Journal of Logical and Algebraic Methods in Programming 84 (2015) 238–258

Contents lists available at ScienceDirect

Journal of Logical and Algebraic Methods in Programming

www.elsevier.com/locate/jlamp

A unified framework for differential aggregations in Markovian process algebra

Max Tschaikowski, Mirco Tribastone *

Electronics and Computer Science, University of Southampton, United Kingdom

ARTICLE INFO

Article history:
Received 19 February 2014
Received in revised form 28 October 2014
Accepted 28 October 2014
Available online 4 November 2014

Keywords:
Markovian process algebra
Fluid semantics
Lumpability

ABSTRACT

Fluid semantics for Markovian process algebra have recently emerged as a computationally attractive approximate way of reasoning about the behaviour of stochastic models of large-scale systems. This interpretation is particularly convenient when sequential components characterised by small local state spaces are present in many independent copies. While the traditional Markovian interpretation causes state-space explosion, fluid semantics is independent from the multiplicities of the sequential components present in the model, just associating a single ordinary differential equation (ODE) with each local state. In this paper we analyse the case of a process algebra model inducing a large ODE system. Previous work, known as *exact* fluid lumpability, requires two symmetries: ODE aggregation is possible for processes that i) are isomorphic and that ii) are present with the same multiplicities. We first relax the latter requirement by introducing the notion of *ordinary* fluid lumpability, which yields an ODE system where the sum of the aggregated variables is preserved exactly. Then, we consider *approximate* variants of both notions of lumpability which make nearby processes symmetric after a perturbation of their parameters. We prove that small perturbations yield nearby differential trajectories. We carry out our study in the context of a process algebra that unifies two synchronisation semantics that are well studied in the literature, useful for the modelling of computer systems and chemical networks, respectively. In both cases, we provide numerical evidence which shows that, in practice, many heterogeneous processes can be aggregated with negligible errors.

1. Introduction

State-space explosion is a well-known problem that hinders our capability to analyse models for large-scale reactive systems that are based on an explicit discrete-state representation. The focus of this paper is on models of concurrent and parallel systems defined by process algebra, where this problem manifests itself with an exponential growth of the cardinality of the state space as a function of the number of interacting processes. Much ingenuity has gone into finding appropriate symmetries at the model level which induce a suitable coarsening of the state space that retains some information about the original one. In this respect, the classical results on bisimilarity allow to relate processes of possibly different state space sizes which are however equivalent with respect to an external observer [1].

* Corresponding author.
E-mail addresses: m.tschaikowski@soton.ac.uk (M. Tschaikowski), m.tribastone@soton.ac.uk (M. Tribastone).

http://dx.doi.org/10.1016/j.jlamp.2014.10.004
2352-2208/

Analogous notions have been lifted to the quantitative setting, when process algebra is enriched with a stochastic semantics that maps onto a discrete-state Markov process rather than onto a labelled transition system. Here, a fundamental research question has been to define appropriate process-algebraic behavioural equivalences that induce a *meaningful* aggregation at the level of the underlying Markov chain. This has already been established for a number of Markovian process algebra. For instance, the notions of Markovian bisimulation for MTIPP [2] and strong equivalence for PEPA [3] are equivalence relations which are shown to yield a *lumped* Markov chain [4]. This is a Markov chain obtained from the original one by appropriately partitioning states, such that the probability of being in each partition block is equal to the sum of all probabilities of being in each state within the block. (See also [5] for a review of Markovian behavioural equivalences.)

With these approaches, little information is lost in the aggregation and the two stochastic processes are related exactly. Perhaps unexpectedly, however, the equivalence relations that permit such aggregations require strong symmetries at the process-algebraic level, and are essentially able to relate processes that perform the same actions at identical rates. However, there is evidence of criticism on this assumption when we insist on using models to represent actual systems. In this case, confronting reality entails that model parameterisation is difficult because measurements may be subject to noise, thereby inevitably introducing estimation errors (e.g., [6,7]). This observation has motivated a body of research into approximate reasoning about formal models of probabilistic and stochastic systems (e.g., [8–11]).

This paper follows the same line of research, by studying exact as well as approximate notions of aggregation for Markovian process algebra. However, differently from all the aforementioned literature, we target *fluid semantics*. This has recently emerged as an alternative to the classical Markovian semantics, describing the model dynamics in terms of a system of ordinary differential equations (ODEs) [12,13]. These can be interpreted as a deterministic approximation to the expectation of the Markov chain (e.g., [14–17]). When the model under consideration consists of many copies of processes in parallel, the ODE system size is independent from the multiplicities of such copies, unlike the Markovian representation, which, as discussed, suffers from state explosion.

Unfortunately, not every process algebra model enjoys a compact ODE description (e.g., [18]). Indeed, the problem of aggregating large-scale models based on ODEs has attracted the attention of researchers in a variety of other disciplines including control theory [19], theoretical ecology [20], and chemical engineering [21]. Arguably, in the process algebra literature the topic of ODE aggregation has so far received less attention than its stochastic counterpart. In [22,23] we studied a notion of behavioural equivalence for ODEs for PEPA [3]; ODE aggregations have also been investigated in [24] for rule-based models such as Kappa [25] and BioNetGen [26], for the modelling and analysis of biomolecular networks.

The goal of this paper is to extend the toolkit of ODE aggregations available for Markovian process algebra, by making the following contributions.

A unified framework for synchronisation semantics. Let us start with the observation that, in general, the development of theories in Markovian process algebra may depend on the choice of the quantitative semantics of the synchronisation operator, since this is typically defined as a function of the rates at which the operands (i.e., the synchronising processes) exhibit some behaviour. The actual choice of the function is motivated by the kinds of the systems that the language is meant to target. PEPA, for instance, can be useful for the performance evaluation of computing systems, since it has been shown to capture the semantics of other formalisms and methods such as queueing networks [27] or stochastic Petri nets [28]. Thus, the previously cited results of ODE aggregations find applicability in that domain. On the other hand, rule-based models essentially employ a semantics based on the *law of mass action*. As discussed, this is well known to be at the basis of biochemical reaction networks, though it has also been employed in epidemiological models (e.g., [29]) as well as in certain wireless networks (e.g. [30]).

Our initial starting point is to consider a unified process-algebraic framework that encompasses both kinds of interaction, in such a way that any result developed in this context finds immediate applicability to either kind of target system/semantics. To do so, we introduce and study *Fluid Extended Process Algebra* (FEPA), a lightweight extension of *Fluid Process Algebra* presented in [22]. Its characteristic trait is a generic parallel operator that is parameterised using a function of two arguments, defining the actual instance of the law of interaction to be used in the synchronisation. To show that FEPA is a conservative extension, we will first establish that the notion of *exact fluid lumpability* (EFL) presented [22] carries over. This will be used to set the stage for approximate notions of ODE aggregations, discussed later in the paper, which are defined in terms of their exact counterparts.

We take EFL as the starting point of our investigation, with the objective of extending it along two orthogonal directions. On the one hand, we define a new notion of lumpability, called *ordinary fluid lumpability* (OFL), which relaxes assumptions on certain symmetries whilst still guaranteeing exactness of the aggregated system. On the other hand, we consider approximate versions of both EFL and OFL which can yield coarser aggregations, at the cost of losing exactness. To be concrete yet informal for the purpose of overviewing our results, let us consider the process

$$\left(P_1[N_1] \parallel_K P_2[N_2] \parallel_K \cdots \parallel_K P_D[N_D]\right) \parallel_L Q[M] \tag{1}$$

where, for all $1 \leq i \leq D$, P_i is some sequential component that is replicated N_i times, and $\parallel_K$ is the (generic) parallel operator, parameterised by an action set K, in a CSP-like fashion. EFL may essentially reduce the analysis of such a model by considering the fluid trajectory of a *representative* P_i, which is shown to be equal to that of any other P_j if, for all $1 \leq i, j \leq D$, it holds that $N_i = N_j$ and P_i and P_j are isomorphic. Thus, denoting by $V_S(t)$ the ODE solution related to the