

The Virtue of Gentleness: Improving Connection Response Times with SYN Priority Active Queue Management

Tristan Braud¹, **Martin Heusse**, and Andrzej Duda



Grenoble Informatics Lab. — Grenoble Institute of Technology

¹T. Braud is now at the Hong Kong University of Science and Technology

Numerous factors may impact TCP transfers responsiveness

- Available throughput, concurrent traffic
- Bufferbloat / Upload-download interference on asymmetric lines
A buffer too deep at the bottleneck may cause delays in the order of seconds
- (Tail losses)

⇒ Losses during **connection establishment**

- ✓ Conservative **initial** RTO → long retransmission delays

Outline

- How prevalent and damaging are TCP **SYN and SYN/ACK losses**?

- How harmful in reality?

Traces analysis

- What can we do about them?

SYN protection

- How effective are various counter measures?

Testbed experiments

TCP SYN and SYN/ACK losses

- SYN packet retransmission occurs whenever
 - The packet is lost ;
 - or
 - The SYN/ACK is lost...


Loss probability on the forward and return paths
add to each other

- Connection phase duration for SYN loss probability p_{SI} :

$$t_{syn} = RTT + RTO_0 \sum_{k=1}^{\infty} (p_{SI})^k$$

- ✓ RTO_0 is **1 to 3 seconds** !
- ✓ $p_{SI} \approx 2 \times p_I$ (as per the argument above)

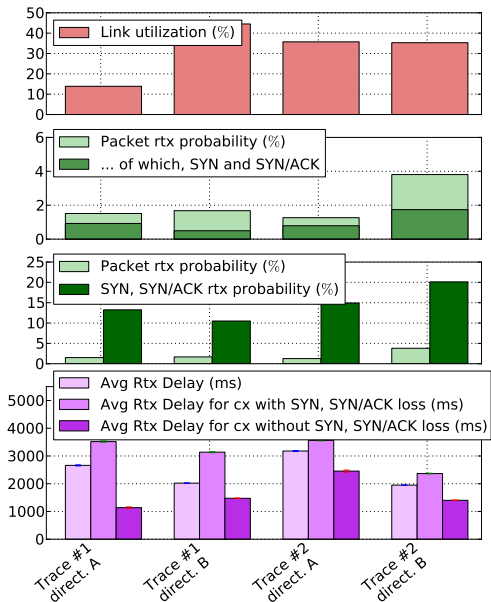
SYN retrans. impact — Simple model²




Trace analysis

- CAIDA and MAWI public traces
 - ✓ Hundred of thousand of connections / trace
 - ✓ CAIDA: 10 Gigabit Ethernet backbone link of a Tier 1 ISP
 - ✓ MAWI: 1Gb/s transit link between WIDE and an upstream ISP
- MAWI trace is bidirectional: it allows us to filter out SYN flood attacks

CAIDA trace



MAWI trace



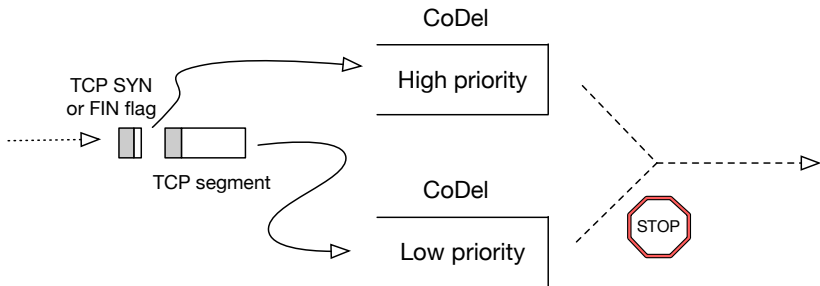
(Statistics for only the connections with an effective data transfer)

Model and Trace analysis conclusions

- SYN and SYN/ACK losses are not unusual
- ... in fact they are more numerous than TCP segment losses!
- Noticeable impact, since the initial RTO is large
(Except for very large transfers)

Avoiding SYN and SYN/ACK losses

SPA: SYN priority AQM

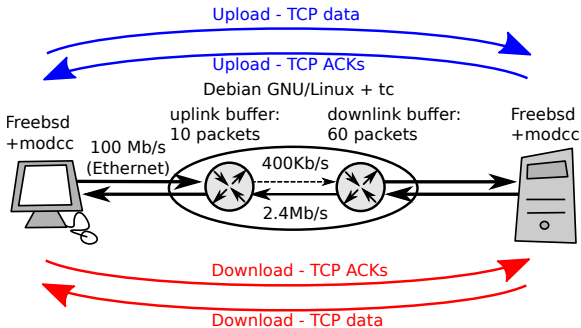


- We use two CoDel³ queues, and prioritize SYN and FIN packets over all other packets
- Much simpler than e.g. FQ CoDel⁴

³K. Nichols and V. Jacobson, “Controlling Queue Delay,” ACM Queue, May 2012


⁴draft-hoeiland-joergensen-aqm-fq-codel-00

Test bench



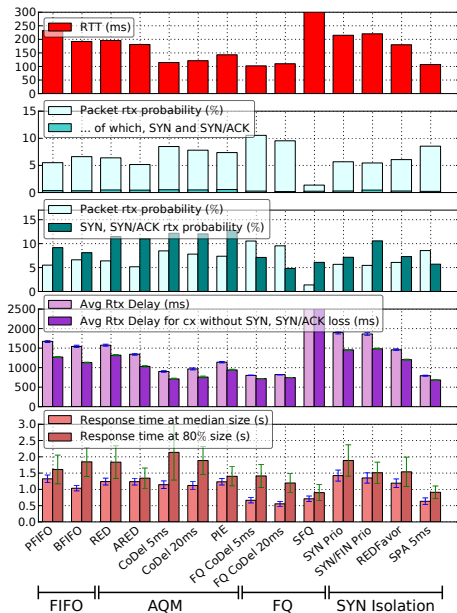
Downlink: Packet FIFO - Byte FIFO - RED - ARED - REDFavor - CoDel - FQ CoDel - PIE - SFQ - SYN Prio - SYN/FIN Prio - SPA

Test bench results – baseline



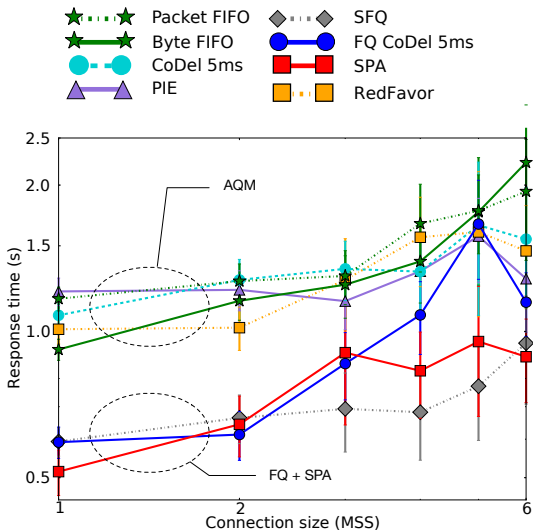
- Realistic heavy tailed transfer sizes
- **Unidirectional** traffic
- FQ CoDel, SFQ, RedFavor, and SPA achieve the shortest response times

Test bench results – with reverse traffic



- Realistic heavy tailed transfer sizes
- **Bulk upload**
- SYN retransmissions increase
- FQ CoDel, SFQ, SPA provide the shortest response times

Test bench results – with reverse traffic (cont.)



Conclusion

- Discarding SYN packets is not a good idea, and it does not seem as though they are treated with much care, out there!
- SPA is the simple combination of 2 CoDel queues (no fair queueing)
- It performs similarly to combining an AQM with fair queueing
- ... In fact, FQ schemes are effective mostly because they protect the connection establishment!

Future research

- There are more SYN losses in the traces than on the testbed.

Why?

What does our synthetic load and/or trace analysis not capture? (CAIDA trace analysis does not filter out SYN flood attacks, though)

N.B.: even with few losses,
SYN losses still have a dramatic impact

- Investigate other factors detrimental to the response time:
DNS?
- What about SYN flood attacks?