Distributed Event Management

Sasu Tarkoma

16.12.2005

Contents

- Introduction
- Challenges
- Fuego Event System
- Overview of Demonstrations
  - PosetBrowser
  - Mobility simulator

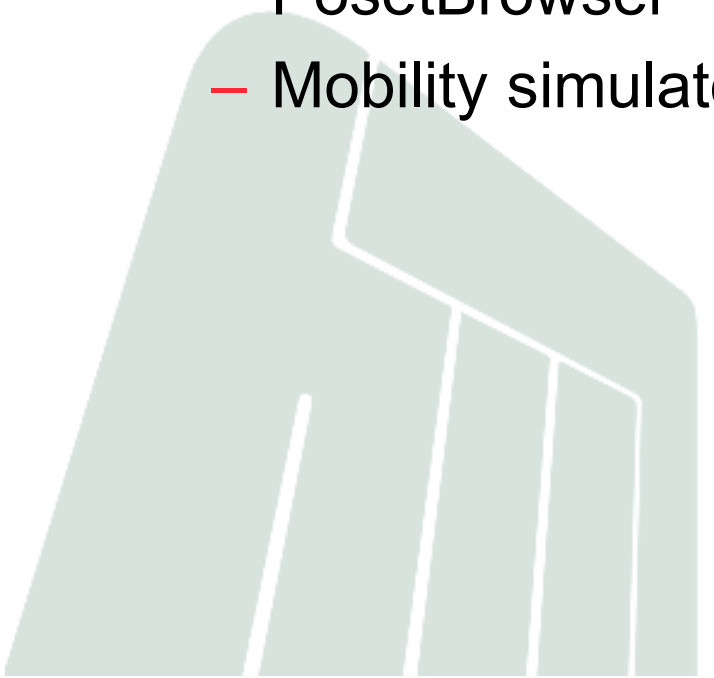

Introduction

- Event-based systems and publish/subscribe
  - Events (or notifications) are propagated from publishers to subscribers (sinks).
- A frequently used communication paradigm
  - Proposed for mobile computing
  - Current Web lacks publish/subscribe support
    - RSS feeds use HTTP and polling.
- Subscriptions are described using filters.
  - Filter: a stateless boolean function.
  - Selects a subset of events.
  - Expressive interest definition and content-matching
- Content model typically typed tuples or XML.
- The event service is a logically centralized service
  - Basic primitives: *sub*, *unsub*, *pub*
  - Various routing topologies and semantics
    - Reverse-path routing

Example of Publish/Subscribe

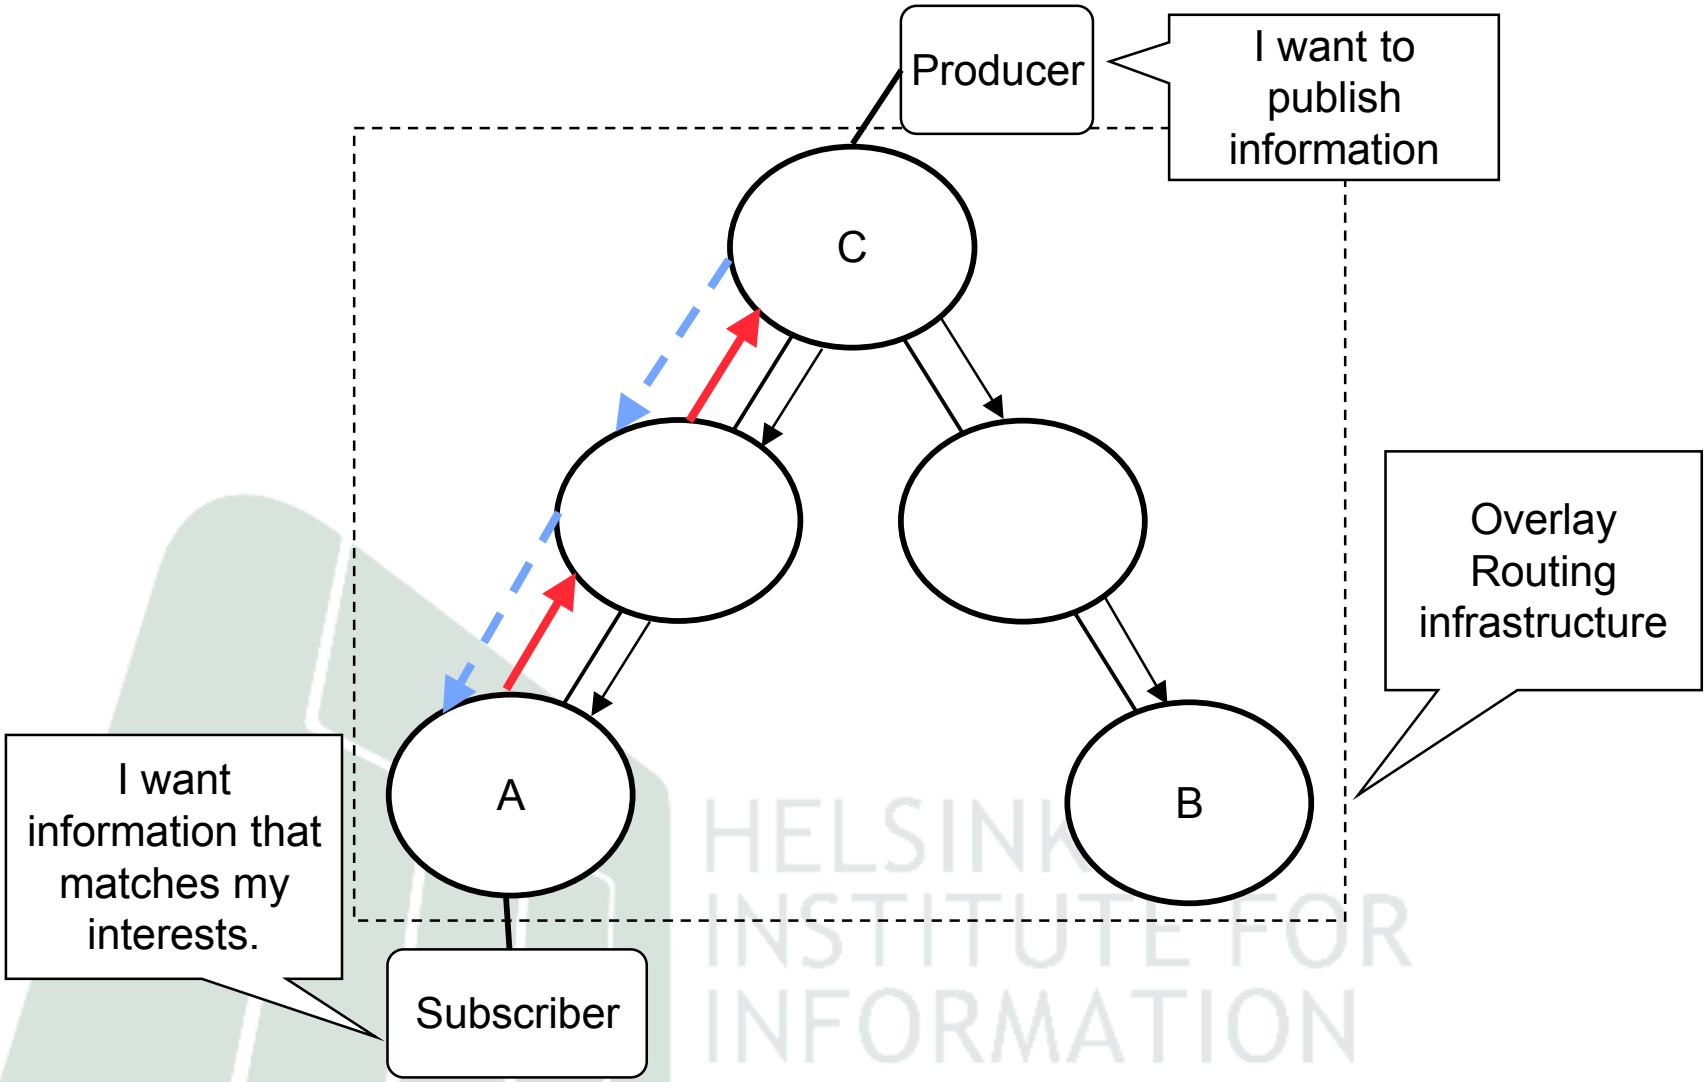

Challenges

- How to cope with mobile users?
  - Disconnected operation
    - Buffering and queue management
  - Mobile subscribers / producers
    - Handover protocol for relocating subscriptions and updating the topology
    - Multiple indirection points
- How to manage large numbers of filters?
  - Covering relations, filter merging
    - [0,10] covers [2,5]
    - [0,10] or [9,20] can be merged to [0,20]
  - These optimizations complicate mobility support
- General requirements
  - fast convergence of the subscription topology
  - mobility-safety: no false negatives

Example Handover

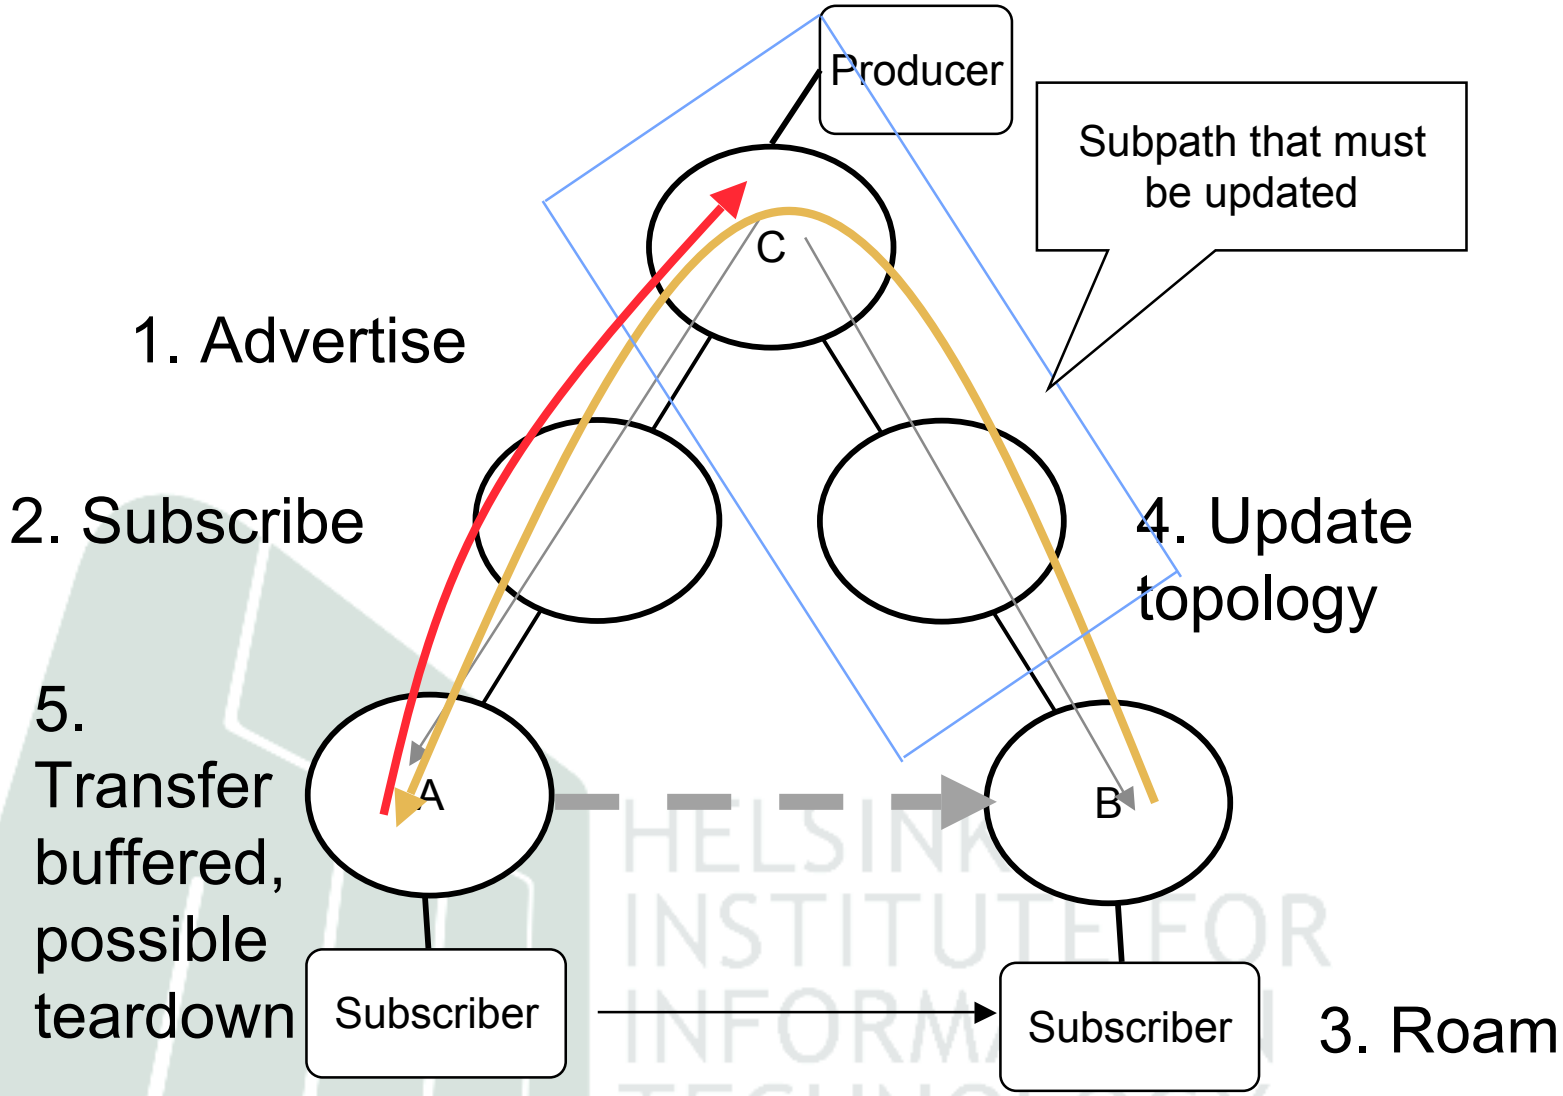

Fuego Event System

- Scalable distributed event framework for mobile computing
- The Fuego event router consists of two parts:
  - access server functionality with buffering and handover support for mobile clients, and
  - extensible routing core for distributed operation
- New data structures for efficient content-based routing:
  - poset (partially ordered set)-derived forest
  - the forest is considerably more efficient than dag (directed acyclic graph) - based structures
- Rendezvous-based mobility support for fast handovers and subscription topology updates
  - RP or paths to RP are updated instead of the whole topology

Summary of Data Structures / Models

- Data structures for cover-based routing
  - Poset-derived forest and variants
- A formal framework for filter merging
  - Both dynamic and static merging
- A discrete model for pub/sub mobility
  - Costs for handovers, simulation results
- Data structures for context-aware operation
  - DoubleForest for matching between context queries and context profiles (cover, overlap)
- The algorithms are useful for different filter and cover-based applications
- We use the DoubleForest in context-based collection and object synchronization
- Live demos on the web: www.hiit.fi/fuego/fc/demos

PosetBrowser

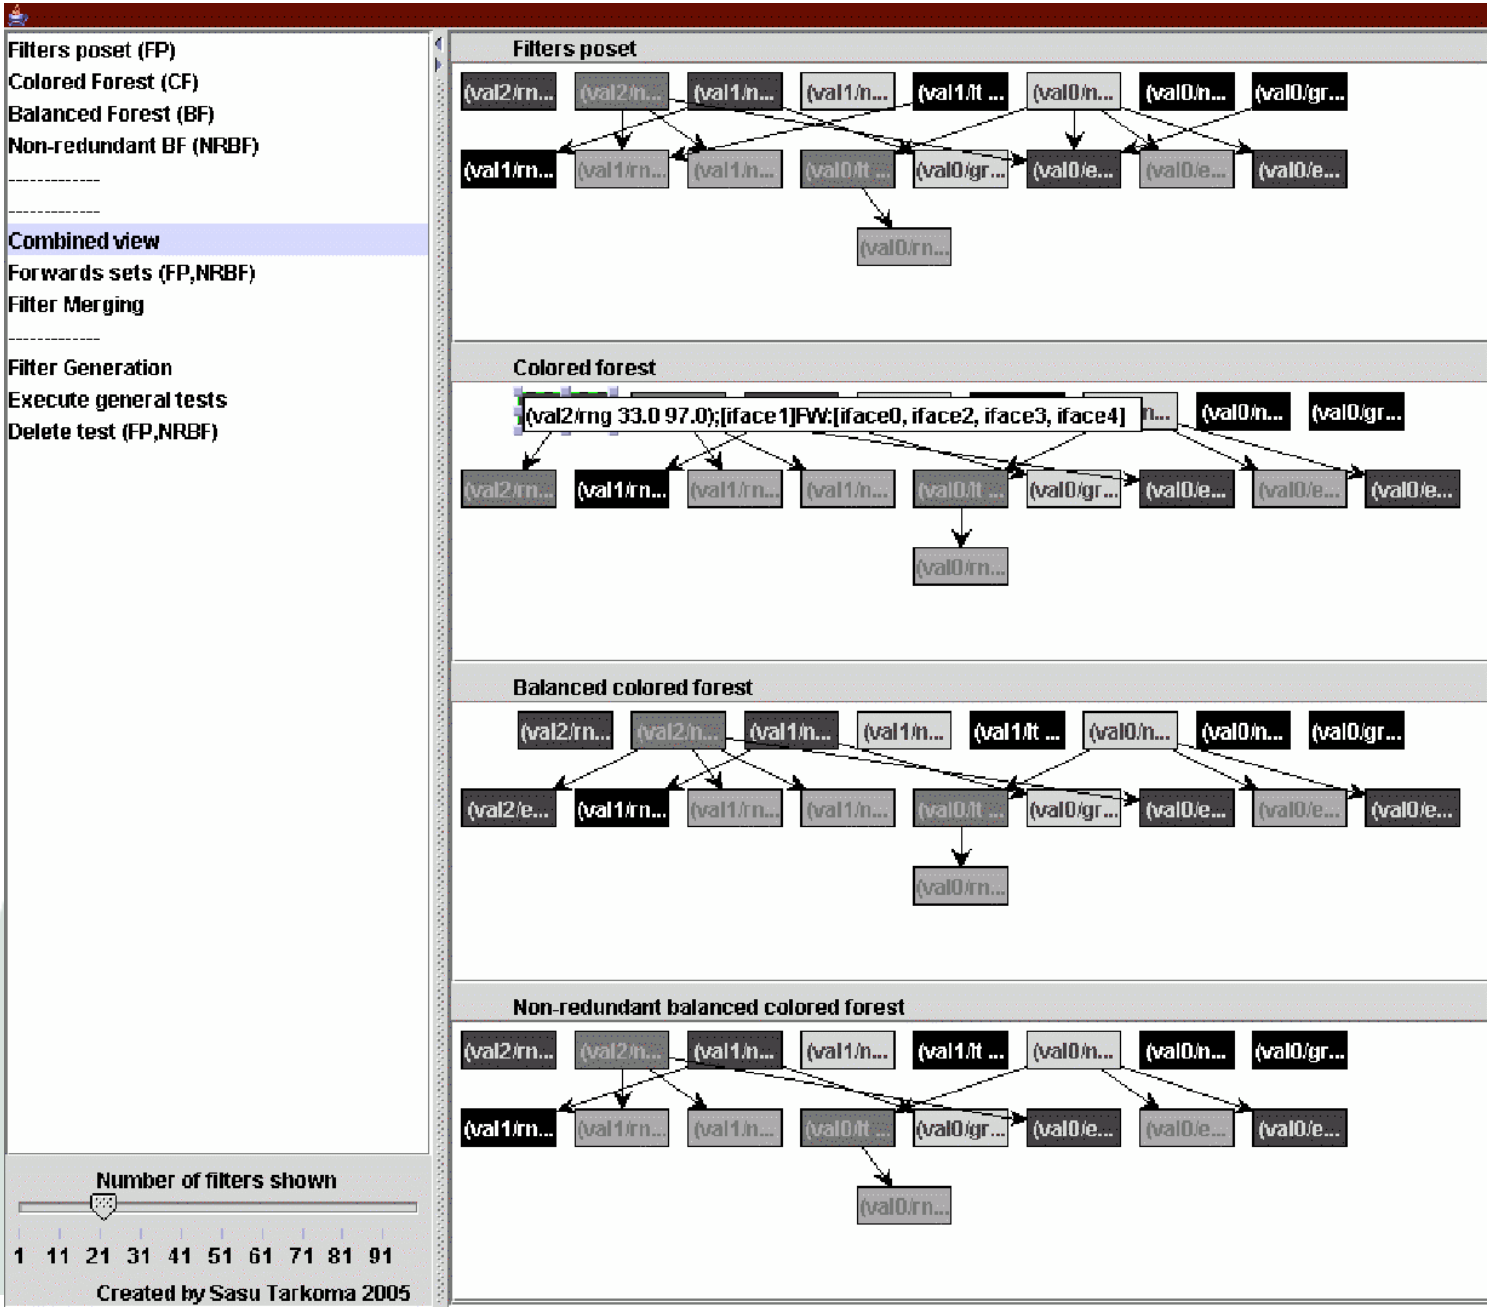

Mobility Simulator

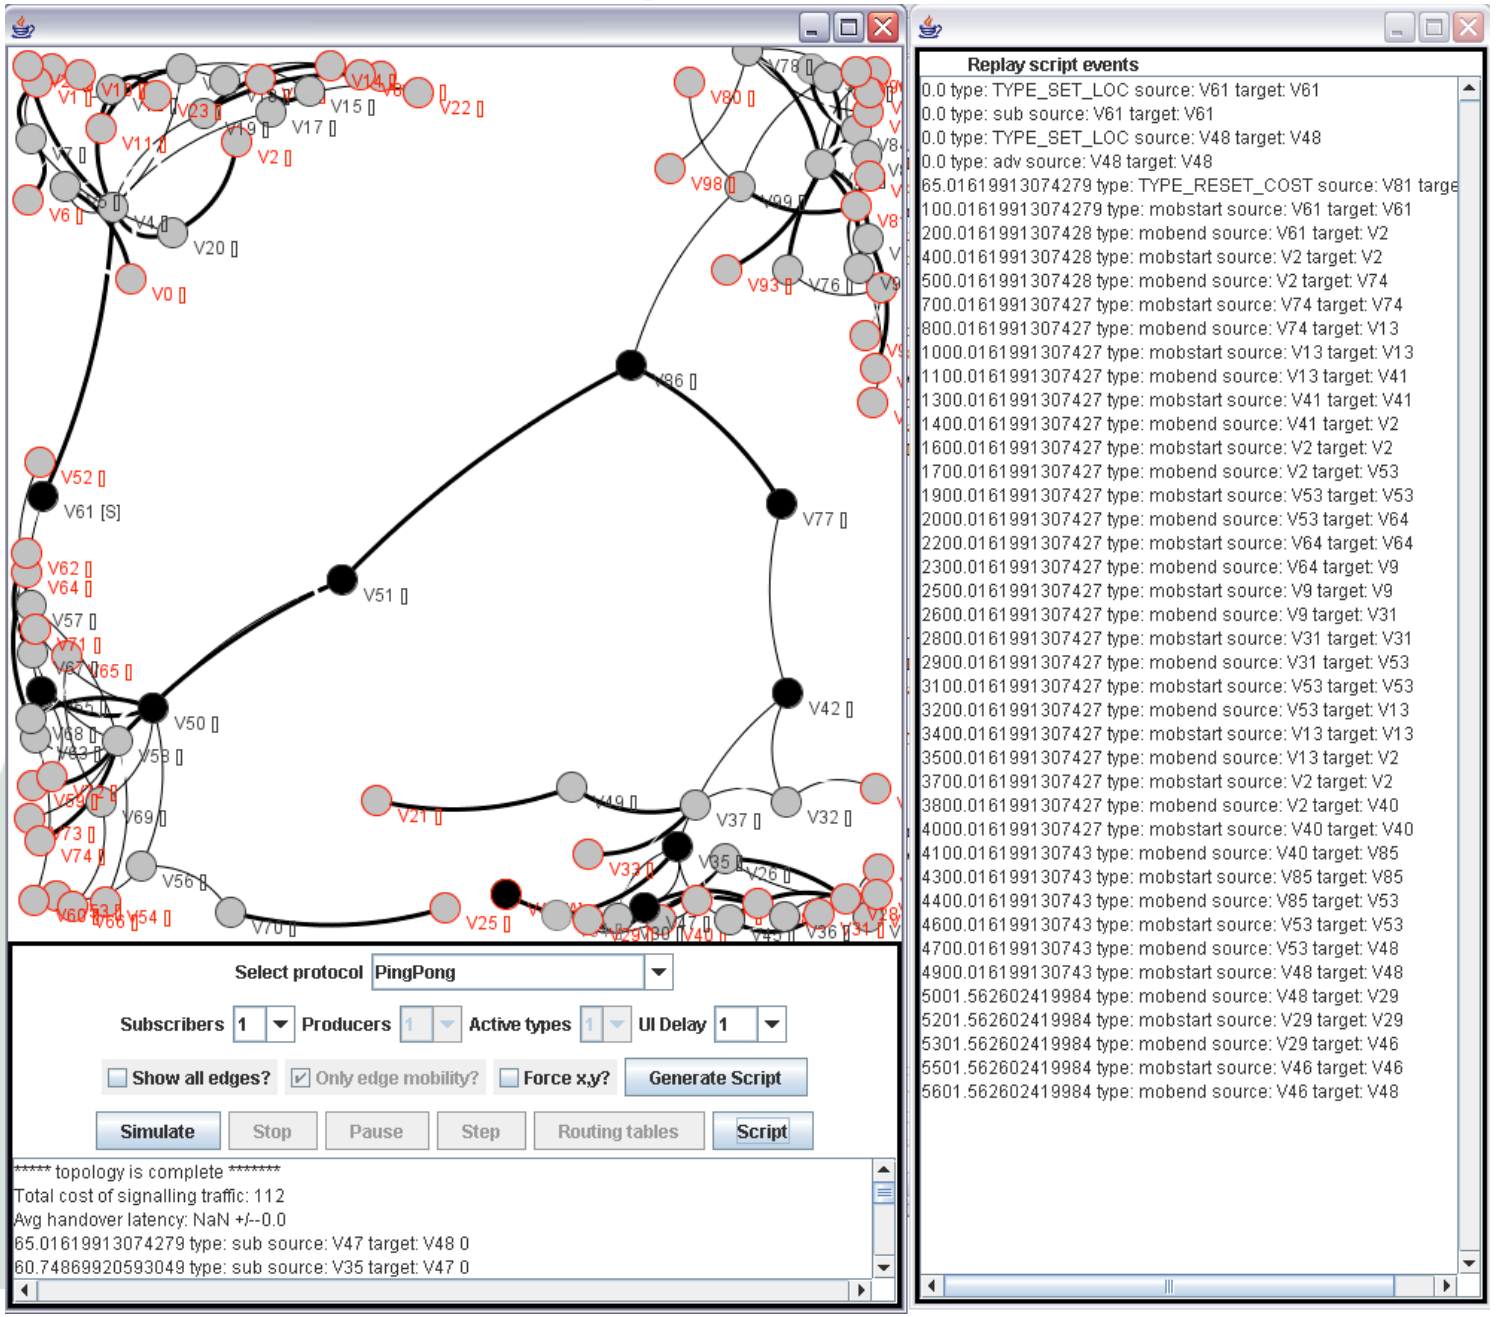

Future Work

- Ad hoc and peer-to-peer routing
- Massive scalability
- Applications
- Load balancing

Recent Papers on Events

- Sasu Tarkoma, Tancred Lindholm, and Jaakko Kangasharju. Collection and Object Synchronization Based on Context Information. In Second IEEE/IFIP International Workshop on Mobility Aware Technologies and Applications, October 2005.
- Sasu Tarkoma and Jaakko Kangasharju. Handover Cost and Mobility-Safety of Content Streams. In Eighth ACM/IEEE International Symposium on Modeling, Analysis and Simulation of Wireless and Mobile Systems, October 2005.
- Sasu Tarkoma and Jaakko Kangasharju. Filter Merging for Efficient Information Dissemination. In 13th International Conference on Cooperative Information Systems, October 2005.
- Sasu Tarkoma and Thalainayar Balasubramanian Ramya. A Gateway for SIP Event Interworking. In Third Workshop on Applications of Wireless Communications, August 2005.
- Sasu Tarkoma and Jaakko Kangasharju. A Data Structure for Content-based Routing. In IASTED EuroIMSA 2005, February 21-23, 2005, Grindelwald, Switzerland.
- Sasu Tarkoma and Jaakko Kangasharju. Mobility and Completeness in Publish/Subscribe Topologies. In IASTED International Conference on Networks and Communication Systems, ACTA Press, April 2005