

Model Checking Interactive Markov Chains^{*}

Lijun Zhang¹ and Martin R. Neuhäuser^{2,3}

¹ Oxford University Computing Laboratory, UK

² Software Modeling and Verification Group, RWTH Aachen University, Germany

³ Formal Methods and Tools Group, University of Twente, The Netherlands

Abstract. Hermanns has introduced interactive Markov chains (IMCs) which arise as an orthogonal extension of labelled transition systems and continuous-time Markov chains (CTMCs). IMCs enjoy nice compositional aggregation properties which help to minimize the state space incrementally. However, the model checking problem for IMCs remains unsolved apart from those instances, where the IMC can be converted into a CTMC. This paper tackles this problem: We interpret the continuous stochastic logic (CSL) over IMCs and define the semantics of probabilistic CSL formulas with respect to the class of fully time and history dependent schedulers. Our main contribution is an efficient model checking algorithm for verifying CSL formulas on IMCs. Moreover, we show the applicability of our approach and provide some experimental results.

1 Introduction

The success of Markovian models for quantitative performance and dependability evaluation is based on the availability of efficient and quantifiably precise solution methods for continuous-time Markov chains (CTMCs) [3]. On the specification side, the continuous stochastic logic (CSL) [2,3] allows to specify a wide variety of performance and dependability measures of interest. A CTMC can be conceived as a labelled transition system (LTS) whose transitions are delayed according to an exponential distribution. Opposed to classical concurrency theory, CTMCs neither support compositional modelling [19] nor do they allow nondeterminism in the model. Several efforts have been undertaken to overcome this limitation, including formalism like the stochastic Petri box calculus [22], statecharts [7] and process algebras [20,17].

Interactive Markov chains (IMCs) [18] conservatively extend process algebras with exponentially distributed delays and comprise most of the other approaches' benefits [10]: As they strictly separate *interactive* from *Markovian* transitions, IMCs extend LTSs with exponential delays in a fully orthogonal way. This enables compositional modelling with intermittent weak bisimulation minimization [17] and allows to augment existing untimed process algebra specifications with random timing [7]. Moreover, the IMC formalism is not restricted to exponential delays but allows to encode arbitrary phase-type distributions such as hyper- and hypoexponentials [26].

Since IMCs smoothly extend classical LTSs, the model has received attention in academic as well as in industrial settings [8,14,15]. In practice however, the theoretical

^{*} Supported by the NWO projects QUPES (612.000.420), by the EU grant FP7-ICT-2007-1 (QUASIMODO) and the DFG as part of SFB/TR 14 AVACS.

benefits have partly been foiled by the fact that the analysis of IMCs is restricted to those instances, where the composed IMC could be transformed into a CTMC. However, IMCs support nondeterminism which arises both implicitly from parallel composition and explicitly by the deliberate use of underspecification in the model [18]. Therefore IMCs are strictly more expressive than CTMCs. As a result, model checking IMCs is an unexplored topic thus far.

In this paper, we overcome this limitation and propose an efficient model checking algorithm to verify CSL formulas on arbitrary IMCs. In our analysis, we use fully time and history dependent schedulers to resolve all of the IMC's nondeterministic choices.

The crucial point in model checking CSL is to compute the maximum (and minimum) probability to visit a set of goal states in some time interval I . We characterize this probability as the least fixed point of a higher-order operator which involves integration over the time domain. Then we use *interactive probabilistic chains* (IPCs) [15] to define a discretization which reduces the time interval bounded reachability problem in IMCs to the problem of computing step-interval bounded reachability probabilities in IPCs. More precisely, we approximate the quantitative behaviour of the IMC up to an a priori specified error bound $\varepsilon > 0$ by its induced IPC and prove that its maximum step-interval bounded reachability coincides (up to ε) with the achievable time-interval bounded reachability probability in the underlying IMC. The resulting IPC is then subject to a modified value iteration algorithm [5], which maximizes the step-interval bounded reachability probability. The time complexity of our approach is in $\mathcal{O}(|\Phi| \cdot (n^{2.376} + (m + n^2) \cdot (\lambda b)^2 / \varepsilon))$, where $|\Phi|$ is the size of the formula, and n, m are the number of states and transitions of the IMC, respectively. Further, $b = \sup I$ is the upper time interval bound and λ is the maximal exit rate in the IMC.

Although we present all results only for maximum time-bounded reachability probabilities, all proofs can easily be adapted to the dual problem of determining the minimum time-bounded reachability probability.

Most of the technical details have been omitted from the paper. However, all proofs and the technicalities that are necessary to establish the error bounds that are stated within the paper can be found in [23, Chapter 6].

Organisation of the paper. The paper proceeds by first giving necessary definitions and background in Section 2. Section 3 presents algorithms for computing the time-interval bounded reachability for IMCs. Section 4 focuses on model checking algorithms for CSL, followed by experimental results in Sec. 5. Section 6 discusses related work and concludes the paper.

2 Preliminaries

Let \mathcal{X} be a finite set. Probability distributions over \mathcal{X} are functions $\mu : \mathcal{X} \rightarrow [0, 1]$ with $\sum_{x \in \mathcal{X}} \mu(x) = 1$. If $\mu(x) = 1$ for some $x \in \mathcal{X}$, μ is *degenerate*, denoted $\mu = \{x \mapsto 1\}$; in this case, we identify μ and x . The set of all probability distributions over \mathcal{X} is denoted $\text{Distr}(\mathcal{X})$. Accordingly, $\mu(X) = \sum_{x \in X} \mu(x)$ for all $X \subseteq \mathcal{X}$.

2.1 Interactive Markov Chains

We recall the definition of interactive Markov chains (IMCs) given in [17]:

Definition 1 (Interactive Markov chain). An interactive Markov chain is a tuple $\mathcal{M} = (\mathcal{S}, Act, IT, MT, \nu)$ where \mathcal{S} and Act are nonempty sets of states and actions, $IT \subseteq \mathcal{S} \times Act \times \mathcal{S}$ is a set of interactive transitions and $MT \subseteq \mathcal{S} \times \mathbb{R}_{>0} \times \mathcal{S}$ is a set of Markovian transitions. Further, $\nu \in Distr(\mathcal{S})$ is the initial distribution.

We distinguish *external* actions in Act_e from *internal* actions in Act_i and set $Act = Act_e \cup Act_i$. Several IMCs may be composed via synchronisation over the set Act_e of external actions, yielding again an IMC. For details, we refer to [17]. In this paper, we consider *closed* IMCs [21], that is, we focus on the IMC \mathcal{M} that is obtained after composition. Accordingly, \mathcal{M} is not subject to any further synchronisation and all remaining external actions can safely be hidden. Therefore, we assume that $Act_e = \emptyset$ and identify the sets Act and Act_i .

For Markovian transitions, $\lambda, \mu \in \mathbb{R}_{>0}$ denote rates of exponential distributions. $IT(s) = \{(s, \alpha, s') \in IT\}$ is the set of interactive transitions that leave state s ; similarly, for Markovian transitions we set $MT(s) = \{(s, \lambda, s') \in MT\}$. A state $s \in \mathcal{S}$ is *Markovian* iff $MT(s) \neq \emptyset$ and $IT(s) = \emptyset$; it is *interactive* iff $MT(s) = \emptyset$ and $IT(s) \neq \emptyset$. Further, s is a *hybrid state* iff $MT(s) \neq \emptyset$ and $IT(s) \neq \emptyset$; finally, s is a *deadlock state* iff $MT(s) = IT(s) = \emptyset$. $MS \subseteq \mathcal{S}$ and $IS \subseteq \mathcal{S}$ denote the sets of Markovian and interactive states in \mathcal{M} . We define $post^M(s) = \{s' \in \mathcal{S} \mid \mathbf{R}(s, s') > 0\}$.




Fig. 1. Example IMC

Example 1. Let \mathcal{M} be the IMC depicted in Fig. 1. Then s_0 is a Markovian state with a transition $(s_0, 0.3, s_2) \in MT(s)$ (depicted by a solid line) to state s_2 with rate $\lambda = 0.3$. The transition's delay is exponentially distributed with rate λ ; hence, it executes in the next $z \in \mathbb{R}_{\geq 0}$ time units with probability $\int_0^z \lambda e^{-\lambda t} dt = (1 - e^{-0.3z})$. As state s_0 has two Markovian transitions, they compete for execution and the IMC moves along the transition whose delay expires first. Clearly, in such a *race*, the *sojourn time* in s_0 is determined by the first transition that executes. As the minimum of exponential distributions is exponentially distributed with the sum of their rates, the sojourn time in a state s is determined by the *exit rate* $E(s) = \sum_{s' \in \mathcal{S}} \mathbf{R}(s, s')$ of state s , where $\mathbf{R}(s, s') = \sum \{\lambda \mid (s, \lambda, s') \in MT(s)\}$. In general, the probability to move from a state $s \in MS$ to a successor state $s' \in \mathcal{S}$ equals the probability that (one of) the Markovian transitions that lead from s to s' wins the race. Therefore, the *discrete branching probability* to move to s' is given by $\mathbf{P}(s, s') = \frac{\mathbf{R}(s, s')}{E(s)}$. Accordingly, for state s_0 of our example, we have $\mathbf{R}(s_0, s_2) = 0.3$, $E(s_0) = 0.3 + 0.6 = 0.9$ and $\mathbf{P}(s_0, s_2) = \frac{1}{3}$.

For interactive transitions, we adopt the *maximal progress assumption* [17, p. 71] which states that internal transitions (i.e. interactive transitions labelled with internal actions) trigger instantaneously. This implies that they take precedence over all Markovian transitions whose probability to execute immediately is 0. Therefore all Markovian transitions that emanate a hybrid state can be removed without altering the IMC's semantics. We do so and assume that $MT(s) \cap IT(s) = \emptyset$ for all $s \in \mathcal{S}$.

To ease the development of the theory, we assume w.l.o.g. that each internal action $\alpha \in Act_i$ has a unique successor state, denoted $succ(\alpha)$; note that this is no restriction, for if $(s, \alpha, u), (s, \alpha, v) \in IT(s)$ are internal transitions with $u \neq v$, we may replace them by new transitions (s, α_u, u) and (s, α_v, v) with fresh internal actions α_u and α_v .

We assume that entering a deadlock state results in a time lock. Therefore, we equip deadlock states $s \in \mathcal{S}$ with internal self-loop (s, α, s) . However, our approach also allows for a different deadlock state semantics, where time continues; in this case, we would add a Markovian instead of an internal self-loop. The *internal successor relation* $\rightsquigarrow_i \subseteq \mathcal{S} \times \mathcal{S}$ is given by $s \rightsquigarrow_i s'$ iff $(s, \alpha, s') \in IT$; further, the *internal reachability relation* \rightsquigarrow_i^* is the reflexive and transitive closure of \rightsquigarrow_i . Accordingly, we define $post^i(s) = \{s' \in \mathcal{S} \mid s \rightsquigarrow_i s'\}$ and $Reach^i(s) = \{s' \in \mathcal{S} \mid s \rightsquigarrow_i^* s'\}$.

2.2 Paths and Events in IMCs

We use a special action $\perp \notin Act$ and let σ range over $Act_\perp = Act \cup \{\perp\}$. A finite *path* is a sequence $\pi = s_0 \xrightarrow{t_0, \sigma_0} s_1 \xrightarrow{t_1, \sigma_1} \dots \xrightarrow{t_{n-1}, \sigma_{n-1}} s_n$ where $s_i \in \mathcal{S}$, $t_i \in \mathbb{R}_{\geq 0}$ and $\sigma_i \in Act_\perp$ for $i \leq n$; n is the length of π , denoted $|\pi|$. We use $\pi[k] = s_k$ and $\delta(\pi, k) = t_k$ to refer to the $(k+1)$ -th state on π and its associated sojourn time. Accordingly, $\Delta(\pi, i) = \sum_{k=0}^i t_k$ is the total time spent on π until (including) state $\pi[i]$. If π is finite with $|\pi| = n$, then $\Delta(\pi) = \Delta(\pi, n-1)$ is the total time spent on π ; similarly, $\pi \downarrow = s_n$ is the last state on π .

Internal transitions occur immediately. Thus an IMC can traverse several states at one point in time. We use $\pi @ t \in (\mathcal{S}^* \cup \mathcal{S}^\omega)$ for the sequence of states traversed on π at time $t \in \mathbb{R}_{\geq 0}$: Formally, let i be the smallest index s.t. $t \leq \Delta(\pi, i)$; if no such i exists, we set $\pi @ t = \langle \rangle$. Otherwise, if $t < \Delta(\pi, i)$ we define $\pi @ t = \langle s_i \rangle$; if $t = \Delta(\pi, i)$, let j be the largest index (or $+\infty$, if no such finite index exists) such that $t = \Delta(\pi, j)$. Then $\pi @ t = \langle s_i \dots s_j \rangle$. We write $s \in \langle s_i \dots s_j \rangle$ if $s \in \{s_i, \dots, s_j\}$; further, if $s \in \langle s_i \dots s_j \rangle$ we define $Pref(\langle s_i \dots s_j \rangle, s) = \langle s_i, \dots, s_k \rangle$, where $s = s_k$ and k minimal. If $s \notin \langle s_i \dots s_j \rangle$, we set $Pref(\langle s_i \dots s_j \rangle, s) = \langle \rangle$. The definitions for *time-abstract paths* are similar.

A path π (time-abstract path π') is a concatenation of a state and a sequence of *combined transitions* (*time-abstract combined transitions*) from the set $\Omega = \mathbb{R}_{\geq 0} \times Act_\perp \times \mathcal{S}$ ($\Omega_{abs} = Act_\perp \times \mathcal{S}$); hence, $\pi = s_0 \circ m_0 \circ m_1 \circ \dots \circ m_{n-1}$ with $m_i = (t_i, \sigma_i, s_{i+1}) \in \Omega$ ($m_i = (\sigma_i, s_{i+1}) \in \Omega_{abs}$). Thus $Paths^n(\mathcal{M}) = \mathcal{S} \times \Omega^n$ is the set of paths of length n in \mathcal{M} ; further, $Paths^*(\mathcal{M})$, $Paths^\omega(\mathcal{M})$ and $Paths(\mathcal{M})$ are the sets of finite, infinite and all paths in \mathcal{M} . To refer to time-abstract paths, we add the subscript *abs*; further the reference to \mathcal{M} is omitted wherever possible.

The measure-theoretic concepts are mentioned only briefly; we refer to [21] for an in-depth discussion. Events in \mathcal{M} are measurable sets of paths; as paths are Cartesian products of combined transitions, we define the σ -field $\mathfrak{F} = \sigma(\mathfrak{B}(\mathbb{R}_{\geq 0}) \times \mathfrak{F}_{Act_\perp} \times \mathfrak{F}_{\mathcal{S}})$ on subsets of Ω where $\mathfrak{F}_{\mathcal{S}} = 2^{\mathcal{S}}$ and $\mathfrak{F}_{Act_\perp} = 2^{Act_\perp}$. Then we derive the product σ -field $\mathfrak{F}_{Paths^n} = \sigma(\{S_0 \times M_0 \times \dots \times M_{n-1} \mid S_0 \in \mathfrak{F}_{\mathcal{S}}, M_i \in \mathfrak{F}\})$ of measurable subsets of $Paths^n$. The cylinder-set construction [1] extends this to infinite paths in the usual way.

2.3 Resolving Nondeterminism by Schedulers

An IMC \mathcal{M} is *nondeterministic* iff there exists $(s, \alpha, u), (s, \beta, v) \in IT(s)$ with $u \neq v$: If both internal transitions (to states s_1 and s_4) in state s_2 of Fig. 1 execute instantaneously,

the successor state is not uniquely determined. To resolve this nondeterminism, we use *schedulers*: If \mathcal{M} reaches state s_2 along a *history* $\pi \in Paths^*$, a scheduler yields a probability distribution over the set $Act_i(\pi \downarrow) = \{\alpha, \beta\}$ of *enabled actions* in s_2 .

Definition 2 (Generic measurable scheduler). A generic scheduler on an IMC $\mathcal{M} = (\mathcal{S}, Act, IT, MT, \nu)$ is a partial mapping $D : Paths^* \times \mathfrak{F}_{Act_i} \rightarrow [0, 1]$ with $D(\pi, \cdot) \in Distr(Act_i(\pi \downarrow))$ for all $\pi \in Paths^*$ with $\pi \downarrow \in IS$. A generic scheduler D is measurable (*GM scheduler*) iff for all $A \in \mathfrak{F}_{Act}$, $D^{-1}(A) : Paths^* \rightarrow [0, 1]$ is measurable.

Measurability states that $\{\pi \mid D(\pi, A) \in B\} \in \mathfrak{F}_{Paths^*}$ holds for all $A \in \mathfrak{F}_{Act}$ and $B \in \mathfrak{B}([0, 1])$; intuitively, it excludes schedulers which resolve the nondeterminism in a way that induces non-measurable sets. Recall that no nondeterminism occurs if $\pi \downarrow \in MS$. However, we slightly abuse notation and assume that $D(\pi, \cdot) = \{\perp \mapsto 1\}$ if $\pi \downarrow \in MS$ so that D yields a distribution over Act_\perp . A *GM scheduler* D is *deterministic* iff $D(\pi, \cdot)$ is degenerate for all $\pi \in Paths^*$. We use *GM* (and *GMD*) to denote the class of generic measurable (deterministic) schedulers. Further, a *GM scheduler* D_{abs} is *time-abstract* (GM_{abs}) iff $abs(\pi) = abs(\pi')$ implies $D_{abs}(\pi, \cdot) = D_{abs}(\pi', \cdot)$.

Example 2. If state s_2 in Fig. 1 is reached along path $\pi = s_0 \xrightarrow{0.4, \perp} s_2$, then $D(\pi)$ might yield the distribution $\{\alpha \mapsto \frac{1}{2}, \beta \mapsto \frac{1}{2}\}$, whereas for history $\pi' = s_0 \xrightarrow{1.5, \perp} s_2$, it might return a different distribution, say $D(\pi) = \{\alpha \mapsto 1\}$.

2.4 Probability Measures for IMCs

In this section, we define the probability measure [21] induced by D on the measurable space $(Paths^\omega, \mathfrak{F}_{Paths^\omega})$. We first derive the probability of measurable sets of combined transitions, i.e. of subsets of Ω :

Definition 3. Let $\mathcal{M} = (\mathcal{S}, Act, IT, MT, \nu)$ be an IMC and $D \in GM$. For all $\pi \in Paths^*$, we define the probability measure $\mu_D(\pi, \cdot) : \mathfrak{F} \rightarrow [0, 1]$ by:

$$\mu_D(\pi, M) = \begin{cases} \sum_{\alpha \in Act_i(\pi \downarrow)} \mathbf{1}_M(\alpha, 0, succ(\alpha)) \cdot D(\pi, \{\alpha\}) & \text{if } s \in IS \\ \int_{\mathbb{R}_{\geq 0}} E(s) e^{-E(s)t} \cdot \sum_{s' \in \mathcal{S}} \mathbf{1}_M(\perp, t, s') \cdot \mathbf{P}(s, s') dt & \text{if } s \in MS. \end{cases} \quad (1)$$

Here, $\mathbf{1}_M$ denotes an indicator, i.e. $\mathbf{1}_M(\sigma, t, s') = 1$ if $(\sigma, t, s') \in M$ and 0, otherwise. Intuitively, $\mu_D(\pi, M)$ is the probability to continue along one of the combined transition in the set M . For an interactive state $s \in IS$, it is the probability of choosing $\alpha \in Act_i(\pi \downarrow)$ such that $(\alpha, 0, succ(\alpha))$ is a transition in M ; if $s \in MS$, $\mu_D(\pi, M)$ is given by the density for the Markovian transition to trigger at time t and the probability that a successor state is chosen respecting M . As paths are inductively defined using combined transitions, we can lift the probability measure $\mu_D(\pi, \cdot)$ to \mathfrak{F}_{Paths^n} :

Definition 4 (Probability measure). Let $\mathcal{M} = (\mathcal{S}, Act, IT, MT, \nu)$ be an IMC and $D \in GM$. For $n \geq 0$, we define the probability measures $Pr_{\nu, D}^n$ inductively on the measurable space $(Paths^n, \mathfrak{F}_{Paths^n})$:

$$Pr_{\nu, D}^0 : \mathfrak{F}_{Paths^0} \rightarrow [0, 1] : \Pi \mapsto \sum_{s \in \Pi} \nu(s) \quad \text{and for } n > 0$$

$$Pr_{\nu, D}^n : \mathfrak{F}_{Paths^n} \rightarrow [0, 1] : \Pi \mapsto \int_{Paths^{n-1}} Pr_{\nu, D}^{n-1}(d\pi) \int_{\Omega} \mathbf{1}_\Pi(\pi \circ m) \mu_D(\pi, dm).$$

Observe that $Pr_{\nu,D}^n$ measures a set of paths Π of length n by multiplying the probabilities $Pr_{\nu,D}^{n-1}(d\pi)$ of path prefixes π (of length $n-1$) with the probability $\mu_D(\pi, dm)$ of a combined transition $m \in M$ which extends π to a path in Π . Together, the measures $Pr_{\nu,D}^n$ extend to a unique measure on $\mathfrak{F}_{Paths^\omega}$: if $B \in \mathfrak{F}_{Paths^n}$ is a measurable base and $C = Cyl(B)$, we define $Pr_{\nu,D}^\omega(C) = Pr_{\nu,D}^n(B)$. Due to the inductive definition of $Pr_{\nu,D}^n$, the Ionescu–Tulcea extension theorem [1] applies, which yields a unique extension of $Pr_{\nu,D}^\omega$ to arbitrary sets in $\mathfrak{F}_{Paths^\omega}$.

2.5 Interactive Probabilistic Chains

Interactive probabilistic chains (IPCs) [15] are the discrete-time analogon of IMCs:

Definition 5 (Interactive probabilistic chain). *An interactive probabilistic chain (IPC) is a tuple $\mathcal{P} = (\mathcal{S}, Act, IT, PT, \nu)$, where \mathcal{S} , Act , IT and ν are as in Def. 1 and $PT : \mathcal{S} \times \mathcal{S} \rightarrow [0, 1]$ is a transition probability function s.t. $\forall s \in \mathcal{S}. PT(s, \mathcal{S}) \in \{0, 1\}$.*

A state s in an IPC \mathcal{P} is *probabilistic* iff $\sum_{s' \in \mathcal{S}} PT(s, s') = 1$ and $IT(s) = \emptyset$; PS denotes the set of all probabilistic states. The sets of interactive, hybrid and deadlock states are defined as for IMCs, with the same assumption imposed on deadlock states. Further, we assume any IPC to be closed, that is $(s, \alpha, s') \in IT$ implies $\alpha \in Act_i$. As for IMCs, we adopt the *maximal progress assumption* [17, p. 71]; hence, internal transitions take precedence over probabilistic transitions.

Definition 6 (IPC scheduler). *Let $\mathcal{P} = (\mathcal{S}, Act, IT, PT, \nu)$ be an IPC. A function $D : Paths_{abs}^* \rightarrow Distr(Act_i)$ with $D(\pi) \in Distr(Act_i(\pi \downarrow))$ is a time abstract history dependent randomized (GM_{abs}) scheduler.*

Note that in the discrete-time setting, measurability issues do not arise. To define a probability measure on sets of paths in \mathcal{P} , we define the probability of a single transition:

Definition 7 (Combined transitions in IPCs). *Let $\mathcal{P} = (\mathcal{S}, Act, IT, PT, \nu)$ be an IPC, $s \in \mathcal{S}$, $\sigma \in Act_\perp$, $\pi \in Paths_{abs}^*$ and $(\sigma, s) \in \Omega_{abs}$ a time abstract combined transition. For scheduler $D \in GM_{abs}$, we define*

$$\mu_D^{abs}(\pi, \{(\sigma, s)\}) = \begin{cases} \mathbf{P}(\pi \downarrow, s) & \text{if } \pi \downarrow \in PS \wedge \sigma = \perp \\ D(\pi, \{\sigma\}) & \text{if } \pi \downarrow \in IS \wedge succ(\sigma) = s \\ 0 & \text{otherwise.} \end{cases}$$

is the probability of the combined transition (σ, s) . For a set of combined transitions $M \subseteq \Omega_{abs}$, we set $\mu_D^{abs}(\pi, M) = \sum_{(\sigma,s) \in M} \mu_D^{abs}(s, \{(\sigma, s)\})$.

The measures μ_D^{abs} extend to a unique measure on sets of paths in \mathcal{P} in the same way as it was shown for the IMC case in Sec. 2.4.

3 Interval Bounded Reachability Probability

We discuss how to compute the maximum probability to visit a given set of *goal states* during a given time interval. Therefore, let \mathcal{I} be the set of nonempty intervals over the

nonnegative reals and let \mathcal{Q} be the set of nonempty intervals with nonnegative rational bounds. For $t \in \mathbb{R}_{\geq 0}$ and $I \in \mathcal{I}$, we define $I \ominus t = \{x - t \mid x \in I \wedge x \geq t\}$ and $I \oplus t = \{x + t \mid x \in I\}$. Obviously, if $I \in \mathcal{Q}$ and $t \in \mathbb{Q}_{\geq 0}$, this implies $I \ominus t \in \mathcal{Q}$ and $I \oplus t \in \mathcal{Q}$.

3.1 A Fixed Point Characterization for IMCs

Let \mathcal{M} be an IMC. For a time interval $I \in \mathcal{I}$ and a set of goal states $G \subseteq \mathcal{S}$, we define the event $\diamond^I G = \{\pi \in \text{Paths}^\omega \mid \exists t \in I. \exists s' \in \pi \circ t. s' \in G\}$ as the set of all paths that are in a state in G during time interval I . The maximum probability induced by $\diamond^I G$ in \mathcal{M} is denoted $p_{max}^{\mathcal{M}}(s, I)$. Formally, it is obtained by the supremum under all GM schedulers:

$$p_{max}^{\mathcal{M}}(s, I) = \sup_{D \in GM} Pr_{\nu_s, D}^\omega(\diamond^I G).$$

Theorem 1 (Fixed point characterization for IMCs). *Let \mathcal{M} be an IMC as before, $G \subseteq \mathcal{S}$ a set of goal states and $I \in \mathcal{I}$ such that $\inf I = a$ and $\sup I = b$. The function $p_{max}^{\mathcal{M}} : \mathcal{S} \times \mathcal{I} \rightarrow [0, 1]$ is the least fixed point of the higher-order operator $\Omega : (\mathcal{S} \times \mathcal{I} \rightarrow [0, 1]) \rightarrow (\mathcal{S} \times \mathcal{I} \rightarrow [0, 1])$ which is defined as follows:*

1. For Markovian states $s \in MS$: $\Omega(F)(s, I)$ equals

$$\begin{cases} \int_0^b E(s) e^{-E(s)t} \cdot \sum_{s' \in \mathcal{S}} \mathbf{P}(s, s') \cdot F(s', I \ominus t) dt & \text{if } s \notin G \\ e^{-E(s)a} + \int_0^a E(s) e^{-E(s)t} \cdot \sum_{s' \in \mathcal{S}} \mathbf{P}(s, s') \cdot F(s', I \ominus t) dt & \text{if } s \in G. \end{cases}$$

2. For interactive states $s \in IS$: $\Omega(F)(s, I)$ equals 1 if $s \in G$ and $0 \in I$, and otherwise, $\Omega(F)(s, I) = \max\{F(s', I) \mid s' \in \text{post}(s)\}$.

Example 3. The fixed point characterization suggests to compute $p_{max}^{\mathcal{M}}(s, I)$ analytically: Consider the IMC \mathcal{M} depicted in Fig. 1 and assume that $G = \{s_3\}$. For $I = [0, b]$, $b > 0$ we have $p_{max}^{\mathcal{M}}(s_3, I) = 1$, $p_{max}^{\mathcal{M}}(s_4, I) = 1 - e^{-0.1b}$ and $p_{max}^{\mathcal{M}}(s_1, I) = \int_0^b e^{-t} \left(\frac{2}{5} \cdot p_{max}^{\mathcal{M}}(s_2, I \ominus t) + \frac{1}{5} \cdot p_{max}^{\mathcal{M}}(s_3, I \ominus t) + \frac{2}{5} \cdot p_{max}^{\mathcal{M}}(s_4, I \ominus t) \right) dt$. For interactive state s_2 , we derive $p_{max}^{\mathcal{M}}(s_2, I) = \max\{p_{max}^{\mathcal{M}}(s_4, I), p_{max}^{\mathcal{M}}(s_1, I)\}$, which yields $p_{max}^{\mathcal{M}}(s_0, I) = \int_0^b 0.9e^{-0.9t} \cdot \left(\frac{2}{3} \cdot p_{max}^{\mathcal{M}}(s_1, I \ominus t) + \frac{1}{3} \cdot p_{max}^{\mathcal{M}}(s_2, I \ominus t) \right) dt$. Hence, an IMC generally induces an integral equation system over the maximum over functions, which is not tractable. Moreover, the iterated integration is numerically unstable [3].

Therefore, we resort to a discretization approach: Informally, we divide the time horizon into small time slices. Then we consider a discrete-time model whose steps correspond to the IMC's behaviour during a single time slice. First, we develop a fixed-point characterization for step bounded reachability on interactive probabilistic chains (IPCs); then we reduce the maximum time interval bounded reachability problem in IMCs to the step interval bounded reachability problem in the discretized IPC. Finally, we show how to solve the latter by a modified value iteration algorithm.

3.2 A Fixed Point Characterization for IPCs

Similar to the timed paths in IMCs, we define $\pi@n \in \mathcal{S}^* \cup \mathcal{S}^\omega$ for the time abstract paths in IPCs: Let $\#^{PS}(\pi, k) = |\{i \in \mathbb{N} \mid 0 \leq i \leq k \wedge \pi[i] \in MS\}|$; then $\#^{PS}(\pi, k)$ is the number of probabilistic transitions that occur up to the $(k+1)$ -th state on π . For fixed $n \in \mathbb{N}$, let i be the smallest index such that $n = \#^{PS}(\pi, i)$. If no such i exists, we set $\pi@n = \langle \rangle$; otherwise i is the index of the n -th probabilistic state that is hit on path π . Similarly, let $j \in \mathbb{N}$ be the largest index (or $+\infty$ if no such finite index exists) such that $n = \#^{PS}(\pi, j)$. Then j denotes the position on π directly before its $(n+1)$ -th probabilistic state. With these preliminaries, we define $\pi@n = \langle s_i, s_{i+1}, \dots, s_{j-1}, s_j \rangle$ to denote the state sequence between the n -th and the $(n+1)$ -th probabilistic state of π . To define step-interval bounded reachability for IPCs, let $k, k' \in \mathbb{N}$ and $k \leq k'$: Then

$$\diamond^{[k, k']}G = \{\pi \in Paths_{abs}^\omega \mid \exists n \in \{k, k+1, \dots, k'\} . \exists s' \in \pi@n. s' \in G\}$$

is the set of paths that visit G between discrete time-step k and k' in an IPC \mathcal{P} .

Accordingly, we define the maximum probability for the event $\diamond^{[k, k']}G$:

$$p_{max}^{\mathcal{P}}(s, [k, k']) = \sup_{D \in GM_{abs}} Pr_{\nu_s, D}^\omega(\diamond^{[k, k']}G).$$

Theorem 2 (Fixed point characterisation for IPCs). *Let $\mathcal{P} = (S, Act, IT, PT, \nu)$ be an IPC, $G \subseteq S$ a set of goal states and $I = [k, k']$ a step interval. The function $p_{max}^{\mathcal{P}}$ is the least fixed point of the higher-order operator $\Omega : (S \times \mathbb{N} \times \mathbb{N} \rightarrow [0, 1]) \rightarrow (S \times \mathbb{N} \times \mathbb{N} \rightarrow [0, 1])$ where*

1. for probabilistic states $s \in PS$:

$$\Omega(F)(s, [k, k']) = \begin{cases} 1 & \text{if } s \in G \wedge k = 0 \\ 0 & \text{if } s \notin G \wedge k = k' = 0 \\ \sum_{s' \in S} PT(s, s') \cdot F(s', [k, k'] \ominus 1) & \text{otherwise;} \end{cases}$$

2. for interactive states $s \in IS$: $\Omega(F)(s, [k, k']) = 1$ if $s \in G$ and $k = 0$. Otherwise, $\Omega(F)(s, [k, k']) = \max_{s' \in post(s)} F(s', [k, k'])$.

Observe that for IMCs, the recursive expression of the probabilistic reachability does not decrease the time interval I for interactive states, whereas for IPCs, the recursive expression does not decrease the corresponding step interval $[k, k']$.

3.3 A Discretization That Reduces IMCs to IPCs

For an IMC \mathcal{M} and a step duration $\tau > 0$, we define the discretized IPC \mathcal{M}_τ of \mathcal{M} :

Definition 8 (Discretization). *An IMC $\mathcal{M} = (S, Act, IT, MT, \nu)$ and a step duration $\tau > 0$ induce the discretized IPC $\mathcal{M}_\tau = (S, Act, IT, PT, \nu)$, where*

$$PT(s, s') = \begin{cases} (1 - e^{-E(s)\tau}) \cdot \mathbf{P}(s, s') & \text{if } s \neq s' \\ (1 - e^{-E(s)\tau}) \cdot \mathbf{P}(s, s') + e^{-E(s)\tau} & \text{if } s = s'. \end{cases} \quad (2)$$




Fig. 2. Interval bounded reachability in IMCs with lower interval bounds

In \mathcal{M}_τ , each probabilistic transition $PT(s, s') > 0$ corresponds to one *time step* of length τ in the underlying IMC \mathcal{M} : More precisely, $PT(s, s')$ is the probability that a transition to state s' occurs within τ time units. In case that $s' = s$, the first summand in $PT(s, s')$ is the probability to take a self-loop back to s , i.e. a transition that leads from s back to s executes; the second summand denotes the probability that no transition occurs within the next τ time units and thus, the systems stays in state $s = s'$.

Now we state the correctness of the discretization: To compute the probability $p_{max}^{\mathcal{M}}(s, [a, b])$, we analyze step-interval bounded reachability in the discretized IPC \mathcal{M}_τ , where each step *approximately* corresponds to τ time units. First we show that $p_{max}^{\mathcal{M}_\tau}(s, [0, \lceil \frac{b}{\tau} \rceil])$ converges from below to $p_{max}^{\mathcal{M}}(s, [0, b])$ if $\tau \rightarrow 0$:

Theorem 3. *Let $\mathcal{M} = (\mathcal{S}, Act, IT, MT, \nu)$ be an IMC, $G \subseteq \mathcal{S}$ a set of goal states, $I = [0, b] \in \mathcal{Q}$ a time interval with $b > 0$ and $\lambda = \max_{s \in MS} E(s)$. Further, let $\tau > 0$ be such that $b = k_b \tau$ for some $k_b \in \mathbb{N}_{>0}$. For all $s \in \mathcal{S}$ it holds:*

$$p_{max}^{\mathcal{M}_\tau}(s, [0, k_b]) \leq p_{max}^{\mathcal{M}}(s, I) \leq p_{max}^{\mathcal{M}_\tau}(s, [0, k_b]) + k_b \cdot \frac{(\lambda\tau)^2}{2}.$$

Example 4. Consider the IMC \mathcal{M} and its discretized IPC \mathcal{M}_τ in Fig. 2(a) and Fig. 2(b), resp. Assume that $G = \{s_2\}$ and fix some $\tau > 0$, $k \in \mathbb{N}_{>0}$. Further, let $I = [0, k\tau]$. In the IMC \mathcal{M} , it holds that $p_{max}^{\mathcal{M}}(s_0, I) = \int_0^{k\tau} \lambda e^{-\lambda t} \cdot p_{max}^{\mathcal{M}}(s_1, I \ominus t) dt = 1 - e^{-\lambda k\tau}$. In \mathcal{M}_τ , we obtain $p_{max}^{\mathcal{M}_\tau}(s_0, [0, k]) = \sum_{i=1}^k (e^{-\lambda\tau})^{i-1} (1 - e^{-\lambda\tau}) = 1 - e^{-\lambda k\tau}$, which is the geometric distribution function for parameter $p = 1 - e^{-\lambda\tau}$.

So far, we only considered intervals of the form $I = [0, b]$, $b > 0$. In what follows, we extend our results to arbitrary intervals. However, this is slightly involved:

If $s \in MS$ is a Markovian state and $b > 0$, then $p_{max}^{\mathcal{M}}(s, (0, b]) = p_{max}^{\mathcal{M}}(s, [0, b])$. However this is not true for interactive states: If s_1 (instead of s_0) is made the only initial state in \mathcal{M} and \mathcal{M}_τ of Fig. 2, the probability to reach s_2 within interval $[0, b]$ is 1 whereas it is 0 for the right-semiclosed interval $(0, b]$. Further, the discretization is imprecise for point intervals: To see this, note that if $I = [\tau, \tau]$, then $p_{max}^{\mathcal{M}}(s_0, I) = 0$, whereas $p_{max}^{\mathcal{M}_\tau}(s_0, [1, 1]) = 1 - e^{-\lambda\tau}$.

Now, let $I = [k_a\tau, k_b\tau]$ be a *closed* interval with $k_a, k_b \in \mathbb{N}$ and $0 < k_a < k_b$. In the IMC \mathcal{M} in Fig. 2(a), we obtain $p_{max}^{\mathcal{M}}(s_0, I) = \int_{k_a\tau}^{k_b\tau} \lambda e^{-\lambda t} \cdot p_{max}^{\mathcal{M}}(s_1, I \ominus t) dt = e^{-\lambda k_a\tau} - e^{-\lambda k_b\tau}$, whereas for its discretized IPC \mathcal{M}_τ (see Fig. 2(b)), we derive

$$p_{max}^{\mathcal{M}_\tau}(s_0, [k_a, k_b]) = \sum_{i=k_a}^{k_b} (e^{-\lambda\tau})^{i-1} \cdot (1 - e^{-\lambda\tau}) = e^{-\lambda(k_a-1)\tau} - e^{-\lambda k_b\tau}.$$

Clearly, the two probabilities differ in the first term by a factor of $e^{\lambda\tau}$. To see the reason, let $k_a = 2$ and $k_b = 3$: We have $p_{max}^{\mathcal{M}}(s, [2\tau, 3\tau]) = e^{-2\lambda\tau} - e^{-3\lambda\tau}$; however, in \mathcal{M}_τ it holds $p_{max}^{\mathcal{M}_\tau}(s, [2, 3]) = e^{-\lambda\tau} \cdot (1 - e^{-\lambda\tau}) + e^{-2\lambda\tau} \cdot (1 - e^{-\lambda\tau}) = e^{-\lambda\tau} - e^{-3\lambda\tau}$. As each step in \mathcal{M}_τ corresponds to a time interval of length τ (cf. Fig. 3), the interval bounds 2τ and 3τ fall in different discretization steps. Hence in the discretization, we add two steps which leads to an error. If instead we compute $p_{max}^{\mathcal{M}}(s, (2\tau, 3\tau])$, we obtain $p_{max}^{\mathcal{M}_\tau}(s, (2, 3]) = p_{max}^{\mathcal{M}_\tau}(s, [3, 3]) = e^{-2\lambda\tau} - e^{-3\lambda\tau}$, as desired.




Fig. 3. Discretization steps

Based on these observations, we extend Thm. 3 to intervals with positive lower bounds. To avoid some technicalities, we first restrict to right-semiclosed intervals:

Theorem 4. *Let $\mathcal{M} = (\mathcal{S}, Act, IT, MT, \nu)$ be an IMC, $G \subseteq \mathcal{S}$ a set of goal states, $I = (a, b] \in \mathcal{Q}$ a time interval with $a < b$ and $\lambda = \max_{s \in MS} E(s)$. If $\tau > 0$ is such that $a = k_a\tau$ and $b = k_b\tau$ for some $k_a, k_b \in \mathbb{N}$, then it holds for all $s \in \mathcal{S}$:*

$$p_{max}^{\mathcal{M}_\tau}(s, (k_a, k_b]) - k_a \cdot \frac{(\lambda\tau)^2}{2} \leq p_{max}^{\mathcal{M}}(s, I) \leq p_{max}^{\mathcal{M}_\tau}(s, (k_a, k_b]) + k_b \cdot \frac{(\lambda\tau)^2}{2} + \lambda\tau.$$

The error bounds for the case of lower interval bounds that are stated in Thm. 4 are derived using double induction over k_a and k_b , respectively.

Theorem 5. *If \mathcal{M}, G and τ are as in Thm. 4 and $I \in \mathcal{Q}$ is a time interval with $\inf I = a$ and $\sup I = b$ such that $a < b$ and $a = k_a\tau, b = k_b\tau$ for $k_a, k_b \in \mathbb{N}$ and $0 \notin I$, then*

$$p_{max}^{\mathcal{M}_\tau}(s, (k_a, k_b]) - k_a \cdot \frac{(\lambda\tau)^2}{2} \leq p_{max}^{\mathcal{M}}(s, I) \leq p_{max}^{\mathcal{M}_\tau}(s, (k_a, k_b]) + k_b \cdot \frac{(\lambda\tau)^2}{2} + \lambda\tau.$$

For the remaining cases, note that for all states $s \in \mathcal{S}$ and intervals $I = \emptyset$ or $I = [a, a]$ with $a > 0$ it holds that $p_{max}^{\mathcal{M}}(s, I) = 0$. Finally, for the case that $I = [0, 0]$, an interactive reachability analysis suffices to compute $p_{max}^{\mathcal{M}}(s, I)$, which is either 1 or 0.

3.4 Solving the Problem on the Reduced IPC

Let $\mathcal{P} = (\mathcal{S}, Act, IT, PT, \nu)$ be an IPC, $G \subseteq \mathcal{S}$ a set of goal states and $[k_a, k_b]$ a step interval. In this section, we discuss how to compute $p_{max}^{\mathcal{P}}(s, [k_a, k_b])$ via a modification of the well known *value iteration* algorithm [5]. The adaptation is non-trivial, as we consider step intervals that correspond to the number of *probabilistic steps* that are taken. This is reflected in our algorithm which only decreases the step counter for probabilistic, but not for internal transitions. We discuss step bounded reachability first:

Step Bounded Reachability: We aim at computing $p_{max}^{\mathcal{P}}(s, [0, k])$ for $0 \leq k$. This works as follows: In each step $i = 0, 1, \dots, k$ of the iteration, we use two vectors $\vec{v}_i \in [0, 1]^{\mathcal{S}}$ and $\vec{u}_i \in [0, 1]^{\mathcal{S}}$, where \vec{v}_i is the probability vector obtained from \vec{u}_{i-1} by one step in the classical value iteration algorithm and \vec{u}_i is obtained by computing the backwards closure along interactive transitions w.r.t. \vec{v}_{i-1} .

Each of the k value iteration steps consists of two phases: First, \vec{v}_i is computed: If $s \in PS \cap G$, then $\vec{v}_i(s) = 1$. If $s \in PS \setminus G$, then $\vec{v}_i(s)$ is the weighted sum of the probabilistic successor states s' of s , multiplied by the result $\vec{u}_{i-1}(s')$ of the previous step. In the second phase, \vec{u}_i is obtained by the backward closure of \vec{v}_i along internal transitions. Initially, we set $\vec{v}_0(s) = 1$ if $s \in G$, and $\vec{v}_0(s) = 0$, otherwise. Then: $\forall i \in \{0, \dots, k\}. \vec{u}_i(s) = \max \{ \vec{v}_i(s') \mid s \rightsquigarrow_i^* s' \}$ and for \vec{v}_i :

$$\forall i \in \{1, \dots, k\}. \vec{v}_i(s) = \begin{cases} \sum_{s' \in \mathcal{S}} PT(s, s') \cdot \vec{u}_{i-1}(s') & \text{if } s \in PS \setminus G \\ 1 & \text{if } s \in PS \cap G \\ \vec{u}_{i-1}(s) & \text{if } s \in IS. \end{cases}$$

For efficiency reasons the set $\{s' \in \mathcal{S} \mid s \rightsquigarrow_i^* s'\}$ can be precomputed by a backwards search in the interactive reachability graph of \mathcal{P} .

After k value iteration steps $p_{max}^{\mathcal{P}}(s, [0, k])$ is obtained as the probability in $\vec{u}_k(s)$.

Step-Interval Bounded Reachability: In this part, we compute $p_{max}^{\mathcal{P}}(s, [k_a, k_b])$, for interval bounds $0 < k_a < k_b$. Again, we compute a sequence $\vec{v}_0, \vec{u}_0, \dots, \vec{v}_{k_b}, \vec{u}_{k_b}$. As $k_a > 0$, we split the value iteration in two parts: In the first $k_b - k_a$ value iteration steps, we proceed as before and compute the probability vectors $\vec{v}_0, \vec{u}_0, \dots, \vec{v}_{k_b - k_a}, \vec{u}_{k_b - k_a}$. Thus, we compute the probabilities $p_{max}^{\mathcal{P}}(s, [0, k_b - k_a])$ for all $s \in \mathcal{S}$.

The vector $\vec{v}_{k_b - k_a}$ provides the initial probabilities of the second part: In the remaining $i \in \{k_b - k_a + 1, \dots, k_b\}$ value iteration steps, we set $\vec{v}_i(s) = 0$ if $s \in IS$ and $\vec{v}_i(s) = \sum_{s' \in \mathcal{S}} PT(s, s') \cdot \vec{u}_{i-1}(s')$ if $s \in PS$. The vectors \vec{u}_i are as before. To see why, note that the value iteration algorithm proceeds in a backward manner, starting from the goal states. We do not set $\vec{v}_i(s) = 1$ if $s \in G$ in the last k_a iteration steps, as in the first k_a transitions, reaching a goal state does not satisfy our reachability objective. To avoid that the probabilities of interactive states $s \in IS$ erroneously propagate in the vectors $\vec{u}_i(s)$ from the first to the second part, in the second part we define $\vec{v}_i(s) = 0$ for all $s \in IS$ (instead of $\vec{v}_i(s) = \vec{u}_{i-1}(s)$ as in the first part). Let us illustrate this:

Example 5. We compute $p_{max}^{\mathcal{P}}(s, [1, 2])$ in the IPC \mathcal{P} in Fig. 4 for initial state s_0 and goal state s_3 : In the first part, apply the value iteration to compute \vec{u}_1 : $\vec{v}_0(s) = 1$ if $s = s_3$ and 0, otherwise. By the backwards closure, $\vec{u}_0 = (1, 0, 0, 1)$. Thus $p_{max}^{\mathcal{P}}(s_0, [0, 0]) = 1$, as s_0 can reach G by the interactive α -transition. For \vec{v}_1 , we have $\vec{v}_1(s_0) = \vec{u}_0(s_0) = 1$ and $\vec{v}_1(s_1) = \frac{1}{2}\vec{u}_0(s_3) + \frac{1}{2}\vec{u}_0(s_2) = \frac{1}{2}$. In this way, we obtain $\vec{v}_1 = (1, \frac{1}{2}, \frac{1}{4}, 1)$ and $\vec{u}_1 = (1, \frac{1}{2}, \frac{1}{4}, 1)$. With the probabilities \vec{u}_1 , the first part ends after $k_b - k_a = 1$ value iteration steps. As $k_a = 1$, one iteration for the lower step bound follows. Here $\vec{v}_2(s_0) = \vec{v}_2(s_3) = 0$ as $s_0, s_3 \in IS$; further $\vec{v}_2(s_1) = \frac{1}{2}\vec{u}_1(s_3) + \frac{1}{2}\vec{u}_1(s_2) = \frac{5}{8}$ and $\vec{v}_2(s_2) = \frac{1}{2}\vec{u}_1(s_2) + \frac{1}{4}\vec{u}_1(s_3) + \frac{1}{4}\vec{u}_1(s_1) = \frac{1}{2}$. Finally, $\vec{u}_2 = (\frac{5}{8}, \frac{5}{8}, \frac{1}{2}, \frac{1}{2})$. Therefore, we obtain that $p_{max}^{\mathcal{P}}(s_0, [1, 2]) = \vec{u}_2(s_0) = \frac{5}{8}$.

3.5 Algorithm and Complexity

Let $\mathcal{M}, G, \varepsilon$ and I as before, with $b = \sup I$. For $\varepsilon > 0$, choose k_b such that $k_b \cdot \frac{(\lambda\tau)^2}{2} + \lambda\tau \leq \varepsilon$. With $\tau = \frac{b}{k_b}$, the smallest such k_b is $k_b = \lceil \frac{\lambda^2 b^2 + 2\lambda b}{2\varepsilon} \rceil$. Then the step duration τ induces the discretized IPC \mathcal{M}_τ . By Thm. 5, $p_{max}^{\mathcal{M}}(s_0, I)$ can be approximated (up to ε) by $p_{max}^{\mathcal{M}_\tau}(s_0, (k_a, k_b])$. Let $n = |\mathcal{S}|$ and $m = |IT| + |MT|$ be the number of states and

transitions of \mathcal{M} , respectively. In the worst case, \mathcal{M}_τ has n states, and $m + n$ transitions. In each value iteration step, the update of the vector \vec{v}_i takes at most time $m + n$; for \vec{u}_i , the sets $Reach^i(s)$ are recomputed. In the general case, the best theoretical complexity for computing the reflexive transitive closure is in $\mathcal{O}(n^{2.376})$, as given by [13]. As $m^* \subseteq \mathcal{S} \times \mathcal{S}$, the number of transitions in the closure m^* is bounded by n^2 . Hence, with an appropriate precomputation of m^* , updating \vec{u}_i takes time $\mathcal{O}(n^2)$. Therefore, with k_b value iteration steps, the worst case time complexity of our approach is in $n^{2.376} + (m + n + n^2) \cdot (\lambda b) \cdot (\lambda b + 2) / (2\varepsilon) \in \mathcal{O}(n^{2.376} + (m + n^2) \cdot (\lambda b)^2 / \varepsilon)$.




Fig. 4. Example IPC

4 Model Checking the Continuous Stochastic Logic

For model checking, we consider a finite set $AP = \{a, b, c, \dots\}$ of *atomic propositions* and *state labelled IMCs*: A *state labelling function* $L : \mathcal{S} \rightarrow 2^{AP}$ assigns to each state the set of atomic propositions that hold in that state. To specify quantitative properties, we extend the continuous stochastic logic (CSL) [3,12], which reasons about qualitative and quantitative properties of CTMCs to the nondeterministic setting:

Definition 9 (CSL syntax). For $a \in AP$, $p \in [0, 1]$, $I \subseteq \mathcal{Q}$ an interval and $\trianglelefteq \in \{<, \leq, \geq, >\}$, CSL state and CSL path formulas are defined by

$$\Phi ::= a \mid \neg\Phi \mid \Phi \wedge \Phi \mid \mathcal{P}_{\trianglelefteq p}(\varphi) \quad \text{and} \quad \varphi ::= \mathcal{X}^I\Phi \mid \Phi \mathcal{U}^I \Phi.$$

Intuitively, a path $\pi \in Paths^\omega$ satisfies the formula $\mathcal{X}^I\Phi$ ($\pi \models \mathcal{X}^I\Phi$) if the first transition on π occurs in time-interval I and leads to a successor state in $Sat(\Phi)$. Similarly, π satisfies the until formula $\Phi \mathcal{U}^I \Psi$ if a state in $Sat(\Psi)$ is reached at some time point $t \in I$ and before that, all states satisfy state formula Φ .

Definition 10 (CSL semantics). Let $\mathcal{M} = (\mathcal{S}, Act, IT, MT, AP, L, \nu)$ be a state labelled IMC, $s \in \mathcal{S}$, $a \in AP$, $I \in \mathcal{Q}$, $\trianglelefteq \in \{<, \leq, \geq, >\}$ and $\pi \in Paths^\omega$. For state formulas, we define $s \models a$ iff $a \in L(s)$, $s \models \neg\Phi$ iff $s \not\models \Phi$ and $s \models \Phi \wedge \Psi$ iff $s \models \Phi$ and $s \models \Psi$. Further, $s \models \mathcal{P}_{\trianglelefteq p}(\varphi)$ iff for all $D \in GM$ it holds that $Pr_{\nu_s, D}^\omega \{\pi \in Paths^\omega \mid \pi \models \varphi\} \trianglelefteq p$. For path formulas, we define

$$\begin{aligned} \pi \models \mathcal{X}^I\Phi &\iff \pi[1] \models \Phi \wedge \delta(\pi, 0) \in I \\ \pi \models \Phi \mathcal{U}^I \Psi &\iff \exists t \in I. \exists s \in \pi @ t. s \models \Psi \wedge \forall s' \in Pref(\pi @ t, s). s' \models \Phi \\ &\quad \wedge \forall t' \in [0, t). \forall s'' \in \pi @ t'. s'' \models \Phi. \end{aligned}$$

To model check an IMC w.r.t. a CSL state formula Φ , we successively consider the state subformulas Ψ of Φ and calculate the sets $Sat(\Psi) = \{s \in \mathcal{S} \mid s \models \Psi\}$. For atomic propositions, conjunction and negation, this is easy as $Sat(a) = \{s \in \mathcal{S} \mid a \in L(s)\}$, $Sat(\neg\Psi) = \mathcal{S} \setminus Sat(\Psi)$ and $Sat(\Psi_1 \wedge \Psi_2) = Sat(\Psi_1) \cap Sat(\Psi_2)$. Therefore we only discuss the probabilistic operator $\mathcal{P}_{\trianglelefteq p}(\varphi)$ for next and bounded until formulas. To decide $Sat(\mathcal{P}_{\trianglelefteq p}(\varphi))$, it suffices to maximize (or minimize, which can be done

similarly) $Pr_{\nu_s, D}^\omega(\{\pi \in Paths^\omega \mid \pi \models \varphi\})$ w.r.t. all schedulers $D \in GM$. We define $p_{max}^M(s, \varphi) = \sup_{D \in GM} Pr_{\nu_s, D}^\omega(\{\pi \in Paths^\omega \mid \pi \models \varphi\})$ and consider both types of path formulas:


The Next Formula. Computing $p_{max}^M(s, \mathcal{X}^I \Phi)$ is easy: We proceed inductively on the structure of the formula and assume that $Sat(\Phi)$ is already computed. Let $a = \inf I$, $b = \sup I$ and $s \in MS$. Then $p_{max}^M(s, \mathcal{X}^I \Phi) = \int_a^b E(s) e^{-E(s)t} \cdot \sum_{s' \in Sat(\Phi)} \mathbf{P}(s, s') dt = \mathbf{P}(s, Sat(\Phi)) \cdot (e^{-E(s)a} - e^{-E(s)b})$, where $\mathbf{P}(s, Sat(\Phi)) = \sum_{s' \in Sat(\Phi)} \mathbf{P}(s, s')$ is the probability to move to a successor state $s' \in Sat(\Phi)$. If $s \in IS$, $0 \in I$ and $post(s) \cap Sat(\Phi) \neq \emptyset$, then $p_{max}^M(s, \mathcal{X}^I \Phi) = 1$; otherwise $p_{max}^M(s, \mathcal{X}^I \Phi) = 0$.

The Until Formula. Let $\varphi = \Phi U^I \Psi$ with $I \in \mathcal{Q}$ and assume that $Sat(\Phi)$ and $Sat(\Psi)$ are already computed. We reduce the problem to compute $p_{max}^M(s, \varphi)$ to the maximum interval-bounded reachability problem: Therefore, define $\mathcal{S}_{=0}^\varphi = \{s \in \mathcal{S} \mid s \models \neg \Phi\}$. In the next step, we turn all states $s \in \mathcal{S}_{=0}^\varphi$ into absorbing states by replacing all its outgoing transitions by a single interactive self loop. This is similar to the approach taken in [3,6] for model checking CTMCs and MDPs. Formally, a state $s \in IS$ is *absorbing* iff $post^i(s) = \{s\}$. Hence, as soon as a path enters an absorbing state, it cannot reach a different state anymore. Moreover, due to the maximal progress assumption, time does not progress any further in absorbing states. Intuitively, making $\mathcal{S}_{=0}^\varphi$ -states absorbing is justified as follows. If a path π enters a state $s \in \mathcal{S}_{=0}^\varphi$, it can be decided immediately whether $\pi \models \Phi U^I \Psi$, or not: If $s \models \Psi$ holds and if state s is entered at some time in the interval I , then $\pi \models \Phi U^I \Psi$. Otherwise $\pi \not\models \Phi U^I \Psi$ holds.

Theorem 6 (Time-bounded until). Let $\mathcal{M} = (\mathcal{S}, Act, IT, MT, AP, L, \nu)$ be a state labelled IMC, $\varphi = \Phi U^I \Psi$ a CSL path formula with $I \in \mathcal{Q}$ and $G = Sat(\Psi)$ the set of goal states. Further, assume that all states $s \in \mathcal{S}_{=0}^\varphi$ are made absorbing. Then

$$p_{max}^M(s, \Phi U^I \Psi) = p_{max}^M(s, I) \quad \text{for all } s \in \mathcal{S}.$$

Theorem 6 reduces the problem to compute $p_{max}^M(s, \Phi U^I \Psi)$ of the until formula to the problem of computing the interval bounded reachability probability $p_{max}^M(s, I)$ with



(a) Time-bounded reachability in \mathcal{M}

problem	states	ε	λ	b	prob.	time
<i>Erl</i> (30, 10)	35	10^{-3}	10	4	0.672	50s
<i>Erl</i> (30, 10)	35	10^{-3}	10	7	0.983	70s
<i>Erl</i> (30, 10)	35	10^{-4}	10	4	0.6718	268s
ws-cl, $N=4$	820	10^{-6}	$2 \cdot 10^1$	$3.3 \cdot 10^{-5}$		2d
ws-cl, $N=4$	820	10^{-4}	$2 \cdot 10^2$	$4 \cdot 10^{-4}$		15h
ws-cl, $N=4$	820	10^{-3}	$2 \cdot 10^3$	$5 \cdot 10^{-3}$		6d

(b) Computation time for different parameters

Fig. 5. Experimental results for *Erl*(30, 10) and the workstation cluster from [16]

respect to the set of goal states $G = Sat(\Psi)$. The latter can be computed efficiently by the discretization approach introduced in Sec. 3.3.

For CSL state-formula Φ , let $|\Phi|$ be the number of state subformulas of Φ . In the worst case, the interval bounded reachability probability is computed $|\Phi|$ times. Hence the model checking problem has time complexity $\mathcal{O}(|\Phi| \cdot (n^{2.376} + (m + n^2) \cdot (\lambda b)^2 / \varepsilon))$.

5 Experimental Results

We consider the IMC in Fig. 6, where $Erl(30, 10)$ denotes a transition with an Erlang (k, λ) distributed delay: This corresponds to $k = 30$ consecutive Markovian transitions each of which has rate λ . The mean time to move from s_2 to the goal s_4 is $\frac{k}{\lambda} = 3$ with a variance of $\frac{k}{\lambda^2} = \frac{3}{10}$. Hence, with very high probability we move from s_2 to s_4 after approximately 3 time units. The decision that maximizes the probability to reach s_4 in time interval $[0, b]$ in state s_1 depends on the sojourn in state s_0 . Fig. 5(a) depicts the computed maxima for time dependent schedulers and the upper part of Tab. 5(b) lists some performance measurements.

If $AP = \{g\}$ and s_4 is the only state labelled with g , we can verify the CSL formula $\Phi = \mathcal{P}_{\geq 0.5}(\diamond^{[3,4]}g)$ by computing $p_{max}^{\mathcal{M}}(s_0, [3, 4])$ with the modified value iteration. The result $p_{max}^{\mathcal{M}}(s_0, [3, 4]) = 0.6057$ meets the bound ≥ 0.5 in Φ , implying that $s_0 \models \Phi$.

Finally, the lower part of Tab. 5(b) lists the performance of our approach for a large scale example [16], where we conduct a dependability analysis of a cluster of $2N$ workstations to estimate its failure probability over a finite time horizon. This rather stiff model has a high computational complexity in our prototypical implementation, as the failure events are very rare which leads to a large time horizon.

All measurements were carried out on a 2.2GHz Xeon CPU with 16GB RAM.

6 Related Work and Conclusions

In the setting of stochastic games, the time-bounded reachability problem has been studied extensively in [11], with extensions to timed automata in [9]. Closely related to ours is the work in [7], where globally uniform IMCs — which require the sojourn times in all Markovian states to be equally distributed — are transformed into continuous-time Markov decision processes (CTMDPs). Subsequently, the algorithm in [4] is used to compute the maximum time-bounded reachability probability in the resulting globally uniform CTMDP. However, the applicability of this approach is severely restricted, as global uniformity is hard (and often impossible) to achieve.

Further, the above approaches rely on time-abstract schedulers which are proved to be strictly less powerful than the time-dependent ones that we consider here [4,24].

In [25], we relax the restriction to global uniformity and consider locally uniform CTMDPs for which we propose a discretization that computes maximum time-bounded




Fig. 6. The $Erl(30, 10)$ model \mathcal{M}

reachability probabilities under *late schedulers*: In locally uniform CTMDPs, late schedulers outperform *early schedulers* [24], which are the largest class of history and time dependent schedulers definable on general CTMDPs [21].

The discretization approach in this paper resembles that of [25]. However, the results are complementary: In general, transforming IMCs to CTMDPs as done in [21] does not yield locally uniform CTMDPs. Hence, the approach in [25] is inapplicable for the analysis of IMCs. However, we expect to solve the problem of computing time-interval bounded reachability in CTMDPs by analysing the CTMDP's induced IMC.

By providing an efficient and quantifiably precise approximation algorithm to compute interval bounded reachability probabilities, this paper solves a long standing open problem in the area of performance and dependability evaluation. Moreover, we solve the CSL model checking problem on arbitrary IMCs.

Acknowledgement. We thank Holger Hermanns and Joost-Pieter Katoen for their comments and for many fruitful discussions about earlier versions of this work.

References

1. Ash, R., Doléans-Dade, C.: Probability & Measure Theory, 2nd edn. Academic Press, London (2000)
2. Aziz, A., Sanwal, K., Singhal, V., Brayton, R.K.: Verifying continuous time Markov chains. In: Alur, R., Henzinger, T.A. (eds.) CAV 1996. LNCS, vol. 1102, pp. 269–276. Springer, Heidelberg (1996)
3. Baier, C., Haverkort, B.R., Hermanns, H., Katoen, J.-P.: Model-checking algorithms for continuous-time Markov chains. IEEE TSE 29, 524–541 (2003)
4. Baier, C., Hermanns, H., Katoen, J.-P., Haverkort, B.R.: Efficient computation of time-bounded reachability probabilities in uniform continuous-time Markov decision processes. Theor. Comp. Sci. 345, 2–26 (2005)
5. Bertsekas, D.: Dynamic Programming and Optimal Control, vol. II. Athena Scientific, Belmont (1995)
6. Bianco, A., de Alfaro, L.: Model checking of probabilistic and nondeterministic systems. In: Thiagarajan, P.S. (ed.) FSTTCS 1995. LNCS, vol. 1026, pp. 499–513. Springer, Heidelberg (1995)
7. Böde, E., Herbstritt, M., Hermanns, H., Johr, S., Peikenkamp, T., Pulungan, R., Rakow, J., Wimmer, R., Becker, B.: Compositional dependability evaluation for STATEMATE. IEEE Trans. Software Eng. 35, 274–292 (2009)
8. Boudali, H., Crouzen, P., Haverkort, B.R., Kuntz, M., Stoelinga, M.: Architectural dependability evaluation with Arcade. In: DSN, pp. 512–521. IEEE, Los Alamitos (2008)
9. Bouyer, P., Forejt, V.: Reachability in stochastic timed games. In: Albers, S., Marchetti-Spaccamela, A., Matias, Y., Nikolettseas, S., Thomas, W. (eds.) ICALP 2009. LNCS, vol. 5556, pp. 103–114. Springer, Heidelberg (2009)
10. Bravetti, M., Hermanns, H., Katoen, J.-P.: YMCA: Why Markov chain algebra? In: Essays on Algebraic Process Calculi. Electronic Notes in Theoretical Computer Science, vol. 162, pp. 107–112. Elsevier, Amsterdam (2006)
11. Brazdil, T., Forejt, V., Krcal, J., Kretinsky, J., Kucera, A.: Continuous-time stochastic games with time-bounded reachability. In: FSTTCS. LIPIcs (2009) (to appear)
12. Cerotti, D., Donatelli, S., Horváth, A., Sproston, J.: CSL model checking for generalized stochastic Petri nets. In: QEST, pp. 199–210. IEEE, Los Alamitos (2006)

13. Coppersmith, D., Winograd, S.: Matrix multiplication via arithmetic progressions. In: ACM Symposium on Theory of Computing. ACM, New York (1987)
14. Coste, N., Garavel, H., Hermanns, H., Hersemeule, R., Thonnart, Y., Zidouni, M.: Quantitative evaluation in embedded system design: Validation of multiprocessor multithreaded architectures. In: DATE, pp. 88–89. IEEE, Los Alamitos (2008)
15. Coste, N., Hermanns, H., Lantreibeq, E., Serwe, W.: Towards performance prediction of compositional models in industrial GALS designs. In: Bouajjani, A., Maler, O. (eds.) Computer Aided Verification. LNCS, vol. 5643, pp. 204–218. Springer, Heidelberg (2009)
16. Haverkort, B.R., Hermanns, H., Katoen, J.-P.: On the use of model checking techniques for dependability evaluation. In: Reliable Distributed Systems, pp. 228–239. IEEE, Los Alamitos (2000)
17. Hermanns, H. (ed.): Interactive Markov Chains: The Quest for Quantified Quality. LNCS, vol. 2428. Springer, Heidelberg (2002)
18. Hermanns, H., Herzog, U., Katoen, J.-P.: Process algebra for performance evaluation. *Theor. Comp. Sci.* 274, 43–87 (2002)
19. Hermanns, H., Katoen, J.-P.: Automated compositional Markov chain generation for a plain-old telephone system. *Sci. Comput. Program.* 36, 97–127 (2000)
20. Hillston, J.: A Compositional Approach to Performance Modelling. Cambridge University Press, Cambridge (1996)
21. Johr, S.: Model Checking Compositional Markov Systems. PhD thesis, Saarland University, Saarbrücken, Germany (2007)
22. Maciá, H., Valero, V., Cuartero, F., Ruiz, M.C.: sPBC: A Markovian extension of Petri box calculus with immediate multiactions. *Fundamenta Informaticae* 87, 367–406 (2008)
23. Neuhäuser, M.R.: Model Checking Nondeterministic and Randomly Timed Systems. PhD thesis, RWTH Aachen University, Aachen, Germany (2010)
24. Neuhäuser, M.R., Stoelinga, M., Katoen, J.-P.: Delayed nondeterminism in continuous-time Markov decision processes. In: de Alfaro, L. (ed.) FOSSACS 2009. LNCS, vol. 5504, pp. 364–379. Springer, Heidelberg (2009)
25. Neuhäuser, M.R., Zhang, L.: Time-bounded reachability in continuous-time Markov decision processes. Technical report, RWTH Aachen University (2009)
26. Pulungan, R.: Reduction of Acyclic Phase-Type Representations. PhD thesis, Universität des Saarlandes, Saarbrücken, Germany (2009)