Critical search diameter in distributed information systems

Cedric Westphal

Nokia Research Center

Mountain View, CA

Basic Set-Up

- ad hoc networks consisting of nodes with limited capacity to store information.
- distributed networks with no centralized management nor infrastructure
- Each node has constraints which prevent a node from maintaining exhaustive information:
  - due to storage costs
  - due to volatility of the information
  - due to on/off duty cycles
  - due to power costs
- On the other hand, network as a whole requires the information to be stored somewhere to protect the integrity of the network.

Example: service discovery

- A node requests the location of some service in an infrastructure-less network.
- Service mappings stored in service table
- Service location is requested using an extension of AODV: a service request is a route request with a service option.
- Service mapping is given a TTL value.
- Once the lifetime has been reached, entry is flushed from the table.

Service Discovery

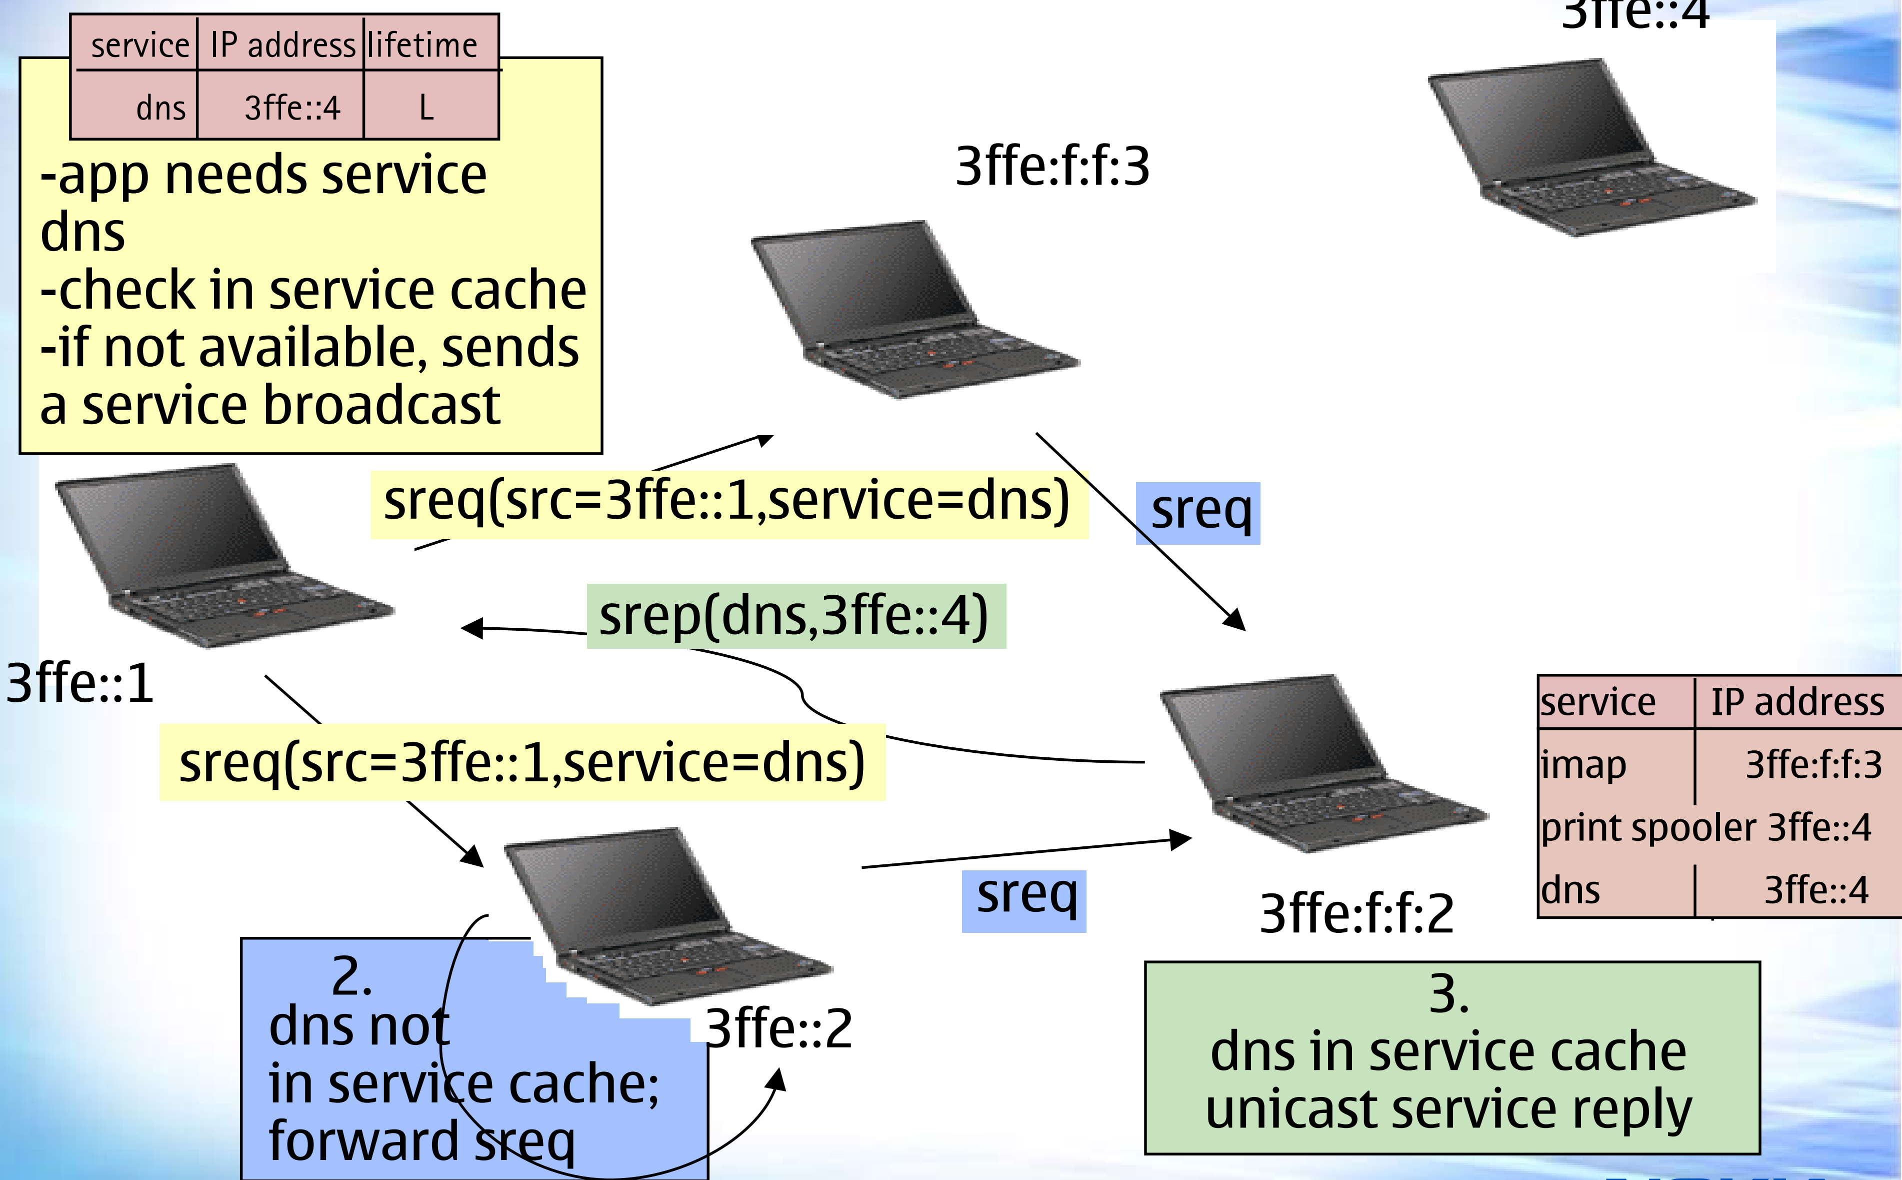

Example: peer-to-peer network

- Consider a peer-to-peer overlay, a la kazaa, where each participating nodes forwards traffic to and from other nodes.
- Distributed copies of a data object are replicated in the nodes participating in the p2p network
  - for instance, mp3s files
- Each node can only contain a partial view of the data objects.
  - my ipod = 4GB memory
  - my music files = 20 times this amount.
  - as I download new files, I flush out older files: each file stay only a limited time on the device.
- Network connected nodes can participate in the network (cf. Motorola itune phones)
- For people to participate in the network, the network must be able to provide them with files which are more rare: incentive for the network to keep copies of all files.

Example: peer-to-peer

"viral" transmission of games on gaming platforms such as Nokia NGage (or Sony PSP)

Example: Sensor Network

- Limited capacity: distributed storage over the network
- Information stored in a mote can be flushed when the mote turns itself on and off.
- Some information critical to the network (redundancy required)
- Broadcast mechanism preferred communication mode when addressing limited/not available.

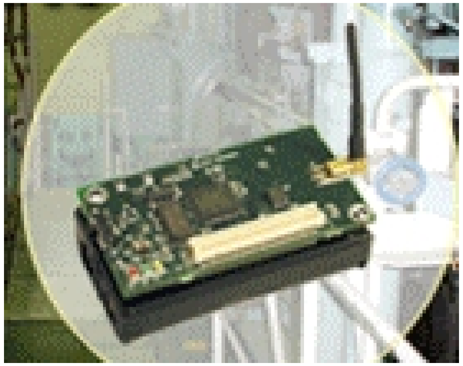

NOKIA

Information Location Protocol

- Information is distributed and not permanent: how to locate the information.
- Simplest way: a request reply mechanism similar to AODV.
  - Send an object request, have other nodes relay the request.
  - If a node is able to respond to the request, then they issue a reply containing the requested object.
  - Intermediate nodes do not access the traffic they relay (typically request at application layer, forwarding at network layer), only end nodes get the object in the process.
- Other mechanisms possible (DHT for instance) but are not considered here.
- AODV parameter: max_hop_count, ie. the number of times a node can be relayed. This sets a maximum value D for the search diameter. Goal: localized search, extended ring search, prevent infinite routing loops

Two competing mechanisms: dissemination and disparition

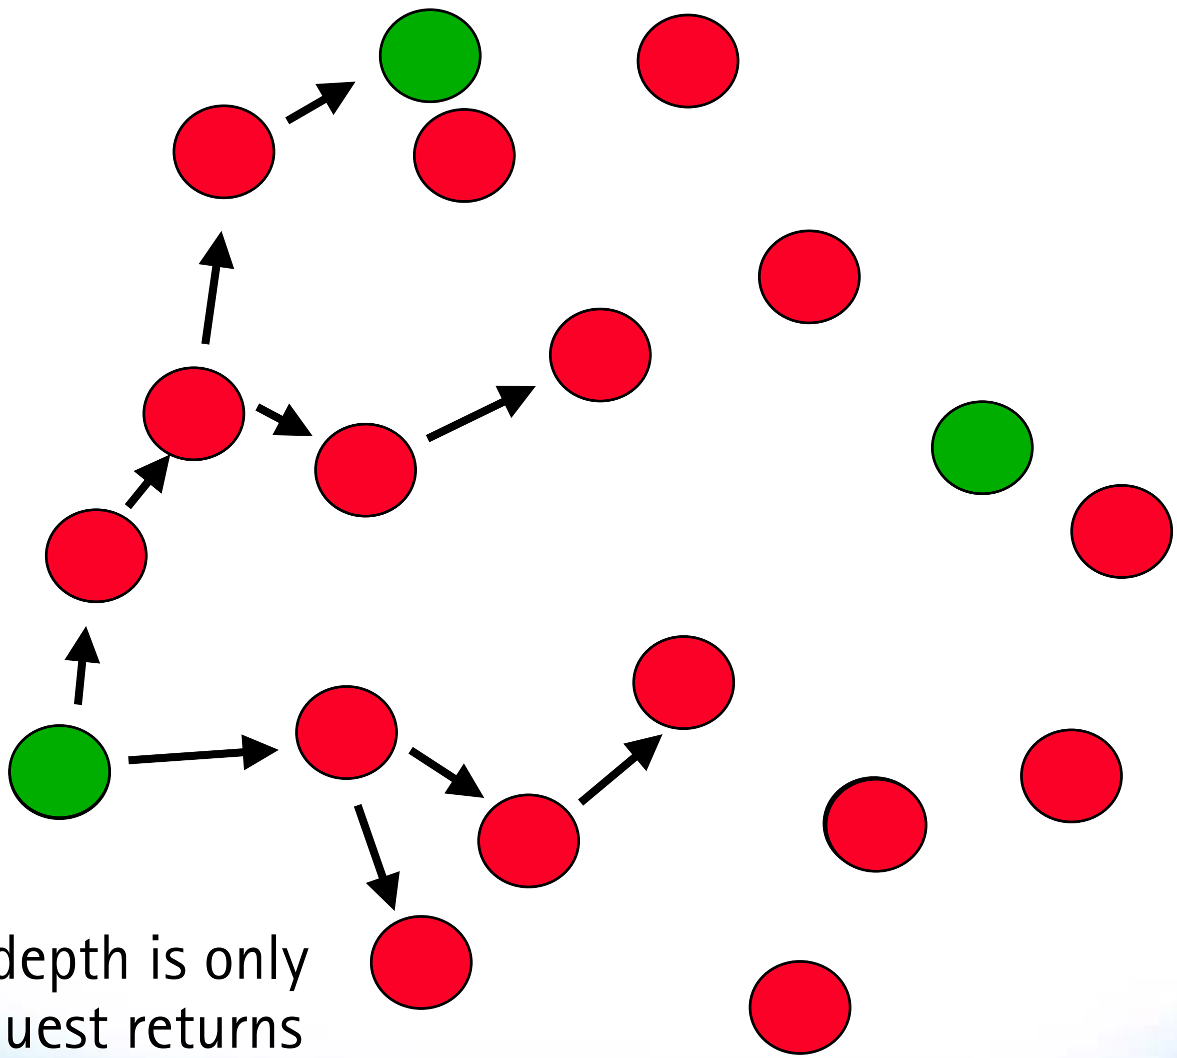

if search depth is only 3, the request returns no result

Model

- Consider a network with only one replicated object.
  - Several objects can be added independently.
- Nodes are distributed according to a Poisson boolean model. We assume for anything to make sense that it is supercritical.
  - if not, lifetime is finite a.s.
- Each node that does not have the object will request it with rate λ.
- The request travels D hops. If a requests finds a node with the live object within D hops of the requestor, then requestor gets the object. X(node, t): 0 -> 1.
- The object disappears from each node at rate 1. X(node, t):1->0.
- This is a first step, so consider infinite graph and survivability only.

Basic Questions:

- What is the relationship between λ and D so that at least one copy of the object stays alive indefinitely?
  - in practice, λ is set a priori and not a parameter.
  - in some settings, lifetime can be increased by say increasing storage (reflected in λ).
- Can one maintain copies of the object alive indefinitely at a "reasonable cost" to the network.
- Answers are relatively obvious…
- Survivability with some probability: define $\tau = \inf\{t: X=0\}$
- We are interested in the case $P(\tau = \infty) > 0$.
- τ depends on the initial conditions.
- Initial condition: we seed the system with a finite set of nodes for which $X(node, 0)=1$.

Related Work.

- Contact process (Harris, 1974) on a lattice.

green: infected

red: cured

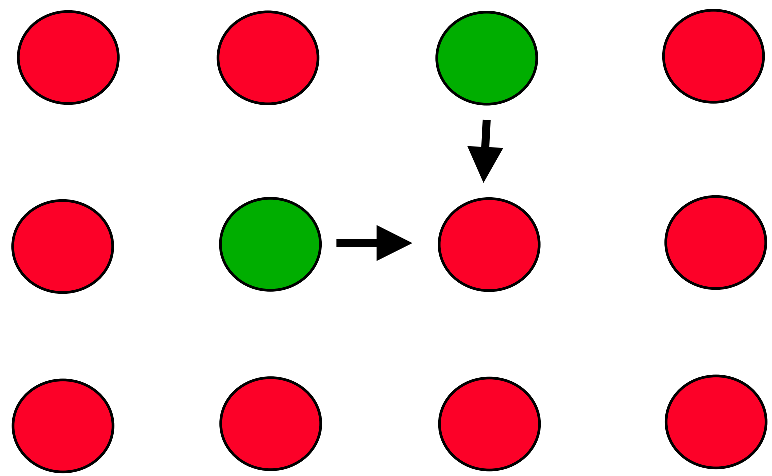

•infected nodes infect their neighbors at rate λ and cure themselves at rate 1

•similar but push instead of pull, and nodes add up their effort to infect.

NOKIA

Is there a critical diameter D_c for survivability?

- Necessary condition on the diameter for the dissemination process to survive.
- At least would tell us that there is no hope for D less than D_c
- Answer is true in the lattice
  - proof in MobiHoc'05 paper.
  - basic idea is to couple the dissemination process on the lattice to a contact process. This entails a relationship between D_c and the critical rate in the contact process.
- Answer is true in the homogeneous tree
  - but two different rates!
- Here: extend this to the random graph.

Poisson boolean model

nodes distributed with intensity ζ

ζ and R such that supercritical

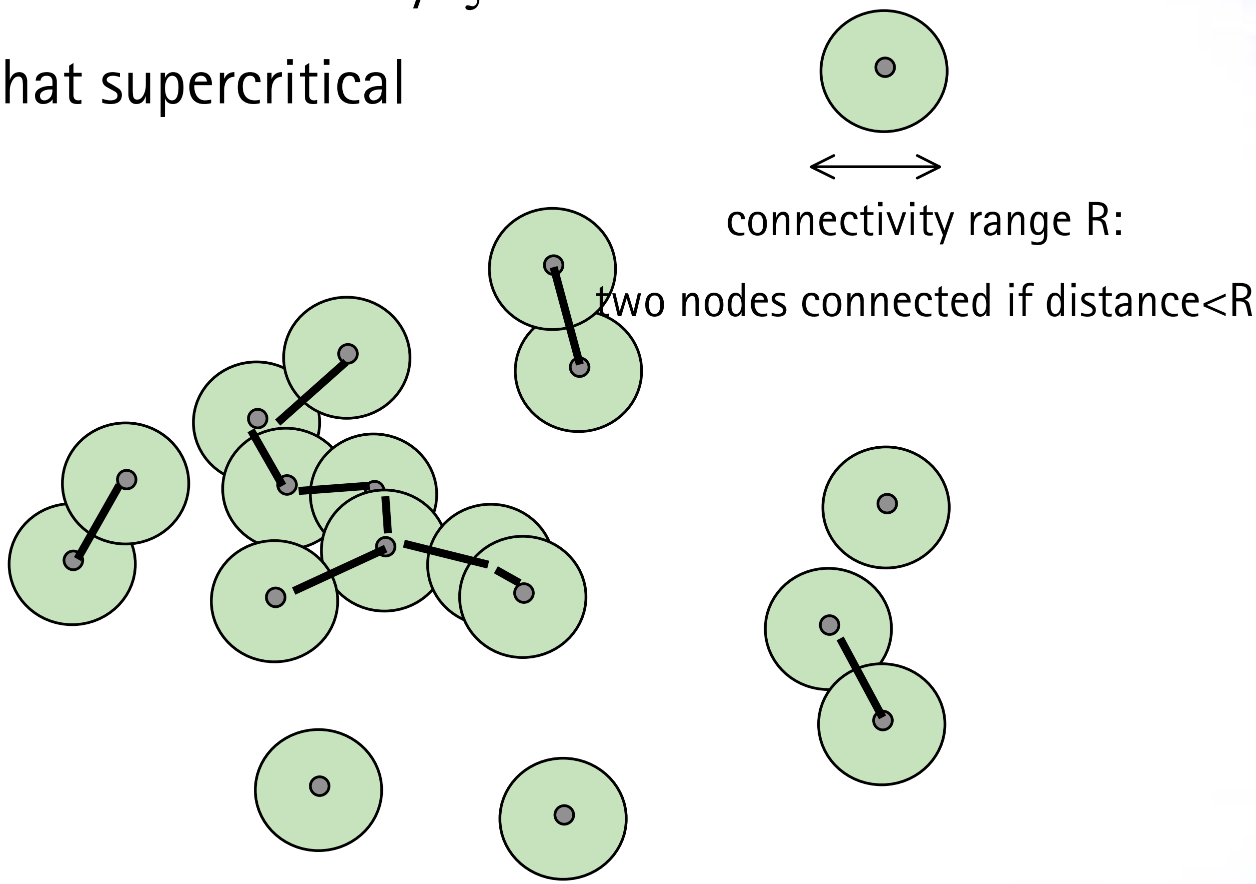

Existence of a critical diameter in the Poisson boolean model

- Theorem: For a fixed λ, there exists a critical diameter D_c such that, for $D > D_c$, the dissemination process with search depth D satisfies:

$$P[\tau = \infty] = \theta > 0$$

Existence of a critical diameter in the Poisson boolean model: proof

- Idea: create a lattice structure over the random graph by considering squares such that there is a "cross of Lorraine" connected path over two neighboring paths (cf. Dousse et al, MobiHoc 04).

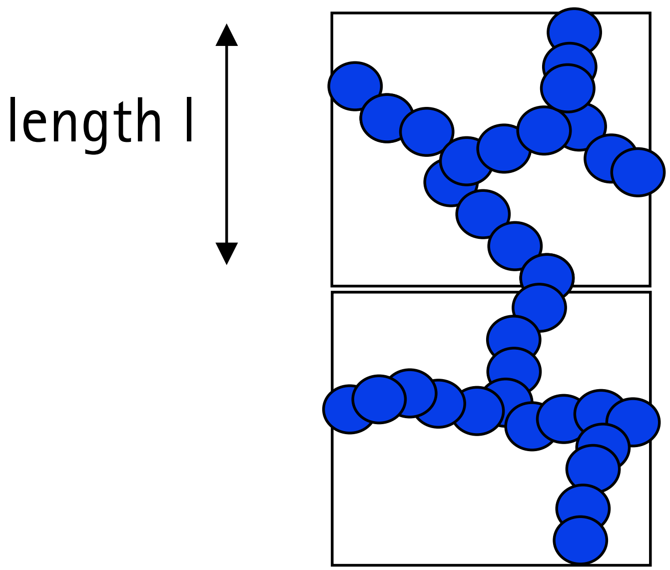

- For l large enough, one can map the random model to a bond percolation on the lattice with parameter p.

Existence of a critical diameter in the Poisson boolean model

- On the bond percolation model, one can find a critical diameter Δ depending on p and l.
- The number of hops can be bounded above by Δ times the maximum possible number of hops within each square, which is bounded above by O(l^2).
- Taking a D bigger than these values ensures survivability of the dissemination process in the original random graph.

Cost of Survivability

- We considered the dumbest mechanism to ensure survivability.
- How efficient? Define the cost as the number of sent messages per node per unit of time.
- Our dissemination process gives us an upper bound on the cost to maintain the information alive in the network, namely:

$$\lambda(\zeta \pi R^2)^D$$

- How much can we improve on it?
  - If we only care about survivability (at the expense of each node accessing the information), then we can improve it a lot!

Dissemination Cost

- We use the square lattice as previously with squares of side l.
- One node in the infinite cluster in each square we designated as a store-and-request node
  - only nodes that are allowed either to store the object, or to disseminate the request by issuing a request or a reply.
  - nodes on the cross of Lorraine path are forwarding nodes, and relay the messages between the store-and-request nodes.
  - other nodes are idle.
- One can imagine a mechanism so that each node alternate roles within the three types above.
- The cost can be reduced to less than ε for any $\varepsilon>0$.

Dissemination Cost

- Theorem: Define the cost c of maintaining the information alive to be the number of messages transmitted per node per unit of time. For each $\varepsilon>0$, we can construct a dissemination mechanism such that

$$c < \varepsilon$$

NOKIA

Reduce the cost?

- We only need to see that the number of relaying nodes vs. the total number of nodes in a square is decaying as l increases
- One can show (similarly as in Franceschetti 2005) that each crossing path in the square can be contained in a l x log(l) rectangle.
- In each rectangle, O(log(l)l) relays at most, vs. $O(l^2)$ in the square, thus the decay of the dissemination cost.

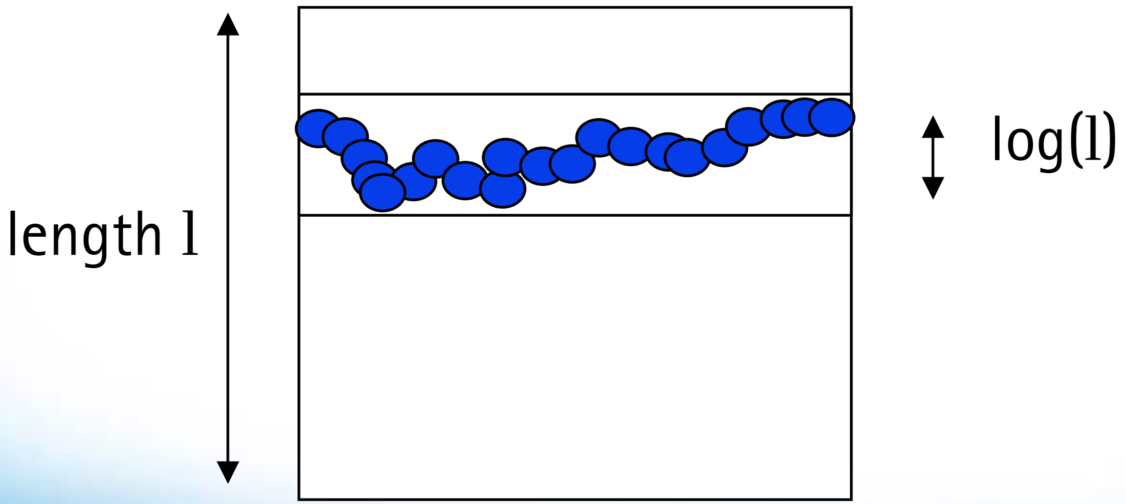

Conclusions

- Preliminary results, but so far few results considering distributed information with finite lifetime.
- Cost saving mechanism is crude, hierarchical schemes may permit to keep all nodes involved at lower cost.
- Infinite graph assumption unrealistic, it would be interesting to study the asymptotic behavior for a large network of size $N \to \infty$
- First order results. How to improve survivability by control/distributed network management?
- Questions?