Mean Field Methods for Computer and Communication Systems

Jean-Yves Le Boudec
EPFL

July 2010

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 $X^N(t/I(N))$
- In discrete time
  - I(N) = O(1): mean field limit is in discrete time
  - I(N) = O(1/N): mean field limit is in continuous time

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

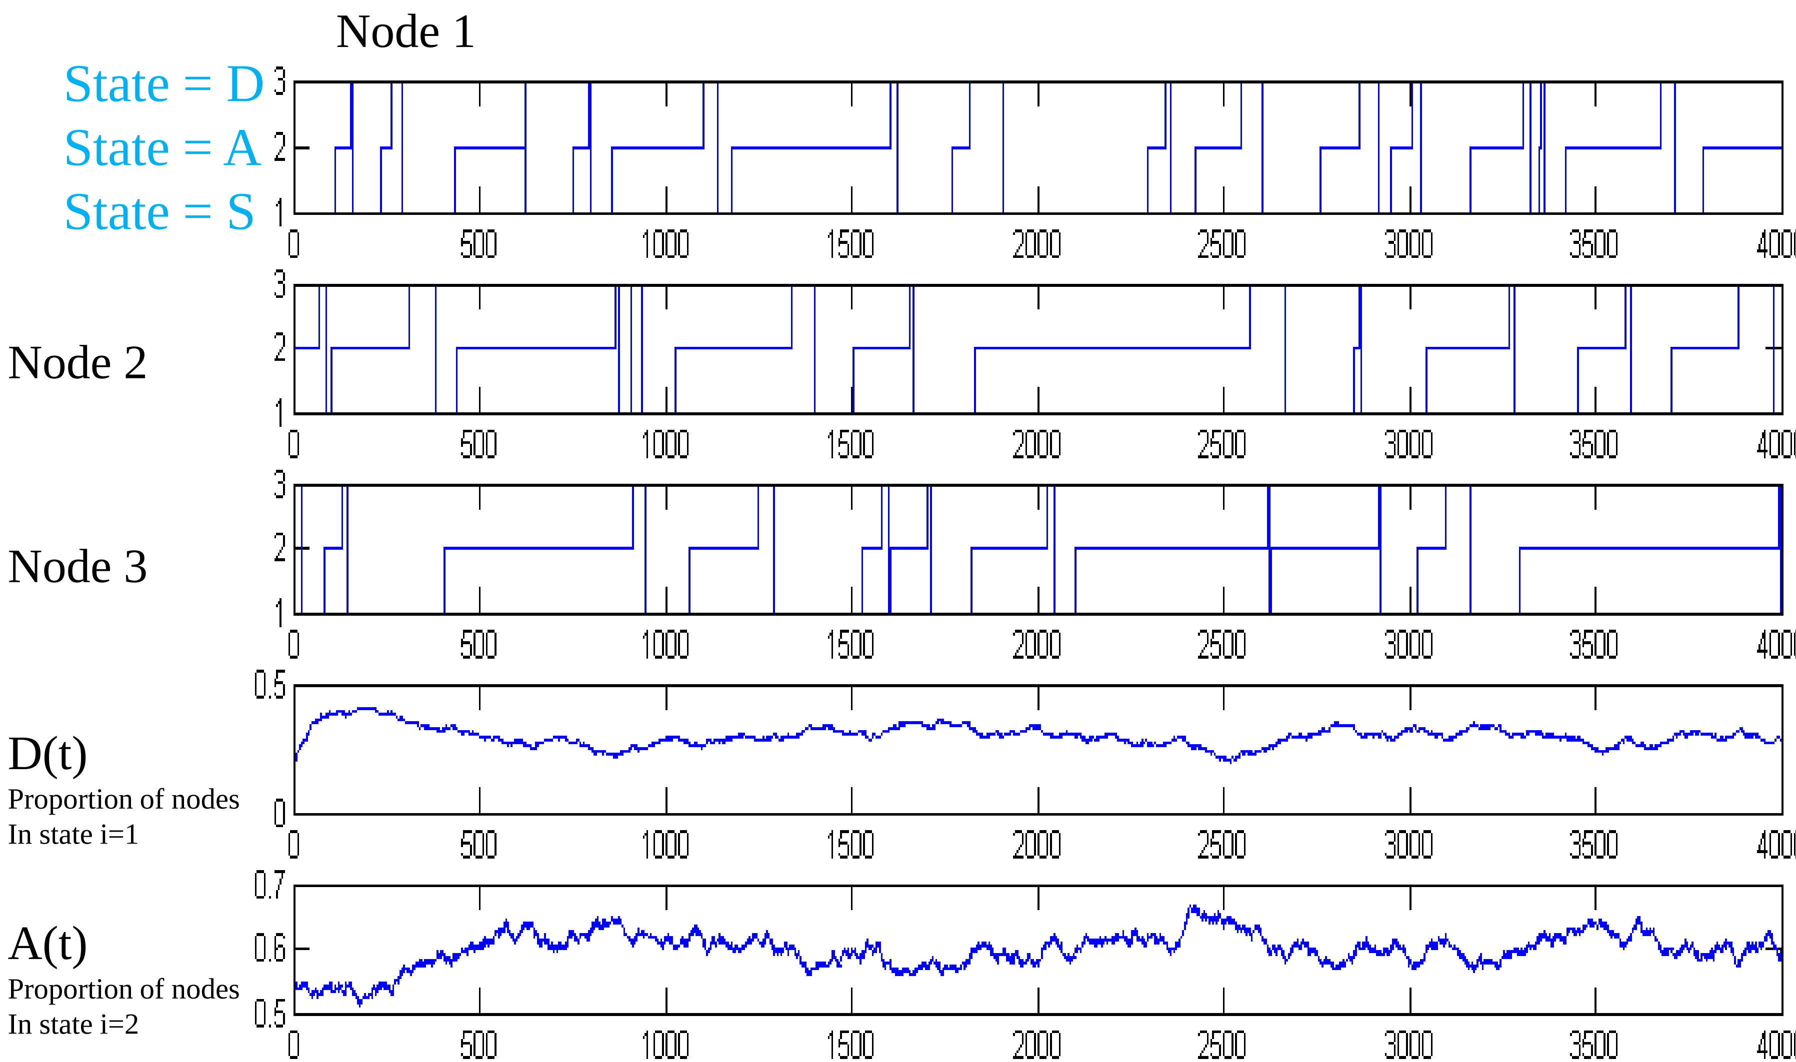

$$\beta = 0.01, \delta_A = 0.005, \delta_D = 0.0001, \alpha_0 = \alpha = 0.0001, h = 0.3, r = 0.1, \lambda = 0.0001$$

Sample Runs with N = 1000

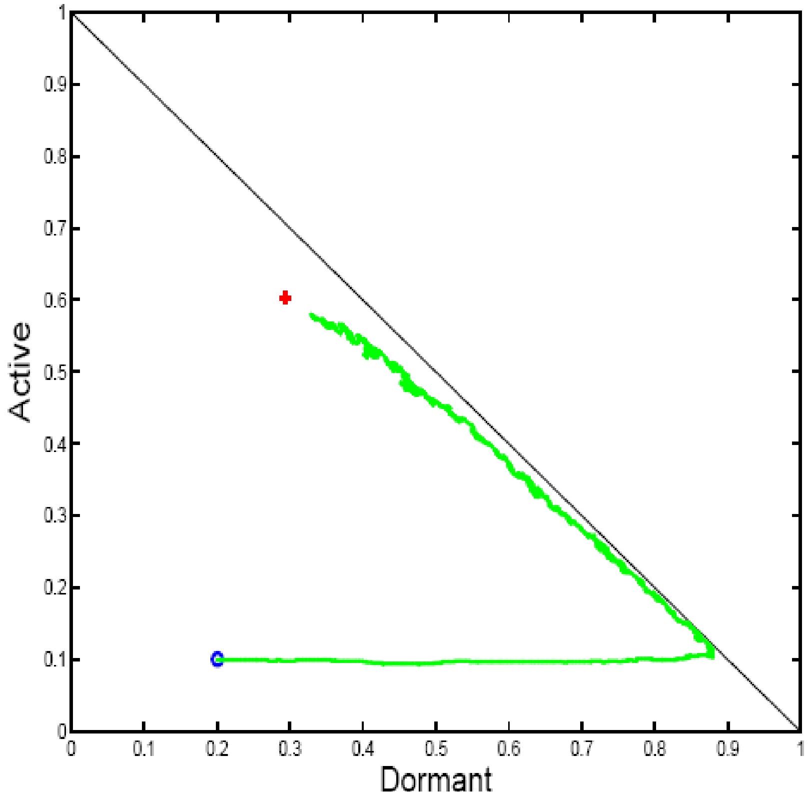

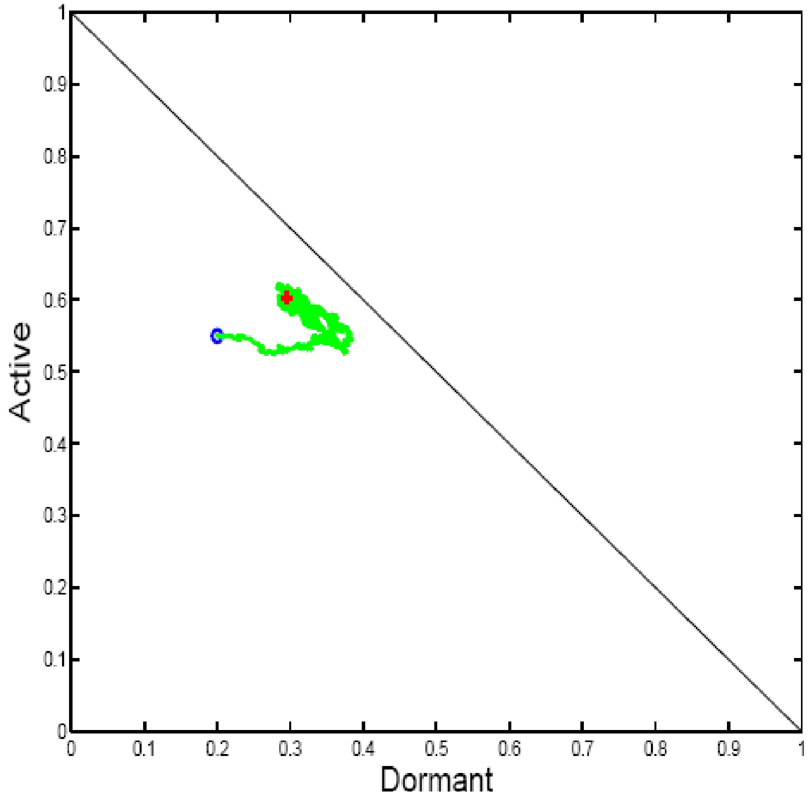

$$\beta = 0.01, \delta_A = 0.005, \delta_D = 0.0001, \alpha_0 = \alpha = 0.0001, h = 0.3, r = 0.1, \lambda = 0.0001$$

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: HTTP Metastability

- N flows between hosts and servers
- Flow *n* is OFF or ON
- Time is discrete, occupancy measure = proportion of ON flows
- At every time step, every flow switches state with proba matrix that depends on the proportion of ON flows

- I(N) = 1; Mean field limit is an iterated map (discrete time)

[Baccelli et al.(2004)Baccelli, Lelarge, and McDonald]

- Other Examples where the mean field limit is in discrete time:
  TCP flows with a buffer in [Tinnakornsrisuphap and Makowski(2003)]

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

Example: Age of Gossip

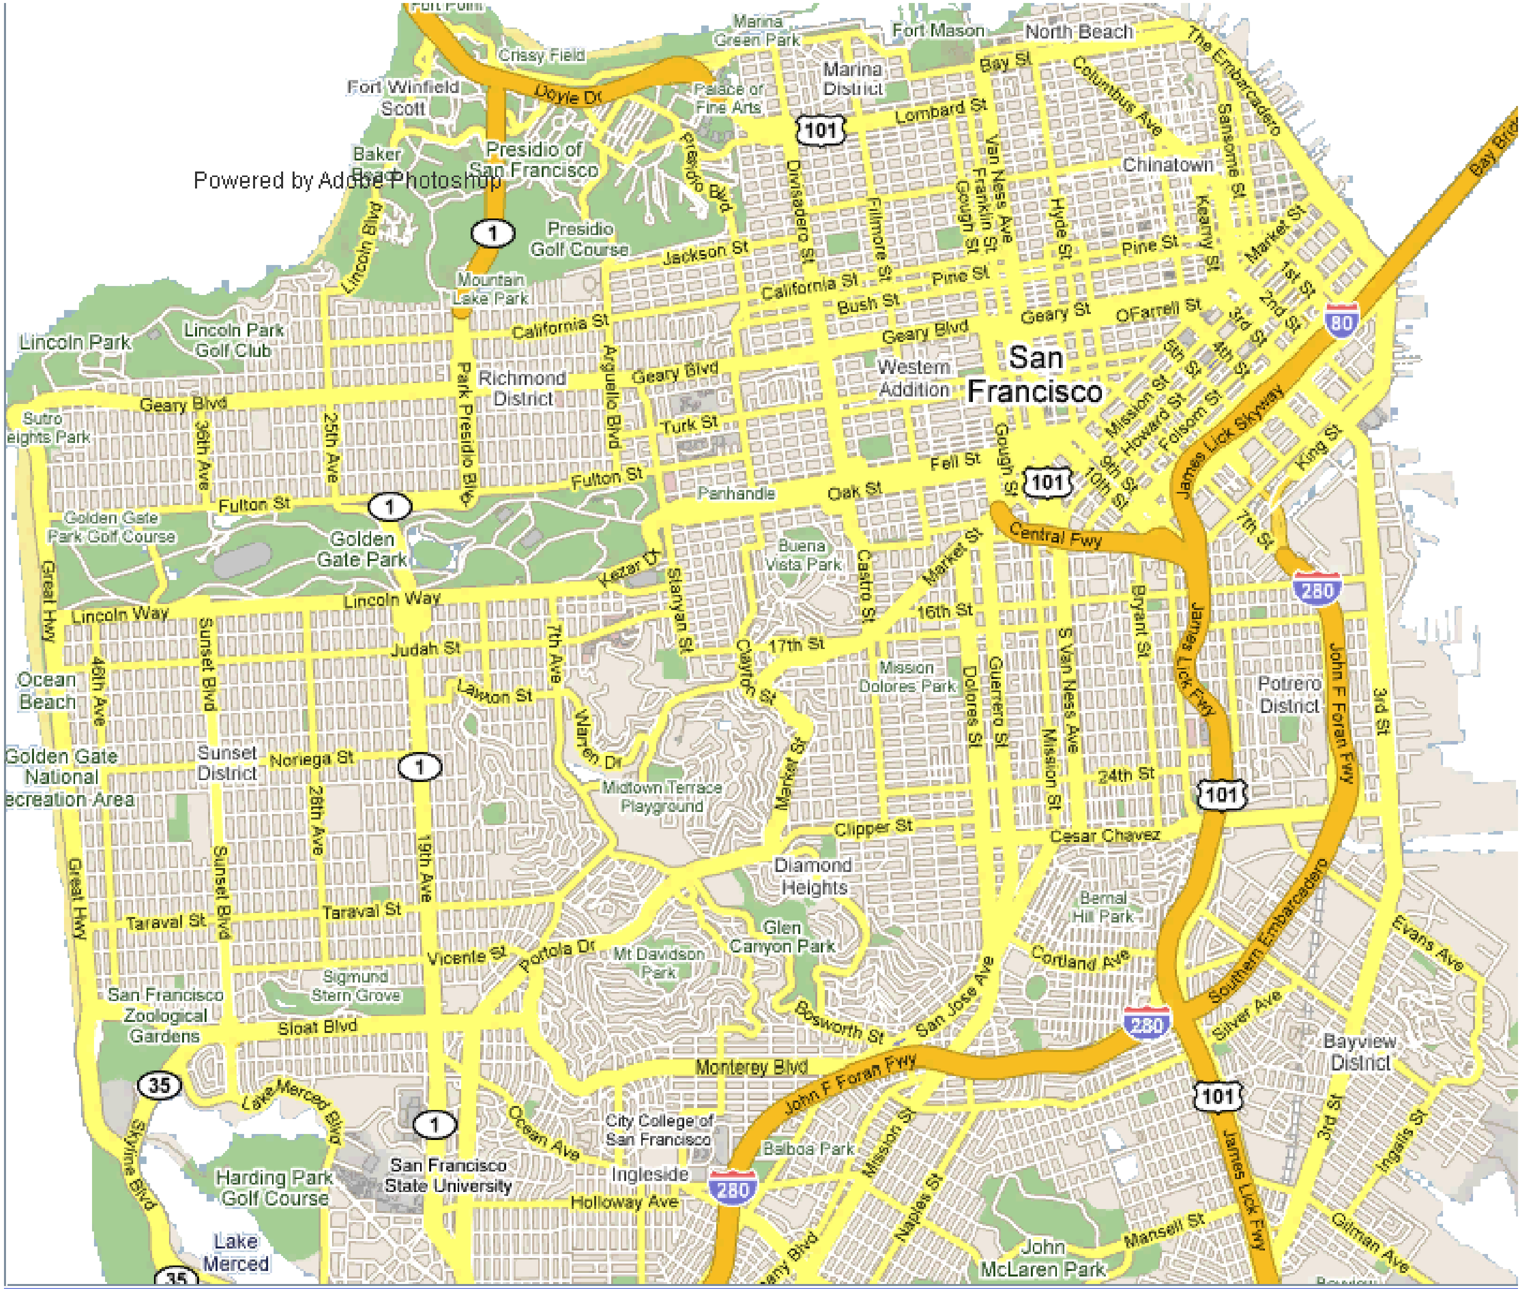

Example: Age of Gossip

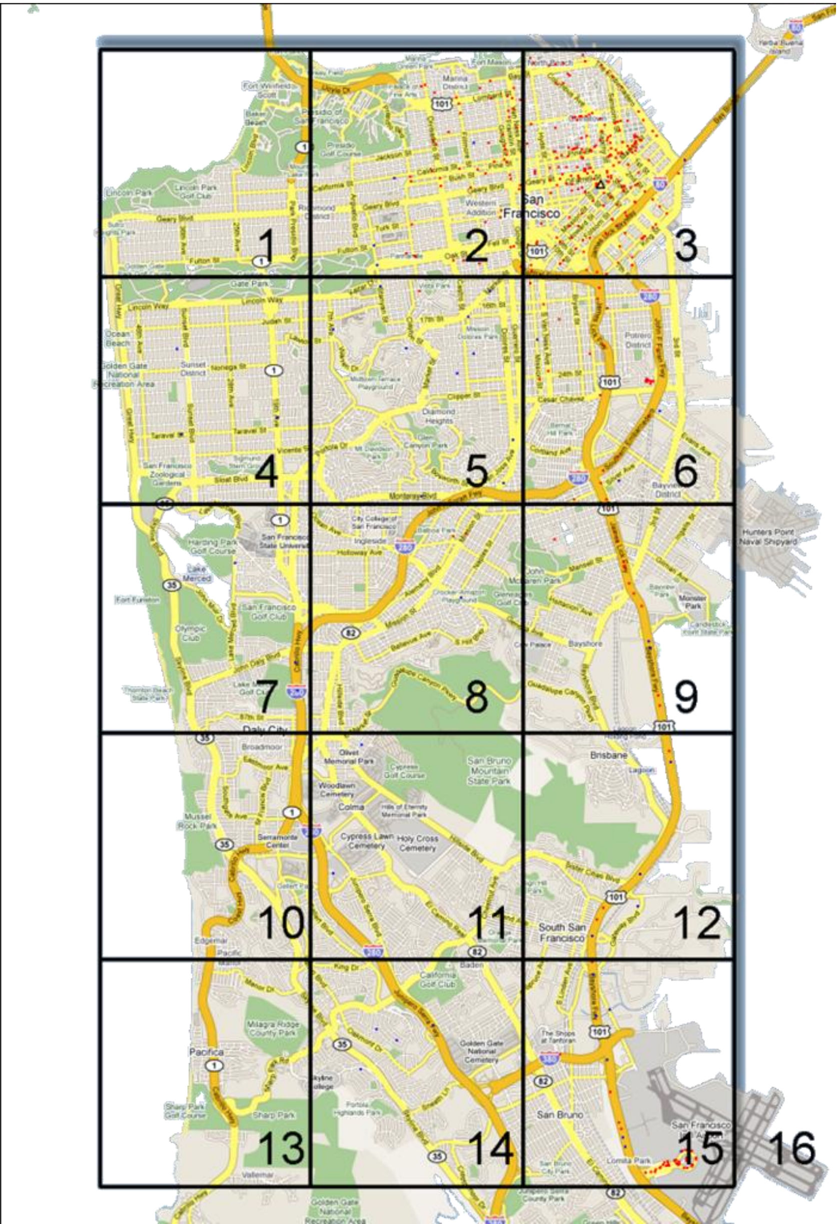

- Mobile node state = (c, t)
  c = 1 … 16 (position)
  $t \in R^+$ (age)

- 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)Chaintreau, Le Boudec, and Ristanovic]

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)]

Contents

- Mean Field Interaction Model
- The Mean Field Limit
- Convergence to Mean Field
- The Decoupling Assumption
- Optimization

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

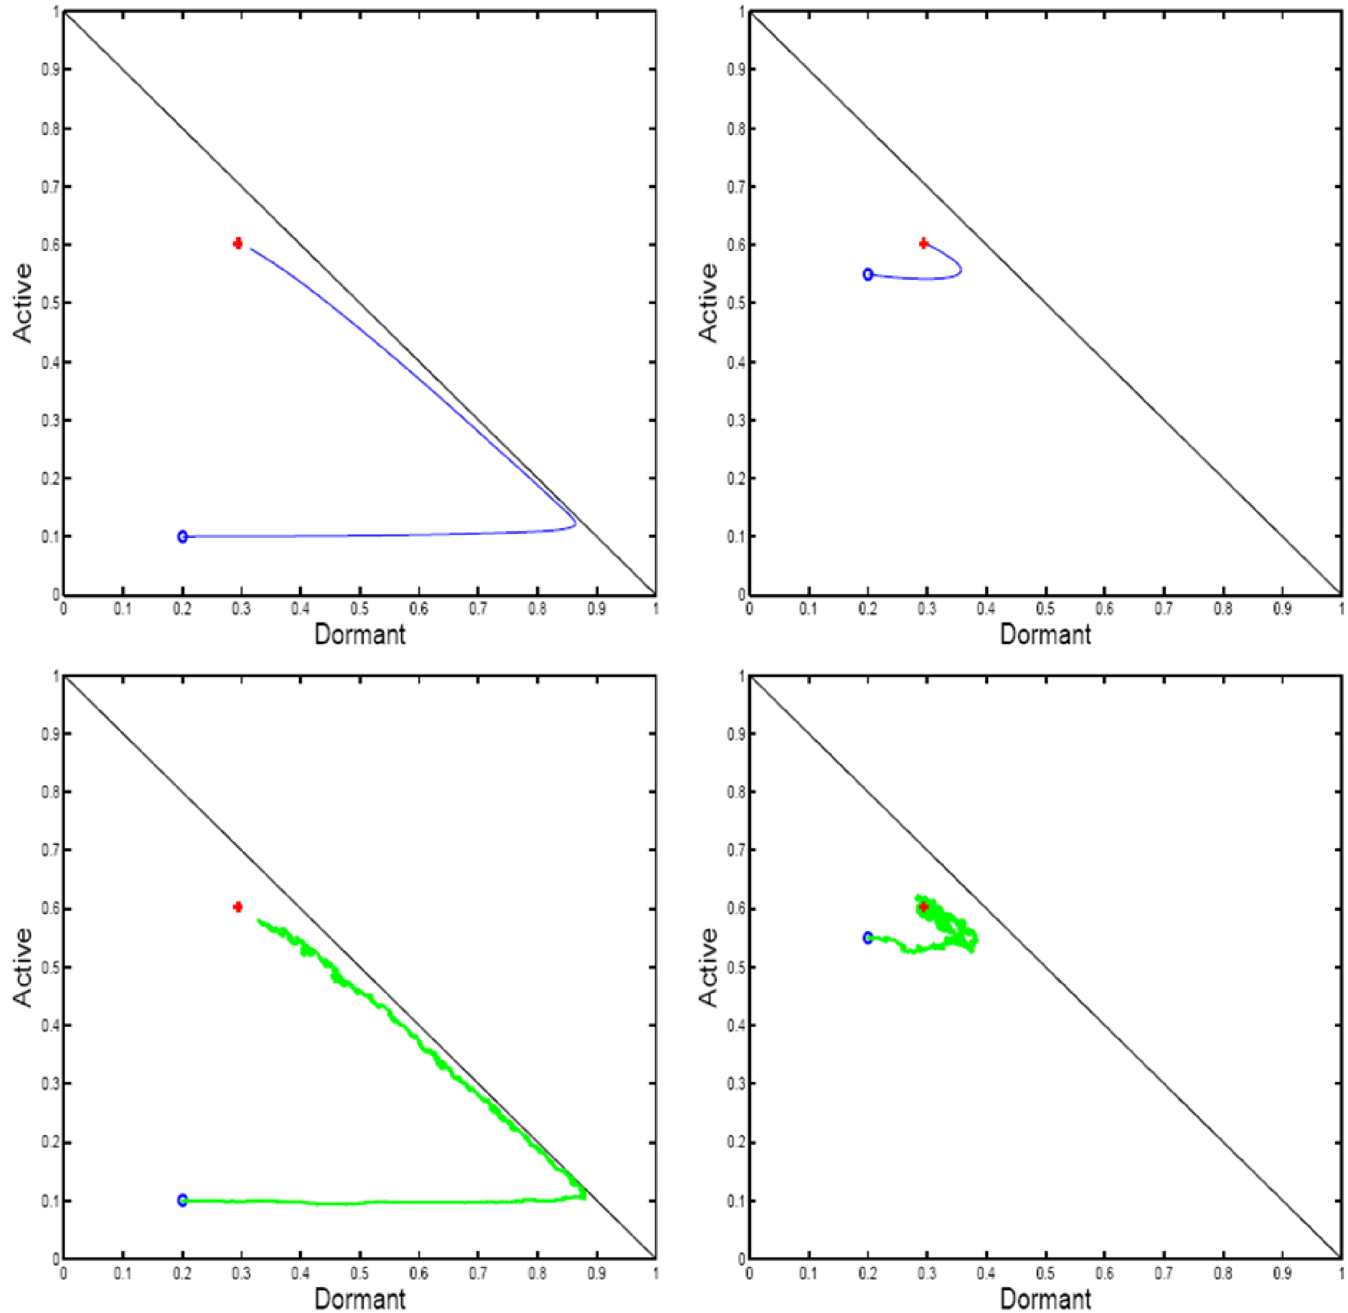

Propagation of Chaos is Equivalent to Convergence to a Deterministic Limit

Definition

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 ([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

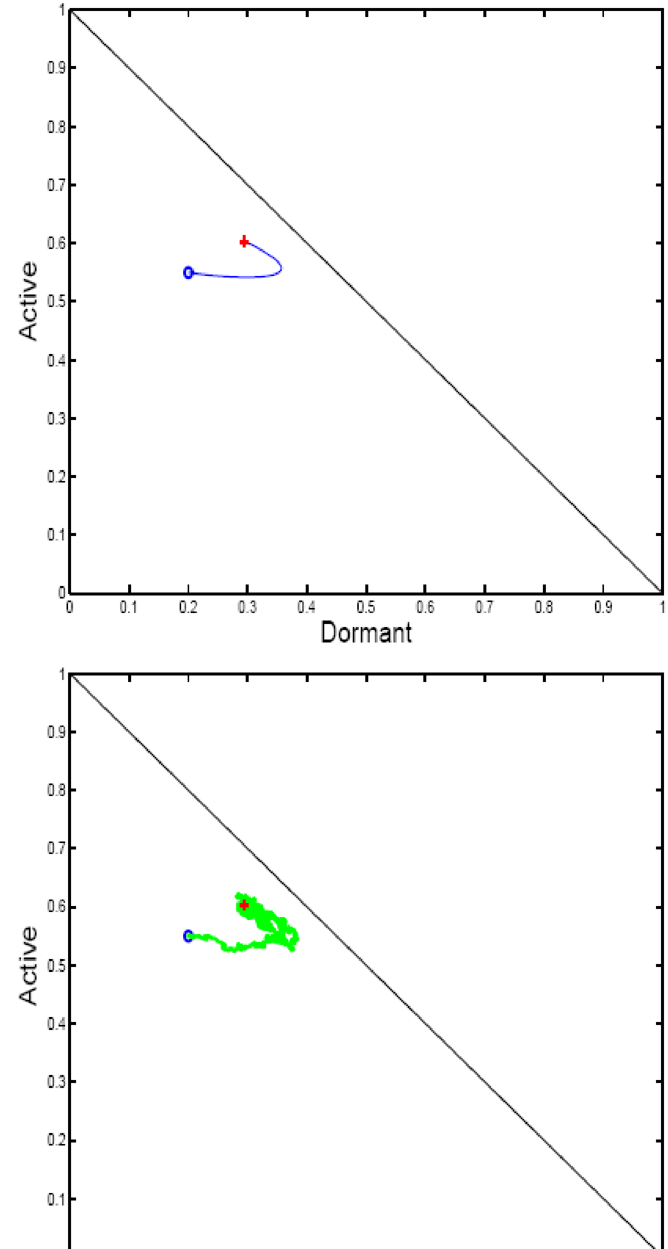

Example: Decoupling Assumption

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

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

- The decoupling assumptions says that

$$p_j^N(t/N|i) \approx p_j(t|i)$$

where *p(t|i)* is a continuous time, non homogeneous process

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

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

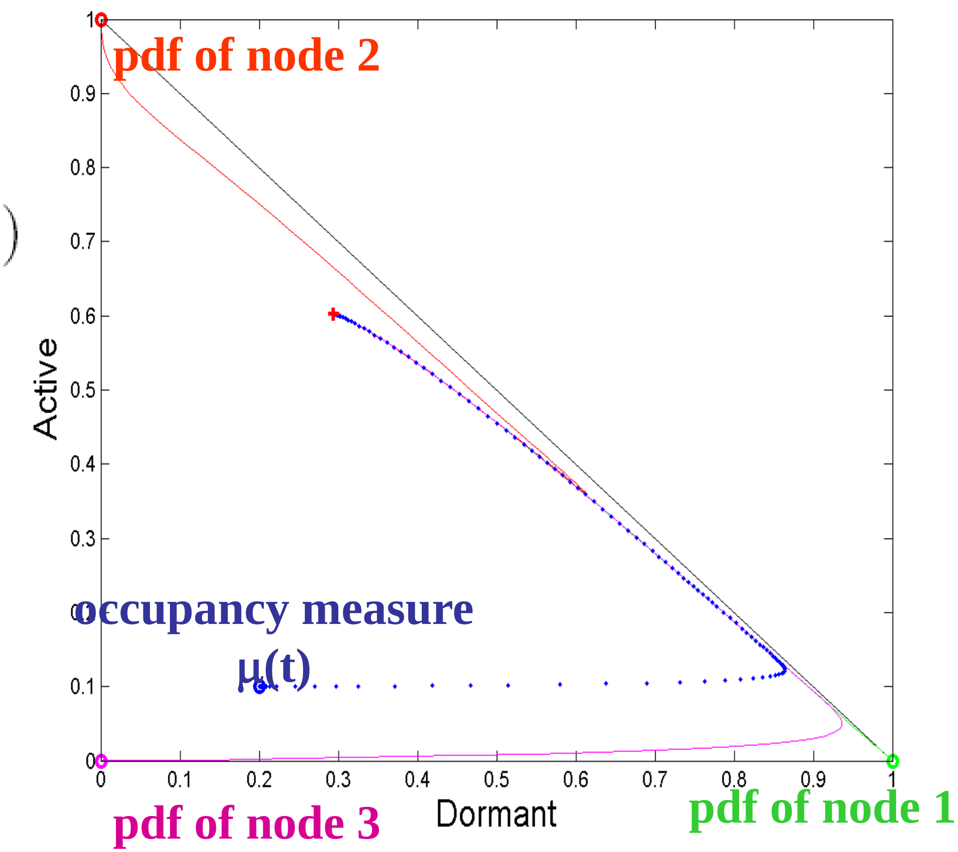

[Tembine et al.(2009)Tembine, Le Boudec, El-Azouzi, and Altman]
[Le Boudec et al.(2007)Le Boudec, McDonald, and Mundinger]

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

Contents

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)]

Kurtz's Theorem

- 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, 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]

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
- $(X_1^N(0), ..., X_N^N(0))$ 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 $\sup_{0 \leq k \leq K} \mathbb{P}\left(\|M^N(k) - m(k)\|\right) \to 0$ 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)]
$I(N) = 1/N$, 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 $\mathcal{G}\cdot\varphi(x,x') = \int \varphi(y,y')f(y,y'|x,x')dydy'$ then
  $\int |f(y,y'|x,x')|\, dydy' \leq \Lambda$, 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

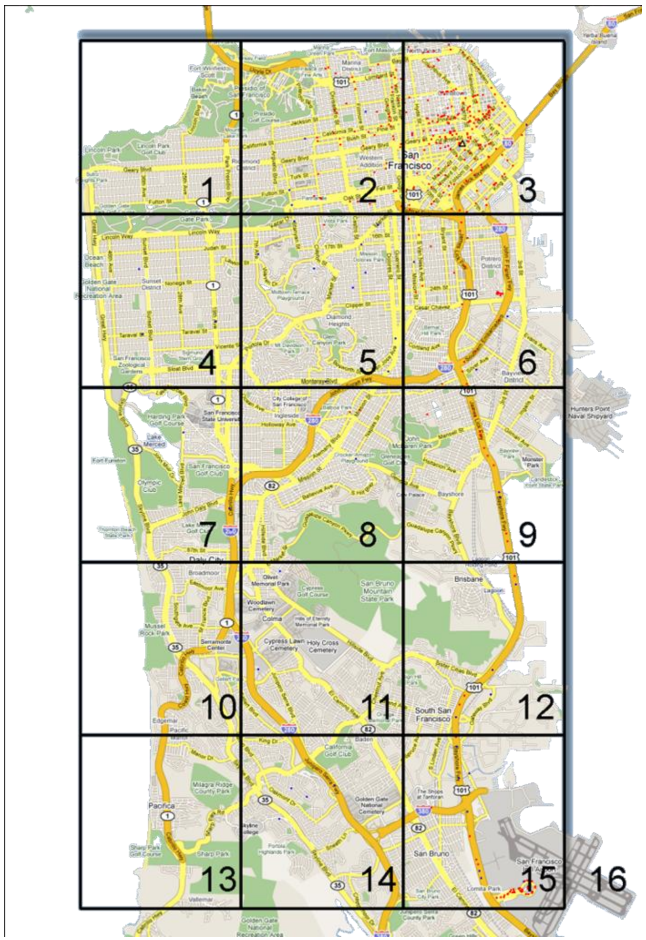

- [Graham and Méléard 1997] applies, i.e. mean field convergence occurs for iid initial conditions
- [Chaintreau et al.(2009)Chaintreau, Le Boudec, and Ristanovic] shows more, i.e. weak convergence of initial condition suffices

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

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

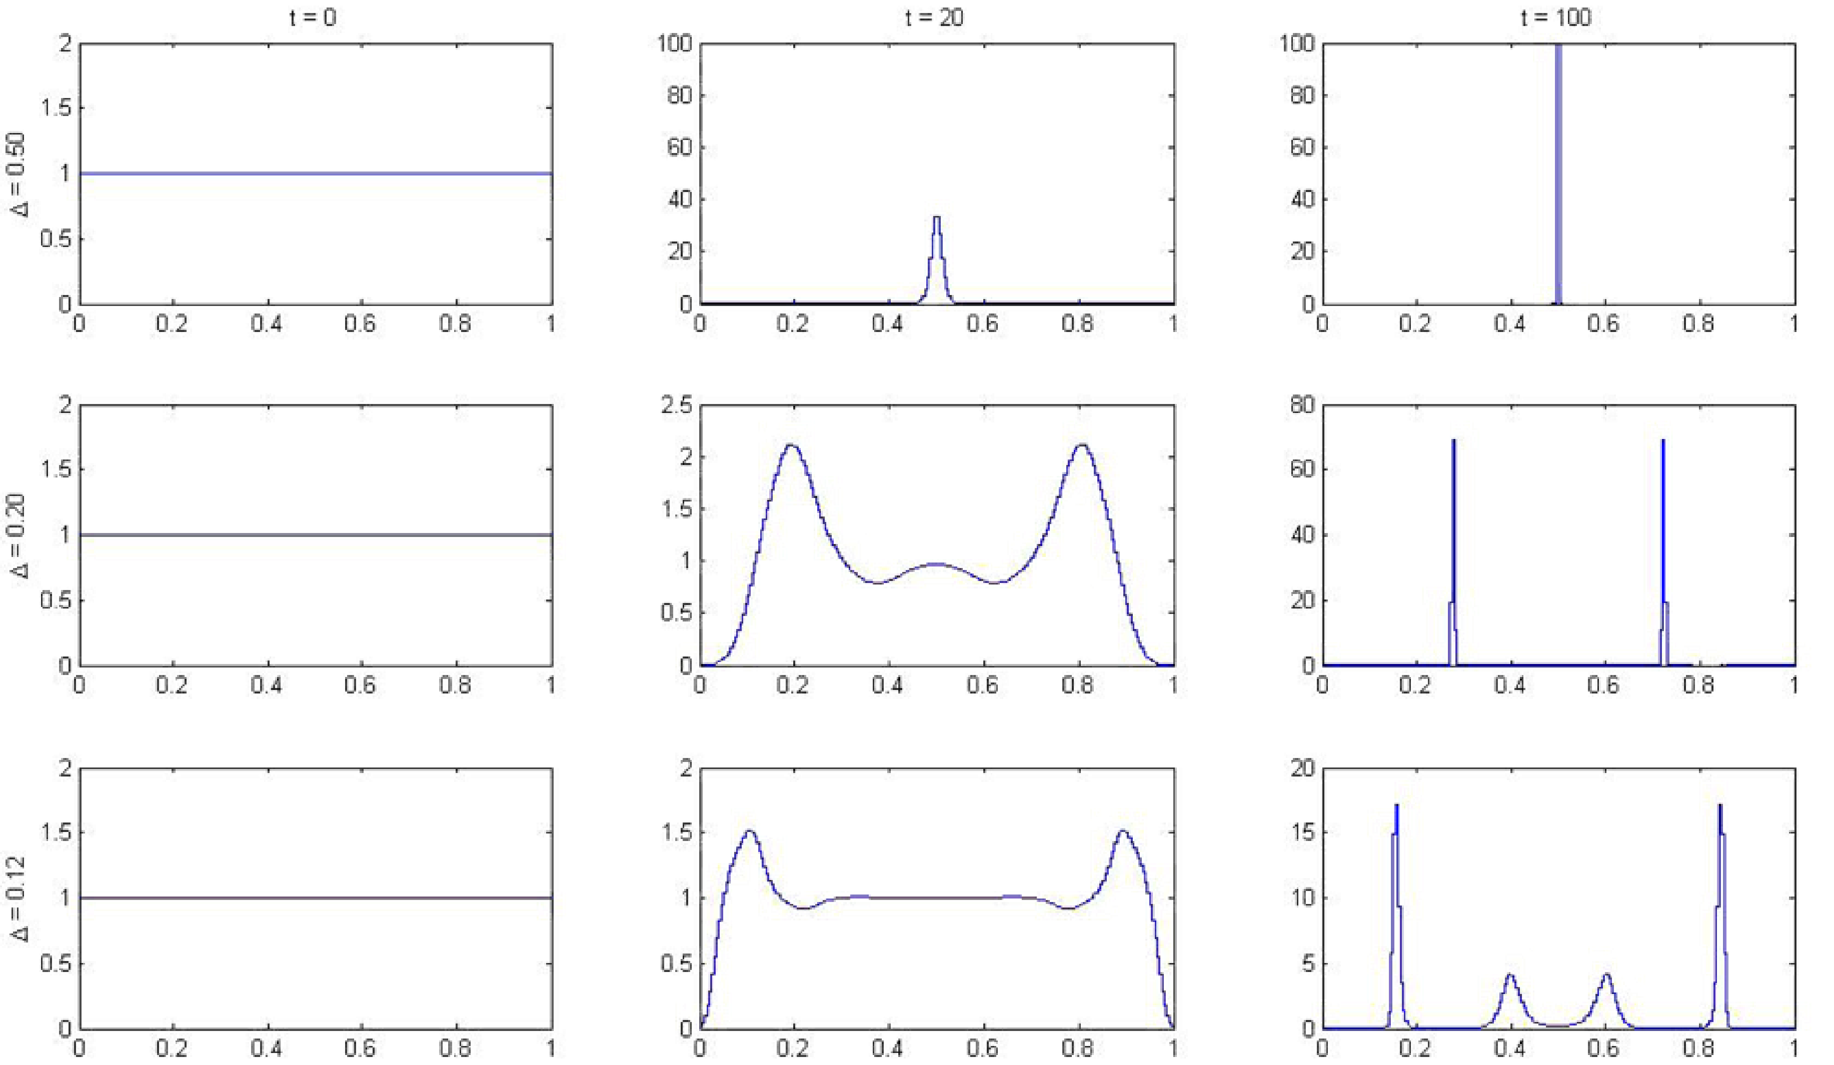

Is There Convergence to Mean Field ?

- Intuitively, yes
- Discretized version of the problem:
  - Make set of ratings discrete
  - Generic results apply: number of meetings is upper bounded by 2
  - There is convergence for any initial condition such that $M^N(0) \to m_0$
- This is what matlab does.
- However, there can be no similar result for the real version of the problem
  - There are some initial conditions such that $M^N(0) \to m_0$ while there is not convergence to the mean field
  - There is convergence to mean field if initial condition is iid from m_0

Contents

Decoupling Assumption

- Is true when mean field convergence holds, i.e. almost always
- It is often used in stationary regime

Example

- 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
- We are in the good case: the diagram commutes

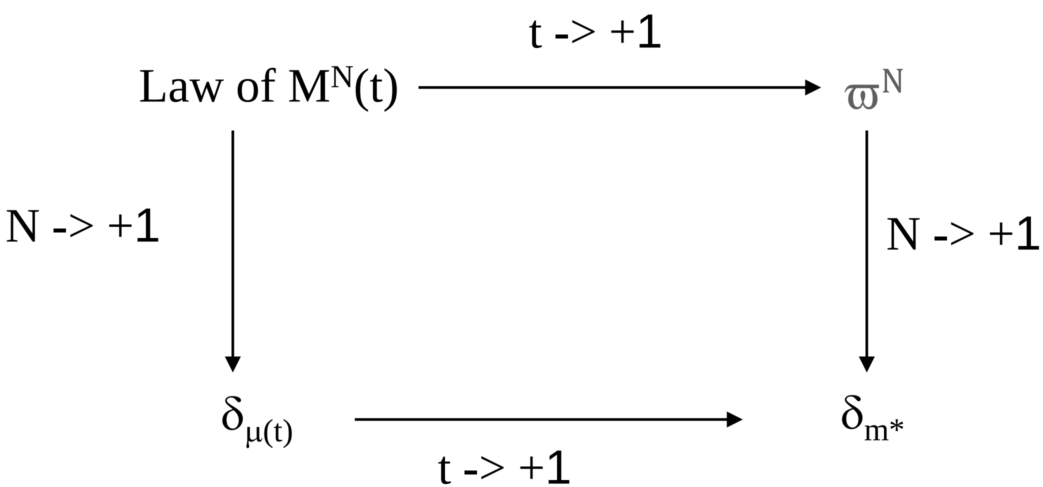

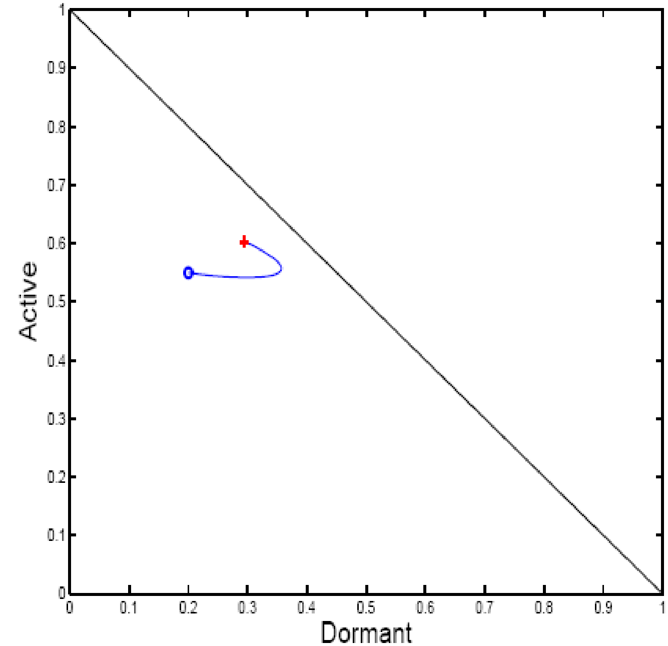

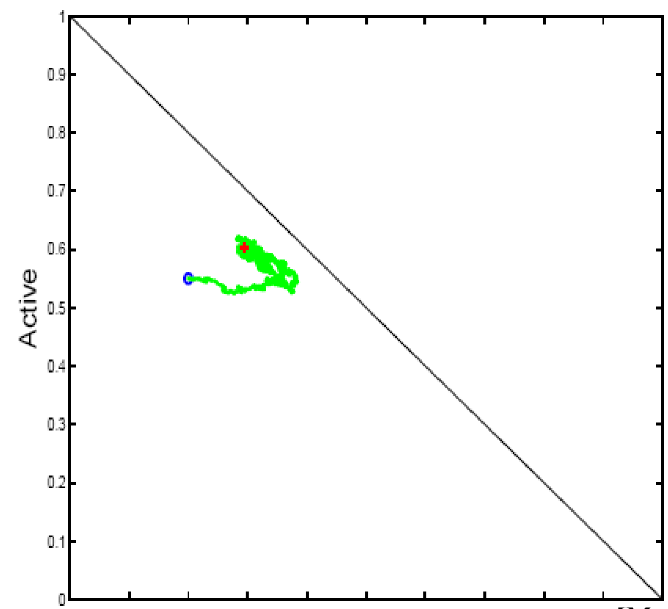

Counter-Example

- The ODE does not converge to a unique attractor (limit cycle)
- Assumption H does not hold; does the decoupling assumption still hold ?

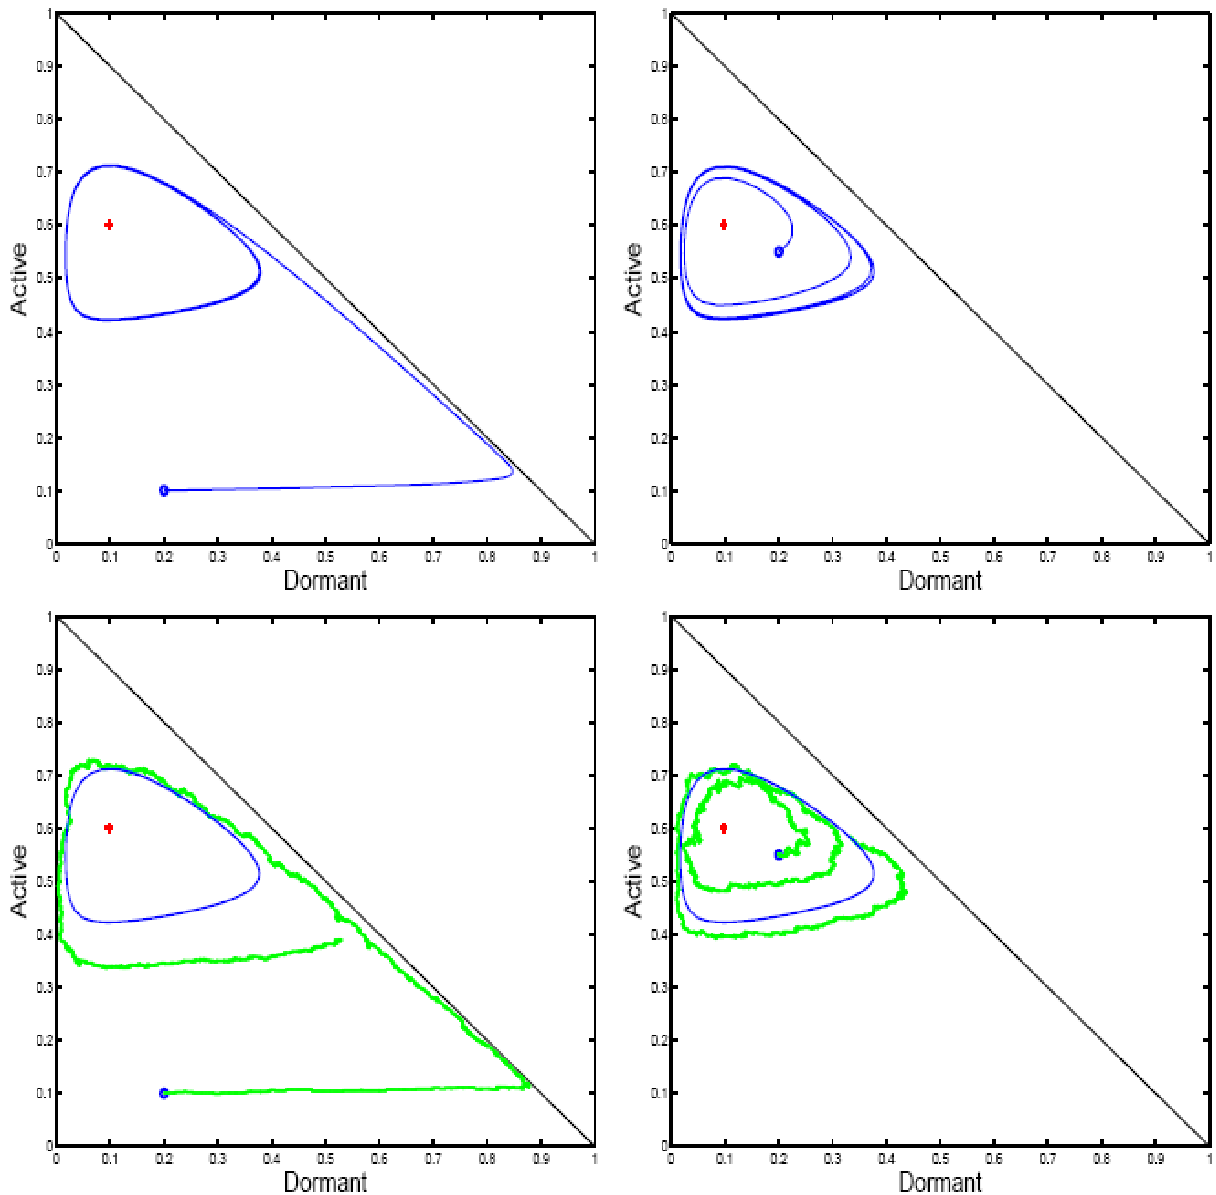

Same as before Except for one parameter value

$h = 0.1$ instead of 0.3

Decoupling Assumption Does Not Hold Here In Stationary Regime

- 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

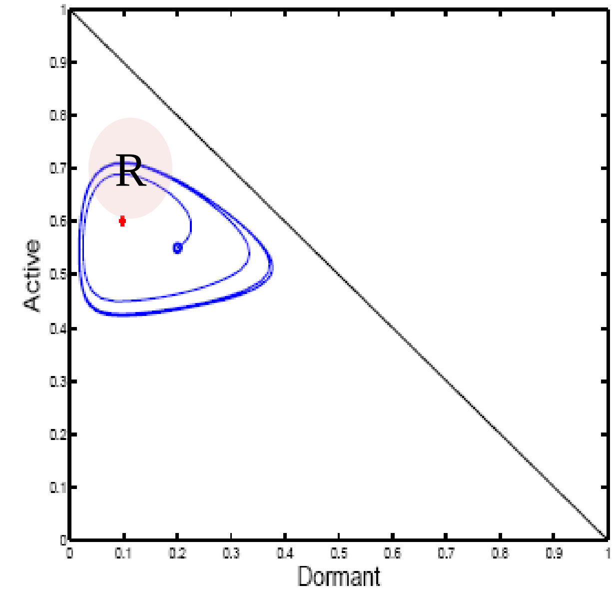

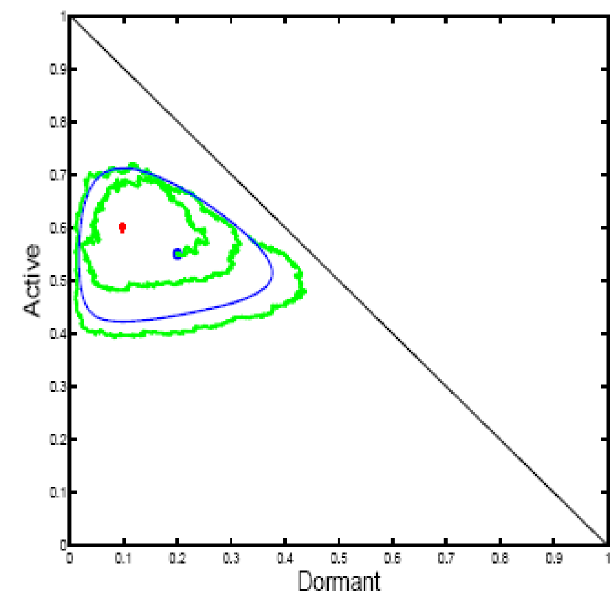

Numerical Example

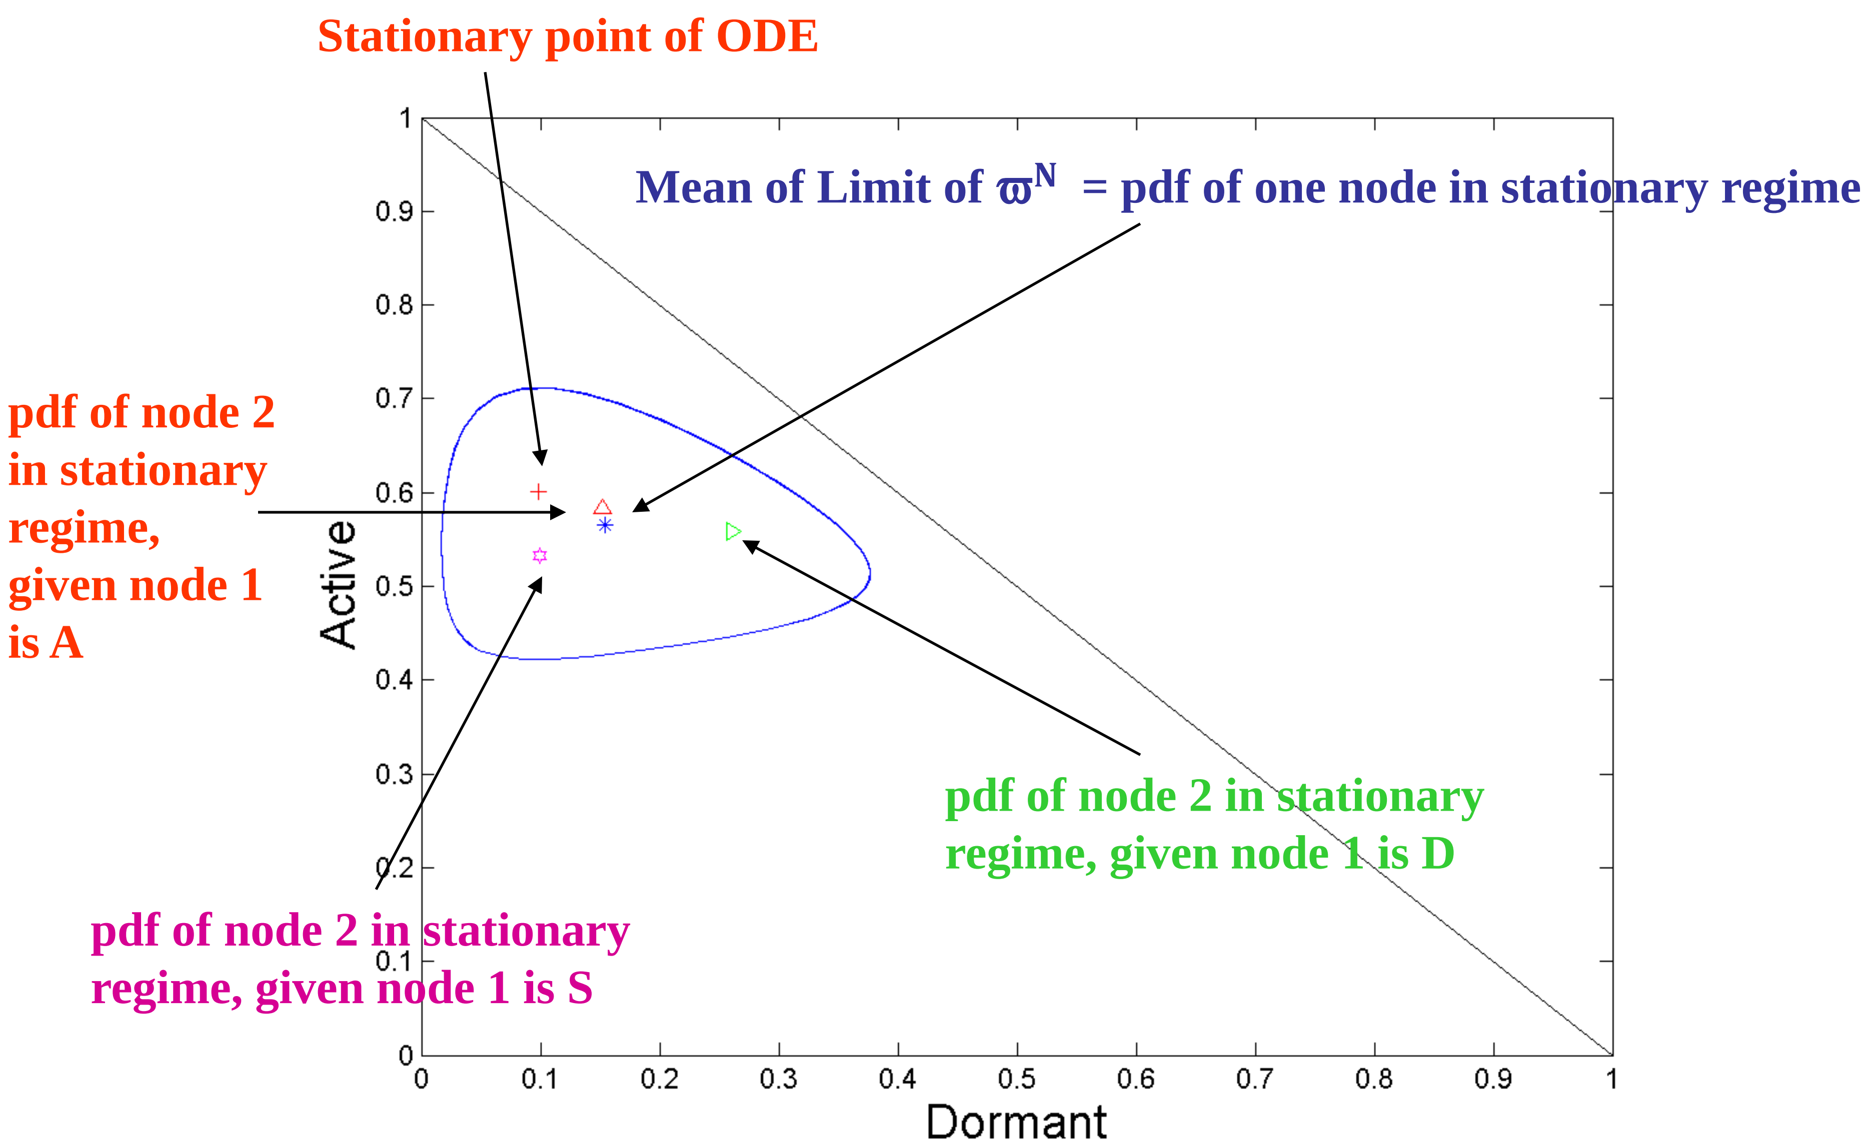

Numerical Results ($h = 0.1$).

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

Simplified Analysis 2
Decoupling Assumption (Stationary Regime)

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

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

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

Where is the Catch ?

- Fluid approximation and fast simulation result say that nodes *m* and *n* are asymptotically independent

- But we saw that nodes may not be asymptotically independent

… is there a contradiction ?

The Diagram Does Not Commute

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

- 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

Generic Result for Stationary Regime

- Original system (stochastic):
  - $(X^N(t))$ is Markov, finite, discrete time
  - Assume it is irreducible, thus has a unique stationary proba ν^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

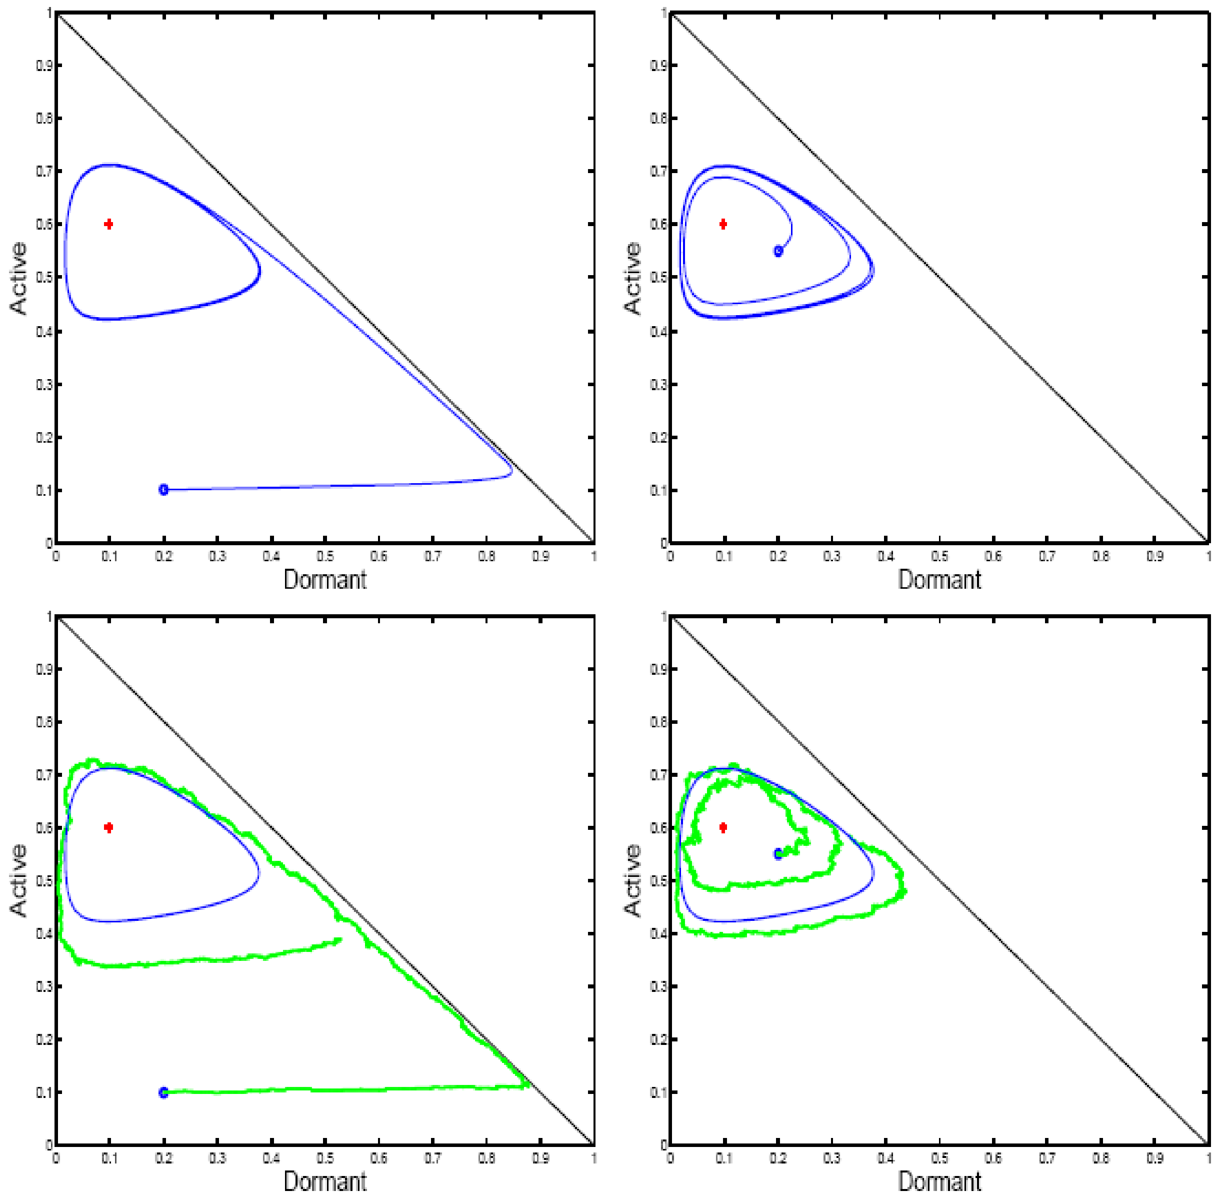

Quiz

- $M^N(t)$ is a Markov chain on E={(a, b, c) ¸ 0, a + b + c =1, a, b, c 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.

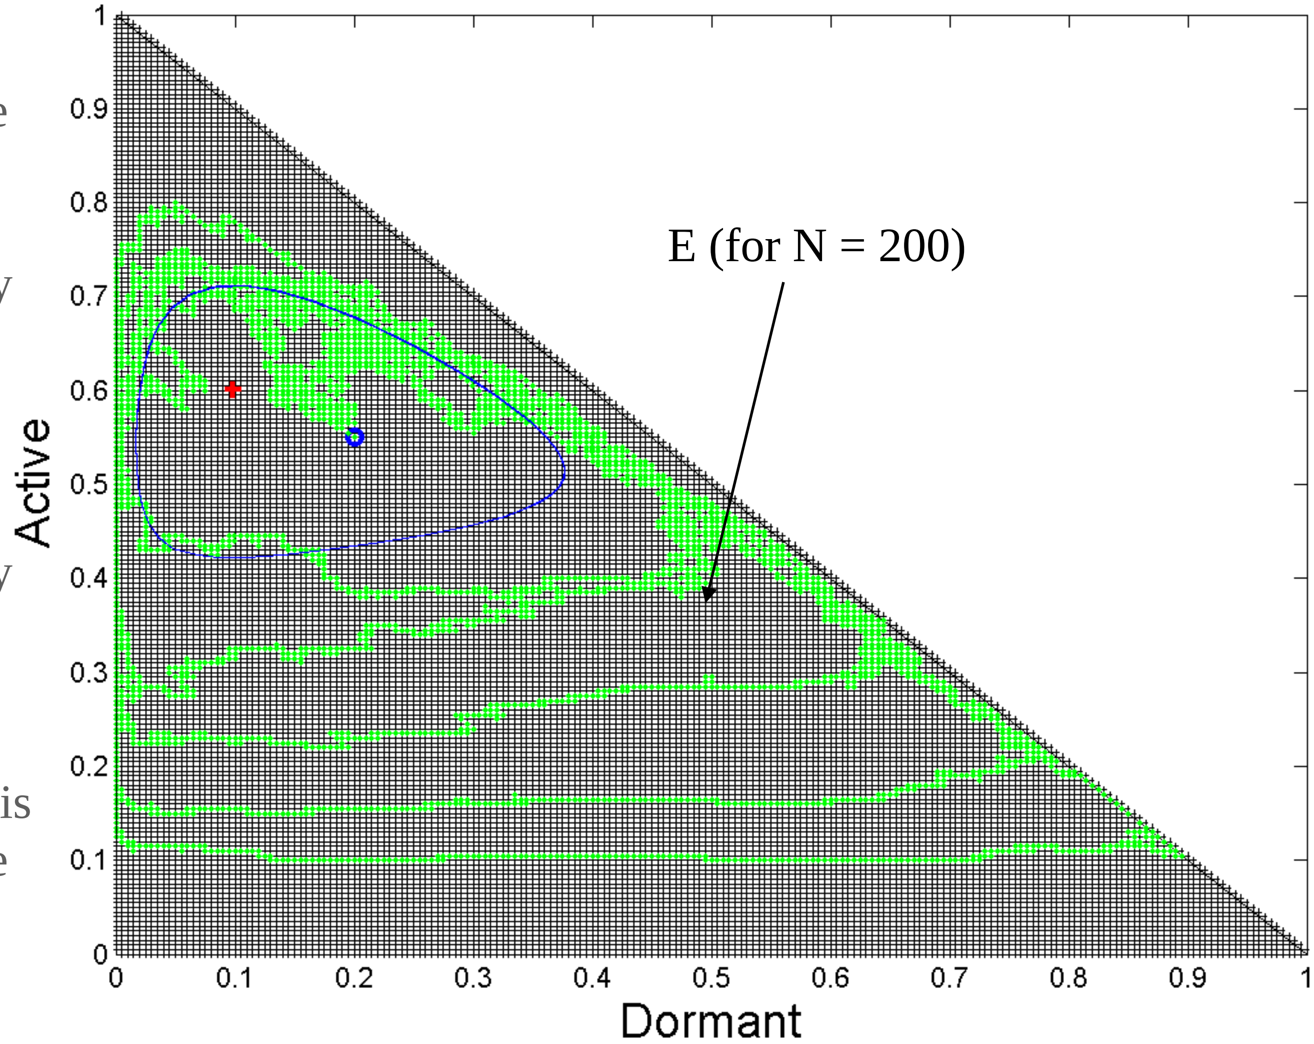

Decoupling Assumption in Stationary Regime Holds under (H)

- For large N the *decoupling assumption* holds at any fixed time t
- It holds in stationary regime under assumption (H)
  - (H) ODE has a unique global stable point to which all trajectories converge
- Otherwise the *decoupling assumption* may not hold in stationary regime

- It 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

- Study the ODE !

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

- Essential assumption is

(H) $\mu(\tau)$ 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*

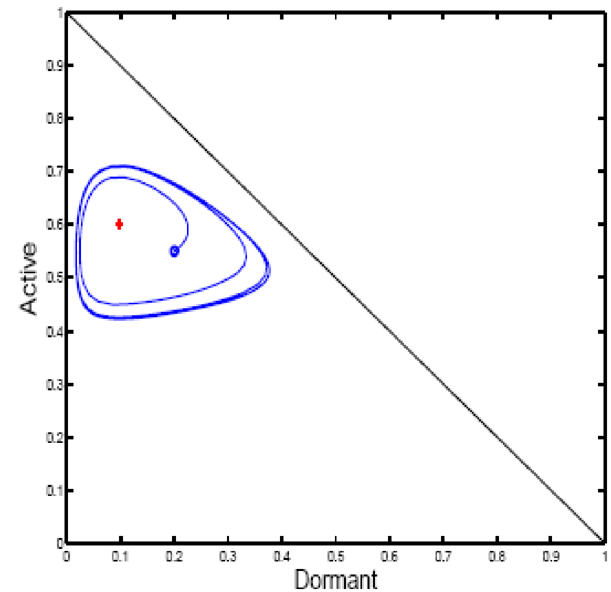

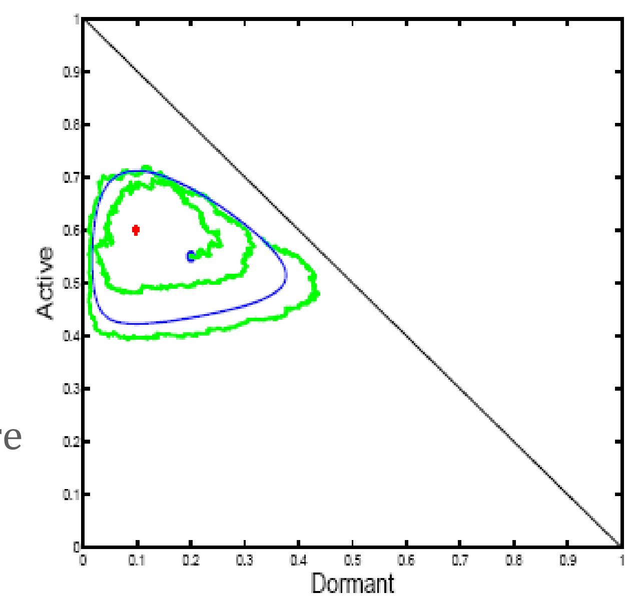

Example: 802.11 with Heterogeneous Nodes

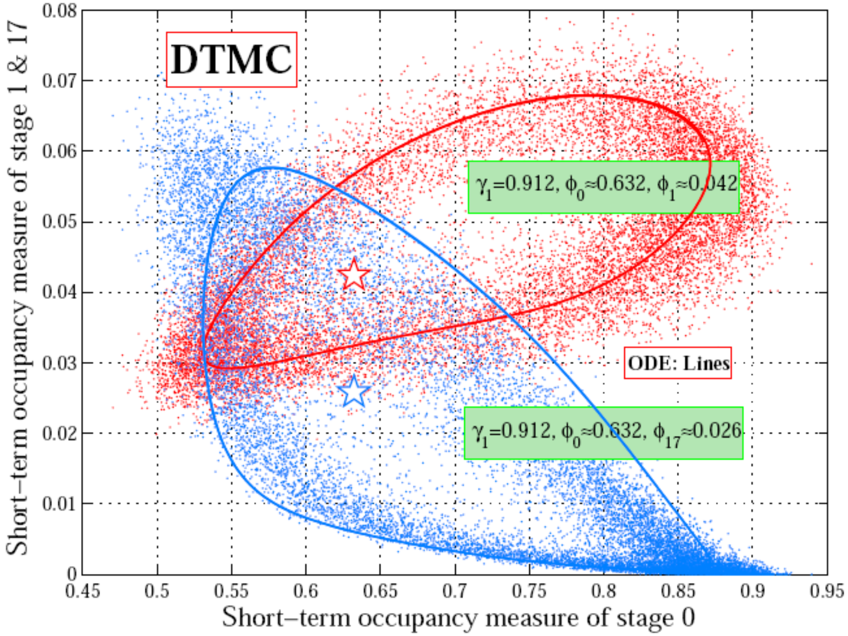

- [Cho2010]

  Two classes of nodes with heterogeneous parameters (restransmission probability)

  Fixed point equation has a unique solution

  There is a limit cycle

Contents

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)Tembine, Le Boudec, El-Azouzi, and Altman] 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 $N \longrightarrow \infty$ under uniform convergence and continuity of $R^N \rightarrow R$. 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 $\varepsilon-$ optimal strategies (resp. $\varepsilon-$equilibrium) for game with infinite N.*

Centralized Control

- [Gast et al.(2010)Gast, Gaujal, and Le Boudec]
- Markov decision process
  - Finite state space per object, discrete time, N objects
  - Transition matrix depends on a control policy
  - For large N the system without control converges to mean field
- 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
- Writing the mean field equations is simple and provides a first order approximation
- Mean field is much more than a fluid approximation: decoupling assumption / fast simulation
- Decoupling assumption in stationary regime is not necessarily true.
- Mean field equations may reveal emerging properties
- 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.

[Benaïm and Le Boudec(2008)] M. Benaïm 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.

[Benaïm et al.(2006)Benaïm, Hofbauer, and Sorin] M. Benaïm, J. Hofbauer, and S. Sorin. Stochastic approximations and differential inclusions. *SIAM Journal on Control and Optimization*, 44(1):328–348, 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.

[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.

[Gomez, Graham, Le Boudec 2010] Gomez-Serrano J., Graham C. and Le Boudec, JY, « The Bounded Confidence Model of Opinion Dynamics », preprint, http://infoscience.epfl.ch/record/149416

[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. http://arxiv.org/abs/1004.2342v1, 2010.

[Graham and Méléard(1994)] C. Graham and S. Méléard. Chaos hypothesis for a system interacting through shared resources. *Probability Theory and Related Fields*, 100(2):157–174, 1994.

[Graham and Méléard(1997)] C. Graham and S. Méléard. Stochastic particle approximations for generalized Boltzmann models and convergence estimates. *The Annals of probability*, 25(1):115–132, 1997.

[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*, 2009.

[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 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.