HyStart++: Modified Slow Start for TCP

draft-balasubramanian-tcpm-hystartplusplus-03

TCPM interim meeting
April 29, 2020

Praveen Balasubramanian, Yi Huang, Matt Olson

Microsoft

HyStart++ Recap

- Slow Start can overshoot ideal send rate and cause massive packet loss
- HyStart: Exit slow start early based on Delay Increase algorithm
  - Inter-Packet Arrival algorithm does not perform well due to ACK compression
- Compensate for premature slow start exit
  - Congestion Avoidance algorithm can take time to ramp up
- Use maximum of cwnd computed by Limited Slow Start (RFC3742) and Congestion Avoidance, until next congestion signal

Performance data

- Large scale A/B test covering billions of flows on Windows systems
  - Reduction in retransmission timeouts
    - 99% of connections have fewer than 2 RTOs over lifetime
    - 0.64% connections moved from 1 RTO to 0 RTOs over lifetime
    - 0.7% connections moved from 2 RTOs to 1 RTO over lifetime
  - Working on getting more production data and metrics

- Lab data
  - Fair towards non-HyStart++ flows
  - 100 Mbps bandwidth, BDP size bottleneck buffer
  - For large RTT flows (100 msec)
    - Up to 39% improvement in average and P90 goodput for short flows
    - Up to 14% improvement in average and P90 goodput for long flow
  - No noticeable improvement for small RTT flows (50 msec, 25 msec)
  - Across all tests
    - Number of bytes retransmitted reduced by 50%
    - Number of RTOs reduced by 36%
    - Loss recovery success rate improves 43.48% -> 52%

Changes in draft-03

- Incorporated review feedback from Neal Cardwell, Martin Duke, Ilpo Järvinen, Christoph Paasch, and Junho Choi
  - Thanks for the reviews!

- Summary of changes
  - Clarified relationship with Appropriate Byte Counting
  - Clarified when HyStart++ ends
  - Fixed some equations that used bytes versus segments
  - Variable name changes in pseudocode

Status & Next Steps

- HyStart++ is deployed on by default for all connections
  - Windows 10 May 2019 Update onwards
  - Windows Server 2019 1903 version onwards
- Look into usage of bandwidth or throughput estimate
- Future: compare HyStart++, BBR STARTUP phase, and Paced Chirping
- Adopt document in tcpm