Optimal Video Encoding @ Scale: Dynamic Optimizer

Ioannis Katsavounidis
Video Algorithms,
Netflix

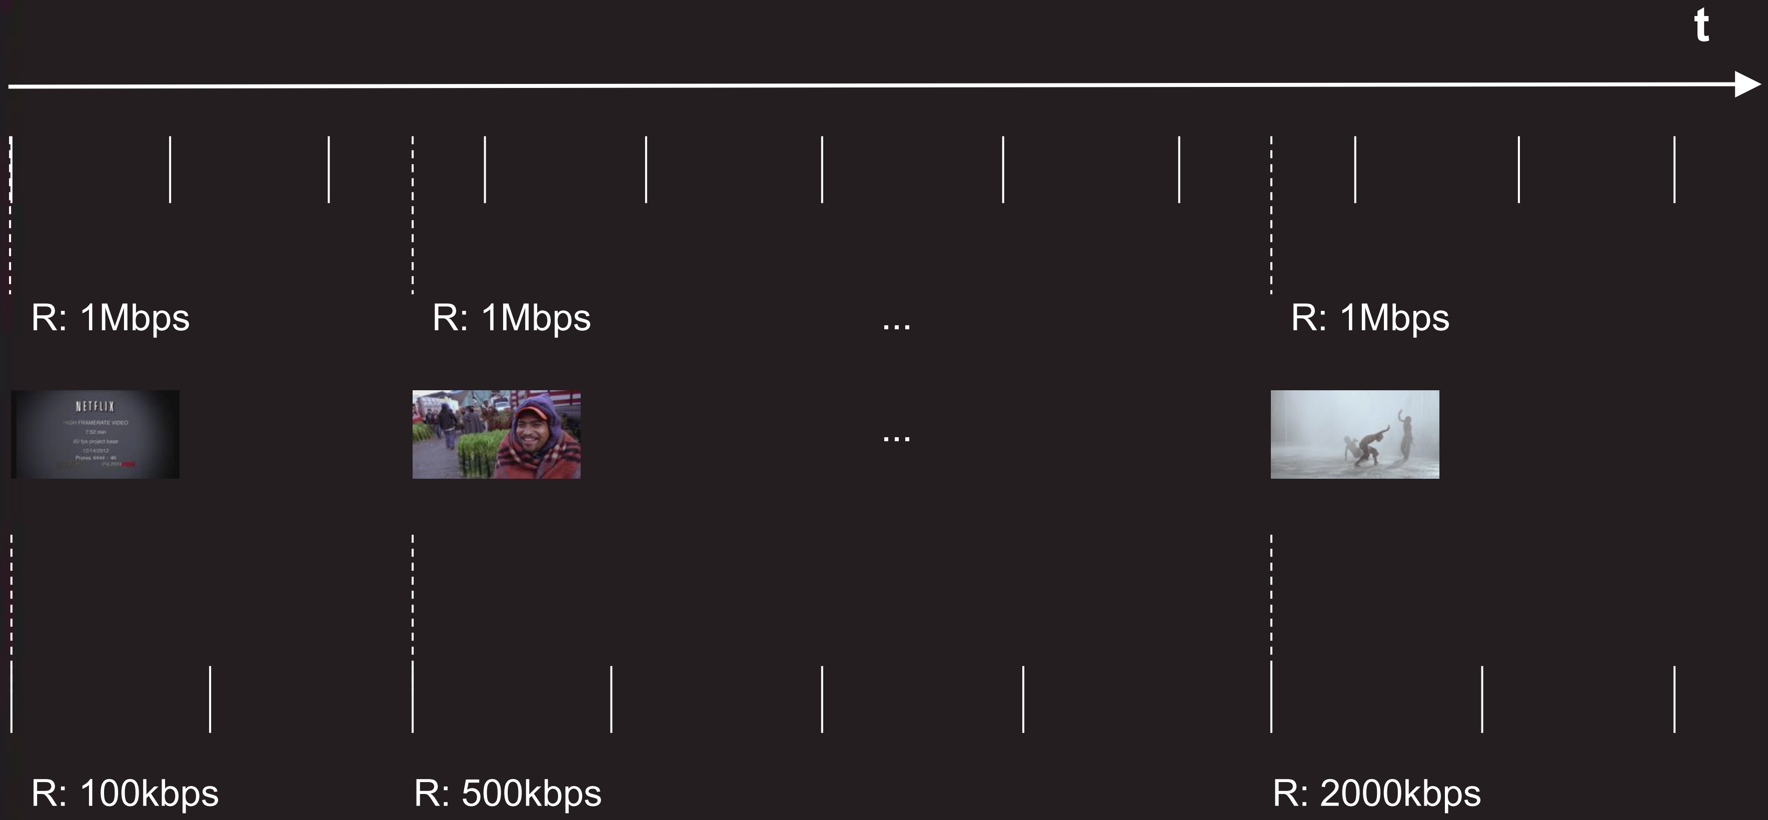
t
R: 1Mbps
R: 1Mbps
...
R: 1Mbps
NETFLIX
...
R: 100kbps
R: 500kbps
R: 2000kbps


BIENVENIDOS


BIENVENIDOS


BIENVENIDOS


BIENVENIDOS

Motivation

1. Why use same encoding parameters throughout a long and diverse video sequence?
2. Why impose a fixed Intra-frame interval?
3. Why consider (only) compression artifacts in video quality?
4. Why use MSE (PSNR)?
5. How to choose optimal combination of encoding parameters for a long and diverse video sequence?
6. How can we obtain the entire convex hull of optimal encodes?

Framework: Single shot processing

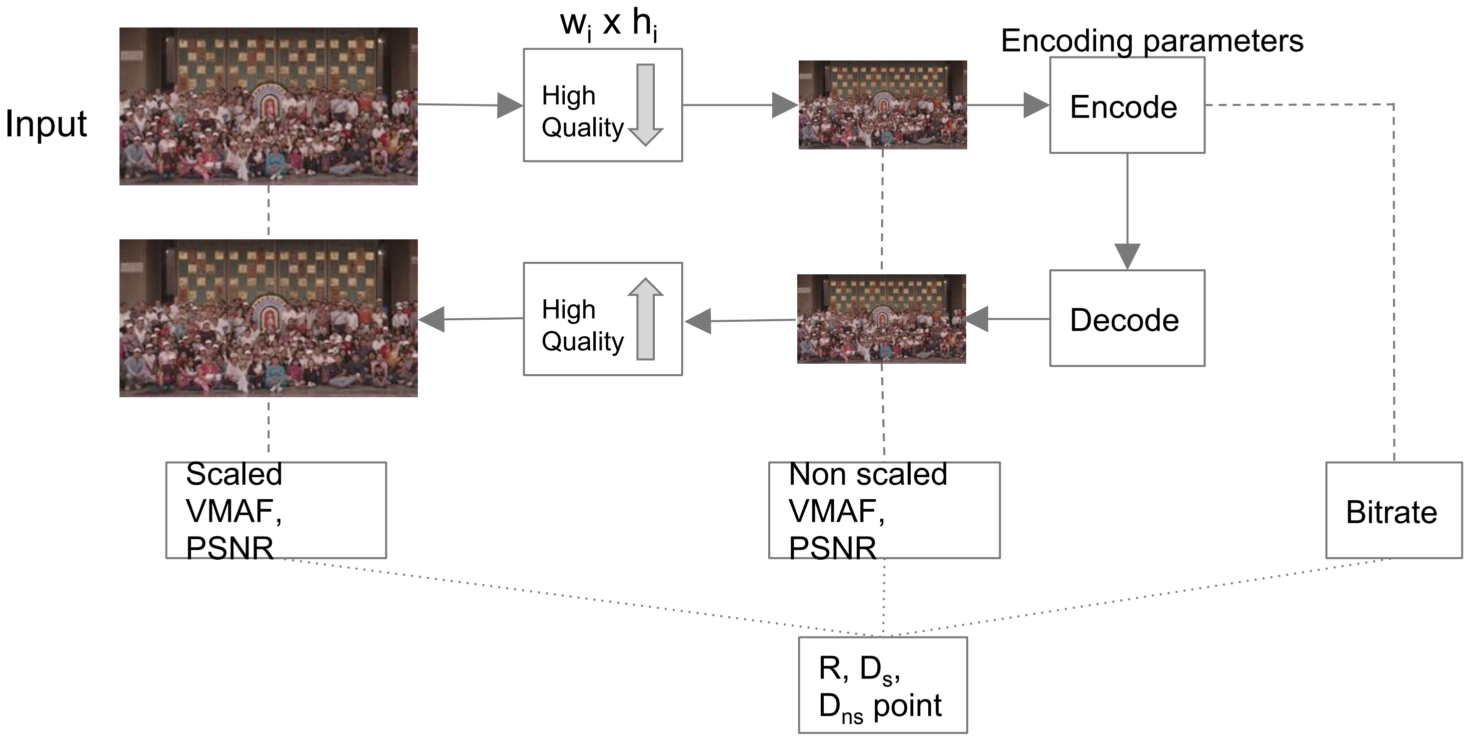

Framework: Convex hull of optimal shot encodes

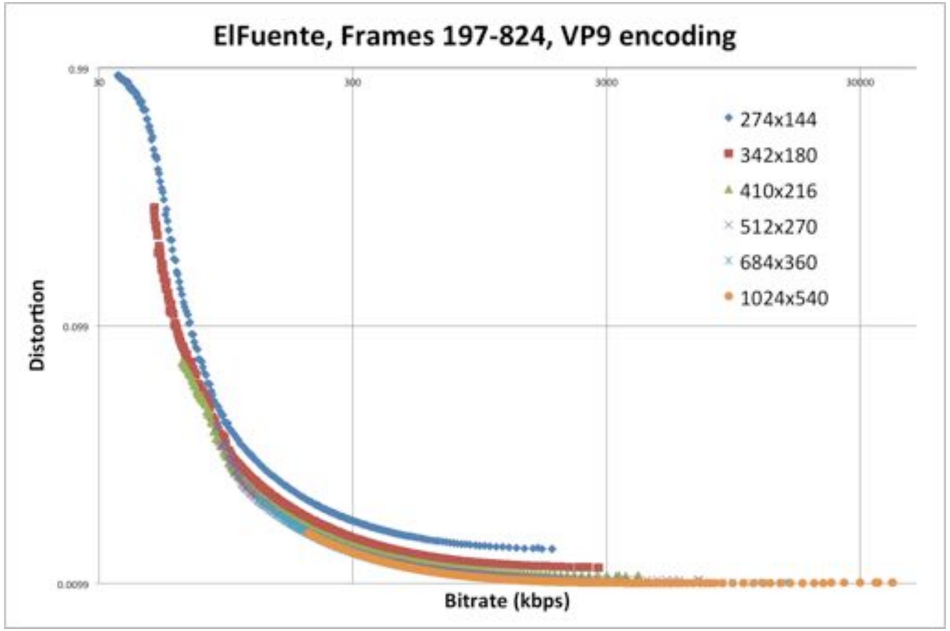

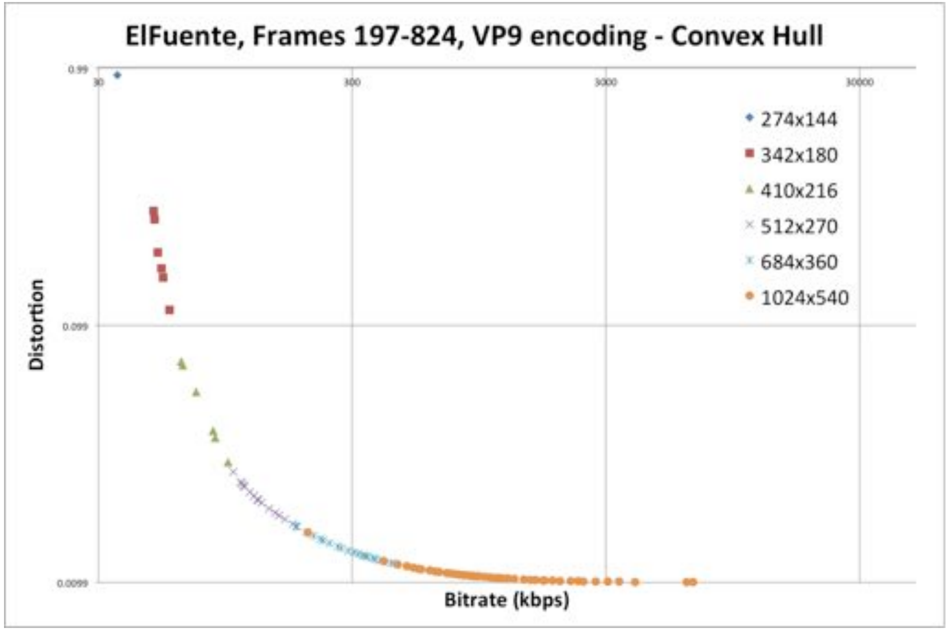

Framework: Trellis

Number of pre-encodes = number of QPs (n) x number of shots (m)

VP9, 63 QPs

Framework: Trellis optimal path

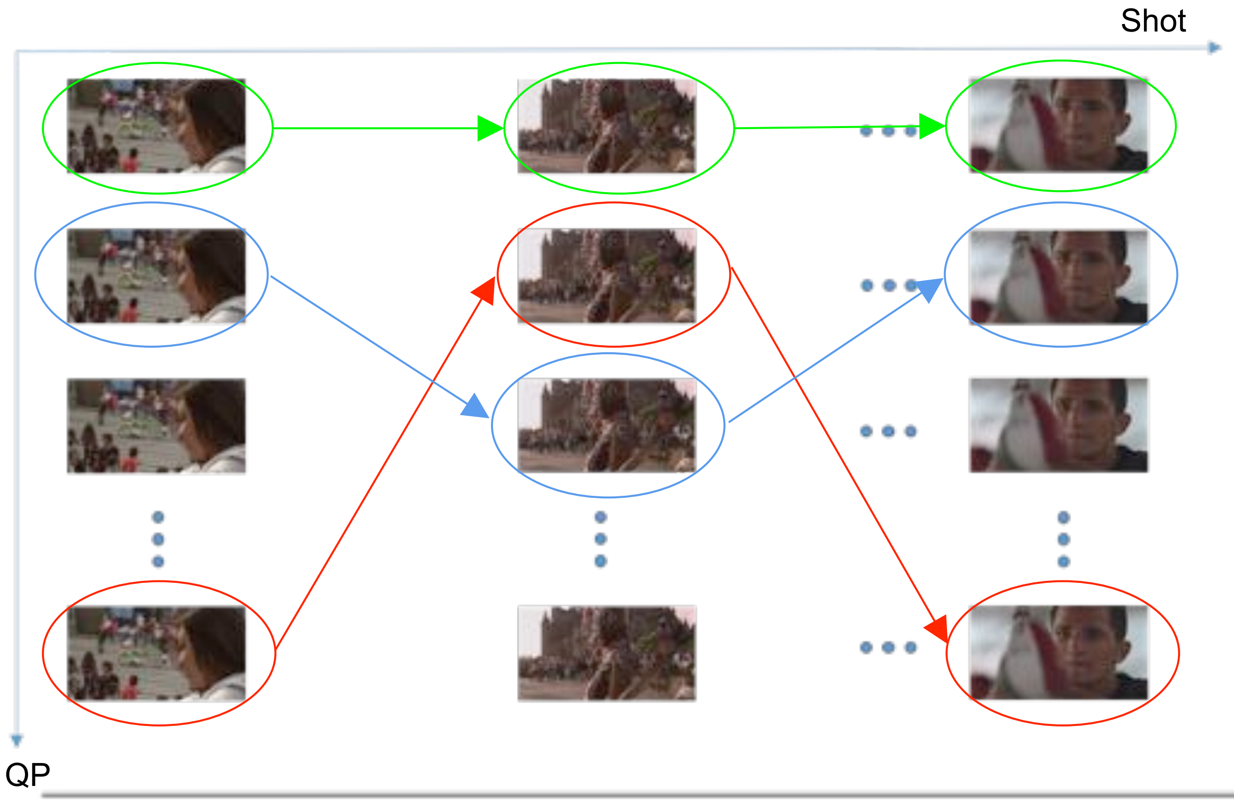

Fixed QP encode

Highest (average) quality encode, with bitrate x kbps

Lowest (average) bitrate encode, with quality y

Framework: Resolutions

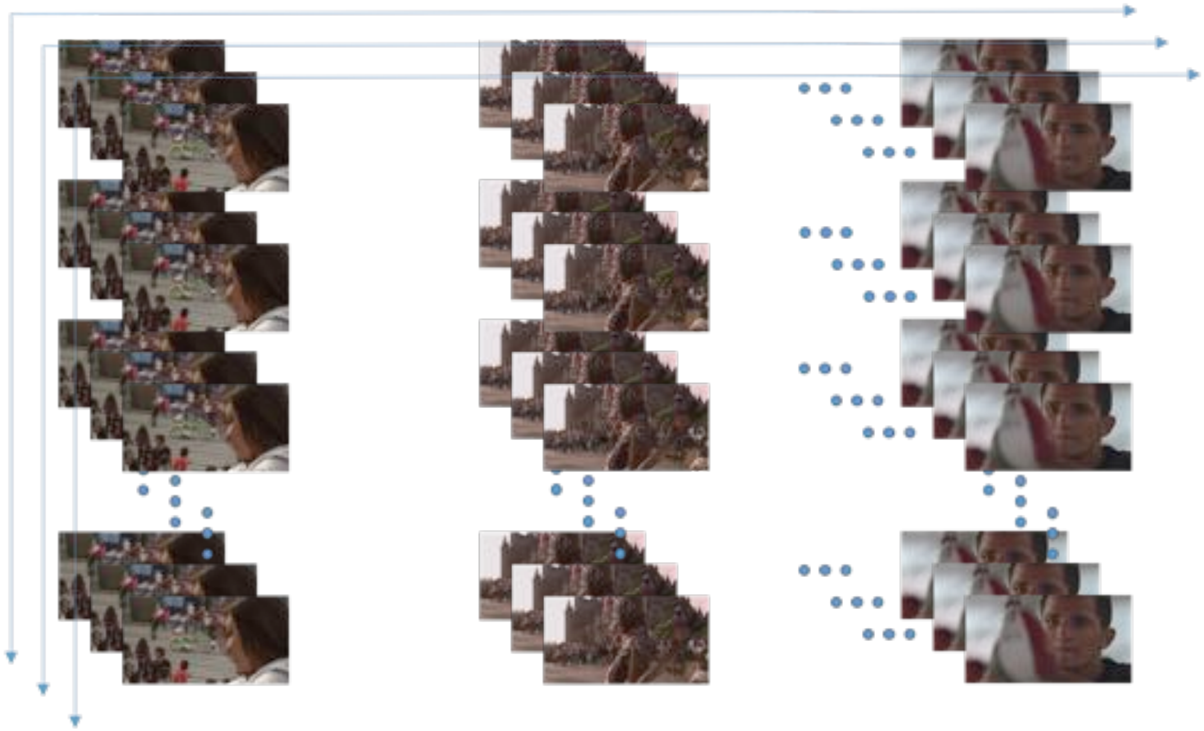

Number of pre-encodes = number of QPs (n) x number of shots (m) x number of resolutions (r)

VP9, 63 QPs, 7 resolutions

Results: Video Content

- **10 titles**
  - 8 representative titles from NETFLIX catalog
  - 2 publicly available (“El Fuente” and “Meridian”)


Breaking
Bad


NETFLIX
ORANGE is the new BLACK


NETFLIX
HOUSE of CARDS


MARVEL'S
AVENGERS
MARVEL

NETFLIX


NETFLIX
BLOODLINE


NETFLIX
WET HOT
AMERICAN SUMMER
FIRST DAY OF CAMP


NETFLIX
MARVEL
DAREDEVIL


NETFLIX
BoJACK
HORSEMAN
HOLLYWOO

Results: DO VP9 vs. Per-Title Optimal QP

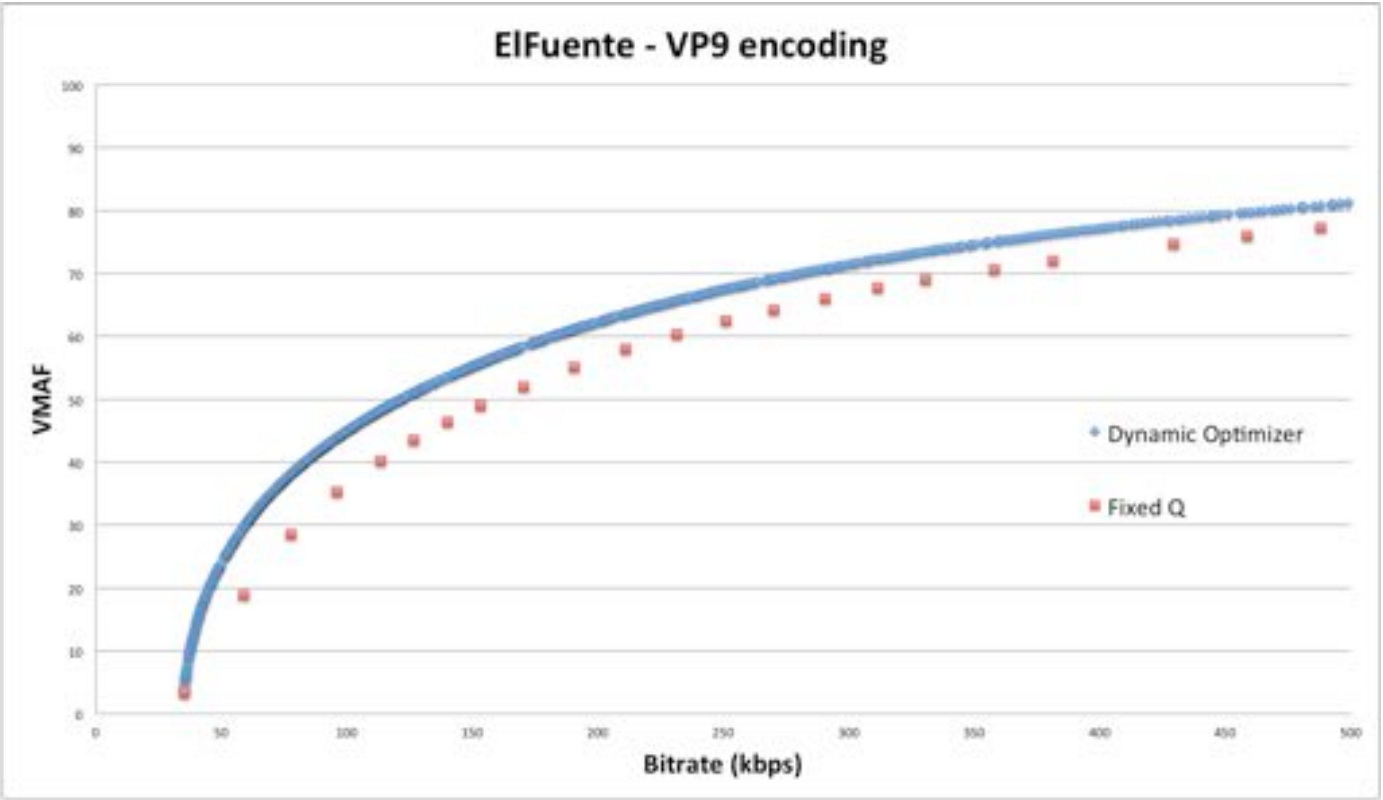

NETFLIX

Results: DO vs. Fixed Q BD-rate

Title	VMAF	PSNR
Bloodline	-18.88%	-31.26%
BoJack	-48.39%	-60.18%
BreakingBad	-22.81%	-38.10%
Daredevil	-20.07%	-34.42%
ElFuente	-25.93%	-45.58%
HoC	-25.85%	-34.52%
Meridian	-38.82%	-70.11%
OINB	-19.22%	-31.92%
TheAvengers	-25.61%	-37.32%
WHAS	-14.45%	-38.13%
AVERAGE	**-25.72%**	**-41.96%**

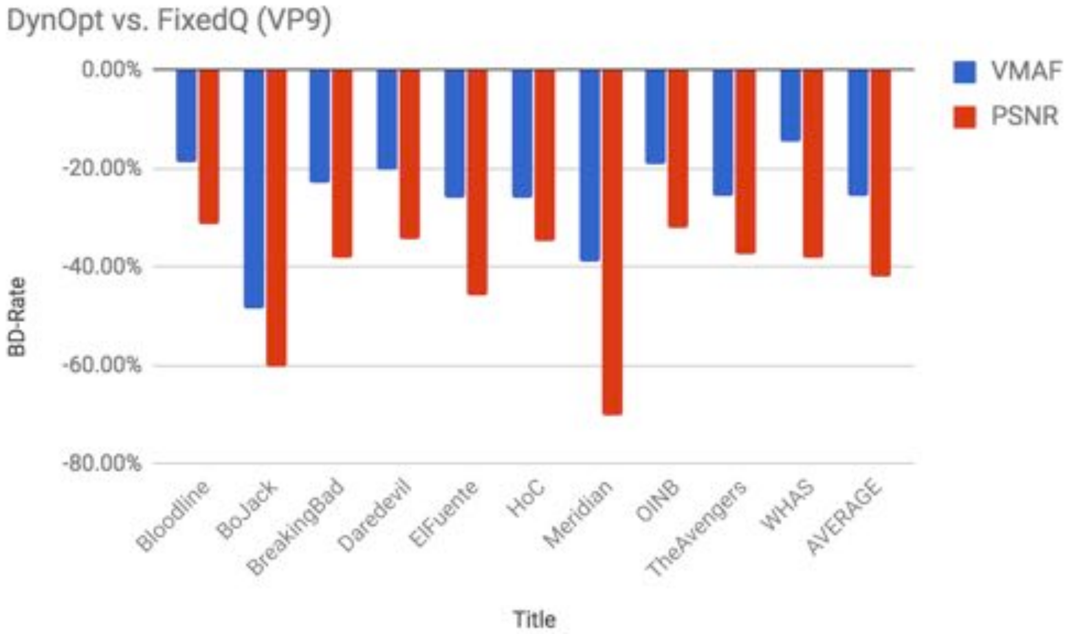

Results: Applied to AVC-High, VP9 and HEVC

- Run at scale (production pipeline)
- Multiple titles (30)
- About 10 min. each
- Reasonably slow speed settings

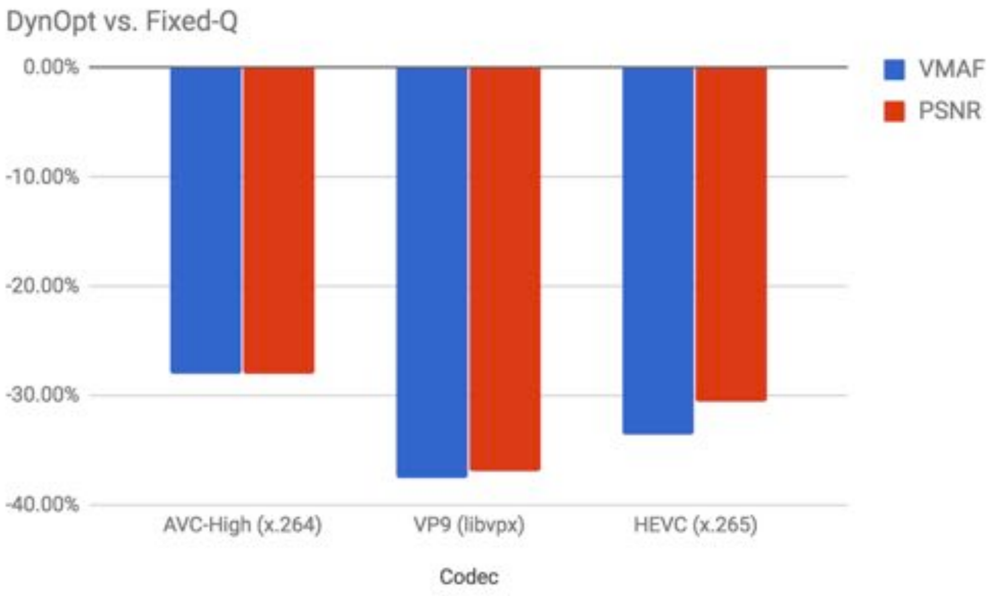

Codec	VMAF	PSNR
AVC-High (x264)	-28.04%	-27.99%
VP9 (libvpx)	-37.61%	-36.97%
HEVC (x265)	-33.51%	-30.52%

Parallel Encoding

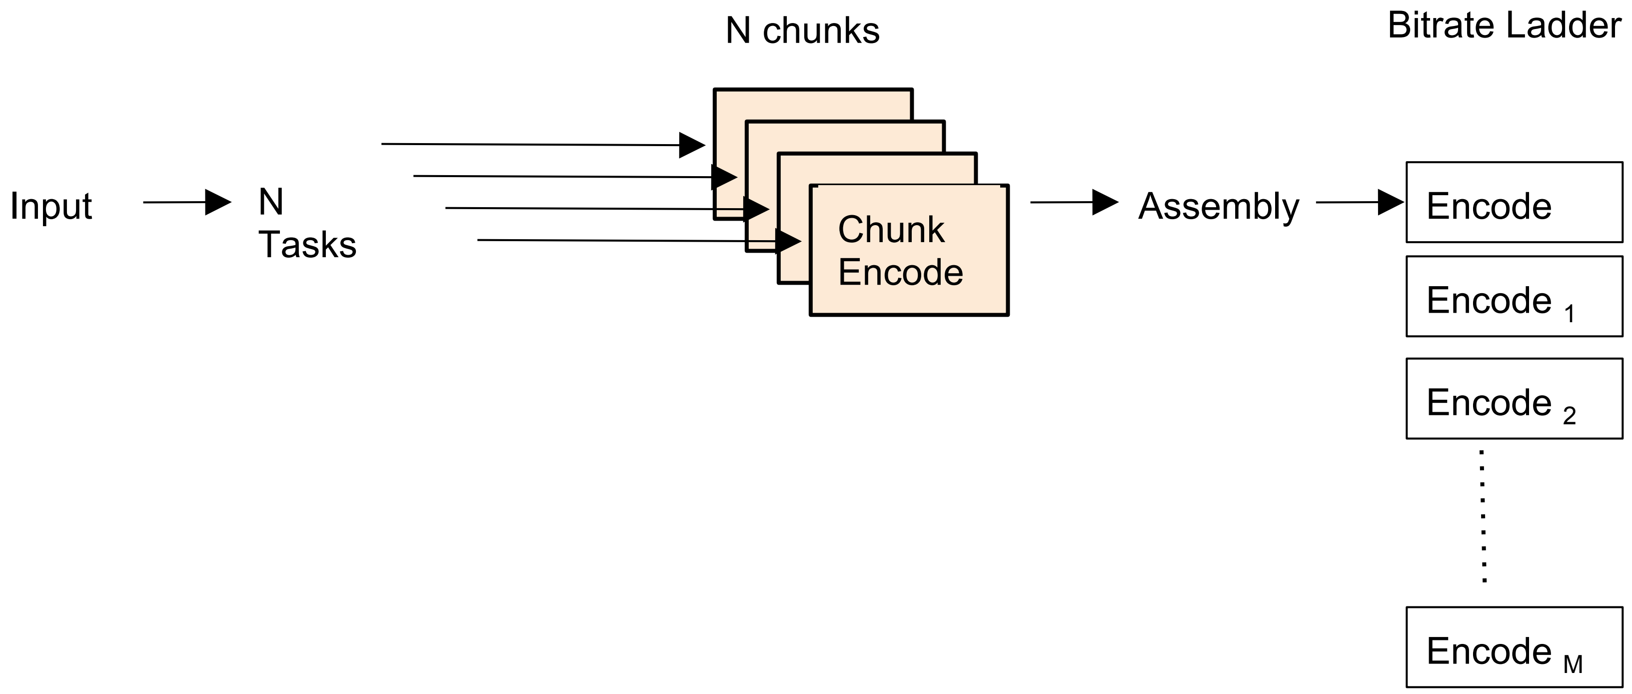

Tasks = chunks * encodes

Implementation challenges

	~3 min chunks	Encodes	Tasks
Avengers	48	10	480
El Fuente	2	10	20

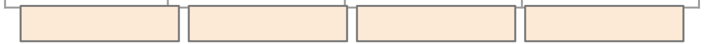

VP9, 63 QPs, 7 resolutions

	Shots	Encodes	Tasks
Avengers	2915	441	1,285,515
El Fuente	96	441	42,336

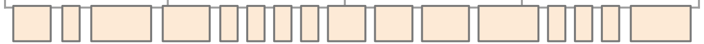

Number of actual encodes

Number of tasks

Tasks = chunks * encodes

Number of encodes

	Shots	Encodes	Tasks
Avengers	2915	441	1,285,515
El Fuente	96	441	42,336

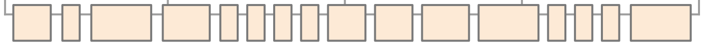

	Shots	Encodes	Tasks
Avengers	2915	35	102,025
El Fuente	96	35	3,360

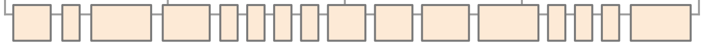

The point of diminishing returns

Find subset of operating points that produce “equivalent” performance to the “full” optimizer

Constrained dynamic optimizer

Tasks = chunks * encodes

Number of tasks

	Shots	Encodes	Tasks
Avengers	2915	35	102,025
El Fuente	96	35	3,360

	Chunks	Encodes	Tasks
Avengers	46	35	1610
El Fuente	2	35	70

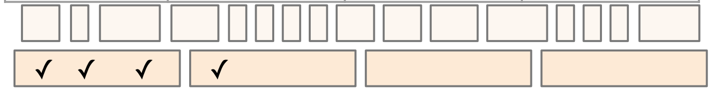

Collation

Combine multiple shots into chunks

Checkpoints after every shots

Tasks = chunks * encodes

Internal Spot market

- Borrow unused instances
- Daily peak of ~12,000 instances

115 movies, runtime between 2 - 3 hrs, numshots range from 725 to 3973, total shots 235,017

Codec	Total CPU time	Avg CPU time	Avg wall clock time
H264 AVC	15,281 days	132 days	7 days
VP9	38,284 days	332 days	9 days

So far...

Shots	CPU time
18,436,049	1,466,311 days

“Compute complexity is the currency we use to buy video quality”
David Ronca, Director of Encoding Technologies, VQEG meeting @NETFLIX, May 2017

Demo

Summary

- Joint optimization of shots
- Codec agnostic and object metric agnostic framework
- Orthogonal to I/P/B quality optimization by codecs
- Upper bound to compare rate control mechanisms within and between codecs
- Provides ~25% bitrate savings at same quality
- Streams are 100% compliant; ready to be consumed by existing clients

Video Algorithms Team @100M party

Academic research partners

University of Texas Austin

Prof. Al Bovik
Todd Goodall
Christos Bampis
Zeina Sinno

Université de Nantes

Prof. Patrick Le Callet
Lukáš Krasula

University of Southern California

Prof. C.-C. Jay Kuo
Joe Yuchieh Lin
Haiqiang Wang

University of Bristol

Prof. David Bull
Felix Mercer Moss
Mariana Afonso

Open problems

Perceptual quality metrics

- VMAF for HDR and 4K
- Temporal features
- Temporal pooling

Dynamic optimizer complexity reduction

Next-generation video codecs

- AV1 speed-up
- Integrating perceptual metrics

Next-generation image compression


Picture Coding Symposium 2018
June 24-27, San Francisco
Paper deadline: January 10, 2018

Special session: “System-level perceptual video coding optimization”