dirceu@ndrc.kyutech.ac.jp *Dirceu Cavendish*

UCLA /KIT

Lossless Congestion Control
Delay based LCC
LCC Requirements/Candidates
Capacity/Congestion Probing
CCP Simulations
CCP Experiments
LCC and IETF (help)

Lossless Congestion Control

Motivation Control packet retransmissions, which is undesirable for networks and applications alike.

Benefits APPLICATIONS:
- Fresher packets/segments are delivered.
- Shallower sender/receiver buffers can be used.
- Old data delivery is avoided.

NETWORKS:
- Higher resource utilization and aggregate goodput.

How -Most popular TCPs are packet loss driven. We need delay based congestion control protocols, to shift TCP operating point away from buffer overflow.
-Lossless congestion control (LCC) protocols should avoid operating on near packet loss point.
-LCC protocols should be conservative towards throughput, limiting it to "safe" levels for the network AND appropriate levels for application.

Delay based LCC

Delay based TCPs

-Listens to segment rtts. Most OSs support at least microsecond rtt measurement accuracy.
-Regulate transmission rate to keep segments' rtts at an acceptable level.
-Disambiguates between loss and congestion

Delay based LCC

- Buffer filling levels are kept low
- Network buffers are used to cope with excessive in flight segments during network transients
- Focuses on network utilization with packet loss control.

CC operating points

-Full buffer (losses)
-Empty buffer (throughput degradation)
-Anything in between (loss/throughput tradeoff)

LCC Requirements & Candidates

Premisses

- Retransmissions are undesirable for both applications and networks
- Throughput at any cost is undesirable (fairness, discard at receiver).

Delay based LCC

- Senders monitor rtts.
- Senders regulate their TX rate so as to keep rtts at a given operating point. Queues are kept away from their overflow levels.
- Most delay based TCPs do not operate at "knee of the congestion curve", but much above, incurring high losses, as a trade-off for high throughput.

LCC Candidates

[Leith07] D. Leith, R. Shorten, G. McCullagh, J. Heffner, L. Dunn, F. Baker, "Delay-based AIMD Congestion Control", in PFLDnet, February 2007.

[Cavendish07] D. Cavendish, C. Marcondes, M. Gerla, "Capacity and Congestion Probing: Towards a Stable and Lossless TCP", Submitted to Infocom 2008.

dirceu@ndrc.kyutech.ac.jp
Dirceu Cavendish
UCLA /KIT

Lossless Congestion Control
Delay based LCC
LCC Requirements/Candidates
Capacity/Congestion Probing
CCP Simulations
CCP Experiments
LCC and IETF (help)

Capacity & Congestion Probing TCP

Proportional + Integral controller

- Based on control theoretical approach [Cavendish04]
- Estimate session path bottleneck capacity and storage space
- cwnd(k) = f (storage(k), inFlight(k));

CCP control properties

-Timeout driven window regulation
-Guaranteed window convergence
-Allows throughput vs loss tradeoff tuning

TCP-CCP Protocol

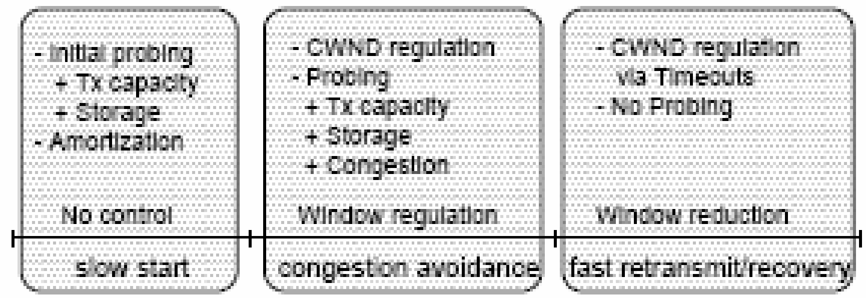

Fig. 5. CCP protocol

UCLA /KIT

- Lossless Congestion Control
- Delay based LCC
- LCC Requirements/Candidates
- Capacity/Congestion Probing
- CCP Simulations
- CCP Experiments
- LCC and IETF (help)

Parking Lot Simulation Results

Estimators' accuracy

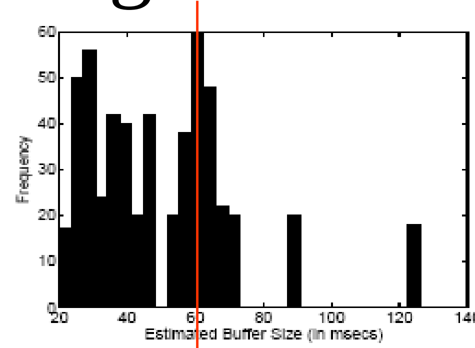

a) Buffer Estimation (rttmax - rttmin)

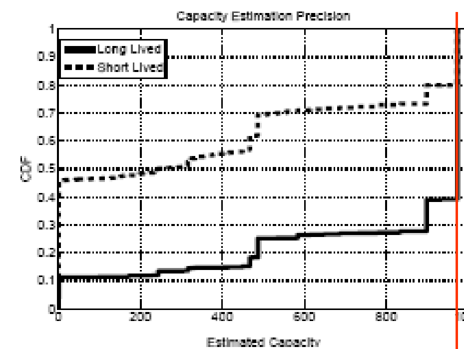

b) CapacityEstimation

Performance comparison
CCP, NewReno, FAST

CCP: 40/50 % less gput
20/200x less loss

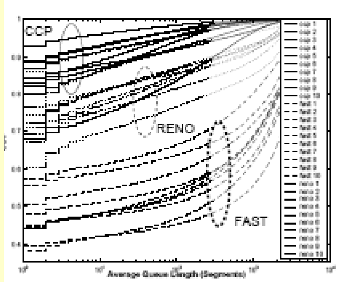

a) Buffer Utilization

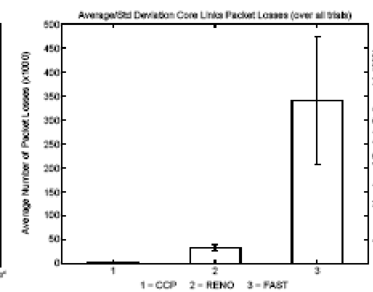

b) Segment Losses

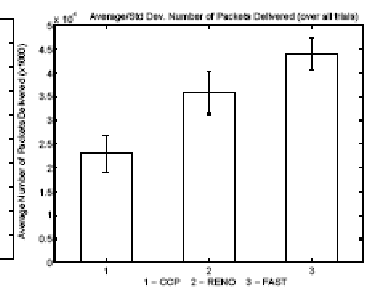

c) Goodput

Dynamics

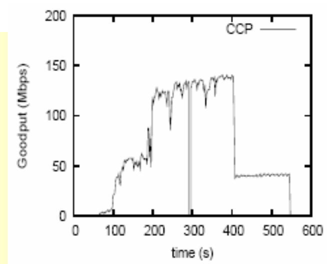

a) CCP

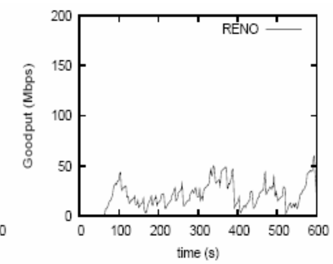

b) Reno

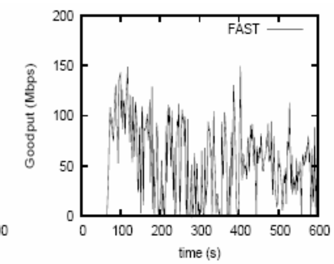

c) Fast

Fig. 10. Goodput Dynamics

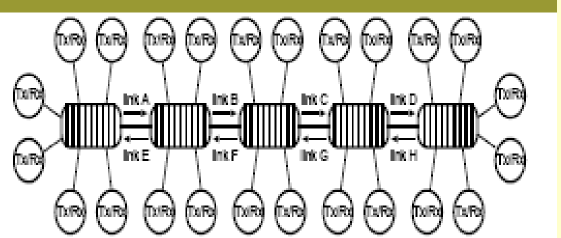

All links: 1Gbps; 15msec Exp.Dist. 4MB drop-tail
b) Parking lot
Fig. 6. Simulated Networks

NETWORK SCENARIO
Parking Lot topology
1Gbps all links, 15msec delays
140 flows
- 40 long lived (4Gfiles)
- 100 short lived (1MB Pareto)
800Mbps load on core links

dirceu@ndrc.kyutech.ac.jp
Dirceu Cavendish

UCLA /KIT

Transoceanic Experiments

Lossless Congestion Control

Delay based LCC

LCC Requirements/Candidates

Capacity/Congestion Probing

CCP Simulations

CCP Experiments

LCC and IETF (help)

BIC/Highspeed/CCP

Dynamics

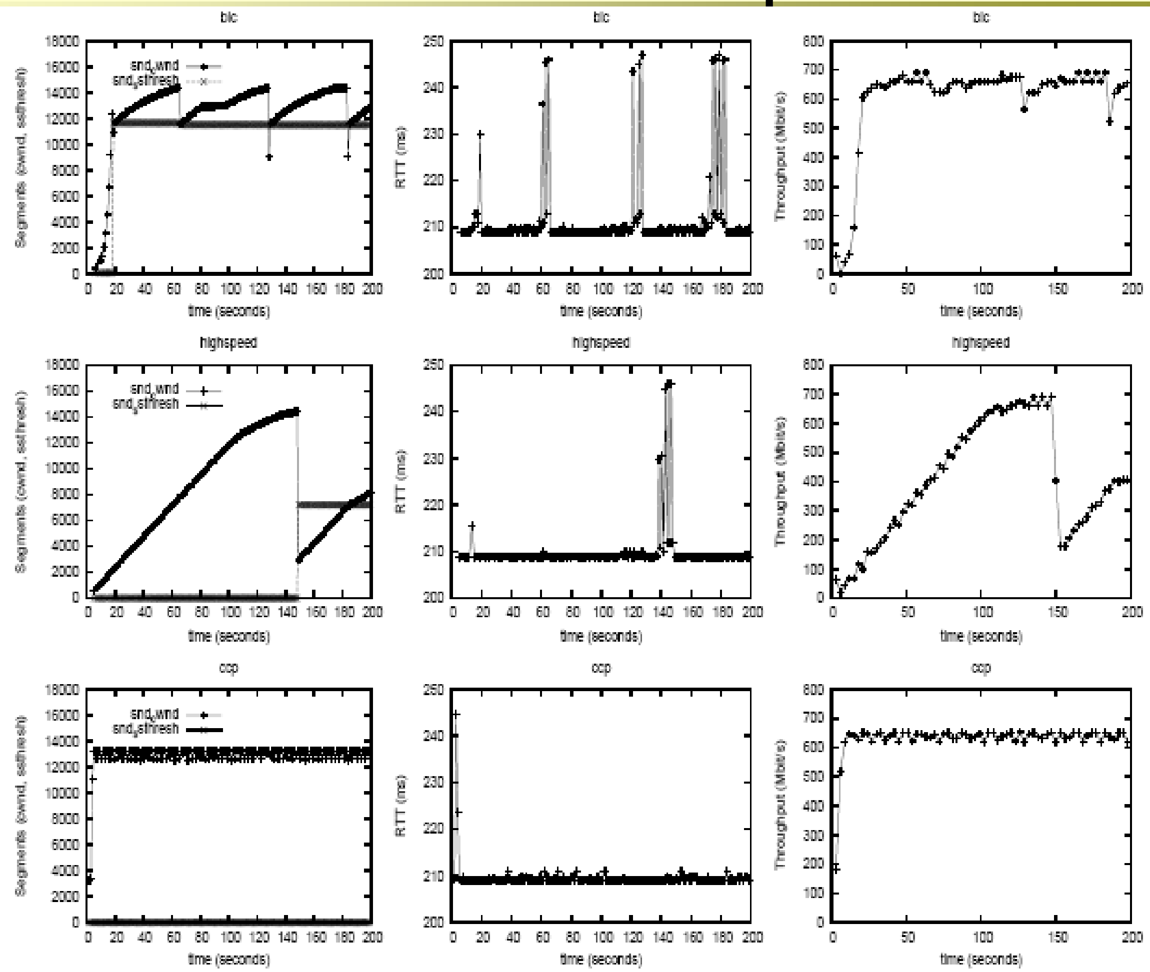

Fig. 11. TCP Transients - Transatlantic Scenario

NETWORK SCENARIO
Clean Pipe
1Gbps narrower link
208msec rtt UCLA/KIT
Pathrate/pathload tested
Large socket buffers
Iperf application
All (9) Linux supported algos

UCLA /KIT

LCC and IRTF

DCCP
- Offers multiple congestion control options:
  + TCP-Like
    LCC could fit here
  + TFRC
    Equation based rate control: r(t) = f(loss_rate)
- Active on accommodating applications such as RealAudio, Internet Telephony, and Interactive Games into a congestion control framework.

LCC
-Sequence numbers are useful for rtt tracking purposes
-Nanosecond level accuracy is useful for certain path scenarios

Next Steps?
-Our actions?
-Volunteers?

UCLA /KIT

Thank you !

Collaborators
Cesar Marcondes – UCLA
Mario Gerla – UCLA
Yuji Oie - KIT

References

[Cavendish07] D. Cavendish, C. Marcondes, M. Gerla, "Capacity and Congestion Probing: Towards a Stable and Lossless TCP", Submitted to Infocom 2008.

[Leith07] D. Leith, R. Shorten, G. McCullagh, J. Heffner, L. Dunn, F. Baker, "Delay-based AIMD Congestion Control", in PFLDnet, February 2007.

[Cavendish04] D. Cavendish, M. Gerla, S. Mascolo, "A Control Theoretical Approach to Congestion Control in Packet Networks", In Transactions on Networking, Vol. 42, Issue 5, pp. 893-906, Oct. 2004.