Queueing Systems Vol. 100
Special Issue: 100 views on queues

Open Problems in Refined Mean Field Approximations

Tim Hellemans · **Benny Van Houdt**

February 2, 2022

1 Introduction

Mean field models are a powerful and popular tool to approximate large-scale stochastic systems in computer science, biology and beyond [9, 1, 2, 5]. Many mean field models are characterized by a set of ordinary differential equations (ODEs), the unique fixed point of which approximates the steady state performance when the stochastic system becomes large. For instance, consider a system consisting of N homogeneous FIFO queues with service rate one, where jobs arrive at rate λN and are dispatched among the N queues by a dispatcher. Assume further that the dispatcher selects d queues at random when a job arrives and assigns the job to the server with the least number of jobs among the d selected servers (meaning the dispatcher uses the power-of-d choices dispatching policy denoted as $SQ(d)$), then the set of ODEs in case that the jobs have exponential durations with mean 1, is given by

$$\frac{d}{dt}u_i(t) = \lambda(u_{i-1}(t)^d - u_i(t)^d) - (u_i(t) - u_{i+1}(t)),$$

for $i \geq 1$, where $u_i(t)$ represents the fraction of the queues with i or more jobs (meaning $u_0(t) = 1$ for all t). For this particular model one can prove [10, 11] that the stationary measures of the stochastic systems weakly converge as N tends to infinity to the Dirac measure of the unique fixed point of the set of ODEs given by $\pi_i(\lambda, d) = \lambda^{\frac{d^i-1}{d-1}}$. In other words the probability that there are i or more jobs in a queue converges to $\pi_i(\lambda, d)$ as N tends to infinity.

Convergence of course does not guarantee high accuracy for finite N. For instance, for the above model the mean field approximation is known to worsen as λ approaches one. In [3] it was proven that a broad class of mean field models is $1/N$-accurate in

Tim Hellemans
ORTEC, Boortmeerbeek, Belgium. E-mail: timhellemanstim@gmail.com

Benny Van Houdt
Department of Computer Science, University of Antwerp, Belgium. E-mail: benny.vanhoudt@uantwerpen.be

	$N=10$	$N=20$	$N=30$	$N=50$	$N=100$	$\lambda^2/(2\cdot(1-\lambda))$
$\lambda=0.7$	0.8930	0.8680	0.8490	0.8500	0.8300	0.8167
$\lambda=0.9$	4.5130	4.2760	4.1400	4.0850	4.0400	4.0500
$\lambda=0.95$	10.8130	10.0420	9.6270	9.3150	9.0800	9.0250

Table 1 Table of the value $N\cdot(W_N^{SQ(2)}(\lambda)-W^{SQ(2)}(\lambda))$ for different choices of λ and N.

	$N=10$	$N=20$	$N=30$	$N=50$	$N=100$	$\lambda^2/(2\cdot(1-\lambda))$
$\lambda=0.7$	1.0912	1.0068	0.9880	1.0243	0.8825	0.8167
$\lambda=0.9$	4.9676	4.5711	4.4835	4.4857	4.2606	4.0500
$\lambda=0.95$	11.8948	10.7204	10.2907	9.9708	9.4035	9.0250

Table 2 Table of the value $N\cdot(W_N^{LL(2)}(\lambda)-W^{LL(2)}(\lambda))$ for different choices of λ and N.

expectation. More specifically, denote by $W_N^{SQ(d)}(\lambda)$ the average queue length for the system of N homogeneous servers defined above. If we denote the expected queue length for the mean field approximation as $W^{SQ(d)}(\lambda)=\sum_{i\geq 1}\pi_i(\lambda,d)$, the difference $|W^{SQ(d)}(\lambda)-W_N^{SQ(d)}(\lambda)|$ is of order $O(1/N)$. Furthermore, in [6] it was shown that a refined mean field model can be developed such that the approximation error is of the order $O(1/N^2)$. This means that for the above model there exists a (computable) value $\zeta(\lambda,d)$ such that $|W^{SQ(d)}(\lambda)+\zeta(\lambda,d)/N-W_N^{SQ(d)}(\lambda)|=O(1/N^2)$. One can even go even further to obtain a more accurate refined approximation at a higher computational cost [4]. It was further noted that $\zeta(\lambda,2)\approx\frac{\lambda^2}{2(1-\lambda)}$ in this specific model. We repeat the numeric experiment carried out in that paper, but we present the value of

$$N\cdot(W_N^{SQ(2)}(\lambda)-W^{SQ(2)}(\lambda)),$$

in Table 1. We observe that $N\cdot(W_N^{SQ(2)}(\lambda)-W^{SQ(2)}(\lambda))$ indeed converges to a constant close to $\lambda^2/2(1-\lambda)$.

2 Problem Statement and Discussion

While the refined mean field approximation in [6] can be applied to many systems, these systems all have a discrete state space. There is however also a strong interest in mean field models that involve continuous variables. For example, suppose that we consider the same system as in the previous section, but now the dispatcher assigns each incoming job to the server with the smallest workload among d randomly selected servers (instead of with the least number of jobs). We refer to this dispatching policy as LL(d). Then, it was shown in [8] that $\bar{F}_N(w)$, the probability that the workload exceeds w in a server, converges to

$$\bar{F}(w)=\left(\lambda+(\lambda^{1-d}-\lambda)e^{(d-1)w}\right)^{\frac{1}{1-d}},$$

as N tends to infinity and $\int_{w>0}\bar{F}(w)dw=\sum_{n\geq 0}\frac{\lambda^{dn+1}}{1+n(d-1)}$.

To study the accuracy of this mean field approximation for finite N, we repeat the experiment performed for SQ(d) for LL(d). Denote $W_N^{LL(d)}(\lambda)$ as the mean workload in the stochastic system with N servers and let $W^{LL(d)}(\lambda) = \int_{w>0} \bar{F}(w)dw$ be the mean field approximation. Table 2 presents the value of $N \cdot (W_N^{LL(2)}(\lambda) - W^{LL(2)}(\lambda))$ for various combinations of λ and N. We observe the same phenomenon as for SQ(d), that is, the mean field approximation also appears to be $1/N$-accurate. In fact, we have performed similar experiments for other workload dependent load balancing policies such as Redundancy-d (see [7]) and also observed a $1/N$-accuracy. This brings us to the following two open problems regarding refined mean field approximations:

1. Is it possible to prove that certain types of mean field models involving continuous variables are $1/N$-accurate as was done in [3] for broad classes of mean field models with a discrete state space?
2. If so, can we develop efficient numerical methods to compute a refined mean field approximation as was done in [6] in the discrete case?

We believe these are challenging open problems that are of high interest to the performance evaluation community.

References

1. M. Benaim and J.-Y. Le Boudec. A class of mean field interaction models for computer and communication systems. *Performance Evaluation*, 65(11-12):823–838, 2008.
2. L. Bortolussi, J. Hillston, D. Latella, and M. Massink. Continuous approximation of collective system behaviour: A tutorial. *Performance Evaluation*, 70:317–349, 2012.
3. N. 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. N. Gast, L. Bortolussi, and M. Tribastone. Size expansions of mean field approximation: Transient and steady-state analysis. *Performance Evaluation*, 129:60–80, 2019.
5. N. Gast and B. Gaujal. Markov chains with discontinuous drifts have differential inclusion limits. *Performance Evaluation*, 69(12):623–642, 2012.
6. N. Gast and B. Van Houdt. A refined mean field approximation. *Proc. ACM Meas. Anal. Comput. Syst.*, 1(2):33:1–33:28, Dec. 2017.
7. T. Hellemans, T. Bodas, and B. Van Houdt. Performance analysis of workload dependent load balancing policies. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 3(2):35, 2019.
8. T. Hellemans and B. Van Houdt. On the power-of-d-choices with least loaded server selection. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 2(2):1–22, 2018.
9. T. Kurtz. *Approximation of population processes*. Society for Industrial and Applied Mathematics, 1981.
10. M. Mitzenmacher. The power of two choices in randomized load balancing. *IEEE Transactions on Parallel and Distributed Systems*, 12(10):1094–1104, 2001.
11. N. Vvedenskaya, R. Dobrushin, and F. Karpelevich. Queueing system with selection of the shortest of two queues: an asymptotic approach. *Problemy Peredachi Informatsii*, 32:15–27, 1996.