

On the Suitability of BBR Congestion Control for QUIC over GEO SATCOM Networks



Aitor Martin

Naeem Khademi

Department of Electrical Engineering and Computer Science

University of Stavanger, Norway



Introduction





- Geosynchronous Satellite Communication (GEO SATCOM) networks are becoming popular candidates for providing broadband Internet connectivity in novel 5G/6G use-cases
- In parallel, we are witnessing major breakthrough on the Internet transport layer;
 - 1) standardization and deployment of QUIC, a general-purpose transport protocol that is
 - a) fully encrypted
 - b) deployed on the user space over UDP
 - 2) development of modern congestion control (CC), i.e. BBR, aiming to optimize bandwidth utilization while minimizing network latency
- To this date, around 8,3% of websites are already using QUIC [1]. Therefore, we expect an increase of QUIC traffic over SATCOM

Problem statement





mainly caused by

- 1. long propagation delay
 - 2. propagation errors
- 3. bandwidth asymmetry

HOW TO MITIGATE THE IMPACT OF THESE CHALLENGES?

TCP traffic is usually optimized with **Performance-Enhancing Proxies (PEPs)**

However, **QUIC's** full encryption disables PEP optimizations!

Several studies have shown that **TCP-PEP outperforms QUIC greatly** [2], even with QUIC's fast handshake

Raised interest in boosting QUIC performance over SATCOM through protocol mechanisms We investigate the use of BBR congestion control

[2] Nicolas Kuhn et al., 2020. QUIC: Opportunities and threats in SATCOM. In 2020 10th Advanced Satellite Multimedia Systems Conference and the 16th Signal Processing for Space Communications Workshop (ASMS/SPSC)



In this work, we investigate the following items:

- 1. the **general performance** of QUIC with BBRv2 over GEO SATCOM
- 2. other important aspects of **CC performance** (e.g., intra- and inter-protocol fairness, latecomer fairness, etc.) over long-haul satellite links
- 3. the impact of **packet loss** and **bandwidth asymmetry**
- 4. the impact of **QUIC implementation choice**





Background

Transport Layer over SATCOM

Challenges introduced by SATCOM links

1. Long Round-Trip Time (RTT) ~ 600 ms

- Long protocol feedback
 - Slower CC convergence
 - Delayed loss detection and recovery
- High Bandwidth-Delay Product (BDP) \rightarrow Larger buffers needed

2. Propagation errors due to e.g. rain fading

• Can be problematic for loss-based CC (e.g. NewReno or CUBIC)

3. Bandwidth asymmetry

• ACK congestion in the return path can limit forward throughput



Transport Layer over SATCOM

Challenges introduced by SATCOM links

1. Long Round-Trip Time (RTT) ~ 600 ms

- Long protocol feedback
 - Slower CC convergence
 - Delayed loss detection and recovery
- High Bandwidth-Delay Product (BDP) \rightarrow Larger buffers needed

2. Propagation errors due to e.g. rain fading

• Can be problematic for loss-based CC (e.g. NewReno or CUBIC)

3. Bandwidth asymmetry

Satellite-optimized CC (larger IW, faster slow start)

Proposed Solutions

BDP Frame extension [3]

FEC for QUIC

Model-based CC (e.g. BBR)

ACK Frequency Extension [4]



Challenges introduced by SATCOM links

1. Long Round-Trip Time (RTT) ~ 600 ms

- Long protocol feedback
 - Slower CC convergence
 - Delayed loss detection and recovery
- High Bandwidth-Delay Product (BDP) \rightarrow Larger buffers needed

2. Propagation errors due to e.g. rain fading

• Can be problematic for loss-based CC (e.g. NewReno or CUBIC)

3. Bandwidth asymmetry

• ACK congestion in the return path can limit forward throughput

[3] Nicolas Kuhn et al. BDP Frame Extension. Internet-Draft draft-kuhn-quic-bdpframe-extension-00, IETF, March 2022[4] Jana Iyengar and Ian Swett. QUIC Acknowledgement Frequency. Internet-Draft draft-ietf-quic-ack-frequency-01, IETF, October 2021.



Satellite-optimized CC (larger

IW, faster slow start)

Proposed Solutions

BDP Frame extension [3]

FEC for QUIC

Model-based CC (e.g. BBR)









- The <u>goal of BBR</u>: find **optimal pacing rate** to maximize link utilization while keeping path RTT as low as possible
 - How? Measuring path Bottleneck Bandwidth and RTT.
- Previous studies show that BBR tends to beat CUBIC over lossy paths [5]

However, three main issues were found with the first version of BBR [6]

- 1. Unfairness between parallel BBR flows
- 2. Aggresiveness against parallel loss-based CC flows
- 3. RTT unfairness





2019 - an <u>update to BBR is proposed</u>, named **BBRv2**, introducing a more complex bandwidth probing mechanism, aiming to solve the previously mentioned issues

BBRv2 also reacts to packet loss and ECN



Related studies



Year	Study	Transport	Scenario	Congestion Control		
2016	Cardwell et al.		Torrostrial			
2017	Hock et al.	ТСР				
2018	Scholz et al.	ICP	Terrestria	BBRVI, CUBIC		
2019	Jaeger et al.					
2020	Gomez et al.		Terrestrial			
2020	Song et al.	ТСР				
2021	Song et al.			DDRVZ, DDRVI, CUDIC		
2022	Yang et al.					
2021	Claypool et al.	ТСР	SATCONA	BBRv1, CUBIC		
2022	Zhao et al.	ICP	SATCOIVI	BBRv1, CUBIC, PCC, Hybla		
2018	Wang et al.	QUIC	SATCOM	BBRv1, CUBIC		
2022	Our paper	QUIC	SATCOM	BBRv2, BBRv1, CUBIC		

A high amount of QUIC implementations out there, developed by different agents:

- Web and CDN service providers, e.g. LiteSpeed, Akamai, Cloudflare
- Big technological companies, e.g. Google, Facebook, Apple, Microsoft
- IETF/IRTF contributors, e.g., ngtcp2, picoquic



- Robin Marx et al. (2020) [7], Sebastian Endres et al. (2022) [8] have reported high heterogeneity among implementations
- Even though RFC9002 specifies a CC mechanism similar to NewReno for QUIC, many stacks implement CUBIC and BBR

[7] Robin Marx et.al. 2020. Same Standards, Different Decisions: A Study of QUIC and HTTP/3 Implementation Diversity. en. In Proceedings of the Workshop on the Evolution, Performance, and Interoperability of QUIC. ACM, Virtual Event USA, (August 2020)
[8] Sebastian Endres et.al. 2022. Performance of QUIC Implementations Over Geostationary Satellite Links





Experimental Testbed Setup







UiS Testbed based on $\ensuremath{\mathsf{TEACUP}}$

Flexible experiment design and test automation

For this work, TEACUP was extended for QUIC and QLOG support

QUIC clients

QUIC servers



- Link emulation with **netem/tc**

	SAT	TERR
One Way Delay (OWD)	300 ms	50 ms
Downlink Bandwidth	20 Mbps	20 Mbps
Uplink Bandwidth	20/2 Mbps	20 Mbps
Bottleneck Buffer Size	0.25 0.5 1.0	2.0 x BDP
Packet Loss Ratio (PLR)	0%, 0.1	%, 1%

- Two QUIC implementations:
 - ngtcp2 allows to experiment with BBRv2
 - picoquic reported good performance over SATCOM [8]

of Stavange



Results





We set up four different experimental scenarios:

- 1. Single-Flow Bulk Download
- 2. Multi-Flow Fairness
- 3. Latecomer Issue
- 4. Mice versus Elephant flows

All the experiments are run <u>10 times</u>



We set up four different experimental scenarios:

- 1. Single-Flow Bulk Download
- 2. Multi-Flow Fairness
- 3. Latecomer Issue

Results

4. Mice versus Elephant flows



Run duration: 120 seconds Implementations: ngtcp2 and picoquic

Results: Bulk Download



Symmetric link (20/20), no packet loss



Results: Bulk Download



Symmetric link (20/20), no packet loss



Results: Bulk Download with packet loss

 \rightarrow ngtcp2 - bbr2 \leftrightarrow ngtcp2 - bbr1 # ngtcp2 - cubic \leftrightarrow picoquic - bbr1 # picoquic - cubic



Results: Bulk Download with packet loss

 \rightarrow ngtcp2 - bbr2 \leftrightarrow ngtcp2 - bbr1 # ngtcp2 - cubic \leftrightarrow picoquic - bbr1 # picoquic - cubic



Results: Bulk Download with packet loss

 \rightarrow ngtcp2 - bbr2 \leftrightarrow ngtcp2 - bbr1 # ngtcp2 - cubic \leftrightarrow picoquic - bbr1 # picoquic - cubic



Results: Bulk Download with uplink traffic



Now, we introduce **cross-traffic in the uplink**, and measure forward goodput over symmetric and asymmetric bandwidth SATCOM setups

ngtcp2	Symmetric 20/20			Asymmetric 20/2		
Down Up	BBRv2	BBRv1	CUBIC	BBRv2	BBRv1	CUBIC
BBRv2	14.73	14.94	13.56	5.31	8.74	6.36
BBRv1	15.10	14.97	14.66	8.03	8.25	5.25
CUBIC	14.90	13.64	11.70	7.38	8.64	7.74
picoquic	Symmetric 20/20			Asymmetric 20/2		
Down Up	BBRv2	BBRv1	CUBIC	BBRv2	BBRv1	CUBIC
BBRv2	-	-	-	-	-	-
BBRv1	-	18.36	18.34	-	18.35	18.35
CUBIC	-	18.38	18.25	-	18.34	18.35

Average Forward Goodput (Mbps)

Results: Bulk Download with uplink traffic



Now, we introduce **cross-traffic in the uplink**, and measure forward goodput over symmetric and asymmetric bandwidth setups.

ngtcp2	Symmetric 20/20			Asymmetric 20/2		
Down Up	BBRv2	BBRv1	CUBIC	BBRv2	BBRv1	CUBIC
BBRv2	14.73	14.94	13.56	5.31	8.74	6.36
BBRv1	15.10	14.97	14.66	8.03	8.25	5.25
CUBIC	14.90	13.64	11.70	7.38	8.64	7.74
picoquic	Symmetric 20/20			Asymmetric 20/2		
Down Up	BBRv2	BBRv1	CUBIC	BBRv2	BBRv1	CUBIC
BBRv2	-	-	-	-	-	-
BBRv1	-	18.36	18.34	-	18.35	18.35
CUBIC	-	18.38	18.25	-	18.34	18.35

Average Forward Goodput (Mbps)

ngtcp2 performance drops on asymmetric links

Results: Bulk Download with uplink traffic



Now, we introduce **cross-traffic in the uplink**, and measure forward goodput over symmetric and asymmetric bandwidth setups.

ngtcp2	Symmetric 20/20		Asymmetric 20/2			
Down Up	BBRv2	BBRv1	CUBIC	BBRv2	BBRv1	CUBIC
BBRv2	14.73	14.94	13.56	5.31	8.74	6.36
BBRv1	15.10	14.97	14.66	8.03	8.25	5.25
CUBIC	14.90	13.64	11.70	7.38	8.64	7.74
picoquic	Symmetric 20/20			Asymmetric 20/2		
Down Up	BBRv2	BBRv1	CUBIC	BBRv2	BBRv1	CUBIC
BBRv2	-	-	-	-	-	-
BBRv1	-	18.36	18.34	-	18.35	18.35
CUBIC	-	18.38	18.25	-	18.34	18.35

Average Forward Goodput (Mbps)

picoquic remains stable on asymmetric links



University of Stavanger

We set up four different experimental scenarios:

- 1. Single-Flow Bulk Download
- 2. Multi-Flow Fairness
- 3. Latecomer Issue

Results

4. Mice versus Elephant flows

We use Jain's Fairness Index (JFI) to measure fairness between parallel flows



Flows 2,4,6... (even flows)

Run duration: 300 seconds Implementation: ngtcp2

Results: Intra-protocol fairness



Results: Inter-protocol fairness



Results: Inter-protocol fairness





Results: Inter-protocol fairness



Average goodput ratio achieved by each flow





We set up four different experimental scenarios:

- 1. Single-Flow Bulk Download
- 2. Multi-Flow Fairness
- 3. Latecomer Issue

Results

4. Mice versus Elephant flows



Run duration: 300 seconds Implementation: ngtcp2



CUBIC latecomers converge very slowly on long RTT paths

CUBIC

SATCOM (RTT = 600 ms)

-O-Flow 1 ---- Flow 2 ---- Flow 3 ---- Aggregate



BBR latecomers converge faster!

SATCOM (RTT = 600 ms)

-O-Flow 1 -Flow 2 -Flow 3 -Flow 4 ----- Aggregate





BBRv2 latecomers join the link less aggressively

SATCOM (RTT = 600 ms)

-O-Flow 1 -Flow 2 -Flow 3 -Flow 4 ----- Aggregate





BBRv2 achieves <u>better</u> <u>long-term fairness</u>

SATCOM (RTT = 600 ms)







We set up four different experimental scenarios:

client1

- 1. Single-Flow Bulk Download
- 2. Multi-Flow Fairness
- 3. Latecomer Issue

Results

4. Mice versus Elephant flows



Background traffic: elephant flow

Implementation: ngtcp2



CC: BBRv1, CUBIC

server1

Results: Mice vs Elephant Flows



Background Traffic BBRv1



Discussion and Conclusion





Impact of Congestion Control choice

- BBR provides better performance overall under lossy links

- BBRv1 provides the best performance
- BBRv2 provides better fairness towards itself and towards CUBIC
- BBRv2 latecomers are less aggressive and still converge fast
- BBRv2 seems to be the on the right path for fairer coexistence with other flows
 - But BBRv2 performance suffers from the long RTT + packet loss present in SATCOM links
 - Further BBR iterations could contemplate these long RTT scenarios





Impact of bandwidth asymmetry

- A **1:10 asymmetry** has proven to be a great challenge for ngtcp2, with great performance drops

- But results have shown that an ACK policy such as picoquic's can maintain performance
 - picoquic sends around 10 times less ACK frames in our experiment results
 - This stresses **further research into optimized ACK strategies for SATCOM networks**





- picoquic outperforms ngtcp2 across CC algorithms

- Better resilience to packet loss
- Better performance with bandwidth asymmetry
- Possible reasons
 - Flow control window mechanism
 - ACK policies





To summarize:

- **BBR** seems to be a good candidate to yield better performance SATCOM networks
 - BBRv2 adds great improvements to fairness
 - But BBRv2 fairness and performance could be further improved for SATCOM-like scenarios (i.e. high BDP and packet loss)
- **Bandwidth asymmetry** is a problem in the abscence of satellite-optimized ACK policies
- Picoquic's satelite-optimizations seem to be key for QUIC





Future directions:

- Improve the experimental setup:
 - Introduce a more realistic satellite model (L1-L2 mechanisms, packet loss models)
 - Use a wider set of QUIC implementations
- Propose and study different ACK strategies under various asymmetric setups
 - Could this be better implemented using MASQUE?
 - How do different CC coexist with these ACK policies?



Thank you! Questions?

Aitor Martin Naeem Khademi

Department of Electrical Engineering and Computer Science

University of Stavanger, Norway