

Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: 31 August 2023

P. Balasubramanian  
Confluent  
Y. Huang  
M. Olson  
Microsoft  
27 February 2023

HyStart++: Modified Slow Start for TCP  
draft-ietf-tcpm-hystartplusplus-14

Abstract

This document describes HyStart++, a simple modification to the slow start phase of congestion control algorithms. Slow start can overshoot the ideal send rate in many cases, causing high packet loss and poor performance. HyStart++ uses increase in round-trip delay as a heuristic to find an exit point before possible overshoot. It also adds a mitigation to prevent jitter from causing premature slow start exit.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 31 August 2023.

Copyright Notice

Copyright (c) 2023 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components

extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

## Table of Contents

1. Introduction . . . . .	2
2. Terminology . . . . .	3
3. Definitions . . . . .	3
4. HyStart++ Algorithm . . . . .	3
4.1. Summary . . . . .	3
4.2. Algorithm Details . . . . .	4
4.3. Tuning constants and other considerations . . . . .	6
5. Deployments and Performance Evaluations . . . . .	7
6. Security Considerations . . . . .	8
7. IANA Considerations . . . . .	8
8. Acknowledgements . . . . .	8
9. References . . . . .	8
9.1. Normative References . . . . .	8
9.2. Informative References . . . . .	8
Authors' Addresses . . . . .	9

## 1. Introduction

[RFC5681] describes the slow start congestion control algorithm for TCP. The slow start algorithm is used when the congestion window (cwnd) is less than the slow start threshold (ssthresh). During slow start, in absence of packet loss signals, TCP increases cwnd exponentially to probe the network capacity. This fast growth can overshoot the ideal sending rate and cause significant packet loss which cannot always be recovered efficiently.

HyStart++ uses increase in round-trip delay as a signal to exit slow start before potential packet loss occurs as a result of overshoot. This is one of two algorithms specified in [HyStart]. After the slow start exit, a new Conservative Slow Start (CSS) phase is used to determine whether the slow start exit was premature and to resume slow start. This mitigation improves performance in presence of jitter. HyStart++ reduces packet loss and retransmissions, and improves goodput in lab measurements and real world deployments.

While this document describes Hystart++ for TCP, it can also be used for other transport protocols which use slow start such as QUIC [RFC9002] or SCTP [RFC9260].

## 2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

## 3. Definitions

We repeat here some definition from [RFC5681] to aid the reader.

**SENDER MAXIMUM SEGMENT SIZE (SMSS):** The SMSS is the size of the largest segment that the sender can transmit. This value can be based on the maximum transmission unit of the network, the path MTU discovery [RFC1191], [RFC4821] algorithm, RMSS (see next item), or other factors. The size does not include the TCP/IP headers and options.

**RECEIVER MAXIMUM SEGMENT SIZE (RMSS):** The RMSS is the size of the largest segment the receiver is willing to accept. This is the value specified in the MSS option sent by the receiver during connection startup. Or, if the MSS option is not used, it is 536 bytes [RFC1122]. The size does not include the TCP/IP headers and options.

**RECEIVER WINDOW (rwnd):** The most recently advertised receiver window.

**CONGESTION WINDOW (cwnd):** A TCP state variable that limits the amount of data a TCP can send. At any given time, a TCP MUST NOT send data with a sequence number higher than the sum of the highest acknowledged sequence number and the minimum of cwnd and rwnd.

## 4. HyStart++ Algorithm

### 4.1. Summary

[HyStart] specifies two algorithms (a "Delay Increase" algorithm and an "Inter-Packet Arrival" algorithm) to be run in parallel to detect that the sending rate has reached capacity. In practice, the Inter-Packet Arrival algorithm does not perform well and is not able to detect congestion early, primarily due to ACK compression. The idea of the Delay Increase algorithm is to look for spikes in RTT (round-trip time), which suggest that the bottleneck buffer is filling up.

In HyStart++, a TCP sender uses traditional slow start and then uses the "Delay Increase" algorithm to trigger an exit from slow start. But instead of going straight from slow start to congestion avoidance, the sender spends a number of RTTs in a Conservative Slow

Start (CSS) phase to determine whether the exit from slow start was premature. During CSS, the congestion window is grown exponentially like in regular slow start, but with a smaller exponential base, resulting in less aggressive growth. If the RTT reduces during CSS, it's concluded that the RTT spike was not related to congestion caused by the connection sending at a rate greater than the ideal send rate, and the connection resumes slow start. If the RTT inflation persists throughout CSS, the connection enters congestion avoidance.

#### 4.2. Algorithm Details

The following pseudocode uses a limit,  $L$ , to control the aggressiveness of the cwnd increase during both standard slow start and CSS. While an arriving ACK may newly acknowledge an arbitrary number of bytes, the Hystart++ algorithm limits the number of those bytes applied to increase the cwnd to  $L \cdot \text{SMSS}$  bytes.

lastRoundMinRTT and currentRoundMinRTT are initialized to infinity at the initialization time. currRTT is the RTT sampled from the latest incoming ACK and initialized to infinity.

```
lastRoundMinRTT = infinity
currentRoundMinRTT = infinity
currRTT = infinity
```

Hystart++ measures rounds using sequence numbers, as follows: Define windowEnd as a sequence number initialized to SND.NXT. When windowEnd is ACKed, the current round ends and windowEnd is set to SND.NXT.

At the start of each round during standard slow start ([RFC5681]) and CSS, initialize the variables used to compute last round and current round's minimum RTT:

```
lastRoundMinRTT = currentRoundMinRTT
currentRoundMinRTT = infinity
rttSampleCount = 0
```

For each arriving ACK in slow start, where  $N$  is the number of previously unacknowledged bytes acknowledged in the arriving ACK:

Update the cwnd:

```
cwnd = cwnd + min(N, L * SMSS)
```

Keep track of minimum observed RTT:

```

currentRoundMinRTT = min(currentRoundMinRTT, currRTT)
rttSampleCount += 1

```

For rounds where at least N\_RTT\_SAMPLE RTT samples have been obtained and currentRoundMinRTT and lastRoundMinRTT are valid, check if delay increase triggers slow start exit:

```

if ((rttSampleCount >= N_RTT_SAMPLE) AND
    (currentRoundMinRTT != infinity) AND
    (lastRoundMinRTT != infinity))
  Compute a RTT Threshold clamped between MIN_RTT_THRESH and MAX_RTT_THRESH
  RttThresh = max(MIN_RTT_THRESH, min(lastRoundMinRTT / MIN_RTT_DIVISOR, MAX_RTT_THRESH))
  if (currentRoundMinRTT >= (lastRoundMinRTT + RttThresh))
    cssBaselineMinRtt = currentRoundMinRTT
    exit slow start and enter CSS

```

For each arriving ACK in CSS, where N is the number of previously unacknowledged bytes acknowledged in the arriving ACK:

Update the cwnd:

```

cwnd = cwnd + (min(N, L * SMSS) / CSS_GROWTH_DIVISOR)

```

Keep track of minimum observed RTT:

```

currentRoundMinRTT = min(currentRoundMinRTT, currRTT)
rttSampleCount += 1

```

For CSS rounds where at least N\_RTT\_SAMPLE RTT samples have been obtained, check if current round's minRTT drops below baseline indicating that HyStart exit was spurious:

```

if (currentRoundMinRTT < cssBaselineMinRtt)
  cssBaselineMinRtt = infinity
  resume slow start including HyStart++

```

CSS lasts at most CSS\_ROUNDS rounds. If the transition into CSS happens in the middle of a round, that partial round counts towards the limit.

If CSS\_ROUNDS rounds are complete, enter congestion avoidance by setting ssthresh to current cwnd.

```

ssthresh = cwnd

```

If loss or ECN-marking is observed anytime during standard slow start or CSS, enter congestion avoidance by setting ssthresh to current cwnd.

```
ssthresh = cwnd
```

#### 4.3. Tuning constants and other considerations

It is RECOMMENDED that a HyStart++ implementation use the following constants:

```
MIN_RTT_THRESH = 4 msec
MAX_RTT_THRESH = 16 msec
MIN_RTT_DIVISOR = 8
N_RTT_SAMPLE = 8
CSS_GROWTH_DIVISOR = 4
CSS_ROUNDS = 5
L = infinity if paced, L = 8 if non-paced
```

These constants have been determined with lab measurements and real world deployments. An implementation MAY tune them for different network characteristics.

The delay increase sensitivity is determined by MIN\_RTT\_THRESH and MAX\_RTT\_THRESH. Smaller values of MIN\_RTT\_THRESH may cause spurious exits from slow start. Larger values of MAX\_RTT\_THRESH may result in slow start not exiting until loss is encountered for connections on large RTT paths.

MIN\_RTT\_DIVISOR is a fraction of RTT to compute delay threshold. A smaller value would mean a bigger threshold and thus less sensitive to delay increase, and vice versa.

While all TCP implementations are REQUIRED to take at least one RTT sample each round, implementations of HyStart++ are RECOMMENDED to take at least N\_RTT\_SAMPLE RTT samples. Using lower values of N\_RTT\_SAMPLE will lower the accuracy of the measured RTT for the round; higher values will improve accuracy at the cost of more processing.

The minimum value of CSS\_GROWTH\_DIVISOR MUST be at least 2. A value of 1 results in the same aggressive behavior as regular slow start. Values larger than 4 will cause the algorithm to be less aggressive and maybe less performant.

Smaller values of CSS\_ROUNDS may miss detecting jitter and larger values may limit performance.

Packet pacing [ASA00] is a possible mechanism to avoid large bursts and their associated harm. A paced TCP implementation SHOULD use L = infinity. Burst concerns are mitigated by pacing and this setting allows for optimal cwnd growth on modern networks.

For TCP implementations that pace to mitigate burst concerns, L values smaller than INFINITY may suffer performance problems due to slow cwnd growth in high speed networks. For non-paced TCP implementations, L values smaller than 8 may suffer performance problems due to slow cwnd growth in high speed networks; L values larger than 8 may cause an increase in burstiness and thereby loss rates, and result in poor performance.

An implementation SHOULD use HyStart++ only for the initial slow start (when ssthresh is at its initial value of arbitrarily high per [RFC5681]) and fall back to using traditional slow start for the remainder of the connection lifetime. This is acceptable because subsequent slow starts will use the discovered ssthresh value to exit slow start and avoid the overshoot problem. An implementation MAY use HyStart++ to grow the restart window ([RFC5681]) after a long idle period.

In application limited scenarios, the amount of data in flight could fall below the bandwidth-delay product (BDP) and result in smaller RTT samples which can trigger an exit back to slow start. It is expected that a connection might oscillate between CSS and slow start in such scenarios. But this behavior will neither result in a connection prematurely entering congestion avoidance nor cause overshooting compared to slow start.

## 5. Deployments and Performance Evaluations

As of February 2023, HyStart++ as described in this document has been default enabled for all TCP connections in the Windows operating system for over two years with pacing disabled and an actual L = 8.

In lab measurements with Windows TCP, HyStart++ shows both goodput improvements as well as reductions in packet loss and retransmissions compared to traditional slow start. For example, across a variety of tests on a 100 Mbps link with a bottleneck buffer size of bandwidth-delay product, HyStart++ reduces bytes retransmitted by 50% and retransmission timeouts (RTOs) by 36%.

In an A/B test where we compare HyStart++ draft 01 to traditional slow start across a large Windows device population, out of 52 billion TCP connections, 0.7% of connections move from 1 RTO to 0 RTOs and another 0.7% connections move from 2 RTOs to 1 RTO with HyStart++. This test did not focus on send-heavy connections and the impact on send-heavy connections is likely much higher. We plan to conduct more such production experiments to gather more data in the future.

## 6. Security Considerations

HyStart++ enhances slow start and inherits the general security considerations discussed in [RFC5681].

An attacker can cause Hystart++ to exit slow start prematurely and impair the performance of a TCP connection by, for example, dropping data packets or their acknowledgements.

The ACK division attack outlined in [SCWA99] does not affect Hystart++ because the congestion window increase in Hystart++ is based on the number of bytes newly acknowledged in each arriving ACK rather than by a particular constant on each arriving ACK.

## 7. IANA Considerations

This document has no actions for IANA.

## 8. Acknowledgements

During the discussions of this work on the TCPM mailing list, in working group meetings, helpful comments, critiques, and reviews were received from (listed alphabetically by last name): Mark Allman, Bob Briscoe, Neal Cardwell, Yuchung Cheng, Junho Choi, Martin Duke, Reese Enghardt, Christian Huitema, Ilpo Järvinen, Yoshifumi Nishida, Randall Stewart, and Michael Tuexen.

## 9. References

### 9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC5681] Allman, M., Paxson, V., and E. Blanton, "TCP Congestion Control", RFC 5681, DOI 10.17487/RFC5681, September 2009, <<https://www.rfc-editor.org/info/rfc5681>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

### 9.2. Informative References

- [ASA00] Aggarwal, A., Savage, S., and T. Anderson, "Understanding the Performance of TCP Pacing", Proceedings IEEE INFOCOM 2000, DOI 10.1109/INFCOM.2000.832483, 2000, <<https://doi.org/10.1109/INFCOM.2000.832483>>.
- [HyStart] Ha, S. and I. Ree, "Taming the elephants: New TCP slow start", Computer Networks vol. 55, no. 9, pp. 2092-2110, DOI 10.1016/j.comnet.2011.01.014, 2011, <<https://doi.org/10.1016/j.comnet.2011.01.014>>.
- [RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, DOI 10.17487/RFC1122, October 1989, <<https://www.rfc-editor.org/info/rfc1122>>.
- [RFC1191] Mogul, J. and S. Deering, "Path MTU discovery", RFC 1191, DOI 10.17487/RFC1191, November 1990, <<https://www.rfc-editor.org/info/rfc1191>>.
- [RFC4821] Mathis, M. and J. Heffner, "Packetization Layer Path MTU Discovery", RFC 4821, DOI 10.17487/RFC4821, March 2007, <<https://www.rfc-editor.org/info/rfc4821>>.
- [RFC9002] Iyengar, J., Ed. and I. Swett, Ed., "QUIC Loss Detection and Congestion Control", RFC 9002, DOI 10.17487/RFC9002, May 2021, <<https://www.rfc-editor.org/info/rfc9002>>.
- [RFC9260] Stewart, R., Tüxen, M., and K. Nielsen, "Stream Control Transmission Protocol", RFC 9260, DOI 10.17487/RFC9260, June 2022, <<https://www.rfc-editor.org/info/rfc9260>>.
- [SCWA99] Savage, S., Cardwell, N., Wetherall, D., and T. Anderson, "TCP congestion control with a misbehaving receiver", ACM Computer Communication Review, 29(5), DOI 10.1145/505696.505704, 1999, <<https://doi.org/10.1145/505696.505704>>.

## Authors' Addresses

Praveen Balasubramanian  
Confluent  
899 West Evelyn Ave  
Mountain View, CA 94041  
United States of America  
Email: [pravb.ietf@gmail.com](mailto:pravb.ietf@gmail.com)

Yi Huang  
Microsoft  
One Microsoft Way  
Redmond, WA 94052  
United States of America  
Phone: +1 425 703 0447  
Email: huanyi@microsoft.com

Matt Olson  
Microsoft  
Phone: +1 425 538 8598  
Email: maolson@microsoft.com