

Mean field and fluid approaches to Markov chain analysis

Jeremy T. Bradley *

jb@doc.ic.ac.uk

Department of Computing, Imperial College London, UK

Representing the explicit state space of performance models has inherent difficulties. Just as the state-space explosion effects functional correctness evaluation, so it can also be easily a problem in performance models. In particular, classical Markov chain analysis of any variety requires exploration of the global state space and, even for a simple system, this quickly becomes computationally infeasible. Fluid and mean-field analysis techniques attempt to side-step the state-space explosion and provide a computationally cheap way of analysing certain features of Markov chains.

1 Introduction

In recent years, there has been a substantial interest in the *fluid* or *mean-field* approach to analysing stochastic problems in computing [1, 2, 3, 4, 5, 6, 7] and stochastic process algebra models in particular [8, 9, 10, 11, 12]. This style of analysis of the underlying Markov process offers an attractive computational alternative to traditional explicit state space analysis techniques. However, there are certain limitations which need to be kept in mind. Often fluid and mean-field analysis require the construction of an aggregate state space that means that specific model features are abstracted away. Additionally, in constructing a fluid model of a large Markov chain which incorporates synchronisation features, approximations may need to be made and it is important that a modeller is made aware of those approximations when analysing their model. If we are prepared to live with these aspects of the analysis, then the computational benefits are such that huge analysis tasks can be carried out in a relatively short period of time.

Quantitative analysis of Markov processes by means of the fluid or mean-field approach exploits substantial parallelism in the original model to construct a series of ordinary differential equations (ODEs) which capture the transient evolution of the system. Analysis of substantial Markov models using variety of fluid techniques [1, 13, 10, 8] is made possible and analyses of state spaces of 10^{100} states and beyond are not uncommon. Explicit state-space performance techniques which analyse the underlying continuous-time Markov chain directly (for example, [14, 15]) are typically limited to 10^{11} states for a steady-state style of analysis. It is perhaps, slightly unfair to make this comparison directly, as fluid and mean-field analyses do not scale with the number of global states in a model, but instead with the local state space of the constituent Markov processes.

In discrete time, mean-field techniques have been well documented by Benaïm and Le Boudec [2]. McCaig *et al.* [12] developed a discrete-time mean-field framework around the synchronous process algebra WSCCS, based on original work by Sumpter [16]. While Bakhshi *et al.* [4, 7] have developed some discrete-time model-specific analysis techniques for gossip protocols using the mean-field technique. In the continuous-time domain, Hillston developed *fluid-flow* analysis [1] to make first-order approximations of massively parallel PEPA models. Bortolussi [8] has presented a formulation for the stochastic constraint programming language, sCCP and Cardelli has a first-order fluid analysis translation to ordinary differential equations (ODEs) for the π -calculus [9].

*The author is funded in part by the EPSRC on the AMPS project (reference EP/G011737/1).

2 Fluid passage-time analysis

The key impact of fluid and mean-field analysis techniques has been that it allows the modeller to trial many different system configurations and parameterisations easily whereas, as has been noted, a single system instance might take many hours or days to analyse using classical techniques. Latterly, fluid techniques have been used to extract useful passage-time measures [17, 18, 19] from Markov processes specified using stochastic process algebra. This allows a modeller to focus on the duration of key transactions in the higher-level system. With the rapid computation of fluid analysis, parameter sweeping can be efficiently performed to find which model variables (exponential rates or scale parameters) have greatest effect on a key passage-time measure. Often these passage times will take the form of requirements or service level objectives (SLOs) for the system, for example, 97.8% of search queries should be responded to within 0.5 seconds. Understanding how a key passage time is sensitive to changes in model parameters can help improve the design of the system. In particular how a passage time reacts to changes in the scale of component deployment within the system, so-called scalability analysis, is critical to the system engineering process.

References

- [1] J. Hillston, “Fluid flow approximation of PEPA models,” in *QEST’05, Proceedings of the 2nd International Conference on Quantitative Evaluation of Systems*, (Torino), pp. 33–42, IEEE Computer Society Press, September 2005.
- [2] M. Benaïm and J.-Y. Le Boudec, “A class of mean field interaction models for computer and communication systems,” *Performance Evaluation*, vol. 65, no. 11-12, pp. 823–838, 2008.
- [3] L. Massoulié and M. Vojnovic, “Coupon replication systems,” *IEEE/ACM Transactions on Networking*, vol. 16, pp. 603–616, June 2008.
- [4] R. Bakhshi, L. Cloth, W. Fokkink, and B. Haverkort, “Mean-field analysis for the evaluation of gossip protocols,” in *QEST’09, Proceedings of the 5th IEEE Conference on the Quantitative Evaluation of Systems*, pp. 247–256, IEEE Computer Society, September 2009.
- [5] A. Ganesh, S. Lilienthal, D. Manjunath, A. Proutiere, and F. Simatos, “Load balancing via random local search in closed and open systems,” *ACM SIGMETRICS Performance Evaluation Review*, vol. 38, p. 287, June 2010.
- [6] S. Shakkottai and R. Johari, “Demand-aware content distribution on the Internet,” *IEEE/ACM Transactions on Networking*, vol. 18, pp. 476–489, Apr. 2010.
- [7] 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.
- [8] L. Bortolussi and A. Policriti, “Stochastic concurrent constraint programming and differential equations,” in *QAPL’07, 5th Workshop on Quantitative Aspects of Programming Languages*, vol. 190 of *Electronic Notes in Theoretical Computer Science*, pp. 27–42, Elsevier, September 2007.
- [9] L. Cardelli, “On process rate semantics,” *Theoretical Computer Science*, vol. 391, pp. 190–215, February 2008.
- [10] R. Hayden and J. T. Bradley, “A fluid analysis framework for a Markovian process algebra,” *Theoretical Computer Science*, vol. 411, pp. 2260–2297, May 2010.

- [11] M. Tribastone, S. T. Gilmore, and J. Hillston, “Scalable differential analysis of process algebra models,” *IEEE Transactions on Software Engineering*, vol. 38, pp. 205–219, Jan–Feb 2012.
- [12] C. McCaig, R. Norman, and C. Shankland, “From individuals to populations: A mean field semantics for process algebra,” *Theoretical Computer Science*, vol. 412, no. 17, pp. 1557–1580, 2011.
- [13] L. Cardelli, “On process rate semantics,” *Theoretical Computer Science*, vol. 391, pp. 190–215, February 2008.
- [14] W. J. Knottenbelt, P. G. Harrison, M. S. Mestern, and P. S. Kritzinger, “A probabilistic dynamic technique for the distributed generation of very large state spaces,” *Performance Evaluation*, vol. 39, pp. 127–148, February 2000.
- [15] M. Kwiatkowska, G. Norman, and D. Parker, “PRISM: Probabilistic symbolic model checker,” in *TOOLS’02, Proceedings of the 12th International Conference on Modelling Techniques and Tools for Computer Performance Evaluation* (A. J. Field et al., ed.), vol. 2324 of *Lecture Notes in Computer Science*, (London), pp. 200–204, Springer-Verlag, 2002.
- [16] D. J. T. Sumpter and D. S. Broomhead, “Relating individual behaviour to population dynamics,” *Proceedings of the Royal Society: Series B*, vol. 268, pp. 925–932, 2001.
- [17] A. Clark, A. Duguid, S. T. Gilmore, and M. Tribastone, “Partial evaluation of PEPA models for fluid-flow analysis,” in *Proceedings of the 5th European Performance Engineering Workshop on Computer Performance Engineering (EPEW)* (N. Thomas and C. Juiz, eds.), vol. 5261 of *Lecture Notes in Computer Science*, pp. 2–16, Springer Berlin Heidelberg, Aug. 2008.
- [18] R. Hayden, A. Stefanek, and J. T. Bradley, “Fluid computation of passage time distributions in large Markov models,” *Theoretical Computer Science*, vol. 413, pp. 106–141, January 2012.
- [19] R. Hayden, J. T. Bradley, and A. Clark, “Performance specification and evaluation with Unified Stochastic Probes and fluid analysis,” *IEEE Transactions on Software Engineering*, 2012. (In press).