Studying Multi-rate Multicast Congestion Control with Explicit Router Feedback

Kiyohide NAKAUCHI, NICT, Japan
Katsushi KOBAYASHI, NICT, Japan

PFLDnet 2006
Feb. 2, 2006

Explicit Router Feedback

- Promising way to enhance end-to-end performance
- Challenging networks require more sophisticated fine-grained router feedback
  - Link capacity (PTP)
  - Available bandwidth (TCP Quick-Start, XCP)
  - Queue length (VCP)
  - Queue size
  - Loss rate (ETEN)

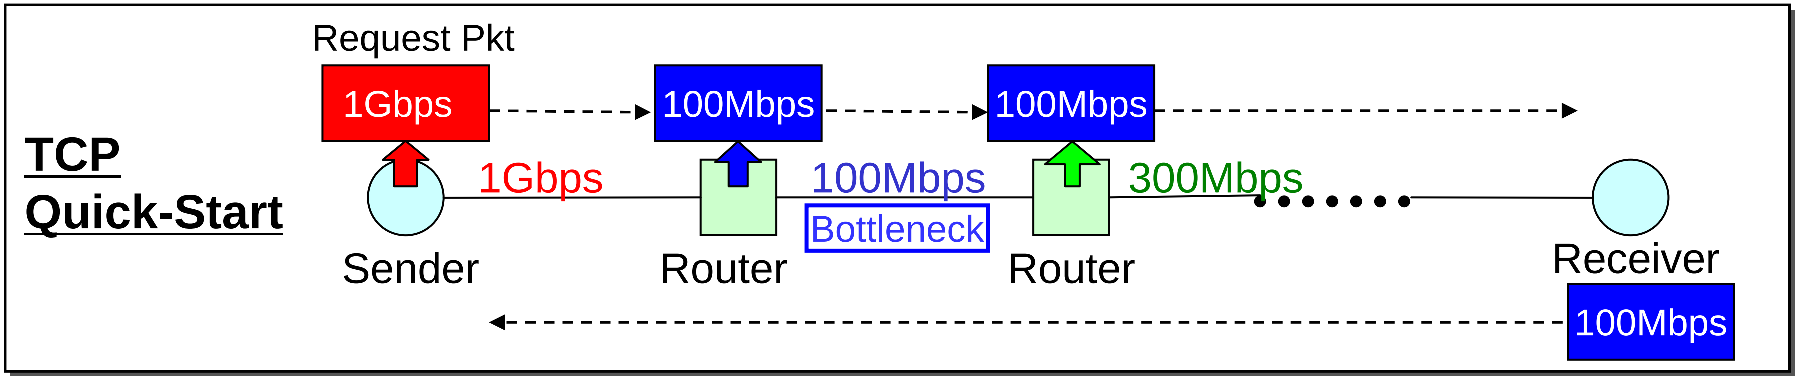

Per-path Feedback Is Not Enough

- Per-path feedback
  - Existing schemes can *independently* provide only the maximum or minimum value along the path
  - Combining separate feedbacks for multiple status
- Per-hop feedback
  - "Queue size (of the router) on the bottleneck link"
  - Needed for precise parameter configuration

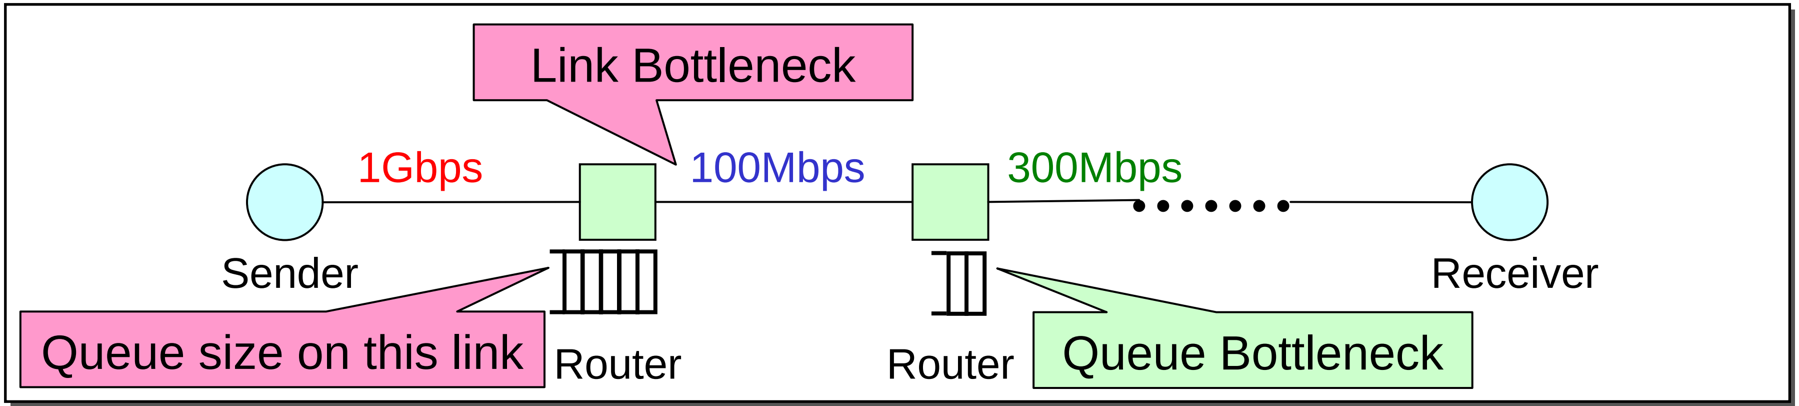

Outline

- SIRENS: a fine-grained and per-hop explicit router feedback framework
- A Case for Multi-rate Multicast
- Simulations
- Implementation Status
- Conclusions

SIRENS Framework

- Fine granularity: three "profiles"
  - LINK: link capacity, available bandwidth
  - LOSS: packet loss rate, link error rate
  - QUEUE: queue size, link delay
- Per-hop feedback
  - Captures a snapshot of the status or statistics specified by a profile on each hop
  - Interpretation of QoS semantics is posed on end-hosts
- SIRENS is only the notification scheme
  - How to use feedback information depends on end-hosts

Protocol Behavior (LINK Profile)

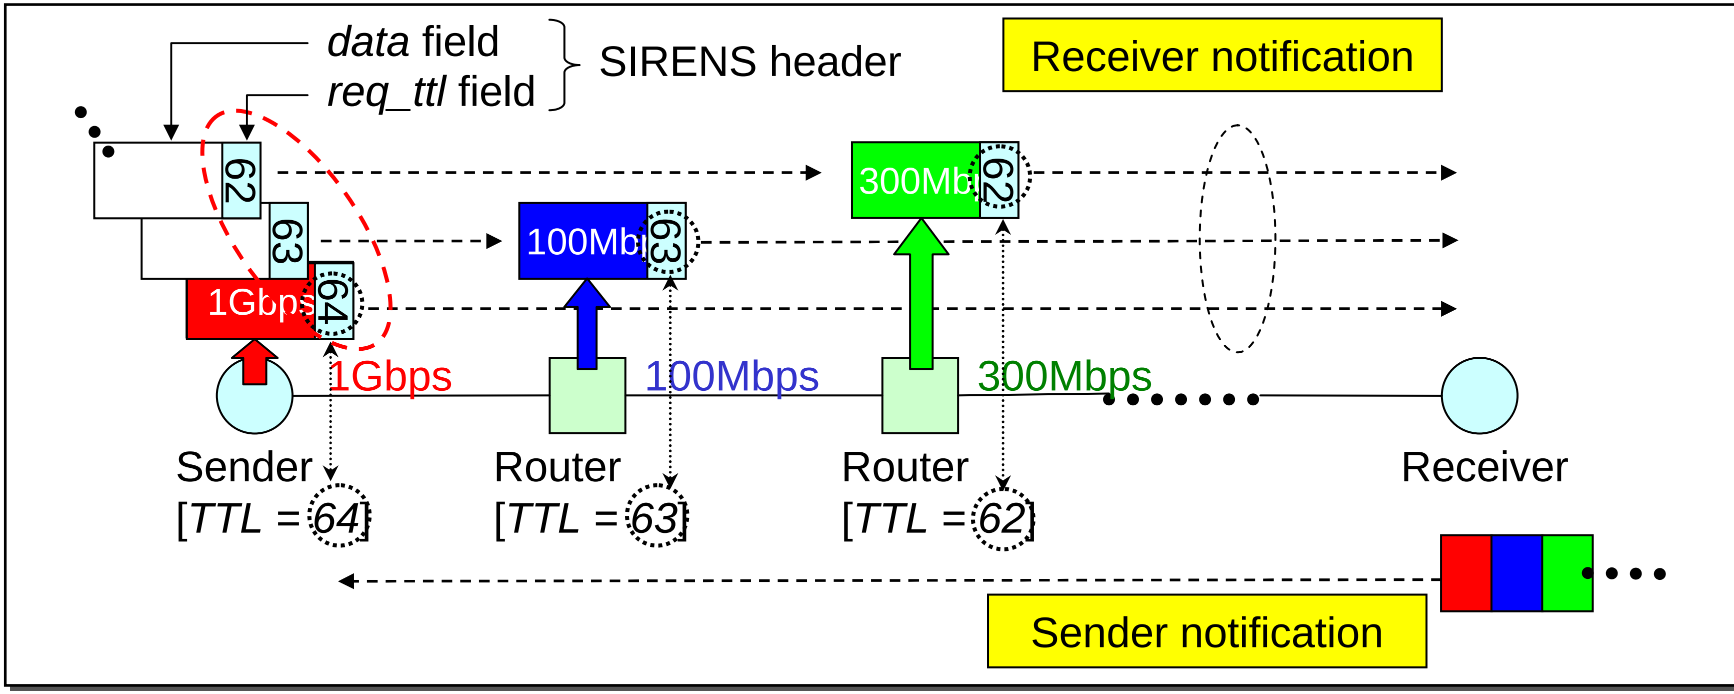

- Sender:
  - Sets *req_ttl* from the same value as IP *TTL* in the decreasing order
- Router:
  - Writes status specified by a profile into the corresponding packet
- Receiver:
  - Assembles cumulative feedbacks (and returns them to the sender)

Applications

- TCP variants for high-speed networks
  - Parameter configuration for TCP Limited Slow-start (PFLDnet 2005)
  - Parameter configuration for TCP variants
- Congestion control for wireless/mobile networks
- Multi-rate multicast congestion control (PFLDnet 2006)
- …

A Case for Multi-rate Multicast

- Multi-rate multicast
  - Heterogeneous receivers
  - Ex) layered multicast, simulcast
- Receiver-driven
  - Several multicast servers with different rate
  - Receivers independently determine the optimal multicast group

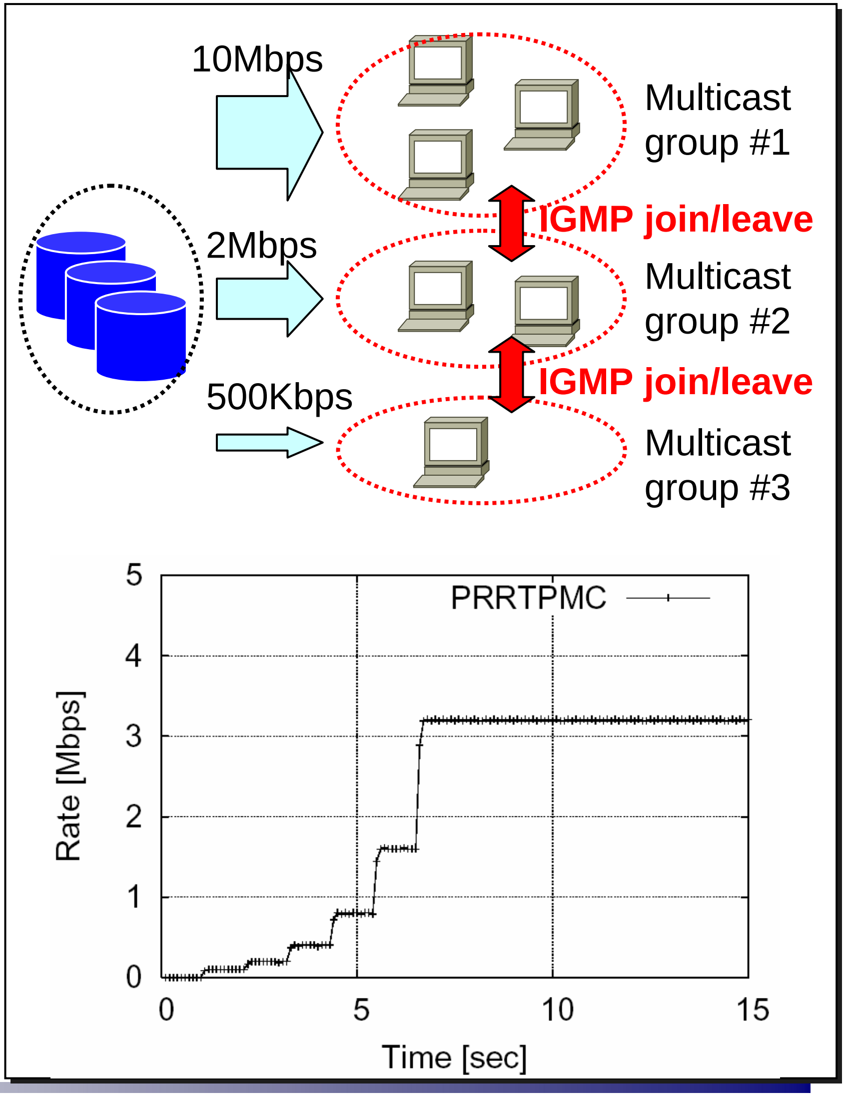

Traditional Issues

- Low responsiveness to dynamic traffic changes
  - "trial-join" and timer control
  - Incremental join / decremental leave
  - Limitation of bandwidth estimation mechanisms
  - IGMP leave delay
- Difficulty in TCP-Friendliness
  - Sending rate is CBR and course-grained

Explicit Multi-rate Multicast (EMcast)

- Feedback information for each receiver
  - Available bandwidth (SIRENS LINK profile)
- Measurement at each receiver
  - Loss rate

- Congestion control (case for simulcast)
  - Each receiver directly joins the *i*-th multicast group (instead of incrementally or decrementally)
  - Based on “TCP steady-state throughput”

EMcast Congestion Control

- Each receiver calculates TCP steady-state throughput at every time interval *T*
  - TCP steady-state throughput = 1.3 x MTU / (RTT x sqrt(R))

- Then, each receiver directly joins the *i*-th multicast group

> *i* = max(*j*), where *j* meets
> *j*-th sending rate <
> min(*AvailableBandwidth*, TCP steady-state throughput)

Simulation Setup (*ns2*)

- # of multicast groups : 10
- Rate of *i*-th multicast group : 100x2^(*i*-1) [Kbps]
  (Maximum rate = 51.2Mbps, *i*=10)
- Queue size : 100 [pkts]

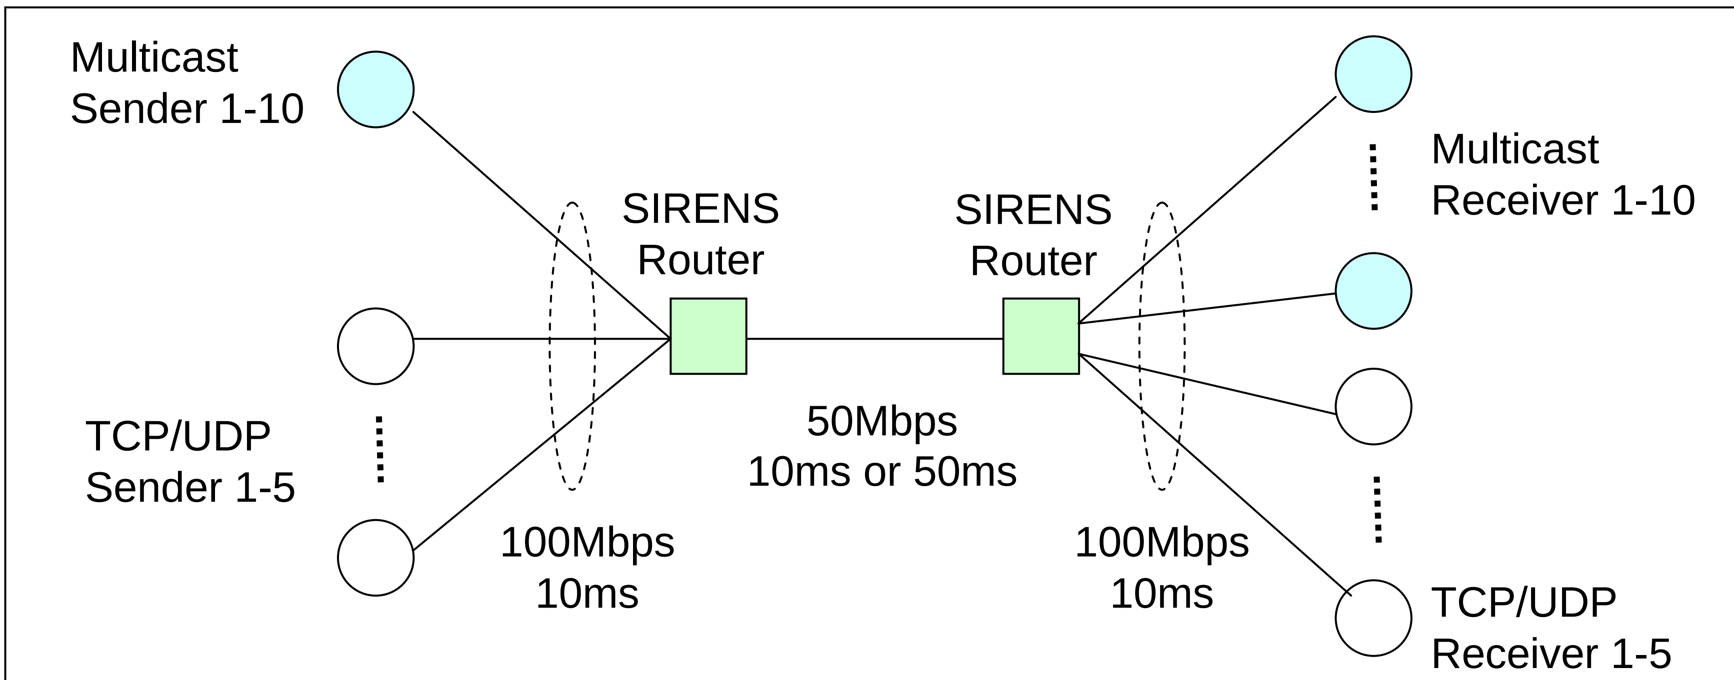

TCP-friendliness (1/2)

- EMcast with a single receiver competes with TCP flows

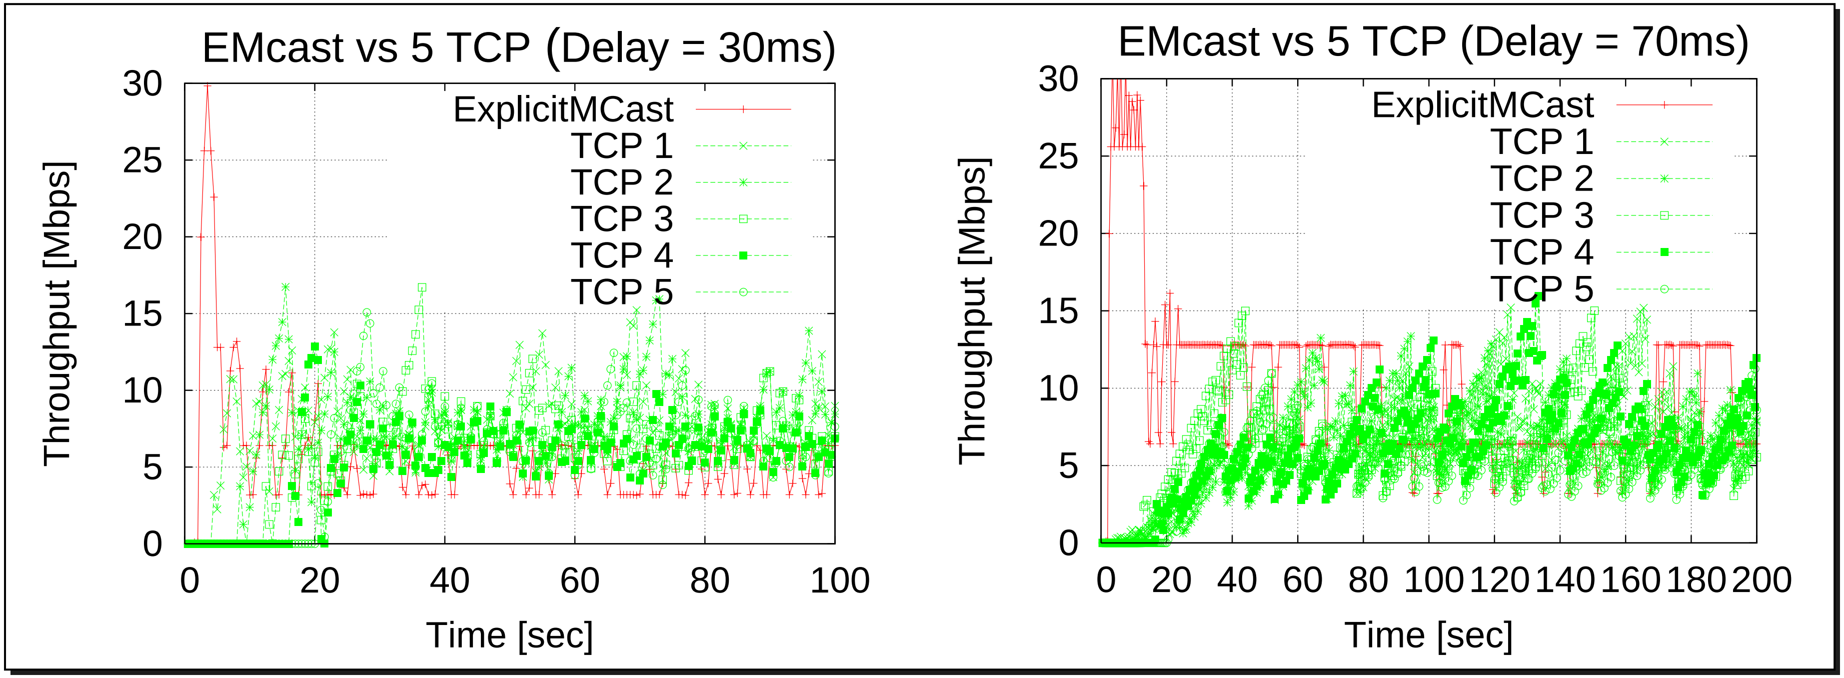

TCP-friendliness (2/2)

- EMcast with 10 receivers competes with TCP flow(s)
- EMcast receivers join at 2-20 sec
- Slow convergence in low loss rate environments

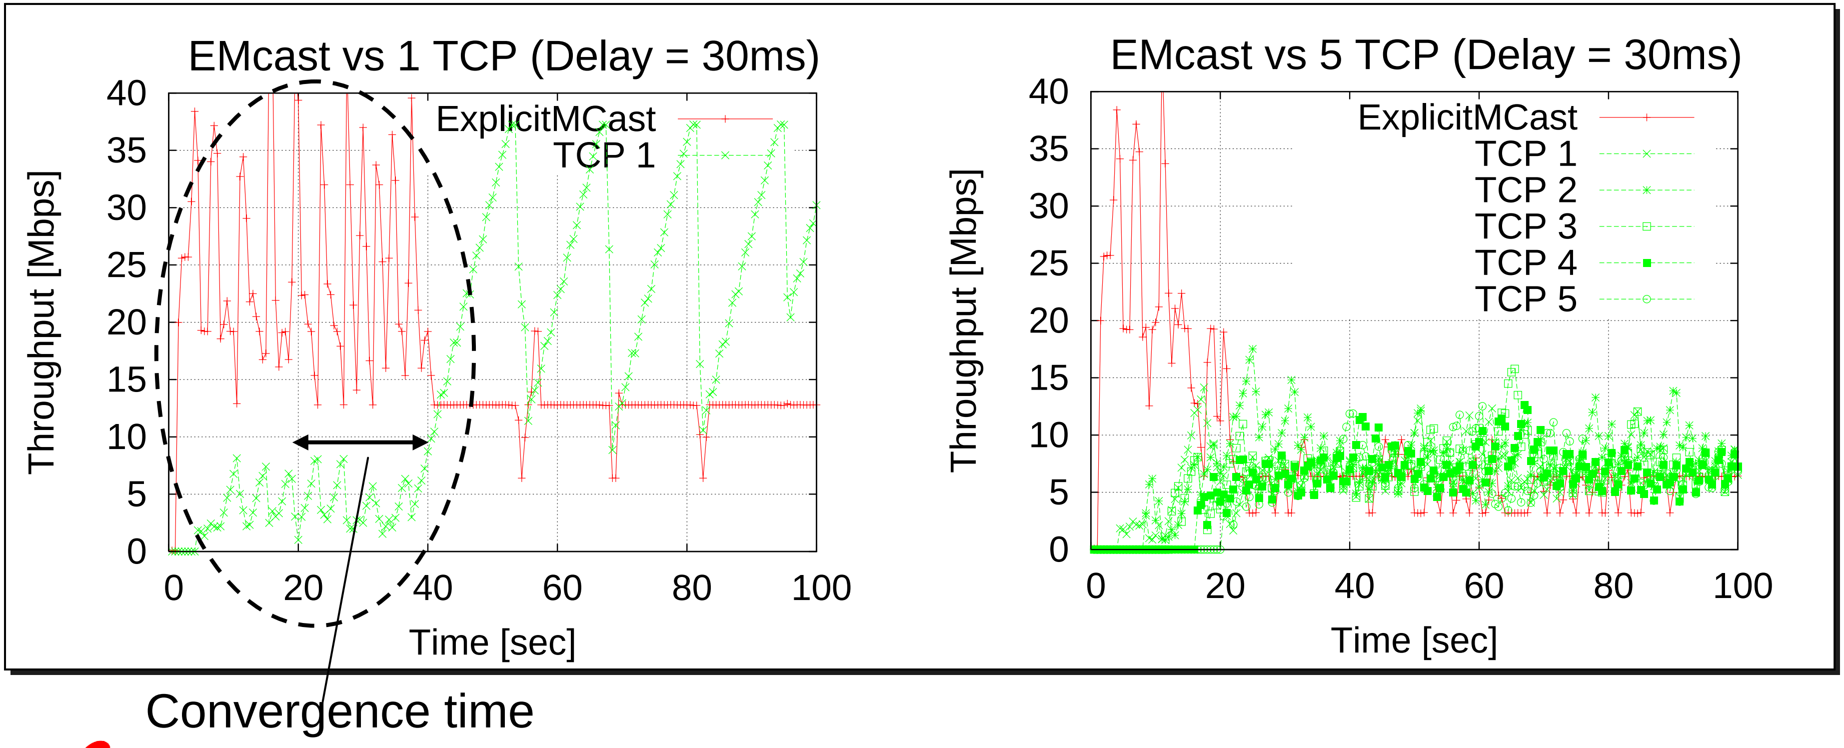

Responsiveness

- Comparing with Receiver-driven Layered Multicast (RLM)
- CBR cross traffic is generated at 50-80 sec

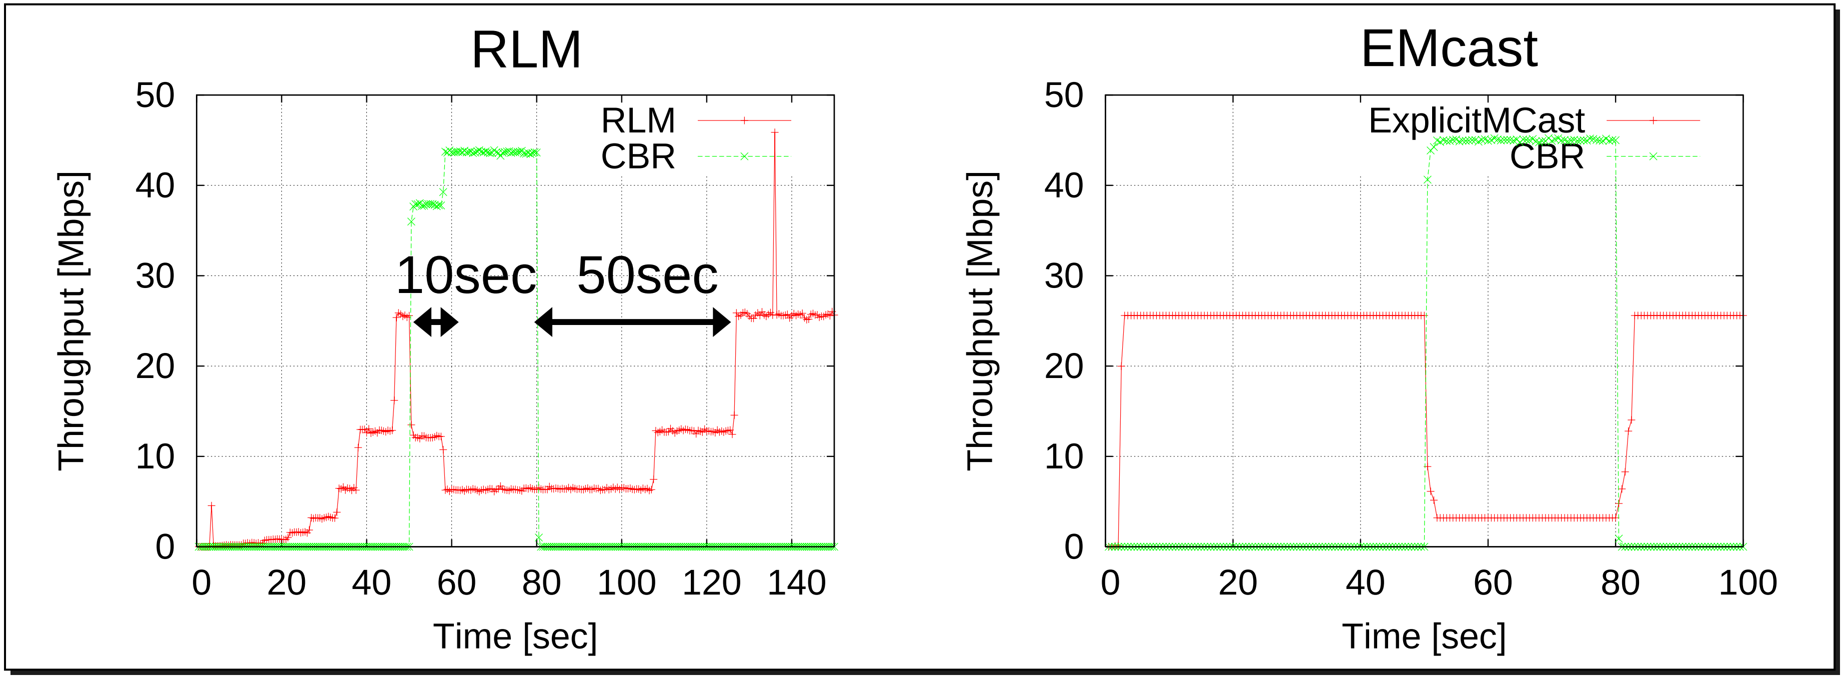

Implementation Status

- Iperf client for each multicast group with different sending rate
- Iperf server with EMcast congestion control

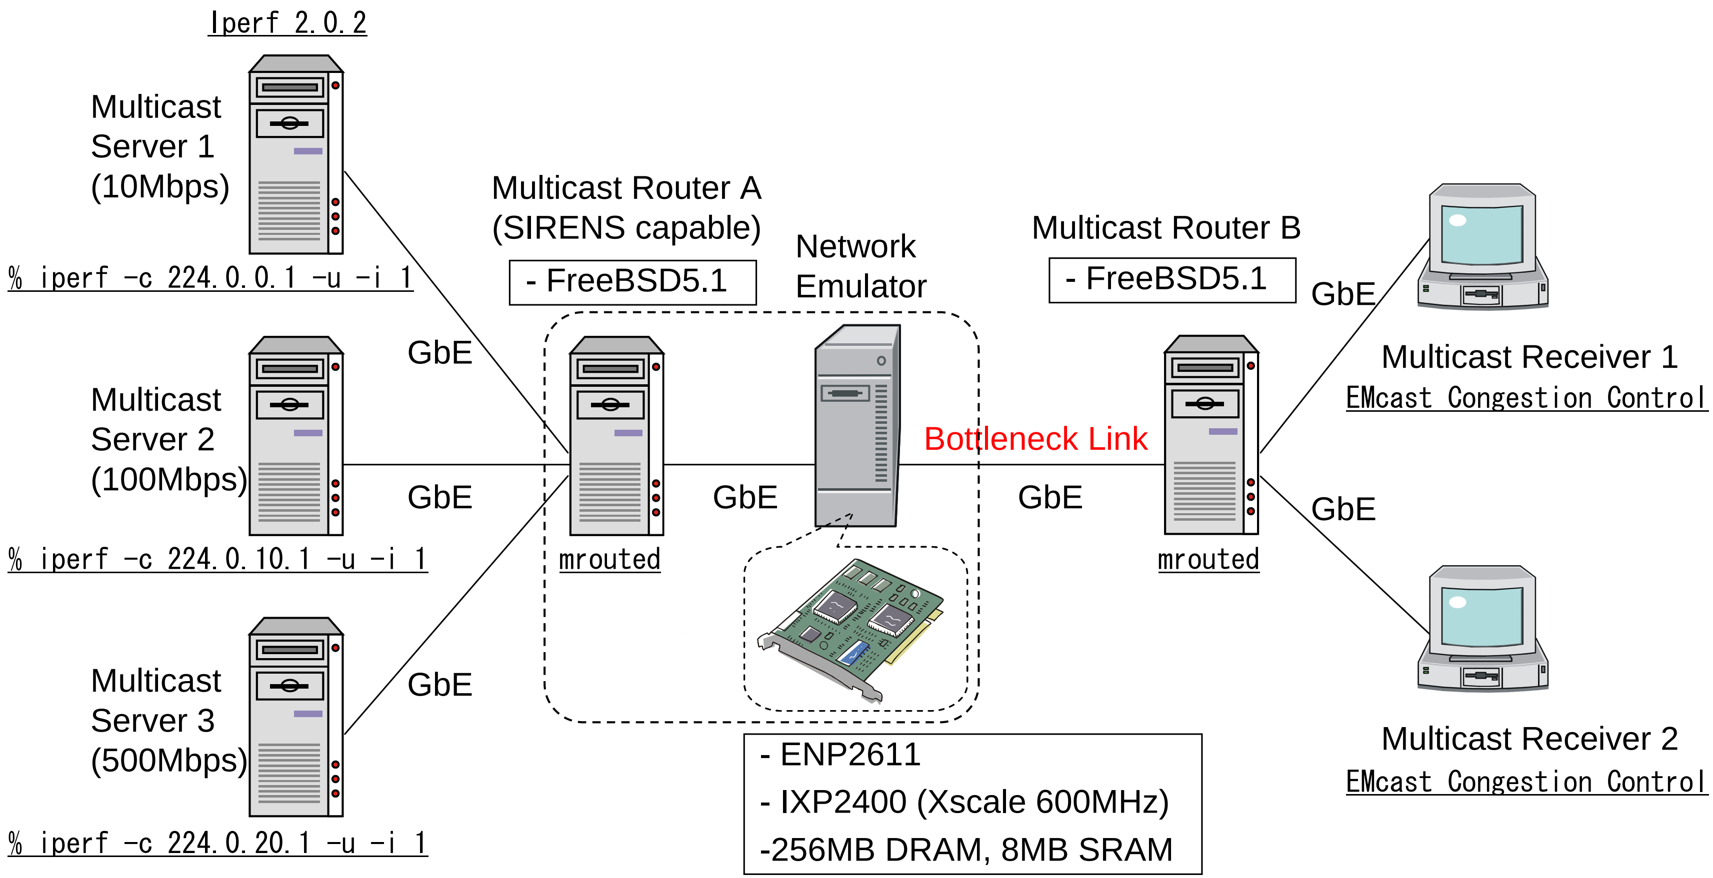

Conclusions

- We overviewed SIRENS, a fine-grained and per-hop explicit router feedback framework
- We proposed EMcast, multi-rate multicast congestion control with SIRENS
- We evaluated the performance by simulations
  - EMcast achieved TCP-friendliness using TCP steady-state throughput
  - EMcast achieved fast responsiveness using SIRENS (available bandwidth)