The video codecs behind modern remote display protocols

Rody Kossen

PQR

Rody Kossen - Consultant

Rody.Kossen@pqr.nl

@r_kossen

rody-kossen-186b4b40

www.rodykossen.com

PQR

Agenda

- The basics
- H264 – The magician
- Other codecs
- Hardware support (NVIDIA)

The basics

Let’s start with some terms

RGB

HEVC

Bit Depth

YCbCr

YUV 4:4:4

BT. 2020

H264

H265

Chroma Subsampling

Lossy

Color Gamut

LossLess

VP9

YUV 4:2:0

AV1

REC 709

YUV

Colors

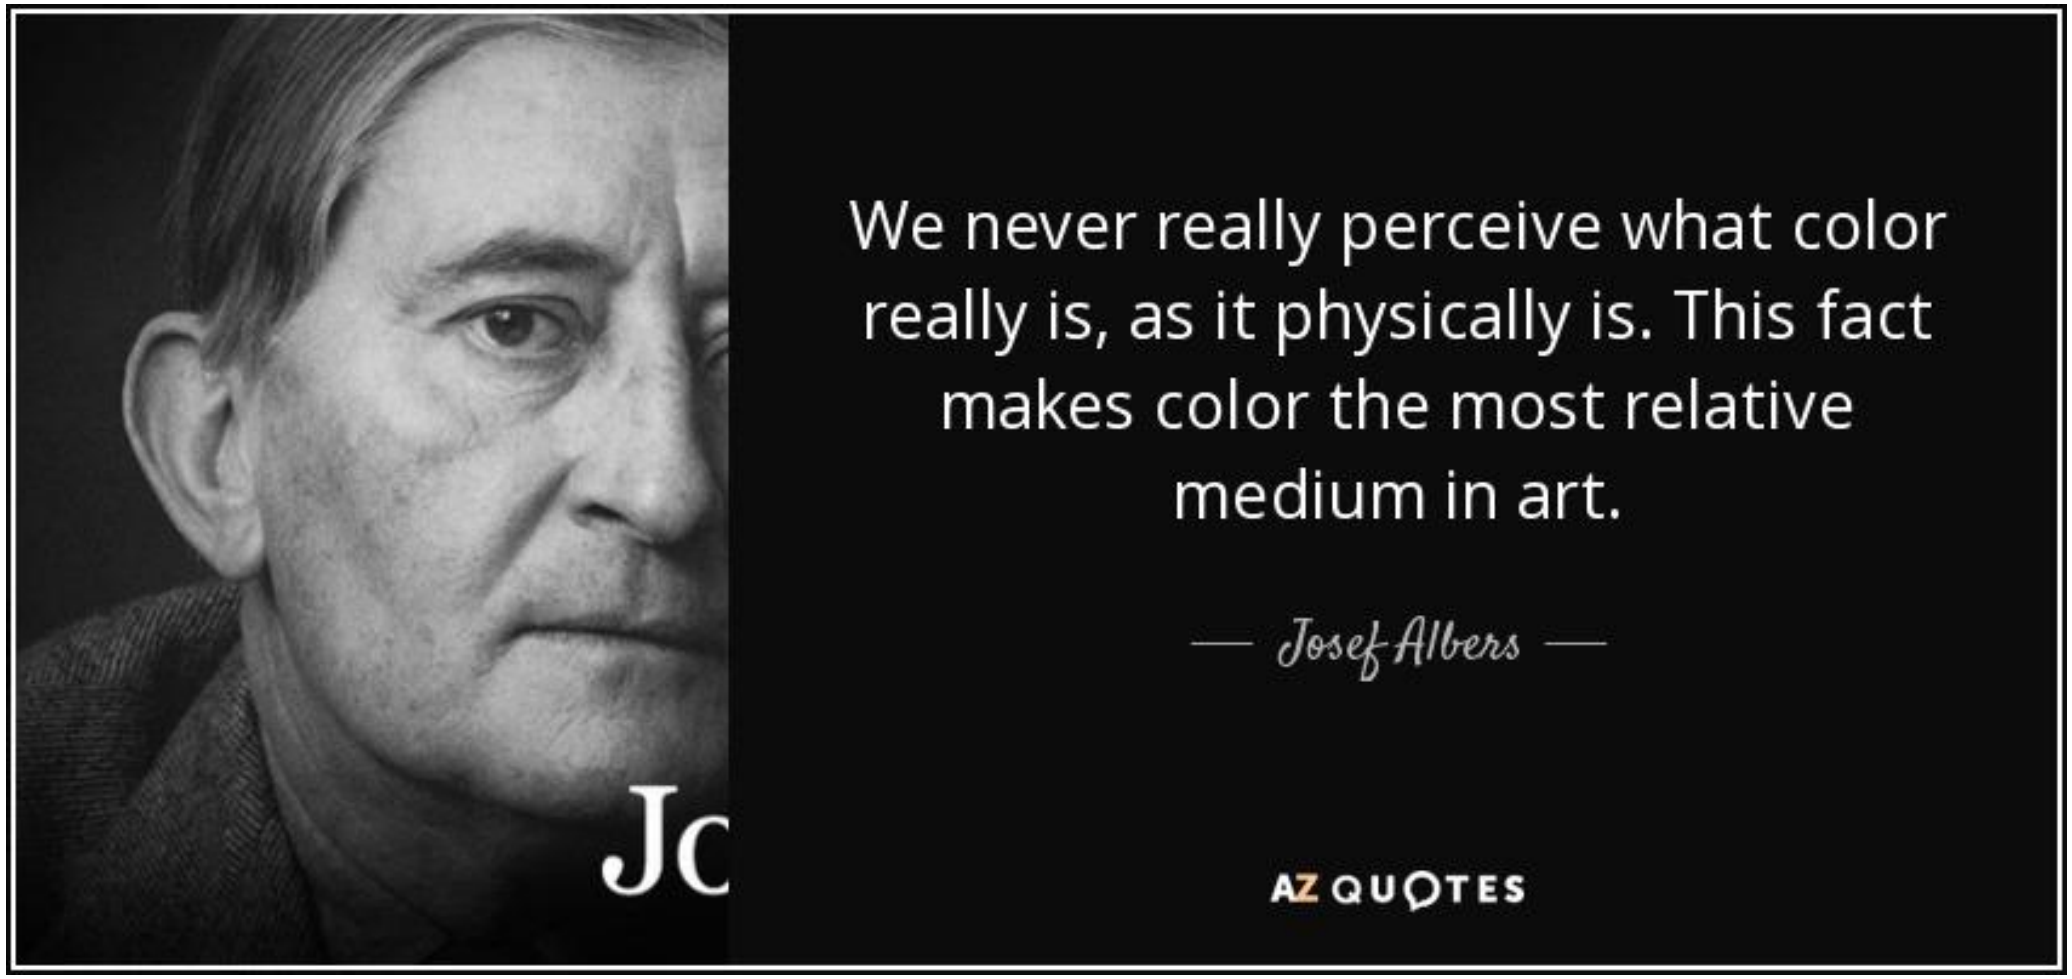
We never really perceive what color really is, as it physically is. This fact makes color the most relative medium in art.
— Josef Albers —
AZ QUOTES

PQR

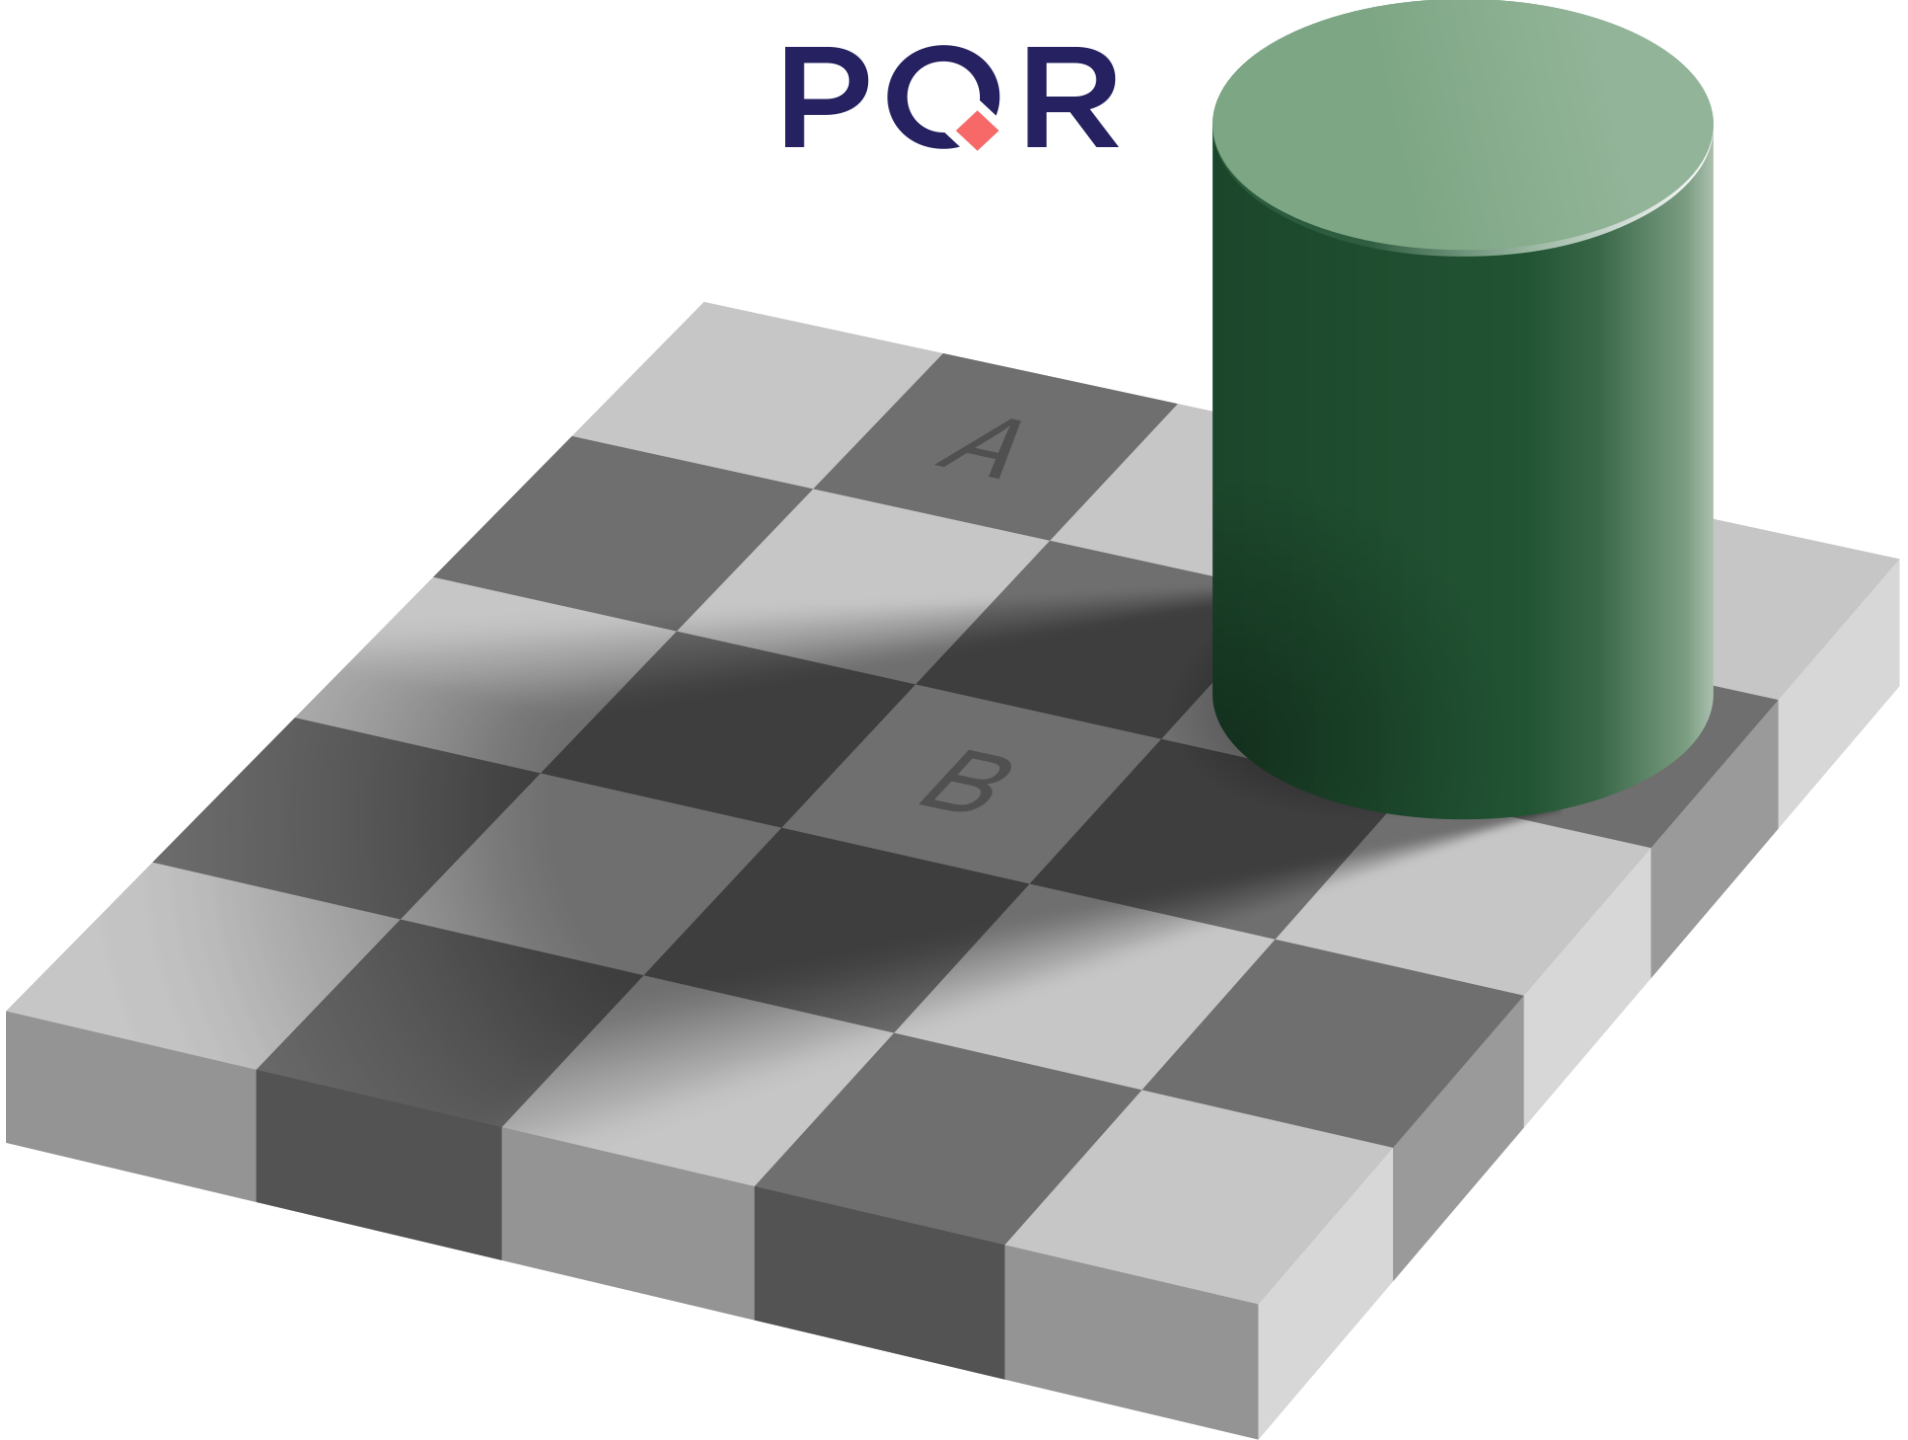

PQR

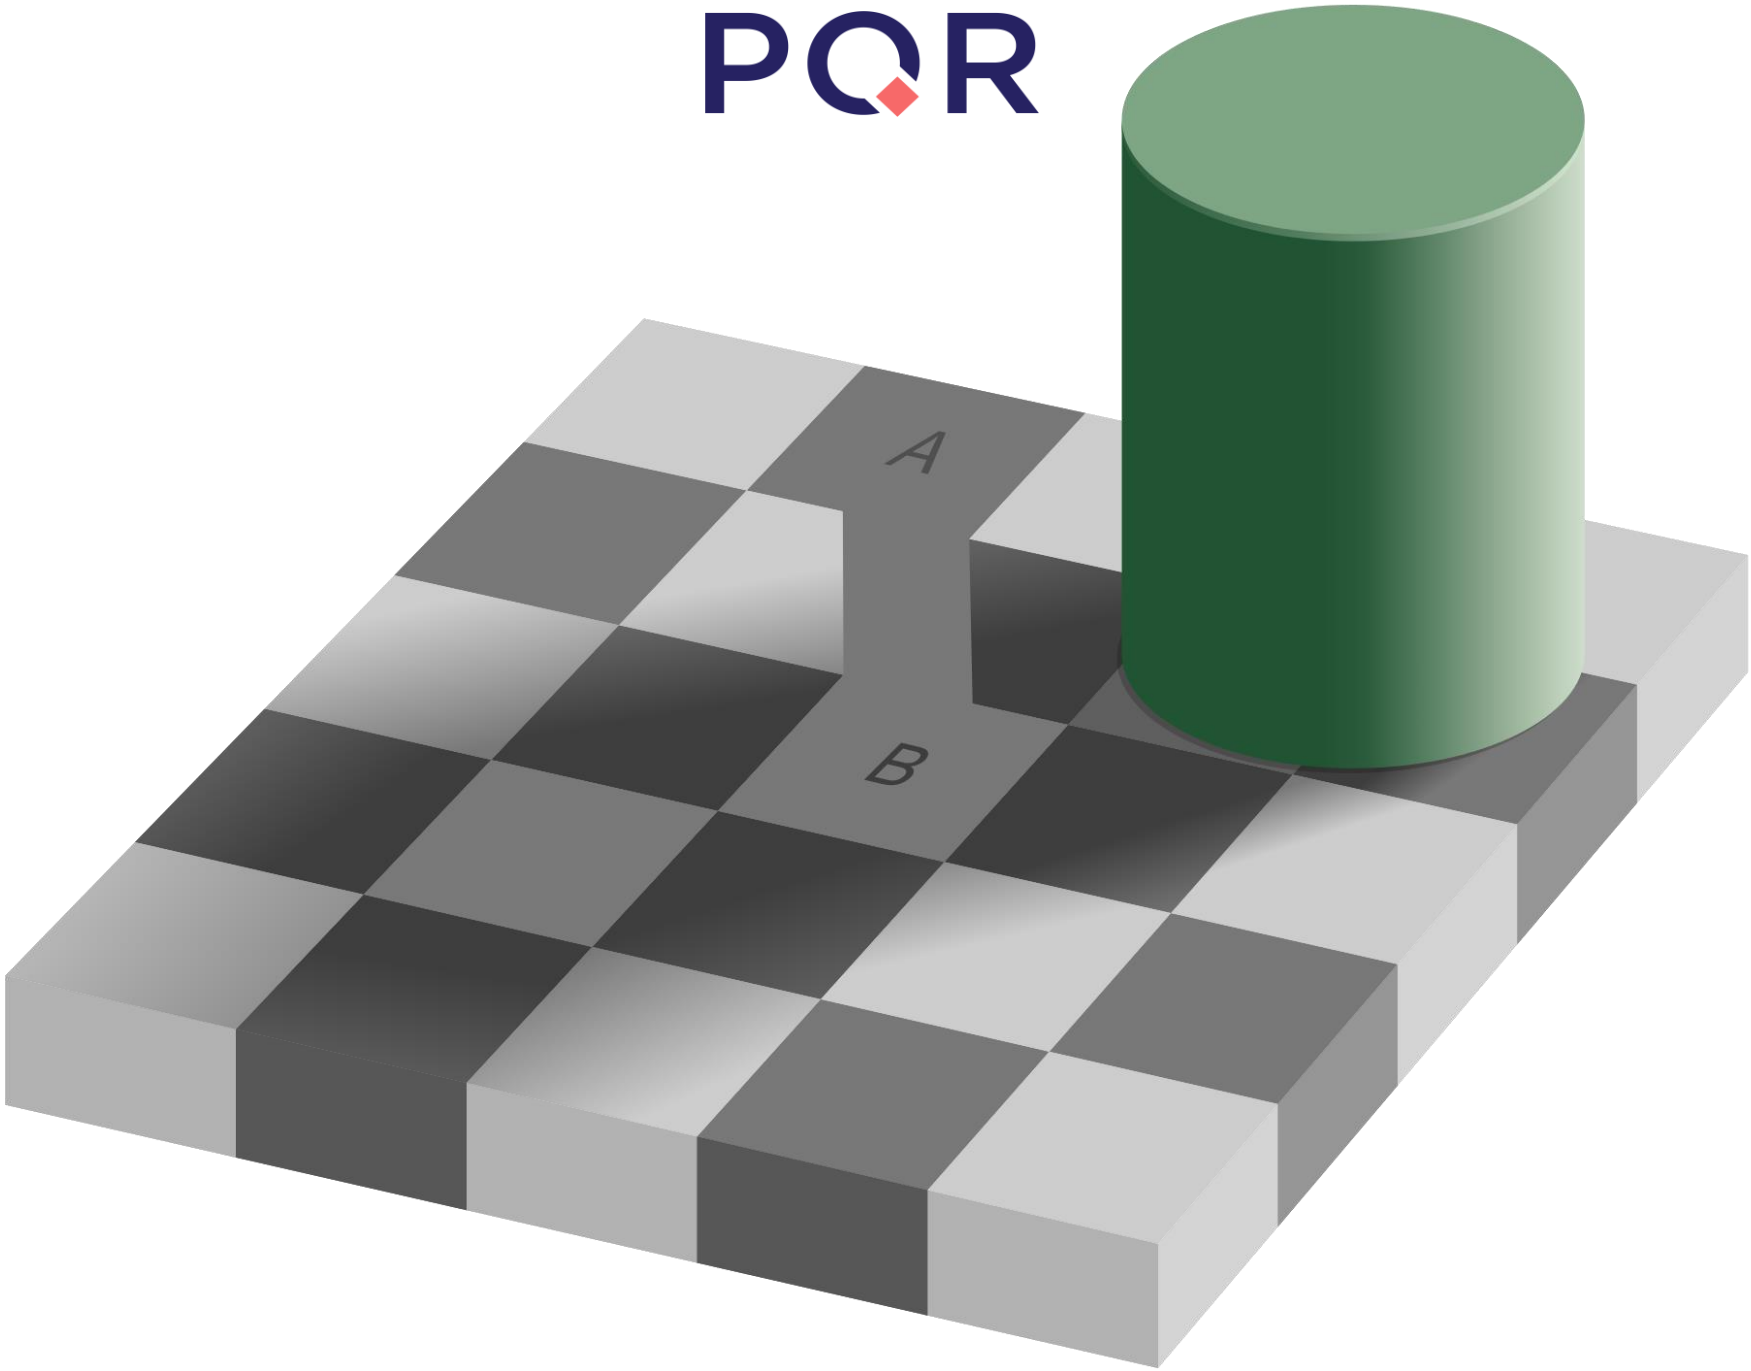

Visible colors

Color vision table

State	Types of cone cells	Approx. number of colors perceived	Carriers
Monochromacy	1	200	Marine mammals, owl monkey, Australian sea lion, achromat primates
Dichromacy	2	40,000	Most terrestrial non-primate mammals, color blind primates
Trichromacy	3	10 million[47]	Most primates, especially great apes (such as humans), marsupials, some insects (such as honeybees)
Tetrachromacy	4	100 million	Most reptiles, amphibians, birds and insects, rarely humans
Pentachromacy	5	10 billion	Some insects (specific species of butterflies), some birds (pigeons for instance)

Color Bit-Depth = Amount of different colors

- Used in 2 ways:
  - The number of bits used to indicate the color of a single pixel, for example 24 bit
  - Number of bits used for each color component of a single pixel (**R**ed **G**reen **B**lue), for example 8 bit

- Can be combined with Alpha, which describes transparency

- Most commonly used:
  - High Color – **16** bit = **5** bits per color + 1 unused bit = 32.768 colors = 2^15
  - True Color – **24** bit Color + 8 bit Alpha = **8** bits per color = 16.777.216 colors
  - Deep color – **30** bit Color + 10 bit Alpha = **10** bits per color = 1.073 billion colors

Color Bit-Depth

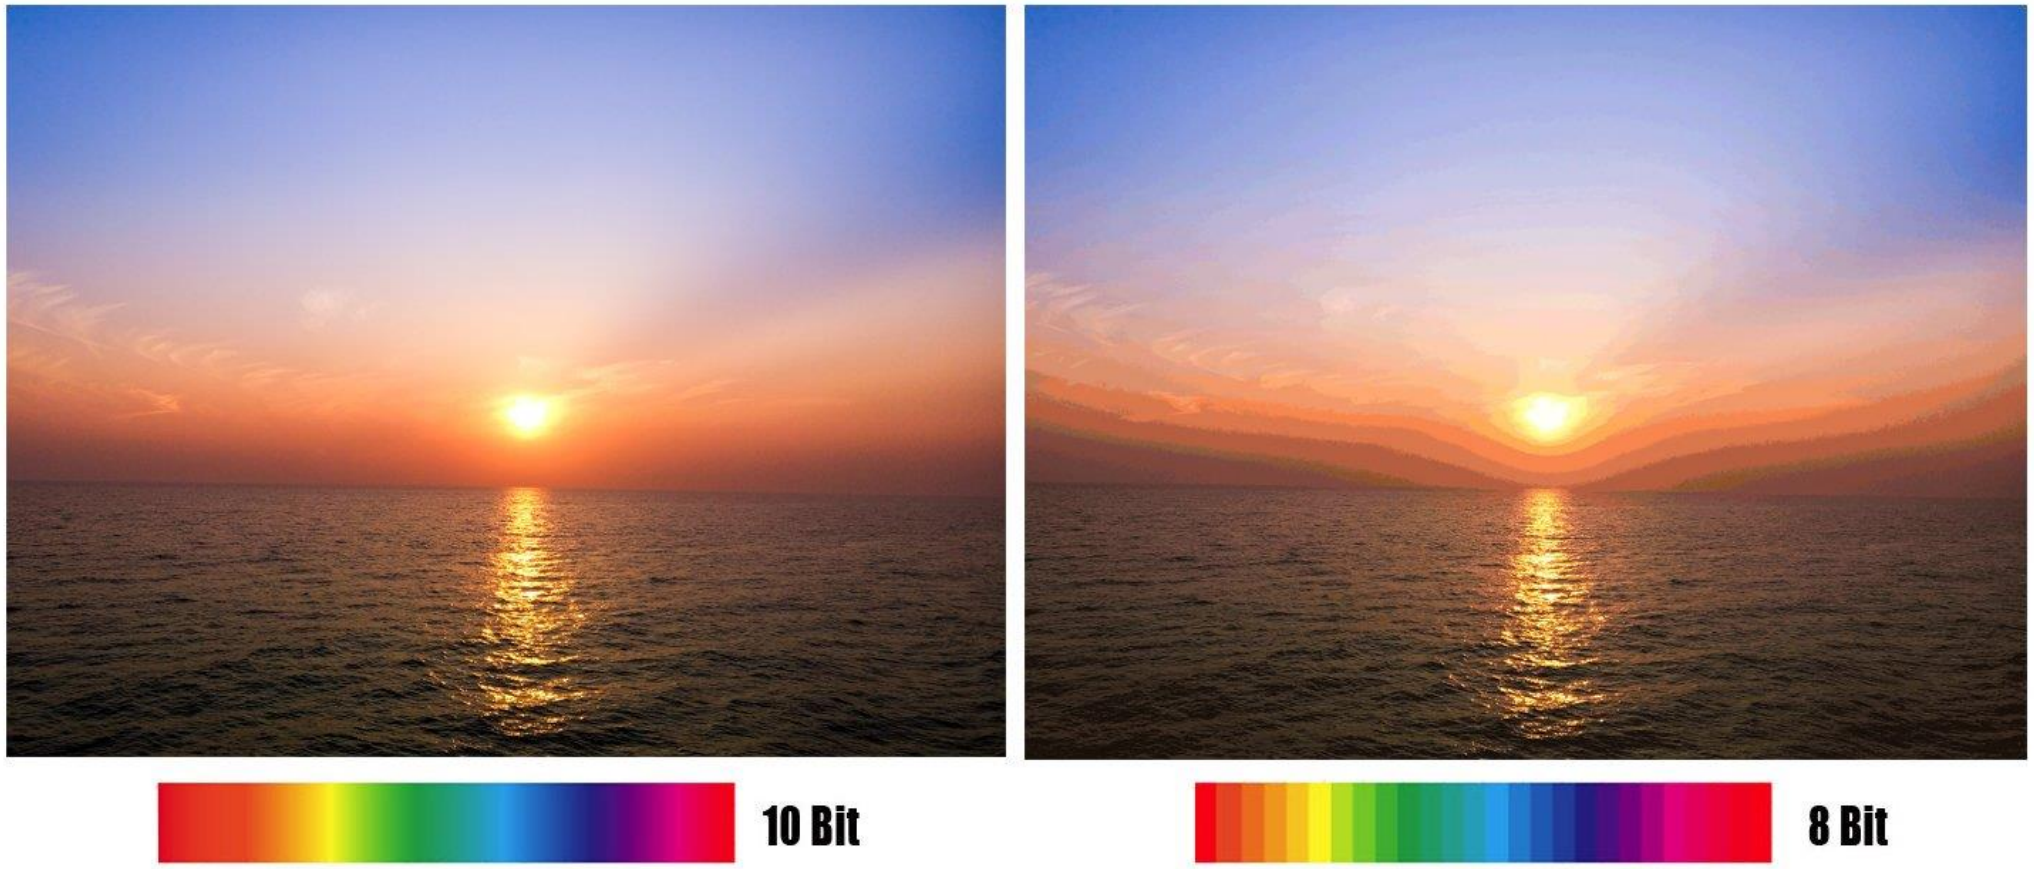

Color Gamut

- Describes the subset of colors which can be displayed in relation to the human eye
- Standardized gamut:
  - Rec.709 (Blu-ray)
  - Rec.2020 (Ultra HD Blu-ray)
  - Adobe RGB
  - sRBG

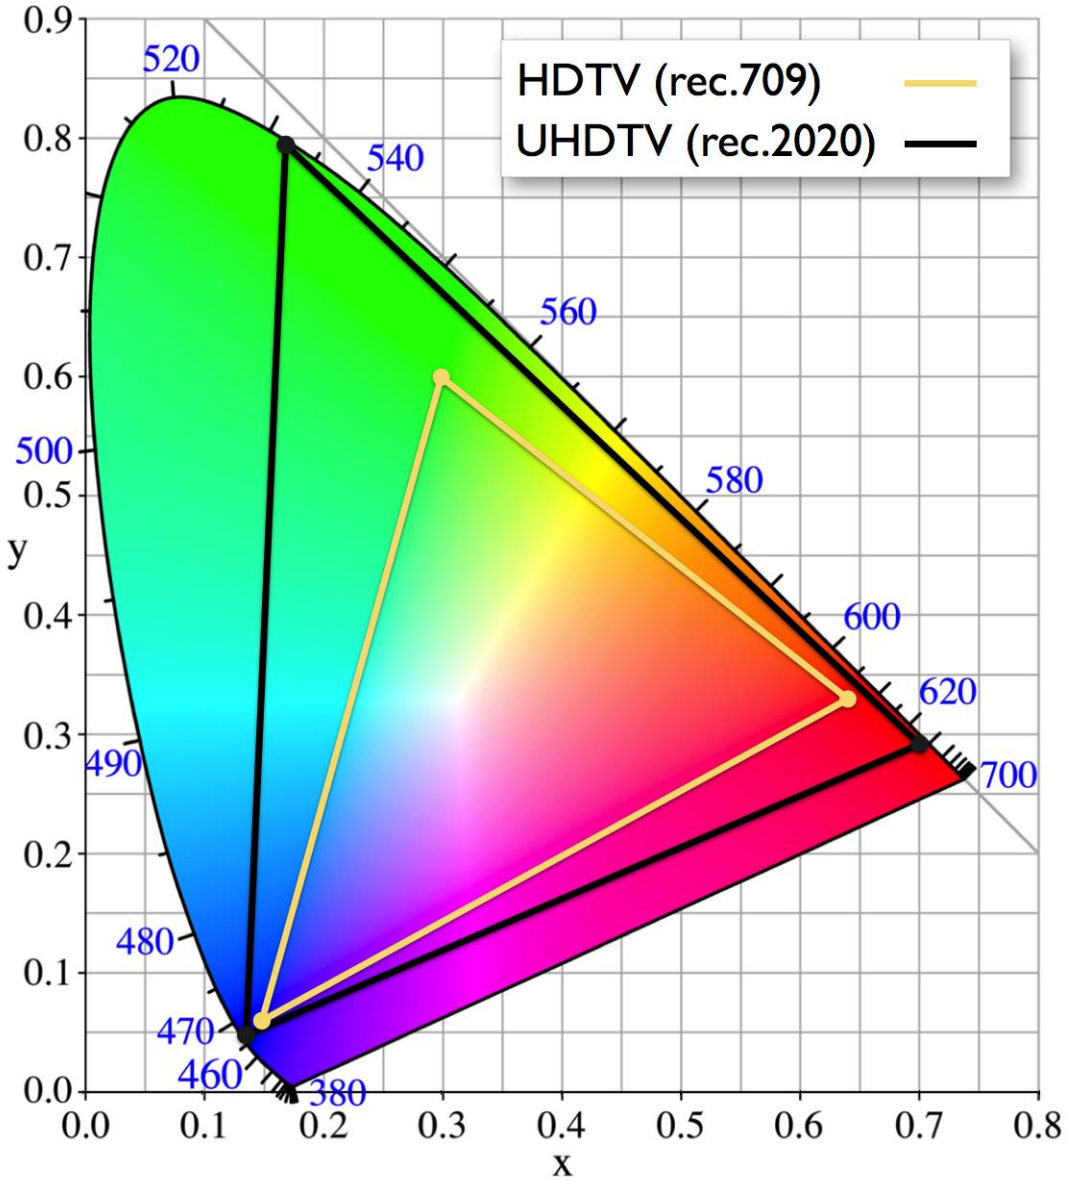

UHDTV vs HDTV Color Gamuts

Compression

- Lossy
  - Irreversible compression
  - Used to reduce data size
  - Examples: MP3, JPEG, MPEG-4
- Lossless
  - Reversable compression
  - Examples: ZIP, PNG, FLAC or Dolby TrueHD
- Visually Lossless
  - Irreversible compression
  - Difference can’t be seen by the Human Eye

YCbCr Encoding

- RGB uses **R**ed **G**reen **B**lue values to describe a color
- YCbCr is a different way of storing color:
  - Y = Luma or Brightness of the Color
  - Cr = Chroma difference for Red
  - Cb = Chroma difference for Blue

- Based on the idea that our visual system is **less** sensitive for **Color** than **Luminance**
- RGB can be converted LossLess to YCbCr
- YCbCr is also known as YUV

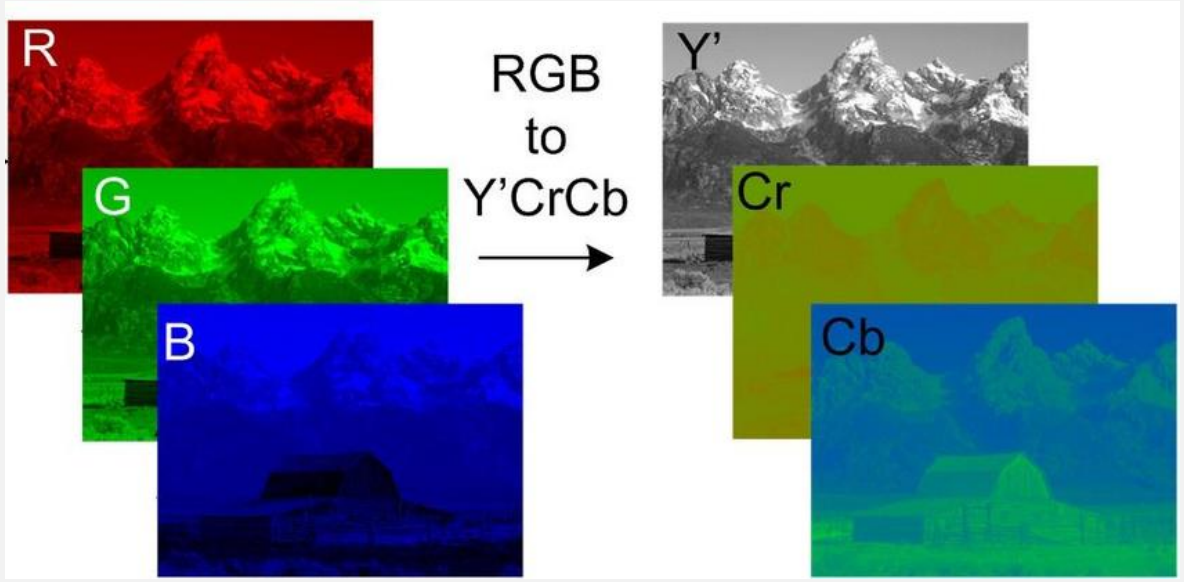

Chroma Subsampling

- Reduces the amount of Color information
- Allows picture clarity to be maintained
- Can reduce the file size up to 50%

Chroma Sub-sampling

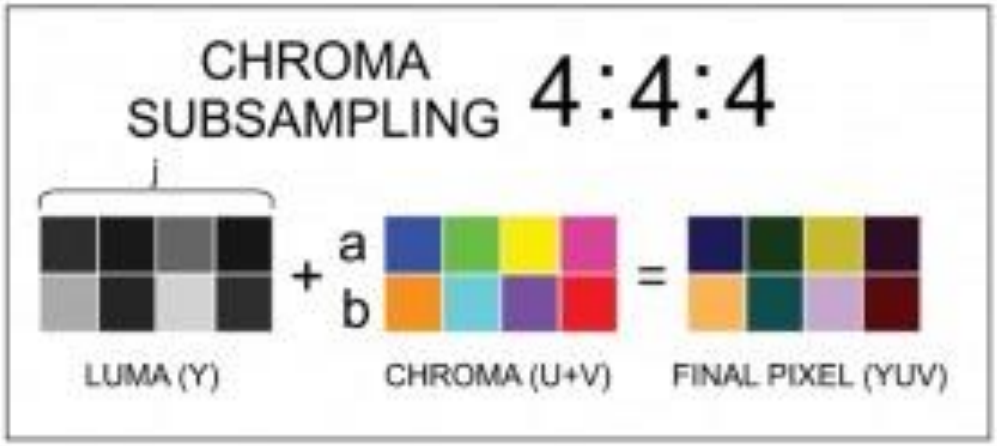

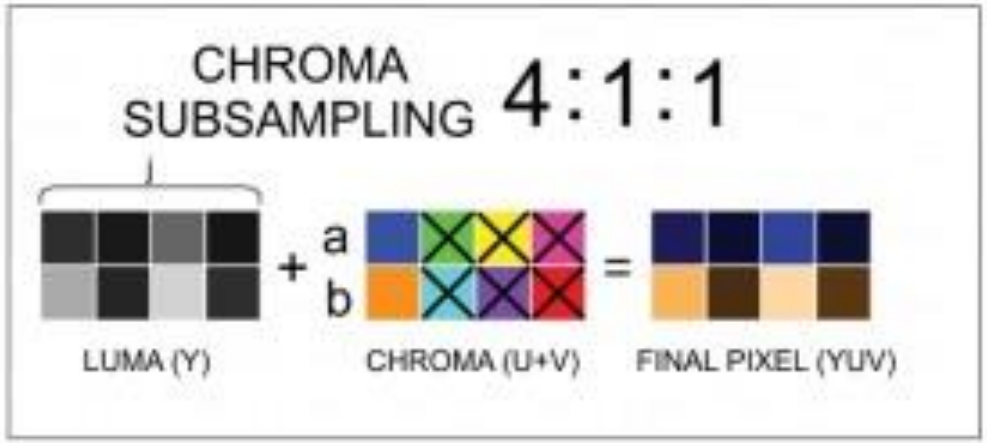

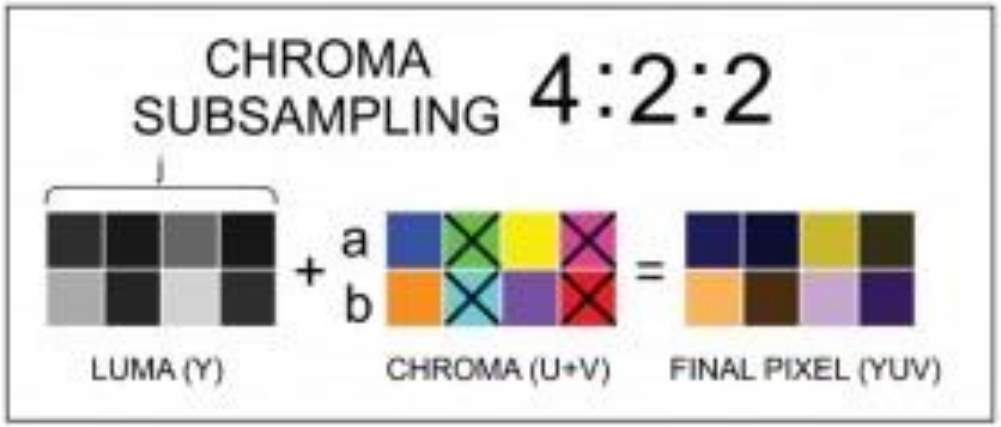

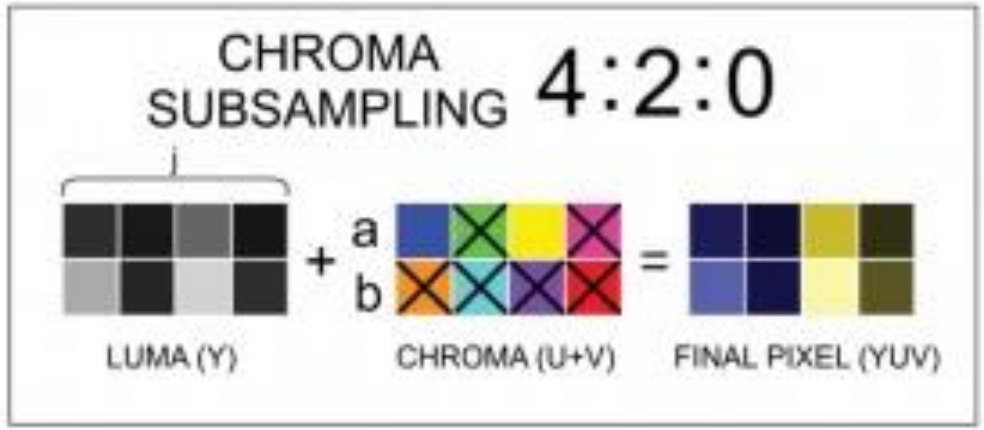

Chroma Sub-sampling – The downside

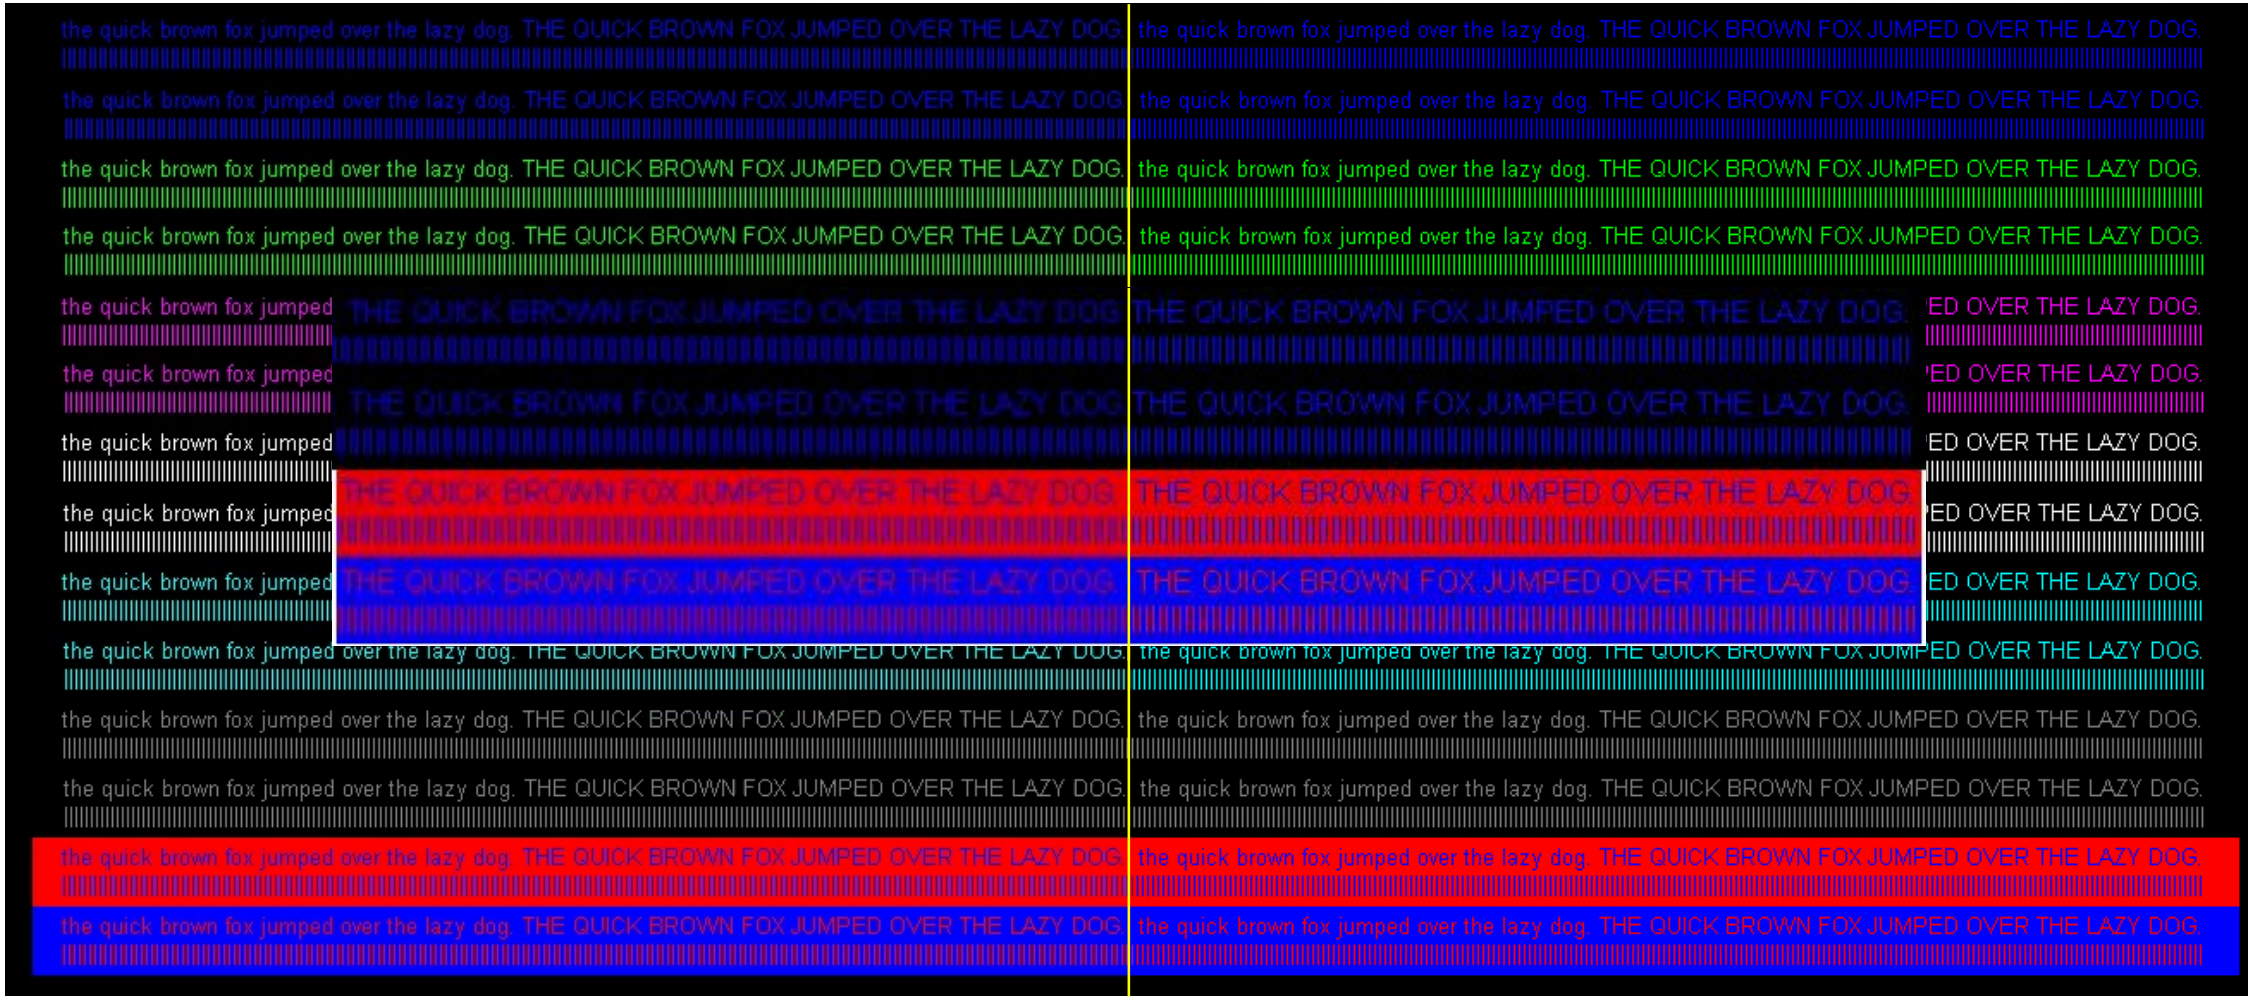

Chroma Sub-sampling – The downside

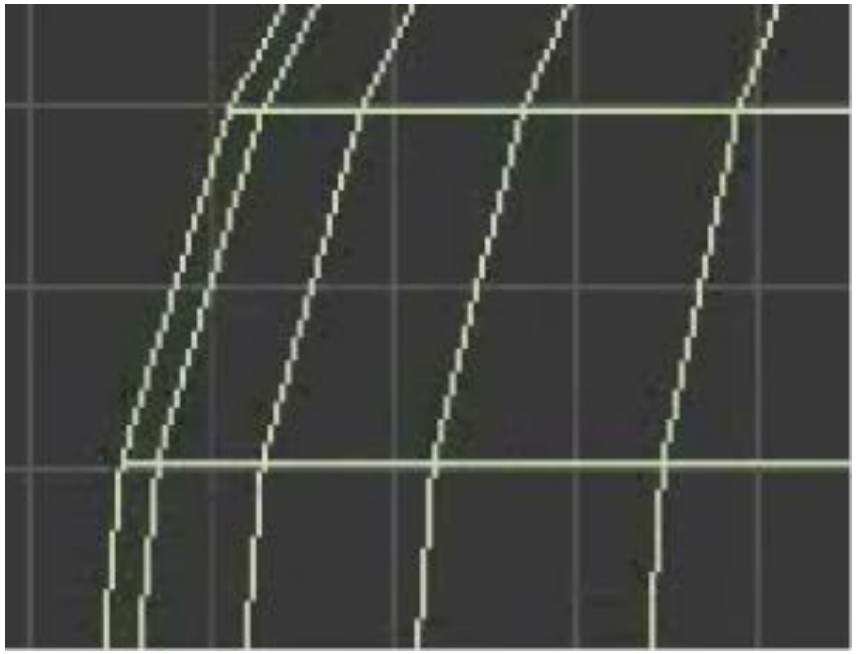

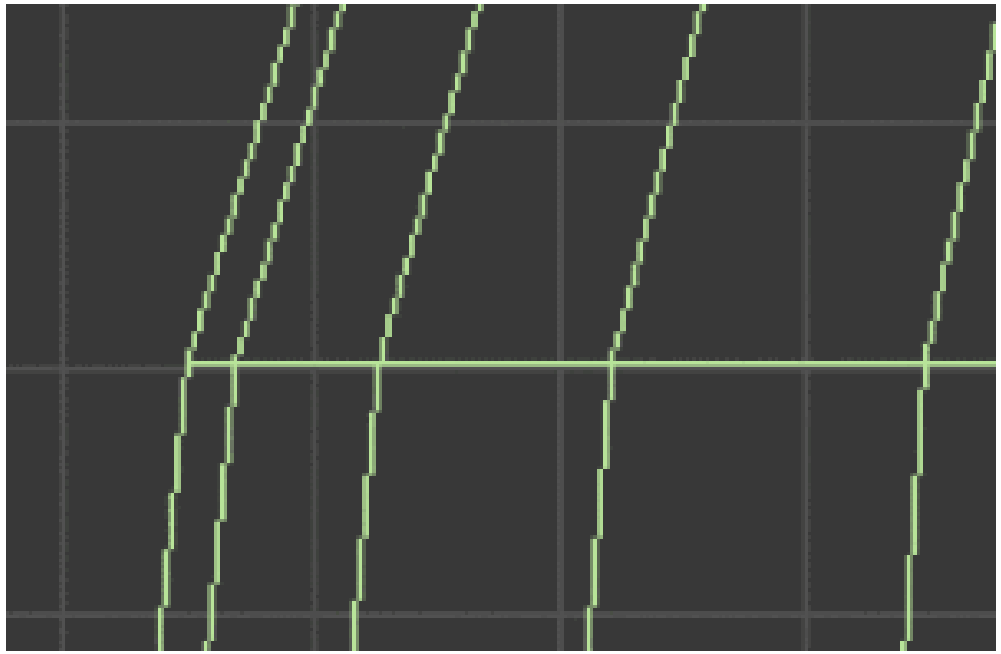

H264 – The magician

Used source: https://sidbala.com/h-264-is-magic/

Why a video codec

- It’s all about size
  - Pixels are built up by **R**ed **G**reen **B**lue values
  - 8 bits per Channel = 24 Bit Color (16,7 Million colors)
  - 10 bits per Channel = 30 Bit Color (1 Billion colors)
- Full HD:
  - 1920 * 1080 * 24 bit = **5.9** Mbyte per frame
  - 30 FPS = **177** MB/Sec
  - 60 FPS = **354** MB/Sec
- 4K:
  - 3840 * 2160 * 3 Bytes = **23.7** Mbyte per frame
  - 30 FPS = **711** MB/Sec
  - 60 FPS = **1422** MB/Sec

The history

Timeline of international video compression standards

Year	Standard	Publisher	Popular implementations
1984	H.120	ITU-T	
1988	H.261	ITU-T	Videoconferencing, videotelephony
1993	MPEG-1 Part 2	ISO, IEC	Video-CD
1995	H.262/MPEG-2 Part 2	ISO, IEC, ITU-T	DVD Video, Blu-ray, DVB, ATSC, SVCD
1996	H.263	ITU-T	Videoconferencing, videotelephony, video on mobile phones (3GP)
1999	MPEG-4 Part 2	ISO, IEC	Video on Internet (DivX, Xvid ...)
2003	H.264/MPEG-4 AVC	ISO, IEC, ITU-T	Blu-ray, HD DVD, DVB, ATSC, iPod Video, Apple TV, videoconferencing
2006	VC-1	SMPTE	Blu-ray, video on Internet
2013	H.265/MPEG-H HEVC	ISO, IEC, ITU-T	Ultra HD Blu-ray, DVB, ATSC 3.0, UHD streaming, High Efficiency Image Format, macOS High Sierra, iOS 11|-
2018	AV1	Alliance for Open Media	HTML5 video

H264 / AVC1

- Introduced as standard in 2003
- Used on most Blu-Ray videos
- Supports rec. 709 color space
- Hardware support by almost every device
    - TVs / Mobile Phones / Tablets
    - Raspberry Pi
    - Intel
    - NVIDIA

The H264 Magic

- Size comparison of capture of Tweakers.net homepage

- PNG: 886 Kb
- H264 **5 seconds:** ~400 Kb

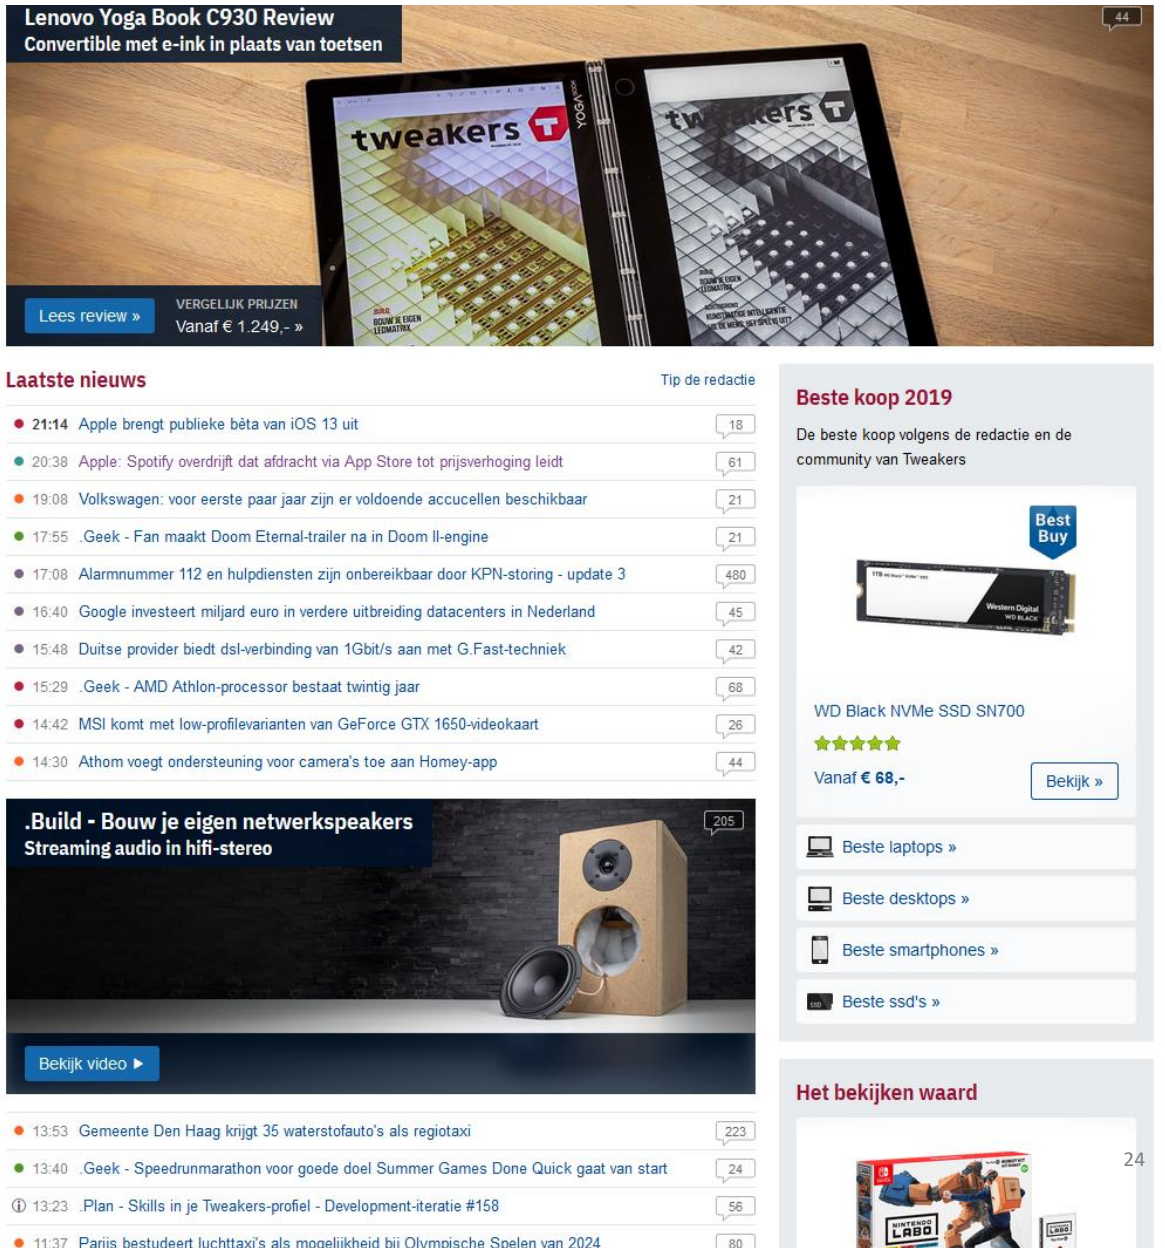

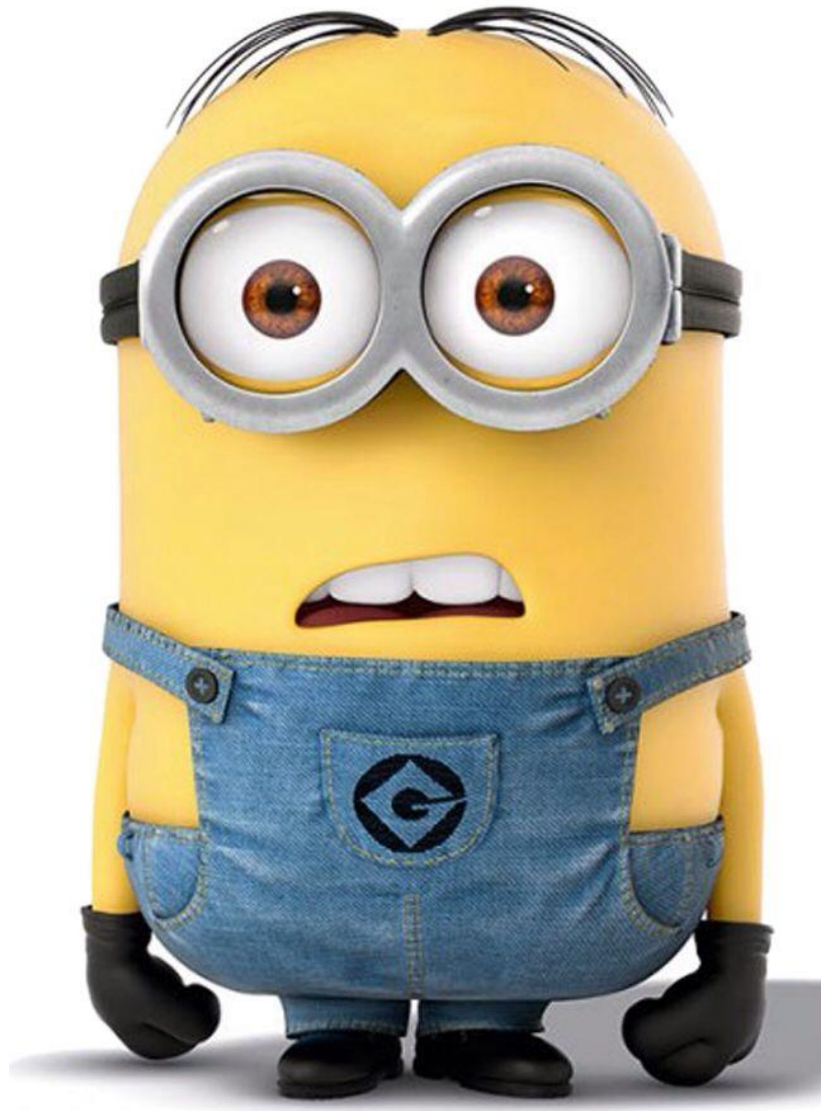

The tricks

- Entropy
  - > Information entropy is the number of bits required to represent some information
  - > Entropy is used for data deduplication
  - > For example if you toss a coin 12 times you can describe it as follows:
    - 1) HHHHHHTTTTTT
    - 2) 6xH & 6xT
- Chroma Subsampling
  - > YUV 4:2:0

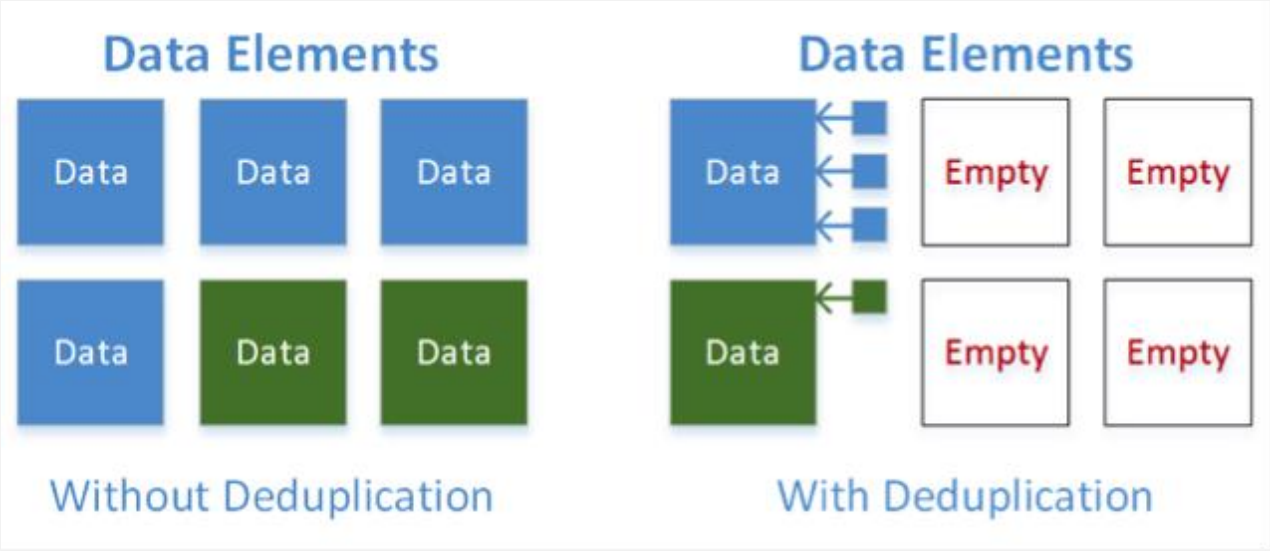

The tricks

- Frequency domain
  - Transforms an image to a representation of “level of detail”
  - High Frequency = High Detail = Border of representation
  - Low Frequency = Gentle Details = Center of representation
- Quantization
  - Filters out the high-detailed areas
  - Can be lossless if you leave the Frequency (level of detail) high enough
  - Massively reduces the image size

Original Image

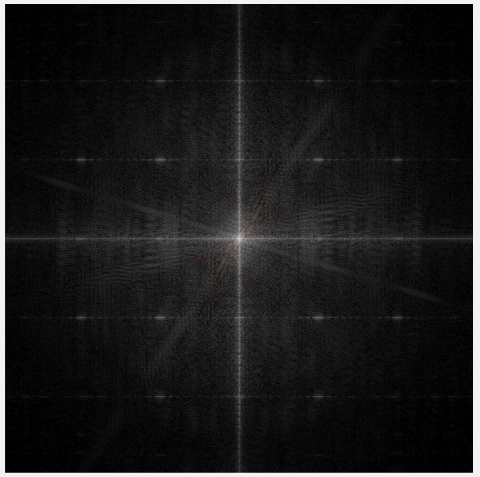

Frequency Domain Representation

PQR

Frequency Domain Masks

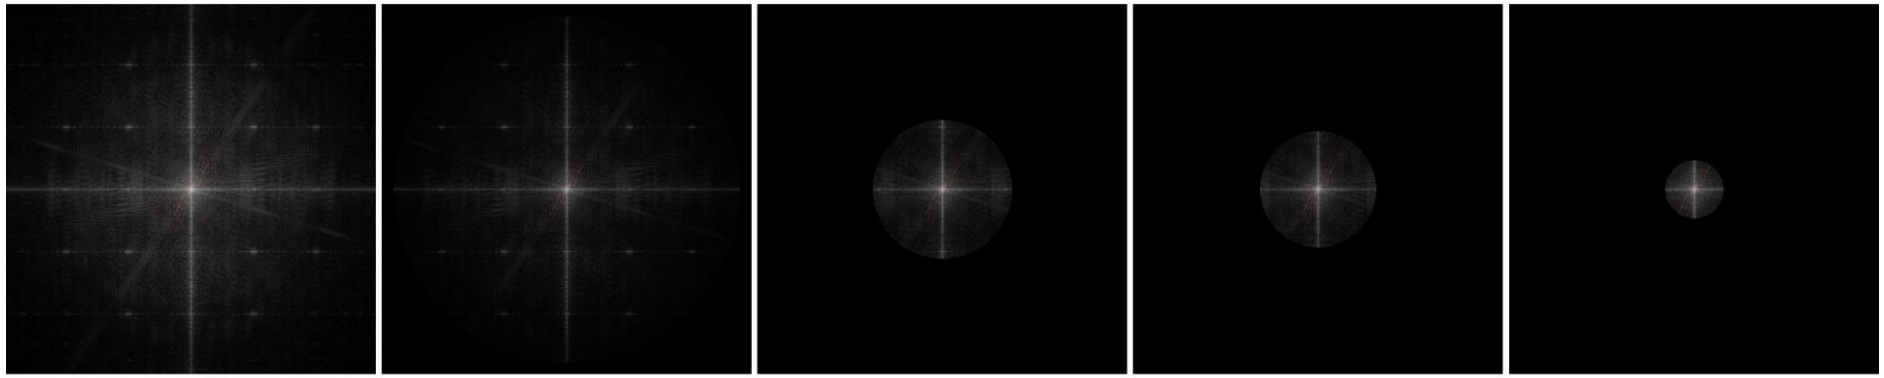

Spatial Domain Output Images

100% 69% 30% 11% 2%

The tricks

- Motion Compensation
  - > “Predicts” the frame by the previous or future frames
  - > Accounts motion of camera or objects
  - > Image is cut in slices (16x16 – 4x4)
  - > Consists of 3 types of frames:
  1. I-Frame = Complete image (Key Frames)
  2. P-Frame = Predicted, holds only changes from previous frame
  3. B-Frame= Bidirectional Predicted, difference between current frame and preceding and following frames

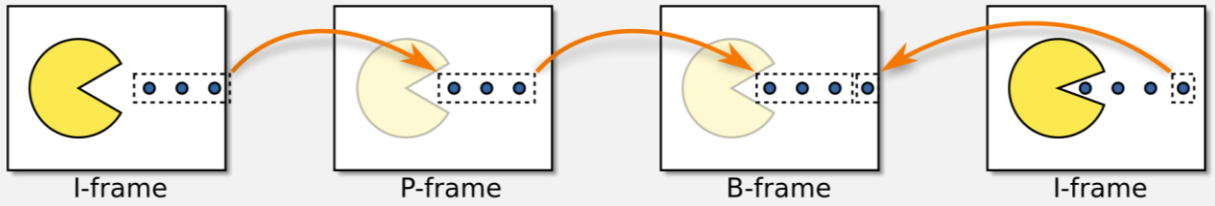

The result

- Resolution: 1920 * 1080 (FullHD)
- RAW: 890 Mb
- H264: 400 Kb
- Reduction: > 2200 times!

Other codecs

PQR

H265 / HEVC

Video coding standard	Average bit rate reduction compared with H.264/MPEG-4 AVC HP			
	480p	720p	1080p	2160p
HEVC	52%	56%	62%	64%

- Introduced as standard in 2013
- Used on Ultra-HD Blu-ray for 4K / HDR
- Up to 8K resolution
- Supports BT.2020 color space
- Build to **reduce** bit rate to facilitate higher resolutions at the **cost** of performance (uses 8x more CPU than H264!)
- Supported by many devices:
  - NVIDIA since 2015
  - Intel since Skylake / Kaby Lake
  - Apple since iPhone 6
  - Android Lollipop (5.0)

VP9

- Developed by Google for **YouTube**
- Released in 2013
- Open-Source & Royalty free
- Supported by most common browsers (FireFox / Chrome / Edge)
- Hardware support by most devices
  - NVIDIA since Pascal (Only Decoding)
  - Intel since Apollo Lake (decode) and Kaby Lake (encode / decode)

PQR

AV1

- Developed by Alliance for Open Media
- Released in 2018
- Open-Source & Royalty free
- 20%-40% higher data compression than VP9 & H265
- Hardware Encoding / Decoding is not commonly available yet
- 80 FPS Encode speed on 2x Intel Xeon 8280 @4 Ghz (56 Cores/112 Thread)
  - In comparison H265 = 365 FPS with same hardware

Hardware support

Hardware support for Codecs on NVIDIA cards

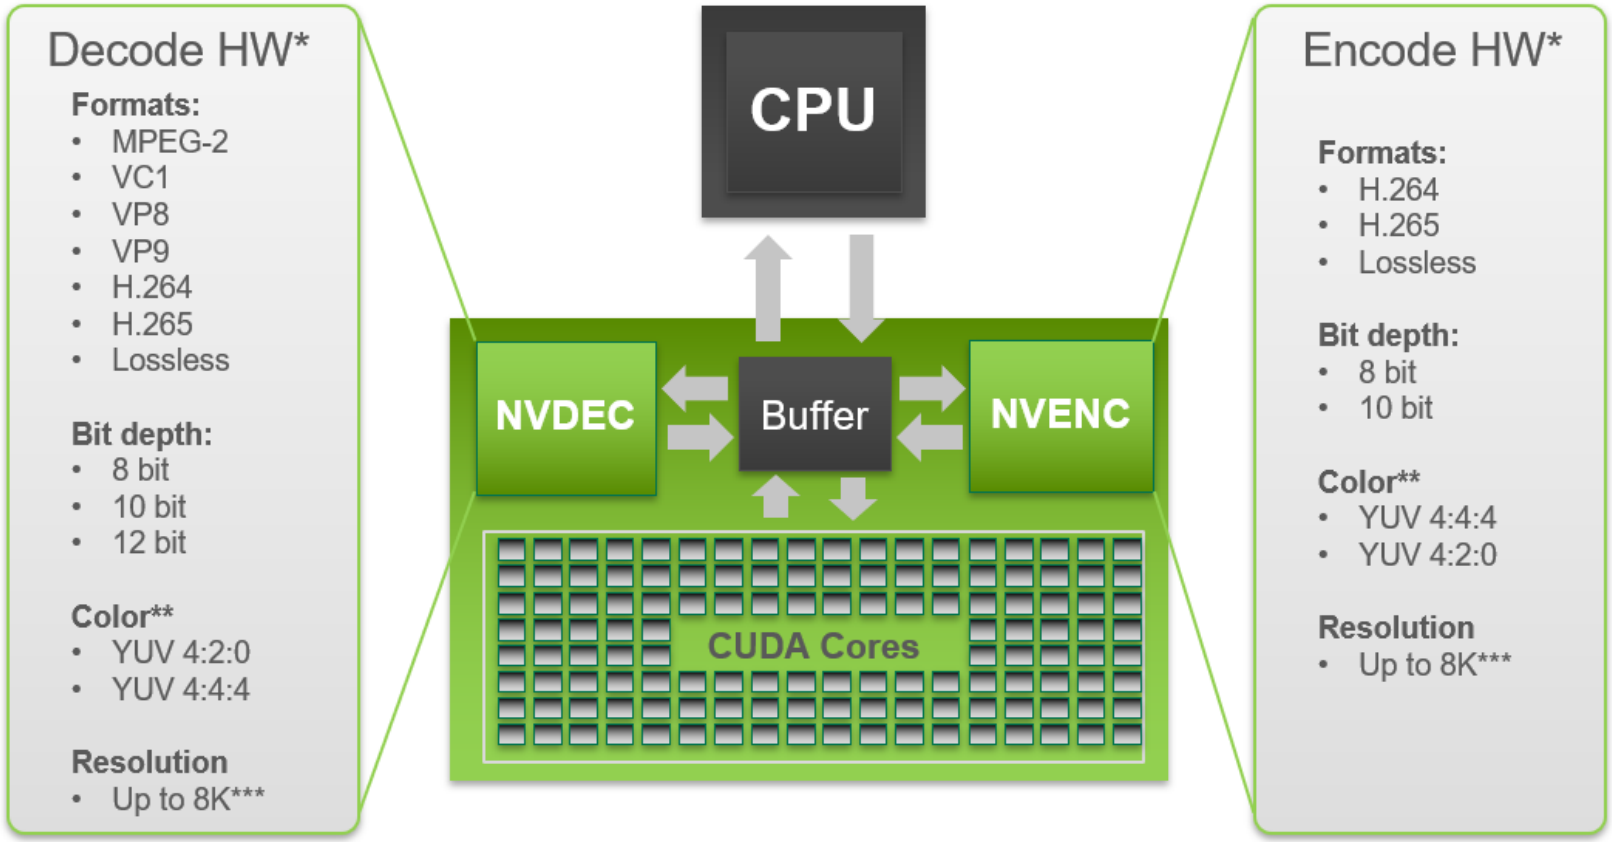

NVIDIA Hardware Encode

GPU	H.264 (AVCHD) YUV 4:2:0		H.264 (AVCHD) YUV 4:4:4		H.264 (AVCHD) LOSSLESS		H.265 (HEVC) YUV 4:2:0		H.265 (HEVC) YUV 4:4:4		H.265 (HEVC) LOSSLESS	
	MAX Color	MAX Res.	MAX Color	MAX Res.	MAX Color	MAX Res.	MAX Color	MAX Res.	MAX Color	MAX Res.	MAX Color	MAX Res.
Kepler	8-bit	4096 x 4096	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Maxwell (1st Gen)*	8-bit	4096 x 4096	8-bit	4096 x 4096	8-bit	4096 x 4096	N/A	N/A	N/A	N/A	N/A	N/A
Maxwell (2nd Gen)	8-bit	4096 x 4096	8-bit	4096 x 4096	8-bit	4096 x 4096	8-bit	4096 x 4096	N/A	N/A	N/A	N/A
Maxwell (GM206)	8-bit	4096 x 4096	8-bit	4096 x 4096	8-bit	4096 x 4096	8-bit	4096 x 4096	8-bit	4096 x 4096	8-bit	4096 x 4096
Pascal	8-bit	4096 x 4096	8-bit	4096 x 4096	8-bit	4096 x 4096	10-bit	8192 x 8192**	10-bit	8192 x 8192**	10-bit	8192 x 8192**
Volta	8-bit	4096 x 4096	8-bit	4096 x 4096	8-bit	4096 x 4096	10-bit	8192 x 8192	10-bit	8192 x 8192	10-bit	8192 x 8192
Turing	8-bit	4096 x 4096	8-bit	4096 x 4096	8-bit	4096 x 4096	10-bit	8192 x 8192	10-bit	8192 x 8192	10-bit	8192 x 8192

NVIDIA Hardware Encode

BOARD	FAMILY	CHIP	# OF CHIPS	# OF NVENC /CHIP	Total # of NVENC	Max # of concurrent sessions	H.264 (AVCHD) YUV 4:2:0	H.264 (AVCHD) YUV 4:4:4	H.264 (AVCHD) Lossless	H.265 (HEVC) 4K YUV 4:2:0	H.265 (HEVC) 4K YUV 4:4:4	H.265 (HEVC) 4K Lossless	H.265 (HEVC) 8k	HEVC B Frame support
GRID														
GRID K1	Kepler	GK107	4	1	4	Unrestricted	YES	NO	NO	NO	NO	NO	NO	NO
GRID K2	Kepler	GK104	2	1	2	Unrestricted	YES	NO	NO	NO	NO	NO	NO	NO
GRID K340	Kepler	GK107	4	1	4	Unrestricted	YES	NO	NO	NO	NO	NO	NO	NO
GRID K520	Kepler	GK104	2	1	2	Unrestricted	YES	NO	NO	NO	NO	NO	NO	NO
TESLA														
Tesla K10	Kepler	GK104	2	1	2	Unrestricted	YES	NO	NO	NO	NO	NO	NO	NO
Tesla K20X	Kepler	GK110	1	1	1	Unrestricted	YES	NO	NO	NO	NO	NO	NO	NO
Tesla K40	Kepler	GK110B	1	1	1	Unrestricted	YES	NO	NO	NO	NO	NO	NO	NO
Tesla K80	Kepler (2nd Gen)	GK210	2	1	2	Unrestricted	YES	NO	NO	NO	NO	NO	NO	NO
Tesla M10	Maxwell (1st Gen)	GM107	4	1	4	Unrestricted	YES	YES	YES	NO	NO	NO	NO	NO
Tesla M4	Maxwell (GM206)	GM206	1	1	1	Unrestricted	YES	YES	YES	YES	NO	NO	NO	NO
Tesla M40	Maxwell (2nd Gen)	GM200	1	2	2	Unrestricted	YES	YES	YES	YES	NO	NO	NO	NO
Tesla M6	Maxwell (2nd Gen)	GM204	1	2	2	Unrestricted	YES	YES	YES	YES	NO	NO	NO	NO
Tesla M60	Maxwell (2nd Gen)	GM204	2	2	4	Unrestricted	YES	YES	YES	YES	NO	NO	NO	NO
Tesla P4	Pascal	GP104	1	2	2	Unrestricted	YES	YES	YES	YES	YES	YES	YES	NO
Tesla P6	Pascal	GP104	1	2	2	Unrestricted	YES	YES	YES	YES	YES	YES	YES	NO
Tesla P40	Pascal	GP102	1	2	2	Unrestricted	YES	YES	YES	YES	YES	YES	YES	NO
Tesla P100	Pascal	GP100	1	3	3	Unrestricted	YES	YES	YES	YES	YES	YES	NO	NO
Tesla V100	Volta	GV100	1	3	3	Unrestricted	YES	YES	YES	YES	YES	YES	YES	NO
Tesla T4	Turing	TU104	1	1	1	Unrestricted	YES	YES	YES	YES	YES	YES	YES	YES

NVIDIA Performance Enhancements

ENCODE PERFORMANCE

H.264 1080p (1920x1080) 4:2:0 8bit 30fps (SINGLE NVENC)

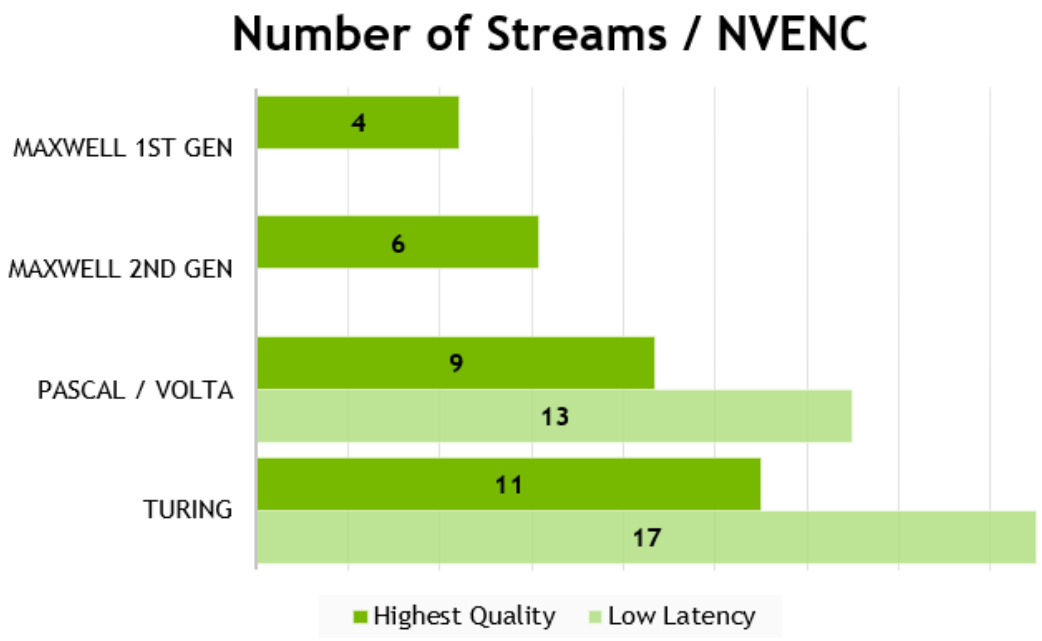

#NVENC	GPUs
x1	Maxwell Quadro K2200 (1st Gen)/M2000 (2nd Gen) Maxwell (2nd Gen) Tesla M4 Pascal Quadro P2000/P4000 Turing Quadro & Tesla
x2	Maxwell (2nd Gen) Quadro M4000/M5000/M6000 Maxwell (2nd Gen) Tesla M6/M40 Pascal Quadro P5000/P6000 Pascal Tesla P4/P40
x3	Pascal Quadro GP100 Pascal Tesla P100 Volta Quadro GV100 Volta Tesla V100
x4	Maxwell (2nd Gen) Tesla M60

Note: All GPUs not featured above are limited to 2 simultaneous sessions

NVIDIA Hardware Decode

BOARD	FAMILY	CHIP	# OF CHIPS	# OF NVDEC /CHIP	Total # of NDEC	MPEG-1	MPEG-2	VC-1	VP8	VP9			H.264 (AVCHD)	H.265 (HEVC) 4:2:0			*H.265 (HEVC) 4:4:4		
										8 bit	10 bit	12 bit		8 bit	10 bit	12 bit	8 bit	10 bit	12 bit
GRID																			
GRID K1	Kepler	GK107	4	1	4	YES	YES	YES	NO	NO	NO	NO	YES	NO	NO	NO	NO	NO	NO
GRID K2	Kepler	GK104	2	1	2	YES	YES	YES	NO	NO	NO	NO	YES	NO	NO	NO	NO	NO	NO
GRID K340	Kepler	GK107	4	1	4	YES	YES	YES	NO	NO	NO	NO	YES	NO	NO	NO	NO	NO	NO
GRID K520	Kepler	GK104	2	1	2	YES	YES	YES	NO	NO	NO	NO	YES	NO	NO	NO	NO	NO	NO
TESLA																			
Tesla K10	Kepler	GK104	2	1	2	YES	YES	YES	NO	NO	NO	NO	YES	NO	NO	NO	NO	NO	NO
Tesla K20X	Kepler	GK110	1	1	1	YES	YES	YES	NO	NO	NO	NO	YES	NO	NO	NO	NO	NO	NO
Tesla K40	Kepler	GK110B	1	1	1	YES	YES	YES	NO	NO	NO	NO	YES	NO	NO	NO	NO	NO	NO
Tesla K80	Kepler (2nd Gen)	GK210	2	1	2	YES	YES	YES	NO	NO	NO	NO	YES	NO	NO	NO	NO	NO	NO
Tesla M10	Maxwell (1st Gen)	GM107	4	1	4	YES	YES	YES	YES	NO	NO	NO	YES	NO	NO	NO	NO	NO	NO
Tesla M4	Maxwell (GM206)	GM206	1	1	1	YES	YES	YES	YES	YES	NO	NO	YES	YES	YES	NO	NO	NO	NO
Tesla M6	Maxwell (2nd Gen)	GM204	1	1	1	YES	YES	YES	YES	NO	NO	NO	YES	NO	NO	NO	NO	NO	NO
Tesla M60	Maxwell (2nd Gen)	GM204	2	1	2	YES	YES	YES	YES	NO	NO	NO	YES	NO	NO	NO	NO	NO	NO
Tesla M40	Maxwell (2nd Gen)	GM200	1	1	1	YES	YES	YES	YES	NO	NO	NO	YES	NO	NO	NO	NO	NO	NO
Tesla P4 / P6	Pascal	GP104	1	1	1	YES	YES	YES	YES	YES	NO	NO	YES	YES	YES	YES	NO	NO	NO
Tesla P40	Pascal	GP102	1	1	1	YES	YES	YES	NO	YES	YES	YES	YES	YES	YES	YES	NO	NO	NO
Tesla P100	Pascal	GP100	1	1	1	YES	YES	YES	YES	YES	NO	NO	YES	YES	YES	YES	NO	NO	NO
Tesla V100	Volta	GV100	1	1	1	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	NO	NO	NO
Tesla T4	Turing	TU104	1	2	2	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES	YES