GZ06 Adaptive and Mobile Systems
Dr. Cecilia Mascolo

Predictive Publish/Subscribe for Delay Tolerant Mobile Ad Hoc Networks

P. Costa Vrjie Universiteit, NL

C. Mascolo UCL, UK

M. Musolesi Dartmouth College, USA

G. P. Picco, Trento University, Italy

DTNs: What ?

- Network with intermittent connectivity
- Data are disseminated exploiting host mobility
  - long and variable delays
  - high error rates
  - heterogeneity

DTNs: Why ?

- To cover regions with no infrastructure
- Networks with pedestrians
- Intelligent highways
- Emergency rescue
- Sensor networks
- V[illegible]g

DTNs: How ?

- Delay Tolerant Networks (DTNs) introduce novel challenges to system designers
- High decoupling among hosts is mandatory
- Publish/Subscribe recently emerged as a promising approach in this domain
  - the data consumer (*subscriber*) specifies the filtering of relevant data based on its content
  - producers and consumers are fully agnostic of each other

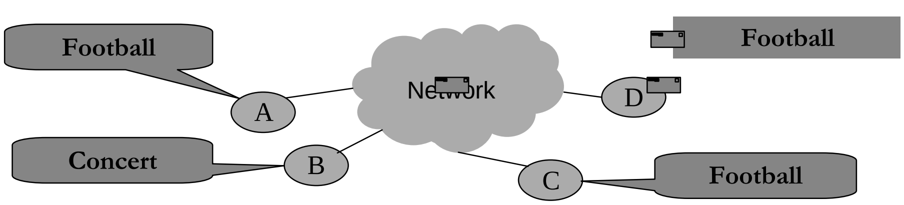

Research Challenges

- Publish/Subscribe looks very appropriate for mobile environments
- Nevertheless, mobility greatly complicates its deployment
  - only few works address publish/subscribe in mobile networks
  - none of them takes the social nature of human networks into account
- Observation #1: people do *not* move randomly
- Observation #2: people with similar interests tend to be co-located quite regularly
  - ✓ this characteristic can be exploited to drive forwarding decision

Our Proposal: SocialCast

- We assume that the mobility of users is driven by their social behavior that, in turn, is determined by their common interests
- We exploited predictions based on metrics of social interaction to identify the best information carriers
- *SocialCast* complements the information about the receivers' interests, necessary to routing information, with data about the social ties of people and their predicted movements

Overview

- Exploitation of *store-and-forward* mechanisms based on hosts acting as carriers for asynchronous delivery of messages to final recipients
- Choice of the *best carrier(s)* based on the evaluation of *context information*
  - Host colocation
  - Host mobility
  - Battery level
  - …
- We use Time Series Analysis based on State Space Models to keep history into account and to predict the evolution of DTN scenarios and Utility Theory for combining the different dimensions of the context

Protocol Description

- γ replicas are created and disseminate only to “good” neighbors
  - good: its utility is higher than mine

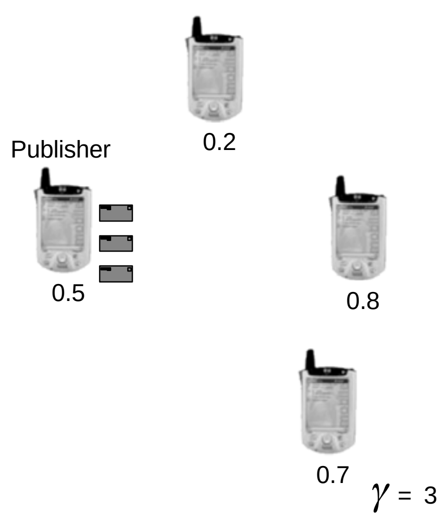

Protocol Description

- γ replicas are created and disseminate only to "good" neighbors
  - good: its utility is higher than mine
- Messages are kept until:
  - a better carrier is found
  - a subscriber is found
- Only carriers actively participate in message dissemination
  - replicas never increase
  - TTL is used to prevent infinite propagation
- A subscriber can act as carrier as well (if its utility is high enough)

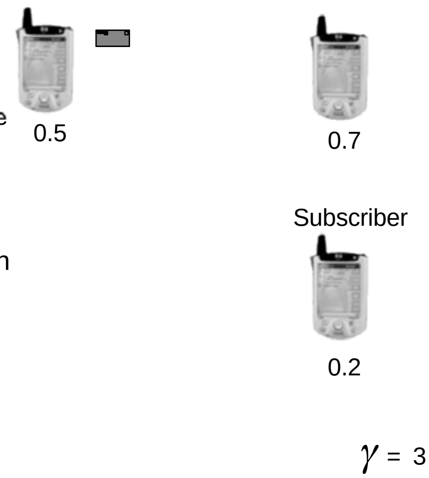

Calculation of the Utility

- Host utility calculated using multi-criteria decision theory
  - A utility is associated to each context attribute (i.e., a utility associated to host colocation)
  - Utilities are then composed using a weighted functions

$$U(a_1, a_2, ..., a_n) = \sum_{k=1}^{n} w_{a_k} \widehat{U}_{a_k}$$

- Host utilities are calculated locally and then periodically broadcasted to 1-hop neighbors

Calculation of the Utility

- We tested the algorithm considering two attribute: colocation and change degree of connectivity:

$$U_{h,i} = w_{cdc_h} \widehat{U}_{cdc_h} + w_{col_{h,i}} \widehat{U}_{col_{h,i}}$$

- Rationale:
  - if a node has been colocated with a subscriber, it will do again with high probability (*colocation*)
  - if a node exhibits high mobility, it is more likely to meet other good carries and/or subscribers (*change degree of connectivity*)
- Other dimensions can be included as well
  - e.g., battery level (only powerful nodes are used as carriers)

Predicting The Future

- Knowledge about the current values of these social attributes is helpful, but only to a limited extent
  - what matters is the value they will assume in the near future
- We compute these predicted values using forecasting techniques based on *Kalman filter*
  - only information about the current state must be maintained (no history is needed)

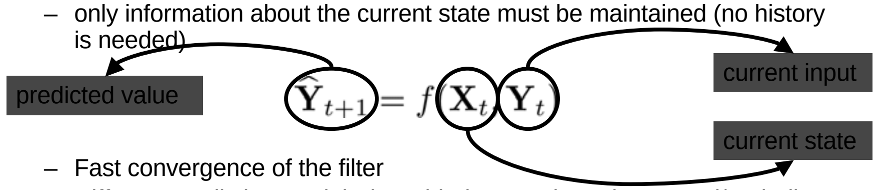

$$\widehat{\mathbf{Y}}_{t+1} = f(\mathbf{X}_t, \mathbf{Y}_t)$$

  - Fast convergence of the filter
  - Different prediction models (considering trends and seasonal/periodic behaviour)
  - suitable for resource-constrained devices

Simulation Parameters

- Scenario:
  - 100 nodes in a 4 km x 4 km area
  - Transmission range=250 m, omnidirectional antenna
  - Speed 1-6 m/s
  - Community mobility model (see next)
- Traffic:
  - 28800 sec (8 hours) of simulated time
  - messages are published between 3000 and 3500 sec
  - 10 different interests
- Protocol parameters:
  - Retransmission interval=60 sec
  - Routing tables dissemination interval=60 sec
  - utility weights $w_i = 0.5$
- Metrics
  - *delivery* (fraction of subscribers reached by the message)
  - *traffic* (overall number of forwarded messages)
  - comparison against a protocol with random carrier selection

Community-based Mobility Model

- Traditional mobility models assume random movements
  - inadequate to capture the social nature of human networks
- Mobile devices are usually carried by humans
  - therefore, mobile networks can be modeled as social networks after all
- Community-based Mobility Model [MM'07]
  - establishment of the group based on the community structure of the social network of the individuals carrying the devices
  - movement based on the relationships between individuals
  - model evaluated using real traces provided by Intel

[MM'07] ***Designing Mobility Models Founded on Social Networks Theory***
Mirco Musolesi and Cecilia Mascolo
ACM Mobile Computing and Communications Review. ACM Press. July 2007.

Delivery vs. Replicas

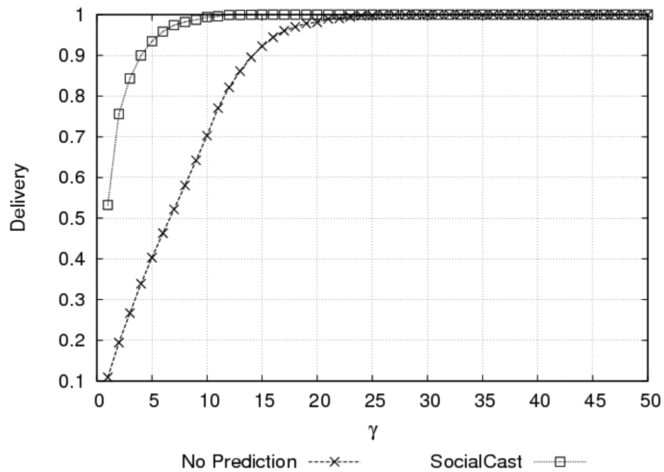

Overhead vs. Replicas

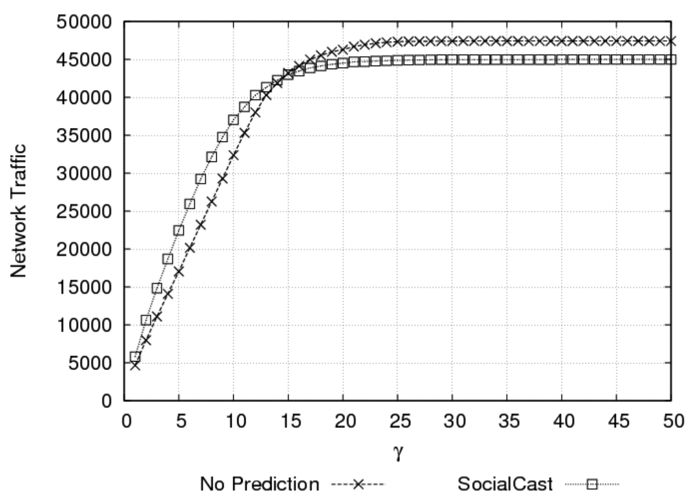

Number of Copies vs Latency

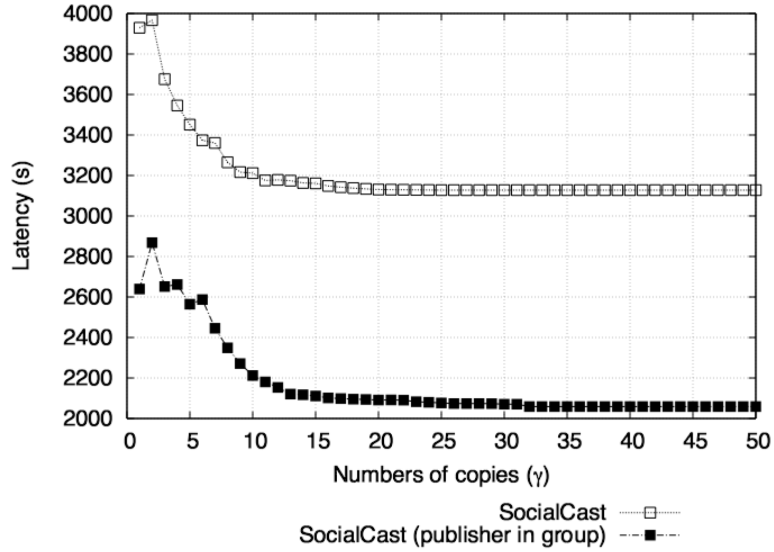

Delivery vs. TTL

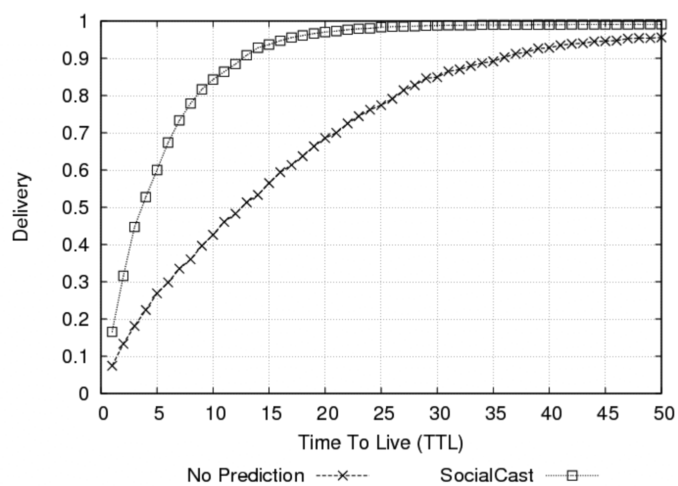

UCL

Overhead vs. TTL

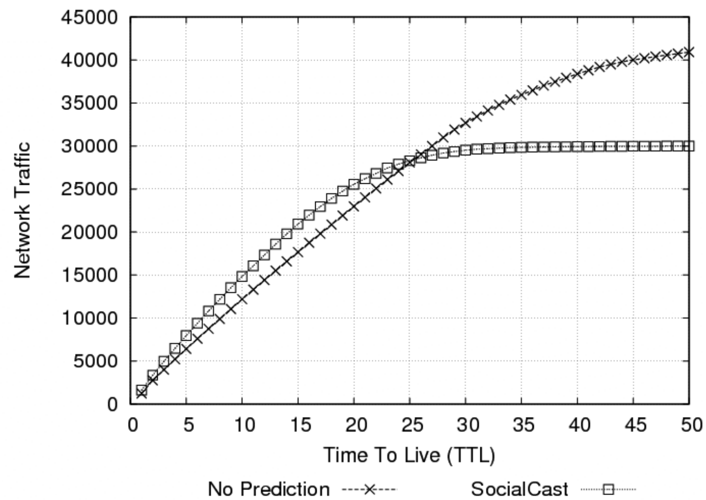

UCL

Hop Count Distribution

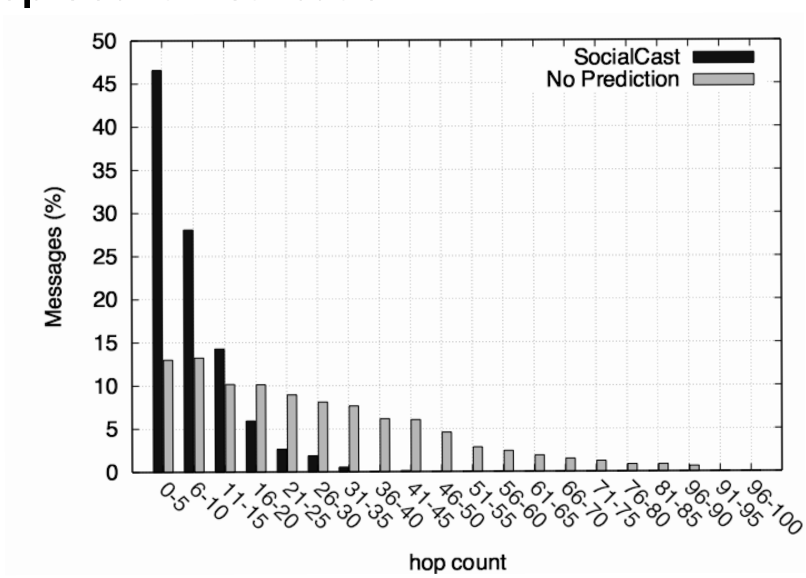

Buffer Size Distribution

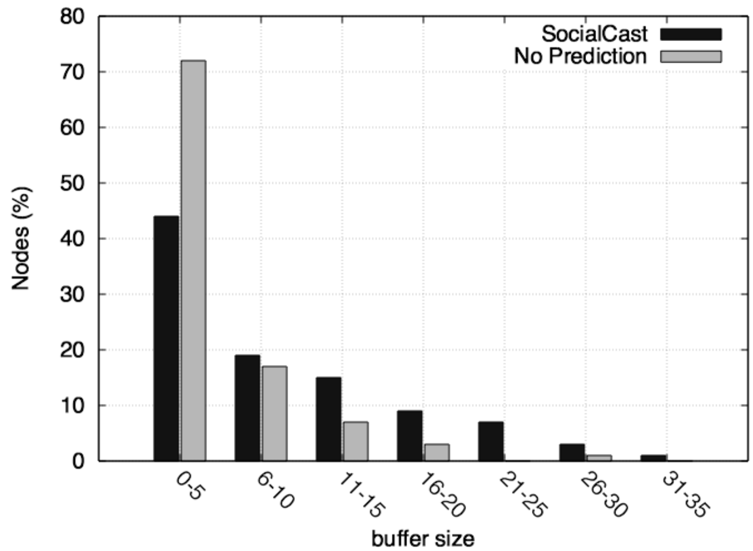

Conclusions

- Publish/Subscribe appears as a natural paradigm for communication in DTNs
- Socially-bound hosts are likely to be co-located regularly
- SocialCast exploits these colocation patterns to efficiently route messages from publisher to subscribers