Mean-field models for interacting battery-powered devices

Richard A. Hayden
Department of Computing
Imperial College London
rh@doc.ic.ac.uk

1 Introduction

Markovian fluid models [e.g. 6, 11], including *fluid stochastic Petri nets (FSPNs)* [e.g. 14, 18, 19] and *Markovian reward models* [e.g. 2, 7, 15, 17] allow the performance analysis of systems which include *a priori* continuous in addition to discrete state elements. In the area of performance–energy modelling, a key example of such a system would be a battery-powered device. In addition to the discrete operating state of the device, the level of remaining charge in the battery must also be modelled. This could be achieved in the context of a standard Markov chain by splitting the possible charge levels into discrete stages, but would usually be much more accurately captured through a continuous real variable [8].

However, if, for example, many such interacting battery-powered devices are to be considered together in a model of, say, a wireless sensor network, the well-known problem of state-space explosion is quickly a barrier to scalable analysis. In the case of purely discrete-state Markov chain models, so-called *mean-field* and *fluid-analysis* techniques based on ordinary differential equations (ODEs) [e.g. 1, 3, 4, 9, 10, 12, 13] have recently become very popular as a means to achieve a scalable approximate analysis of such massively-parallel models.

In this paper, we show briefly how such techniques can be extended to systems consisting of many interacting fluid models, that is, where individual agents in the model each have associated one or more continuous variables. In contrast to mean field for discrete-state Markov chains, the mean-field limit will be expressed as the solution to a small system of non-linear partial (functional) differential equations (P(F)DEs)[1]. Similarly to the Markov chain case, the size of this system will be independent of population size, which is the key property that guarantees a scalable analysis.

A similar approach has been applied in [5] to a system of interacting mobile nodes which each node had associated an 'age variable' keeping track of the last time they were updated. In this case, the value of the continuous variable does not impact on the discrete-state evolution of the nodes. This allows the mean-field analysis to be decoupled so that the discrete-state evolution can be obtained first using traditional mean-field methods in terms of a system of ODEs. The solutions to these ODEs are then used to construct a simpler system of PDEs without the need for non-local terms on the right-hand sides. Furthermore, the fact that the continuous variable in [5] always decreases at the same rate in each discrete state means that the PDEs admit relatively straightforward solutions as transformations of simpler ODE systems. Neither of these assumptions are true for our purposes resulting in a more complicated analysis. Indeed, we desire that the battery level may directly impact the evolution of individuals, not least when it runs out. Furthermore, any realistic battery model must allow for different discharge (or charge) rates in different discrete states.

[1] *Functional* here means that the right-hand side of the differential equations may depend non-locally on the dependent functions, i.e. may include integral terms or reference the dependent functions at specific values of the independent variables.

Submitted to:
ICEP 2012

2 A simple example

We introduce the approach in this paper by means of a very simple worked example, but it should be clear how it could be extended to much more detailed scenarios. Specifically, we consider here a large number N of wireless nodes each powered by their own local battery, which is modelled by a continuous variable in $[0,1]$ so that 1 represents a fully-charged battery and 0 an empty battery. Nodes behave very simply — each node can be in one of two discrete states: *active* (`a`) or *idle* (`i`). We assume an idle node is collecting data or waiting for stimulus and an active node is exchanging information with a neighbour. In both states, nodes consume energy from their battery, potentially at different (deterministic) rates. As a power-saving measure, each node implements a threshold control which becomes active when its battery level decreases below a given level, say $A > 0$. We say a node is *high* if its battery level is greater than A, *low* if its battery level is less than or equal to A but strictly greater than zero, and *empty* if the battery level is zero. When a node is low, it adapts the transmission power of its wireless radio so that the energy consumption of an active low node is $\gamma_l \geq 0$ and of an active high node is $\gamma_h \geq 0$. In the idle state, energy consumption is $\gamma_i \geq 0$ independent of battery level, and, presumably parameters are chosen so that $\gamma_i \leq \gamma_l \leq \gamma_h$ (although this is not a technical requirement). When a node is empty, it ceases to move between states or partake in the exchange of information with other nodes.

A node moves from state `i` to state `a` at Markovian rate λ. A node moves from state `a` to state `i` after successfully exchanging information with a neighbour node. The rate at which this happens depends on the current transmission power of the node's radio (and thus the node's battery level) and on the number of other currently active nodes and the transmission power at which their respective radios are operating. The Markovian rate we use for this transition is then:

$$\mathbf{1}_{\{0<z\leq A\}}\beta_l\frac{(\beta_l(\mathbf{a}_l-1)+\beta_h\mathbf{a}_h)}{N}+\mathbf{1}_{\{z>A\}}\beta_h\frac{(\beta_l(\mathbf{a}_l)+\beta_h(\mathbf{a}_h-1))}{N}$$

where z is the node's current battery level, $\mathbf{a}_h$ is the number of active high nodes and $\mathbf{a}_l$ is the number of active low nodes. We follow the usual methodology for the stochastic modelling of epidemics in that the chance of a node finding a neighbour in a small period of time is proportional to the number of other active nodes. We have adapted this so that it depends also on the transmission power of the other nodes' radios as captured through the parameters β_l and β_h. Indeed, if $\beta_l = \beta_h$, this rate is simply $\frac{\beta(\mathbf{a}-1)}{N}$, where $\mathbf{a} := \mathbf{a}_l + \mathbf{a}_h$ and $\beta := \beta_l^2 = \beta_h^2$.

2.1 Derivation of mean-field differential equations

The mean-field approximations will be given by the following functions: $F_a(t,z)$ will be the mean-field approximation to the proportion of active nodes at time t with battery level less than or equal to z and $F_i(t,z)$ will capture the same quantity for idle nodes. For suitable initial conditions, we will see that these functions admit densities on $(0,1]$ with respect to z at all times t, which we write as $f_a(t,z)$ and $f_i(t,z)$, respectively. Lifting the case of standard fluid models [11] to the mean-field level, we would also anticipate a collection of mass at $z = 0$ as batteries run out, so we will also require functions $e_a(t)$ and $e_i(t)$ which give the mean-field approximations to the proportions of empty active and empty idle nodes, respectively. The mean-field equations will be given in terms of these four quantities and we will then recover $F_a(t,z) = e_a(t) + \int_0^z f_a(t,v)\,\mathrm{d}v$ and $F_i(t,z) = e_i(t) + \int_0^z f_i(t,v)\,\mathrm{d}v$.

We now show how the mean-field differential equations can be derived informally. Let $t, \delta t \in \mathbb{R}_+$

and $z \in (0,1)$, then, heuristically for large N, we have:

$$f_a(t+\delta t,z) \approx f_a(t,z) + \underbrace{\left(f_a\left(t, z+[\mathbf{1}_{\{z\le A\}}\gamma_l + \mathbf{1}_{\{z>A\}}\gamma_h]\delta t\right) - f_a(t,z)\right)}_{\text{Discharging of batteries in } [t,t+\delta t]} + \underbrace{\lambda\,\delta t\, f_i(t,z)}_{\text{Discrete transitions } \mathtt{i}\to\mathtt{a} \text{ in } [t,t+\delta t]}$$
$$\underbrace{-(\mathbf{1}_{\{z\le A\}}\beta_l + \mathbf{1}_{\{z>A\}}\beta_h)\delta t\, f_a(t,z)\left(\beta_l\int_0^A f_a(t,v)\,\mathrm{d}v + \beta_h\int_A^1 f_a(t,v)\,\mathrm{d}v\right)}_{\text{Discrete transitions } \mathtt{a}\to\mathtt{i} \text{ in } [t,t+\delta t]} + o(\delta t)$$

Then dividing by δt and taking the limit $\delta t \to 0$, we obtain the following PFDE, accurate for large N:

$$\frac{\partial f_a(t,z)}{\partial t} - (\mathbf{1}_{\{z\le A\}}\gamma_l + \mathbf{1}_{\{z>A\}}\gamma_h)\frac{\partial f_a(t,z)}{\partial z} =$$
$$\lambda f_i(t,z) - (\mathbf{1}_{\{z\le A\}}\beta_l + \mathbf{1}_{\{z>A\}}\beta_h) f_a(t,z)\left(\beta_l\int_0^A f_a(t,v)\,\mathrm{d}v + \beta_h\int_A^1 f_a(t,v)\,\mathrm{d}v\right)$$

Treating the evolution of $f_i(t,z)$ similarly, we obtain:

$$\frac{\partial f_i(t,z)}{\partial t} - \gamma_i\frac{\partial f_i(t,z)}{\partial z} =$$
$$(\mathbf{1}_{\{z\le A\}}\beta_l + \mathbf{1}_{\{z>A\}}\beta_h) f_a(t,z)\left(\beta_l\int_0^A f_a(t,v)\,\mathrm{d}v + \beta_h\int_A^1 f_a(t,v)\,\mathrm{d}v\right) - \lambda f_i(t,z)$$

Due to the discontinuities at $z = A$, we consider here a weak form of solution that is not necessarily smooth at such points. In addition, we have:

$$e_a(t+\delta t) \approx e_a(t) + \underbrace{\int_0^{\gamma_l\delta t} f_a(t,v)\,\mathrm{d}v}_{\text{Discharging of batteries in } [t,t+\delta t]} + o(\delta t)$$

Dividing by δt, taking the limit $\delta t \to 0$ and making the analogous considerations for $e_i(t)$ gives:

$$\frac{\mathrm{d}e_a(t)}{\mathrm{d}t} = \gamma_l f_a(t,0) \qquad \frac{\mathrm{d}e_i(t)}{\mathrm{d}t} = \gamma_l f_i(t,0)$$

where $f_a(t,0) := \lim_{z\to 0} f_a(t,z)$ and $f_i(t,0) := \lim_{z\to 0} f_i(t,z)$.

The model considered here could, in theory, be analysed exactly in terms of a very large (exponential in N) system of *linear* PDEs in $N+1$ variables[2], by applying the standard theory for the transient analysis of fluid models [11]. Analogously to the case of mean field for continuous-time Markov chains (CTMCs) in terms of ODEs, the mean-field approach reduces the required analysis to that of a very small system (independent of N) of *non-linear* hyperbolic PDEs, which, in this case, additionally have functional (integral) terms.

Since the equations for $f_a(t,z)$ and $f_i(t,z)$ involve both derivatives with respect to t and to z, we need to supply both initial and boundary conditions, which we give as follows: $f_a(0,z) := g(z)$ and $f_i(0,z) := 0$ for $z \in (0,1]$, where $g : (0,1] \to \mathbb{R}_+$ is the initial density of battery levels and is taken here to be truncated normal with parameters $\mu = 1.0$ and $\sigma = 0.05$. We have assumed here that all nodes start active (other choices for initial conditions are of course possible) and also $f_a(t,1) := f_i(t,1) := 0$

[2]One variable for time and N to keep track of each node's battery level.

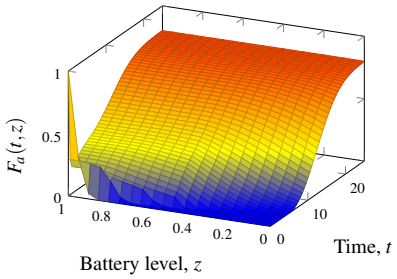

(a) Mean-field solution for $F_a(t,z)$

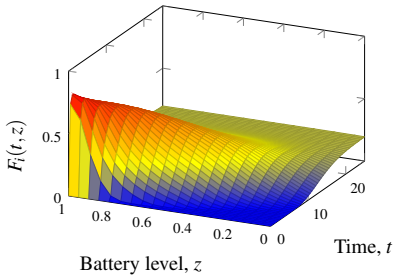

(b) Mean-field solution for $F_i(t,z)$

Compared to single trace

Pop. size	Avg. error	Max. error
$N = 10$	0.1425	0.3540
$N = 100$	0.0264	0.1212
$N = 1000$	0.0138	0.0345
$N = 10000$	0.0042	0.0336

Compared to mean of 1000 traces

Pop. size	Avg. error	Max. error
$N = 10$	0.0307	0.0577
$N = 100$	0.0042	0.0234
$N = 1000$	*Simulation too costly*	
$N = 10000$	*Simulation too costly*	

(c) Mean-field compared to simulations over $(t,z) \in [0,25] \times [0,1]$

Figure 1: Mean-field solutions and stochastic simulation comparisons. Parameters: $\gamma_h = 0.35, \gamma_l = 0.1, \gamma_i = 0.03, \beta_h = 2.5, \beta_l = 1.0, \lambda = 0.3, A = 0.5$.

for $t > 0$ since no batteries can be full after some time has elapsed. Finally, we set $e_a(0) := e_i(0) := 0$ and thus assume that no batteries start completely empty. In order to solve this system of four coupled equations numerically, we employ a two-step version of the Lax–Friedrichs method [16]. Some example results and comparisons with stochastic simulation are given in Figure 1. The results in Figure 1c suggest mean-field convergence empirically (as $N \to \infty$), but, for reasons of brevity, we do not give the proof here.

3 Conclusion

In this paper we have briefly illustrated a mean-field methodology based on partial functional differential equations for the analysis of massive systems of interacting fluid models. We have focussed on one important application area — systems of interacting battery-powered devices and shown how the approach yields accurate results.

References

[1] Jonatha Anselmi & Maaike Verloop (2011): *Energy-aware capacity scaling in virtualized environments with performance guarantees*. Performance Evaluation 68(11), pp. 1207–1221, doi:10.1016/j.peva.2011.07.004.

[2] Christel Baier, Lucia Cloth, Boudewijn R. Haverkort, Holger Hermanns & Joost-Pieter Katoen (2010): *Performability assessment by model checking of Markov reward models*. Formal Methods in System Design 36(1), pp. 1–36, doi:10.1007/s10703-009-0088-7.

[3] Rena Bakhshi, Lucia Cloth, Wan Fokkink & Boudewijn R. Haverkort (2011): *Mean-field framework for performance evaluation of push-pull gossip protocols*. Performance Evaluation 68(2), pp. 157–179, doi:10.1016/j.peva.2010.08.025.

[4] Michel Benaïm & Jean-Yves Le Boudec (2008): *A class of mean field interaction models for computer and communication systems*. Performance Evaluation 65(11-12), pp. 823–838, doi:10.1016/j.peva.2008.03.005.

[5] Augustin Chaintreau, Jean-Yves Le Boudec & Nikodin Ristanovic (2009): *The age of gossip*. In: *Eleventh international joint conference on measurement and modeling of computer systems (SIGMETRICS)*, ACM Press, New York, New York, USA, pp. 109–120, doi:10.1145/1555349.1555363.

[6] Dongyan Chen, Yiguang Hong & Kishor S. Trivedi (2002): *Second-order stochastic fluid models with fluid-dependent flow rates*. Performance Evaluation 49(1-4), pp. 341–358, doi:10.1016/S0166-5316(02)00113-X.

[7] Lucia Cloth (2006): *Model checking algorithms for Markov reward models*. Ph. D., University of Twente.

[8] Lucia Cloth, Marijn R. Jongerden & Boudewijn R. Haverkort (2007): *Computing Battery Lifetime Distributions*. In: *37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, IEEE, pp. 780–789, doi:10.1109/DSN.2007.26.

[9] Nicolas Gast & Gaujal Bruno (2010): *A mean field model of work stealing in large-scale systems*. In: *ACM SIGMETRICS international conference on Measurement and modeling of computer systems*, 38, ACM Press, New York, New York, USA, p. 13, doi:10.1145/1811039.1811042.

[10] Marco Gribaudo (2009): *Analysis of Large Populations of Interacting Objects with Mean Field and Markovian Agents*. In Jeremy T. Bradley, editor: *6th European Performance Engineering Workshop (EPEW)*, *Lecture Notes in Computer Science* 5652, Springer Berlin Heidelberg, pp. 218–219, doi:10.1007/978-3-642-02924-0.

[11] Marco Gribaudo & Miklós Telek (2007): *Fluid models in performance analysis*. In Marco Bernardo & Jane Hillston, editors: *Formal Methods for Performance Evaluation (SFM)*, *Lecture Notes in Computer Science* 4486, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 271–317, doi:10.1007/978-3-540-72522-0.

[12] Richard A. Hayden & Jeremy T. Bradley (2010): *A fluid analysis framework for a Markovian process algebra*. *Theoretical Computer Science* 411(22-24), pp. 2260–2297, doi:10.1016/j.tcs.2010.02.001.

[13] Jane Hillston (2005): *Fluid flow approximation of PEPA models*. In: *Second International Conference on the Quantitative Evaluation of Systems (QEST)*, IEEE, pp. 33–42, doi:10.1109/QEST.2005.12.

[14] Graham Horton, Vidyadhar G. Kulkarni, David M. Nicol & Kishor S. Trivedi (1998): *Fluid stochastic Petri nets: Theory, applications, and solution techniques*. *European Journal of Operational Research* 105(1), pp. 184–201, doi:10.1016/S0377-2217(97)00028-3.

[15] Andrew Reibman, Roger Smith & Kishor Trivedi (1989): *Markov and Markov reward model transient analysis: An overview of numerical approaches*. *European Journal of Operational Research* 40(2), pp. 257–267, doi:10.1016/0377-2217(89)90335-4.

[16] Lawrence F. Shampine (2005): *Two-step Lax–Friedrichs method*. *Appl. Math. Letters*, pp. 1134–1136.

[17] Edmundo de Souza e Silva & H. Richard Gail (1998): *An algorithm to calculate transient distributions of cumulative rate and impulse based reward*. *Communications in Statistics. Stochastic Models* 14(3), pp. 509–536, doi:10.1080/15326349808807486.

[18] Kishor S. Trivedi & Vidyadhar G. Kulkarni (1993): *FSPNs: Fluid stochastic Petri nets*. In: *Application and Theory of Petri Nets*, *Lectures Notes in Computer Science* 691, Springer Berlin / Heidelberg, pp. 24–31, doi:10.1007/3-540-56863-8_38.

[19] Bruno Tuffin, Dong S. Chen & Kishor S. Trivedi (2001): *Comparison of Hybrid Systems and Fluid Stochastic Petri Nets*. *Discrete Event Dynamic Systems* 11(1), pp. 77–95, doi:10.1023/A:1008387132533.