FFMS2: Indexing, Edge Cases, and Insanity

Hiding the Complexity of Frame-Accurate Seeking

Derek Buitenhuis

@daemon404

derek@videolan.org

EBU
OPERATING EUROVISION AND EURORADIO

15 September 2018
Amsterdam, NL

$ whoami

- Senior Video Encoding Engineer at Vimeo
  - Distributed transcoder
  - On-the-fly image service (before it was cool)
  - HLS/DASH on-the-fly packager
  - Subtitles / captions ingest and output
- Open Source Multimedia Developer
  - FFmpeg, L-SMASH, FFMS2, etc.
  - Daemon404 on freenode
- VideoLAN Non-Profit Association Member

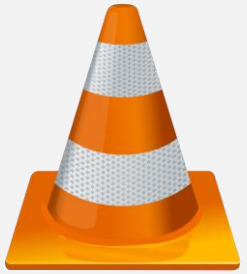

Frame / Sample Accurate Seeking is Hard

- Not everyone is lucky enough to get known-format deliverables (iTunes, IMF, DPP, etc.)
  - User-generated content can, and will by anything

Frame / Sample Accurate Seeking is Hard

- Not everyone is lucky enough to get known-format deliverables (iTunes, IMF, DPP, etc.)
  - User-generated content can, and will by anything
- Streams can be (extremely non-exhaustive):
  - Non-indexed (MPEG-TS, MPEG-PS, etc.)
  - Corrupt
  - Contain virtual timelines (edit lists, ordered chapters)
  - Have invisible frames
  - Crazy timestamps

Frame / Sample Accurate Seeking is Hard

- Not everyone is lucky enough to get known-format deliverables (iTunes, IMF, DPP, etc.)
  - User-generated content can, and will by anything
- Streams can be (extremely non-exhaustive):
  - Non-indexed (MPEG-TS, MPEG-PS, etc.)
  - Corrupt
  - Contain virtual timelines (edit lists, ordered chapters)
  - Have invisible frames
  - Crazy timestamps
- Only feasible solution is FFmpeg (libavcodec / libavformat)
  - Doesn’t seek frame accurately out of the box
  - Unstable, verbose, massive API (much of which is missing, and actually in ffmpeg.c)

FFMS2: We Go Crazy So You Don’t Have To

- Indexes files or streams and builds packet-to-frame map
  - Non-trivial due to things like NVOPs, alt-refs, PAFF, libavformat’s implementation of edit lists, etc.
  - Proper concept of timestamps (each frame has multiple types!)

FFMS2: We Go Crazy So You Don’t Have To

- Indexes files or streams and builds packet-to-frame map
  - Non-trivial due to things like NVOPs, alt-refs, PAFF, libavformat’s implementation of edit lists, etc.
  - Proper concept of timestamps (each frame has multiple types!)
- Stable, mature, threadsafe, easy to use C API
  - Fully documented (whack me if not!)
  - Stable ABI (10 years!)
  - AviSynth and VapourSynth plugins; built into x264
  - Proper support for mid-stream changes, real colorspace info (HDR included), etc.
  - Proper / useful error handling
  - MIT licensed

```
const FFMS_Frame *curframe = FFMS_GetFrame(videosource, framenumber, &errinfo);
```

FFMS2: Coming Soon

- Custom I/O Callbacks
- Caller-owned Frame Buffers
- Mid-stream codec changes (“soon”)
- Virtual Timeline Metadata Export (need to finish upstreaming new APIs to FFmpeg first)

Links and Info!

- FFMS2 (source, docs, binaries): https://github.com/ffms/ffms2/
- FFMS2 Index Parsing in Go (and JSON dump util): https://github.com/dwbuiten/dumpindex/
- VapourSynth: https://vapoursynth.com **/** https://github.com/vapoursynth/vapoursynth/
- avformat Timeline API RFC v1: http://ffmpeg.org/pipermail/ffmpeg-devel/2018-March/227437.html
  - v2 + implementation in the works