

Daala Update

IETF 97 (서울시)


Progress Since Berlin

- Fixed-point PVQ version completed
- PVQ ported to AV1
- Daala entropy coder now the default in AV1
- Improvements in metrics tools
 - 10 bit support for all metrics
 - Found and fixed an overflow for FastSSIM


FastSSIM overflow

- Identified and fixed an integer overflow in FastSSIM
- Only affected video around 1080p and larger
 - That includes some videos in ntt-short-1
 - Still image sets were ~1 megapixel, so not affected


Previously Reported Fast MS-SSIM January 2014 to April 2016



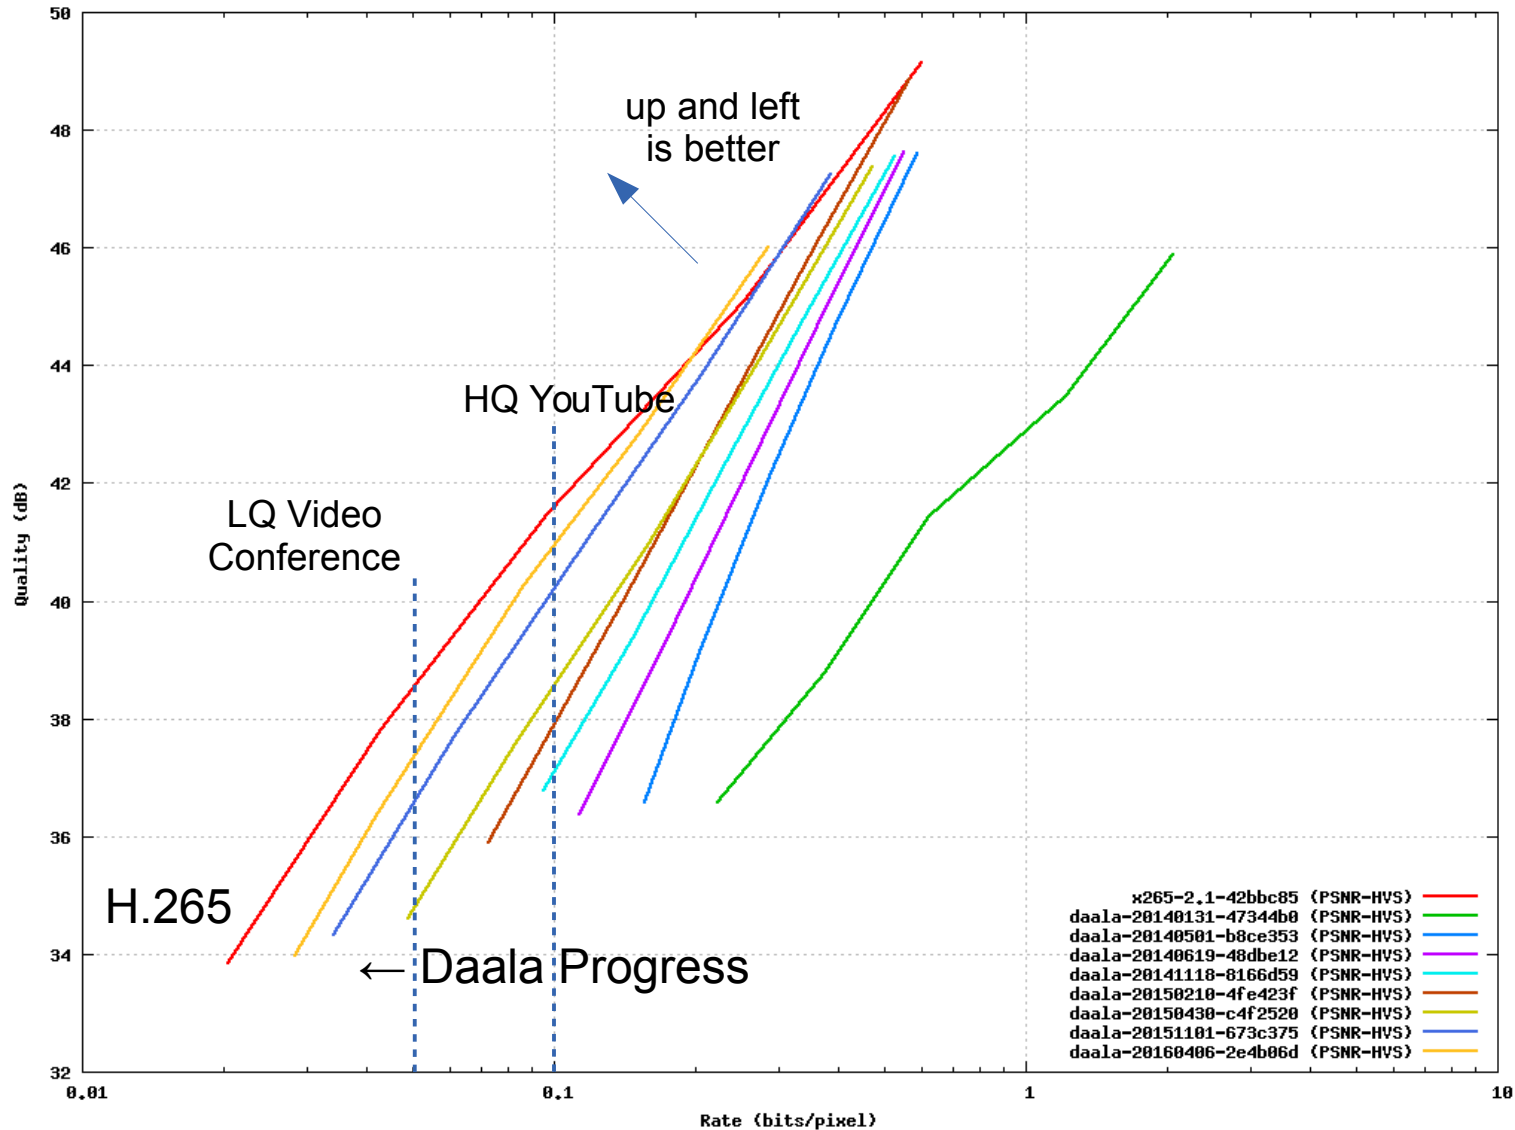
Corrected Fast MS-SSIM January 2014 to April 2016



objective-1-fast MS-SSIM January 2014 to April 2016



objective-1-fast PSNR-HVS January 2014 to April 2016



Summary Since Berlin

- 60 commits
- 6-9% encoder speed improvements
- Aggregate BD-rate results

	RATE (%)
PSNR	-0.2433
PSNRHVS	-0.2610
SSIM	-0.2055
MSSSIM	-0.1850
CIEDE2000	-0.2087

(objective-1-fast, default options)

Changes

Fixed-Point PVQ

- 27 of the 60 commits (45%)
- Overall BD-rate changes very small
- Encoder still uses doubles for search
 - Not critical for standardization (non-normative)
 - We may always use them on some platforms
 - E.g., x86 does not have integer SIMD reciprocal square root approximation (but ARM does), helpful for RDO

PVQ Encoder Speed Improvements

- Using approximate rate for making RDO decisions
 - 0.3% regression
- Order gain/theta candidates by K, start search for next candidate from result of previous
 - And skip the search entirely when K doesn't change
 - Small (<0.1%) improvement
- Don't consider gain/theta candidates that can't possibly have lower distortion than "skip"
 - No metrics change

Deringing Filter Changes (AV1)

- Use shorter (5-tap) filters for chroma
 - Big improvements for the Minecraft sequence
 - Overall CIEDE change 0.01%
- Now works with 4:2:2 input
 - By disabling chroma deringing
- Simpler threshold calculation for the second filter
 - Based on total change made by first filter
 - No longer need the per-pixel change in the second pass
 - Small (0.05%...0.1%) regression

Deringing Speed-Ups (AV1)

- Landed SIMD code from Daala
- Greatly reduced the amount of copying and buffering
- Stop running at all when threshold == 0
- Decoder speed impact went from ~50% → 8.5%
 - 3.2% for the actual filter
 - 5.3% for buffering/copying/management
- Still room for improvement
 - Buffer copies are not SIMD
 - Could combine with deblocking to reduce cache misses

Q15 Entropy Coder Adaptation

- Wrote down what we described in Berlin in `draft-terriberry-netvc-codingtools-01`
- Showed it to some of our hardware partners
 - No one had any heartburn
- Still may change/improve going forward

Future Plans

- Enable activity masking, quantization matrices for PVQ
 - Including further tuning for visual quality
- Porting Chroma-from-Luma to AV1
- Porting Daala's rate control to AV1
- Improved integration of Daala tools in AV1

Questions?