IEEE CQR 2012
May 17, 2012 - San Diego, CA

An Efficient and Scalable Engine for Large Scale Multimedia Overlay Networks

Pier Luca Montessoro

Department of Electric, Managerial and Mechanical Engineering (DIEGM)

University of Udine
Italy

Stefan Wieser
Laszlo Böszörmenyi

Institute of Information Technology (ITEC)

Klagenfurt University
Austria

Motivations

- Increasing demand of multimedia streaming and remote storage
- No control on the network infrastructure and limited cooperation from ISPs
- To make overlays a feasible solution we must provide:
  - Scalability
  - Flexibility through dynamic self-organization
  - Performance (fast packet forwarding)

Architecture

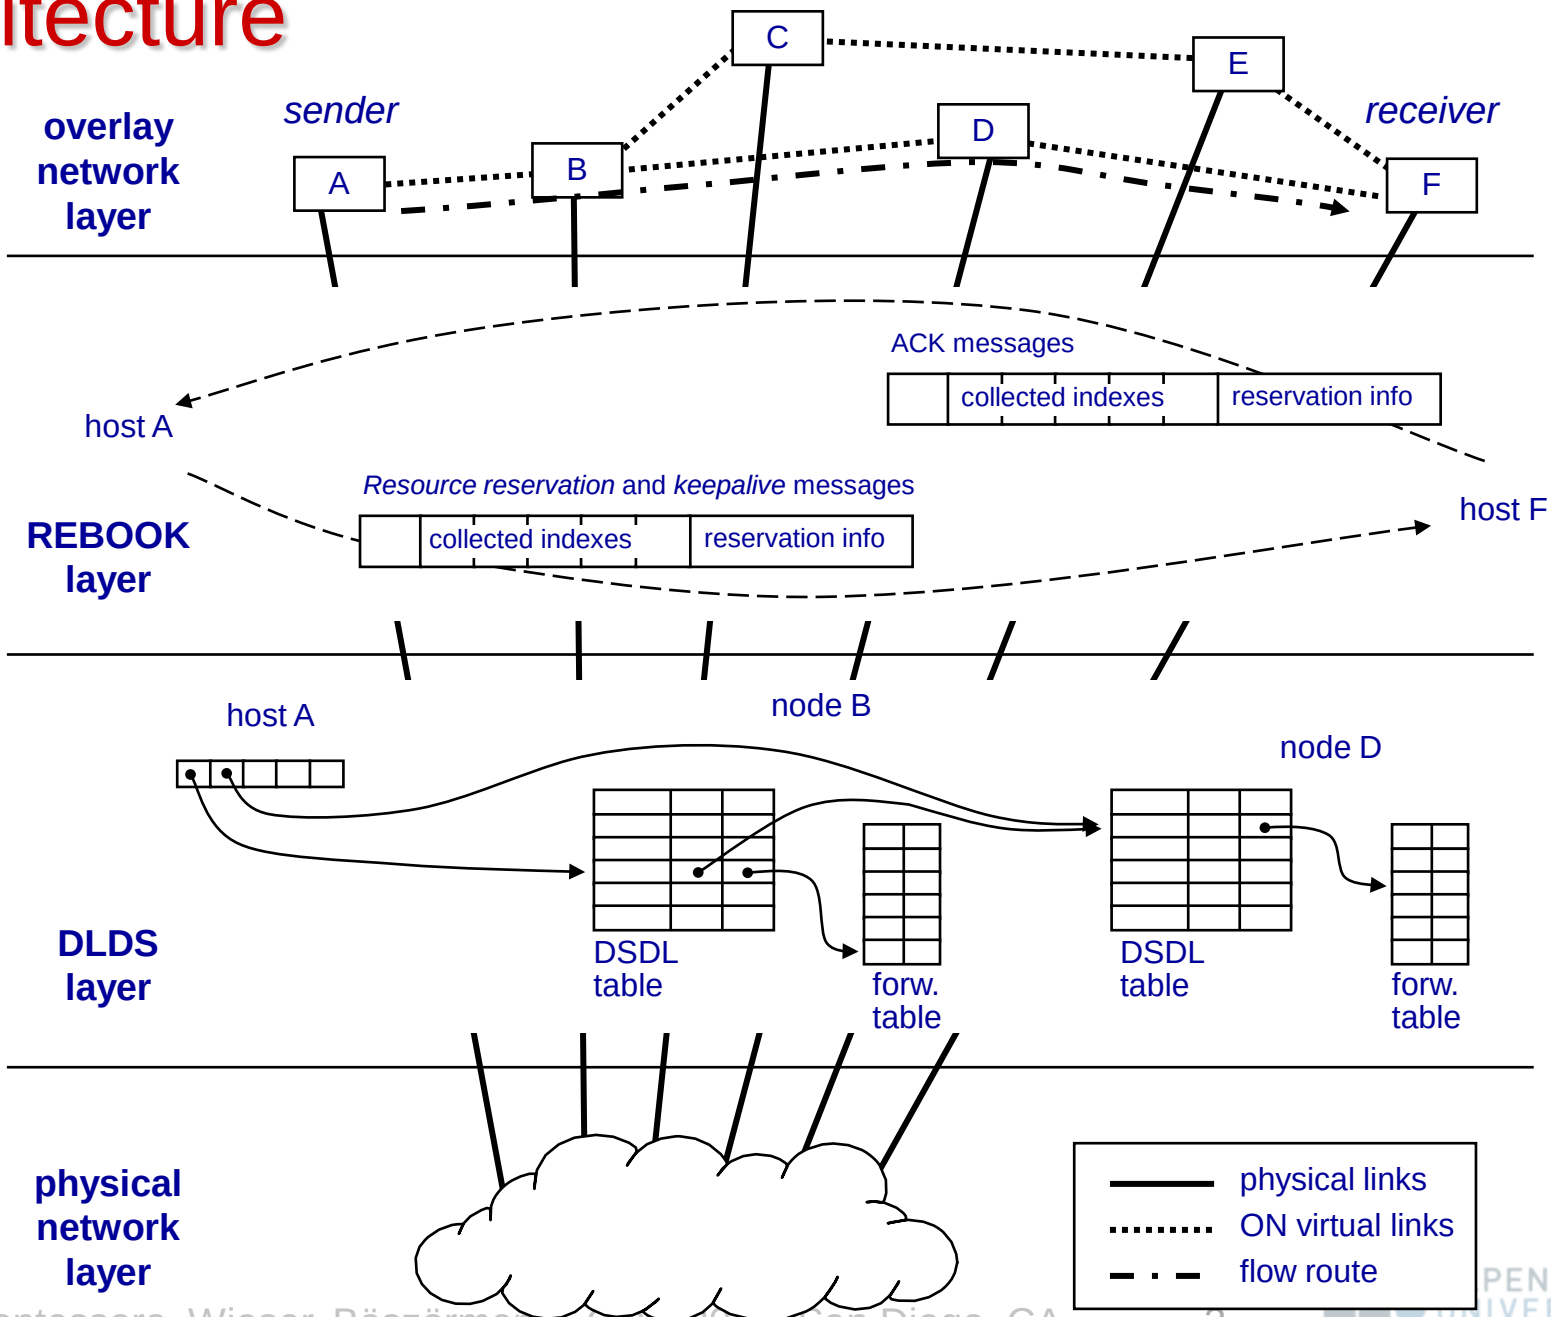

REBOOK/DLDS

- REBOOK
  - deterministic, dynamic, per-flow resource reservation
  - It IS NOT another reservation protocol
  - It IS a distributed algorithm for efficient status information handling within intermediate nodes
- DLDS (Distributed Linked Data Structure)
  - the enabling algorithm

REBOOK: a deterministic, robust and scalable resource booking algorithm. (JONS 2010)

A novel algorithm for dynamic admission control of elastic flows. (FITCE 2011)

Distributed Linked Data Structures for Efficient Access to Information within Routers. (ICUMT 2010)

Efficient Management and Packets Forwarding for Multimedia Flows. (JONS 2012)

DLDS (Distributed Linked Data Structure)

During setup: “pointers” collection

- store resource reservation information in routers

  AND

- keep track of pointers (memory addresses or indexes in tables) along the path

Afterwards

- use the pointers to access status information without searching

→ *constant cost to access the per flow resource reservation info*

→ *virtual circuit performance for packet forwarding*

Resource reservation and pointers collection

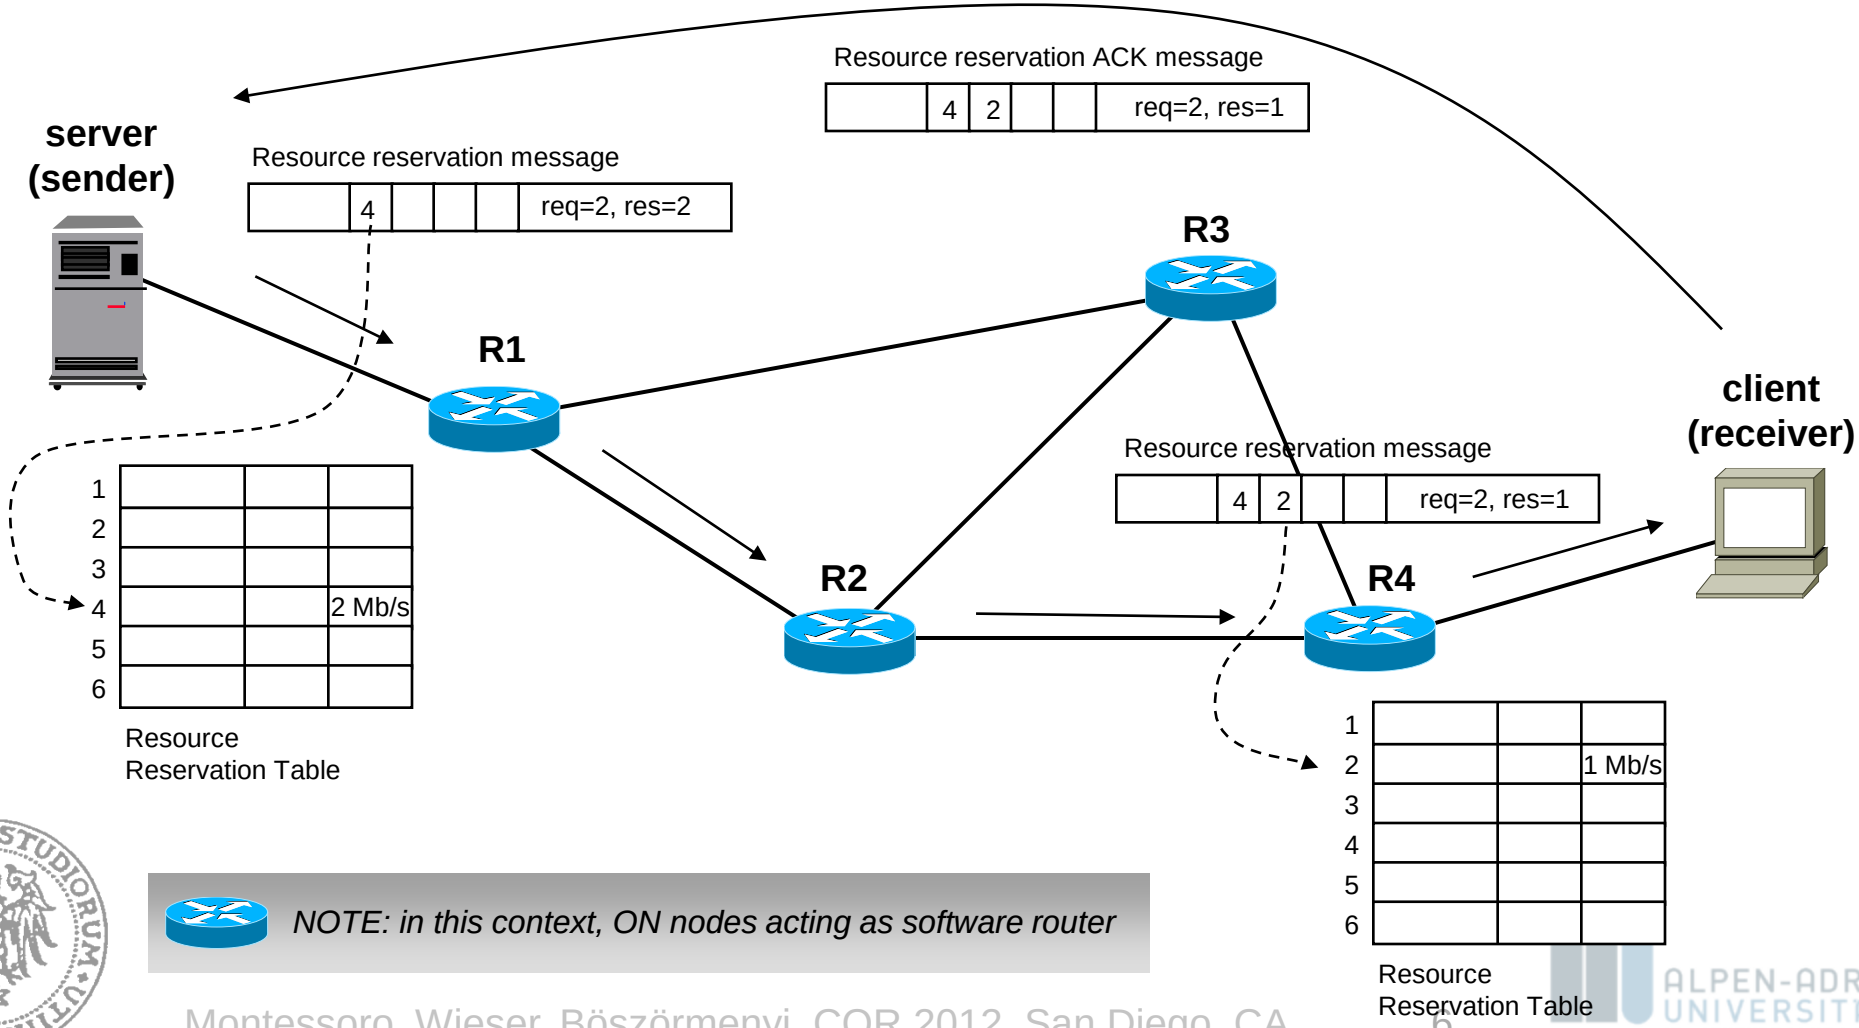

NOTE: in this context, ON nodes acting as software router

Fast packet forwarding

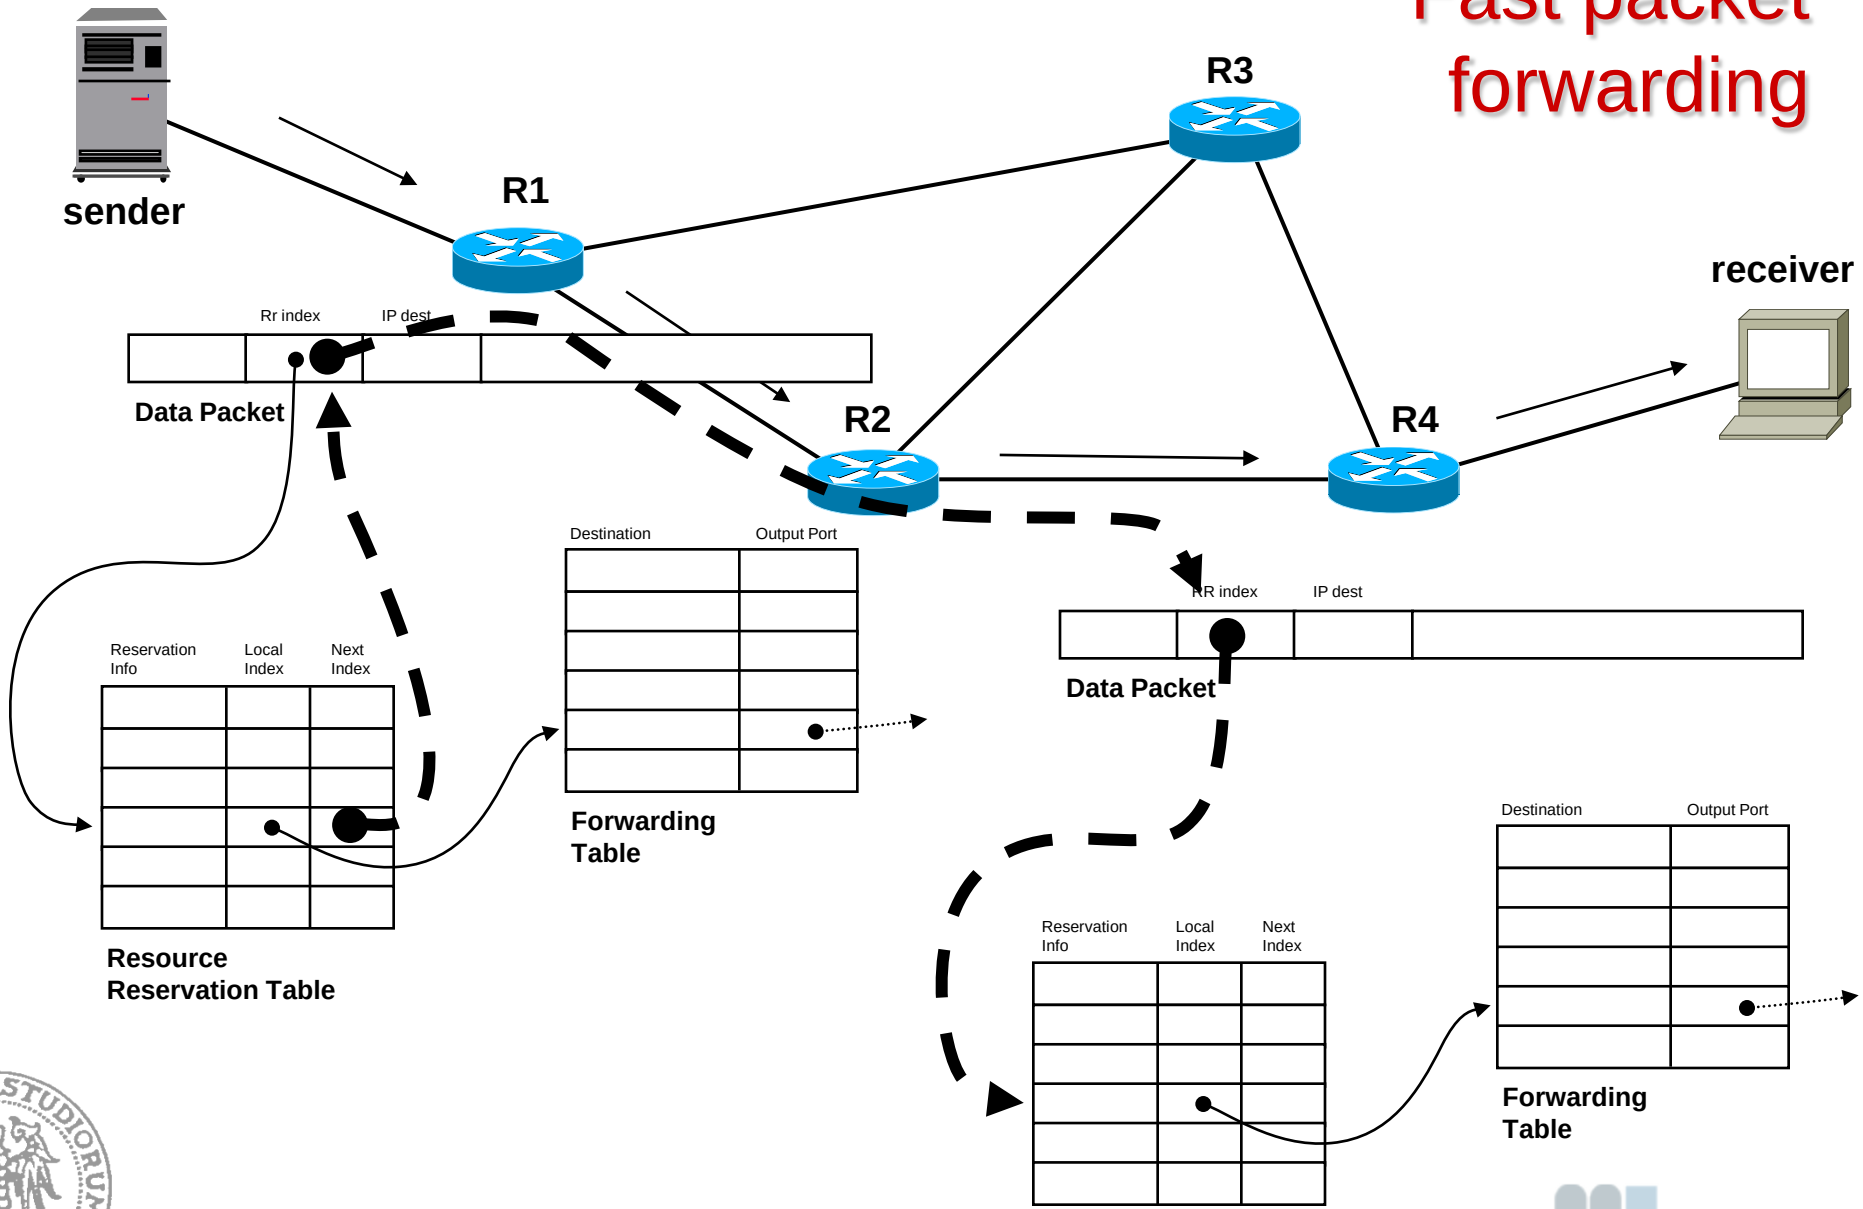

A Few Problems

- Route changes, disappearing flows, end nodes or routers faults
  - High speed consistency check
  - Low priority table cleanup process

- Need to dynamically change assigned resource amounts
  - Partial release
  - Distributed control function for optimality and fairness

Does it work?

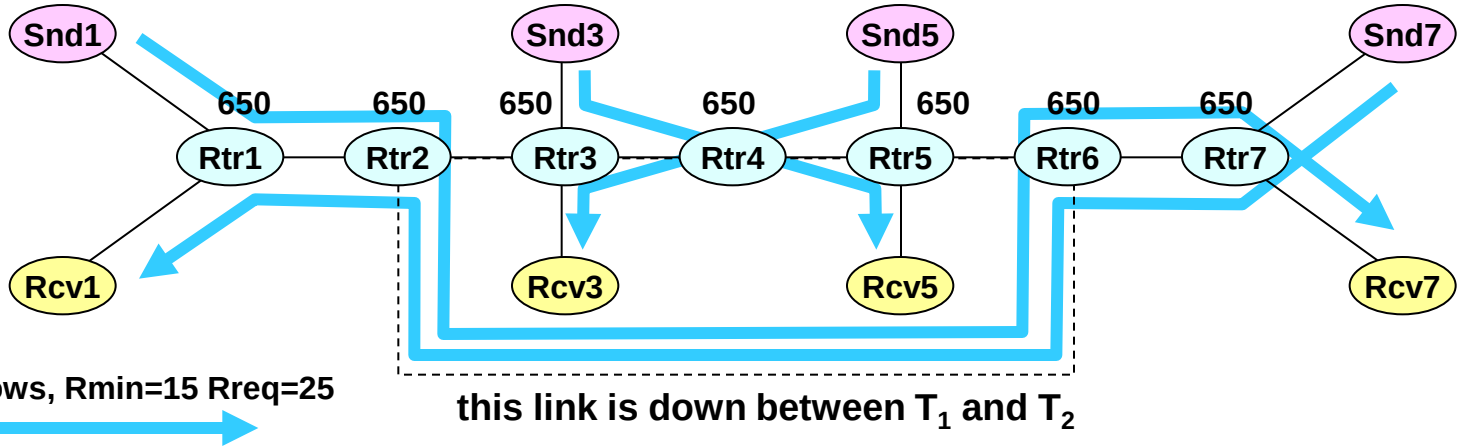
Snd1
Snd3
Snd5
Snd7
650
650
650
650
650
650
650
Rtr1
Rtr2
Rtr3
Rtr4
Rtr5
Rtr6
Rtr7
Rcv1
Rcv3
Rcv5
Rcv7
10 UDP flows, Rmin=15 Rreq=25
this link is down between T_1 and T_2

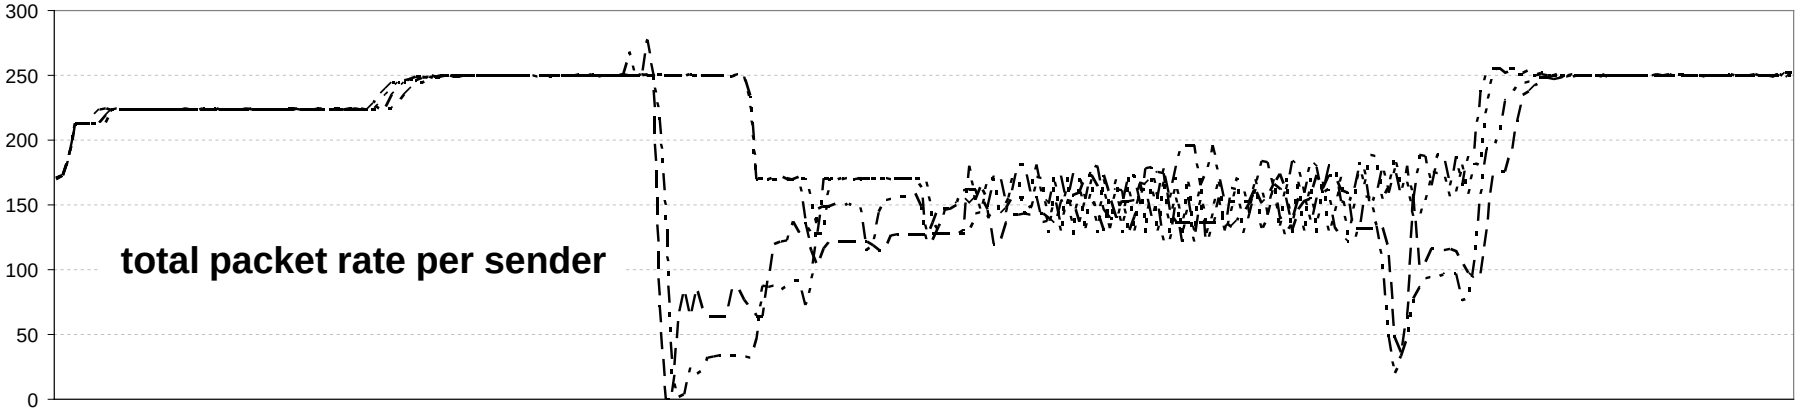
300
250
200
150
100
50
0
total packet rate per sender

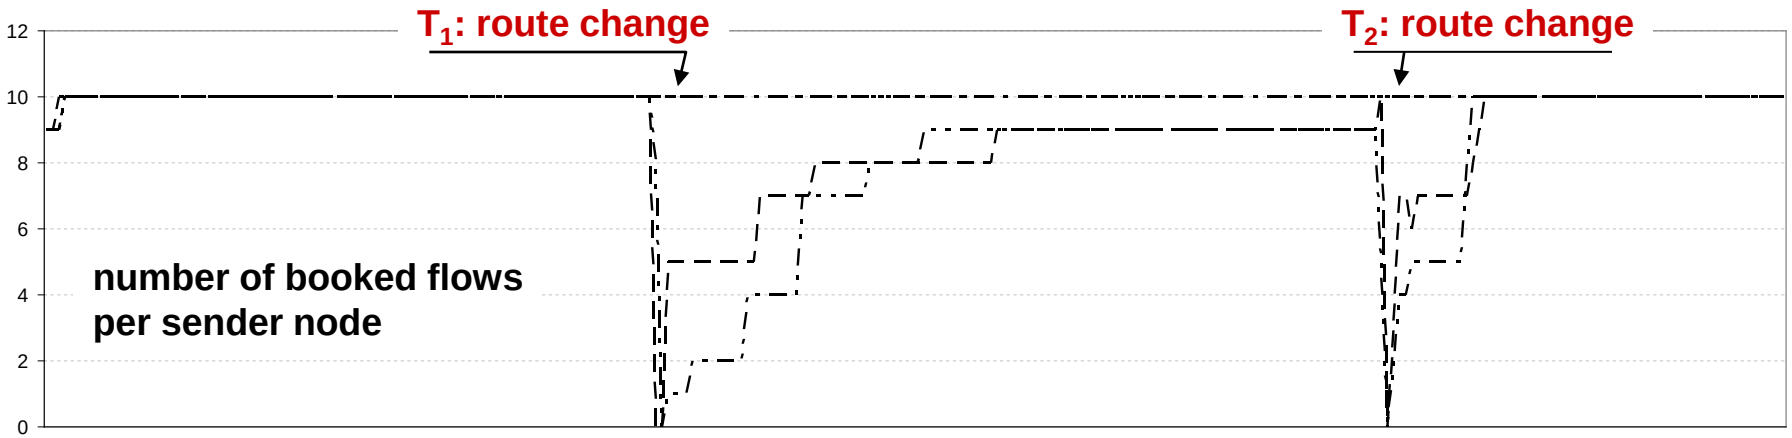
T_1: route change
T_2: route change
12
10
8
6
4
2
0
number of booked flows per sender node

Does it work? *(cont'd)*

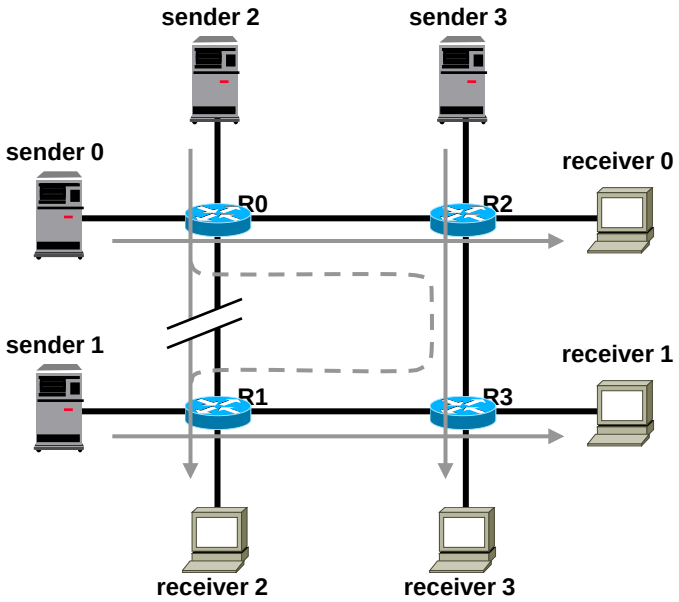

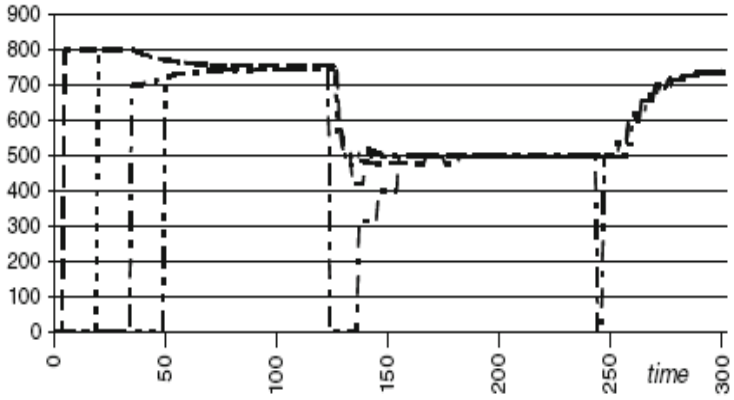

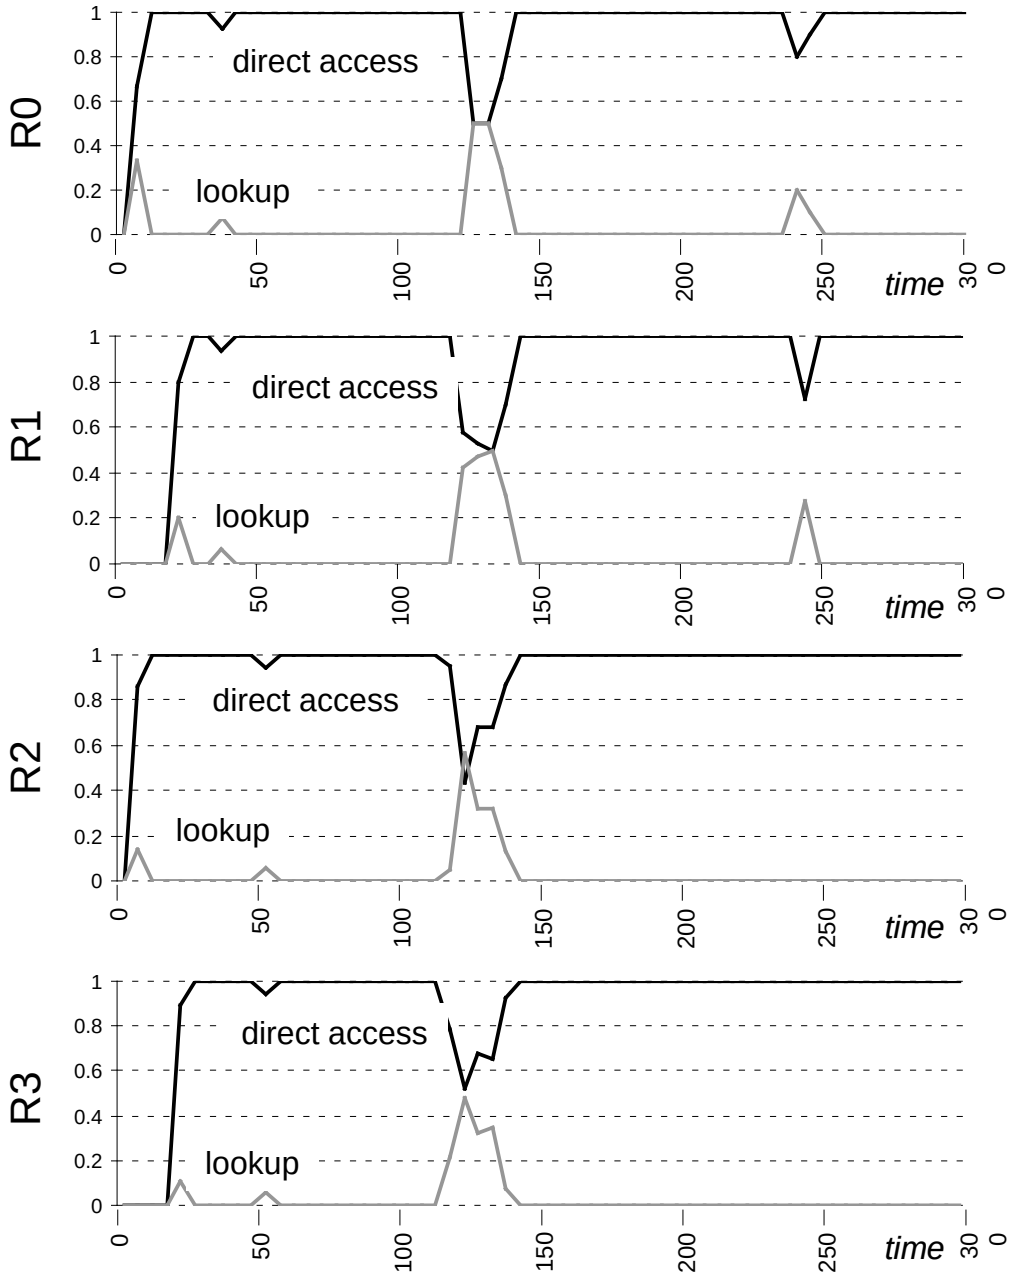

Flocks: Flexible and Self-Organizing Overlay

- Gossip-based protocol
  - Uses “Interest” to build the overlay
  - Idea: Interested nodes “stick” together
  - Interest in neighbours computed locally
- QoS estimation and QoS-aware routing
  - QoS-estimate to any node, fast and scalable to large networks
- Topology aggregation to generate a hierarchical representation for scalability
  - Nodes only need a small local view

Flocks: Interest-based construction of overlay networks. (MMEDIA10)

Decentralized topology aggregation for QoS estimation in large overlay networks. (NCA2011)

Combined Overlay

- Why combine Flocks with REBOOK?
  - Flocks provide a flexible and scalable overlay
  - REBOOK allows to keep track of bandwidth allocation and avoid overcommitting links
  - DLDS enable fast routing table access and higher performance once routes are established (O(1))

Evaluation: Overlay Setup

- Overlay with 160 Flock-REBOOK nodes
  - Nodes only have a small local view
  - Two random groups connected by bottleneck
  - Within each group: Interest in high bandwidth

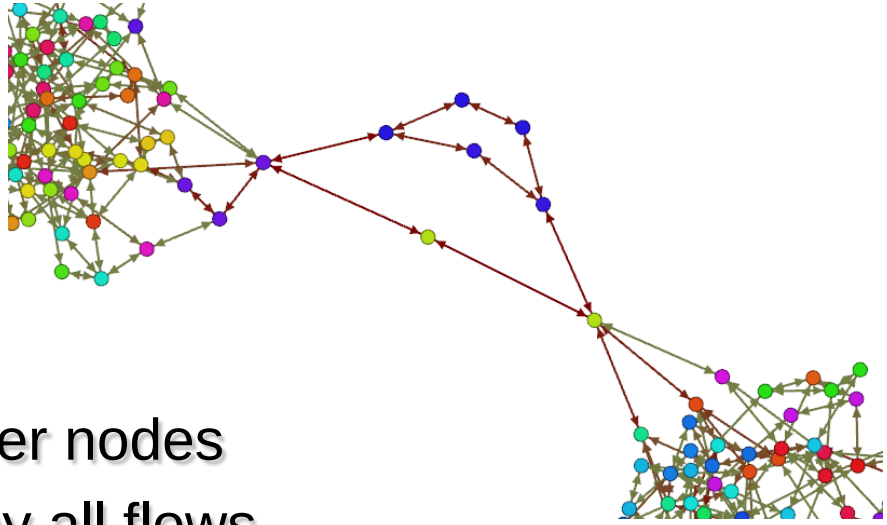

- Small local view of only 4 other nodes
- Artificial bottleneck crossed by all flows
  - Alternative paths outside local view
  - Must refer to aggregated information
  - REBOOK tracks allocated resources
- New reservation
  - Use path with the highest estimated bandwidth

Experimental Results: Admitted Flows

- Admitted flows increased significantly
  - Overlay avoids overcommitting links
  - Discovers alternative path w. QoS estimation

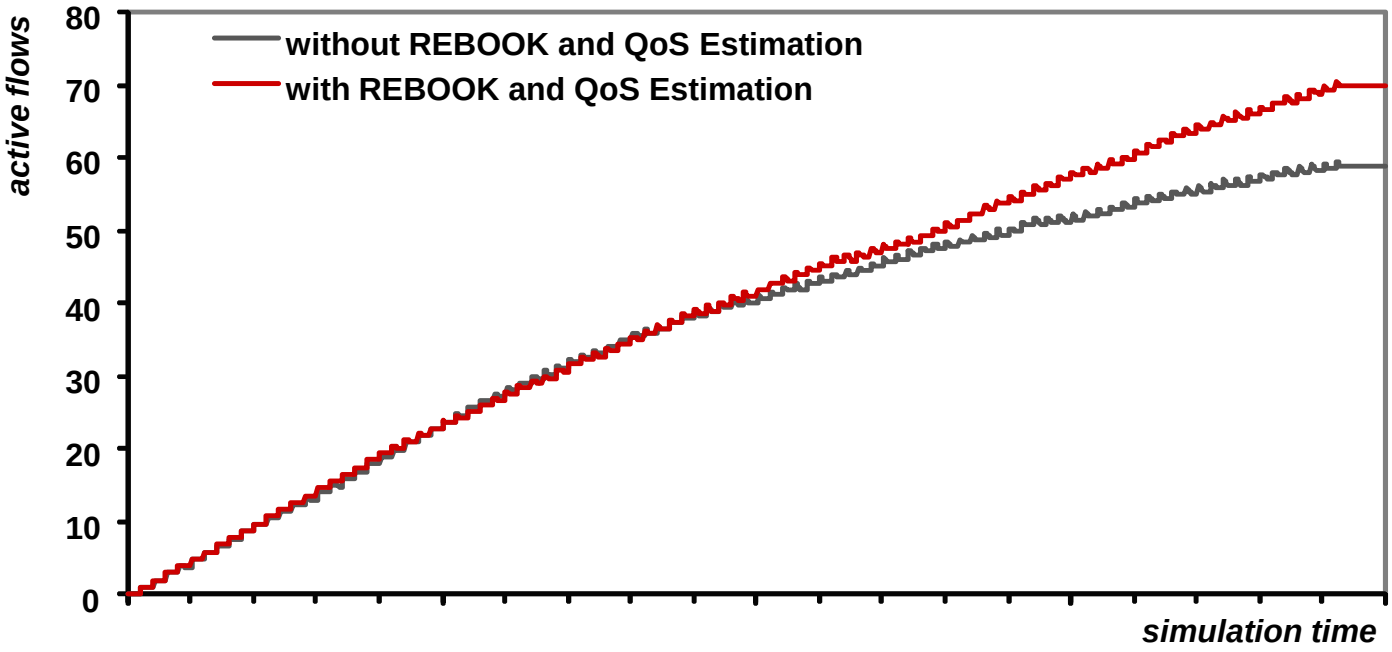

Experimental Results: Admitted Flows

- Bandwidth per flow increases as well
  - Overlay uses alternative paths used before REBOOK needs to reduce resources

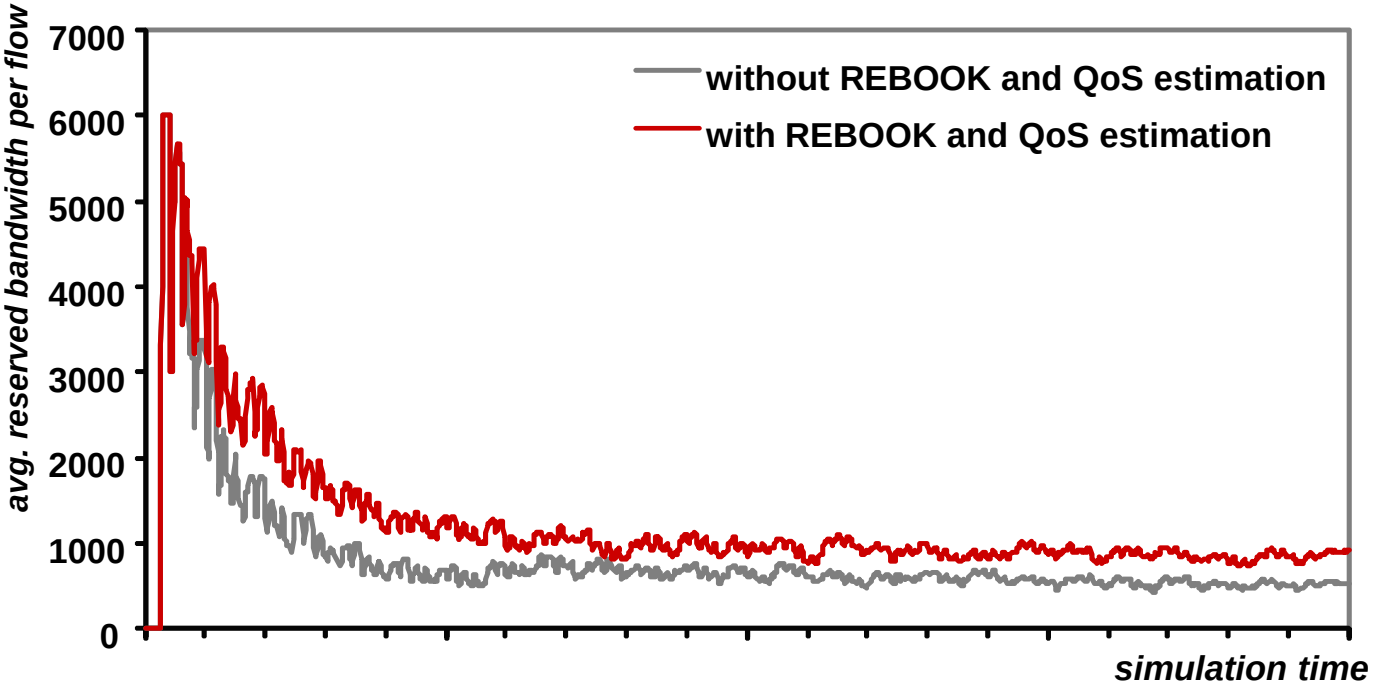

Experimental results *(cont'd)*

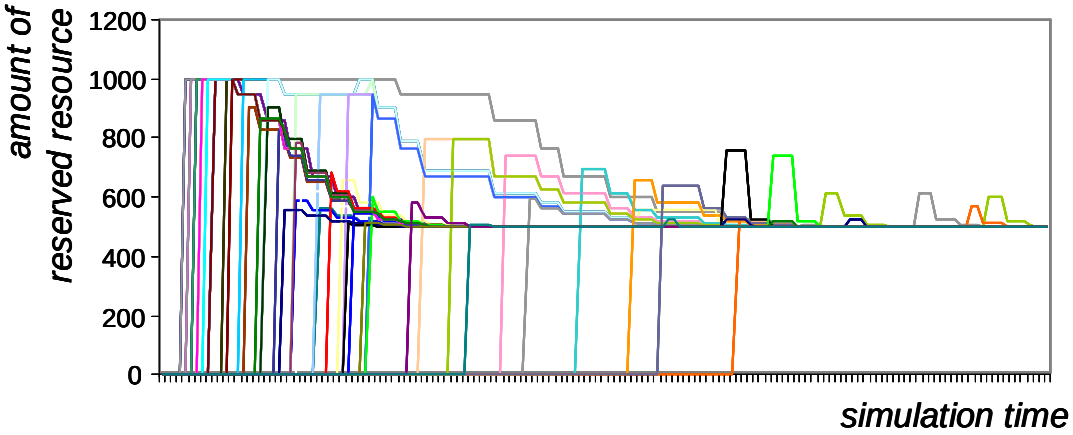

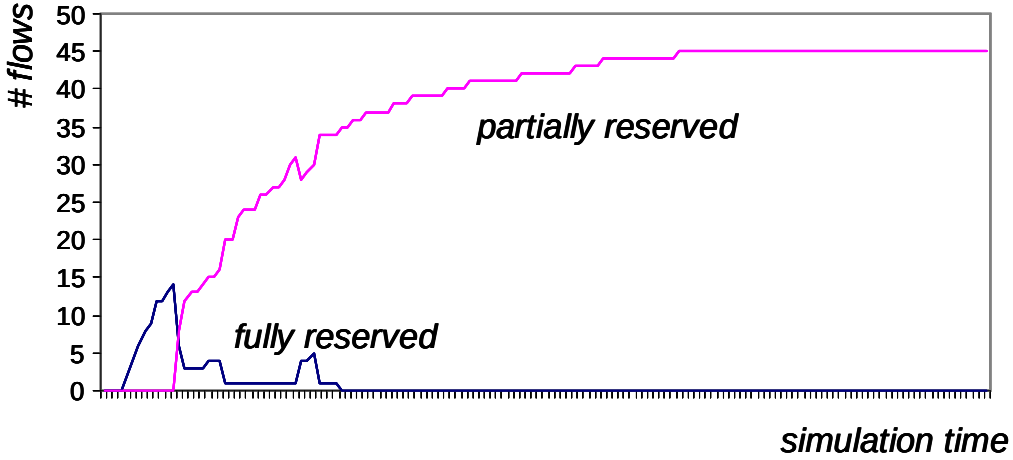

Resource reservation in a small congested network

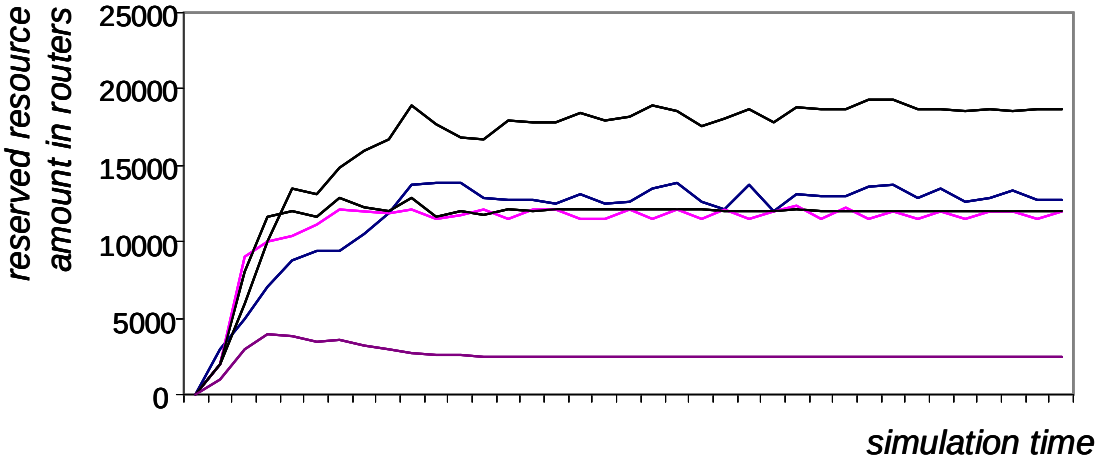

Performance

Activity	CPU time	
Resource reservation setup	200 ns	per flow
Keepalive message handling	100 ns	per flow
Resource reservation release	25 ns	per flow
Forwarding table access	10.6 ns	per packet

Conclusion

- Flocks provide a flexible and scalable overlay for multimedia delivery
- Hierarchical aggregation and QoS estimation of Flocks allow discovering alternative paths using small, local views
- REBOOK allows a scalable way for overlays to keep track of aware of resource utilization
- REBOOK and DLDS enable efficient and fast performing routing for overlays

Thank you!

www.montessoro.it

Flocks: Interests

- Each node has two sets of name-value pairs
  - May differ from node to node, and over time
  - Used to rank neighbours
- Property-Set (“what I have”)
  - Shared with neighbours: virtual, uncertain, inheritable
- Interest-Set (“what I want”)
  - Not shared, evaluated locally

FLOCKS PERFORMANCE: TRAFFIC

- Average incoming traffic per node (kbit/s)

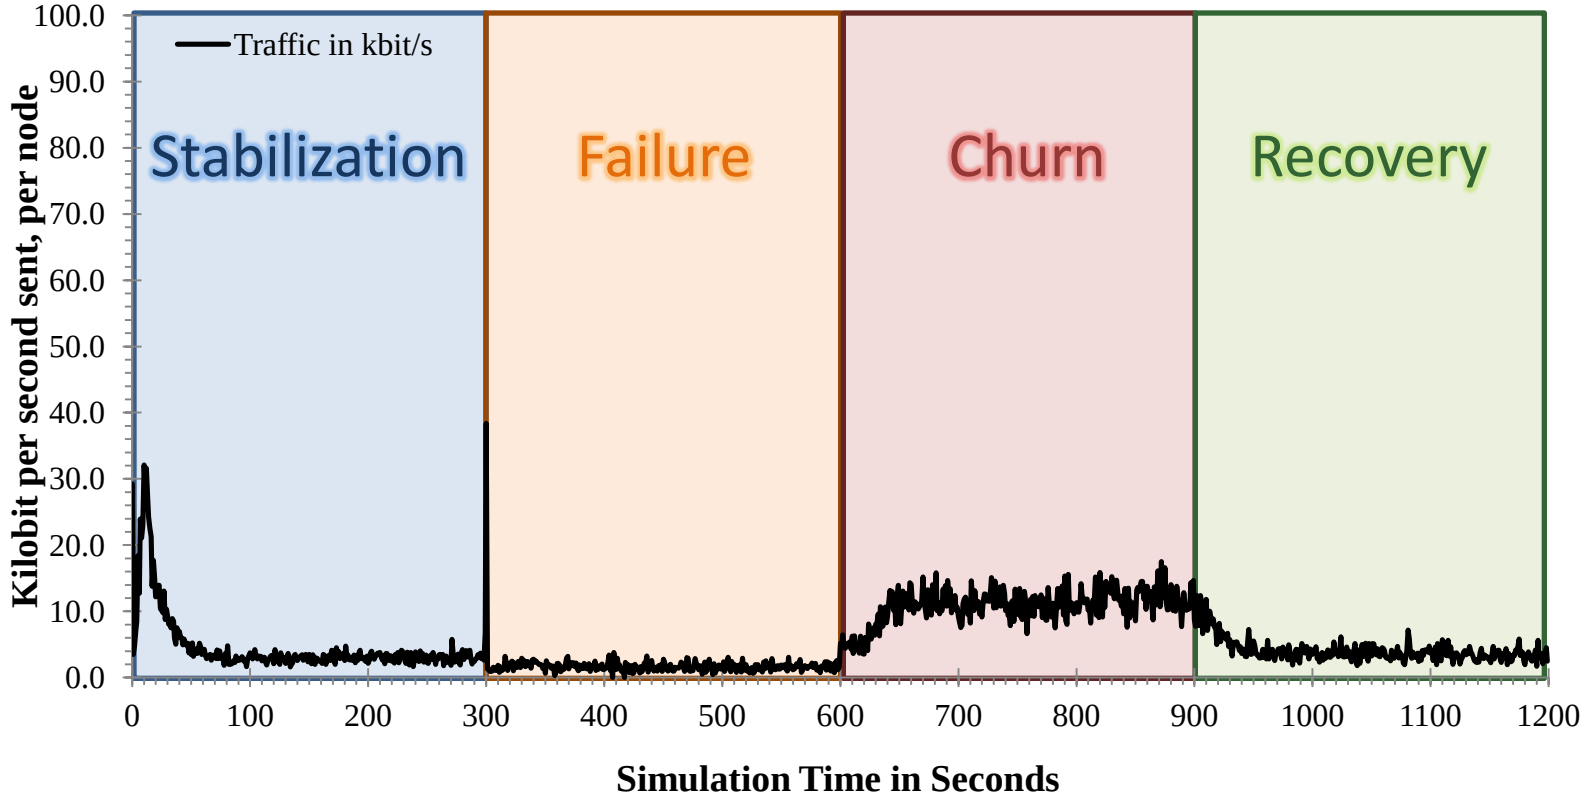

PERFORMANCE: QOS

- Measures closeness to the global optimum

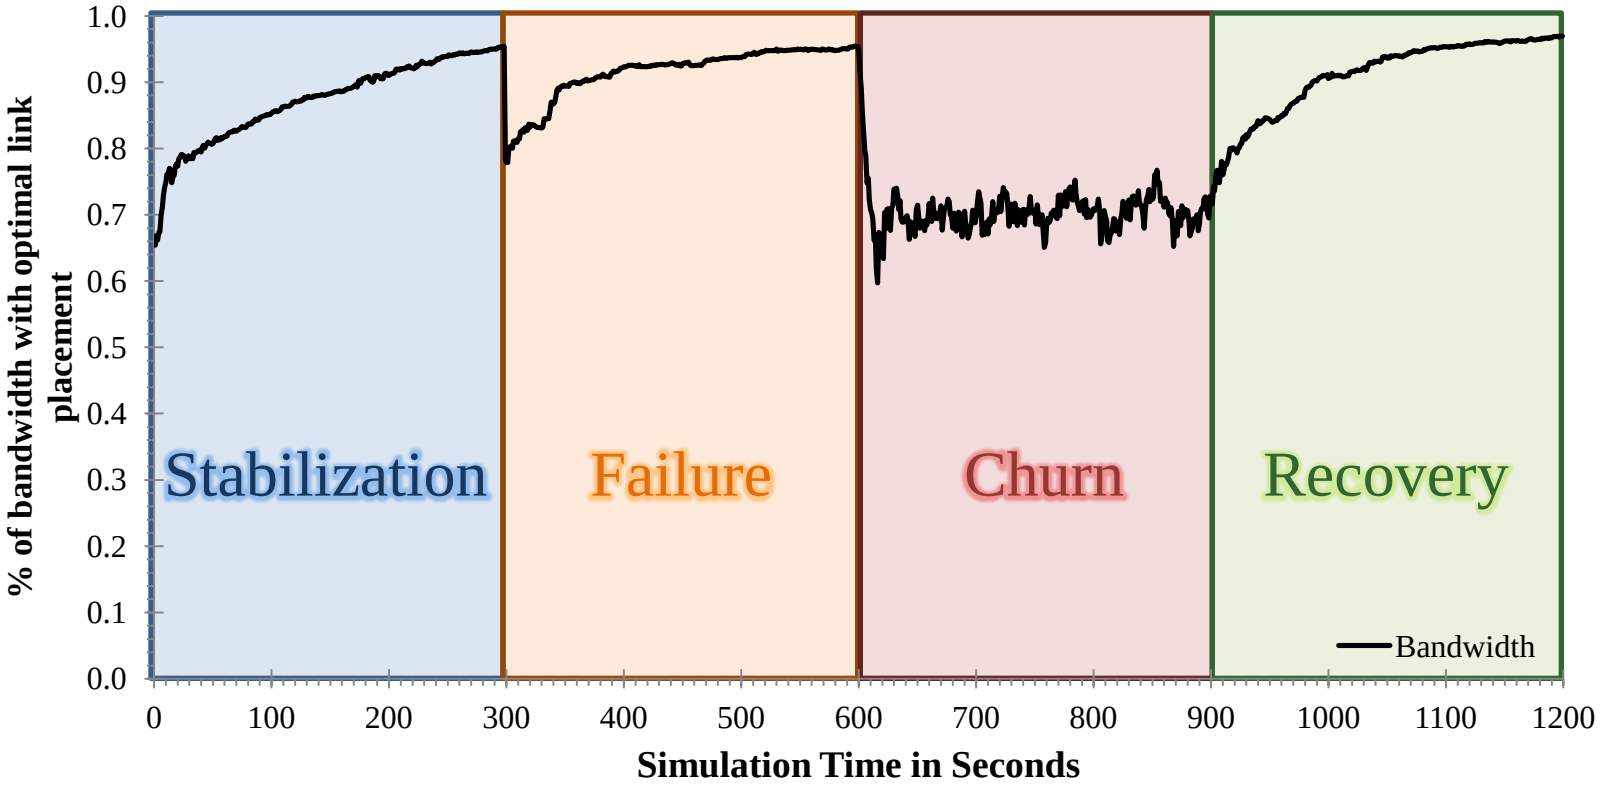

Thank you!

www.montessoro.it