Comparing MPEG AVC and SVC for Adaptive HTTP Streaming

Hari KALVA, Velibor ADZIC, and Borko FURHT

***Abstract*—Adaptive video streaming over HTTP is fast becoming a preferred way for delivering on-demand as well as live video service. The MPEG committee is finalizing a standard for Dynamic Adaptive Streaming over HTTP (DASH) and there is significant industry interest in this solution. DASH assumes/specifies MPEG AVC for content delivery. We analyze the use of MPEG scalable video coding (SVC) for adaptive streaming applications taking the end-to-end service delivery costs into consideration. Our models show that service delivery with SVC costs more and the costs may not outweigh the single bit stream representation and flexibility of SVC for most applications.**

I. Introduction

There has been strong industry interest in delivering video over HTTP. Unlike traditional video on demand (VoD) where video is delivered at a precise bitrate and the underlying network has to guarantee certain level of service, video services over HTTP use the public Internet and hence need to adapt to varying bandwidth and network conditions during playback. The MPEG committee is developing a standard known as Dynamic Adaptive Streaming over HTTP (DASH) to enable such adaptive services over the public Internet [1]. Video services using DASH essentially pre-fetch segments of video and play the downloaded segments continuously. Such prefetching cushions the services against fluctuations in the available bandwidth and also increase the capacity of the delivery system [2]. Since the content is pre-fetched using HTTP, there are no packet losses and hence there is no need to provide error resilience in the bitstream. Furthermore, there is also no need to encode video at a constant bitrate (CBR) as is necessary in services such as broadcasting and realtime video streaming. The duration of segments used in a service is application dependent and the standard requires a bandwidth specification that allows continuous delivery of segments for uninterrupted playback. A stream encoded with an average bitrate over a segment of given duration provides better quality and does not affect playback when delivered using DASH.

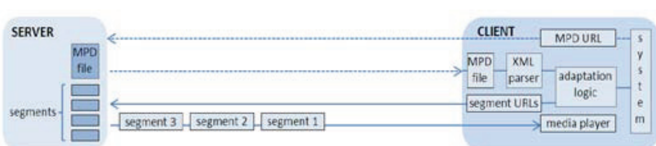

Figure 1. Segmentation and Delivery for DASH

II. Content Encoding for Adaptive Streaming

A. Multi bitrate Video

As network conditions change during a session, the video session has to adapt to the changing conditions. This dynamic adaptation is supported by encoding video segments at multiple bitrates and fetching segments at appropriate bitrate. Segment length plays a role in determining how fast the system can adapt. Shorter segments allow the system to adapt to bandwidth changes faster. However, since an intra (I) frame is needed when segments are switched, shorter segments tend to become more inefficient because of the cost of more frequent I frames. One candidate to consider when supporting multiple bitrates is the MPEG scalable video coding (SVC) [3]. SVC encodes video as a base layer and set of enhancement layers, with each layer increasing the quality over the previous layer. SVC is designed to support multiple bitrates and resolutions efficiently. However, supporting scalability increases the overhead. Compared to MPEG AVC/H.264, SVC uses about 20% more bits to achieve the same quality and this increase is content dependent. While multiple layers of a video coded using SVC can support adaptive bitrates, a separate AVC video bitstream has to be encoded for each supported bitrate.

B. Encoding Complexity

SVC encoding is more complex than AVC encoding. However, this increased complexity can be ignored in the final analysis since the video is encoded just once for DASH like service. Appropriate bitrates are extracted on demand and the extraction complexity is negligible. Similarly, the complexity of encoding AVC video at multiple bitrates can be ignored as this is done once per video. Since a full AVC bitstream is encoded at each bitrate, the cost of storing AVC coded video increases as the number of supported bitrates increase.

III. Cost Analysis and Comparison

The cost of adaptive HTTP service using SVC and AVC is presented in this section. The costs are modeled based on the following assumptions: 90 minute videos are encoded at multiple bitrates using SVC and AVC. For AVC representations, 16 bitrates from 500 Kbps to 8000 Kbps in steps of 500 Kbps are used. For SVC, 16 layers are coded with base layer coded at 500 Kbps and each enhancement layer coded to give equivalent quality of the AVC bitstream. We use a 20% bitrate overhead over AVC. The model does not consider video resolutions or spatial scalability separately as the bitrate is the key factor in evaluating the cost. We acknowledge that the SCV overhead varies depending on the content but the primary relationship of SVC requiring more bits than AVC for the same quality will always hold. A 90 minute video length is used to model Netflix like movie delivery service.

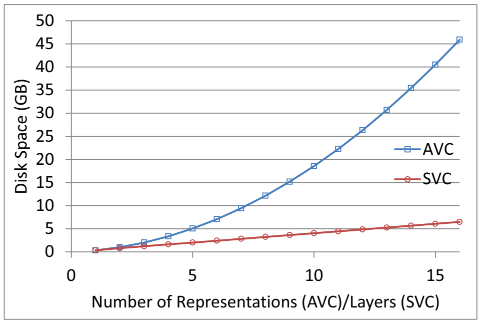

Figure 2. Storage requirements of SVC vs. AVC

Figure 2 shows the storage requirements of SVC and AVC as a function of the number of bitrates supported by the service. The disk space requirements for AVC quickly increase with the number of representations.

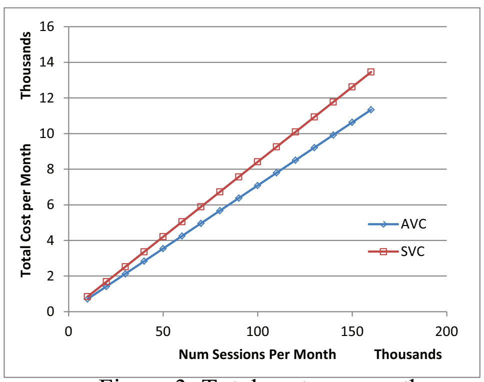

Figure 3. Total cost per month

The cost of a video delivery service is comprised of the cost of storing content, cost of bandwidth, and computing cost of the server. To model the costs we used the current price of storage, bandwidth, and computing charges by the Amazon cloud services (EC2 and S3). The storage cost is $0.125 GB per month and the bandwidth cost is $0.07 per GB. Based on these costs, the cost of delivering video services per month is shown in Figure 3. The cost of SVC increases as the number of sessions per month increase. This model assumes that 5 representations from 500 Kbps to 2500 kbps are used and the sessions are uniformly distributed across bitrates. Since the cost of a session is dominated by the bandwidth costs, SVC sessions cost more. The gains from the smaller storage costs of SVC bit streams are negligible for large number of sessions.

The cost of SVC sessions stays low if the number of sessions remain small. Figure 4 shows the costs of SVC and AVC when using 16 representations. As shown in Figure 4, if the total number of sessions is less than 64, the cost of SVC stays lower than the cost of AVC sessions. SVC is effective for real time conversational services (star topology) where bandwidth among the participants determines the bitrates delivered [4]. In such cases the overhead of SVC over AVC (for the same quality video) becomes negligible when compared to the bitrates required by AVC simulcasting. For on-demand scenarios (one-to-one topology), such as the ones envisioned for DASH, receiver requests the desired bitrate from the server and, hence, the overhead of the scalable video is an unnecessary cost. This remains true even in the case of the live video supported by DASH where the user access still remains one-to-one.

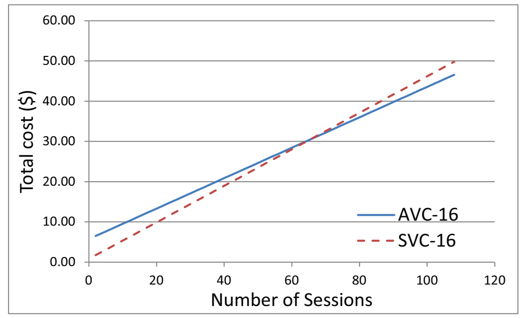

Figure 4. Costs for small number of sessions

IV. CONCLUSIONS

This paper presents a brief introduction to adaptive video services over HTTP and evaluates the cost of MPEG AVC/H.264 and MPEG SVC. While SVC has lower storage requirements than AVC, total costs are lower for AVC. The cost of video delivery today is dominated by the cost of bandwidth. The cost of storage becomes a factor when the number of sessions is very small as is the case for the least popular movies in a catalogue. This presents an interesting opportunity to reduce the costs of a video service by using SVC for videos with fewer sessions. Alternatively, the low-demand videos can be stored using a single representation and real time transcoding used when content is requested.

V. REFERENCES

[1] T. Stockhammer, P. Fröjdh, I. Sodagar, S. Rhyu,(ed.) “Information technology — MPEG systems technologies — Part 6: Dynamic adaptive streaming over HTTP (DASH)”,ISO/IEC, MPEG Draft International Standard, 2011.

[2] H. Kalva and B. Furht, “Techniques for Improving the Capacity of Video-on-Demand Systems,” Proceedings of the 29th Hawaii International Conference on System Sciences HICSS-29), Vol. 2, Jan 3-6 1996.

[3] H. Schwarz, D. Marpe, and T. Wiegand, “Overview of the Scalable Video Coding extension of the H.264/AVC standard,” IEEE Transactions on Circuits Systems Video Technol., vol. 17, no. 9, pp. 1103–1120, Sept. 2007.

[4] A. Eleftheriadis, R. Civanlar, and O. Shapiro, “Multipoint videoconferencing with scalable video coding,” in Journal of Zhejiang University – Science A, vol. 7, no. 5, pp. 696–705, May 2006.