Normal and inhomogeneous moment closures for stochastic process algebras

Anton Stefanek, Marcel C. Guenther, Jeremy T. Bradley
Department of Computing
Imperial College London
London, United Kingdom
{as1005,mcg05,jb}@doc.ic.ac.uk

This paper discusses the application of moment closures to continuous Markov chains derived from process algebras such as GPEPA and MASSPA. Two related approaches are being investigated. Firstly we re-formulate normal moment closure in a process algebra framework using Isserlis' theorem. Secondly we apply a mixture of this normal closure and less precise moment closures for the purpose of reducing coupling between ordinary differential equations (ODE) derived from the underlying Markov chain. We present three case-studies to show how both normal and inhomogeneous moment closures can significantly improve the numerical accuracy of ODE moment approximations.

1 Introduction

Fluid analysis of continuous Markov chains is a popular technique that deals with the state space explosion problem in models of systems exhibiting massive parallelism. The technique has been derived and used in numerous applications, such as performance evaluation [1, 2, 3], systems biology and chemistry [4, 5, 6].

The main idea of the technique is that in continuous time Markov chains where the states are finite vectors of integer valued *populations*, so–called *Population CTMCs* (PCTMCs), moments, such as mean and variance, of the populations are continuous functions over time. Using the Chapman–Kolmogorov equations, it is possible to derive a system of ordinary differential equations, ODEs, with solution equal to these moments. Although this system of ODEs describes the moments exactly, it can contain infinitely many ODEs and also involve ODEs of auxiliary quantities that have unknown right hand sides. In case of a PCTMC derived from a model written in the PEPA stochastic process algebra, an intuitive approximation can lead to a finite closed system of ODEs describing mean populations [1] and a family of systems of ODEs describing higher/joint moments of populations up-to any finite order [2].

However, in many models coming from systems biology and chemistry, the exact systems of ODEs are infinite. This often results from moments depending on moments of higher order. A simple example is a system of ODEs where ODEs describing the mean depend on 2^{nd} order moments, 2^{nd} order moment ODEs on 3^{rd} order moments and so on. A common reason for this dependence is the non-linear mass-action dynamics, such as in Systems Biology [4, 7, 5] and Ecology [8] and also complex spatial dynamics for example in Markovian Agent Models (MAM)s [9], an agent based spatial performance analysis framework. One way to solve such infinite systems of ODEs numerically is to *close* the equations, i.e. approximate the higher-order dependence so that a finite system of ODEs can be obtained. Many

Submitted to:
PASTA 2011, Sicily

different *moment closures* have been suggested in the past [10, 4, 11], but mostly applying to very specific models. In this paper, we present a novel *inhomogeneous moment closure*, based on the *normal moment closure* [4, 5], that can be used to obtain accurate and efficient moment approximations in a large class of models. We illustrate this closure on spatial models described in the MASSPA process algebra.

In Section 2 we re-formulate the *normal moment closure* using Isserlis' theorem, in the framework of Hayden *et al.* [2]. We confirm that this closure can produce good quantitative and qualitative results without any need for a priori information about the model. However, despite closing the system of ODEs, the normal moment closure still maintains a high degree of coupling between the individual equations and requires a large number of different combinations of joint moments. This limitation is especially apparent in discrete spatial models that need to handle extremely large numbers of populations. In Section 3 we introduce *inhomogeneous moment closures* - a new class of moment closures providing balance between accuracy and computational cost. Finally Section 4 compares the numerical accuracy of the above mentioned moment closures.

2 Normal Moment Closure

As mentioned in the previous section, it is possible to derive a system of ODEs exactly describing the evolution of moments of populations in PCTMCs. For example, taking a hypothetical PCTMC with states $(X_1(t),\dots,X_n(t)) \in \mathbb{N}^n$ and at most quadratic propensities, we can use the method of Engblom [4] (or equivalently that of Gillespie [5] and others) to get a system of ODEs containing equations such as

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbb{E}[X_1(t)] = \dots + C \times \mathrm{Cov}[X_2(t),X_3(t)] + \dots \tag{1}$$

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathrm{Cov}[X_2(t),X_3(t)] = \dots + D \times \mathrm{CM}[X_4(t),X_5(t),X_6(t)] + \dots \tag{2}$$

where $\mathrm{CM}[\cdot]$ is a third joint central moment and C,D some constants. The suggested *normal moment closure*, dating to at least the 1957 work of Whittle [10] reasons that if X_i were multivariate normal, the third central moment would be zero. Ignoring the third moments in the above equation, we can obtain a closed system of ODEs that often provides high accuracy. See Fig. 1 for a PCTMC of a circadian clock model defined in [4]. The framework of Hayden [2] derives systems of ODEs in describing *raw* higher moments instead of central moments. Therefore, we cannot simply ignore terms such as above, but need to use basic properties of expectation to replace third order terms with a linear combination of second order terms. In general, considering closures at order higher than two, we can use Isserlis' theorem: If $X_1,X_2,\dots,X_{2n+1}$ are multivariate normal with mean $\vec{\mu}$ and covariance matrix (σ_{ij})

$$\mathbb{E}[(X_1-\mu_1)\cdots(X_{2n+1}-\mu_{2n+1})] = 0$$
$$\mathbb{E}[(X_1-\mu_1)\cdots(X_{2n}-\mu_{2n})] = \sum\prod \mathbb{E}[(X_i-\mu_i)(X_j-\mu_j)]$$

where $\sum\prod$ sums through all the distinct partitions of $1,\dots,2n$ into disjoint sets of pairs i,j. For example, instead of including an ODE for the third order joint moment $\mathbb{E}[X_1(t)X_2(t)^2]$ we can close the expansion by using the approximation

$$\begin{aligned}&\mathbb{E}[X_1(t)X_2(t)^2]\\ &\quad\approx 2\times\mathbb{E}[X_2(t)]\mathbb{E}[X_2(t)X_1(t)] + \mathbb{E}[X_1(t)]\mathbb{E}[X_2(t)^2] - 2\times\mathbb{E}[X_1(t)]\times\mathbb{E}[X_2(t)]^2\end{aligned}$$

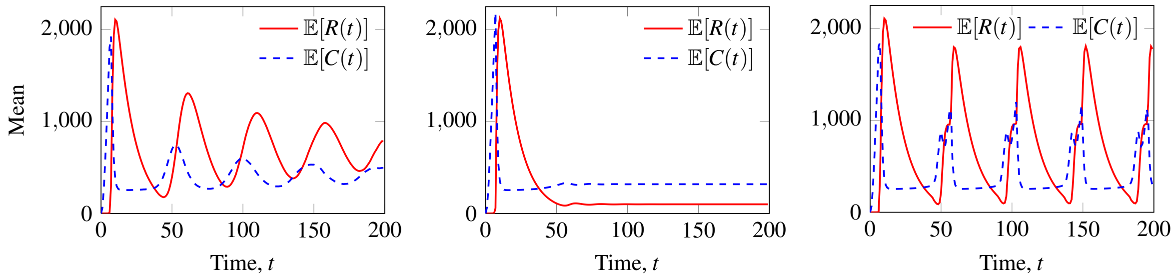

(a) Simulation (b) First order closure (c) Second order closure

Figure 1: Normal moment closure illustrated on a model of circadian clock. Figure (a) shows the mean of two component populations as obtained from 10^4 replications of stochastic simulation. Figure (b) shows an estimate of these means when using a first order ODE (or mean field) approximation. Figure (c) shows the notable qualitative improvement when ODEs for the second order moments are used and closed under the normal moment closure.

3 Inhomogeneous Moment Closure

Despite the fact we reach the lowest error when applying the normal closure to all ODEs (cf. Sections 2 and 4), there may be situations where it is infeasible to apply the normal closure globally. One reason could be that the application of a normal closure to all ODEs would be too expensive due to its tendency to couple ODEs heavily. Another one might be that we want to close some ODEs at lower orders than others to get more precise results for some important populations. *Inhomogeneous moment closures* can help to overcome these problems. In general these are closure strategies that apply a combination of various moment closure techniques to a system of ODEs. In contrast to the generic closures such as the normal closure, inhomogeneous closures are usually model or framework specific and often need a priori reasoning about the model they are applied to. Their benefit, however, is that they can produce highly accurate results at lower cost than for instance a standard normal closure. In Section 4 we give a detailed example of a spatially motivated inhomogeneous moment closure.

4 Worked Examples

In this section we investigate a spatially motivated inhomogeneous 2^{nd} order moment closure for the evaluation of MASSPA models. MASSPA models produce unclosed systems of ODEs due to quadratic terms in the agent communication rates [12]. These terms originate from message induced agent communication and take the general form of $\alpha\mathbb{E}[XZ]$, where X and Z stand for the component counts of the receiving and sending agent's population respectively. In earlier research [9] mean field approximation $\mathbb{E}[XZ] \approx \mathbb{E}[X]\mathbb{E}[Z]$ was applied as a 1^{st} order closure. In [12] the authors suggest a simple 2^{nd} order closure of the form $\mathbb{E}[XYZ] \approx \mathbb{E}[XY]\mathbb{E}[Z]$, where Z is again symbolising the sending agent's population. We term this closure a 2^{nd} order mean field closure. While this closure was found to give qualitatively good results, it was also shown to produce poor quantitative results in some cases. Furthermore, the accuracy of the mean population counts was hardly improved by taking into account the variance and covariance information when applying this closure. In the following we introduce a spatially motivated

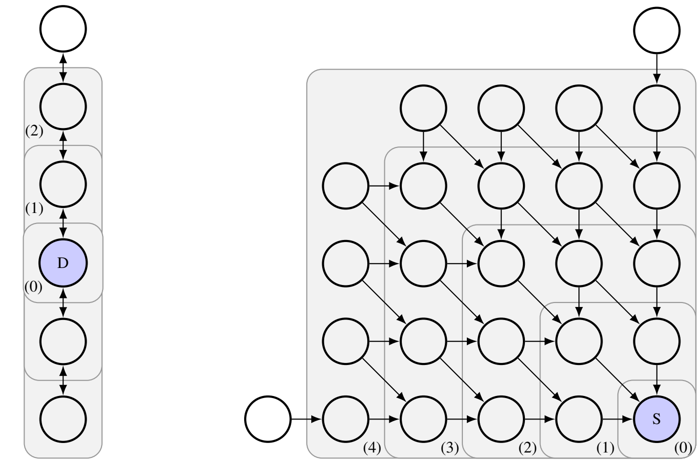

(a) On/Off (b) WSN

Figure 2: Topologies of the On/Off and the WSN model. The On/Off model uses a simple inter-communication pattern, the WSN model a source to sink pattern. The (i)-labels represent the neighbourhood areas captured by the respective MASSPA(i) closure.

inhomogeneous moment closure which overcomes these deficits.

To improve the 2nd order ODE approximation we first applied the normal closure (cf. Section 2) to all second order ODEs. This approach worked well and we found that both 1st and 2nd order ODEs became more accurate in both the simple *On/Off* and the *WSN* model [12]. However, applying the normal closure to all ODEs in a MASSPA model can be computationally expensive as it couples the ODEs much more heavily than the 2nd order mean field closure. As a consequence we devised an inhomogeneous 2nd order closure for MASSPA models, which uses a combination of the normal and the 2nd order mean field closure. Suppose we want to evaluate the mean of population $X@l$, the population of an agent in state X at location l. Let $ds(X@l)$ denote the derivative states of this agent population, assuming an ergodic agent chain. We define the MASSPA(i) closure to be a 2nd order closure that applies the normal closure to any 3rd order term that contains a derivative population of any ith order neighbour of the populations that we want to evaluate. All other 3rd order terms in 2nd order ODEs are approximated by $\mathbb{E}[XYZ] \approx \mathbb{E}[XY]\mathbb{E}[Z]$. Any agent population that $X@l$ can communicate with directly, is considered a 1st order neighbour. Any 1st order neighbour of $X@l$'s 1st order neighbours is considered a 2nd order neighbour of $X@l$ and so on. This naturally implies that the closure becomes more expensive if we evaluate different population counts in different locations, as this increases the number of 3rd order terms that we approximate using a normal closure. Hence, this closure is most economical when analysing quantities in a particular spatial region. Two special cases are MASSPA(0) which only uses the normal closure for 3rd order terms involving $ds(X@l)$ and MASSPA() which does not use any normal closure. The MASSPA($\cdot$) closure is motivated by the assumption that covariances of component counts, which are spatially close to the component counts to be evaluated, have a higher impact on the numerical accuracy. However, similar inhomogeneous closures could be applied to ODEs resulting from non-spatial SPAs, too.

The following analysis has been conducted on the two models mentioned in [12]. While agent definitions were not altered, we increased the number of locations in both models. Fig. 2 shows the new

Closure	ODEs	$KS(\mathbb{E}[On@D_1])$	$KS(DEV[On@D_1])$	$KS(\mathbb{E}[On@D_{10}])$	$KS(DEV[On@D_{10}])$
MASSPA()	29	0.00194847	0.17180182	0.00019674	0.17218328
MASSPA(0)	60	0.00099163	0.04948122	0.00010303	0.04980903
MASSPA(1)	87	0.00025405	0.00571611	0.00004656	0.00567927
MASSPA(2)	90	0.00012499	0.00176882	0.00003852	0.00189467

(a) On/Off model with $Off@D_1 = 300$ and $Off@D_{10} = 3000$ initially.

Closure	ODEs	$KS(\mathbb{E}[WSN\_0@S])$	$KS(DEV[WSN\_0@S])$	$KS(\mathbb{E}[WSN\_1@S])$	$KS(DEV[WSN\_1@S])$
MASSPA()	1070	0.00060328	0.00190883	0.00034313	0.00058631
MASSPA(0)	1190	0.00037395	0.00031677	0.00023814	0.00018996
MASSPA(1)	1840	0.00024354	0.00012443	0.00015451	0.00010978
MASSPA(2)	3370	0.00019028	0.00013081	0.00012511	0.00010570
MASSPA(3)	7740	0.00018925	0.00013081	0.00012509	0.00010527
Closure	ODEs	$KS(\mathbb{E}[WSN\_2@S])$	$KS(DEV[WSN\_2@S])$	$KS(\mathbb{E}[WSN\_3@S])$	$KS(DEV[WSN\_3@S])$
MASSPA()	1070	0.00034717	0.00061081	0.00061779	0.00235754
MASSPA(0)	1190	0.00024066	0.00015023	0.00044477	0.00048434
MASSPA(1)	1840	0.00022213	0.00015306	0.00026074	0.00013934
MASSPA(2)	3370	0.00022068	0.00015469	0.00019027	0.00014406
MASSPA(3)	7740	0.00022084	0.00015489	0.00019756	0.00014448

(b) WSN model with $WSN\_0@S = 100$ initially.

Figure 3: The two tables show the Kolmogorov-Smirnov (KS) error divided by the agent population size. To compute the K-S statistic we compared the ODE solutions with the ensemble statistics of a Gillespie simulation with 500,000 replications. $DEV[\cdot]$ represents the standard deviation. The ODEs column shows the number of ODEs needed to compute $\mathbb{E}[\cdot]$ and $DEV[\cdot]$ for the chosen populations.

topologies for the *On/Off* and the *WSN* model. In Fig. 3 we compare the accuracy of the inhomogeneous MASSPA(i) moment closure for various values of i. The difference between the normalised K-S errors obtained for the MASSPA() and the MASSPA(i) closures shows that introducing the normal closure always improves the accuracy of both mean and standard deviation. The benefit becomes most apparent in the standard deviation approximation, especially for $DEV[On@D_1]$ and $DEV[On@D_{10}]$. Our examples also show that higher values of i may not necessarily improve the accuracy. This can be seen in the *WSN* model where MASSPA(2) is just as accurate as MASSPA(3) and MASSPA(4), the latter being obmitted from the table. Furthermore we investigated the impact of scaling on MASSPA($\cdot$) closures. The model from which $On@D_{10}$ was computed had a ten times larger population than the model we used to obtain the $On@D_1$ moments. As expected, scaling increased the accuracy of the mean approximation. However, even in the larger model, MASSPA(i) closures improved the accuracy significantly. Finally a comparison between the number of ODEs needed by the MASSPA(i) closures and the MASSPA() closure shows that the computational overhead is reasonable compared to the accuracy gained. In any case the ODE analysis was much faster than the Gillespie simulation no matter which MASSPA($\cdot$) closure we used.

5 Conclusions

We have applied the normal moment closure to unclosed ODEs derived from stochastic process algebra definitions. Our examples show that the closure can improve the accuracy of the numerical result without any need for prior knowledge about the structure of the underlying PCTMC. Furthermore we have introduced inhomogenous closures and shown in the MASSPA example that there are meaningful inhomogeneous closures which can achieve similar accuracy compared to a normal closure at lower computational cost. In the future we will look at further inhomogeneous moment closures. One particularly promising candidate is an order dependent closure, which varies the order at which ODEs of certain population moments are closed. Moreover, we intend to investigate the impact of the normal and the inhomogenous closure presented in this paper on further models.

References

[1] J. Hillston, "Fluid flow approximation of PEPA models," *Second International Conference on the Quantitative Evaluation of Systems QEST 2005*, pp. 33–42, 2005.

[2] R. A. Hayden and J. T. Bradley, "A fluid analysis framework for a Markovian process algebra," *Theoretical Computer Science*, vol. 411, no. 22-24, pp. 2260–2297, 2010.

[3] R. Bakhshi, L. Cloth, W. Fokkink, and B. R. Haverkort, "Mean-field framework for performance evaluation of push-pull gossip protocols," *Performance Evaluation*, vol. 68, no. 2, pp. 157 – 179, 2011. Advances in Quantitative Evaluation of Systems - QEST 2009.

[4] S. Engblom, "Computing the moments of high dimensional solutions of the master equation," *Applied Mathematics and Computation*, vol. 180, no. 2, pp. 498–515, 2006.

[5] C. S. Gillespie, "Moment-closure approximations for mass-action models.," *IET systems biology*, vol. 3, no. 1, pp. 52–58, 2009.

[6] D. T. Gillespie, *Markov Processes: An Introduction for Physical Scientists*. Academic Press, 1991.

[7] L. Bortolussi, "On the Approximation of Stochastic Concurrent Constraint Programming by Master Equation," *Electronic Notes in Theoretical Computer Science*, vol. 220, no. 3, pp. 163–180, 2008.

[8] D. J. Murrell, U. Dieckmann, and R. Law, "On moment closures for population dynamics in continuous space.," *Journal of Theoretical Biology*, vol. 229, no. 3, pp. 421–432, 2004.

[9] M. Gribaudo, D. Cerotti, and A. Bobbio, "Analysis of On-off policies in Sensor Networks Using Interacting Markovian Agents," *6th IEEE International Conference on Pervasive Computing and Communications PerCom (2008)*, pp. 300–305, 2008.

[10] P. Whittle, "On the use of the normal approximation in the treatment of stochastic processes," *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 19, no. 2, pp. pp. 268–281, 1957.

[11] I. Nsell, "Moment closure and the stochastic logistic model," *Theoretical Population Biology*, vol. 63, no. 2, pp. 159 – 168, 2003.

[12] M. C. Guenther and J. T. Bradley, "Higher moment analysis of a spatial stochastic process algebra.," in *8th European Performance Engineering Workshop - EPEW 2011*, 2011.