Local Failure Detection in Mobile Distributed Systems

Nigamanth Sridhar
Electrical and Computer Engineering
Cleveland State University

Collaborators: Hamza Zia (CSU), Shiva Sastry (Akron)

Context

Application Area: Wireless Sensor Networks

Goal of WSN: Dense instrumentation of physical world with sensors, actuators.

Consequences:

- Each node has to be *extremely cheap and dispensable*.
- Failure is the norm, not an exception

Failure Locality

Failure cannot be avoided. Next best strategy is to localize the effects of a failure.

Failure Locality measures distance (in hops) at which the effect of a failure is "felt".

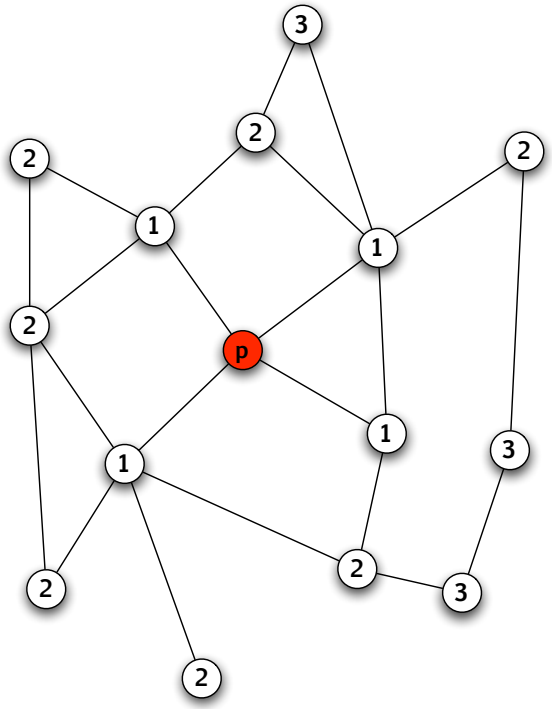

Failure Containment

Basic idea: Failed node's neighbors "detect" the failure, and quarantine the failure; the rest of the network is protected.

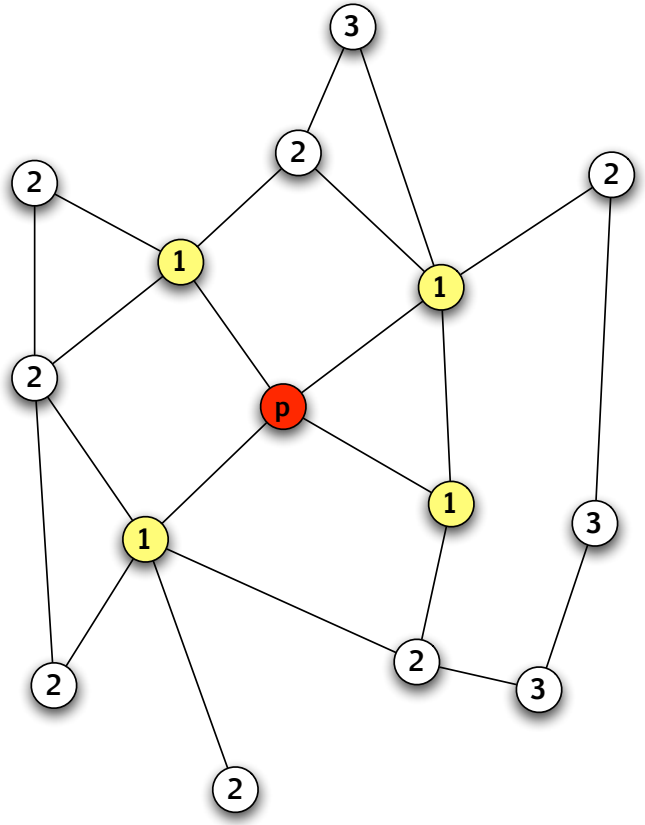

How Failure Detectors Work

Basic idea: p waits for some timeout period for some communication from q, and then begins to suspect q.

Several strategies:

- Simple timeouts
- Adaptive timeouts
- Ping + timeout
- Leases

Note that suspicion may not be well-founded!

Example: Dining Philosophers

Problem of resource allocation in a graph.

Each node shares some resource(s) with neighbors. In order to enter critical section, must have permission from all neighbors.

Specification:

Safety: No two neighbors eat simultaneously
Progress: Every hungry node eats eventually

A Solution: Hygienic Dining

1. Every edge in graph represents a fork
2. A node has to have all forks to eat
3. Priority established by partial order

Problem: Poor Failure Locality

Failure Locality measures how far in the graph a failure is "felt".

Hygienic Dining has failure locality d.

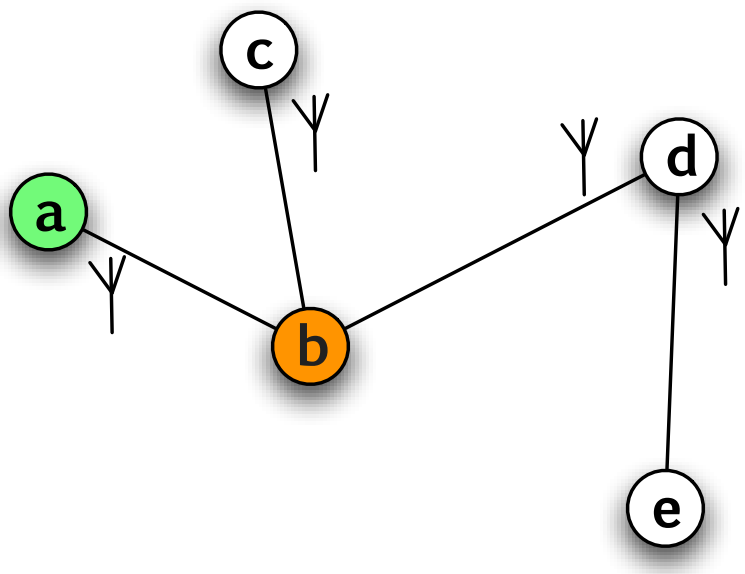

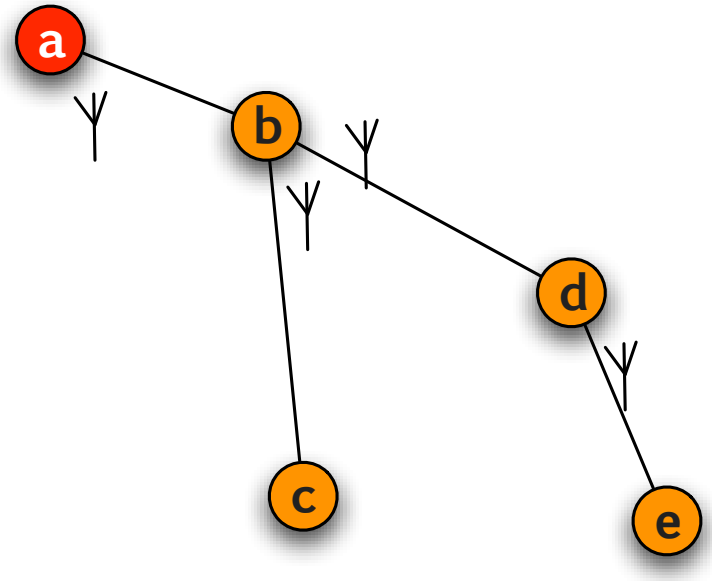

Local Failure Detection to the Rescue

If a node suspects one of its neighbors (using a local failure detector), it *"shields"* the rest of the network from this failure.

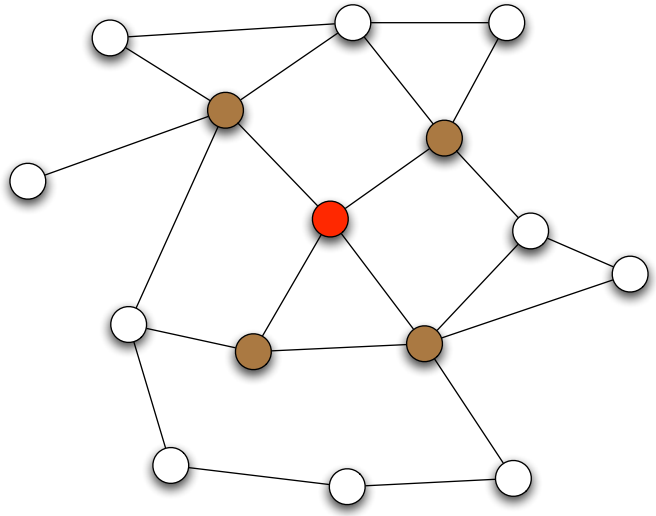

Failure locality is 1.

Dynamic Topologies

But here's a new problem: change in network topology!

A local failure detector cannot distinguish between a failed neighbor and a node that's no longer a neighbor; mistakes happen!

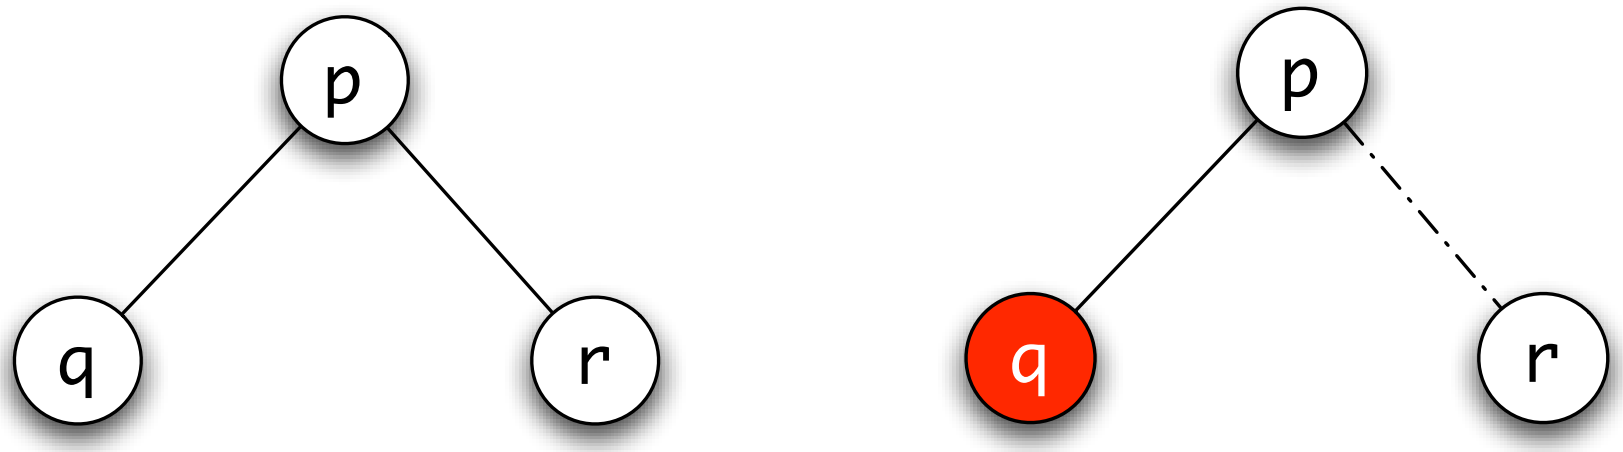

Mistakes are Expensive

A single mobile node nullifies the transformation.
Failure locality is back to d.

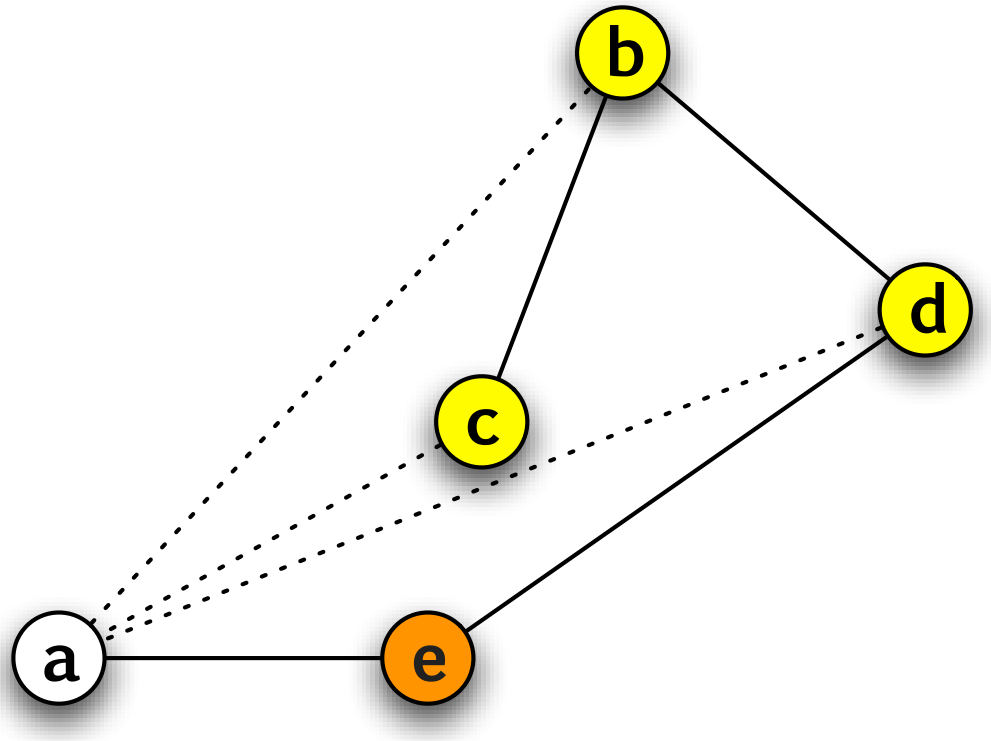

Solution: Share Information

Our approach: once in a while, each node in the network lets others in the network know about who it suspects currently.

$\Diamond\mathcal{P}_{\ell}^{m}$ — Eventually Perfect Local Failure Detector that tolerates mobility.

Design of $\Diamond\mathcal{P}_\ell^m$

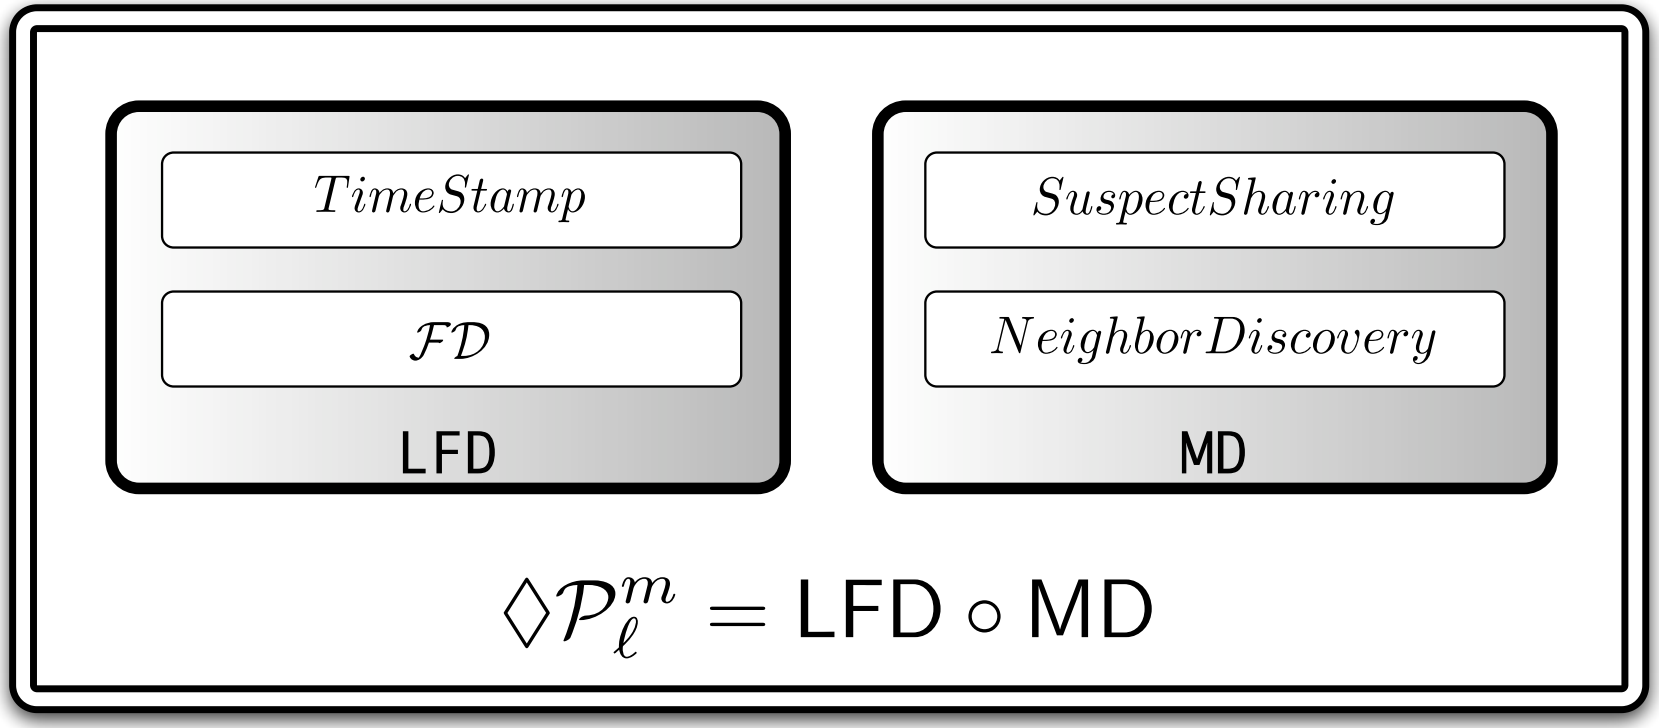

We won't talk about LFD today. Rest of the talk is focussed on $SuspectSharing$ (MD).

$SuspectSharing$ Algorithm

Initiator sends out suspect group ($\mathcal{SG}$) to its neighbors to begin a *diffusing computation*.

Suspect group contains:

- The set of suspects x that the process u maintains ($\mathcal{S}_u$)
- The durations (ts_x) for which each process x has been suspected
- The id of the process that suspects x (denoted by σ, u in this case)
- The number of hops for which this process x has been in $\mathcal{SG}$ (denoted by d_σ, 0 initially)

SuspectSharing Algorithm (contd.)

When a node w receives $\mathcal{SG}$, it looks at it, and if it finds any live nodes, w exonerates them.

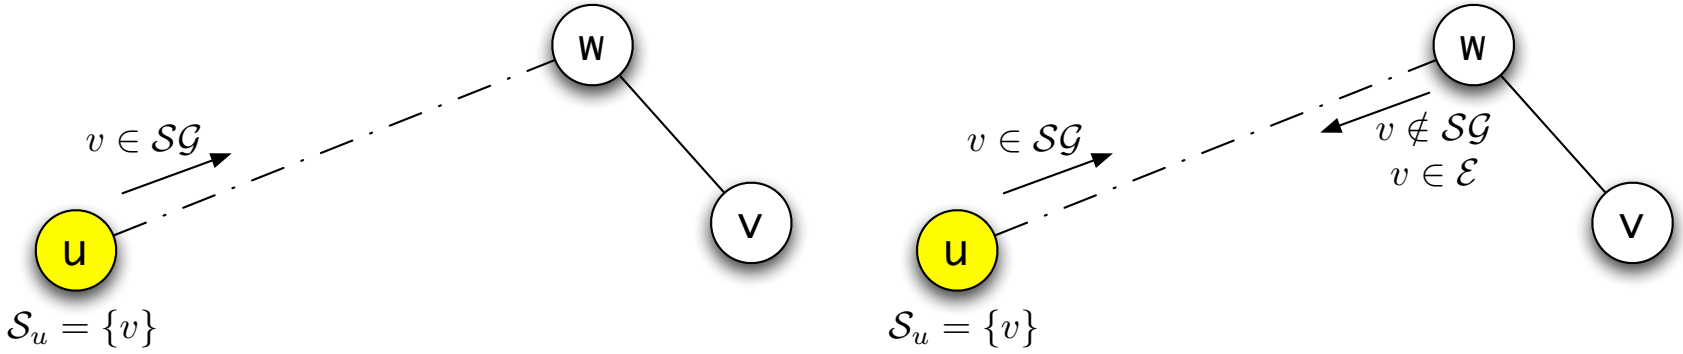

SuspectSharing Algorithm (contd.)

Each node also adds its local suspects to $\mathcal{SG}$.

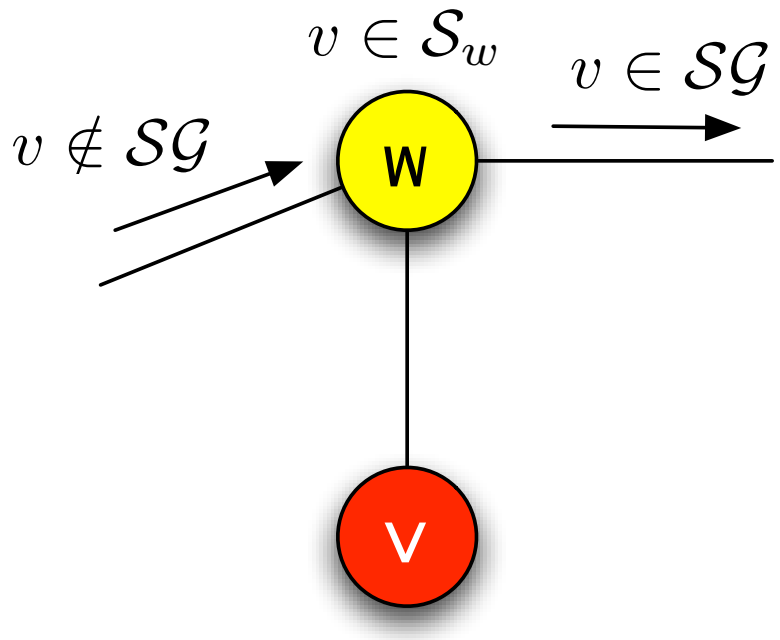

SuspectSharing Algorithm (contd.)

In the shrinking phase of diffusing computations, nodes "correct their view of the world".

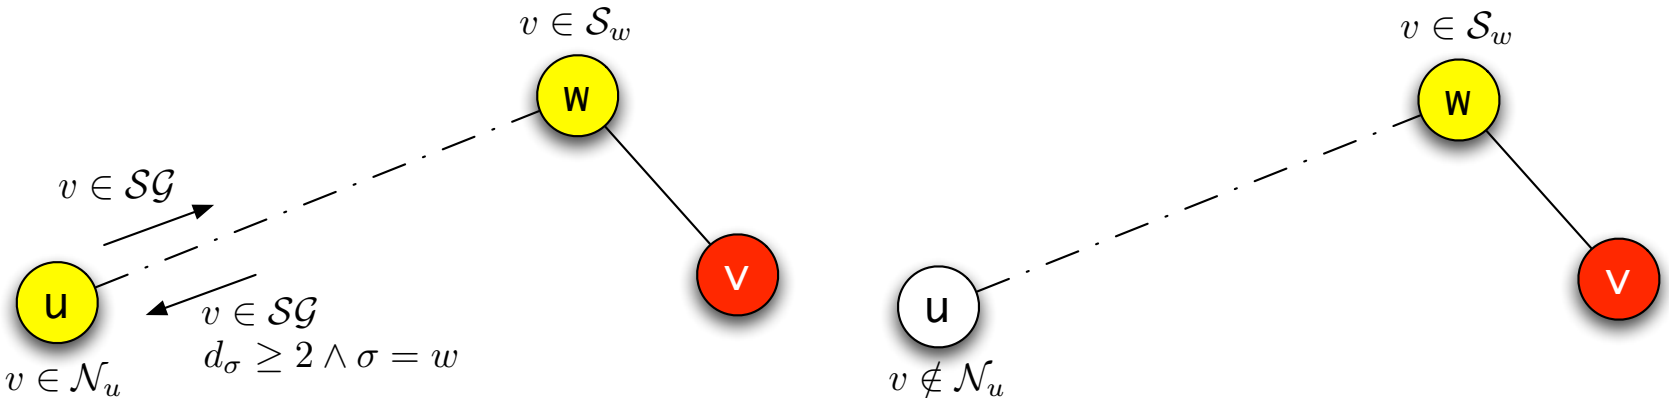

Implementation

Implementation for "motes" (TinyOS/nesC).

Component is implemented as a middleware service that applications can use.

Local failure detection uses Lease strategy.

Parameters to the failure detector:

- Average lease duration
- Gossip recurrence time

Average Message Overhead/Node

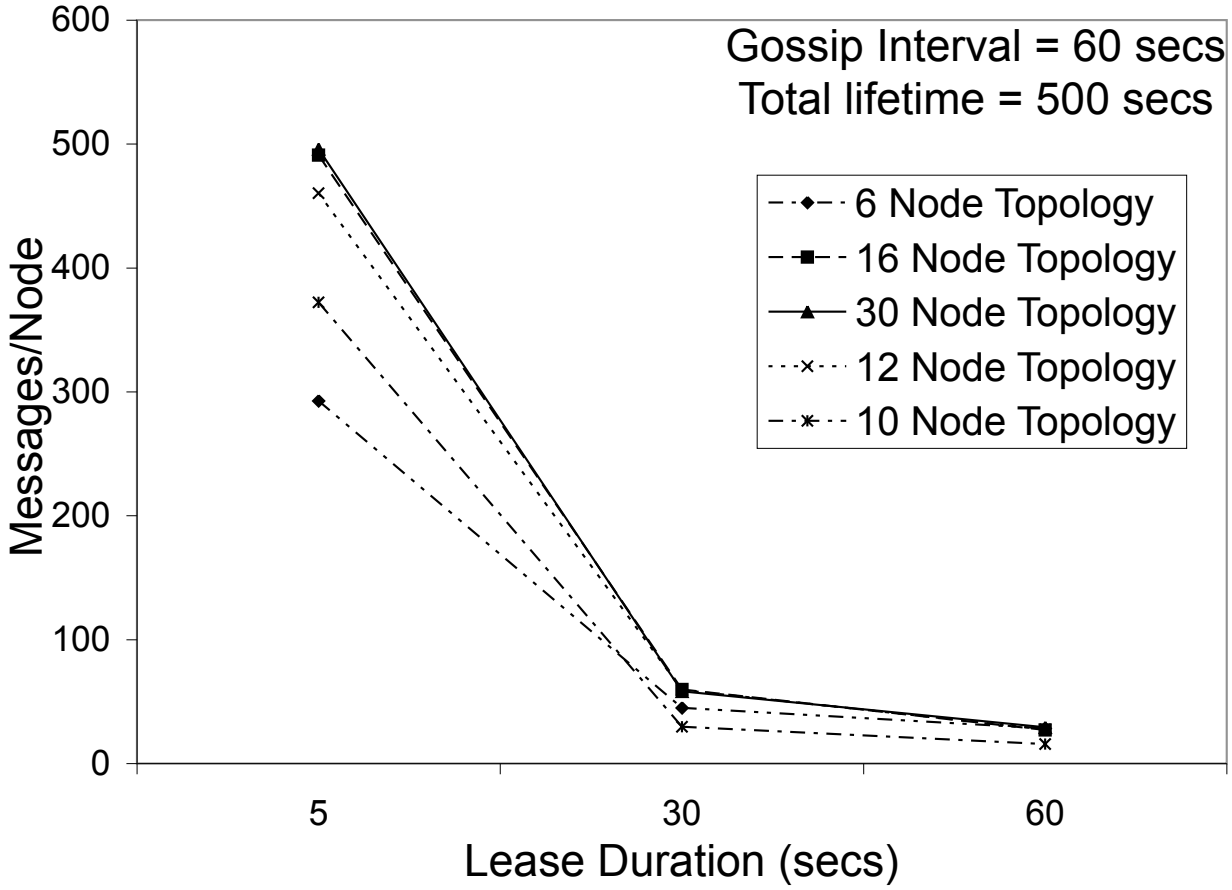

Message Overhead in Entire Network

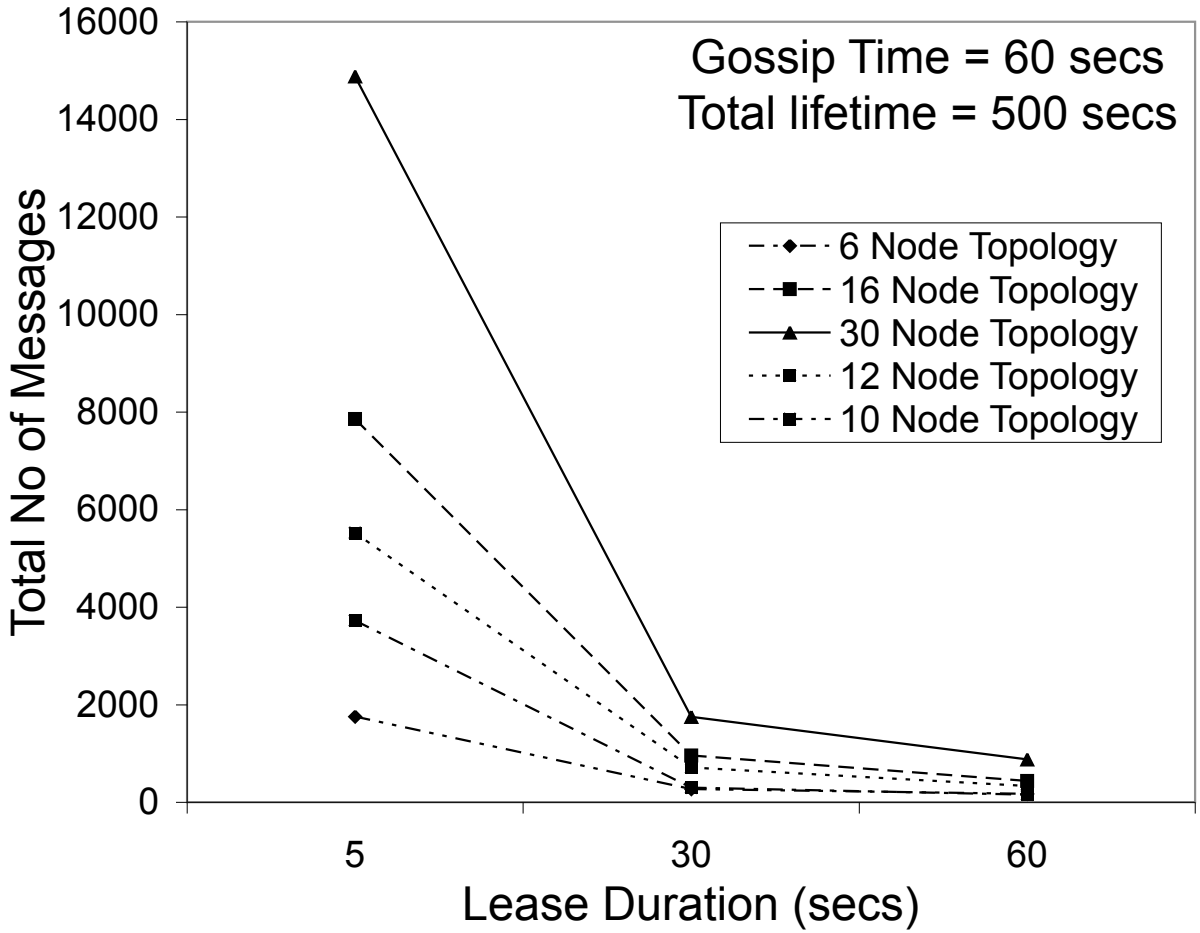

Mistake Duration

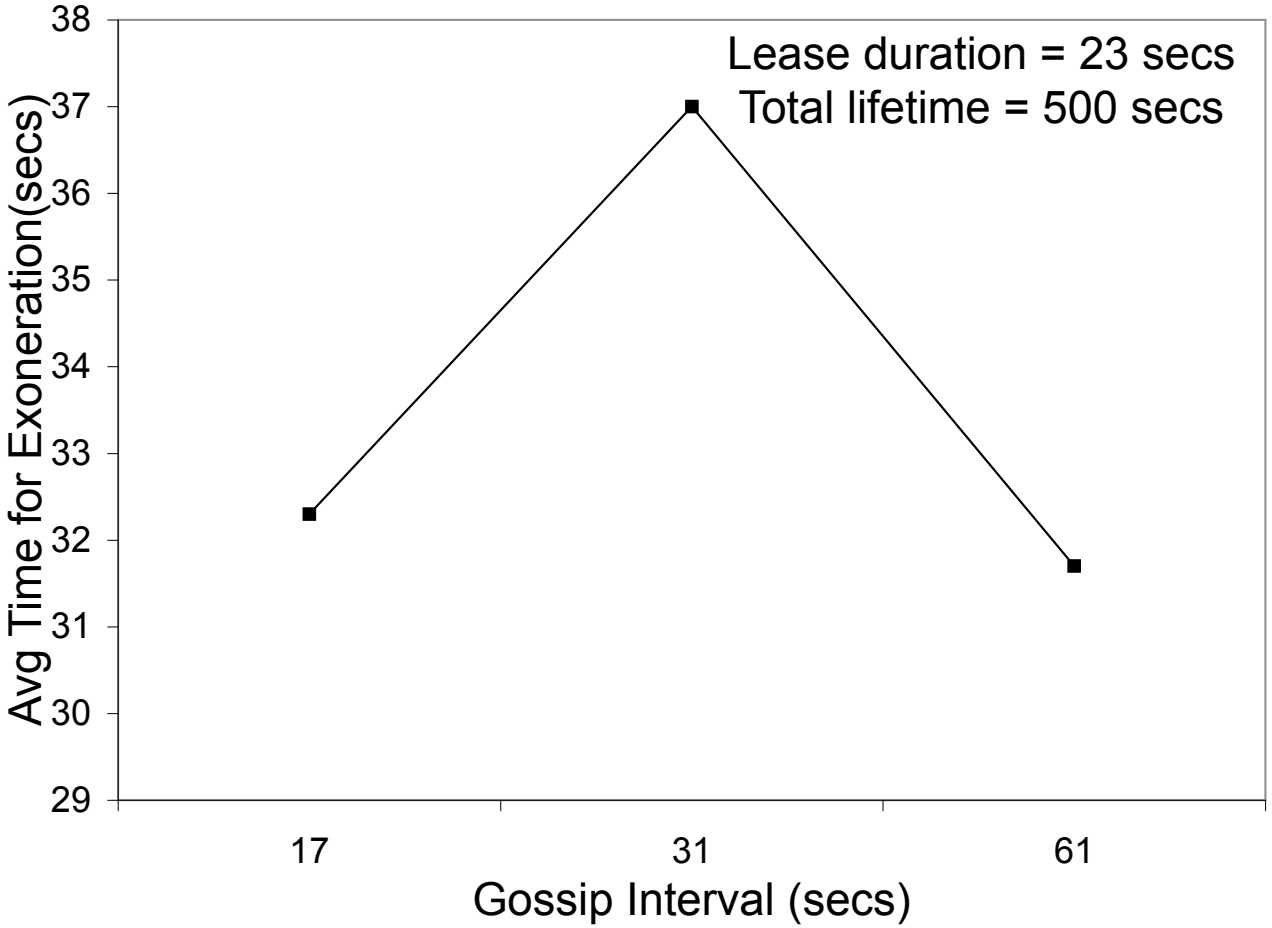

Summary

- Local failure detection is important in sensor network context
- Dynamic topologies are a reality
- $\Diamond\mathcal{P}_{\ell}^{m}$ performs similar to a local failure detector, and functions like a global failure detector