Available online at www.sciencedirect.com

SCIENCE DIRECT®

Theoretical Computer Science 345 (2005) 2–26

Theoretical Computer Science

www.elsevier.com/locate/tcs

Efficient computation of time-bounded reachability probabilities in uniform continuous-time Markov decision processes☆

Christel Baier[a], Holger Hermanns[b, c, *], Joost-Pieter Katoen[c], Boudewijn R. Haverkort[c]

[a] *University of Bonn, Römerstraße 164, D-53117 Bonn, Germany*
[b] *Saarland University, Im Stadtwald, D-66123 Saarbrücken, Germany*
[c] *University of Twente, P.O. Box 217, 7500 AE Enschede, The Netherlands*

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. It furthermore proves that these probabilities coincide for (time-abstract) history-dependent and Markovian schedulers that resolve nondeterminism either deterministically or in a randomized way.

Keywords: Continuous-time; Markov decision process; Temporal logic; Model checking; Time-bounded reachability

1. Introduction

A continuous-time Markov decision process (CTMDP) [10,20, 31,34] is a generalization of a continuous-time Markov chain (CTMC) in which both probabilistic and

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

* Corresponding author. Saarland University, Department of Computer Science, IM Stadtwald, D-66123 Saarbrücken, Germany. Tel.: +49 681 302 5631; fax: +49 681 302 5636.

E-mail address: hermanns@cs.uni-sb.de (H. Hermanns).

0304-3975/$ - see front matter © 2005 Published by Elsevier B.V.
doi:10.1016/j.tcs.2005.07.022

nondeterministic choices co-exist. CTMDPs are a natural modeling formalism applicable in many contexts, ranging from stochastic control theory [20] and scheduling [13,1] to dynamic power management [32].

Importance of CTMDPs: The class of CTMDPs is particularly interesting, because it can be viewed as a common semantic model for various performance and dependability modelling formalisms including generalized stochastic Petri nets [2], Markovian stochastic activity networks [33], and interactive Markov chains (IMC) [23]. 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 [16,17,19,23]. All these notions are semantic notions. They are usually checked by an exhaustive exploration of the state space associated with a given model. A model is discarded if the check fails. In other words, no specification-level check is available, and the offered analysis algorithms are actually partial algorithms.

Model checking: Model checking of CTMCs [6] has received remarkable attention in recent years. Various model checkers exist [25,27,15], 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 [5]—to time-bounded reachability problems [6], which in turn can be approximated efficiently using a stable numerical technique called uniformization [26]. 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. With the notable exception of De Alfaro [3,4], who studied long-run properties of semi-Markov decision processes, we are not aware of any model checking algorithm for CTMDPs. This stands in sharp contrast to discrete-time Markov decision processes, for which model checking algorithms are well-understood [12,9] and, for instance, implemented in tools like PRISM [30] or RAPTURE [18].

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 [12,9]—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 probabilities for timed reachability in uniform CTMDPs under *all* time-abstract schedulers. The time complexity is in $\mathcal{O}(t \cdot E \cdot N^2 \cdot M)$ and the space complexity in $\mathcal{O}(N^2 \cdot M)$ where t is the time bound, E is the uniform exit rate of the CTMDP under