Dino López Pacheco
University of Nice
Sophia-Antipolis, Fr.
dino.lopez@unice.fr

Tuan Tran Thai
University of Toulouse
Toulouse, Fr.
tuan.tran-thai@isae.fr

Emmanuel Lochin
University of Toulouse
Toulouse, Fr
emmanuel.lochin@isae.fr

Fabrice Arnal
Tales Alenia Space
Toulouse, Fr.
fabrice.arnal@thalesaleniaspace.com

Towards an incremental deployment of ERN protocols: a proposal for an E2E-ERN hybrid protocol

Controlling Congestion

- Congestion can lead to data losses and to a decrease of the performance of the flows
  - Prevent congestion collapse
- TCP (in its New Reno variant) has been the most widely deployed congestion control protocol
  - AIMD algorithm enables fairness
  - TCP is an End-to-End protocol
    - Only implemented at end hosts
    - Easy to deploy over different technologies

Current Challenges

- Internet is heterogeneous
  - Hardware
    - Slow/fast paths
    - Links with low/high capacity
    - wired/wireless links
    - Layer-3/Layer-2 equipments
  - Software
    - Different policies in routers
    - Different policies in end host

- Standard TCP performs poorly over paths with Large Bandwidth-Delay Product (LBDP)

- High speed TCP variants fulfill LBDP path but potentially increase unfairness

Controlling congestions in the core network

- Active Queue Management (AQM)
  - Routers do not communicate with end hosts
- Binary signaling (e.g. ECN).
  - Binary congestion signal from routers to end hosts
- Explicit Rate Notification (ERN) protocols
  - Each packet carries up useful information for feedback computation ($cwnd\_, RTT$)
  - Routers indicate to senders the more adapted emitting rate *(feedback)*
  - No states per flow

ERN protocols in a full compliant ERN networks

- No sawtooth behavior
- High link utilization
- Low buffer occupancy
- Equally share bandwidth

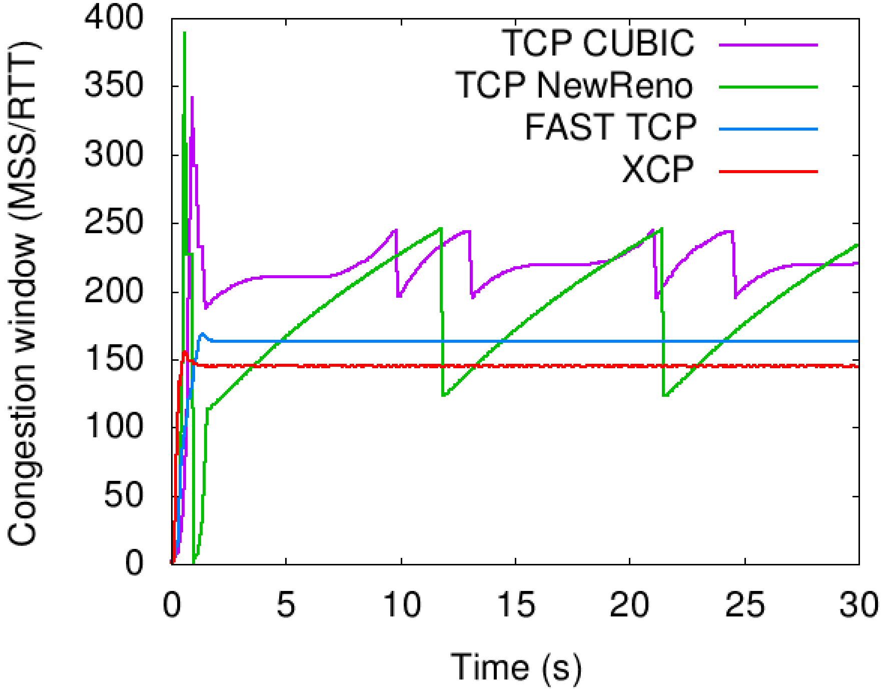

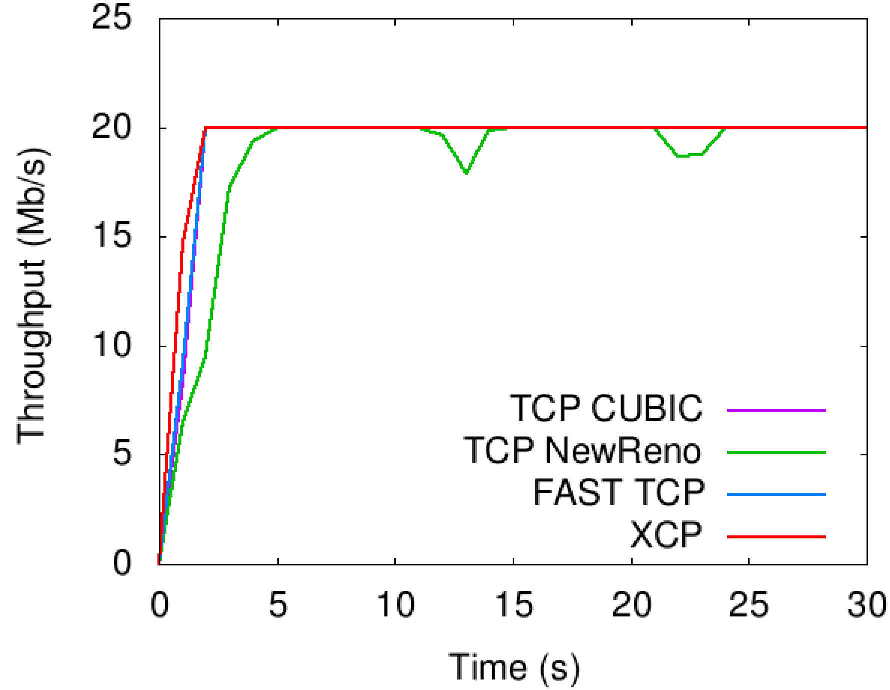

Interoperability of ERN protocols

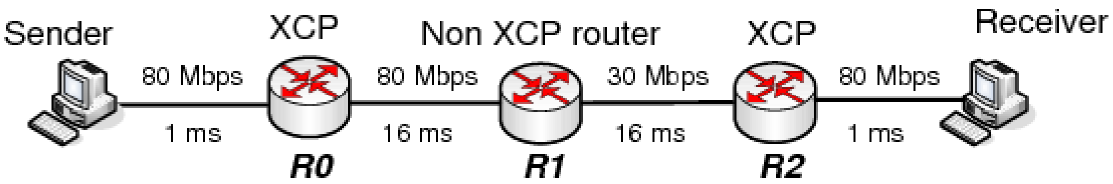

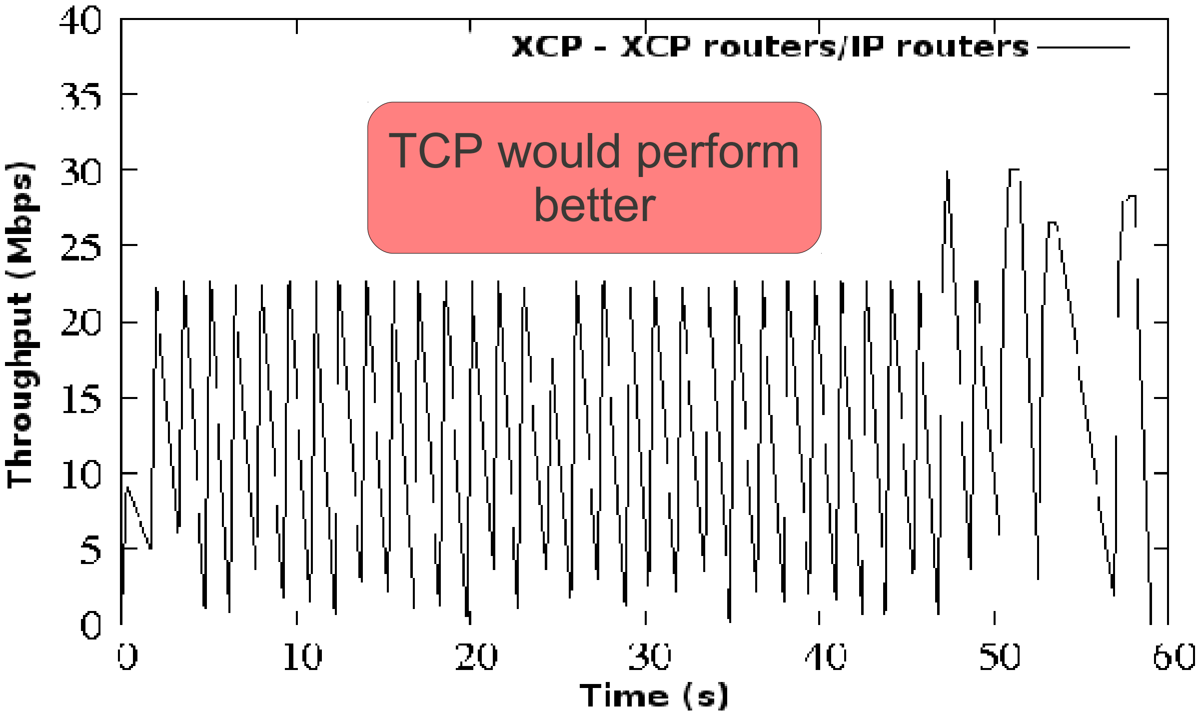

- If the bottleneck is not ERN router => wrong view of the network capacity

Interoperability of ERN protocols

- When non-ERN flows are present
  - Approach O/N
    - No room for foreign protocols
  - Approach (O-(I+Q))/N
    - Too conservative
    - I is originated by any flow
    - ERN flows cannot compete against foreign protocols

Ex. Conservative Approach

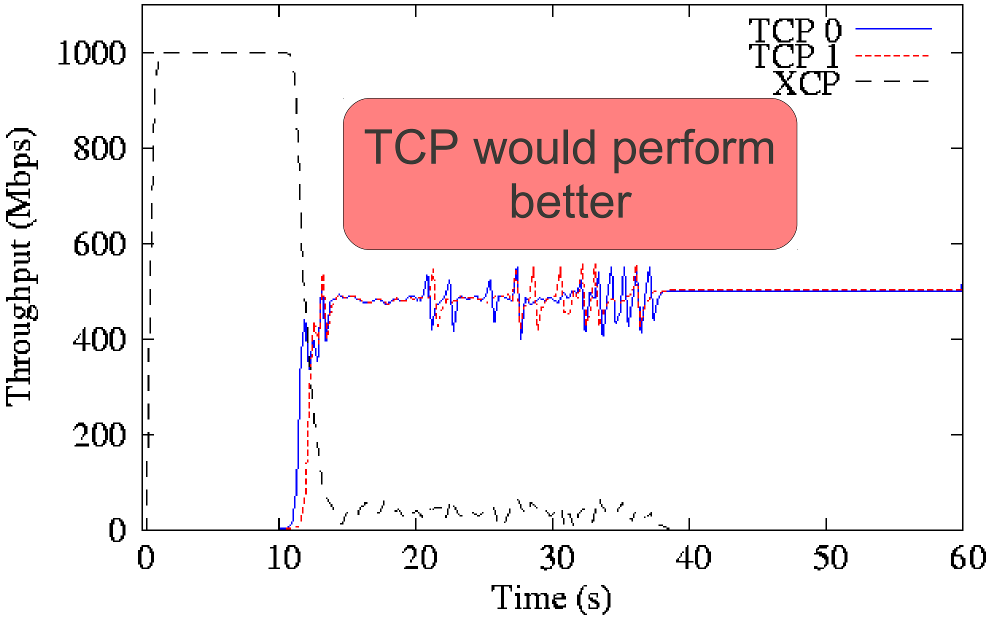

Strategies to deploy ERN protocols in real networks

- In satellite-based networks, [Kapoor05] proposes to isolate XCP with splitting proxies
  - Splitting proxies are unsuitable
- TCP-friendliness mechanisms (to be implemented in routers)
  - Probabilistically estimate the number of flows and probabilistically drop TCP packets [Lopez07]
    - Accuracy of the estimator
    - Problem of reactiveness. Best case: 1 period of estimation + 1 period of execution + time to affect sender behavior = $2*t_{est} + RTT$
  - Additionally, estimate the aggressiveness of TCP [Chia08]
    - Problem of reactiveness
    - Lowest aggressiveness != estimated aggressiveness

Strategies to deploy ERN protocols in real networks

- At the end hosts Detect non-ERN bottlenecks and switch to TCP if needed
  - RTT = 2*RTT0
    - Should RTT reach the threshold before congestion collapse?
  - Predicted ERN rate != receiving rate
    - This test might lead to false-positive results
  - No possibility to switch back to ERN protocol if the bottleneck moves to an ERN router

The IP-ERN architecture

Rationale

- Benefit from ERN protocols in heterogeneous networks
  - Enable a gradual deployment
  - Avoid heuristics with a behavior difficult to predict
- In ERN bottlenecks, ERN protocols are the best option
- In non-ERN bottlenecks, TCP can perform better than ERN

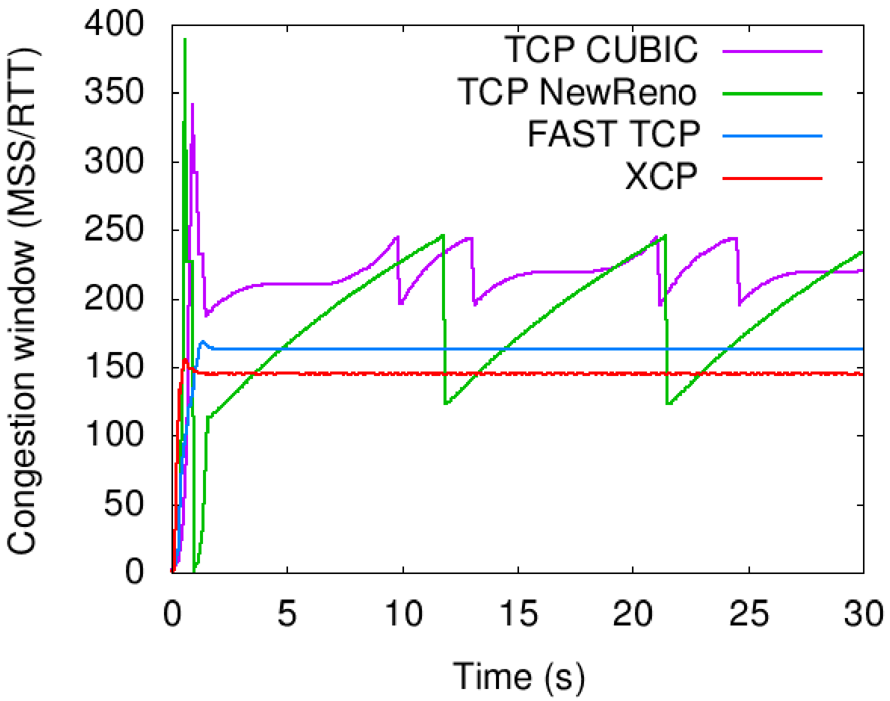

$$cwnd\_ = min(cwnd\_tcp\_, cwnd\_ern\_)$$

The IP-ERN architecture

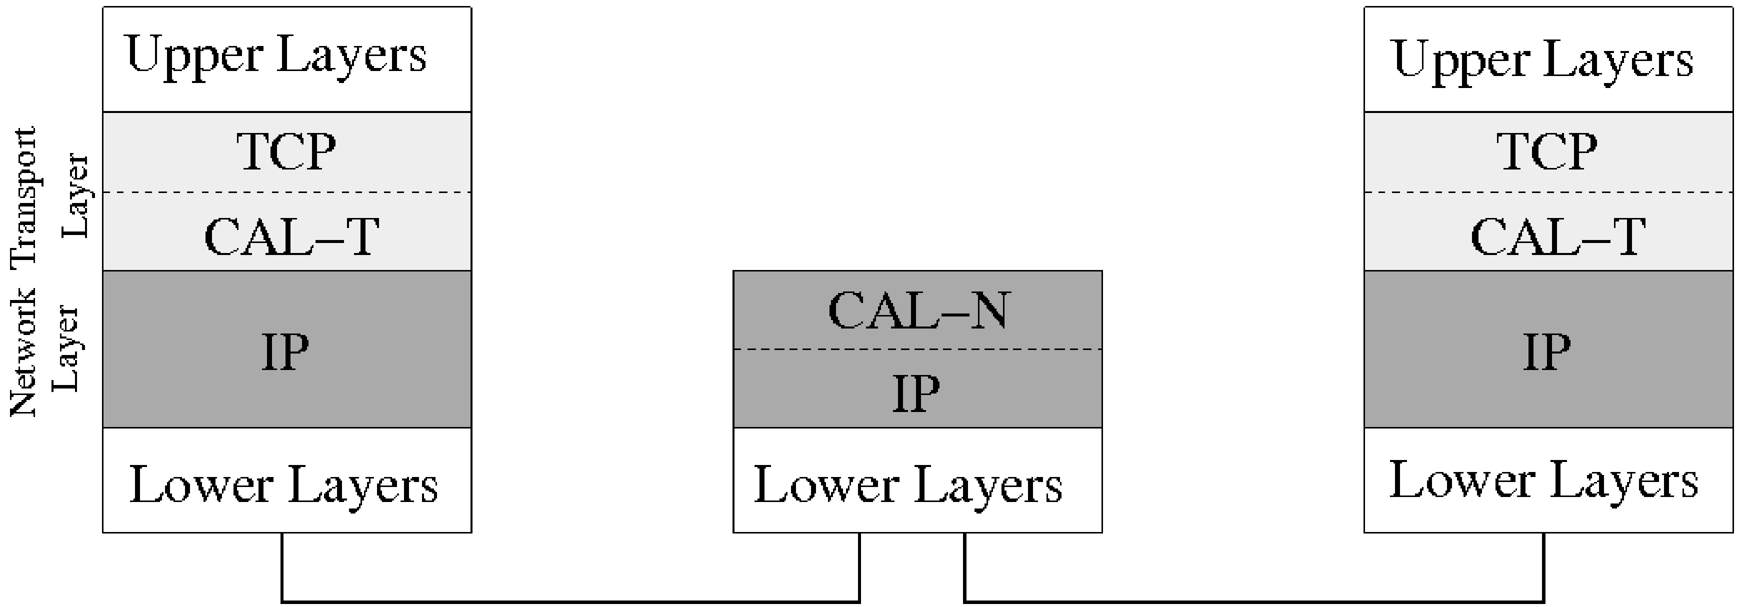

- Transport layer
  - TCP sub-layer
    - TCP protocol stack (including congestion control)
  - Congestion Awareness sub-layer at the Transport Level (CAL-T)
    - ERN protocol
    - Light algorithm – possible having two congestion control algorithms

The IP-ERN architecture

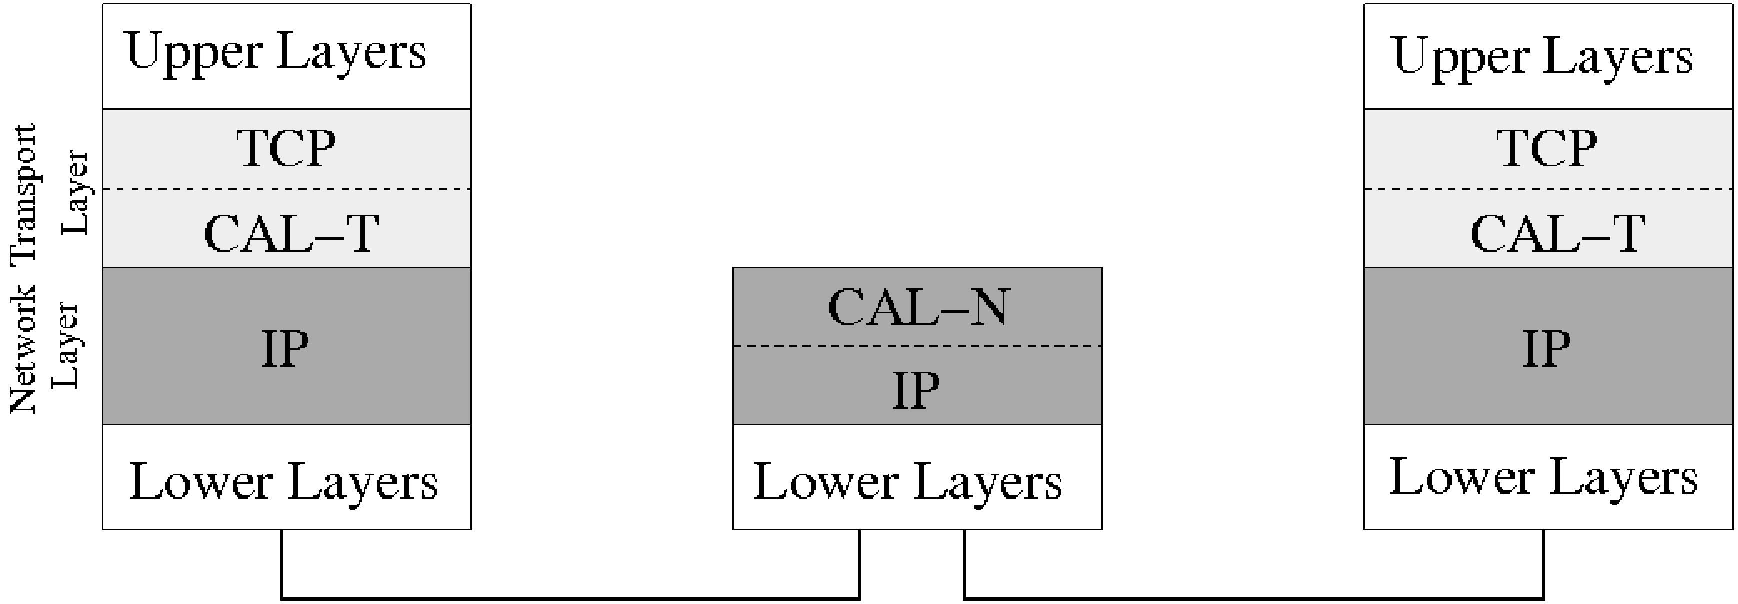

- Network Layer
  - IP sub-layer
  - Congestion Awareness sub-layer at the Network Level (CAL-N)
    - ERN protocol

Setting-up IP-ERN connections

- Sender
  - ERN header is inserted just after the IP-header
  - SYN and SYN-ACK packets indicate ERN availability in the TCP option field

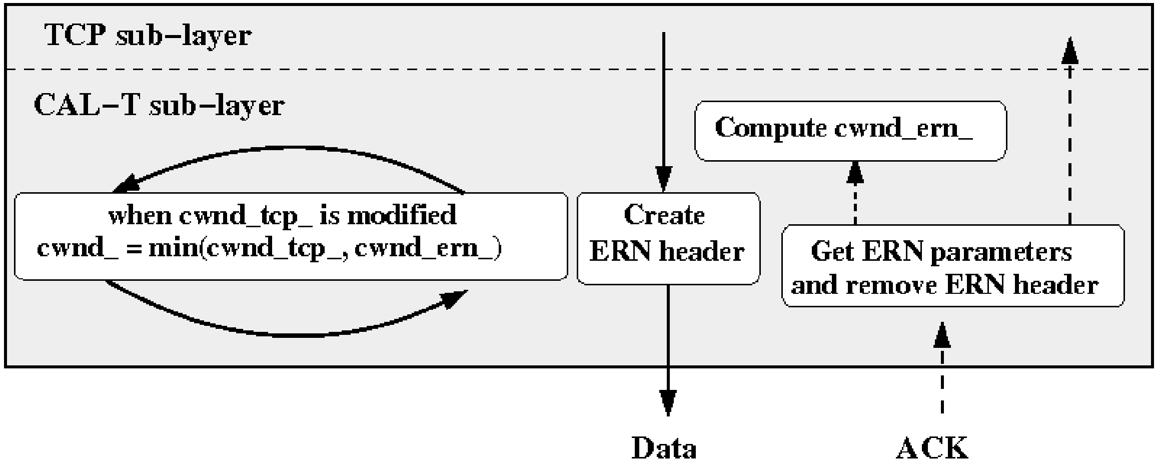

- Router
  - One small buffer for TCP and one another for ERN
  - Easy buffer occupancy computation to preserve fairness

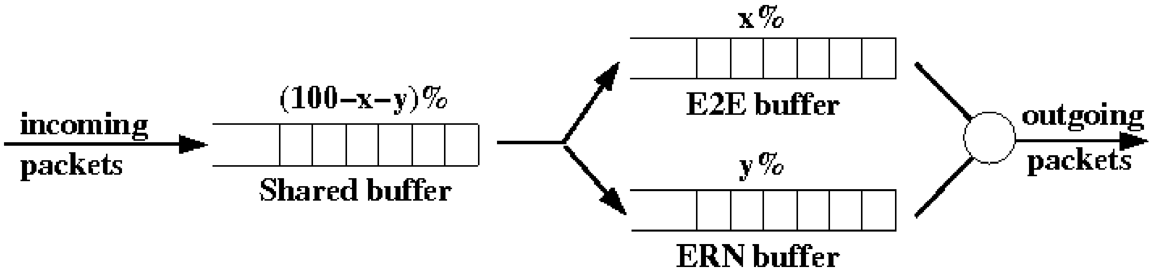

Simulation Results with *ns-2*
IP-ERN: CUBIC [Rhee05] + XCP [Katabi02]

Correctness

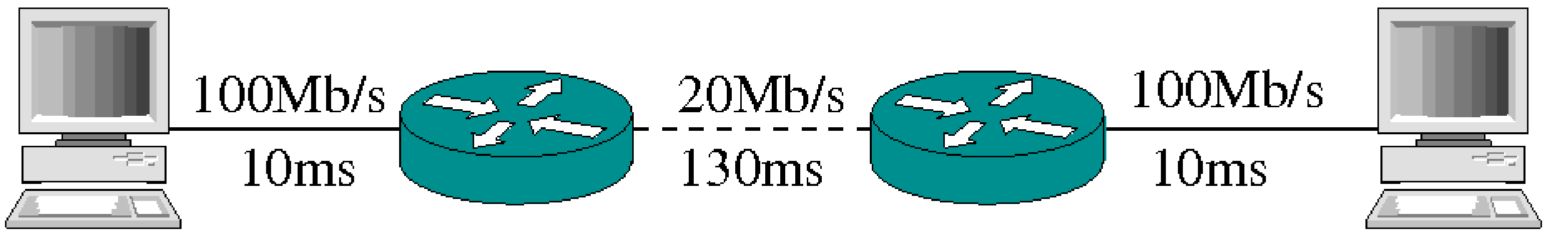

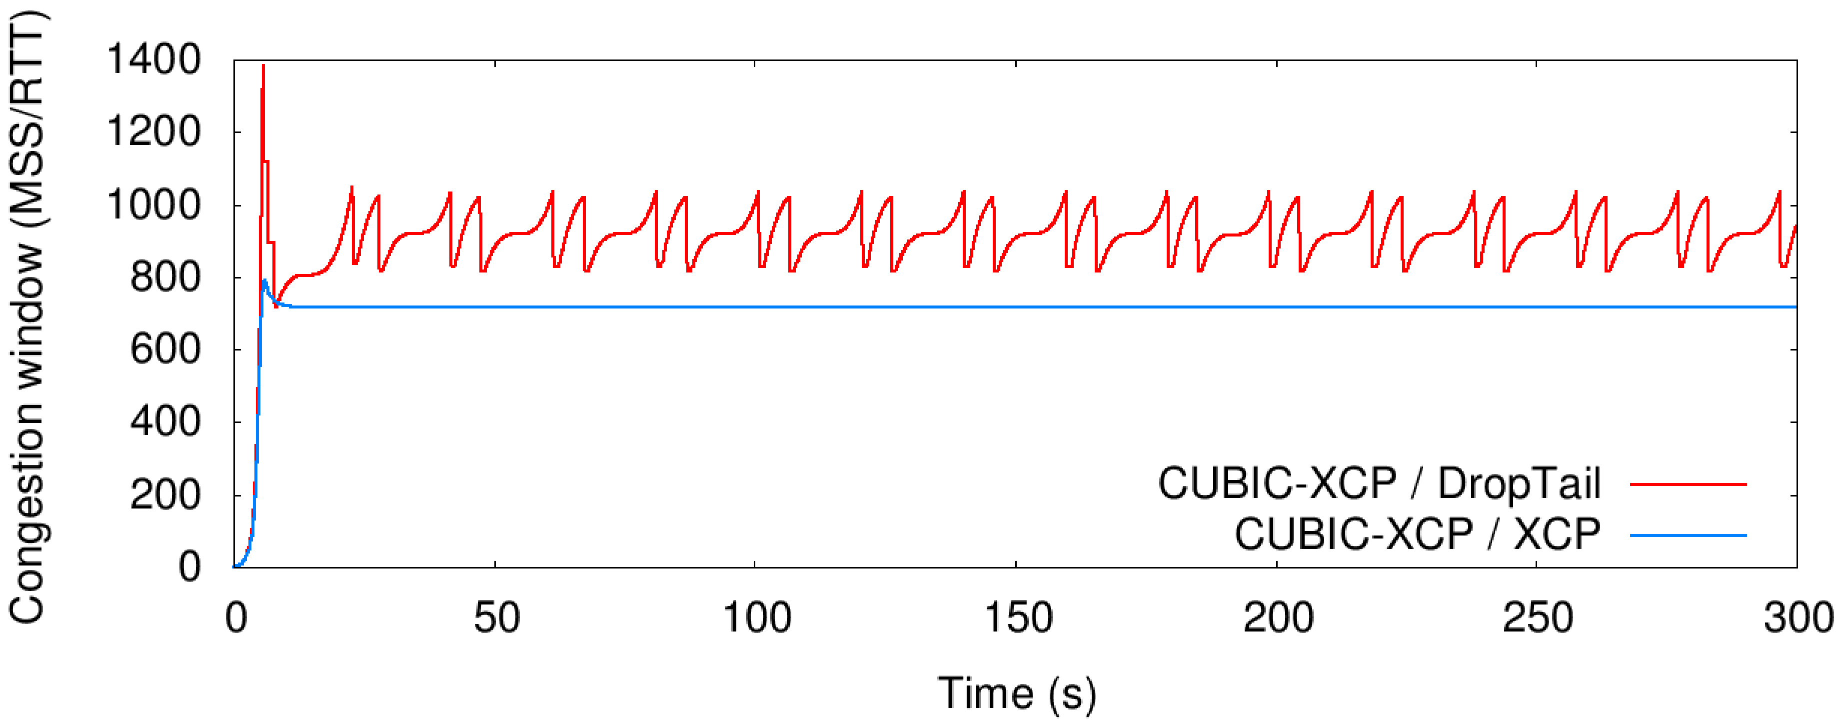

Bandwidth sharing (XCP routers)

Only CUBIC-XCP flows

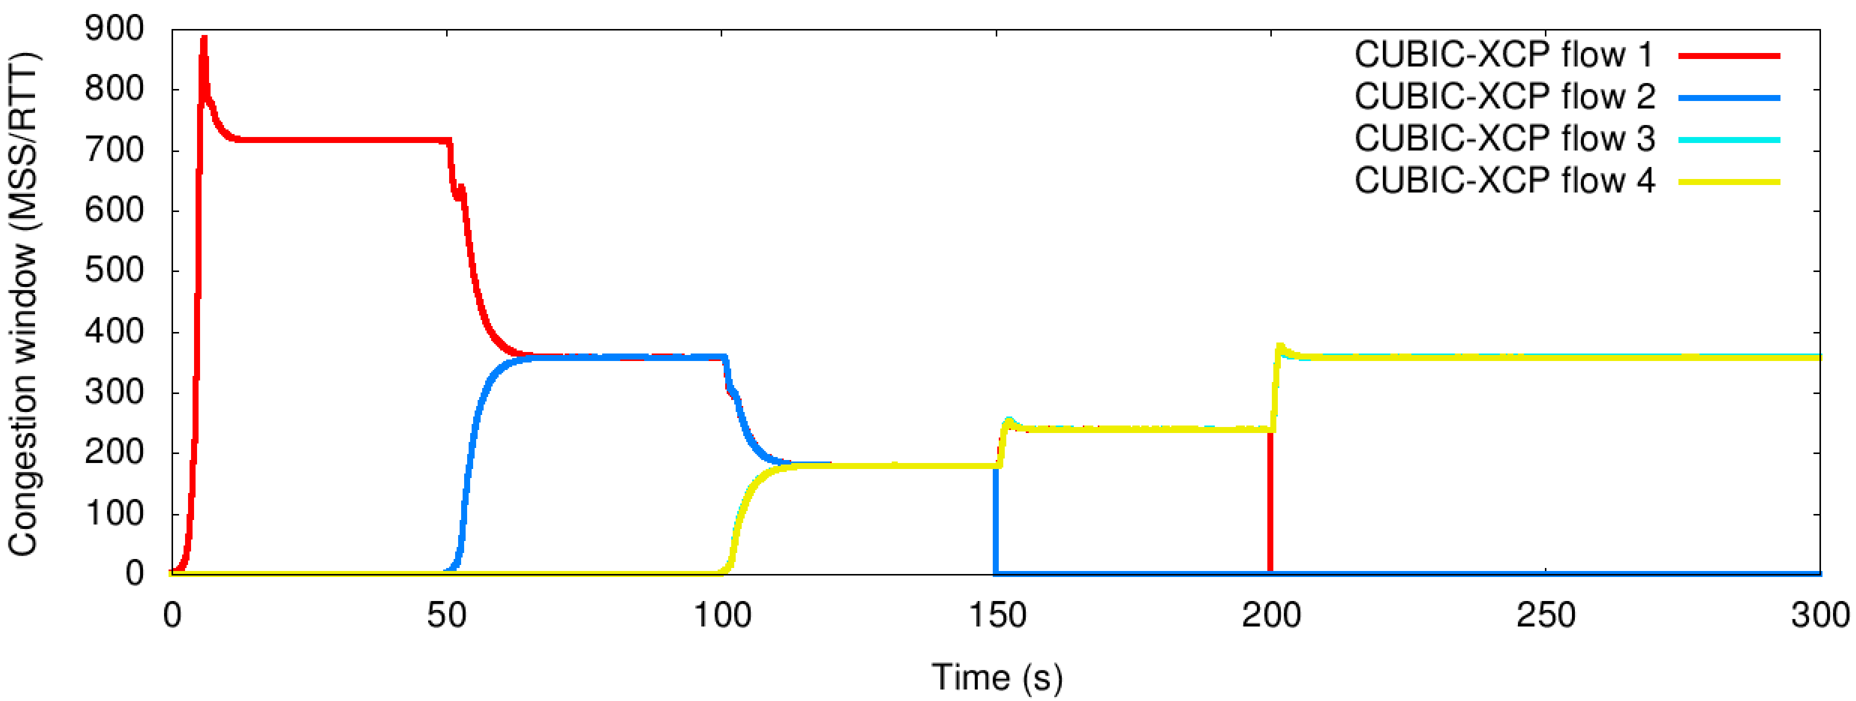

Bandwidth sharing (XCP routers)

One CUBIC vs two CUBIC-XCP flows

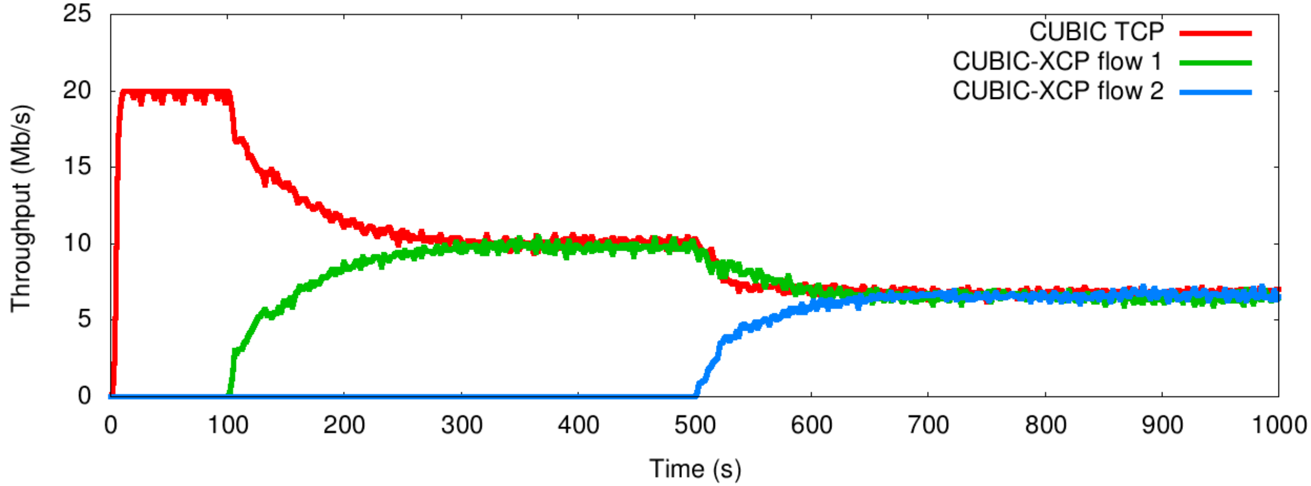

One CUBIC-XCP vs two CUBIC flows

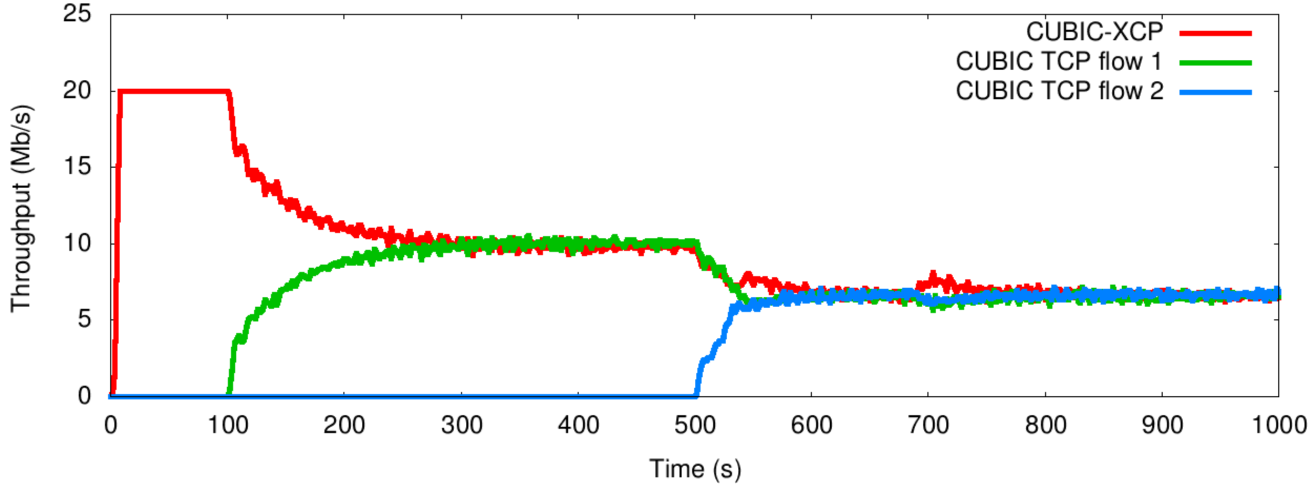

Reactiveness

- ERN → TCP
  - 1 RTT (Time to detect losses)
- TCP → ERN
  - 1 RTT (Time to get the feedback)

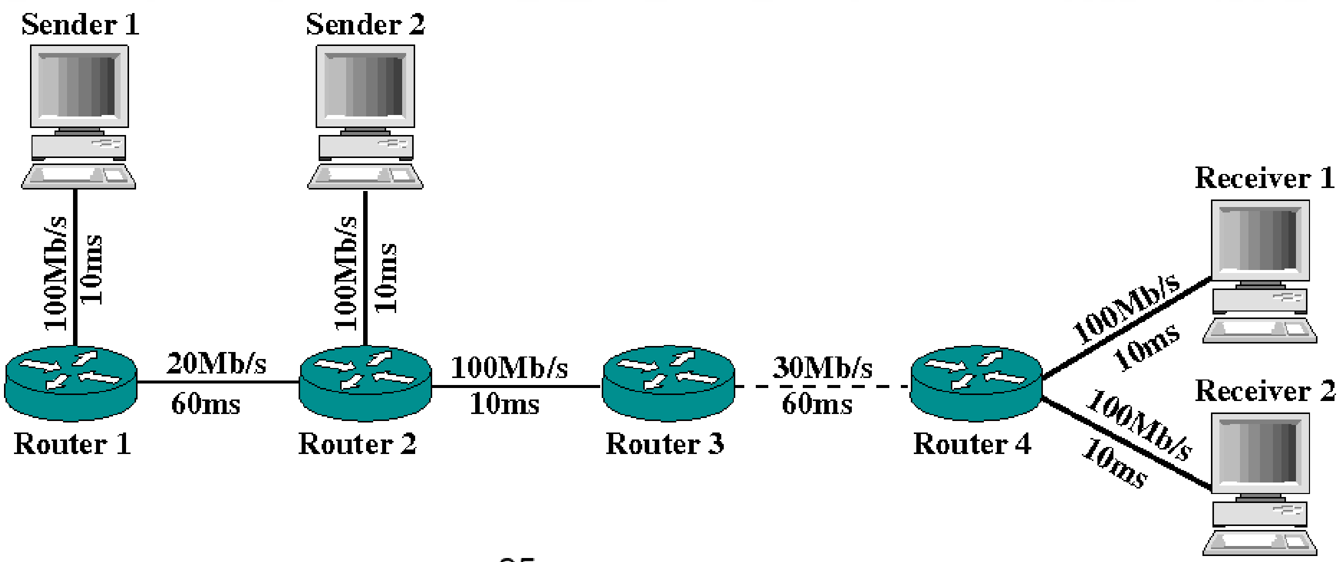

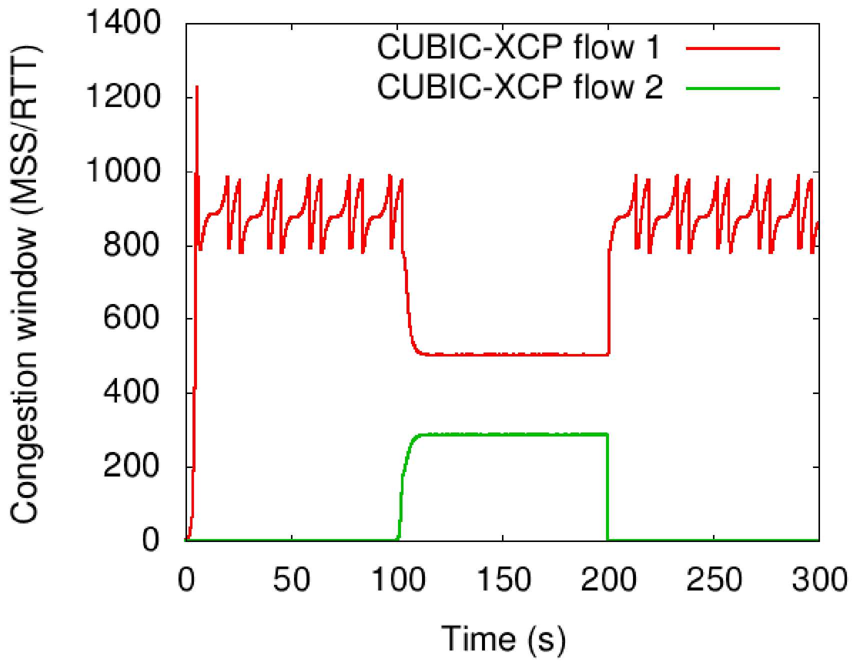

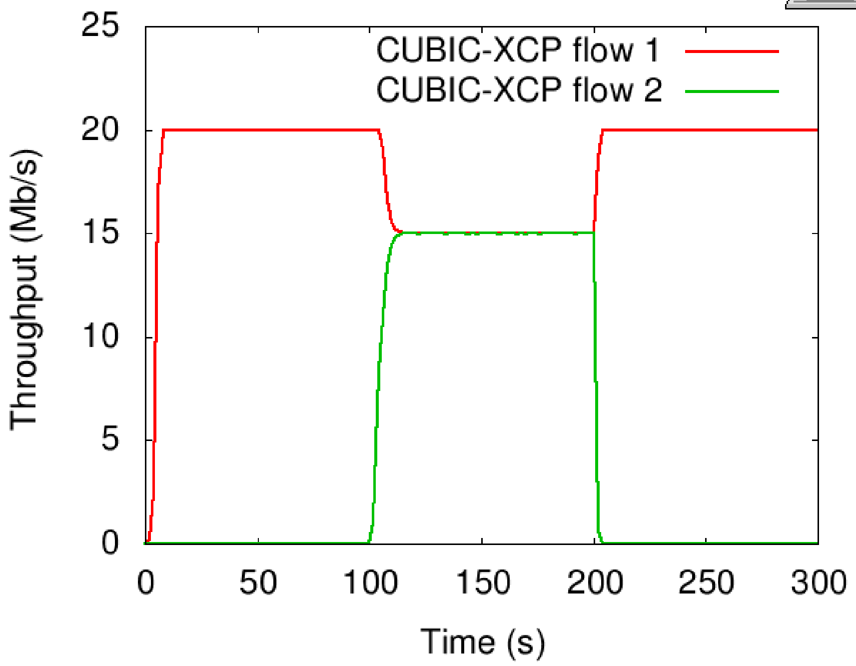

Gradual Deployment Scenarios

- Bottleneck is ERN-capable and only ERN flows are present
  - Senders will fully benefit from ERN's capabilities
- Bottleneck is not ERN-capable
  - Senders will not benefit at all from ERN's capabilities, but they completely benefit from TCP
- Bottleneck is ERN-capable and ERN flows share the resources with non-ERN flows
  - Fairness between ERN flows can be higher than fairness between non-ERN flows
    - Limited non-ERN flows

IP-ERN's benefits from now on

Virtual Private Networks

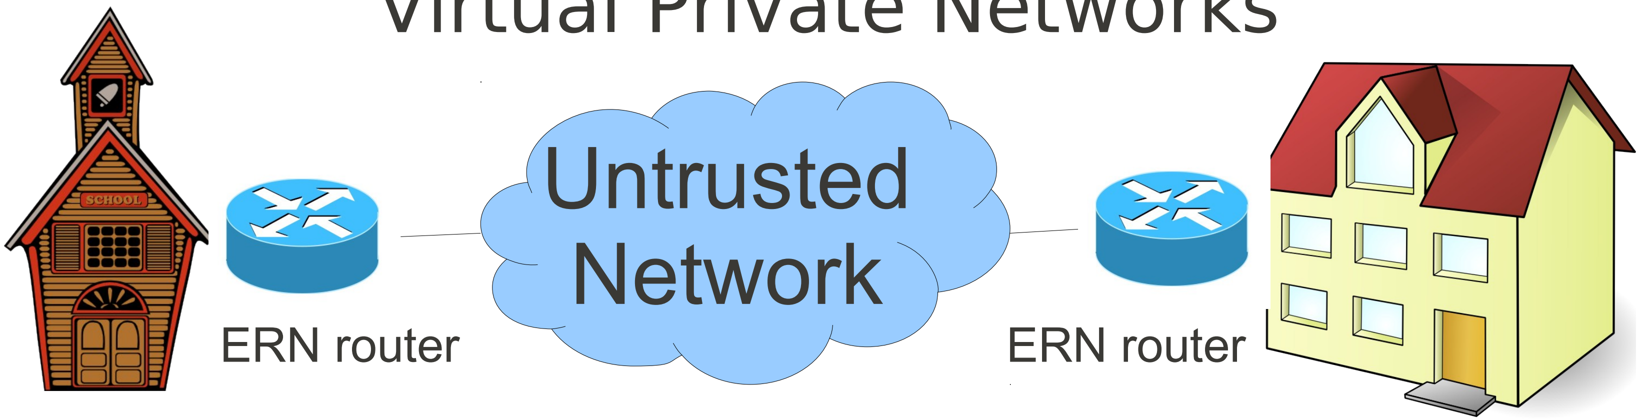

IP-based Satellite Networks

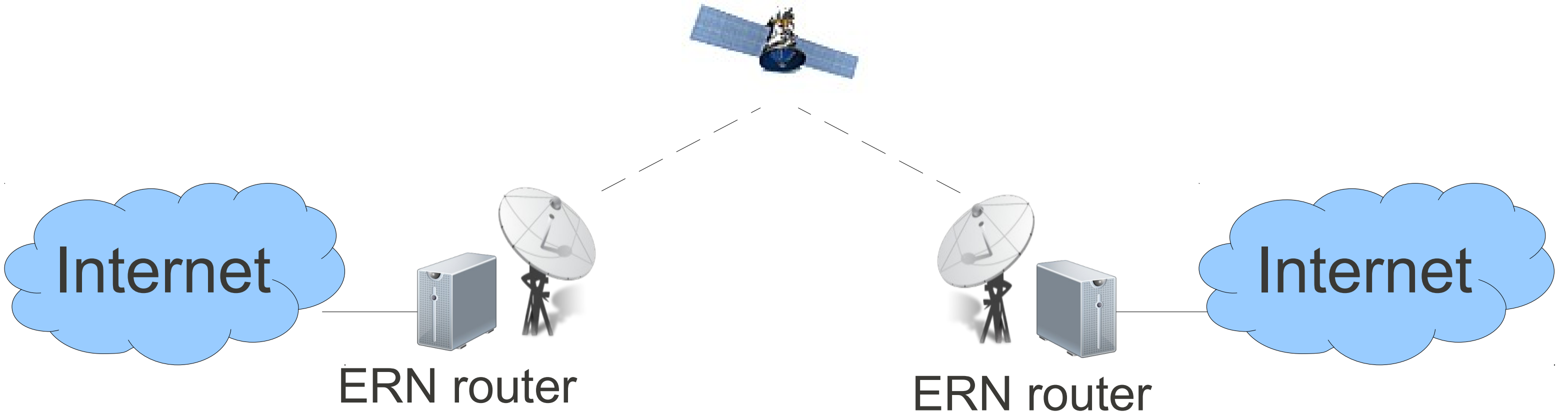

Conclusion

- Security considerations
  - Cohabitation between the ERN header and IPsec
  - ERN header protection
- Build a prototype
  - Test on short/long paths
  - Interplay between IP-ERN and transport layer middle boxes
- Propose an Internet Draft
  - We need your comments :)

References

[Chia08] Chia-Hui Tai, Jiang Zhu, and N. Dukkipati, “Making large scale deployment of rcp practical for real networks,” in IEEE Infocom, April 2008.

[Kapoor05] Aman Kapoor, Aaron Falk, Ted Faber, Yuri Pryadkin, "Achieving Faster Access to Satellite Link Bandwidth", 8th IEEE Global Internet Symposium, Miami, FL, March 2005.

[Katabi02] D. Katabi, M. Handley, and C. Rohrs, “Congestion Control for High Bandwidth-Delay Product Networks,” in ACM Sigcomm, 2002.

[Lopez07] Dino Lopez Pacheco, Congduc Pham, and Laurent Lefèvre, “Fairness issues when transferring large volume of data on high speed networks with router-assisted transport protocols,” in High Speed Networks Workshop, in conjunction with IEEE Infocom, May 2007.

[Rhee05] Injong Rhee and Lisong Xu, “CUBIC: A New TCP-Friendly High-Speed TCP Variant,” in PFLDNet, February 2005.

Thank you

Questions

ERN protocols in a nutshell

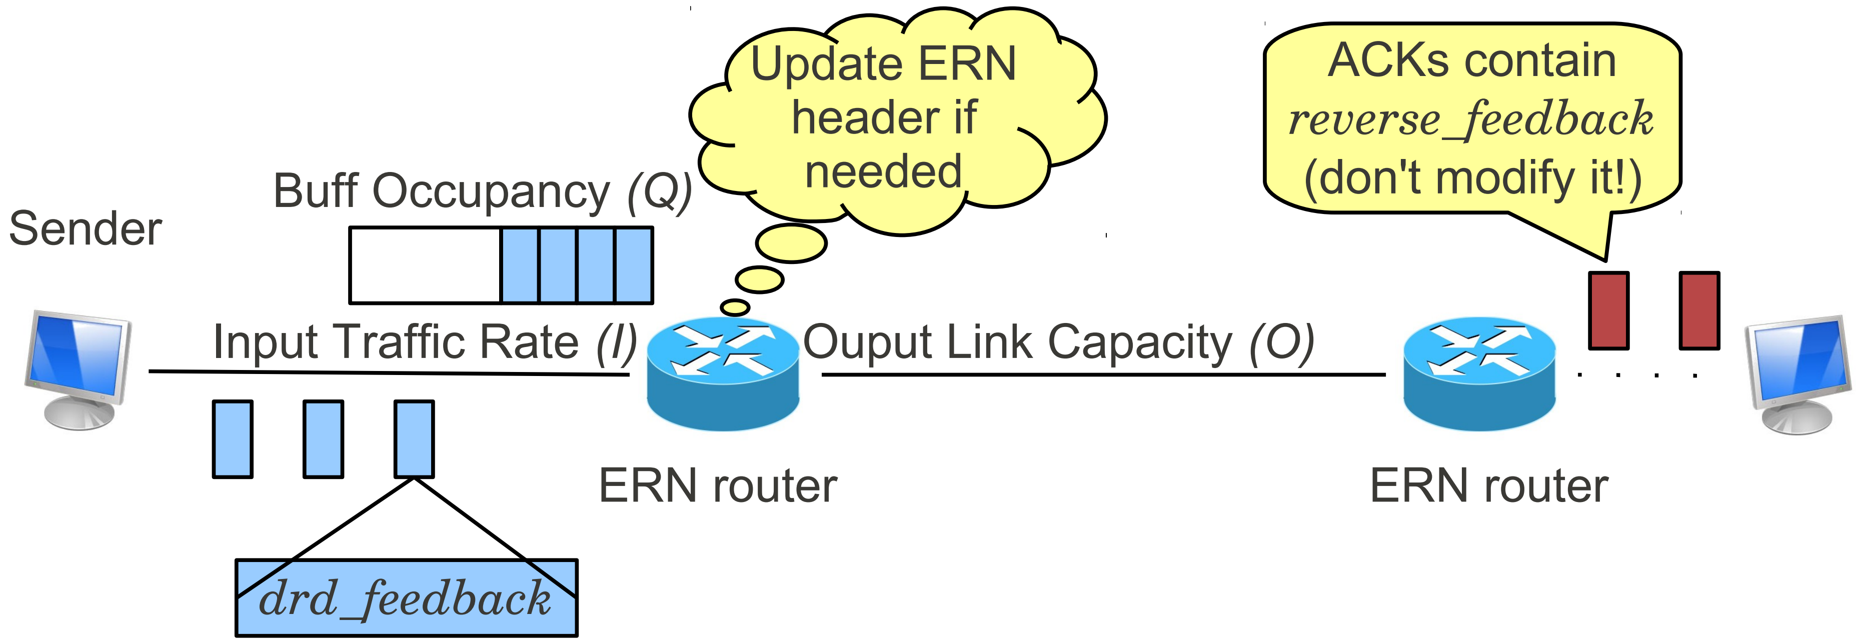

- ERN header's parameters (cwnd, RTT) allow routers to compute the number of flows N
- Resources are fairly shared between sources
  - Approach 1. O / N
  - Approach 2. $(O - (I + Q)) / N$