QOE ISSUES RELEVANT TO VIDEO STREAMING IN CABLE NETWORKS

Jeremy Bennington
Praveen Mohandas

Cheetah Technologies, Sunnyvale, CA

Abstract

This paper explores how monitoring video quality in a streaming service environment is different from a traditional video network. It discusses technical differences as well as how quality assurance data must be correlated and reported so that it becomes operationally relevant and actionable for an operator. The paper also proposes a set of best practices and architecture for operators to consider for monitoring video quality in a streaming service environment.

INTRODUCTION

Emerging Internet video streaming platforms like Hulu, Netflix, YouTube, Amazon and iTunes are challenging the traditional video service business offered by Cable MSO's. In addition to providing on-demand service, they provide access to TV shows and events aired previously – providing a cloud-based video functionality across multiple mobile platforms, an experience appealing to many users. User generated content aggregators like Justin-TV, YouTube, Ustream, LiveStream and several other providers are also trying to get into the premium content space. Netflix currently has about 20 million customers; in few years they expect this number to rise to 60 million. MSOs, like Comcast through its Xfinity TV application, TimeWarner Cable and Cablevision with their iPad apps, are rapidly innovating their service offerings as well to provide a multiple screen service to subscribers.

While subscribers expect good video quality for multi-screen streaming services just like they do for traditional video services, the radically different architecture, distribution model and consumption behaviors of these new services calls for a new approach to qualifying and ensuring a high Quality of Experience (QoE). New ways of measuring, correlating and aggregating QoE data are required to take into account, at the most basic level, the disappearance of the household as the ultimate frontier. With streaming services, users are not limited to a single device, to a single operator network, or even to a device hard-wired to the network.

Monitoring video QoE of streaming services involves different codecs (MPEG4), wrappers (Adobe Flash, HTTP5, MS Silverlight) and distribution protocols. Note, with recent announcements, the adoption of HTTP Live Streaming (HLS / HTTP5) may become the dominant method. All these techniques pose a significant challenge due to segmentation of the video into small chunks with multiple bit rates for transmission. Segmentation, however, only represents a portion of the issue at hand.

The biggest challenge lies in the representation of the data and in making it operationally relevant to the operator in a rapidly changing business and technical climate. Some of the questions operators are trying to address include: how was the user's experience across multiple assets since the idea of a channel does not directly apply? How was the experience across multiple screens? How was the experience of the subscribers within a single household (ie. revenue generating unit)? How was the experience of a specific asset across a geographic region or across a type of device? How does the streaming experience of a broadcast correlate with the traditional delivery of a given channel?

STREAMING VIDEO ARCHITECTURE

There are several architectures for streaming of video. In general they all start with either a video broadcast stream from a programmer that may also be broadcast over a traditional distribution like CATV, or the streaming video starts with a fixed file asset like VOD or user generated content. This part of the architecture can be named "Content Ingest" (see figure 1).

The next step in the streaming video architecture is the "Content Encoding" or transcoding of the video into the codec and bit rate desired by the operator. Typically the encoding is MPEG4 and includes multiple bit rates for each program so a client can adapt to the network bandwidth available. There are two to four bitrates of each program in typical deployments. The next step in the "Content Encoding" is the segmentation of the video into small "chunks" of video typically about 2 seconds long. These segments are progressively requested and downloaded by the client application based on the m3u8 "playlist" file.

After the encoding and segmentation of the video the programs arrive at the "Origin Servers" where the information about the programs is made accessible by the clients. The client negotiate and communicate with the Origin Servers to gain access to the program and the different bit rates as bandwidth constraints are determined by the HTTP/TCP protocols and client interaction. From the Origin Servers the video is transmitted over a managed or unmanaged network. It may be a Content Delivery Network (CDN) specifically or purely third

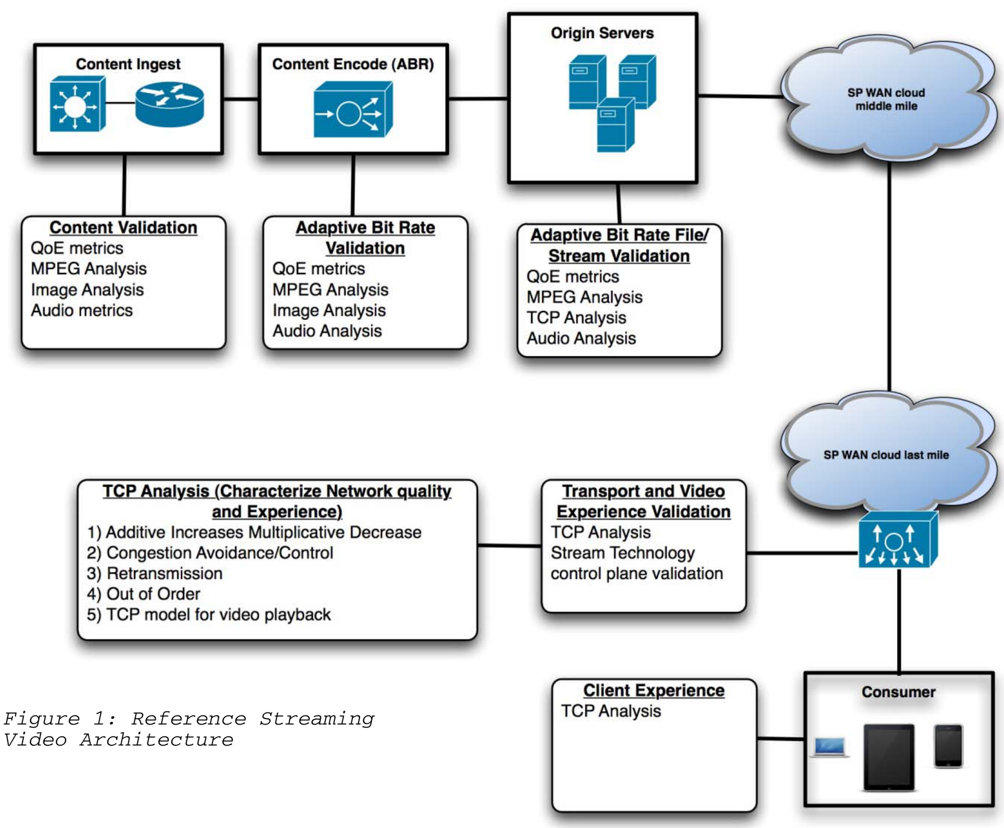

Figure 1: Reference Streaming Video Architecture

party unmanaged networks or any mixture. Once the video has traversed the network it is consumed by a client. The client may reside on a TV, PC, tablet or mobile device.

COMMON QOE METRICS FOR STREAMING & CATV

In both streaming video and traditional video in a cable network the content acquired from a programmer or user must be validated for quality before encoding. This "Content Validation" provides quality assurance and a baseline for the operator to compare QoE at all the subsequent steps of video processing and distribution. Content Validation at the core is not highly dependent on the MPEG and transport information and instead measures the decompressed video quality as if a subscriber were actually viewing the video.

Using human vision system models we can computationally model the subjective score of video quality, but instead of repeating subjective scores using personnel watching mosaics of video monitors one can automate the process. Thus making the video fidelity (image analysis) scores objective, repeatable, consistent and scalable to ensure that an operator always knows what the video quality was to begin with. This avoids finger pointing between programmer and operator, as well as solving answering the age old question of junk in / junk out or good in / junk out?

When video fidelity measurements are tracked over time operators and programmers can ascertain the overall quality ranking of programs, monitor changes in quality and short term events where the quality diverges. If this is then correlated with information about the MPEG coding and transport the programmer and operator can ensure the Content is validated before an operator processes the video for distribution.

Once the content has been validated, quality trends identified and quality baselines established the encoding/transcoding of the video must be monitored. Monitoring of the encoding process for streaming video and traditional video is very similar. The output of the encoder must measured for MPEG-TS metrics like PCR jitter and other ETSI TR101290 type metrics. The MPEG VCL (video coding layer) should be analyzed to determine the GOP structure, slice types, quantizer values, bit rates and other attributes of how the video was encoded. Last, but most important, the video content must be inspected for image artifacts commonly introduced by encoders like blockiness, frozen video, blackout and jerkiness.

All of the QoE measurements from the Content Encode can be correlated to provide quality assurance of the encoding/transcoding process across the service offering. This information should also be compared and correlated to the content ingest measurements to provide the operator with a measurement of the degradation of each program introduced by their processing of the video.

At this point in a traditional cable network the video would be multiplexed with other programs and sent over the network for transmission (with possibly ad-insertion). In a streaming network the video is not multiplexed. Instead it is segmented into two second chucks (typically) and transmitted to clients via HTTP/TCP by the origin server. At the origin server an operator should emulate a client to monitor the various bit rates of each service and the availability.

Another vast difference between streaming and traditional video distribution in cable networks is the end device and sometimes the presentation of the content. In traditional environments a tightly controlled STB decodes the video (at a fixed bit rate) and provides one of three or four fixed resolutions to the television. Even though there are many television suppliers it is more constrained than the display types and capabilities of mobile devices! The presentation of the video is also tightly constrained by a channel lineup and guide typically provided by the cable operator unless a CableCard device is being used. In the streaming environment there are multiple bit rates, a huge type of display types, presentation, etc.

UNIQUE METRICS FOR STREAMING

Unlike traditional video distribution, in the streaming architecture, video will be encoded into multiple bit rates. This is done to allow a client to adapt to bandwidth availability thus providing high quality video when bit rates are high and lower quality video when bit rates are low, but reducing the number of stalls which are more annoying to a user than image quality (in most cases). With multiple bit rates of video being produced by the encoder they each must be correlated with the content ingest and between each other to ensure consistency. This can also be used to ascertain the end-user quality experience when they look at a program and shift between different bit rates of the same program.

Unlike the traditional distribution and reception of video by the user in a streaming environment the quality may be changing on purpose! Traditionally a user tunes to a HD or SD channel and receives a fixed resolution and bit rate service. An operator monitors to ensure that the service meets that expectation to reduce support costs and maintain high customer satisfaction. In a streaming environment a user will have a variable quality experience. The causes of the variability may be within the control of the subscriber, operator or a third party. Regardless of the cause an operator should monitor several attributes to ascertain the subscriber QoE. Primarily they should look at the bit rates that are streamed for each asset to each subscriber. This correlated with the number of stalls provided a very good indication of the subscriber QoE. It is also very insightful to correlate this with the percentage of the asset viewed for a VOD asset or the length of view versus average length of view for a broadcast asset.

Unlike in a traditional distribution in a streaming architecture the transmission of video is over TCP/HTTP. This is a very flexible method, but it requires constant communication between the client and the server to provide the service. This increases the upstream bandwidth requirements of the service. If a cable operator chooses to provide a streaming video service they must consider not only the large downstream bandwidth requirements for their DOCSIS network, but also the upstream requirements. And, would it even make sense to build the upstream transactions into the QoS mechanisms of DOCSIS? Similarly a cable operator could provide a different SLA/QoS for streaming service, assuming the business case and law permits.

In m3u8 files there are bandwidth descriptors, however, monitoring the actually video bandwidth upstream and downstream provides a more accurate view of the network requirements. Moreover, aggregating the bandwidth usage based on bit rate types, program/asset and locations provides a better understanding of what programs are generating the most bandwidth demand, what locations are consuming the most bandwidth per stream, etc. Another important correlation is the bandwidth usage, stalls/buffering events, QoE and the duration/percentage of a program watched. From this an operator can gain insight into the affects of QoE on user behavior. For instance if a channel is only viewed for short periods or a VOD assent is not played completely is it because of the QoE and number of stalls, or was it ok and that is the normal behaviors for he particular piece of content?

BUSINESS MODEL IMPACTS ON QOE

The choice of business model and the choice of broadcast, VOD short-form and VOD long-form have a large impact on the QoE for streaming. For instance the user expectation of QoE for a subscription service is higher than a free service, but lower than that of a PPV service. Similarly the QoE expectation of long form video is higher than short form even though it may be more difficult to meet the short form QoE requirement (a conundrum).

The choice of a business model that adds streaming to an existing subscription versus a unique service also impacts the user expectations of QoE. In some cases linking the traditional subscription to a streaming one may create an expectation of the same QoE, which may be difficult to meet.

The choice of business model whereby the network/CDN is wholly owned and QoS can be guaranteed and QoE measured at all points is vastly different than an environment where the Origin Servers are the last asset that is owned. If the network is owned and tightly constrained by the operator they can monitor the origin server, client and also at the edge of the CDN where it interfaces with the broadband network. In this model it is easier to determine the quality across the service and the cause of issues when they arise. If the Origin Servers are the last owned part of the network then there is a greater need for monitoring within the client player/device to ascertain what the end user QoE is.

CONCLUSIONS

Video streaming in cable networks is a valuable service to expand a cable operator's value to subscribers by reaching beyond the TV to all media capable devices. Moreover it offers a different value and level of interaction with subscribers not available in a traditional TV model. When deploying a video streaming service QoE must be monitored to ensure that subscribers are receiving the QoE expected by the operator. Moreover, the operator needs to monitor QoE and be able to identify the root cause of any issues and determine if they are issues that are internal or external. Some of the QoE monitoring methods are the same between streaming and traditional video distribution, while there are several methods that are unique to streaming.

As operators launch a video streaming service the ability to correlate quality to viewing behaviors and across new and old distribution is vital to not only providing a good service and reducing cost, but also in finding new revenue opportunities. For instance the cable operator will now be able to differentiate different users within a "subscriber" address. In the traditional video environment a subscriber is a household with TVs. In the new paradigm of providing both a traditional TV service and streaming there is the possibility that each person in the household will consume their own content together or independently. Targeting ads to each person in the household becomes possible. Enabling interaction between the streaming services and the TV becomes possible as well. All of these possibilities hinge around ensuring a specific QoE for each service and correlating the data to reduce costs and ensure that behaviors are truly related to the content and not a poor QoE.

BIBLIOGRAPHY

[1] S. Winkler: *Digital Video Quality*. John Wiley & Sons, 2005.

[2] J. Deutscher: *Smooth Streaming Primer*. http://learn.iis.net/page.aspx/941/smooth-streaming-primer/