High quality in a low latency world

AV1 for live streaming and RTC

Thomas Davies thomas@visionular.com
Dmitriy Sychov dmitriy@visionular.com

Visionular

MHV2024, Denver

Low latency high performance AV1 (LD AV1)

- Specialised RTC codebase
  - Low memory + codesize + sw overheads
  - Configurable speed settings
- Intelligent content classification and mode decision
- Adaptive resolution
- Fine-grained parallelisation

Low latency high performance AV1 (LD AV1)

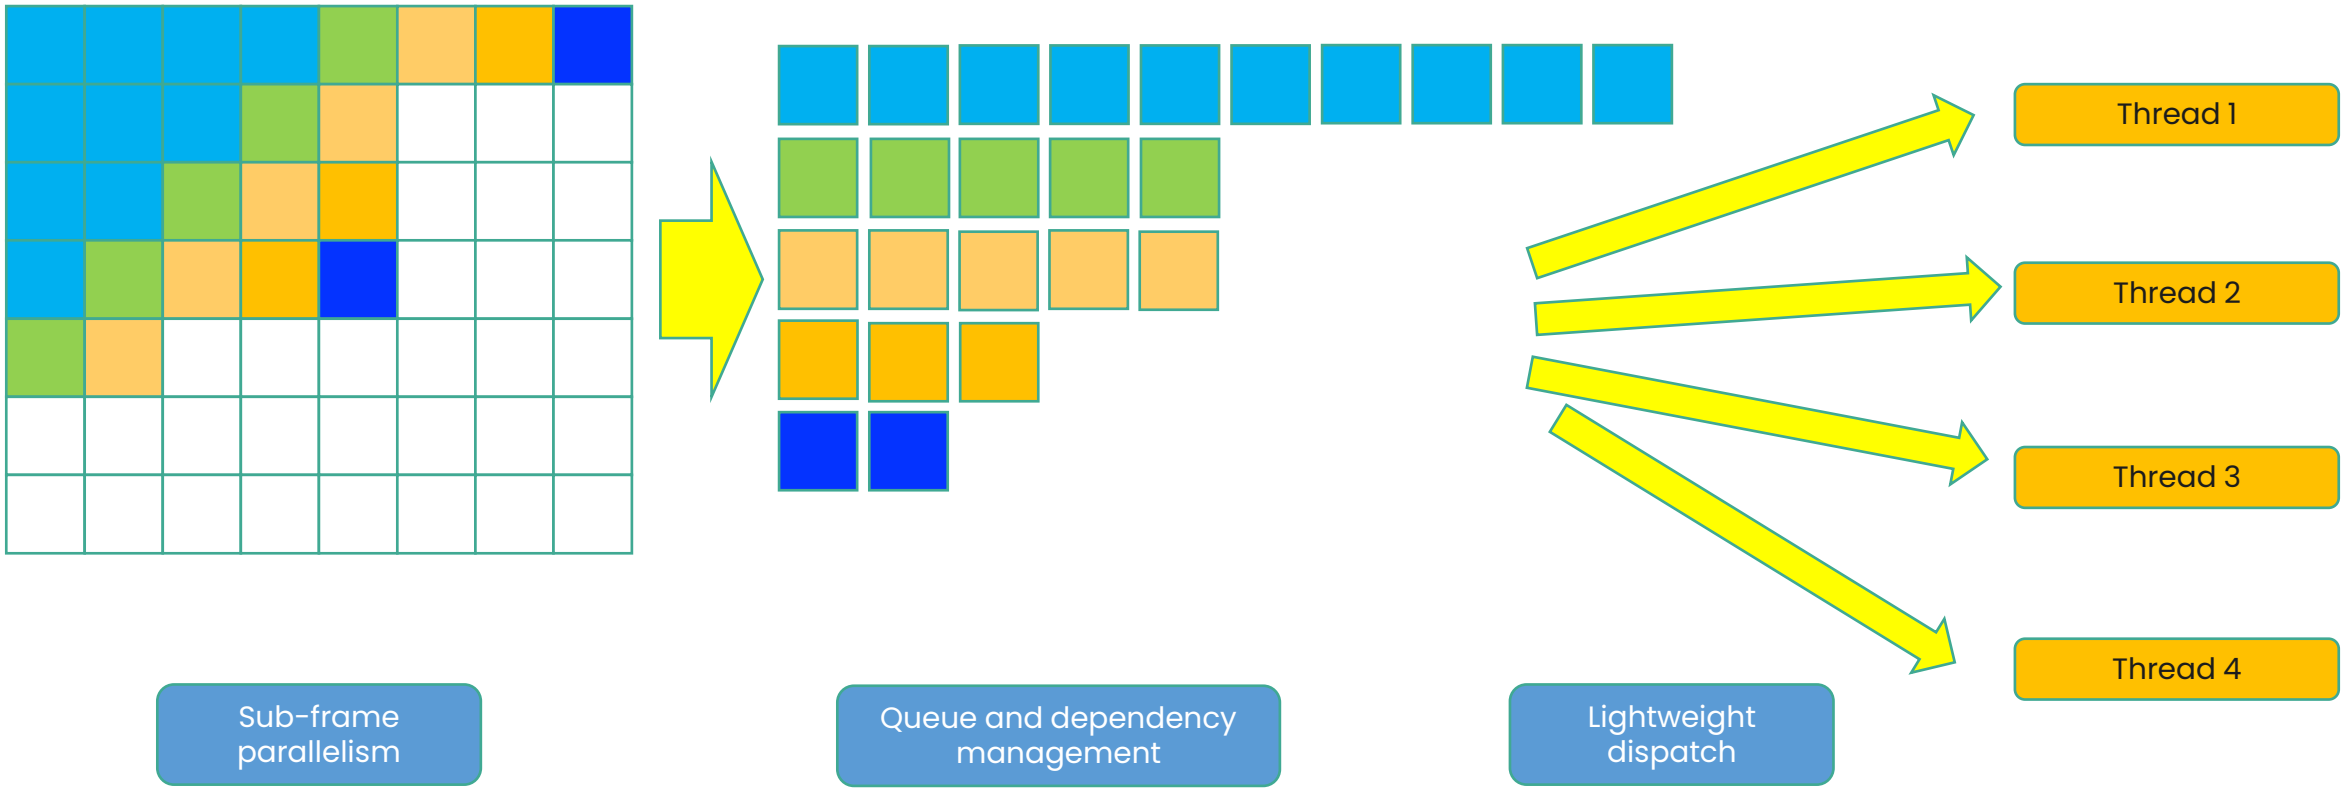

Three low delay encoding scenarios

Scenario 1

Mobile ARM encoding on iPhone 12

Video Conferencing and screenshare

Scenario 2

Mobile ARM encoding on M1 processor (MacMini)

Video Conferencing and screenshare

Compare with libaom-rt and HW HEVC

Scenario 3

x86 encoding on AMD Ryzen server

Live streaming: mixed content

Compare with x264

Scenarios 1&2: Mobile RTC

Resolutions and video characteristics	1080p25/30, 720p30 Portrait + landscape
	Camera: conferencing low to high illumination, low to high activity Screen content: PPT, word, excel, editing and similar
Speed settings	LD AV1: veryfast, superfast, ultrafast, ultrafast2 libaom-rt: speeds 6-10
Threads	1, 2, 4
Bitrates	1000 - 2400 kb/s

Performance versus libaom (single threaded)

libaom speed	LD AV1 speed	Content type	Combined BDR %	Encoding speedup %†
6	Superfast	Camera	-6.09	63
7	Ultrafast	Camera	-2.77	87
10	Ultrafast2	Camera	-2.63	40
libaom speed	**LD AV1 speed**	**Content type**	**Combined BDR %**	**Encoding speedup %**
6	Veryfast	Screen	-75.0	83
7	Superfast	Screen	-63.3	170
8	Ultrafast	Screen	-71.70	220
10	Ultrafast2	Screen	-59.21	93

†$100 * (FPS - FPS_{LIBAOM})/FPS_{LIBAOM}$

Performance vs HW HEVC

Speed setting	Content type	BD-VMAF	BD-PSNR (dB)
Veryfast	Camera	12.85	3.50
Superfast	Camera	12.71	3.38
Ultrafast	Camera	12.48	3.22
Ultrafast2	Camera	12.09	2.97

Speed setting	Content type	Combined BDR %
Veryfast	Screen	-75.74
Superfast	Screen	-74.69
Ultrafast	Screen	-72.92
Ultrafast2	Screen	-67.23

MacMini M1

Speed setting	Platform	FPS 1 thread	FPS 2 threads	FPS 4 threads
Veryfast	Camera	**42.34**	78.31	131.50
Superfast	Camera	**54.46**	101.51	172.05
Ultrafast	Camera	**78.12**	142.94	237.76
Ultrafast2	Camera	**101.69**	182.57	297.10
Speed setting	**Platform**	**FPS 1 thread**	**FPS 2 threads**	**FPS 4 threads**
Veryfast	Screen	**126.98**	201.39	260.00
Superfast	Screen	**186.57**	283.69	351.52
Ultrafast	Screen	**211.87**	332.45	417.52
Ultrafast2	Screen	**271.18**	397.18	501.92

iPhone 12

Speed setting	Platform	FPS 1 thread	FPS 2 threads	FPS 4 threads
Veryfast	Camera	25.57	34.66	46.09
Superfast	Camera	33.24	**46.17**	58.80
Ultrafast	Camera	**43.58**	**58.37**	77.68
Ultrafast2	Camera	**58.59**	**80.10**	94.39
Speed setting	**Platform**	**FPS 1 thread**	**FPS 2 threads**	**FPS 4 threads**
Ultrafast2	Screen	**163.66**	200.68	213.97

Scenario 3: Live Streaming vs x264

Threads	4/8
x264 presets	Medium, fast
x264 tune	Zerolatency
Speed settings tested	Veryfast, Superfast, Ultrafast
LD AV1 tune	Camera mode, single tile, zero lag
Keyframe interval	120 frames
Resolutions	1080p, 720p 30fps – 60fps
Bitrates	BDR: 1mb/s – 6mb/s Concurrency: 6mb/s
CPU	AMD Ryzen 9 7950X 16-Core Processor

Scenario 3: Live Streaming vs x264

Zero latency, multi-threading, 4 threads

Test set: 1080p60/720p of various characteristics

Settings	Overall PSNR	PSNR-Y	SSIM-AVG	SSIM-Y	VMAF	Encoding Speedup 100 * (FPS - FPS_{x264}) / FPS_{x264}
x264 medium vs AV1 veryfast	-48.74	-54.23	-45.79	-47.16	-51.01	37%
x264 medium vs AV1 superfast	-45.35	-51.13	-42.33	-43.60	-48.42	64%
x264 medium vs AV1 ultrafast	-39.82	-45.96	-36.30	-37.45	-44.35	104%

Scenario 3: Live Streaming vs x264

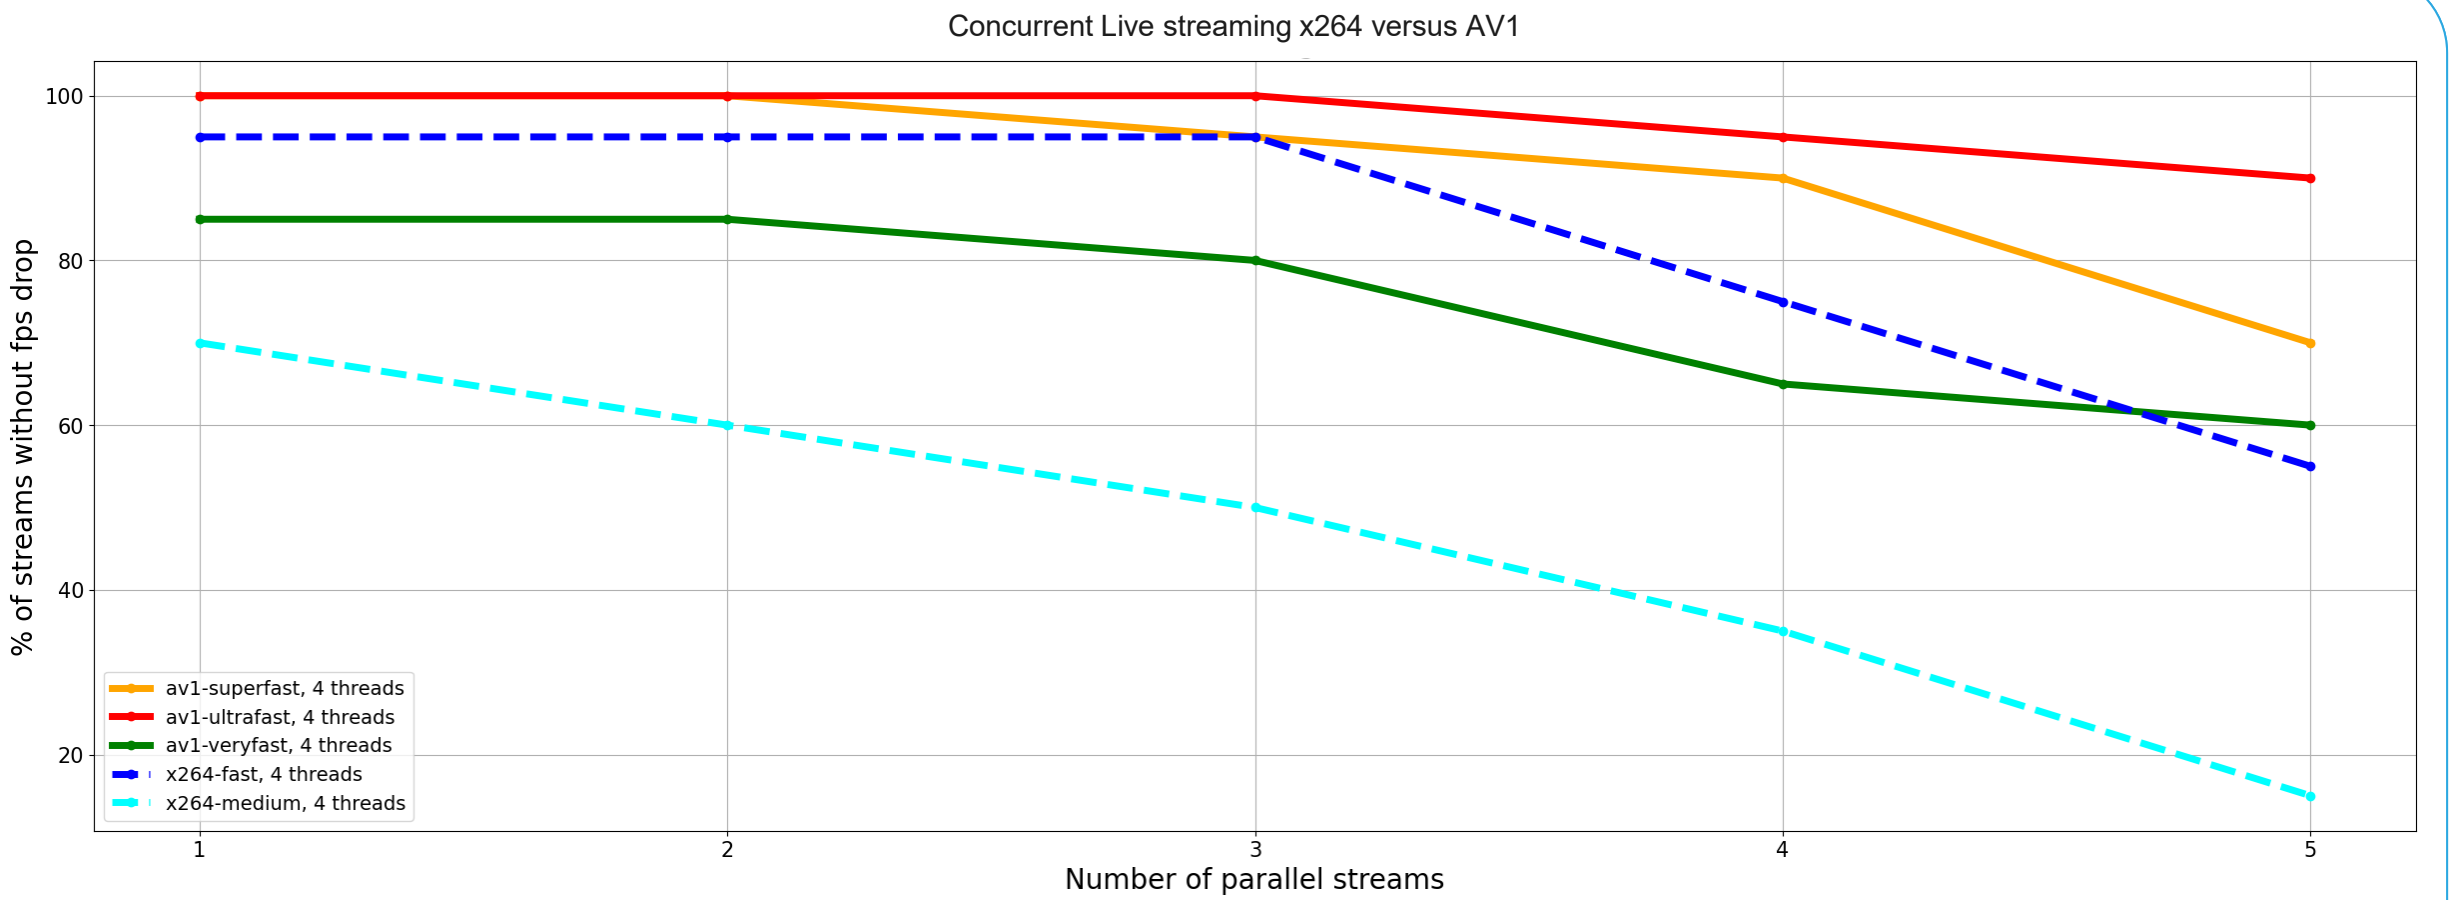

Takeaways

- LD AV1: a high performance low-latency AV1 encoder, featuring:
  - Advanced sub-frame parallelism
  - Intelligent pre-analysis, content analysis and mode decision
- For mobile, we have:
  - enabled real time AV1 encoding on mid-range iPhones
  - obtained huge quality gains over HW HEVC
  - with large speed gains versus libaom-rt
- For live streaming:
  - ~Double the speed for half the bitrate of x264 medium
  - Significant scaling of encoder density

info@visionular.com

[Silicon Valley] 280 2nd Street #C, Los Altos, CA 94022 USA

THANK YOU!

Thomas Davies thomas@visionular.com
Dmitriy Sychov dmitriy@visionular.com

Visionular