Mean Field Methods for Computer and Communication Systems
Part 1: Finite Horizon

Jean-Yves Le Boudec
EPFL

ROCKS Automn School
October 2012

Contents

1. Mean Field Interaction Model
2. Convergence to Mean Field Limit *made easy*
3. Formulating the Mean Field Limit
4. Fast Simulation, Decoupling assumption
5. Convergence to Mean Field Limit
   – general case

1

MEAN FIELD INTERACTION MODEL

Mean Field

- A *model* introduced in Physics
  - interaction between *particles* is via distribution of states of all particle
- An *approximation* method for a large collection of particles
  - assumes *independence* in the master equation
- Why do we care in information and communication systems ?
  - Model interaction of many objects:
  - Distributed systems, communication protocols, game theory, self-organized systems

A Few Examples Where Applied

[1] L. Afanassieva, S. Popov, and G. Fayolle. Models for transporation networks. *Journal of Mathematical Sciences, 1997 – Springer*.

[2] F. Baccelli, A. Chaintreau, D. De Vleeschauwer, and D. R. McDonald. Http turbulence, May 2004.

[3] F. Baccelli, M. Lelarge, and D. McDonald. Metastable regimes for multiplexed tcp flows. In *Proceedings of the*

[5] M.-D. Bordenave, Charles and A. Proutiere. A particle system in interaction with a rapidly varying environment: Mean field limits and applications. arXiv:math/0701363v2.

[11] S. Kumar and L. Massoulié. Integrating streaming and file-transfer internet traffic: Fluid and diffusion approximations. MSR-TR-2005-160.

[16] Y. M. Suhov and N. D. Vvedenskaya. Dobrushin's mean-field approximation for a queue with dynamic routing. *Markov Processes and Related Fields*, 3(4):493–526, 1997.

[17] P. Tinnakornsrisuphap and A. M. Makowski. Limit behavior of ecn/red gateways under a large number of tcp flows. In *Proceedings IEEE INFOCOM 2003, The 22nd Annual Joint Conference of the IEEE Computer and Communications Societies, San Franciso, CA, USA*, March 30 - April 3 2003.

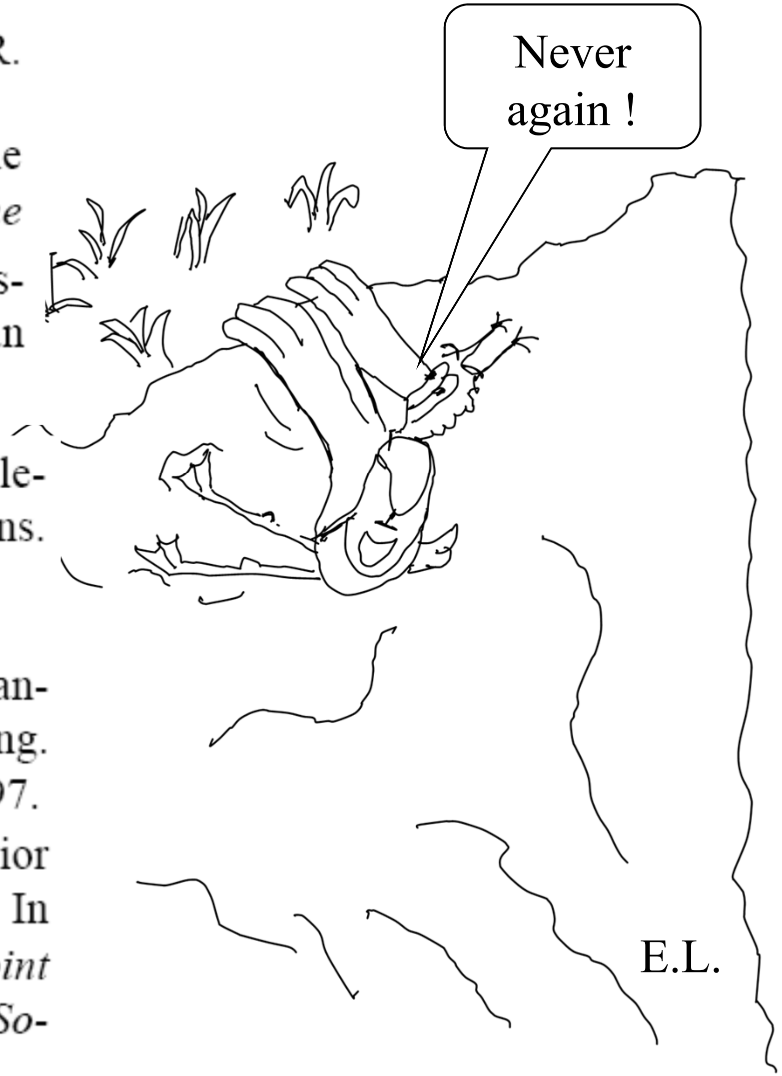

Mean Field Interaction Model

- Time is discrete or continuous
- N objects, N large
- Object n has state $X_n(t)$
- $(X^N_1(t), \ldots, X^N_N(t))$ is Markov
- Objects are observable only through their state

- "Occupancy measure" $M^N(t)$ = distribution of object states at time t
- ***Theorem*** [Gast (2011)] $M^N(t)$ is Markov

Example: 2-Step Malware

- Mobile nodes are either
  - `S' Susceptible
  - `D' Dormant
  - `A' Active
- Time is discrete
- Transitions affect 1 or 2 nodes

- State space is finite
  = {`D' , `A' ,`S'}

- Occupancy measure is
  M(t) = (D(t), A(t), S(t)) with
  S(t)+ D(t) + A(t) =1

D(t) = proportion of nodes in state `D'

[Benaïm and Le Boudec(2008)]

1. Recovery
   - D -> S
2. Mutual upgrade
   - D + D -> A + A
3. Infection by active
   - D + A -> A + A
4. Recovery
   - A -> S
5. Recruitment by Dormant
   - S + D -> D + D

   Direct infection
   - S -> D
6. Direct infection
   - S -> A

2-Step Malware – Full Specification

At every time step, pick one node unif at random

- If node is in state D:
  - With proba δ_D mutate to S
  - With proba $\lambda \frac{ND-1}{N}$, meet another D node and both mutate to A
- If node is in state A:
  - With proba $\beta \frac{D}{h+D}$ change one D node to A
  - With Proba δ_A mutate to S
- If node is in state S
  - With proba rD meet a D node and become infected D
  - With proba α_0 become infected D
  - With proba α become infected A

1. Recovery
   - D -> S
2. Mutual upgrade
   - D + D -> A + A
3. Infection by active
   - D + A -> A + A
4. Recovery
   - A -> S
5. Recruitment by Dormant
   - S + D -> D + D

   Direct infection
   - S -> D
6. Direct infection
   - S -> A

case	prob
1	$D\delta_D$
2	$D\lambda\frac{ND-1}{N-1}$
3	$A\beta\frac{D}{h+D}$
4	$A\delta_A$
5	$S(\alpha_0 + rD)$
6	$S\alpha$

Simulation Runs, N=1000 nodes

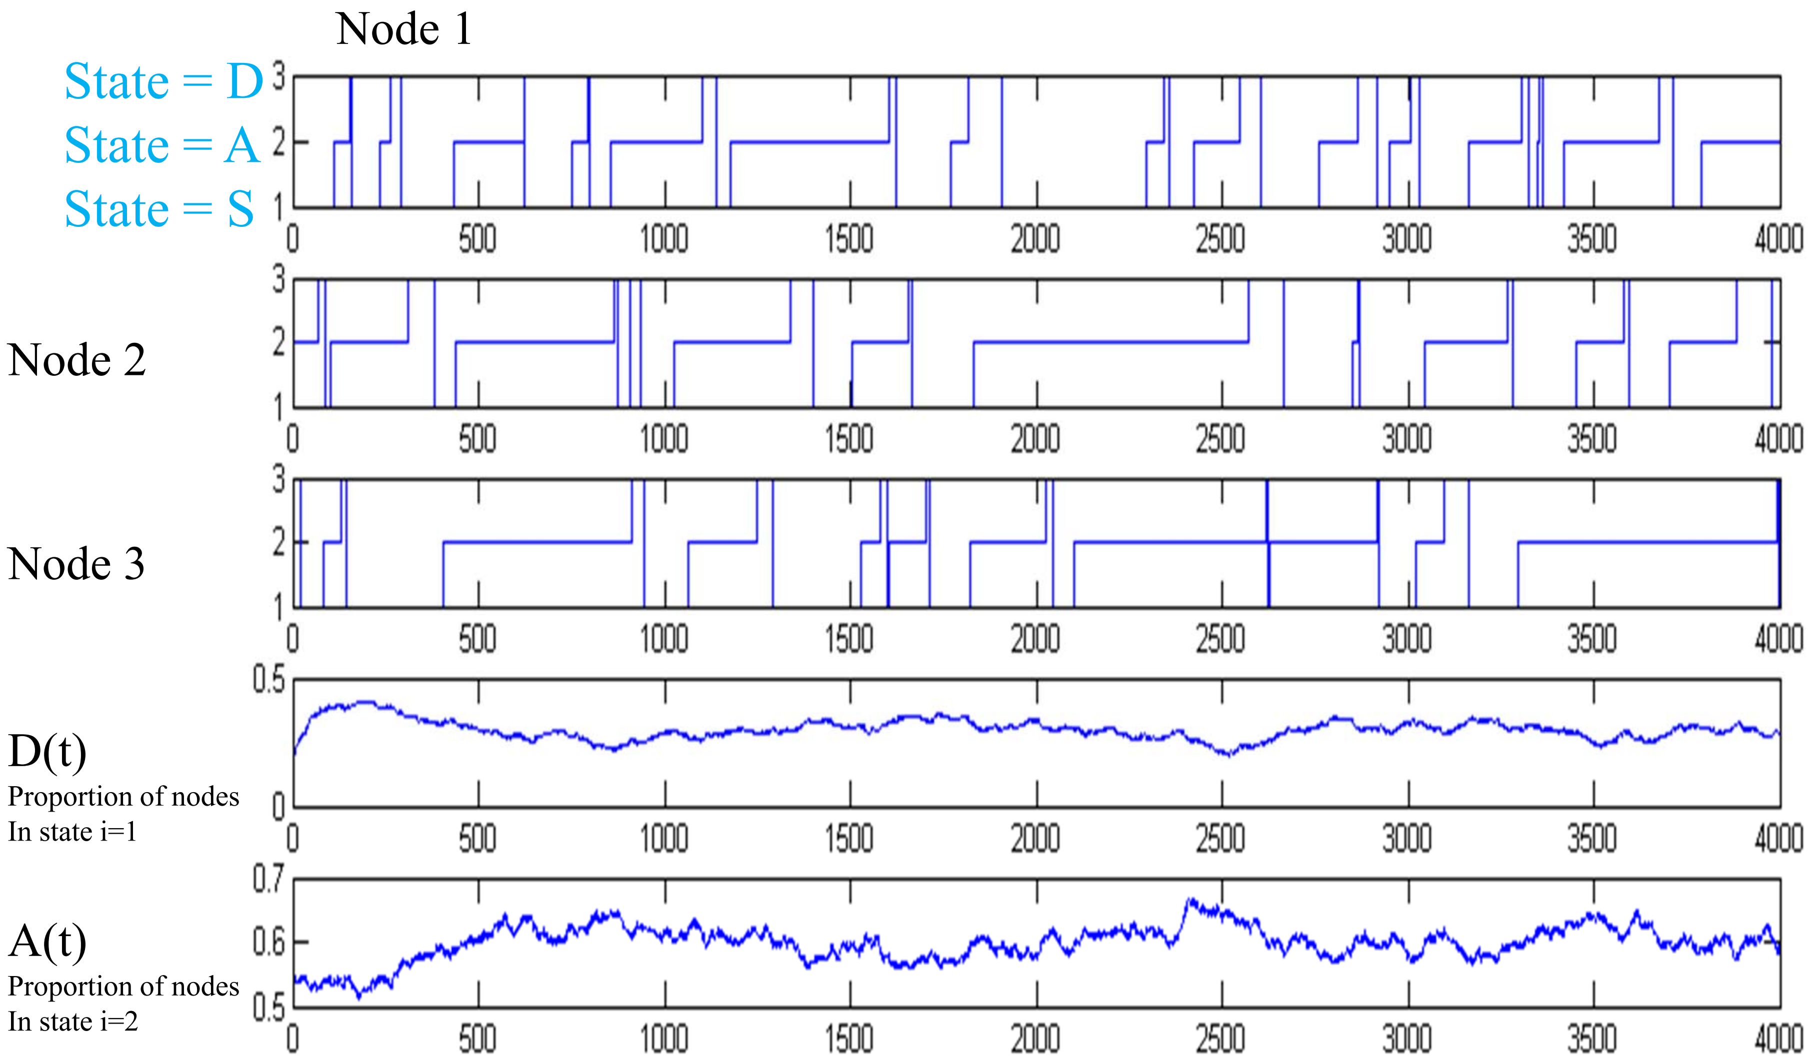

$$\beta = 0.01, \delta_A = 0.005, \delta_D = 0.0001, \alpha_0 = \alpha = 0.0001, h = 0.3, r = 0.1, \lambda = 0.0001$$

Sample Runs with N = 1000

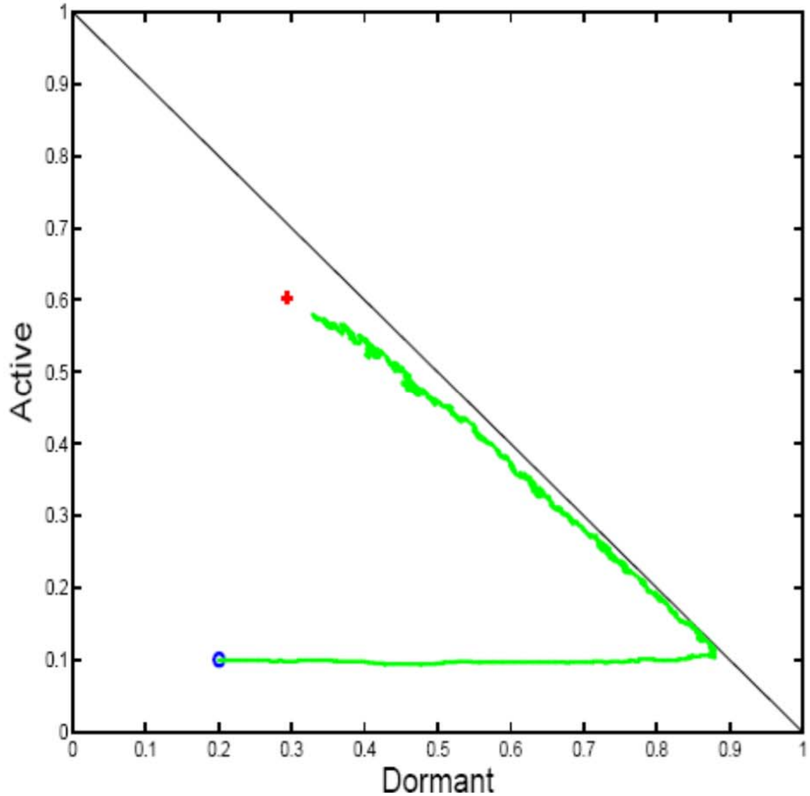

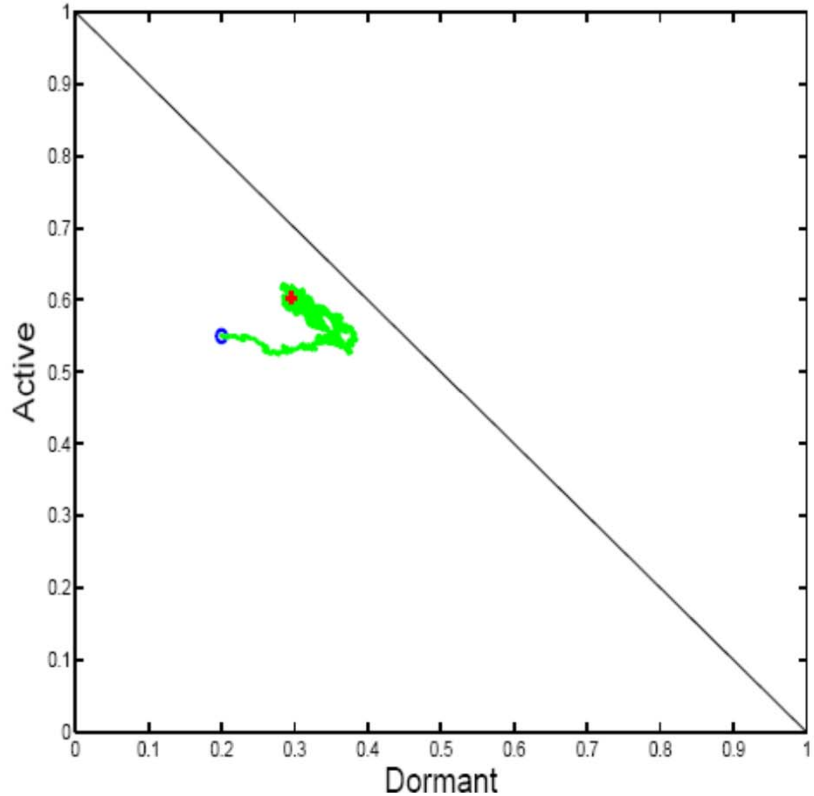

$$\beta = 0.01, \delta_A = 0.005, \delta_D = 0.0001, \alpha_0 = \alpha = 0.0001, h = 0.3, r = 0.1, \lambda = 0.0001$$

Example: WiFi Collision Resolution Protocol

- N nodes, state = retransmission stage k
- Time is discrete, $I(N) = 1/N$; mean field limit is an ODE
- Occupancy measure is $M(t) = [M_0(t),...,M_K(t)]$ with $M_k(t)$ = proportion of nodes at stage k
- [Bordenave et al.(2008)Bordenave, McDonald, and Proutiere, Bordenave et al.(2007)Bordenave, McDonald, and Proutiere]

Example: Dissemination in a Vehicle Fleet

- Nikodin Ristanovic's PhD thesis

- Without Taxi to Taxi Dissemination →

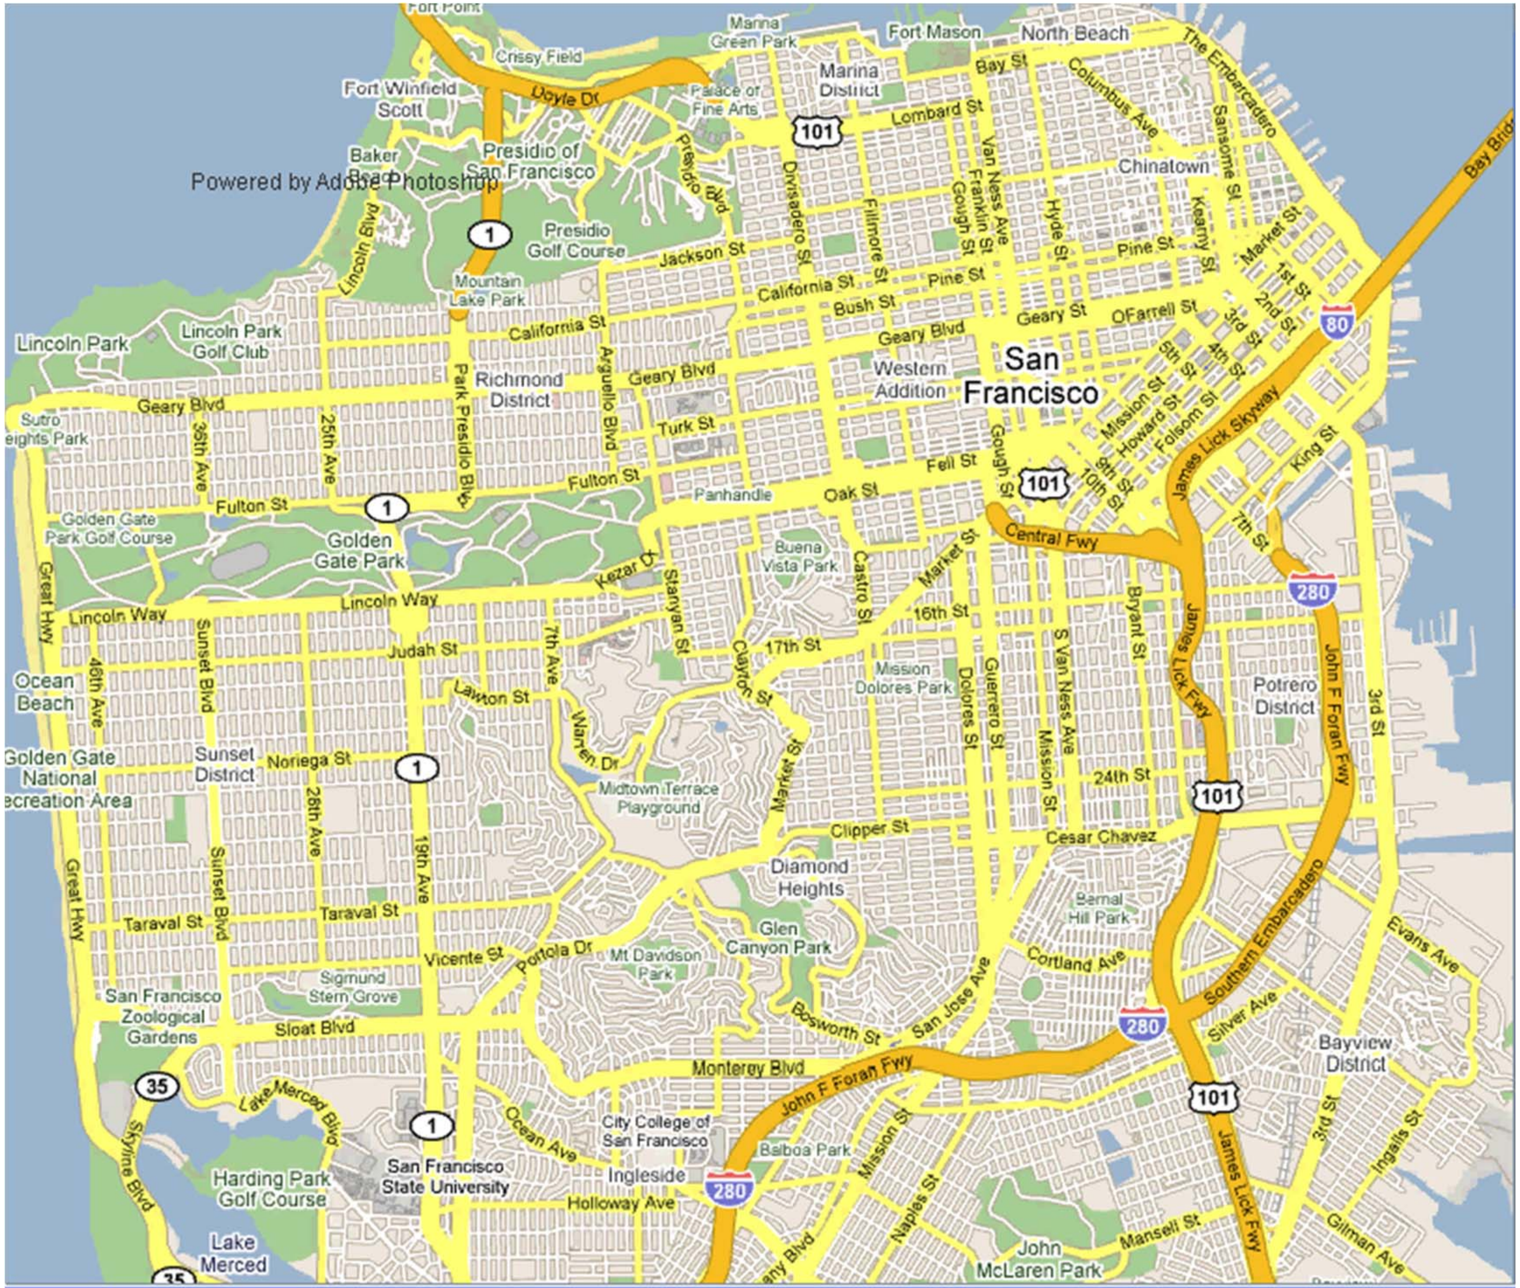

With Taxi to Taxi Dissemination

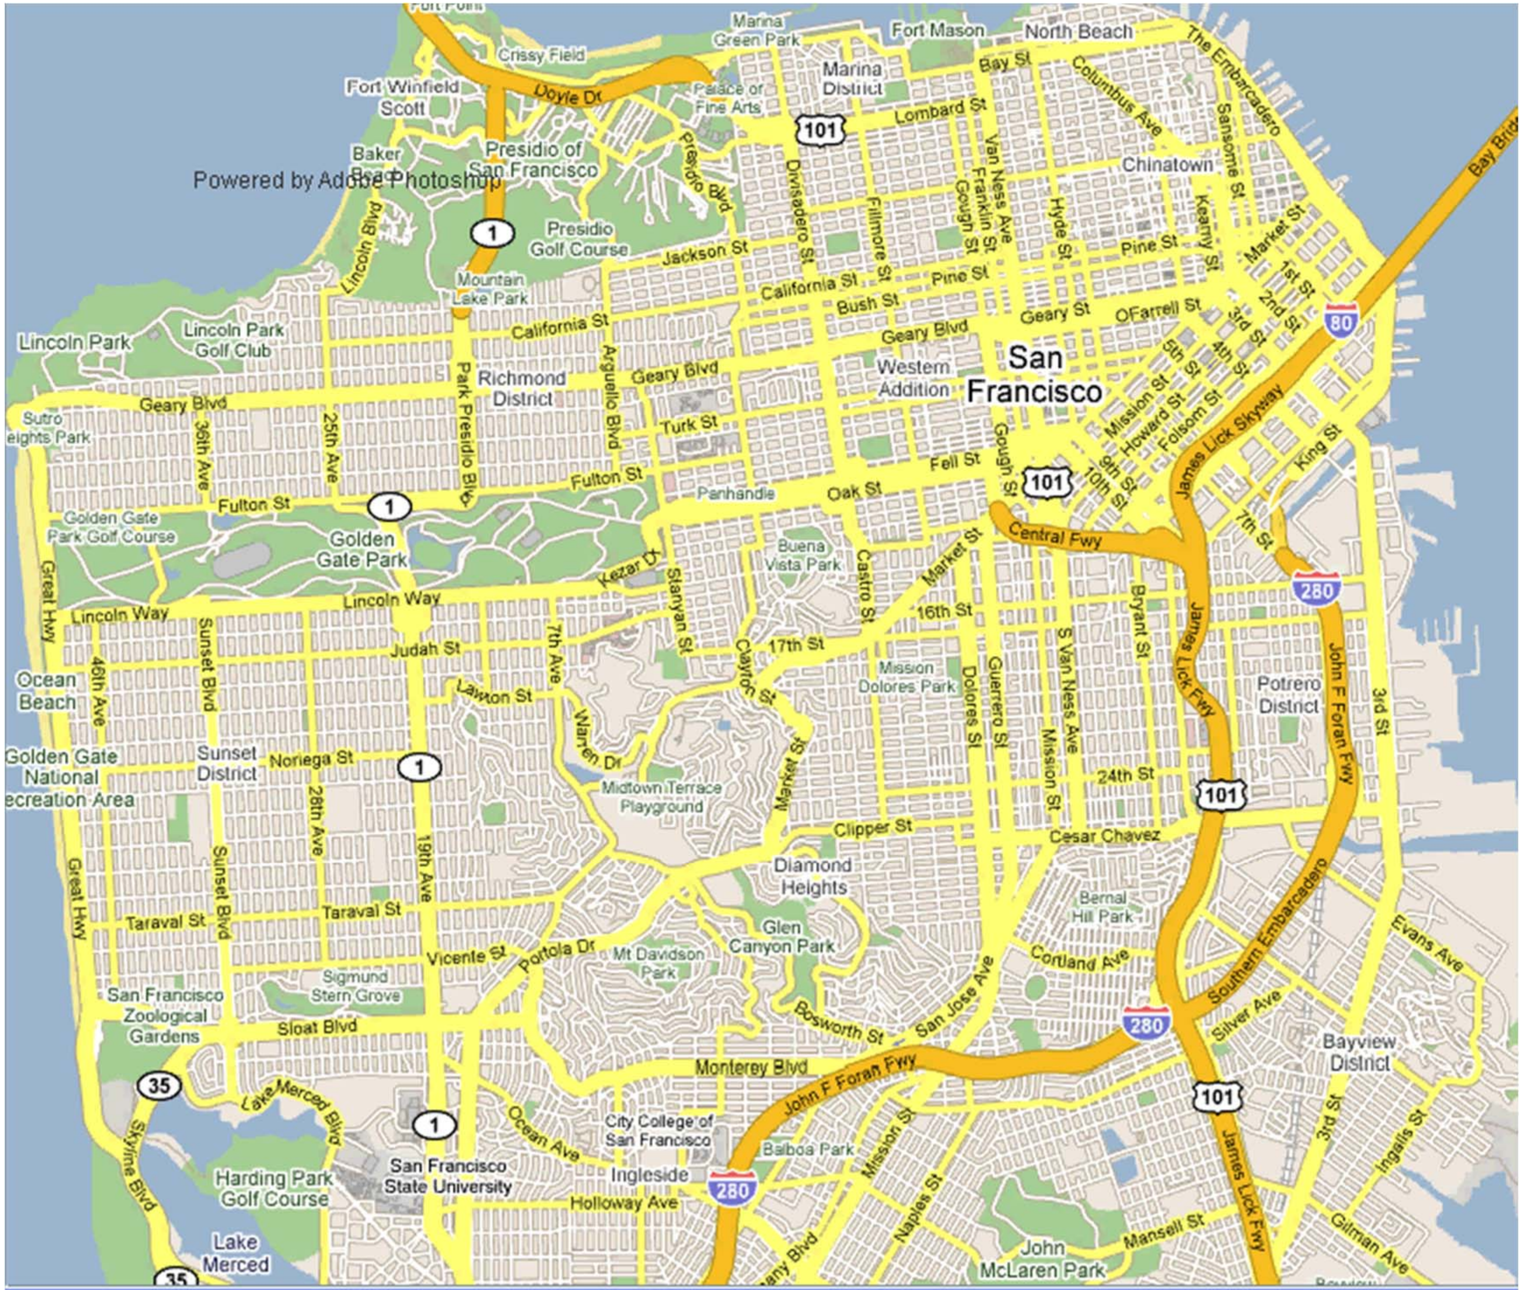

The Importance of Being Spatial

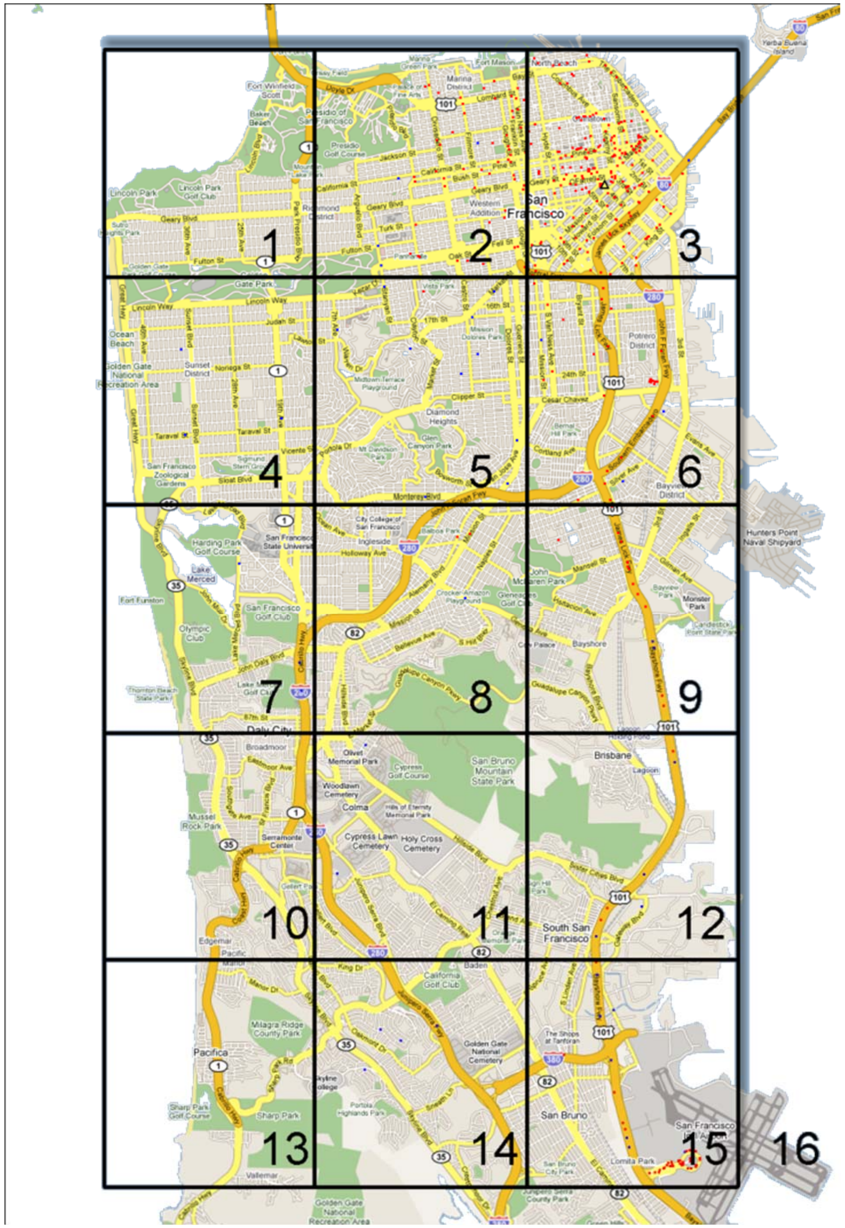

- Mobile node state = (c, t)
  c = 1 … 16 (position)
  t $\in R^+$ (age of gossip)

- Time is continuous
- Occupancy measure is
  $F_c(z,t)$ = proportion of nodes that at location *c* and have age ≤ *z*

  [Age of Gossip, Chaintreau et al.(2009)]

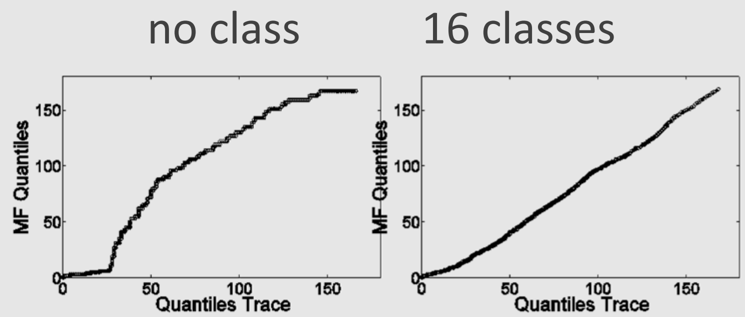

Qqplots simulation vs mean field

What can we do with a Mean Field Interaction Model ?

- Large N asymptotics, Finite Horizon
  - fluid limit of occupancy measure (ODE)
  - decoupling assumption
    (fast simulation)
- Issues
  - When valid
  - How to formulate the fluid limit

- Large t asymptotic
  - Stationary approximation of occupancy measure
  - Decoupling assumption
- Issues
  - When valid

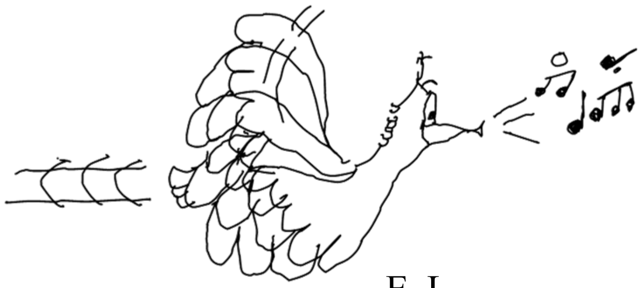

E. L.

2.

CONVERGENCE TO MEAN FIELD *MADE EASY*

To Obtain a Mean Field Limit we Must Make Assumptions about the Intensity *I(N)*

- $I(N)$ = (order of) expected number of transitions per object per time unit

- A mean field limit occurs when we re-scale time by $I(N)$ i.e. one time slot $\approx I(N)$
  i.e. we consider $X^N(t/I(N))$

- I(N) = O(1/N): mean field limit is in continuous time [Benaïm and Le Boudec (2008)]

  I(N) = O(1): mean field limit is in discrete time
  [Le Boudec et al (2007)]

Intensity for 2-step malware model is $1/N$

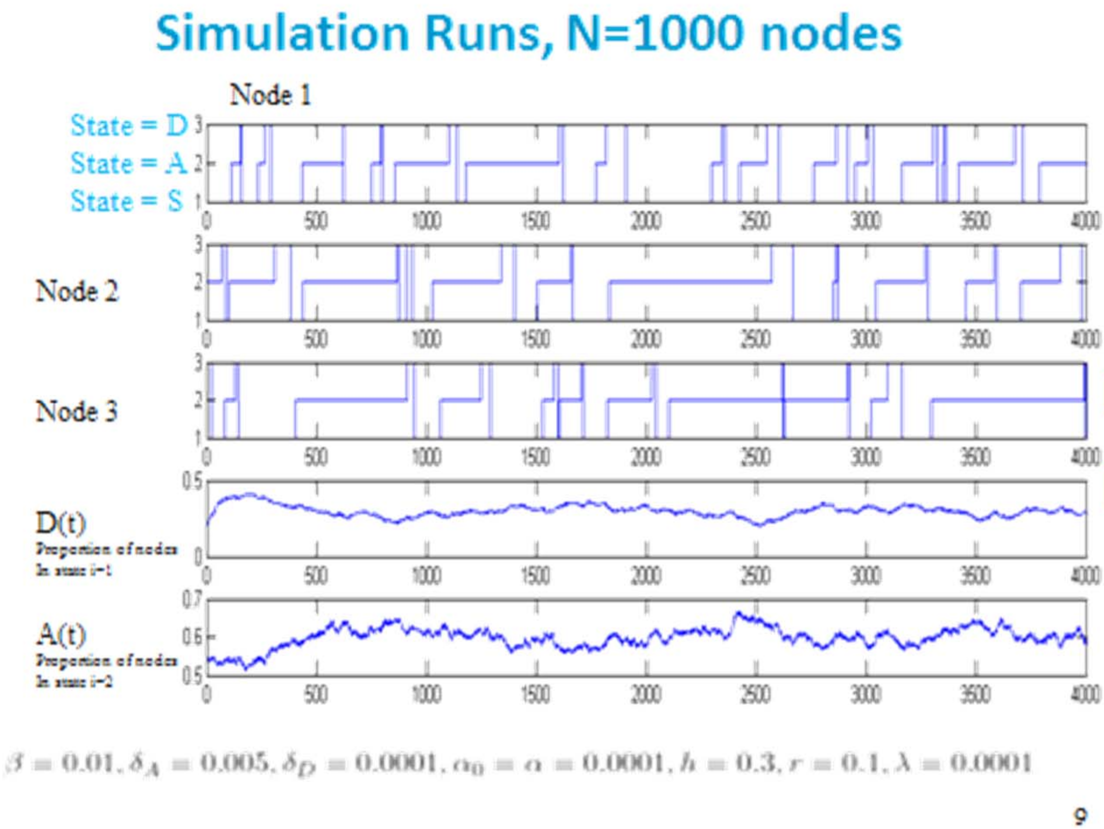

- In one time step, the number of objects affected by a transition is 0, 1 or 2; mean number of affected objects is $O(1)$
- There are N objects
- Expected number of transitions per time slot per object is $O\left(\frac{1}{N}\right)$

The Mean Field Limit

- Under very general conditions (given later) the occupancy measure converges, in law, to a deterministic process, *m(t),* called the *mean field limit*

$$M^N\left(\frac{t}{I(N)}\right) \to m(t)$$

- Finite State Space
  $+$ Vanishing intensity $\left[\, I(N) = O\left(\frac{1}{N}\right)\right]$
  $\Rightarrow$ mean field limit is ODE

Mean Field Limit
N = +∞

$m(t) \approx M^N(t\,N)$

Stochastic system
N = 1000

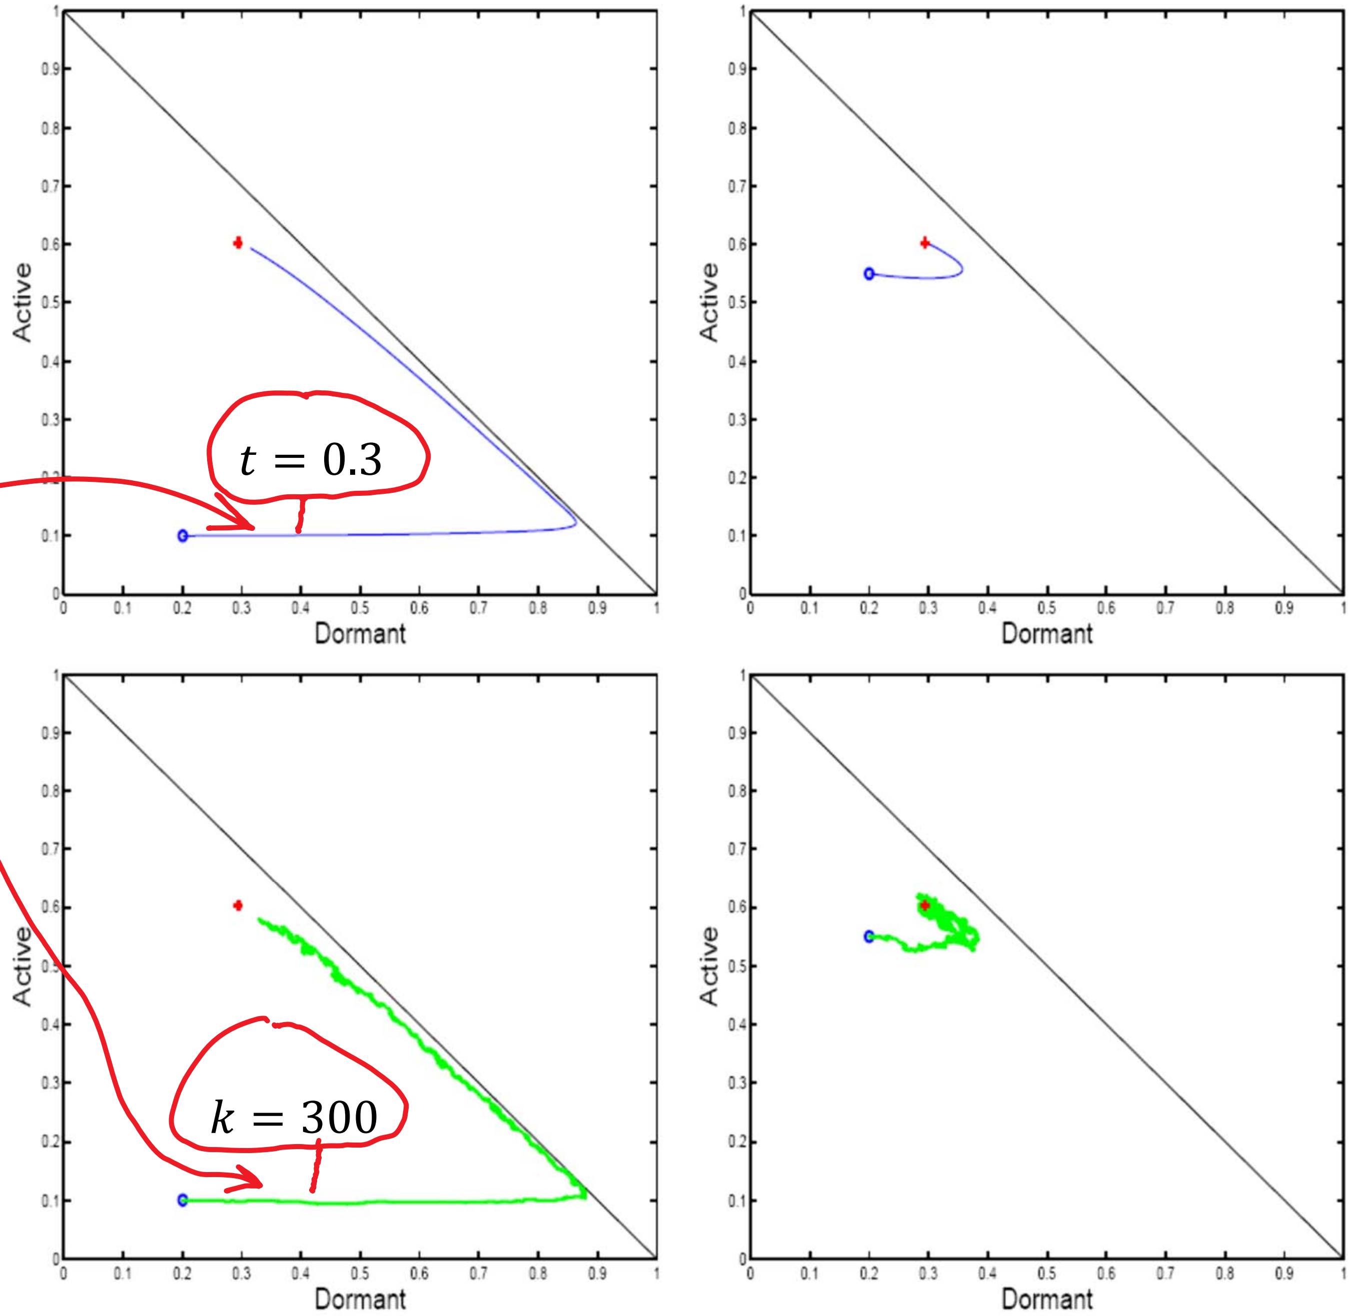

Sufficient Conditions for Convergence verifiable by inspection

- Condition 1: state space (for one object) is finite
- and Condition 2: $I(N) \to 0$
- and Condition 3: probabilities at every time slot depend smoothly (C^1) on all parameters and have a limit when $N \to \infty$
- and Condition 4 : Second moment of number of objects affected in one timeslot $\leq$ a constant

Example: Convergence to Mean Field; the 4 Conditions Apply

Mean Field Limit
N = +∞

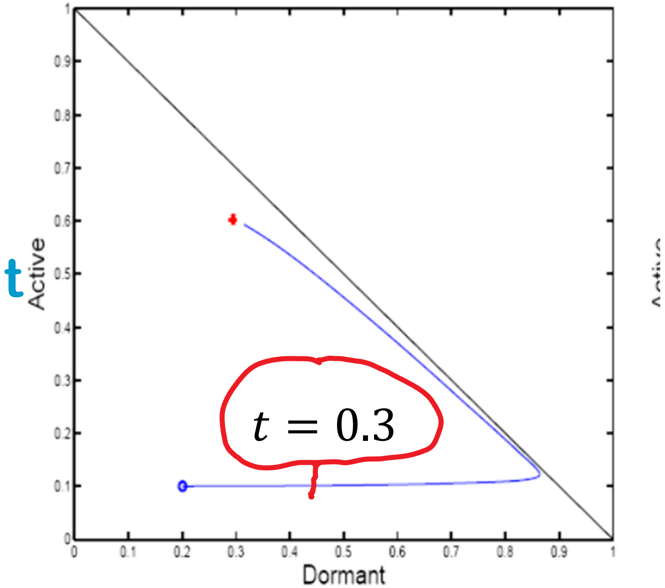

Stochastic system
N = 1000

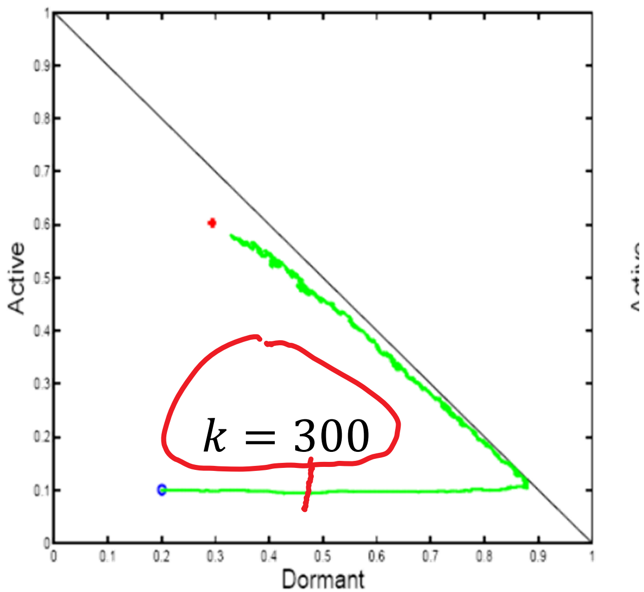

1. 3 states
2. $I(N) = \frac{1}{N}$
3. See table
4. Number of transitions per time step is bounded by 2

The convergence theorem

[Benaïm and Le Boudec(2008), Ioannidis and Marbach(2009)]

- Let $W^N(k)$ be the number of objects that do a transition in time slot k. Note that $\mathbb{E}\left(W^N(k)\right) = NI(N)$, where $I(N) \stackrel{\text{def.}}{=}$intensity. Assume

$$\mathbb{E}\left(W^N(k)^2\right) \leq \beta(N) \quad \text{with} \quad \lim_{N\to\infty} I(N)\beta(N) = 0$$

- $M^N(0) \to m_0$ in probability
- regularity assumption on the drift (generator)

Then $\sup_{0\leq t\leq T} \mathbb{P}\left(\|M^N(t) - m(t)\|\right) \to 0$ in probability.

3

FORMULATING THE MEAN FIELD LIMIT

A key concept to write the mean field limit is the *drift*

- Assume you have a model with a mean field limit as in the previous section
- The mean field limit is an ODE
- How can we write the ODE without error ?
- Solution: study first the *drift* of the original model

Drift of a Markov Process

- Given some discrete time Markov process $Z(k)$ on some state space $E \subset R^d$
- the drift f of the process is the mapping $E \to E$ defined by:

$$f(z) \coloneqq E\big(Z(k+1) - Z(k) \big| Z(k) = z\big)$$

- Example: 2-step malware with N objects:

$$Z(k) = M^N(k) = \big(D(k), A(k), S(k)\big)$$

$$f^N\begin{pmatrix} d \\ a \\ s \end{pmatrix} = E\left(\begin{pmatrix} D(k+1) - D(k) \\ A(k+1) - A(k) \\ S(k+1) - S(k) \end{pmatrix} \middle| \begin{pmatrix} D(k) \\ A(k) \\ S(k) \end{pmatrix} = \begin{pmatrix} d \\ a \\ s \end{pmatrix}\right)$$

$$=: \begin{pmatrix} f_1^N(d, a, s) \\ f_2^N(d, a, s) \\ f_3^N(d, a, s) \end{pmatrix}$$

Let's compute $f_3^N(s,a,d)$ $:= E(S(k+1)-s|(S(k)=s, A(k)=a, D(k)=d)$

1. Recovery
   - D -> S
2. Mutual upgrade
   - D + D -> A + A
3. Infection by active
   - D + A -> A + A
4. Recovery
   - A -> S
5. Recruitment by Dormant
   - S + D -> D + D

   Direct infection
   - S -> D
6. Direct infection
   - S -> A

case	prob
1	$D\delta_D$
2	$D\lambda\frac{ND-1}{N-1}$
3	$A\beta\frac{D}{h+D}$
4	$A\delta_A$
5	$S(\alpha_0 + rD)$
6	$S\alpha$

$$f_3^N(s,a,d) =$$

Let's compute $f_3^N(d,a,s)$ $:= E(S(k+1)-s \mid (D(k)=d, A(k)=a, S(k)=s)$

1. Recovery
   - D -> S
2. Mutual upgrade
   - D + D -> A + A
3. Infection by active
   - D + A -> A + A
4. Recovery
   - A -> S
5. Recruitment by Dormant
   - S + D -> D + D

   Direct infection
   - S -> D
6. Direct infection
   - S -> A

case	prob
1	$D\delta_D$
2	$D\lambda\frac{ND-1}{N-1}$
3	$A\beta\frac{D}{h+D}$
4	$A\delta_A$
5	$S(\alpha_0+rD)$
6	$S\alpha$

$$f_3^N(d,a,s) = \frac{1}{N}\left(D\delta_D + A\,\delta_A - S(\alpha_0+rD) - S\alpha \right)$$

The drift for the 2-step malware example with *N* objects is

$$\text{drift} = f(D, A, S) = \frac{1}{N}\begin{pmatrix} -D\delta_D - 2D\lambda\frac{ND-1}{N-1} - A\beta\frac{D}{h+D} + S(\alpha_0 + rD) \\ 2D\lambda\frac{ND-1}{N-1} + A\beta\frac{D}{h+D} - A\delta_A + S\alpha \\ D\delta_D + A\delta_A - S(\alpha_0 + rD) - S\alpha \end{pmatrix}$$

The mean field limit is derived from the drift

- Given some discrete time Markov process $Z(k)$ on some state space $E \subset R^d$ with drift f :

$$\underbrace{Z(k+1) = Z(k) + f\big(Z(k)\big)}_{\text{Deterministic evolution}} + \underbrace{\xi(k)}_{\substack{\text{Stochastic evolution} \\ \text{Martingale noise}}}$$

- Application to mean field model $Z = M^N$:

$$M^N(k+1) = M^N(k) + f^N(M^N(k)) + \xi^N(k)$$
$$= M^N(k) + \underbrace{I(N)\left[\frac{f^N(M^N(k))}{I(N)}\right]}_{\substack{\text{has a limit } f \text{ under} \\ \text{conditions 1 to 4}}} + \underbrace{\xi^N(k)}_{\substack{\to 0 \text{ under} \\ \text{conditions 1 to 4}}}$$

Interpretation of the Mean Field limit as a stochastic approximation of an ODE

- Let $f(m) := \lim_{N\to\infty} \frac{f^N(m)}{I(N)}$ (re-scaled drift)
  This limit exists by Conditions 1 to 4.

- $M^N(k+1) \approx M^N(k) + I(N) f(M^N(k)) + noise$

  i.e. $M^N(k)$ is an approximation of the ODE
  $$\frac{dm}{dt} = f(m)$$

  with time step $\Delta t = I(N)$

The ODE for the 2-step malware example

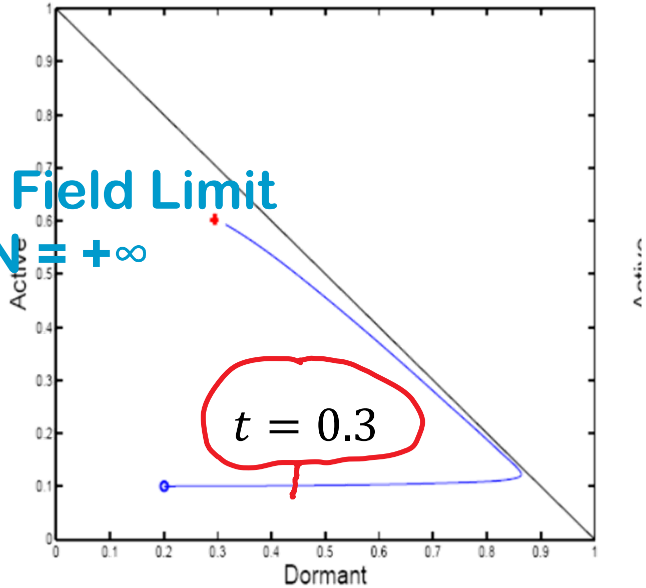

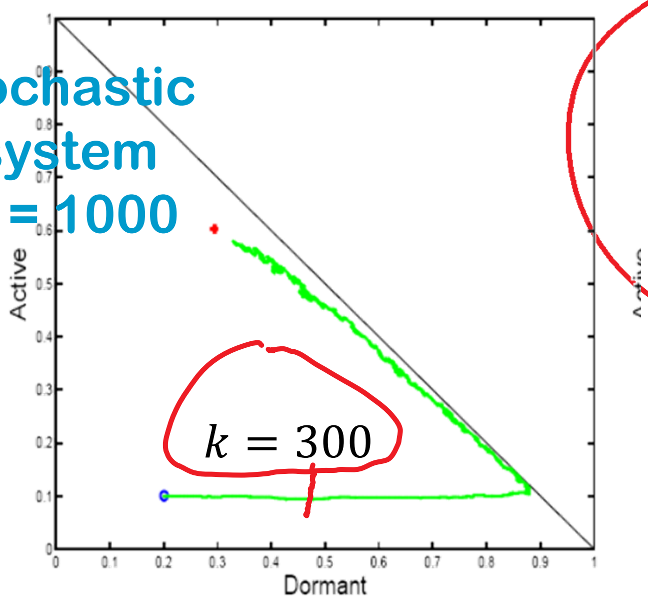

$$\text{drift} = f^N(D, A, S) =$$

$$\frac{1}{N}\begin{pmatrix} -D\delta_D - 2D\lambda\frac{ND-1}{N-1} - A\beta\frac{D}{h+D} + S(\alpha_0 + rD) \\ 2D\lambda\frac{ND-1}{N-1} + A\beta\frac{D}{h+D} - A\delta_A + S\alpha \\ D\delta_D + A\delta_A - S(\alpha_0 + rD) - S\alpha \end{pmatrix}$$

ODE

$$\begin{aligned} \frac{\partial D}{\partial t} &= -\delta_D D - 2\lambda D^2 - \beta A\frac{D}{h+D} + (\alpha_0 + rD)S \\ \frac{\partial A}{\partial t} &= 2\lambda D^2 + \beta A\frac{D}{h+D} - \delta_A A + \alpha S \\ \frac{\partial S}{\partial t} &= \delta_D D + \delta_A A - (\alpha_0 + rD)S - \alpha S \end{aligned}$$

Formulating the Mean Field Limit: Automation

- *Drift* = sum over all transitions of

  proba of transition

  x

  Delta to system state $M^N(t)$

- Re-scale drift by intensity

- Equation for mean field limit is

  $\frac{dm}{dt}$ = limit of

  rescaled drift

- Can be automated using reaction language

  http://icawww1.epfl.ch/IS/tsed

case	prob	effect on (D, A, S)
1	$D\delta_D$	$\frac{1}{N}(-1, 0, 1)$
2	$D\lambda\frac{ND-1}{N-1}$	$\frac{1}{N}(-2, +2, 0)$
3	$A\beta\frac{D}{h+D}$	$\frac{1}{N}(-1, +1, 0)$
4	$A\delta_A$	$\frac{1}{N}(0, -1, +1)$
5	$S(\alpha_0 + rD)$	$\frac{1}{N}(+1, 0, -1)$
6	$S\alpha$	$\frac{1}{N}(0, +1, -1)$

4.

FAST SIMULATION AND DECOUPLING ASSUMPTION (PERF TUT)

The Decoupling Assumption

- Often used in analysis of complex systems
- Says that k objects are asymptotically mutually independent (k is fixed and $N \to \infty$)
- What is the relation to mean field convergence ?

The Decoupling Assumption

- Often used in analysis of complex systems
- Says that k objects are asymptotically mutually independent (k is fixed and $N \to \infty$)
- What is the relation to mean field convergence ?

- [Sznitman 1991] [For a mean field interaction model:]

Decoupling assumption
$\Leftrightarrow$
$M^N(t)$ converges to a deterministic limit

- Further, if decoupling assumption holds, $m(t) \approx$ state proba for any arbitrary object

The Two Interpretations of the Mean Field Limit

- At any time t

$$P(X_n(t) = 'A') \;\approx\; A\left(\frac{t}{N}\right)$$

$$P(X_m(t) = 'D', X_n(t) = 'A') \;\approx\; D\left(\frac{t}{N}\right) A\left(\frac{t}{N}\right)$$

where (D, A, S) is solution of ODE

- Thus for $N = 1000$ and simulation step $k = 300$:
  - Prob (node n is dormant) ≈ 0.48
  - Prob (node n is active) ≈ 0.19
  - Prob (node n is susceptible) ≈ 0.33

- *m(t)* approximates both

1. the occupancy measure $M^N(t)$
2. the state probability for one object at time t, drawn at random among N

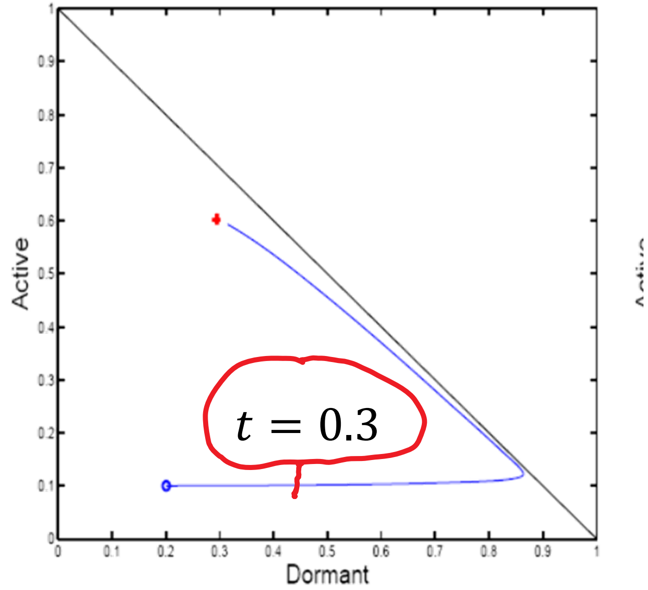

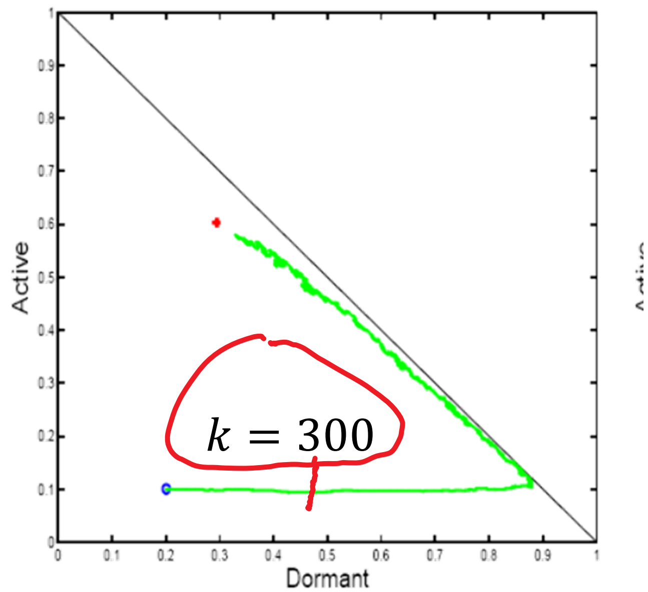

Fast Simulation

- The evolution for one object as if the other objects had a state drawn randomly and independently from the distribution *m(t)*

- Is valid *over finite horizon* whenever mean field convergence occurs

- Can be used to perform «fast simulation», i.e., simulate in detail only one or two objects, replace the rest by the mean field limit (ODE)

$$p_j^N(t|i) = \mathrm{P}(X_n^N(t) = j \,|X_n^N(0) = i)$$

$$p_j^N\left(\frac{t}{N}|i\right) \approx p_j(t|i)$$

where $\vec{p}(t|i)$ is the (transient) probability of a continuous time nonhomogeneous Markov process

$$\frac{d}{dt}\vec{p}(t|i) = \vec{p}(t|i)^T A(\vec{m}(t))$$

- Same ODE as mean field limit, with different initial condition

$$\frac{d}{dt}\vec{m}(t) = \vec{m}(t)^T A(\vec{m}(t)) = F(\vec{m}(t))$$

We can fast-simulate one node, and even compute its PDF at any time

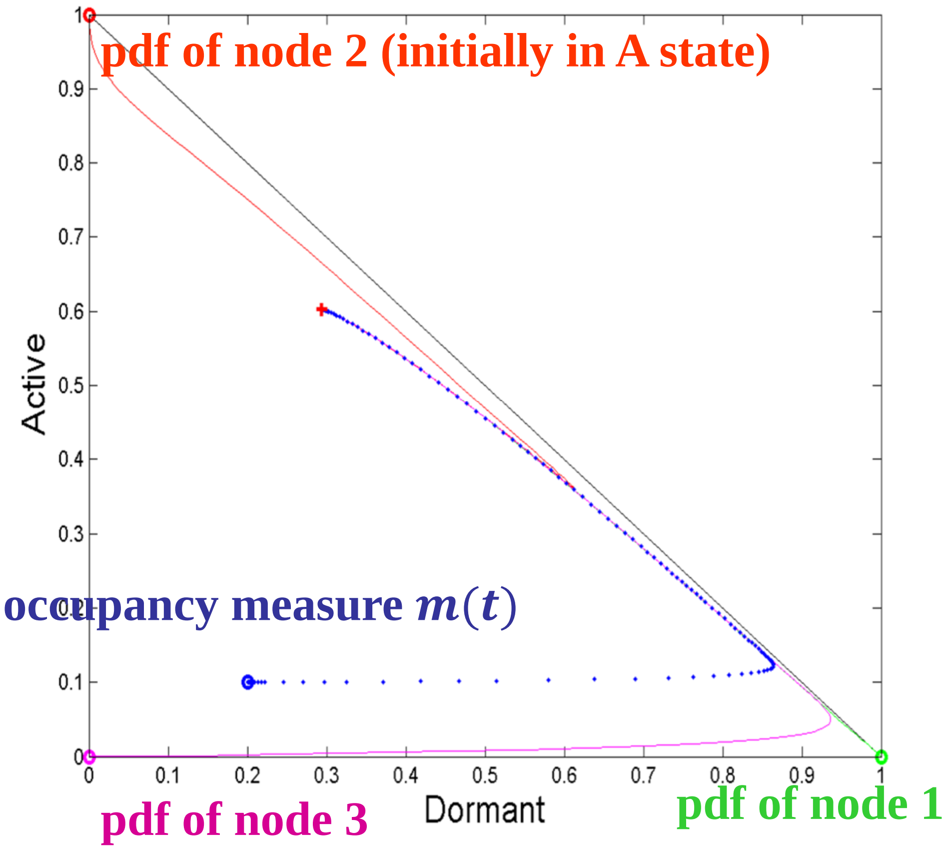

The Two Interpretations of the Mean Field Limit

m(t) is the approximation for large *N* of

1. the occupancy measure $M^N(t)$
2. the state probability for one object at time *t*, drawn at random among *N*

The state probability for one object at time t, known to be in state i at time 0, follows the same ODE as the mean field limit, but with different initial condition

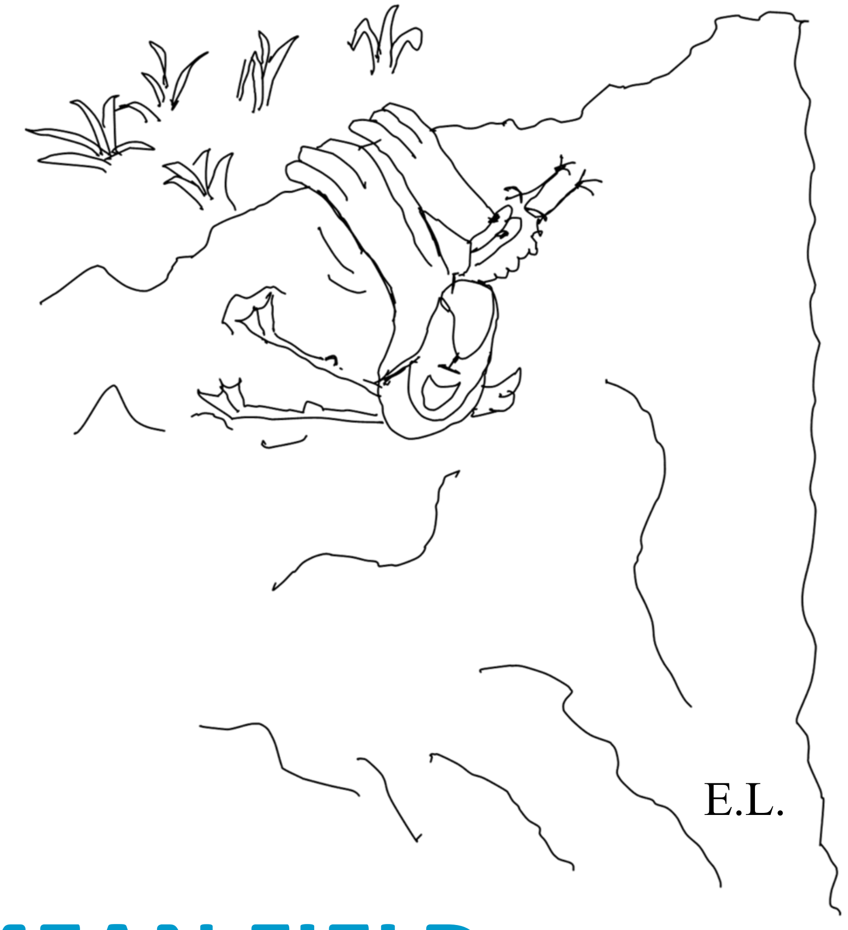

5.

CONVERGENCE TO MEAN FIELD – GENERAL CASE

There are many variants of the mean field convergence result of Section 2

- As long as state space is finite, results remain simple
- Example: «Kurtz's theorem»: time is discrete and state space is finite [Kurtz(1970), Sandholm(2006)] Let

$$f^N(m) \stackrel{\text{def}}{=} \frac{1}{I(N)} \mathbb{E}\left(M^N(k+1) - m \middle| M^N(k) = m\right)$$

$$A^N(m) \stackrel{\text{def}}{=} \frac{1}{I(N)} \mathbb{E}\left(\left\|M^N(k+1) - m\right\| \middle| M^N(k) = m\right)$$

$$B^N(m) \stackrel{\text{def}}{=} \frac{1}{I(N)} \mathbb{E}\left(\left\|M^N(k+1) - m\right\| \mathbf{1}_{\{\|M^N(k+1)-m\|>\delta_N\}} \middle| M^N(k) = m\right)$$

- $\lim_N \sup_m \left\|f^N(m) - f(m)\right\| = 0$ for some f,
  $\sup_N \sup_m A^N(m) < \infty$
  $\lim_N \sup_m \left\|B^N(m)\right\| = 0$ with $\lim_{N\to\infty} \delta_N = 0$
- $M^N(0) \to m_0$ in probability

Then $\sup_{0\le t\le T} \mathbb{P}\left(\left\|M^N(t) - m(t)\right\|\right) \to 0$ in probability.

«Kurtz's Theorem» is another Classical Result for Convergence to Mean Field

- Original Sytem is in discrete time and I(N) -> 0; limit is in continuous time
- State space for one object is finite

[Kurtz(1970), Sandholm(2006)] Let

$$f^N(m) \stackrel{\text{def}}{=} \frac{1}{I(N)} \mathbb{E}\left(M^N(k+1) - m \middle| M^N(k) = m \right)$$

$$A^N(m) \stackrel{\text{def}}{=} \frac{1}{I(N)} \mathbb{E}\left(\left\| M^N(k+1) - m \right\| \middle| M^N(k) = m \right)$$

$$B^N(m) \stackrel{\text{def}}{=} \frac{1}{I(N)} \mathbb{E}\left(\left\| M^N(k+1) - m \right\| \mathbf{1}_{\{\|M^N(k+1)-m\| > \delta_N\}} \middle| M^N(k) = m \right)$$

- $\lim_N \sup_m \left\| f^N(m) - f(m) \right\| = 0$ for some f,
  $\sup_N \sup_m A^N(m) < \infty$
  $\lim_N \sup_m \left\| B^N(m) \right\| = 0$ with $\lim_{N\to\infty} \delta_N = 0$
- $M^N(0) \to m_0$ in probability

Then $\sup_{0\le t\le T} \mathbb{P}\left(\left\| M^N(t) - m(t) \right\| \right) \to 0$ in probability.

Discrete Time, Discrete Time Limit when I(N)=O(1)

[Le Boudec et al.(2007)Le Boudec, McDonald, and Mundinger, Tinnakornsrisuphap and Makowski(2003)]
$\lim_N I(N) = 1$

- Object i draws next state at time k independent of others with transition matrix $K^N(M^N)$
- $M^N(0) \to m_0$ a.s. [in probability]
- regularity assumption on the drift (generator)

Then $\sup_{0 \le k \le K} \mathbb{P}\left(\left\|M^N(k) - m(k)\right\|\right) \to 0$ a.s. [in probability]

Extension to a Resource

- Model can be complexified by adding a global resource *R(t)*

- Slow: *R(t)* is expected to change state at the same rate *I(N)* as one object

  ⇒ call it an object of a special class

- Fast: *R(t)* changes state at the aggregate rate $N\ I(N)$

  ⇒ (easy) extensions of the theory

  [Benaïm and Le Boudec(2008)] [Bordenave et al.(2007)Bordenave, McDonald, and Proutiere]

General State Space: The Mean Field Limit is no longer an ODE

- Every taxi has a state
  - Position in area c = 0 … 16
  - Age of last received info

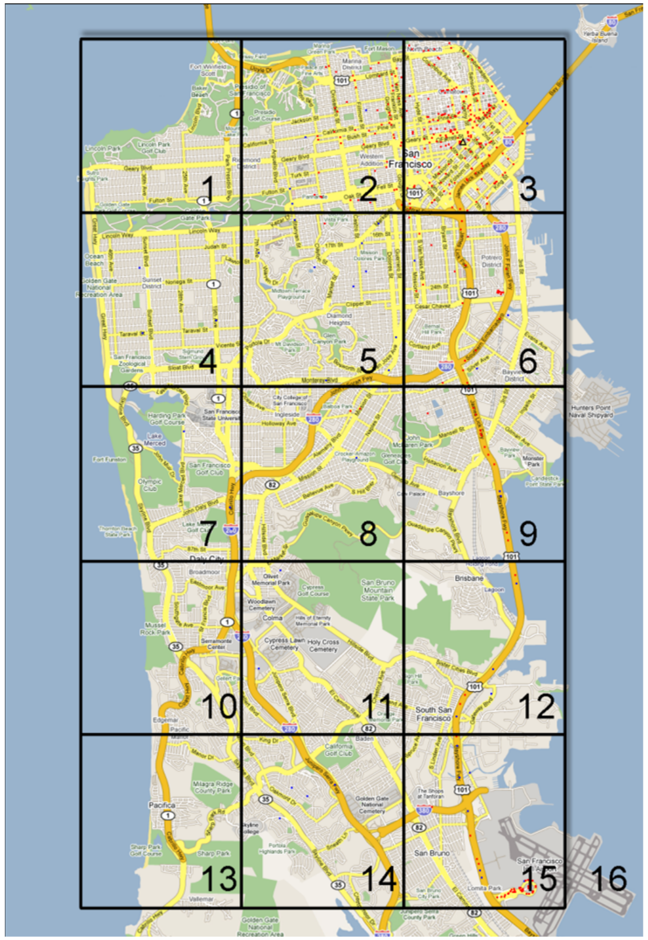

- Occupancy measure is $F_c(z,t)$ = proportion of nodes that at location c and have age ≤ z
- Mean Field Equations:

$$
\begin{cases}
\forall c \in \mathcal{C}\,, & \dfrac{\partial F_c(z,t)}{\partial t} + \dfrac{\partial F_c(z,t)}{\partial z} = \\
& \displaystyle\sum_{c' \neq c} \rho_{c',c} F_{c'}(z,t) - \left(\sum_{c' \neq c} \rho_{c,c'} \right) F_c(z,t) \\
& + \left(u_c(t|d) - F_c(z,t)\right)\left(2\eta_c F_c(z,t) + \mu_c\right) \\
& \displaystyle + \left(u_c(t|d) - F_c(z,t)\right) \sum_{c' \neq c} 2\beta_{\{c,c'\}} F_{c'}(z,t) \\
\forall c \in \mathcal{C}\,, & \forall t \geq 0\,,\ F_c(0,t) = 0 \\
\forall c \in \mathcal{C}\,, & \forall z \geq 0\,,\ F_c(z,0) = F_c^0(z)\,.
\end{cases}
$$

General State Space: Convergence to Mean Field

- There *is* convergence to mean field

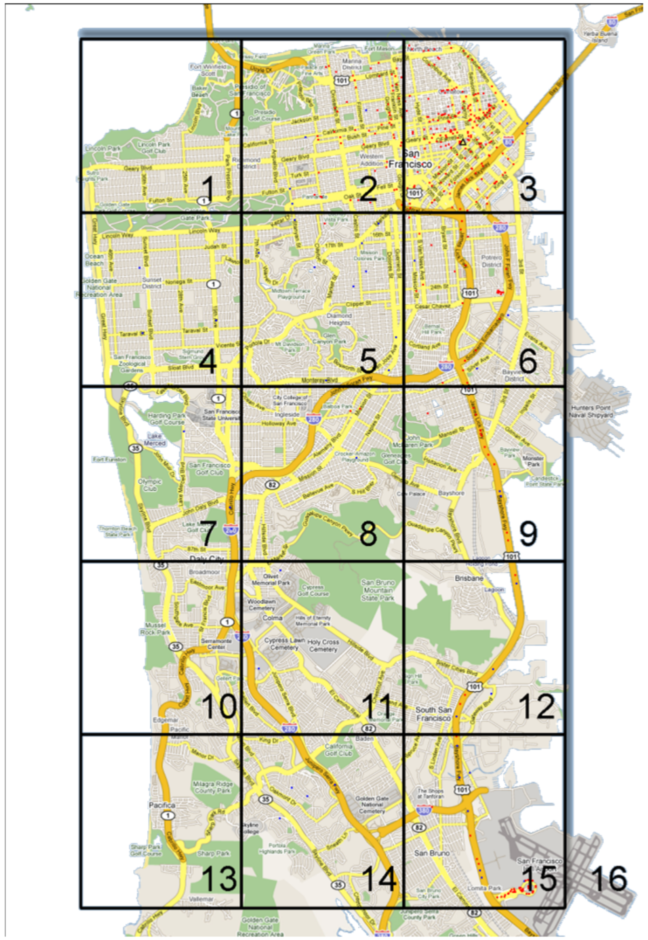

- [Graham and Méléard 1997] applies, i.e. mean field convergence occurs for iid initial conditions
- [Chaintreau et al.(2009)Chaintreau, Le Boudec, and Ristanovic] for arbitrary initial conditions

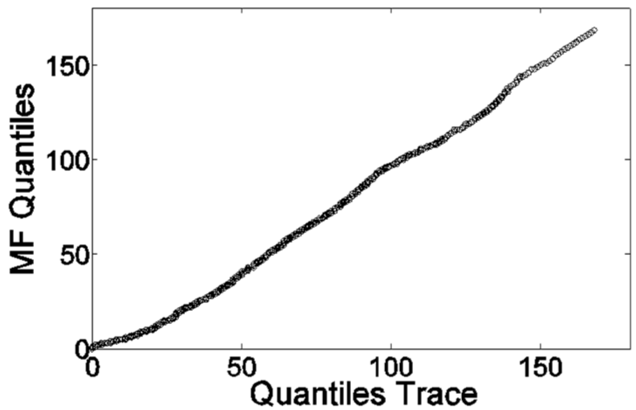

Qqplots simulation vs mean field

General State Space : A Generic Mean Field Convergence Result

- «Graham and Méléard: A generic result for **general** state space (in particular non enumerable).

[Graham and Méléard(1997), Graham and Méléard(1994)]
$I(N) = 1/N$, continuous time.

- Object i has a free evolution plus pairwise interactions.
- $X_n^N(0)_{n=1...N}$ is iid with common law m_0
- Generator of pairwise meetings is uniformly bounded in total variation norm
  e.g. if $\mathcal{G} \cdot \varphi(x, x') = \int \varphi(y, y') f(y, y' | x, x') dy dy'$ then
  $\int |f(y, y'|x, x')| \, dy dy' \leq \Lambda$, for all x, x'

Then there is propagation of chaos with explicit bounds in total variation over finite time intervals. Mean field independence holds.

When things get (surprisingly hard): The Bounded Confidence Model

- Introduced in [Deffuant et al (2000)], used in mobile networks in [Buchegger and Le Boudec 2002]; Proof of convergence to Mean Field in [Gomez, Graham, Le Boudec 2010]

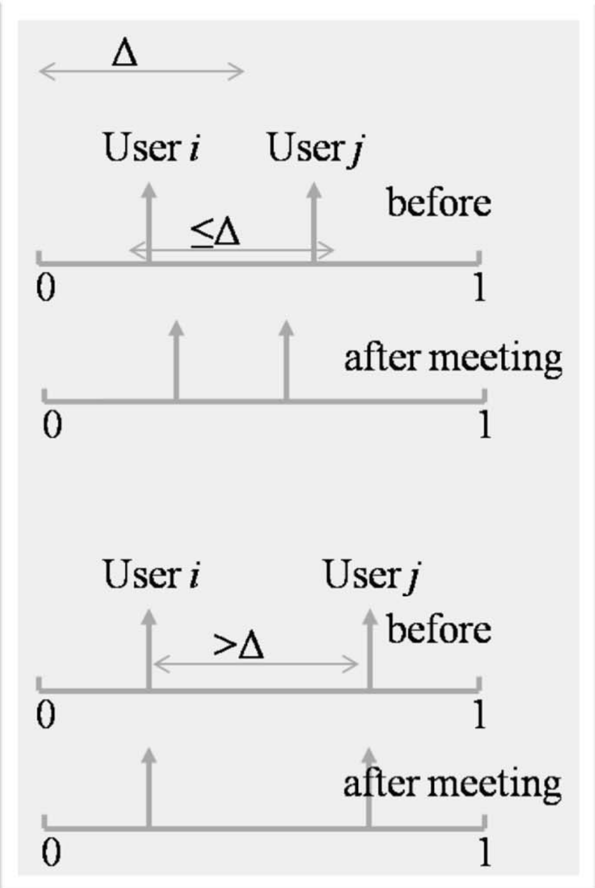

- Discrete time. State space =[0, 1].
  $X_n^N(k) \in [0,1]$ rating of common subject held by peer n
- Two peers, say i and j are drawn uniformly at random.
  If $\left|X_i^N(k) - X_j^N(k)\right| > \Delta$ no change; else

$$X_i^N(k+1) = wX_i^N(k) + (1-w)X_j^N(k),$$
$$X_j^N(k+1) = wX_j^N(k) + (1-w)X_i^N(k),$$

PDF of Mean Field Limit

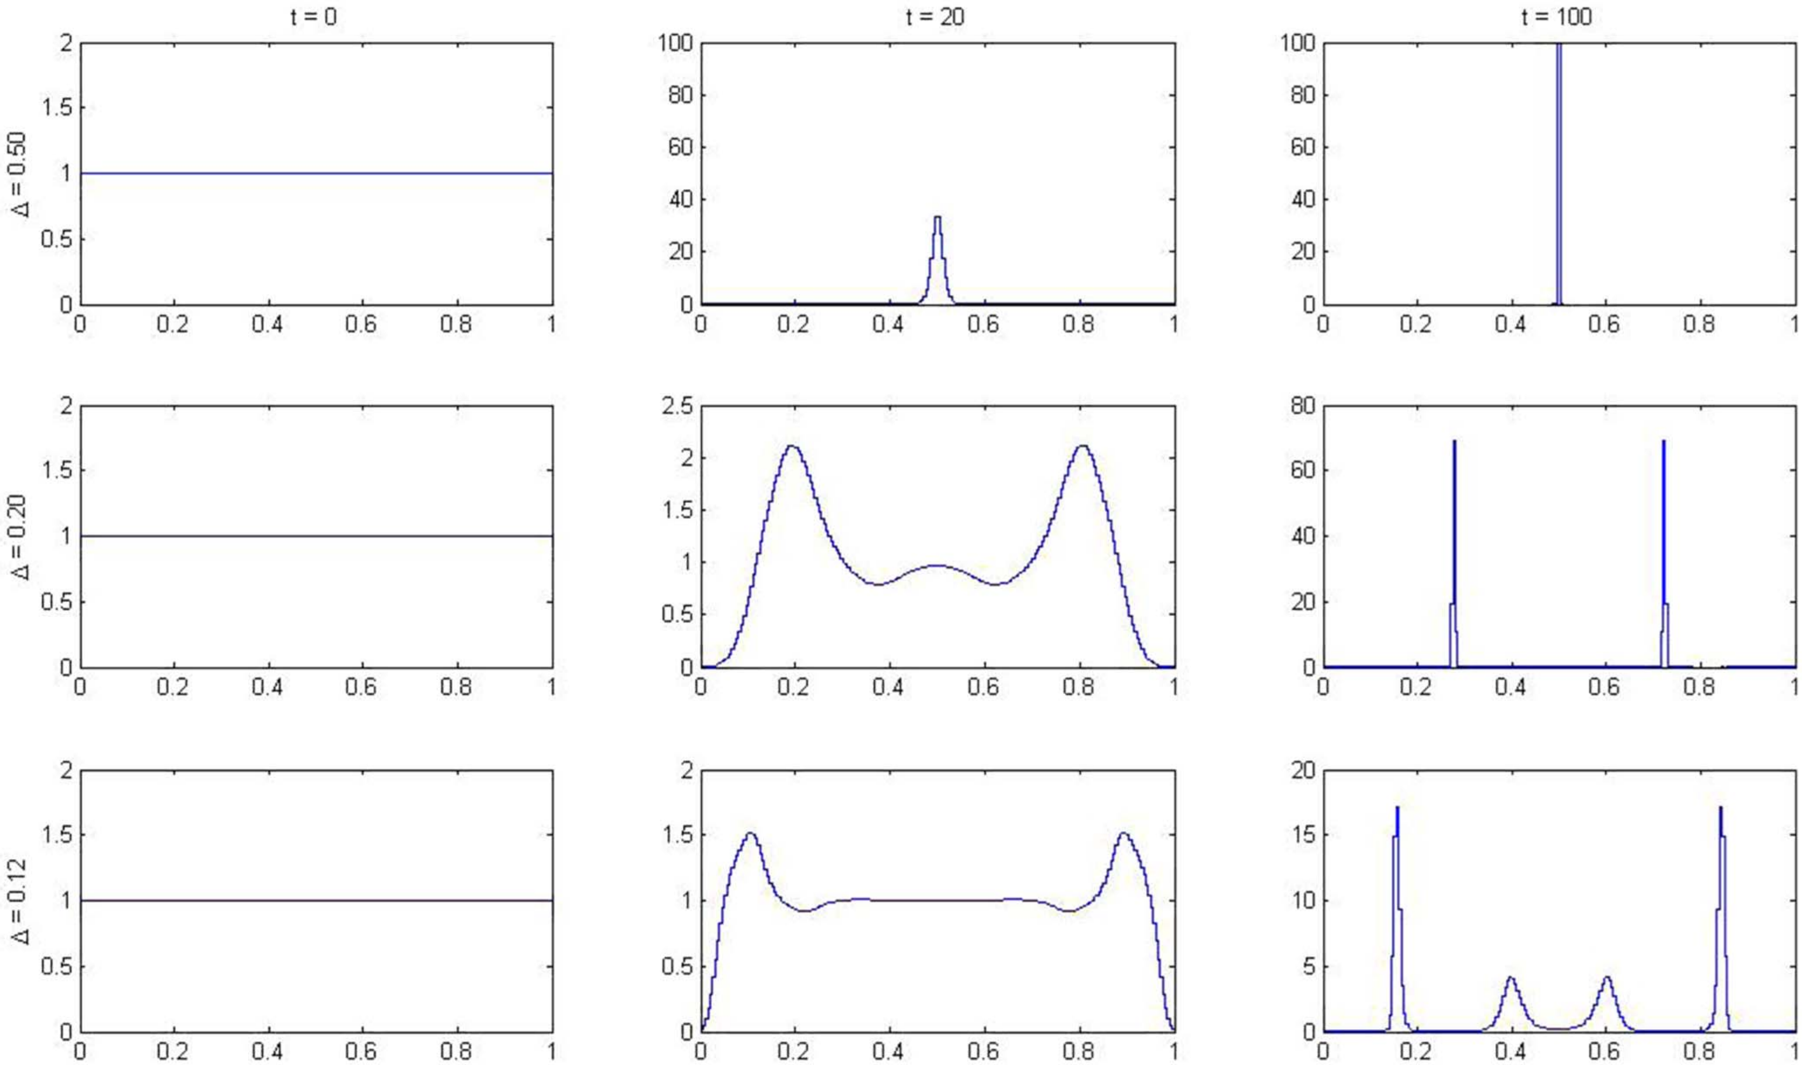

Is There Convergence to Mean Field ?

- Yes for the discretized version of the problem
  - Replace ratings in [0,1] by fixed point real numbers on *d* decimal places
  - The number of meetings is upper bounded by a constant, here 2 (Section 3)
  - There is convergence for any initial condition such that $M^N(0) \to m_0$
- This is what any simulation implements

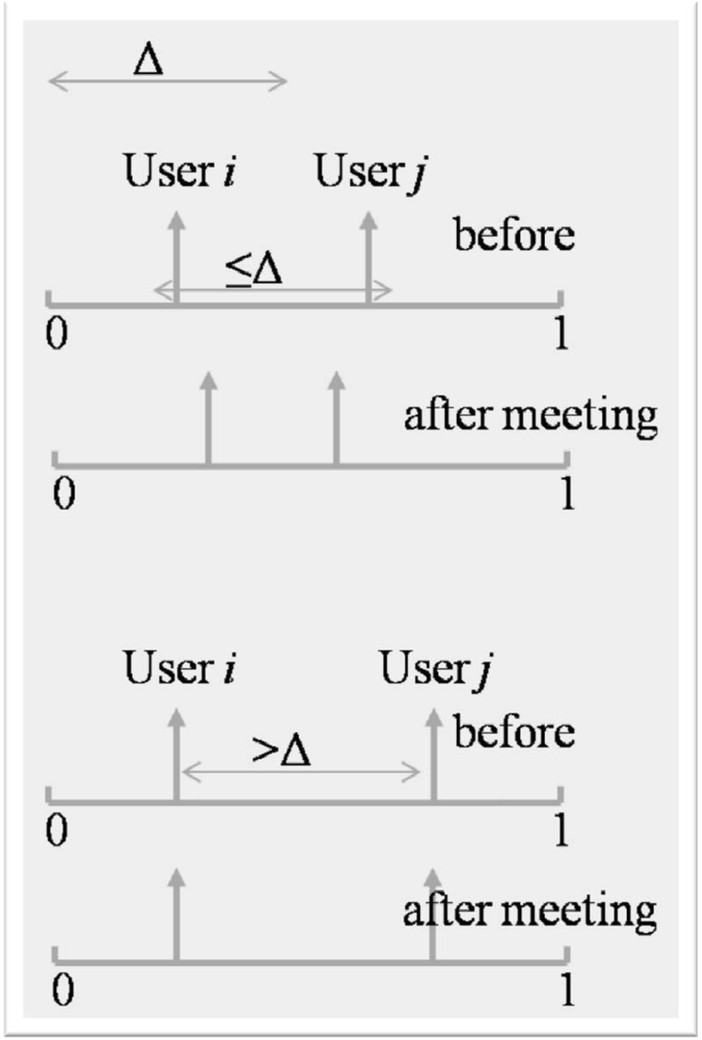

Is There Convergence to Mean Field ?

- There can be no similar result for the real version of the problem
  - Counter Example: $M^N(0)$ -> $m(0)$ (in the weak topology) but $M^N(t)$ does not converge to $m(t)$

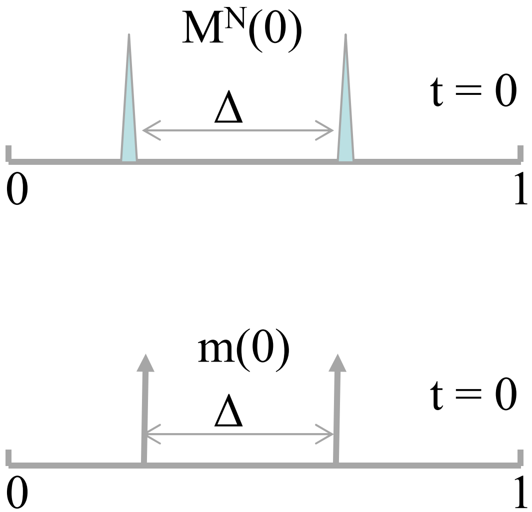

- There *is* convergence to mean field if initial condition is iid from m_0 [Gomez et al, 2010]

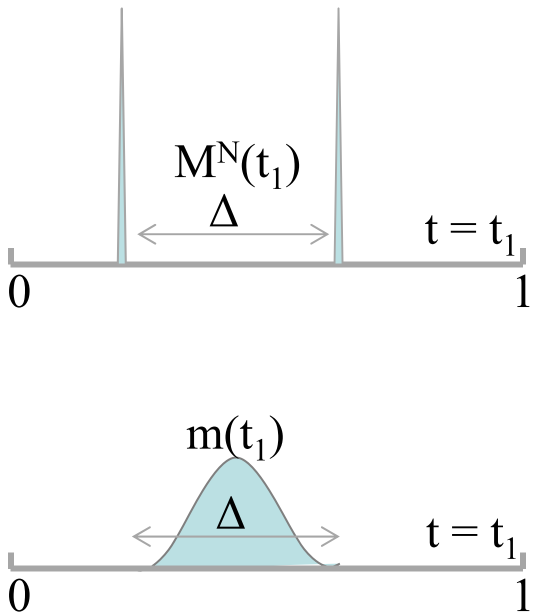

Convergence to Mean Field

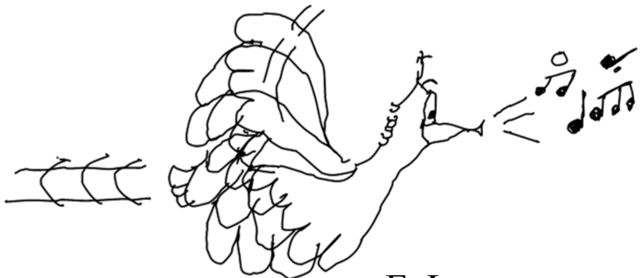

E. L.

- Thus:

For the finite state space case, most cases are verifiable by inspection of the model

- For the general state space, things may be more complex fluid limit is not an ODE there may be no convergence to mean field

E.L.

Thank You …

References

[Baccelli et al.(2004)Baccelli, Lelarge, and McDonald] F. Baccelli, M. Lelarge, and D McDonald. Metastable regimes for multiplexed tcp flows. In *Proceedings of the Forty-Second Annual Allerton Conference on Communication, Control, and Computing, Allerton House, Monticello, Illinois, USA, University of Illinois at Urbana-Champaign*, October 2004.

[Benaim and Le Boudec(2008)] 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.

[Benaïm and Weibull(2003)] M. Benaïm and J. Weibull. Deterministic approximation of stochastic evolution. *Econometrica*, 71:873–904, 2003.

[Benaim et al.(2006)Benaim, Hofbauer, and Sorin] M. Benaim, J. Hofbauer, and S. Sorin. Stochastic approximations and differential inclusions ii: Applications. 2006.

[Bordenave et al.(2007)Bordenave, McDonald, and Proutiere] C. Bordenave, D. McDonald, and A. Proutiere. A particle system in interaction with a rapidly varying environment: Mean field limits and applications. *Arxiv preprint math/0701363*, 2007.

[Bordenave et al.(2008)Bordenave, McDonald, and Proutiere] C. Bordenave, D. McDonald, and A. Proutiere. Performance of random medium access control, an asymptotic approach. In *Proceedings of the 2008 ACM SIGMETRICS international conference on Measurement and modeling of computer systems*, pages 1–12. ACM, 2008.

[Buchegger and Le Boudec(2002)] S. Buchegger and J.-Y. Le Boudec. Performance analysis of the confidant protocol (cooperation of nodes - fairness in dynamic ad-hoc networks). In *Proceedings of MobiHoc'02*, June 2002.

[Chaintreau et al.(2009)Chaintreau, Le Boudec, and Ristanovic] A. Chaintreau, J.Y. Le Boudec, and N. Ristanovic. The age of gossip: spatial mean field regime. In *Proceedings of the eleventh international joint conference on Measurement and modeling of computer systems*, pages 109–120. ACM, 2009.

[Darling and Norris(2008)] RWR Darling and J.R. Norris. Differential equation approximations for Markov chains. *Probability surveys*, 5:37–79, 2008.

[Deffuant et al.(2000)Deffuant, Neau, Amblard, and Weisbuch] G. Deffuant, D. Neau, F. Amblard, and G. Weisbuch. Mixing beliefs among interacting agents. *Advances in Complex Systems*, 3:87–98, 2000.

[Ethier and Kurtz(2005)] Stewart N. Ethier and Thomas G. Kurtz. *Markov Processes, Characterization and Convergence*. Wiley, 2005.

[Gast et al.(2010)Gast, Gaujal, and Le Boudec] Nicolas Gast, Bruno Gaujal, and Jean-Yves Le Boudec. Mean field for Markov Decision Processes: from Discrete to Continuous Optimization. Technical Report arXiv:1004.2342v2, 2010.

[Gomez-Serrano et al, 2012] Gomez-Serrano J., Graham C. and Le Boudec J.-Y.
The Bounded Confidence Model Of Opinion Dynamics
Mathematical Models and Methods in Applied Sciences, Vol. 22, Nr. 2, pp. 1150007-1--1150007-46, 2012.

[Graham and Méléard(1994)] Carl Graham and Sylvie Méléard. Chaos hypothesis for a system interacting through shared resources. *Probab. Theory Related Fields*, 100(2):157–173, 1994. ISSN 0178-8051. doi: 10.1007/BF01199263. URL `http://dx.doi.org/10.1007/BF01199263`.

[Graham and Méléard(1997)] Carl Graham and Sylvie Méléard. Stochastic particle approximations for generalized Boltzmann models and convergence estimates. *Ann. Probab.*, 25(1):115–132, 1997. ISSN 0091-1798. doi: 10.1214/aop/1024404281. URL `http://dx.doi.org/10.1214/aop/1024404281`.

[Ioannidis and Marbach(2009)] S. Ioannidis and P. Marbach. Absence of Evidence as Evidence of Absence: A Simple Mechanism for Scalable P2P Search. In *INFOCOM 2009, IEEE*, pages 576–584. IEEE, 2009.

[Kelly(1991)] F.P. Kelly. Loss networks. *The annals of applied probability*, 1 (3):319–378, 1991.

[Kurtz(1970)] T.G. Kurtz. Solutions of ordinary differential equations as limits of pure jump Markov processes. *Journal of Applied Probability*, 7(1): 49–58, 1970.

[Le Boudec(2010)] Jean-Yves Le Boudec. The Stationary Behaviour of Fluid Limits of Reversible Processes is Concentrated on Stationary Points. Technical Report arxiv:1009.5021.v2, 2010.

[Le Boudec et al.(2007)Le Boudec, McDonald, and Mundinger] Jean-Yves Le Boudec, David McDonald, and Jochen Mundinger. A Generic Mean Field Convergence Result for Systems of Interacting Objects. In *QEST'07*, 2007.

[McDonald(2007)] David McDonald. Lecture Notes on Mean Field Convergence, March 2007.

[Sandholm(2006)] W.H. Sandholm. Population games and evolutionary dynamics. *Unpublished manuscript, University of Wisconsin*, 2006.

[Sznitman(1991)] A.S. Sznitman. Topics in propagation of chaos. In P.L. Hennequin, editor, *Springer Verlag Lecture Notes in Mathematics 1464, Ecole d'Eté de Probabilités de Saint-Flour XI (1989)*, pages 165Ű–251, 1991.

[Tembine et al.(2009)Tembine, Le Boudec, El-Azouzi, and Altman] Hamidou Tembine, Jean-Yves Le Boudec, Rachid El-Azouzi, and Eitan Altman. Mean Field Asymptotic of Markov Decision Evolutionary Games and Teams. In *Gamenets 2009*, 2009. Invited Paper.

[Tinnakornsrisuphap and Makowski(2003)] Peerapol Tinnakornsrisuphap and Armand M. Makowski. Limit behavior of ecn/red gateways under a large number of tcp flows. In *Proceedings IEEE INFOCOM 2003, The 22nd Annual Joint Conference of the IEEE Computer and Communications Societies, San Franciso, CA, USA*, March 30 - April 3 2003.