QBR™

QBR BENEFITS
ENHANCING EXISTING ABR STREAMING

DECEMBER 2017

50 Francisco Street, San Francisco, CA 94133
sales@mediamelon.com
+1 415-967-1256

QBR BENEFITS

QBR from MediaMelon enhances existing ABR or adaptive bitrate streaming methods, reducing the volume of data delivered while increasing the quality of experience. Depending on application, this can reduce data requirements by up to 35%, with no change to the video encoder or decoder. This efficiency gain can optionally be used to increase perceived visual quality in scenes where it matters most.

Reducing the volume of data delivered not only lowers expenditure on content delivery networks and other infrastructure for online video services, it also extends their reliable reach over constrained networks, notably for mobile use. This can considerably improve the business case for some services. The relative simplicity and comparatively low cost of implementing QBR can provide a rapid return on investment, while improving the user experience, which extends usage and customer loyalty.

Conventional approaches

There are various possible ways to improve the quality of online video delivery. One option is to increase the efficiency of compression. This might be achieved by modifying the compression parameters, employing an improved implementation of an existing compression standard, or moving to a different compression scheme.

Video compression is highly technical and optimising the balance of visual quality and data rate can be an art as much as a science, typically involving iterative subjective and objective evaluation.

The optimum data rate at which to encode video varies significantly according to the type of material, even from scene to scene. More detail and motion typically requires more data to encode the scene at a given quality. For relatively simple scenes there can be diminishing returns in encoding at higher data rates, with little subjective benefit for the viewer.

It is possible to measure differences between the original source and the encoded output at different data rates to determine objectively the point at which a given visual quality can be achieved. This can vary not only according to the type of material but even from scene to scene or shot to shot.

Packaged recorded media are typically encoded using two passes to determine which scenes can be allowed more data to achieve a given average and enable it to fit on a disc. Variable bit-rate encoding, commonly capped so it does not exceed a maximum data rate, allows the data rate to vary according to complexity.

In the broadcast world it is also common to allow the data rate to vary significantly according to the visual complexity of the programme, spreading the allocated data rate dynamically across multiple channels.

Adaptive Bitrate

In the world of online video, it is common to encode media in short segments at a range of different data rates to enable the player client to adapt to local network conditions and request each segment at the most appropriate rendition.

The encoder is required to produce a set of parallel encodings at different data rates, with relatively little variation around each specified step in average data rate. Each step on the ladder, from low to higher data rates, is determined the configuration of the encoder. Determining which data rates to encode requires careful judgement, often based on best practices that have been generally established.

For instance, media may be encoded at a six different data rates, in steps ranging from 0.4 megabits per second to 4.0 megabits per second. In this case, the lowest data rate produces ten times less data than the highest rate rendition.

The player then selects segments to download at the most appropriate data rate based on an analysis of local conditions. One of the main criteria is the ability to maintain a local storage buffer, based on the rate at which the player is able to receive media segments. If the buffer falls below a certain threshold the player will switch to segments of a specified lower data rate, otherwise it will generally request higher data rate segments.

This adaptive approach has proved very useful in enabling uninterrupted playback of video when the available data rate across the network can vary over time, as is often the case with best efforts delivery over an unmanaged network, where quality of service is not guaranteed.

Notably, these adaptation decisions are based on the quality of the network rather than the quality of the video. The player simply follows a set of fairly fundamental rules or heuristics to maintain its local buffer. The player is entirely unaware of the quality of the picture. All it knows is the nominal average data rate of each of the available renditions.

Conventional adaptive bitrate methods are relatively simple but remarkably inefficient.

Delivering a high quality picture involves making available encoded renditions at appropriately high data rates and hoping that the player will have an opportunity to select these segments whenever it can.

This is clearly inefficient for sections of video where there is relatively little perceptual benefit in encoding at a higher data rate. Depending on the type of material, this may apply to large sections of video, meaning that much of data that is delivered is effectively wasted, as these sections could have been delivered at similar quality at a lower data rate.

Digital compression aims to reduce redundancy, yet the fixed steps of adaptive bitrate encoding introduce the delivery of data at higher rates than may be strictly necessary.

This results in unnecessary overhead for the service provider, network provider and the end user. It means higher content distribution network costs, requires more network infrastructure and affects users with limited data plans or caps, particularly on mobile networks.

Major online video providers have started to tune encoding parameters for particular programmes or genres. For instance, an animation may not require such a high data rate as an action movie. Therefore the set of data rates at which video is encoded can be tuned on a genre-by-genre or event title-by-title basis. Optimising encoding parameters in this way may yield significant savings in data, but might only be practical for the largest service providers. Furthermore, it is not generally practical to tune the different data rates on a scene-by-scene or shot-by-shot basis.

QBR approach

An alternative approach is to optimise the efficiency with which the multi-bitrate outputs of existing encoding methods are delivered. This can be achieved automatically on a second-by-second basis.

This is the approach proposed with QBR, which provides backward compatible optimisation of existing ABR streams. This can work with existing encoders, across a mix of different encoders, and even with pre-encoded material.

The enhancement involves automatically comparing visual characteristics across each of the encoded renditions, without reference to the original source, in order to determine the degree to which any increase in data rate contributes to a perceivable increase in visual quality. This can be performed over a group of pictures in an encoded sequence or for each segment of output.

From this it is possible to predict where a decrease in data rate will not impair the viewing experience, or conversely whether and to what extent an increase in data rate will result in a subjectively noticeable improvement in perceived visual quality.

For sections of video that are hard to compress because they represent complex motion or detail, an increase in data rate may be highly beneficial in maintaining consistent visual quality, while a lower data rate will conversely compromise quality in ways that are very apparent to the viewer.

Some sections of video may be relatively easy to compress because they involve less motion or detail. These can be satisfactorily represented at significantly lower data rates without affecting the apparent visual quality.

QBR effectively provides the benefits of variable bit rate encoding to an adaptive bitrate system.

The variation in data rate can extend from the bottom to the top rung of the ladder of encoded renditions, rather than within relatively narrow bands for each step.

The key to this enhanced approach is to compute the variation in quality between the different bitrate renditions on the server side, where they can be compared to one another. This information can then be used to help steer the player to receive the most appropriate rendition for the visual complexity of that section.

Deployment

QBR can be implemented in different ways, depending on the system architecture.

One way is to shuffle encoded data between the resulting renditions on the server side, substituting sections at higher or lower data rates, at the group of pictures or segment level. This can be achieved either in be repackaging the segments or redirecting requests for particular segments.

Alternatively, the information about the differences in encoding quality between each rendition can be communicated to the player in relatively few bytes, either in the data stream, as part of a manifest file, or as a separate hint file. This allows the player to make better decisions about which segments to request, based not only on local buffer conditions but also on the relative encoding quality of each segment.

In the case of recorded material, this allows the player to plan ahead, deciding where to request segments at lower data rates in order to afford more data for segments of higher visual complexity.

For live media it is evidently not possible to know about the complexity of sections in the future, which have yet to occur, but it may be possible to take advantage of any latency or time offsets to provide information about the visual complexity of whatever coming up next.

Summary

By knowing about the differences in visual quality between encoding renditions at different data rates, it is possible to optimise the overall visual quality for a given average data rate. This can result in dramatic reductions in the data required to represent scenes of low visual complexity, while preserving or enhancing the visual quality of scenes of higher visual complexity.

Unlike other server-side solutions, such as changing encoding parameters or upgrading encoders, this can be implemented transparently while maintaining existing workflows and systems.

QBR automatically optimises visual quality for the available data rate, without requiring changes to encoding on a per title basis.

Furthermore, QBR can achieve this on a per session basis. So different behaviours could be applied for different devices or end users, depending on business rules. For instance, the data rate might be reduced for advertising supported users, while quality might be optimised for a premium service or paying subscriber.

QBR introduces the notion of maintaining and preserving visual quality into adaptive bitrate encoding. It combines the efficiency of variable data rate encoding with the benefits of adaptive bitrate approaches.

Rather than simply providing different renditions at different data rates and letting the player adapt between them based on local conditions, QBR takes into account differences in the perceived quality at different data rates to optimise the volume of data delivered and the resulting user experience.

Quite simply, QBR from MediaMelon is the next generation of adaptive bitrate streaming.

mediamelon.com

50 Francisco Street, San Francisco, CA 94133
sales@mediamelon.com
+1 415-967-1256