UCL

GZ06 Adaptive and Mobile Systems
Dr. Cecilia Mascolo

Semi-Probabilistic Content-Based Publish-Subscribe

Paolo Costa, Gian Pietro Picco

Politecnico di Milano, Italy

[slides from the ICDCS presentation of the authors]

UCL

Publish-Subscribe Middleware

- Applications publish asynchronous event notifications and declare their interest in event classes by issuing a subscription
- Subscriptions are collected by an event dispatcher component, responsible for routing events to matching subscribers
- High degree of decoupling among components

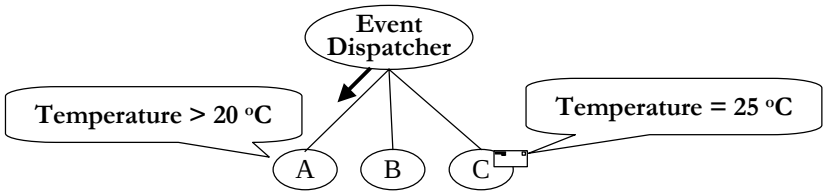

Motivation of this Paper

- Publish/subscribe systems for mobile ad hoc networks suffer from scalability problems
- However the ability to be able to publish and notify hosts of events despite mobility is very useful

Distributed Content-based Publish-Subscribe

- **Content-based:**
  - Subscriptions contain expressions (event patterns) that allow clients to filter events based on their content
  - More difficult to implement than *subject-based,* where the set of subjects is determined a priori
- **Distributed:**
  - Tree-based routing (e.g., subscription forwarding):
    - Every dispatcher propagates subscriptions to the others
    - Events follow the routes laid by subscriptions

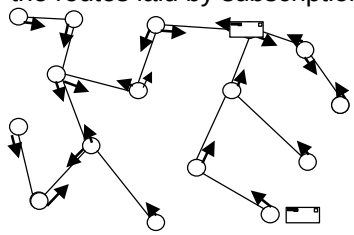

Content-based routing on a dynamic topology

- The pub/sub paradigm is appropriate for dynamic environments
- Unfortunately, most current systems do not tolerate topological reconfiguration
- Topological reconfiguration of the dispatching infrastructure can occur in:
  - peer-to-peer: nodes connect or disconnect unexpectedly
  - MANETs: physical mobility of hosts
  - sensor networks: sensors switching on / off

Semi Probabilistic Approach

- Deterministically (propagate and log subscriptions):
  - ✓ Routing is optimal
  - × How to avoid loops on a graph?
  - × Deterministic information get obsolete very soon in dynamic environments
- Probabilistically (epidemic):
  - ✓ No state required
    - resilient to reconfiguration
  - × Propagation extends to all the network
  - × How to limit the infection ?

The Idea

- A semi-probabilistic approach combining:
  - the efficiency of deterministic routing
  - the resilience to reconfiguration and inherent simplicity of probabilistic approaches
- Deterministic information is maintained only in the proximity of the subscriber
- Routing is performed:
  - deterministically (if possible)
  - probabilistically (otherwise)

Subscription Propagation

- The *subscription horizon* parameter f determines how far a subscription is propagated

✓ Subscriptions propagated only close to the subscriber
  - Loops are less likely to occur
  - Lesser obsolete information

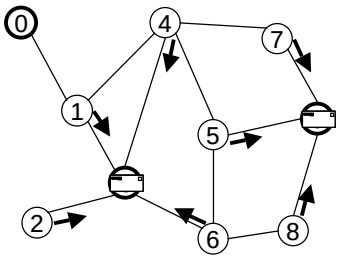

Publisher
Subscriber
Subscription
f = 1

Event Propagation

- An event received is always forwarded only along a subset of the outgoing links (*fanout*)
- The selected links may change from time to time but the percentage of links selected is fixed and is determined by the *event propagation threshold* (t)
- How links are selected ?
  - deterministic links are selected first
  - the remaining ones are chosen probabilistically

Event Propagation

- t=0.5 means each event is forwarded on the half of outgoing links
- Obviously, an event is never forwarded on the same link it comes from

✓ Low number of infected nodes

- propagation is limited by t and integrated with deterministic information

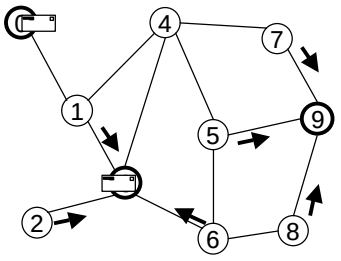

Publisher

Subscriber

Infected with event

f=1 (subscription horizon)

t=0.5 (propagation threshold)

Route Selection

- Independently from the available information, the propagation threshold cannot be passed
- Subscription selection is prioritized according to f:
  - subscriptions at distance *d* are preferred over those at distance *d + 1*
  - subscriptions at *d = 1* are always used (no matter what the propagation threshold is)
- An event is never forwarded to a leaf dispatcher unless it is a subscriber

Dealing with Dynamicity

- Unsubscriptions are handled similarly to subscriptions
- Additionally, they are used to handle topology changes as well
  - When a new link appears, the dispatcher sends a subscription message along it
  - When it vanishes, the dispatcher behaves as if it received an unsubscription message for all subscriptions associated with that link

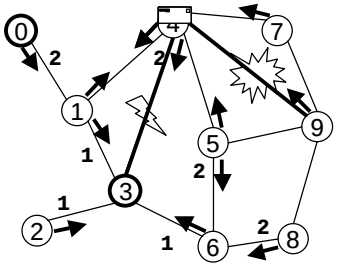

Simulation Settings

- We evaluated the performance of our approach in two different scenarios
  - Overlay Network Scenario:
    - Controlled graph with reconfiguration (an existing link vanishes and a new link is provided elsewhere) each *0.03s*
  - Mobility Scenario:
    - Planar graph with nodes moving using *Random Direction Model*
- The metrics we analyzed are:
  - event delivery (received events / expected events)
  - overhead (event and subscription messages)
- Simulations were developed with *OmNet++*

Overlay Network Scenario

Network Size	50—400
Graph Degree	5±1
Publishing Load (per dispatcher)	5 publish/s
Receivers	10
Available Patterns	512
Patterns per Node	2

- An existing link vanishes and a new link is provided each *0.03s*
- We use combinations of the following parameters:
  - subscription horizon f:
    - 0 (*purely probabilistic*), 1, 2, 3
  - event propagation threshold t (fanout = $t \cdot 5$):
    - 0.25 (fanout = 1), 0.5 (fanout = 2), 1 (*flooding*)

Event Delivery

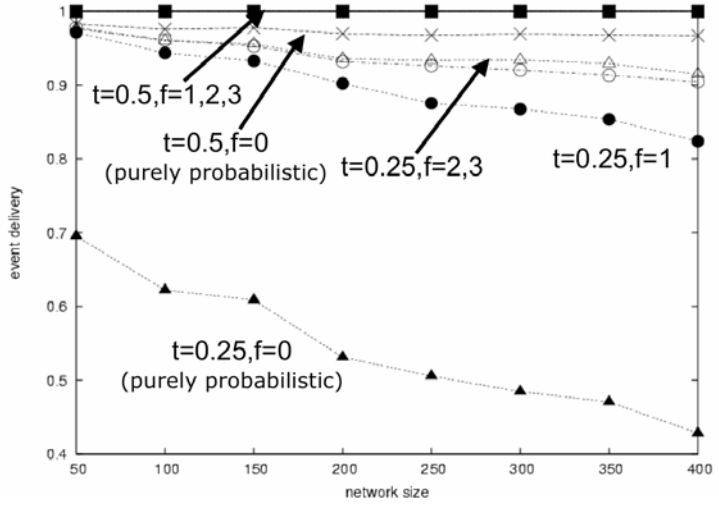

Overhead

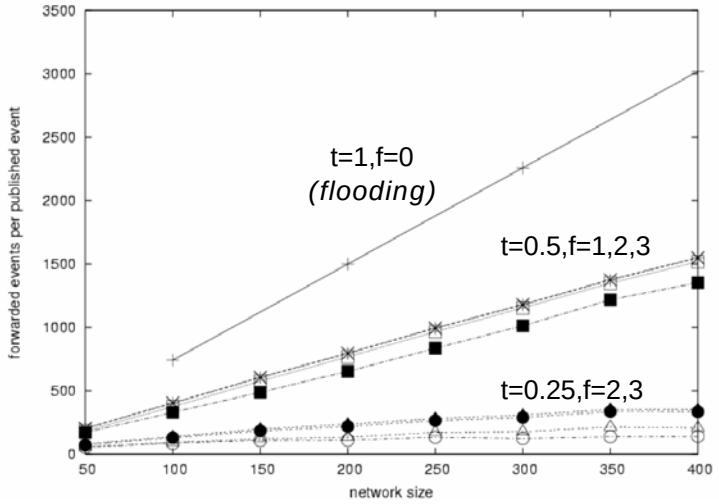

Mobility Scenario

Network size	200
Area	4 Km^2
Wireless range	250 m
Velocity	2 m/s
Publishing Load (per dispatcher)	5 publish/s
Receivers	20
Available Patterns	512
Patterns per Node	2

- Scenario was generated using a MANET topology generator (*ANSim*)
- The mobility model we choose is *Random Direction Model*, where nodes move along a direction until they reach one of the edges.

Mobility Scenario

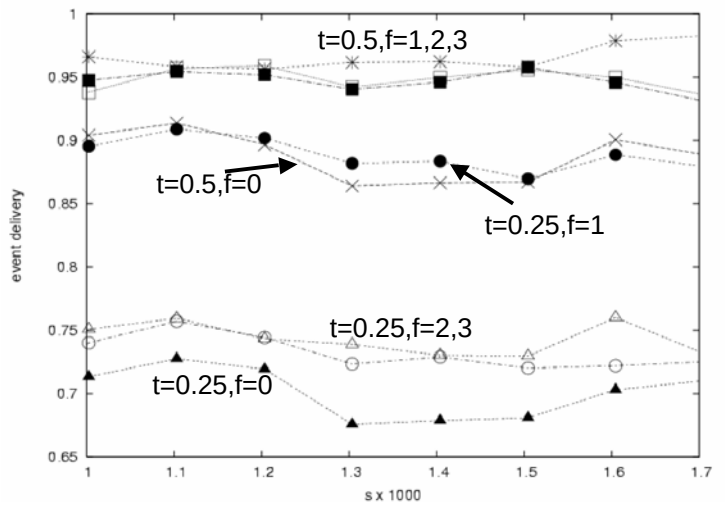

Conclusions and Future Work

- Combining deterministic and probabilistic routing provides high event delivery and low overhead, without sacrificing scalability
- The resulting algorithm is very simple
  - small footprint software
  - easy deployment on resource-constrained devices
- Applied it also to sensor networks scenario (techrep. available at **http://www.elet.polimi.it/~picco**)
  - Different issues:
    - *wireless broadcast transmission*
    - *network collisions*
    - *sleeping nodes*
- Our ultimate goal is to devise a mechanism that adapts to network conditions, by properly modifying t and f

Future work

- Applied to sensors:**Publish-Subscribe on Sensor Networks: A Semi-probabilistic Approach"** with Paolo Costa and Silvana Rossetto. In *Proceedings of the 2nd IEEE International Conference on Mobile Ad Hoc and Sensor Systems (MASS'05)*, November 7-10, 2005, Washington DC (USA).
- Extensions for DTN:
- **Adaptive Routing for Intermittently Connected Mobile Ad Hoc Networks** Mirco Musolesi, Stephen Hailes and Cecilia Mascolo In Proc. of IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WOWMOM05). Taormina, Italy. June 2005.
- Use a different Mobility Model:
- **A Community Based Mobility Model for Ad Hoc Network Research** Mirco Musolesi and Cecilia Mascolo UCL Research Note RN/05/31. Submitted for Publication. December 2005.