Performance Evaluation

Lecture 2: Epidemics

Giovanni Neglia
INRIA – EPI Maestro
16 December 2013

Epidemics on a graph: SI model

One of the neighbours
is randomly selected
and infected

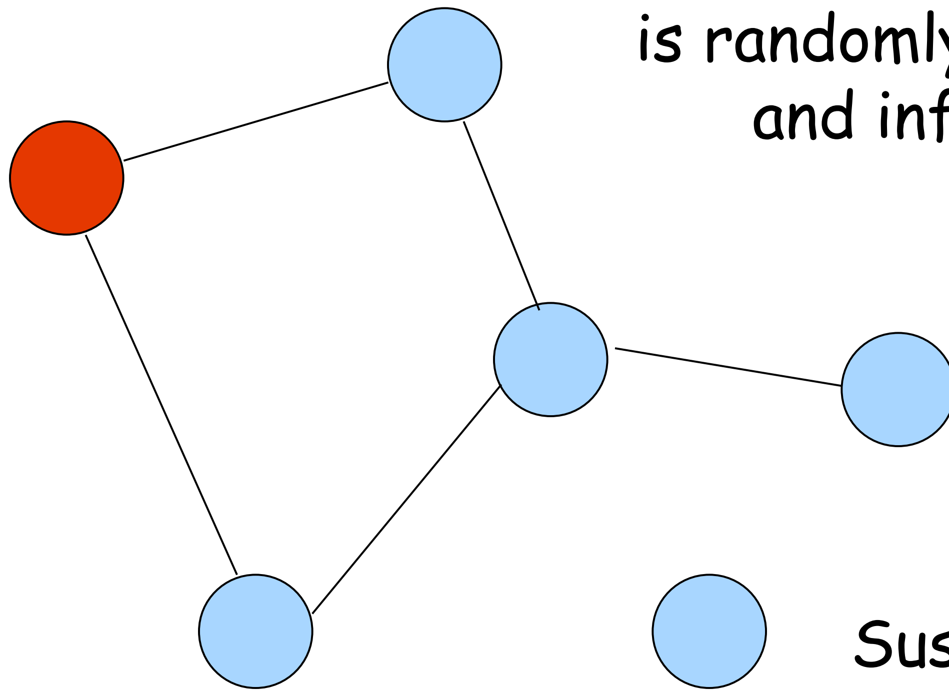

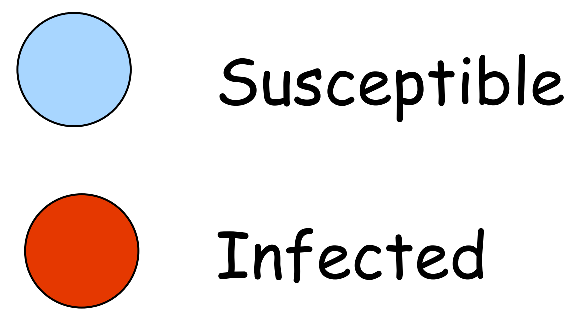

Epidemics on a graph: SI model

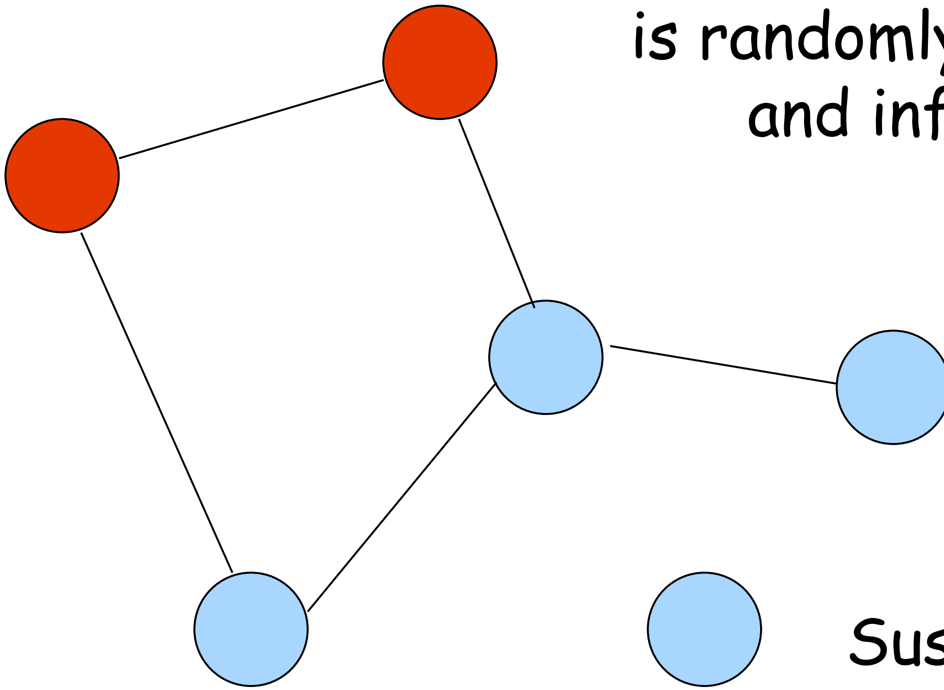

One of the neighbours
is randomly selected
and infected

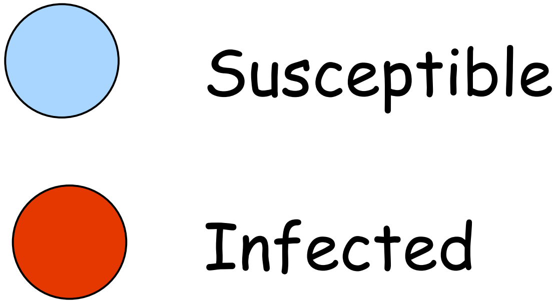

Any interest for Computer Networks?

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

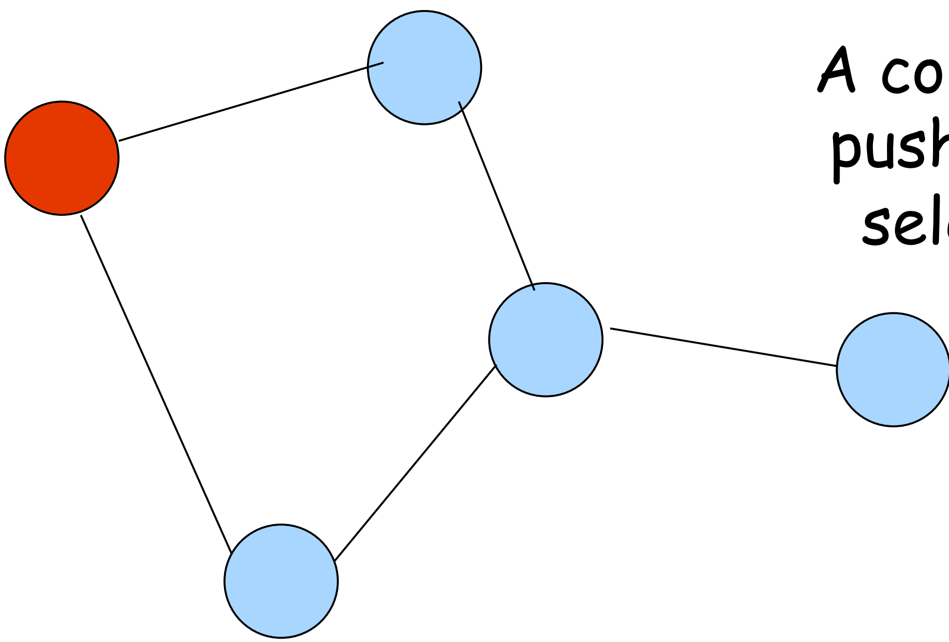

A copy of the chunk is pushed to a randomly selected neighbour

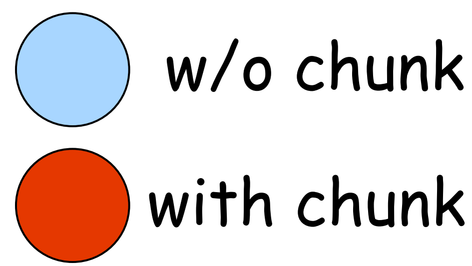

Time-slotted synchronous behaviour

Every T an infected node transmits the disease to a randomly selected neighbour

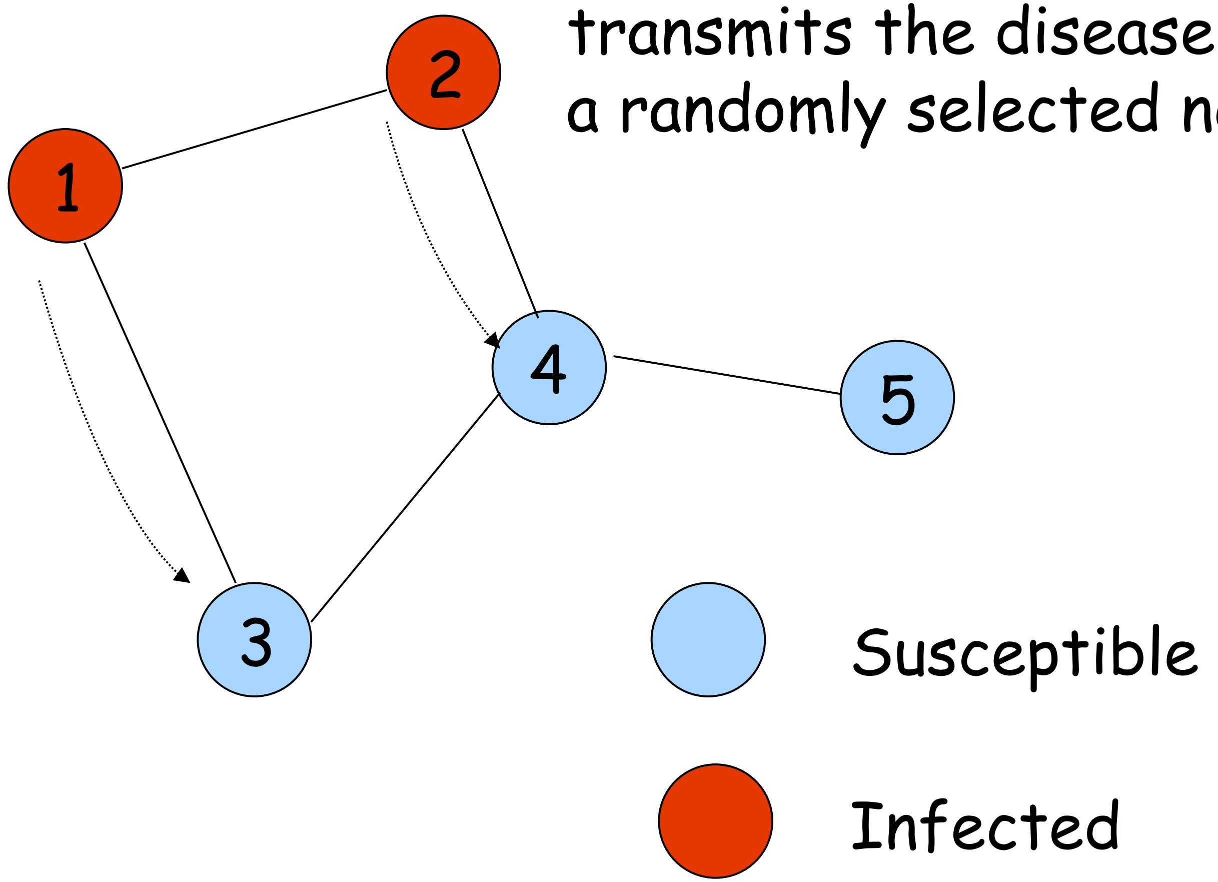

Time-slotted synchronous behaviour

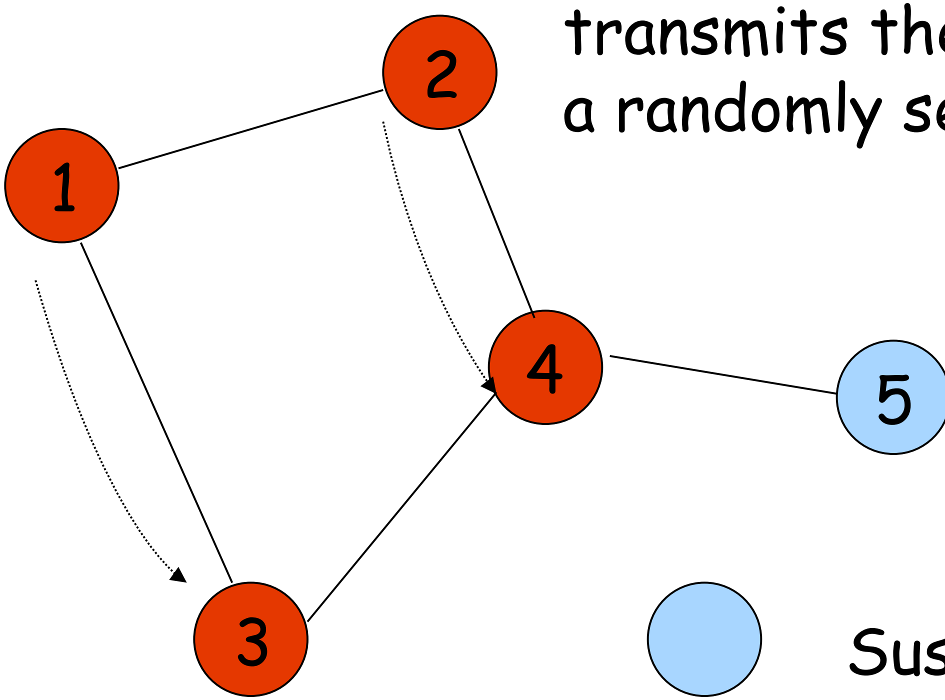

Every T an infected node transmits the disease to a randomly selected neighbour

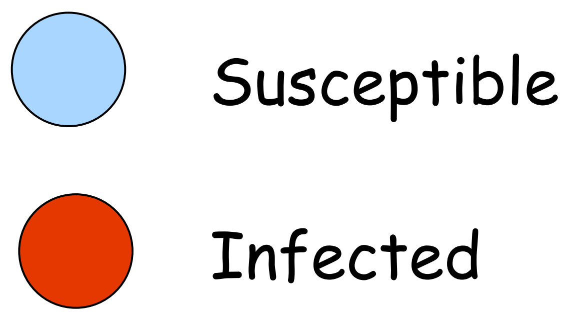

Time-slotted synchronous behaviour

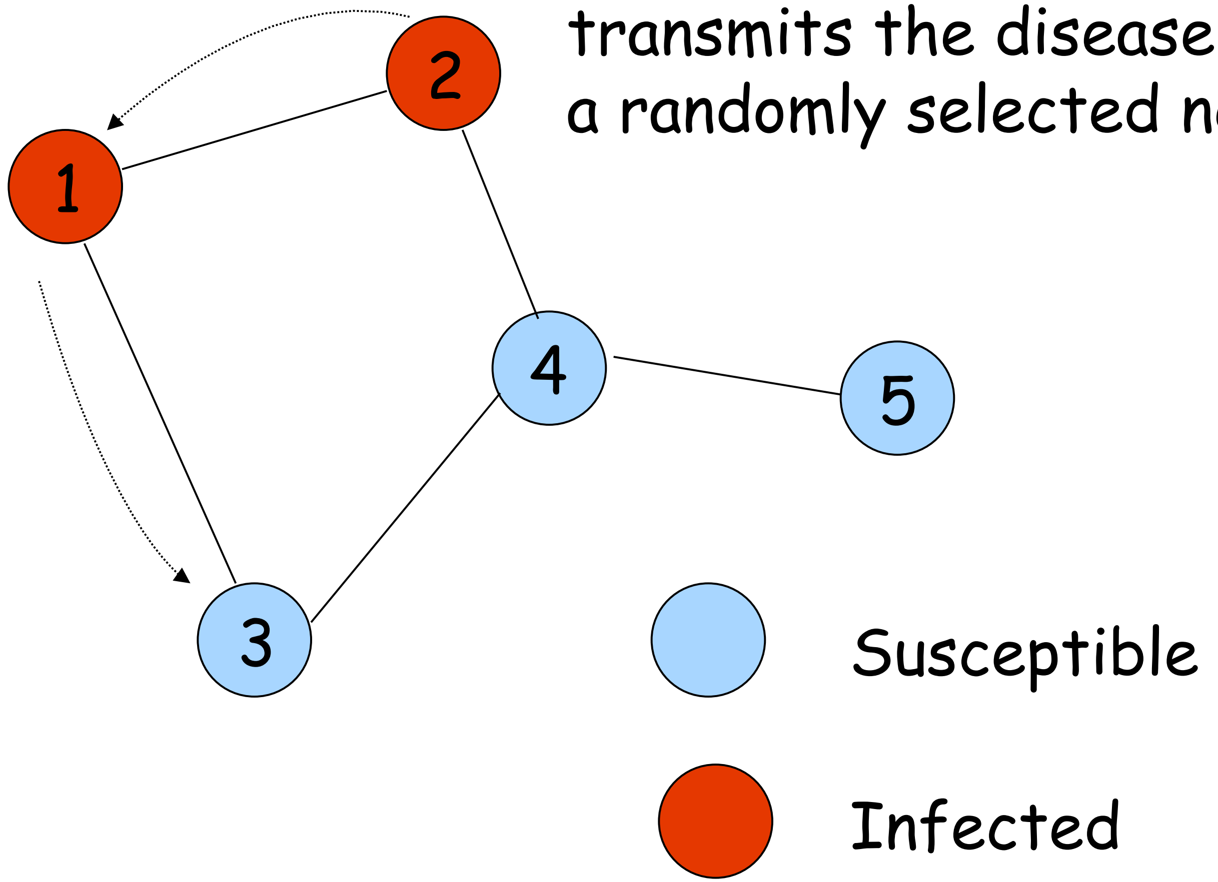

Every T an infected node transmits the disease to a randomly selected neighbour

Time-slotted synchronous behaviour

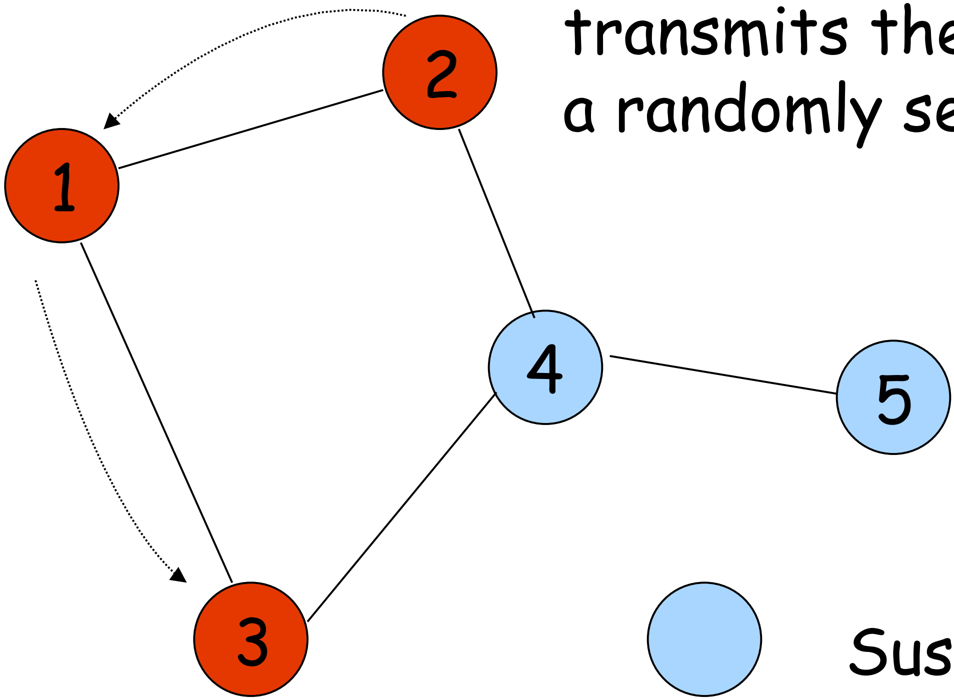

Every T an infected node transmits the disease to a randomly selected neighbour

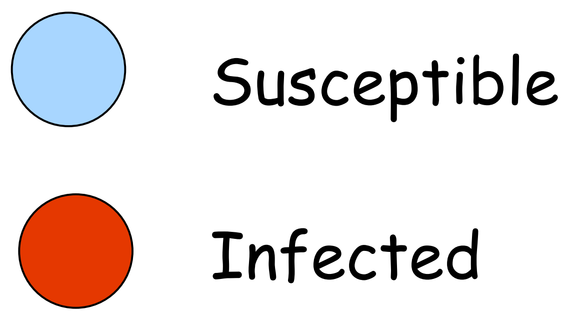

How do you model it?

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

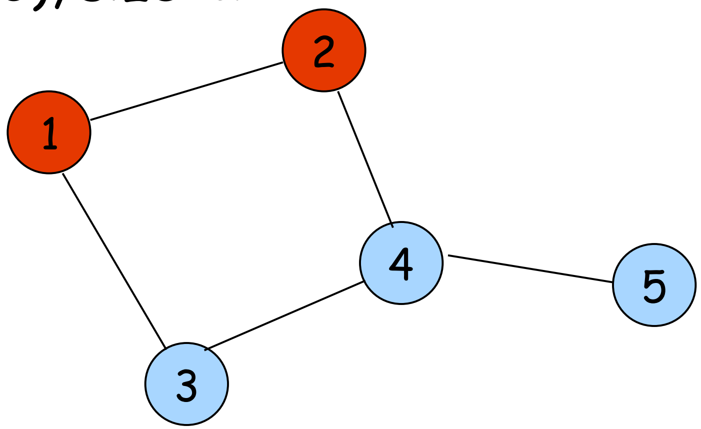

  - Probability transitions among states
    - e.g. Prob((1,1,0,0,0)->(1,1,1,0,0))=1/4

Asynchronous behaviour

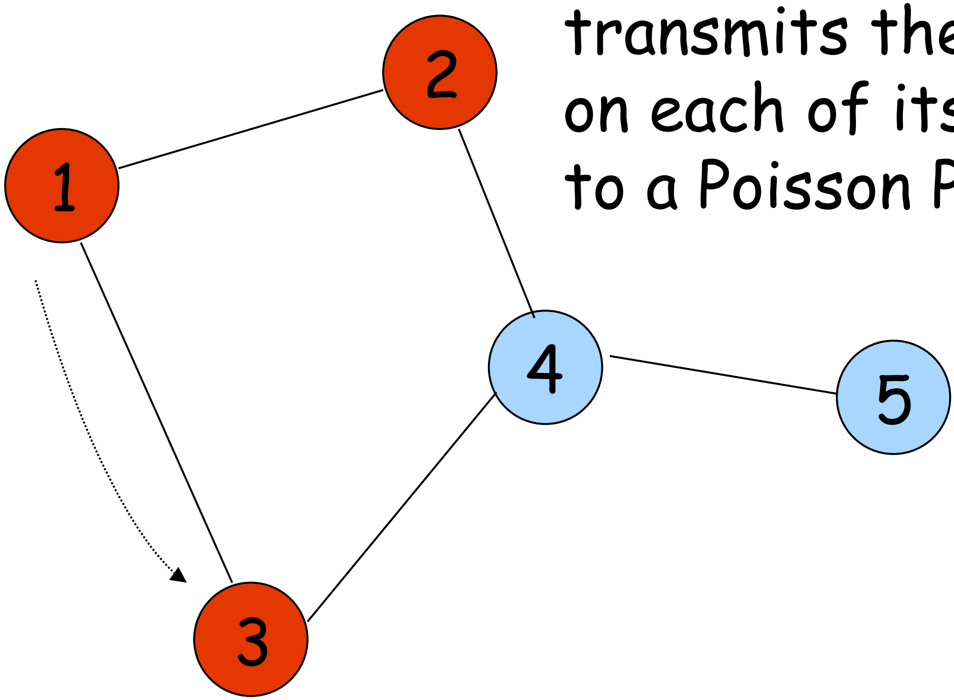

Every infected node
transmits the disease
on each of its links according
to a Poisson Process with rate β

How to model it?

- A Continuous-time Markov process/Chain (C-MC)
  - System state at time t is a vector specifying if every node is infected (1) or not (0)
  - Rate transitions between state pairs
    - e.g, q((1,1,1,0,0)->(1,1,1,1,0))=2β

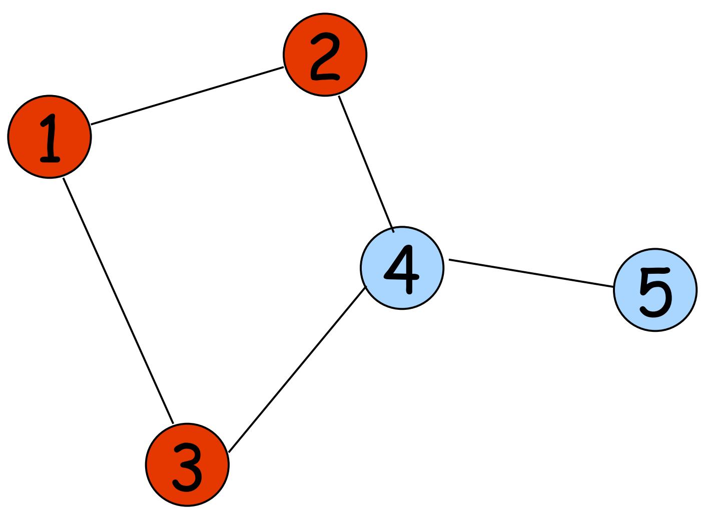

What to study and how

- P the transition matrix ($2^N x 2^N$)
- 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
    - $O(K\ A)$ where A is the number of non-null entries in P if we adopt the iterative procedure (K is the number of iterations), in our case $A \in O(|E|)$

Similar for C-MC

- Stationary distribution (equilibrium)
  - $\pi = \pi P$, D-MC
  - $\pi Q = 0$, C-MC

- Transient analysis
  - $\pi(k+1) = \pi(k)P$, D-MC
  - $d\pi(t)/dt = \pi(t)Q$, C-MC

Outline

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

A motivating example: epidemics

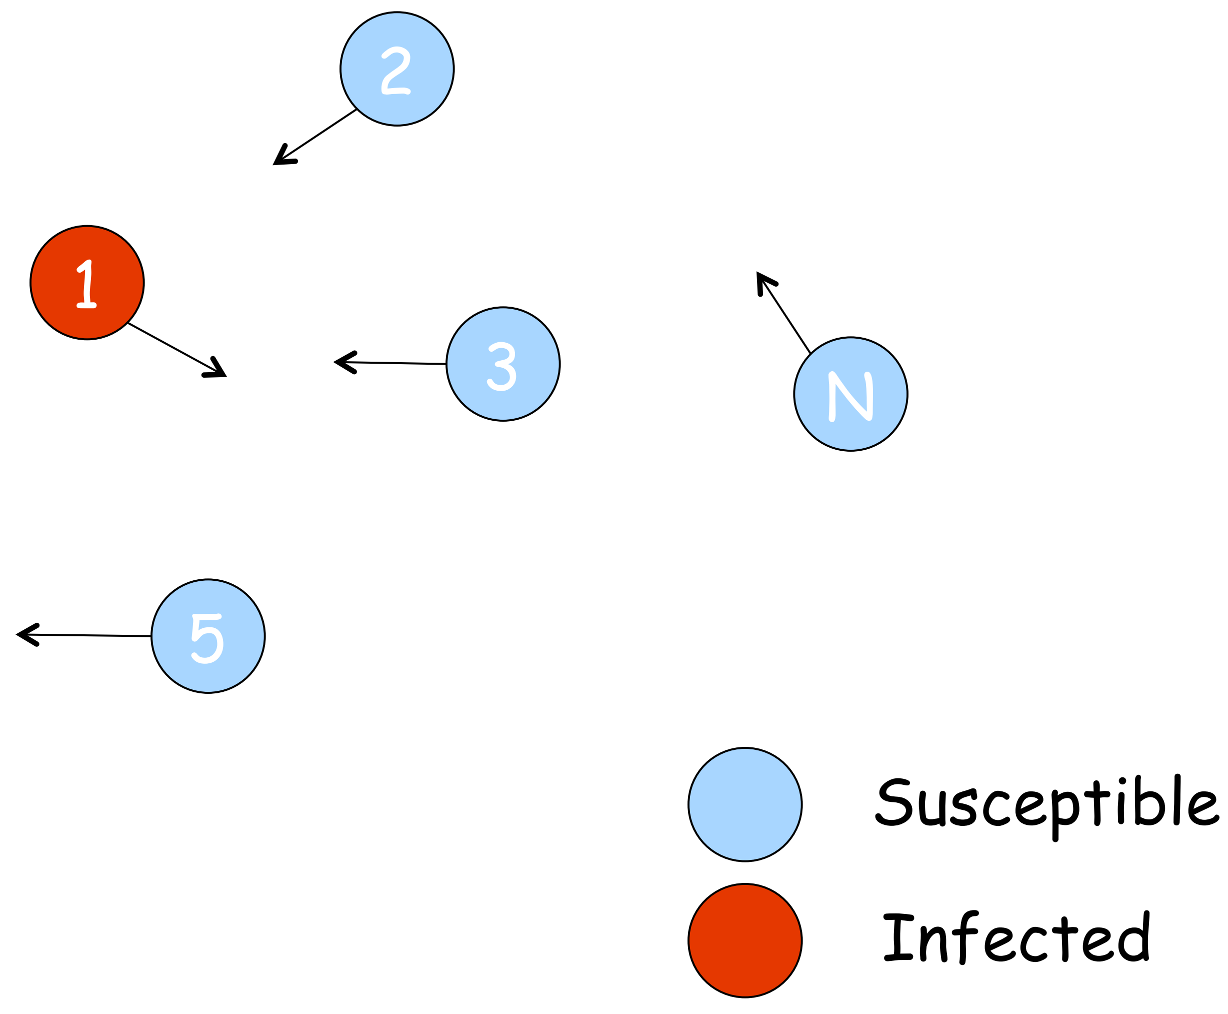

A motivating example: epidemics

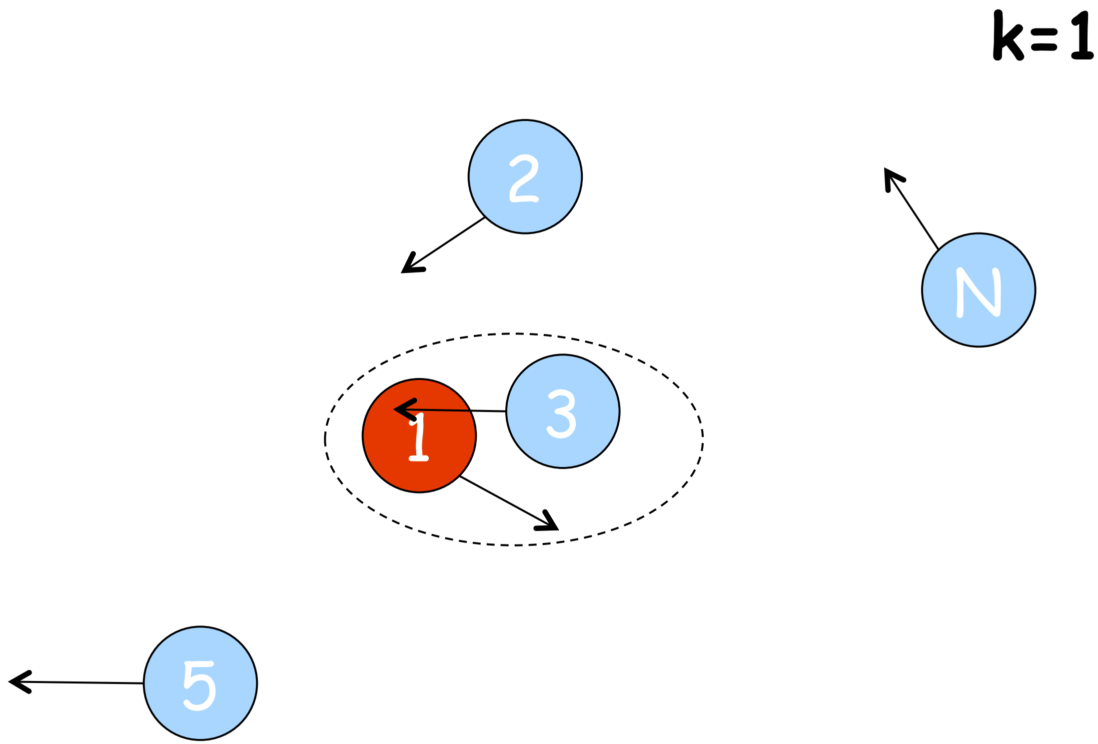

At each slot there is a probability p that two given nodes meet.
Assume meetings to be independent.

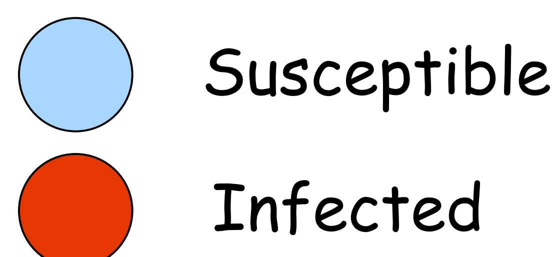

A motivating example: epidemics

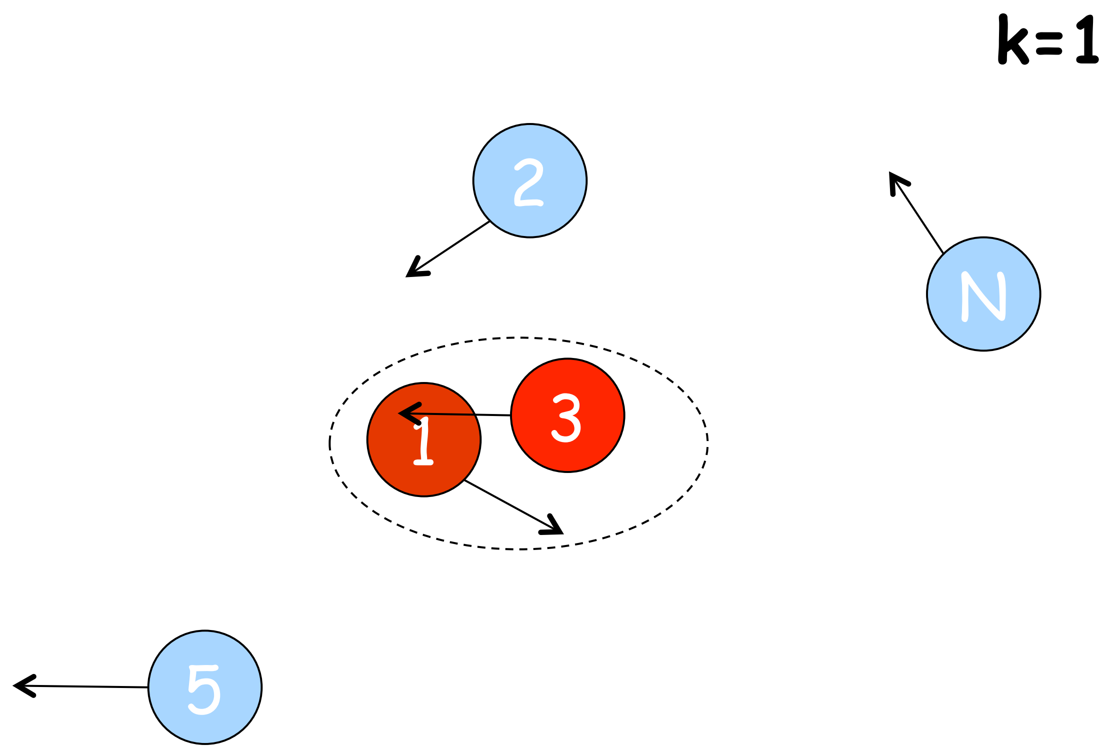

At each slot there is a probability p that two given nodes meet.
Assume meetings to be independent.

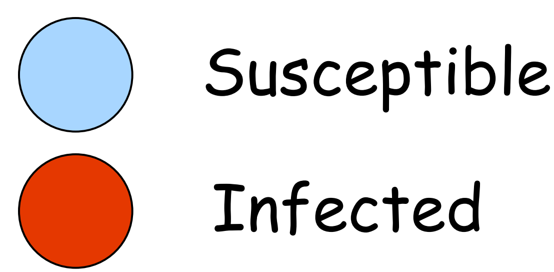

A motivating example: epidemics

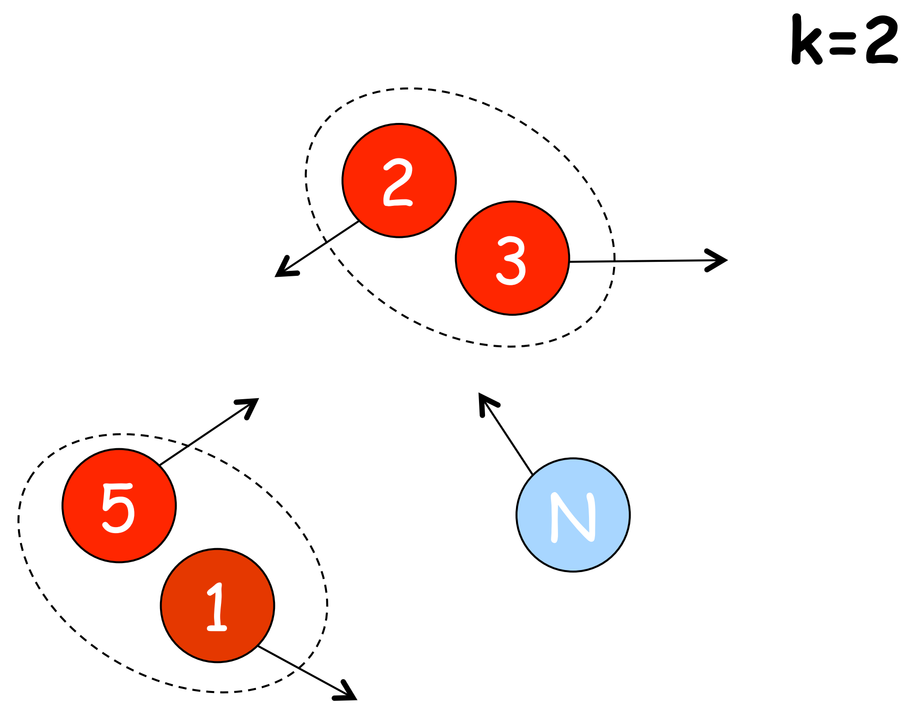

At each slot there is a probability p that two given nodes meet.
Assume meetings to be independent.

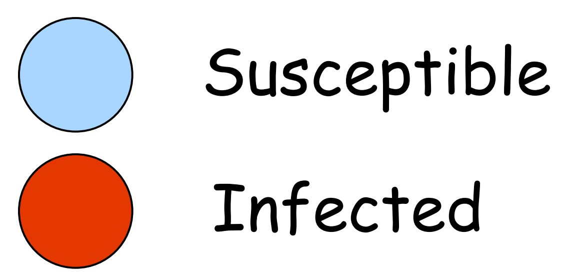

Delay Tolerant Networks
(a.k.a. Intermittently Connected Networks)

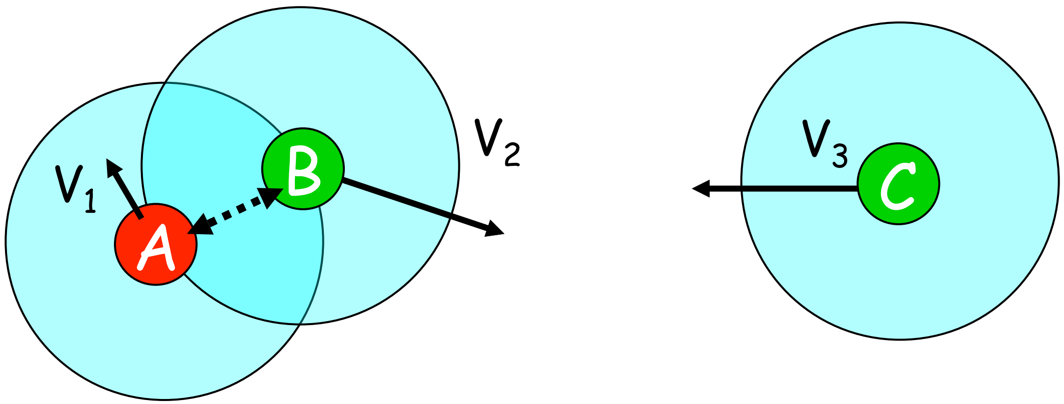

- mobile wireless networks
- no path at a given time instant between two nodes
  - because of power contraint, fast mobility dynamics
  - maintain capacity, when number of nodes (N) diverges
    - Fixed wireless networks: $C = \Theta(\text{sqrt}(1/N))$ [Gupta99]
    - Mobile wireless networks: $C = \Theta(1)$, [Grossglauser01]
- a really challenging network scenario
  - No traditional protocol works

Some examples

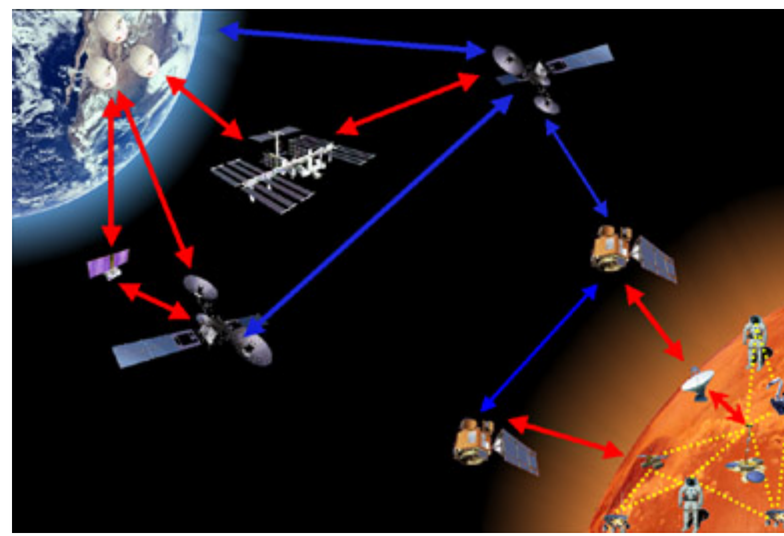

Inter-planetary backbone

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

Epidemic Routing

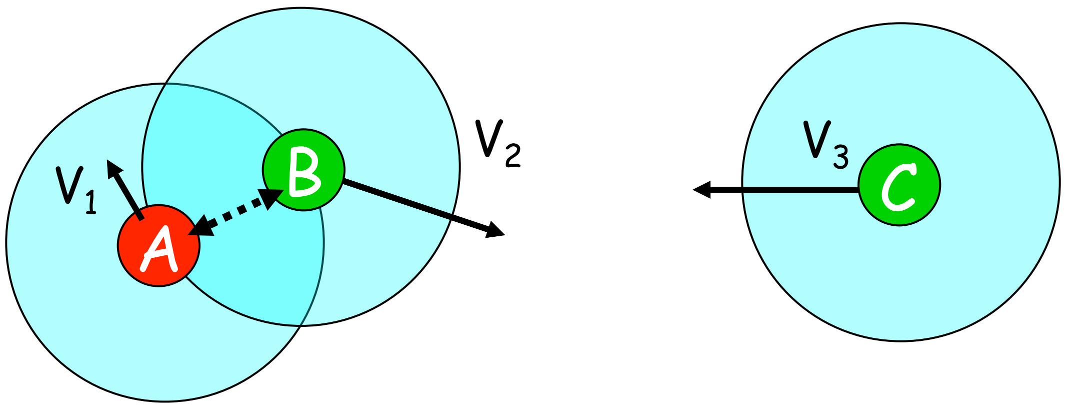

- Message as a disease, carried around and transmitted

Epidemic Routing

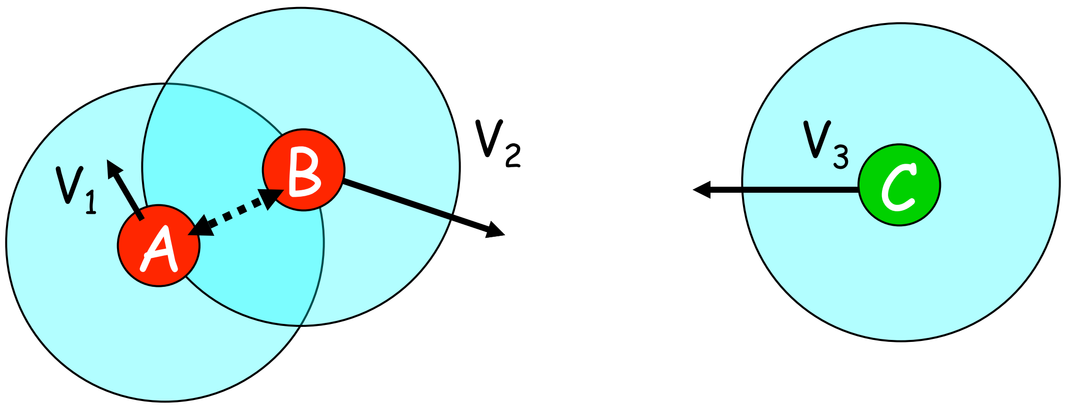

- Message as a disease, carried around and transmitted
  - Store, Carry and Forward paradigm

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

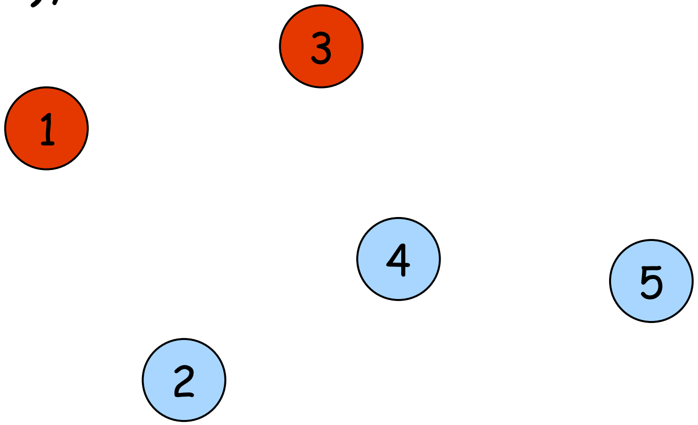

  - Probability transitions among states
    - e.g. Prob((1,0,1,0,0)->(1,1,1,0,0))=?

Transition probabilities

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

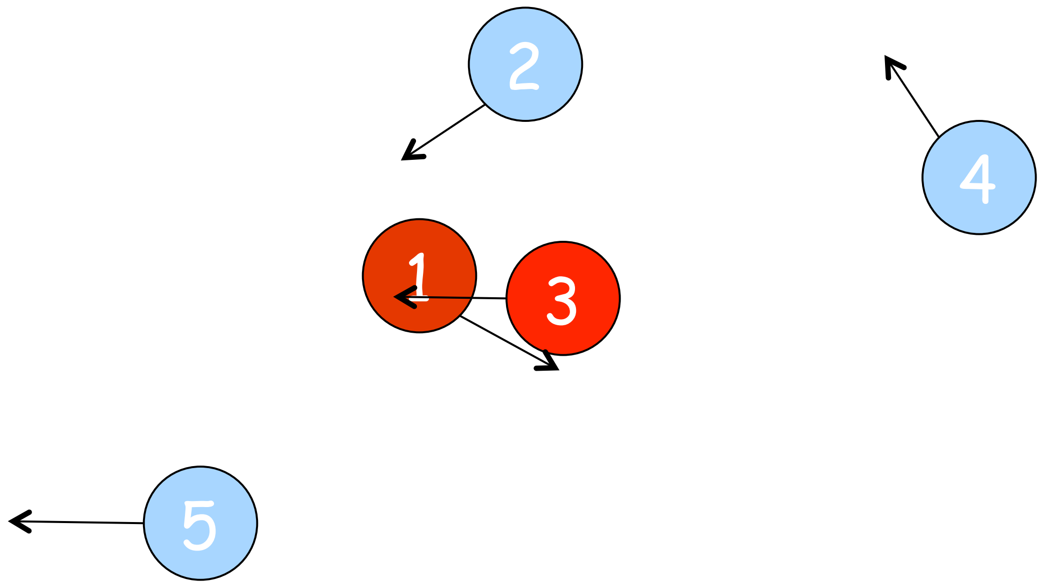

At slot k, when there are I(=I(k)) infected nodes, the prob. that node 2 gets infected is: $q_I=1-(1-p)^I$

Transition probabilities

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

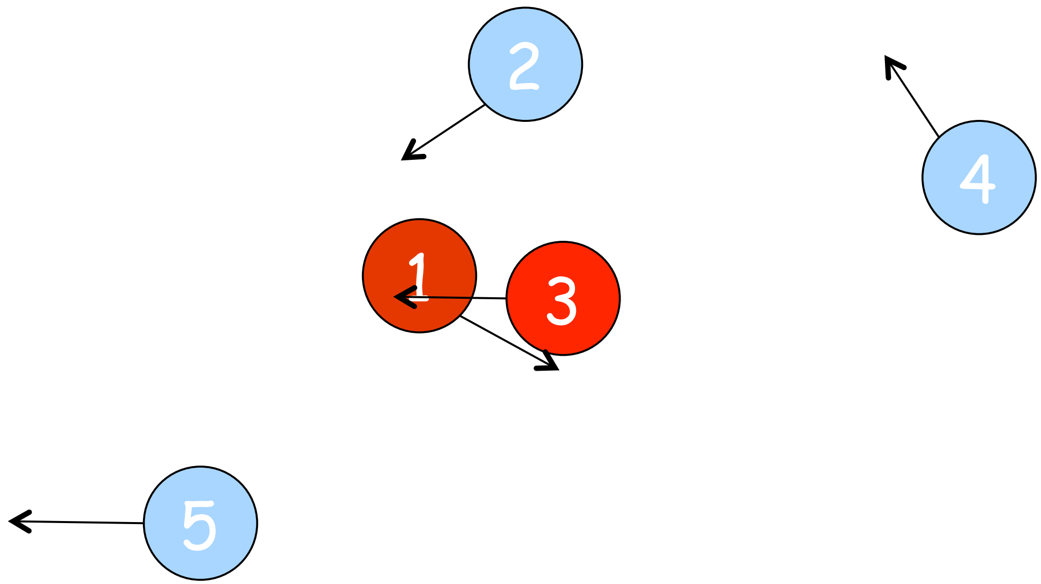

Prob((1,0,1,0,0)->(1,1,1,0,0))=$q_2(1-q_2)^2$
Where $q_I=1-(1-p)^I$

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
    - $O(K\ A)$ where A is the number of non-null entries in P if we adopt the iterative procedure (K is the number of iterations), in our case $M \varepsilon O((2^N)^2)$

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
  - $Prob(I(k+1)=I+n \mid I(k)=I) = C^n_{N-I}\, q_I^n\, (1-q_I)^{N-n-I}$
    - $(I(k+1)-I(k) \mid I(k)=I) \sim Bin(N-I, q_I)$
    - $q_I=1-(1-p)^I$

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

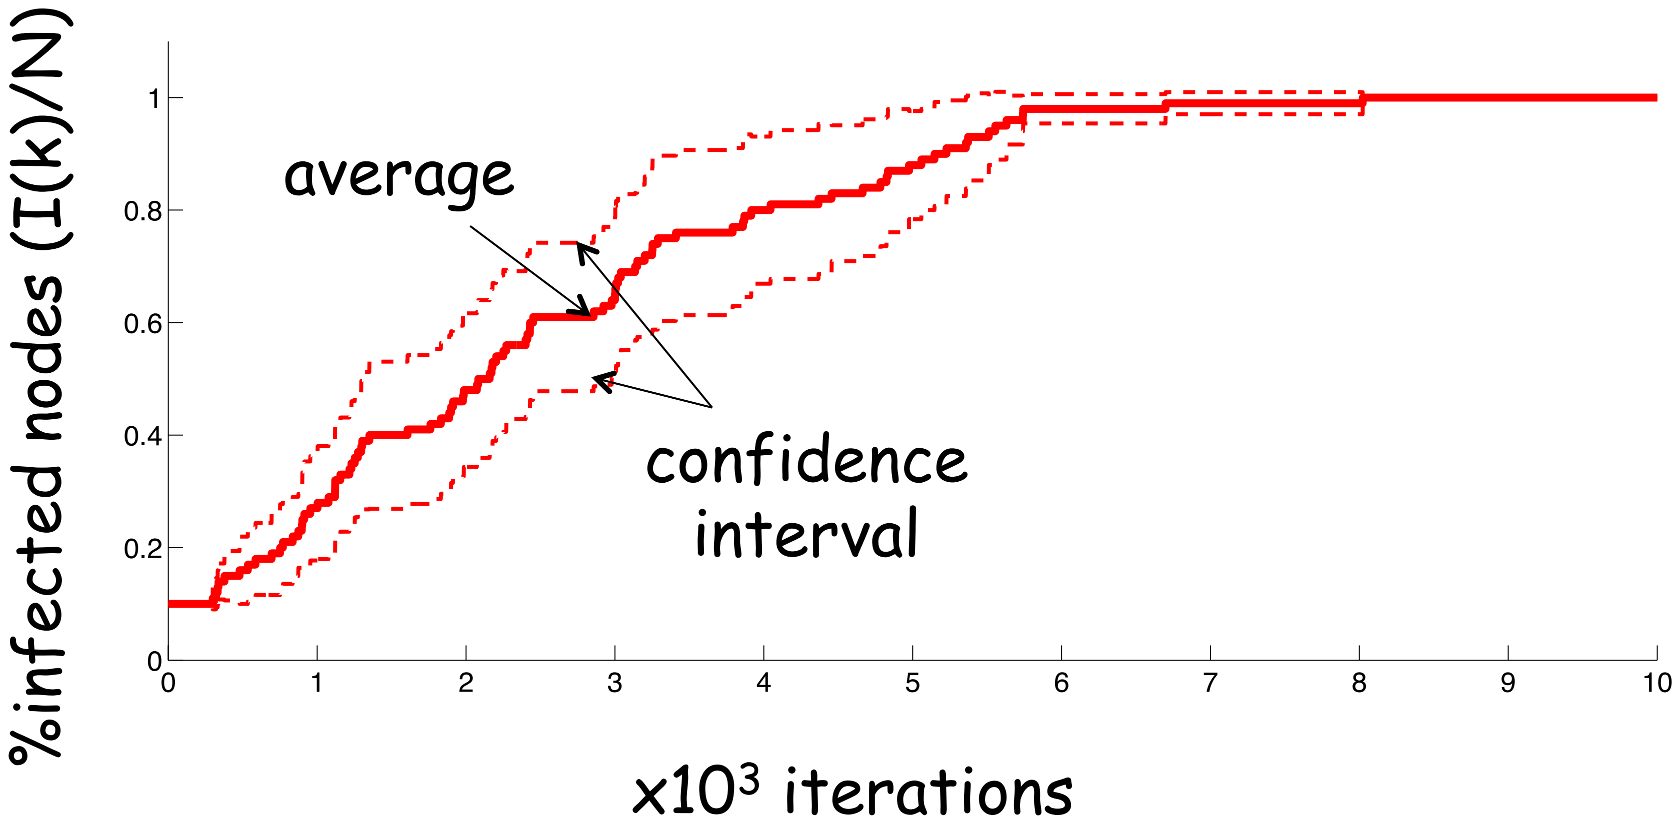

Some numerical examples
$p=10^{-4}$, $N=100$, $I(0)=N/10$, 10 runs

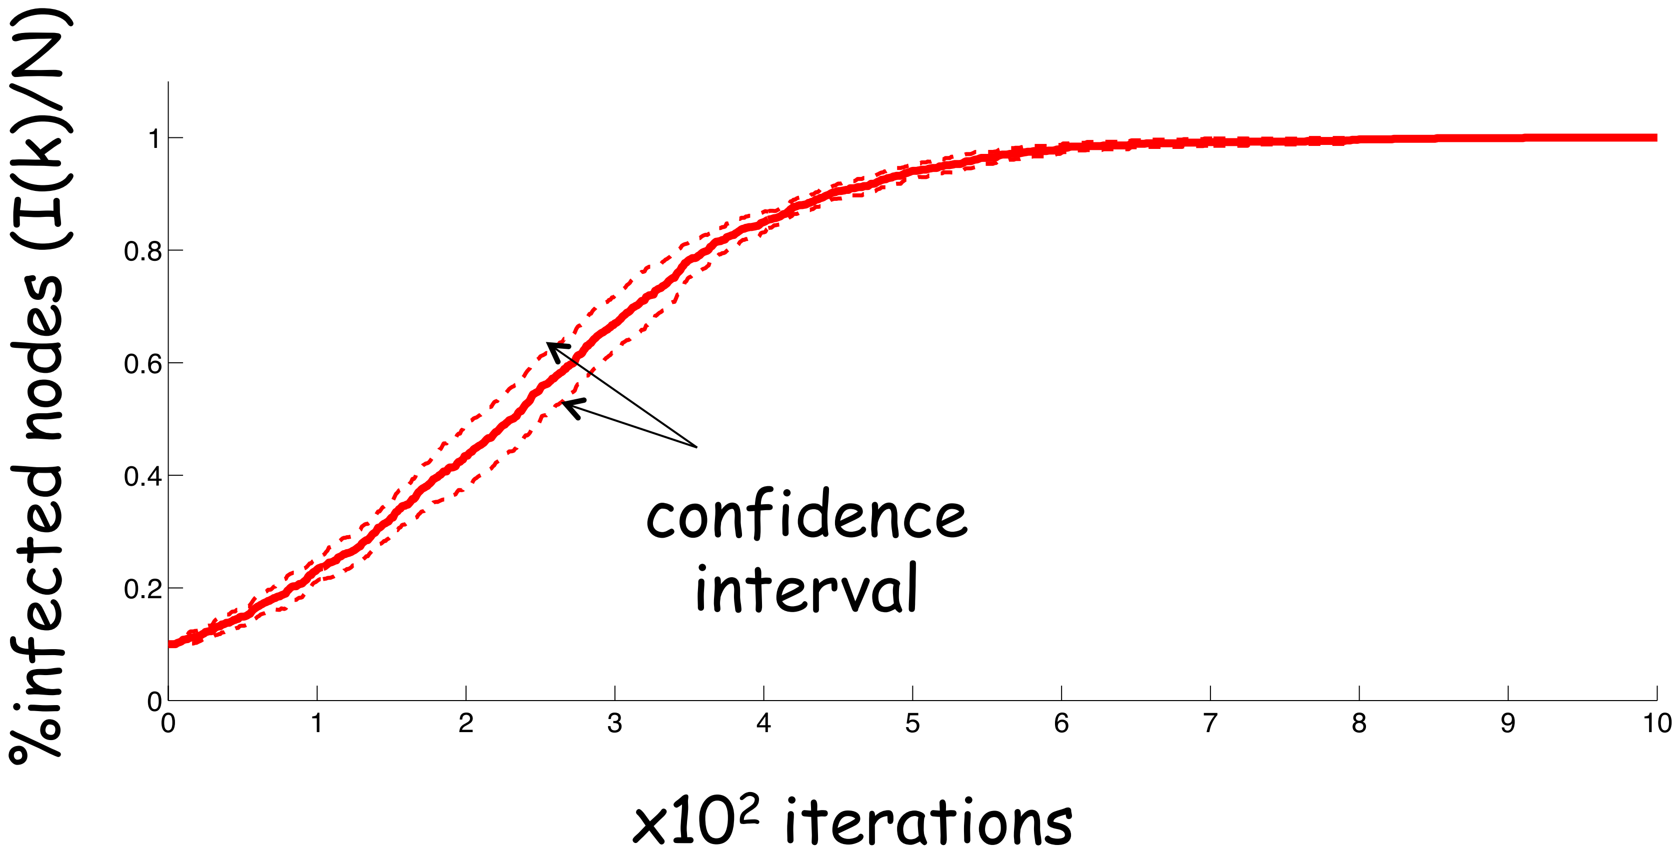

Some numerical examples $p=10^{-4}$, $N=10000$, $I(0)=N/10$, 10 runs

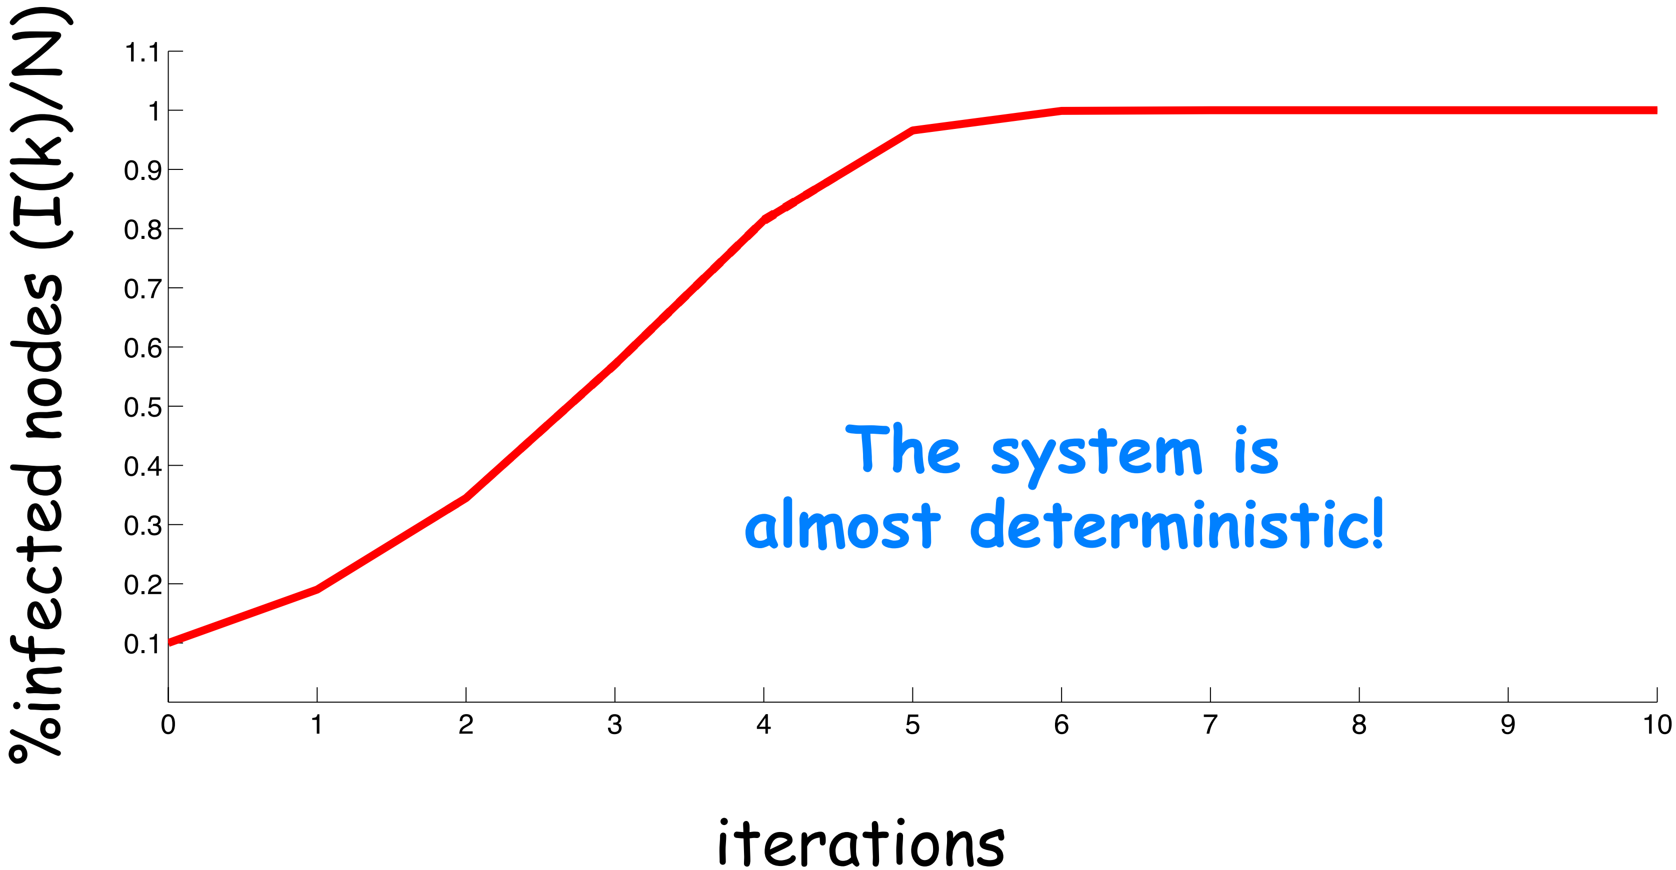

Summary

- 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

Spoiler

- $i^{(N)}(k)$, fraction of infected nodes at time k
- Solve
  $di(t)/dt=i(t)(1-i(t))$,
  with $i=i_0$
  - Solution: $i(t)=1/((1/i_0-1)\, e^{-t}+1)$

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

Outline

- Limit of Markovian models
- Mean Field (or Fluid) models
  - exact results
  - extensions to graphs
  - 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:
  - $P_{\sigma,\sigma'}=\mathrm{Prob}((v_1,v_2,\ldots v_N)\to(u_1,u_2,\ldots u_N))=$
    $\mathrm{Prob}((v_{\varphi(1)},v_{\varphi(2)},\ldots v_{\varphi(N)})\to(u_{\varphi(1)},u_{\varphi(2)},\ldots u_{\varphi(N)}))$

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
  - #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 $1/N$ on $[0,1]^V \times [0,\beta]$, with $\beta>0$

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 $\mathbf{M}^{(N)}(0) \to \mathbf{m}_0$ in probability (/mean square) as $N \to \infty$, then $\sup_{0\leq t\leq T}||\underline{\mathbf{M}}^{(N)}(t)-\boldsymbol{\mu}(t)|| \to 0$ 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 $\varepsilon(N)$

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

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

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

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

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

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

Application to the SI model

- Out of the impasse: introduce a scaling for p
  - If $p^{(N)}=p_0/N^a$ $a>1 \Rightarrow (1-i)(1-(1-p^{(N)})^{i\,N})\to 0$
  - 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

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

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

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

    $E[W^N(k)^2]=((N-I(k))q_I)^2 + (N-I(k))q_I(1-q_I)$
    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
    - $f_2^{(N)}(\mathbf{m})/\varepsilon(N) =$
      $=(1-i)\,(1 - (1-(p_0/N^2))^{i\,N})/(p_0/N)$
    - continuous derivatives in i and in 1/N (not evident)

Practical use of the convergence result

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

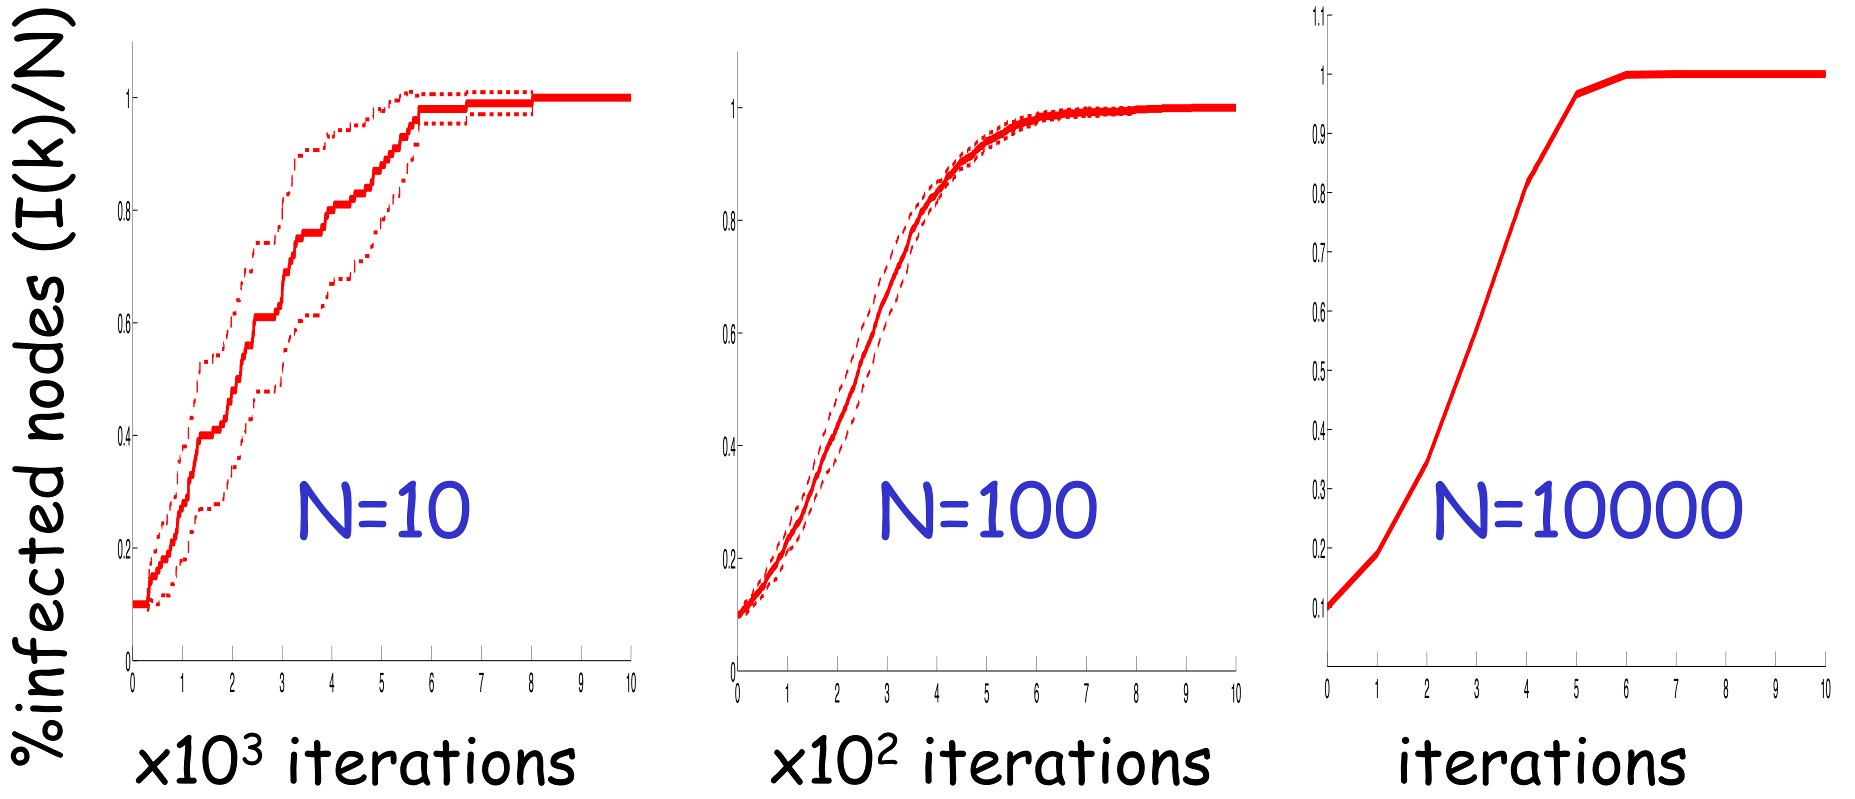

Advantage of Mean Field

- 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