BBR’s Sharing Behavior with CUBIC and Reno

Validating Influential Models in Updated Settings

Fatih Berkay Sarpkaya, Ashutosh Srivastava, Fraida Fund, Shivendra Panwar

IFIP Networking 2025

NYU TANDON SCHOOL OF ENGINEERING

2016
BBRv1 is introduced

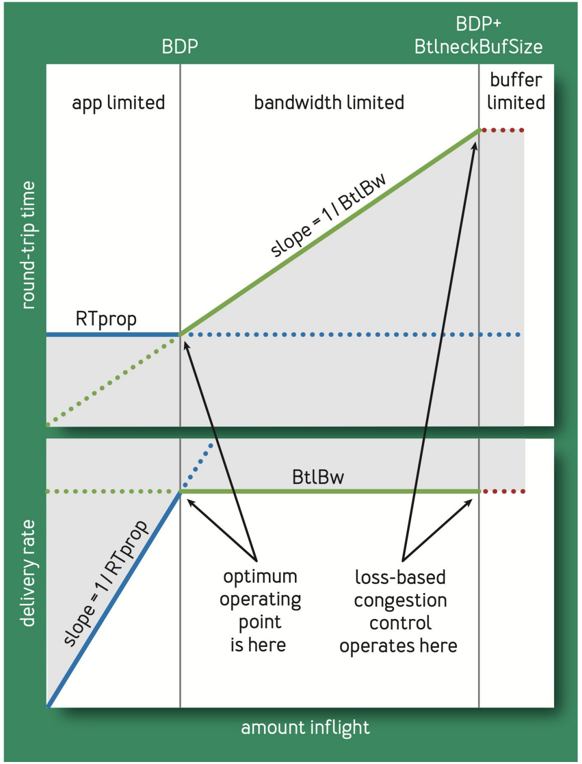

BBR: Bottleneck Bandwidth and Round-trip propagation time

BDP (Bandwidth-delay product):
Max in-flight data before queuing = Bandwidth × RTT

Source: Neal Cardwell, Yuchung Cheng, C. Stephen Gunn, Soheil Hassas Yeganeh, and Van Jacobson. 2016. BBR: Congestion-Based Congestion Control: Measuring bottleneck bandwidth and round-trip propagation time. Queue 14, 5 (September-October 2016), 20–53. https://doi.org/10.1145/3012426.3022184

2016
BBRv1 is introduced

2017-2018
Hock et al. & Scholz et al. note fairness issues

- unfairness to flows with loss-based congestion control
- increased queuing delays
- RTT unfairness within large buffers
- a massive amount of packet losses

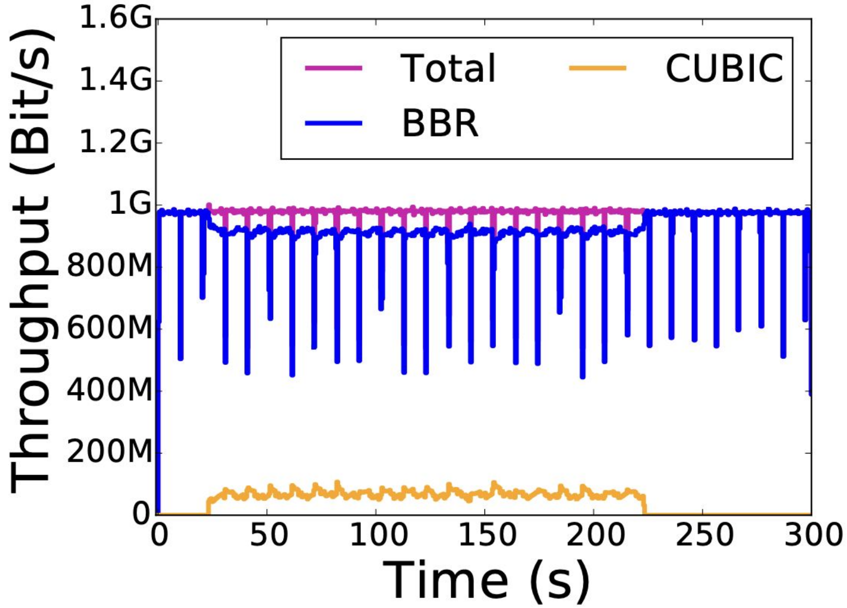

Source: M. Hock, R. Bless and M. Zitterbart, "Experimental evaluation of BBR congestion control," 2017 IEEE 25th International Conference on Network Protocols (ICNP), Toronto, ON, Canada, 2017, pp. 1-10, doi: 10.1109/ICNP.2017.8117540. keywords: {Delays;Bandwidth;Packet loss;Linux;Throughput;Internet},

2016	2017-2018	2019
BBRv1 is introduced	Hock et al. & Scholz et al. note fairness issues	Steady-state model by Ware et al.

- ***Modeling BBR's Interactions with Loss-Based Congestion Control***
  IMC 2019 – R. Ware, M.K. Mukerjee, S. Seshan, J. Sherry
- Key insight: In competition, BBRv1 becomes window-limited and sends at its in-flight cap, aggressively pushing out loss-based flows.

$$BBR_{frac} = \left(1 - \frac{1}{2} + \frac{1}{2X} + \frac{4N}{q}\right) \times \left(1 - \left(\frac{q}{c} + .2 + l\right) \times \frac{1}{10}\right)$$

X → buffer size (BDP) & q → buffer size (packets) & l → propagation delay (seconds)
N → # of BBR flows & c → bottleneck link capacity (packets / seconds)

2016
BBRv1 is introduced

2017-2018
Hock et al. & Scholz et al. note fairness issues

2019
Steady-state model by Ware et al.

2019
BBRv2 is introduced

BBR v2 improvements: coexistence with Reno/CUBIC

- BBR v1: low throughput for Reno/CUBIC flows sharing some paths
- BBR v2: adapts bandwidth probing for better coexistence with Reno/CUBIC [IETF 102]

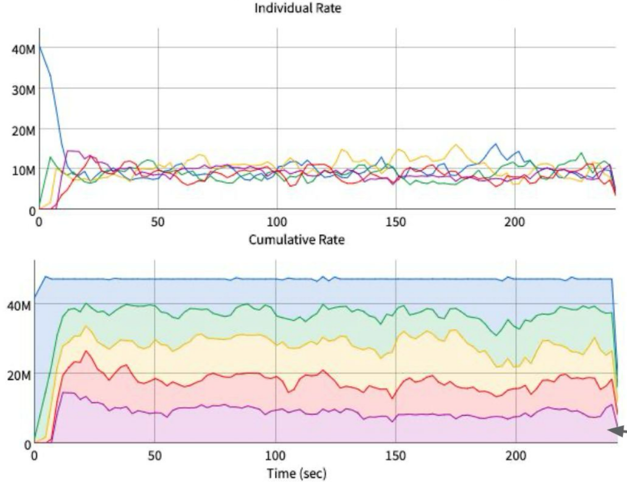

4 CUBIC, 1 BBR v2, 50M, 40ms,
buffer = 1xBDP
start time {0, 2, 4, 6, 8} secs

	bw	retrans
CUBIC 1	10.5 M	0.3%
CUBIC 2	9.1 M	0.1%
CUBIC 3	10.4 M	0.1%
CUBIC 4	8.7 M	0.1%
BBR v2	9.3 M	0.1%

BBRv2
A Model-based Congestion Control

IETF 104 & ICCRG session
Prague, Mar 2019

Background

2016
BBRv1 is introduced

2017-2018
Hock et al. & Scholz et al. note fairness issues

2019
Steady-state model by Ware et al.

2019
BBRv2 is introduced

2022
Fluid model by Scherrer et al.

- ***Model-Based Insights on the Performance, Fairness, and Stability of BBR***
  IMC 2022 – S. Scherrer, M. Legner, A. Perrig, S. Schmid

- Captures BBRv1 and v2's rate evolution using differential equations.

- In the model, the sending rate is formulated as follows, accounting for the switching between the ProbeRTT and ProbeBW phases:

$$x_i = m_i^{\mathrm{prt}} \cdot \frac{w_i^{\mathrm{prt}}}{\tau_i} - \left(1 - m_i^{\mathrm{prt}}\right) \cdot x_i^{\mathrm{pbw}}$$

2016
BBRv1 is introduced

2017-2018
Hock et al. & Scholz et al. note fairness issues

2019
Steady-state model by Ware et al.

2019
BBRv2 is introduced

2022
Fluid model by Scherrer et al.

2023
BBRv3 is introduced

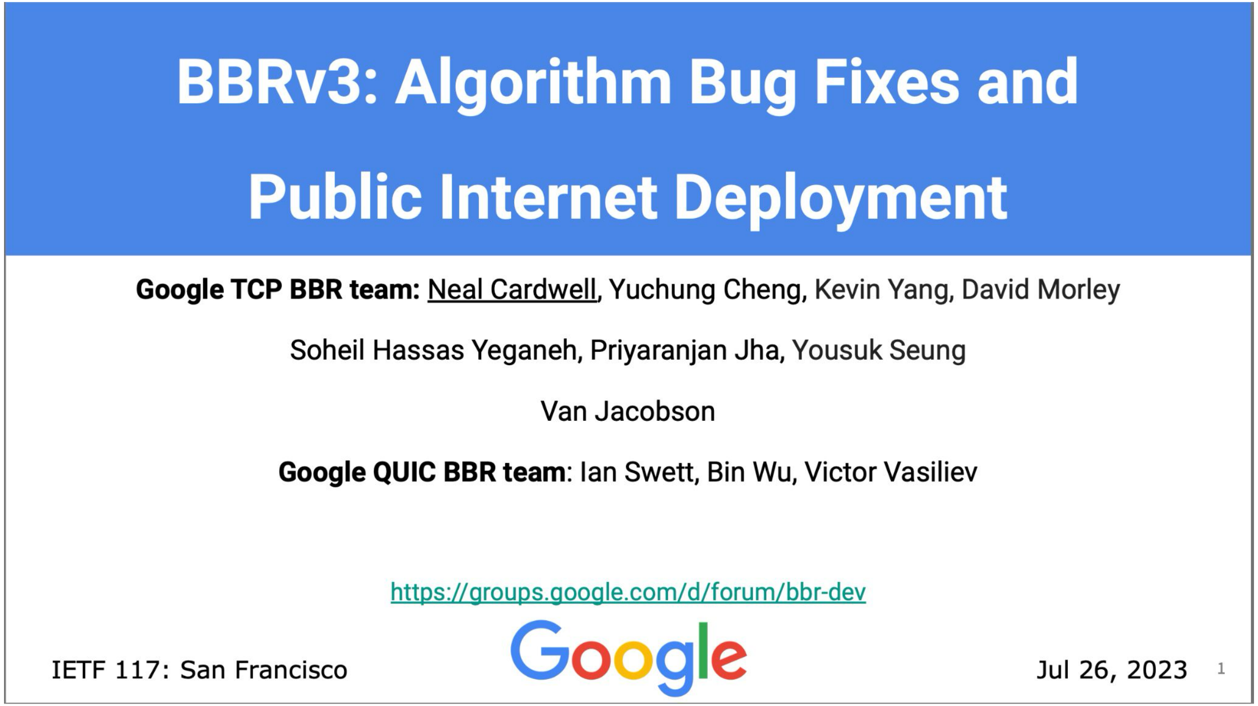

BBRv3:

BBRv2 + bug fixes and performance tuning

2016
BBRv1 is introduced

2017-2018
Hock et al. & Scholz et al. note fairness issues

2019
Steady-state model by Ware et al.

2019
BBRv2 is introduced

2022
Fluid model by Scherrer et al.

2023
BBRv3 is introduced

2024
Extent of BBR deployment

Websites	Traffic share [56]	CCA
google domains	13.85%	BBRv3
netflix.com	13.74%	Reno
facebook.com	6.45%	CUBIC
apple.com	4.59%	AkamaiCC
disneyplus.com	4.49%	CUBIC
amazon.com	4.24%	BBRv1
tiktok.com	3.93%	AkamaiCC
primevideo.com	2.67%	BBRv2
hulu.com	2.44%	AkamaiCC

Source: Ayush Mishra, Lakshay Rastogi, Raj Joshi, and Ben Leong. 2024. Keeping an Eye on Congestion Control in the Wild with Nebby. (ACM SIGCOMM '24). &
https://datatracker.ietf.org/meeting/121/materials/slides-121-iccrg-promises-and-potential-of-bbrv3-00.pdf

2016
BBRv1 is introduced

2017-2018
Hock et al. & Scholz et al. note fairness issues

2019
Steady-state model by Ware et al.

2019
BBRv2 is introduced

2022
Fluid model by Scherrer et al.

2023
BBRv3 is introduced

2024
Extent of BBR deployment

2025
Do we need new models?

Validation Scenarios

Network Settings	Steady-State Model (2019)	Fluid Model (2022)
Multiple BBR vs 1 loss-based flow	✅	
Multiple BBR vs multiple loss-based flows		✅
Shallow - moderate buffer size	✅	✅
Deep buffer scenarios (> 7BDP)	✅	
BBR version	✅ BBRv1 Model developed before BBRv2 and BBRv3 release	✅ BBRv1, ✅ BBRv2 Model developed before BBRv3 release

Research Questions

- Which model better predicts BBR's behavior when competing with a single loss-based flow?
- How does the fluid model perform in networks with very large buffers?
- To what extent do either of the existing models (developed for BBRv1/v2) capture the behavior of BBRv3? Do we need new models to describe BBRv3 behavior?

Experiment Setup

- We systematically replicate and extend the steady-state and fluid model evaluations of BBR.

a) BBR vs 1 loss-based flow

- iperf3 flows for 1000 seconds (last 200 seconds reported)

b) BBR vs multiple BBR or multiple loss-based flows

- iperf3 flows for 300 seconds (last 120 seconds reported)

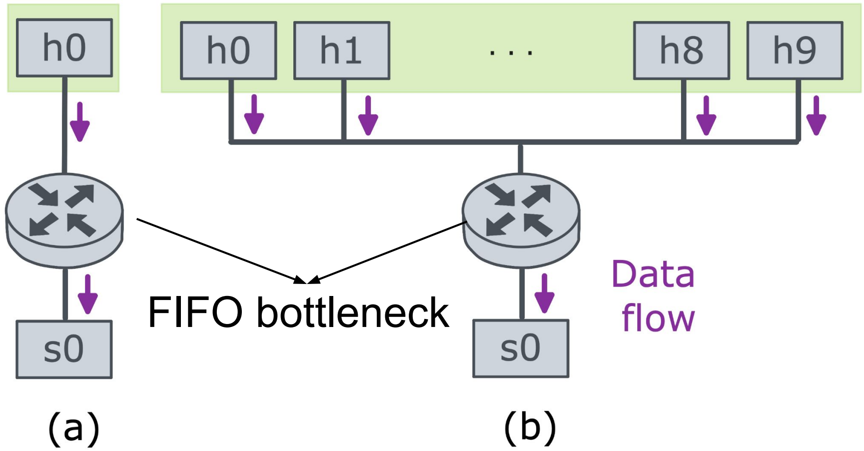

RTT: 10 or 40 ms
Capacity: 40 or 50 Mbps

RTT: Between 30-40 ms
Capacity: 100 Mbps

FABRIC

For BBRv1, Steady-State Works, Fluid Model Struggles (against a single loss-based flow)

- Steady-state model accurately predicts BBRv1 behavior. (MSE: 0.006)
- The fluid model is less effective with a large number of flows and deep buffers. (MSE: 0.058)

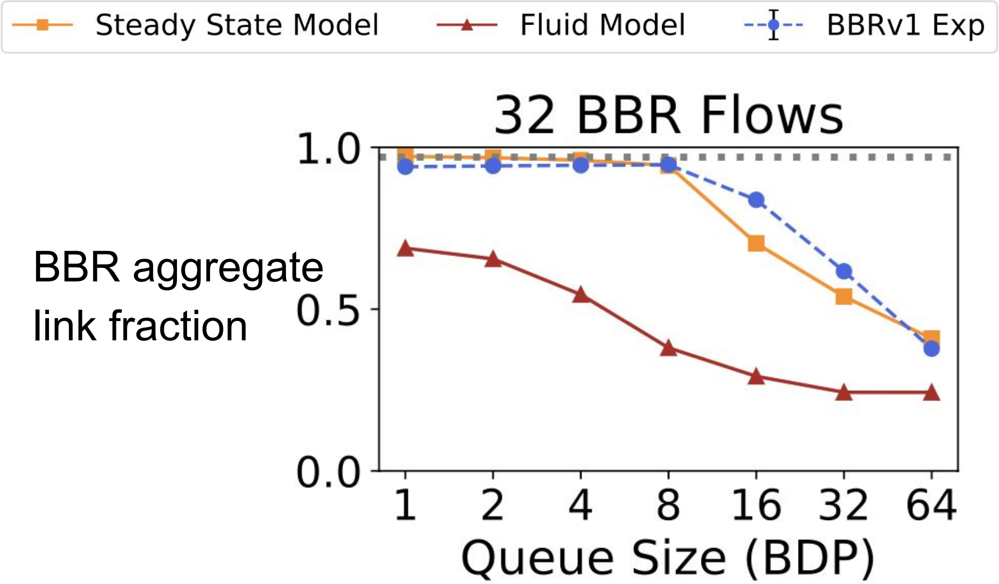

40ms RTT × 10 Mbps
BBRv1 vs 1 CUBIC

The fluid model partially captures BBRv2's behavior (against a single loss-based flow)

- The fluid model successfully captures the reduced aggressiveness of BBRv2 compared to BBRv1.
- However, it overestimates BBRv2's link share against CUBIC and underestimates it against Reno.

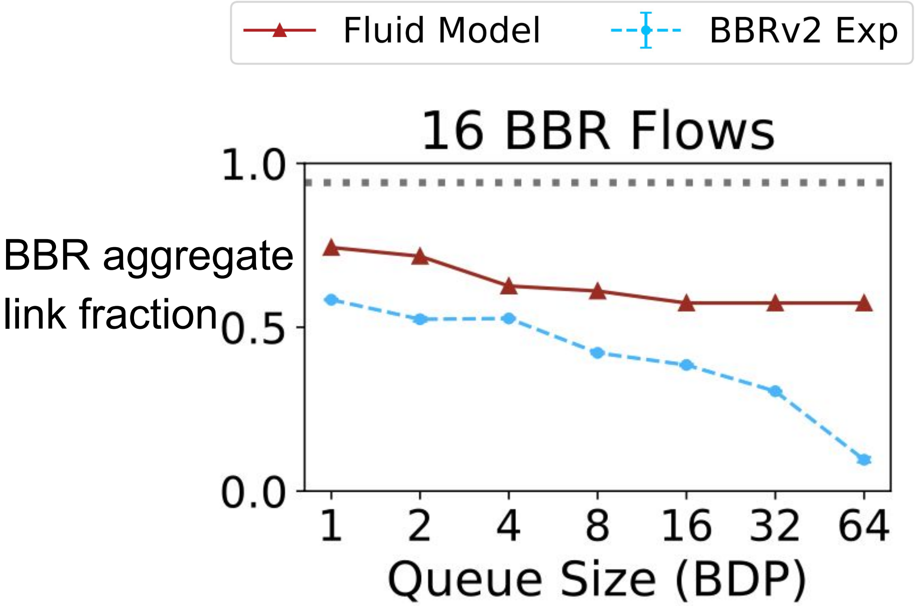

30ms RTT × 50 Mbps
BBRv2 vs 1 CUBIC

The fluid model partially captures BBR behavior (against multiple flows)

- The fluid model aligns with our results for loss and inter-flow fairness for BBRv1 and BBRv2.
- It becomes less effective at capturing critical behaviors - particularly intra-flow fairness, buffer occupancy, and utilization - as buffer sizes increase.

Neither model fully captures BBRv3 behavior (against a single loss-based flow)

- BBRv3 strikes a balance between the aggressiveness levels of BBRv1 and BBRv2 for shallow buffers with a few flows.

- Model predictions aligns with this observation.

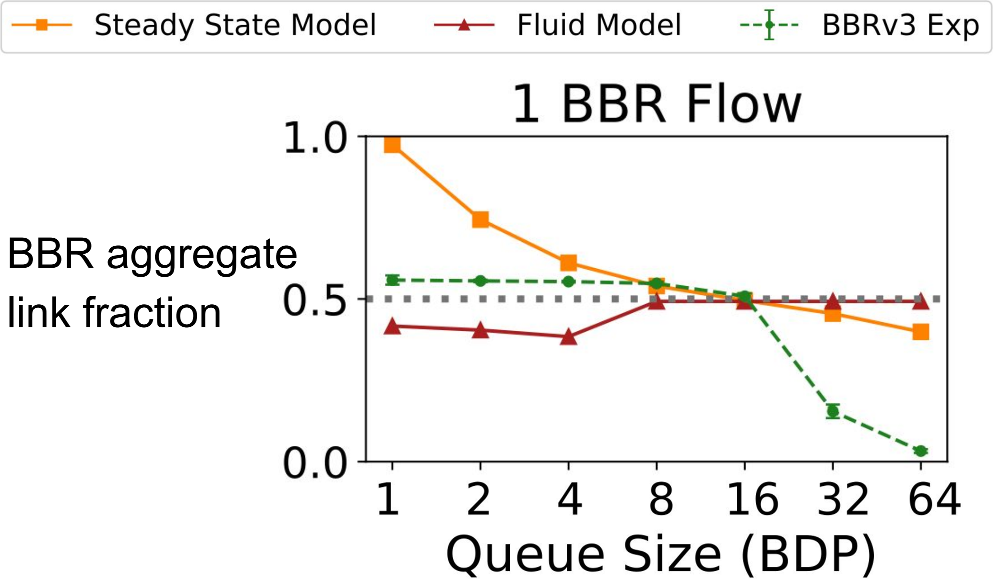

30ms RTT × 50 Mbps
BBRv3 vs 1 CUBIC

Neither model fully captures BBRv3 behavior (against a single loss-based flow)

- Neither steady-state model nor fluid model is designed for BBRv3.
- However, the Steady-state model of BBRv1 is a reasonable predictor when many flows compete in shallow buffers.

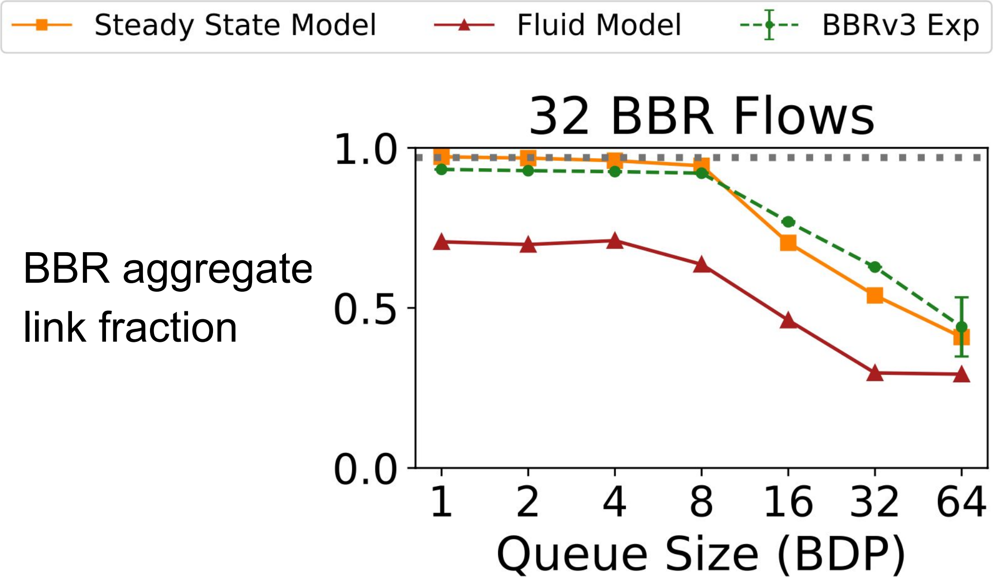

30ms RTT × 50 Mbps

BBRv3 vs 1 CUBIC

Neither model fully captures BBRv3 behavior (against multiple flows)

- While fluid models for BBRv1 and BBRv2 partially explain BBRv3's shallow buffer behavior, they are not a good predictor in deep buffers—especially for intra- flow fairness, utilization, and buffer occupancy.

- Existing models do not fully capture BBRv3's performance, indicating the need for a new model.

Existing models partially capture BBR behavior, but updates needed for BBRv3

- BBRv1 Steady-state model is more accurate than the fluid model at predicting link share for BBR flows against a single loss based flow.
- BBRv1/v2 fluid models works well for shallow/moderate buffers but has limitations with deep buffers or many flows.
- Existing theoretical models may require updates—or entirely new modeling approaches—to accurately represent BBRv3's behavior.

BBR's Sharing Behavior with CUBIC and Reno

Fatih Berkay Sarpkaya, Ashutosh Srivastava, Fraida Fund, Shivendra Panwar

Thank You for Listening !

Experiment Artifacts:
https://github.com/fatihsarpkaya/bbr-shared-bottleneck

Contact: fbs6417@nyu.edu