Theoretical Computer Science ••• (••••) •••–•••

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

Exact fluid lumpability in Markovian process algebra

Max Tschaikowski, Mirco Tribastone *

Department for Informatics, Ludwig-Maximilians University of Munich, Germany

ARTICLE INFO

Keywords:
Ordinary differential equations
State-space explosion
Aggregation
Process algebra

ABSTRACT

Quantitative analysis by means of discrete-state stochastic processes is hindered by the well-known phenomenon of state-space explosion, whereby the size of the state space may have an exponential growth with the number of agents of the system under scrutiny. When the stochastic process underlies a Markovian process algebra model, this problem may be alleviated by suitable notions of behavioural equivalence that induce *lumping* at the underlying continuous-time Markov chain, establishing an exact relation between a potentially much smaller aggregated chain and the original one. For the analysis of massively parallel systems, however, lumping techniques may not be sufficient to yield a computationally tractable problem. Recently, much work has been directed towards forms of *fluid* techniques that provide a set of ordinary differential equations (ODEs) approximating the expected path of the stochastic process. Unfortunately, even fluid models of realistic systems may be too large for feasible analysis. This paper studies a behavioural relation for process algebra with fluid semantics, called *projected label equivalence*, which is shown to yield an *exactly fluid lumpable model*, i.e., an aggregated ODE system which can be related to the original one without any loss of information. Project label equivalence relates sequential components of a process term. In general, for any two sequential components that are related in the fluid sense, nothing can be said about their relationship from the stochastic viewpoint. We define and study a notion of well-posedness which allows us to relate fluid lumpability to the stochastic notion of *semi-isomorphism*, which is a weaker version of the common notion of isomorphism between the doubly labelled transition systems at the basis of the Markovian interpretation.

1. Introduction

Markov processes have proven useful in the quantitative modelling of systems. In the case of discrete-state Markov chains, either in (homogeneous) discrete or in continuous time, the solution for transient and steady-state analysis is cast into a classical linear algebra problem for which a wide range of efficient numerical techniques exist [1]. Markov chain models are however intrinsically based on an interleaving semantics, which causes the infamous problem of state-space explosion: the state space size may grow exponentially with the number of elements of the system. A number of approaches to state-space reduction have been devised, including product-form solutions (e.g., [2,3]), decompositions (e.g., [4]) and matrix analytics methods (e.g. [5]). Particularly relevant for the present paper are *lumping techniques*, where a (hopefully much smaller) chain may be defined which preserves the system's original stochastic behaviour, either exactly, as in the case of ordinary and exact lumpability (e.g., [6]), or in an approximate manner (e.g., *near* [6] and *quasi* lumpability [7]).

In this paper we are concerned with continuous-time Markov chains (CTMCs) inferred by Markovian process algebra (e.g., [8–10]). In this context, lumpability has been at the basis of suitable notions of behavioural equivalence between processes

* Corresponding author.
E-mail addresses: tschaikowski@pst.ifi.lmu.de (M. Tschaikowski), tribastone@pst.ifi.lmu.de (M. Tribastone).

0304-3975/$ – see front matter
http://dx.doi.org/10.1016/j.tcs.2013.07.029

Please cite this article in press as: M. Tschaikowski, M. Tribastone, Exact fluid lumpability in Markovian process algebra, Theoretical Computer Science (2013), http://dx.doi.org/10.1016/j.tcs.2013.07.029

which induce an aggregated Markov chain; in the literature, these relations are known as *strong equivalence* in PEPA [8, Chap. 8], *strong Markovian bisimilarity* [11], and *strong performance equivalence* [12]. Although, in general, lumping techniques require the availability of the full CTMC to be aggregated, exploiting results of congruence of such relations with respect to parallel composition has allowed for methods that are able to construct the lumped CTMC *on-the-fly*, i.e., directly without constructing the original CTMC first [13–15]. Although these methods may not yield optimal lumping, they have proven successful in exploiting simple symmetries. For instance, parallel compositions of processes which are syntactically equal up to a permutation, e.g., $P \mid Q$ and $Q \mid P$, can be lumped into the same macro-state of the aggregated CTMC. This typically reduces the computational complexity of the state-space size from exponential to polynomial in the number of sequential processes. However, it still represents a severe limitation when analysing massively distributed systems.

To cope with these difficulties, more recent work has been carried out toward equipping stochastic process algebra with *fluid* semantics. The underlying mathematical object for the analysis is now a system of ordinary differential equations (ODEs) which gives an approximate time-course evolution of the *population* of processes exhibiting a particular local state. Especially in the case of large *population processes*, this approximation is very accurate but typically much more compact than the lumped CTMC, as it is independent of the actual population sizes but is only dependent on the number of local behaviours of the distinct sequential processes. The relationship between the CTMC and the fluid semantics has been studied in the context of PEPA [16], an extension called PEPA$+\Pi$, useful for the analysis of biochemical systems [17], Bio-PEPA [18], Cardelli's stochastic interacting processes [19], and stochastic Concurrent Constraint Programming [20]. The significant computational savings provided by differential analysis, together with its widespread use in computational systems biology, have also stimulated the development of process algebra directly equipped with an ODE semantics [21].

Unfortunately, ODE models of realistic complex systems may still be too large for feasible analysis. This problem has motivated work on ODE aggregations in diverse contexts such as control theory [22], theoretical ecology [23], and economics [24]. In this paper, we study aggregation of ODEs induced by a stochastic process algebra. We carry out this investigation in the context of PEPA, the first process algebra to be equipped with a fluid semantics [16]. In principle, however, with suitable syntactical changes our approach is applicable to other calculi. We consider a distilled version of a recent fluid framework for PEPA, *grouped PEPA* (GPEPA) [25], of which we omit parts that are not essential in the development and understanding of the ideas presented in this paper. Specifically, we consider PEPA without the usual *hiding* operator. The reason for this choice is twofold. First, hiding does not pose technical difficulties in the proofs that we derive in the sequel, hence our results carry over straightforwardly if it is taken into account. Second, although GPEPA supports the hiding operator, in effect its operational semantics carries out a transformation of action types such that *silent* (i.e., τ-type) actions that originate through hiding are relabelled so as to keep track of the original nonsilent action that was performed. Another simplification that we apply to GPEPA is that we consider a kind of *normal form*. Let us briefly sketch the intuition here. In standard GPEPA, a collection of N independent processes to be analysed with the same group of ODEs would be defined as

$$H\{\underbrace{P \parallel P \parallel \cdots \parallel P}_{N \text{ times}}\}, \tag{1}$$

where H is a label that uniquely identifies the group of processes and $\parallel$ denotes CSP-style parallel composition over an empty action set. In our simplified calculus, which we call *Fluid Process Algebra* (FPA), the term would simply correspond to $H\{P\}$, whereas its initial multiplicity is injected by an external function that maps labelled terms into nonnegative integers—e.g., we would write $V_{(H,P)}(0) = N$ to express that, at time $t = 0$, there are N components in state P. This approach is similar to the notion of *reduced context* presented in [26] or to the *species-oriented* view of the system in Bio-PEPA. Formally showing that such a normal form can be considered without loss of generality requires some considerable machinery which however appears to be unnecessary with respect to the theory presented in this paper. The interested reader will find such a discussion in Appendix B; otherwise, our FPA can simply be interpreted as a stand-alone process algebra which considers large replicas of PEPA-like processes, directly yielding an underlying system of ODEs for the analysis. The necessary definitions required for the development of the results on aggregation of ODEs induced by process algebra are instead summarised in Section 2.

Paper contributions. Our study follows a programme which is analogous in spirit to that to be followed for obtaining aggregation of CTMCs induced by process algebra. However, we point out that the techniques used to prove the necessary results are starkly different because of the intrinsically different operational models that underlie the calculi, namely a doubly labelled transition system for the CTMC semantics and an ODE system for the fluid interpretation. Nevertheless, it is interesting to relate the novel concepts introduced in this paper to those that are typically employed for aggregation in the CTMC semantics.

In the CTMC realm, there is a one-to-one mapping between the states of the labelled transition system and the CTMC state. Similarly, one may think of a specific function solution to one differential equation of the overall ODE system as a (continuous) state of the system. A lumpable partition is a partition over the set of the CTMC states such that some requirement over the transition rates is satisfied. In the case of exact lumpability, if it holds that the states in the same partition element are initially equiprobable, then they will be equiprobable at all future time points. In this paper, we introduce the analogous notion of *exact fluid lumpability* (Section 3) which is intuitively defined as a partition over the

Please cite this article in press as: M. Tschaikowski, M. Tribastone, Exact fluid lumpability in Markovian process algebra, Theoretical Computer Science (2013), http://dx.doi.org/10.1016/j.tcs.2013.07.029