Measuring Network Quality for End-Users, 2021

An Internet Architecture Board virtual workshop

Day 2

Agenda

Metrics 2

14:00 Chairs' Intro

14:10 Jonathan Foulkes. Metrics helpful in assessing Internet Quality

14:17 Vijay Sivaraman, Sharat Madanapalli, Himal Kumar. Measuring Network Experience Meaningfully, Accurately, and Scalably

14:24 Dave Reed, Levi Perigo. Measuring ISP Performance in Broadband America: a Study of Latency Under Load

14:31 Discussion

Metrics 3

15:00 Kyle MacMillian, Nick Feamster. Beyond Speed Test: Measuring Latency Under Load Across Different Speed Tiers

15:07 Gregory Mirsky, Xiao Min, Gyan Mishra, Liuyan Han. Error Performance Measurement in Packet-Switched Networks

15:14 Gino Dion. Focusing on latency, not throughput, to provide better internet experience and network quality

15:21 Praveen Balasubramanian. Transport Layer Statistics for Network Quality

15:28 Discussion

16:00 Break

Cross-Layer 1

16:10 Jari Arkko, Mirja Kuehlewind. Observability is needed to improve network quality

16:17 Robin Marx, Joris Herbots. Merge Those Metrics: Towards Holistic (Protocol) Logging

16:24 Rajat Ghai. Measuring & Improving QoE on the Xfinity Wi-Fi Network

16:31 Discussion

Cross-Layer 2

17:00 Koen De Schepper, Olivier Tilmans, Gino Dion. Challenges and opportunities of hardware support for Low Queuing Latency without Packet Loss

17:07 Ken Kerpez, Jinous Shafiei, John Cioffi, Pete Chow, Djamel Bousaber. State of Wi-Fi Reporting

17:14 Mikhail Liubogoshchev. Cross-layer Cooperation for Better Network Service

17:21 Francois Michel, Olivier Bonaventure. Packet delivery time as a tie-breaker for assessing Wi-Fi access points

17:28 Discussion

18:00 End of Day 2

Metrics helpful in assessing Internet Quality

Jonathan Foulkes

Helpful Internet Quality Metrics

- Line capacity and latency under working loads
  - Capacity variability - 36% of lines vary by more than 10% over time
  - Accounting for ALL capacity usage on the link is important
  - Reporting capacity usage and latency statistics over time
- Dynamic variability

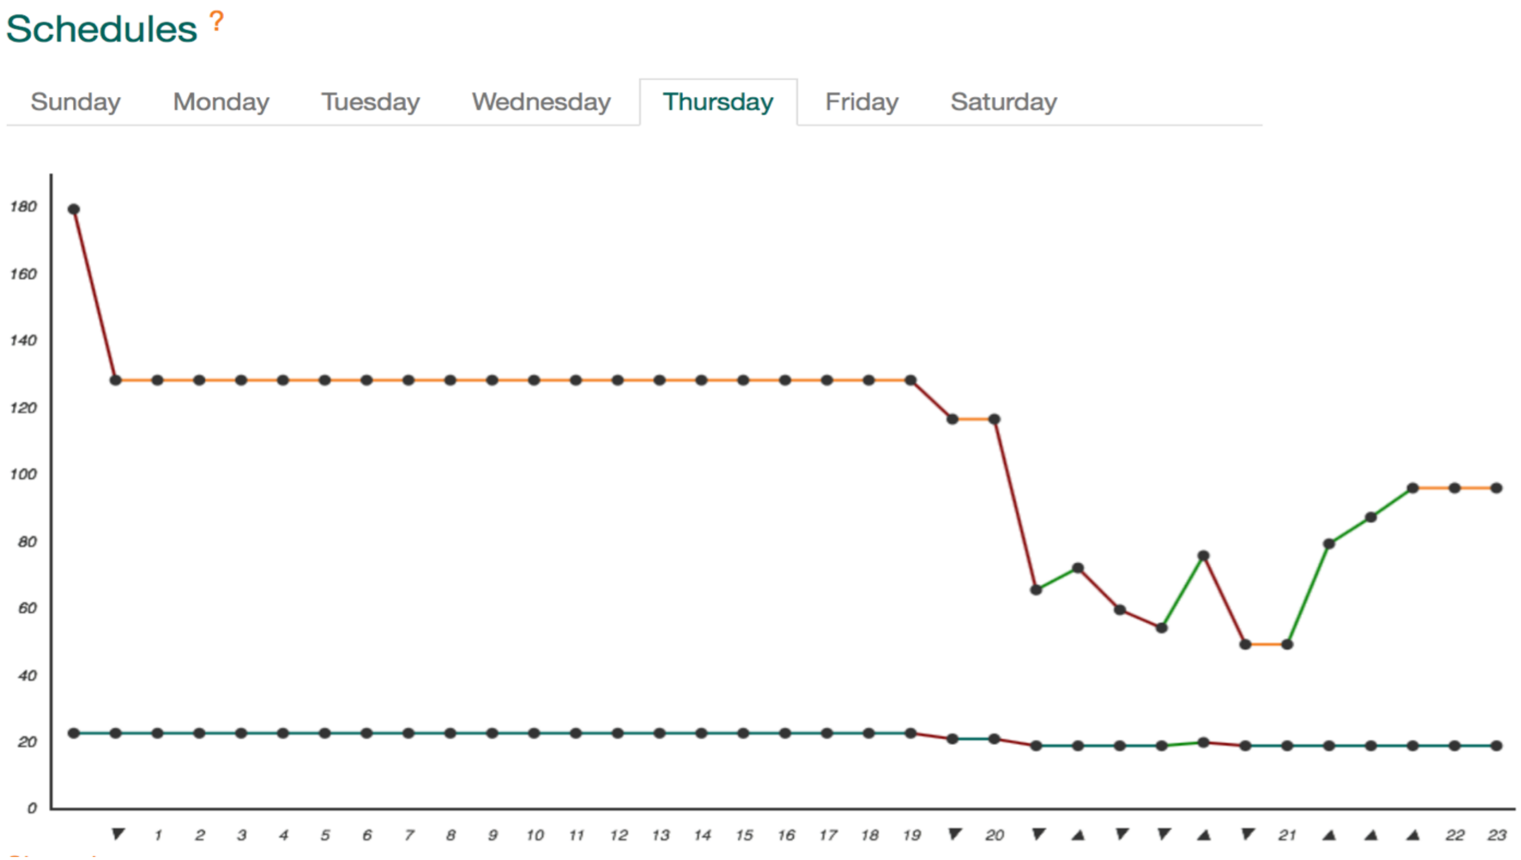

Line stability metrics

- Link loss
  - Date time & duration
  - Segment local vs link losses
- Line instability events
  - 27% have this
  - 16% have >50 events

Unstable line definition:

An increase in latencies with little to no load

Typical causes:
- Link errors due to modem or local loop
- Backhaul congestion

Missing – Local metrics

- Quantify local link performance
  - Capacity and latencies under load
- Contrast to external link metrics
  - Understand where loss of quality happens
- Local network service
  - Router hosted or App on wired system

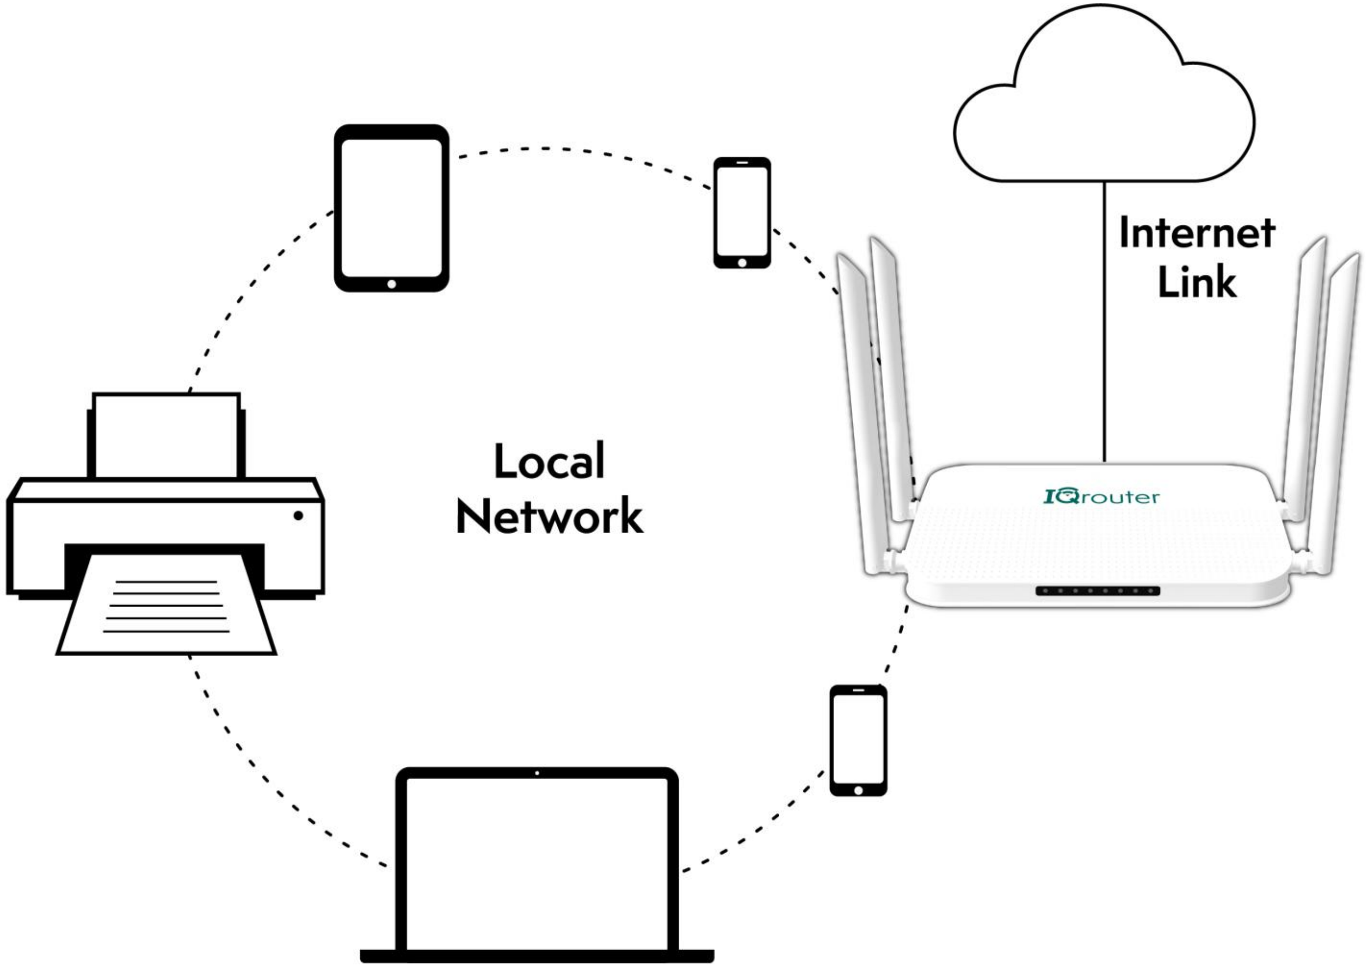

Measuring Network Experience Meaningfully, Accurately, and Scalably

Vijay Sivaraman

Measuring Network Experience Meaningfully, Accurately, and Scalably

Vijay Sivaraman, Sharat Madanapalli, Himal Kumar

UNSW SYDNEY

CANOPUS NETWORKS

15 Sep 2021

Meaningfully: What and Why?

- Speed has diminishing returns for consumers
  - … and the gap across ISPs is narrowing
  - … while racing ISP economics to the bottom
- Latency, loss, etc. may not be directly perceived
  - Applications can absorb these
- Let’s measure what the consumer perceives:
  - Streaming video (on-demand or live):
    - Is video at best resolution? Is it freezing?
  - Gaming:
    - Is latency low and consistent (i.e. low jitter)?
  - Conferencing:
    - Are there stutters and dropouts?

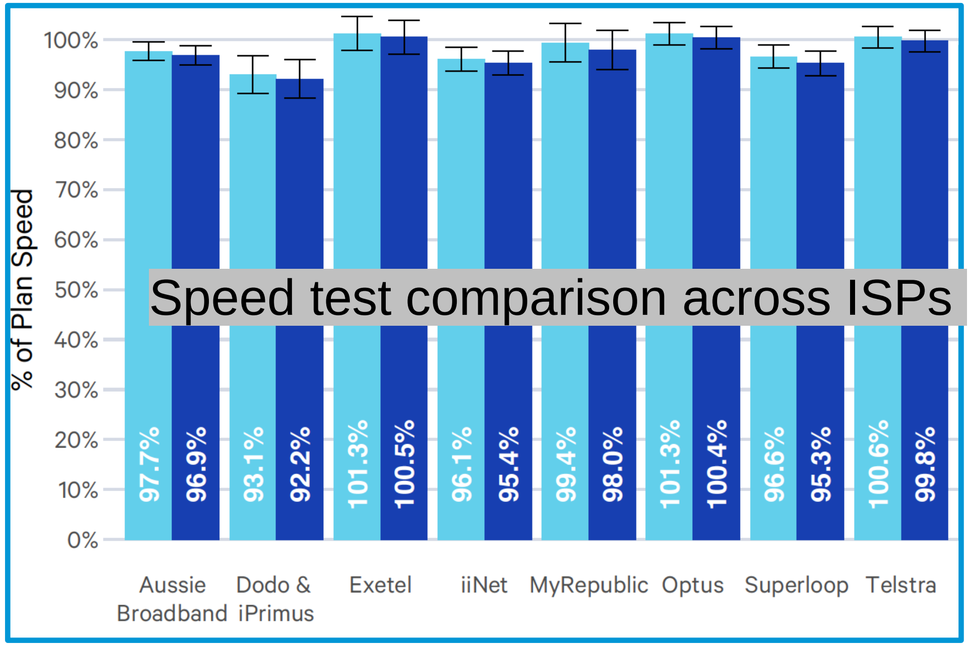

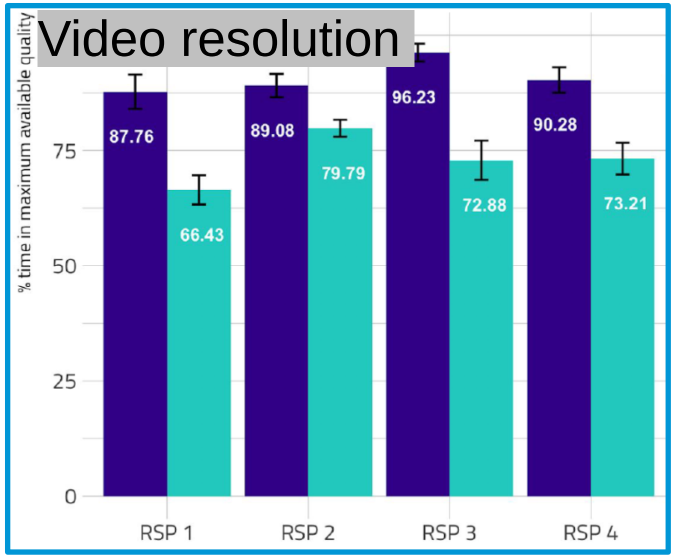

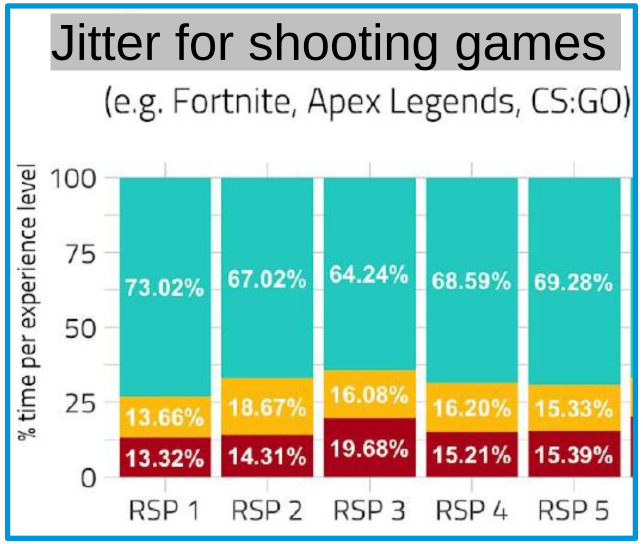

UNSW SYDNEY

CANOPUS NETWORKS

Accurately and Scalably: How?

- Terabit speed Programmable Network chip
  - Scalable to multi-Tbps
- Push telemetry to “pulse” each flow
  - Accurate to sub-100msec (tunable)
- AI behavioral models for application experience
  - Video: chunk fetch patterns for resolution & buffer state
  - Games: latency (TCP), jitter spikes (UDP)
  - Methods are engineered to application dynamics
  - … and agnostic to data (& header) encryption

- Commercial trials underway with Telcos

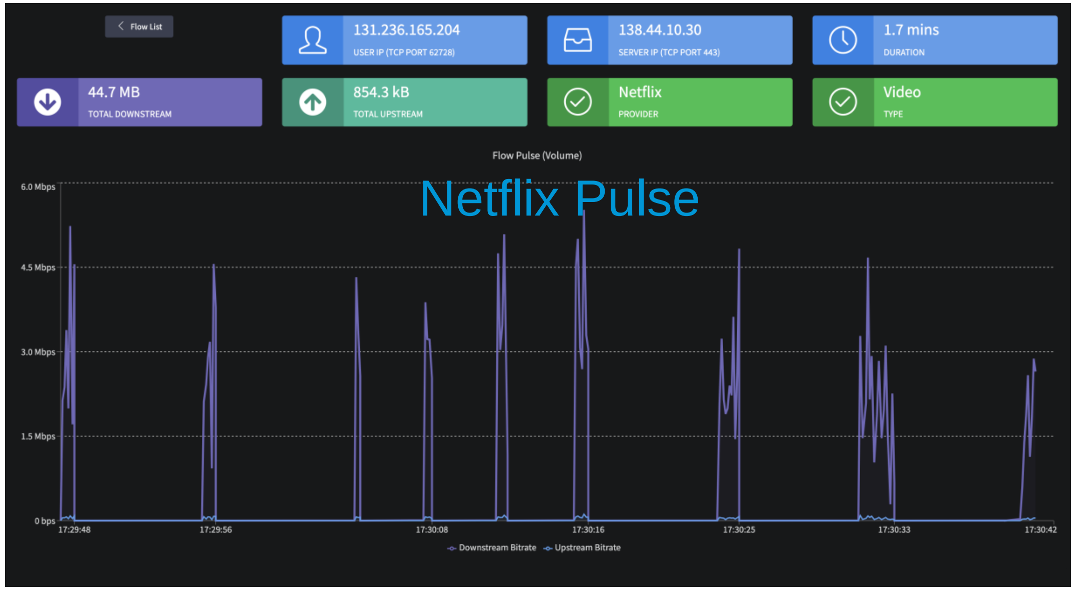

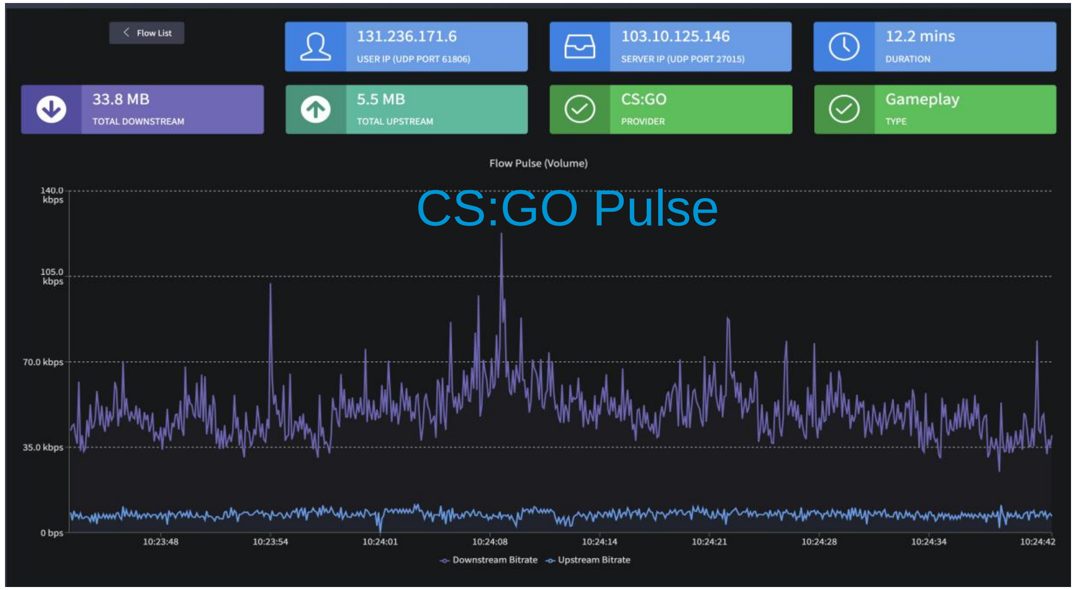

Measuring ISP Performance in Broadband America: a Study of Latency Under Load

David Reed

MEASURING ISP PERFORMANCE IN BROADBAND AMERICA

A Study of Latency Under Load

Dr. David P. Reed & Dr. Levi Perigo

Computer Science Department

University of Colorado Boulder

FCC's Downstream Latency Under Load Data By Technology (Average RTT, 2020)

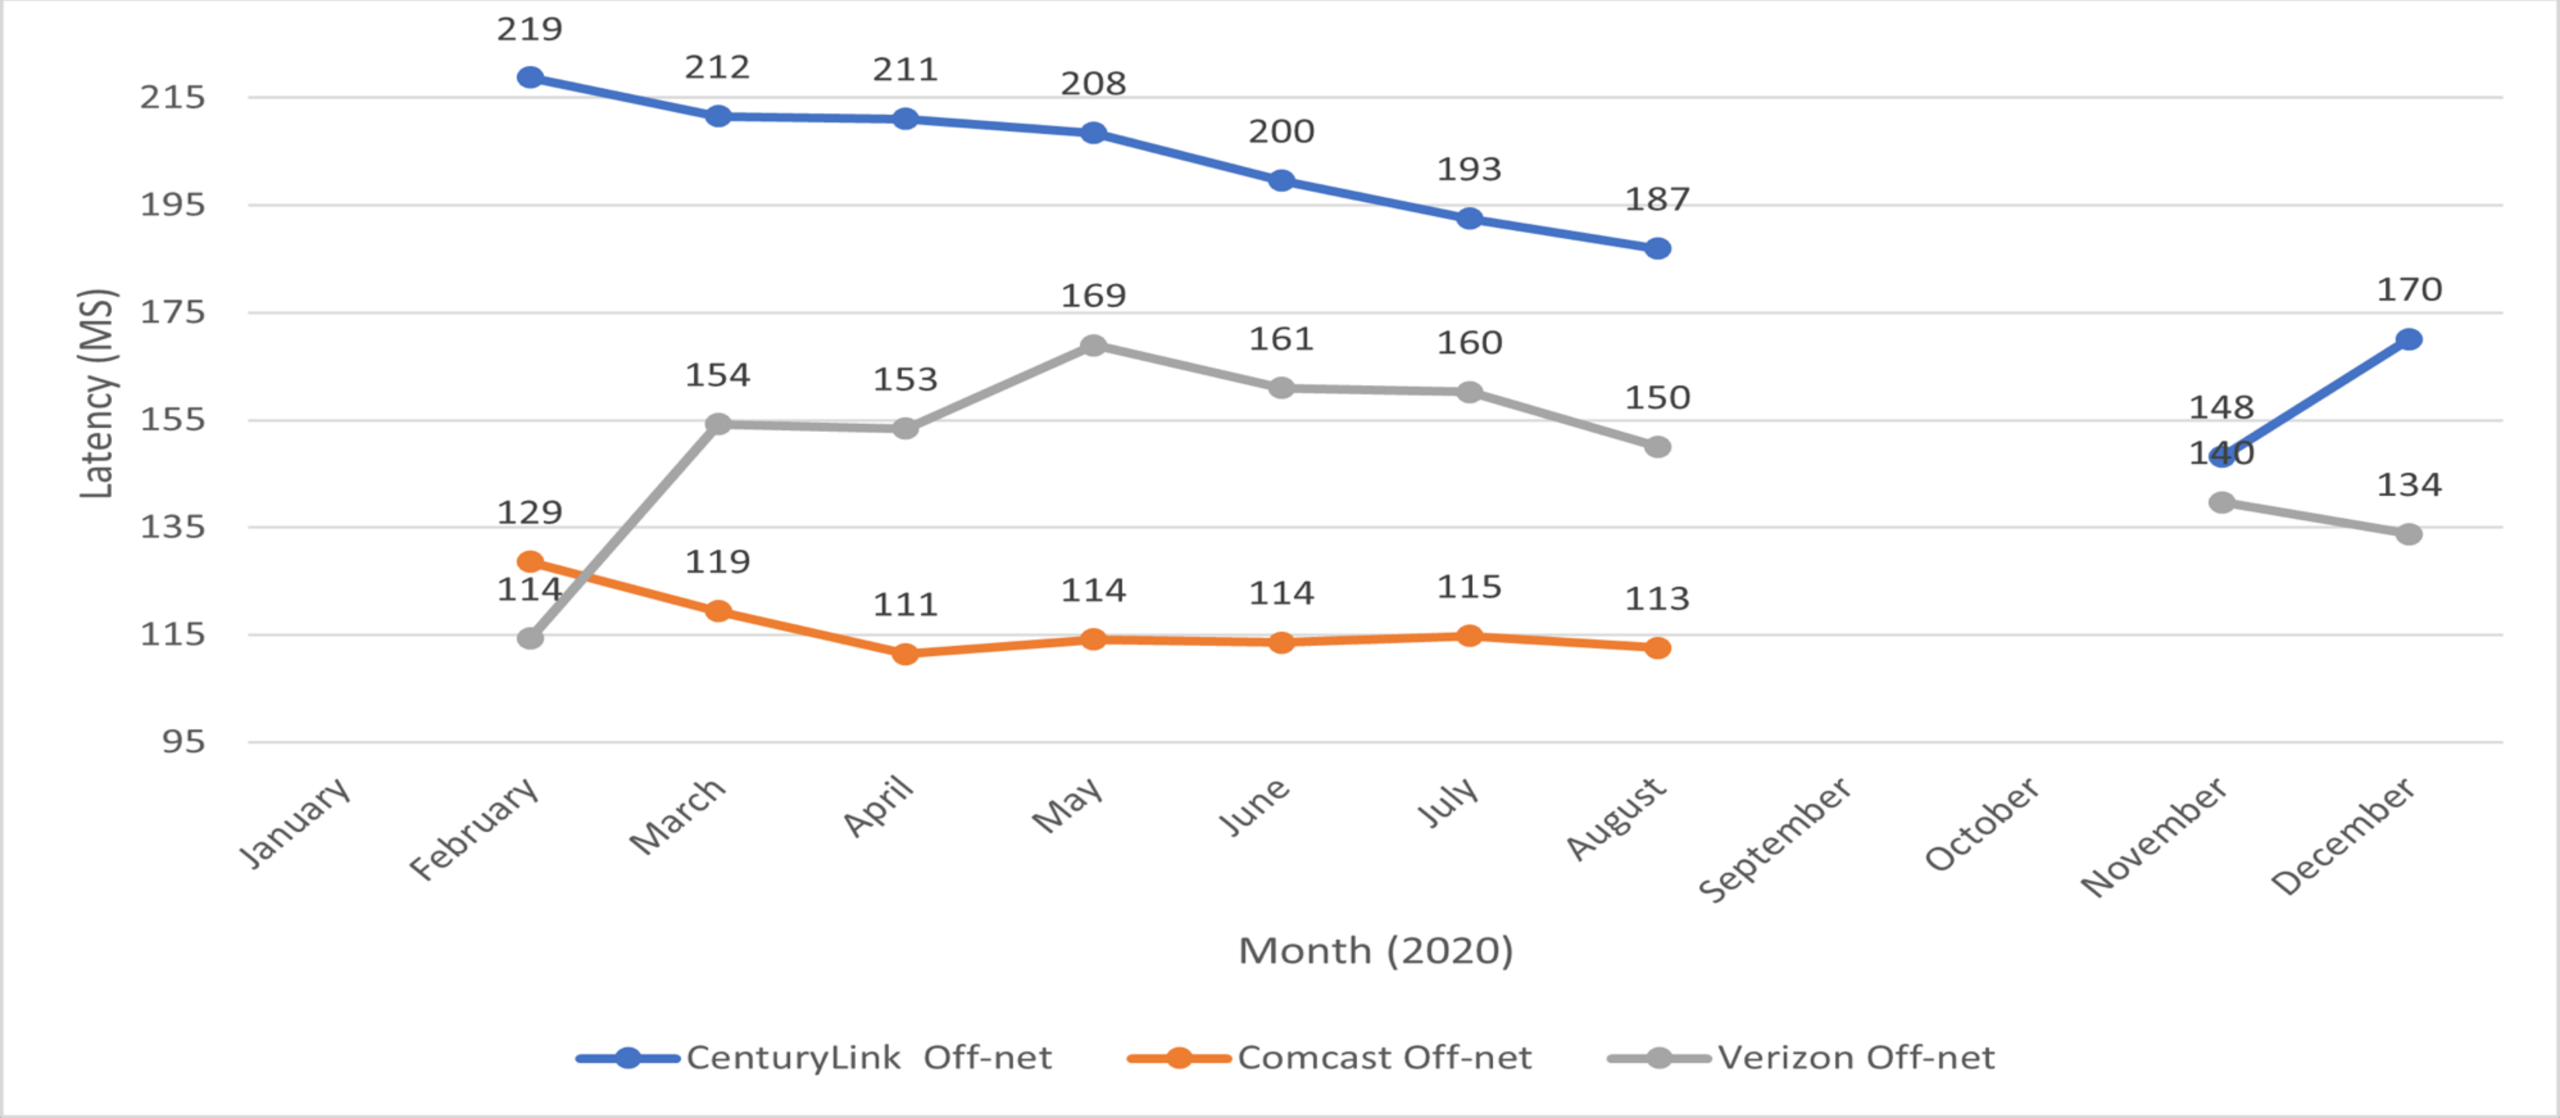

FCC's Downstream Latency Under Load Data By Technology (Maximum RTT, 2020)

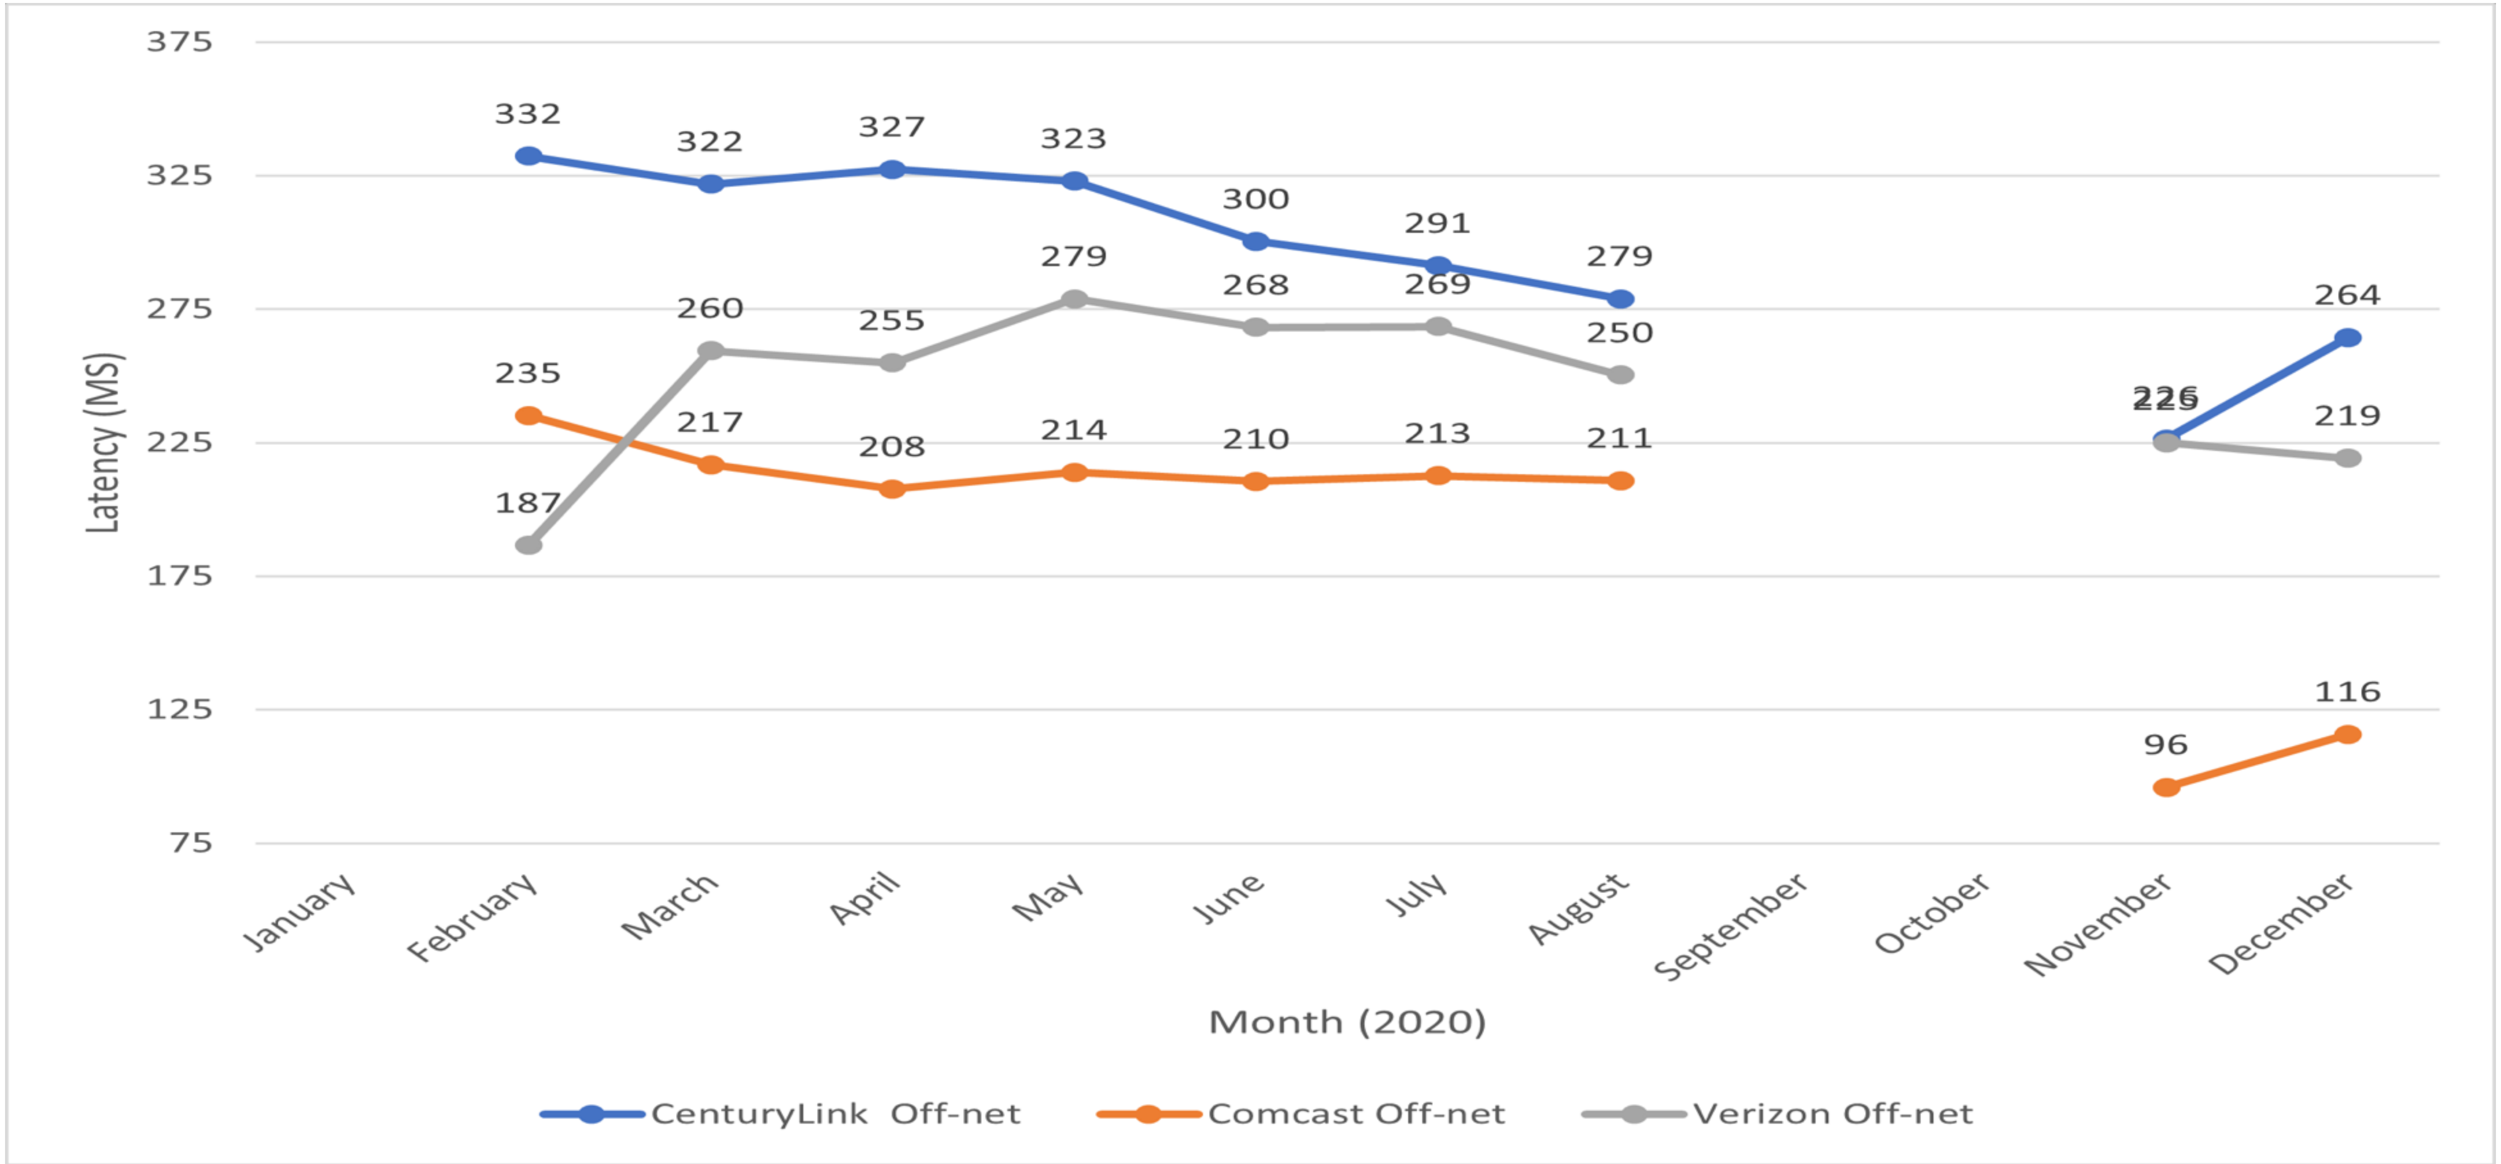

FCC's Historical Latency Under Load Data Over Cable (Longitudinal Average RTT, Comcast, 2020)

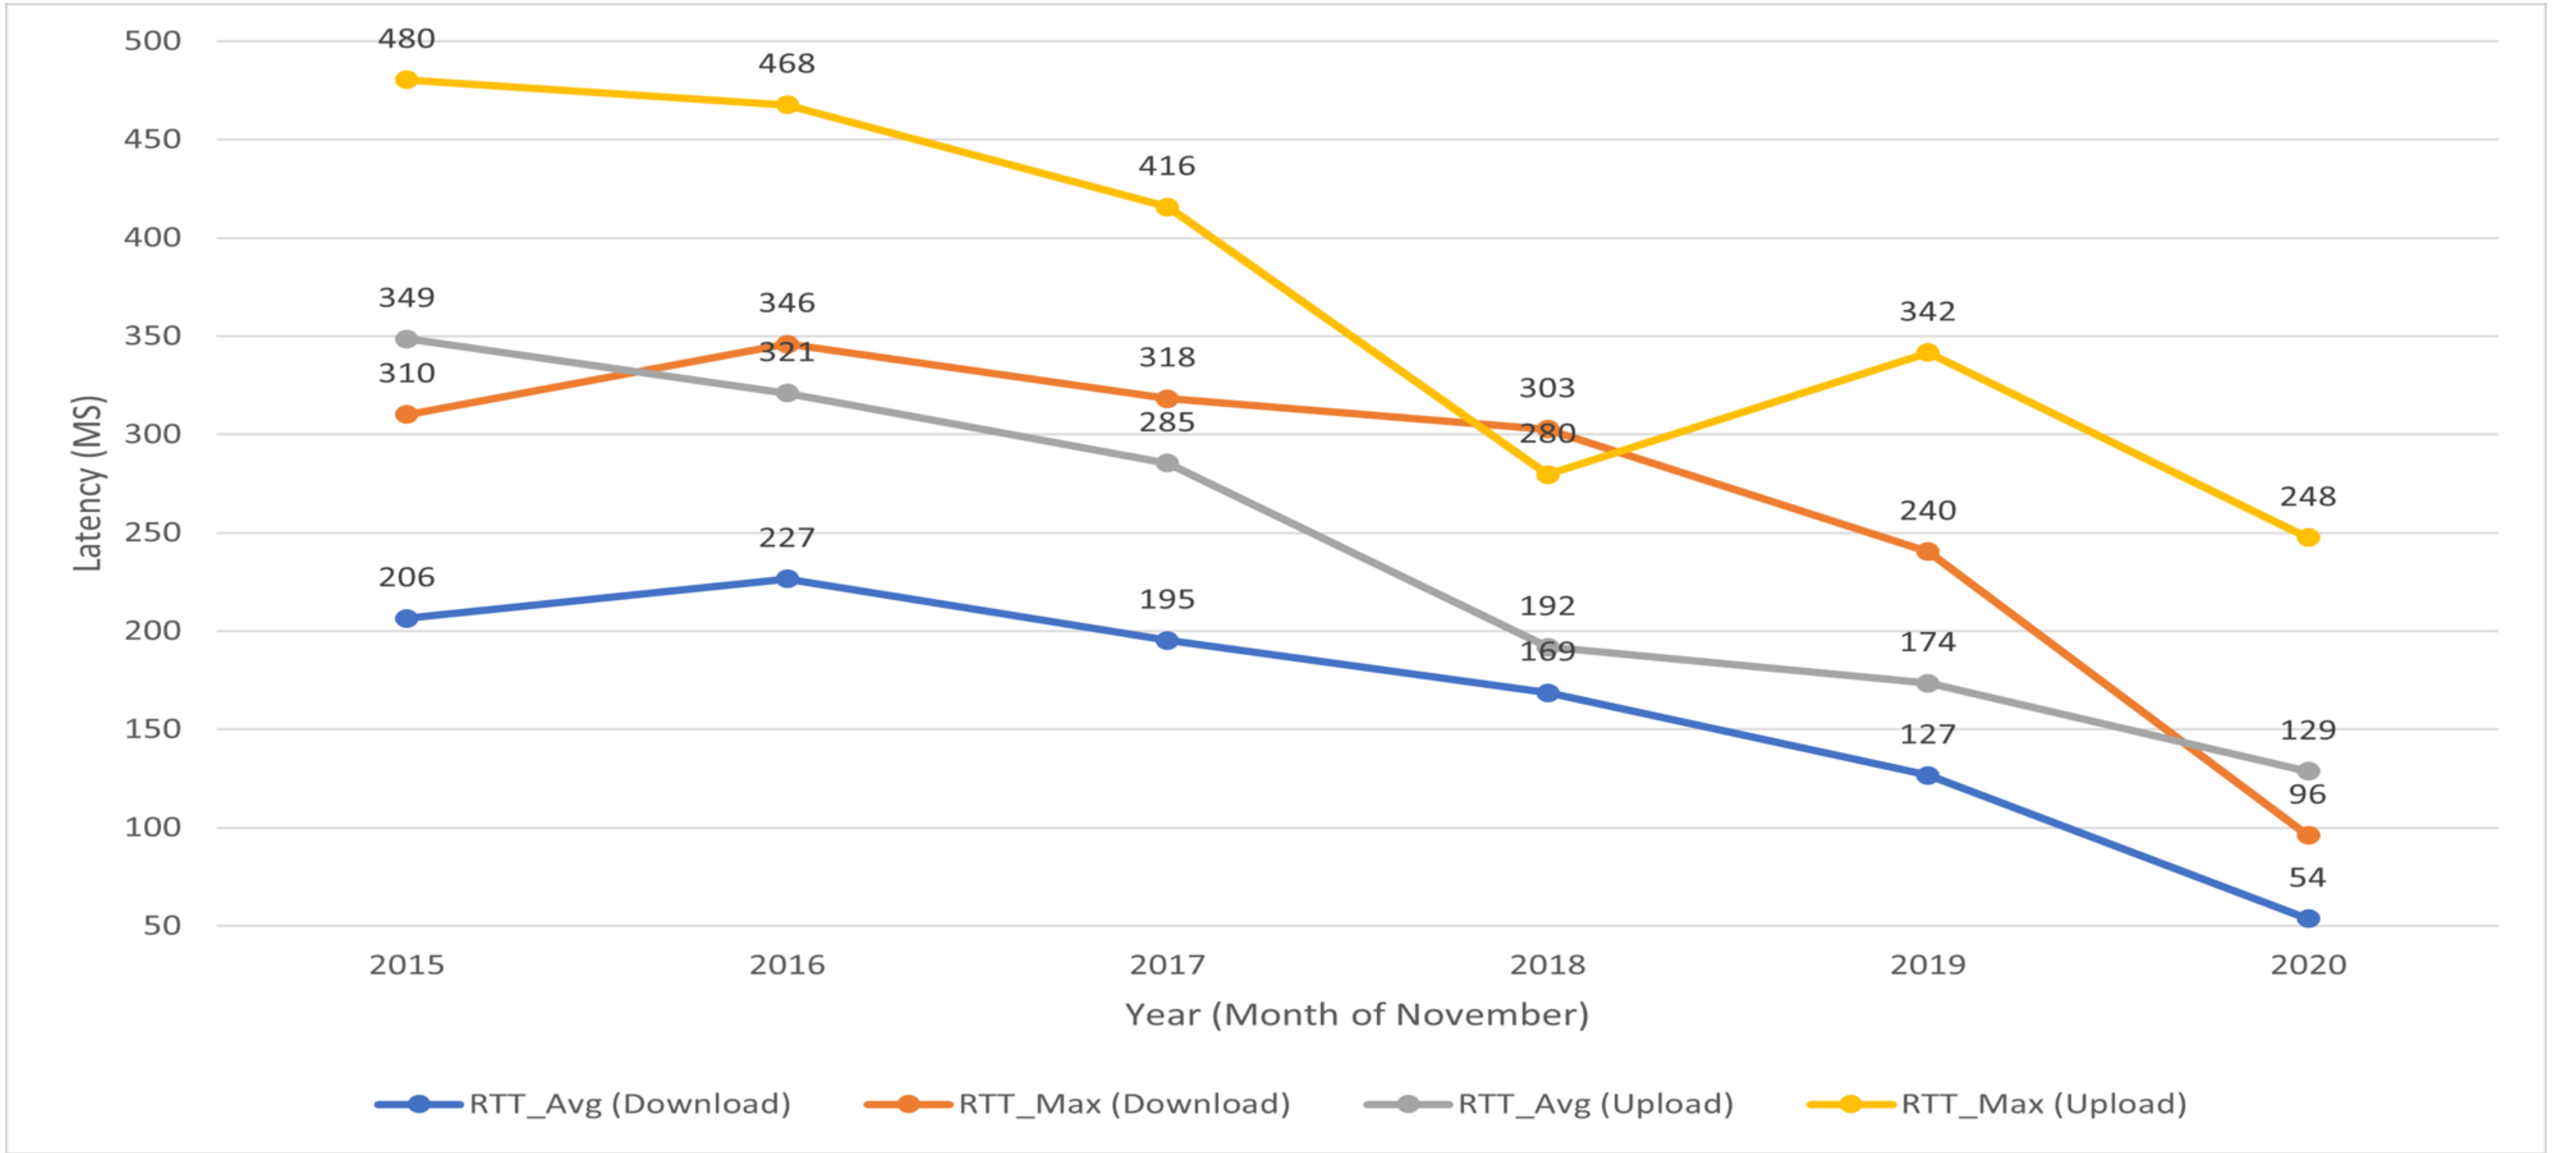

FCC's Historical Latency Under Load Data Over Cable (Comcast, 2020)

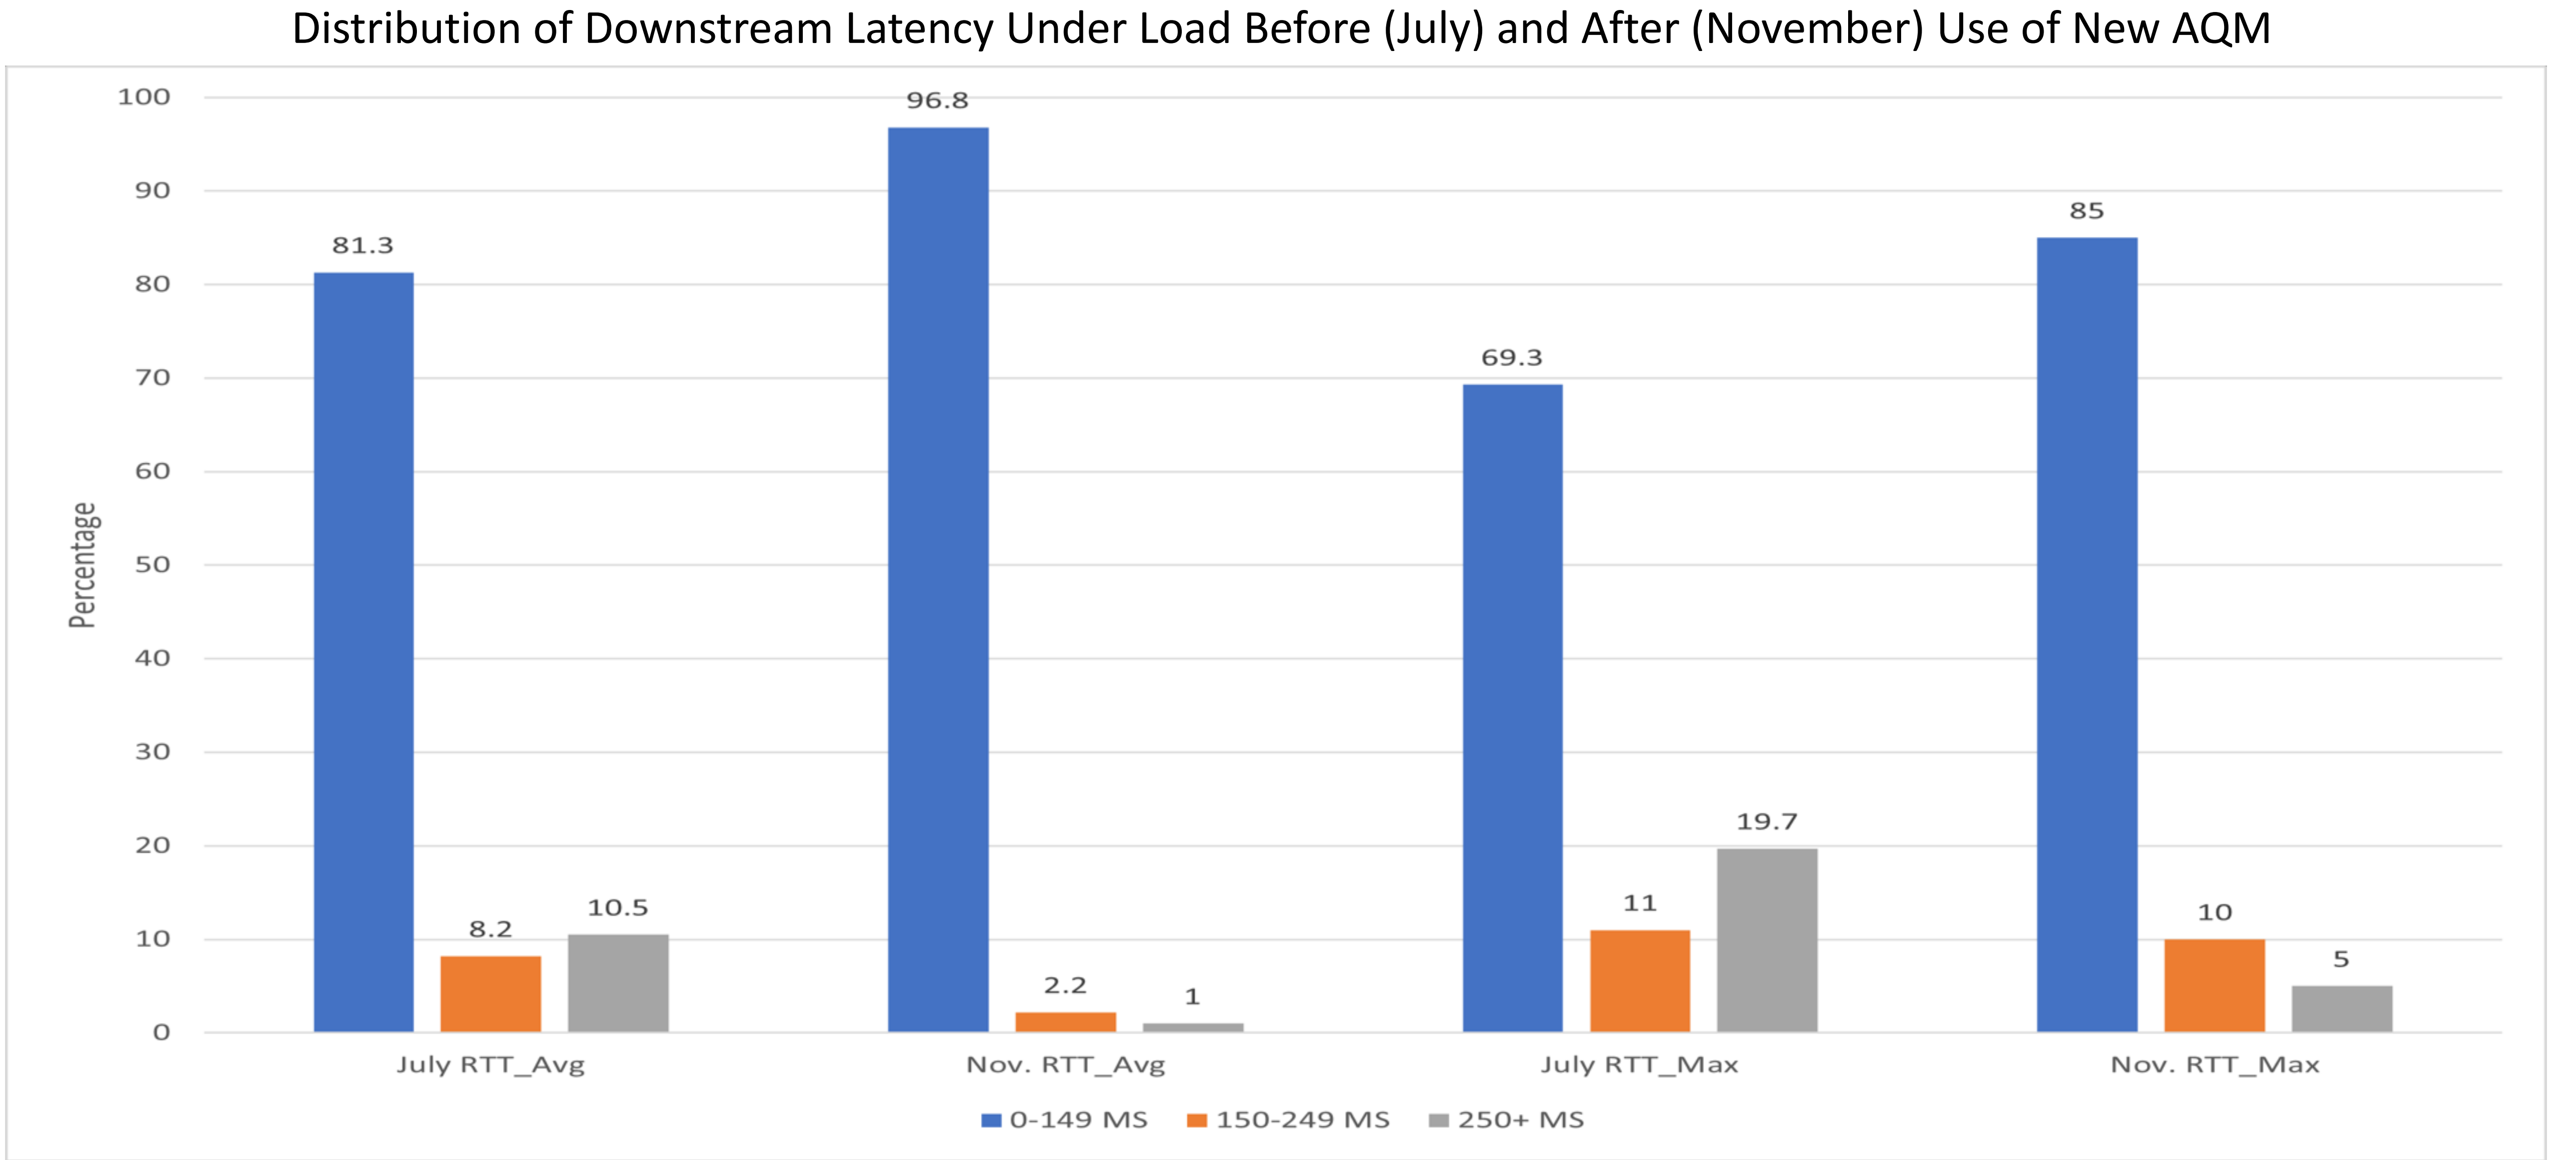

Findings

- LUL data shows network performance changes under increased traffic loads – *useful additional information for future FCC MBA Reports*
- Substantial difference downstream and upstream LUL for all technologies – due to *lower upstream speed, multiple access to shared upstream bandwidth*
- Substantial variation in downstream and upstream LUL by technology type – *due to varying service speeds, cable/FTTH latency much lower than DSL*
- Max RTT shows upper bound on latency experienced by users – *enough for users to experience notable degradation of real-time applications*
- LUL performance by cable ISP (Comcast) improved average of ~10% per year – *due to increasing service speeds and IP protocol/device improvements*
    - New AQM immediately improved average LUL performance by 48% in 2020
- LUL data shows relationship between latency and service speed – *forecasts improved latency performance as trend continues to higher speeds*

Discussion

To be enqueued, please write ‘+q’ in the chat

To cancel being enqueued, please write ‘-q’ in the chat

The duration of each comment is limited by 60 seconds

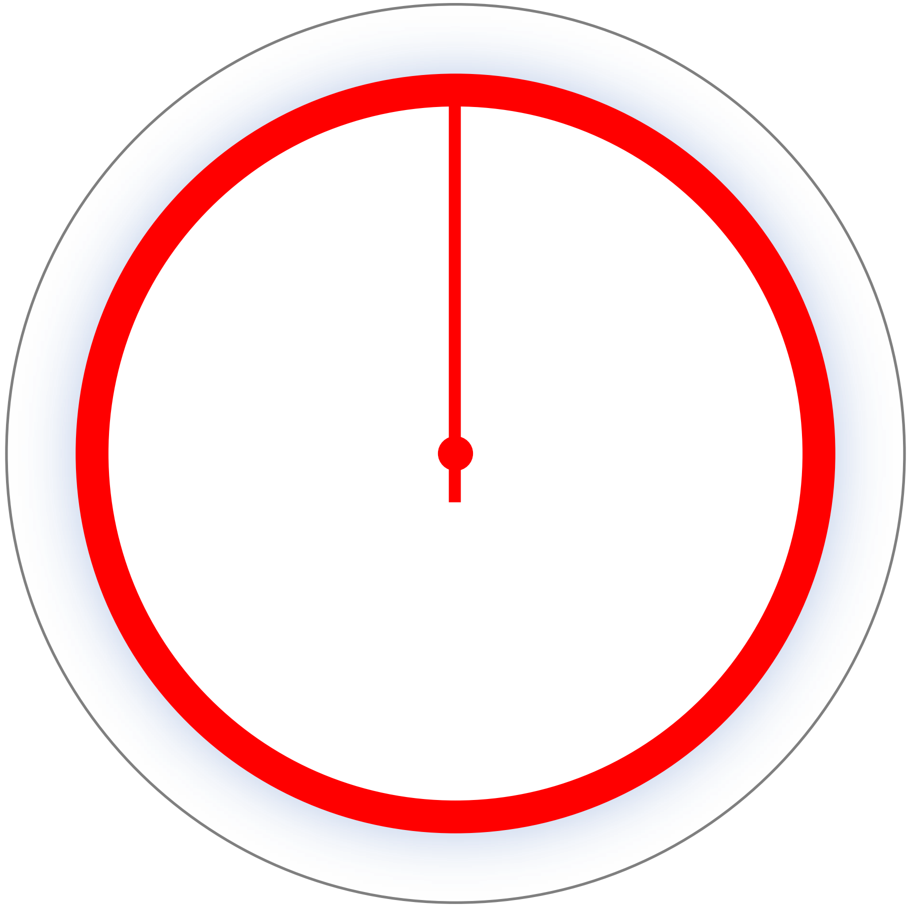

Beyond Speed Test: Measuring Latency Under Load Across Different Speed Tiers

Kyle MacMillan

Measuring Latency Under Load Across Different Speed Tiers

- Preliminary results from latency under load tests conducted across tens of homes across Chicago.
- Latency under load is given by the RTT of an ICMP packet to a given destination when either uplink or downlink is saturated.
- Link is saturated using iPerf3 TCP speed test.
- Tests are conducted every 2 hours.
- Data collected from at least 30 days of tests.

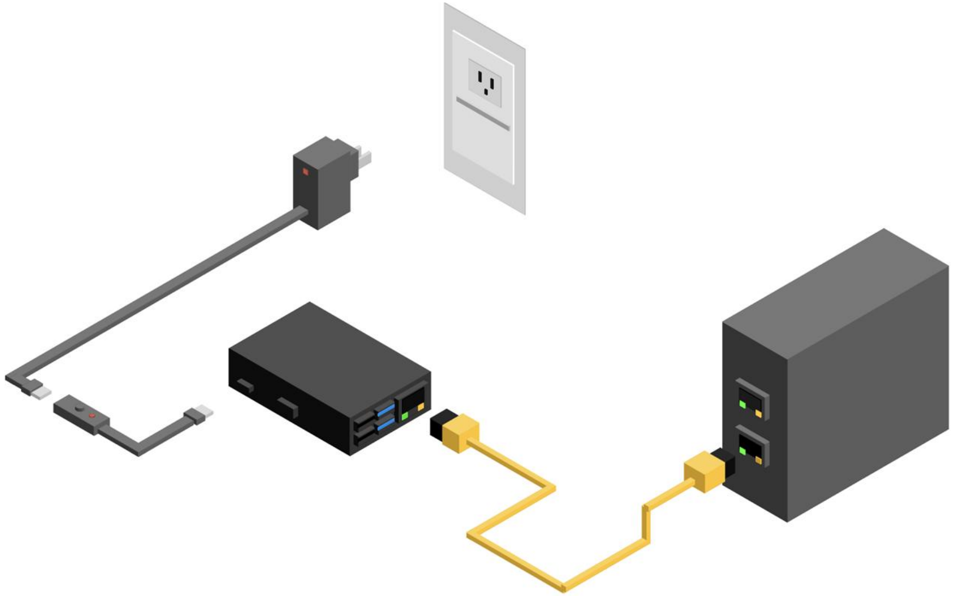

Figure 1: Measurement collection setup. A raspberry pi on a wired connection runs our measurement suite in each participant's home.

Kyle MacMillan, Nick Feamster

University of Chicago

Observation #1: Latency Under Load Related to Speed Tier

- Users subscribed to internet service plans with lower throughput generally experience higher latency under load than those on higher speed plans.

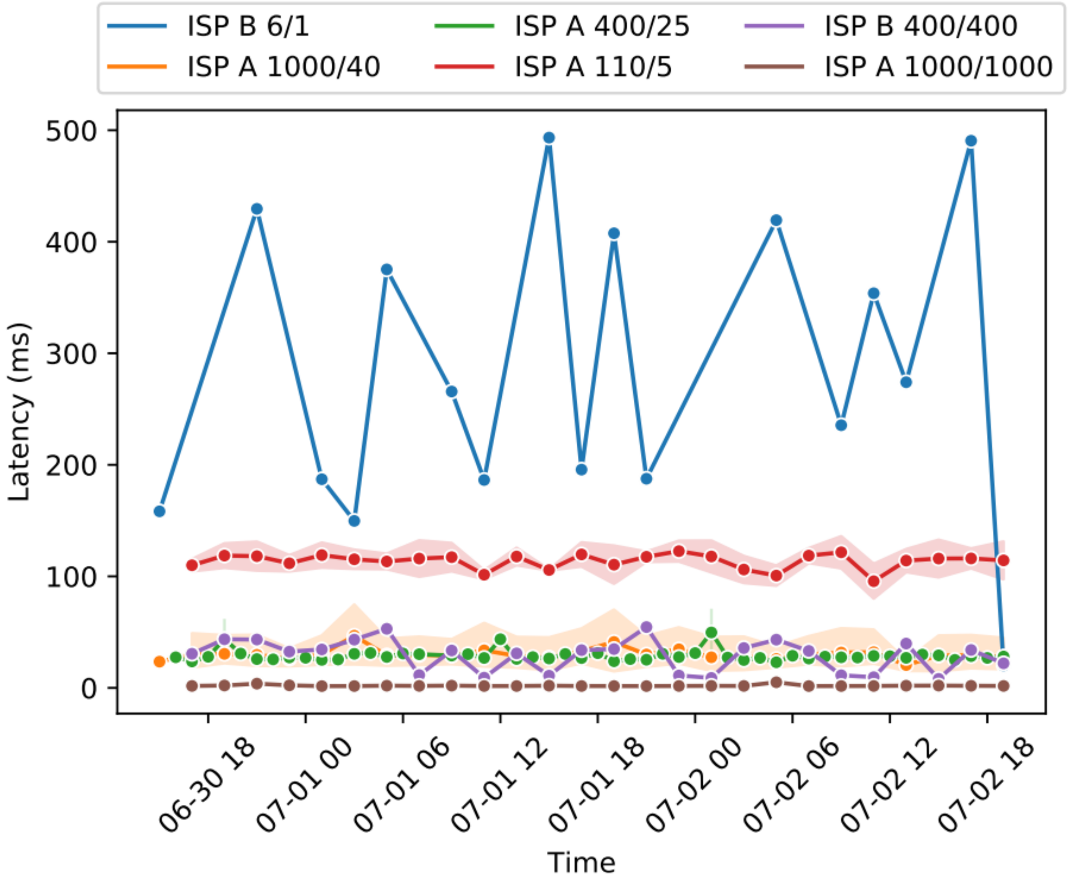

Figure 2: Mean latency under upstream load over time, aggregated by subscriber plan and ISP, with 95% confidence intervals.

Observation #2:

Same Service Plan + Same ISP ≠ Same Latency Under Load

- Users subscribed to the same service plan on the same ISP experience different latency under load.
- While the differences are clear, the causes are challenging to uncover.

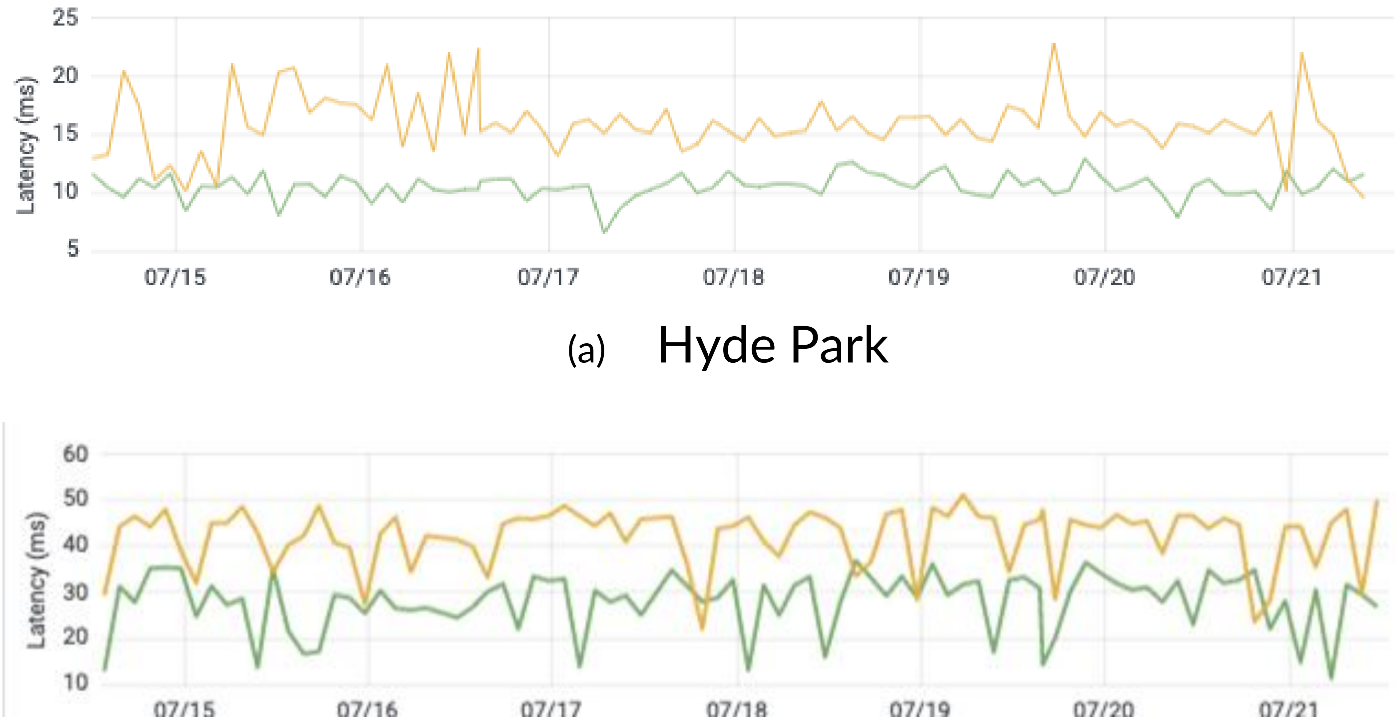

(a) Hyde Park

(b) South Shore

Figure 3: Latency Under Load for two subscribers to ISP A's 1 Gbps service plan, in two different Chicago neighborhoods.

Error Performance Measurement in Packet-Switched Networks

Gregory Mirsky

Error Performance Measurement in Packet-switched Networks

Greg Mirsky (Ericsson)
Xiao Min (ZTE)
Gyan Mishra (Verizon)
Liuyan Han (China Mobile)

What is Error Performance?

- OAM toolset includes methods to detect defects and measure performance
- Defect is an inability to communicate. Defect in PSN is Loss of path continuity, i.e., there’s no path through the network to get a packet from the source node to the destination node
- Defect state is the state of 100% packet loss – bridge Fault Management and Performance Monitoring OAM
- Packet Loss is an infinite delay of a packet
- Error Performance – quantitative characterization of the network condition between endpoints

EPM is Active OAM

- EPM is well-known in constant bit-rate, e.g., TDM, communication technologies (ITU-T G.826 and G.827)
  - based on the guaranteed presence of data, several EPM states and metrics defined, including state of path availability and unavailability
- A packet-switched network is based on the principle of statistical multiplexing and does not provide a predictable, guaranteed rate of receiving packets in the specified flow
- Without predictable flow, the operational state of a PSN cannot be characterized with certainty. If the state to be determined using only data traffic, how to differentiate pause in receiving data packets caused by the nature of the application from caused by the network failure?
- Only active OAM can create a sub-flow with a predictable rate of packets that EPM OAM can use

EPM Apparatus

- Consider using G.826/G.827 EPM parameters:
  - Errored Interval (second)
  - Severely Errored Interval (second)
  - Error-free Interval (second)
- Consecutive intervals form a period of:
  - Availability
  - Unavailability
- To make it stable, the definition of a period includes hysteresis. For example:
  - Ten consecutive Severely Errored intervals determine that a path is in an unavailable period that started at the beginning of the first Severe Errored interval.
  - A sequence of Errored and Error-free intervals shorter than ten does not change the state of the path, i.e., it is still in unavailable period.
- Other metrics:
  - Errored Interval Ratio = Errored Intervals / Total Number of Intervals
  - Severely Errored Interval Ratio = Severely Errored Intervals / Total Number of Intervals

Focusing on latency, not throughput, to provide better internet experience and network quality

Gino Dion

Low Latency vs Consistent Latency

Cloud Gaming Latency Budget Example

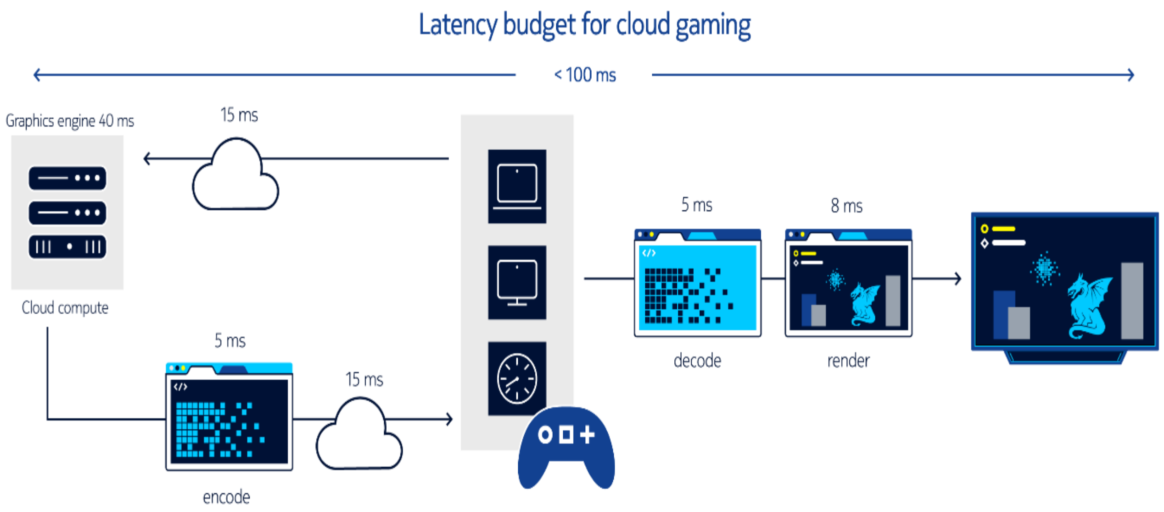

"How many of the attendees are competitive gamers? Operate your own Twitch stream? If not, can you really relate to application outcomes and benefits?"

"You can't easily fix low latency consistency with just better peering and edge compute, or better gaming "netcode", or more efficient video codecs; you need an end-or to-end latency strategy."

"Instead of application(s) as unique entities, think of use cases and services. CSP don't sell better application packages, they sell better class of services"

	Median Latency	Jitter Potential (99th percentile)
In-Home Ethernet	1-2 ms	<1 ms
In-Home Wi-Fi (sub 100 Mb/s traffic load, without AQM)	~10-20 ms	~10-4000 ms
xDSL	~20-50 ms	~5 ms
DOCSIS (3.0 under load, with buffer control)	~10 ms	~10-200 ms
LLD DOCSIS (estimated)	~1 ms	~1-5 ms
Access - FTTH	<1 ms	~1-2 ms
4G/5G (RAN latency, non-URLLC based profile, QCI 9)	~15-35 ms	~10-500 ms
IP Transport & Peering	~40 ms (near-edge) ~225 ms (intercontinental)	~5-10 ms ~10-50 ms
Applications & platform	~15-40 ms	~10-500 ms

Cloud Gaming Real-World Results – PI2 AQM and L4S

FTTH in a dense residential area, over WIFI, 80Mb/s of sustained Internet traffic

Real-world metrics need to be based on real-world benefits

- First graph is a residential gateway (integrated WIFI) on FTTH, without any advanced form of AQM.
- The second graph is a CPE with PI2 and L4S enabled, under traditional home network load. You are looking at 30% improvement in the average latency, but more importantly the latency is very consistent with very little jitter above 100ms
- Some of the jitter spikes observed do not originate from the in-home network itself, but from the general internet and gaming service platform.

*Without PI2 and L4S

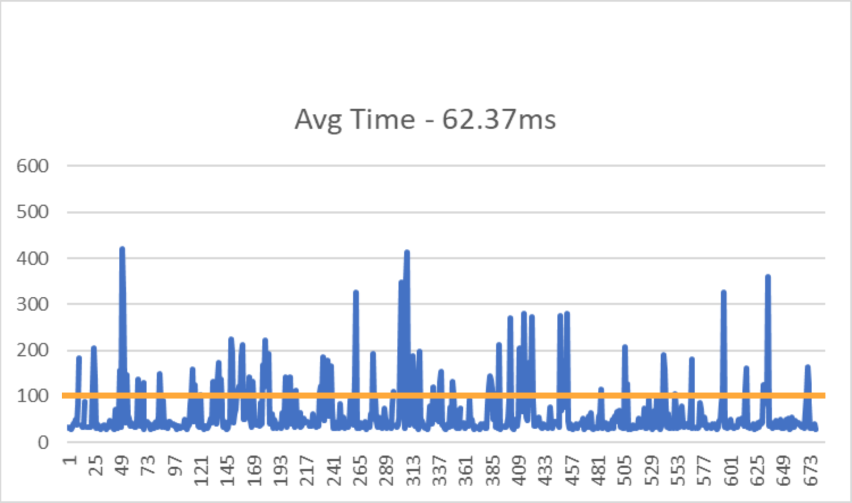

*With PI2 and L4S

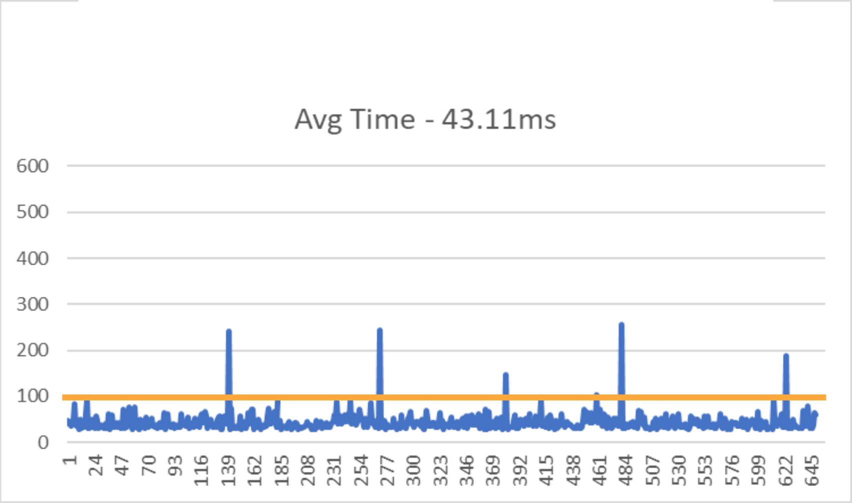

My personal K/D ratio improved by 0.5 in less than 4 weeks!

If I see a “lag spike” as a gamer, it’s already too late, but please tell me where it came from!

Source : Measurement and results for video-conferencing courtesy of DOMOS, 2021

Last-Mile & CPE Latency Video-Conferencing Benchmarking Results

4G Fixed-Wireless-Access with in-home WIFI

“Legacy” CPE measurements

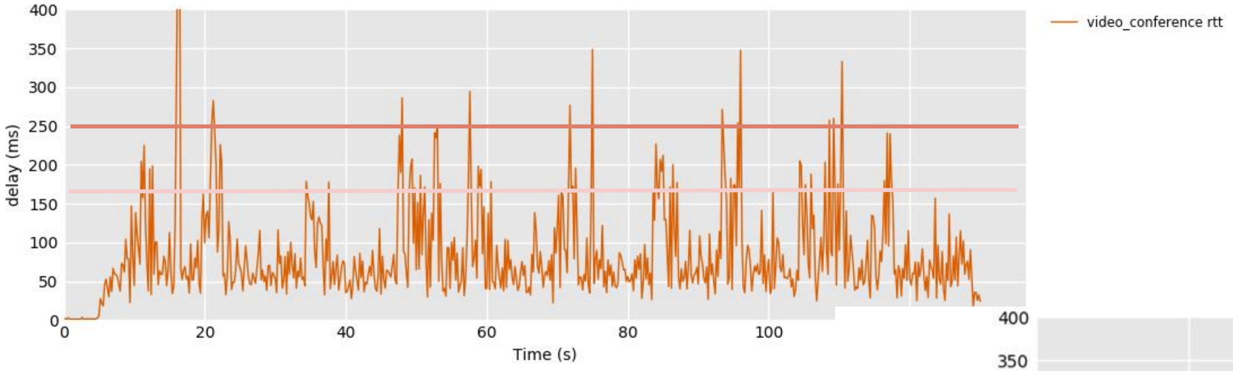

With PI2 and end-user managed latency

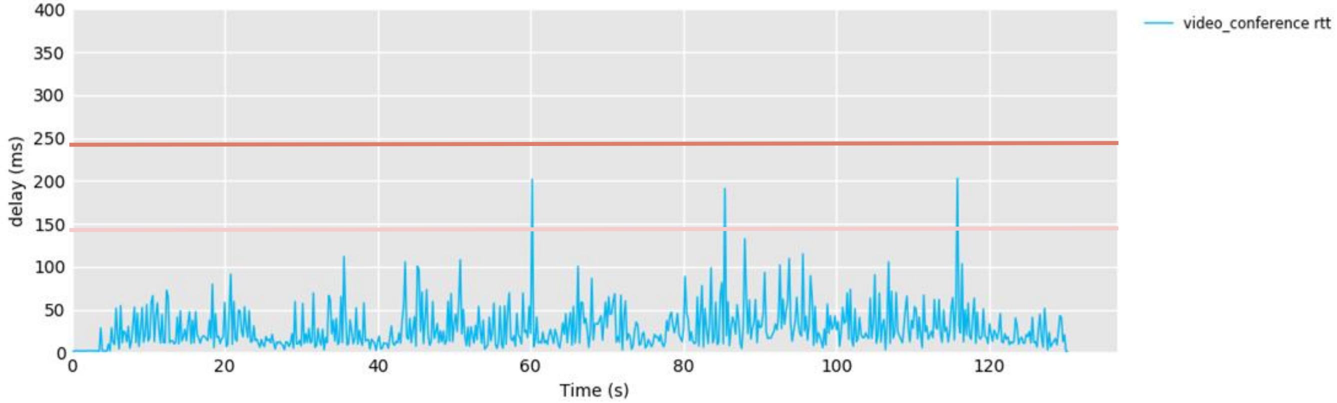

KPIs		Default	Domos + PI2	Change
Latency:	Peak	647	203	-69%
	99 percentile	266	96	-64%
	90 percentile	165	56	-66%
	Median	60	14	-77%
	Mean	78	23	-71%
Packet loss:	Rate:	0.00019	0.00019	0%

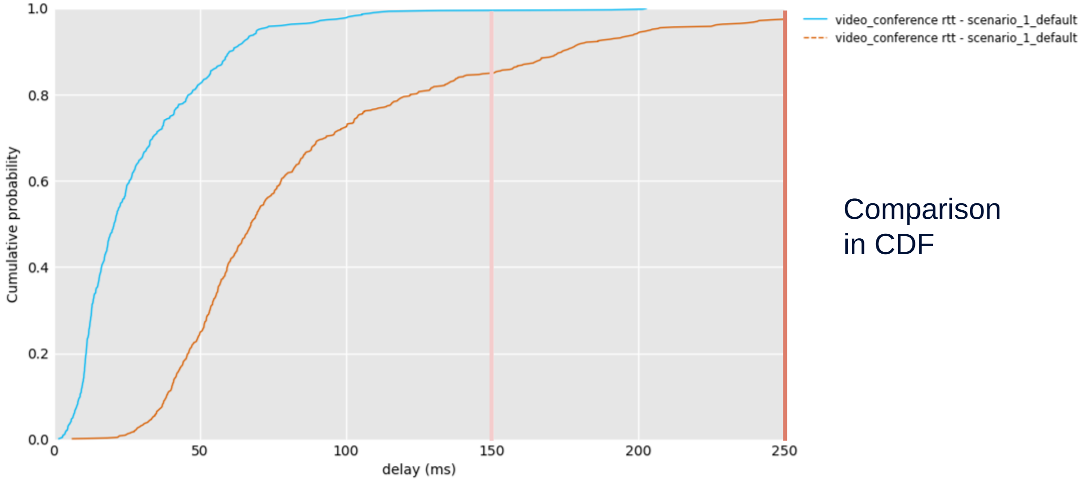

Comparison in CDF

Source : Measurement and results for video-conferencing courtesy of DOMOS, 2021

Cooperative approach to latency improvements

It’s no longer about agreement, but about alignment

- Focus on tangible, real-world benefits and results from the consumer point-of-view. Consumers don’t care about CDF curves, they care about K/D, FPS, etc. Measurements, comparisons and improvements of TCP stacks and queuing algorithms are important but will not drive consumer demand and service provider adoption. Focus should be placed on high-value use cases : Working-from-home, gaming, Enterprise services in a residential context, video/content streaming, etc

- Ping times are often misleading, and the “speedtest” has reached it’s end-game, we already provide nearly 300x more peak capacity in the last mile than the average sustained usage from broadband subs. You can’t easily sell more bandwidth since it doesn’t solve existing latency issues, nor can the consumer actually put it to any good use. Speedtest should be redefined as a function of what latency can be sustained under a given speedtest. Give latency a context.

- Applications and network elements need to have the same access and visibility to congestion markers so they can adapt natively, instead of probing the network and trying to outsmart it. We are at the point of diminishing returns if we continue the silo based approach to solving latency. L4S provides a proven means to deliver this today.

- The reality is that the latency consistency issues aren’t from the overall internet, they are largely confined to the last mile and in-home. Providing better tools to the end-users and the service providers than “ping” times for latency monitoring are required, there is little value in highlighting a lab spike if we can’t pinpoint the root cause.

- Platform and webscale providers such as Apple, Google, Netflix, etc. have been trying to tackle and solve latency issues for years (for their own services), without implicit support or help from the networks. But today’s networks and consumer CPEs are capable and ready to support this (ie. L4S).

Is the consumer market ready? Yes

How much would you be willing to pay as a consumer for an enhanced and improved "Working-from-home" networking experience, or a better "gaming" experience?	
1. 0$/ month	22.20%
2. 5$/ month	37.00%
3. 10$/ month	27.80%
4. 15$/ month	5.60%
5. 20$/ month	7.40%

Source : Nokia sponsored LightReading webinar, April 2021

Poll Question	Would you rather have the service provider offer a custom broadband package tailored for certain latency needs (ie Gaming, working-from-home, etc), or should they provide the end-user the ability to self-manage this capability?	
	service provider controlled latency	35%
	end-user managed latency	65%

Source : Nokia sponsored LightReading webinar, September 2021

Will CSP see value, and will the consumer?

- Average NPS score increase:
  - Enterprise sub: 10
  - Pilot participant: 20
- 100% of pilot participant retention rate
- How much would you miss feature (0 - 5):
  - 1 Prioritisation video conferences: 4.4
  - 2 Prioritisation specific device: 4.4
  - 3 Troubleshooting : 4.0
  - 4 Information about wi-fi issues: 3.6
  - 5 recommendations for wi-fi issues: 4.6

Survey results from Enterprise customers, live field trial, EU FWA provider, 2021. Offering at $10Euro more per month "extra"

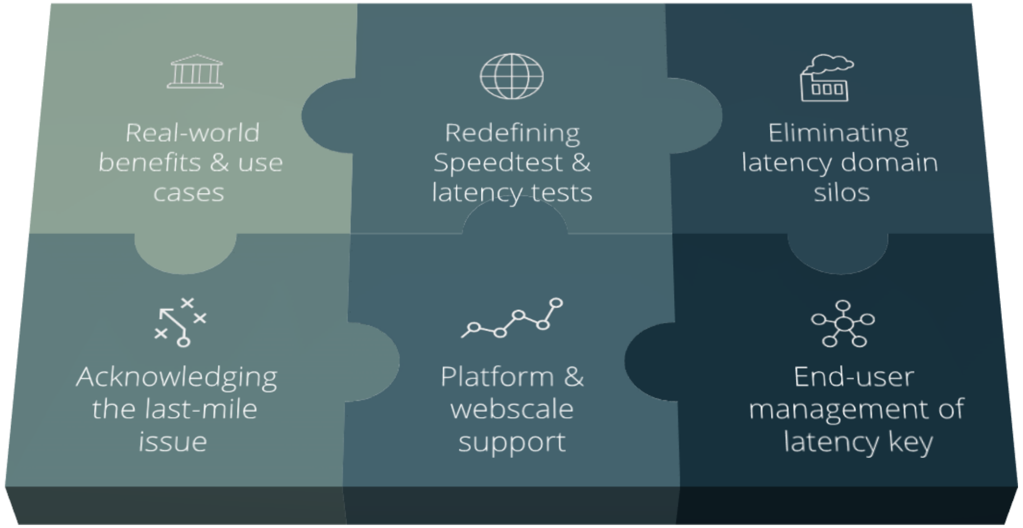

Cooperative approach to latency improvements

It’s no longer about agreement, but about alignment

There isn’t a day that goes by that I don’t get a Tier1/2 CSP from around the world asking how they can offer a better latency service (gaming, working, from home, enterprise services in a residential content with SLA, etc)

I’m putting it out there, just ask, and we can connect the dots and make this happen!

Transport Layer Statistics for Network Quality

Praveen Balasubramanian

Transport Layer Metrics for Network Quality

- Transport layer metrics
  - RTT, Loss, Reordering, Bandwidth
  - Estats / TCP_INFO
  - Push model?
  - QUIC has additional complexity

- Correlation with application layer metrics
  - Time to first byte, Response time, Delivery rate, Jitter
  - Connection reuse problems
  - Multi-stream problems
  - Support for stats diff @ request start/end?

Challenges and Research Ideas

- Challenges
    - Multi-platform world
    - Different administrative domains
    - Application intent
        - Not expressed today
        - Difficult to infer
    - User intent
        - Low versus high priority
    - Receive side transport has less visibility

- Research Ideas
    - Design application layer mechanisms to exchange transport metrics
        - Privacy preserving
    - API improvements for expressing application intent
    - Mechanisms for inferring user intent
    - Passive and active measurements

Discussion

To be enqueued, please write ‘+q’ in the chat

To cancel being enqueued, please write ‘-q’ in the chat

The duration of each comment is limited by 60 seconds

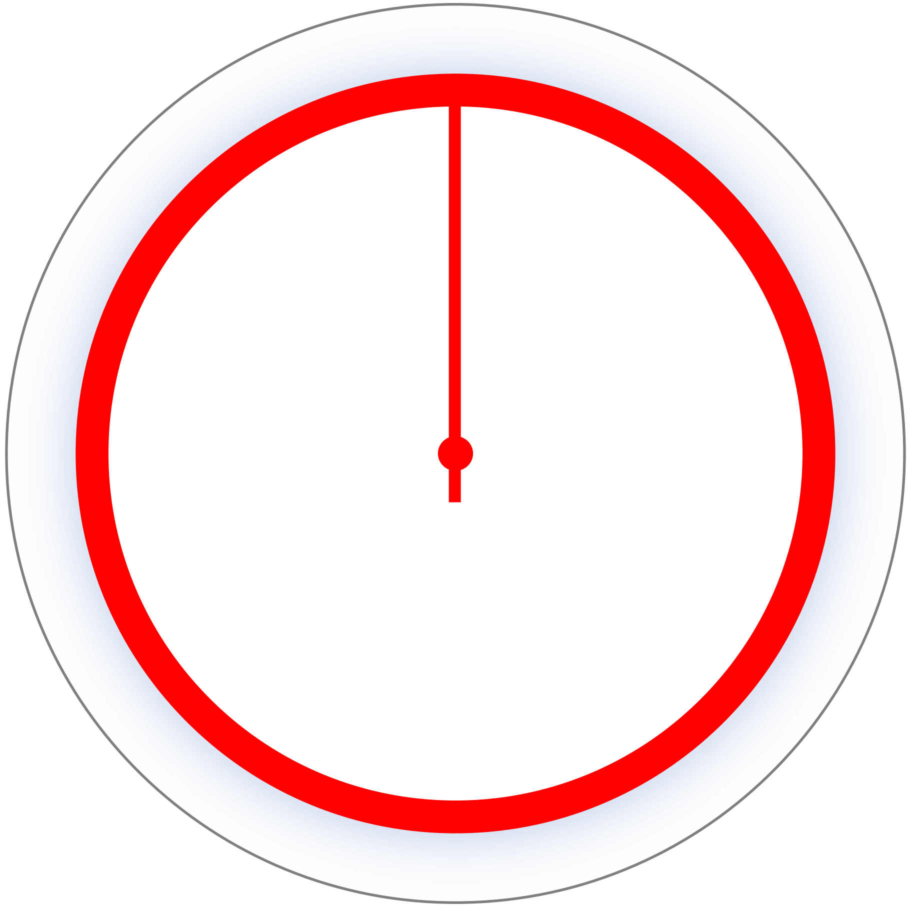

Observability is needed to improve network quality

Jari Arkko

Session: Cross Layer
Topic: Observability

Arkko & Kühlewind, 2021

The problem:
Unsatisfactory application behaviour. But why? Wifi? Cloud? Operator?
- Limited observability
- Problem isolation and debugging is difficult
- Many aspects of desirable behaviour are not directly visible
- Dynamically changing situation is not shared to others

Direction for solutions:
- Better collaboration among parties
- Observability, explicit identification of the situation

Constraints:
- Only for mutually beneficial cases (RFC 8558)
- Needs standards
- Avoid data with privacy, filtering, etc. impacts

Examples
- Network assist for measurements
- Endpoint assist for measurements
- Probing
- Capability discovery
- Security indications

Merge Those Metrics: Towards Holistic (Protocol) Logging

Robin Marx

Merge Those Metrics:
Towards Holistic (Protocol) Logging

Measuring Network Quality for End-Users. IAB Workshop 2021

Robin Marx
Joris Herbots

robin.marx@kuleuven.be
joris.herbots@uhasselt.be

KU LEUVEN
UHASSELT

Cross-Layer and Cross-Endpoint log aggregation

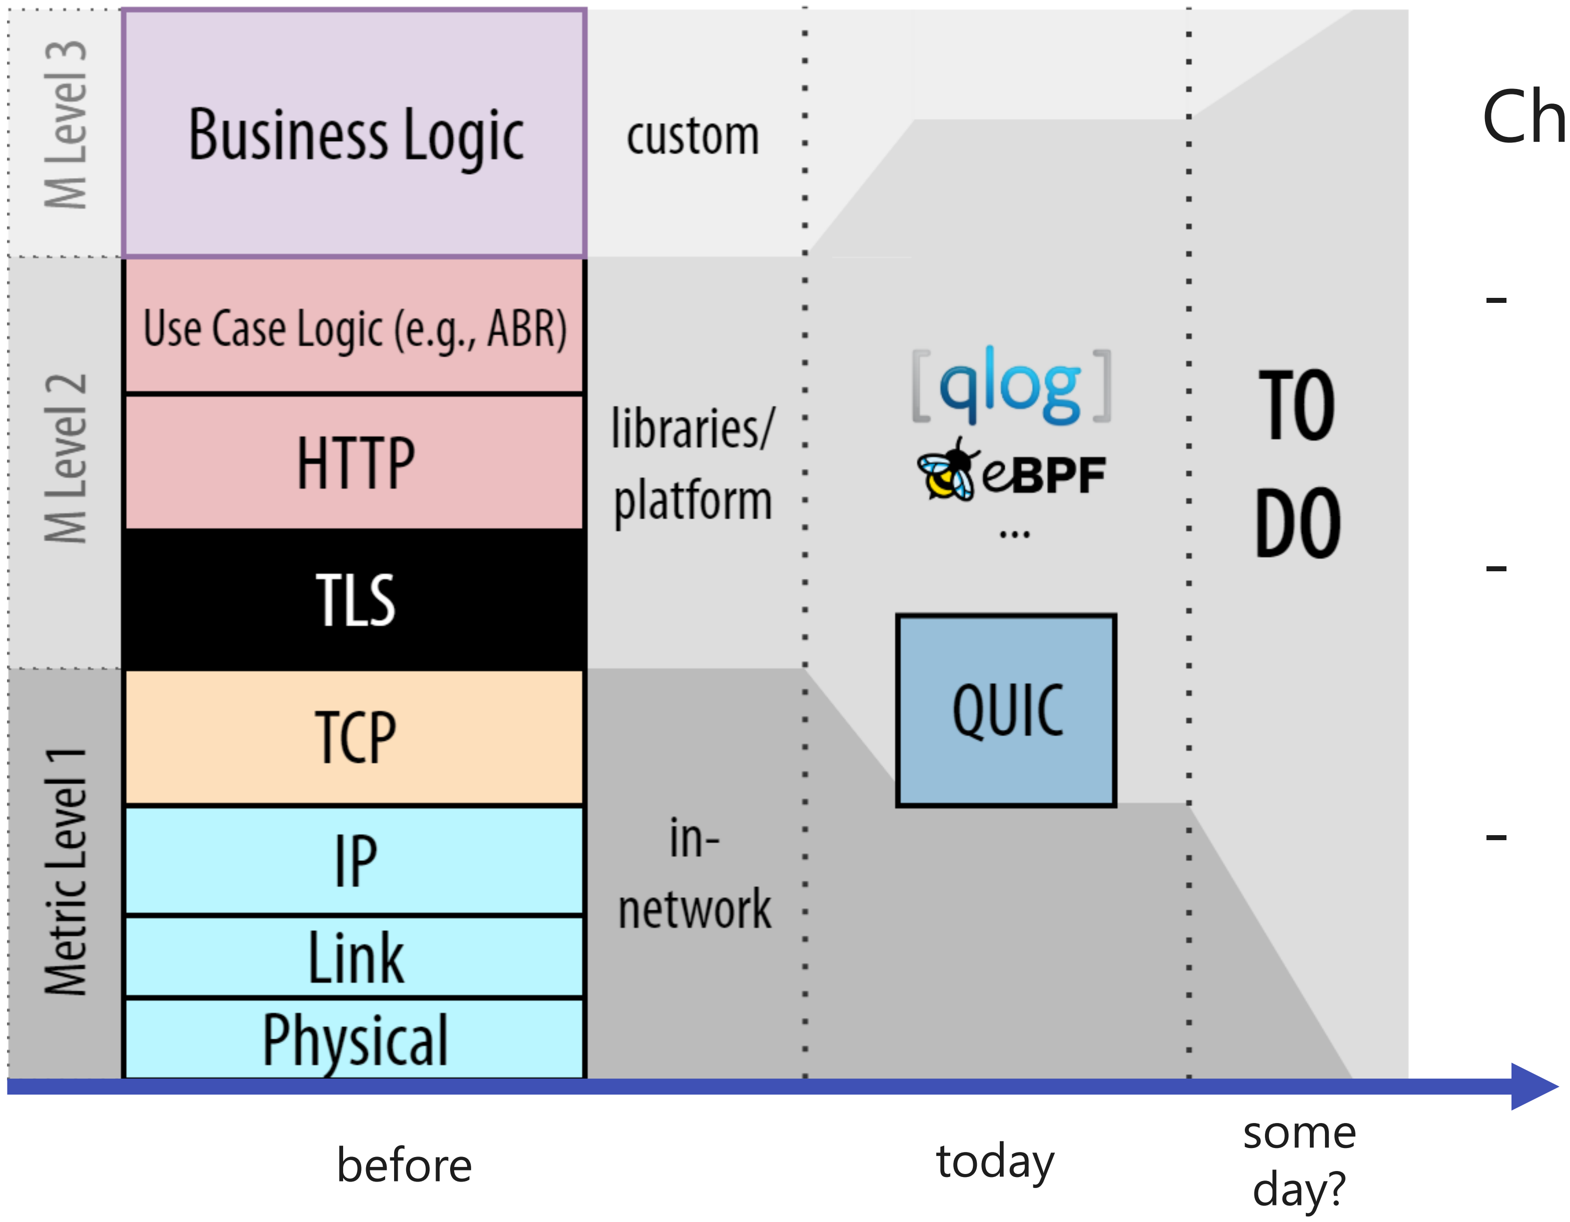

Challenges aplenty:

- Store, aggregate and synchronize post-hoc
- Analyze and correlate with tooling
- Network emulation / (partial) trace playback

Share That Data (Please)!

Especially for low-layer-encrypted future,

endpoint logging is the only thing (?) that scales + deals with privacy

Observability

- Share metrics with (privileged) network operators?

Industry vs Academia: datasets

- Real loss? Real MTUs? Real geographical metric distributions? ...
- Privacy? Intellectual Property/Competitive Edge? *Afraid we might laugh?*

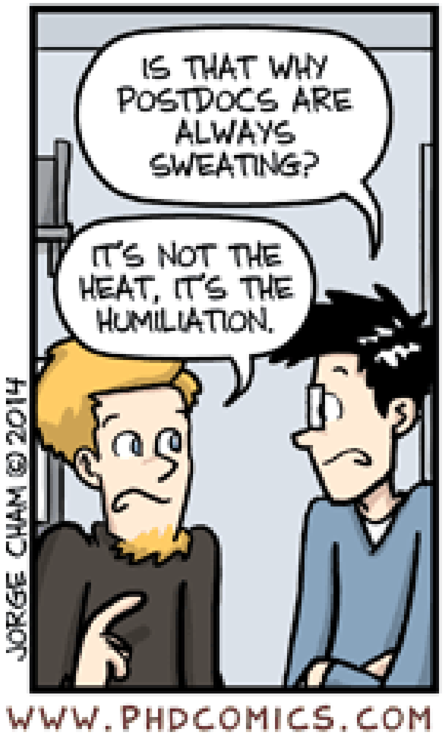

Measuring & Improving QoE on the Xfinity Wi-Fi Network

Rajat Ghai

Using TCP Connect Latency for measuring CX and Network Optimization

Network Quality workshop

Comcast

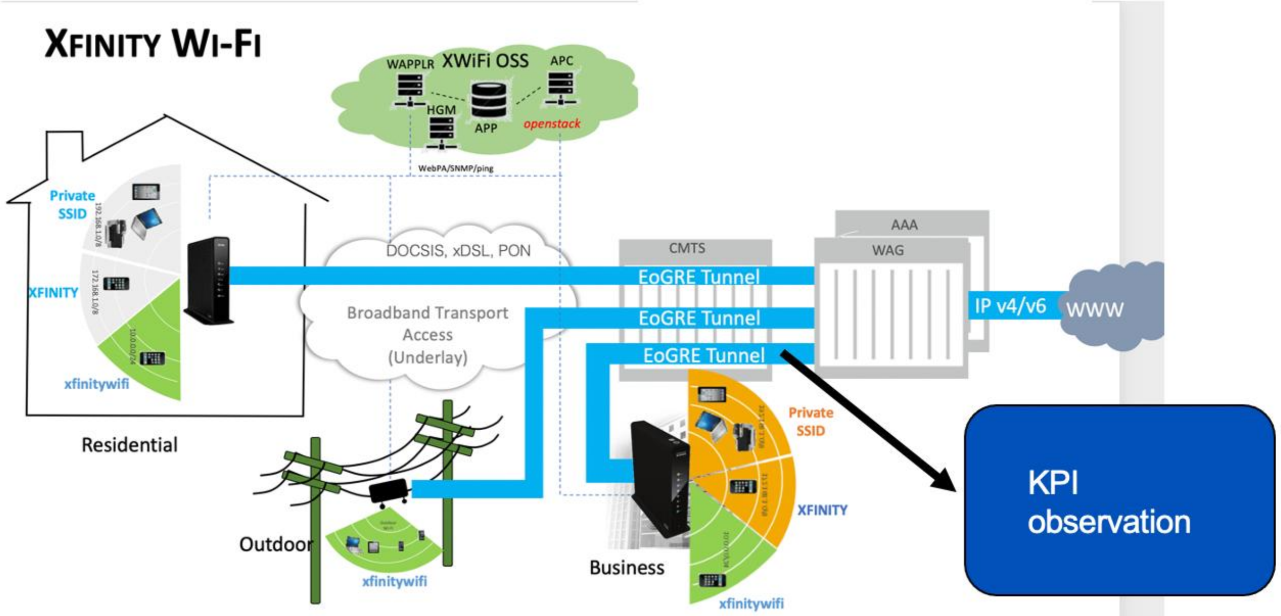

QoE KPI

TCP conn: 150 msec

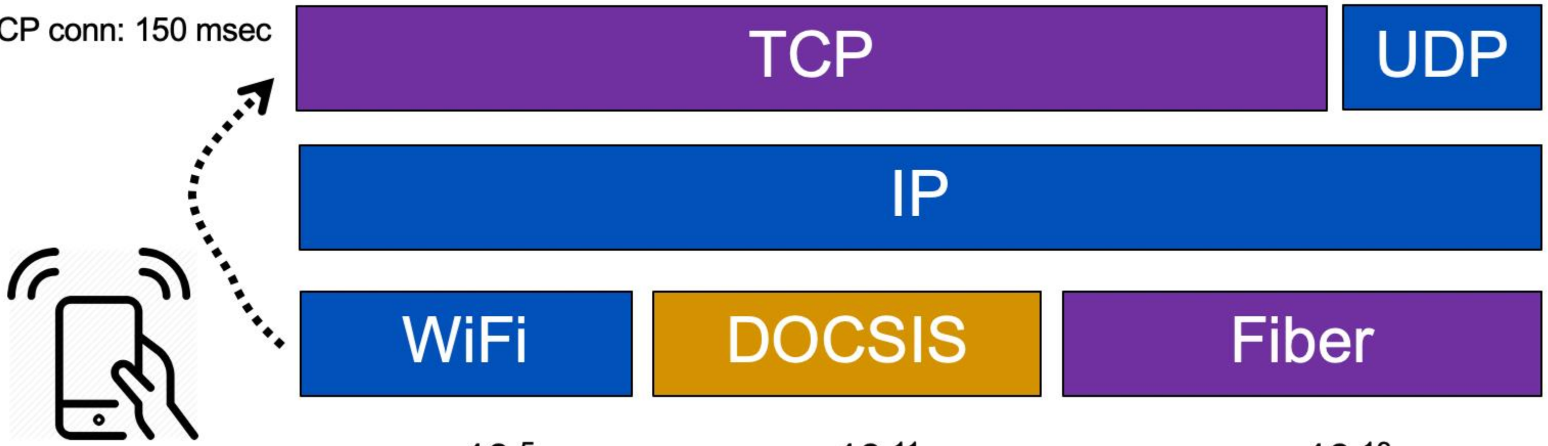

$\sim 10^{-5}$ $\sim 10^{-11}$ $\sim 10^{-13}$ BER

WiFi connection quality can be observed at Layer 4
by observing TCP connect latency

CX KPI Predictors

- **<u>TCP connect Latency</u>**: This is our Primary Latency predictor KPI. It is measures from SACK→ACK, as round trip from network to client and back.

- **<u>TCP Retransmission</u>**: This is our secondary KPI that we are evaluating as a predictor of characterizing good/bad sessions. This is measured for the **<u>entire lifetime</u>** of tcp connection NOT just the start. This marker at 4% is candidate for the CX predictor (specially for mobility / nomadic scenarios where wireless channel quality variability due to motion creates excessive L2 delays which in turn creates retransmissions at TCP layer.

- **<u>Throughput</u>**: Throughput is measured as actual consumed in octets on a 1 second interval granularity. {*Not to be confused with a synthetic speedtest that tests the max capacity capability of a connection*}

TCP Connect Latency

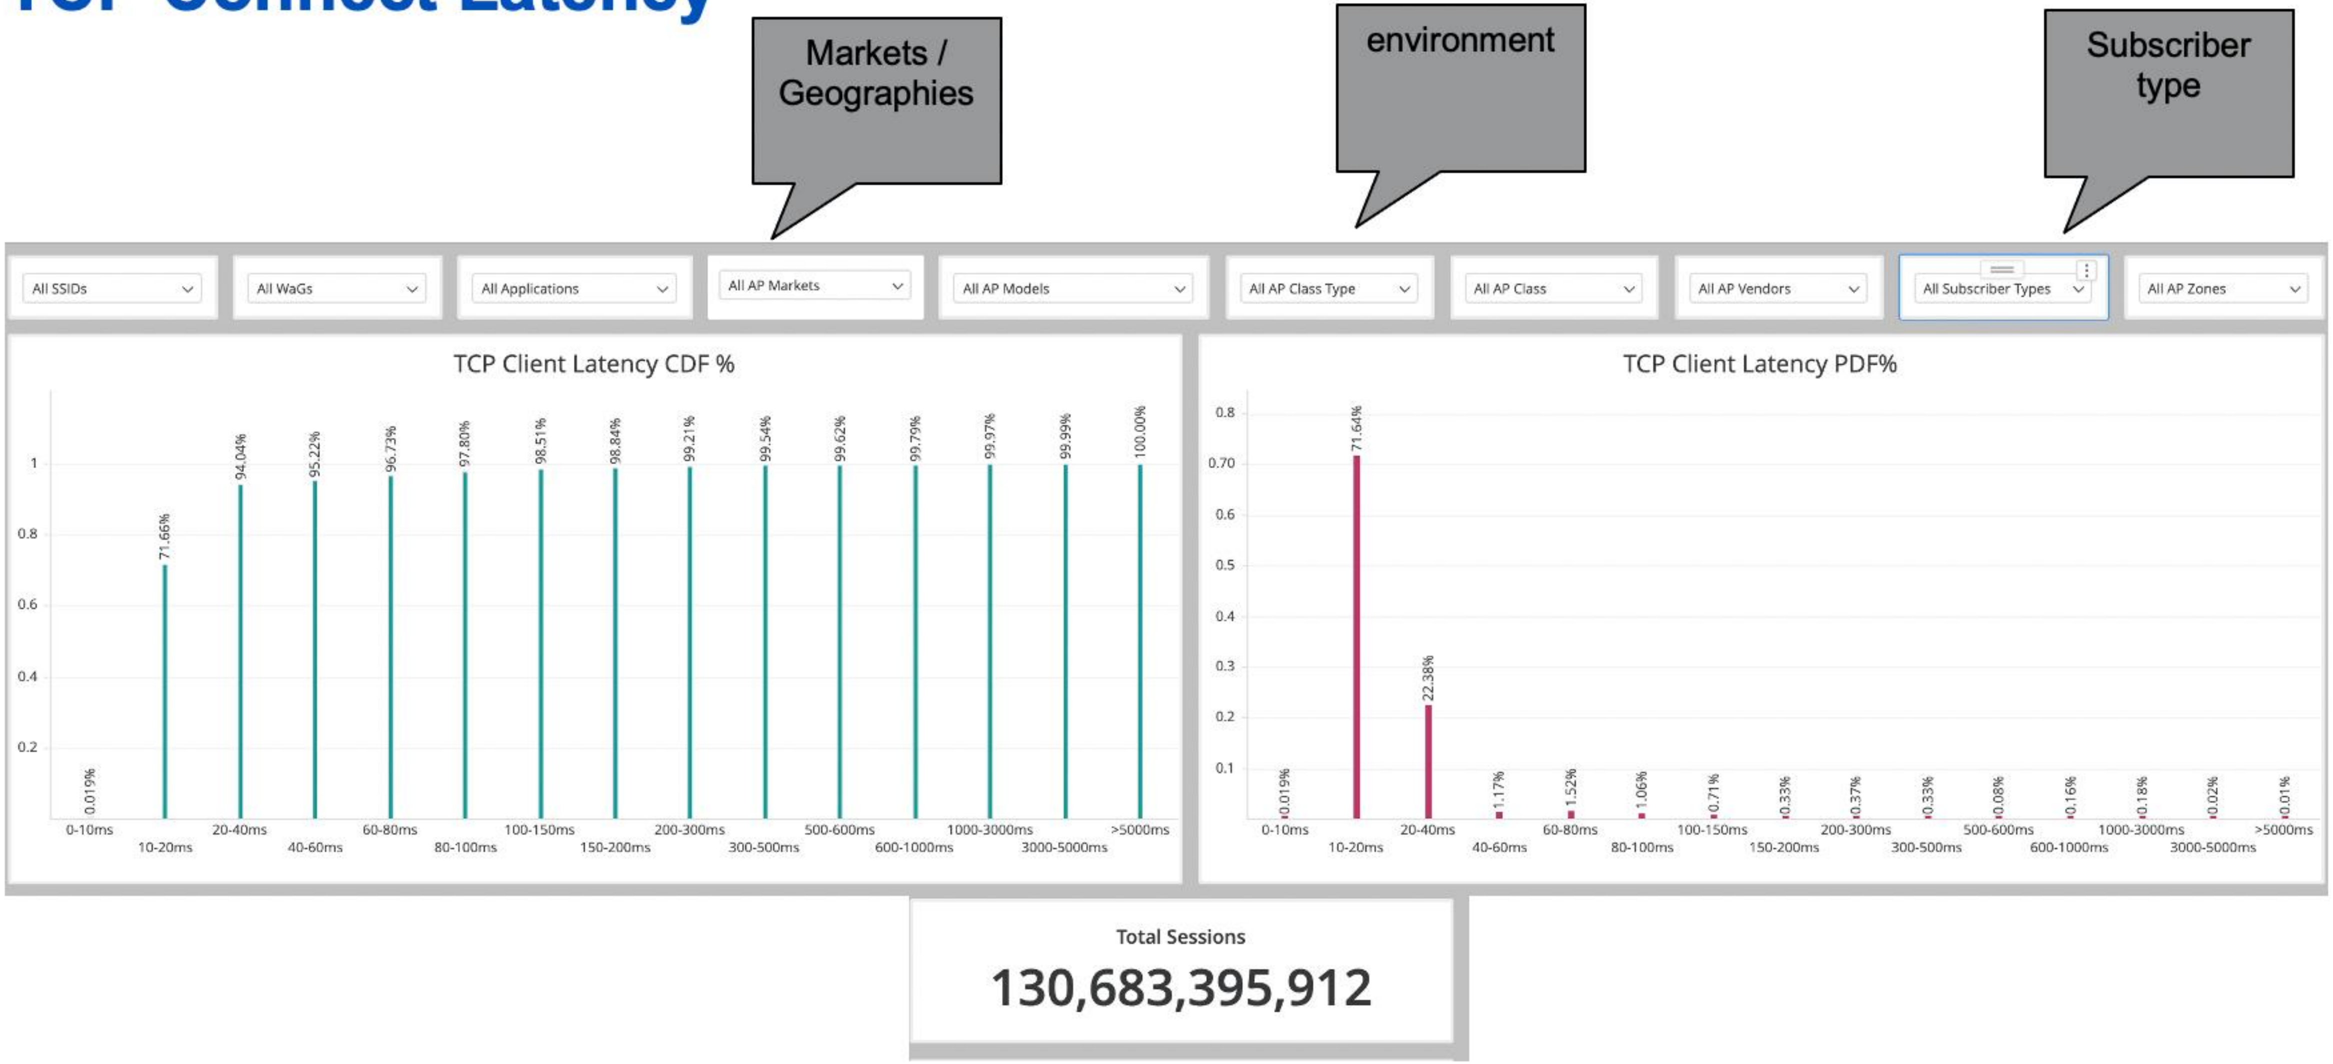

"We collect, store, and use all data in accordance with our privacy disclosures to users and applicable laws".

TCP Retransmissions

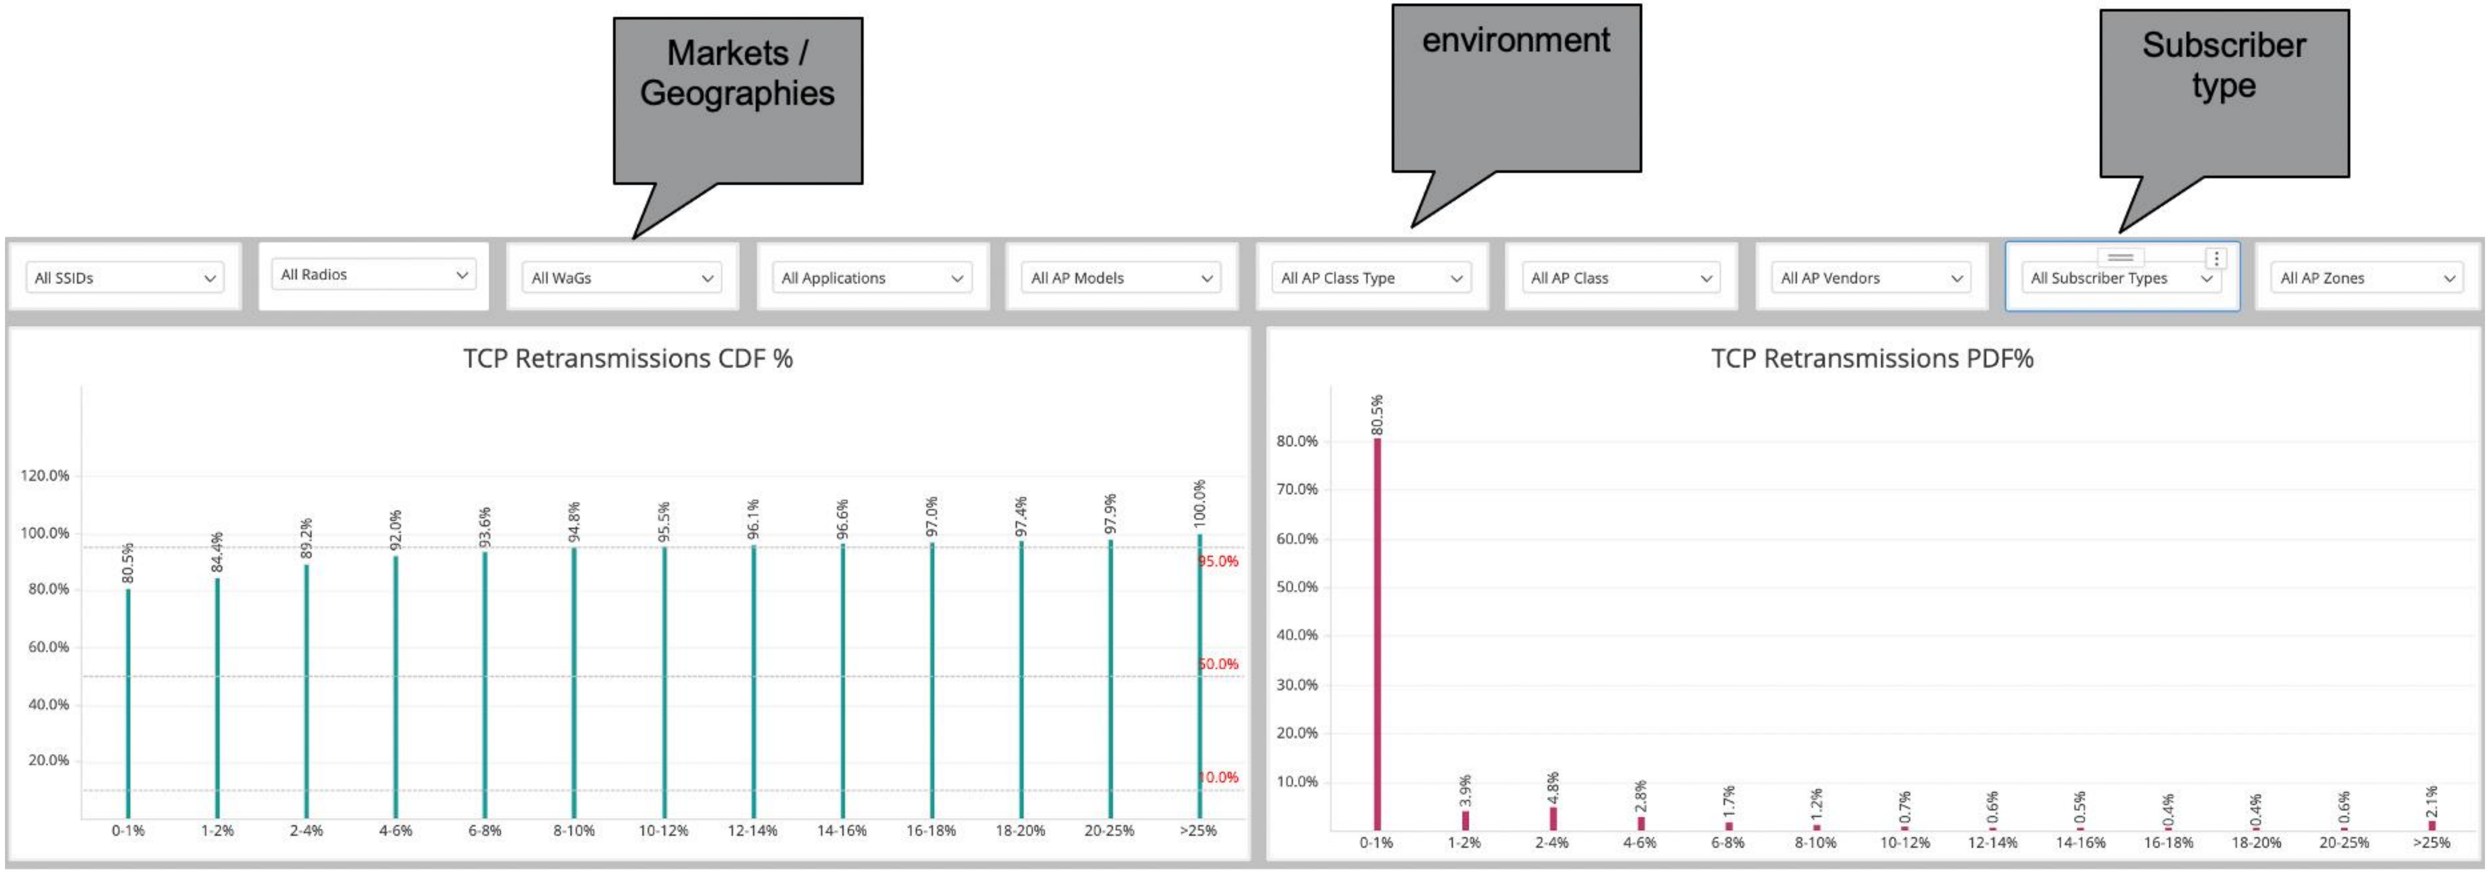

"We collect, store, and use all data in accordance with our privacy disclosures to users and applicable laws".

CX / Latency improvement methodology

- Conduct Network A/B tests for network optimization
  - Pick A & B sites of very similar traffic and usage profiles.
  - Site A is the Control site
  - Site B is the Test site
  - Baseline and trend KPIs on both sites for fixed duration of time (usually 2 weeks)
  - Make the CX based Network Optimization change only on the Test site
  - Let the Test site (site B) soak the change for a fixed duration of time (usually 1-2 weeks)
  - Start trending KPIs on both control and test sites during post baseline and soak phase for a fixed duration of time (usually 2 weeks)
  - Analyze the network optimization KPI outcome on test site as compared to control site.
  - If KPI trends are favorable on various KPI statistical models, then deploy the change nationwide.

  - Wash, rinse , repeat.

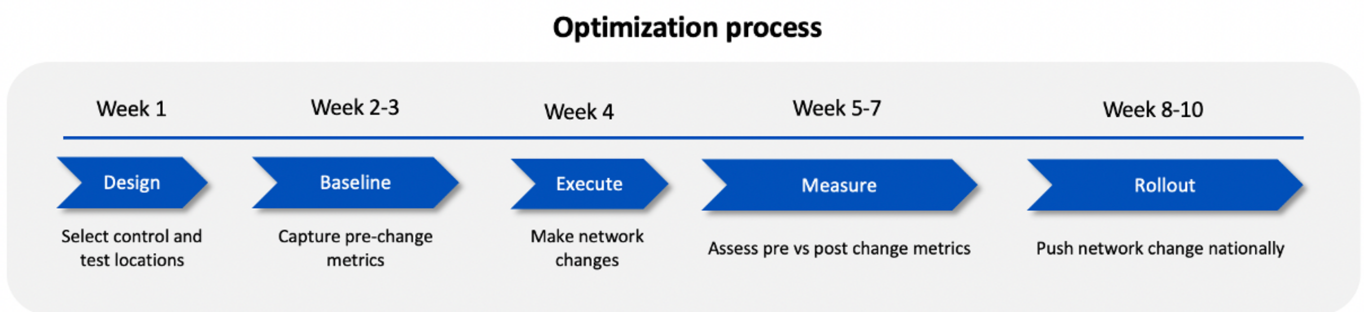

Discussion

To be enqueued, please write ‘+q’ in the chat

To cancel being enqueued, please write ‘-q’ in the chat

The duration of each comment is limited by 60 seconds

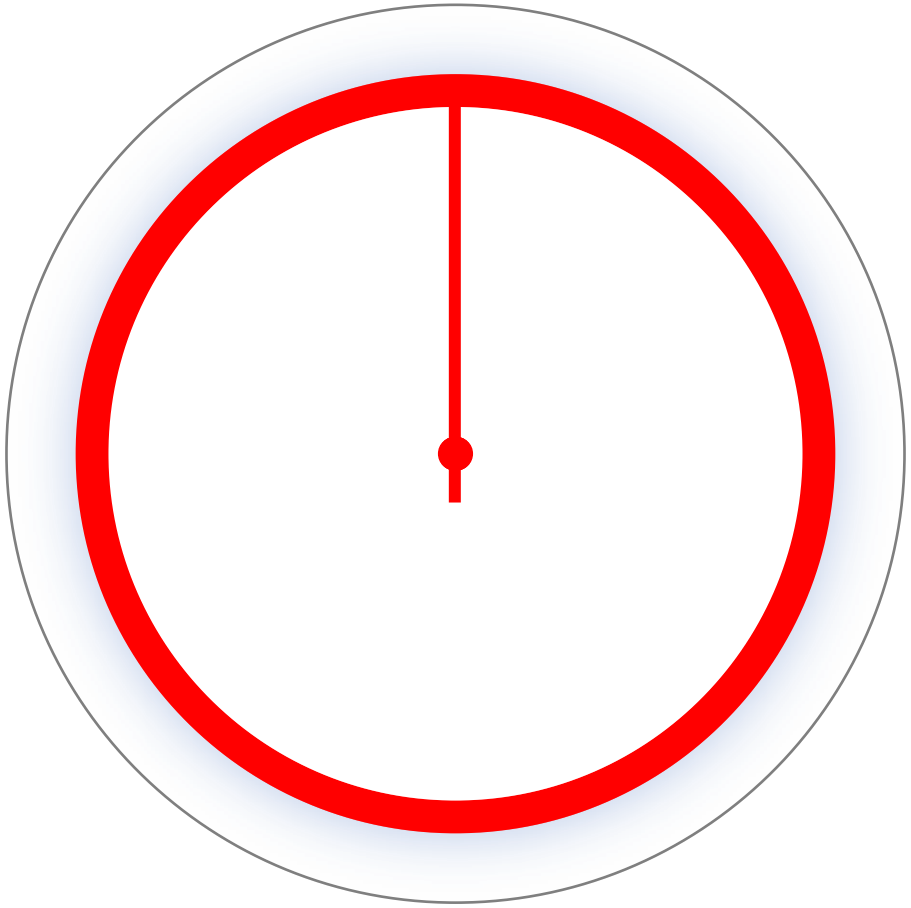

Challenges and opportunities of hardware support for Low Queuing Latency without Packet Loss

Koen De Schepper

Challenges and opportunities of hardware support for Low Queuing Latency without Packet Loss

Koen De Schepper, Olivier Tilmans, Gino Dion

IAB Workshop on Measuring Network Quality for End-Users

15 September 2021

Low latency with rate priority Low latency without rate priority

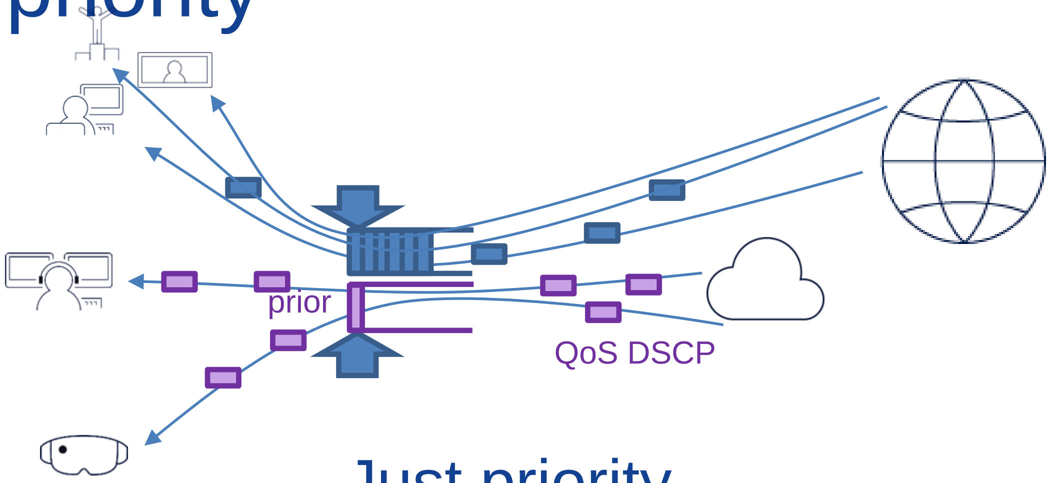

Just priority

Give apps all BW they need

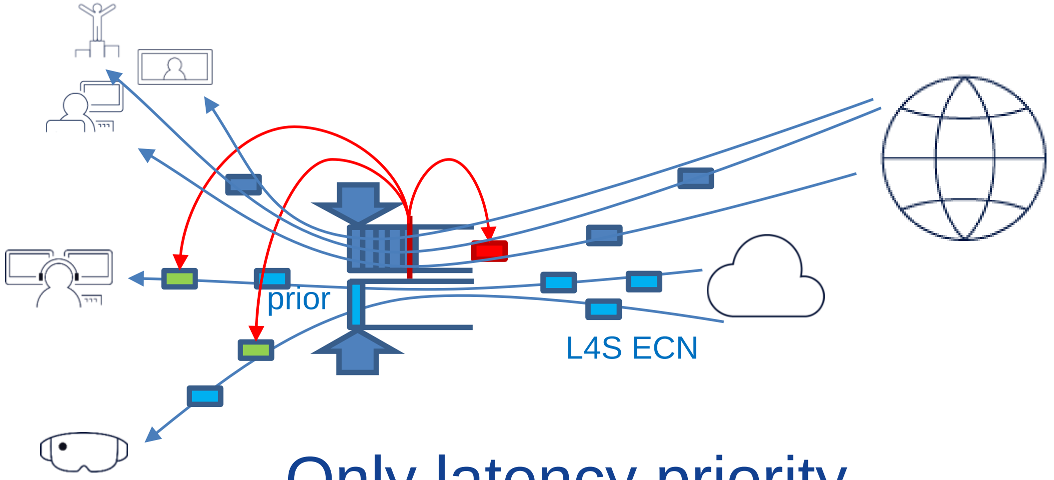

Only latency priority

Control rate of all flows equally

Low latency with rate priority □ Low latency without rate priority

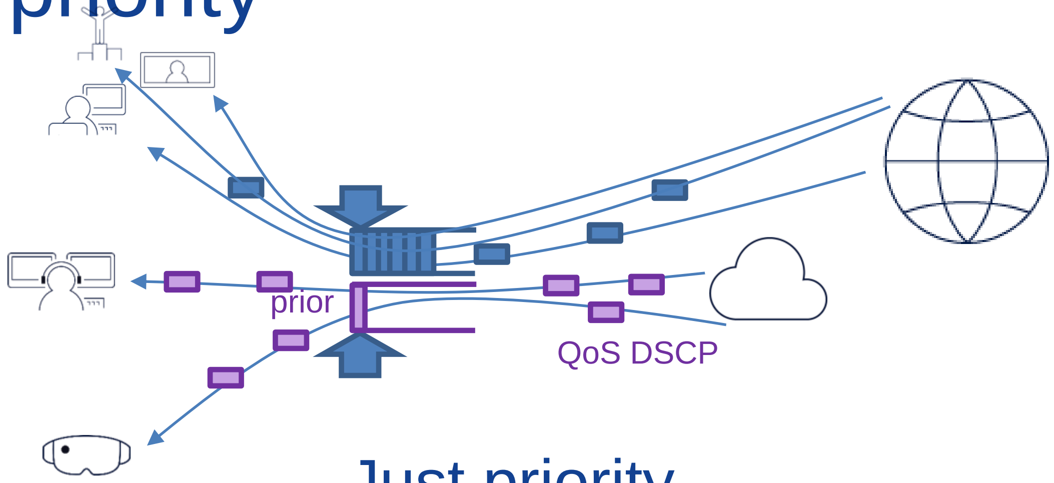

Just priority
Give apps all BW they need

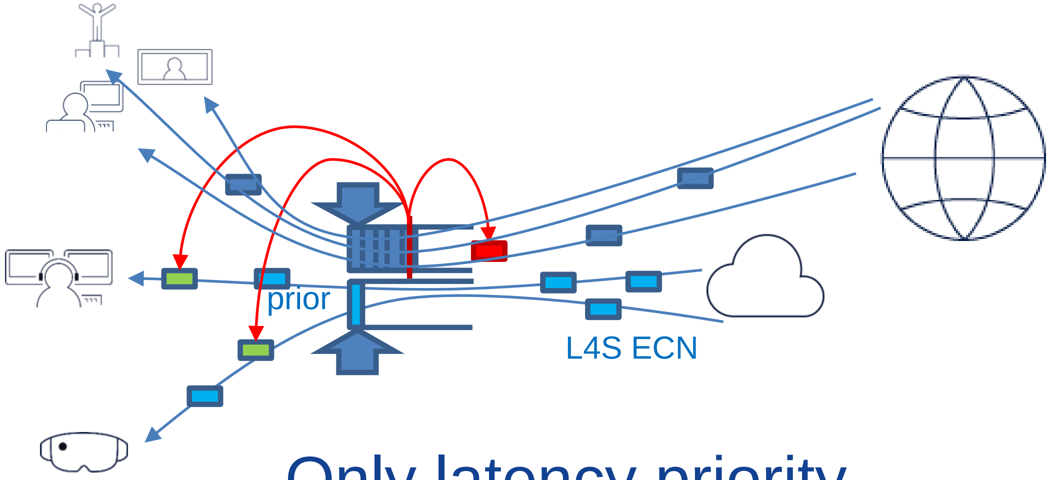

Only latency priority
Control rate of all flows equally

New mindset for lower layers:

Low Latency ≠ priority ➔ Low Latency = priority + rate control

Priority **NEEDS** rate coupling

Bigger Queue ≠ needs more bandwidth

Design reorientation needed from speed towards latency

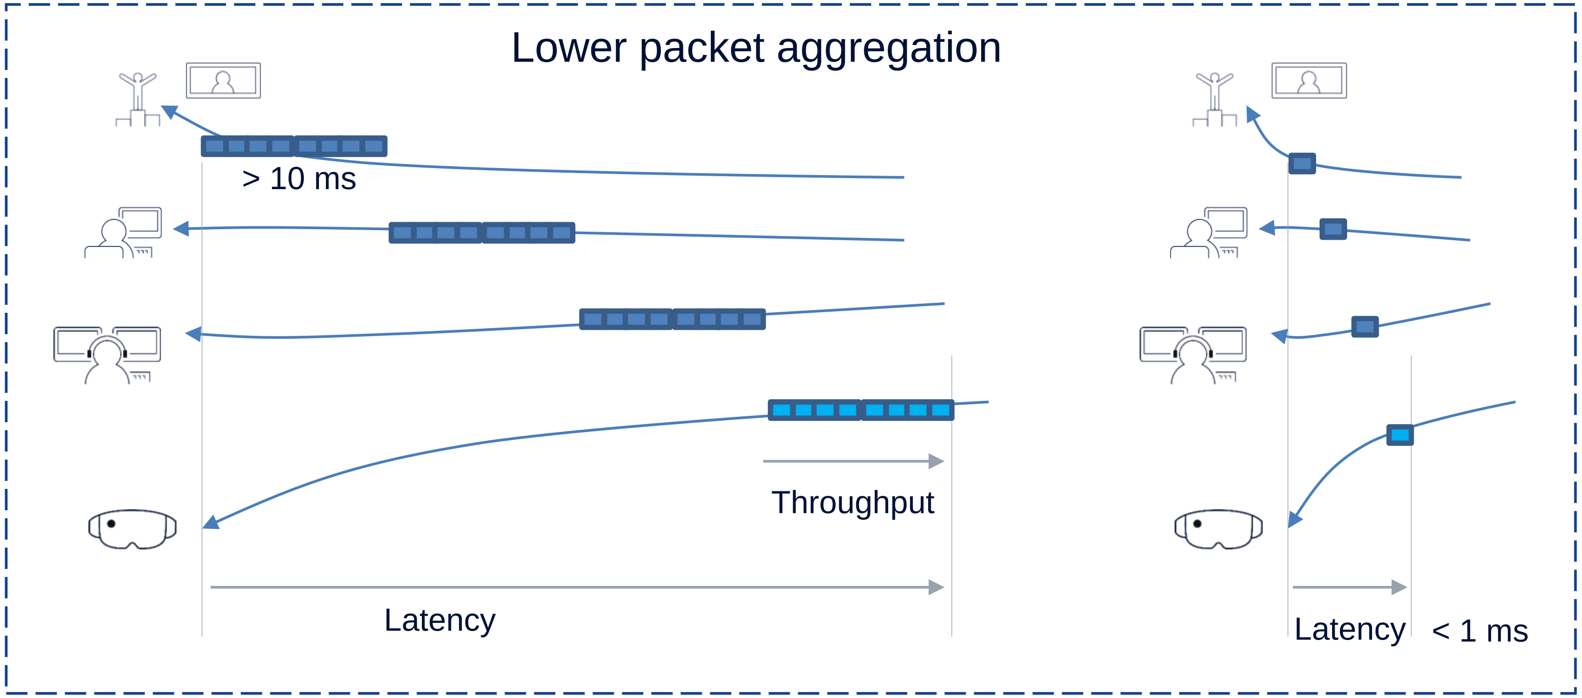

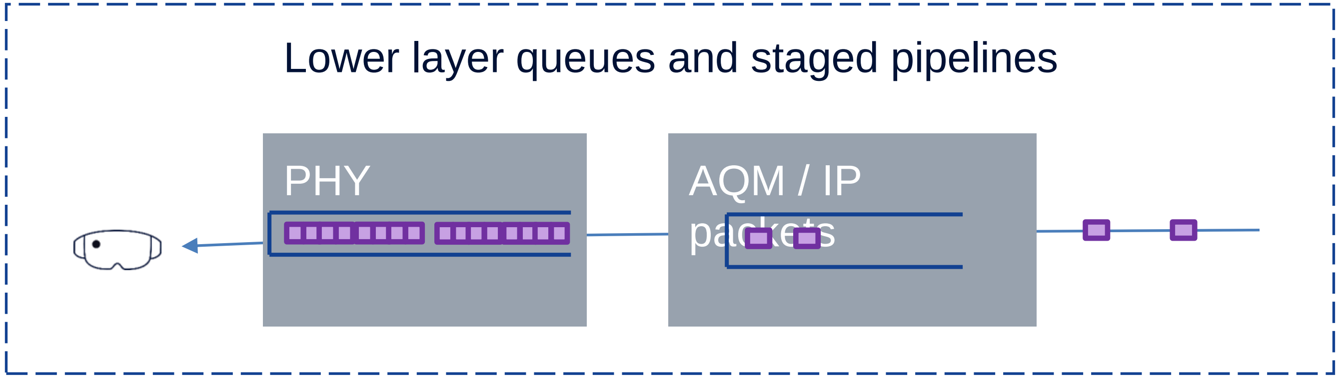

Design reorientation needed from speed towards latency

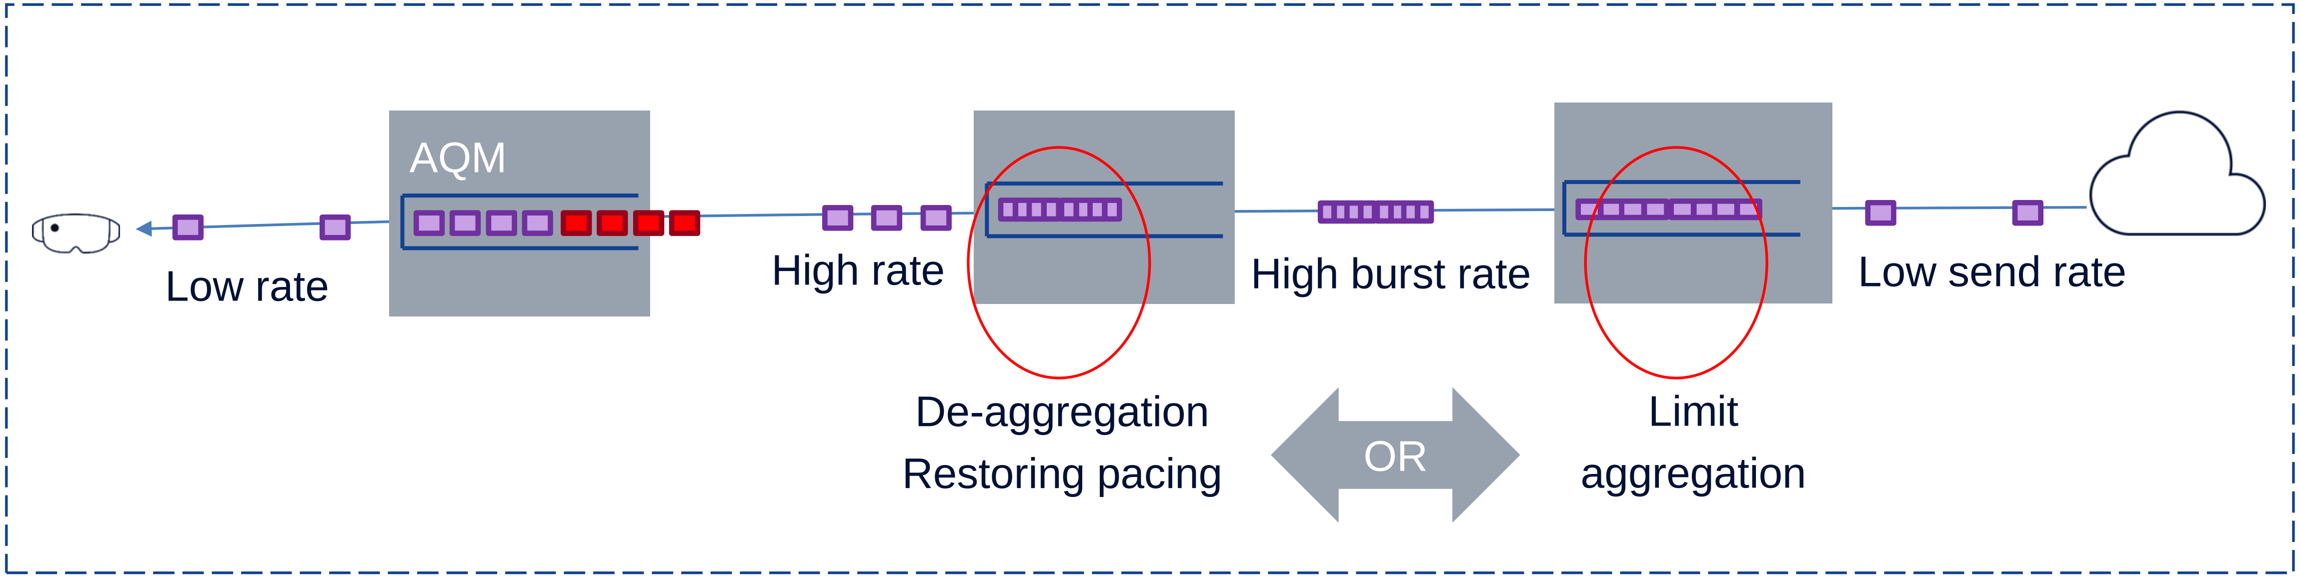

Optimize ECN usage:

- Access ECN in IP headers
- Support ECN in lower layers
- …

Measuring latencies

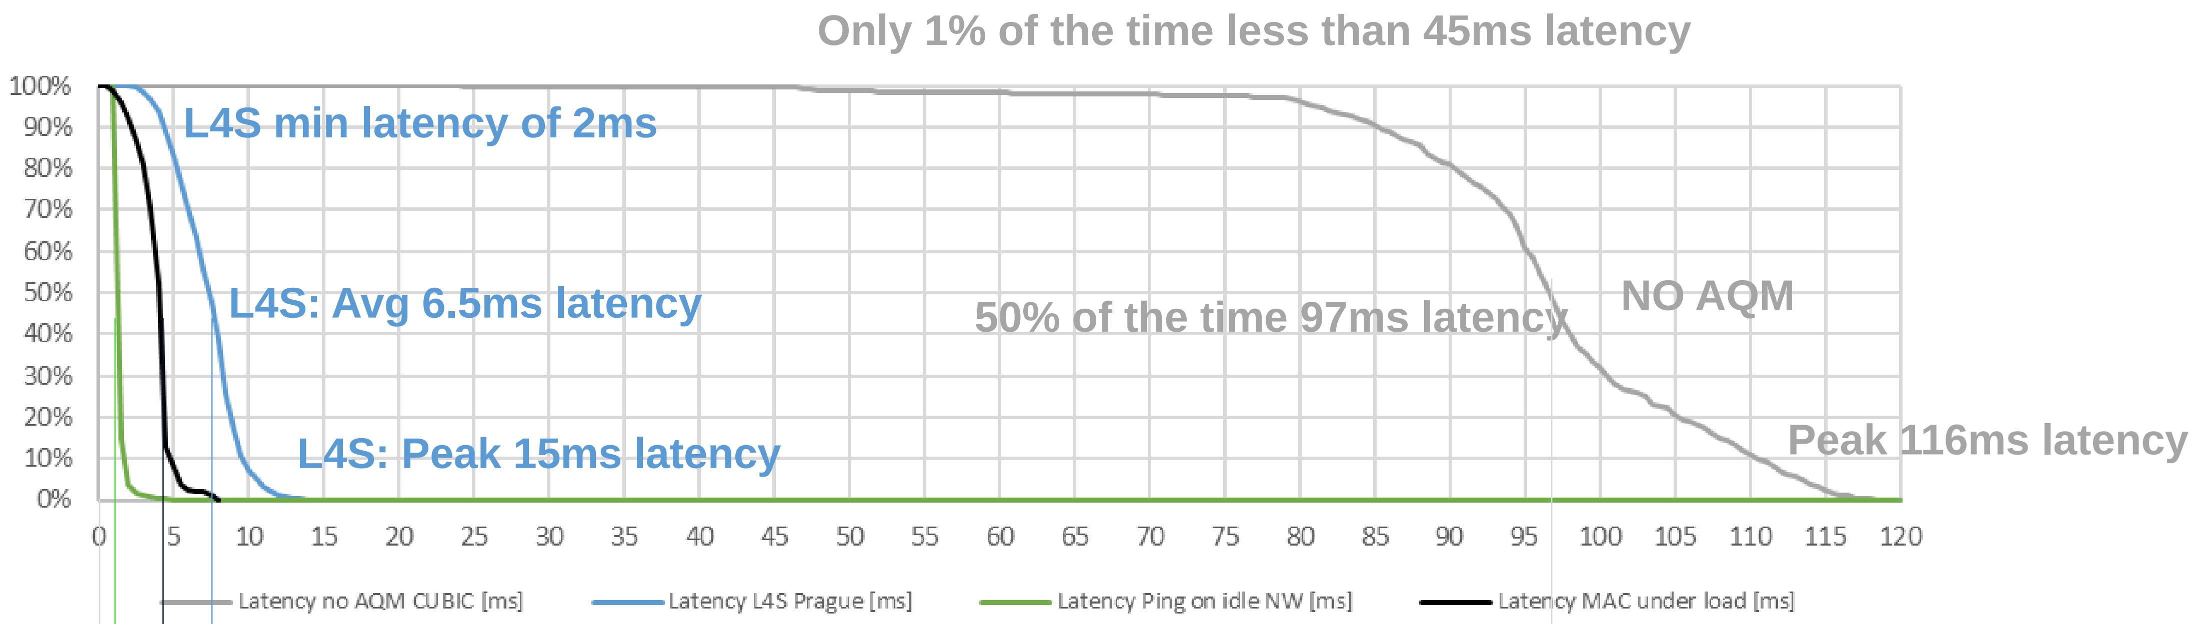

Bufferbloat MUST be avoided

1 Extra Packet aggregation time due to multistage pipelining can be avoided

Packet aggregation / serialization time can be reduced

Base processing time

Should Network report latency per application?

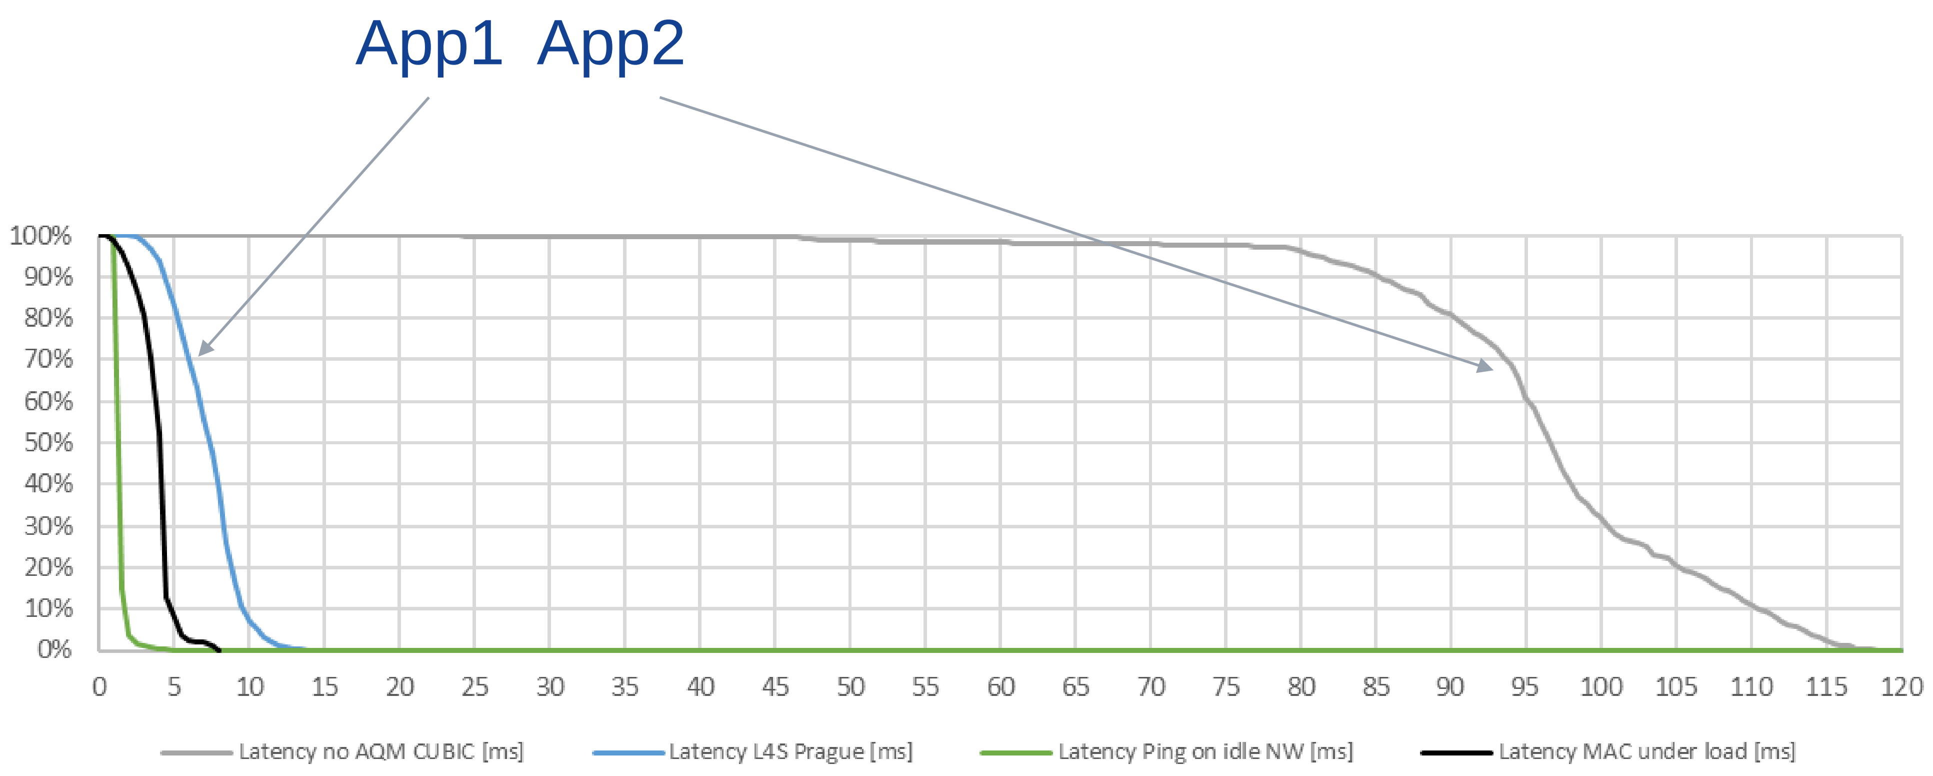

State of Wi-Fi Reporting

Ken Kerpez

State of Wi-Fi Reporting, IAB workshop on Measuring Network Quality for End-Users, 2021

John Cioffi, Jinous Shafiei, Ken Kerpez, Pete Chow, and Djamel Bousaber; ***ASSIA Inc.***

- ASSIA collects and anonymizes many data parameters
  - Millions of lines in both *North America* and *Europe*
  - Performance metrics, Diagnostic parameters, Network status & test

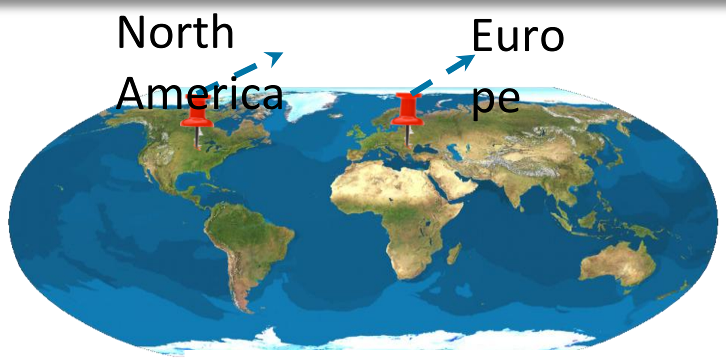

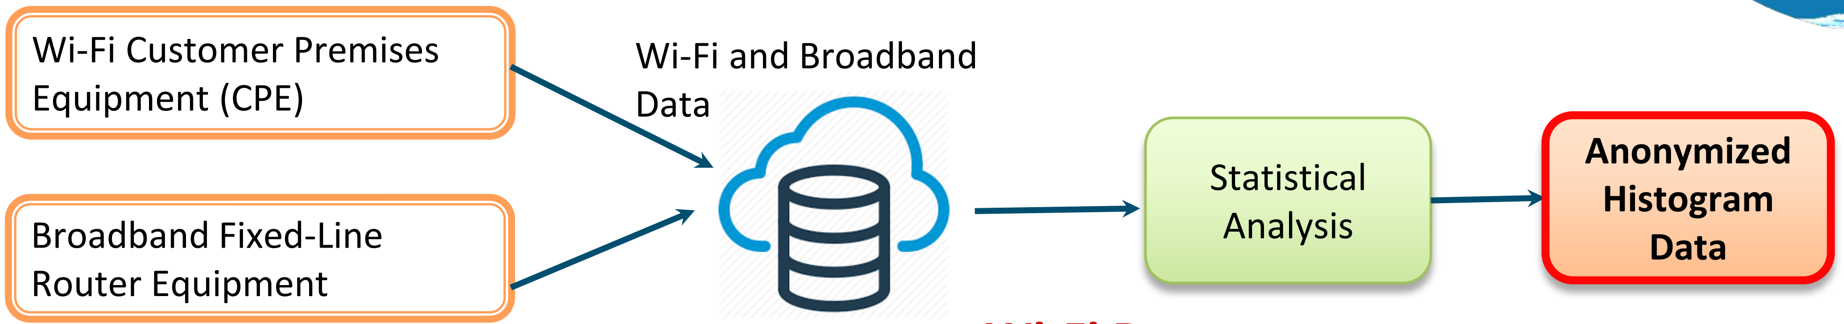

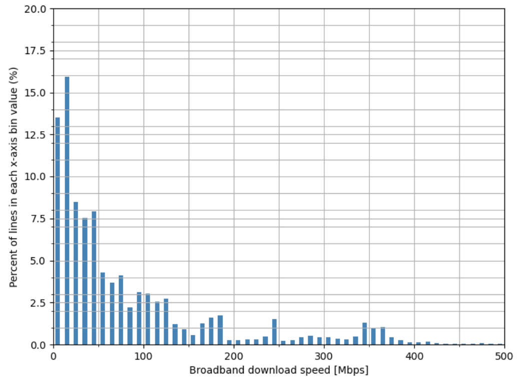

Wi-Fi Data

Wi-Fi Throughput (speed)	Daily, 2.4 and 5 GHz bands
Wi-Fi Transmit Rate	Daily, 2.4 and 5 GHz bands
Wi-Fi Throughput to transmit rate ratio	Daily, 2.4 and 5 GHz bands
Wi-Fi Congestion	Daily and max hour, 2.4 and 5 GHz bands
Wi-Fi Interference	Daily and hourly, 2.4 and 5 GHz bands
Wi-Fi Traffic	Daily and hourly, upstream and downstream, 2.4 and 5 GHz bands
Wi-Fi Latency	Daily, 2.4 and 5 GHz bands

Broadband Data

Broadband Traffic	Daily and hourly, upstream and downstream
Broadband Throughput (speed)	Daily, upstream and downstream
Broadband Latency (round-trip)	Daily

Wi-Fi Traffic increases & Wi-Fi limiting internet access

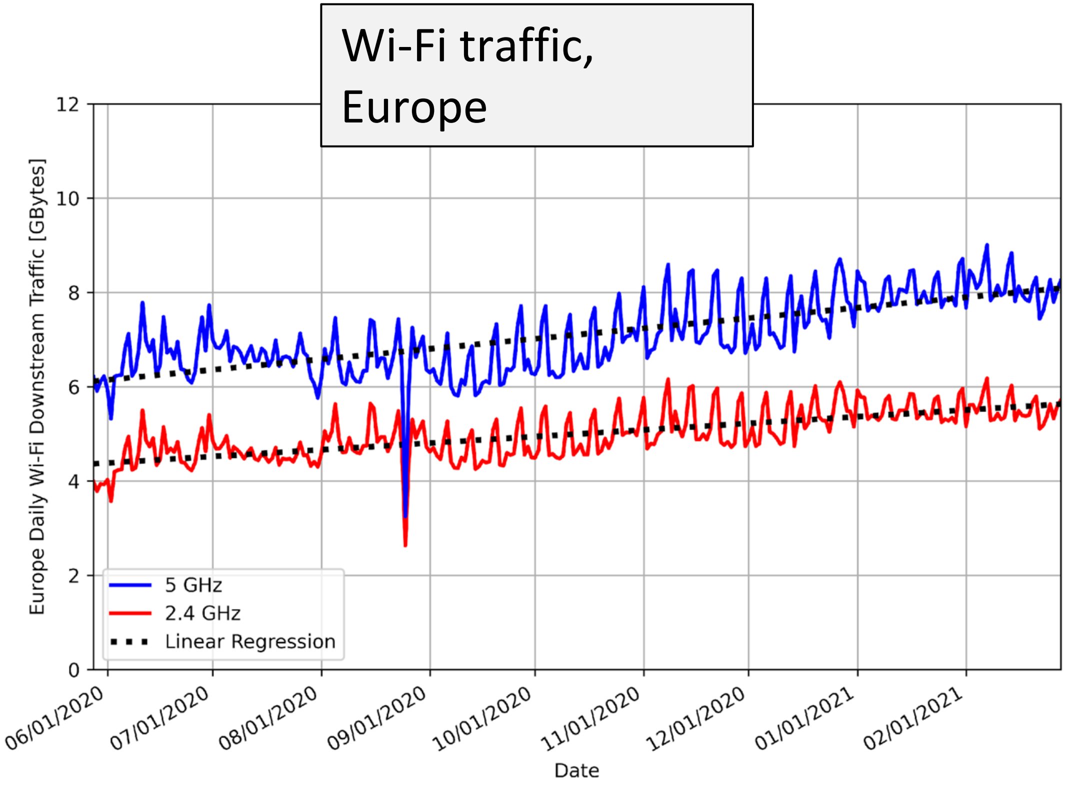

Wi-Fi traffic doubles about every 3 years, in both 5 GHz and 2.4 GHz

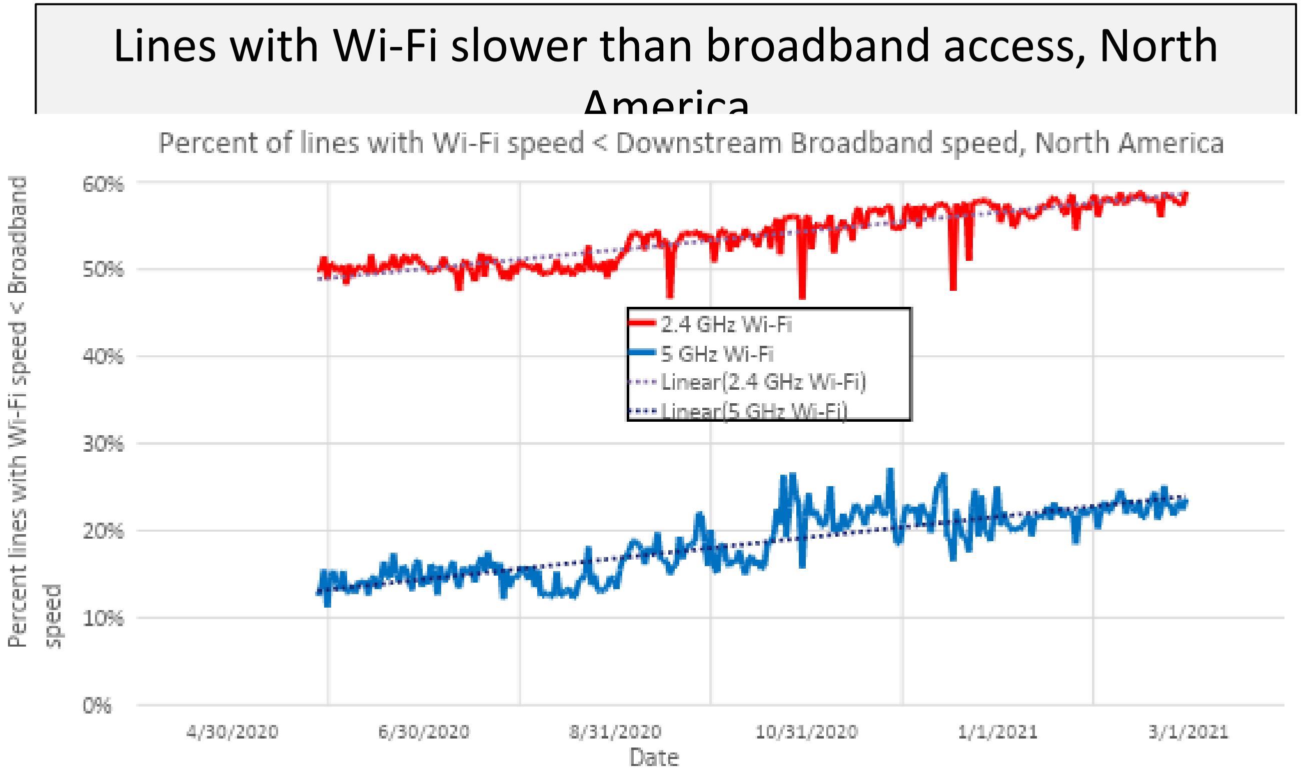

Wi-Fi slower than broadband	Annual increase
2.4 GHz Wi-Fi	13.0%
5 GHz Wi-Fi	14.4%

Internet access is increasingly limited by Wi-Fi

ASSIA

Annualized Percent Change in Wi-Fi Data, found by linear regression

Individual Wi-Fi parameter trends

	2.4 GHz	5 GHz
North America		
Wi-Fi traffic, downstream	4.4%	30.2%
Wi-Fi traffic, upstream	5.5%	22.5%
Wi-Fi interference	7.1%	18.3%
Wi-Fi congestion in busy hour	-3.6%	760.9%
Wi-Fi latency	13.4%	21.7%
Wi-Fi throughput / transmit rate	-7.3%	-18.8%
Europe		
Wi-Fi traffic, downstream	42.0%	42.0%
Wi-Fi traffic, upstream	14.4%	21.8%
Wi-Fi interference	3.7%	5.4%
Wi-Fi congestion in busy hour	64.0%	28.6%
Wi-Fi latency	29.9%	5.7%
Wi-Fi throughput / transmit rate	-8.7%	-8.4%

Overall Wi-Fi trend

Spectrum-need score = 0.2*ds traffic + 0.2*us traffic + 0.2*interference + 0.2*latency - 0.2*tput/tx rate

Continent, Wi-Fi Band	% Annual increase in spectrum-need score (linear regression)
North America, 2.4 GHz	13.2%
North America, 5 GHz	37.1%
Europe, 2.4 GHz	24.8%
Europe, 5 GHz	25.3%

Overall, Wi-Fi spectrum need rises about 25% annually

Wi-Fi **growth is rapid**

Maintenance of good QoE **is increasingly challenging**

 Need more spectrum 6 GHz will help

Full report: [ASSIA-DSA-Summit-Presentation-v7.8.pdf (dynamicspectrumalliance.org)]

Cross-layer Cooperation for Better Network Service

Mikhail Liubogoshchev

IITP RAS

Wireless Networks Lab

KHARKEVICH INSTITUTE FOR INFORMATION TRANSMISSION PROBLEMS
OF THE RUSSIAN ACADEMY OF SCIENCES

Cross-layer Cooperation for Better Network Service

Mikhail Liubogoshchev
Institute for Information Transmission Problems, Russian Academy of Sciences, Moscow, Russia
liubogoshchev@wireless.iitp.ru
https://wireless.iitp.ru

Application-side Challenges

Wireless Networks Lab
IITP RAS

Hard to perform fine-grained measurements required for latency sensitive applications: audio/videoconferencing, AR/VR, cloud gaming, etc.

"Single packet" measurements:

- Random protocol-induced delays in wireless networks
  - HARQ retransmission timeout in LTE (approx. 8 ms), channel access procedure in Wi-Fi
  - Present at any load
  - Non-linearly depend on load

"Multiple packets" measurements

- High wireless throughput requires high aggregation:
  - 1 ms scheduling/transmission period in LTE (0.5ms in current 5G deployments)
  - Dozens of kB transmissions in Wi-Fi
- High probability of an unsuccessful packet decoding:
  - Target PER=10%
  - For a dozen of packets in a single transmission, the retransmission is required with >70% probability
  - Retransmission of a single packet prevents all the following packets from forwarding up
- Packets often arrive at the receiver "simultaneously"

Need for Cross-Layer Communication

Wireless Networks Lab
IITP RAS

- End-nodes know "what they want":
  - Traffic type
  - Current session state (e.g., video buffering)
  - QoS requirements (incl. the requirements "at the moment")
- Network devices know "what they can offer":
  - The number and capabilities of clients
  - The state of (wireless) medium: channel throughput, PER, etc.
  - The applied policy and the typical processing delays
  - If some client will soon enter low-coverage area
- We can make the Internet much more efficient by exchanging this information:
  - Intelligent QoE-aware L2 management: provide service as much as and when it is required
  - Accurate and robust L7 (and probably L4) rate adaptation: do not send more data than the network can process
  - Enhanced end-user privacy: applications reveal metadata vs. ISPs deploy DPI

- Examples of such communication protocols:
  - ISO, ISO/IEC 23009-5:2017 Information technology – Dynamic adaptive streaming over HTTP(DASH) – Part 5: Server and network assisted DASH (SAND). ISO.2
  - I. F. Akyildiz, E. Khorov, A. Kiryanov, D. Kovkov, A. Krasilov, M. Liubogoshchev, D. Shmelkin, and S. Tang, "xStream: A new platform enabling communication between applications and the 5g network," in 2018 IEEE Globecom Workshops (GC Wkshps), Dec 2018, pp. 1–6.
  - Nikaein, Navid, Chia-Yu Chang, and Konstantinos Alexandris. "Mosaic5G: Agile and flexible service platforms for 5G research." ACM SIGCOMM Computer Communication Review 48.3 (2018): 29-34.
  - 3GPP TS 26.348 Northbound Application Programming Interface (API) for Multimedia Broadcast/Multicast Service (MBMS) at the xMB reference point; (Release 16). 3GPP, 2020
  - 3GPP TS 21.222:Common API Framework for 3GPP Northbound APIs (Release 15). 3GPP, 2021.

Packet delivery time as a tie-breaker for assessing Wi-Fi access points

Francois Michel

Packet delivery time as a tie-breaker for Wi-Fi access points

François Michel

francois.michel@uclouvain.be

Olivier Bonaventure

olivier.bonaventure@uclouvain.be

Diverse devices, diverse use-cases

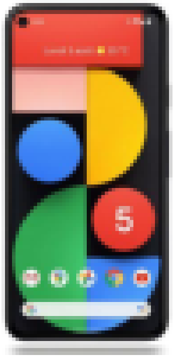

File transfer

Video on demand

Live streaming

Audio/video conference

Online gaming

Cloud gaming

2.4GHz/5GHz, private APs,
campus APs, ...

Diverse needs

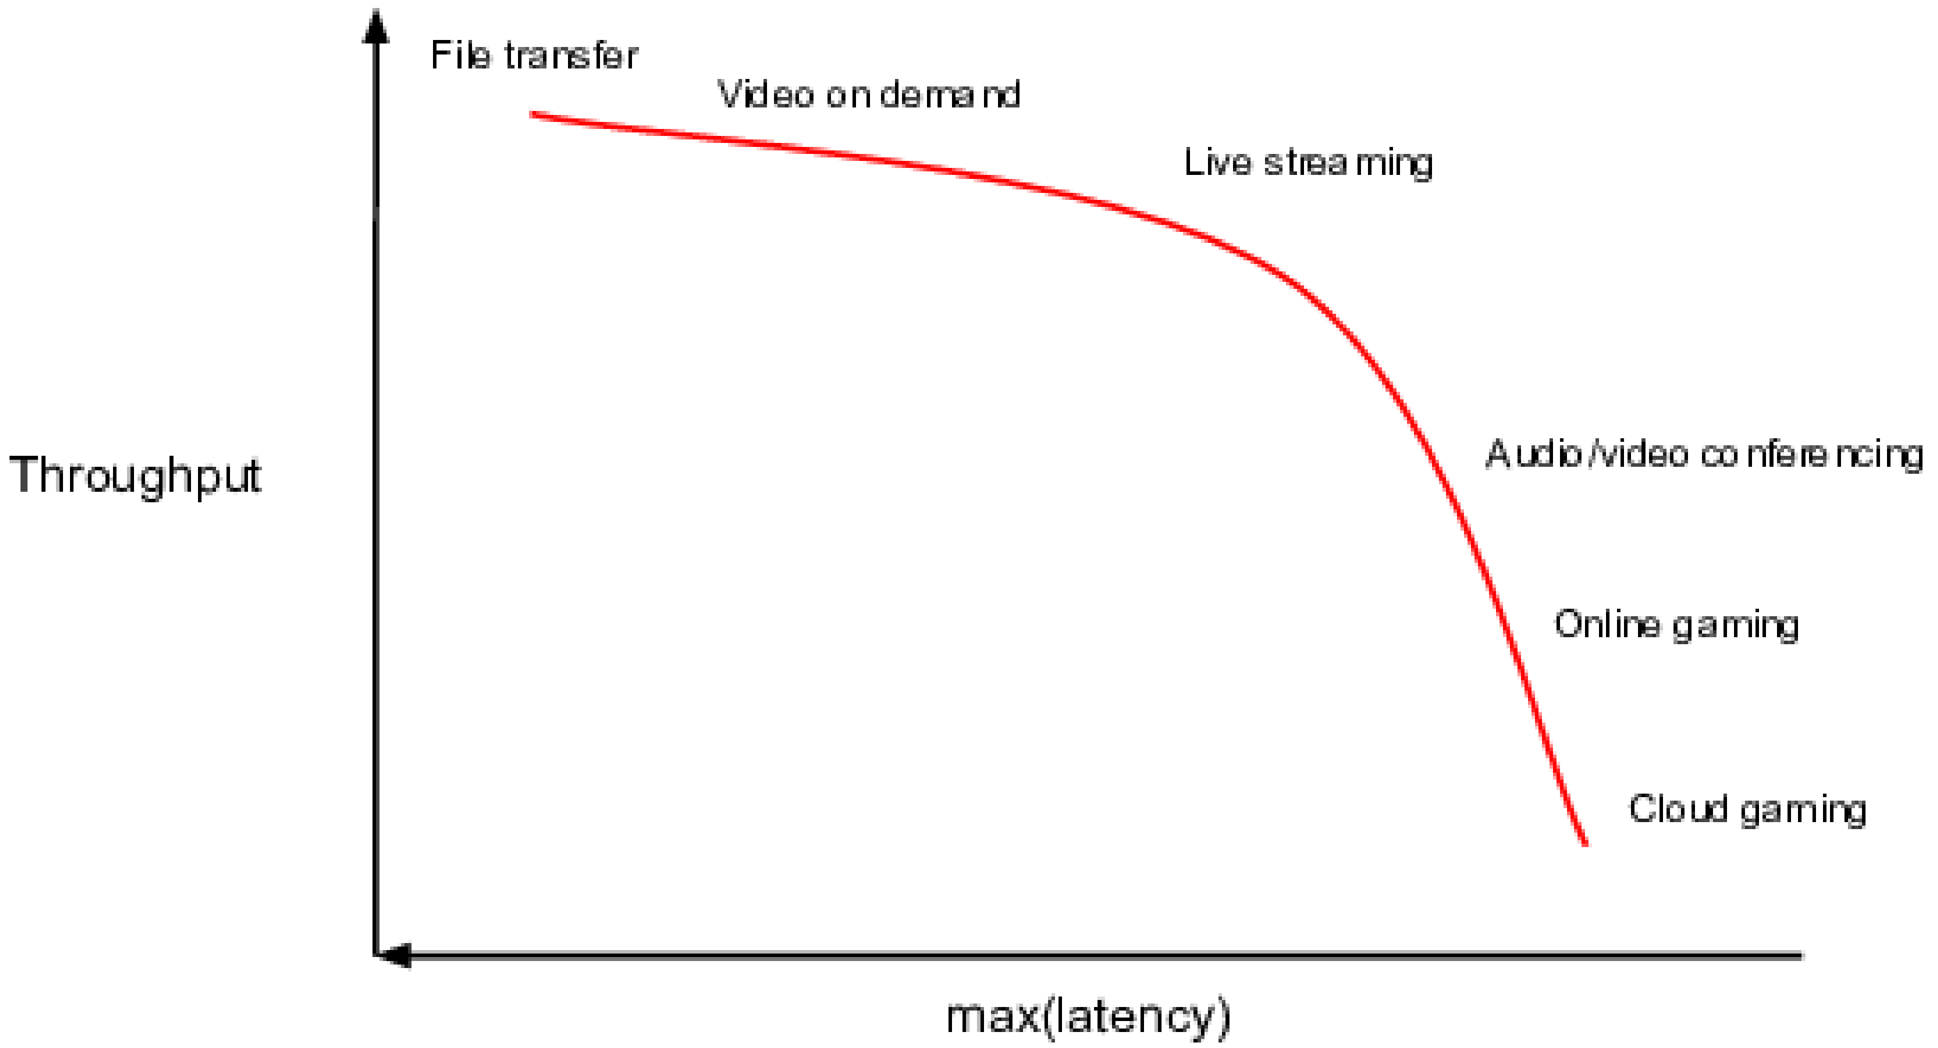

Sometimes several access points are available

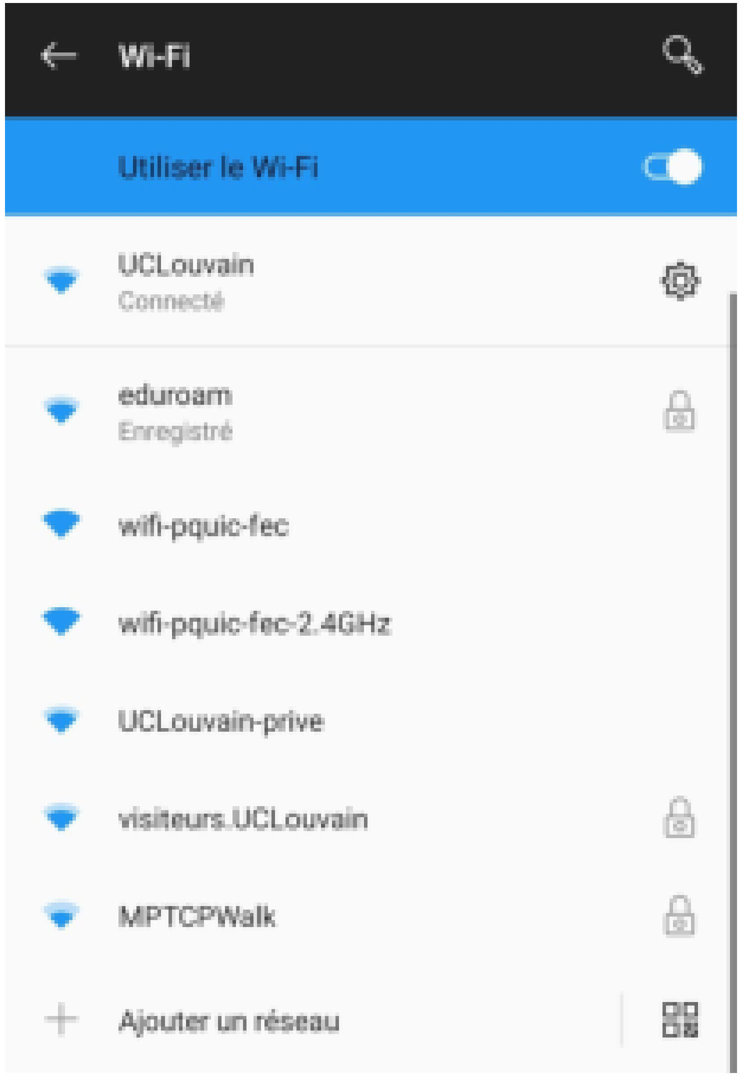

How to choose the AP ?

We have signal and theoretical BW informations but no information about the latency/jitter

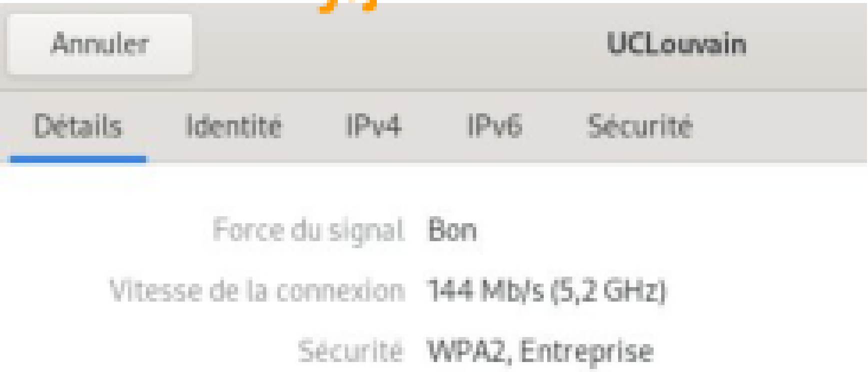

Latency can vary even at same distance to the AP

Same AP
Same distance
Same selected bitrate
(144Mb/s)
Different channels

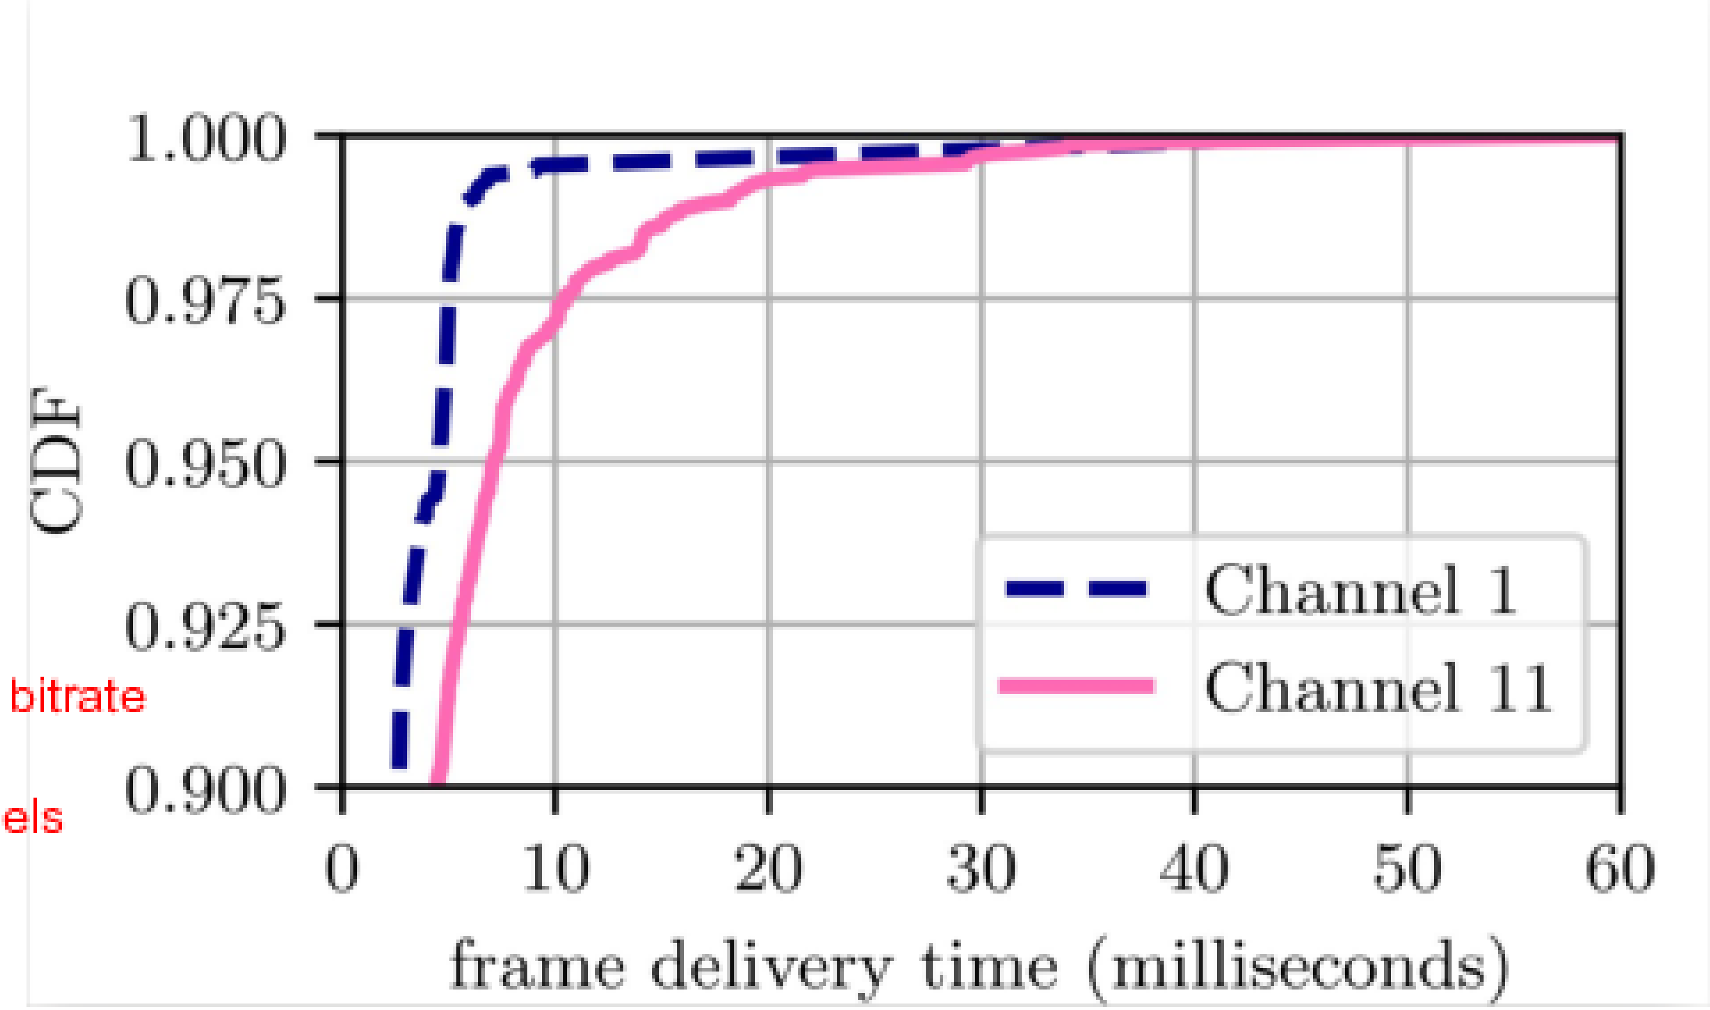

Help the user selecting the AP for their needs

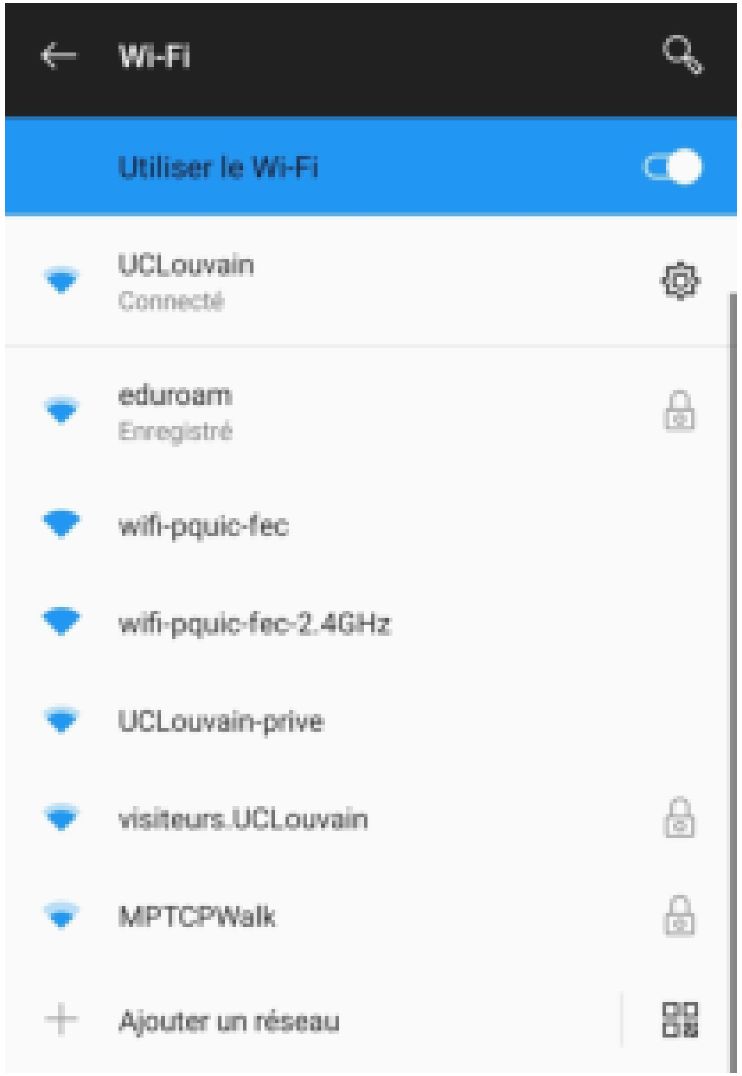

Help the user selecting the AP for their needs

Determined while using the channel

Could be cached by the device (or Apple/Google's servers ?) or announced by the AP

Adapt the application to the current conditions

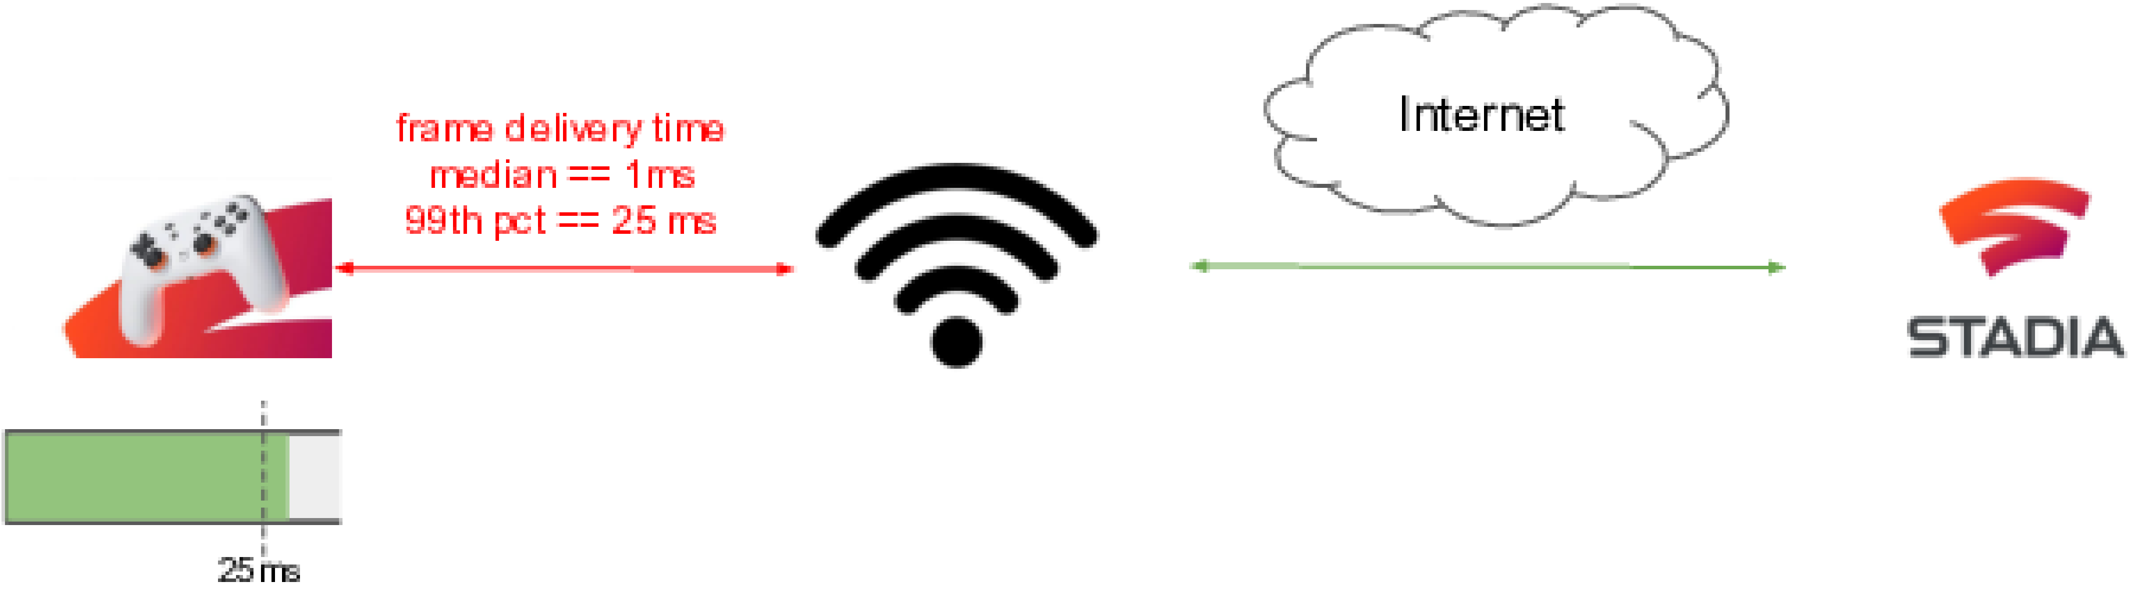

Start the game with a playback buffer of >25ms for less latency sensitive games. Discourage the use of other games ?

Exposing Wi-Fi frame delivery time statistics

Applications, users and maintainers may benefit from having information on the AP latency. Avg or percentiles of Wi-Fi frames delivery time could be exposed

- Help the user selecting the correct AP for their needs
- Identify the origin of the latency/jitter
- Adapt the application to the current conditions (use of playback buffers if needed)
- Help the application to choose the correct interface for initiating the transfer
- Help network maintainers reducing latency

Discussion

To be enqueued, please write ‘+q’ in the chat

To cancel being enqueued, please write ‘-q’ in the chat

The duration of each comment is limited by 60 seconds