Routing Problems in Decentralized Networks

Nicolas Schabanel
joint work with Emmanuelle Lebhar

CNRS - École Normale Supérieure de Lyon
http://perso.ens-lyon.fr/nicolas.schabanel/

Oberwolfach 2004

About "real" networks

Share particular and "usual" properties, such as:

- Usually very large (billions of nodes) with linear number of edges
  $\Longrightarrow$ Individuals only have a partial view
- Small diameter
- Power law distributions (on degree,...)
- High clustering coefficient

Smallworld?

- 1967: Milgram experiment
  - every one is at distance ≤ 7 from every other... and is able to find, based on its local information, a short path to a random unknown person !

The problem

- Explain how in spite of very partial & local view of the network, individuals manage to route efficiently in "undesigned" real network?
- Need some a priori on the network, i.e., some reliable information

Smallworld models

- Watts & Strogatz (1998)
- Kleinberg (2000)
- Other “real” graph models:
  - the configuration model
  - preferential attachment

Watts & Strogatz (1998)

- Smallworld = small diameter & high clustering
- Increase randomness (random rewiring) in a mesh:
  - few randomness: large diameter
  - intermediate randomness: a smallworld (high clustering & small diameter)
  - lot of randomness: no clustering, disconnected
- "Edge of chaos" spirit

Watts & Strogatz (1998)

- Fails to capture the specific nature of smallworlds:

  Ability to route with (very) partial knowledge

Kleinberg (2000)

- An augmented d-dimensional mesh n x … x n
- Underlying mesh = geographical (global) knowledge
- Extra long range link = private (local & random) knowledge
- Local view of the network from a node =
  geographical position + private contacts

Kleinberg (2000)

- Model of individuals: Decentralized algorithm
  - Route only to known contacts either local or long range, i.e. to contacts of already visited nodes.
  - i.e., only use their partial knowledge of the network

Kleinberg (2000)

- Long range contact v of u is randomly chosen with probability $1/d(u,v)^s$
- A simple greedy decentralized algorithm routes in $O(\log^2 n)$ time iff $s = d$!
- If $s \neq d$, any decentralized algorithm takes poly(n) time.
- When $s \leq d$: logarithm diameter
  When $d < s < 2d$: polylog diameter
  When $s > 2d$: polynomial diameter

Questions:

- How to model partial knowledge in smallworld? What is a “good” graph model?
- How to model individual behavior? Is the greedy algorithm the “best” choice?

This question is debated next.

- How would this affect network design?

Algorithms on "Kleinberg-like" networks

- Changing the topology:
  - k long range neighbors:
    - $O(\log^2 n / k)$ greedy algorithm [K00]
  - increase awarness:
    - $O(\log^2 n / k \log k)$ NoN-greedy [CGS02 & MNW04]
    - $O(\log^{1+1/d} n)$ greedy algorithm [FGP04 & MV04]
- New decentralized algorithm:
  - $O(\log n (\log \log n / \log k)^2)$ semi-greedy [LS04]

Algorithms on "Kleinberg-like" networks

- Lower bounds:
  - Whatever the topology is, greedy takes $\Omega(\log^2 n / k \log \log n)$ [ADS02]
  - Any decentralized algorithm visits at least as many nodes as the greedy $\Omega(\log^2 n / k)$ [MNW04]

Our algorithm

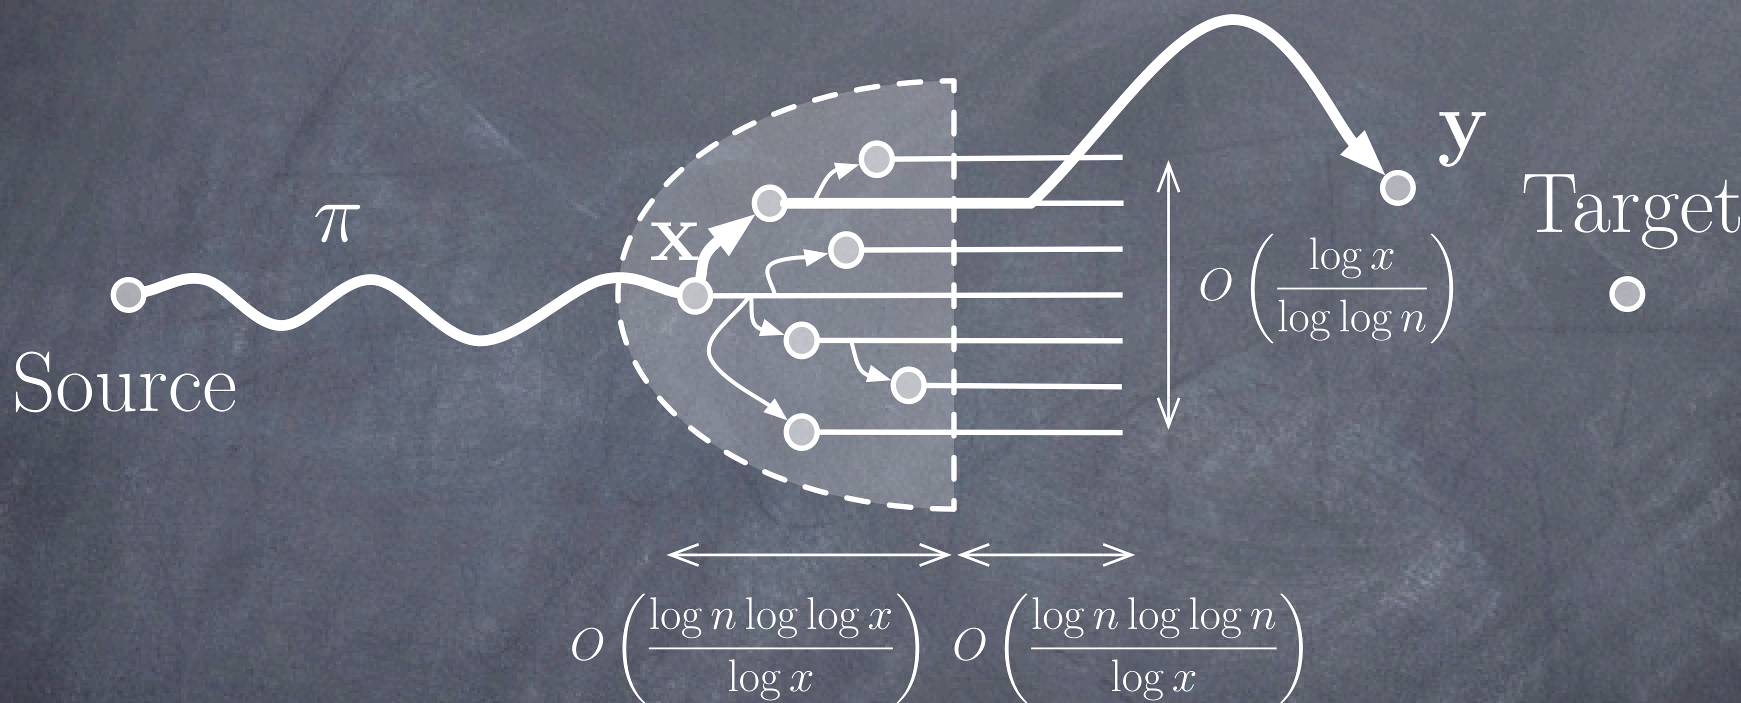

- Expected length of the path =

$$\sum_{i=1}^{\log n} \frac{\log n \log\log n}{\log 2^i} = O(\log n (\log\log n)^2)$$

Open problems

- Bias of the measure of the network?
- Should algorithms ignore the particular nature of real graphs? Would it change considerably their design?