

Size Expansions of Mean Field Approximation: Transient and Steady-State Analysis

Nicolas Gast, Luca Bortolussi, Mirco Tribastone

► **To cite this version:**

Nicolas Gast, Luca Bortolussi, Mirco Tribastone. Size Expansions of Mean Field Approximation: Transient and Steady-State Analysis. 2018 - 36th International Symposium on Computer Performance, Modeling, Measurements and Evaluation, Dec 2018, Toulouse, France. pp.1-2, 10.1016/j.peva.2018.09.005 . hal-01891636

HAL Id: hal-01891636

<https://hal.inria.fr/hal-01891636>

Submitted on 9 Oct 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Size Expansions of Mean Field Approximation: Transient and Steady-State Analysis

Extended abstract*

Nicolas Gast
Inria, France
Univ. Grenoble Alpes, LIG
Grenoble, France
nicolas.gast@inria.fr

Luca Bortolussi
DMG, University of Trieste
Italy
lbortolussi@units.it

Mirco Tribastone
IMT School for Advanced
Studies Lucca
Italy
mirco.tribastone@imtlucca.it

ABSTRACT

Mean field approximation is a powerful tool to study the performance of large stochastic systems that is known to be exact as the system's size N goes to infinity. Recently, it has been shown that, when one wants to compute expected performance metric in steady-state, mean field approximation can be made more accurate by adding a term in $1/N$ to the original approximation. This is called the *refined* mean field approximation in [7]. In this paper, we show how to obtain the same result for the transient regime and we provide a further refinement by expanding the term in $1/N^2$ (both for transient and steady-state regime). Our derivations are inspired by moment-closure approximation. We provide a number of examples that show this new approximation is usable in practice for systems with up to a few tens of dimensions.

1. INTRODUCTION

Mean field approximation is a widely used technique in the performance evaluation community (*e.g.*, to study load-balancing strategies [9, 10] to mention a popular application). The focus of this approximation is to study the performance of systems composed of a large number of interacting objects. This approximation can be used to study transient (for example the time to fill a cache [6]) or steady-state properties (for example the steady-state response time of a system [9, 10]). One of the reasons of the success of mean field approximation is that it is often very accurate as soon as N , the number of objects in the system, exceeds a few hundreds. In fact, this approximation can be proven to be asymptotically exact as N goes to infinity, see for example [8, 1] and explicit bounds for the convergence rate exist [2, 3, 11].

Recently, the authors of [7] proposed what they call a *refined* mean field approximation that can be used to characterize steady-state performance metrics more precisely. Their refinement uses that for many models, a steady-state expected performance metric of a system with N objects $\mathbb{E}[h(X)]$ is equal to its mean field approximation $h(\pi)$ plus

a term in $1/N$:

$$\mathbb{E}[h(X)] = h(\pi) + \frac{1}{N}V_{(h)} + o\left(\frac{1}{N}\right), \quad (1)$$

where π is the fixed point of the ODE that describes the mean field approximation and $V_{(h)}$ is a constant that can be easily evaluated numerically. By using a number of examples, they show that the refined approximation $h(\pi) + \frac{1}{N}V_{(h)}$ is much more accurate than the mean field approximation for moderate system sizes (*i.e.*, a few tens of objects).

2. MAIN RESULTS

In [5], we extend this method in two directions: First we generalize Equation (1) to the transient behavior; second we establish the existence of a second order term in $1/N^2$ (both in transient and steady-state regimes). More precisely, we establish conditions such that for any smooth function h , there exist constants $V_{(h)}$ and $A_{(h)}$ such that for any time $t \in [0; \infty) \cup \{\infty\}$:

$$\mathbb{E}[h(X(t))] = h(x(t)) + \frac{1}{N}V_{(h)}(t) + \frac{1}{N^2}A_{(h)}(t) + o\left(\frac{1}{N^2}\right).$$

We show that for the transient regime $t \in [0, \infty)$, $V_{(h)}(t)$ and $A_{(h)}(t)$ satisfy a linear time-inhomogeneous differential equation that can be easily integrated numerically. For the steady-state, the constants $V_{(h)}(\infty)$ and $A_{(h)}(\infty)$ correspond to the fixed point of this differential equation.

We use the above expansion to propose two new approximations that depend on the system size N and that are expansions of the classical mean field approximation to the order $1/N$ and $1/N^2$. This gives three approximations:

- The mean field approximation: $h(x(t))$;
- $1/N$ -expansion: $h(x(t)) + V_{(h)}(t)/N$;
- $1/N^2$ -expansion: $h(x(t)) + V_{(h)}(t)/N + A_{(h)}(t)/N^2$.

We derive our results for the classical model of density-dependent population process [8].

3. NUMERICAL COMPARISON

In [5], we study numerically the accuracy of the approximation by considering three examples: two malware propagation models and the classical supermarket models of [9, 10]. Our numerical results shows that the two expansions

*The full version is available at [5].

very accurately capture the transient behavior of such a system even when $N \approx 10$. Moreover, they are generally much more accurate than the classical mean field approximation for small values of N (for transient and steady-state regimes). Our experiments also confirm the good accuracy of the $1/N$ -expansion approximation that was observed for steady-state values in [7]: In most cases, the largest gain in accuracy comes from the $1/N$ -term (both for the transient and steady-state values). The $1/N^2$ -term does improve the accuracy but only marginally. We also study the limit of the method by studying an unstable mean field model that has an unstable fixed point.

The Supermarket Model. To give a flavor of the results, let us consider the supermarket model (more examples can be found in the full paper). The system is composed of N identical servers. Jobs arrive at a central broker according to a Poisson process of rate ρN and are dispatched towards the servers by using the JSQ(2) policy: for each incoming job, the broker samples 2 servers at random and sends the jobs to the server that has the smallest number of jobs in its queue (ties are broken at random). The time to process a job is exponentially distributed with mean 1.




Figure 1: Supermarket model : expected queue length as a function of time.

In Figure 1, we plot how the expected queue length evolves with time for the supermarket model with $N = 10$ servers and $\rho = 0.9$. We start in a system where 8 queues have 3 jobs and 2 queues have two jobs. We observe that the two expansions are much more accurate than the classical mean field approximation and that the $1/N^2$ -expansion provides a slightly better approximation than the $1/N$ -expansion.

In Table 1, we show the average queue length for different values of ρ and N . Again, we observe that the values coming from the two expansions are much closer to the values obtained by simulation than the classical mean field approximation. Again, most of the gain seems to come from the $1/N$ -term.

N	k	ρ	Mean field	$1/N$ -expansion	$1/N^2$ -expansion	Simulation
10	2	0.9	2.3527	2.7513	2.8045	2.8002
20	2	0.9	2.3527	2.5520	2.5653	2.5662
10	2	0.95	3.2139	4.1017	4.3265	4.2993
20	2	0.95	3.2139	3.6578	3.7140	3.7124

Table 1: Supermarket model, steady-state average queue length : comparison of the value computed by simulation with the three approximations.

4. CONCLUSION

To summarize, we show in [5] how mean field approximation can be refined by a first term in $1/N$ and a second term $1/N^2$ where N is the size of the system. We exhibit conditions that ensure that this asymptotic expansion can be applied for the transient as well as the steady-state regimes. In the transient regime, these constants satisfy ordinary differential equations that can be easily integrated numerically. We provide a few examples that show that the $1/N$ and $1/N^2$ expansions are much more accurate than the classical mean field approximation. We also study the limitations of the approach and show that, when the mean field approximation does not have an attractor, these new approximations might be unstable for large time horizons. Obtaining a better approximation in this case remains a challenge that we leave for future work.

Finally, we show that, despite the complexity of the formulas, it is relatively easy to compute the $1/N$ and $1/N^2$ terms (in the transient and steady-state regimes) for realistic models such as the supermarket model. The developed method is generic and is implemented in a tool [4].

Reproducibility. The code to reproduce the paper – including simulations, figures and text – is available at [5].

5. REFERENCES

- [1] Michel Benaïm and Jean-Yves Le Boudec. A class of mean field interaction models for computer and communication systems. *Performance Evaluation*, 65(11):823–838, 2008.
- [2] Luca Bortolussi and Richard A Hayden. Bounds on the deviation of discrete-time markov chains from their mean-field model. *Performance Evaluation*, 70(10):736–749, 2013.
- [3] Nicolas Gast. Expected values estimated via mean-field approximation are $1/n$ -accurate. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 1(1):17, 2017.
- [4] Nicolas Gast. Refined mean field tool. https://github.com/ngast/rmf_tool, 2018.
- [5] Nicolas Gast, Luca Bortolussi, and Mirco Tribastone. Size expansions of mean field approximation: Transient and steady-state analysis. <https://github.com/ngast/sizeExpansionMeanField>, 2018.
- [6] Nicolas Gast and Benny Van Houdt. Transient and steady-state regime of a family of list-based cache replacement algorithms. In *Proceedings of the 2015 ACM SIGMETRICS*, pages 123–136. ACM, 2015.
- [7] Nicolas Gast and Benny Van Houdt. A refined mean field approximation. *Proc. ACM Meas. Anal. Comput. Syst*, 1, 2017.
- [8] Thomas G Kurtz. Solutions of Ordinary Differential Equations as Limits of Pure Jump Markov Processes. *Journal of Applied Probability*, 7:49–58, 1970.
- [9] Michael David Mitzenmacher. *The Power of Two Random Choices in Randomized Load Balancing*. PhD thesis, PhD thesis, Graduate Division of the University of California at Berkeley, 1996.
- [10] Nikita Dmitrievna Vvedenskaya, Roland L’vovich Dobrushin, and Fridrikh Izrailevich Karpelevich. Queueing system with selection of the shortest of two queues: An asymptotic approach. *Problemy Peredachi Informatsii*, 32(1):20–34, 1996.
- [11] Lei Ying. On the approximation error of mean-field models. In *Proceedings of the 2016 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Science*, pages 285–297. ACM, 2016.