A Study on Performances Behaviors of TCP BBR and CUBIC TCP in Deep Buffer Network

Kanon Sasaki†
Electrical Eng. and Electronics
Kogakuin Univ. Graduate School
Tokyo, Japan
cm19019@ns.kogakuin.ac.jp

Saneyasu Yamaguchi
Dept. of Inf. and Com. Eng.
Kogakuin Univerity
Tokyo, Japan
sane@cc.kogakuin.ac.jp

KEYWORDS
TCP BBR, CUBIC TCP, TCP Fairness

1 INTRODUCTION
A new TCP congestion control algorithm, called TCP BBR, was proposed [1]. CUBIC TCP is widely used for this decade. In this paper, we evaluate the performances of TCP BBR and CUBIC TCP when they work concurrently in a deep buffer network, show their surprising long-term behaviors, and reveal the reason for the peculiar behaviors with insight into the queue behavior.

2 EVALUATION AND ANALYSES
We performed the iperf benchmarks using TCP BBR and CUBIC TCP concurrently. Each TCP established one connection. The connections shared the bottleneck link with a queue size of which was 16Ki or 64Ki packets. The one-way propagation delay time was 5 ms. All the windows sizes were 100 MiB.
Fig. 1 and 3 show the throughputs with the queue size 16KiB and 64 Kib, respectively. We can see that the performances behaved in the very long-term. The results in Fig. 3 shows the throughputs were severely unfair at the beginning and changed their behaviors to converging into fair-share at a time suddenly. Fig. 1 presences more surprising results. They repeated a cycle of unfair and converging and their behaviors changed suddenly.
Fig. 2 and 4 show the transition of the queue length at the bottleneck link. Fig. 2 indicates that the TCP BBR started increasing its windows after the CUBIC TCP stopped increasing its window size. Namely, a TCP BBR connection cannot increase its window size while another TCP is increasing its window size and queue length. This is the reason why throughputs suddenly go into the converging mode. Comparing Fig. 2 and 4, we can see that the queue length reached and did not reach the maximum queue length, which is the buffer size, respectively. In Fig. 2, the buffer reached full and some packets were lost, then the CUBIC TCP connection decreased its window size, i.e. the number of packets in the queue. This is the reason why two connections repeated the cycle. In the case of Fig. 4, the buffer was large enough and packets were not lost even when both TCPs reached the maximum size. Consequently, they converged into fair-share.

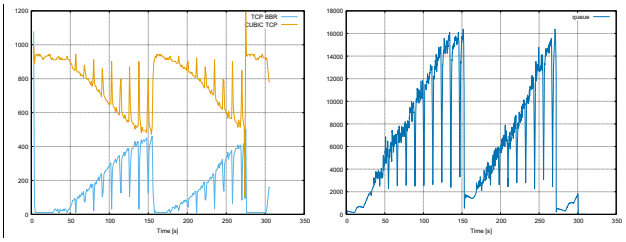

Figure 1: Throughput (queue size 16384)

Figure 2: Queue length (queue size 16384)

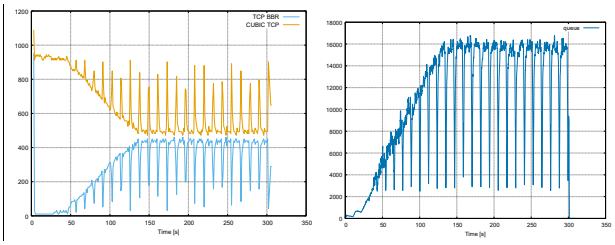

Figure 3: Throughput (queue size 65536)

Figure 4: Queue length (queue size 65536)

3 CONCLUSION
In this paper, we evaluated the behaviors of TCP BBR and CUBIC TCP and showed their peculiar behaviors. We then revealed the behavior of the queue and showed the reason for the behaviors.

REFERENCE
[1] Neal Cardwell, Yuchung Cheng, C. Stephen Gunn, Soheil Hassas Yeganeh, and Van Jacobson, “BBR: Congestion-Based Congestion Control,” *Queue* 14, 5, pages 50 (October 2016), 34 pages, 2016. DOI: https://doi.org/10.1145/3012426.3022184

ACKNOWLEDGMENTS
This work was supported by JSPS KAKENHI Grant Numbers 17K00109, 18K11277.This work was supported by JST CREST Grant Number JPMJCR1503, Japan.