Performance Evaluation

Second Part
Lecture 3

Giovanni Neglia
INRIA – EPI Maestro
23 January 2012

Outline

- Limit of Markovian models
- Mean Field (or Fluid) models
  - exact results
  - extensions
  - Applications
    - Bianchi's model
  - (Heterogeneous networks)
  - Applications
    - Epidemic routing

Heterogeneous Networks

- Denote P(d) the probability that a node has degree d
- If the degree does not change much, we can replace d with <d>
  - what we have done for ER graphs (N,p)
    - Binomial with parameters (N-1,p)
- How should we proceed (more) correctly?
  - Split the nodes in degree classes
  - Write an equation for each class
- Remark: following derivation will not be as rigorous as previous ones

Heterogeneous Networks

- N_d number of nodes with degree d (=N*P(d))
- I_d: number of infected nodes with degree d
- Given node i with degree d and a link e_{ij}, what is the prob. that j has degree d'?
  - P(d')? NO
- and if degrees are uncorrelated? i.e. Prob(neighbour has degree d'|node has a degree d) independent from d,
  - P(d')? NO
  - Is equal to $d'/\langle d \rangle P(d')$

Heterogeneous Networks

- Given node i with degree d and a link e_{ij}
- Prob. that j has degree d' is
  - d'/<d> P(d')
- Prob. that j has degree d' and is infected
  - d'/<d> P(d') $I_{d'}/N_{d'}$
  - more correct (d'-1)/<d> P(d') $I_{d'}/N_{d'}$
- Prob. that i is infected through link e_{ij} is
  - $p = p_g \Sigma_{d'} (d'-1)/\langle d\rangle P(d') I_{d'}/N_{d'}$
- Prob. that i is infected through one link
  - $1-(1-p)^d$

Heterogeneous Networks

- $E[(I_d(k+1)-I_d(k)|\mathbf{I}(k)=\mathbf{I})] = (N_{d'}-I_d)(1-(1-p)^d)$
  - $p = p_g \Sigma_{d'} (d'-1)/\langle d\rangle P(d')\, I_{d'}/N_{d'}$
- $f_d^{(N)}(i)=(1-i_d)(1-(1-p)^d)$
  - $i_d = I_d/N_d$
  - if we choose $p_g = p_{g0}/N$
  - $f_d(\mathbf{i})=(1-i_d)\, p_{g0}\, d \underbrace{\Sigma_{d'}(d'-1)/\langle d\rangle P(d')\, i_{d'}}_{\Theta}$
- $di_d(t)/dt=f_d(\mathbf{i}(t))=p_{g0}\,(1-i_d(t))\,\Theta(t)$

Heterogeneous Networks

- $di_d(t)/dt = f_d(i(t)) = p_{g0}\,(1 - i_d(t))\, d\, \Theta(t)$,
  - for d=1,2…
  - $\Theta(t) = \Sigma_{d'} (d'-1)/\langle d\rangle\, P(d')\, i_{d'}(t)$
  - $i_d(0) = i_{d0}$, for d=1,2…
- If $i_d(0) \ll 1$, for *small* t
  - $di_d(t)/dt \approx p_{g0}\, d\, \Theta(t)$
  - $d\Theta(t)/dt = \Sigma_{d'} (d'-1)/\langle d\rangle\, P(d')\, di_{d'}(t)/dt$

$$\approx p_{g0}\, \Sigma_{d'} (d'-1)/\langle d\rangle\, P(d')\, d'\, \Theta(t) =$$

$$= p_{g0}\, (\langle d^2\rangle - \langle d\rangle)/\langle d\rangle\, \Theta(t)$$

Heterogeneous Networks

- $d\Theta(t)/dt \approx p_{g0}(\langle d^2\rangle - \langle d\rangle)/\langle d\rangle\ \Theta(t)$
  - Outbreak time: $\langle d\rangle/((\langle d^2\rangle - \langle d\rangle)\ p_{g0})$
    - For ER $\langle d^2\rangle = \langle d\rangle(\langle d\rangle + 1)$, we find the previous result, $1/(\langle d\rangle p_{g0})$
    - What about for Power-law graphs, $P(d) \sim d^{-\gamma}$?
- For the SIS model:
  - $d\Theta(t)/d \approx p_{g0}(\langle d^2\rangle - \langle d\rangle)/\langle d\rangle\ \Theta(t) - r_0\ \Theta(t)$
  - Epidemic threshold: $p_{g0}\ (\langle d^2\rangle - \langle d\rangle)/(\langle d\rangle r_0)$

Outline

- Limit of Markovian models
- Mean Field (or Fluid) models
  - exact results
  - extensions
  - Applications
    - Bianchi's model
  - (Heterogeneous networks)
  - Applications
    - Epidemic routing

Outline and references

Introduction to epidemic routing

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

Standard Epidemic Routing

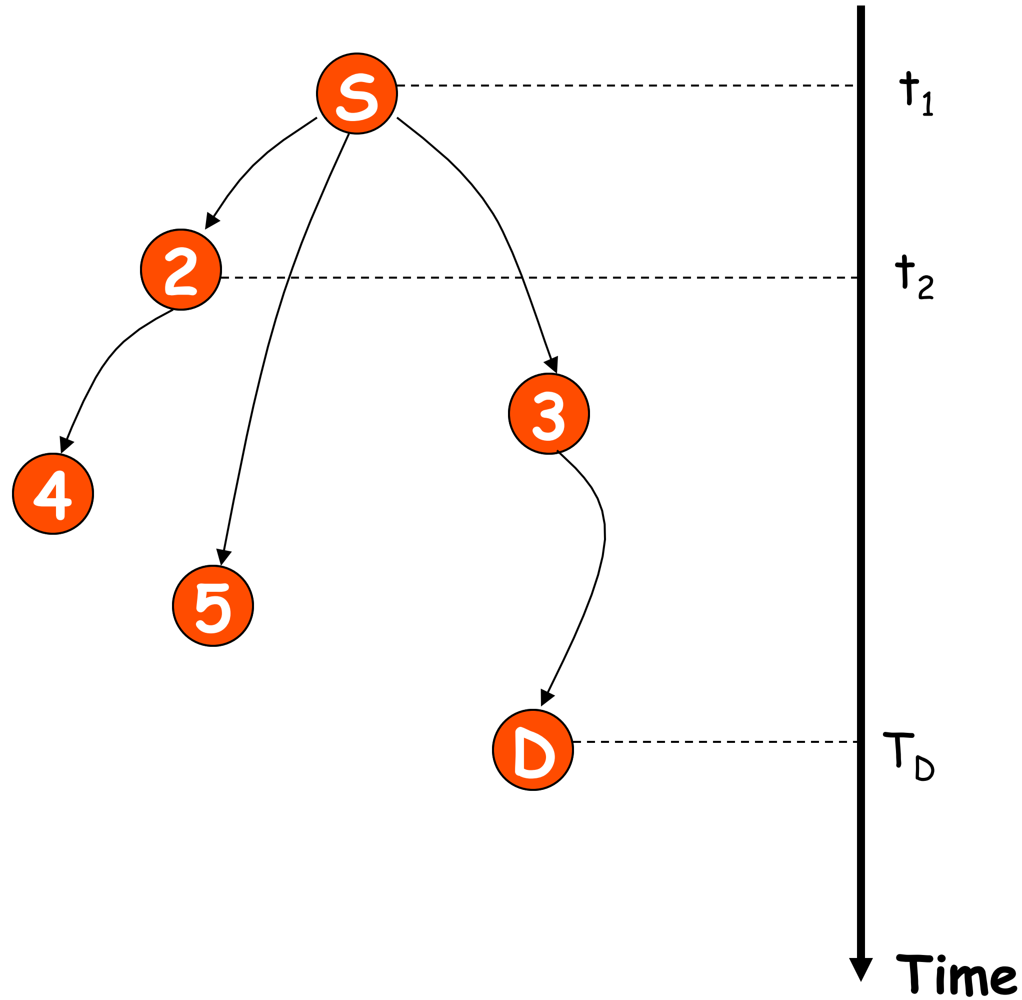

Epidemic Style Routing

- Epidemic Routing [Vahdat&Becker00]
  - Propagation of a pkt -> Disease Spread
  - achieve min. delay, at the cost of transm. power, storage
- trade-off delay for resources
  - K-hop forwarding, probabilistic forwarding, limited-time forwarding, spray and wait…

2-Hop Forwarding

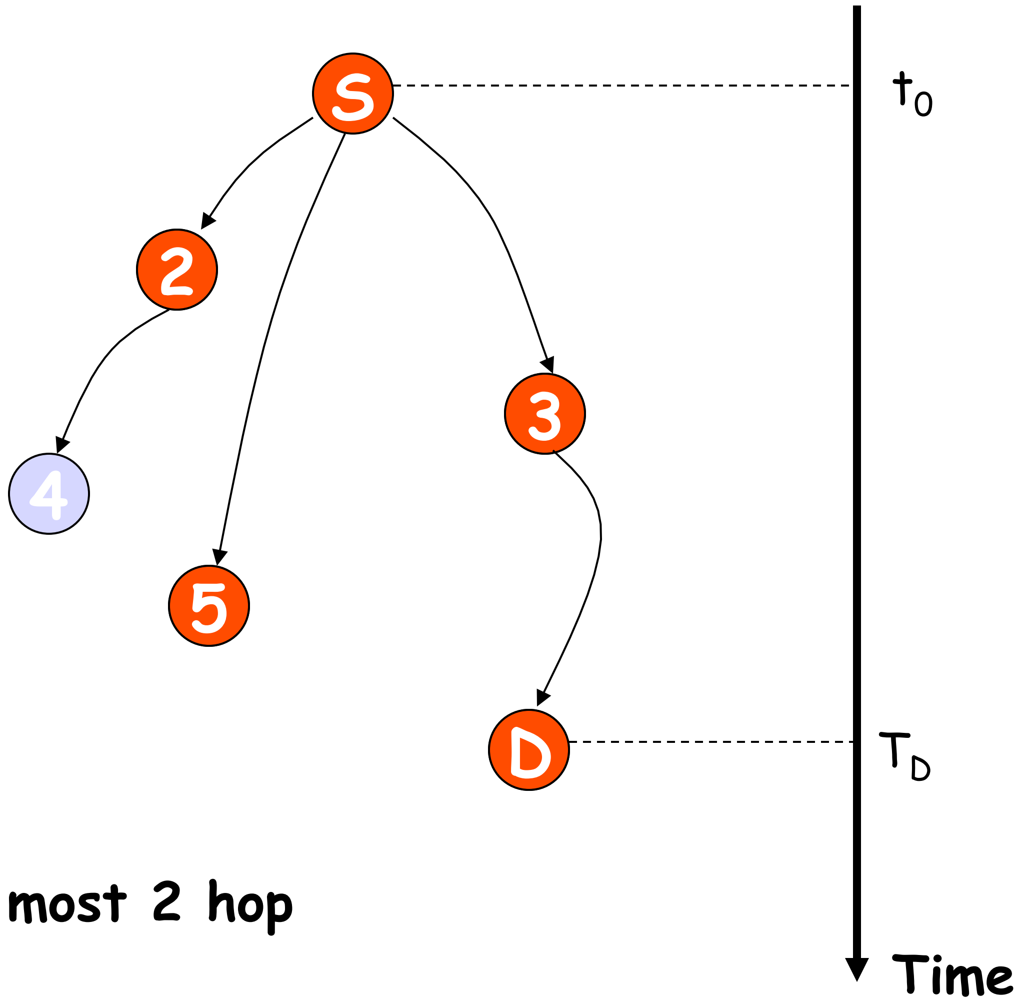

At most 2 hop

Limited Time Forwarding

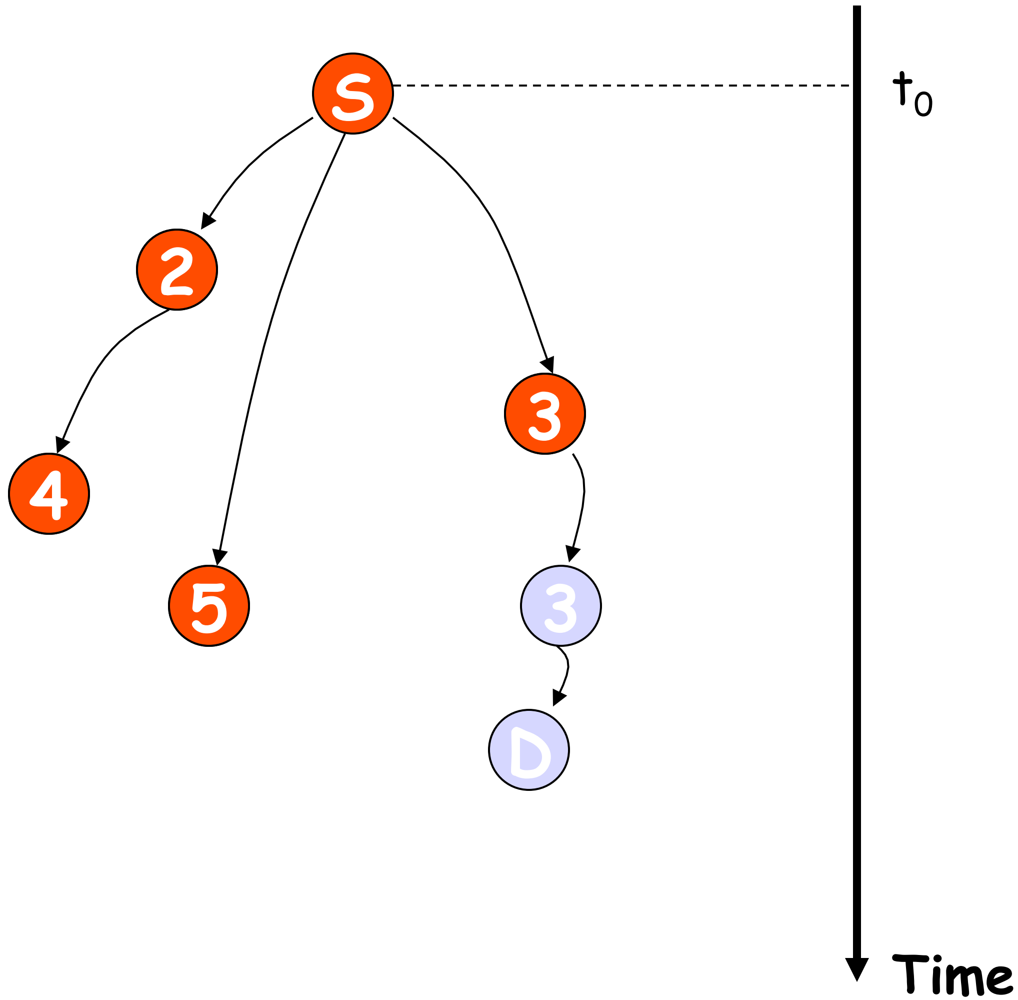

Epidemic Style Routing

- Epidemic Routing [Vahdat&Becker00]
  - Propagation of a pkt -> Disease Spread
  - achieve min. delay, at the cost of transm. power, storage
- trade-off delay for resources
  - K-hop forwarding, probabilistic forwarding, limited-time forwarding, spray and wait…
- Recovery: deletion of obsolete copies after delivery to dest., e.g.,
  - TIMERS: when time expires all the copies are erased
  - IMMUNE: dest. cures infected nodes
  - VACCINE: on pkt delivery, dest propagates anti-pkt through network

IMMUNE Recovery

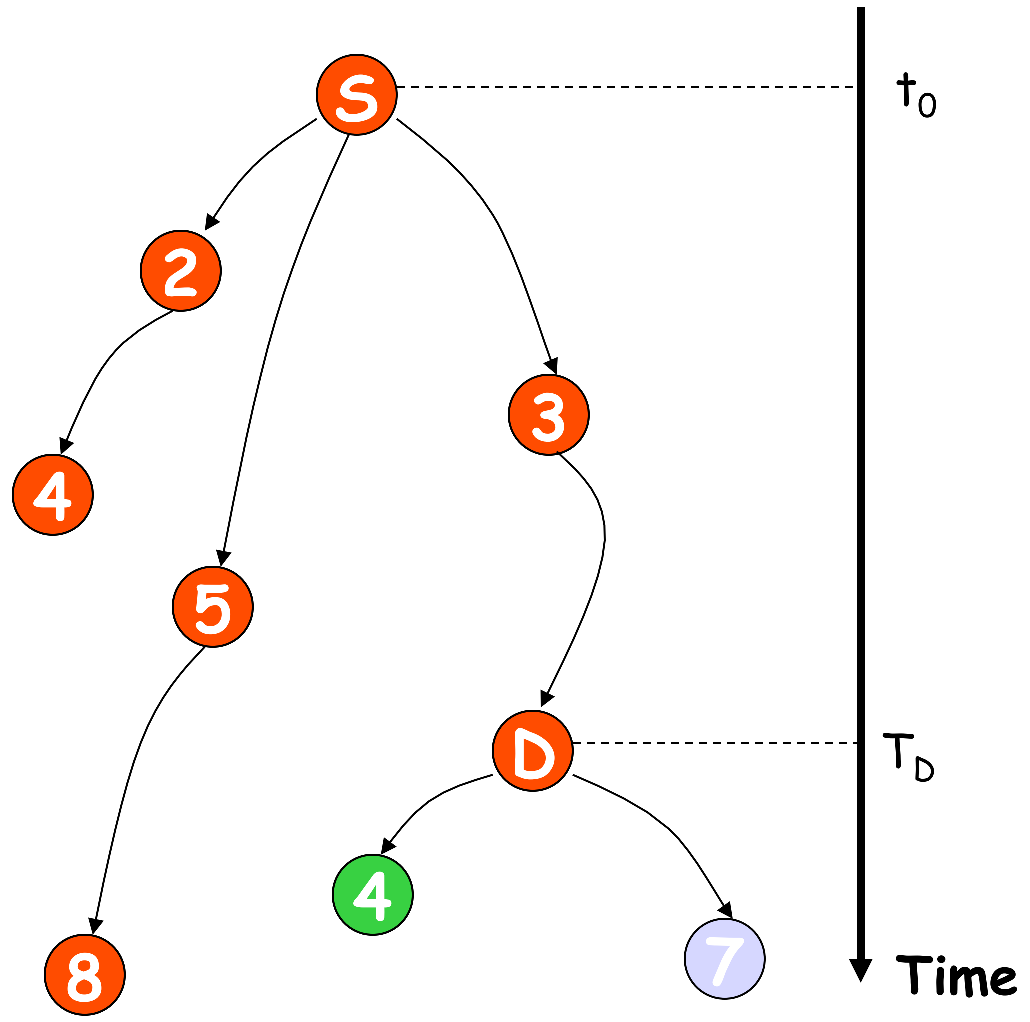

VACCINE Recovery

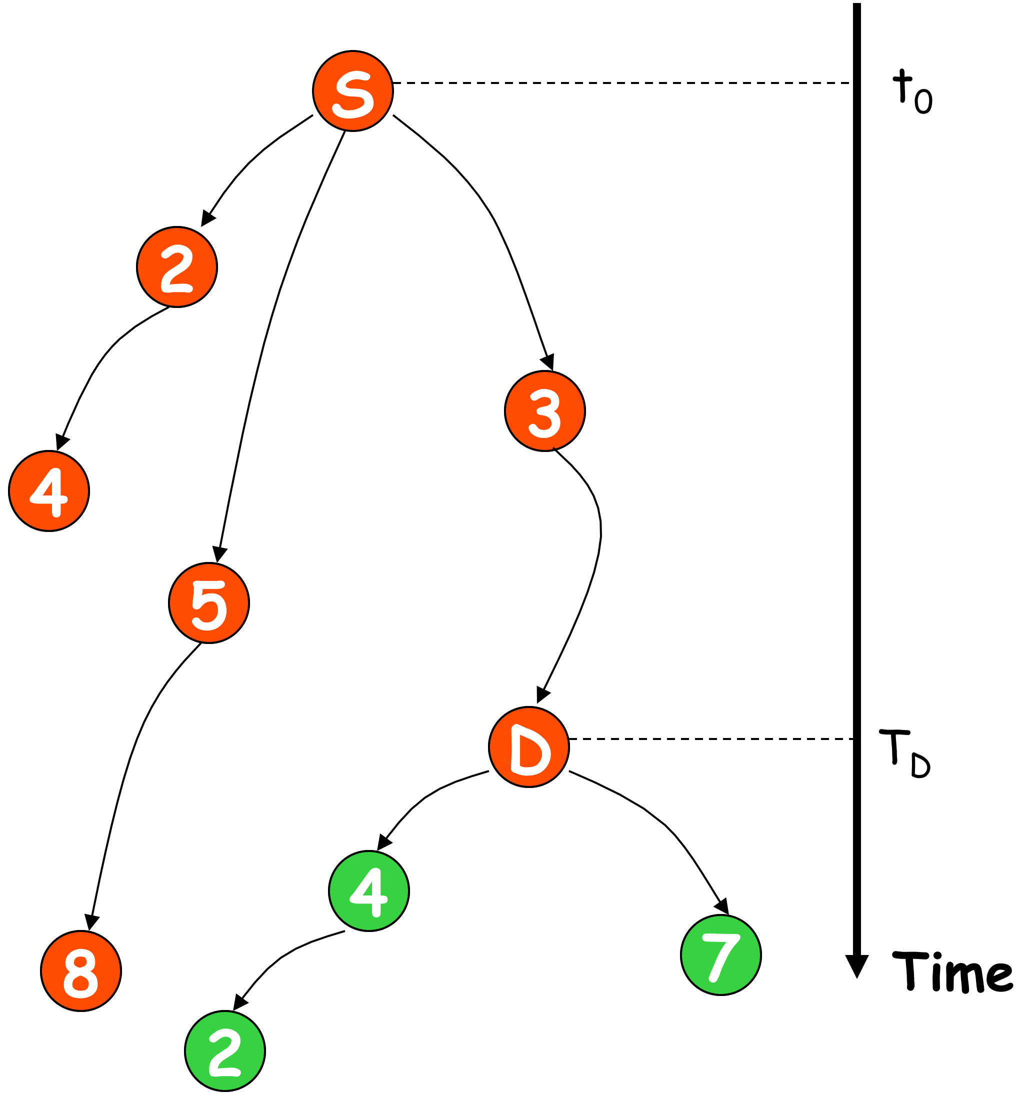

Outline

- Introduction to Epidemic Routing
- Markovian models
  - the key to Markov model
  - Markovian analysis of epidemic routing
- Fluid models

The setting we consider

- N+1 nodes
  - moving independently in an finite area A
  - with a fixed transmission range r and no interference
  - 1 source, 1 destination
- Performance metrics:
  - Delivery delay T_d
  - Avg. num. of copies at delivery C
  - Avg. total num. of copies made G
  - Avg. buffer occupancy

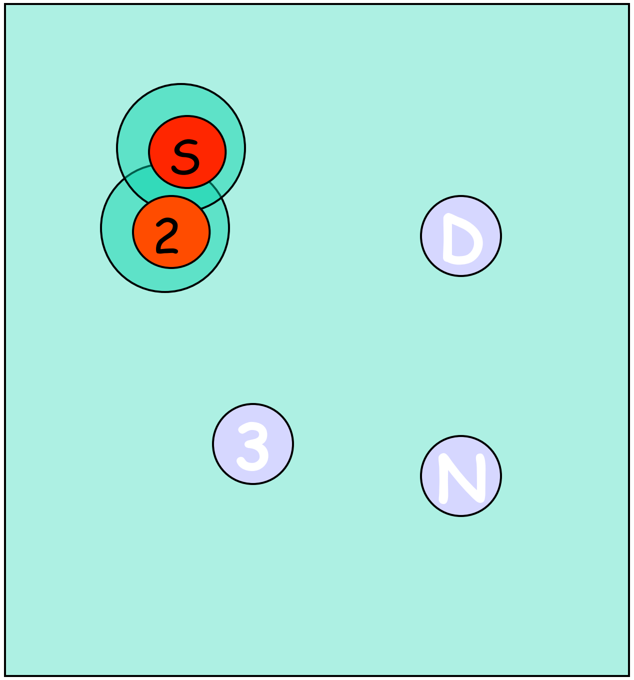

Standard random mobility models

Random Waypoint model (RWP)

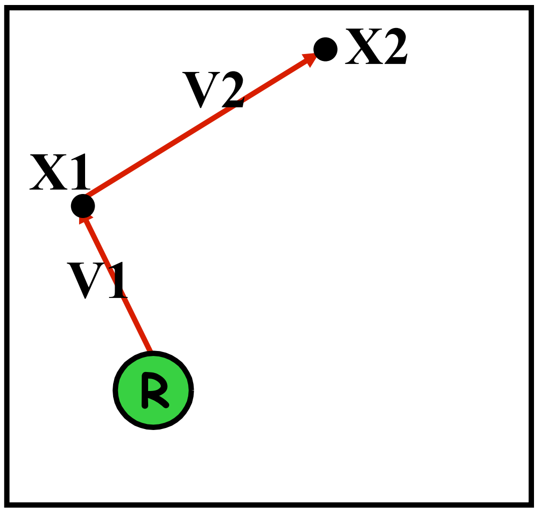

- Next positions (Xi)s are uniformly distributed
- Speeds (Vi)s are uniformly distributed (Vmin,Vmax)

Random Direction model (RD)

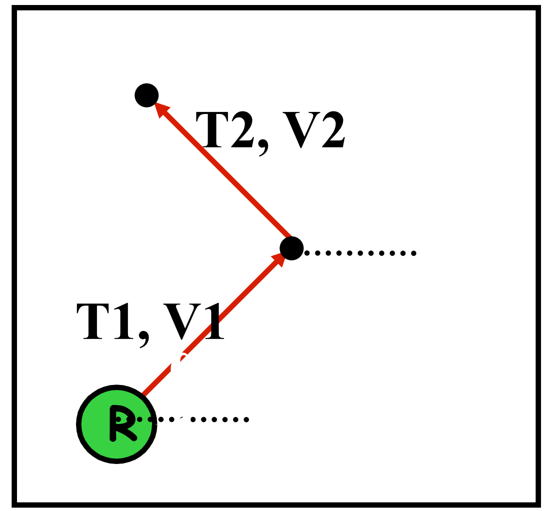

- Directions (αi) are uniformly distributed (0, 2π)
- Speeds (Vi) are uniformly distributed (Vmin,Vmax)
- Travel times (Ti) are exponentially / generally distributed

The key to Markov model

[Groenevelt05]

if nodes move according to standard random mobility model (random waypoint, random direction) with average relative speed E[**V***],

and if **Nr²** is small in comparison to ***A***

pairwise meeting processes are ***almost*** independent Poisson processes with rate:

$$\lambda \approx \frac{2wrV^*}{A}$$

w: mobility specific constant

Intuitive explanation

Exponential distribution finds its roots in the independence assumptions of each mobility model:

- Nodes move independently of each other
- Random waypoint: future locations of a node are independent of past locations of that node.
- Random direction: future speeds and directions of a node are independent of past speeds and directions of that node.

There is some probability q that two nodes will meet before the next change of direction. At the next change of direction the process repeats itself, **almost** independently.

Why “almost”?

pairwise meeting processes are ***almost*** independent Poisson processes with rate:

$$\lambda \approx \frac{2wrV^*}{A}$$

w: mobility specific constant

1. inter-meeting times are not exponential
    - if N1 and N2 have met in the near past they are more likely to meet (they are close to each other)
        - the more the bigger it is r^2 in comparison to *A*
2. meeting processes are not independent
    - if in [t,t+τ] N1 meets N2 and N2 meets N3, it is more likely that N1 meets N3 in the same interval
        - the more the bigger it is r^2 in comparison to *A*
    - moreover if Nr^2 is comparable with *A* (dense network) a lot of meeting happen at the same time.

Examples

Nodes move on a square of size 4x4 km^2 (L=4 km)

Different transmission radii (R=50,100,250 m)

Random waypoint and random direction:

no pause time

$[v_{min},v_{max}]=[4,10]$ km/hour

Random direction: travel time ~ exp(4)

Pairwise Inter-meeting time

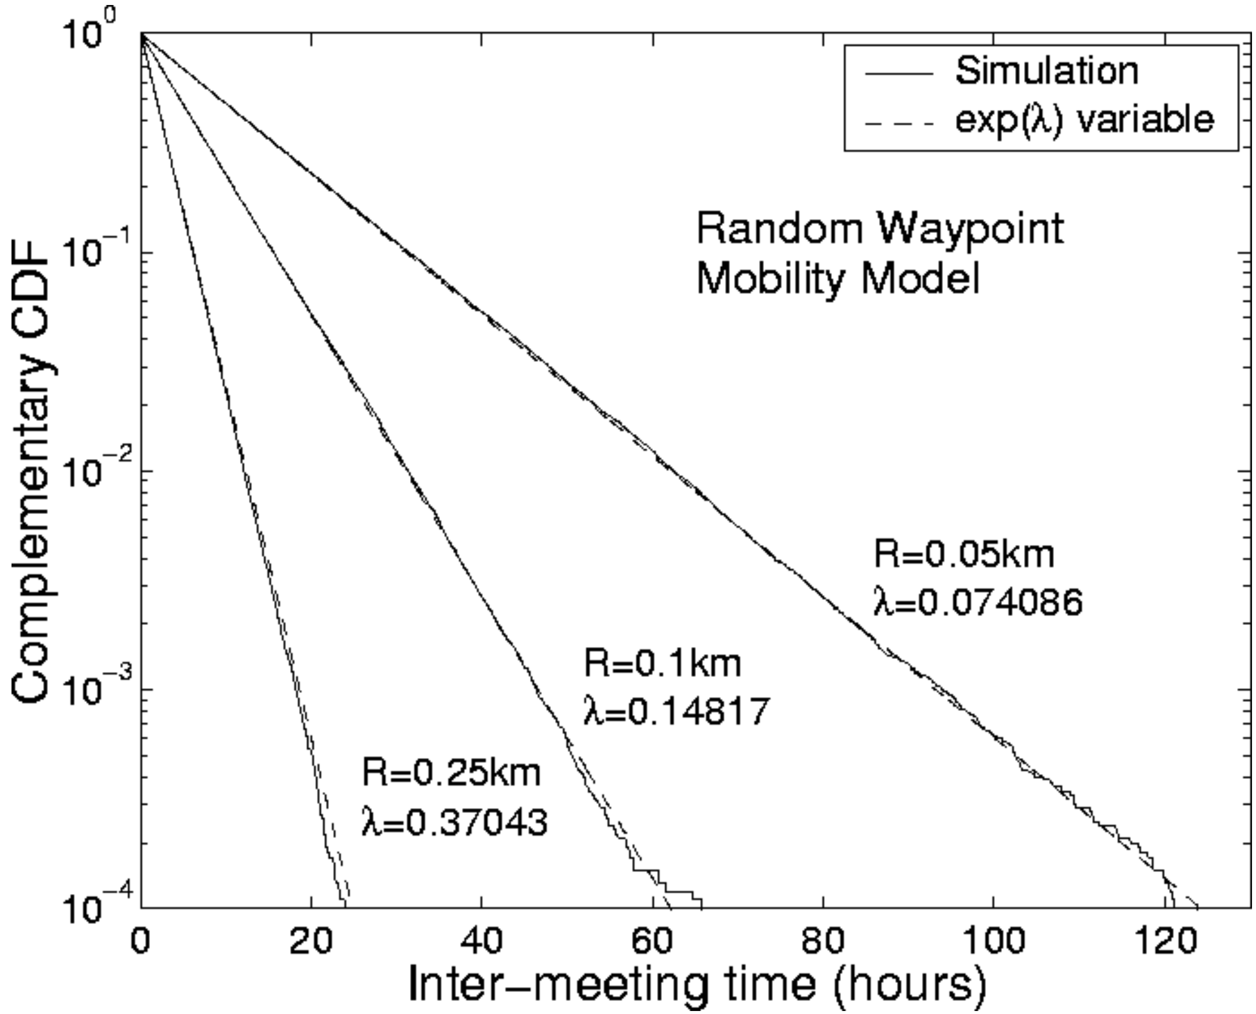

$$P(T > t) = e^{-\lambda t}$$

$$\log P(T > t) = -\lambda t$$

Pairwise Inter-meeting time

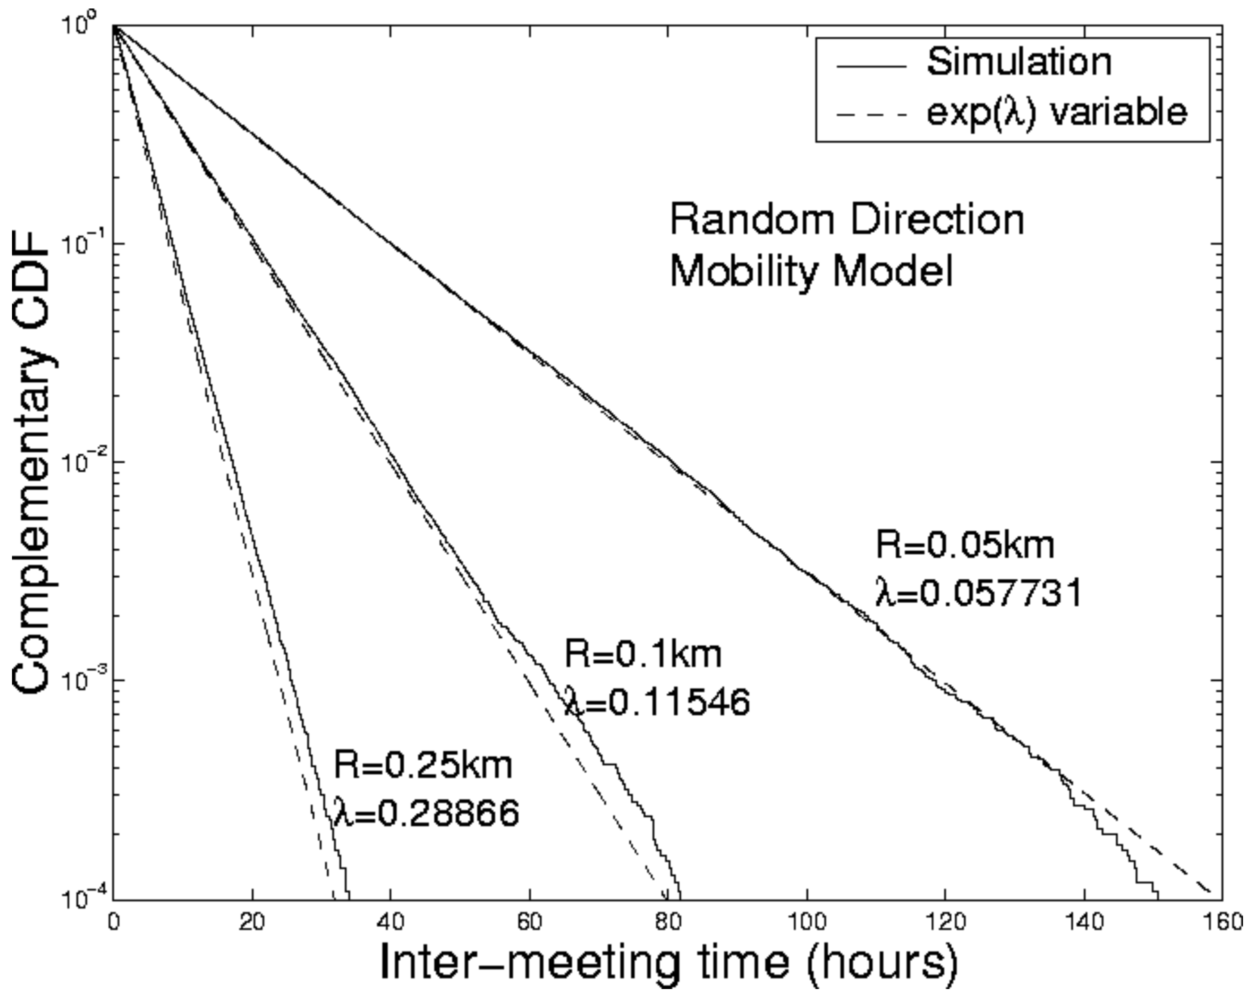

$$P(T > t) = e^{-\lambda t}$$

$$\log P(T > t) = -\lambda t$$

Pairwise Inter-meeting time

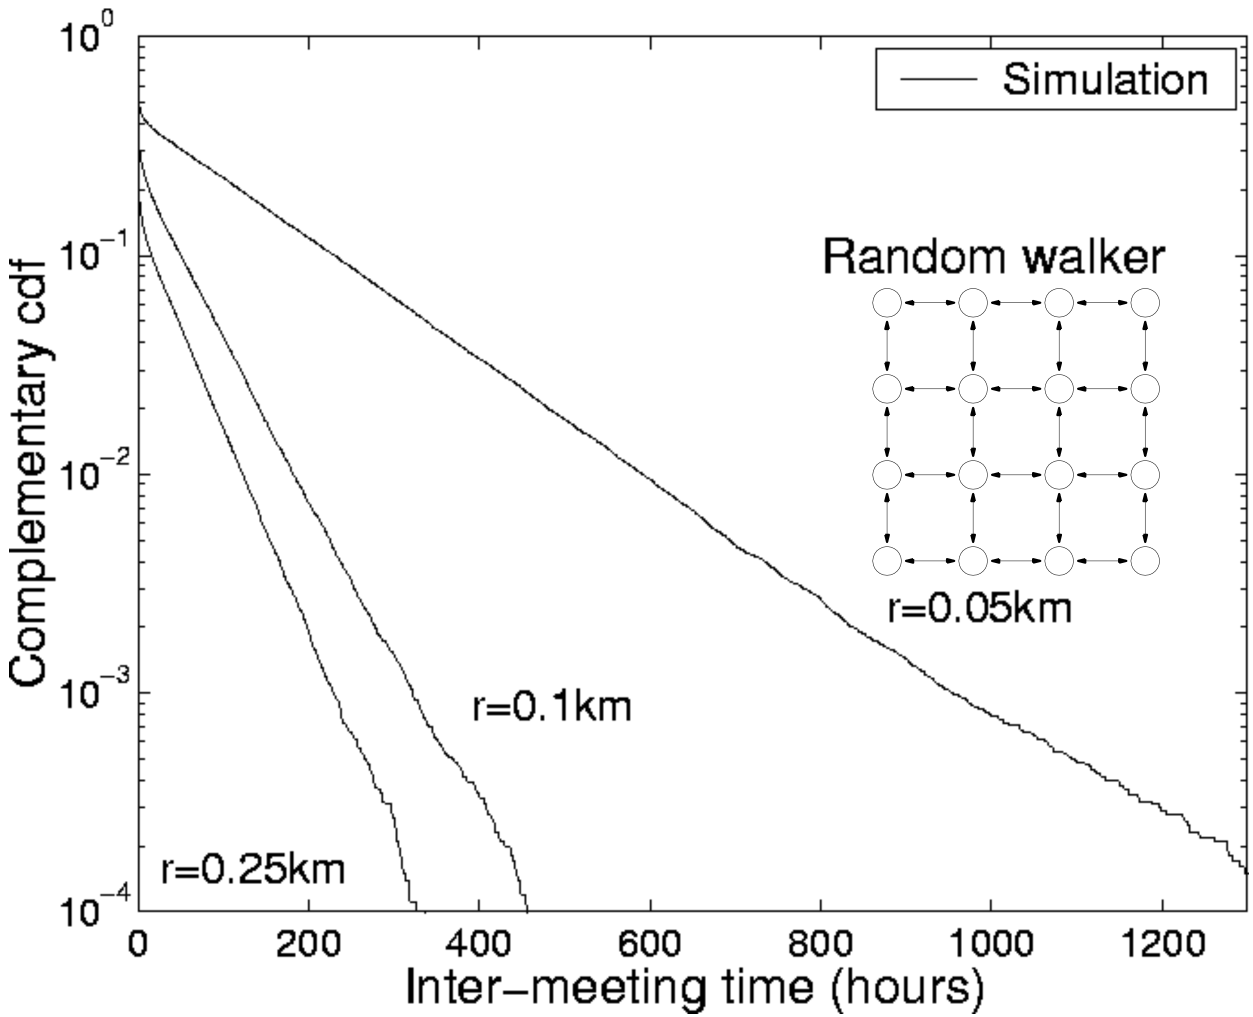

$$P(T > t) = e^{-\lambda t}$$

$$\log P(T > t) = -\lambda t$$

The derivation of λ

Assume a node in position (x_1,y_1) moves in a straight line with speed V_1.

Position of the other node comes from steady-state distribution with pdf $\pi(x,y)$.

Look at the area A covered in Δ_t time:

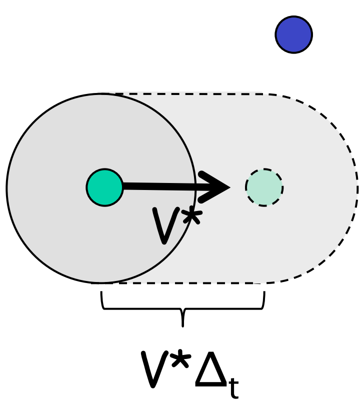

The derivation of λ

Probability that nodes meet given by

$$p_{x_1,y_1} = \iint_A \pi(x,y)\,dx\,dy.$$

For small *r* the points in $\pi(x,y)$ in *A* can be approximated by $\pi(x_1,y_1)$ *to give*

$$p_{x_1,y_1} \approx 2r \cdot V_1^* \cdot \Delta_t \cdot \pi(x_1,y_1).$$

Unconditioning on (x_1,y_1) gives

$$p = \int_0^L \int_0^L p_{x_1,y_1} \cdot \pi(x_1,y_1)\,dx_1 y_1$$

$$\approx 2r \cdot V_1^* \cdot \Delta_t \cdot \int_0^L \int_0^L \pi^2(x_1,y_1)\,dx_1 y_1$$

The derivation of λ

Proposition: Let r<<L. The inter-meeting time for the random direction and the random waypoint mobility models is approximately exponentially distributed with parameter

$$\lambda \approx 2\,r \cdot E[V^*] \cdot \int_0^L \int_0^L \pi^2(x,y)\,dx\,dy,$$

Here *E[V*]* is the average relative speed between two nodes and $\pi(x,y)$ is the pdf in the point *(x,y)*.

The derivation of λ

Proposition: Let r<<L. The inter-meeting time for the random direction and the random waypoint mobility models is approximately exponentially distributed with parameter

$$\lambda_{RD} \approx \frac{2rE[V^*]}{L^2}, \qquad \lambda_{RW} \approx \frac{2r\omega E[V^*]}{L^2}.$$

Here *E[V*]* is the average relative speed between two nodes and ω ≈ 1.3683 is the Waypoint constant.

If speeds of the nodes are constant and equal to *v*,

$$\lambda_{RD} \approx \frac{8rv}{\pi L^2}, \qquad \lambda_{RW} \approx \frac{8\omega rv}{\pi L^2}.$$

Summary up to now

- First steps of this research
  - a good intuition
  - some simulations validating the intuition for a reasonable range of parameters
- What could have been done more
  - prove that the results is asymptotically (r->0) true "in some sense"
- What can be built on top of this?
  - Markovian models for routing in DTNs

2-hop routing

Model the number of occurrences of the message as an absorbing Continuous Time Markov Chain (C-MC):

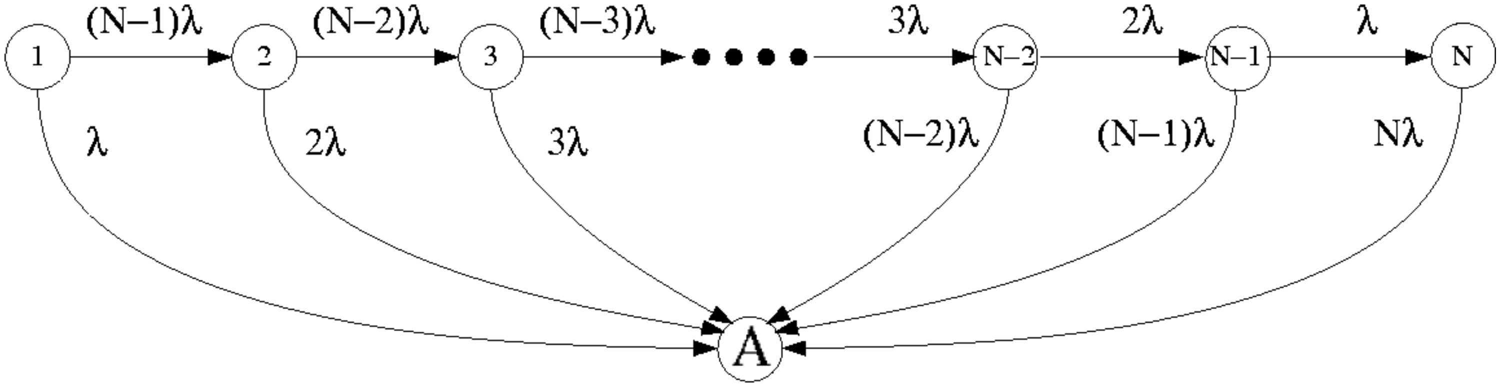

- State $i \in \{1,\ldots,N\}$ represents the number of occurrences of the message in the network.
- State *A* represents the destination node receiving (a copy of) the message.

Epidemic routing

Model the number of occurrences of the message as an absorbing C-MC:

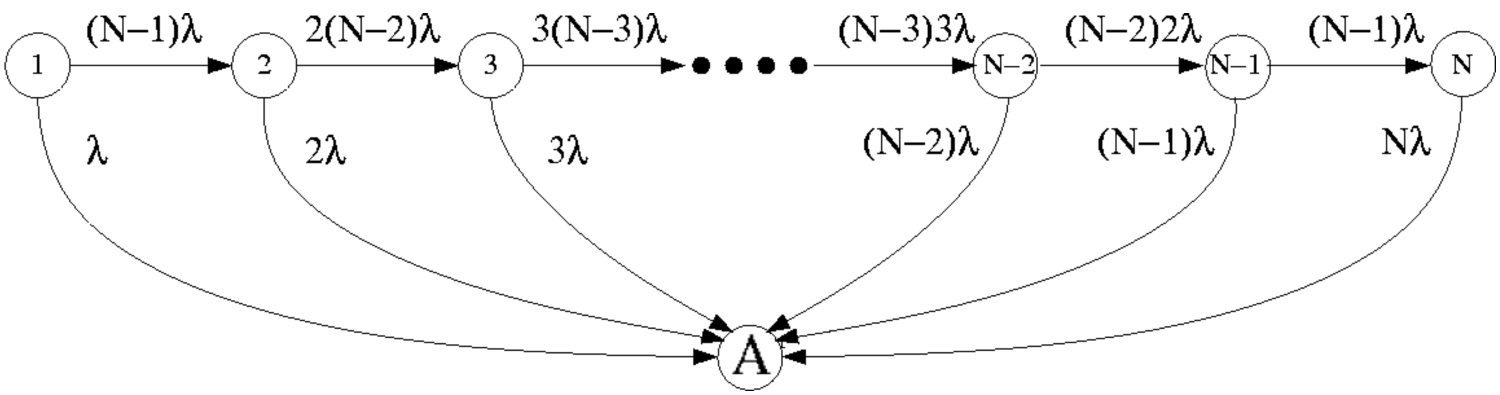

- State i∈{1,…,N} represents the number of occurrences of the message in the network.
- State *A* represents the destination node receiving (a copy of) the message.