Enabling Large Data Transfers on Dynamic, Very High-Speed Network Infrastructures

Dino M. Lopez P. — INRIA RESO/LIP — ENS at Lyon, France

Congduc PHAM — LIUPPA — University of Pau, France

ICN 2006

Index

- Introduction
- The XCP protocol
- XCP-r: a new more robust XCP version
- Comparing XCP-r with others protocols
- Conclusion
- Futures works

Introduction

TCP & High Speed Networks

High Speed Networks (HSN) with the most up-to-date networking technologies are currently being deployed to support

- a large variety of new scientific applications.
- an increasing demand for high-speed Internet access and bandwidth-consuming applications.

The TCP congestion control algorithm has been successful in handling low to medium speed networks.

The challenge now is to make TCP operational and efficient on these new high-speed network infrastructures.

Dynamics High Speed Networks (DHSN)

Inside a High-Speed Network, the available best-effort bandwidth can vary over time for many reasons:

- non regulated flows (UDP).
- QoS Mechanisms.
- Resources reservation.

In these DHSNs, TCP congestion avoidance mechanism is not efficient and not responsive enough to recover from packet losses.

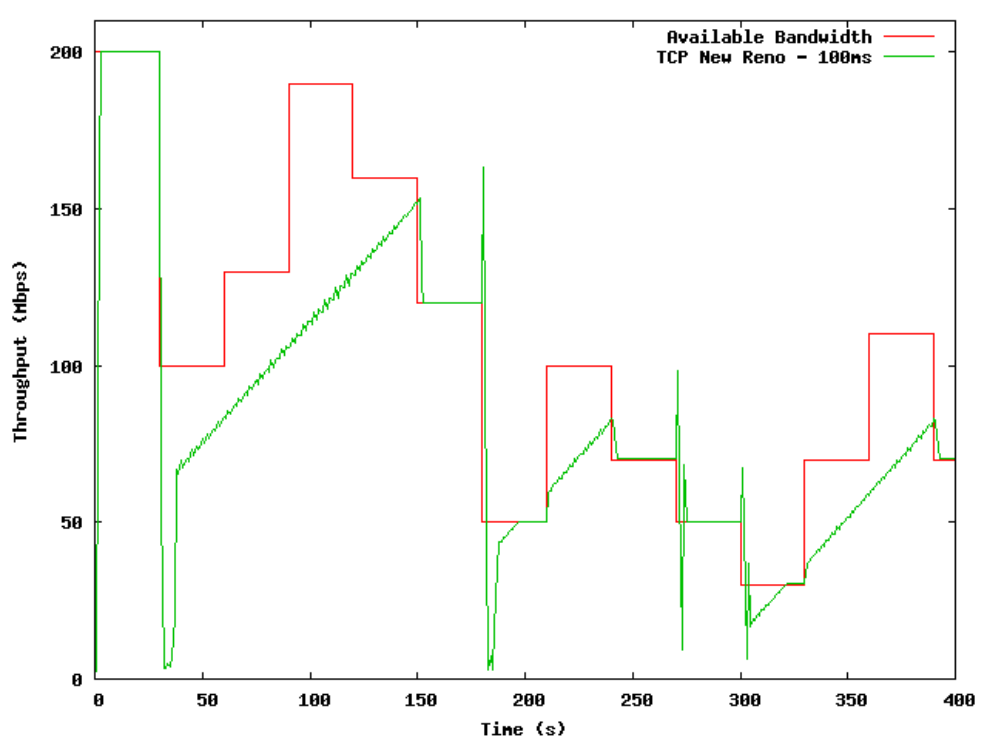

Optimizing TCP

In order to solve the problem of TCP in Dynamics High Speed Networks, some strategies have been proposed, for example:

- Tuning the TCP parameters for a high speed network (eg. web100). However, configuring TCP with optimal parameters for a Dynamic HSN is a difficult task.

- Creating parallel flows to increase the total throughput (eg. GridFTP). However, determining the number of parallel flows to get the best performance in DHSNs is almost impossible.

New TCP version for High Speed Networks and New Transport Protocols

At the moment, there are many TCP versions that take advantage of the resources in a High Speed Network, for example:

- FAST TCP
- TCP Westwood+
- SCTP, HSTCP
- etc.

Other propositions very different from TCP (based in the use of assisted routers), such as:

- CADPC/PTP
- XCP

The XCP Protocol

The XCP Protocol

- XCP is a router-assisted solution (XCP routers)
- Routers can compute the available bandwidth by monitoring the input traffic rate I_r, the output link capacity O_r, and the persistent queue size Q.

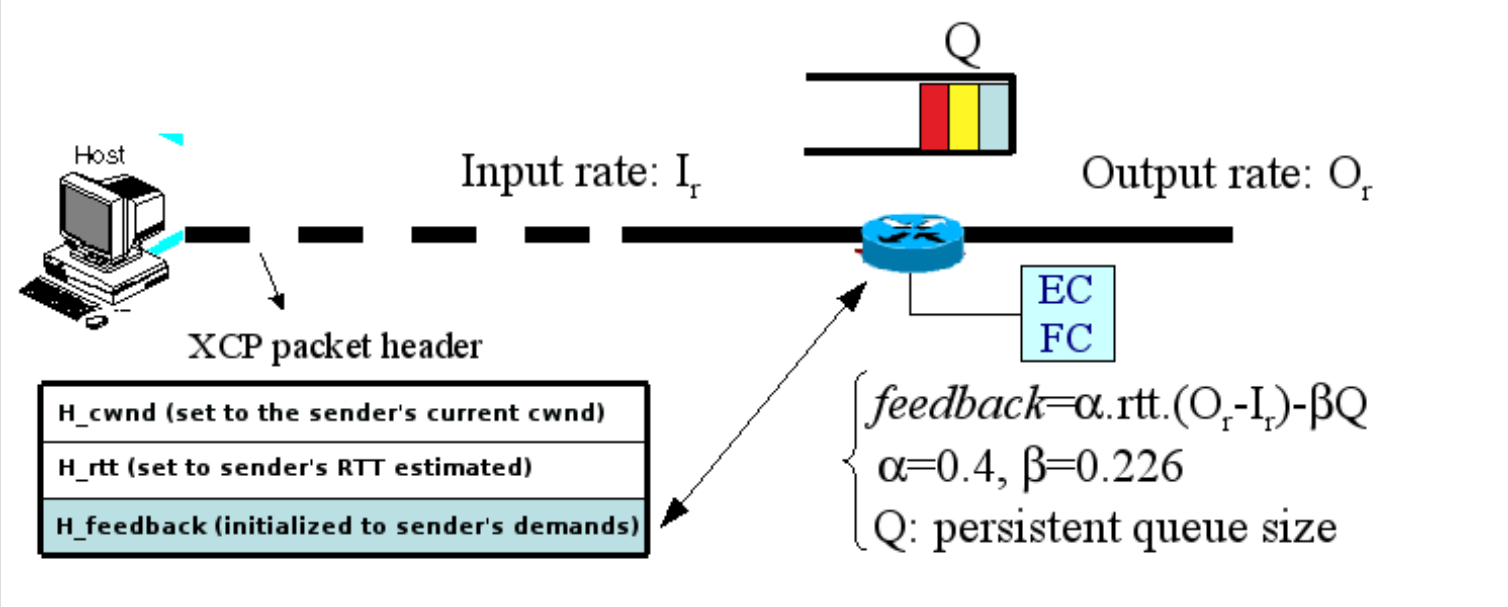

The XCP Protocol (2)

In this way:

- The feedback is computed by the XCP routers and sent back to the sender in the ACK packets.
- The sender receives the feedback and adds this value to his congestion window size.
- The sender does not make any assumption about the state of the network. This information is provided by the XCP routers.

XCP could be efficient in a Dynamic High Speed Network.

XCP in a Dynamic High Speed Network

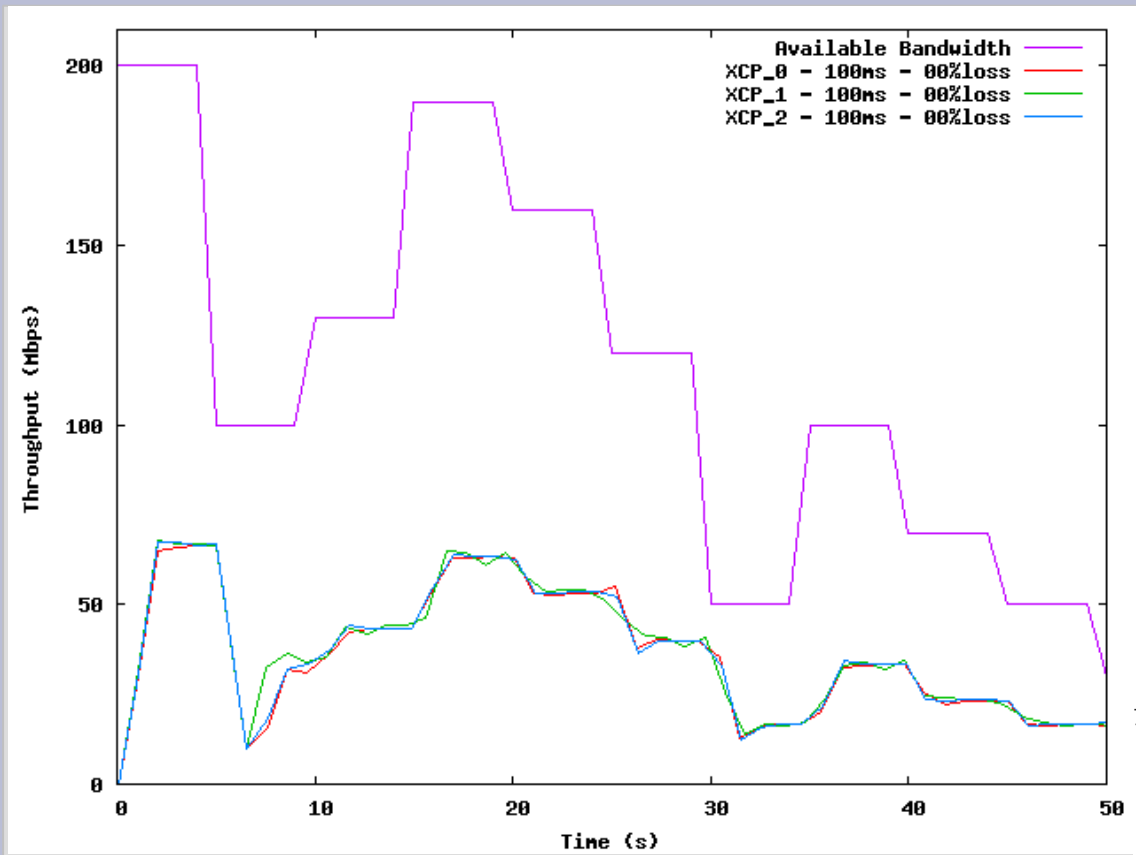

Good fairness and stability even in unfriendly environments

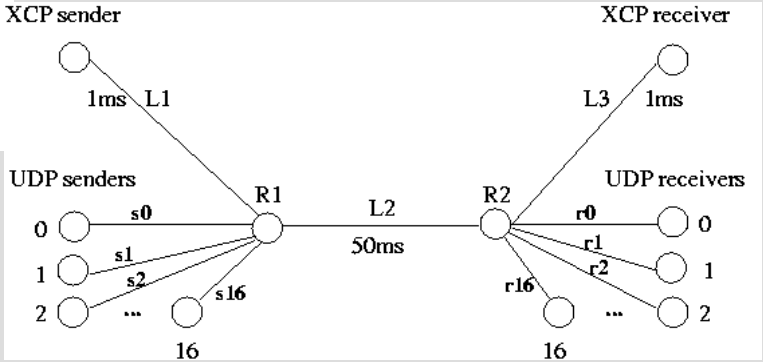

Loss of ACK packets in XCP networks

The loss of packets is common in networks. In case of ACK losses (feedbacks) inside of an XCP network:

- if the bandwidth does not vary over time, there is no problem because:
  - when the throughput is small, the feedback will be large enough to get the maximum capacity.
  - when the throughput is optimal, the loss of very few feedbacks will not have any effect.

- if the bandwidth varies over time, and the ACKs are lost when the bandwidth decreases, these losses will produce a wrong calculus in the sender's congestion window size, increasing the probability of data packet losses.

XCP (12% ACK loss rate – 1Gb link – 100ms RTT)

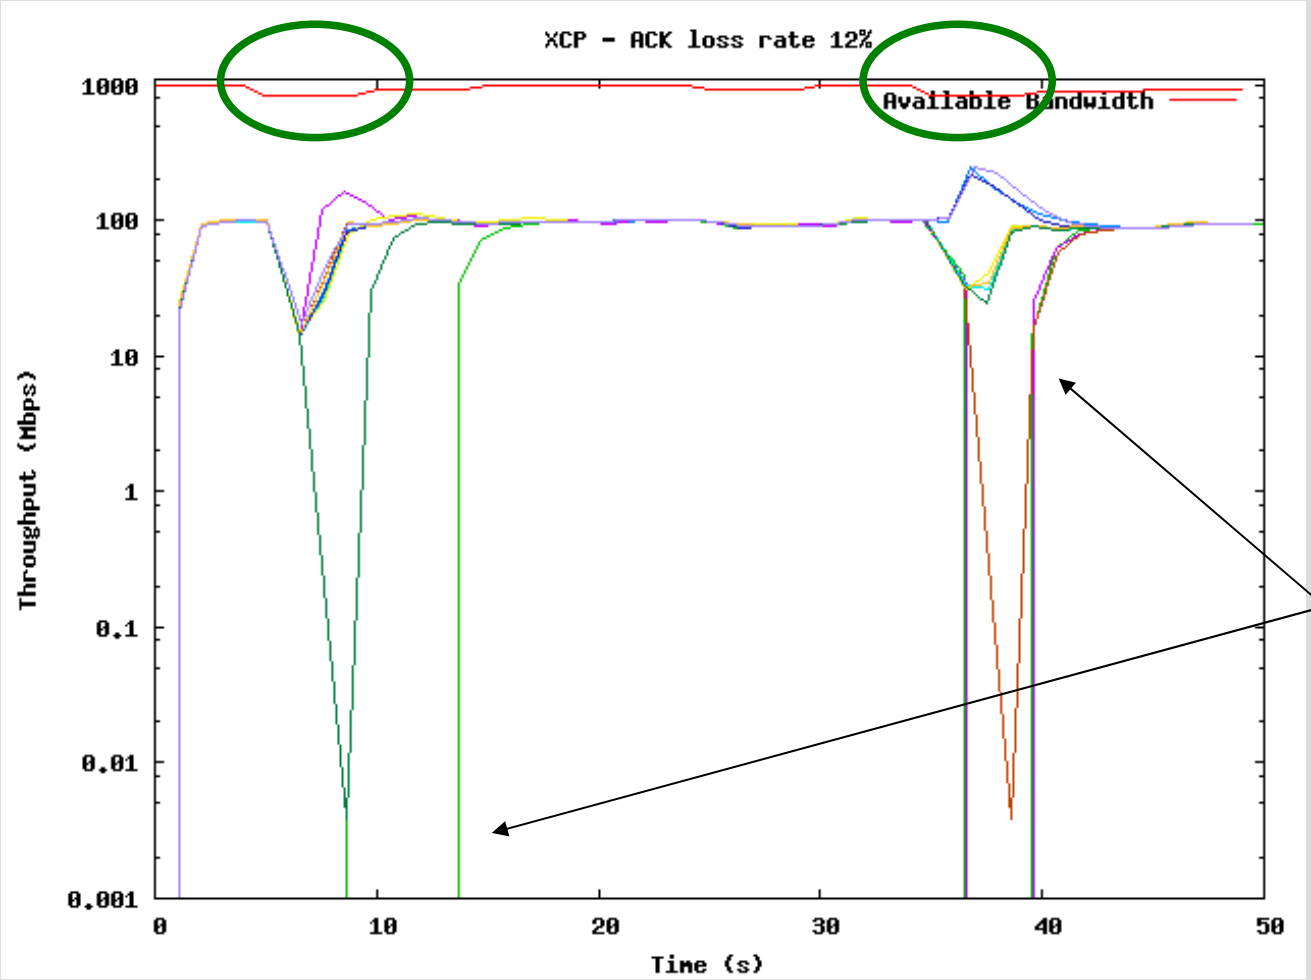

Small changes in the bandwidth can introduce many problems in XCP

Timeouts problems, Unfairness and difficulty in restarting the connection

XCP-r

XCP-r, a more robust XCP version

Since the ACK losses produce a wrong calculus of the congestion window size in the sender, XCP-r computes this value in the receiver side and send it to the sender inside the ACK packets.

Some technical aspects in XCP-r

- The receiver will create an equivalent variable to *cwnd (cwnd')* for each connection established.
- After a congestion problem, if the sender modifies its congestion window size, it will notify the change and the receiver will update the cwnd' variable .

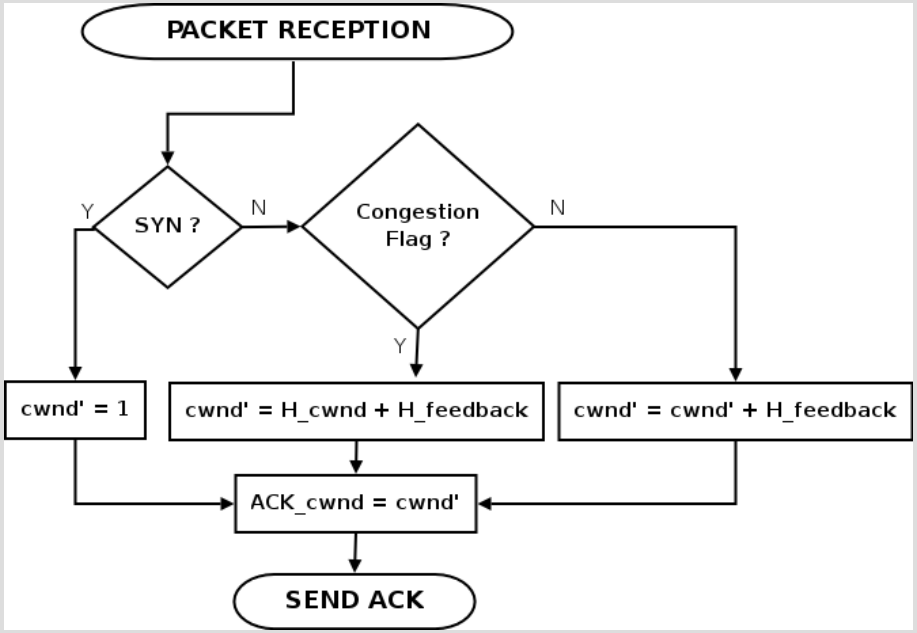

New algorithm in the receiver side

XCP-r (12% ACK loss rate – 1Gb link – 100ms RTT)

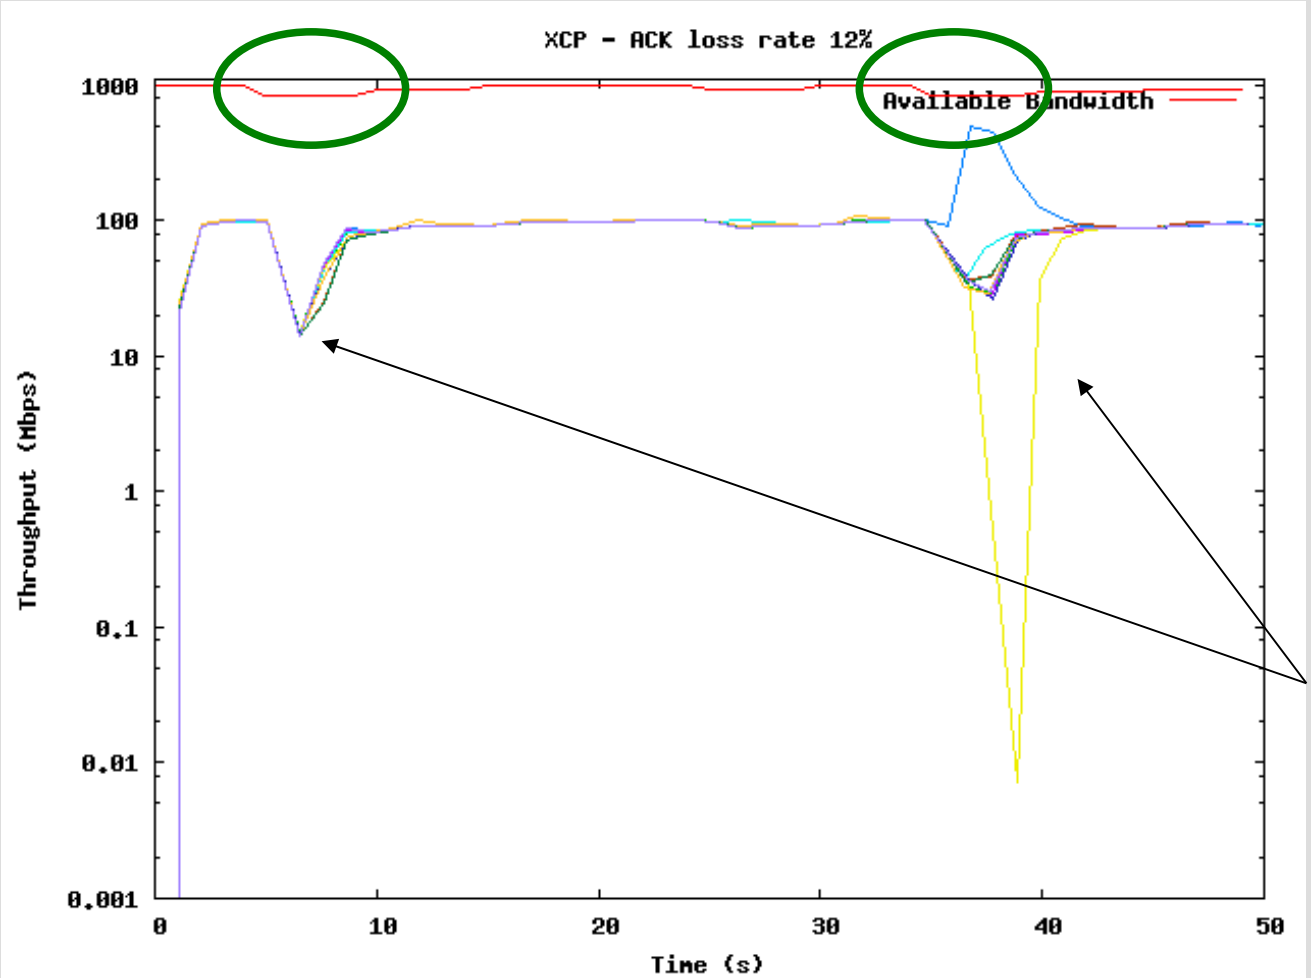

XCP-r improves the performance!

Fast recovery after the timeouts problems and better fairness level

Comparing XCP-r to other protocols

XCP & XCP-r in a very High Speed Network (12% ACK loss rate – 1Gb link – 200ms RTT)

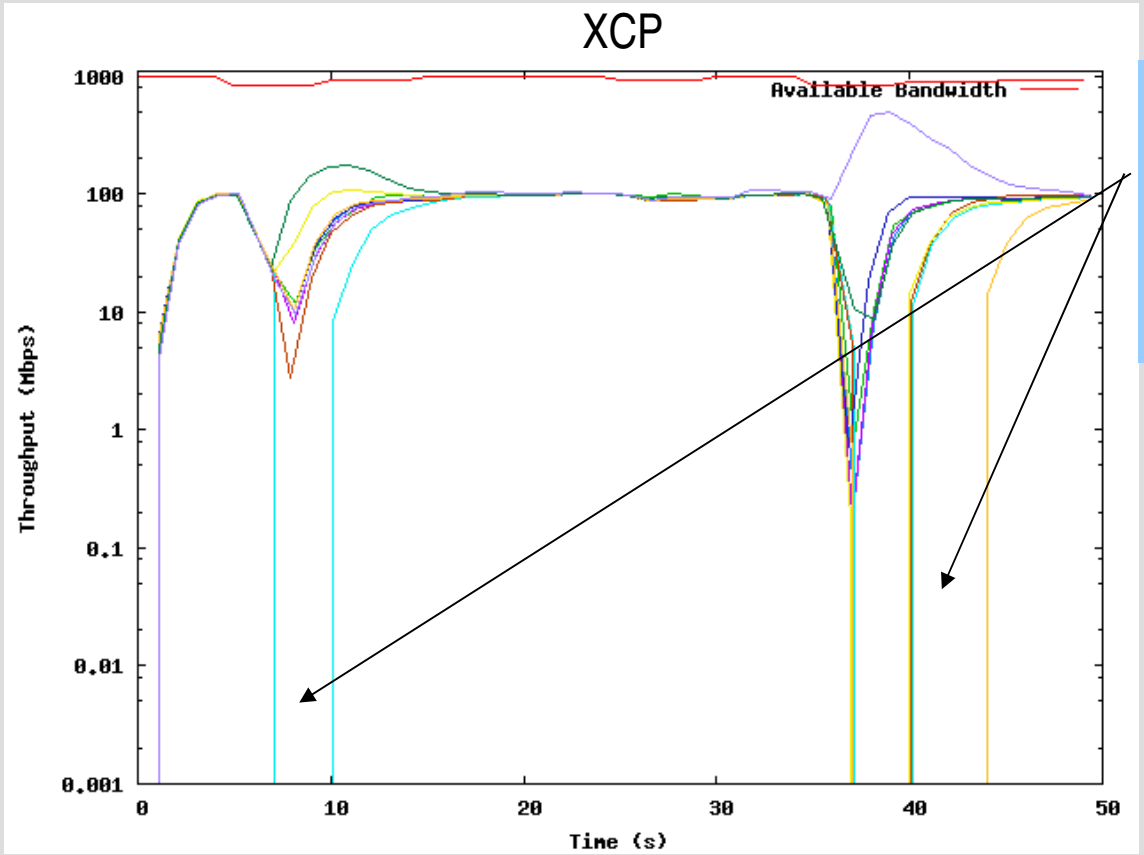

Timeouts introduced by the changes in the bottleneck bandwidth

Many problems to restore the connections...

XCP & XCP-r in a very High Speed Network (12% ACK loss rate – 1Gb link – 200ms RTT)

XCP-r

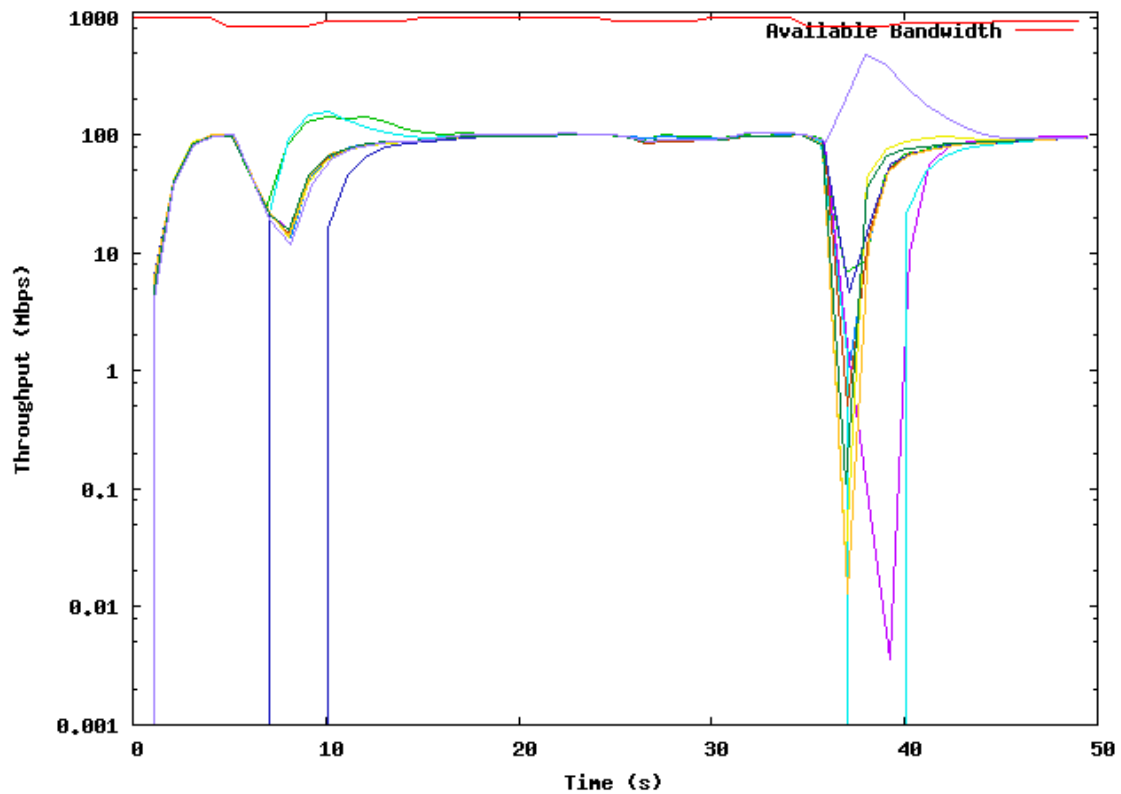

the connections are restored really quickly!

XCP

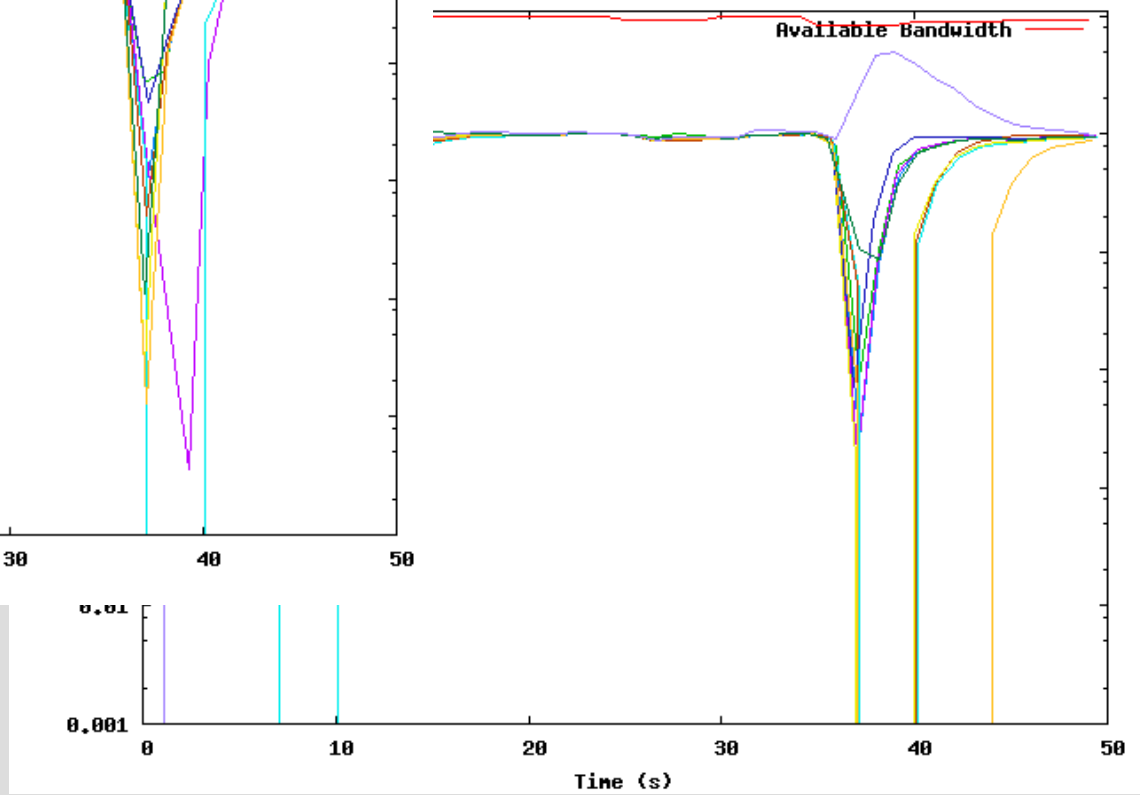

XCP-r does not avoid all the timeouts, however

XCP-r vs XCP (Mbytes transferred per flow and Fairness)

XCP-r gets an excellent fairness level even under the most difficult conditions

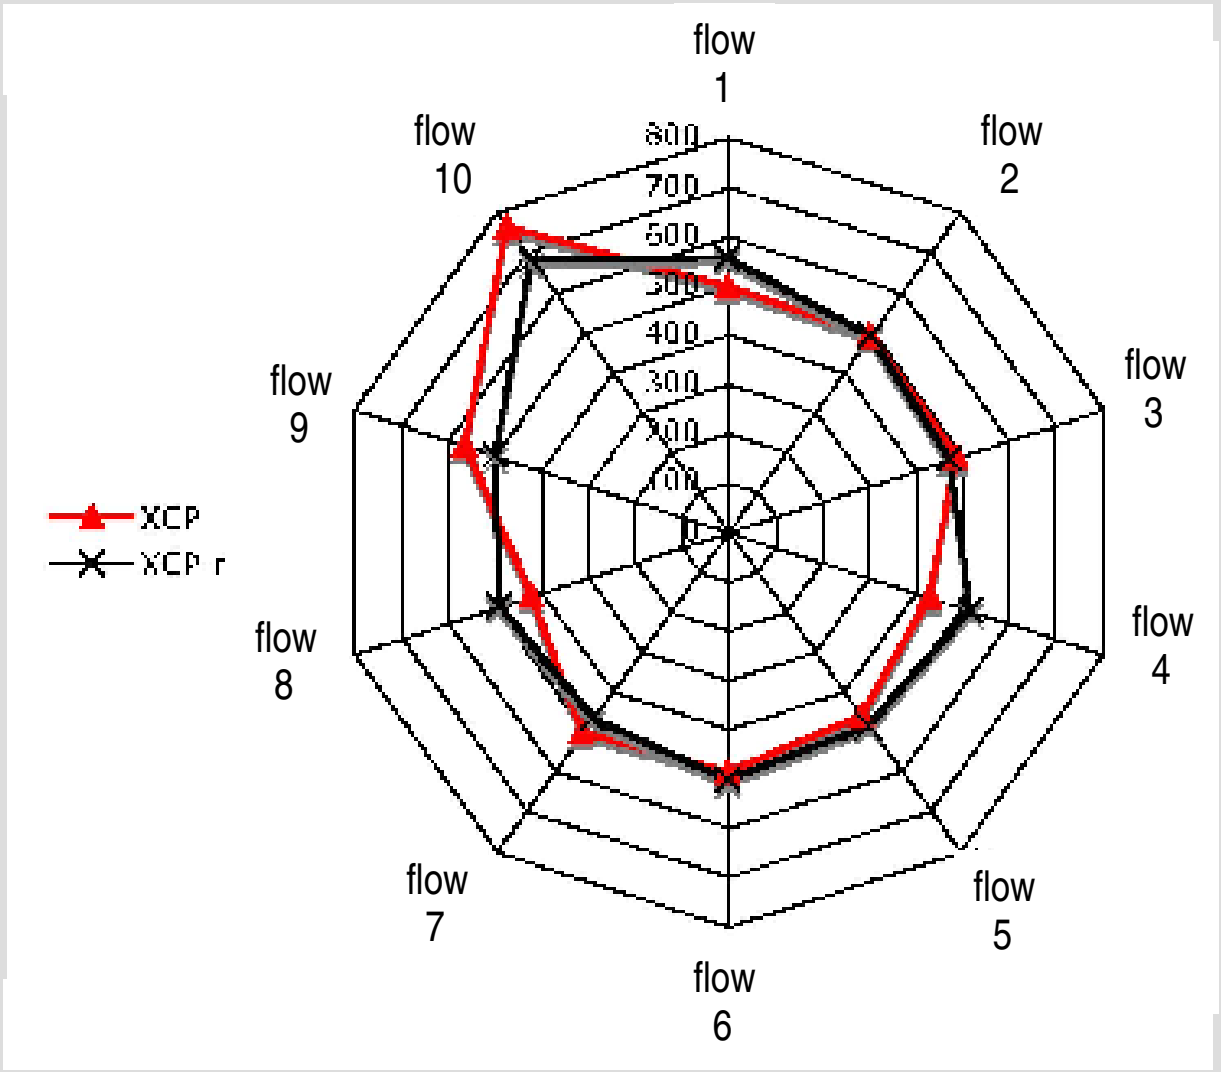

Standard Deviation

XCP : 91.9

XCP-r : 59.24

XCP-r vs XCP, TCP and High Speed TCP (Total Mbytes transferred)

the performance of TCP is really low :'-(

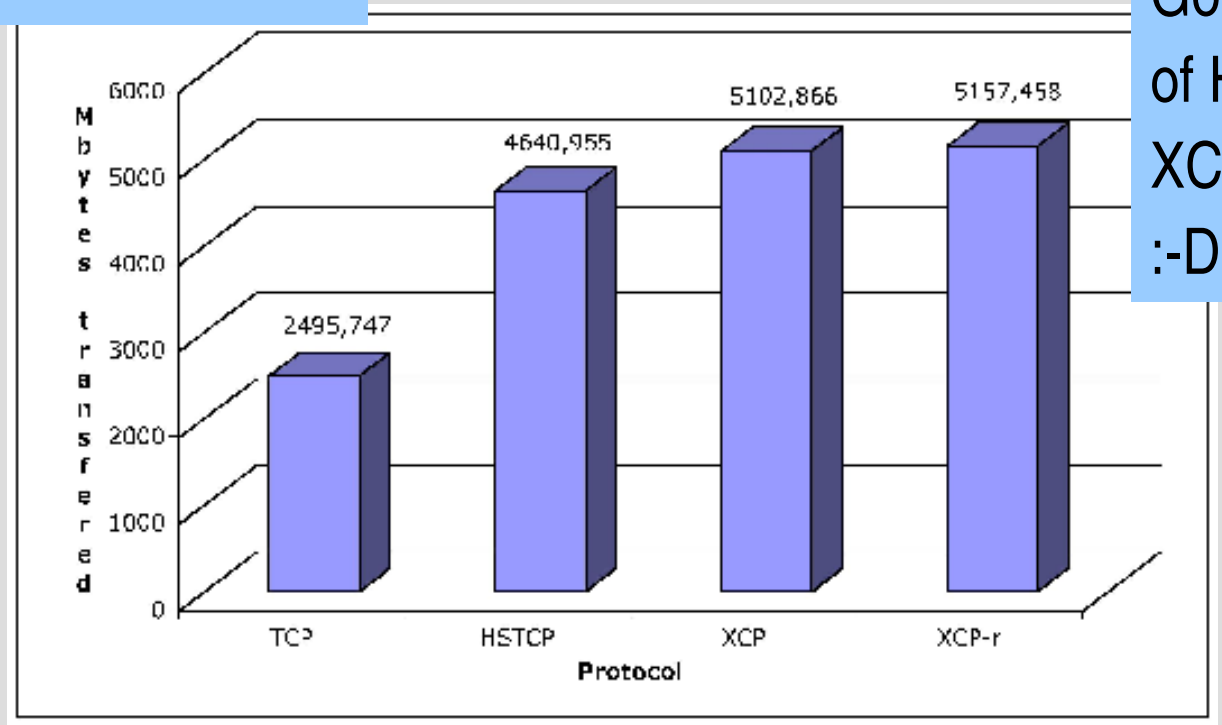

Good Performance of High Speed TCP, XCP, and XCP-r :-D

XCP vs TCP and High Speed TCP (Mbytes transferred per flow and Fairness)

TCP and HSTCP are not really fair...

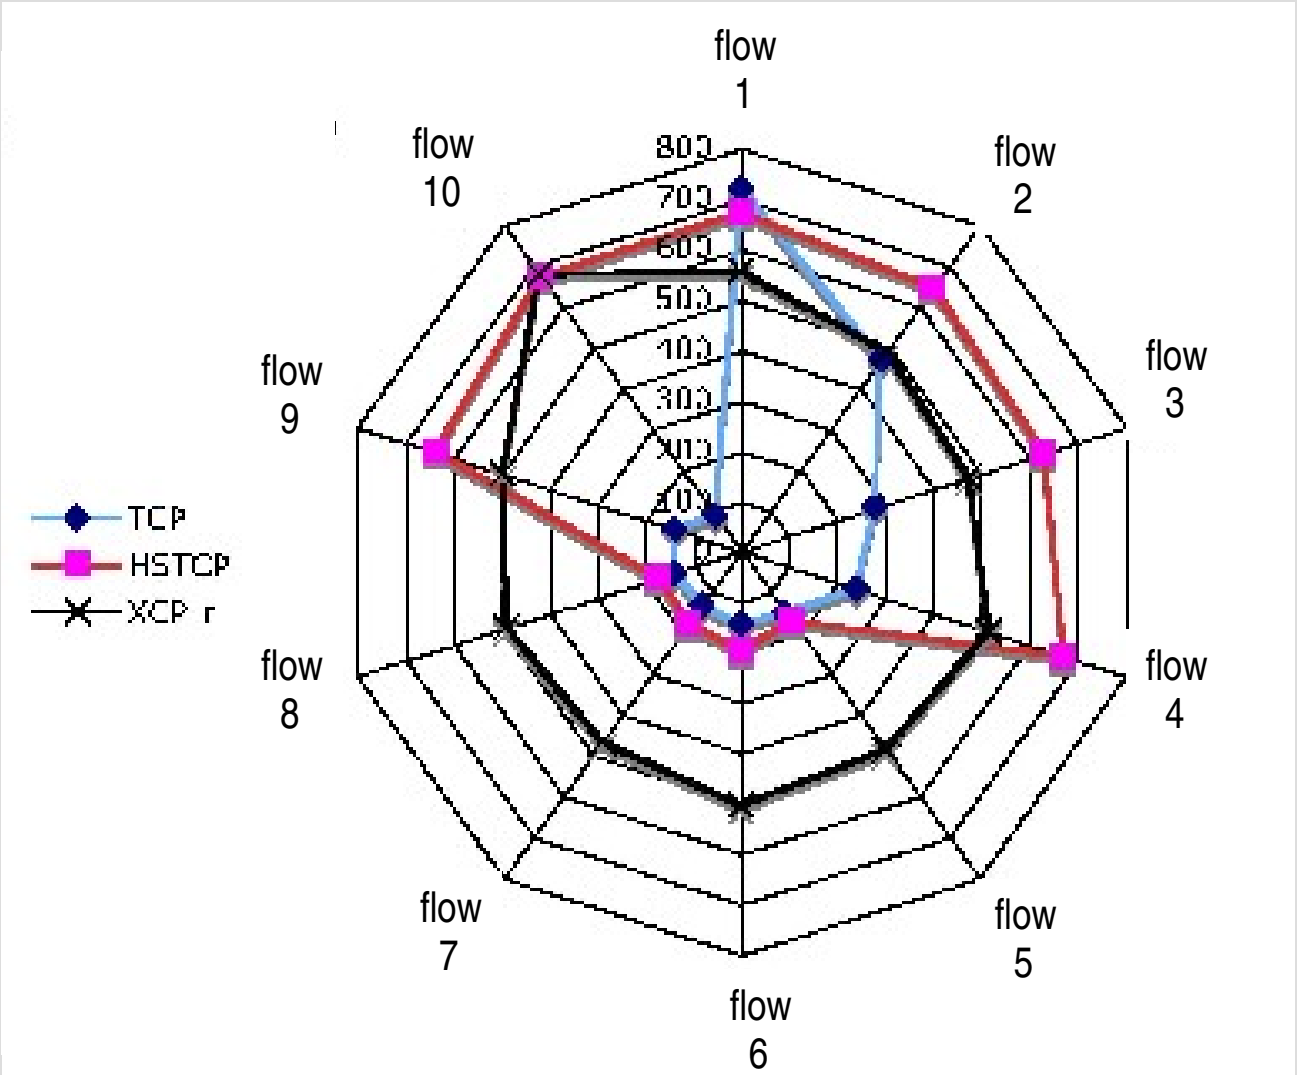

Standard Deviation

TCP :188.67

HSTCP : 233.86

XCP-r : 59.24

Conclusions

Conclusion

- The end-to-end solutions (like TCP, HSTCP, etc.) are inefficient and unfair in DHSN.

- Router assisted approachs, like XCP and XCP-r, show their superiority in term of fairness and efficiency.

- The XCP behavior is seriously affected in unfriendly HSN.

- XCP-r is a more robust XCP version, that provides
  - a high performance,
  - fast recovery after a problem,
  - fairness and high stability for large data transfers

  in a broader range of network conditions.

Future works

- Inter-operability between XCP routers & non assisted routers.

- Fairness between XCP and TCP.

Thank you!

dmlopezp@ens-lyon.fr