Science of Computer Programming 112 (2015) 58–74

Contents lists available at ScienceDirect

Science of Computer Programming

www.elsevier.com/locate/scico

Improving time bounded reachability computations in interactive Markov chains

Hassan Hatefi *, Holger Hermanns

Saarland University, Department of Computer Science, Campus Saarbrücken, D-66123 Germany

ARTICLE INFO

Article history:
Received 9 June 2014
Received in revised form 16 March 2015
Accepted 14 May 2015
Available online 15 June 2015

Keywords:
Interactive Markov chains
Time bounded reachability
Discretisation
Root finding

ABSTRACT

Interactive Markov Chains (IMCs) are compositional behaviour models extending both Continuous Time Markov Chain (CTMC) and Labelled Transition System (LTS). They are used as semantic models in different engineering contexts ranging from ultramodern satellite designs to industrial system-on-chip manufacturing. Different approximation algorithms have been proposed for model checking of IMCs, with time bounded reachability probabilities playing a pivotal role. This paper addresses the accuracy and efficiency of approximating time bounded reachability probabilities in IMCs, improving over the state-of-the-art in both efficiency of computation and tightness of approximation. Moreover, a stable numerical approach, which provides an effective framework for implementation of the theory, is proposed. Experimental evidence demonstrates the efficiency of the new approach.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

Why IMCs? Over the last decade, a formal approach to quantitative performance and dependability evaluation of concurrent systems has gained maturity. At its root are continuous-time Markov chains for which efficient and quantifiably precise solution methods exist [1]. On the specification side, continuous stochastic logic (CSL) [2,1] enables the specification of a large spectrum of performance and dependability measures. A CTMC can be viewed as a labelled transition system (LTS) whose transitions are delayed according to exponential distributions. Opposed to classical concurrency theory models, CTMCs neither support compositional modelling [3] nor do they allow nondeterminism in the model. Among several formalisms that overcome these limitations [4–7], interactive Markov chains (IMCs) [8] stand out. IMCs conservatively extend classical concurrency theory with exponentially distributed delays, and this induces several further benefits [9]. In particular, it enables compositional modelling with intermittent weak bisimulation minimisation [5] and allows to augment existing untimed process algebra specifications with random timing [4]. Moreover, the IMC formalism is not restricted to exponential delays but allows to encode arbitrary phase-type distributions such as hyper- and hypoexponentials [10]. Since IMCs smoothly extend classical LTSs, the model has received attention in academic as well as in industrial settings [11–14].

Why time bounded reachability? The principles of model checking IMCs are by now well understood. One analysis strand, implemented for instance in CADP [15], resorts to CSL model checking of CTMCs. But this is only applicable if the weak bisimulation quotient of the model is indeed a CTMC, which cannot be guaranteed always. This is therefore a partial solution

* Corresponding author.
E-mail addresses: hhatefi@cs.uni-saarland.de (H. Hatefi), hermanns@cs.uni-saarland.de (H. Hermanns).

http://dx.doi.org/10.1016/j.scico.2015.05.003
0167-6423/© 2015 Elsevier B.V. All rights reserved.

technique, albeit it integrates well with compositional construction and minimisation approaches, and is the one used in industrial applications.

The time bounded reachability problem has been successfully addressed [16–18] for Continuous-Time Markov Decision Processes (CTMDPs) with respect to the class of so-called *late* schedulers. These schedulers can change their decision while residing in a state. CTMDPs are very closely related to IMCs, but are not compositional. In this paper we however focus on *early* schedulers which fix a decision upon arriving at a state and are unable to change that decision afterwards. Early schedulers arise naturally from compositional model construction using IMCs.[1] Notably, the adaptation of optimal time bounded reachability computation for CTMDPs under late schedulers to IMCs under early schedulers has not been studied. Therefore the existing techniques for CTMDPs are not directly applicable to IMCs.

The approximate CSL model checking problem for IMCs has been solved by Neuhäusser and Zhang [19,20]. Most of the solutions resort to untimed model checking [21]. The core innovation lies in the solution of the time bounded model checking problem, that is needed to quantify a *bounded until formula* subject to a (real-valued) time interval. The problem is solved by splitting the time interval into equally sized discretisation steps, each small enough to carry at most one Markov transition with high probability. However, the practical efficiency and accuracy of this approach to evaluate time bounded reachability probabilities turns out substantially inferior to the one known for CTMCs, and this limits its applicability to real industrial cases. As a consequence, model checking algorithms for other, less precise, but still highly relevant properties have been coined [22], including expected reachability and long run average properties.

Our contribution We revisit the approximation of time bounded reachability probabilities so as to arrive at an improved computational approach. For this, we generalise the discretisation approach of Neuhäusser and Zhang [19,20] by considering the effect of multiple Markov transition firings in each discretisation step. We show that this can be exploited by a tighter error bound, and thus a more accurate computation. We put the theoretical improvement into practice by proposing a stable numerical approach, which provides ϵ-optimal time bounded reachability together with its corresponding scheduler (choice decisions in the model) for IMCs. Empirical results demonstrate that the improved algorithm can gain more than one order of magnitude speedup.

Organisation of the paper. In Section 2 we define IMC and provide the notations and concepts that we use throughout the paper. Section 3 explains the state-of-the-art method in time bounded reachability computations of IMCs. We introduce the theoretical foundation of our improvement in Section 4. Section 5 develops a stable numerical approach with strict error bound for the implementation. Section 6 reports on an empirical evaluation of our method applied to an industrial case study. Finally, Section 7 concludes the paper.

This paper is an extended version of a conference paper that appeared in FSEN 2013 [23]. All proofs are included as Appendices A–D. Section 5 contains novel contributions on a stable numerical approximation scheme, and this implies that the experimental results presented in Section 6 are revised in their entirety.

2. Interactive Markov chain

An Interactive Markov Chain (IMC) is a model that generalises both CTMC and LTS. In this section, we provide the definition of IMC and the necessary concepts relating to it.

Definition 1 *(IMC).* An IMC [5] is a tuple $\mathcal{M} = (S, Act, \longrightarrow, \dashrightarrow, s_0)$, where S is a finite set of states, Act is a set of actions, including internal invisible action τ, $\longrightarrow \subset S \times Act \times S$ is a set of interactive transitions, $\dashrightarrow \subset S \times \mathbb{R}_{\geq 0} \times S$ is a set of Markov transitions, and s_0 is the initial state.

Maximum progress vs. urgency States of an IMC are partitioned into *interactive*, *Markov* and *hybrid*. Interactive (Markov) states have only interactive (Markov) outgoing transitions, while hybrid states have transitions of both types. Let S_I, S_M and S_H be the set of interactive, Markov and hybrid states, respectively. An IMC might have states without any outgoing transition. For the purpose of this paper, any such state is turned into a Markov state by adding a self loop with an arbitrary rate.

Depending on whether or not they can interact with the environment, IMCs are classified into closed and open models. An open IMC can interact with the environment and in particular, can be composed with other IMCs, e.g. via parallel composition. For such models, a *maximal progress* assumption [5] is imposed which implies that τ-transitions take precedence over Markov transitions whenever both are enabled in a state. In contrast, a closed IMC is not subject to any further communication and composition. In this paper, we assume that the models we are going to analyse are closed, and impose the stronger *urgency* assumption which means that any interactive transition has precedence over Markov transitions. In other words, interactive transitions are taken immediately whenever enabled in a state, leaving no chance for enabled Markov transitions. Consequently, in a closed IMC, hybrid states can be regarded as interactive states.

[1] The only schedulers that can be defined for IMCs are early, whereas both early and late schedulers are definable for CTMDPs.

Download English Version:

https://daneshyari.com/en/article/434920

Download Persian Version:

https://daneshyari.com/article/434920

Daneshyari.com