

Promises and Potential of **BBRv3**

Danesh Zeynali, **Emilia N. Weyulu**, Seifeddine Fathalli
Balakrishnan Chandrasekaran, Anja Feldmann



UNIVERSITÄT
DES
SAARLANDES

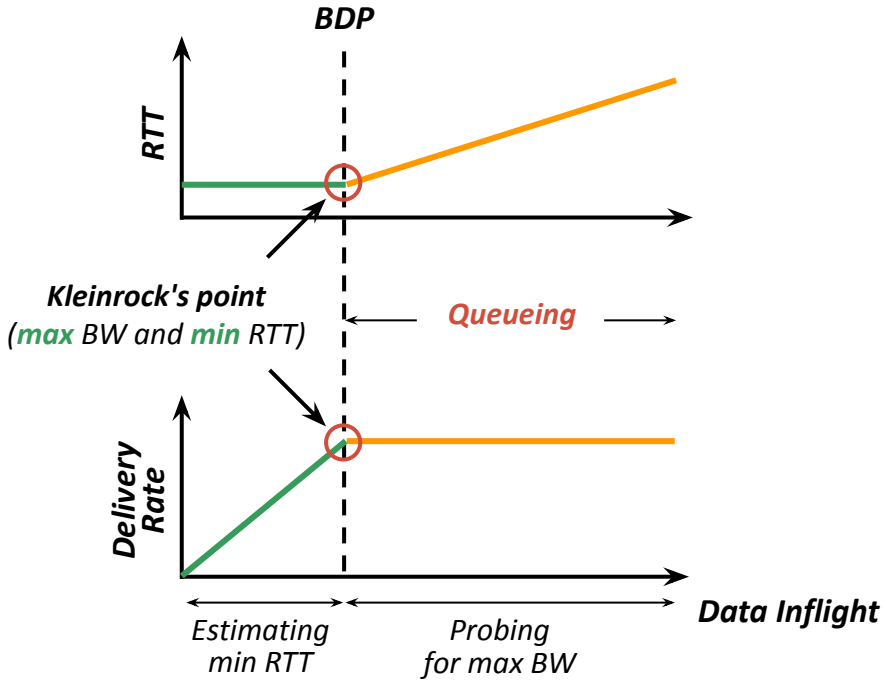


MAX-PLANCK-INSTITUT
FÜR INFORMATIK

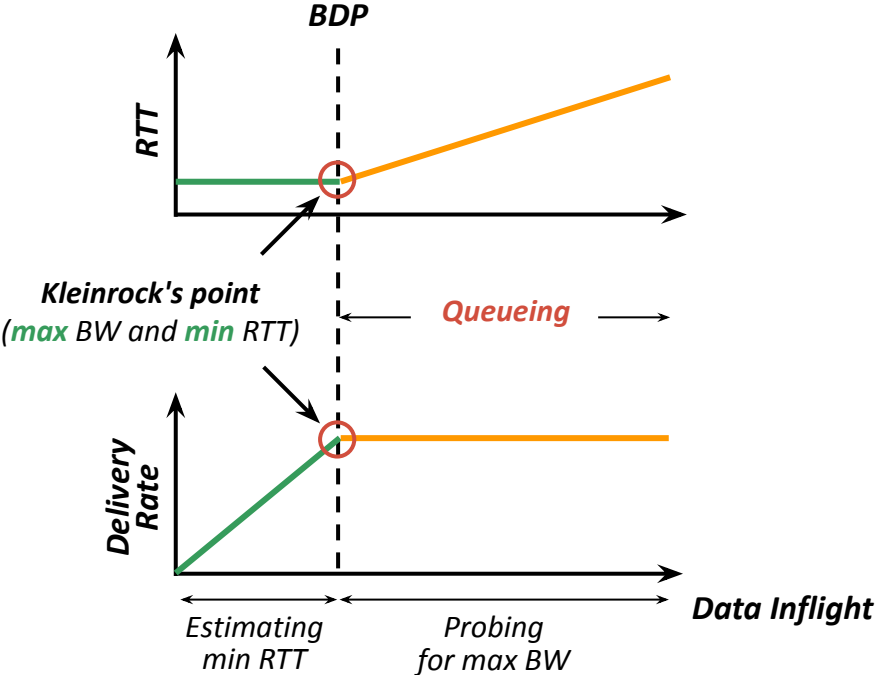


UNIVERSITY
AMSTERDAM

BBR: Bottleneck Bandwidth and Round-trip propagation time



BBR: Bottleneck Bandwidth and Round-trip propagation time



BBR: Bottleneck Bandwidth and Round-trip propagation time

Table 5: CCAs deployed by most popular websites on the Internet by traffic-share.

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

BBR: Bottleneck Bandwidth and Round-trip propagation time

Table 3: Classification results for websites by CDN websites. The values after the slashes are after a clustering step on traces within each CDN.

CDN	BBR	BIC	CDG	CUBIC	Highspeed	HTCP	NV	Reno	Vegas	Westwood	Yeah	Unknown	All Invalid	RTT > 85ms	Unresponsive	Total
Akamai	470/491	0	3/4	4	0	0	0	0	0	0	0	115/91	189	36	233	1050
Cloudflare	1233/1595	0	6/7	5/6	0	0	0	1	0	0	0	824/460	394	55	989	3507
Cloudfront	530/545	0	9/10	7/10	0	0	3/2	0	0	0	0	74/56	78	10	121	832
Fastly	21/25	1/13	3	25/130	0	0	0/1	0	0	0	0	174/52	26	3	30	283
Google	29	0	1	2	0	0	2	0	0	0	1	230	37	18	66	386
Other CDN	28/32	2/0	1	53/92	0	0	1	0	0	0	0	72/31	54	41	226	478
No CDN	116/122	3/0	8/9	89/116	3/5	2	5	4/3	1/0	3	0	146/115	127	1205	1752	3464
Total	2427/2839	6/13	31/35	185/360	3/5	2	11	5/4	1/0	3	1	1635/1035	905	1368	3417	10000

BBRv3: The **new** kid on the block (IETF 117)


Minor evolution of BBRv2 (IETF 105)

- Fix BW convergence with/without loss and/or ECN marks
- Performance tuning

BBRv3: The **new** kid on the block (IETF 117)

Minor evolution of BBRv2 (IETF 105)

- Fix BW convergence with/without loss and/or ECN marks
- Performance tuning



Promises

- Fast throughput convergence
- Better coexistence with CUBIC/Reno
- Lower retransmit rate
- Reduced latency for different buffer configurations

Verifying CCA Performance

Network Simulation



OMNeT++

Network Emulation

Mininet



GNS3®

Verifying CCA Performance


Network Simulation




Network Emulation



Lab Testbed




Network A/B tests




Verifying CCA Performance

What test scenarios to run?





Verifying CCA Performance

What test scenarios to run?




What parameters to use?





Verifying CCA Performance

What test scenarios to run?



What parameters to use?



- **Hard to reproduce earlier results**
- **CCA performance may not hold up in a different network**

Verifying CCA Performance

What test scenarios to run?



What parameters to use?



Modeling BBR's Interactions with Loss-Based Congestion Control

Ranysha Ware
rware@cs.cmu.edu
Carnegie Mellon University

Matthew K. Mukerjee
mukerjee@nefeli.io
Nefeli Networks

Srinivasan Seshan
srini@cs.cmu.edu
Carnegie Mellon University

Justine Sherry
sherry@cs.cmu.edu
Carnegie Mellon University

- Hard to find test scenarios that can differentiate between results
- CCA performance may not hold up in a different network

Verifying CCA Performance

What test scenarios to run?



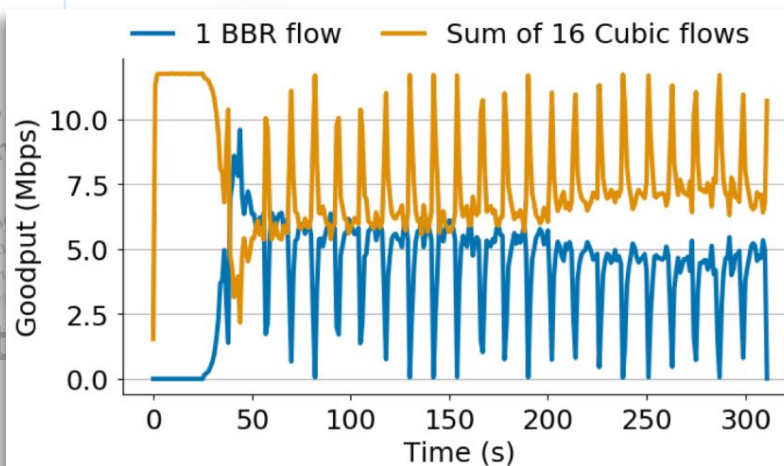
- Hard
- CCA performance may not hold

Ranysha Ware
rware@cs.cmu.edu
Carnegie Mellon University

Matthew K. Mukerjee
mukerjee@nefeli.io
Nefeli Networks

Srinivasan S
srini@cs.cmu.edu
Carnegie Mellon

What parameters to use?



Analysis of CCA Performance

Outline a systematic and principled approach to evaluate a CCA

Analysis of CCA Performance

Outline a systematic and principled approach to evaluate a CCA

Intra-CCA fairness




Analysis of CCA Performance

Outline a systematic and principled approach to evaluate a CCA

Intra-CCA fairness

RTT fairness



Analysis of CCA Performance

Outline a systematic and principled approach to evaluate a CCA

Intra-CCA fairness



RTT fairness

100 ms



5 ms



Coexistence with loss-based CCAs

BBRv3



Cubic




Analysis of CCA Performance

Outline a systematic and principled approach to evaluate a CCA


Intra-CCA fairness




RTT fairness




Coexistence with loss-based CCAs




Real-world network traffic



Evaluation: Testbed setup




Evaluation: Testbed setup



Is BBRv3 fair to itself?



Two competing flows: Deep buffer (16xBDP)

Simultaneous start





Two competing flows: Deep buffer (16xBDP)

Simultaneous start




Two competing flows: Deep buffer (16xBDP)

Simultaneous start





Staggered start





Two competing flows: Deep buffer (16xBDP)

Simultaneous start




Staggered start




Two competing flows: Deep buffer (16xBDP)

Simultaneous start



Staggered start



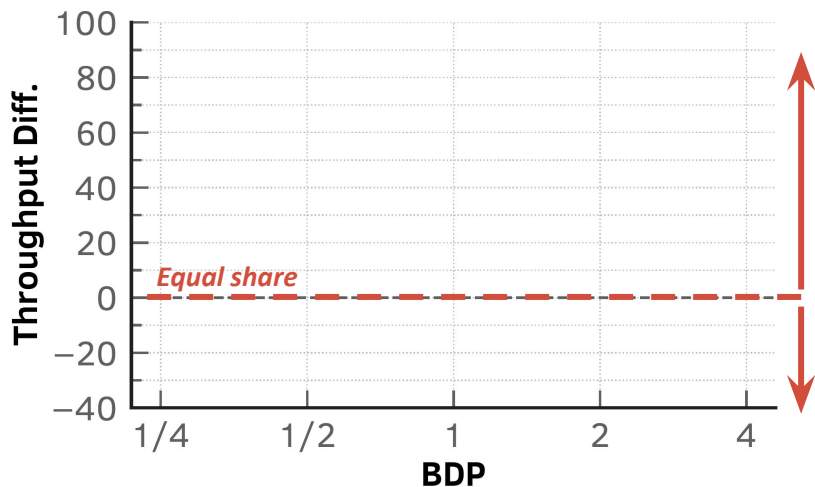
BBRv3: Convergence to an equitable share takes longer when flows do not start at the same time

Is BBRv3 flow fair to itself?

Flows with different RTTs

RTT fairness

160 vs. 10 ms

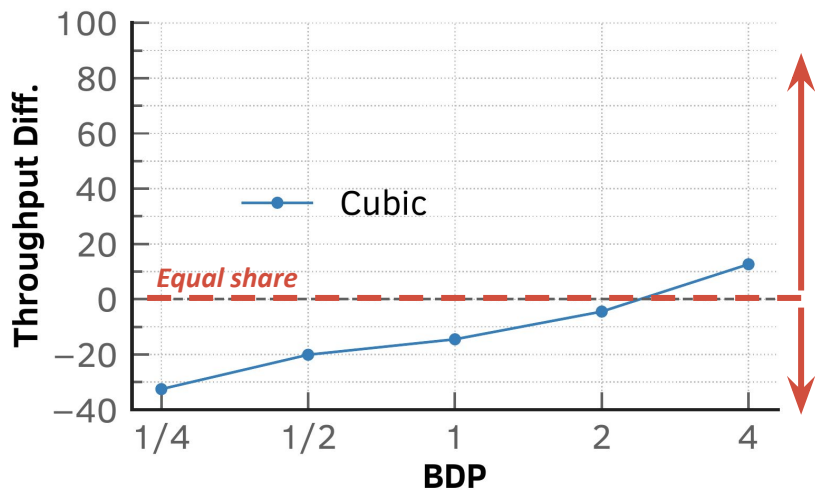


**Higher bandwidth share for the
long-RTT flow**

**Higher bandwidth share for the
short-RTT flow**

RTT fairness

160 vs. 10 ms

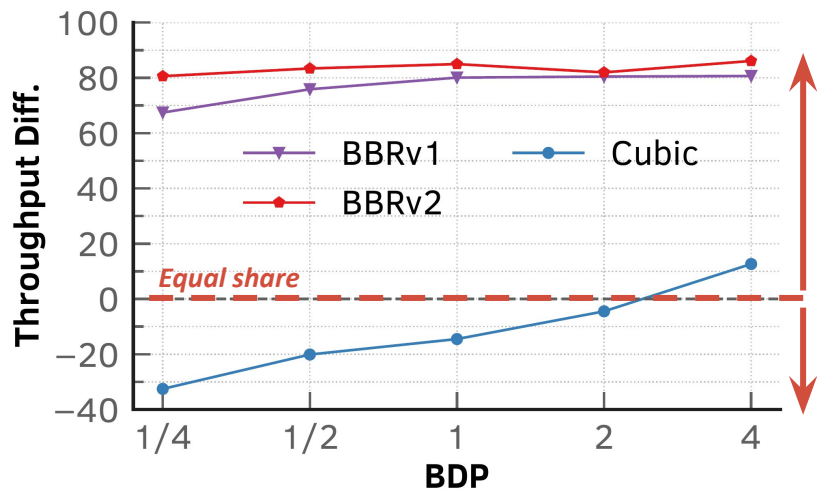


**Higher bandwidth share for the
long-RTT flow**

**Higher bandwidth share for the
short-RTT flow**

RTT fairness

160 vs. 10 ms

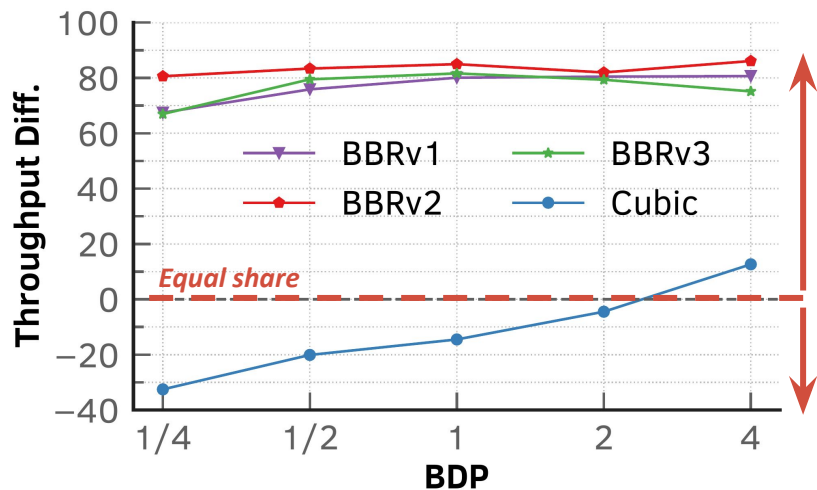


Higher bandwidth share for the *long-RTT* flow

Higher bandwidth share for the *short-RTT* flow

RTT fairness

160 vs. 10 ms




**Higher bandwidth share for the
long-RTT flow**

**Higher bandwidth share for the
short-RTT flow**

RTT fairness


160 vs. 10 ms



Jain's Fairness Index – Min. fairness is 0.5

RTT fairness


160 vs. 10 ms



Jain's Fairness Index – Min. fairness is 0.5

RTT fairness


160 vs. 10 ms




Jain's Fairness Index – Min. fairness is 0.5

RTT fairness

160 vs. 10 ms




160 vs. 80 ms




Jain's Fairness Index – Min. fairness is 0.5

RTT fairness

160 vs. 10 ms




160 vs. 80 ms




Jain's Fairness Index – Min. fairness is 0.5

RTT fairness

160 vs. 10 ms



160 vs. 80 ms




Bias towards long RTT flows persists in BBRv3

Is BBRv3 flow fair to itself?

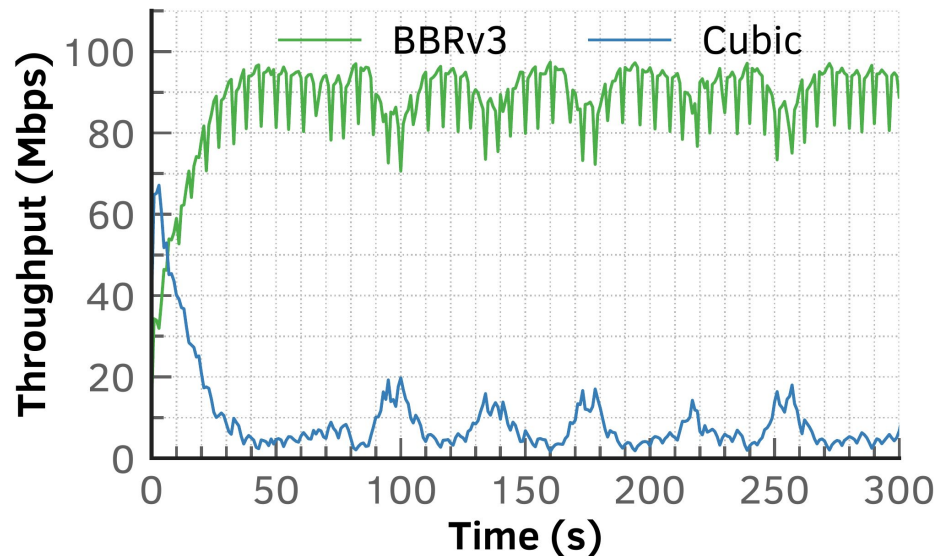

Flows with different RTTs

Does BBRv3 coexist well with Cubic?

Coexistence with Cubic: 1xBDP buffer




Coexistence with Cubic: 1xBDP buffer



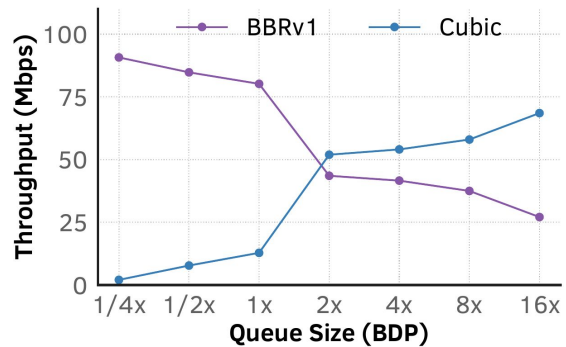
Coexistence with Cubic

BBRv1

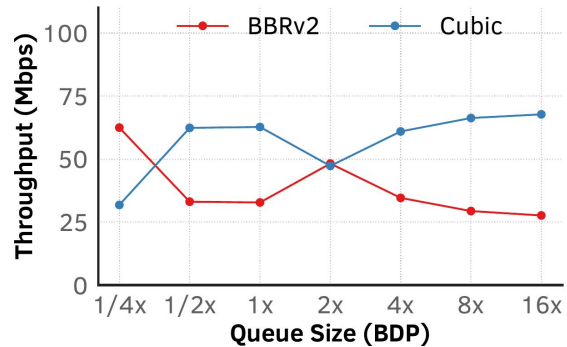


Coexistence with Cubic

BBRv1




BBRv2




Coexistence with Cubic

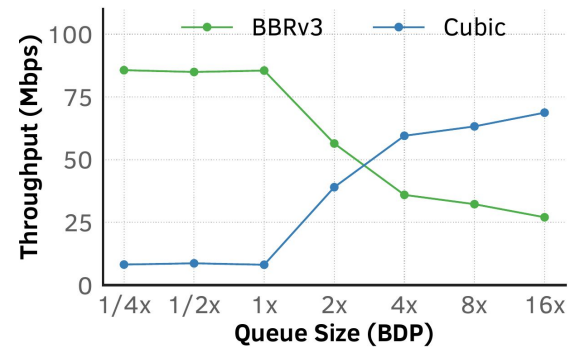
BBRv1



BBRv2




BBRv3




Coexistence with Cubic


BBRv1



BBRv2




BBRv3




BBRv3' unfairness towards Cubic is exacerbated at shallow buffer scenarios

Single BBR flow competing with multiple Cubic flows




Ware, R., et. al, : Modeling BBR's Interactions with Loss-Based Congestion Control. ACM IMC '19

Inter-CCA fairness: 1xBDP buffer



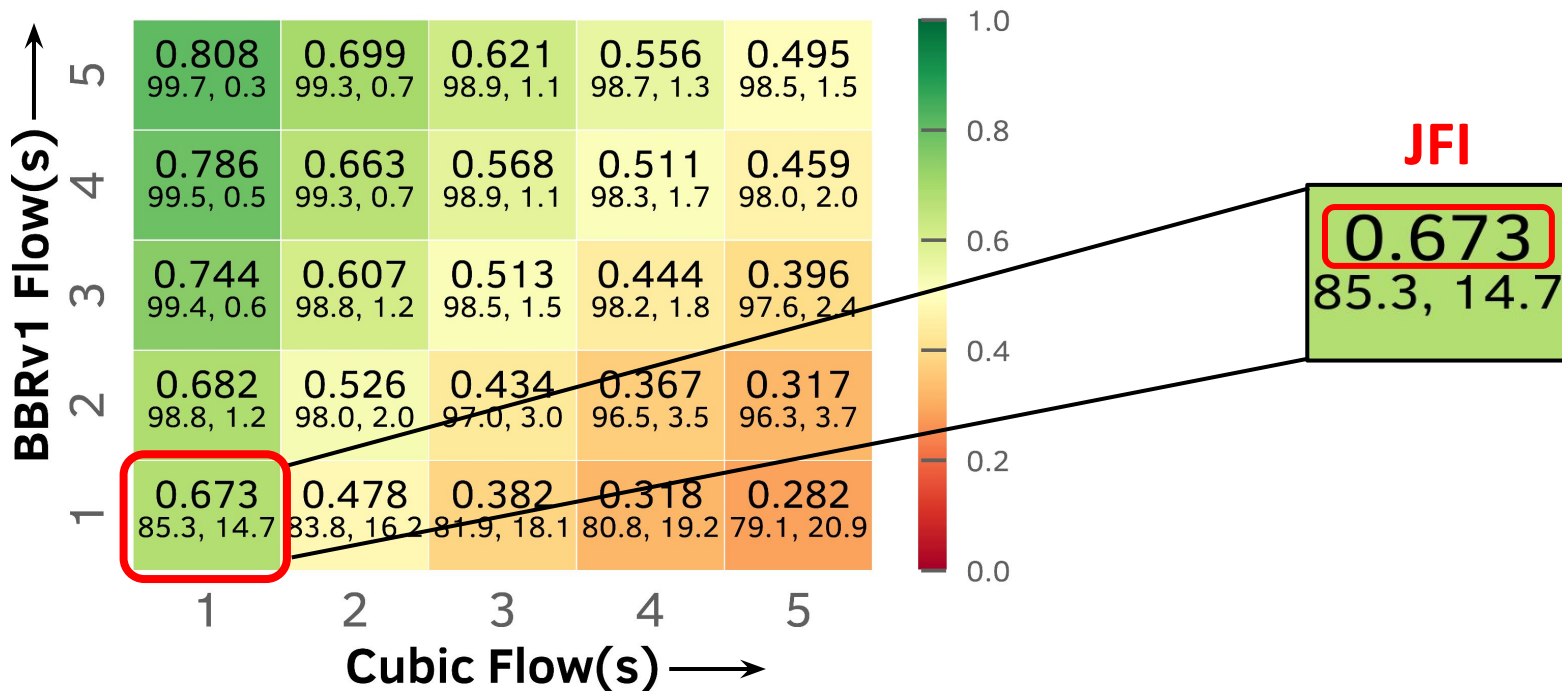
Inter-CCA fairness: 1xBDP buffer

BBRv1



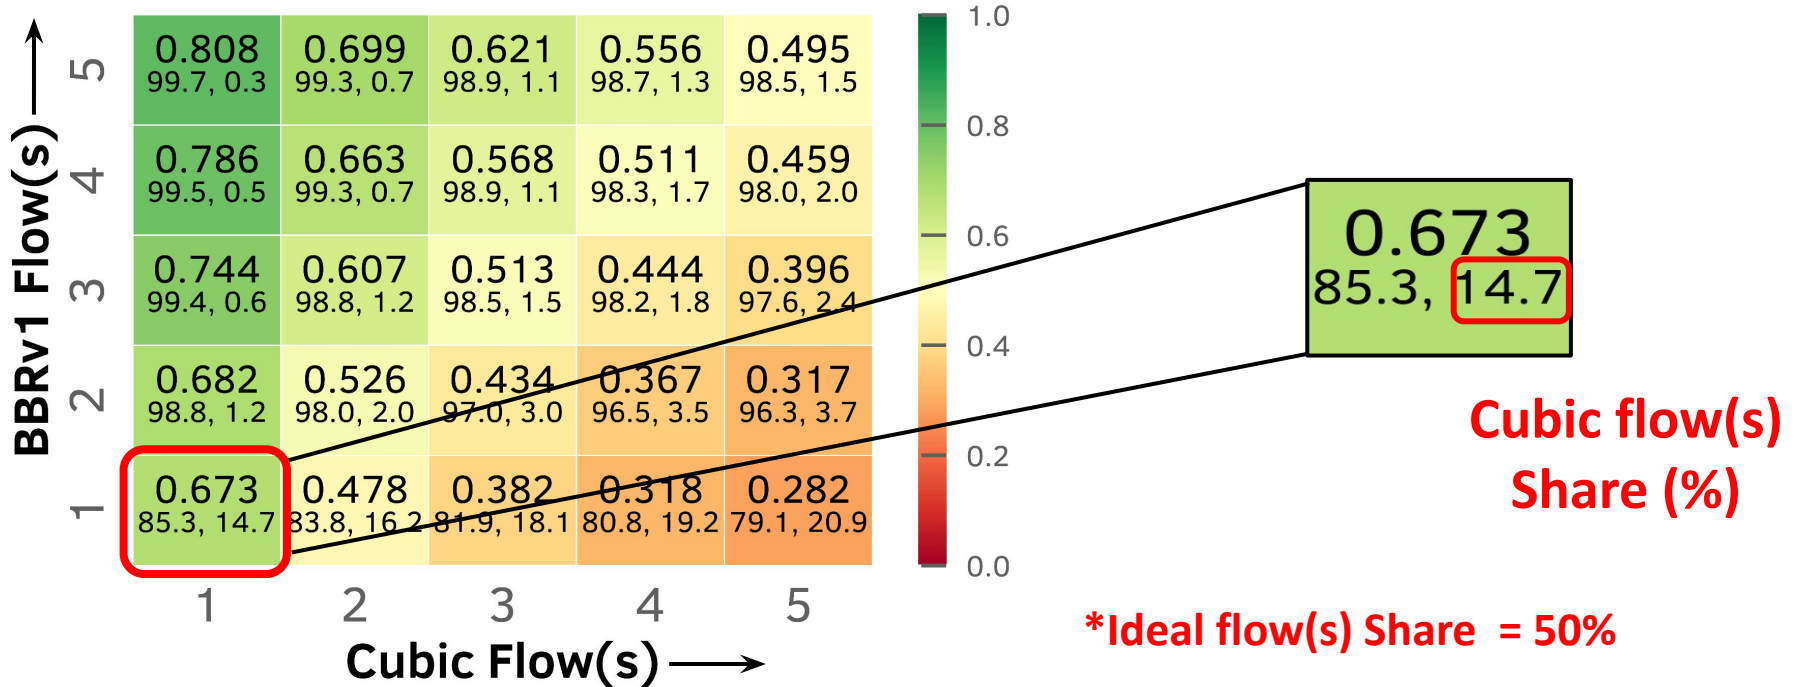
Inter-CCA fairness: 1xBDP buffer

BBRv1



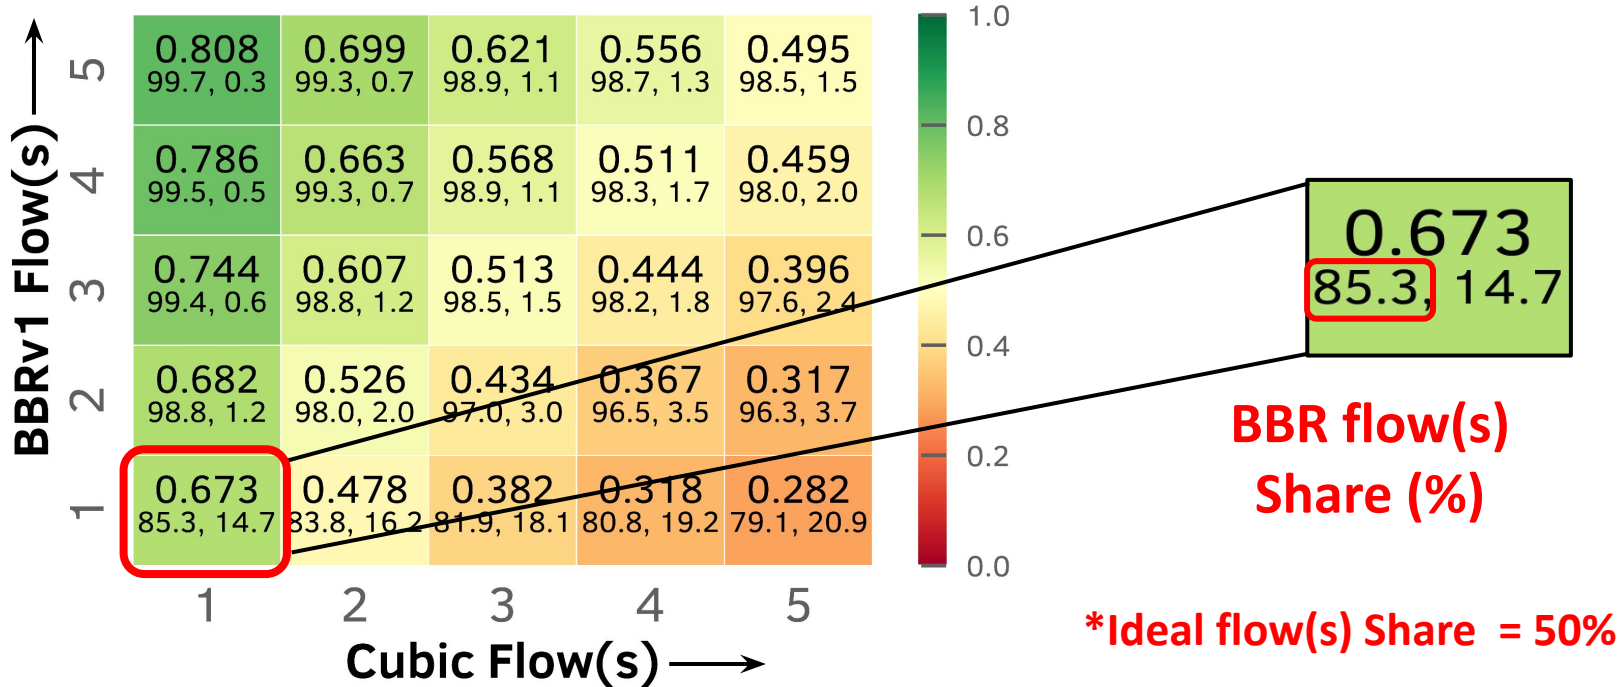
Inter-CCA fairness: 1xBDP buffer

BBRv1



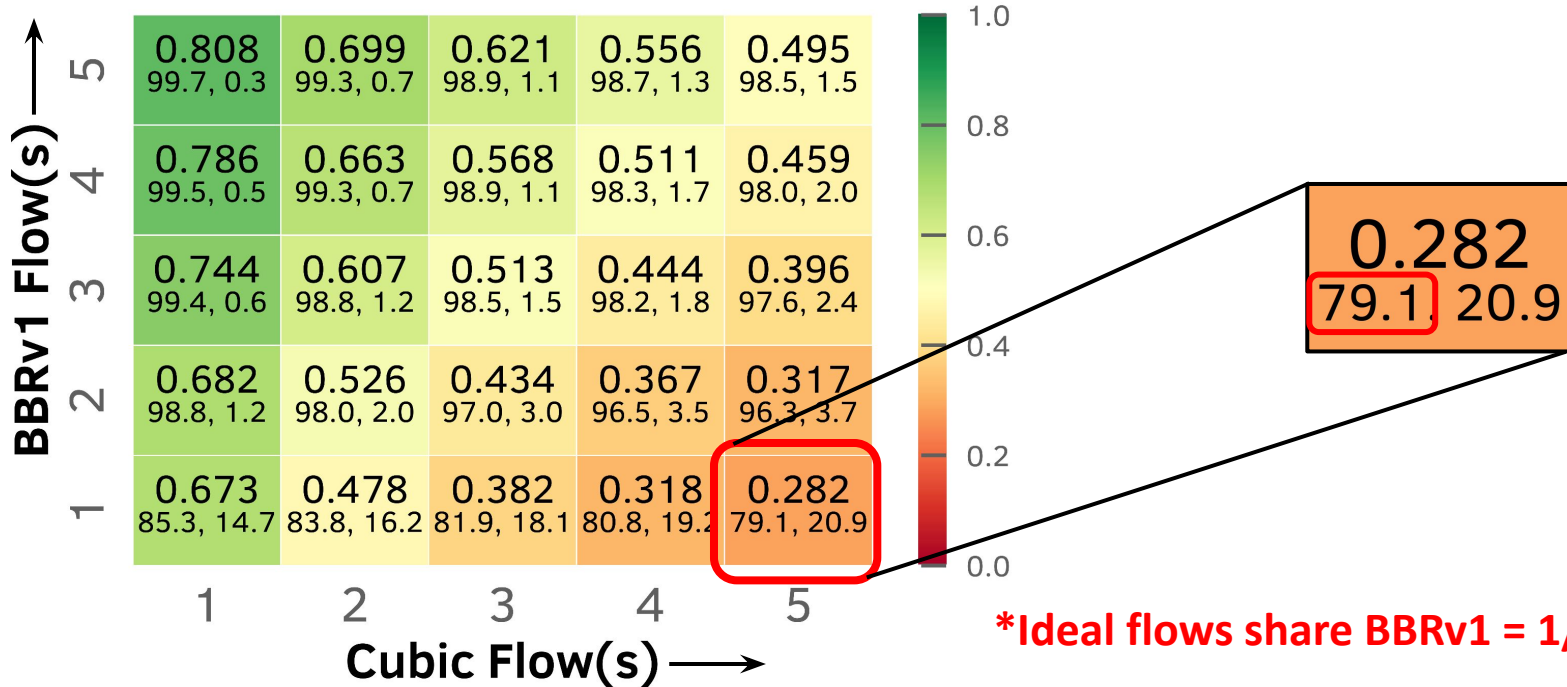
Inter-CCA fairness: 1xBDP buffer

BBRv1




Inter-CCA fairness: 1xBDP buffer

BBRv1




Inter-CCA fairness: 1xBDP buffer

BBRv1




Inter-CCA fairness: 1xBDP buffer

BBRv1



0.282
79.1, 20.9


BBRv2



0.944
21.8, 78.2


Inter-CCA fairness: 1xBDP buffer

BBRv1




0.282
79.1, 20.9

BBRv2



0.944
21.8, 78.2


BBRv3



0.430
61.8, 38.2


Inter-CCA fairness: 1xBDP buffer

BBRv1




0.808
99.7, 0.3

BBRv2



0.554
49.4, 50.6


BBRv3



0.808
96.0, 4.0


Inter-CCA fairness: 1xBDP buffer

BBRv1




0.808
99.7, 0.3

BBRv2



0.554
49.4, 50.6

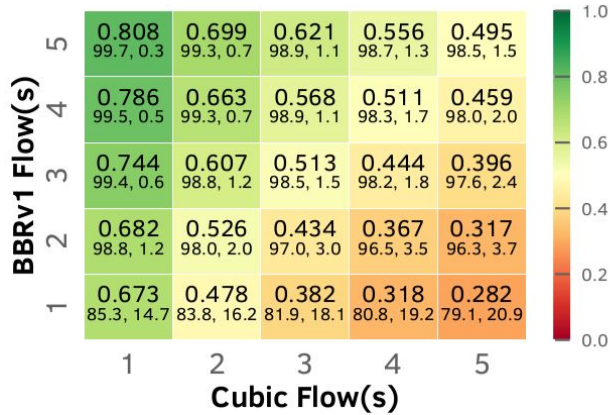
BBRv3



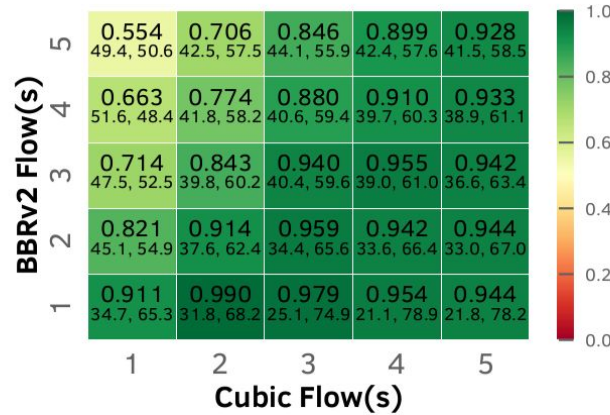
0.808
96.0, 4.0

Inter-CCA fairness: 1xBDP buffer


BBRv1



BBRv2




BBRv3




BBRv3 shows unfairness towards multiple Cubic flows as seen with BBRv1

Inter-CCA fairness: Enabling ECN, 1xBDP buffer



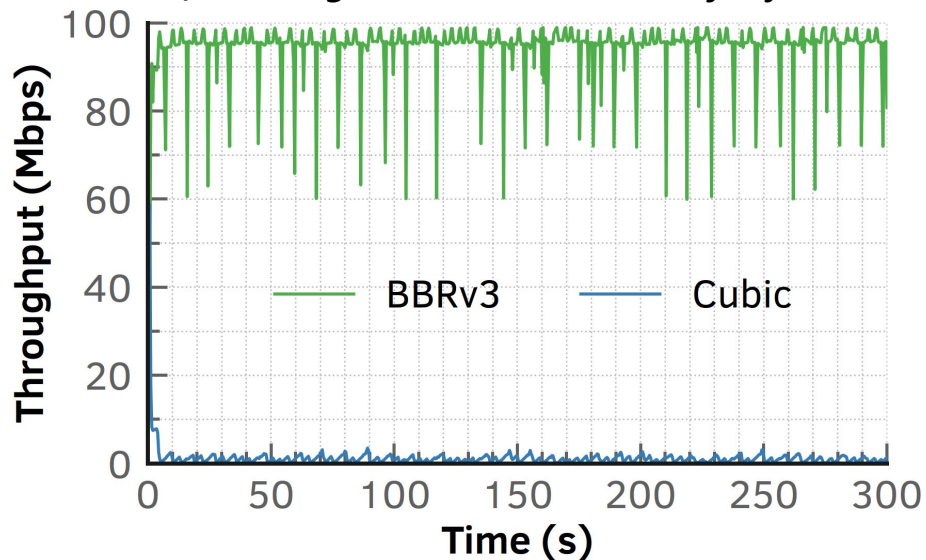
ECN disabled

Inter-CCA fairness: Enabling ECN, 1xBDP buffer



ECN disabled

CoDel AQM, Marking threshold = 24.2 ms of sojourn time



ECN enabled

Is BBRv3 flow fair to itself?

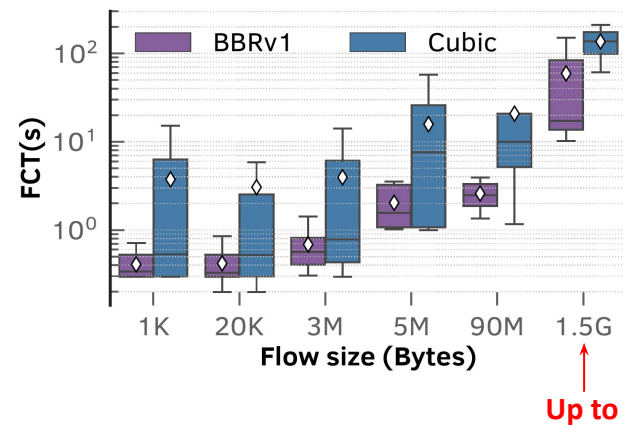
Flows with different RTTs

Does BBRv3 coexist well with Cubic?

BBRv3 with real-world network traffic

Inter-CCA workload


BBRv1 vs. Cubic



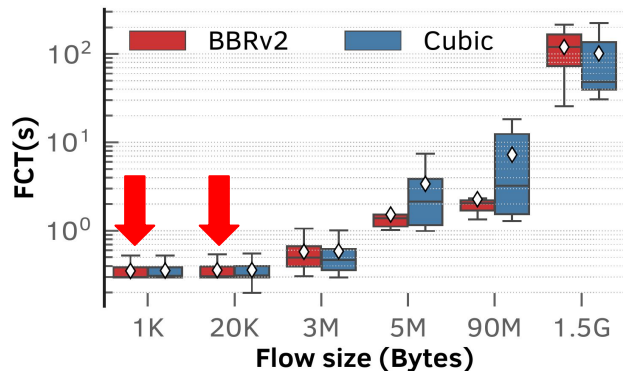
**Based on MAWI trace of 05 January 2023*

Inter-CCA workload

BBRv1 vs. Cubic




BBRv2 vs. Cubic




**Based on MAWI trace of 05 January 2023*

Inter-CCA workload

BBRv1 vs. Cubic




BBRv2 vs. Cubic




**Based on MAWI trace of 05 January 2023*

Inter-CCA workload


BBRv1 vs. Cubic



BBRv2 vs. Cubic



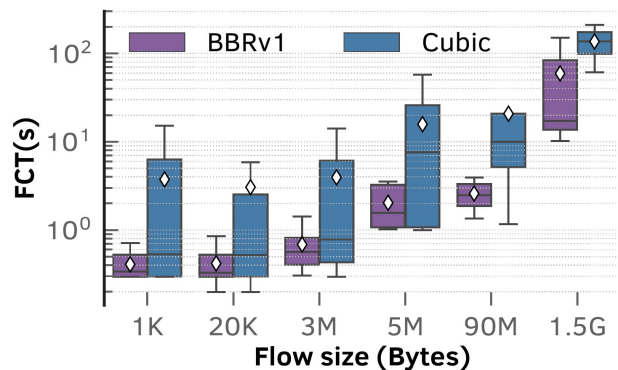
BBRv3 vs. Cubic




**Based on MAWI trace of 05 January 2023*

Inter-CCA workload

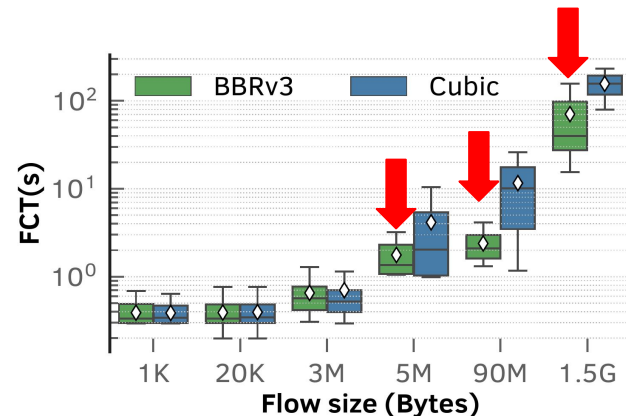
BBRv1 vs. Cubic



BBRv2 vs. Cubic




BBRv3 vs. Cubic



- BBRv2, BBRv3: Fairness improved for short flows
- For long flows, unfairness persists similar to BBRv1


Conclusion

Long Convergence time in deep buffer




Conclusion

Long Convergence time in deep buffer




Bias towards long-RTT flows




Conclusion




Long Convergence time in deep buffer



Bias towards long-RTT flows




Struggles to co-exist with loss-based CCAs





Conclusion

Long Convergence time in deep buffer

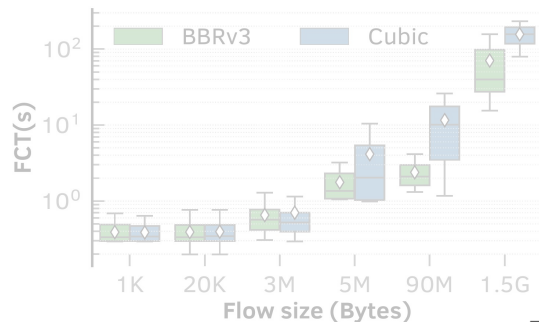




Bias towards long-RTT flows



inet-bbrv3eval.mpi-inf.mpg.de
tinyurl.com/bbrv3eval

Struggles to co-exist with loss-based CCAs



Backup Slides


Cubic Internet Traffic Share

Table 4: Distribution of CCA variants among Alexa Top 20k websites measured from different viewpoints by *Nebby*.


Variant	Ohio		Paris		Mumbai		Singapore		Sao Paulo	
	Websites	Share	Websites	Share	Websites	Share	Websites	Share	Websites	Share
BBRv1	2,594	13%	1,900	9.5%	1,635	8%	2,541	12.7%	1,280	6.4%
BBRv2	515	2.6%	373	1.9%	12	0.1%	251	1.3%	0	0%
BIC	712	3.5%	807	4%	837	4.2%	402	2%	227	1.1%
CUBIC	8,202	41%	8,406	42%	8,822	44.1%	8,673	43.4%	6,982	34.9%
HSTCP	0	0%	0	0%	0	0%	0	0%	0	0%
HTCP	583	2.9%	370	1.9%	522	2.6%	421	2.1%	363	1.8%
Illinois	721	3.6%	684	3.4%	1,121	5.6%	625	3.1%	229	1.1%
New Reno	1,840	9.2%	1,509	7.5%	3,032	15.2%	2,093	10.5%	2,683	13.5%
Vegas	878	4.4%	421	2.1%	301	1.5%	526	2.6%	511	2.5%
Veno	112	0.6%	382	1.9%	52	0.3%	21	0.1%	11	0.1%
Westwood	201	1%	170	0.9%	0	0%	0	0%	0	0%
Scalable	18	0.1%	0	0%	0	0%	6	0%	0	0%
YeAh	123	0.6%	89	0.4%	64	0.3%	0	0%	112	0.6%
Unknown	3,501	17.5%	4,889	24.4%	3,602	18.1%	4,441	22.2%	7,602	38%
Unresponsive	0	0%	0	0%	0	0%	0	0%	0	0%
Total	20,000	100%	20,000	100%	20,000	100%	20,000	100%	20,000	100%

Convergence: BBRv1, BBRv2, 16xBDP




BBRv1




BBRv2



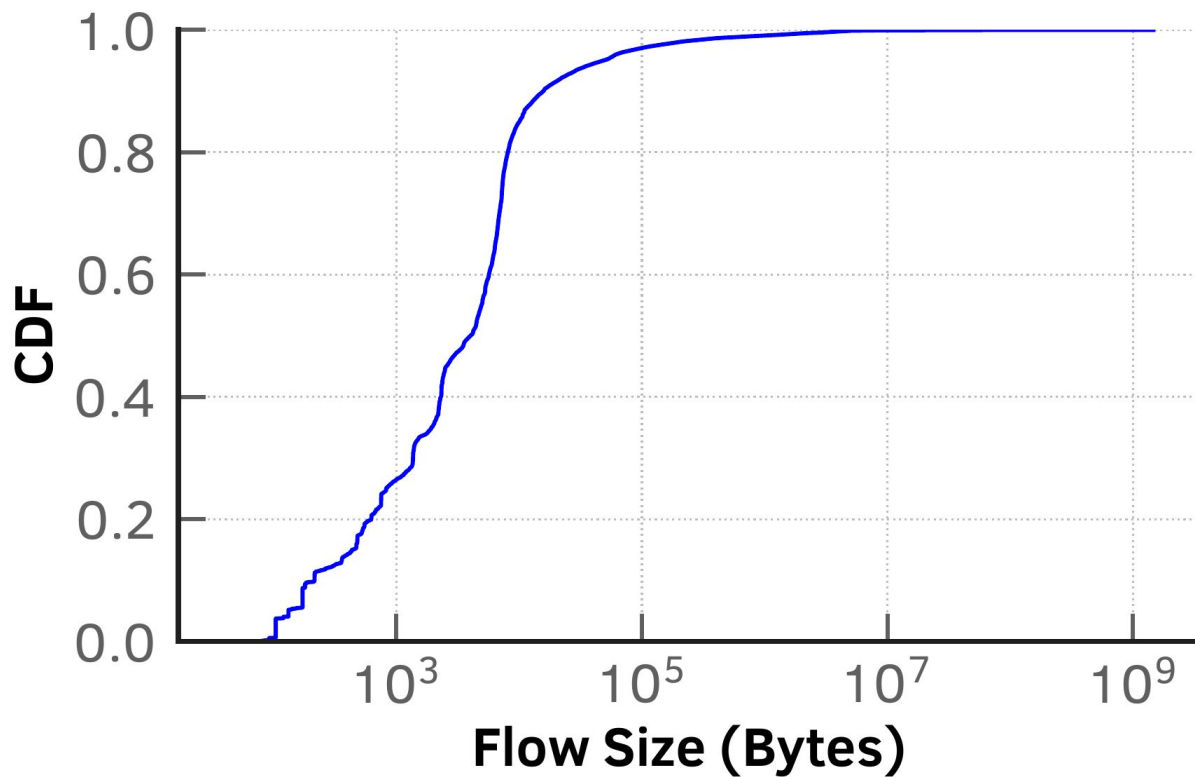
BBR inter-CCA fairness



BBRv3 inter-CCA fairness

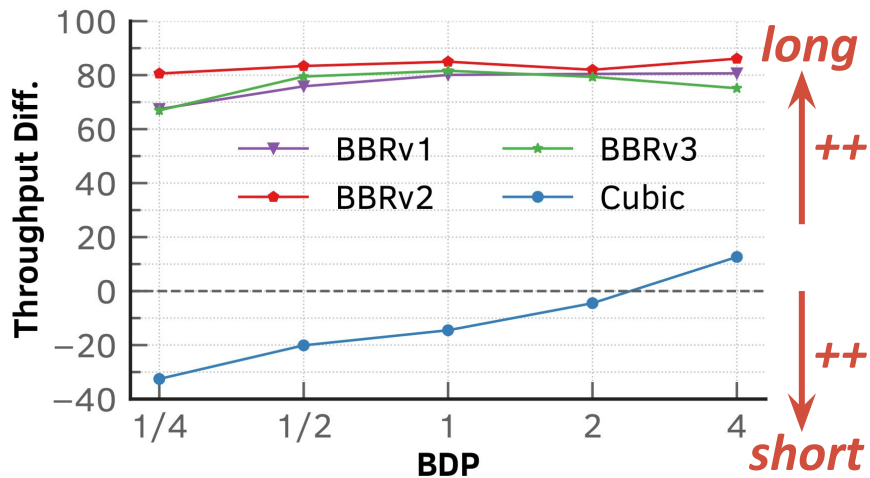


MAWI workload: 05 January 2023




RTT fairness: Long-RTT flow starts first

160 vs. 10ms

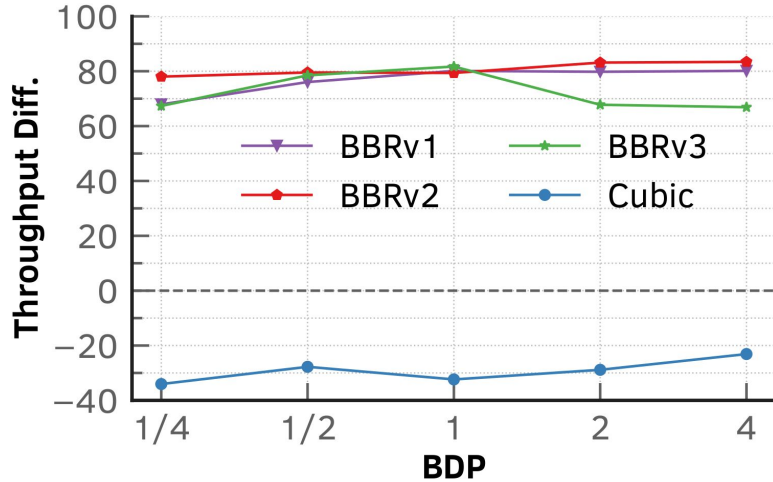


160 vs. 80ms




RTT fairness: short-RTT flow starts first

10 vs. 160ms




80 vs. 160ms




RTT fairness: short-RTT flow starts first

10 vs. 160ms




80 vs. 160ms



RTT fairness: short-RTT flow starts first

10 vs. 160ms



80 vs. 160ms

