

Promises and Potential of BBRv3

Danesh Zeynali¹, Emilia N. Weyulu¹, Seifeddine Fathalli¹, Balakrishnan Chandrasekaran², and Anja Feldmann¹

¹ Max-Planck-Institut für Informatik

{dzeynali,eweyulu,fathalli,anja}@mpi-inf.mpg.de

² Vrije Universiteit Amsterdam b.chandrasekaran@vu.nl

Abstract. The Bottleneck-Bandwidth and Round-trip (BBR) congestion control algorithm was introduced by Google in 2016. Unlike prior congestion-control algorithms (CCAs), BBR does not rely on signals that are weakly correlated with congestion (e.g., packet loss and transient queue delay). Instead, it characterizes a path using two parameters, bottleneck bandwidth and round-trip propagation time, and is designed to converge with a high probability to Kleinrock’s optimal operating point [34]. Essentially, in stable state, BBR maximizes throughput while minimizing delay and loss. Google has used BBR for a significant fraction of its network traffic both within its datacenters and on its WAN since 2017 [15]. BBR’s interaction dynamics with Cubic, the widely used CCA in the Internet, has received intense scrutiny: Some studies observed BBR to be unfair to Cubic, or generally loss-based CCAs. Google, to its credit, has diligently revised BBR’s design to address the criticisms. This paper focuses on characterizing the promises and potential of the third, and most recent, revision of BBR—introduced to the public in July 2023. We empirically evaluate BBRv3’s performance across a range of network scenarios, e.g., considering different buffer sizes, round-trip times, packet losses, and flow-size distributions. We show that despite the improvements and optimizations introduced in BBRv3, it struggles to achieve an equitable sharing of bandwidth when competing with Cubic, the widely used CCA in the Internet, in a wide range of network conditions.

Keywords: Congestion control, BBR, fairness

1 Introduction

Bottleneck Bandwidth and Round-trip (BBR) is a relatively new congestion control algorithm (CCA) that takes a proactive approach in detecting congestion on a network path. It is designed to enable a sender to converge with a high probability to Kleinrock’s optimal operating point [34], which maximizes throughput while minimizing both delay and loss. To this end, BBR estimates the bottleneck bandwidth and round-trip propagation time of a path, and paces the sender appropriately to avoid building any queue along the path. It has fast replaced Cubic for *all* TCP flows on Google’s B4 WAN [9,13]; for context, Google’s WAN

traffic makes up 13.85% of all Internet traffic [41]. Since BBR runs purely on the sender side, a sample of Google’s edge-to-end-user traffic (i.e., content served on google.com and youtube.com) has been delivered using BBR from as early as 2016 [9]. More recently, Google announced that they use BBR(v3) for *all* internal WAN traffic and public Internet traffic served from google.com [8].

It is quite challenging to design a CCA that works reasonably well in a wide range of network conditions. Traffic characteristics, round-trip times, and bottleneck buffers—to name a few parameters—may differ substantially from one network to another. Unsurprisingly, though Google designed BBR by analyzing traces of world-wide network traffic [13], several independent studies and reports have reported it to be highly unfair to other CCAs in the Internet [7,50,26,42,49,28,17,43,2]. In particular, its incompatibility with Cubic, one of the widely used loss-based CCAs in the Internet, has received much scrutiny [28,17,2]. Google, to its credit, has been diligently evolving BBR to address the concerns raised, and its efforts recently culminated with the release of BBRv3, in 2023 [12]. We focus on the performance and fairness claims (or “promises”) of this most recent version of BBR in this work.

BBR builds a model of the network path based on the bottleneck bandwidth and round-trip propagation time measured from each ACK. BBRv1 used this model to pace the sender at a rate equal to the estimated bottleneck bandwidth, while keeping the in-flight data quite close to the bandwidth-delay product (BDP). Google reported that BBRv1 achieves high throughput under random packet losses as high as 15% and maintains small queues *regardless* of buffer size [9]. In-depth evaluations revealed, however, BBRv1 to be extremely aggressive when competing with loss-based CCAs and unfair in shallow buffer scenarios [7,50,26,42]. Besides, studies also observed increased retransmission rate [7,26,42], high RTT unfairness and queue occupancy [7,49,26], and ACK aggregation issues in WiFi and cellular networks [11]. BBRv2 was released in 2019 to address these issues; chief among its changes was the inclusion of loss as a congestion signal. It reacted to explicit congestion notifications (ECNs) and optimized the congestion-window (`cwnd`) update logic. Despite the inter-CCA and RTT fairness improvements that Google reported [11,14], independent evaluations showed it suffering from low link utilization and being unfair to loss-based CCAs in deep buffer scenarios [37,30,45]. BBRv3, the most recent release introduced in July 2023, claimed to address these concerns. In particular, BBRv3 (a) claimed to offer quick bandwidth convergence to fair shares both with and without loss or ECN signals and (b) boasted of optimizations to minimize queuing delays and packet losses both during and shortly after slow start (i.e., `STARTUP`) [8].

Thoroughly evaluating a CCA to determine whether it falls short of its promises and, most importantly, along what dimensions and how, is an arduous task. The environment (e.g., a datacenter or private WAN) in which the CCA is originally designed, or for which it was specifically developed, may introduce implicit assumptions (for instance, about network conditions or traffic scenarios), which may not hold in the environment where the CCA is even-

tually deployed (e.g., public Internet). The literature has several examples of CCAs whose behavior in real-world network conditions significantly deviated from that observed in the controlled environments where they were designed and tested (e.g., [51,20,40]). In case of BBR, the conflicting observations between Google’s internal evaluations and the various external evaluations—of BBRv1 and BBRv2—clearly attest to the challenges in accurately evaluating its performance. We bridge this conflict by designing a set of experiments—based on first principles—that unequivocally demonstrate whether BBRv3 holds its promises and where there is potential for improvement. Along the way, we outline a systematic and principled approach for evaluating a CCA in practice.

In this work, we empirically analyze the performance of BBRv3 in a wide range of network conditions. In particular, we investigate its claims of quick (throughput) convergence to fair shares in deep and shallow buffer scenarios, and of its improved compatibility (compared to BBRv1 and BBRv2) with loss-based CCAs. An accurate evaluation of these claims has substantial implications for BBR’s deployment in the public Internet, since Google is transitioning more and more of its end-user facing traffic to use BBR. Consequently, we evaluate BBR in network scenarios where both long-lived (i.e., elephant) and short-lived (i.e., mouse) flows compete for bandwidth on a shared link. Our choice of using different sizes of flows is based on prior work that demonstrated that while short-lived flows contribute most packets to Internet traffic, long-lived flows contribute most bytes [4,56].

We summarize our contributions as follows.

- ★ We present a first independent empirical evaluation of BBRv3, Google’s newest version of BBR, across a range of network conditions and justify the rationale behind our choice of evaluation settings. We focus our evaluation, in particular, on BBRv3’s promises to (a) share bandwidth equitably with other loss-based CCAs, particularly in deep buffer scenarios, and (b) assure fairness when contending for throughput with flows that experience different RTTs, a common case in the public Internet where BBRv3 is increasingly being deployed.

- ★ Our evaluations show that BBRv3’s throughput unfairness towards loss-based CCAs, i.e., Cubic, is nearly identical to BBRv1’s behavior and is exacerbated in shallow buffers. We find that BBR’s bias towards long-RTT flows persists in BBRv3, with the unfairness magnified when the difference in RTTs of competing flows is significant. Cubic, the widely used CCA in the Internet, in comparison, offers better throughput for flows with short RTTs than long RTTs.

- ★ We release our network testbed configuration, data sets, and scripts for experiment orchestration and analyses [54].

2 Background

BBR enables a sender to maximize delivery rate while minimizing delay and loss, i.e., converge to Kleinrock’s optimal operating point [34]. To this end, it employs a sequential probing state machine that periodically alternates between probing for higher bandwidths (by increasing the delivery rate) and testing for

Table 1: An overview of the key changes between the three BBR versions.

Life cycle phases	Property	BBRv1	BBRv2	BBRv3
<i>Startup</i>	cwnd_gain	$2/\ln 2$ (~ 2.89)	$2/\ln 2$ (~ 2.89)	2.00
	pacing_gain	$2/\ln 2$ (~ 2.89)	$2/\ln 2$ (~ 2.89)	2.77
	Max. cwnd	3xBDP	–	–
	inflight_hi	–	max.cwnd	max(est. BDP, last cwnd)
	Exit	send. rate <25% for 3 consec. RTTs	loss/ECN rate \geq thresh. (8)	loss/ECN rate \geq thresh. (6)
<i>Drain</i>	pacing_gain		0.35	
	Exit		cwnd ≤ 1 xBDP	
<i>ProbeBW</i>	Phases	8 fixed gain cycles	{Cruise, Refill, Up, Down}	{Cruise, Refill, Up, Down}
	*	Cycle=RTprop	cwnd limits={inflight_hi,inflight_lo}	–
	cwnd_gain	–	cwnd_gain(Up)=2.0	cwnd_gain(Up)=2.25
	pacing_gain	[1.25,0.75,1,1,1,1,1,1]	pacing_gain(Down)=0.75	pacing_gain(Down)=0.90
	Exit	cwnd \geq (pacing_gain x BDP) or loss	loss/ECN rate \geq thresh.	loss/ECN rate \geq thresh.
<i>ProbeRTT</i>	Frequency	10 s	5 s	5 s
	cwnd	4	BDP/2	–
	Duration	200 ms + RTprop	–	–

lower RTTs (by draining the bottleneck queue) of a path. It then uses these bandwidth and RTT measurements to build a path model that determines various congestion control parameters of a TCP implementation such as `cwnd` and pacing rate [9]. Since its introduction in 2016, BBR has undergone three major iterations, and Tab. 1 summarizes the key differences between these versions.

BBRv1 The first version of BBR used four phases: *Startup*, *Drain*, *ProbeBW* (or *Probe Bandwidth*), and *ProbeRTT* [9]. The Startup phase grows the sending rate exponentially, similar to NewReno, Vegas, and Cubic, but it uses a window increase factor (i.e., `cwnd_gain`) of $2/\ln 2$. Once the bottleneck bandwidth estimate “plateaus” (i.e., three attempts to double the delivery rate results in increasing it by less than 25%), it exits the Startup phase. To drain the queue that the Startup may have built, BBRv1 then enters the Drain phase, by reducing its `pacing_gain` to 0.75. A BBR flow spends the majority of its time in the ProbeBW phase, where it probes for bandwidth via *gain cycling*. Essentially, BBRv1 cycled through a sequence of eight values for `pacing_gain`— $5/4$, $3/4$, 1, 1, 1, 1—where a gain higher than one indicates a probe for higher bandwidth and a gain lesser than one represents the attempt to drain any queue built earlier. With the average of the eight values being one, the approach allowed ProbeBW to maintain an average pacing rate that is equal to the estimated bandwidth. BBRv1 continuously sampled the bottleneck bandwidth during this phase. Every 10 seconds, BBRv1 stopped the ProbeBW phase and entered the ProbeRTT phase. It reduced its `cwnd` to four packets for as long as 200 ms and one round-trip. The low rate drains the bottleneck queue and allows the sender to sample the minimum propagation RTT of the path; it is only in this phase that BBR updates its minimum RTT estimate. Several studies revealed, however, that BBRv1 was extremely aggressive when competing with loss-based CCAs, and highly unfair in shallow buffer scenarios [7,50,26,42,28].

BBRv2 Google introduced BBRv2 in 2019 to alleviate the problems with BBRv1 [13,14]. BBRv2 split the ProbeBW phase into four new sub-phases: *Down*, *Cruise*, *Refill*, and *Up*. Unlike BBRv1, it reacted to loss or ECN signals to facilitate an equitable sharing of bandwidths with widely used CCAs




Fig. 1: An overview of (a) BBRv3’s high-level implementation architecture & (b) the life cycle of the CCA showing how it transitions phases of congestion control.

such as Cubic and NewReno [14]. In BBRv2, the Startup phase ends not only when the bandwidth estimate “plateaus,” but also if loss or ECN mark rate exceeds a threshold (i.e., 2%). It capped the maximum in-flight data volume to the maximum congestion window observed, prior to entering the Drain phase. During ProbeBW, BBRv2 picks the maximum bandwidth measured in the last two ProbeBW attempts instead of the last 10 RTTs as in the case of BBRv1. BBRv2 maintained three different `pacing_gain` values to bypass the issue of 2xBDP in-flight cap, which was identified in prior work as the primary source of unfairness in BBRv1 [50]: 1.25 for Probe:Down, 0.75 for ProbeBW:Up, and 1 for both ProbeBW:Cruise and ProbeBW:Refill. BBRv2 also replaced the 8-phase cycle for ProbeBW with an adaptive probing time, i.e., a value picked at random between 2-3 seconds, for improving compatibility with Cubic and NewReno.

BBRv3 BBRv3 was introduced in July 2023 at the 117th IETF meeting [8]. Fig. 1 illustrates both the high-level implementation architecture and the life cycle of the CCA, showing how it transitions between the different congestion control phases. BBRv3 is a minor revision of BBRv2, and it addresses two key performance issues of BBRv2. First, an implementation change caused BBRv2 to prematurely exit the ProbeBW phase that in turn prevented it from sharing bandwidth equitably with BBRv1 or loss-based CCAs such as Cubic [8]. The issue resulted in under-utilization of the link even after the link was no longer congested. BBRv3, hence, continues probing for bandwidth until (a) the loss rate or ECN mark rate is below a set threshold (i.e., 2%) or (b) the bandwidth saturates even if the delivery rate is not restricted by the `inflight_hi` threshold. Second, the choice of parameter values (e.g., `cwnd_gain` and `pacing_gain`) used in the ProbeBW phase made BBRv2 highly unfair: Especially in deep buffer scenarios if the sender received no loss or ECN signals, it caused the sender to be unfair to competing flows [8]. To mitigate these fairness concerns, Google adjusted several parameters, e.g., `cwnd_gain` , `pacing_gain` , and `inflight_hi` , used in the different congestion control phases, which per their report resulted in improving the CCA’s performance [8].

3 Methodology

We empirically evaluate the performance of the three versions of BBR and compare them to that of Cubic using a custom network testbed. The testbed consists of six Dell R6515 blade servers, each with 16 cores and 128 GiB of RAM. We installed one or two Broadcom NetXtreme BCM5720 25 GbE NICs, as required, on each server and used 25 Gbps DACs for interconnecting the servers and four APS Networks BF6064X-T switches (as shown in Fig. 2). We configured the servers to run Linux kernel v5.4 (Debian 10). We used two VMs on one of the servers to evaluate different versions of BBR implementations using the same machine: One VM ran BBRv2 with Linux kernel v5.13.12 from [22], while the other ran BBRv3 with Linux kernel v6.4.0 from [23]. We did not enable ECN support in our testbed, and so neither BBRv2 nor BBRv3 benefited from ECN in our evaluations. Since BBRv3 promises quick (bandwidth) convergence to fair share even without loss or ECN signals [8], we assume such a scenario in our setup and leave the evaluations in the presence of ECN to future work.

To evaluate under a wide range of network conditions, we varied the capacity of the bottleneck link (in red in Fig. 2) and one-way delays as needed for the different experiments using the Linux traffic control (`tc`) utility. Specifically, we used `netem` [47] on an intermediate node, designated as the network latency emulator (NLE), between the end hosts to avoid issues with TCP small queues [31,38]. We also used `tc-tbf` [1] on the NLE to configure the bottleneck bandwidth and buffer size.

Unless otherwise stated, we added half of the configured delay before the bottleneck and half after the bottleneck. The base RTT between sender-receiver pairs in the testbed, without any added delay, was 1.5 ms on average.

To emulate traffic representative of real-world network conditions, we generated workloads by using flow sizes from a MAWI trace [32] downloaded on January 5, 2023. We extracted the flows using Zeek [29] and scaled up the inter-arrival times by a factor of 5 to account for the difference in link capacities between the MAWI infrastructure and our testbed. The flow sizes in this data set followed the Pareto distribution, while the inter-arrival times followed the Lognormal distribution. We then used Harpoon [44] to generate traffic using samples from these empirical distributions. Unless otherwise stated, we repeated each experiment three times and reported the statistics across the three runs.

4 Evaluation

We now evaluate the three versions of BBR and Cubic in a wide range of network conditions. Our goal is to ascertain the ability of each CCA to share the bandwidth in a fair or equitable manner with other flows—regardless of whether the




Fig. 2: *Network testbed for evaluating performance of CCAs in a single bottleneck setup.*




Fig. 3: Throughput achieved by two similar, synchronous (i.e., starting at the same time) flows when competing for bandwidth on a bottleneck with a deep buffer of size $16 \times \text{BDP}$.

contending flows use the same CCA or a different one. To this end, we evaluate fairness in two broad settings: *similar* and *dissimilar* flows. Similar flows use the same CCA and experience the same RTT, but may vary from one another in size or arrival times (at a bottleneck link). Dissimilar flows, in contrast, differ from one another based on one or more factors, e.g., choice of CCAs and RTTs experienced by the flows.

4.1 Similar and synchronous flows

We begin our evaluation with a simple scenario where two similar flows, i.e., using the same CCA and experiencing the same RTT (of 100 ms), contend for bandwidth on a bottleneck (refer Fig. 2). We start both flows *at the same time* and let them run for a duration of 300 s and plot the throughput experienced by the flows as a function of time in Fig. 3. Overall, all BBR variants achieve an equitable bandwidth share when competing with a similar and synchronous flow. BBRv3 significantly reduces the throughput oscillations and allows the flows to converge much quicker than either of the earlier two versions. We repeated this experiment by varying the bottleneck queue size from one-fourth to 16 times the BDP, doubling the queue size once for each experiment. Although not shown here, our inferences hold across all these settings.

Fig. 4a indicates whether the flows experience high retransmissions when competing with another similar, synchronous flow.³ The heatmap reports the percentage of retransmissions (i.e., ratio of the aggregate number of retransmitted packets to the total number of packets delivered by both flows) experienced by the two flows; higher values indicate higher retransmissions, and, hence, lower values (in green) imply better throughput or utilization. Per this figure, although BBRv1 experiences high retransmissions when competing with another BBRv1 flow started at the same time in low buffer settings, across all other settings, the CCAs converge quickly without incurring much loss. This observation of high retransmissions when using BBRv1 in shallow buffers is in line with observations

³ We omit the results for $8 \times \text{BDP}$ and $16 \times \text{BDP}$ as we see virtually no retransmissions in such settings.

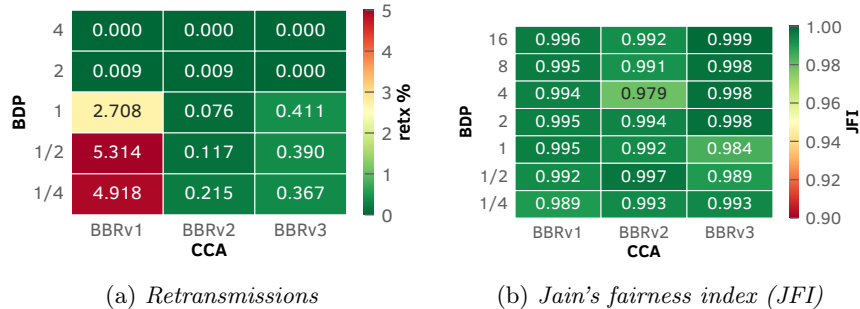


Fig. 4: (a) Retransmissions experienced by and (b) Jain's fairness index (JFI) of two similar, synchronous flows competing for bandwidth on a bottleneck link as a function of bottleneck queue size.




Fig. 5: Throughput achieved by two similar but staggered flows (*i.e.*, second flow starts 15s after the first) when competing for bandwidth on a bottleneck with a deep buffer of size $16 \times BDP$.

made by prior work [7,26,42]: They are due to BBRv1's aggressive probing behavior and its indifference to loss. The Jain's fairness index (JFI) for the flows in Fig. 4b, however, shows that all three versions equitably share the bottleneck with a similar and synchronous flow, regardless of bottleneck buffer size; BBRv3 shows only a marginal improvement, if any, over the prior two versions.

Takeaways. BBR achieves an equitable sharing of bandwidth with similar synchronous flows, and BBRv3 address the well-documented issues of aggressiveness and unfairness in BBRv1.

4.2 Similar but staggered flows

We now repeat the earlier evaluation, but with one change: We *stagger* the flows so that the second flow starts 15s after the first. In the Internet flows may arrive at a bottleneck at random times, when other contending flows on the link have already started or reached their stable state. The staggering of flows simply emulates this typical traffic condition.

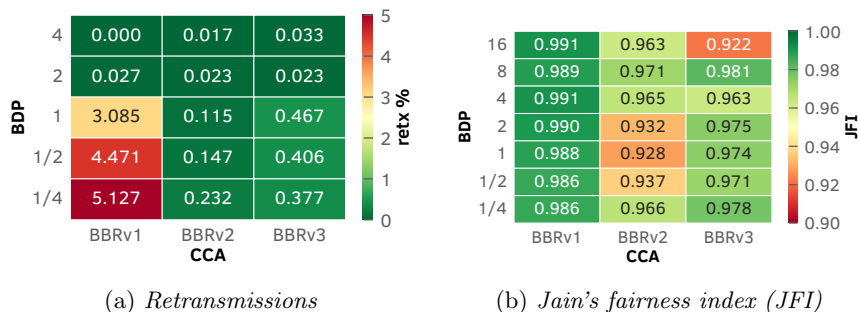


Fig. 6: (a) Retransmissions experienced by and (b) JFI of two similar but staggered flows (*i.e.*, second one starts 15 s after the first flow) competing for bandwidth on a bottleneck link as a function of bottleneck buffer size.

Per Fig. 5, BBR implementations do not converge quickly to an equitable share when flows do not start at the same time—quite unlike the case of similar and synchronous flows (§4.1). For instance, when a BBRv3 flow joins 15 s later at a bottleneck link with a $16 \times \text{BDP}$ buffer, it takes *more than 4 minutes* for it to achieve an equitable bandwidth sharing. Such large buffers are not uncommon in the internet: Router manufacturers may (by default) provide large buffers [19], and administrators may configure larger buffers on transcontinental links to cope with the high RTTs [16] or improve video QoE [46].

We observe that when flow starts are staggered by an interval of 15 s, the flows experience higher retransmissions (Fig. 6a) than when they are started at the same time.⁴ The heatmap in the plot reports, as in §4.1, the percentage of retransmissions experienced by the two flows; higher values (in yellow and red) indicate higher retransmissions. BBRv1 in particular suffers high retransmissions in shallow buffer settings, again owing to its aggressive probing behavior and indifference to losses. A key change in BBRv2 (compared to BBRv1) was the addition of using loss as a congestion signal; this change enables it to reduce the loss rates in shallow buffer settings and avoid unnecessary retransmissions [10]. BBRv2 experiences, as a result, the lowest retransmissions of all three implementations. BBRv3 inherits BBRv2’s changes and, as a consequence, experiences low retransmissions, which matches Google’s claims about BBRv3 [8]. We observe that BBRv3 experiences more retransmissions than BBRv2 in shallow buffer settings, but the former achieves better fairness than the latter across a range of buffer sizes (Fig. 6b), with the exception of the largest buffer size (of 16-times BDP). BBRv1, which does not use packet loss to infer congestion, leads to higher fairness than BBRv2 and BBRv3, both of which take loss into consideration. These observations could partly be explained by BBRv3’s less aggressive be-

⁴ As before, we omit the results for 8 times BDP and 16 times BDP as we see virtually no retransmissions in such settings.




Fig. 7: When we stagger the arrivals and departures of five flows at a bottleneck link of 100 Mbps and buffer size $32 \times \text{BDP}$ by 100 s, (a) BBRv1 flows achieve fair throughput shares faster than (b) BBRv2 and (c) BBRv3.




Fig. 8: When we stagger the arrivals and departures of five flows at a bottleneck link of 100 Mbps by 100 s, (a) BBRv1 flows achieve better shares of throughput than (b) BBRv2 and (c) BBRv3.

havior compared to BBRv1 (compare BBRv3’s `cwnd_gain` and `pacing_gain` to those of BBRv1 in Tab. 1).

Takeaways. While BBR’s refinements seem to allow it reduce losses, with recent versions taking losses into account to control their probing, even the most recent version of BBR struggles to achieve high fairness when flow arrivals are staggered by a few seconds. This observation has crucial implications for BBRv3’s adoption in the Internet, since in the public Internet arrivals of competing flows might often be staggered with respect to one another.

4.3 Co-existence in deep buffer scenarios

To evaluate how quickly different CCAs converge to fair bandwidth share, we run a staggered-flow experiment (similar to that of HPCC [35]), where we introduce five flows, one after another. Again, the flows are similar and we add a new flow to the testbed every 100 s until we reach five concurrent flows, and then remove one flow every 100 s until we have only one left. This experiment enables us to compare and contrast how quickly different CCAs react to dynamic changes in traffic conditions and converge to the fair share of the bottleneck (100 Mbps). In this evaluation, we specifically focus on deep buffer scenarios to verify Google’s claim that BBRv3 can equitably share bandwidth in such conditions [8].

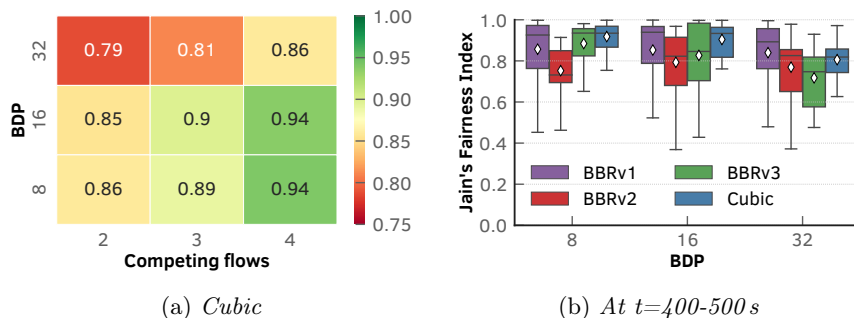


Fig. 9: (a) Cubic offers improved fairness as the number of flows increases, and (b) when all five flows compete at the bottleneck with a deep ($32 \times \text{BDP}$) buffer, BBRv3 has the worst fairness than any other CCA.

Fig. 7 shows the timelines of flows in the staggered-flow experiment in a deep-buffer (i.e., $32 \times \text{BDP}$) scenario for all three BBR versions. Per this figure, BBRv3 struggles to achieve a fair share between the flows; when all five flows compete at the bottleneck, BBRv3 performs the worst compared to both the previous versions. BBRv1 flows converge quickly to fair share both when new flows arrive and old ones leave; its quick convergence to fair share is perhaps largely due to its aggressive probing behavior, which can efficiently discover bandwidth changes. The average JFI across three runs (Fig. 8) shows that BBRv3 performs better than BBRv2, but poorer than BBRv1. Cubic, in contrast to BBRv3, performs poorly with a small numbers of flows, although it improves with increasing number of flows (Fig. 9a). Unlike Cubic, BBRv3's performance degrades with increasing number of flows. If we analyze fairness during the period when all five flows are competing for bandwidth (Fig. 9b), BBR typically performs poorer than Cubic on average, exhibiting large variations in performance. Particular in deep-buffer settings, BBRv3 performs worst compared to both Cubic and earlier versions of BBR, perhaps because of its less aggressive choice of probing parameters.

Takeaways. While BBRv3's intra-protocol fairness in deep buffer scenarios is better than that of BBRv2, its fairness is poorer than that of BBRv1 and seems to worsen as the number of flows increases (in deep-buffer scenarios). BBRv3 choice of less aggressive probing parameters compared to BBRv1 seems to be the reason behind the observed behavior.

4.4 Dissimilar flows: RTT fairness

Evaluations concerning fairness of CCAs in prior work (e.g., [50,7]) usually consider flows with similar RTTs. Flows in the Internet typically have different RTTs, which make it challenging in terms of fairly sharing bandwidth when such flows interact at a bottleneck link, even if they all use the same CCA. The RTT dictates a CCA's reaction to bandwidth limitations or to other flows, since




Fig. 10: Average throughput difference (TD); 160 ms flow starts first, then 15 s later the (10 ms, 20 ms, 40 ms) flow starts.





Fig. 11: Average Jain's fairness index (JFI); 160 ms flow starts first, then 15 s later the (10 ms, 20 ms, 40 ms) flow starts.

the signals typically used by a sender (e.g., packet loss or delay) for determining how much it can send and/or how fast, must traverse the path from the bottleneck to the receiver, which then echoes it back to the sender.

With this experiment, we evaluate the bandwidth share when flows using the same CCA traverse different routing paths, thus experiencing different RTTs, and sharing a common bottleneck along the way. We run two flows using the same CCA, where the first flow is our base flow with a fixed RTT of 160 ms and the second flow experiences a different RTT—one of 10, 20, 40, 80 ms, each constituting a different experiment. In the first experiment scenario, the base flow, 160 ms, starts first and the second flow joins after 15 seconds (Fig. 10, Fig. 11 and Fig. 12). In the second scenario, the base flow joins 15 seconds after the flow with the shorter RTT (Fig. 13, Fig. 14 and Fig. 15). These scenarios help us characterize how the different CCAs behave when the short (or long) RTT flow joins the network, while the other flow has already reached a steady state; we can safely assume a flow has finished slow start after 15 seconds.⁵ Specifically, we calculate the throughput difference between the base flow and the second flow, as follows: $BaseFlow_{AvgT_{put}} - SecondFlow_{AvgT_{put}}$. A positive value implies that the base flow obtains the majority of the bandwidth, while a negative value indicates that the second flow receives higher bandwidth than

⁵ We do not disturb a flow while in slow start to ensure that it does not exit slow start prematurely, which would result in link underutilization.

(a) *TD*; 160 ms vs 80 ms (b) *JFI*; 160 ms vs 80 msFig. 12: Average throughput difference (*TD*) and Jain's fairness index (*JFI*); 160 ms flow starts first, then 15 s later the 80 ms flow starts.(a) *TD*; 10 ms vs 160 ms (b) *TD*; 20 ms vs 160 ms (c) *TD*; 40 ms vs 160 msFig. 13: Average throughput difference (*TD*); short *RTT* flow (10 ms, 20 ms, 40 ms) starts first, then 15 s later the 160 ms flow starts.

the first. Additionally, we also compute the Jain's fairness index to quantify the CCA's fairness.

When we start the long-RTT flow first, all versions of BBR favor the long-RTT flow (Fig. 10); it receives much higher bandwidth than the second flow. This bias towards the long-RTT flow (starting first) diminishes a bit when the difference in flow RTTs decreases (Fig. 10c); BBRv2 and BBRv3 (which is largely similar to BBRv2) perform much better than BBRv1 when RTT differences between the base flow and second flow decreases (Fig. 12), perhaps owing to their less aggressive design. These observations also hold if we start the short-RTT flow first and then the long-RTT (or base) flow (Fig. 13 and Fig. 15). Our inferences are in agreement with prior work that demonstrated BBR favoring long-RTT flows over short-RTT flows [49,26], which is quite different from what has been observed in case of AIMD mechanisms [6,25]; our results confirm that BBRv3's behavior does not vary substantially from BBRv2 with respect to preferring long-RTT flows over short-RTT flows. Cubic, unlike BBR, favors the second (short-RTT) flow over the base (long-RTT) flow, even though the second flow starts after the base flow, across almost all buffer settings.

All BBR versions exhibit high RTT unfairness when the RTT difference between the base and second flow is significant, regardless of which flow starts first (Fig. 11, Fig. 12, Fig. 14, and Fig. 15). Unlike BBR, Cubic achieve high fairness across all settings, regardless of how we size the bottleneck buffer. BBR's bias



Fig. 14: Average Jain's fairness index (*JFI*); short *RTT* flows (10 ms, 20 ms, 40 ms) start first, then 15 s later the 160 ms flow starts.




Fig. 15: Average throughput difference (*TD*), 80 ms flow starts first, then 15 s later the 160 ms flow starts.

towards long-RTT flows is less pronounced in shallow buffers (if the RTTs of the flows do not differ substantially), where the short-RTT flow is able to compete for a better bandwidth share with the long-RTT flow.

Fairness is improved for all BBR versions if the short RTT flow starts first, as it is able to grow its `cwnd` before the long RTT flow starts. Additionally, the smaller the gap between the two flow's RTT values, the better the fairness as both flows are able to estimate the bottleneck bandwidth and RTT at the same rate (Fig. 12 and Fig. 15). When BBR flows with different RTTs compete for bandwidth, the short-RTT flow is likely to observe the increase in queuing faster than the long-RTT flow (since the rate at which a sender can observe such signals depend on the RTT). Furthermore, since BBR determines the delivery rate based on such signals, the short-RTT flow may consistently slow down in response to any observed queuing before the long-RTT flow reacts. When flows eventually synchronize, short-RTT flows regain some of the bandwidth share, but the situation reverts to benefitting the long-RTT flow again as they probe for bandwidth.

Takeaways. BBR's bias towards long-RTT flows is well-known, and our evaluations confirm that BBRv3 improvements or optimizations do not change the status quo. BBR offers poor fairness across flows that differ substantially in RTT, which might be typical in the Internet. Cubic, the widely used CCA in the




Fig. 16: Throughput of a BBR flow when competing for bandwidth with a Cubic flow on a bottleneck link with a $1 \times \text{BDP}$ buffer.




Fig. 17: Jain's fairness index (JFI) of two flows using different CCAs and competing on a bottleneck with a $1 \times \text{BDP}$ buffer.

Internet, unlike BBR, offers higher fairness than BBR in nearly all our evaluations scenarios with flows with different RTTs.

4.5 Dissimilar flows: Inter-CCA fairness

We now evaluate BBRv3's ability to share bandwidth equitably when competing with flows using CCAs other than BBRv3. With more than 40% of current Internet traffic estimated to be delivered using BBR [36], it is quite important to characterize how BBRv3 behaves when competing with non-BBRv3 flows already prevalent in the Internet. Such an inter-CCA characterization is also timely and crucial for network operators and generally the networking community, since Google intends to submit BBRv3 soon for inclusion in the Linux kernel [8]. To characterize BBRv3 flows' ability to co-exist with non-BBRv3 flows, we run two (co-existing) long-lived flows, each using a different CCA; we start these flows at the same time. We set the RTT to 100 ms and the bandwidth to 100 Mbps, and we vary the bottleneck buffer size as a function of BDP.

BBRv1 has long been reported to be extremely unfair to loss-based CCAs such as Cubic [50,7,49,26]. Fig. 16a, hence, simply confirms this well-established claim. BBRv2 seems to fare well than BBRv1 (Fig. 16b), since unlike the former, the latter takes packet loss into consideration. BBRv2 improves upon BBRv1 with respect to its ability to share bandwidth with loss-based CCAs, even allowing Cubic to obtain more than its share of fair bandwidth when competing




Fig. 18: Retransmissions when (a) a *BBRv3* flow and (b) a *Cubic* flow compete against flows using other CCAs.

with *BBRv2* flows. *BBRv3*, which Google claims will quickly converge to fair share, performs even slightly worse than *BBRv1* (Fig. 16c). Even if flows experience same RTTs, *BBRv3*, in its current form, offers no room for *Cubic* flows to co-exist; any RTT differences between the flows may only exacerbate this situation.

Fig. 17 shows the Jain’s fairness index for the competing flows to characterize *BBR*’s inter-CCA fairness. In shallow-buffer scenarios, *BBRv1* is quite unfair to all others. No other CCA, not even other versions of *BBR*, is able to obtain a fair share when competing with *BBRv1* (Fig. 17a); increasing bottleneck buffer sizes, nevertheless, seems to alleviate this situation substantially. This behavior is presumably due to *BBRv1* only backing off when its large in-flight data cap ($3 \times \text{BDP}$) is exceeded (refer Tab. 1), ignoring packet loss unlike *BBRv2* and *BBRv3*. Unlike *BBRv1*, *BBRv2* performs quite well (with high JFI values) across all bottleneck buffer settings (Fig. 17b) primarily because of its ability to reach to packet loss. It can equitably share bandwidth with *Cubic*, making it safer than *BBRv1* for deployment in the public Internet where the latter is quite prevalent. *BBRv3*, despite being only a minor revision of *BBRv2*, performs quite poorly, especially against *Cubic* (Fig. 17c). In deep-buffer settings, *Cubic* is able to send more data into the network before experiencing congestion (i.e., packet loss); naturally, it obtains a higher bandwidth compared to shallow-buffer settings and, as a result, the fairness index improves.

Overall, while the design of *BBRv3* (which it inherits from *BBRv2*) represents a substantial progress (in creating a fair and safe CCA), the performance tweaks [8] or optimizations (refer Tab. 1) that Google has introduced in *BBRv3* are exactly the opposite. With regards to Google’s claims about reduced packet loss, we observe frequent retransmissions (Fig. 18) when *BBRv3* competes against *Cubic*, comparable almost to those when *BBRv1* competes against *Cubic* in both deep and shallow-buffer settings.

Thus far we focused on a scenario where a single *BBR* flow competes with a single *Cubic* flow for bandwidth on a bottleneck. Previous studies showed that




Fig. 19: *JFI* when 1-5 Cubic flows compete against 1-5 (a) *BBRv1*, (b) *BBRv2* and (c) *BBRv3* flows.

even when competing with several Reno or Cubic flows a single *BBRv1* flow was able to grab most of the link bandwidth [40,50]. We revisit this evaluation where we pit a single *BBR* flow against several Cubic flows. More specifically, we conduct a series of experiments where we vary the number of *BBR* flows as well as competing Cubic flows from 1 through 5, and characterize the fairness (measured via *JFI*) across the competing flows in each experiment (Fig. 19). As before, we set the link bandwidth to 100 Mbps, the RTT to 100 ms, and the buffer size to $1 \times \text{BDP}$ (i.e., 1250 KB). We start all the flows at the same time and run them for 360 s. Fig. 19 shows the *JFI* of the experiments with the three *BBR* versions competing against Cubic. The top value in each cell is the mean *JFI* value over time, concatenated over three experiment runs, while the bottom value show the mean throughput shares of *BBR* and Cubic.

In Fig. 19a, *JFI* values become smaller (i.e., fairness worsens) as we move from top to bottom or right to left. This observation confirms the observations from prior work that even a single *BBRv1* flow can outcompete multiple Cubic flows [40,50]. *BBRv2*'s design completely reverses this behavior (Fig. 19b); when we increase the number of Cubic flows fairness improves, emphasizing once again that we can safely deploy *BBRv2* in the public Internet. *BBRv3*, in contrast, resurrects *BBRv1*'s aggressive behavior (Fig. 19c); it offers lower fairness than *BBRv2*, although the Cubic flows seem to be able to obtain a higher bandwidth when competing against *BBRv3* than when competing against *BBRv1*.

Takeaways. While *BBRv2* makes significant improvements towards achieving equitable bandwidth sharing when competing against loss-based CCAs, *BBRv3* seems to undo most, if not all, such improvements. *BBRv3* does not equitably share bandwidth with loss-based CCAs such as Cubic, and its behavior in some instances is even worse than that of *BBRv1*. Despite its ability to react to packet loss, it remains highly unfair to Cubic flows, in shallow buffers.

4.6 Dissimilar flows: Real-world workloads

Until now, we used a small, fixed number of (long-lived) flows in our evaluations. Now, we turn to using realistic traffic workloads to quantify *BBR*'s performance in real-world network conditions. Specifically, we use a distribution




Fig. 20: (a) Distribution of flow sizes in the MAWI trace; Flow completion times (FCTs) and FCT “slowdown” for various flow sizes across different CCAs.

of flows comprising both short-lived (i.e., mouse) and long-lived (i.e., elephant) flows. We sample the size of these flows from an empirical distribution of flow sizes observed in the Internet (Fig. 20a). The CDF in Fig. 20a depicts the flow size distribution from a MAWI trace (refer §3). We use the Harpoon [44] traffic generator to sample the flow sizes and connection times from this trace data. We fixed the random seed used in the sampling to ensure that the sampled distribution of flow sizes remains the same across the evaluations of all CCAs; we can then compare the performance of the same set of flows across different CCAs and determine which CCA performs the best. We conduct two distinct experiments using the topology shown in Fig. 2 with the bottleneck bandwidth set to 100 Mbps, RTT to 100 ms, and the buffer size to BDP. In the first, we set all flows to use the same CCA, and in the second, we configure half of the flows to use one of the BBR versions and the other half to use Cubic. We then analyze the flow completion times (FCTs) as well as the FCT “slowdowns” [53] computed by normalizing the measured FCT by the theoretically optimal FCT obtained by taking into account the flow size, the bandwidth, and the RTT.

When all flows use the same CCA. In case of BBRv1, the FCTs of short flows, i.e., flows smaller than 300 KB, are comparatively longer than those of the long flows (Fig. 20b). The FCT slowdowns for the longest flows in Fig. 20c when compared to those of the smaller flows clearly show that BBRv1 prioritizes long flows over their shorter counterparts. BBRv2, in contrast, offers lower FCT for short flows than for long flows. Lastly, BBRv3 offers smaller FCT slowdowns for short flows than for long flows, but its FCT slowdowns are consistently higher than those of BBRv2.

When flows are split equally between two CCAs. To evaluate how BBR behaves when competing with Cubic (similar to the earlier evaluations in §4.5) in realistic traffic conditions, we configure half of the flows in the workload to use one of the BBR versions and the other half to use Cubic. As before, we plot the FCTs and FCT slowdowns of the flows grouped into buckets of different sizes in Fig. 21. The FCTs and FCT slowdown values corresponding to the experiment where BBRv1 flows compete with Cubic (Fig. 21a and Fig. 21d) are not surprising; BBRv1 is quite unfair to Cubic, as amply demonstrated in §4.5, and its ability to seize

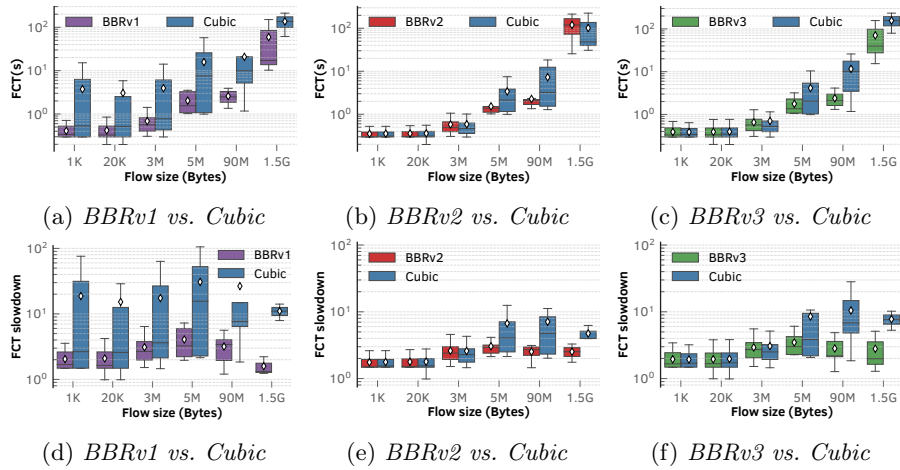


Fig. 21: Flow completion times (FCTs) and FCT “slowdowns” for flows of varying sizes split between two CCAs.

bandwidth aggressively from Cubic flows result overall in substantially smaller FCTs and FCT slowdowns for BBRv1 flows than Cubic flows, regardless of flow size. BBRv2 is far less aggressive to Cubic than BBRv1 (Fig. 21b and Fig. 21e). Short flows of both CCAs experience similar FCT slowdowns, while long flows using BBRv2 fare better than those using Cubic. BBRv3’s performance is similar to that of BBRv2, except in case of long flows. (Fig. 21c and Fig. 21f) Long flows using Cubic experience higher FCTs and FCT slowdowns when competing against BBRv3 than against BBRv1.

Takeaways. The evaluations confirm the highly aggressive and unfair behavior of BBRv1 towards loss-based CCAs such as Cubic. BBRv2’s design significantly improves that status quo, making Cubic flows achieve nearly their fair share when competing with BBRv2 flows. While BBRv3 behaves similar to BBRv2 for short flows allowing both Cubic and BBRv3 flows to experience similar FCT slowdowns, long flows using BBRv3 experience much smaller FCT slowdowns than those using Cubic.

5 Related Work

There is an extensive body of prior work on congestion control in the Internet. Google’s BBR is a relatively new addition and its design intelligently combines several good ideas from the literature. Several independent studies have examined BBR’s performance in various scenarios and compared it with that of other CCAs, partly perhaps because of how quickly Google has been transitioning large volumes of traffic to use BBR [8]. Below, we briefly review these prior work.

Many prior work such as [50,49,7,42,26] conducted comprehensive evaluations of BBRv1, focusing on its interactions with other well-known and widely

used CCAs, e.g., NewReno, Cubic, and Vegas. Past evaluations of BBRv1 considered both similar and dissimilar flows (i.e., flows with varying RTTs or using different CCAs). Some investigated BBRv1’s performance in both shallow and deep buffer scenarios [50,26,7], and some others analyzed the impact of extremely long delays, which are typical in satellite networks [18]. Prior studies identified several shortcomings with BBRv1, including unfairness to loss-based CCAs in shallow buffers, frequent retransmissions, high RTT unfairness as well as high queue occupancy. The scope of virtually all such evaluations was, however, limited to long-lived flows. Hurtig et al. tested BBRv1 with various buffer sizes and link bandwidths and used a workload comprising a mix of bulk as well as short transfers [27], but they used fixed sizes (instead of sampling them from observed traffic distributions) for long-lived and short-lived flows.

Active queue management (AQM) techniques are designed for mitigating congestion [3], and several prior work have studied the impact of deploying AQMs on the performance of different CCAs (e.g., [24]). In the case of BBR, studies have shown that AQMs such as FQ-CoDel improve the performance of both BBRv1 and BBRv2 in deep buffer scenarios [55,21,33]. We consider the analyses of AQM deployments on BBRv3 as orthogonal to this work.

BBRv2 uses ECN signals to detect queuing, and Tahiliani et al. [48] showed that BBRv2 reduces queue occupancy at the bottleneck, when ECN was enabled. Multiple studies demonstrated that BBRv2 has lower link utilization than BBRv1 owing to its conservative behavior during bandwidth probing and that both versions were unfair to loss-based CCAs in deep buffer scenarios [37,30,45,52]. We did not evaluate the benefits of ECN for BBRv3 in this study, but we leave that for future work.

Some recent work also attempted to address the BBR’s shortcomings. Bi et al. [5], for instance, focused on BBRv1’s frequent retransmissions issue. They proposed dynamically adjusting the in-flight data cap based on the estimate bottleneck buffer size from RTT samples as well as packet losses. BBRv2+ used path delay information for tuning the aggressiveness of bandwidth probing [52]. They used Mahimahi [39] for evaluating BBRv2+ in emulated WiFi and LTE networks. While we do not propose solutions for addressing the shortcomings in BBRv3, we hope our comprehensive evaluations and detailed analyses will pave the way for designing and validating different solutions.

6 Concluding Remarks

The bottleneck bandwidth and round-trip propagation time (BBR) algorithm is a relatively new congestion control algorithm (CCA). BBR’s design eschews the typical congestion signals (e.g., transient queue delays and losses) used in the vast majority of prior work. It measures instead bottleneck bandwidth and round-trip times on each ACK and paces the sender to retain the delivery rate close to the bottleneck bandwidth and avoid any queue build up. Several studies, however, demonstrated this initial design to be highly unfair to loss-based CCAs such as Cubic. Google responded to the criticisms by evolving the design and

releasing newer versions of BBR. BBRv3, released in July 2023, is the most recent version in the evolution of BBR.

Given the increasing volume of traffic that Google has been transitioning to use BBR and its current efforts to include BBRv3 in the Linux kernel, we turned our attention to performing a systematic and rigorous evaluation of BBRv3 in a range of realistic network conditions and traffic scenarios. Specifically, we checked whether Google’s claims of BBRv3’s fairness towards other CCAs hold water in our evaluations. While BBRv3 has evolved substantially compared to the earlier versions, in our evaluations BBRv3 struggles to achieve an equitable bandwidth sharing with competing flows. Even when competing with similar flows (i.e., flows using BBRv3) a small difference in arrival times causes BBRv3 to incur substantial delays in bandwidth convergence. Despite the revisions, optimizations, and bug fixes, BBRv3’s highly competitive behavior stifles Cubic on a bottleneck link, particularly in shallow buffer settings. These results have crucial implications for BBRv3’s adoption and deployment in the public Internet. We release our testbed configuration and scripts for running and analyzing the experiments as open source artifacts [54], and we hope our efforts encourage the community to think about an objective template or framework for evaluating modern CCAs.

References

1. Alexey N. Kuznetsov and Bert hubert: tc-tbf(8) - Linux man page. <https://linux.die.net/man/8/tc-tbf> (2023), [Last accessed: October 01, 2023]
2. Ayush, Jingzhi, Melodies, Sean, Raj, Ben: Conjecture: Existence of Nash Equilibria in Modern Internet Congestion Control. In: Proceedings of the 5th Asia-Pacific Workshop on Networking. APNet '21 (2022)
3. Baker, F., Fairhurst, G.: IETF Recommendations Regarding Active Queue Management. RFC 7567 (Jul 2015)
4. Bauer, S., Jaeger, B., Helfert, F., Barias, P., Carle, G.: On the evolution of Internet flow characteristics. In: Applied Networking Research Workshop (ANRW) (2021)
5. Bi, P., Xiao, M., Yu, D., Zhang, G.: oBBR: Optimize Retransmissions of BBR Flows on the Internet. In: USENIX Annual Technical Conference (ATC) (2023)
6. Brown, P.: Resource sharing of TCP connections with different round trip times. In: IEEE Int. Conference on Computer Communications (INFOCOM) (2000)
7. Cao, Y., Jain, A., Sharma, K., Balasubramanian, A., Gandhi, A.: When to Use and When Not to Use BBR: An Empirical Analysis and Evaluation Study. In: ACM Internet Measurement Conference (IMC) (2019)
8. Cardwell, N., Cheng, Y., Gunn, C.S., Yang, S., Morley, D., Yeganeh, S.H., Swett, I., Wu, B., Vasiliev, V., Jha, P., Seung, Y., Jacobson, V.: BBRv3: Algorithm Bug Fixes and Public Internet Deployment. Tech. rep., Internet Engineering Task Force (IETF) (Jul 2023), <https://datatracker.ietf.org/meeting/117/materials/slides-117-ccwg-bbrv3-algorithm-bug-fixes-and-public-internet-deployment-00>
9. Cardwell, N., Cheng, Y., Gunn, C.S., Yeganeh, S.H., Jacobson, V.: BBR: Congestion-Based Congestion Control. ACM Queue 14(5) (Oct 2016)
10. Cardwell, N., Cheng, Y., Gunn, C.S., Yeganeh, S.H., Swett, I., Iyengar, J., Vasiliev, V., Jacobson, V.: BBR Congestion Control: IETF 100 Update: BBR

- in shallow buffers. Tech. rep., Internet Engineering Task Force (IETF) (Nov 2017), <https://www.ietf.org/proceedings/100/slides/slides-100-iccr-g-a-quick-bbr-update-bbr-in-shallow-buffers-00.pdf>
11. Cardwell, N., Cheng, Y., Gunn, C.S., Yeganeh, S.H., Swett, I., Iyengar, J., Vasiliev, V., Jha, P., Seung, Y., Jacobson, V.: BBR Congestion Control Work at Google IETF 101 Update. Tech. rep., Internet Engineering Task Force (IETF) (Mar 2018), <https://datatracker.ietf.org/meeting/101/materials/slides-101-iccr-g-an-update-on-bbr-work-at-google-00>
 12. Cardwell, N., Cheng, Y., Yang, K., Morley, D., Yeganeh, S.H., Jha, P., Seung, Y., Jacobson, V., Swett, I., Wu, B., Vasiliev, V.: BBR Congestion Control: Fundamentals and Updates. <https://research.cec.sc.edu/files/cyberinfra/files/BBR%20-%20Fundamentals%20and%20Updates%202023-08-29.pdf> (August 2023)
 13. Cardwell, N., Cheng, Y., Yeganeh, S.H., Swett, I., Jacobson, V.: BBR Congestion Control. Internet-Draft draft-cardwell-iccr-g-bbr-congestion-control-02, Internet Engineering Task Force (IETF) (Mar 2022), <https://datatracker.ietf.org/doc/draft-cardwell-iccr-g-bbr-congestion-control/02/>
 14. Cardwell, N., Cheng, Y., Yeganeh, S.H., Swett, I., Vasiliev, V., Bin Wu, M.M., Jha, P., Seung, Y., Jacobson, V.: BBR v2: A Model-based Congestion Control IETF 105 Update. Tech. rep., Internet Engineering Task Force (IETF) (Jul 2019), <https://datatracker.ietf.org/meeting/105/materials/slides-105-iccr-g-bbr-v2-a-model-based-congestion-control-00>
 15. Cardwell, N., Yuchung, C.: TCP BBR congestion control comes to GCP - your Internet just got faster. <https://cloud.google.com/blog/products/networking/tcp-bbr-congestion-control-comes-to-gcp-your-internet-just-got-faster/> (2017), [Last accessed: February 04, 2024]
 16. Chandrasekaran, B., Smaragdakis, G., Berger, A., Luckie, M., Ng, K.C.: A Server-to-Server View of the Internet. In: ACM CoNEXT (December 2015)
 17. Claypool, S.: Sharing but not Caring - Performance of TCP BBR and TCP CUBIC at the Network Bottleneck (March 2019)
 18. Claypool, S., Chung, J., Claypool, M.: Comparison of TCP Congestion Control Performance over a Satellite Network. In: Passive and Active Measurement (PAM) (2021)
 19. Corporation, I.: Intel® Tofino 6.4Tbps, 4 pipelines. <https://www.intel.com/content/www/us/en/products/sku/218643/intel-tofino-6-4-tbps-4-pipelines/specifications.html> (2017), [Last accessed: October 01, 2023]
 20. Dong, M., Meng, T., Zarchy, D., Arslan, E., Gilad, Y., Godfrey, B., Schapira, M.: PCC Vivace: Online-Learning Congestion Control. In: USENIX Symposium on Networked Systems Design and Implementation (NSDI) (2018)
 21. Fejes, F., Gombos, G., Laki, S., Nádas, S.: Who Will Save the Internet from the Congestion Control Revolution? In: Workshop on Buffer Sizing (2019)
 22. Google: TCP BBR v2 Alpha/Preview Release. <https://github.com/google/bbr/tree/v2alpha> (2021), [Last accessed: October 01, 2023]
 23. Google: TCP BBR v3 Release. <https://github.com/google/bbr/tree/v3> (2023), [Last accessed: October 01, 2023]
 24. Grazia, C.A., Patriciello, N., Klapez, M., Casoni, M.: A cross-comparison between TCP and AQM algorithms: Which is the best couple for congestion control? In: International Conference on Telecommunications (2017)
 25. Ha, S., Rhee, I., Xu, L.: CUBIC: a new TCP-friendly high-speed TCP variant. ACM SIGOPS Operating Systems Review **42**(5) (2008)
 26. Hock, M., Bless, R., Zitterbart, M.: Experimental evaluation of BBR congestion control. In: IEEE Int. Conference on Network Protocols (ICNP) (2017)

27. Hurtig, P., Haile, H.K., Grinnemo, K.J., Brunström, A., Atxutegi, E., Liberal, F., Arvidsson, Å.: Impact of TCP BBR on CUBIC Traffic: A Mixed Workload Evaluation. In: International Teletraffic Congress (2018)
28. Huston, G.: BBR TCP. <https://labs.ripe.net/author/gih/bbr-tcp/> (May 2017)
29. (ICSI), I.C.S.I.: Zeek: An Open Source Network Security Monitoring Tool. <http://www.zeek.org> (2022), [Last accessed: September 20, 2023]
30. Ivanov, A.: Evaluating BBRv2 on the Dropbox Edge Network. In: Netdev, The Technical Conference on Linux Networking (2020)
31. Jonathan Corbet: TCP small queues. <https://lwn.net/Articles/507065/> (2012), [Last accessed: October 01, 2023]
32. Kenjiro, C., Koushirou, M., Akira, K.: Traffic data repository at the WIDE project. In: USENIX Annual Technical Conference (ATC) (2000)
33. Kfoury, E.F., Gomez, J., Crichigno, J., Bou-Harb, E.: An emulation-based evaluation of TCP BBRv2 Alpha for wired broadband. *Computer Communications* (2020)
34. Kleinrock, L.: Power and deterministic rules of thumb for probabilistic problems in computer communications. In: IEEE International Conference on Communications (1979)
35. Li, Y., Miao, R., Liu, H.H., Zhuang, Y., Feng, F., Tang, L., Cao, Z., Zhang, M., Kelly, F., Alizadeh, M., Yu, M.: HPCC: High Precision Congestion Control. In: Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM) (2019)
36. Mishra, A., Sun, X., Jain, A., Pande, S., Joshi, R., Leong, B.: The Great Internet TCP Congestion Control Census. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* **3**(3) (2019)
37. Nandagiri, A., Tahiliani, M.P., Misra, V., Ramakrishnan, K.K.: BBRv1 vs BBRv2: Examining Performance Differences through Experimental Evaluation. *IEEE International Symposium on Local and Metropolitan Area Networks (LANMAN)* (2020)
38. Neal Cardwell: BBR evaluation with netem. https://groups.google.com/g/bbr-dev/c/8LYkNt17V_8 (2017), [Last accessed: October 01, 2023]
39. Netravali, R., Sivaraman, A., Das, S., Goyal, A., Winstein, K., Mickens, J., Balakrishnan, H.: Mahimahi: accurate Record-and-Replay for HTTP. In: USENIX Annual Technical Conference (ATC) (2015)
40. Philip, A.A., Ware, R., Athapathu, R., Sherry, J., Sekar, V.: Revisiting TCP Congestion Control Throughput Models & Fairness Properties at Scale. In: *ACM Internet Measurement Conference (IMC)* (2021)
41. Sandvine: The Global Internet Phenomena Report January 2023. Tech. rep., Sandvine Corporation (Jan 2023), https://www.sandvine.com/hubfs/Sandvine_Redesign_2019/Downloads/2023/reports/Sandvine%20GIPR%202023.pdf?hsCtaTracking=3cbc04da-fd44-481b-ad03-811d23b7b2c5%7C131df09f-dbdd-41a0-9dce-aac8879403ff
42. Scholz, D., Jaeger, B., Schwaighofer, L., Raumer, D., Geyer, F., Carle, G.: Towards a Deeper Understanding of TCP BBR Congestion Control. In: *IFIP Networking* (2018)
43. Shah, A.: BBR Evaluation at a Large CDN. <https://edg.io/tw/resources/technical-article/bbr-evaluation-at-a-large-cdn/> (September 2019)
44. Sommers, J., Kim, H., Barford, P.: Harpoon: A Flow-Level Traffic Generator for Router and Network Tests. *ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Science* **32**(1) (jun 2004)

45. Song, Y.J., Kim, G.H., Mahmud, I., Seo, W.K., Cho, Y.Z.: Understanding of BBRv2: Evaluation and Comparison With BBRv1 Congestion Control Algorithm. *IEEE Access* (2021)
46. Spang, B., Walsh, B., Huang, T.Y., Rusnock, T., Lawrence, J., McKeown, N.: Buffer sizing and Video QoE Measurements at Netflix. In: *Workshop on Buffer Sizing* (2019)
47. Stephen Hemminger and Fabio Ludovici and Hagen Paul Pfeifer: tc-netem(8) — Linux manual page. <https://man7.org/linux/man-pages/man8/tc-netem.8.html> (2023), [Last accessed: October 01, 2023]
48. Tahiliani, M.P., Misra, V.: A Principled Look at the Utility of Feedback in Congestion Control. In: *Workshop on Buffer Sizing* (2019)
49. Turkovic, B., Kuipers, F.A., Uhlig, S.: Interactions between Congestion Control Algorithms. In: *Network Traffic Measurement and Analysis Conference (TMA)* (2019)
50. Ware, R., Mukerjee, M., Seshan, S., Sherry, J.: Modeling BBR’s Interactions with Loss-Based Congestion Control. In: *ACM Internet Measurement Conference (IMC)* (2019)
51. Yan, F.Y., Ma, J., Hill, G.D., Raghavan, D., Wahby, R.S., Levis, P., Winstein, K.: Pantheon: the training ground for Internet congestion-control research. In: *USENIX Symposium on Networked Systems Design and Implementation (NSDI)* (2018)
52. Yang, F., Wu, Q., Li, Z., Liu, Y., Pau, G., Xie, G.: BBRv2+: Towards balancing aggressiveness and fairness with delay-based bandwidth probing. *Computer Networks* (2022)
53. Zapletal, A., Kuipers, F.: Slowdown as a Metric for Congestion Control Fairness. In: *SIGCOMM Workshop on Hot Topics in Networking (HotNets)*. *HotNets ’23* (2023)
54. Zeynali, D., N. Weyulu, E., Fathalli, S., Chandrasekaran, B., Feldmann, A.: Promises and Potential of BBRv3: Testbed configuration and evaluation artifacts. <https://dzeynali.gitlab.io/> (2024)
55. Zhang, H., Zhu, H., Xia, Y., Zhang, L., Zhang, Y., Deng, Y.: Performance Analysis of BBR Congestion Control Protocol Based on NS3. In: *International Conference on Advanced Cloud and Big Data (CBD)* (2019)
56. Zhang, Y., Breslau, L., Paxson, V., Shenker, S.: On the Characteristics and Origins of Internet Flow Rates. In: *Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)* (2002)