Information and Computation 260 (2018) 99–125

Contents lists available at ScienceDirect

Information and Computation

www.elsevier.com/locate/yinco

Lumping-based equivalences in Markovian automata: Algorithms and applications to product-form analyses

Giacomo Alzetta [b], Andrea Marin [a,*], Carla Piazza [b], Sabina Rossi [a]

[a] *Università Ca' Foscari Venezia, Italy*
[b] *Università degli Studi di Udine, Italy*

A R T I C L E   I N F O

Article history:
Received 7 February 2017
Received in revised form 13 February 2018

Keywords:
Stochastic automata
Quantitative analysis
Behavioural equivalences
Product-forms stochastic models

A B S T R A C T

In this paper we consider two relations over stochastic automata, named *lumpable bisimulation* and *exact equivalence*, that induce a strong and an exact lumping, respectively, on the underlying Markov chains. We show that an exact equivalence over the states of a non-synchronising automaton is indeed a lumpable bisimulation for the corresponding reversed automaton and then it induces a strong lumping on the time-reversed Markov chain underlying the model. This property allows us to prove that the class of quasi-reversible models is closed under exact equivalence. Quasi-reversibility is a pivotal property to study product-form models. Hence, exact equivalence turns out to be a theoretical tool to prove the product-form of models by showing that they are exactly equivalent to models which are known to be quasi-reversible. Algorithms for computing both lumpable bisimulation and exact equivalence are introduced. Case studies as well as performance tests are also presented.

1. Introduction

Stochastic models play a key role in reliability and performance analysis providing a sound framework for real improvements of software and hardware architectures, including telecommunication systems. Continuous Time Markov Chains (CTMCs) constitute the underlying semantics model of a plethora of modelling formalisms such as Stochastic Petri nets [29], Stochastic Automata Networks (SAN) [32], queueing networks [5] and a class of Markovian process algebras (MPAs), e.g., [18,16]. The aim of these formalisms is to provide a high-level description language for complex real-time systems and automatic analysis techniques. Modularity in the model specification is an important feature of both MPAs and SANs that allows one to describe large systems in terms of interactions of simpler components. Nevertheless, one should notice that a modular specification does not lead to a modular analysis, in general. Thus, although the intrinsic compositional properties of such formalisms are extremely helpful in the specification of complex systems, in many cases carrying out an exact analysis for those models (e.g., those required by quantitative model checking) may be extremely expensive from a computational point of view.

The use of equivalence relations for quantitative models is an important formal approach that allows one to compare different systems as well as to improve the efficiency of some analysis [22,13,12]. To give an example, if we can prove that

* Corresponding author.
E-mail addresses: alzetta.giacomo@spes.uniud.it (G. Alzetta), marin@dais.unive.it (A. Marin), carla.piazza@uniud.it (C. Piazza), srossi@dais.unive.it (S. Rossi).

https://doi.org/10.1016/j.ic.2018.04.002
0890-5401/

a model P is in some sense equivalent to Q and Q is much simpler than P, then we can carry out an analysis of the simplest component to derive the properties of the original one.

Bisimulation based relations on stochastic systems inducing the notions of ordinary (or strong) and exact lumpability for the underlying Markov chains have been studied in [8,2,18,11,34]. In this paper, we first recall the notions of *lumpable bisimulation* [19] and *exact equivalence* [10,11] which have been both proved to be a congruence for Markovian process algebras and stochastic automata whose synchronisation semantics is defined as the master/slave synchronisation of the Stochastic Automata Networks (SAN) and comply with the ordinary and exact, respectively, lumping for Markov processes.

Interestingly, we show that an exact equivalence over a non-synchronising stochastic automaton is indeed a lumpable bisimulation on the reversed automaton (see [26] for a similar result in the context of Markov chains instead of stochastic automata) and then it induces a *strong lumping* on the time-reversed Markov chain underlying the model. This important property, allows us to prove that the class of quasi-reversible [20] stochastic networks is closed under exact equivalence. Quasi-reversibility is one of the most important and widely used characterisations of product-form models, i.e., models whose equilibrium distribution can be expressed as the product of functions depending only on the local state of each component. Informally, we can say that product-forms project the modularity in the model definition to the model analysis, thus drastically reducing the computational costs of the derivation of the quantitative indices. Basically, a composition of quasi-reversible components whose underlying chain is ergodic has a product-form solution, meaning that one can check the quasi-reversibility modularly for each component, without generating the whole state space.

In this paper we provide a new methodology to prove or disprove that a stochastic automaton is quasi-reversible: it is sufficient to show that a model is exactly equivalent to another one which is known to be (or to be not) quasi-reversible. In practice, this approach is useful because proving the quasi-reversibility of a model may be a hard task since it requires one to reverse the underlying CTMC and check some conditions on the reverse process, see, e.g., [20,14,27,25]. Conversely, by using exact equivalence, one can prove or disprove the quasi-reversibility property by considering only the forward model, provided that it is exactly equivalent to another (simpler) quasi-reversible model known in the wide literature of product-forms. Moreover, while automatically proving quasi-reversibility is in general unfeasible, checking the exact equivalence between two automata can be done algorithmically by exploiting a partition refinement strategy, similar to that of Paige and Tarjan's algorithm for bisimulation [30].

We prove that both the notion of lumpable bisimulation and that of exact equivalence can be reduced to a labeled weighted compatibility problem and we generalize the algorithm for compatibility presented in [36] in order to deal with labels without increasing its computational complexity.

This paper is an extended version of the work published in [24]. We extended our previous work by presenting rigorous proofs for all the results stated in the paper and proposing an efficient algorithm for computing both lumpable bisimulations and exact equivalences. Moreover, we introduce two case studies and show a set of performance tests for an implementation of the algorithms.

The paper is structured as follows. Section 2 introduces the notation and recalls the basic definitions on Markov chains. In Section 3 we give the definition of stochastic automata and specify their synchronisation semantics. Section 4 presents the definition of quasi-reversibility for stochastic automata. Lumpable bisimulation and exact equivalence are introduced in Section 5. In this section we prove that the class of quasi-reversible automata is closed under the exact equivalence relation. Algorithms for computing lumpable bisimulation and exact equivalence are presented in Section 6. In Section 7 we describe two case studies and present some performance tests. Finally, Section 8 concludes the paper.

2. Markov chains, reversibility and lumpability

In this section we review the theoretical background on continuous-time Markov chains and the concepts of reversibility and lumpability.

2.1. Continuous-Time Markov Chains

A Continuous-Time Markov Chain (CTMC) is a stochastic process $X(t)$ for $t \in \mathbb{R}^+$ taking values into a discrete state space $\mathcal{S}$ such that (1) $X(t)$ is *stationary*, i.e., $(X(t_1), X(t_2), \dots, X(t_n))$ has the same distribution as $(X(t_1+\tau), X(t_2+\tau), \dots, X(t_n+\tau))$ for all $t_1, t_2, \dots, t_n, \tau \in \mathbb{R}^+$; (2) $X(t)$ has the *Markov property*, i.e., the conditional (on both past and present states) probability distribution of its future behaviour is independent of its past evolution until the present state:

$$Prob(X(t_{n+1}) = s_{n+1} \mid X(t_1) = s_1, X(t_2) = s_2, \dots, X(t_n) = s_n) = Prob(X(t_{n+1}) = s_{n+1} \mid X(t_n) = s_n).$$

A CTMC $X(t)$ is said to be *time-homogeneous* if the conditional probability $Prob(X(t+\tau) = s \mid X(t) = s')$ does not depend upon t, and is *irreducible* if every state in $\mathcal{S}$ can be reached from every other state. A state in a Markov process is called *recurrent* if the probability that the process will eventually return to the same state is one. A recurrent state is called *positive-recurrent* if the expected number of steps until the process returns to it is finite. A CTMC is *ergodic* if it is irreducible and all its states are positive-recurrent. In the case of finite Markov chains, irreducibility is sufficient for ergodicity.