ORBX 2 Technical Introduction

October 2013

Summary

The ORBX 2 video codec is a next generation video codec designed specifically to fulfill the requirements for low latency real time video streaming. It provides state of the art performance in both speed and quality, while providing many unique features such as high dynamic range that are not available in existing widely used video codecs.

When compared to the industry standard H.264 video codec, ORBX2 is able to achieve **less than half** the bitrate for similar quality, whilst simultaneously supporting higher speed encoding and decoding.

ORBX2 is supported by ORBX.js, allowing the video to be decoded on modern standard browsers without any plugin or extension.

Features

Low Delay Encoding

ORBX2 was designed from the ground up to support low delay encoding and real time video streaming over the Internet. The encoder has specific optimizations in order to reduce the jitter associated with video transmission over a constrained bandwidth link such as that found in most homes.

Parallel Encoding

The design of the ORBX2 encoder is specially tailored to take advantage of modern massively parallel compute architectures, such as those provided by GPUs and compute accelerators. Much of the encoding process can be implemented on these devices, allowing for extremely fast encoding. This can be directly contrasted to the industry standard H.264, where the design of the codec precludes an implementation on GPUs without significant losses in quality and performance.

Parallel Decoding

The ORBX2 decoder supports multi-threading, allowing it to scale easily on modern multi-core processor architectures. Much like the encoder, the decoder also supports implementation on massively parallel compute architectures.

Low Complexity Decoder

The decoder for ORBX2 has intentionally been kept as low complexity as possible, without sacrificing on quality. This allows the codec to support decoding of HD resolutions even on low end platforms, such as cellphones and tablets.

Next Gen Resolutions

With GPU based encoding, ORBX2 is able to achieve 4K and even 8K encoding in real time with low latency and low jitter. This is something that is simply not achievable with any other codec.

High Bitdepth

ORBX2 supports channels with up to 12 bits of data, without requiring any extensions to the format, or special builds.

Flexible Multi-Channel Support

ORBX2 was designed with applications beyond just simple color video streaming. The codec already supports depth + alpha streams such as those produced by Kinect, and motion point cloud data. The built in support for flexible sets of image planes with independent attributes, makes ORBX2 able to be directly leveraged in many new and exciting applications.

ORBX.js HTML5 Decoder

One of the most unique features of ORBX2 is that decoding this codec is supported by ORBX.js. The ORBX.js library is a pure javascript library that supports decoding of video streams, and managing connections within standards compliant HTML5 compatible browsers. This means, ORBX2 video can be streamed directly to a modern browser without the need for Flash, or any sort of plugin or extension.

Benchmark Results

The following section provides results comparing ORBX2 with H.264 in low-delay encoding using a short clip from the well known Big Buck Bunny video. *(See appendix for detailed benchmark setup information)*

PSNR Comparison

The first test is using the well known PSNR metric. The test was performed across the entire quality and bitrate range that each codec was able to achieve.

The first graph shows the results for PSNR calculated across all color image planes, providing an indication of color reproduction quality.

PSNR-YUV for ORBX2 vs H.264

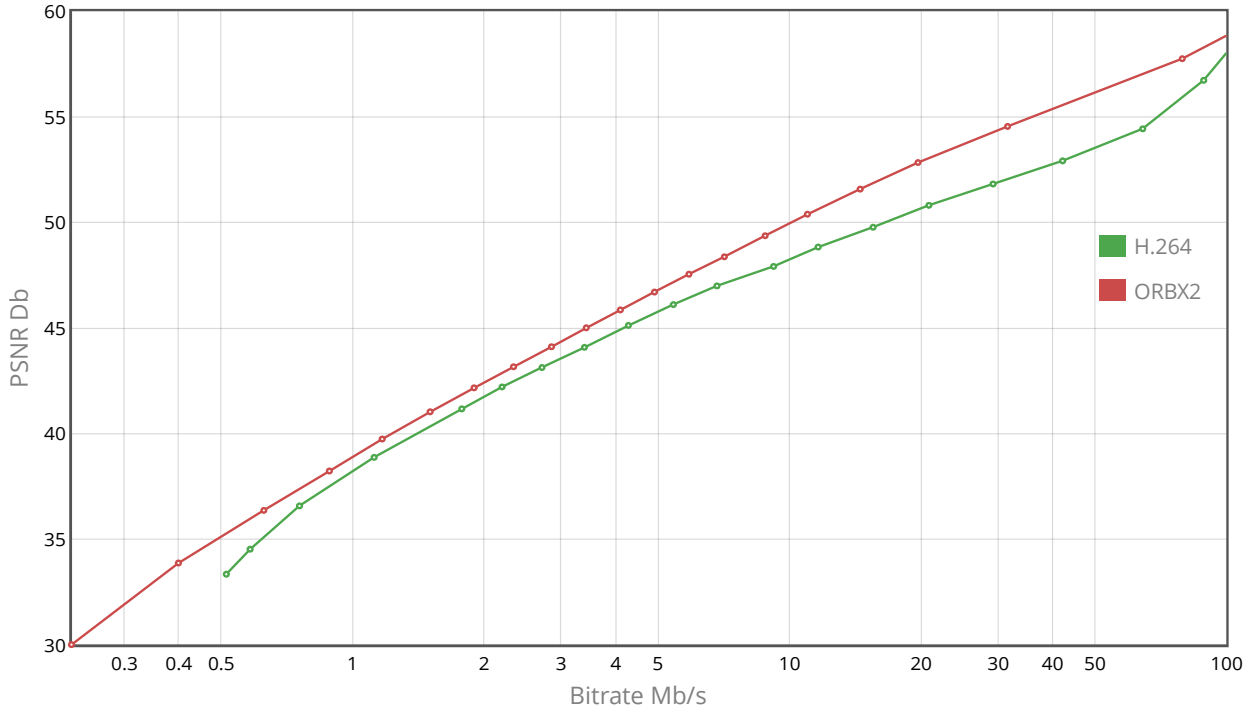

The next graph has PSNR calculated only on the luminance channel.

PSNR-Y for ORBX2 vs H.264

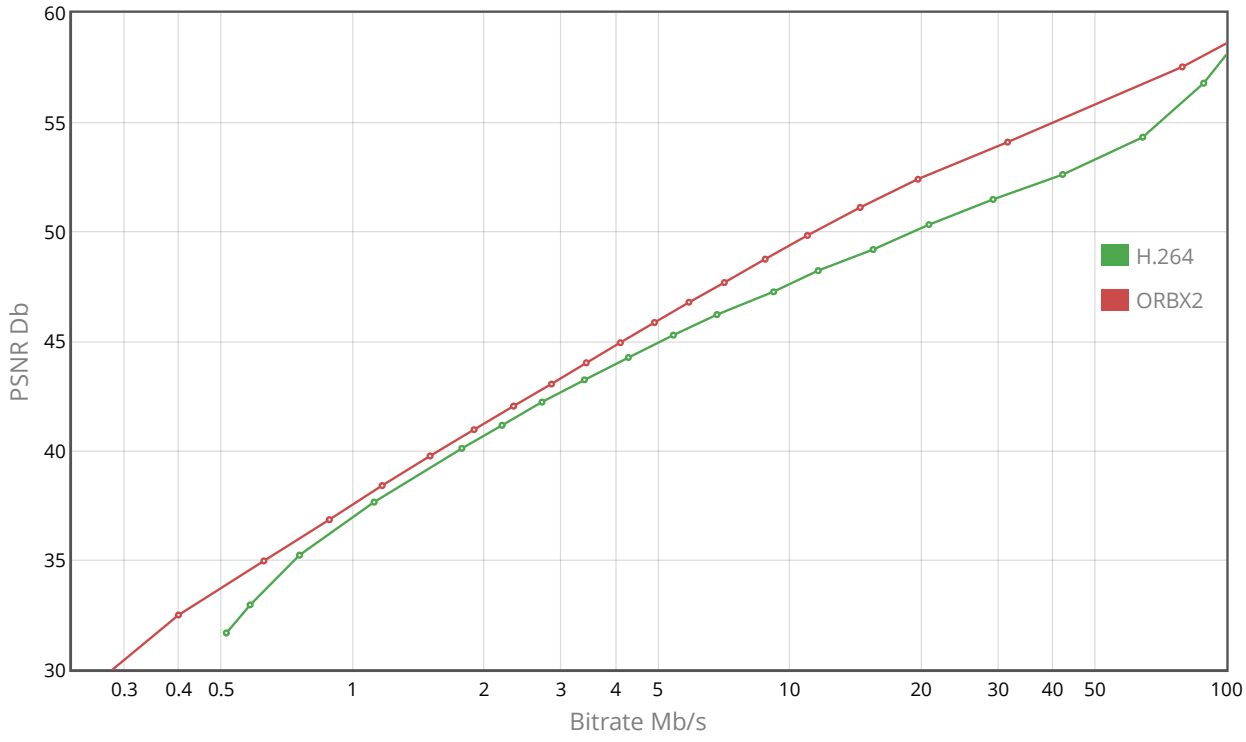

Analysis

From these graphs, it can be seen that for high quality streams of 10Mb/s or greater, ORBX2 provides as little as **half** the bitrate for the same objective quality. At its very worst, ORBX2 still achieves **85%** the bitrate of H.264 for the same quality.

PSNR-HVS Comparison

The PSNR metric is widely used for objective comparisons in video and image compression mostly due to its simplicity. However it is well known that it doesn't correlate well with human visual perception as a measure of visual quality. For this reason, researchers came up with an alternative metric called PSNR-HVS, which is modeled on the Human Visual System (http://www.ponomarenko.info/psnrhvsm.htm).

The next graph shows the metric PSNR-HVS calculated on the luminance channel.

PSNR-HVS-Y for ORBX2 vs H.264

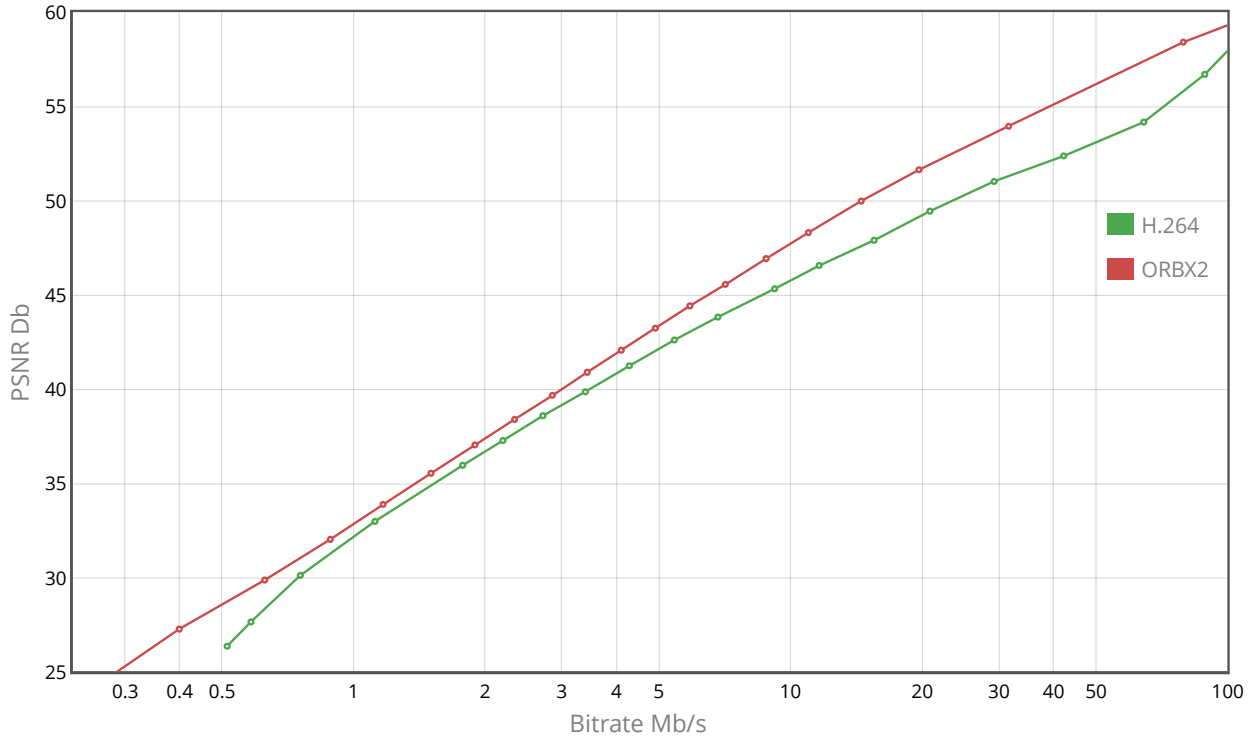

Analysis

The results here mirror those from the PSNR metric. Because ORBX2 is designed to achieve high subjective visual quality, rather than just objective quality, ORBX2 is able to maintain its lead. The relative success of H.264 in this test can be seen as a testament to the many years of tuning applied to the x264 encoder.

SSIM Comparison

The SSIM metric has become popular recently as an alternative metric to PSNR for comparing video codecs. As a metric, it is generally regarded as being better than PSNR in representing visual quality, although recent research indicates that it is not as good as PSNR-HVS. It should be noted that different SSIM implementations can produce markedly different results, due to the choice of parameters and the precision of floating point arithmetic.

This next graph shows the comparison using the SSIM metric calculated on the luminance channel.

SSIM-Y for ORBX2 vs H.264

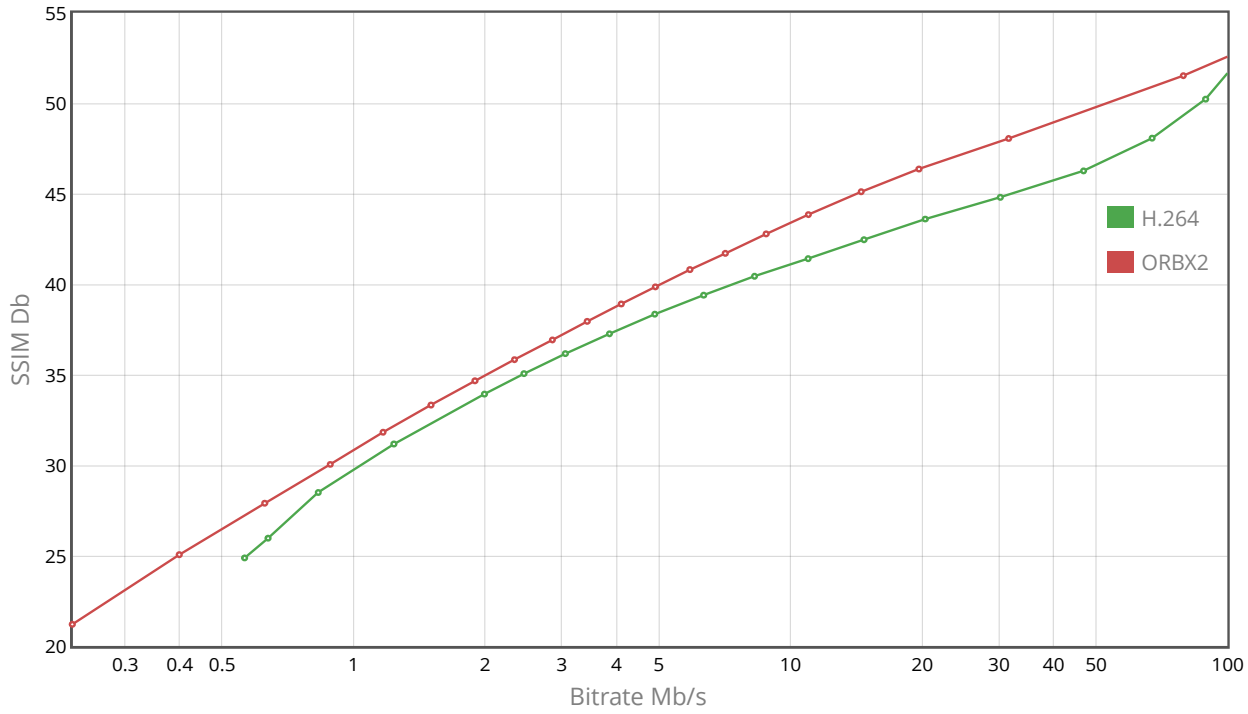

Analysis

The results again mirror those provided by the other metrics, but this time with ORBX2 further ahead of H.264. ORBX2 is able to achieve **less than half** the bitrate for the same quality. A 10Mb/s ORBX2 stream produces equivalent quality to that of a 20Mb/s H.264 stream.

Image Quality Comparison

The following figure compares the image quality of ORBX2 with that of H.264. Both images were encoded at a low quality, and to the same size in order to illuminate the differences.

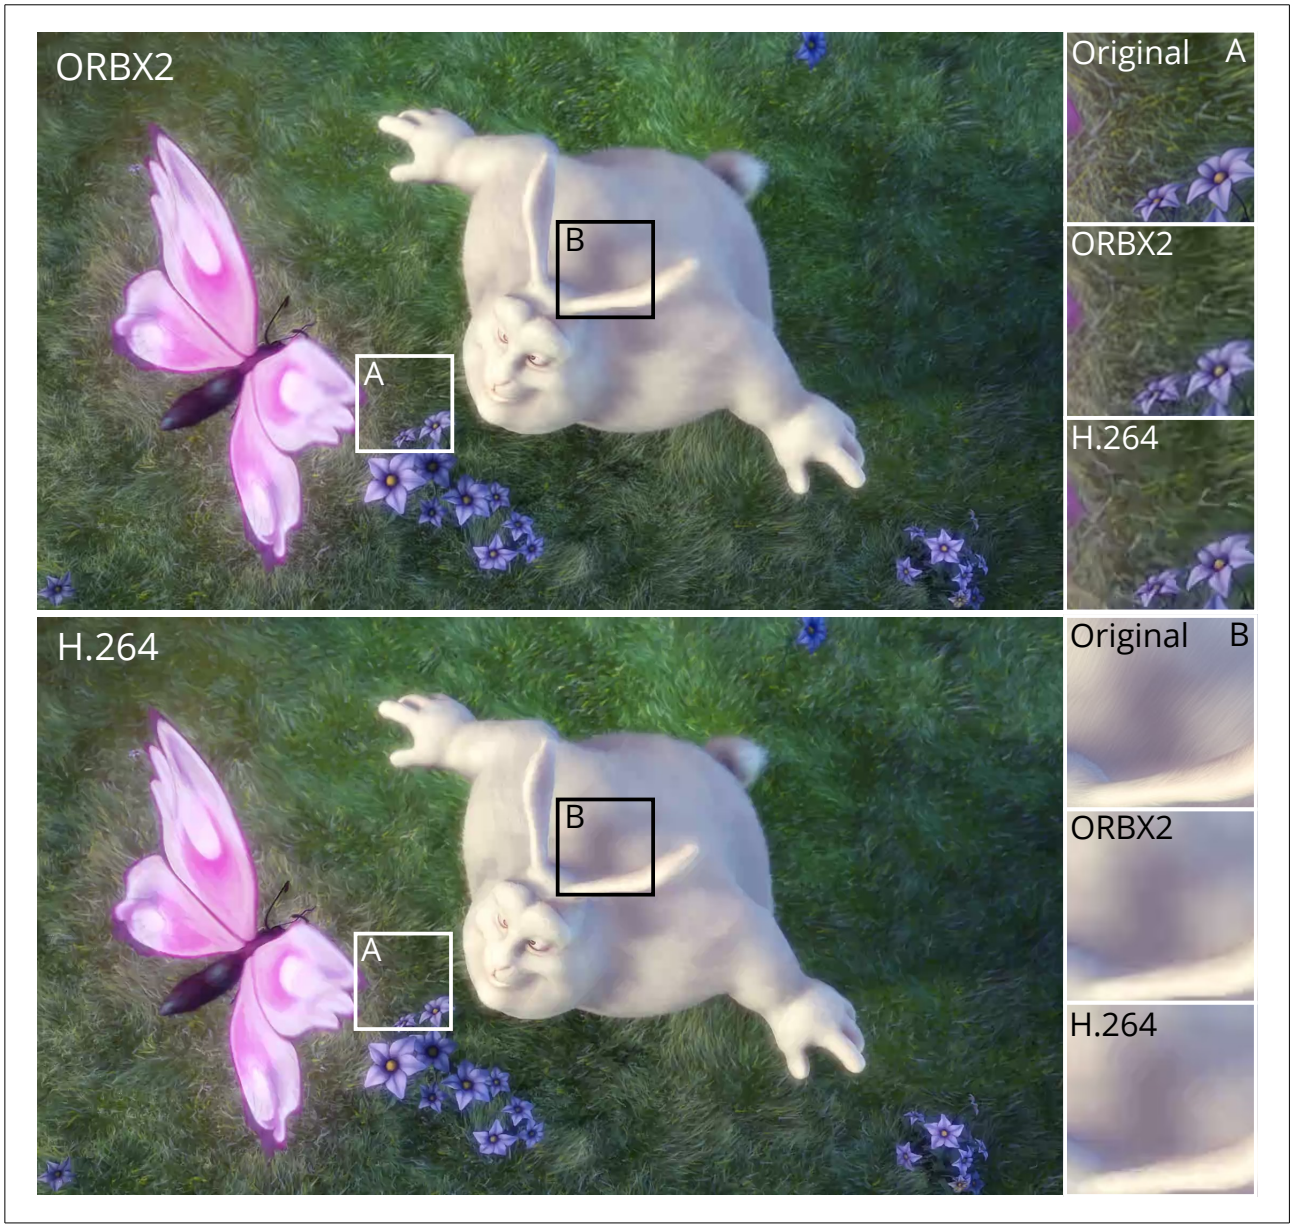

Better Texture Detail

The rectangle A provides an example of the performance of the codecs on a textured region of the image. The rectangle can be seen magnified on the top right edge of the figure, showing the original image and the decoded results from both ORBX2 and H.264. This example shows that ORBX2 retains a greater amount of texture detail than H.264 despite the low bitrate for the encoding.

Better Smooth Gradients

The rectangle B shows an example of a smooth region of the image. In the magnified rectangles at the bottom right, H.264 can be seen to exhibit banding, whereas ORBX2 retains the smooth appearance of

the original.

Overall Subjective Image Quality

Subjective image quality is difficult to measure accurately, which is why this example is necessary. However, considering the images as a whole, it is clear to see that ORBX2 provides both better overall detail, as well as a more uniform, natural and visually pleasing image. Because of the design of the codec, H.264 often causes lower bitrate content to look as though it were a water color painting. ORBX2 is able to avoid this effect completely, producing more naturally looking images.

Appendix: Benchmark Setup

The following sections provide detailed disclosure of the benchmark setup used in producing the benchmark results from the previous section.

Source Video

The video source is the well known Big Buck Bunny clip at 1080p24, with all the tests performed on 500 frames starting with frame number 1500 in the sequence. The sequence can be obtained at http://media.xiph.org/BBB/BBB-1080-png/

The source video was converted to YV12 before encoding so as to remove any possible color space conversion confusion from the testing.

Encoding Parameters for H.264

The H.264 encoding was performed by a build of x264 dated 16th May 2013. The 'slow' profile was chosen, with b-frames disabled to achieve low-delay encoding. For the SSIM comparison, the SSIM tune was chosen, and for all other tests the PSNR tune was chosen.

Encoding Parameters for ORBX2

The ORBX2 encodings used the default high quality encoding profile for all tests. There are no special PSNR or SSIM tunings used.

Metric Calculation

All metrics are calculated individually on each frame, and the average of the metric over all frames in the sequence is reported.

When computing the PSNR-YUV metric, a weighting of 1/3 given to the chrominance channels, to be conservatively in line with the human visual system color sensitivity.