TAMPERE UNIVERSITY OF TECHNOLOGY

ULTRA VIDEO GROUP

Performance Evaluation of Kvazaar HEVC Intra Encoder on Xeon Phi Many-core Processor

Ari Koivula
Marko Viitanen
Ari Lemmetti
Dr. Jarno Vanne
Prof. Timo D. Hämäläinen

Tampere University of Technology
FINLAND

GlobalSIP 2015
Dec 16, 2015
Orlando, Florida, USA

Outline

- Kvazaar intro
  - Digest & statistics
  - Supported coding tools
- Motivation for this work
  - Kvazaar HEVC intra encoder
  - Xeon Phi many-core processor
  - Other state-of-the-art HEVC encoders
- Performance setup & analysis
  - Kvazaar vs. x265
  - On Xeon Phi & Core i7
- Conclusions

HEVC ENCODER

HEVC/H.265

- HEVC: High Efficiency Video Coding
- Developer: JCT-VC (VCEG + MPEG)
- Published as: ITU-T H.265 | ISO/IEC 23008-2
- 1st edition: Jan. 2013
- Three profiles: Main Profile
  Main 10 Profile
  Main Still Picture Profile
- Reference: HEVC Test Model (HM)

Kvazaar Digest

KVAZAAR HEVC ENCODER

- Main developer: TAMPERE UNIVERSITY OF TECHNOLOGY
- Source codes: github.com/ultravideo
- License: GNU LGPLv2.1
- Language: C/ASM
- OS: Windows/Linux/Mac
- Processors: x86, x64, PowerPC, ARM
- Compatibility: FFmpeg libav

#1 academic open-source HEVC encoder

Project Statistics Dec 2015

In a Nutshell, kvazaar HEVC encoder...

... has had 1,487 commits made by 15 contributors representing 19,524 lines of code

... is mostly written in C with a well-commented source code

... has a young, but established codebase maintained by a large development team with stable Y-O-Y commits

... took an estimated 5 years of effort (COCOMO model) starting with its first commit in May, 2012 ending with its most recent commit 7 days ago

Languages

C	88%	Assembly	7%
4 Other	5%		

Lines of Code

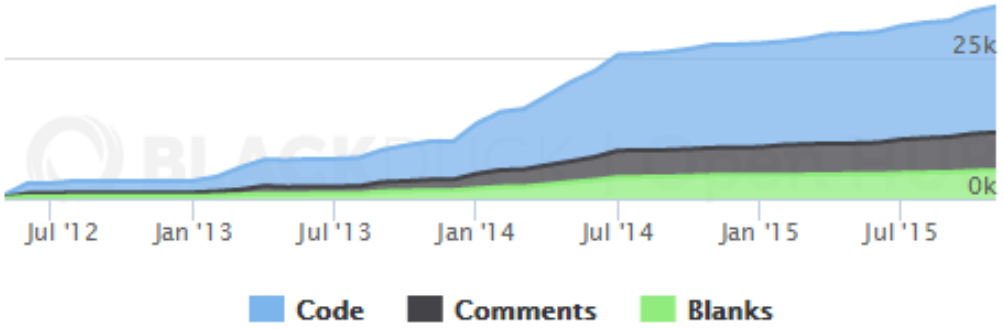

Visitors **(Nov 27 – Dec 10, 2015)**

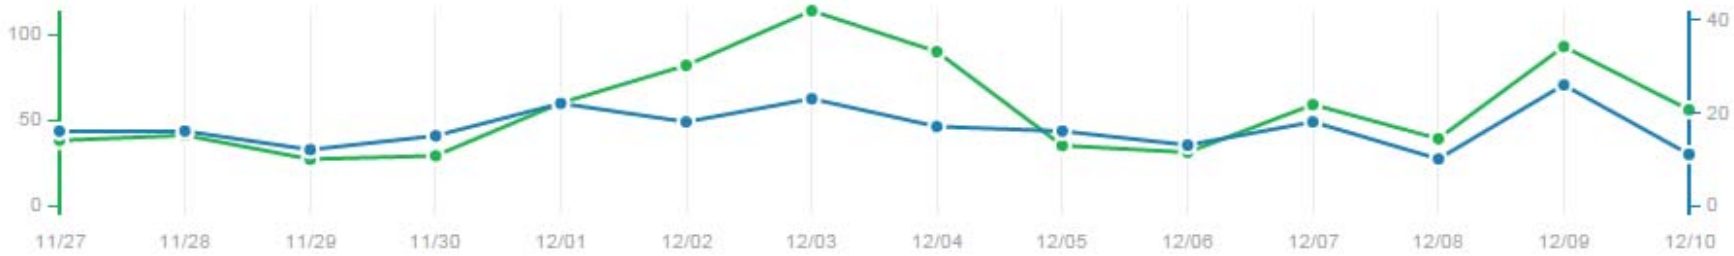

794	167
Views	Unique visitors

KVAZAAR HEVC ENCODER

Supported Coding Tools

Feature	Supported tools	Tools to be added
Profiles	Main, Main 10, Main Still Picture	
Input data format	YUV 4:2:0	4:2:2
YUV dimensions	Divisible by 2	
Output data format	HEVC bytestream, Annex B	
Picture hash	checksum	
Internal bit depth	8, 10	
Color format	4:2:0	4:2:2
Interlaced coding	Yes	
PCM coding	No	
Lossless coding mode	No	Yes
Coding configurations	AI, LP, RA	LB
Slice types	I, P, B	
Tiles	Yes	
Wavefront processing	Yes	
Slices	No	Yes
Sizes of CUs	64×64, 32×32, 16×16, 8×8	
Sizes of TUs	32×32, 16×16, 8×8, 4×4	
Sizes of PUs, Intra	32×32, 16×16, 8×8, 4×4	
Sizes of PUs	Square	Symmetric, Asymmetric
Intra prediction modes	35 (DC, planar, 33 angular), Chroma from luma	
Luma MV accuracy	1/4 pel	
Chroma MV accuracy	1/8 pel	
Coding modes	Intra, Inter, Skip, Merge	
IME algorithm	HEXB, TZ	
Search range	Not limited	
# of reference pictures	Up to 15	
IME metric	SAD	
FME metric	SATD	
Mode decision metrics	SATD, SAD, SSD	
RDO	Partial (includes RDOQ)	
Entropy coding	CABAC	
Loop filtering	DF, SAO (Edge + Band mode)	
Residual coding	Coefficients, 4x4 transform skip	Transquant bypass
GOP	Constant QP, B-frame pyramid	LP, configurable


Motivation

Research Questions

How well Kvazaar scales to many-core processors?

What is the maximum coding speed of Kvazaar on Xeon Phi?

What is the performance of Kvazaar over other state-of-the-art encoders?

Kvazaar Implementation Status

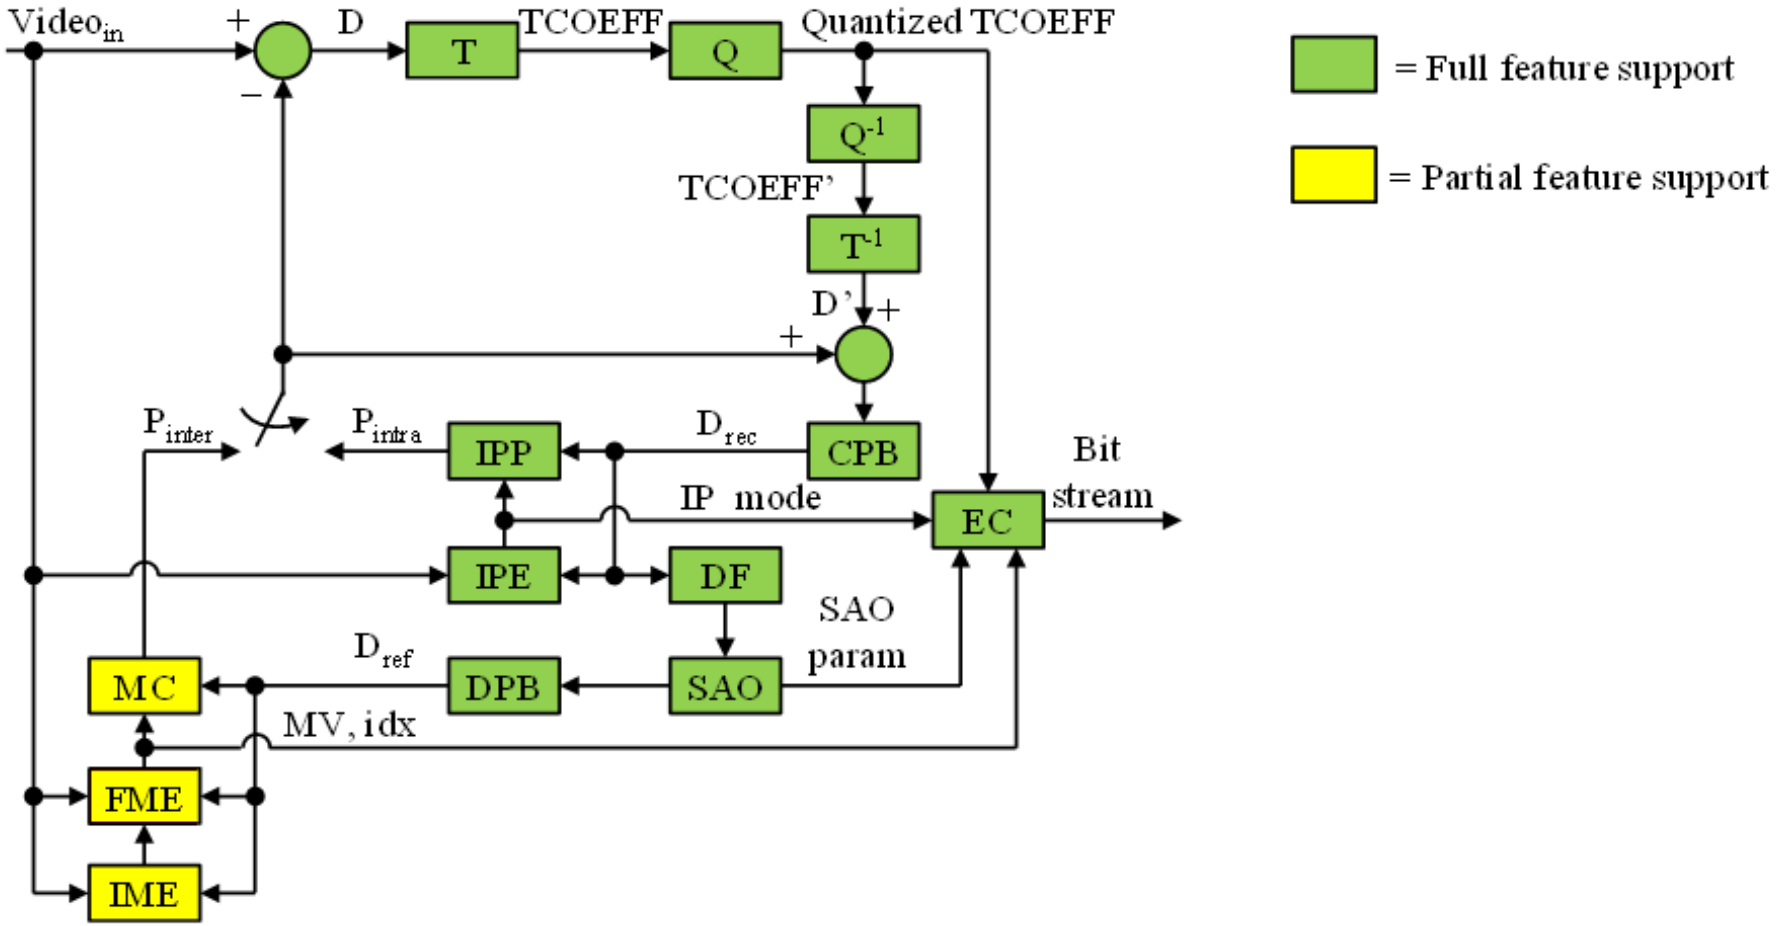

Kvazaar intra encoder considered

Kvazaar Intra Encoder

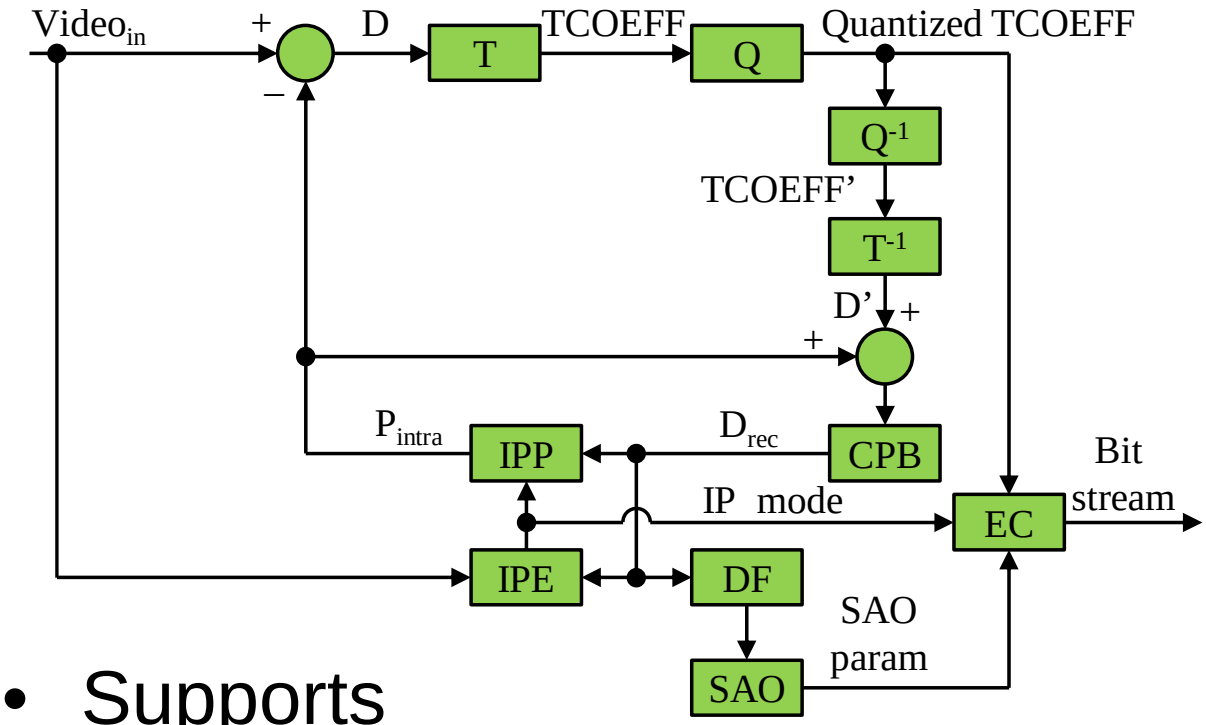

- Supports
  - Complete HEVC block partitioning structure
  - Practically all intra coding tools

Functionality almost identical to HM

Intel Xeon Phi Platform

Family	First-generation Xeon Phi (Knights Corner)
Announced	2011
Process	22 nm
Processor	Intel Xeon Phi 7120X
# of cores	61 (244 logical)
Frequency	1.2 GHz
L1 cache	61 × 32 KB (instruction) + 61 × 32 KB (data)
L2 cache	61 × 512 KB
Memory	16 GB
Compiler	Intel C++ Compiler 2013.3.163
Operating system	64-bit Linux 2.6 µOS with busybox

Other HEVC Encoders

- Many commercial encoders
  - Source codes/operating principles kept confidential
- Active open-source projects:

Encoder	Coordinator	License	Code	Parallelization		
				WPP	Tiles	Picture-level
HM	JCT-VC	BSD	C++	No	No	No
x265	MulticoreWare	GPLv2.0	C++/ASM	Yes	No	Yes
Kvazaar	TUT	LGPLv2.1	C/ASM	Yes	Yes	Yes

Kvazaar evaluated against x265

ANALYSIS SETUP

Encoder Parameters

Encoder	Version	Code*	Compiler	Presets		Parallelization
				High-quality	High-speed	
Kvazaar	0.5.0	C	Intel C	RD2	RD1	WPP, Tiles
x265	1.6	C++	Intel C	veryslow	ultrafast	WPP

* Kvazaar / x265 ASM code not supported on Xeon Phi

Encoder	Parameters
Kvazaar	--input-res=(res) -n (frames) --cpuid=0 -p 1 -q (qp) --owf=31 --rd=(preset) --threads=(threads) --no-transform-skip
wpp	--wpp
tiles	--tiles-width-split=u(x tiles) --tiles-height-split=u(y tiles)
x265	--tune psnr --psnr --hash 3 --log-level debug --no-progress --frame-threads 16 -I 1 --no-open-gop --no-scenecut --rc-lookahead 0 --lookahead-slices 0 --bframes 0 --fps 30 --pmode --pool (threads) --preset (preset) --frames (frames) -q (qp) --ipratio 1 --input-res (res)

Test Sequences

3840×2160 (2160p)

Beauty

Honeybee

Jockey

2560×1600 (1600p)

Traffic

PeopleOnStreet

1920×1080 (1080p)

Kimono

ParkScene

BasketballDrive


Results

Kvazaar vs. x265: 2160p

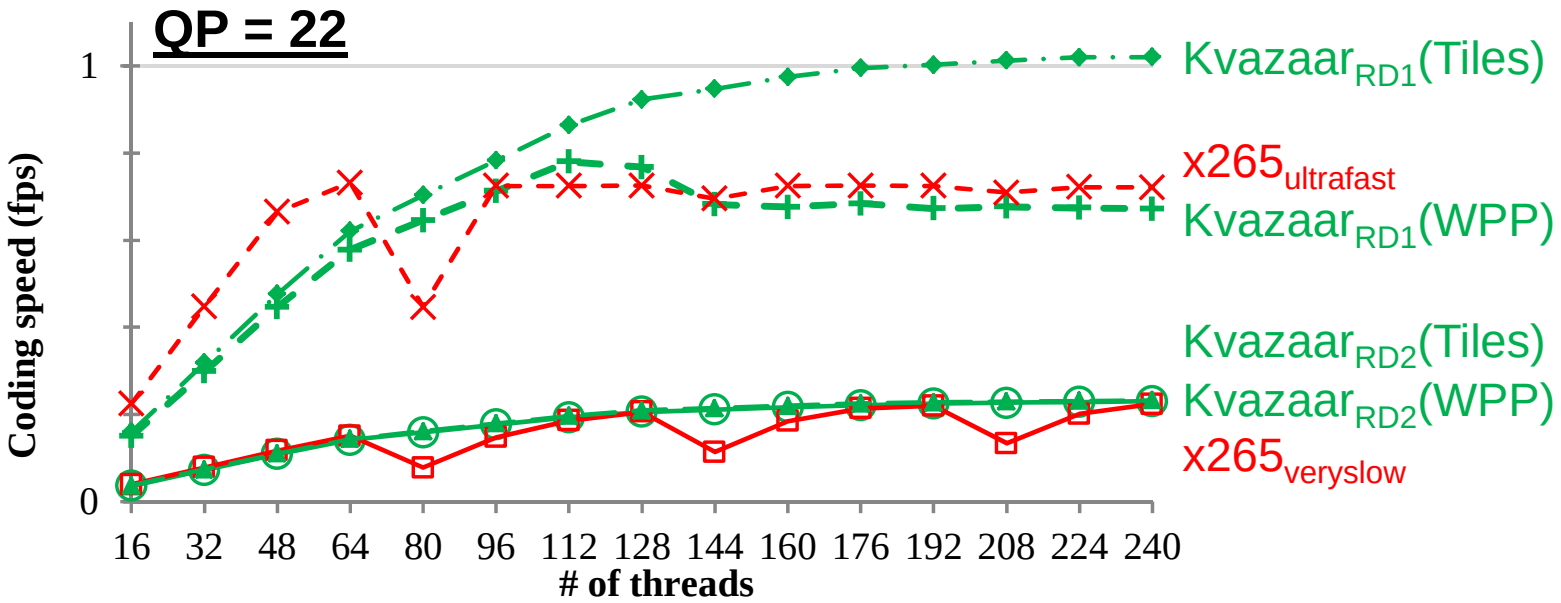

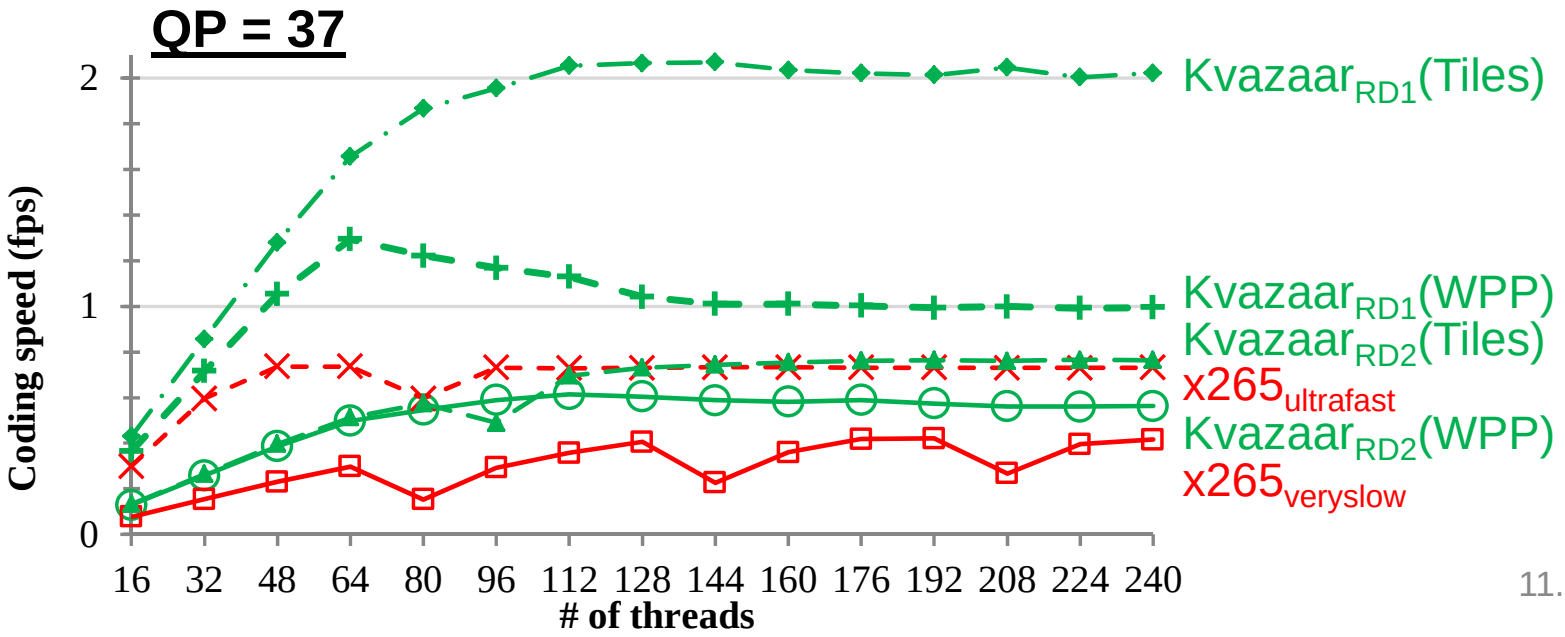

Kvazaar vs. x265: 1080p

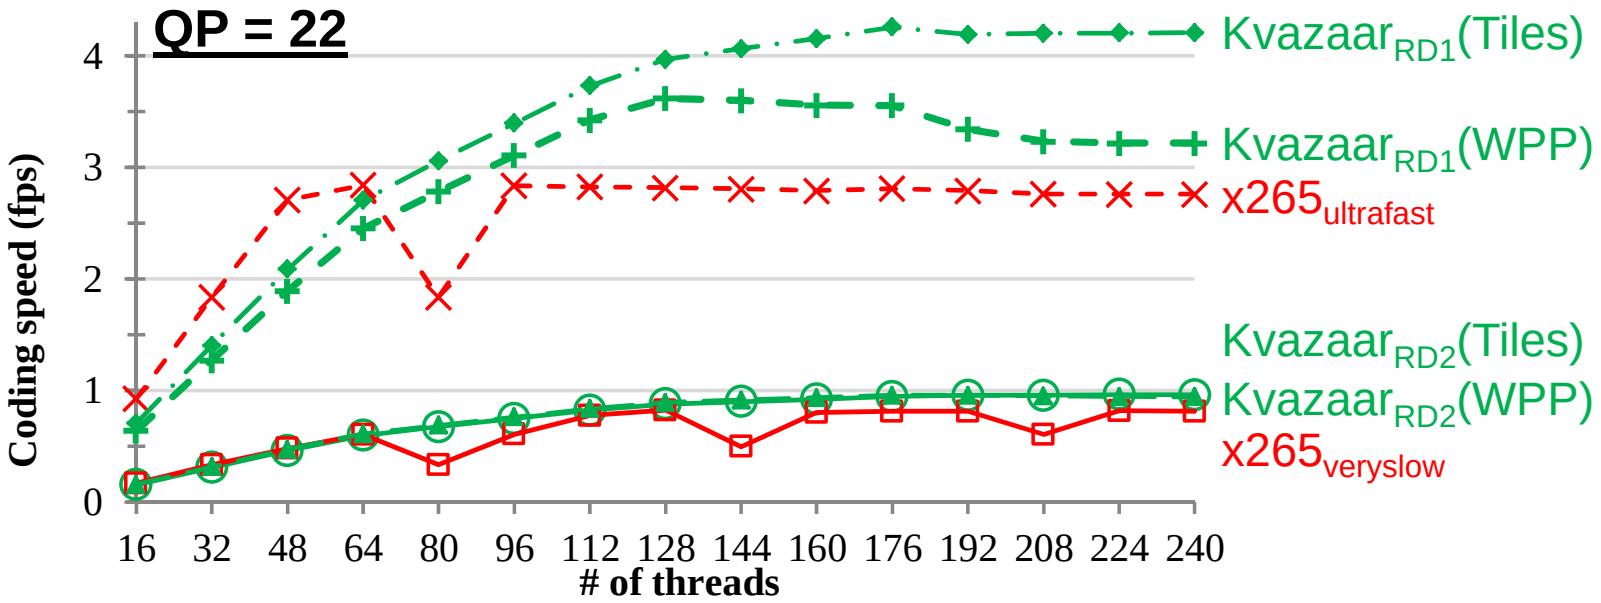

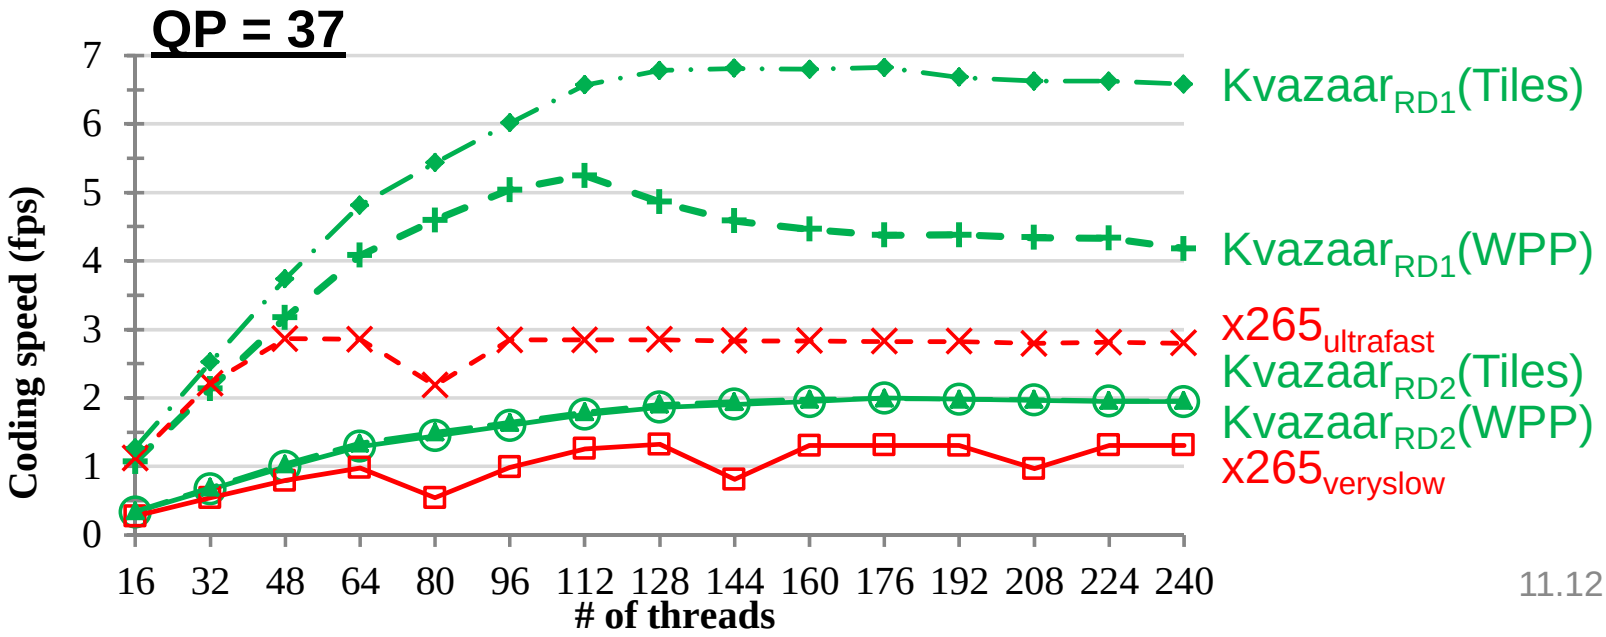

What Makes Kvazaar Better?

Kvazaar combines three design approaches:

1) Individual algorithm optimizations
2) Efficient threading model built on right-sized work units
3) C implementation that minimizes heap memory allocation

Kvazaar's Edge over x265

High-quality presets

Format	WPP		Tiles	
	Speedup	BD-rate	Speedup	BD-rate
2160p	1.3×	-1.8%	1.6×	-0.6%
1600p	1.2×	0.1%	1.2×	0.9%
1080p	1.4×	-0.5%	1.4×	0.5%
AVERAGE	**1.3×**	-0.7%	**1.4×**	0.3%

High-speed presets

Format	WPP		Tiles	
	Speedup	BD-rate	Speedup	BD-rate
2160p	1.3×	-8.7%	2.1×	-7.9%
1600p	1.3×	-13.5%	1.8×	-12.8%
1080p	1.5×	-11.4%	1.9×	-10.2%
AVERAGE	**1.4×**	-11.2%	**1.9×**	-10.3%

Kvazaar's Edge over x265

High-quality presets

Format	WPP		Tiles	
	Speedup	BD-rate	Speedup	BD-rate
2160p	1.3×	-1.8%	1.6×	-0.6%
1600p	1.2×	0.1%	1.2×	0.9%
1080p	1.4×	-0.5%	1.4×	0.5%
AVERAGE	**1.3×**	-0.7%	**1.4×**	0.3%

High-speed presets

Format	WPP		Tiles	
	Speedup	BD-rate	Speedup	BD-rate
2160p	1.3×	-8.7%	2.1×	-7.9%
1600p	1.3×	-13.5%	1.8×	-12.8%
1080p	1.5×	-11.4%	1.9×	-10.2%
AVERAGE	**1.4×**	-11.2%	**1.9×**	-10.3%

Recommended Kvazaar Presets

High-quality presets

Format	WPP	
	Speedup	BD-rate
2160p	1.3×	-1.8%
1600p	1.2×	0.1%
1080p	1.4×	-0.5%
AVERAGE	**1.3×**	-0.7%

High-speed presets

Format	Tiles	
	Speedup	BD-rate
2160p	2.1×	-7.9%
1600p	1.8×	-12.8%
1080p	1.9×	-10.2%
AVERAGE	**1.9×**	-10.3%

Kvazaar Speed: Xeon Phi

High-quality presets

Format	WPP		Xeon Phi
	Speedup	BD-rate	Speed
2160p	1.3×	-1.8%	0.4 fps
1600p	1.2×	0.1%	0.7 fps
1080p	1.4×	-0.5%	1.5 fps
AVERAGE	**1.3×**	-0.7%	

High-speed presets

Format	Tiles		Xeon Phi
	Speedup	BD-rate	Speed
2160p	2.1×	-7.9%	1.5 fps
1600p	1.8×	-12.8%	2.6 fps
1080p	1.9×	-10.2%	5.5 fps
AVERAGE	**1.9×**	-10.3%	

Kvazaar Speed: Xeon Phi vs. i7

13-30%

High-quality presets

Format	WPP		Xeon Phi		Core i7
	Speedup	BD-rate	Speed		Speed
2160p	1.3×	-1.8%	0.4 fps	<	0.5 fps
1600p	1.2×	0.1%	0.7 fps	<	0.8 fps
1080p	1.4×	-0.5%	1.5 fps	<	1.7 fps
AVERAGE	**1.3×**	-0.7%			

19-35%

High-speed presets

Format	Tiles		Xeon Phi		Core i7
	Speedup	BD-rate	Speed		Speed
2160p	2.1×	-7.9%	1.5 fps	<	2.1 fps
1600p	1.8×	-12.8%	2.6 fps	<	3.1 fps
1080p	1.9×	-10.2%	5.5 fps	<	6.6 fps
AVERAGE	**1.9×**	-10.3%			

Conclusions

- Kvazaar intra coding speed on Xeon Phi << real-time
  - Core i7 13-35% faster

→ Xeon Phi not an optimal platform for Kvazaar

- Kvazaar scales almost linearly to the number of cores

→ Real-time speed on more efficient many-core processor

- Recommendations for Kvazaar parallelization:

Preset	Parallelization	BD-rate over x265	Speed over x265
High-quality	WPP	-0.7%	1.3×
High-speed	Tiles	-10.3%	1.9×

CONTACT INFORMATION

Dr. Jarno Vanne

jarno.vanne@tut.fi

+358 40 576 3497

ULTRA VIDEO GROUP

ultravideo.cs.tut.fi
ultravideo@cs.tut.fi