Speeding up the Internet: Fast Startup Congestion Control

Michael Scharf <michael.scharf@googlemail.com>

This work was performed at the Institute of Communication Networks and Computer Engineering (IKR) at the University of Stuttgart.

A Note on these slides

The following diagrams are taken from or based on:

M. Scharf, “Fast Startup Internet Congestion Control for Broadband Interactive Applications”, PhD thesis (under submission), University of Stuttgart, Nov. 2009

Please refer to the PhD thesis when citing these results.

Scope

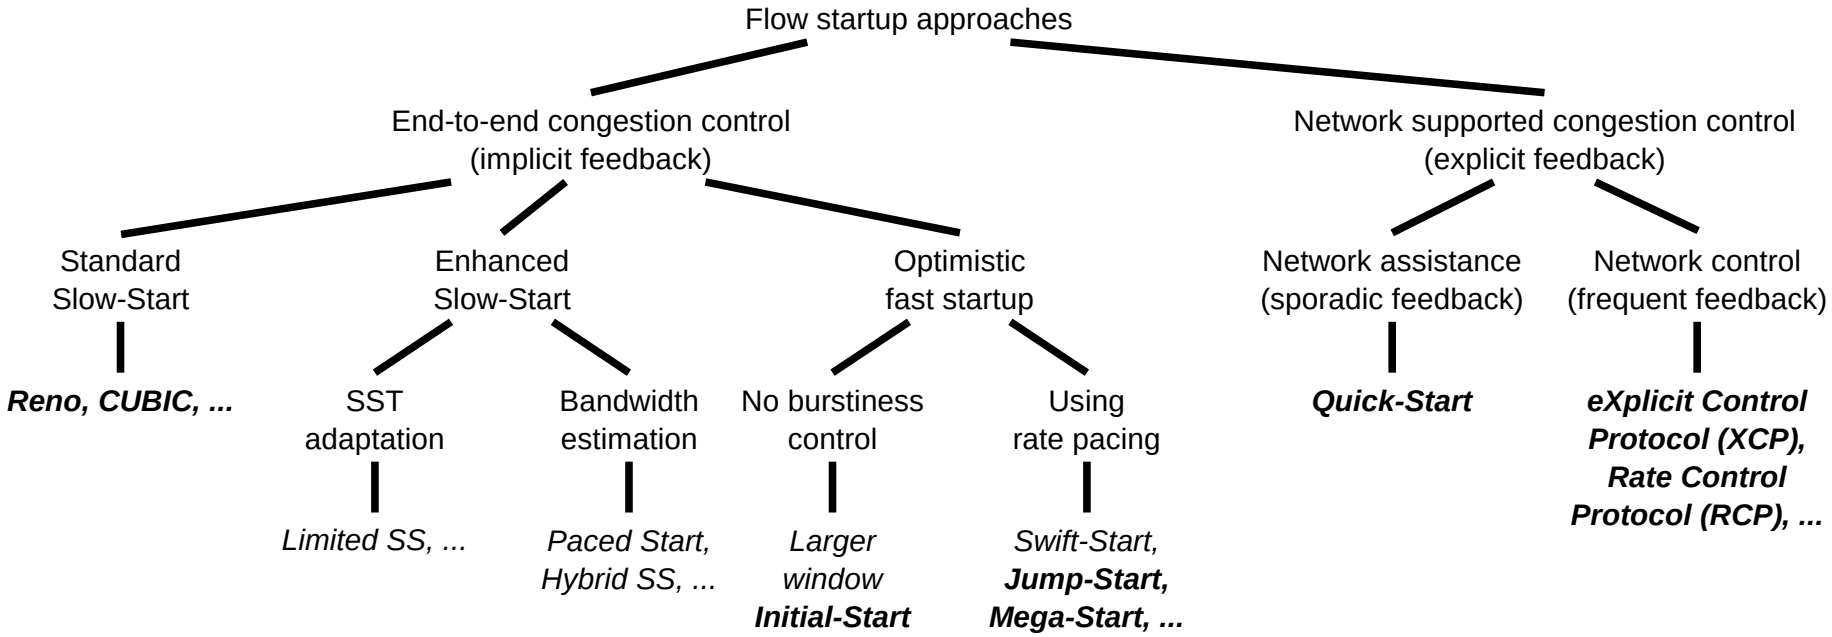

- TCP's standard Slow-Start with CUBIC (**SS**)
- **Initial congestion window of 10 MSS**, in the diagrams named Initial-Start (**IS**)
- Jump-Start of M. Allman et al., slightly modified to reduce aggressiveness (**JS**)
- Quick-Start TCP extension according to RFC 4782 (**QS**)
- Rate Control Protocol (**RCP**)

Evaluation methodology

- **Simulations**
  - Simulation with **Linux code** using the NSC framework
  - Own Linux patches for all TCP extensions, and an own tool for RCP
- **Considered scenarios**
  - Subset of the **TCP evaluation suite**
  - **Dumbbell topology** with 9 different RTTs
  - Bottleneck typically 10 Mbit/s, 50 packets buffer, drop tail
  - Replay of measured **Internet traces** in a-b-t format as recommended in TCP evaluation suite
- **Correctness of implementations** verified by simple testbed measurements

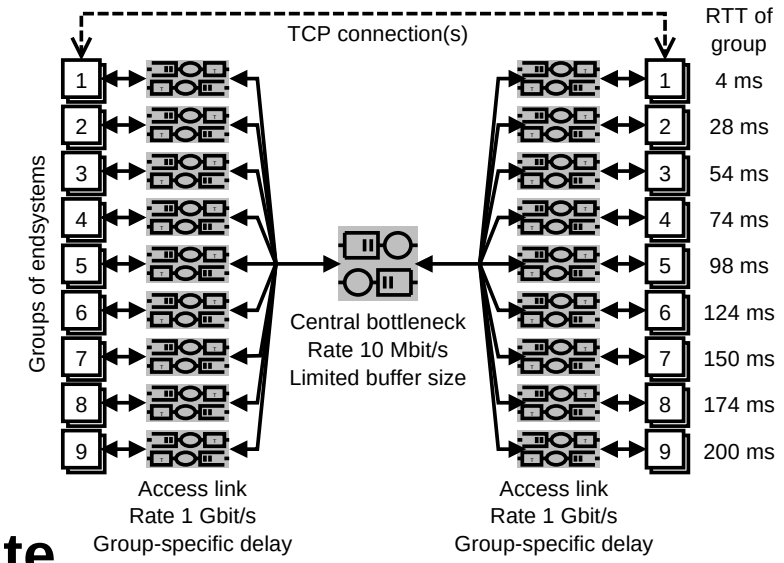

Possible speedup of the different variants

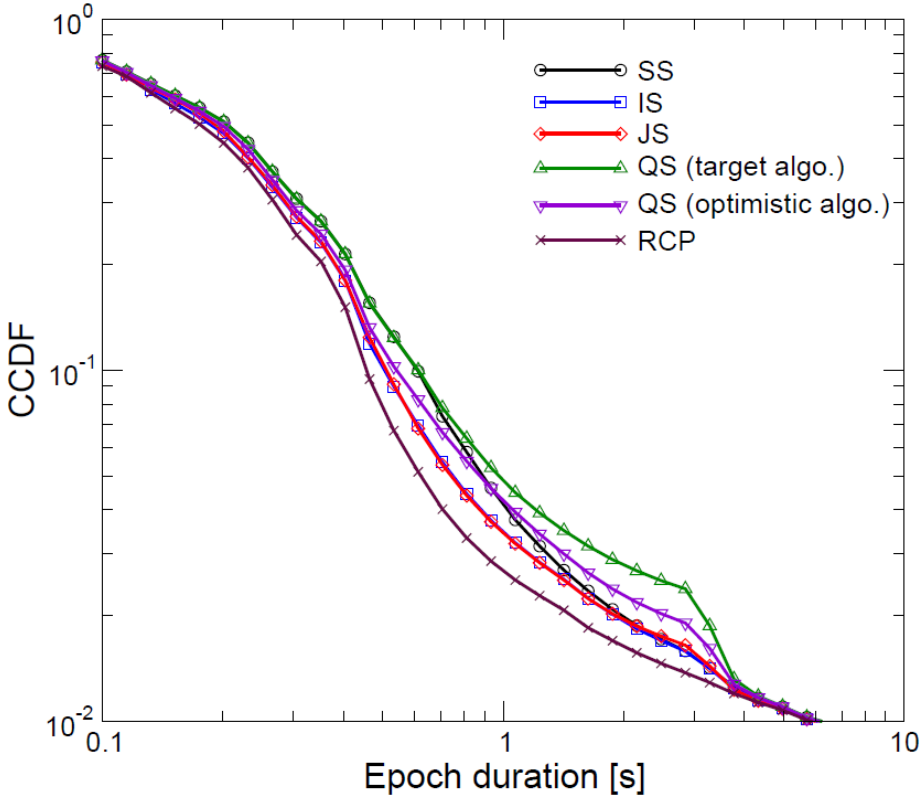

- Performance metric: **Response time** of a-b-t transfers (“epoch duration”)
- **Speedup of mid-sized transfers** by larger initial window
- **Overall benefit is rather small**: Many short transfers, many small RTTs

Trade-off between speedup and packet loss

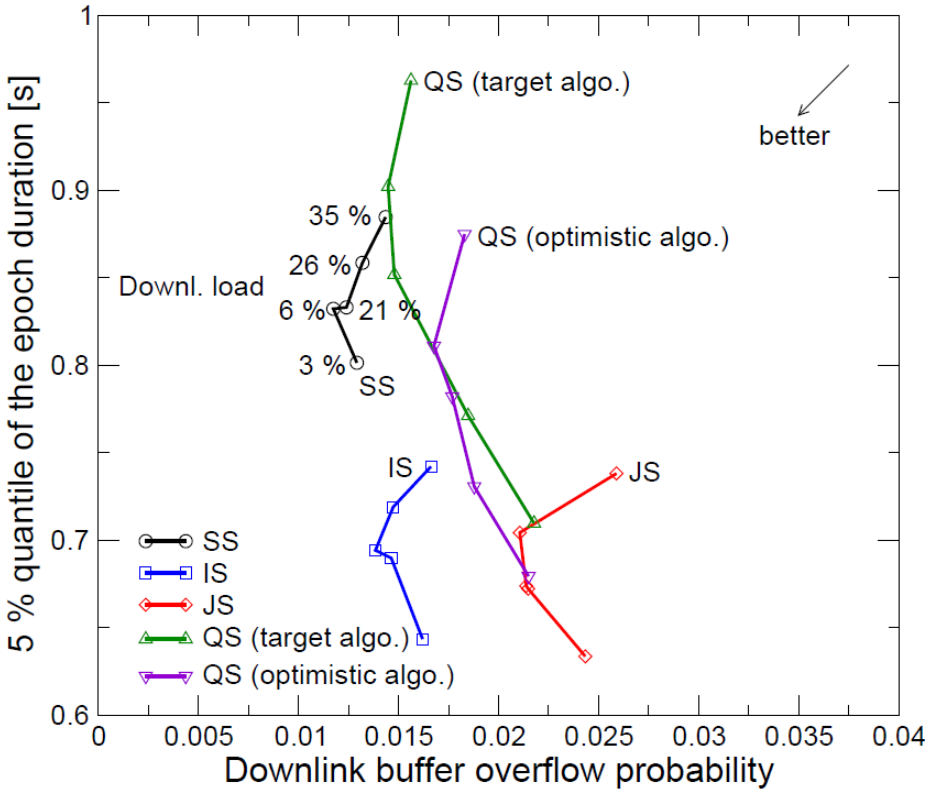

- Bottleneck load up to ca. 40% (due to tool limitation to ca. 1000 stacks)
- Absolute **loss probability increases by ca. 0.5%** by IW=10 MSS
- Somehow surprisingly, **other alternatives have a worse performance**

Fairness to unmodified stacks

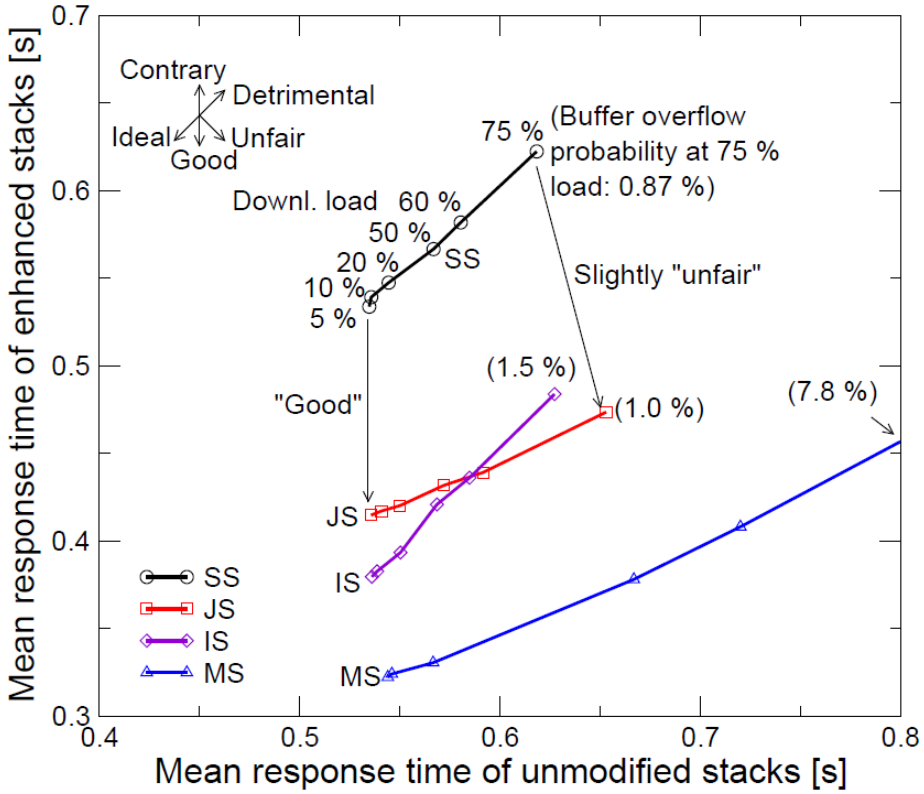

- **Scenario**: 50% of stacks use fast startup, 50% unchanged (CUBIC)
- Slightly changed workload model to allow simulations **up to 75% load**
- **Larger initial window** is rather fair and hardly impacts other flows

Conclusion

Results

- **Moderate benefit** for larger transfers
- Initial window of 10 MSS **works rather well and is quite fair**
- **More sophisticated schemes** tend to be worse
- **Network support** such as Quick-Start can overcome some limitations, but it has **problems of its own**

Recommendations for further work

- Study more extensively the use of **rate pacing**, even if results suggests that it may not be needed for 10 MSS
- Rethink **error recovery algorithms** after fast startup, since there are many degrees of freedom there, too

Selected references

Performance of fast startup congestion control schemes

- M. Scharf. Performance evaluation of fast startup congestion control schemes. In Proc. IFIP Networking 2009, LNCS 5550, Springer-Verlag, pp. 716–727, 2009
- M. Proebster, M. Scharf, and S. Hauger. Performance comparison of router assisted congestion control protocols: XCP vs. RCP. In Proc. 2nd International Workshop on the Evaluation of Quality of Service through Simulation in the Future Internet, 2009
- M. Scharf and H. Strotbek. Performance evaluation of Quick-Start TCP with a Linux kernel implementation. In Proc. IFIP Networking 2008, LNCS 4982, Springer-Verlag, pp. 703–714, 2008
- M. Scharf. Performance analysis of the Quick-Start TCP extension. In Proc. IEEE Broadnets, 2007

Implementation aspects and applicability case studies

- M. Scharf, M. Eissele, C. Mueller, and T. Ertl. Speeding up the 3D Web: A case for fast startup congestion control. In Proc. PFLDNeT, 2009
- S. Hauger, M. Scharf, J. Kögel, and C. Suriyajan. Quick-Start and XCP on a network processor: Implementation issues and performance evaluation. In Proc. IEEE HPSR 2008, 2008.
- M. Scharf, S. Hauger, and J. Kögel. Quick-Start TCP: From theory to practice. In Proc. PFLDnet, 2008

IETF contributions

- D. Papadimitriou, M. Welzl, M. Scharf, and B. Briscoe. Open research issues in Internet congestion control. IRTF Internet Draft, work in progress, August 2009
- M. Scharf, S. Floyd, and P. Sarolathi. TCP flow control for fast startup schemes. IETF Internet Draft, work in progress, July 2008