

Revised CUBIC

TCPM, IETF 112

Nov 2021

Lisong Xu, Sangtae Ha, Injong Rhee, Vidhi Goel, Lars Eggert

Recent changes

Markku Kojo's extensive review

- Updates [RFC 5681] to allow CUBIC's more aggressive sending behavior
- Recommend Hystart++ for slow-start to avoid overshoot
- Extra care must be taken so that end of slow start and the first multiplicative decrease work well together
- In case of packet loss (only), sender MAY use PRR [RFC 6937] to reduce sending rate
- CUBIC, like Reno has slow adaptation in wireless environments

Recent changes

Markku Kojo's extensive review

- Proper queue sizing and management can mitigate risk of high queuing delay as CUBIC fills buffer faster than Reno.
- Reduce *cwnd* in response to ECE until it reaches 1 *MSS* and then use retransmit timer with exponential backoff
- Use *FlightSize* instead of *cwnd* after a congestion event
- For rate-limited apps, recommend RFC 7661 to mitigate the issue of *FlightSize* being significantly smaller than *cwnd*
- Some implementations currently use *cwnd* and may continue to do so

Recent changes

Markku Kojo's extensive review

- Fix a bug in average CUBIC window equation in Discussion
- Rephrase text around spurious retransmission detection algorithms
- CUBIC's response to sudden increase and decrease in capacity due to transient events
- Replace ACK by "new ACK" wherever applicable and define it clearly
- Rename AIMD (TCP) -> Reno
- Moved "MUST NOT" requirement about app-limited period from discussion to main congestion avoidance section

Open Issue

Spurious congestion events

CUBIC modifies RFC 4015 response to spurious RTOs

Restores state for spurious fast retransmits

- A path change can cause reordering and available network capacity may get reduced

Executes UNDO even if it received an ACK with ECE

Does not adjust RTO and RTO calculation to avoid further spurious RTOs

Open Issue

Spurious congestion events

CUBIC modifies RFC 4015 response to spurious RTOs

- $cwnd < - FlightSize + \min(bytes_acked, IW)$
- $ssthresh < - \max(FlightSize, ssthresh)$

Restores state for spurious fast retransmits

- A path change can cause reordering and available network capacity may get reduced

Executes UNDO even if it received an ACK with ECE

Does not adjust RTO etc to avoid further spurious events

Open Issue

Spurious congestion events

CUBIC modifies RFC 4015 response to spurious RTOs

- $cwnd < -FlightSize + \min(bytes_acked, IW)$ \longrightarrow Too conservative?
- $ssthresh < -\max(FlightSize, ssthresh)$

Restores state for spurious fast retransmits

\longrightarrow More common than spurious RTO

- A path change can cause reordering and available network capacity may get reduced

Executes UNDO even if it received an ACK with ECE

\longrightarrow Can be fixed

Does not adjust RTO etc to avoid further spurious events

\longrightarrow Separate draft

Next steps

- Address open issues
- Request chairs to conclude WGLC and request publication