

Codec updates & comparisons

High Efficiency, Moderate Complexity

Video Codec using only RF IPR

(<https://datatracker.ietf.org/ipr/2636/>)

draft-fuldseth-netvc-thor-03

draft-samuelsson-netvc-xvc-01

Steinar Midtskogen (Cisco)

IETF 102 – Montreal, CA – July 2018

Thor status

- No bitstream changes since IETF101
- Intrinsic library overhaul
 - Added optimisations for ARMv8/aarch64
 - Various minor SIMD improvements for x86 and ARM
 - One bugfix which caused incorrect ARM behaviour
 - Some new intrinsics added
 - Fixes to remove warnings or errors generated by some compilers
- Still missing: Daala entropy coder
Decoder robustness

SIMD optimisations (recap)

- The ideas behind the intrinsics library used in Thor are:
 - write SIMD optimised code once for all architectures and compilers
 - hide low level compiler specific code in the library
- Intrinsics that are specific for one architecture and difficult to replicate on others are avoided
 - The final assembly may not always be optimal, but usually “good enough”.

SIMD optimisations

- The limited support for highly specialised intrinsics in the library forces the codec implementor to write code that runs reasonably well on all architectures
- If WebAssembly (wasm) gets support for integer SIMD, it should be straight-forward to port the intrinsics library to use this
- The decoder itself can then be transmitted with the bitstream and run at near native speed in browsers
 - allows quick deployment of codec changes
 - IPR issues can be addressed quickly

XVC vs XVC RF vs Thor

- The following slides compare XVC with the royalty free profile of XVC (draft-samuels-son-netvc-xvc-01)
- 8 sequences, 10 seconds each, 4 qp values/bitrates:
 - Kimono (1920x1080 24Hz)
 - BasketballDrive (1920x1080 50Hz)
 - BQTerrance (1920x1080 60 Hz)
 - FourPeople (1280x720 60 Hz)
 - Johnny (1280x720 60 Hz)
 - ChangeSeats (1920x1080 60Hz)
 - HeadAndShoulder (1920x1080 60Hz)
 - TelePresence (1920x1080 60Hz)
- Same test set used to benchmark Thor before AWCY
- High latency setting, single pass

XVC vs XVC RF vs Thor

- Thor configuration (445c7c4):
 - config_HDB16_high_efficiency +
-enable_cfl_inter 1 -bitdepth 12 -cdef 1
 - qp 25 30 35 40
- XVC configuration (283009f):
 - speed-mode 0,1,2 -tune 1-chroma-qp-offset-u -1
-chroma-qp-offset-v -1
 - qp 25 30 35 40
 - 51 restriction flags set to 1 for RF profile

XVC (anchor) vs XVC RF

- speed-mode 1:

Sequence	BDR-Y	BDR-U	BDR-V
Kimono	11.592	8.000	10.072
BasketballDrive	13.429	15.385	14.833
BQTerrace	14.798	23.722	20.657
FourPeople	10.386	17.720	17.392
Johnny	14.641	22.099	19.632
ChangeSeats	15.068	5.776	7.032
HeadAndShoulder	18.550	4.767	14.743
TelePresence	18.903	15.073	9.927
Average	14.671	14.068	14.286

XVC (anchor) vs XVC RF

- speed-mode 2:

Sequence	BDR-Y	BDR-U	BDR-V
Kimono	11.342	7.914	9.769
BasketballDrive	10.546	12.489	11.344
BQTerrace	13.882	21.718	23.671
FourPeople	11.262	18.319	17.101
Johnny	17.091	22.515	21.543
ChangeSeats	11.492	4.415	7.260
HeadAndShoulder	18.469	-0.951	19.058
TelePresence	15.829	15.336	14.203
Average	13.739	12.719	15.493

Thor (anchor) vs XVC RF

- Speed-mode 0 (514 % more running time than Thor):

Sequence	BDR-Y	BDR-U	BDR-V
Kimono	-15.823	-24.150	-19.094
BasketballDrive	-29.579	-34.379	-34.402
BQTerrace	-19.543	-26.591	-26.627
FourPeople	-11.727	-17.792	-16.819
Johnny	-9.672	-16.701	-18.153
ChangeSeats	-14.867	-14.100	-18.553
HeadAndShoulder	1.086	10.059	-8.424
TelePresence	-8.285	-20.560	-11.610
Average	-13.551	-18.027	-19.210

Thor: 1.81 fpm (108 fph) XVC: 0.29 fpm (17.6 fph)

Thor (anchor) vs XVC RF

- speed-mode 1 (3% more running time than Thor):

Sequence	BDR-Y	BDR-U	BDR-V
Kimono	-13.287	-23.957	-18.512
BasketballDrive	-25.192	-31.079	-30.263
BQTerrace	-12.492	-20.994	-21.380
FourPeople	-7.110	-12.844	-11.659
Johnny	-4.381	-8.767	-11.392
ChangeSeats	-9.176	-9.158	-16.860
HeadAndShoulder	7.249	18.751	0.136
TelePresence	-3.561	-16.383	-7.136
Average	-8.494	-13.054	-14.633

Thor: 1.81 fpm

XVC: 1.76 fpm

Thor (anchor) vs XVC RF

- speed-mode 2 (62% less running time than Thor):

Sequence	BDR-Y	BDR-U	BDR-V
Kimono	-1.778	-7.362	-5.466
BasketballDrive	-16.119	-18.268	-20.432
BQTerrace	4.683	-5.454	-5.611
FourPeople	5.898	1.595	0.143
Johnny	25.249	19.158	23.398
ChangeSeats	5.606	11.199	3.166
HeadAndShoulder	43.380	78.651	42.618
TelePresence	20.174	19.649	26.415
Average	10.887	12.396	8.029

Thor: 1.81 fpm

XVC: 4.75 fpm

Codebase complexity

- Thor has 34,836 lines of code (21,584 without SIMD)
- XVC has 25,593 lines of code
- AV1 has 211,200 lines of code (195,341 without SIMD)


AV1 status

- The spec is frozen
- AV1 (2018-07-15) vs. VP9 (2017-11-02):
 - PSNR-Y: -31.7%
 - PSNR-U: -38.4%
 - PSNY-V: -40.6%
 - CIEDE2000: -34.0%
 - APSNR: -32.1%
 - MS SSIM: -31.0%
 - **150x encoding time**

AV1 compression history

- Compression/speed relationships measured using AWCY
 - Mixed content: objective-1-fast
 - <1% BDR improvement since IETF101 but ~2x speed-up
- Low delay configuration
- BDR anchor is AV1 in July 2016, roughly equivalent to VP9
- Note that the speed axis is logarithmic

AV1 compression history



AV1 complexity history

