Measuring Starlink Protocol Performance

Geoff Huston
APNIC

From NANOG 92…

WHAT'S UP WITH STARLINK SPEEDS?

KEMAL ŠANJTA
PRINCIPAL INTERNET ANALYST
KEMALS@CISCO.COM

MIKE HICKS
PRINCIPAL SOLUTIONS ANALYST
MHICKS@CISCO.COM

CISCO ThousandEyes

NANOG™

From NANOG 92…

WHAT'S UP WITH STARLINK SPEEDS?

KEMAL ŠANJTA
PRINCIPAL INTERNET ANALYST
KEMALS@CISCO.COM

MIKE HICKS
PRINCIPAL SOLUTIONS ANALYST
MHICKS@CISCO.COM

CISCO ThousandEyes

N

CONCLUSION

- Underlying TCP mechanisms play a significant role in Starlink performance
- Congestion avoidance algorithm BBR shows significantly better results:
  - Between 1.4X and 18.4X difference when it comes to Download
  - Between 1.2X and 3.4X difference when it comes to Upload
- Time to switch to BBR, as default congestion avoidance algorithm, on the major operating systems?

From NANOG 92…

WHAT'S UP WITH STARLINK SPEEDS?

KEMAL ŠANJTA
PRINCIPAL INTERNET ANALYST
KEMALS@CISCO.COM

MIKE HICKS
PRINCIPAL SOLUTIONS ANALYST
MHICKS@CISCO.COM

CISCO ThousandEyes

CONCLUSION

WHY does BBR outperform CUBIC on Starlink systems?

… in Starlink

… avoidance algorithm BBR shows significantly better results:

- Between 1.4X and 18.4X difference when it comes to Download
- Between 1.2X and 3.4X difference when it comes to Upload

- Time to switch to BBR, as default congestion avoidance algorithm, on the major operating systems?

Low Earth Orbit

- LEO satellites are stations between 160km and 2,000km in altitude. Starlink operates at 550km (V2 includes additional orbital shells at 525km, 530km and 535km)
- High enough to stop it slowing down by “grazing” the denser parts of the earth’s ionosphere
- Not so high that it loses the radiation protection afforded by the Inner Van Allen belt.
- At a height of 550km, the minimum signal propagation delay to reach the satellite and back is **3.7ms**, at 25° elevation it’s **7.5ms**.

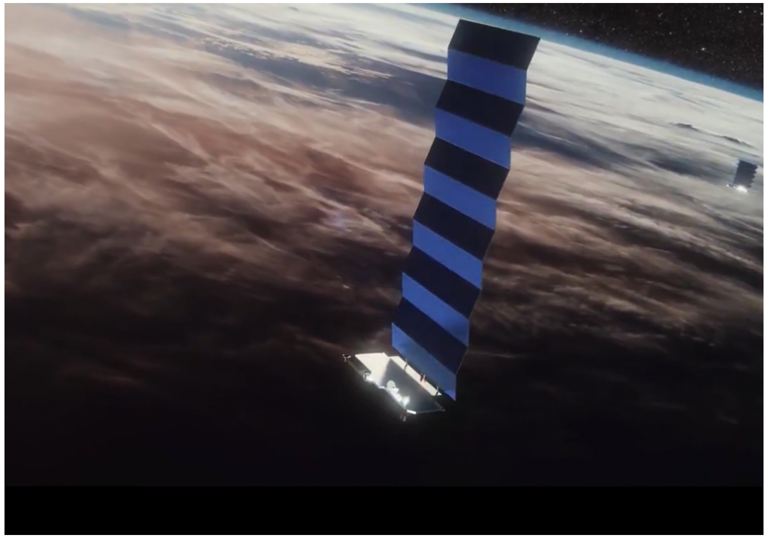
Image - spacex

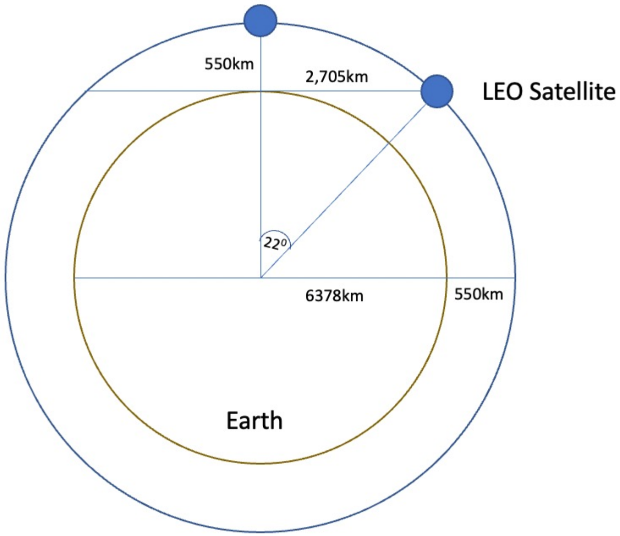

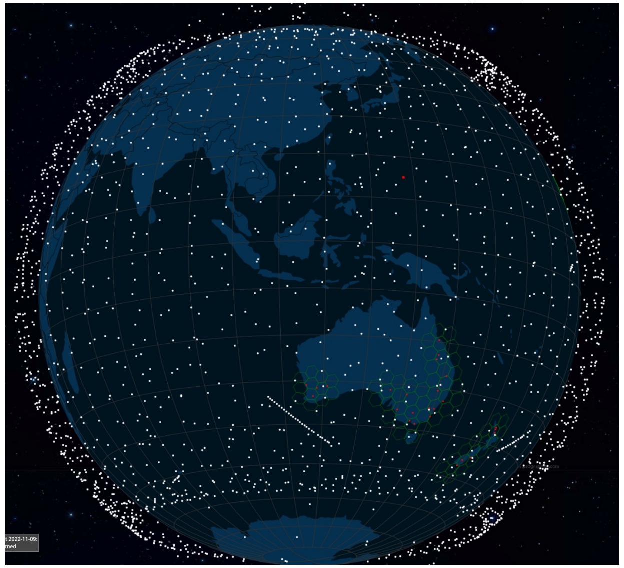
screenshot from starwatch app

Starlink is "interesting"!

- The spacecraft are very close to the Earth – which means:
  - They don’t need specialised high-power equipment to send and receive signals
    - Even hand-held mobile devices can send and receive signals with a LEO!
  - They can achieve very high signal speeds *
    - It’s an array of highly focussed signal beams
- But you need a large number of satellites to provide a continuous service
- The extremely high cost of launching a large constellation of LEO spacecraft has been the major problem with LEO service until recently

  (Which was a major reason why Motorola’s Iridium service went bankrupt soon after launch)

* Mobiles can’t achieve high speeds, but if you use an antenna then things improve dramatically!

Starlink Architecture

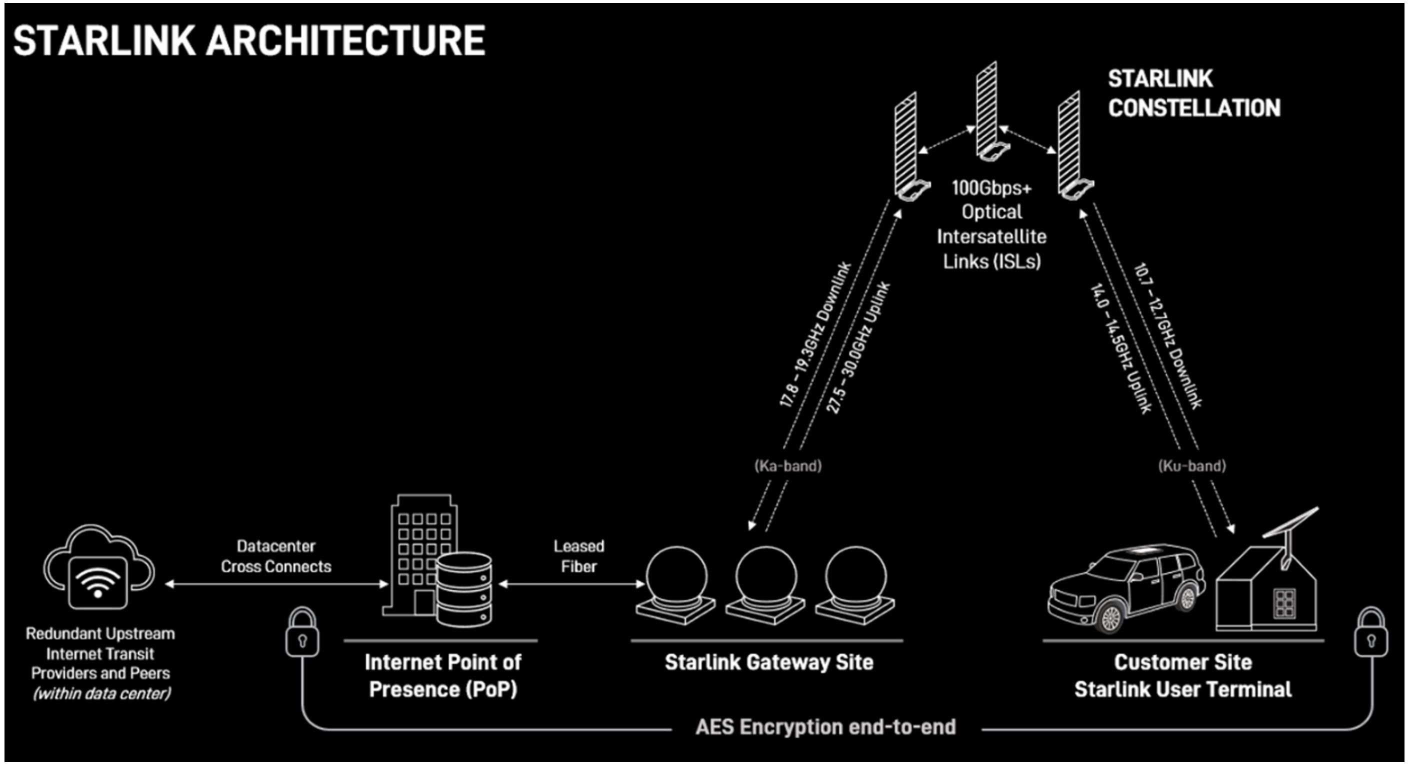

Tracking a LEO satellite

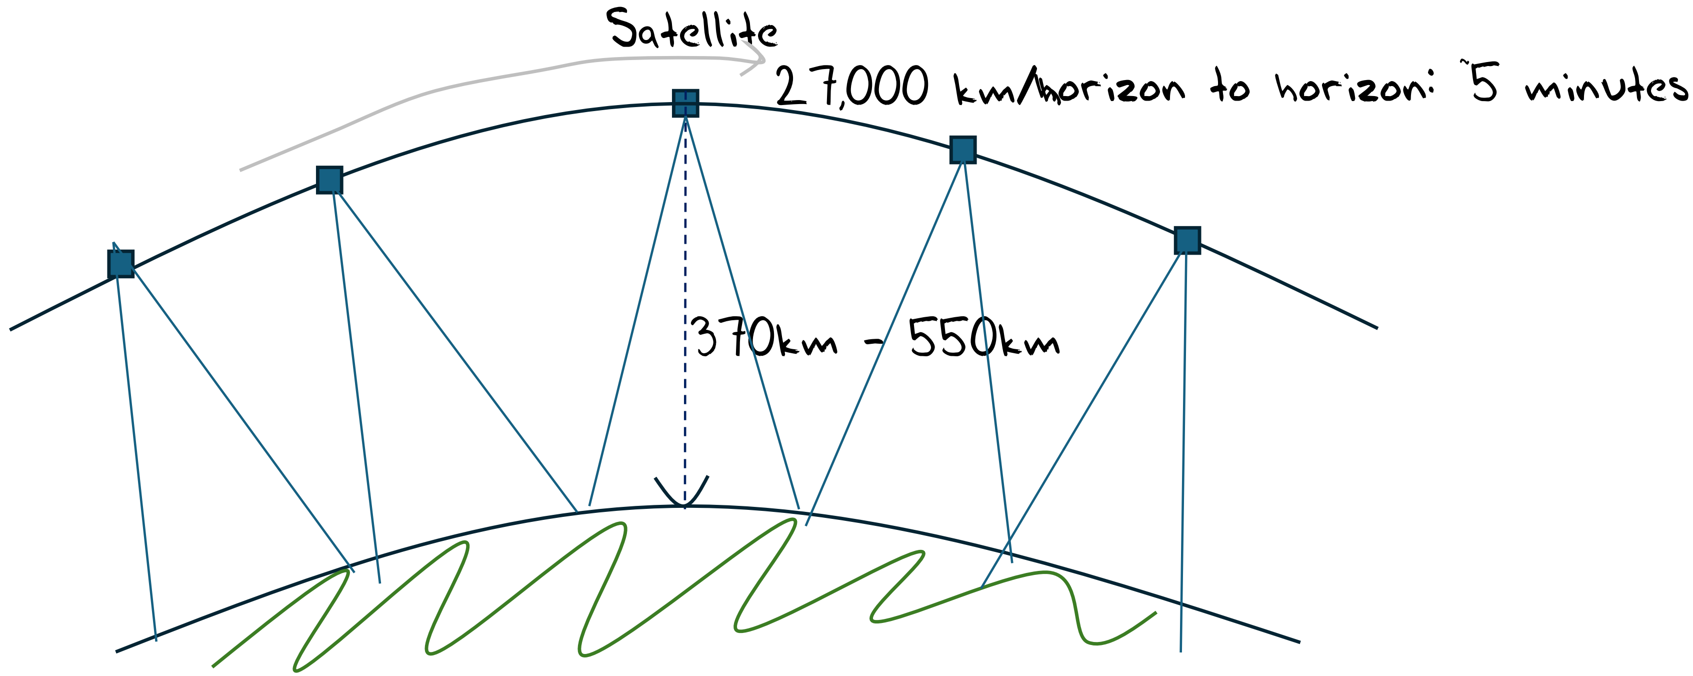

Looking Up

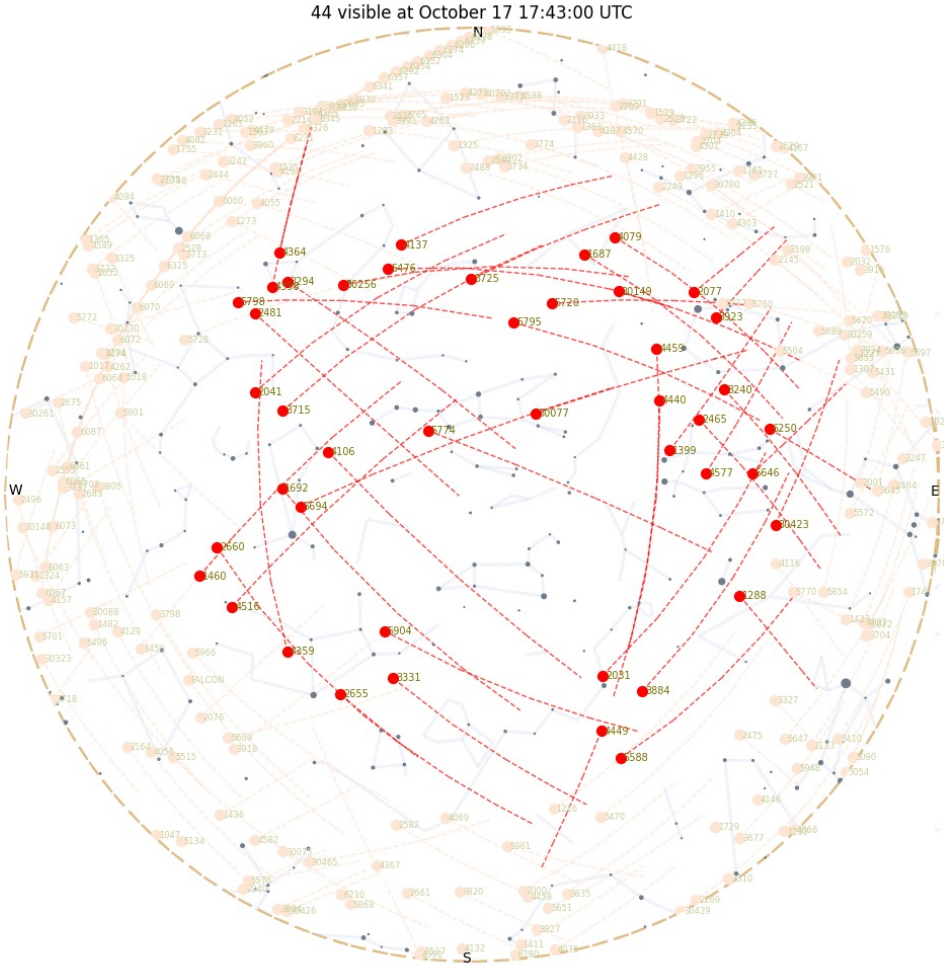

Starlink tracks satellites with a minimum elevation of 25°.

There are between 30 – 50 visible Starlink satellites at any point on the surface between latitudes 56° North and South

Each satellite traverses the visible aperture for a maximum of ~3 minutes

Starlink Scheduling

- A satellite is assigned to a user terminal in 15 second time slots
- Tracking of a satellite (by phased array focussing) works across 11 degrees of arc per satellite in each 15 second slot

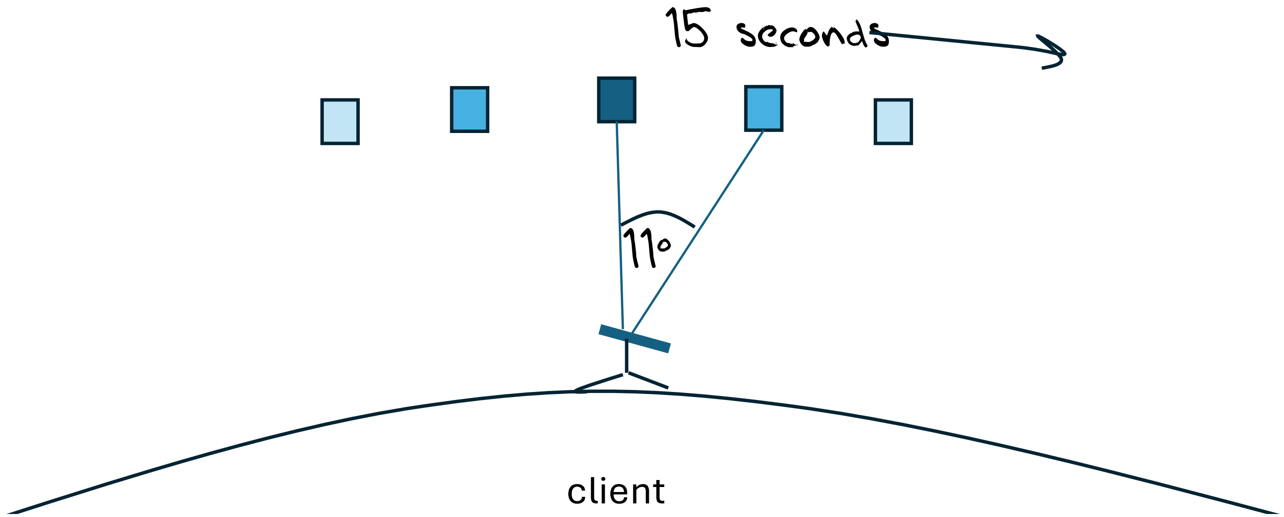

Starlink Spot Beams

- Each spacecraft uses 2,000 MHz of spectrum for user downlink and splits it into 8x channels of 250 MHz each
- Each satellite has 3 downlink antennas and 1 uplink antennas, and each can do 8 beams x 2 polarizations, for a total of 48 beams down and 16 up.

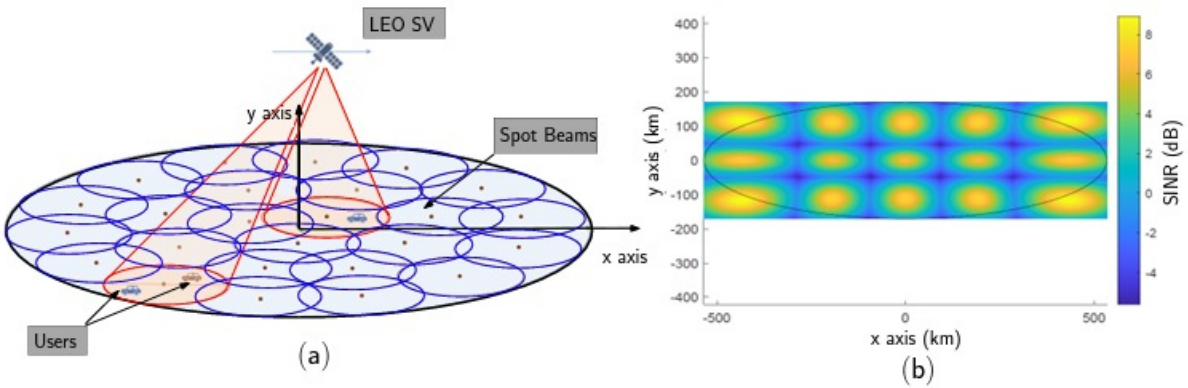

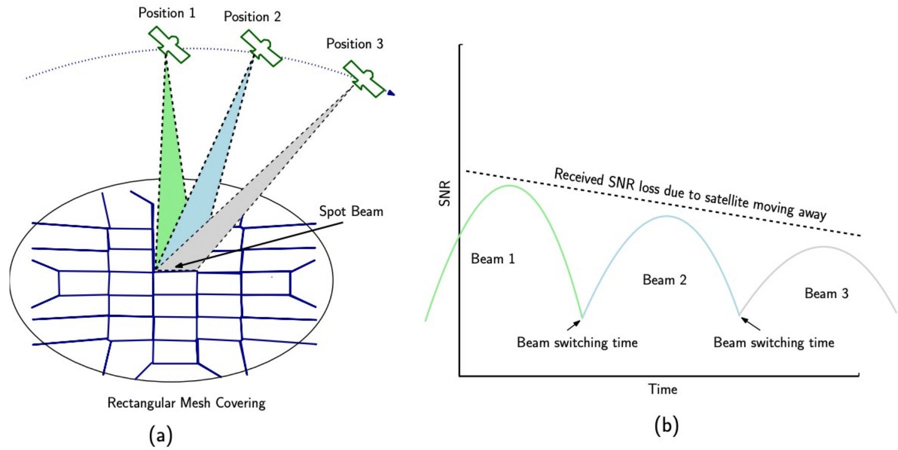

“Unveiling Beamforming Strategies of Starlink LEO Satellites”

https://people.engineering.osu.edu/sites/default/files/2022-10/Kassas_Unveiling_Beamforming_Strategies_of_Starlink_LEO_Satellites.pdf

Starlink's Reports

```
$ starlink-grpc-tools/dish_grpc_text.py -v  status
id:                      ut01000000-00000000-005dd555
hardware_version:        rev3_proto2
software_version:        5a923943-5acb-4d05-ac58-dd93e72b7862.uterm.release
state:                   CONNECTED
uptime:                  481674
snr:
seconds_to_first_nonempty_slot: 0.0
pop_ping_drop_rate:      0.0
downlink_throughput_bps: 16693.330078125
uplink_throughput_bps:   109127.3984375
pop_ping_latency_ms:     49.5
Alerts bit field:        0
fraction_obstructed:     0.041490074247l2181
currently_obstructed:    False
seconds_obstructed:
obstruction_duration:    1.9579976797103882
obstruction_interval:    540.0
direction_azimuth:       -42.67951583862305
direction_elevation:     64.61225128173828
is_snr_above_noise_floor: True
```

Reported Capacity & Latency

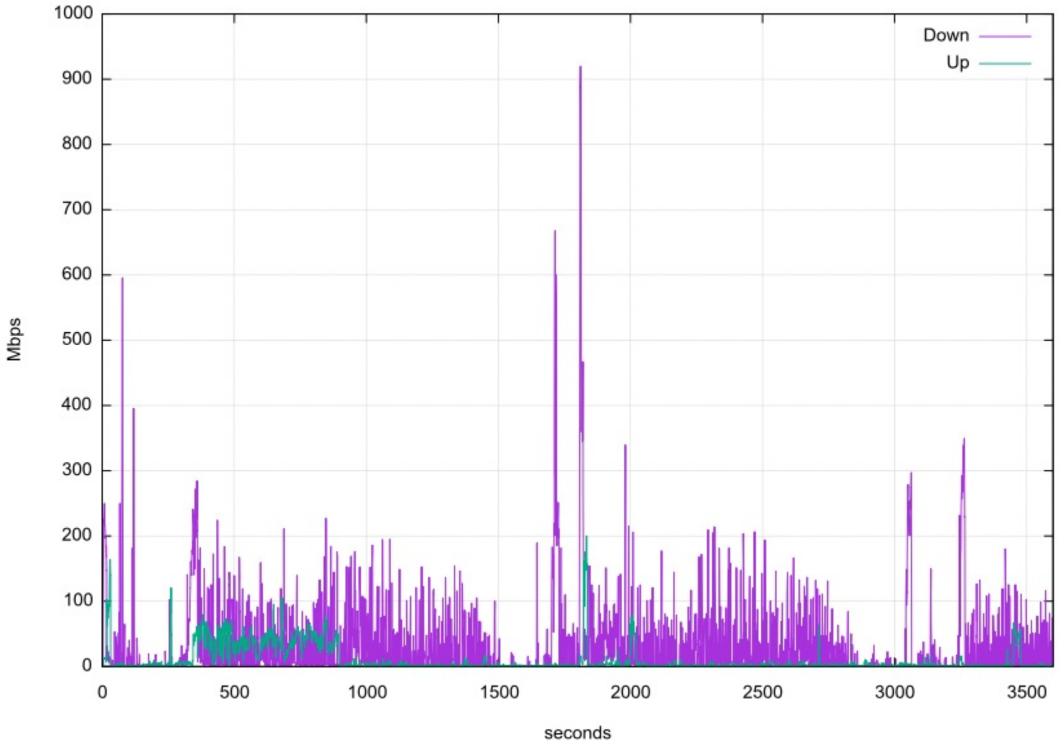

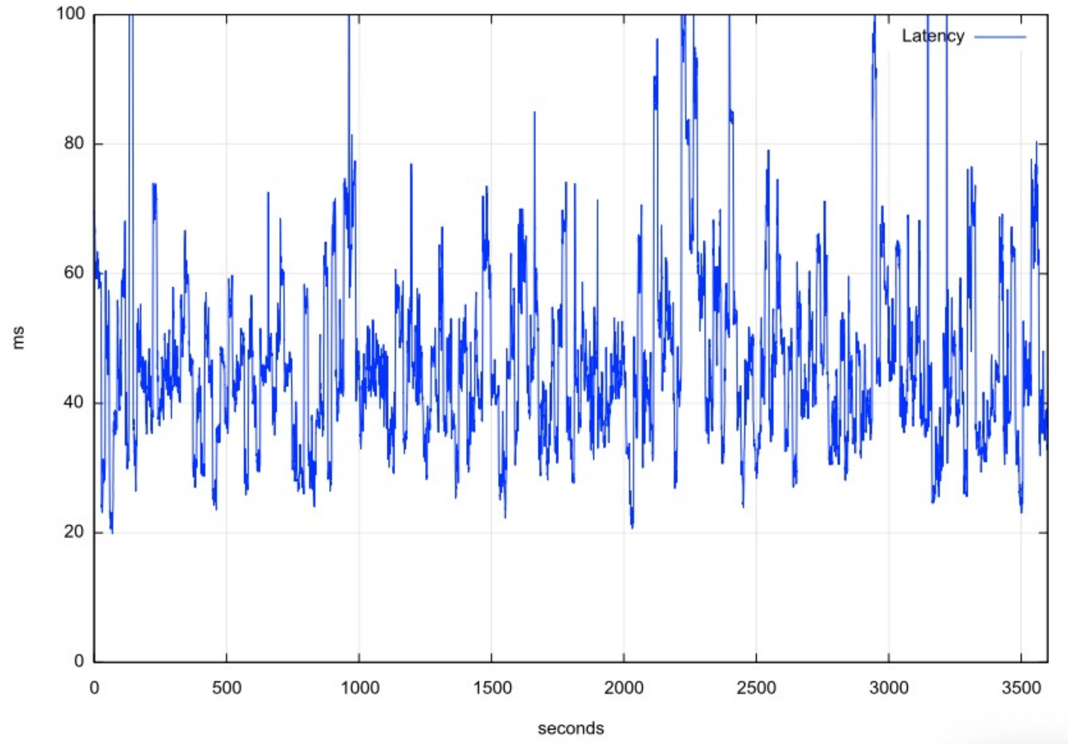

Reported Capacity & Latency

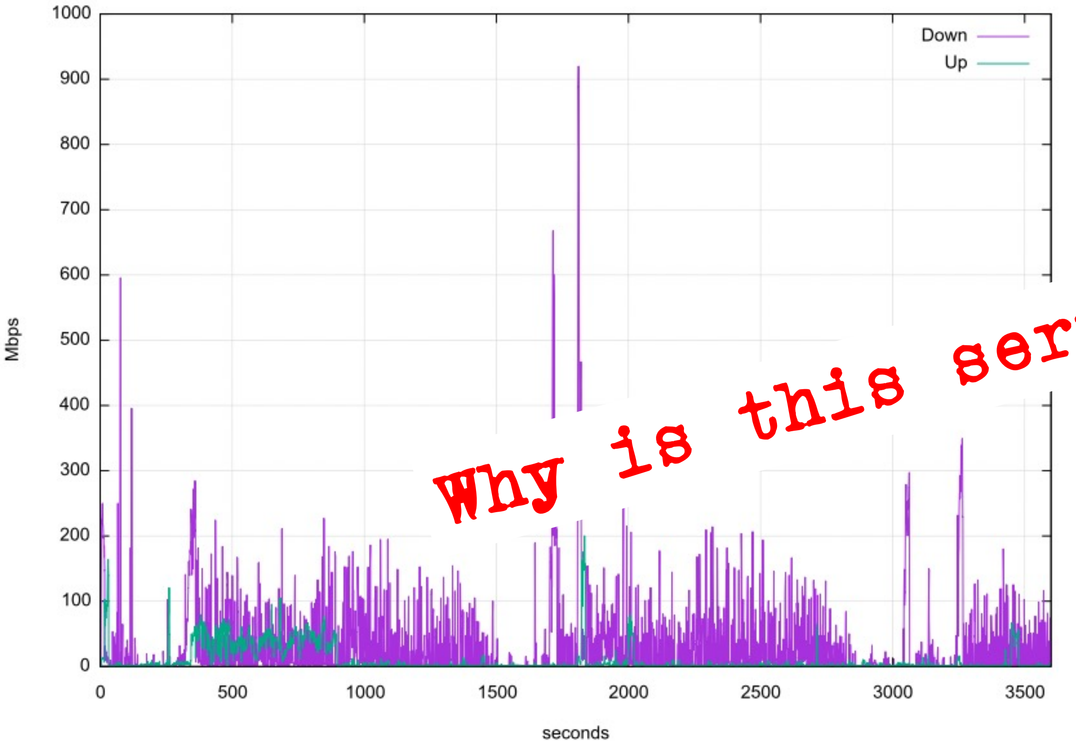

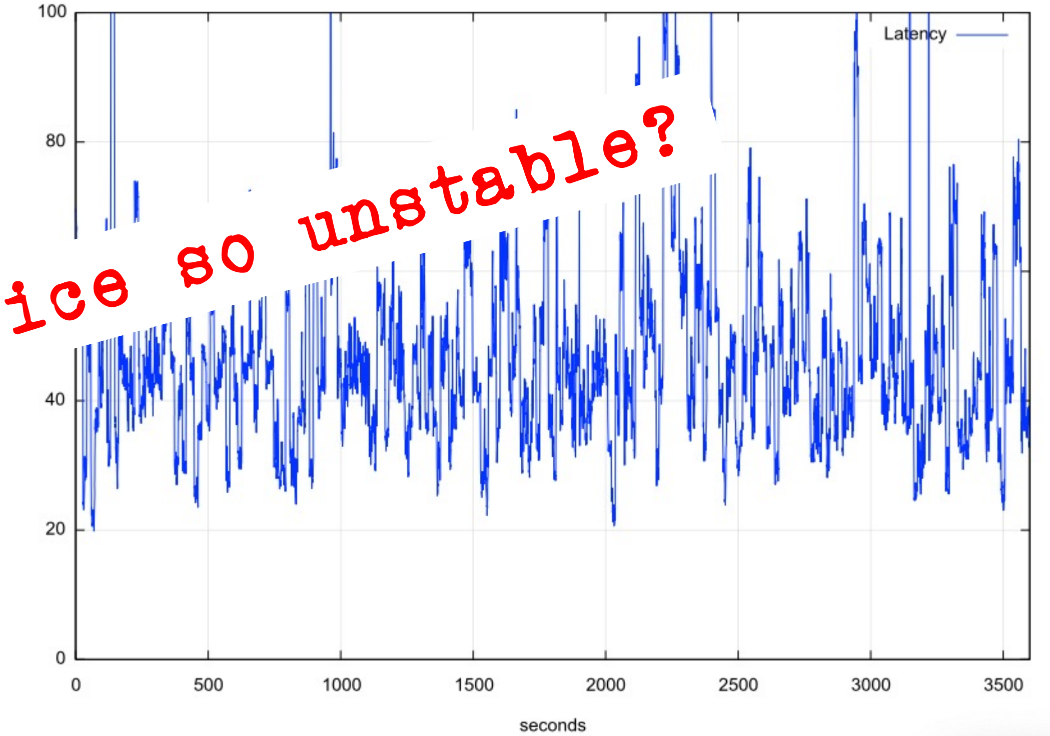

Why is this service so unstable?

Starlink Scheduling

- Latency changes on each satellite switch
- If we take the minimum latency on each 15 second scheduling interval, we can expose the effects of the switching interval on latency
- Across the 15 second interval there will be a drift in latency according to the satellite's track and the distance relative to the two earth points
- Other user traffic will also impact on latency, and also the effects of a large buffer in the user modem

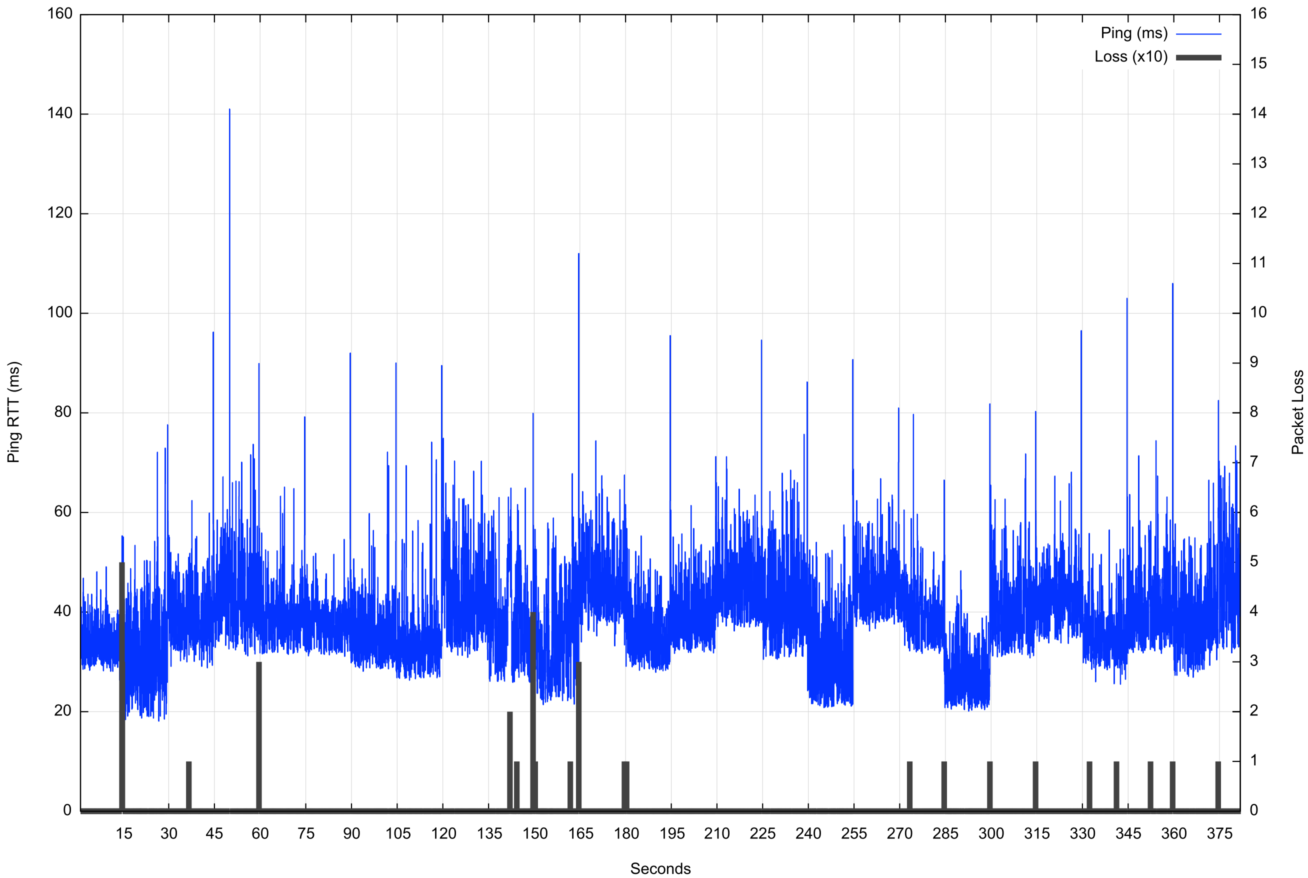

Satellite handover

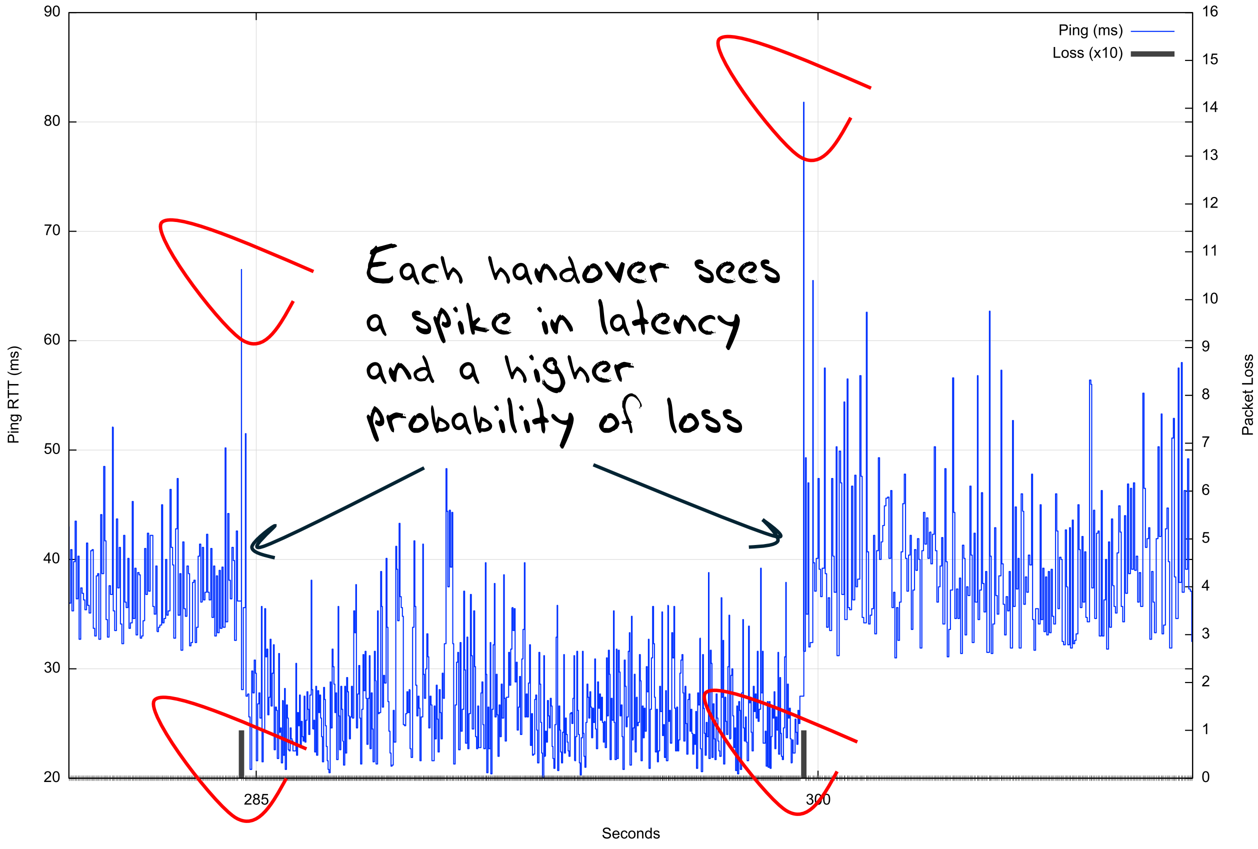

Satellite handover

- Packet loss occurs most frequently during handover events, and if confined to small set of packets
- This is NOT congestion-based tail-drop loss - so the packet loss can be generally repaired by a TCP SACK mechanism without needing a TCP session restart

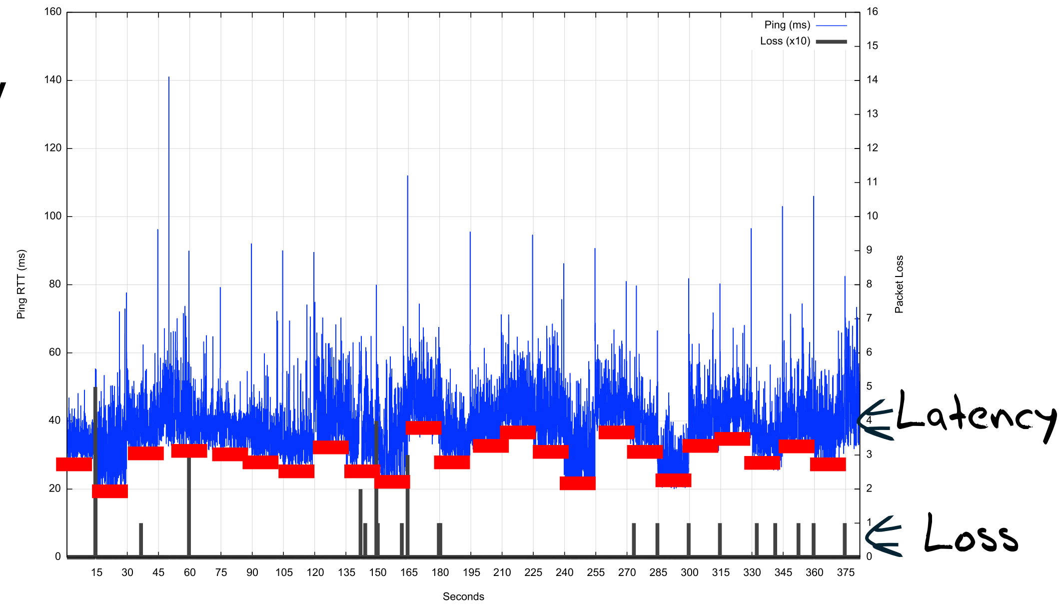

Latency

- The signal propagation latency to push a signal up and back to a Starlink spacecraft varies between 3.7ms and 7.5ms – which is an RTT of 7.4ms to 15ms
- But the minimum RTTs we see are 20ms!
- And the jitter is at an extremely high frequency (as is the change in channel capacity)

Varying SNR means varying signal modulation

- Starlink likely uses IEEE 802.11ac dynamic channel rate control, adjusting the signal modulation to match the current SNR
- This continual adjustment causes continual shift in the available capacity and imposes a varying latency on the round-trip time
- It also adds a few ms to the RTT times

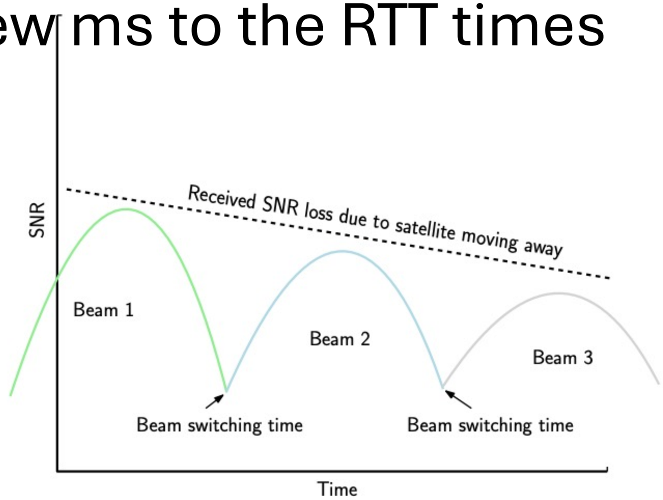

So What?

- The variation in latency and capacity occurs at high frequency, which means that TCP control is going to struggle to optimize itself against a shifting target
- TCP uses **ACK pacing** which means it attempts to optimize its sending rate over multiple RTT intervals

Frames

- Starlink does NOT provide each user with a dedicated frequency band
- The system uses multiplexing to divide a channel into frames and sends 750 frames per second. Each frame is divided into 302 intervals.
- Each frame carries a header that contains satellite, channel and modulation information

Starlink Characteristics

- Varying SNR produces varying modulation, which is expressed as varying capacity and delay
- Relative motions of earth and spacecraft add to varying latency
- 15 second satellite handover generates regular loss and latency extension
- Contention for common transmission medium leads to queuing delays

In Summary…

- The profile of a Starlink connection is governed by:
  - Regular handovers at precise 15 second intervals when the connection will experience:
    - a change in base latency
    - very short-term additional queuing latency of ~80 – 100ms
    - a high probability of non-congestion packet loss
  - Adaptive dynamic rate control by constantly changing signal modulation
    - High frequency jitter
    - No consistent upper bound to channel capacity
  - Common channel contention with other users

How well does all this work?

Speedtest measurements:

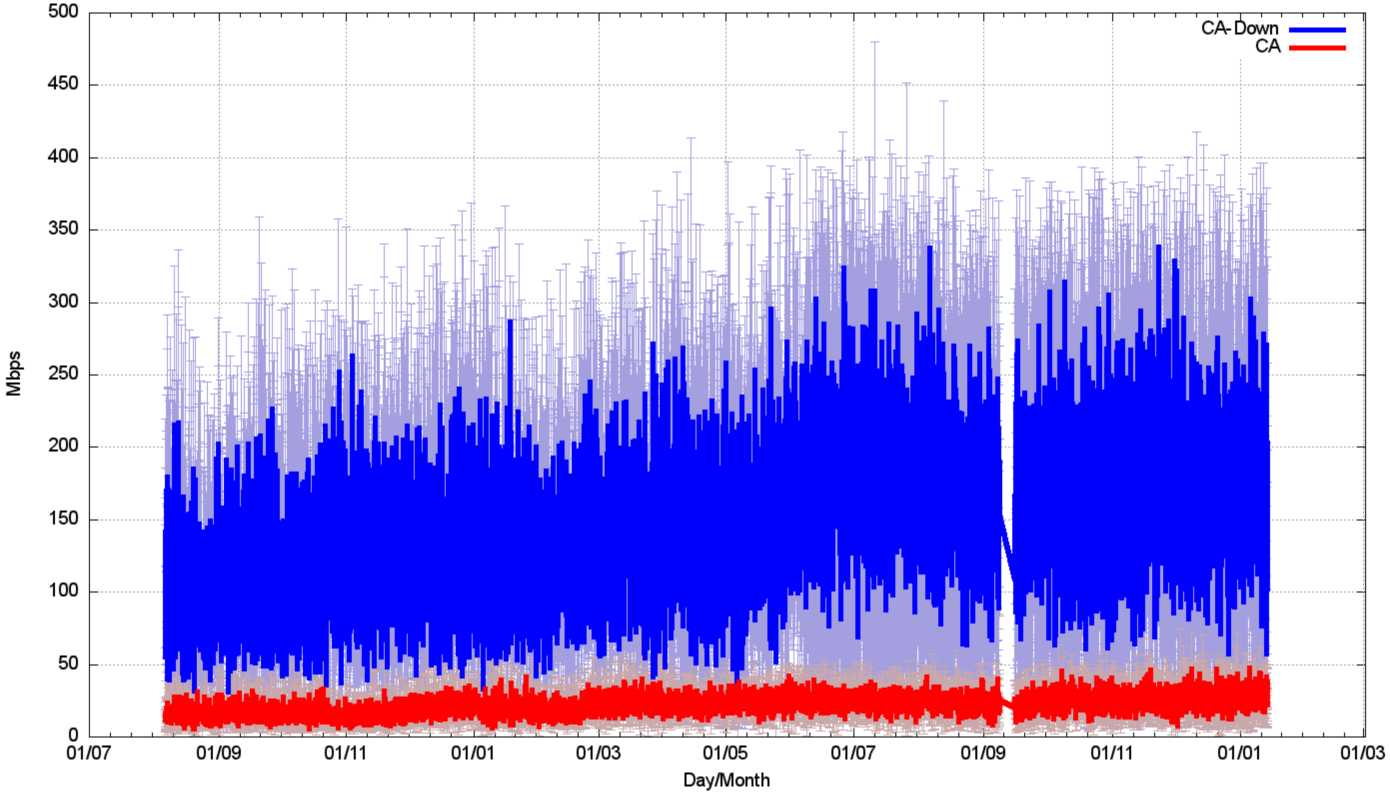

We should be able to get ~160Mbps out of a Starlink connection.

TCP Flow Control Algorithms

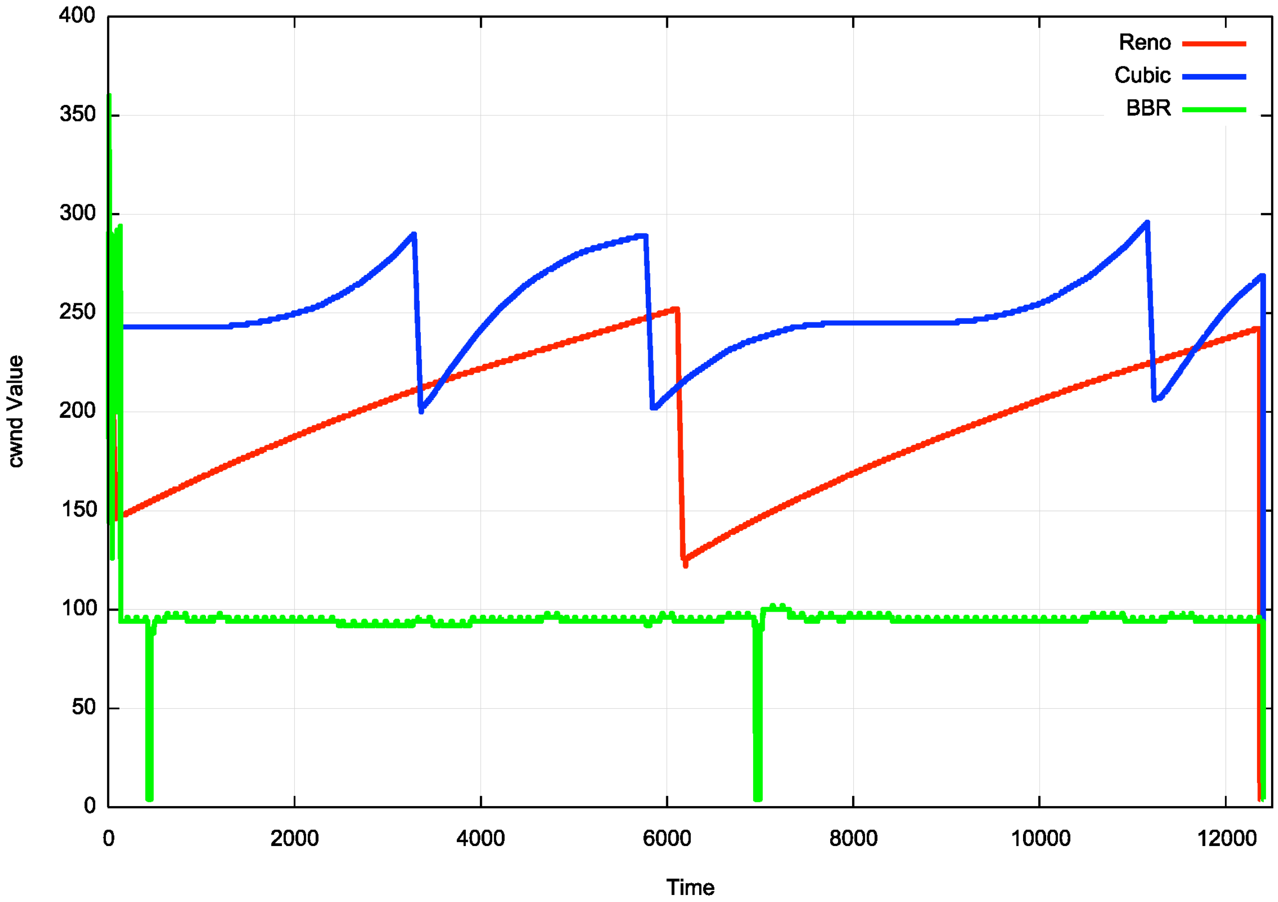

"Ideal" Flow behaviour for each protocol

iperf3 - cubic, 40 seconds

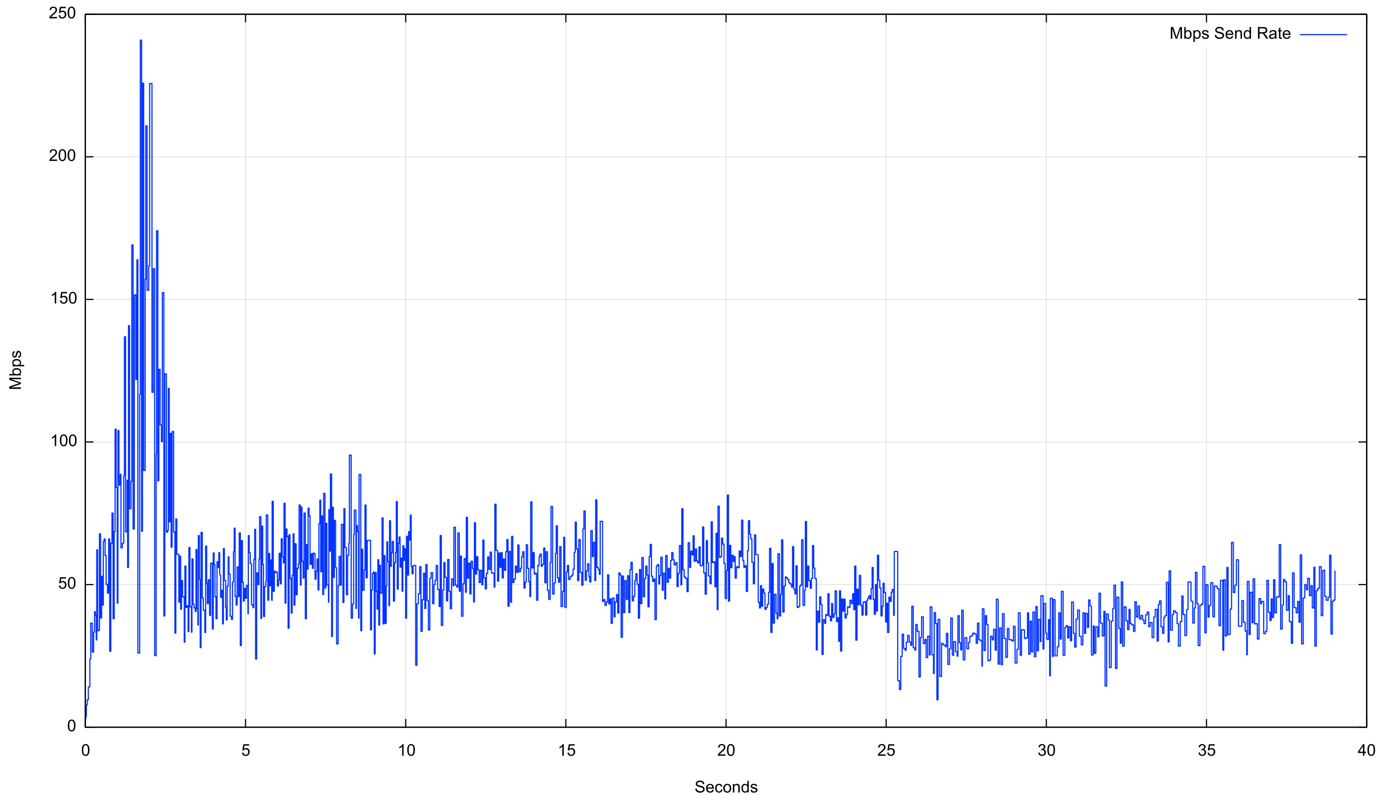

iperf3 - cubic, 40 seconds

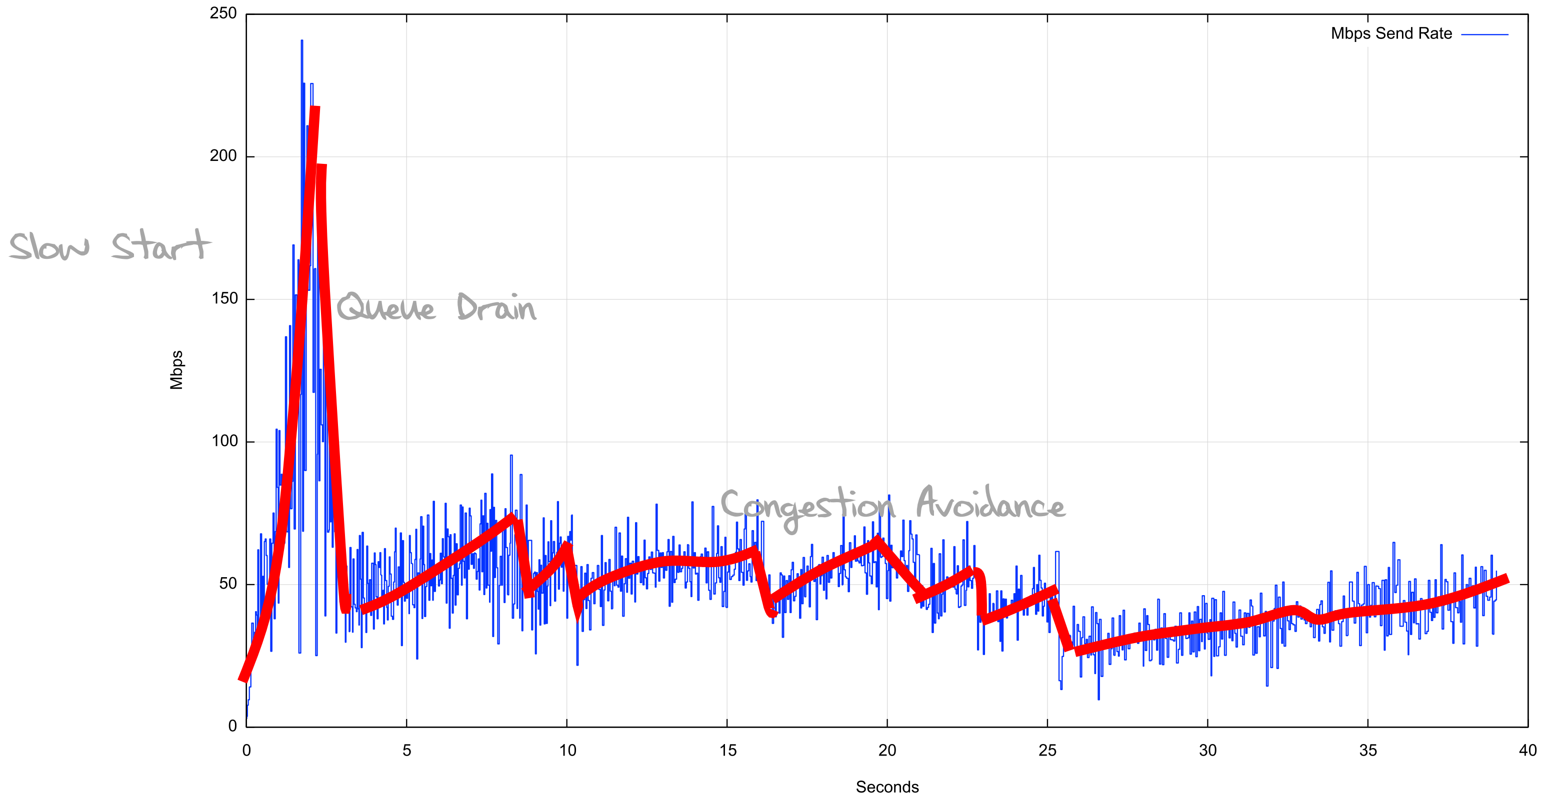

iperf3 - bbr

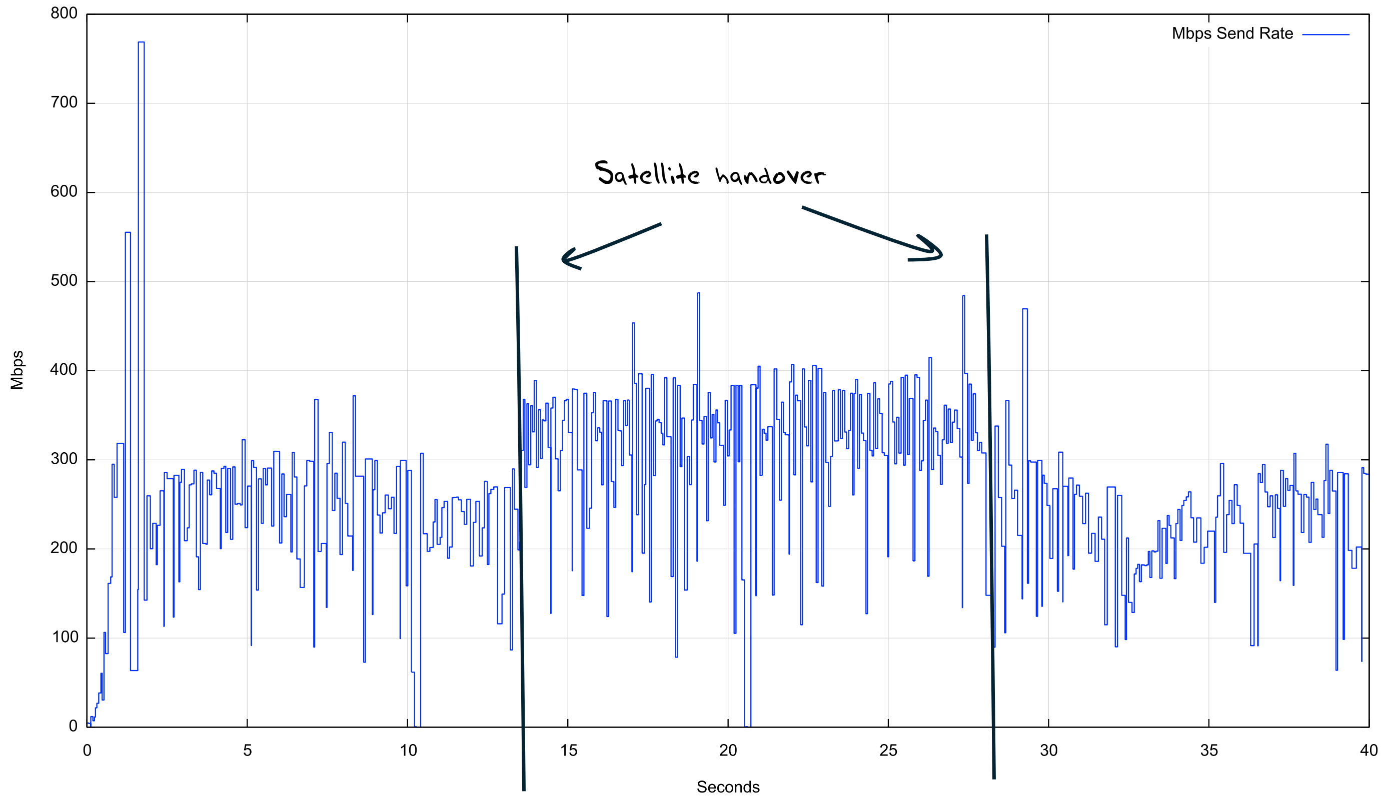

Mbps Send Rate
Satellite handover
Mbps
Seconds
0
100
200
300
400
500
600
700
800
0
5
10
15
20
25
30
35
40

Transport Protocol Considerations

- Starlink services have three issues:
  - Very high jitter rates – varying signal modulation
  - High levels of micro-loss (1.4%) – largely due to loss on satellite handover events (every 15 seconds)
  - Common bearer contention between users
- Loss-based flow control algorithms will over-react and pull back the sending rate over time
  - Short transactions work very well
  - Paced connections (voice, zoom, video streaming) tend to work well most of the time
- To obtain better performance you need to move to flow control algorithms that are not as loss-sensitive, such as BBR

Other considerations

- Senders should use **fair queuing** to pace sending rates and avoid bursting and tail drop behaviours
- **SACK** (selective acknowledgement) for TCP can help in rapid repair to multiple lost packets
- It's likely that **ECN** would also be really helpful to the protocol to disambiguate latency changes due to satellite behaviours and network queue buildup - i.e. only shrink the congestion window in response to ECN signals, not packet loss!

Starlink Performance

Starlink is perfectly acceptable for:

- short transactions
- video streaming
- conferencing

- The service can sustain 40 – 50Mbps delivery for long-held sessions during local peak use times in high density use scenarios
  - The isolated drop events generally do not intrude into the session state
- In off-peak and/or low-density contexts it can deliver 200-300Mbps
- Or, if the server uses BBR, then higher throughput is possible!
- It can be used in all kinds of places where existing wire and mobile radio systems either under-perform or aren't there at all!
- Its probably not the best trunk infrastructure service medium, but it's a really good high speed last mile direct retail access service, particularly for remote locations!

Making Starlink Faster

- Increase antennae transmitter power
- Use higher gain antennae with narrower beams
- Drop the orbital altitude to 340Km
- Drop the minimum elevation angle from 25º to 20º
- Use more bands (Ka-, V-, and E- bands)

(Proposed measures described in an October 2024 FCC application by Starlink)


Questions?