AV1 and Arm

Vibhoothi
Luca Barbato

Who are we?

Vibhoothi

- **rav1e** contributor
- Member and mentor at amFOSS.
- Research Assistant, at Trinity College, Dublin.
- Contact:
  - mindfreeze@xiph.org
  - @vibhoothi on Twitter.

Who are we?

Luca Barbato

- **rav1e** and **dav1d** contributor among many other open source software.
- Contact
  - lu_zero@gentoo.org

Intro

- We will talk a little about AV1
- We will talk a little about Arm
- We will talk a little about dav1d
- We will talk a lot about rav1e, libaom and SVT-AV1
- We will discuss on what's missing to make AV1 shine on Arm.

What is AV1?

- AV1 is a next-generation video codec standard published by AOMedia
  - Offers best compression out of all available formats
    - Designed to be used in a broad set of use cases
  - Royalty-free and Open-source

What is AV1?

- AV1 is a next-generation video codec standard published by AOMedia
  - Offers best compression out of all available formats
    - Designed to be used in a broad set of use cases
  - Royalty-free and Open-source
  - Support from many largest technology companies in the world

What is AV1?

Tencent · SAMSUNG · moz://a · facebook · Chips&Media™ · BBC R&D · Adobe · XILINX®

NVIDIA · NETFLIX · intel · Apple · BROADCOM · ALLEGRO Digital Video Technology · hulu · Vidyo

Microsoft · Google · arm · CableLabs · ateme · CoSMo Software Consulting Pte Ltd · iCue

cisco · amazon · 金山云 · Amlogic · agora.io · gfycat · Visionular · videolan.org

BITMOVIN Solving Complex Video Problems · ARGON DESIGN · AMD · Ittiam · NGCODEC Next Generation Video · iQIYI 爱奇艺 悦享品质 · VeriSilicon · socionext

Alibaba Group 阿里巴巴集团 · Polycom · V-SILICON · REALTEK · vimeo

What is AV1?

- AV1 is a next-generation video codec standard published by AOMedia
    - Offers best compression out of all available formats
        - Designed to be used in a broad set of use cases
    - Royalty-free and Open-source
    - Support from many largest technology companies in the world
        - All of whom have pledged their patents toward AV1
    - It is already fairly ubiquitous, at least if we think about **decoding**.

Arm

- Arm is an architecture famous for its ubiquitous deployment in the low-power, mobile devices
- There is a bright future for Arm on laptops and in the cloud.

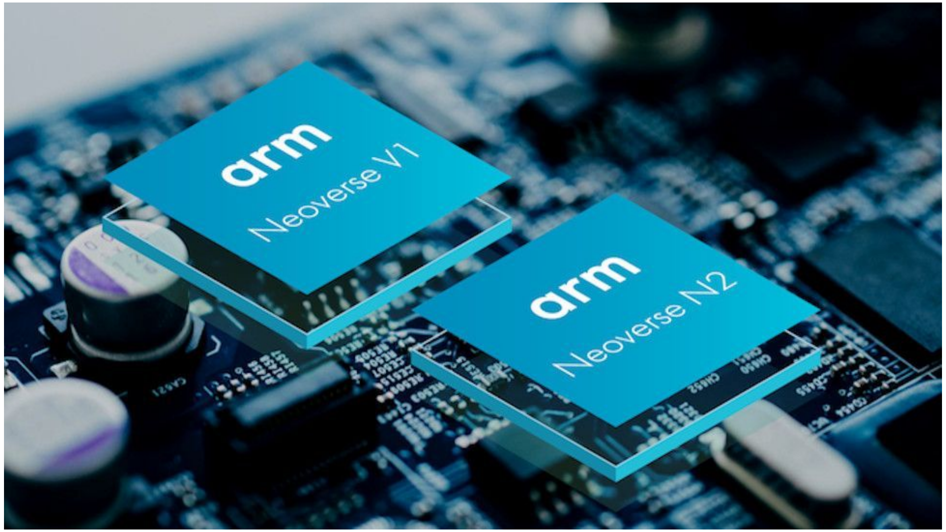

Source: Google Images

Arm: The Industry's Architecture of Choice

Extraordinary growth – from sensors to server

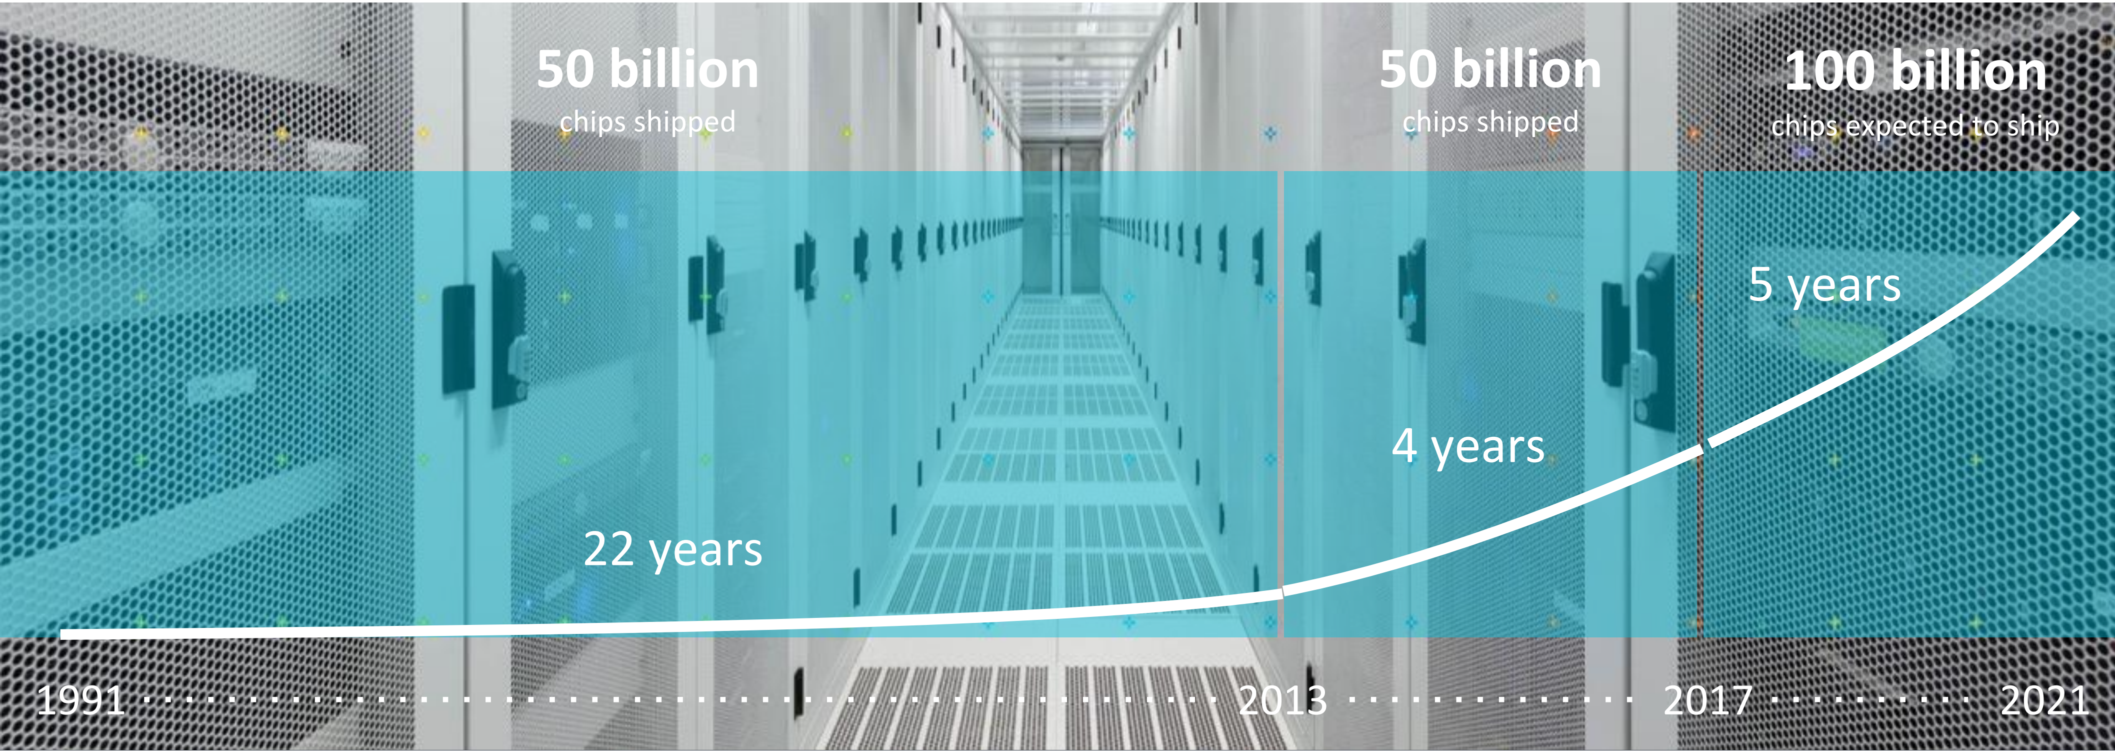

The Arm Ecosystem: 1000+ partners

Arm Neoverse Powered AWS Compute Instances

Announced at AWS Re:Invent 2019 - GA in May 2020

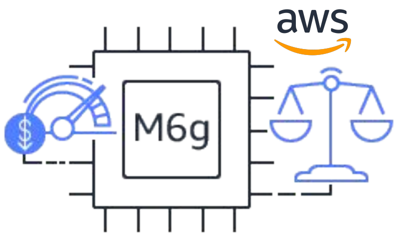

⬆ 40%

Price/Performance compared to x86 M5 Instances

⬇ 20%

Lower cost compared to x86 M5 Instances

64-bit Neoverse Cores with Custom 7nm Silicon

M6g Instances Currently Deployed At

NETFLIX

honeycomb.io

Mobiuspace

arm TREASURE DATA

redbox.

RAYGUN

nielsen

hotelbeds

…and more…

Enabling a Frictionless Cloud-Native Developer Experience

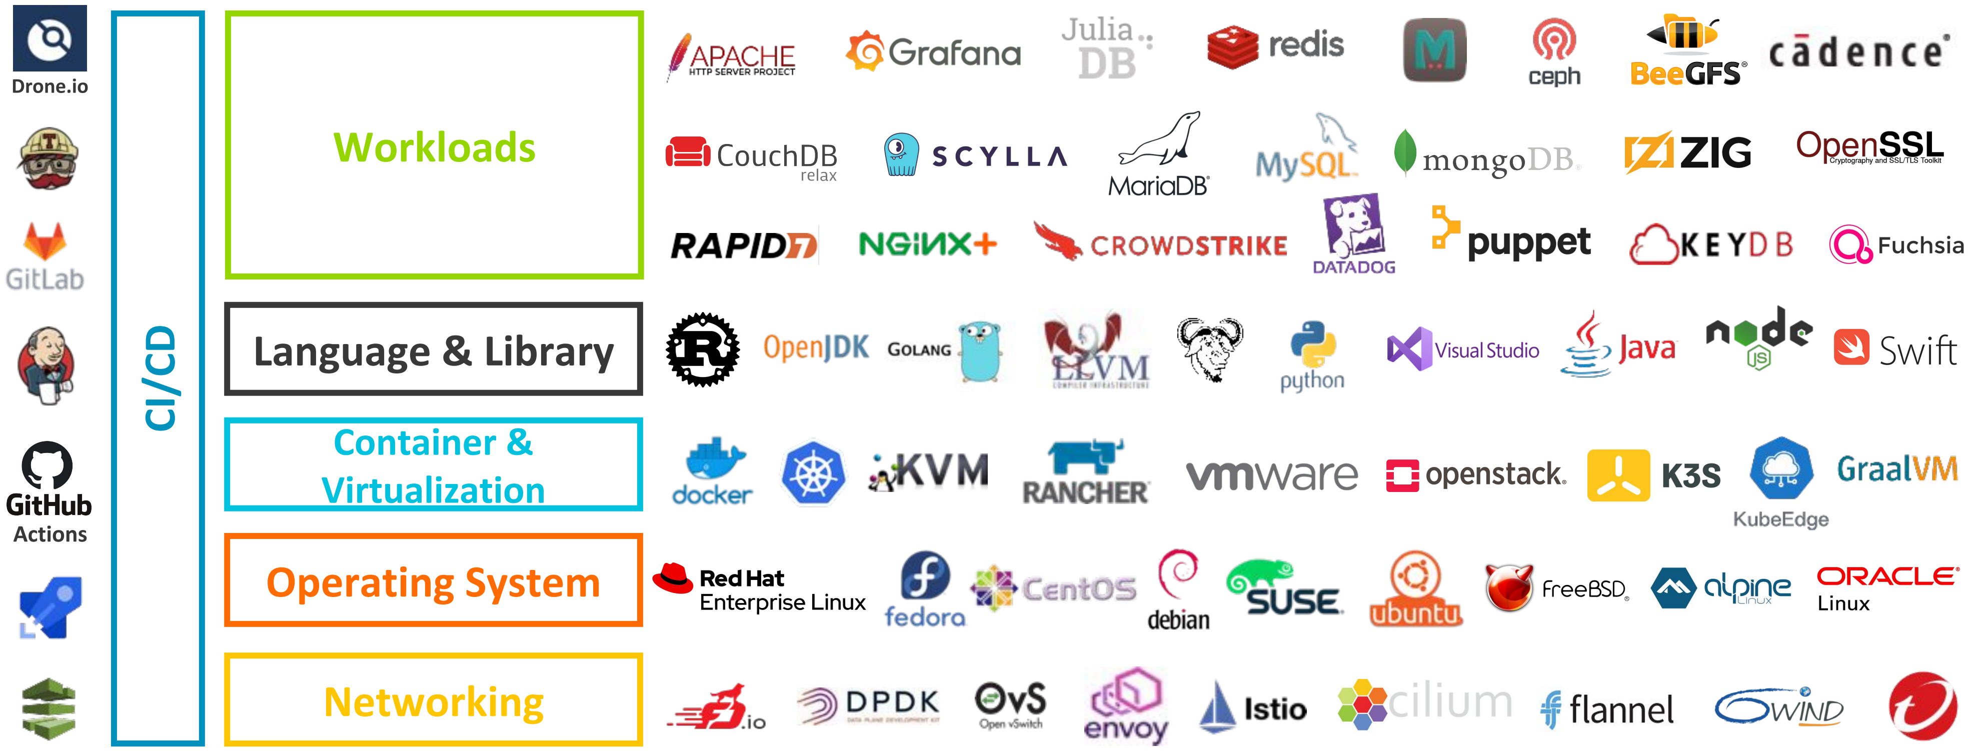

CI/CD

Workloads

Language & Library

Container & Virtualization

Operating System

Networking

Arm and AV1 Hardware Producers

BROADCOM

MEDIATEK

REALTEK

Amlogic

Rockchip

and lot more...

Arm and AV1 Hardware Producers

BROADCOM®

MEDIATEK

Roku®

mi

oppo

REALTEK

DUNE HD

HUAWEI

amlogic

realme

ZTE

vivo

LG
Life's Good

SAMSUNG

Rockchip

TCL

motorola
A Lenovo Company

and lot more...

Decoding AV1?

AV1 decoding capability is nearly ubiquitous already

- AV1 decoding capability is built-in Android Q
- Microsoft has a media extension for it.
- VLC supports AV1 Linux & macOS (using dav1d)
- Fast AV1 decoding in most browsers
  - Firefox 69
  - Chrome 77
  - Edge
  - Safari 12.1.12
    - WASM-compiled dav1d

* Using experimental ffmpeg patch to enable SVT-AV1 options.

Building the Future of Visual Experiences — Together!

* ogv.js WebAssembly AV1 shim is experimental and limited in resolution.

Decoding AV1?

AV1 decoding capability is nearly ubiquitous already

- AV1 Decoding support in various CPUs and GPUs:
  - NVIDIA GPUs
  - AMD GPUs
  - Intel Tiger Lake CPUs
- Google Duo adds support for AV1 decoding
- Netflix Streams AV1 on Android Mobile devices using dav1d.

AV1 STREAMING LIVE TO ALL BROWSERS
SW ENCODING ON 2ND GENERATION INTEL® XEON® SCALABLE PROCESSORS
ffmpeg* + SVT-AV1
56C/112T server
AV1 "live" speed
1080p30
720p30
360p30
DASH segment
DASH segment
DASH segment
BROWSER PLAYBACK
Firefox (native)
Chrome (native)
Edge Insider (native)
Safari (WebAssembly*)
moz://a
WIKIMEDIA FOUNDATION
* Using experimental ffmpeg patch to enable SVT-AV1 options
* ogv.js WebAssembly AV1 shim is experimental and limited in resolution.
Building the Future of Visual Experiences — Together!
intel

dav1d

- Announced in 2018 in VideoLAN Developer Days 2018 and Demuxed 2018 with 0.1.0, now at **0.7.1** (0.8.0 to be out soon).
- Used as default AV1 Decoder for almost all places => Production-ready.
- dav1d Goals
  - “AV1 needs a good decoder”
  - Fast decoder everywhere
  - Very cross-platform
  - Small binary size

dav1d

- Faster on Arm
- dav1d's
  - C Version is fast
  - Threading is better
  - Low-level with hand-written asm and no intrinsics
- Supports 10 bit Assembly for Arm
  - yet to be available for x86

Software decoding AV1 is already VERY fast

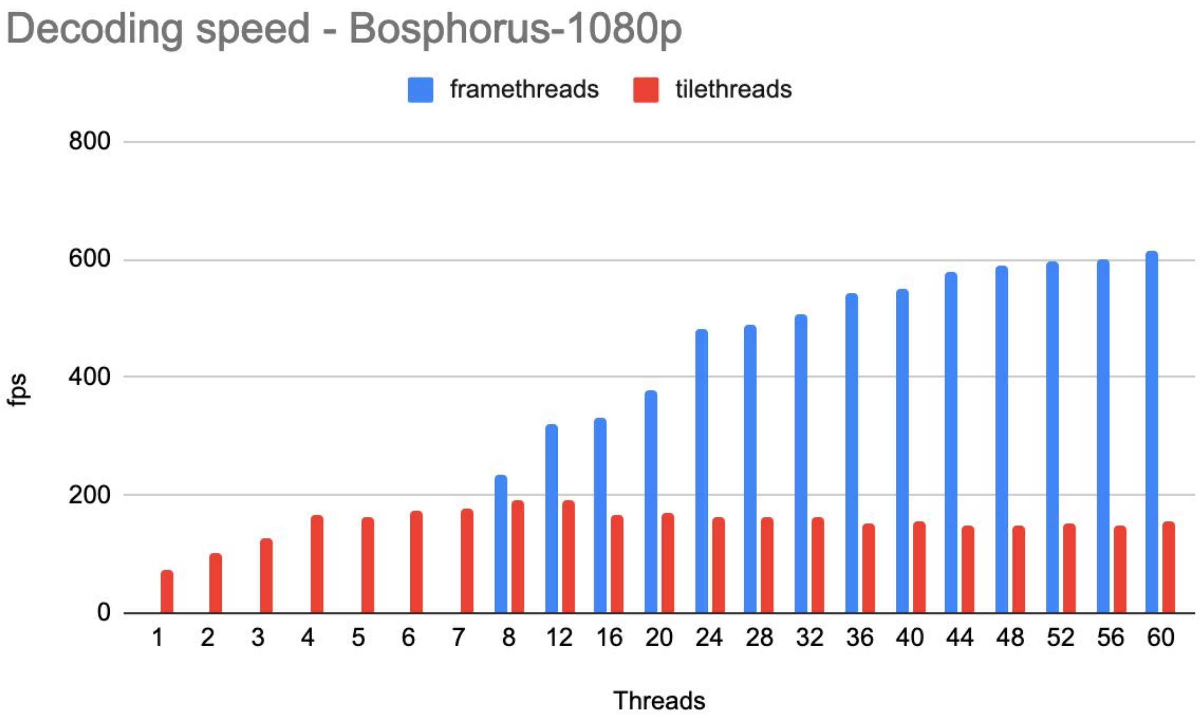

```
frametherads: dav1d --framethreads $thread -i ~/Encodes/aom-bosphorus-cpu-6.ivf -o /dev/null
tilethreads: dav1d --tilethreads $thread -i ~/Encodes/aom-bosphorus-cpu-6.ivf -o /dev/null
```

Software decoding AV1 is already VERY fast

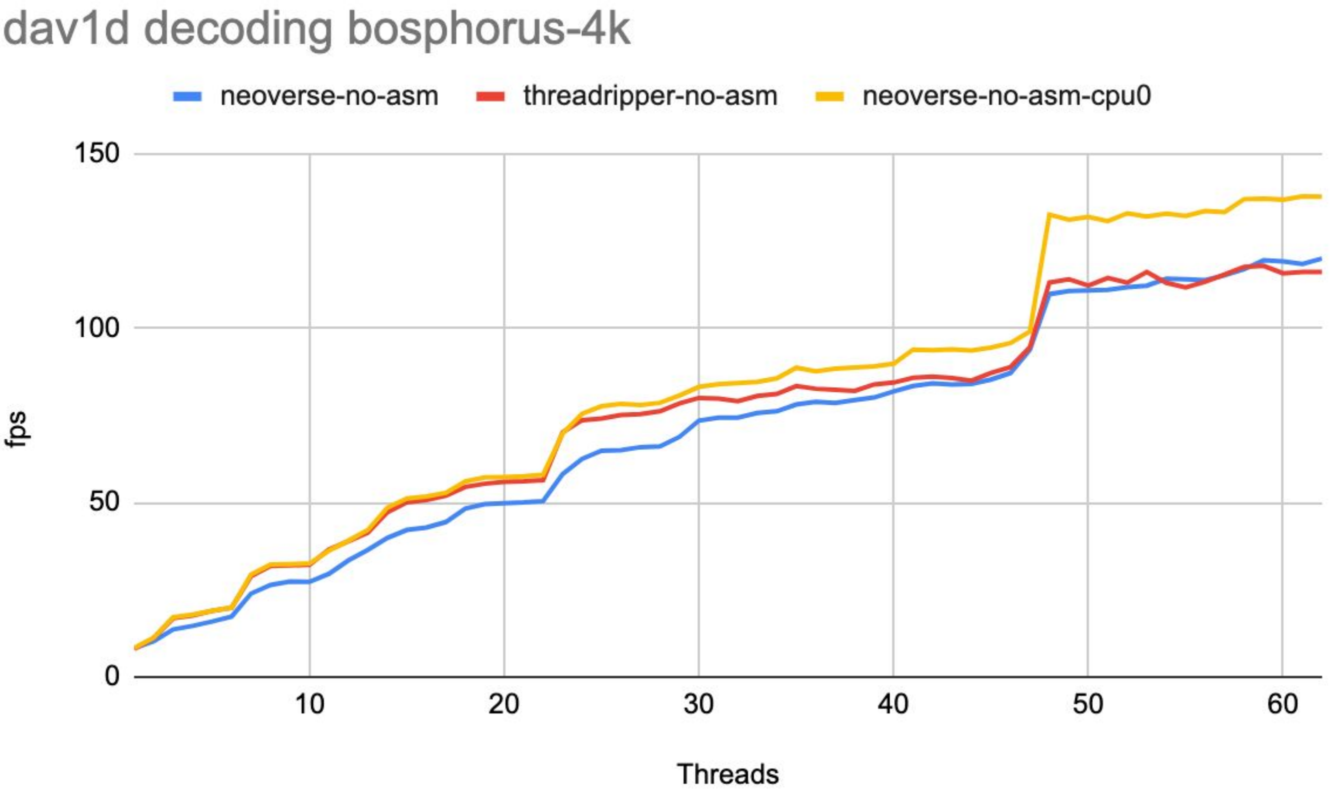

```
neoverse-no-asm: ./dav1d --cpumask 0 --framethreads $thread -i ~/rav1e-bosphorus-4k-32tiles.ivf -o /dev/null
threadripper-no-asm: ./dav1d --cpumask 0 --framethreads $thread -i ~/Encodes/rav1e-bosphorus-4k-32tiles.ivf -o /dev/null
neoverse-no-asm-cpu0: numactl --cpubind=0 --membind=0 dav1d --cpumask 0 --framethreads $thread -i ~/Encodes/rav1e-bosphorus-4k-32tiles.ivf
-o /dev/null
```

Software decoding AV1 is already VERY fast

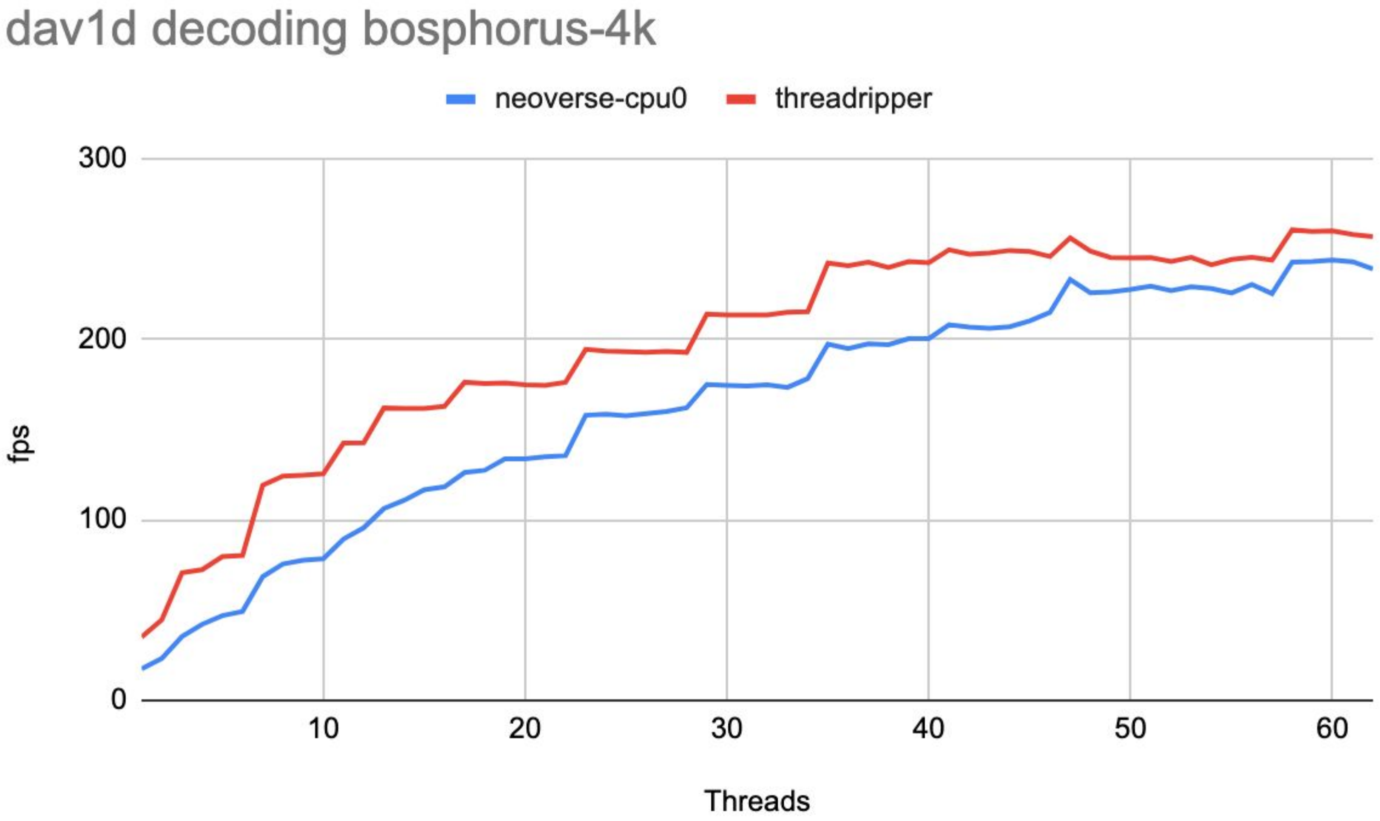

```
neoverse-cpu0: numactl --cpubind=0 --membind=0 dav1d --framethreads $thread -i ~/Encodes/rav1e-bosphorus-4k-32tiles.ivf -o /dev/null
threadripper: ./dav1d --framethreads $thread -i ~/rav1e-bosphorus-4k-32tiles.ivf -o /dev/null
```

Software decoding AV1 is already VERY fast

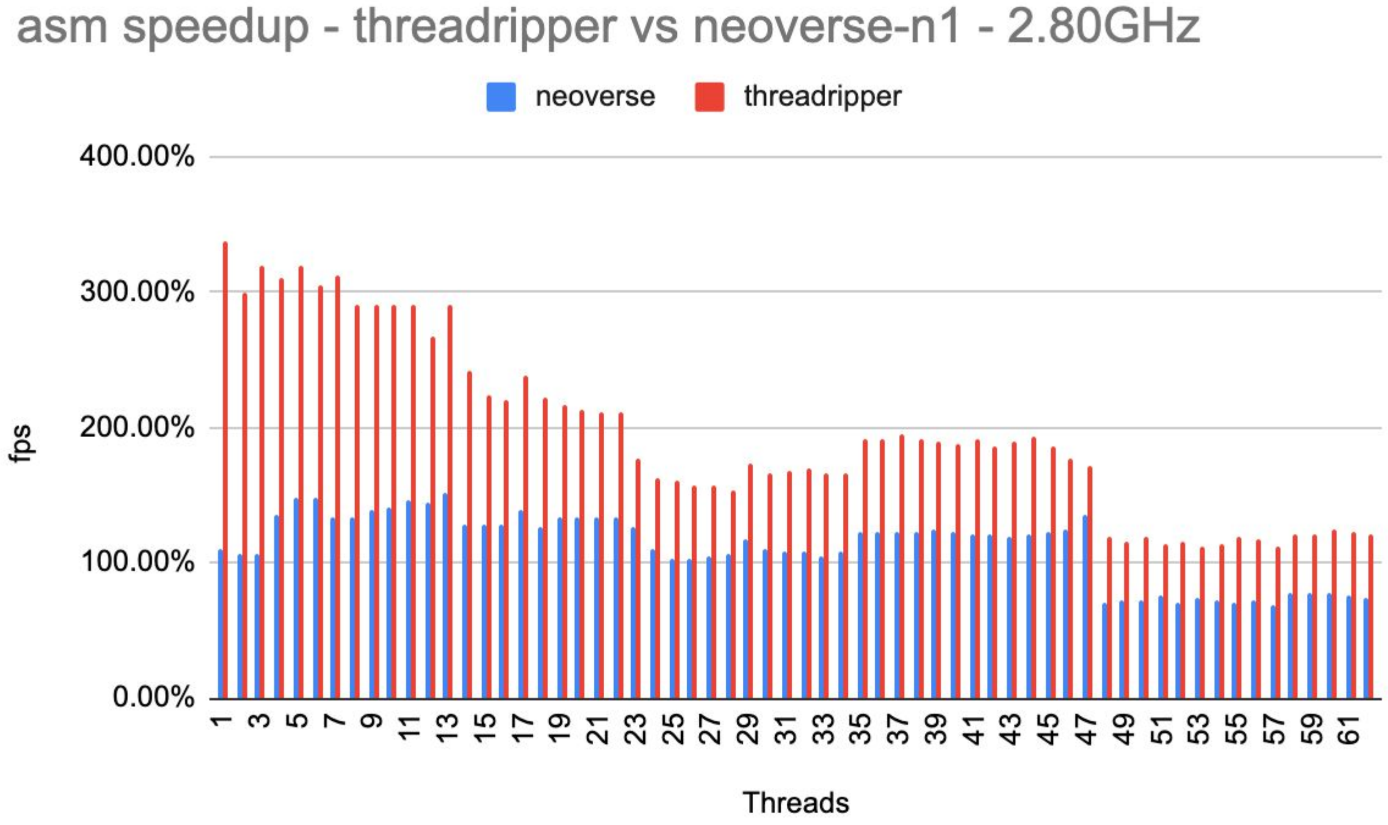

Software decoding AV1 is already VERY fast

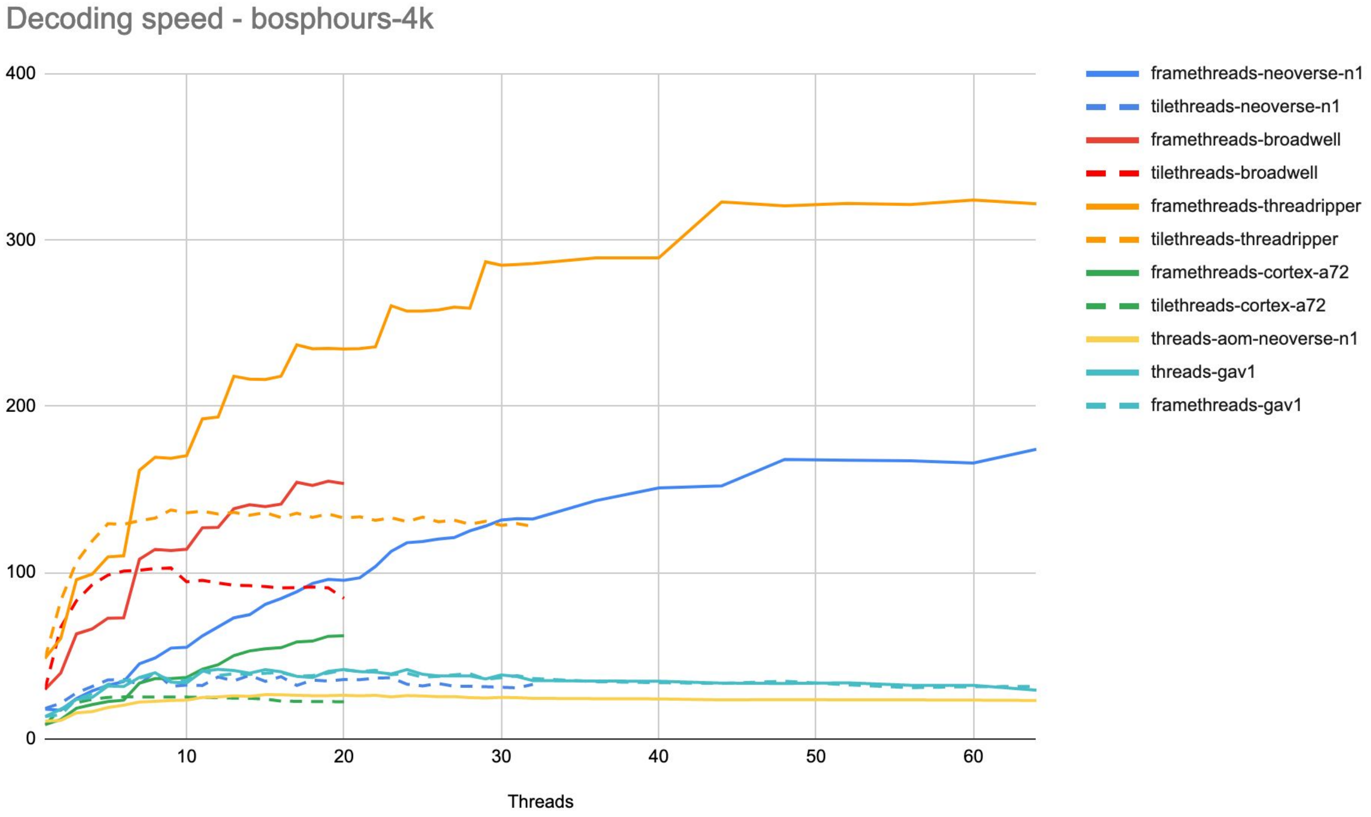

Encoding AV1

Encoding is always harder

- **x264** took about 7 years to become great
- **x265** needed nearly the same time to become a good competitor
  - It managed to leverage a good deal of experience
  - But making effective use of the HEVC features takes more effort.

Encoding AV1

Encoding AV1 is **not** different

- **libaom** and **SVT-AV1** are built upon years of past experience
  - **libvpx** and the whole **SVT** family of encoders
- Lots of effort is being spent in using well what makes AV1 amazing.

Encoding AV1 in Arm

Encoding AV1 since the beginning had been excruciating **slow**

- **libaom** proved you can have amazing quality
  - But at the cost of spending a huge amount of time for encoding
- **SVT-AV1** proved you can be fast
  - As long you have hardware and you are on x86_64

Encoding AV1 in Arm

Encoding AV1 since the beginning had been excruciating **slow**

- **rav1e** is efficient
  - It is written from scratch in Rust
  - It tries to cover more use-cases
    - Focuses on exploring different solutions and algorithms
    - Leverage what we learned from daala to improve perceptual quality over PSNR
  - Shares dav1d's hand written assembly

rav1e is an AV1 encoder

- rav1e is written in Rust
  - With a fair amount of arch-specific assembly for x86_64 (and aarch64)

What is rav1e?

rav1e is an AV1 encoder

- rav1e is written in Rust
  - With a fair amount of arch-specific assembly for x86_64 (and aarch64)
- rav1e can be used as a command line tool
  - available in many distros

What is rav1e?

rav1e is an AV1 encoder

- rav1e is written in Rust
  - With a fair amount of arch-specific assembly for x86_64 (and aarch64)
- rav1e can be used as a command line tool
  - available in many distros

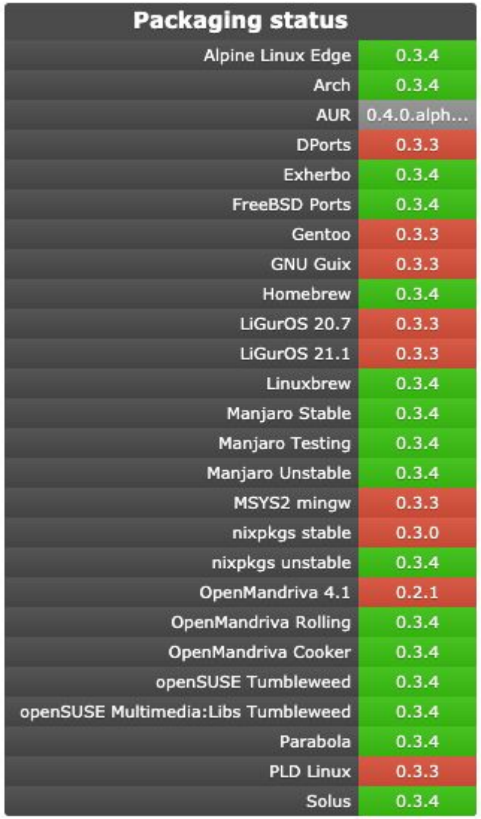

Packaging status

Alpine Linux Edge	0.3.4
Arch	0.3.4
AUR	0.4.0.alph...
DPorts	0.3.3
Exherbo	0.3.4
FreeBSD Ports	0.3.4
Gentoo	0.3.3
GNU Guix	0.3.3
Homebrew	0.3.4
LiGurOS 20.7	0.3.3
LiGurOS 21.1	0.3.3
Linuxbrew	0.3.4
Manjaro Stable	0.3.4
Manjaro Testing	0.3.4
Manjaro Unstable	0.3.4
MSYS2 mingw	0.3.3
nixpkgs stable	0.3.0
nixpkgs unstable	0.3.4
OpenMandriva 4.1	0.2.1
OpenMandriva Rolling	0.3.4
OpenMandriva Cooker	0.3.4
openSUSE Tumbleweed	0.3.4
openSUSE Multimedia:Libs Tumbleweed	0.3.4
Parabola	0.3.4
PLD Linux	0.3.3
Solus	0.3.4

What is rav1e?

rav1e is an AV1 encoder

- Or as a normal library using the common open source frameworks
  - VideoLAN's VLC,
  - GStreamer,
  - FFmpeg
  - libavif
  - Anything else that can consume a C or a Rust API.
- Used for Encoding Staff Pick Videos in Vimeo.

What to expect on Arm?

The encoding speed on a specific platform depends on its support

- **libaom** has lots of NEON-specific codepaths
- **SVT-AV1** has none
- **rav1e** has some, and more is slated to be integrated in the close future.

Encoding AV1 in Arm

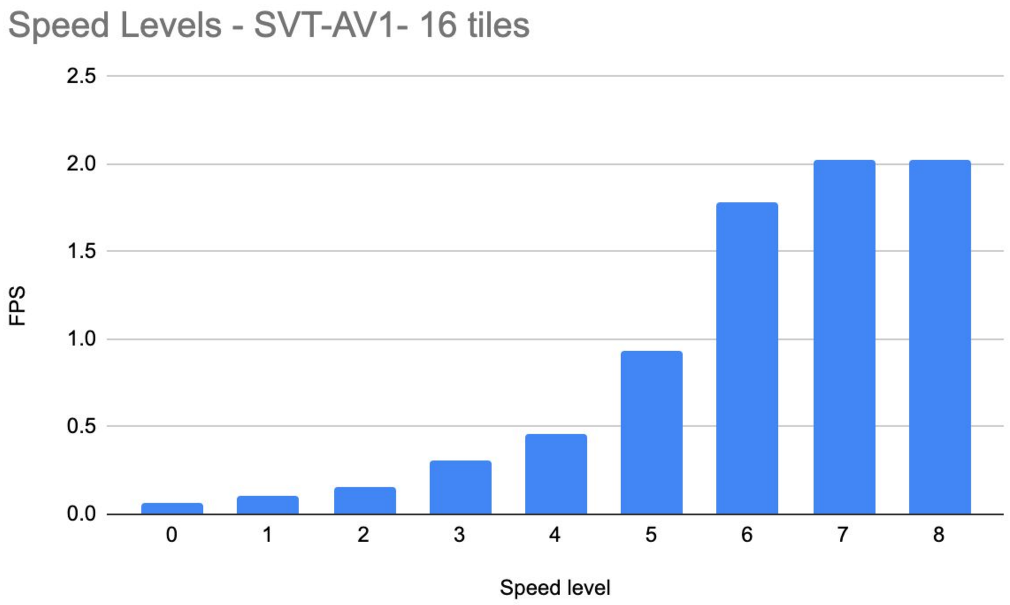

```
numactl --cpubind=0 --membind=0 ./SvtAv1EncApp -i ~/Samples/Bosphorus_1920x1080_120fps_420_8bit_YUV.y4m --preset 0 -b
~/Encodes/svt-bosphorus-1080p-p0-tiles16.ivf --tile-rows 2 --tile-columns 2 --lp 60 -n 96
```

Encoding AV1 in Arm

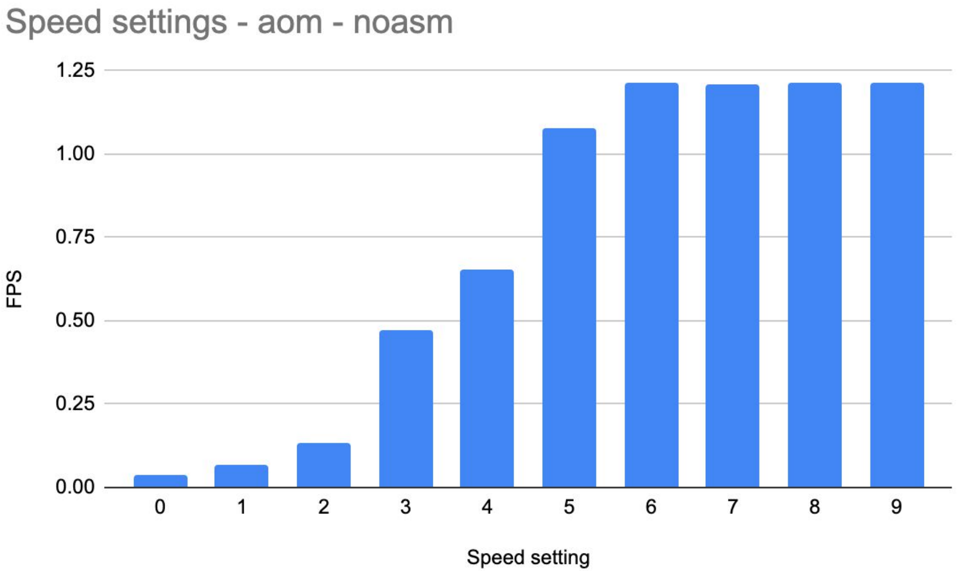

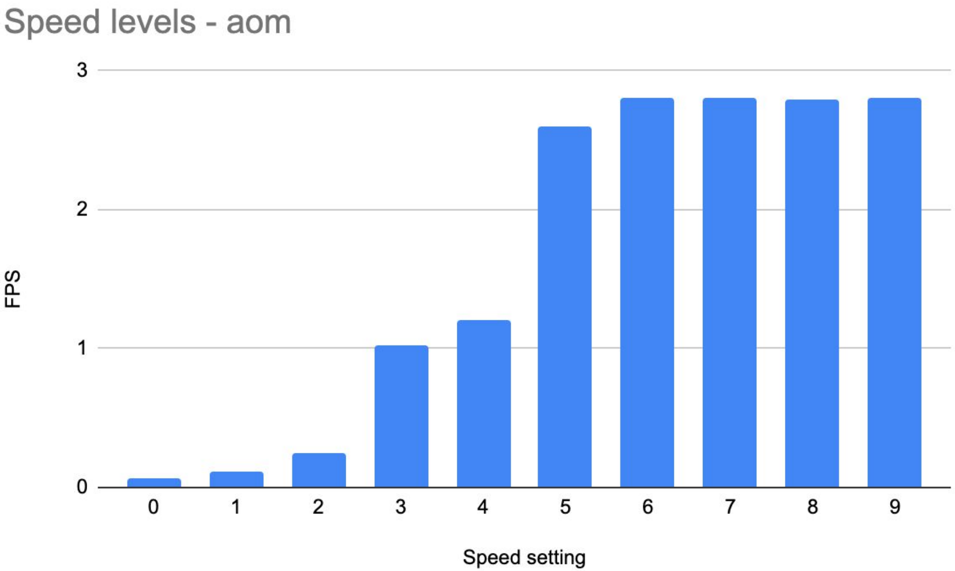

```
numactl --cpubind=0 --membind=0 ./aomenc --kf-max-dist=32 --tile-rows=2 --tile-columns=2 --cpu-used=0 --threads=60
~/Samples/Bosphorus_1920x1080_120fps_420_8bit_YUV.y4m --ivf -o ~/Encodes/aom-bosphorus-1080p-16tiles-cpu-used-0.ivf --limit=96
```

Encoding AV1 in Arm

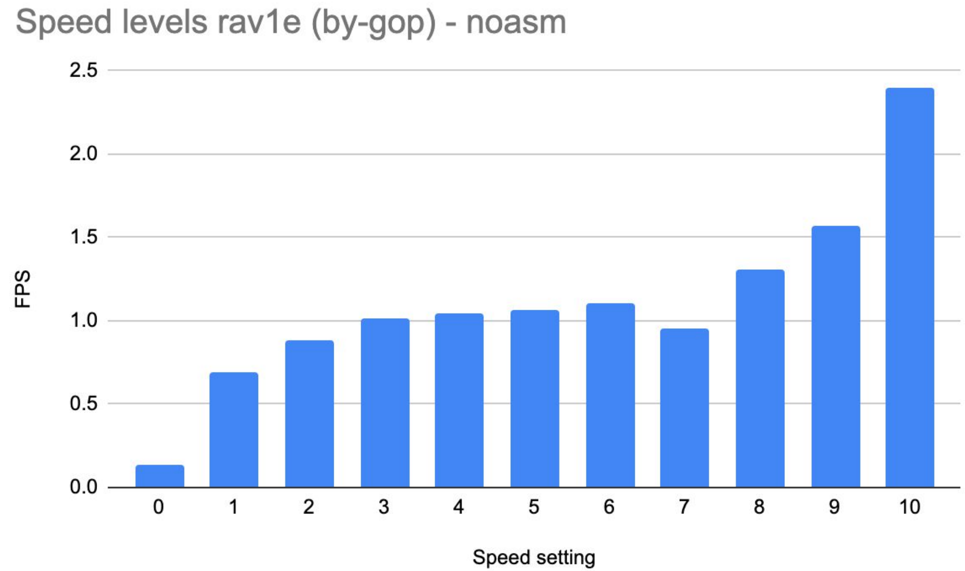

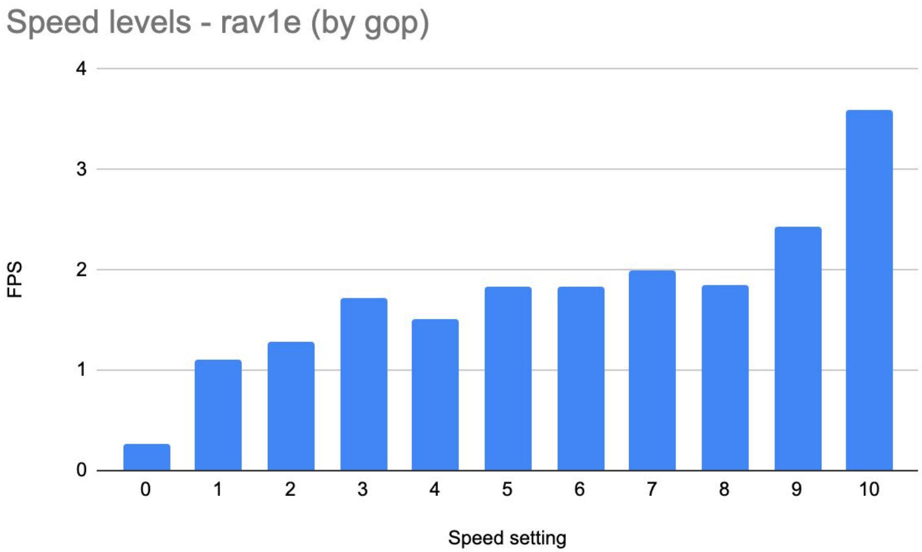

```
numactl --cpubind=0 --membind=0 ../target/release/rav1e-by-gop ~/Samples/Bosphorus_1920x1080_120fps_420_8bit_YUV.y4m --memory
unlimited --threads 60 -I 32 --tiles 16 -o ~/Encodes/rav1e-bosporus-1080p-s0-tiles16.ivf -s 0 --frames 96
```

Encoding AV1 in Arm

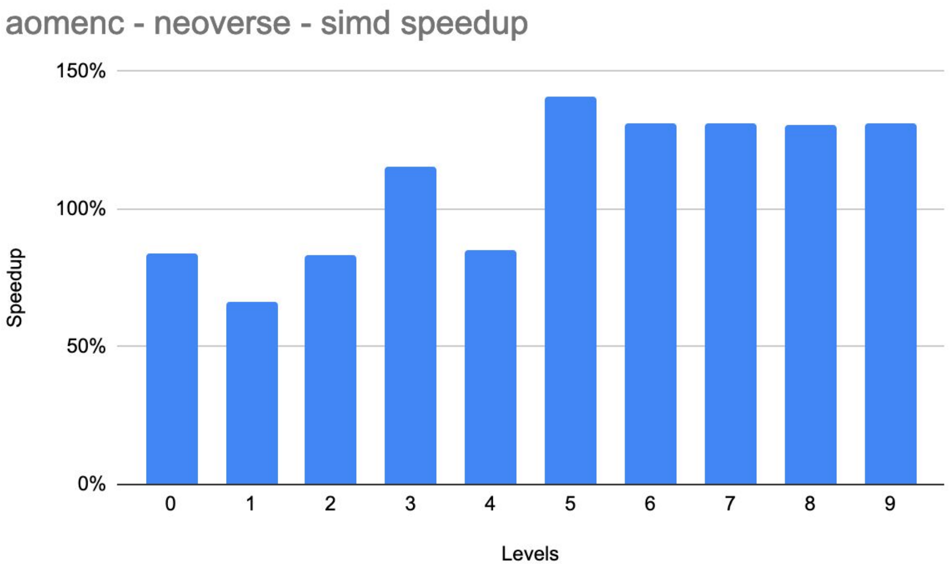

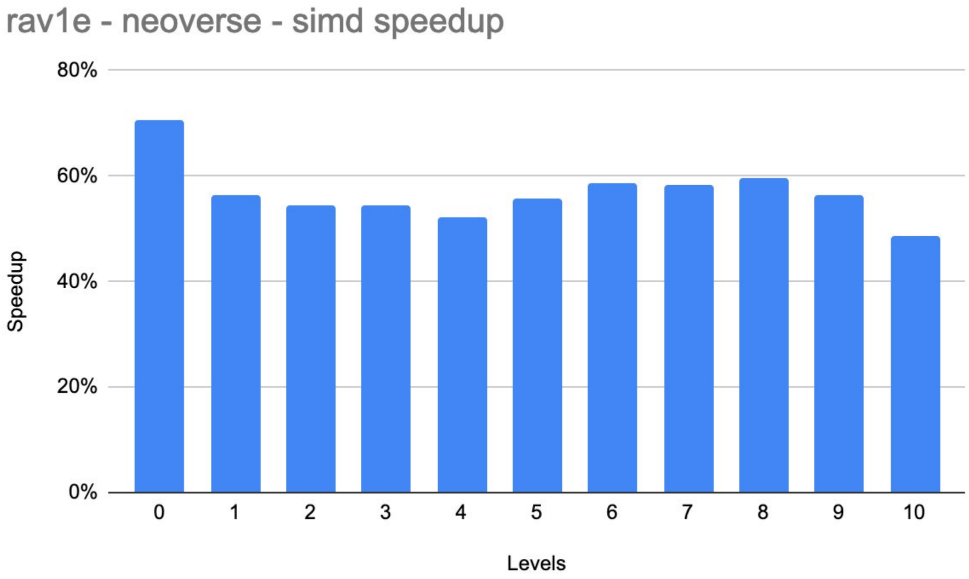

What is missing?

- **SVT-AV1** suffers the total lack of SIMD support
- **libaom** has the best SIMD support so far
- **rav1e** has partial SIMD support, but can leverage additional parallelism

What is missing in Arm

- More Hardware decoders

What is missing in Arm

- More Hardware decoders
- More available high performance server platforms
- Encoder Specific Optimisations

What is missing in Arm

- More Hardware decoders
- More available high performance server platforms
- Encoder Specific Optimisations
- Wider Adoption of Arm as laptop and desktop platforms

What is missing in Arm

- More Hardware decoders
- More available high performance server platforms
- Encoder Specific Optimisations
- Wider Adoption of Arm as laptop and desktop platforms
  - Honeycomb-lx2k is the best you can buy right now

What is missing in Arm

- More Hardware decoders
- More available high performance server platforms
- Encoder Specific Optimisations
- Wider Adoption of Arm as laptop and desktop platforms
  - Honeycomb-lx2k is the best you can buy right now
  - But the future is promising
    - Ampere Computing Altra
    - Apple Silicon
    - And much more

Questions?

Thank You