Performance Evaluation

Second Part

Lecture 1

Giovanni Neglia
INRIA – EPI Maestro
9 January 2012

Course organization

- Part 2.A
  - Fluid models to overcome the limitations of Markov Processes analysis
  - A specific networking problem
    - Epidemic Routing in Delay Tolerant Networks

- Part 2.B
  - Introduction to game theory

Material

- Slides
- References part A
  - Mean Field
    - Mean Field Methods for Computer and Communication Systems: A Tutorial, Jean-Yves Le Boudec
    - A class of mean field interaction models for computer and communication systems, Benaïm, Le Boudec, Journal Performance Evaluation, Vol. 65 Issue 11-12, Nov., 2008
  - A survey with pointers to continuous time Markov processes and links to stochastic approximation and propagation of chaos
    - Ch. 2 of Nicolas Gast's PhD thesis "Optimization and Control of Large Systems, Fighting the Curse of Dimensionality"

Material

- Slides
- References part A
  - Dynamical Processes on Complex Networks, Barrat, Barthélemy, Vespignani, Cambridge Press
    - Random graphs models, ch.3
    - Methodological approaches, ch. 4
    - Epidemiological models, ch. 9

Material

- References part A
  - Routing in DTNs
    - Markovian models
      - Message Delay in Mobile Ad Hoc Networks, R. Groenevelt, G. Koole, and P. Nain, Performance, Juan-les-Pins, October 2005
      - Impact of Mobility on the Performance of Relaying in Ad Hoc Networks, A. Al-Hanbali, A.A. Kherani, R. Groenevelt, P. Nain, and E. Altman, IEEE Infocom 2006, Barcelona, April 2006
    - Fluid models
      - Performance Modeling of Epidemic Routing, X. Zhang, G. Neglia, J. Kurose, D. Towsley, Elsevier Computer Networks, Volume 51, Issue 10, July 2007, Pages 2867-2891

Material

- References part B
  - Game Theory and Strategy, Straffin, Mathematical Association,
    - Two-person zero-sum games
      - Matrix games
        - » Pure strategy equilibria (dominance and saddle points), ch 2
        - » Mixed strategy equilibria, ch 3
      - Game trees, ch 7
      - About utility, ch 9

Material

- References part B
  - Game Theory and Strategy, Straffin, Mathematical Association,
    - Two-person non-zero-sum games
      - Nash equilibria…
        - …And its limits (equivalence, interchangeability, Prisoner's dilemma), ch. 11 and 12
      - Strategic games, ch. 14
      - Evolutionary games, ch. 15

Evaluation

- 80% final exam
- 20% assignments (every two weeks)

Contacts

- giovanni.neglia@inria.fr
- INRIA, Lagrange building, last floor, L108
- For slides, assignements, etc.
  - www-sop.inria.fr/members/Giovanni.Neglia/perf11/

Performance Evaluation

Fluid Models

Giovanni Neglia
INRIA – EPI Maestro
9 January 2012

Outline

- Limit of Markovian models
- Mean Field (or Fluid) models
  - exact results
  - extensions
  - applications

A motivating example: epidemics

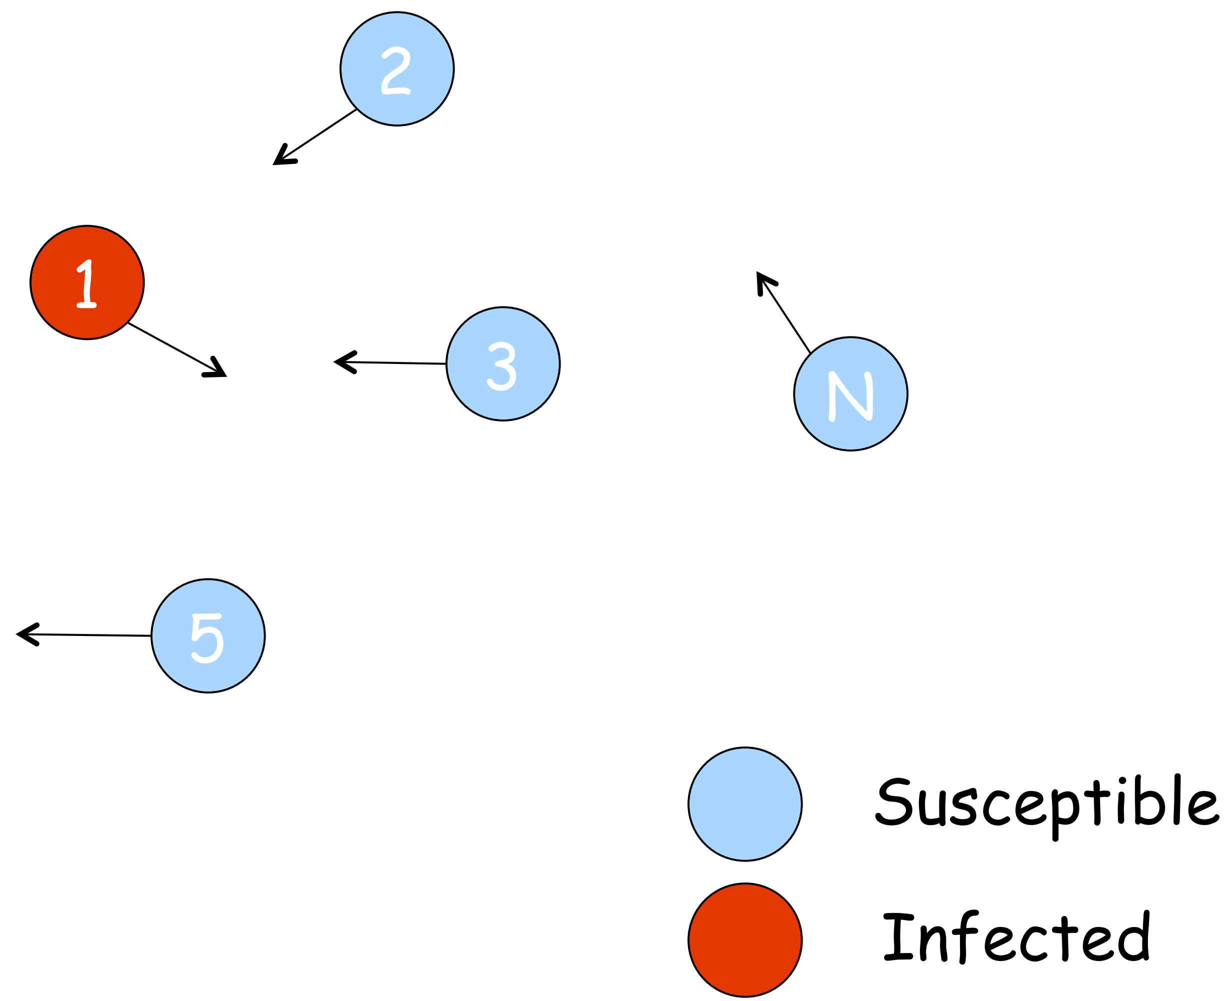

A motivating example: epidemics

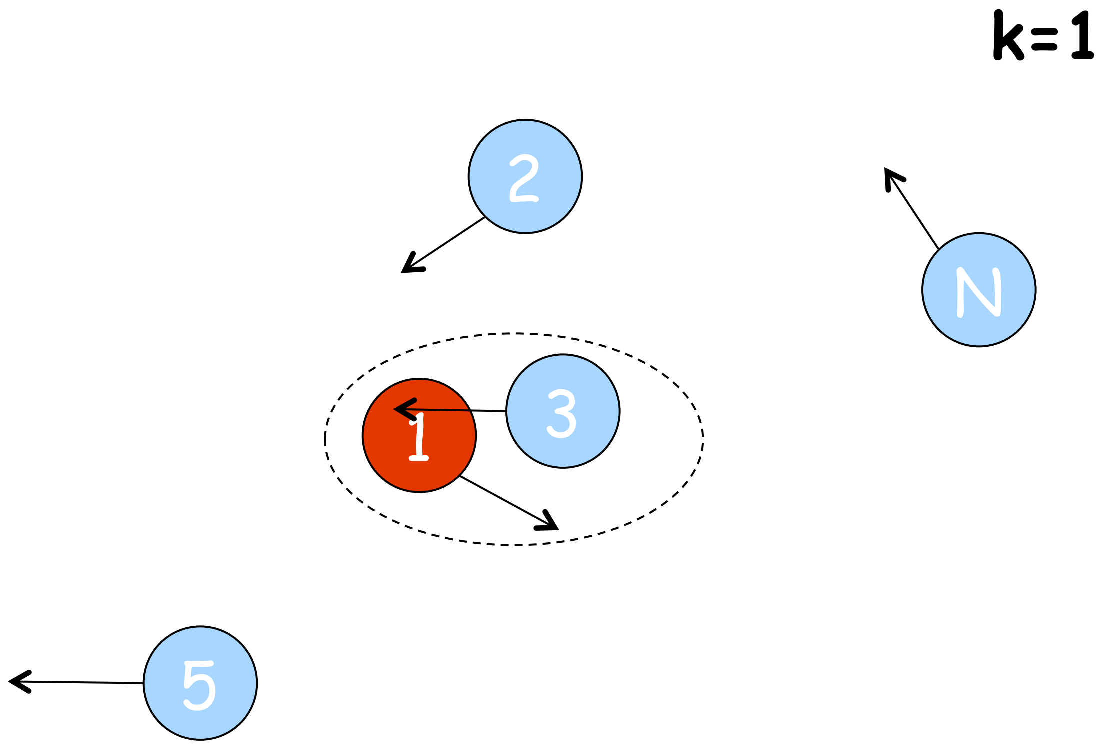

At each slot there is a probability p that two given nodes meet.
Assume meetings to be independent.

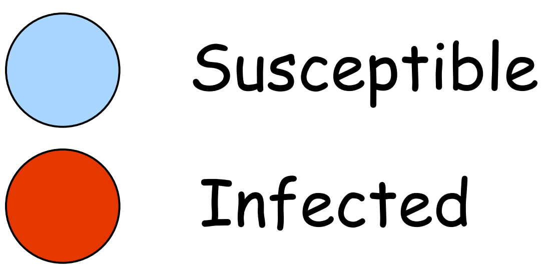

A motivating example: epidemics

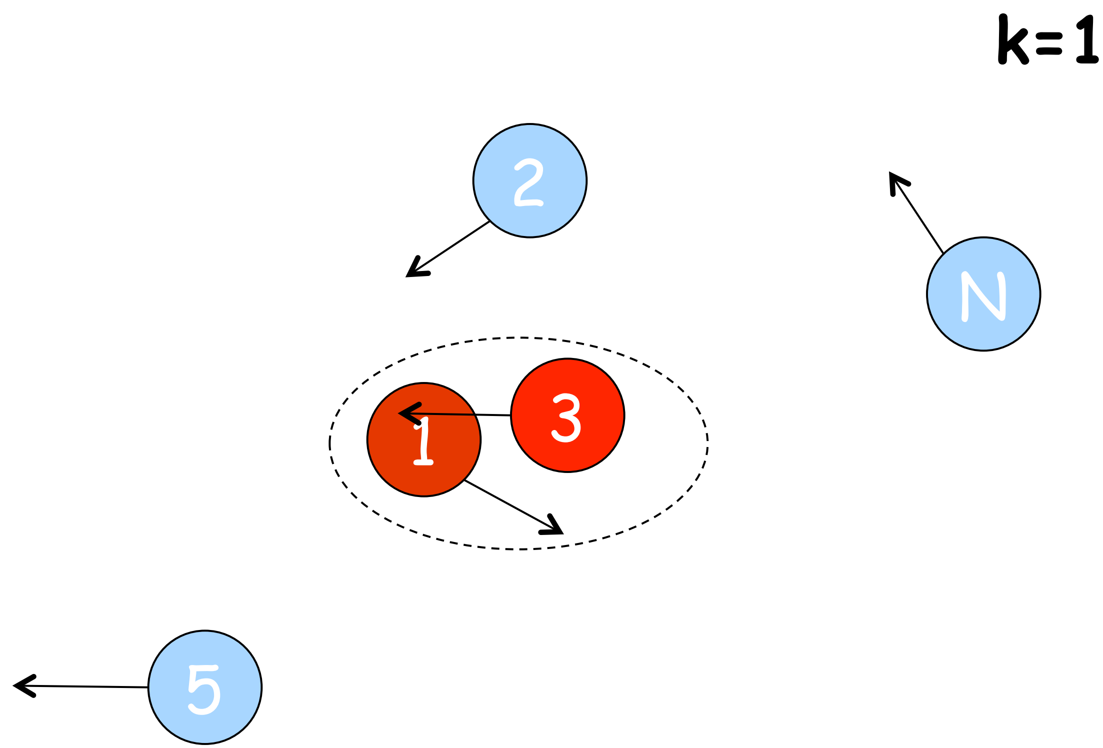

At each slot there is a probability p that two given nodes meet.
Assume meetings to be independent.

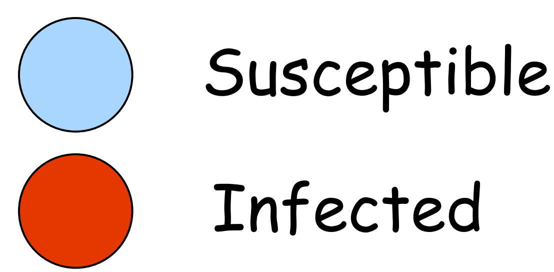

A motivating example: epidemics

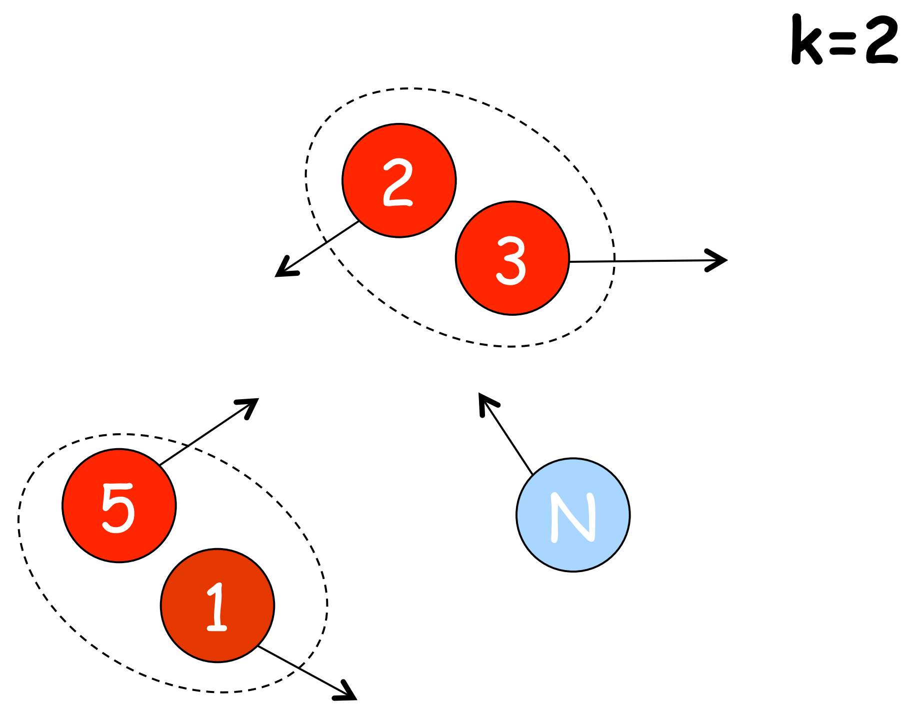

At each slot there is a probability p that two given nodes meet.
Assume meetings to be independent.

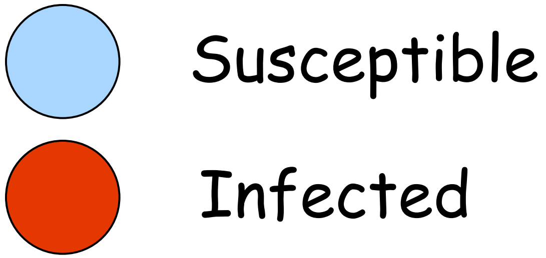

How do you model it?

- A Markov Chain
  - System state at time k is a vector specifying if every node is infected (1) or not (0)
    - e.g. (1,0,1,0,0), size: 2^5

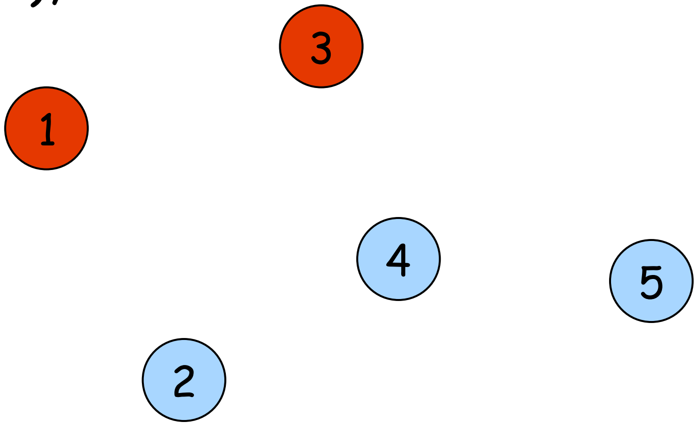

  - Probability transitions among states
    - e.g. Prob((1,0,1,0,0)->(1,1,1,0,0))=?

Transition probabilities

Prob((1,0,1,0,0)->(1,1,1,0,0))=?

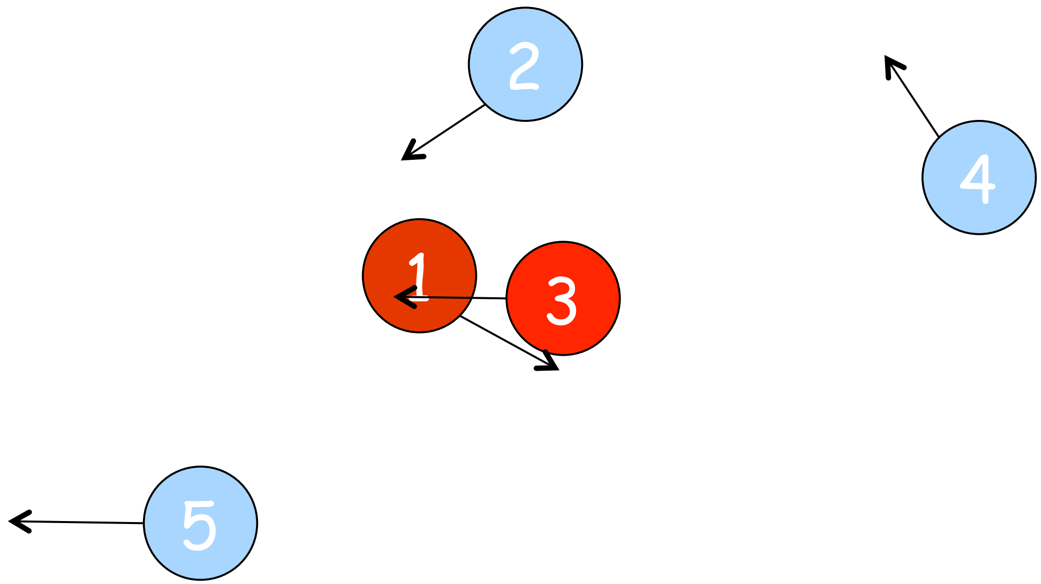

At slot k, when there are I(=I(k)) infected nodes, the prob. that node 2 gets infected is: $q_I=1-(1-p)^I$

Transition probabilities

Prob((1,0,1,0,0)->(1,1,1,0,0))=?

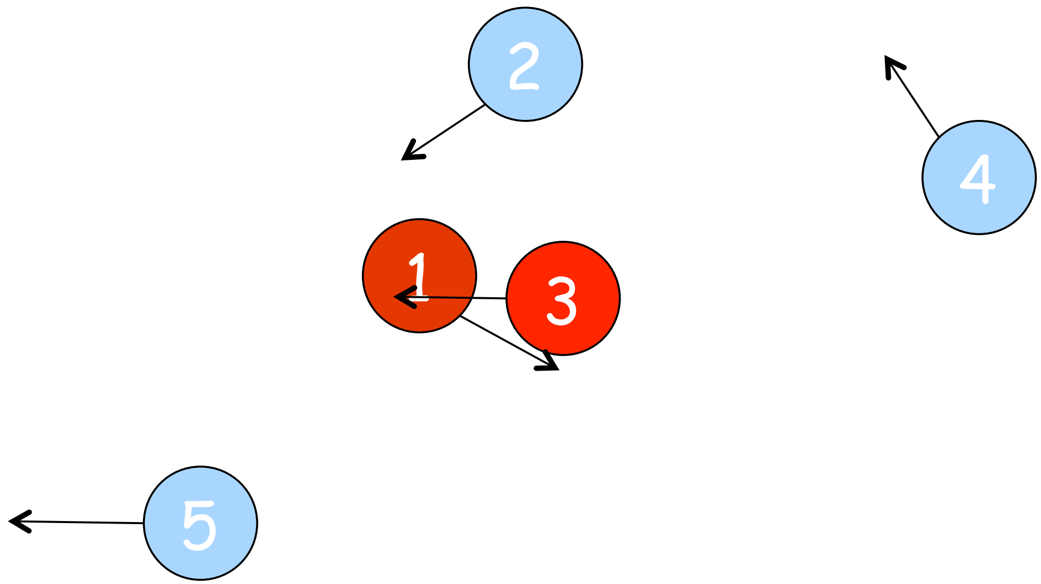

Prob((1,0,1,0,0)->(1,1,1,0,0))=$q_2(1-q_2)^2$
Where $q_I=1-(1-p)^I$

What to study and how

- P the transition matrix (2^Nx2^N)
- Transient analysis
  - $\pi(k+1)=\pi(k)P$,
  - $\pi(k+1)=\pi(0)P^{k+1}$,
- Stationary distribution (equilibrium)
  - $\pi=\pi P$
  - If the Markov chain is irreducible and aperiodic
  - Computational cost:
    - $O((2^N)^3)$ if we solve the system
    - O(K M) where M is the number of non-null entries in P if we adopt the iterative procedure (K is the number of iterations), in our case $M=O((2^N)^2)$

Can we simplify the problem?

- all the nodes in the same state (infected or susceptibles) are **equivalent**
- If we are interested only in the number of nodes in a given status, we can have a more succinct model
  - state of the system at slot k: I(k)
  - it is still a MC
  - $Prob(I(k+1)=I+n \mid I(k)=I) = C^n_{N-I}\, q_I^n\, (1-q_I)^{N-n-I}$
    - $(I(k+1)-I(k) \mid I(k)=I) \sim Bin(N-I, q_I)$
    - $q_I=1-(1-p)^I$

Any interest for Computer Networks?

- Flooding
  - Epidemic Routing in Delay Tolerant Networks

Delay Tolerant Networks
(a.k.a. Intermittently Connected Networks)

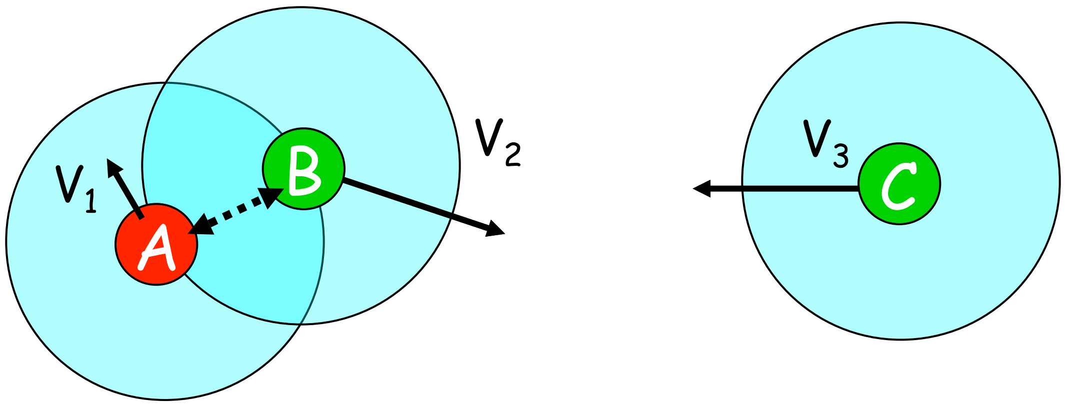

mobile wireless networks

no path at a given time instant between two nodes

- because of power contraint, fast mobility dynamics
- maintain capacity, when number of nodes (N) diverges
  - Fixed wireless networks: $C = \Theta(\mathrm{sqrt}(1/N))$ [Gupta99]
  - Mobile wireless networks: $C = \Theta(1)$, [Grossglauser01]

a really challenging network scenario

- No traditional protocol works

Some examples

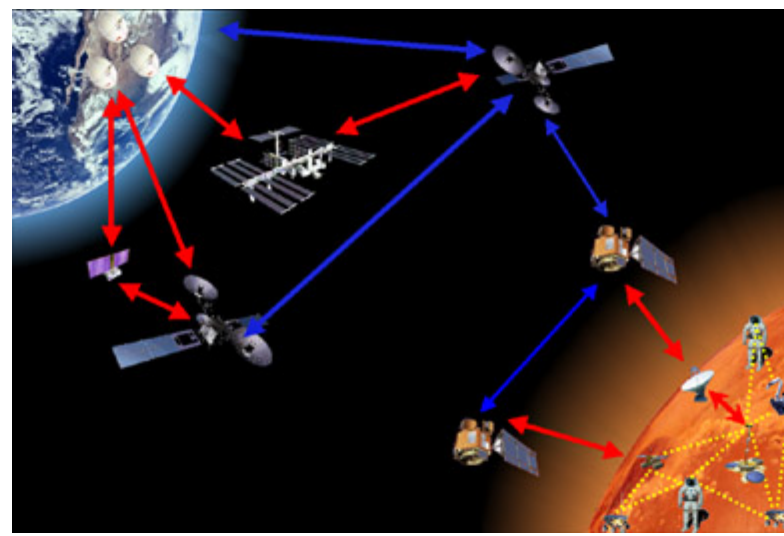

Inter-planetary backbone

- Network for disaster relief team
- Military battle-field network
- …

Epidemic Routing

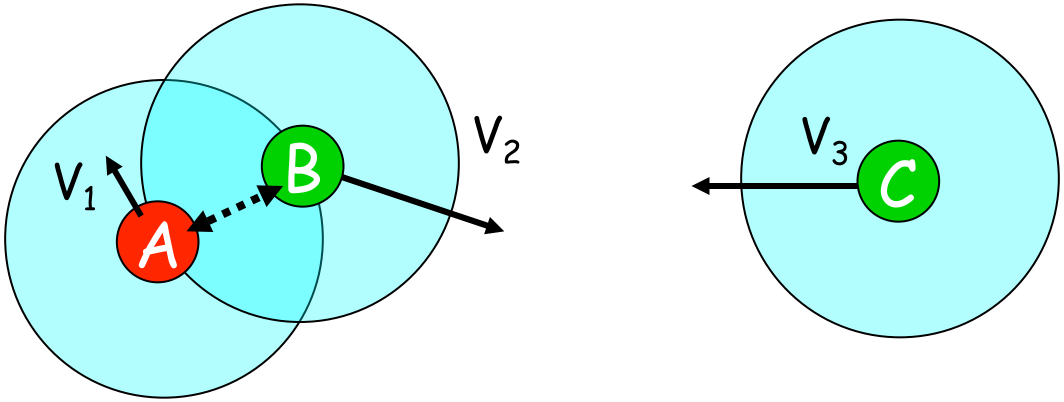

Epidemic Routing

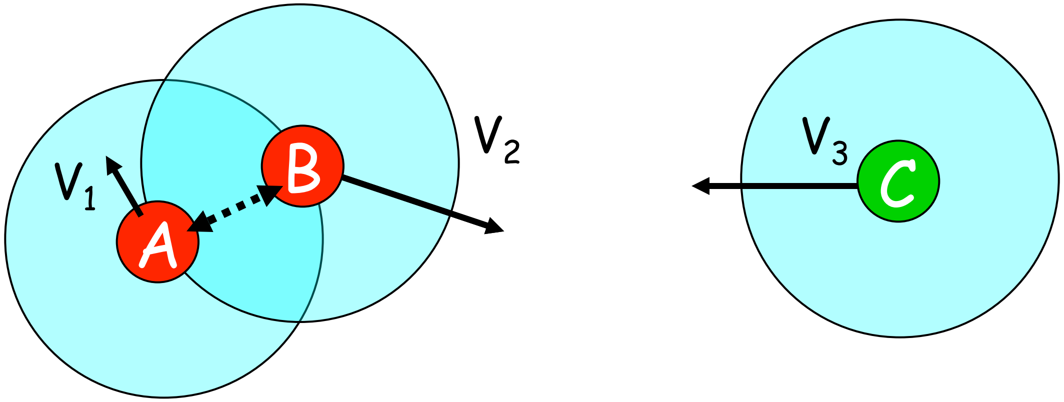

Any interest for Computer Networks?

- Flooding
  - Epidemic Routing in Delay Tolerant Networks
- Chunk distribution in a P2P streaming system (push algorithms)

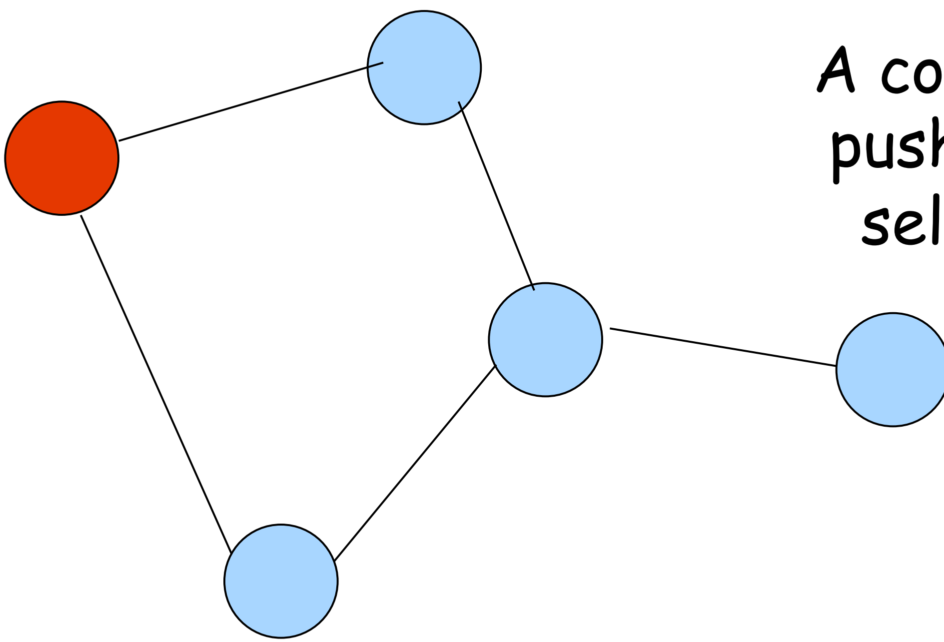

A copy of the chunk is pushed to a randomly selected neighbour

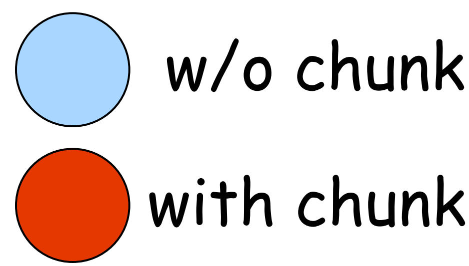

Some numerical examples
$p=10^{-4}$, N=10, I(0)=N/10, 10 runs

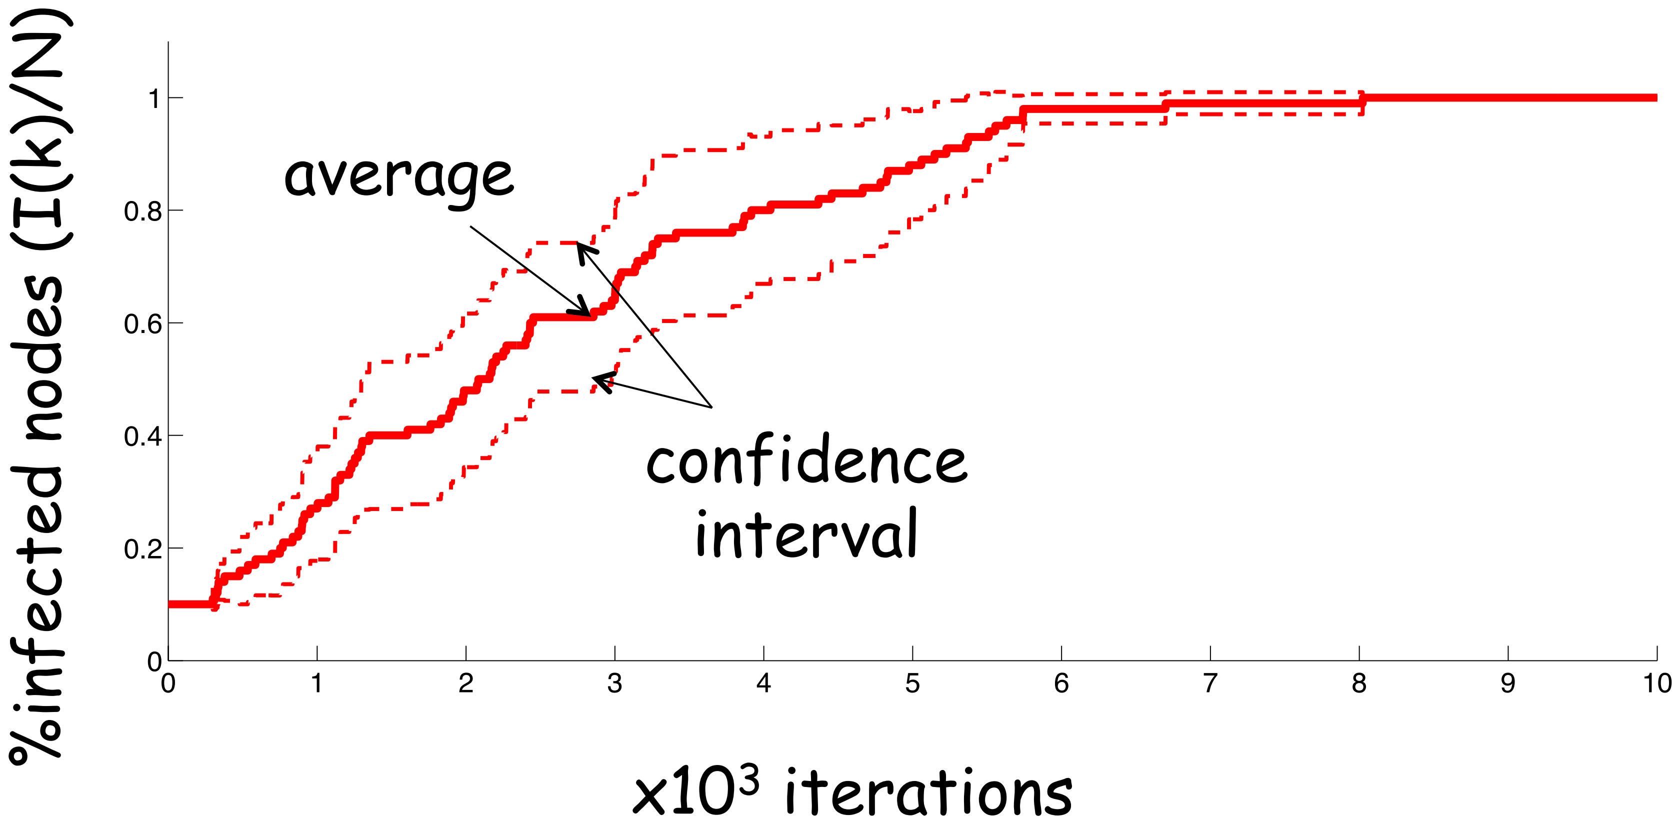

Some numerical examples
$p=10^{-4}$, N=100, I(0)=N/10, 10 runs

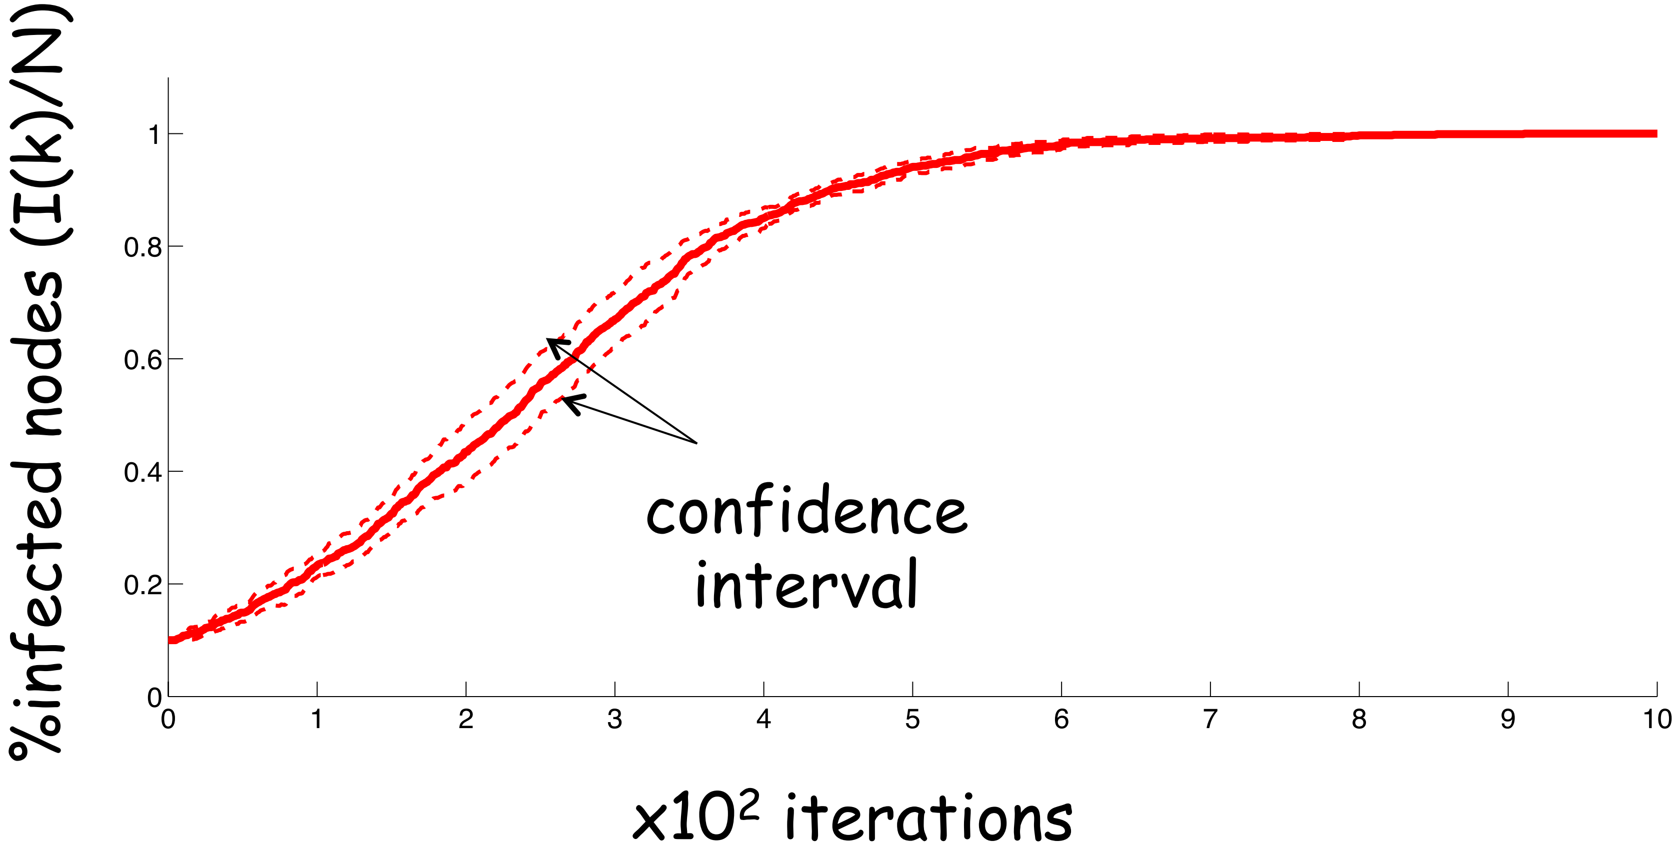

Some numerical examples
$p=10^{-4}$, N=10000, I(0)=N/10, 10 runs

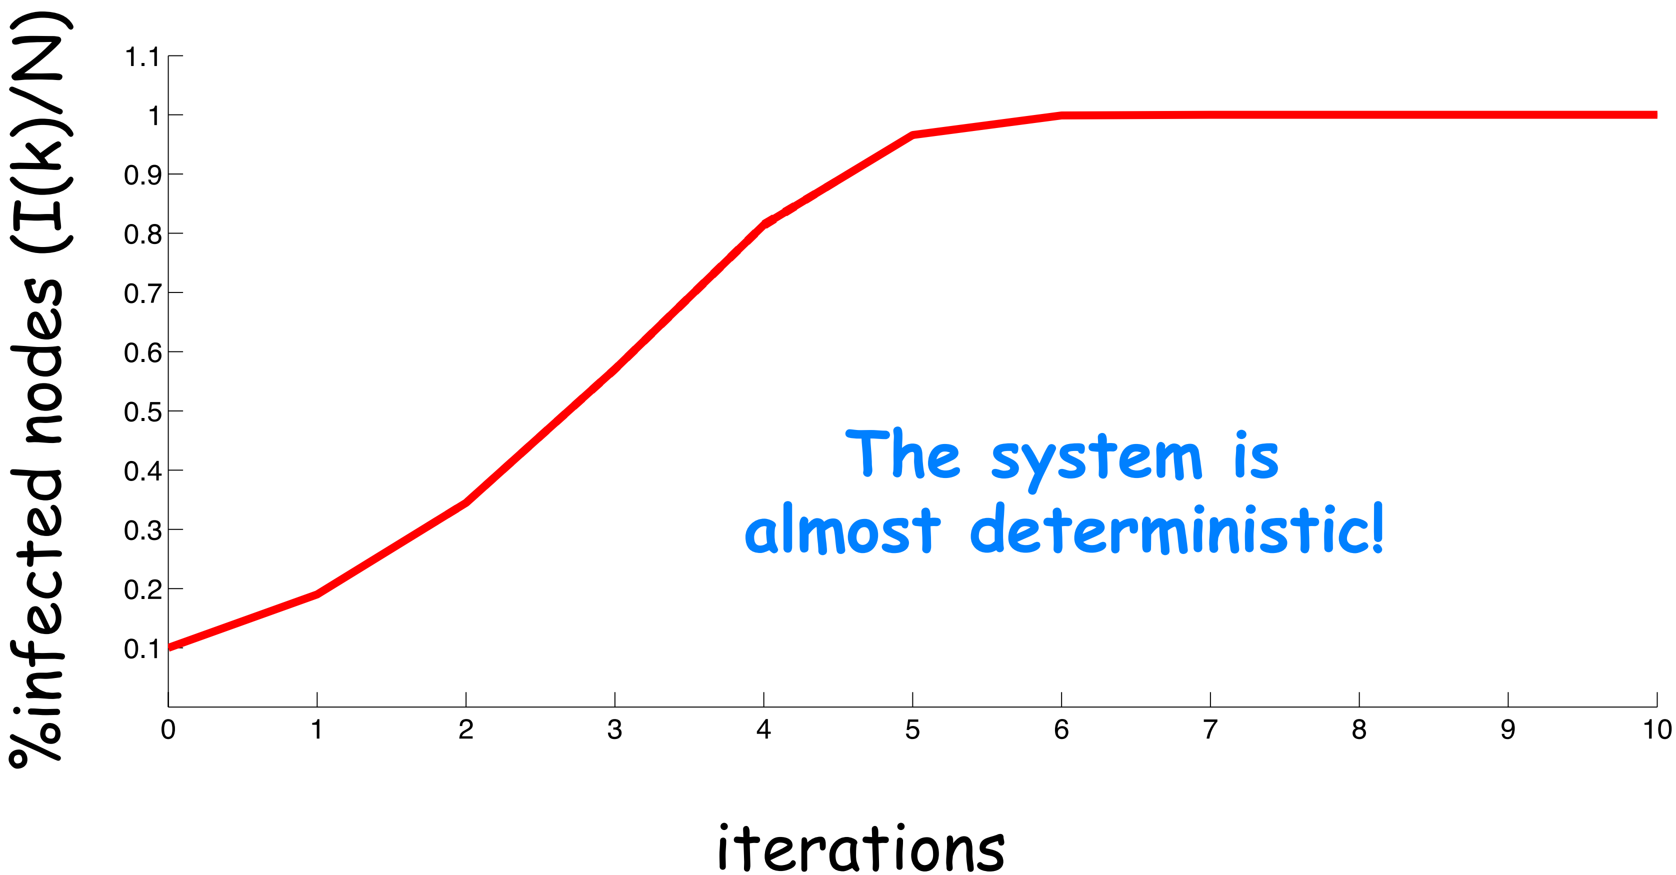

Summary

- For a large system of interacting equivalent objects, the Markov model can be untractable…
- but a deterministic description of the system seems feasible in terms of the empirical measure (% of objects in each status)
  - intuition: kind of law of large numbers
- Mean field models describe the deterministic limit of Markov models when the number of objects diverges

Outline

- Limit of Markovian models
- Mean Field (or Fluid) models
  - exact results
  - extensions
  - applications

References

- Results here for discrete time Markov Chains
  - Benaïm, Le Boudec "A Class of Mean Field Interaction Models for Computer and Communication Systems", LCA-Report-2008-010
- A survey with pointers to continuous time Markov processes and links to stochastic approximation and propagation of chaos
  - Ch. 2 of Nicolas Gast's PhD thesis "Optimization and Control of Large Systems, Fighting the Curse of Dimensionality"

Necessary hypothesis: Objects' Equivalence

- $\pi(k+1)=\pi(k)P$
- A state $\boldsymbol{\sigma} = (v_1, v_2, \ldots v_N)$, $v_j \in V$ ($|V| = V$, finite)
  - E.g. in our example $V=\{0,1\}$
- P is invariant under any label permutation φ:
  - $P_{\sigma,\sigma'} = Prob((v_1, v_2, \ldots v_N) \rightarrow (u_1, u_2, \ldots u_N)) = Prob((v_{\varphi(1)}, v_{\varphi(2)}, \ldots v_{\varphi(N)}) \rightarrow (u_{\varphi(1)}, u_{\varphi(2)}, \ldots u_{\varphi(N)}))$

Some notation and definitions

- $X_n^{(N)}(k)$: state of node n at slot k
- $M_v^{(N)}(k)$: occupancy measure of state v at slot k
  - $M_v^{(N)}(k)=\Sigma_n \mathbf{1}(X_n^{(N)}(k)=v)/N$
  - SI model: $M_2^{(N)}(k)=I^{(N)}(k)/N=i^{(N)}(k)$, $M_1^{(N)}(k)=S^{(N)}(k)/N=s^{(N)}(k)=1-i^{(N)}(k)$
- $\mathbf{M}^{(N)}(k)=(M_1^{(N)}(k),M_2^{(N)}(k),\ldots M_V^{(N)}(k))$
  - SI model: $(1-i^{(N)}(k),i^{(N)}(k))$
- $\mathbf{f}^{(N)}(\mathbf{m})=E[\mathbf{M}^{(N)}(k+1)-\mathbf{M}^{(N)}(k)|\mathbf{M}^{(N)}(k)=\mathbf{m}]$
  - Drift or intensity, it is the **mean field**

Other hypotheses

- Intensity vanishes at a rate $\varepsilon(N)$
  - $\lim_{N\to\infty} \mathbf{f}^{(N)}(\mathbf{m})/\varepsilon(N)=\mathbf{f}(\mathbf{m})$
- Second moment of number of object transitions per slot is bounded
  - $\#\text{transitions} < W^N(k)$,
    $E[W^N(k)^2 | \mathbf{M}^{(N)}(k)=\mathbf{m}] < cN^2\varepsilon(N)^2$
- Drift is a smooth function of $\mathbf{m}$ and $1/N$
  - $\mathbf{f}^{(N)}(\mathbf{m})/\varepsilon(N)$ has continuous derivatives in $\mathbf{m}$ and in $1/N$

Convergence Result

- Define $\underline{\mathbf{M}}^{(N)}(t)$ with t real, such that
  - $\underline{\mathbf{M}}^{(N)}(k\,\varepsilon(N))=\mathbf{M}^{(N)}(k)$ for k integer
  - $\underline{\mathbf{M}}^{(N)}(t)$ is affine on $[k\,\varepsilon(N),(k+1)\varepsilon(N)]$
- Consider the Differential Equation
  - $d\boldsymbol{\mu}(t)/dt=\mathbf{f}(\boldsymbol{\mu})$, with $\boldsymbol{\mu}(0)=\mathbf{m}_0$
- Theorem
  - For all T>0, if $\mathbf{M}^{(N)}(0) \rightarrow \mathbf{m}_0$ in probability (/mean square) as $N \rightarrow \infty$, then $\sup_{0\leq t\leq T}||\underline{\mathbf{M}}^{(N)}(t)-\boldsymbol{\mu}(t)|| \rightarrow 0$ in probability (/ mean square)

Convergence of random variables

- The sequence of random variables $X^{(N)}$ converges to X in probability if
  - for all $\delta>0$ $\mathrm{Lim}_{N\rightarrow\infty}\mathrm{Prob}(|X^{(N)} - X|>\delta)=0$
- The sequence of random variables X^N converges to X in mean square if
  - $\mathrm{Lim}_{N\rightarrow\infty}E[|X^{(N)} - X|^2]=0$
- Convergence in mean square implies convergence in probability

Application to the SI model

- Assumptions' check
  - ✓ Nodes are equivalent
  - – Intensity vanishes at a rate $\varepsilon(N)$

    $\mathbf{f}^{(N)}(\mathbf{m})=E[\mathbf{M}^{(N)}(k+1)-\mathbf{M}^{(N)}(k)|\mathbf{M}^{(N)}(k)=\mathbf{m}]$

    $M_2^{(N)}(k)=I^{(N)}(k)/N=i^{(N)}(k), M_1^{(N)}(k)=1-M_2^{(N)}(k)$

    $(I^{(N)}(k+1)-I^{(N)}(k) \mid I^{(N)}(k)=I) \sim Bin(N-I, q_I) \Rightarrow$

    $E[I^{(N)}(k+1)-I^{(N)}(k) \mid I^{(N)}(k)=I] = q_I (N-I)$

    $E[i^{(N)}(k+1)-i^{(N)}(k)|i^{(N)}(k)=i] = (1-i)\, q_I$

    $= (1-i)(1-(1-p)^{i\,N}) \to (1-i)$ when N diverges!

Application to the SI model

- Out of the impasse: introduce a scaling for p
  - If $p^{(N)}=p_0/N^a$ $a>1$ => $(1-i)(1-(1-p^{(N)})^{i\,N})\to 0$
  - Consider a=2
    - $(1-i)(1-(1-p^{(N)})^{i\,N}) \sim (1-i)\, i\, p_0/N$ (for N large)
  - $\varepsilon(N)=p_0/N$
  - $f_2(\mathbf{m}) = f_2((s,i)) = s\, i = i\,(1-i)$
- Lesson to keep: often we need to introduce some parameter scaling

Application to the SI model

- Assumptions' check
  - ✓ Nodes are equivalent
  - ✓ Intensity vanishes at a rate $\varepsilon(N)=p_0/N$
  - – Second moment of number of object transitions per slot is bounded

    $\#transitions < W^N(k)$,

    $E[W^N(k)^2 | \mathbf{M}^{(N)}(k)=\mathbf{m}] < cN^2\varepsilon(N)^2$

    $W^N(k)=\#trans. \sim Bin(N-I(k),q_I)$

    $E[W^N(k)^2]=((N-I(k))q_I)^2 + (N-I(k))q_I(1-q_I)$
    is in $O(N^2\,\varepsilon(N)^2)$

Application to the SI model

- Assumptions' check
  - ✓ Nodes are equivalent
  - ✓ Intensity vanishes at a rate $\varepsilon(N)=p_0/N$
  - ✓ Second moment of number of object transitions per slot is bounded
  - ✓ Drift is a smooth function of **m** and 1/N
    - $f_2^{(N)}(\mathbf{m}) = (1-i)(1-(1-p^{(N)})^{i\,N})$
      $= (1-i)\,(1 - (\Sigma_{n=0,\ldots N}\, C^n_N\, (p_0/N^2)^n)^i)$
    - continuous derivatives in i and in 1/N

Practical use of the convergence result

- Theorem
  - For all T>0, if $\mathbf{M}^{(N)}(0) \rightarrow \mathbf{m}_0$ in probability (/mean square) as $N \rightarrow \infty$, then $\sup_{0 \le \tau \le T} || \underline{\mathbf{M}}^{(N)}(t) - \boldsymbol{\mu}(t) || \rightarrow 0$ in probability (/ mean square)
  - Where $\boldsymbol{\mu}(t)$ is the solution of

    $d\boldsymbol{\mu}(t)/dt = \mathbf{f}(\boldsymbol{\mu})$, with $\boldsymbol{\mu}(0) = \mathbf{m}_0$

- $\mathbf{M}^{(N)}(0) = \mathbf{m}_0$, $\mathbf{M}^{(N)}(k) = \underline{\mathbf{M}}^{(N)}(k\varepsilon(N)) \approx \boldsymbol{\mu}(k\varepsilon(N))$

Application to the SI model

- $f_2(\mathbf{m})=f_2((s,i))=i(1-i)$
- $d\mu_2(t)/dt=f_2(\mu_2(t))=\mu_2(t)(1-\mu_2(t))$, with $\mu_2(0)=\mu_{0,2}$
  - Solution: $\mu_2(t)=1/((1/\mu_{0,2}-1)\, e^{-t}+1)$

- If $i^{(N)}(0)=i_0$,
  $i^{(N)}(k) \approx \mu_2(k\varepsilon(N))=1/((1/i_0-1)\exp(-k\, p_0/N)+1)$
  $=1/((1/i_0-1)\exp(-k\, N\, p)+1)$

Back to the numerical examples
$p=10^{-4}$, $I(0)=N/10$, 10 runs

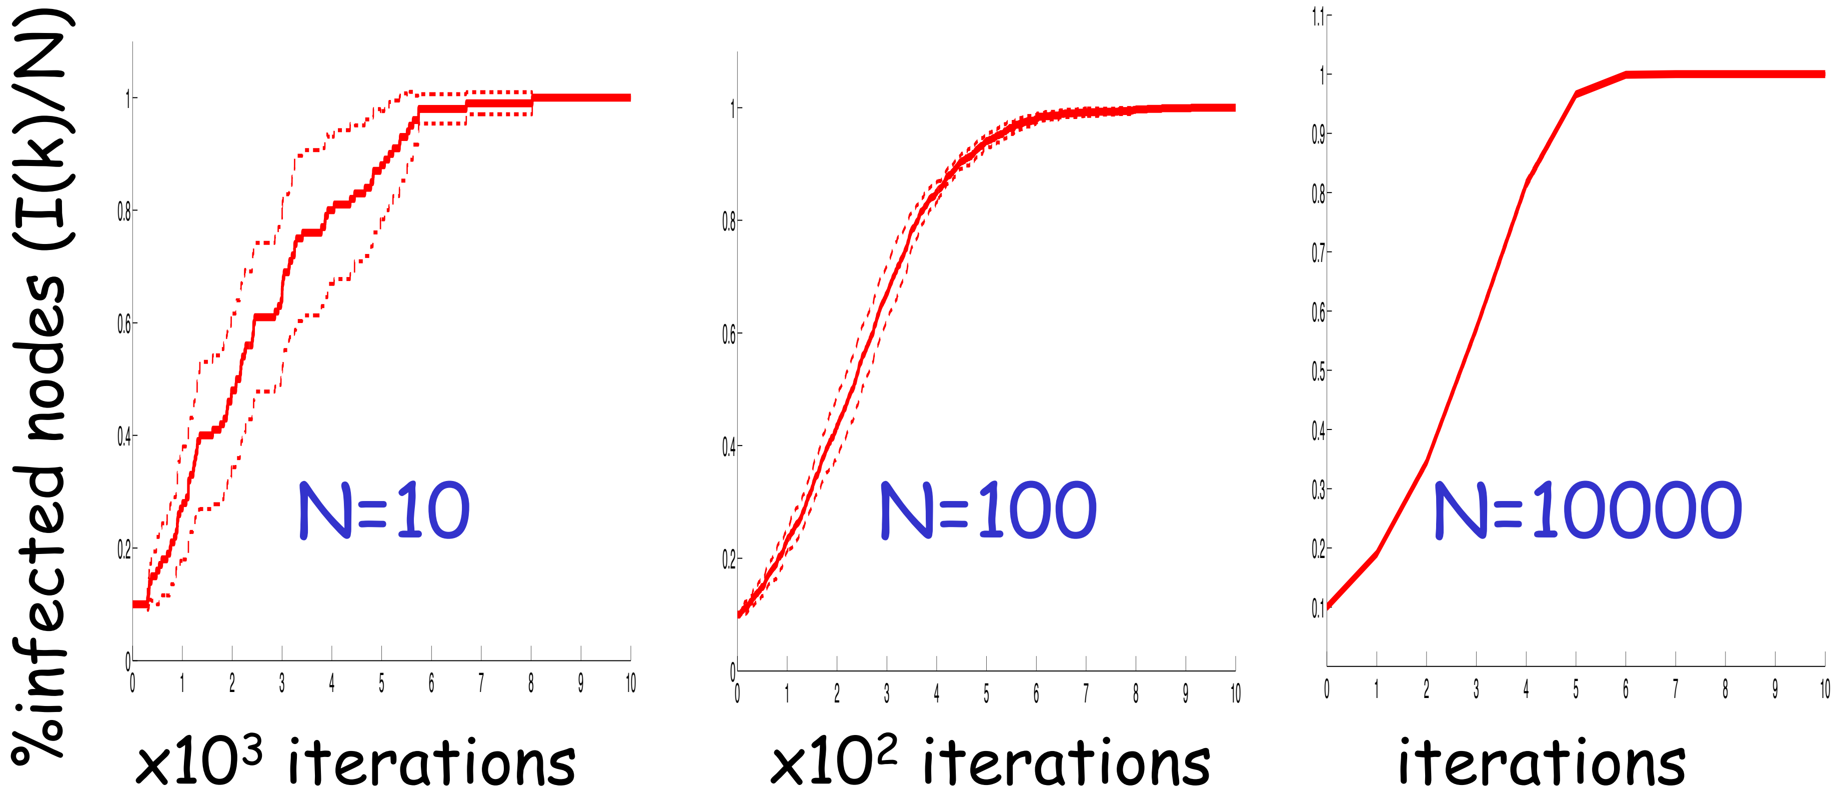

Advantage of Mean Field

- If $i^{(N)}(0)=i_0$,
  $i^{(N)}(k) \approx \mu_2 (k\varepsilon(N))=1/((1/i_0-1) \exp(-k\, p_0/N)+1)$
  $=1/((1/i_0-1) \exp(-k\, N\, p)+1)$
  - solved for each N with negligible computational cost
- In general: solve numerically the solution of a system of ordinary differential equations (size = #of possible status)
  - simpler than solving the Markov chain

SIS model

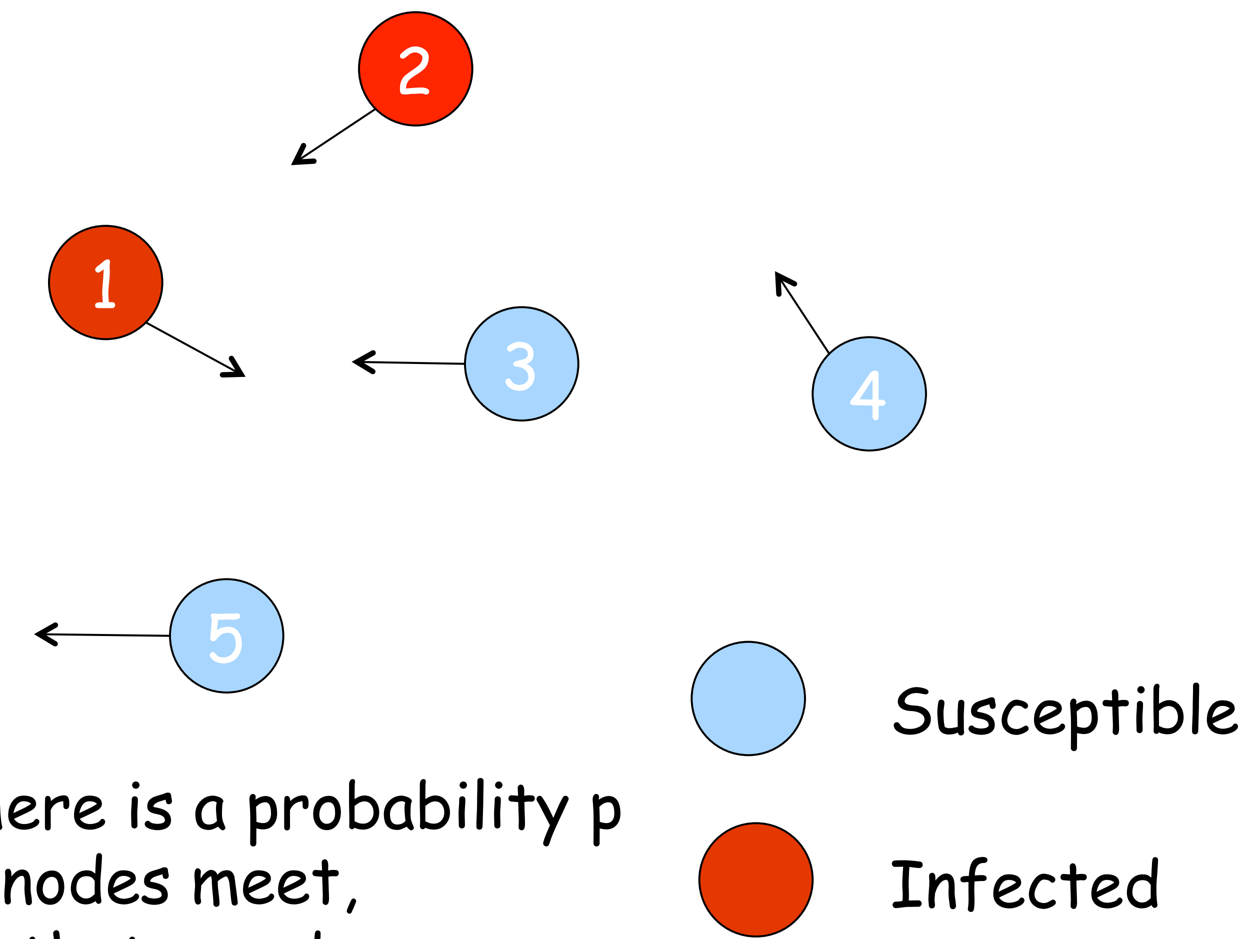

At each slot there is a probability p
that two given nodes meet,
a probability r that a node recover.

SIS model

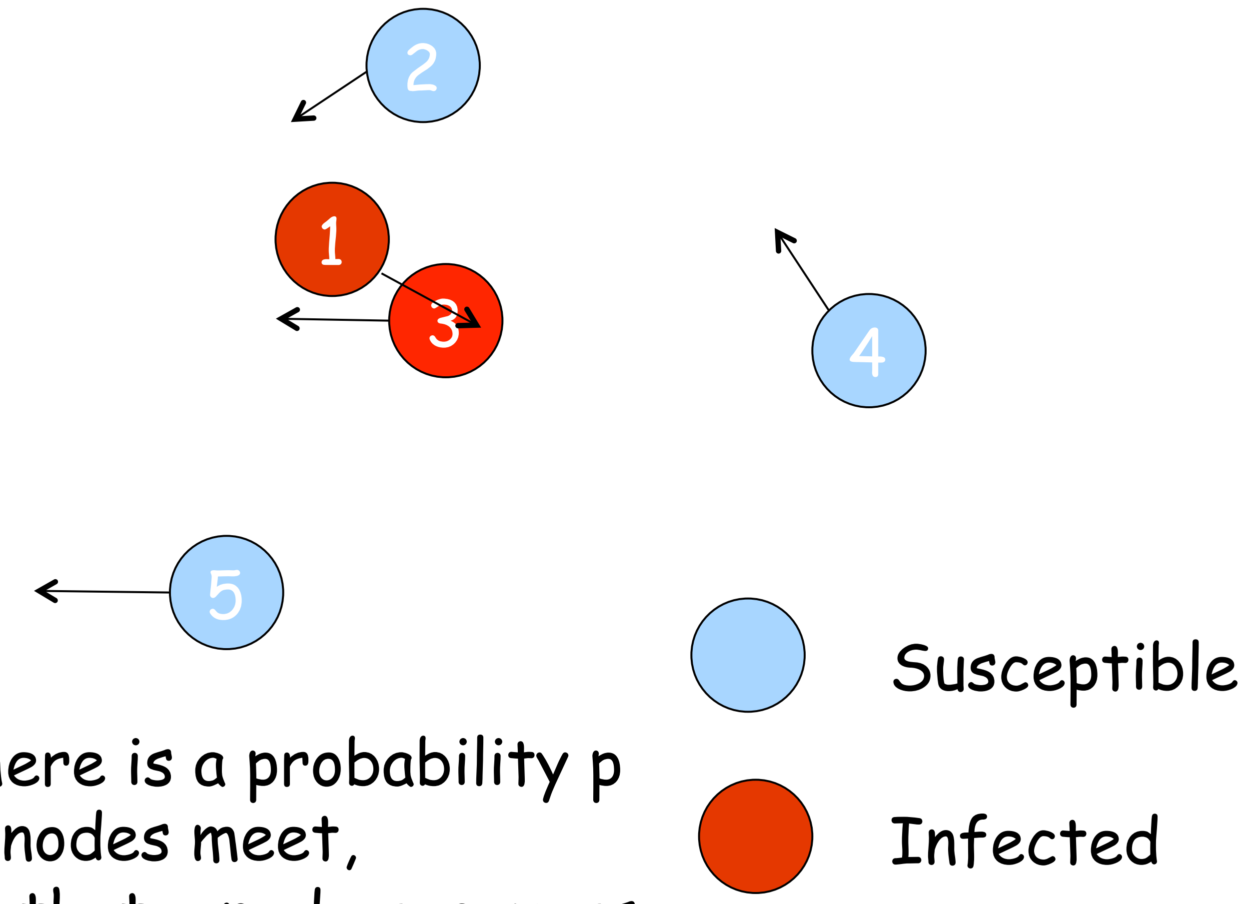

At each slot there is a probability p that two given nodes meet,
a probability r that a node recovers.

Let's practise

- Can we propose a Markov Model for SIS?
  - No need to calculate the transition matrix
- If it is possible, derive a Mean Field model for SIS
  - Do we need some scaling?

Study of the SIS model

- We need $p^{(N)}=p_0/N^2$ and $r^{(N)}=r_0/N$
- If we choose $\varepsilon(N)=1/N$, we get
  - $di(t)/dt = p_0\, i(t)(1-i(t)) - r_0\, i(t)$

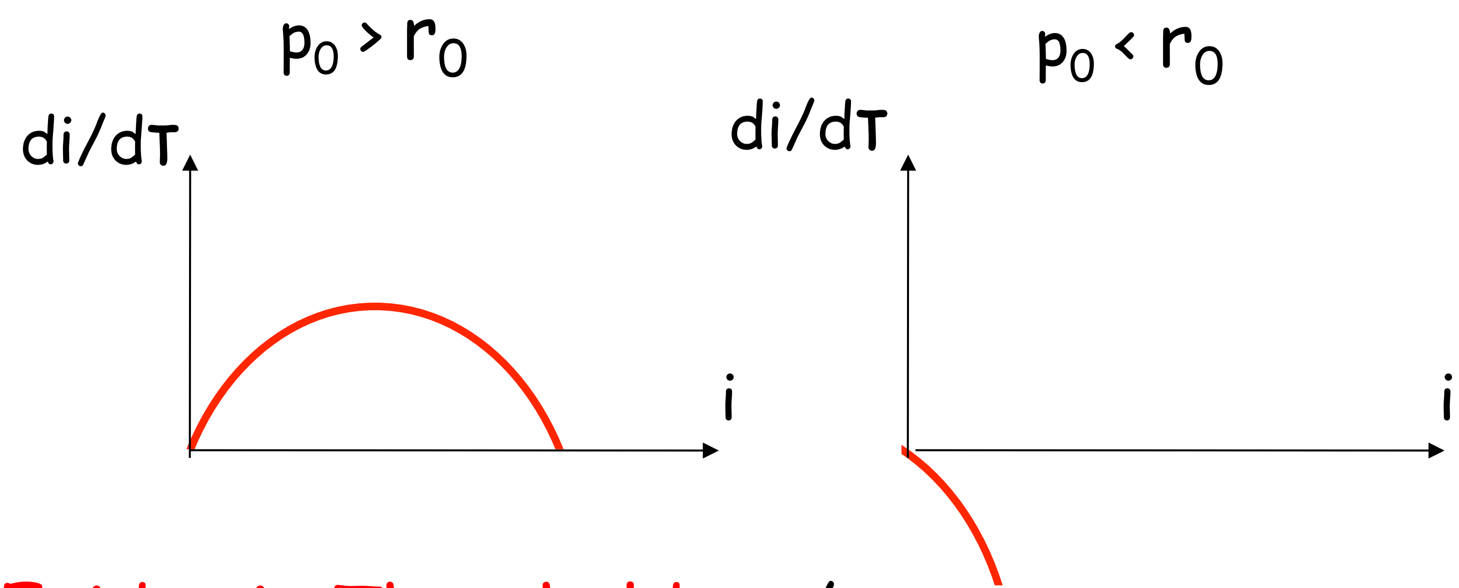

Epidemic Threshold: p_0/r_0

$N=80, p_0=0.1$

$r_0 = 0.05$

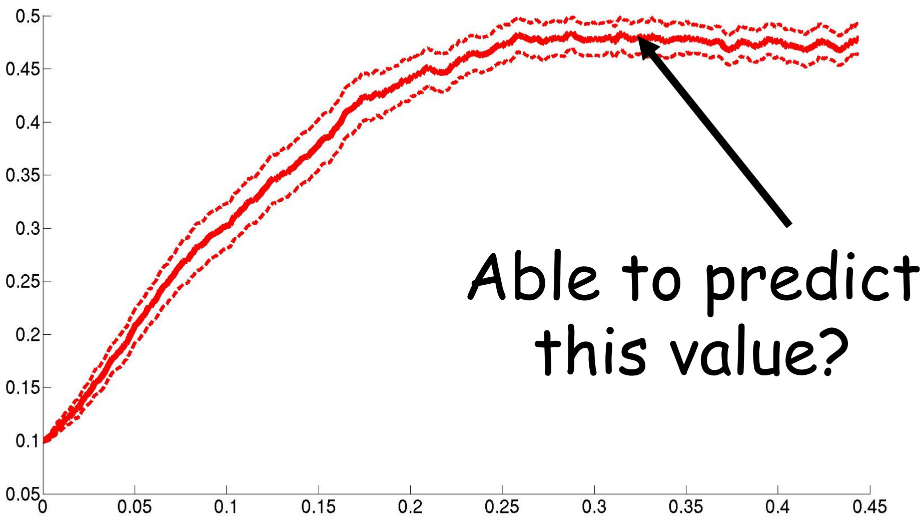

$r_0 = 0.125$

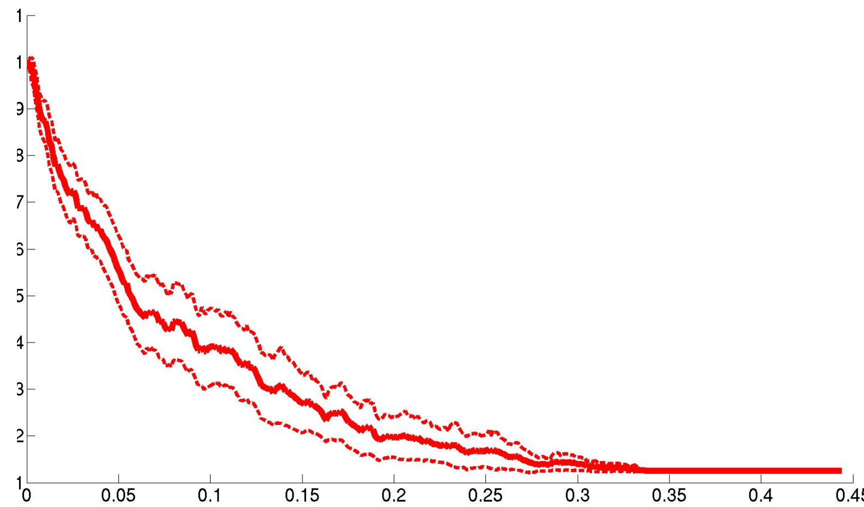

Study of the SIS model

- $d\mu_2(t)/dt = p_0\, \mu_2(t)(1-\mu_2(t)) - r_0\, \mu_2(t)$
- Equilibria, $d\mu_2(t)/dt=0$
  - $\mu_2(\infty)=1-r_0/p_0$ or $\mu_2(\infty)=0$

There is more: Independence

- Theorem 2
  - Under the assumptions of Theorem 1, and that the collection of objects at time 0 is exchangeable $(X_1^N(0), X_2^N(0), \ldots X_N^N(0))$,

  then for any fixed n and t:

  $$\lim_{N\to\infty} \text{Prob}(\underline{X}_1^N(t)=i_1, \underline{X}_2^N(t)=i_2, \ldots \underline{X}_n^N(t)=i_n)=$$
  $$=\mu_{i1}(t)\mu_{i2}(t)\ldots\mu_{in}(t)$$
- MF Independence Property, a.k.a. Decoupling Property, Propagation of Chaos

Remarks

- $(X_1^N(0), X_2^N(0), \dots X_N^N(0))$ exchangeable
  - Means that all the states that have the same occupancy measure $\mathbf{m}_0$ have the same probability
- $\lim_{N\to\infty} \mathrm{Prob}(\underline{X}_1^N(t)=i_1, \underline{X}_2^N(t)=i_2, \dots \underline{X}_n^N(t)=i_n)=$

  $$=\mu_{i1}(t)\mu_{i2}(t)\dots\mu_{in}(t)$$

  - Application

    $\mathrm{Prob}(X_1^N(k)=i_1, X_2^N(k)=i_2, \dots X_k^N(k)=i_k) \approx$

    $$\approx \mu_{i1}(k\varepsilon(N))\mu_{i2}(k\varepsilon(N))\dots\mu_{ik}(k\varepsilon(N))$$

Probabilistic interpretation of the occupancy measure
(SI model with $p=10^{-4}$, $N=100$)

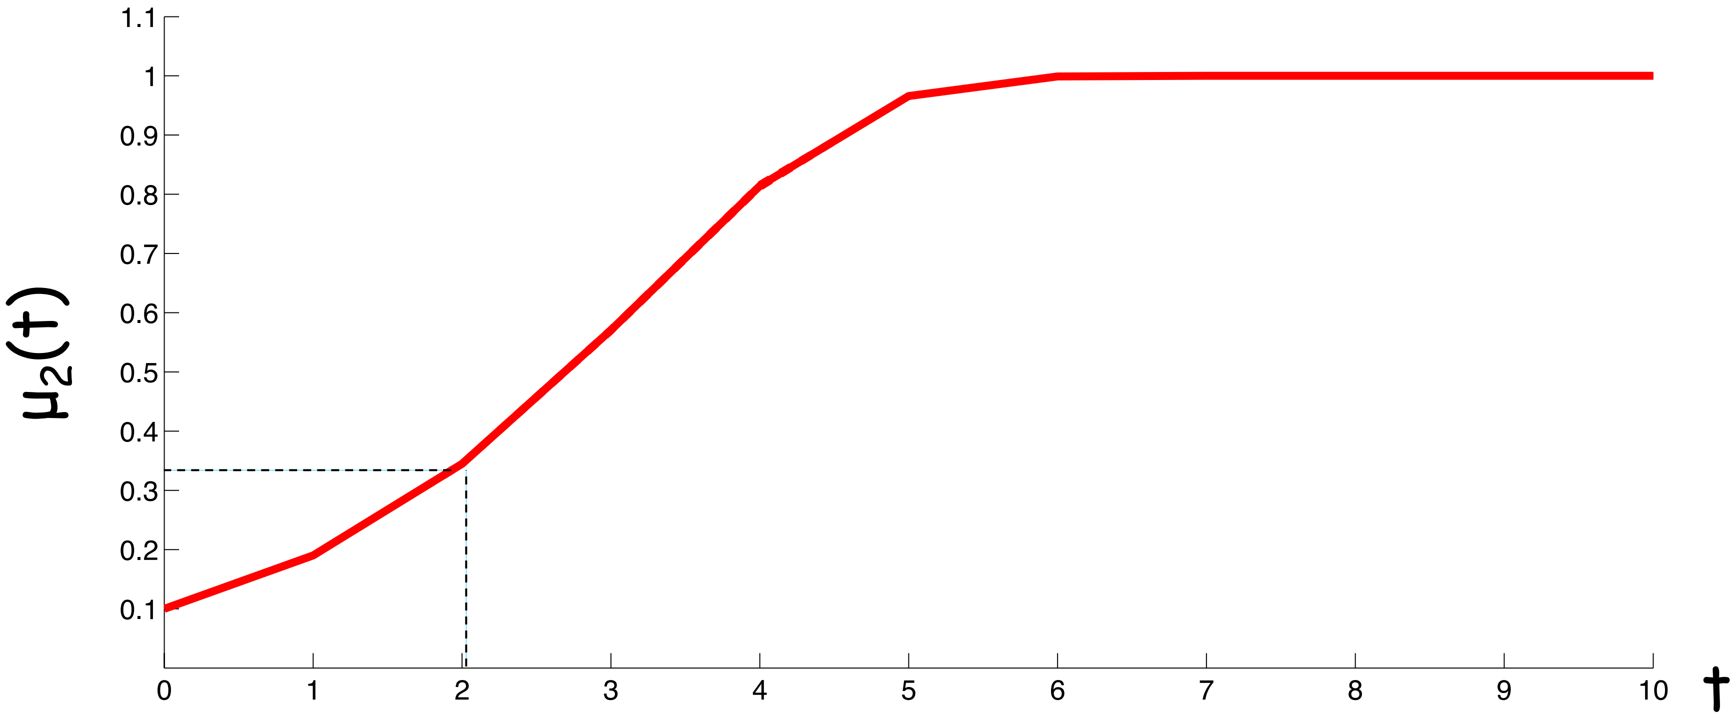

Prob(nodes 1,17,21 and 44 infected at k=200)=

$=\mu_2(k\ p\ N)^4=\mu_2(2)^4\approx(1/3)^4$

What if 1,17,21 and 44 are surely infected at k=0