Supporting Mobile Streaming Services in Future Publish/Subscribe Networks

Konstantinos Katsaros, Nikos Fotiou,
George C. Polyzos, George Xylomenos

Mobile Multimedia Laboratory
Department of Computer Science
Athens University of Economics and Business

polyzos@aub.gr

Outline

- Internet and mobility
- Mobile IP
- Multicast assisted mobility
- Internet Clean-Slate Design
- Publish-Subscribe Networking
- Overlay multicast architecture
  - Pastry
  - Scribe
- Overlay multicast assisted mobility (OMAM)
- OMAM vs. MIPv6: case studies
- Performance Evaluation
- Preliminary results
- Limitations and Future work

Internet and mobility

- Internet not designed with mobility in mind
  - No distinction between Location & End-point identifiers
- "Add-on" solutions
  - Mobile IP and optimizations
  - Micro-mobility protocols e.g. Cellular IP
  - Signaling delays, inefficient routing
- IP Multicast assisted mobility
  - —Localize route changes
  - —IP multicast failed to gain momentum!

Mobile IP

- Updates routing information so that the MN can be reached
  - at 1 node (HA, every move),
    or 2 nodes (& CN, every move if active connection)
  - Updates not local to the move
  - Binding Updates
- Bidirectional tunneling
  - All traffic passing through Home Network
  - Sub-optimal routing
- Route optimization
  - Binding update towards CN
  - Return Routability procedure
    - —Excessive signaling

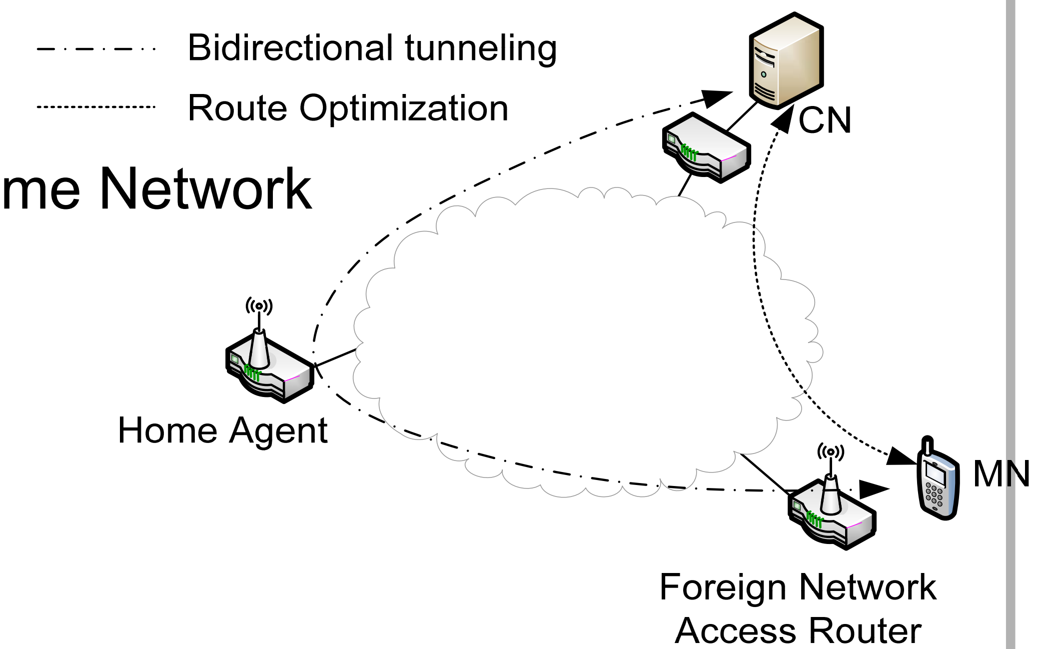

Multicast assisted mobility

- Localize routing updates
  - Not necessary to inform the source (CN)
- Multicast tree per user
  - Multiple users may share a single tree
- Proactive
  - Data can delivered to multiple locations
    - —All locations around the current location
    - —Predicted locations
  - Resource consumption
- Reactive
  - Data redirection upon handoff
- Application dependent

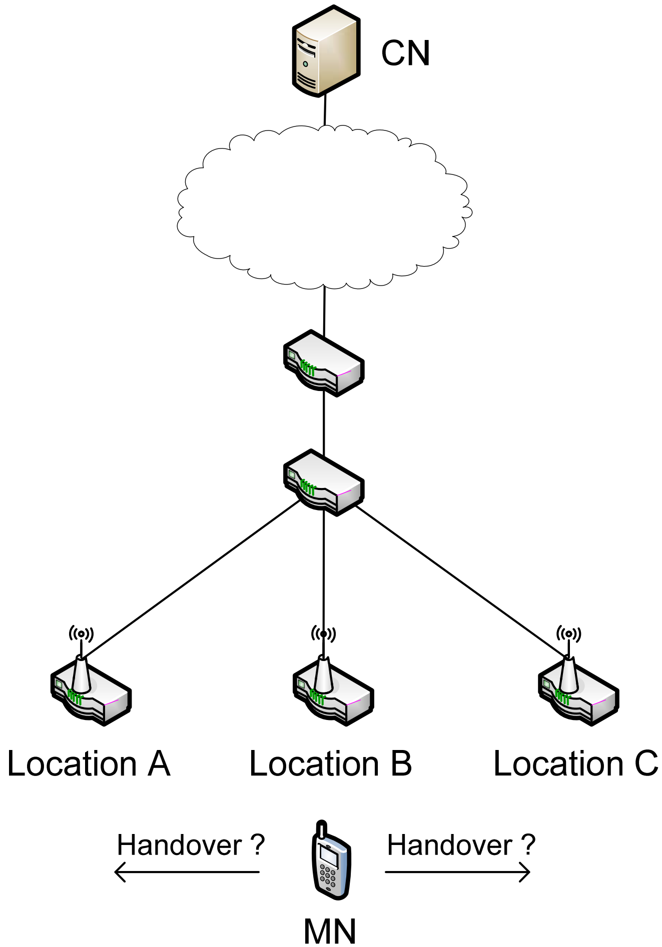

Are Internet Fundamentals Still Valid?

Fundamentals of the Internet

- Collaboration
  - Reflected in forwarding and routing
- Cooperation
  - Reflected in trust among participants
- Endpoint-centric services
  - (mail, FTP, even web)
  - Reflected in E2E principle
- **⇒ IP, full end-to-end reachability**

Reality in the Internet Today

- Phishing, spam, viruses
  - There is no trust any more!
- Current economics favor senders
  - Receivers are forced to carry the cost of unwanted traffic
- Information-centric services
  - Do endpoints really matter?
  - Endpoint-centric services move towards information retrieval through, e.g., CDNs

⇒ IP with middleboxes & significant decline in trust in the Internet

Vision

- **Envision a system that dynamically adapts to evolving concerns and needs of its participating users**

- Publish–subscribe based internetworking architecture restores the balance of network economics **incentives between the sender and the receiver**
- Recursive use of publish-subscribe paradigm enables dynamic change of roles between actors

Information-centric Network

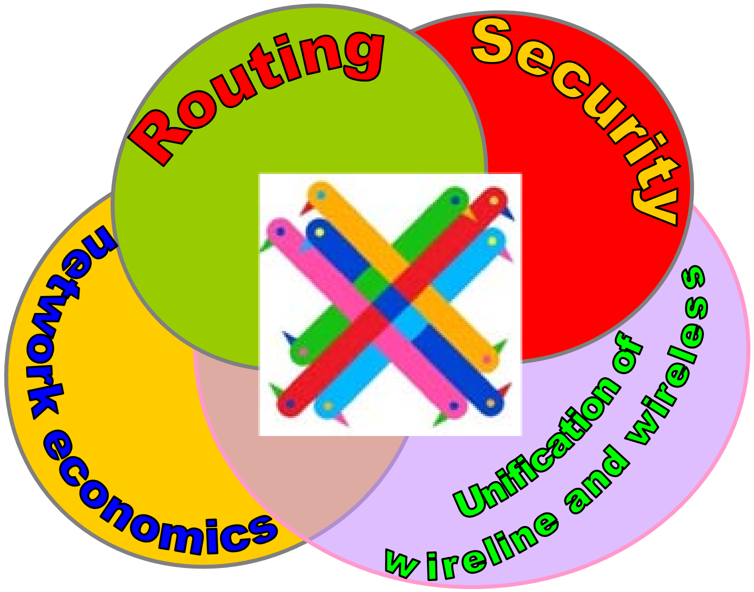

Main PSIRP design principles

- Information is multi-hierarchically organised
  - Higher-level information semantics are constructed in the form of directed acyclic graphs (DAGs), starting with meaningless forwarding labels towards higher level concepts (e.g., ontologies).
- Information scoping
  - Mechanisms are provided that allow for limiting the reachability of information to the parties having access to the particular mechanism that implements the scoping.
- Scoped information neutrality
  - Within each scope of information, data is only forwarded based on the given (scoped) identifier.
- The architecture is receiver-driven
  - No entity shall be delivered data unless it has agreed to receive those beforehand, through appropriate signalling methods.

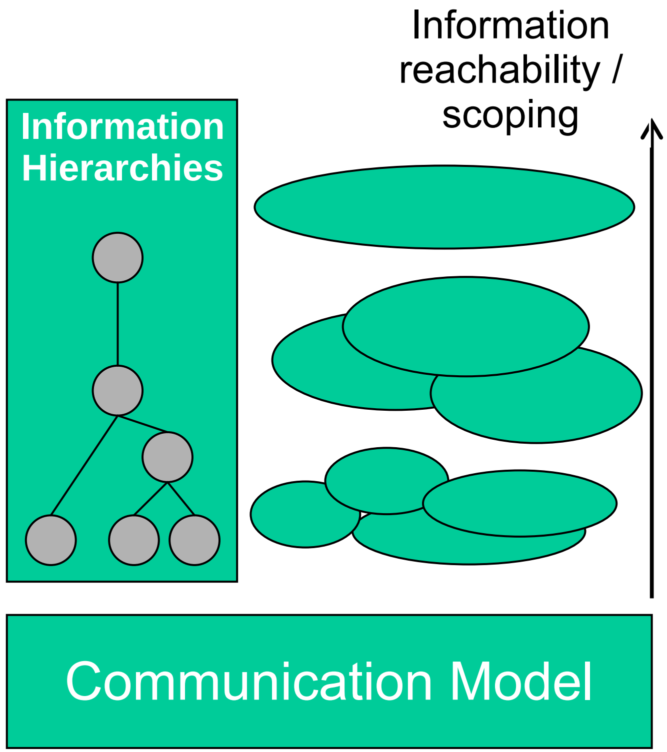

Overlay multicast architecture

- Considering an overlay publish/subscribe architecture
  - Access routers participate in a DHT (Pastry)
  - Also providing multicast routing (Scribe)
  - (Mobile) end-nodes directly connected to an overlay access router (OAR)
    - —Neither participate Pastry, nor carry an IP address
- Easier to deploy
  - Incremental/partial deployment
- At the cost of *signaling* and *stretch*
- Special care must be taken for inter-domain routing
  - Hierarchical DHTs (e.g. Canon)

Pastry overview

- DHT based, overlay routing infrastructure
- Nodes have unique 128-bit IDs
- Nodes are responsible for the keys (numerically) closer to their ID
- Proximity awareness (e.g. RTT)
- Routing of incoming message with key *K:*
  - Send to the (physically) closest of nodes that share one more digit with *K*
  - Increasing length of overlay (hops)
- Logarithmic number of overlay hops to any destination

Scribe overview

- Fully distributed pub/sub multicast infrastructure
  - Topic-based
    - —RV point = Pastry owner of Overlay Key(topic)
  - Subscribers send JOIN messages towards RV
    - —Each node forwards the message in the overlay UNLESS it has already subscribed to the group
  - Any node can send data to the RV point.
- Route convergence
    - —"Group members that are close in the network tend to be children of the same parent in the multicast tree that is also close to them."
    - —Simulations have demonstrated that the distance between the parent and every child is approximately equal to the distance between the children (in the proximity space).

Overlay multicast assisted mobility (OMAM)

- Whenever a MN wishes to receive a publication and/or upon handoff it sends a subscribe(RId) message to its OAR (Reactive)
- The OAR issues a Scribe JOIN message towards the RV
- OARs schedule a LEAVE Scribe message for a specific group when the last (mobile) member of that group has disassociated from the AP *+ delay*
  - —***delay**:* else the tree may have collapsed before handoff
  - —A mobile node may return to its original OAR
- In addition to multicast in general:
  - Route convergence: neighboring access points expected to have a close by common ancestor (CA)
- No end-to-end signaling: fast re-routing
- At the cost of path stretch!

polyzos@aueb.gr

OMAM vs. MIPv6: case study

- Packet flow considered already established
- MN initially attached to its Home Network
- What happens upon handoff?
  - MIPv6: Route Optimization (RO)
    - —Return Routability procedure
  - OMAM: newly visited OAR joins the tree
    - —Single JOIN message to OAR
      - – Propagates until lowest common ancestor of current and previous OAR
- Pastry signaling omitted
  - DHT assumed already available

polyzos@aueb.gr

OMAM vs. MIPv6: signalling

Required for Step 5 *or* No Route Optimization

- MIPv6: Return Routability procedure overhead

1.	Binding Update (BU),	MN→HA	$(d_{MN \to HA})$
2.	Binding Acknowledgement (BA),	HA→MN	$(d_{MN \to HA})$
3.	Home Test init (HoTi),	MN → HA→CN	$(d_{MN \to HA} + d_{HA \to CN})$
4.	Care-of-Test init (CoTi),	MN → CN	$(d_{MN \to CN})$
5.	Home Test (HT),	CN → HA → MN	$(d_{MN \to HA} + d_{HA \to CN})$
6.	Care-of Test (CT),	CN → MN	$(d_{MN \to CN})$
7.	Binding Update (BU),	MN → CN	$(d_{MN \to CN})$
8.	Binding Acknowledgement (BA),	CN → MN	$(d_{MN \to CN})$

In parallel (1–2 and 3–4)

In parallel (5–6)

- OMAM: newly visited OAR_k joins the tree

1.	Scribe JOIN msg,	MN → OAR_k → CA
2.	Scribe LEAVE msg,	OAR_{k-1} → CA

OMAM vs. MIPv6 with RO

- Resume Time (*RT* ~handoff)
  - MIPv6:

$$RT_{MIPv6} = 4\mathrm{d}_{\mathrm{MN}\to\mathrm{HA}} + 2\mathrm{d}_{\mathrm{MN}\to\mathrm{CN}} + 2\mathrm{d}_{\mathrm{HA}\to\mathrm{CN}}$$

  - OMAM:

$$RT_{OMAM} = \mathrm{d}_{\mathrm{MN}\to\mathrm{OAR_k}} + \mathrm{d}_{\mathrm{OAR_k}\to\mathrm{CA}}$$

  - Route Convergence:

$$\mathrm{d}_{\mathrm{OAR_k}\to\mathrm{CA}} = \alpha \times \mathrm{d}_{\mathrm{OAR_k}\to\mathrm{OAR_{k\text{-}1}}}, \alpha \to 1$$

- OMAM faster when: $RT_{MIPv6} > RT_{OMAM}$

$$\Rightarrow \quad \boxed{\alpha < 4 + 2\frac{2\mathrm{d}_{\mathrm{OAR_{k\text{-}1}}\to\mathrm{HA}} + \mathrm{d}_{\mathrm{OAR_k}\to\mathrm{CN}} + \mathrm{d}_{\mathrm{HA}\to\mathrm{CN}}}{\mathrm{d}_{\mathrm{OAR_{k\text{-}1}}\to\mathrm{OAR_k}}}}$$

But, according to route convergence property:

$$\alpha \to 1$$

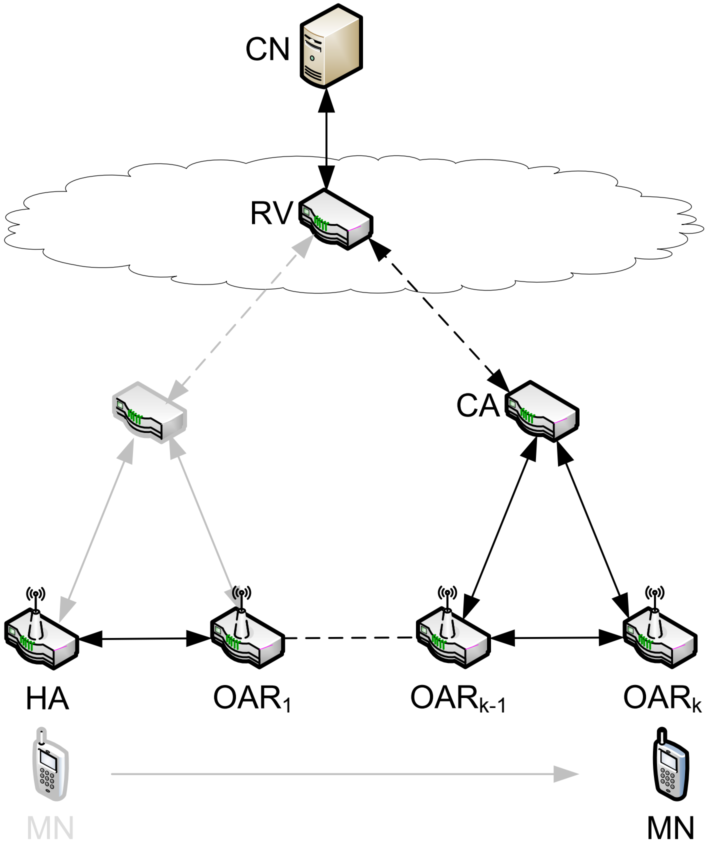

OMAM vs. MIPv6 simple BU

– MN simply sends a BU message to HA

- MIPv6:

$$RT'_{MIPv6} = \mathrm{d}_{\mathrm{MN}\rightarrow\mathrm{OAR_k}} + \mathrm{d}_{\mathrm{OAR_k}\rightarrow\mathrm{HA}}$$

- OMAM:

$$RT_{OMAM} = \mathrm{d}_{\mathrm{MN}\rightarrow\mathrm{OAR_k}} + \mathrm{d}_{\mathrm{OAR_k}\rightarrow\mathrm{CA}}$$

– OMAM faster when: $RT_{MIPv6} > RT_{OMAM}$

$$\Rightarrow \boxed{\alpha < \frac{\mathrm{d}_{\mathrm{OAR_k}\rightarrow\mathrm{HA}}}{\mathrm{d}_{\mathrm{OAR_{k-1}}\rightarrow\mathrm{OAR_k}}}}$$

But, usually: $\mathrm{d}_{\mathrm{OAR_k}\rightarrow\mathrm{HA}} > \mathrm{d}_{\mathrm{OAR_{k-1}}\rightarrow\mathrm{OAR_k}}$

and $\alpha \rightarrow 1$

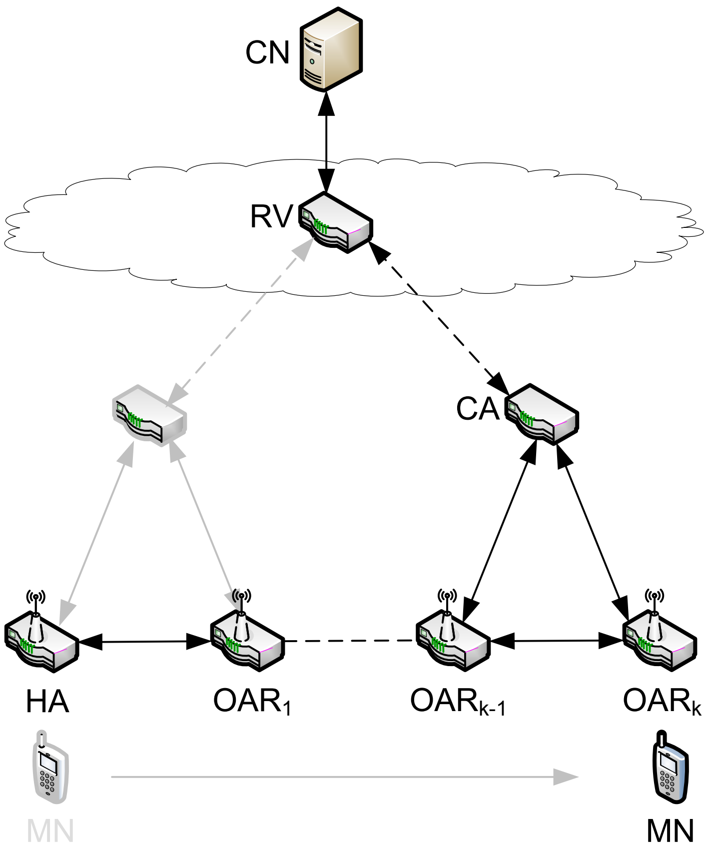

Performance Evaluation

- Performance metrics
  - Packet loss
    - —Lost connectivity + signaling delay
    - —Depicts handover speed (depends on acceptable delay)
  - End-to-end packet delay
    - —Time required for a packet to reach its destination
    - —Depicts overlay stretch!
  - Resume time
    - —Time required for the first packet to be received after a handoff
    - —Depicts handover speed
- One-way communication, e.g. video streaming
- Simulation environment
  - OMNeT++, xMIPv6, OverSim

Topology

- Grid topology
  - IEEE 802.11b APs
- Full coverage
  - Focus on signaling-based disruption
- UDP stream: H.264, Level 1 SQCIF video stream, 30.9 fps

Parameter	Value
Grid size	30 x 30
Number of MNs	1
Number of CNs	1
Wired connections type	100Mps Ethernet
Propagation delay (ms)	0.5
Data rate (Kbps)	64
Packet size (bytes)	26
Total number of packets sent	556200

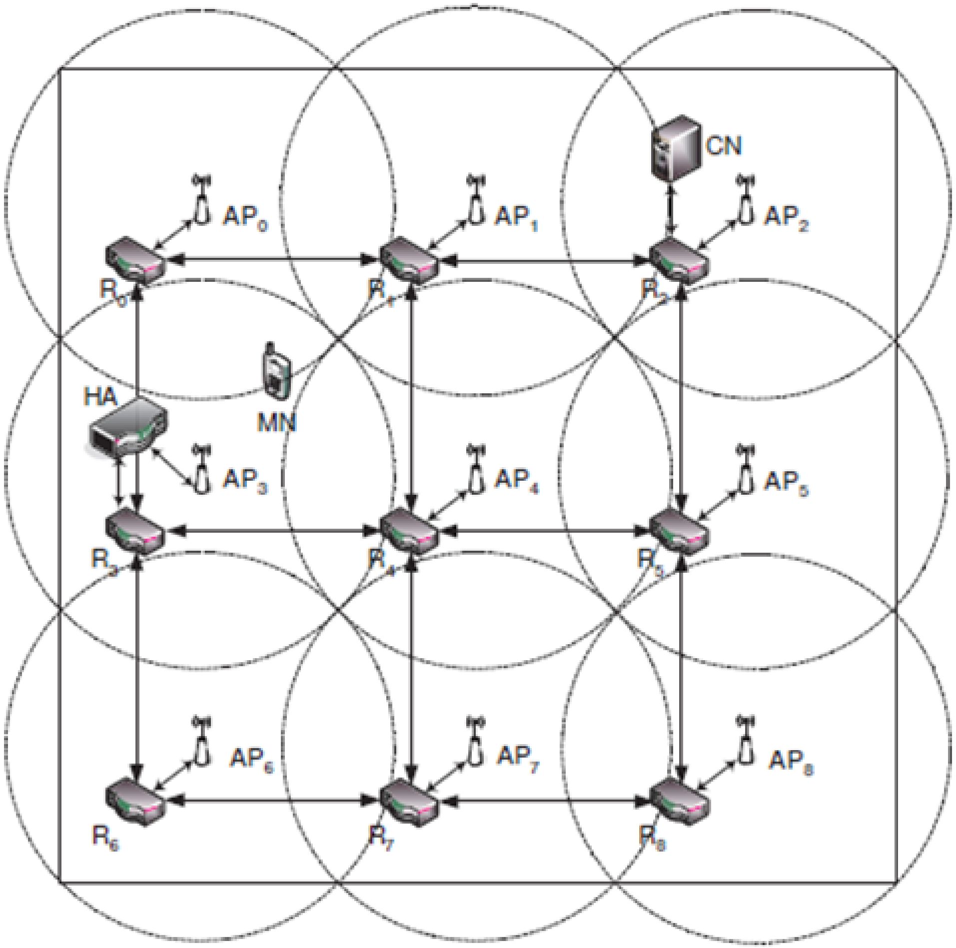

Preliminary results

- Significant gains in signaling overhead
  - Service disruption greatly improved
- At the cost of increased end-to-end delay
  - Impact of overlay routing i.e. *stretch*
  - Acceptable for non-interactive streaming application

	MIPv6	**Mobile Scribe**
Packet loss	2.002%	1.059%
End-to-end delay	12ms	17ms
Resume time	1.208 sec	0.007 sec

Conclusions & Future work

- Multicast presents significant advantages in supporting mobility
  - Enabled/revisited in an overlay context
- DHT substrate properties further enhance multicast tree properties
- Promising preliminary results,
  - Especially for streaming applications

- Comparison with micro-mobility protocols
  - Hierarchical Mobile IPv6
- Measure/quantify *route convergence* property!
- Simplistic topology & mobility model
  - Incorporate campus-wide wireless traces, e.g. Dartmouth campus traces
- Two-way, reliable communication
  - *Lag-behind/get-ahead,* reverse path vs. distinct trees, *etc.*

Thanks!

George C. Polyzos

Mobile Multimedia Laboratory

Department of Informatics/Computer Science
Athens University of Economics and Business
47A Evelpidon, 11362 Athens, Greece

polyzos@aueb.gr, http://mm.aueb.gr/
Tel.: +30 210 8203 650, Fax: +30 210 8203 325