Deep Space IP
Some simulation results

IETF 120, Vancouver, Canada

Marc Blanchet, marc.blanchet@viagenie.ca, July 2024

Last IETF meeting

- Presented results:
  - HTTP request to Voyager (aka 36 hours RTT)
  - QUIC handling reordering
  - SNMP with 4h RTT working
  - NTP with 8h RTT working
- Done with a « heavy » testbed using wall clock time
- Since then…

10 days RTT!!!

- 5 days delay one way (432000s), 10 days RTT (864000s)
- If connection is already established and is not closed, then only one RTT for request and response

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000	10.132.239.239	10.132.239.121	QUIC	1242	Initial, DCID=31e20cceeacc6ddb3903704db5c9be687df54441, SCID=05e7e0f8a63866b8, PK…
2	432000.0…	10.132.239.121	10.132.239.239	QUIC	1380	Handshake, DCID=05e7e0f8a63866b8, SCID=c2d1484de334b668, PKN: 0, CRYPTO
3	864000.0…	10.132.239.239	10.132.239.121	QUIC	1242	Handshake, DCID=c2d1484de334b668, SCID=05e7e0f8a63866b8, PKN: 0, ACK_ECN, CRYPTO
4	864000.0…	10.132.239.239	10.132.239.121	QUIC	200	Protected Payload (KP0), DCID=c2d1484de334b668, PKN: 1, NCI, NCI, NCI, NCI, STREA…
5	1296000.…	10.132.239.121	10.132.239.239	QUIC	689	Protected Payload (KP0), DCID=9e463c45aa076440
6	1728000.…	10.132.239.239	10.132.239.121	QUIC	77	Protected Payload (KP0), DCID=c2d1484de334b668, PKN: 2, ACK_ECN
7	1728000.…	10.132.239.239	10.132.239.121	QUIC	84	Protected Payload (KP0), DCID=c2d1484de334b668, PKN: 3, ACK_ECN, CC
8	2160000.…	10.132.239.121	10.132.239.239	QUIC	81	Protected Payload (KP0), DCID=9e463c45aa076440, PKN: 4, ACK_ECN, CC

HTTP Download

- 4M file downloaded, 600s delay
  - including connection open and close, which are not needed if the QUIC connection was previous established
- At 1200s (packet #3): HTTP GET
- At 1800s (packets #4-590: HTTP Response (all packets for the 4M file)
- At 2400s: acknowledgements of reception. (We shall be able to decrease the number of ACKs packets. TBD)
- File/Media transfer over HTTP: ✅

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000	10.188.64.179	10.188.64.80	QUIC	1242	Initial, DCID=09fdc0dd933e591c218bf2167a6a036fce8d8108, SCID=e179fd94eca15817, PKN: 0, CRYPTO, PADDING
2	600.002288	10.188.64.80	10.188.64.179	QUIC	1380	Handshake, DCID=e179fd94eca15817, SCID=e6b7c120c85bfed2, PKN: 0, CRYPTO
3	1200.005105	10.188.64.179	10.188.64.80	QUIC	200	Protected Payload (KP0), DCID=e6b7c120c85bfed2, PKN: 1, NCI, NCI, NCI, NCI, STREAM(0)
4	1800.027403	10.188.64.80	10.188.64.179	QUIC	420	Protected Payload (KP0), DCID=15a5305b42284179, PKN: 3, DONE, ACK_ECN, AF, CRYPTO, RC
5	1800.027404	10.188.64.80	10.188.64.179	QUIC	156…	Protected Payload (KP0), DCID=15a5305b42284179
6	1800.027467	10.188.64.80	10.188.64.179	QUIC	4842	Protected Payload (KP0), DCID=15a5305b42284179
7	1800.027636	10.188.64.80	10.188.64.179	QUIC	8442	Protected Payload (KP0), DCID=15a5305b42284179
8	1800.027749	10.188.64.80	10.188.64.179	QUIC	6042	Protected Payload (KP0), DCID=15a5305b42284179
9	1800.027919	10.188.64.80	10.188.64.179	QUIC	9642	Protected Payload (KP0), DCID=15a5305b42284179
10	1800.028107	10.188.64.80	10.188.64.179	QUIC	108…	Protected Payload (KP0), DCID=15a5305b42284179
588	1800.095089	10.188.64.80	10.188.64.179	QUIC	6042	Protected Payload (KP1), DCID=15a5305b42284179
589	1800.095102	10.188.64.80	10.188.64.179	QUIC	3642	Protected Payload (KP1), DCID=15a5305b42284179
590	1800.095217	10.188.64.80	10.188.64.179	QUIC	4842	Protected Payload (KP1), DCID=15a5305b42284179
591	2400.106510	10.188.64.179	10.188.64.80	QUIC	82	Protected Payload (KP0), DCID=e6b7c120c85bfed2, PKN: 2, ACK_ECN
592	2400.107677	10.188.64.179	10.188.64.80	QUIC	85	Protected Payload (KP0), DCID=e6b7c120c85bfed2, PKN: 3, ACK_ECN
593	2400.114945	10.188.64.179	10.188.64.80	QUIC	114	Protected Payload (KP0), DCID=e6b7c120c85bfed2, PKN: 4, ACK_ECN, NCI
594	2400.116011	10.188.64.179	10.188.64.80	QUIC	88	Protected Payload (KP0), DCID=e6b7c120c85bfed2, PKN: 5, ACK_ECN
595	2400.117155	10.188.64.179	10.188.64.80	QUIC	89	Protected Payload (KP0), DCID=e6b7c120c85bfed2, PKN: 6, ACK_ECN
596	2400.118305	10.188.64.179	10.188.64.80	QUIC	89	Protected Payload (KP0), DCID=e6b7c120c85bfed2, PKN: 8, ACK_ECN
597	2400.119293	10.188.64.179	10.188.64.80	QUIC	91	Protected Payload (KP0), DCID=e6b7c120c85bfed2, PKN: 9, ACK_ECN
598	2400.120164	10.188.64.179	10.188.64.80	QUIC	92	Protected Payload (KP0), DCID=e6b7c120c85bfed2, PKN: 10, ACK_ECN
623	2400.146343	10.188.64.179	10.188.64.80	QUIC	97	Protected Payload (KP1), DCID=e6b7c120c85bfed2, PKN: 35, ACK_ECN
624	2400.151962	10.188.64.179	10.188.64.80	QUIC	105	Protected Payload (KP1), DCID=e6b7c120c85bfed2, PKN: 36, ACK_ECN, CC
625	3000.149095	10.188.64.80	10.188.64.179	QUIC	82	Protected Payload (KP1), DCID=15a5305b42284179
626	3000.153007	10.188.64.80	10.188.64.179	QUIC	82	Protected Payload (KP1), DCID=15a5305b42284179
627	3600.149501	10.188.64.179	10.188.64.80	QUIC	107	Protected Payload (KP1), DCID=e6b7c120c85bfed2, PKN: 37, ACK_ECN, CC

Link interruptions

- Three nodes network: Client - link - Forwarder - link - Server
- Ping
- 30 seconds delay on each outgoing interface of the forwarder => 60s RTT
- Link down for 5 minutes
  - During link down, forwarding node stores packets
- when link up, data is de-stored and forwarded (still with 60s RTT)

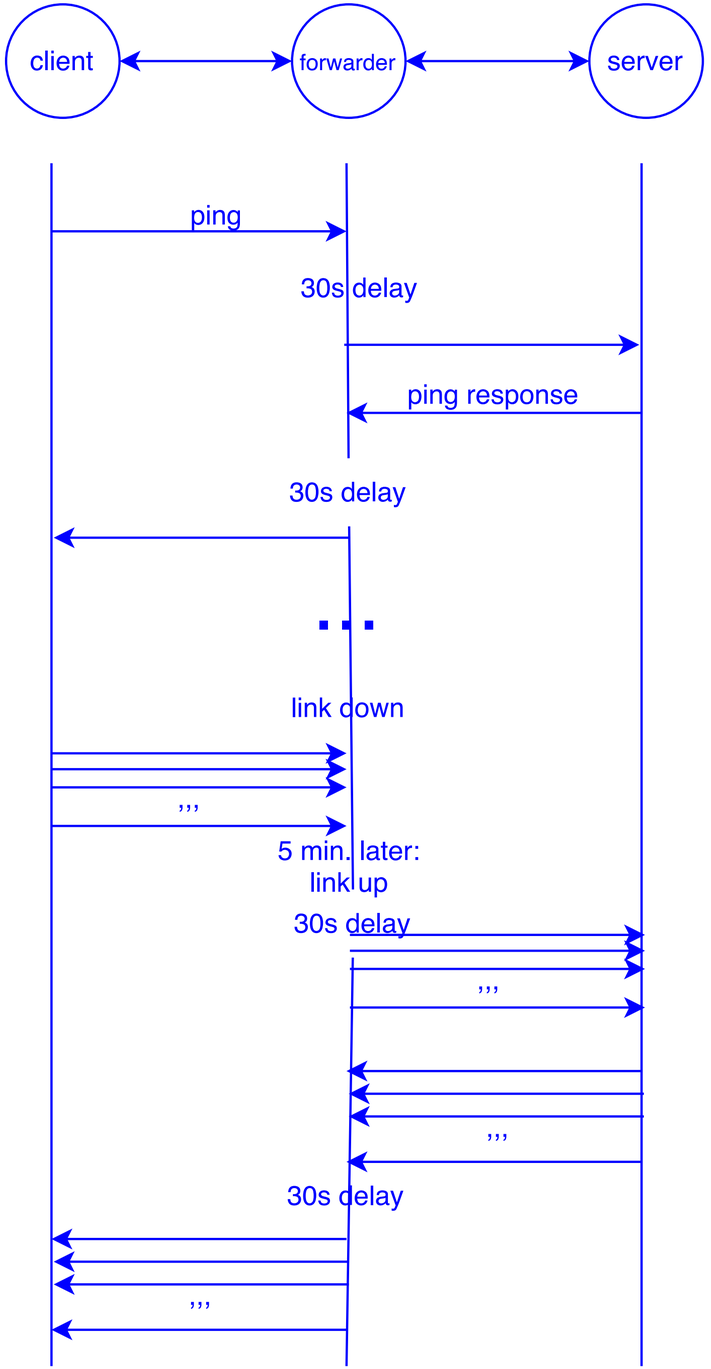

Link Interruptions

- ping -v -c 400 -W 100000 fc00:1::3
- PING fc00:1::3 (fc00:1::3) 56 data bytes
- 64 bytes from fc00:1::3: icmp_seq=1 ident=12259 ttl=62 time=60015 ms
- 64 bytes from fc00:1::3: icmp_seq=2 ident=12259 ttl=62 time=60004 ms
- 64 bytes from fc00:1::3: icmp_seq=3 ident=12259 ttl=62 time=60004 ms
- ...
- 64 bytes from fc00:1::3: icmp_seq=35 ident=12259 ttl=62 time=60004 ms
- 64 bytes from fc00:1::3: icmp_seq=36 ident=12259 ttl=62 time=60004 ms
- **64 bytes from fc00:1::3: icmp_seq=37 ident=12259 ttl=62 time=360184 ms**
- 64 bytes from fc00:1::3: icmp_seq=38 ident=12259 ttl=62 time=359160 ms
- 64 bytes from fc00:1::3: icmp_seq=39 ident=12259 ttl=62 time=358137 ms
- 64 bytes from fc00:1::3: icmp_seq=40 ident=12259 ttl=62 time=357113 ms
- ...
- 64 bytes from fc00:1::3: icmp_seq=317 ident=12259 ttl=62 time=74110 ms
- 64 bytes from fc00:1::3: icmp_seq=318 ident=12259 ttl=62 time=73087 ms
- 64 bytes from fc00:1::3: icmp_seq=319 ident=12259 ttl=62 time=72065 ms
- ...
- 64 bytes from fc00:1::3: icmp_seq=398 ident=12259 ttl=62 time=60004 ms
- 64 bytes from fc00:1::3: icmp_seq=399 ident=12259 ttl=62 time=60004 ms
- 64 bytes from fc00:1::3: icmp_seq=400 ident=12259 ttl=62 time=60004 ms
- --- fc00:1::3 ping statistics ---
- 400 packets transmitted, 400 received, **0% packet loss**, time 407896ms
- rtt min/avg/max/mdev = 60003.980/170727.305/360183.665/99770.604 ms, pipe 353

- 60 seconds RTT
- 5 minutes interruption
- link back
- Implemented forwarding policy is FIFO, so the « oldest» is forwarded first.
- After all stored packets forwarded, then the packets are forwarded directly

New Simulation Environment

- Simulations were time consuming and progressing slow:
  - Wall clock with 10 days RTT, means ~1 month to finish the simulation (connection establishment, request-response, connection close)
  - Had to change code in API each time.
- Wanted to do more QUIC simulations, « faster ».
- Adolfo Ochagavia implemented a Quinn workbench that:
  - Put in a JSON files the transport parameters and the network characteristics one wants to simulate
  - Do time warping: a 20 days RTT simulation is run in 2 seconds.
  - Set seeds so that it is fully repeatable with the exact same results
  - Creates the pcap file and associated SSLKEYLOG file to decrypt the packets
  - Limitation: only test the Quinn QUIC stack and only testing HTTP.
    - The « heavy » testbed is still useful to test other protocols and to do system testing
    - Does not seem easy to port it to other QUIC stacks.

Delay 1 day, loss 5%, repeat 10, QUIC-large values

- One way 1 day delay, packet loss 5%, 10 times repeat HTTP request and response, QUIC stack set with large values (space compatible)
- Total time: 1987200s (same as without packet loss, since loss was recovered using the next packets)
- Client packets sent: 22 (3087 bytes)
- Server packets sent: 20 (12313 bytes)
  - Server packets dropped: 2
  - (by the network simulation)
- Conclusion: QUIC recovered successfully and all data were properly sent reliably ✅

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000	1.1.1.1	88.88.88.88	QUIC	1228	Initial, DCID=22a3467b8c1180a3eeba67d7dfc1fe9b8e9111ff, PKN: 0, CRYPTO, PADDING
2	86400.000000	88.88.88.88	1.1.1.1	QUIC	1228	Handshake, PKN: 0, CRYPTO
3	172800.000000	1.1.1.1	88.88.88.88	QUIC	1228	Handshake, PKN: 0, ACK_ECN, CRYPTO
4	172800.000000	1.1.1.1	88.88.88.88	QUIC	64	Protected Payload (KP0), PKN: 1, STREAM(0)
5	259200.000000	88.88.88.88	1.1.1.1	QUIC	427	Protected Payload (KP0), PKN: 1, DONE, AF, CRYPTO
6	259200.000000	88.88.88.88	1.1.1.1	QUIC	1074	Protected Payload (KP0), PKN: 2, STREAM(0)
7	345600.000000	1.1.1.1	88.88.88.88	QUIC	54	Protected Payload (KP0), PKN: 2, ACK_ECN
8	345600.000000	1.1.1.1	88.88.88.88	QUIC	64	Protected Payload (KP0), PKN: 3, STREAM(4)
9	432000.000000	88.88.88.88	1.1.1.1	QUIC	54	Protected Payload (KP0), PKN: 3, ACK_ECN
10	432000.000000	88.88.88.88	1.1.1.1	QUIC	1074	Protected Payload (KP0), PKN: 4, STREAM(4)
11	518400.000000	1.1.1.1	88.88.88.88	QUIC	64	Protected Payload (KP0), PKN: 4, STREAM(8)
12	604800.000000	88.88.88.88	1.1.1.1	QUIC	1074	Protected Payload (KP0), PKN: 5, STREAM(8)
13	691200.000000	1.1.1.1	88.88.88.88	QUIC	61	Protected Payload (KP0), PKN: 5, ACK_ECN
14	691200.000000	1.1.1.1	88.88.88.88	QUIC	64	Protected Payload (KP0), PKN: 6, STREAM(12)

...

33	1555200.000000	1.1.1.1	88.88.88.88	QUIC	61	Protected Payload (KP0), PKN: 15, ACK_ECN
34	1555200.000000	1.1.1.1	88.88.88.88	QUIC	64	Protected Payload (KP0), PKN: 16, STREAM(32)
35	1641600.000000	88.88.88.88	1.1.1.1	QUIC	61	Protected Payload (KP0), PKN: 16, ACK_ECN
36	1641600.000000	88.88.88.88	1.1.1.1	QUIC	1074	Protected Payload (KP0), PKN: 17, STREAM(32)
37	1728000.000000	1.1.1.1	88.88.88.88	QUIC	61	Protected Payload (KP0), PKN: 17, ACK_ECN
38	1728000.000000	1.1.1.1	88.88.88.88	QUIC	64	Protected Payload (KP0), PKN: 18, STREAM(36)
39	1814400.000000	88.88.88.88	1.1.1.1	QUIC	61	Protected Payload (KP0), PKN: 18, ACK_ECN
40	1814400.000000	88.88.88.88	1.1.1.1	QUIC	1074	Protected Payload (KP0), PKN: 19, STREAM(36)
41	1900800.000000	1.1.1.1	88.88.88.88	QUIC	61	Protected Payload (KP0), PKN: 19, ACK_ECN
42	1900800.000000	1.1.1.1	88.88.88.88	QUIC	57	Protected Payload (KP0), PKN: 20, ACK_ECN, CC
43	1987200.000000	88.88.88.88	1.1.1.1	QUIC	61	Protected Payload (KP0), PKN: 20, ACK_ECN
44	1987200.000000	88.88.88.88	1.1.1.1	QUIC	58	Protected Payload (KP0), PKN: 21, ACK_ECN, CC

60 days RTT!, packet loss 5%

cargo run --release -- --config example-configs/dtn-delay30d-loss5.json

--- Params ---

* Quinn seed: 0

* Network seed: 42

* Transport config path: example-configs/dtn-delay30d-loss5.json

* Delay: 2592000.00s (5184000.00s RTT)

* Extra delay (10.00% chance): 0.20s

5184000s = 60 days

* **Packet loss ratio: 5.00%**

* Packet duplication ratio: 0.00%

--- Requests ---

0.00s CONNECT

5184000.20s GET /index.html

10368000.20s GET /index.html

15552000.20s GET /index.html

20736000.20s GET /index.html

25920000.40s GET /index.html

25920000.40s WARN Client packet lost (#15)!

42120001.02s WARN Client packet lost (#19)!

52488001.02s GET /index.html

57672001.02s GET /index.html

62856001.02s GET /index.html

68040001.22s GET /index.html

73224001.22s GET /index.html

78408001.22s Done sending 10 requests

81000001.22s Connection closed

81000001s > 900 days!

remember: 10 requests/response

--- Stats ---

* Time from start to connection closed: 81000001.22s (15.63 RTT)

* Client packets successfully sent: 17 (2895 bytes)

  * From the above packets, 0 were duplicates (0 bytes)

  * From the above packets, 0 were received out of order by the peer (0 bytes)

* Client packets dropped: 2 (58 bytes)

* Server packets successfully sent: 21 (12650 bytes)

  * From the above packets, 0 were duplicates (0 bytes)

  * From the above packets, 0 were received out of order by the peer (0 bytes)

* Server packets dropped: 0 (0 bytes)

60 days RTT (cont.)

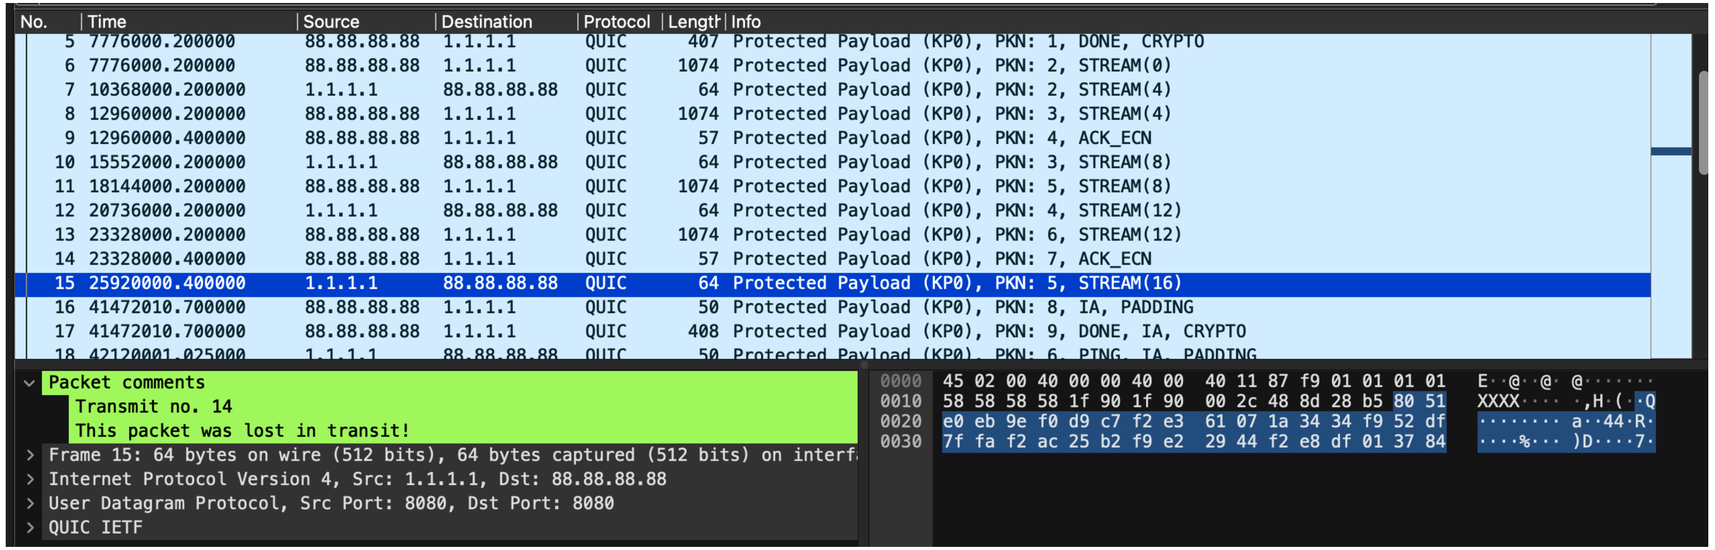

No.	Time	Source	Destination	Protocol	Length	Info
5	7776000.200000	88.88.88.88	1.1.1.1	QUIC	407	Protected Payload (KP0), PKN: 1, DONE, CRYPTO
6	7776000.200000	88.88.88.88	1.1.1.1	QUIC	1074	Protected Payload (KP0), PKN: 2, STREAM(0)
7	10368000.200000	1.1.1.1	88.88.88.88	QUIC	64	Protected Payload (KP0), PKN: 2, STREAM(4)
8	12960000.200000	88.88.88.88	1.1.1.1	QUIC	1074	Protected Payload (KP0), PKN: 3, STREAM(4)
9	12960000.400000	88.88.88.88	1.1.1.1	QUIC	57	Protected Payload (KP0), PKN: 4, ACK_ECN
10	15552000.200000	1.1.1.1	88.88.88.88	QUIC	64	Protected Payload (KP0), PKN: 3, STREAM(8)
11	18144000.200000	88.88.88.88	1.1.1.1	QUIC	1074	Protected Payload (KP0), PKN: 5, STREAM(8)
12	20736000.200000	1.1.1.1	88.88.88.88	QUIC	64	Protected Payload (KP0), PKN: 4, STREAM(12)
13	23328000.200000	88.88.88.88	1.1.1.1	QUIC	1074	Protected Payload (KP0), PKN: 6, STREAM(12)
14	23328000.400000	88.88.88.88	1.1.1.1	QUIC	57	Protected Payload (KP0), PKN: 7, ACK_ECN
15	25920000.400000	1.1.1.1	88.88.88.88	QUIC	64	Protected Payload (KP0), PKN: 5, STREAM(16)
16	41472010.700000	88.88.88.88	1.1.1.1	QUIC	50	Protected Payload (KP0), PKN: 8, IA, PADDING
17	41472010.700000	88.88.88.88	1.1.1.1	QUIC	408	Protected Payload (KP0), PKN: 9, DONE, IA, CRYPTO
18	42120001.025000	1.1.1.1	88.88.88.88	QUIC	50	Protected Payload (KP0), PKN: 6, PING, IA, PADDING

```
Packet comments
    Transmit no. 14
    This packet was lost in transit!
> Frame 15: 64 bytes on wire (512 bits), 64 bytes captured (512 bits) on interf
> Internet Protocol Version 4, Src: 1.1.1.1, Dst: 88.88.88.88
> User Datagram Protocol, Src Port: 8080, Dst Port: 8080
> QUIC IETF
```

```
0000  45 02 00 40 00 00 40 00  40 11 87 f9 01 01 01 01   E··@··@· @·······
0010  58 58 58 58 1f 90 1f 90  00 2c 48 8d 28 b5 80 51   XXXX···· ·,H·(··Q
0020  e0 eb 9e f0 d9 c7 f2 e3  61 07 1a 34 34 f9 52 df   ········ a··44·R·
0030  7f fa f2 ac 25 b2 f9 e2  29 44 f2 e8 df 01 37 84   ····%··· )D····7·
```

- Pcap also contains the loss packets, which is very useful to follow where the data has been resent later on. It is clearly identified in the packet comments (pcapng feature) and in the stdout log.

Next Steps

- Continue more simulations:
  - QUIC
  - System, other protocols, larger networks
- Look at different approaches to CC and various parameters
  - Current simulations are using a simplified CC which only does flow control: managing the window.
  - Might be possible to use BBR, but more work to be done