Overlay Multicast Assisted Mobility for 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
Greece

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 routing inefficiencies

- 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!

At the same time:

- Lack of multicast support & shift to information-centric services resulted in excessive traffic
  - P2P, file-sharing applications dominate traffic (e.g. BitTorrent)
  - End-to-end Internet semantics neglect network resource consumption
    - Redundant transmissions
  - IP multicast would prove beneficial but again …not available!

Are Internet Fundamentals Still Valid?

Fundamentals of the Internet

- Cooperation
  - Reflected in trust among participants
- Collaboration
  - Reflected in forwarding and routing
- Endpoint-centric services
  - (mail, FTP, even web)
  - Reflected in E2E principle
- Stationary endpoints

- ⇒ **IP, full end-to-end reachability**

vs.

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
  - Endpoint-centric services move towards information retrieval through, e.g., CDNs
  - Cloud computing
- Mobility

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

It's the new ways Internet is used; that was not designed for…

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

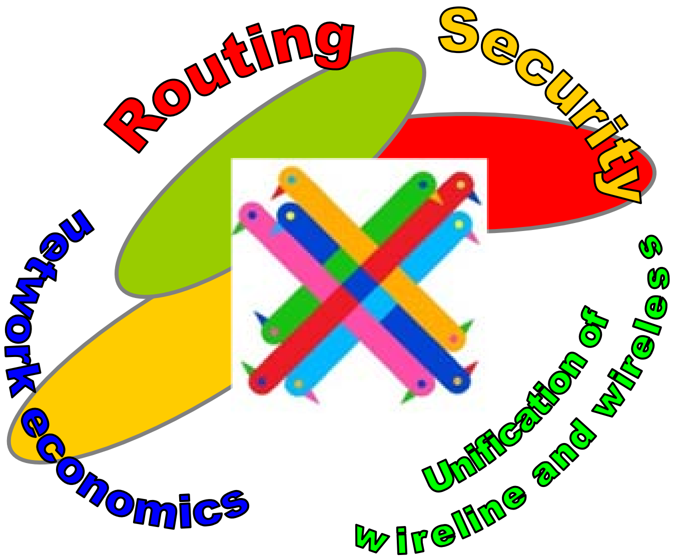

Information-centric Network

PSIRP PUBLISH-SUBSCRIBE INTERNET ROUTING PARADIGM
ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS

The Publish/Subscribe approach

- **Endpoints:**
  - Publishers: data owners
    - Provide pieces of information in the form of *publications*
  - Subscribers (data consumers)
    - Express interest in pieces of information via *subscriptions*
- **Network:**
  - Event notification service (broker substrate): matching *publications* and *subscriptions*

⇒

- *End-to-end decoupling*
  - Publishers/Subscribers need not be aware of corresponding Subscribers/Publishers
  - Asynchronous communication
- *Multicast*
  - Multiple subscriptions can be grouped, brokers merge data streams
  - Norm in pub/sub
- *Caching*
  - Pub/sub state and multicast suitable for in-network caching

⇓

Ideal for mobility!

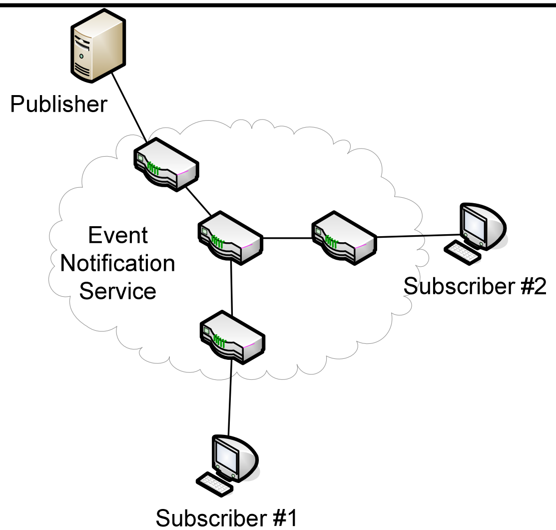

Overlay multicast architecture

- Considering an overlay publish/subscribe architecture
  - Access routers participate in a DHT (Pastry)
  - Also providing overlay multicast routing (Scribe)
  - (Mobile) end-nodes directly connected to an overlay access router (OAR)
    - Neither participate Pastry, nor carry an IP address
- Overlay approach: easier to deploy
  - Incremental/partial deployment
  - Not only for mobility support …!
- At the cost of extra *signaling* and *stretch*
- Special care must be taken for inter-domain routing
  - Hierarchical DHTs (e.g. Canon)

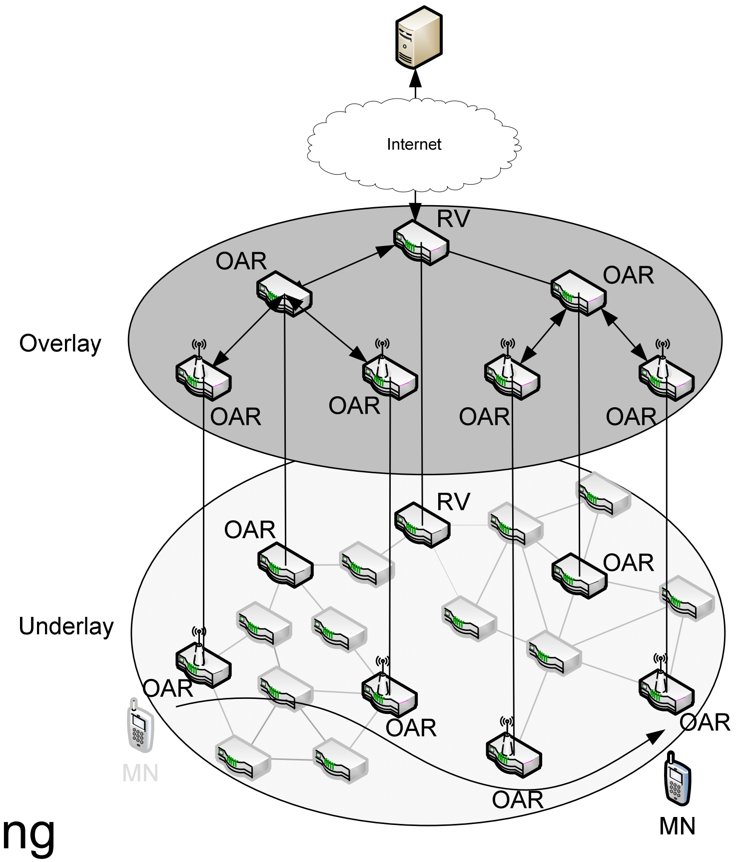

Multicast assisted mobility revisited

- User movement tends to be localized
  i.e. trajectory visiting neighboring network entities
- Target: localize routing updates too!
  - Not necessary to inform the source (CN)
- Multicast tree per user
  - Multiple users may share a single tree
    - Especially in a PSIRP architecture!
- Proactive
  - Data can delivered to multiple locations
    - All locations around the current location
    - Predicted locations
  - Resource consumption
- Reactive
  - Data redirection upon handoff
- Application dependent

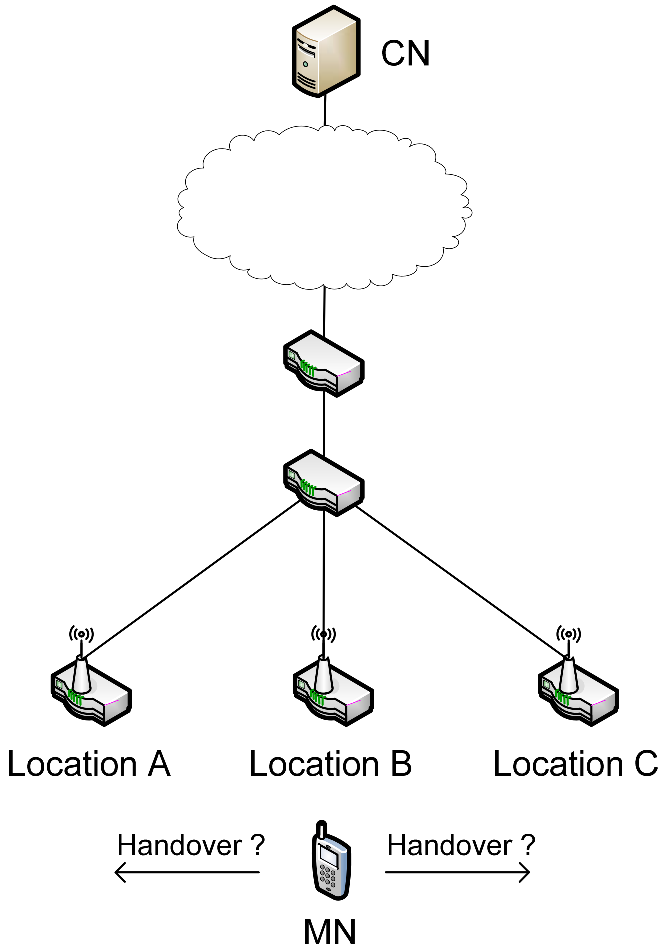

Overlay Multicast Assisted Mobility (OMAM)

- Overlay realization:
  - MN sends a subscription message to its OAR (Reactive) to receive a publication
    - Also upon handoff
    - Translated by the OAR into 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
- **Route convergence:** neighboring access points expected to have a close by common ancestor (CA)
  - In favor of localized routing updates!
- No end-to-end signaling: fast re-routing
- At the cost of path stretch!

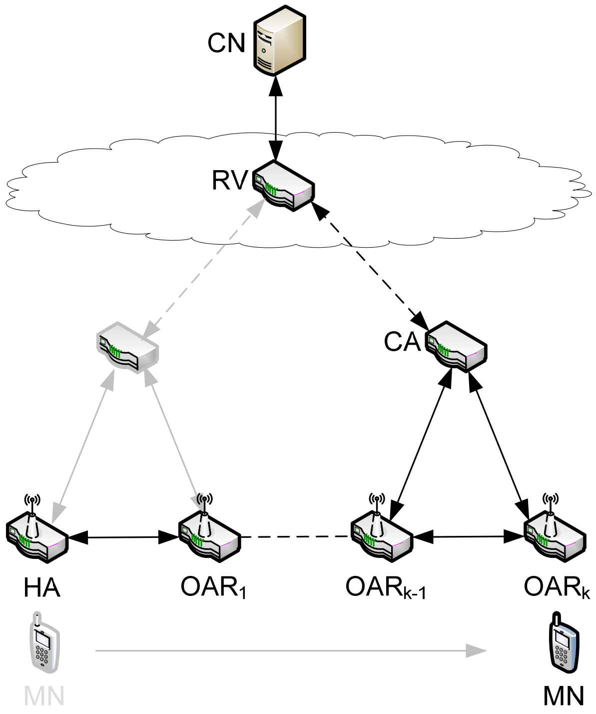

The available solution: Mobile IP

- 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

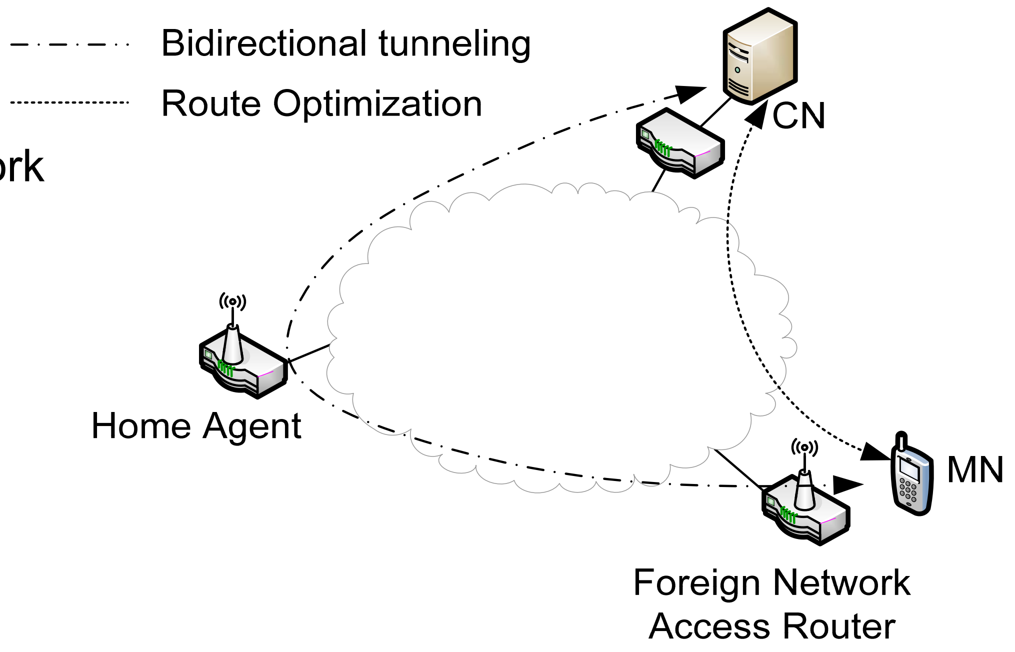

OMAM vs. MIPv6: case study

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

OMAM vs. MIPv6 simple BU

- MN simply sends a BU message to HA or CN
  - MIPv6:

$$RT_{MIPv6} = d_{MN \rightarrow OAR_k} + d_{OAR_k \rightarrow HA/CN}$$

  - OMAM:

$$RT_{OMAM} = d_{MN \rightarrow OAR_k} + d_{OAR_k \rightarrow CA}$$

  - Route Convergence:

$$d_{OAR_k \rightarrow CA} = \alpha \times d_{OAR_k \rightarrow OAR_{k-1}}, \alpha \rightarrow 1$$

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

$$\Rightarrow \boxed{\alpha < \frac{d_{OAR_k \rightarrow HA/CN}}{d_{OAR_{k-1} \rightarrow OAR_k}}}$$

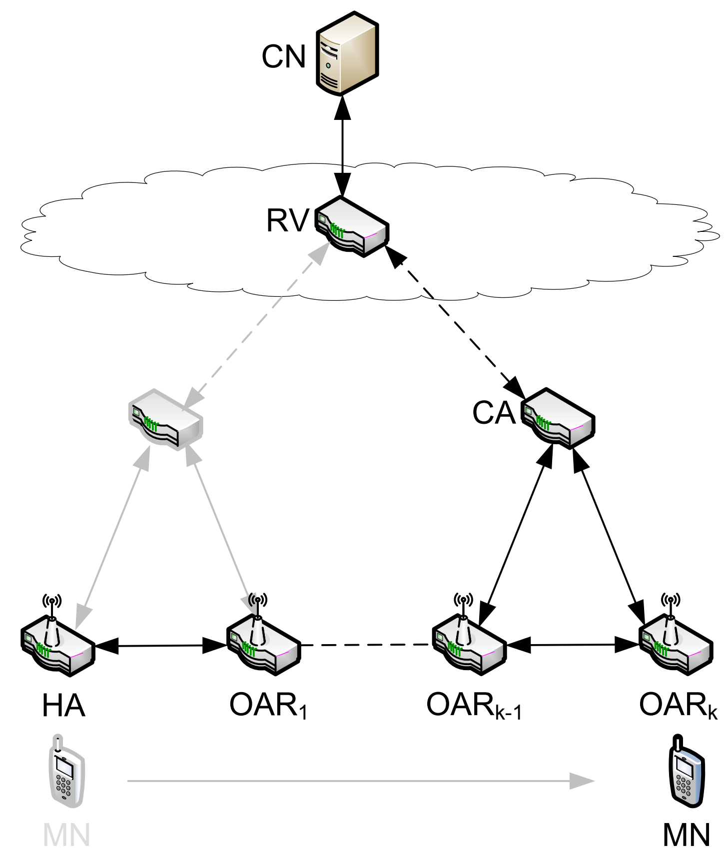

But, usually: $d_{OAR_k \rightarrow HA/CN} > d_{OAR_{k-1} \rightarrow OAR_k} \Leftrightarrow \frac{d_{OAR_k \rightarrow HA/CN}}{d_{OAR_{k-1} \rightarrow OAR_k}} > 1$

and $\alpha \rightarrow 1$

Localized character of movement + Multicast + Route Convergence ⇒ Reduced handoff delay!

($d_{x \rightarrow y}$: delay of message sent from network entity x to y)

PSIRP PUBLISH-SUBSCRIBE INTERNET ROUTING PARADIGM
ATHENS UNIVERSITY OF ECONOMICS AND BUSINESS

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

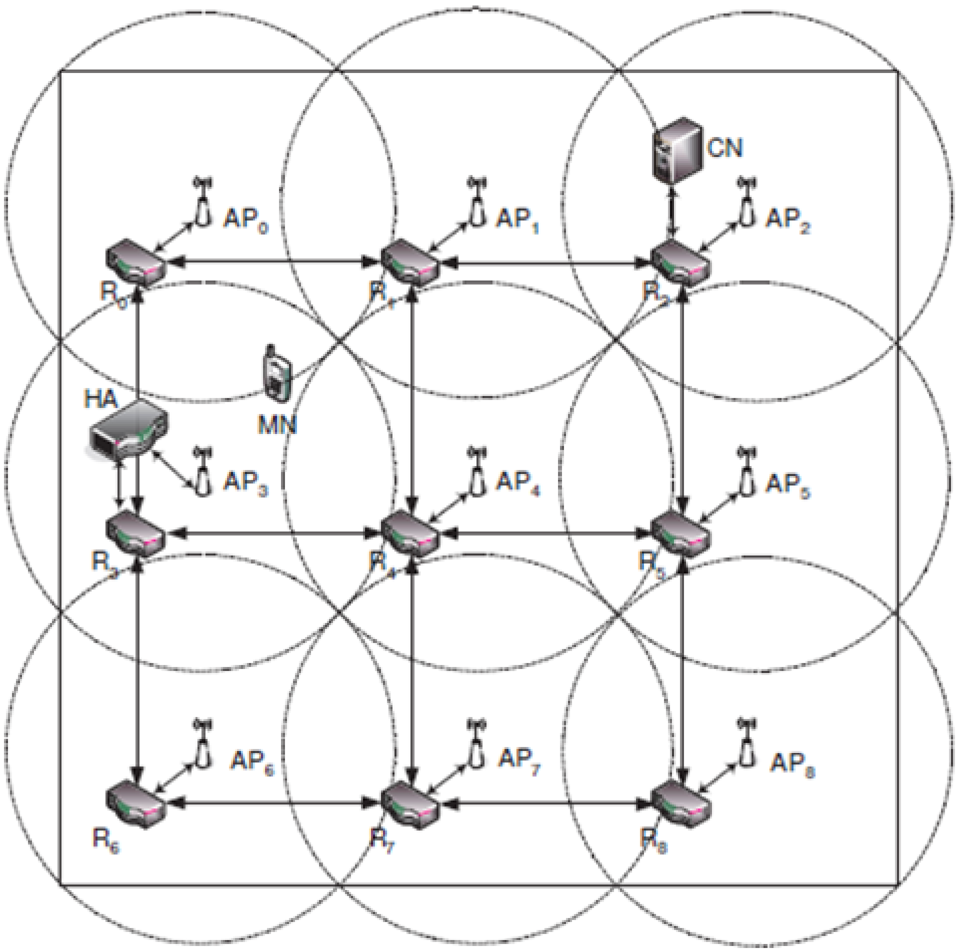

Parameter	Value
Grid size	30 x 30
Number of MNs	1
Number of CNs	1
Wired connections type	100Mps Ethernet
Propagation delay (ms)	0.5
Data rate (Kbps)	64
Packet size (bytes)	26
Total number of packets sent	556200

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!

Konstantinos Katsaros

Mobile Multimedia Laboratory

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

ntinos@aueb.gr, http://mm.aueb.gr/
Tel.: +30 210 8203 693, Fax: +30 210 8203 686

Back up slides ...

OMAM vs. MIPv6: signalling

Required for Step 5 *or* No Route Optimization

- MIPv6: Return Routability procedure overhead

In parallel (steps 1–2, 3–4)

In parallel (steps 5–6)

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	$(d_{MN \to CN})$

- OMAM: newly visited OAR_k joins the tree

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

($d_{x \to y}$: delay of message sent from network entity x to y)

OMAM vs. MIPv6 with RO

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

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

  - OMAM:

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

  - Route Convergence:

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

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

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

- But, according to route convergence property:

$$\boxed{\alpha \to 1}$$