

Kyber

Interactive Streaming

(near) Real Time Video and Controls streaming

FRnOG 39






Jean-Baptiste Kempf

Resume of My Video Life


- **VideoLAN**: being working on the VideoLAN project for 17 years, president of the VideoLAN NPO since its creation (2008)
Doing most of the non-developer tasks of VideoLAN
- **VLC**: Active developer since 2006, notably on GUI, Windows, Android ports, codecs and demuxers, packaging and releases
- **FFmpeg**: Active community member and peacekeeper of FFmpeg. De facto involved in releases and roadmaps.
- **Shadow**: ex-CTO of Cloud Gaming/Desktop company
- **Technical Consultant**: Video startups, scaleups and e-Commerce business





Lower Latency

Lower Latency Streaming



We're talking about encoded latency, else we would talk in lines :)

Why do we need lower latency?

Interactivity

Remote Desktop

VM / VDI / DaaS
Cloud Gaming
Cloud Desktop



Robots & Drones

Robots, Drones
AR Supervision
Cars?



Safety & Critical

Remote Monitor


Virtual Monitor
SlingBox
Industrial Supervision



App Streaming

Remote Video Production
Trial / Demo of Apps
Visual Cloud Apps





Kyber

Watizit?

Kyber

Open Source

Real Time

Control of Machines

Solution SDK

Client, Server and Networking stack
Streaming video, audio, *subtitles* unidir
Streaming inputs bidirectionally
Modular SDK and application
Quic protocol & ~~Web~~RTCP

Multi-platform Client (+ Web)
Multiple platforms for the Servers
All Codecs (H.264, HEVC, VP9, AV1)
Multiple Hardware & Software Encoders
Based on VLC and FFmpeg

Demo




Windows

Linux Server

The image shows a Linux desktop environment with a blue border. The main window is a Google Chrome browser displaying a YouTube video. The video title is "On ne pourrait pas refaire VLC aujourd'hui. JB Kempf (VLC)". The video content shows two men in a recording studio, one wearing headphones and using a laptop, and the other wearing glasses and headphones, sitting at a desk with a mixing console and microphones. The video player interface includes a progress bar at 0:54 / 54:53, a play button, and a volume icon. Below the video, the channel name "Burst" with 3.2K subscribers and a "Subscribe" button are visible. To the right of the video, there are icons for likes (1.8K), comments, share, and save. A small thumbnail for another video titled "OSAP VLC : LE START-UPPER QUI NE VOULAIT PAS ÊTRE RICHE" is also present. The desktop taskbar at the bottom shows various application icons, including a search bar with the text "Taper ici pour rechercher", and the system tray on the right shows the date "17:10 19/10/2023" and the language "FRA".

How does it work?



Features

Video

Audio-Video Server based on FFmpeg libraries, through txproto

Pushed-based Streaming server, graph-based and multi-threaded per node
Video Server can composite GPU overlays

Player based on libVLC, tuned for O-latency (*push-based approach*)

Video Codecs tested: *H.264, HEVC, AV1, VP9*

Audio Codecs tested: *PCM, Flac, Opus*

Hardware and Software encoders

4:2:0 -> 4:4:4 upsampling on client

Input

New Input Server written in Rust from scratch

Push-based Input streaming server, graph-based, able to filter and merge inputs

I/O Support:

Keyboard, Mouse (+*Cursor*),
Gamepads (+*Rumble*), Copy-Paste,
File-Transfer, USB/IP

Virtual Video, Mouse, Keyboard
and Gamepad Drivers

Cross-Platform, Client = Server

Network

Multiplexer Server written in Rust from scratch

Multi Protocol: Quic and WebTransport

Opens only one port (*TCP+UDP*)

TLS and Security handled at connection

Multi-user support (*Main, Student*)

Selectable features (audio, video, inputs)

Input latency is independent from Video latency

Separate process

Measures

Desktop @ 60 Hz

libVLC H.264: ~16ms - 1 frame

Desktop @ 120 Hz

libVLC H.264: ~12ms - 1 ½ frames

libVLC HEVC: ~12ms - 1 ½ frames

Desktop @ 240 Hz

libVLC HEVC: ~10ms - ~2 ½ frames

libVLC H.264: ~10ms - ~2 ½ frames

Web @ 60 Hz

Kyber H.264 Soft: ~33ms - 2 frames

Kyber HEVC Hard: ~33ms - 2 frames

Web @ 120 Hz

Kyber HEVC Hard: ~16ms - ~2 frames

Kyber H.264 Soft: ~24ms - ~3 frames



We can do better!

Extra Low



We can do better!

Protocol

Quic / WT

Multiple protocols supported by the muxer

Big focus on Quic, because TLS, Uni-Socket, Multi-Stream, Bi-Directional and Datagrams

Audio / Video data can use Datagrams
Inputs are reliable and Bi-Directional

Use of WebTransport in very similar way than Quic, including DataGrams inside WebBrowser

Multiple Modes

Reliable: sends each channel on one QUIC/WT stream

GOPstream: 1 QUIC/WT stream per GoP

Unreliable:

Video Packets are sent in DataGrams mode
Config and Control packets are always sent in reliable streams

Unreliable_fec:

Use of FEC (RaptorQ) to recover info without needing retransmissions when loss of packets

Unreliable Protocol

Channels

One connection with multiple channels, similar to MoQ tracks

Request to Server to subscribe to the right channels (*channel_id*)

Stream vs Datagram


Some channels are either on Datagram or Stream mode

One channel is composed of multiple Streams or multiple Datagrams

Video can be datagrams and Inputs are always BiDi streams

Groups

Packets are grouped in groups for Media, to keep config (SPS/PPS) + data together




Unreliable Protocol + FEC

FEC

Use of RaptorQ (*for now*) but other schemes are possible

RaptorQ source symbols of size configured to maximize the max Quic Datagram size.

Hopefully, the RaptorQ encoded symbols are going on their own UDP datagrams



Web Version

Wasm

All the desktop Rust code is running in the web browser using Webassembly with same codebase

Notably Muxer and Input Server/Client

WebTransport

The Web version is using, in Rust/Wasm the same codebase and the same protocol

WebTransport instead of Quic

FEC running in Wasm

WebCodec

Video Decoding is done through WebCodec, in Rust/Wasm

Rendering is done through a Canvas
Compositing time is controlled by the app

Audio is done through WebCodec

VLC.wasm

Video Decoding can also be done through VLC.wasm

No WebRTC :)



Thanks



Do you have any questions?
jb@videolan.org
kyber.media



Processing duration

