Refinements of Mean Field Approximation

HDR defense – Nicolas Gast

Inria

January 30, 2020. Grenoble

The objective of my work is to provide tools to:

- Describe distributed systems;
- Optimize their behavior.

Good system design needs performance evaluation.

Tools:

- Stochastic modeling, Markov Chains.
- Dynamical systems.
- Optimization, optimal control.

Example 1: Networking and Congestion Control

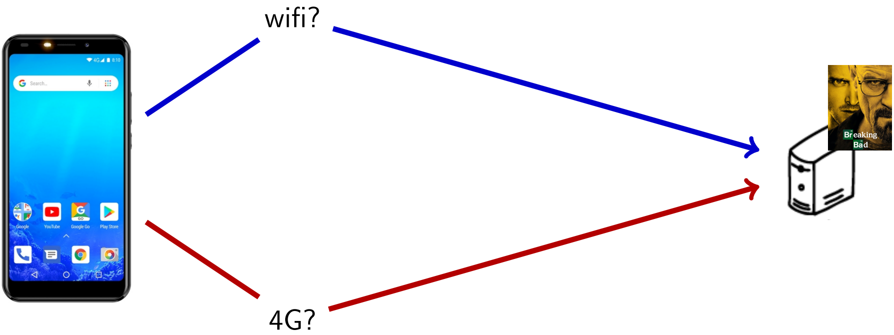

[1] "MPTCP is not pareto-optimal: performance issues and a possible solution". (CoNext 2012 (best paper), ToN 2013) by Khalili, Gast, Popovic, and Le Boudec.

Example 1: Networking and Congestion Control

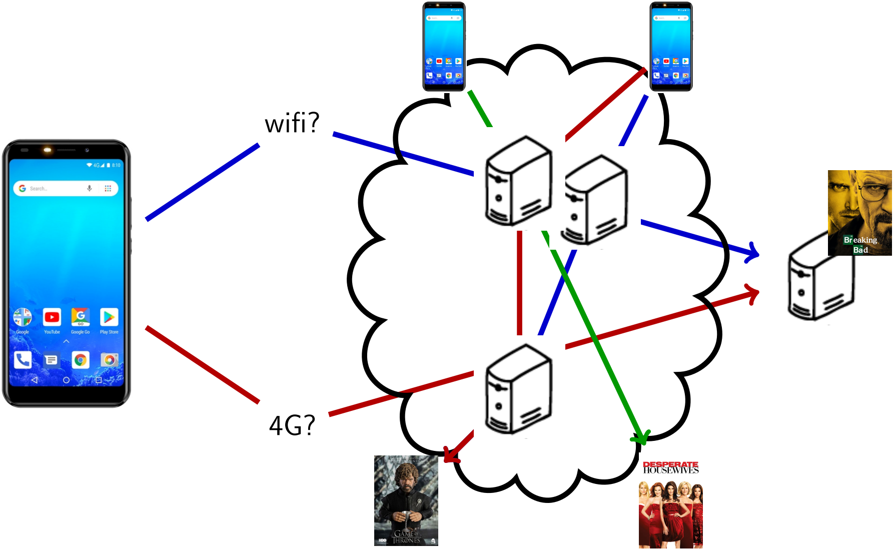

[1] “MPTCP is not pareto-optimal: performance issues and a possible solution”. (CoNext 2012 (best paper), ToN 2013) by Khalili, Gast, Popovic, and Le Boudec.

Example 2: Load balancing

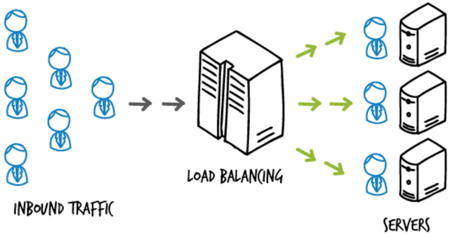

N servers

Which allocation policy?

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

[2] A mean field model of work stealing in large-scale systems (Sigmetrics 2010) by Gast and Gaujal, A refined mean field approximation (Sigmetrics 2018 (best paper)) by Gast and Van Houdt

Example 2: Load balancing

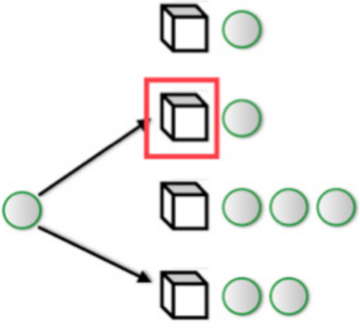

N servers

Which allocation policy?

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

[2] A mean field model of work stealing in large-scale systems (Sigmetrics 2010) by Gast and Gaujal, A refined mean field approximation (Sigmetrics 2018 (best paper)) by Gast and Van Houdt

This leads us to study markovian models of interacting agents.

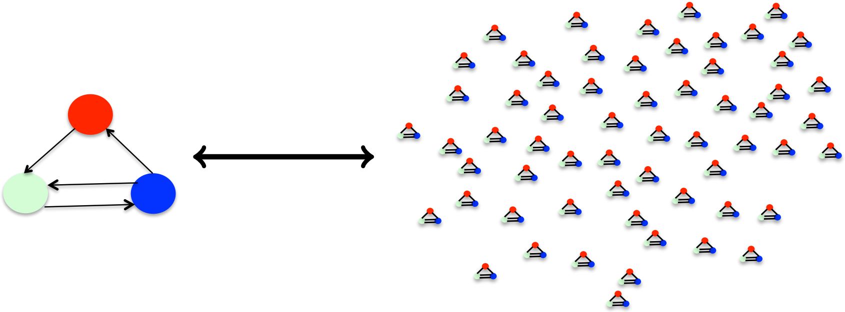

This leads us to study markovian models of interacting agents.

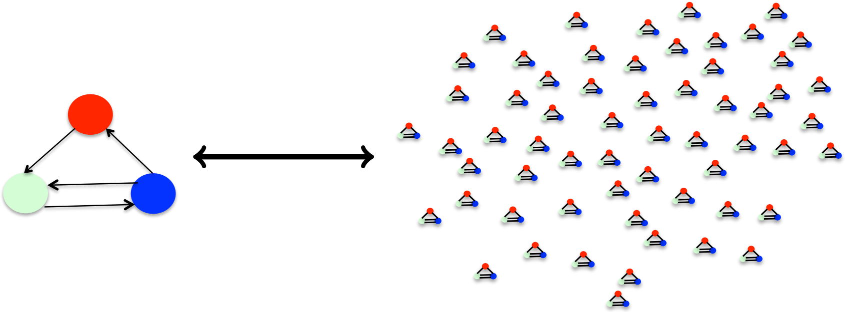

Problem: state space explosion.
S states per object, N objects $\Rightarrow$ S^N states.

In my work, I develop and use models that scale.

Today, I will mostly focus on mean field approximation.

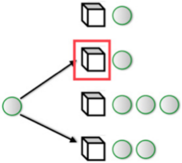

$N = 10?$ $\quad$ $N = 100$ $\quad$ $N = 1000?$ $\quad$ N servers

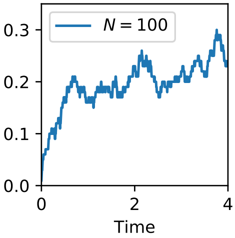

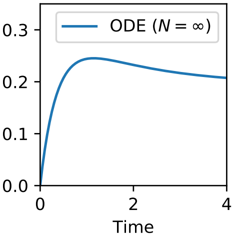

Overview of my contributions

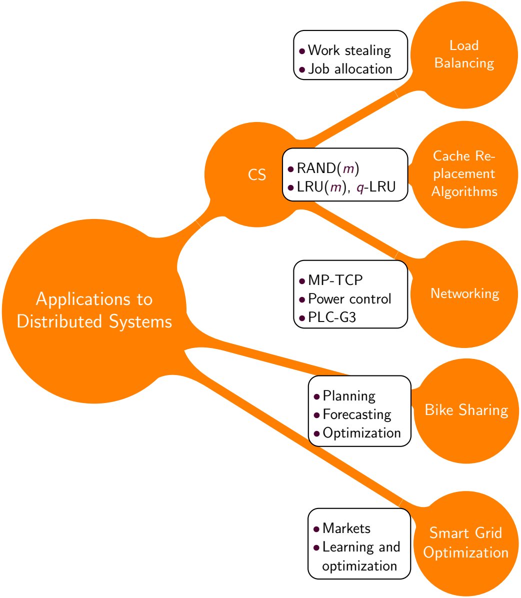

Overview of my contributions

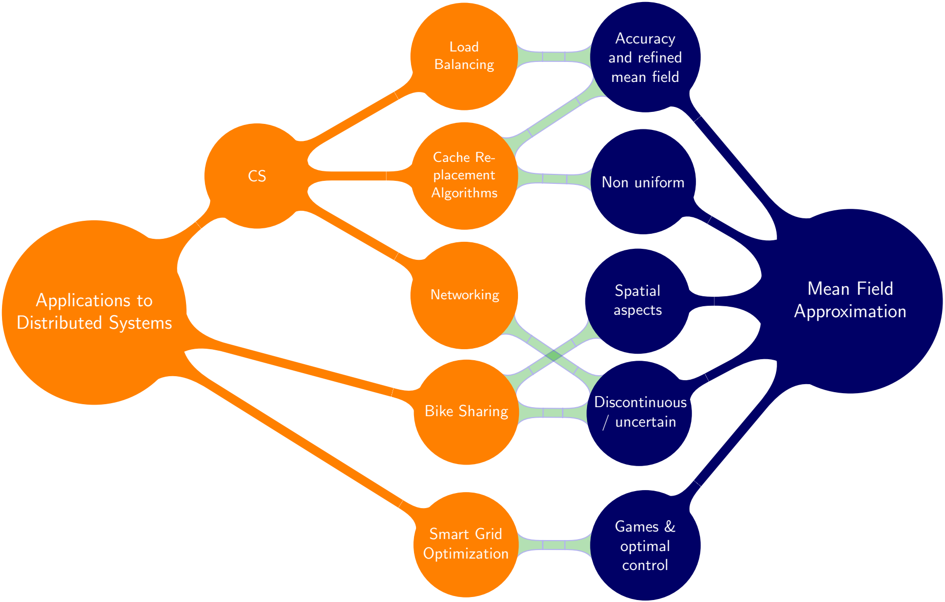

Overview of my contributions

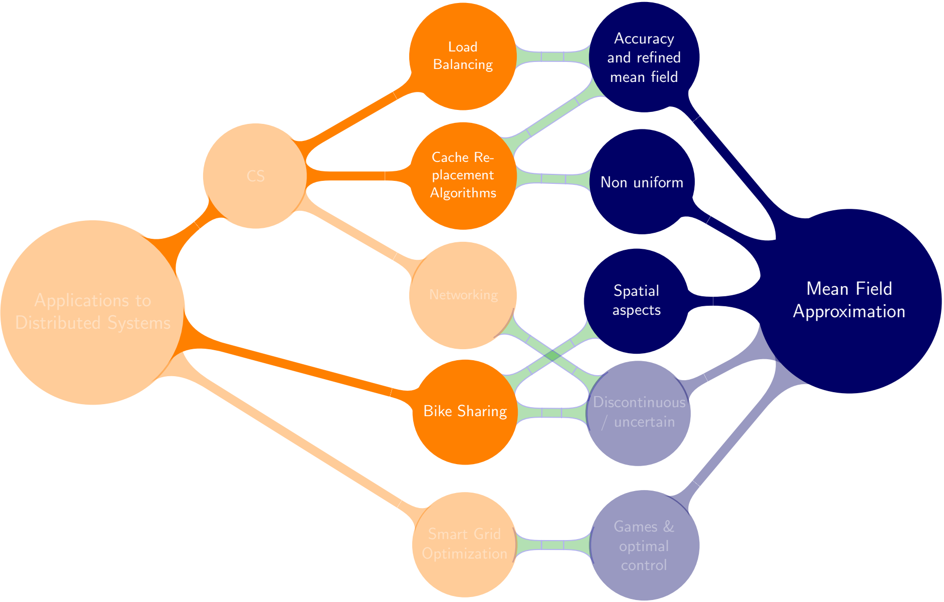

Overview of my contributions

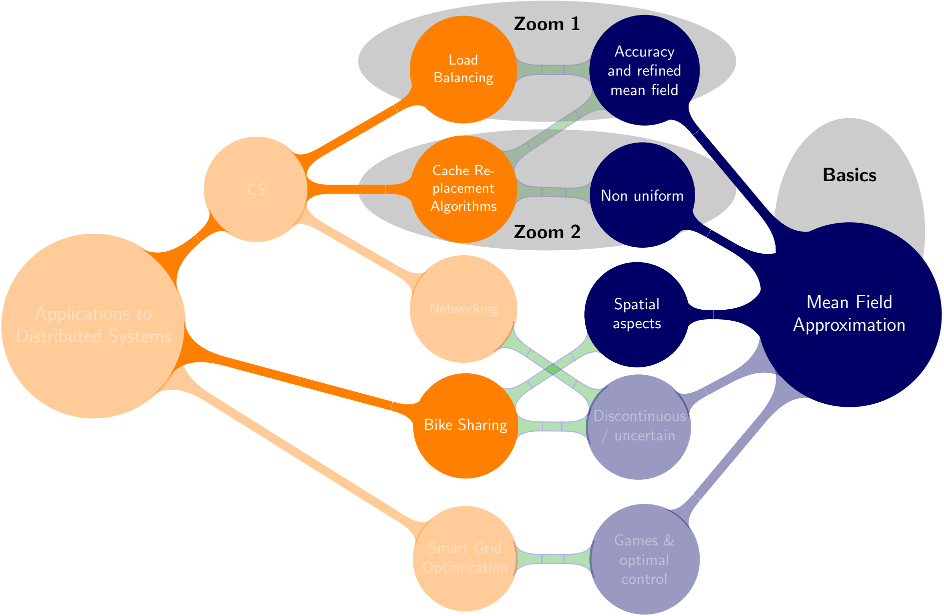

Idea of mean field approximation

Each individual interacts with the mass while having a negligible effect on the mass.

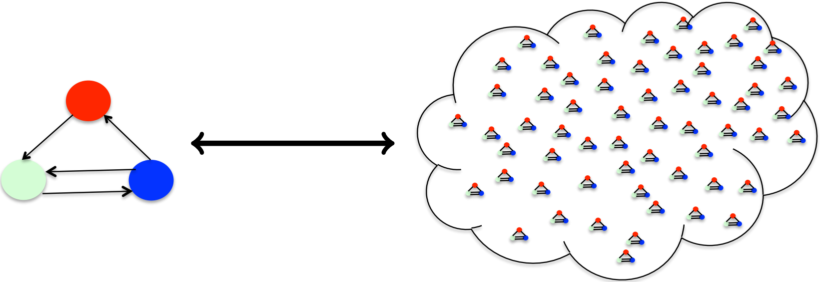

- Mean field theory (Statistical mechanics, 1800s-1900s)
- Theoretical biology
- Computer modeling (ex: Baccelli 92, Vvedenskaya 96)
- Mean field games (Lasry-Lions 2007)

Idea of mean field approximation

Each individual interacts with the mass while having a negligible effect on the mass.

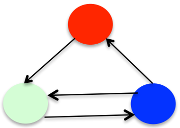

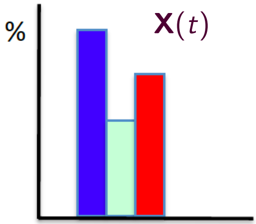

- Mean field theory (Statistical mechanics, 1800s-1900s)
- Theoretical biology
- Computer modeling (ex: Baccelli 92, Vvedenskaya 96)
- Mean field games (Lasry-Lions 2007)

Example: The supermarket model (SQ(2))

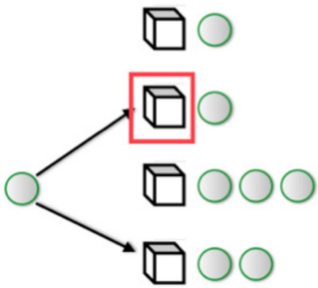

Randomly choose two, select one.

N identical queues.

- Arrival rate $N\lambda$;
- Service rate μ.

State = queue sizes.

$$Q(t) = (1, 1, 2, 3).$$

Example: The supermarket model (SQ(2))

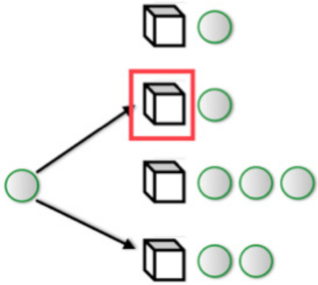

N identical queues.

- Arrival rate $N\lambda$;
- Service rate μ.

State = queue sizes.

$$Q(t) = (1, 1, 2, 3).$$

If all queues are identical, we can simplify: $X = (X_1, X_2, \dots)$

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

Above: $X(t) = (1, .5, .25, 0, 0, 0, \dots)$.

$SQ(2)$: Transitions and Mean Field Approximation.

The fractions of queue with i jobs or more changes when:

$$\text{A job arrives: } X_i \mapsto X_i + \frac{1}{N} \qquad \text{(at rate } N\lambda(X_{i-1}^2 - X_i^2))$$

$$\text{A job departs: } X_i \mapsto X_i - \frac{1}{N} \qquad \text{(at rate } N\mu(X_i - X_{i+1})).$$

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

$$\dot{x}_i = \underbrace{\lambda(x_{i-1}^2 - x_i^2)}_{\text{Arrival}} - \underbrace{\mu(x_i - x_{i+1})}_{\text{Departure}}$$

$SQ(2)$: Transitions and Mean Field Approximation.

The fractions of queue with i jobs or more changes when:

A job arrives: $X_i \mapsto X_i + \frac{1}{N}$ (at rate $N\lambda(X_{i-1}^2 - X_i^2)$)

A job departs: $X_i \mapsto X_i - \frac{1}{N}$ (at rate $N\mu(X_i - X_{i+1})$).

$x \mapsto x + \frac{1}{N}\ell$

Rate $Nr(\ell, x)$

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

$$\dot{x}_i = \underbrace{\lambda(x_{i-1}^2 - x_i^2)}_{\text{Arrival}} - \underbrace{\mu(x_i - x_{i+1})}_{\text{Departure}}$$

Drift $f(x) = \sum_\ell \ell r(\ell, x)$.

This is a density dependent population process.

Mean field is asymptotically exact

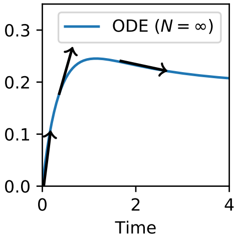

$N = 10$ $N = 100$ $N = 1000$ $N = \infty$

Mean field is asymptotically exact

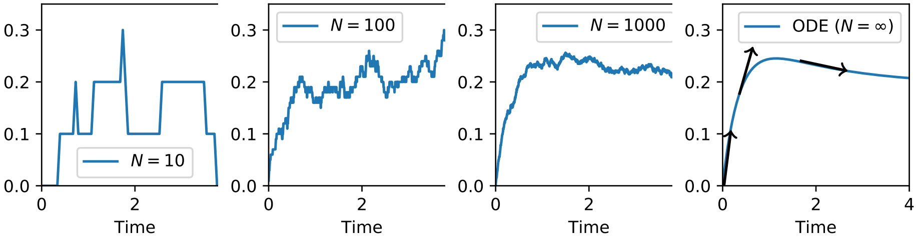

$N = 10$ $N = 100$ $N = 1000$ $N = \infty$

The stochastic approximation approach:

$$X\left(t + \frac{1}{N}\right) = X(t) + \frac{1}{N}\Big(f(X(t)) + \text{noise}\Big).$$

Mean field is asymptotically exact

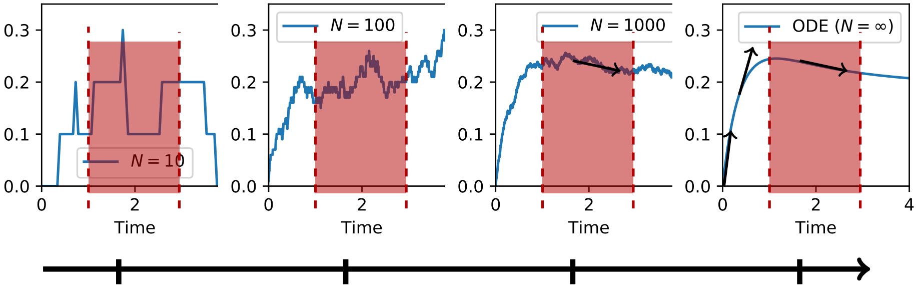

$N = 10$ $N = 100$ $N = 1000$ $N = \infty$

The stochastic approximation approach:

$$X\left(t + \frac{1}{N}\right) = X(t) + \frac{1}{N}\left(f(X(t)) + \text{noise}\right).$$

A few contributions

Mean field approximation is exact as $N \to \infty$. We extended this methodology to ...

- ...study discontinuous or imprecise systems;
  - Differential inclusions [PEVA 2012, DSN 2016]
  - Numerical algorithms. [QEST 2019]
- ...simplify optimal control problems.
  - Discrete time [DEDS 2011]
  - Continuous time [TAC 2012]
  - Mean field game [JDG 2019]

Key question

How large should N be?

A few contributions

Mean field approximation is exact as $N \to \infty$. We extended this methodology to ...

- ...study discontinuous or imprecise systems;
  - Differential inclusions [PEVA 2012, DSN 2016]
  - Numerical algorithms. [QEST 2019]
- ...simplify optimal control problems.
  - Discrete time [DEDS 2011]
  - Continuous time [TAC 2012]
  - Mean field game [JDG 2019]

Key question

How large should N be?

A few contributions

Mean field approximation is exact as $N \to \infty$. We extended this methodology to ...

- ...study discontinuous or imprecise systems;
  - Differential inclusions [PEVA 2012, DSN 2016]
  - Numerical algorithms. [QEST 2019]

- ...simplify optimal control problems.
  - Discrete time [DEDS 2011]
  - Continuous time [TAC 2012]
  - Mean field game [JDG 2019]

Key question

How large should N be?

Outline

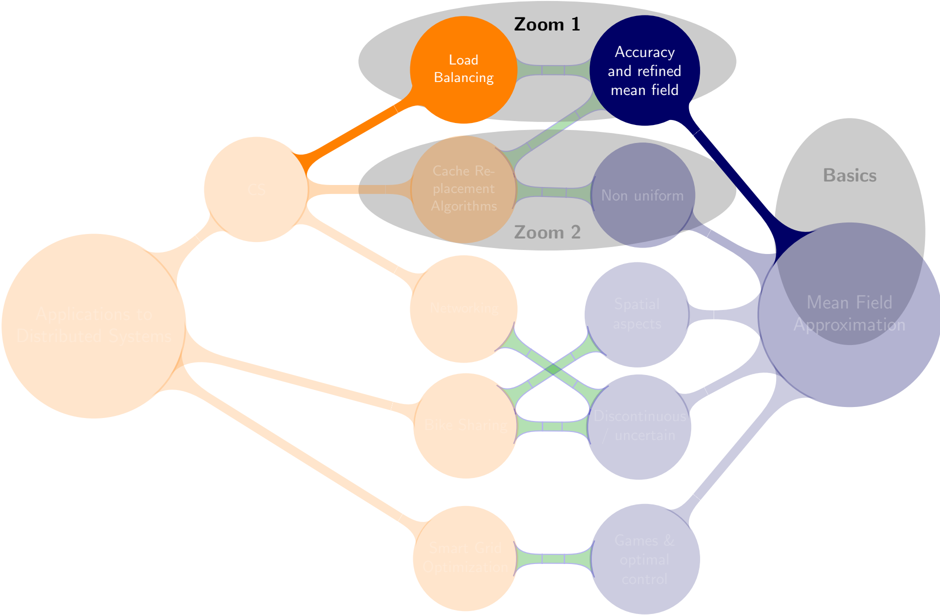
Zoom 1
Load Balancing
Accuracy and refined mean field
CS
Cache Replacement Algorithms
Non uniform
Basics
Zoom 2
Networking
Spatial aspects
Mean Field Approximation
Applications to Distributed Systems
Bike Sharing
Discontinuous / uncertain
Smart Grid Optimization
Games & optimal control

Mean field approximation is a law of large number

Theorem (Kurtz 1970s…Ying 2016) For homogeneous and regular models

$$\mathbb{E}\left[\|X(t) - x(t)\|\right] = O(1/\sqrt{N})$$

where x is the solution of the ODE $\dot{x} = f(x)$ and f is the drift.
(Valid for $t > 0$ and $t = +\infty$ if exponentially stable attractor.)

Mean field approximation is a law of large number

Theorem (Kurtz 1970s…Ying 2016) For homogeneous and regular models

$$\mathbb{E}\left[\|X(t) - x(t)\|\right] = O(1/\sqrt{N})$$

where x is the solution of the ODE $\dot{x} = f(x)$ and f is the drift.
(Valid for $t > 0$ and $t = +\infty$ if exponentially stable attractor.)

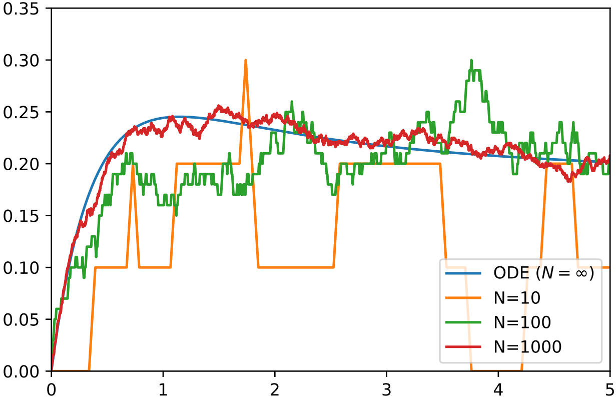

In practice, mean field works well even for "small" systems

Why?

N	10	100	1000	$+\infty$
Average queue length for SQ(2)	3.81	3.39	3.36	3.35

In practice, mean field works well even for "small" systems

Why?

N	10	100	1000	$+\infty$
Average queue length for SQ(2)	3.81	3.39	3.36	3.35

$$\mathbb{E}\left[\|X - x\|\right] = \Theta(1/\sqrt{N})$$

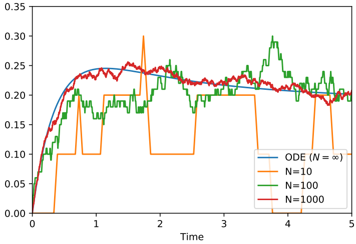

In practice, mean field works well even for "small" systems

Why?

N	10	100	1000	$+\infty$
Average queue length for SQ(2)	3.81	3.39	3.36	3.35
Error	0.45	0.039	0.004	0

$$\mathbb{E}\left[\|X - x\|\right] = \Theta(1/\sqrt{N}) \quad \gg \quad \|\mathbb{E}\left[X\right] - x\| = \Theta(1/N)$$

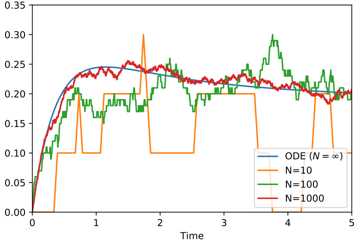

Often, what matters for performance is the bias $\mathbb{E}\left[X\right] - x$.

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

Theorem (Kolokoltsov 2012, G. 2017, G. and Van Houdt 2018). For a density dependent population process X, i f the drift f is twice differentiable, then for any $t > 0$:

1. There exists a (deterministic) vector $V(t)$ such that:

$$\mathbb{E}[X(t)] = \underbrace{x(t) + \frac{V(t)}{N}}_{\text{refined mean field}} + O(1/N^2)$$

2. $V(t)$ can be easily computed numerically

If it has a unique exponentially stable attractor, this is true for $t = \infty$.

Refined mean field is designed to study finite system

$N = 10$ $\quad$ $N = 100$ $\quad$ $N = 1000$ $\quad$ $N = \infty$

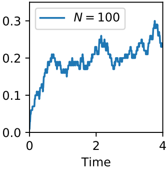

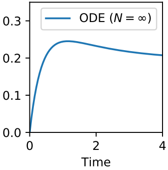

Mean field: $\quad Perf(N) \quad \approx \quad Perf(\infty)$

Refined mean field is designed to study finite system

$N = 10$ $N = 100$ $N = 1000$ $N = \infty$

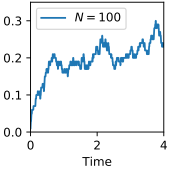

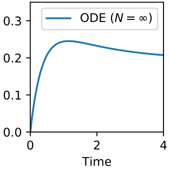

Mean field: $Perf(N) \approx Perf(\infty)$

Refined Mean field: $Perf(N) \approx \frac{V}{N} + Perf(\infty)$

Our results on refined mean field methods

- Mean field is $1/N$-accurate [Sigmetrics 2017]
  - Stein's method (inspired by Braverman et al 2015 / Ying 2016)
- Refinement for steady-state [Sigmetrics 2018[3]]
  - It is very accurate even for $N = 10$.
  - We can quantify $1/N$ variants such as choosing with/without replacement.
- Extension to finite-time and $1/N^2$ approximation [Performance 2018]
  - Moment-closure approach, tensor decomposition
  - Numerical tool[4]
- Synchronous systems [PEVA 2019]

[3] best paper award

[4] https://github.com/ngast/rmf_tool

Our results on refined mean field methods

- Mean field is $1/N$-accurate [Sigmetrics 2017]
  - Stein's method (inspired by Braverman et al 2015 / Ying 2016)
- Refinement for steady-state [Sigmetrics 2018[3]]
  - It is very accurate even for $N = 10$.
  - We can quantify $1/N$ variants such as choosing with/without replacement.
- Extension to finite-time and $1/N^2$ approximation [Performance 2018]
  - Moment-closure approach, tensor decomposition
  - Numerical tool[4]
- Synchronous systems [PEVA 2019]

[3] best paper award

[4] https://github.com/ngast/rmf_tool

Our results on refined mean field methods

- Mean field is $1/N$-accurate [Sigmetrics 2017]
  - Stein's method (inspired by Braverman et al 2015 / Ying 2016)
- Refinement for steady-state [Sigmetrics 2018[3]]
  - It is very accurate even for $N = 10$.
  - We can quantify $1/N$ variants such as choosing with/without replacement.
- Extension to finite-time and $1/N^2$ approximation [Performance 2018]
  - Moment-closure approach, tensor decomposition
  - Numerical tool[4]
- Synchronous systems [PEVA 2019]

[3] best paper award

[4] https://github.com/ngast/rmf_tool

Our results on refined mean field methods

- Mean field is $1/N$-accurate [Sigmetrics 2017]
  - Stein's method (inspired by Braverman et al 2015 / Ying 2016)
- Refinement for steady-state [Sigmetrics 2018[3]]
  - It is very accurate even for $N = 10$.
  - We can quantify $1/N$ variants such as choosing with/without replacement.
- Extension to finite-time and $1/N^2$ approximation [Performance 2018]
  - Moment-closure approach, tensor decomposition
  - Numerical tool[4]
- Synchronous systems [PEVA 2019]

[3] best paper award

[4] https://github.com/ngast/rmf_tool

Our results on refined mean field methods

- Mean field is $1/N$-accurate [Sigmetrics 2017]
  - Stein's method (inspired by Braverman et al 2015 / Ying 2016)
- Refinement for steady-state [Sigmetrics 2018[3]]
  - It is very accurate even for $N = 10$.
  - We can quantify $1/N$ variants such as choosing with/without replacement.
- Extension to finite-time and $1/N^2$ approximation [Performance 2018]
  - Moment-closure approach, tensor decomposition
  - Numerical tool[4]
- Synchronous systems [PEVA 2019]

[3] best paper award

[4] https://github.com/ngast/rmf_tool

The approximation is very accurate, even for $N = 10$.

Average queue length for supermarket SQ(2) model.

N	10	20	30	50	100	∞
$\rho = 0.9$						
Simulation ("exact")	2.804	2.567	2.491	2.434	2.393	–
Refined mean field	2.751	2.552	2.486	2.432	2.393	2.353
Error	0.053	0.015	0.005	0.002	$<$0.001	

The approximation is very accurate, even for $N = 10$.

Average queue length for supermarket SQ(2) model.

N	10	20	30	50	100	∞
$\rho = 0.9$						**Mean field approximation**
Simulation ("exact")	2.804	2.567	2.491	2.434	2.393	
Refined mean field	2.751	2.552	2.486	2.432	2.393	2.353
Error	0.053	0.015	0.005	0.002	<0.001	

The accuracy of the classical mean field degrades as ρ approaches 1:

$$\text{average queue length} \approx \log_2 \frac{1}{1-\rho} + \frac{1}{N}\frac{\rho^2}{2(1-\rho)} + \frac{1}{N^2}\frac{1}{20(1-\rho)^2}$$

Where does the $1/N$ comes from?

The moment closure approach

Consider a system for which X becomes $X + 1/N$ at rate NX^2. We have:

$$\frac{d}{dt}\mathbb{E}[X] = \mathbb{E}\left[X^2\right]$$

This equation is not closed

Where does the $1/N$ comes from?

The moment closure approach

Consider a system for which X becomes $X + 1/N$ at rate NX^2. We have:

$$\frac{d}{dt}\mathbb{E}[X] = \mathbb{E}\left[X^2\right] \qquad \approx \mathbb{E}[X]^2 \text{ (mean field approx.)}$$

This equation is ~~not~~ closed

Where does the $1/N$ comes from?

The moment closure approach

Consider a system for which X becomes $X + 1/N$ at rate NX^2. We have:

$$\frac{d}{dt}\mathbb{E}[X] = \mathbb{E}\left[X^2\right] \qquad \cancel{\approx \mathbb{E}[X]^2 \text{ (mean field approx.)}}$$

$$\frac{d}{dt}\mathbb{E}\left[X^2\right] = 2\mathbb{E}\left[X^3\right] + \frac{1}{N}\mathbb{E}\left[X^2\right]$$

This equation is not closed

Where does the $1/N$ comes from?

The moment closure approach

Consider a system for which X becomes $X + 1/N$ at rate NX^2. We have:

$$\frac{d}{dt}\mathbb{E}[X] = \mathbb{E}[X^2] \qquad \cancel{\approx \mathbb{E}[X]^2 \text{ (mean field approx.)}}$$

$$\frac{d}{dt}\mathbb{E}[X^2] = 2\mathbb{E}[X^3] + \frac{1}{N}\mathbb{E}[X^2] \qquad \approx 2(3\mathbb{E}[X^2]\mathbb{E}[X] - 2\mathbb{E}[X]^2) + \frac{1}{N}\mathbb{E}[X^2]$$

This equation is ~~not~~ closed

Where does the $1/N$ comes from?

The moment closure approach

Consider a system for which X becomes $X + 1/N$ at rate NX^2. We have:

$$\frac{d}{dt}\mathbb{E}[X] = \mathbb{E}\left[X^2\right] \qquad \cancel{\approx \mathbb{E}[X]^2 \text{ (mean field approx.)}}$$

$$\frac{d}{dt}\mathbb{E}\left[X^2\right] = 2\mathbb{E}\left[X^3\right] + \frac{1}{N}\mathbb{E}\left[X^2\right] \qquad \cancel{\approx 2(3\mathbb{E}\left[X^2\right]\mathbb{E}[X] - 2\mathbb{E}[X]^2) + \frac{1}{N}\mathbb{E}\left[X^2\right]}$$

$$\frac{d}{dt}\mathbb{E}\left[X^3\right] = \mathbb{E}\left[\frac{3X^4}{N} + \frac{4X^3}{N^2} + \frac{X^2}{N^3}\right]$$

$$\vdots$$

These equations are never closed.

- They can be closed by assuming $\mathbb{E}\left[(X - \mathbb{E}[X])^d\right] \approx 0$
- This gives a $O(1/N^{\lfloor (d+1)/2 \rfloor})$-accurate approximation.

Outline

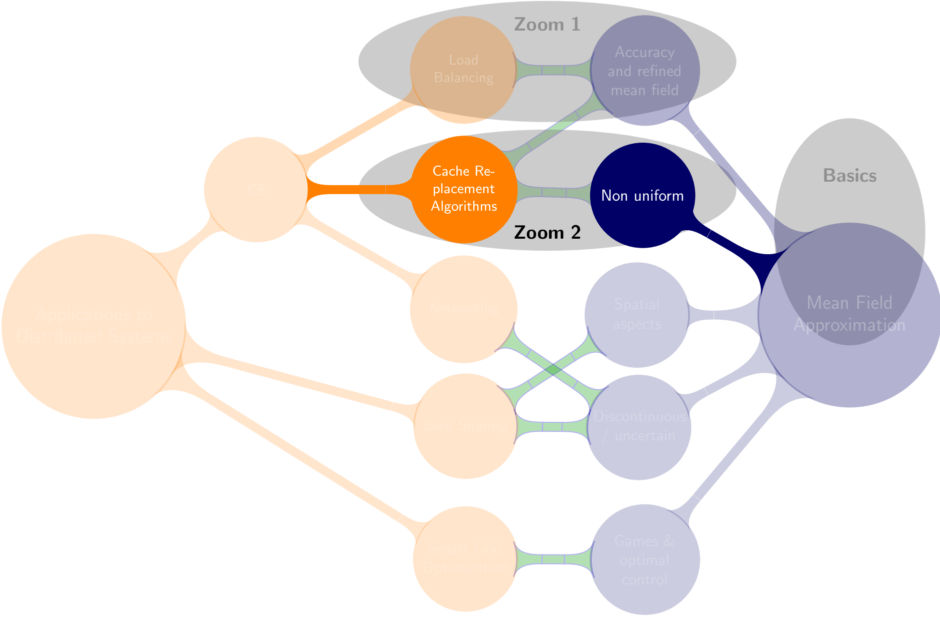

Zoom 1
Load Balancing
Accuracy and refined mean field
Basics
CS
Cache Replacement Algorithms
Non uniform
Zoom 2
Applications to Distributed Systems
Networking
Spatial aspects
Mean Field Approximation
Bike Sharing
Discontinuous / uncertain
Smart Grid Optimization
Games & optimal control

Analysis of Cache (Re)placement Policies

Analysis of Cache (Re)placement Policies

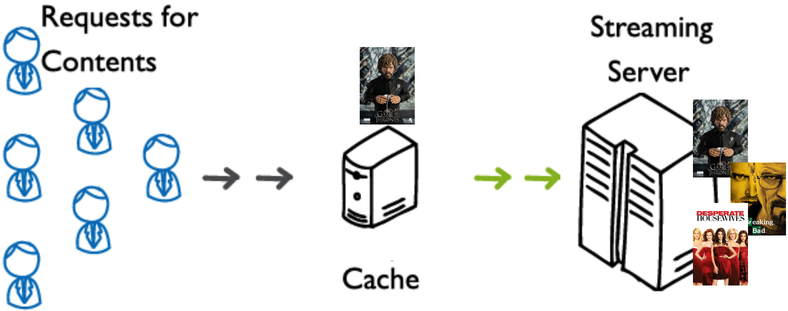

- Popularity-oblivious policies (LRU, RANDOM)

The RAND($\vec{m}$) policy

- RANDOM: exchange the requested item with a random item

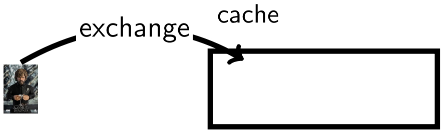

The RAND($\vec{m}$) policy

- RAND($\vec{m}$): exchange the requested item with an item from next list

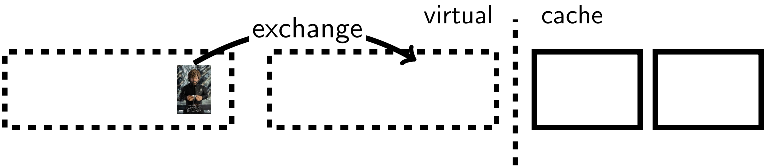

Design question: how many lists?

Conjecture from the 80s: if popularities do not vary, adding more lists is always better.

Our contributions

- RAND(m) model [Sigmetrics 2015]
  - Disprove a 30 year old conjecture (more lists is not always better).
  - Prove accuracy of a heterogeneous mean field model.
- Extension LRU(m) and q-LRU variants [ITC 2016, PEVA 2017]
  - Asymptotically exact analysis.
- Refined approximation for RAND(m) variants [Submitted, 2019]
  - We can define a refined approximation for an heterogeneous model.

Our contributions

- RAND(m) model [Sigmetrics 2015]
  - Disprove a 30 year old conjecture (more lists is not always better).
  - Prove accuracy of a heterogeneous mean field model.
- Extension LRU(m) and q-LRU variants [ITC 2016, PEVA 2017]
  - Asymptotically exact analysis.
- Refined approximation for RAND(m) variants [Submitted, 2019]
  - We can define a refined approximation for an heterogeneous model.

Our contributions

- RAND(m) model [Sigmetrics 2015]
  - Disprove a 30 year old conjecture (more lists is not always better).
  - Prove accuracy of a heterogeneous mean field model.
- Extension LRU(m) and q-LRU variants [ITC 2016, PEVA 2017]
  - Asymptotically exact analysis.
- Refined approximation for RAND(m) variants [Submitted, 2019]
  - We can define a refined approximation for an heterogeneous model.

Our contributions

- RAND(m) model [Sigmetrics 2015]
  - Disprove a 30 year old conjecture (more lists is not always better).
  - Prove accuracy of a heterogeneous mean field model.
- Extension LRU(m) and q-LRU variants [ITC 2016, PEVA 2017]
  - Asymptotically exact analysis.
- Refined approximation for RAND(m) variants [Submitted, 2019]
  - We can define a refined approximation for an heterogeneous model.

Heterogeneous mean field approximation

Let $H_i(t)$ be the popularity in list i.

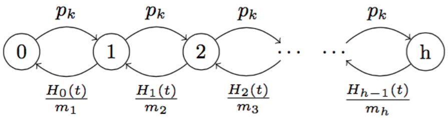

Heterogeneous mean field approximation

Let $H_i(t)$ be the popularity in list i.

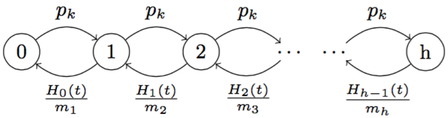

If $x_{k,i}(t)$ is the probability that item k is in list i at time t, we approximately have:

$$\dot{x}_{k,i}(t) = p_k x_{k,i-1}(t) - \overset{\text{Popularity in cache i-1}}{\boxed{\sum_j p_j x_{j,i-1}(t)}} \frac{x_{k,i}(t)}{m_i} + \mathbf{1}_{\{i<h\}} \left(\underset{\text{Popularity in cache i}}{\boxed{\sum_j p_j x_{j,i}(t)}} \frac{x_{k,i+1}(t)}{m_{i+1}} - p_k x_{k,i}(t) \right)$$

This approximation is of the form $\dot{x} = xQ(x)$.

The mean field ODE is asymptotically exact

"Theorem": For any T, there exists C such that for any popularities p and any list sizes $m_1 \dots m_h$

$$\mathbb{E}\left[\sup_{i,t<T/(\max_k p_k + \max_i 1/m_i)} \|H_i(t) - \delta_i(t)\|\right] \leq C\sqrt{\max_k p_k + \max_i 1/m_i}.$$

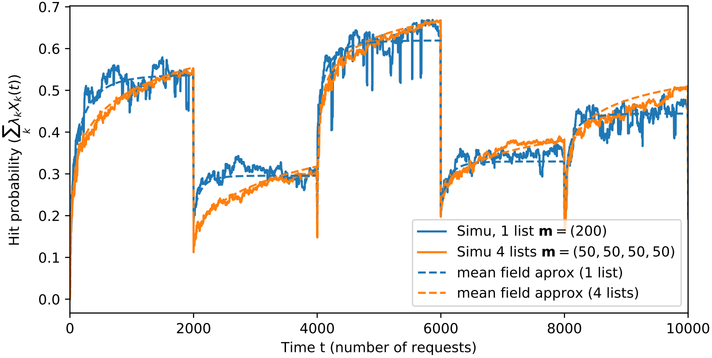

1000 items (Zipf). Popularities change every 2000 requests

We can refine the approximation

- We define a scaled process with N copies of each item.
- We use $Perf(N) \approx Perf(\infty) + \frac{V}{N}$ with $N = 1$.

We can refine the approximation

- We define a scaled process with N copies of each item.
- We use $Perf(N) \approx Perf(\infty) + \frac{V}{N}$ with $N = 1$.

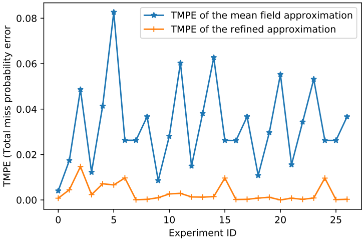

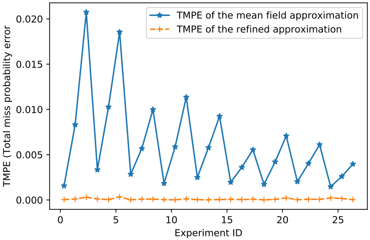

$K = 10$ items		$K = 100$ items
3%	MF error	0.6%
0.3%	RMF error	0.002%

Conclusion: Mean field is a powerful tool to study stochastic systems

- Works not only for large, homogeneous systems.

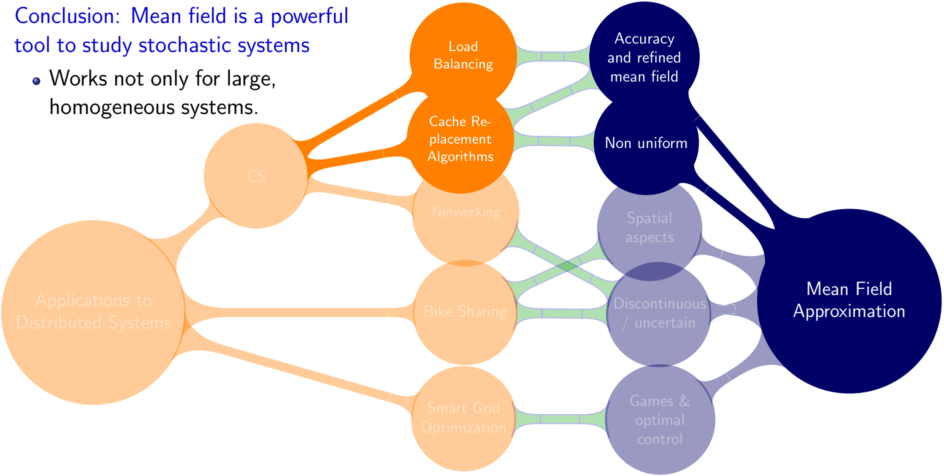

Conclusion: Mean field is a powerful tool to study stochastic systems

- Works not only for large, homogeneous systems.

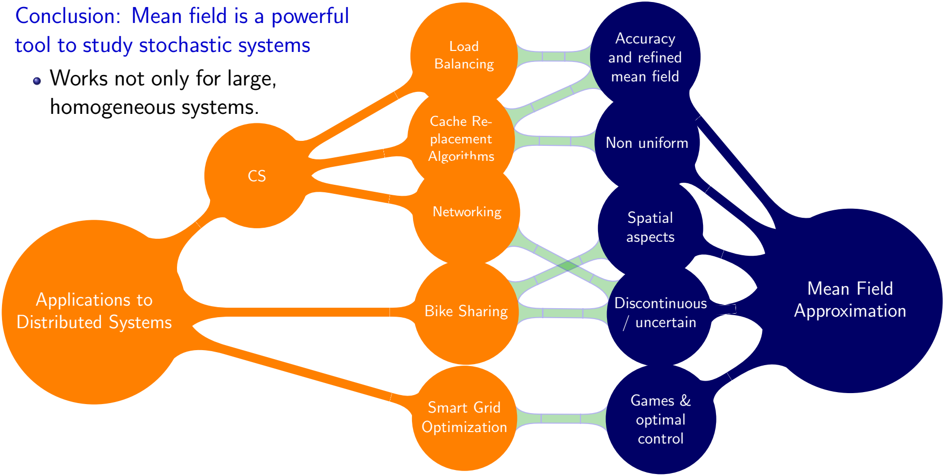

Future research directions:

- Online optimization algorithms

Conclusion: Mean field is a powerful tool to study stochastic systems

- Works not only for large, homogeneous systems.

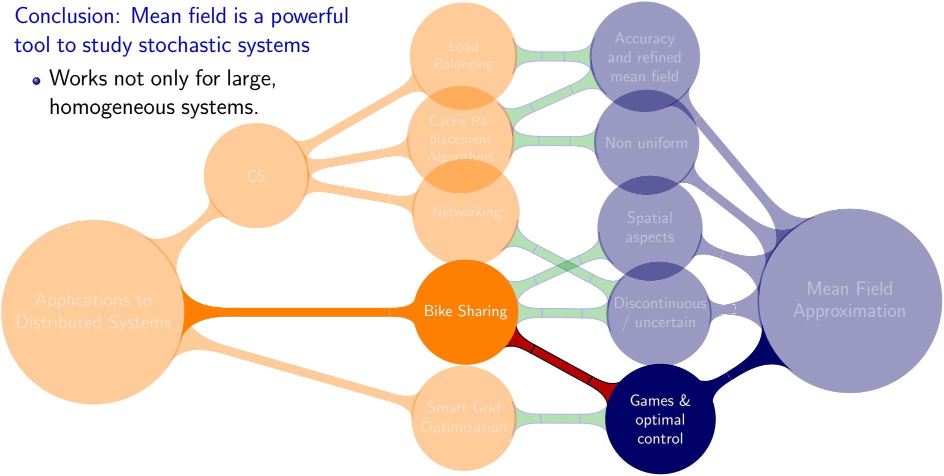

Future research directions:

- Online optimization algorithms

Conclusion: Mean field is a powerful tool to study stochastic systems

- Works not only for large, homogeneous systems.

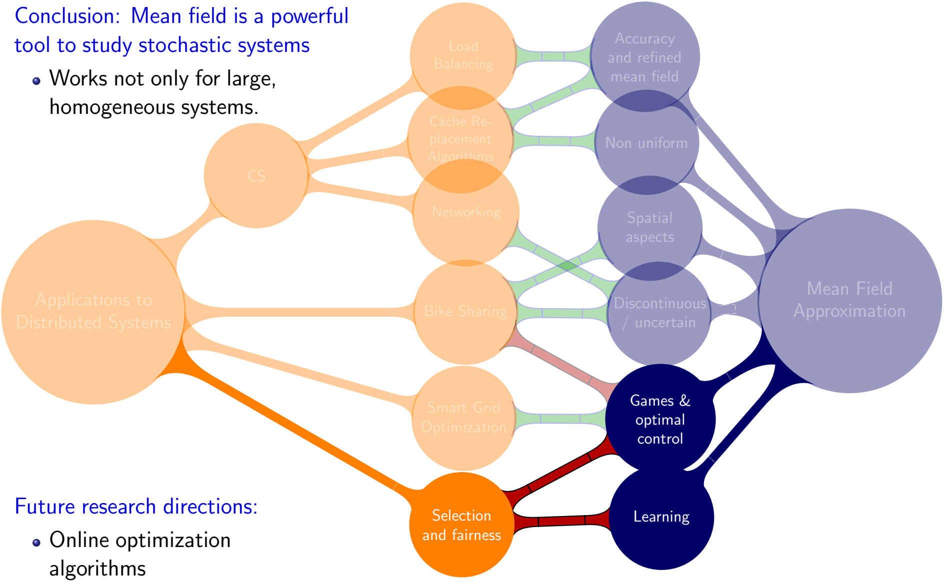

Future research directions:

- Online optimization algorithms

Conclusion: Mean field is a powerful tool to study stochastic systems

- Works not only for large, homogeneous systems.

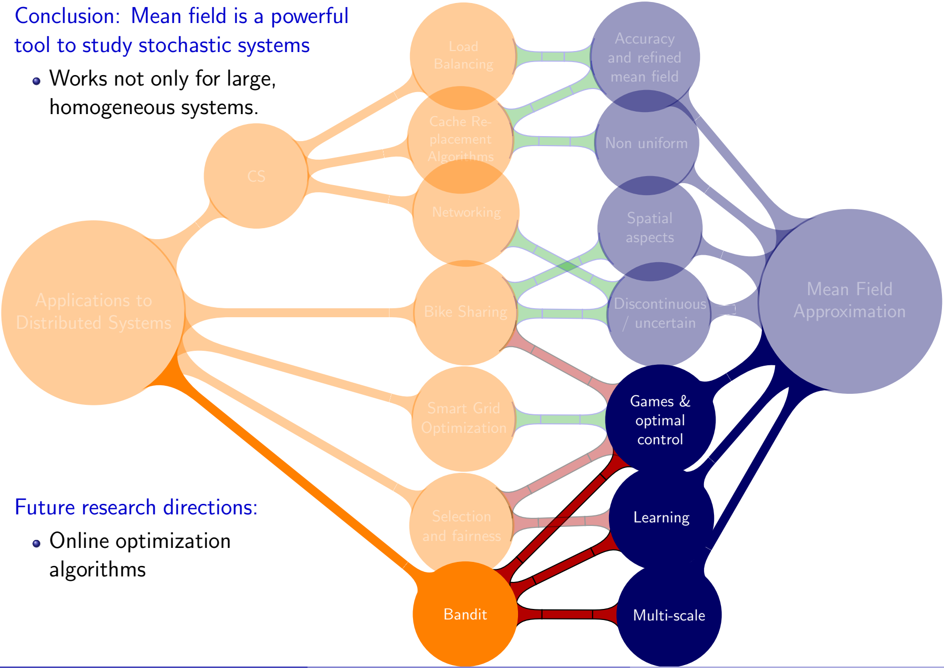

Future research directions:

- Online optimization algorithms

Conclusion: Mean field is a powerful tool to study stochastic systems

- Works not only for large, homogeneous systems.

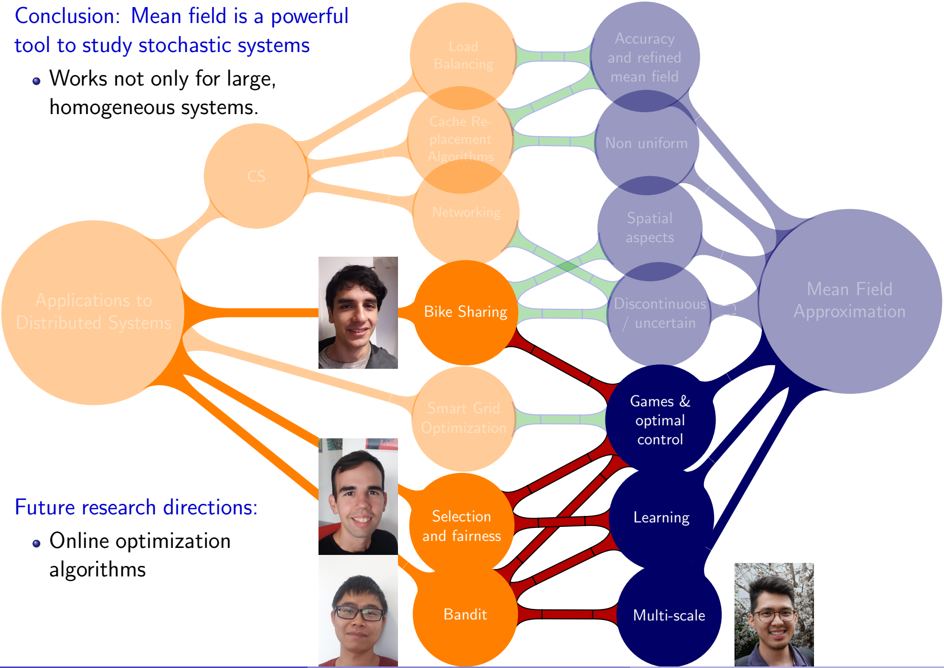

Future research directions:

- Online optimization algorithms

Thank you

Advisors

Former PhD students / post-docs

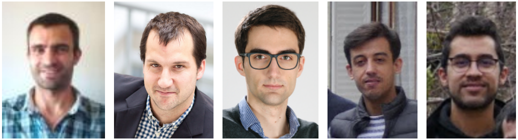

Phd students

Phd students

The refined approximation can also account for behaviors that are indistinguishable by classical mean field methods

Example: choosing with or without replacement

Let x_i be the fraction of queues of size i or more. Pick two queues, what is the probability that smallest has size i?

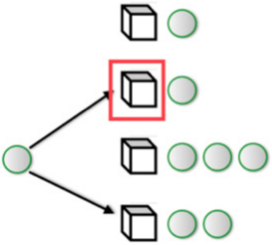

The refined approximation can also account for behaviors that are indistinguishable by classical mean field methods

Example: choosing with or without replacement

Let x_i be the fraction of queues of size i or more. Pick two queues, what is the probability that smallest has size i?

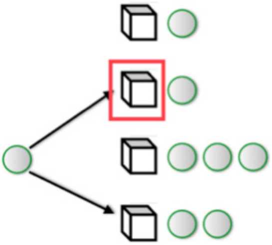

- If picked with replacement: $x_i^2 - x_{i+1}^2$.
- If picked without replacement: $x_i \dfrac{Nx_i - 1}{N-1} - x_{i+1} \dfrac{Nx_{i+1} - 1}{N-1}$

The refined approximation can also account for behaviors that are indistinguishable by classical mean field methods

Example: choosing with or without replacement

Let x_i be the fraction of queues of size i or more. Pick two queues, what is the probability that smallest has size i?

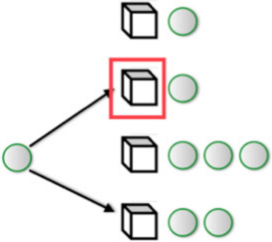

- If picked with replacement: $x_i^2 - x_{i+1}^2$.
- If picked without replacement: $x_i \dfrac{Nx_i - 1}{N-1} - x_{i+1}\dfrac{Nx_{i+1} - 1}{N-1}$

For $\rho = 0.9$:

$$\text{Average queue length (with rep)} \approx 2.353 + \frac{4}{N}$$

$$\text{Average queue length (without rep)} \approx 2.353 + \frac{4}{N} - \frac{1}{N}$$

How does the expected queue length evolve with time?

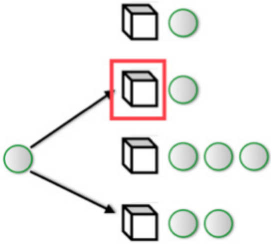

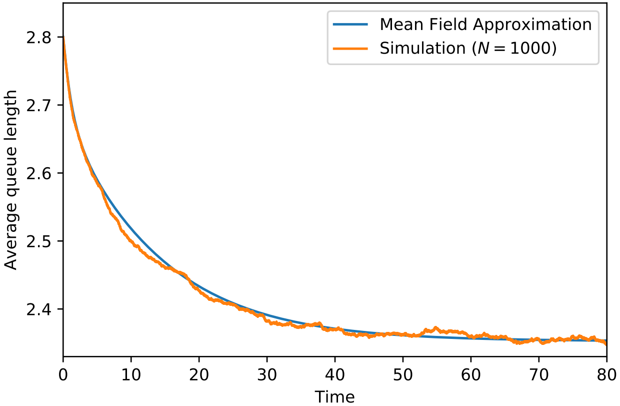

How does the expected queue length evolve with time?

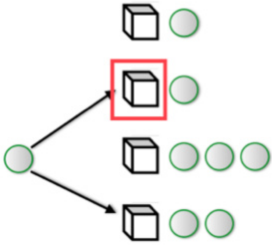

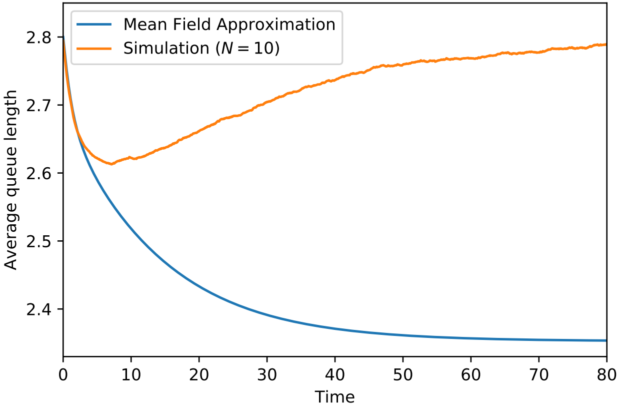

How does the expected queue length evolve with time?

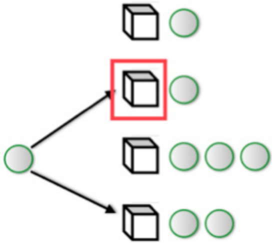

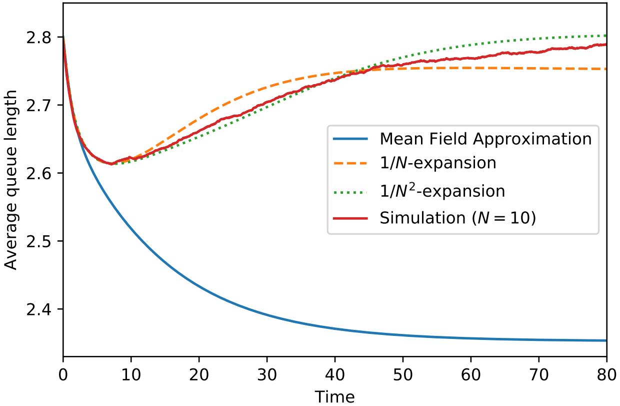

Remark about computation time :

- 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`

Does it always work?

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

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

Does it always work?

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

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

NO: