


# TCP and Starlink

Geoff Huston AM  
APNIC

APNIC 58



screenshot from starwatch app


# LEOs in the News

## Together Optus and SpaceX Plan to Cover 100% of Australia

12 July 2023, 04:00 PM

Science / Entertainment / More +

- Optus' collaboration with SpaceX aims to provide regional Australia with a new way to connect starting in late 2024.
- Optus plans to roll out SMS from late 2024, with voice and data also on the horizon from late 2025.



### TELECOMMUNICATION

## Elon Musk's Starlink launches satellite internet service in Japan

Company offers high-speed access to remote areas

Screenshot: <https://asia.nikkei.com/Business/Telecommunication/Elon-Musk-s-Starlink-launches-satellite-internet-service-in-Japan>

TECH / MOBILE / T-MOBILE

## T-Mobile and SpaceX Starlink say your 5G phone will connect to satellites next year

Screenshot - <https://www.theverge.com/2022/8/25/23320722/spacex-starlink-t-mobile-satellite-internet-mobile-messaging>

Home > News > Technology > Telco/ISP

## Telstra goes live with Starlink for homes

By Richard Chirgwin  
Mar 26 2024 12:10PM

10 tweets, 1 share, 1 email


Lower cost, marginally lower performance than buying direct.

<https://www.itnews.com.au/news/telstra-goes-live-with-starlink-for-homes-606423#:~:text=Telstra%20has%20kicked%20off%20its,the%20end%20of%20the%20year.>




# Newtonian Physics

- If you fire a projectile with a speed greater than 11.2Km/sec it will not fall back to earth, and instead head away from earth never to return
- On the other hand, if you incline the aiming trajectory and fire it at a critical speed it will settle into an orbit around the earth
- The higher the altitude, the lower the orbital speed required to maintain orbit




THAT by means of centripetal forces, the Planets may be retained in certain orbits, we may easily understand, if we consider the motions of projectiles. For a stone projected is by the pressure of its own weight forced out of the rectilinear path, which by the projection alone it should have pursued, and made to describe a curve line in the air; and through that crooked way is at last brought down to the ground. And the greater the velocity is with which it is projected, the farther it goes before it falls to the Earth. We may therefore suppose the velocity to be so encreased, that it would describe an arc of 1, 2, 5, 10, 100,

The effects of centripetal forces.



Title page of *Principia*, first edition (1687)

# Solar Radiation Physics





- The rotating iron core of the Earth produces a strong magnetic field
- This magnetic field deflects solar radiation – the Van Allen Belt
- Sheltering below the Van Allen Belt protects the spacecraft from the worst effects of solar radiation, allowing advanced electronics to be used in the spacecraft

# Low Earth Orbit

- LEO satellites are stations between 160km and 2,000km in altitude.
- 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° it’s 7.5ms.

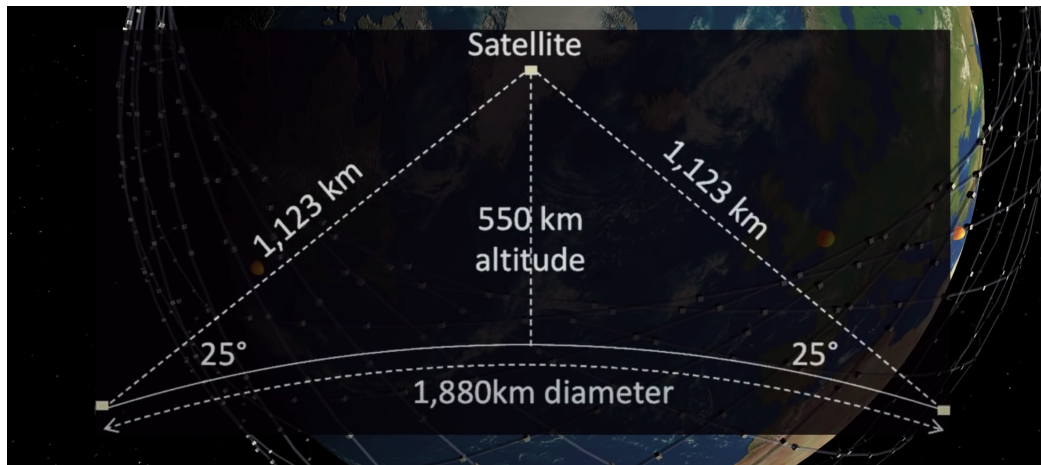


Image - spacex



screenshot from starwatch app

# Starlink Constellation





If you use a minimum angle of elevation of  $25^\circ$  then at an altitude of 550km each satellite spans a terrestrial footprint of no more than  $\sim 900\text{km}$  radius, or  $2\text{M K}^2$

At a minimum, a LEO satellite constellation needs 500 satellites to provide coverage of all parts of the earth's surface

For high quality coverage the constellation will need 6x-20x that number (or more!)

# Starlink Constellation

- 6,231 in-service operational spacecraft, operating at an altitude of 550km



<https://satellitemap.space/>




# So LEOs are "interesting"!

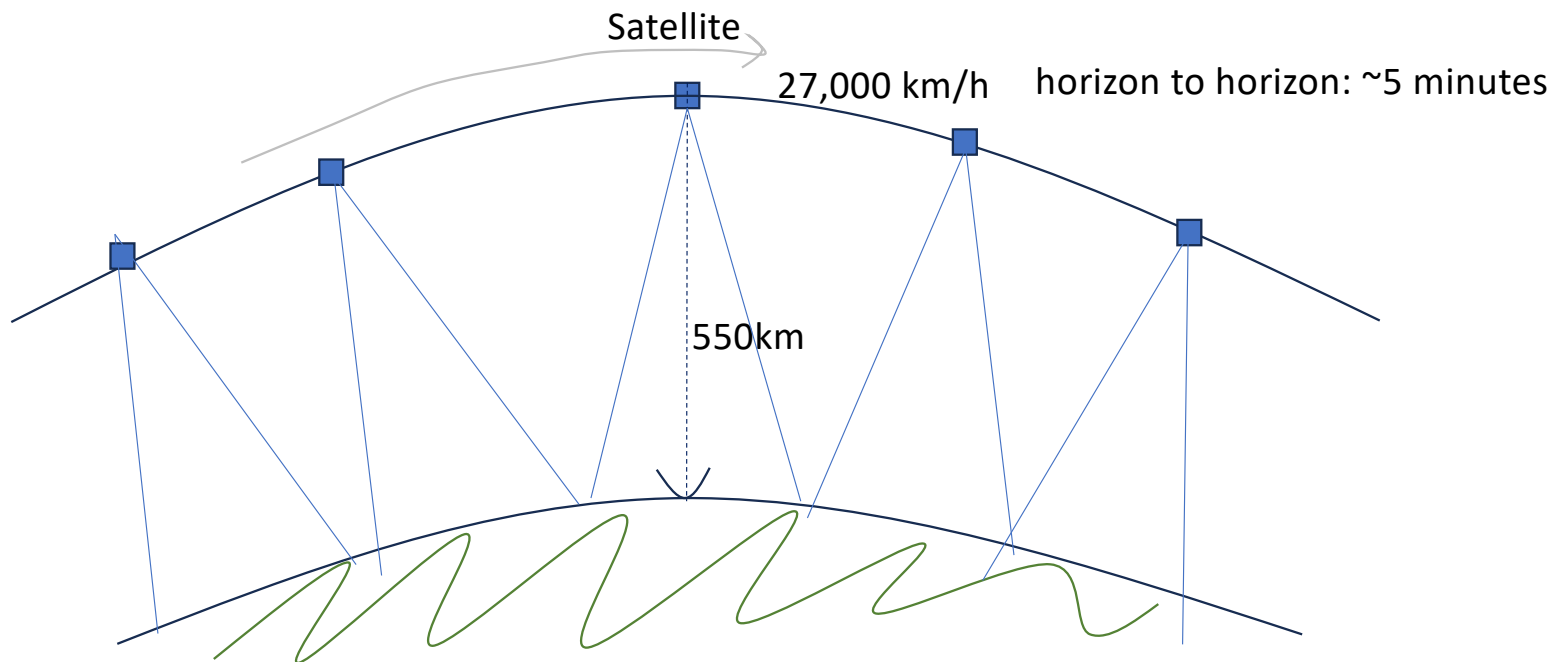
- They are very close to the Earth – which means:
  - They can achieve very high signal speeds
    - It's a highly focussed signal beam
  - They are harder to disrupt by external interference
  - They don't need specialised high-power equipment to send and receive signals
    - Even hand-held mobile devices can send and receive signals with a LEO (slowly!)
- But you need a large number of them 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 is why Motorola's Iridium service went bankrupt soon after launch



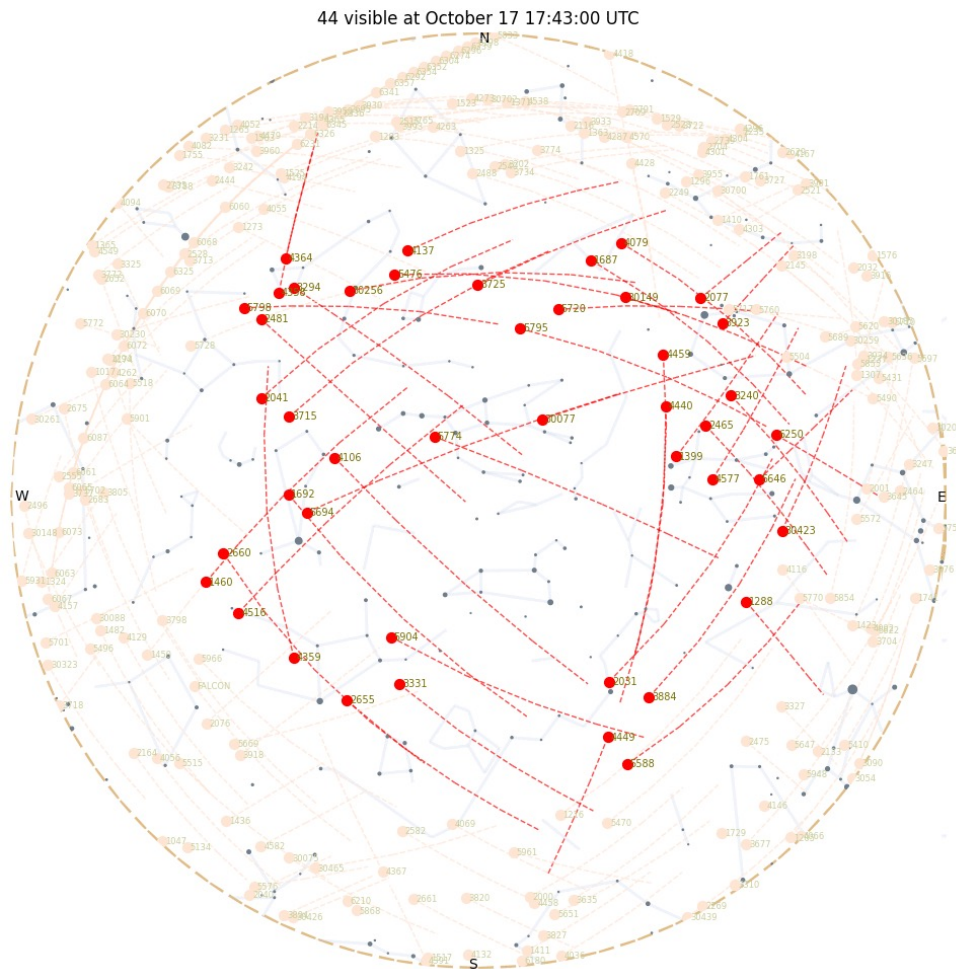
# Starlink Architecture



# 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

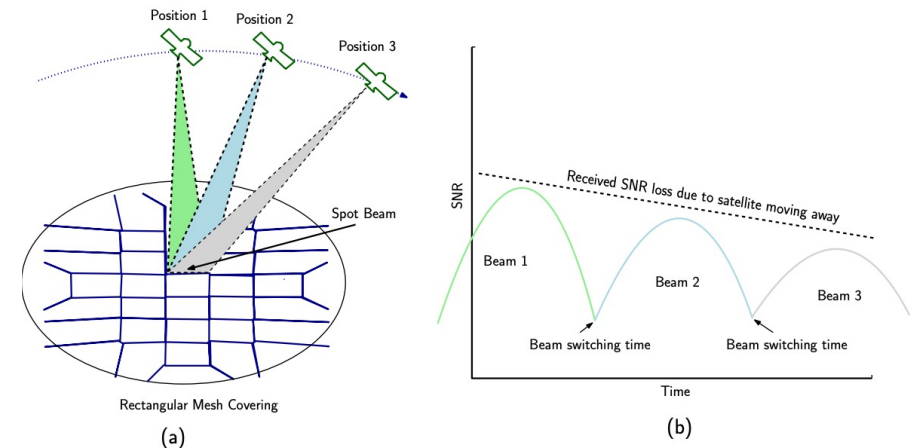
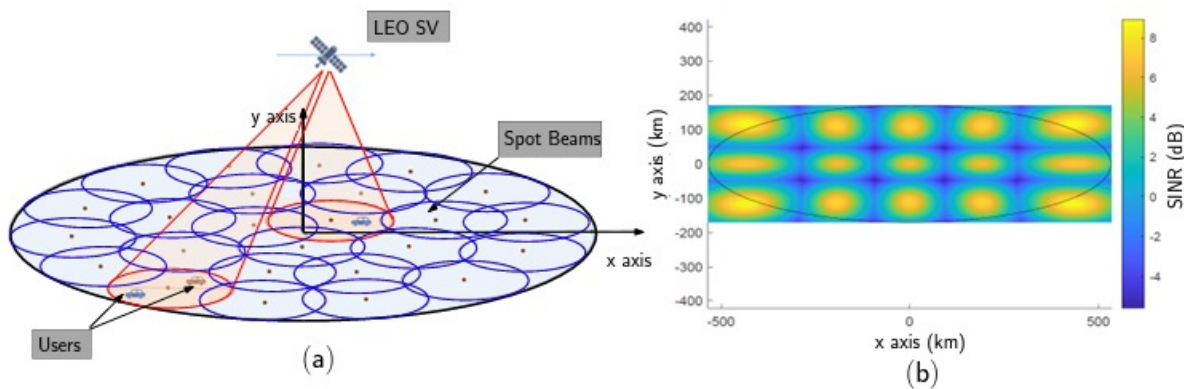
- A satellite is assigned to a user terminal in 15 second time slots
- Tracking of a satellite (by phased array focussing) works across 11 degrees of arc per satellite in each 15 second slot





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

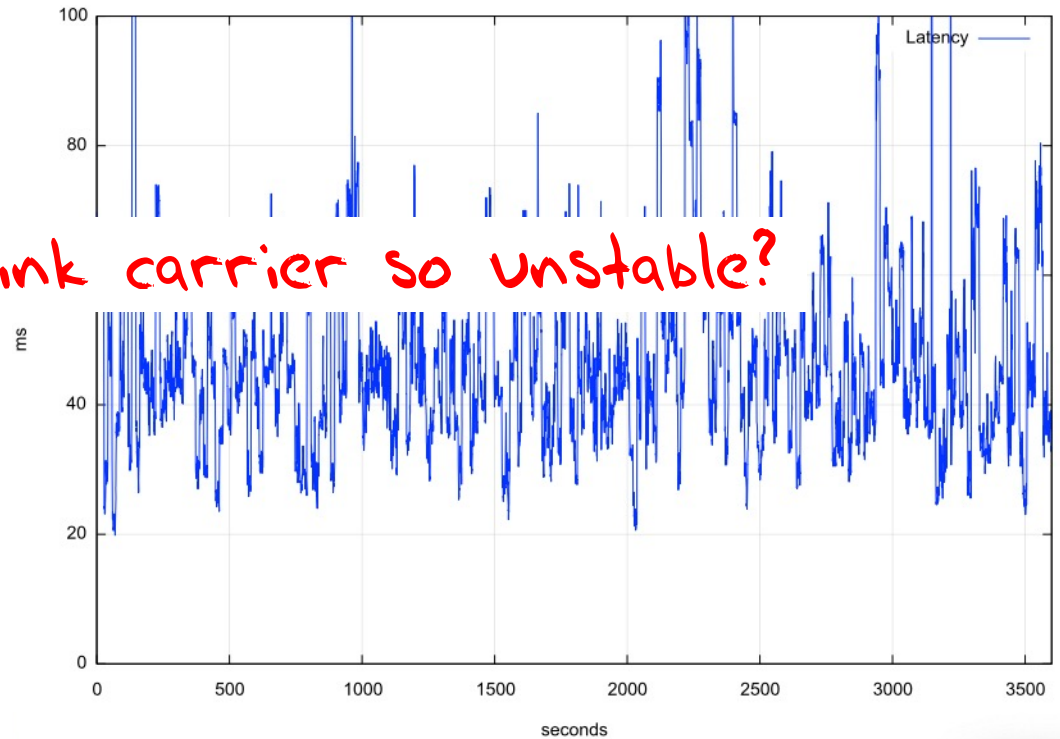
# Reported Capacity & Latency



# Reported Capacity & Latency

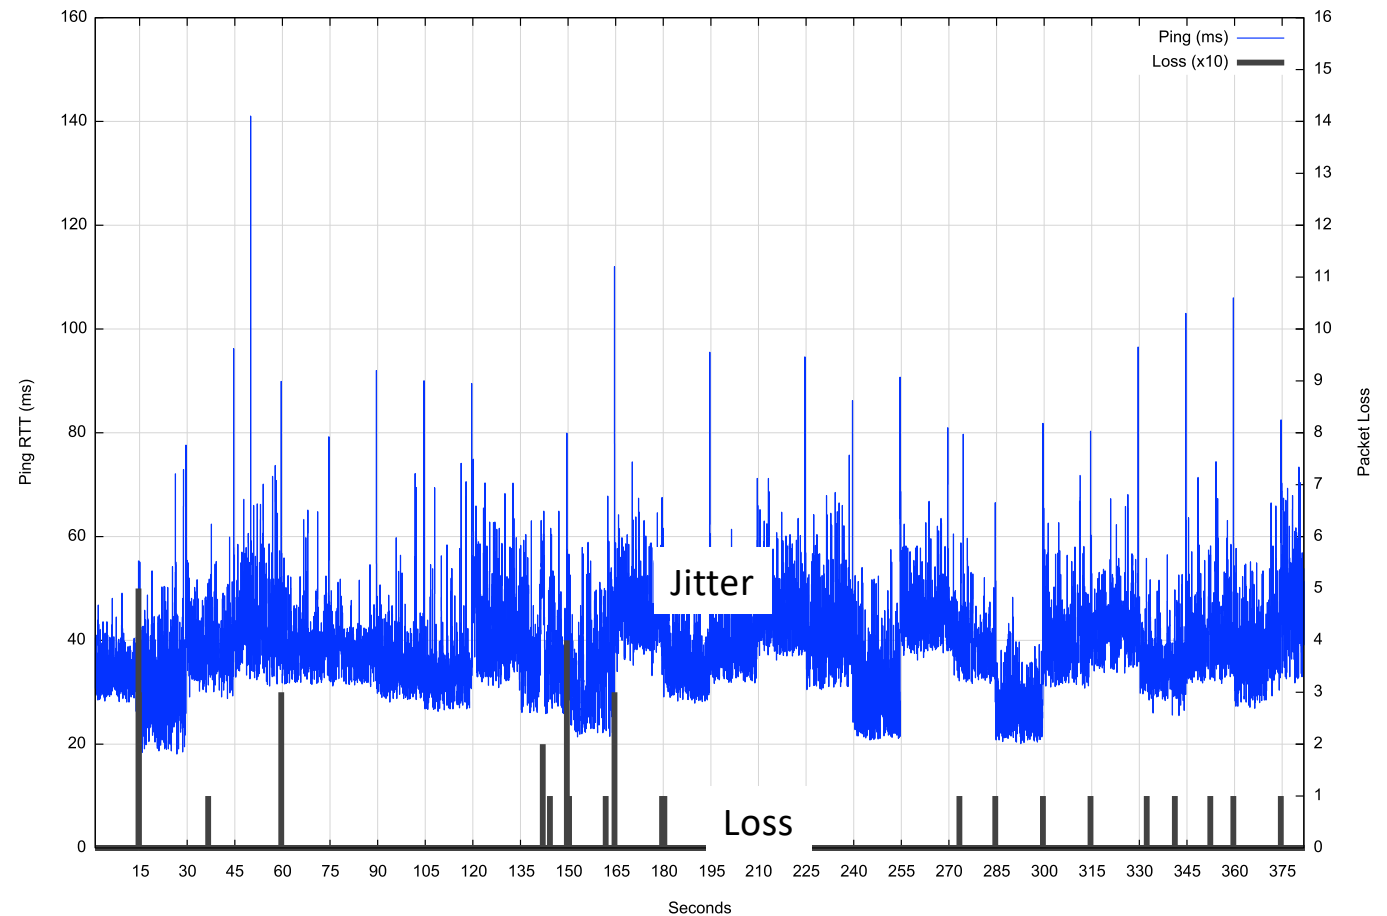


Why is the Starlink carrier so unstable?






# 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



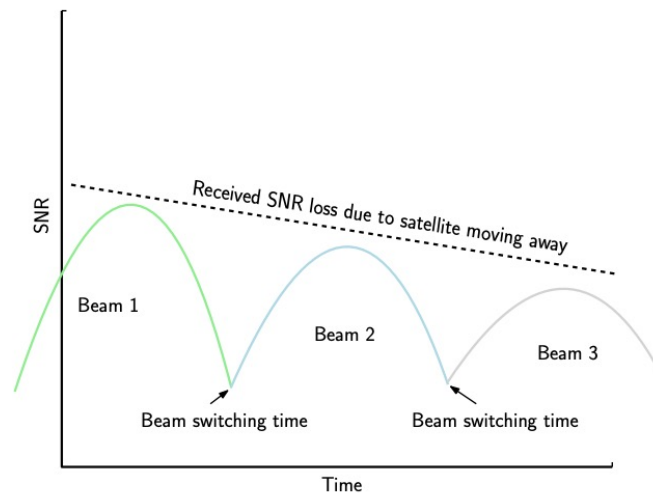


# Satellite handover



# Varying SNR

- 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



# 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 carrier 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



# TCP is the Internet

- The Transmission Control Protocol is an end-to-end protocol that creates a reliable stream protocol from the underlying IP datagram device
- This single protocol is the “beating heart” at the core of the Internet
- TCP operates as an adaptive rate control protocol that attempts to operate **efficiently** and **fairly**

# TCP Performance Objectives

To maintain an average flow which is both **Efficient** and **Fair**

## **Efficient:**


- Minimise packet loss
- Minimise packet re-ordering
- Do not leave unused path bandwidth on the table!

## **Fair:**

- Do not crowd out other TCP sessions
- Over time, take an average  $1/N$  of the path capacity when there are  $N$  other TCP sessions sharing the same path

# It's a Flow Control process

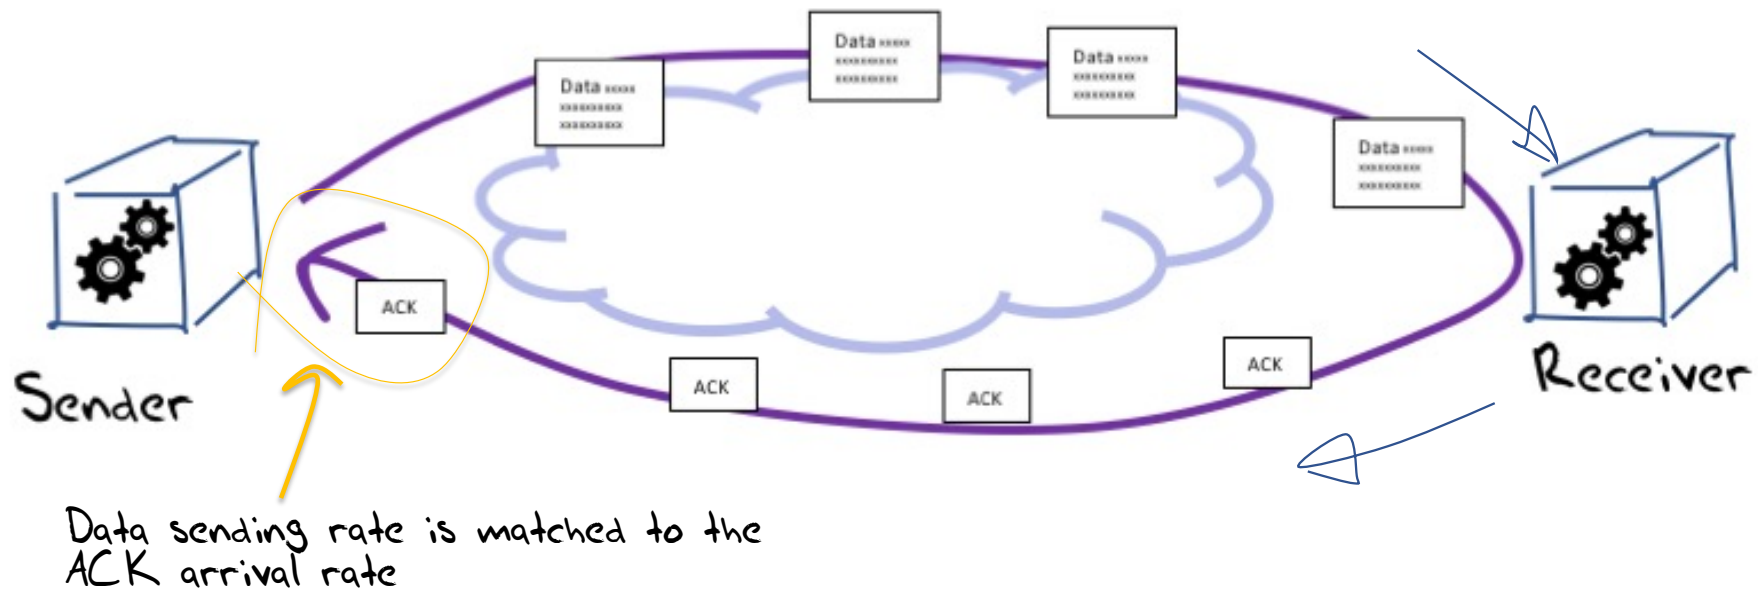
- Think of this as a multi-flow fluid dynamics problem
- Each flow has to gently exert pressure on the other flows to signal them to provide a fair share of the network, and be responsive to the pressure from all other flows



# TCP Control

TCP is an **ACK Pacing** protocol

If the sender sends one packet each time it receives an ACK, then the sender will maintain a steady number of packets in flight within the network



# TCP and Starlink

- TCP uses **ACK pacing** which means it attempts to optimize its sending rate over multiple RTT intervals
  - TCP assumes a stable carrier with low jitter and a stable channel capacity
  - When this is not the case TCP tends to reduce its sending rate to achieve stability
- 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




# How well does Starlink work?

Speedtest measurements:

We should be able to get  
~180Mbps out of a Starlink  
connection.


But Speedtest is NOT  
TCP – so lets look at TCP  
performance



# "Classic TCP" - TCP Reno


- Additive Increase Multiplicative Decrease (AIMD)
  - While there is no packet loss, increase the sending rate by one segment (MSS) each RTT interval
  - If there is packet loss (detected by duplicate ACKs) pause for 1xRTT and decrease the sending rate by 50% over the next RTT Interval by halving the sender's send window
- Start Up
  - Each RTT interval, double the sending rate
  - We call this "slow start" – probably because its anything but slow!!!

# The Classic TCP Picture




# CUBIC

- CUBIC is designed to be useful for high-speed sessions while still being 'fair' to other sessions and also efficient even at lower speeds
- Rather than probe in a linear manner for the sending rate that triggers packet loss, CUBIC uses a non-linear (cubic) search algorithm



# CUBIC and Queue formation



# CUBIC assessment

- Can react quickly to available capacity in the network
- Tends to sit for extended periods in the phase of queue formation
- Can react efficiently to long fat pipes and rapidly scale up the sending rate
- Operates in a manner that tends to exacerbate 'buffer bloat' conditions




# And there's a whole lot more...

| <i><b>TCP Variant</b></i>    | <i><b>Feedback</b></i> |   |
|------------------------------|------------------------|---|
| <i><b>RENO</b></i>           | Loss                   | AIMD                                    |
| <i><b>Vegas</b></i>          | Delay                  |   |
| <i><b>High Speed TCP</b></i> | Loss                   |   |
| <i><b>BIC</b></i>            | Loss                   | Binary Increase                         |
| <i><b>CUBIC</b></i>          | Loss                   | Cubic function increase - Linux-Adopted |
| <i><b>Agile-TCP</b></i>      | Loss                   | High Speed - Low Delay                  |
| <i><b>H-TCP</b></i>          | Loss                   | High Speed                              |
| <i><b>Fast</b></i>           | Delay                  | Akamai Proprietary                      |
| <i><b>Compound TCP</b></i>   | Loss/Delay             | Microsoft Adopted                       |
| <i><b>Westwood</b></i>       | Loss                   | Dynamic setting of Slow Start Threshold |
| <i><b>Elastic TCP</b></i>    | Loss/Delay             | High Speed - High Delay                 |


# Optimising Flow State

- There are three 'states' of flow management:
  - **Under-Utilised** – where the flow rate is below the link capacity and no queues form
  - **Over-Utilised** – where the flow rate is greater than the link capacity and queues form
  - **Saturated** – where the queue is filled and packet loss occurs
- Loss-based control systems probe upward to the Saturated point, and back off quickly to what they guess is the Under-Utilised state in order to let the queues drain
- But the optimal operational point for any flow is at the point of state change from Under to Over-utilised, not at the Saturated point
- We can detect this point by careful handling of delay – the onset of queuing causes additional delay in the observed round trip time

# RTT and Delivery Rate with Queuing



# TCP Flow Control Algorithms




“Ideal” Flow behaviour  
for each protocol


RENO and CUBIC are  
loss-based control  
algorithms where the  
upper limit in the send  
rate is established by  
encountering packet loss

BBR is a delay based  
control algorithm where  
the upper send limit is  
established by the onset  
of increased delay

# Starlink using iperf3 - cubic, 40 seconds





# Starlink using iperf3 - cubic, 40 seconds






# Peak vs off-Peak - CUBIC




Channel contention has a big impact on performance

# Starlink with iperf3 - bbr, 40- seconds



# Starlink with iperf3 - bbr, 40- seconds



# 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



---

## CONCLUSION

- N
  - 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?**

- 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?

# BBR Characteristics

- BBR is not sensitive to packet loss, so the regular packet loss events every 15 seconds do not impact BBR performance
- BBR uses **even pacing** of sent packets, so does not use network buffers to smooth out sender bursts
- BBR does not perform continuous delay monitoring, but instead “spikes” the sending rate every 8 RTT intervals by a massive 25%
  - BBR only checks for a change in delay during this spike interval
- This allows BBR to operate an internal model of channel capacity that is based on averaging across  $8 \times \text{RTT}$  intervals, reducing its sensitivity to jitter and high frequency capacity changes



# Protocol Considerations

- Starlink services have three issues:
  - Very high jitter rates – varying signal modulation
  - High levels of micro-loss (1.4%) – largely due to 15s satellite handover events
  - 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 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
- Its likely that ECN would also be really helpful to disambiguate latency changes due to satellite behaviours and network queue buildup

# 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
- 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 25° to 20°
- Use more bands (Ka-, V-, and E- bands)

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

Questions?

