Zixi Technical Overview

Streaming Video over the Internet and Zixi

White Paper

Date: May 2015

Preface

The Video delivery landscape is evolving rapidly and is radically more diverse than just a few years ago. In the past, video was delivered to TV screens at home via managed networks like Hybrid Fiber Coax (HFC) or via Direct to Home (DTH) Satellite operators who took great strides to ensure good video quality. Today, video is delivered to many more devices and screens using unmanaged networks (AKA the Internet).

Viewers demand a TV experience that is delivered to mobile devices to be viewed from almost anywhere. The impact of poor video quality is significant in terms of how fast a consumer will give up on trying to view buffered, jittery, or poor video quality, especially for a live event where they as a consumer are emotionally invested. This has created a significant operational challenge for broadcasters, digital media companies, and service providers alike as the Internet does not offer the same quality of service consumers are used to seeing on TV and the instant available bandwidth and latency is uncertain or fluctuating. The impact of poor video quality, video dropouts, and jitter which translates to lost revenue in the form of digital ads. To overcome this challenge, Adaptive-Bit-Rate (ABR) streaming technology was developed and deployed. This white paper briefly reviews some of the popular ABR implementation protocols and outlines the advantages of Zixi protocol enabling the delivery of high quality video over the internet while maximizing link utilization with error-free reception, low fixed latency regardless of distance, jitter, or packet loss to any device, anywhere.

What is Adaptive Bit Rate Streaming?

Adaptive bitrate streaming is a technique used in streaming multimedia over computer server-based networks. While in the past most video streaming technologies utilized streaming protocols such as RTP with RTSP, today's adaptive streaming technologies are almost exclusively based on HTTP and designed to work efficiently over large distributed HTTP networks such as the Internet. Prior to ABR streaming, web-based or mobile video delivery was typically done by encoding a stream at a fixed bit rate and frame size. Viewers could buffer some of the video, and then simultaneously download and play it back. This delivery model was similar to cable transmission, where a single bit rate is transmitted over a reliable medium.

The three main factors that affect streaming video performance over the IP network are packet loss, latency and jitter. Data is sent via "packets" from one computer to another. Packet errors occur when small packages of data get lost or delayed traveling across a network. When the data packets get dropped or delayed, performance issues occur that make applications such as video streaming difficult. Jitter is the variability of latency between received packets.

Network latency is the term used to indicate any kind of delay that happens in data communication over a network. Network connections in which small delays occur are called low-latency networks whereas network connections which suffer from long delays are called high-latency networks. High latency creates bottlenecks in any network communication. It prevents the data from taking full advantage of the network pipe and effectively decreases the communication bandwidth. The impact of latency on network bandwidth can be temporary or persistent based on the source of the delays.

Unlike managed networks, Internet bandwidth can vary widely causing viewers with low bandwidth availability to suffer from excessive buffering. To compensate, video is encoded at lower bit rates, lowering quality which negatively affects quality of experience (QoE). Even then, fluctuations in bandwidth cause buffering delays.

To solve this problem, streaming content is encoded into multiple layers, each at a different bit rate, frame size and/or frame rate. These layers are combined into a single package that represents the original content. ABR players switch between layers depending upon the device and available bandwidth, to ensure consistent high-quality playback.

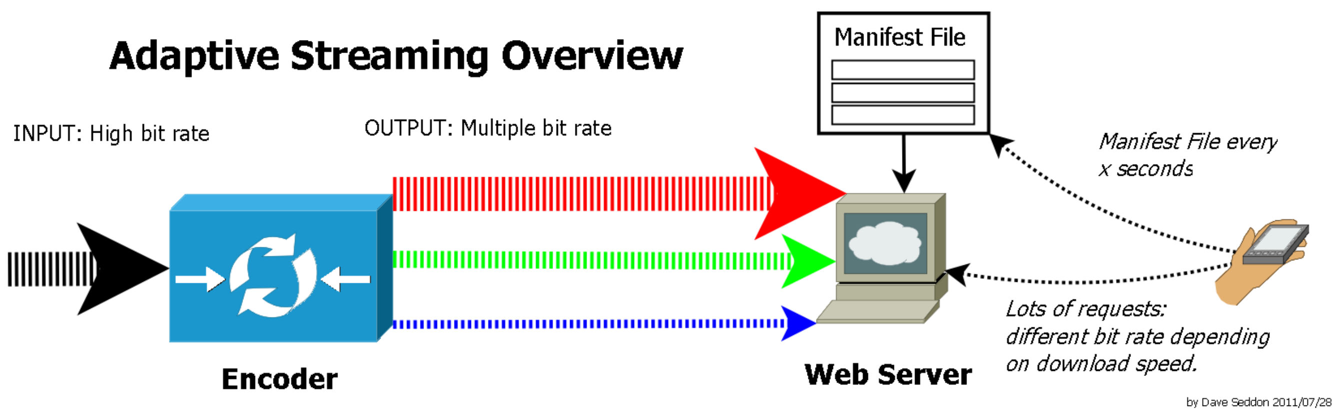

Figure 1: Adaptive Streaming Overview (source Wikipedia)

How Does Adaptive Bit Rate Streaming Work?

ABR streaming works by detecting a client's bandwidth (and CPU capacity) and adjusting the quality of a video stream in real time. The player in the client device switches between the different bit rates depending on available resources. More specifically, adaptive bitrate streaming is a method of video streaming over HTTP where the source content is encoded at multiple bit rates, then each of the different bit rate streams are segmented into small multi-second parts. The streaming client is made aware of the available streams at differing bit rates, and segments of the streams by a manifest file. When starting, the client requests the segments from the lowest bit rate stream. If the client finds the download speed is greater than the bit rate of the segment downloaded, then it will request the next higher bit rate segments. Later, if the client finds the download speed for a segment is lower than the bit rate for the segment, and therefore the network throughput has deteriorated, then it will request a lower bit rate segment. The segment size can vary depending on the particular implementation, but they are typically between two (2) and ten (10) seconds. All current implementations of ABR require significant buffering (latency) typically ranging from 10 to 30 seconds and are most suitable for file based (i.e. not live) video streaming (VOD).

Streaming Protocols and Popular Implementations of ABR

Adobe Real Time Messaging Protocol (AKA RTMP)

RTMP is a TCP-based protocol which maintains persistent connections and allows low-latency communication. To deliver streams smoothly and transmit as much information as possible, it splits streams into fragments and their size is negotiated dynamically between the client and server: The default fragment sizes are 64-bytes for audio data and 128 bytes for video data. Fragments from different streams may then be interleaved and multiplexed over a single connection. With longer data chunks the protocol carries only a one-byte header per fragment, thereby incurring very little overhead. However, in practice individual fragments are not typically interleaved. Instead, the interleaving and multiplexing is done at the packet level, with RTMP packets across several different active channels being interleaved in such a way as to ensure that each channel meets its bandwidth, latency, and other quality-of-service requirements. Packets interleaved in this fashion are treated as indivisible, and are not interleaved on the fragment level.

MPEG-DASH

MPEG-DASH, is an adaptive bitrate streaming technique that enables high quality streaming of media content over the Internet delivered from conventional HTTP web servers. Similar to Apple's HTTP Live Streaming (HLS) solution, MPEG-DASH works by breaking the content into a sequence of small HTTP-based file segments, each segment containing a short interval of playback time of a content

MPEG-DASH is the only adaptive bit-rate HTTP-based streaming solution that is an international standard. Work on DASH started in 2010; it became a Draft International Standard in January 2011 and an International Standard in November 2011.

HLS

HTTP Live Streaming (HLS) is an HTTP-based media streaming communications protocol implemented by Apple Inc. as part of QuickTime X and iOS. HLS supports both live and video on demand content. It works by breaking down streams or video assets into several small MPEG2-TS files (video chunks) of varying bit rates and set duration using a stream or file segmenter. One such segmenter implementation is provided by Apple. The segmenter is also responsible for producing a set of index files in the M3U8 format which acts as a playlist file for the video chunks. Each playlist pertains to a given bitrate level, and contains the relative or absolute URLs to the chunks with the relevant bitrate. The client is then responsible for requesting the appropriate playlist depending on the available bandwidth.

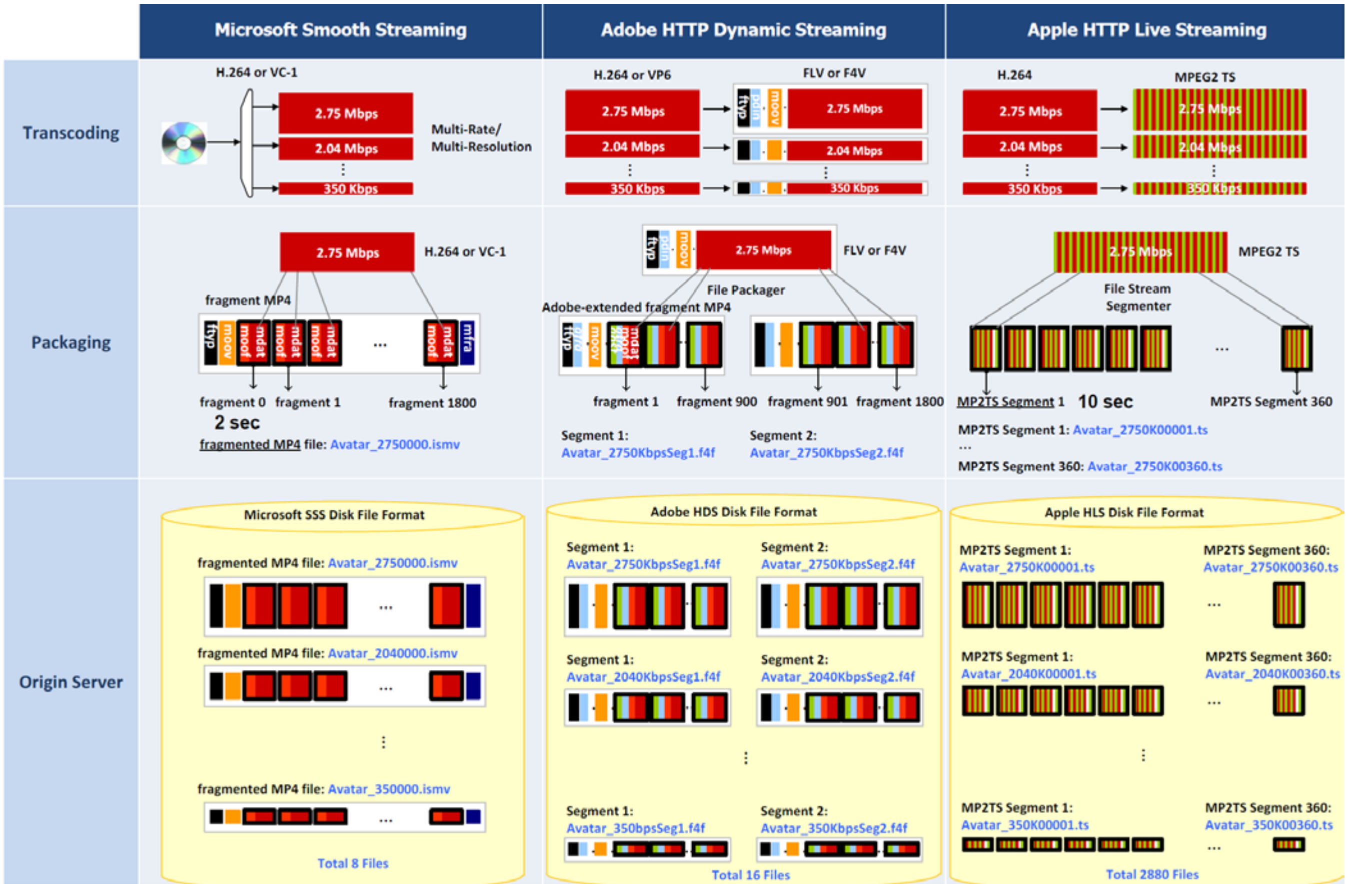

Figure 2: HTTP Adaptive Streaming comparison

RTSP

The Real Time Streaming Protocol (RTSP) is a network control protocol designed for use in entertainment and communications systems to control streaming media servers. The protocol is used for establishing and controlling media sessions between end points. Clients of media servers issue VCR-style commands, such as play and pause, to facilitate real-time control of playback of media files from the server. The transmission of streaming data itself is not a task of the RTSP protocol. Most RTSP servers use the Real-time Transport Protocol (RTP) in conjunction with Real-time Control Protocol (RTCP) for media stream delivery; however some vendors implement proprietary transport protocols.

Figure 3: Performance comparison of 3 different streaming protocols (RTMP, HLS and RTSP)

Zixi Streaming Protocol

Zixi Streaming Protocol delivers outstanding performance (at low and fixed latency), superior reliability (no packet loss) and broadcast video quality at the physical end-to-end bandwidth available (SD, HD, and UHD), at the same time with no tradeoffs to delay, resolution and stutter. In spite of the varying network conditions of the public Internet, where the amount of network errors, packet loss, jitter and out-of-order packets fluctuate "every second" - Zixi's proven and patented technology dynamically resolves these issues in sub- millisecond response times to deliver pristine, uninterrupted low latency video.

With minimum overhead to physical bandwidth, this dynamic mechanism provides low end-to-end latency, removes jitter, recovers and re-orders packets, smooth video delivery and regenerates video in its original form all in real-time. The Zixi streaming protocol is the best method to maximize link utilization with error-free reception, low fixed latency regardless of distance, jitter, or packet loss:

How does Zixi do that?

The Zixi streaming protocol combines multiple techniques to provide:

- Hybrid intelligent error correction mechanism:
    - Content-aware Forward Error Correction (CA-FEC)
    - Content-aware Automatic Repeat Request (ARQ) packet loss recovery
- Dynamic assessment of link quality
- Dynamic de-jitter buffer and null bit sequence compression option
- Dynamic adaptive bit rate control with encoder backpressure feedback based on available bandwidth
- Network bonding
- Seamless Adaptive bitrate over UDP, unicast or multicast

The above capabilities enable high quality, reliable video streaming over the internet with:

- Low latency with a settable latency control
- Efficient utilization of available link bandwidth
- Highest possible video quality
- Dynamic adjustment to network changes
- Distance agnostic
- Multicasting to reduce network bottlenecks

Zixi's transport stream starts at content acquisition where the Zixi Feeder encapsulates a video stream and initiates delivery over IP either directly to a Zixi Receiver (point-to-point) over any distance or to a Zixi Broadcaster installed in the cloud or an on-premise server for multi-point. Zixi is able to maximize physical bandwidth as the protocol requires very low overhead (2-3%). This allows for high bandwidth utilization when combined with Zixi error correction while delivering a high quality stream in even severely degraded network conditions. Zixi also has the ability to support network bonding - to aggregate bandwidth across multiple network connections - for even higher bit rate delivery over IP. In addition, Zixi is able to dynamically throttle the source encoder bitrate for even greater flexibility and resiliency.

The Broadcaster, which is Zixi's management hub, enables the implementation of complex, multipoint workflows simply and reliably. The Zixi Broadcaster supports clustering for scalability and offers built-in redundancy for protected failover. Zixi Broadcaster supports recording (in the cloud or to local storage on a deployed Zixi appliance), time-delayed playback, repackaging and transcoding. Zixi enables delivery to Zixi-enabled devices (best case scenario for high quality delivery) and non-Zixi enabled devices in popular protocols (RTMP, HTTPs, HLS, etc.). Zixi can transcode and deliver to any digital platform in a variety of different bit rates and protocols simultaneously.

Zixi Compared With Popular Streaming Protocols

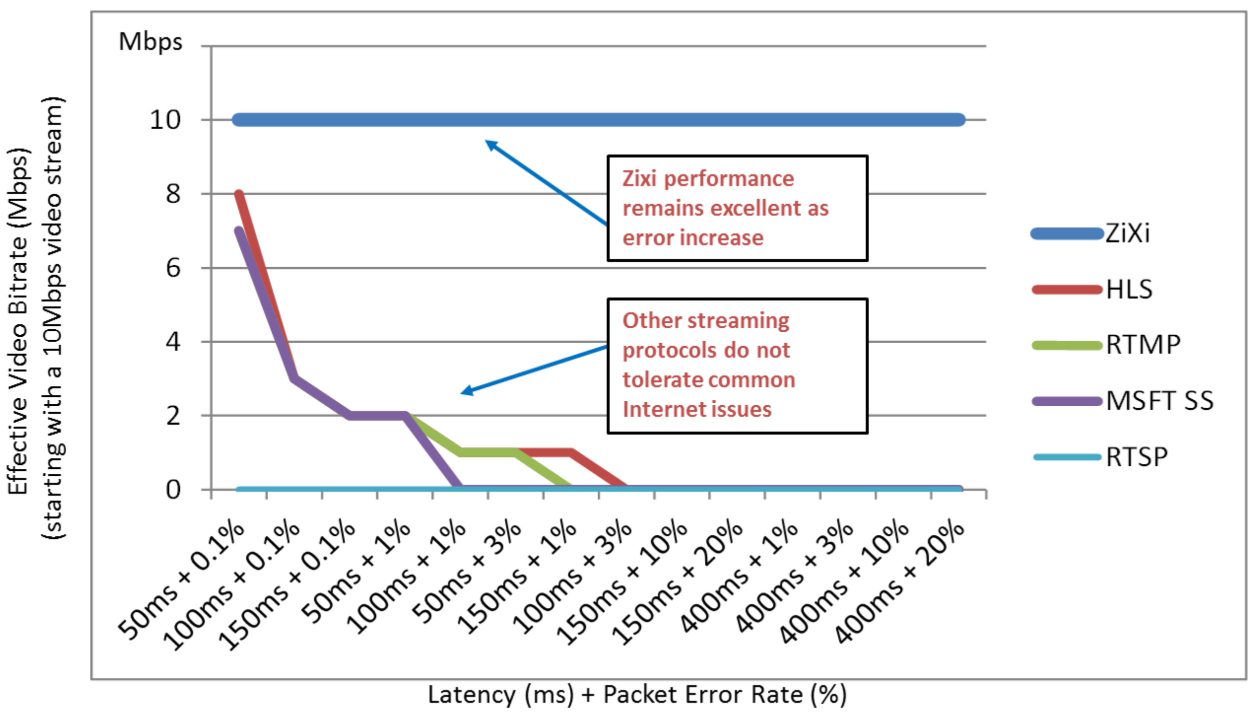

Figure 4: Zixi technology performance compared with many of the popular streaming protocols with increasing latency and packet loss

Key Advantages of Zixi Protocol:

- An end-to-end network and content-aware software-based protocol
- Able to stream at predictable low and fixed latency over standard internet connections
- Distance-agnostic and low overhead
- Maximizes the available bandwidth to deliver high quality video (SD, HD, UHD) to any device anywhere
- Flexible implementation - encapsulates the encoded stream and reliably and securely delivers that stream either to a Zixi Receiver (Point to Point or Multi-point to Multi-point) or sends that stream to the Zixi Broadcaster (on premise or in the cloud (AWS)), where the file can be transcoded to a variety of common protocols and bitrates
- Agnostic to the encoded stream type – works with any video format
- Platform neutral - Zixi SW can be embedded in any device (encoders/decoders (HW or SW), camera, mobile devices (phone/tablet), set top box, game console
- Wide acceptance - Zixi has developed a growing ecosystem of partners who have embedded Zixi in to their devices making adding Zixi to an existing live event workflow seamless and easy

Summary

Video Quality of Experience (VQoE) impacts viewer engagement, especially for a live event where viewers are emotionally invested (live sports or concert event), which has a direct impact on ad revenue for content rights holders, broadcasters, digital media companies (new media), service providers and more. Delivering live content from anywhere to any device at any time where consumers expect a TV like experience on their mobile device is the standard that companies now need to deliver. When it comes to delivering live video over a standard internet connection, that challenge is magnified even further creating an operational challenge for media companies looking to meet this demand for consumers for content delivered anywhere at any time, while driving new ad revenue across multiple platforms. Delivering a TV-like experience to a mobile device for live and VOD content is critical to driving ad revenue as a quality experience means viewers will watch the entire event including the commercials.

Companies employ different tools that leverage adaptive bit rates techniques to deliver live and VOD content over standard internet connections in order to deal with the challenges outlined in this white paper for video delivery over IP. The problem, even with these various protocols, is that the experience can still become compromised as network interference drives the use of the lowest available bit rate adversely affecting the QoS. As network availability becomes scarce, the adaptive bit rate protocols deliver a lower quality to make up for loss of bandwidth. This can translate not just to lower quality but to loss of signal altogether which for a live event is unacceptable for consumers.

Zixi solves this problem by encapsulating the stream in the Zixi transport stream protocol and, using a low overhead, maximizes the available bandwidth from the source to deliver a higher quality experience. The Zixi protocol then leverages a host of advanced error correction techniques to manage jitter and packet loss, automatically defining which technique to use depending on what is happening on the network at a given time. It is this network awareness that makes Zixi's end-to-end platform unique, able to deliver video (SD, HD, and UHD), from any device anywhere over any distance to any device anywhere error-free at the highest bitrate and lowest latency. The result is a high quality user experience delivered to any device that positively impacts the ability for broadcasters, digital media companies, and service providers' ability to drive ad revenue.

Zixi does all this leveraging the availability of the Internet while driving down cost for delivery vs. traditional video delivery methods like fiber, satellite, and MPLS. Zixi is not only more cost-effective but easier to implement saving further on operational costs as it requires fewer people to run a live event. Zixi gathers network data and presents that data inside of the application showing how many packets are lost and how many Zixi is recovering in the live stream. This is an added layer of visibility that is not available from traditional delivery methods.

Zixi offers an end-to-end transport protocol that delivers a high end user experience, at a low cost, from anywhere to anywhere. Content providers can employ Zixi in live event productions including sports, concerts, theater, and more either leveraging the growing Zixi ecosystem of HW and SW tools to originate a Zixi stream or using Zixi running on almost any appliance from non-Zixi embedded devices. The net result is a higher user experience that positively affects revenue while driving down operational costs.