



COLLABORA

Linux Stateless Video Decoder Support

Nicolas Dufresne
Principal Software
Engineer



**Embedded Linux
Conference**
North America





COLLABORA

➤ **About me**


- Over 10 years at Collabora
- Core GStreamer developer
- Contribute to Linux Media



2011


The beginning of Linux CODECs

- Google partners with Samsung and Asus
- Produce the first ARM Chromebook
- Based on Exynos 5 SoC
- Includes Samsung MFC Decoder
- MFC V4L2 M2M driver landed mainline





V4L2 M2M



- A V4L2 output queue is used for the bitstream
- A V4L2 capture queue is used for the decoded pictures
- Additional control flow are added to support draining, flushing
- Inter-queue configuration flow is needed



Pros

- Minimal per CODEC code needed

Cons

- Requires a firmware
- Harder to multiplex



COLLABORA

CODA Driver

- CODA driver was added
 - Design from Chips&Media
 - CODA Hx4 and 960 support
- Enabling i.MX51 and i.MX6
- Reversed engineering

2014


© 2015 FREESCALE



The beginning of Linux State-Less CODECs

➤ **2015**

- Google partners with Rockchip
- 2nd gen of ARM Chromebook
- New type of CODEC hardware
- Rockchip VDPU ?





Reference(s)



Bitstream



Parameters



Pictur
e

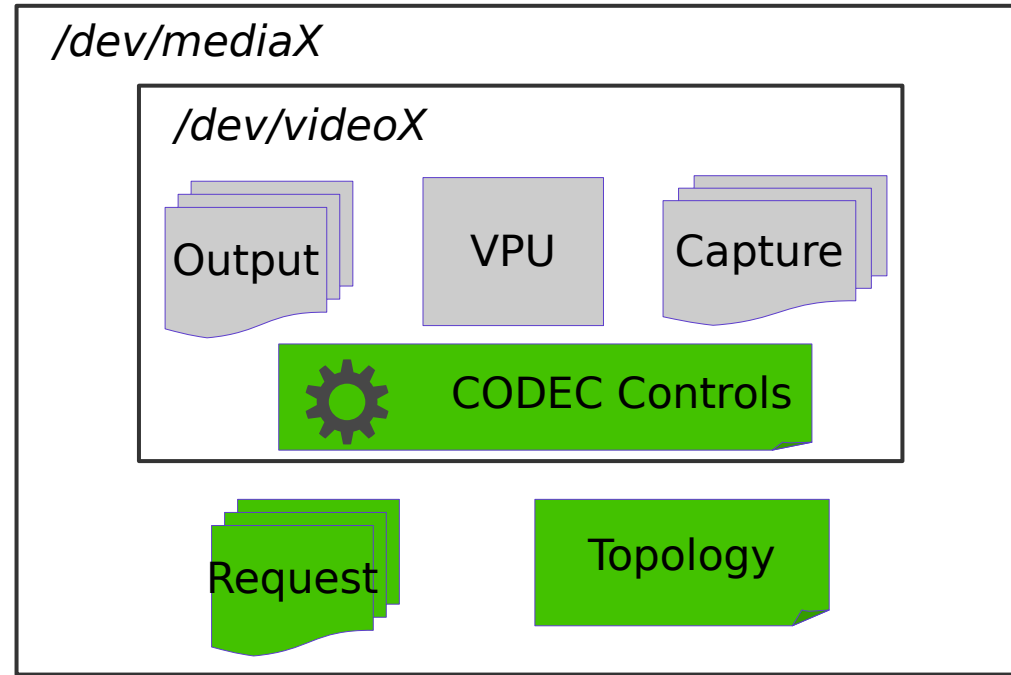


GPU

- Registers are replaced with a command as part of the command stream
- Crafting command stream is HW specific and is only implemented in userspace drivers (Mesa)
- Commands are scheduled by the GPU driver
- Exposing VPU in Mesa is done through standard APIs (VA API, VDPAU, DXVA2, NVDEC), but without a GPU, using these APIs can be cumbersome and overkill
- Using multiple GPU hardware in the same application remains tedious




**V4L2 M2M
+
Request**



- Per-frame (or slice) controls are associated with bitstream buffers using requests
- References are signaled using user defined timestamps
- Better identification can be done using the Media Controller topology



NALU Sequence



Annex B NALU



AVCc NALU





Decoding process

- Locate and parse NAL headers
- Parse non-VLC and VCL NAL (SPS, PPS, Slice Headers)
- Calculate frame_num (and handle gaps)
- Calculate POC and pic_num
- (Sliced Base) prepare reference lists
-



Decoding process (continued)

- Fill SPS / PPS, Decode parameters, Slice params V4L2 structures
- (Slice Based) Modify reference lists
- Decode the slices/frame
- Do DPB management as per spec
- Output frames that could be re-ordered



V4L2 Specific Process

- Allocate a Request (an FD)
- Set per-frame/slice params for this request
- Queue a v4l2_buffer for the request
- Queue the request
- Poll the request FD for completion



MediaTek VPU

➤ **2016**

- State-full H.264, VP8 and VP9 decode
- Tiled output only (requires HW converter)



Qualcomm Venus

➤ **2017**

- State-full MPEG4, MPEG2, H.264, VC1, H.264, VP8, VP9 and HEVC decode
- MPEG4, H.263, H.264, VP8, HEVC



Upstreaming Stalled

- Could not settle on the Request / Job API
- Low knowledge of CODEC decoding process by the linux-media maintainers
- Only one hardware to test the API design
- No formal specification (not that state-full CODEC had any either)



COLLABORA

Pocket
CHIP

Allwinner VPU support Kickstarter by Bootlin

2018

- Request API is finalized
- MPEG2 Support landed in Staging
- H.264 support was progressing (but only sliced based)
- Reversed engineered from binary userspace blob
- VAAPI userspace drivers



➤ **2019**

- The crowd funding had gain good momentum
- A formal specification was merged
- H264, VP8, HEVC uAPI added as staging control API
- Cedrus gained H264 and HEVC support
- RK3288 driver was mainlined (MPEG2, H264, VP8)
- RK3288 driver was renamed !?!



The Hantro Driver

- I.MX8M Quad, using Hantro G1/G2
- Registry compatible with the RK3288





COLLABORA


**Hantro
Company**

Hantro
VISIBLY BETTER

on2
technologies

Google

VeriSilicon



- ``stmmac`` is an ethernet driver, that was thought to be an STM design
- It was later found to be DesignWare design, shared across numerous SoC
- Still called ``stmmac`` to maintain kernel API stability



Testing and Fixing

- FFMPEG support
- Bug fixing
- Interlaced Content Support



2020

- RK3399 JPEG, MPEG2, H.264 and VP9 support landed
- GStreamer gained base classes for state-less CODEC, with already merge DXVA2 and NVDEC support
- **GStreamer** H.264 and VP8 V4L2 support landed
- Embedded World Conference 2020 **canceled**
- VA V4L2 Request driver was abandoned
- And a lot more coming ...
- (at 2m distance of course)



COLLABORA





COLLABORA

