Mean Field Methods for Computer and Communication Systems: A Tutorial

Jean-Yves Le Boudec

EPFL

Performance 2010

Namur, November 16-19 2010

References

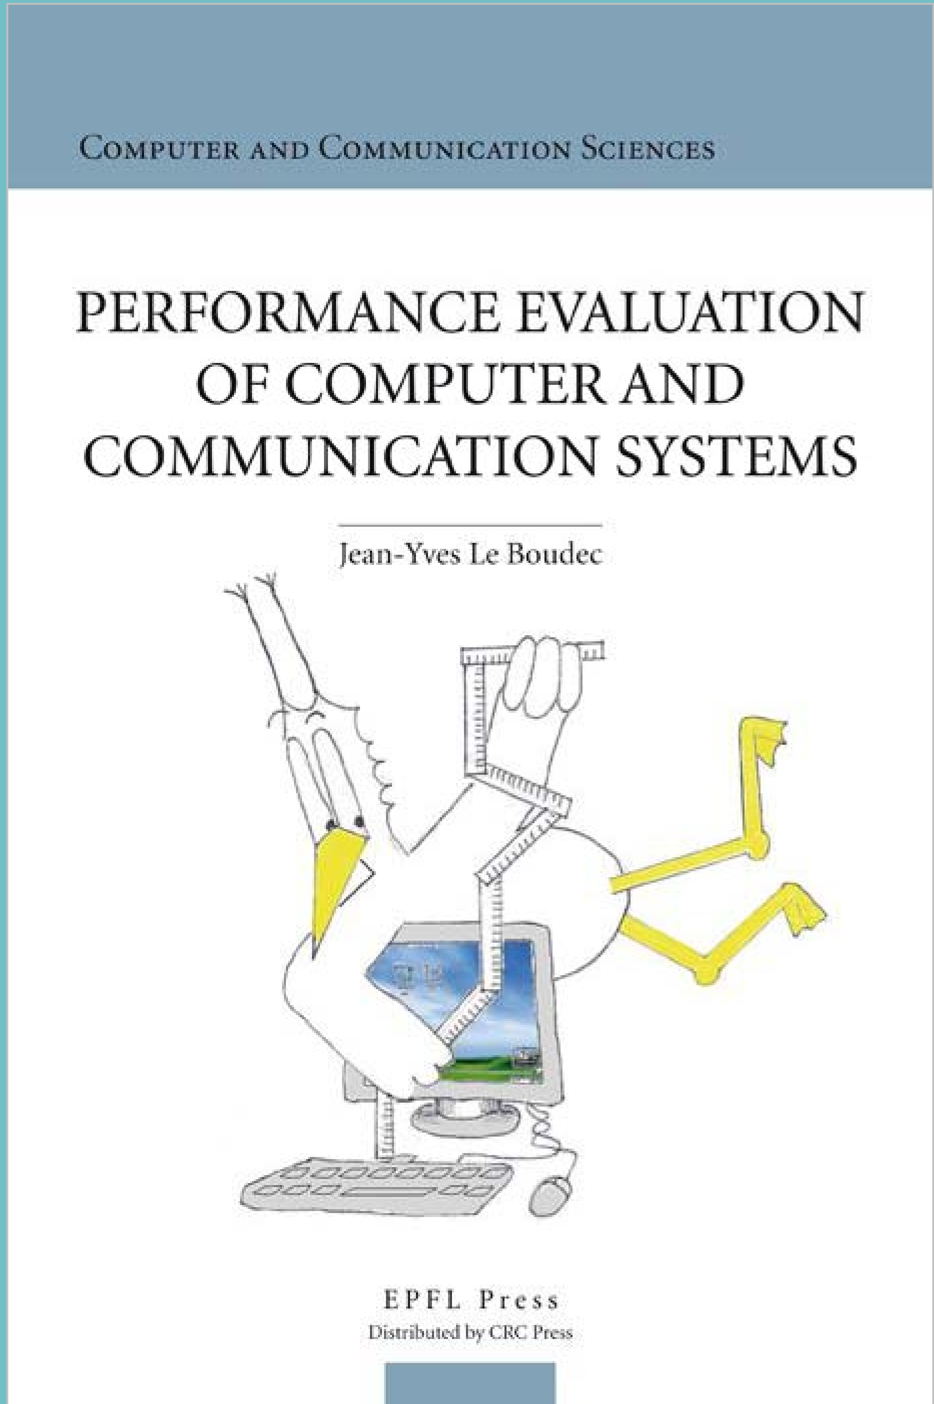

published Oct 2010

perfeval.epfl.ch

Contents

MEAN FIELD INTERACTION MODEL

Common Assumptions

- Time is discrete or continuous

- N objects
- Object n has state $X_n(t)$
- $(X^N{}_1(t), \ldots, X^N{}_N(t))$ is Markov
  => $M^N(t)$ = occupancy measure process is also Markov

- Objects can be observed only through their state

- N is large

Called "*Mean Field Interaction Models*" in the Performance Evaluation community [McDonald(2007), Benaïm and Le Boudec(2008)]

Intensity *I(N)*

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

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

- If time is discrete for X^N
  - I(N) = O(1): mean field limit is in discrete time [Le Boudec et al (2007)]
  - I(N) = O(1/N): mean field limit is in continuous time [Benaïm and Le Boudec (2008)]

Example: 2-Step Malware

- Mobile nodes are either
  - \`S' Susceptible
  - \`D' Dormant
  - \`A' Active
- Time is discrete
- Nodes meet pairwise (bluetooth)
- One interaction per time slot, $I(N) = 1/N$; mean field limit is an ODE
- State space is finite = {\`S' , \`A' ,\`D'}
- Occupancy measure is $M(t) = (S(t), D(t), A(t))$ with $S(t)+ D(t) + A(t) =1$

$S(t)$ = proportion of nodes in state \`S'

[Benaïm and Le Boudec(2008)]

- Possible interactions:

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

Simulation Runs, N=1000 nodes

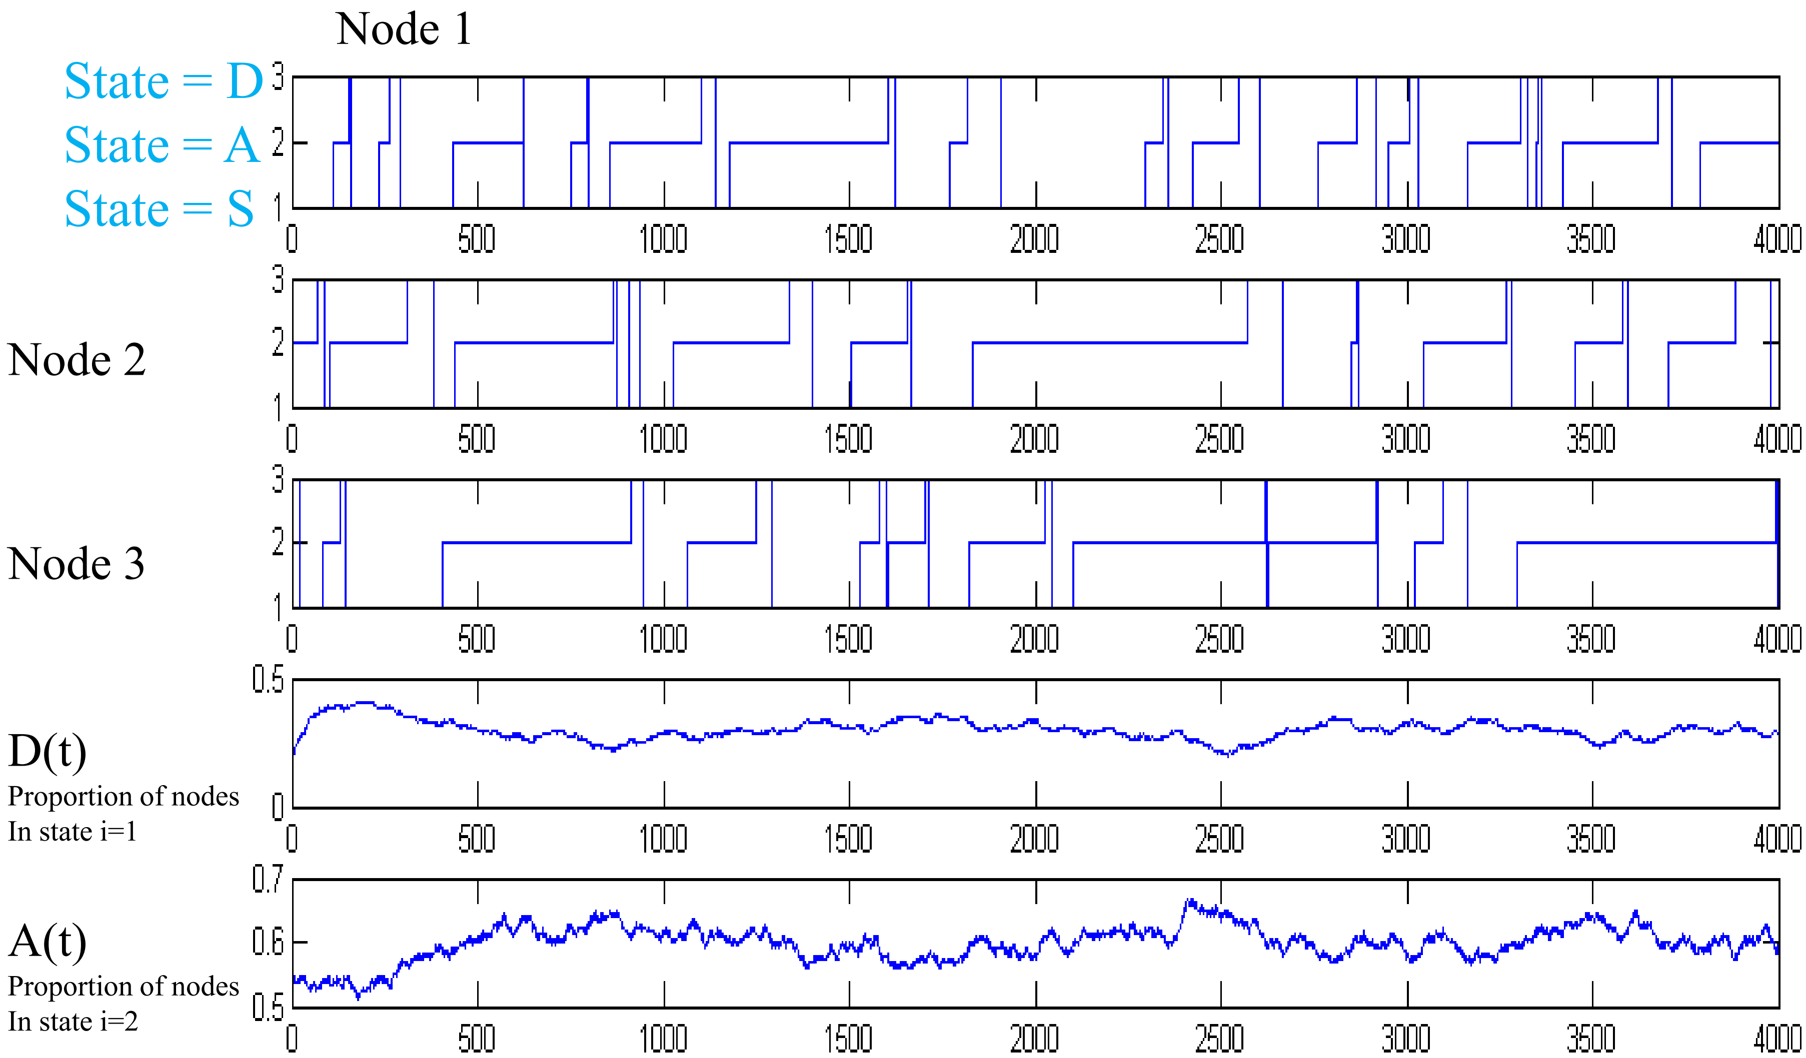

$$\beta = 0.01, \delta_A = 0.005, \delta_D = 0.0001, \alpha_0 = \alpha = 0.0001, h = 0.3, r = 0.1, \lambda = 0.0001$$

Sample Runs with N = 1000

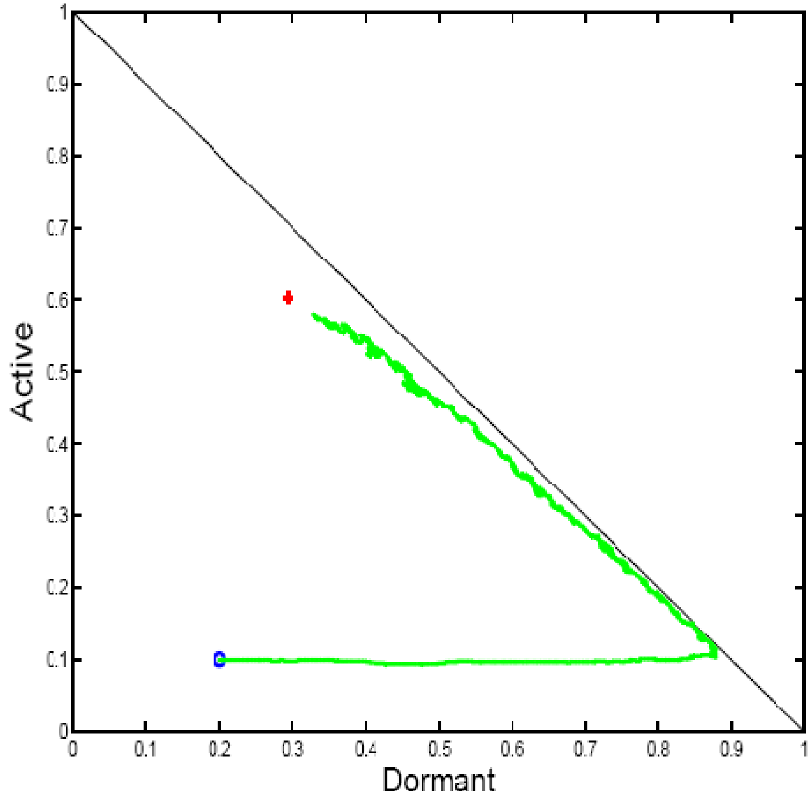

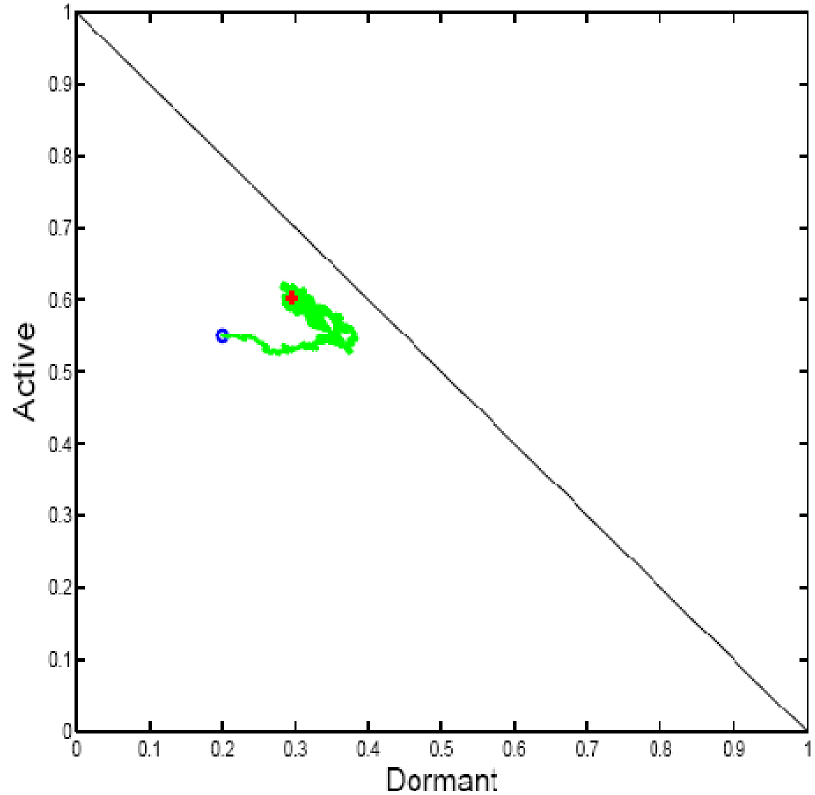

$$\beta = 0.01, \delta_A = 0.005, \delta_D = 0.0001, \alpha_0 = \alpha = 0.0001, h = 0.3, r = 0.1, \lambda = 0.0001$$

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: TCP and ECN

- [Tinnakornsrisuphap and Makowski(2003)]

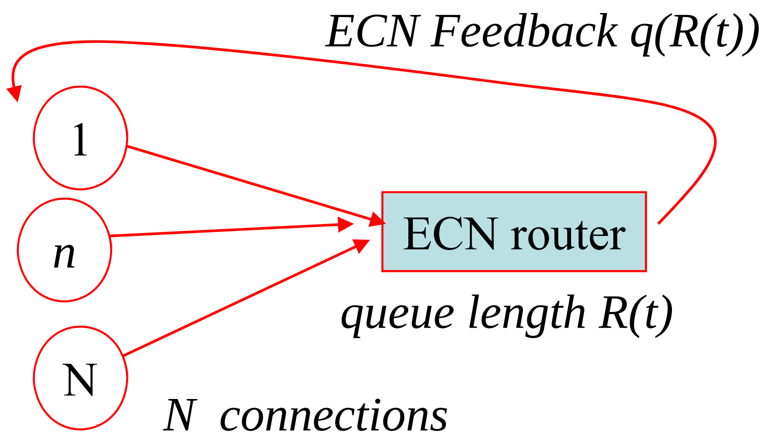

At, every time step, all connections update their state: I(N)=1

- Time is discrete, mean field limit is also in discrete time (iterated map)

- Similar examples:
  HTTP Metastability [Baccelli et al.(2004)Baccelli, Lelarge, and McDonald]

  Reputation System [Le Boudec et al.(2007)Le Boudec, McDonald, and Mundinger]

Example: Age of Gossip

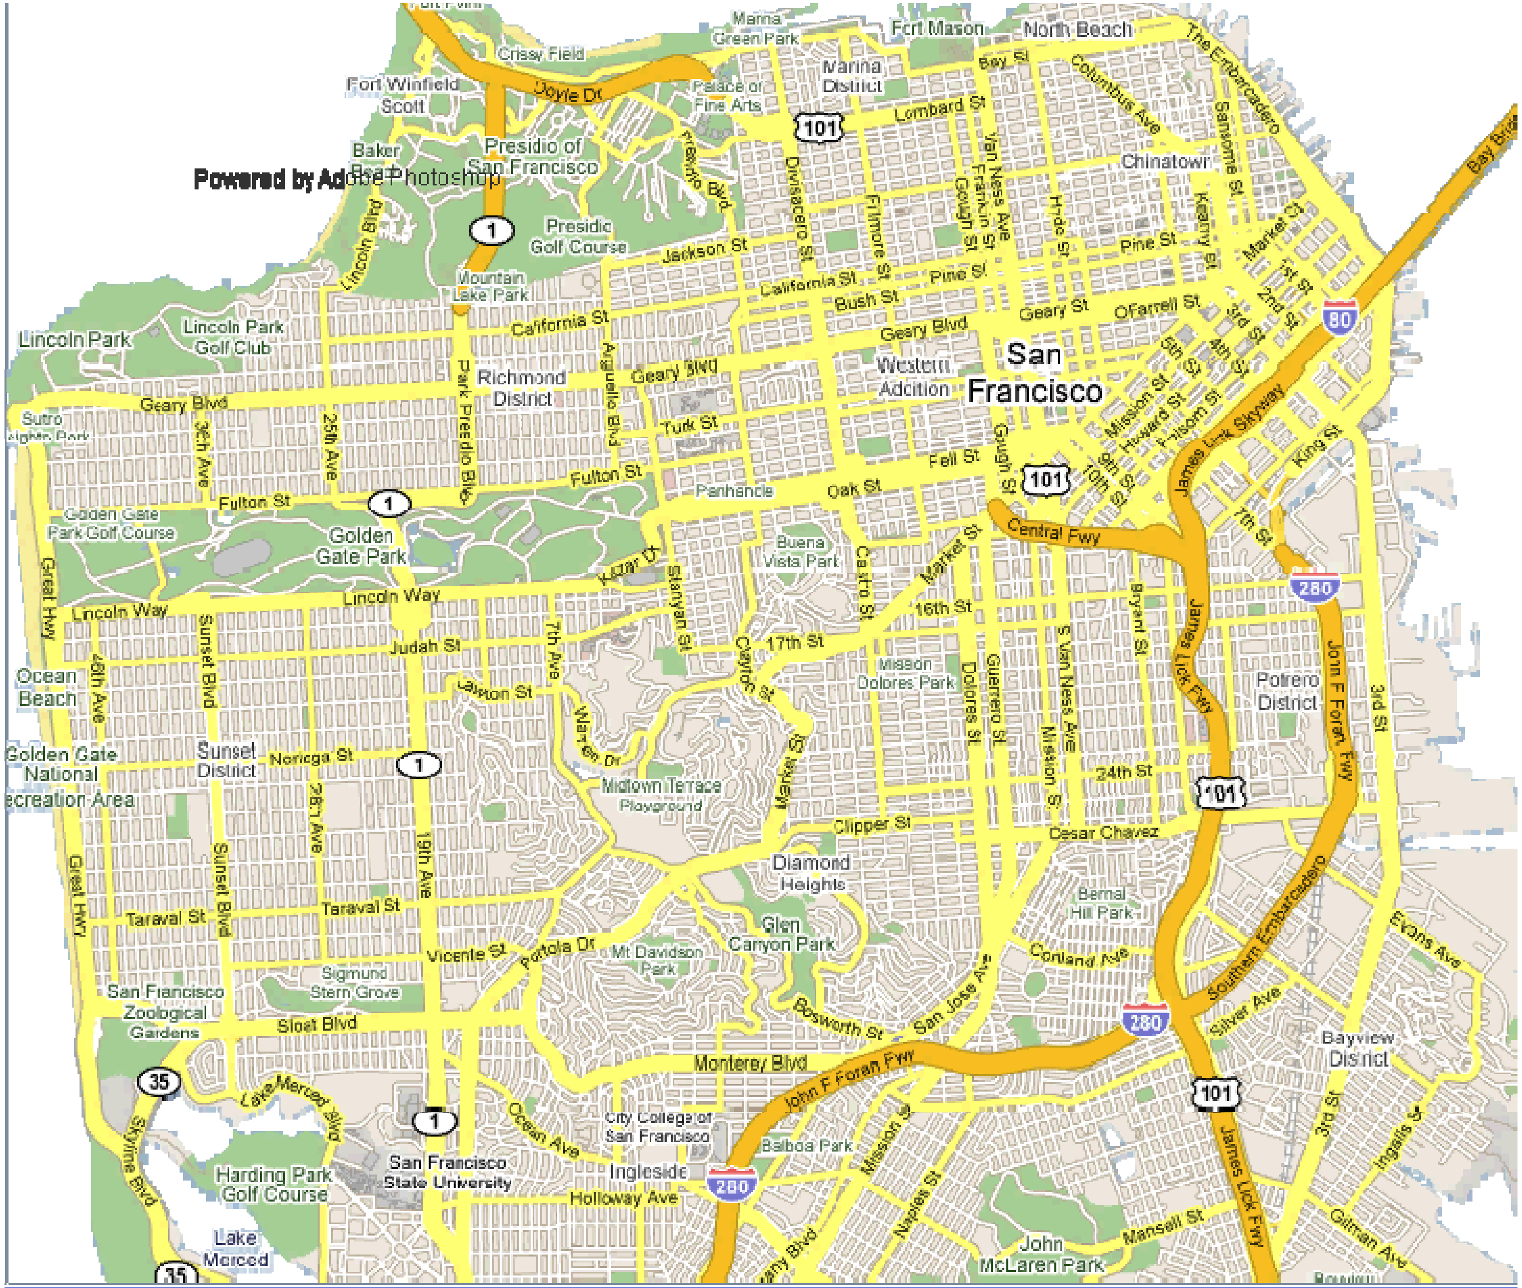

Example: Age of Gossip

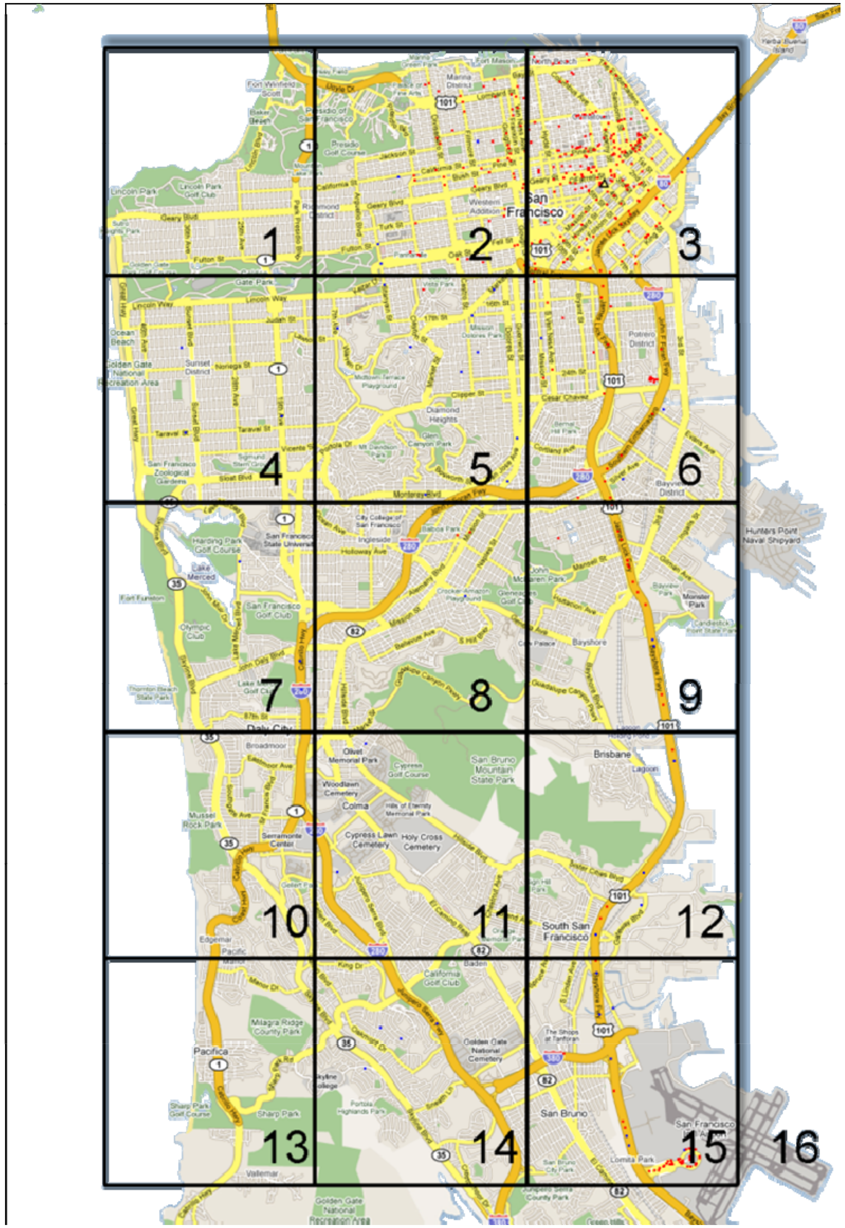

- Mobile node state = (c, t)
  c = 1 … 16 (position)
  $t \in R^+$ (age)

- Time is continuous, I(N) = 1
- Occupancy measure is
  $F_c(z,t)$ = proportion of nodes that at location c and have age ≤ z

[Chaintreau et al.(2009)]

Spatial Representation

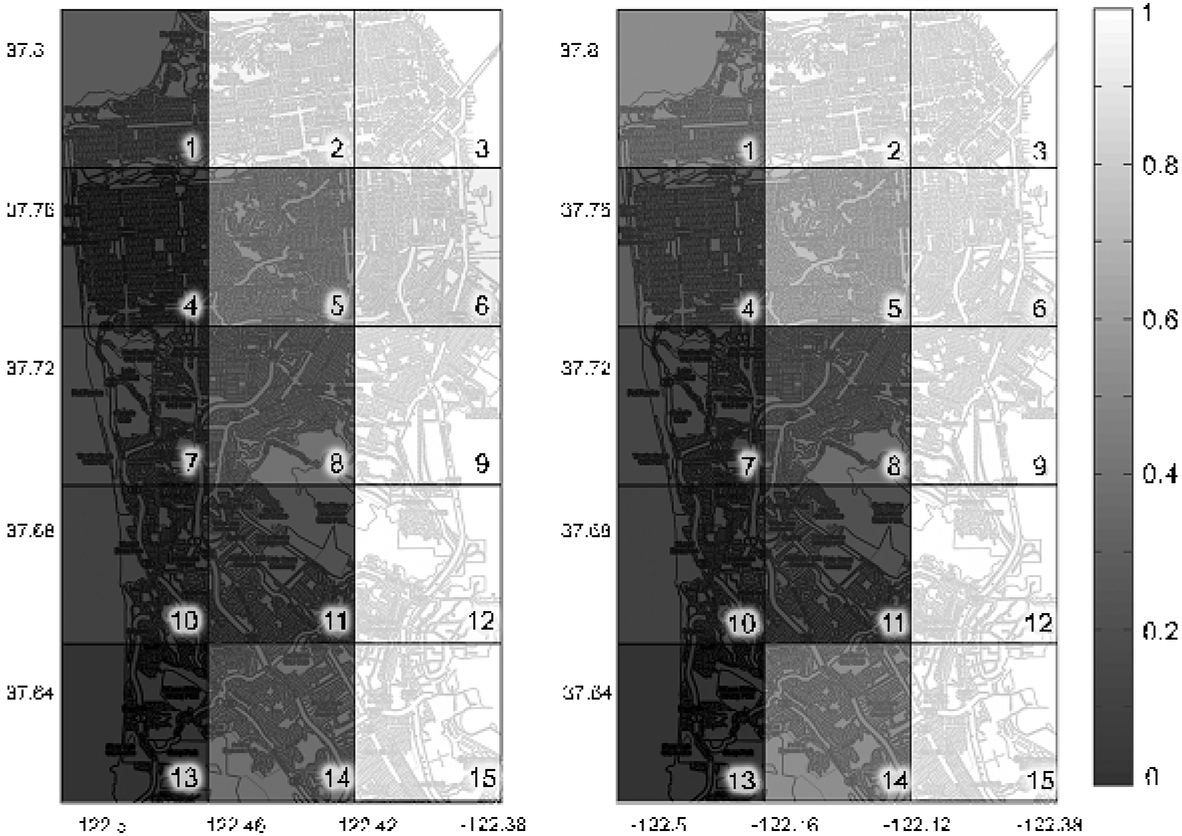

Comparison between the mean-field limit and the trace. Percentages of mobile nodes in classes 1-15 with age z<20mn at time t=300mn (1 p.m.).

The Importance of Being Spatial

- We compare the previous 16 class case with a simple 2 class case (C=2)
- The first figure suggests that for the case C=16, trace and MF data samples come from the same distribution
- For the case C=2 we observe the strong bias present for both low and high age

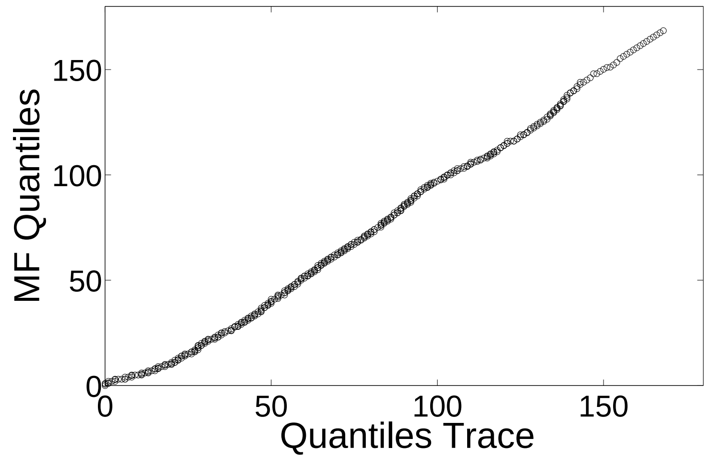

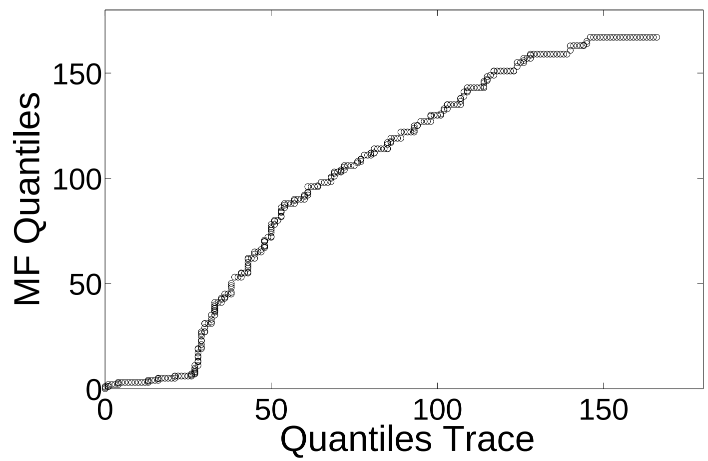

QQ plots, comparing the age distribution of trace data and data artificially obtained from the mean-field CDF, for 16 class and 2 class scenarios. Time period observed 5 p.m.-6 p.m.

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) is change state at the aggregate rate N I(N)

  -> requires special extensions of the theory

  [Bordenave et al.(2007)Bordenave, McDonald, and Proutiere]

  [Benaïm and Le Boudec(2008)]

What can we do with a Mean Field Interaction Model ?

- Large N asymptotics
  - = fluid limit
  - Markov chain replaced by a deterministic dynamical system
  - ODE
  - Fast Simulation
- Issues
  - When valid
  - Don't want do devote an entire PhD to show mean field limit
  - How to formulate the ODE

- Large t asymptotic
  - ≈ stationary behaviour
  - Useful performance metric
- Issues
  - Is stationary regime of ODE an approximation of stationary regime of original system ?
  - Does this justify the "Decoupling Assumption"?

FINITE HORIZON

MEAN FIELD LIMIT

The Mean Field Limit

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

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

- [Graham and Méléard(1994)] consider the occupancy measure L^N in path space

$$M^N(t) \stackrel{\text{def}}{=} \frac{1}{N}\sum_n \delta_{X_n^N(t)}$$

$$L^N \stackrel{\text{def}}{=} \frac{1}{N}\sum_n \delta_{X_n^N}$$

Mean Field Limit
N = +∞

Stochastic system
N = 1000

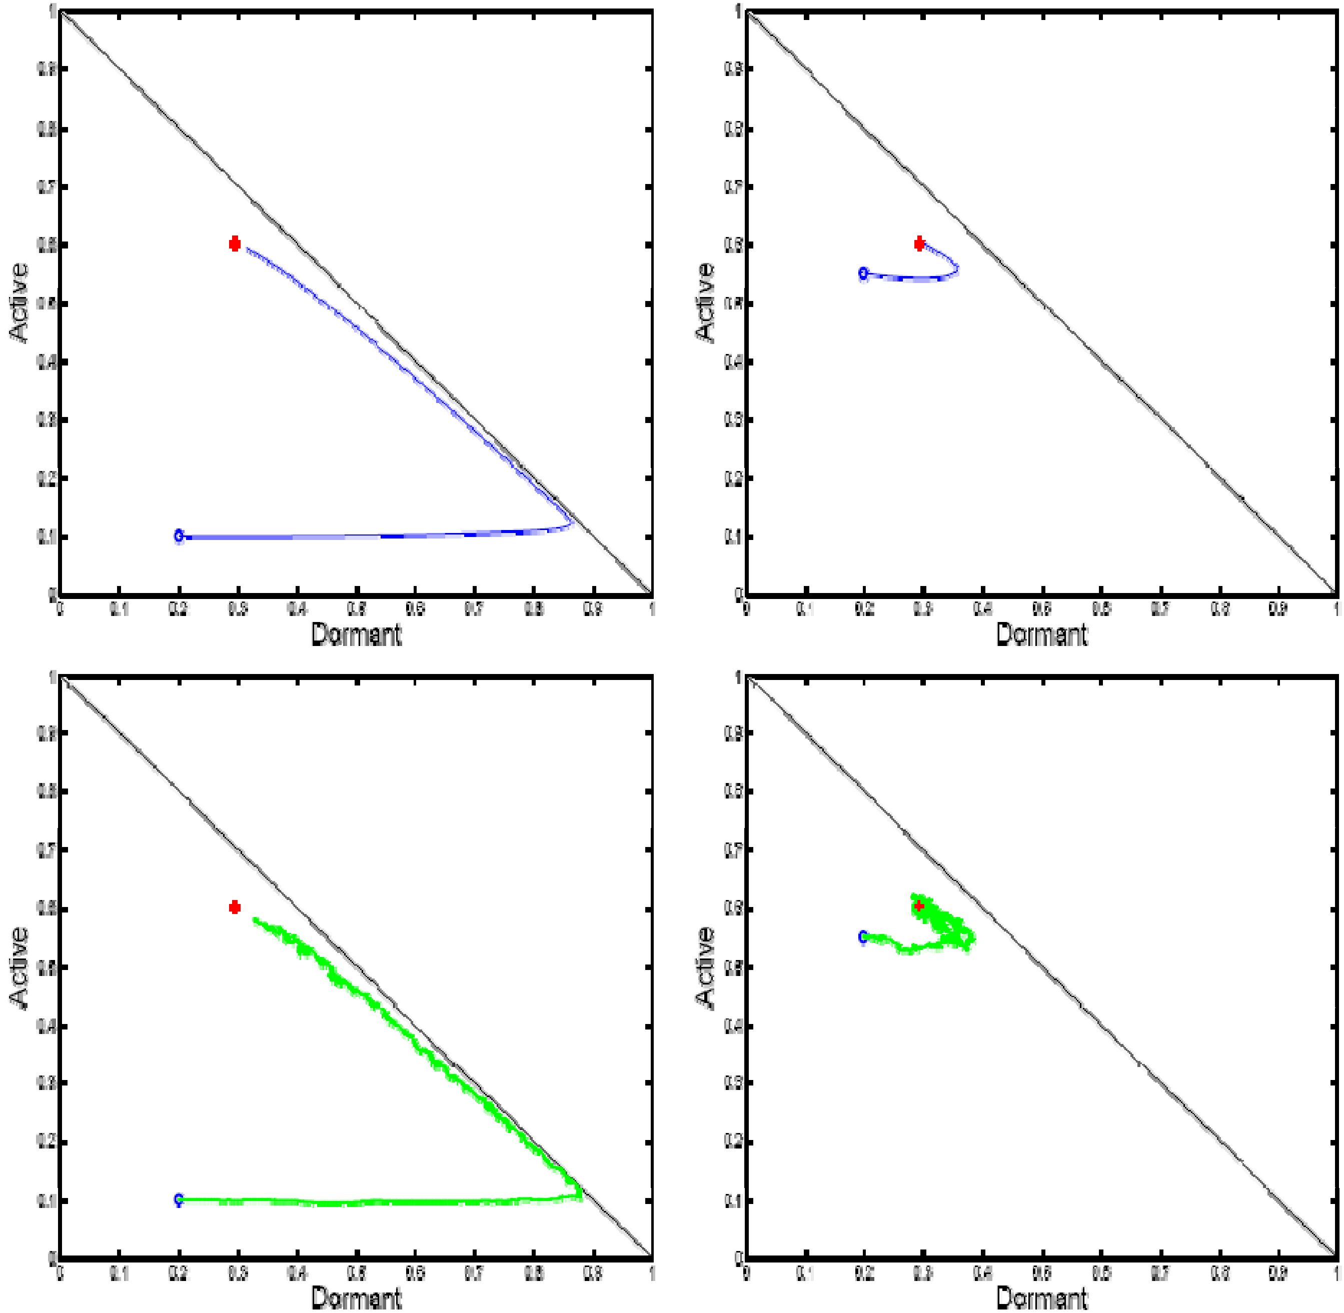

Mean Field Limit Equations

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	$S\alpha$

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

$$\frac{\partial D}{\partial t} \approx -\delta_D D - 2\lambda D^2 - \beta A\frac{D}{h+D} + (\alpha_0 + rD)S$$

$$\frac{\partial A}{\partial t} \approx 2\lambda D^2 + \beta A\frac{D}{h+D} - \delta_A A + \alpha S$$

$$\frac{\partial S}{\partial t} \approx \delta_D D + \delta_A A - (\alpha_0 + rD)S - \alpha S$$

Propagation of Chaos is Equivalent to Convergence to a Deterministic Limit

Definition 1.1 *Let $X^N = (X_1^N, ..., X_N^N)$ be an exchangeable sequence of processes in $\mathcal{P}(S)$ and $m \in \mathcal{P}(S)$ where S is metric complete separable. $(X^N)_N$ is m-chaotic iff for every k: $\mathcal{L}(X_1^N, .., X_k^N) \to m \otimes ... \otimes m$ as $N \to \infty$.*

Theorem 1.1 ([Sznitman(1991)]) *$(X^N)_N$ is m-chaotic then the occupancy measure $M^N \stackrel{\text{def}}{=} \frac{1}{N}\sum_{n=1}^{N} \delta_{X_n^N}$ converges in probability (and in law) to m.*

If the occupancy measure converges in law to m then $(X^N)_N$ is m-chaotic.

Propagation of Chaos
Decoupling Assumption

- **(Propagation of Chaos)**
  If the initial condition $(X^N_n(0))_{n=1\ldots N}$ is exchangeable and there is mean field convergence then the sequence $(X^N_n)_{n=1\ldots N}$ indexed by *N* is *m*-chaotic

  k objects are asymptotically independent with common law equal to the mean field limit, for any fixed *k*

$$\mathcal{L}\left(X_1\left(\frac{t}{I(N)}\right), \ldots, X_k\left(\frac{t}{I(N)}\right)\right) \to m(t) \otimes \ldots \otimes m(t)$$

- **(Decoupling Assumption)**
  (also called Mean Field Approximation, or Fast Simulation)
  The law of one object is asymptotically as if all other objects were drawn randomly with replacement from *m(t)*

Example: Propagation of Chaos

- 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 large t :
  - Prob (node n is dormant) ≈ 0.3
  - Prob (node n is active) ≈ 0.6
  - Prob (node n is susceptible) ≈ 0.1

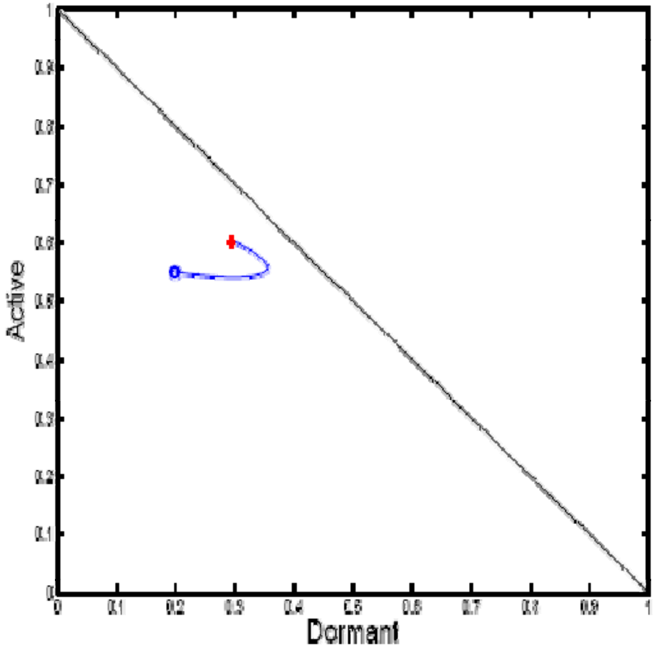

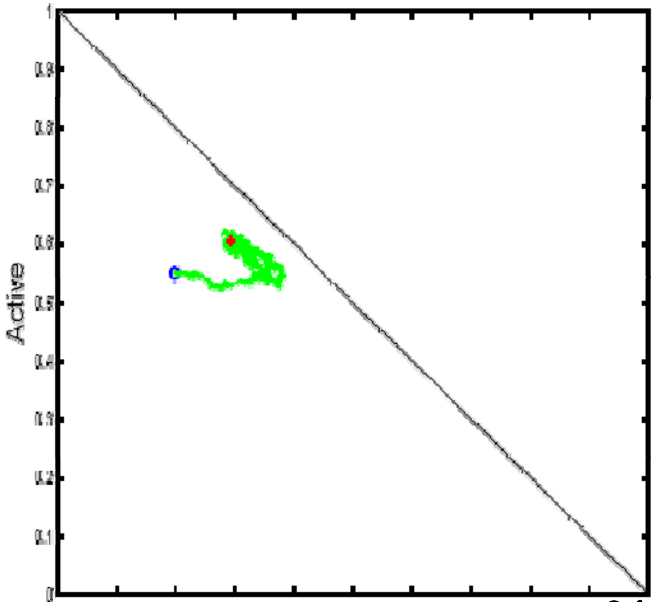

The Two Interpretations of the Mean Field Limit

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 Mean Field Approximation

- Common in Physics
- Consists in pretending that $X^N{}_m(t), X^N{}_n(t)$ are independent in the time evolution equation
- It is asymptotically true for large N, at fixed time t, for our model of interacting objects, when convergence to mean field occurs.
- Also called "decoupling assumption" (in computer science)

FINITE HORIZON

CONVERGENCE TO MEAN FIELD LIMIT

The General Case

- Convergence to the mean field limit is very often true
- A general method is known [Sznitman(1991)]:
  - Describe original system as a markov system; make it a martingale problem, using the generator
  - Show that the limiting problem is defined as a martingale problem with unique solution
  - Show that any limit point is solution of the limitingmartingale problem
  - Find some compactness argument (with weak topology)

- Requires knowing [Ethier and Kurtz(2005)]

Finite State Space per Object : Kurtz's Theorem

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

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

$$
\begin{aligned}
f^N(m) &\stackrel{\text{def}}{=} \frac{1}{I(N)}\mathbb{E}\left(M^N(k+1)-m \,\middle|\, M^N(k)=m\right) \\
A^N(m) &\stackrel{\text{def}}{=} \frac{1}{I(N)}\mathbb{E}\left(\|M^N(k+1)-m\| \,\middle|\, M^N(k)=m\right) \\
B^N(m) &\stackrel{\text{def}}{=} \frac{1}{I(N)}\mathbb{E}\left(\|M^N(k+1)-m\|\, 1_{\{\|M^N(k+1)-m\|>\delta_N\}} \,\middle|\, M^N(k)=m\right)
\end{aligned}
$$

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

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

Discrete Time, Finite State Space per Object

- Refinement + simplification, with a fast resource

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

When limit is non continuous:

[Benaim et al.(2006)Benaim, Hofbauer, and Sorin]

Example: Convergence to Mean Field

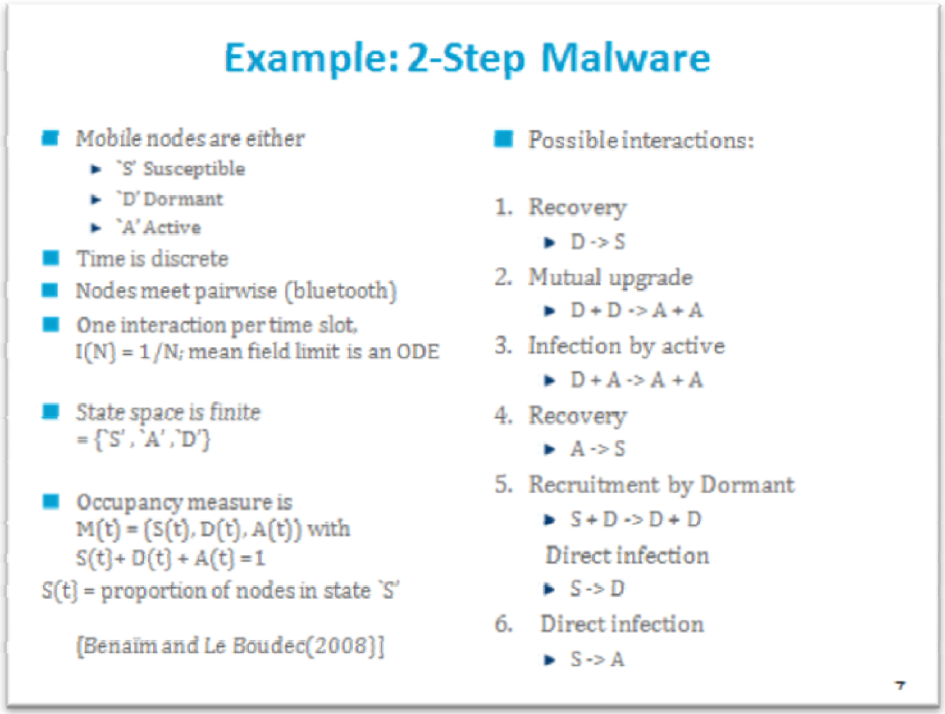

- Rescale time such that one time step = 1/N
- Number of transitions per time step is bounded by 2, therefore there is convergence to mean field

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

Discrete Time, Enumerable State Space per Object

- State space is enumerable with discrete topology, perhaps infinite; with a fast resource

  [Bordenave et al.(2007)Bordenave, McDonald, and Proutiere]

  - Probability that objects i and j do a transition in one time slot is $o(1/N)$
  - $M^N(0) \to m(0)$ in probability for the weak topology
  - $(X_1^N(0), \dots, X_N^N(0))$ is exchangeable at time 0
  - regularity assumption on the drift (generator)

  Then M^N is m-chaotic.

- Essentially : same as previous plus exchangeability at time 0

Discrete Time, Discrete Time Limit

- Mean field limit is in discrete time

[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]

Continuous Time

- « Kurtz's theorem » also holds in continuous time (finite state space)
- Graham and Méléard: A generic result for **general** state space (in particular non enumerable).

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

- Object i has a free evolution plus pairwise interactions.
- $X_n^N(0)_{n=1\ldots N}$ 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.

Age of Gossip

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

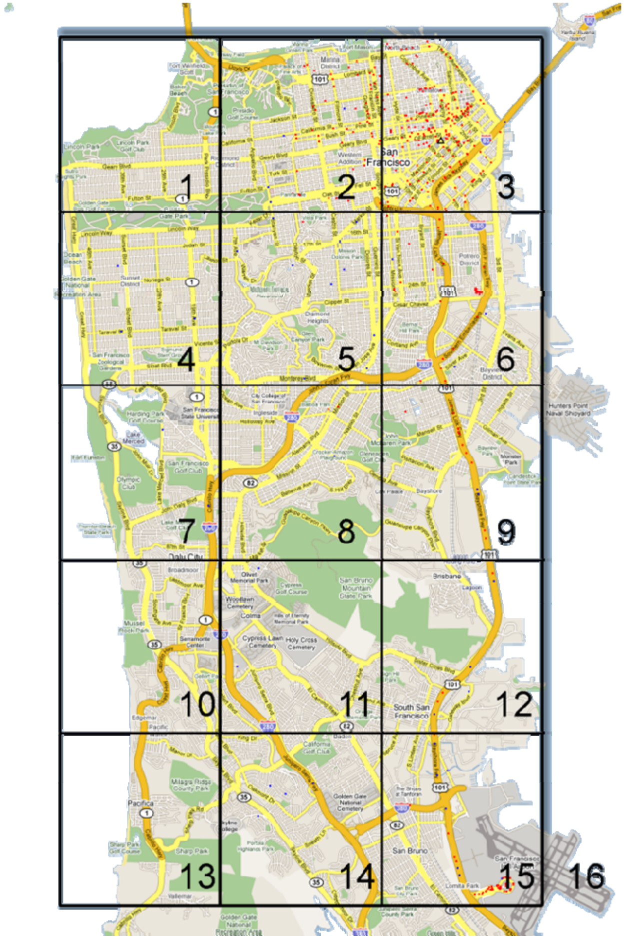

- [Graham and Méléard 1997] applies, i.e. mean field convergence occurs for iid initial conditions
- [Chaintreau et al.(2009)Chaintreau, Le Boudec, and Ristanovic] shows more, i.e. weak convergence of initial condition suffices

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

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]

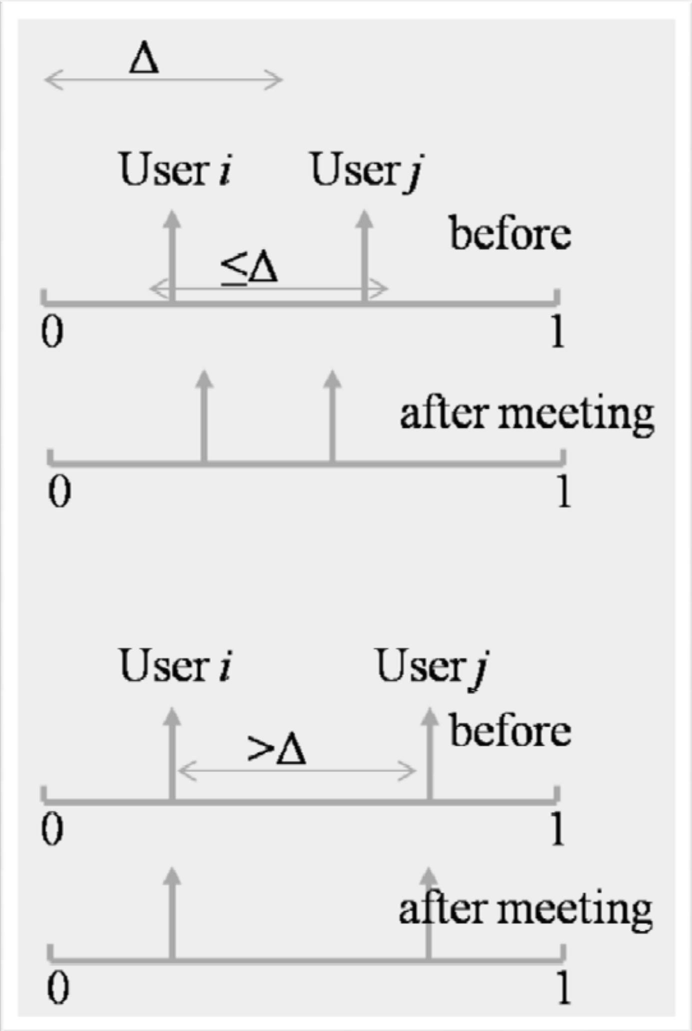

Discrete time. State space =[0, 1].

$X_n^N(k) \in [0,1]$ 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

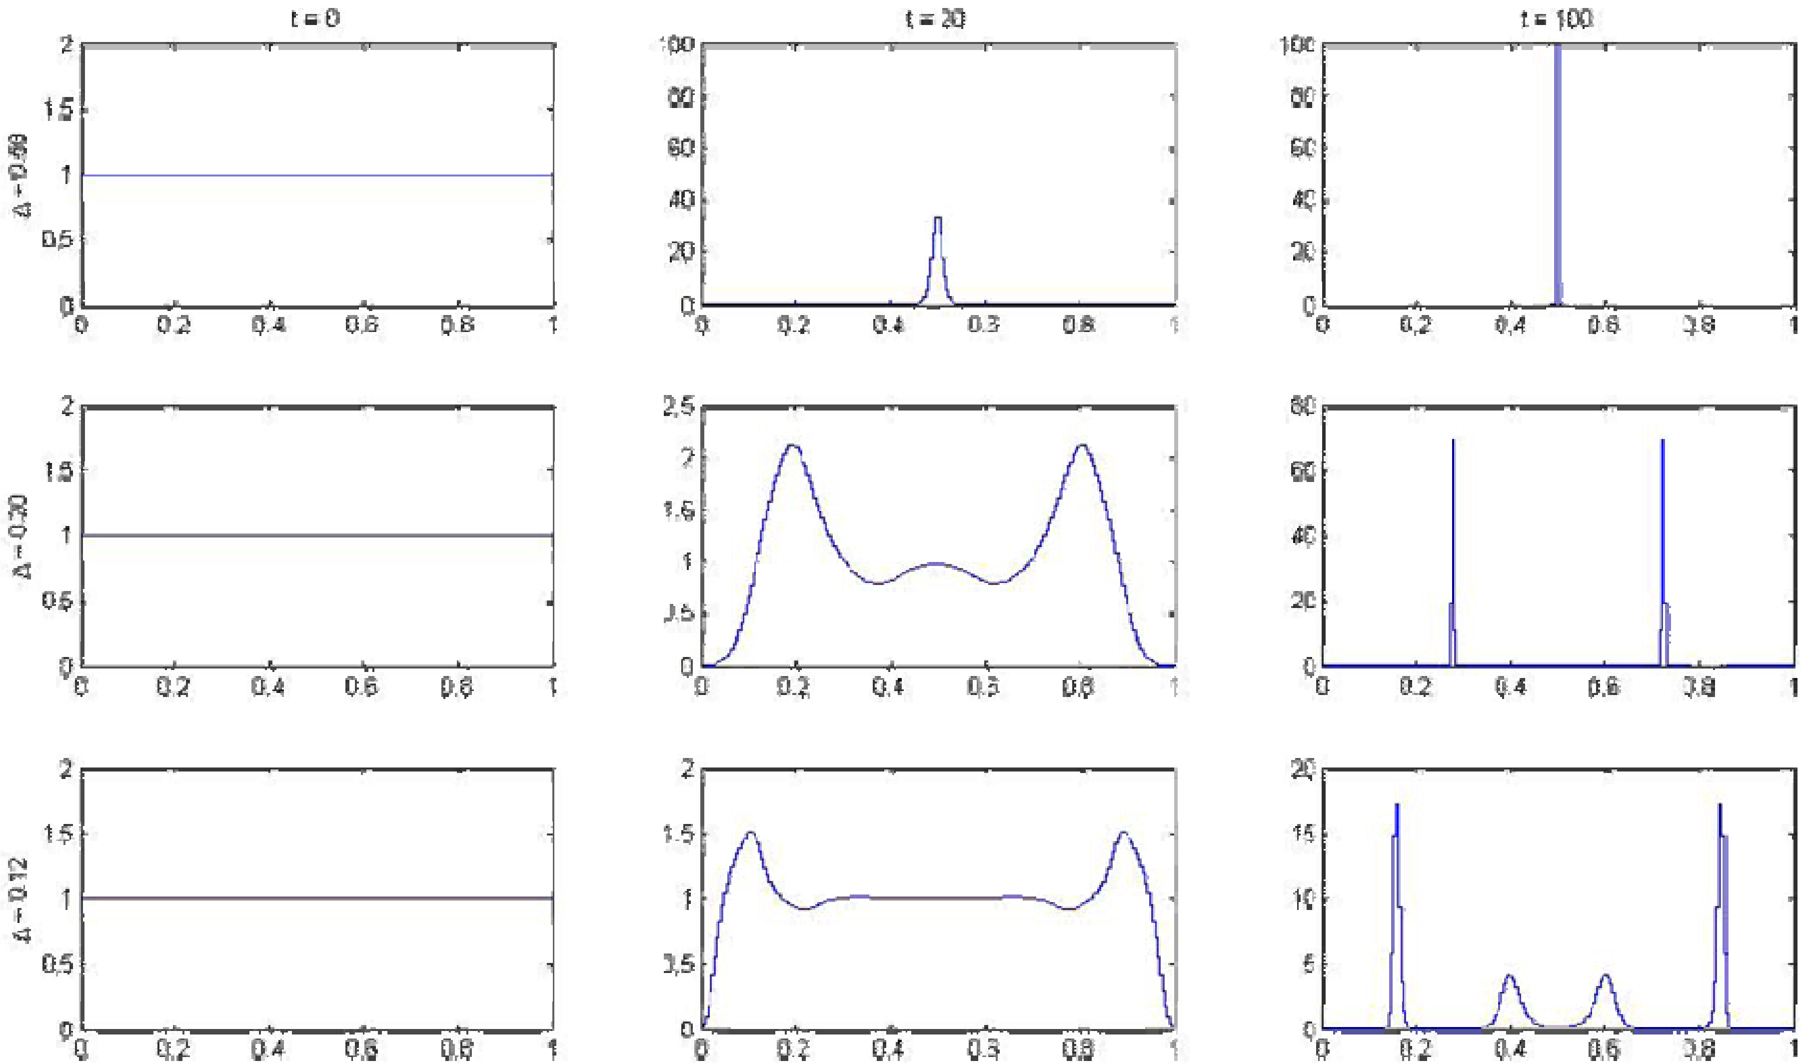

Is There Convergence to Mean Field ?

- Yes for the discretized version of the problem

  - Replace ratings in [0,1] by fixed point real numbers on d decimal places
  - Generic result says that mean field convergence holds (use [Benaim Le Boudec 2008], the number of meetings is upper bounded by a constant, here 2).
  - There is convergence for any initial condition such that $M^N(0) \to m_0$

- This is what any simulation implements

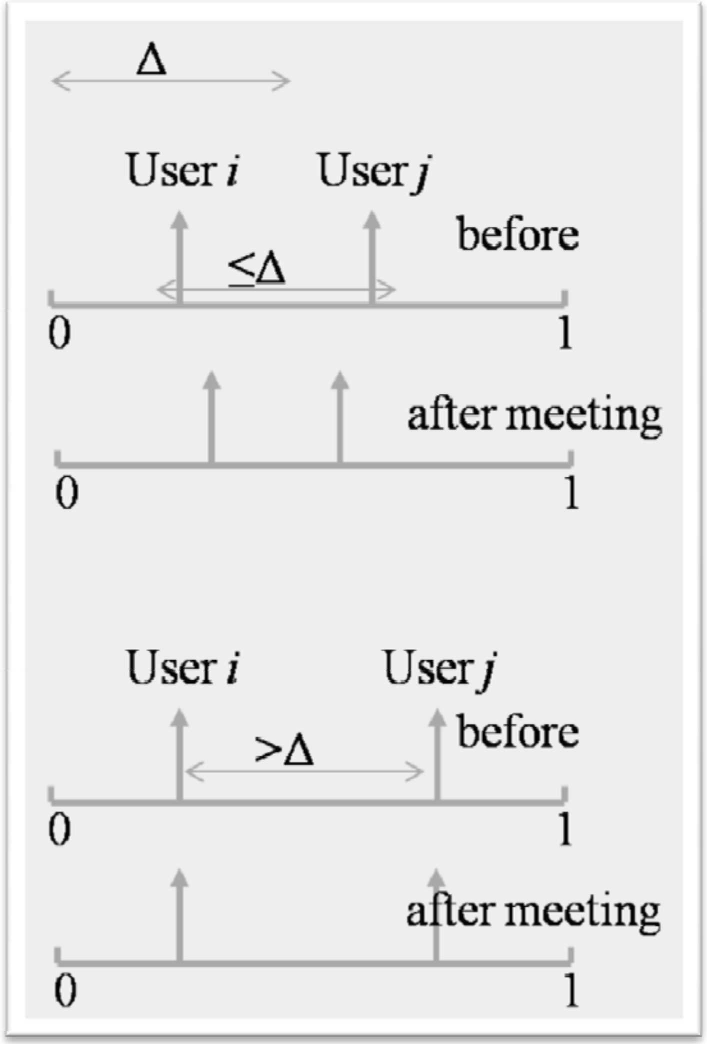

Is There Convergence to Mean Field ?

- There can be no similar result for the real version of the problem
  - Counter Example: $M^N(0) \to m(0)$ (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]

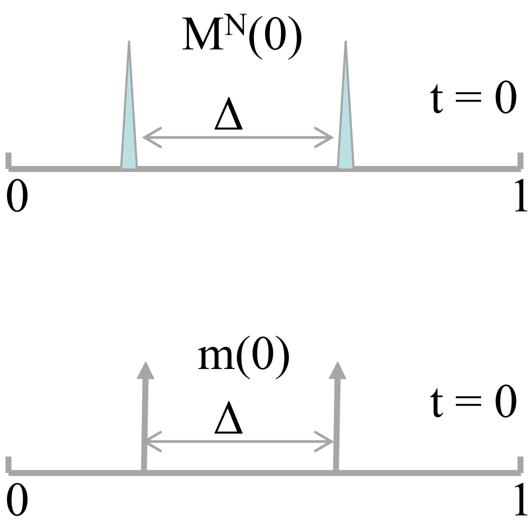

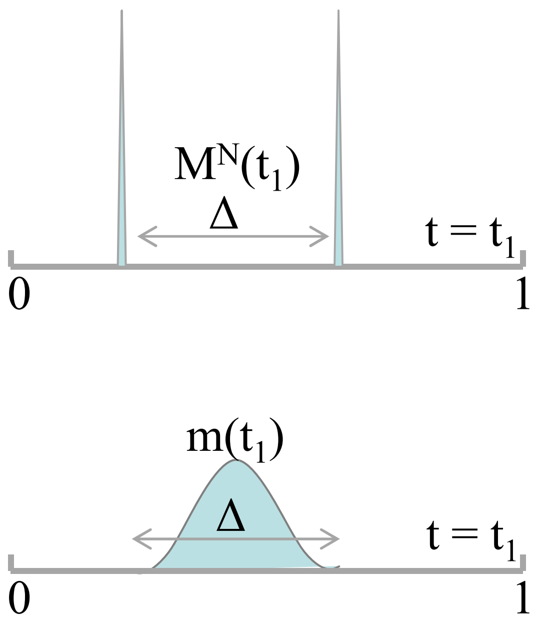

Convergence to Mean Field

- For the finite state space case, there are many simple results, often verifiable by inspection

  For example [Kurtz 1970] or [Benaim, Le Boudec 2008]

- For the general state space, things may be more complex

FINITE HORIZON

RANDOM PROCESS MODULATED BY MEAN FIELD LIMIT

Fast Simulation = Random Process Modulated by Mean Field Limit

Assume we know the state of *one tagged object* at time 0; we can approximate its evolution by replacing all other objects collectively by the mean field limit (e.g. the ODE)

The state of this object is a jump process, with transition matrix driven by the ODE [Darling and Norris, 2008]

A stronger result than propagation of chaos – does not require exchangeability

2-Step Malware Example

- $p^N_j(t|i)$ is the probability that a node that starts in state i is in state j at time t:

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

- Then $p_j^N(t/N|i) \approx p_j(t|i)$ where *p(t/i)* is a continuous time, non homogeneous process

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

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

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

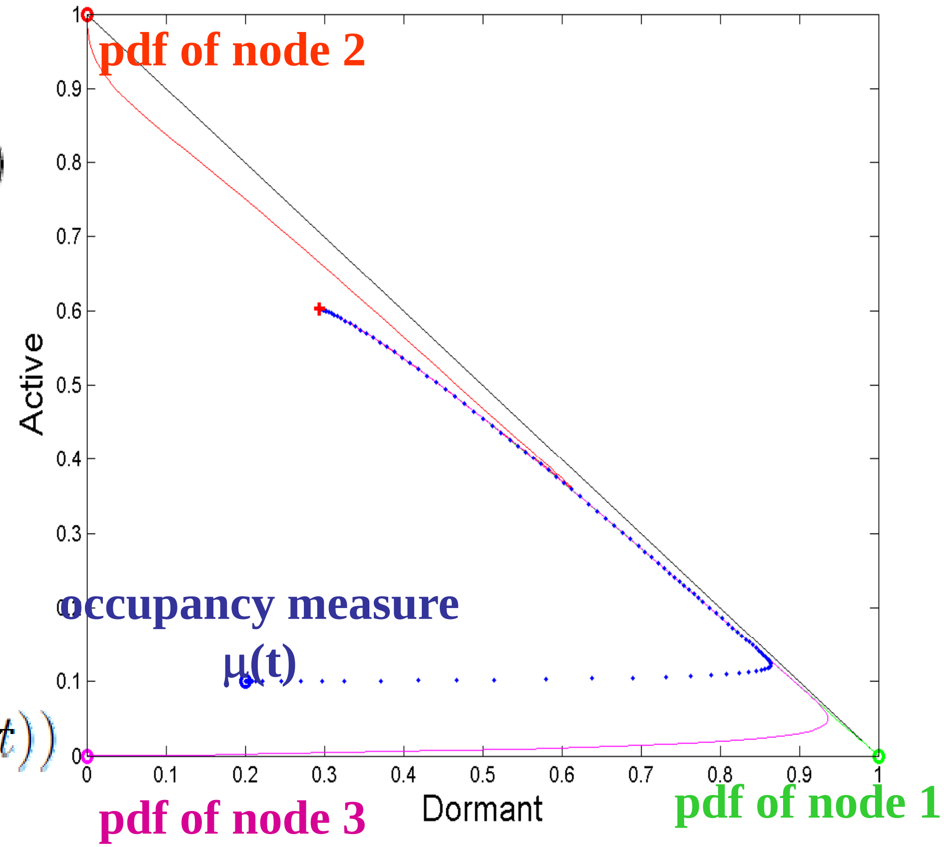

Details of the 2-Step Malware Example

- $P^N_{i,j}$ (m) is the marginal transition probability for one object, given that the state of the system is m

$$P^N(\vec{m}) = I + \frac{1}{N}\begin{pmatrix} -\frac{A}{h+D}\beta - 2\lambda\frac{ND-1}{N-1} - \delta_D & \frac{A}{h+D}\beta + 2\lambda\frac{ND-1}{N-1} & \delta_D \\ 0 & -\delta_A & \delta_A \\ \alpha_0 + Dr & \alpha & -\alpha_0 - Dr - \alpha \end{pmatrix}$$

$$= I + \frac{1}{N}A^N(\vec{m})$$

$$\vec{m} = (D, A)$$

- Note: Knowing the transition matrix P^N (m) is not enough to be able to simulate (or analyze) the system with *N* objects
  - Because there may be simultaneous transitions of several objects (on the example, up to 2)
- However, the fast simulation says that, in the large *N* limit, we can consider one (or *k*) objects as if they were independent of the other *N*-*k*
  - $(X^N_1(t/N), M^N(t/N))$ can be approximated by the process $(X_1(t), m(t))$ where m(t) follows the ODE and $X_1(t)$ is a jump process with time-dependent transition matrix A(m(t)) where $A^N(\vec{m}) \to A(\vec{m})$

$$
\begin{aligned}
P^N(\vec{m}) &= I + \frac{1}{N} \overset{A^N}{\boxed{\begin{pmatrix} -\frac{A}{h+D}\beta - 2\lambda\frac{ND-1}{N-1} - \delta_D & \frac{A}{h+D}\beta + 2\lambda\frac{ND-1}{N-1} & \delta_D \\ 0 & -\delta_A & \delta_A \\ \alpha_0 + Dr & \alpha & -\alpha_0 - Dr - \alpha \end{pmatrix}}} \\
&= I + \frac{1}{N} A^N(\vec{m})
\end{aligned}
$$

- The state of one object is a jump process with transition matrix:

$$
A(\vec{m}) = \begin{pmatrix} -\frac{A}{h+D}\beta - 2\lambda D - \delta_D & \frac{A}{h+D}\beta + 2\lambda D & \delta_D \\ 0 & -\delta_A & \delta_A \\ \alpha_0 + Dr & \alpha & -\alpha_0 - Dr - \alpha \end{pmatrix}
$$

where m = (D, A, S) depends on time (is solution of the ODE)

Computing the Transition Probability

- $P^N_{i,j}$ (m) is the transition probability for one object, given that the state if m

$$
\begin{aligned}
P^N(\vec{m}) &= I + \frac{1}{N}\begin{pmatrix} -\frac{A}{h+D}\beta - 2\lambda\frac{ND-1}{N} - \delta_D & \frac{A}{h+D}\beta + 2\lambda\frac{ND-1}{N} & \delta_D \\ 0 & -\delta_A & \delta_A \\ \alpha_0 + Dr & \alpha & -\alpha_0 - Dr - \alpha \end{pmatrix} \\
&= I + \frac{1}{N}A^N(\vec{m})
\end{aligned}
$$

where I is the identity matrix and $\vec{m} = (D, A, S)$.

$P^N_{1,3}$ is the probability that one node in state $i = 1$, i.e. 'D' moves to state $j = 3$, i.e. 'S'. This corresponds to case 1 in the table. The probability that this case occurs in one time slot is $D\delta_D$ and the probability that the transition affects precisely the node of interest is $\frac{D\delta_D}{ND}$ since there are ND nodes in the 'D' state. Thus $P^N_{1,3} = \frac{1}{N}\delta_D$.

$P^N_{1,2}$ is the probability that one node in state $i = 1$, i.e. 'D' moves to state $j = 2$, i.e. 'S'. This corresponds to cases 2 and 3. The probability is the sum of the probabilities for each of these two cases, as they are mutually exclusive. The probability that case 2 occurs is $D\lambda\frac{ND-1}{N-1}$ (given by the table). The probability that this node is affected, given that case 2 occurs is $\frac{2}{ND}$ since case 2 affects 2 nodes that are in state 'D'. Thus the probability that this node does a transition of case 2 is $\frac{2}{N}\lambda\frac{ND-1}{N-1}$. Similarly, the probability that this node does a transition of case 3 is $\frac{AN}{h+D}\beta$. Thus $P^N_{1,2} = \frac{1}{N}\left(\frac{A}{h+D}\beta + 2\lambda\frac{ND-1}{N-1}\right)$.

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

STATIONARY REGIME

STATIONARY REGIME OF MEAN FIELD LIMIT

Stationary Regimes

- Original process is random, assume it has a unique stationary regime

- The mean field limit is deterministic;

Q: What is the stationary regime for a deterministic process ?

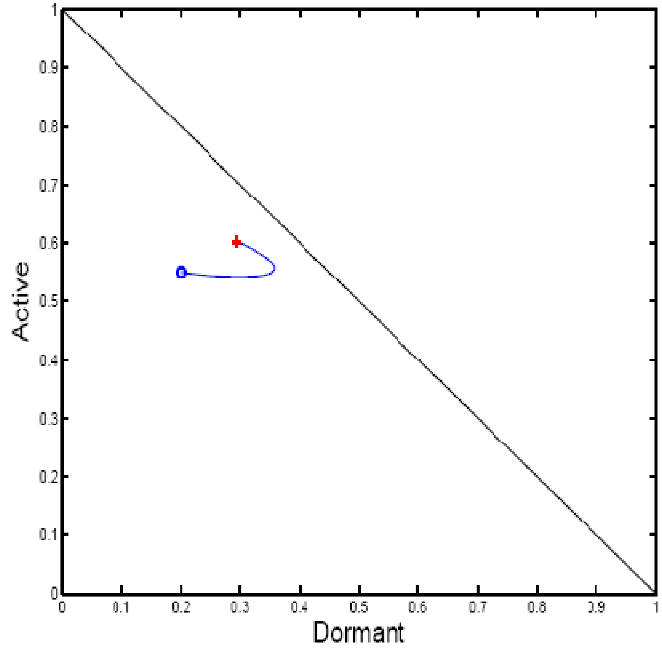

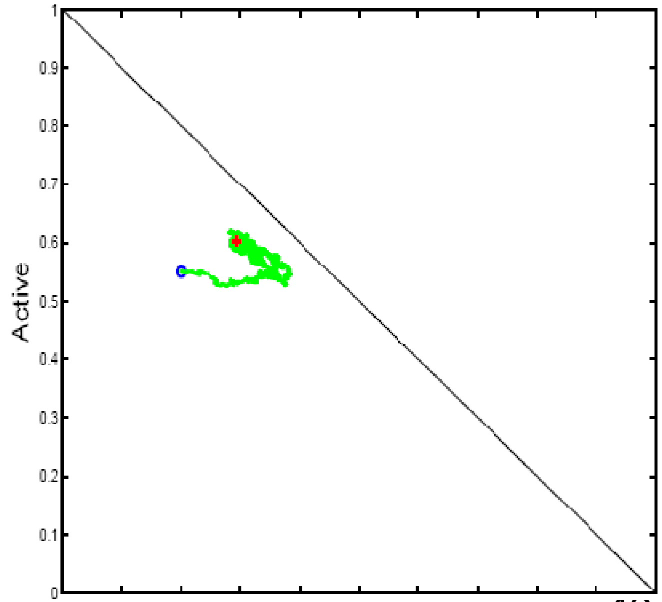

Frequent Answer

- Mean field limit :

$$\frac{d\vec{m}}{dt} = F(\vec{m})$$

- Stationary regime

$$F(\vec{m}) = \vec{0}$$

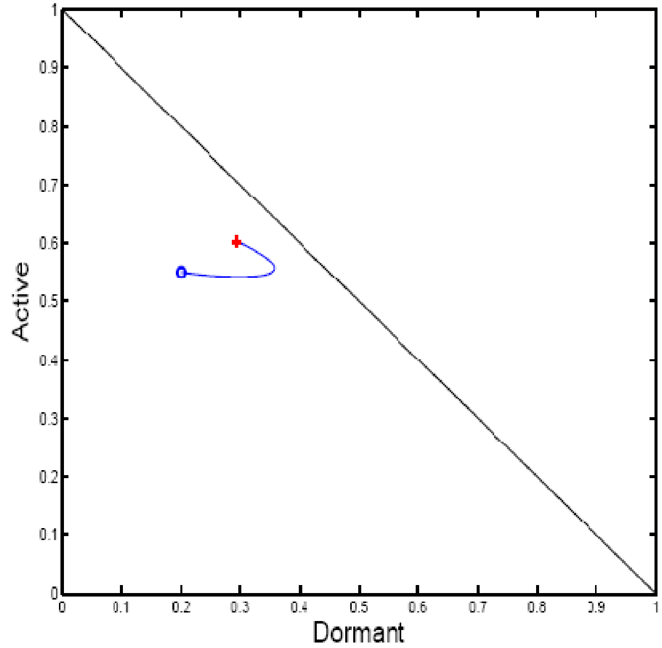

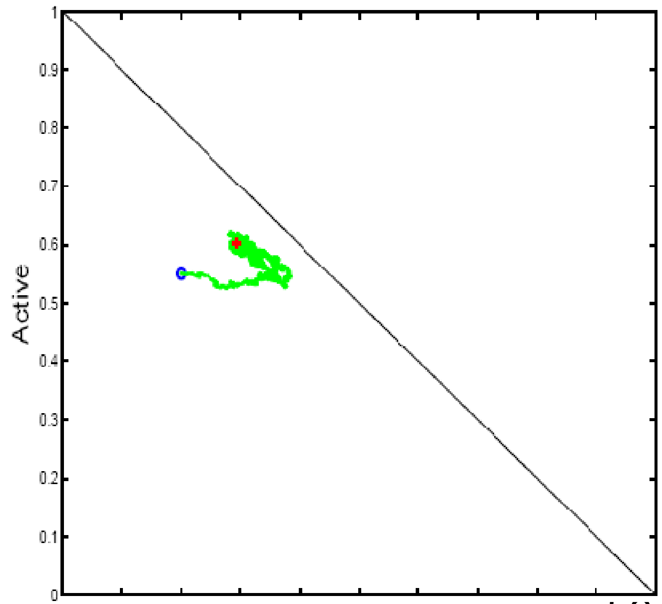

Example

- Same as before except for one parameter value : h = 0.1 instead of 0.3
- The ODE does not converge to a unique attractor (limit cycle)
- The equation F(m) = 0 has a unique solution (red cross)

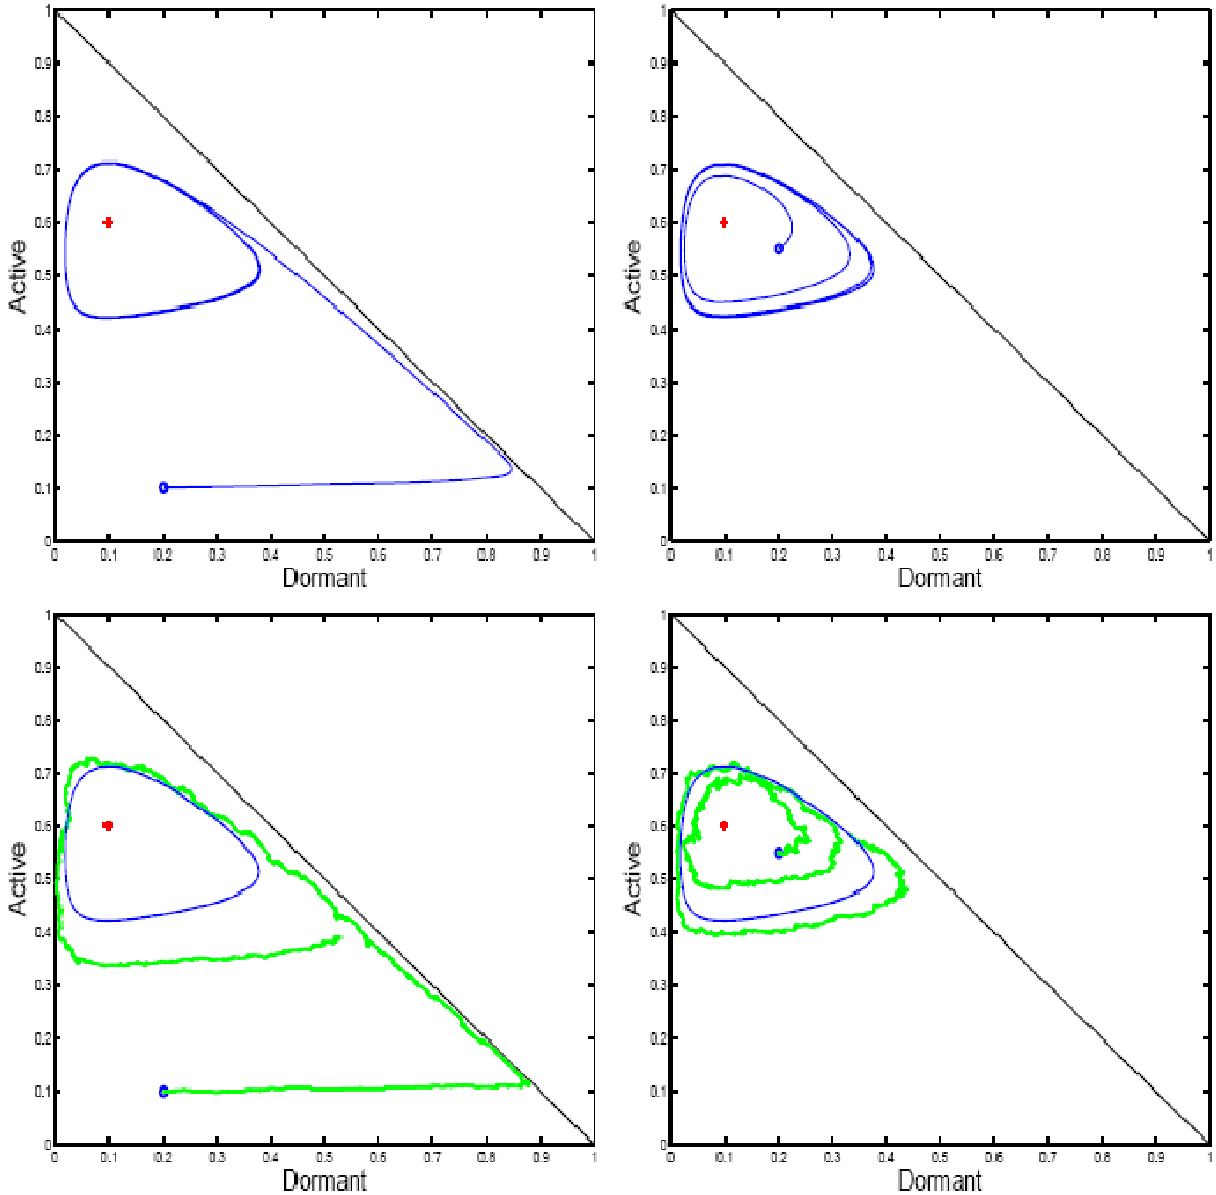

STATIONARY REGIME

CRITIQUE OF FIXED POINT METHOD

The Fixed Point Method

- A generic method, sometimes implicitly used
- Method is as follows:
  - Assume many interacting objects, focus on one object
  - Pretend this and other objects have a state distributed according to some proba m
  - Pretend they are independent
  - Write the resulting equation for m (a fixed point equation) and solve it, assumption

- Can be interpreted as follows
  - Assume a mean field interaction model, converges to mean field
  - Propagation of chaos => objects are asymptotically independent

Example: 802.11 Analysis, Bianchi's Formula

802.11 single cell

- m_i = proba one node is in backoff stage I
- β= attempt rate
- γ = collision proba

$$\frac{dm_0}{d\tau} = -m_0 q_0 + \beta(\vec{m})\left(1-\gamma(\vec{m})\right) + q_K m_K \gamma(\vec{m})$$

$$\frac{dm_i}{d\tau} = -m_i q_i + m_{i-1} q_{i-1} \gamma(\vec{m}) \qquad i = 1, ..., K$$

$$\beta(\vec{m}) = \sum_{i=0}^{K} q_i m_i$$

$$\gamma(\vec{m}) = 1 - e^{-\beta(\vec{m})}$$

See [Benaim and Le Boudec , 2008] for this analysis

Solve for Fixed Point:

$$m_i = \frac{\gamma^i}{q_i} \frac{1}{\sum_{k=0}^{K} \frac{\gamma^k}{q_k}}$$

Bianchi's Fixed Point Equation [Bianchi 1998]

$$\boxed{\begin{aligned} \gamma &= 1 - e^{-\beta} \\ \beta &= \frac{\sum_{k=0}^{K} \gamma^k}{\sum_{k=0}^{K} \frac{\gamma^k}{q_k}} \end{aligned}}$$

Example: Kelly's Alternate Routing [Kelly, 1991]

- $N = K(K-1)/2$ links, each of capacity C calls
- Arrival of calls to link ab with rate λ
- If link is saturated ($X_{ab}(t) = C$), arriving call attempts one two-hop alternate route (ac, cb); if either link on chosen alternate route is saturated, call is lost
- Call duration is expo(1)
- $X_{ab}(t)$ = number of calls using link ab; $Y_{ab}^{c}(t)$ = number of calls diverted via c
- System state = $(X_{ab}(t), Y_{ab}^{c}(t))_{a,b,c}$

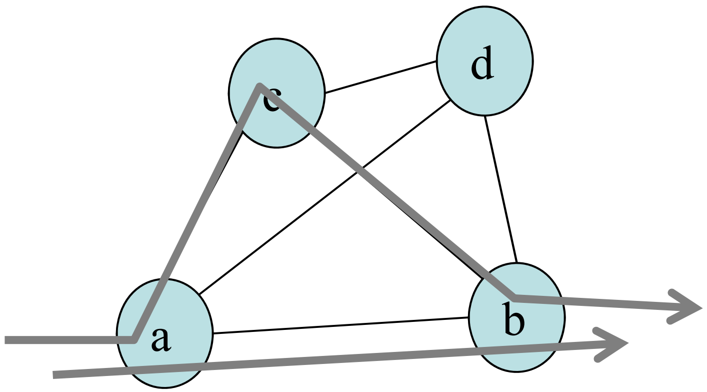

- This is *not* a mean field interaction model
  - If we rename object ab we need to rename obejct abc accordingly
- However,there is convergence to a deterministic occupancy measure and propagation of chaos [e.g. Graham and Méléard 1997]

Kelly's Alternate Routing Simplified Model

- $N = K(K-1)/2$ links, each of capacity C calls
- Arrival of calls to link n with rate λ
- If link is saturated ($X_n(t) = C$), arriving call attempts one alternative pair (n_1, n_2) of links; if either link on chosen alternate route is saturated, call is lost.
- If call is accepted on two hop route, both legs of the call become independent
- Call duration is expo(1)

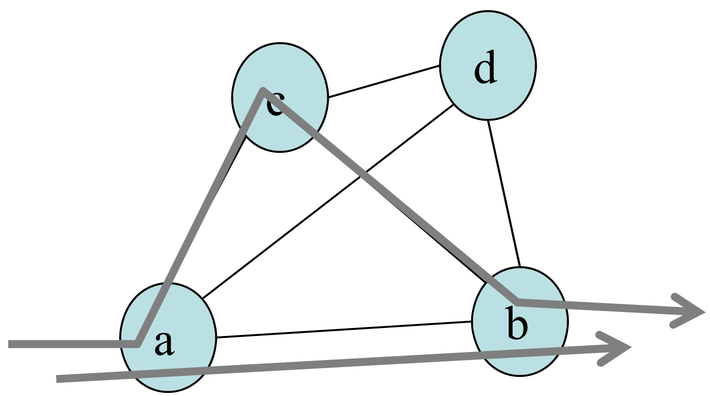

- This *is* a mean field interaction model, has same limiting equations as original limit.

Mean field equations:

$$\begin{aligned} X_n^N(t) \in \{0,1,2...,C\} &= \text{state of link } n \\ \sum_{k=0}^{n} \dot{m}_k(t) &= (n+1)m_{n+1}(t) - \gamma(t)m_n(t), \;\; n = 0,1,....C-1 \\ \gamma(t) &= \lambda\left\{1 + 2m_C(t)\left[(1 - m_C(t)\right]\right\} \end{aligned}$$

Fixed point: solve for m_n and γ

$$\begin{aligned} (n+1)m_{n+1} &= \gamma m_n \\ \gamma &= \lambda\left\{1 + 2m_C(t)\left((1 - m_C\right)\right\} \end{aligned}$$

Which gives

$$m_n = \frac{\gamma^n}{n!} \Big/ \left(\sum_{k=0}^{C} \frac{\gamma^k}{k!}\right)$$

the stationary points are obtained by solving for m_C and γ in

$$\begin{aligned} m_C &= E(\gamma, C) \\ \gamma &= \lambda\left[1 + 2m_C(1 - m_C)\right] \end{aligned}$$

with

$$E(\gamma, C) \stackrel{\text{def}}{=} \frac{\gamma^C}{C!} \Big/ \left(\sum_{k=0}^{C} \frac{\gamma^k}{k!} \right)$$

which is equivalent to

$$\boxed{m_C = E(\lambda\left[1 + 2m_C(1 - m_C)\right], C)}$$

Fixed Point Equation for saturation prob m_C

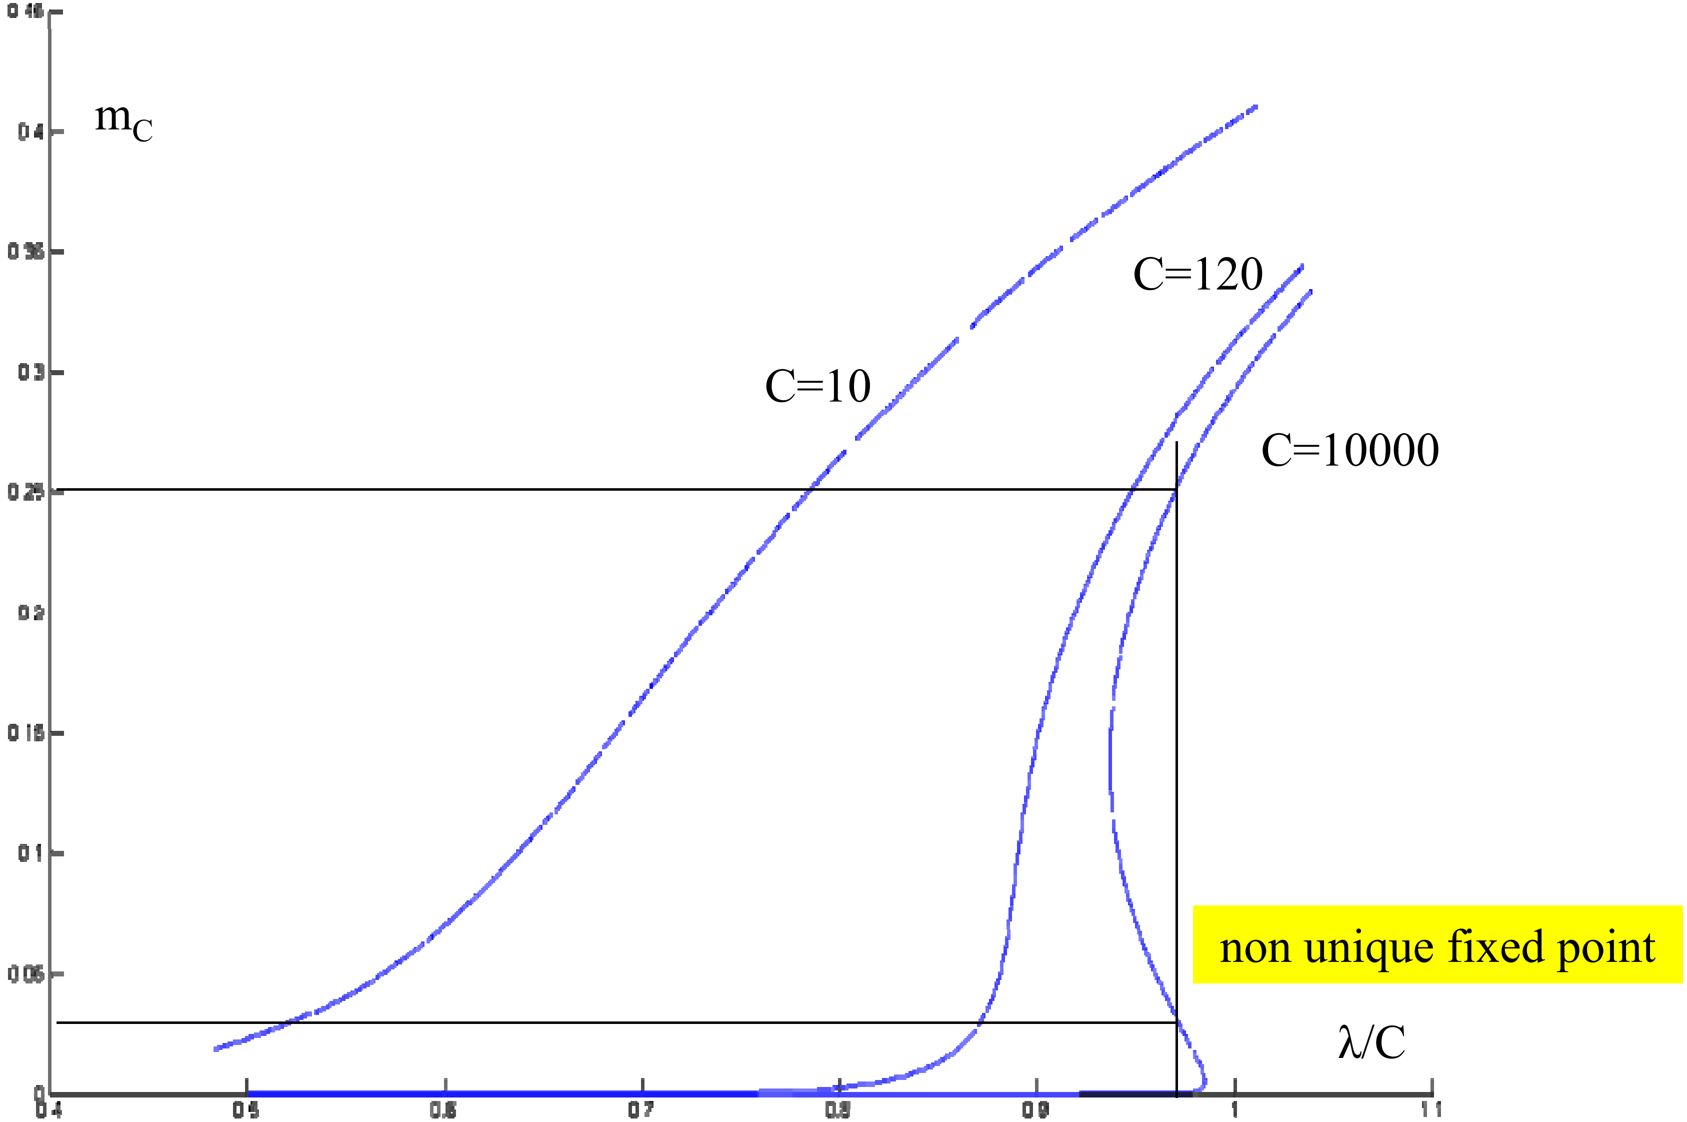

m_C
C=10
C=120
C=10000
non unique fixed point
λ/C

Fixed Point Method Applied to 2-Step Malware Example

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	$S\alpha$

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
6. Direct infection
   - S -> A

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

- Solve for (D,A,S)
- Has a unique solution

Example Where Fixed Point Method Succeeds

- In stationary regime:
  - Prob (node *n* is dormant) ≈ 0.3
  - Prob (node *n* is active) ≈ 0.6
  - Prob (node *n* is susceptible) ≈ 0.1

  - Nodes *m* and *n* are independent

- The diagram commutes

$$\begin{tikzcd}[column sep=huge, row sep=huge]
\text{Law of } M^N(t) \arrow[r, "t \to +\infty"] \arrow[d, "N \to +\infty"'] & \varpi^N \arrow[d, "N \to +\infty"] \\
\delta_{\mu(t)} \arrow[r, "t \to +\infty"'] & \delta_{m^*}
\end{tikzcd}$$

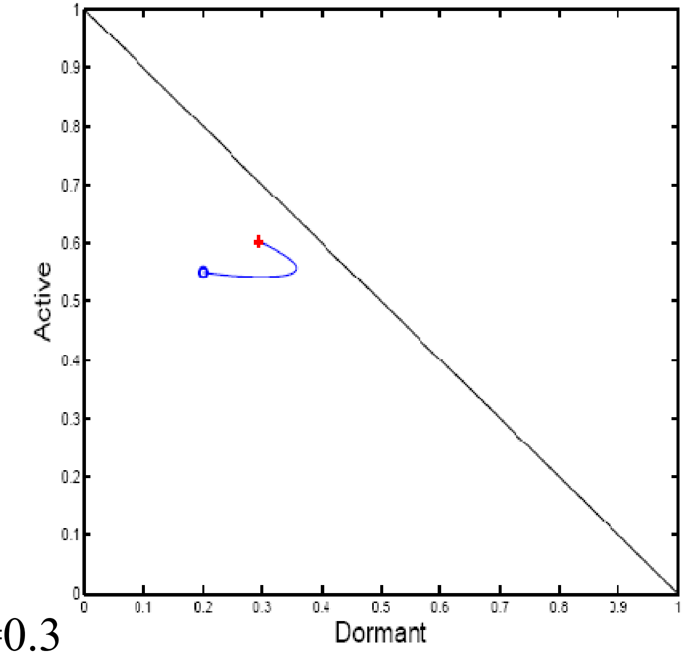

h=0.3

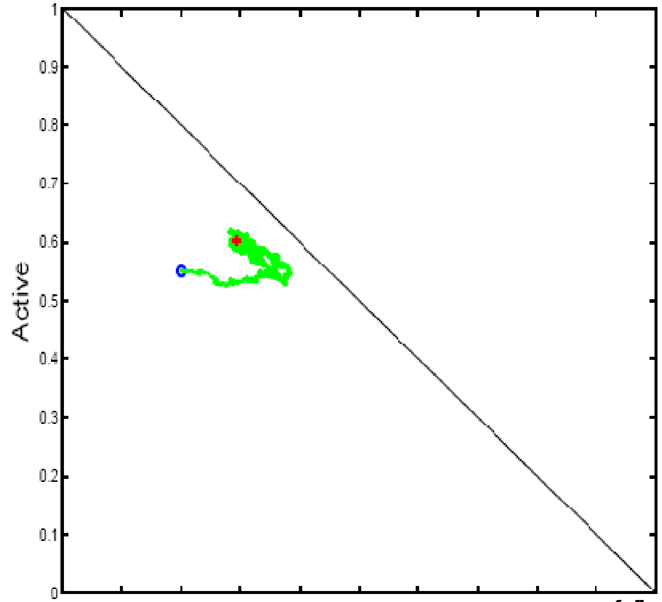

Example Where Fixed Point Method Fails

- In stationary regime, m(t) = (D(t), A(t), S(t)) follows the limit cycle
- Assume you are in stationary regime (simulation has run for a long time) and you observe that one node, say n=1, is in state 'A'
- It is more likely that m(t) is in region R
- Therefore, it is more likely that some other node, say n=2, is also in state 'A'
- This is synchronization

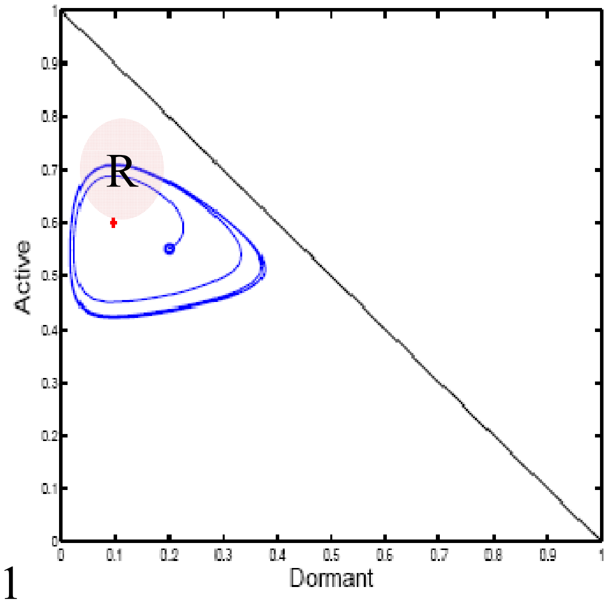

h=0.1

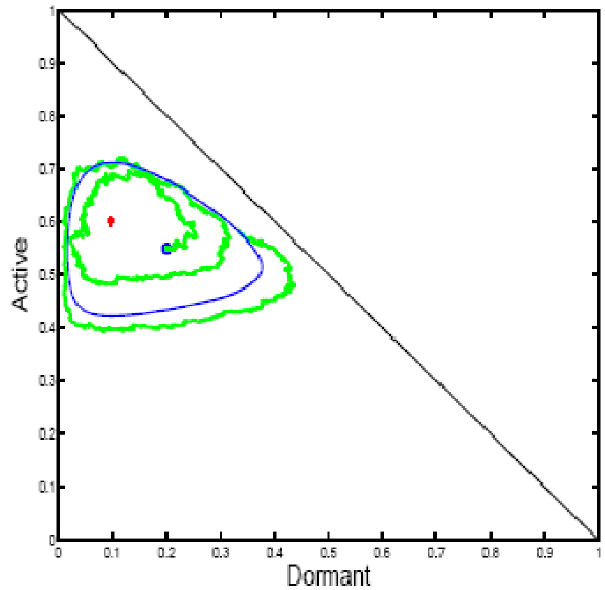

Joint PDFs of Two Nodes in Stationary Regime

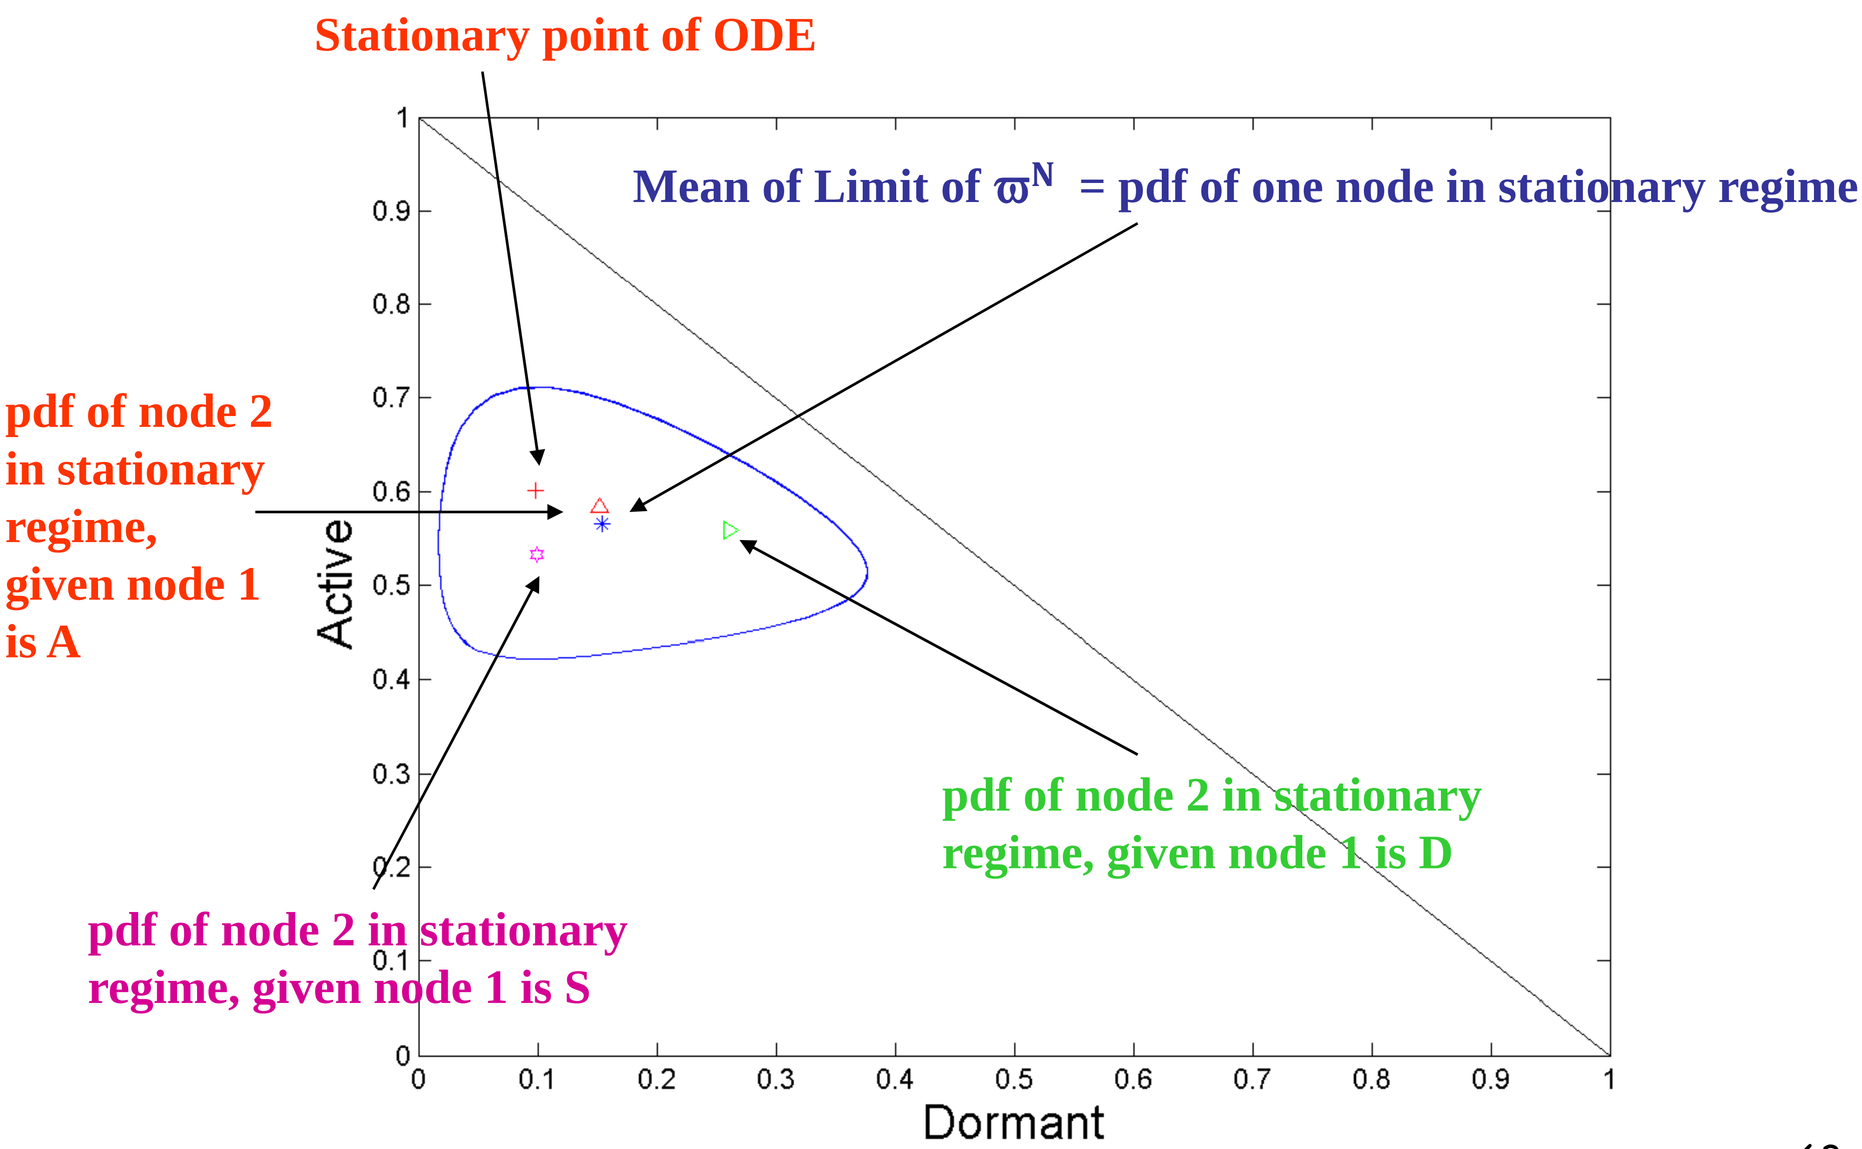

Numerical Results ($h = 0.1$).

prob of state	D	A	S
given D	0.261	0.559	0.181
given A	0.152	0.583	0.264
given S	0.099	0.533	0.368
unconditional	0.154	0.565	0.281

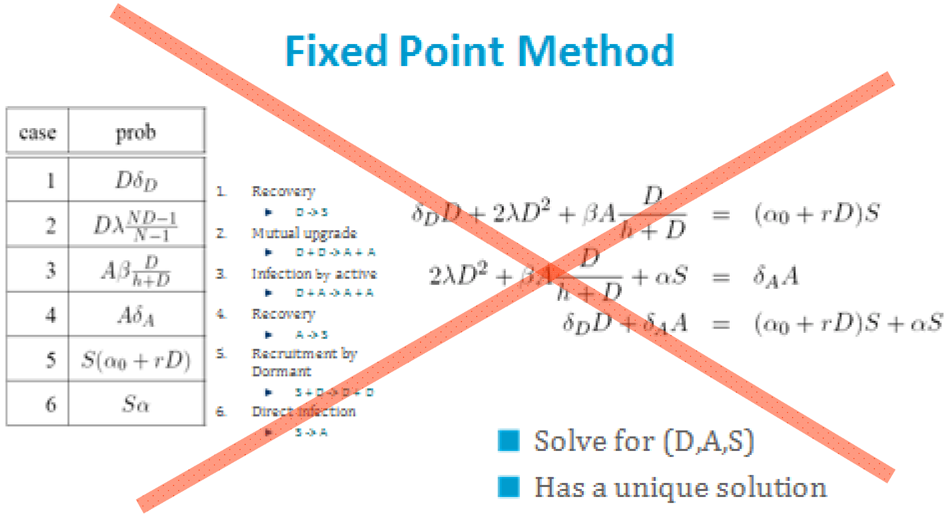

Where is the Catch ?

- Mean field convergence implies that nodes m and n are asymptotically independent
- There *is* mean field convergence for this example
- But we saw that nodes may not be asymptotically independent

… is there a contradiction ?

Markov chain is ergodic

$$\begin{tikzcd} \mathbb{P}(X_1^N(t/N) = i \text{ and } X_1^N(t/N) = j) \arrow[r, "t\to\infty"] \arrow[d, "N\to\infty"'] & \pi_{i,j}^N \arrow[d, "N\to\infty"] \\ \mu_i(t)\mu_j(t) & ??? \end{tikzcd}$$

Mean Field Convergence

- Mean Field convergence implies asymptotic Independence in Transient Regime, but says nothing about Stationary Regime
- We have three general results

Result 1: Fixed Point Method Holds under (H)

- Assume that

(H) ODE has a unique global stable point to which all trajectories converge

- Theorem [e.g. Benaim et al 2008] : The limit of stationary distribution of M^N is concentrated on this fixed point
  - i.e., under (H), the fixed point method and the decoupling assumptions are justified

- Uniqueness of fixed point is not sufficient
- (H) has nothing to do with the properties at finite N
  - In our example, for h=0.3 the decoupling assumption holds in stationary regime, for h=0.1 it does not
  - In both cases the Markov chain at finite N has the same graph.
- Study the ODE !

The Diagram Does Not Always Commute

- h=0.1

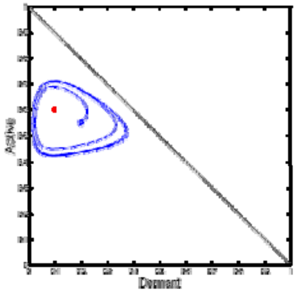

$$\begin{array}{ccc} \mathbb{P}(X_1^N(t/N) = i \text{ and } X_1^N(t/N) = j) & \xrightarrow{t\to\infty} & \pi_{i,j}^N \\ \Big\downarrow {\scriptstyle N\to\infty} & & \Big\downarrow {\scriptstyle N\to\infty} \\ \mu_i(t)\mu_j(t) & & \frac{1}{T}\int_0^T \mu_i(t)\mu_j(t)dt \end{array}$$

- For large *t* and *N*:

$$\begin{aligned} \mathbb{P}(X_1^N(t/N) = i \text{ and } X_1^N(t/N) = j) &\approx \frac{1}{T}\int_0^T \mu_i(t)\mu_j(t)dt \\ &\neq \left(\frac{1}{T}\int_0^T \mu_i(t)dt\right)\left(\frac{1}{T}\int_0^T \mu_j(t)dt\right) \end{aligned}$$

where *T* is the period of the limit cycle

Result 2 for Stationary Regime

- Original system (stochastic):
  - $(X^N(t))$ is Markov, finite, discrete time
  - Assume it is irreducible, thus has a unique stationary proba $\nu^{\mathcal{N}}$
  - Let ϖ^N be the corresponding stationary distribution for $M^N(t)$, i.e.

    $P\left(M^N(t)=(x_1,\dots,x_I)\right) = \varpi^N(x_1,\dots,x_I)$ for x_i of the form k/n, k integer
- Theorem [Benaim]

Theorem 3 *The support of any limit point of ϖ^N is a compact set included in the Birkhoff center of Φ.*

Birkhoff Center: closure of set of points s.t. $m \in \omega(m)$
Omega limit: $\omega(m)$ = set of limit points of orbit starting at m

- Here:
  Birkhoff center = limit cycle ∪ fixed point

- The theorem says that the stochastic system for large N is close to the Birkhoff center,

  i.e. the stationary regime of ODE is a good approximation of the stationary regime of stochastic system

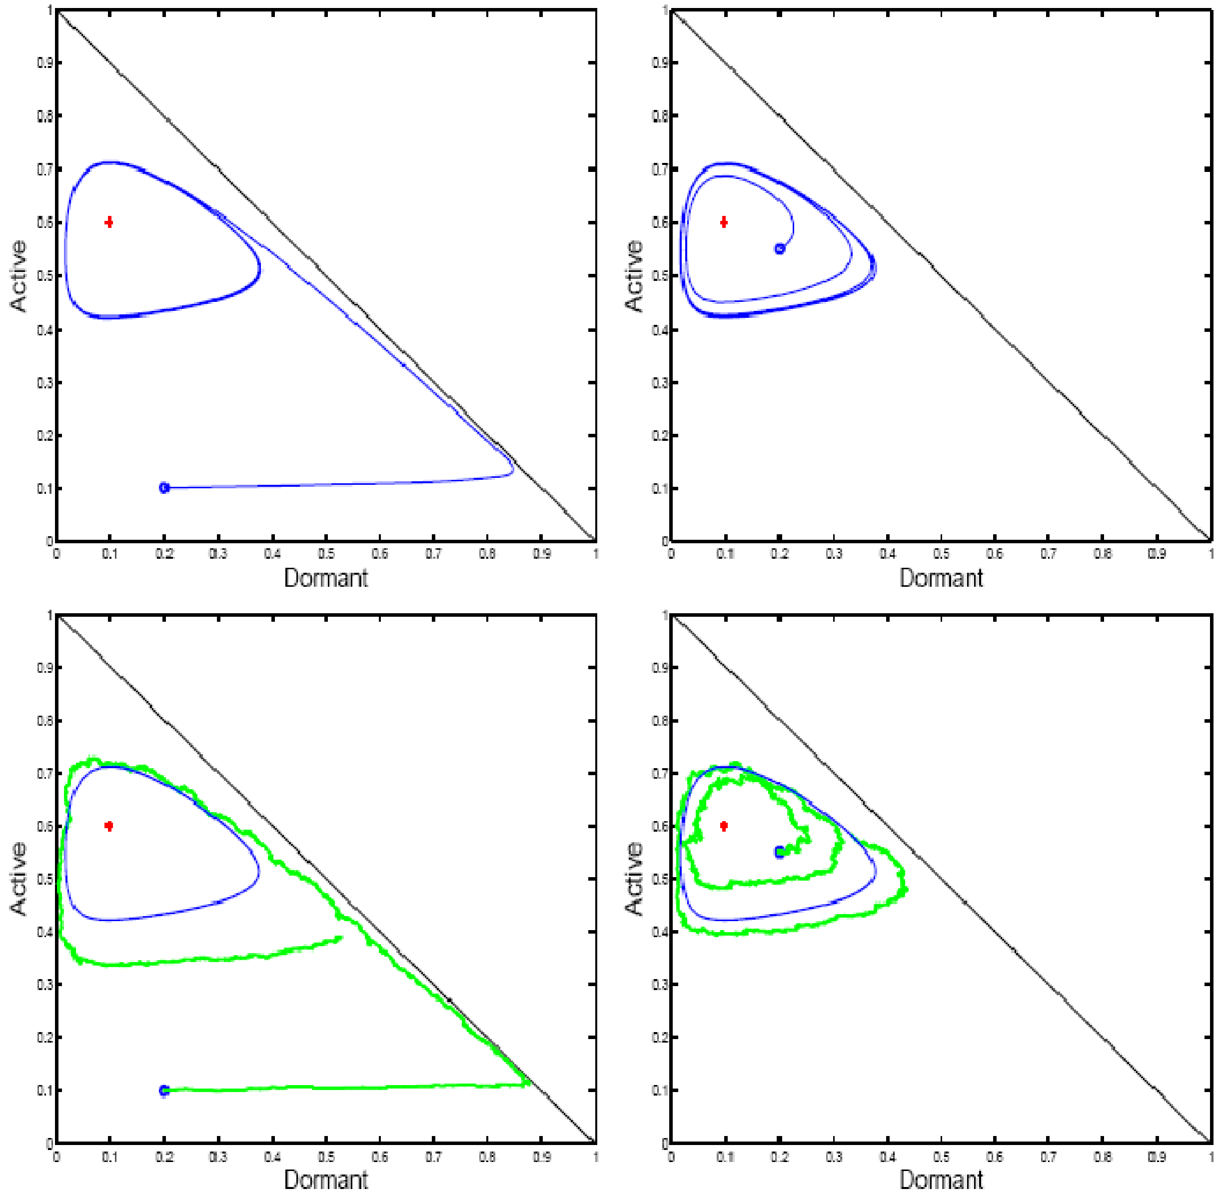

Existence and Unicity of a Fixed Point are not Sufficient for Validity of Fixed Point Method

- Essential assumption is

(H) $m(t)$ converges to a unique m^*

- It is not sufficient to find that there is a unique stationary point, i.e. a unique solution to $F(m^*)=0$
- Counter Example on figure
  - $(X^N(t))$ is irreducible and thus has a unique stationary probability η^N
  - There is a unique stationary point (= fixed point) (red cross)
    - $F(m^*)=0$ has a unique solution
    - but it is not a stable equilibrium
  - The fixed point method would say here
    - Prob (node n is dormant) ≈ 0.1
    - Nodes are independent
  - ... but in reality
    - We have seen that nodes are not independent, but are correlated and *synchronized*

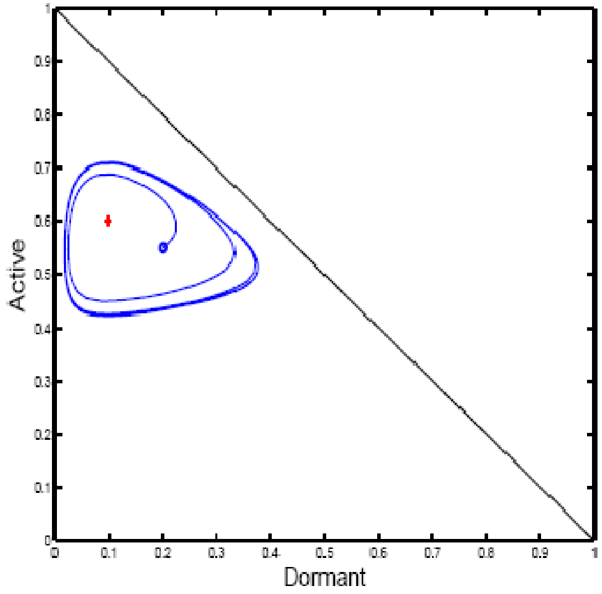

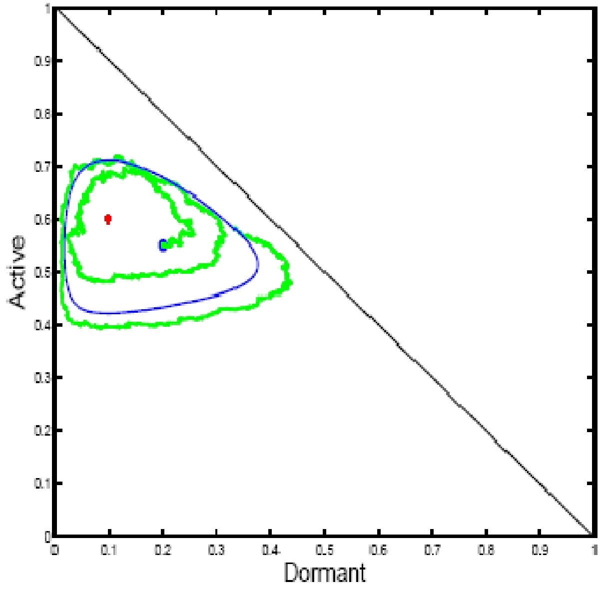

Example: 802.11 with Heterogeneous Nodes

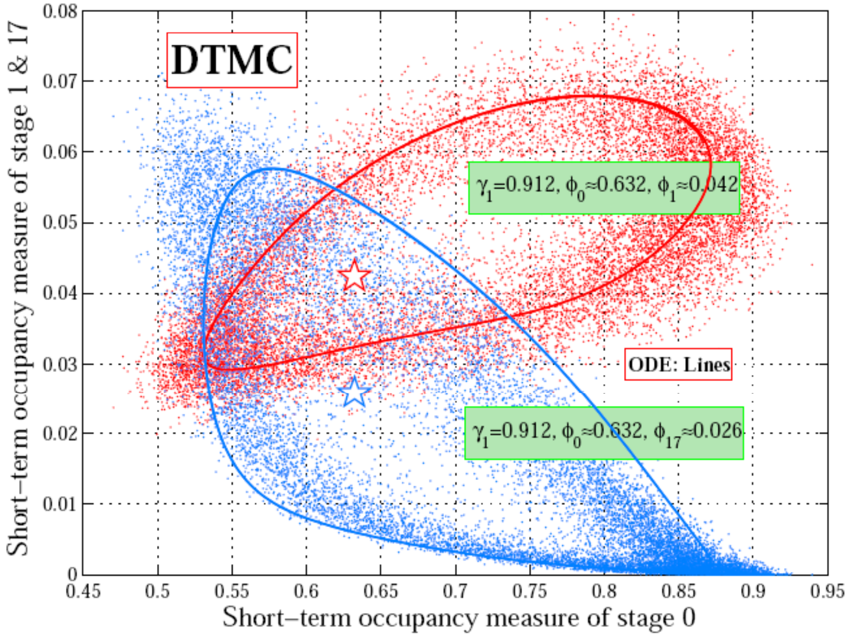

- [Cho2010]

  Two classes of nodes with heterogeneous parameters (restransmission probability)

  Fixed point equation has a unique solution

  There is a limit cycle

Quiz

- $M^N(t)$ is a Markov chain on $E=\{(a, b, c) \geq 0, a + b + c = 1, \ a, b, c \text{ multiples of } 1/N\}$

A. $M^N(t)$ is periodic, this is why there is a limit cycle for large N.

B. For large N, the stationary proba of M^N tends to be concentrated on the blue cycle.

C. For large N, the stationary proba of M^N tends to a Dirac.

D. $M^N(t)$ is not ergodic, this is why there is a limit cycle for large N.

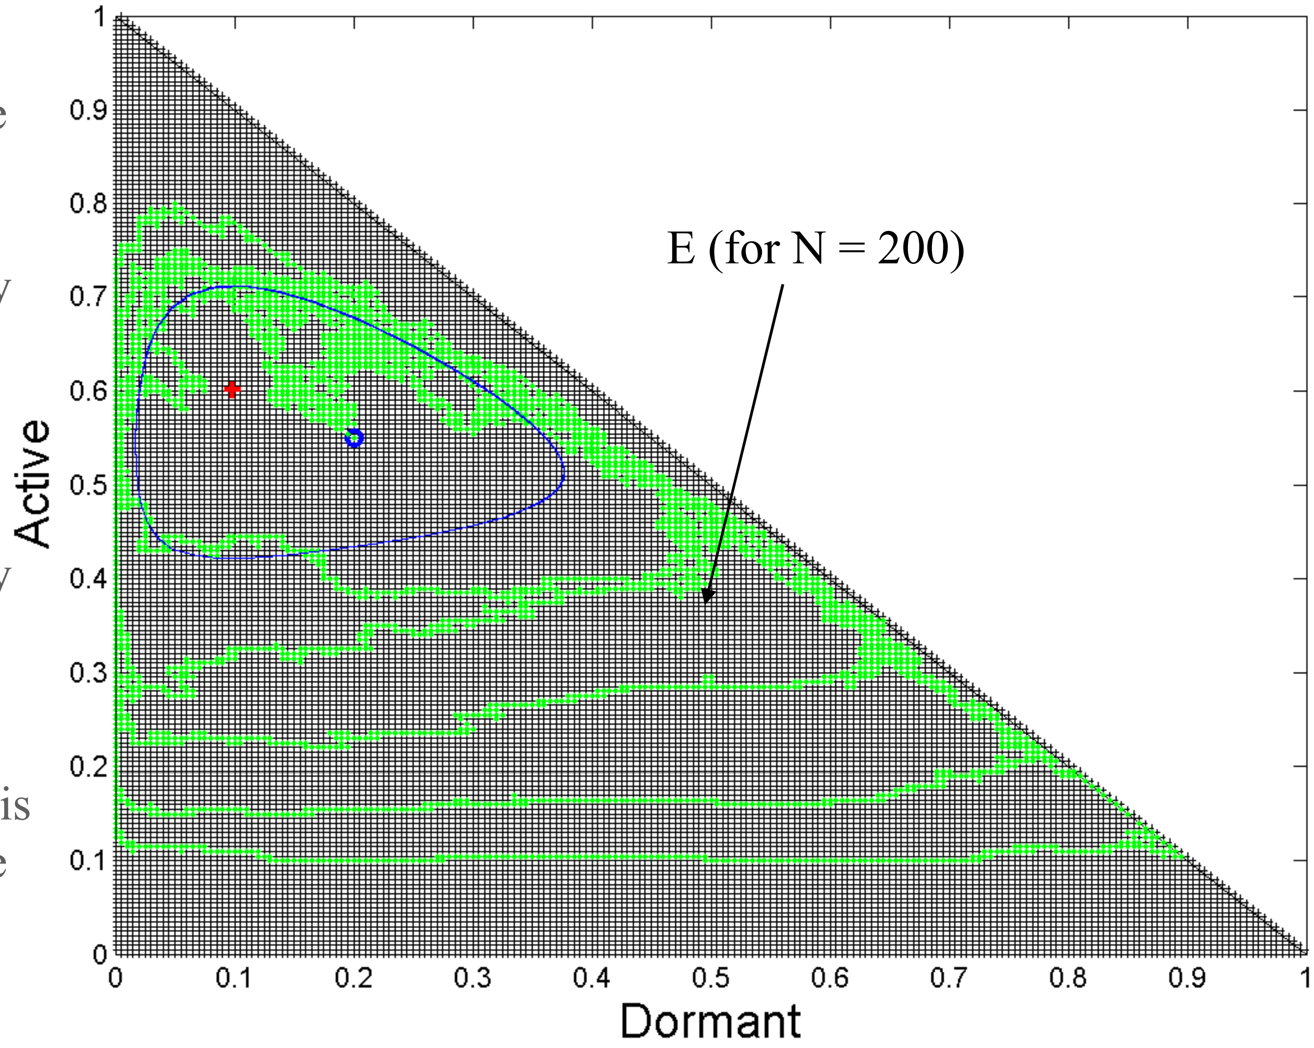

STATIONARY REGIME

REVERSIBLE CASE

Result 3: Reversible Case

- **Definition** Markov Process X(t) on enumerable state E space, with transition rates $q(i,j)$ is reversible iff
  1. It is ergodic
  2. There exists some probability distribution p such that, for all i, j in E

$$p(i)\, q(i,j) = p(j)\, q(j,i)$$

- If X(t) is reversible iff
  1. It is stationary (strict sense)
  2. It has same process law under reversal of time
- Most processes are not reversible, but some interesting cases exist:
  - Product form queuing networks with reversible routing matrix (e.g, on a bus)
  - Kelly's alternate routing models

Result 3: Reversible Case

Theorem 1.2 ([Le Boudec(2010)]) *Assume some process $Y^N(t)$ converges at any fixed t to some deterministic system $y(t)$ at any finite time. Assume the processes Y^N are reversible under some probabilities Π^N. Let $\Pi \in \mathcal{P}(E)$ be a limit point of the sequence Π^N. Π is concentrated on the set of stationary points S of the fluid limit $y(t)$*

- Stationary points = fixed points
- If process with finite *N* is reversible, the stationary behaviour is determined only by fixed points.
- Even if (H) does not hold

Example: Kelly's Alternate Routing

- System with N nodes is reversible
- Kelly's analysis looks for fixed points only
- Justified by reversibility

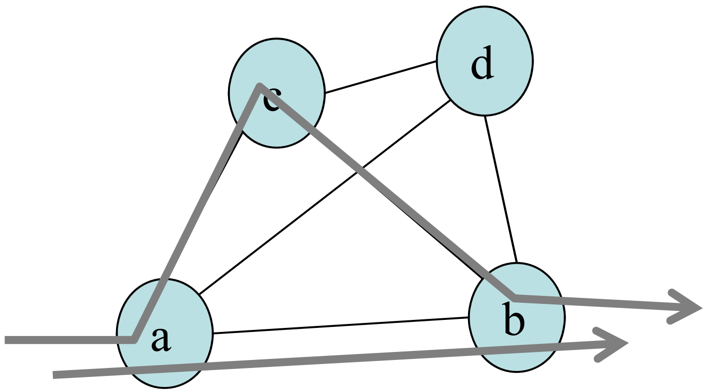

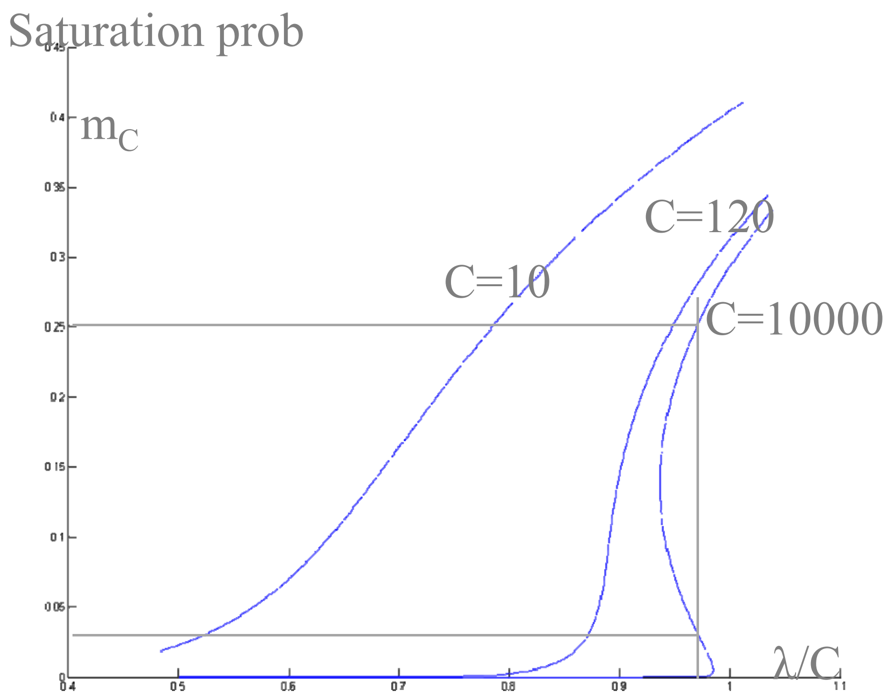

OPTIMIZATION

Decentralized Control

- Game Theoretic setting; *N* players, each player has a class, each class has a policy; each player also has a state;
  - Set of states and classes is fixed and finite
  - Time is discrete; a number of players plays at any point in time.
  - Assume similar scaling assumptions as before.
- [Tembine et al.(2009)]
  For large *N* the game converges to a single player game against a population;

Theorem 3.6.2 (Infinite N). *Optimal strategies (resp. equilibrium strategies) exist in the limiting regime when $N \to \infty$ under uniform convergence and continuity of $R^N \to R$. Moreover, if $\{U^N\}$ is a sequence of ε_N−optimal strategies (resp. ε_N−equilibrium strategies) in the finite regime with $\varepsilon_N \longrightarrow \varepsilon$, then, any limit of subsequence $U^{\phi(N)} \longrightarrow U$ is an ε− optimal strategies (resp. ε−equilibrium) for game with infinite N.*

Optimal, Centralized Control

- [Gast et al.(2010)]
- Markov decision process (MDP)
  - Finite state space per object, discrete time, N objects
  - Transition matrix depends on a control policy
  - For large N the system control converges to mean field, under any control
- Mean field limit
  - ODE driven by a control function
- **Theorem:** under similar assumptions as before, the optimal value function of MDP converges to the optimal value of the limiting system
- The result transforms MDP into fluid optimization, with very different complexity

Conclusion

- Mean field models are frequent in large scale systems
- Mean field is much more than a fluid approximation: decoupling assumption / fast simulation / random process modulated by fluid limit
- Decoupling assumption holds at finite horizon; may not hold in stationary regime.
- Stationary regime is more than stationary points, in general
  (except for reversible case)
- Control on mean field limit may give new insights

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.(2010)Gomez-Serrano, Graham, and Le Boudec] Javier Gomez-Serrano, Carl Graham, and Jean-Yves Le Boudec. The Bounded Confidence Model Of Opinion Dynamics. Technical Report arxiv:1006.3798.v1, 2010.

[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.