

Efficient Computation of Time-Bounded Reachability Probabilities in Uniform Continuous-Time Markov Decision Processes^{*}

Christel Baier¹, Boudewijn Haverkort²,
Holger Hermanns^{2,3}, and Joost-Pieter Katoen²

¹ Institut für Informatik I, University of Bonn, Germany

² Faculty of Electrical Engineering, Mathematics and Computer Science,
University of Twente, The Netherlands

³ Department of Computer Science, Saarland University, Germany

Abstract. A continuous-time Markov decision process (CTMDP) is a generalization of a continuous-time Markov chain in which both probabilistic and nondeterministic choices co-exist. This paper presents an efficient algorithm to compute the maximum (or minimum) probability to reach a set of goal states within a given time bound in a uniform CTMDP, i.e., a CTMDP in which the delay time distribution per state visit is the same for all states. We prove that these probabilities coincide for (time-abstract) history-dependent and Markovian schedulers that resolve nondeterminism either deterministically or in a randomized way.

1 Introduction

Why continuous-time Markov decision processes? A continuous-time Markov decision process (CTMDP) [6,14,21,24] is a generalization of a continuous-time Markov chain (CTMC) in which both probabilistic and nondeterministic choices co-exist. CTMDPs occur in many contexts, ranging from stochastic control theory [14] to dynamic power management [22]. We are particularly interested in this class of models because they can be viewed as a common semantic model for various performance and dependability modelling formalisms including generalised stochastic Petri nets [1], Markovian stochastic activity models [23], interactive Markov chains (IMC) [17] and TIPP process algebra [16].

So far, the analysis of models developed in these and related formalisms was restricted to the subset that corresponds to CTMCs, usually referred to as 'non-confused', 'well-defined', or 'well-specified' models [11,12,13,17]. All these notions are semantic notions, usually checked by an exhaustive exploration of the state space, with models being discarded if the check fails. In other words, no specification-level check is available, and the offered analysis algorithms are actually partial algorithms.

^{*} This work is supported by the NWO-DFG bilateral project Validation of Stochastic Systems.

Why time-bounded reachability? Model checking of CTMCs [4] has received remarkable attention in recent years. Various model checkers exist [18,20,10], answering questions such as: *Is the probability to hop along Φ -states, until reaching a Ψ -state within 5 to 10 time units greater than 0.95?* The core algorithmic innovation allowing to answer such questions is a mapping from interval-bounded until-formulae – specified in the continuous stochastic logic CSL [2] – to time-bounded reachability problems, which in turn can be approximated efficiently using a stable numeric technique called uniformization [19]. To enable the same kind of questions being answered for models specified in any of the above mentioned formalisms, the key problem is how to compute time-bounded reachability probabilities in CTMDPs. This is the problem we address in this paper. While model checking algorithms for discrete-time Markov decision processes are well-understood [8,5], we are not aware of any model checking algorithm for continuous-time Markov decision processes.

Our contribution. Given a CTMDP, our aim is to compute the maximum (or minimum) probability to reach – under a given class of schedulers – a certain set of states within t time units, given a starting state. We consider this problem for uniform CTMDPs, a class of CTMDPs in which the delay time distribution per state visit is the same for all states, governed by a unique exit rate E . We show that an efficient greedy algorithm can be obtained using truncated Markovian deterministic (MD)-schedulers, that is, step-dependent schedulers which schedule up to a limited depth. The algorithm computes the maximum (or minimum) probabilities for timed reachability. It is then shown that these probabilities for truncated MD-schedulers coincide with the maximum (or minimum) probabilities for timed reachability for Markovian and history-dependent schedulers (both deterministic and randomized). We show that stationary Markovian schedulers – as opposed to the discrete case [8,5] – yield a smaller maximum, whereas timed history-dependent schedulers may yield a higher probability.

The main result of this paper is a computationally efficient approximation algorithm for computing maximum (or minimum) probabilities for timed reachability in uniform CTMDPs under all time-abstract schedulers. The time complexity is in $\mathcal{O}(E \cdot t \cdot N^2 \cdot M)$ and the space complexity in $\mathcal{O}(N^2 \cdot M)$ where t is the time bound, N is the number of states, and M is the number of actions in the CTMDP under consideration.

Organization of the paper. Section 2 introduces the necessary background. Section 3 presents the algorithm for uniform CTMDP. Section 4 places the result of the algorithm in the context of other classes of schedulers. Section 5 discusses the problem of uniformizing arbitrary CTMDPs. Section 6 concludes the paper. Proofs of all theorems can be found in [3].

2 Background and Problem Statement

2.1 Continuous-Time Markov Decision Processes

A *continuous-time Markov decision process* \mathcal{M} , is a tuple (S, Act, \mathbf{R}) with S a finite set of states, Act a finite set of actions, and $\mathbf{R} : (S \times Act \times S) \rightarrow \mathbb{R}_{\geq 0}$

the three-dimensional rate matrix. For each state $s \in S$ we require the existence of a pair $(\alpha, s') \in Act \times S$ with $\mathbf{R}(s, \alpha, s') > 0$. Note that this can easily be established by adding self-loops, i.e., having $\mathbf{R}(s, \alpha, s) > 0$ for some $\alpha \in Act$. If Act is a singleton, we can project \mathbf{R} on an $(S \times S)$ matrix, and \mathcal{M} is a CTMC.

For $B \subseteq S$, let $\mathbf{R}(s, \alpha, B)$ denote the total rate to move from state s via action α to some state in B , i.e.,

$$\mathbf{R}(s, \alpha, B) = \sum_{s' \in B} \mathbf{R}(s, \alpha, s').$$

The behavior of a CTMDP is as follows. $\mathbf{R}(s, \alpha, s') > 0$ intuitively means that there is a transition from s to s' under action α . If state s has outgoing transitions for distinct actions, one of these actions is selected nondeterministically where we assume that the nondeterminism is resolved by means of a scheduler (also called policy or adversary). Given that action α has been chosen, $1 - e^{-\mathbf{R}(s, \alpha, s') \cdot t}$ is the probability that the α -transition $s \rightarrow s'$ can be triggered within t time units. Thus, the delay of α -transition $s \rightarrow s'$ is governed by the negative exponential distribution with rate $\mathbf{R}(s, \alpha, s')$. If $\mathbf{R}(s, \alpha, s') > 0$ for more than one state s' , a competition between the α -transitions originating in s exists, known as the race condition. The discrete probability of selecting α -transition $s \rightarrow s'$ is determined by the embedded discrete-time Markov decision process (DTMDP, for short) of \mathcal{M} , denoted $emb(\mathcal{M})$, a tuple (S, Act, \mathbf{P}) with S and Act as before and

$$\mathbf{P}(s, \alpha, s') = \frac{\mathbf{R}(s, \alpha, s')}{E(s, \alpha)} \text{ if } E(s, \alpha) > 0$$

and 0 otherwise. Here, $E(s, \alpha) = \mathbf{R}(s, \alpha, S)$, i.e., $E(s, \alpha)$ is the exit rate of state s via some α -transition. Note that $\mathbf{P}(s, \alpha, s')$ is the time-abstract probability for the α -transition from s to s' when action α is chosen. For $B \subseteq S$ let $\mathbf{P}(s, \alpha, B) = \sum_{s' \in B} \mathbf{P}(s, \alpha, s')$ denote the probability to move from s to some state in B via an α -transition. An alternative formulation of the requirement that in every state at least one action is enabled, can be stated using E as:

$$Act(s) = \{ \alpha \in Act \mid E(s, \alpha) > 0 \} \neq \emptyset \text{ for any state } s.$$

A CTMDP (S, Act, \mathbf{R}) is *uniform* if for some $E > 0$ it holds $E(s, \alpha) = E$ for any state $s \in S$ and $\alpha \in Act(s)$. Note that $E(s, \alpha) = 0$ (whence $\alpha \notin Act(s)$ follows) is possible in uniform CTMDPs. Stated in words, in a uniform CTMDP the exit rates for all states and all enabled actions are equal.

2.2 Paths

A (timed) *path* σ in CTMDP \mathcal{M} is a finite or infinite sequence

$$\sigma \in (S \times Act \times \mathbb{R}_{>0})^* \times S \cup (S \times Act \times \mathbb{R}_{>0})^\omega.$$

For infinite path $\sigma = s_0, \alpha_0, t_0, s_1, \alpha_1, t_1, s_2, \alpha_2, t_2, \dots$ we require time-divergence, i.e., $\sum t_i = \infty$. We write

$$s_0 \xrightarrow{\alpha_0, t_0} s_1 \xrightarrow{\alpha_1, t_1} s_2 \xrightarrow{\alpha_2, t_2} \dots$$

rather than $s_0, \alpha_0, t_0, s_1, \alpha_1, t_1, s_2, \alpha_2, t_2, \dots$. The corresponding *time-abstract* path is: $s_0 \xrightarrow{\alpha_0} s_1 \xrightarrow{\alpha_1} s_2 \xrightarrow{\alpha_2} \dots$. We write $\sigma \rightarrow s$ for the time- and action-abstract path σ that is followed by state s . In the remainder of this paper we use the term path for timed, time-abstract, and time- and action-abstract paths whenever the kind of path is clear from the context. Let $first(\sigma)$ denote the state in which σ starts. For finite path σ , $last(\sigma)$ denotes the last state of σ .

2.3 Schedulers

Nondeterminism in a CTMDP is resolved by a *scheduler*. For deciding which of the next nondeterministic actions to take, a scheduler may have access to the current state only or to the path from the initial to the current state (either with or without timing information). Schedulers may select the next action either (i) *deterministically*, i.e., depending on the available information, the next action is chosen in a deterministic way, or (ii) in a *randomized* fashion, i.e., depending on the available information the next action is chosen probabilistically. Accordingly, the following classes of schedulers D are distinguished [21], where $Distr(Act)$ denotes the collection of all distributions on Act :

- stationary Markovian deterministic (SMD, also called simple schedulers),

$$D : S \rightarrow Act \text{ such that } D(s) \in Act(s);$$

- stationary Markovian randomized (SMR),

$$D : S \rightarrow Distr(Act) \text{ such that } D(s)(\alpha) > 0 \text{ implies } \alpha \in Act(s);$$

- Markovian deterministic (MD, also called step-dependent schedulers),

$$D : S \times \mathcal{I} \rightarrow Act \text{ such that } D(s, n) \in Act(s);$$

- Markovian randomized (MR),

$$D : S \times \mathcal{I} \rightarrow Distr(Act) \text{ such that } D(s, n)(\alpha) > 0 \text{ implies } \alpha \in Act(s);$$

- (time-abstract) history-dependent, deterministic (HD),

$$D : (S \times Act)^* \times S \rightarrow Act \text{ such that } D(\underbrace{s_0 \xrightarrow{\alpha_0} s_1 \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_{n-1}}}_{\text{time-abstract history}}, s_n) \in Act(s_n);$$

- (time-abstract) history-dependent, randomized (HR),

$$D : (S \times Act)^* \times S \rightarrow Distr(Act)$$

such that $D(s_0 \xrightarrow{\alpha_0} s_1 \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_{n-1}} s_n)(\alpha) > 0$ implies $\alpha \in Act(s_n)$.

All these schedulers are time-abstract; time-dependent schedulers will be briefly discussed in Section 4. We write X to denote the class of all X -schedulers over a fixed CTMDP \mathcal{M} .¹

¹ Strictly speaking we should write $X(\mathcal{M})$, but \mathcal{M} is omitted as it should be clear from the context.

Note that for any HD-scheduler, the actions in the history can be ignored, i.e., HD-schedulers may be considered as functions $D : S^+ \rightarrow Act$, as for any sequence s_0, s_1, \dots, s_n the relevant actions α_i are given by $\alpha_i = D(s_0, s_1, \dots, s_i)$. Hence, any state-action sequence $s_0 \xrightarrow{\alpha_0} s_1 \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_{n-1}} s_n$ where $\alpha_i \neq D(s_0, s_1, \dots, s_i)$ for some i , does not describe a path fragment that can be obtained from D .

The scheduler-types form a hierarchy, e.g., any SMD-scheduler can be viewed as a MD-scheduler (by ignoring parameter n) which, in turn, can be viewed as a HD-scheduler (by ignoring everything from the history except its length). A similar hierarchy exists between SMR, MR, and HR schedulers. Moreover, deterministic schedulers can be regarded as trivial versions of their corresponding randomized schedulers that assign probability 1 to the actions selected.

2.4 Induced Stochastic Process

Given a scheduler D (of arbitrary type) and a starting state, D induces a stochastic process \mathcal{M}_D on CTMDP \mathcal{M} . For deterministic schedulers (HD, MD, and SMD), the induced process \mathcal{M}_D is a continuous-time Markov chain (CTMC), referred as \mathcal{C}_D in the sequel. For MD- and HD-schedulers, though, the state space of \mathcal{C}_D will in general be infinitely large (but countable). Formally, HD-scheduler $D : S^+ \rightarrow Act$ on CTMDP $\mathcal{M} = (S, Act, \mathbf{R})$ induces the CTMC $\mathcal{C}_D = (S_D, \mathbf{R}_D)$ with $S_D = S^+$ as state space, $\mathbf{R}_D(\sigma, \sigma \rightarrow s) = \mathbf{R}(last(\sigma), D(\sigma), s)$ and 0 otherwise. The embedded discrete time Markov chain (DTMC) $emb(\mathcal{C}_D)$ is a tuple (S_D, \mathbf{P}_D) where

$$\mathbf{P}_D(\sigma, \sigma') = \frac{\mathbf{R}_D(\sigma, \sigma')}{E_D(\sigma)} \quad \text{if } E_D(\sigma) > 0$$

and 0 otherwise. Here, $E_D(\sigma) = \sum_{\sigma' \in S_D} \mathbf{R}_D(\sigma, \sigma')$, i.e., the exit rate of σ in \mathcal{C}_D .

States in CTMC \mathcal{C}_D are state sequences $s_0 \rightarrow s_1 \rightarrow \dots s_{n-1} \rightarrow s_n$ corresponding to time-abstract, unlabeled path fragments in the original CTMDP \mathcal{M} . State s_n stands for the current state in the CTMDP whereas states s_0 through s_{n-1} describe the history. Intuitively, the stochastic process induced by HD-scheduler D on CTMDP \mathcal{M} results from unfolding \mathcal{M} into an (infinite) tree while resolving the nondeterministic choices according to D . For SMD-schedulers the induced CTMC is guaranteed to be finite. More precisely, for SMD-scheduler D , \mathcal{C}_D can be viewed as a CTMC with the original state space S , as all sequences that end in s , say, are lumping equivalent [9].

2.5 Maximum Probability for Timed Reachability

Given CTMDP \mathcal{M} , our aim is to compute the maximum (or minimum) probability to reach – under a given class of schedulers – a certain set B of states within t time units, when starting from a given state s . That is, we want to calculate for time $t > 0$, $B \subseteq S$, $s \in S$ and class of X -schedulers:

$$\sup_{D \in X} \Pr_D(s, \overset{\leq t}{\rightsquigarrow} B)$$

up to some a priori given accuracy ε . Here \Pr_D denotes the induced probability measure in \mathcal{C}_D as formally defined by a standard cone construction in [4]. Intuitively, if B is considered as the set of “bad” states, then the value to be computed is the sharpest bound p for which it is guaranteed that the probability to reach a bad state from s in the next t time units is at most p under all “relevant” schedulers, i.e., all schedulers of type X .

3 An Algorithm for Uniform CTMDPs

In the sequel, unless otherwise stated, let \mathcal{M} be uniform and E be its unique exit rate. Note that CTMC \mathcal{C}_D which is obtained from the uniform CTMDP \mathcal{M} by HD-scheduler D is also uniform [19].

3.1 Approximation

For HD-scheduler D , the (infinite) vector of the probabilities $\Pr_D(\sigma, \overset{\leq t}{\rightsquigarrow} B)$ for all states σ in the CTMC \mathcal{C}_D (i.e., all $\sigma \in S^+$) is given by:

$$\left(\Pr_D(\sigma, \overset{\leq t}{\rightsquigarrow} B) \right)_{\sigma \in S^+} = \sum_{n=0}^{\infty} e^{-E \cdot t} \cdot \frac{(E \cdot t)^n}{n!} \cdot \mathbf{P}_{D,B}^n \cdot \mathbf{i}_B$$

where $\mathbf{i}_B = (\mathbf{i}_B(\sigma))_{\sigma \in S^+}$ with $\mathbf{i}_B(\sigma) = 1$ if $\text{last}(\sigma) \in B$, and 0 otherwise, and

$$\mathbf{P}_{D,B}(\sigma, \sigma') = \begin{cases} \mathbf{P}_D(\sigma, \sigma') & \text{if } \text{last}(\sigma) \in S \setminus B \\ 1 & \text{if } \sigma' = \sigma \text{ and } \text{last}(\sigma) \in B \\ 0 & \text{otherwise.} \end{cases}$$

$\mathbf{P}_{D,B}$ is the (infinite) transition probability matrix of the CTMC $\mathcal{C}_{D,B} = (S_D, \mathbf{R}_{D,B})$ that is obtained from \mathcal{C}_D by equipping any B -state (i.e., any path $\sigma \in S^+$ with $\text{last}(\sigma) \in B$) with a self-loop and removing all its other outgoing transitions: $\mathbf{R}_{D,B}(\sigma, \sigma') = \mathbf{R}_D(\sigma, \sigma')$ if $\text{last}(\sigma) \notin B$, $\mathbf{R}_{D,B}(\sigma, \sigma) = E$ if $\text{last}(\sigma) \in B$, and 0 otherwise. The justification of this transformation is as follows. As the aim is to compute the probability to reach a B -state before a certain time bound, it is not of importance what happens once such a state has been visited, and therefore its outgoing transitions can be replaced by a self-loop.

For the sake of brevity let

$$\psi(n) = e^{-E \cdot t} \cdot \frac{(E \cdot t)^n}{n!}$$

denote the Poisson probabilities, i.e., $\psi(n)$ is the probability of n events occurring within t time units in a Poisson process with rate E . Note that, for $s \in S$:

$$\Pr_D(s, \overset{\leq t}{\rightsquigarrow} B) = \left(\sum_{n=0}^{\infty} \psi(n) \cdot \mathbf{P}_{D,B}^n \cdot \mathbf{i}_B \right)(s) = \psi(0) \cdot \mathbf{i}_B(s) + \left(\sum_{n=1}^{\infty} \psi(n) \cdot \mathbf{P}_{D,B}^n \cdot \mathbf{i}_B \right)(s)$$

Later we will exploit that for $s \notin B$:

$$\Pr_D(s, \overset{\leq t}{\rightsquigarrow} B) = \left(\sum_{n=1}^{\infty} \psi(n) \cdot \mathbf{P}_{D,B}^n \cdot \mathbf{i}_B \right) (s)$$

Rather than computing the precise maximum probabilities we use an approximation in the following way. For any state s , the value $\Pr_D(s, \overset{\leq t}{\rightsquigarrow} B)$ will be approximated, up to a given accuracy ε , by

$$\widetilde{\Pr}_D(s, \overset{\leq t}{\rightsquigarrow} B) = \left(\sum_{n=0}^k \psi(n) \cdot \mathbf{P}_{D,B}^n \cdot \mathbf{i}_B \right) (s)$$

where $k = k(\varepsilon, E, t)$ depends on ε , E and t , but neither on state s nor on scheduler D . This can be seen as follows. Let $\| \cdot \|$ denote the row-sum norm. Then, for sufficiently large $k = k(\varepsilon, E, t)$:

$$\begin{aligned} & \left\| \sum_{n=0}^{\infty} \psi(n) \cdot \mathbf{P}_{D,B}^n \cdot \mathbf{i}_B - \sum_{n=0}^k \psi(n) \cdot \mathbf{P}_{D,B}^n \cdot \mathbf{i}_B \right\| = \left\| \sum_{n=k+1}^{\infty} \psi(n) \cdot \mathbf{P}_{D,B}^n \cdot \mathbf{i}_B \right\| \\ & \leq \sum_{n=k+1}^{\infty} \psi(n) \cdot \underbrace{\|\mathbf{P}_{D,B}^n\|}_{\leq 1} \cdot \underbrace{\|\mathbf{i}_B\|}_{\leq 1} \leq \sum_{n=k+1}^{\infty} \psi(n) \leq \varepsilon \end{aligned}$$

Note that $\sum_{n=0}^{\infty} e^{-E \cdot t} \frac{(E \cdot t)^n}{n!} = \sum_{n=0}^{\infty} \psi(n)$ converges for all E and t . Hence, for any scheduler D and state s :

$$\widetilde{\Pr}_D(s, \overset{\leq t}{\rightsquigarrow} B) = \left(\sum_{n=0}^k \psi(n) \cdot \mathbf{P}_{D,B}^n \cdot \mathbf{i}_B \right) (s) \leq \Pr_D(s, \overset{\leq t}{\rightsquigarrow} B)$$

Our strategy is to construct some HD-scheduler D_0 such that for any state $s \in S$:

$$\widetilde{\Pr}_{D_0}(s, \overset{\leq t}{\rightsquigarrow} B) \geq \sup_{D \in \text{HD}} \widetilde{\Pr}_D(s, \overset{\leq t}{\rightsquigarrow} B). \quad (1)$$

This yields:

$$\sup_{D \in \text{HD}} \underbrace{\Pr_D(s, \overset{\leq t}{\rightsquigarrow} B) - \varepsilon}_{\leq \widetilde{\Pr}_D(s, \overset{\leq t}{\rightsquigarrow} B)} \leq \widetilde{\Pr}_{D_0}(s, \overset{\leq t}{\rightsquigarrow} B) \leq \Pr_{D_0}(s, \overset{\leq t}{\rightsquigarrow} B) \leq \sup_{D \in \text{HD}} \Pr_D(s, \overset{\leq t}{\rightsquigarrow} B).$$

Since $\mathbf{P}_{D,B}^n(s, \sigma) = 0$ for any σ containing more than n transitions, i.e., more than $n+1$ states, the value

$$\widetilde{\Pr}_{D_0}(s, \overset{\leq t}{\rightsquigarrow} B) = \left(\sum_{n=0}^k \psi(n) \cdot \mathbf{P}_{D_0,B}^n \cdot \mathbf{i}_B \right) (s)$$

only depends on the k -th truncation of D_0 , i.e., the function

$$D_0 \Big|_k : \bigcup_{0 < n \leq k} S^n \rightarrow \text{Act}, \quad D_0 \Big|_k(\sigma) = D_0(\sigma).$$

Intuitively speaking, only the first k decisions of D_0 are relevant (and not “later” ones) for determining the value $\widetilde{\text{Pr}}_{D_0}(s, \overset{\leq t}{\rightsquigarrow} B)$. There are only finitely many such truncations when ranging over all HD-schedulers. The brute-force approach would consider all of them in order to determine the maximum. This technique works, but is highly inefficient because the total number of such truncations, $\prod_{s \in S} |Act(s)|^k$, grows exponentially in the number of states s with $|Act(s)| > 1$. Note that

$$\prod_{s \in S} |Act(s)|^k \geq 2^{|T|k} \quad \text{if } |Act(s)| \geq 2 \text{ for all } s \in T \subseteq S$$

i.e., the total number of truncations to be considered is exponential in k .

3.2 A Greedy Algorithm to Compute Scheduler D_0

Consider truncated MD-schedulers of the form $D : S \times \{1, \dots, k\} \rightarrow Act$. (Later on, it is shown that considering such schedulers suffices.)

The actions $act(s, i) \in Act(s)$ for $0 < i \leq k$ will be determined such that the truncated MD-scheduler D_0 with $D_0(s, i) = act(s, i)$ fulfils equation (1). Let \mathbf{P}_i denote the probability matrix of cardinality $|S| \times |S|$ where the row $\mathbf{P}_i(s, \cdot) = \mathbf{P}(s, act(s, i), \cdot)$ if $s \notin B$ and $\mathbf{P}_i(s, \cdot) = \mathbf{i}_s$ if $s \in B$, where \mathbf{i}_s denotes the identity vector for state s . \mathbf{P}_i thus denotes the probability matrix induced by the scheduler D_0 at step i .

For $s \notin B$, the actions $act(s, i)$ will be determined in a backward manner, i.e., starting from $i=k$. For $i=k$, the selected action $act(s, k) \in Act(s)$ satisfies:

$$\mathbf{P}_k(s, B) = \mathbf{P}(s, act(s, k), B) = \max_{\alpha \in Act(s)} \mathbf{P}(s, \alpha, B)$$

That is, $\mathbf{P}_k(s, \cdot)$ is determined such that for any state s the probability to move to a B -state within at most one step is maximized. Generalizing this strategy, for $i < k$, $act(s, i)$ is chosen such that the probability to move to a B -state within at most $k-i+1$ steps is maximal under the truncated MD-scheduler $D : S \times \{i-1, \dots, k\} \rightarrow Act$ defined by:

$$D(s, j) = act(s, i+j-1), \text{ for } 0 < j \leq k-i+1.$$

That is, \mathbf{P}_i is constructed such that for $i \geq 1$ the vector

$$q_i = \sum_{n=i}^k \psi(n) \cdot \mathbf{P}_i \cdot \mathbf{P}_{i+1} \cdot \dots \cdot \mathbf{P}_n \cdot \mathbf{i}_B$$

is state-wise maximal under all vectors of the form

$$\sum_{n=i}^k \psi(n) \cdot \mathbf{P}_* \cdot \mathbf{P}_{i+1} \cdot \dots \cdot \mathbf{P}_n \cdot \mathbf{i}_B$$

where \mathbf{P}_* is an $|S| \times |S|$ -matrix with $\mathbf{P}_*(s, \cdot) = \mathbf{P}(s, \alpha, \cdot)$ for some action $\alpha \in Act(s)$ if $s \notin B$ and $\mathbf{P}_*(s, \cdot) = \mathbf{i}_s$ if $s \in B$. In the above equations, $\mathbf{i}_B =$

$(i_B(s))_{s \in S}$ stands for the bit-vector that represents the characteristic function of B (as a subset of the original state space S). I.e., $i_B(s) = 1$ if $s \in B$ and $i_B(s) = 0$ if $s \in S \setminus B$.²

Informally, $q_i(s)$ is the maximum conditional probability to reach B taking i to k steps within t time units, given that state s is occupied before the i -th step. We let $q = \psi(0) \cdot \mathbf{i}_B + q_1$, which for the $(S \setminus B)$ -states agrees with the desired probability vector to reach a B -state within at most k steps when the time bound to reach B is t . For $s \in B$ we have $\Pr_D(s, \overset{\leq t}{\rightsquigarrow} B) = 1$. Moreover, for $s \notin B$ it holds $q(s) = \psi(0) \cdot \mathbf{i}_B(s) + q_1(s) = q_1(s)$ (as $\mathbf{i}_B(s) = 0$). In the sequel, we are therefore only interested in the calculation of the vector $q_1(s)$.

The main steps of our procedure are summarized in Algorithm 1. A stable and efficient algorithm to compute the Poisson probabilities $\psi(i)$ has been proposed in [15] and can be adopted here. Note that for the computation of the values $\sup_{D \in HD} \widetilde{\Pr}_D(s, \overset{\leq t}{\rightsquigarrow} B)$ there is no need to compute (and store) the matrices \mathbf{P}_i . Instead, it suffices to compute the vectors

$$\begin{aligned} q_i &= \sum_{n=i}^k \psi(n) \cdot \mathbf{P}_i \cdot \mathbf{P}_{i+1} \cdot \dots \cdot \mathbf{P}_n \cdot \mathbf{i}_B \\ &= \psi(i) \cdot \mathbf{P}_i \cdot \mathbf{i}_B + \sum_{n=i+1}^k \psi(n) \cdot \mathbf{P}_i \cdot \mathbf{P}_{i+1} \cdot \dots \cdot \mathbf{P}_n \cdot \mathbf{i}_B \\ &= \psi(i) \cdot \mathbf{P}_i \cdot \mathbf{i}_B + \mathbf{P}_i \cdot \sum_{n=i+1}^k \psi(n) \cdot \mathbf{P}_{i+1} \cdot \dots \cdot \mathbf{P}_n \cdot \mathbf{i}_B \\ &= \psi(i) \cdot \mathbf{P}_i \cdot \mathbf{i}_B + \mathbf{P}_i \cdot q_{i+1} \end{aligned}$$

where $q_{k+1} = \mathbf{0}$ is the 0-vector. For $s \notin B$, we have $(\mathbf{P}_i \cdot \mathbf{i}_B)(s) = \mathbf{P}(s, \alpha, B)$ if $\alpha = \text{act}(s, i)$.

3.3 Complexity of the Algorithm

Algorithm 1 can be implemented with a space complexity in $\mathcal{O}(|S|^2 \cdot |\text{Act}| + |S|)$, where the term $|S|^2 \cdot |\text{Act}|$ stands for the representation of the uniform CTMDP \mathcal{M} while the term $|S|$ stands for the vectors q_{i+1} and q_i . Note that there is no need to store q_{i+1} once q_i has been computed. The values $q_i(s, \alpha)$ are only needed temporarily, and as mentioned before, there is no need to compute and store the matrices \mathbf{P}_i . Inspection of the pseudo-code of Algorithm 1 reveals that the worst-case time complexity is asymptotically bounded by:

$$k \cdot \sum_{s \in S \setminus B} \sum_{\alpha \in \text{Act}(s)} |\{s' \in S \mid \mathbf{R}(s, \alpha, s') > 0\}|$$

² At several other places, we shall use the same notation \mathbf{i}_B for the bit-vector $(i_B(\sigma))_{\sigma \in S^+}$ that represents the characteristic function of B viewed as subset of the state space of the CTMC induced by a HD-scheduler. Here, we identify B with the set of finite paths σ where $\text{last}(\sigma) \in B$. Whenever the notation \mathbf{i}_B occurs in our formulae the dimension of \mathbf{i}_B should be clear from the context.

Algorithm 1 Greedy approximation algorithm for computing $\sup_{D \in HD} \Pr_D(s, \overset{\leq t}{\rightsquigarrow} B)$

```

k :=  $k(\varepsilon, E, t)$ ; (* determine number of required steps *)
for all  $s \in S$  do  $q_{k+1}(s) := 0$ ; od (* initialize  $q_{k+1}$  as null-vector *)
for all  $i = k, k-1, \dots, 1$  do
  for all  $s \in S \setminus B$  do
     $m := -1$ ;
    for all  $\alpha \in Act(s)$  do (* search the optimal row  $\mathbf{P}_i(s, \cdot)$  *)
       $m := \max \left( m, \psi(i) \cdot \mathbf{P}(s, \alpha, B) + \sum_{s' \in S} \mathbf{P}(s, \alpha, s') \cdot q_{i+1}(s') \right)$ ;
    od
     $q_i(s) := m$ ; (* choose maximum *)
  od
  for all  $s \in B$  do  $q_i(s) := \psi(i) + q_{i+1}(s)$ ; od (*  $\mathbf{P}_i(s, \cdot) := \mathbf{i}_s$  for all states  $s \in B$  *)
od
for all  $s \in S$  do
  if  $s \notin B$  then  $q(s) := q_1(s)$ ; else  $q(s) := 1$ ; fi
od
return the vector  $q$ .

```

which is in $\mathcal{O}(E \cdot t \cdot |S|^2 \cdot |Act|)$. Note that $k = k(\varepsilon, E, t)$ grows proportionally with $E \cdot t$. This bound can be improved by performing a reachability analysis (as a preprocessing phase of Algorithm 1) to determine the set T of states from which a B -state can be reached. The main iteration then only needs to be performed for all states in $T \setminus B$ rather than $S \setminus B$. For the other states we have, for any scheduler D , $\Pr_D(s, \overset{\leq t}{\rightsquigarrow} B) = 0$ for $s \in S \setminus T$, and $\Pr_D(s, \overset{\leq t}{\rightsquigarrow} B) = 1$ for $s \in B$.

3.4 Correctness of the Algorithm

Although our greedy algorithm is based on a truncated MD-scheduler – only the first k steps are memorized – it approximates the maximum probability to reach the set of states B within t time units under *all* HD-schedulers. This is shown by the following result where $q(s)$ is the s -component of the vector q as returned by Algorithm 1.

Theorem 1. $\sup_{D \in HD} \Pr_D(s, \overset{\leq t}{\rightsquigarrow} B) - \varepsilon \leq q(s) \leq \sup_{D \in HD} \Pr_D(s, \overset{\leq t}{\rightsquigarrow} B)$.

As a result, the vector computed by Algorithm 1 is state-wise optimal under all HD-schedulers, up to the accuracy ε .

4 Other Scheduling Disciplines

By Theorem 1 it follows that our greedy algorithm computes the maximum probability for timed reachability under all HD-schedulers. In this section, we show that this also applies to any MR-, MD-, and, more importantly, to any

HR-scheduler. In addition, we will show that this does neither hold for SMD-schedulers nor for schedulers that can base their decision on the timing of actions. Finally, it is shown that adding a simple notion of fairness is invariant under these maximum probabilities for HD-schedulers.

Markovian deterministic schedulers. In the sequel, let $s \in S$ be a state, $t \geq 0$ a time point and $B \subseteq S$ a set of states. Theorem 1 states that the vector computed by Algorithm 1 is state-wise optimal under all HD-schedulers, up to a given accuracy ε . As Algorithm 1 calculates, in fact, a truncation of an MD-scheduler it directly follows that the suprema under MD- and HD-schedulers agree:

$$\textbf{Theorem 2.} \quad \sup_{D \in MD} \Pr_D(s, \overset{\leq t}{\rightsquigarrow} B) = \sup_{D \in HD} \Pr_D(s, \overset{\leq t}{\rightsquigarrow} B).$$

History-dependent randomized schedulers. The next results yields that the supremum under HD- and HR-schedulers coincides:

$$\textbf{Theorem 3.} \quad \sup_{D \in HD} \Pr_D(s, \overset{\leq t}{\rightsquigarrow} B) = \sup_{D \in HR} \Pr_D(s, \overset{\leq t}{\rightsquigarrow} B).$$

A few remarks are in order. Theorems 2 and 3 show that the suprema for the probabilities to reach a set of goal states within a given time bound under the classes of scheduler MD, HD, MR and HR coincide. (For MR-schedulers this stems from the fact that $MD \subseteq MR \subseteq HR$.) For probabilities of some other events, however, such correspondence may not be established. That is, in general, randomized schedulers can be better than deterministic schedulers. This observation was made by Beutler and Ross [7] who showed that the maximum of time-average rewards under randomized schedulers might be larger than under deterministic schedulers. In fact, the crux of the proof of Theorem 3 is the observation that the values $\Pr_D(s, \overset{\leq t}{\rightsquigarrow}_{\leq n} B)$ converge to $\Pr_D(s, \overset{\leq t}{\rightsquigarrow} B)$, where the subscript $\leq n$ denotes that B has to be reached within at most n steps. This property is not guaranteed for other events.

Stationary Markovian deterministic schedulers. Different from the discrete time setting, where SMD-schedulers suffice for maximum probabilities to reach a set of goal states within a given number of steps [8,5], this does not hold for the corresponding question – interpreting the number of steps in the discrete case as elapse of time – on CTMDPs. A counterexample is given in Fig. 1(a). Here, states are represented as circles and there is an edge between states s and s' labeled with action α if and only if $\mathbf{R}(s, \alpha, s') > 0$. Action labels and rates are indicated at each edge. Let $B = \{s_2\}$, and consider the only two relevant SMD-schedulers, D_α , selecting action α in state s_0 , and D_β , selecting action β . Comparing them with $D_{\beta\alpha}$, i.e., the scheduler that after selecting β once switches to selecting α in state s_0 , we find that for a certain range of time bounds t , $D_{\beta\alpha}$ outperforms both D_β and D_α . Intuitively, the chance of stuttering in state s_0 (by choosing β initially) may influence the remaining time to reach B to an extent that it becomes profitable to continue choosing α . For $t = 0.5$, for instance, $\Pr_{D_{\beta\alpha}}(s_0, \overset{\leq 0.5}{\rightsquigarrow} B) = 0.4152$, whereas for D_α and D_β these probabilities




Fig. 1. Uniform CTMDPs where (a) SMD-schedulers are less powerful, and (b) where THD schedulers are more powerful than HD-schedulers.

are 0.3935 and 0.3996, respectively. Thus, SMD-schedulers are not expressive enough for maximum probabilities to reach a set of goal states within a given time bound under all HD/HR-schedulers.³

Timed schedulers. This paper only considers schedulers that do not take the timing information into account. It is however worth noticing that timed history-dependent (THD) schedulers are more powerful than time-abstract history dependent schedulers (class HD and HR), in the sense that it is possible that:

$$\sup_{D \in THD} \Pr_D(s, \overset{\leq t}{\rightsquigarrow} B) > \sup_{D \in HD} \Pr_D(s, \overset{\leq t}{\rightsquigarrow} B).$$

Here, *THD* refers to the class of schedulers given by functions $D : (S \times Act \times \mathbb{R}_{>0})^* \times S \rightarrow Act$ (only choosing from $Act(s)$ for any path ending in state s), i.e., THD-schedulers are able to observe the time points of state changes. To see that they may yield a higher probability, consider for example the uniform CTMDP in Fig. 1(b). In this example, it depends on the time instance of entering s_1 whether it is more profitable to continue choosing α or β . To be more precise, consider the only relevant HD-schedulers, D_α (choosing α in s_1) and D_β (choosing β). Fig. 2 plots the probability to reach B starting from state s_1 if choosing D_α , respectively D_β , given by

$$\Pr_{D_\alpha}(s_1, \overset{\leq t}{\rightsquigarrow} B) = 1 - e^{-t}, \quad \text{and} \quad \Pr_{D_\beta}(s_1, \overset{\leq t}{\rightsquigarrow} B) = 1 - e^{-2t} \cdot (1 + 2t).$$

Let t_0 be the time instance satisfying $e^{t_0} = 1 + 2t_0$, i.e., the time point where both plots cross. The THD-scheduler D defined by $D(s_0 \xrightarrow{\gamma, u} s_1) = \alpha$ if $t - u < t_0$ and β otherwise, maximizes the probability to reach $B = \{s_3\}$ from state s_0 within t time units, and obviously outperforms both D_α and D_β .

Fairness. We conclude this section by considering a simple notion of fairness for schedulers. Let $\sigma = s_0 \xrightarrow{\alpha_0, t_0} s_1 \xrightarrow{\alpha_1, t_1} \dots$ be an infinite path. Infinite path

³ For SMR-schedulers this is an open issue.




Fig. 2. Functions $1 - e^{-t}$ and $1 - e^{-2t} \cdot (1 + 2t)$ for $t \geq 0$

σ is called *fair* if and only if for each state s that occurs infinitely often in σ and each action $\alpha \in Act(s)$, there are infinitely many indices n such that $(s_n, \alpha_n) = (s, \alpha)$. Stated in words, for any state that is visited infinitely often, each of its outgoing actions cannot have been selected only a finite number of times. (Note that this notion of fairness is rather weak; for instance, a scheduler that finitely many times selects the same action in a state that is visited finitely often – without ever considering one of the other possibilities – is considered to be fair.) Scheduler D (of some class) is called fair if and only if

$$\Pr_D \{ \sigma \in Path(s) \mid \sigma \text{ is fair} \} = 1$$

for all states $s \in S$, where $Path(s)$ denotes the set of paths that start in s . Let FHD denote the set of all fair HD-schedulers. The following result states that maximum probabilities under HD-schedulers and their fair counterparts coincide:

Theorem 4. $\sup_{D \in HD} \Pr_D(s, \overset{\leq t}{\rightsquigarrow} B) = \sup_{D \in FHD} \Pr_D(s, \overset{\leq t}{\rightsquigarrow} B)$.

5 The Uniformization Problem

Algorithm 1 assumes that the CTMDP under consideration is uniform. We now discuss the case in which the CTMDP is not uniform, i.e., the exit rates $E(s, \alpha)$ are not guaranteed to be identical for any state s and any $\alpha \in Act(s)$.

In the setting of CTMCs, uniformization [19] can be employed to transform a CTMC into a uniform one while keeping transient probabilities invariant⁴. For CTMDPs, a similar recipe might be followed. However, a simple adaptation of the uniformization approach for CTMCs (as proposed, for instance, in [6,21]) to CTMDPs is not adequate for our purpose. The problem with this approach is that the correspondence between schedulers on uniform CTMDP \mathcal{M}' and its

⁴ Although uniformization is traditionally considered as a transformation from a CTMC into a discrete-time Markov chain, it can equally well be considered as a transformation between CTMCs.



Fig. 3. An example illustrating why uniformization on CTMDPs is not obvious

original CTMDP \mathcal{M} is lost. (A similar observation has been made by Beutler and Ross [7] when comparing MD- and MR-schedulers for computing time-average rewards.) This can be illustrated as follows. Applying “standard” uniformization to CTMDP $\mathcal{M} = (S, Act, \mathbf{R})$ with $E \geq \max_{s \in S, \alpha \in Act} E(s, \alpha)$ would yield the CTMDP $\mathcal{M}' = (S, Act, \mathbf{R}')$ with $\mathbf{R}'(s, \alpha, s') = \mathbf{R}(s, \alpha, s')$ for $s \neq s'$, and $\mathbf{R}'(s, \alpha, s) = \mathbf{R}(s, \alpha, s) + E - E(s, \alpha)$ if $\alpha \in Act(s)$ and 0 otherwise. That is, each state s is equipped with a self-loop for each action $\alpha \in Act(s)$ if E exceeds the total exit rate to take an α -transition from s . Applying this recipe to the CTMDP \mathcal{M} depicted in Fig. 3(a) results in the CTMDP \mathcal{M}' in Fig. 3(b). The latter has appeared in Fig. 1(a) already. It is not difficult to see that for any X -scheduler on \mathcal{M} there exists a corresponding X -scheduler on \mathcal{M}' , as any choice in \mathcal{M} can be matched by the same choice in \mathcal{M}' . The reverse, however, does not hold. For instance, the MD-scheduler $D_{\beta\alpha}$ on \mathcal{M}' discussed in Section 4 does not correspond to any MD-scheduler D on \mathcal{M} , since the self-loop in state s_0 in \mathcal{M}' cannot be mimicked by \mathcal{M} . Recall from Section 4 that $\Pr_{D_{\beta\alpha}}(s_0, \overset{\leq 0.5}{\rightsquigarrow} \{s_2\})$ is higher than the respective probabilities for D_α and D_β in \mathcal{M}' . The latter in turn correspond to the only relevant HD-scheduler on \mathcal{M} . As a consequence, the maximum probability (obtained for some MD-scheduler generated by Algorithm 1) to reach the set $\{s_2\}$ from state s_0 in 0.5 time units on \mathcal{M}' is higher than the probability for any HD-scheduler in \mathcal{M} .

We are experimenting with other forms of uniformization to overcome this problem. As yet, it is open whether a variation of the basic concept of uniformization can be used to reduce the timed reachability problem for general CTMDPs to that of uniform CTMDPs.

6 Concluding Remarks

This paper considered the problem of computing the maximum probability to reach a set of goal states within a given time bound in a uniform CTMDP. It is shown that truncated Markovian deterministic schedulers suffice for approximating a solution to this problem in an efficient manner for (time-abstract) history-dependent and Markovian schedulers, both deterministic and randomized ones. This does neither apply to timed history-dependent schedulers nor

to Markovian stationary (i.e., simple) schedulers. The question whether SMR-schedulers may yield the same optimum (or a smaller optimum) is open.

Although all results in this paper have been presented for maximum probabilities, the same results can be obtained for minimal probabilities, i.e.,

$$\inf_{D \in X} \Pr_D(s, \overset{\leq t}{\rightsquigarrow} B)$$

up to some accuracy ε .⁵ Instead of a greedy policy that maximizes the probability to reach the set of goal states in each step of the computation, the algorithm in this case minimizes this quantity in each step.

The presented numerical algorithm is remarkably efficient. Its worst-case time complexity is in $\mathcal{O}(E \cdot t \cdot N^2 \cdot M)$ where E is the unique exit rate of the uniform CTMDP, t is the time bound, N is the number of states, and M is the number of actions. Thus, compared to CTMCs, the increase in computational effort is linear in the number of actions in the CTMDP, i.e., the amount of nondeterminism, but no more than that. This is the best we can hope for, since the time complexity of computing the corresponding probability in a CTMC is in $\mathcal{O}(E \cdot t \cdot N^2)$ [4].

Acknowledgments. The authors thank Luca de Alfaro (University of California at Santa Cruz) for insightful discussions. David Jansen (MPI Saarbrücken) provided various helpful comments on a draft of this paper.

References

1. M. Ajmone Marsan, G. Balbo, G. Conte, S. Donatelli, and G. Franceschinis. *Modelling with Generalized Stochastic Petri Nets*. John Wiley & Sons, 1995.
2. A. Aziz, K. Sanwal, V. Singhal and R. Brayton. Model checking continuous time Markov chains. *ACM TOCL*, 1(1): 162–170, 2000.
3. C. Baier, B. Haverkort, H. Hermanns, and J.-P. Katoen. Efficient computation of time-bounded reachability probabilities in uniform continuous-time Markov decision processes. CTIT Technical Report 03-50, University of Twente, 2003.
4. C. Baier, B. Haverkort, H. Hermanns, and J.-P. Katoen. Model-checking algorithms for continuous-time Markov chains. *IEEE Transactions on Software Engineering*, 29(6): 524–541, 2003.
5. C. Baier and M.Z. Kwiatkowska. Model checking for a probabilistic branching time logic with fairness. *Distributed Computing*, 11: 125–155, 1998.
6. D.P. Bertsekas. *Dynamic Programming and Optimal Control. Volume II*. Athena Scientific, 1995.
7. F. Beutler and K.W. Ross. Uniformization for semi-Markov decision processes under stationary policies. *Journal of Applied Probability*, 24: 644–656, 1987.
8. A. Bianco and L. de Alfaro. Model checking of probabilistic and nondeterministic systems. In *Foundations of Software Technology and Theoretical Computer Science*, LNCS 1026: 499–513, 1995.
9. P. Buchholz. Exact and ordinary lumpability in finite Markov chains. *Journal of Applied Probability*, 31: 59–75, 1994.

⁵ Only Theorem 4 turns wrong when the supremum over all fair schedulers is replaced by the infimum over all fair schedulers. See [5] for a counterexample for DTMDPs.

10. P. Buchholz, J.-P. Katoen, P. Kemper, and C. Tepper. Model-checking large structured Markov chains. *Journal of Logic and Algebraic Programming*, 56:69–97, 2003.
11. G. Chiola, S. Donatelli, and G. Franceschinis. GSPNs versus SPNs: What is the actual role of immediate transitions? In *Petri Nets and Performance Models 1991*, IEEE CS Press, 1991.
12. G. Ciardo and R. Zijal. Well-defined stochastic Petri nets. In *Proc. 4th International Workshop on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS'96)*, pages 278–284, 1996.
13. D.D. Deavours, and W.H. Sanders. An efficient well-specified check. In *Petri Nets and Performance Models*, IEEE CS Press, 1999.
14. E.A. Feinberg. Continuous time discounted jump Markov decision processes: a discrete-event approach. 1998.
<http://www.ams.sunysb.edu/~feinberg/public/>.
15. B.L. Fox and P.W. Glynn. Computing Poisson probabilities. *Communications of the ACM*, 31(4): 440–445, 1988.
16. H. Hermans, U. Herzog, and V. Mertsiotakis. Stochastic process algebras – Between LOTOS and Markov chains. *Computer Networks and ISDN Systems*, 30(9–10):901–924, 1998.
17. H. Hermans. *Interactive Markov Chains and the Quest for Quantified Quality*. LNCS 2428, Springer, 2002.
18. H. Hermans, J.-P. Katoen, J. Meyer-Kayser, and M. Siegle. A tool for model-checking Markov chains. *Int. Journal on Software Tools for Technology Transfer*, 4(2):153–172, 2003.
19. A. Jensen. Markov chains as an aid in the study of Markov processes. *Skand. Aktuarietidskrift*, 3: 87–91, 1953.
20. J.-P. Katoen, M.Z. Kwiatkowska, G. Norman and D. Parker. Faster and symbolic CTMC model checking. In *Process Algebra and Probabilistic Methods*, LNCS 2165: 23–38, Springer, 2001.
21. M.L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, 1994.
22. Q. Qiu and M. Pedram. Dynamic power management based on continuous-time Markov decision processes. In *Design Automation Conference 99*, pp. 555–561, 1999.
23. W.H. Sanders and J.F. Meyer. Reduced base model construction methods for stochastic activity networks. *IEEE Journal on Selected Areas in Communications*, 9(1):25–36, January 1991.
24. L. Sennot. *Stochastic Dynamic Programming and the Control of Queueing Systems*. John Wiley & Sons, 1999.