The Quest to Send 4K Video Over Wi-Fi Networks

The VoW Factor – (*that's **V**ideo **o**ver **W**i-Fi*)

A Technical Paper prepared for the Society of Cable Telecommunications Engineers
By

Charles Cheevers
CTO Customer Premises Solutions
ARRIS Group Inc
Charles.Cheevers@arrisi.com

Ian Wheelock – Ian.Wheelock@arris.com
Carol Ansley – Carol.Ansley@arrisi.com
Kurt Lumbatis – Kurt.Lumbatis@arrisi.com

Contents

Overview

The Quest to send 4K Video content over Wi-Fi – the VoW era

The evolution to video over Wi-Fi (as the authors like to call the VoW era) is probably one of the most significant changes in the cable industry since we went from analog video distribution to all-digital on QAM technology.

The focus of this paper is **not** on the deep science of the technology of Wi-Fi and the abilities to send packets over increasingly congested Wi-Fi networks – but rather to mix in the softer elements of what consumers really need of their Wi-Fi network in the home and how to make the Wi-Fi delivery network **deterministic** in all scenarios of interference, over subscription and congestion.

It assumes we are leveraging the latest and greatest in 802.11 wireless MAC and PHY technologies as the basic fundamental building blocks and also that we have currently relatively clear 5GHz spectrum to begin our journey of VoW but instead focuses on the key areas that work to increase the likelihood that the residential consumer will always get the desired performance, quality of service and determinism in behavior within their own home – and specific user, device and service behavior and preferences.

The thesis for this paper is that Wi-Fi delivery can fail – and fail hugely – yet there are ways to leverage the dynamics of the consumer home and service consumption preferences to ensure that the Quality of Experience of Wi-Fi delivered video is always to a consumer's desired expectation.

.

Figure 1 - 4K over Wi-Fi - the cable industry's own moon landing

The delivery of 4K resolution video is the cable industry's 'Moon Landing' target. We are already on the quality border line for Wi-Fi delivery of Over-The-Top (OTT) video streaming solutions – with 53%[1] of users in North America reporting streaming issues and 8%[2] of them having regular issues. The 2.4 GHz spectrum is full, particularly in dense Multiple Dwelling Unit (MDU) environments and we are seeing regular problems with consumers trying to stream OTT content at modest bandwidths. In almost all of these issues, the problem is not with their access bandwidth tier – it's with the congestion in the 2.4 GHz Wi-Fi domain (Figure 2). We are just about able to reliably deliver these relatively low bitrate video streams, yet our thirst for High Definition (HD) quality streaming continues to grow.

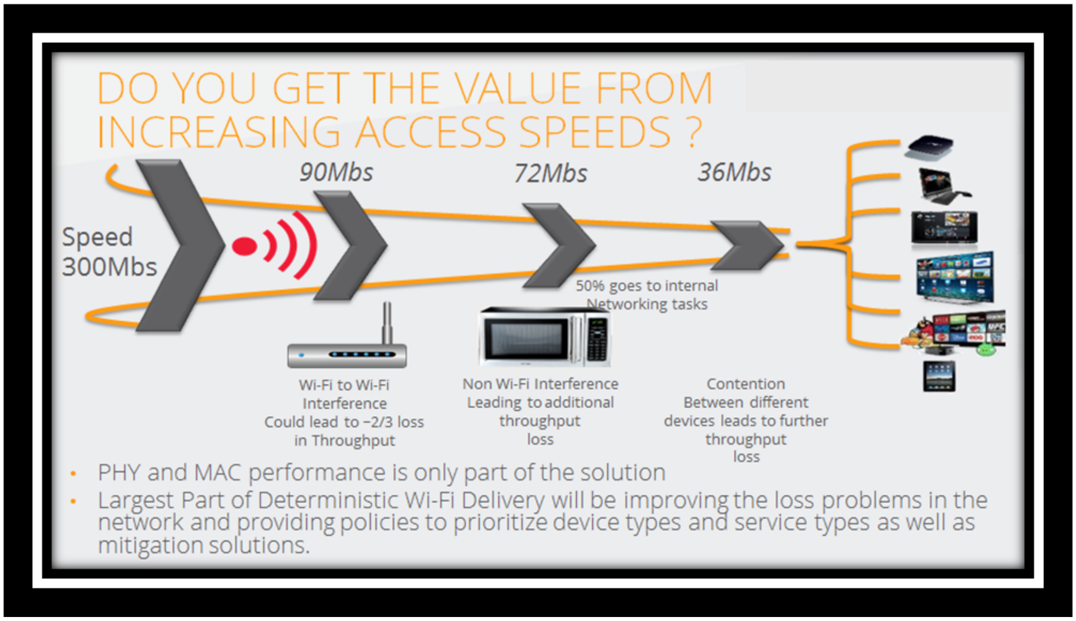

Figure 2 - High Access SLA speeds cut back by 2.4 Ghz Wi-Fi LAN delivery abilities

Enter 4K video and the network bandwidth demands increase by 3-6 times the current HD 5 Mbps delivery rate. With 4kp60 and 10 bit color gamut now being the norm for aspiring 4K content delivery – bandwidth requirements can be pushed to 30 Mbps even with the latest High Efficiency Video Codec (HEVC). We need to get ahead of this increased bandwidth demand and lay the foundation for Quality of Experience-driven service delivery in the home.

[1] ARRIS Consumer Entertainment Index
[2] ARRIS Consumer Entertainment Index

We also know that we are about to move our managed video delivery to IP video and concurrently shift Wi-Fi to become the primary connection medium for new set-top boxes (STBs) enabling self-install for video services.

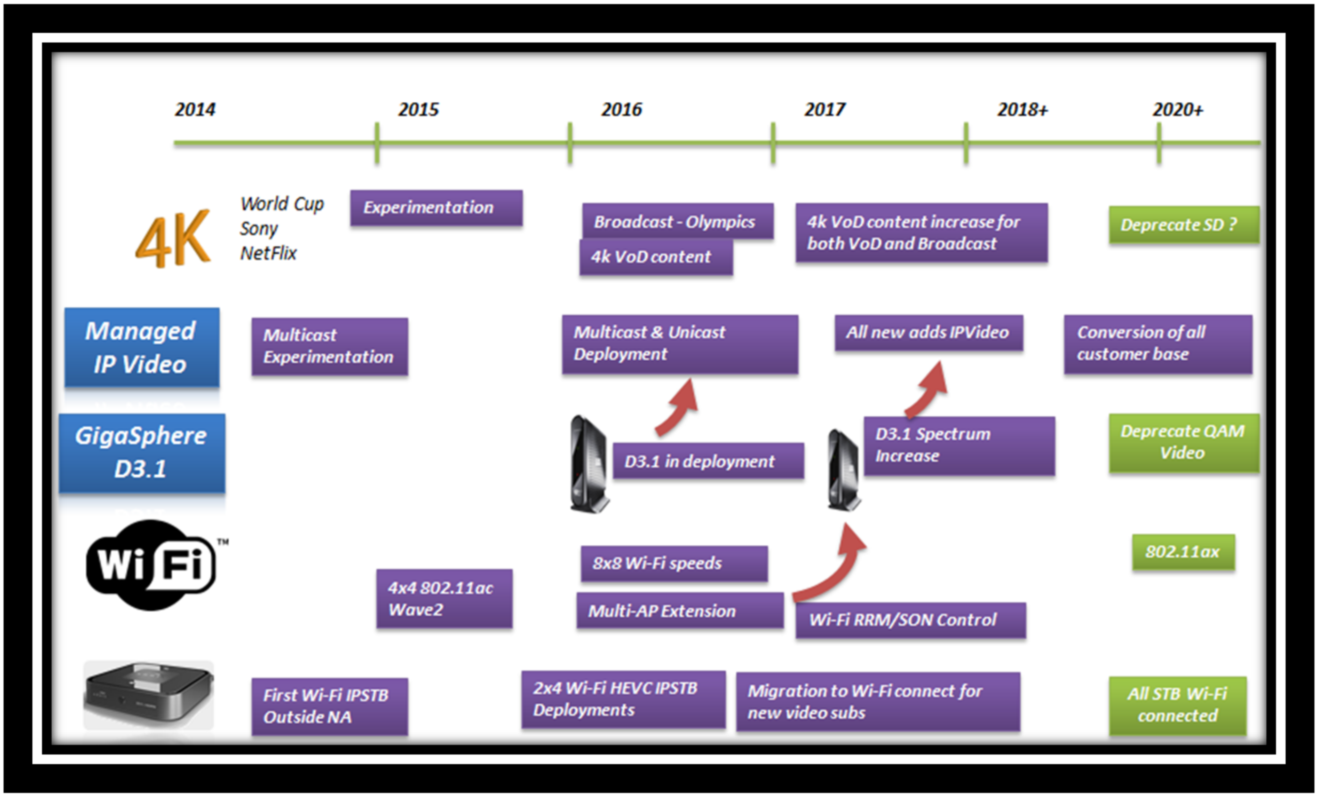

Figure 3 - The EcoSytem and sample timing of moving to All IP VoW

The window for this transition seems to be from now to the rise of DOCSIS 3.1-enabled networks in 2016. DOCSIS 3.1 will be the platform for both IP video delivery and more critically 4k video delivery on the Hybrid Fiber Coax (HFC) access network. If we try to extrapolate progress for the next 10 year cycle, Figure 3 illustrates a potential path for the 5 interlinked areas. The stepping stones drive towards everything IP over D3.1 and Wi-Fi enabled IP STB deployments – with eventual integration of 802.11ax in the 2018+ timeframe. As this hardware-enabled capability rolls in – the shift to more 4K content will happen slowly as 4K TV penetration increases, content becomes available and we continue the visual journey with Color Gamut, High Dynamic Range and Frames per Second to create compelling visual experiences.

Let us look at the toolkit we can develop to ensure the Wi-Fi video and high speed data delivery systems meet the needs of the consumer. Due to Wi-Fi's non-determinism – the key to solve the problem is to assume that all services and devices need to be prioritized to use the available airtime. This needs to be policed and mapped to the best possible result for the consumer at every changing use interval of the day.

We will outline the foundation, building blocks and some cement to allow successful architectures and solutions to make Wi-Fi airtime deterministic and allow the home user to have transparent, well understood Wi-Fi behavior in the home network. Some of this approach may require changing the way we view the delivery of cable services to the home user – where we take the customer through a journey similar to Cellular Voice technology where there is some acceptance of dips in quality (voice call drops) with these small compromises not diminishing the flexibility of the service. Additionally, we want to stretch the customer a little bit, by assuming some understanding of the nature of wireless networking and to allow the operator to help them improve wireless networking performance in their home. Of course these have to be very simple and direct in their implementation and cater to the majority of the capabilities of subscribers.

Determinism – enabling the 4K VoW factor.

Is there really a problem currently streaming video over Wi-Fi?

Before we look to solve the problem, we need first to agree that there is a problem with the current home network's ability to stream IP video to the user. This is typically an OTT service solution although many Multiple System Operators (MSOs) also have OTT-like streaming services. While there can be issues with streaming services from the Data Center through the Content Distribution Network (CDN) and Internet Service Provider (ISP) Access network – typically most of the streaming issues come from congestion or issues in the Wi-Fi home network.

In 2014, the average number of Wi-Fi connected devices is at 8 devices with many users well over 20 – we are getting ready for some challenging Wi-Fi times.

With ever increasing bitrates of OTT video has been a constant increase in the access bandwidth speeds delivered – to keep up with the bandwidth demand from users accessing the OTT MPEG-4 compressed video streams. Many of the OTT providers have also moved from progressive download to Adaptive Bitrate (ABR) technologies to try and cater for the vagaries of Wi-Fi quality and avoid the 'streaming video' killer – the buffering symbol.

Streaming technologies like HLS, Smooth, DASH, and HDS have all been architected around the vagaries of Wi-Fi performance and 'grab' segments of video to ensure that they have some insurance against the network, the Wi-Fi network in particular, not being able to get the next segment to the decoder. These technologies are greedy compared to the frame per second perfect model of QAM based video delivery and the VBV model of the world's deployed QAM STBs.

Let us just bring the elephant into the room… Our Wi-Fi networks are congested.
Let me rephrase that – our 2.4 GHz Wi-Fi networks are congested.

Using devices capable of 5 GHz with 802.11n and 802.11ac provides some relief to be able to offer increasing bandwidth levels of non-buffering IP video. We have to deal with being able to support both 2.4 GHz and 5 GHz Wi-Fi networks and increasingly focus on 5 Ghz as our in house video distribution network solution.

We need to approach Wi-Fi networking completely differently from wired networking and to think differently to create the right Quality of Experience (QoE) for the end user. This paradigm shift - admitting that we, as an industry, cannot guarantee quality levels to all devices all the time – is a hard one to accept as quality of service is key to customer experience and retention. However, user experience and quality has some latitude and if we provide consumers with the best experience almost all of the time – then occasional shortfall do not affect overall satisfaction levels.

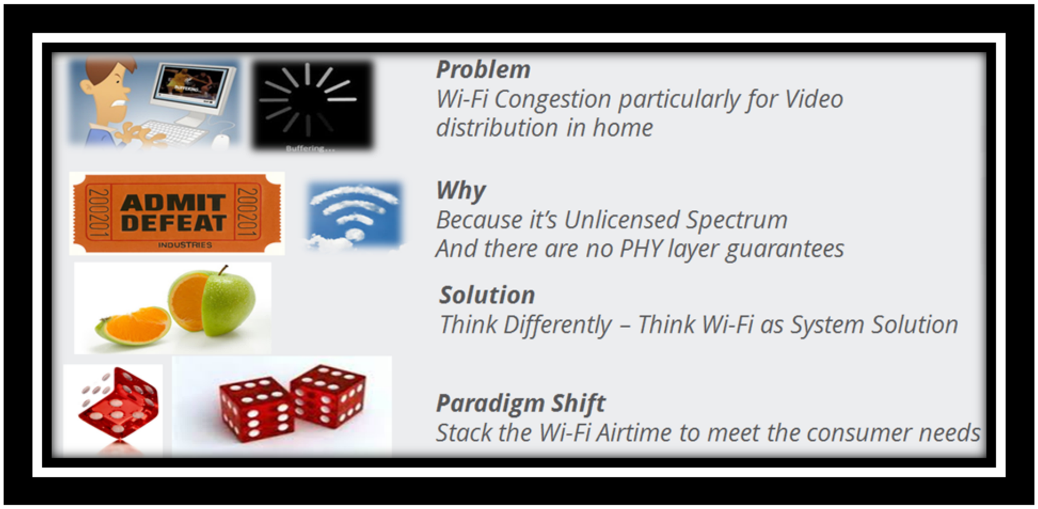

Figure 4 - Unlicensed Wi-Fi Spectrum requires different solution to wired networks

Consumer Behavior

Consumers are watching a lot of screens and lots of TV including in multiple rooms and increasingly in new modes like binge viewing of series content. This section pulls data from ARRIS 2014 Consumer Entertainment Index to illustrate current consumer behavior findings that impact the consumer's video usage and Wi-Fi QoE expectations.

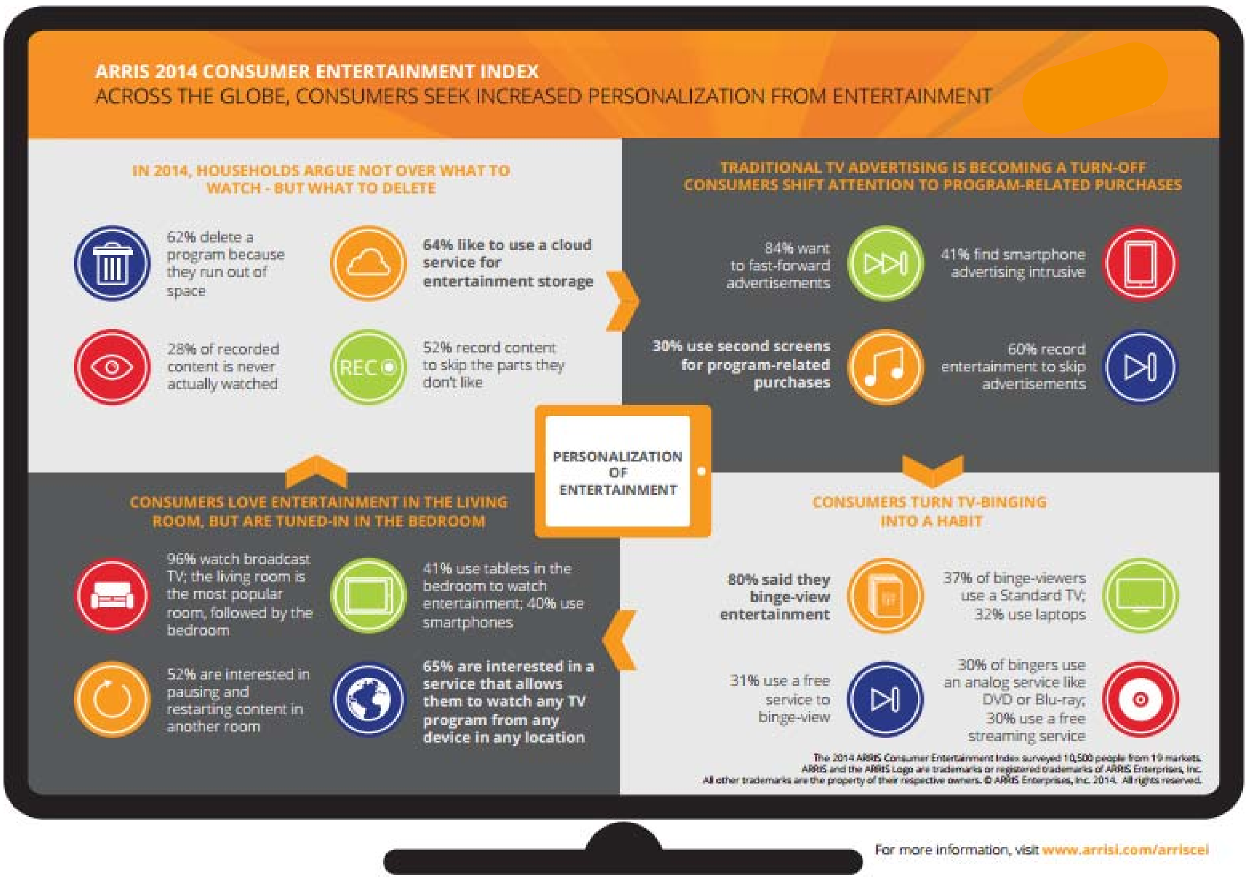

Figure 5 - ARRIS Consumer Entertainment Index 2014 Summary.

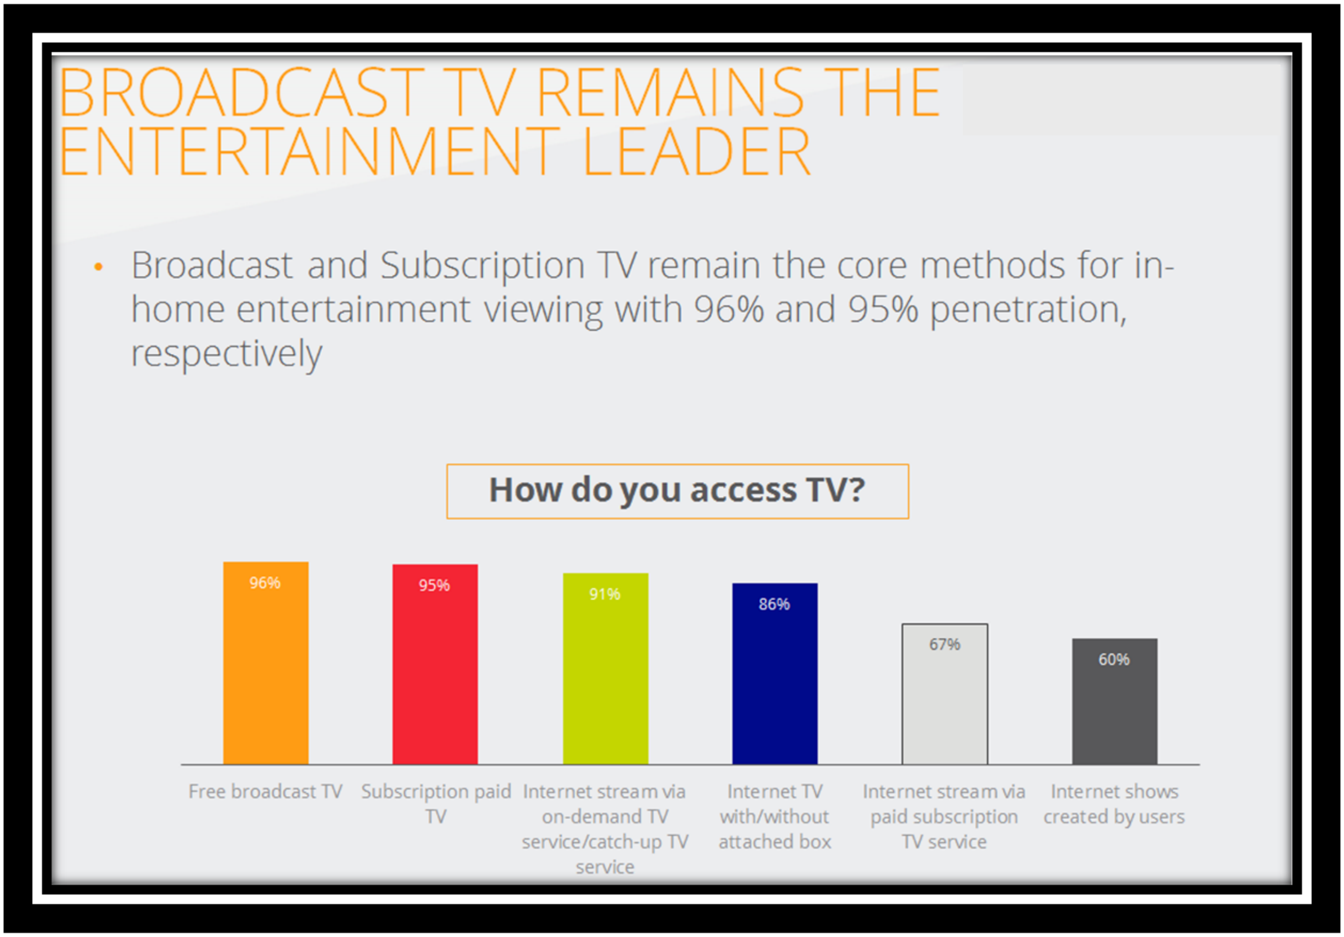

Figure 6 - Broadcast TV still key service for users[3]

Consumers are looking at other sources and screens but still have broadcast and Subscription TV at the core of our home entertainment experience. OTT and other viewing habits have added to the time spent watching video entertainment and not cannibalized core broadcast viewing habits. Consumer Digital Video Recording (DVR) of broadcast content has increased, accounting for 15% of viewing – continuing to grow – and still hitting the saturation point as 27% of content recorded is never watched.

3 ARRIS Consumer Entertainment Index

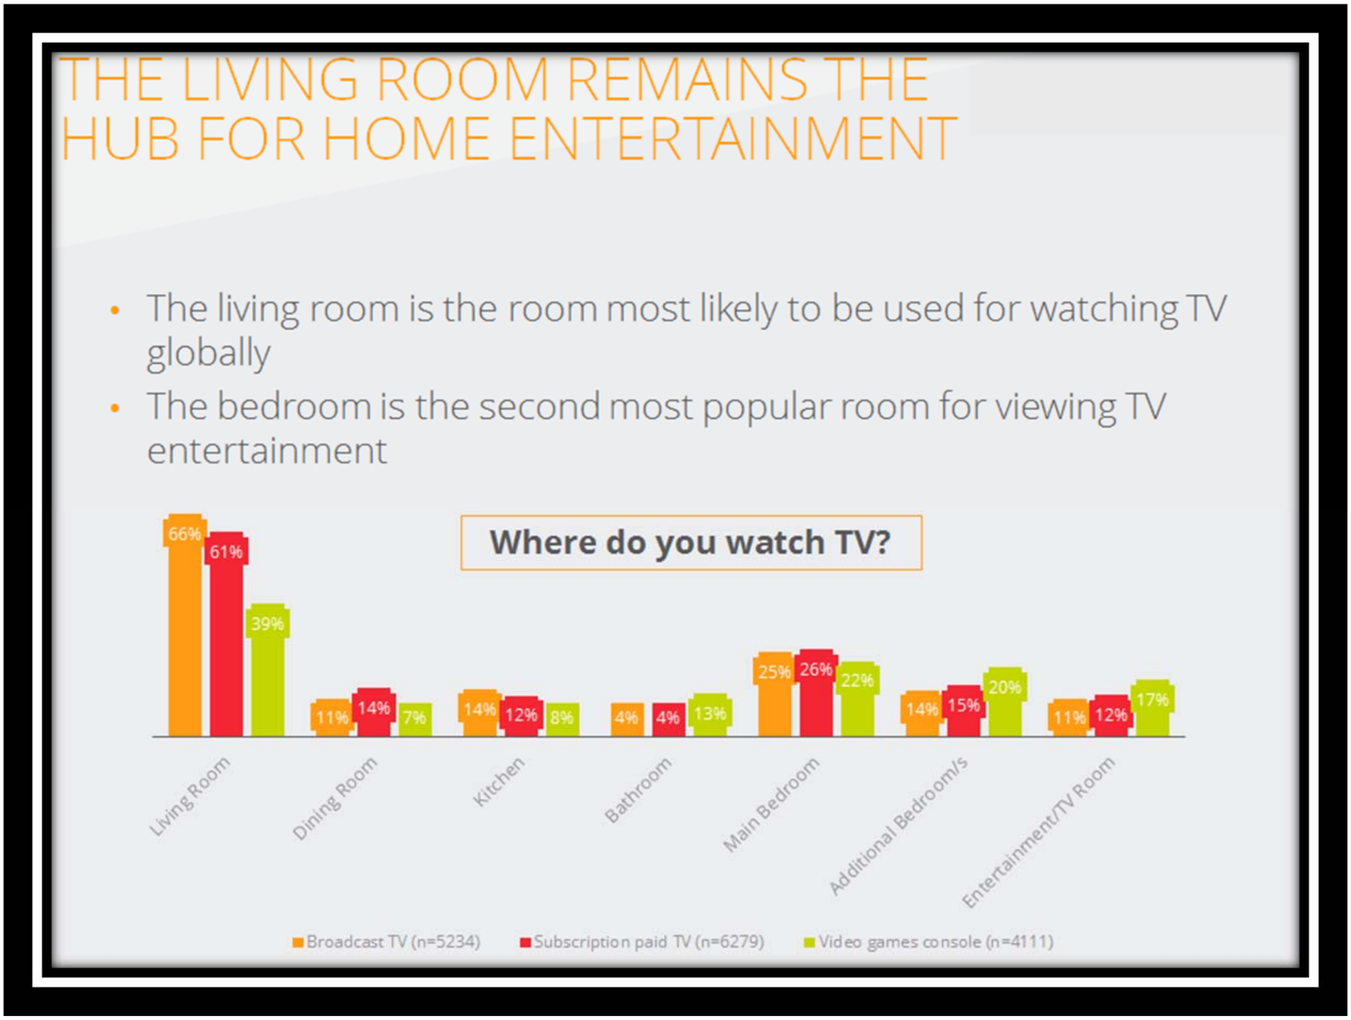

Figure 7 - Screens everywhere in the home – and increasingly Wi-Fi connected[4]

The further from the living room we are, the more likely we are to use a wireless connection for the connected device or screen. Using laptops and tablets in rooms like the bathroom and bedroom is driving significant performance challenges for video transmission over Wi-Fi.

As operators start to market wireless video connections more to facilitate large screen TV connection in places that don't have coaxial outlets, we feed consumers' requirements to have fixed large resolution screens in new areas of their home - but create a future problem for 4K content delivery. We are advertising TV on the patio and making TV's mobile. While these TVs are typically 1080p, it sets an expectation for the consumer that wireless TV service can be 4K in the future.

[4] ARRIS Consumer Entertainment Index

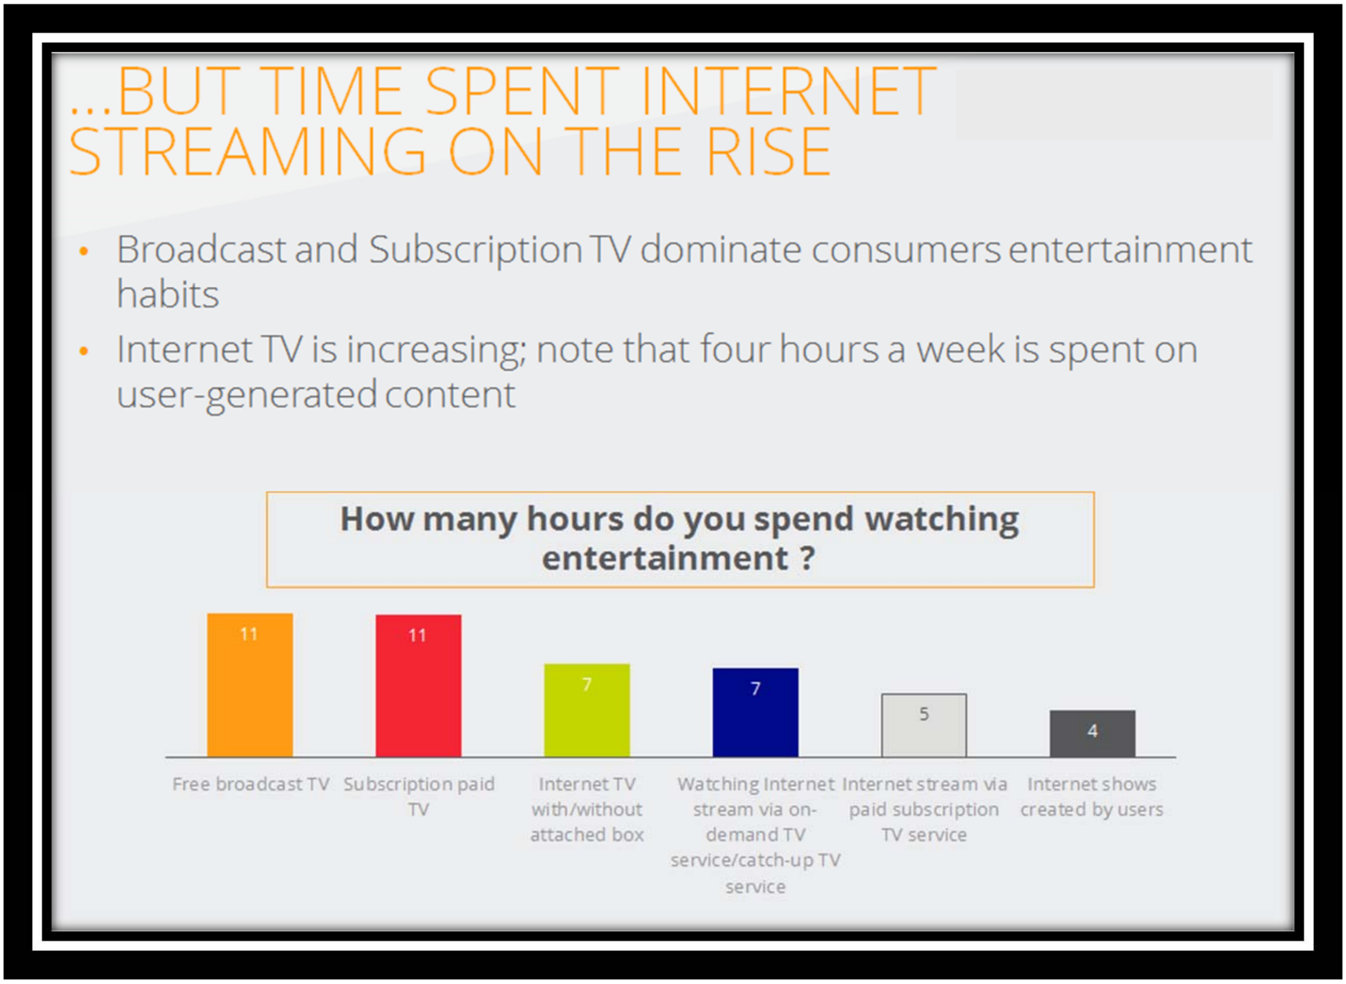

Figure 8 – Time spent Internet streaming is on the rise – and typically Wi-Fi connected[5]

Consumers value all services and are now vigilant on value for money and quality as they compare the relative cost of Broadband plus OTT Video Packages versus Pay TV service. They constantly question the ratio of cost to value for money. Every MSO must work to drive value to the consumer and differentiate their services to keep the 'stickiness' of 'always there' reliable video service.

A QAM to Coax outlet-based delivery system can send 4K content to an infinite amount of concurrent home TV devices in the largest of homes – current Wi-Fi solutions cannot. In fact, in many homes, a single reliable 15 Mbps session can have problems. Four

[5] ARRIS Consumer Entertainment Index

concurrent 15 Mbps is often not supported by the HSD access tier purchased by the consumer, but moreover the Wi-Fi airtime needed to stream 60Mbps is not available.

True, much of this relates to 2.4 Ghz – and 5 Ghz bandwidths and multi-antenna-based Access Points (APs) and client devices will help to realize high quality high bandwidth delivery over 5GHz Wi-Fi. However, getting ready for congestion in 5Ghz is the primary purpose of this paper – and solutions to keep us away from the dreaded buffering icon.

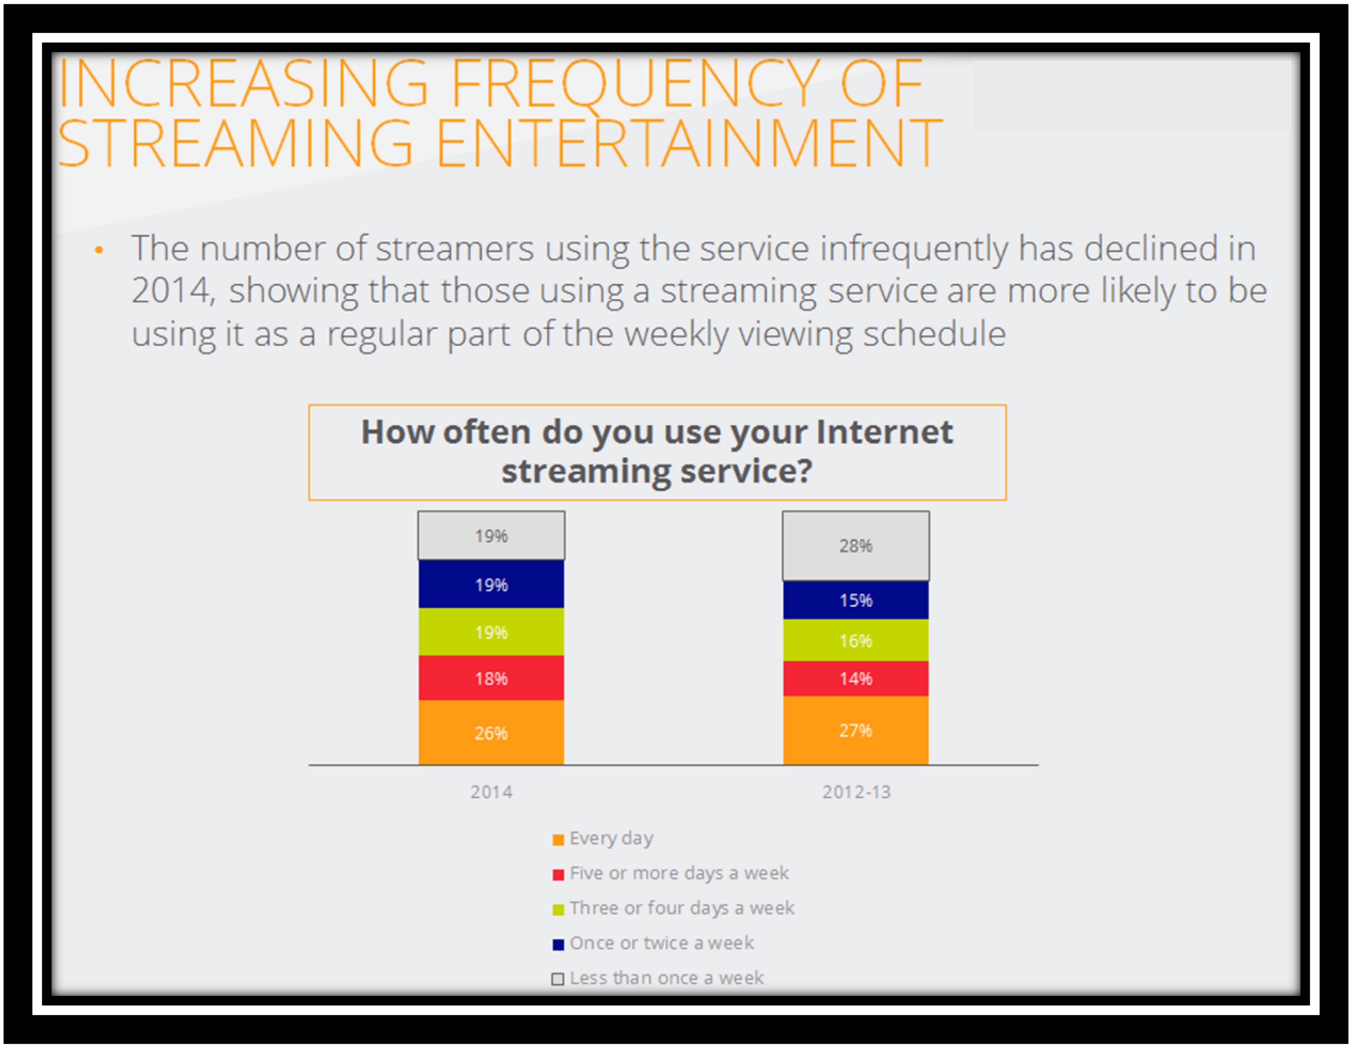

Figure 9 – Increasing Frequency of Internet Streaming – and Typically Wi-Fi Connected[6]

[6] ARRIS Consumer Entertainment Index

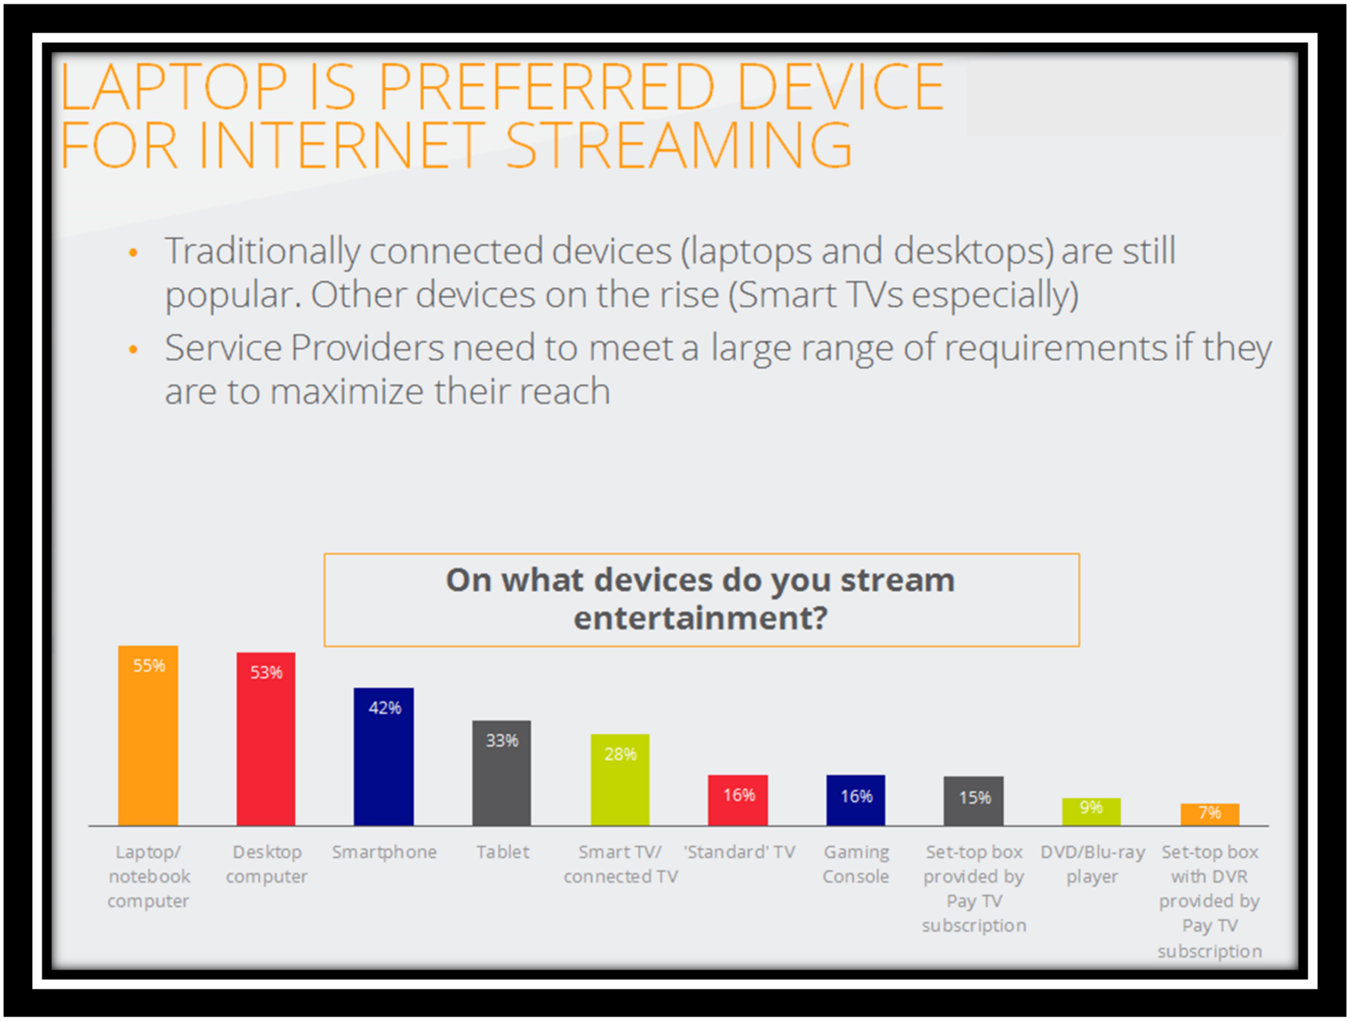

Figure 10 - Devices used to watch streamed IP video content[7]

For streaming services, the laptop and desktop devices are still used more than mobile devices. There are reasons for this:

- People still invest in a having a laptop for work and other typing intensive activities
- Tablets are being bought at higher frequencies to a point where more family members have their own

It is also important to note here that when we explore our toolkit to improve VoW, the number of antenna on both the gateway and the client device forms part of the solution

[7] ARRIS Consumer Entertainment Index

matrix. PCs and tablets are gradually increasing the number of transmit (Tx) and receive (Rx) antenna to be able to support more RF diversity and MIMO functionality.

Conversely, the 1x1 Tx/Rx solutions in smartphones cause the most problems in the home Wi-Fi network – for both coverage (APs eventually can't see the smartphone at edge of network) and for airtime efficiency (at the edge of the network forcing lowest modulation schemes and taking up more airtime slots).

A multi-antenna solution is also important for MSO-provided IP STBs with convergence on solutions offering 2x4 (Tx/Rx), designed to optimize range and receive throughput in the home Wi-Fi network – particularly for those 4K video streams of the future.

The cumulative use of devices like Smart TVs, Blu-ray players, and tablets (Figure 11) in some homes can exceed the use of single device types.

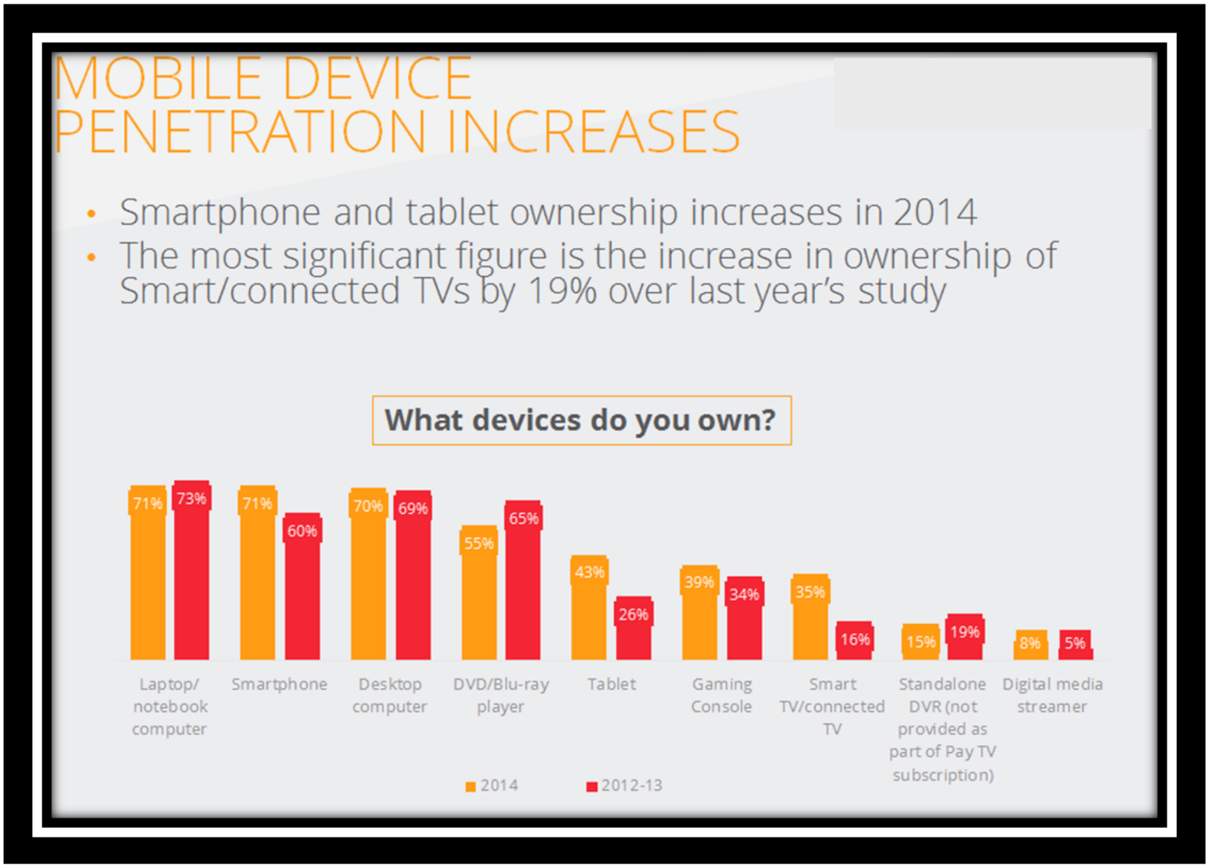

Figure 11 - Smart TV and tablets growing fastest[8]

[8] ARRIS Consumer Entertainment Index

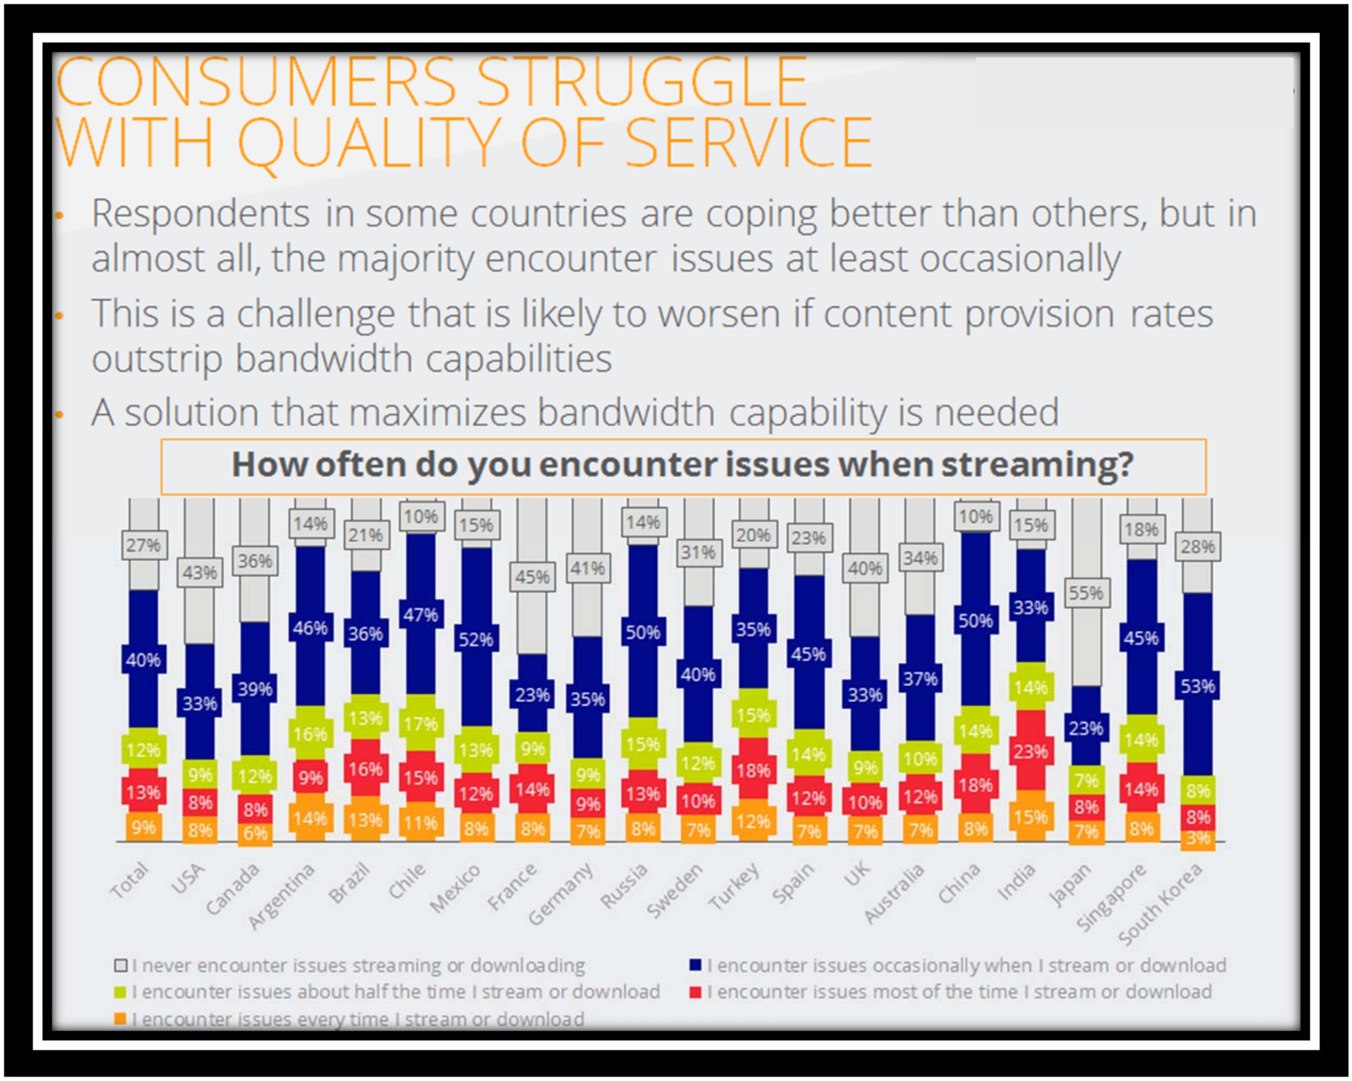

Figure 12 - Smart TV and tablets growing fastest[9]

People are watching more rather than less video, in every corner of the house and garden and increasingly on Wi-Fi rather than wired connections. Not surprisingly then the statistics show that users are seeing streaming quality problems. As an example, in the United States, 57%[10] of people have regular streaming problems with 8%[11] of users experiencing problems every time they stream or download video content. There are factors here which are not always Wi-Fi related, but Wi-Fi network congestion is a significant contributing factor.

[9] ARRIS Consumer Entertainment Index
[10] ARRIS Consumer Entertainment Index
[11] ARRIS Consumer Entertainment Index

The Stepping Stones and Toolkit to make Wi-Fi and VoW Deterministic

There are five main areas to target to make VoW Quality of Experience optimal at every time slice of home user consumption. These are outlined below in Figure 13 and we will explore each one of these throughout the main portion of this paper.

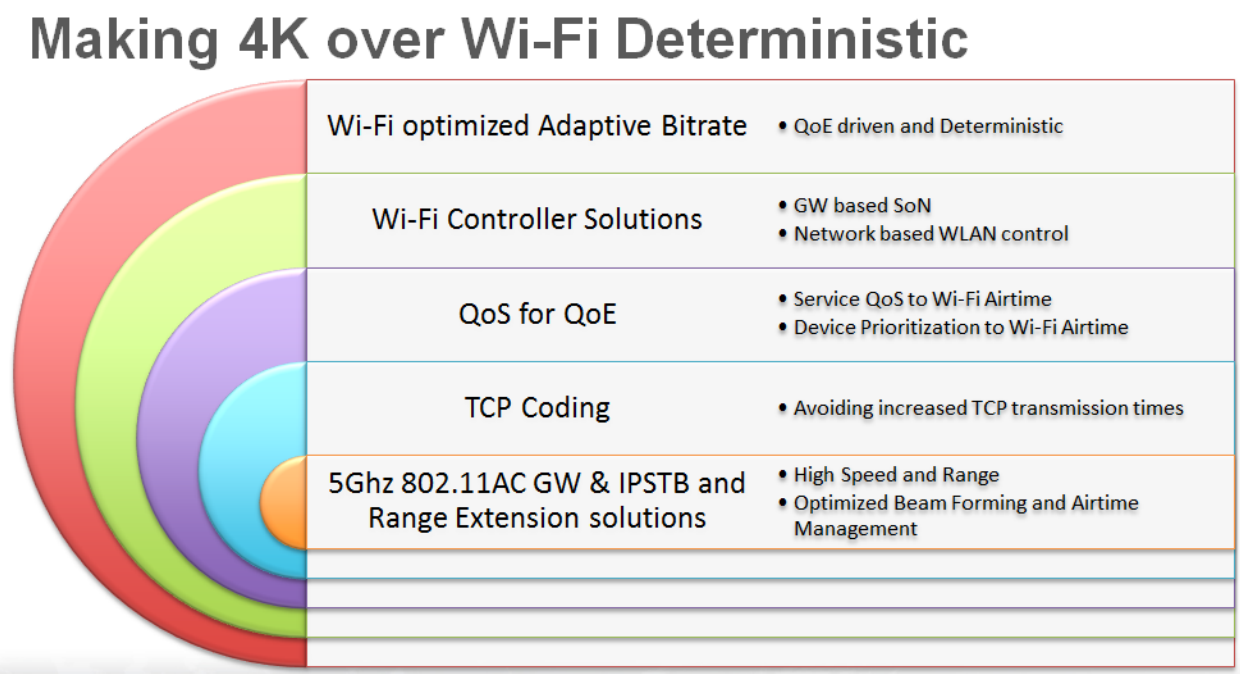

Figure 13 - The 5 Stepping Stone Toolkit items of Deterministic VoW

The main philosophy in this approach is to ensure that the fundamental hardware Wi-Fi capabilities of home devices are wrapped in software, policy, and cloud solutions that add deterministic elements to maximize closed loop service delivery to home devices. The integrated package of core hardware elements with software solutions can succeed where generic solutions cannot completely cover all reliability issues.

Core of the VoW Solution – the Trifecta of Wi-Fi devices

At the core of the solution for deterministic VoW is a trifecta of devices that interoperate to get a more powerful delivery than each one separately.

Figure 14 - The Home Gateway, Wi-Fi Range Extender and Wi-Fi STB

The main premise of this new architecture is that the customer needs the best Wi-Fi signal strength possible to ensure at the most fundamental of levels - the PHY layer – we have a chance of getting the end device good service. However, while we increase range and signal strength, those same range extension devices can also act as interferers taking up spectrum themselves – so using Wi-Fi controller and/or self-organizing techniques to manage operator-owned APs may also improve the use of the Wi-Fi airtime.

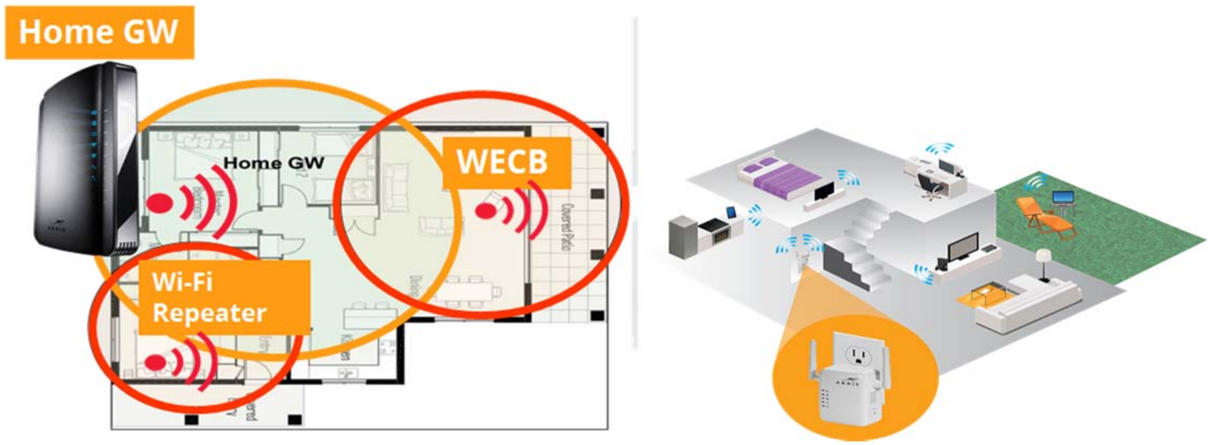

Figure 15 - Range Extension adds its own interference that needs to be managed

A balance between cost and complexity of solution for the subscriber always needs to be considered. The following key features and directions are now considered table stakes for this core element of the home Wi-Fi solution:

- Home Gateway Device
  - Engineered with highest RF output power to country-defined limits
  - Now targeting 5 GHz frequency for video delivery solutions but maintaining 2.4 GHz dual band concurrent operation
  - Engineered with multiple radios and antennas to offer RF diversity and supporting MIMO
  - Delivering 802.11ac performance
  - Offering advanced Beam forming and Beam steering technologies
  - Offering range extension support via physical interfaces
    - Gigabit Ethernet, MoCA 2.0, G.Hn, Wi-Fi
  - Potential future solutions utilizing increased channel and MIMO to 8x8 levels
- Wi-Fi Range Extension
  - Two typical flavors – Wi-Fi Repeater or Wired to Wi-Fi extension
    - Wired to Wi-Fi - typically backbone of Ethernet, MoCA or Powerline
    - Wi-Fi to Wi-Fi
  - Auto-Discovery / Auto-configuration and Meshing Technologies are being used to make the addition of the device simple
  - Effective throughput on the Wi-Fi spectrum can be 50% lower as the repeater may use some of its Receive time to capture packets and retransmit into the same spectrum. The problem is worse at 2.4 Ghz rather than 5 Ghz where there is currently more available spectrum.
  - Cost of this additional device and the ergonomics of location have been additional challenges
- Wi-Fi Enabled IP STB

- MSOs looking to move their managed video service to Wi-Fi enabled STB
- Cost, quality and size of Wi-Fi additions are a key element of the solution decision
  - Converging to 2x2 and 2x4 like solutions
- Need to ensure that good communications can be maintained in both directions

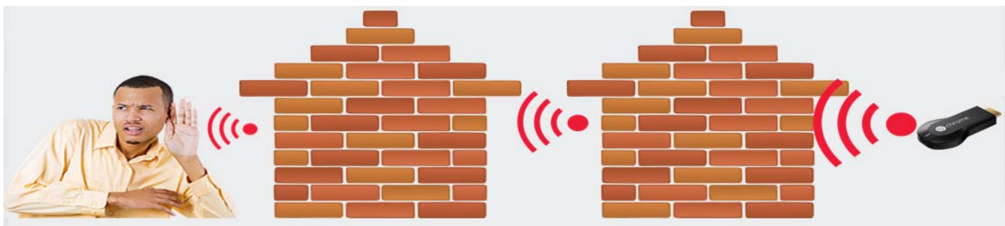

Figure 16 - Ultra-Low Power Wi-Fi STB may have range problem of the AP hearing it

The Home Gateway Wi-Fi Essentials

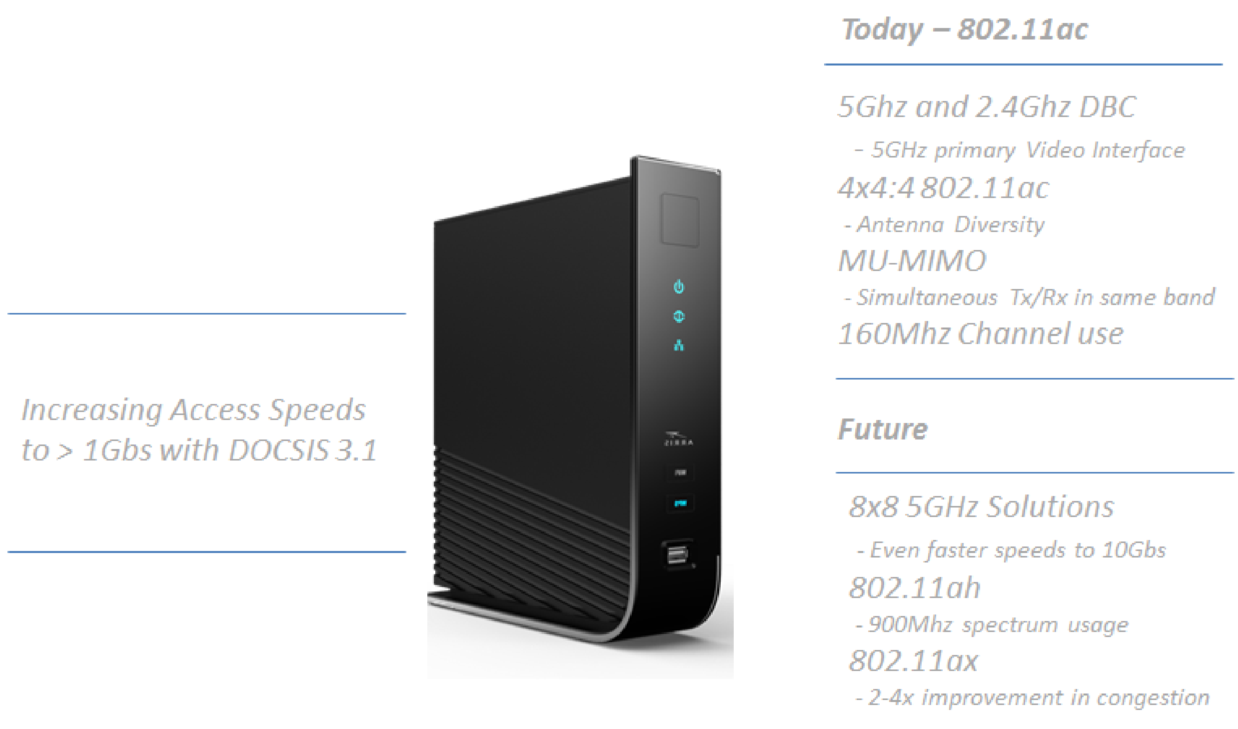

Figure 17 - The Home Wi-Fi Gateway – Big Powerful Wi-Fi Tx with sensitive Rx

The home Wi-Fi gateway is the core of any solution and includes an array of features to ensure that the MAC and PHY capabilities of the hardware are utilized. It also contains features to ensure use of cleanest channels with scanning and detection algorithms and also more sophisticated Beam forming and Beam steering solutions. These attributes are the subject of other papers so will not be covered in detail in this paper.

One additional item worth mentioning is that the 2.4 GHz band is sometimes written off for managed video transmission. This should not necessarily be the case. With its superior range propagation characteristics, it can be leveraged when 2.4 GHz spectrum has clean channels that can sustain the needed throughput rates.

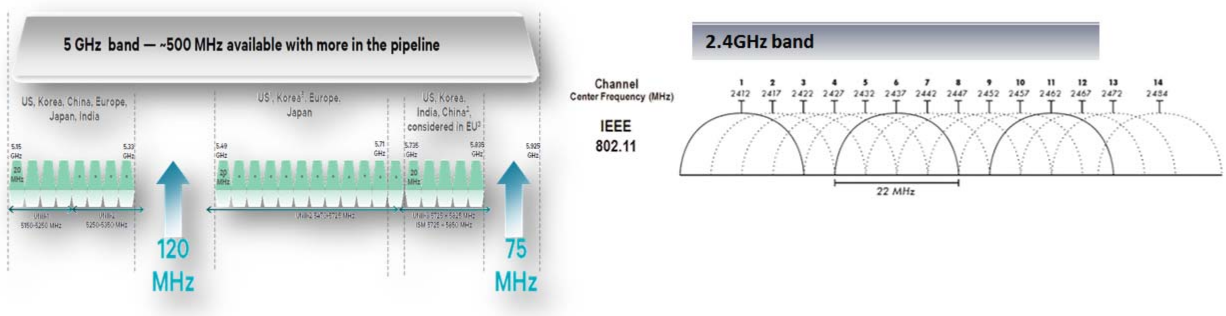

Figure 18 - 5GHz spectrum new hope for video transmission

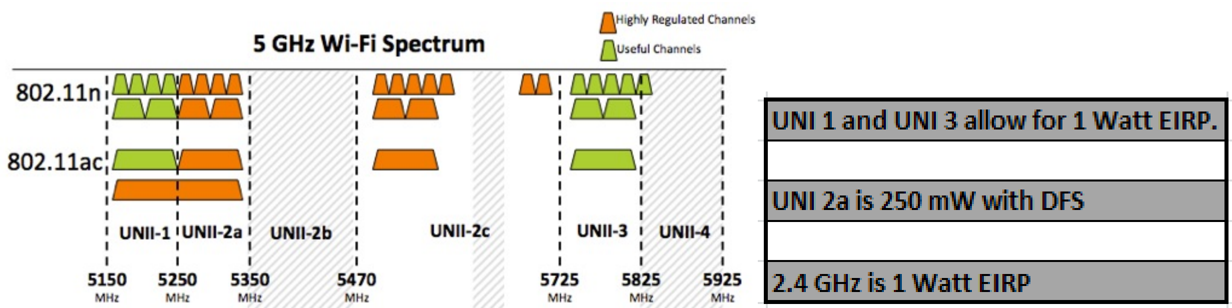

Figure 19 - 5GHz Specific Usage Bands and Regulatory

The Range Extension Wi-Fi Essentials

Figure 20 - Range Extension Solutions

We are moving more and more to a multiple access point home with strong signal strength throughout the household. We are now also trying to get good coverage even outside the home. The addition of devices for range improvements can act as additional interference, so the control of these additional APs is hugely important for overall airtime efficiency in the home. The use of a multistage controller function in the gateway and the network is explored later on this paper.

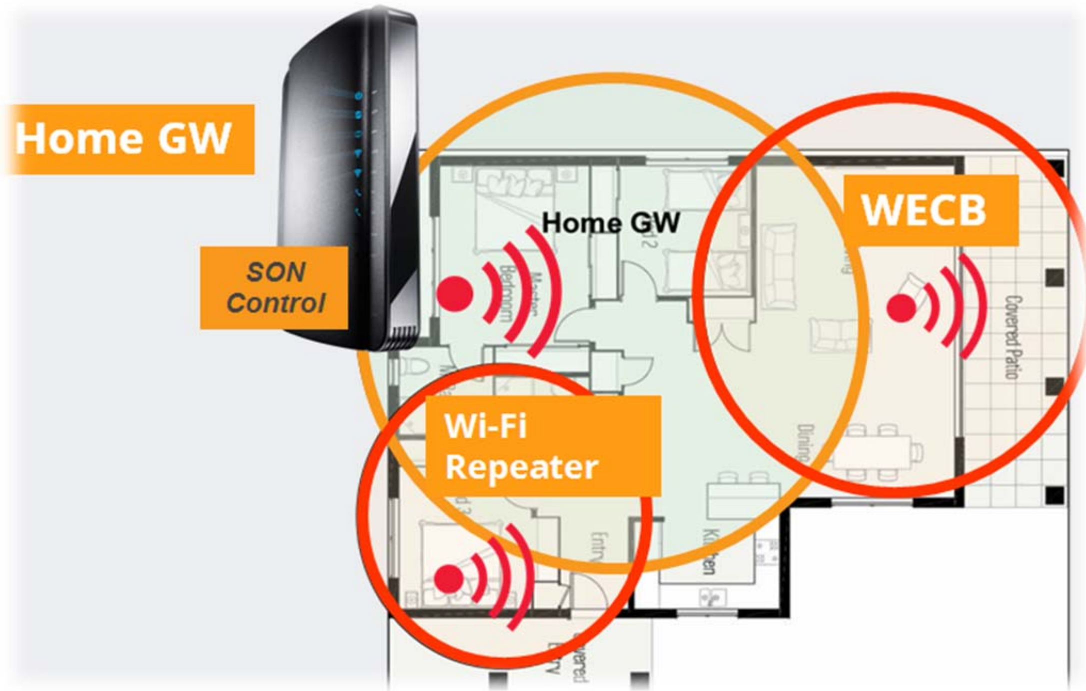

Figure 21 - Gateway as µController in the multi Access Point home

One item for future discussions and consideration is the potential change of Wi-Fi AP integration with the move to DOCSIS 3.1. DOCSIS 3.1 gateways (GWs) are likely to drift more and more to the entry point of the home RF network and may optionally not have integrated Wi-Fi. These GWs will certainly have options for range extension through high speed Gigabit Ethernet or Dual Band MoCA or G.Hn to other locations in the house. An example of how this could look is shown below and one extreme version of this could be to harden the D3.1 GW to the entry point and potentially move the AP to a first floor location in the home.

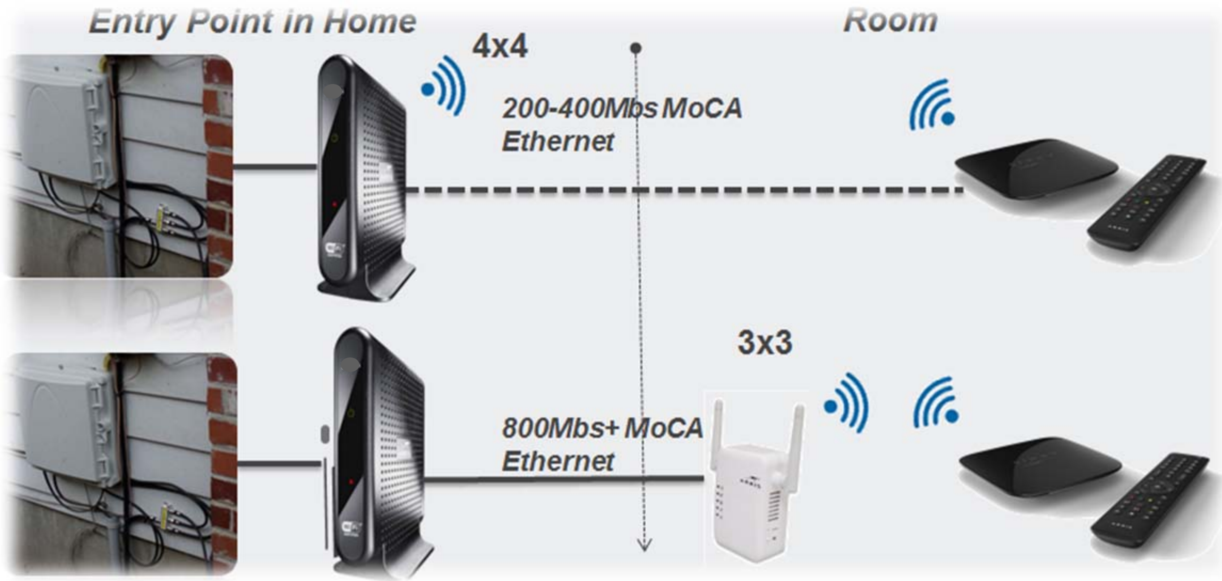

Figure 22 - Potential ideas for better location of Wi-Fi AP

The Wi-Fi enabled IPSTB Wi-Fi Essentials

Figure 23 - The new star – Wi-Fi enabled Cable IP STB

Important aspects of the 4K video service configuration

The type of video use case, real-time or off-line, influences the expected video bitrates. Also, since sporting events are likely to be highlighted as driving subscriber adoption of 4K, most Wi-Fi solutions should consider higher ends of the bitrate window as likely and important use cases for the solution.

The delivery of fragmented IP video has a typical implementation that assures that the STB buffers 1 to 3 fragments ahead of beginning play out of the video. A new fragment is requested when the buffer begins to run low and typically the delivery of the next fragment can be accomplished across 1-3 seconds without any interruption being noticed by the end user. The fragment's delivery speed or bitrate is not important as long as all of the fragment's packets arrive within the delivery window. Also, if any packets have been lost, the TCP protocol automatically requests a retransmission from the server. These retransmissions are in addition to Wi-Fi retransmission mechanisms built into Wi-Fi itself. Because of the temporal fluctuations in bandwidth common to Wi-Fi, this robustness against instantaneous jitter and delay is very important for reliable delivery.

Another area of characterization for the IP video streams is the profile of the encoded video stream. For a fragmented stream, the profile is set by the capabilities of the streaming server and the fragment size. Based on the server's capabilities, for example, the peak stream rate can be characterized which, in turn, can inform how the AP needs to ration its limited bandwidth. If the AP knows that the stream is a fragmented delivery, it can apply a less restrictive priority setting than if the stream is a live linear feed. For a live streaming linear feed, knowledge of the expected peak to average bitrate profile (PAR) is significant. Some encoders may generate streams with high peak to average bitrates which may place stress on the AP as it seeks to balance the airtime allocated to several clients. If more than one STB is being served by a single AP and the PAR is high, there may be instances where the available Wi-Fi bitrate may be temporarily oversubscribed, even though the average bitrate across a few seconds may still be within acceptable limits. For best performance, the buffers in the STBs for live linear streams with high PAR may need to be greater than the average bitrate might otherwise indicate to allow for robust performance when multiple high bitrate streams are available.

One other area to consider is the additional flexibility provided by an Adaptive Bit Rate (ABR) delivery system using fragmented content. On a 4K display screen, degradations of resolution are apparent in short order, but in the event that there is simply not enough bandwidth available to provide the 4K bit stream, downshifting through an ABR protocol can allow the viewer to continue to view their program without interruption. For live programming in particular, it is probably preferable to continue viewing a decreased resolution program rather than halt the program session entirely.

The following table gives a rough view of the variation in bitrates between different resolutions, and also highlights that even “4K” streams can vary widely in their bitrate depending upon the frame rate as well as other aspects such as color depth.

Resolution	Bandwidth (HEVC encoding)
4k p60 (60frames/second)	30 Mb/s
4k p30 (30 frames/second)	15 Mb/s
1080 p60 (60 frames/second)	7.5 Mb/s
1080 p30 (30 frames/second)	3.75 Mb/s
720 p60 (60 frames/second)	3.33 Mb/s
960x540 p60 (60 frames/second)	1.875 Mb/s
768x432 p60 (60 frames/second)	0.9 Mb/s

What does this mean in the air?

Typically, the 1448 byte MTU of Ethernet interface becomes the TCP fragment size used on client and AP side unless changed in drivers for both AP (e.g. 11n/ac embedded in gateway box) and client (thin client at TV), or we can use jumbo Ethernet frames on Wi-Fi interface to the TCP stack. Air transfers of the 1448-byte fragments are done as MPDUs or an A-MPDU. For example, a 2-second HLS media chunk carrying, e.g., AVC+HE-AAC at 1.6 Mbps for a mobile app consists of ~200kbytes or 138 x 1448-byte TCP packets. For 40 MHz 11n single spatial stream at rate 5/6 64QAM MCS (150 Mbps PHY rate), this requires about 22 11n OFDM symbol durations per 1448 byte MTU or 77.2 µsec airtime per MTU or 10.65 msecs to send the whole 200 k-byte chunk, not including MAC overhead for inter-frame spacing’s for channel access and other overheads. To mobile devices (such as tablets and smartphones) given their small size and limited antenna, typical speeds rarely exceed more than about 40-50 Mbps 11n MAC throughput even sitting next to a 5 GHz router with 20 MHz channels.

For 11ac at 256QAM rate 5/6, 80 MHz channel and single spatial stream (433 Mbps PHY), the PHY sends 195 bytes per OFDM symbol so a 1448-byte packet would be transmitted as a single MPDU in about 8 symbols or 28.8 µsecs. A 30 Mbps 4kp60 2-second media chunk would contain about 7.5 Mbytes or 5180 TCP packet requiring 0.15 seconds to deliver, neglecting overhead. The 11ac MAC allows a max A-MPDU size of 1 MByte (64 MPDUs at 16kBytes) and an 11ac AP will perform aggregation of the 1448 byte MTUs if the TCP connection is running at maximum speed (in congestion avoidance), but the Wi-Fi air interface can take advantage of the Block ACK feature to reduce the number of ACKs needed and make more efficient use of the limited airtime resource.

Important Aspects of the Wireless Environment

Because of the expanded requirements of 4K video on the wireless channel, the AP needs to evaluate whether the current Wi-Fi channel is likely to have sufficient bandwidth to support a new 4K stream. The channel data bandwidth for Wi-Fi can be influenced by many factors. The AP first may be supporting other clients with their own data requirements. The AP may be able to develop traffic profiles for those other clients that can be used to predict whether the addition of a 4K stream will be problematic. The configuration of the AP may also affect the bandwidth availability, for example, each SSID added to the AP can reduce the available bandwidth by 1% to 3%. To mitigate these issues, a dual band concurrent AP may be able to separate non-video clients from video clients by assigning non-video clients to a separate frequency band from that used by the video clients.

Many factors may influence the Wi-Fi environment outside of the AP's control. There may be many other APs with clients surrounding the AP in question. These other units can cause congestion in the airwaves for the video delivery AP when their transmissions cause the video AP to have to back off and wait for a clear transmission time. Non-Wi-Fi interference is also prevalent in the 2.4 GHz band. The best solution for a video AP facing Wi-Fi and non-Wi-Fi interference is to at least attempt to find a clean channel or at least a channel with fewer interference sources. This solution can also improve the performance of the non-video clients.

Over longer time frames, the AP may be able to determine traffic patterns that can be used to predict performance issues. If the AP finds that congestion tends to happen every evening around primetime, it may be able to proactively change channel or proactively shift to a more aggressive CCA algorithm during that time. If the AP is integrated into the video delivery subsystem and the STB device has local storage, it may also be able to request an advance download of material in off-peak times.

Channel/Spectrum Management

Airtime analysis with co-channel interference detection of both Wi-Fi and non-Wi-Fi sources enables the AP to make intelligent channel selection choices. The selection of the appropriate channel bandwidth may also be influenced by an airtime analysis. For example, if the environment sweep detects an existing 802.11n system operating with a 20MHz channel, the video AP may be able to define a channel bandwidth that skirts that interfering signal entirely. Similar to the advantages to sweeping the airspace for other interferers, supporting radar signal detection is also advantageous allowing an AP to utilize the lightly used UNI-II midband spectrum.

The 802.11 standards also offer many opportunities to conserve bandwidth for other services. For example, if the video AP is required to offer multiple SSIDs for other services, the AP may be instructed to use beacon bursting to allow a single beacon burst to communicate information about multiple SSIDs.

Real-time Link Management

Once a 4K video delivery session has begun, the AP must diligently monitor the session for possible degradation. If the AP has information about a required bitrate to sustain the service, then it can track the current signal strength and quality to ensure that the bitrate can be met. This task becomes more complicated when there are several streams to be balanced against one another. The detection of potential issues that might lead to a robustness action must be made as soon as possible to avoid any user visible impact. A potential issue might be detected by reports for the client of excessive errors or link degradation. If the AP can receive indications from higher layers within the 4K STB, the STB may also indicate to the AP that its video buffer is becoming alarming low. In most cases the AP would already be aware that data was backing up at the AP because of retransmission or other actions, but if the delay is being caused higher up in the network, it may be useful for the AP to be notified that the STB may switch to a lower bitrate stream to attempt to avoid having to pause the play out.

Video Link Management and Admission Control

As was discussed earlier in this paper, the AP may operate independently or it may be in communication with higher level video management entities. The management entities may inquire of an AP as to whether a new 4K session can be supported, or the AP may just have a new session begin without any notice beforehand. The AP can characterize the current state of the network, using tools discussed earlier, and use that information to decide whether or not a new 4K session can be supported without impacting other services.

For example, an AP may determine that a new 4K session can only be supported by restricting the bandwidth available to other best effort services. A more complicated decision could be posed by the existence of several high priority sessions, such as other video sessions, that result in insufficient bandwidth remaining for a successful 4K video delivery session. In that case, the AP could have a logic table that defines which streams get preference, or the AP could send a message to a higher layer of software asking for a decision on which stream to interrupt. If that higher layer of software included a user interface, that user might be asked which stream(s) will be allowed to continue playout, similar to the current multi-tuner DVR STBs asking which programs to terminate when its tuner supply is exhausted.

Similar problems can also arise when the end user changes programs from a lower resolution program to a higher resolution program. A robust video capable AP must have a series of rules to allow it to provide deterministic behavior even in challenging circumstances.

Fast Efficient Channel Change

An important feature for a video ready AP is the ability to move clients to new channels when the current channel deteriorates. As mentioned earlier, that channel impairment may be caused by other Wi-Fi devices or by interference from other users of the unlicensed bands. The AP must be able to direct clients to the optimum interface for their expected use cases. The AP must gather information during idle periods to keep its understanding of the surrounding channels up to date. When the AP is selecting a new channel for video delivery, it should preferentially select high power channels to enable the widest range. In the 5 GHz band, it is worth noting that several channels fall into this category, yet the expanded dwell time requirements of DFS bands make it difficult to utilize them successfully for channel change during times of active video sessions.

Coded TCP – Enhancing Delivery of VoW

Video Delivery using TCP or UDP

Compared to wired communication systems, wireless systems are prone to channel errors and interference, both leading to packet errors. Two approaches to dealing with these errors are Forward Error Correction (FEC) and Automatic Repeat-reQuest (ARQ).

Video over Wi-Fi (VoW) can be delivered in the home using many different approaches, including delivering transport streams over UDP/RTP and over TCP/HTTP. In most cases the video delivery involves the traditional store and forward techniques of routers and switches, transmitting the video content from the source all the way through to the receiver devices.

UDP/RTP Video Delivery and Application Layer FEC

In the case of UDP/RTP delivery it has been possible for several years to incorporate the use of Application Layer (AL) FEC with UDP/RTP delivery. Multiple AL-FEC schemes exist such as proMPEG COP#3 2D parity FEC, Reed-Solomon over GF(2^8) and LDPC-Staircase codes. The FEC schemes generally do not require any feedback mechanism, and can immediately recover from packet losses without requiring feedback. The lack of need of a feedback or return channel makes FEC particularly attractive for IP Multicasting of UDP/RTP streams. ARQ can also be incorporated with the UDP/RTP delivery scheme to cope with the situation when FEC fails to recover missing data (typically when packet loss exceeds the capability of the FEC scheme).

Most AL-FEC schemes for UDP/RTP rely on significant stream buffering in order to provide a large enough data set to generate sufficient and efficient FEC parity information, with this buffering introducing latency on the overall transmission. In the case of proMPEG COP#3 FEC 2D parity FEC with a matrix of 10x10 (or 100 media packets), the scheme can correct for both consecutive losses as well as individual packet losses. However, on a 2 Mbps video flow, such a scheme introduces a latency of 526.4ms and overhead of 20%, while being able to recover up to 10 media packets from 120 transmitted.

As the name implies, Application Layer FEC relies upon the video sender and receiver introducing specific application layer processing to be able to cope with packet loss, as well as backup operations to recover if the FEC scheme fails. In the case of proMPEG COP#3, the application needs to be aware of the main media packet UDP flows and up to two additional UDP flows carrying FEC information, and understand the FEC scheme and how each of the flows relate to the media packet flow (typically via RTP sequence numbers).

Such an end-to-end scheme relies upon a “one size fits all” FEC protection approach, with the consequence of having to factor in the worst possible network packet loss from the source to the receivers even though some network segments may only have minimal loss. The scheme relies on the traditional notion of store and forward networking, where data remains untouched from sender to receiver. The only way to modify the FEC is to receive the entire FEC protected block and run a replacement FEC over the received data, a process that incurs further latency.

TCP/HTTP Video Delivery

In the case of TCP/HTTP video delivery, video transport streams are delivered either end-to-end using something like Adaptive Bit Rate (ABR) Streaming or from a local DLNA media server using HTTP streaming. In any of these cases no FEC scheme is currently used, so the streaming relies completely on the ARQ capabilities of TCP to recover from packet loss or congestion conditions.

Wi-Fi Link Layer ARQ

When streaming video over wireless (802.11a/b/g/n/ac) regardless of the transport (UDP or TCP), packet loss is currently addressed through a combination of link layer ARQ and down shifting of Modulation and Coding Scheme (MCS). The change in MCS can have a significant impact due to the reduction in available bandwidth. Link layer ARQ applies to unicast data transmission, it is not applicable to multicast data transmission. Note that multicast over wireless does not have the same throughput capacity as unicast due to it relying on the lowest common MCS between receivers.

The link layer ARQ operates by determining if a transmitted packet (or aggregated set of packets) was successfully received by the far end transmitting an explicit acknowledgment. If the ACK is not received by the sender, the low level Wi-Fi driver (or hardware) retransmits the data. This process is continued until either an ACK is received by the sender or the limit of retries is reached. In the case of an ACK being received, a response is returned to the higher-level network driver to indicate successful transmission and the next packet is transmitted. If the retry limit is reached then the Wi-Fi driver indicates the network packet has dropped. This can be passed to the higher layer protocols in order to perform a higher layer transmission. Any retries use up the shared Wi-Fi airtime and tend to increase the Round Trip Time (RTT) for the specific data session, which for TCP sessions can limit the total throughput achieved.

Some Wi-Fi drivers will initiate a downshift from the current MCS used with the remote receiver in the event that a lot of packet loss is being observed (i.e. the retransmission limit is being constantly reached).

In the event that packet loss occurs, the current Wi-Fi scheme of ARQ tends to consume airtime (through the necessary back-off and retransmission scheme) reducing the overall capacity shared between all Wi-Fi users.

The link-layer ARQ mechanism of Wi-Fi helps significantly with maintaining high throughput of TCP sessions, as the link-layer retransmissions prevent packet loss from manifesting at the far end TCP receiver.

TCP Congestion Control Algorithm

The TCP protocol operates a congestion control algorithm that attempts to increase the TCP transmission rate relative to how reliable the network is performing, using a so-called “Additive Increase, Multiplicative Decrease” (AIMD) approach. TCP uses a sliding/congestion window scheme that provides a reliable, in-order delivery of packets. TCP controls the growth of the congestion window capacity using the AIMD approach, allowing the window to grow for every ACK received. However in the event of congestion being observed through duplicate ACKs transmitted by the receiver to the TCP sender, the AIMD algorithm kicks in and significantly reduces the size of the congestion window capacity and in turn significantly reduces the TCP throughput. Duplicate ACKs indicate to the sender the last successfully received TCP packet.

On a lossy network, such as wireless, the loss of a packet triggers the reduction in the congestion window. As a result, packet loss has a significant impact on the TCP session throughput , even though the wireless network continues to offer plenty of capacity (i.e. it is not congested).

The TCP congestion control algorithm triggers ARQ on receipt of the duplicate ACKs; unfortunately when this ARQ is triggered and the congestion/sliding window ‘collapses’, there is a direct impact on the throughput bandwidth of the TCP connection, even if the Wi-Fi channel is rated at a significantly higher bandwidth. Unfortunately TCP cannot make any distinction between congestion and packet loss. As a result, packet loss has a significant impact on the TCP session throughput, even though the wireless network continues to offer plenty of capacity (i.e. it is not congested). The following diagram illustrates this collapse (remember window size has a direct impact on throughput).

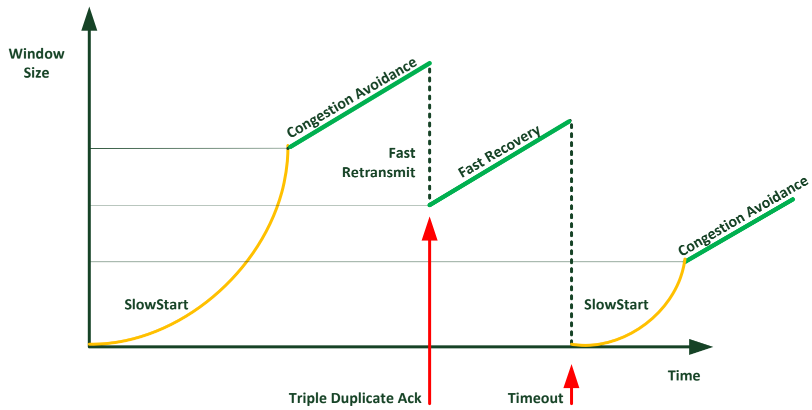

Figure 24 - Performance of TCP in Presence of Congestion and Packet Loss

As mentioned the use of Link Layer ARQ in Wi-Fi helps to mask the dropped packets to TCP, at the cost of wasted airtime and a potential increase in the RTT (both effect the throughput of TCP). With the increase in available Wi-Fi capacity in 802.11ac, some of the ARQ issues with Wi-Fi are probably more acceptable than when operating a 20MHz wide channel in 2.4GHz with 802.11n.

Robust TCP over Wireless Networks

Research Areas

For many years, research teams have been investigating how to make TCP robust over Wi-Fi networks, with a lot of focus on the link-layer ARQ handling, FEC and modifications to TCP congestion control. A team from UC Berkeley presented a paper [4] in 1996 titled "*A Comparison of Mechanisms for Improving TCP Performance over Wireless Links*" that described how a reliable link layer protocol, with some knowledge of TCP provides very good performance. Even though this was before the first official 802.11 network (802.11a – Sep 1999) it established the groundwork for other studies. Another later study, presented by University of Massachusetts in 2002, titled "*TCP-Cognizant Adaptive Forward Error Correction in Wireless Networks*", identified how wireless links were categorized by high bit error rates and intermittent connections that directly impact the performance of wireless due to TCP misinterpreting non-congestion packet loss as indications of packet loss, which drives down the TCP throughput. They suggested the use of TCP with adaptive FEC (AFEC), and presented how this approach out performed schemes such as TCP-SACK, Snoop and some other physical layer. Sometime later in 2010 the University of Trento, Italy, presented its study, titled "*TCP-Aware Forward Error Correction for Wireless Networks*", that again identified the benefits of TCP using FEC, but this time driven by TCP semantics. They highlighted how the FEC strength (of Reed Solomon encoding) could be increased or decreased depending on the state of congestion windows and TCP state. They concluded that the approach provides an excellent tradeoff between the offered error protection and the amount of total redundancy added to the TCP flow for FEC. They proposed implementations covering end to end for sender and receiver TCP stacks, or local implementations covering the wireless part of the connection only.

Coded TCP (CTCP)

A significant number of other research papers were also developed over time on these subjects (too many to list). One of the most recent papers on this topic was published in the Proceedings of the IEEE, titled "Network Coding Meets TCP: Theory and Implementation" by Sundararajan et al. It extends upon the years of earlier research in the area (including a modified TCP congestion window algorithm) but includes the use of an alternative error correction scheme (random linear network block coding) and a novel transport layer solution. An additional paper, "Network Coded TCP (CTCP)", by Kim et al., with some overlapping authors, restated the earlier paper and introduced more quantitative results from a real implementation of the scheme. The second paper describes a new transport protocol approach that introduces network coding directly as opposed to an indirect shim layer (as per original paper).

The rest of this document summarizes the network coding transport scheme (based on UDP) and describes the results identified in the paper and in some other protocol testing.

Coded TCP provides a reliable connection oriented transport over UDP with the added benefit of erasure correction. It follows the basic principles of TCP connection management and handling, and shares available bandwidth fairly with other real TCP connections. The typical implementation for Coded TCP is to act as a network proxy for real TCP connections. In this instance, a proxy exists at both the sender and receiver, diverting TCP connections over the point-to-point Coded TCP (UDP transport) connection. The Coded TCP proxy connection implements all of the listed features.

Coded TCP relies on a modification of the existing TCP congestion management algorithm, particularly in the case of AIMD handling. It eschews the use of the existing TCP sliding window principle in preference for a systematic block coding approach to enable the correction of missing packets (known as "erasure correction"). ACKs continue to play a key role in terms of identifying received packets, Round Trip Time and packet loss versus congestion. In the absence of packet loss Coded TCP defaults to a TCP-like protocol "without" coding, where all the congestion management controls apply.

Unlike TCP and the previous paper on TCP/NC which both use a sliding-window approach, Coded TCP uses Random Linear Coding (RLC). The basis of RLC is to assemble a collection of *p* data packets, split them into fixed size blocks of *n* bytes (typical 1460 bytes), and using a random set of *n* coefficients generate an additional set of *q* coded packets. The system transmits all *p* and *q* packets, and relies on the fact that if any *p* packets from the total *p+q* packets sent are transmitted, the original *p* packets can be recovered. If the system receives <*p* packets, then it must request additional packets to recover the original *p* packets. However, the number of *q* additional packets to include can be made proportional to the packet loss rate.

Flexibility is applied to how the block codes are generated. Using systematic block codes allows for totally un-coded data packets to be transmitted before sending any coded packets. A coded packet is generated by randomly coding all packets in the block. This approach is most effective in terms of erasure correction, where with high probability the coded packet will correct for a single erasure in the transmitted block. The approach also means that un-coded data can be accessed immediately without requiring decoding.

The low level details of network coding and specifically random linear network coding are too involved to be presented in this paper. Suffice it to say, the mathematics behind the erasure correction capabilities of RLNC have been proven in countless academic papers, and have been shown to be as effective as typical block codes such as Reed Solomon, LDPC, Raptor Codes and Luby Codes in correcting data.

In the case of comparison between TCP and Coded TCP the following diagram shows the basic "goodput" capable under varying packet loss rates for a 25 Mbps, 10ms Round Trip Time (RTT) network connection.

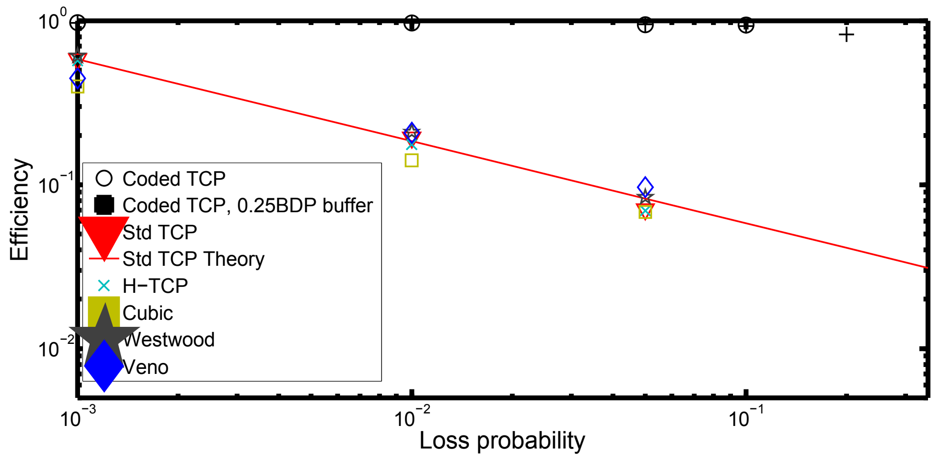

(a) Link 25Mbps, RTT 20ms

Figure 25 - Goodput with varying packet loss rates for 25 Mbps, 10ms RTT network (2012-CTCP.PDF)

The diagram compares multiple TCP congestion control algorithms to the Coded TCP implementation in a lossy network. The red diagonal line identifies the expected performance of standard TCP in a lossy network. The loss probability is shown at 10^{-3}(0.1%), 10^{-2}(1%) and 10^{-1}(10%). The efficiency axis provides a measure of how efficient (percentage wise) each of the congestion control algorithms can be in the presence of packet loss. The 100 value equates to 25 Mbps. The diagram shows for a 1% packet loss the majority of the existing TCP algorithms manage about 20% link efficiency, indicating that 80% of the available capacity is just not used. The Coded TCP algorithm markers on the graph show a 90% or greater link efficiency in the same packet loss conditions.

The following graph presents a closer view of the top 90% of the above graph, where Coded TCP is delivering the most "goodput".

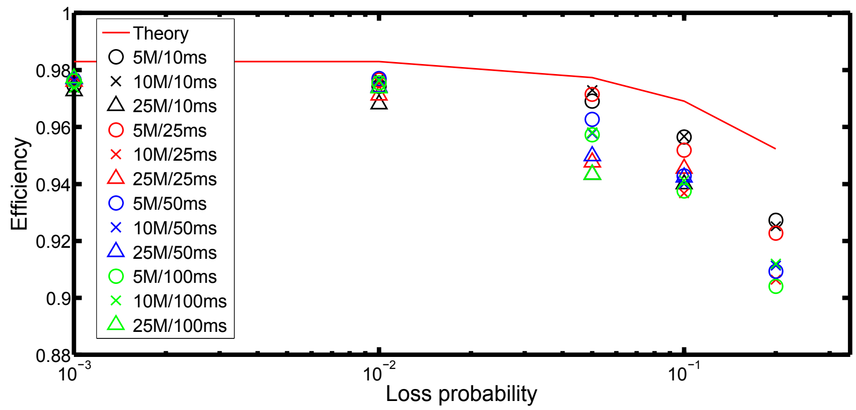

(b) CTCP

Figure 26 - Coded TCP Performance (2012-CTCP.PDF)

The above graph expands on the earlier graph and highlights the same loss probability and efficiency comparison for Coded TCP, by including different link speeds and RTTs. As can be seen, Coded TCP manages to maintain a link efficiency of greater than 90% for all the presented cases.

In terms of implementation for devices like cable gateways the following is proposed,

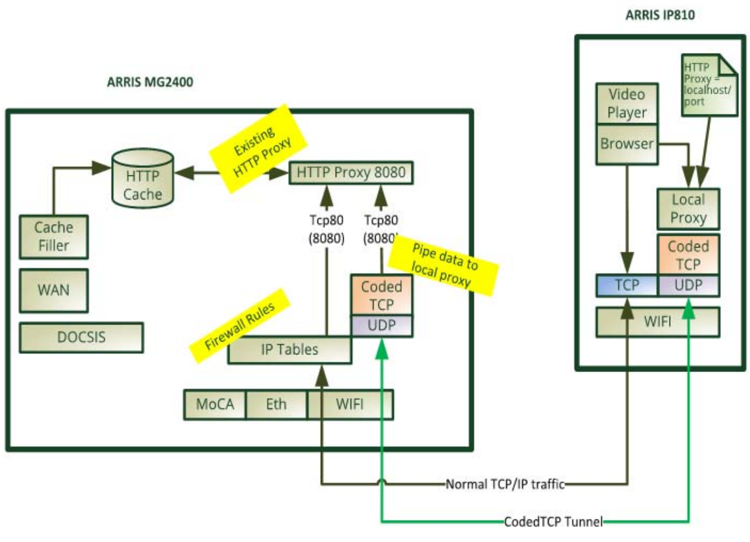

Figure 27 - Gateway and STB enabled Coded TCP architecture

Based on the previous graphs, it's easy to tell that TCP packet loss can have a big impact on the playout of video over TCP. The following examples of a 60second video clip being played out over TCP with multiple packet loss situations shows how well Coded TCP copes with the poor network conditions compared to a standard TCP implementation.

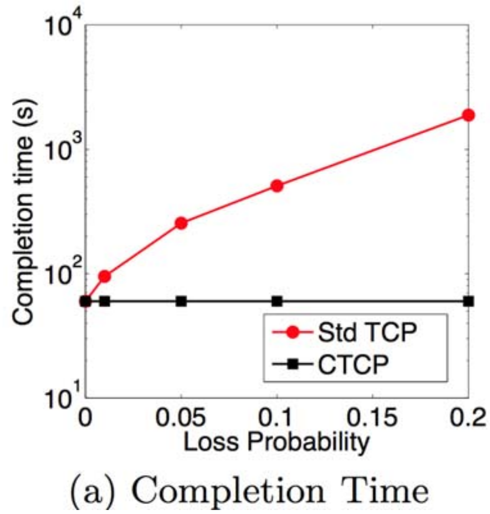

Figure 28 - 60 second video clip over TCP with coded and non-coded performance (2012-CTCP.pdf)[12]

The graph plots completion time of the 60 second playout compared to loss probability ranging. Coded TCP delivers a constant playout of 60s through the presented loss rates. Standard TCP however, even at 1% packet loss, increases the playout time from 60s up to 95s. Packet loss greater than 1% causes significantly longer playout. The jump for 60s to 95s means that the standard TCP playout is stuttering and stalling. Increasing to 5% packet loss shows a marked increase in playout time to 250seconds, and 10% packet loss slows to a 500second playout time.

The capability of Coded TCP and Random Linear Network Coding provide a way of delivering video over demanding wireless networks.

12 (Various images used with permission from 2012-CTCP.PDF, Network Coded TCP (CTCP) from Kim et. al, see bibliography)

Quality of Service for Quality of Experience

There are two key areas to target for making decisions on mapping packets into airtime when

Services_To_Deliver_Packets > Airtime_Capacity

These areas are: defining Service Priorities and Device priorities to ensure that the best packets get access to the airtime and maximizeing the user experience. This definition of what the ‘best’ packets to slot into the Wi-Fi airtime is something that can be defined logically for both the services and the devices in all of our homes. The philosophy here is to leverage the typical home service and device usage behavior. Here are some of the premises that this approach leverages:

- Customer tolerance to mobile device Wi-Fi poorer performance is higher than static device performance
- Customer tolerance to smaller screen devices with problems is higher than with larger screen TV devices
- Bill payer experience is more important than non-bill payer
- Moms and Dads (Bill Payers) don’t worry if their kids’ mobile Wi-Fi devices have problems particularly in bedrooms and particularly when streaming video
- Edge of range devices should not be allowed to impact overall service levels and the customer made aware of extension solutions
- Some events like the purchase of PPV event or VoD movie can raise the temporal value of the device involved
- Some linear viewing events have much higher priorities than other services

So, for example, the following service level priorities could be established:

(1) Voice > (2) Managed video > (3) HSD

Within video service delivery, there are opportunities to provide another layer granularity to prioritizing different services. A sample scheme for doing this is illustrated below

Video Type	Priority	Bandwidth
PPV	1	
Multicast	2	
4K VoD	3	
HD VoD	4	
SD VoD	5	
Free VoD	6	-

Figure 29 - Sample Service Prioritization Scheme to map to Wi-Fi scheduler

Service Classification Information

The system must be given enough information to reliably identify packets for Quality of Service (QoS) treatment. For example, a table-based solution can be used, if the packets are tagged using DSCP.

DSCP Bits (hex)	QoS treatment	Comment
0x38	Highest priority, Low latency	AC_VO WMM
0x30	Other levels TBD	AC_VO WMM
0x28		AC_VI WMM
0x20		AC_VI WMM
0x18		AC_BE WMM
0x10		AC_BK WMM
0x08		AC_BK WMM
0x00		AC_BE WMM

Device prioritization is the second most important element of QoS mapping to airtime. All devices are not created equal in the home and we need to leverage the following premises to maximize best user experience:

- IP STBs connected to TVs are higher priority devices
- Tablets and particularly smartphones at the edge of the Wi-Fi network are low priority devices
- Devices on the guest network are lower priority than the private and video networks
- Some Internet of Things devices may have very high priority like sensors and security devices
- Parents may choose to lower the priority of kids' devices and in particular to force them to no traffic mode after certain times (Bed time)

There are certain temporal phenomena that also can be regulated by this device=based control:

- Edge of Network devices – particularly ones that are constantly at edge of network
- Device Swarms – Party events where tens of users connect to guest Wi-Fi network

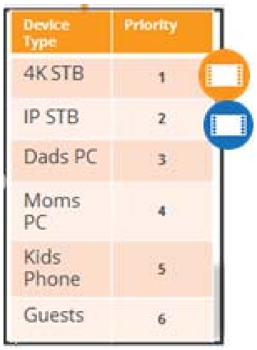

Figure 30 - Sample Device Prioritization Scheme to map to Wi-Fi scheduler

Device Classification Information

Traffic prioritization based on device type would be in addition to any system configuration optimization, such as preferential channel usage. These prioritization solutions could be done as opt-in scheme with users to ensure that user not the MSO provides the device priority. A simple scheme to allow the user to tag a device to a High, Medium and Low priority bucket can be implemented on the gateways web interface or even as a simple application on the TV screen.

Example Packet Classification

When a packet arrives from the HFC network, it can be classified into a queue for that traffic type either based upon an explicit packet tag or based on ACLs. The ACLs might use the Destination MAC address, or other information. If there are no applicable QoS indications for that packet, it gets best effort treatment.

For example, the following QoS table could be set up:

Indicator	QoS treatment	Comment
MAC Address – 001DCE79845	Strict priority, highest queue	IP STB - Video
VPN traffic	Strict Priority, second queue	User - VPN traffic
DSCP bits – 0x38,0x30	Strict Priority, second queue	Video traffic for non-STB
DSCP – 0x18, 0x00, 0x10,0x08	Best Effort queue	HSD traffic, non-priority
No Markings recognized	Best Effort queue	HSD traffic, non-priority

As packets come in, they are compared first to any ACLs or packet markings requiring priority treatment. After any strict priority traffic has been identified, the remaining traffic is queued up for best effort treatment.

Airtime Fairness

An area related to QoS is airtime fairness. This feature requires an AP to distribute airtime fairly across the attached clients. The intent is that clients who have a poor connection are limited in their ability to affect the airtime available to other clients.

Current schemes provide a mixture of airtime fairness and strict priority to satisfy the competing needs of dedicated services, such as IP STBs that require strict priority over other devices, and protecting the overall level of service provided to the entire ecosystem of clients subtending from an AP.

Community Wi-Fi Considerations

With the advent of Community Wi-Fi (hotspot) deployments, the airtime usage model for delivery of 4K video becomes even more complex. The channel overhead to manage additional service(s), and airtime demands of associated hotspot data must be considered to deliver reliable video. Areas of consideration must include:

Hotspot Service Management

The overhead for managing a Hotspot Service includes the transmission of management frames. Beacons, probes, probe responses, authentication (including 802.1x), association, and action frames are added for each service. Since these are high priority 802.11 frames, they must be considered and mitigated in a video deployment where such services are offered. Mitigation areas should include beacon bursting, not responding to probes when the service is completely utilized (client limiting) and/or adaptively enabling the service only when sufficient bandwidth is free to support Hotspot clients.

Hotspot Service Data

Associated Client Data – the amount of airtime consumed by associated hotspot clients must also be managed in order to provide acceptable hotspot services while not impacting video performance. Airtime consumption for hotspots can be effectively managed by limiting the number of hotspot clients allowed, via airtime management of hotspot data, and through tight control of hotspot clients.

Policy and Rules based interface to the Wi-Fi Airtime Scheduler

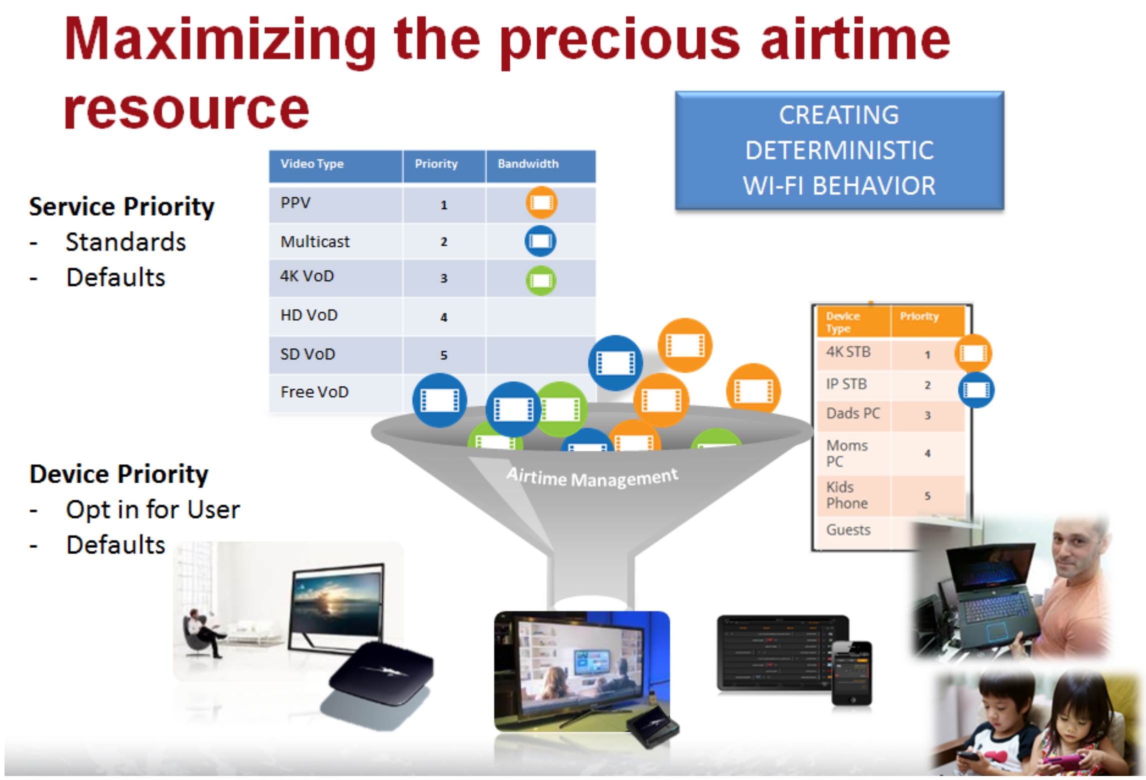

Figure 31 - Service and Device Prioritization mapped to Airtime Scheduler

Wi-Fi Controller Solutions

The concept of a Wi-Fi or Wireless LAN (WLAN) controller is certainly not a new one and indeed there are efforts already underway in Cablelabs to create standardization around RRM/SON solutions. Enterprise networks have long relied on WLAN controllers to control, secure and policy the massive spike in Bring Your Own Device (BYOD) user needs. The importance of signal quality and range in making VoW deterministic and the introduction of additional APs in the home – makes Wi-Fi controller solutions a potential important toolkit now in the home – and in particular a 2 stage Hybrid Controller approach of (i) the primary gateway managing the Range extensions devices and (ii) the network based Controller managing interesting domains of Wi-Fi. There are other papers that describe this approach in detail – so we will just touch on the basics here to acquaint the reader with this solution. .

The Hybrid Controller

To build the necessary Wi-Fi radio and client relationships, the creation of heterogeneous Wi-Fi networks for each and every subscriber is essential. To do in a residential deployment with its unique challenges and scaling becomes the hurdle which needs a unique solution.

The Hybrid Controller concept becomes the cornerstone of the building block to overcome the obstacles unique to these deployments. The key is to segment the RF layer control and the need for real and near-real time resource management such that the DOCSIS gateway in the premise is able to act autonomously on conditions it is closest to understand.

Specifically, the gateway is provided knowledge of other managed access points in a set of policies from the cloud controller to allow for semi-autonomous decision making by the gateway as Wi-Fi RF conditions continually change. This now changes the relationship in a Dual AP model to a Gateway-Child managed Wi-Fi network model managed locally and semi-autonomously in the home.

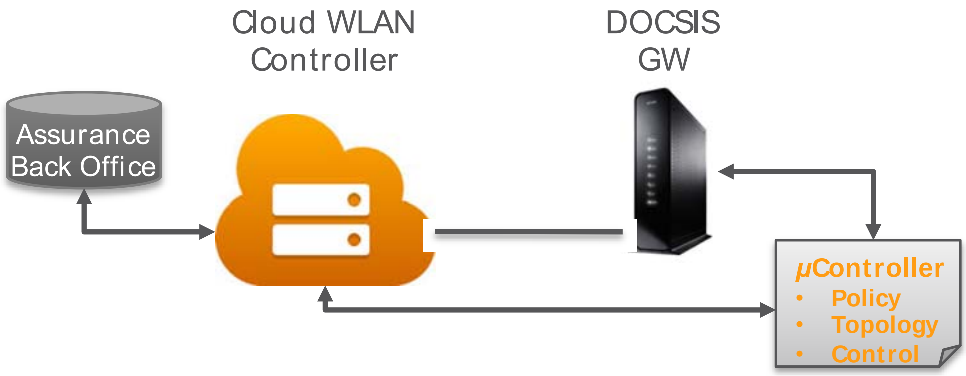

Figure 31 Hybrid Controller

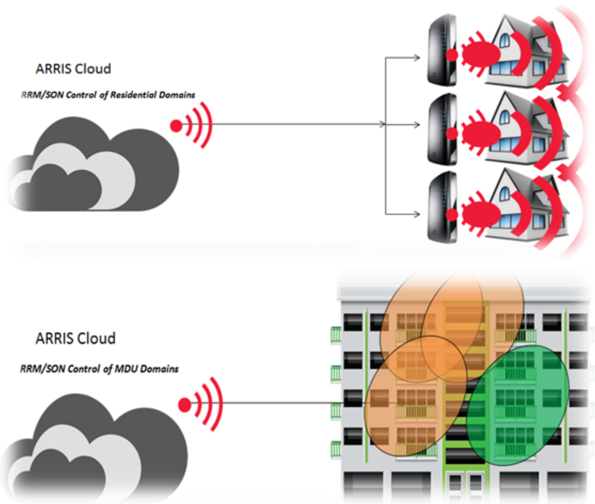

Figure 32 - Hybrid Controller

Through ongoing channel information and client statistics gathering observed both on the DOCSIS gateway and from managed in-home range-extending access points, the μController is able to determine the best client association to access point and whether RF power may need to be adjusted or RF channels need to be changed. Some of this information is derived through the use of 802.11k and 802.11v client reports and 802.11h control mechanisms. These components of the 802.11 standard enables Wi-Fi station information reporting and communication of radio topology and real-time assessment of localized channel conditions.

With the μController extending the DOCSIS gateway now into a home based Radio Resource Management (RRM) capable platform, the DOCSIS gateway embeds appropriate Information Elements (IE) into beacons. This signifies to clients the DOCSIS gateway access point is an RRM network element. Clients may now request neighbor information from the DOCSIS gateway, and the DOCSIS gateway may obtain information from compliant Wi-Fi clients.

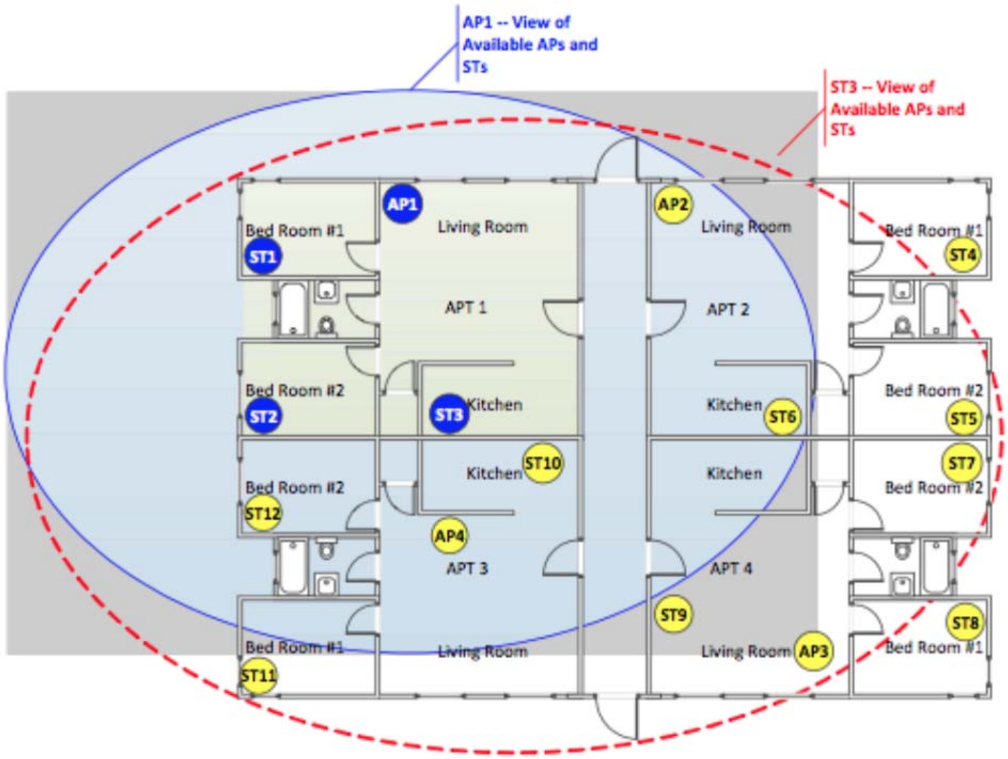

Figure 33 - Apartments - MDU Topology Mapping

*μ*Controller

WiFi DataBase -- APs and STs (as seen by AP1, or one of the ST in the Home Network)

Station Statistics

SSID	BSSID	RxRCPI	RxRSNI	FrmCnt	Av RCPI	Lst RCPI	Lst RSNI	Load	CCA	TxRCPI	TxRSNI	RtyCnt	MRtyCnt	FrmDupCnt	RTSFailCnt	ACKFailCnt
APT_2	01:02:03:04:02:00 (AP2)	-50	xyz													
APT_3	01:02:03:04:03:00 (AP2)	-60	xyz													
APT_4	01:02:03:04:04:00 (AP2)	-60	xyz													
APT_1	01:02:03:04:01:01 (ST1)	-40	xyz	5000	abc	def	ghi	5%	5%	-40	pdq	100	10	7	52	8
APT_1	01:02:03:04:01:02 (ST2)	-50	xyz	10000	abc	def	ghi	10%	10%	-50	pdq	100	20	24	152	56
APT_1	01:02:03:04:01:03 (ST3)	-60	xyz	50000	abc	def	ghi	50%	50%	-60	pdq	2345	171	321	1452	456

Figure 34 - μController Station Statistics

This architecture also provides a method to 'self-heal' a home network if any of the links in a multi-in-home AP deployment becomes unusable. The client device is supplied with the information on where to go and be 'told' to move using an 802.11r fast transition message. Between these two steps, the *μ*Controller has exchanged any 802.11i temporal key information such that when the client associates to the new access point, no re-authentication delay is encountered.

Having built a local Wi-Fi topology, the μController in the gateway is able to communicate the neighbor list of access points, current RF channel state and client states learned to the Cloud Controller.

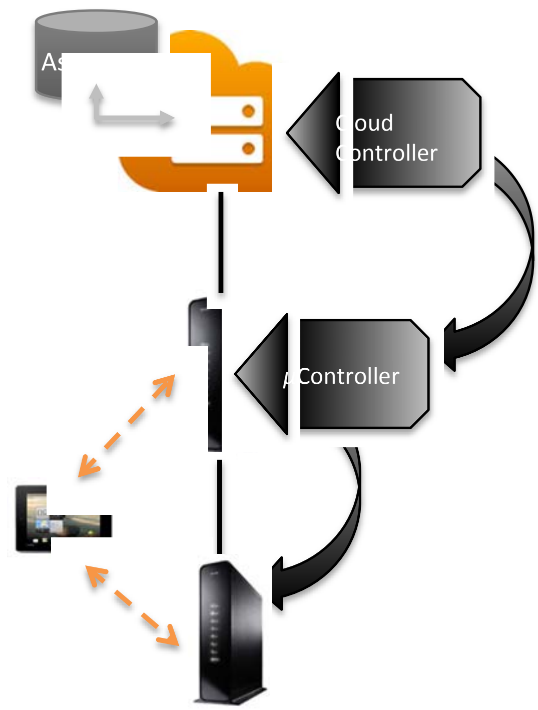

Figure 35 - µController - Cloud Controller Relationship

As these direct and indirect topology objects are learned, they are sent to the WLAN Cloud Controller for analysis. The Cloud Controller has the role of defining policy and controls for the μControllers, and to correlate received topology data to manage Wi-Fi quality across multiple customer premise networks.

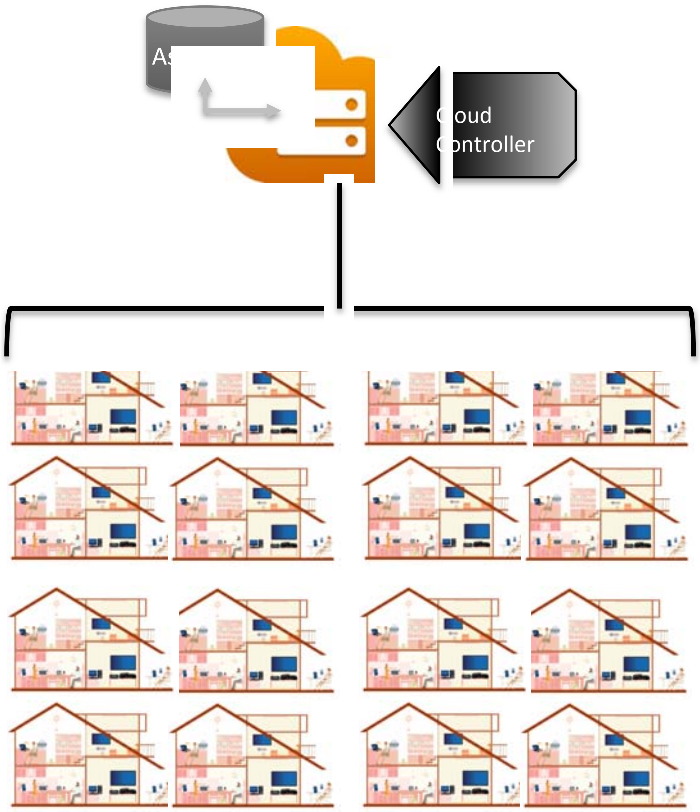

Figure 36 - Control across networks

Adaptive Bit Rate Solutions for both Multicast and Unicast IP Video

Finally, the last consideration for the deterministic VoW toolkit is Adaptive Bit Rate (ABR) and its application to both multicast and unicast video. ABR should be leveraged as the final fallback ensuring that IP video streams for 4K transmissions do not buffer and destroy user experience.

As described in an earlier section, when the delivery system cannot deliver the desired profile, a client can get a lower Video profile feed to keep the client buffer playing video. The upscaling abilities of the STB and TV can often provide enough compensation to make the reduced resolution indiscernible to the average user – maintaining their level of enjoyment in the programming.

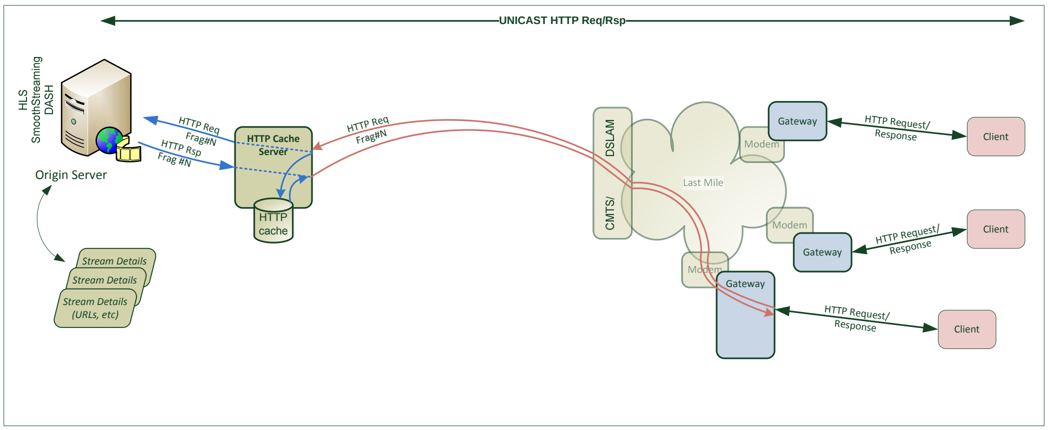

Figure 37 - Unicast Delivery of IP ABR Video

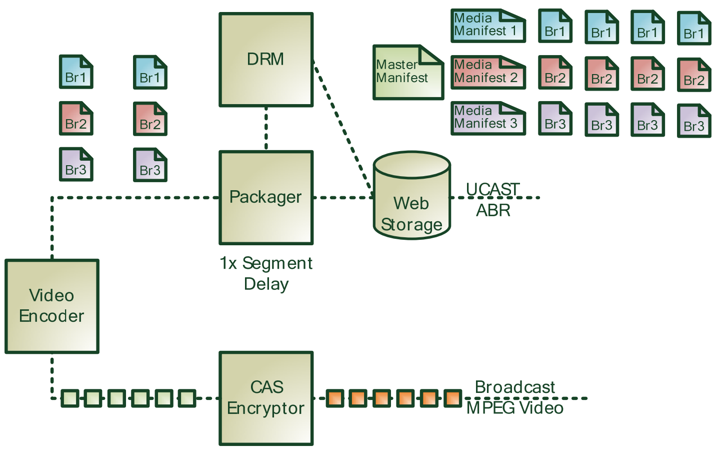

Figure 38 - Simple ABR Streaming Architecture

Summary

Wi-Fi is inherently non-deterministic in nature but is the key connection media of our future home. While the Hardware elements of the Core Wi-Fi solution are key to the fundamental quality of Wi-Fi – to ensure that we can control Wi-Fi airtime at all times within home service delivery model – there are additional features that need to be added to ensure the same consistent deterministic behavior within different Wi-Fi airtime budgets. This paper has tried to outline some of these features and architectural changes that can make the difference on consistent utilization of airtime. It has suggested that additional features particularly around the areas of Service and Device level priorities be employed to make sure that there is a precise determinism to end client quality at all times in Wi-Fi airtime budgets. Additionally, the ability to leverage coding solutions on the TCP Wi-Fi network may also add enough capability in marginal Wi-Fi conditions to keep the home a buffer free zone for 4K video.

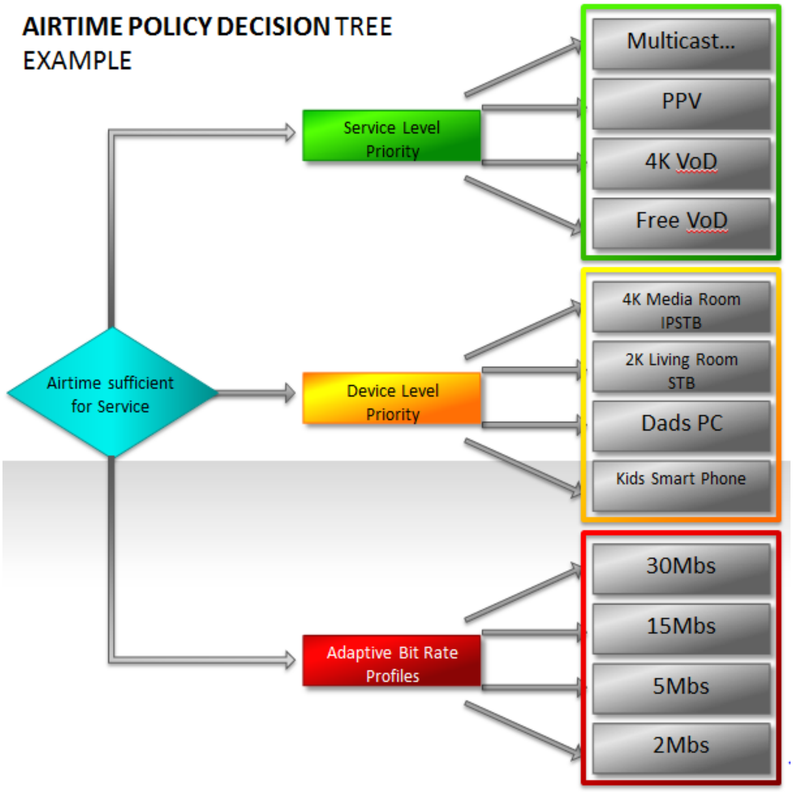

Figure 39 - Airtime Decision Tree – Example

As all devices are not created equal and all services command different user experiences, we propose that this tiered service/device priority decision tree should be leveraged to get the overall Mean Time Between Failures of non-buffered 4K Video streams to the expected levels of a pay for TV service (Figure 36). Adding additional features like Wi-Fi control of overlapping Wi-Fi domains will also improve chances of delivery in MDU environments in particular and in multi AP single home solutions. Finally, Adaptive Bit Rate solutions were designed to help balance quality with buffering for higher bandwidth video flows – and with Wi-Fi the typical culprit for congestion on a managed Video over IP network- ABR will be the final fallback and a key part of the decision tree for ensuring 60fps video to 4K TV's. Here's to the VoW experience and keeping that 4K Video tamed in the home.

Bibliography

1 Network Coding Meets TCP: Theory and Implementation, March 2011
Jay Kumar Sundararajan, Devavrat Shah, Muriel M´edard, Fellow IEEE, Szymon Jakubczak, Michael Mitzenmacher, and Jo˜ao Barros

2 Network Coded TCP (CTCP), ACM Mobicom 2013,
MinJi Kim∗, Jason Cloud∗, Ali ParandehGheibi∗, Leonardo Urbina∗, Kerim Fouli∗, Douglas Leith†, Muriel Médard∗ (MIT/NUIM)

3 TCP-Aware Forward Error Correction for Wireless Networks, MobiLight 2010, LINCST 45,
Dzmitry Kliazovich, Magda Bendazzoli, and Fabrizio Granelli, DISI - University of Trento

4 TCP-Cognizant Adaptive Forward Error Correction in Wireless Networks, IEEE INFOCOM 2002,
Benyuan Liu, Dennis L. Goeckel, Don Towsley, University of Massachusetts

5 A Comparison of Mechanisms for Improving TCP Performance over Wireless Links, SIGCOMM 96,
Computer Science Division, University of California at Berkeley

6 The Future of Wi-Fi Service Management – Hybrid Controller Approach
Chris Busch, Vishal Dhruv, ARRIS – SCTE 2014 Paper

6 Cablelabs RRM/SON Specifications

- WR-SP-WiFi-GW-I02-120216
- WR-SP-WiFi_MGMT-I04-140311
- WR-SP-WiFi-ROAM-I03-140311

Abbreviations & Acronyms

ABR	Adaptive Bit Rate protocol
ACK	ACKnowledgement
ACL	Access Control List
AL	Application Layer
A-MPDU	Aggregated MAC Protocol Data Unit
AP	Access Point
AVC	Advanced Video Codec, H.264/MPEG-4 Part 10
BYOD	Bring Your Own Device
CCA	Clear Channel Assessment
CDN	Content Distribution Network
DASH	Dynamic Streaming over HTTP
DBC	Dual Band Concurrent radios
DFS	Dynamic Frequency Selection
DLNA	Digital Living Network Alliance
DOCSIS	Data Over Cable Service Interface Specification
DSCP	DiffServ Code Point
DVR	Digital Video Recorder
FEC	Forward Error Correction
GW	GateWay
HD	High Definition video resolution
HDS	HTTP Dynamic Streaming
HE-AAC	High Efficiency Advanced Audio Coding
HEVC	High Efficiency Video Codec
HFC	Hybrid Fiber Coax outside plant
HLS	HTTP Live Streaming
HSD	High Speed Data
HTTP	HyperText Transfer Protocol
IP	Internet Protocol
ISP	Internet Service Provider
MAC	Medium Access Control
MCS	Modulation and Coding Scheme
MDU	Multiple Dwelling Unit
MIMO	Multiple Input Multiple Output
MPDU	MAC Protocol Data Unit
MSO	Multiple System Operator
MTU	Maximum Transmission Unit (Packet Size)
OFDM	Orthogonal Frequency Division Multiplexing
OTT	Over The Top video service

PAR	Peak to Average Ratio
PC	Personal Computer
PHY	Physical layer
RF	Radio Frequency
RLC	Random Linear Coding
RTP	Real-time Transport Protocol
RTT	Round Trip Time
SSID	Service Set IDentifier
QAM	Quadrature Amplitude Modulation
QoE	Quality of Experience
QoS	Quality of Service
TCP	Transmission Control Protocol
TV	Television
UDP	User Datagram Protocol
VBV	Video Buffer Verifying MPEG buffer model
VoW	Video over Wi-Fi
VPN	Virtual Private Network
WLAN	Wireless Local Area Network