Unmanaged ABR:

How to Control those Unruly Teenagers

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

John Ulm
Fellow of Technical Staff
ARRIS
john.ulm@arrisi.com

Amit Eshet
VP, Video Solutions Strategy
ARRIS
Amit.Eshet@arrisi.com

Nir Radian
Sr. Mgr., Video Algorithm
ARRIS
Nir.Radian@arrisi.com

Abstract

As a parent, we've often learned that the last resort is to take the car keys away from those unruly teenagers that we're struggling to control. Operators will be finding themselves in a similar situation as they try to scale managed video services over an Adaptive Bit Rate (ABR) solution. ABR protocols have become the mainstay of multi-screen devices like tablets, smart phones, game controllers and Smart TVs. This will force ABR to become part of an operator's managed video service as well.

This paper details our continued research into Adaptive protocols. Pure unmanaged ABR creates several key issues that impact Quality of Experience (QoE). These include system issues around instability, unfairness and lost capacity. All of these take a toll on QoE, which we quantify with measured Video Quality (VQ) data and other lab results.

Industry research is diving deeper into these issues. We took a closer look at several options in our research labs including our approach to this evolution called Managed Adaptive Streaming (MAS). The paper details the results found.

With MAS, the operator regains control to provide a first rate video service with exceptional Quality of Experience while retaining the underlying benefits of Adaptive protocols. Operators can significantly increase their IP Video capacity while gracefully handling congestion and improving user experience. Managed ABR will quickly become the future of IP Video.

INTRODUCTION

Adaptive Bit Rate (ABR) protocols have become the mainstay of multi-screen devices like tablets, smart phones, gaming devices and Smart TVs for accessing Over-The-Top (OTT) video content. Because of their explosive popularity, it is highly desirable for an operator to provide existing video services to these devices. However, consumers will expect the same Quality of Experience (QoE) to which they are accustomed with today's primary TV screen delivery.

The ABR protocols have been optimized to operate over an erratic internet connection. However, the ABR client based control with no insights into system behaviors has demonstrated many inappropriate behaviors. The ABR client's greedy behavior leads to significant unfairness, instability and inefficiencies. These traits are not suitable for an operator to offer a true managed ABR video service with the associated QoE.

Operators have a number of challenges in offering a "managed" ABR video service. The limited bandwidth makes it challenging to support the demand of a large number of concurrent users while maintaining good video quality for each user. In addition, existing implementations of ABR client controlled distribution mechanisms are not very efficient. They tend to under utilize the available bandwidth and provide uneven visual qualities to the clients. Therefore, understanding the issues around delivery of ABR will be crucial for MSO's video service delivery, and for their ongoing profitability.

Research into ABR has led to the evolution of Managed Adaptive Streaming (MAS) or sometimes also referred to as Smart ABR (SABR). By adding some cloud based intelligence back into the system, the operator can regain control to provide a first rate video service with exceptional Quality of Experience while retaining the key underlying benefits of Adaptive protocols. MAS/SABR is a server controlled system that can manage the bit rates and video quality that each client receives.

Intelligence in the MAS server maintains client state, available client download bandwidth or channel capacity, and a measure of "reasonable" client video quality. "Reasonable" could be dependent on client attributes such as client display size or the type of video content being watched. Based on that intelligence, the MAS server controls what bit rate each client gets. This can prevent oscillations in network bandwidth utilization and increase network utilization.

The paper discusses lab results showing the optimization and Video Quality achieved. The MAS system is compared in detail to traditional unmanaged ABR delivery. With MAS, operators can significantly increase their IP Video capacity while gracefully handling congestion and providing an improved user experience.

ABR Overview

Adaptive Bit Rate (ABR) is a delivery method for streaming video over IP. Adaptive streaming uses HTTP as the transport for small video chunks of approximately 2-10 seconds each. This enables the content to easily traverse firewalls, and the system scales exceptionally well as it leverages traditional HTTP caching mechanisms in the CDN.

Adaptive streaming was developed for video distribution over the Internet. In order to deal with the unpredictable performance characteristics typical of this environment, ABR includes the ability to switch at chunk boundaries between different encodings of the same content. This is illustrated in Figure 1. Depending upon available bandwidth, an ABR client can choose the optimum encoding to maximize the user experience.

The server stores several chunk sizes for each segment in time. The ABR client predicts the available bandwidth and requests the best chunk size using the appropriate URI. Since the ABR client is controlling when the content is requested, this is seen as a client-pull mechanism, compared to traditional streaming where the server pushes the content. Using URIs to create the playlist enables very simple client devices using web browser-type interfaces. A more in-depth discussion of ABR video delivery can be found in [ULM2010].

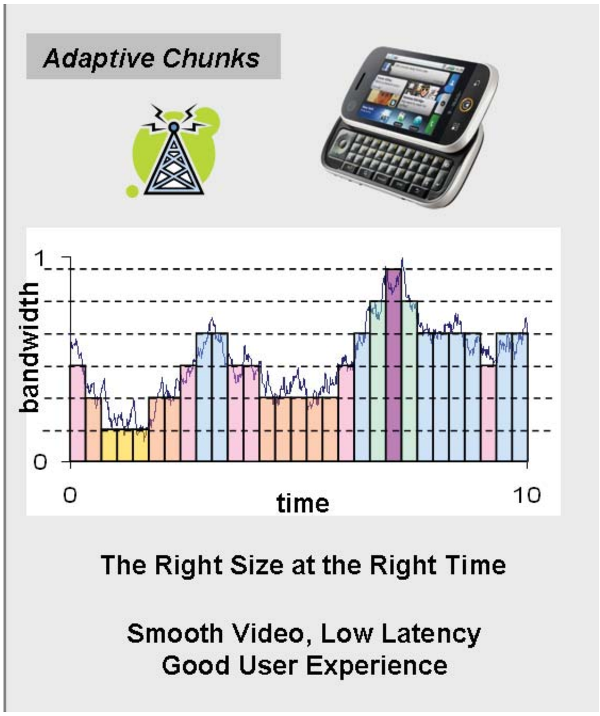

Figure 1 – Adaptive Streaming Basics

Importance of ABR – 2nd and 3rd Screens

ABR based video streaming has become the de-facto standard for video delivery to IP devices such as PCs, tablets, smart-phones, gaming devices and Smart TVs. ABR clients are typically shipped with (or are available for download to) these devices as soon as they are released. Given the short lifetime of this class of device this is a key enabler, especially compared to the time required to deploy software to traditional cable STB devices. As mentioned previously, ABR delivery simply requires an HTTP connection with sufficient bandwidth so that it is available both on net and off net. With these advantages, essentially all video delivery to second and third screen devices uses this mechanism.

ABR vs. Current Managed Video Delivery

ABR video delivery has a number of very significant differences to both MPEG video delivery and streamed IP video delivered over Real-time Transport Protocol/User Datagram Protocol (RTP/UDP) as used in a Telco TV system. Foremost, ABR has been developed to operate autonomously over an unmanaged generic IP network.

- *The client device decides on bit rate (i.e. bandwidth) decisions based on its interpretation of network conditions.*

This is fundamentally different from the approaches used for existing MPEG or conventional streamed UDP video delivery, where devices under the direct control of the network operator make the important decisions relating to bandwidth. Thus, in

MPEG delivery, the encoding, statistical multiplexing and streaming devices determine the bit rate for a given video stream. These devices are under control of the service provider. In contrast, the behavior of ABR clients is specified by the CPE developer which, in general, will be a third party outside the service provider's control.

An ABR client selects a file chunk with a bit rate that it believes to be most appropriate based on a number of factors including network congestion (as perceived by the client) and the depth of its playback buffer.

- *Thus the load presented to the network can fluctuate dramatically.*

Operators in a controlled network can guarantee that adding new user sessions do not impact existing users. Once resources are exhausted, any additional session requests will be denied, introducing a probability of blocking into the system.

ABR clients join and leave the network as users start and stop applications. From a network perspective, there is no concept of a session with reserved resources or admission control. Again this is the antithesis of MPEG or UDP video in which the control plane operates to request and reserve network resources and determines when to admit users. In a pure ABR model with network congestion, each new session will reduce the bandwidth available to all existing sessions rather than be denied.

- *Thus, all users may see a variation in video quality as other ABR clients start or change bit rates.*

With MPEG or UDP streaming video delivery, congestion control is not relevant as the control plane provides admission control to ensure it does not occur. When ABR is used for video delivery, congestion control is a potential issue. The situation is complex in that three levels of congestion control mechanisms are involved operating at different layers in the protocol stack. At the media access control (MAC) level, the CMTS is responsible for scheduling downstream DOCSIS traffic. Operating at the transport level is standard Transmission Control Protocol (TCP) flow control based on window sizes and ACKs. Finally, at the application level the client can select the video bit rate to request. The latter two levels of control (TCP and application) are the responsibility of the ABR clients and as such are outside the control of the network operator. Interaction between these flow control mechanisms is not well documented and may have unforeseen impacts.

In summary, ABR clients base their decisions on what to request based on their local knowledge and observed conditions rather than on an overall view of the network conditions. This is in contrast to MPEG or UDP streaming where the network operator provisions the video bit rates based on knowledge of the end-to-end network and expected loads.

Potential Managed ABR Solutions

Smart ABR (SABR)

An approach to managed ABR solutions called Smart ABR (SABR) was previously discussed in [ULM2013]. This paper highlighted lab results from a Motorola Mobility Advanced Technology team that showed some significant issues with unmanaged ABR including instability, unfairness and inefficiencies. The paper goes on to show how a Smart ABR approach can alleviate these challenges. Here are some of the highlights from that paper:

Operators have a number of challenges in offering a “managed” ABR video service. The limited bandwidth makes it challenging to support the demand of a large number of concurrent users while maintaining good video quality for each user. In addition, existing ABR client controlled distribution mechanism creates issues with QoE, network utilization, fairness and system stability. Our lab results presented confirms all of these conditions. In fact, fairness was an issue in multiple areas such as screen resolutions (HD vs. VGA), client software revision, shared home bandwidth and even the timing of when each client starts. While some clients were stable, others oscillated between bit rates and suffered video buffer under-runs.

A second series of tests were run leveraging enhanced CMTS QoS to try and fix the above issues. The results show only marginal improvements over the unmanaged ABR tests and the underlying issues of unfairness and instability still exists. Even had this approach fixed the problems, there are still concerns that the control plane could scale to support this solution.

Finally, our tests focused on a cloud based approach called Smart ABR (SABR) where the operator can regain control to provide a first rate video service with exceptional Quality of Experience while retaining the underlying benefits of Adaptive protocols. Our lab results show that all major shortcomings of unmanaged ABR can be addressed. The system becomes stable, network utilization improves and video quality fairness is provided across the entire client population while gracefully handling congestion and providing an improved user experience. Further tests suggest that SABR delivered comparable video quality in almost half the bandwidth while providing fairness and stability.

Managed Adaptive Streaming (MAS)

Half way around the globe from the SABR team, another R+D team was independently investigating a very similar approach called Managed Adaptive Streaming (MAS). MAS is also a method of delivery of ABR video in which the cloud based server, and not the client, decides on the best version of a video asset to download at any time under changing network conditions.

The MAS Server is responsible for delivering video to multiple clients. It bases its decisions on two main factors:

- Video quality of each segment of video delivered to each client

- Available bandwidth for all clients.

The MAS server acts as both a web server for video segment requests and as a controlling agent. Specifically, it receives requests from clients for video segments of different assets at different times, and overrides the client's choice of which bit rate version to download, and delivers that video segment in a controlled bit rate.

It is aware of the following:

- All requests for segments at any given time by all clients
- The available total channel bandwidth
- Information regarding the requested video
    - Segment sizes
    - Segment quality

In conventional ABR video distribution, the ABR client determines the bit rate of the next file to download from the options in the playlist and retrieves this directly from the Content Delivery Network (CDN). By adding intelligence into the cloud, this decision could potentially be overridden from the network in a number of ways. This is referred to as MAS.

The playlist file provides the bit rate options specified by the service provider. Normally this selection would be statically provisioned and implemented by the encoding and packaging processes as the video asset was processed. It is conceivable that the playlist may be manipulated and the network can regain control of what bit rates are available to the ABR client.

During peak utilization, existing managed video delivery uses admission control to block new users from accessing the system. With MAS, an operator could gracefully handle congestion during peak times with no blocking of users, but rather a slight degradation of video quality. This reduction in quality during peak times is analogous to statistical multiplexing in legacy MPEG video. During peak times, the stat mux reduces bit rates across the various video streams to fit within its channel. The MAS system has an advantage in that it will be multiplexing over a larger channel using DOCSIS bonding.

In a MAS system, all the clients are controlled from the server side. The system level intelligence in the server understands the state for every client; the available bandwidth for each client; and a "reasonable" visual quality of the video for a given size of display and attributes of video etc. Based on that intelligence the server controls what bit rate each client gets. This avoids the oscillations and increases network utilization.

MSOs can leverage their managed network to greatly optimize the scarce and costly bandwidth resources and dramatically improve the subscriber QoE without allocation of extra bandwidth. How? Control has to move back to the network to ensure an intelligent allocation of bandwidth to the various sessions, as well as optimization of the QoE.

For the managed ABR approach to be effective, it should leverage the following information:

1. Network topology, including available bandwidth for multi-screen video service;
2. Session information, including association of subscriber to a SG, device type, and service level agreement (SLA);
3. Content information, including format (e.g., SD/ HD), quality information.

An ABR Proxy entity is added to the delivery flow. The role of the proxy is to manage the transmission of the ABR traffic. To explain the behavior of the proxy, let's focus on a single DOCSIS SG. The proxy is statically configured with the bandwidth available to the managed multi-screen video service on that SG. When a client initiates a session, the proxy is configured with information on the subscriber, the device type (e.g., iPad), and the content (e.g., available versions). Upon session initiation or termination, the proxy will re-divide the aggregated bandwidth between the active sessions in a way that equalizes the QoE.

Defining QoE per session can be made using either static or dynamic methods. The static method relies on information like resolution, bit rate, and quality grade for each of the versions available for the content, as well as the device type. The dynamic method is based on QoE grade associated with each segment. When the dynamic method is used, the proxy can alter the bandwidth allocated to each session every segment. Quality information can be created as part of the process of content generation. This can take place either in real-time (for live content) or offline.

The functionality of the ABR proxy can be introduced in several ways:

- It can be integrated into the CMTS or be provided on a standalone platform.
- In case of a standalone platform, the ABR proxy can be used to control the bit rate of the CMTS's service flows, or to multiplex the ABR traffic guaranteeing it fits into the allocated bandwidth. Multiplexing of the ABR traffic allows for expediting play-out start as well as improved mitigation for interference in the home wireless network.
- It can also be used to override the requests coming from the clients. This allows for further optimization of the bandwidth utilization by mitigating the conservative behavior of the client.

Regardless of the implementation the ABR proxy meets the following critical requirements:

1. Support for standard ABR clients;
2. Seamless integration into the multi-screen video architecture;
3. Scalability – the capacity of the proxy can easily grow with the service;
4. Cost-effectiveness – the proxy provides a straightforward and effective solution for bandwidth and QoE optimization.

It was discussed how HTTP ABR paves the path for MSOs to offer a video streaming multi-screen service. However, this comes at the price of bandwidth and QoE inefficiencies. To overcome these, MSOs can augment their CDN with an ABR proxy, thus offering a differentiated managed service optimizing spectrum, as well as the subscriber's experience on a per screen and content basis.

Unlike the unmanaged ABR system, in which the TCP congestion control makes sure that all clients receive equal bandwidth, the MAS server uses the video quality data to make sure that all clients receive equal video quality, which is not necessarily equal bandwidth. As part of the control mechanism, the server also makes sure that the available bandwidth is divided in such a way that no client reaches low levels of playback buffer, and all clients are equally robust to the turmoil of the network.

MAS Test Methodology

Our goals were to research ABR behaviors in a working environment. First item was to replicate and quantify existing unmanaged ABR characteristics. Then test these same conditions for a cloud based MAS Server solution. Previous work [CLOONAN] discussed results from a simulator. Our goal was to capture lab based results with live client interaction. The MAS testing continues from the SABR testing. While SABR testing was for a limited number of clients, the MAS tests focused on scaling the clients to a much larger network capacity. Emphasis was kept on steady state operation while the system was relatively stable.

Experiment Setup

To scale the tests to a large number of clients, a simulator was built that could emulate dozens or even hundreds of clients. The lab setup is shown below:

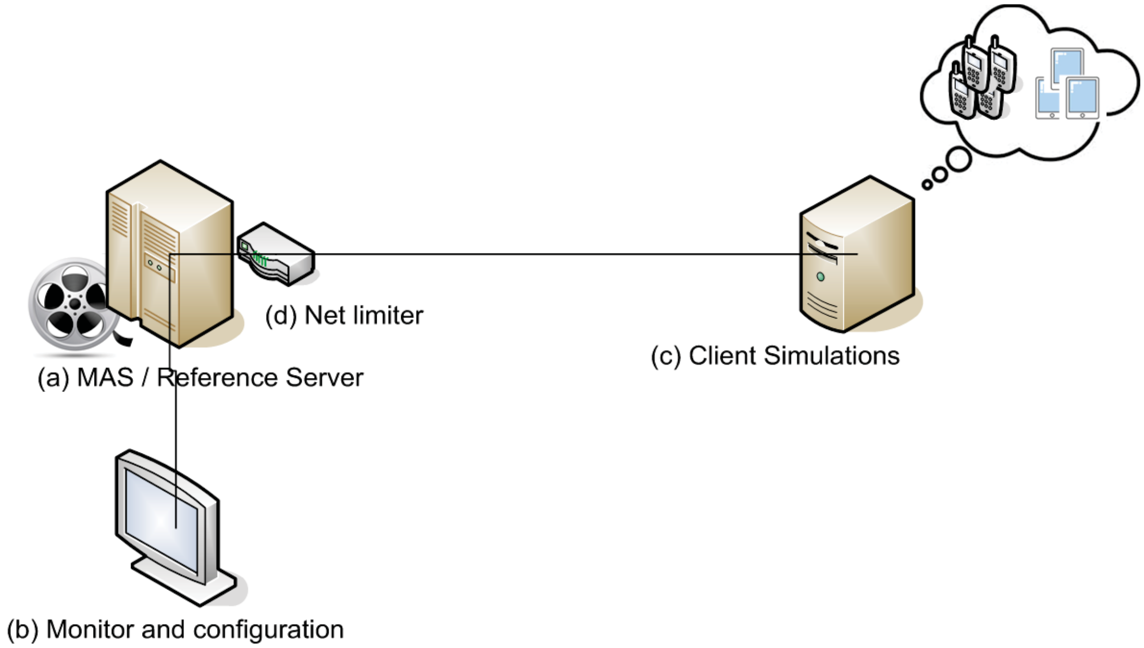

Figure 2 – MAS Test Experiment Setup

The experiment includes a Linux based web server that can behave either as a reference unmanaged ABR system or as a managed system (MAS). This server has access to a database containing a set of assets that are described below. This server is configured and monitored via a remote terminal. An additional windows process runs on a separate machine and serves as a client simulation which can generate requests for ABR video for multiple clients. The Linux NetEm kernel module limits the output bandwidth of the web server so the tests can mimic various network capacities.

Content

A set of 11 short movie clips (3 – 6 minutes long) has been chosen for this experiment. Movies differ in complexity (i.e. sports, nature scenes, news broadcast).

The following table shows the profiles chosen for this experiment:

Bit Rate (Kbps)	Resolution
360	512x288
1000	640x360
1500	768x432
2800	1280x720

Table 1 – Available Bit Rates

Devices

A client simulator tool was used to model a standard HLS video client.

3 different devices were used to model client behavior:

- Hand-held (smartphone, typically 4 inch screen, short viewing distance)
- Tablet (typically 9-10 inch screen, short viewing distance)
- Large screen HDTV (typically 50 inch screen, long viewing distance)

Hand-held devices were allowed to download only the first two versions (max resolution of 640x360) while the tablets and TVs were allowed to download all 4 versions.

The percentage of each device from the total lineup is summarized in the table below.

Device	Percentage	Max Bit Rate
HDTV	50%	2800 Kbps
Tablet	40%	2800 Kbps
Hand-held	10%	1000 Kbps

Table 2 – Client Device Distribution, Max Bit Rate

Video Quality Measurements

Each movie was segmented into 2 second segments. Each segment was given a video quality grade by the VQ objective estimation algorithm between 0 – 100. The following table expands the grades into a subjective human experience.

VQ Range	Quality	Impairment
85-100	Excellent	Imperceptible
70-85	Good	Perceptible but not annoying
55-70	Fair	Slightly annoying
40-55	Poor	Annoying
0-40	Bad	Very annoying
Video Freeze	Terrible	Unacceptable

Table 3 – Video Quality Metrics

The four bit rate versions of all movies have segments within this range, but it is important to note that typically the best version is **Good** or **Excellent** barring a few segments in which the best version is only **Fair**.

Before being available to the web server, the video is ingested, and a VQ (Video Quality) grade is produced for each segment of video in each version. This quality represents the complexity of the video as well as other encoding artifacts that may be present and reduce the viewer's experience. To summarize: the server has knowledge of the estimated video quality that the user will perceive for each segment of video at each version, for each client that is downloading from the web server, taking into account the device type (hand-held, tablet or HDTV).

With this knowledge, the managed server (MAS) decides on the best quality that can be viewed by the system, given the amount of clients and bandwidth constraints. If more clients join, the target quality is decreased. If clients leave, then bandwidth is freed and the remaining clients can enjoy better quality.

Result Graphs

For each test scenario, several graphs were produced to help demonstrate the QoE delivered. These graphs include:

- Video Quality vs. Time: This graph depicts the minimum video quality between all clients at every second interval as well as the 90th percentile at every second interval. The 90th is a measurement of the video quality that is being experienced by the majority of clients, hold a few clients which may either be suffering from momentary buffering or are simply viewing a movie segment with inherently low video quality.
- Video Quality histogram – this graph depicts the dispersion of quality among all the clients throughout the entire experiment.
- Playback Buffer Fullness histogram – this graph depicts the amount of video playback at the client buffer in all clients throughout the experiment. Samples in the “1000 msec” bin point to severely low buffer levels which translates to cessation in playback.
- Percentage of Clients vs. Buffer Fullness – this histogram shows the percentage of clients which had at any point in the experiment below the category buffer fullness (for example, 30% of clients had at some point in the experiment less than 1 second of playback which – this will cause the video to freeze at the clients device)

Note: in each experiment, graphs do not include an initial “start-up time” which takes the system to reach relative stability. A complete set of these results are available in the Appendix.

MAS Test Results

Unmanaged ABR – Reference Baseline

Our reference scenario consisted of an unmanaged ABR setup starting with 18 clients operating in a 40Mbps network. This is enough load to completely fill the network to capacity. The figure below provides insights into the video quality achieved.

The timeline provides two data points for each sample:

- 90th percentile VQ – point where 90% of segments exceed this VQ
  - o Should stay in the Good or Excellent VQ range
- Minimum VQ – lowest VQ measured for each time sample
  - o Should stay mostly in the Fair range, with a rare Poor occurrence for the small screen handheld device

The 90th percentile VQ gives an indication of overall video quality for the majority of video segments. In this test it drops into the Fair range a handful of times. The Minimum VQ shows the worst client. In this test there are still a number of video segments in the Bad range. The total VQ percentage for all video segments is shown in the lower left. It turns out that 1.3% of video segments fell into the Poor VQ range and 0.3% is in the Bad range.

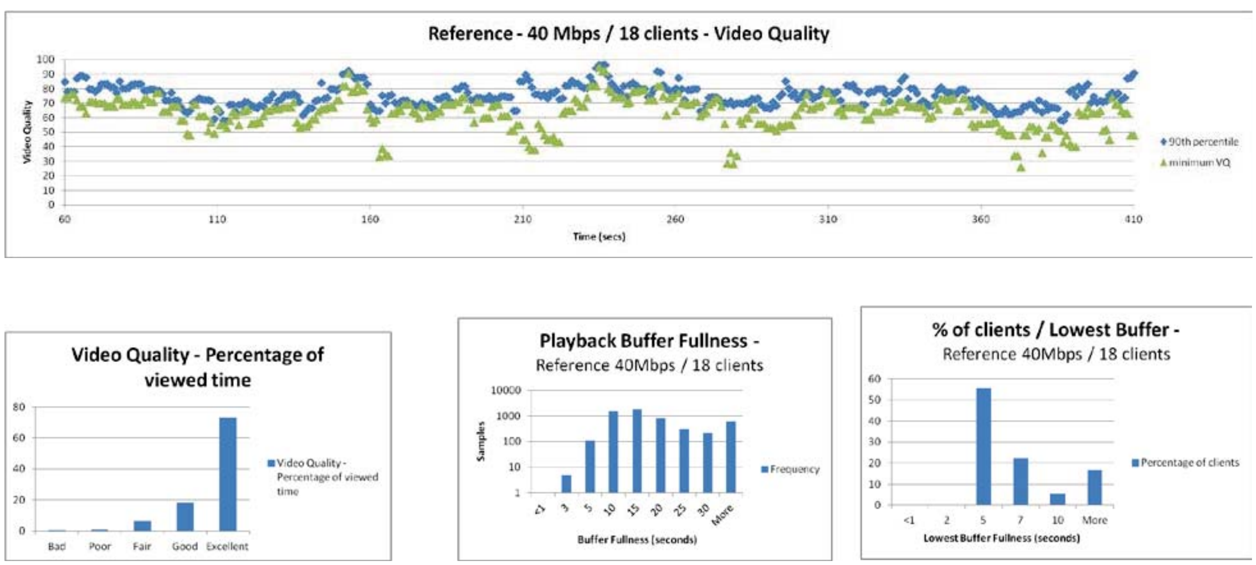

Figure 3 – Unmanaged ABR VQ: 40Mbps, 18 clients

The other two smaller bar charts above provide an insight into the client's playback buffer. If the buffer gets too low then we run the risk of a buffer under-run that will cause a frame freeze. This is much worst then receiving a video segment with poor or bad VQ. In this test, the playback buffer varies a good bit but there were no buffer under-runs. Most clients maintained at least a 5 second playback buffer. Overall, this QoE might be acceptable for Over The Top (OTT) content, but would be very marginal for an operator trying to offer a managed Video service.

For the next unmanaged ABR test, the number of clients was increased by almost 40% to 25 clients. The resulting data in the timeline graph shows that the system has become more volatile. Not only has the 90th percentile VQ dropped into the Fair range quite often, but the number of Poor video segments jumped to 3.0% and the Bad video segments increased fourfold to 1.1%.

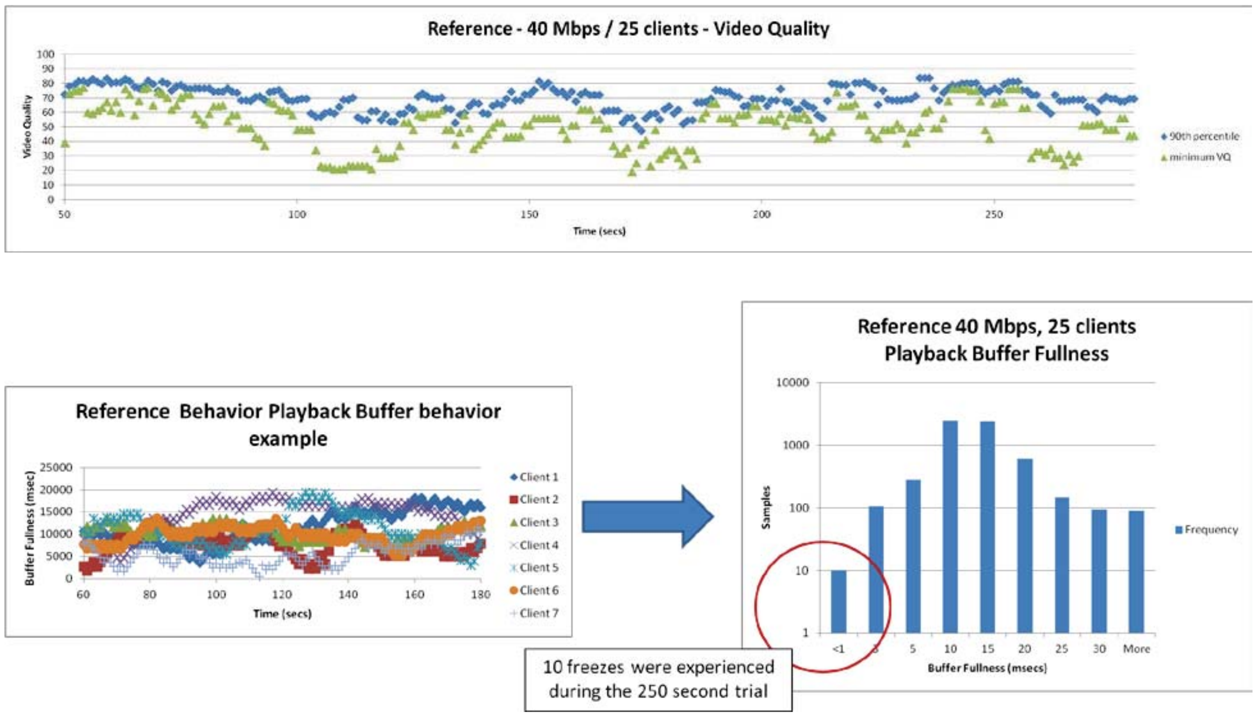

Figure 4 – Unmanaged ABR VQ: 40Mbps, 25 clients

Taking a closer look at the playback buffer shows much higher volatility in the playback buffer as well. This time, a total of 10 video freezes occurred across 7 clients during the 250 second trial. Thus 28% of the clients were impacted with video freezes. From an operator's perspective, the QoE has now become very unacceptable.

MAS vs. Unmanaged ABR

Starting from the above Reference baseline, the next test ran with the MAS server delivering to 25 clients. As can be seen in the time line graph below, the MAS test with 25 clients was much better than the unmanaged ABR test with only 18 clients. With MAS, the 90th percentile VQ stayed in the Good/Excellent ranges almost the entire time and there were absolutely no Bad video segments. The Poor video segments came in at 0.3% but this is attributed to a complex video segment being delivered to the bandwidth limited small screen handheld device.

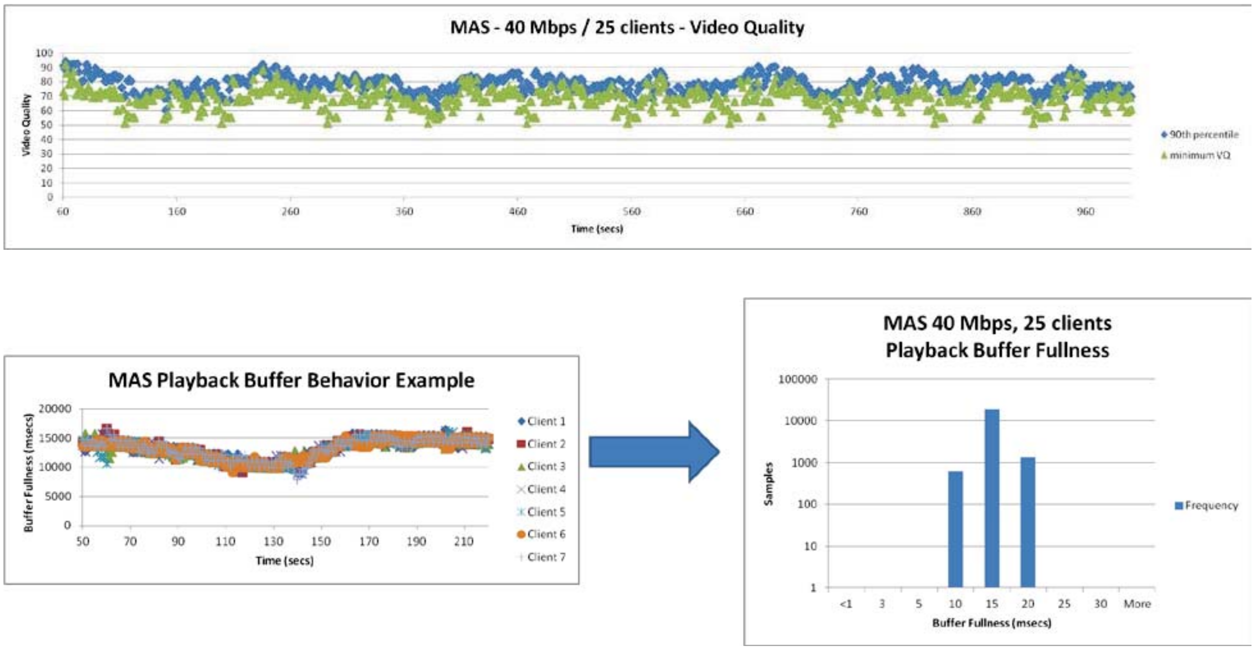

Figure 5 – MAS VQ: 40Mbps, 25 clients

The Playback buffer behavior for the MAS trial is even more impressive. The buffers remained rock solid with very little variation. This bodes very well for supporting live linear content where buffers need to be minimized.

Other MAS tests were run with increased clients to see if MAS had a breaking point. With 30 clients, the number of Fair video segments nudged up from 4% to 8%, but there was no change in Poor or Bad video segments. At 40 clients, there is a system wide degradation in video quality as the Fair video segments shot up to 32% while Poor segments nudged slightly to 0.9%. Meanwhile, there were still zero Bad segments and the playback buffer remained solid with zero video freezes. The MAS system with 40 clients was much better behaved than the unmanaged ABR system with 18 clients.

Scaling with Network Capacity

For the next series of tests, the network capacity was doubled to 80Mbps and then doubled again to 160Mbps. As the network capacity is doubled, so is the number of clients. Some folks have claimed that there is a statistical multiplexing gain that video will experience with the larger capacity network seen from more DOCSIS bonded channels. So, conventional wisdom would have that system behavior should improve for the same client to capacity ratio.

For unmanaged ABR, the behavior actually got slightly worse as the network capacity increases. This is seen below on the left. It turns out that 100 clients @ 160Mbps are higher than both 50 clients @ 80Mbps and 25 clients @ 40Mbps for Fair, Poor and Bad video segments. The playback buffer behavior also gets worst with increasing capacity. The number of video freezes goes from 10 (25@40) to 30 (50@80) to 112 (100@160).

The percentage of clients experiencing these video frame freezes jumps from 28% to 38% then all the way to 64% of all clients. By all indications, the unmanaged ABR system is becoming more unstable with growing capacity with given client ratio.

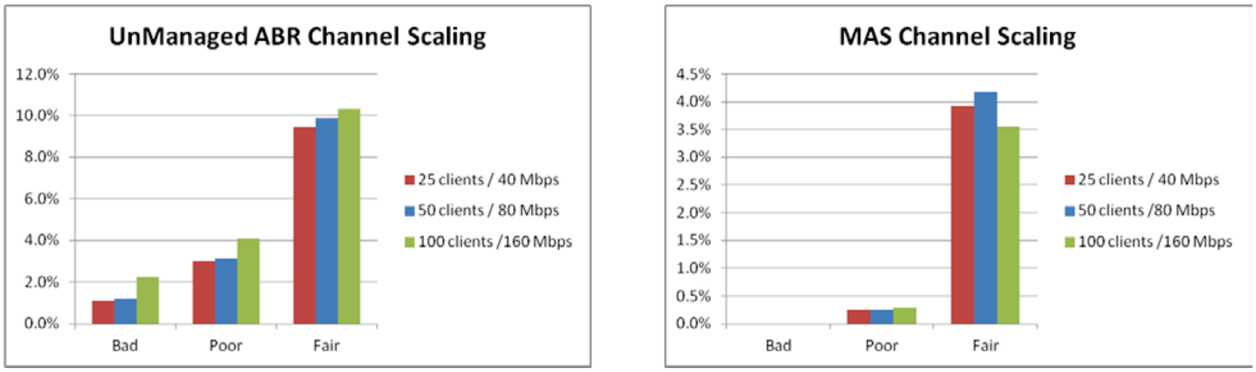

Figure 6 – Network Capacity Scaling: Unmanaged ABR, MAS

The MAS results on the right show a different tale. The system remains stable and varies very little as network capacity increases. There is a slight decrease in Fair video segments once the capacity is increased to 160Mbps. Otherwise, MAS system behavior is the same for each capacity.

An interesting note is that the mythical stat mux gain is not present as capacity increases from 40Mbps to 160Mbps. The rationale here is that MAS has already provided significant stat mux gains for the first 25 clients. It appears that the system reaches a point of diminishing gains beyond that point, but rather scales linearly.

The charts below show some additional comparisons between unmanaged ABR and a cloud based MAS system. The 40Mbps test shows 18 & 25 client results for unmanaged ABR and 25 & 30 clients for MAS. The 160Mbps test shows 72 & 100 client results for unmanaged ABR and 100 & 120 clients for MAS.

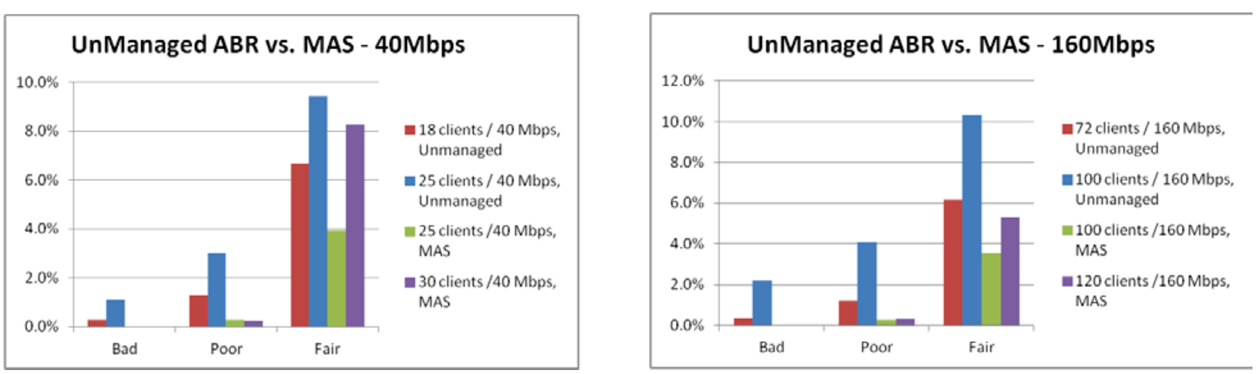

Figure 7 – Unmanaged ABR vs. MAS @ 40Mbps, 160Mbps

In both the 40Mbps and 160Mbps trials, MAS continues to have no video freezes, zero Bad segments and a very small 0.3% residual of Poor segments associated with the handheld device. The other interesting data point from above is that the MAS 120@160Mbps shows a noticeable decline in Fair segments from ~8% down to ~5%. This may indicate that there may be some advantages for MAS with a larger network capacity when the system is overloaded.

MAS – Behavior under Severe Client Overload

With that observation, the next MAS trial pushed up to 160 clients @ 160Mbps. With the system extensively overloaded, the number of Fair video segments shot up to 32% as everyone had to share the limited bandwidth. The number of Poor video segments had a slight nudge up to 0.8% and there were still zero Bad segments. The playback buffer behavior also remained quite stable with zero video freezes.

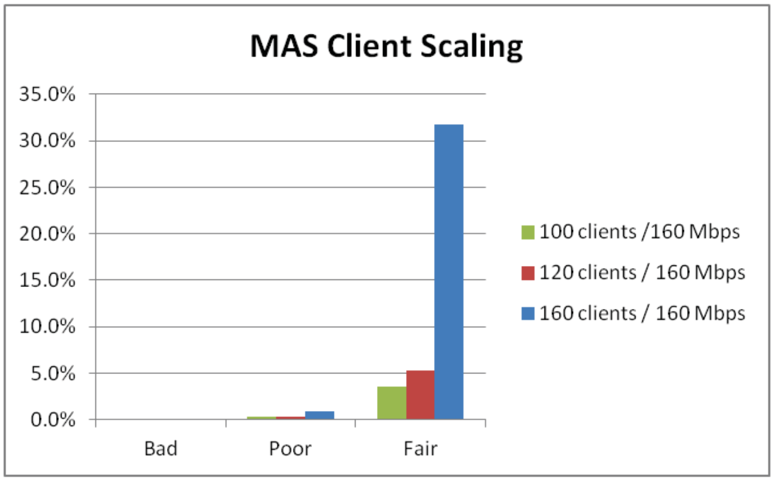

Figure 8 – MAS Client Scaling

This is a significant attribute about the MAS system. In traditional video systems, a complex admission control system is put into place in order to guarantee video quality to all clients. As the system becomes oversubscribed, then blocking starts to occur as clients are denied access to the system. With MAS, the operator can take a completely different tack. The above test show that even with nearly 2:1 oversubscription, the MAS system not only remains stable, but delivers a respectable video quality with two thirds of the video segments being delivered still in the Good or Excellent VQ range.

Managed ABR Conclusion

Operators have a number of challenges in offering a "managed" ABR video service to multi-screen devices such as tablets, smartphones, gaming devices and Smart TV's. The limited bandwidth makes it challenging to support the demand of a large number of concurrent users while maintaining good video quality for each user. In addition, existing unmanaged ABR client controlled distribution mechanism creates issues with QoE, network utilization, fairness and system stability. Previous SABR lab results presented in [ULM2013] confirmed all of these unmanaged ABR issues and highlights the advantages of using a cloud based managed ABR approach. SABR results also shows that DOCSIS QoS cannot fix the underlying issues, only make moderate improvements.

This paper continues the research into managed ABR video services and highlights the results from our Managed Adaptive Streaming (MAS) system. In particular, results show the impact as the number of clients and the network bandwidth scales to much larger numbers.

The initial baseline tests for unmanaged ABR for a moderately full system (i.e. 18 clients in 40Mbps) confirmed some of the SABR findings. At this level, a noticeable number of video segments were delivered with Poor or even Bad Video Quality (VQ). Overall, this QoE might be acceptable for Over The Top (OTT) content, but would be very marginal for an operator trying to offer a managed Video service. As the client load is increased by almost 40% to 25 clients, the unmanaged ABR system degrades significantly. There was a fourfold increase in Bad & Poor VQ segments and a total of 28% of the clients even experienced video freezes from playback buffer under-runs. From an operator's perspective, the QoE has now become totally unacceptable.

The MAS system was extremely well behaved at 40Mbps. Even with 25 clients, MAS did not experience any video freezes and caused no Bad or Poor VQ video segments. Even at 30 clients, the MAS system had no problems and was behaving much better than the unmanaged ABR with 18 clients. Finally in a severely overloaded MAS system with 40 clients, it still remained stable. A third of the video segments dropped from Good/Excellent down to a Fair VQ as clients had to share bandwidth, but Poor VQ segments was still less than 1% and there were no Bad VQ segments or video freezes. The MAS system at 40 clients is still better behaved than unmanaged ABR with 18 clients.

The next set of results take a look at how an ABR system scales with network capacity. Conventional wisdom suggests that things will improve with statistical multiplexing as one goes from one to two to four bonded DOCSIS channels. The percentage of clients experiencing these video frame freezes jumps from 28% to 38% then all the way to 64% of all clients and the proportionate number of Bad, Poor and Fair VQ video segments increases. By all indications, the unmanaged ABR system is becoming more unstable with growing capacity which is contrary to the conventional wisdom.

Meanwhile, the MAS system scales nicely with network capacity. In fact there was a very modest improvement with the increased network capacity, but nothing significant.

The MAS team concluded that most of the statistical multiplexing gain occurs with the first two dozen clients, and the system scales linearly beyond that.

The final set of results provided showed how MAS behaves under significant loads: 160 clients in a 160Mbps network. For these severely overloaded systems, there appears to be some additional gains from the larger network capacity. This is a significant attribute about the MAS system. In traditional video systems, a complex admission control system is put into place in order to guarantee video quality to all clients, which causes blocking when system is overloaded. Our results show that even with nearly 2:1 oversubscription, the MAS system not only remains stable, but delivers a respectable video quality with two thirds of the video segments being delivered in the Good or Excellent VQ range, less than 1% in the Poor VQ range and zero Bad VQ segments.

Our results show that a cloud based approach such as MAS enables the operator to regain control in providing a first rate video service with exceptional Quality of Experience while retaining the underlying benefits of Adaptive protocols. Our lab results show that all major shortcomings of unmanaged ABR can be addressed. The system becomes stable, network utilization improves and video quality fairness is provided across the entire client population while gracefully handling congestion and providing an improved user experience. These tests suggest that MAS delivered comparable video quality with more than twice the client load while providing fairness and stability. So, taking the car keys away from those unruly teenagers does appear to be the way for the operator to take control again.

References

[ULM2013] – "Smart ABR – The Future of Managed IP Video Services"; J. Ulm, A. Luthra, P. Moorthy, M. Schmidt, H. Xu, The Cable Show NCTA/SCTE Technical Sessions spring 2013

[CLOONAN] – "Competitive Analysis of Adaptive Video Streaming Implementations"; T. Cloonan Ph.D., Jim Allen Ph.D., The Cable Show NCTA/SCTE Technical Sessions spring 2011

[ULM2010] – "Adaptive Streaming – New Approaches for Cable IP Video Delivery"; J. Ulm, T. du Breuil, G. Hughes, S. McCarthy, The Cable Show NCTA/SCTE Technical Sessions spring 2010

[WHITE] – "Reclaiming Control of the Network from ABR Video Clients"; G. White, J. Ulm, The Cable Show NCTA/SCTE Technical Sessions 2012

Appendix – MAS Test Results

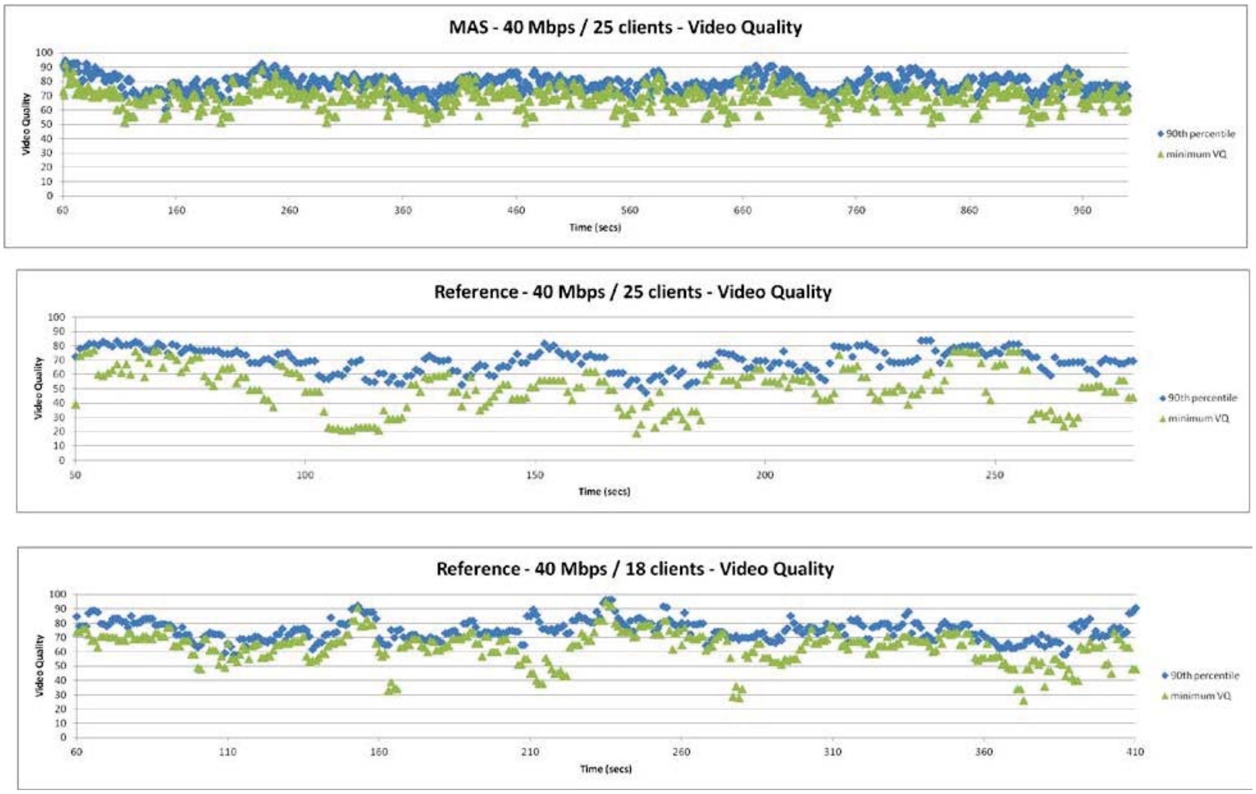

Figure A1 – VQ Time Line: 40Mbps Unmanaged & MAS

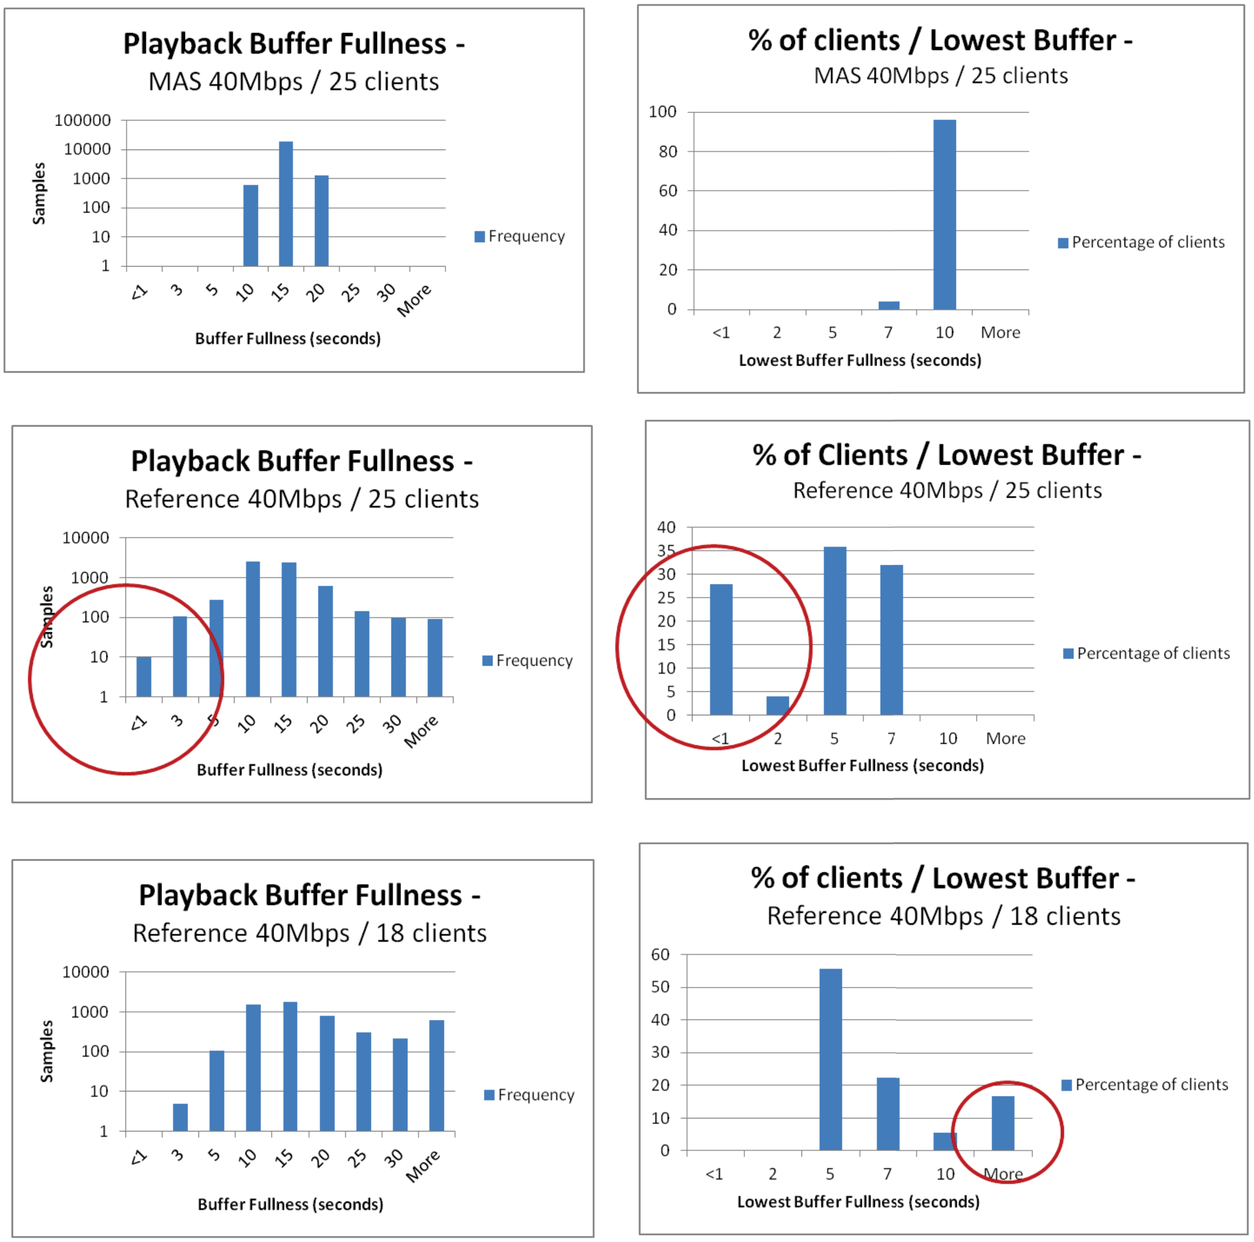

Figure A2 – Playback Buffer statistics: 40Mbps, Unmanaged & MAS

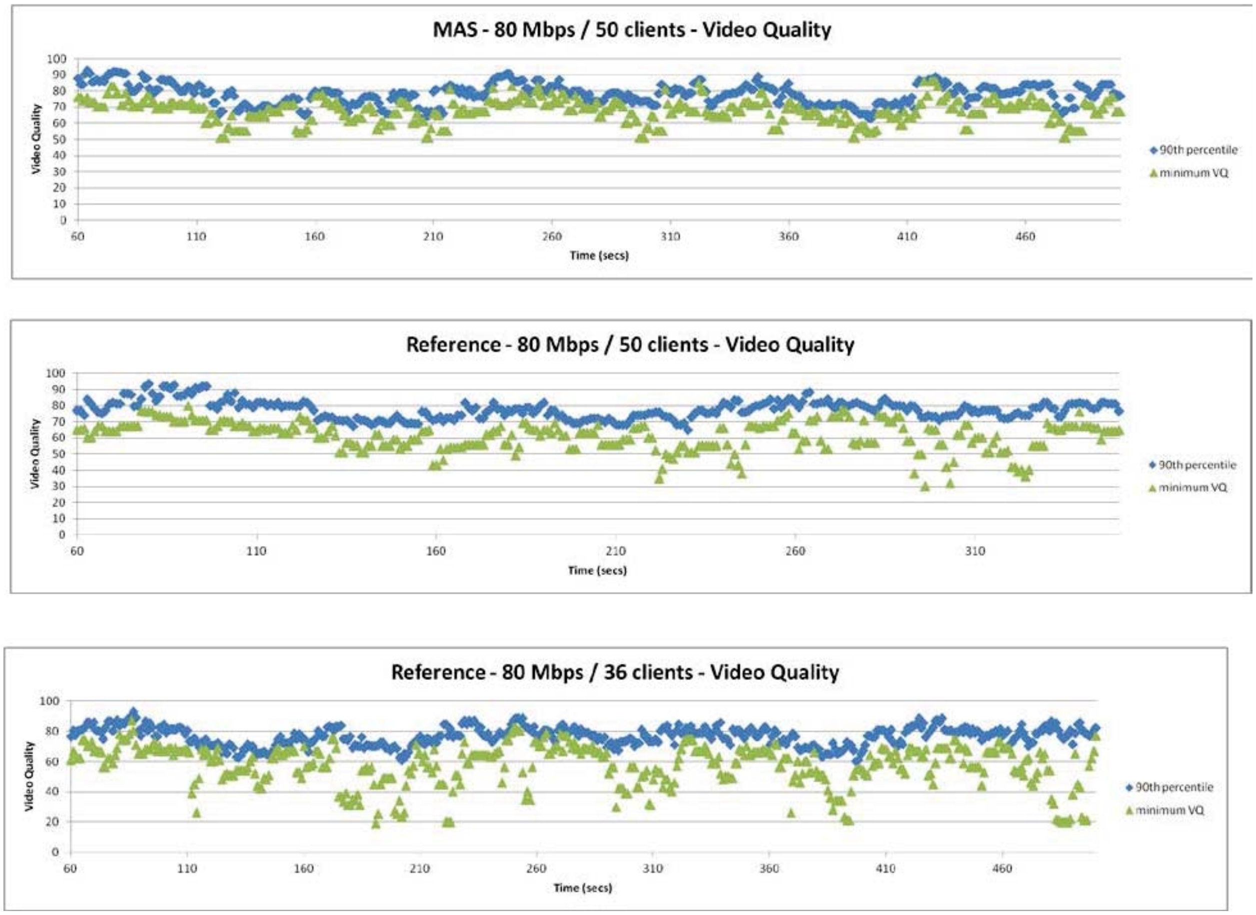

Figure A3 – VQ Time Line: 80Mbps Unmanaged & MAS

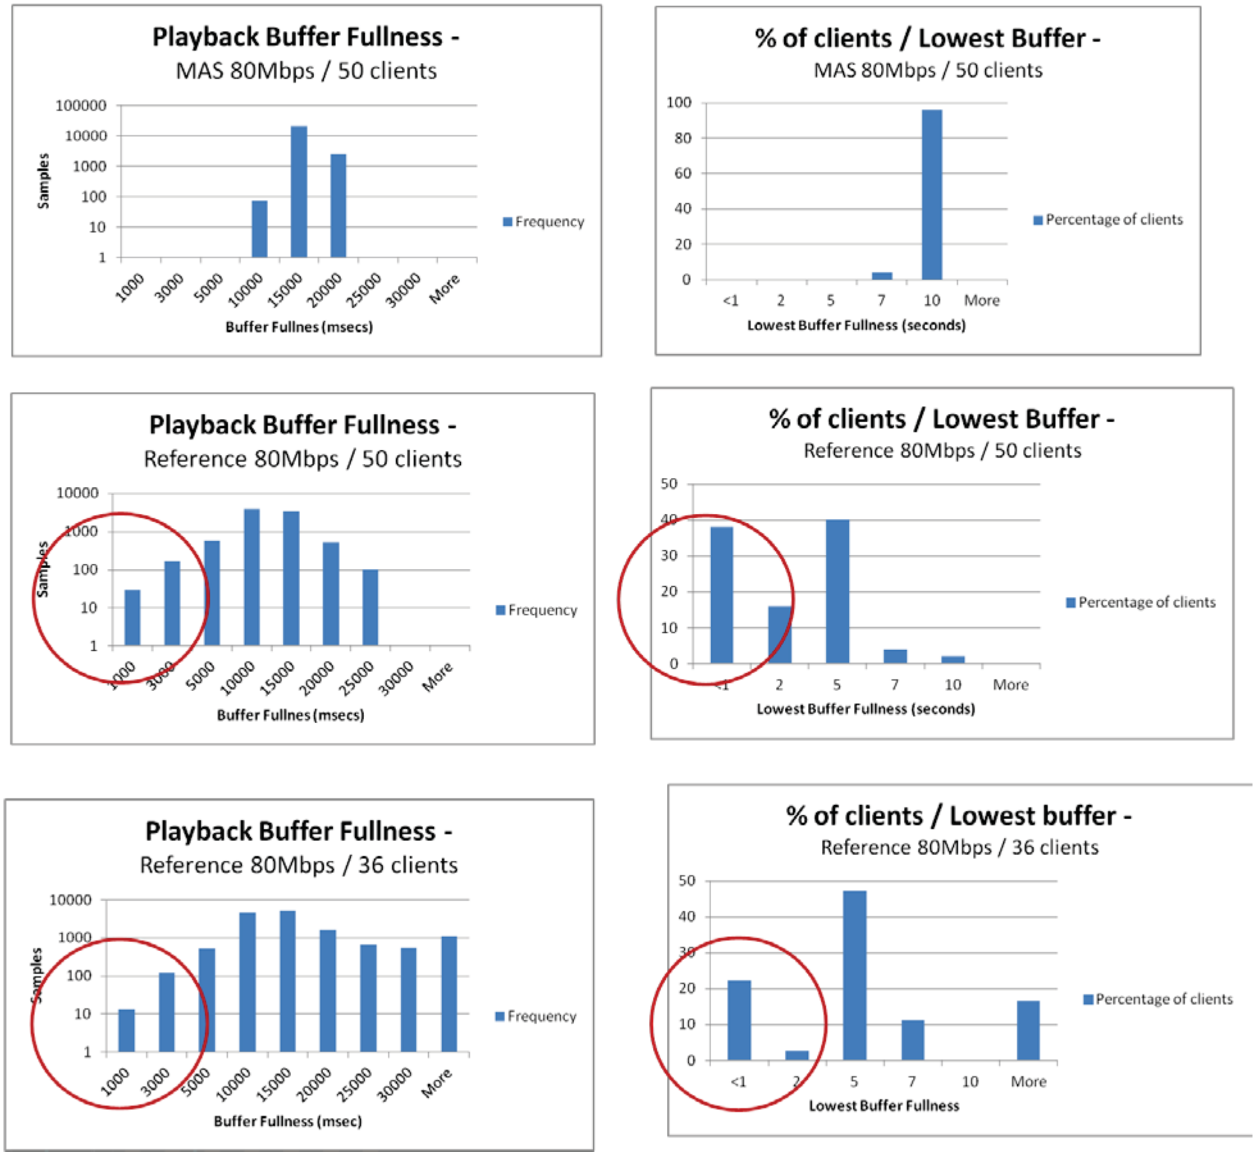

Figure A4 – Playback Buffer statistics: 80Mbps, Unmanaged & MAS

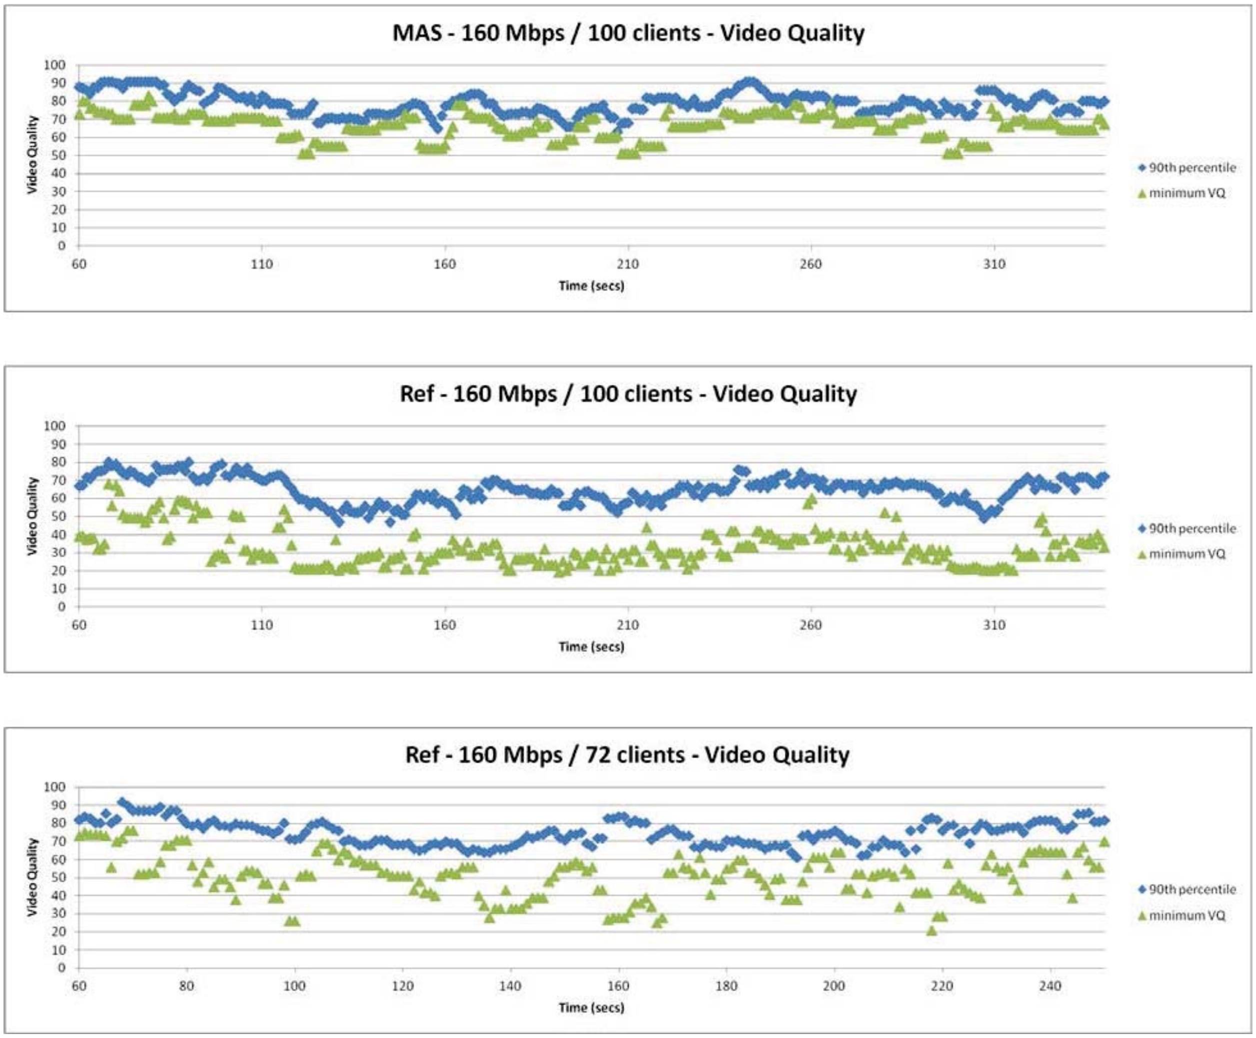

Figure A5 – VQ Time Line: 160Mbps Unmanaged & MAS

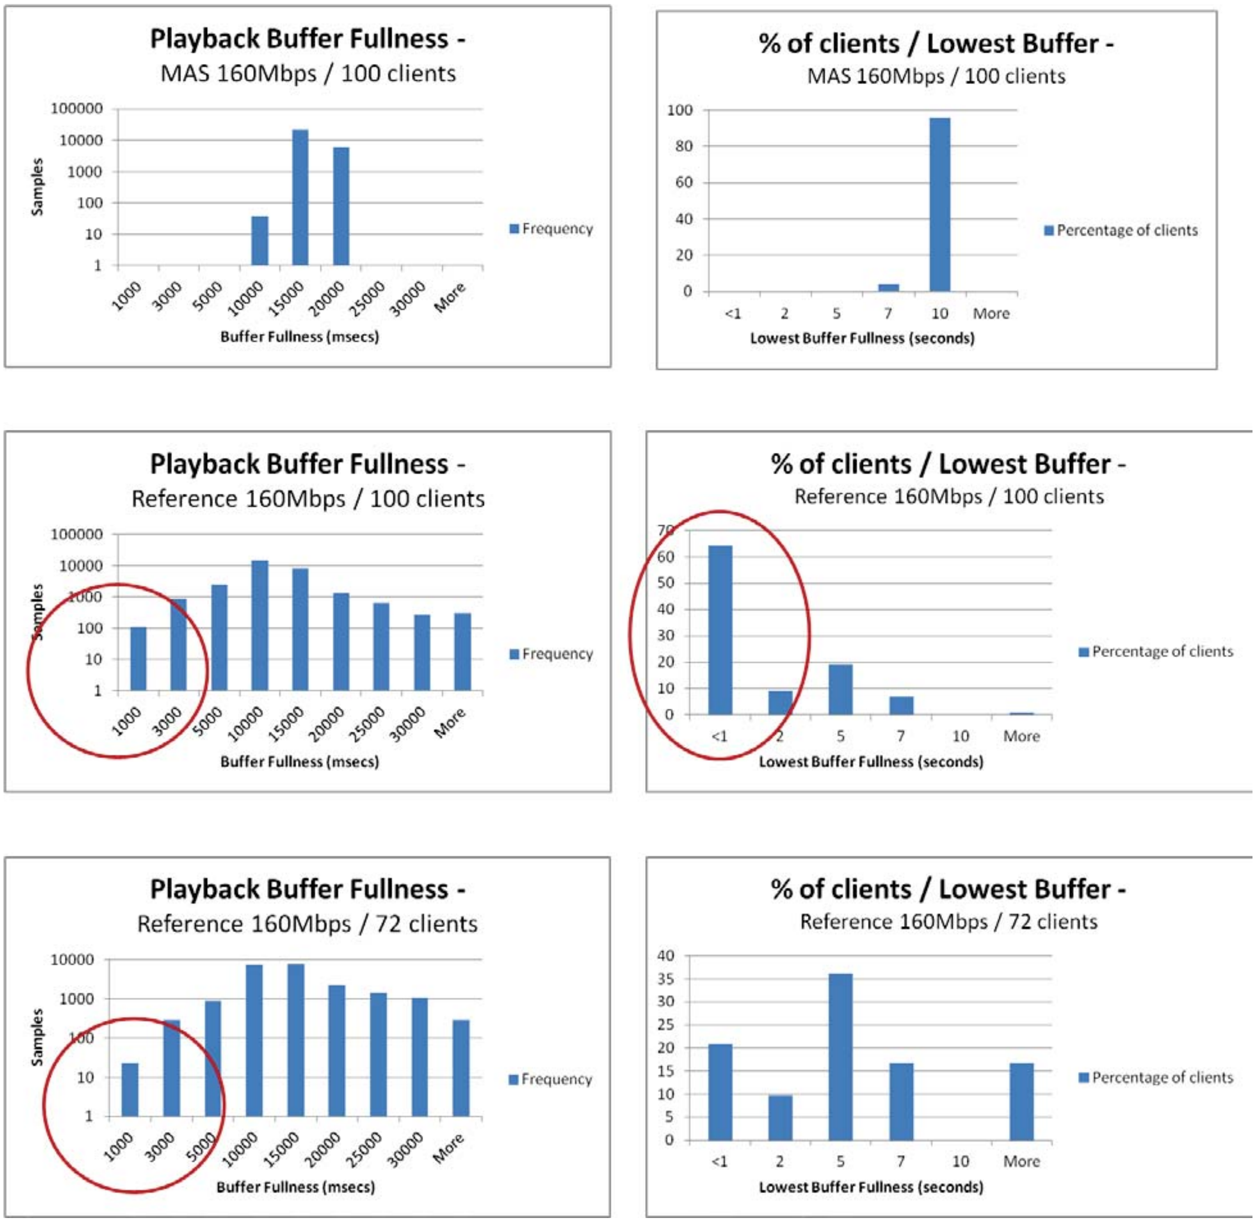

Figure A6 – Playback Buffer statistics: 160Mbps, Unmanaged & MAS

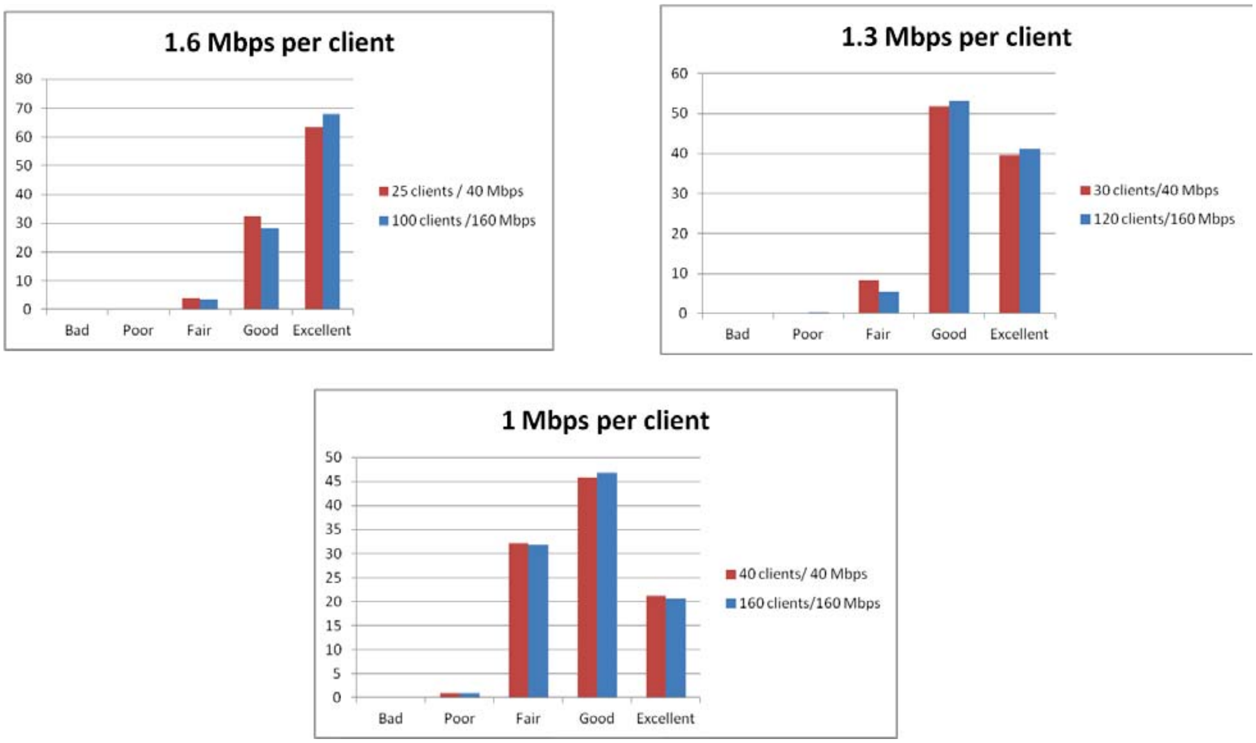

Figure A7 – MAS Capacity per Client