

Logic and Model Checking by Imprecise Probabilistic Interpreted Systems

Alberto Termine^{1*}, Alessandro Antonucci², Giuseppe Primiero^{1**}, and
Alessandro Facchini^{2***}

¹ Department of Philosophy, University of Milan, Milan, Italy.
`alberto.termine@unimi.it`, `giuseppe.primiero@unimi.it`

² Istituto Dalle Molle di Studi sull'Intelligenza Artificiale (IDSIA), USI-SUPSI,
Lugano, Switzerland.
`alessandro@idsia.ch`, `alessandro.facchini@idsia.ch`

Abstract Stochastic multi-agent systems raise the necessity to extend probabilistic model checking to the epistemic domain. Results in this direction have been achieved by epistemic extensions of Probabilistic Computation Tree Logic and related Probabilistic Interpreted Systems. The latter, however, suffer of an important limitation: they require the probabilities governing the system's behaviour to be *fully* specified. A promising way to overcome this limitation is represented by *imprecise probabilities*. In this paper we introduce imprecise probabilistic interpreted systems and present a related logical language and model-checking procedures based on recent advances in the study of imprecise Markov processes.

Keywords: Probabilistic Interpreted Systems · Imprecise Markov Chains · Imprecise Probabilities · Model Checking.

1 Introduction

Probabilistic model checking arises in connection with the specification and verification of computational systems of stochastic nature. Broadly speaking, it includes a series of languages for specifying probabilistic properties of stochastic systems and relative semantics based on Markov models [3]. Notable examples of logical languages for property specification are PCTL [14], its extensions (PCTL*, PRCTL) and CSL [3]. Probabilistic model checking has been applied to many different fields, such as software verification [7], communication protocols [6], and even computational biology [8,5].

In recent years, the increasing relevance of stochastic multi agent systems (MAS for short) has raised the necessity of extending probabilistic model checking to languages endowed with epistemic modalities. Given its popularity, it has

* Corresponding Author.

** Supported by the project “Departments of Excellence 2018-2022”, Ministry of Education, University and Research (MIUR).

*** Supported by the Hasler foundation grant n. 20061.

been natural to propose *epistemic* extensions of probabilistic model checking, in particular within the field of MAS verification. Regarding property specification, these extensions result in a series of languages merging PCTL and standard epistemic operators for both single and multi-agent knowledge and belief (e.g. in [26,9]). Regarding model specification, these extensions exploit the formalism of so-called *Probabilistic Interpreted Systems* (PIS), a class of structures obtained by merging standard interpreted systems [24] with Markov models [3]. For example, [26] introduces the logic PCTLK to represent probabilistic knowledge in stochastic MASs. The logic is conceived to merge PCTL and epistemic operators, the former modelled through a probabilistic state-transition matrix, as in standard PCTL, the latter modelled by epistemic accessibility relations as in canonical interpreted systems. Furthermore, the paper proves the reducibility of relevant model-checking tasks for PCTLK to more standard model-checking procedures implementable in *PRISM*, the canonical software tool to model-check probabilistic systems [18].

Despite their success, standard probabilistic model-checking and its epistemic extensions suffer the limitation of requiring the probabilities describing the system behaviour to be precisely defined. This represents a problem especially for epistemic domains, because it is tantamount to assuming that an agent *always* knows *precisely* all the probability values describing each possible state-transition of the system. In other words, it is impossible to model agents with an high-order uncertainty about transition probabilities. This is the case, for instance, when agents in a MAS (partially) ignore the stochastic behaviour of other agents in the system. A possible way to overcome this limitation is represented by so-called parametric Markov models [2,10], which replace precise probabilities with unknown parameters. In [2] for instance, the authors introduce an extension of PCTL specific for parametric Markov chains. The complexity of the corresponding model-checking procedure, based on *fraction-free Gaussian elimination*, is however exponential in the number of states of the models, hence limiting its applicability only to models of small size.

An alternative, but poorly explored, approach is offered by the formalism of imprecise probabilities [25] and, for what interests us here, related imprecise Markov models such as imprecise Markov chains (IMC) [23,13,16]. Roughly, IMCs are the imprecise counterparts of precise Markov chains obtained by replacing precise probability distributions with so-called Credal sets, i.e., sets of probability distributions describing the model and compatible with some specific constraints given by the agents [11].

A first attempt to extend probabilistic model-checking to imprecise probabilities has been proposed in [23], which introduces an *imprecise* PCTL with a semantics based on discrete-time imprecise Markov chains (IMC). The language for properties specification is obtained by replacing the standard probability operator with an operator for representing lower and upper bounds of imprecise probability distributions. Model checking with respect to the new probability operators is reduced to the computation of lower and upper bounds of marginal probabilities on an IMC. These bounds are computed efficiently by means of

specific transition operators whose applicability is an optimisation task solvable through linear programming. This approach enables the authors to verify that shifting from precise to imprecise probabilistic models does not affect the overall complexity of the most relevant model-checking procedures. Another example of imprecise probabilistic model-checking is offered in [21], which proposes a semantics and corresponding model-checking procedures based on imprecise Markov reward models. Differently from [23], the model-checking procedures outlined by [21] are based on an algorithm proposed by [22] for computing, among others, lower and upper bounds of *hitting probabilities*. The present work extends the results presented in [21] with the as yet unexplored application of model-checking with *imprecise* probabilistic models to the *epistemic domain*.

The paper is structured as follows. In Section 2 we recap some background knowledge about Markov models and their imprecise counterparts. In Section 2.2 we introduce *imprecise probabilistic interpreted systems*, a new kind of structures conceived as the imprecise counterparts of the probabilistic interpreted systems proposed in [26]. In Section 3 we introduce a new language, called *epistemic imprecise PCTL* (EIPCTL) extending standard PCTL with new *imprecise-probabilistic*, *epistemic* and *doxastic* operators. In Section 3.2 we introduce a proper semantics for EIPCTL based on imprecise probabilistic interpreted systems. In Section 4 we discuss relevant procedures for model checking imprecise probabilistic interpreted systems against EIPCTL formulae. Interestingly, we verify that shifting to imprecise probabilities does not affect the overall computational complexity of the relevant model-checking tasks, which therefore remains polynomial in the number of states of the models. In Section 5, we propose a simple illustrative example. Finally, in Section 6 we underline some conclusive remarks about future extensions.

2 Background

Let \mathcal{S} denote a finite non-empty set of possible states and S a variable taking its values from \mathcal{S} . A *probability mass-function* (PMF) over S , denoted as $P(S)$, is a non-negative normalized real map defined over \mathcal{S} . Furthermore, given a real-values function f of S , its *expectation* based on $P(S)$, denoted $E_P[f]$, is defined as: $E[f] := \sum_{s \in \mathcal{S}} f \cdot P(S)$. A *joint* PMF $P(S', S)$ is a PMF that gives for each pair of states s, s' the probability that s and s' jointly occur. A *conditional* PMF instead, defined as $P(S'|s) := \{P(s', s)/P(s)\}_{\forall s' \in \mathcal{S}}$, is a PMF that assigns to each $s' \in \mathcal{S}$ the probability of s' to occur given that s occurred. Furthermore, if $P(s'|s) = P(s')$ for each $s' \in \mathcal{S}$, we say that S' is (stochastically) *irrelevant* to s . It is easy to check that stochastic irrelevance and independence are equivalent.

A *Credal Set* (CS) over S , denoted by $K(S)$, is a collection of PMFs over S compatible with some given constraints. We consider here only finitely generated CSs, i.e., CSs whose convex hull has only a finite number of extreme points. Given a function f of S , its *upper expectation* with respect to $K(S)$ is defined as $\bar{E}_P[f] := \sup_{P(S) \in K(S)} E_P[f]$ while the lower expectation is defined as $\underline{E} := \inf_{P(S) \in K(S)} E_P[f]$. Furthermore, as stated by [11], suprema (infima)

of upper (lower) expectations can be equivalently reduced to maxima (mimima) over the extreme points of the CS convexification. Consequently, we can identify a CS with the extreme points of its convex hull. Given a joint CS $K(S, S')$, defined as a collection of joint PMFs $P(S, S')$, the (marginal) CS $K(S')$ is obtained by element-wise marginalization of S on its elements. At the same time, if $P(s) > 0$ for each $P(S) \in K(S)$, the conditional CS $K(S'|s)$ can be obtained by element-wise conditioning. Finally, we say that S is *epistemically* irrelevant to S' if and only if $K(S'|s) = K(S')$ for each $s \in \mathcal{S}$. It is easy to check that epistemic irrelevance is the generalization to the formalism of CSs of the standard notion of stochastic irrelevance. Although in the present paper we decide to adopt *epistemic irrelevance* as the standard generalization of stochastic irrelevance, other possible ways of representing it are based on stochastic independence of the CS elements, or of the CS extreme points (usually called, respectively, *strict* and *strong* irrelevance) [12]. All these notions are equivalent for unconditional queries [17], as it is the case for most of the inferential tasks considered in this work.

2.1 Markovian Models

Precise Markov Chains. A precise discrete-time Markov Chain (MC, for short) is a family of categorical stochastic variables $\{S_t\}_{t \in \mathbf{N}}$ taking their values from \mathcal{S} , that satisfies the *Markov* property, i.e., $P(S_{t+1}|S_t) = P(S_{t+1}|S_t, \dots, S_0)$, and the *stationarity* assumption, i.e. $P(S_{t+1}|S_t)$ is the same for each $t \in \mathbf{N}$. Given the Markov property and the stationarity assumption, a MC can be fully described by a single *transition matrix* $T : \mathcal{S}^2 \mapsto [0, 1]$ such that $T(s, s') := P(S_{t+1} = s' | S_t = s)$ for each $(s, s') \in \mathcal{S}^2$ and a $t \in \mathbf{N}$ whose choice is arbitrary because of stationarity.

To compute relevant inferences in MCs it is useful to introduce a *transition operator* \hat{T} and its *dual* \hat{T}^\dagger . The former maps a non-negative real function f defined over \mathcal{S} to its left scalar product, i.e.:

$$(\hat{T}f)(s) := \sum_{s' \in \mathcal{S}} T(s', s) \cdot f(s'); \quad (1)$$

for each $s \in \mathcal{S}$; while the latter maps the same function on the right scalar product, i.e.:

$$(\hat{T}^\dagger f)(s) := \sum_{s' \in \mathcal{S}} T(s, s') \cdot f(s'); \quad (2)$$

for each $s \in \mathcal{S}$. It is easy to check that $\hat{T}P(S_t) = P(S_{t+1})$ hence, by the well-known total probability theorem, $\hat{T}^t P(S_0) = P(S_t)$. Similarly, it is also easy to check that $\hat{T}^\dagger f(S_t) = E[f(S_{t+1})|S_t]$ hence, by definition of conditional expectation, $((\hat{T}^\dagger)^t f(S_0))(s) = E[f(S_t)|S_0 = s]$.

Notice that, given an *event* $B \subseteq \mathcal{S}$ and the indicator vector \mathbf{I}_B of B , the \hat{T} operator allows to efficiently compute the marginal probability of B , i.e., $P(S_t \in B) := \hat{T}^t \mathbf{I}_B$. Similarly, its dual \hat{T}^\dagger is useful for computing the *hitting probability* vector h_B , i.e. the vector that provides for each $s \in \mathcal{S}$ the probability

of reaching a given event $B \subseteq \mathcal{S}$ from s *eventually* in the future. According to standard literature [19,3], h_B is defined as the vector of the *minimal* solutions, for each $s \in \mathcal{S}$, of the following system of linear equations:

$$h_B(s) := \begin{cases} 1 & \text{if } s \in B, \\ \sum_{s' \in \mathcal{S}} h_B(s') & \text{otherwise} \end{cases} \quad (3)$$

For MCs with a finite time-horizon $t \in \mathbf{N}$ an efficient algorithm to solve the above system in time polynomial in $|\mathcal{S}|$ is the following. Let h_B^t denote the *bounded-time* hitting probability vector providing, for each $s \in \mathcal{S}$, the probability of reaching B until a number of time-steps less than or equal to t . Let \mathbf{I}_B denote the indicator function of B returning, for each $s \in \mathcal{S}$, 1 if $s \in B$ and 0 otherwise. For $t = 0$, h_B^t is given by

$$h_B^{t=0} = \mathbf{I}_B. \quad (4)$$

The probability of reaching B in 0 time-steps can correspond only to 1 or 0 depending on whether the actual state is included in the event B or not. Given \mathbf{I}_B and \mathbf{I}_{B^c} , respectively the indicator vector of B and of the complement of B , the algorithm proceeds by computing h_B^τ for increasing values of $\tau := 0, 1, \dots, t$ as follows:

$$h_B^\tau = \mathbf{I}_B + \mathbf{I}_{B^c} \cdot [\hat{T}^\dagger h_B^{\tau-1}], \quad (5)$$

where the sums and the products are intended as element-wise operations on the vector arrays.

Obviously, this procedure allows to compute h_B^t only for finite time horizons $t \in \mathbf{N}$. The standard definition of hitting probability, however, is conceived for MCs of possibly infinite time-length and refers to the probability of reaching B *eventually* in the future. Intuitively, the computation of h_B corresponds to compute $\lim_{\tau \rightarrow +\infty} h_B^\tau$. As under the *stationary* hypothesis it has been proved that $\lim_{\tau \rightarrow +\infty} h_B^\tau$ is always defined [19], it is possible to approximate the values of h_B through Equation 5 by iterating h_B^τ over increasing values of τ until convergence.

Imprecise Markov Chains. An *imprecise Markov chain* (IMC) is the imprecise counterpart of a MC. It is obtained replacing the initial PMF $P(S_0)$ with a CS $K(S_0)$ and all the conditional PMFs $\{P(S_{t+1}|s_t)\}_{s_t \in \mathcal{S}}$ with conditional CSs $\{K(S_{t+1}|s_t)\}_{s_t \in \mathcal{S}}$. The imprecise counterpart of the *stationarity* hypothesis consists of assuming the specification of the collections of CSs $K(S_{t+1}|S_t)$ independent of t . As for standard MCs, under the stationarity hypothesis, a compact specification of the CSs can be achieved in terms of an initial CS $K(S_0)$ and a collection $\{K(S'|s)\}_{s \in \mathcal{S}}$ of *transition* CSs. The collection of CSs can be seen as an *imprecise* transition matrix $\mathcal{T} := \{K(S'|s)\}_{s \in \mathcal{S}}$ whose rows consist of the transition CSs $K(S'|s)$ for each $s \in \mathcal{S}$. This matrix provides a full specification of the stochastic behaviour of the system modelled by the IMC [12]. Similarly, the linear transition operator in Equations (1) and its dual in (2) are replaced in IMCs with analogous *non-linear* operators for modelling, respectively, the *lower*

and *upper* bounds of transition probabilities. Following [23, Definition 14], in particular, the *upper* transition operator, denoted by $\overline{\mathcal{T}}$, is defined as follows:

$$(\overline{\mathcal{T}}f)(s) := \sup_{T(S',s) \in K(S'|s)} \sum_{\forall s' \in \mathcal{S}} T(s',s) \cdot f(s') \quad (6)$$

while its dual [22, Eq. 1], denoted by $\overline{\mathcal{T}}^\dagger$, is defined as follows:

$$(\overline{\mathcal{T}}^\dagger f)(s) := \sup_{T(s,S') \in K(S'|s)} \sum_{\forall s' \in \mathcal{S}} T(s,s') \cdot f(s') \quad (7)$$

The analogous *lower* operators, denoted respectively by $\underline{\mathcal{T}}$ and $\underline{\mathcal{T}}^\dagger$ can be defined by replacing the *supremum* in equations (6) and (7) with an *infimum*. Notice that, the optimization for (7) is a linear programming task whose feasible region is the convex hull generated by $K(S'|s)$ that can be described by a finite number of linear constraints, see [25].

Similarly to precise MCs, these operators can be used to compute lower and upper bounds of, respectively, *marginal* and *hitting* probabilities. In particular, as recently proved in [23, Eq. 34,35], the lower $\underline{P}(S_t \in B)$ and upper $\overline{P}(S_t \in B)$ bounds of the marginal probability $P(S_t \in B)$ for a number $t \in \mathbf{N}$ of time-steps can be computed by t application of the lower (upper) transition operator to the indicator vector \mathbf{I}_B of B , i.e.:

$$\underline{P}(S_t \in B) := \underline{\mathcal{T}}^t \mathbf{I}_B \quad (8)$$

$$\overline{P}(S_t \in B) := \overline{\mathcal{T}}^t \mathbf{I}_B \quad (9)$$

Similarly, [16, Lemma 14] and [22] proved that a recursive schema analogous to 5 can be used to compute, respectively, the *lower* \underline{h}_B^t and the *upper* \overline{h}_B^t hitting probability vectors for IMCs of a finite time length $t \in \mathbf{N}$. As in the precise case, the initialization for both $\underline{h}_B^{t=0}$ and $\overline{h}_B^{t=0}$ is given by the indicator vector of B , while the recursive steps are defined as follows:

$$\underline{h}_B^\tau = \mathbf{I}_B + \mathbf{I}_{B^c} \underline{\mathcal{T}}^\dagger \underline{h}_B^{\tau-1}, \quad (10)$$

$$\overline{h}_B^\tau = \mathbf{I}_B + \mathbf{I}_{B^c} \overline{\mathcal{T}}^\dagger \overline{h}_B^{\tau-1}; \quad (11)$$

These definitions, similarly to their analogous for precise MCs in equation (5), allow to compute the lower and upper hitting probability vectors only for IMCs of finite time length. As in the precise case, however, the generalization to IMCs of infinite time length can be obtained by computing the respective limits: $\lim_{\tau \rightarrow +\infty} \underline{h}_B^\tau$ and $\lim_{\tau \rightarrow +\infty} \overline{h}_B^\tau$. As proved in [16, Prop 16], these limits are defined. Consequently, the values of the lower \underline{h}_B and the upper \overline{h}_B hitting probability vectors for IMCs of infinite time length can be approximated by iterating both (10) and (11) for increasing values of τ until converge.

Labelled Markov Chains When dealing with model-checking tasks, we need to refer to *labelled* MCs (respectively, labelled IMCs), which are standard MCs (respectively, IMCs) augmented with a set of atomic propositions $AP := \{p, q, \dots\}$ and a labelling function $l : \mathcal{S} \mapsto 2^{AP}$ that assigns to each $s \in \mathcal{S}$ a set $l(s) \subseteq AP$. From now on, when talking about MCs (respectively, IMCs) we always refer to their labelled extensions.

2.2 Probabilistic Interpreted Systems

Precise PISs. In computational logic, *probabilistic interpreted systems* (PIS) are usually considered the reference formalism for representing knowledge and beliefs in stochastic MASs [26]. A PIS is a tuple:

$$M_{\text{PIS}} := \langle \mathcal{S}, \mathcal{A}, \{\sim^i\}_{i \in \mathcal{A}}, \{T^i\}_{i \in \mathcal{A}}, AP, l(s) \rangle; \quad (12)$$

consisting of a finite non-empty set of states \mathcal{S} , a finite non-empty set of agents $\mathcal{A} := \{i, j, \dots, n\}$, a set of atomic propositions AP , a labelling function $l : \mathcal{S} \mapsto 2^{AP}$, a transition matrix T^i for each $i \in \mathcal{A}$ describing the stochastic behaviour of the single agent i and an *epistemic equivalence relation* (EER) $\sim^i \subseteq 2^{\mathcal{S} \times \mathcal{S}}$ for each agent $i \in \mathcal{A}$ such that \sim^i associates to each $s \in \mathcal{S}$ all the states $s' \in \mathcal{S}$ that are *epistemically equivalent* (or *indistinguishable*) from s according to i . Given a state $s \in \mathcal{S}$, the set of all the states $s' \in \mathcal{S}$ such that $s \sim^i s'$ is called the *equivalence class* of s for i , denoted as $Eq^i(s)$. Given a group of agents $\Gamma \subseteq \mathcal{A}$, specific EERs for different kinds of multi-agent knowledge can be defined, including:

- *Everybody Knows*: $\sim_E^\Gamma := \bigcup_{i \in \Gamma} \sim^i$
- *Common Knowledge*: $\sim_C^\Gamma := \text{it}(\bigcup_{i \in \Gamma} \sim^i)$, where *it* denotes the *iterative closure*
- *Distributed Knowledge*: $\sim_D^\Gamma := \bigcap_{i \in \Gamma} \sim^i$.

Each EER induces a respective epistemic equivalence class (EEC) for Γ . In the following, by Eq_E^Γ , Eq_C^Γ , Eq_D^Γ , we denote the equivalence classes respectively for *Everybody Knows*, *Common Knowledge* and *Distributed Knowledge*. Finally, while each individual transition matrix T^i , $i \in \mathcal{A}$ describes the stochastic behaviour of a single agent, a global transition matrix T_{PIS} describing the stochastic behaviour of the whole MAS can be generated computing, for each $s, s' \in \mathcal{S} \times \mathcal{S}$, the *logarithmic pooling* of the transitions:

$$T_{\text{PIS}}(s, s') := \eta \prod_{i \in \mathcal{A}} T^i(s, s') \quad (13)$$

where η is a *normalizing factor* given by

$$\eta := \frac{1}{\sum_{s'' \in \mathcal{S}: s'' \neq s'} T^i(s, s'')}$$

that forces the transitions to satisfy the condition: $\sum_{s' \in \mathcal{S}} T_{\text{PIS}}(s, s') = 1$ for each $s \in \mathcal{S}$. The global transition matrix T_{PIS} also describes a specific MCs. This is typically called the *embedded* MC of the PIS and describes the overall stochastic behaviour of the whole MAS.

Imprecise Probabilistic Interpreted Systems. An imprecise probabilistic interpreted systems (IPIS) is the imprecise counterpart of a PIS. For each agent $i \in \mathcal{A}$, let $\{K^i(S'|s)\}_{s \in \mathcal{S}}$ denote a family of (credal) sets including, for each $s \in \mathcal{S}$, all the transition PMFs $P^i(S'|s)$ that are compatible with some agent's beliefs. By replacing in a standard PIS, for each $i \in \mathcal{A}$, the transition matrices T^i with the imprecise transition matrices $\mathcal{T}^i := \{K^i(S'|s)\}_{s \in \mathcal{S}}$, whose rows correspond to the transition CSs: $K^i(S'|s)$, $s \in \mathcal{S}$, we obtain an IPIS. Since transition CSs are *sets* of PMFs, admitting such sets is tantamount to admit agents' *high-order* (non-quantified) uncertainty about transition probabilities.

Similarly to PISs, in the case of IPISs a global *imprecise* transition matrix $\mathcal{T}_{\text{IPIS}}$ can be obtained computing, for each transition $s, s' \in \mathcal{S} \times \mathcal{S}$, the *credal logarithmic pooling* of the family of conditional CSs $\{K^i(S'|s) : i \in \mathcal{A}\}$ defined as the element-wise application of the standard logarithmic pooling to the elements of the credal sets. This element-wise approach, however, might comport exponential complexity with respect to the number of agents in the model. A similar problem also occurs when considering alternative strategies, such as the one proposed in [1] within the framework of general credal networks. An efficient way to overcome the problem, here adopted, consists of considering a so-called *outer approximation* of the lower and upper bounds of the credal logarithmic pooling. This is achieved by defining:

$$\underline{\mathcal{T}}_{\text{IPIS}}(s, s') := \frac{\prod_{i \in \mathcal{A}} \underline{\mathcal{T}}_{\text{IPIS}}^i(s, s')}{\prod_{i \in \mathcal{A}} \underline{\mathcal{T}}_{\text{IPIS}}^i(s, s') + \sum_{s'' \neq s'} \prod_{i \in \mathcal{A}} \underline{\mathcal{T}}_{\text{IPIS}}^i(s, s'')} ; \quad (14)$$

and similarly for the upper bound:

$$\overline{\mathcal{T}}_{\text{IPIS}}(s, s') := \frac{\prod_{i \in \mathcal{A}} \overline{\mathcal{T}}_{\text{IPIS}}^i(s, s')}{\prod_{i \in \mathcal{A}} \overline{\mathcal{T}}_{\text{IPIS}}^i(s, s') + \sum_{s'' \neq s'} \prod_{i \in \mathcal{A}} \overline{\mathcal{T}}_{\text{IPIS}}^i(s, s'')} ; \quad (15)$$

The obtained global transition matrix $\mathcal{T}_{\text{IPIS}}$ consists of an *interval-valued* transition matrix $\mathcal{T}_{\text{IPIS}}$ whose entries are intervals $(a, b) \subseteq [0, 1]$ with a and b representing, respectively, the lower $\underline{\mathcal{T}}_{\text{IPIS}}(s, s')$ and the upper $\overline{\mathcal{T}}_{\text{IPIS}}(s, s')$ bounds of the transition probabilities. Similarly to the precise case, the global matrix describes a specific IMC called the *embedded* IMC of the IPIS. As we show in the following, this can be used to compute inferences arising with the overall stochastic behaviour of the whole MAS modelled by the IPIS.

3 Imprecise Epistemic PCTL

3.1 EIPCTL Syntax

The EIPCTL syntax is defined as follows:

$$\begin{aligned}
 \mathcal{A} &:= \{i, j, \dots, n\}, \\
 \nabla &:= \{<, \leq, =, \geq, >\}, \\
 \phi &:= \top \mid p \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid P_J\psi \mid \underline{P}_{\nabla b}\psi \mid \overline{P}_{\nabla b}\psi \mid K^i\phi \mid E^\Gamma\phi \mid C^\Gamma\phi \mid D^\Gamma\phi, \\
 \psi &:= \bigcirc\phi \mid \phi_1 \bigcup \phi_2 \mid \phi_1 \bigcup_{\leq t} \phi_2, \\
 \epsilon &:= B_{\nabla b}^i\phi \mid B_{\nabla \bar{b}}^i\phi;
 \end{aligned}$$

The language is an epistemic extension of the well-known PCTL able to cope with agents' *high-order* uncertainty about transition probabilities. To this aim, the canonical PCTL probability operator is replaced with three new operators, for the following formulas with $b \in [0, 1]$, $J \subseteq [0, 1]$:

1. $\underline{P}_{\nabla b}\psi$: The lower bound of the probability of reaching a path that satisfies ψ is ∇b ;
2. $\overline{P}_{\nabla b}\psi$: The upper bound of the probability of reaching a path that satisfies ψ is ∇b ;
3. $P_J\psi$: The probability of reaching a path that satisfies ψ belongs to the interval J .

In addition, we extend standard PCTL language including a (single-agent) knowledge operator and canonical multi-agent operators for *Everybody Knows*, *Common Knowledge* and *Distributed Knowledge*. Finally, we also include in the language two weighted-belief operators:

- $B_{\nabla b}^i\phi$: The agent i believes that the lower bound of the probability to reach ϕ eventually in the future is ∇b ;
- $B_{\nabla \bar{b}}^i\phi$: The agent i believes that the upper bound of the probability to reach ϕ eventually in the future is ∇b .

In the following, the doxastic formulae including these operators are called *imprecise probabilistic beliefs*.

3.2 Semantics of EIPCTL

Semantics of Boolean formulae. Given an IPIS $\mathcal{M}_{\text{IPIS}}$ and a state $s \in \mathcal{S}$, the following conditions hold:

$$\begin{aligned}
 \mathcal{M}_{\text{IPIS}}, s &\models p \text{ iff } p \in l(s), \\
 \mathcal{M}_{\text{IPIS}}, s &\models \neg\phi \text{ iff } \mathcal{M}_{\text{IPIS}}, s \not\models \phi, \\
 \mathcal{M}_{\text{IPIS}}, s &\models \phi_1 \wedge \phi_2 \text{ iff } \mathcal{M}_{\text{IPIS}}, s \models \phi_1 \text{ and } \mathcal{M}_{\text{IPIS}}, s \models \phi_2.
 \end{aligned}$$

Semantics of ψ formulae. Given an IPIS $\mathcal{M}_{\text{IPIS}}$ and a path π , the following conditions hold:

$$\begin{aligned} \mathcal{M}_{\text{IPIS}}, \pi &\models \bigcirc \phi \text{ iff } \mathcal{M}, \pi(1) \models \phi, \\ \mathcal{M}_{\text{IPIS}}, \pi &\models \phi_1 \bigcup^{\leq t} \phi_2 \text{ iff } \exists \tau \leq t : \mathcal{M}_{\text{IPIS}}, \pi(\tau) \models \phi_2 \\ &\quad \forall \tau' : 0 \leq \tau' < \tau : \mathcal{M}_{\text{IPIS}}, \pi(\tau') \models \phi_1, \\ \mathcal{M}_{\text{IPIS}}, \pi &\models \phi_1 \bigcup \phi_2 \text{ iff } \exists \tau \geq 0 : \mathcal{M}_{\text{IPIS}}, \pi(\tau) \models \phi_2 \\ &\quad \forall \tau' : 0 \leq \tau' < \tau \mathcal{M}_{\text{IPIS}}, \pi(\tau') \models \phi_1. \end{aligned}$$

Semantics of Probabilistic Formulae. Given a model $\mathcal{M}_{\text{IPIS}}$ and a state $s \in \mathcal{S}$, let $\text{Paths}(s)$ denote the set of all the paths $\pi := (\pi(0), \pi(1), \dots)$ such that $\pi(0) = s$. We denote by $P_{\text{IPIS}}(\pi \in \text{Paths}(s) | \pi \models \psi)$ the *overall*³ probability that a path π satisfying the property ψ belongs to $\text{Paths}(s)$. Similarly, we denote by $\underline{P}_{\text{IPIS}}(\pi \in \text{Paths}(s) | \pi \models \psi)$ and $\overline{P}_{\text{IPIS}}(\pi \in \text{Paths}(s) | \pi \models \psi)$, respectively the *lower* and *upper* bounds of $P_{\text{IPIS}}(\pi \in \text{Paths}(s) | \pi \models \psi)$.

The satisfiability conditions for *probabilistic-until* formulae are hence defined as follows:

$$\begin{aligned} \mathcal{M}_{\text{IPIS}}, s &\models \underline{P}_{\nabla b} \psi \text{ iff } \underline{P}_{\text{IPIS}}(\pi \in \text{Paths}(s) | \pi \models \psi) \nabla b, \\ \mathcal{M}_{\text{IPIS}}, s &\models \overline{P}_{\nabla b} \psi \text{ iff } \overline{P}_{\text{IPIS}}(\pi \in \text{Paths}(s) | \pi \models \psi) \nabla b, \\ \mathcal{M}_{\text{IPIS}}, s &\models P_J \psi \text{ iff } \mathcal{M}_{\text{IPIS}}, s \models \underline{P}_{=(\inf J)} \psi \text{ and } \mathcal{M}_{\text{IPIS}}, s \models \overline{P}_{=(\sup J)} \psi. \end{aligned}$$

Notice that, similarly to standard PCTL [3], the computation of the lower and upper bounds of $P_{\text{IPIS}}(\pi \in \text{Paths}(s) | \pi \models \psi)$ varies depending on ψ . We analyse further this point in the next section focused on model-checking procedures.

Semantics of Epistemic Formulae. Given an IPIS $\mathcal{M}_{\text{IPIS}}$, an agent $i \in \mathcal{A}$ or a group of agents $\Gamma \subseteq \mathcal{A}$ and a state $s \in \mathcal{S}$, the following conditions hold:

$$\begin{aligned} \mathcal{M}_{\text{IPIS}}, s &\models K^i \phi \text{ iff } \forall s', s \sim^i s' : s' \models \phi, \\ \mathcal{M}_{\text{IPIS}}, s &\models E^\Gamma \phi \text{ iff } \forall s', s \sim_E^\Gamma s' : s' \models \phi, \\ \mathcal{M}_{\text{IPIS}}, s &\models C^\Gamma \phi \text{ iff } \forall s', s \sim_C^\Gamma s' : s' \models \phi, \\ \mathcal{M}_{\text{IPIS}}, s &\models D^\Gamma \phi \text{ iff } \forall s', s \sim_D^\Gamma s' : s' \models \phi. \end{aligned}$$

Semantics of Imprecise Probabilistic Beliefs. The weighted-belief operators of EIPCTL model the lower and upper bounds of the probability that a single agent $i \in \mathcal{S}$ eventually reaches a state satisfying ϕ . Following the *probabilistic until* semantics (Section 3.2), this probability can be written as $P^i(\pi \in \text{Paths}(s) | \pi \models \top \bigcup \phi)$, that is, the probability that ϕ is satisfied eventually in the future according to agent i . This probability is computed analogously to $P(\pi \in \text{Paths}(s) | \pi \models \top \bigcup \phi)$, see Section 3.2, but replacing the global imprecise transition matrix $\mathcal{T}_{\text{IPIS}}$, describing the overall stochastic behaviour of the whole MAS, with the transition matrix \mathcal{T}^i that describes the specific behaviour of the agent $i \in \mathcal{A}$. As here

³ i.e., computed through the global transition matrix $\mathcal{T}_{\text{IPIS}}$.

we consider imprecise models, as usual, we are interested in computing the lower and upper bounds of $P^i(\pi \in \text{Paths}(s) | \pi \models \top \cup \phi)$ that we denote, respectively, by $\underline{P}(\pi \in \text{Paths}(s) | \pi \models \top \cup \phi)$ and $\overline{P}(\pi \in \text{Paths}(s) | \pi \models \top \cup \phi)$. The procedure to compute those bounds is further detailed in Section 4. Here we limit to introduce the satisfiability conditions for imprecise probabilistic belief formulae as follows:

$$\begin{aligned} \mathcal{M}_{\text{IPIS}}, s \models B_{\nabla b}^i \phi &\text{ iff } \forall s' : s \sim^i s' \underline{P}^i(\pi \in \text{Paths}(s') | \pi \models \top \cup \phi) \nabla b, \\ \mathcal{M}_{\text{IPIS}}, s \models B_{\nabla \bar{b}}^i \phi &\text{ iff } \forall s' : s \sim^i s' \overline{P}^i(\pi \in \text{Paths}(s') | \pi \models \top \cup \phi) \nabla b. \end{aligned}$$

4 Model Checking

The present section describes specific procedures to model-check an IPIS against properties specified in the EIPCTL language. In particular, it aims to prove that relevant model-checking tasks can be solved using slightly modified versions of the algorithms described in 2.1. Here we consider only procedures relevant for probabilistic and epistemic formulae of EIPCTL. The checking procedures against Boolean and CTL formulae are standard and a detailed explanation of them can be found in [3].

Probabilistic Formulae. Given the semantics introduced in Section 3.2, to check whether a model satisfies a given probabilistic formula requires to compute the lower and upper bounds of $P_{\text{IPIS}}(\pi \in \text{Paths}(s) | M, \pi \models \psi)$. These, in turn, vary depending on ψ .

Probabilistic Next. We first consider the case when $\psi := \bigcirc \phi$. Let Φ denote the set of all the states satisfying ϕ . Thus, $P_{\text{IPIS}}(\pi \in \text{Paths}(s) | M, \pi \models \psi)$ corresponds to the marginal probability $P_{\text{IPIS}}(S_1 \in \Phi | S_0 = s)$. The first step for determining such probability consists in generating the indicator vector \mathbf{I}_Φ computing, for each $s \in \mathcal{S}$, the indicator function

$$\mathbf{I}_\Phi(s) := \begin{cases} 1 & \text{if } s \in \Phi, \\ 0 & \text{else.} \end{cases}.$$

This step requires a time linear in $|\mathcal{S}|$. The second step requires to introduce the lower $\underline{\mathcal{T}}_{\text{IPIS}}$ or the upper $\overline{\mathcal{T}}_{\text{IPIS}}$ transition operator. Given the global transition matrix $\mathcal{T}_{\text{IPIS}}$, these can be defined following Equation (6) and used for computing the lower and upper bounds of $P_{\text{IPIS}}(S_1 \in \Phi | S_0 = s)$ as follows:

$$\underline{P}_{\text{IPIS}}(S_1 \in \Phi | S_0 = s) = (\underline{\mathcal{T}}_{\text{IPIS}} \mathbf{I}_\Phi)(s). \quad (16)$$

$$\overline{P}_{\text{IPIS}}(S_1 \in \Phi | S_0 = s) = (\overline{\mathcal{T}}_{\text{IPIS}} \mathbf{I}_\Phi)(s). \quad (17)$$

As stated in Section 2.1, each application of either the lower or the upper transition operator requires to solve, for each $s \in \mathcal{S}$, a linear programming task whose feasible region is the convex hull obtained by convex closure of the local

CS $K(S'|s)$. By definition, the time complexity of each linear programming task is, at most, polynomial in $|\mathcal{S}|$. Hence, since the computation requires to solve at most $|\mathcal{S}|$ linear programming tasks, the overall time complexity is at most polynomial in $|\mathcal{S}|$.

Probabilistic Bounded Until. For $\psi = \phi_1 \bigcup^{\leq t} \phi_2$, let us first define Φ_1 and Φ_2 as the subsets of \mathcal{S} satisfying, respectively, ϕ_1 and ϕ_2 . The probability in Equations (16) and (17) can be seen as a bounded-time hitting probability of Φ_2 with the additional condition that all the states visited before reaching Φ_2 are in Φ_1 . We denote such *conditional* hitting probability by $h_{\Phi_2|\Phi_1}^t(s)$, $s \in \mathcal{S}$. A recursive algorithm analogous to equations (10) and (11) can be formulated to compute the values of the lower $h_{\Phi_2|\Phi_1}^t$ and upper $\bar{h}_{\Phi_2|\Phi_1}^t$ hitting probability vectors. Let \mathbf{I}_{Φ_2} denote the indicator vector of Φ_2 . Let $\mathbf{I}_{\Phi_1 \setminus \Phi_2}$ denote the indicator vector giving 1 to all the states that are in Φ_1 but not in Φ_2 and 0 otherwise. Finally, let $\mathcal{T}_{\text{IPIS}}$ denote the *global* transition matrix describing the overall behaviour of the whole MAS, and let $\mathcal{T}_{\text{IPIS}}^\dagger, \bar{\mathcal{T}}_{\text{IPIS}}^\dagger$ denote, respectively, the lower and the upper *dual* transition operators obtained as by definitions (10) and (11). A slightly modified version of the algorithms in (10), (11) for computing, for each $0 < \tau \leq t$, the above lower and upper hitting probability vectors can be achieved as follows:

$$h_{\Phi_2|\Phi_1}^\tau := \mathbf{I}_{\Phi_2} + \mathbf{I}_{\Phi_1 \setminus \Phi_2} (\mathcal{T}_{\text{IPIS}}^\dagger h_{\Phi_2|\Phi_1}^{\tau-1}), \quad (18)$$

$$\bar{h}_{\Phi_2|\Phi_1}^\tau := \mathbf{I}_{\Phi_2} + \mathbf{I}_{\Phi_1 \setminus \Phi_2} (\bar{\mathcal{T}}_{\text{IPIS}}^\dagger \bar{h}_{\Phi_2|\Phi_1}^{\tau-1}). \quad (19)$$

Notice that, exactly as in Equations (10) and (11) the initialization is given by the indicator function of Φ_2 while the recursive steps consist of iterated applications of the lower (upper) transition operator to the hitting vector computed at the precedent time-step $\tau - 1$, for each $0 < \tau \leq t$. The only relevant difference with the analogous scheme presented in Section 2.1 consists of the indicator vector $\mathbf{I}_{\Phi_1 \setminus \Phi_2}$ that replaces \mathbf{I}_{B^c} , i.e., the indicator vector of the complement of the hitting event B . In the general scheme, \mathbf{I}_{B^c} limits the iteration considering only paths that have not already visited an $s \in B$. Here, by $\mathbf{I}_{\Phi_1 \setminus \Phi_2}$ we limit the iteration considering only paths whose actual and previous states are all in Φ_1 and that have not already reached a state $s \in \Phi_2$. This constraint follows from the semantics of probabilistic until, see Section 4.

The time complexity remains polynomial with respect to $|\mathcal{S}|$. In fact, the solution of the schema (18) requires: (i) to generate the indicator vectors for Φ_2 and $\Phi_1 \setminus \Phi_2$, (ii) to execute element-wise sums and products on the vectors arrays, (iii) to execute $t - 1$ applications of the lower (respectively, upper) transition operator. (i) requires a number of time-steps linear in $|\mathcal{S}|$ and (ii) requires a number of time-steps linear in $|\mathcal{S}|$ for each recursive step $0 < \tau \leq t$. Both, hence, do not affect the overall time complexity of the procedure. Finally, remember that each application of the lower (upper) dual operator has a time complexity at most polynomial in $|\mathcal{S}|$, see Section 2.1. As the overall procedure requires a finite number $t - 1 \in \mathbf{N}$ of successively recursive applications of the respective transition operator, one for each recursive step $0 < \tau \leq t$, the overall time

complexity results being polynomial in $|\mathcal{S}|$. To conclude, notice that here we consider the CSs $K_{\text{IPIS}}(S'|s)$, $s \in \mathcal{S}$ describing the *global* transitions of the whole MAS. By the definition of *global* transition matrix $\mathcal{T}_{\text{IPIS}}$ in Section 2.2, these CSs are given by the rows of $\mathcal{T}_{\text{IPIS}}$.

Probabilistic Until. The strategy to solve model-checking tasks for probabilistic (unbounded) until formulae, i.e. formulae composed by a probabilistic operator ranging over $\psi := \phi_1 \bigcup \phi_2$, is the same as for probabilistic bounded until but considering, in place of IMCs of a finite time length $t \in \mathbf{N}$, IMCs of infinite time length. Recall that the existence of the limits $\lim_{\tau \rightarrow +\infty} \underline{h}_B^\tau$ and $\lim_{\tau \rightarrow +\infty} \bar{h}_B^\tau$ defined in section 2.1 has been proved, see [16, Lemma 14]. The slightly modification introduced in 10 and 11, consisting of replacing \mathbf{I}_B in the general schema with the indicator vector $\mathbf{I}_{\phi_1 \setminus \phi_2}$, does not affect the validity of the proof outlined in [16]. Consequently, the values of the lower and upper hitting probability vectors $\underline{h}_{\phi_2|\phi_1}$ and $\bar{h}_{\phi_2|\phi_1}$ can be approximated by iterating the schema 10, respectively, the schema 11, over increasing values of τ until convergence. Regarding the overall time complexity of the procedure, the same reasoning outlined above for probabilistic bounded until formulae holds.

Epistemic Formulae. The model-checking for epistemic formulae requires an iterative procedure. In practice, it consists of computing the *epistemic equivalence class* (EEC) relative to the specified agent (respectively, group of agents) and the actual state of the system, hence checking for each state in the EEC whether it satisfies the formula nested by the epistemic operator.

Let $\kappa := K^i, E^\Gamma, C^\Gamma, D^\Gamma, \kappa\phi$ and let $Eq_{E,C,D}^{i,\Gamma}$ be a generic notation for one of the possible EEC $Eq^i(s)$, $Eq_E^\Gamma(s)$, $Eq_C^\Gamma(s)$, $Eq_D^\Gamma(s)$. Given a state $s \in \mathcal{S}$, our task consists of defining a procedure for checking whether $s \models \kappa\phi$. Step (i) consists of deciding which states $s' \in \mathcal{S}$ belong to the EEC $Eq_{E,C,D}^{i,\Gamma}(s)$. This can be obtained in time linear in $|\mathcal{S}|$ by simply computing the *characteristic function* of $Eq_{E,C,D}^{i,\Gamma}$, i.e.:

$$\mathbf{I}_{Eq_{E,C,D}^{i,\Gamma}}(s) := \begin{cases} 1 & \text{if } s' \in Eq_{E,C,D}^{i,\Gamma}(s), \\ 0 & \text{otherwise.} \end{cases} \quad (20)$$

Step (ii) consists of selecting the appropriate model-checking procedure for ϕ , else check, for each $s' \in Eq_{E,C,D}^{i,\Gamma}(s)$, whether $s' \models \phi$. If ϕ is an epistemic formula, (i.e., $\phi := \kappa'\phi'$) we return to step (i) and repeat the same procedure for each $s' \in Eq_{E,C,D}^{i,\Gamma}$. Steps (i) and (ii) are successively iterated until all the resting nested formulae ϕ are all non-epistemic formulae. The time complexity of the whole procedure is polynomial in $|\mathcal{S}|$. In fact, (i) consists of computing an EEC, a task that can be solved in time linear in $|\mathcal{S}|$ while (ii) requires the execution of a checking procedure for each state in the EEC. As all the checking tasks relative to any kind of EIPCTL non-epistemic formulae are solvable in time polynomial in $|\mathcal{S}|$, also (ii) will be solvable in time at most polynomial in $|\mathcal{S}|$.

Imprecise Probabilistic Belief. The model-checking procedure for imprecise probabilistic beliefs formulae requires to compute $\underline{P}^i(\pi \in \text{Paths}(s)|\pi \models \top \cup \phi)$ and $\overline{P}^i(\pi \in \text{Paths}(s)|\pi \models \top \cup \phi)$. As the semantics of *imprecise probabilistic belief* formulae is analogous to the semantics of *probabilistic until* formulae, we can reduce the computation of $\underline{P}^i(\pi \in \text{Paths}(s)|\pi \models \top \cup \phi)$ and $\overline{P}^i(\pi \in \text{Paths}(s)|\pi \models \top \cup \phi)$ to the computation of, respectively, \underline{h}_ϕ^i and \overline{h}_ϕ^i . These are the lower and upper bounds of the hitting probability of Φ computed through the i 's transition matrix \mathcal{T}^i instead of the *global* transition matrix $\mathcal{T}_{\text{IPIS}}$. Let $\underline{\mathcal{T}}_i^\dagger$ and $\overline{\mathcal{T}}_i^\dagger$ denote respectively the lower and the upper dual transition operators obtained from \mathcal{T}^i as by definitions (10), (11). The model-checking procedure is analogous to (18) with the only difference that we replace the dual transition operators $\underline{\mathcal{T}}_{\text{IPIS}}^\dagger, \overline{\mathcal{T}}_{\text{IPIS}}^\dagger$ with the analogous dual transition operators $\underline{\mathcal{T}}_i^\dagger, \overline{\mathcal{T}}_i^\dagger$. Regarding the time complexity of the procedure, the same reasoning outlined above for probabilistic until formulae holds.

5 Example

We now validate the applicability of the proposed model-checking tasks on a simple example. Consider an IPIS $\mathcal{M}_{\text{IPIS}}$ made of three agents $\mathcal{A} := \{i, j, k\}$ and three states $\mathcal{S} := \{1, 2, 3\}$ labelled as: $p := \{1, 2\}$, $q := \{3\}$. The EER for each agent are defined as:

$$\sim^i: \{1, 2\}, \{3\}; \quad \sim^j: \{1, 2\}, \{3\}; \quad \sim^k: \{1\}, \{2\}, \{3\}. \quad (21)$$

while the stochastic behaviours of the single agents are described by the following *imprecise*, interval-valued, transition matrices:

$$i := \begin{bmatrix} 0 & 0.4-0.9 & 0.1-0.6 \\ 0.2-0.8 & 0 & 0.2-0.8 \\ 0.3-0.5 & 0.7-0.5 & 0 \end{bmatrix}, \quad j := \begin{bmatrix} 0 & 0.45-0.95 & 0.05-0.55 \\ 0.25-0.88 & 0 & 0.12-0.75 \\ 0.32-0.5 & 0.5-0.78 & 0 \end{bmatrix},$$

$$k := \begin{bmatrix} 0 & 0.55-0.95 & 0.05-0.45 \\ 0.15-0.95 & 0 & 0.05-0.85 \\ 0.32-0.5 & 0.5-0.78 & 0 \end{bmatrix}.$$

The task consists in checking whether the following formula holds for $s = 1$:

$$\mathcal{M}, s \models K^i P_{0.99-1} \top \bigcup^{\leq 150} (3); \quad (22)$$

First, we derive the global transition matrix $\mathcal{T}_{\text{IPIS}}$ from the subjective transition matrices stated above. Applying Equations (14) and (15) to each row of the matrix, and rounding to the fourth digit, we obtain:

$$\begin{bmatrix} 0 & 0.4-0.9997 & 0.0003-0.6 \\ 0.0145-0.9982 & 0 & 0.0018-0.9855 \\ 0.0673-0.5 & 0.5-0.9327 & 0 \end{bmatrix}.$$

Second, we compute for each $s' \in \mathcal{S} : 1 \sim^i s'$ whether:

$$\mathcal{M}_{\text{IPIS}}, s' \models \underline{P}_{\geq 0.99} \top \bigcup^{\leq 150} (3).$$




Figure 1. Hitting probability ranges w.r.t. increasing time horizon $0 < \tau \leq 150$ for starting state 1 (left) and 2 (right). The horizontal dotted line denote the threshold level $\underline{h}_{(3)} = 0.99$, while the vertical dotted line show the first time step for which such threshold is exceeded.

Since the only states equivalent to 1 for i are 1 and 2, we check the above formula with respect to both these two states. Following Equation (16), this corresponds to calculate the lower $\underline{h}_{(3)}^{150}$ and upper $\bar{h}_{(3)}^{150}$ bounded-time hitting probability vectors relative to event $\{3\}$ by means of the recursive scheme in Equations (10) and (11) for a time horizon $t = 150$. These lower and upper bounds are reported in Figure 1. Finally, we extract from the vector the respective values for states 1 and 2 and check for both if the probability conditions specified in the formula are satisfied. As showed in Figure 1, both for states 1 and 2 the lower bound reached value 0.99 as well as the upper bounds converges to 1 before $t = 150$. The above formula is hence satisfied for $s = 1$.

6 Conclusions

In this paper we define a logic and relative model-checking procedures to model-check stochastic MASs characterized by agents' high-order (not quantified) uncertainty about transition probabilities. Here we limit to consider agents whose stochastic behaviour can be described by discrete-time models. Recent developments [15] in the study of imprecise *continuous-time* Markov chains (ICTMC) strongly suggest that an analogous extension for agents whose behaviour is described by continuous-time models is possible. These models are particularly relevant for applications in fields like computational and systems biology, see [8,5]. Other important extensions we would consider are multi-agent imprecise Markov decision-processes because they offer a natural connection with the field of Reinforcement Learning [20]. Finally, we aim to consider possible connections with other logical formalism developed within the theory of imprecise probabilities, such as depth-bounded belief functions [4].

References

1. Antonucci, A., Huber, D., Zaffalon, M., Luginbühl, P., Chapman, I., Ladouceur, R.: Credo: A military decision-support system based on credal networks. In: Proceedings of the 16th Conference on Information Fusion (FUSION 2013). Istanbul, Turkey (2013)
2. Baier, C., Hensel, C., Hutschenreiter, L., Junges, S., Katoen, J., Klein, J.: Parametric markov chains: PCTL complexity and fraction-free gaussian elimination. *Information and Computation* **272**, 104504 (2020). <https://doi.org/10.1016/j.ic.2019.104504>, <https://doi.org/10.1016/j.ic.2019.104504>
3. Baier, C., Katoen, J.: Principles of Model Checking. MIT Press (2008)
4. Baldi, P., Hosni, H.: Depth-bounded belief functions. *International Journal of Approximate Reasoning* **123**, 26–40 (2020). <https://doi.org/10.1016/j.ijar.2020.05.001>, <https://doi.org/10.1016/j.ijar.2020.05.001>
5. Benes, N., Brim, L., Pastva, S., Safránek, D.: Model checking approach to the analysis of biological systems. In: Liò, P., Zuliani, P. (eds.) *Automated Reasoning for Systems Biology and Medicine, Computational Biology*, vol. 30, pp. 3–35. Springer (2019)
6. Bentahar, J., Moulin, B., Meyer, J.C.: A new model checking approach for verifying agent communication protocols. In: Proceedings of the Canadian Conference on Electrical and Computer Engineering, CCECE 2006, May 7–10, 2006, Ottawa Congress Centre, Ottawa, Canada. pp. 1586–1590. IEEE (2006). <https://doi.org/10.1109/CCECE.2006.277640>, <https://doi.org/10.1109/CCECE.2006.277640>
7. Bérard, B., Bidoit, M., Finkel, A., Laroussinie, F., Petit, A., Petrucci, L., Schnoebelen, P., McKenzie, P.: *Systems and Software Verification, Model-Checking Techniques and Tools*. Springer (2001)
8. Brim, L., Ceska, M., Safránek, D.: Model checking of biological systems. In: Bernardo, M., de Vink, E.P., Pierro, A.D., Wiklicky, H. (eds.) *Formal Methods for Dynamical Systems - 13th International School on Formal Methods for the Design of Computer, Communication, and Software Systems, SFM 2013*, Bertinoro, Italy, June 17–22, 2013. *Advanced Lectures. Lecture Notes in Computer Science*, vol. 7938, pp. 63–112. Springer (2013)
9. Chen, T., Primiero, G., Raimondi, F., Rungta, N.: A computationally grounded, weighted doxastic logic. *Studia Logica* **104**(4), 679–703 (2016)
10. Daws, C.: Symbolic and parametric model checking of discrete-time markov chains. In: Zhiming, L., Araki, K. (eds.) *Theoretical Aspects of Computing - ICTAC 2004, First International Colloquium*, Guiyang, China, September 20–24, 2004, Revised Selected Papers. *Lecture Notes in Computer Science*, vol. 3407, pp. 280–294. Springer (2004)
11. De Cooman, G., De Bock, J., Lopatzidis, S.: Imprecise stochastic processes in discrete time: global models, imprecise markov chains and ergodic theorems. *Int. J. Approx. Reason.* **76**, 18–46 (2016)
12. De Cooman, G., Hermans, F., Quaeghebeur, E.: Imprecise markov chains and their limit behavior. *Probability in the Engineering and Informational Sciences* **23**(4), 597–635 (2009)
13. Delgado, K.V., De Barros, L.N., Dias, D.B., Sanner, S.: Real-time dynamic programming for markov decision processes with imprecise probabilities. *Artificial Intelligence* **230**, 192–223 (2016)

14. Hansson, H., Jonsson, B.: A logic for reasoning about time and reliability. *Formal Aspects of Computing* **6**(5), 512–535 (1994)
15. Krak, T., De Bock, J., Siebes, A.: Imprecise continuous-time markov chains. *International Journal of Approximate Reasoning* **88**, 452–528 (2017)
16. Krak, T.E., T’Joens, N., Bock, J.D.: Hitting times and probabilities for imprecise markov chains. In: Bock, J.D., de Campos, C.P., de Cooman, G., Quaeghebeur, E., Wheeler, G.R. (eds.) *International Symposium on Imprecise Probabilities: Theories and Applications, ISIPTA 2019, 3-6 July 2019, Thagaste, Ghent, Belgium. Proceedings of Machine Learning Research*, vol. 103, pp. 265–275. PMLR (2019), <http://proceedings.mlr.press/v103/krak19a.html>
17. Mauá, D.D., de Campos, C.P., Benavoli, A., Antonucci, A.: Probabilistic inference in credal networks: New complexity results. *Journal of Artificial Intelligence Research* **50**, 603–637 (2014)
18. Parker, D., Norman, G., Kwiatkowska, M.: Prism: Probabilistic symbolic model checker, 4.6 version (2020), <https://www.prismmodelchecker.org>
19. Revuz, D.: Markov chains. Elsevier (2008)
20. Sutton, R.S., Barto, A.G.: Reinforcement learning - an introduction. Adaptive computation and machine learning, MIT Press (1998)
21. Termine, A., Antonucci, A., Facchini, A., Primiero, G.: Robust model checking with imprecise markov reward models. In: *International Symposium on Imprecise Probabilities: Theories and Applications, ISIPTA 2021, July 6-9 2021, Granada, Spain. Proceedings of Machine Learning Research*, PMLR (2021)
22. T’Joens, N., Krak, T., De Bock, J., De Cooman, G.: A recursive algorithm for computing inferences in imprecise markov chains. In: Kern-Isberner, G., Ognjanovic, Z. (eds.) *Symbolic and Quantitative Approaches to Reasoning with Uncertainty, 15th European Conference, ECSQARU 2019, Belgrade, Serbia, September 18-20, 2019, Proceedings. Lecture Notes in Computer Science*, vol. 11726, pp. 455–465. Springer (2019)
23. Troffaes, M.C.M., Skulj, D.: Model checking for imprecise markov chains. In: Cozman, F., Denoeux, T., Destercke, S., Seidenfeld, T. (eds.) *ISIPTA ’13 : proceedings of the eighth international symposium on imprecise probability : theories and applications July 2-5 2013, Compiègne, France.*, pp. 337–344. Society for Imprecise Probability: Theories and Applications (SIPTA) (July 2013)
24. van Ditmarsch, H., van der Hoek, W., Halpern, J., Kooi, B. (eds.): *Handbook of Epistemic Logic*. College Publications (2015)
25. Walley, P.: *Statistical reasoning with imprecise probabilities*. Chapman and Hall (1991)
26. Wan, W., Bentahar, J., Hamza, A.B.: Model checking epistemic-probabilistic logic using probabilistic interpreted systems. *Knowledge Based Systems* **50**, 279–295 (2013). <https://doi.org/10.1016/j.knosys.2013.06.017>, <https://doi.org/10.1016/j.knosys.2013.06.017>