal{I}^P = (S, s_0, \varphi_P, A, V, P)$, a state $s \in S$ and a set $S' \subseteq \text{Succ}(s)$, we have that for any parameter valuation ψ ,*

$$\psi \in LC(s, S') \iff \exists \rho \in \text{Dist}(S) \text{ s.t. } \begin{cases} \forall s' \in S \setminus S', \rho(s') = 0 \text{ and} \\ \forall s' \in S', \rho(s') \in \psi(\varphi_P(s, s')) \end{cases}$$

We remark that the intervals associated with transitions to states outside of S' are not taken into account in this proposition. Indeed, we only ensure that there exists a distribution ρ such that the intervals of probability going from s to S' are satisfied and $\rho(S \setminus S') = 0$, but we do not ensure that 0 is an admissible probability value for transitions going from s to $S \setminus S'$. Therefore S' has to be well chosen, i.e. such that $(\text{Succ}(s) \setminus S') \subseteq Z(s)$, and $LC(s, S')$ has to be accompanied with other constraints in order to ensure that 0 is an admissible probability value for transitions going outside of S' .

3.2 The notion of n -consistency for IMCs

We now introduce the notion of n -consistency in the IMC setting and then adapt this notion to pIMCs. Informally, a state s is n -consistent in IMC $\mathcal{I} = (S, s_0, \varphi, A, V)$ if there exists an unfolding of depth n starting from s for which each node admits a probability distribution satisfying all of its outgoing probability intervals. Intuitively, if one can find a sufficiently deep unfolding satisfying this property from s_0 , then the IMC is consistent. Finding the optimal depth for this unfolding is an issue, but we prove later in the section that we do not need

to go deeper than $|S|$. In practice, n -consistency is defined by induction over the structure of $\mathcal{I}$. The intuition is that state $s \in S$ is n -consistent iff there exists a distribution ρ matching its outgoing intervals, and if $n > 0$ then $\rho(s') > 0$ implies that s' is $(n - 1)$ -consistent. Unfortunately, this intuitive definition raises an issue: it may be the case that some state s' appears several times in the unfolding from s and we cannot ensure that the same outgoing distribution is chosen every time s' appears. This is problematic as we want use this unfolding in order to build an implementation respecting the structure of $\mathcal{I}$, and we therefore need to provide a unique distribution for each reachable state in S . We thus propose an alternative definition that first fixes an outgoing distribution for all states via a function $D : S \rightarrow \text{Dist}(S)$ and then enforces this distribution in the induction.

Definition 5 (n -consistency). *Let $\mathcal{I} = (S, s_0, \varphi, A, V)$ be an IMC and let $D : S \rightarrow \text{Dist}(S)$ be a function that assigns a distribution on S to each state of $\mathcal{I}$. State $s \in S$ is (n, D) -consistent iff for all $s' \in S$, $D(s)(s') \in \varphi(s, s')$, and, for $n > 0$, $D(s)(s') > 0$ implies s' is $(n - 1, D)$ -consistent.*

We say that s is n -consistent if there exists $D : S \rightarrow \text{Dist}(S)$ such that s is (n, D) -consistent.

We start with the following intuitive observation: whenever a given state is (n, D) -consistent, then it is also $(n - 1, D)$ -consistent.

Lemma 1. *Given an IMC $\mathcal{I} = (S, s_0, \varphi, A, V)$, a function $D : S \rightarrow \text{Dist}(S)$ and a state $s \in S$, for all $n > 0$, $s \in S$ is (n, D) -consistent implies $s \in S$ is $(n - 1, D)$ -consistent.*

Although the definition of n -consistency introduced above requires that a unique distribution is fixed *a priori* for all states in the IMC, we show in the following lemma that this is in fact not necessary and that the function $D : S \rightarrow \text{Dist}(S)$ can be constructed on-the-fly.

Lemma 2. *Given an IMC $\mathcal{I} = (S, s_0, \varphi, A, V)$ and a state $s \in S$, we have that for all $n > 0$, if there exists $\rho \in \text{Dist}(S)$ such that $\rho(s') \in \varphi(s, s')$ for all s' and $\rho(s') > 0$ implies that s' is $(n - 1)$ -consistent, then there exists a function $D : S \rightarrow \text{Dist}(S)$ such that $D(s) = \rho$ and s is (n, D) -consistent.*

Definition 5 is thus equivalent to the following intuitive inductive definition: a state s is n -consistent iff there exists a distribution ρ satisfying all of its outgoing probability intervals and such that for all $s' \in S$, $\rho(s') > 0$ implies that s' is $(n - 1)$ -consistent.

Example 3. Consider pIMC $\mathcal{I}^P$ from Figure 2 and two of its instances $\psi_1(\mathcal{I}^P)$ and $\psi_2(\mathcal{I}^P)$, with $\psi_1(p) = \psi_1(q) = 0.3$ and $\psi_2(p) = \psi_2(q) = 0.5$. In both IMCs, state 4 is not 0-consistent as one cannot find any distribution satisfying its outgoing intervals. On the other hand, State 2 is 0-consistent in both IMCs. State 2 is also 1-consistent in $\psi_2(\mathcal{I}^P)$ as there exists a distribution matching its intervals and avoiding State 4, but not in $\psi_1(\mathcal{I}^P)$ as any distribution satisfying the outgoing intervals of State 2 in $\psi_1(\mathcal{I}^P)$ must assign a positive probability to the transition to State 4, which is not 0-consistent.

As explained above, the intuition is that an IMC $\mathcal{I} = (S, s_0, \varphi, A, V)$ is consistent whenever one can find a sufficiently deep unfolding starting in its initial state and such that every node in this unfolding admits a probability distribution that satisfies its outgoing intervals. We show in the following lemma that the notion of n -consistency admits a fixpoint in the sense that there is a bound N for which being N -consistent is equivalent to being k -consistent for any $k \geq N$. In fact, we show that $|S|$ is an upper bound for the value of N .

Lemma 3. *Given an IMC $\mathcal{I} = (S, s_0, \varphi, A, V)$, a function $D : S \rightarrow \text{Dist}(S)$ and a state $s \in S$, for all $n \geq |S|$, s is (n, D) -consistent implies that s is $(n + 1, D)$ -consistent.*

As a consequence to Lemmas 1 and 3, we say that state s is D -consistent if it is (n, D) -consistent for some $n \geq |S|$. Similarly, we say that state s is consistent if it is D -consistent for some D .

We now propose two lemmas that link the notion of $(|S|, D)$ -consistency of the initial state of a given IMC $\mathcal{I} = (S, s_0, \varphi, A, V)$ to the existence of an implementation $\mathcal{M}$ respecting the structure of $\mathcal{I}$. The intuition of the following lemma is that the transition matrix defined in $\mathcal{M}$ is a candidate function for the $(|S|, D)$ -consistency of s_0 .

Lemma 4. *Given an IMC $\mathcal{I} = (S, s_0, \varphi, A, V)$, if (S, s_0, M, A, V) is an implementation of $\mathcal{I}$ then s_0 is $(|S|, D)$ -consistent, where $D : S \rightarrow \text{Dist}(S)$ is defined by $\forall s, s' \in S, D(s)(s') = M(s, s')$.*

Reversely, the next lemma shows that whenever s_0 is $(|S|, D)$ -consistent, then D is a candidate transition matrix for an implementation of $\mathcal{I}$ respecting its structure.

Lemma 5. *Given an IMC $\mathcal{I} = (S, s_0, \varphi, A, V)$, if s_0 is $(|S|, D)$ -consistent, then the Markov Chain (S, s_0, M, A, V) , where M is defined by $\forall s, s' \in S, D(s)(s') = M(s, s')$, is an implementation of $\mathcal{I}$.*

The following theorem follows directly from Theorem 1 and Lemmas 4 and 5 and concludes our section by stating one of our main results: a new characterisation of consistency for IMCs based on the notion of n -consistency.

Theorem 2. *Given an IMC $\mathcal{I} = (S, s_0, \varphi, A, V)$, $\mathcal{I}$ is consistent iff s_0 is $|S|$ -consistent.*

3.3 Consistency of pIMCs

We now move to the problem of consistency of pIMCs. As said earlier, our aim in this case is not only to decide whether a given pIMC is consistent, but also to synthesise all parameter valuations that ensure consistency of the resulting IMC. For this purpose, we adapt the notion of n -consistency defined above to pIMCs.

Given a pIMC $\mathcal{I}^P = (S, s_0, \varphi_P, A, V, P)$, we say that $s \in S$ is n -consistent iff there exists an IMC $\mathcal{I} = (S, s_0, \varphi, A, V)$ such that $\mathcal{I}$ is an instance of $\mathcal{I}^P$ and in which s is n -consistent. The set of parameter valuations ensuring that s is n -consistent is $\{\psi \mid s \text{ is } n\text{-consistent in } \psi(\mathcal{I}^P)\}$. We now propose a construction for the set of parameter valuations $\mathbf{Cons}_n(s)$ ensuring that a given state s in $\mathcal{I}^P$ is n -consistent. As in the previous section, this set is defined by induction on n . The intuition is as follows: a given parameter valuation ψ is in $\mathbf{Cons}_n(s)$ iff there exists a distribution ρ that matches the outgoing probability intervals of s in $\psi(\mathcal{I}^P)$ and such that it only leads to $(n-1)$ -consistent states. Because of Lemma 2, this ensures that s is indeed n -consistent in $\psi(\mathcal{I}^P)$. The existence of a distribution such as ρ is then conditioned by the set of potential successor states that can be reached from s in $\psi(\mathcal{I}^P)$. We thus start by fixing a set of states X that we want to avoid and then compute the set of valuations $\mathbf{Cons}_n^X(s)$ that ensure n -consistency of s through a distribution ρ that avoids states from X . Formally, we define $\mathbf{Cons}_n^X(s)$ as follows: let $\mathbf{Cons}_0^X(s) = LC(s, \text{Succ}(s) \setminus X) \cap \left[\bigcap_{s' \in X} \mathbf{Low}(\varphi_P(s, s')) = 0 \right]$ and for $n \geq 1$,

$$\mathbf{Cons}_n^X(s) = \left[\bigcap_{s' \in \text{Succ}(s) \setminus X} \mathbf{Cons}_{n-1}(s') \right] \cap [LC(s, \text{Succ}(s) \setminus X)] \cap \left[\bigcap_{s' \in X} \mathbf{Low}(\varphi_P(s, s')) = 0 \right]$$

The set of valuations ensuring n -consistency is then the union, for all potential choices of X , of $\mathbf{Cons}_n^X(s)$. Recall that, because of the definition of LC given at the end of Section 3.1, we need to choose X as a subset of $Z(s)$. Therefore, we define $\mathbf{Cons}_n(s) = \bigcup_{X \subseteq Z(s)} \mathbf{Cons}_n^X(s)$. We first observe that the choice of X has no impact on 0-consistency.

Lemma 6. *Let $\mathcal{I}^P = (S, s_0, \varphi_P, A, V, P)$ be a pIMC and let $s \in S$. For all $X \subseteq Z(s)$, we have $\mathbf{Cons}_0^X(s) \subseteq \mathbf{Cons}_0^\emptyset(s)$.*

As a consequence of Lemma 6 above, we have $\mathbf{Cons}_0(s) = LC(s, \text{Succ}(s))$. We illustrate the construction for $\mathbf{Cons}_n$ in the following example.

Example 4. Consider the pIMC $\mathcal{I}^P$ given in Figure 2. The computation of $\mathbf{Cons}_n$ for states 0, 1, 2 is illustrated in Figure 4. We start with computing the parameter valuations ensuring 0-consistency of all states: $\mathbf{Cons}_0(0) = \mathbf{Cons}_0(3)$ and both allow all possible parameter valuations, $\mathbf{Cons}_0(4) = (p \leq 0.3) \cap (p \geq 0.5) = \emptyset$, $\mathbf{Cons}_0(2) = (p + q + 0.5 \geq 1)$ and $\mathbf{Cons}_0(1) = (q + 0.3 \leq 1) \cap (q + 1 \geq 1) \cap (q \geq 0.3) = (q \leq 0.7) \cap (q \geq 0.3)$. Observe that for all n , we have $\mathbf{Cons}_n(s) = \mathbf{Cons}_0(s)$ for $s = 1, 3, 4$ since the value of $\mathbf{Cons}$ for their successors remains the same. We now reason on 1-consistency for state 2. By construction, its set of possibly avoidable successors is $Z(2) = \{1, 2, 4\}$, and $\mathbf{Cons}_1^X(2) = \emptyset$ when $4 \notin X$ because $\mathbf{Cons}_0(4) = \emptyset$, and also when $X = \{1, 2, 4\}$. For the other values of X , we obtain $\mathbf{Cons}_1^{\{1,4\}}(2) = \mathbf{Cons}_0(2) \cap (q \geq 1) = (p + q + 0.5 \geq 1) \cap (q \geq 1) = (q = 1)$, $\mathbf{Cons}_1^{\{2,4\}}(2) = \mathbf{Cons}_0(1) \cap (p \geq 1) = (q \leq 0.7) \cap (q \geq 0.3) \cap (p \geq 1)$ and

Fig. 4: Illustration of the construction of Cons_n in pIMC $\mathcal{I}^P$ from Figure 2

$\text{Cons}_1^{\{4\}}(2) = \text{Cons}_0(1) \cap \text{Cons}_0(2) \cap (p + q \geq 1) = (q \leq 0.7) \cap (q \geq 0.3) \cap (p + q + 0.5 \geq 1) \cap (p + q \geq 1) = (q \leq 0.7) \cap (q \geq 0.3) \cap (p + q \geq 1)$. Hence $\text{Cons}_1(2) = \bigcup_{X \subseteq Z(2)} \text{Cons}_1^X(2) = (q = 1) \cup [(q \leq 0.7) \cap (q \geq 0.3) \cap (p \geq 1)] \cup [(q \leq 0.7) \cap (q \geq 0.3) \cap (p + q \geq 1)] = (q = 1) \cup [(q \leq 0.7) \cap (q \geq 0.3) \cap (p + q \geq 1)]$. Furthermore, we can show that $\text{Cons}_n(2) = \text{Cons}_1(2)$ for all $n \geq 1$. Similarly, we can show that $\text{Cons}_1(0) = (p + q \geq 0.5) \cup [(q \leq 0.7) \cap (q \geq 0.3)]$, and $\text{Cons}_n(0) = \text{Cons}_2(0) = [(q \leq 0.7) \cap (q \geq 0.3)] \cup (q = 1)$ for all $n \geq 2$.

Our aim is now to prove that $\text{Cons}_n(s)$ contains exactly all parameter valuations ensuring that s is n -consistent. We first show that $\text{Cons}_n^X(s)$ works as intended, i.e. contains exactly all parameter valuations ψ ensuring that s is n -consistent in $\psi(\mathcal{I}^P)$ while using a distribution that avoids X .

Lemma 7. *Given a pIMC $\mathcal{I}^P = (S, s_0, \varphi_P, A, V, P)$, a state $s \in S$, a set $X \subseteq Z(s)$ and a parameter valuation $\psi : P \rightarrow [0, 1]$, we have $\psi \in \text{Cons}_n^X(s)$ iff there exists a function $D : S \rightarrow \text{Dist}(S)$ such that $\forall s, s', s' \in X$ implies $D(s)(s') = 0$ and state s is (n, D) -consistent in the IMC $\psi(\mathcal{I}^P)$.*

A direct consequence of Lemma 7 above is that $\text{Cons}_n(s)$ contains exactly all parameter valuations ensuring that s is n -consistent.

Proposition 2. *Given a pIMC $\mathcal{I}^P = (S, s_0, \varphi_P, A, V, P)$, a state $s \in S$ and a parameter valuation $\psi : P \rightarrow [0, 1]$, we have $\psi \in \text{Cons}_n(s)$ iff s is n -consistent in the IMC $\psi(\mathcal{I}^P)$.*

It directly follows from Lemma 1 and Proposition 2 that for all $n \geq 1$ and $s \in S$, $\text{Cons}_n(s) \subseteq \text{Cons}_{n-1}(s)$, i.e. that each computation step restricts the sets of parameter valuations.

We conclude this section by our main result, which follows directly from Proposition 2 and Theorem 2: the set $\mathbf{Cons}_{|S|}(s_0)$ contains exactly all parameter valuations ensuring that the pIMC $\mathcal{I}^P = (S, s_0, \varphi_P, A, V, P)$ is consistent.

Theorem 3. *Given a pIMC $\mathcal{I}^P = (S, s_0, \varphi_P, A, V, P)$ and a parameter valuation $\psi : P \rightarrow [0, 1]$, we have $\psi \in \mathbf{Cons}_{|S|}(s_0)$ iff the IMC $\psi(\mathcal{I}^P)$ is consistent.*

One can therefore compute the set of parameter valuations ensuring that a given pIMC $\mathcal{I}^P = (S, s_0, \varphi_P, A, V, P)$ is consistent by computing $\mathbf{Cons}_{|S|}(s_0)$. If the parameters are chosen inside $\mathbf{Cons}_{|S|}(s_0)$, the resulting IMC is consistent: it admits at least one implementation that avoids all inconsistent states.

Example 5. In our running example, $\mathbf{Cons}_5(0) = (0.3 \leq q \leq 0.7) \cup (q = 1)$. Hence, the IMC is consistent for all values of q satisfying this condition and any value of p .

Regarding complexity, if, for instance, we represent the sets of parameters by finite unions of systems of linear inequalities, basic set operations like intersection are polynomial in the number of parameters. Then computing $\mathbf{Cons}_0(s)$ for all $s \in S$ is polynomial in the number of parameters, as well as, given some X, n and s , computing $\mathbf{Cons}_n^X(s)$. There are $|S|$ states and here n can also take at most $|S|$ successive values. Set X however is chosen in $Z(s)$ for each s . So there are up to $2^{|Z(s)|}$ possible choices for X . Now, remark that $|Z(s)|$ is typically small compared to $|S|$ but, in the worst case, it can be equal to $|S|$. So the worst case asymptotic complexity of the algorithm is exponential in the number of states of the pIMC.

In the following, we write $\mathbf{Cons}(s)$ (resp. $\mathbf{Cons}^X(s)$) for the sets $\mathbf{Cons}_{|S|}(s)$ (resp. $\mathbf{Cons}_{|S|}^X(s)$).

4 Consistent reachability

Another interesting problem for IMCs and pIMCs is consistent reachability. This problem can be declined in two flavours: existential and universal. Given an IMC $\mathcal{I} = (S, s_0, \varphi, A, V)$ and a target set of states $G \subseteq S$, existential consistent reachability amounts to deciding whether there exists an implementation $\mathcal{M}$ respecting the structure of $\mathcal{I}$ in which G is reachable from s_0 with a non-zero probability. Dually, universal consistent reachability amounts to deciding whether the set G is reachable from s_0 with a non-zero probability in *all implementations* respecting the structure of $\mathcal{I}$. When moving to pIMCs, as in the previous section, we are interested in synthesising all parameter valuations ensuring that a given set of states is universal/existential consistent reachable in the resulting IMCs. In this section, we first focus on the existential problem and start with providing a construction that allows for deciding the existential consistent reachability problem for IMCs. We then adapt this construction to the pIMC setting and finally discuss how this construction can be adapted in order to solve the universal consistent reachability problem for IMCs/pIMCs.

4.1 Existential Consistent Reachability for IMCs

Given an IMC $\mathcal{I} = (S, S_0, \varphi, A, V)$, we say that a target set $G \subseteq S$ is *existential consistent reachable* in $\mathcal{I}$ iff there exists an implementation $\mathcal{M} = (S, s_0, M, A, V)$ of $\mathcal{I}$ in which the probability of reaching G from s_0 is strictly greater than 0. Formally, there must exist a path $s_0 \rightarrow \dots \rightarrow s_n$ in $\mathcal{M}$ where $M(s_i, s_{i+1}) > 0$ for all $0 \leq i < n$ and $s_n \in G$. We insist on the word *consistent* because it is not only important that there exists a sequence of transitions with positive probability matching the intervals in $\mathcal{I}$ and reaching G , but also that this sequence can be mimicked in an implementation, i.e. that the chosen probability distributions do not violate other intervals or do not impose that inconsistent states are also reached. In the following, when clear from the context, we sometimes omit the words “existential consistent” and only say that G is reachable in $\mathcal{I}$.

Notice that our definition of existential consistent reachability only takes into account implementations that respect the structure of $\mathcal{I}$. Although this looks like a limitation, the following theorem shows that any implementation $\mathcal{M}$ of $\mathcal{I}$ can be turned into an implementation $\tilde{\mathcal{M}}$ that respects the structure of $\mathcal{I}$ and that is equivalent to $\mathcal{M}$ with respect to consistent reachability.

Theorem 4. *Let $\mathcal{I} = (S, s_0, \varphi, A, V)$ be an IMC and $G \subseteq S$ be a target set of states. For all MC $\mathcal{M} = (T, t_0, M, A, V_M) \in \llbracket \mathcal{I} \rrbracket$, there exists an MC $\tilde{\mathcal{M}} = (S, s_0, \tilde{M}, A, V) \in \llbracket \mathcal{I} \rrbracket$ such that G is reachable in $\mathcal{M}$ iff $\{t \in T \mid \exists s \in G, t \models s\}$ is reachable in $\tilde{\mathcal{M}}$.*

Since the problem of existential consistent reachability mixes the notions of consistency and reachability, we cannot separate these two notions. For consistency of a given state s , one has to show that there exists a distribution matching the outgoing intervals of s and reaching only consistent states. On the other hand, for reachability of G , one has to show that there exists a distribution that reaches a state s' from which G is reachable. The difficulty here is that we have to make sure that the same distribution is chosen for both problems, not only in state s but also in all the states that are reached both through the unfolding inherent to consistency and through the path inherent to reachability. As for consistency, we thus propose to start by fixing a unique outgoing distribution for all states in S with a function $D : S \rightarrow \text{Dist}(S)$ and enforce that these distributions have to be chosen in our inductive definition of consistent existential reachability.

Formally, we say that $G \subseteq S$ is $(0, D)$ -reachable from $s \in S$ iff s is D -consistent and $s \in G$. For $n > 0$, G is (n, D) -reachable from s iff s is D -consistent and either $s \in G$ or there exists s' such that $D(s)(s') > 0$ and G is $(n - 1, D)$ -reachable from s' . The intuition is that G is (n, D) -reachable from s if s is consistent and G can be reached in at most n steps from s using distributions from D . We then say that G is n -reachable from s if there exists a function $D : S \rightarrow \text{Dist}(S)$ such that G is (n, D) -reachable from s .

As for consistency, we can also provide another equivalent definition for n -reachability in which the function $D : S \rightarrow \text{Dist}(S)$ is constructed on the fly: $G \subseteq S$ is n -reachable from $s \in S$ iff either $s \in G$ and s is consistent, or there exists a

distribution matching the outgoing intervals of s , reaching only consistent states and at least one state s' from which G is $(n - 1)$ -reachable.

We start with the following intuitive observation: whenever G can be reached in at most n steps from s through D , it can also be reached in at most k steps for any $k \geq n$. This is formalised in the following lemma.

Lemma 8. *Let $\mathcal{I} = (S, s_0, \varphi, A, V)$ be an IMC, $G \subseteq S$ a target set of states and $D : S \rightarrow \text{Dist}(S)$ a function that associates a distribution on S with all states. We have that for all $n \geq 0$ and $s \in S$, if G is (n, D) -reachable from s then G is $(n + 1, D)$ -reachable from s .*

From our definitions, we can say that G is reachable in $\mathcal{I}$ iff there exists N such that G is N -reachable from the initial state s_0 . Intuitively, we expect that $N \leq |S|$, i.e. that if a path of length at most $|S|$ leading to G cannot be found, then there is no hope of finding a longer path leading to G . This result is formalised in the following lemma.

Lemma 9. *Given an IMC $\mathcal{I} = (S, s_0, \varphi, A, V)$ and a target set $G \subseteq S$, G is existential consistent reachable in $\mathcal{I}$ iff G is $|S|$ -reachable from s_0 .*

We thus conclude that our construction for n -reachability allows deciding in a linear number of iterations whether a given set G is reachable in the IMC $\mathcal{I}$.

4.2 Existential Consistent Reachability for pIMCs

We now move to the pIMC setting. As said previously, given a pIMC $\mathcal{I}^P = (S, s_0, \varphi_P, A, V, P)$ and a target set of states $G \subseteq S$, our aim is to compute the set of parameter valuations ψ ensuring that there exists an implementation of IMC $\psi(\mathcal{I}^P)$ in which G is reachable. We proceed as for the consistency problem presented in the previous section: we propose a construction based on the notion of n -reachability for IMCs that, for each state $s \in S$, inductively constructs a set of parameter valuations $\text{Reach}_n^G(s)$ that eventually converges to the desired set. The intuition is similar to the construction for $\text{Cons}_n(s)$: we first select a set $X \subseteq Z(s)$ of states that we want to avoid and define the set of valuations $\text{Reach}_n^{G,X}(s)$ that ensure that G can be reached from s in at most n steps with a distribution that avoids X while preserving consistency. In the rest of the section, we use the constraint on parameters $(s \in G)$ with the following meaning: $(s \in G)$ is empty if $s \notin G$ and universal otherwise. We formally define $\text{Reach}_n^{G,X}(s)$ for all $s \in S$, $n \geq 0$ and $X \subseteq Z(s)$ inductively as follows: $\text{Reach}_0^{G,X}(s) = \text{Cons}^X(s) \cap (s \in G)$, and for $n > 0$

$$\text{Reach}_n^{G,X}(s) = \text{Cons}^X(s) \cap \left[(s \in G) \cup \bigcup_{s' \in \text{Succ}(s) \setminus X} \text{Reach}_{n-1}^G(s') \cap \text{Up}(\varphi_P(s, s')) > 0 \cap \sum_{s'' \neq s'} \text{Low}(\varphi_P(s, s'')) < 1 \right]$$

Informally, $\text{Reach}_0^{G,X}(s)$ is empty if $s \notin G$ and contains exactly all parameter valuations ensuring that s is consistent while avoiding X otherwise. For $n > 0$,

$\text{Reach}_n^{G,X}(s)$ either contains exactly all parameter valuations ensuring that s is consistent while avoiding X if $s \in G$ or all parameter valuations ensuring that s is consistent while avoiding X and that G is reachable in at most $n-1$ steps from at least one potential successor s' of s not in X that can be reached in one step from s with a strictly positive probability. In some sense, choosing a given set X constrains the structure of the implementations we are looking for. Since we are attacking the problem of *existential* consistent reachability, we therefore need to explore every possible choice for X , and return all parameter valuations ensuring the property for at least one set X . We thus define $\text{Reach}_n^G(s)$ as the union, for all potential choices of X , of $\text{Reach}_n^{G,X}(s)$: $\text{Reach}_n^G(s) = \bigcup_{X \subseteq Z(s)} \text{Reach}_n^{G,X}(s)$. Remark that, for $n = 0$, we obviously have $\text{Reach}_0^G(s) = \text{Cons}(s) \cap (s \in G)$.

We show in the following lemma that the definition of $\text{Reach}_n^{G,X}(s)$ is faithful to our intuition and contains exactly all parameter valuations ψ ensuring that G is n -reachable from s while avoiding X in the IMC $\psi(\mathcal{I}^P)$.

Lemma 10. *Given a pIMC $\mathcal{I}^P = (S, s_0, \varphi_P, A, V, P)$, a state $s \in S$, a target set of states $G \subseteq S$, $X \subseteq Z(s)$ and $n \geq 0$, $\psi \in \text{Reach}_n^{G,X}(s)$ iff there exists a function $D : S \rightarrow \text{Dist}(S)$ such that $D(s)(s') = 0$ for all $s' \in X$ and G is (n, D) -reachable from s in the IMC $\psi(\mathcal{I}^P)$.*

A direct consequence of Lemma 10 is the following proposition, stating that $\text{Reach}_n^G(s)$ contains exactly all the parameter valuations ψ ensuring that G is n -reachable from s in the IMC $\psi(\mathcal{I}^P)$.

Proposition 3. *Given a pIMC $\mathcal{I}^P = (S, s_0, \varphi_P, A, V, P)$, a state $s \in S$, a target set of states $G \subseteq S$ and $n \geq 0$, $\psi \in \text{Reach}_n^G(s)$ iff G is n -reachable from state s in the IMC $\psi(\mathcal{I}^P)$.*

Based on Proposition 3 and Lemma 9, we conclude with the following theorem that shows that the set of parameter valuations ensuring existential consistent reachability can be computed in a linear number of iterations using our construction.

Theorem 5. *Given a pIMC $\mathcal{I}^P = (S, s_0, \varphi_P, A, V, P)$ and a target set $G \subseteq S$, $\text{Reach}_{|S|}^G(s_0)$ is the exact set of parameter values such that G is reachable in $\mathcal{I}^P$.*

4.3 Consistent Avoidability and Universal Consistent Reachability

We now briefly show how the results presented in this paper can be adapted to universal consistent reachability, i.e. the problem of synthesising all parameter valuations ensuring that a set G is reachable in *all implementations* of the corresponding instances of a given pIMC $\mathcal{I}^P$. We first start with a related problem, consistent avoidability, and then build a solution to the universal consistent reachability problem from the proposed solution.

Consistent Avoidability. Given an IMC $\mathcal{I} = (S, s_0, \varphi, A, V)$, we say that a set $G \subseteq S$ is consistent avoidable in $\mathcal{I}$ iff $\mathcal{I}$ is consistent and there exists an

implementation $\mathcal{M}$ respecting the structure of $\mathcal{I}$ in which G is not reachable from s_0 . Given a pIMC $\mathcal{I}^P = (S, s_0, \varphi_P, A, V, P)$, we want to synthesise all parameter valuations ψ such that $G \subseteq S$ is consistent avoidable in $\psi(\mathcal{I}^P)$. The construction for consistent avoidability resembles the construction for consistency presented in Section 3. Intuitively, consistency is an avoidability property, in which we want to avoid the locally inconsistent states. We therefore need only to update our notion of local consistency: formally, we say that G is 0-avoidable from s if $s \notin G$ and s is 0-consistent. For $n > 0$, we say that G is n -avoidable from s if $s \notin G$ and there exists a distribution ρ satisfying the outgoing intervals of s and reaching only states from which G is $(n - 1)$ -avoidable. Following the same reasoning as in Section 3, we can show that, given an IMC $\mathcal{I} = (S, s_0, \varphi, A, V)$ and a set $G \subseteq S$, G is avoidable in $\mathcal{I}$ iff G is $|S|$ -avoidable from s_0 .

In the pIMC setting, we proceed similarly: we directly use the formula for $\text{Cons}_n(s)$ replacing all occurrences of $LC(s, S')$, for any s and S' , with $LC(s, S') \cap (s \notin G)$. We thus define the new operator $\text{Avoid}_n^G(s)$, for all $n \geq 0$ and all states s of the pIMC. It is then easy to show that the set $\text{Avoid}_{|S|}^G(s_0)$, hereafter written just $\text{Avoid}^G(s_0)$, represents the desired set of parameter valuations, i.e. exactly all parameter valuations ψ ensuring that G is consistent avoidable in $\psi(\mathcal{I}^P)$.

Universal Consistent Reachability. Given an IMC $\mathcal{I} = (S, s_0, \varphi, A, V)$ and a target set of states $G \subseteq S$, we say that G is *universal consistent reachable* in $\mathcal{I}$ iff G is reachable from s_0 in *all implementations* respecting the structure of $\mathcal{I}$. In the pIMC setting, our aim is to synthesise all parameter valuations ensuring that a given target set of states G is universal consistent reachable in the resulting IMCs. This set can be directly derived from the constructions proposed in the previous sections. Indeed, the complement set of Avoid^G as presented above represents all the parameter valuations ensuring either that the resulting IMC is inconsistent or that the set G is reachable in all implementations of the resulting IMC. Therefore, given a pIMC $\mathcal{I}^P = (S, s_0, \varphi_P, A, V, P)$ and a target set of states $G \subseteq S$, we can define $\text{uReach}^G(s_0) = \text{Cons}(s_0) \cap \text{Avoid}^G(s_0)$ and show that $\text{uReach}^G(s_0)$ contains exactly all parameter valuations ψ ensuring that G is universal consistent reachable in $\psi(\mathcal{I}^P)$.

5 Conclusion and future work

In this paper, we have explored the problem of consistency of pIMCs, an extension of Interval Markov Chains that allows parameters as endpoints of the intervals. Indeed, parameter valuations must satisfy constraints so that all the outgoing intervals of reachable states are well-formed and the sum of their endpoints surround 1. We show that such consistency constraints can be iteratively explored, solved and combined, thus synthesising all parameter values ensuring consistency. A similar approach also applies to consistent reachability and avoidability problems.

The properties in this paper give a good view of how to proceed to synthesise parameters in order to guarantee consistency and reachability. Future work will aim at providing efficient algorithms and heuristics for pIMCs exploration.

References

1. André, É., Fribourg, L., Sproston, J.: An extension of the inverse method to probabilistic timed automata. *Formal Methods in System Design* 42(2), 119–145 (2013)
2. Barbuti, R., Levi, F., Milazzo, P., Scatena, G.: Probabilistic model checking of biological systems with uncertain kinetic rates. *Theor. Comput. Sci.* 419(0), 2 – 16 (2012)
3. Benedikt, M., Lenhardt, R., Worrell, J.: LTL model checking of interval Markov chains. In: *TACAS. LNCS*, vol. 7795, pp. 32–46. Springer (2013)
4. Bertrand, N., Fournier, P.: Parameterized verification of many identical probabilistic timed processes. In: *FSTTCS. LIPIcs*, vol. 24, pp. 501–513 (2013)
5. Bertrand, N., Fournier, P., Sangnier, A.: Playing with probabilities in reconfigurable broadcast networks. In: *FoSSaCS. LNCS*, vol. 8412, pp. 134–148. Springer (2014)
6. Biondi, F., Legay, A., Nielsen, B., Wasowski, A.: Maximizing entropy over Markov processes. In: *LATA. LNCS*, vol. 7810, pp. 128–140. Springer (2013)
7. Caillaud, B., Delahaye, B., Larsen, K., Legay, A., Pedersen, M., Wasowski, A.: Constraint Markov chains. *Theor. Comput. Sci.* 412(34), 4373–4404 (2011)
8. Chakraborty, S., Katoen, J.P.: Model checking of open interval Markov chains. In: *ASMTA. LNCS*, vol. 9081, pp. 30–42. Springer (2015)
9. Chamseddine, N., DufLOT, M., Fribourg, L., Picaronny, C., Sproston, J.: Computing expected absorption times for parametric determinate probabilistic timed automata. In: *QEST*. pp. 254–263. IEEE Computer Society (2008)
10. Chatterjee, K., Sen, K., Henzinger, T.A.: Model-checking omega-regular properties of interval Markov chains. In: *FOSSACS. LNCS*, vol. 4962, pp. 302–317. Springer (2008)
11. Daws, C.: Symbolic and parametric model checking of discrete-time Markov chains. In: *ICTAC. LNCS*, vol. 3407, pp. 280–294. Springer (2004)
12. Dehnert, C., Junges, S., Jansen, N., Corzilius, F., Volk, M., Brintjes, H., Katoen, J.P., Abraham, E.: Prophecy: A probabilistic parameter synthesis tool. In: *CAV. LNCS*, Springer (2015)
13. Delahaye, B., Katoen, J., Larsen, K., Legay, A., Pedersen, M., Sher, F., Wasowski, A.: Abstract probabilistic automata. *Inf. Comput.* 232, 66–116 (2013)
14. Delahaye, B., Larsen, K., Legay, A., Pedersen, M., Wasowski, A.: Consistency and refinement for interval Markov chains. *J. Log. Algebr. Program.* 81(3), 209–226 (2012)
15. Delahaye, B.: Consistency for parametric interval Markov chains. In: *SynCoP. OASICS*, Schloss Dagstuhl (2015)
16. Delahaye, B., Lime, D., Petrucci, L.: Parameter Synthesis for Parametric Interval Markov Chains. HAL research report hal-01219823 (2015)
17. Fioriti, L.F., Hahn, E., Hermanns, H., Wachter, B.: Variable probabilistic abstraction refinement. In: *ATVA. LNCS*, vol. 7561, pp. 300–316. Springer (2012)
18. Gori, R., Levi, F.: An analysis for proving probabilistic termination of biological systems. *Theor. Comput. Sci.* 471(0), 27 – 73 (2013)
19. Hahn, E., Han, T., Zhang, L.: Synthesis for PCTL in parametric Markov decision processes. In: *NSFM. LNCS*, vol. 6617, pp. 146–161. Springer (2011)
20. Hahn, E., Hermanns, H., Wachter, B., Zhang, L.: PARAM: A model checker for parametric Markov models. In: *CAV. LNCS*, vol. 6174, pp. 660–664. Springer (2010)

21. Hahn, E., Hermanns, H., Zhang, L.: Probabilistic reachability for parametric Markov models. *Software Tools for Technology Transfer* 13(1), 3–19 (2011)
22. Jonsson, B., Larsen, K.: Specification and refinement of probabilistic processes. In: *LICS*. pp. 266–277. IEEE Computer (1991)
23. Lanotte, R., Maggiolo-Schettini, A., Troina, A.: Decidability results for parametric probabilistic transition systems with an application to security. In: *SEFM*. pp. 114–121. IEEE Computer Society (2004)
24. Lanotte, R., Maggiolo-Schettini, A., Troina, A.: Parametric probabilistic transition systems for system design and analysis. *Formal Aspects of Computing* 19(1), 93–109 (2007)
25. Sen, K., Viswanathan, M., Agha, G.: Model-checking Markov chains in the presence of uncertainties. In: *TACAS. LNCS*, vol. 3920, pp. 394–410. Springer (2006)