QUIC FOR SATELLITE COMMUNICATIONS

JCSA - AFNIC

31/08/2021

Nicolas KUHN

1

QUIC FOR SATELLITE COMMUNICATIONS

INTRODUCTION TO SATELLITE COMMUNICATIONS

Myth #1: SATCOM systems are quite specific

Indeed:

- Limited frequency resource (regulation, etc.)
- Dish alignment
- No standards for network infrastructure (lack of interoperability)

BUT:

- High level architecture similar to other access networks

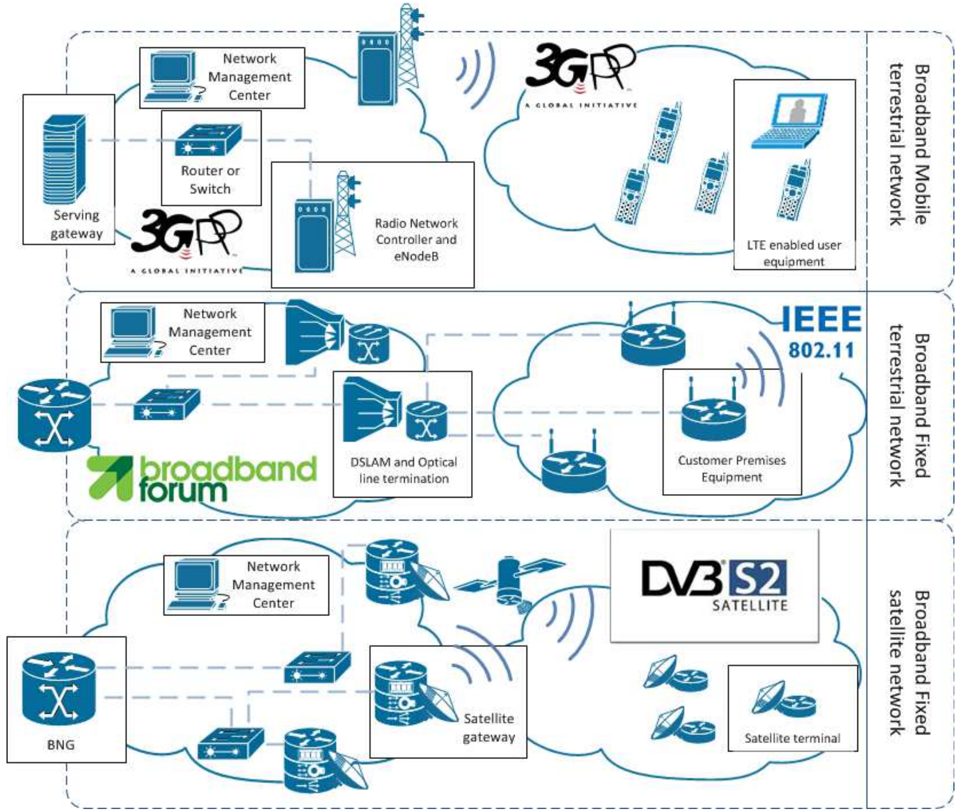

Myth #2: Latency is huge with SATCOM access

Indeed:

- For geostationary accesses, there is an important propagation delay (RTT of 500ms)

BUT:

- End-to-end latency is not just about signal propagation delay (e.g. Bufferbloat in cellular networks) [1]
- (honestly) it is not that bad

B. Briscoe; A. Brunstrom; A. Petlund; D. Hayes; D. Ros; I. J. Tsang; S. Gjessing; G. Fairhurst; C. Griwodz; M. Welzl, "Reducing Internet Latency: A Survey of Techniques and their Merits," in IEEE Communications Surveys & Tutorials

Myth #2: Latency is huge with SATCOM access

Light page – Wikipedia type

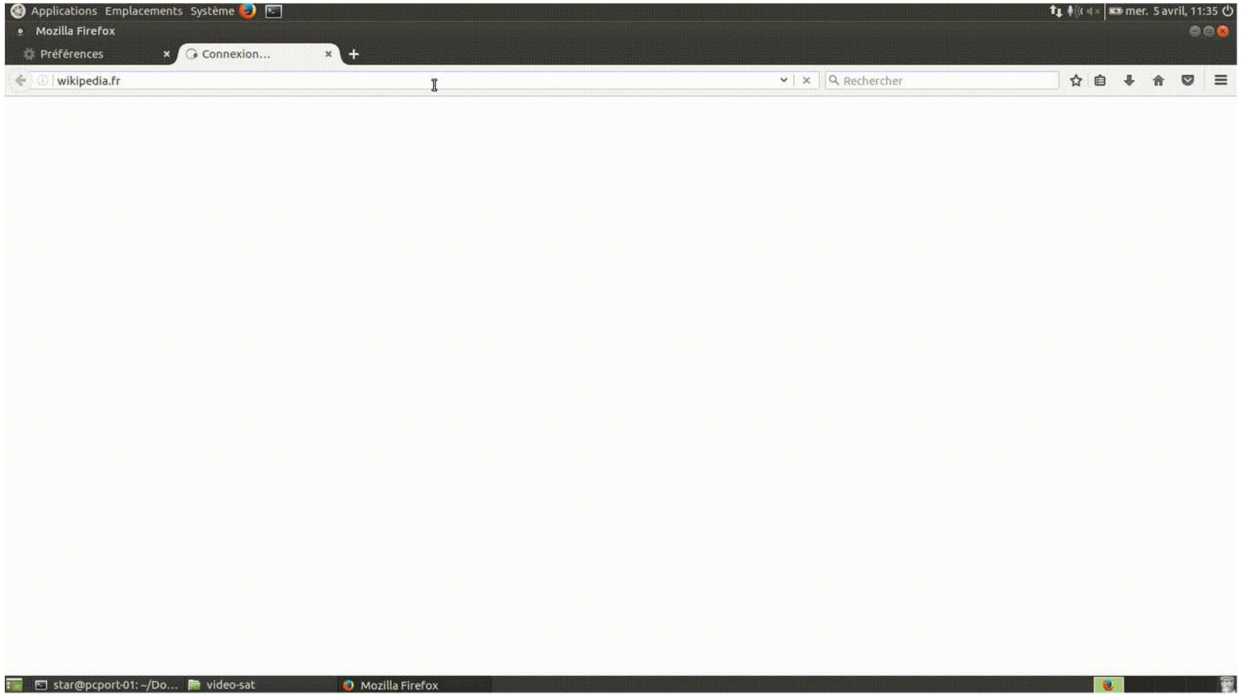

TOOWAY satellite Internet access :

- Solution furnished by ISP ALSATIS with EUTELSAT operator
- 20Mbps download / 6 Mbps upload

Myth #2: Latency is huge with SATCOM access

Heavy page – news media type

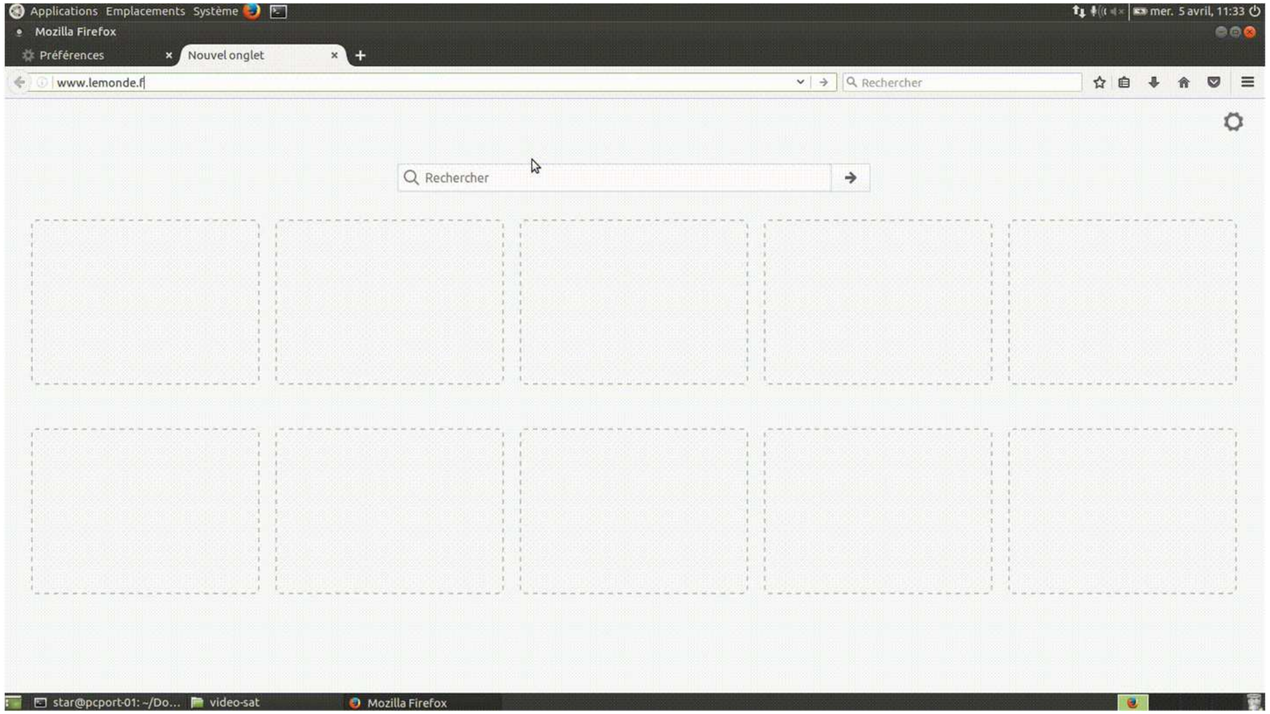

TOOWAY satellite Internet access :

- Solution furnished by ISP ALSATIS with EUTELSAT operator
- 20Mbps download / 6 Mbps upload

Not a Myth #3: SATCOM systems require ‘middleboxes’

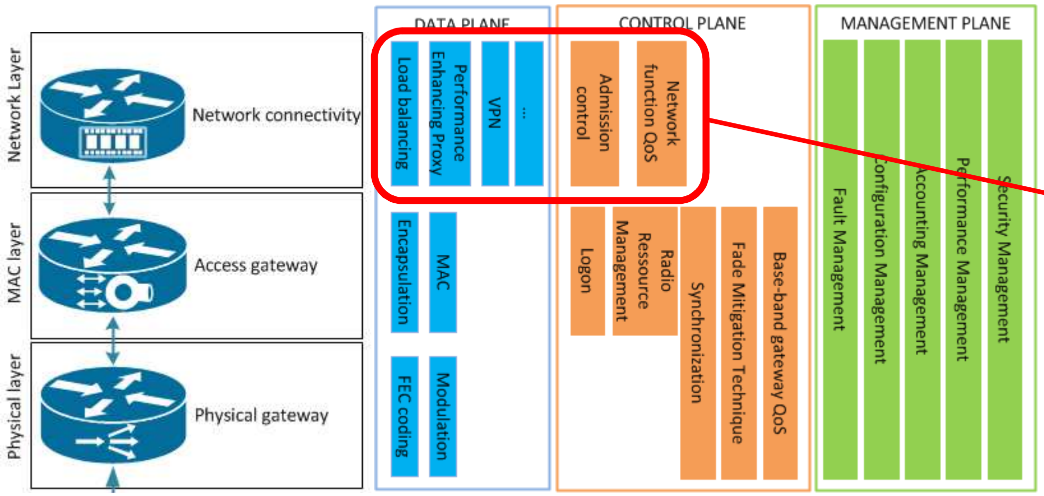

Performance Enhancing Proxy (PEP) – RFC 3135
“magic” mix of transport technologies

- Split TCP connections
- Transparent compression

No support of the most recent improvements at the servers or clients

2

QUIC FOR SATELLITE COMMUNICATIONS

PEP-TCP : A LIMITED SOLUTION WITH QUIC TRAFFIC

Transport layer issues in GEO systems

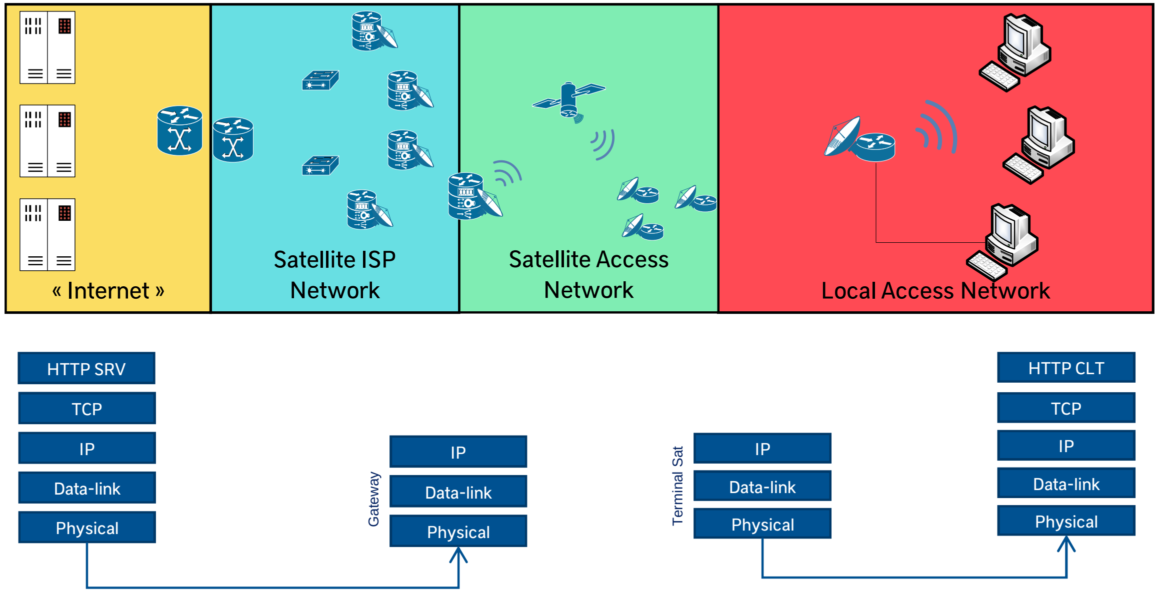

- Connection initialization:
  - Setting up the connection requires three round trips, impacting the moment from which the actual data can be transmitted
- Required window size:
  - To fully exploit the available capacity, it is necessary to increase the sending buffers are the client and the server
- Reliability:
  - Packet loss detection and correction is slow (end-to-end retransmission performance is also affected on GEO access)
- Convergence of congestion control:
  - The exponential increase in data rate is considerably slowed down for a GEO satellite.

Transport layer issues in GEO systems

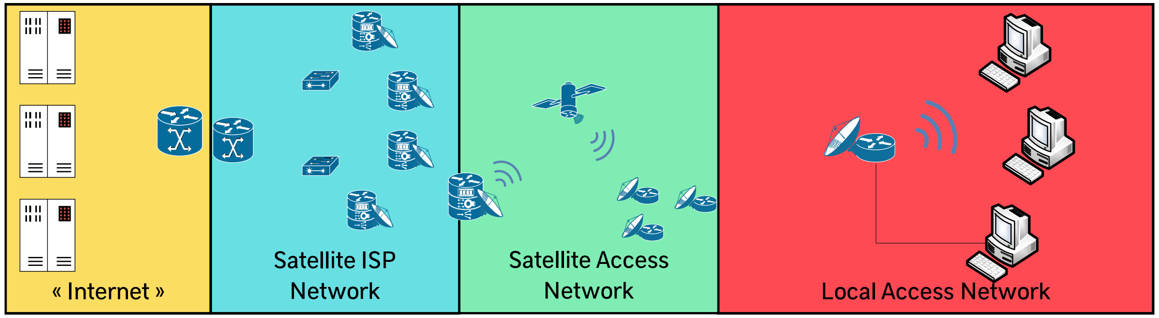

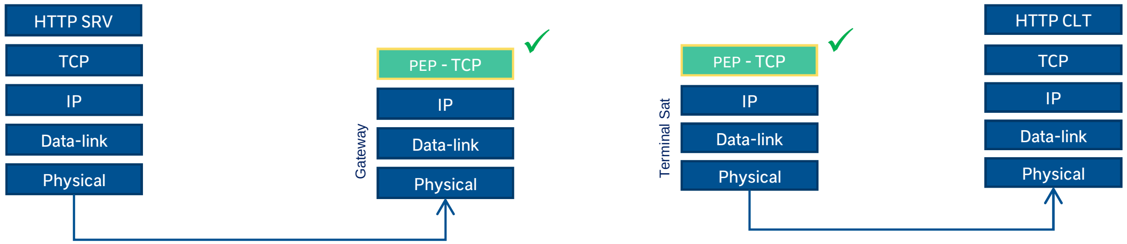

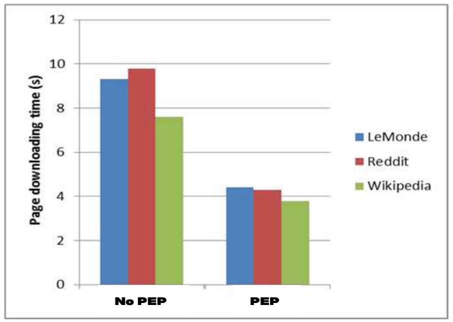

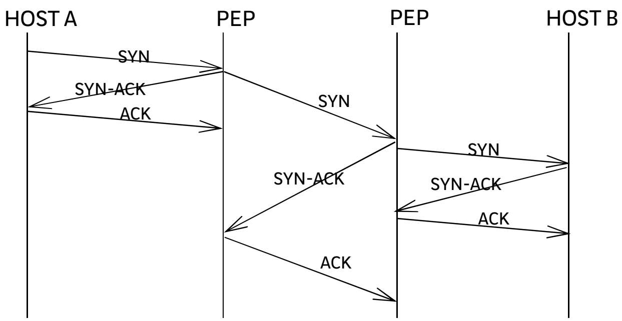

- Connection initialization:
  - Setting up the connection requires three round trips, impacting the moment from which the actual data can be transmitted
  - ➢ [PEP-TCP] Can enable TCP Fast-Open
- Required window size:
  - To fully exploit the available capacity, it is necessary to increase the sending buffers are the client and the server
  - ➢ [PEP-TCP] Improved by custom TCP buffers in TCP – PEP
- Reliability:
  - Packet loss detection and correction is slow (end-to-end retransmission performance is also affected on GEO access)
  - ➢ [PEP-TCP] Loss recovery in splitted in three segments
- Convergence of congestion control:
  - The exponential increase in data rate is considerably slowed down for a GEO satellite.
  - ➢ [PEP-TCP] Improved by custom TCP AIMD in TCP – PEP
  - ➢ [PEP-TCP] Improved by custom TCP initial windows in TCP – PEP

PEP-TCP : A limited solution with QUIC traffic

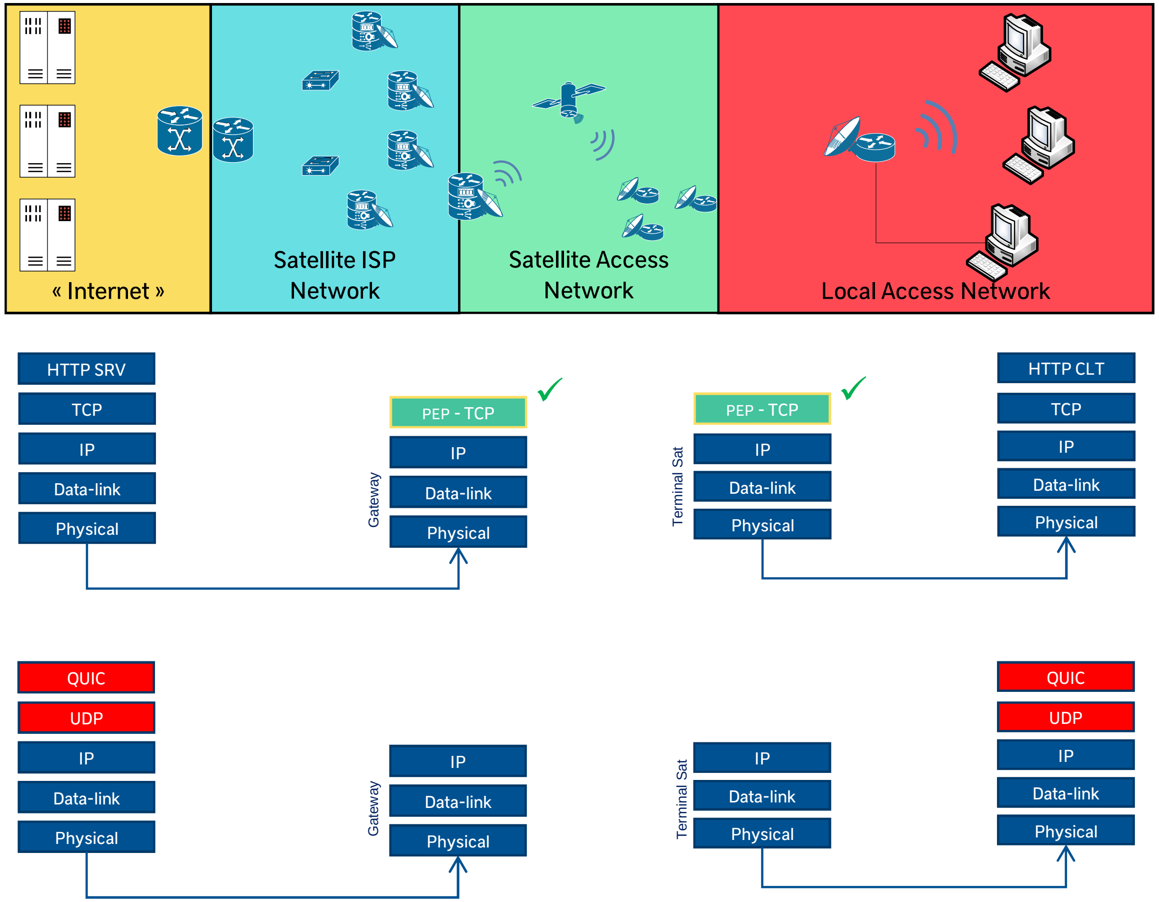

- Connection initialization:
  - Setting up the connection requires three round trips, impacting the moment from which the actual data can be transmitted
  - ➢ [PEP-TCP] Can enable TCP Fast-Open
  - ➢ [QUIC] Saving one (or two) round trip
- Required window size:
  - To fully exploit the available capacity, it is necessary to increase the sending buffers are the client and the server
  - ➢ [PEP-TCP] Improved by custom TCP buffers in TCP – PEP
  - ➢ [QUIC] Limited by end points
- Reliability:
  - Packet loss detection and correction is slow (end-to-end retransmission performance is also affected on GEO access)
  - ➢ [PEP-TCP] Loss recovery in splitted in three segments
  - ➢ [QUIC] Loss recovery is end-to-end
- Convergence of congestion control:
  - The exponential increase in data rate is considerably slowed down for a GEO satellite.
  - ➢ [PEP-TCP] Improved by custom TCP AIMD in TCP – PEP
  - ➢ [PEP-TCP] Improved by custom TCP initial windows in TCP – PEP
  - ➢ [QUIC] End points congestion control may not be adapted

Overview of a QUIC packet

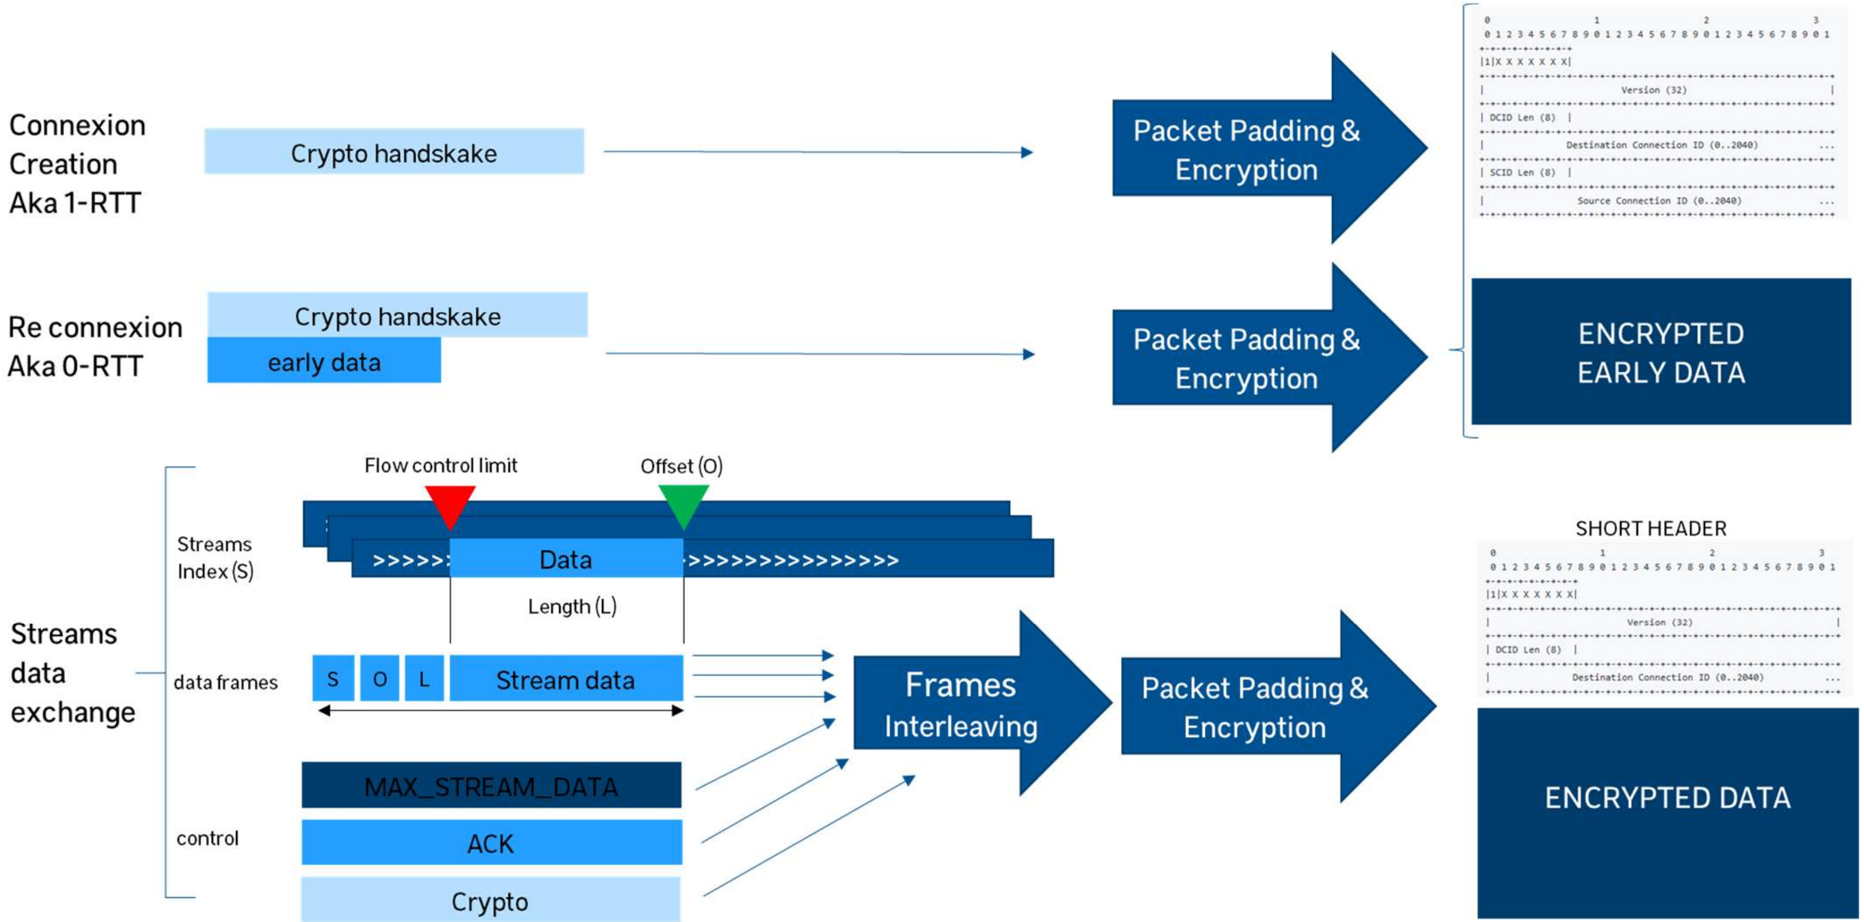

Nicolas Kuhn, Emile Stephan
Internet End-to-end evolution
NetSatDay meets Nof

QUIC vs PEP-TCP

	Who is the winner ?	Why ?	What can we do to help QUIC ?	What are our chances ?
Connection initialization				
Required window size				
Reliability				
Convergence of congestion control				

Disclaimer : Assumes end-to-end QUIC not specifically deployed in a SATCOM access network

3

QUIC FOR SATELLITE COMMUNICATIONS

CONNECTION INITIALIZATION

Connection initialization

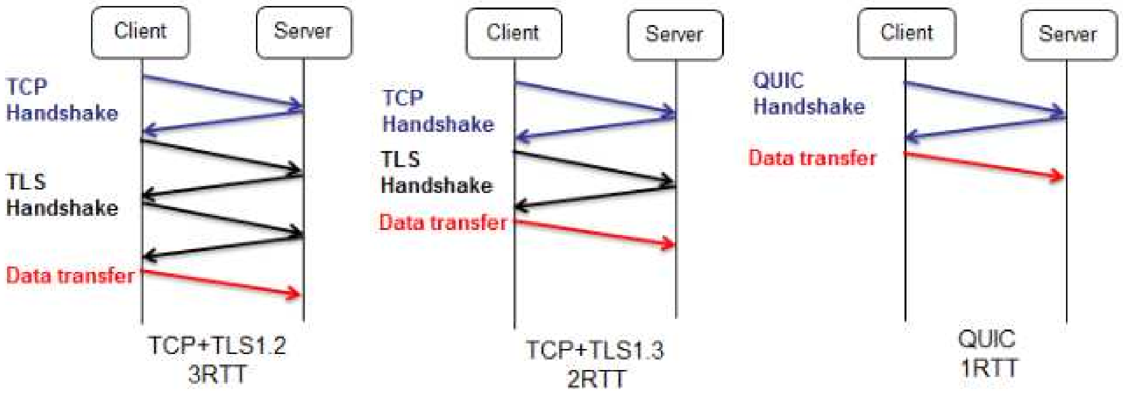

Y. Cui, T. Li, C. Liu, X. Wang and M. Kühlewind, "Innovating Transport with QUIC: Design Approaches and Research Challenges," in IEEE Internet Computing, vol. 21, no. 2, pp. 72-76, Mar.-Apr. 2017.
doi: 10.1109/MIC.2017.44

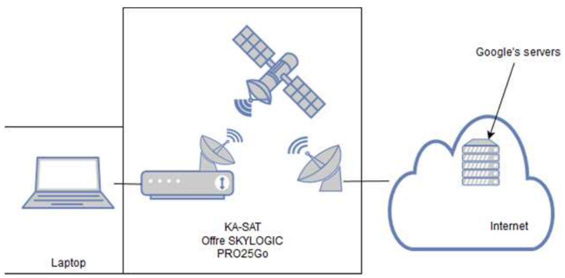

Target (3 objects, 11 kB)

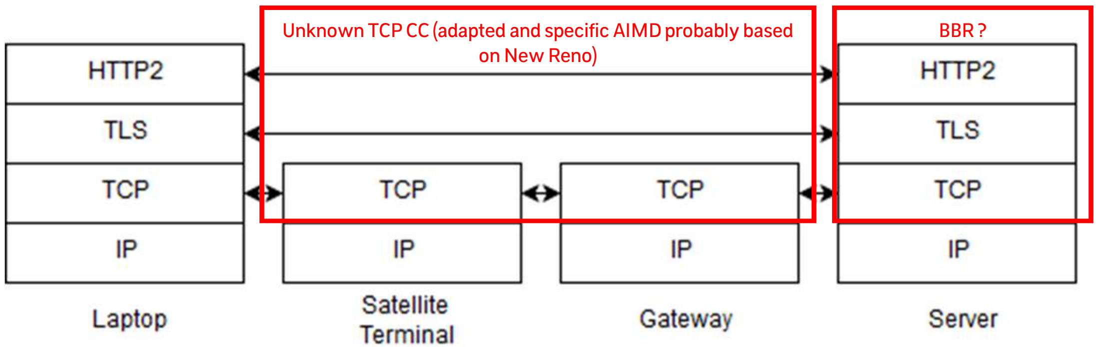

Google QUIC performance over a public SATCOM access
International Journal of Satellite Communications and Networking
THOMAS, L. ; DUBOIS, E. ; KUHN, N. ; LOCHIN, E. 2019

Connection initialization

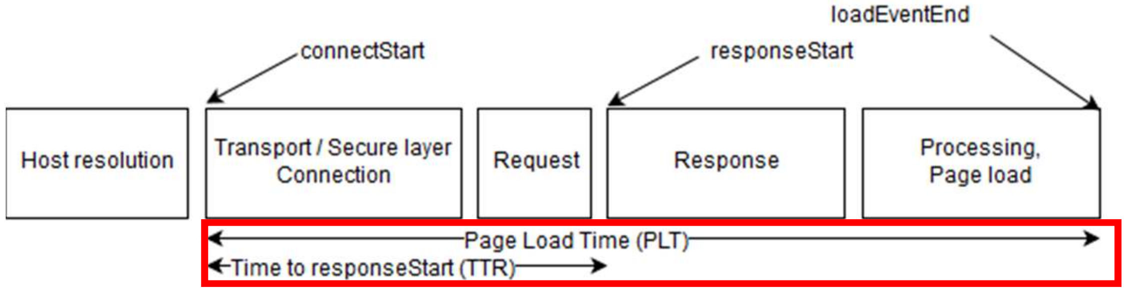

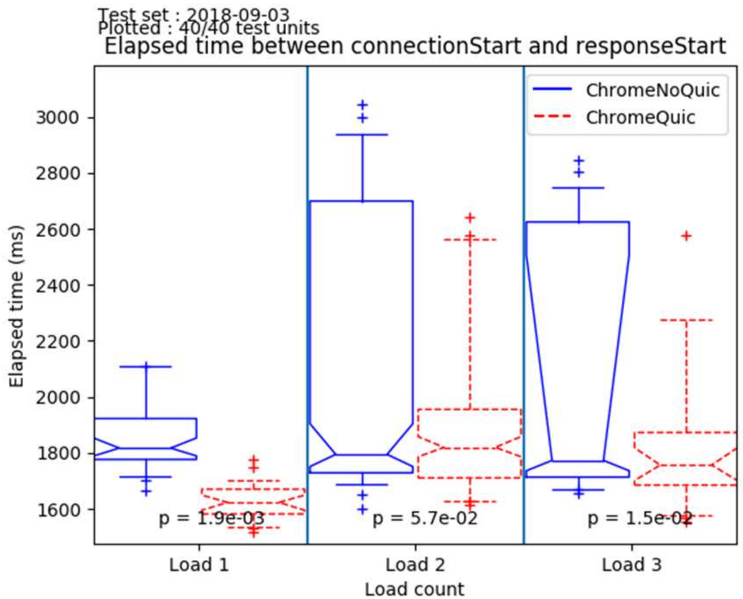

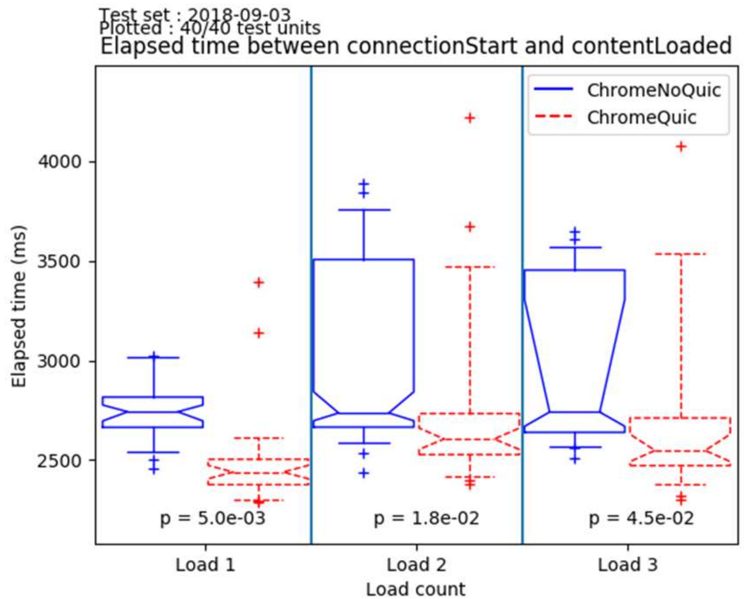

Google QUIC performance over a public SATCOM access
International Journal of Satellite Communications and Networking
THOMAS, L. ; DUBOIS, E. ; KUHN, N. ; LOCHIN, E. 2019

QUIC vs PEP-TCP

	Who is the winner ?	Why ?	What can we do to help QUIC ?	What are our chances ?
Connection initialization	QUIC	Reduced handshake Either deployment of 0RTT	No need !	
Required window size				
Reliability				
Convergence of congestion control				

Disclaimer : Assumes end-to-end QUIC not specifically deployed in a SATCOM access network

4

QUIC FOR SATELLITE COMMUNICATIONS

REQUIRED WINDOW SIZE

Required window size

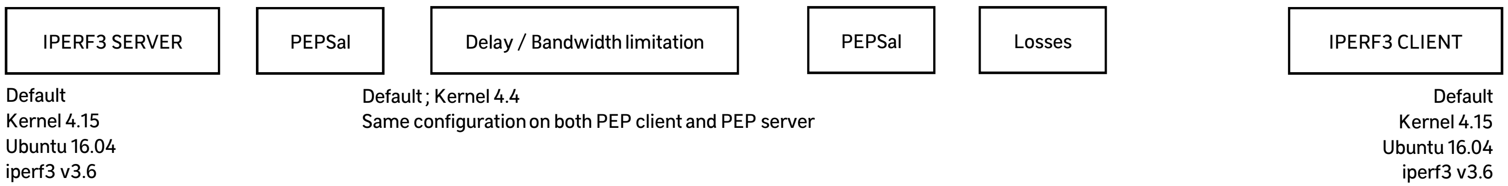

	TCP_WMEM_MAX (MB)	TCP_RMEM_MAX (MB)	CORE_WMEM_MAX CORE_RMEM_MAX (MB)	ICWND IRWND (packets)
NO PEP	4	6	0,2	10
PEP A	4	6	0,2	10
PEP B	4	6	0,2	100
PEP C	33	33	33	10
PEP D	33	33	33	100

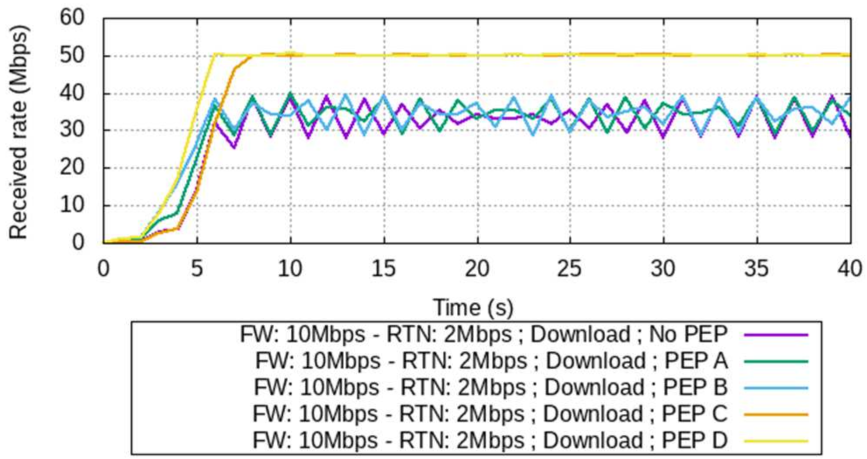

Required window size

- Client and servers use GO-QUIC
  - https://github.com/lucas-clemente/quic-go
  - Modified to support the download of objects in sequence or in parallel
- Object size
  - Short (40 KB), medium (290 KB), long (4 MB), large (66 MB)

Layer	Parameter	Default	H-BDP	Unit
Connection	DefaultMaxCongestionWindowPackets	1000	2500	packets
	InitialCongestionWindow	32	32	tcpMSS
	MaxTrackedSkippedPackets	10	50	packets
	MaxTrackedSentPackets	2 500	10 000	packets
	MinPacingDelay	100	10	microsec
Stream	InitialMaxStreamData	512	6000	KB
	DefaultMaxReceiveStream FlowControlWindow	6	12	MB

Required window size

RTT : 600ms

Object	Mode	Default	H-BDP	(H-BDP - Default)/Default
Short: 10x40KB	Seq	7,65	7,45	-3%
	Par	3,65	2,54	-30%
Med 10x 292KB	Seq	11,13	10,91	-2%
	Par	5,20	4,08	-21%
Long 10x4149KB	Seq	31,57	20,77	-34%
	Par	21,36	10,78	-50%
Large 2x66.390KB	Seq	61,74	29,35	-52%
	Par	60,17	26,79	-55%
All 31 Obj: 111.211KB	All	51,03	26,16	-49%

RTT : 100ms

Obj	Mode	Default	H-BDP	(H-BDP - Default)/Default
Short: 10x40KB	Seq	1,44	1,39	-4%
	Par	,64	,54	-16%
Med 10x 292KB	Seq	2,58	4,00	55%
	Par	1,10	,95	-13%
Long 10x4149KB	Seq	9,12	8,46	-7%
	Par	7,65	7,49	-2%
Large 2x66.390KB	Seq	23,42	23,43	0%
	Par	23,19	23,23	0%
All 31 Obj: 111.211KB	All	19,55	19,52	0%

QUIC vs PEP-TCP

	Who is the winner ?	Why ?	What can we do to help QUIC ?	What are our chances ?
Connection initialization	QUIC	Reduced handshake Either deployment of 0RTT	No need !	
Required window size	PEP-TCP	Specifically sized windows	Convince big players to increase window size	
Reliability				
Convergence of congestion control				

Disclaimer : Assumes end-to-end QUIC not specifically deployed in a SATCOM access network

5

QUIC FOR SATELLITE COMMUNICATIONS

RELIABILITY

Loss in SATCOM SYSTEMS

End to end measurements on a real satellite public access

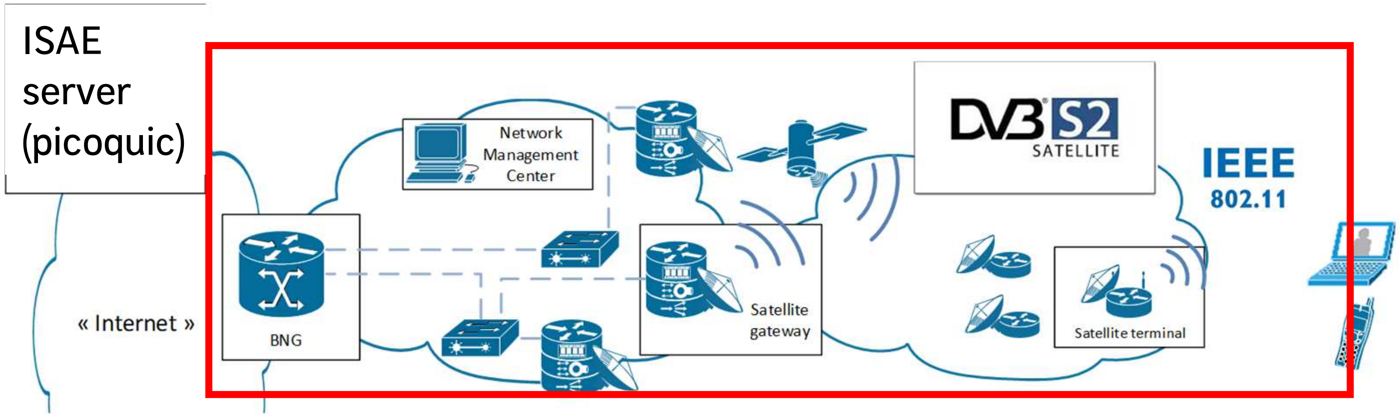

Loss identified by missing QUIC packets are the receiver

- Gilbert-Elliot model
- Probability to go from « good » to « bad » state = 0.018 !

Losses in SATCOM systems : identification and impact
MAPRG – IETF 106
Nicolas KUHN, Emmanuel DUBOIS, Alexandre FERRIEUX, François MICHEL, Emmanuel LOCHIN

Impact of transmission losses on a TCP connection

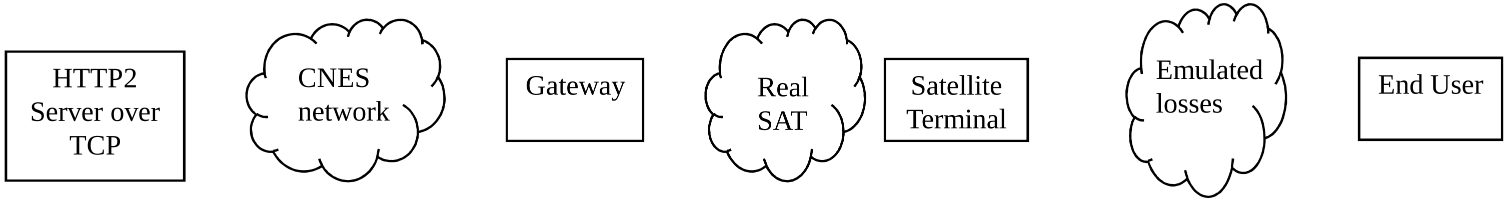

Loss ratio	Time needed to download 1 GB (s)	Goodput (Mbps)	Loss impact (1- Goodput-loss/Goodput-noloss)
0	797	10	0
0.0001	935	8.5	0.15
0.0005	1528	5.2	0.48
0.001	1863	4.2	0.58
0.005	7140	1.1	0.89

- Experimental evaluations of QUIC showed good performance for short flows with public accesses
- For long flows, the E2E losses can have a huge impact

Losses in SATCOM systems : identification and impact
MAPRG – IETF 106
Nicolas KUHN, Emmanuel DUBOIS, Alexandre FERRIEUX, François MICHEL, Emmanuel LOCHIN

Where are the losses ?

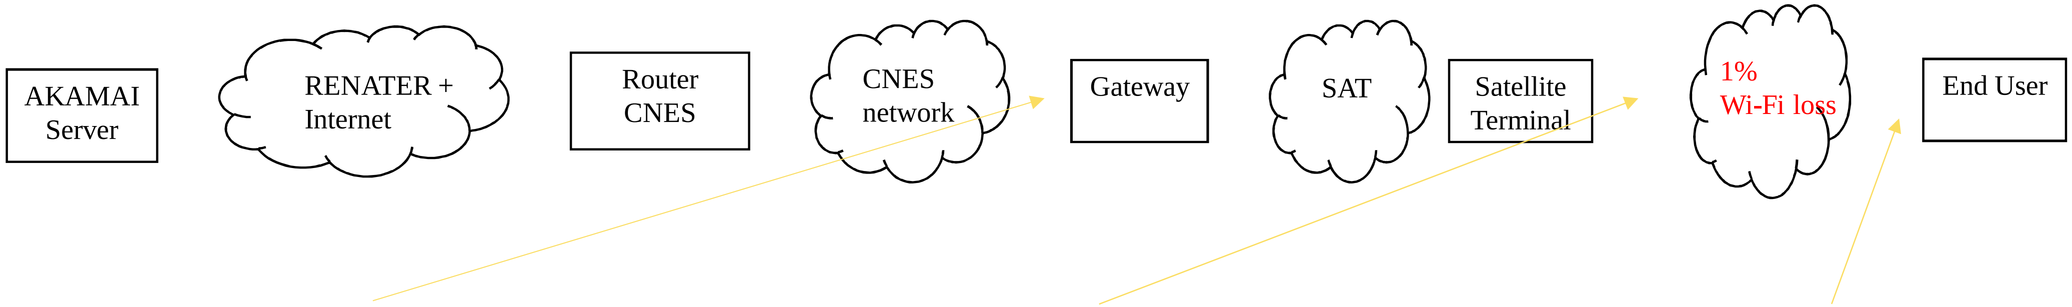

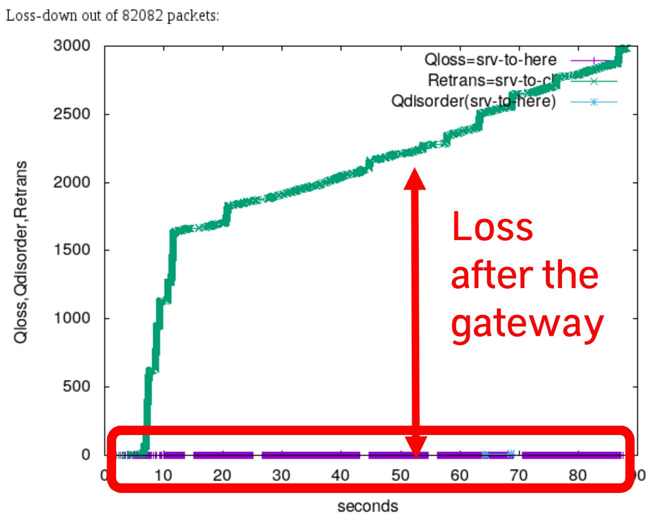

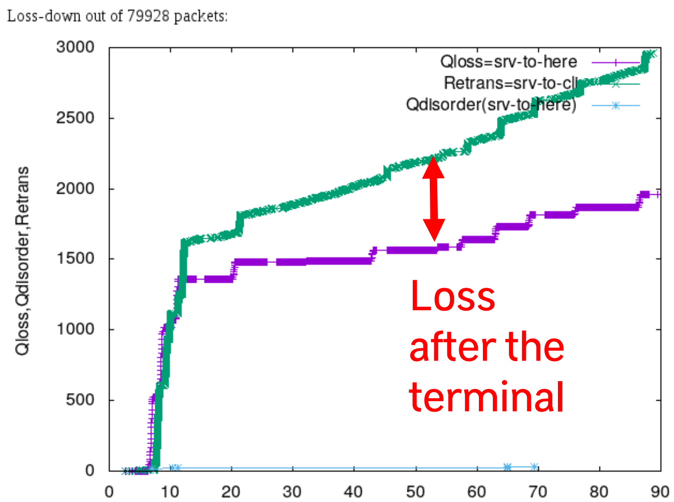

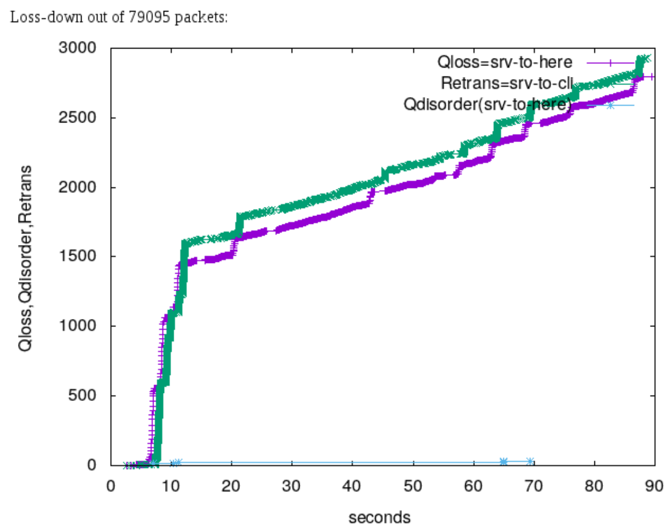

Losses in SATCOM systems : identification and impact
MAPRG – IETF 106
Nicolas KUHN, Emmanuel DUBOIS, Alexandre FERRIEUX, François MICHEL, Emmanuel LOCHIN

What can be done about this ?

- Add coding / retransmission with an encapsulation tunnel

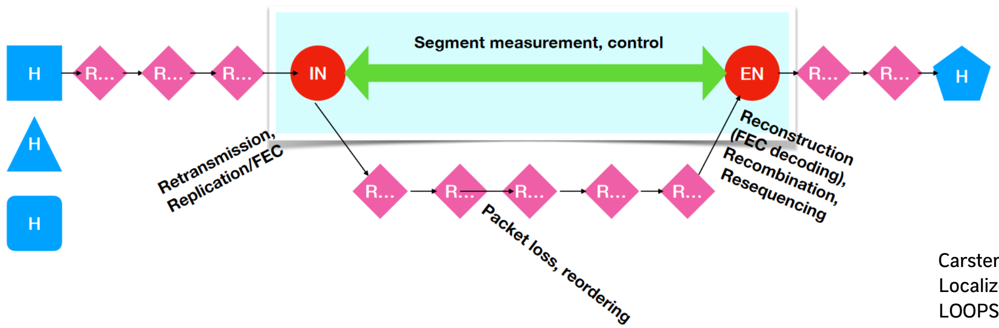

Carsten Bormann
Localized Optimizations over Path Segments
LOOPS BOF @ IETF 108 July 31, 2020

Side note :

'soon to be' RFC on the interaction between congestion control and coding :

Coding and congestion control in transport – Nicolas Kuhn , Emmanuel Lochin , François Michel , Michael Welzl
https://datatracker.ietf.org/doc/draft-irtf-nwcrg-coding-and-congestion/

- Deployment of congestion control 'ignoring' losses

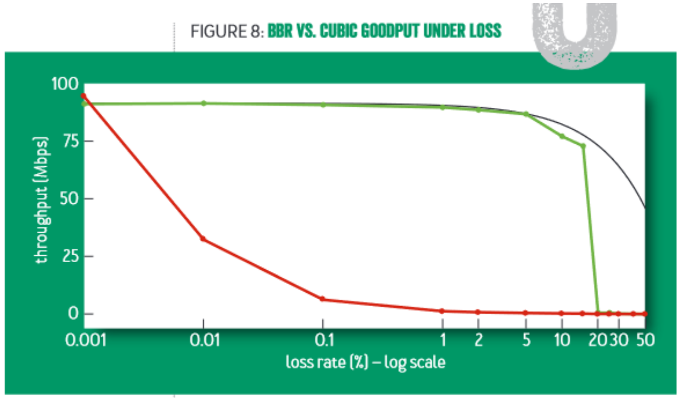

Neal Cardwell, Yuchung Cheng, C. Stephen Gunn, Soheil Hassas Yeganeh, and Van Jacobson. 2016.

BBR: Congestion-Based Congestion Control: Measuring bottleneck bandwidth and round-trip propagation time. Queue 14, 5 (September-October 2016), 20–53. DOI:https://doi.org/10.1145/3012426.3022184

QUIC vs PEP-TCP

	Who is the winner ?	Why ?	What can we do to help QUIC ?	What are our chances ?
Connection initialization	QUIC	Reduced handshake Either deployment of 0RTT	No need !	
Required window size	PEP-TCP	Specifically sized windows	Convince big players to increase window size	
Reliability	PEP-TCP	Loss recovery includes a large RTT	Add coding with an encapsulation tunnel Deployment of congestion control ignoring losses	
Convergence of congestion control				

Disclaimer : Assumes end-to-end QUIC not specifically deployed in a SATCOM access network

6

QUIC FOR SATELLITE COMMUNICATIONS

CONVERGENCE OF CONGESTION CONTROL

On the need for a quick convergence to data-rate

Tests based on open platform

https://forge.net4sat.org/kuhnn/openbach-example-simple

- QUIC :
    - PICOQUIC implementation
    - BBR (PICOQUIC implementation of BBR)
- Variable RTT (100 ms -> 500 ms)
- File size : 500 KB, 1MB, 10 MB, 100MB
- Bottleneck (Forward/Return)
    - 1 Mbps / 100 kbps
    - 10 Mbps / 2 Mbps
    - 50 Mbps / 25 Mbps
    - 200 Mbps / 100 Mbps

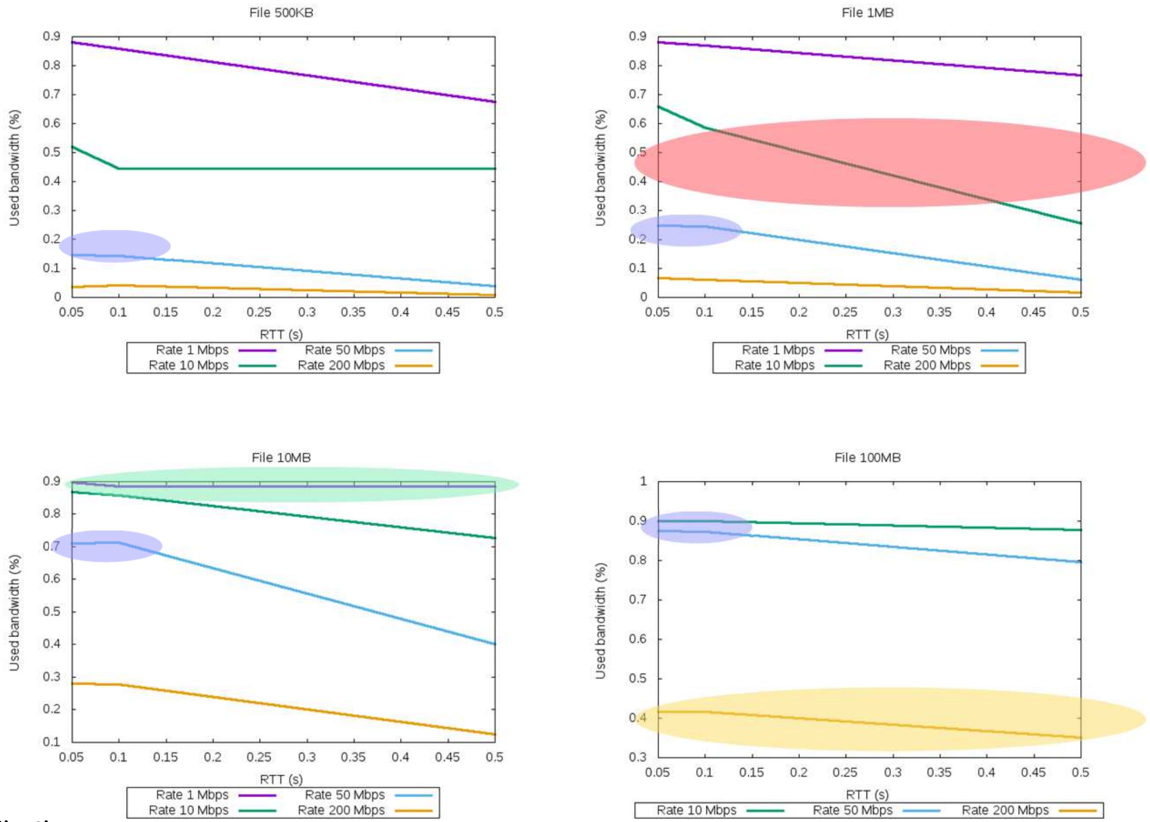

- With a 10 MB file and 1 Mbps, the link is used for all RTT
- For shorter files (1MB), increasing the RTT severely impacts link utilization
- When the data rate is high (250 Mbps), even a 100 MB transfer does not utilize the link
- Increasing the file size increases the link utilization

QUICK connection establishment with QUIC

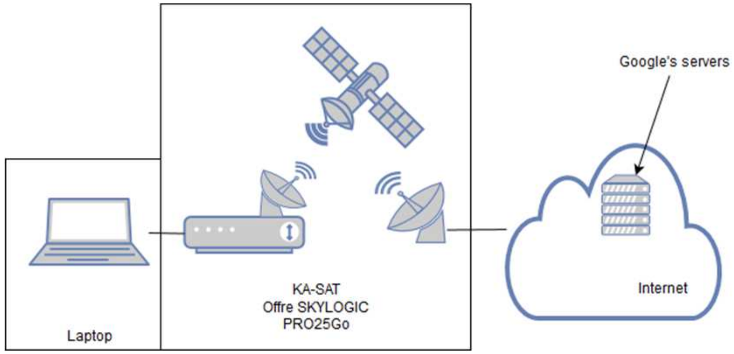

Target (1 object, 5.3MB)

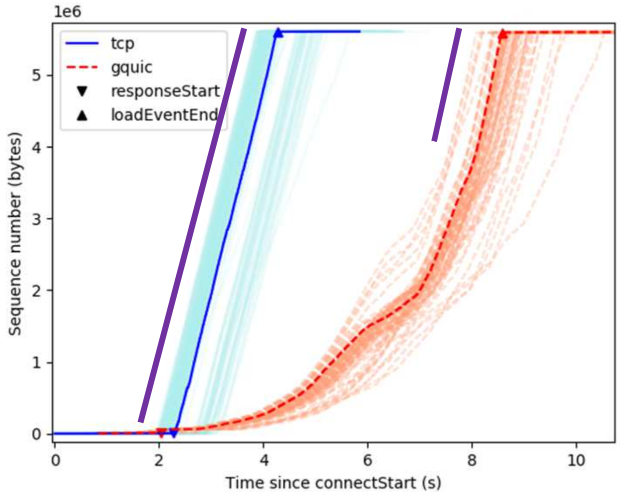

Google QUIC performance over a public SATCOM access
International Journal of Satellite Communications and Networking
THOMAS, L. ; DUBOIS, E. ; KUHN, N. ; LOCHIN, E. 2019

0-RTT-BDP extension : Description

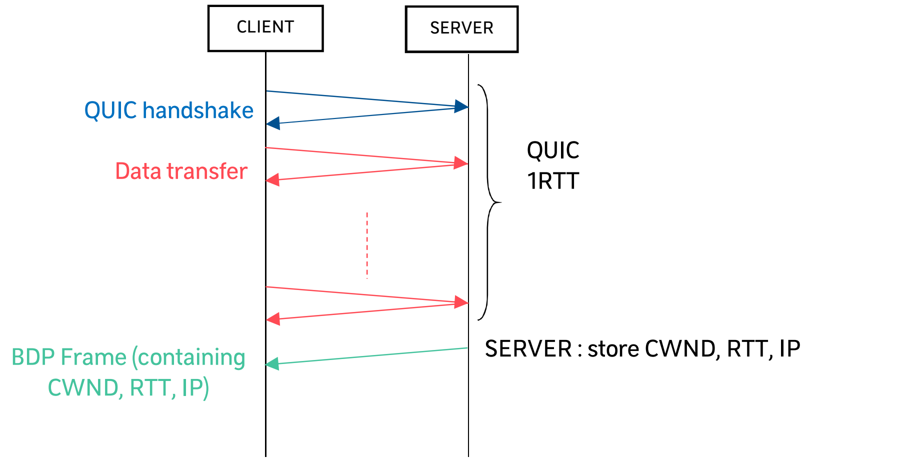

1. During a previous session, current RTT (current_rtt), CWND (current_cwnd) and client's current IP (current_client_ip) are stored as saved_rtt, saved_cwnd and saved_client_ip;

2. When resuming a session, the server might set the current_rtt and the current_cwnd to the saved_rtt and saved_cwnd of a previous connection.

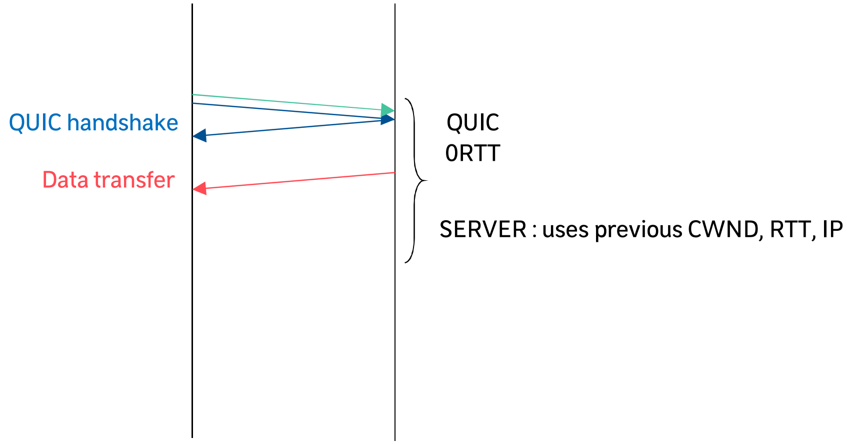

https://datatracker.ietf.org/doc/html/draft-kuhn-quic-0rtt-bdp

Implemented in PICOQUIC by F. Simo, D. Pradas
https://github.com/private-octopus/picoquic/pull/1209

0-RTT-BDP extension : QLOG example of QUIC 0RTTBDP connection

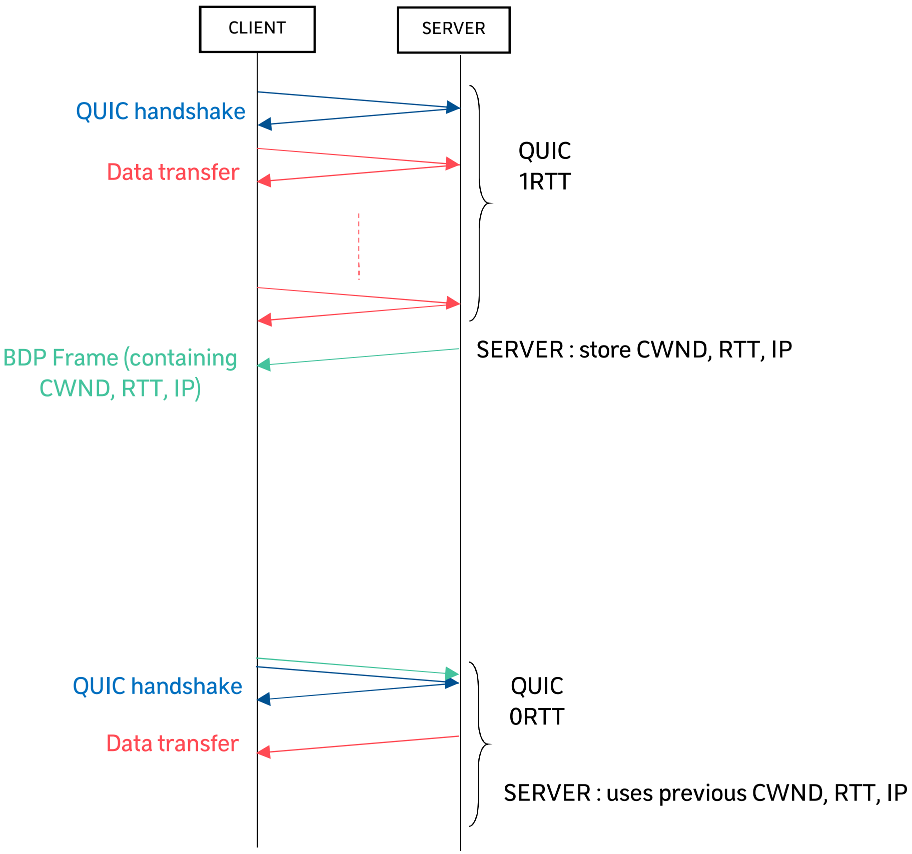

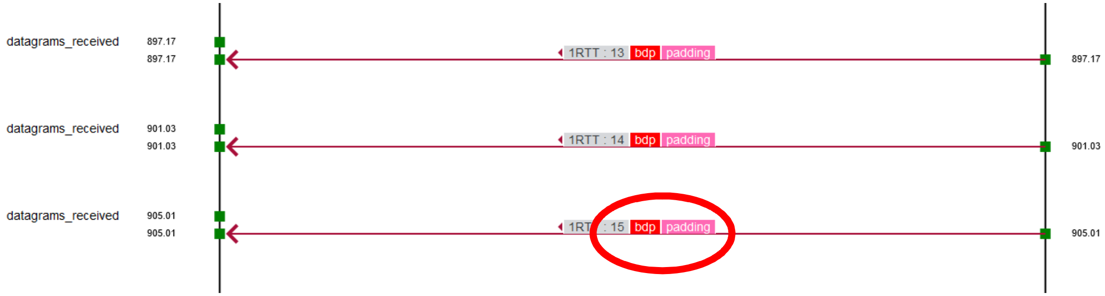

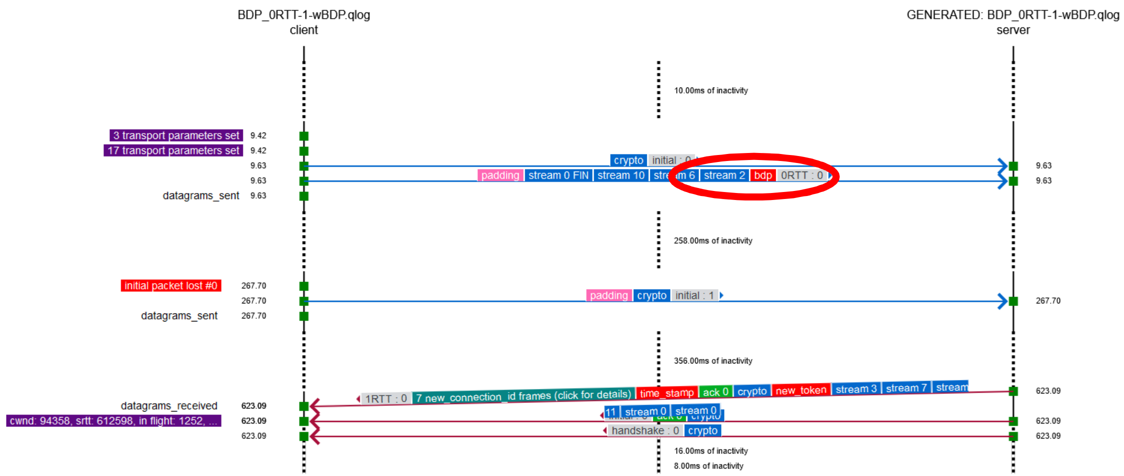

https://datatracker.ietf.org/doc/html/draft-kuhn-quic-0rtt-bdp

0-RTT-BDP extension : Emulated performance

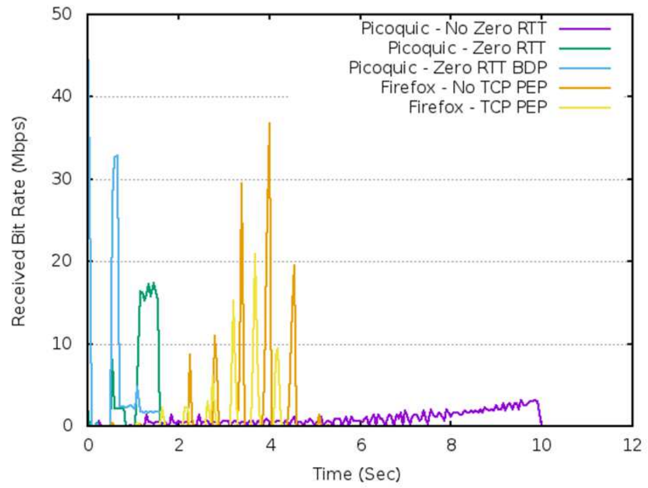

- Evaluations based on
  - draft-kuhn-quic-4-sat-06 scenarios
  - Implementation of draft-kuhn-quic-0rtt-bdp-07
  - Picoquic https://github.com/private-octopus/picoquic/pull/1073
- Network characteristics:
  - 50 Mbps download / 10 Mbps upload
  - RTT : 650 ms
- Congestion control
  - CUBIC
  - 0-RTT-BDP reaction:
    - jump to preciously measured capacity (not recommended but “easy to implement” as a first step)
    - Beware the potential issue in using bytes_in_flight metric
- Application level
  - 2 MB transfer - median

Without 0-RTT	With 0-RTT	With 0-RTT-BDP
4,3 s	3,4 s	2,9 s

F. Simo, D. Pradas

0-RTT-BDP extension : Real satellite access Performance

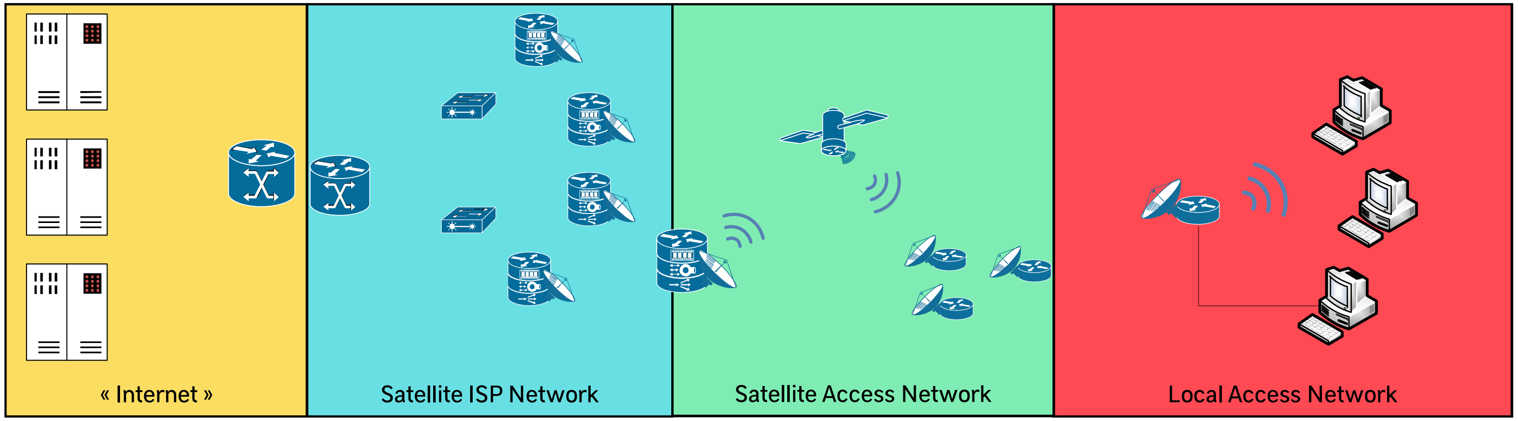

PICOQUIC CLIENT
- Default and 0-RTT : commit : 6b917d69bf4ac69d4ab43c425 54bb702ed844561
- 0RTTBDP : https://github.com/private-octopus/picoquic/pull/1209

Satellite : GEO - KaSAT
Offer : PRO25

PICOQUIC SERVER
- Default and 0-RTT : commit : 6b917d69bf4ac69d4ab43c425 54bb702ed844561
- 0RTTBDP : https://github.com/private-octopus/picoquic/pull/1209

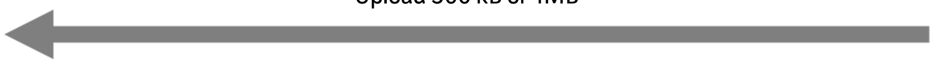

0-RTT-BDP extension : Real satellite access Performance

Without 0-RTT

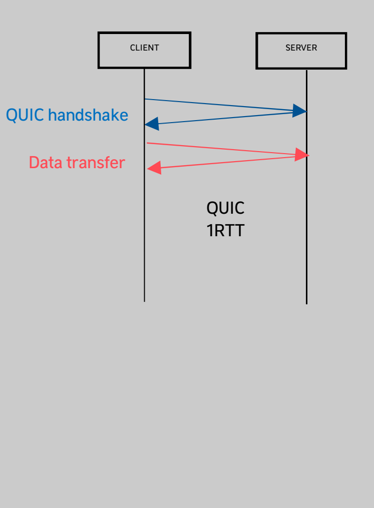

With 0-RTT

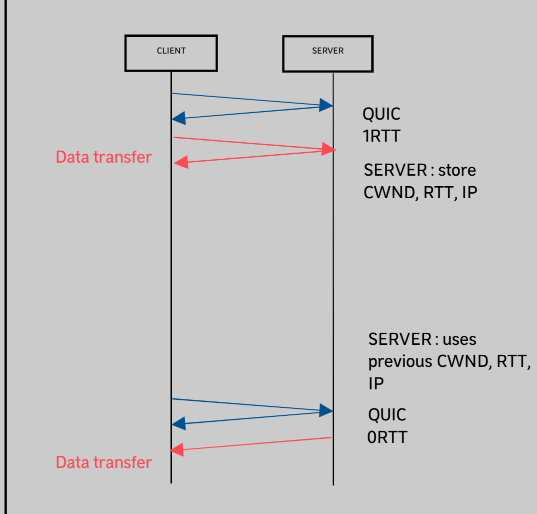

With 0-RTT-BDP

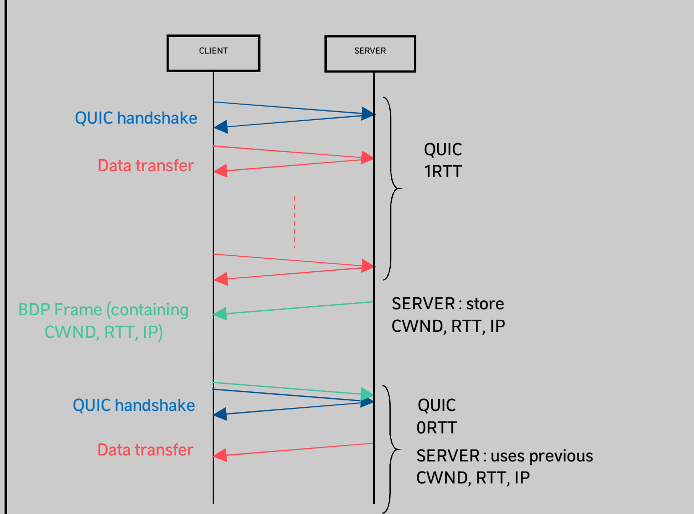

	Without 0-RTT		With 0-RTT		With 0-RTT-BDP	
	500 kB	1 MB	500 kB	1 MB	500 kB	1 MB
Min transfer time	3,12 s	3,87 s	2,43 s	3,19 s	1,87 s	2,78 s
Average transfer time	11,34 s	17,15 s	7,59 s	10,24 s	4,24 s	6,47 s
Max transfer time	47,82 s	61,43 s	33,69 s	33,92 s	15,55 s	23,88 s

Gain for a full 500kB or 1MB transfer (including HTTP GET)
Gain due to the arrival of HTTP request one RTT in advance (and large satellite RTT !)

IETF Status

draft-kuhn-quic-0rtt-bdp includes 3 methods

2 methods are implemented in picoquic

BDP frame - https://github.com/private-octopus/picoquic/pull/1209

local storage of CWND, RTT parameters - https://github.com/private-octopus/picoquic/pull/1204

Next

Looking for other implementers

Integration in QUIC interop matrix

QUIC vs PEP-TCP

	Who is the winner ?	Why ?	What can we do to help QUIC ?	What are our chances ?
Connection initialization	QUIC	Reduced handshake Either deployment of 0RTT	No need !	
Required window size	PEP-TCP	Specifically sized windows	Convince big players to increase window size	
Reliability	PEP-TCP	Loss recovery includes a large RTT	Add coding with an encapsulation tunnel Deployment of congestion control ignoring losses	
Convergence of congestion control	PEP-TCP	Slow ramp up to available data-rate	Deployment of congestion control with aggressiv ramp up Deployment of 0RTTBDP !	

Disclaimer : Assumes end-to-end QUIC not specifically deployed in a SATCOM access network

7

QUIC FOR SATELLITE COMMUNICATIONS

CONCLUSION

Conclusion

	Who is the winner ?	Why ?	What can we do to help QUIC ?	What are our chances ?
Connection initialization	QUIC	Reduced handshake Either deployment of 0RTT	No need !	
Required window size	PEP-TCP	Specifically sized windows	Convince big players to increase window size	
Reliability	PEP-TCP	Loss recovery includes a large RTT	Add coding with an encapsulation tunnel Deployment of congestion control ignoring losses	
Convergence of congestion control	PEP-TCP	Slow ramp up to available data-rate	Deployment of congestion control with aggressiv ramp up Deployment of 0RTTBDP !	

QUIC : dangerous opportunity for SATCOM systems

- Good performance for short files
- Cheaper ground segments (no more middleboxes)
- Sensitivity to losses
- Complex ‘in the middle’ optimization

To guarantee end user performances

- Risks of blocking QUIC traffic
- TCP is not dead !

Conclusion

Big thanks to AFNIC

Questions ?

Going further

On going IETF documents

- draft-kuhn-quic-0rtt-bdp : Transport parameters for 0-RTT connections. N. Kuhn, E. Stephan, G. Fairhurst, T. Jones and C. Huitema.
- draft-jones-tsvwg-transport-for-satellite : Enhancing Transport Protocols over Satellite Networks. T. Jones, G. Fairhurst, N. Kuhn, J. Border and E. Stephan.

IETF mailing list

- Encrypted Transport over Satellite (EToSat)
- https://www.ietf.org/mailman/listinfo/Etosat

Some pointers

- Thomas, L, Dubois, E, Kuhn, N, Lochin, E. Google QUIC performance over a public SATCOM access. Int J Satell Commun Network. 2019; 37: 601– 611. https://doi.org/10.1002/sat.1301
- Ahmed, T., Dubois, E., Dupé, J.-B., Ferrús, R., Gélard, P., and Kuhn, N. (2018) Software-defined satellite cloud RAN. Int. J. Satell. Commun. Network., 36: 108– 133. doi: 10.1002/sat.1206.
- N. Kuhn, F. Michel, L. Thomas, E. Dubois and E. Lochin, "QUIC: Opportunities and threats in SATCOM," 2020 10th Advanced Satellite Multimedia Systems Conference and the 16th Signal Processing for Space Communications Workshop (ASMS/SPSC), 2020, pp. 1-7, doi: 10.1109/ASMS/SPSC48805.2020.9268814.
- J. Deutschmann, K. Hielscher and R. German, "Satellite Internet Performance Measurements," 2019 International Conference on Networked Systems (NetSys), 2019, pp. 1-4, doi: 10.1109/NetSys.2019.8854494.
- J. Border, B. Shah, C. Su and R. Torres, "Evaluating QUIC's Performance Against Performance Enhancing Proxy over Satellite Link," 2020 IFIP Networking Conference (Networking), 2020, pp. 755-760.
- A. Custura, T. Jones and G. Fairhurst, "Impact of Acknowledgements using IETF QUIC on Satellite Performance," 2020 10th Advanced Satellite Multimedia Systems Conference and the 16th Signal Processing for Space Communications Workshop (ASMS/SPSC), 2020, pp. 1-8, doi: 10.1109/ASMS/SPSC48805.2020.9268894.

Contact :

- nicolas.kuhn.ietf@gmail.com