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

Nicolas Gast Inria, France
Luca Bortolussi Univ. Trieste, Italy
Mirco Tribastone IMT Lucca, Italy

IFIP Performance 2018, Toulouse

Good system design needs performance evaluation

Example : load balancing

N servers

Which allocation policy?

- Random
- Round-robin
- *JSQ*
- *JSQ*(d)
- *JIQ*

Good system design needs performance evaluation

Example : load balancing

N servers

Which allocation policy?

- Random
- Round-robin
- *JSQ*
- *JSQ*(d)
- *JIQ*

We need methods to characterize emerging behavior starting from a stochastic model of interacting objects

- We can use mean field approximation.

Mean Field and Refined Mean Field Approximations

For the steady-state performance of many systems,[1]:

$$Perf(N) \approx \underbrace{Perf(\infty)}_{\text{mean field approximation}} + \frac{V}{N}$$

We provided analytical and numerical methods to compute V.

Example: steady-state average queue length ($\rho = 0.9$)

Policy	Mean Field ($N = \infty$)	$N = 100$ Simu.	$N = 10$ Simu
SQ(2)	2.35	2.39	2.80
Pull-push	1.64	1.70	2.30

[1] Ref: "A Refined Mean Field Approximation" by G. and Van Houdt (SIGMETRICS 2018)

Mean Field and Refined Mean Field Approximations

For the steady-state performance of many systems,[1]:

$$Perf(N) \approx \underbrace{\underbrace{Perf(\infty)}_{\text{mean field approximation}} + \frac{V}{N}}_{\text{refined mean field approximation}}$$

We provided analytical and numerical methods to compute V.

Example: steady-state average queue length ($\rho = 0.9$)

Policy	Mean Field ($N=\infty$)	$N=100$		$N=10$	
		Simu.	R.M.F.	Simu	R.M.F.
SQ(2)	2.35	2.39	2.39	2.80	2.75
Pull-push	1.64	1.70	1.70	2.30	2.29

[1] Ref: "A Refined Mean Field Approximation" by G. and Van Houdt (SIGMETRICS 2018)

Question addressed in this work

Our contributions

- We can compute the next term in the expansion.
- We can do the same analysis for the transient regime?
- We study the cost (computation) and the benefit (accuracy).

$$Perf(N,t) \approx Perf(\infty,t) + \frac{V(t)}{N} + \frac{A(t)}{N^2} + \dots$$

Outline

Outline

"Mean field approximation" simplifies many problem

But how to apply it?

$\lim_{N\to\infty}$

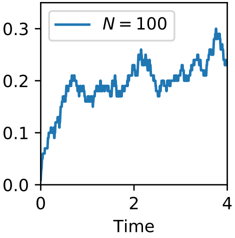

$=$

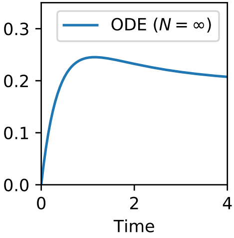

Mean field approximation

Applications :

- Performance of load balancing / caching algorithms
- Communication protocols (CSMA, MPTCP, Simgrid)
- Mean field games (evacuation, Mexican wave)
- Stochastic approximation / learning
- Theoretical biology

The supermarket model (SQ(2))

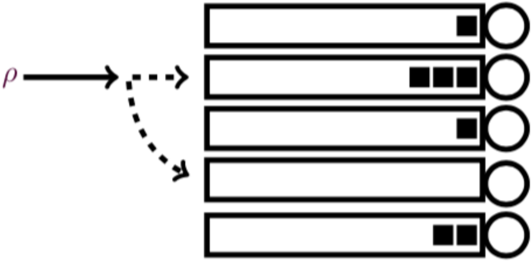

Arrival at each server ρ.

- Sample $d-1$ other queues.
- Allocate to the shortest queue

Service rate$=1$.

$SQ(d)$: state representation

The state space is $X = (X_1, X_2, \dots)$ where

$$X_i(t) = \text{fraction of queues with queue length} \geq i.$$

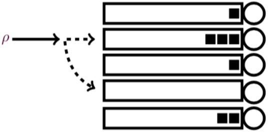

$$X = (1, 0.8, 0.4, 0.2, 0, 0, 0, \dots)$$

State transitions and Mean Field Approximation

State changes on x:

$$x \mapsto x + \frac{1}{N}\mathbf{e_i} \text{ at rate } N\rho(x_{i-1}^d - x_i^d)$$
$$x \mapsto x - \frac{1}{N}\mathbf{e_i} \text{ at rate } N(x_i - x_{i+1})$$

The mean field approximation is to consider the ODE associated with the drift (average variation):

$$\dot{x}_i = \underbrace{\rho(x_{i-1}^d - x_i^d)}_{\text{Arrival}} - \underbrace{(x_i - x_{i+1})}_{\text{Departure}}$$

Density dependent population process (Kurtz, 70s)

A population process is a sequence of CTMCs $X^N(t)$ indexed by the population size N, with state space $E^N \subset E$ and transitions (for $\ell \in \mathcal{L}$):

$$X \mapsto X + \frac{\ell}{N} \qquad \text{at rate } N\beta_\ell(X).$$

Density dependent population process (Kurtz, 70s)

A population process is a sequence of CTMCs $X^N(t)$ indexed by the population size N, with state space $E^N \subset E$ and transitions (for $\ell \in \mathcal{L}$):

$$X \mapsto X + \frac{\ell}{N} \qquad \text{at rate } N\beta_\ell(X).$$

The Mean field approximation

The drift is $f(x) = \frac{d}{dt}\mathbb{E}\left[X(t) \mid X(0) = x\right] = \sum_\ell \ell\beta_\ell(x)$.

The mean field approximation is the solution of the ODE $\dot{x} = f(x)$.

Density dependent population process (Kurtz, 70s)

A population process is a sequence of CTMCs $X^N(t)$ indexed by the population size N, with state space $E^N \subset E$ and transitions (for $\ell \in \mathcal{L}$):

$$X \mapsto X + \frac{\ell}{N} \qquad \text{at rate } N\beta_\ell(X).$$

The Mean field approximation

The drift is $f(x) = \frac{d}{dt}\mathbb{E}\left[X(t) \mid X(0) = x\right] = \sum_\ell \ell\beta_\ell(x)$.

The mean field approximation is the solution of the ODE $\dot{x} = f(x)$.

Example: SQ(d) load balancing

$$\dot{x}_i = \rho(x_{i-1}^d - x_i^d) - (x_i - x_{i+1})$$

It has a unique attractor: $\pi_i = \rho^{(d^i-1)/(d-1)}$.

Accuracy of the mean field approximation

Numerical example of SQ(d) load balancing ($d = 2$)

N	Simulation (steady-state average queue length)					Fixed Point
	10	20	30	50	100	∞ (mean field)
$\rho = 0.7$	1.2194	1.1735	1.1584	1.1471	1.1384	1.1301
$\rho = 0.9$	2.8040	2.5665	2.4907	2.4344	2.3931	2.3527
$\rho = 0.95$	4.2952	3.7160	3.5348	3.4002	3.3047	3.2139

Fairly good accuracy for $N = 100$ servers.

Accuracy of the mean field approximation

Numerical example of SQ(d) load balancing ($d = 2$)

	Simulation (steady-state average queue length)					Fixed Point
N	10	20	30	50	100	∞ (mean field)
$\rho = 0.7$	1.2194	1.1735	1.1584	1.1471	1.1384	1.1301
$\rho = 0.9$	2.8040	2.5665	2.4907	2.4344	2.3931	2.3527
$\rho = 0.95$	4.2952	3.7160	3.5348	3.4002	3.3047	3.2139

Fairly good accuracy for $N = 100$ servers.

Outline

Expected values estimated by mean field are 1/N-accurate

Some experiments (for SQ(2) with $\rho = 0.9$):

N	10	100	1000	∞
Average queue length (simulation)	2.8040	2.3931	2.3567	2.3527
Error of mean field	0.4513	0.0404	0.0040	0

Error decreases as $1/N$

System Size Expansion Approach

Recall that the transitions are $X \mapsto X + \frac{\ell}{N}$ at rate $N\beta_\ell(x)$.

$$\frac{d}{dt}\mathbb{E}[X] = \mathbb{E}\left[\sum_\ell \beta_\ell(X)\ell\right] = \mathbb{E}[f(X)] \qquad \text{(Exact)}$$

$$\frac{d}{dt}x = f(x) \qquad \text{(Mean Field Approx.)}$$

System Size Expansion Approach

Recall that the transitions are $X \mapsto X + \frac{\ell}{N}$ at rate $N\beta_\ell(x)$.

$$\frac{d}{dt}\mathbb{E}[X] = \mathbb{E}\left[\sum_\ell \beta_\ell(X)\ell\right] = \mathbb{E}[f(X)] \qquad \text{(Exact)}$$

$$\frac{d}{dt}x = f(x) \qquad \text{(Mean Field Approx.)}$$

We can now look at the second moment:

$$\begin{aligned}\mathbb{E}[(X-x)\otimes(X-x)] &= \mathbb{E}[(f(X)-f(x))\otimes(X-x)] \qquad \text{(Exact)}\\ &+ \mathbb{E}[(X-x)\otimes(f(X)-f(x))]\\ &+ \frac{1}{N}\mathbb{E}\left[\sum_{\ell\in\mathcal{L}} \beta_\ell(X)\ell\otimes\ell\right]\end{aligned}$$

System Size Expansion Approach

Recall that the transitions are $X \mapsto X + \frac{\ell}{N}$ at rate $N\beta_\ell(x)$.

$$\frac{d}{dt}\mathbb{E}[X] = \mathbb{E}\left[\sum_\ell \beta_\ell(X)\ell\right] = \mathbb{E}[f(X)] \qquad \text{(Exact)}$$

$$\frac{d}{dt}x = f(x) \qquad \text{(Mean Field Approx.)}$$

We can now look at the second moment:

$$\begin{aligned}\mathbb{E}[(X-x)\otimes(X-x)] &= \mathbb{E}[(f(X)-f(x))\otimes(X-x)] \qquad \text{(Exact)}\\ &+ \mathbb{E}[(X-x)\otimes(f(X)-f(x))]\\ &+ \frac{1}{N}\mathbb{E}\left[\sum_{\ell\in\mathcal{L}}\beta_\ell(X)\ell\otimes\ell\right]\end{aligned}$$

... We can also look at higher order moments

$$\begin{aligned}\mathbb{E}\left[(X-x)^{\otimes 3}\right] &= 3\mathrm{Sym}\mathbb{E}[(f(X)-f(x))\otimes(X-x)\otimes(X-x)]\\ &+ \frac{3}{N}\mathrm{Sym}\mathbb{E}\left[\sum_{\ell\in\mathcal{L}}\beta_\ell(X)\ell\otimes\ell\otimes(X-x)\right] + \frac{1}{N}\mathbb{E}\left[\sum_{\ell\in\mathcal{L}}\beta_\ell(X)\ell\otimes\ell\otimes\ell\right]\end{aligned}$$

Using this approach, we can derive linear ODEs

Theorem. Assume that f is C^2 and let x be the solution of $\frac{d}{dt}x = f(x)$.

$$\frac{d}{dt}\mathbb{E}\left[X(t)\right] = x(t) + O(1/N).$$

Using this approach, we can derive linear ODEs

Theorem. Assume that f is C^2 and let x be the solution of $\frac{d}{dt}x = f(x)$.

$$\frac{d}{dt}\mathbb{E}\left[X(t)\right] = x(t) + O(1/N).$$

Let $Y(t) = X(t) - x(t)$. Then :

$$\mathbb{E}\left[Y(t)\right] = \frac{1}{N}V(t) + O(1/N^2)$$
$$\mathbb{E}\left[Y(t)\otimes Y(t)\right] = \frac{1}{N}W(t) + O(1/N^2)$$

where

$$\frac{d}{dt}V^i = f^i_j V^j + f^i_{j,k} W^{j,k}$$
$$\frac{d}{dt}W^{j,k} = f^j_\ell W^{\ell,k} + f^k_\ell W^{j,\ell}$$

Using this approach, we can derive linear ODEs

Theorem. Assume that f is C^2 and let x be the solution of $\frac{d}{dt}x = f(x)$.

$$\frac{d}{dt}\mathbb{E}\left[X(t)\right] = x(t) + O(1/N).$$

Let $Y(t) = X(t) - x(t)$. Then :

$$\mathbb{E}\left[Y(t)\right] = \frac{1}{N}V(t) + \frac{1}{N^2}A(t) + O(1/N^3)$$

$$\mathbb{E}\left[Y(t)\otimes Y(t)\right] = \frac{1}{N}W(t) + \frac{1}{N^2}B(t) + O(1/N^3)$$

$$espY(t)^{\otimes 3} = \frac{1}{N^2}C(t) + O(1/N^3)$$

$$espY(t)^{\otimes 4} = \frac{1}{N^2}D(t) + O(1/N^3)$$

where

$$\frac{d}{dt}V^i = f^i_j V^j + f^i_{j,k}W^{j,k}$$

$$\frac{d}{dt}W^{j,k} = f^j_\ell W^{\ell,k} + f^k_\ell W^{j,\ell}$$

$$\frac{d}{dt}A^i = f^i_j A^j + f^i_{j,k}B^{j,k} + f^i_{j,k,\ell}C^{j,k,\ell} + f^i_{j,k,\ell,m}D^{j,k,\ell,m}$$

$$\frac{d}{dt}B^{i,j} = f^i_k B^{k,j} + f^j_k B^{k,j} + \frac{3}{2}\left[f^i_{k,\ell}C^{k,\ell,j} + f^j_{k,\ell}C^{k,\ell,i}\right] + 2(f^i_{k,\ell,m}D^{k,\ell,m,j} + f^j_{k,\ell,m}D^{k,\ell,m,i}) + \frac{1}{2}Q^{i,j}_k V^k + \frac{1}{2}Q^{i,j}_{k,\ell}W^{k,\ell}$$

$$\dots$$

Computational issues

Recall that $x(t)$ be the mean field approximation and $Y(t) = X(t) - x(t)$.

You can close the equations by assuming that $Y^{(k)} = 0$ for $k > K$.

- For $K = 0$, this gives the mean field approximation ($1/N$-accurate)
- For $K = 2$, this gives the refined mean field ($1/N^2$-accurate).
- For $K = 4$, this gives a second order expansion ($1/N^3$-accurate).

For a system of dimension d, $Y(t)^{(k)}$ has d^k equations.

Computational issues

- The mean field is a system of non-linear ODE of dimension d.
- The $1/N$ term adds two systems of **time-inhomogeneous linear** ODEs of dimension d^2 and d.
- The $1/N^2$ term adds four systems of **time-inhomogeneous linear** ODEs of dimension d^4, d^3, d^2 and d.

To compute, you essentially need up to the second (for the $1/N$-term) or the fourth (for the $1/N^2$-term) derivatives of the drifts.

Outline

The SQ(2)

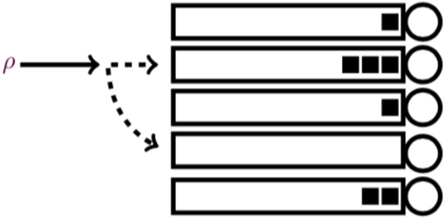

Arrival at each server ρ.

- Sample $d-1$ other queues.
- Allocate to the shortest queue

Service rate$=1$.

	$N=10$	$N=20$	$N=50$	$N=100$
Mean Field	2.3527	2.3527	2.3527	2.3527
$1/N$-expansion	2.7513	2.5520	2.4324	2.3925
$1/N^2$-expansion	2.8045	2.5653	2.4345	2.3930
Simulation	2.8003	2.5662	2.4350	2.3931

$SQ(2)$: **Steady-state average queue length** ($\rho=0.9$).

How does the expected queue length evolve with time?

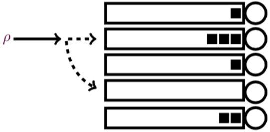

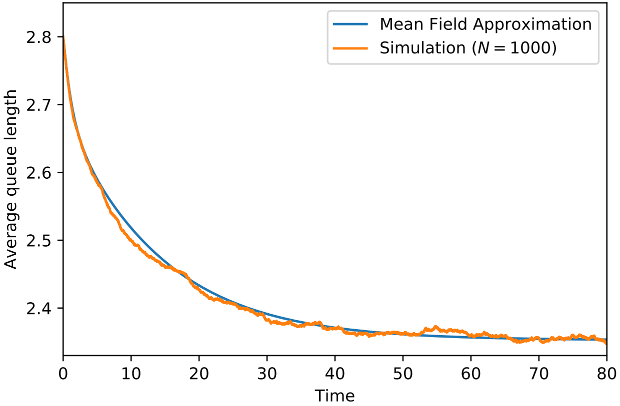

How does the expected queue length evolve with time?

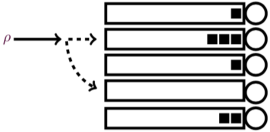

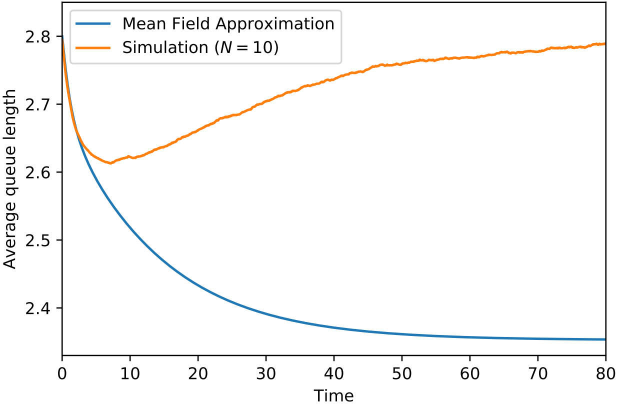

How does the expected queue length evolve with time?

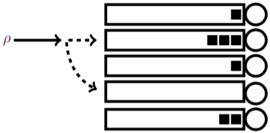

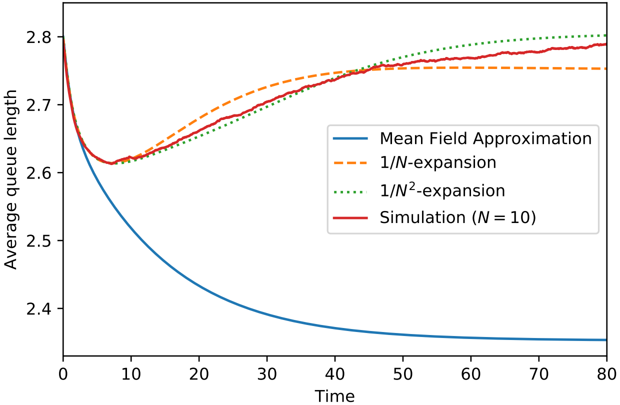

Remark about computation time :

- 10min/1h (simulation $N = 1000/N = 10$), C++ code. Requires many simulations, confidence intervals,...
- 80ms (mean field), 700ms ($1/N$-expansion), 9s ($1/N^2$-expansion), Python `numpy`

Analysis of the computation time

For the numerical examples of $SQ(2)$, I used a bounded queue size d.

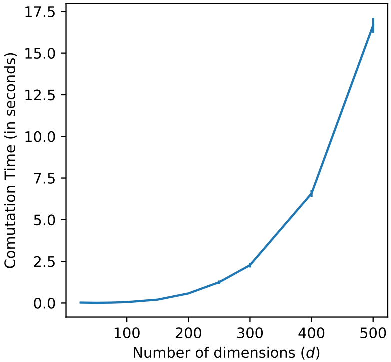

Time to compute the $1/N$-expansion

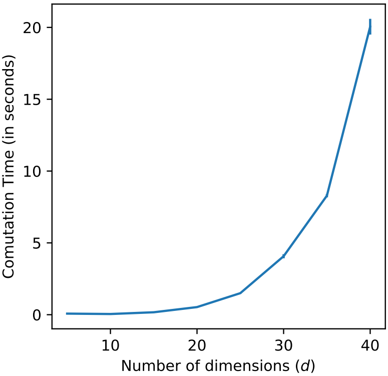

Time to compute the $1/N^2$-expansion

Analysis of the computation time (Python numpy implementation)

Does it always work?

Can I exchange the limits $N \to \infty$, $k \to \infty$, $t \to \infty$?

$$\mathbb{E}\left[X(t)\right] = x(t) + \frac{1}{N}V(t) + \frac{1}{N^2} + \cdots + O(\frac{1}{N^{k+1}})$$

Does it always work?

Can I exchange the limits $N \to \infty$, $k \to \infty$, $t \to \infty$?

$$\mathbb{E}[X(t)] = x(t) + \frac{1}{N}V(t) + \frac{1}{N^2} + \dots + O(\frac{1}{N^{k+1}})$$

NO:

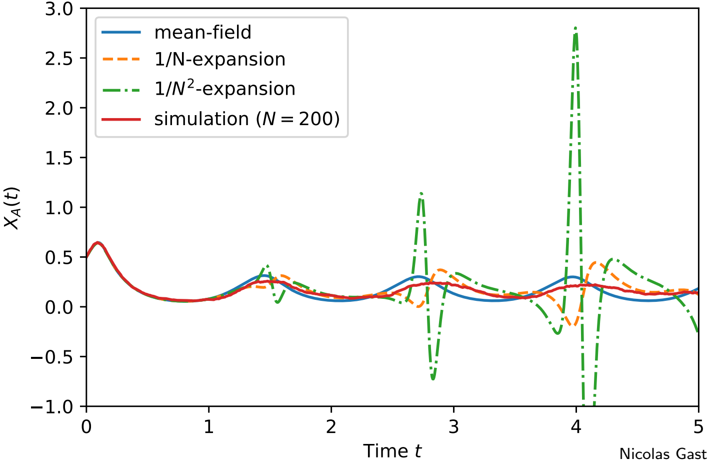

Outline

Recap and extensions

For a mean field model with four differentiable drift

$$\mathbb{E}[X(t)] = x(t) + \frac{V(t)}{N} + \frac{A(t)}{N^2} + \dots$$

- We can build expansion in $1/N$

From a computational point of view:

- The $1/N$-term involves d^2 linear equations.
- The $1/N^2$-term involves d^4 linear equations.
- Most of the gain seems to come from the $1/N$-term.

Some References

Paper (simulation, slides) is reproducible!

https://github.com/ngast/sizeExpansionMeanField/

nicolas.gast@inria.fr

http://mescal.imag.fr/membres/nicolas.gast

- A Refined Mean Field Approximation by Gast and Van Houdt. SIGMETRICS 2018 (best paper award)
- Size Expansions of Mean Field Approximation: Transient and Steady-State Analysis Gast, Bortolussi, Tribastone
- Expected Values Estimated via Mean Field Approximation are O(1/N)-accurate by Gast. SIGMETRICS 2017.
- https://github.com/ngast/rmf_tool/