Improving The Latency Of An MSO Network For Gaming And Real Time Applications

A Technical Paper prepared for SCTE•ISBE by

Colin Dearborn
Sr. Network Architect II
Shaw Cablesystems G.P.
2728 Hopewell Place. NE, Calgary, AB, T1Y 7J7
(403) 538-5297
colin.dearborn@sjrb.ca

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

List of Tables

Title **Page Number**

1. Introduction

The gaming industry continues to grow, eclipsing the movie entertainment industry in revenues. In 2019, the digital video game market garnered $120.1 billion (SuperData, 2020), while the worldwide box office earned $42.5 billion (Comscore, 2020). In a 2018 survey, it was estimated that more than 23 million Canadians played video games (The Entertainment Software Association of Canada, 2018). This means more than 60% of Canadians could be considered gamers. Gamers are demanding higher speeds and lower latencies from their providers to achieve competitive and enjoyable multiplayer gaming experiences. These customers are choosing their service provider on either real - or perceived - advantages that the network technologies offered by providers. Additionally, real time applications such as voice and video conferencing require low latency and jitter to perform with high reliability and quality. While it is impossible for any single service provider to control end to end latency of internet traffic beyond their own network, there are steps that can be taken to achieve the lowest possible latency. Multiple Service Operators (MSOs) are constantly developing and deploying the latest technologies and strategies to improve latency and jitter on their networks.

ISPs advertise technologies that reduce latency, and third party measurements that indicate lower latency than competitors. For example, TELUS, a Canadian telco ISP, has a website dedicated to explaining latency and claiming thatFibre To The Home (FTTH) is better for gaming.

Figure 1 – Sample ISP Gaming/Latency Website (TELUS, 2020)

Speed test sites also are starting to show latency scores as another tool for measuring ISP quality (Speedtest.net / Ookla, DSLReports speedtest, and bing.com's speedtest to name a few).

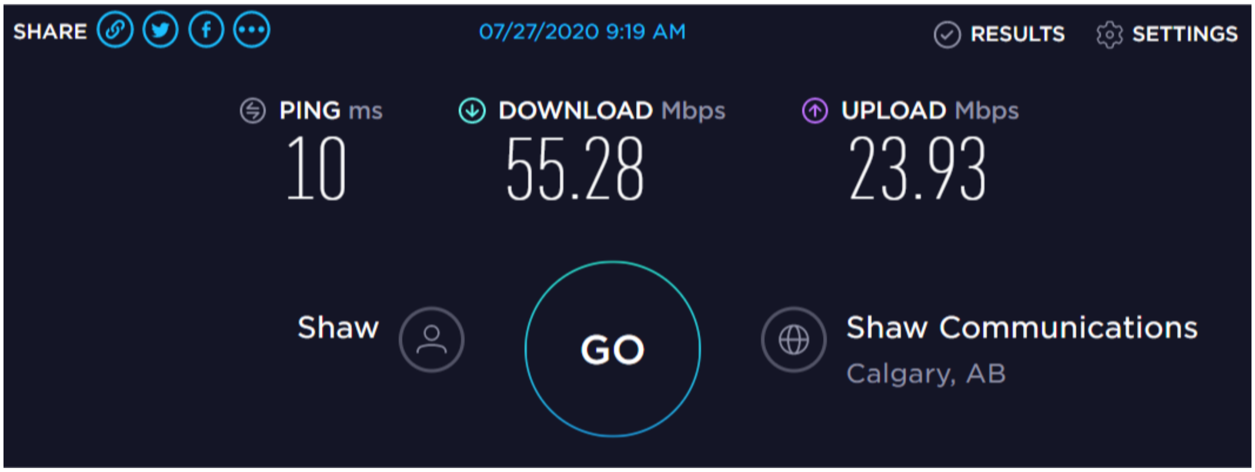

Figure 2 – Speedtest.net's Result Page Showing Latency Result (Ookla, 2020)

League of Legends publishes ISP connectivity scores based on latency, jitter and packetloss. Their North American site for this is located at https://lagreport.na.leagueoflegends.com/en/ (Figure 3 – League of Legends' Lag Report) and shows not only your current provider's score, but a "leaderboard" of providers in your area (Figure 4 – League of Legends' ISP Leaderboard for Calgary, AB), and many games have in-game displays of latency statistics.

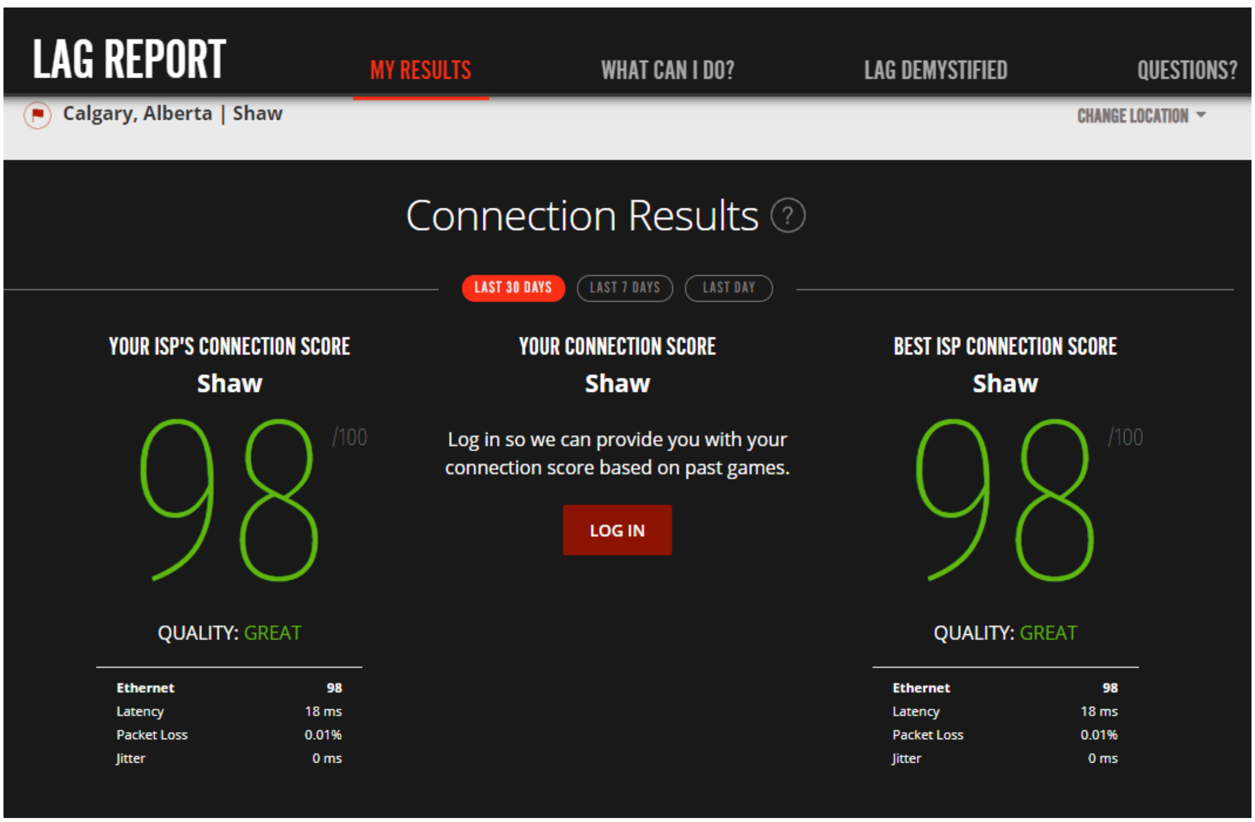

Figure 3 – League of Legends' Lag Report (Riot Games, 2020)

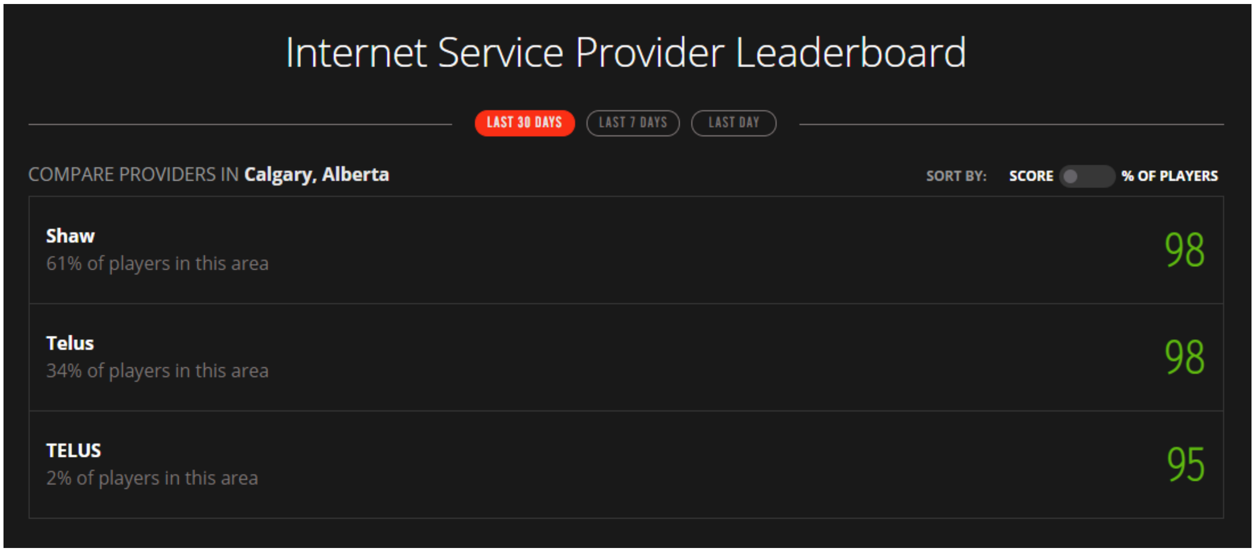

Figure 4 – League of Legends' ISP Leaderboard for Calgary, AB (Riot Games, 2020)

Gamers are very interested in latency and jitter because of its impact on targeting and player location within the game – which can mean the difference between winning and losing. Jitter, which is often unpedictable, can be especially challenging for game servers to compensate for.

While average latency is the amount of time (on average) a packet takes to traverse the network, it is measureable and relatively stable (on average). Jitter is the difference in latency between different packets and is the variability of latency. The more variability, the harder it is for a game to be accurate and consistent between players.

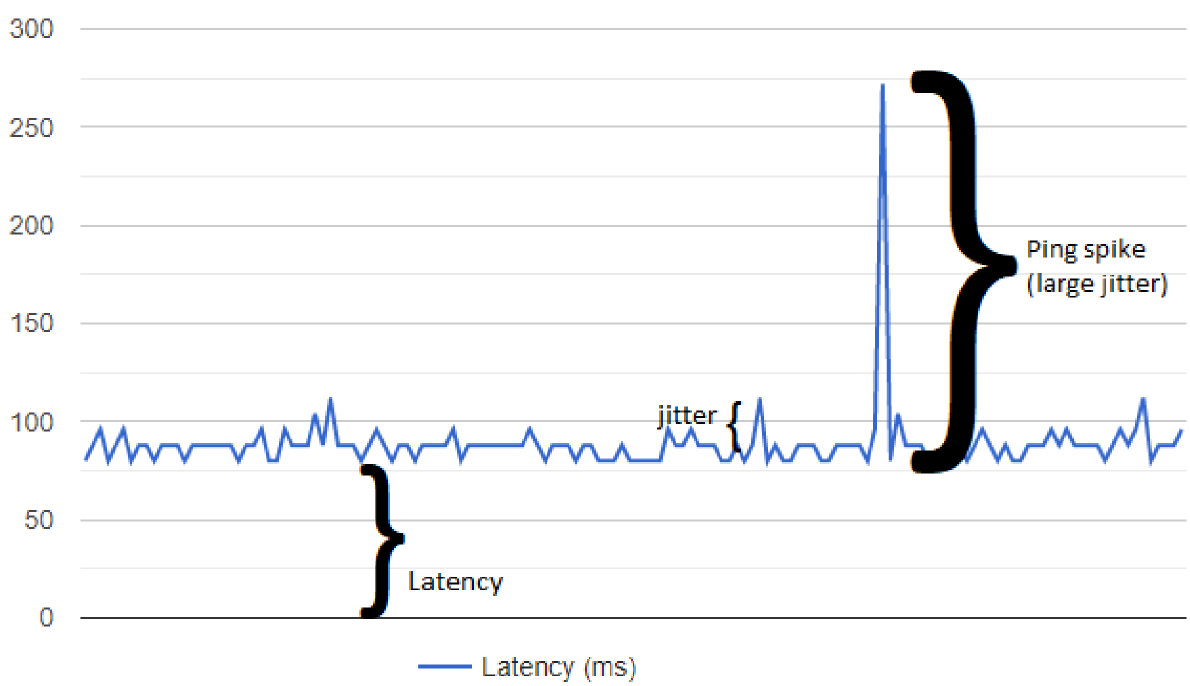

Figure 5 – Average Latency And Jitter

This paper will discuss current and future methods that help improve latency and jitter at several points in the network: in the home, in the Data Over Cable Service Interface Specification (DOCSIS) network, and in the core ISP network. Technologies such as WiFi 6 Orthogonal Frequency-Division Multiple Access (OFDMA), CableLabs' Dual Channel WiFi, Low Latency DOCSIS (LLD) and routing optimizations will be discussed. We will show the level of improvement available at each part of the network, and the perceived and actual performance gains that can be expected, as well as the promise of LLD in the DOCSIS network.

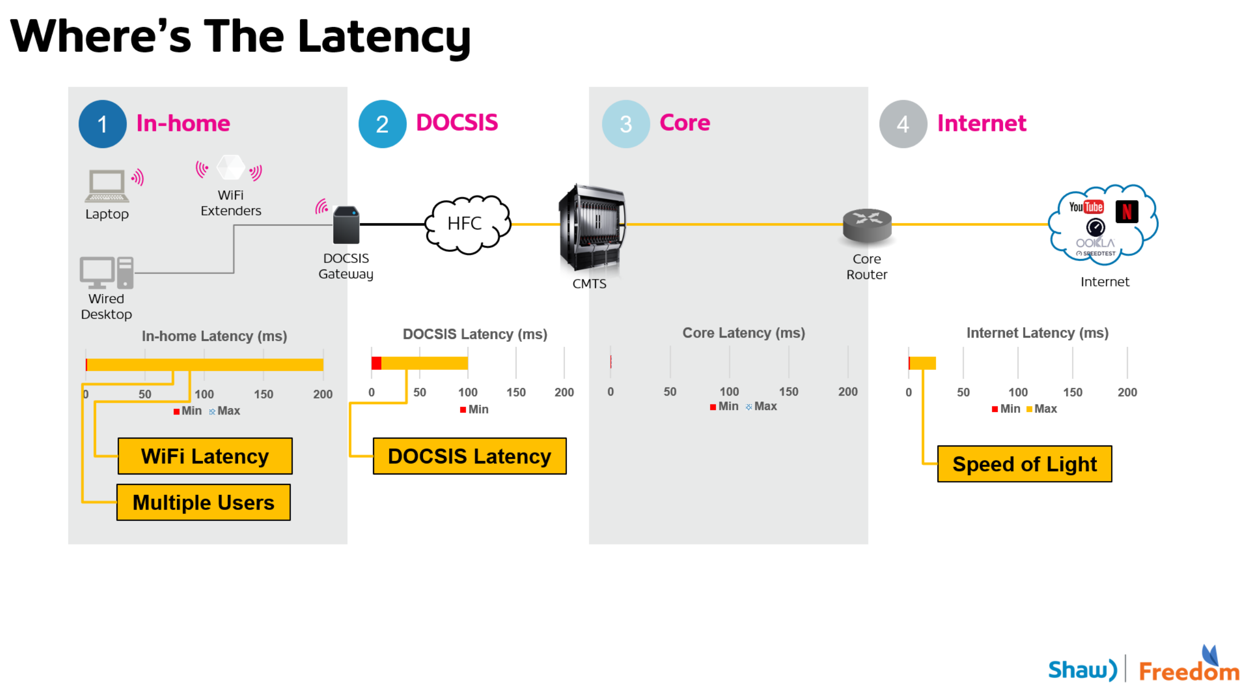

Figure 6 – Where's the Latency?

We will discuss three of the main areas shown in Figure 5; In-Home, the DOCSIS network, and the Core network, discussing some currently-available methods for reducing latency in these areas. Then, this paper will look at some technologies coming in the near future for each of these areas.

2. Current Methods For Improving Latency and/or Jitter

2.1. In-Home Technologies

Latency in the home is usually caused by WiFi usage. When wired, most home networks perform in the millisecond range, and quite frequently under a millisecond of latency is added traversing the home network. Over WiFi, conditions can vary wildly in different locations within the home, sometimes even in the same location. Yet customers continue to use WiFi as a primary method for connecting devices due to ease of use and the lack physical wiring in the home.

2.1.1. WiFi Multimedia

WiFi Multimedia (WMM) is a WiFi specification that was built to prioritize voice and video traffic over best effort internet traffic and also allows to de-prioritize background, non-latency sensitive traffic. When congestion happens WMM can ensure prioritized traffic will continue to get the bandwidth it requires. This also means it spends less time in the queue, avoiding latency and jitter penalties. We have seen applications such as Zoom employ WMM in the wild today, ensuring video calls perform well even in hostile WiFi environments. However, this is all at the cost of applications and traffic that don't use WMM, as their traffic will be penalized in WiFi congestion/contention scenarios. Any router supporting 802.11n or later should also have WMM enabled by default.

2.1.2. Summary

For applications that utilize WMM, and during congestion, the improvement can be tens to hundreds of milliseconds, as non-WMM WiFi queues can be quite large.

2.2. Access Network – DOCSIS

Technologies that can be employed in today's DOCSIS network include enabling Active Queue Management (AQM), adding orthogonal frequency-division multiplexing (OFDM) for DOCSIS 3.1 modems, and reducing DOCSIS Upstream Bandwidth Allocation Map (MAP) message intervals to 1ms or less. For the following tests, the test suite used defines jitter as IP Packet Delay Variation (IPDV) and is the average of all differences between each packet sent.

2.2.1. OFDM

Adding even a relatively small OFDM channel (64MHz) can allow for a small (0.5ms) end to end latency improvement. Based on Shaw's testing to date, this improvement doesn't seem to vary much with the size of the OFDM channel; it is only a function of whether the OFDM channel is present (and operational) or not.

Table 1 – Latency And Jitter In Relation To OFDM Channel Presence And Size

OFDM size (MHz)	Average Latency (ms)	Average Jitter (ms)
0 (not configured)	9.62627	2.19518
64	8.93366	2.22067
192	8.93708	2.21159

2.2.2. MAP Intervals

The MAP messages control when modems can request bandwidth, and also tells modems when they can transmit the data they have. Reducing the time between these messages can have an impact on latency as it gives more opportunities for modems to request bandwidth. Reducing the MAP interval to 1 ms or less is a requirement of LLD, but it is also something that is tunable today on some CCAP hardware.

Based on Shaw's testing to date, moving to < 1 ms MAP intervals can lower average latency a further 0.2ms, however it negatively impacts jitter by about about 0.2 ms as well.

Table 2 – Latency And Jitter when MAP Interval Is Changed

MAP Interval (ms)	Average Latency (ms)	Average Jitter (ms)
3.2	10.06525	1.93527
0.8	9.86053	2.20513

2.2.3. Upstream AQM

Active Queue Management helps protect any network queue (wherever it is enabled) from ping spikes and latency due to network queuing. Commonly referred to as "buffer bloat", this problem occurs when a bottleneck on the network receives data faster than it can send it out; it causes network latency to climb. Some of these queues can hold a large amount of data, allowing them to hold packets for very long times; sometimes into the range of multiple seconds. AQM uses the TCP mechanism of dropping packets to slow down TCP streams and actively monitors the time packets are spending in the queue. DOCSIS cable modems utilize the DOCSIS PIE algorithm to determine which packets are required to be dropped. Figure 7 shows the improvement with AQM enabled vs. disabled. As mentioned, since TCP dropping is the mechanism used to improve latency, when the upstream is congested with pure UDP traffic, the improvement is negligible. However, in a mixed or pure TCP traffic scenario, the latency can be improved dramatically, as seen in Table 3.

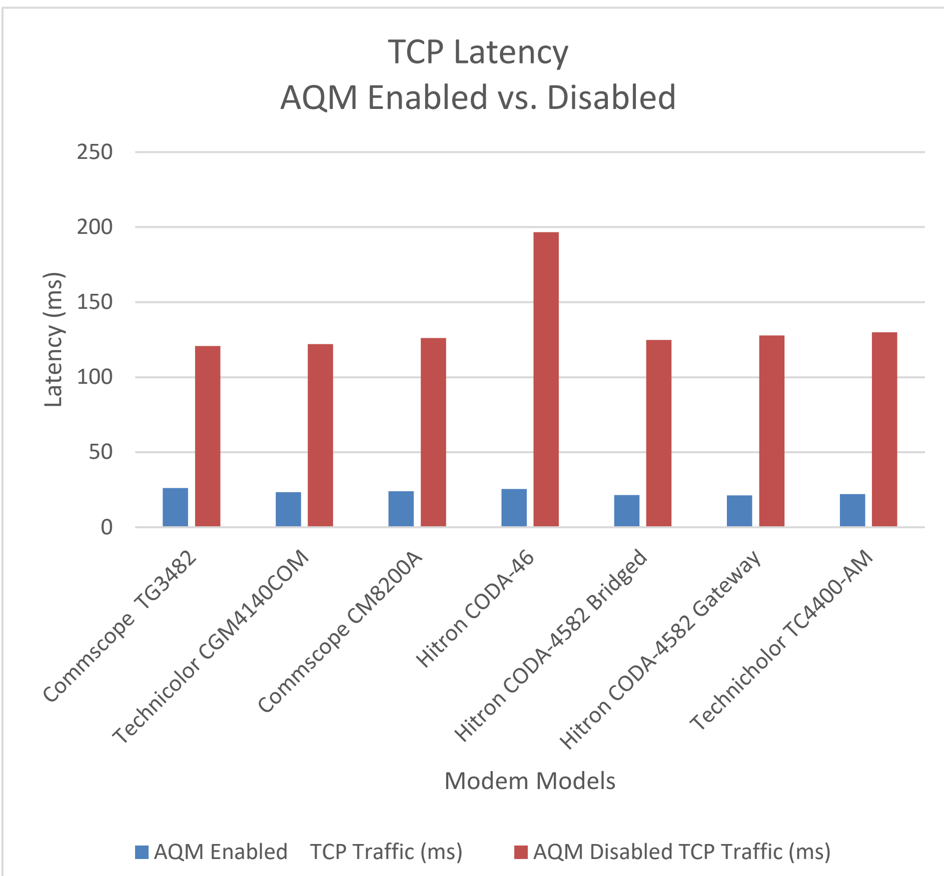

Figure 7 – TCP Latency When AQM Is Enabled Or Disabled

Table 3 – Latency When AQM Is Enabled/Disabled

	Commscope TG3482	Technicolor CGM4140COM	Commscope CM8200A	Hitron CODA-46	Hitron CODA-4582 Bridged	Hitron CODA-4582 Gateway	Technicholor TC4400-AM
AQM Enabled UDP Traffic (ms)	154.1	181.0	174.6	168.3	421.1	8.5	172.9
AQM Disabled UDP Traffic (ms)	201.8	216.2	213.2	194.1	22212.7	14.7	209.0
AQM Enabled TCP Traffic (ms)	26.2	23.3	24.0	25.6	21.5	21.3	22.2
AQM Disabled TCP Traffic (ms)	120.8	122.0	126.1	196.6	124.8	127.9	129.9

2.2.4. Downstream AQM

One of our CCAP vendors implements the Weighted Random Early Detection (WRED) AQM algorithm for downstream traffic. However, by default, it is turned off, so when customers saturate their downstream service flows with downstream traffic, latency increases due to bufferbloat on the service flow. We configure our downstream to use the WRED AQM with a minimum threshold of 40 ms, a maximum threshold of 60 ms, and drop probability of 50%. Figure 8 shows the difference in latency over a short (about 1 minute) download between WRED being enabled and disabled.

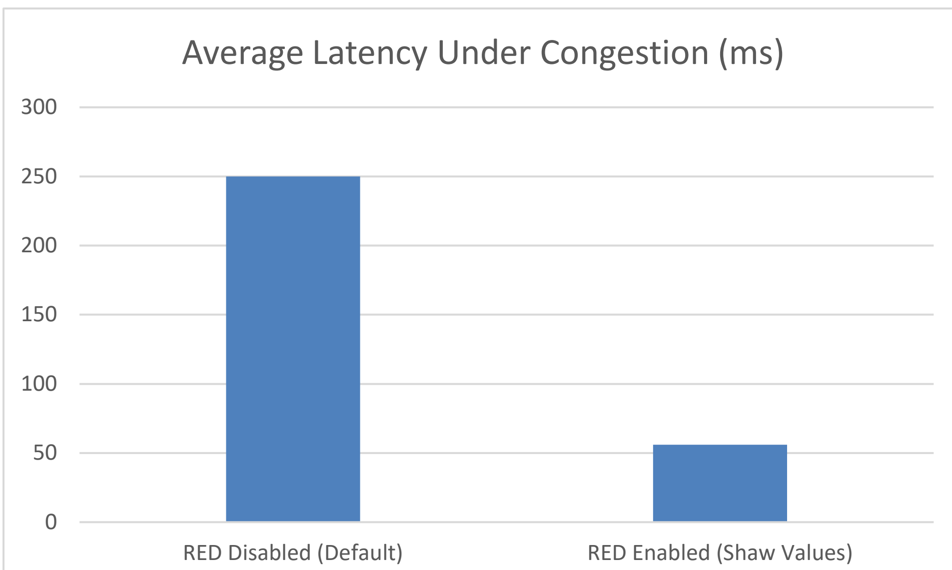

Figure 8 – Average Latency Under Congestion

2.2.5. Extra UDP Traffic

As more traffic bandwidth is injected through a set of service flows, latency and jitter appear to start to at first increase, and then decrease for a final improvement of over 3.5ms in latency, and over 1.6ms in jitter. This experiment added a second 75Mb/s UDP stream, and compared to the results to cases where no extra data streams were added.

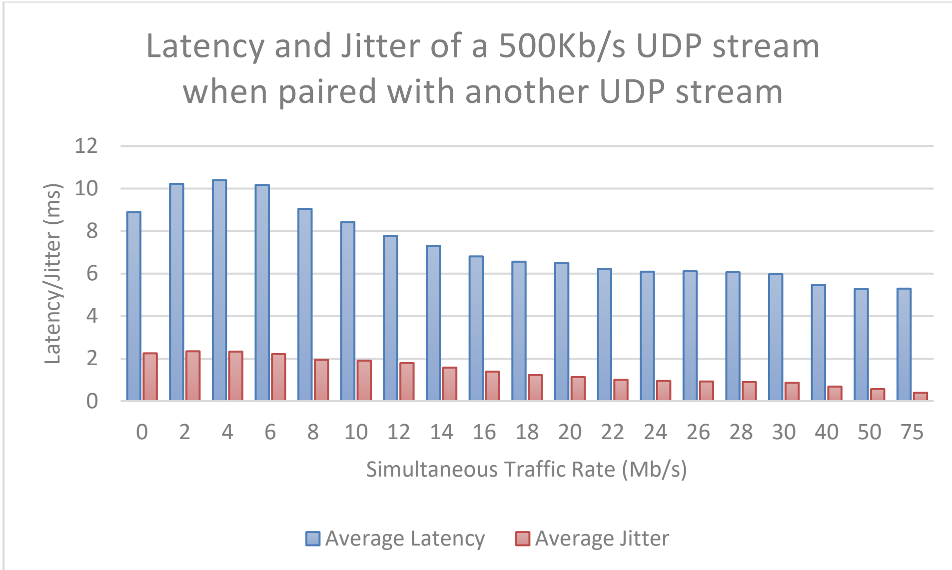

Figure 9 – Latency And Jitter In The Presence Of A Second UDP Stream

This is clearly not practical for a production deployment, as it locks up a lot of upstream bandwidth. It may also just be an artifact of the scheduler used by the specific CCAP chassis used in this test, but is nonethless interesting for its results in lowering both latency and improving jitter.

2.2.6. Summary

Utilizing DOCSIS 3.1 channels (OFDM) and reducing map intervals to 1 ms or less can bring latency down by about 0.7 ms combined, at the cost of some jitter (increase of an average jitter of about 0.2 ms). Adding AQM can also ensure that latency does not grow unreasonably due to congestion, and it can also save over 100 ms of latency during congested periods.

2.3. Core Network

2.3.1. Route and Metric Tuning

This method can have a significant impact on latency. Results can have a greater than 10 ms improvement over non-optimal routes. Mapping is done by using ICMP to find the latencies of different routes between different paths through the network. An example is shown in Figure 10.

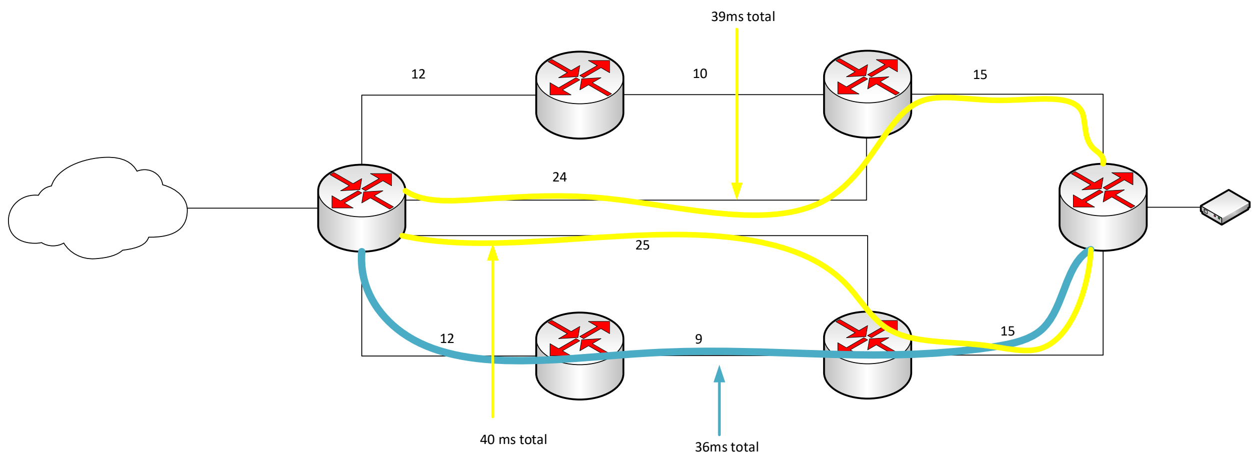

Figure 10 – Latency Of Multiple Paths Through A Network

The yellow path represents the network paths when no metric tuning has been done, and all links are equal cost. In this case, there are two paths. Depending on the load balancing algorithm the first router on the right uses, or could use one or the other path (39 or 40 ms), or could load balance packets between them for a single session resulting in a 1 ms jitter on the path.

When metrics are tuned to take the latency of each link into account, the resulting path is the blue path, giving the customer 36 ms of latency to traverse the network to the peer/transit point, and no jitter.

2.3.1. Disabling ECMP For Peering and Transit Links (Latency Stability)

Border Gateway Protocol (BGP) will sometimes advertise equal costs for multiple peering or transit points for a given network. In this case, latencies can vary wildly from session to session as traffic moves between these different peering or transit links. End users would complain that latency was "good" on one day, but was "bad" on another day as their gaming session would change the exit point from our network.

In Figure 11, with Equal Cost Multipath (ECMP) enabled, a customer could reach a server through two network paths in the local network, if BGP had the same metric at each edge router. If the total Return Trip Time (RTT) for the yellow path is noticeably different than the green path, the customer will notice that for different gaming sessions to this server, he could see different latencies. Allowing the IGP metrics from metric tuning to be the tiebreaker for routing means that the network path (in this case) would always be the green path due to the lower latency internally. Depending on the latency to the server outside the network, the total RTT could actually be higher through the green path, but the latency being stable across sessions was more acceptable to customers, even if it was higher for specific servers.

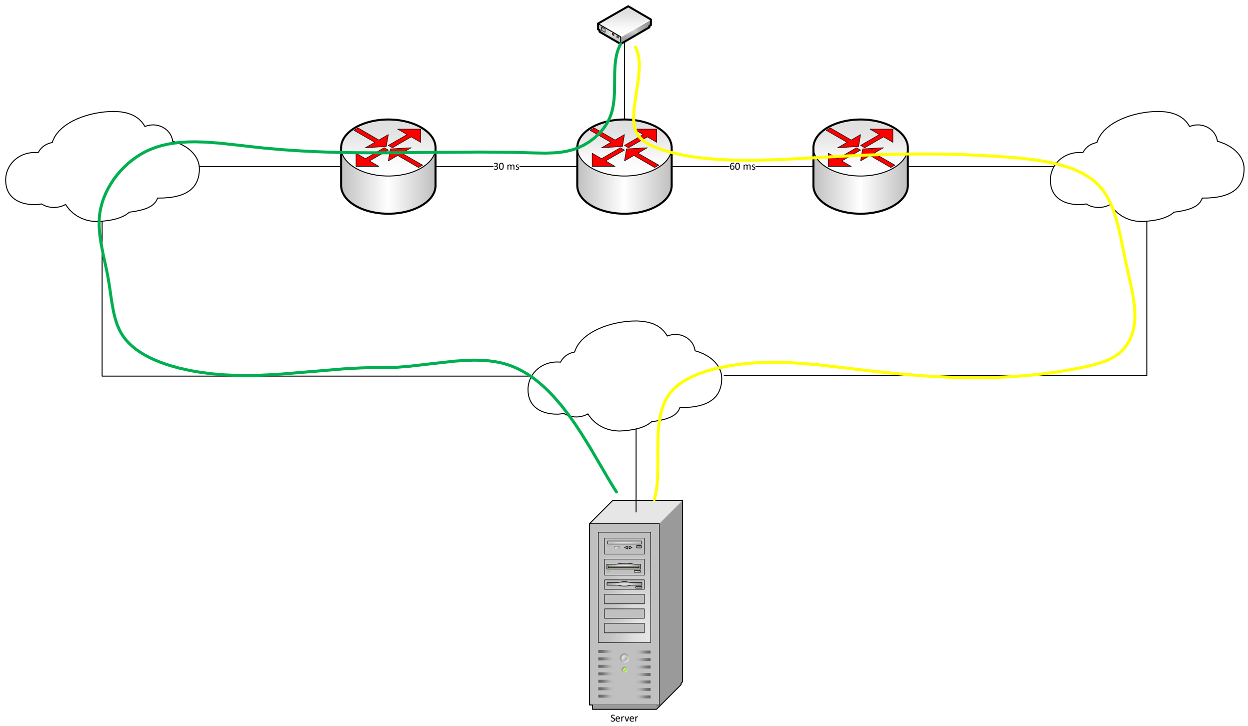

Figure 11 – Network Exit Points

2.3.2. Summary

These methods in the Core Network can improve latencies from 1 to 50 ms or more, depending on size and complexity of the network.

3. Upcoming Technologies And Methods For Improving Latency And Jitter

3.1. In-Home Technologies

WiFi 6 includes Orthogonal Frequency-Division Multiple Access (OFDMA). This feature improves bandwidth sharing with large numbers of clients to prevent clients talking over each other, thereby reducing latency and jitter in busier networks. In order to take advantage of OFDMA, however, all clients must be WiFi 6 enabled with OFDMA capability.

Dual Channel WiFi™ is a technology developed by CableLabs to provide a separate (or multiple separate) downstream channels in combination with a "legacy" channel. This allows large downloads, and streaming to be shunted to this separate channel, keeping the legacy channel open for upstream and legacy clients, reducing latency and jitter by reducing contention and congestion when clients are requesting large amounts of traffic.

3.2. Access Network – DOCSIS

3.2.1. Low Latency DOCSIS

The two technologies in LLD that will contribute the most to latency and jitter improvements are dual-queue and the Proactive Grant Service (PGS). LLD AQM will also ensure that latency is held down during congestion and Coupled AQM will ensure that bandwidth is available to both the low latency queue as well as the classic queue during congestion

3.2.1.1. Dual-Queue

Dual-queue provides a separate set of DOCSIS service flows, so time sensitive traffic doesn't have to compete with non time sensitive traffic within a single queue. This means that even when there is a large amount of traffic on the legacy service flow, it will not affect the queuing behaviour of the low latency service flow.

Dual-queue also allows the DOCSIS scheduler to schedule latency sensitive traffic differently to provide low latency and low jitter.

So, why don’t we have all traffic low latency? Queuing delays are actually caused by some applications sending traffic in a manner that results in a build up of packets in the network. The biggest source of this type of traffic is any application that uses TCP today. Those applications need the network to provide a deep buffer to absorb bursts of traffic, and don’t perform well if the network has shallow buffers. But games and latency sensitive applications typically deal with small packet loss better than latency variation, and large buffers allow for large variation or spikes in latency, as well as large latency if the buffers are kept full for a long time. Having small buffers just for latency sensitive applications will ensure packets are not held for too long. The dual-queue feature of Low Latency DOCSIS equipment gives both types of traffic the appropriate buffer for their needs: a shallow buffer for Low Latency traffic, and a deep buffer for Classic traffic.

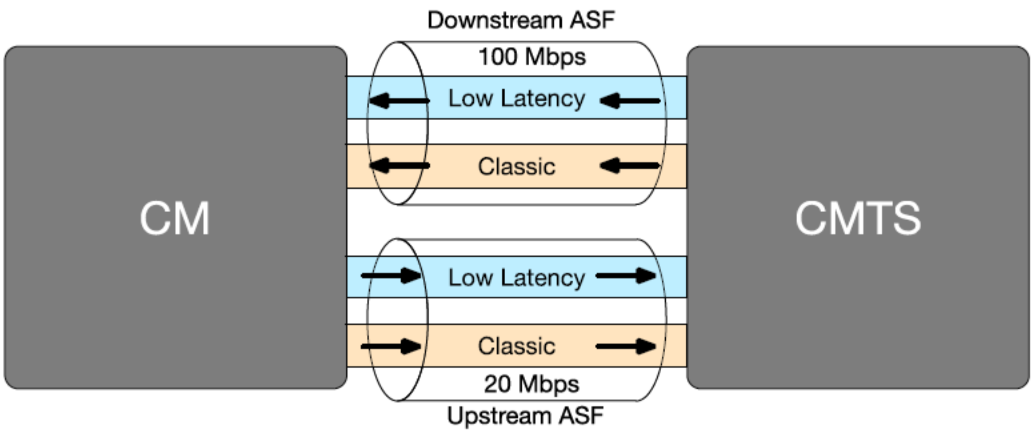

Figure 12 – Dual Queue Service Flows (White, Sundaresan, & Briscoe, 2019)

This does mean that we will need to differentiate the traffic via some means. We do have access to DOCSIS packet classifiers, but these would be hard to manage as it would require classification of every type of traffic and every different game to be put into the low-latency service flow. LLD will use several methods to classify traffic into the low latency service flow; a Differentiated Services (DiffServ) value or Explicit Congestion Notification (ECN). CableLabs has proposed a DiffServ value of 0x2A be defined as Non-Queue-Building (NQB), and LLD would use this value, as well as ECN values. As ECN is supported, Low-Latency Low-Loss Scalable throughput (L4S) will also be supported for applications that need both high bandwidth and low latency. CableLabs has been working with game developers to inform them how to mark their packets through gaming conferences and through the website pingspikeskill.com.

3.2.1.2. Proactive Grant Service

Proactive Grant Service (PGS) is a new DOCSIS upstream scheduling service included in LLD. PGS is like the Unsolicited Grant Service (UGS) used primarly for voice services, but where UGS is very static with the size of packets and the number of grants given, PGS is much more dynamic, allowing for different size packets and different amounts of bandwidth. PGS helps reduce the grant-delay cycle inherent in the DOCSIS protocol (Figure 13 – DOCSIS Grant Delay Cycle).

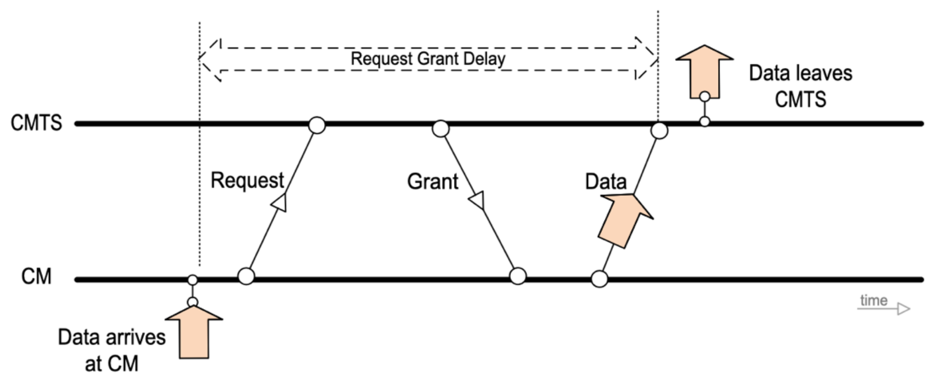

Figure 13 – DOCSIS Grant Delay Cycle (White, Sundaresan, & Briscoe, 2019)

As the name indicates, PGS proactively tries to predict and give grants for the modem to transmit data based on past behaviour, so that the modem doesn't have to request the grants (effectively removing the grant delay cycle latency); but the downside is that if the modem doesn't have any data to transmit, that bandwidth is now wasted.

3.2.1.3. Upstream Scheduling Improvements

Typically it takes more time than a single MAP interval to process a request for bandwidth, so reducing the MAP interval alone does not provide a huge boost to latency reduction. If the turnaround for processing the request can be reduced along with the MAP interval reduction, this should allow for a good reduction of latency and jitter even without PGS.

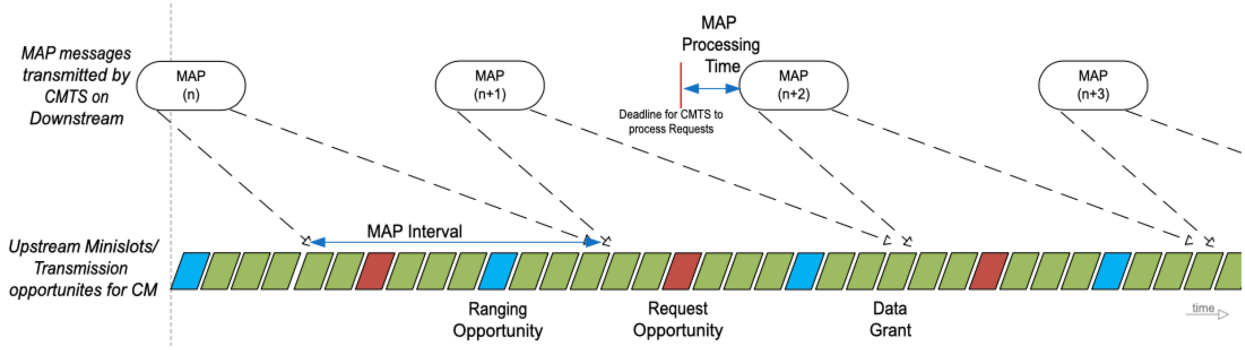

Figure 14 – MAP Interval and MAP Processing Time (White, Sundaresan, & Briscoe, 2019)

3.2.1.4. Coupled AQM

Coupled AQM ensures that there is shared bandwidth for both the low latency service flow and the classic service flow within the aggregate service flow during times of queuing. The total bandwidth consumed by the coupled service flows is limited by the AQM configuration settings. This feature ensures that the low latency service flow does not starve out the classic service flow under this condition, and aims to ensure that classic TCP sessions and future L4S TCP sessions all receive a fair bandwidth allocation.

3.2.1.5. Summary

Total potential improvements with LLD will run in the range of 5-10 ms of latency (more in cases with congestion/queuing), as LLD can deliver DOCSIS latency of ~1ms (White, Sundaresan, & Briscoe, 2019) and jitter will most certainly be improved as well with both PGS and upstream scheduling improvements. This will significantly close the “latency gap” that last mile fibre products use in their marketing, even though 5-10 ms may seem like a small improvement overall.

3.3. Core Network

Currently the optimization of routing configurations for specific destinations is very manual and labour intensive, and it has only been done a handful of times when teams have the cycles to perform the required tuning. Going forward we are considering launching a project to keep these up to date and ensure they don’t cause problems as our network evolves and changes. We are also looking at enabling partial automation for these optimisations, but that is further in the future.

CIN networks can add another level of complexity to these calculations especially for more remote sites due to the diverse path requirements and distances involved. Even within a single metropolitan area, fiber paths and switching delays can contribute to a slight difference in latency between paths. Paying attention to the different latencies between each path will continue to be an ongoing process.

Cloud Gaming, Augmented Reality (AR) and Virtual Reality (VR) applications will start to require lower latency, and since the networks between users and where servers are hosted can make up a large portion of the latency budget, edge computing or edge cloud, where servers are moved into ISP facilities, or very close to them, can also play a part in the ability for an ISP to lower latency for their customers.

4. Conclusion

Carefully focusing on the parts of the network over which you have control provides ample opportunity for reducing inefficiencies. Enabling AQM and having devices and applications that utlise WMM can yield many latency benefits. In addition, wiring as many devices as possible in the home can help keep latency down in times of congestion. Enabling DOCSIS 3.1 downstream channels (OFDM) and reducing upstream MAP intervals to 1 millisecond or less can bring DOCSIS access latency down a small but measureable amount. Finally, ensuring your core network is optimized for latency via continual measuring and adjusting of IGP metrics can give you an edge over even FTTH.

Some technologies look impressive, but need to be framed within their use case. For example, AQM, WMM, WiFi6 OFDMA, and Dual Channel WiFi all perform impressively during congestion or on busy networks, but during normal operations may not yield much improvement for latency or jitter. That doesn't mean they aren't important, as even home networks can have short bursts of extreme traffic or a large number of devices requiring network resources at the same time. Some of the technologies may seem like they don't contribute much to a lower RTT on the network, such as 1 ms MAP Intervals, OFDM, or even LLD, but these can lower the difference of latencies seen between FTTH and DOCSIS, and in the case of LLD even eliminate that difference altogether.

In the future WiFi 6 within the home will further hold latency down as more and more devices are added to home networks. LLD may bring latency down to as low as 1 ms in the access network. Further optimisations on the Core network can ensure latency on a DOCSIS network is as low as even the best FTTH network.

Figure 15 below shows the rough amount of latency that can be shaved off the RTT with each technology or optimization. Some are additive, some (such as LDD, which includes MAP Intervals and AQM) includes other technologies listed. Some may seem small, but work under most conditions (MAP Intervals/OFDM) while others (marked with asterisks) seem like large gains, but only work under some conditions (AQM). LLD has technologies that will improve both areas, as seen with its two entries. These will all work together to improve latency and jitter on the network keeping latency sensitive applications and gamers satisfied and happy.

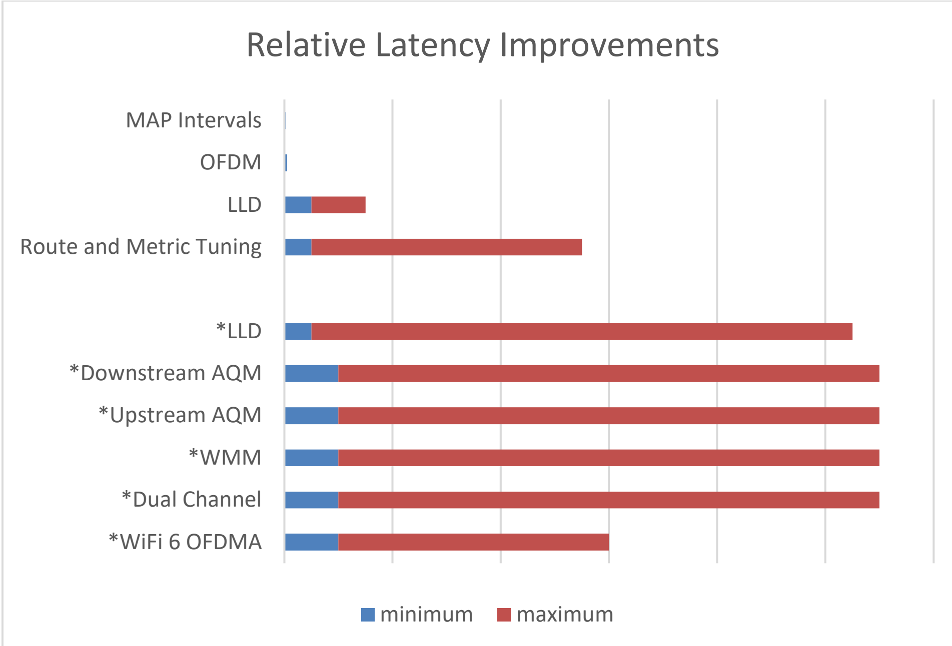

Figure 15 – Relative Latency Improvements

Abbreviations

AQM	Active Queue Management
AR	Augmented Reality
BGP	Border Gateway Protocol
DOCSIS	Data Over Cable Service Interface Specification
FTTH	Fiber To The Home
IGP	Interior Gateway Protocol
ISP	Internet Service Provider
LLD	Low Latency DOCSIS
MAP	Upstream Bandwidth Allocation Map
OFDM	Orthogonal Frequency-Division Multiplexing
OFDMA	Orthogonal Frequency-Division Multiple Access
PGS	Proactive Grant Service
RTT	Round Trip Time
WRED	Weighted Random Early Detection
UGS	Unsolicited Grant Service
VR	Virtual Reality
WMM	WiFi Multimedia

Bibliography & References

Comscore. (2020, 01 10). *Comscore Reports Highest Ever Worldwide Box Office*. Retrieved from Comscore: https://www.comscore.com/Insights/Press-Releases/2020/1/Comscore-Reports-Highest-Ever-Worldwide-Box-Office

Ookla. (2020, 07 27). *Speedtest by Ookla results*. Retrieved from Speedtest by Ookla : https://www.speedtest.net/result/9820878391

Riot Games. (2020, 08). *League of Legends Lag Report*. Retrieved from League of Legends: https://lagreport.na.leagueoflegends.com/en/

SuperData. (2020). *2019 Year In Review.* Retrieved from SuperData Research: https://www.superdataresearch.com/2019-year-in-review/

TELUS. (2020, 08 01). *Best Internet for Gaming - Fibre Internet | TELUS*. Retrieved from TELUS: https://www.telus.com/en/internet/best-internet-for-gaming

The Entertainment Software Association of Canada. (2018). *Essential Facts About The Canadian Video Game Industry.* Retrieved from The Entertainment Software Association of Canada: http://theesa.ca/wp-content/uploads/2018/10/ESAC18_BookletEN.pdf

White, G., Sundaresan, K., & Briscoe, B. (2019). *Low Latency DOCSIS: Overview And Performance Characteristics.* SCTE.

