A *Refined* Mean Field Approximation

Nicolas Gast   Benny Van Houdt

Inria   Univ. Antwerp

ACM SIGMETRICS, Irvine, CA, June 2018

Good system design needs accurate performance evaluation

Example : load balancing with N server

N servers

Which allocation policy?

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

Good system design needs accurate performance evaluation

Example : load balancing with N server

N servers

Which allocation policy?

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

Model with finite N is difficult to analyze.

Many systems are analyzed via mean field approximation

It can be shown that some systems simplify as N goes to infinity

"Theorem". $\lim_{N\to\infty}$ (

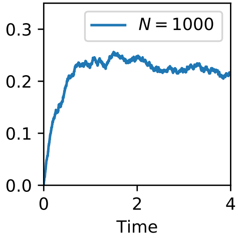

) =

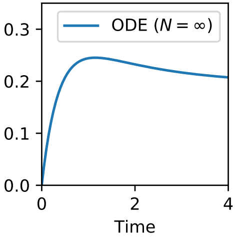

Mean field approximation

- Theoretical biology, statistical mechanics
- Game theory (Mean field games : evacuation, Mexican wave)
- Performance of computer systems : Load balancing (power of two-choice), Wireless (CSMA), Caching,...

Mean-field approximation is widely used in our community

A few examples of recent SIGMETRICS papers...

- 2018 **The PDE Method for the Analysis of Randomized Load Balancing Networks** – Aghajani et al.
- 2018 **Asymptotically Optimal Load Balancing Topologies** – Mukherjee et al.
- 2018 **On the Power-of-d-choices with Least Loaded Server Selection** – Hellemans and Van Houdt
- 2018 **Delay Scaling in Many-Sources Wireless Networks without Queue State Information** – Borst and Zubeldia
- 2017 **Analysis of a Stochastic Model of Replication in Large Distributed Storage Systems: A Mean-Field Approach** – Sun et al.
- 2017 **Optimal Service Elasticity in Large-Scale Distributed Systems** – Mukherjee et al
- 2017 **Stein's Method for Mean Field Approximations in Light and Heavy Traffic Regimes** – Ying
- 2017 **Expected Values Estimated via Mean-Field Approximation are 1/N-Accurate** – G
- 2016 **Asymptotics of Insensitive Load Balancing and Blocking Phases** – Jonckheere - Prabhu
- 2016 **On the Approximation Error of Mean-Field Models** – Ying
- 2015 **Power of d Choices for Large-Scale Bin Packing: A Loss Model** – Xie et al
- 2015 **Transient and Steady-state Regime of a Family of List-based Cache Replacement Algorithms** – G, Van Houdt
- 2014 **Data Dissemination Performance in Large-Scale Sensor Networks** – Meyfroyt et al.
- 2013 **Queueing system topologies with limited flexibility.** – Tsitsiklis, Xu
- 2013 **A mean field model for a class of garbage collection algorithms inflash-based solid state drives.** – Van Houdt
- 2012 **Fluid limit of an asynchronous optical packet switch with shared per link full range wavelength conversion.** – Van Houdt, Bortolussi
- 2011 **On the power of (even a little) centralization in distributed processing.** – Xu and Tsitsiklis
- 2010 **Randomized load balancing with general service time distributions.** – Bramson et al.
- 2010 **Incentivizing peer-assisted services: a fluid shapley value approach.** – Misra et al
- 2010 **A mean field model of work stealing in large-scale systems.** – G, Gaujal
- 2009 **The age of gossip: spatial mean field regime.** – Chaintreau et al.

Common steps in many of these papers:

1. Prove the convergence to a limit (the mean field approximation)
2. Analyze the limit
3. Evaluate numerically models with finite N.

Mean field is for (very) large systems. What about moderate sizes?

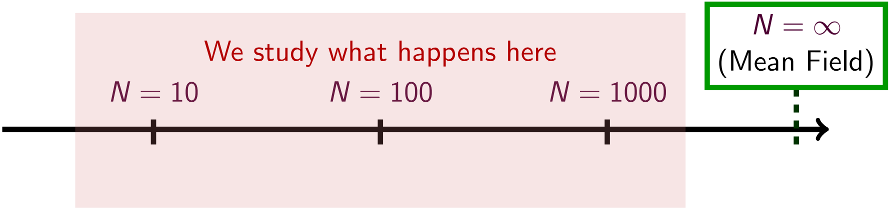

Mean field is for (very) large systems. What about moderate sizes?

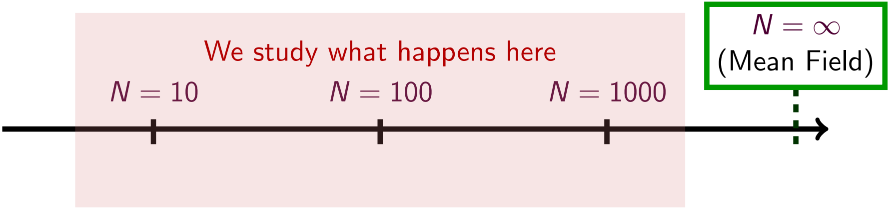

For many systems, asymptotically:

$$Perf(N) \approx \underbrace{Perf(\infty)}_{\text{Mean field approximation}} + \frac{1}{N}V$$

Mean field approximation

Refined mean field approximation: $Perf(\infty) + \frac{1}{N}V$

V can be computed by an analytical method

By studying what happens when $N \to \infty$, we get a very accurate approximation even for $N = 10$

	Coupon	Supermarket	Pull/push
Simulation ($N = 10$)	1.530	2.804	2.304
Refined mean field ($N = 10$)	1.517	2.751	2.295
Mean field ($N = \infty$)	1.250	2.353	1.636

Outline

Outline

1. Mean field and refined mean field approximations
2. Numerical experiments : how (more) accurate is the refined approximation?
3. Conclusion

We study a population of N interchangeable objects.

X denotes the empirical measure.

$$X_i(t) = \text{ fraction of objects in state } i$$

Framework: Density dependent population processes (Kurtz 70s)[1]

A population process is a sequence of CTMC $\mathbf{X}^N$, indexed by the population size N, with state spaces $\mathbf{E}^N \subset \mathbf{E}$, with initial state x_0 and with transitions (for $\ell \in \mathcal{L}$):

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

[1] Our results can also be applied to the discrete-time model of (Benaim, Le Boudec 2008).

Framework: Density dependent population processes (Kurtz 70s)[1]

A population process is a sequence of CTMC $\mathbf{X}^N$, indexed by the population size N, with state spaces $\mathbf{E}^N \subset \mathbf{E}$, with initial state x_0 and with transitions (for $\ell \in \mathcal{L}$):

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

The drift (average variation) is $f(x) = \sum_\ell \ell \beta_\ell(x)$.

The mean field approximation is : $\dot{x} = f(x)$.

[1] Our results can also be applied to the discrete-time model of (Benaim, Le Boudec 2008).

Example : supermarket model, JSQ(2)[2]

More examples in the paper

N servers

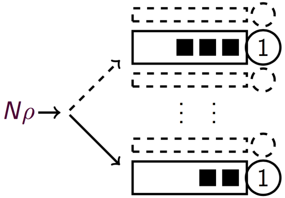

Randomly choose two, and select one

[2]Vvedenskaya et al. 96, Mitzenmacher 98.

Example : supermarket model, JSQ(2)[2]

More examples in the paper

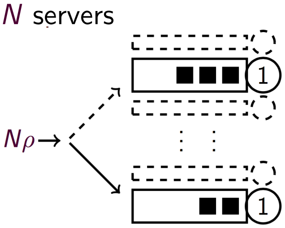

Randomly choose two, and select one

X_i = fractions of servers with i or more jobs.

The transitions are:

$$X \mapsto X + \frac{1}{N}\mathbf{e}_i \text{ at rate } N\rho(X_{i-1}^2 - X_i^2)$$

$$X \mapsto X - \frac{1}{N}\mathbf{e}_i \text{ at rate } N(x_i - x_{i+1})$$

The mean field approximation is given by the (infinite) system of ODE:

$$\dot{x}_i = \underbrace{\rho(x_{i-1}^2 - x_i^2)}_{\text{arrivals}} - \underbrace{(x_i - x_{i+1})}_{\text{departures}}$$

[2]Vvedenskaya et al. 96, Mitzenmacher 98.

Steady-state analysis : main assumptions

(A0) $\sup_x \sum_\ell |\ell|^2 |\beta_\ell(x)| < \infty$.

(A1) The stochastic process is a density dependent population process.

(A2) The drift f is twice-differentiale

(A3) The ODE has a globally stable attractor π, *i.e.*, for any solution x of the ODE $\dot{x} = f(x)$:

$$\|x(t) - \pi\| \leq Ce^{-\alpha t} \|x(0) - \pi\|.$$

(A4) For each N, the population process has a unique stationary distribution.

The constant is defined as a function of the first two derivatives of the drift at π

Let π be the fixed point of the mean field approximation and

$$A = Df(\pi) \qquad B = D^2 f(\pi) \qquad Q_{ij} = \sum_{\ell} \ell_i \ell_j \beta_\ell(\pi).$$

Let W be the unique solution of the Lyapunov equation

$$AW + (AW)^T = Q$$

THEOREM 3.1. *Assume that the model satisfies (A0–A4). Let $h : \mathcal{E} \to \mathbb{R}$ be a twice-differentiable function that has a uniformly continuous second derivative. Then,*

$$\lim_{N\to\infty} N \left(\mathbf{E}^{(N)} \left[h(X^{(N)}) \right] - h(\pi) \right) = \sum_i \frac{\partial h}{\partial x_i}(\pi) V_i + \frac{1}{2} \sum_{i,j} \frac{\partial^2 h}{\partial x_i \partial x_j}(\pi) W_{ij}, \tag{2}$$

where the matrices A, C and W are defined above and V_i is equal to:

$$V_i = -\sum_j (A^{-1})_{i,j} \left[C_j + \frac{1}{2} \sum_{k_1,k_2} (B_j)_{k_1,k_2} W_{k_1,k_2} \right]. \tag{3}$$

The constant is defined as a function of the first two derivatives of the drift at π

Let π be the fixed point of the mean field approximation and

$$A = Df(\pi) \qquad B = D^2 f(\pi) \qquad Q_{ij} = \sum_{\ell} \ell_i \ell_j \beta_\ell(\pi).$$

Let

To compute V, you need to :

- Evaluate derivatives at π
- Solve a Lyapunov equation (linear algebra)

THEOREM 3.1. *Assume that the model satisfies (A0–A4). Let $h : \mathcal{E} \to \mathbb{R}$ be a twice-differentiable function that has a uniformly continuous second derivative. Then,*

$$\lim_{N\to\infty} N \left(\mathbf{E}^{(N)} \left[h(X^{(N)}) \right] - h(\pi) \right) = \sum_i \frac{\partial h}{\partial x_i}(\pi) V_i + \frac{1}{2} \sum_{i,j} \frac{\partial^2 h}{\partial x_i \partial x_j}(\pi) W_{ij}, \tag{2}$$

where the matrices A, C and W are defined above and V_i is equal to:

$$V_i = -\sum_j (A^{-1})_{i,j} \left[C_j + \frac{1}{2} \sum_{k_1,k_2} (B_j)_{k_1,k_2} W_{k_1,k_2} \right]. \tag{3}$$

Main ideas of the proof

Stein's method (1); comparison of generators; (2); perturbation theory (3).

Let G_h be the function $G_h(x) = \int_0^\infty (h(\Phi_t(x)) - h(\pi))dt$, where $\Phi_t(x)$ is the solution of the ODE $\dot{x} = f(x)$ starting in x at time 0.

$$
\begin{aligned}
N\mathbb{E}\left[h(X^N) - h(\pi)\right] &= N\mathbb{E}\left[\Lambda G_h)(X^N)\right] \\
&= N\mathbb{E}\left[(\Lambda - L^{(N)})(G_h)(X^N)\right] && (1) \\
&= \frac{1}{2}\mathbb{E}\left[\sum_\ell \beta_\ell(X^N) D^2 G_h(X^N)\cdot(\ell,\ell)\right] + O(\frac{1}{N}) && (2) \\
&\to \frac{1}{2}\sum_\ell \beta_\ell(\pi) D^2 G_h(\pi)\cdot(\ell,\ell). && (3)
\end{aligned}
$$

The computation of $D^2 G_h(\pi)$ gives you the result (perturbation theory).

Main ideas of the proof

Stein's method (1); comparison of generators; (2); perturbation theory (3).

Let G_h be the function $G_h(x) = \int_0^\infty (h(\Phi_t(x)) - h(\pi))dt$, where $\Phi_t(x)$ is the solution of the ODE $\dot{x} = f(x)$ starting in x at time 0.

$$
\begin{aligned}
N\mathbb{E}\left[h(X^N) - h(\pi)\right] &= N\mathbb{E}\left[\Lambda G_h)(X^N)\right] \\
&= N\mathbb{E}\left[(\Lambda - L^{(N)})(G_h)(X^N)\right] && (1)\\
&= \frac{1}{2}\mathbb{E}\left[\sum_\ell \beta_\ell(X^N) D^2 G_h(X^N)\cdot(\ell,\ell)\right] + O(\frac{1}{N}) && (2)\\
&\to \frac{1}{2}\sum_\ell \beta_\ell(\pi) D^2 G_h(\pi)\cdot(\ell,\ell). && (3)
\end{aligned}
$$

The computation of $D^2 G_h(\pi)$ gives you the result (perturbation theory).

Outline

How hard is the computation of the refined model?

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

How to compute V?

- V can sometimes be computed in closed (not often)
- Numerical evaluation is easy (linear algebra)
  https://github.com/ngast/rmf_tool/

The supermarket model (JSQ(2))

N	10	20	30	50	100	∞
$\rho = 0.7$						
Simulation	1.2194	1.1735	1.1584	1.1471	1.1384	–
Refined mf	1.2150	1.1726	1.1584	1.1471	1.1386	1.1301
$\rho = 0.9$						
Simulation	2.8040	2.5665	2.4907	2.4344	2.3931	–
Refined mf	2.7513	2.5520	2.4855	2.4324	2.3925	2.3527
$\rho = 0.95$						
Simulation	4.2952	3.7160	3.5348	3.4002	3.3047	–
Refined mf	4.1017	3.6578	3.5098	3.3915	3.3027	3.2139

Two-choice model : Average queue length for various values of ρ and N. We compare simulation with the refined mean field approximation

The supermarket model (JSQ(2))

N	10	20	30	50	100	∞
$\rho = 0.7$						
Simulation	1.2194	1.1735	1.1584	1.1471	1.1384	–
Refined mf	1.2150	1.1726	1.1584	1.1471	1.1386	1.1301
$\rho = 0.9$						
Simulation	2.8040	2.5665	2.4907	2.4344	2.3931	–
Refined mf	2.7513	2.5520	2.4855	2.4324	2.3925	2.3527
$\rho = 0.95$						
Simulation	4.2952	3.7160	3.5348	3.4002	3.3047	–
Refined mf	4.1017	3.6578	3.5098	3.3915	3.3027	3.2139

Two-choice model : Average queue length for various values of ρ and N.
We compare simulation with the refined mean field approximation

The supermarket model (JSQ(2))

N	10	20	30	50	100	∞
$\rho = 0.7$						
Simulation	1.2194	1.1735	1.1584	1.1471	1.1384	–
Refined mf	1.2150	1.1726	1.1584	1.1471	1.1386	1.1301
$\rho = 0.9$						
Simulation	2.8040	2.5665	2.4907	2.4344	2.3931	–
Refined mf	2.7513	2.5520	2.4855	2.4324	2.3925	2.3527
$\rho = 0.95$						
Simulation	4.2952	3.7160	3.5348	3.4002	3.3047	–
Refined mf	4.1017	3.6578	3.5098	3.3915	3.3027	3.2139

Mean field approximation

Two-choice model : Average queue length for various values of ρ and N.
We compare simulation with the refined mean field approximation

Pull-push model (servers with ≥ 2 jobs push to empty)

N	10	20	50	100	∞
$\rho = 0.8$					
Simulation	1.5569	1.4438	1.3761	1.3545	–
Refined mean field	1.5473	1.4403	1.3761	1.3547	1.3333
$\rho = 0.90$					
Simulation	2.3043	1.9700	1.7681	1.7023	–
Refined mean field	2.2945	1.9654	1.7680	1.7022	1.6364
$\rho = 0.95$					
Simulation	3.4288	2.6151	2.1330	1.9720	–
Refined mean field	3.4369	2.6232	2.1350	1.9723	1.8095

Mean field approximation

Push-pull model : Mean queue length under pull/push with $r = 1/(1-\rho)$: simulation vs refined mean field approximation

Comparison of policy

Is pull-push or JSQ(2) better for $\rho = 0.9$ and $N = 10$?

- Mean field predicts that pull-push reduces the average queue length by 30%.
- Refined mean field predicts : the reduction is only 17%.
- Simulation : the reduction is about 16.5%.

Other example of result : the impact of choosing with or without replacement (power of two-choice, $N = 10$ servers)

$$\Delta\{\text{Avg queue length (with-without)}\} \approx \frac{1}{N}\sum_{j=1}^{\infty}\sum_{i=1}^{\infty}(\rho^{2^{i+j}-2^j} - \rho^{2^{i+j}-1})2^{i-1}$$

Other example of result : the impact of choosing with or without replacement (power of two-choice, $N = 10$ servers)

$$\Delta\{\text{Avg queue length (with-without)}\} \approx \frac{1}{N}\sum_{j=1}^{\infty}\sum_{i=1}^{\infty}(\rho^{2^{i+j}-2^j} - \rho^{2^{i+j}-1})2^{i-1}$$

		Simulation	Refined mean field	Mean field
$\rho = 0.7$	with	1.215	1.215	1.1301
	without	1.173	1.169	1.1301
	with-without	0.042	0.046	–
$\rho = 0.9$	with	2.820	2.751	2.3527
	without	2.705	2.630	2.3527
	with-without	0.115	0.121	–
$\rho = 0.95$	with	4.340	4.102	3.2139
	without	4.169	3.923	3.2139
	with-without	0.171	0.179	–

Outline

Recap

1. We can use the rate of convergence to define a refined approximation. The main ideas are:
   - The mean field approximation is $x = \lim_{N\to\infty} X^N$
   - Using linear algebra, we can compute $V = \lim_{N\to\infty} N(X^N - \pi)$
   - The refined approximation is $x + V/N$.

2. The refined approximation is often very accurate even for $N = 10$:

	Coupon	Supermarket	Pull/push
Simulation ($N = 10$)	1.530	2.804	2.304
Refined mean field ($N = 10$)	1.517	2.751	2.295
Mean field ($N = \infty$)	1.250	2.353	1.636

Potential applications

- More examples in the paper.
- Variant of this model can be studied
- Application to queuing systems
- Some assumptions can be relaxed

Main references :

- A Refined Mean Field Approximation by G and Van Houdt. SIGMETRICS 2018 https://hal.inria.fr/hal-01622054/ https://github.com/ngast/rmf_tool/
- Expected Values Estimated via Mean Field Approximation are O(1/N)-accurate by G SIGMETRICS 2017. https://github.com/ngast/meanFieldAccuracy