

All your multimedia are belong to us

Linux Conference Australia 2015
Auckland, 2015-01-16

Diego Biurrun

\$ whoami

Interest in free software at the end of the 90s, downhill since then:

- Mozilla QA in 2000
- MPlayer in 2002
- FFmpeg in 2004
- Libav since 2011

Outline

- 1 Technical Overview
- 2 Legal Issues
- 3 The Split
- 4 Questions

What is Libav/FFmpeg all about?

- FFmpeg started by Fabrice Bellard in 2000
- trivia: “FF” stands for “fast forward”
- collection of libraries for all things multimedia
- bunch of command line utilities
- ~600kloc in Git
- LGPL licensed (with some GPL video filters)

Community

- mostly volunteers
- dislike
 - C++
 - non-C99-compliant compilers
 - autotools
- like
 - speed
 - standards-compliance (C99, POSIX)
 - portability (POSIX)
 - *BSD, *libc, Cygwin, MingGW*, Haiku, OS/2, OS X, ...
 - gcc, clang, icl, MSVC, Sun CC, ...
- flames
- troll-driven development

Total World Domination

- free world: VLC, GStreamer, MPlayer, Blender, XBMC/Kodi, Audacity, . . .
- proprietary world: shareware multimedia conversion tools
- web services: Vimeo, YouTube, Facebook, . . .

Command Line Tools

- avconv(ert) - transcoding tool
- avplay - minimal SDL-based player
- avprobe - multimedia file analyzer

Libraries

- libavutil
- libavformat
- libavdevice
- libswscale
- libavfilter
- libavresample
- libavcodec

- ~19kloc
- helper code
- CPU detection, math, crypto, logging, strings, buffering
- general purpose structures and data types

- ~92kloc
- containers (AVI, ASF, MP4/MOV, Matroska, Ogg, ...)
- network protocols ((S)RT(S)P, HTTP(S), HLS, MMS(T), SCTP, Gopher, ...)

- ~5kloc
- input devices (ALSA, X11, Video4Linux, PulseAudio, ...)

libwscale

- ~16kloc
- colorspace conversion
- scaling
- legacy ugliness

- ~17kloc
- video filters (deinterlace, scale, crop, effects, ..)
- audio filters (channel or volume transformations, effects, ...)

libavresample

- ~6kloc
- audio resampling and audio channel mapping

- ~411kloc
- decoders and encoders, especially decoders
- hardware accel (VDPAU, VA-API, ...)
- subtitles (SRT, ASS/SSA, ...)
- the project's crown jewels

libavcodec Codecs

- audio (~130)
 - MP2, MP3, (E)AC-3, DTS, AAC, Vorbis, Opus, (A)(D)PCM, ALAC, AMR, RealAudio, FLAC, NellyMoser, Windows Media Audio, ATRAC 1/3, ...
- video (~180)
 - H.261, H.263(+), MPEG-1/2, MPEG-4, DivX, H.264, HEVC, Flash, DV, Theora, VP3/5/6/7/8/9, MJPEG, JPEG-2000, Indeo 2/3/4/5, SVQ1/3, RealVideo 1/2/3/4, Apple ProRes, Windows Media Video 1/2/3, ...
- fringe formats


Architecture-Specific Optimizations

- x86, ARM, PowerPC assembly (SIMD) - all in 32 or 64 bit variants
- core routines exist as C and assembly variants
- optimized code accessed through function pointers set at runtime

Legacy Spaghetti Code


much of the old code monolithic

- *MPEG* code
- MpegEncContext (308 lines structure)
- dsputil (now deceased)
- libswscale
- libavformat depends on libavcodec



- huge attack surface due to input handling
- old code written with a cowboy attitude
- tools: fuzzing, Coverity, grep

Future



Future

- multimedia world standardizing, format zoo legacy only
- more codecs (DTS-HD, GoToMeeting)
- refactoring
- stability work
- hardware acceleration

Wishlist

- Cineform codec
- libswscale replacement
- new hwaccel framework
- paid maintainers

Build System

- POSIX shell and GNU Make only
- fully featured (dependencies, cross-compilation, toolchains)
- fast
- it works (TM)
- possibly future standalone project

- FATE Automatic Testing Environment
- ~2000 individual tests
- mostly full system tests, not unit tests
- clients push results to collector
- results aggregated on a web page
- FATE sandbox: Oracle (will tell your FATE)

FATE

FATE Automated Test Environment

[master](#)
[release/11](#) 53
 [release/10](#) 48
 [release/9](#) 46
 [release/0.8](#) 28

[Architectures](#) 8
 [Variants](#) 16
 [Operating systems](#) 16
 [Compilers](#) 26
 [Configurations](#) 262
 [Contributors](#) 11

Time	Rev	Arch	OS	Compiler	Comment	Warnings	Tests
1 minute ago	v12_dev0-669-g402fb55	x86_64	linux	gcc 4.7.4 (Gentoo Hardened 4.7.4 p1.2, ple-0.5.5)	CPUFLAGS=xop	128 ±	1949 / 1949
2 minutes ago	v12_dev0-669-g402fb55	x86_64	linux	gcc 4.7.4 (Gentoo Hardened 4.7.4 p1.2, ple-0.5.5)	CPUFLAGS=fma3	128 ±	1949 / 1949
4 minutes ago	v12_dev0-669-g402fb55	x86_64	linux	gcc 4.7.4 (Gentoo Hardened 4.7.4 p1.2, ple-0.5.5)	CPUFLAGS=fma4	128 ±	1949 / 1949
5 minutes ago	v12_dev0-669-g402fb55	x86_32	mingw32	gcc 4.7.4 (GCC)		134 ±	1949 / 1949
6 minutes ago	v12_dev0-669-g402fb55	x86_64	linux	gcc 4.7.4 (Gentoo Hardened 4.7.4 p1.2, ple-0.5.5)	--disable-mmxx	143 ±	1949 / 1949
8 minutes ago	v12_dev0-669-g402fb55	x86_64	sunos	gcc 4.6.4 (GCC)		362 ±	1858 / 1858
8 minutes ago	v12_dev0-669-g402fb55	x86_32	openbsd	clang version 3.3 (tags/RELEASE_33 /final)	--disable-inline-asm	264 ±	1847 / 1938 ▶
8 minutes ago	v12_dev0-669-g402fb55	x86_64	linux	gcc 4.7.4 (Gentoo Hardened 4.7.4 p1.2, ple-0.5.5)	--disable-mmxxext	232 ±	1949 / 1949
8 minutes ago	v12_dev0-669-g402fb55	x86_32	mingw32	gcc 4.6.4 (GCC)		230 ±	1949 / 1949
9 minutes ago	v12_dev0-669-g402fb55	ppc	aix	gcc 4.9.2 (GCC)		271 ±	1947 / 1949 ▶

FFmtech foundation

- 501(c)(3)
- sponsors development of codecs, optimizations, reverse engineering, infrastructure cost
- non-profit status rejected

Outline

- 1 Technical Overview
- 2 Legal Issues
- 3 The Split
- 4 Questions

Legal Landscape

multimedia is a legal minefield

- entities collecting patent royalties
 - MPEG-LA (MPEG-2, H.264, HEVC)
 - AAC
 - Fraunhofer (MP3)
 - DTS
 - Dolby
- reverse engineering
- trademarks (DTS, AC-3)
- sometimes murky copyrights / licensing

Legal Landscape

multimedia is a legal minefield

- entities collecting patent royalties
 - MPEG-LA (MPEG-2, H.264, HEVC)
 - AAC
 - Fraunhofer (MP3)
 - DTS
 - Dolby
- reverse engineering
- trademarks (DTS, AC-3)
- sometimes murky copyrights / licensing

—> clearly, jail and doom await

Legal Issues (2000)

Legal Issues (2001)

Legal Issues (2002)

Legal Issues (2003)

Legal Issues (2004)



Breaking Eggs And Making Omelettes

Topics On Multimedia Technology and Reverse Engineering

Legal Threat #00001

© May 26th, 2005 by ☿ Multimedia Mike

Do you have any idea how long I have been involved in multimedia hacking and reverse engineering? About 5 years now. All that while, folks have warned me sternly, and constantly, that this type of work would get me sued to death. I am pleased to announce that today I received my first legal threat. I feel that my work has finally been validated!



Well, it was not necessarily a legal threat, like those notorious "nastygram" cease & desist letters. It was more like a veiled reference to a possible future legal threat. Someone identifying himself as the assistant general counsel for On2 said that the company took exception to the fact that I was posting decompilations of their Java decoder.

And just when I was starting to feel that no one cared about my work...

Naturally, this raises some pressing questions. First and foremost, why was I contacted by the *assistant* general counsel? Why doesn't my case warrant the attention of the lead/primary/head general counsel? Maybe if I went after their latest generation codec, VP7, my actions would merit an escalation.

For the time being, I have decided to not post the Java decompilations on my [Practical](#)

Links

- [About Multimedia Mike...](#)
- [About This Blog...](#)
- [Clone2727's Blog](#)
- [Coding Horror](#)
- [Diary Of An x264 Developer](#)
- [Email Multimedia Mike \(mike-at-multimedia.cx\)](#)
- [Flameeye's Weblog](#)
- [Gaming Pathology](#)
- [ginger's thoughts](#)
- [Hardwarebug](#)
- [Kostya's Wild Codec World](#)
- [Lair Of The Multimedia Guru](#)
- [Multimedia Exploration Journal](#)
- [Penguin.SWF](#)

On2 Technologies and the libvp6 incident



Legal Issues (2006)

Legal Issues (2007)

Nellymoser, Inc and ASAO voice codec IP

The logo for Nellymoser, featuring the word "nellymoser" in a lowercase, sans-serif font. The letter "n" is light gray, and the dot above it is a solid orange circle. The remaining letters "ellymoser" are also in a light gray color.

Legal Issues (2009)

Legal Issues (2010)

letter from the MPEG-LA

Legal Issues (2011)

Legal Issues (2012)

Legal Issues (2013)

Legal Issues (2014)

Srsly?

Real legal exposure does exist, but

- trademarks are easy to avoid,
- entities go after (large) commercial users,
- don't care if software is FOSS or proprietary,
- there's always the j-b defence.

Srsly?

Real legal exposure does exist, but

- trademarks are easy to avoid,
- entities go after (large) commercial users,
- don't care if software is FOSS or proprietary,
- there's always the j-b defence.

—> Don't be a coward, just hack!

Copyright Situation

- fully free software, LGPL v2.1 or later
- all nonfree pieces excised
- GPL bits mostly replaced

Outline

- 1 Technical Overview
- 2 Legal Issues
- 3 The Split
- 4 Questions

Prelude

- by 2010/2011 morale in FFmpeg low
- project dictatorial and centralized
- development stagnating, patches stalled
- many developers unhappy and frustrated
- by beginning of 2011 many (core) devs have dropped out

Revolution

- rebels begin to gather on IRC, decide to fork
- notice they are a large majority
- declaration of independence signed by 18 devs
- new Git tree
- leadership team of 7 developers

Result



Counterrevolution

- (un)surprisingly the loyalists are shocked and angry
- a vote is started, only loyalists participate, 7 people in favor
- loyalists receive control of ffmpeg.org domain from Fabrice
- set up fresh project server and new infrastructure
- revolutionaries move to libav.org domain

Aftermath

- both groups of developers stay separate
- run separate projects with different policies
- development speed picks up

- organized as a community of equal peers
- strict four eyes principle
- code of conduct
- goal: quality and maintainability
- real world presence
- no public statements about FFmpeg

FFmpeg

- evil overlord
- merge everything from Libav and other sources
- free speech mailing list
- goal: more features
- mail and IRC presence
- public insults for Libav (developers)

My Personal Opinion

My Personal Opinion

Linux Conference Australia 2015 code of conduct:

All presentation material should be suitable for people aged 12 and above. All public presentations are subject to this code of conduct and thus may not contain:

- *sexual or violent imagery*
- *exclusionary language*
- *language which is not appropriate for an all-ages audience*

Outline

- 1 Technical Overview
- 2 Legal Issues
- 3 The Split
- 4 Questions

Libav and You

- patches welcome
- paid maintainers needed

Questions

?