Journal of Computing & Biomedical Informatics
ISSN: 2710 - 1606

Volume 09 Issue 02
2025

Research Article
https://doi.org/10.56979/902/2025

BBR v3-Cubic smackdown: A Fairness and Convergence Quantitative Evaluation

Muhammad Ahsan[1*], Sadia Abbas Shah[1], Farrukh Nadeem[1], and Raybal Akhtar[2]

[1]Department of Software Engineering, University of Management and Technology, Lahore, Punjab, Pakistan.
[2]Department of Informatics and Systems, University of Management and Technology, Lahore, Punjab, Pakistan.
[*]Corresponding author: Muhammad Ahsan. Email: mahsan@alumni.usc.edu

Received: July 03, 2025 **Accepted:** August 28, 2025

Abstract: Congestion control has been an important component of TCP for nearly four decades now. Bottleneck Bandwidth and Round-Trip Time (BBR) has brought a major change in the ways the congestion onset can be monitored, and proactive measures can be taken instead of a reactive technique being used in loss-based traditional algorithms, such as Cubic and Reno, for a long time. With the introduction of the latest iteration of BBR, BBR v3, it has been claimed that it has improved its co-existence with Cubic flows with better fairness, but convergence issues within BBR-v3 have been reported. No existing study quantifies BBR v3's fairness with Cubic, and the convergence of BBR v3 for intra-protocol streams needs to be quantified for each stream's throughput as well. The simulations and emulations generally don't bring the true picture of the performance of a congestion control algorithm. In this paper, we have evaluated BBR v3 with cubic using our real-time physical testbed using Jain's Fairness Index to bring a more accurate fairness analysis of BBR v3 and Cubic streams. For convergence, a well-established statistical metric that measures the relative stability of throughput, known as the Coefficient of Variation (CoV), has been calculated for BBR v3/Cubic flows. We used Flent (a FLExible Network Tester) to perform rigorous tests using various pairs of streams in upload, and the results, along with the metadata, have been saved for reproducibility and validation. Our thorough testing on both wired (Ethernet) and wireless (Wi-Fi 4) testbeds confirms that fairness issues between BBR v3 and Cubic streams persist. These issues are particularly serious as the number of streams increases. Similarly, our convergence tests confirm that BBR v3 flows, especially the first stream, obtain a larger share of the bandwidth, and the throughput of each stream remains highly volatile.

Keywords: BBR; Congestion Control; Convergence; Fairness; CoV; Cubic; JFi

1. Introduction

Congestion control has been an important part of TCP for the last four decades. Active work on it is going on in the scientific community with a common goal of mitigating congestion to ensure the network pipes are not overloaded with data that results in packet loss and delays due to piled-up buffers in the switches and routers. For many years, loss-based congestion control algorithms (CCAs) have been the popular choice, with Cubic [1] leading from the front. Then came Google's Bottleneck Bandwidth and Round-trip time (BBR) [1], its working principle is totally different from loss-based, as it is model-based. It tries to use the golden rule of Kleinrock [2] to work at its defined optimal operating point and tries to keep the pipe just full but no fuller [3]. Revisions were made in BBR, and now we have its latest iteration, BBR v3 [4]. Cubic, which was introduced a lot earlier, has already made ground, and many clients and server machines are already using it. It is the default congestion control in the mainline Linux kernel and in Microsoft Windows 10/11 as well. Making it the most widely deployed congestion control in the current era of the internet. So, BBR v3 flows must be fair to any Cubic flows going over the same path. It should yield a fair bandwidth share among its own and Cubic streams. The early versions of BBR were reported with this fairness issue, that BBR is not giving the Cubic flows their fair share, and consequently, revisions were made in the BBR that resulted in the latest version of BBR v3 in July 2023.

Along with this issue of fairness with Cubic, it has been observed that it also faces an issue of convergence with its own streams. When multiple BBR streams in upload are sent from a server, their convergence is very poor. The situation gets worse when staggered streams (streams with a delay between them) are sent. The first stream takes a lot more bandwidth than the other streams, and all these streams don't converge at all. Convergence is important as it tells how much time it takes for streams to settle down and come to a stable point where they share fair bandwidth. The high oscillations in stream throughputs and never settling down to an equitable share create not only unnecessary delays, but the spikes may cause excessive queuing delays.

In this paper, we have investigated the issue of fairness with Cubic and intra convergence issues of BBR v3 streams. We have quantitatively evaluated fairness using the Jain Fairness Index (JFi) [5] and used the statistical measure Coefficient of Variation (CoV) [6] to quantify the convergence of each BBR v3 stream. To the best of our knowledge, there is no work done in evaluating BBR v3 with respect to these two metrics except [7], in which Danesh et al. have shared their findings for BBR v3 using a simulated testbed. In our case, it is a state-of-the-art, atomic testbed for wired Ethernet and Wireless. We have used Flent [8] to get the results. The choice of the tool comes with its versatility and flexibility in performing network tests smoothly, and also provides reproducibility and validation of the tests with the inclusion of metadata in the output compressed JavaScript Object Notation (JSON) file.

The rest of the paper is structured with Section 2 providing the details on related works done in this regard. Section 3 elaborates on the methodology involved to get the results. The real-time physical testbed details are also presented in this section. Section 4 focuses on the results and discussion, and Section 5 gives the conclusion about our overall findings and the way forward.

2. Related Works

To the best of our knowledge, there is a gap in existing literature regarding the quantitative evaluation of BBR v3, and if we include a real testbed like the one in this study, then it becomes a totally novel assessment of BBR v3.

Danesh et al. [7] in their work evaluated BBR v3 in a simulated environment and found that BBR v3 still has fairness issues, although it is better than the previous versions of BBR. It is right now the only study that evaluated BBR v3, but it lacks tests on different network scenarios, such as a wireless testbed. As the wireless networks behave totally differently from the wired networks, and often have more packet losses and comparatively higher latency due to multipath interference and signal attenuation issues. Their work focuses more on identifying fairness and convergence issues in BBR v3, but does not provide any solution for it.

Emilia et al. [9] showed that BBR v3, when enabled with explicit congestion notification (ECN), severely chokes the other Cubic flows. They explored the impact of ECN enabling on bandwidth sharing, addressing modern network conditions and strong practical relevance for Internet traffic mix scenarios. This is an important finding, but they didn't provide any solution to address this fairness.

Jose et al. [10] evaluated BBR v3 in wired broadband network scenarios. Their testing was through simulations, and they showed that BBR v3, although better than BBR v1 and BBR v2 in terms of fairness with Cubic, still exhibits fairness issues. They evaluated BBR v3 using various metrics such as propagation delays, loss rates, number of flows, and different buffer sizes. Active queue management (AQM) was also used to control RTT unfairness. Their work focuses mainly on wired broadband networks, providing limited insights into wireless or mixed network conditions. While highlighting fairness improvements with multiple BBR v3 flows, solutions to fairness challenges lack detailed exploration or a proposal.

Another work by Fateh et al. [11] is about BBR v3 sharing with Cubic as well as Reno [12]. Their paper systematically evaluates existing theoretical models for TCP BBR congestion control against experimental data, including the newly introduced BBR v3, revealing significant model limitations and the need for updated theoretical frameworks. It provides a systematic comparison between steady-state and fluid models across different BBR versions. They used the FABRIC [13] national-scale programmable testbed, ensuring reproducibility and enabling future research validation. Their study focuses exclusively on bulk-transfer TCP traffic, which may not represent realistic application behaviors. Also, it lacks empirical verification with real or simulated traces.

Our work provides a quantitative analysis of fairness issues of BBR v3 and cubic using the Jain Fairness Index (JFi) and the intra convergence issues of BBR v3 streams using the statistical metric, coefficient of variation (CoV). The tests are performed via Flent for different time durations with a step size of 0.1. Our real-time physical testbed for both wired and wireless setups is better, as only simulation results may not fully represent real-world network conditions.

3. Methodology

This section outlines the physical testbeds employed for real-time quantitative evaluation of BBR v3 across both wired and wireless environments. The client system runs Linux Ubuntu 21, while the server operates on Debian 12 (Bookworm). The client is specially configured with bbr3, compiled from its latest GitHub branch [14], and is loaded as a kernel module for Linux kernel version 6.13.7 on demand. It uses a custom-compiled kernel 6.13.7 on the Ubuntu 21 client machine. The wired connection is provided via a Gigabit RealTek USB-based Ethernet controller, while the wireless connection utilizes a Qualcomm QCA9377 PCIe adapter supporting Wi-Fi 4 and 5. The wireless router is a Huawei EchoLife, a Gigabit Passive Optical Network (GPON) terminal functioning as both an Optical Network Terminal (ONT) and a Wi-Fi access point.

For the wired setup, the server is powered by an AMD Ryzen 7 machine running Linux kernel version 6.1.0-17-amd64. It hosts a Netperf 3 server connected via a Realtek RTL8411 PCI Express Gigabit Ethernet controller, with the server process listening on port 50,000 Fig 1. On the client side, the generic segmentation offload (GSO) quantum for BBR is initially set to 2 Maximum Segment Size (MSS) to control the maximum data aggregate passed to TCP small queues (TSQ), which is configured with an 8 ms timeout aligned with a pacing shift of 7 in Linux 6.x kernels [15]. These testbeds, illustrated in Figs 1 and 2, reflect typical network configurations commonly found in home and office environments, where clients may be connected either wirelessly or through wired links, while servers are predominantly wired. These setups serve to assess the convergence and fairness for Cubic and BBR v3 CCAs.

All experiments utilize the Flent testing tool [16], which automates tests with various traffic flows while measuring throughput and latency. Flent also gathers key metadata to ensure result authenticity and reproducibility. Tests conducted include TCP Upload with varying numbers of streams (1, 2, 4, 8, 12, and 16) over both wired and wireless testbeds. Test results with eight TCP streams are shared in this paper for the sake of brevity. The rest of the test details and scripts can be sought from our online repository [17].

Image Source: Physical Testbed, "Client, Wired Router, and Server icons" are imported from https://www.pngegg.com/, which is a free hosting site. CC BY 4.0.

Image Source: Physical Testbed, "Client, Wireless Router, and Server icons" are imported from https://www.pngegg.com/, which is a free hosting site. CC BY 4.0

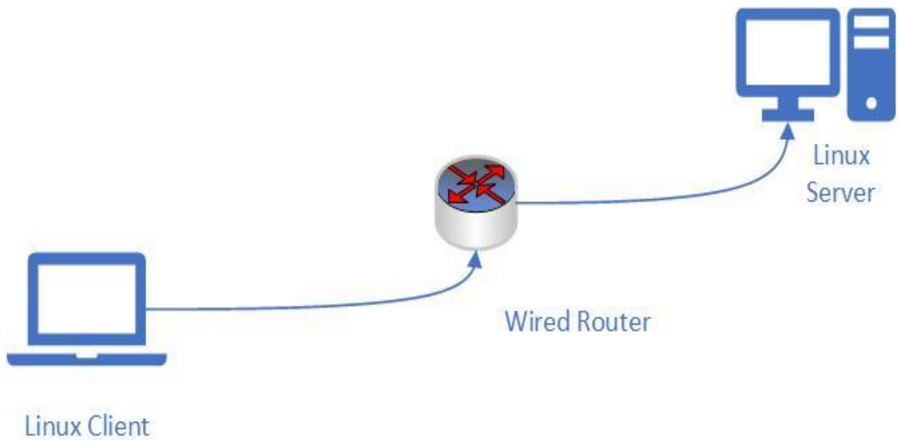

Figure 1. Linux-Ubuntu client via a wired router to a Linux Debian server

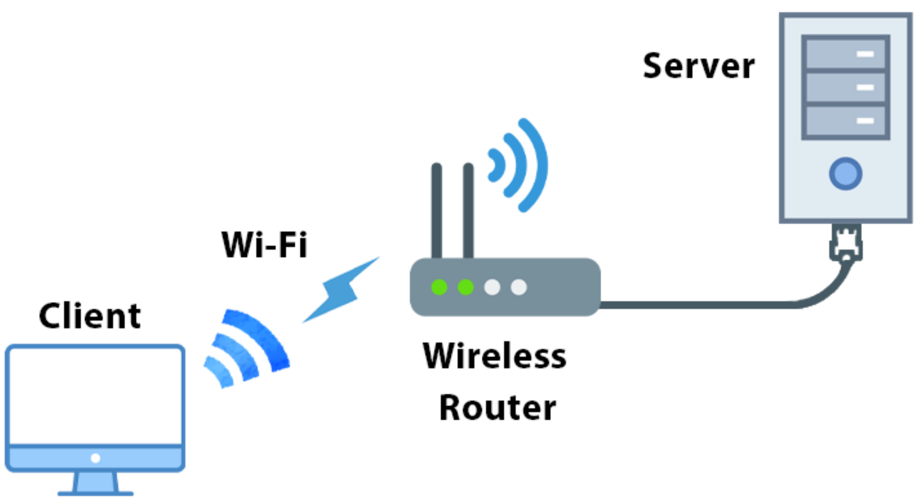

Figure 2. Linux-Ubuntu wireless client via Wi-Fi 4/5 router to a wired Linux Debian server

4. Results and discussion

In this section results for tests performed on the physical testbed are discussed. It is divided into three sub-sections. First and second sub-section discusses BBR v3 convergence test results with staggered streams for 200 and 300 seconds durations using the wired testbed. The fairness test and the resulting JFi have been discussed and calculated in the third sub-section.

4.1. BBR-v3 convergence testing with staggered streams in upload in wired (Scenario for a 200-second duration)

"Table 1." The COV for all streams is consistently high, ranging from 0.76 to 0.90. This confirms that the streams did not converge to a stable, low-variability state during the 200-second test period. The standard deviation for all streams is very high, indicating significant fluctuations in throughput.

The mean throughput for each stream generally decreases as more streams are added. Stream 1, which starts first, has the highest average throughput, while subsequent streams have lower averages due to increased competition for network resources.

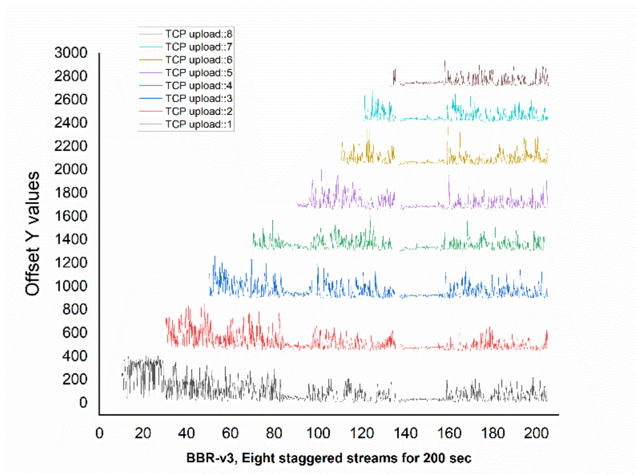

Figure 3. Eight BBR v3 staggered streams in upload

Table 1. BBR v3 convergence statistics for 200 sec

Stream	Start Time (s)	Mean (μ) (Mb/s)	Standard Deviation (σ) (Mb/s)	COV
Stream 1	10.1	100.41	90.61	0.9
Stream 2	30.2	80.85	70.09	0.87
Stream 3	50.2	71.85	58.71	0.82
Stream 4	70.3	56.79	43.25	0.76
Stream 5	90.4	62.71	50.86	0.81
Stream 6	110.8	56.98	46.56	0.82
Stream 7	120.8	52.51	43.67	0.83
Stream 8	132.7	47.46	36.87	0.78

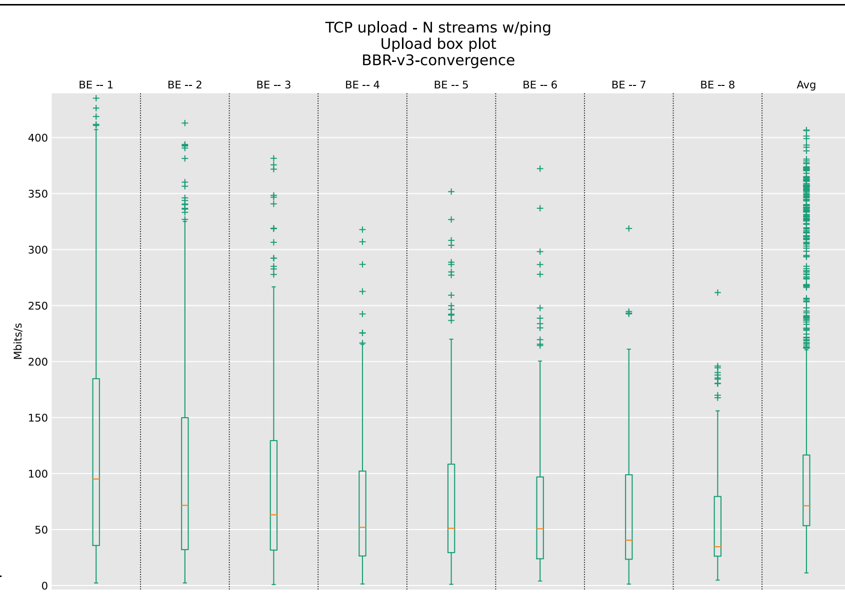

Figure 4. Eight streams median 200 seconds duration

4.2. BBR-v3 convergence testing with staggered streams in upload in wired (Scenario for a 300-second duration)

To further test the BBR-v3 convergence, eight streams were sent in upload at different times 5, 25,45,65,85,105,115,125 seconds, for a total test duration of 300 seconds at a step duration of 0.1 seconds. Fig 5 shows the eight staggered streams graphically with added Offset Y values so that they can be viewed easily.

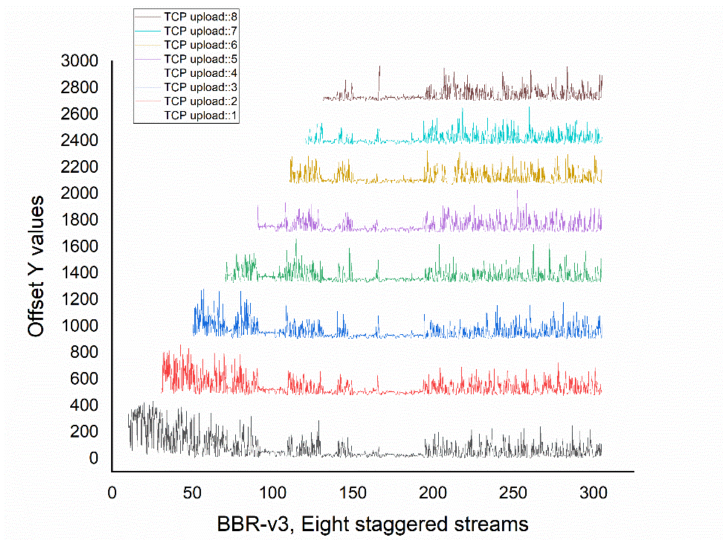

Figure 5. Eight BBR v3 staggered streams in upload

"Table 2" presents three key metrics for each stream: the mean, standard deviation, and Coefficient of Variation (COV). Analyzing these values provides insights into how the network throughput is being shared and how the streams are behaving over time. COV=σ/μ is a unitless measure of the variability of the data (σ) relative to its mean (μ). It is the most important metric for assessing convergence. A COV approaching zero would indicate a highly stable and converged stream. The COV values are relatively high for all streams, ranging from 0.37 to 0.44. The values shown in "Table 2" are not trending towards zero, showing the BBR-v3 streams are not converging to an equitable bandwidth share within 300 seconds.

Table 2. BBR v3 convergence statistics for 300 se

Stream	Start Time (s)	Mean (μ) (Mb/s)	Standard Deviation (σ) (Mb/s)	COV
Stream 1	5	82	36.08	0.44
Stream 2	25	69.52	29.2	0.42
Stream 3	45	63.2	24.02	0.38
Stream 4	65	55.65	20.69	0.37
Stream 5	85	56.11	21.32	0.38

Stream 6	105	53.38	20.29	0.38
Stream 7	115	50.52	20.2	0.4
Stream 8	125	49.06	19.62	0.4

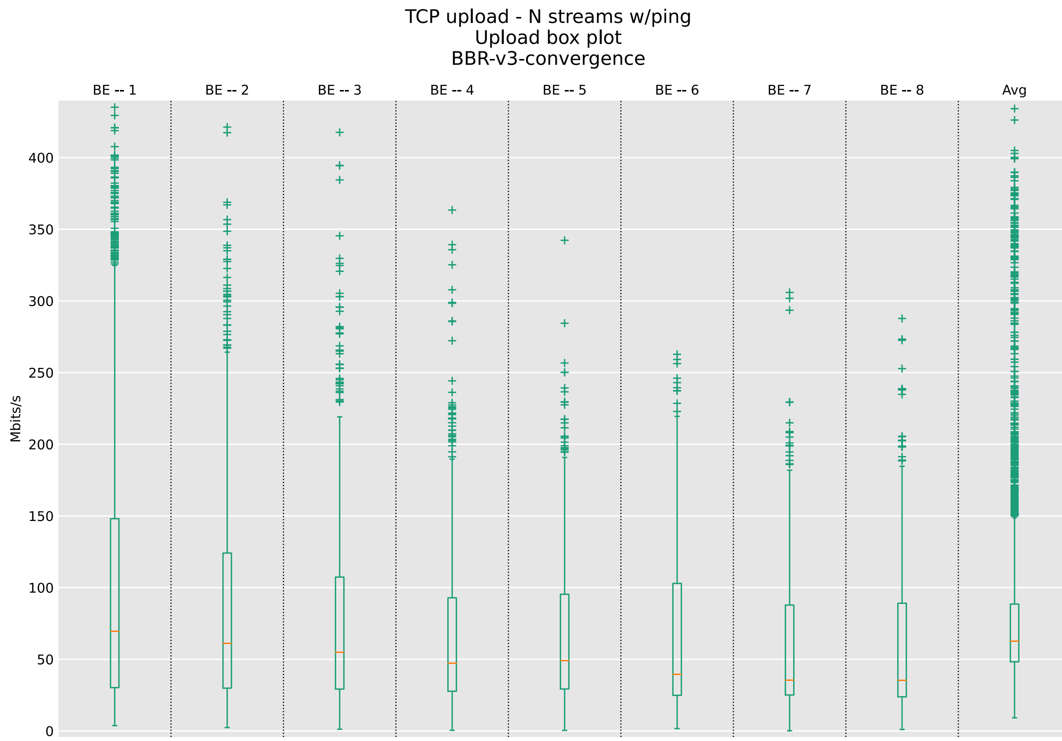

Figure 6. Eight streams' medians for a 300-second duration

The mean throughput consistently decreases as more streams are added. Stream 1 has the highest mean at 72.00 Mb/s, while Stream 8 has the lowest at 45 Mb/s. This is the expected behavior of a shared network link. As each new TCP stream begins, it competes for the available bandwidth, causing the average throughput for all streams to be reduced. The initial streams capture a larger share of the bandwidth because there is less competition.

4.3. BBR-v3 Cubic JFI in a wireless scenario

To evaluate fairness between BBR v3 and Cubic streams, we used Jain's fairness Index. Jain's Fairness Index is a widely used metric to assess the fairness of resource distribution among multiple data flows, assuming each flow has identical data rate requirements. The instantaneous fairness index JFi is defined using the data rates R_i as follows:

$$\mathrm{JFi}\big(R_1, R_{2,} \dots \dots \dots R_n\big) = \frac{\left(\sum_{i=1}^{n} R_i\right)^2}{n.\sum_{i=1}^{n} R_i^2}$$

Where n is the number of active flows, Ri is the instantaneous date rate of flow i and JFi is a real number in the interval [1/n, 1] with a best-case value of 1 if the data rate is qual for all flows, i.e. the available bandwidth has been fairly shared, and a minimum case of 1/n, if only one aggressive flow is monopolizing the available bandwidth. Fig 7 shows the combined Jain's fairness Index plot for BBR v3 and Cubic streams (2, 3, 4, 8, 10, 12, 15) each in upload directions together. The JFi index was initially high when only 2 and 3 streams each of BBR v3 and Cubic were sent using the TCP upload test. It was 0.88 and 0.8, but when the streams increased to 4, 8, 10, 12, and 15, we saw the JFi started to drop and was only 0.64 for 15 streams of BBR v3 and 15 streams of Cubic upload test. This is a clear indication that BBR v3 still has fairness issues with Cubic when the number of flows is increased.

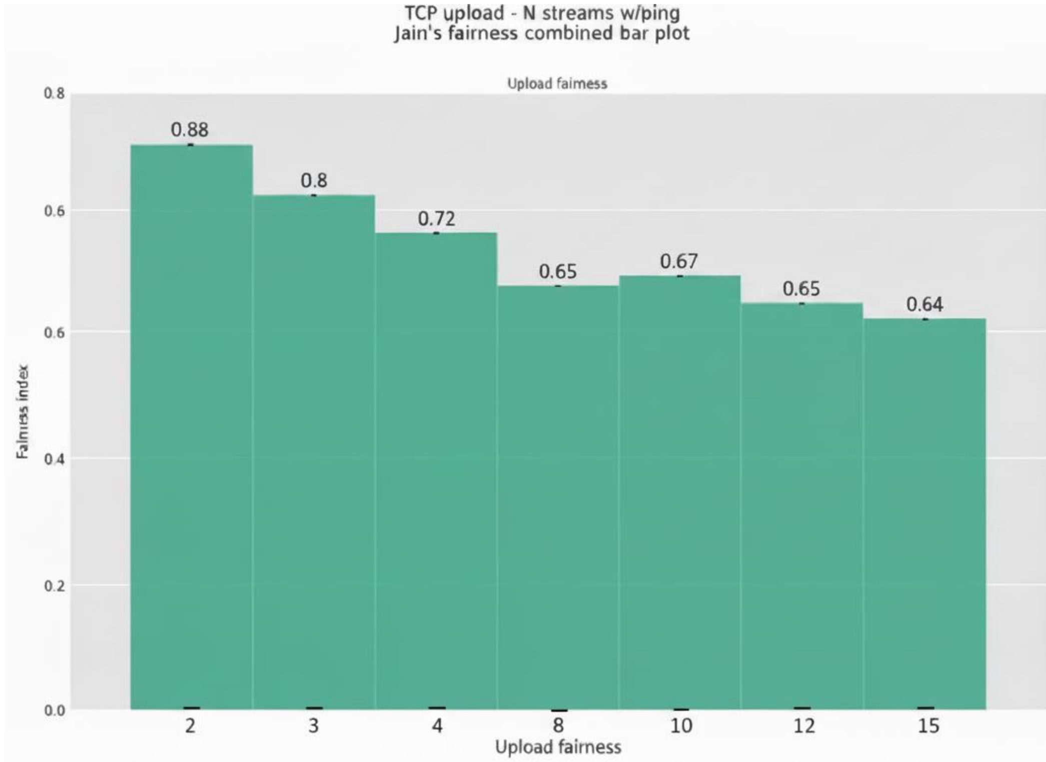

Figure 7. Jain's fairness Index for various BBR-v3, Cubic streams

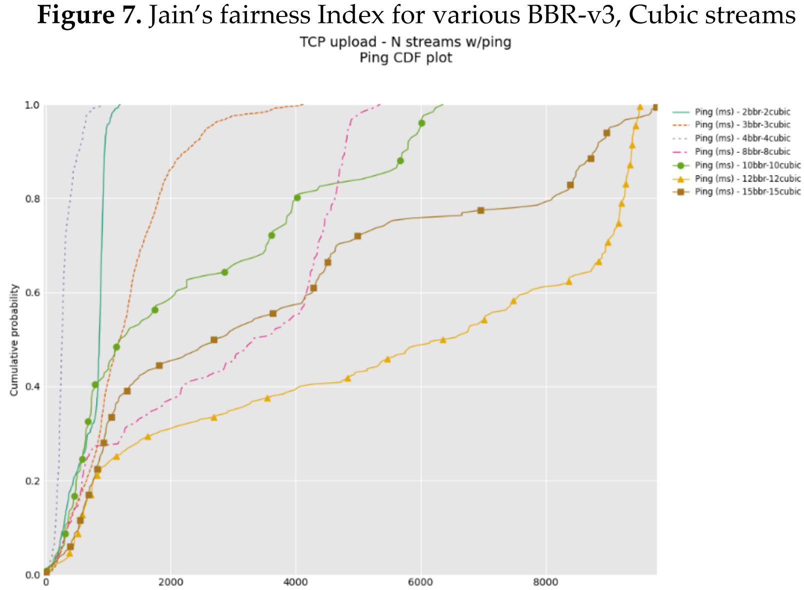

Figure 8. Ping CDF for various BBR-v3/Cubic streams

Fig 8 presents the cumulative distribution function (CDF) of ping latency for various configurations of BBR v3 and Cubic TCP upload stream pairs, ranging from 2 to 15 streams each. The results reveal a clear trend of increasing latency and variability as the number of concurrent streams grows. Lower stream counts (e.g., 2bbr-2cubic) exhibit tight latency distributions with minimal queuing delay, indicating fair coexistence. In contrast, higher stream counts (e.g., 15bbr-15cubic) show significantly increased ping times and broader distributions, suggesting increased queuing and potential unfairness.

These findings highlight that BBR v3's aggressive bandwidth probing may dominate Cubic under high concurrency, leading to degraded latency performance and fairness issues in mixed deployments.

5. Conclusion

In this paper, we have provided a quantitative analysis of BBR v3 fairness and convergence. The fairness is evaluated with the popular loss-based congestion control Cubic, and convergence is primarily investigated with BBR v3 itself. Metrics such as JFi and CoV were used to quantify the findings of our results. The convergence test done with eight staggered streams in upload for 200 and 300 seconds gave the CoV in the range of 0.76 to 0.9 and 0.37 to 0.44, showing the convergence issue still exists in BBR v3, as the streams didn’t settle for a fair share after a considerable amount of time. For fairness, we did rigorous testing with BBR v3 and Cubic streams on different sets (2, 3, 4, 8, 10, 12, 15) in upload, and the corresponding JFi of 0.88, 0.8, 0.72, 0.65, 0.67, 0.65, 0.64 proved that as we increase the number of stream sets, the JFi dropped. These results clearly show that the fairness issue with Cubic still exists in BBR v3, especially when the number of streams is increased.

We believe that our quantitative analysis will help the scientific community to probe further on these lines and to propose a revised version of BBR v3 that is fair with Cubic and converges well. This will surely help BBR to be included in the current Linux mainline kernel 6.xx in the coming years.

Data Availability Statement:

Data is available at the GitHub repository https://github.com/mahsan76/

Funding:

No funding was received for this research

Conflict of Interests:

The author declares no conflict of interest.

References

1. Ha S, Rhee I, Xu L. CUBIC: a new TCP-friendly high-speed TCP variant. ACM SIGOPS operating systems review. 2008;42(5):64-74.
2. Kleinrock L, editor Power and deterministic rules of thumb for probabilistic problems in computer communications. ICC 1979; International Conference on Communications, Volume 3; 1979.
3. Kleinrock L. Internet congestion control using the power metric: Keep the pipe just full, but no fuller. Ad hoc networks. 2018;80:142-57.
4. Cardwell N, Cheng Y, Yang K, Morley D, Yeganeh SH, Jha P, et al. BBRv3: Algorithm bug fixes and public internet deployment. Proceedings of the IETF-117 Internet Engineering Task Force, San Francisco, CA, USA. 2023;26.
5. Sediq AB, Gohary RH, Schoenen R, Yanikomeroglu H. Optimal tradeoff between sum-rate efficiency and Jain's fairness index in resource allocation. IEEE Transactions on Wireless Communications. 2013;12(7):3496-509.
6. Abdi H. Coefficient of variation. Encyclopedia of research design. 2010;1(5):169-71.
7. Zeynali D, Weyulu EN, Fathalli S, Chandrasekaran B, Feldmann A, editors. Promises and potential of BBRv3. International Conference on Passive and Active Network Measurement; 2024: Springer.
8. Høiland-Jørgensen T, Grazia CA, Hurtig P, Brunstrom A, editors. Flent: The flexible network tester. Proceedings of the 11th EAI International Conference on Performance Evaluation Methodologies and Tools; 2017.
9. Zeynali D, Weyulu EN, Fathalli S, Chandrasekaran B, Feldmann A, editors. BBRv3 in the public Internet: a boon or a bane? Proceedings of the 2024 Applied Networking Research Workshop; 2024.
10. Gomez J, Kfoury EF, Crichigno J, Srivastava G. Evaluating TCP BBRv3 performance in wired broadband networks. Computer Communications. 2024;222:198-208.
11. Sarpkaya FB, Srivastava A, Fund F, Panwar S. BBR's Sharing Behavior with CUBIC and Reno. arXiv preprint arXiv:250507741. 2025.
12. Ahsan M, Awan MJ, Yasin A, Bahaj SA, Shehzad HMF. Performace evaluation of TCP cubic, compound TCP and NewReno under Windows 20H1, via 802.11 n Link to LTE Core Network. Annals of the Romanian Society for Cell Biology. 2021;25(6):5357-69.
13. Baldin I, Nikolich A, Griffioen J, Monga IIS, Wang K-C, Lehman T, et al. Fabric: A national-scale programmable experimental network infrastructure. IEEE Internet Computing. 2020;23(6):38-47.
14. Cardwell N. Linux TCP BBRv3 GitHub2025 [Available from: https://github.com/google/bbr.
15. Grazia CA, Patriciello N, Høiland-Jørgensen T, Klapez M, Casoni M, Mangues-Bafalluy J, editors. Adapting TCP small queues for IEEE 802.11 networks. 2018 IEEE 29th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC); 2018: IEEE.
16. Ahsan M, Nabeel M, Khan MMA, Hasan MZ, Ashiq W. Performance Evaluation of BBR-v3 with Cubic and Reno in a Ubiquitous Wired/Wi-Fi Channel. Journal of Computing & Biomedical Informatics. 2025.
17. Ahsan M, Nabeel M, Ahmad A, Javed I, Ashiq W. A Real-World Testbed Evaluation of BBR-n+ with Other Major TCP Congestion Control Algorithms. Journal of Computing & Biomedical Informatics. 2025.