

# **Exposing the Technical and Commercial Factors Underlying Internet Quality of Experience**

Don Bowman

NANOG 60

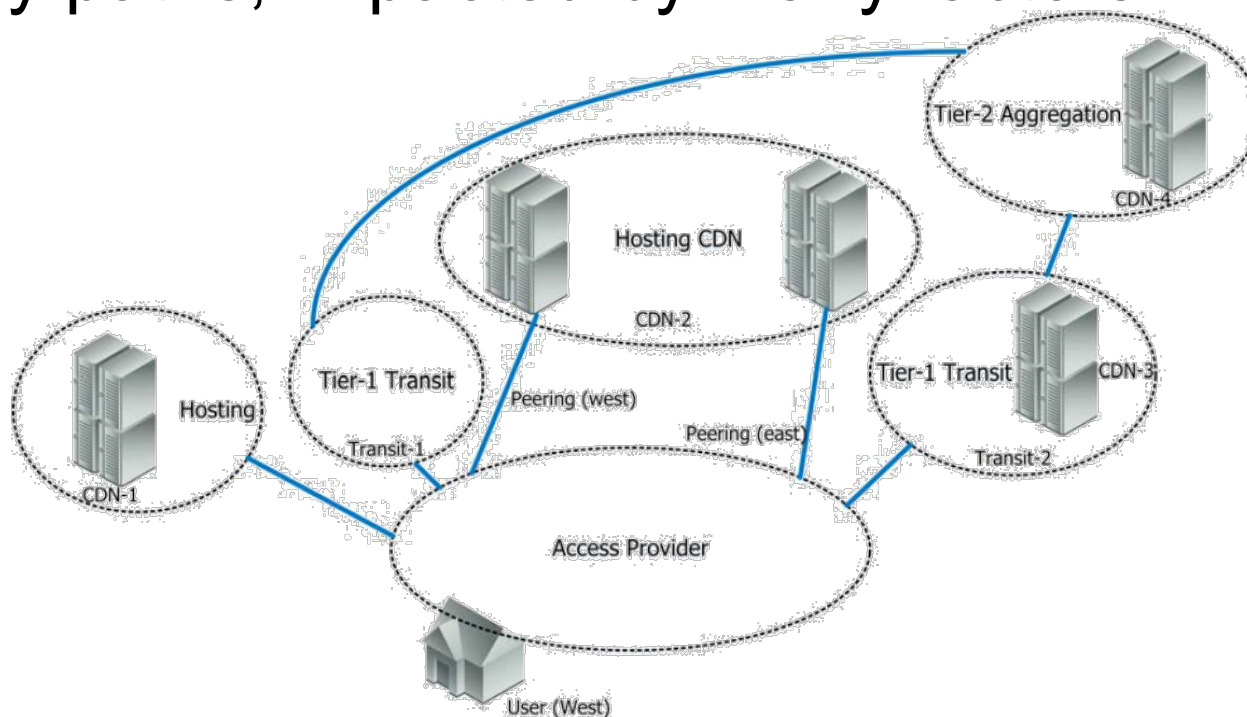
February 11, 2014

# Research Goals

- 6 US participants (5 MSO, 1 LEC)
- Depict how and where does data flows, who are the types of players
  - Who has what incentive
- Show that Quality is an end to end concept
  - Some actions by one player can be corrected at expense by downstream player (e.g. routing)
  - Some cannot (e.g. origin encoding, device limit)
- Demonstrate that Capacity & Demand are different
  - Existing benchmarks are poor


# Streaming Video as a Proxy for Quality

- Streaming video is sensitive and prevalent, so is a common proxy for quality
- In this network, video can reach the user by many paths, impacted by many factors

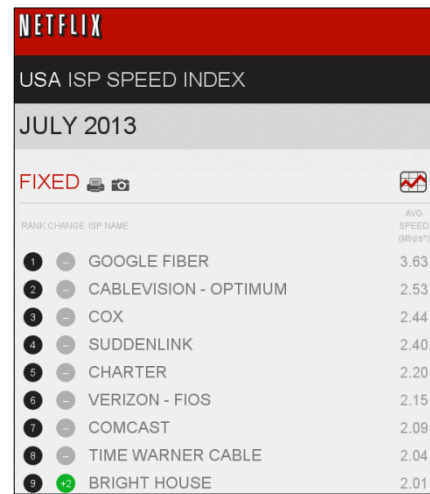


# Popular Quality Benchmarks

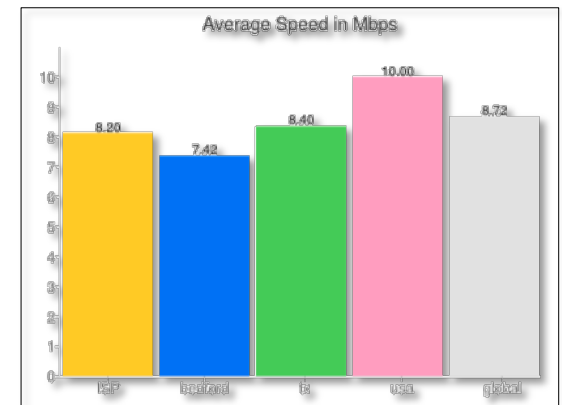
- Several services have gained popularity as credible sources of quality metrics
  - But are they accurate? We took a look.



Ookla's  
Speedtest.net  
attempts to  
measure capacity



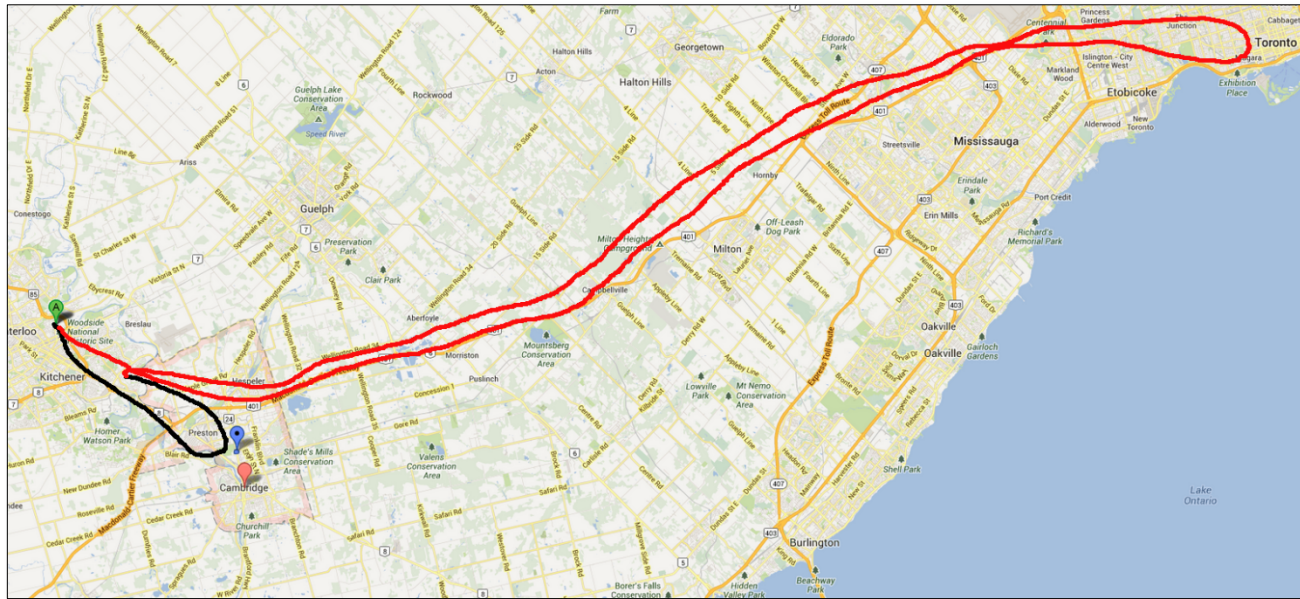
Netflix measures  
delivered  
bandwidth as a  
proxy for demand



YouTube attempts  
to approximate  
both capacity and  
demand

# Quality Benchmarks: Speedtest

- In practice, the results reported by Speedtest showed enormous variation, dependent upon the server used for the test
  - In this image, both servers are in the same building (in Kitchener, Ontario), but have different routes






# Quality Benchmarks: Speedtest




- In practice, the results reported by Speedtest showed enormous variation, dependent upon the server used for the test
- Consistently variable in every country we tested: Singapore, Hong Kong, South Africa, Australia, Brazil, Canada, United States

# Quality Benchmarks: Speedtest

- Comcast's 105 Mbps service






- AT&T's U-Verse





# Quality Benchmarks: Speedtest

- Singapore, famed for high-speed fibre, has a 2:1 bandwidth difference



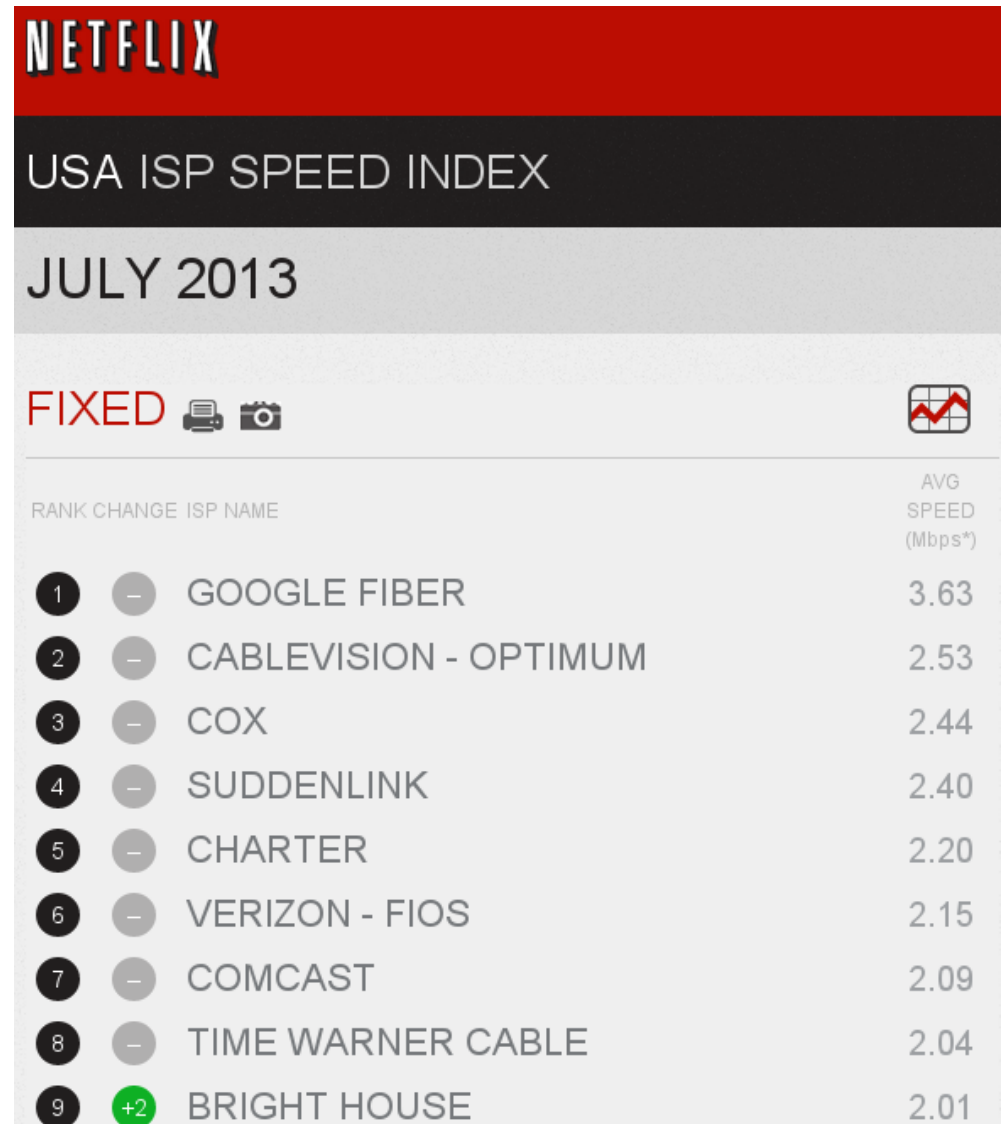


# Quality Benchmarks: Speedtest

- In practice, the results reported by Speedtest showed enormous variation, dependent upon the server used for the test
- Consistently variable in every country we tested: Singapore, Hong Kong, South Africa, Australia, Brazil, Canada, United States
- Speedtest is not an accurate measurement of quality, as it is far too dependent upon server location and characteristics

# Quality Benchmarks: Netflix

- With every update of the Netflix ISP Speed Index, network operators either rejoice or scratch their heads




# Quality Benchmarks: Netflix

- When we looked a little deeper at Netflix, we observed a few things
  - Each ISP experienced a peak in OpenConnect bandwidth every day in the early morning
  - Each ISP showed variation in volume per CDN and the time of demand
  - Some observed quality dips on a CDN at some time, but none observed dips in all

# Quality Benchmarks: Netflix


- Example: Netflix bitrate by CDN over a day



# Quality Benchmarks: Netflix


- We also used the server latency as a proxy for server location, and found that traditional CDNs showed little time-of-day variation, but OpenConnect was strongly correlated to UTC

**Latency, Internet side, Netflix by CDN**



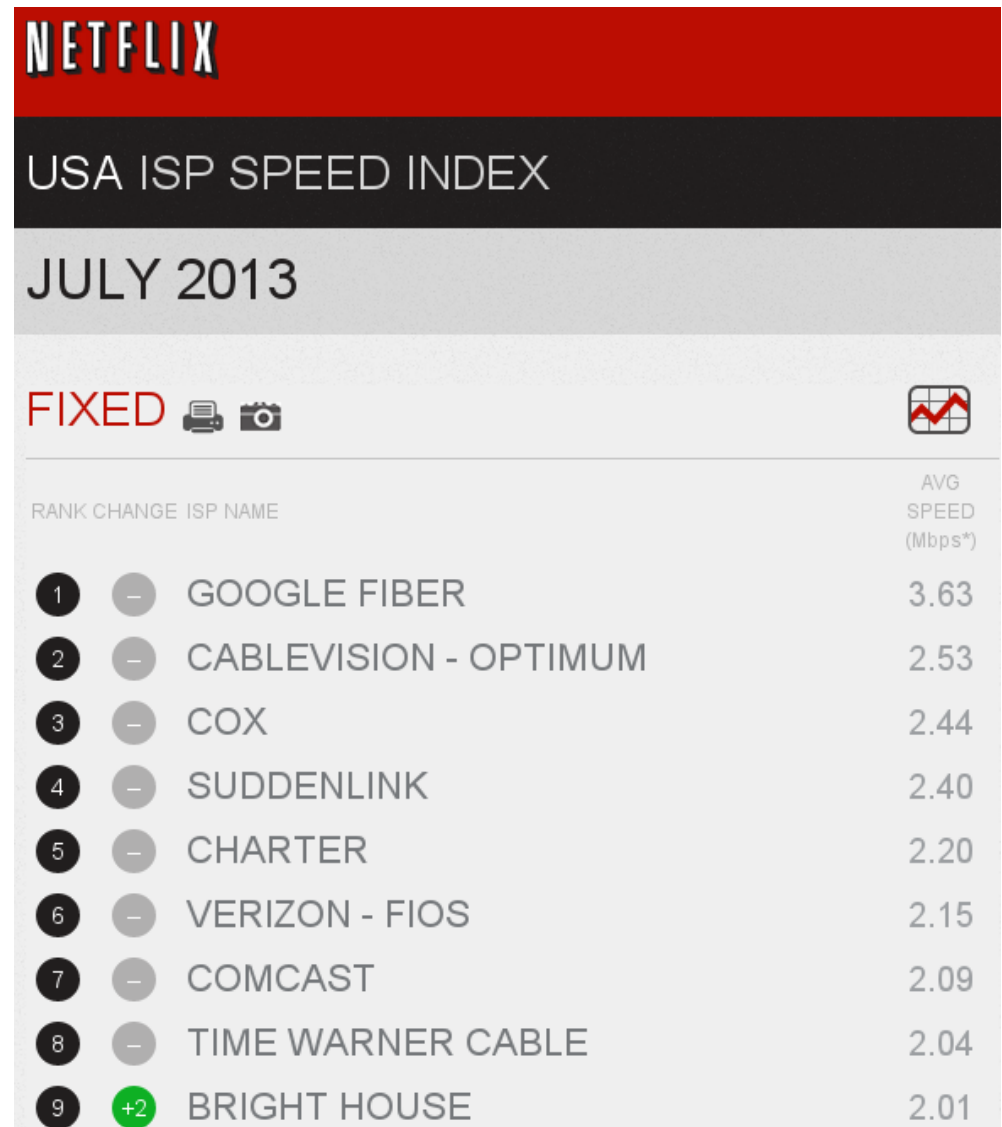
# Quality Benchmarks: Netflix

- In our observations, Netflix was the only video provider to have the latency scale with load, and this was the case only on OpenConnect



# Quality Benchmarks: Netflix


- We can conclude that the Netflix ISP Speed Index is a flawed measure
- Too dependent upon OpenConnect locations and characteristics






# Quality Benchmarks: YouTube

- YouTube measures the average delivered speed and reports against relevant comparators
  - In this image, “ISP” is Time Warner Cable’s 50 Mbps service, measured in Dallas




# Quality Benchmarks: YouTube

- But this isn't a true measure of connection speed to YouTube
- To deliver video, YouTube bursts on for ~2 seconds, then switches off for ~2 seconds




# Quality Benchmarks: YouTube

- The modem in this case can sustain a 40 Mbps connection, but the average is <20
- And the average video bitrate is between 6 Mbps and 8 Mbps





# Quality Benchmarks: YouTube

- We observed that YouTube experienced a dip in delivered bandwidth around 12pm and 9pm



# Quality Benchmarks: YouTube

- We observed that YouTube experienced a dip in delivered bandwidth around 12pm and 9pm
- Comparison to other video providers shows that this issue is isolated to YouTube



# Quality Benchmarks: YouTube

- YouTube measures the average delivered speed and reports against relevant comparators
- We find that YouTube's measurement is also flawed
  - Doesn't measure actual connection speed to YouTube
  - YouTube's servers seem to experience congestion, even when the network has excess capacity


# Our Own: Latency of Top 100 Domains

- The top 100 web domains combine to form an illustrative proxy for both ‘user experience’ and ‘congestion’
  - We measure the round-trip time
  - Top 100 is determined by observation and measurement
  - Provides a consistent method of comparing worldwide performance
  - Can be used to differentiate between access network and transit/peer networks



# Our Own: Latency of Top 100 Domains

- The top 100 web domains combine to form an illustrative proxy for both 'user experience' and 'congestion'




This graph shows little congestion on the access side, since the round-trip time relatively constant.

# Our Own: Latency of Top 100 Domains

- The top 100 web domains combine to form an illustrative proxy for both ‘user experience’ and ‘congestion’

This graph shows congestion on the transit/peer side, illustrated by a rise in round-trip time during the evening hours.



# Summary

- Traffic flow is impacted by many independent decisions
  - Some technical, many commercial
- Common quality benchmarks (e.g., Ookla's Speedtest.net and Netflix ISP Index) are misleading
  - None are very accurate, but all are widely believed
- The end user's quality of experience (QoE) is fundamentally dependent upon both technical and commercial factors