AMD's high-end Family 15h (Bulldozer-based) processor lines

Dezső Sima

October 2018

(Ver. 2.0)

Contents

1. Overview of AMD's Family 15h processor lines based on high performance oriented Bulldozer modules

1. Overview of AMD's Family 15h lines, based on high-performance oriented Bulldozer modules

AMD's Family 15h processor lines are based on the Bulldozer microarchitecture and include four generations, as follows:

AMD's Family 15h processor lines, based on high-performance oriented Bulldozer modules

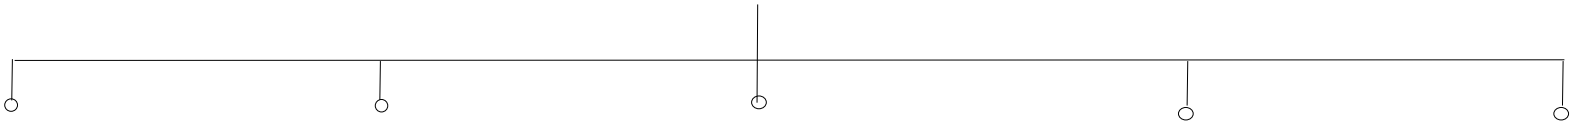

AMD's Family 15h Models 00h-0Fh	**AMD's Family 15h Models 10h-1Fh**	**AMD's Family 15h Models 30h-3Fh**	**AMD's Family 15h Models 60h-6Fh**	**AMD's Family 15h Models 70h-7Fh**
Bulldozer family (2011)	**Piledriver family** (2012)	**Steamroller family** (2014)	**Excavator family v1** (2015)	**Excavator family v2** (2016)
1. Generation Bulldozer family	2. Generation Bulldozer family	3. Generation Bulldozer family	4. Generation Bulldozer family	4. Generation Bulldozer family Refresh
Chapter 2	*Chapter 3*	*Chapter 4*	*Chapter 5 Sections 5.1-5.3*	*Chapter 5 Sections 5.4-5.5*

Four-generations of AMD's Bulldozer architecture [2]

Published: February 2011

Note: This roadmap does not show any figures for the performance increase.

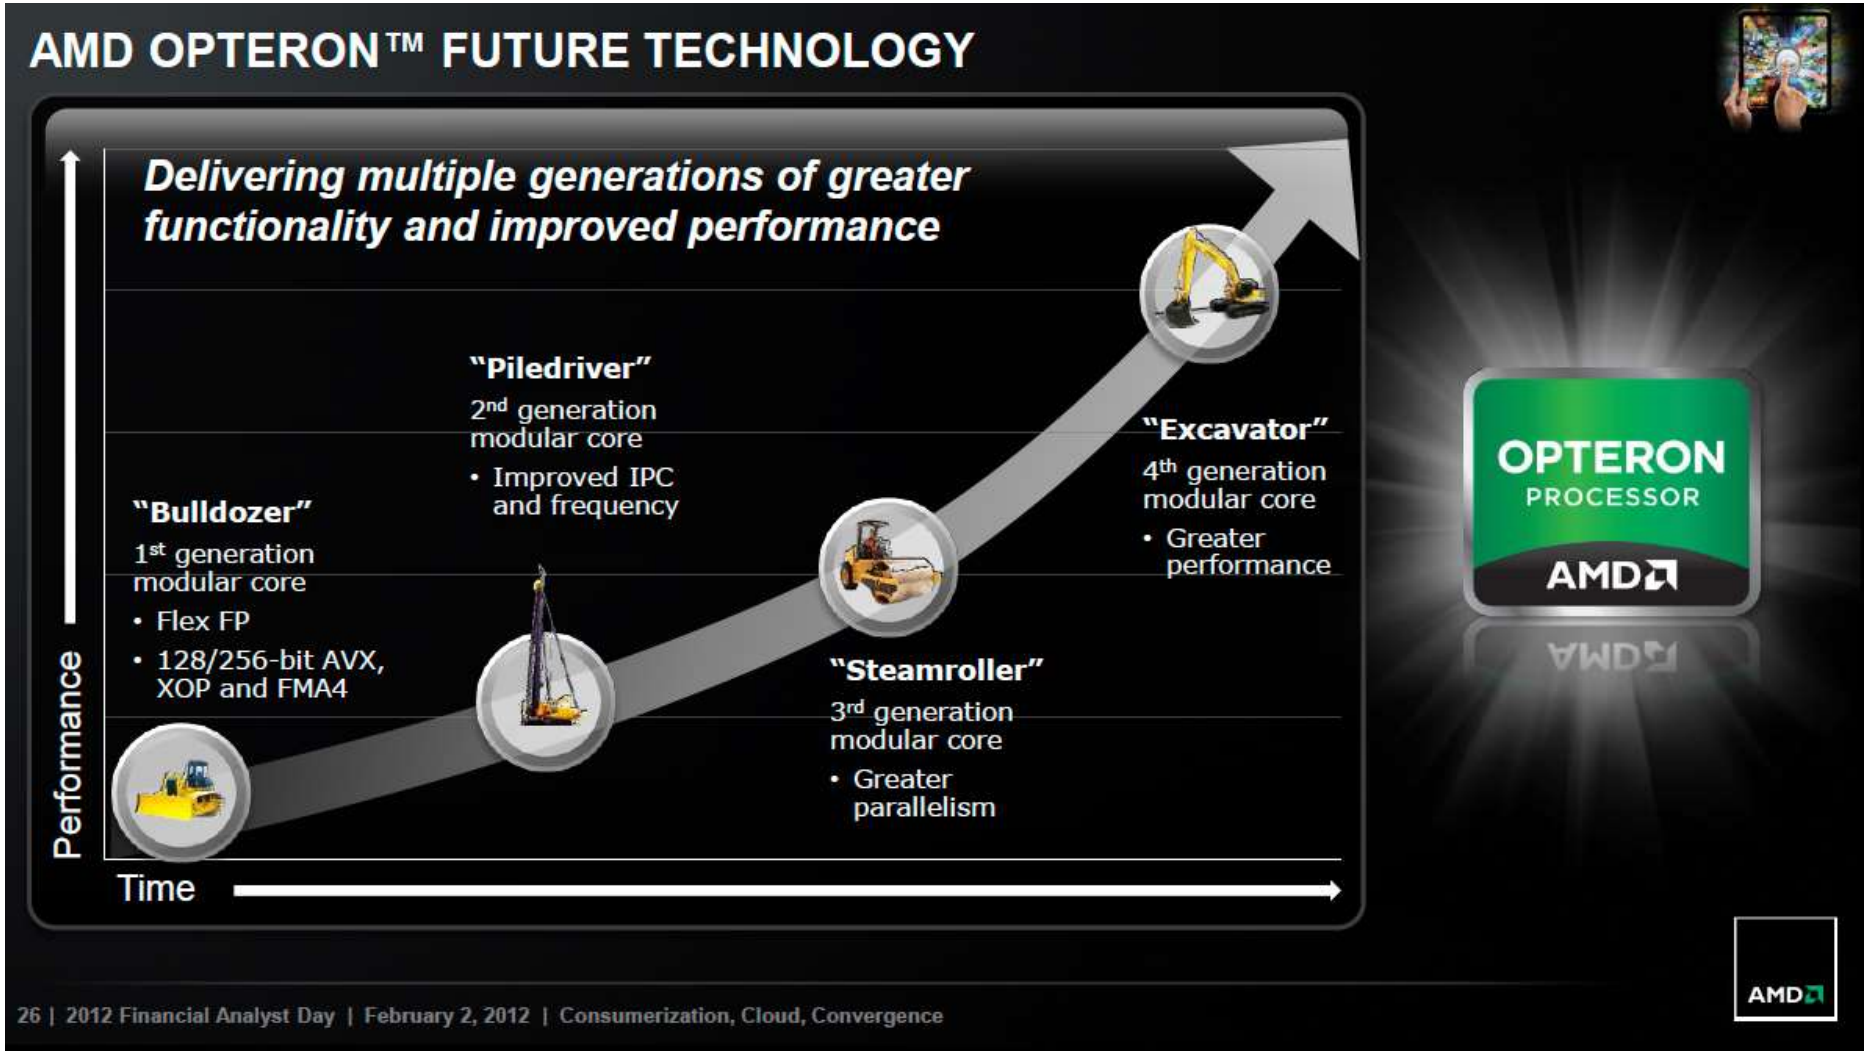

1. Overview of AMD's Family 15h lines, based on high-perf. Bulldozer modules (3)

Note to the terminology

In the literature and also in these slides the term Bulldozer is used in two disrtinct interpretations;

a) The Bulldozer designation typically refers to the 1. generation Family 15h processor lines.

b) In addition, the Bulldozer designation is also used to refer to the whole set of performance oriented Family 15h processor lines.

According to this interpratation we designate

- the Piledriver processor lines as the 2. generation Bulldozer lines and
- the Steamroller processor lines as the 3. generation Bulldozer lines.

Usually, the context clarifies which interpretation fits.

AMD's projection to increase performance in post Bulldozer architectures [19]

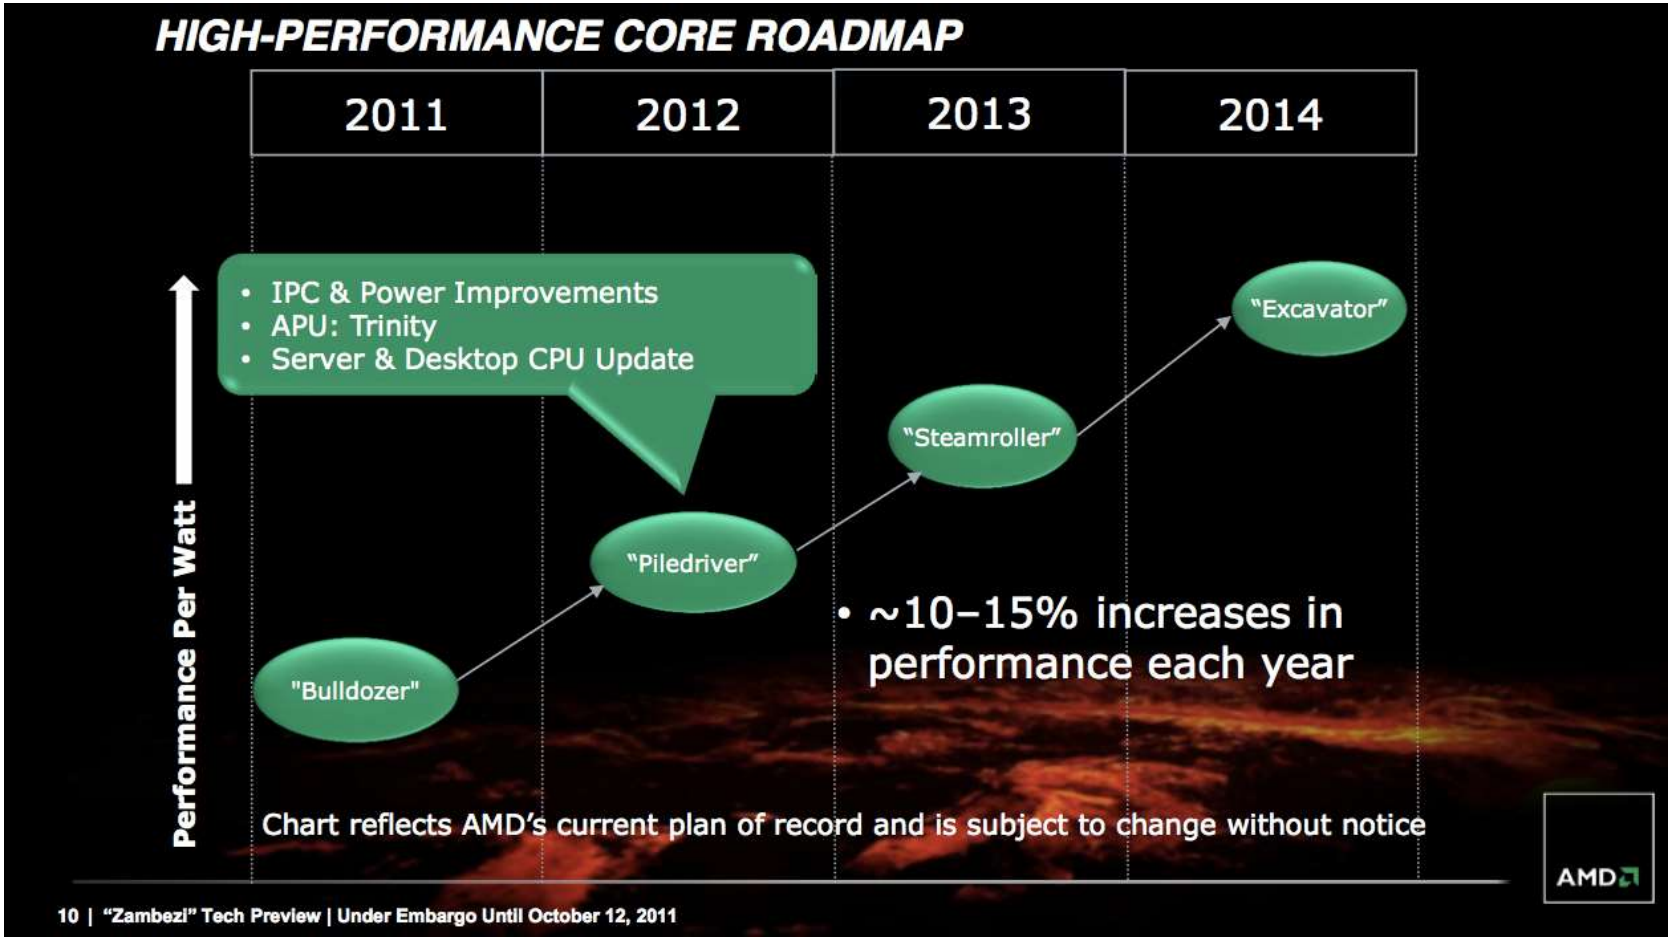

With the above slow rate of performance increase it is strongly questionable whether AMD will able to catch up ever with Intel's future processor lines.

1. Overview of AMD's Family 15h lines, based on high-perf. Bulldozer modules (5)

Performance increase of AMD's DP servers up to the Interlagos server lines [18]

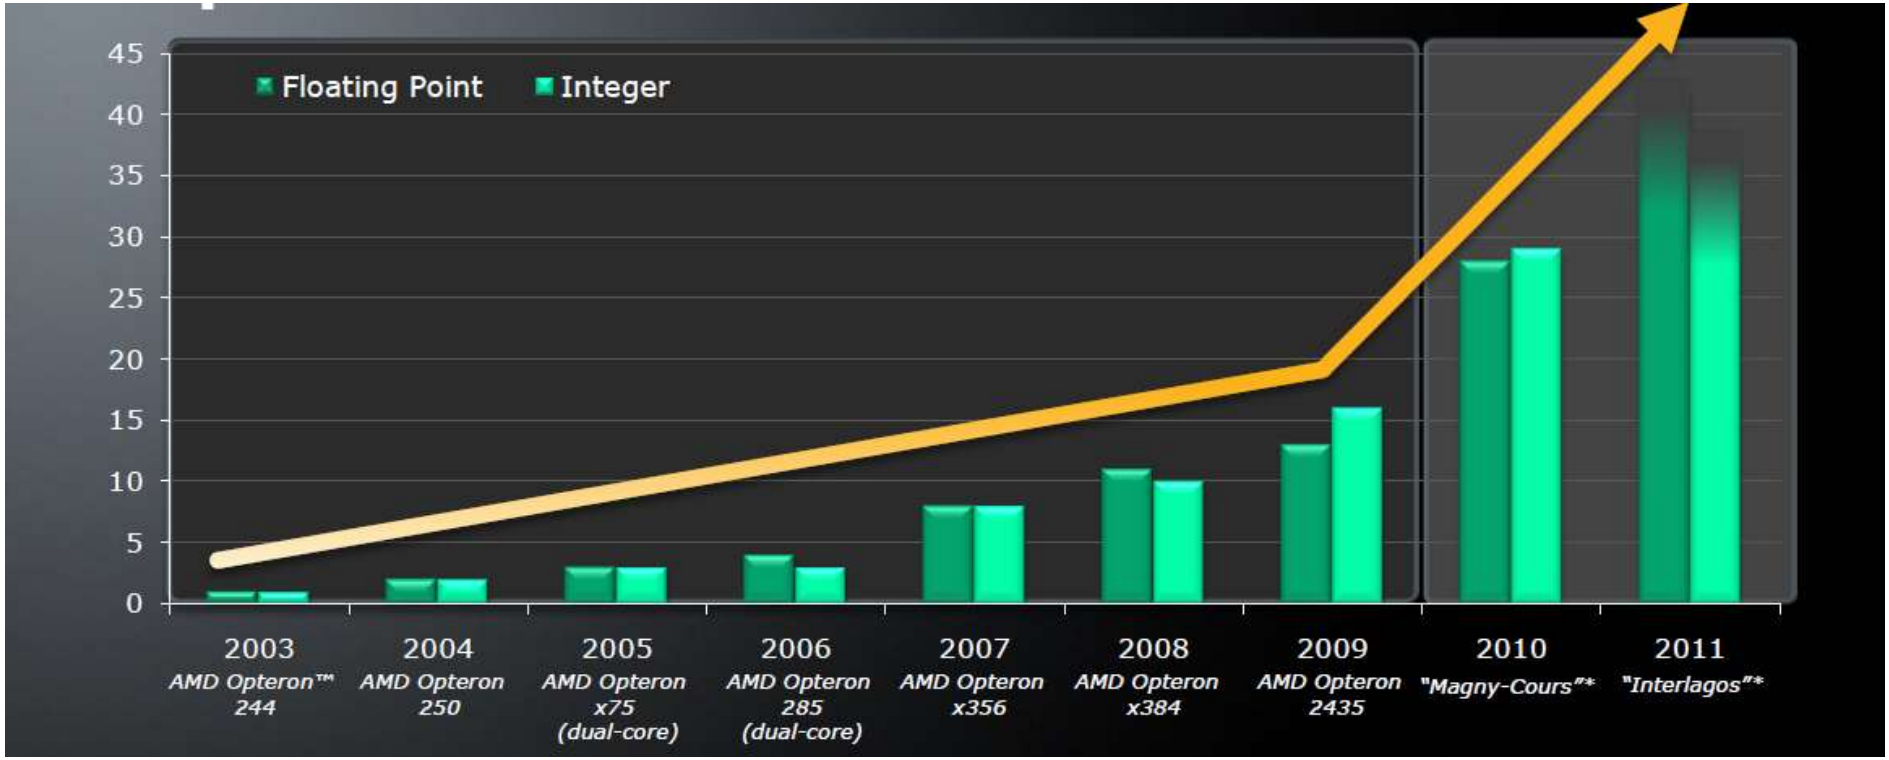

Introduction to the Family 15h lines of processors, designated also as the Bulldozer lines

- The Bulldozer project started in 2005 [4].
- First release of Bulldozer-based desktops: 10/2011
  First release of Bulldozer-based servers: 11/2011
- Bulldozer-based processors are built up of compute modules.

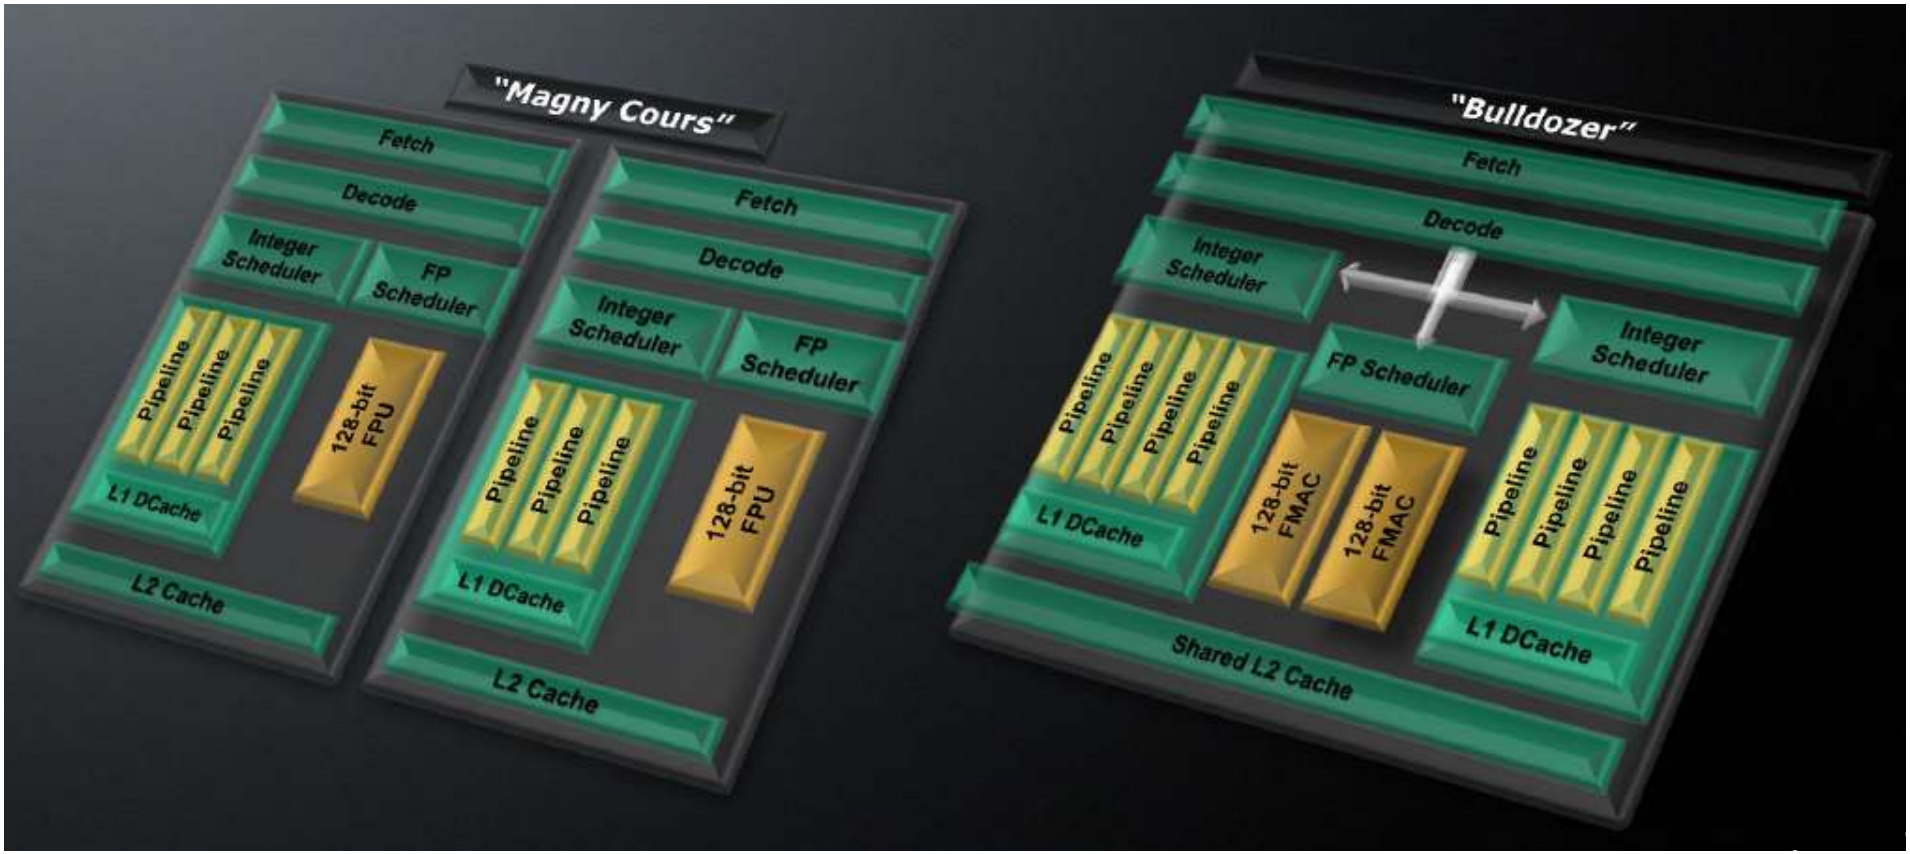

The compute module of the Family 15h processors

It is designated also as the Bulldozer module.

A Bulldozer module can execute two threads in parallel i.e. it can be considered as built up of two cores [15].

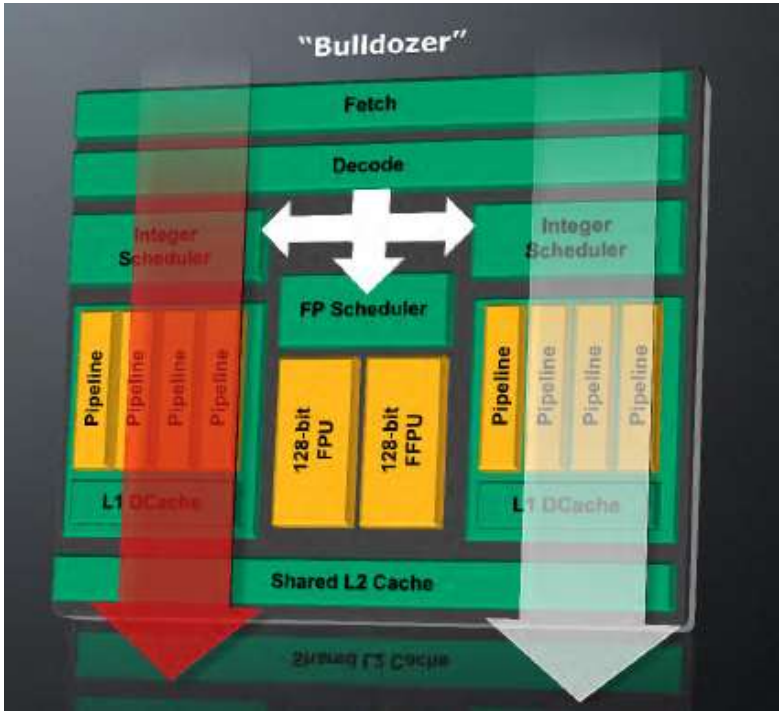

1. Overview of AMD's Family 15h lines, based on high-perf. Bulldozer modules (8)

The difference to traditional cores is that both cores of a Bulldozer module have beyond dedicated also shared resources in order to reduce silicon area and save power [4].

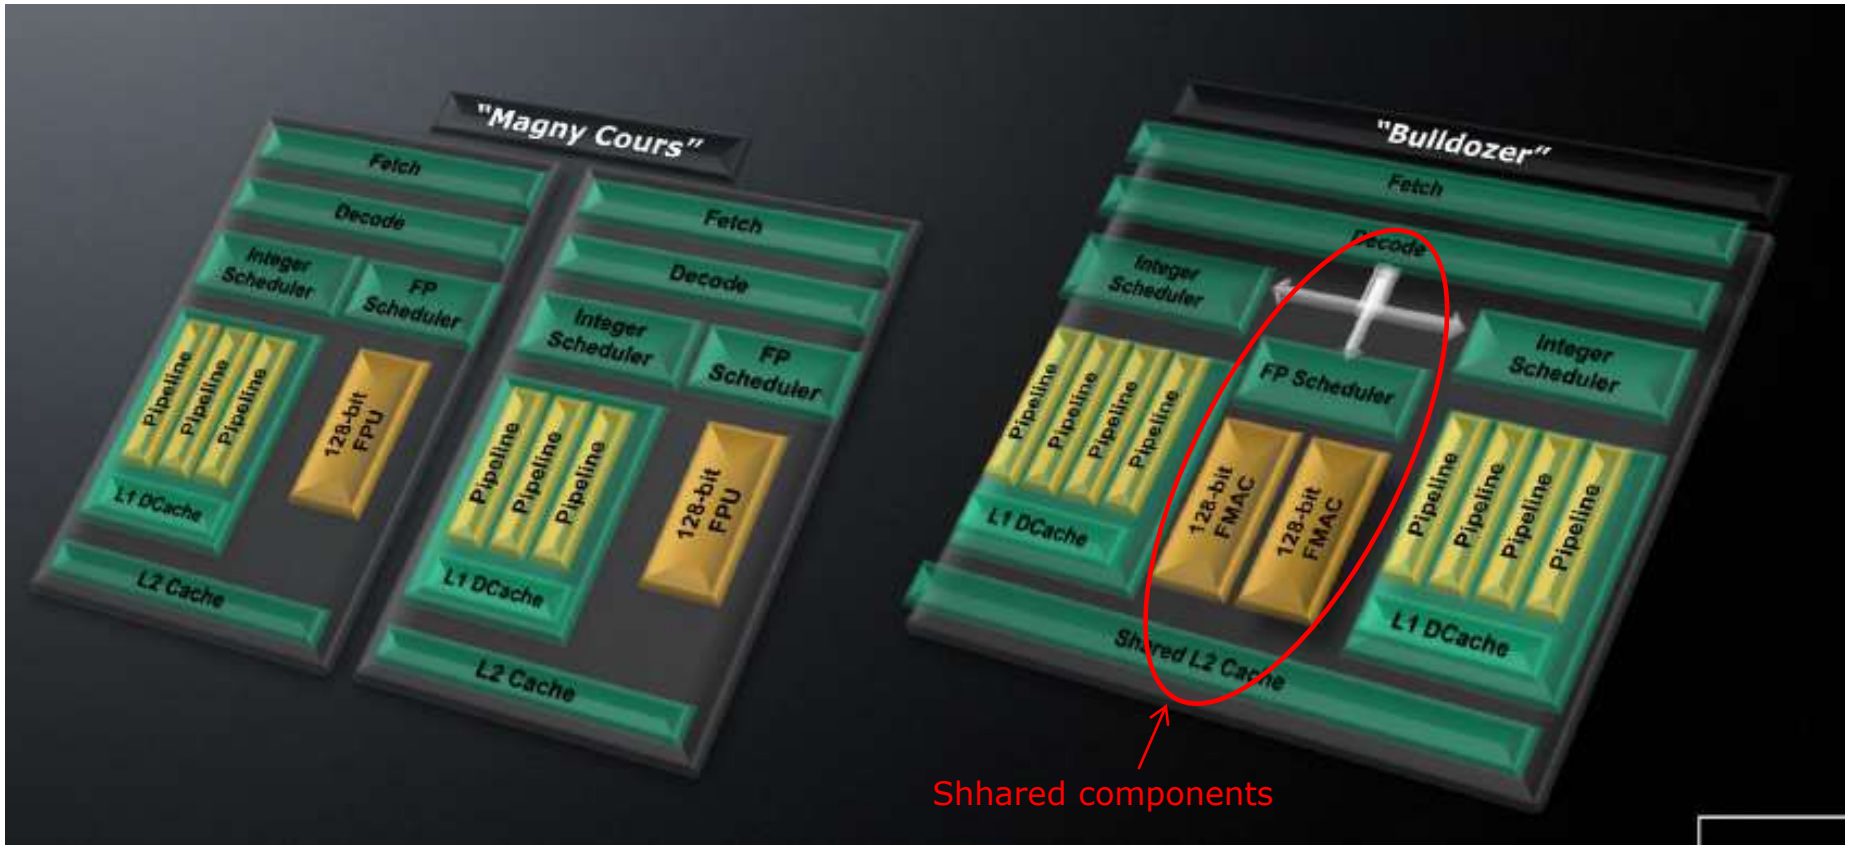

Shared and dedicated components of the Bulldozer cores

Dedicated and shared components of the Bulldozer cores are indicated in the Figure below. Shared components may be shared either at the module level or at the chip level, as the next Figure indicates [4].

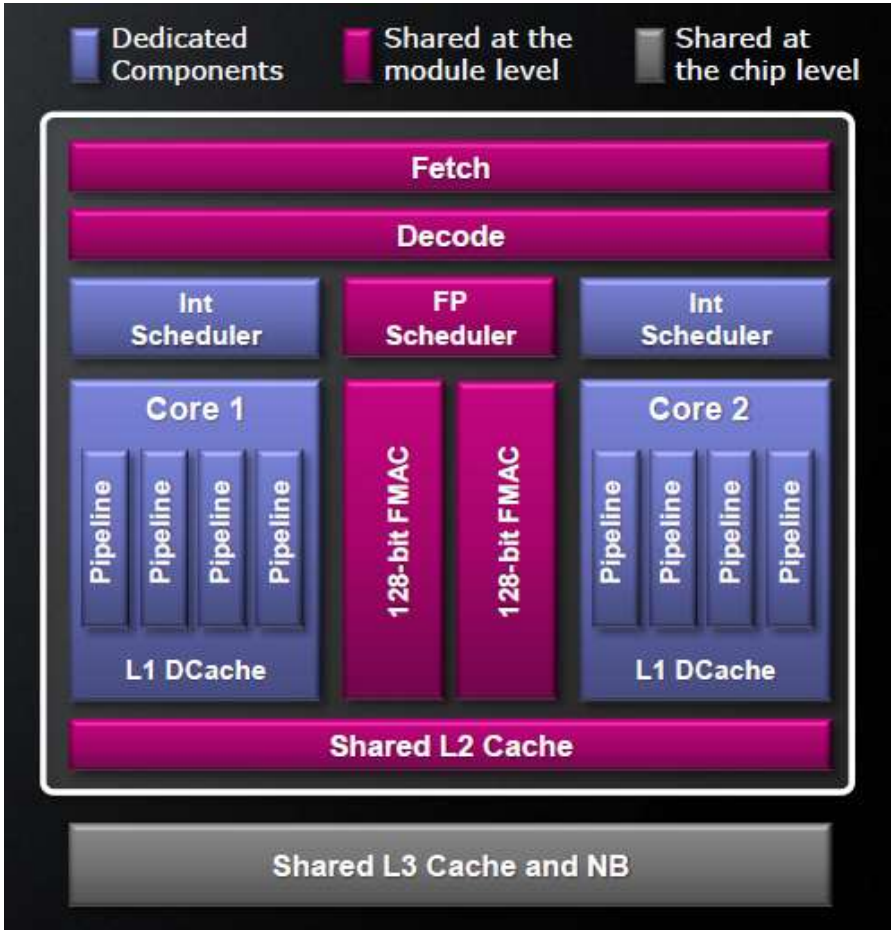

Design philosophy of using compute modules in Bulldozer-based designs

Main design aspects-1 [3]

a) AMD optimized the microarchitecture of their Bulldozer-based processors for multithreaded workloads rather than for single threaded performance.

b) In light of the Fusion system architecture concept AMD's belief is that heavy FP tasks should not be executed on the CPU cores but on an integrated GPGPU.

As a consequence the FP part of the microarchitecture may be designed for a low FP-load.

c) A further key aspect was reducing power consumption.

This goal motivated a number of design decisions related to the microarchitecture and also the utilization of a number of low power techniques, that are discussed in Section 2.4.

AMD's decision to integrate two conventional cores into a Bulldozer module is in line with the aspects a) to c) since

- a module provides two high performance, separate FX cores to support multithreading (aspect a))
- the choice to include a single moderately high performance shared FP unit satisfies aspects b) and c)
- sharing the complex x86 decoding by two cores reduces power consumption (aspect c), nevertheless it reduces the decode bandwidth.

Design philosophy of using compute modules

Main design aspects-2 [3]

d) As far as the single-threaded performance concerns, AMD focused on increasing clock speed rather than ILP.

To increase clock speed AMD lengthened Buldozer's pipeline compared to their previous K10/K10.5 designs (which used 12 FX stages and 17 FP stages).

AMD declined to release the pipeline depth of Bulldozer, nevertheless, according to unofficial sources Bulldozer has 18 pipeline stages [12].

1. Overview of AMD's Family 15h lines, based on high-perf. Bulldozer modules (12)

Remark

Number of pipeline stages in recent Intel and AMD processors

Processor	No. of pipeline stages
K8 to K10.5	Integer pipeline: 12 stages FP pipeline: 17 stages
Bulldozer	18 stages?
Core 2	14 stages
Penryn	14 stages
Nehalem	16 stages
Sandy Bridge/Ivy Bridge	14 stages

Example: Clock speed gain achieved by the 1. generation Bulldozer design vs. the previous K10.5 design-1

As the basic building blocks of the 1. generation Bulldozer-based processors (and also of all futher generations) are 4-module units, called in case of the 1. generation as the Orochi dies, and these 4-module units include 8 cores, (nevertheless with shared resources) actually we will compare clock frequencies of 8-core Family 15h Bulldozer systems with the previous 6-core K10.5 Istambul-based designs.

a) Servers

Comparing clock speeds of

K10.5 Istambul-based 6C DP servers (Lisbon) vs. Family 15h Bulldozer based 8C DP servers (called Valencia).

Main operational parameters of AMD's K10.5 Istambul-based DP servers (Lisbon) [13]

Model Number	Step.	Cores	Freq.	L2 Cache	L3 Cache	HT	Multi	Voltage	ACP	TDP
D0, Quad core										
Opteron 4122	D0	4	2.2 GHz	4x 512 KB	6 MB	3.2 GHz	11x	1.3125	75 W	95 W
Opteron 4130	D0	4	2.6 GHz	4x 512 KB	6 MB	3.2 GHz	13x	1.3125	75 W	95 W
D0, Quad core, high-efficiency										
Opteron 41LE HE	D0	4	2.3 GHz	4x 512 KB	6 MB	2.2 GHz	11.5x	1.1875	50 W	65 W
Opteron 41QS HE	D0	4	2.5 GHz	4x 512 KB	6 MB	2.2 GHz	12.5x	1.1875	50 W	65 W
D1, Six core										
Opteron 4180	D1	6	2.6 GHz	6x 512 KB	6 MB	3.2 GHz	13x	1.35	75 W	95 W
Opteron 4184	D1	6	2.8 GHz	6x 512 KB	6 MB	3.2 GHz	14x	1.35	75 W	95 W
D1, Six core, high-efficiency										
Opteron 41KX HE	D1	6	2.2 GHz	6x 512 KB	6 MB	2.2 GHz	11x	1.1875	50 W	65 W
Opteron 4170 HE	D1	6	2.1 GHz	6x 512 KB	6 MB	3.2 GHz	10.5x	1.1875	50 W	65 W
Opteron 4174 HE	D1	6	2.3 GHz	6x 512 KB	6 MB	3.2 GHz	11.5x	1.1875	50 W	65 W
Opteron 4176 HE	D1	6	2.4 GHz	6x 512 KB	6 MB	3.2 GHz	12x	1.1875	50 W	65 W
D1, Six core, energy-efficient										
Opteron 41GL EE	D1	6	1.8 GHz	6x 512 KB	6 MB	2.2 GHz	9x	0.9625	32 W	40 W
Opteron 4162 EE	D1	6	1.7 GHz	6x 512 KB	6 MB	3.2 GHz	8.5x	0.9625	32 W	35 W
Opteron 4164 EE	D1	6	1.8 GHz	6x 512 KB	6 MB	3.2 GHz	9x	0.9625	32 W	35 W

Main operational parameters of AMD's Family 15h-based DP servers (Valencia) [13]

Model Number	Step.	Cores	Freq.			Cache		HT	Multi	V_{core}	ACP	TDP
			Base	All turbo	Max. turbo	L2	L3					
B2, Six core												
Opteron 4226	B2	6	2.7 GHz	2.9 GHz	3.1 GHz	3 × 2 MB	8 MB	3.2 GHz	13.5×-15.5×		75 W	95 W
Opteron 4234	B2	6	3.1 GHz	3.3 GHz	3.5 GHz	3 × 2 MB	8 MB	3.2 GHz	15.5×-17.5×		75 W	95 W
Opteron 4238	B2	6	3.3 GHz	3.5 GHz	3.7 GHz	3 × 2 MB	8 MB	3.2 GHz	16.5×-18.5×		75 W	95 W
B2, Six core, high-efficiency												
Opteron 4228 HE	B2	6	2.8 GHz	3.1 GHz	3.6 GHz	3 × 2 MB	8 MB	3.2 GHz	14×-18×		50 W	65 W
B2, Eight core												
Opteron 4280	B2	8	2.8 GHz	3.1 GHz	3.5 GHz	4 × 2 MB	8 MB	3.2 GHz	14×-17.5×		75 W	95 W
Opteron 4284	B2	8	3.0 GHz	3.3 GHz	3.7 GHz	4 × 2 MB	8 MB	3.2 GHz	15×-18.5×		75 W	95 W
B2, Eight core, high-efficiency												
Opteron 4274 HE	B2	8	2.5 GHz	2.8 GHz	3.5 GHz	4 × 2 MB	8 MB	3.2 GHz	12.5×-17.5×		50 W	65 W
B2, Eight core, energy-efficient												
Opteron 4256 EE	B2	8	1.6 GHz	1.9 GHz	2.8 GHz	4 × 2 MB	8 MB	3.2 GHz	8×-14×		32 W	35 W

Example: Clock speed gain achieved by the 1. generation Bulldozer design vs. the previous K10.5 design-2

b) Desktops

Comparing clock speeds of

K10.5 Istambul-based 6C desktops (Phenom II X6) vs. Family 15h Bulldozer based 8C desktops (FX).

Main features of AMD's K10.5-based Phenom™ II X6 desktop processors [14]

Model Number	Frequency	Total L2 Cache	L3 Cache	Packaging	Thermal Design Power	CMOS Technology
1100T*	3.3 GHz	3MB	6MB	socket AM3	125W	45nm SOI
1090T*	3.2 GHz	3MB	6MB	socket AM3	125W	45nm SOI
1075T	3.0 GHz	3MB	6MB	socket AM3	125W	45nm SOI
1065T	2.9 GHz	3MB	6MB	socket AM3	95W	45nm SOI
1055T	2.8 GHz	3MB	6MB	socket AM3	125W	45nm SOI
1045T	2.7 GHz	3MB	6MB	socket AM3	95W	45nm SO

Main features of AMD's 1. generation Bulldozer-based FX desktop processors [14]

Model Number	Frequency	Total L2 Cache	L3 Cache	Packaging	Thermal Design Power	CMOS Technology
FX 8150	3.6/4.2 GHz	8MB	8MB	socket AM3+	125W	32nm SOI
FX 8120	3.1/4.0 GHz	8MB	8MB	socket AM3+	125W	32nm SOI
FX 8100	3.1/3.7 GHz	8MB	8MB	socket AM3+	95W	32nm SOI
FX 6200	3.8/4.1 GHz	6MB	8MB	socket AM3+	125W	32nm SOI
FX 6100	3.3/3.9 GHz	6MB	8MB	socket AM3+	95W	32nm SOI
FX 4170	4.2/4.3 Ghz	4MB	8MB	socket AM3+	125W	32nm SOI
FX 4100	3.6/3.8 Ghz	4MB	8MB	socket AM3+	95W	32nm SOI

Example: Clock speed gain achieved by the 1. generation Bulldozer design vs. the previous K10.5 design - Summary

	Clock frequencies of K10.5-based 6C lines	Clock frequencies of Family 15h-based 8C lines
DP lines	2.6-2.8 GHz	2.8-3.0 GHz
Desktops	2.8-3.3 GHz	3.1-3.6 GHz

The achieved clock speed gain of Bulldozer-based designs is about 10 – 20 %.

This speed gain is quite moderate since K10.5 Istambul-based processors are fabricated by 45 nm whereas Family 15h Bulldozer-based processors with 32 nm feature size.

The width of the Bulldozer cores

Bulldozer's cores have a new, 4-wide microarchitecture unlike previous 3-wide K8 Family 12h designs, as detailed in Section 2.2.3.

Remark

With the 4-wide Bulldozer design AMD caught up with Intel's 4-wide Core 2 (2006) and subsequent designs.

Overview of AMD's Family 15h (Bulldozer)-based processor lines

	Launched in	2011	2012	2013	2013	2015	2016
		Family 15h (00h-0Fh) (**Bulldozer**)	**Family 15h** (10h-1Fh) (**Piledriver**)	**Family 15h** (10h-1Fh) **(Piledriver v.2)**	**Family 15h** (30h-3Fh) (**Steamroller**)	**Family 15h** (60h-6Fh) (**Excavator v.1**)	**Family 15h** (77h-3Fh) (**Excavator v.2**)
Servers	**4P servers** (85-140 W)	Interlagos	Abu Dhabi				
	2P servers (85-140 W)	Valencia	Seoul				
	1P servers (85-140 W)	Zurich	Delhi				
Desktops	**High perf.** (~95-125 W)	Zambezi FX-Series	Vishera FX-Series				
	Mainstream (~65-95 W)		Trinity A10-A4	Richland A10/A8/A6/A4	Kaveri A10/A8		
Notebooks	**Mainstream** (~25-35 W)		Trinity A10/A8/A6M	Richland A10/A8/A6M	Kaveri FX/A10/A8P		Bristol Ridge FX/A12/A10P
	Ultra-thin (~10 – 15 W)		Trinity A10/A6M	Richland A10/A8/A6/A4M	A8 Pro/A8(B) A6 Pro/A6(B)	Carrizo FX/A10/A8P	Bristol Ridge FX/A12/A10P Stoney Ridge A9/A6
Tablets (~5 W)							

1. Overview of AMD's Family 15h lines, based on high-perf. Bulldozer modules (22)

Remark

- AMD (and also Intel) designates their notebook (laptop) processors and also tablets as mobile processors.
- Thus we use the terms mobile, notebook and laptop processors interchangeable.

Overview of Family 15h (Bulldozer)-based mainstream desktop lines -1

	Launched in	2011	2012	2013	2013	2015	2016
		Family 15h (00h-0Fh) (**Bulldozer**)	**Family 15h** (10h-1Fh) (**Piledriver**)	**Family 15h** (10h-1Fh) (**Piledriver v.2**)	**Family 15h** (30h-3Fh) (**Steamroller**)	**Family 15h** (60h-6Fh) (**Excavator v.1**)	**Family 15h** (77h-3Fh) (**Excavator v.2**)
Servers	**4P servers** (85-140 W)	Interlagos	Abu Dhabi				
	2P servers (85-140 W)	Valencia	Seoul				
	1P servers (85-140 W)	Zurich	Delhi				
Desktops	**High perf.** (~95-125 W)	Zambezi FX-Series	Vishera FX-Series				
	Mainstream (~65-95 W)		Trinity A10-A4	Richland A10/A8/A6/A4	Kaveri A10/A8		
Notebooks	**Mainstream** (~25-35 W)		Trinity A10/A8/A6M	Richland A10/A8/A6M	Kaveri FX/A10/A8P		Bristol Ridge FX/A12/A10P
	Ultra-thin (~10 – 15 W)		Trinity A10/A6M	Richland A10/A8/A6/A4M	A8 Pro/A8(B) A6 Pro/A6(B)	Carrizo FX/A10/A8P	Bristol Ridge FX/A12/A10P Stoney Ridge A9/A6
Tablets (~5 W)							

Overview of Family 15h (Bulldozer)-based mainstream desktop lines -2

Base arch.	Intro	Desktop family name	Series	Tech n.	Core count	L2 (up to)	L3	GPU	Memory (up to)	TDP [W]	Socket
Family 15h (10h-1Fh) (Piledriver)	9/2012	**Trinity**	A10/8/6/4 5x00(K)	32 nm	2/4 (1/2 CM)	2x2MB	--	Radeon HD7xxxD	DDR3-1866	65/100 W	FM2
Family 15h (10h-1Fh) (Piledriver v.2)	6/2013	**Richland**	A10/8/6/4 6x00(K)	32 nm	2/4 (1/2 CM)	2x2MB	--	Radeon HD8xxxD	DDR3-2133	65/100W	FM2
Family 15h (30h-3Fh) (Steamroller)	1/2014	**Kaveri**	A10 Pro/A10 A8 Pro/A8 7x00(K/B)	28 nm	2/4 (1/2 CM)	2x2MB	--	Radeon HD7xxD	DDR3-2133	65/95 W	FM2+

CM: Computing Module

Overview of Family 15h (Bulldozer)-based mainstream mobile lines -1

	Launched in	2011	2012	2013	2013	2015	2016
		Family 15h (00h-0Fh) (**Bulldozer**)	**Family 15h** (10h-1Fh) (**Piledriver**)	**Family 15h** (10h-1Fh) **(Piledriver v.2)**	**Family 15h** (30h-3Fh) (**Steamroller**)	**Family 15h** (60h-6Fh) (**Excavator v.1**)	**Family 15h** (77h-3Fh) (**Excavator v.2**)
Servers	**4P servers** (85-140 W)	Interlagos	Abu Dhabi				
	2P servers (85-140 W)	Valencia	Seoul				
	1P servers (85-140 W)	Zurich	Delhi				
Desktops	**High perf.** (~95-125 W)	Zambezi FX-Series	Vishera FX-Series				
	Mainstream (~65-95 W)		Trinity A10-A4	Richland A10/A8/A6/A4	Kaveri A10/A8		
Notebooks	**Mainstream** (~25-35 W)		Trinity A10/A8/A6M	Richland A10/A8/A6M	Kaveri FX/A10/A8P		Bristol Ridge FX/A12/A10P
	Ultra-thin (~10 – 15 W)		Trinity A10/A6M	Richland A10/A8/A6/A4M	A8 Pro/A8(B) A6 Pro/A6(B)	Carrizo FX/A10/A8P	Bristol Ridge FX/A12/A10P Stoney Ridge A9/A6
Tablets (~5 W)							

Overview of Family 15h (Bulldozer)-based mainstream mobile lines -2

Base arch.	Intro	Desktop family name	Series	Tech n.	Core count	L2 (up to)	L3	GPU	Memory (up to)	TDP [W]	Socket
Family 15h (10h-1Fh) (Piledriver)	5/2012	**Trinity**	A10/8/6M 4xxxM	32 nm	2/4 (1/2 CM)	2x2 MB	--	Radeon HD7xxxD	DDR3-1600 DDR3L-1600	25 W 35 W	FP2 FS1r2
Family 15h (10h-1Fh) (Piledriver v.2)	2/2013	**Richland**	A10/8/6/4M 5x5xM	32 nm	2/4 (1/2 CM)	2x2MB	--	Radeon HD8xxxD	DDR3-1866 DDR3L-1600	35 W	FS1r2
Family 15h (30h-3Fh) (Steamroller)	6/2014	**Kaveri**	FX-7600P A10-7400P A8-7200P	28 nm	4 (2 CM)	2x2MB	--	Radeon HD7xxxD HD6xxxD HD5xxxD	DDR3-2133 DDR3L-1866	35 W	FP3
Family 15h (60h-6Fh) (Excavator v.2)	5/2016	**Bristol Ridge**	FX-9830P A12-9730P A10-9630P	28 nm	4 (2 CM)	2x1MB	--	Radeon HD7xxxD HD5xxxD	DDR4-2400	35 W	FP4

Overview of Family 15h (Bulldozer)-based ultra-thin mobile lines -1

	Launched in	2011	2012	2013	2013	2015	2016
		Family 15h (00h-0Fh) (**Bulldozer**)	**Family 15h** (10h-1Fh) (**Piledriver**)	**Family 15h** (10h-1Fh) (**Piledriver v.2**)	**Family 15h** (30h-3Fh) (**Steamroller**)	**Family 15h** (60h-6Fh) (**Excavator v.1**)	**Family 15h** (77h-3Fh) (**Excavator v.2**)
Servers	**4P servers** (85-140 W)	Interlagos	Abu Dhabi				
	2P servers (85-140 W)	Valencia	Seoul				
	1P servers (85-140 W)	Zurich	Delhi				
Desktops	**High perf.** (~95-125 W)	Zambezi FX-Series	Vishera FX-Series				
	Mainstream (~65-95 W)		Trinity A10-A4	Richland A10/A8/A6/A4	Kaveri A10/A8		
Notebooks	**Mainstream** (~25-35 W)		Trinity A10/A8/A6M	Richland A10/A8/A6M	Kaveri FX/A10/A8P		Bristol Ridge FX/A12/A10P
	Ultra-thin (~10 – 15 W)		Trinity A10/A6M	Richland A10/A8/A6/A4M	A8 Pro/A8(B) A6 Pro/A6(B)	Carrizo FX/A10/A8P	Bristol Ridge FX/A12/A10P Stoney Ridge A9/A6
Tablets (~5 W)							

Overview of Family 15h (Bulldozer)-based ultra-thin mobile lines -2

Base arch.	Intro	Desktop family name	Series	Tech n.	Core count	L2 (up to)	L3	GPU	Memory (up to)	TDP [W]	Socket
Trinity A10/A8/A6M	5/2012	**Trinity**	A10/6M 4x55M	32 nm	2/4 (1/2 CM)	2x2MB	--	Radeon HD7xxxG	DDR3-1333 DDR3L-1333	17/25 W	FP2
Family 15h (10h-1Fh) (Piledriver v.2)	5/2013	**Richland**	A10/8/6/4-5x45M	32 mm	2/4 (1/2 CM)	2x2MB	--	Radeon HD8xxxG	DDR3-1333	17/19 //25 W	FP2
Family 15h (30h-3Fh) (Steamroller)	6/2014	**Kaveri**	A8 Pro-7150B A8-7100	28 nm	4 (2 CM)	1MB	--	Radeon HD5xxxD	DDR3-1600 DDR3L-1600	19 W	FP3
			A6 Pro-7050B A6-7000		2 (1 CM)			Radeon HD4xxxD		17 W	
Family 15h (60h-6Fh) (Excavator v.1)	6/2015	**Carrizo**	FX8800P A10/8 8xxxP	28 nm	4 (2 CM)	2x1MB	--	Radeon HD8xxxD HD7xxxD HD8xxxD	DDR3-2133	15 W	FP4
Family 15h (60h-6Fh) (Excavator v.2)	5/2016	**Bristol Ridge**	FX-9800P A12-9700P A10-9600P	28 nm	4 (2 CM)	2x1MB	--	Radeon HD7xxxD HD5xxxD	DDR4-1866	15 W	FP4
Family 15h (60h-6Fh) (Excavator v.2)	5/2016	**Stony Ridge**[1]	A9-9420 A9-9410 A6-9220 A6-9210	28 nm	(1 CM)	1 MB	--	Radeon HD5xxxD HD4xxxD	DDR4-2133	15 W	FP4 FT4

1: Stony Ridge processors have only a single memory channel

AMD's APU generations [93]

AMD APU Generations							
	1st	2nd	3rd	4th	5th	6th	7th
Platform Name	Llano	Trinity	Kabini	Kaveri	Beema	Carrizo	Bristol Ridge
Core	K10 / Stars	Steamroller	Jaguar	Piledriver	Puma	Excavator	
Released	Q2 2011	Q2 2012	Q2 2013	Q1 2014	Q2 2014	Q2 2015	Q2 2016
Market	Main	Main	Entry	Main	Entry	Main	Both
Top SKU	A8-3550MX	A10-4657M	A6-5200	FX-7600P	A8-6410	FX-8800P	FX-9830P
Threads	4C / 4T	2M / 4T	4C / 4T	2M / 4T	4C / 4T	2M / 4T	2M / 4T
Peak MHz	2.7 GHz	3.2 GHz	2.0 GHz	3.6 GHz	2.4 GHz	3.4 GHz	3.7 GHz
TDP	45 W	35 W	25 W	35 W	15 W	35W	35 W
IGP Family	HD 6620G	HD 7000	HD 8400	R7	R5	R7	R7
IGP Generation	VLIW-5	VLIW-4	GCN 1.0	GCN 1.1	GCN 1.1	GCN 1.2	GCN 1.2
SPs	400 444 MHz	384 686 MHz	128 600 MHz	512 686 MHz	128 850 MHz	512 800 MHz	512

2. First generation Bulldozer-based (Family 15h Models (00h-0Fh) processor lines

- 2.1 Overview of the Family 15h Bulldozer-based processor lines
- 2.2 The Bulldozer Compute Module
- 2.3 The Orochi die
- 2.4 New power management features of the Bulldozer design
- 2.5 Bulldozer-based server lines
- 2.6 The Bulldozer-based Zambezi DT line

2.1 Overview of the Family 15h Bulldozer-based processor lines

2.1 Overview of the Bulldozer-based processor lines [3]

Officially designated as the Family 15h Models 00h-0Fh processor lines.

They are called also as the 1. generation Bulldozer-based processor lines.

2.1 Overview of the Family 15h Bulldozer-based processor lines (2)

Overview of AMD's Bulldozer-based server and high-performance desktop lines -1 [1]

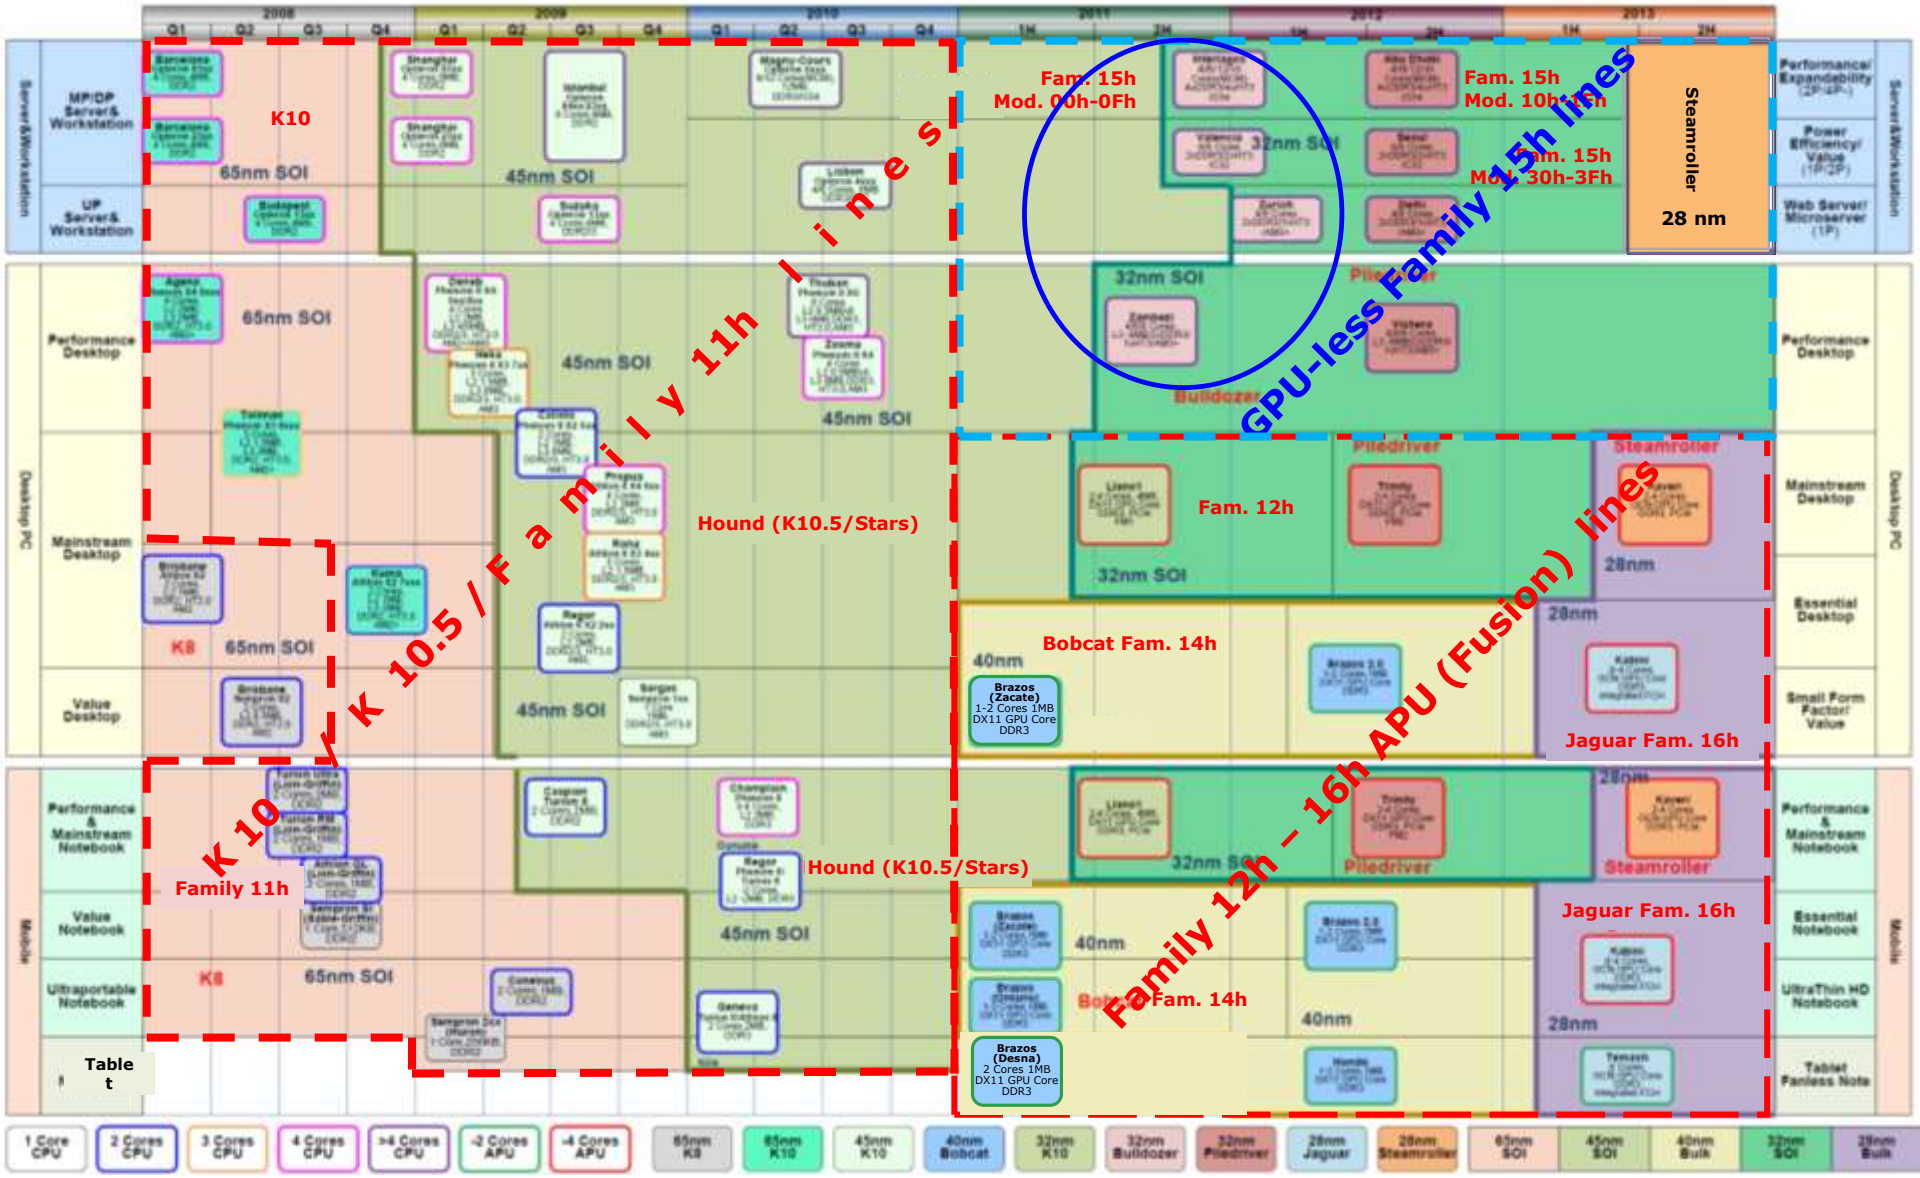

2.1 Overview of the Family 15h Bulldozer-based processor lines (3)

Overview of AMD's Bulldozer-based server and high-performance desktop lines -2

	Launched in	2011	2012	2013	2013	2015	2016
		Family 15h (00h-0Fh) (**Bulldozer**)	**Family 15h** (10h-1Fh) (**Piledriver**)	**Family 15h** (10h-1Fh) **(Piledriver v.2)**	**Family 15h** (30h-3Fh) (**Steamroller**)	**Family 15h** (60h-6Fh) (**Excavator v.1**)	**Family 15h** (77h-3Fh) (**Excavator v.2**)
Servers	**4P servers** (85-140 W)	Interlagos	Abu Dhabi				
	2P servers (85-140 W)	Valencia	Seoul				
	1P servers (85-140 W)	Zurich	Delhi				
Desktops	**High perf.** (~95-125 W)	Zambezi FX-Series	Vishera FX-Series				
	Mainstream (~65-95 W)		Trinity A10-A4	Richland A10/A8/A6/A4	Kaveri A10/A8		
Notebooks	**Mainstream** (~25-35 W)		Trinity A10/A8/A6M	Richland A10/A8/A6M	Kaveri FX/A10/A8P		Bristol Ridge FX/A12/A10P
	Ultra-thin (~10 – 15 W)		Trinity A10/A6M	Richland A10/A8/A6/A4M	A8 Pro/A8(B) A6 Pro/A6(B)	Carrizo FX/A10/A8P	Bristol Ridge FX/A12/A10P Stoney Ridge A9/A6
Tablets (~5 W)							

2.1 Overview of the Family 15h Bulldozer-based processor lines (4)

Positioning AMD's Bulldozer-based server lines

Base arch./ stepping		Intro	4P Server family name	Series	Techn	Cores (up to)	L2 (up to)	L3 (up to)	Memory (up to)	HT/ dir. (up to)	Socket
K8	C0/CG	4/2003	Sledge-hammer	800	130 nm	1C	1 MB	-	DDR-333	HT 1.0: 3.2 GB/s	940
	E4/E6	12/2004	Athens	800	90 nm	1C	1 MB	-	DDR-400	HT 2.0: 4.0 GB/s	940
	E1/E6	4/2005	Egypt	800	90 nm	2C	2*1 MB	-	DDR-400	HT 2.0: 4.0 GB/s	940
	F2/F3	8/2006	Santa Rosa	8200	90 nm	2C	2*1 MB	-	DDR2-667	HT 2.0: 4.0 GB/s	F
K10	BA/B1-B3	8/2007	Barcelona	8300	65 nm	4C	4*1/2 MB	2 MB	DDR2-667	HT 2.0: 4.0 GB/s	F
K10.5	C2/C3	11/2008	Shanghai	8300	45 nm	4C	4*1/2 MB	6 MB	DDR2-800	HT 2.0/3.0: 4.0/8.8 GB/s	F
	CE	6/2009	Istambul	8400	45 nm	6C	6*1/2 MB	6 MB	DDR2-800	HT 3.0: 9.6 GB/s	F
	D1	3/2010	Magny Course (2xIstambul)	6100	45 nm	2x6C	12*1/2 MB	6 MB	DDR3-1333	HT 3.1: 12.8 GB/s	G34
Fam 15h Mod. 00h-0Fh (Bulldozer)		11/2011	Interlagos (2xOrochi die)	6200	32 nm	2x4 CM (2x8 C)	2*4* 2 MB/CM	2* 8MB/ 4 CM	DDR3-1600	HT 3.1: 12.8 GB/s	G34
Fam. 15h Mod. 10h-1Fh (Piledriver)		11/2012	Abu Dhabi (2 dies)	6300	32 nm	2x4 CM (2x8 C)	2*4* 2 MB/CM	2* 8 MB/ 4 CM	DDR3-1866	HT 3.1 12.8 GB/s	G34
Fam. 15h Subsequent models		No lines launched									

2.1 Overview of the Family 15h Bulldozer-based processor lines (5)

Positioning AMD's high-performance desktop lines (except Zen-based lines)

Base arch./ stepping		Intro	High perf. DT family	Series	Techn.	Core count (up to)	L2 (up to)	L3 (up to)	Memory (up to)	HT/ dir. (up to)	Socket
K8	CG	9/2003	Claw-Hammer	Athlon 64	130 nm	1	1 MB	-	DDR-400	HT 2.0: 4.0 GB/s	754/ 939
	E4	4/2005	San Diego	Athlon 64	90 nm	1	1 MB	-	DDR-400	HT 2.0: 4.0 GB/s	939
	E6	5/2005	Toledo	Athlon 64 X2	90 nm	2	2*1 MB	-	DDR-400	HT 2.0: 4.0 GB/s	939
	E2/E3	5/2006	Windsor	Athlon 64 X2	90 nm	2	2*1 MB	-	DDR2-800	HT 2.0: 4.0 GB/s	AM2
K10	B2 B3	11/2007 3/2008	Agena	Phenom X4	65 nm	4	4*½ MB	2 MB	DDR2-1066	HT 3.0: 8.0 GB/s	AM2+
K10.5	C2 C2/C3	1/2009 2/2009	Deneb	Phenom II X4	45 nm	4	4*½MB	6 MB	DDR2-1066 DDR3-1333	HT 3.0: 8.0 GB/s	AM2+ AM3
	E0	4/2010	Thuban	Phenom II X6	45 nm	6	6*½MB	6 MB	DDR2-1066 DDR3-1333	HT 3.0: 8.0 GB/s	AM3
Fam. 11h (Griffin)		-	-	-	-	-	-	-	-	-	-
Fam. 12h (Llano)		6/2011	Llano	Fusion A8	32 nm	4	4*1 M	-	DDR3-1866	UMI: 5 GT/s	FM1
Fam. 14h (Bobcat)		-	-	-	-	-	-	-	-	-	-
Fam. 15h Models 00h-0Fh (Bulldozer)		10/2011	Zambezi	FX-series	32 nm	4 CM (8 C)	4x2 MB/CM	8 MB	DDR3-1866	HT 3.1: 12.8 GB/s	AM3+
Fam. 15h Models 10h-1Fh (Piledriver)		10/2012	Vishera	FX-series	32 nm	4 CM (8 C)	4*2 MB/CM	8 MB	DDR3-1866	HT 3.1: 12.8 GB/s	AM3+
No further Fam. 16h based lines		-	-	-	-	-	-	-	-	-	-

2.2 The Bulldozer Compute Module

- 2.2.1 Overview of the Bulldozer Compute Module
- 2.2.2 ISA extensions introduced in the Bulldozer design
- 2.2.3 The microarchitecture of the Bulldozer Compute Module
- 2.2.4 Assessing the performance potential of the Bulldozer Compute Module

2.2.1 Overview of the Bulldozer Compute Module

2.2.1 Overview of the Bulldozer Compute Module (1)

The Bulldozer Compute module

It includes two cores with dedicated and shared resources, as disscussed already in Chapter 1, but redrawn below [4].

The "Bulldozer" module has shared and dedicated components

The shared components:

- Help reduce power consumption
- Help reduce die space (cost)

The dedicated components:

- Help increase performance and scalability

"Bulldozer" dynamically switches between shared and dedicated components to maximize performance per watt

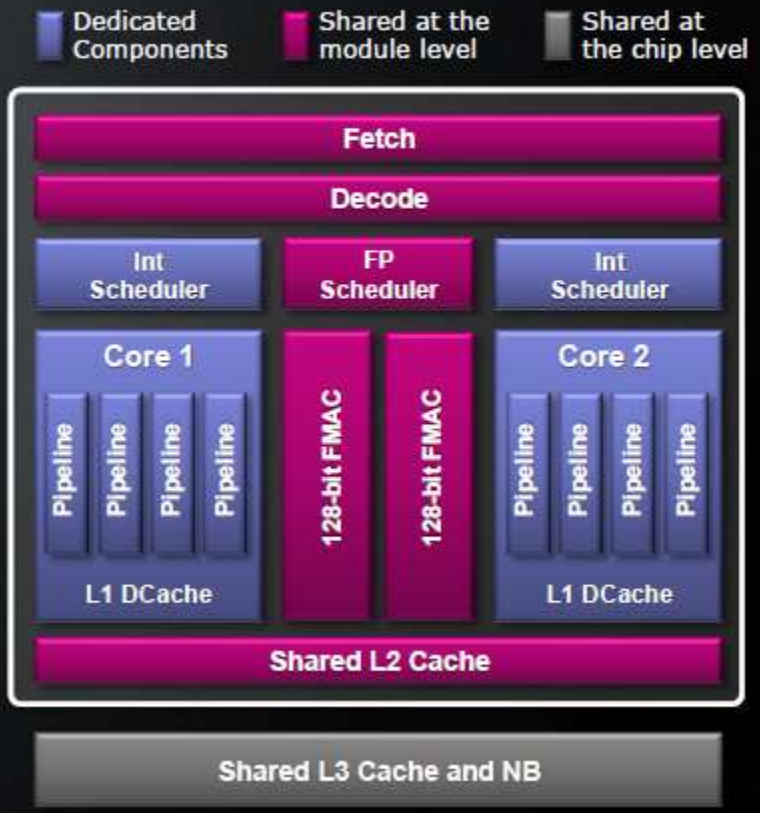

2.2.1 Overview of the Bulldozer Compute Module (2)

Principle of operation of a Bulldozer module [4]

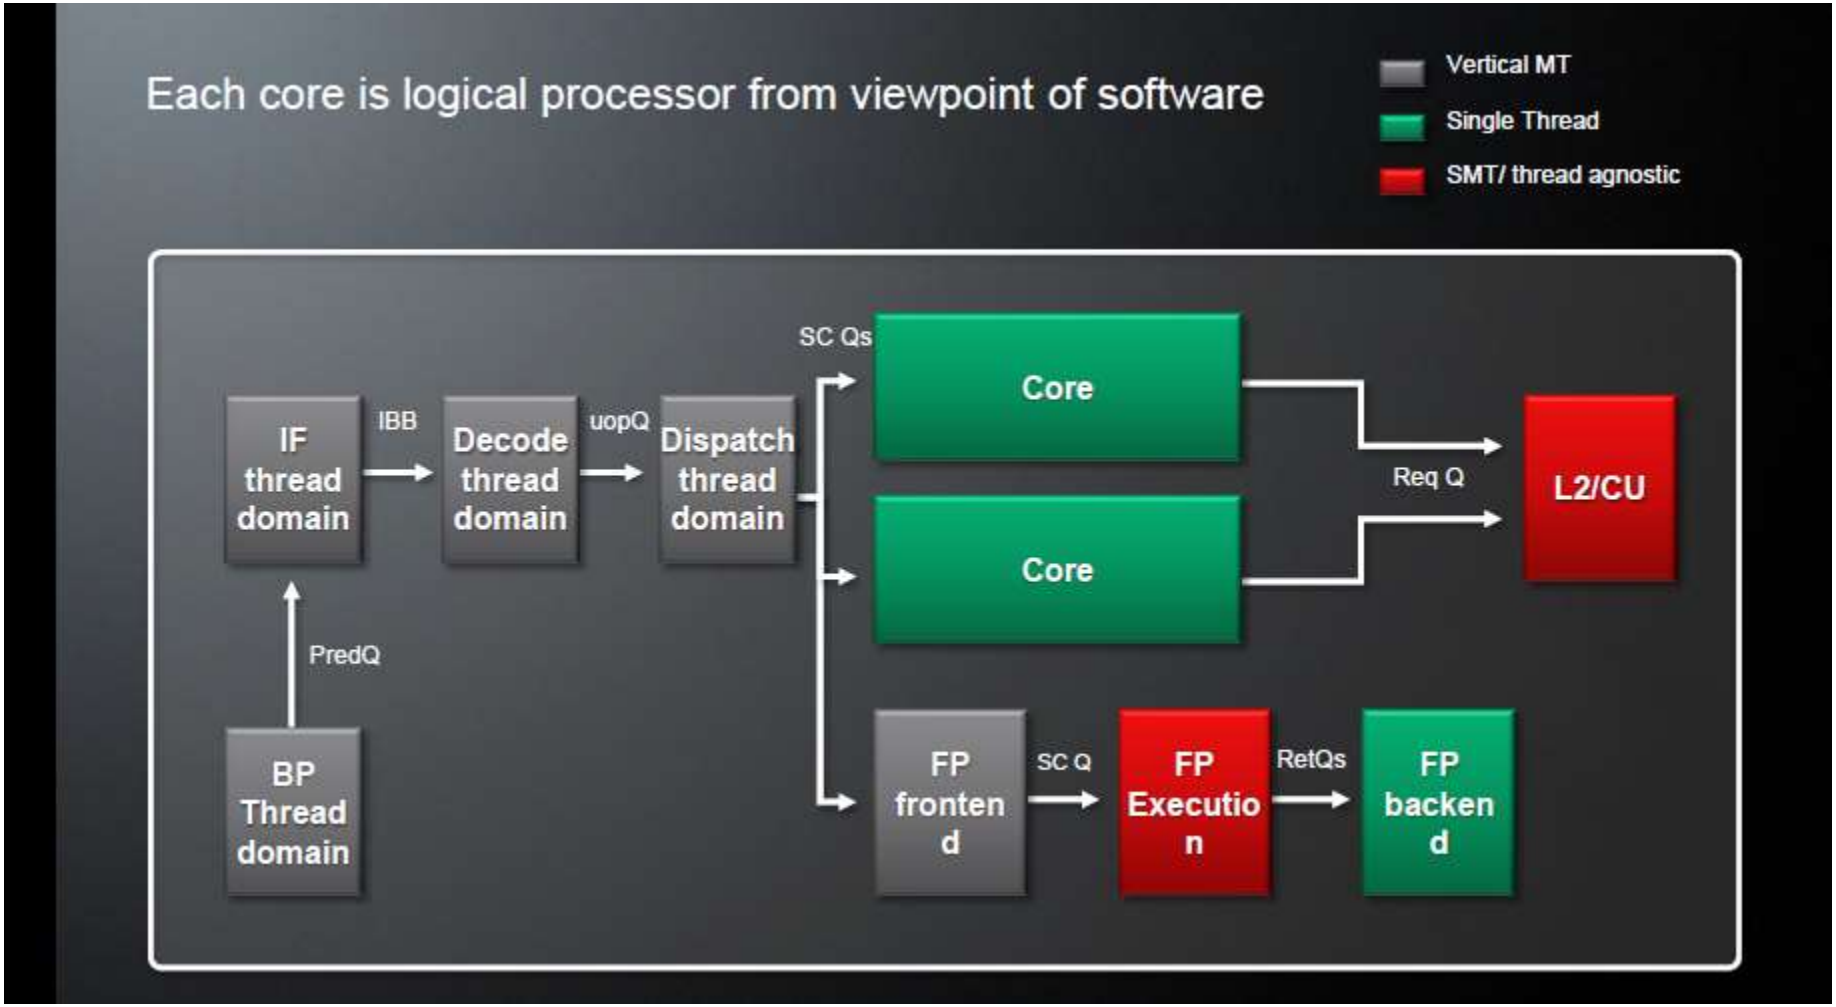

BP: Branch Prediction **Pred**: Prediction **Q**: Queue **Ret**: Return **Req:** Request **SC**: Scheduling

IBB: Instruction Bytes Buffers (see Section 2.2.3) **MT**: MultiThreading

2.2.2 ISA extensions introduced in the Bulldozer design

2.2.2 ISA extensions introduced in the Bulldozer design

New Bulldozer instructions and their possible use: [15]

Instructions	Applications/Use Cases
SSSE3, SSE4.1, SSE4.2 (AMD and Intel)	• Video encoding and transcoding • Biometrics algorithms • Text-intensive applications
AESNI PCLMULQDQ (AMD and Intel)	• Application using AES encryption • Secure network transactions • Disk encryption (MSFT BitLocker) • Database encryption (Oracle) • Cloud security
AVX (AMD and Intel)	Floating point intensive applications: • Signal processing / Seismic • Multimedia • Scientific simulations • Financial analytics • 3D modeling
FMA4 (AMD Unique)	HPC applications
XOP (AMD Unique)	• Numeric applications • Multimedia applications • Algorithms used for audio/radio

2.2.2 ISA extensions introduced in the Bulldozer design (2)

Introduction of ISA x86 extensions by Intel vs. AMD

Extension	Intel		AMD	
MMX	Pentium MMX	1/1997 →	3/1998	K6
3DNow!	-		2/1999	K6-2
Enh. 3DNow!	-		6/1999	K7 Athlon Model 1
3DNow! Professional	-		7/2001	K7 Athlon MP Model 6
SSE	Pentium III Katmai	2/1999 →	4/2003	K8 Sledgehammer
SSE2	Pentium 4 Willamete	12/2000 →	4/2003	K8 Sledgehammer
SSE3	Pentium 4 Presocott	2/2004 →	12/2004	K8 Athens
SSSE3	Core 2	7/2006 →	1/2011	Family 14h Bobcat
SSE4.a	-		8/2007	K10 Barcelona
SSE4.1	Penryn	11/2007 →	11/2011	Family 15h Bulldozer
SSE4.2	Nehalem	3/2009 →	11/2011	Family 15h Bulldozer
AES-NI	Westmere	3/2010 →	11/2011	Family 15h Bulldozer
PCLMULQDQ instr.	Westmere	3/2010 →	11/2011	Family 15h Bulldozer
AVX	Sandy Bridge	1/2011 →	11/2011	Family 15h Bulldozer
FMA4, XOP instrs.	-		11/2011	Family 15h Bulldozer

MMX: Multi Media Extension
SSE: Streaming SIMD extension
SSSE3: Supplemental SSE3 (SSSE3)

AES: Advanced Encryption Standard
AVX: Advanced Vector Extension

2.2.2 ISA extensions introduced in the Bulldozer design (3)

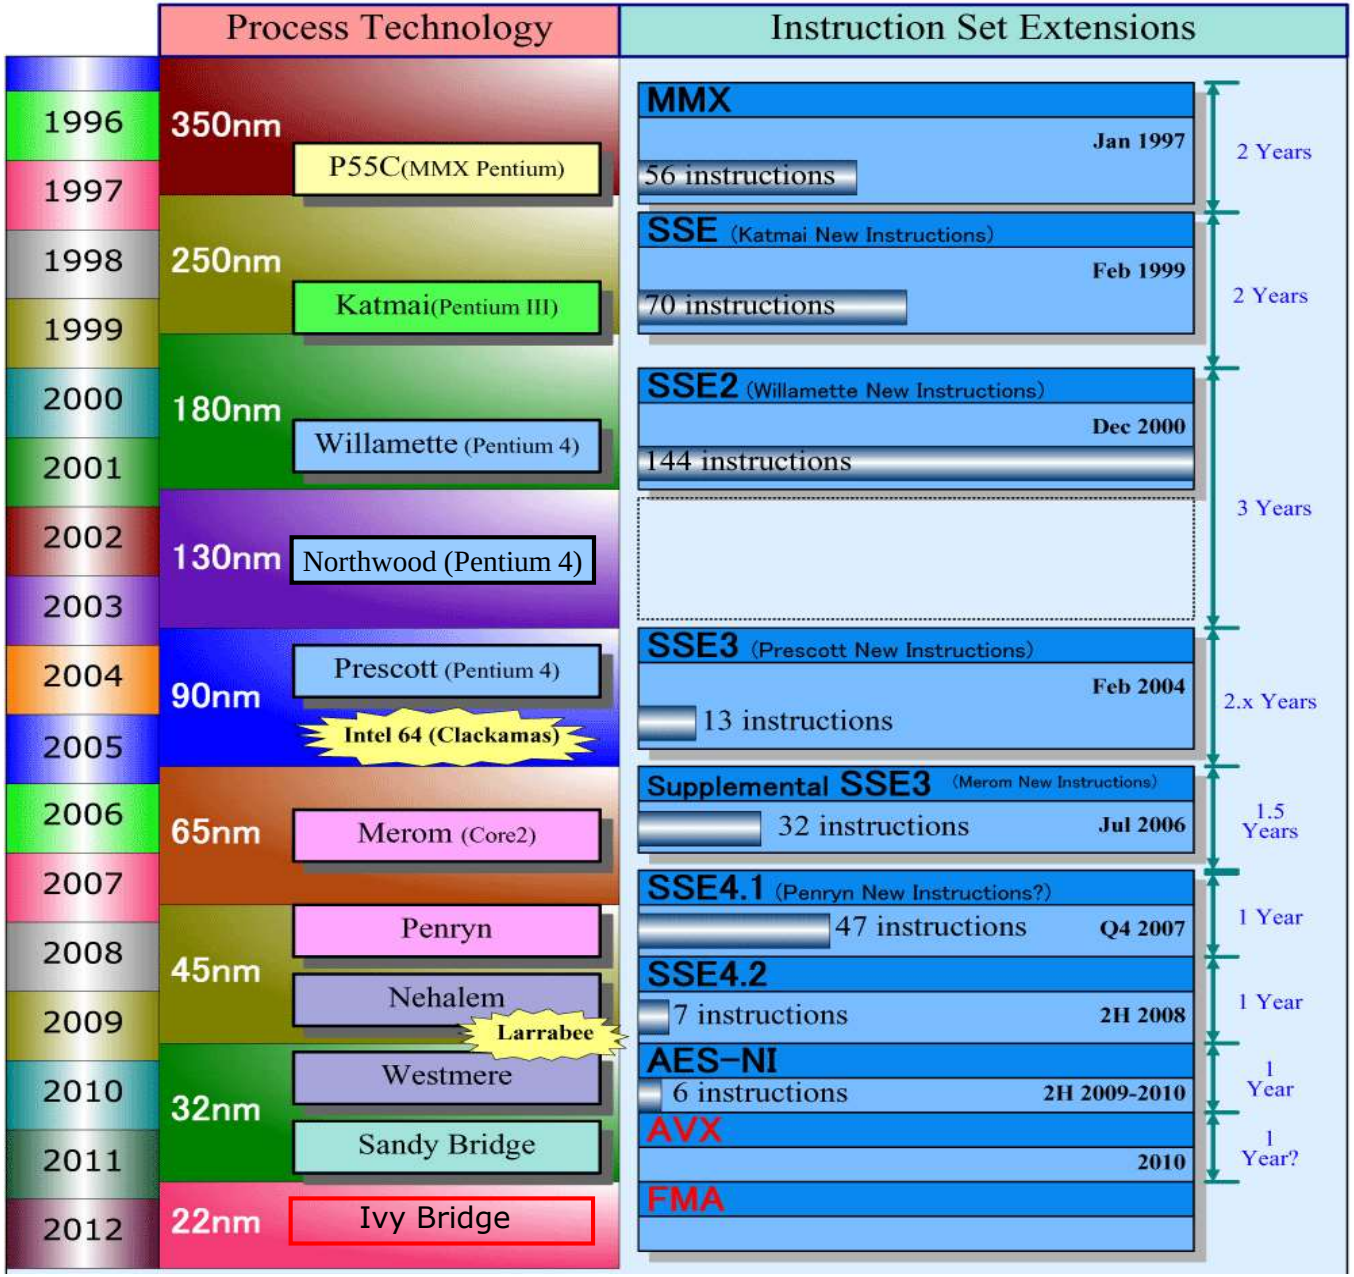

MultiMedia eXtensions

Streaming SIMD Extensions

Advanced Encryption Standard

Advanced Vector Extension

Fused Multiply-Add instr.

Figure: Overview of Intel's x86 ISA extensions (based on [44])

2.2.2 ISA extensions introduced in the Bulldozer design (4)

Comparison of FP-capabilities of Buldozer, Magny Course and Sandy Bridge [16]

Capability	Current AMD FPU	Sandy Bridge	Flex FP
Execute 128-bit FP	✓	✓	✓
Support SSSE3, SSE4.1, SSE4.2		✓	✓
Execute 128-bit AVX		✓	✓
Execute 256-bit AVX		✓	✓
Execute two 128-bit SSE or AVX ADD instructions in 1 cycle			✓
Execute two 128-bit SSE or AVX MUL instructions in 1 cycle			✓
Switch between SSE and AVX instructions without penalty			✓
Execute FMA operations (A=B+C*D)			✓
Supports XOP			✓
FLOPs per cycle (128-bit FP)	48	32	64
FLOPS per cycle (128-bit AVX)	-	32	64
FLOPS per cycle (256-bit AVX)	-	64	64

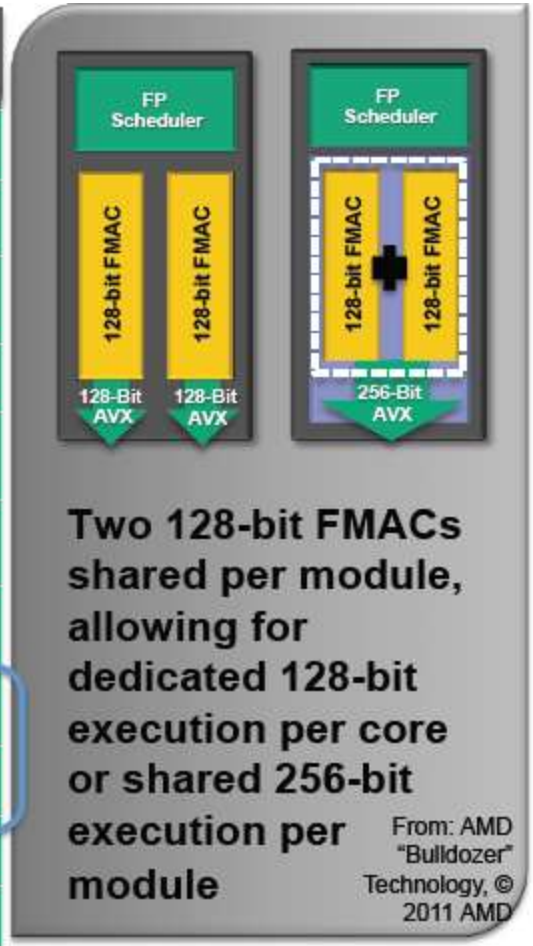

Sandybridge information from http://software.intel.com/en-us/avx/

2.2.2 ISA extensions introduced in the Bulldozer design (5)

Compiler support of Bulldozers new instructions [15]

Compilers	Support
Microsoft Visual Studio 2010 SP1	• All new instructions are supported • Developer has to manually call instructions
GCC 4.5	• All new instructions supported • Can recompile with an option to use new instructions
Open64 4.2.5	• All new instructions supported • Can recompile with an option to use new instructions
PGI 11.1	• AVX and SSE4.2 currently supported • Can recompile with an option to use new instructions • Planned future product release will add support for remaining instructions

2.2.3 The microarchitecture of the Bulldozer Compute Module

2.2.3 Microarchitecture of the Bulldozer Compute Module

Bulldozer based lines are built up on Bulldozer Compute Modules, each of which can be considered as being two conventional cores.

2.2.3 The microarchitecture of the Bulldozer Compute Module (2)

AMD's Bulldozer module contrasted with two cores of Magny Course [4]

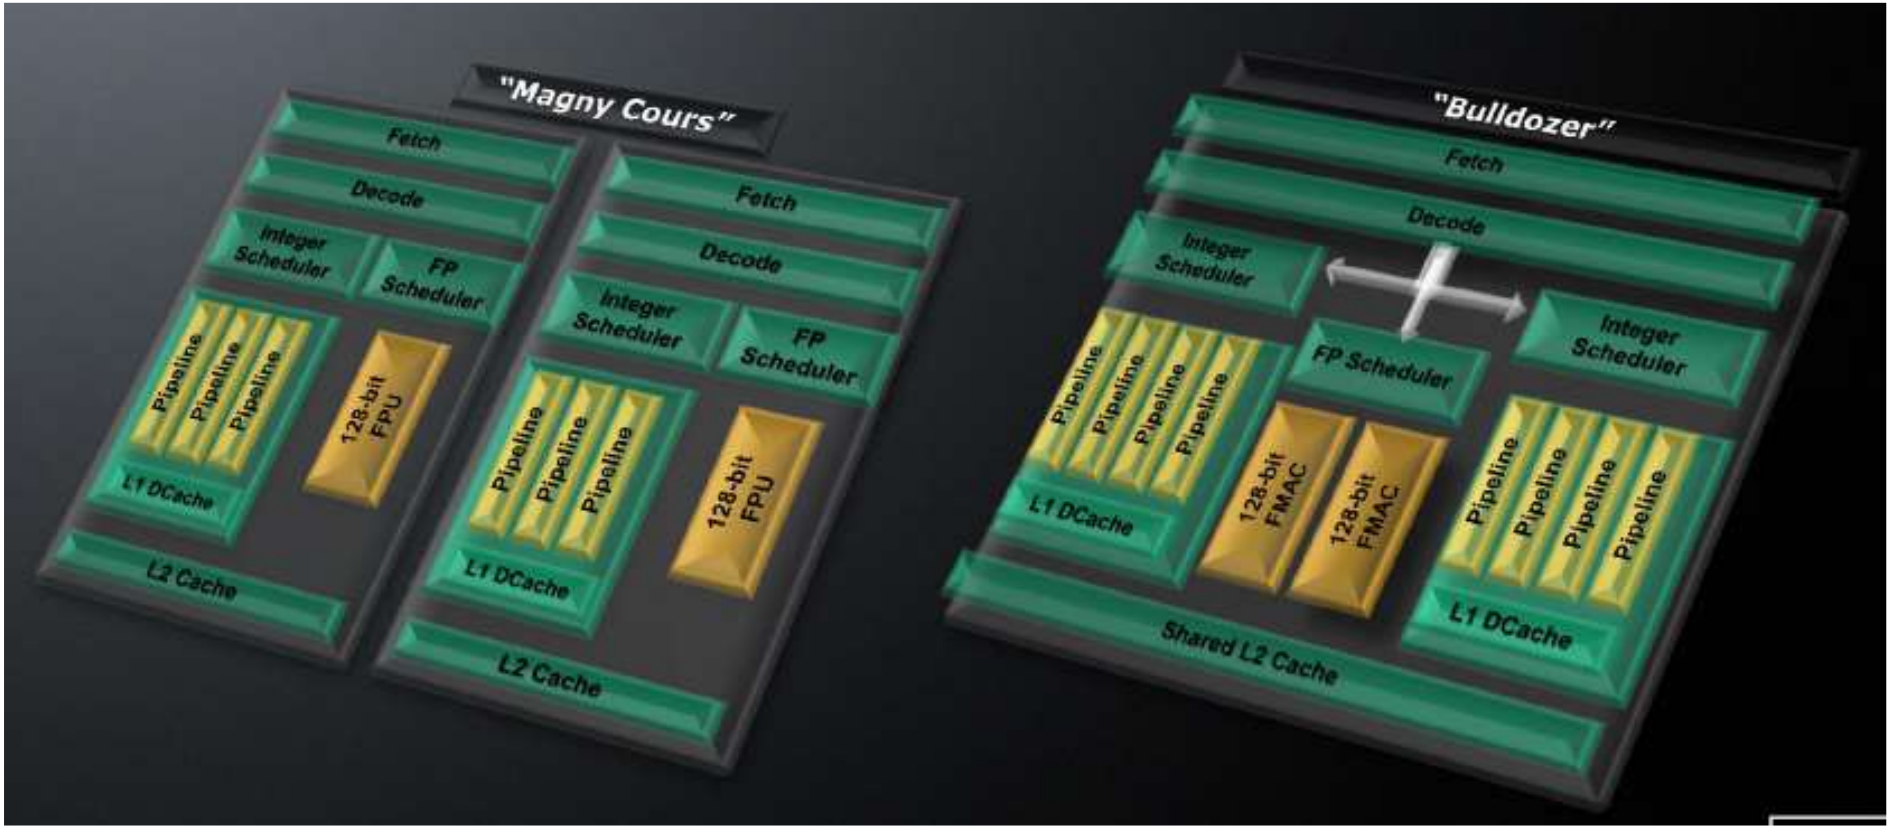

2.2.3 The microarchitecture of the Bulldozer Compute Module (3)

The microarchitecture of a Bulldozer core [10]

- 3. gen. superscalar
- Front-end: 4 wide

Mops: Macro-oerations (actually microoperations Produced by the decoders)

Shared FP unit

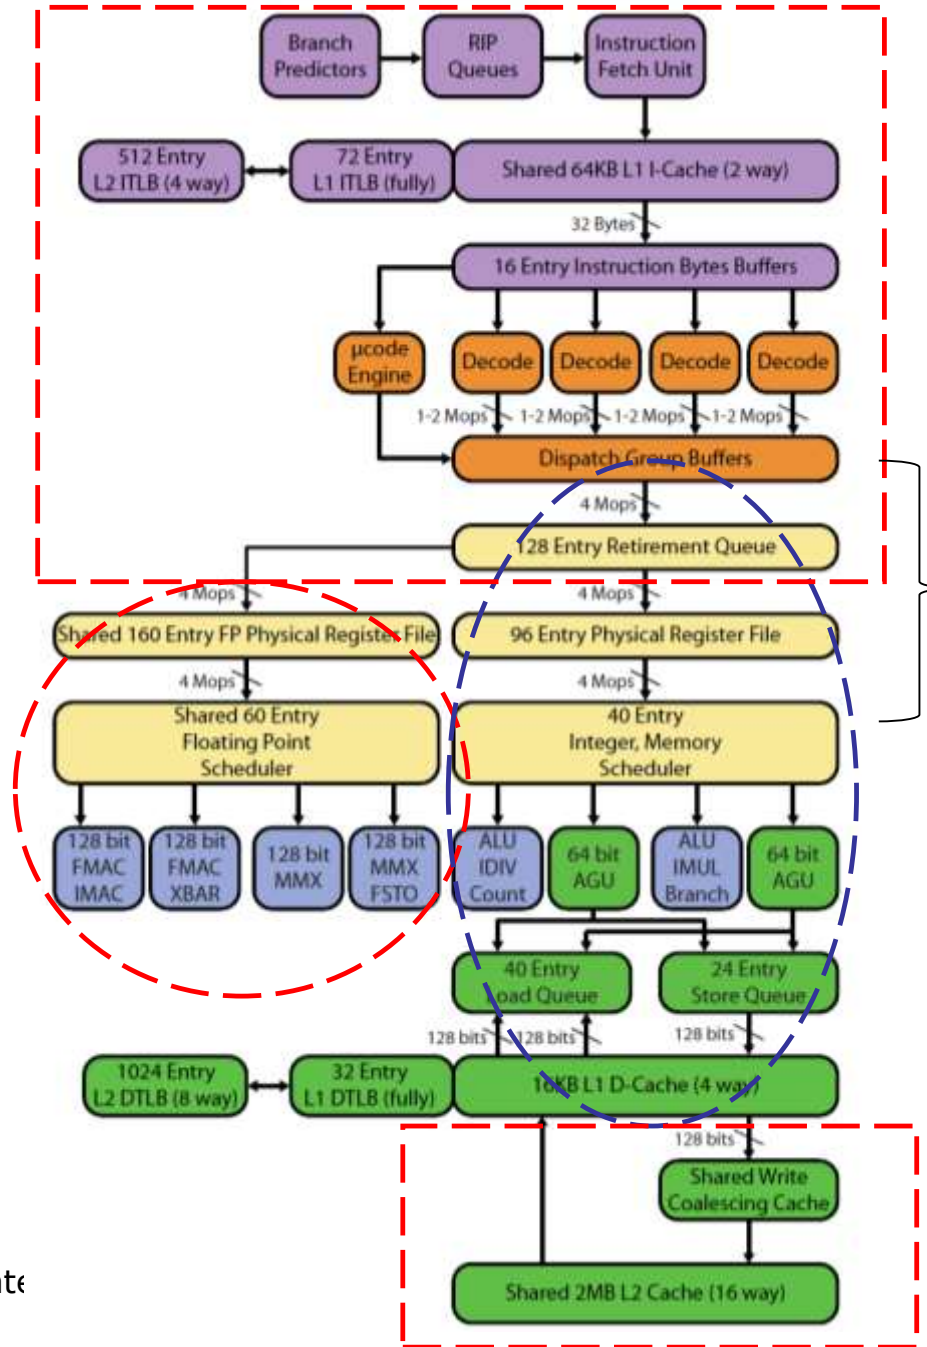

Shared front end

4-wide front end

Per thread dedicated FX unit

Shared L2

RIP: Relative Instruction Pointe

2.2.3 The microarchitecture of the Bulldozer Compute Module (4)

While introducing a 4-wide microarchitecture AMD eliminated their intrinsic drawback vs. Intel that arose with the introduction of Intel's 4-wide Core 2 microarchitecture in 2006 whereas AMD remained stuck with their 3-wide K8 design until Bulldozer.

Block diagram of Intel's Core 2 microarchitecture [11]

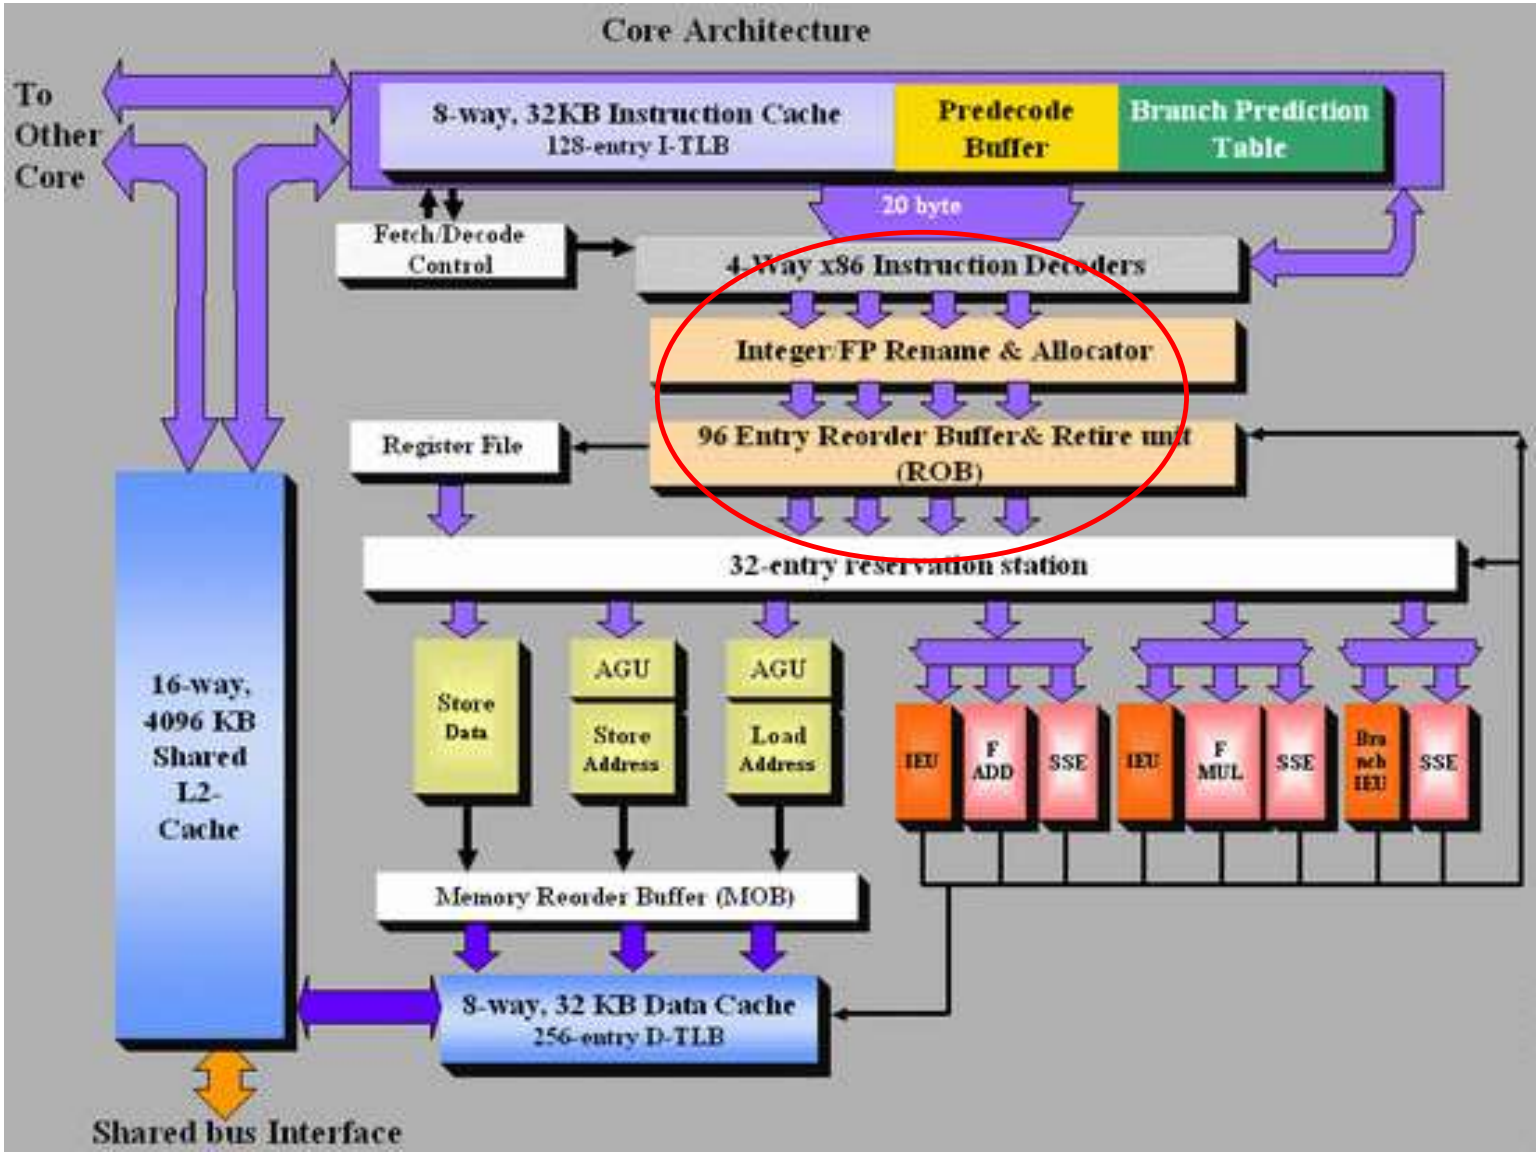

Block diagram of AMD's K8 microarchitecture [11]

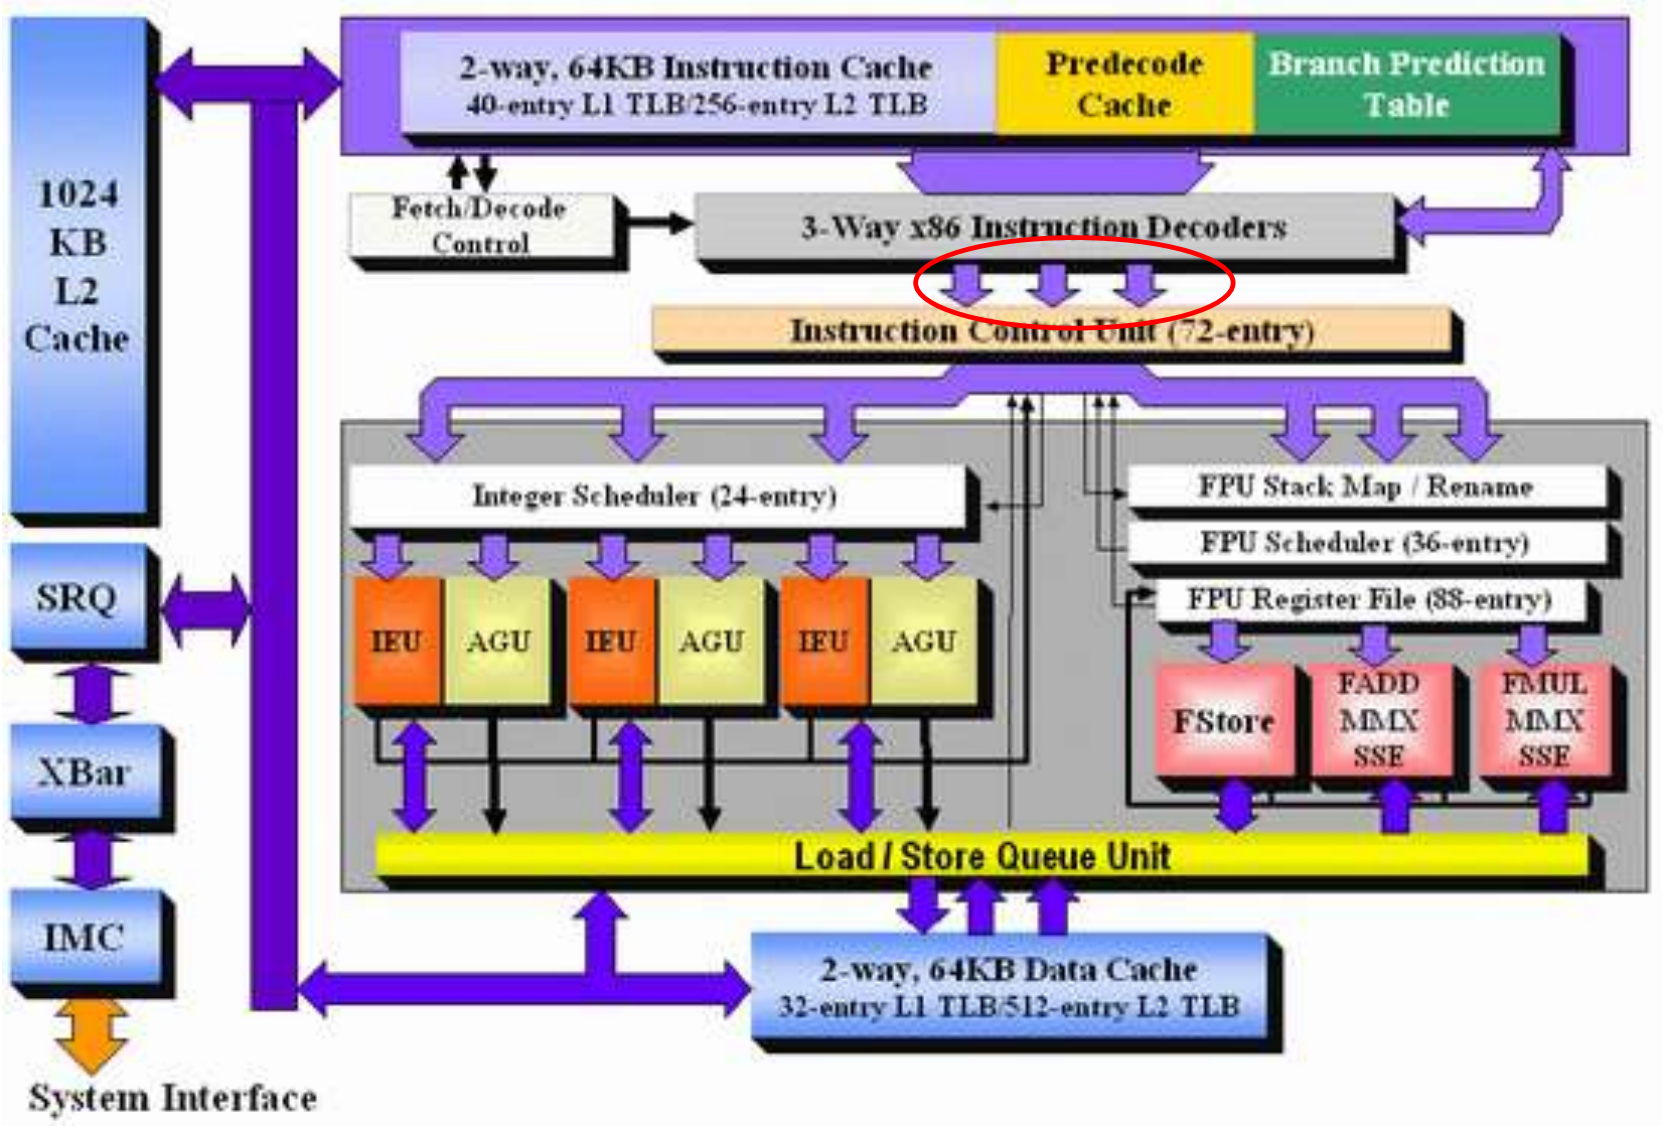

2.2.3 The microarchitecture of the Bulldozer Compute Module (7)

The microarchitecture of a Bulldozer core [10]

- 3. gen. superscalar
- Front-end: 4 wide
- Issue rate to the EUs: 4 + 4 (shared)

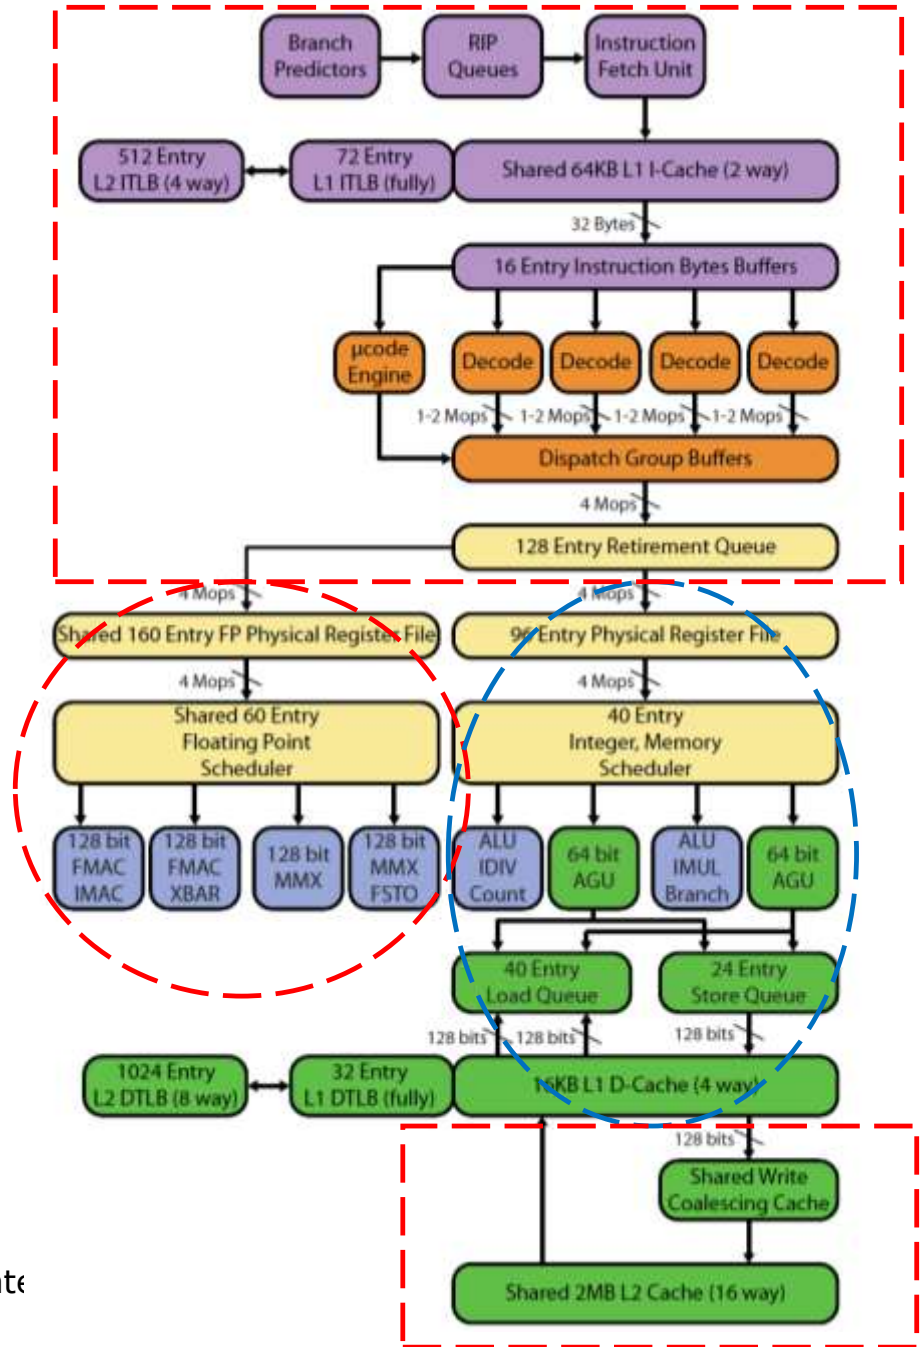

Shared front End

Shared FP unit

Per thread dedicated FX unit

Shared L2

RIP: Relative Instruction Pointe

2.2.3 The microarchitecture of the Bulldozer Compute Module (8)

The microarchitecture of of Intel's Sandy Bridge cores [17]

- 3. gen. superscalar
- Front-end: 4 wide
- Issue rate to the EUs: 6

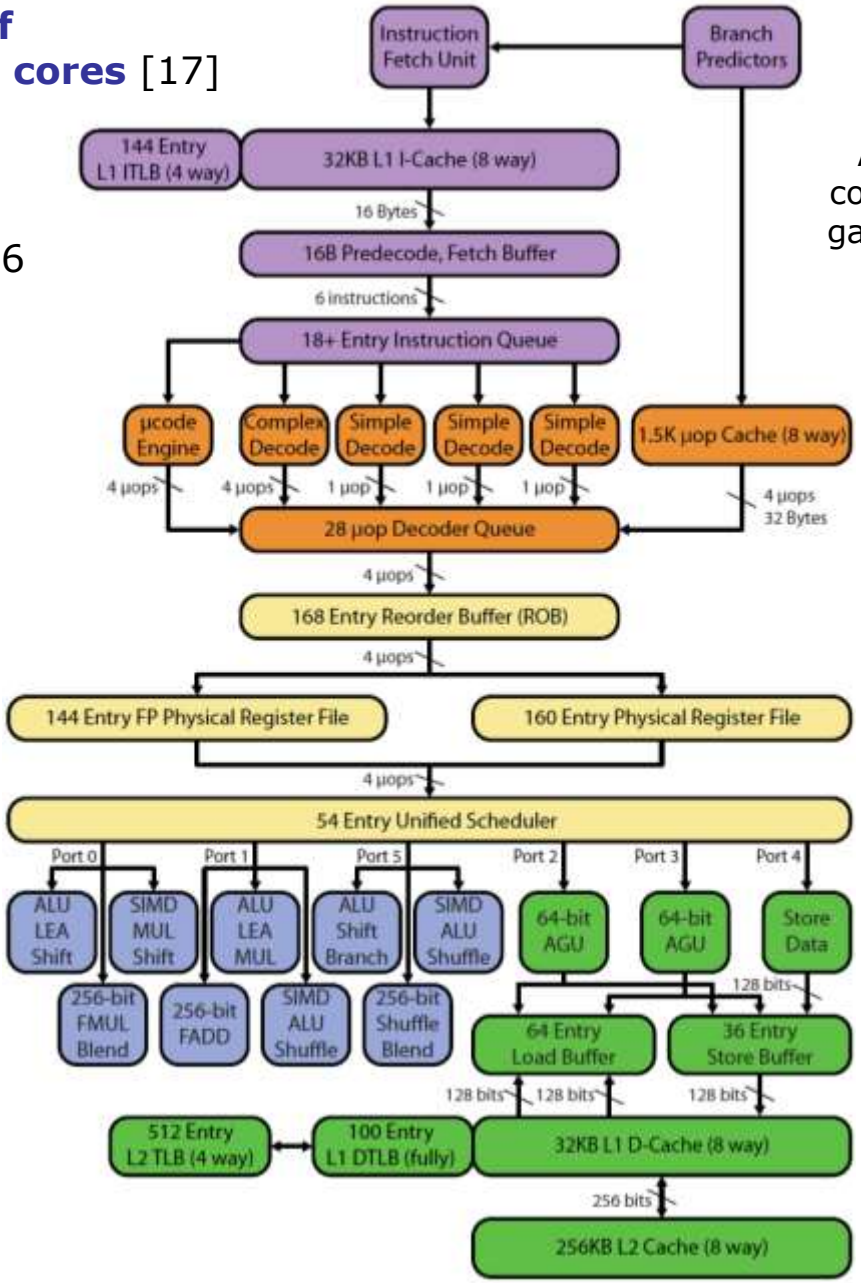

A hit in the uop cache will completely bypass and clock-gate the instruction fetch and decode hardware.

It increases performance and reduces power.

The Decoder Queue performs also loop caching

2.2.3 The microarchitecture of the Bulldozer Compute Module (9)

The microarchitecture of of Intel's Westmere cores [10]

- 3. gen. superscalar
- Front-end: 4 wide
- Issue rate to the EUs: 6

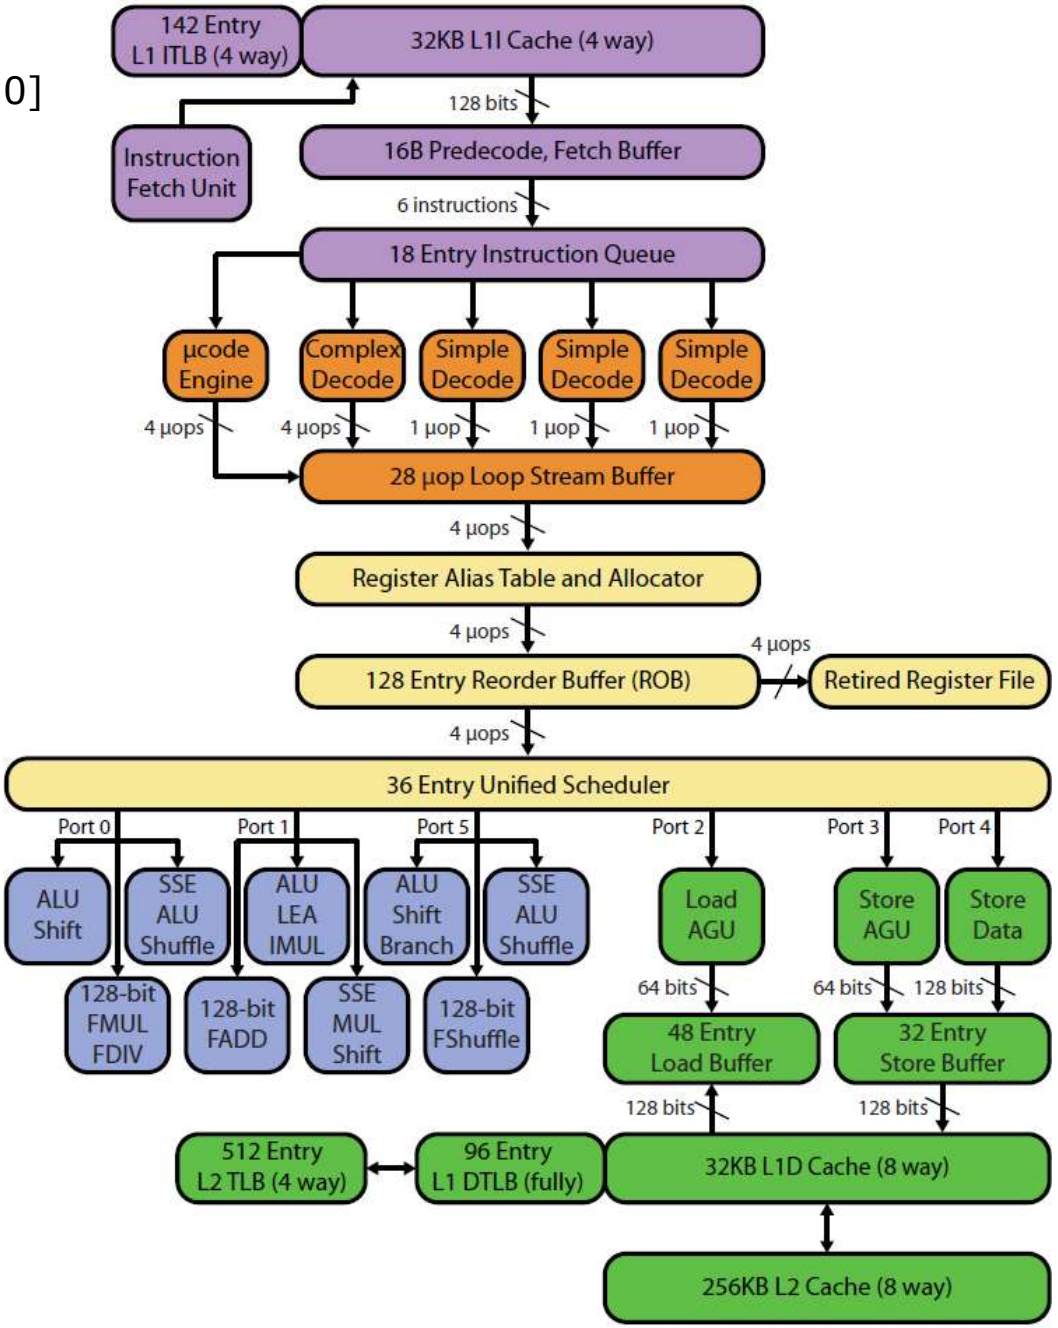

2.2.3 The microarchitecture of the Bulldozer Compute Module (10)

Remark

A very detailed description of Bulldozer's microarchitecture can be found in [10].

2.2.4 Assessing the performance potential of the Bulldozer Compute Module

2.2.4 Assessing the performance potential of the Bulldozer module -1 [3]

1) The FX-part of a Bulldozer core-1

- It can be considered as a peculiar Bulldozer core that shares specific resources (decoding, FP and multimedia processing, L2 cache) with another core.
- A Bulldozer core includes less execution resources than the previous K8-Family 12h cores, as indicated in the next Figure, presumably, in order to reduce power consumption or to optimize performance/power.

Contrasting the execution resources of the Bulldozer core with previous designs

Part of the microarchitecture of a K8/K10/K10.5/Family 12h core [5]

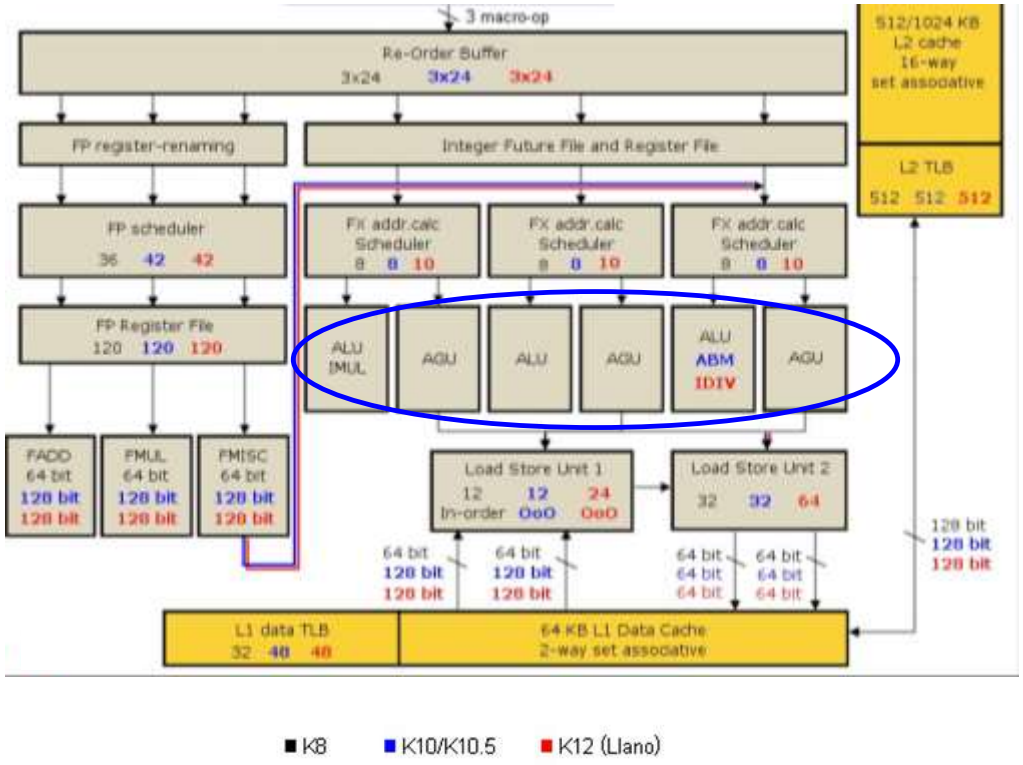

Part of the microarchitecture of a Family 15h Bulldozer core [10]

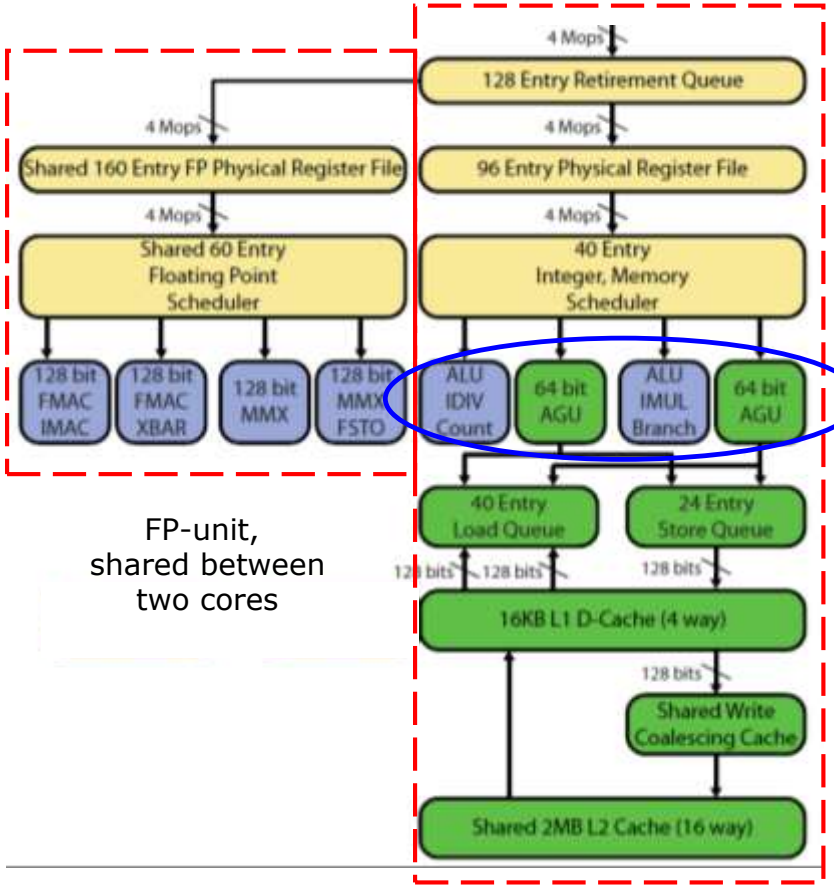

FP-unit, shared between two cores

Per core FX-unit

2.2.4 Assessing the performance potential of the Bulldozer Compute Module (3)

1) The FX-part of a Bulldozer core-2 [3]

Previous K8-Family 12h designs provided basically

- three 64-bit FX ALUs and
- three 64-bit AGUs
  (used as Address Generation Units to calculate memory addresses of load/store operations).

On the other side a Bulldozer core is equipped only with

- two 64-bit FX ALUs and
- two 64-bit AGUs.

As a consequence, a Bulldozer core can execute up to two ALU and up to two AGU operations per cycle, less than previous AMD designs that allowed to perform up to three ALU and up to three AGU operations per cycle.

Assessing the performance potential of the Bulldozer module-2 [3]

2) The FP-part of the Bulldozer module-1

It is shared by two cores and incorporates four 128-bit units.

Contrasting the FP execution resources of the Bulldozer core with previous designs

Part of the microarchitecture of a K8/K10/K10.5/Family 12h core [5]

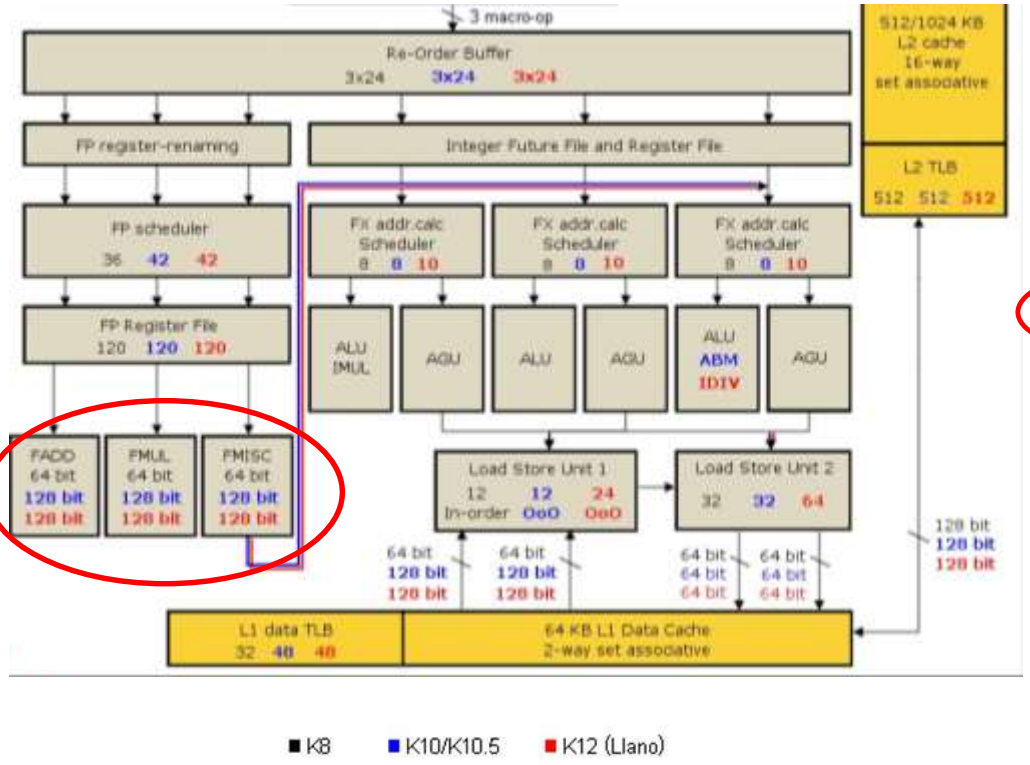

Part of the microarchitecture of a Family 15h Bulldozer core [10]

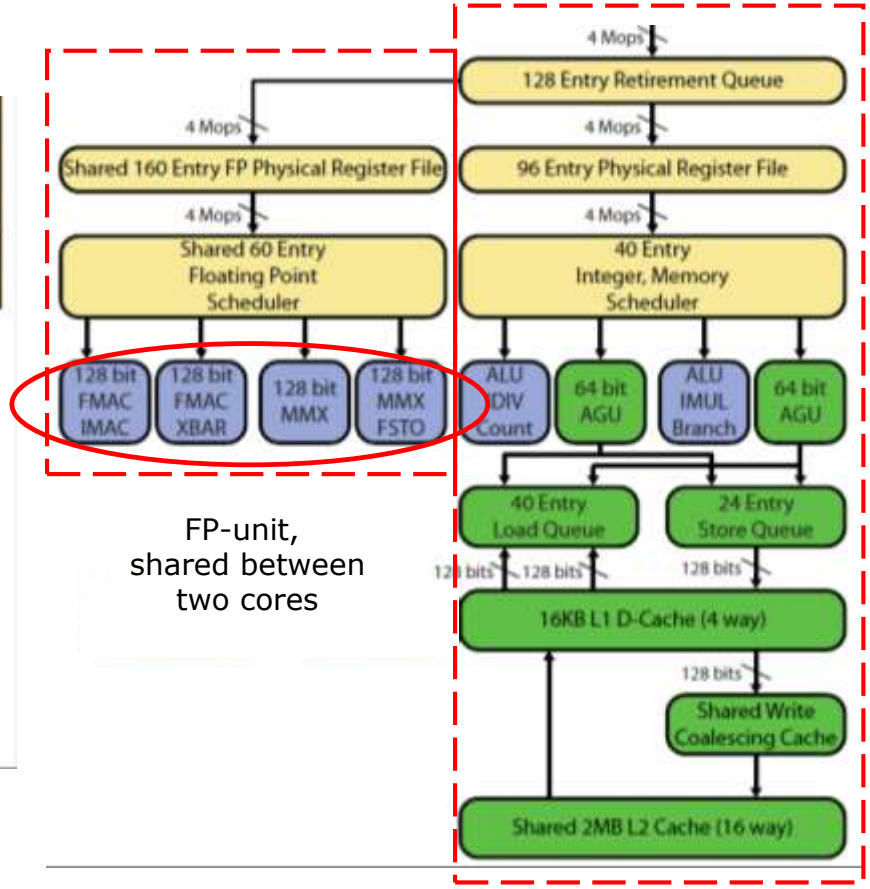

FP-unit, shared between two cores

Per core FX-unit

2.2.4 Assessing the performance potential of the Bulldozer Compute Module (6)

2) The FP-part of the Bulldozer module-2 [3]

From the available four units

- two serve multimedia operations (MMX and SSE) and
- only two can be used for FP operations (FMAC).

The two FMAC (FP Multiply Accumulate) units can be ganged together to execute 256-bit AVX (Advanced Vector Extension) instructions.

2.2.4 Assessing the performance potential of the Bulldozer Compute Module (7)

2) The FP-part of the Bulldozer module-3

On the other hand AMD's K10- Family 12h cores have

- Three 128-bit FP-units, as indicated in the next Figure.

Contrasting the FP execution resources of the Bulldozer core with previous designs

Part of the microarchitecture of a K8/K10/K10.5/Family 12h core [5]

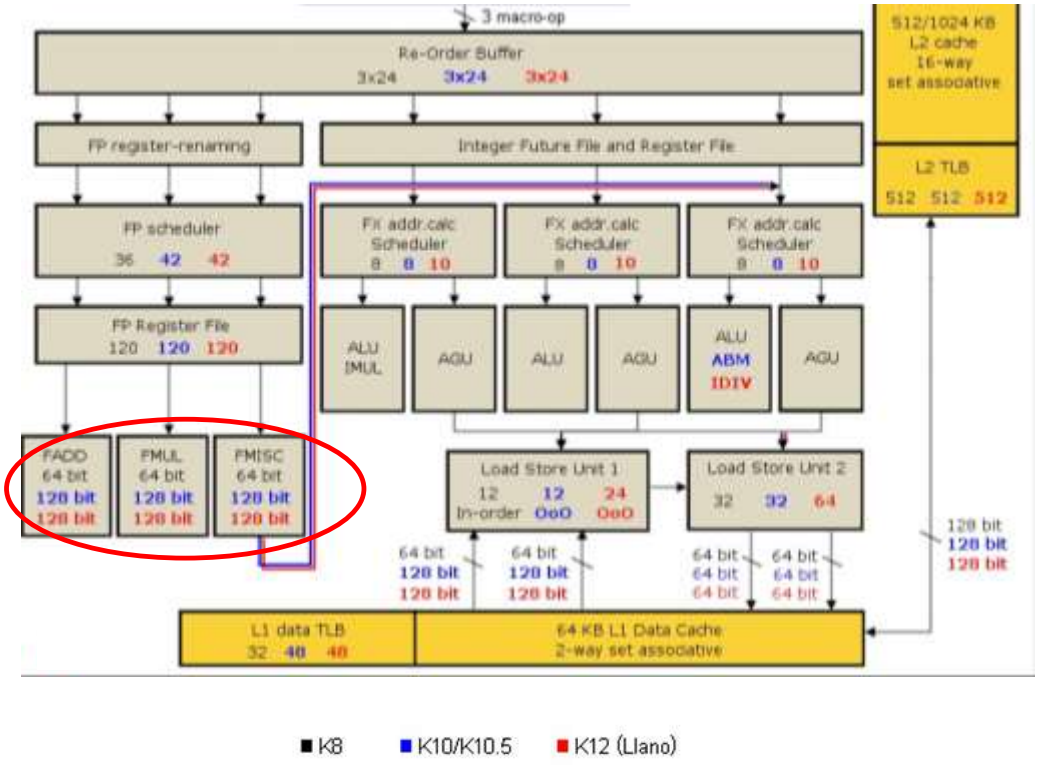

Part of the microarchitecture of a Family 15h Bulldozer core [10]

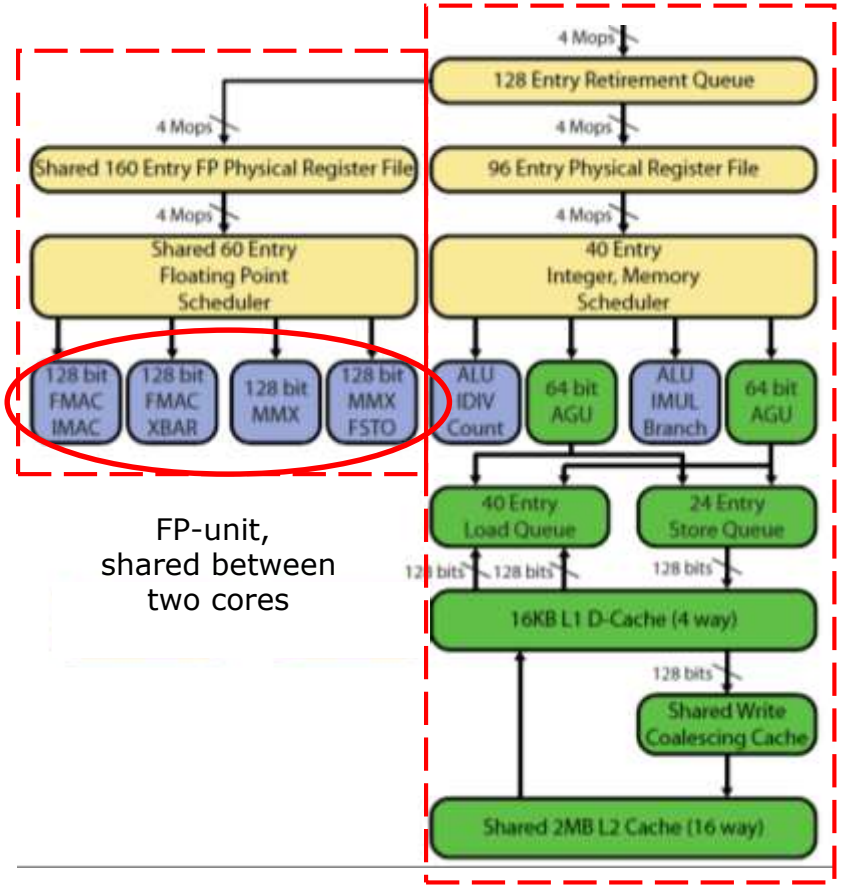

FP-unit, shared between two cores

Per core FX-unit

2.2.4 Assessing the performance potential of the Bulldozer Compute Module (9)

2) The FP-part of the Bulldozer module-4

On the other and AMD's K10- Family 12h cores have

- Three 128-bit FP-units.

Each of AMD's K10-Family 12h cores can perform up to two 64-bit FP operations and beyond that 64-bit MMX or 128 bit SSE operations.

Remark

K8's cores FP-units were only 64-bit wide and each of them could perform only a single FP DP operation.

2.2.4 Assessing the performance potential of the Bulldozer Compute Module (10)

2) The FP-part of the Bulldozer module-5

Comparison of the number of FP DP operations that can be executed per cycle

K10-Family 12h	**Bulldozer**
Up to 3x2 FP DP operations per core	Up to 2x2 FP DP operations per module (two cores)

Obviously, Bulldozer has considerable less per thread available FP execution resources than K10-Family 12h cores, presumably in order to achieve power reduction.

2.2.4 Assessing the performance potential of the Bulldozer Compute Module (11)

3) 256-bit execution resources

Bulldozer makes use of two available 128-bit FMAC units as a 256-bit AVX unit [15] (called by AMD as the FLEX FP).

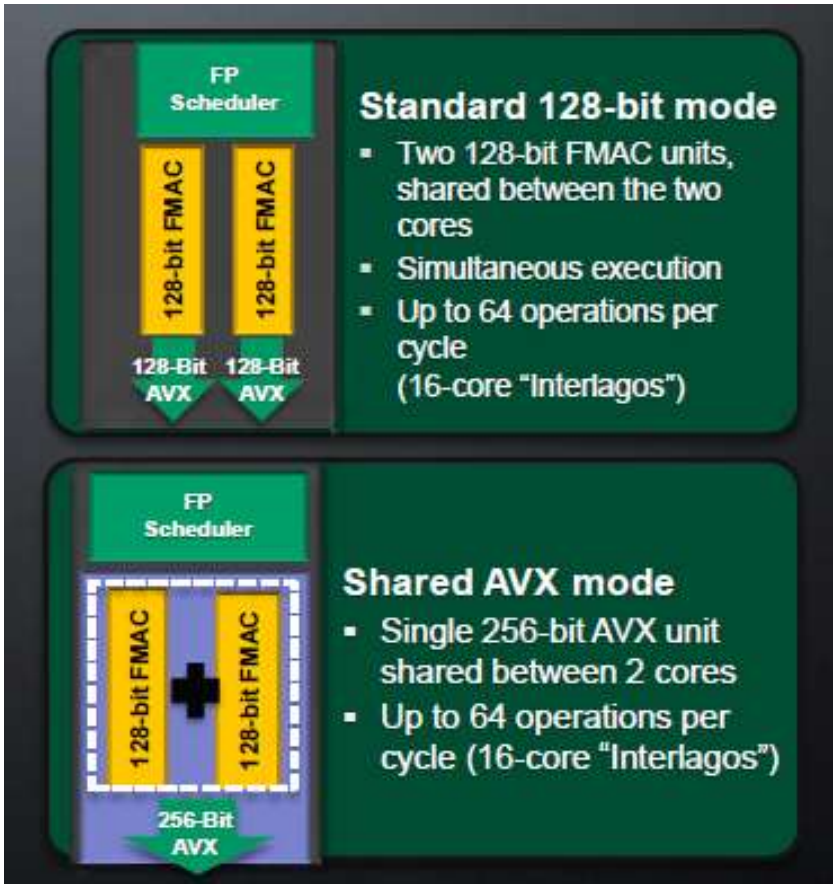

Comparing Bulldozer's per module and Sandy Bridge's per core available 256-bit execution resources-1 [17]

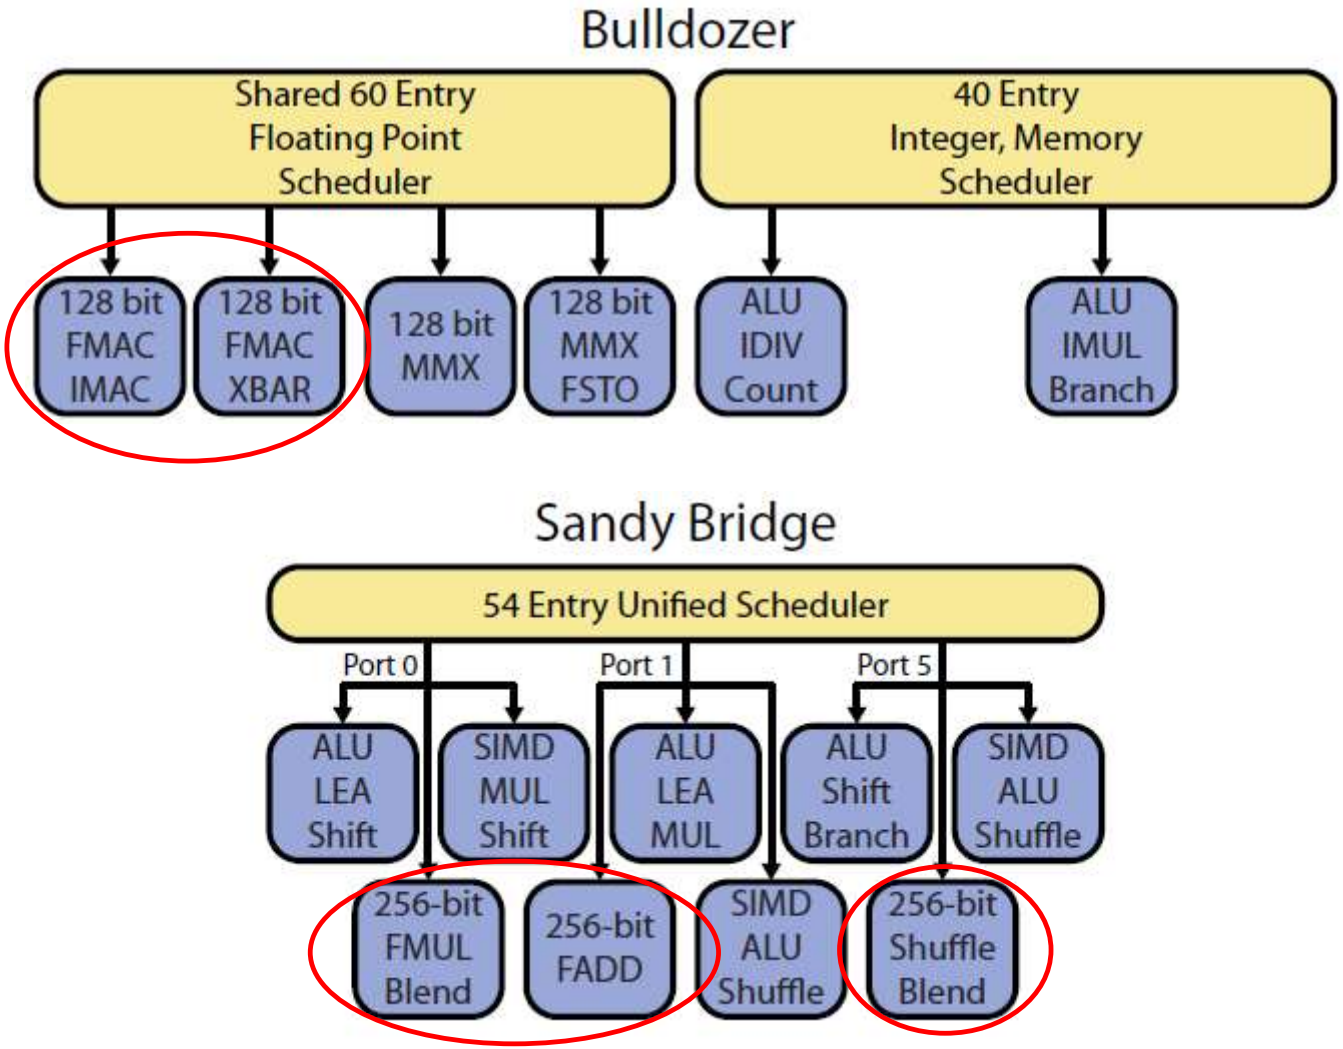

Comparing Bulldozer's per module and Sandy Bridge's per core available 256-bit execution resources-1 [17]

As long as Bulldozer has a single 256-bit execution resource (2 ganged 128-bit FMAC units) per module (two cores)

Intel's Sandy Bridge includes three 256 bit units per core, i.e. it has considerable more 256-bit execution resources.

Assessing the performance potential of the Bulldozer module-3 [3]

4) The pipeline depth of Bulldozer-1

In order to increase single thread performance designers of Bulldozer lengthened its FX pipeline to about 18 or 20 stages compared to 12 stages of the K8-Family 12h designs.

Consequences of the longer pipelines

- Increased penalty of incorrectly guessed branches
- Longer cache and main memory latencies.

2.2.4 Assessing the performance potential of the Bulldozer Compute Module (15)

Cache/main memory latencies of K10/K10.5, Bulldozer and Sandy Bridge processors [3]

Latecny (in cycles)	K10/K10.5	Bulldozer	Sandy Bridge
L1D	3	4	4
L2	14-15	21	11
L3	55-59	65	25
Memory	157-182	195	148

Cache memory latencies can only be assessed however, in relation with cache sizes.

Cache sizes of K10/K10.5, Bulldozer and Sandy Bridge processors

Cache size	K10/K10.5	Bulldozer	Sandy Bridge
L1D	64 KB/core	16 KB/core	32 KB/core
L2 (up to)	512 KB/core	2 MB/module	256 KB/core
L3 (up to)	2/6 MB/proc.	8 MB/4 modules	8 MB

2.2.4 Assessing the performance potential of the Bulldozer Compute Module (16)

4) The pipeline depth of Bulldozer-2 [3]

- As above data shows, Bulldozer has larger L2/L3 caches.
- Bulldozer's larger L2/L3 caches vs. their previous designs as well as its higher clock speed gave rise to higher access latencies.
- Larger caches result in less cache misses but cause higher access latencies that impede IPC.
- AMD's decisions related to the Bulldozer design, including the module concept, the trade-offs concerning pipeline length, cache sizes, and cache and memory latencies are questionable.

In Section 1. we compare the achieved performance of Bulldozer-based server and DT designs with AMD's previous K10.5 (Istambul)-based designs, the results show a moderate about 10 – 20 % increase in clock frequencies despite using 32 nm technology instead of 45 nm.

As far as the module concept is concerned, in their Zen-based processor lines AMD left this solution and makes use of quad-core complexes including private L1 and L2 caches as well as a shared L3 cache (see the related Chapter).

AMD's projection to increase performance per Watt in post Bulldozer architectures [19]

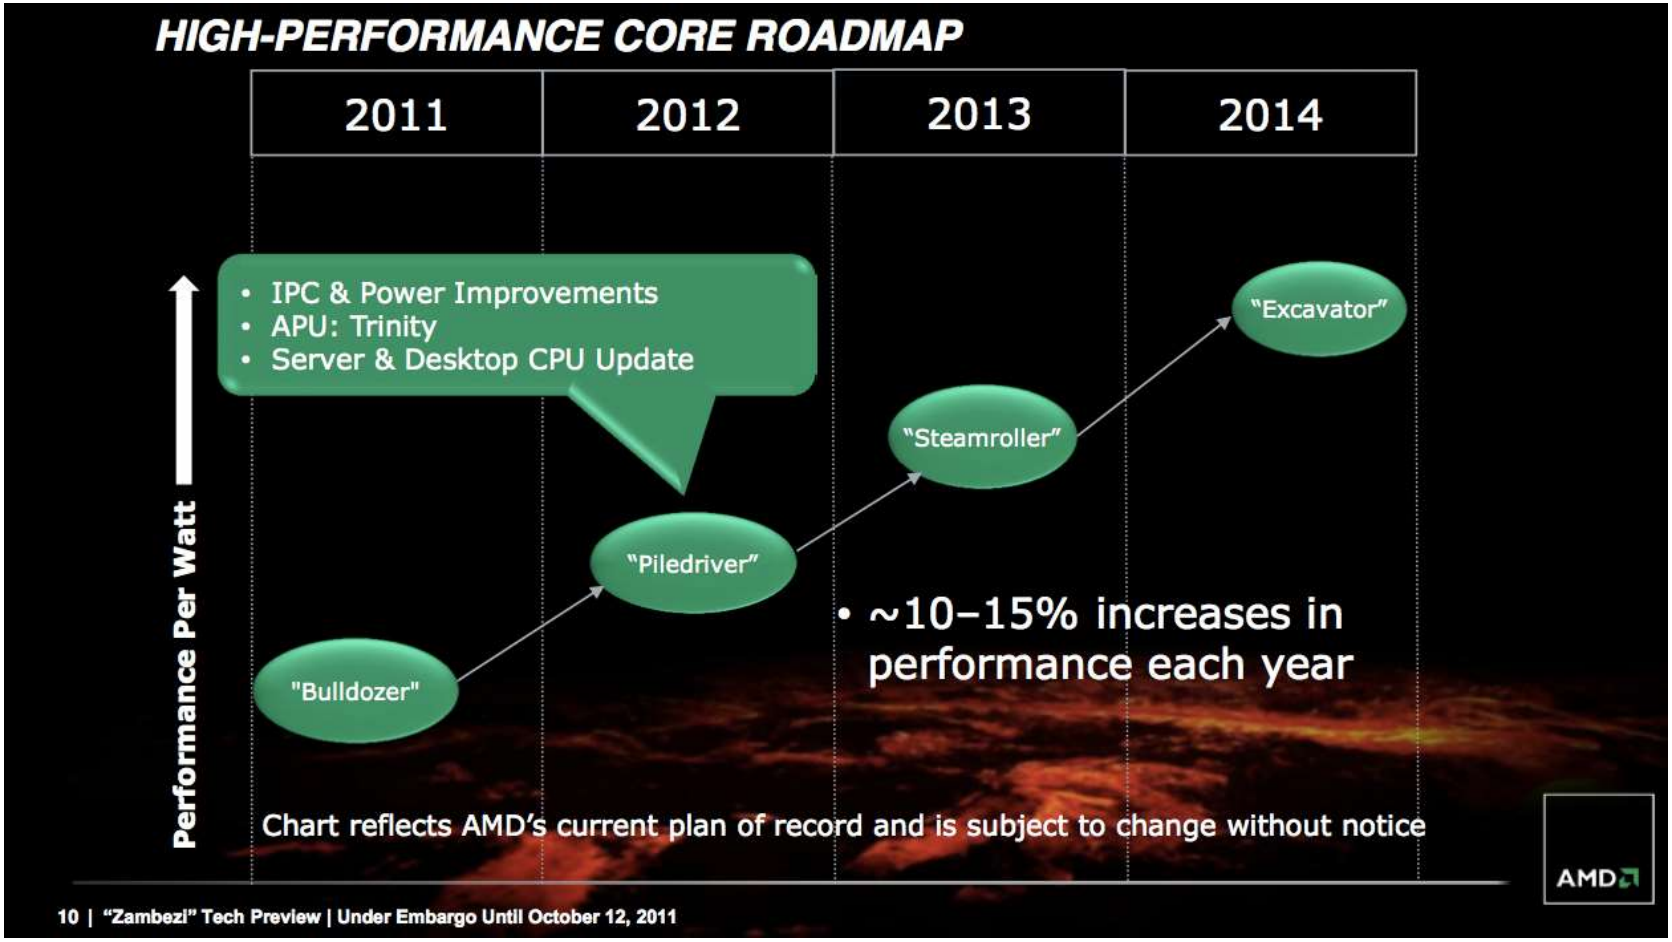

With the above slow rate of performance increase it is strongly questionable whether AMD will able to catch up ever with Intel's future processor lines.

2.3 The Orochi die

2.3 The Orochi die

The high level building block of this family is the Orochi die.
It incorporates 4 Bulldozer modules, as indicated below [4].

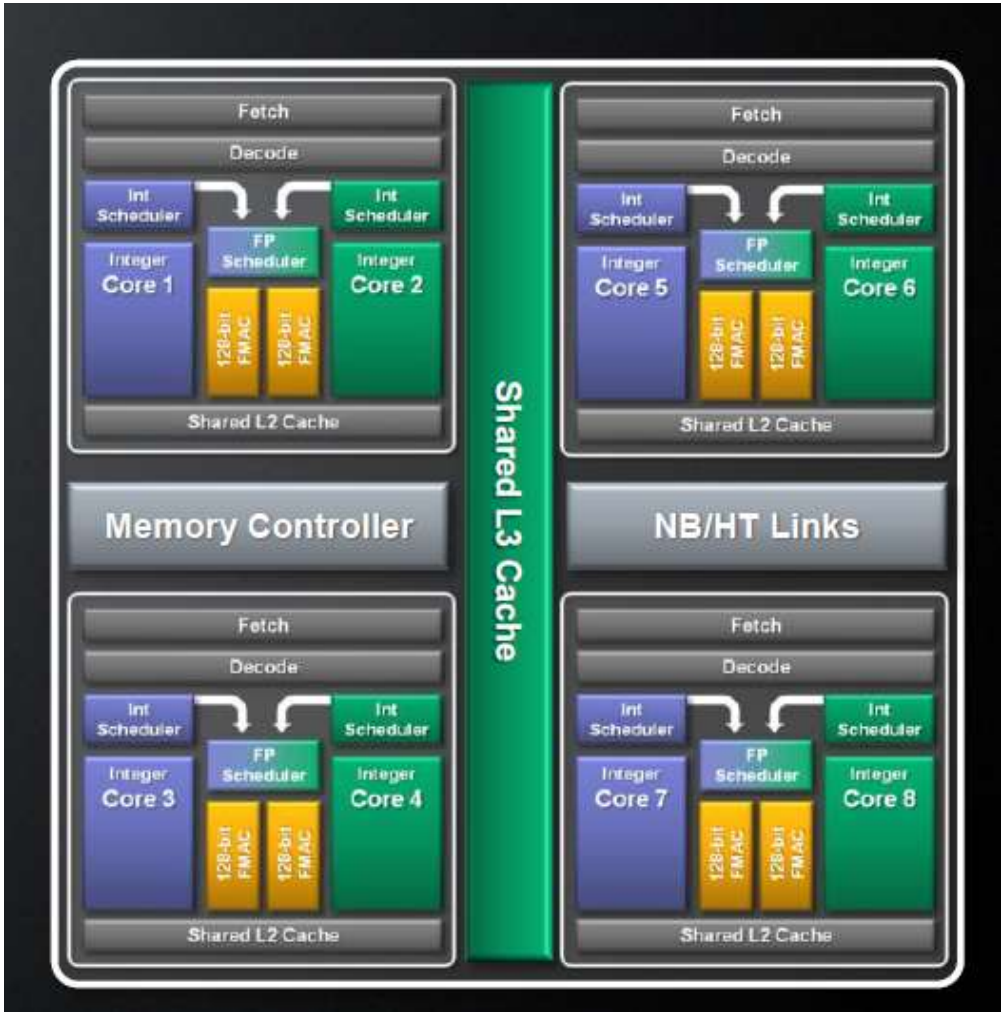

2.3 The Orochi die (2)

Floor plan of the Orochi die

1. gen. Bulldozer-based designs are built up of Orochi dies, each including 4 Bulldozer modules.

The Orochi die [6]

32 nm
1.2 billion transistors
315 mm²
1 MB L2/core
8 MB L3

Servers

Interlagos: 2 Orochi dies
(2x4 modules/2x8 cores)

Valencia: 1 Orochi die
(4 modules/8 cores)

Desktops

Zambezi: 1 Orochi die
(4 modules/8 cores)

2.3 The Orochi die (3)

Main parameters of an Orochi die

- 32 nm feature size
- 1.2 billion transistors
- 315 mm^2
- 1 MB L2/core
- 8 MB L3

Bulldozer-based processors are built up of one or two Orochi dies as follows:

Servers

Interlagos: 2 dies (16 cores) implemented as a Multi-Chip Module (MCM)
Valencia: 1 die (8 cores)

Desktops

Zambezi: 1 die (8 cores)

The Orochi die

- It is the high level building block of Bulldozer-based processors lines.
- It includes 4 Bulldozer module, equaling 8 cores.

AMD's 8 core Orochi die [20]

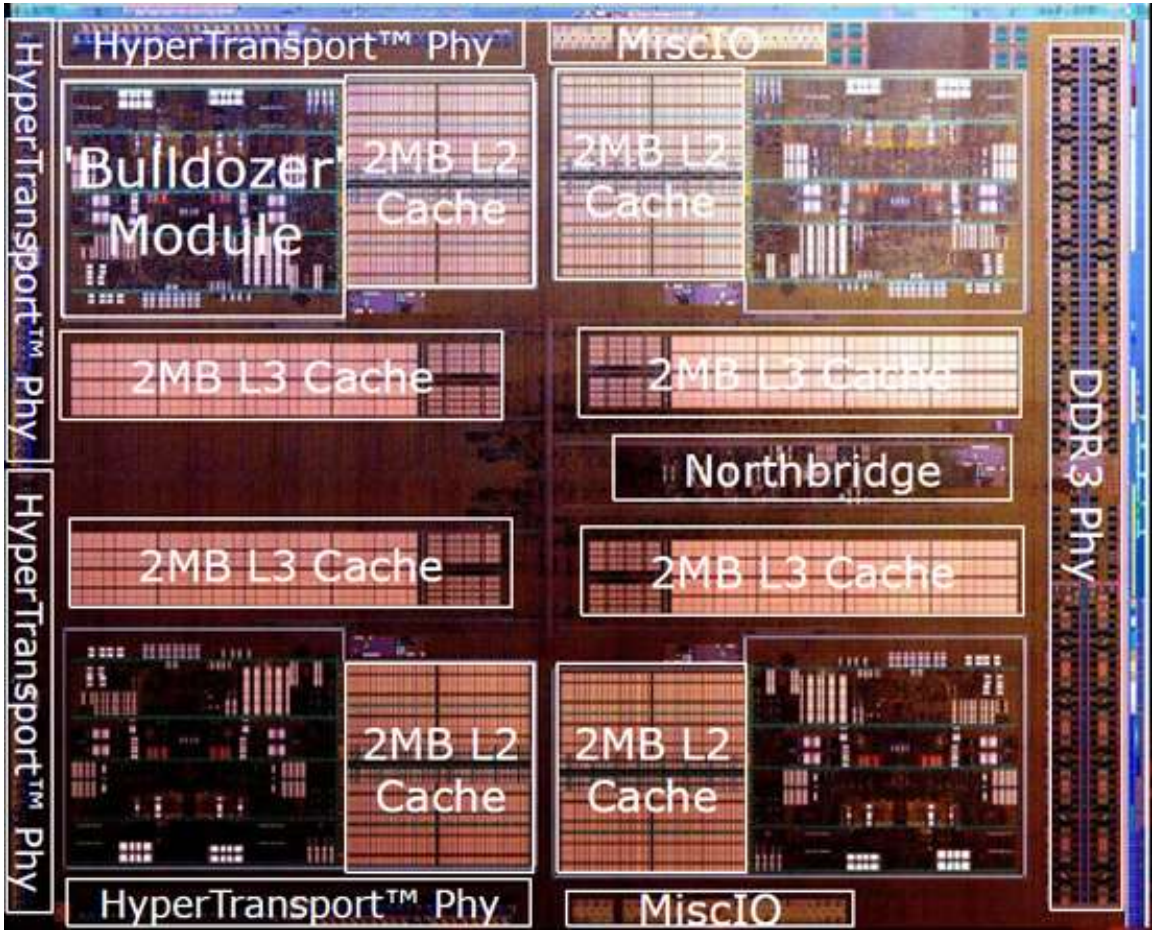

32 nm
1.2 billion transistors
315 mm^2
1 MB L2/core
8 MB L3

The North Bridge of Orochi [21]

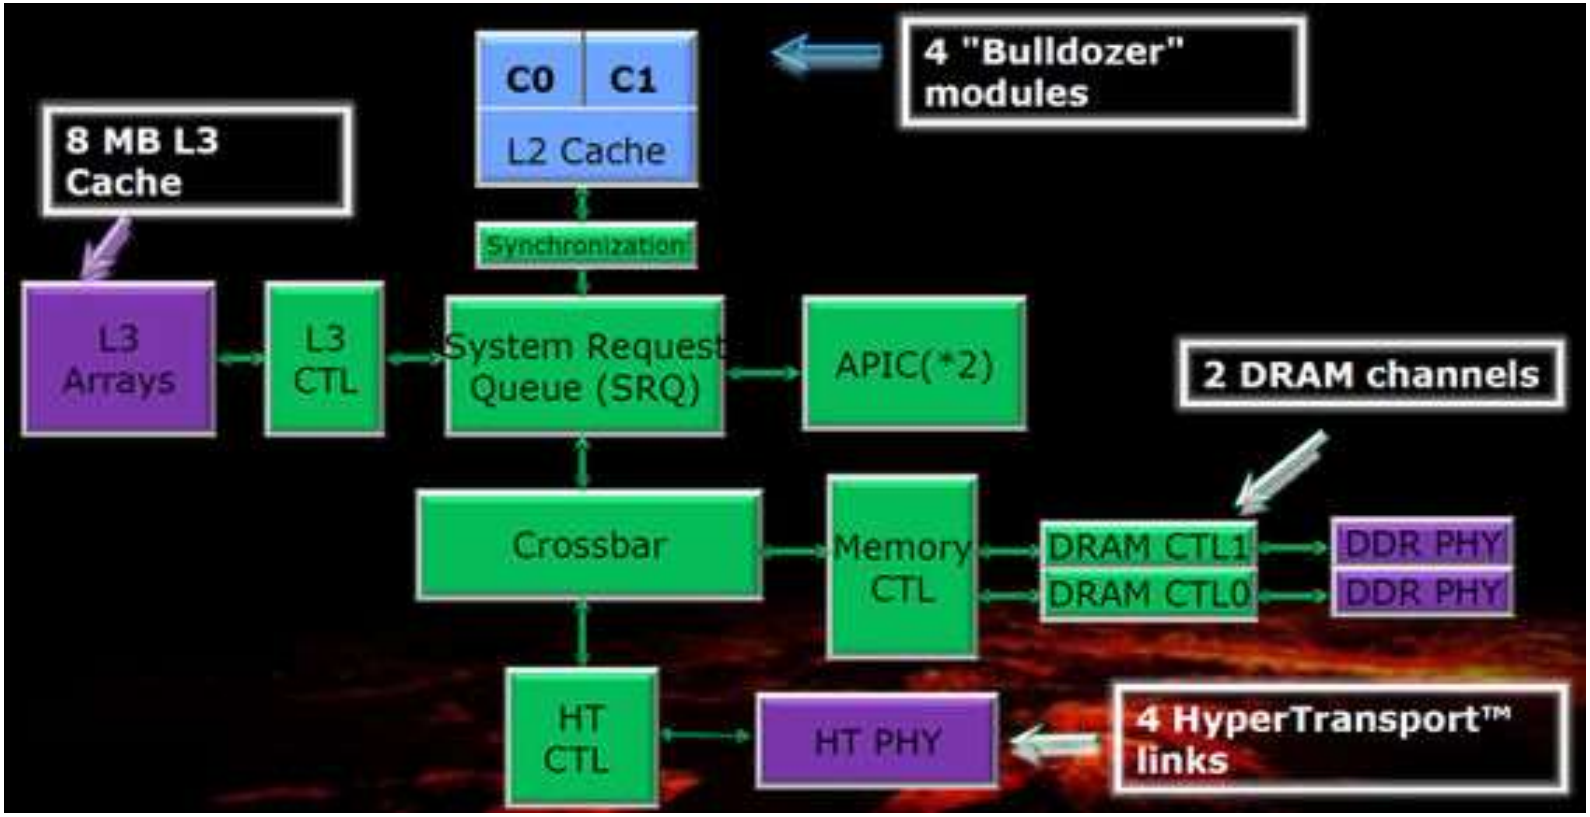

2.3 The Orochi die (6)

Block diagram of the Orochi die

It incorporates 4 Bulldozer-modules (8 cores) [22]

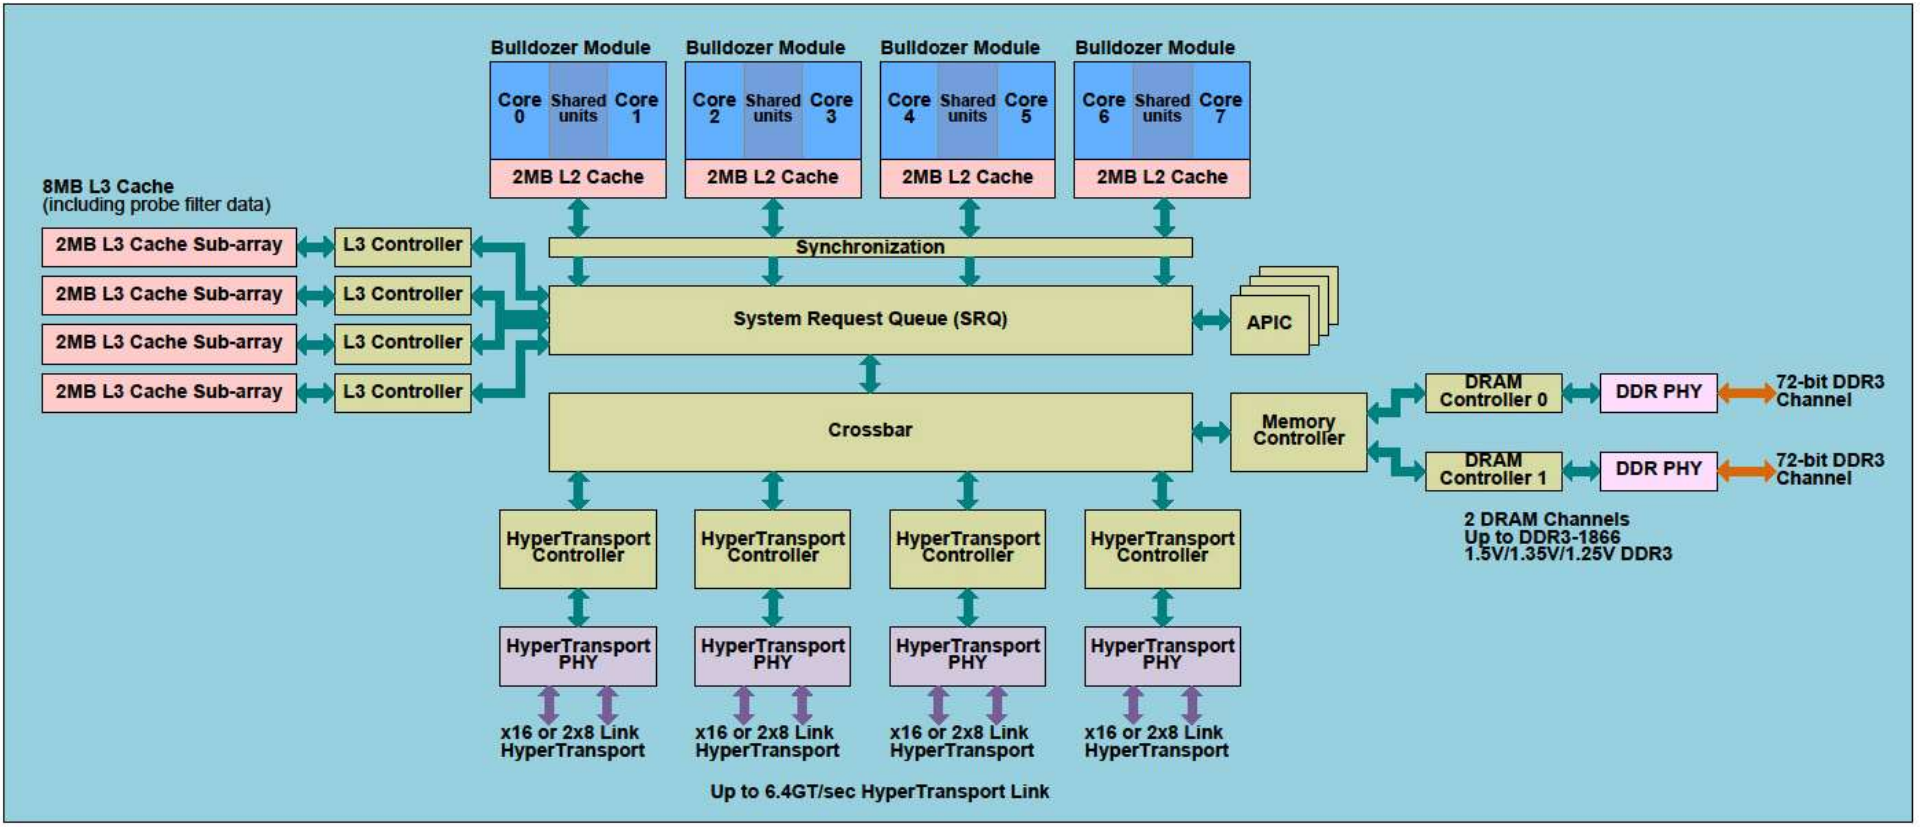

Use of the Orochi die

In servers

- Interlagos (16 cores): 2 dies
- Valencia (8 cores): 1 die

In desktops

- Zambezi (8 cores): 1 die

2.4 New power management features of the Bulldozer design

2.4 New power management features of the Bulldozer design – Overview (based on [4])

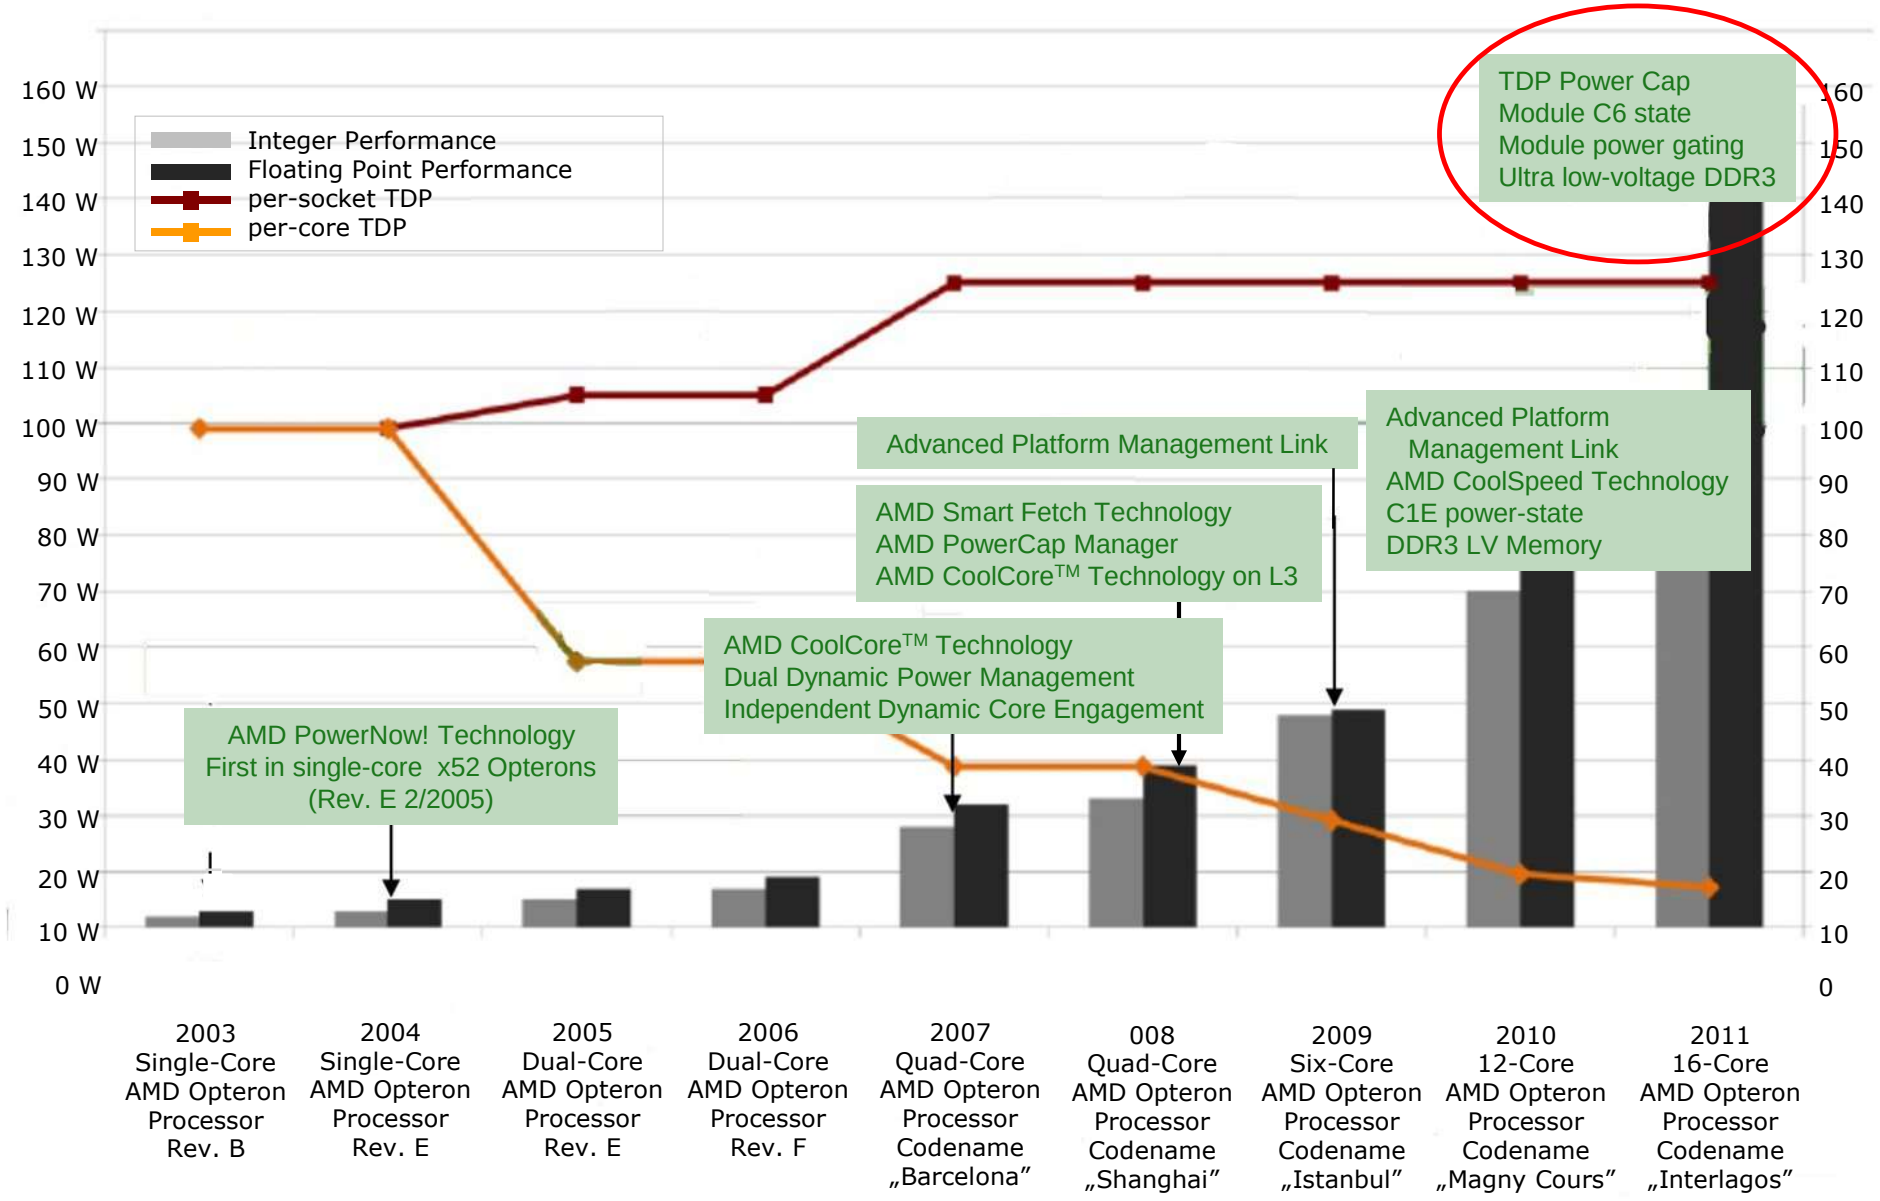

2.4 New power management features of the Bulldozer design (3)

New power management features of the Bulldozer design

- TDP Power Cap
- Module C6 state
- Module level VSS power gating
- Ultra LV-DDR3 support

2.4 New power management features of the Bulldozer design (4)

TDP Power Cap [23]

- Power Capping was introduced in K10.5 Shanghai based servers to set a power limit by setting the max. P-state via BIOS.

  This kind of operation however restricts the processor from using the highest clock frequencies that are associated with the locked out P-states.

  This results in longer response or run times.

- TDP Power Cap however, allows users to restrict power consumption without capping clock frequencies.

  Then while the processor runs under normal circumstances (e.g. at 40-70 % of its full load) the response or run times remain about the same as without power capping.

  The max. TDP can be set either via BIOS or APML.

2.4 New power management features of the Bulldozer design (5)

Module C6 state [24], [6]

(designated as Core C6 state or CC6 state by AMD)

The related BIOS and Kernel Developer's Guide (BKDG) and most AMD literature designates the Module C6 state as Core C6 state.

In the Module C6 state

- the L1 data caches of both cores and the shared L1 instruction cache and the L2 cache of the module are flushed into the L3,
- the module state (register state) is dumped to DRAM, and
- VSS is power gated.

Entering the Module C6 state

When both Bulldozer cores of a module enter an idle state (non C0 state, like C1 to C3) and the condition for flushing the L1/L2 caches remains valid for a preset period of time (checked by a counter)

- the L1/L2 caches of the module are flushed to the L3 cache,
- the internal state of the module is dumped to the DRAM and
- VSS of the module becomes power gated.

Module level VSS power gating results by approximately 95 % reduction of the leakage power [25].

2.4 New power management features of the Bulldozer design (6)

Exiting the Module C6 state

It happens in reverse sequence than entering into the Module C6 state.

2.4 New power management features of the Bulldozer design (7)

Remark [15]

- Entering a Core C6 state with power gating would be possible only for the components which are dedicated for a core, such as the integer unit and the L1 data cache.
- Shared components can be power gated obviously only at the module level.

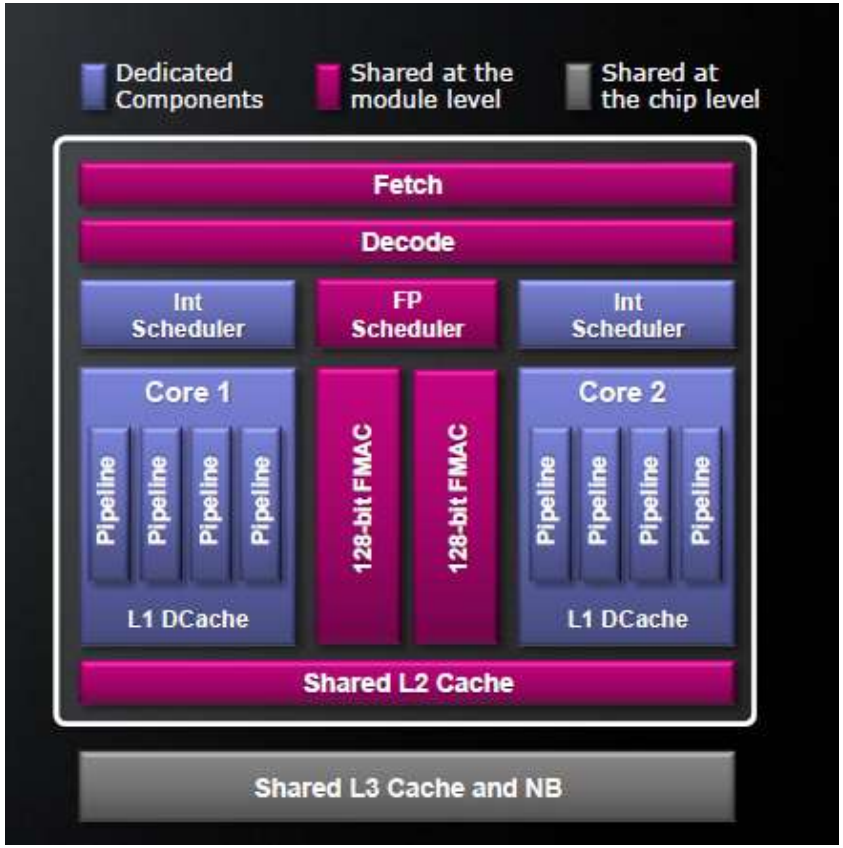

2.4 New power management features of the Bulldozer design (8)

Module level VSS power gating

The last step of entering the Module C6 state is power gating of the module.

A Bulldozer module will be power gated by a dedicated power gating ring that isolates the core VSS from the real VSS [6], as detailed for the Fam. 12h Llano processor in Section 11.

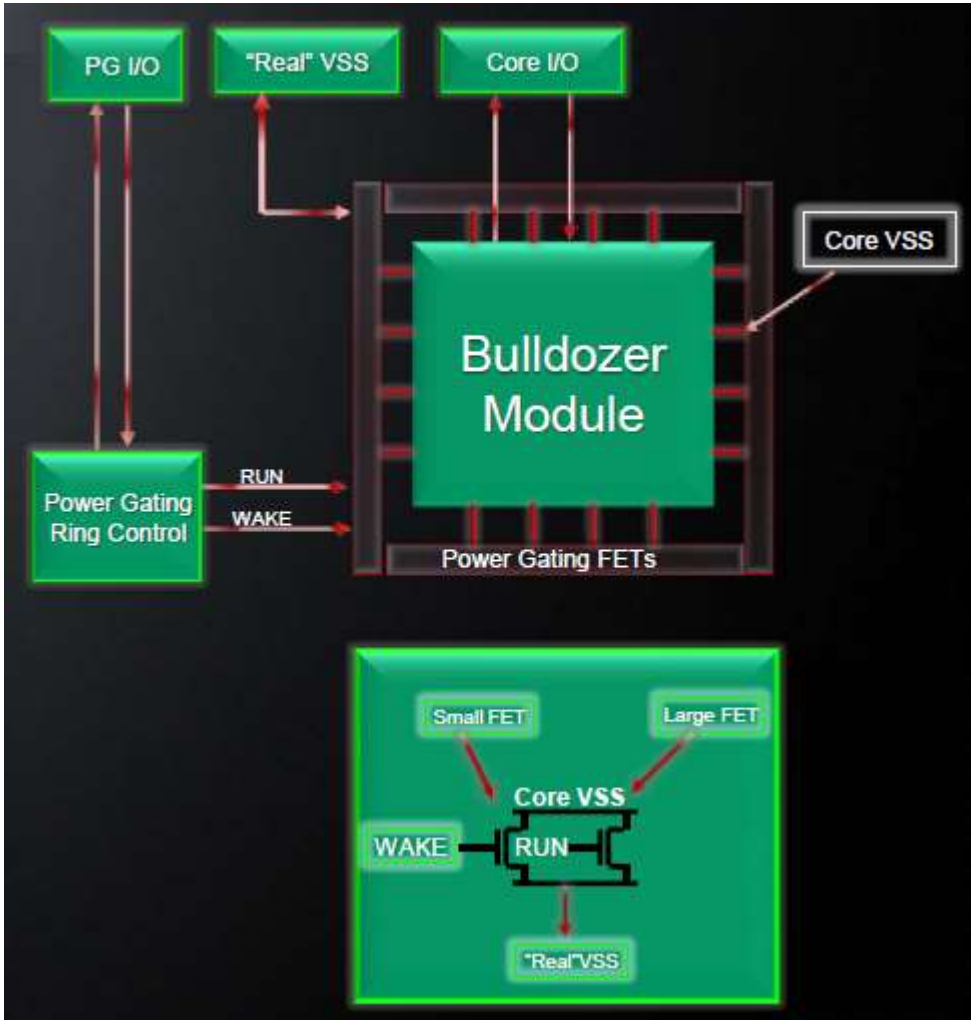

Benefit of module level power gating (C6) vs. C1E state [7]

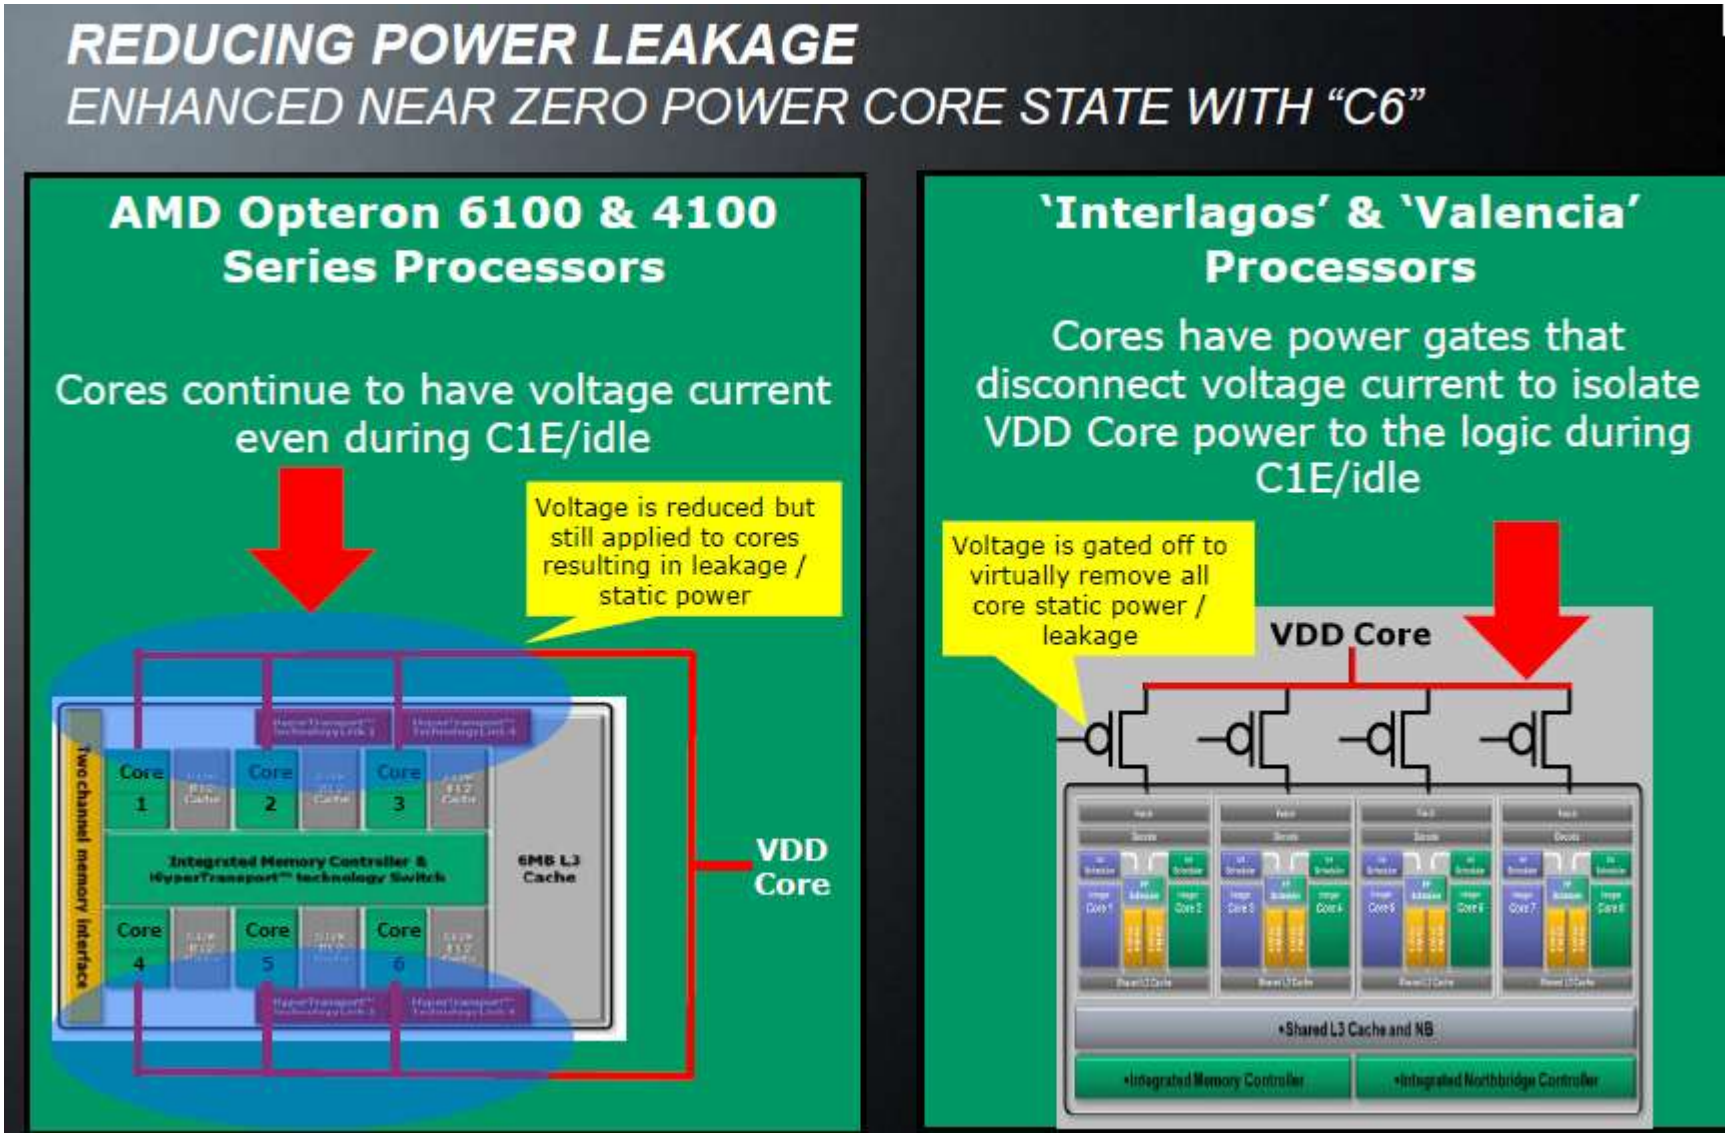

Contrasting the Smart Fetch technique with entering the Module C6 state [7]

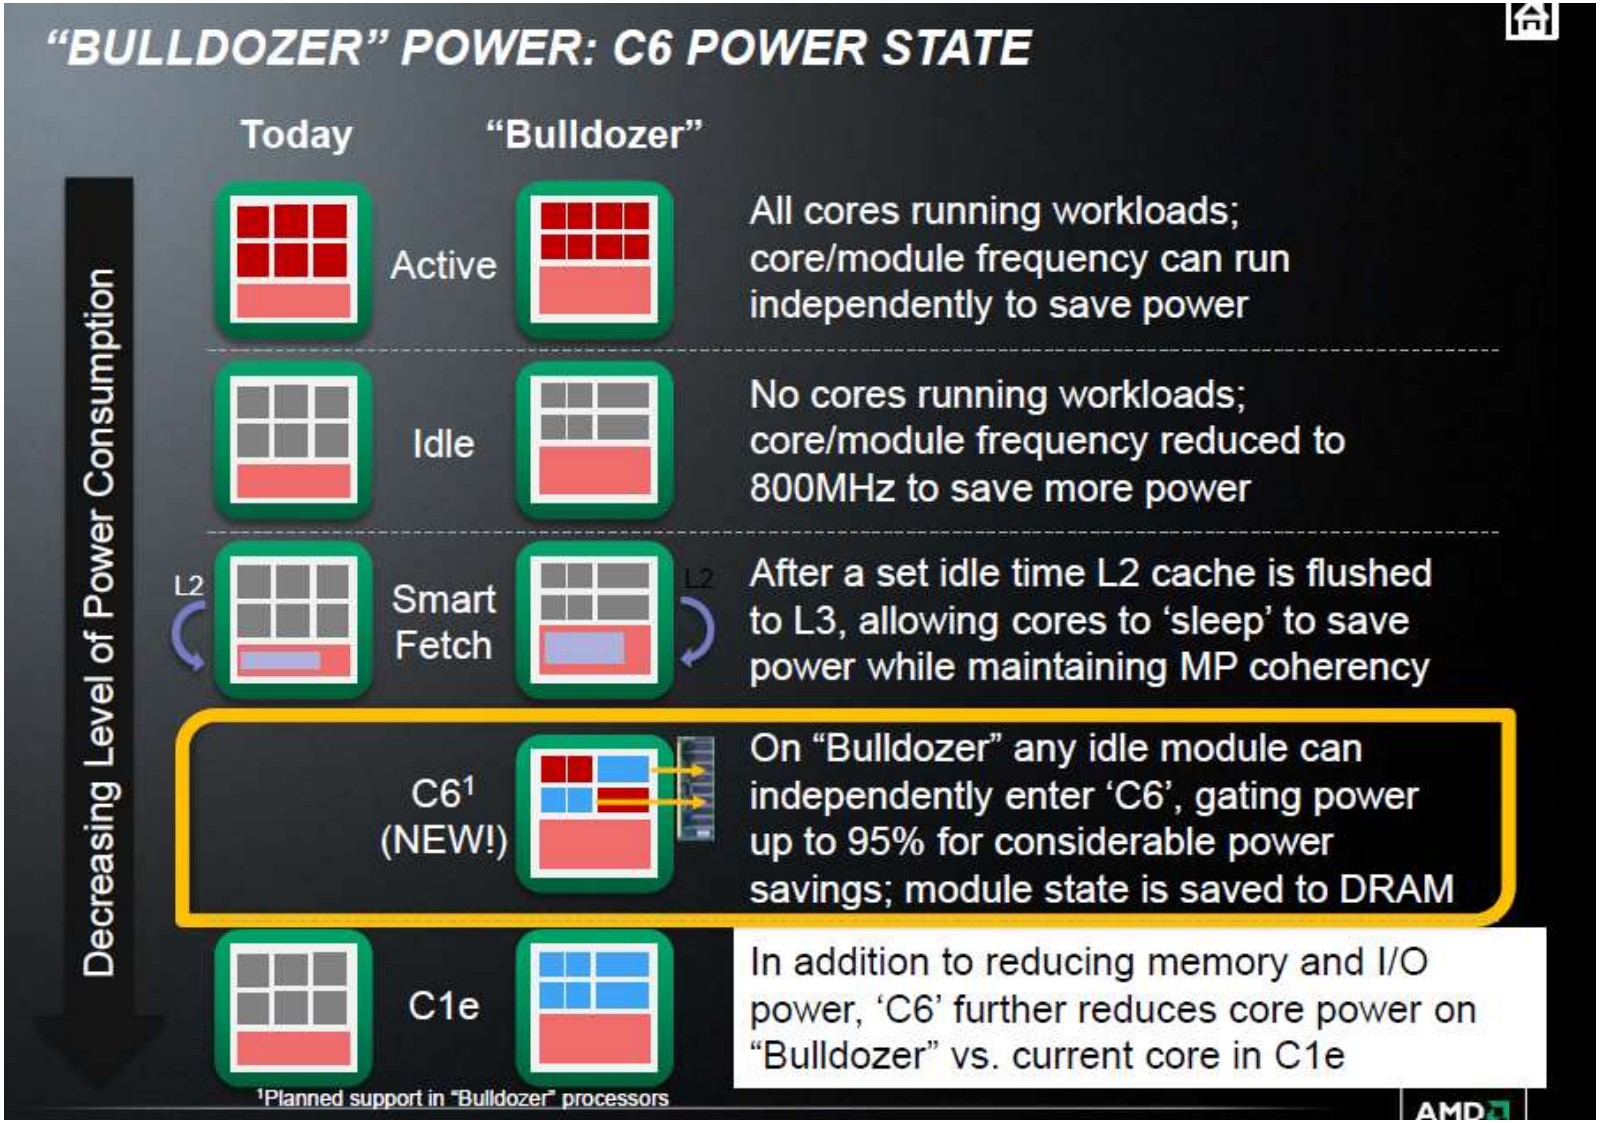

2.4 New power management features of the Bulldozer design (11)

LV-DDR3 support

LV-DDR3 support was already introduced for K10.5 Magny Course servers for 1.35 V low-voltage DDR3 devices.

LV-DDR3 support is now extended for 1.25 V ultra low-voltage DDR3 devices as well.

2.4 New power management features of the Bulldozer design (12)

Remark

Summary of AMD's power management techniques used in the Family 15h [15]

Feature	How it reduces power consumption
TDP Power Cap **New!**	Flexibility to set power limits without capping processor frequency
C6 **New!**	Reduces idle power static leakage up to 95%* by shutting off power to an inactive Bulldozer module
Ultra low power and flexible memory support options **New!**	Supports DDR3ULV 1.25, low power DDR3L 1.35v, and DDR3 1.5v memory technologies
AMD PowerCap Manager	Set fixed performance and power limit on server's processor power consumption
AMD Smart Fetch Technology	Allows idle cores to enter "halt" state
AMD CoolCore™ Technology	Dynamically turning off sections of processor when inactive
AMD CoolSpeed Technology	Highly accurate thermal information & thermal protection
Advanced Processor Management Link	Advanced power control and thermal policies
AMD PowerNow!™ Technology with Independent Dynamic Core Technology	Dynamically operate at lower power and frequencies
Dual Dynamic Power Management	Separates power planes for cores and memory controller
C1E	Reduces memory controller and Hypertransport™ technology links' power

2.5 Bulldozer-based server lines

- 2.5.1 Overview of the Bulldozer-based server lines
- 2.5.2 The Interlagos MP server line
- 2.5.3 The Turbo core technology of Bulldozer-based MP servers
- 2.5.4 The Valencia DP and Zurich UP server lines

2.5.1 Overview of the Bulldozer-based server lines

2.5.1 Overview of the Bulldozer-based server lines-1 [Based on 1]

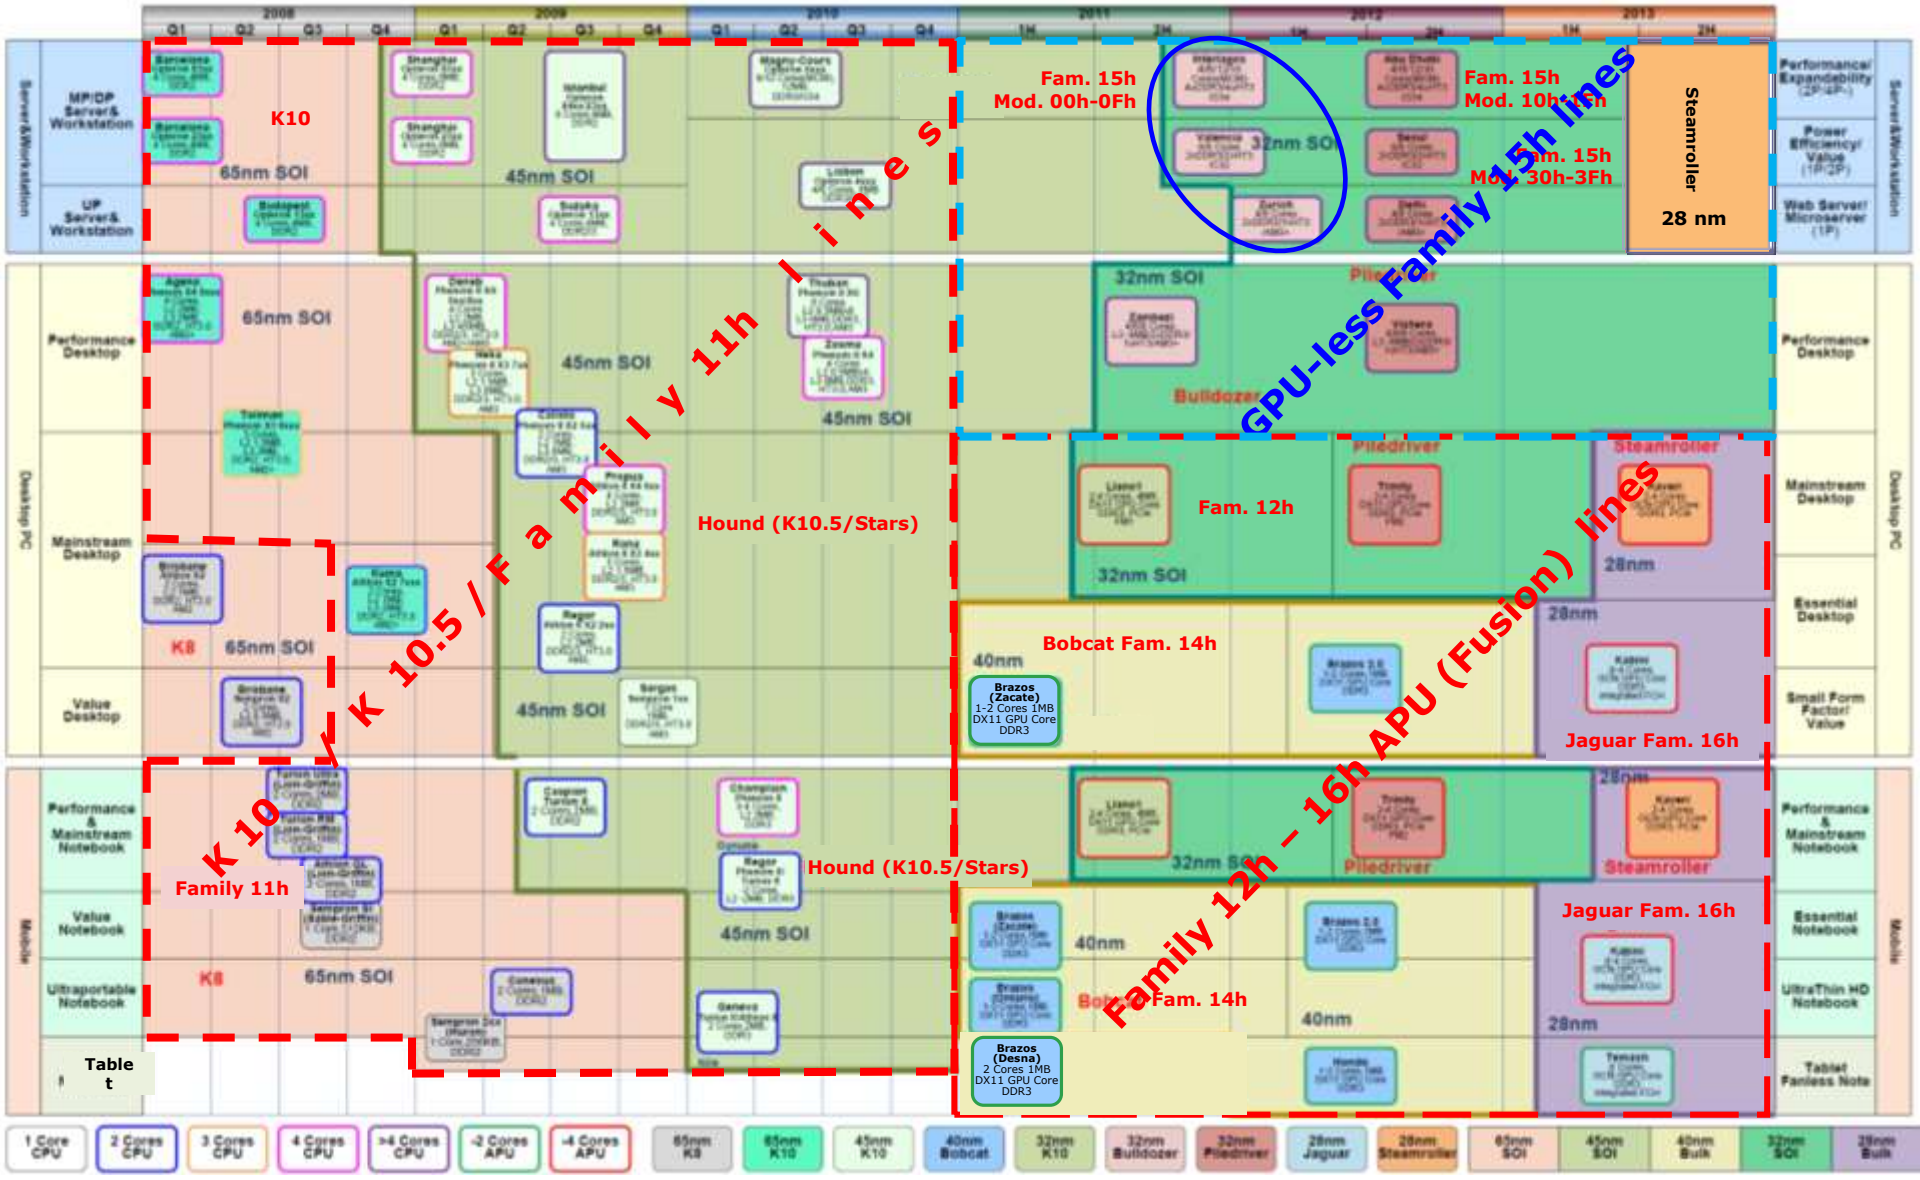

Overview of the Bulldozer-based server lines-2 [Based on 1]

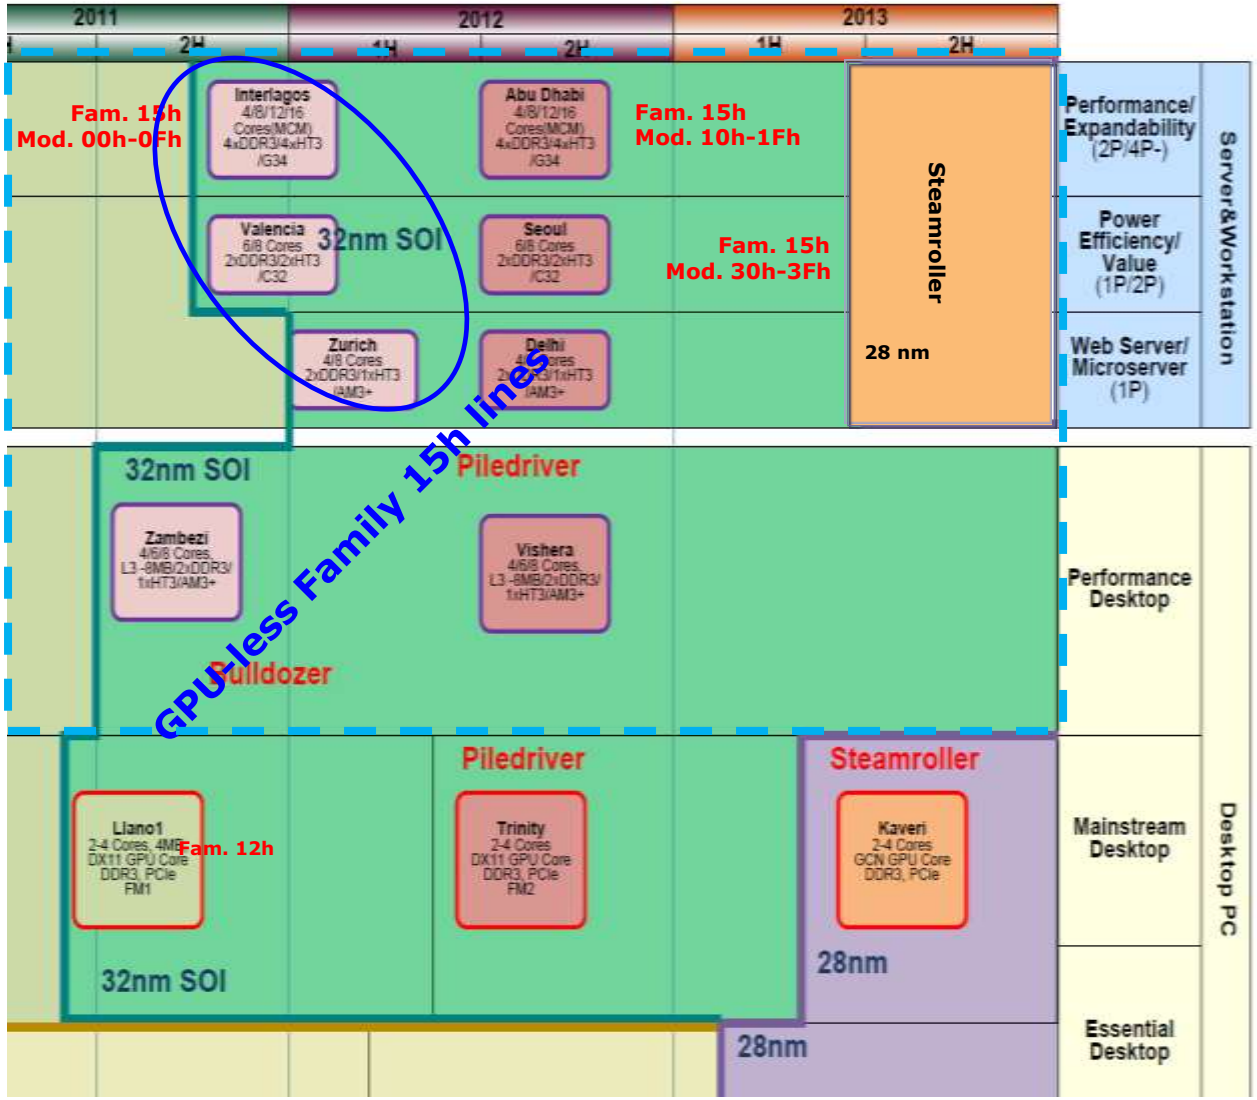

2.5.2 The Interlagos MP server line

2.5.2 The Interlagos MP server line

Base arch./ stepping		Intro	4P Server family name	Series	Techn.	Cores (up to)	L2 (up to)	L3 (up to)	Memory (up to)	HT/ dir. (up to)	Sock et
K8	C0/CG	4/2003	Sledge-hammer	800	130 nm	1C	1 MB	-	DDR-333	HT 1.0: 3.2 GB/s	940
	E4/E6	12/2004	Athens	800	90 nm	1C	1 MB	-	DDR-400	HT 2.0: 4.0 GB/s	940
	E1/E6	4/2005	Egypt	800	90 nm	2C	2*1 MB	-	DDR-400	HT 2.0: 4.0 GB/s	940
	F2/F3	8/2006	Santa Rosa	8200	90 nm	2C	2*1 MB	-	DDR2-667	HT 2.0: 4.0 GB/s	F
K10	BA/B1-B3	8/2007	Barcelona	8300	65 nm	4C	4*1/2 MB	2 MB	DDR2-667	HT 2.0: 4.0 GB/s	F
K10.5	C2/C3	11/2008	Shanghai	8300	45 nm	4C	4*1/2 MB	6 MB	DDR2-800	HT 2.0/3.0: 4.0/8.8 GB/s	F
	CE	6/2009	Istambul	8400	45 nm	6C	6*1/2 MB	6 MB	DDR2-800	HT 3.0: 9.6 GB/s	F
	D1	3/2010	Magny Course (2xIstambul)	6100	45 nm	2x6C	12*1/2 MB	6 MB	DDR3-1333	HT 3.1: 12.8 GB/s	G34
Fam 15h Mod. 00h-0Fh (Bulldozer)		11/2011	Interlagos (2xOrochi die)	6200	32 nm	2x4 CM (2x8 C)	2*4* 2 MB/CM	2* 8MB/ 4 CM	DDR3-1600	HT 3.1: 12.8 GB/s	G34
Fam. 15h Mod. 10h-1Fh (Piledriver)		11/2012	Abu Dhabi (2 dies)	6300	32 nm	2x4 CM (2x8 C)	2*4* 2 MB/CM	2* 8 MB/ 4 CM	DDR3-1866	HT 3.1 12.8 GB/s	G34
Fam. 15h Mod. 30h-3Fh (Steamroller)		2H/ 2013	No further Bulldozer-based server lines								

Block diagram of the Interlagos processor [6]

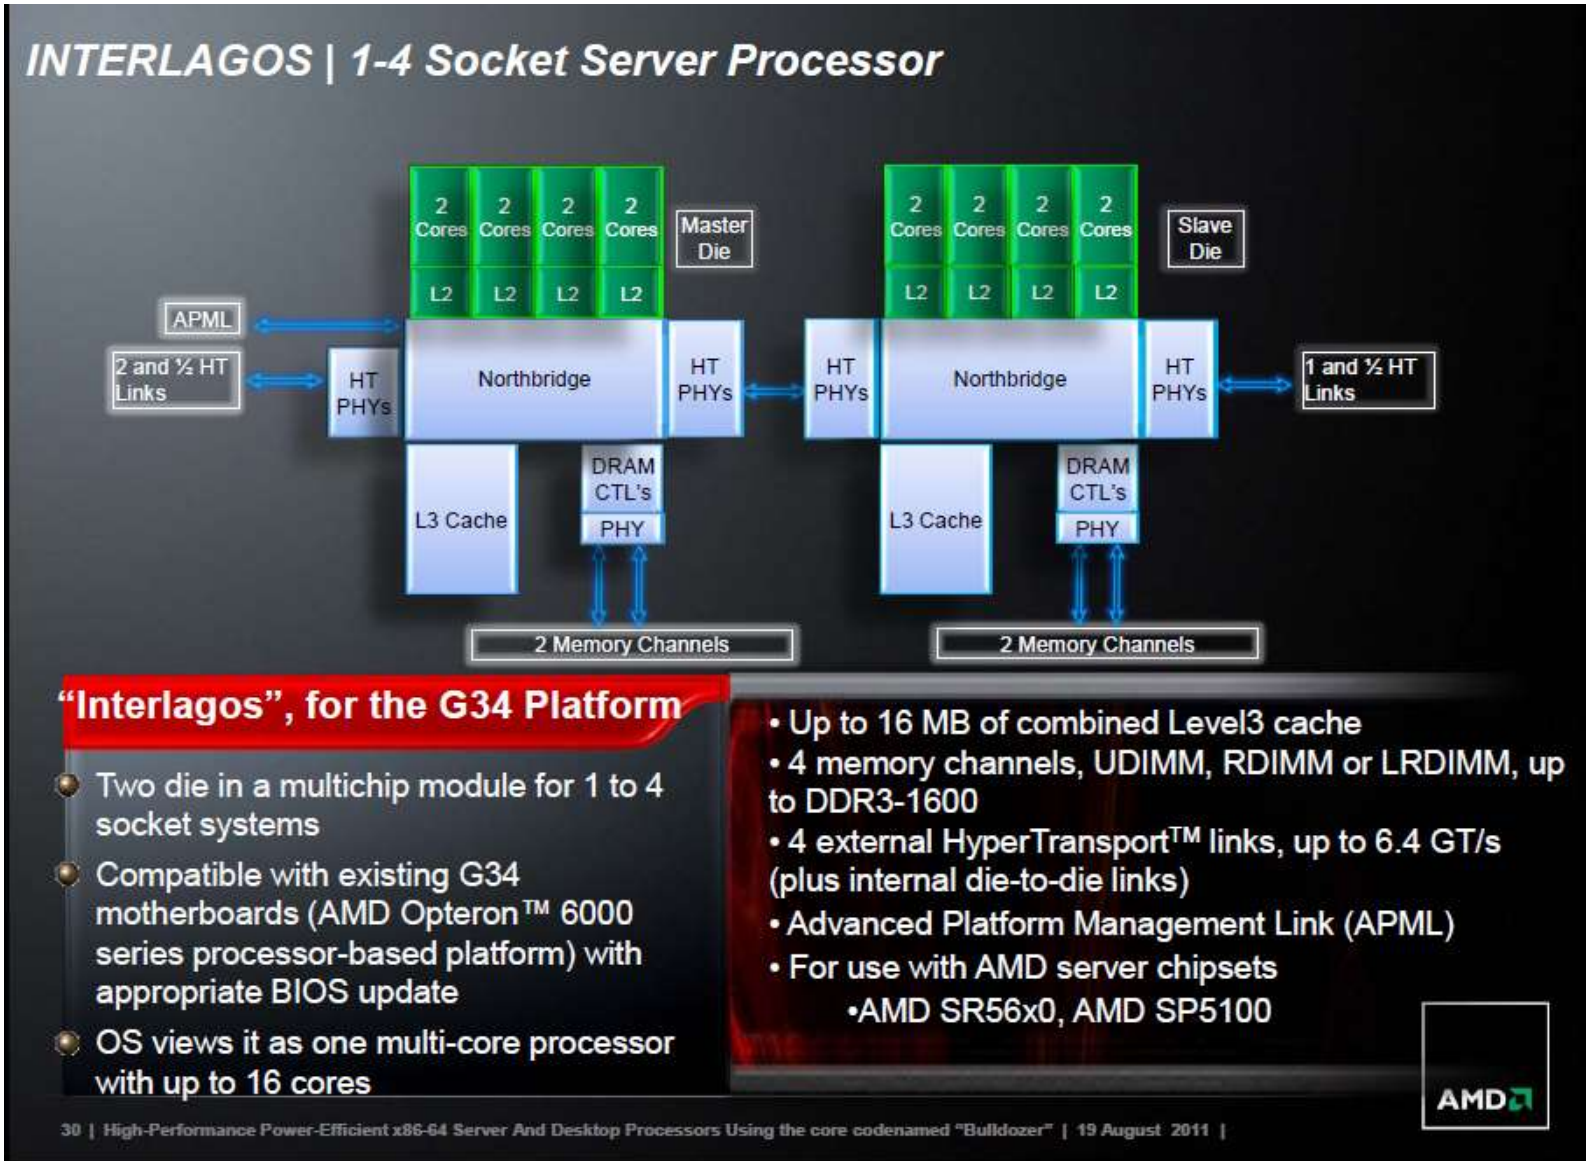

Example: Interlagos-based MP system [6]

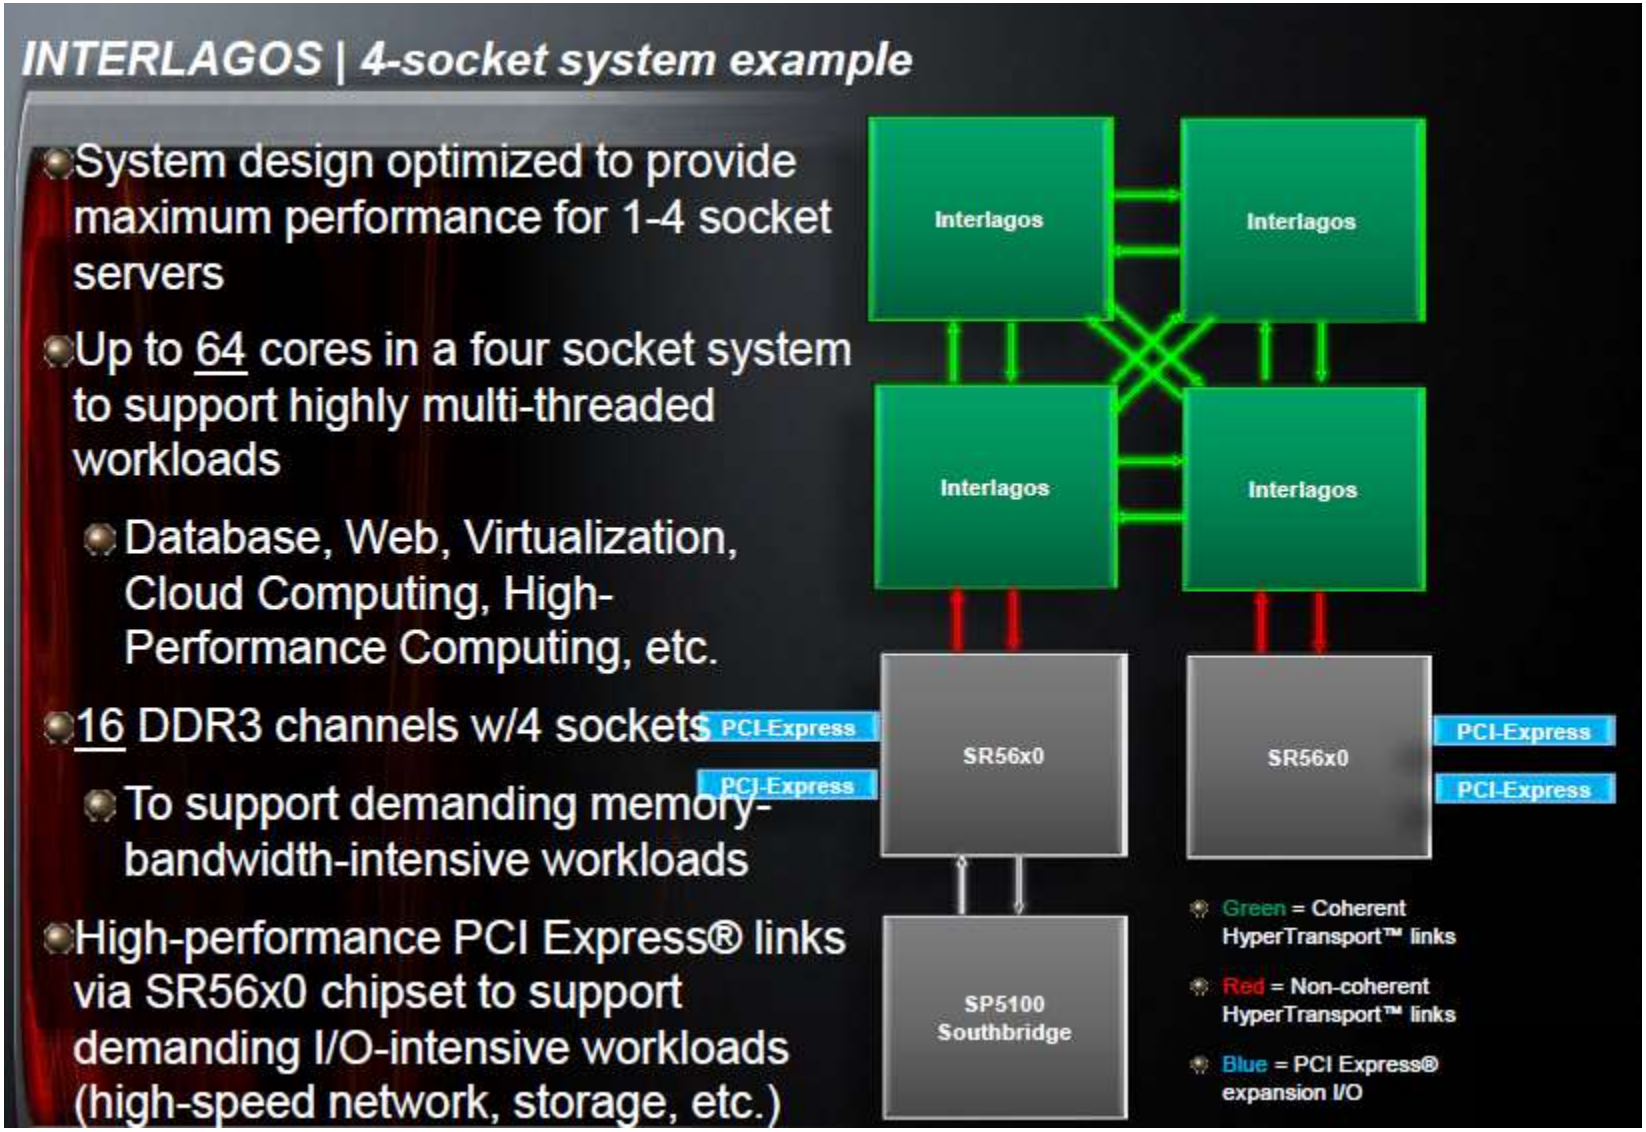

2.5.2 The Interlagos MP server line (4)

Performance increase of AMD's MP servers up to The Interlagos [18]

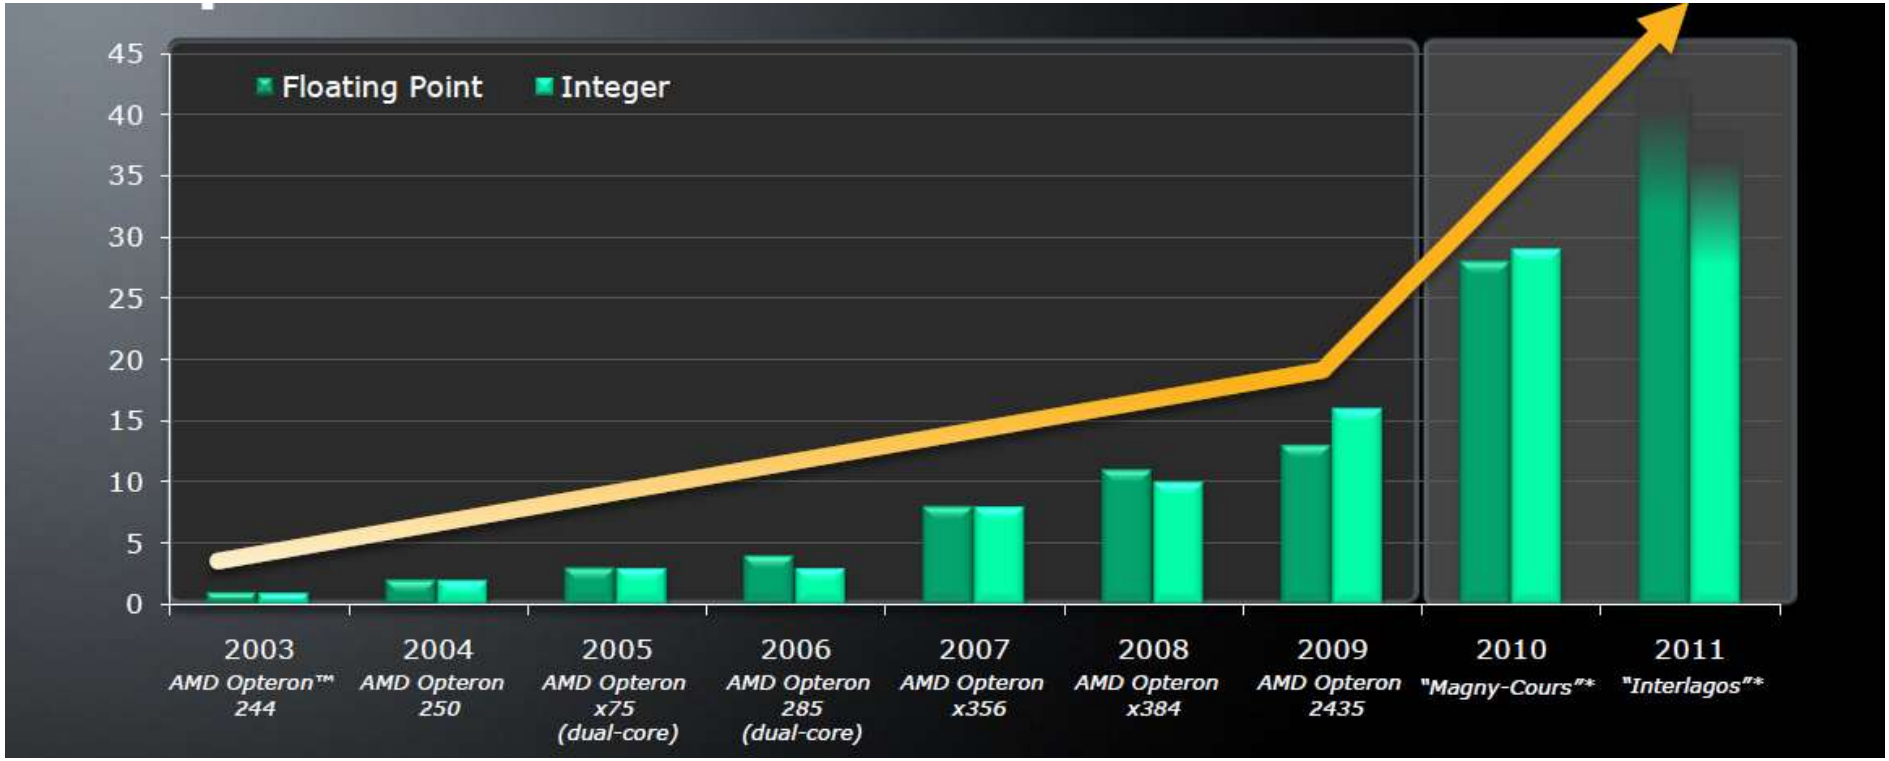

2.5.2 The Interlagos MP server line (5)

Performance/Watt evolution of AMD's server lines [2]

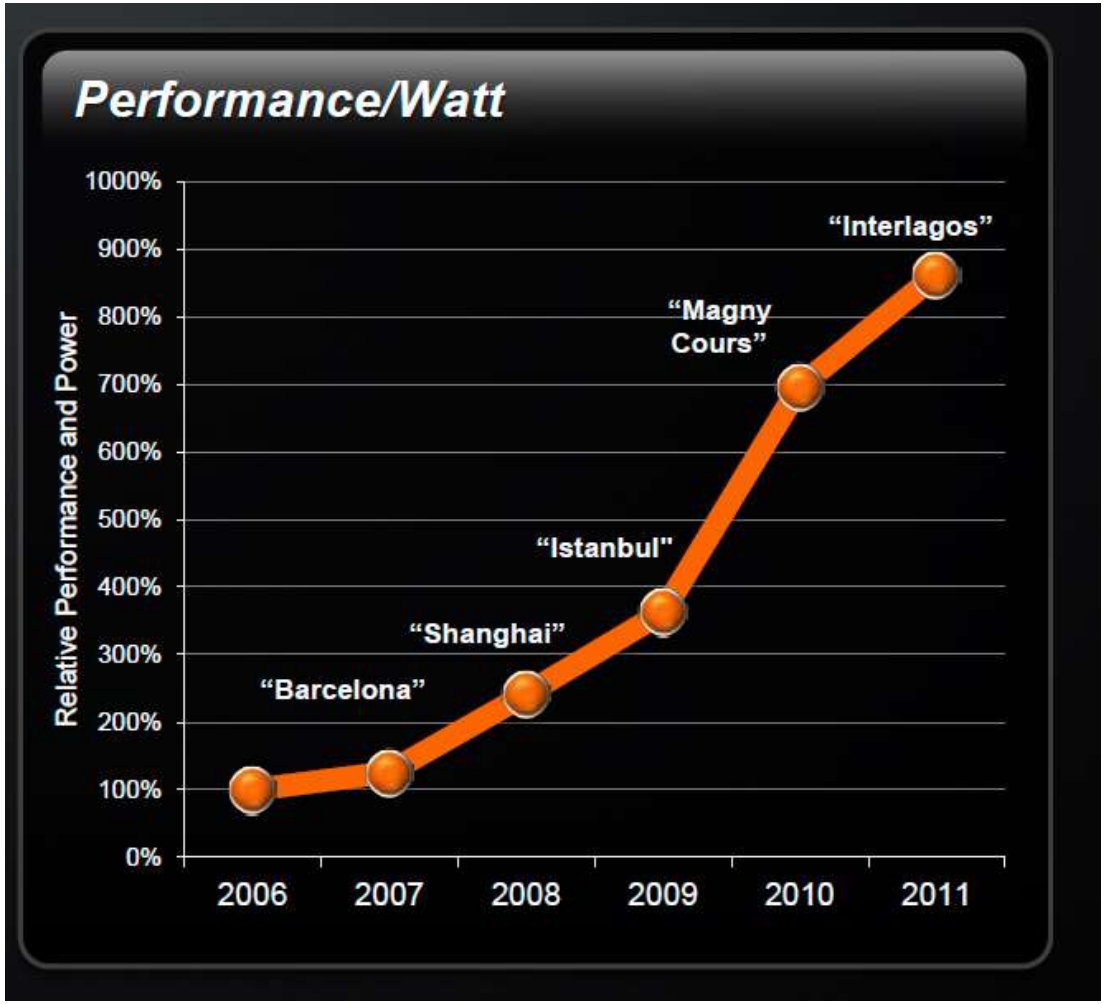

2.5.2 The Interlagos MP server line (6)

Main features of Bulldozer-based Interlagos MP server lines [13]

Released: 11/2011 – Socket G34

Model Number	Step	Cores	Frequency			Cache		HT	Vcore	ACP	TDP
			Base	Full Load turbo	Half Load turbo	L2	L3				
B2, Quad core											
Opteron 6204	B2	4	3.3 GHz	N/A	N/A	2 × 2 MB	2 × 8 MB	3.2 GHz		80 W	115 W
B2, Eight core											
Opteron 6212	B2	8	2.6 GHz	2.9 GHz	3.2 GHz	4 × 2 MB	2 × 8 MB	3.2 GHz		80 W	115 W
Opteron 6220	B2	8	3.0 GHz	3.3 GHz	3.6 GHz	4 × 2 MB	2 × 8 MB	3.2 GHz		80 W	115 W
B2, Twelve core											
Opteron 6234	B2	12	2.4 GHz	2.7 GHz	3.0 GHz	6 × 2 MB	2 × 8 MB	3.2 GHz		80 W	115 W
Opteron 6238	B2	12	2.6 GHz	2.9 GHz	3.2 GHz	6 × 2 MB	2 × 8 MB	3.2 GHz		80 W	115 W
B2, Sixteen core											
Opteron 6272	B2	16	2.1 GHz	2.4 GHz	3.0 GHz	8 × 2 MB	2 × 8 MB	3.2 GHz		80 W	115 W
Opteron 6274	B2	16	2.2 GHz	2.5 GHz	3.1 GHz	8 × 2 MB	2 × 8 MB	3.2 GHz		80 W	115 W
Opteron 6276	B2	16	2.3 GHz	2.6 GHz	3.2 GHz	8 × 2 MB	2 × 8 MB	3.2 GHz		80 W	115 W
Opteron 6282 SE	B2	16	2.6 GHz	3.0 GHz	3.3 GHz	8 × 2 MB	2 × 8 MB	3.2 GHz		105 W	140 W
B2, Sixteen core, high-efficiency											
Opteron 6262 HE	B2	16	1.6 GHz	2.1 GHz	2.9 GHz	8 × 2 MB	2 × 8 MB	3.2 GHz		65 W	85 W

2.5.2 The Interlagos MP server line (7)

Comparing main features of Bulldozer-based lines with the previous generation [4]

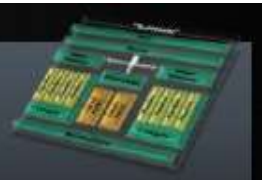

	AMD Opteron™ 4100/6100 Series Processors	"Valencia" / "Interlagos"
Cores	4100: 4 or 6 core; 6100: 8 or 12 core	4200: 6 or 8 core; 6200: 8, 12 or 16 core
Cache (L2 per core / L3 per die)	512KB / 6MB	2MB (shared between 2 cores) / 8MB
Memory Channels and speed	4100: two; 6100: four; up to 1333MHz	4200: two; 6200: four; up to 1600MHz
Floating point capability	128-bit FPU per core (FADD/FMUL)	128-bit dedicated FMAC per core or 256-bit AVX shared between 2 cores
Integer Issues Per Cycle	3	4
Turbo CORE Technology	No	Yes (+500MHz with all cores active)
Power (ACP)	65W, 80W, 105W	TBD (planned 65W, 80W, 105W)
New Instruction Sets		SSSE3, SSE 4.1/4.2, AVX, AES, FMA4, XOP, PCLMULQDQ
Power Gating	AMD CoolCore™, C1E	AMD CoolCore™, C1E, C6
Process / Die Size	45nm SOI	32nm SOI (smaller overall die size)
Performance		Expected up to 50% higher throughput

2.5.2 The Interlagos MP server line (8)

Performance assessment of Family 15h Bulldozer-based MP servers [13]

There are results available for Open Source server workload running on four DP configurations covering competing AMD and Intel server processors.

The Open Source server workload (termed as VApus FOS) was created by Anandtech.
It is is a mix of four Virtual Machines (VM) with open source workloads including

- Apache2,
- MySQL,
- Community server 5.1.37 database,
- VMware's open source groupware Zimbra 7.1.0.

The processors compared in DP configurations are

- AMD Opteron "Bulldozer" based Interlagos 6276 at 2.3GHz -16 cores
- AMD Opteron "Magny-Cours" 6174 at 2.2GHz – 12 cores
- Intel Xeon X5670 "Westmere" at 2.93GHz - 6 cores
- Intel Xeon X5650 "Westmere" at 2.66GHz - 6 cores

These processors have roughly the same price point.

The software environment and the hardware configurations are detailed in [26].

Throughput results of the Open Source server workload runs [26]

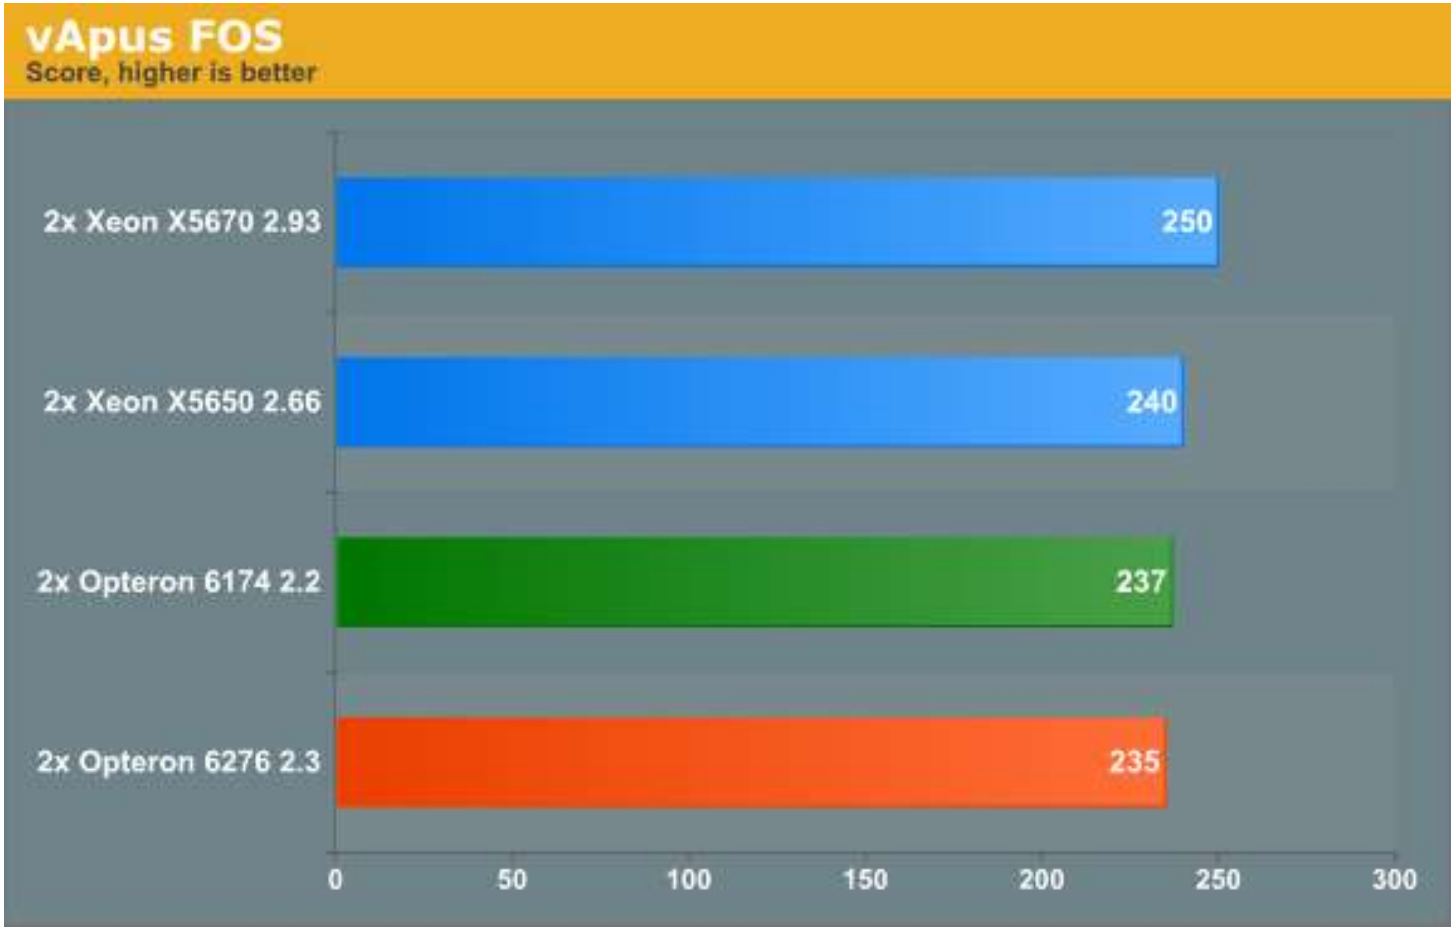

2.5.2 The Interlagos MP server line (10)

Response time results of the Open Source server workload runs [26]

vApus FOS Average Response Times (ms), lower is better!

CPU	PhpBB1	PHPBB2	MySQL OLAP	Zimbra
AMD Opteron 6276	737	587	170	567
AMD Opteron 6174	707	574	118	630
Intel Xeon X5670	645	550	63	593
Intel Xeon X5650	678	566	102	655

Power consumption results of the Open Source server workload runs [26]

2.5.2 The Interlagos MP server line (12)

Assessing the benchmark results gained for the Interlagos 6276 server

- The Bulldozer-based 8-module (16-core) Opteron 6276 (fc = 2.3 GHz) was at writing the report AMD's second highest performing Bulldozer server processor.
  (The flagship model 6282 SE is clocked at 2.6 GHz.)
- The benchmark results show that AMD's 16-core 2.3 GHz Opteron 6276
  - provides only a moderate performance increase over the previous K10.5 Magny Course-based 12-core Opteron 6174, if any and
  - it has lower performance figures than Intel's 6-core Westmere-based Xeon X5650/5670 processors clocked at 2.66 and 2.93 GHz, respectively.

2.5.3 The Turbo core technology of Bulldozer-based MP servers

2.5.3 The Turbo core technology of Bulldozer-based MP servers

Aim of the Turbo core technology

Increase performance of lightly threaded workloads by raising fc if there is a TDP headroom available.

Second generation Turbo core technology

The Turbo core technology of Bulldozer-based MP servers is already AMD's second generation Turbo core technology.

The first generation Turbo tech,nology became introduced with the 6-core K10.5 Istambul based desktop line (Phenom II X6 line) called Thuban (4/2010).

2.5.3 The Turbo core technology of Bulldozer-based MP servers (2)

Principle of operation [6]

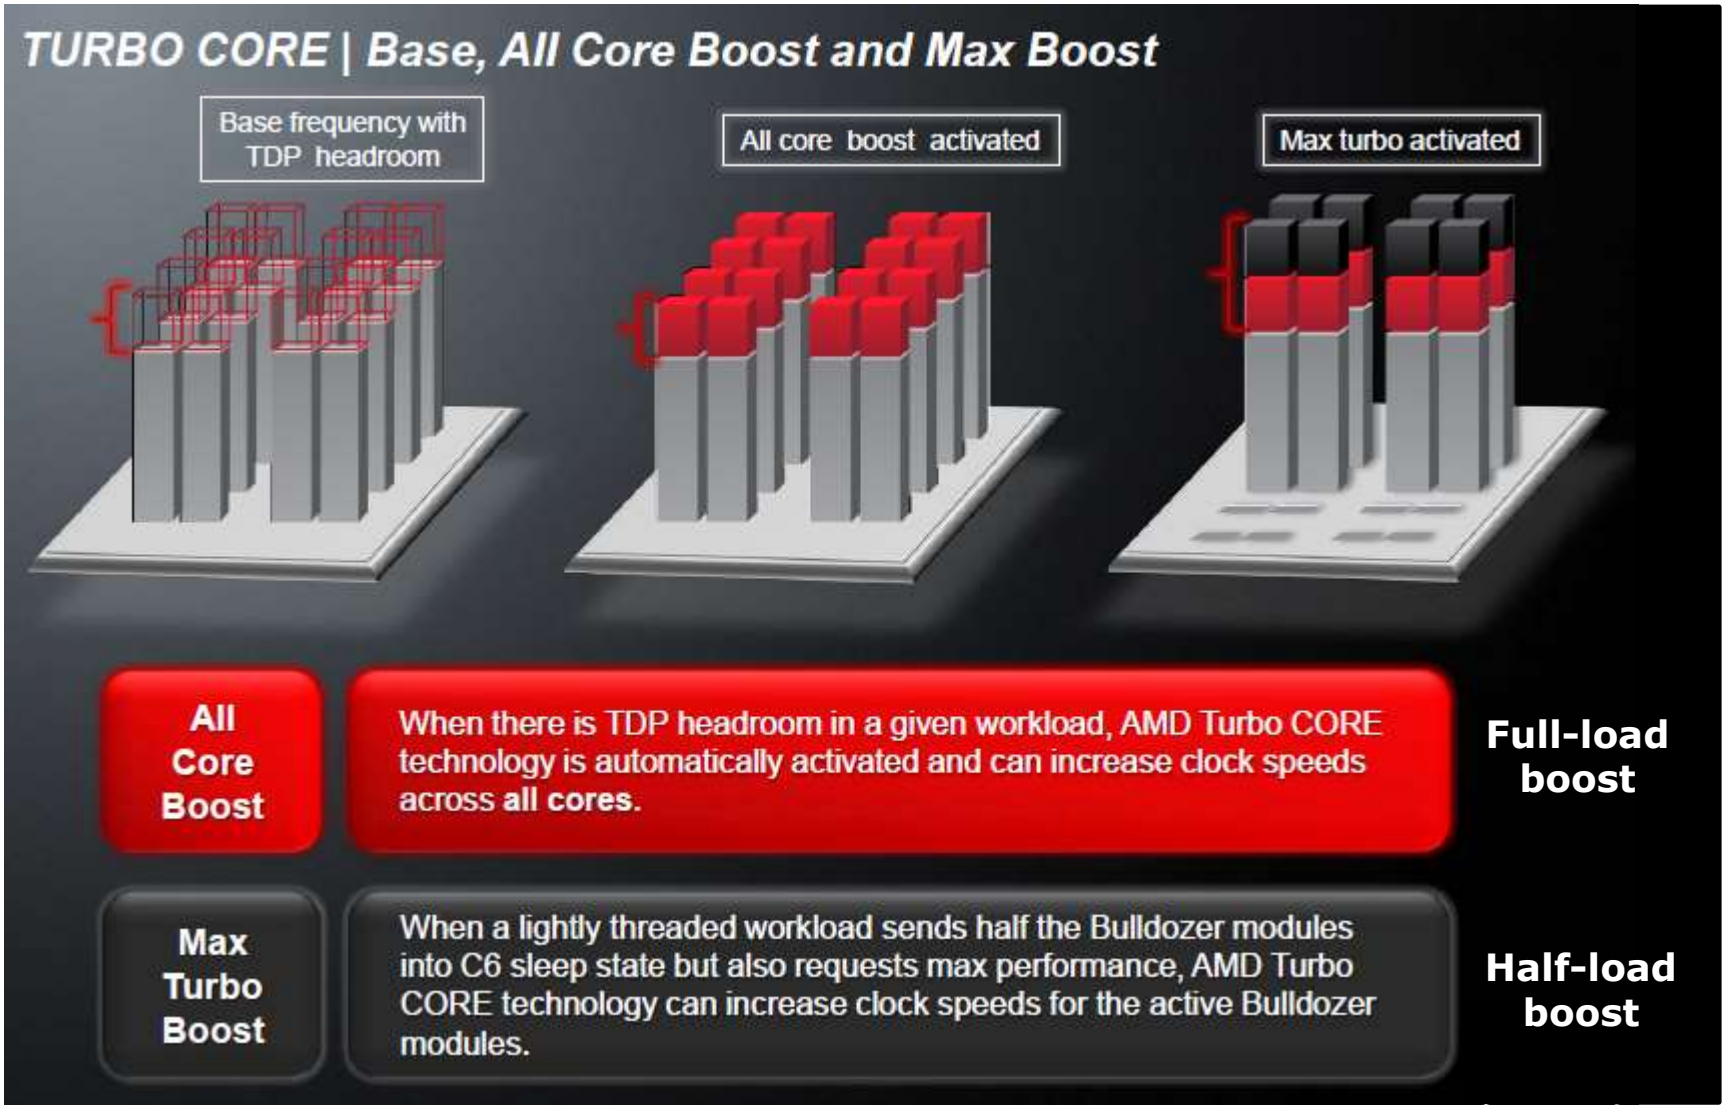

2.5.3 The Turbo core technology of Bulldozer-based MP servers (3)

Full and half load turbo frequencies of Family 15h Bulldozer-based Interlagos MP servers [13]

Model Number	Step	Cores	Frequency			Cache		HT	Vcore	ACP	TDP
			Base	Full Load turbo	Half Load turbo	L2	L3				
B2, Quad core											
Opteron 6204	B2	4	3.3 GHz	N/A	N/A	2 × 2 MB	2 × 8 MB	3.2 GHz		80 W	115 W
B2, Eight core											
Opteron 6212	B2	8	2.6 GHz	2.9 GHz	3.2 GHz	4 × 2 MB	2 × 8 MB	3.2 GHz		80 W	115 W
Opteron 6220	B2	8	3.0 GHz	3.3 GHz	3.6 GHz	4 × 2 MB	2 × 8 MB	3.2 GHz		80 W	115 W
B2, Twelve core											
Opteron 6234	B2	12	2.4 GHz	2.7 GHz	3.0 GHz	6 × 2 MB	2 × 8 MB	3.2 GHz		80 W	115 W
Opteron 6238	B2	12	2.6 GHz	2.9 GHz	3.2 GHz	6 × 2 MB	2 × 8 MB	3.2 GHz		80 W	115 W
B2, Sixteen core											
Opteron 6272	B2	16	2.1 GHz	2.4 GHz	3.0 GHz	8 × 2 MB	2 × 8 MB	3.2 GHz		80 W	115 W
Opteron 6274	B2	16	2.2 GHz	2.5 GHz	3.1 GHz	8 × 2 MB	2 × 8 MB	3.2 GHz		80 W	115 W
Opteron 6276	B2	16	2.3 GHz	2.6 GHz	3.2 GHz	8 × 2 MB	2 × 8 MB	3.2 GHz		80 W	115 W
Opteron 6282 SE	B2	16	2.6 GHz	3.0 GHz	3.3 GHz	8 × 2 MB	2 × 8 MB	3.2 GHz		105 W	140 W
B2, Sixteen core, high-efficiency											
Opteron 6262 HE	B2	16	1.6 GHz	2.1 GHz	2.9 GHz	8 × 2 MB	2 × 8 MB	3.2 GHz		65 W	85 W

2.5.3 The Turbo core technology of Bulldozer-based MP servers (4)

A detailed description of the Bulldozer-based Turbo core technique will be given in connection with the Zambezi desktop processor, in Section 2.6.

2.5.4 The Valencia DP and Zurich UP server lines

2.5.4 The Valencia DP and Zurich UP server lines

AMD's 2012 – 2013 server roadmap [2]

The Family 15h Bulldozer-based DP system (Valencia) [6]

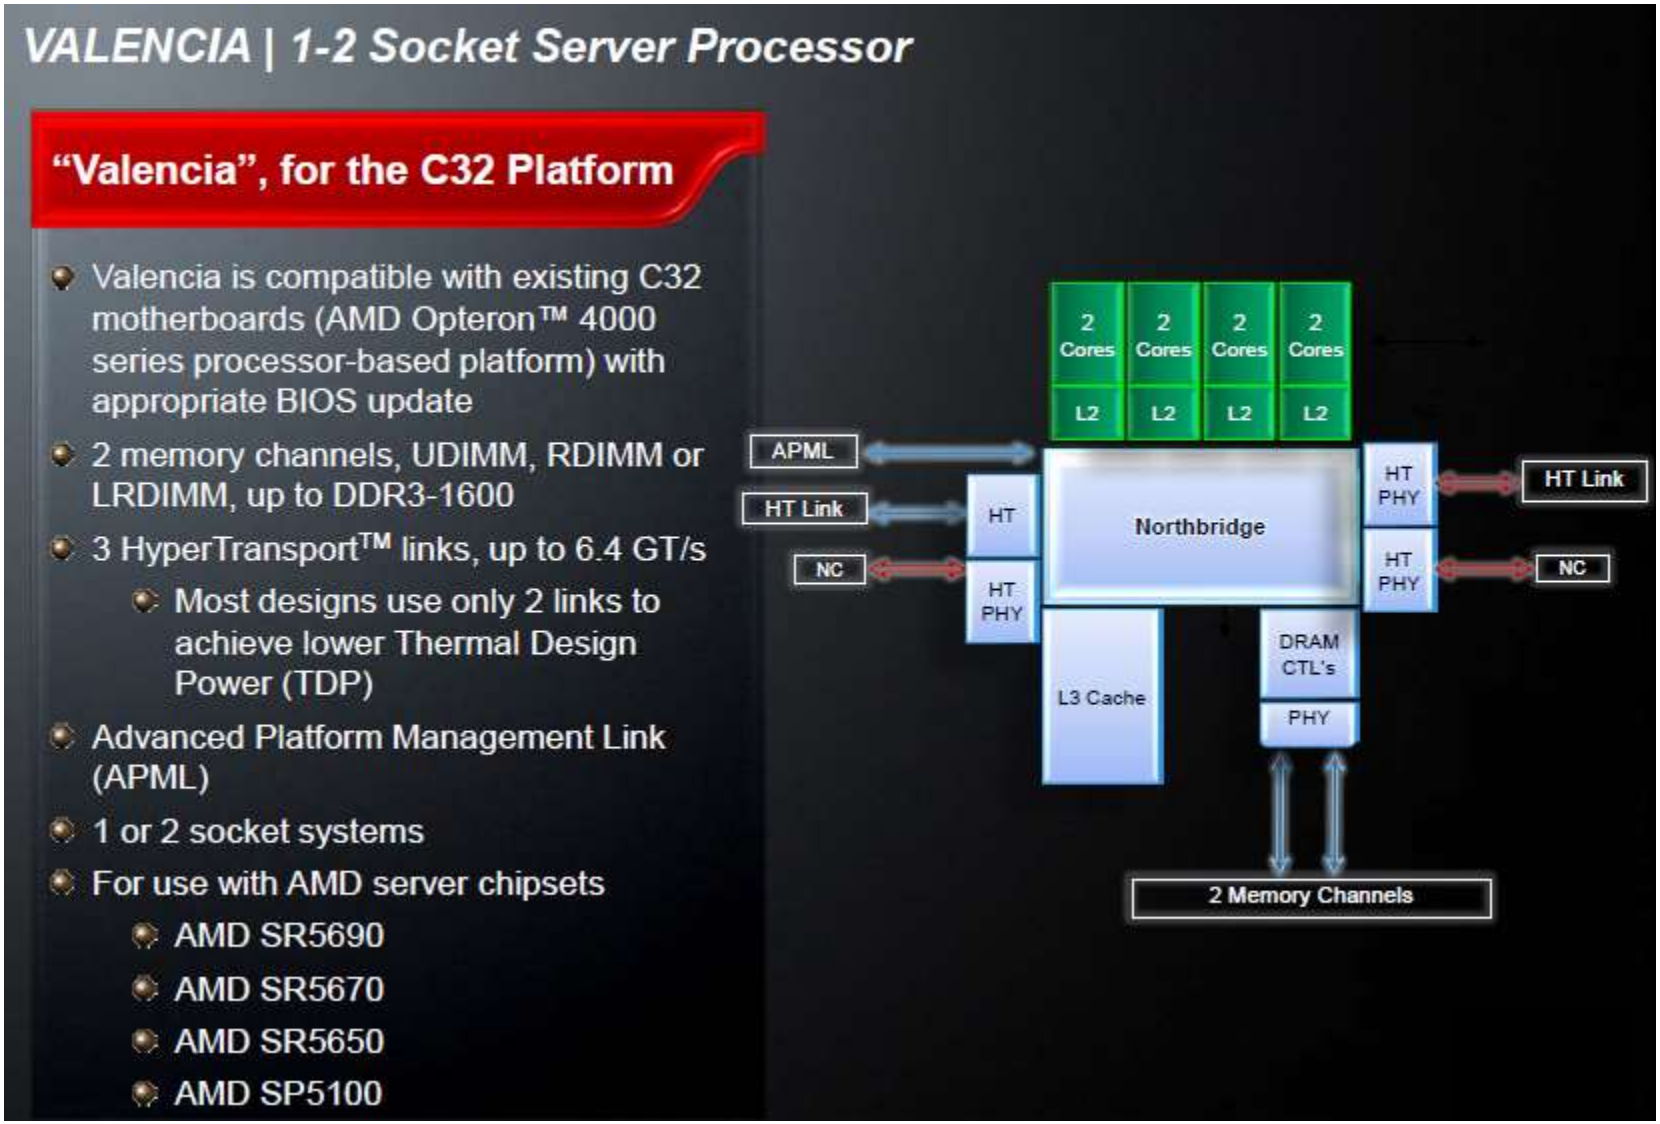

2.5.4 The Valencia DP and Zurich UP server lines (3)

Example Family 15h Bulldozer-based DP system (Valencia) [6]

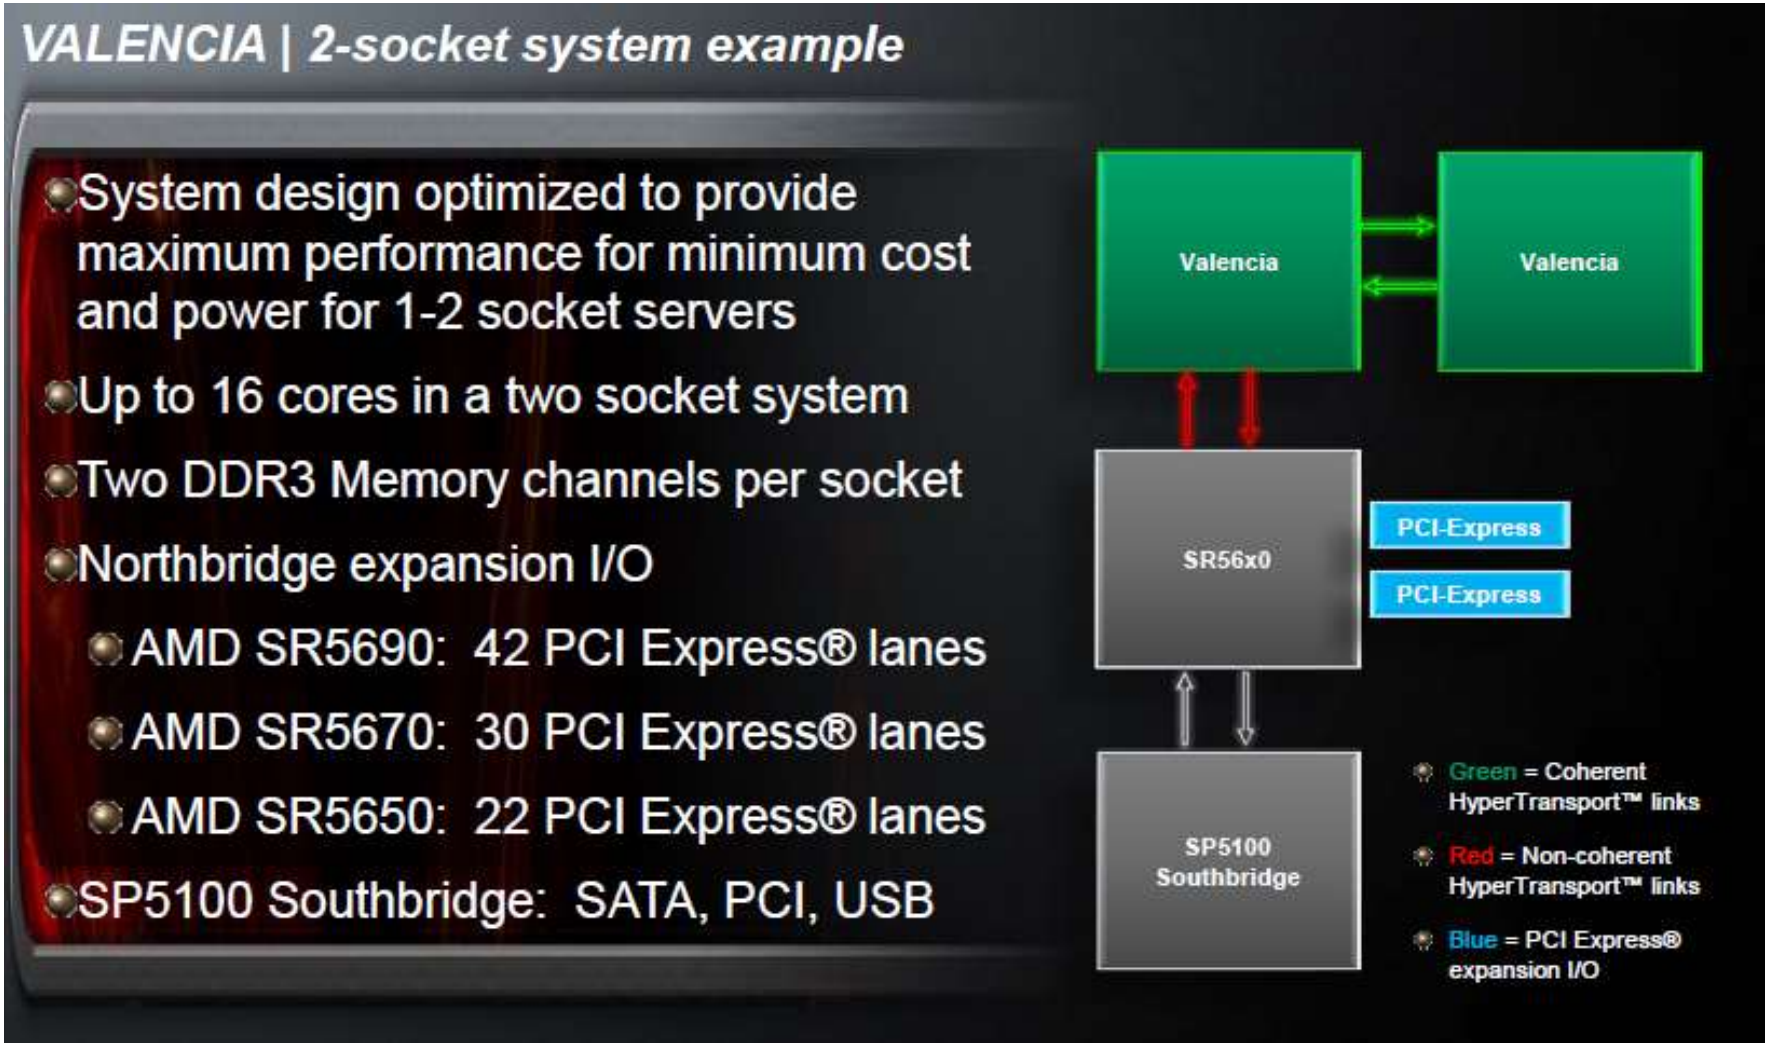

2.5.4 The Valencia DP and Zurich UP server lines (4)

Main parameters of the Family 15h Bulldozer-based Valencia DP server line [13]

Released 11/2011 - Socket C32

Model Number	Step.	Cores	Frequency			Cache		HT	Vcore	ACP	TDP
			Base	Full Load turbo	Half Load turbo	L2	L3				
B2, Quad core, energy-efficient											
Opteron 42DX EE	B2	4	2.2 GHz		3.3 GHz	2 × 2 MB	8 MB	3.2 GHz			40 W
B2, Six core											
Opteron 4226	B2	6	2.7 GHz	2.9 GHz	3.1 GHz	3 × 2 MB	8 MB	3.2 GHz		75 W	95 W
Opteron 4234	B2	6	3.1 GHz	3.3 GHz	3.5 GHz	3 × 2 MB	8 MB	3.2 GHz		75 W	95 W
Opteron 4238	B2	6	3.3 GHz	3.5 GHz	3.7 GHz	3 × 2 MB	8 MB	3.2 GHz		75 W	95 W
B2, Six core, high-efficiency											
Opteron 4228 HE	B2	6	2.8 GHz	3.1 GHz	3.6 GHz	3 × 2 MB	8 MB	3.2 GHz		50 W	65 W
B2, Eight core											
Opteron 4280	B2	8	2.8 GHz	3.1 GHz	3.5 GHz	4 × 2 MB	8 MB	3.2 GHz		75 W	95 W
Opteron 4284	B2	8	3.0 GHz	3.3 GHz	3.7 GHz	4 × 2 MB	8 MB	3.2 GHz		75 W	95 W
B2, Eight core, high-efficiency											
Opteron 42MX HE	B2	8	2.2 GHz		3.3 GHz	4 × 2 MB	8 MB	3.2 GHz			65 W
Opteron 4274 HE	B2	8	2.5 GHz	2.8 GHz	3.5 GHz	4 × 2 MB	8 MB	3.2 GHz		50 W	65 W
B2, Eight core, energy-efficient											
Opteron 4256 EE	B2	8	1.6 GHz	1.9 GHz	2.8 GHz	4 × 2 MB	8 MB	3.2 GHz		32 W	35 W

2.5.4 The Valencia DP and Zurich UP server lines (5)

Main parameters of the Family 15h Bulldozer-based Zurich UP server line [13]

Released 3/2012 – Socket AM3+

Model Number	Step.	Cores	Frequency			Cache		HT	V_{core}	ACP	TDP
			Base	Full Load turbo	Half Load turbo	L2	L3				
B2, Four core, energy-efficient											
Opteron 3250 EE	B2	4	2.5 GHz	2.8 GHz	3.5 GHz	2 × 2 MB	4 MB	3.2 GHz			45 W
Opteron 3260 EE	B2	4	2.7 GHz	3.0 GHz	3.7 GHz	2 × 2 MB	4 MB	3.2 GHz			45 W
B2, Eight core, high-efficiency											
Opteron 3280 HE	B2	8	2.4 GHz	2.7 GHz	3.5 GHz	4 × 2 MB	8 MB	3.2 GHz			65 W

2.5.4 The Valencia DP and Zurich UP server lines (6)

AMD's 2012 – 2013 server roadmap [2]

2.5.4 The Valencia DP and Zurich UP server lines (7)

Subsequent roadmaps of AMD's basic lines [27]

Published: Oct. 2011.

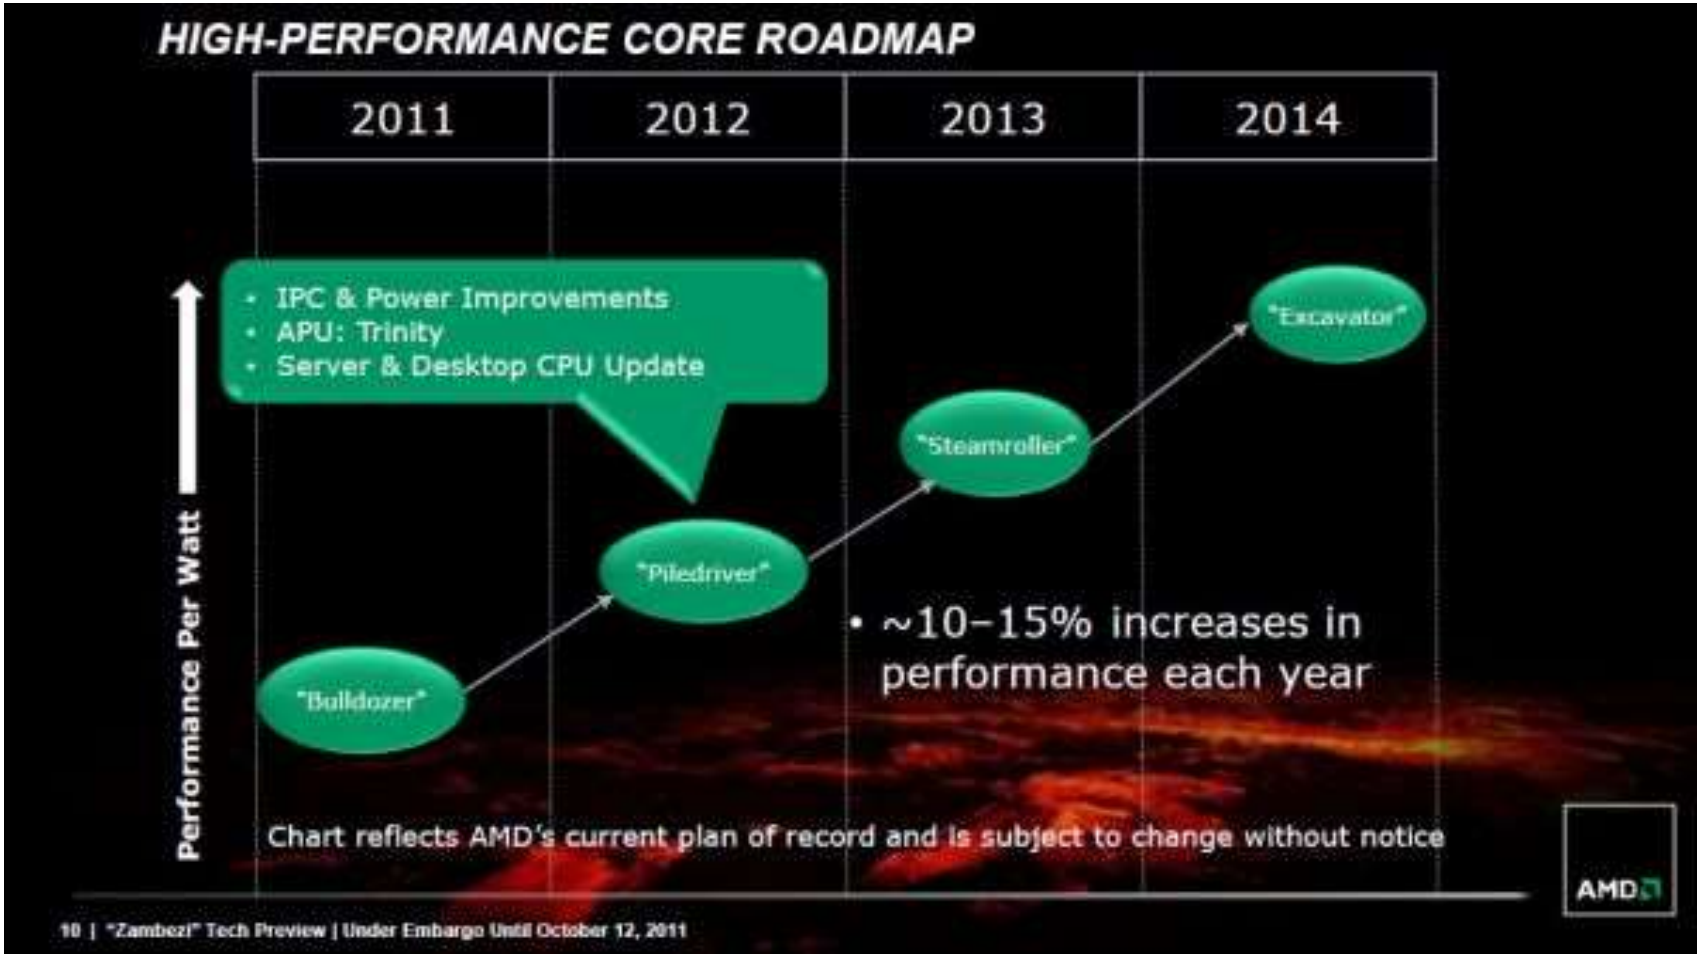

2.5.4 The Valencia DP and Zurich UP server lines (8)

Note

If AMD will achieve only the estimated performance increase of about 10-15 % per year they have no hope to compete with Intel on the high performance segment.

2.6 The Zambezi DT line

- 2.6.1 Overview of the high performance Zambezi desktop line
- 2.6.2 The Turbo core technology of the Zambezi desktop line
- 2.6.3 Performance assessment of the Zambezi desktop line

2.6.1 Overview of the high performance Zambezi desktop line

2.6.1 Overview of the high performance Zambezi desktop line [1]

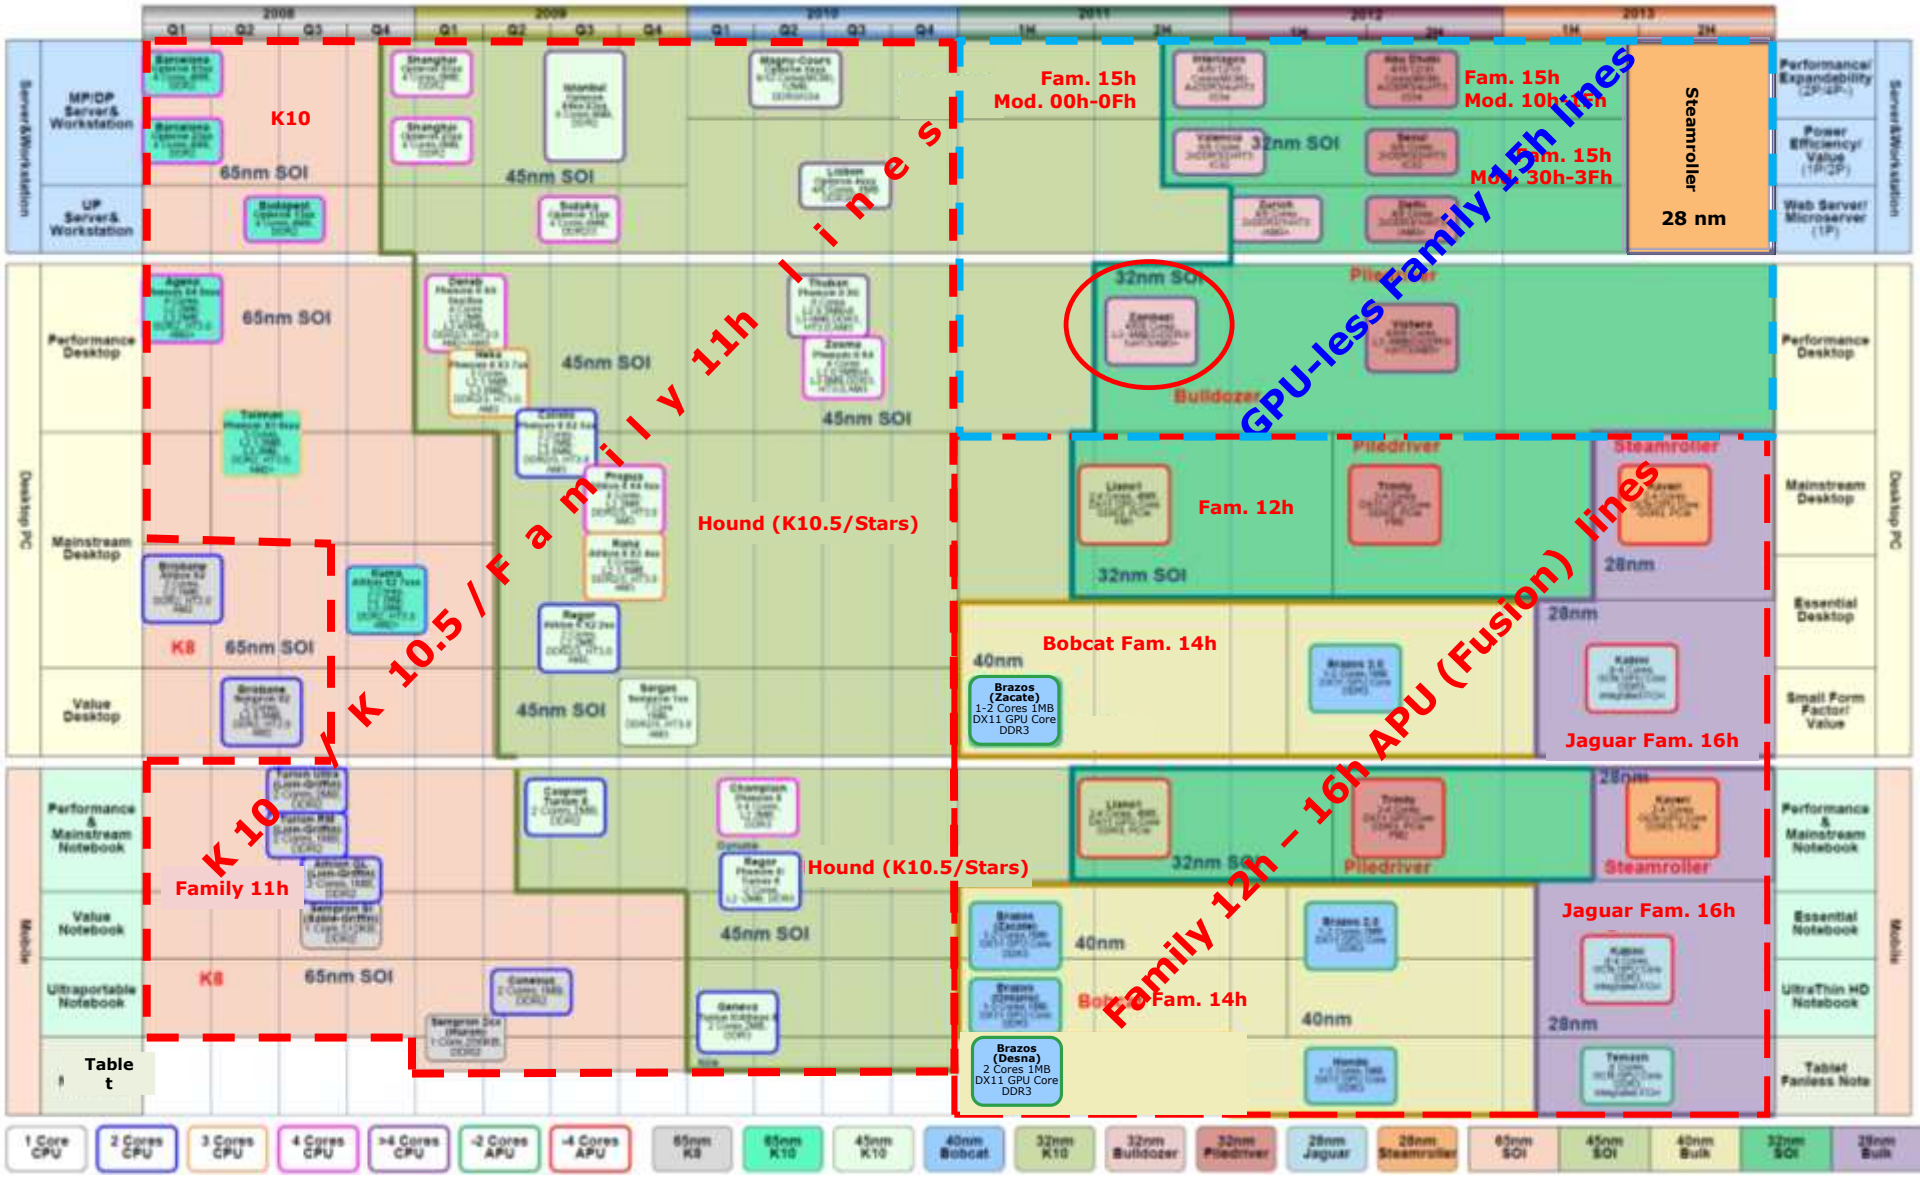

2.6.1 Overview of the high performance Zambezi desktop line (2)

Positioning AMD's Bulldozer-based Zambezi high performance desktop line

	Launched in	2011	2012	2013	2013	2015	2016
		Family 15h (00h-0Fh) (**Bulldozer**)	**Family 15h** (10h-1Fh) (**Piledriver**)	**Family 15h** (10h-1Fh) **(Piledriver v.2)**	**Family 15h** (30h-3Fh) (**Steamroller**)	**Family 15h** (60h-6Fh) (**Excavator v.1**)	**Family 15h** (77h-3Fh) (**Excavator v.2**)
Servers	**4P servers** (85-140 W)	Interlagos	Abu Dhabi				
	2P servers (85-140 W)	Valencia	Seoul				
	1P servers (85-140 W)	Zurich	Delhi				
Desktops	**High perf.** (~95-125 W)	Zambezi FX-Series	Vishera FX-Series				
	Mainstream (~65-95 W)		Trinity A10-A4	Richland A10/A8/A6/A4	Kaveri A10/A8		
Notebooks	**Mainstream** (~25-35 W)		Trinity A10/A8/A6M	Richland A10/A8/A6M	Kaveri FX/A10/A8P		Bristol Ridge FX/A12/A10P
	Ultra-thin (~10 – 15 W)		Trinity A10/A6M	Richland A10/A8/A6/A4M	A8 Pro/A8(B) A6 Pro/A6(B)	Carrizo FX/A10/A8P	Bristol Ridge FX/A12/A10P Stoney Ridge A9/A6
Tablets (~5 W)							

2.6.1 Overview of the high performance Zambezi desktop line (4)

The Family 15h Bulldozer-based high performance Zambezi desktop line [6]

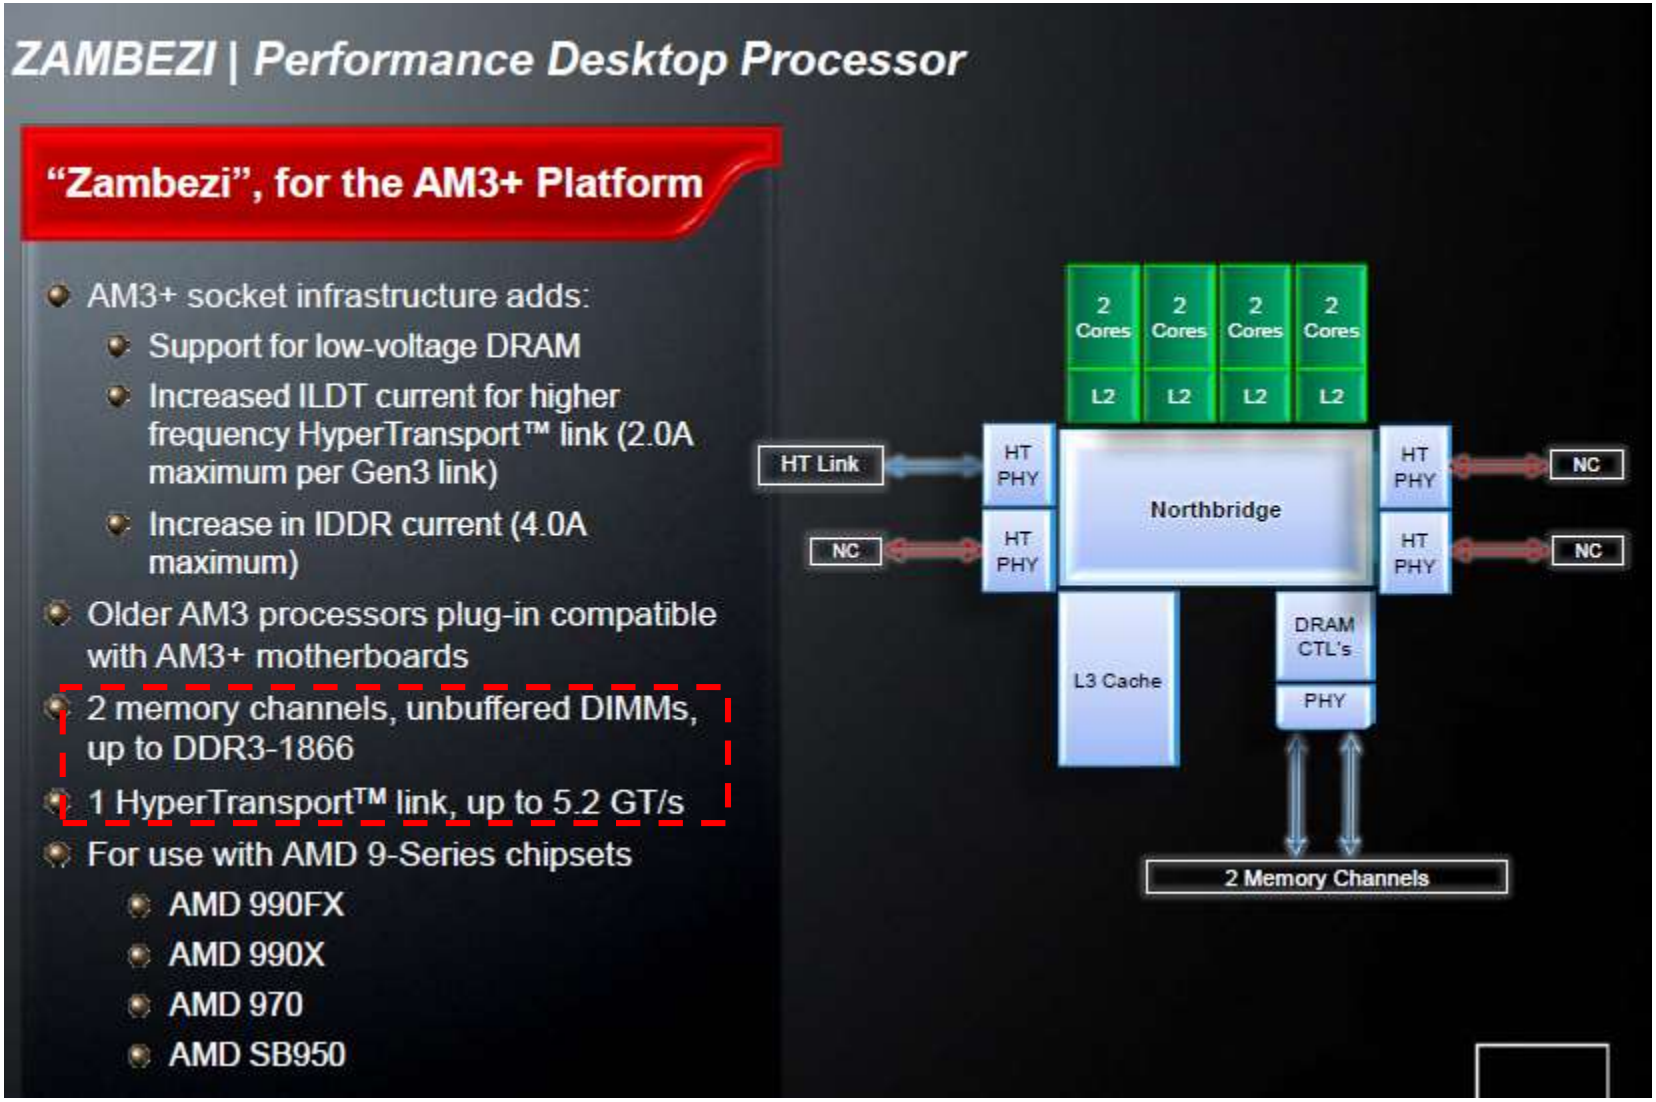

2.6.1 Overview of the high performance Zambezi desktop line (5)

Die plot of Zambezi [28]

Zambezi is based on the Orochi die (it includes 4 Bulldozer modules)

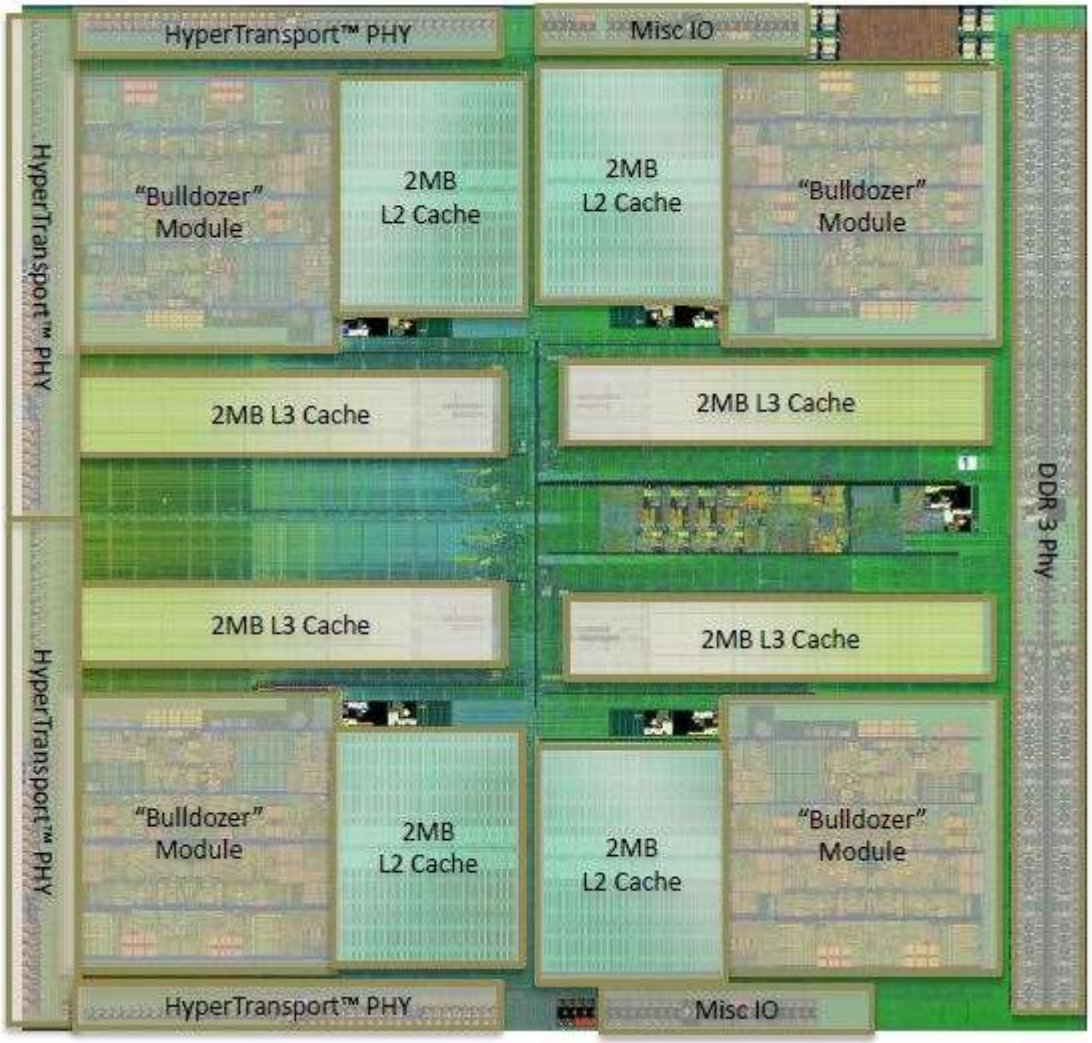

32 nm
315 mm²
1.2 mrd trs

2.6.1 Overview of the high performance Zambezi desktop line (6)

Key parameters of the Family 15h Bulldozer-based Zambezi desktop line [29]

AMD FX-series processors									
Name	**Nominal Frequency**	**Turbo Core Frequency**	**Max Turbo Frequency**	**TDP**	**Cores**	**Level 2 cache**	**Level 3 cache**	**Northbridge Frequency**	**SRP**
FX-8150*	3.6GHz	3.9GHz	4.2GHz	125W	8	8MB	8MB	2.2GHz	$245
FX-8120*	3.1GHz	3.4GHz	4GHz	95W/ 125W	8	8MB	8MB	2.2GHz	$205
FX-8100	2.8GHz	3.1GHz	3.7GHz	95W	8	8MB	8MB	2GHz	Unknown
FX-6100*	3.3GHz	3.6GHz	3.9GHz	95W	6	6MB	8MB	2GHz	$165
FX-4170	4.2GHz	None	4.3GHz	125W	4	4MB	8MB	2.2GHz	Unknown
FX-B4150	3.8GHz	3.9GHz	4GHz	95W	4	4MB	8MB	2.2GHz	Unknown
FX-4100*	3.6GHz	3.7GHz	3.8GHz	95W	4	4MB	8MB	2GHz	$115

2.6.1 Overview of the high performance Zambezi desktop line (7)

System example of a Zambezi desktop system (Scorpius platform) [30]

AMD
Scorpius platform
for FX8150
(Bulldozer based)

8 cores

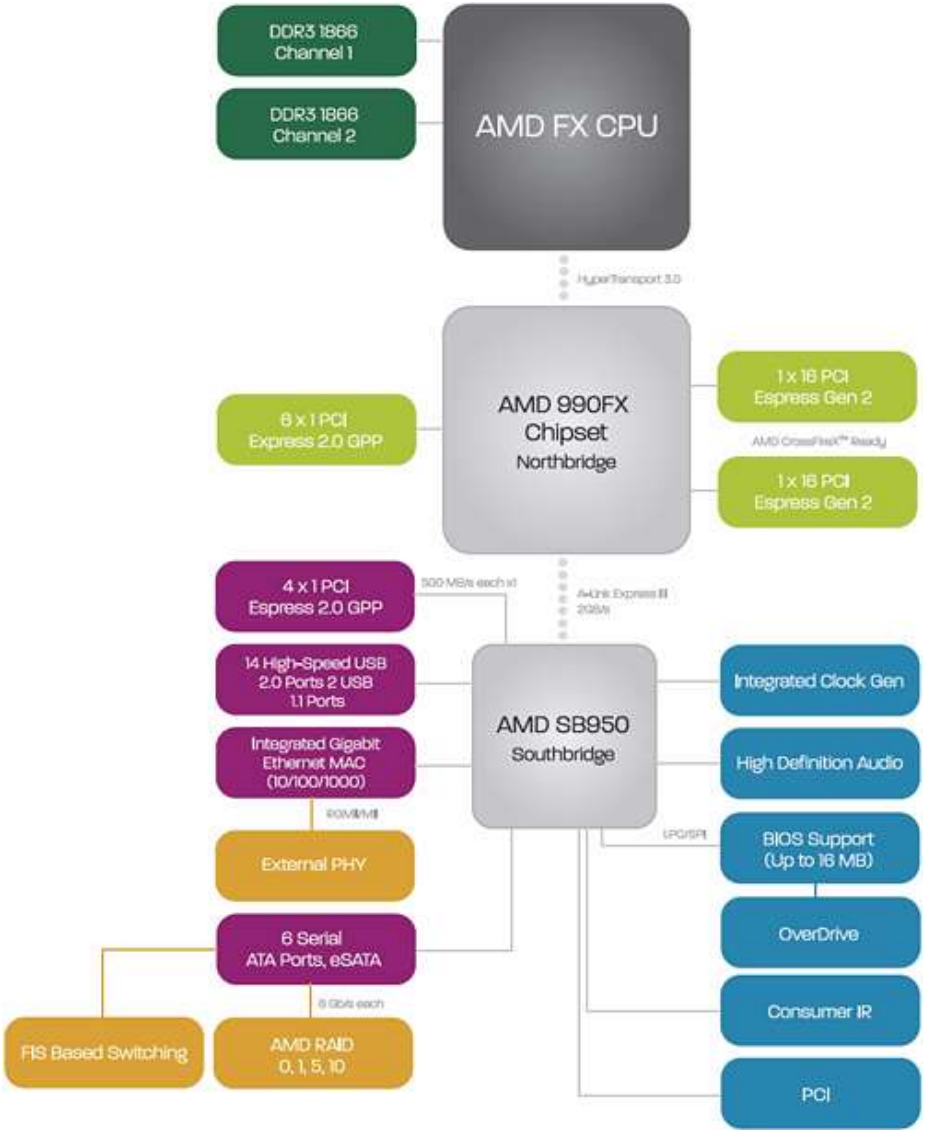

2.6.2 The Turbo core technology of the Zambezi desktop line

2.6.2 The Turbo core technology of the Zambezi desktop line

Contrasting AMD's 1. and 2. gen. Turbo core implementations [36]

AMD's 1. generation Turbo core technology

- It appeared in K10.5 Istambul-based desktops (Phenom II X6, Thuban) in 4/2010.
- This processor did not yet support power-gating.
  → Much less headroom was available for the Turbo core technology.
- Beyond the base clock frequency there was only a single higher frequency value, the turbo frequency.

→ Turbo core became seldom actived and if so, it remained only for short times active.

Example [36]

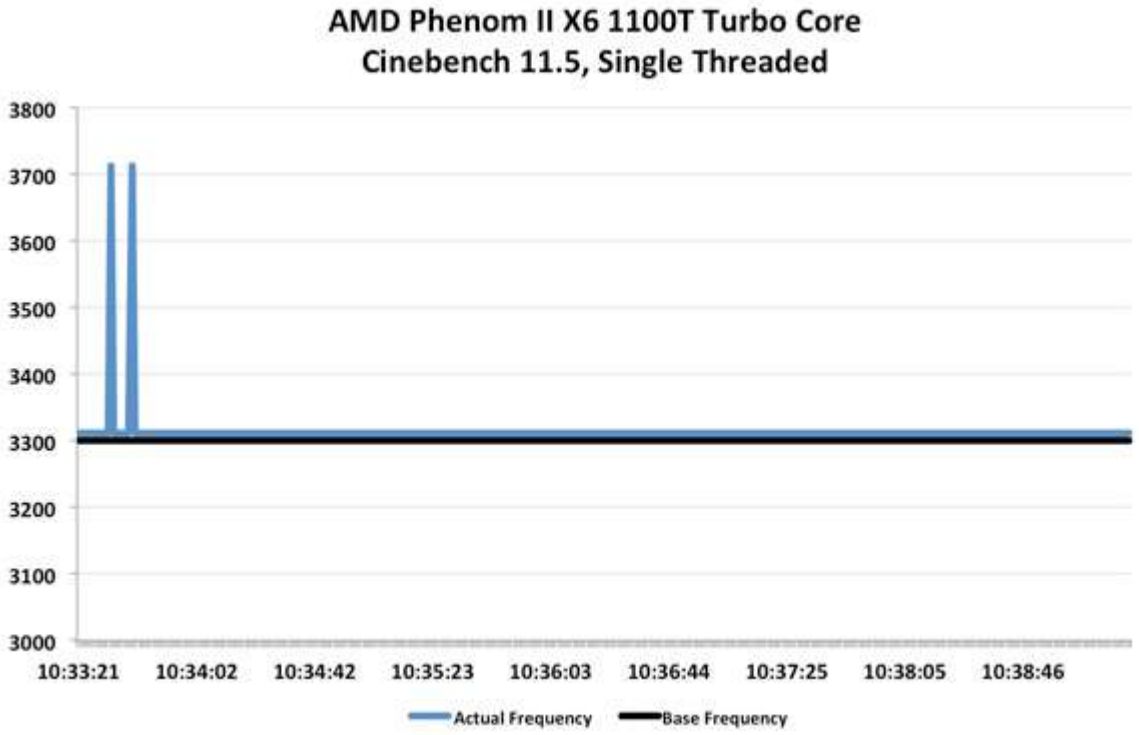

2.6.2 The Turbo core technology of the Zambezi desktop line (3)

AMD's 2. generation Turbo core technology

- The 2. generation Turbo core is introduced in Family 15h Bulldozer-based servers and desktops (Interlagos, Zambezi) in 10/2011.
- These processors do support power-gating.
  → So much more headroom remains for utilizing Turbo core.
- Beyond the base clock frequency there are two turbo levels,

  The 8-core Turbo frequency, that becomes activated if all cores are active but there remains a power headroom up to the TDP, and

  the 4-core Turbo frequency, that can be activated if at least half of the cores are in the CC6 state and the active cores request max. performance.

→ For single threaded applications the active core will run basically at the 8-core Turbo frequency and if there remains enough headroom to the TDP even at the the 4-core Turbo frequency, as demonstrated below.

2.6.2 The Turbo core technology of the Zambezi desktop line (4)

Principle of operation [6]

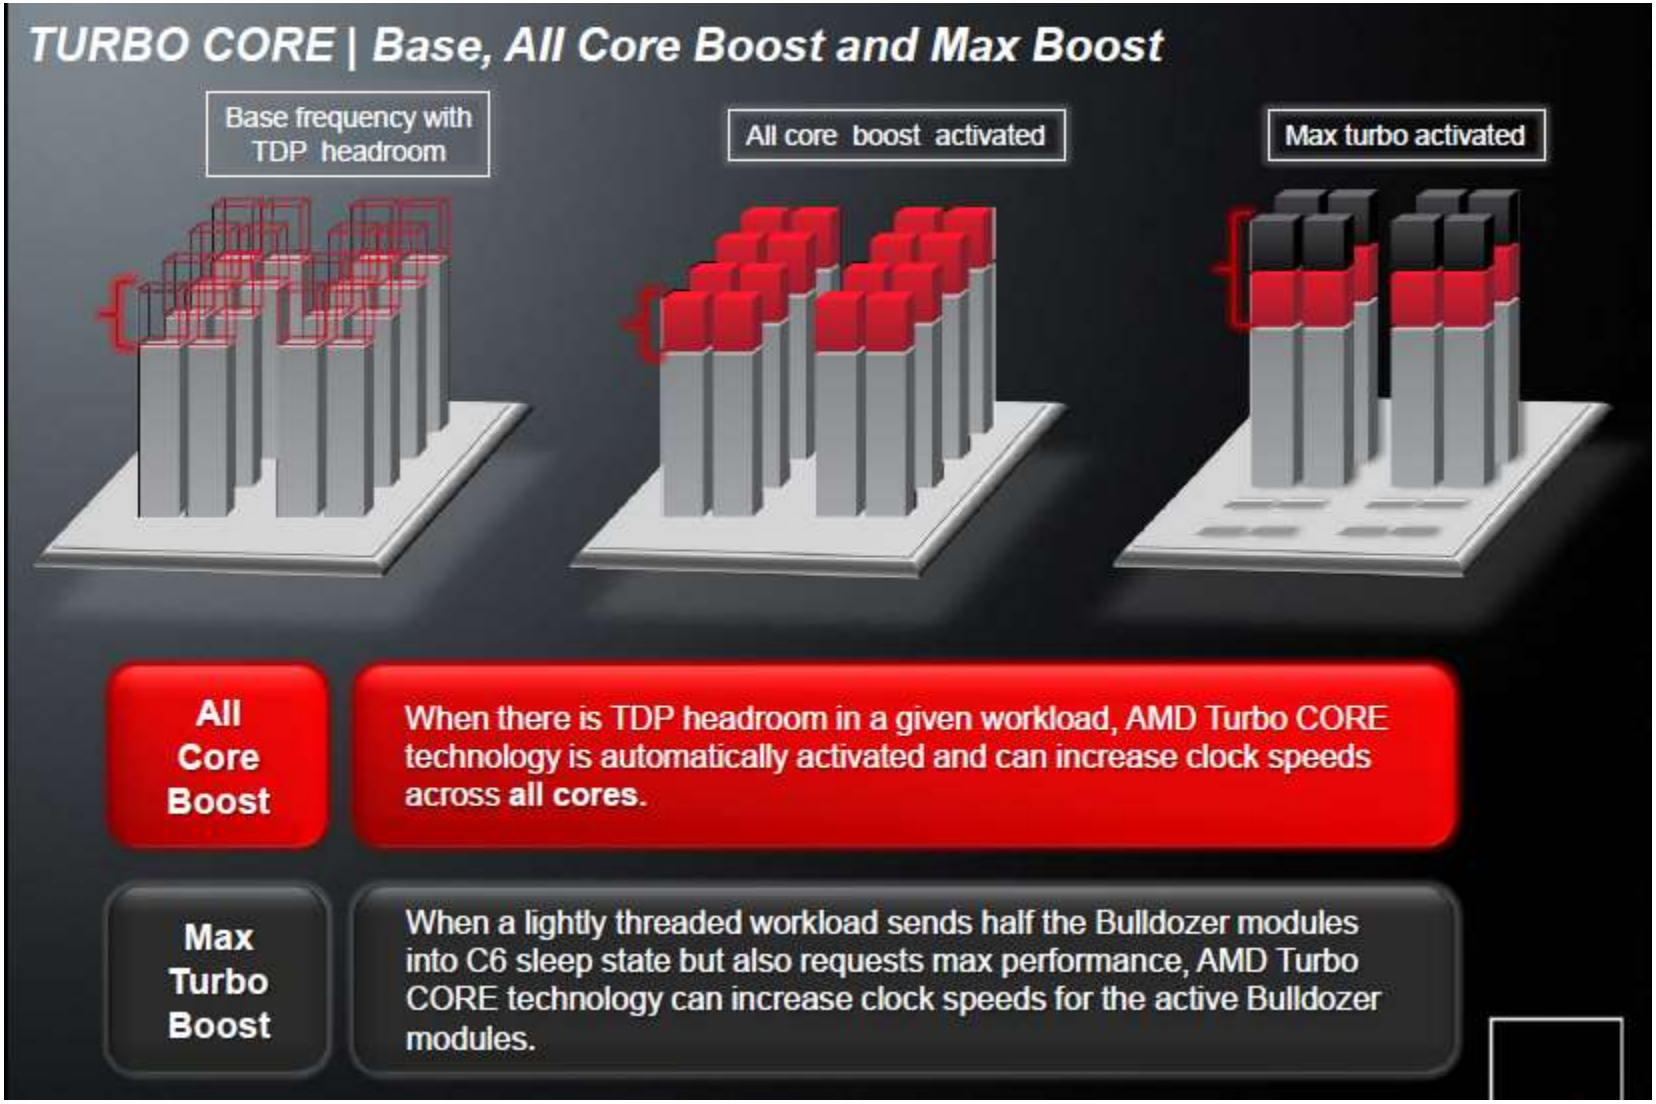

2.6.2 The Turbo core technology of the Zambezi desktop line (5)

Nominal, 8-core Turbo, and 4-core Max Turbo frequencies of Zambezi DT [29]

AMD FX-series processors									
Name	**Nominal Frequency**	**8-core Turbo frequency**	**4-core Max Turbo frequency**	**TDP**	**Cores**	**Level 2 cache**	**Level 3 cache**	**Northbridge Frequency**	**SRP**
FX-8150*	3.6GHz	3.9GHz	4.2GHz	125W	8	8MB	8MB	2.2GHz	$245
FX-8120*	3.1GHz	3.4GHz	4GHz	95W/ 125W	8	8MB	8MB	2.2GHz	$205
FX-8100	2.8GHz	3.1GHz	3.7GHz	95W	8	8MB	8MB	2GHz	Unknown
FX-6100*	3.3GHz	3.6GHz	3.9GHz	95W	6	6MB	8MB	2GHz	$165
FX-4170	4.2GHz	None	4.3GHz	125W	4	4MB	8MB	2.2GHz	Unknown
FX-B4150	3.8GHz	3.9GHz	4GHz	95W	4	4MB	8MB	2.2GHz	Unknown
FX-4100*	3.6GHz	3.7GHz	3.8GHz	95W	4	4MB	8MB	2GHz	$115

2.6.2 The Turbo core technology of the Zambezi desktop line (6)

Example for the operation of AMD's 2. generation Turbo core technology [37]

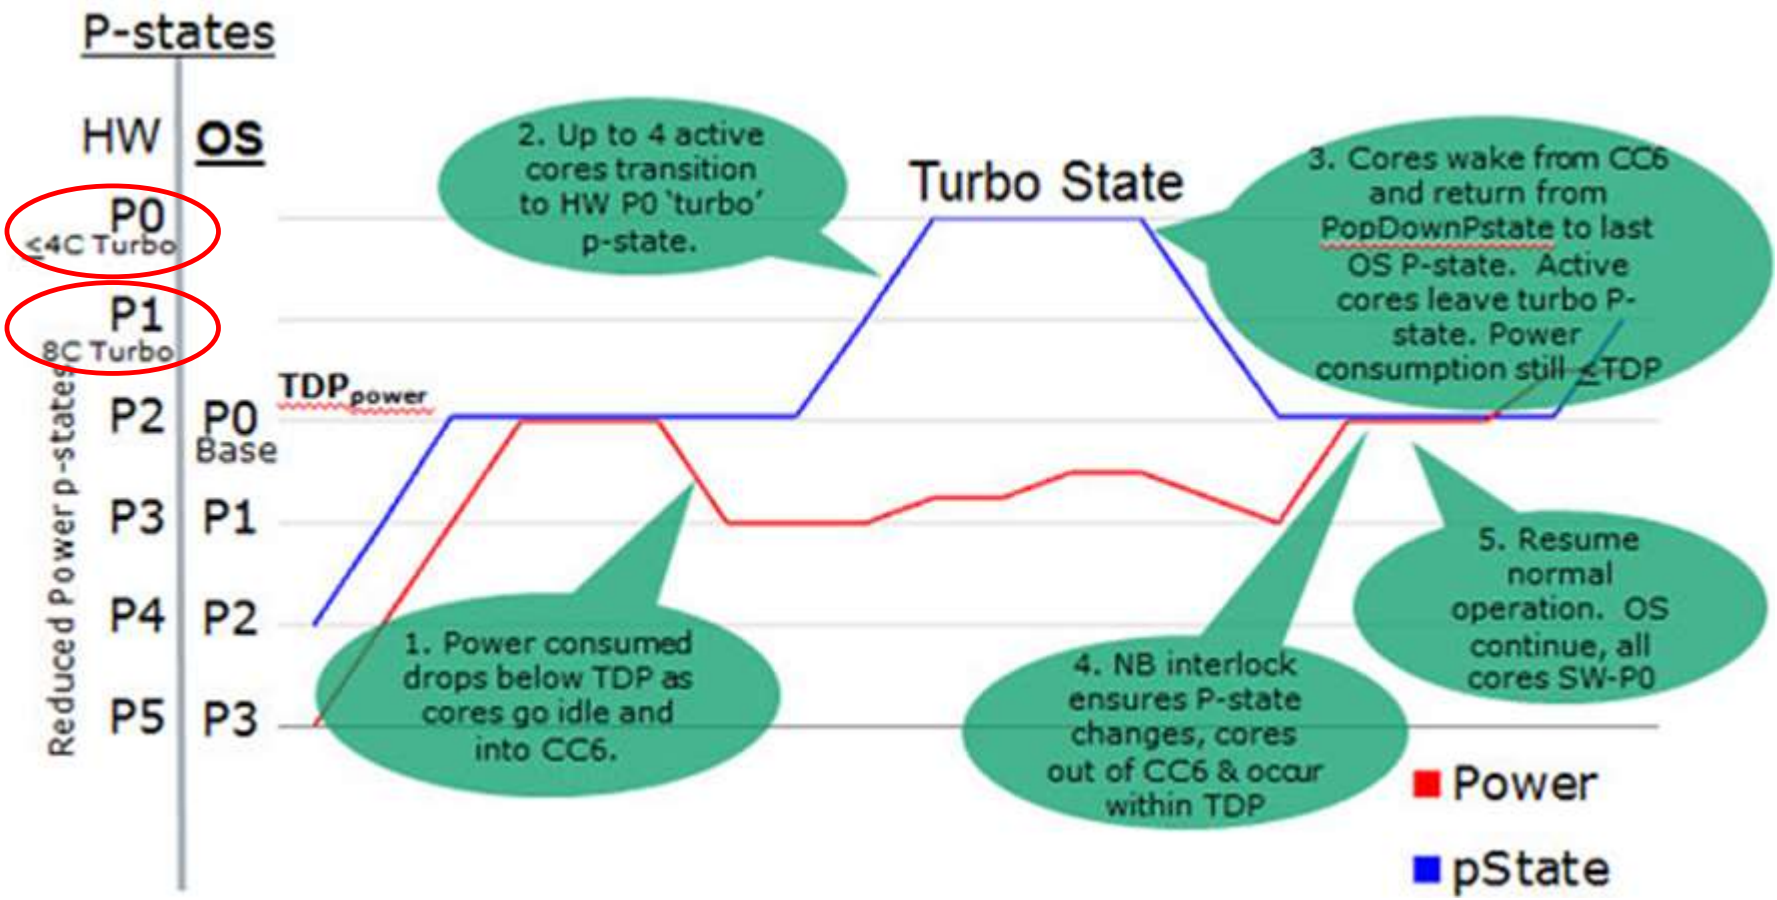

PopDownPstate: Core state saved into the memory when the core enters the CC6 state (Core C6 state)

2.6.2 The Turbo core technology of the Zambezi desktop line (7)

Example: Running a single threaded workload on the 8150 Zambezi DT with Turbo core enabled [36]

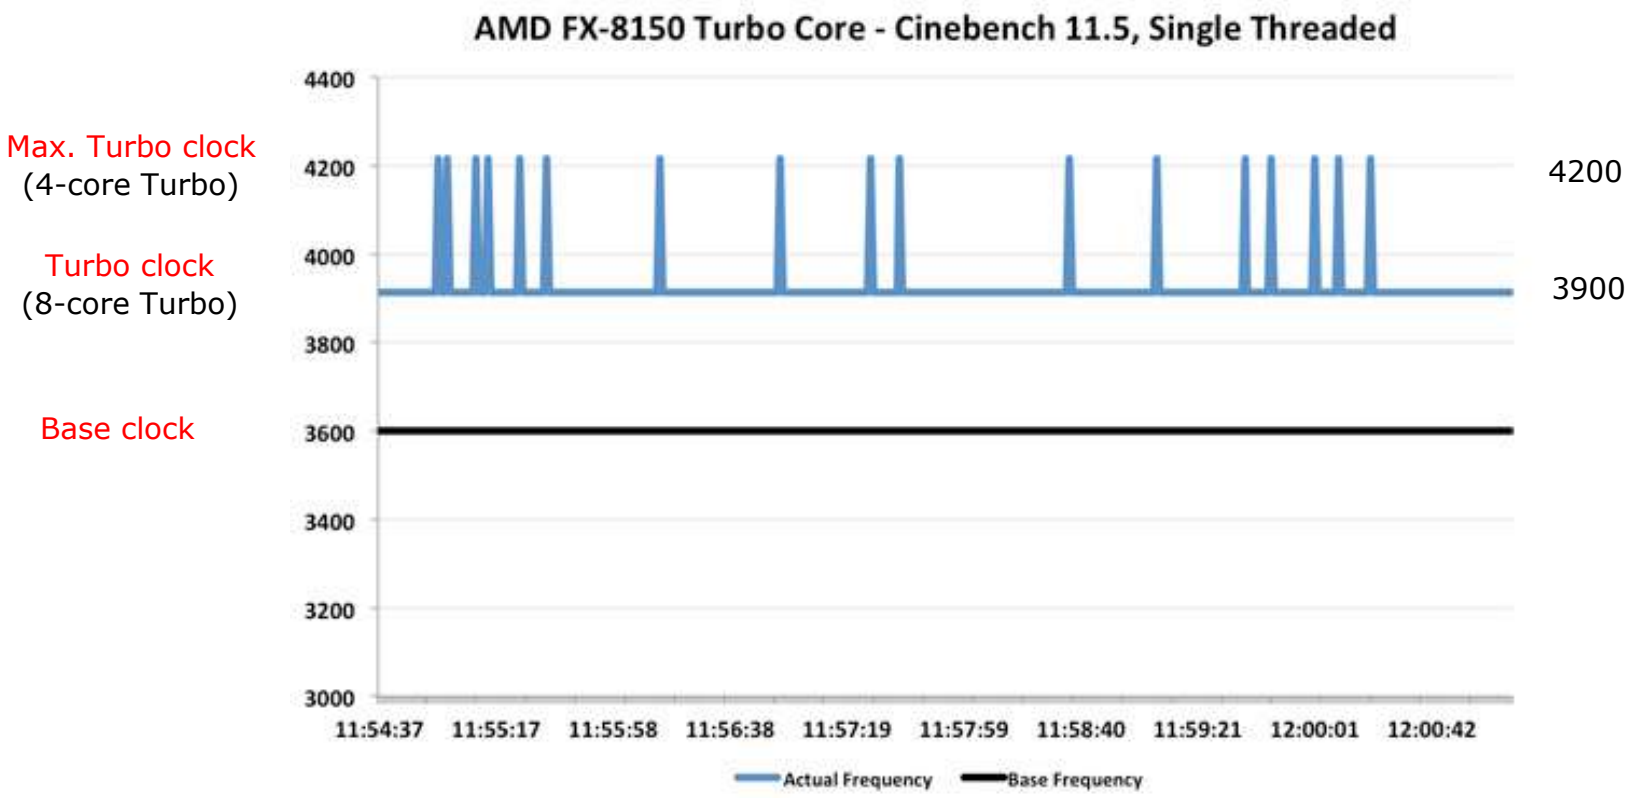

While running a single threaded workload, essentially seven of the 8 cores remain idle.
The processor runs most of the time at the Turbo core frequency (3.9 GHz for the FX-8150).
The average clock speed is 3.93GHz, 9% above the 3.6 GHz base clock of the FX-8150.

2.6.2 The Turbo core technology of the Zambezi desktop line (8)

Run time reduction achieved by enabling Turbo core for a single threaded workload running on an FX-8150 (Zambezi) [38]

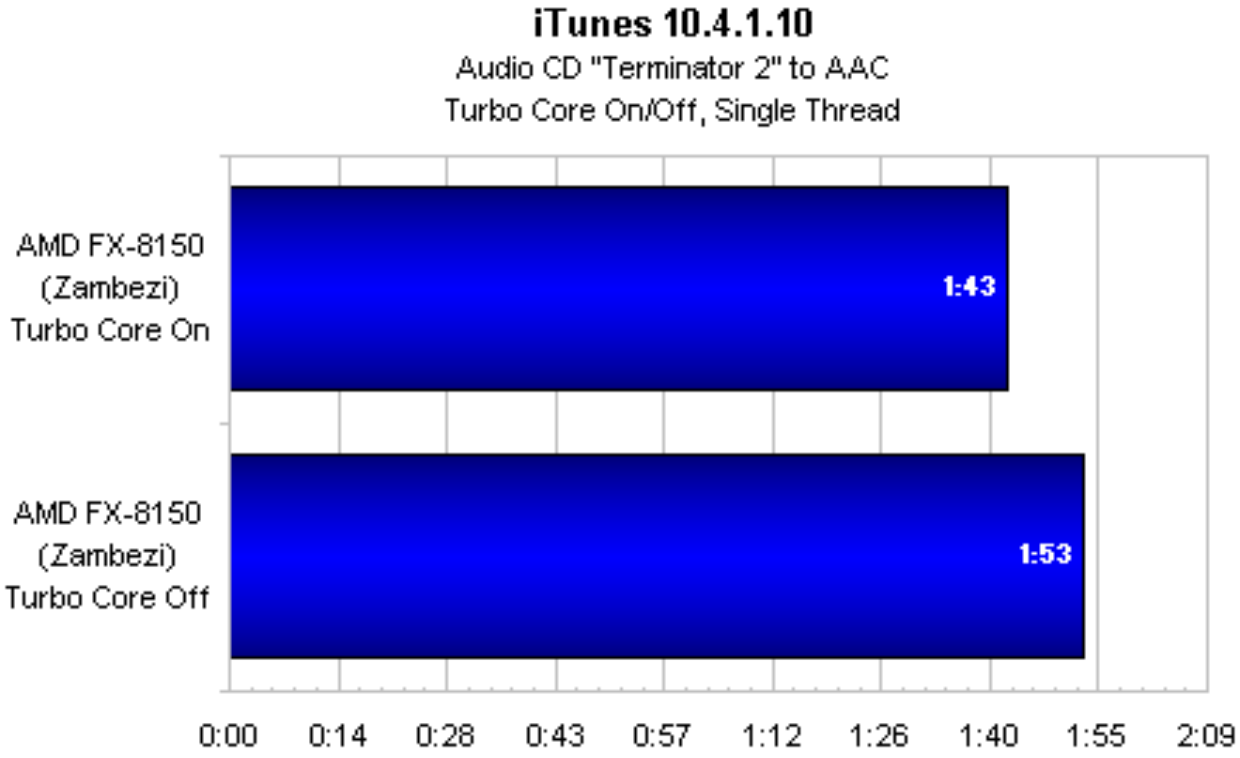

In the single threaded example, 7 of the 8 cores remain typically idle, and with the Turbo core mode enabled, the processor runs mostly at the Turbo frequency, and partly at the Max. Turbo frequency.

This result in a run time reduction of about 10 s (~ 7 %) while the Turbo core mode is activated.

2.6.2 The Turbo core technology of the Zambezi desktop line (9)

Run time reduction achieved by enabling Turbo core for a multi-threaded workload running on an FX-8150 (Zambezi) [38]

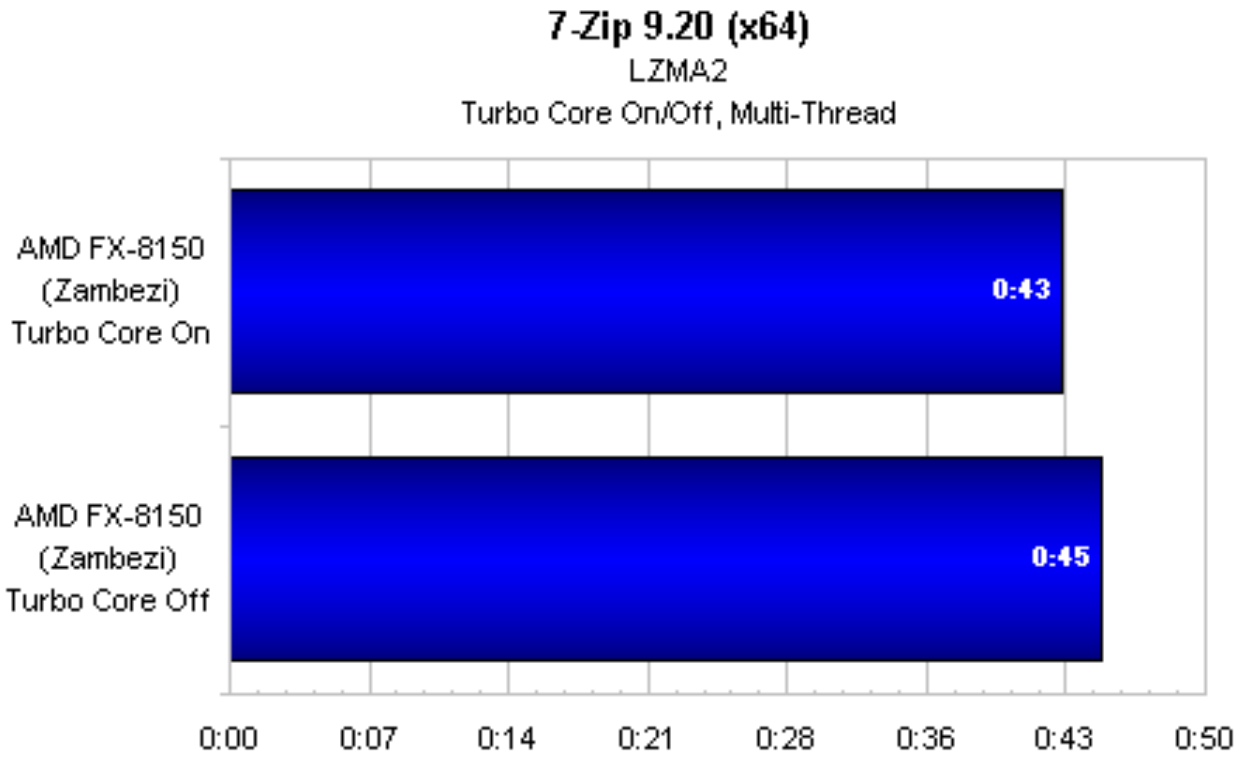

The multi-threaded workload is spread across all 8 cores, and if Turbo core is enabled, clock frequency alternates between the base clock of 3.6 GHz and the (8-core) Turbo frequency of 3.9 GHz.

The resulting run time reduction is about 0.2 s (~ 4 %), much less than for a single threaded workload.

2.6.2 The Turbo core technology of the Zambezi desktop line (10)

Remark

The efficiency of Turbo core is affected also by the scheduler of the OS, as discussed in Section 2.6.3.

2.6.2 The Turbo core technology of the Zambezi desktop line (11)

Contrasting the operation of AMD's 2. gen. Turbo core with that of Intel's Turbo Boost technology, as implemented in Sandy Bridge-based desktops (i5-2500K) [36]

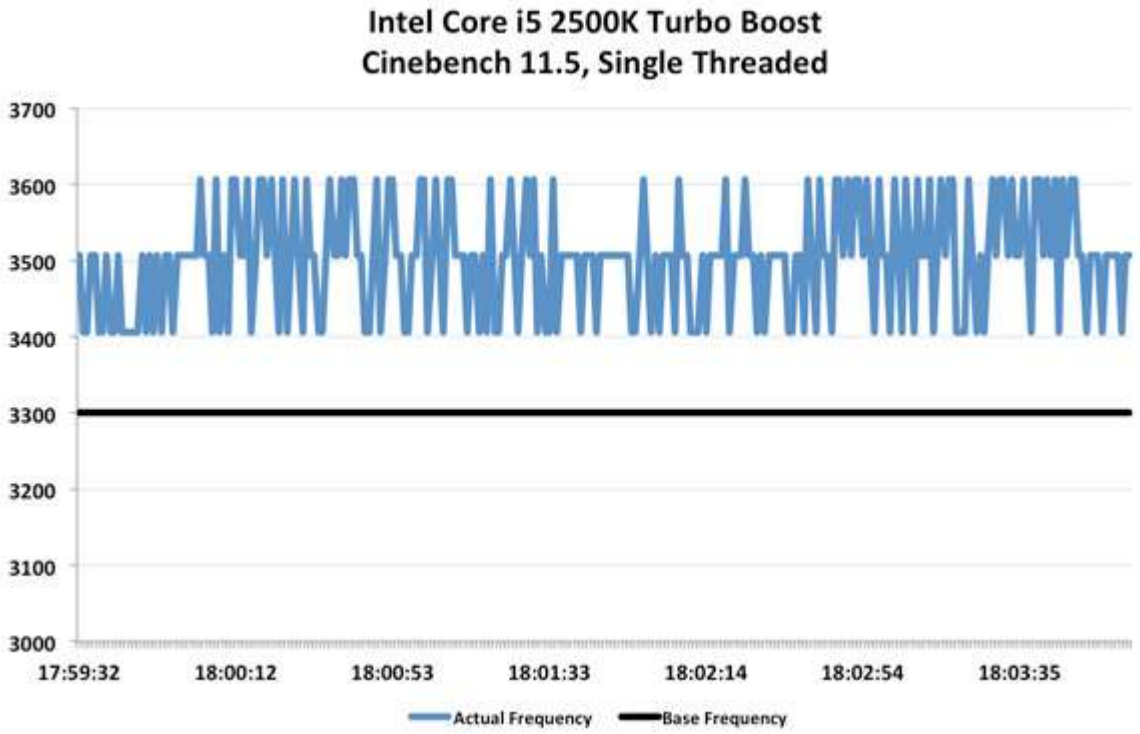

Intel's Turbo Boost implementation gives rise to a more frequent fluctuation than in case of AMD's Turbo core.

The average clock frequency remains at 3.5 GHz, only about 6 % higher over the base frequency.

So it seems that AMD's Turbo core technology, at least in the example shown, is more efficient than Intel's Turbo Boost.

2.6.2 The Turbo core technology of the Zambezi desktop line (12)

Remark

Brief comparison with Intel's Turbo Boost implementations

a) Precursor of Intel's Turbo Boost: EDAT-1

(Enhanced Dynamic Acceleration Technology)

- Introduced in Penryn-based 2-core mobiles in 2008, along with the DPD technology (Deep Power Down Technology).
- The DPD technology is activated by the OS (through the MWAIT API) if a core is „long enough" idle.

"Long enough" will be decided by the OS based on a heuristics to prevent situations when saving and restoring needs more power than gained by entering this state.

When activated the MWAIT API lets flush the L2 cache, save the core state of the idle core into an SRAM that has a private power supply and then reduces core voltage to a very low level.

Thus entering the DPD state assures a power headroom for increasing the clock frequency of the active core.

Principle of operation of Intel's Deep Power Down technology [39]

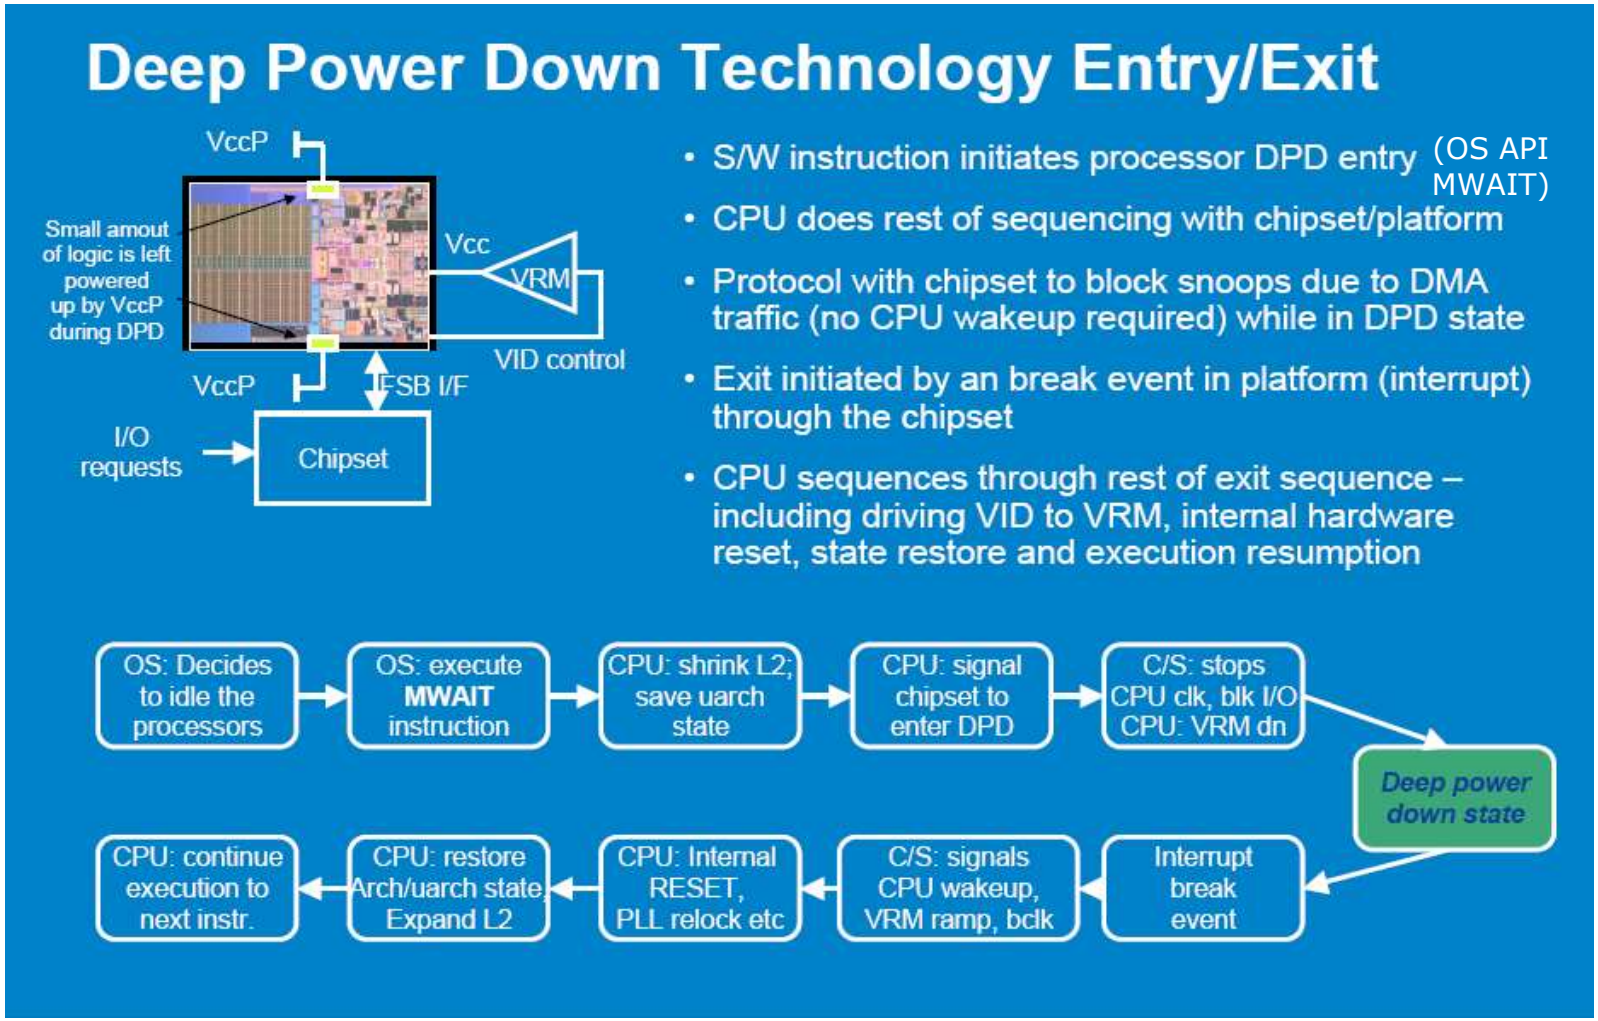

2.6.2 The Turbo core technology of the Zambezi desktop line (14)

a) Precursor of Intel's Turbo Boost: EDAT-2

The EDAT technology

- If one of the cores becomes idle and enters the C3 state or deeper, and the OS requests the highest performance state for the active core,
  the clock frequency of the active core will be raised by a single turbo bin (133 MHz).

2.6.2 The Turbo core technology of the Zambezi desktop line (15)

b) Intel's 1. gen. Turbo Boost

- Introduced in 1. gen. Nehalem processors (such as the 4-core Bloomfield desktops in 2008 for mobiles and desktops), along with
  - Integrated Power Gates (for VCC) to reduce leakage current to near zero, and a
  - Power Control Unit (integrated microcontroller of the complexity of a 486) that has real time sensors for current, voltage and temperatures, samples these values in 5 ms intervals, controls Turbo Boost based on sophisticated algorithms [40], [41].

  - If the OS requests an active core to increase fc beyond the TDP limited maximum frequency (i.e. to enter the PO state),
    and there is available power headroom
    - either by having idle cores
    - or a lightly threaded workload

    the Power Control Unit will increase the core frequency of the active cores

    provided that the power consumption of the socket and junction temperatures of the cores do not exceed the given limits.
  - In turbo mode all active cores in the processor will operate at the same fc and voltage.
  - There are only 2 turbo bins available for boosting the clock frequency (2x133 MHz).

2.6.2 The Turbo core technology of the Zambezi desktop line (16)

c) Intel's enhanced 1. gen. Turbo Boost

- Introduced in 2. gen. Nehalem processors (such as the 4-core Bloomfield desktops in 2009 for all processor categories.
- The enhancement is that there are more than two turbo bins (2x 13 MHz) available for raising core frequency.

2.6.2 The Turbo core technology of the Zambezi desktop line (17)

Available Turbo Boost bins (133 MHz) for the 1. and 2. gen. Nehalem processors [38]

Processor Number	Frequency	4 Cores Active	3 Cores Active	2 Cores Active	1 Core Active
2. gen Nehalem (Lynnfield-based) (2009)					
Core i7-870	2.93 GHz	2	2	4	5
Core i7-860	2.8 GHz	1	1	4	5
Core i5-750	2.66 GHz	1	1	4	4
1. gen. Nehalem (Bloomfield-based) (2008)					
Core i7-975	3.33 GHz	1	1	1	2
Core i7-950	3.06 GHz	1	1	1	2
Core i7-920	2.66 GHz	1	1	1	2

d) Intel's 2. gen. (Next gen.) Turbo Boost (Dynamic Turbo Boost)

- Introduced with the Sandy Bridge line of mobile and desktop processors in 2011.
  These processors incorporate up to 4 cores and a GPU.
- It allows to use the energy budget that became accumulated during idle periods for boosting fc such that the power consumption raises beyond TDP for a short period of time [8].

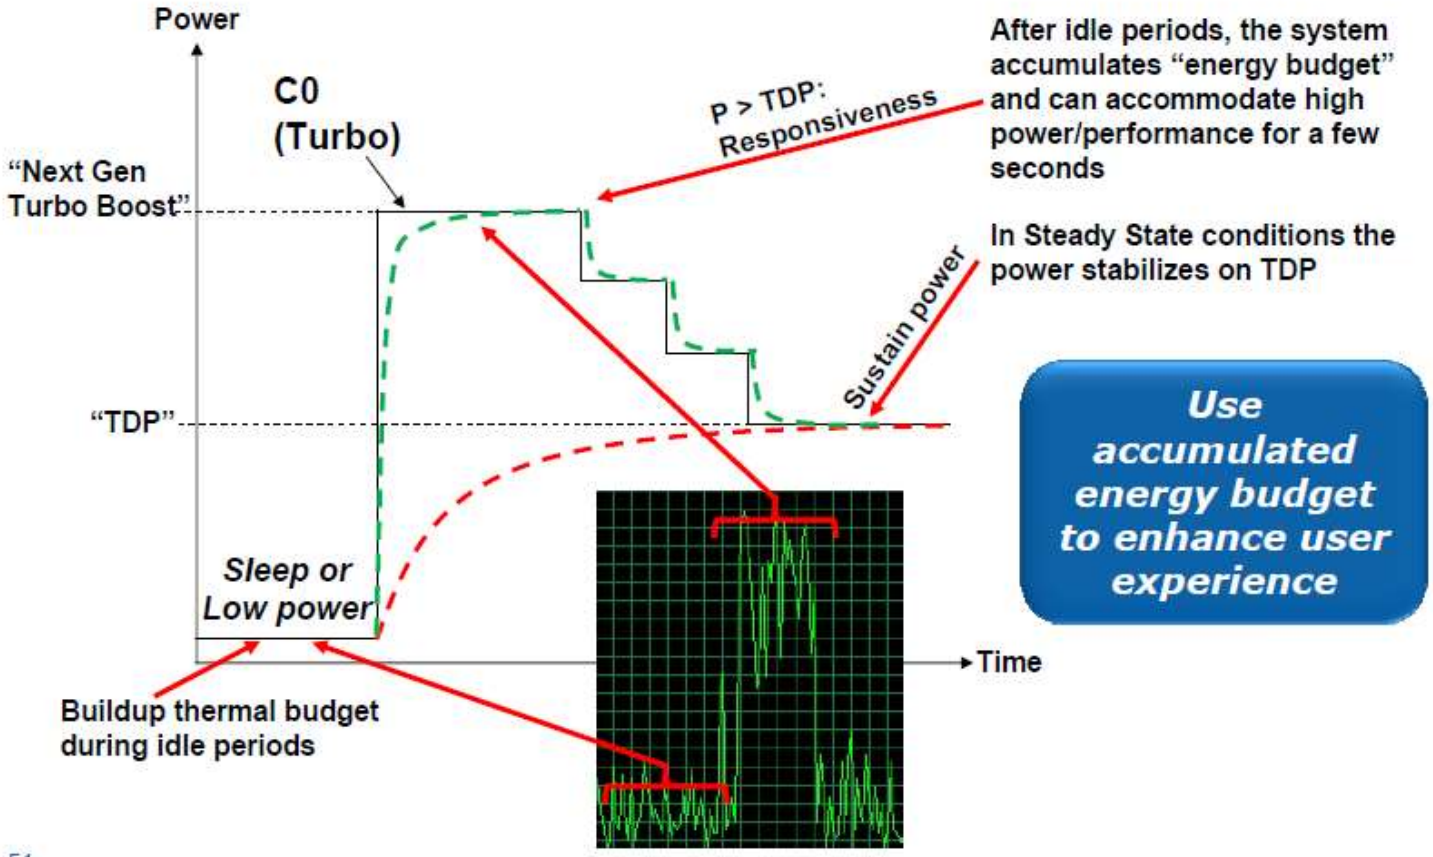

2.6.2 The Turbo core technology of the Zambezi desktop line (19)

Contrasting the introduction of Intel's and AMD's Turbo and Power gating technologies

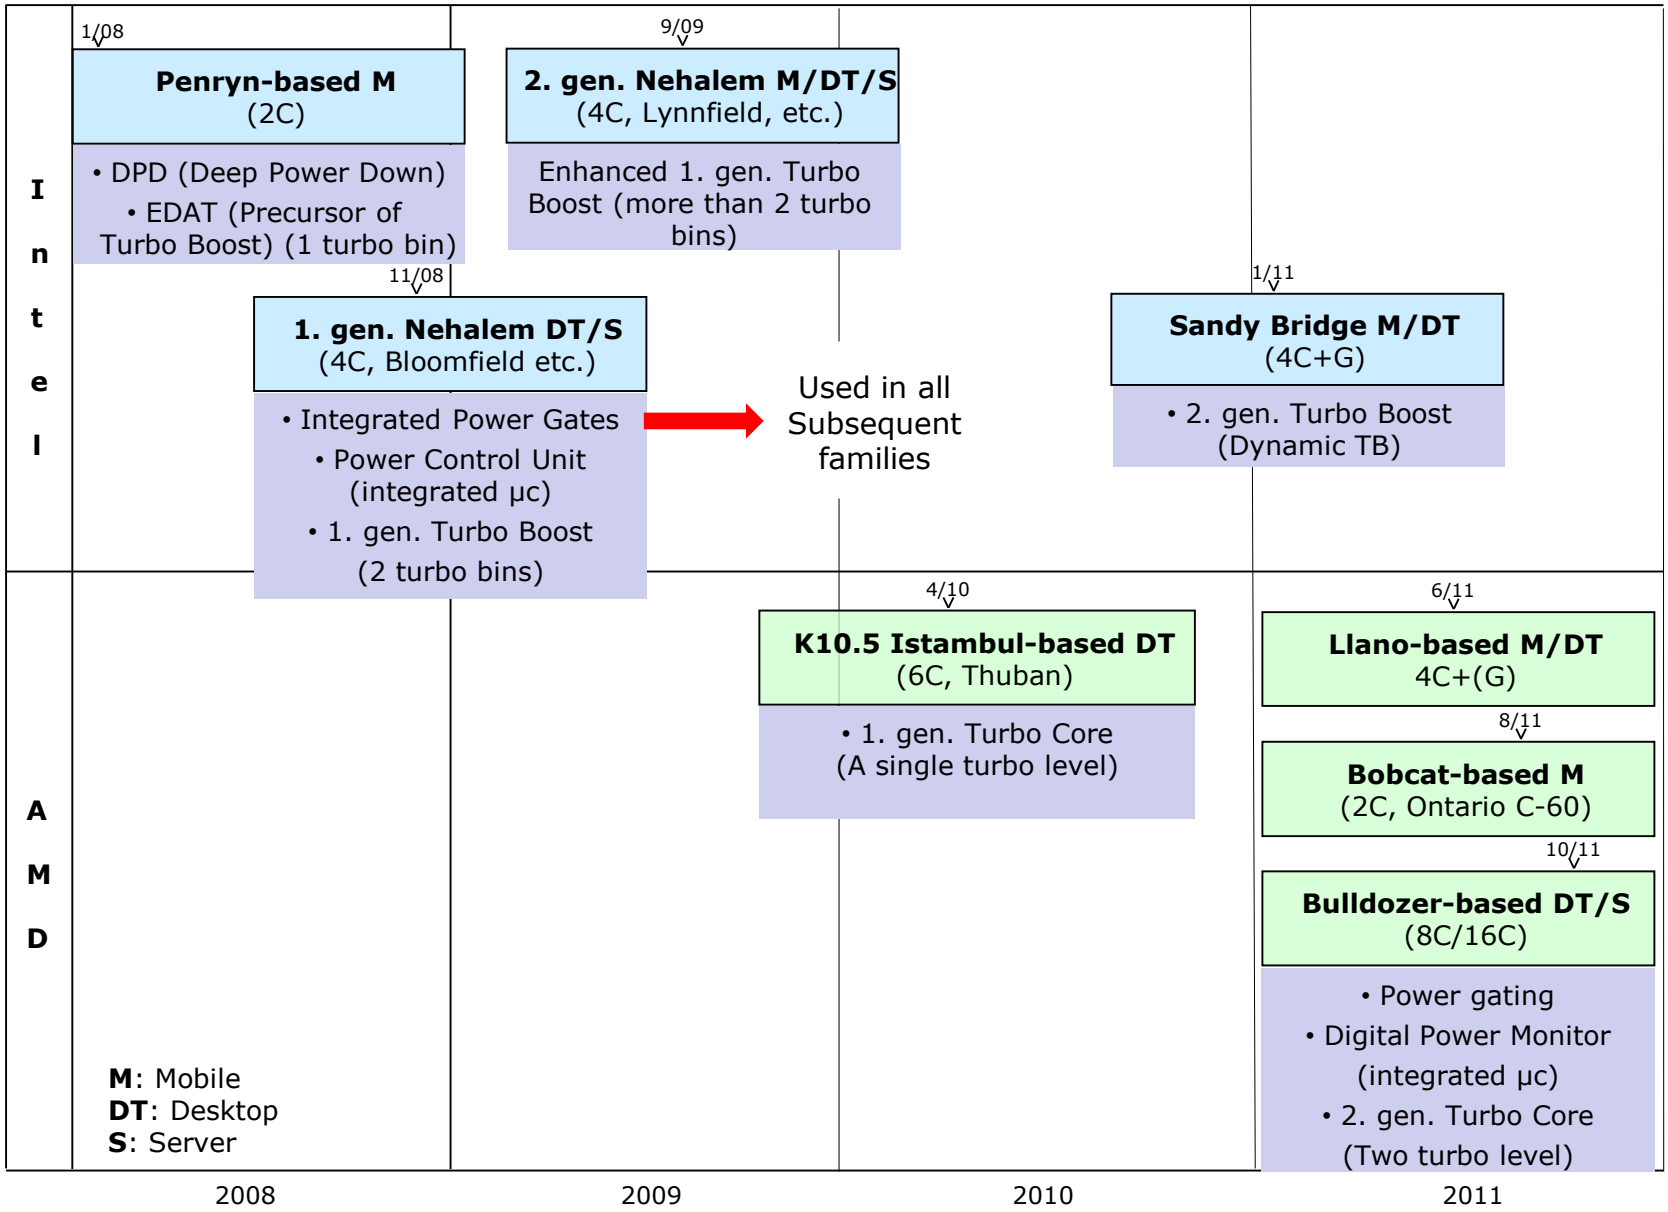

2.6.2 The Turbo core technology of the Zambezi desktop line (20)

Evolution of Intel's Turbo technology [34]

Mobile Desktop	Merom/Penryn (Mobile only)	Nehalem/Westmere		Sandy Bridge
		Clarksfield Lynnfield/Clarkdale	Arrandale	
Control	• CPU Core C-state •Digital power meter	• CPU Core C-states • CPU Power - Platform iMon	• CPU Core C-states • CPU Power- Platform iMon • PG Power- Platform iMon • Package Power	• CPU Core C-states • CPU/ PG/ Package power • Built-in power monitoring •Power Budget Management •Platform Control (EC / VR)
Key New Capabilities	• 1-2 turbo bin when other core is asleep	• Turbo controlled within power limit • Multi-core turbo • More turbo if cores are asleep	• PG dynamic frequency • Driver controlled power sharing between CPU and PG (Mobile)	• HW controlled power sharing between CPU - PG • Brief turbo above TDP → dynamic Turbo •More platform control via PECI 3.0 and SVID
Turbo Behavior **Illustrative only. Does not represent actual number of turbo bins.**	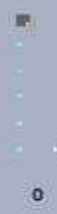	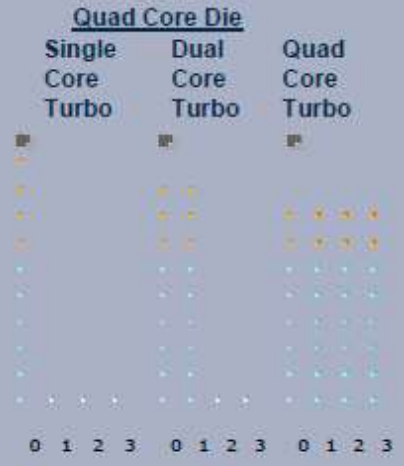	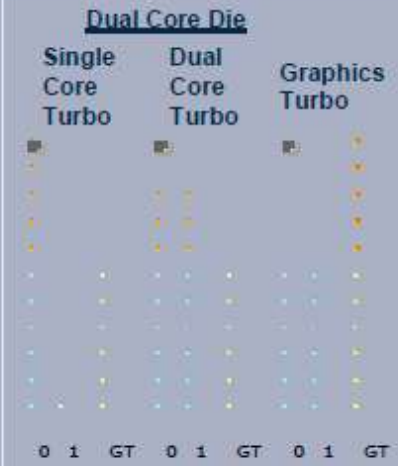	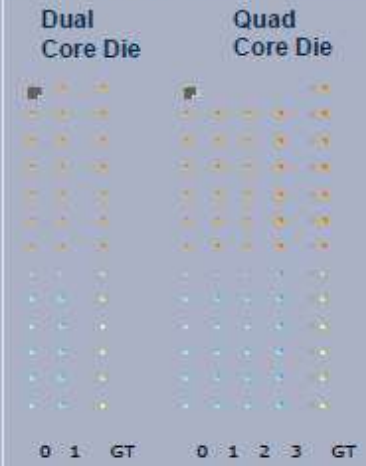

2.6.2 The Turbo core technology of the Zambezi desktop line (21)

As indicated in the previous slide both in the Turbo and the Power gating technologies Intel has a lead of about two years.

2.6.3 Performance assessment of the Zambezi desktop line

2.6.3 Performance assessment of the Zambezi desktop line

There are many benchmark investigations related to AMD's Zambezi, e.g. [31], [32]. Below we show key results of the very extensive report [32] covering a wide range of application areas, including

- synthetic benchmarks
- audio processing
- video processing
- image processing
- packing data
- rendering
- games.

2.6.3 Performance assessment of the Zambezi desktop line (2)

Summary benchmark results including all tests excl. games [32]

Price [US-Dollar]	AMD	Intel
317	-	Core i7 2600K
294	-	Core i7 2600(S)
245	FX-8150	-
216	-	Core i5 2500K
205	FX-8120	Core i5 2500(T/S)
184	-	Core i5 2400
177	-	Core i3 23xx
165	FX-6100	-

HTT: Hyperthreading
CMT: Core-Multithreading (AMD's module concept)
Kerne: Cores

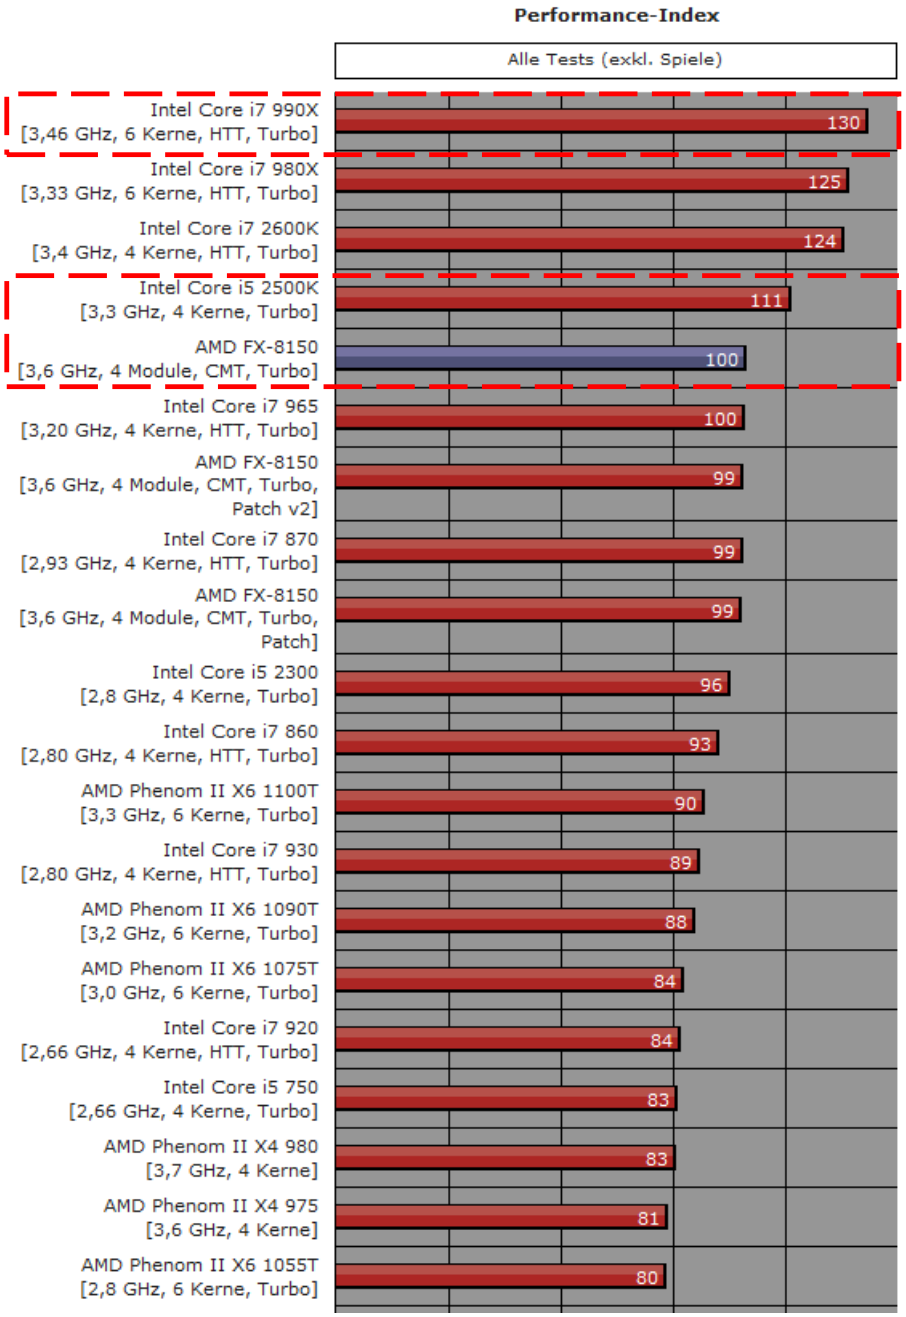

2.6.3 Performance assessment of the Zambezi desktop line (3)

Summary performance assessment of Zambezi-1

a) AMD's Bulldozer-based 4-module, 8-core FX 8150 flagship processor is far away from overtaking the performance leadership from Intel's Sandy Bridge based 6-core i7 990X.

The i7 990X provides about 30 % higher performance across all benchmarks (excl. games) than AMD's Bulldozer based FX 8150 flagship desktop processor, nevertheless for a considerable higher price (~ 1000 $ at the time of publishing the benchmark report).

The fact that Intel has no competition on the high end desktop market implies that Intel can determine high end desktop prices as high as the market it allows.

b) Comparable priced Sandy Bridge based processors have typically higher performance than AMD's FX 8150.

E.g. Although at the time of the cited benchmark review [33] Intel's Sandy Bridge based 4-core i5 2500K costs less than AMD's FX 8150, it performs about 10 % higher than AMD's FX 8150.

Other benchmark investigations reveal also that the Bulldozer-based Zambezi underperforms against Intel's Sandy Bridge-based desktop processors [33], [31].

2.6.3 Performance assessment of the Zambezi desktop line (4)

Remark

In order to take into account AMD's module concept Microsoft released two patches to Windows 7 (patch, patch v2) in cooperation with AMD [32].

Nevertheless, these patches actually did not improve the performance of the FX 8150 [32].

2.6.3 Performance assessment of the Zambezi desktop line (5)

Summary benchmark results including all tests excl. games [32]

Price [US-Dollar]	AMD	Intel
317	-	Core i7 2600K
294	-	Core i7 2600(S)
245	FX-8150	-
216	-	Core i5 2500K
205	FX-8120	Core i5 2500(T/S)
184	-	Core i5 2400
177	-	Core i3 23xx
165	FX-6100	-

HTT: Hyperthreading
CMT: Core-Multithreading (AMD's module concept)
Kerne: Cores

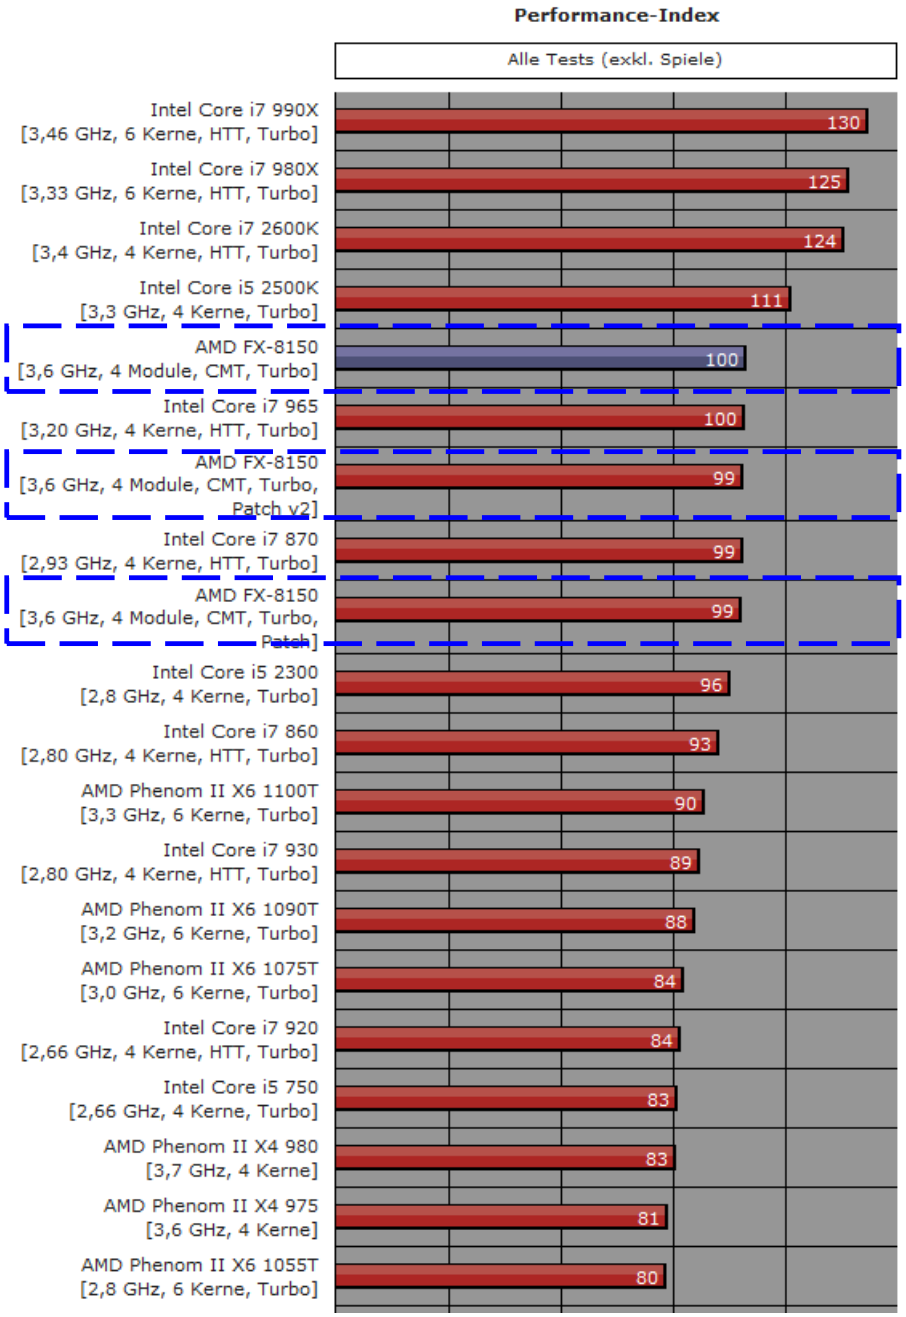

2.6.3 Performance assessment of the Zambezi desktop line (6)

Summary performance assessment of Zambezi-2

c) The Bulldozer-based 8-core (four modules) flagship FX 8150 achieves only a moderate gain (~ 10 – 20 %) vs. AMD's previous 4 to 6-core K10.5 (Phenom X6/Phenom X4) designs.

2.6.3 Performance assessment of the Zambezi desktop line (7)

Summary benchmark results including all tests excl. games [32]

Price [US-Dollar]	AMD	Intel
317	-	Core i7 2600K
294	-	Core i7 2600(S)
245	FX-8150	-
216	-	Core i5 2500K
205	FX-8120	Core i5 2500(T/S)
184	-	Core i5 2400
177	-	Core i3 23xx
165	FX-6100	-

HTT: Hyperthreading
CMT: Core-Multithreading (AMD's module concept)
Kerne: Cores

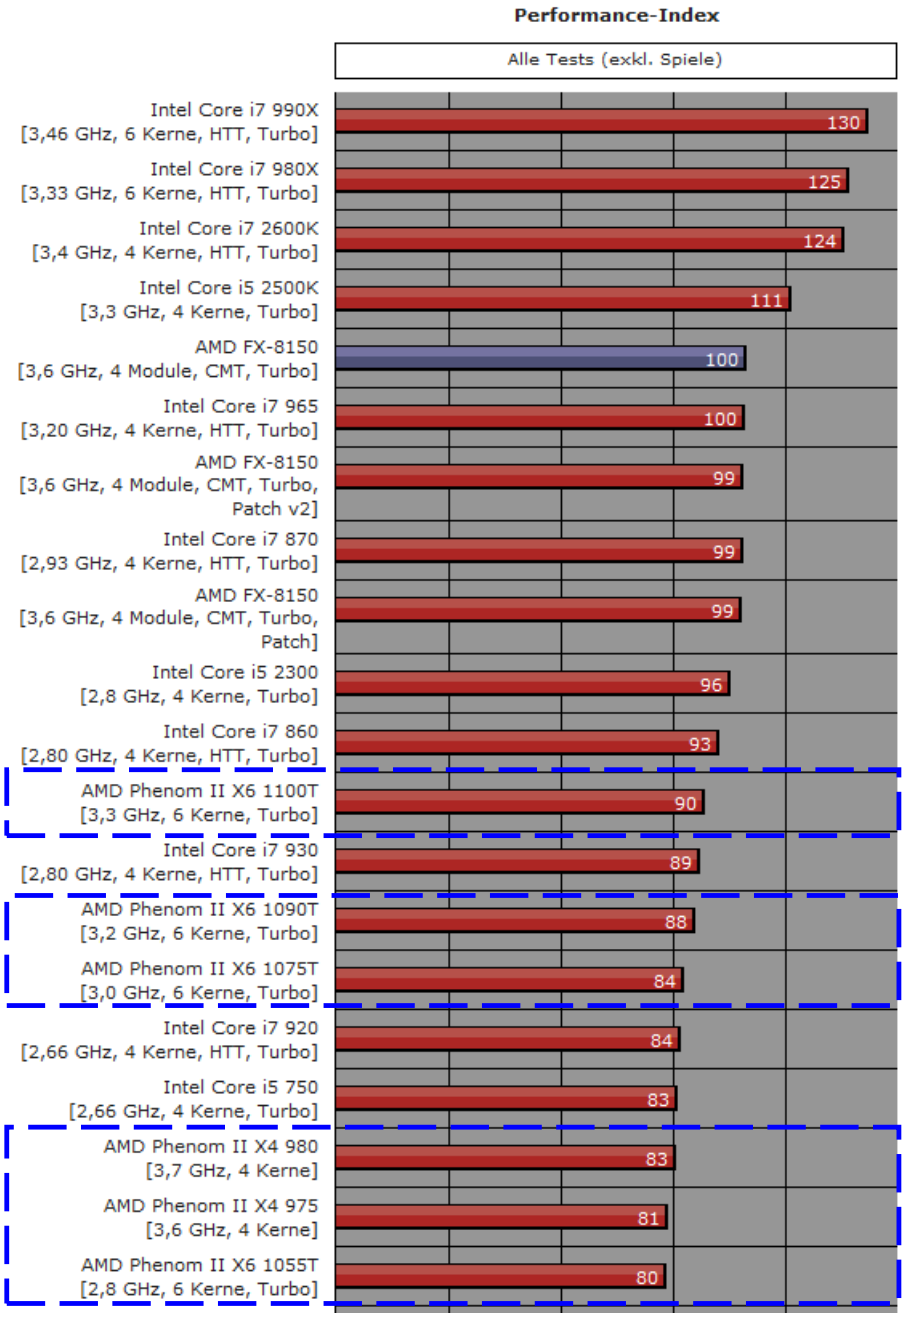

2.6.3 Performance assessment of the Zambezi desktop line (8)

Remarks

a) The scheduling policy of the OS affects the efficiency of the Turbo core technology and thus the achieved performance [3].

Windows 7 does not recognize the module structure of Bulldozer based processors nor the peculiarities of their Turbo core technology.

It will spread threads across modules preventing the activation of the max. turbo speed, since max. turbo speed can only be reached when at least half of the Bulldozer modules are idle (being in the C6 state).

Furthermore, the scheduler of Windows 7 re-schedules threads from time to time.

As a consequence, the processor can not reach it peak performance for workloads that do not utilize all 8 available cores.

2.6.3 Performance assessment of the Zambezi desktop line (9)

Example: Impact of Windows 7's scheduling policy to the activation of Max. Turbo mode [9]

Currently Windows® 7 is unaware of the shared nature contained within the AMD FX-8150 processor. As a result there are possibilities where opportunities for resource sharing or activate higher Turbo Core frequencies are missed.

Sub-Optimal

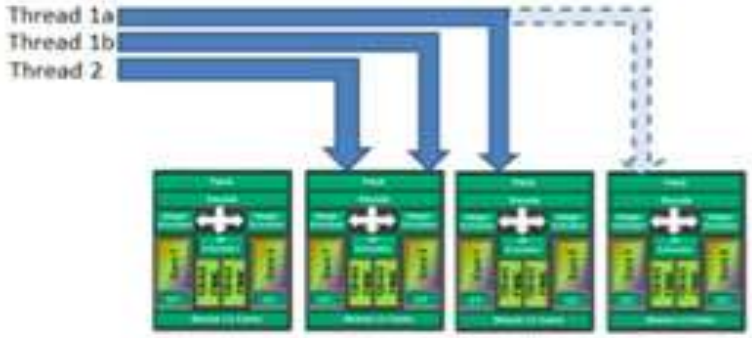

An example where Thread 1b relies on data from Thread 1a, and is scheduled on different Core modules.

Also in this example, the scheduler is assigning a re-iteration of Thread 1a to different Core modules so that max turbo mode can NOT be enabled

Optimal

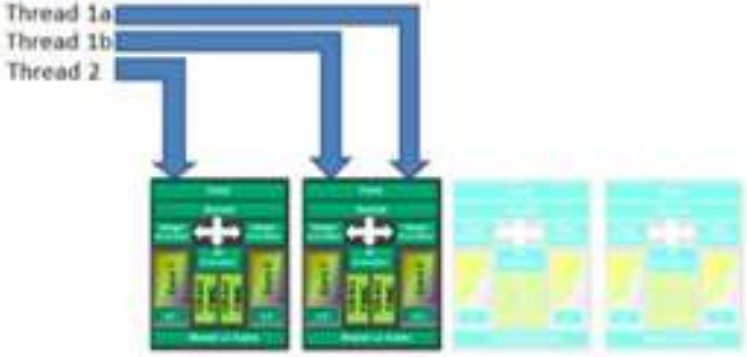

In the optimal scenario – Thread 1a and 1b are scheduled in the same Core module and the unused Cores are parked so that AMD Turbo Core Technology is enabled.

2.6.3 Performance assessment of the Zambezi desktop line (10)

b) There are OS patches available for the FX series – worked out by Microsoft in cooperation with AMD - to remedy this problem [32].

These patches are not too efficient, as benchmark results indicate it [32].

2.6.3 Performance assessment of the Zambezi desktop line (11)

c) Windows 8 has already a redesigned scheduler

The new scheduler takes already into account the module structure and Turbo core features of Bulldozer.

This allows a few percent performance boost, as indicated below for games [35].

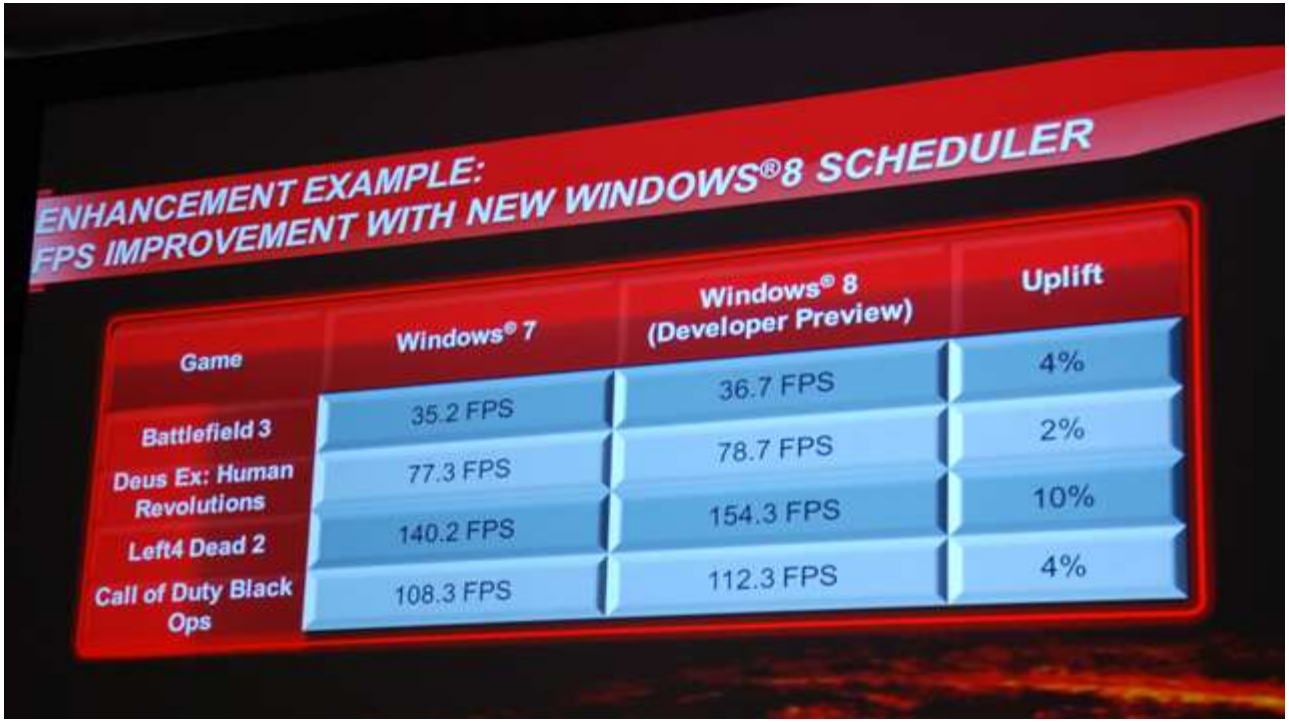

FPS: Frames Per Sec.

2.6.3 Performance assessment of the Zambezi desktop line (12)

Summary assessment of the benchmark results of the Zambezi FX 8150 line [32]

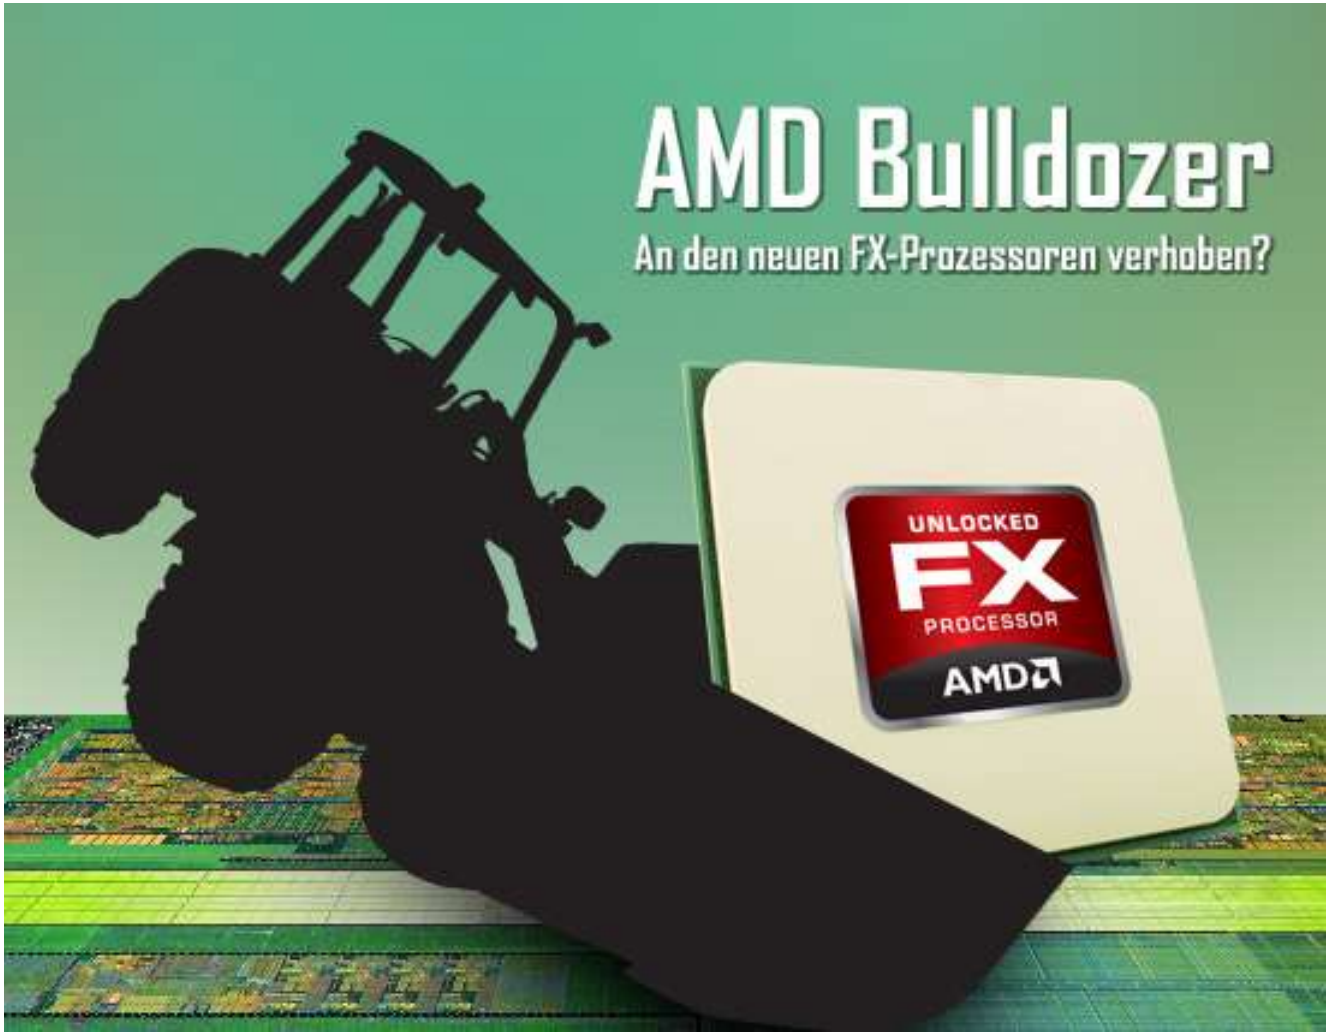

2.6.3 Performance assessment of the Zambezi desktop line (13)

Summary assessment of all Bulldozer based designs

All in all Bulldozer-based server and desktop lines were assessed by many market observers as “disappointing”, e.g. [3], [31], [32].

2.6.3 Performance assessment of the Zambezi desktop line (14)

Remark – **AMD's reorganization after the Bulldozer disaster**

- After Intel announced their Sandy Bridge (1/2011) AMD's Board of Directors pressured AMD's CEO Dirk Meyer to step down.

  Dirk Meyer was originally a processor architect (co-architect of DEC's 21064 and 21264 processors and chief architect of AMD's highly sussessful Athlon (K7) processor).
- In 8/2011 he was followed by Lenovo's former CEO Rory Read.
- Read reorganized AMD and laid off 1400 employees out of a work force of about 40000 in 11/2011 [42], [43].

3. Second generation Piledriver-based (Family 15h Models 30h-3Fh) processor lines

- 3.1 Overview of the Piledriver-based processor lines
- 3.2 The Piledriver Compute Module
- 3.3 Piledriver-based GPU-less processor lines
- 3.4 The Trinity APU lines
- 3.5 The Richland APU lines

3.1 Overview of the Piledriver-based processor lines

3.1 Overview of the Piledriver-based processor lines [based on 1]

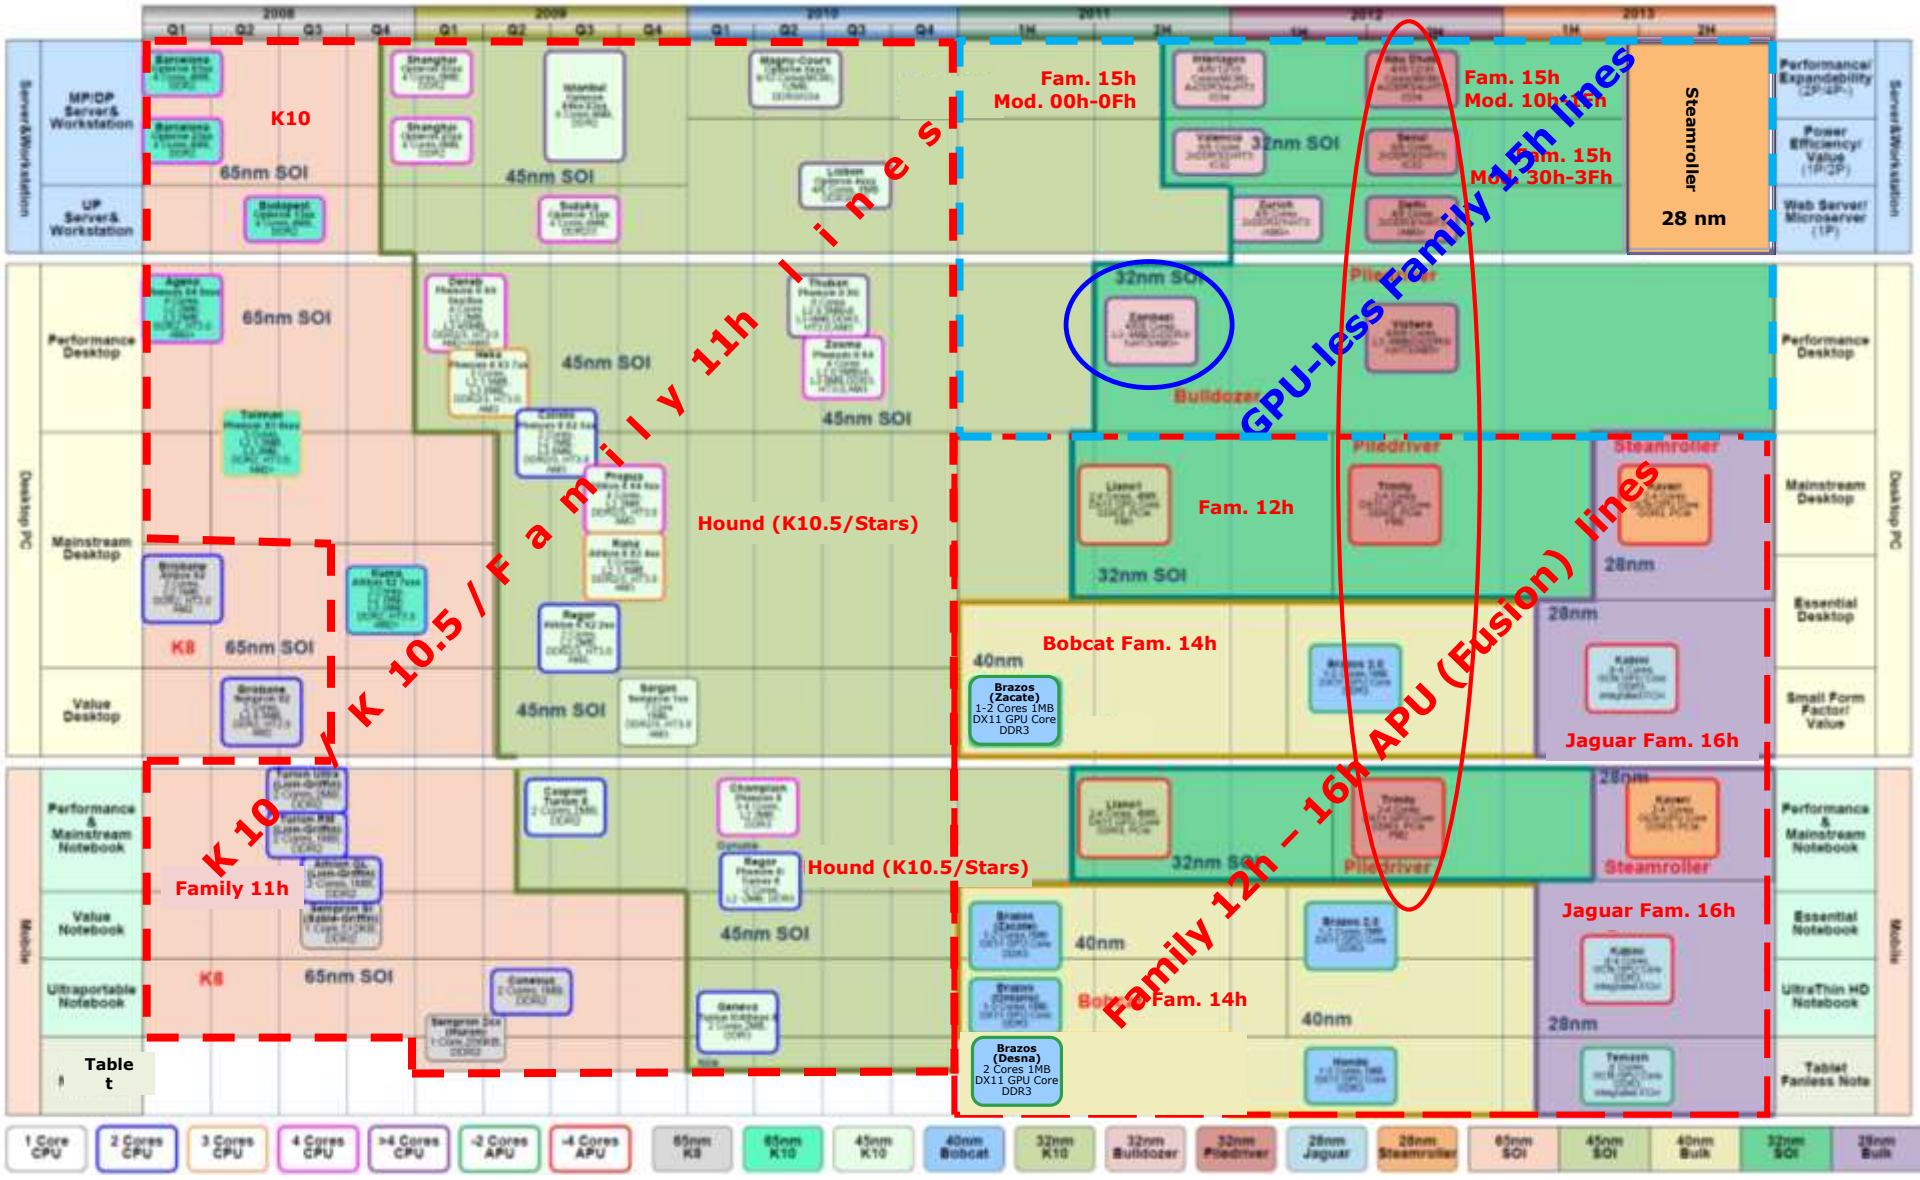

3.1 Overview of the Piledriver-based processor lines (2)

Overview of AMD's Piledriver-based server, desktop and mobile lines

	Launched in	2011	2012	2013	2013	2015	2016
		Family 15h (00h-0Fh) (**Bulldozer**)	**Family 15h** (10h-1Fh) (**Piledriver**)	**Family 15h** (10h-1Fh) (**Piledriver v.2**)	**Family 15h** (30h-3Fh) (**Steamroller**)	**Family 15h** (60h-6Fh) (**Excavator v.1**)	**Family 15h** (77h-3Fh) (**Excavator v.2**)
Servers	**4P servers** (85-140 W)	Interlagos	Abu Dhabi				
	2P servers (85-140 W)	Valencia	Seoul				
	1P servers (85-140 W)	Zurich	Delhi				
Desktops	**High perf.** (~95-125 W)	Zambezi FX-Series	Vishera FX-Series				
	Mainstream (~65-95 W)		Trinity A10-A4	Richland A10/A8/A6/A4	Kaveri A10/A8		
Notebooks	**Mainstream** (~25-35 W)		Trinity A10/A8/A6M	Richland A10/A8/A6M	Kaveri FX/A10/A8P		Bristol Ridge FX/A12/A10P
	Ultra-thin (~10 – 15 W)		Trinity A10/A6M	Richland A10/A8/A6/A4M	A8 Pro/A8(B) A6 Pro/A6(B)	Carrizo FX/A10/A8P	Bristol Ridge FX/A12/A10P Stoney Ridge A9/A6
Tablets (~5 W)							

Piledriver-based processor lines

Piledriver-based processor lines

Piledriver-based GPU-less processor lines
(Section 3.3)

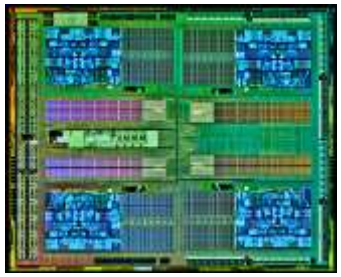

32 nm
315 mm²
1.2 billion trans.

Abu Dhabi
Opteron server line
(Dual chips)

4P: Opteron 6300

G34 socket
(retained)

11/2012

Vishera
high performance
FX desktop line

FX-x3xx

AM3+ socket
(retained)

Volan platform

10/2012

Piledriver-based Trinity APU lines
(Section 3.4)

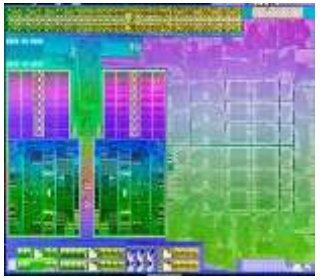

32 nm
226 mm²
1.3 billion trans.

Trinity
mainstream
DT APU line

HD 7000D GPU

A10 - A4 5x00(K)

FM2 socket
(new)

Virgo platform

9/2012

Trinity
mainstr./ultra-thin
mobile APU line

HD 7000G GPU

A10 – A6 4xxxM

FS1r2/FP2 socket
(new)

Comal platform

5/2012

Piledriver 2. gen. based Richland APU line
(Section 3.5)

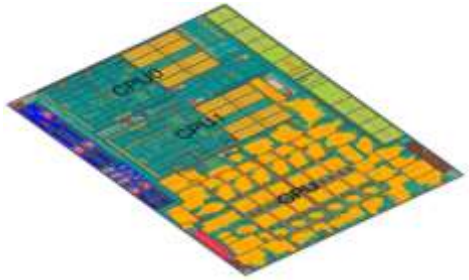

Richland
mainstream
DT APU line

HD 8000D GPU

A10 – A4 6x00(K)

FM2 socket
(retained)

Elite A-Series

6/2013

Richland
mainstr./ultra-thin
mobile APU lines

HD 8000M GPU

A10 – A4 5xxxM

FS1r2 socket
(retained)

Elite performance platform

5/2013

3.2 The Piledriver Compute Module

- 3.2.1 Overview of the Piledriver Compute Module
- 3.2.2 Piledriver’s performance enhancements vs. Bulldozer
- 3.2.3 Piledriver’s power management enhancements vs. Bulldozer
  - 3.2.3.1 A brief introduction into clock distribution networks
  - 3.2.3.2 Principle of the Resonant Clock Mesh (RCM) technology
  - 3.2.3.3 The evolution of implementing RCM

3.2.1 The Piledriver Compute Module

3.2.1 The Piledriver Compute Module

The Piledriver Compute Module includes two cores like the Bulldozer Compute Module, but is a thorough redesign of the ill fated Bulldozer Compute Module [54].

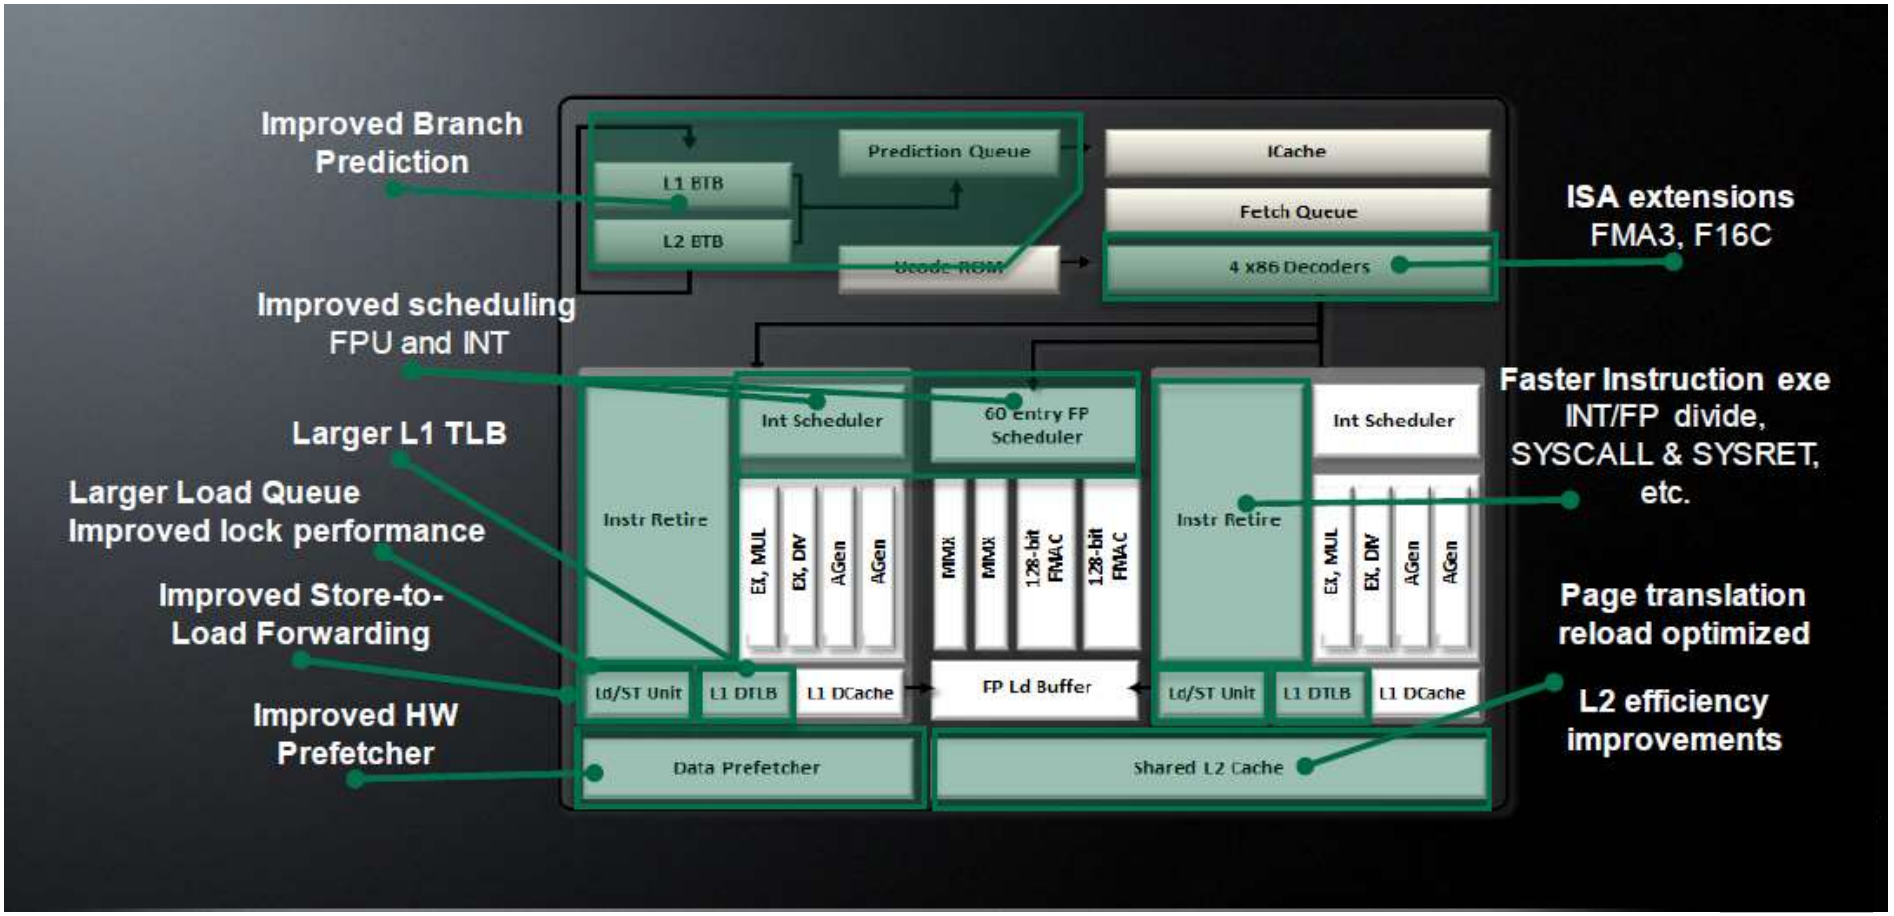

3.2.2 Piledriver’s performance enhancements vs. Bulldozer

3.2.2 Piledriver's performance enhancements vs. Bulldozer [54]

The Piledriver Compute Module includes two cores like the Bulldozer Compute Module, but is a thorough redesign of the ill. fated Bulldozer Compute Module.

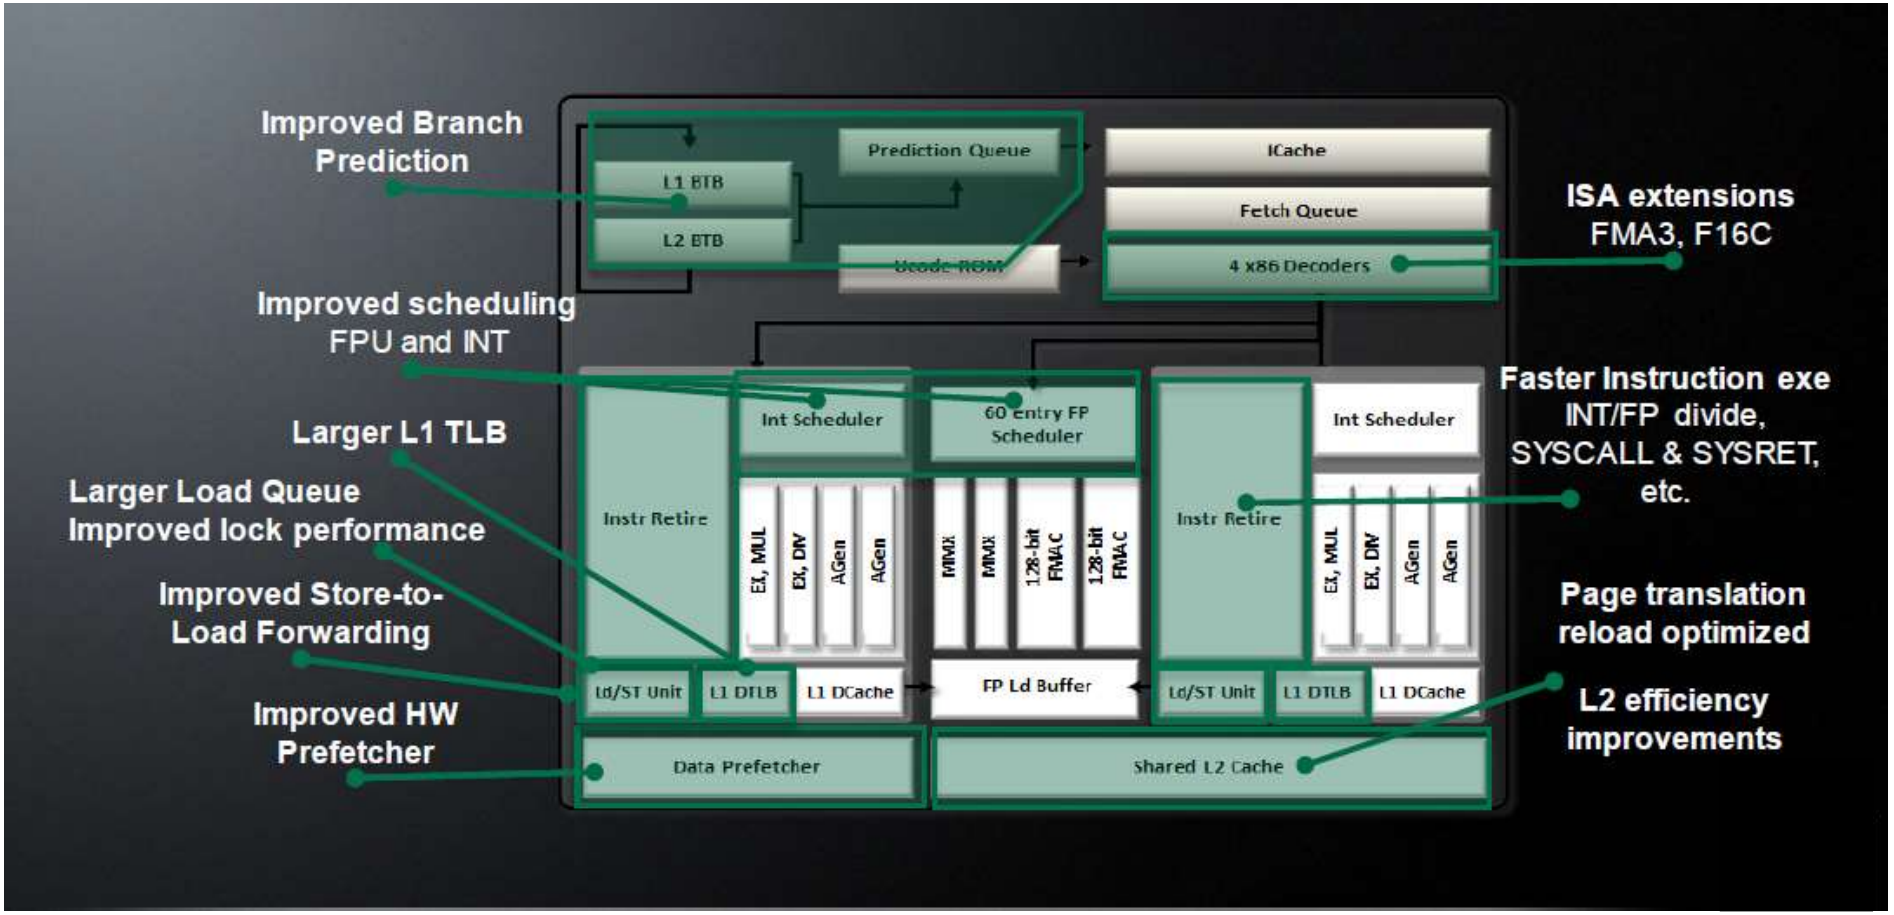

Piledriver's performance enhancements vs. the (Fam. 12h) Husky and Bulldozer cores [55]

32NM "PILEDRIVER" COMPUTE MODULE

X86 CORE REDESIGN

AMD

- Shared fetcher, decoder, floating point unit and L2 within a compute module
- 2 cores and up to 2 MB L2 cache per compute module
- ISA additions: AVX, AVX1.1, FMA3, AES and F16C
- Lightweight profiling support in HW
- "Piledriver" enhancements over "Bulldozer":
  - IPC improvement, leakage reduction, CAC reduction, frequency uplift
- "Piledriver" performance enhancements over "Husky"
  - 26% better system performance for desktop[5]
  - 29% increase in productivity for notebook[2]
  - AMD Turbo Core Technology 3.0

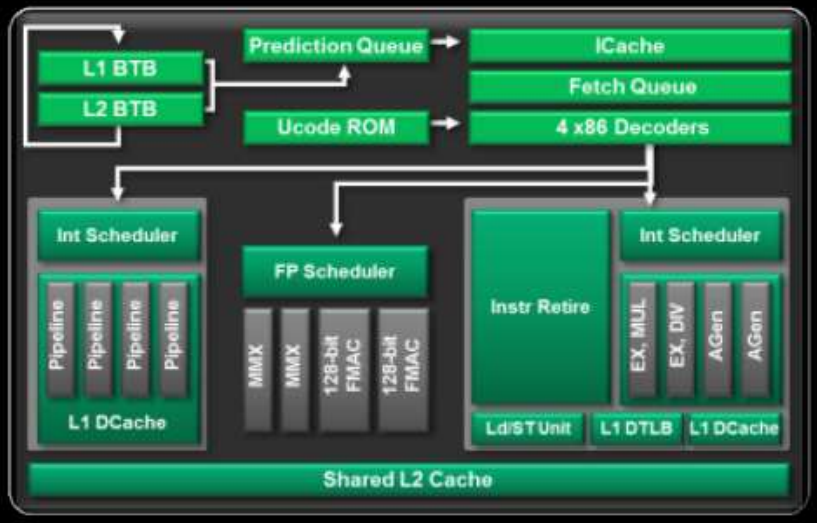

17W, 25W, 35W Notebooks, 65 and 100W Desktops

2.0–3.8 GHz CPU

424–800 MHz GPU

3.2.2 Piledriver's performance enhancements vs. Bulldozer (3)

Remark

A detailed description of Piledriver's improvements and enhancements can be found in [56].

3.2.3 Piledriver’s power management enhancements vs. Bulldozer

3.2.3 Piledriver's power management enhancement vs. Bulldozer [63]

- Along with the Piledriver design AMD introduced the Resonant Clock Mesh technology (RCM) in order to reduce power consumption of the clock distribution network.

  Reduced overall power consumption can be utilized also to increase clock frequency within the same TDP limit.
- Announcement of RCM: in 2/2012 at the ISSCC.
- As the RCM technology aims at reducing the power consumption of clock distribution networks, first we provide a brief overview about them.

  Then we will discuss the principle of operation and the introduction of RCM.

3.2.3.1 A brief introduction into clock distribution networks

3.2.3.1 A brief introduction into clock distribution networks [57]

Along with the increasing number of transistors on a chip and raising clock frequencies clock distribution became a more and more intricate issue and thus a field of intensive research already in the beginning of the 1990's.

Without going into details next we give an overview of the main steps of the evolution of clock distribution networks.

Clock distribution

Tree-based clock distribution

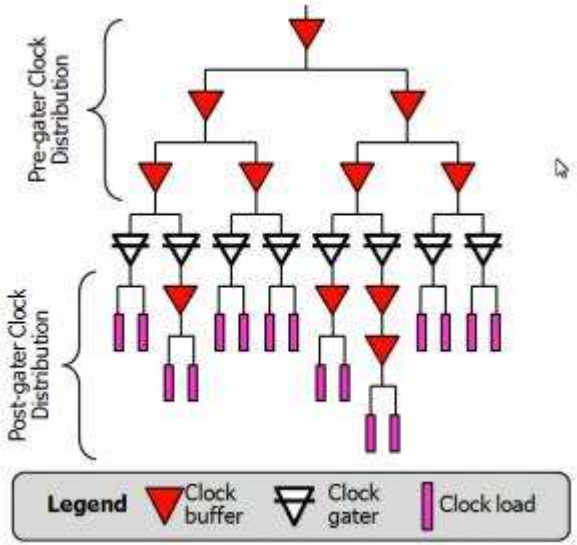

Grid-based clock distribution

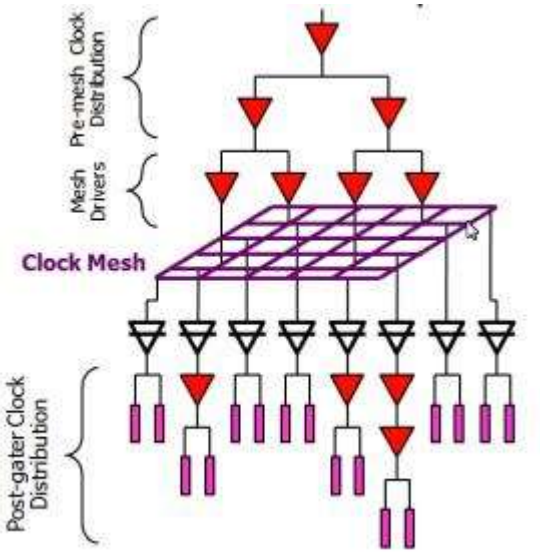

(The grid is actually a low-resistance metal grid fed by a clock driver tree)

Main types of tree-based clock distribution networks [58]

Tree-based clock distribution

Binary-tree based clock distribution

H-tree based clock distribution

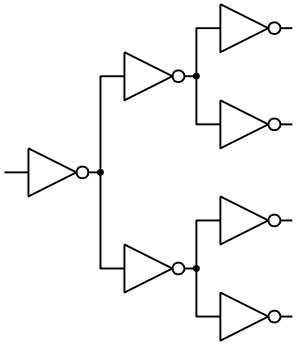

Binary tree

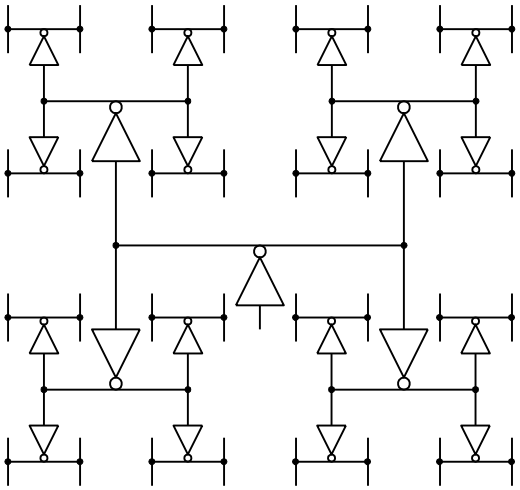

H-tree

Main types of grid-based clock distribution networks

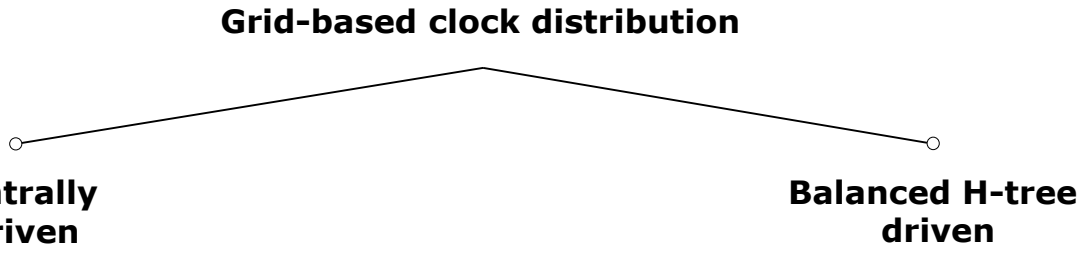

Early grid-based clock distribution networks were centrally driven [59].

Most modern grid-based distribution networks use balanced H-trees to drive the grid [59].

Example: Experimental grid-based clock distribution network with H-tree grid driving

The Figure below is an illustration for an experimental grid-based clock distribution network with H-tree driving [59].

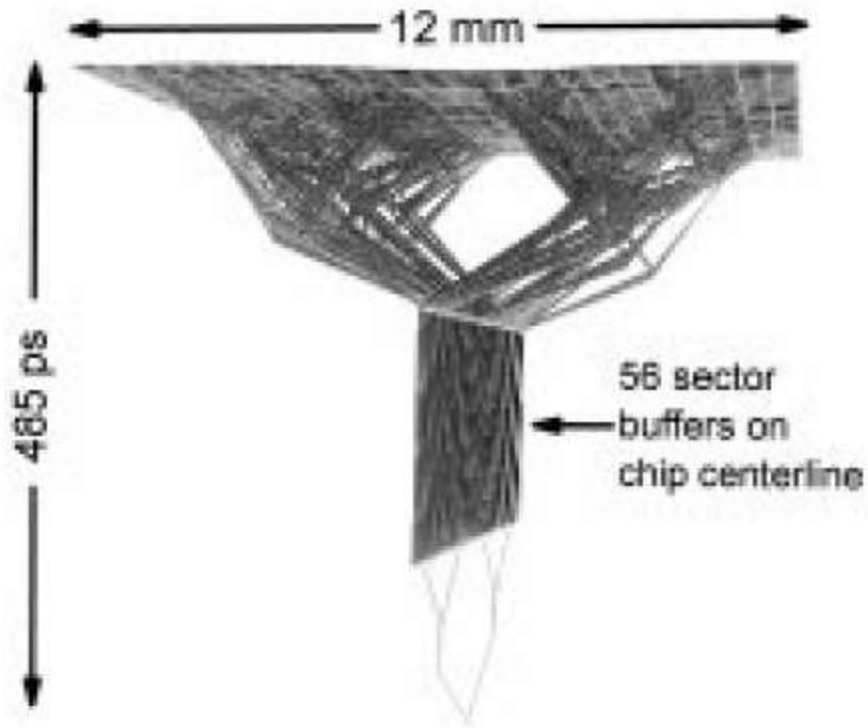

$X-Y$-time rendering of chip C with 56 sector buffers in chip centerline.

Drawback of the grid-based clock distribution

High power consumption due to the buffers needed to drive the grid.

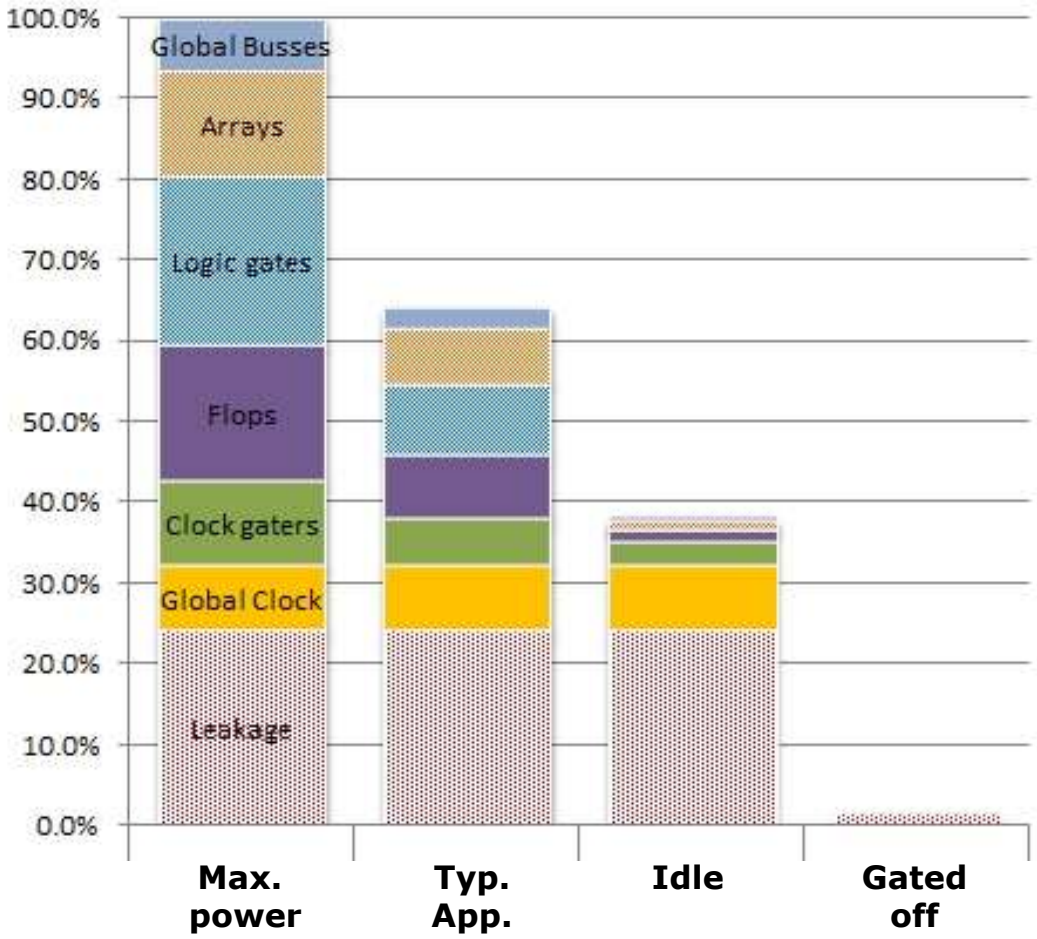

Figure: Distribution of power consumption in a Bulldozer processor [60]

3.2.3.1 A brief introduction into clock distribution networks (6)

Main steps of the evolution of clock distribution networks

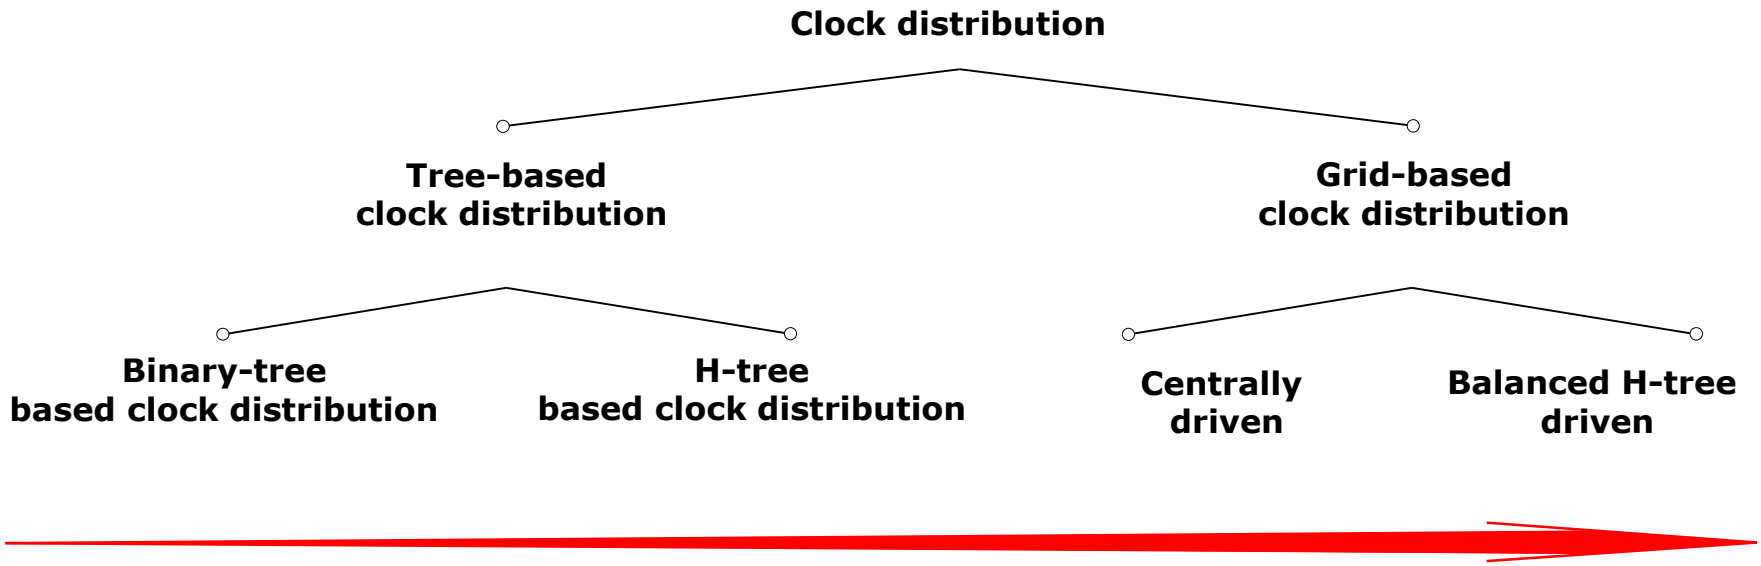

Clock gaters

Clock gating is widely used to reduce power consumption by switching off clocking of temporarily not used parts of the processor.

E.g. already Intel's Pentium 4 (2000) utilized aggressive clock gating, whereas AMD made use of this technique later, presumable beginning with their K8 family (2003).

Clock gaters implement simply an AND function to switch off or on clocking, as indicated below.

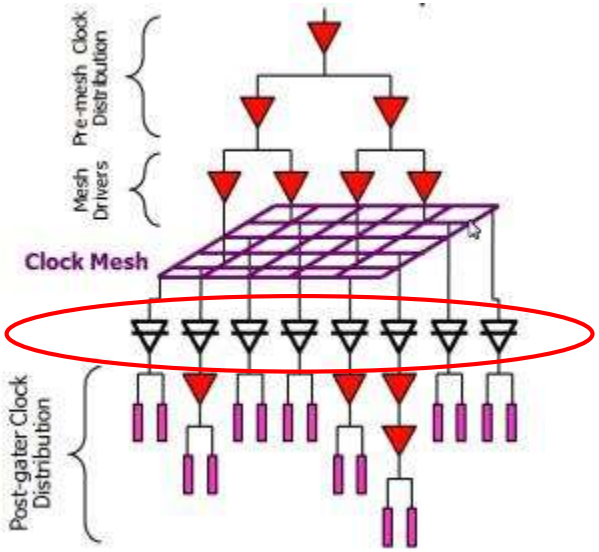

Figure: Use of clock gating to switch off temporarily not used units in a grid-based clock distribution network [57]

3.2.3.1 A brief introduction into clock distribution networks (8)

Resonant clock meshing

It aims at reducing the high power consumption of the grid-based clock distribution network.

Reduced overall power consumption can be utilized to boost clock frequency within a given TDP limit.

3.2.3.2 Principle of the Resonant Clock Mesh (RCM) technology

3.2.3.2 Principle of the Resonant Clock Mesh (RCM) technology

Cyclos provides a very good brief explanation of the RCM technology that we cite subsequently [57].

3.2.3.2 Principle of the Resonant Clock Mesh (RCM) technology (2)

The Power Challenge [57]

A modern SOC can consume up to 30% of its power just on the clock buffers, which really is a big contributor to overall power. Other EDA vendors are focused on reducing power for the areas marked below with red arrows:

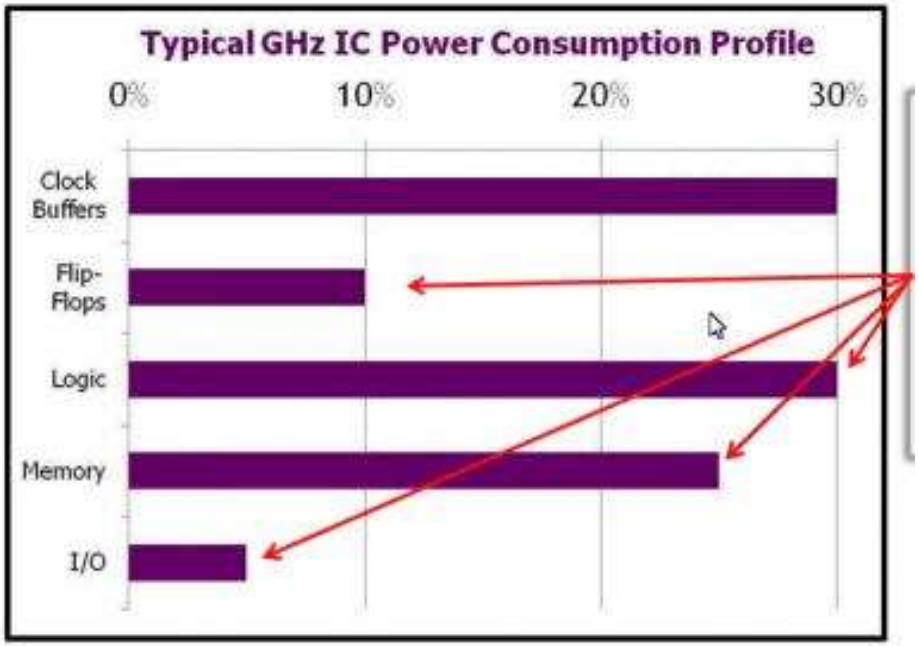

3.2.3.2 Principle of the Resonant Clock Mesh (RCM) technology (3)

The promise of Cyclos technology is to reduce the power consumption on the clock buffers:

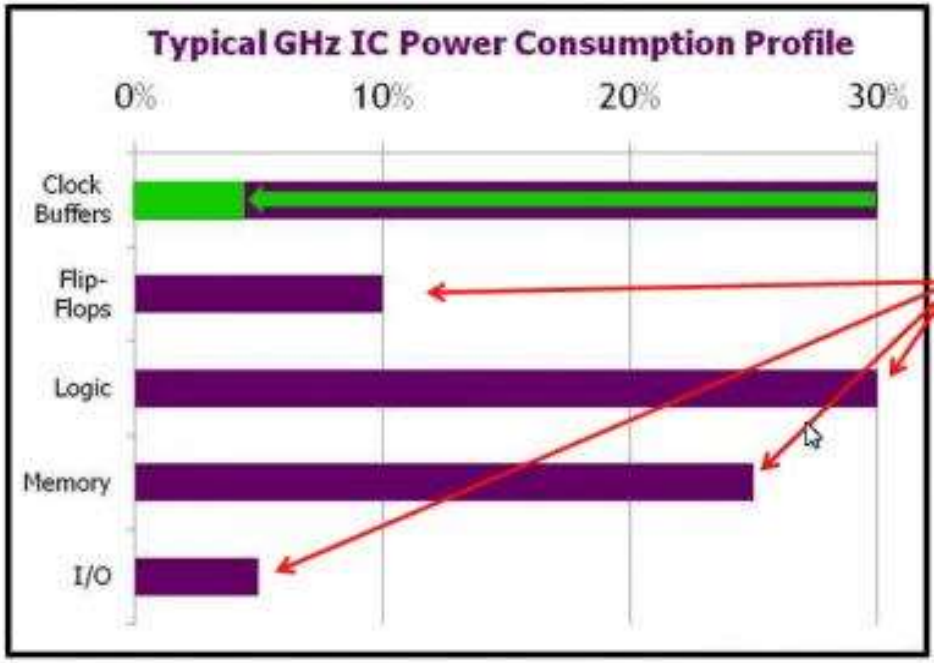

Clock gater: On/Off switch for the clock

3.2.3.2 Principle of the Resonant Clock Mesh (RCM) technology (4)

Many chips today use the familiar clock tree approach to distribute a clock signal across the chip.

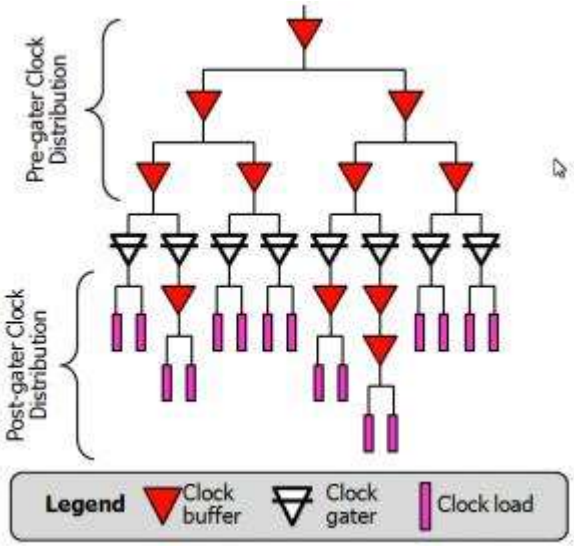

Another clock distribution approach is called the Mesh where a metal layer ties all the distributed clock signals together to form a low resistance Mesh after the initial clock driver cells:

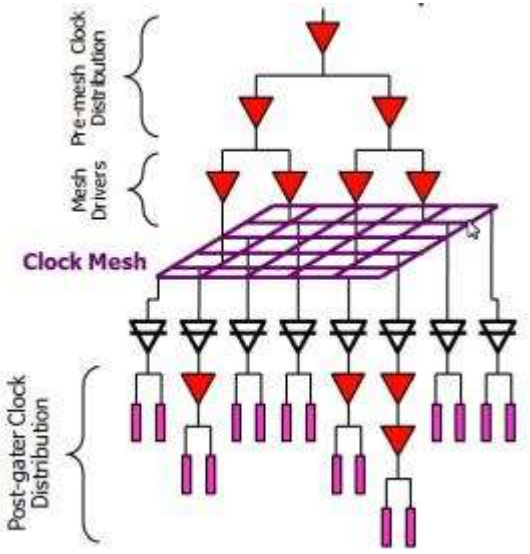

3.2.3.2 Principle of the Resonant Clock Mesh (RCM) technology (5)

The clock mesh gives you a very low skew value however it's capacitance requires increase energy to drive which also increases power consumption. We like the low skew but we don't like increasing the power.

In EE theory classes we all learned about oscilators built out of LC circuits:

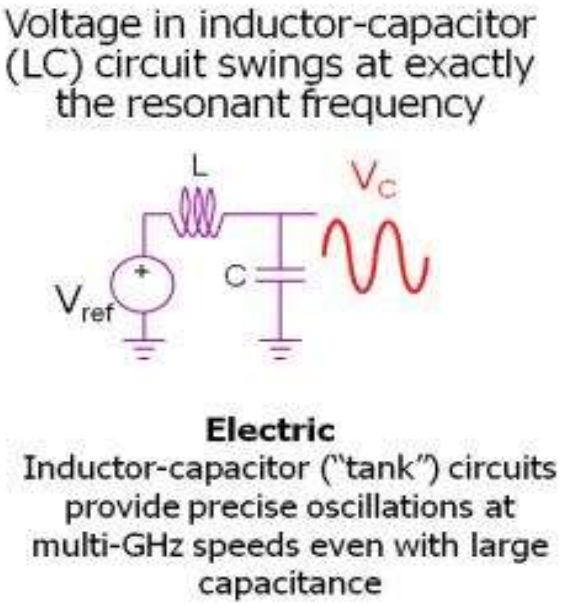

What if we could combine the benefits of the clock mesh topology with the resonance of an oscillator to reduce the energy required to drive a clock network?

3.2.3.2 Principle of the Resonant Clock Mesh (RCM) technology (6)

Hmm, that idea could work in theory:

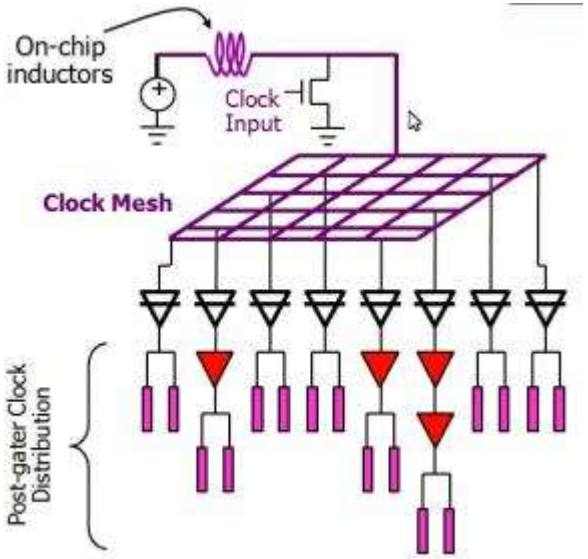

Benefits of such an approach:

- Low clock skews because of the low-resistance mesh
- Metal mesh less impacted by On Chip Variation (OCV) and Process/Voltage/Temperature (PVT) variations
- The Post-gater trees timing are isolated, so ECOs are easier in the design cycle
- Lower power consumed by the clock distribution network

Challenges of this approach:

- EDA tools not commercially developed yet
- Design flow not well understood or built

3.2.3.2 Principle of the Resonant Clock Mesh (RCM) technology (7)

The LC circuit created by the inductors basically helps to recycle clock power thus lowering consumption:

Traditional Clocks

Charge is dissipated one-way from power supply to ground every clock cycle

Resonant Clocks

Like an electric pendulum, power oscillates on-chip between the clock mesh capacitance and Cyclos inductors

A very low-power replenisher is used to "nudge" the resonant clock to restore the energy lost due to metal resistance and to ensure it oscillates at precisely the frequency of the on-chip reference clock

3.2.3.2 Principle of the Resonant Clock Mesh (RCM) technology (8)

OK, lower power consumption for my clock network is always a good thing but are there more benefits? Yes, you even have reduced jitter on your clock edges:

Large Clock Tree Skews Require Significant Guard Banding

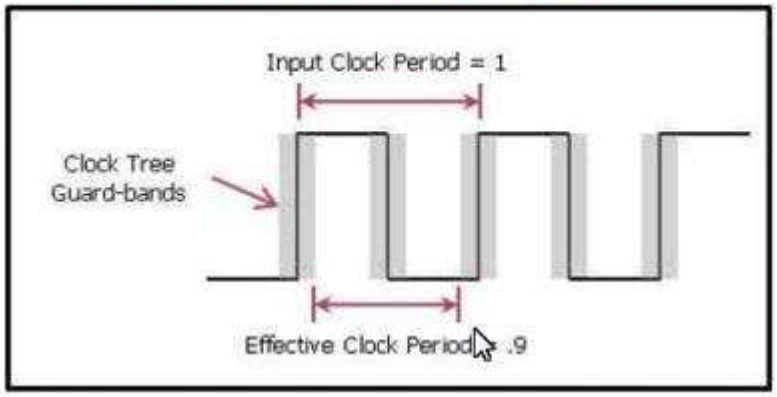

Low Clock Mesh Skews Virtually Eliminate Guard Banding

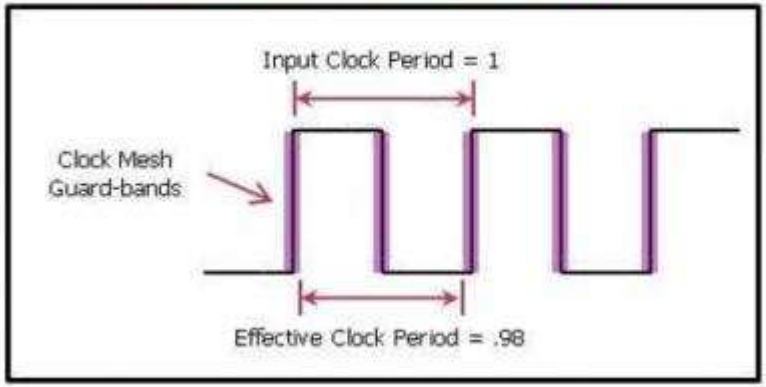

3.2.3.2 Principle of the Resonant Clock Mesh (RCM) technology (9)

The theory of using a resonant mesh for clocks is appealing, but who is really using this in production chip?

I did a quick Google search and found hundreds of articles and patents on the subject, so it looks like the leap from theory to practice has been bridged. A few more side benefits of resonant mesh clock designs are lower RF noise than clock trees, and electromigration reduction from bidirectional current flow in the clock net.

Silicon Confirmation

The Cyclos Semi approach has been used with an ARM926 chip where first silicon showed a 25% to 35% reduction in total power. Several other chips have used this approach with early customers and one DSP chip showed a 75% lower clock power number while using GHz speeds.

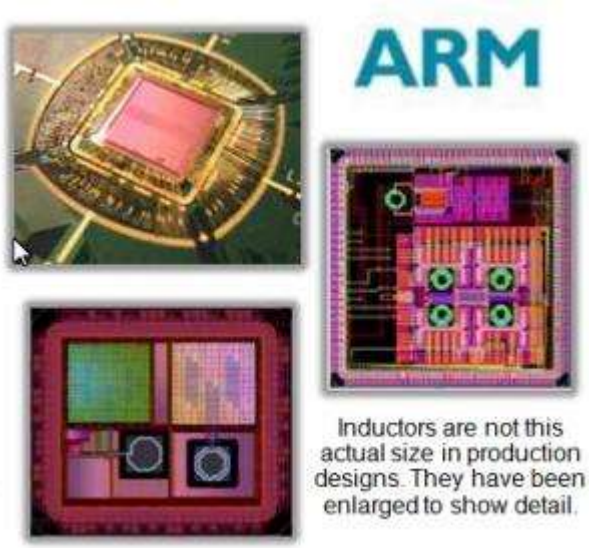

3.2.3.2 Principle of the Resonant Clock Mesh (RCM) technology (10)

Implementation

The theory matches the silicon results, so how do we get inductors onto an IC design? Here's a mesh with distributed inductors built in a top-level of metal using standard processing steps. You don't want circuits underneath these on-chip inductors so that will increase your silicon area up to 5% typically:

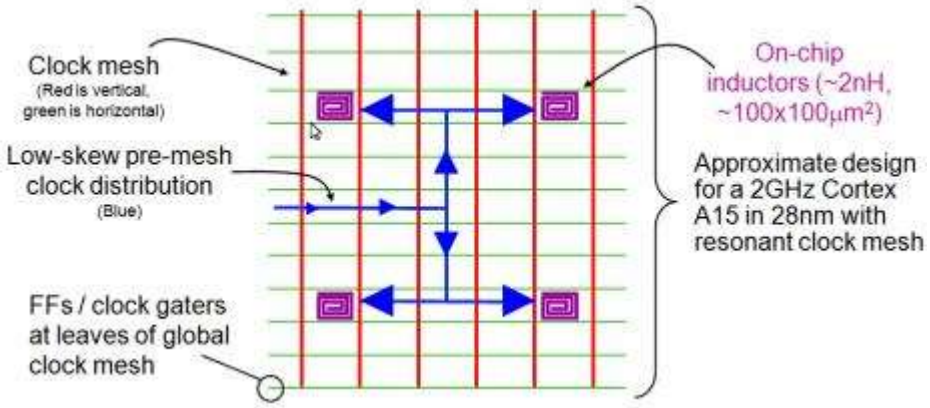

There are at least three clock distribution choices: clock tree, clock mesh, resonant mesh

High-Performance Clock Design Example

	Levels	# Buffers	Delay	Skew	Power
Clock Tree	6	1406	2072 ps	104 ps	380 mW
Clock Mesh	3	698	1546 ps	34 ps	597 mW

Resonant Clock Mesh Power Would be < 250mW

The engineers at Cyclos are promoting the resonant mesh for GHz designs as a way to reduce power and tighten up the clock specs.

EDA Tool Flow for Design

Initially the way to get this resonant clock mesh (RCM) for your chip requires some manual work so you could hire Cyclos as a consulting company or wait until their tool flow is released in 2012. The idea is to create a compiler that automates the layout implementation parts.

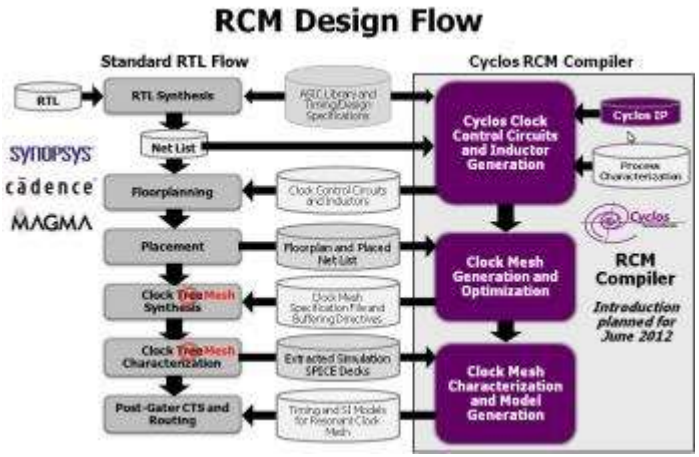

To get the RCM implementation is either $500K as a design service or wait until the RCM Compiler tool is ready around the DAC time frame. An IP license will run you $1M per process node, and finally there's a usage fee.

Compared to the digital libraries from Artisan/ARM there were only usage fees, called royalties.

3.2.3.3 The evolution of implementing RCM

3.2.3.3 The evolution of implementing RCM

- The interest to use RCM for reducing power consumption of clock distribution networks arose more or less in the beginning of the 2010's.
- A number of papers appeared and later also numerous small test chips were developed to demonstrate the potential of RCM [61].

Subsequently, we will briefly review only the large scale approaches to implement RCM in commercial processors.

3.2.3.3 The evolution of implementing RCM (2)

Experimental implementation of RCM in the ARM9EJ-S [62]

- In 2009 Cyclos, a small start up company grounded for marketing the IP of RCM announced that they completed the experimental implementation of RCM in the ARM9EJ-S.

  This was the first experimental implementation of RCM covering the entire clock distribution network of a commercial processor, nevertheless only to demonstrate the feasibility of Cyclos' IP.
- The implementation made use of an off-chip inductor to resonate the parasitic clock capacitance from the root of the clock network.
- The total power savings achieved amounted to 20 % to 35 % depending on the workload.

3.2.3.3 The evolution of implementing RCM (3)

Experimental implementation of RCM in the Cell/B.E. (2009) [61], [62].

- The work aiming at the implementation of the RCM technology into the Cell/B.E. design started when the Cell/B.E. was already in volume production.
- The implementation was limited to the global clock distribution network without including the driving flip-flops resulting in moderate total power savings of about 5%.
- RCM was implemented by using 830 on-chip spiral inductors.

3.2.3.3 The evolution of implementing RCM (4)

Implementation of RCM in AMD's Piledriver-based processor lines (2012) [63]

- It is the first volume production enabled microprocessor that makes use of the RCM technology.
- The clock system operates in two modes: direct-drive (without RCM) and resonant (RCM) mode.
  In direct-drive mode the inductors are shunted by a switch.
- In the chosen implementation a set of five horizontal folded clock trees drive a global clock grid, where each clock tree has up to 25 on-chip inductors.
- Achieved power savings in the clock distribution network is up to 24 %.
- Power savings can be utilized to boost clock speed.
- The implementation of RCM in AMD's Piledriver-based processor lines is restricted however, to be a Rev.2 of an existing clock mesh [60].

  This is in concert with some industry-observers stating that in their first shipped Piledriver-based processors AMD did not implement yet RCM [56].
- Nevertheless, according to a paper the Piledriver-based Trinity A10-4600M processor already includes RCM [64].

  According to this publication the design uses 92 100 µm-wide inductors, spread out over each dual-core processor module.

3.2.3.3 The evolution of implementing RCM (5)

Main features of AMD's Bulldozer- and Piledriver based Opteron server lines [65]

Model	Architecture	CPU Frequency		TDP	Cores	Cache		Unlocked	Stepping	Bus Speed	Launch
		Base	Boost			L2	L3				
FX-8350	Piledriver	4.0GHz	4.20GHz	125W	8	8MB	8MB	Yes	n/a	5200MT/s	Q4 2012
FX-8320		3.50GHz	4.0GHz								
FX-8300		3.30GHz	4.20GHz								
FX-8170	Bulldozer	TBD	TBD						B2		Q1 2012
FX-8150		3.60GHz	4.20GHz								Q4 2011
FX-8120		3.10GHz	4.0GHz								
FX-8100		2.80GHz	3.70GHz								
FX-6350	Piledriver	3.90GHz	4.20GHz	95W	6	6MB			n/a		Q1 2013
FX-6300		3.50GHz	4.10GHz								Q4 2012
FX-6120	Bulldozer	TBD	TBD						B2		Q1 2012
FX-6100		3.30GHz	3.90GHz								Q4 2011
FX-4350		4.20GHz	4.30GHz		4	4MB			n/a		Q1 2013
FX-4320	Piledriver	4.0GHz	4.20GHz								Q4 2012
FX-4300		3.80GHz	4.0GHz								
FX-4170	Bulldozer	4.20GHz	4.30GHz	125W					B2		Q2 2012
FX-4120		TBD	TBD	95W							Q1 2012
FX-4100		3.60GHz	3.80GHz								Q4 2011

3.2.3.3 The evolution of implementing RCM (6)

Remark

In 3/2013 AMD quietly introduced the Piledriver-based 6-core Opteron 635O and 4-core 4350 processors with remarkable increased clock speeds, as shown in the Table before [65].

It can be suspected therefore that these processors are already Rev.2 parts using RCM.

3.2.3.3 The evolution of implementing RCM (7)

Plans to implement Cyclos's RCM in ARM Cortex-A15 [66]

- In 2/2013 Global Foundries and ARM have announced plans to implement Cyclos' RCM in ARM's Cortex-A15 in order to boost clock speed.
- The RCM design will include on-chip inductors.

3.2.3.3 The evolution of implementing RCM (8)

Remark

There is no sign that AMD continued implementing Resonant Clocking in their subsequent microarchitectures (e.g. in the Steamroller or Excavator).

3.3 Piledriver-based GPU-less processor lines

- 3.3.1 Overview of the Piledriver-based GPU-less processor lines
- 3.3.2 The Abu Dhabi Opteron 6300 server line
- 3.3.3 The Vishera high performance FX desktop line

3.3.1 Overview of the Piledriver-based GPU-less processor lines

3.3.1 Overview of the Piledriver-based GPU-less processor lines-1

Piledriver-based processor lines

Piledriver-based GPU-less processor lines

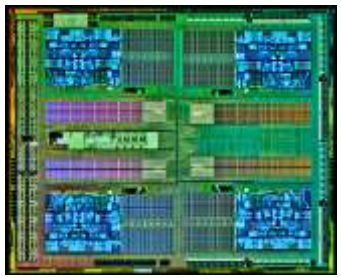

32 nm
315 mm^2
1.2 billion trans.

Abu Dhabi
Opteron server line
(Dual chips)

4P: Opteron 6300

G34 socket
(retained)

11/2012

Vishera
high performance
FX desktop line

FX-x3xx

AM3+ socket
(retained)

Volan platform

10/2012

Piledriver-based Trinity APU lines

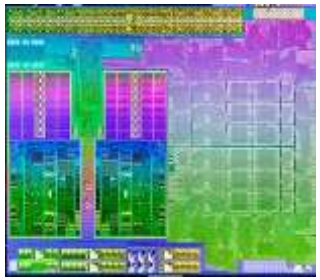

32 nm
226 mm^2
1.3 billion trans.

Trinity
mainstream
DT APU line
HD 7000D GPU

A10 - A4 5x00(K)

FM2 socket
(new)

Virgo platform

9/2012

Trinity
mainstr./ultra-thin
mobile APU line
HD 7000G GPU

A10 – A6 4xxxM

FS1r2/FP2 socket
(new)

Comal platform

5/2012

Piledriver v2-based Richland APU line

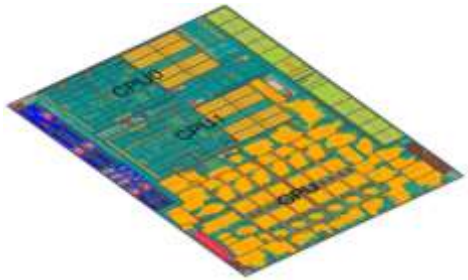

Richland
mainstream
DT APU line
HD 8000D GPU

A10 – A4 6x00(K)

FM2 socket
(retained)

Elite A-Series

6/2013

Richland
mainstr./ultra-thin
mobile APU lines
HD 8000M GPU

A10 – A4 5xxxM

FS1r2 socket
(retained)

Elite performance platform

5/2013

Overview of the Piledrivier based GPU-less processor lines-2

It underlies AMD's Abu Dhabi Opteron 6300 server line and the Vishera FX high performance desktop line.

The first Piledriver-based GPU-less processor line was introduced in 10/2012 as the Vishera high performance FX desktop line.

Key features of the Piledriver –based GPU-less processor die:

- 32 nm feature size,
- 315 mm^2,
- 1.2 billion transistors.

(These are exactly the same figures as those for the related Bulldozer-based Orochi die).

Comparing the Bulldozer-based and Piledriver-based 4-module (8 cores) dies [6], [54]

Bulldozer-based 4-module Orochi die

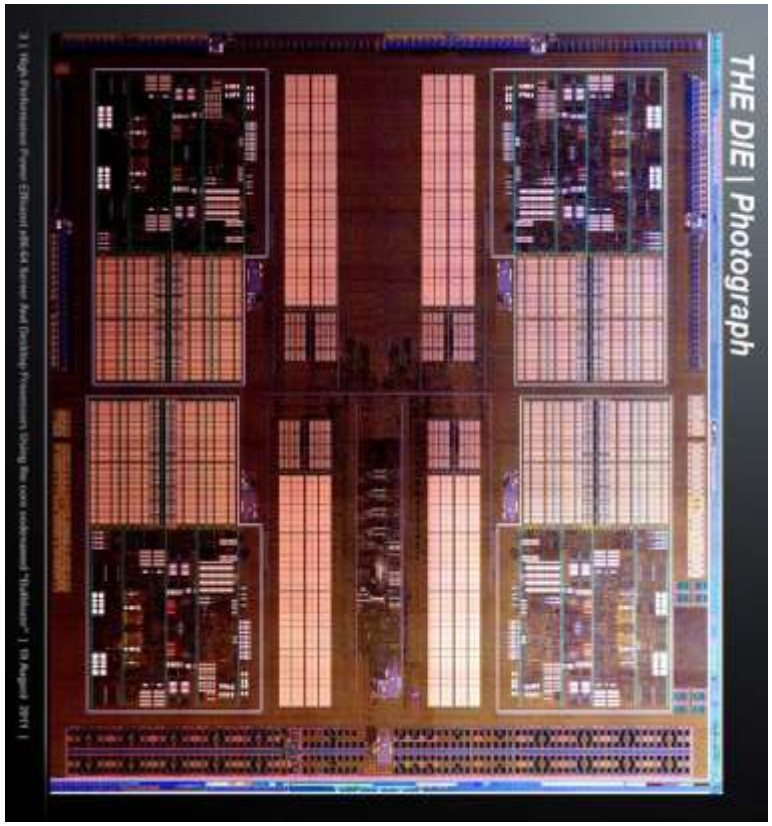

- **32 nm feature size,**
- **315 mm^2,**
- **1.2 billion transistors.**

Piledriver-based 4-modul die

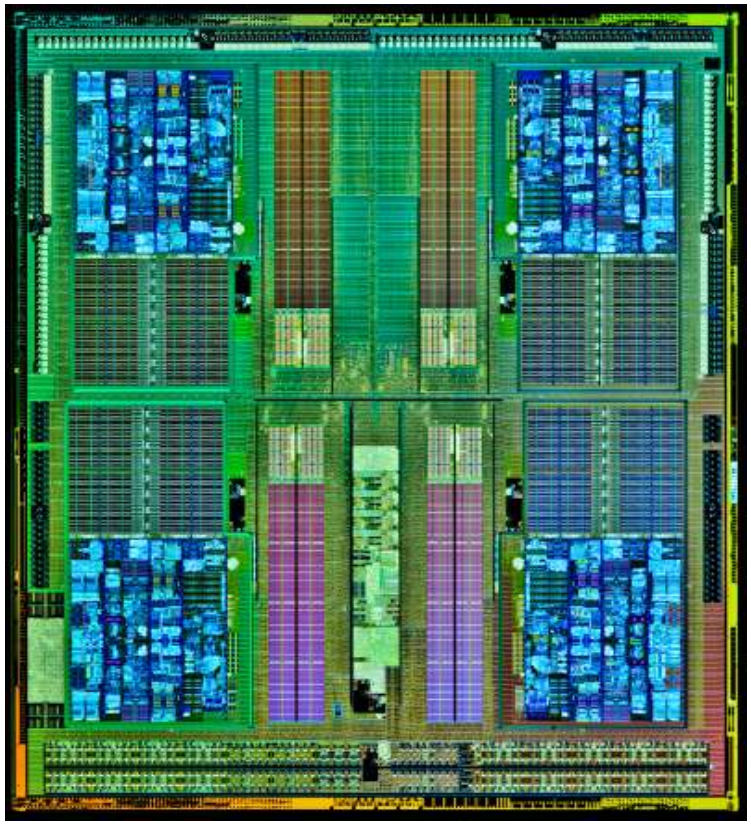

- **32 nm feature size,**
- **315 mm^2,**
- **1.2 billion transistors.**

Main functional blocks of a Piledriver-based GPU-less processor die [54]

It underlies both the Opteron 6300 Abu Dhabi server line and the Vishera high performance desktop line.

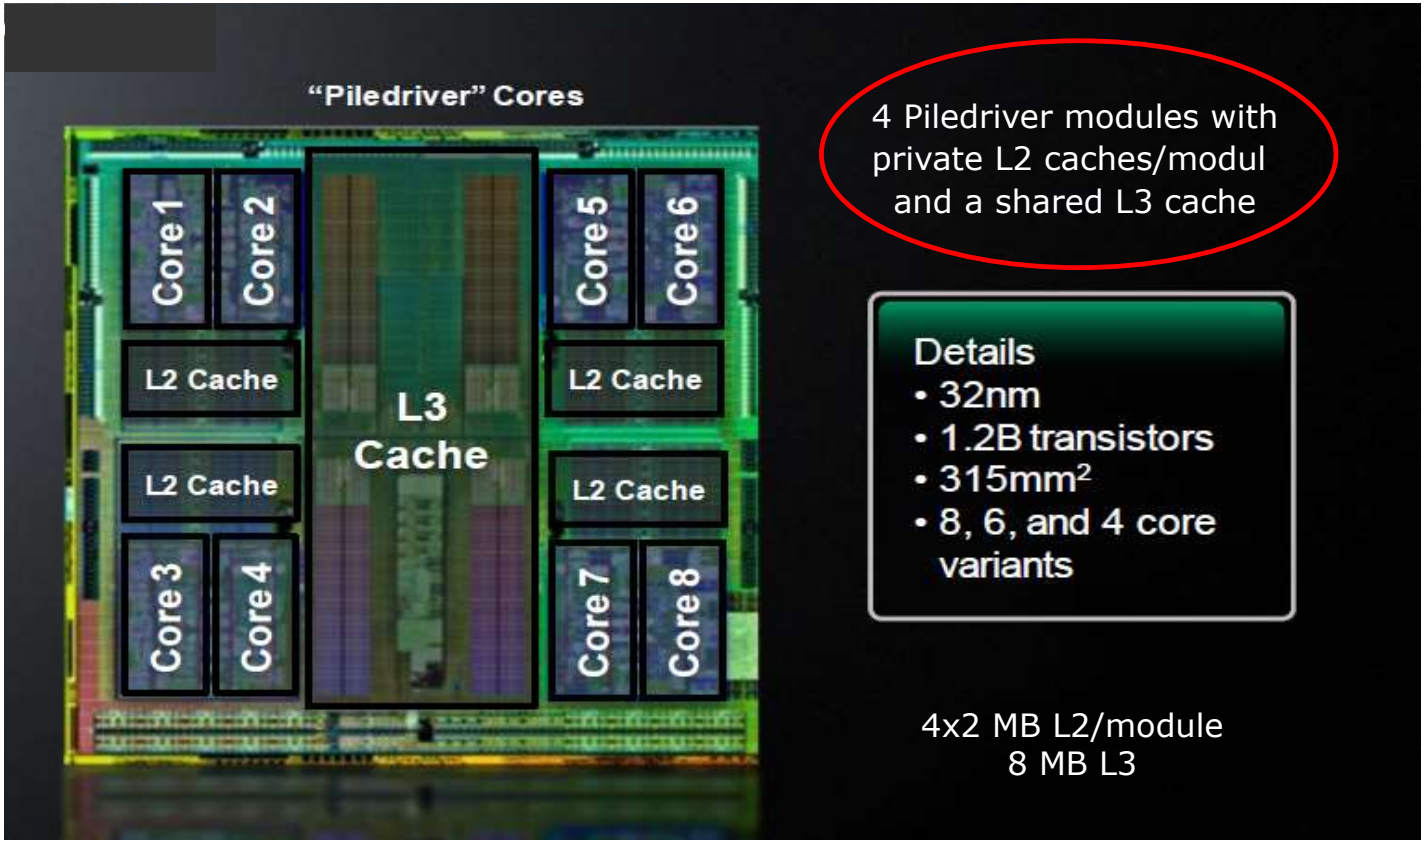

3.3.2 The Abu Dhabi Opteron 6300 server line

3.3.2 The Abu Dhabi Opteron 6300 server line

Piledriver-based processor lines

Piledriver-based GPU-less processor lines

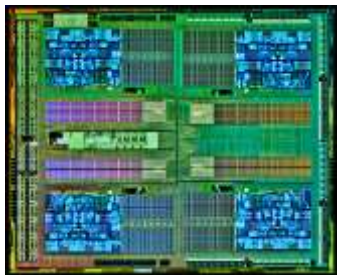

32 nm
315 mm^2
1.2 billion trans.

Piledriver-based Trinity APU lines

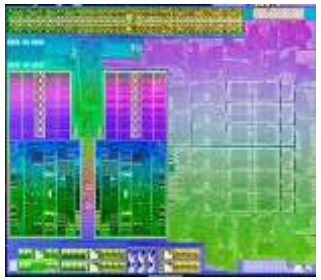

32 nm
226 mm^2
1.3 billion trans.

Piledriver v2-based Richland APU line

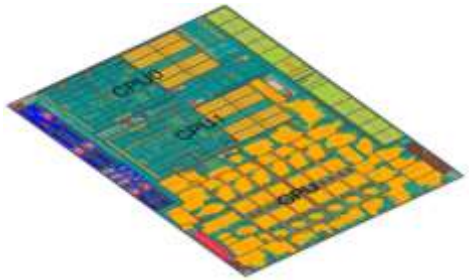

Abu Dhabi	**Vishera**	**Trinity**	**Trinity**	**Richland**	**Richland**
Opteron server line (Dual chips)	high performance FX desktop line	mainstream DT APU line HD 7000D GPU	mainstr./ultra-thin mobile APU line HD 7000G GPU	mainstream DT APU line HD 8000D GPU	mainstr./ultra-thin mobile APU lines HD 8000M GPU
4P: Opteron 6300	*FX-x3xx*	*A10 - A4 5x00(K)*	*A10 – A6 4xxxM*	*A10 – A4 6x00(K)*	*A10 – A4 5xxxM*
G34 socket (retained)	AM3+ socket (retained)	FM2 socket (new)	FS1r2/FP2 socket (new)	FM2 socket (retained)	FS1r2 socket (retained)
	Volan platform	*Virgo platform*	*Comal platform*	*Elite A-Series*	*Elite performance platform*
11/2012	10/2012	9/2012	5/2012	6/2013	5/2013

3.3.2 The Abu Dhabi Opteron 6300 server line (2)

Remark

In fact, only the Opteron 6300 4P server line is dubbed as the Abu Dhabi line, whereas the 2P and 1P models are designated differently and have also different key features, as shown below.

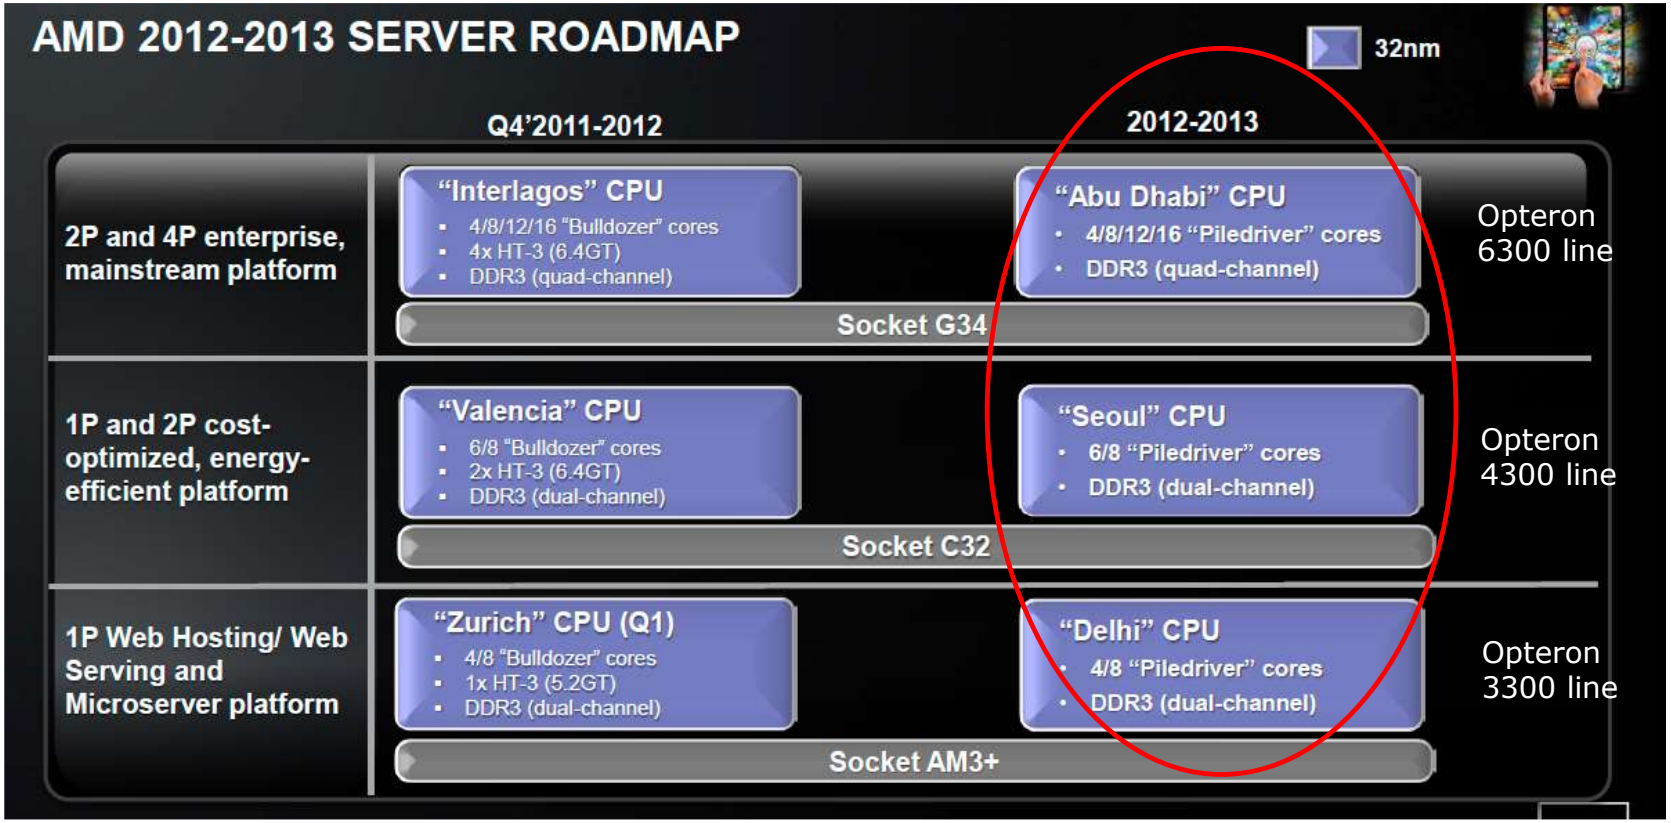

Figure: Sub-families of the Opteron 6300 (Abu Dhabi) server line [51]

Nevertheless, often the whole Opteron 6300 line (4P/2P/1P) is referred to as the Abu Dhabi line.

Main functional blocks of the dual-chip Opteron 6300 (Abu Dhabi) 4P server processor [67]

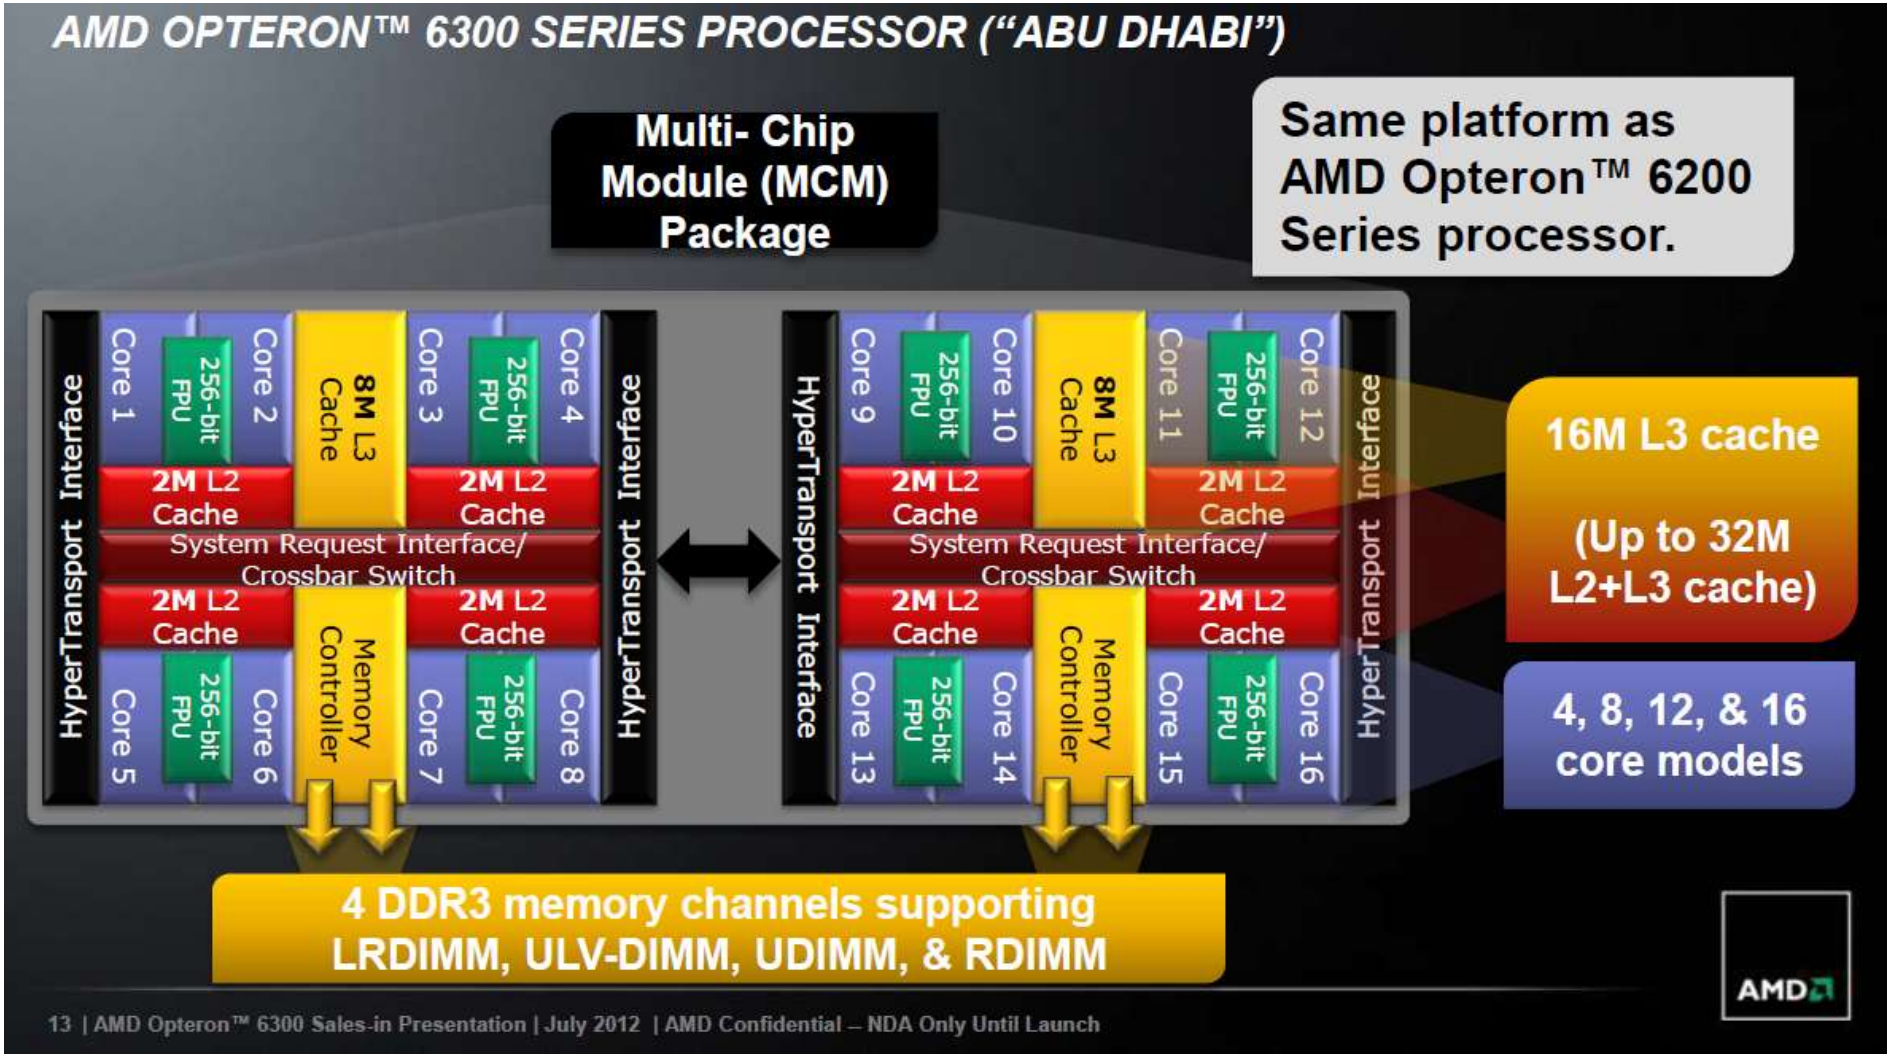

3.3.2 The Abu Dhabi Opteron 6300 server line (4)

Die plot of the dual-chip Opteron 6300 (Abu Dhabi) server processor [68]

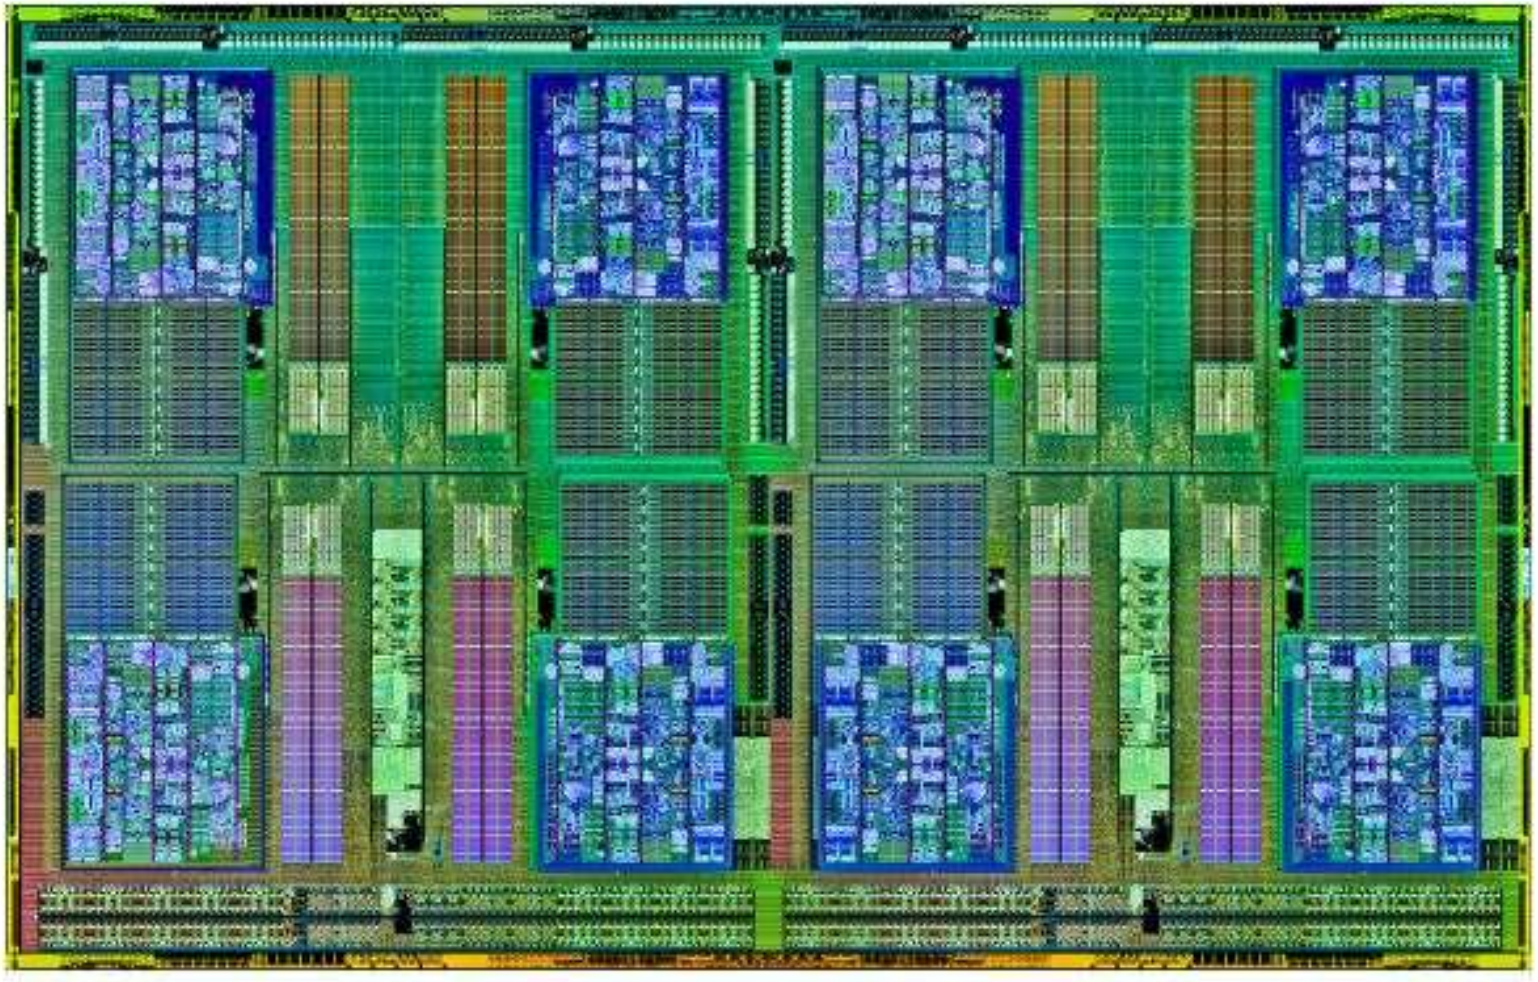

33 nm
315 mm^2

8xPiledriver modules
8x2MB L2
2x8 MB L3

3.3.2 The Abu Dhabi Opteron 6300 server line (5)

Model numbers and main features of the Opteron 6300 (Abu Dhabi) 4P line [69]

Model Number	Core Count	Core Speed	All-Core Turbo Frequency	TDP	1KU Pricing
6386 SE	16	2.8GHz	3.2GHz	140W	$1,392
6380	16	2.5GHz	2.8GHz	115W	$1,088
6378	16	2.4GHz	2.7GHz	115W	$867
6376	16	2.3GHz	2.6GHz	115W	$703
6348	12	2.8GHz	3.1GHz	115W	$575
6344	12	2.6GHz	2.9GHz	115W	$415
6328	8	3.2GHz	3.5GHz	115W	$575
6320	8	2.8GHz	3.1GHz	115W	$293
6308	4	3.5GHz	N/A	115W	$501
6366 HE	16	1.8GHz	2.3GHz	85W	$575

3.3.2 The Abu Dhabi Opteron 6300 server line (6)

Comparison of the Bulldozer-based Opteron 6200 and the Piledriver-based Opteron 6300 server lines [67]

GENERATIONAL COMPARISONS

	AMD Opteron™ 6200 Series Processors	AMD Opteron™ 6300 Series Processors
Cores	4, 8, 12 or 16 core	4, 8, 12 or 16 core
Cache (L2 per core / L3 per die)	2MB (shared between 2 cores) / 8MB	2MB (shared between 2 cores) / 8MB
Memory Channels and speed	four; up to 1600MHz	four; up to 1600MHz
Floating point capability	128-bit dedicated FMAC per core or 256-bit AVX shared between 2 cores	128-bit dedicated FMAC per core or 256-bit AVX shared between 2 cores
Integer Issues Per Cycle	4	4
Turbo CORE Technology	Yes (+500MHz with all cores active and up to 1.3GHz max turbo state)	Yes (+500MHz with all cores active and up to 1.3GHz max turbo state)**
Power (TDP)	85W, 115W, 140W	85W, 115W, 140W
New Instruction Sets		FMA3, F16c, BMI, TBM
Power Gating	AMD CoolCore™, C1E, C6	AMD CoolCore™, C1E, C6
Process / Die Size	32nm SOI	32nm SOI
Performance		Up to 15% higher processing throughput*

The above reflect current expectations regarding features and performance and is subject to change.

3.3.2 The Abu Dhabi Opteron 6300 server line (7)

Adding two new 4S sever models, called Warshaw, to the 6300 line in 01/2014

In 01/2014 AMD launched two new P-tagged models to the 6300 line, the

6370P and the 6338P models.

These models have a TDP of 99 W instead of 115 W consumed by the previous models, as seen in the next Table.

Main features of AMD's Piledriver-based 4S server lines [94]

AMD Opteron™ 6300 Series Processors						
Model Number	**Cores**	**Core Speed**	**Turbo core max. frequency**	**L3 Cache**	**TDP**	**Socket Type**
6386 SE	16	2.8GHz	3.5GHz	16MB	140W	G34
6380	16	2.5GHz	3.4GHz	16MB	115W	G34
6378	16	2.4GHz	3.3GHz	16MB	115W	G34
6376	16	2.3GHz	3.2GHz	16MB	115W	G34
6370P	16	2.0GHz	2.5GHz	16MB	99W	G34
6348	12	2.8GHz	3.4GHz	16MB	115W	G34
6344	12	2.6GHz	3.2GHz	16MB	115W	G34
6338P	12	2.3GHz	2.8GHz	16MB	99W	G34
6328	8	3.2GHz	3.8GHz	16MB	115W	G34
6320	8	2.8GHz	3.3GHz	16MB	115W	G34
6308	4	3.5GHz	N/A	16MB	115W	G34
6366 HE	16	1.8GHz	3.1GHz	16MB	85W	G34

Warsaw models

Abu Dhabi models

3.3.3 The Vishera high performance FX desktop line

3.3.3 The Vishera high performance FX desktop line

Piledriver-based processor lines

Piledriver-based GPU-less processor lines

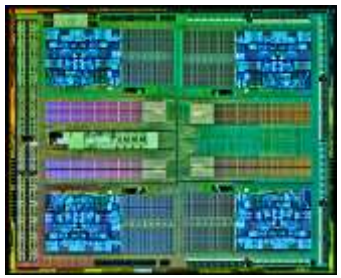

32 nm
315 mm^2
1.2 billion trans.

Abu Dhabi
Opteron server line
(Dual chips)

4P: Opteron 6300

G34 socket
(retained)

11/2012

Vishera
high performance
FX desktop line

FX-x3xx

AM3+ socket
(retained)

Volan platform

10/2012

Piledriver-based Trinity APU lines

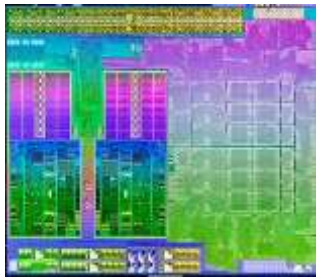

32 nm
226 mm^2
1.3 billion trans.

Trinity
mainstream
DT APU line
HD 7000D GPU

A10 - A4 5x00(K)

FM2 socket
(new)

Virgo platform

9/2012

Trinity
mainstr./ultra-thin
mobile APU line
HD 7000G GPU

A10 – A6 4xxxM

FS1r2/FP2 socket
(new)

Comal platform

5/2012

Piledriver v2-based Richland APU line

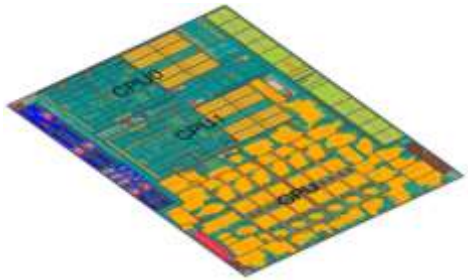

Richland
mainstream
DT APU line
HD 8000D GPU

A10 – A4 6x00(K)

FM2 socket
(retained)

Elite A-Series

6/2013

Richland
mainstr./ultra-thin
mobile APU lines
HD 8000M GPU

A10 – A4 5xxxM

FS1r2 socket
(retained)

Elite performance platform

5/2013

Main functional blocks of the high performance Vishera FX desktop line [54]

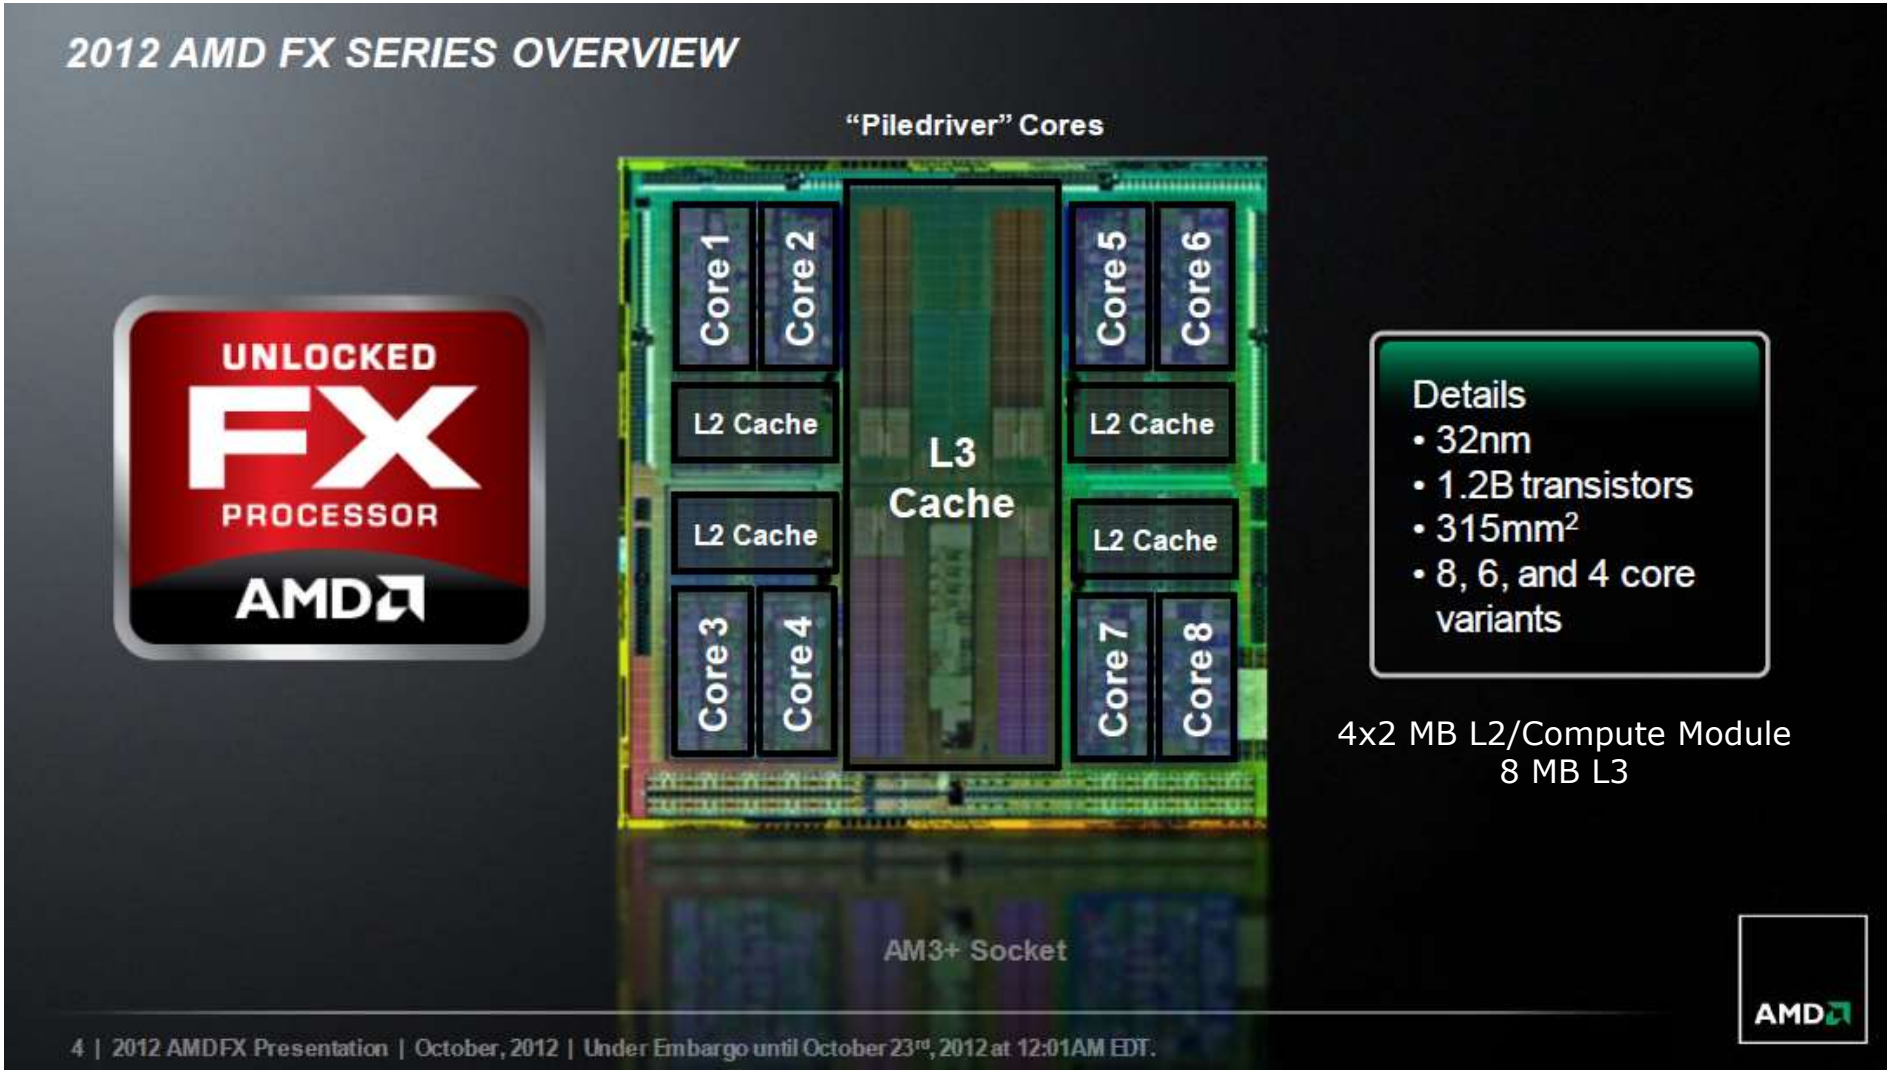

3.3.3 The Vishera high performance FX desktop line (3)

Die plot of the high performance Vishera FX desktop line [54]

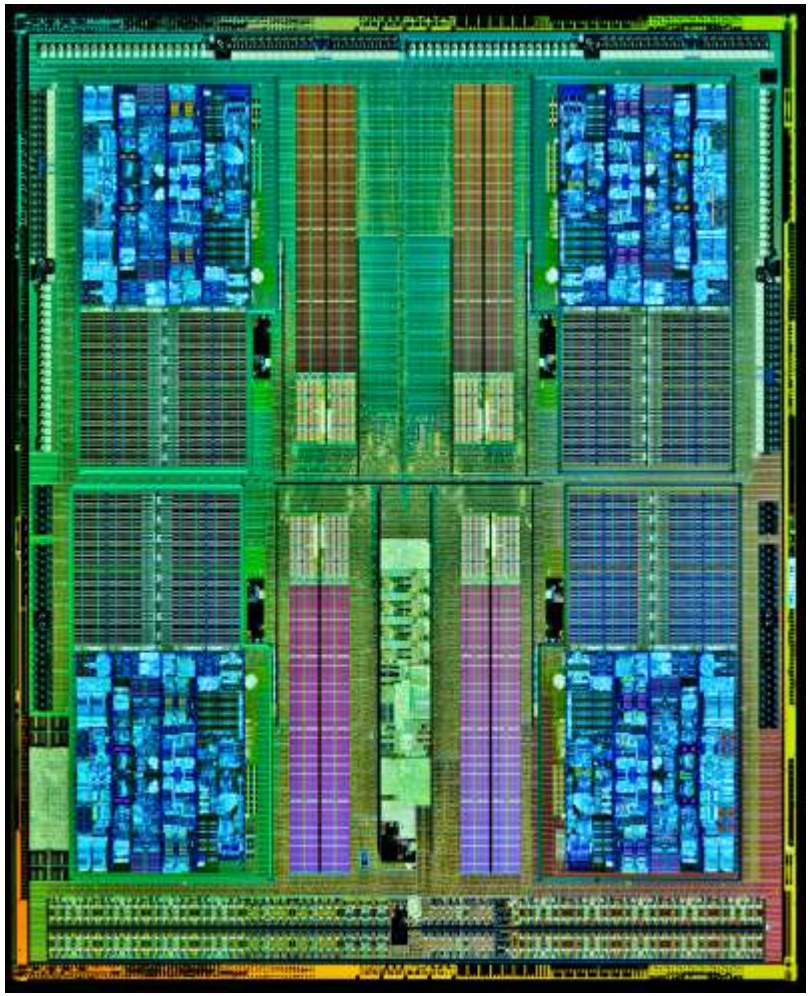

32 nm
315 mm^2
1.2 billion transistors

3.3.3 The Vishera high performance FX desktop line (4)

Model numbers and main features of the high performance Vishera FX desktop line [60]

Típus	Base clock/ /Turbo clock frequencies	L2 cache	L3 cache	TDP	Northbridge clock	List-price $
FX-8350 (8 cores)	4,0/4,1/4,2 GHz	4 x 2 MB	8 MB	125 W	2,2 GHz	195
FX-8320 (8 cores)	3,5/3,7/4,0 GHz	4 x 2 MB	8 MB	125 W	2,2 GHz	169
FX-6300 (6 cores)	3,5/3,8/4,1 GHz	3 x 2 MB	8 MB	95 W	2 GHz	132
FX-4300 (4 cores)	3,8/3,9/4,0 GHz	2 x 2 MB	4 MB	95 W	2 GHz	122

3.3.3 The Vishera high performance FX desktop line (5)

Comparing main features of AMD's Vishera and Zambezi FX desktop lines [49]

CPU Specification Comparison

Processor	Codename	Cores	Clock Speed	Max Turbo	L2/L3 Cache	TDP	Price
AMD FX-8350	Vishera	8	4.0GHz	4.2GHz	8MB/8MB	125W	$199
AMD FX-8150	Zambezi	8	3.6GHz	4.2GHz	8MB/8MB	125W	$183
AMD FX-8320	Vishera	8	3.5GHz	4.0GHz	8MB/8MB	125W	$169
AMD FX-8120	Zambezi	8	3.1GHz	4.0GHz	8MB/8MB	125W	$153
AMD FX-6300	Vishera	6	3.5GHz	4.1GHz	6MB/8MB	95W	$132
AMD FX-6100	Zambezi	6	3.3GHz	3.9GHz	6MB/8MB	95W	$112
AMD FX-4300	Vishera	4	3.8GHz	4.0GHz	4MB/4MB	95W	$122
AMD FX-4100	Zambezi	4	3.6GHz	3.8GHz	4MB/4MB	95W	$101

3.3.3 The Vishera high performance FX desktop line (6)

It can be noticed that high end models of the Piledriver-based Vishera FX line offer about 10 % higher base clock speed than the related models of the previous Bulldozer-based Zambezi line.

3.3.3 The Vishera high performance FX desktop line (7)

Main features of the 9-Series chipset supporting the high performance Vishera DT line [7

The Vishera FX line makes use of the same chipset (9-Series) as the previous Zambezi FX DT line.

AMD 9-SERIES CHIPSETS
THE IDEAL FIT FOR AMD FX-SERIES CPUS

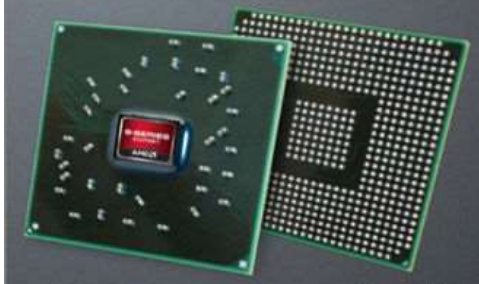

The AMD 9-Series Chipsets unlock the world's first native 8-core desktop processors from AMD with the support of the latest device technologies for an easy, seamless PC experience and the next generation AMD OverDrive™ software for full FX support.

9-SERIES CHIPSET AMD

AMD Chipset	Graphics Support	CPU Compatibility	Socket Support	Memory Support	PCI Express® 2.0	USB Support	SATA 6 GB/s	9-Series Partners
990FX	Up to 4 discrete AMD Radeon™ HD GPUs*	AMD Athlon™ AMD Athlon™ II AMD Phenom™ AMD Phenom™ II AMD FX™	AM3+, AM3	1866 MHz DDR3 memory with AMD Memory Profiles	2x16 or 4x8	Up to 14 USB 2.0 ports	Up to 6 Native	ASRock ASUS BIOSTAR ECS Foxconn GIGABYTE MSI
990X	Up to 2 discrete AMD Radeon™ HD GPUs*				1x16 or 2x8			
970	1 discrete AMD Radeon™ HD GPU				1x16			

AMD 9-Series chipset I/O features are supported through the companion *AMD SB950* chipset.

*with AMD CrossFire™ technology AMD CrossFire™ technology requires an AMD CrossFire Ready motherboard, and may require an AMD CrossFire™ Bridge Interconnect (one for each additional graphics card) and a specialized power supply.

AMD

17 | 2012 AMD FX Presentation | October, 2012 | Under Embargo until October 23rd, 2012 at 12:01AM EDT.

AMD's high-performance processor roadmap from 10/2011 [44]

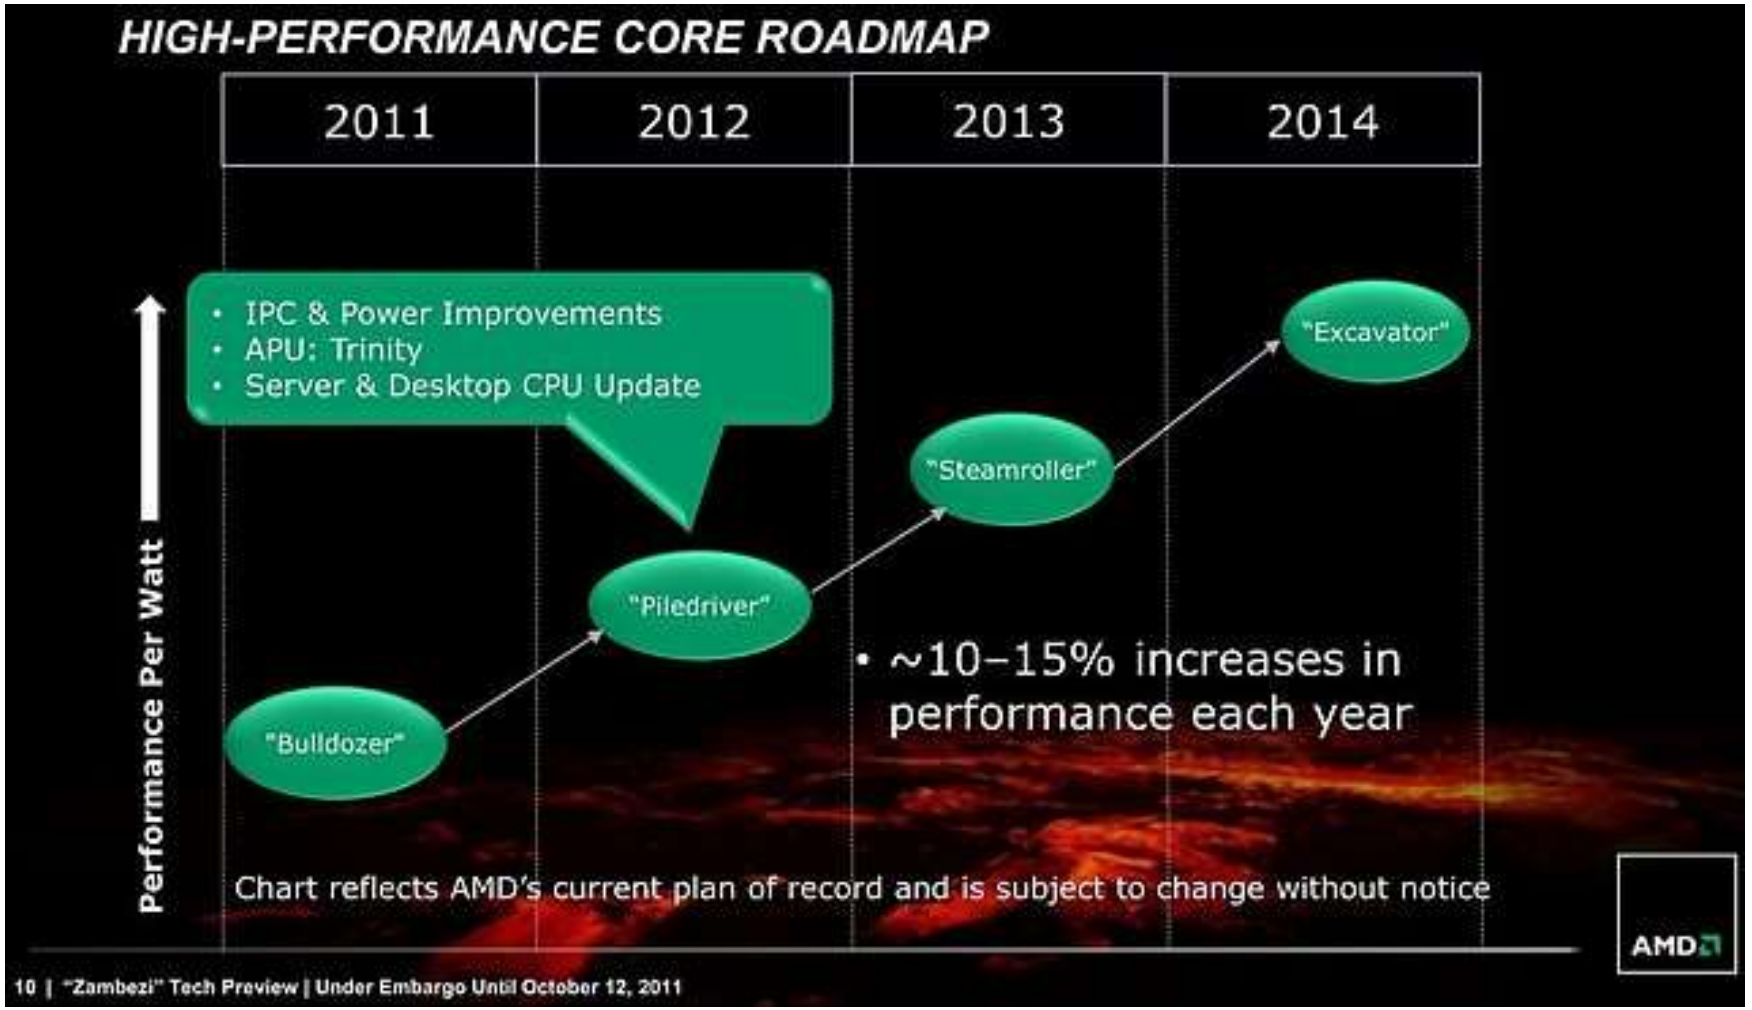

3.4 Piledriver-based Trinity APU lines

- 3.4.1 Overview of the Piledriver-based Trinity APU lines
- 3.4.2 The Trinity APU die
- 3.4.3 The Trinity mainstream desktop APU line
- 3.4.4 The Trinity mobile APU line

3.4.1 Overview of the Piledriver-based Trinity APU lines

3.4.1 Overview of the Piledriver-based Trinity APU lines

Piledriver-based processor lines

Piledriver-based GPU-less processor lines

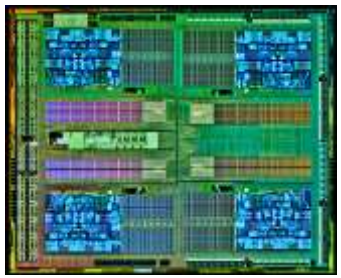

32 nm
315 mm^2
1.2 billion trans.

Abu Dhabi
Opteron server line
(Dual chips)

4P: Opteron 6300

G34 socket
(retained)

11/2012

Vishera
high performance
FX desktop line

FX-x3xx

AM3+ socket
(retained)

Volan platform

10/2012

Piledriver-based Trinity APU lines

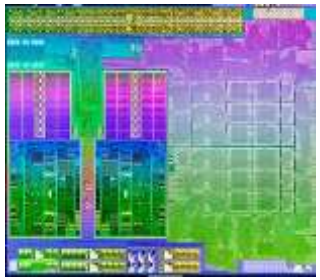

32 nm
226 mm^2
1.3 billion trans.

Trinity
mainstream
DT APU line
HD 7000D GPU

A10 - A4 5x00(K)

FM2 socket
(new)

Virgo platform

9/2012

Trinity
mainstr./ultra-thin
mobile APU line
HD 7000G GPU

A10 – A6 4xxxM

FS1r2/FP2 socket
(new)

Comal platform

5/2012

Piledriver v2-based Richland APU line

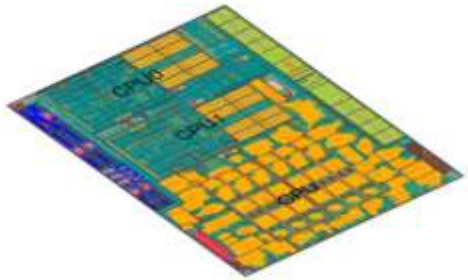

Richland
mainstream
DT APU line
HD 8000D GPU

A10 – A4 6x00(K)

FM2 socket
(retained)

Elite A-Series

6/2013

Richland
mainstr./ultra-thin
mobile APU lines
HD 8000M GPU

A10 – A4 5xxxM

FS1r2 socket
(retained)

Elite performance platform

5/2013

3.4.2 The Trinity APU die

3.4.2 The Trinity APU die

It underlies AMD's mainstream desktop and mobile Trinity APU lines.

The first Trinity APU line was introduced in 5/2012 as the Trinity mobile APU line.

Key features of the Trinity die:

- 32 nm feature size,
- 226 mm^2,
- 1.303 billion transistors.

(These are almost the same figures as those of the Llano die).

AMD's Trinity APU die [71]

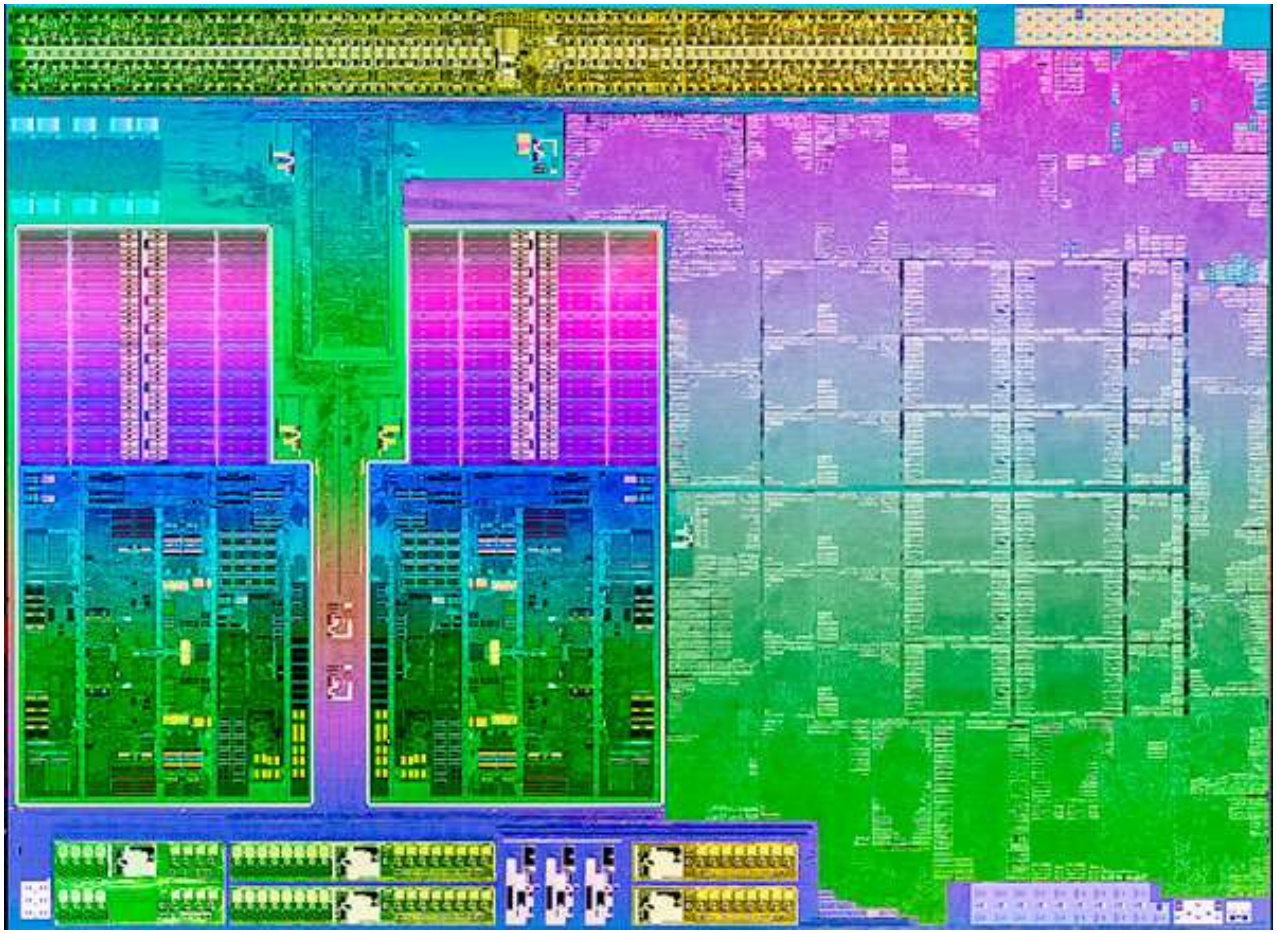

32 nm
226 mm²
1.303 billion transistors

2 Piledriver modules (4 cores)

7000 Series GPU
(Cayman/Northern Islands)
(VLIW4, up to 384 ALUs)

Comparing die plots of AMD's Llano and Trinity dies [72]

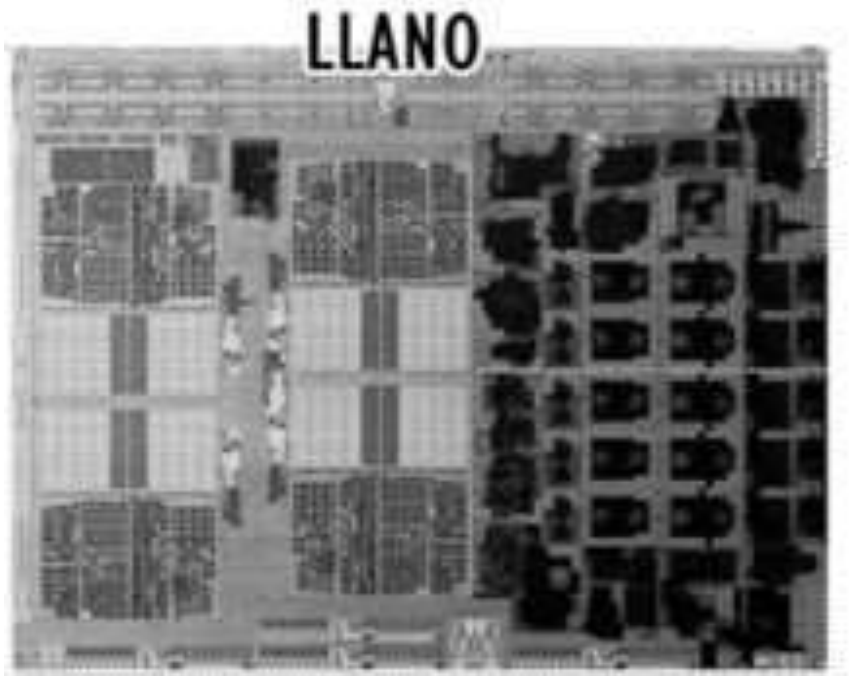

Quad cores

HD 6xxx (Cypress) GPU
(VLIW5, up to 400 ALUs)

32 nm
228 mm^2
1.450 billion transistors

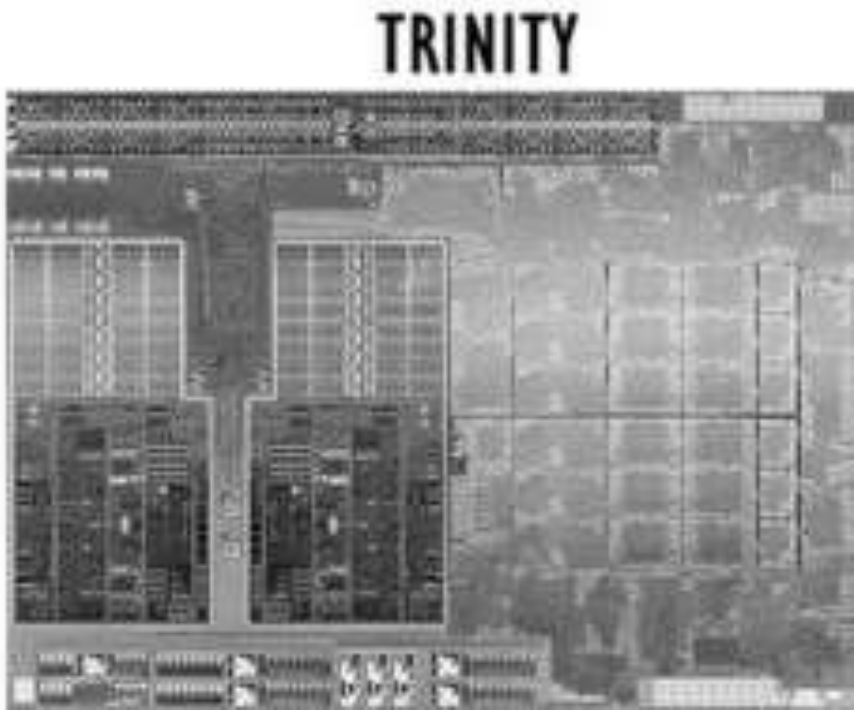

Dual modules (quad cores)

HD 7xxxD GPU (Cayman/Northern Islands)
(VLIW4, up to 384 ALUs)

32 nm
226 mm^2
1.303 billion transistors

3.4.2 The Trinity APU die (4)

Improvements of the Piledriver APU family over the Llano APU family

a) Enhancements of the microarchitecture
b) Improvement of the power management

a) Enhancements of the microarchitecture of the Trinity APU [73]

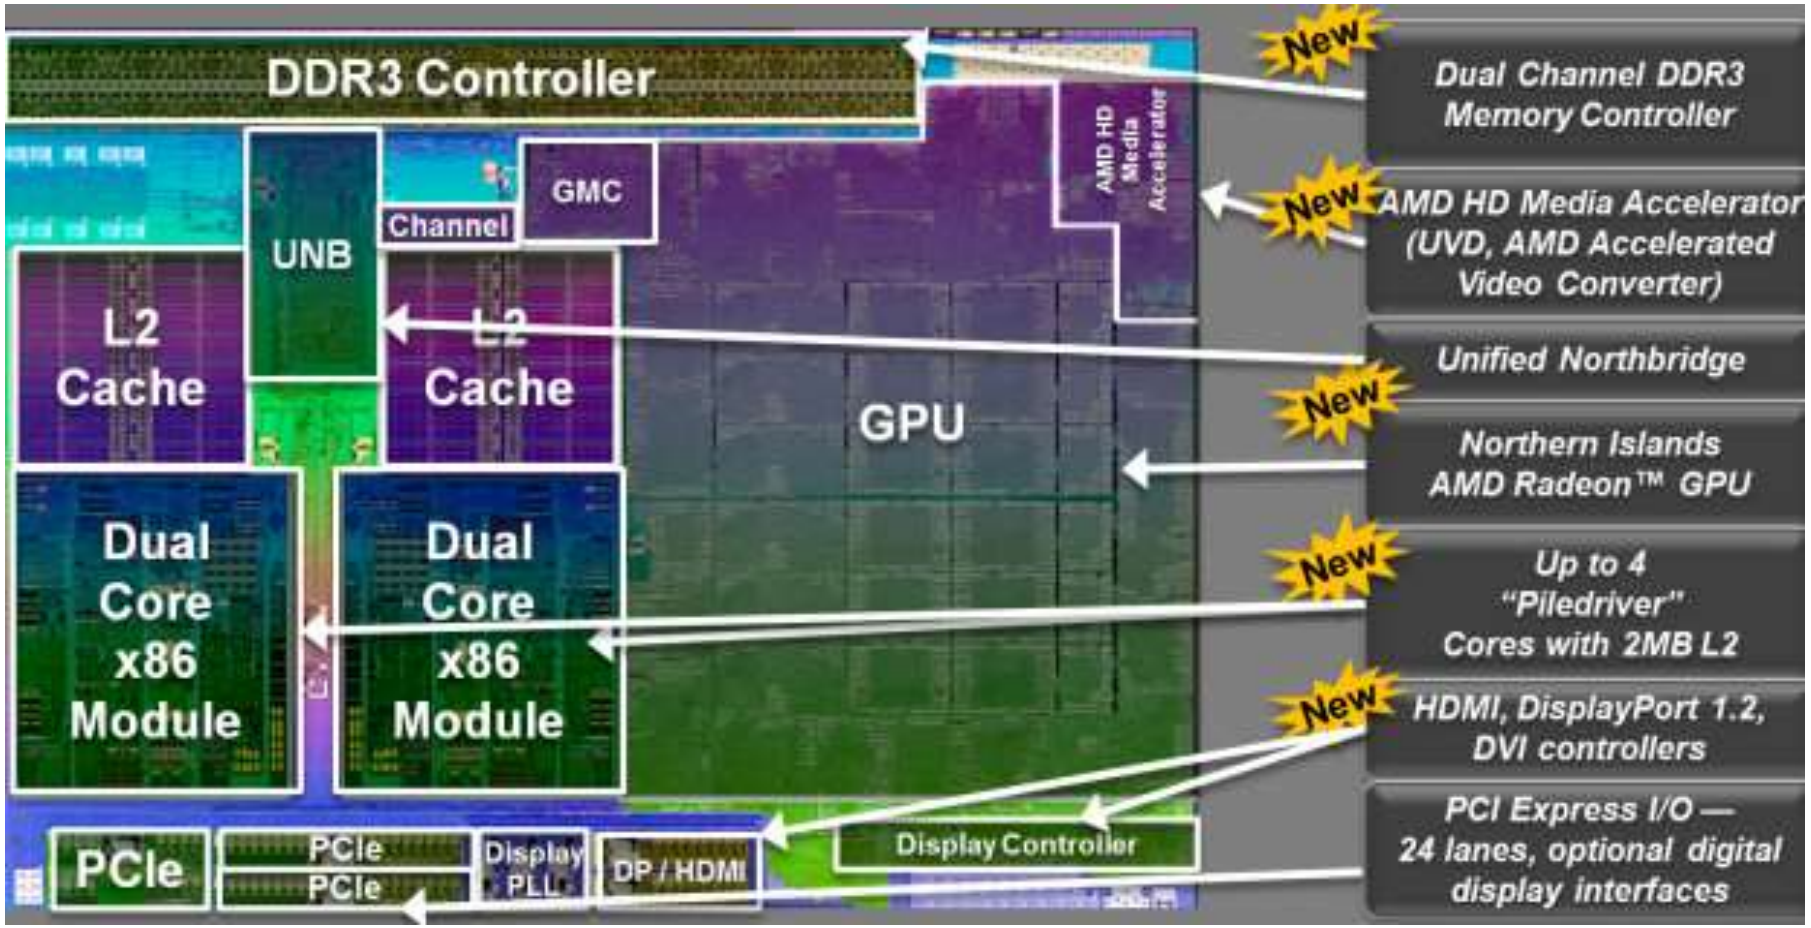

Note: No L3 cache

3.4.2 The Trinity APU die (6)

Here we do not go into details relating to the improvements of the microarchitecture of the Trinity APU but refer e.g. to the following sources [55], [73].

3.4.2 The Trinity APU die (7)

b) Improvement of the power management

In their Trinity APU family AMD introduced the Turbo Core technology 3.0, first in the the Trinity mobile line in 5/2012.

The Turbo Core Technology 3.0 is an improvement of Llano's Turbo Core Technology.

For this reason – before discussing the introduced improved technology - first let us recap the Turbo Core Technology of the Llano APU.

The Turbo Core technology of the Llano APU [74], [75]

Based on a patent filed in 2008 by Naffziger (one of the key processor architects of AMD) [74], Llano became AMD's second processor including the Turbo Core technology (the first one was the K10.5 based high performance Thuban desktop processor (2010)).

Llano digitally monitors a large number (95) of relevant events in each core, such as FX and FP operations, L1/L2 cache accesses etc. to calculate the power dissipation of each unit (4 cores plus the GPU), and also the entire chip, as indicated below.

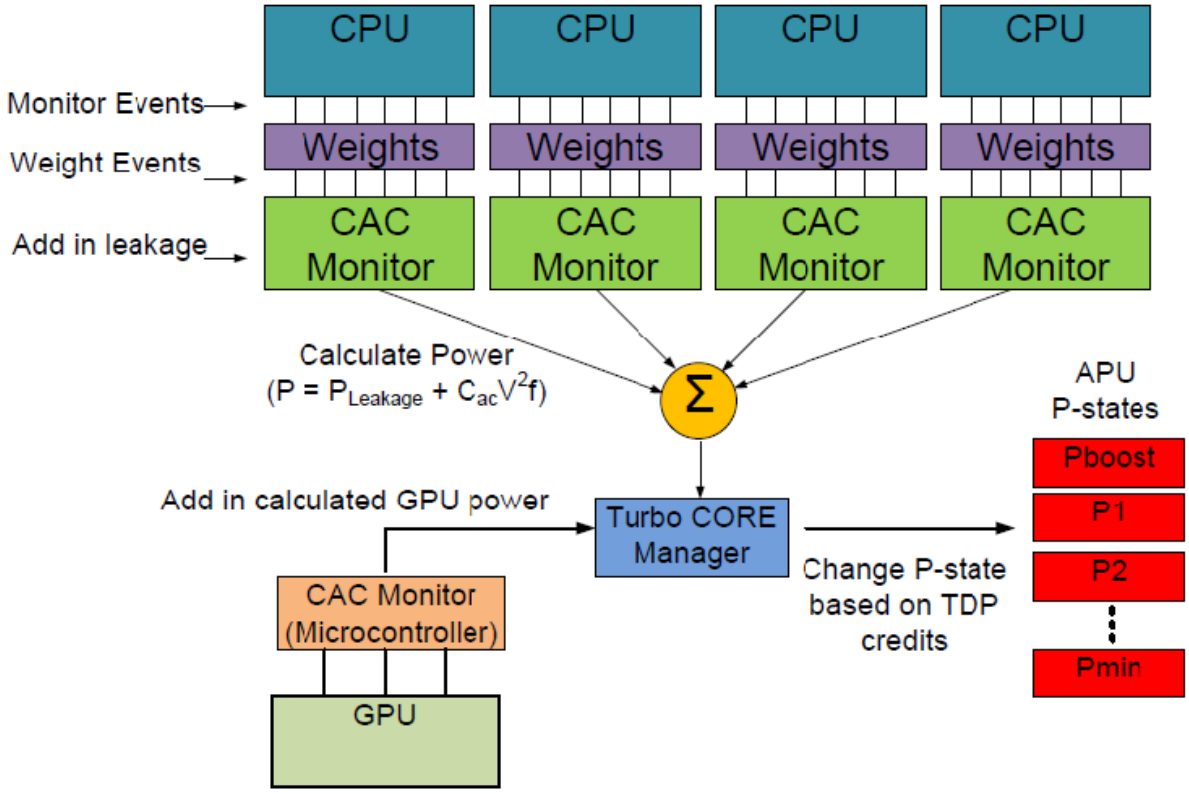

Figure: Simplified layout of the digital power monitoring system of the Llano APU [75]

3.4.2 The Trinity APU die (9)

Based on the calculated power consumption the Turbo Core Manager determines the actual energy margins as the difference between the actual power consumption of the cores and the chip and the related TDP figures.

- Positive margins indicate power headroom
- Negative margins indicate power overage

A power headroom can be utilized for increasing the clock frequency.

Power overages, on the other hand, initiate throttling (clock reduction) of the cores or even the GPU.

If the OS requests higher CPU performance for particular cores and there is a power headroom available, the Turbo Core Manager initiates a clock frequency increase for the related core.

Note that the TDP is considered as a given static value and only the clock frequencies of the cores may be increased, but the clock frequency of the GPU can not be boosted, not even in case when the CPU cores are not fully utilized or are inactive.

3.4.2 The Trinity APU die (10)

The Turbo Core Technology 3.0 of the Trinity APU line [76]

It is an enhancement of Llano's Turbo Core Technology, as indicated below.

Unlike the Llano APU die the Trinity APU die includes two compute modules (CU0 and CU1) and the GPU, rather than 4 cores and the GPU.

Accordingly, the basic layout of the digital power monitoring system has been modified, as follows:

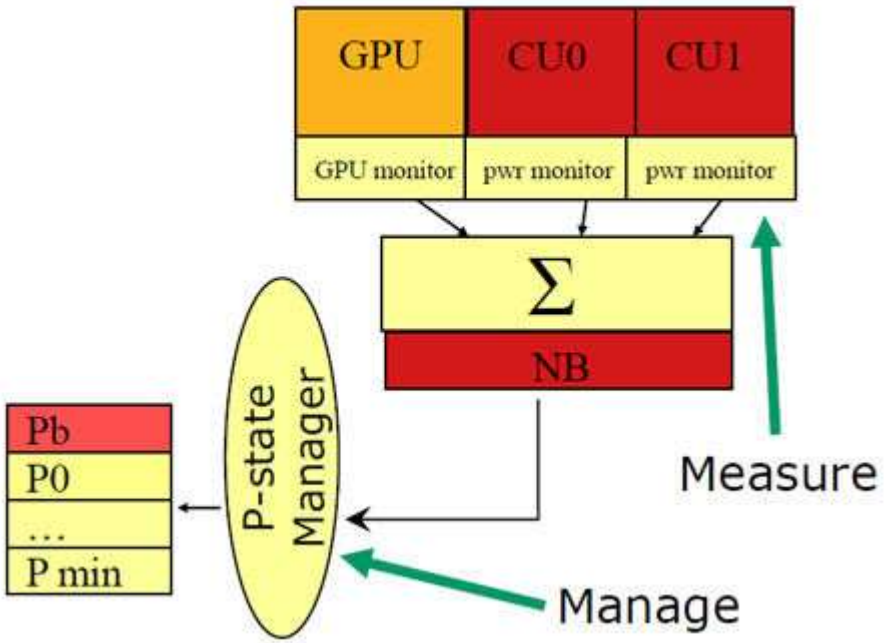

Figure: Simplified layout of the digital power monitoring system of the Trinity APU [76]

3.4.2 The Trinity APU die (11)

The major enhancement of the Turbo Core technology of the Trinity vs. the Llano APU is that the Trinity APU implements a bi-directional turbo management, unlike the Llano APU that could boost only the core frequencies.

This allows now to increase also the clock frequency of the GPU when there is a heavy GPU load and enough power headroom is available, as the following Figure demonstrates it for the Trinity A10-4600M mobile APU.

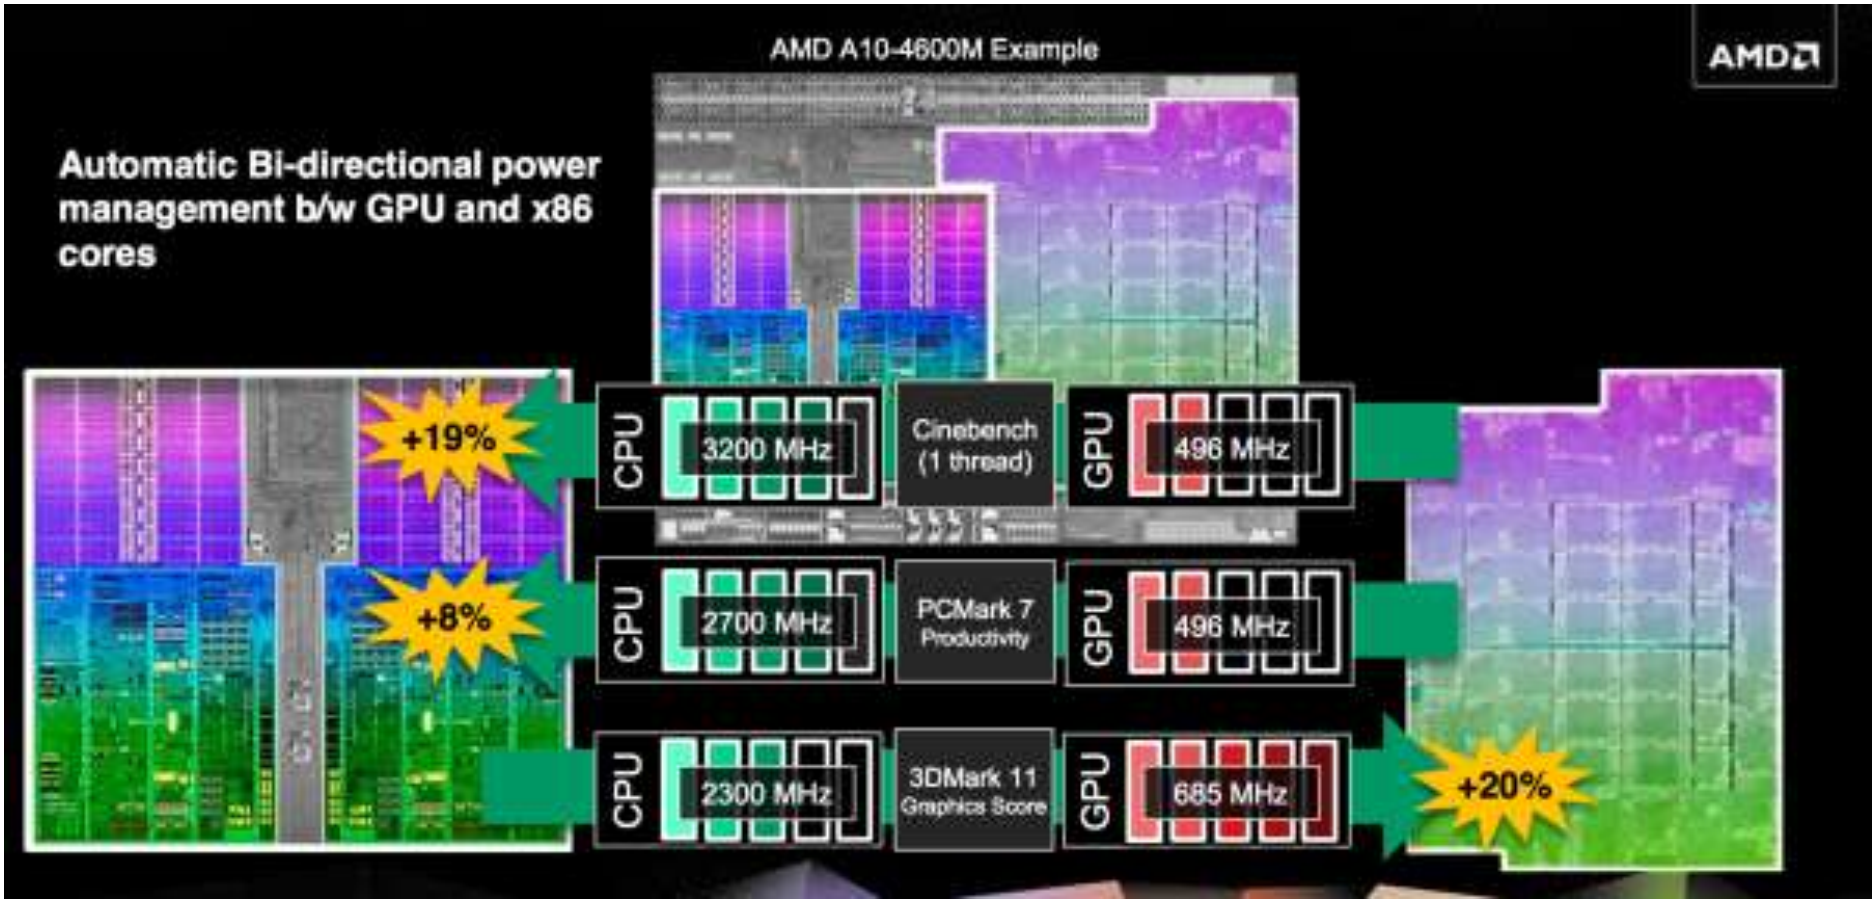

Figure: Example for the operation of the AMD Turbo Core Technology 3.0 [55]

3.4.2 The Trinity APU die (12)

Related to the above Figure we note that in the Trinity A10-4600M mobile APU

- the compute modules have a base clock frequency of 2300 MHz that can be boosted up to 3200 MHz, in steps of 100 MHz, whereas
- the GPU has a base clock frequency of 496 MHz that can be raised to 685 MHz.

Now, according to the actual load pattern, the Turbo core manager (not shown in the Figure) will increase either the core frequency of the compute units or the GPU, when the OS requires a performance increase and enough power headroom is available, as demonstrated in the Figure below.

Illustration of the operation of the Turbo Core Technology 3.0 of the Trinity APU [77]

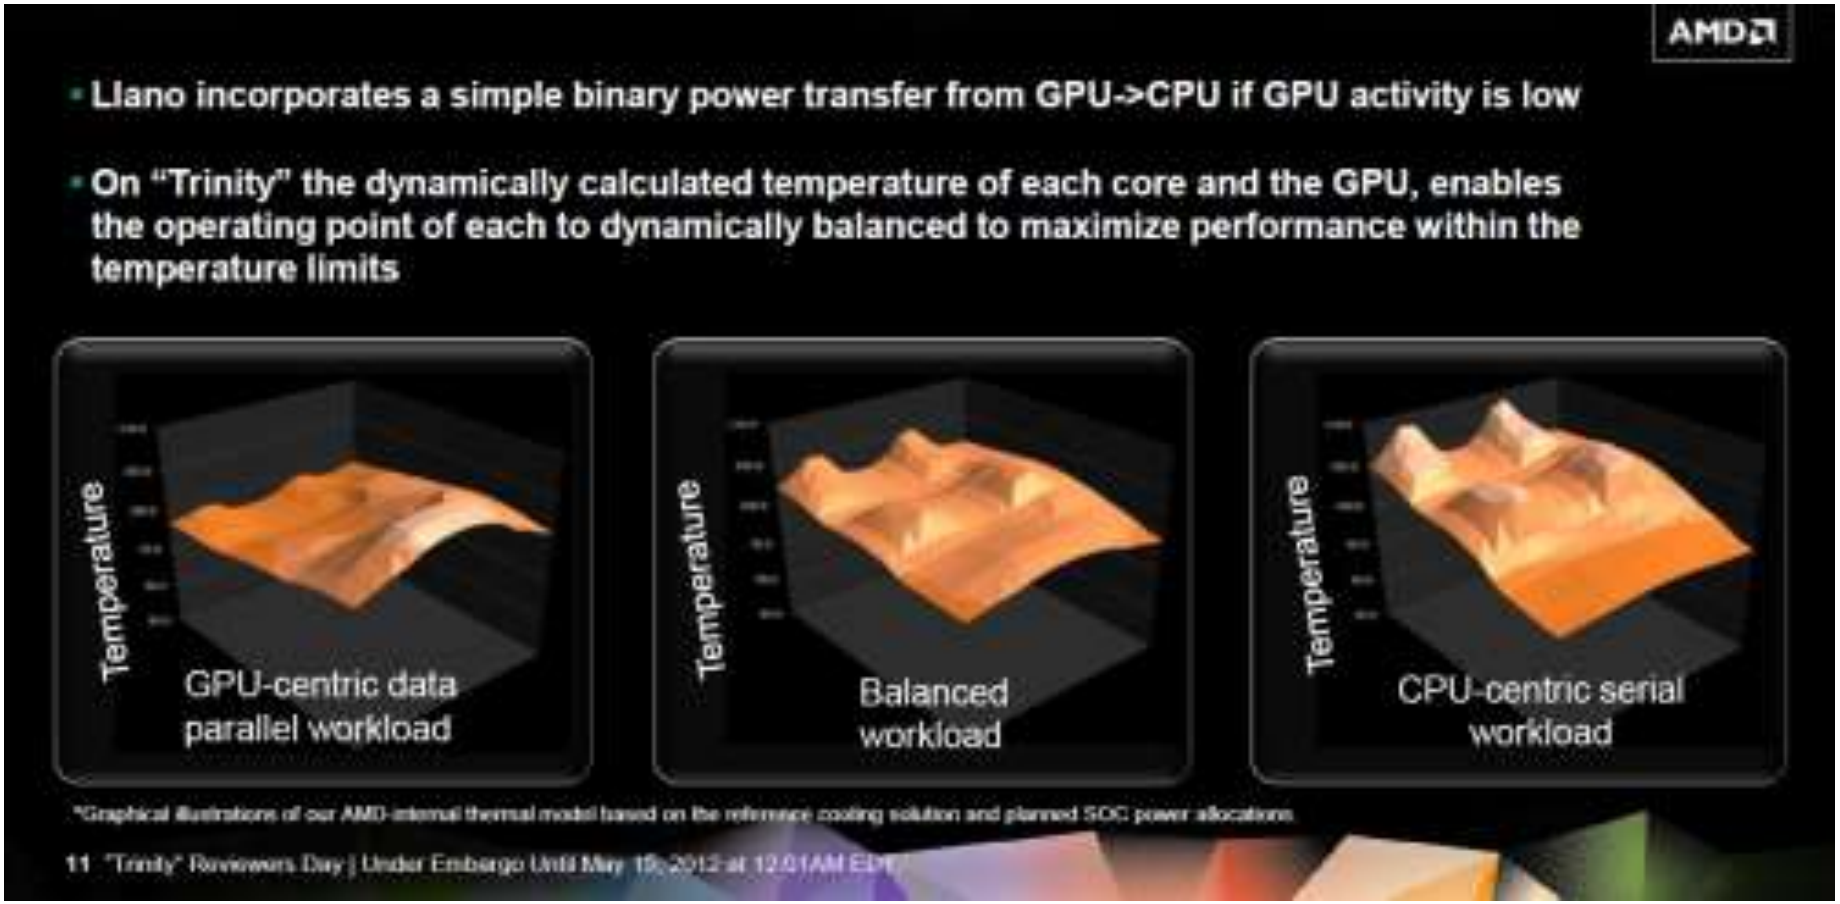

3.4.2 The Trinity APU die (14)

On the other hand, when there is a power overage, the Turbo core manager initiates a clock throttling of the compute units or the GPU to reduce power dissipation below the allowed TDP limit (not shown in the Figure).

3.4.3 The Trinity mainstream desktop APU line

3.4.3 The Trinity mainstream desktop APU line

Piledriver-based processor lines

Piledriver-based GPU-less processor lines

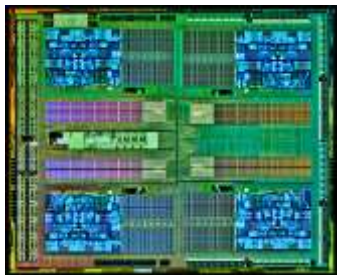

32 nm
315 mm^2
1.2 billion trans.

Piledriver-based Trinity APU lines

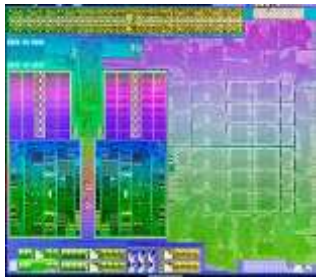

32 nm
226 mm^2
1.3 billion trans.

Piledriver v2-based Richland APU line

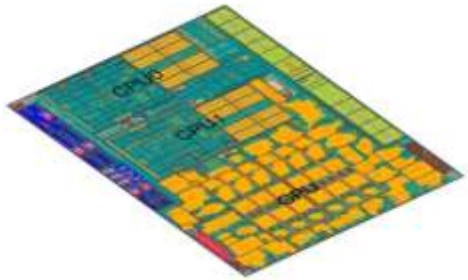

Abu Dhabi	**Vishera**	**Trinity**	**Trinity**	**Richland**	**Richland**
Opteron server line (Dual chips)	high performance FX desktop line	mainstream DT APU line HD 7000D GPU	mainstr./ultra-thin mobile APU line HD 7000G GPU	mainstream DT APU line HD 8000D GPU	mainstr./ultra-thin mobile APU lines HD 8000M GPU
4P: Opteron 6300	*FX-x3xx*	*A10 - A4 5x00(K)*	*A10 – A6 4xxxM*	*A10 – A4 6x00(K)*	*A10 – A4 5xxxM*
G34 socket (retained)	AM3+ socket (retained)	FM2 socket (new)	FS1r2/FP2 socket (new)	FM2 socket (retained)	FS1r2 socket (retained)
	Volan platform	*Virgo platform*	*Comal platform*	*Elite A-Series*	*Elite performance platform*
11/2012	10/2012	9/2012	5/2012	6/2013	5/2013

Positioning of the Trinity mainstream desktop APU line [51]

3.4.3 The Trinity mainstream desktop APU line (3)

Main components of the Trinity mainstream desktop APU [78]

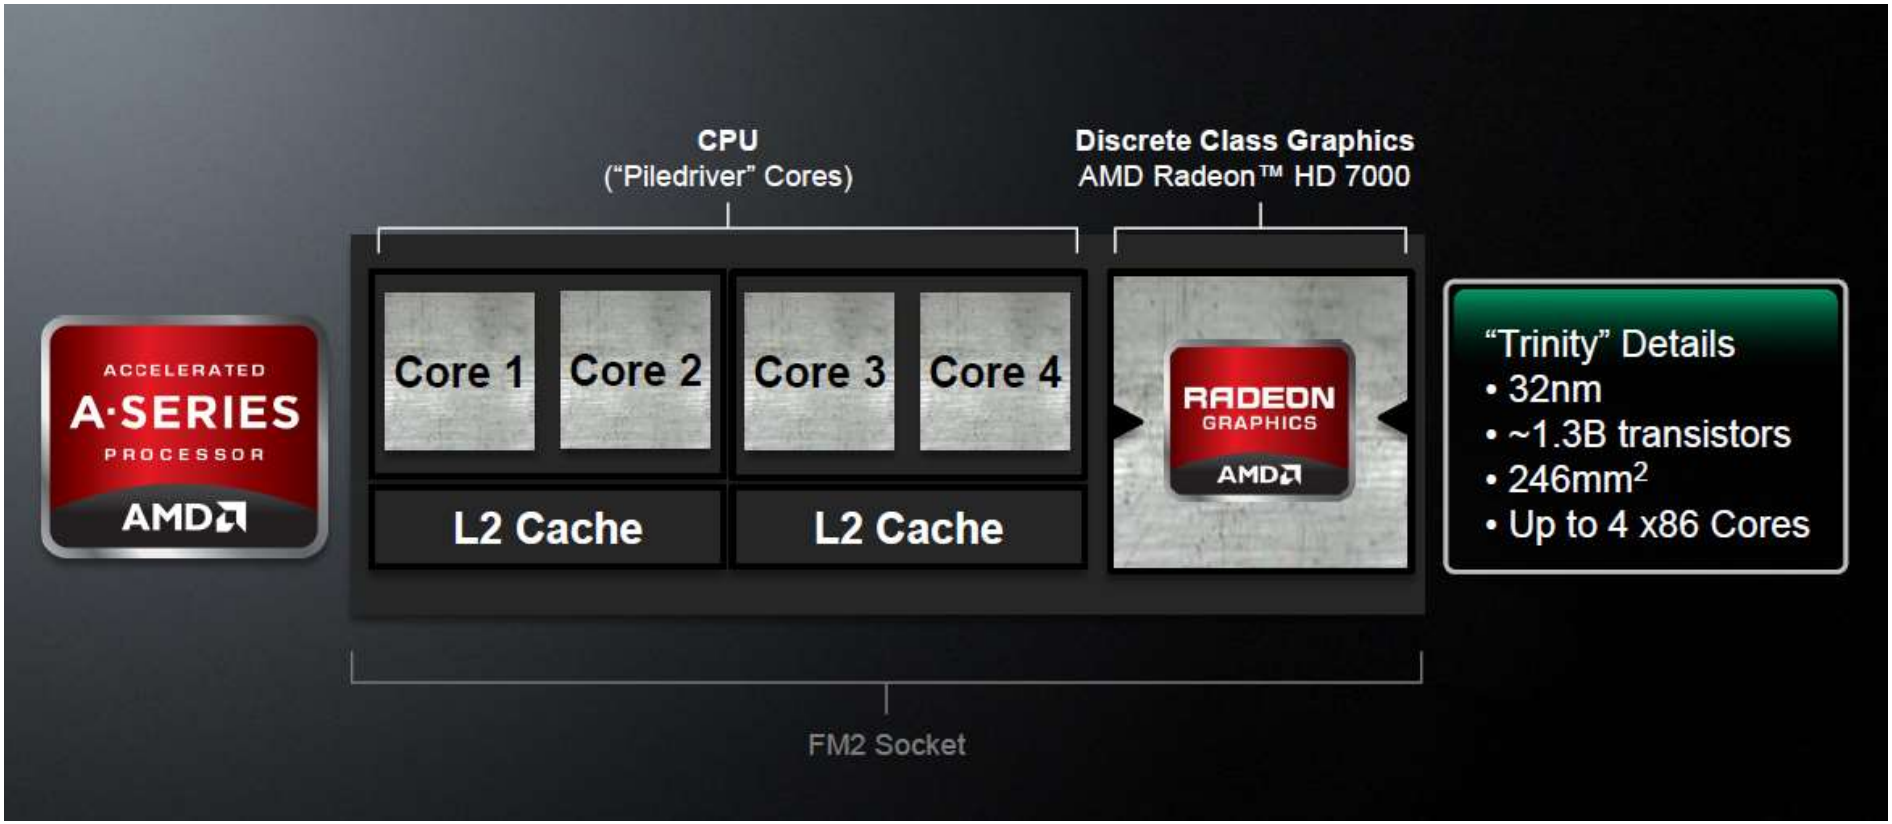

Note: No L3 cache

3.4.3 The Trinity mainstream desktop APU line (4)

Model numbers and main features of the mainstream Trinity desktop APU line [78]
(Virgo platform)

APU Model	A10-5800K	A10-5700	A8-5600K	A8-5500	A6-5400K	A4-5300
AMD Radeon™ Graphics Brand	HD 7660D	HD 7660D	HD 7560D	HD 7560D	HD 7540D	HD 7480D
TDP	100W	65W	100W	65W	65W	65W
AMD Radeon™ Cores	384	384	256	256	192	128
GPU Clock Speed	800 MHz	800 MHz	760 MHz	760 MHz	760 MHz	723 MHz
CPU Cores	4	4	4	4	2	2
CPU Clock (Max Turbo/Base)	4.2/3.8 GHz	4.0/3.4 GHz	3.9/ 3.6 GHz	3.7/ 3.2 GHz	3.8/ 3.6 GHz	3.6/ 3.4 GHz
Total Cache	4MB	4MB	4MB	4MB	1MB	1MB
Max DDR3	1866	1866	1866	1866	1866	1600
AMD Turbo CORE 3.0	Yes	Yes	Yes	Yes	Yes	Yes
Unlocked[1]	Yes	No	Yes	No	Yes	No

3.4.3 The Trinity mainstream desktop APU line (5)

The new FM2 socket of the Trinity mainstream desktop APU line [78]

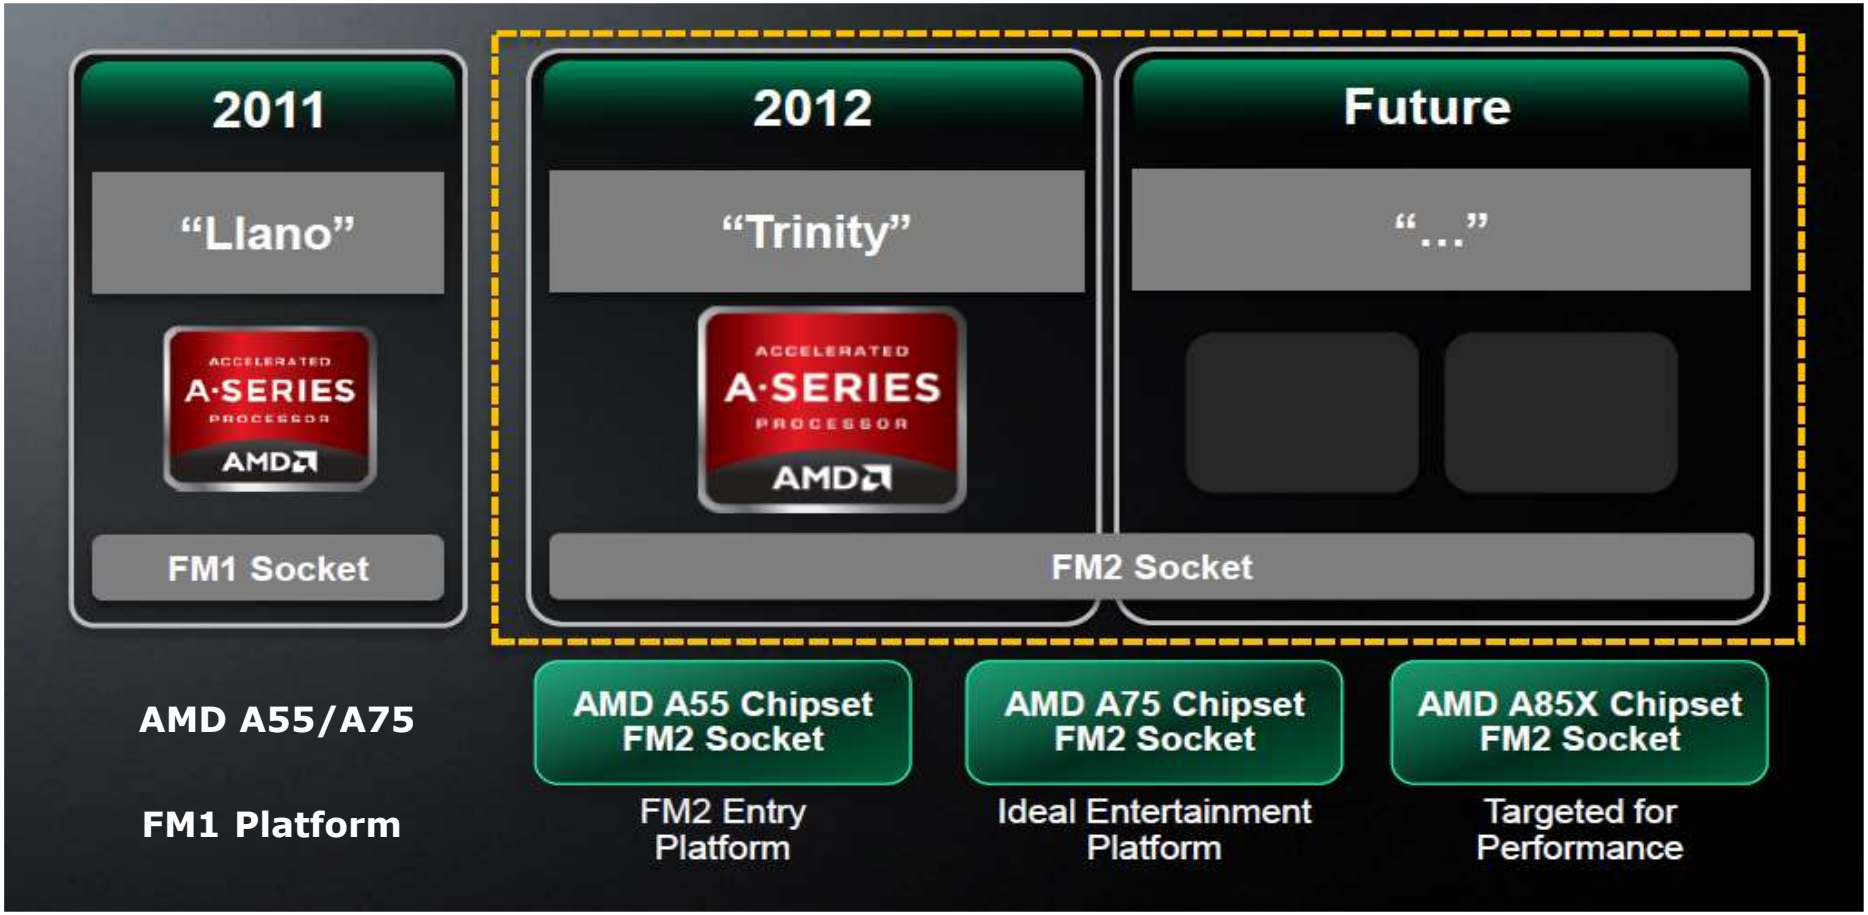

Remark

The A55/A75 FCHs (Fusion Controller Hubs) were already introduced for the Llano A-Series APUs, the A85 is new, it supports the high performance unlocked models of the line.

System architecture of the mainstream Trinity desktop APU with the A85X FCH [79]

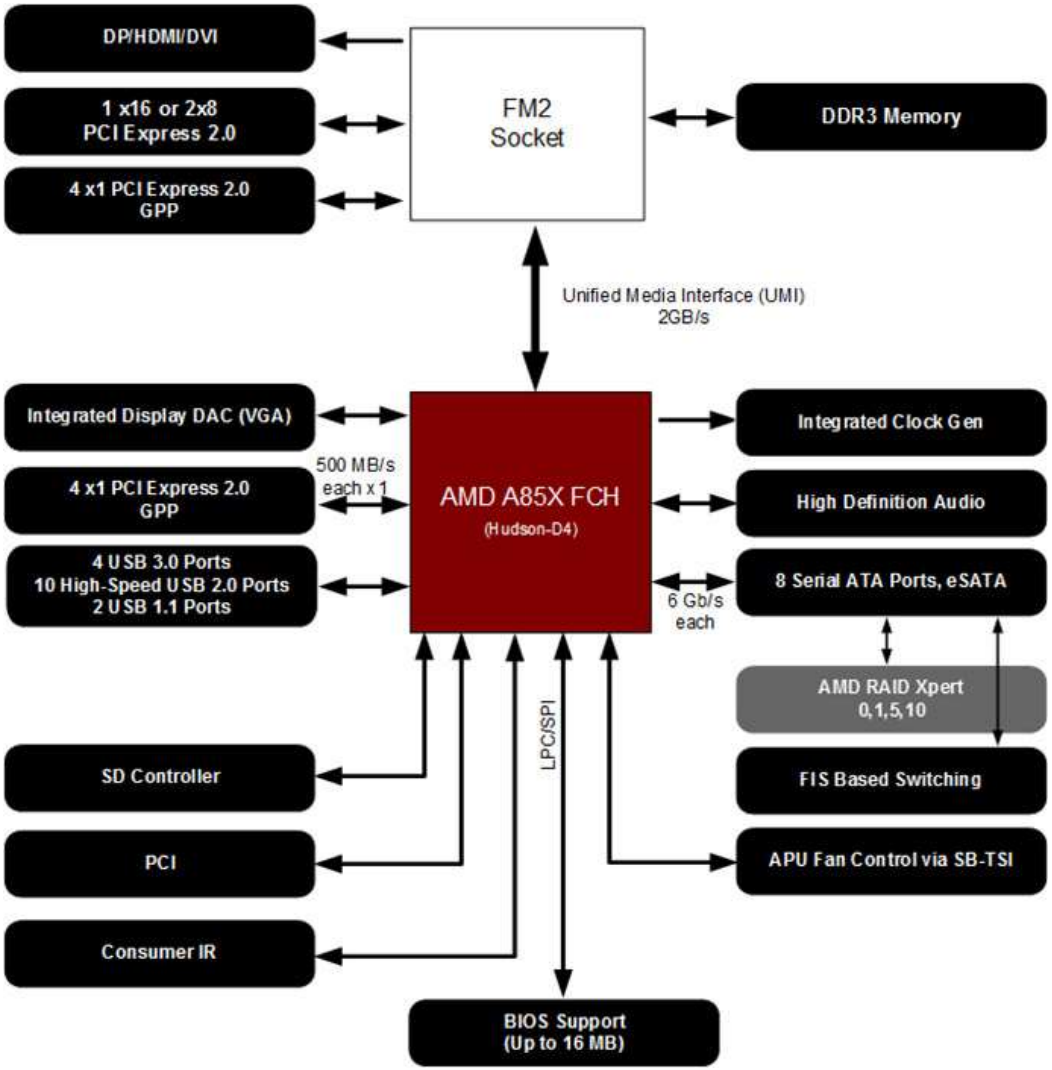

The A85X FCH supports the high performance K models (unlocked models).

Performance increase achieved over the previous A-Series Llano APU line [78]

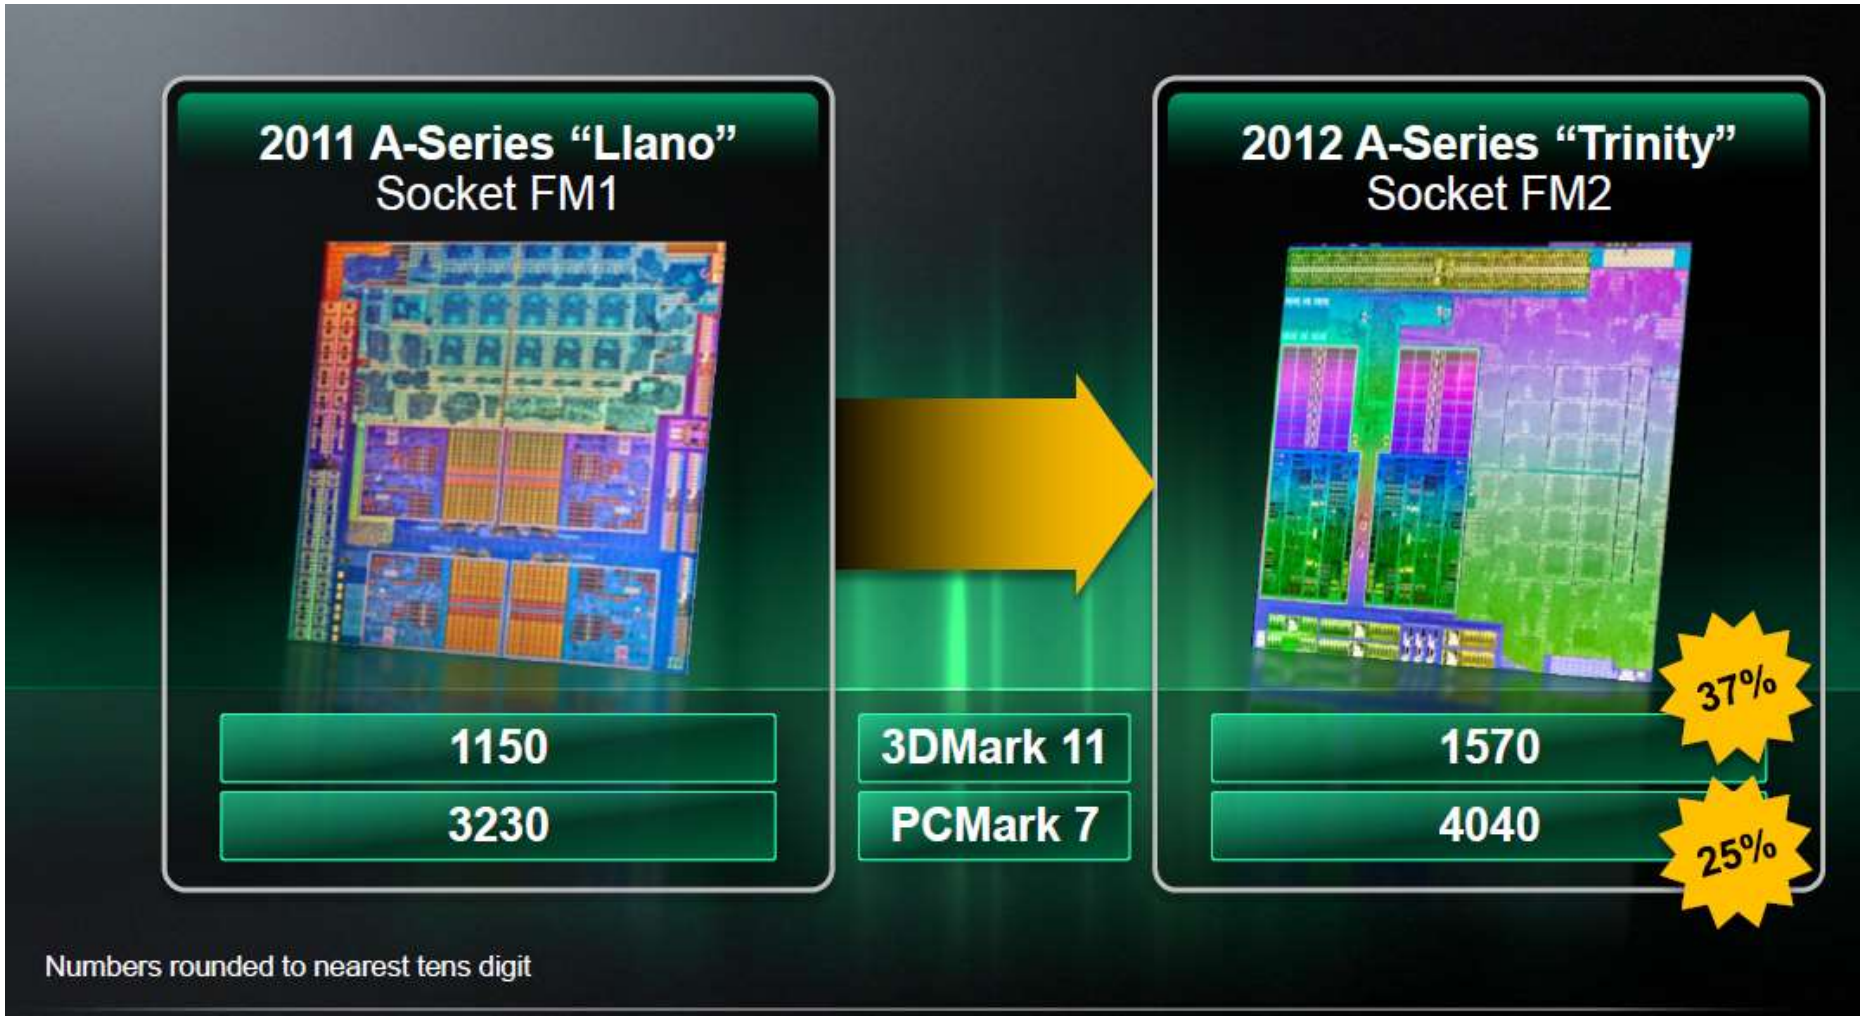

3.4.4 The Trinity mobile APU line

3.4.4 The Trinity mobile APU line

Piledriver-based processor lines

Piledriver-based GPU-less processor lines

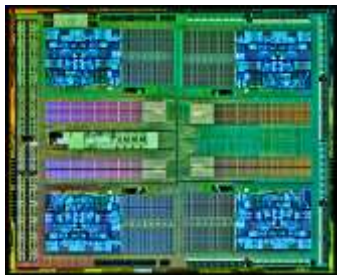

32 nm
315 mm^2
1.2 billion trans.

Piledriver-based Trinity APU lines

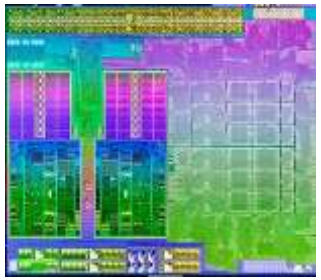

32 nm
226 mm^2
1.3 billion trans.

Piledriver v2-based Richland APU line

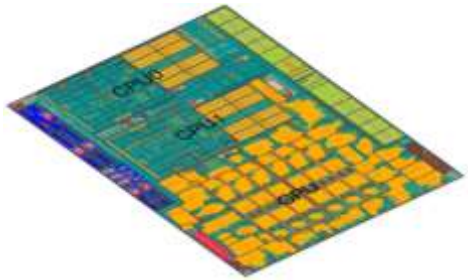

Abu Dhabi	**Vishera**	**Trinity**	**Trinity**	**Richland**	**Richland**
Opteron server line (Dual chips)	high performance FX desktop line	mainstream DT APU line HD 7000D GPU	mainstr./ultra-thin mobile APU line HD 7000G GPU	mainstream DT APU line HD 8000D GPU	mainstr./ultra-thin mobile APU lines HD 8000M GPU
4P: Opteron 6300	*FX-x3xx*	*A10 - A4 5x00(K)*	*A10 – A6 4xxxM*	*A10 – A4 6x00(K)*	*A10 – A4 5xxxM*
G34 socket (retained)	AM3+ socket (retained)	FM2 socket (new)	FS1r2/FP2 socket (new)	FM2 socket (retained)	FS1r2 socket (retained)
	Volan platform	*Virgo platform*	*Comal platform*		*Elite performance platform*
11/2012	10/2012	9/2012	5/2012		5/2013

3.4.4 The Trinity mobile APU line (2)

Positioning of the Trinity mainstream and ultra-thin mobile APU lines -1 [51]

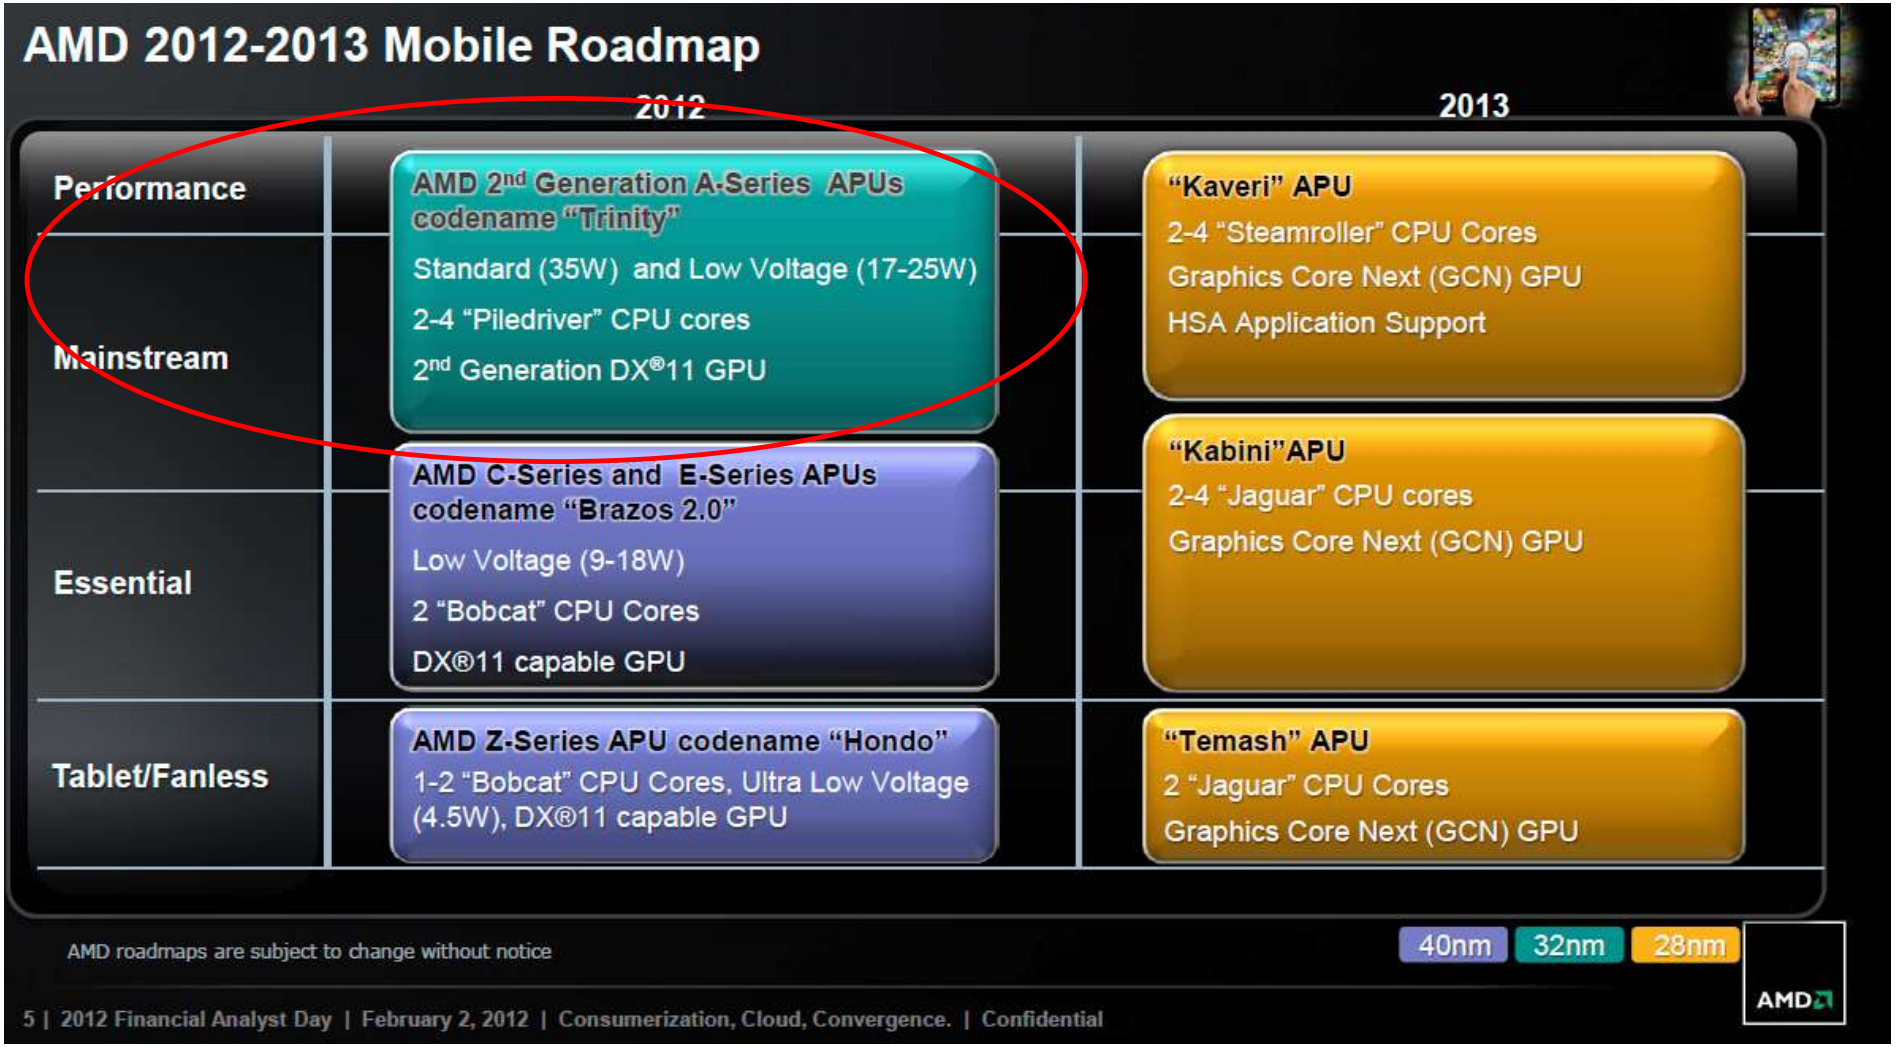

3.4.4 The Trinity mobile APU line (3)

Positioning of the Trinity mainstream and ultra-thin mobile APU lines -2 [52]

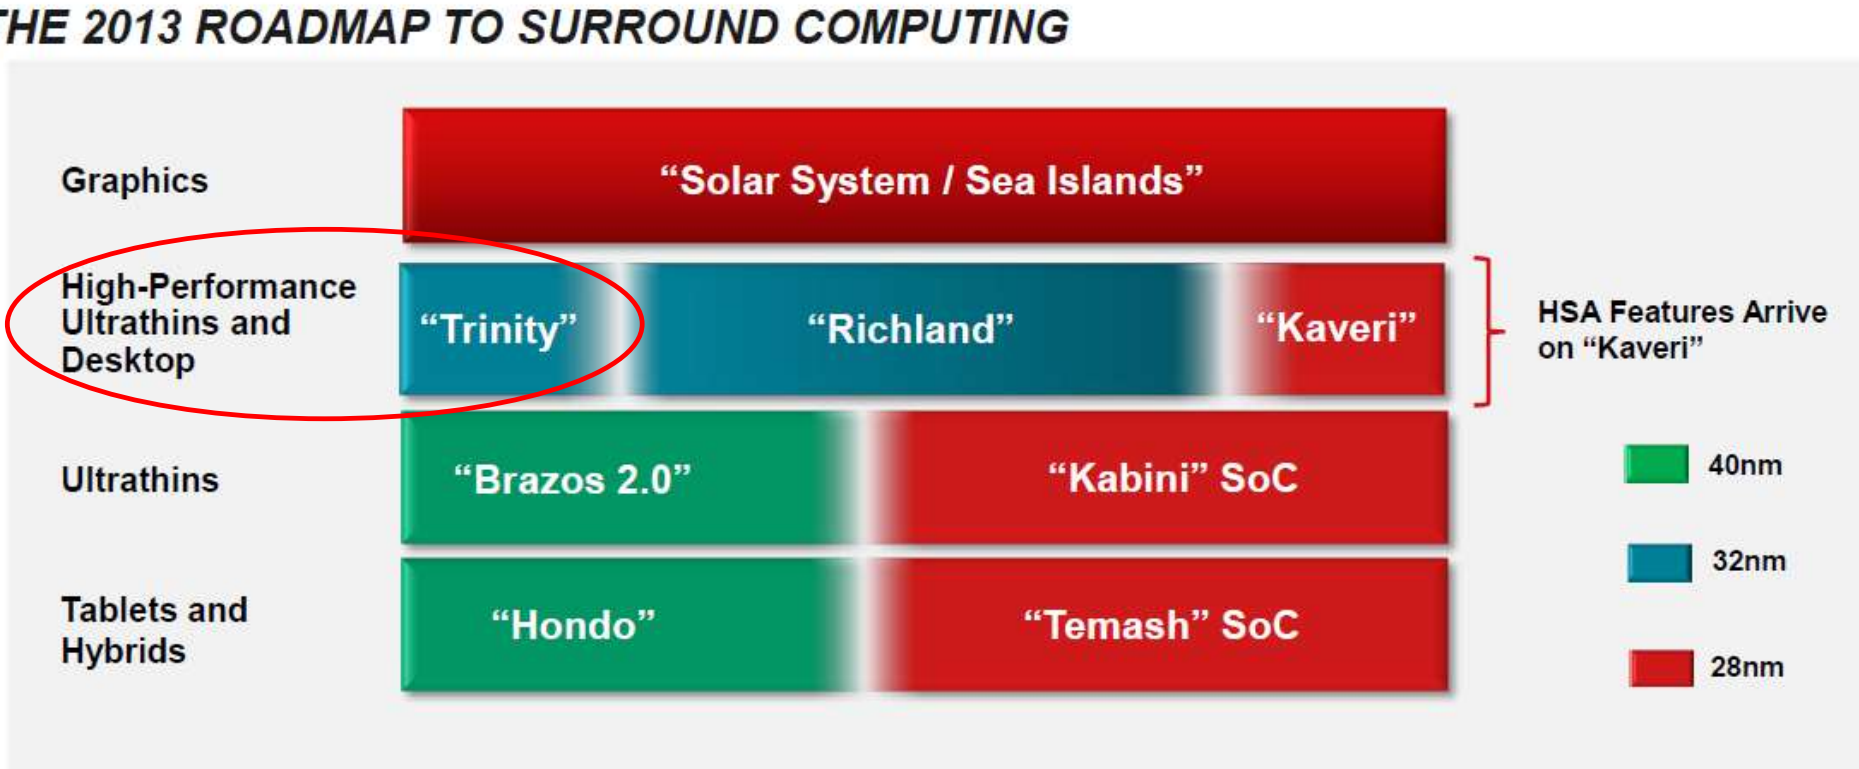

AMD’s Trinity mainstream and ultra-thin mobile APU lines

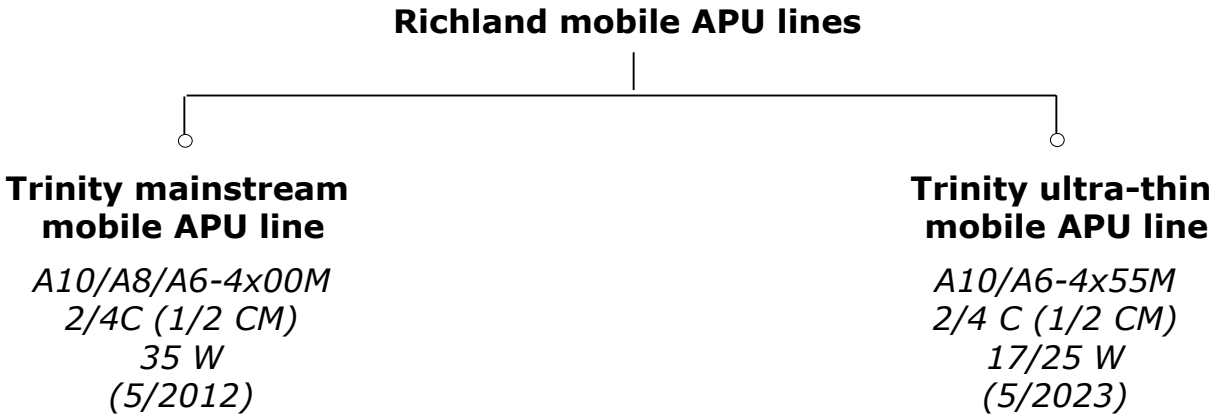

3.4.4 The Trinity mobile APU line (4)

Model numbers and main features of the Trinity mobile APU line [80]

(Comal platform)

Model	Radeon™ Brand	OPN	Package	TDP	CPU Cores	CPU Clock (Max/Base)	L2 Cache	Radeon™ Cores[1]	GPU Clock (Max/Base)	Max DDR3
AMD A-Series Processors										
A10-4600M	HD 7660G	AM4600DEC44HJ	FS1r2	35W	4	3.2GHz/2.3GHz	4MB	384	686MHz/497 MHz	DDR3-1600 DDR3L-1600 DDR3U-1333
A8-4500M	HD 7640G	AM4500DEC44HJ	FS1r2	35W	4	2.8GHz/1.9GHz	4MB	256	655MHz/497 MHz	DDR3-1600 DDR3L-1600 DDR3U-1333
A6-4400M	HD 7520G	AM4400DEC23HJ	FS1r2	35W	2	3.2GHz/2.7GHz	1MB	192	686MHz/497 MHz	DDR3-1600 DDR3L-1600 DDR3U-1333

Model	Radeon™ Brand	OPN	Package	TDP	CPU Cores	CPU Clock (Max/Base)	L2 Cache	Radeon™ Cores[1]	GPU Clock (Max/Base)	Max DDR3
A-Series "Trinity" LV and ULV APUs										
A10-4655M	HD 7620G	AM4655SIE44HJ	FP2	25W	4	2.8GHz/2.0GHz	4MB	384	497MHz/360 MHz	DDR3-1333 DDR3L-1333 DDR3U-1066
A6-4455M	HD 7500G	AM4455SHE24HJ	FP2	17W	2	2.6GHz/2.1GHz	2MB	256	424MHz/327 MHz	DDR3-1333 DDR3L-1333 DDR3U-1066

Trinity based A10/A8 mobiles 5/2012

3.4.4 The Trinity mobile APU line (5)

The Comal mobile platform including the (Piledriver-based) Trinity APU and the A70M/A60M FCH [52]

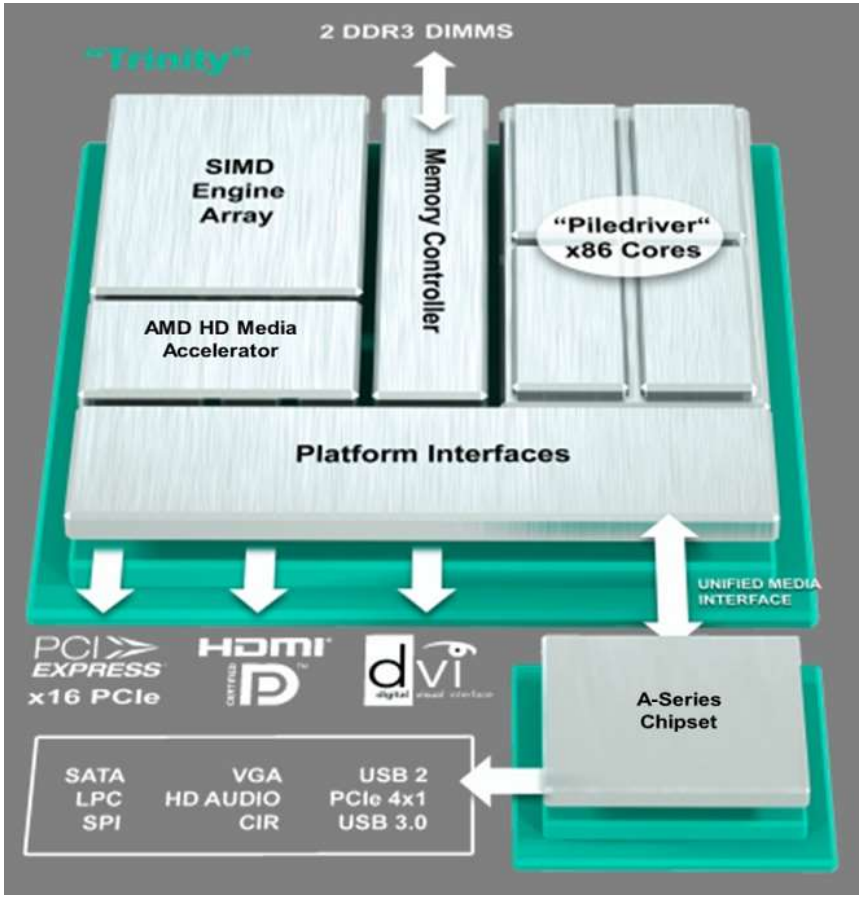

A70M/A60M FCH

3.5 Piledriver v2-based Richland APU lines

- 3.5.1 Overview of the Piledriver v2-based Richland APU lines
- 3.5.2 The Richland mainstream desktop APU line
- 3.5.3 The Richland mobile APU line

3.5.1 Overview of the Piledriver v2-based Richland APU lines

3.5.1 Overview of the Piledriver v2-based Trinity APU lines

Piledriver-based processor lines

Piledriver-based GPU-less processor lines

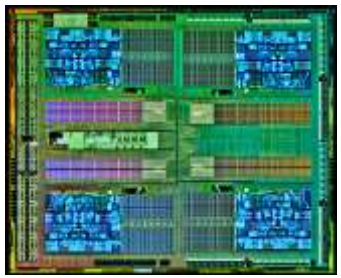

32 nm
315 mm^2
1.2 billion trans.

Piledriver-based Trinity APU lines

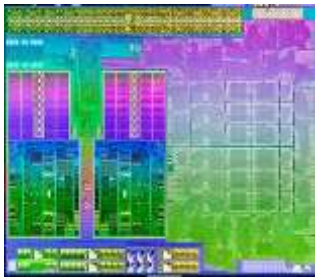

32 nm
226 mm^2
1.3 billion trans.

Piledriver v2-based Richland APU line

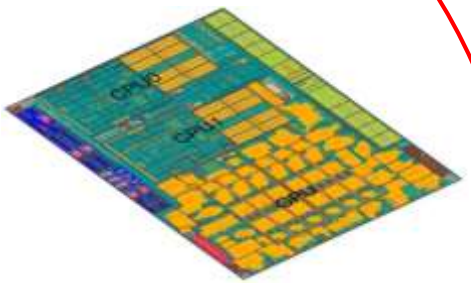

Abu Dhabi	**Vishera**	**Trinity**	**Trinity**	**Richland**	**Richland**
Opteron server line (Dual chips)	high performance FX desktop line	mainstream DT APU line HD 7000D GPU	mainstr./ultra-thin mobile APU line HD 7000G GPU	mainstream DT APU line HD 8000D GPU	mainstr./ultra-thin mobile APU lines HD 8000M GPU
4P: Opteron 6300	*FX-x3xx*	*A10 - A4 5x00(K)*	*A10 – A6 4xxxM*	*A10 – A4 6x00(K)*	*A10 – A4 5xxxM*
G34 socket (retained)	AM3+ socket (retained)	FM2 socket (new)	FS1r2/FP2 socket (new)	FM2 socket (retained)	FS1r2 socket (retained)
	Volan platform	*Virgo platform*	*Comal platform*	*Elite A-Series*	*Elite performance platform*
11/2012	10/2012	9/2012	5/2012	6/2013	5/2013

3.5.1 Overview of the Piledriver-based Richland APU lines (2)

Positioning of the Trinity mainstream desktop and mobile APU lines [52]

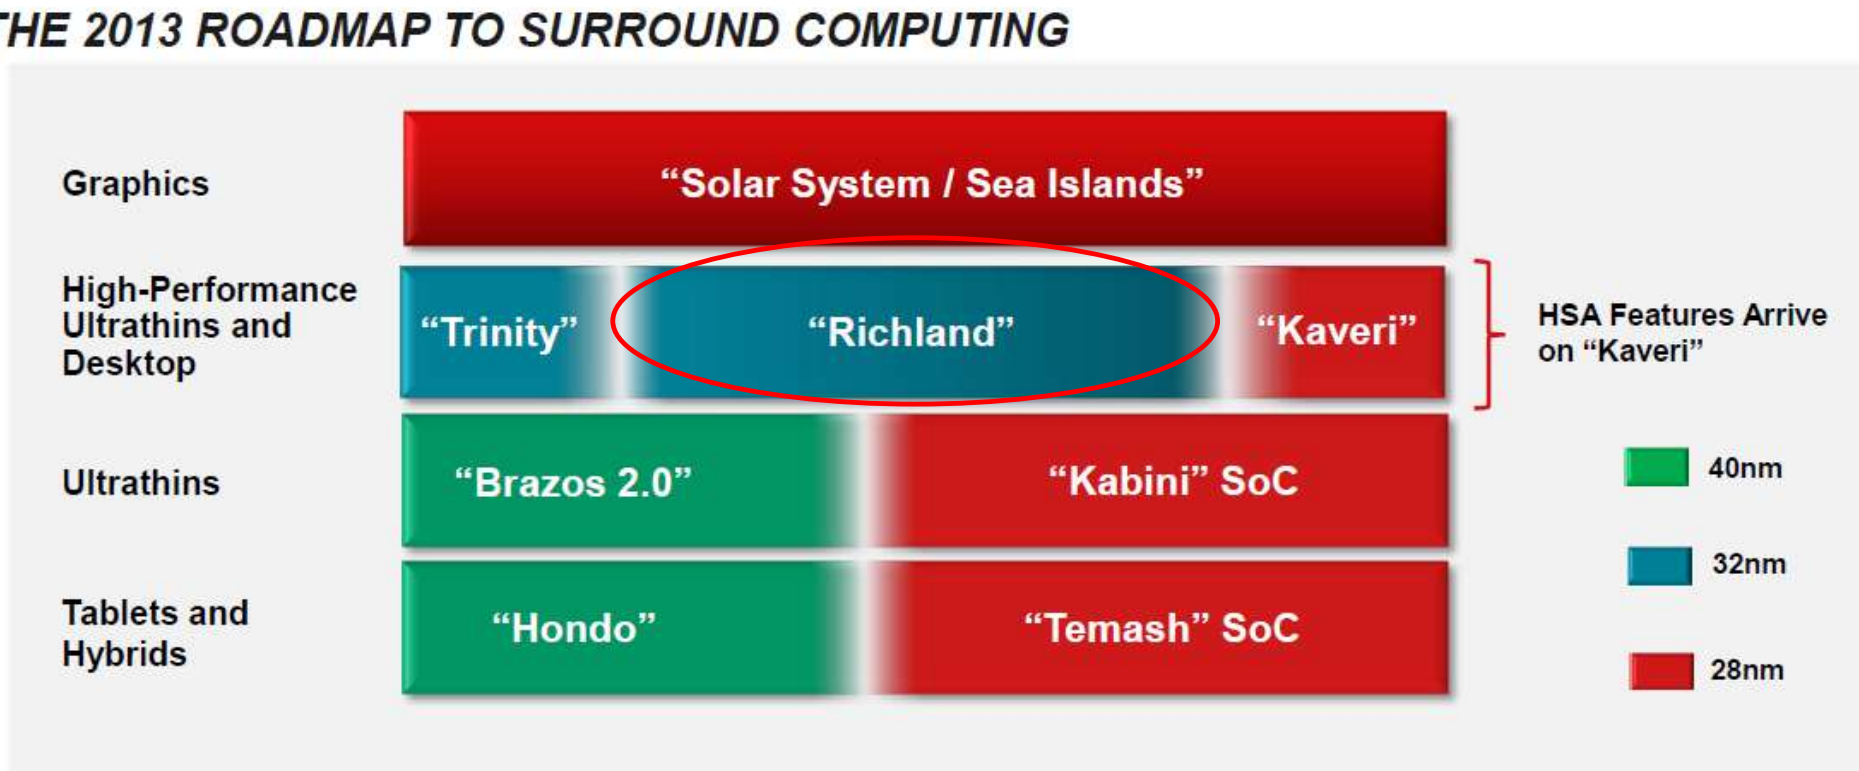

3.5.1 Overview of the Piledriver-based Richland APU lines (3)

Die shot of the Richland APU [81]

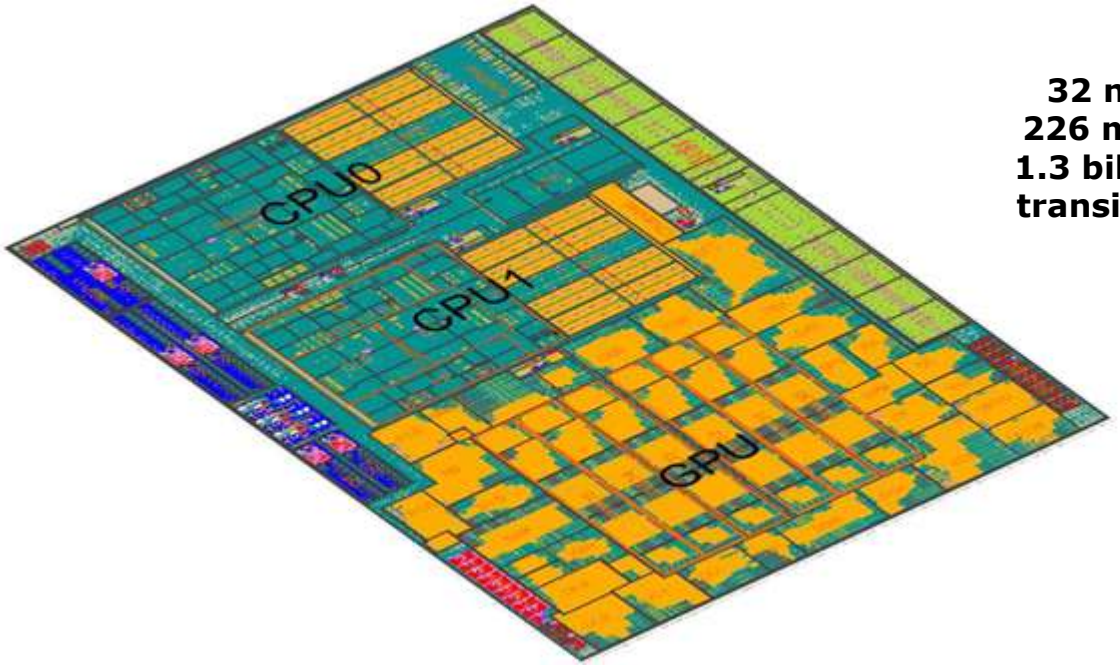

32 nm
226 mm²
1.3 billion transistors

Key features of the Richland mobile APU line as exposed by AMD [82]

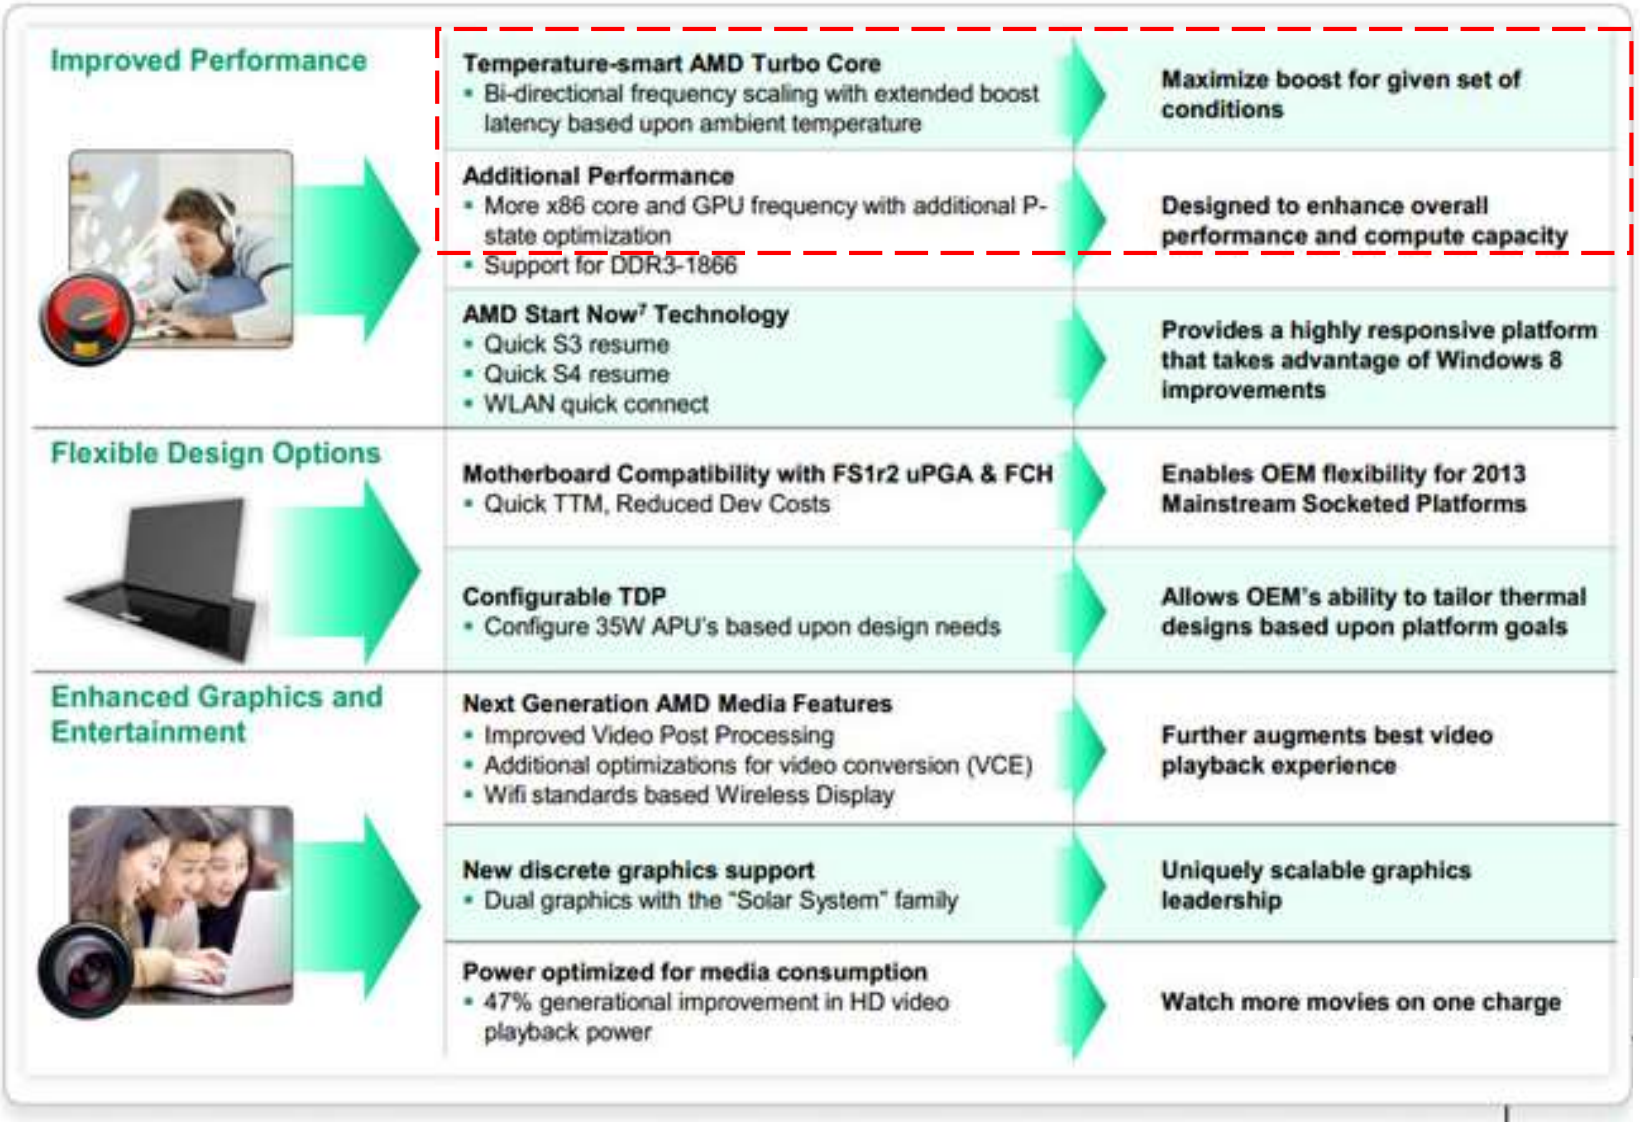

3.5.1 Overview of the Piledriver-based Richland APU lines (5)

Major improvements of the Richland mobile APU line discussed [83], [84]

Richland APUs are based 2. gen. Piledriver cores, are fabricated at the same feature size as the Trinity APUs and incorporate also the same number of transistors.

Their major improvements are as follows:

- about 10 % faster CPU and 4-7 % faster GPU clock speed both in base mode and turbo mode,
- new, HD 8000G series GPUs that are based however further on on the Cayman core (Northern Islands family) including VLIW4 ALUs.

  The new GPUs are claimed to provide 20-40 % more graphics performance in high-end models than the previous HD 7000G series GPUs,
- improved power management, including
  - an enhanced power management technique, called the Temperature Smart Turbo Core (TSTC) that increases battery life (to be detailed subsequentl),
  - introducing additional frequency/voltage operating points (P points) to enhance the efficiency of power management (to be detailed subsequently), and
  - innovative software features (to be detailed subsequently).

Principle of operation of the Temperature Smart Turbo Core (TSTC) technique-1

It enhances turbo core management by including 17 temperature sensors

- 5 on each compute module and
- 7 on the GPU

along with a package sensor (not shown), as indicated in the Figure [82].

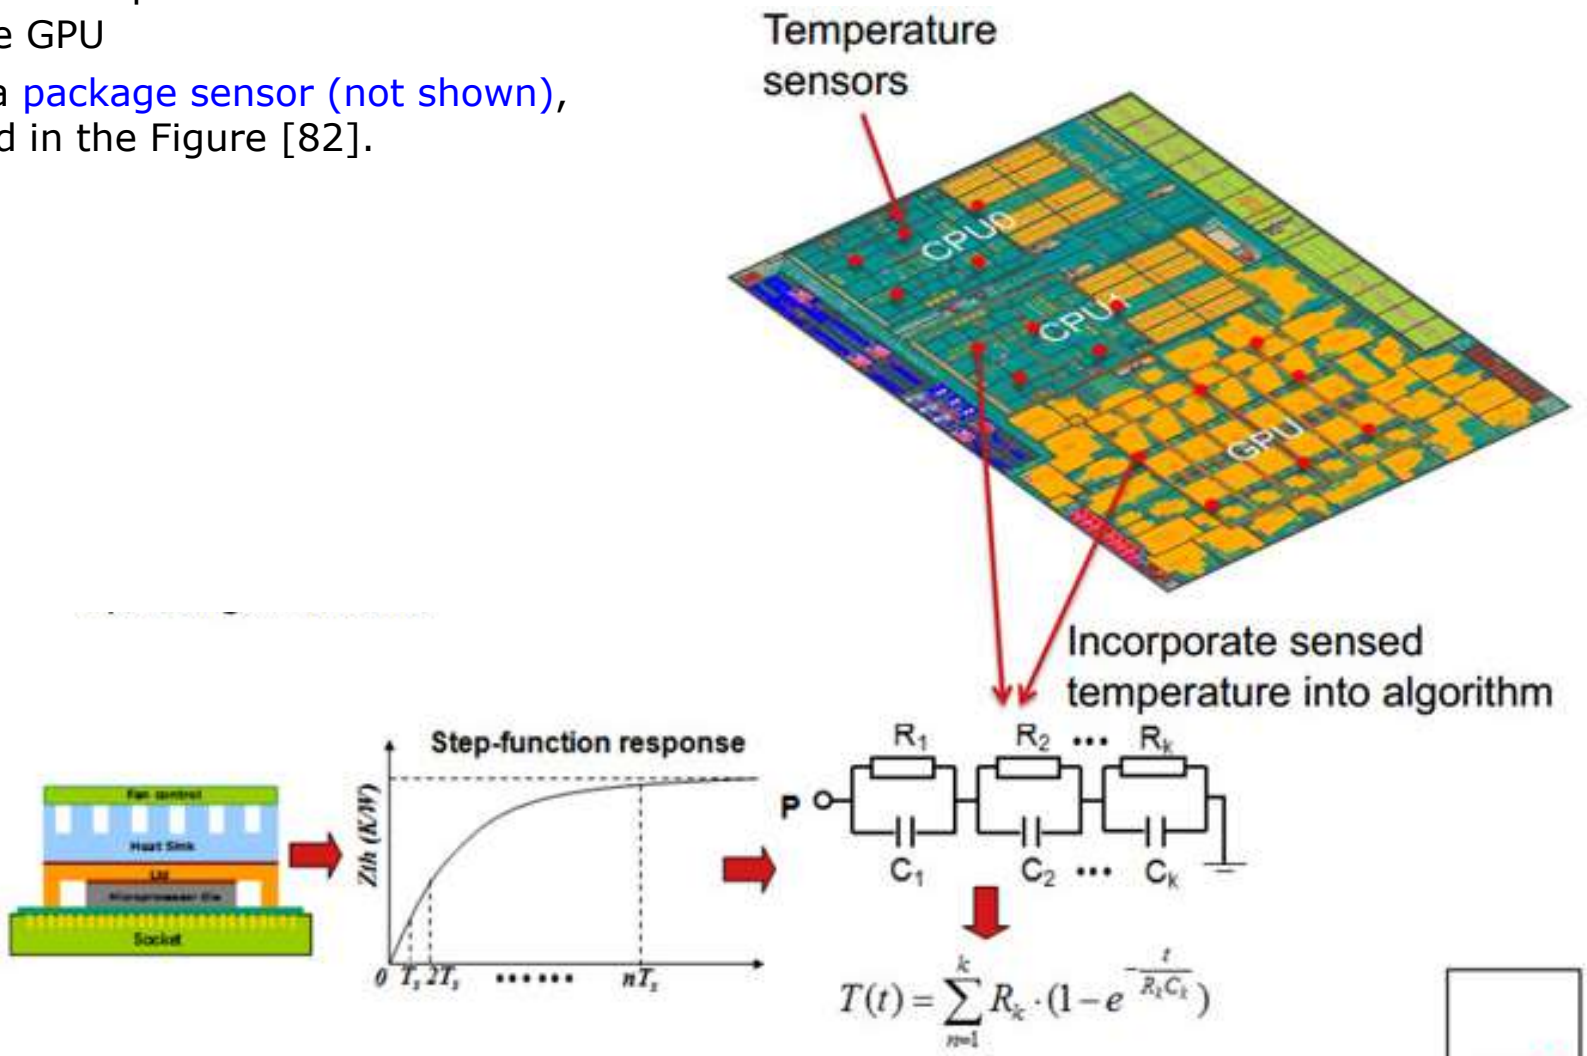

Principle of operation of the Temperature Smart Turbo Core (TSTC) technique-2 [85]

Taking into account temperature data of the compute modules, the GPU and the package, delivered by the respective sensors, allows the Turbo Core Manager more sophisticated real-time clock speed settings of both the compute modules and the GPU according to the actual load pattern while staying within the chip's thermal limits.

This results typically in higher clock speed than granted by the previous Trinity APU implementation, as demonstrated in the next Table.

3.5.1 Overview of the Piledriver-based Richland APU lines (8)

Comparing clock frequencies of the Richland and the Trinity mobile APU lines [86]

A-Series APU Mobile Lineup (2013)									
Model	Code Name	Cores	CPU Clock Base / Turbo	L2 Cache	GPU Core	Radeon Cores	GPU Clock Base / Max	TDP	Max DDR3
A10-5750M	Richland	4	2.5 / 3.5 GHz	4MB	HD 8650G	384	533 / 720 MHz	35W	1866
A10-4600M	Trinity	4	2.3 / 3.2 GHz	4MB	HD 7660G	384	497 / 686 MHz	35W	1600
A8-5550M	Richland	4	2.1 / 3.1 GHz	4MB	HD 8550G	256	515 / 720 MHz	35W	1600
A8-4500M	Trinity	4	1.9 / 2.8 GHz	4MB	HD 7640G	256	497 / 655 MHz	35W	1600
A6-5350M	Richland	2	2.9 / 3.5 GHz	1MB	HD 8450G	192	533 / 720 MHz	35W	1600
A6-4400M	Trinity	2	2.7 / 3.2 GHz	1MB	HD 7520G	192	497 / 686 MHz	35W	1600
A4-5150M	Richland	2	2.7 / 3.3 GHz	1MB	HD 8350G	128	514 / 720 MHz	35W	1600
A4-4300M	Trinity	2	2.5 / 3.0 GHz	1MB	HD 7420G	128	480 / 655 MHz	35W	1600

Principle of operation of the Temperature Smart Turbo Core (TSTC) technique-3 [85]

Handling of possible bottlenecks

A further improvement of the power management algorithm relates to handling of possible bottlenecks.

The previous algorithm granted higher clock speed to the compute modules or the GPU if required, regardless whether or not a higher clock speed could be realized due to possible resource bottlenecks.

The new algorithm takes possible bottlenecks into account, and grants higher clock frequencies only if bottlenecks do not limit the utilization of the higher clock speed granted.

Introducing additional frequency/voltage operating points

With the Richland APU line AMD added more frequency/voltage operating points, termed as P points.

P points are used to adjust dynamically the operating point of the individual compute modules or the GPU to the actual performance need, determined by the OS, as indicated in the Figure.

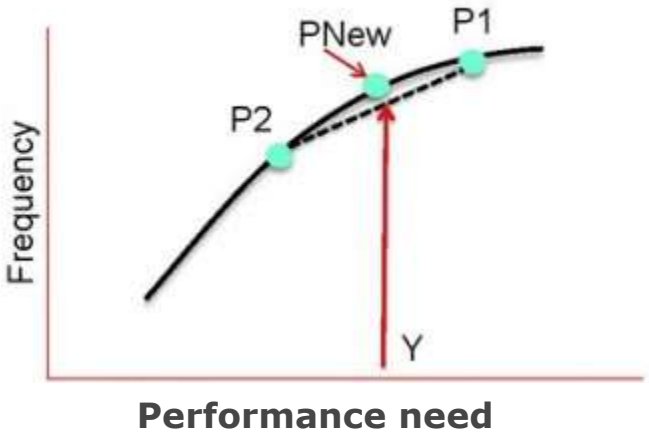

Figure: Additional frequency/voltage points (P points) introduced in the Richland APU [85]

New frequency/voltage operating points (P points) enable an improved adjustment of the chosen operating point to the actual performance need.

This results in a more efficient power management in terms of less power consumption for a given workload (i.e. performance need).

3.5.2 The Richland mainstream desktop APU line

3.5.2 The Richland mainstream desktop APU line

Piledriver-based processor lines

Piledriver-based GPU-less processor lines

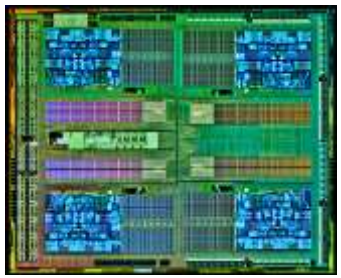

32 nm
315 mm^2
1.2 billion trans.

Abu Dhabi
Opteron server line
(Dual chips)

4P: Opteron 6300

G34 socket
(retained)

11/2012

Vishera
high performance
FX desktop line

FX-x3xx

AM3+ socket
(retained)

Volan platform

10/2012

Piledriver-based Trinity APU lines

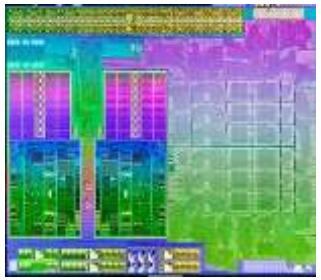

32 nm
226 mm^2
1.3 billion trans.

Trinity
mainstream
DT APU line
HD 7000D GPU

A10 - A4 5x00(K)

FM2 socket
(new)

Virgo platform

9/2012

Trinity
mainstr./ultra-thin
mobile APU line
HD 7000G GPU

A10 – A6 4xxxM

FS1r2/FP2 socket
(new)

Comal platform

5/2012

Piledriver v2-based Richland APU line

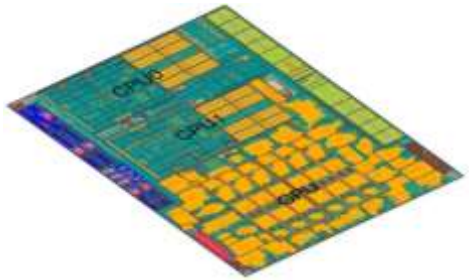

Richland
mainstream
DT APU line
HD 8000D GPU

A10 – A4 6x00(K)

FM2 socket
(retained)

Elite A-Series

6/2013

Richland
mainstr./ultra-thin
mobile APU lines
HD 8000M GPU

A10 – A4 5xxxM

FS1r2 socket
(retained)

Elite performance platform

5/2013

Positioning of the Richland mainstream desktop and mobile APU lines [52]

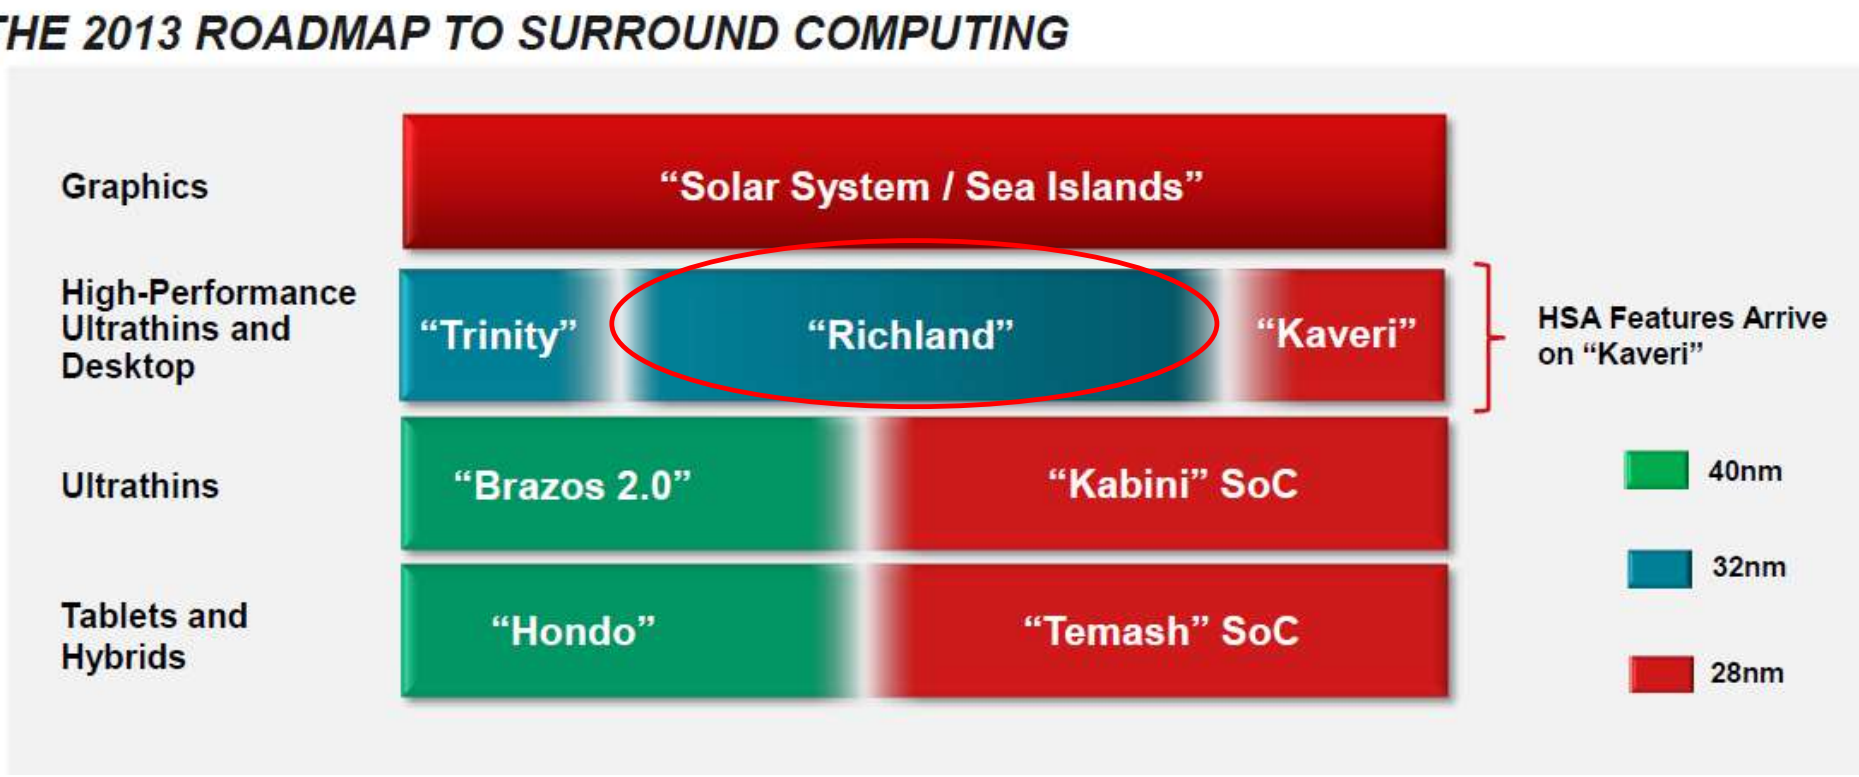

3.5.2 The Richland mainstream desktop APU line (3)

Model numbers and expected key features of the Richland desktop APU line [89]
(Elite Experience platform)

Model	Radeon™ Brand	TDP	Radeon ™ Cores	GPU Clock Speed	CPU Cores	CPU Clock (Max Turbo/ Base)	Total Cache	MAX DDR3	AMD Turbo Core	Unlock	Est Price
A10-6800K	HD 8670D	100W	384	844 MHz	4	4.4/4.1 GHz	4MB	2133	Yes	Yes	$149
A10-6700	HD 8670D	65W	384	844MHz	4	4.3/3.7 GHz	4MB	1866	Yes	No	$149
A8-6600K	HD 8570D	100W	256	844 MHz	4	4.2/3.9 GHz	4MB	1866	Yes	Yes	$119
A8-6500	HD 8570D	65W	256	800 MHz	4	4.1/3.5 GHz	4MB	1866	Yes	No	$119
A6-6400k	HD 8470D	65W	192	800 MHz	2	4.1/3.9 GHz	1MB	1866	Yes	Yes	$77

Remark

Subsequently (in 07/2013) Intel launched also the model A4-6300.

3.5.3 The Richland mobile APU line

3.5.3 The Richland mobile APU line

Piledriver-based processor lines

Piledriver-based GPU-less processor lines

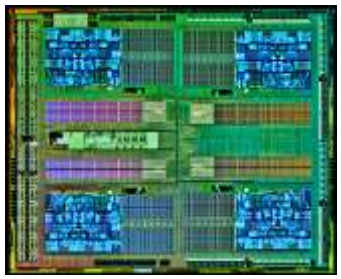

32 nm
315 mm^2
1.2 billion trans.

Piledriver-based Trinity APU lines

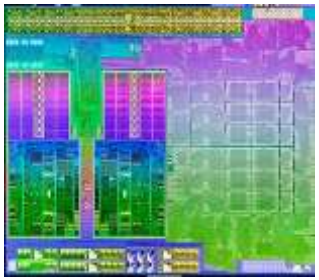

32 nm
226 mm^2
1.3 billion trans.

Piledriver-based Richland APU line

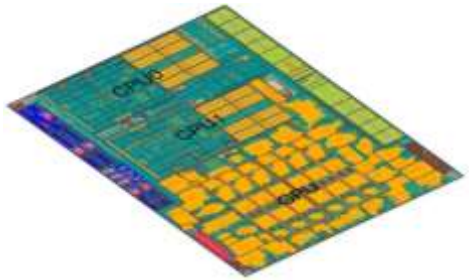

Abu Dhabi	**Vishera**	**Trinity**	**Trinity**	**Richland**	**Richland**
Opteron server line (Dual chips)	high performance FX desktop line	mainstream DT APU line HD 7000D GPU	mainstr./ultra-thin mobile APU line HD 7000G GPU	mainstream DT APU line HD 8000D GPU	mainstr./ultra-thin mobile APU lines HD 8000M GPU
4P: Opteron 6300	*FX-x3xx*	*A10 - A4 5x00(K)*	*A10 – A6 4xxxM*	*A10 – A4 6x00(K)*	*A10 – A4 5xxxM*
G34 socket (retained)	AM3+ socket (retained)	FM2 socket (new)	FS1r2/FP2 socket (new)	FM2 socket (retained)	FS1r2 socket (retained)
	Volan platform	*Virgo platform*	*Comal platform*	*Elite A-Series*	*Elite performance platform*
11/2012	10/2012	9/2012	5/2012	6/2013?	5/2013

3.5.3 The Richland mobile APU line (2)

AMD's Richland mainstream and ultra-thin mobile APU lines

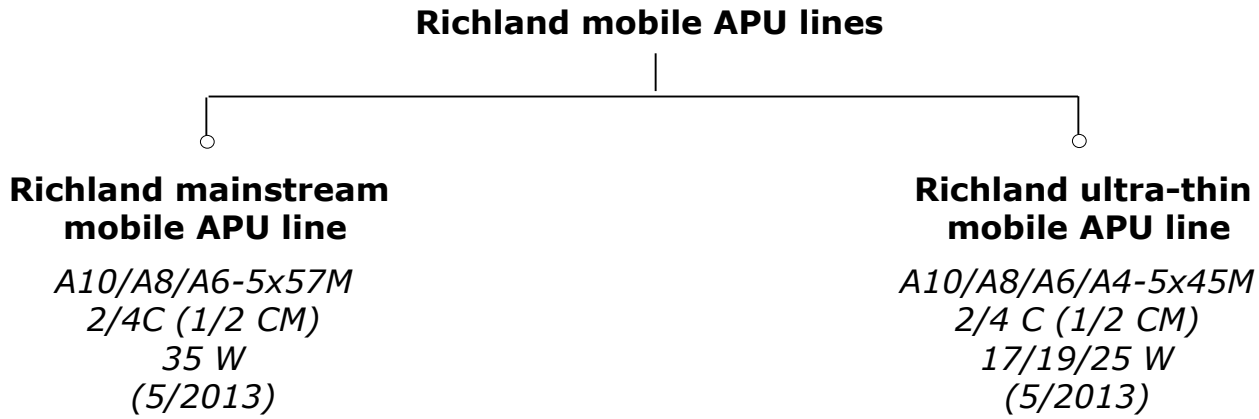

Positioning of the Richland mobile APU lines [52]
(Elite performance APU platform)

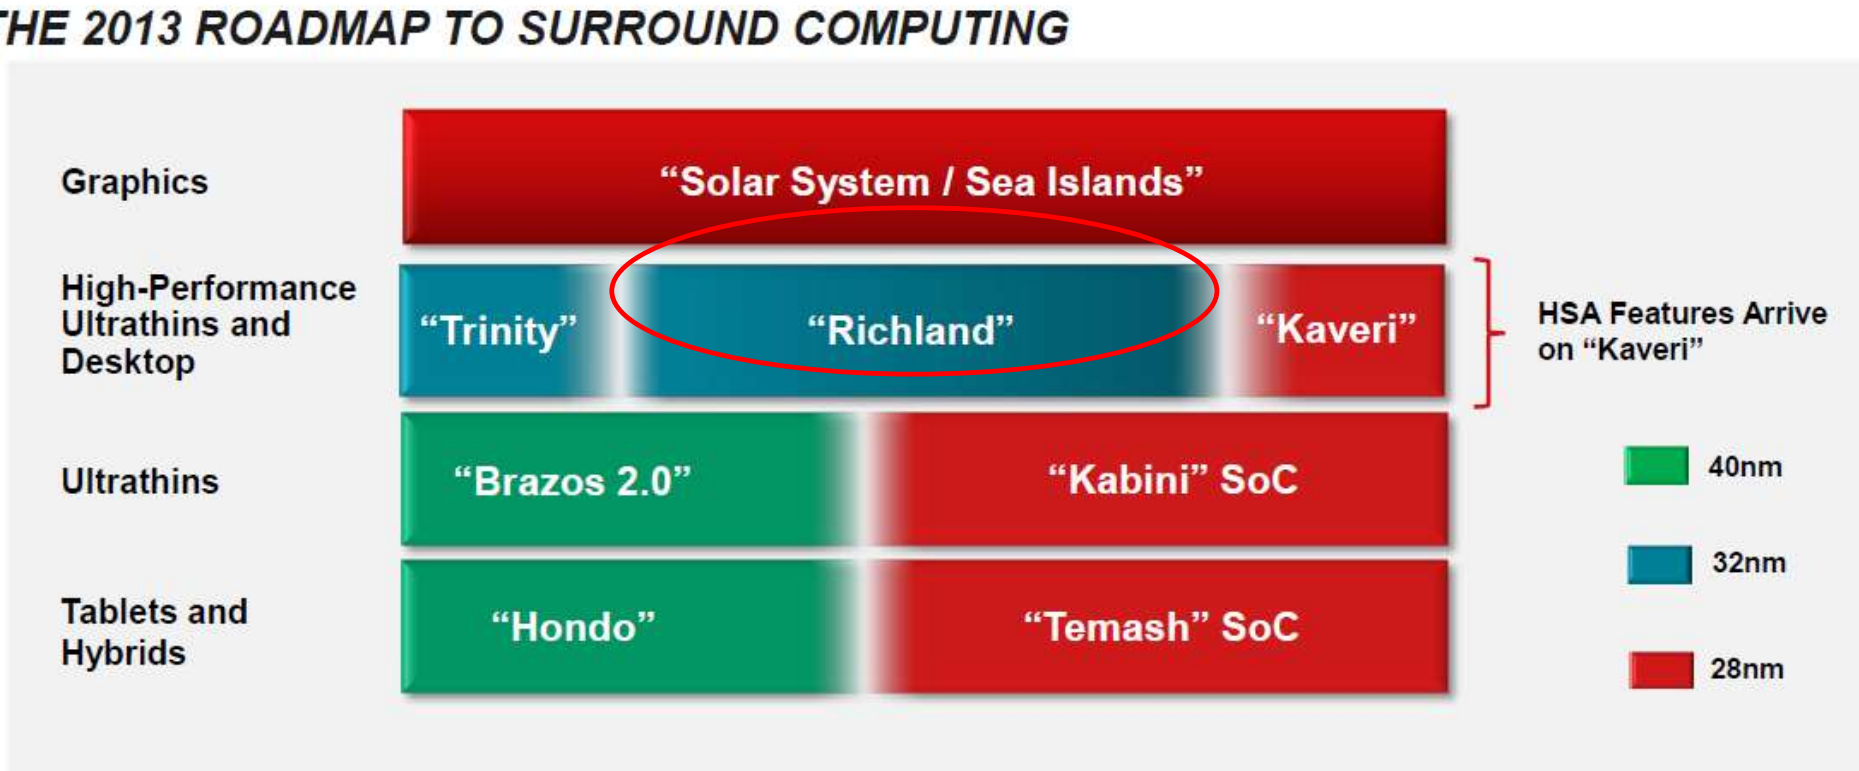

3.5.3 The Richland mobile APU line (4)

Main features of AMD's Richland mainstream mobile APU line [116]

Model	Radeon Core	Package	TDP	CPU Cores	CPU Clock (Max/Base)	L2 Cache	GPU Cores	GPU Clock (Base/Max)	DDR3
A10-5757M	HD 8650G	FP2	35W	4	3.5/2.5 GHz	4 MB	384	600/720 MHz	1600
A8-5557M	HD 8550G	FP2	35W	4	3.1/2.1 GHz	4 MB	256	554/720 MHz	1600
A6-5357M	HD 8450G	FP2	35W	2	3.5/2.9 GHz	1 MB	192	533/720 MHz	1600

3.5.3 The Richland mobile APU line (4b)

Main features of AMD's Richland ultra-thin mobile APU line [116]

Model	Radeon Core	Package	TDP	CPU Cores	CPU Clock (Max/Base)	L2 Cache	GPU Cores	GPU Clock (Base/Max)	DDR3
A10-5745M	HD 8610G	FP2	25W	4	2.9/2.1 GHz	4 MB	384	533/626 MHz	1333
A8-5545M	HD 8510G	FP2	19W	4	2.7/1.7 GHz	4 MB	384	450/554 MHz	1333
A6-5345M	HD 8410G	FP2	17W	2	2.8/2.2 GHz	1 MB	192	450/600 MHz	1333
A4-5145M	HD 8310G	FP2	17W	2	2.6/2.0 GHz	1 MB	128	424/554 MHz	1333

3.5.3 The Richland mobile APU line (5)

Remark

The new chips have the same socket as the previous Trinity mobile APU line (FS1r2 socket), so they are drop-in compatible with the previous platforms and OEMs can quickly ramp up systems based on Richland mobile APUs.

An innovative suite of apps. available on the Richland mobile APU models [87]

AMD Face Login[7]

AMD Screen Mirror[8]

AMD Optimized Games

AMD Gesture Control[9]

3.5.3 The Richland mobile APU line (7)

AMD Face Login [88]

- It is designed as a convenient tool to help log-in to Windows and many popular web sites quickly but should not be used to protect the computer and personal information from unwanted access.
- Only available on the Richland A10 and A8 APUs.
- Requires a webcam, and will only operate on PCs running Windows 7 or Windows 8 operating systems and Internet Explorer version 9 or 10.

AMD Gesture Control [88]

- It is designed to enable gesture recognition as a tool for controlling certain applications on the PC.
- Only available on Richland A10 and A8 APUs.
- Requires a web camera, and will only operate on PCs running Windows 7 or Windows 8.
- Supported Windows desktop apps include: Windows Media Player, Windows Photo Viewer, Microsoft PowerPoint and Adobe Acrobat Reader.
- Supported Windows Store apps include: Microsoft Photos, Microsoft Music, Microsoft Reader and Kindle.

3.5.3 The Richland mobile APU line (8)

AMD Screen Mirror [88]

- It is designed to enable the transmission and display of the PC screen on other compatible networked "mirror" devices.
- Only available on Richland A10, A8 and A6 APUs.
- AMD Screen Mirror supports almost all popular image, audio and video file formats as well as applications, but will not mirror protected content.

AMD optimized games [88]

- Provides driver optimizations for a select set of games.
- The optimized-for-AMD software will be pre-loaded on select Richland A-Series APU-based notebooks or is downloadable from AMD’s website.

3.5.3 The Richland mobile APU line (9)

Support of the new features on mobile Richland processors [87]

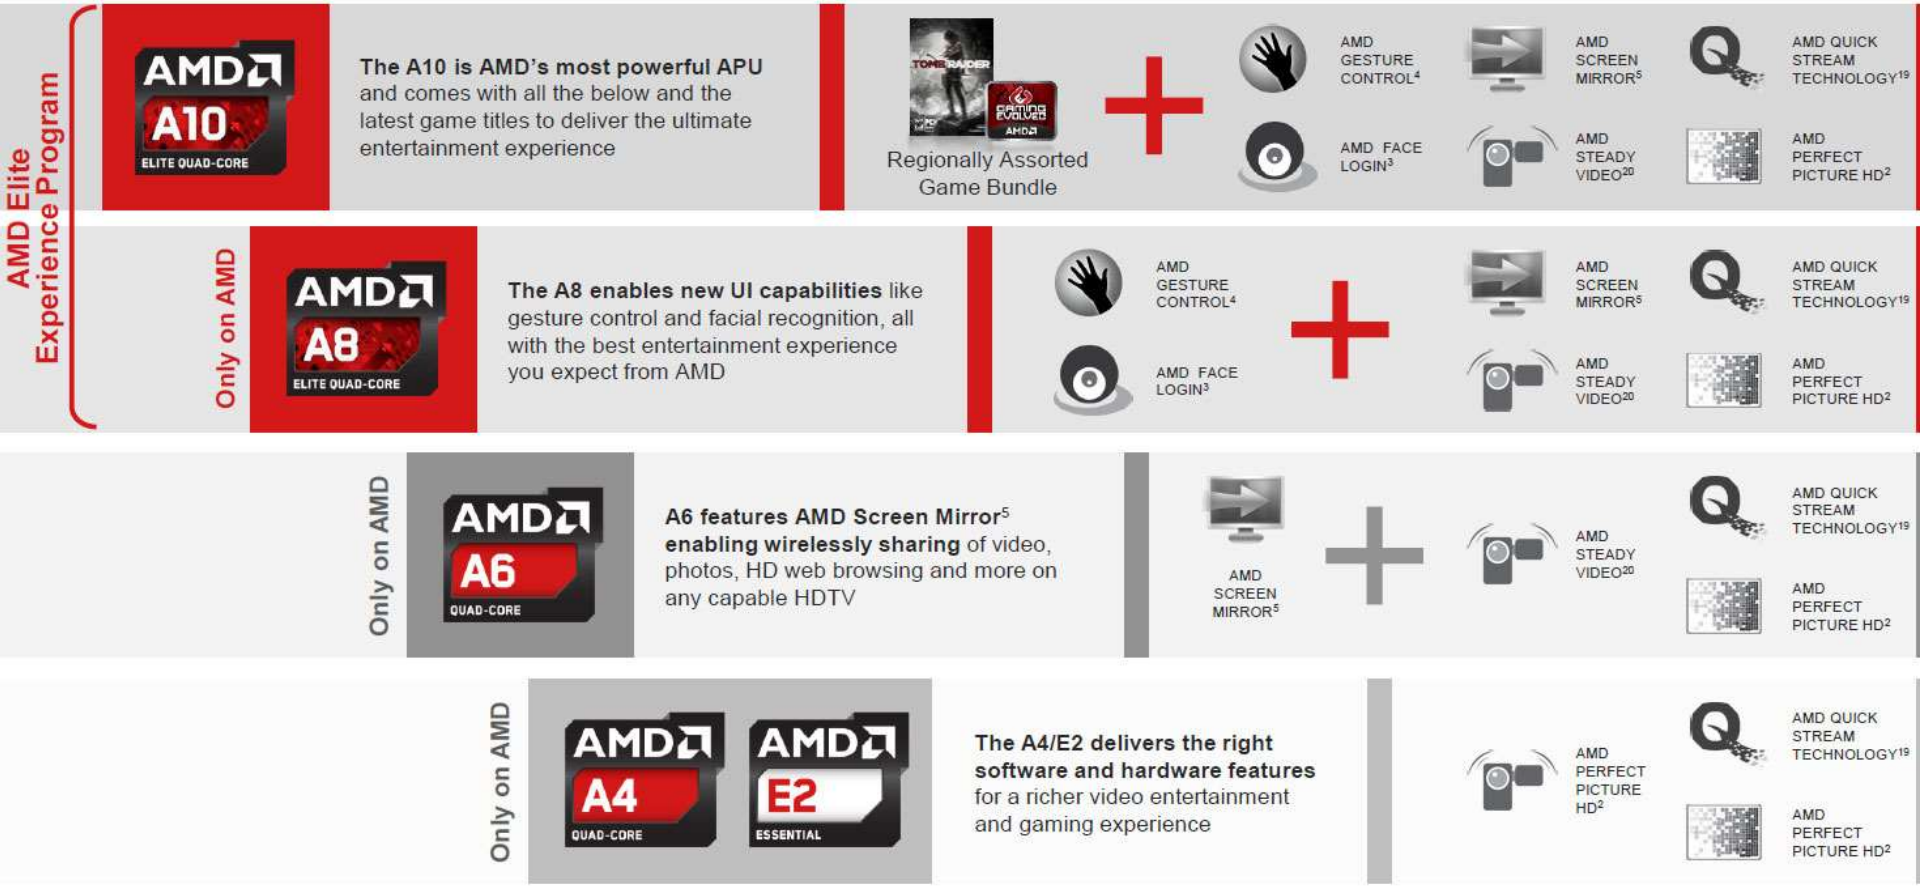

AMD's graphics performance figures of the Richland mobile APU line vs. Intel's Ivy Bridge-based mobile processors [83]

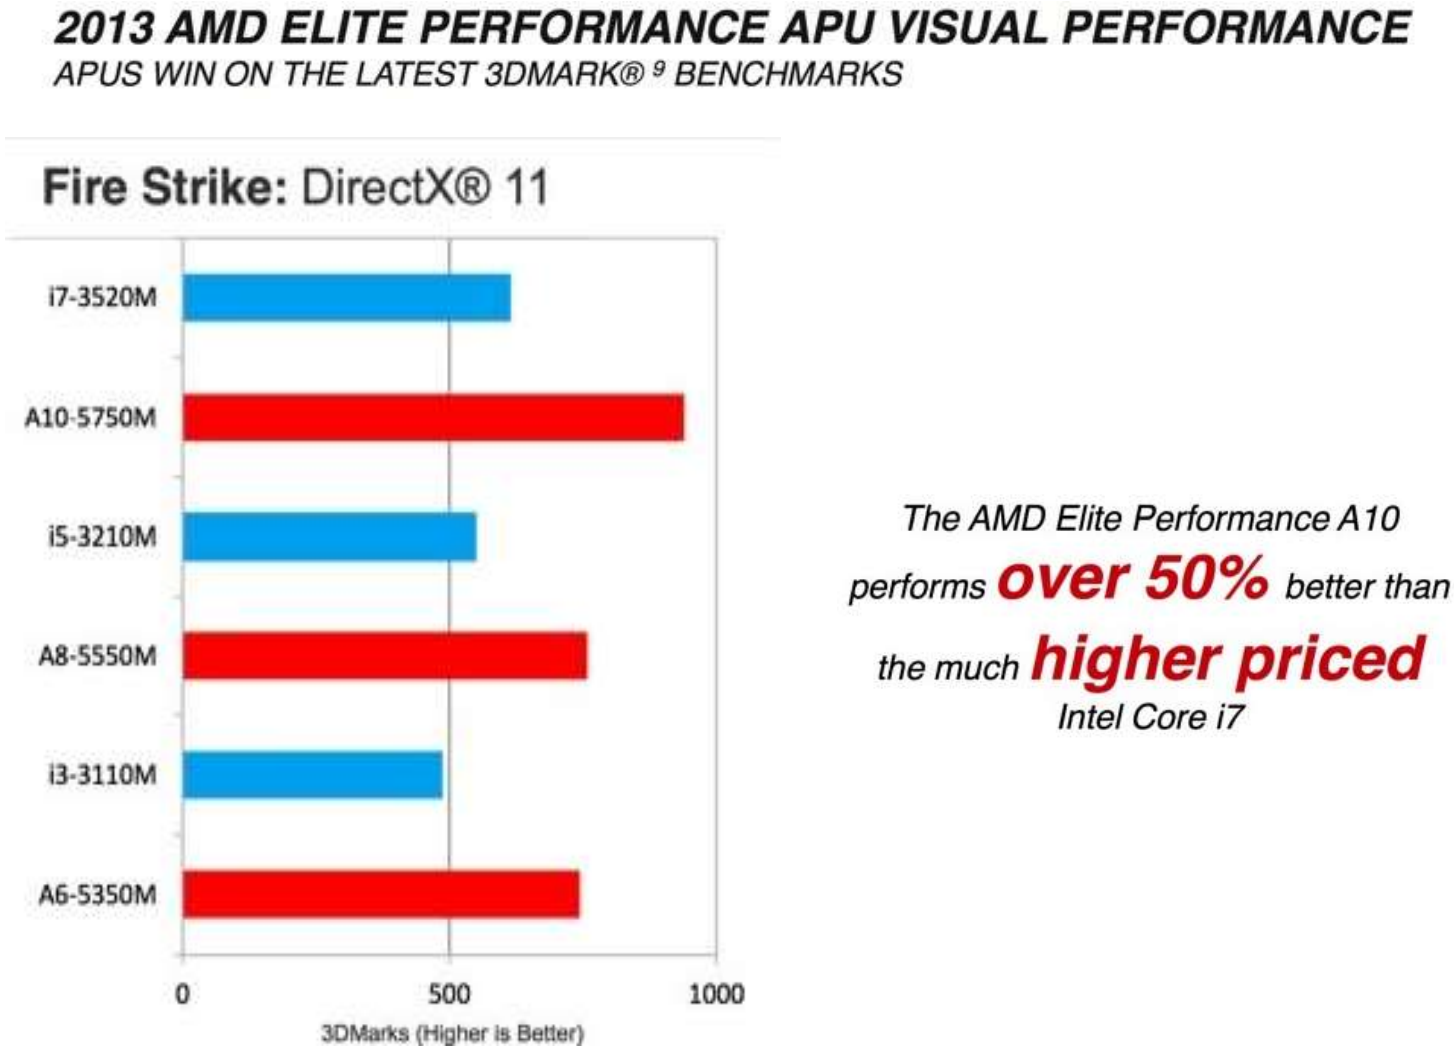

3.5.3 The Richland mobile APU line (11)

Remark

At its introduction in 6/2012 the Intel Core i7-3520M was the fastest dual-core, dual threads/core mobile processor for laptops, based on the Ivy Bridge architecture, with the following key parameters: [90]

Series	Intel Core i7
Codename	Ivy Bridge
Clock rate	2900 - 3600 MHz
Level 1 Cache	128 KB
Level 2 Cache	512 KB
Level 3 Cache	4096 KB
Number of Cores / Threads	2 / 4
Max. Power Consumption (TDP = Thermal Design Power)	35 Watt
Manufacturing Technology	22 nm
Max. Temperature	105 °C
Socket	BGA1023, PGA988
Features	HD Graphics 4000, DDR3-1600 Memory Controller, HyperThreading, AVX, Quick Sync, Virtualization
64 Bit	64 Bit support
Hardware virtualization	VT-x, VT-d
Starting price	$346 U.S.
Announcement date	06/03/2012

4. Third generation Steamroller-based (Family 15h Models 30h-3Fh) processor lines

- 4.1 Overview of the Steamroller-based processor lines
- 4.2 The Steamroller Compute Module
- 4.3 The Steamroller-based Kaveri desktop and mobile APU line

4.1 Overview of the Steamroller-based processor lines

4.1 Overview of the Steamroller-based processor lines [based on 1]

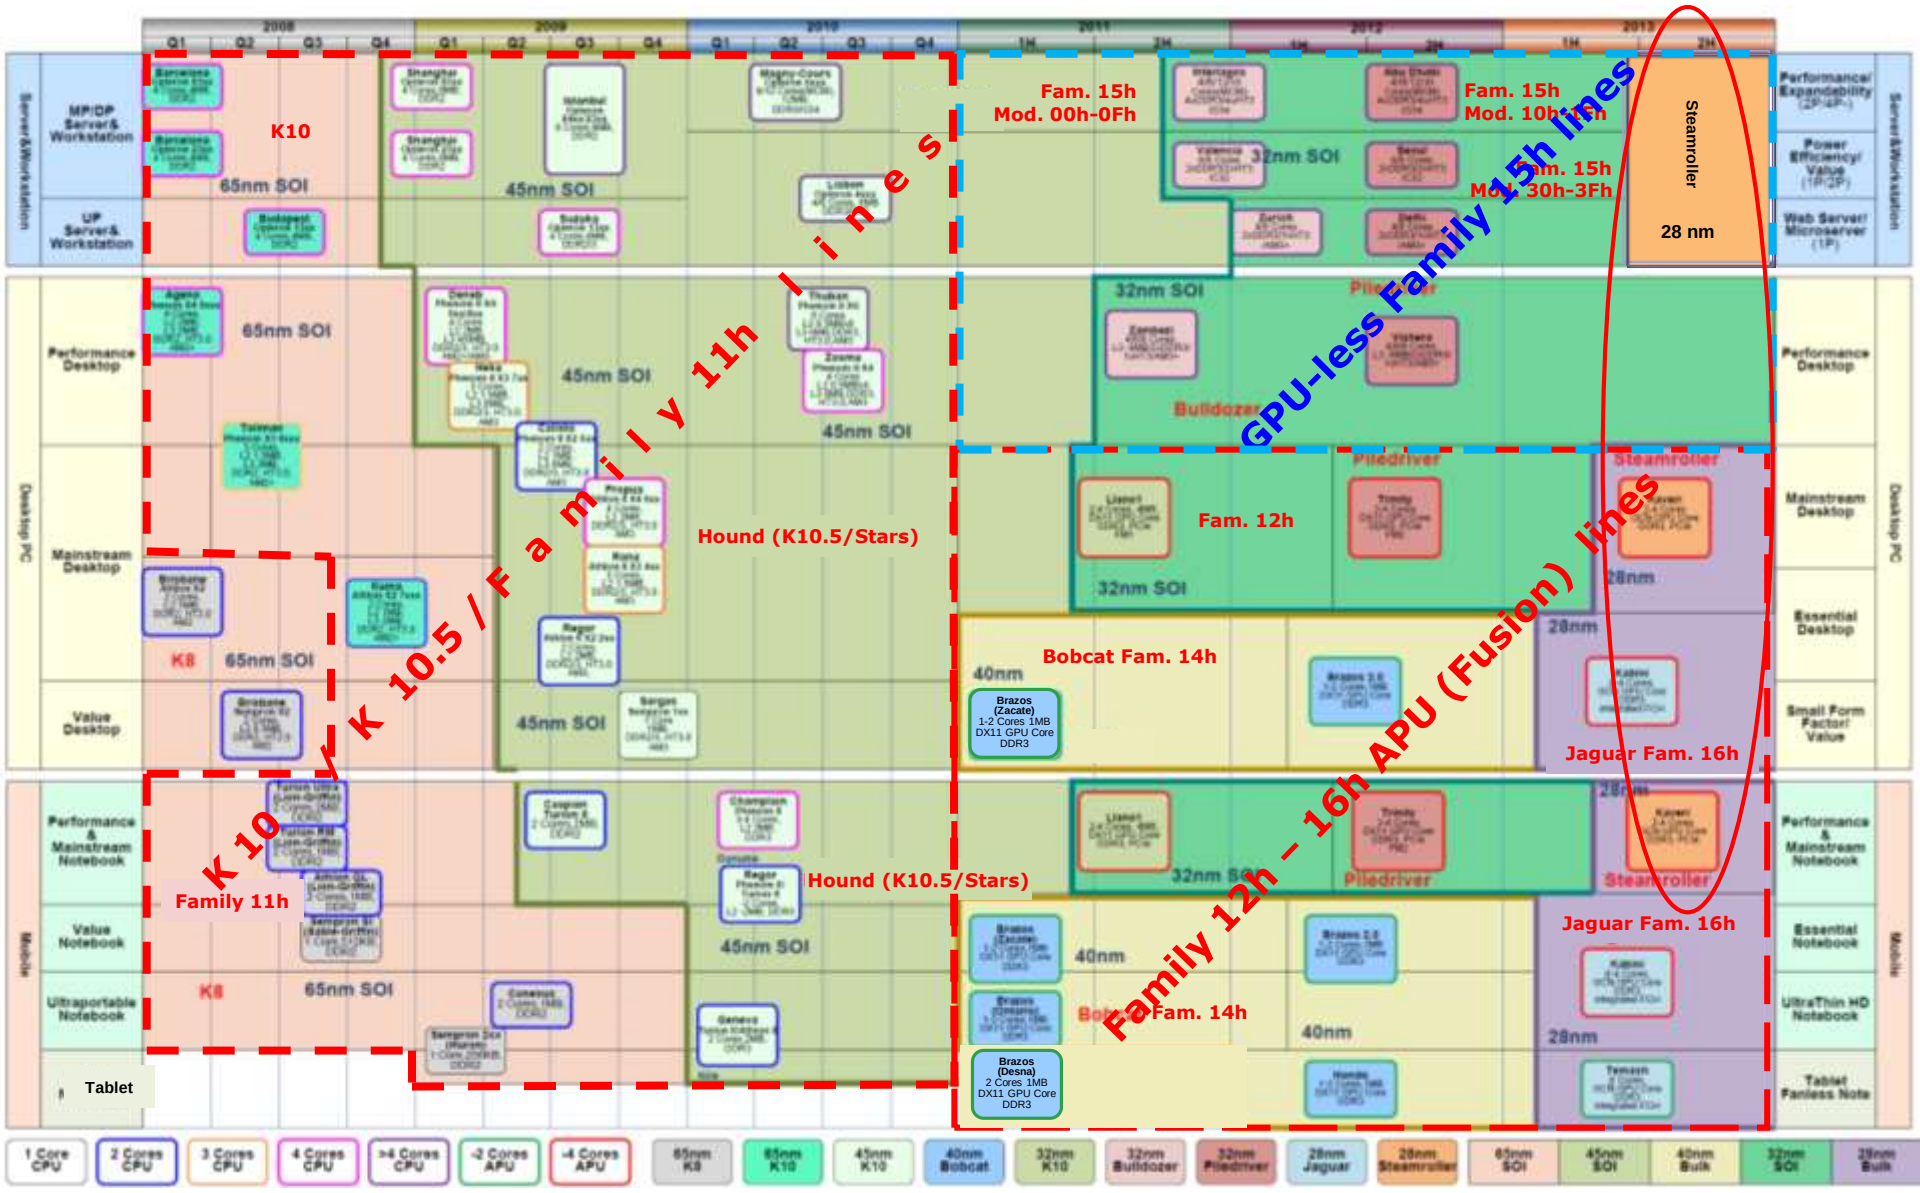

4.1 Overview of the Steamroller-based processor lines (2)

Overview of AMD's Steamroller-based processor lines

	Launched in	2011	2012	2013	2013	2015	2016
		Family 15h (00h-0Fh) (**Bulldozer**)	**Family 15h** (10h-1Fh) (**Piledriver**)	**Family 15h** (10h-1Fh) (**Piledriver v.2**)	**Family 15h** (30h-3Fh) (**Steamroller**)	**Family 15h** (60h-6Fh) (**Excavator v.1**)	**Family 15h** (77h-3Fh) (**Excavator v.2**)
Servers	**4P servers** (85-140 W)	Interlagos	Abu Dhabi				
	2P servers (85-140 W)	Valencia	Seoul				
	1P servers (85-140 W)	Zurich	Delhi				
Desktops	**High perf.** (~95-125 W)	Zambezi FX-Series	Vishera FX-Series				
	Mainstream (~65-95 W)		Trinity A10-A4	Richland A10/A8/A6/A4	Kaveri A10/A8		
Notebooks	**Mainstream** (~25-35 W)		Trinity A10/A8/A6M	Richland A10/A8/A6M	Kaveri FX/A10/A8P		Bristol Ridge FX/A12/A10P
	Ultra-thin (~10 – 15 W)		Trinity A10/A6M	Richland A10/A8/A6/A4M	A8 Pro/A8(B) A6 Pro/A6(B)	Carrizo FX/A10/A8P	Bristol Ridge FX/A12/A10P Stoney Ridge A9/A6
Tablets (~5 W)							

Overview of AMD's Steamroller-based desktop and mobile processor lines

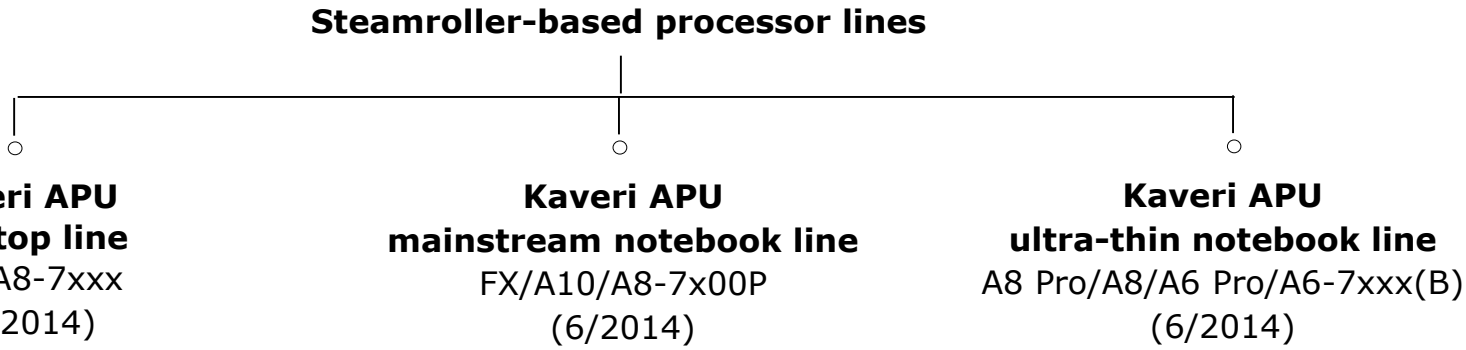

4.2 The Steamroller Compute Module

4.2 The Steamroller Compute Module

Planned introduction of the Steamroller compute module

While introducing their Piledriver-based high performance Zambezi DT line (1í/2011) AMD revealed their plan to introduce Steamroller-based compute modules in 2013, as shown below [44].

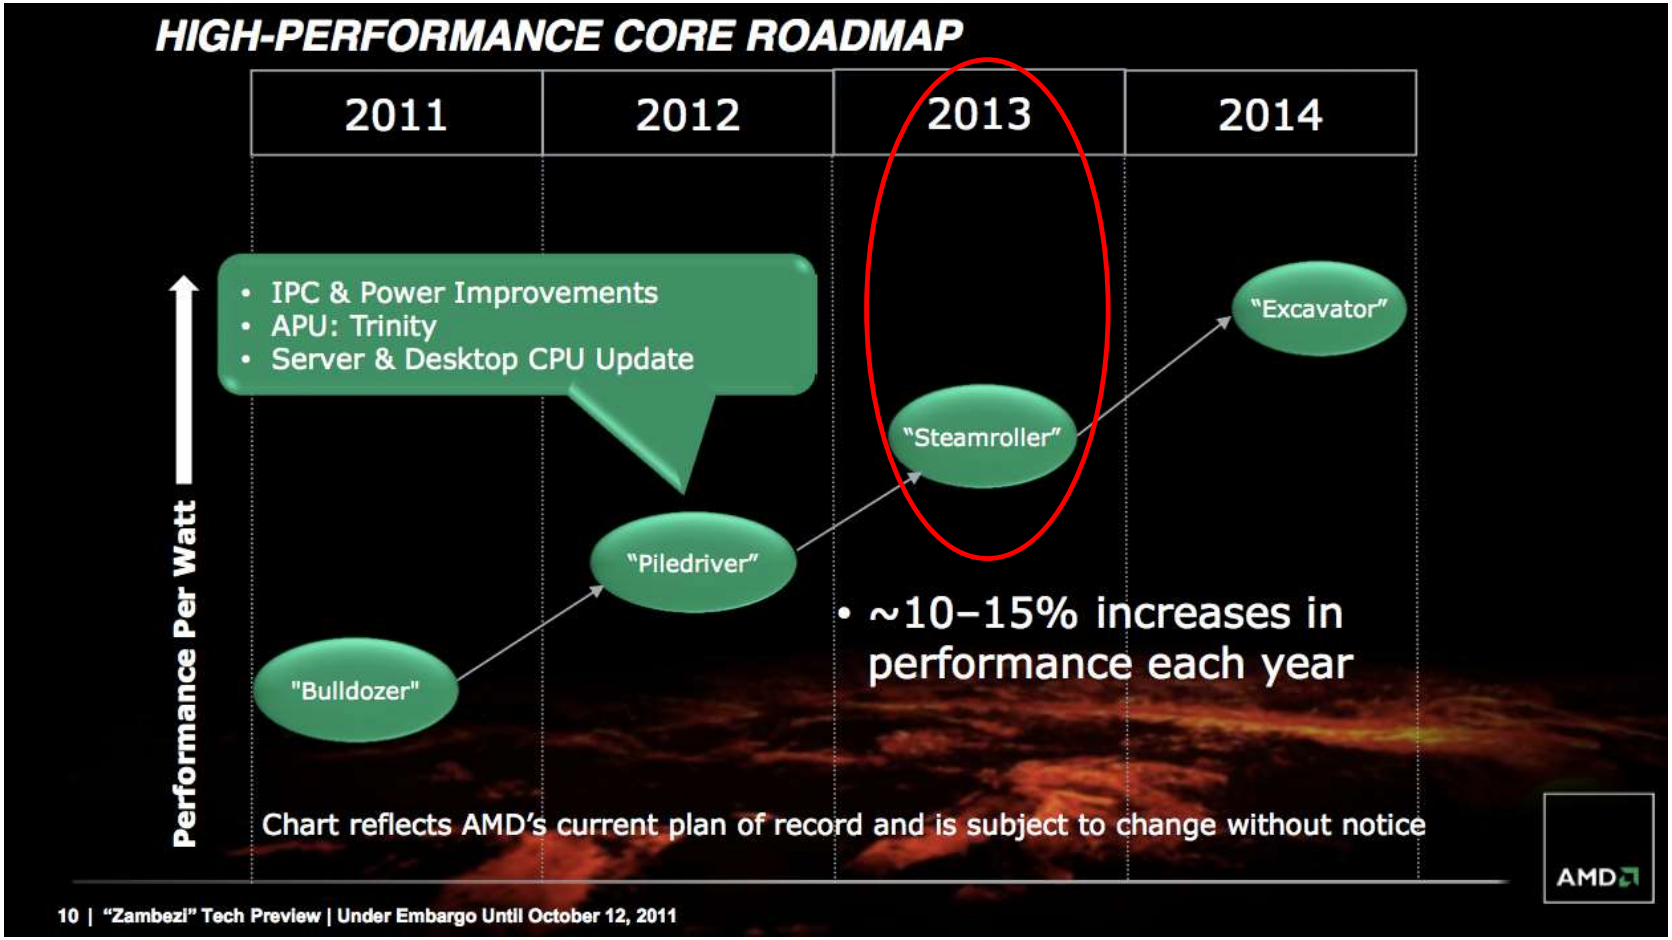

Preview of the Steamroller compute module (CM)

At the Hot Chips 2012 (8/2012) AMD's CTO (Chief Technical Officer) gave a preview of the CM of Steamroller, revealing some high-level details of the microarchitecture, as shown in the next Figures [45].

Block diagram of the Steamroller compute module [45]

AMD "STEAMROLLER" CORE

Multi-threaded microarchitecture

Expands computation efficiency

- Feed the cores faster
- Improve single-core execution
- Push on performance/watt

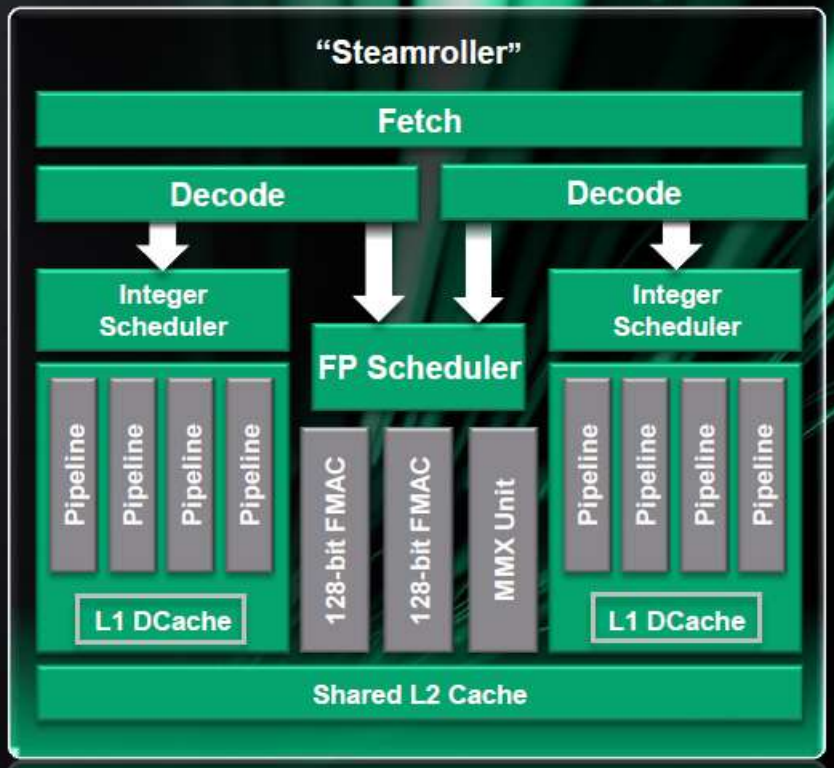

AMD

Improvements of the front-end part of the Steamroller compute module [45]

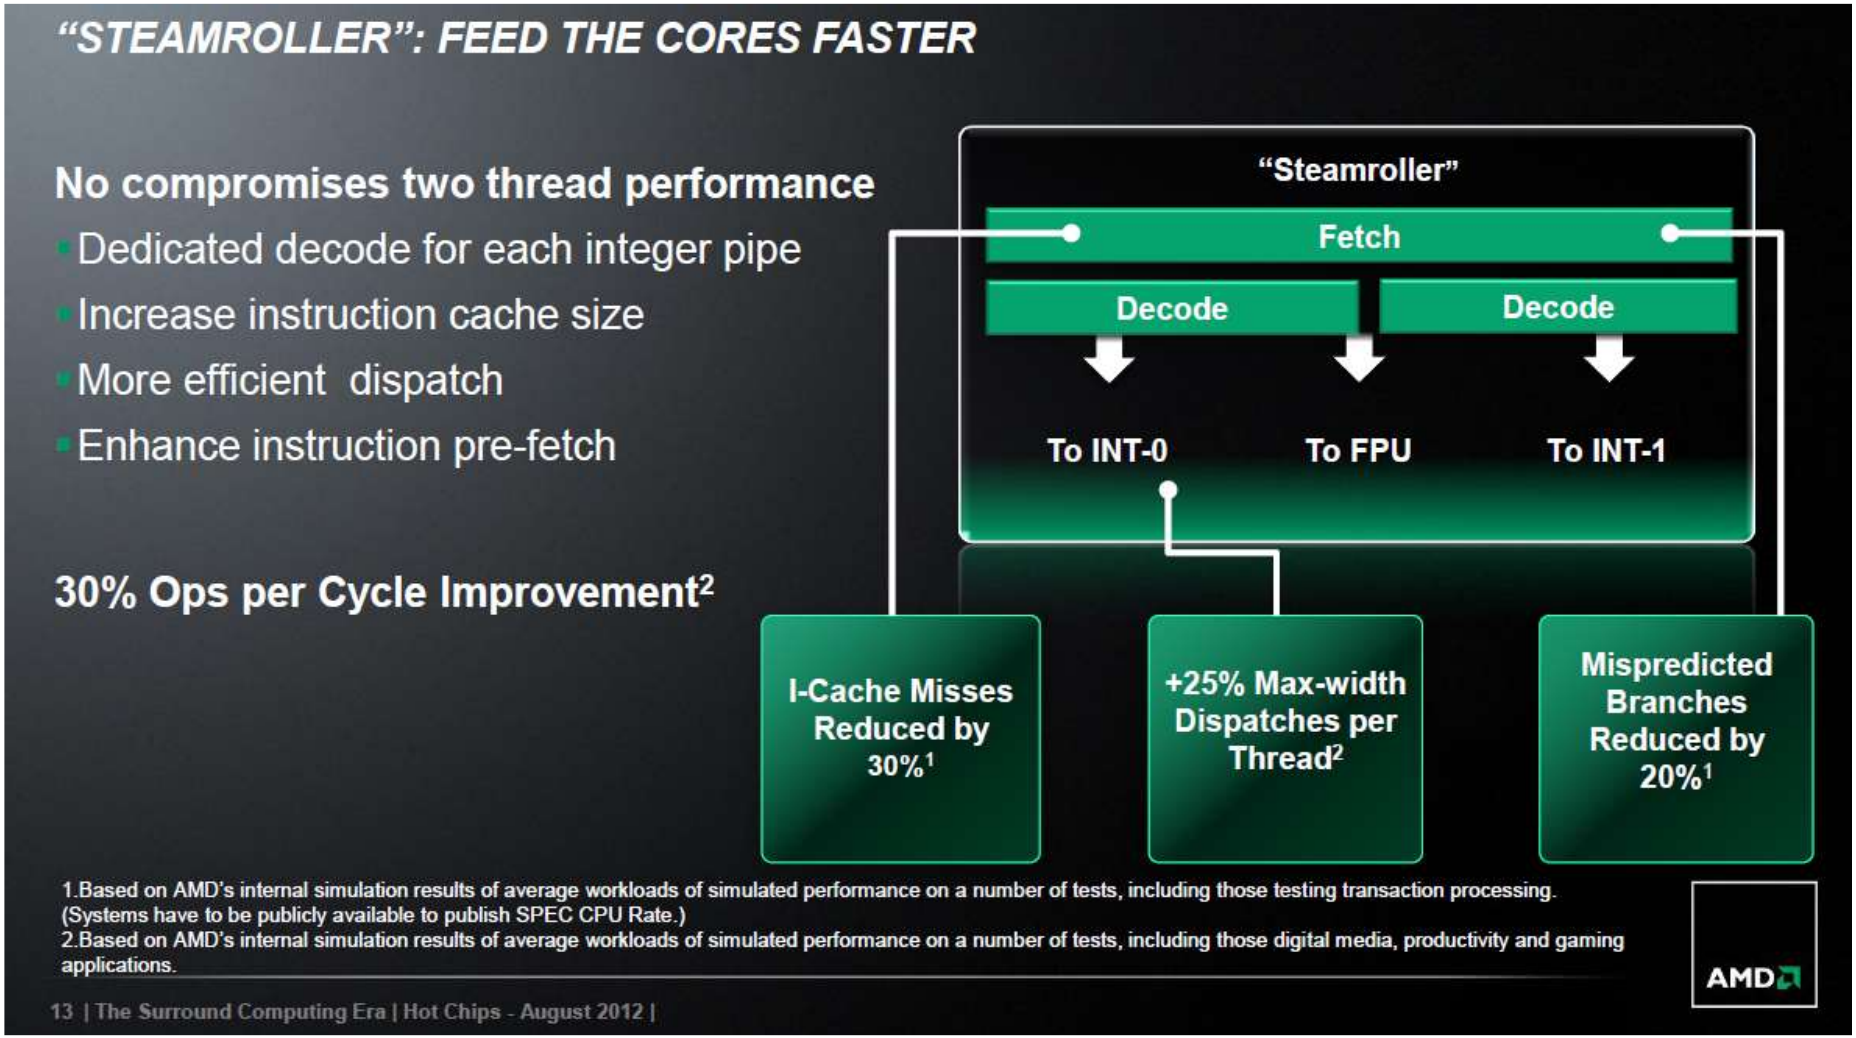

Improving integer scheduling, integer execution and reducing average load latency in the Steamroller compute module [45]

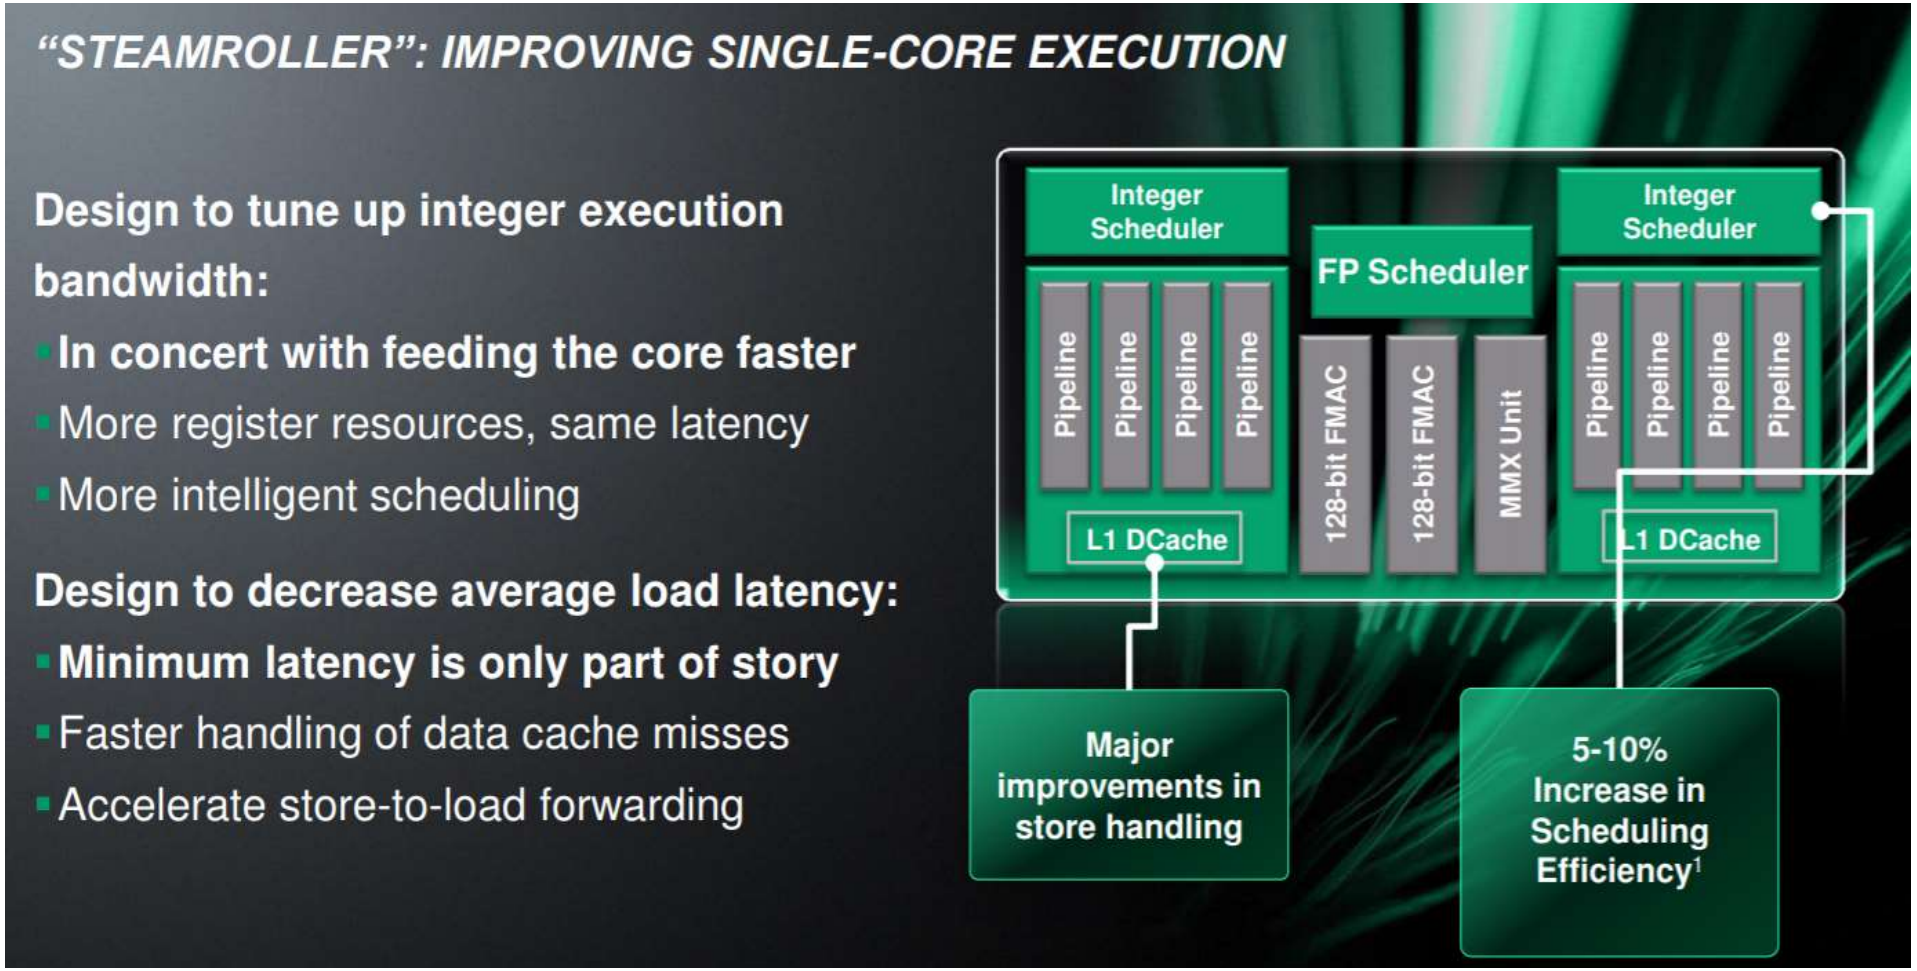

Improving the power efficiency (performance/Watt figure) of the Steamroller compute module [45]

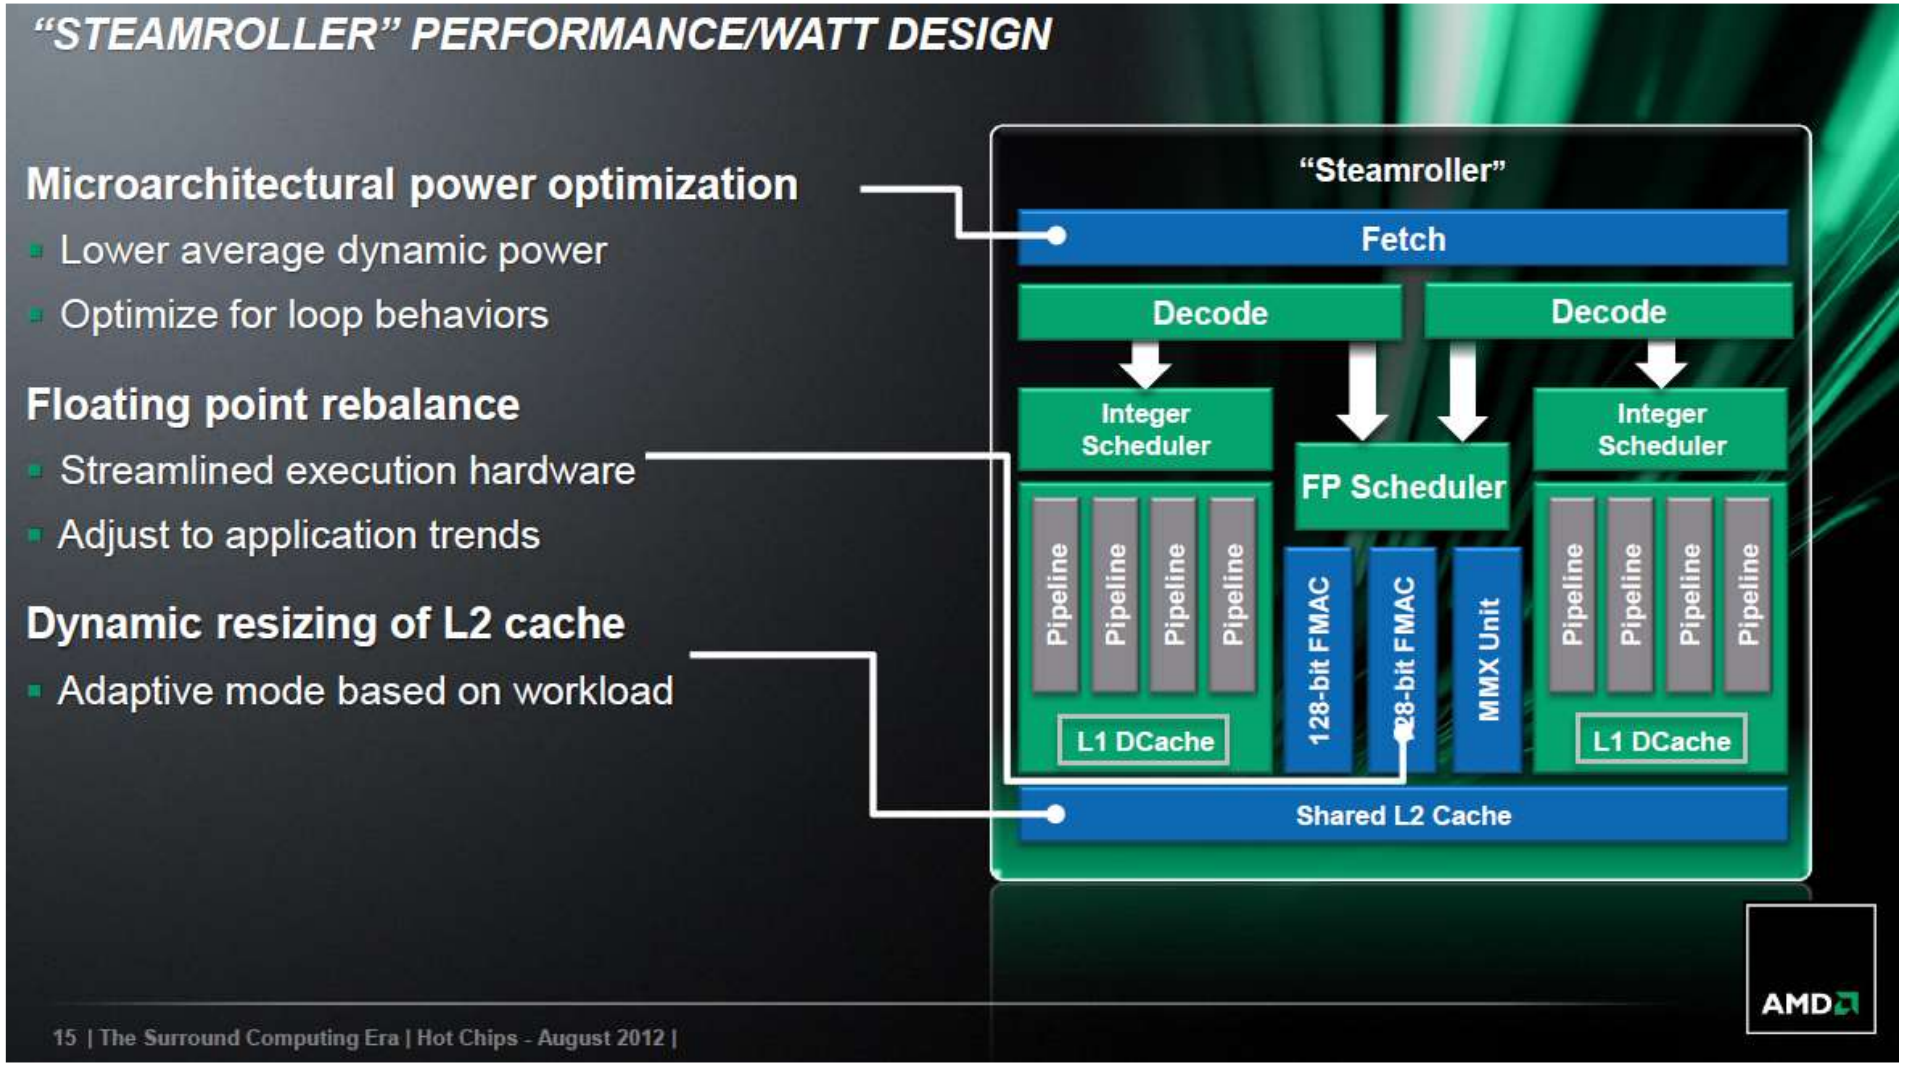

Comparing the block diagrams of three generations of the Family 15h Bulldozer design-1

A comparison of the block diagrams of subsequent three generations of the Family 15h Bulldozer CM design reveal that at the high level block diagram AMD did not made any noticeable change, except of introducing dedicated decoders in the Steamroller core, as shown below.

4.2 The Steamroller Compute Module (8)

Comparing the block diagrams of three generations of the Family 15h Bulldozer design

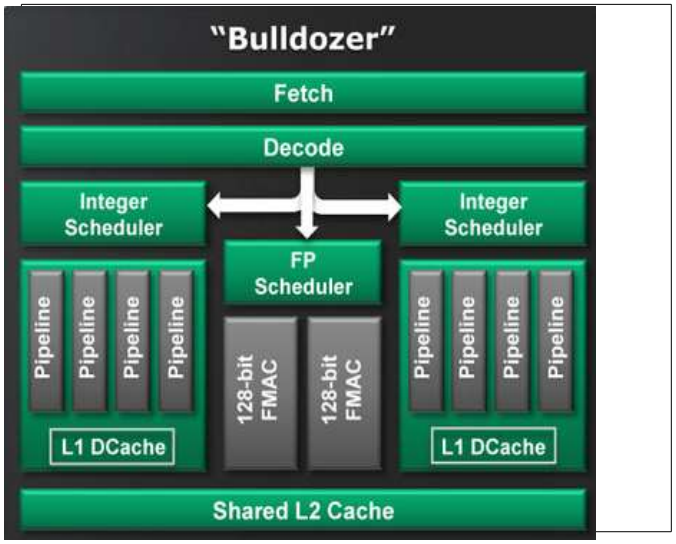

Bulldozer core [95]

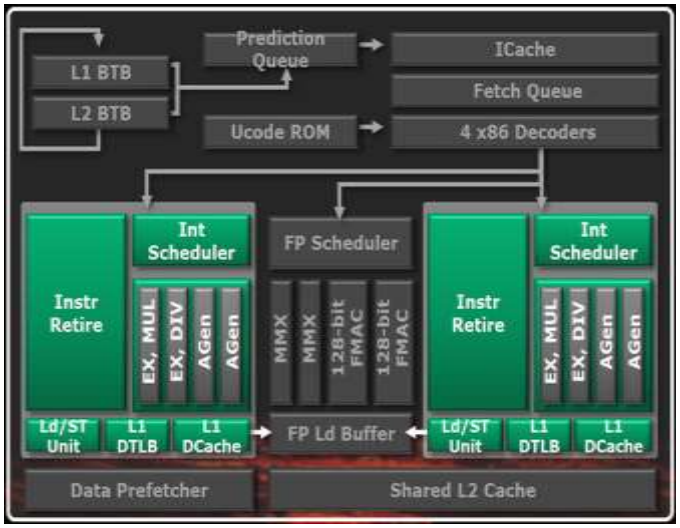

Piledriver core [47]

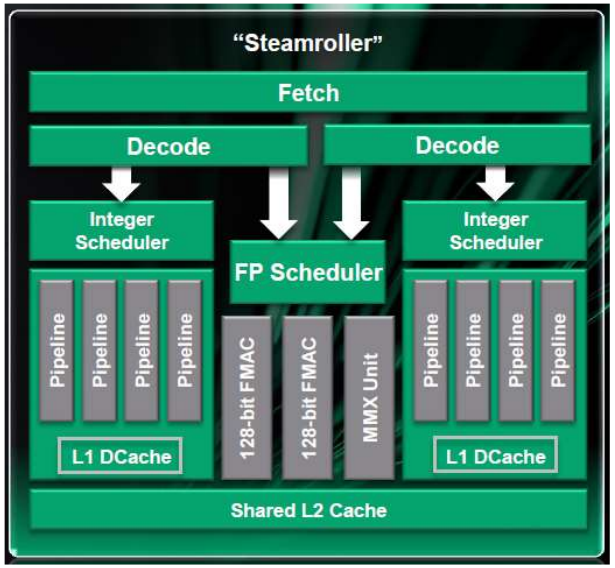

Steamroller core [45]

Improvements made in the microarchitecture of the Steamroller compute module -1

Although not noticeable in the high level block diagram, AMD made a vast number of improvements practically in all parts of the microarchitecture both in their Piledriver design over the Bulldozer CM, and then in their Steamroller design over the Piledriver CM.

As far as the efficiency of the microarchitecture is concerned, these improvements aimed at eliminating bottlenecks brought to light through extensive simulations while using a large number of relevant applications in order to increase IPC.

Major changes of the microarchitecture are revealed in the *"Preliminary BIOS and Kernel Developer's Guide for AMD Family 15h Models 30h-3Fh Processors"* .

From these changes, discussed in [48], we point out the following two:

- In case of integer micro-operations (Cops) increasing the dispatch bandwidth from 4 to 8 while dispatching up to 4 micro-operations per cycle to each core like in previous designs).
- Dispatching and retiring up to 2 stores per cycle instead of just one.

Improvements made in the microarchitecture of the Steamroller compute module -2

Nevertheless, increasing the dispatch bandwidth required further enhancements in the related units to avoid bottlenecks, including

- Increasing the L1 instruction cache size from 64 KB to 96 KB and changing its associativity from 2-way to 3-way
- Increasing the size of associated internal buffers, such as
  - Load queue (LDQ) size increased to 48, from 44.
  - Store queue (STQ) size increased to 32, from 24.
  - Increased L2 BTB size from 5K to 10K and from 8 to 16 banks.
  - Increased PFB (Prefetch Buffer) size from 8 to 16 entries; while the 8 additional entries can be used either for prefetching or as a loop buffer.

Improvements made in the microarchitecture of the Steamroller compute module -3

In addition, AMD introduced a large number of further enhancements to increase IPC, as listed below [48].

- Optimizations of certain features, including
  - Reducing the number of FP pipeline stages from 4 to 3.
  - Optimizing store to load forwarding.
  - Improved loop prediction.
  - Accelerate SYSCALL/SYSRET.
  - Increased snoop tag throughput.
- Enhancing the microarchitecture by
  - Virtualized interrupt controller.
  - Support of the XSA/EOPT instruction.

Remark

In addition to significantly increasing the efficiency of the Steamroller compute unit,
in the overall architecture of the related APU processors AMD made also substantial changes,
as will be briefly discussed in the Section introducing the Kaveri APU (Section 4.3) [48].

4.3 The Kaveri desktop and mobile APU lines

4.3 The Kaveri desktop and mobile APU lines

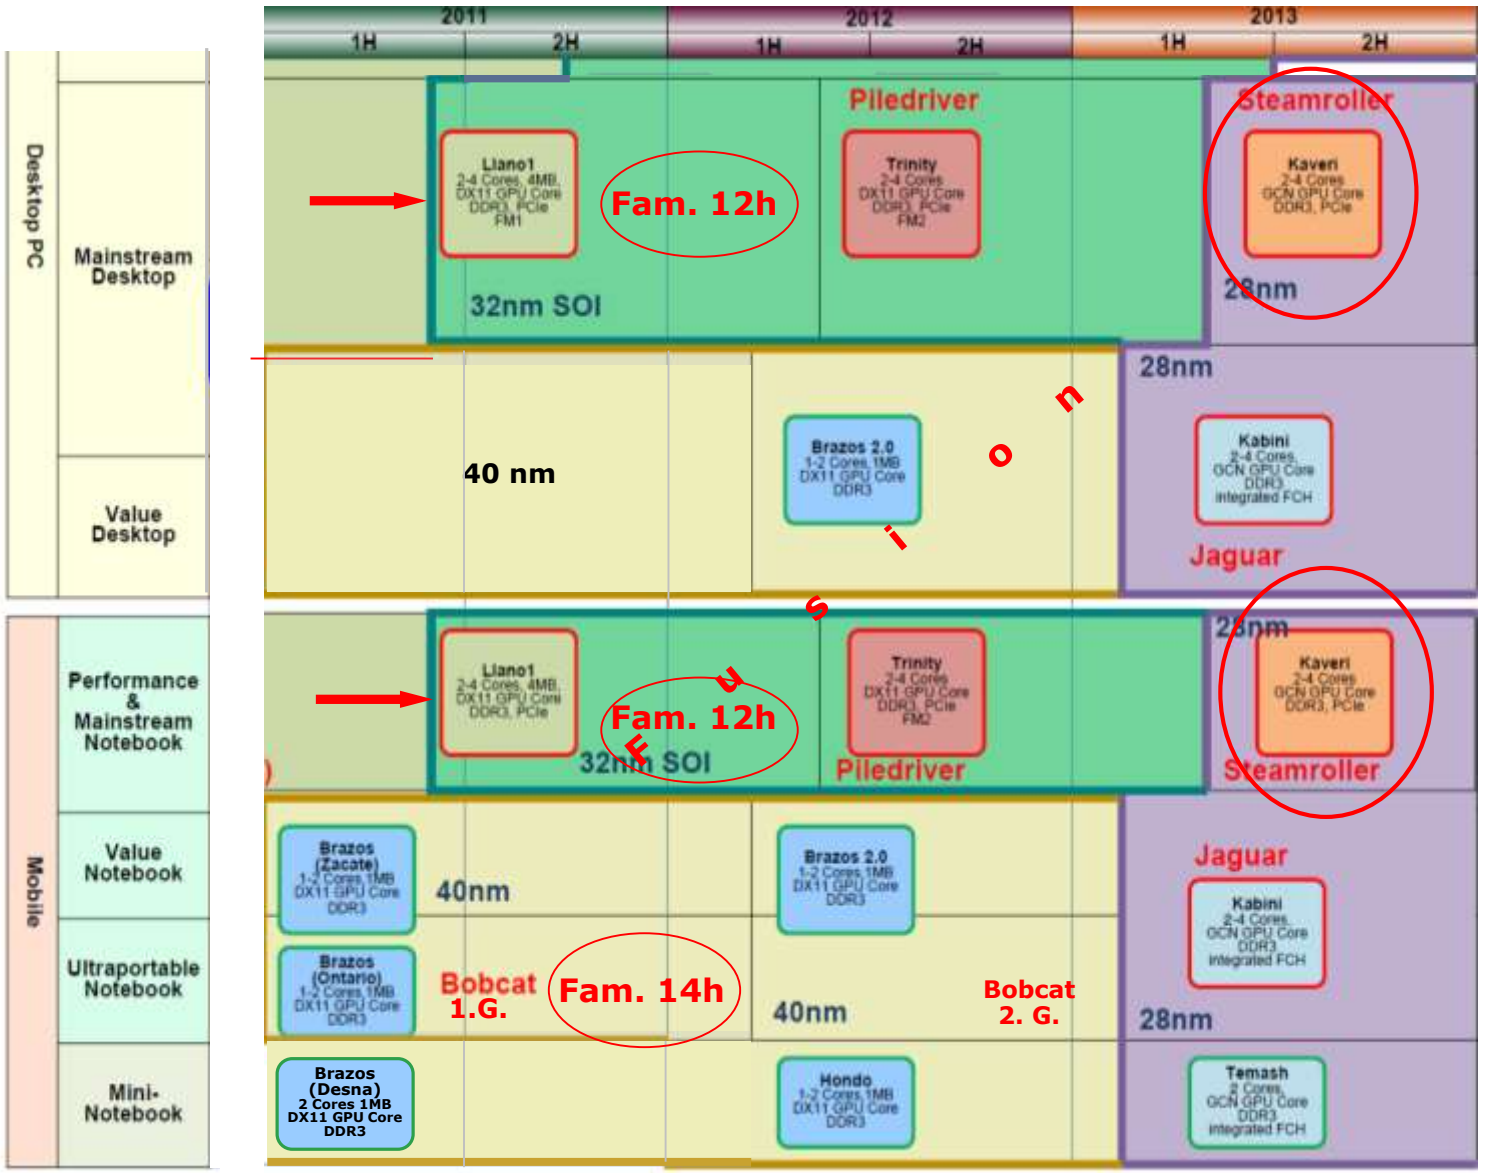

Positioning the Kaveri APU line as mainstream desktop line [51]

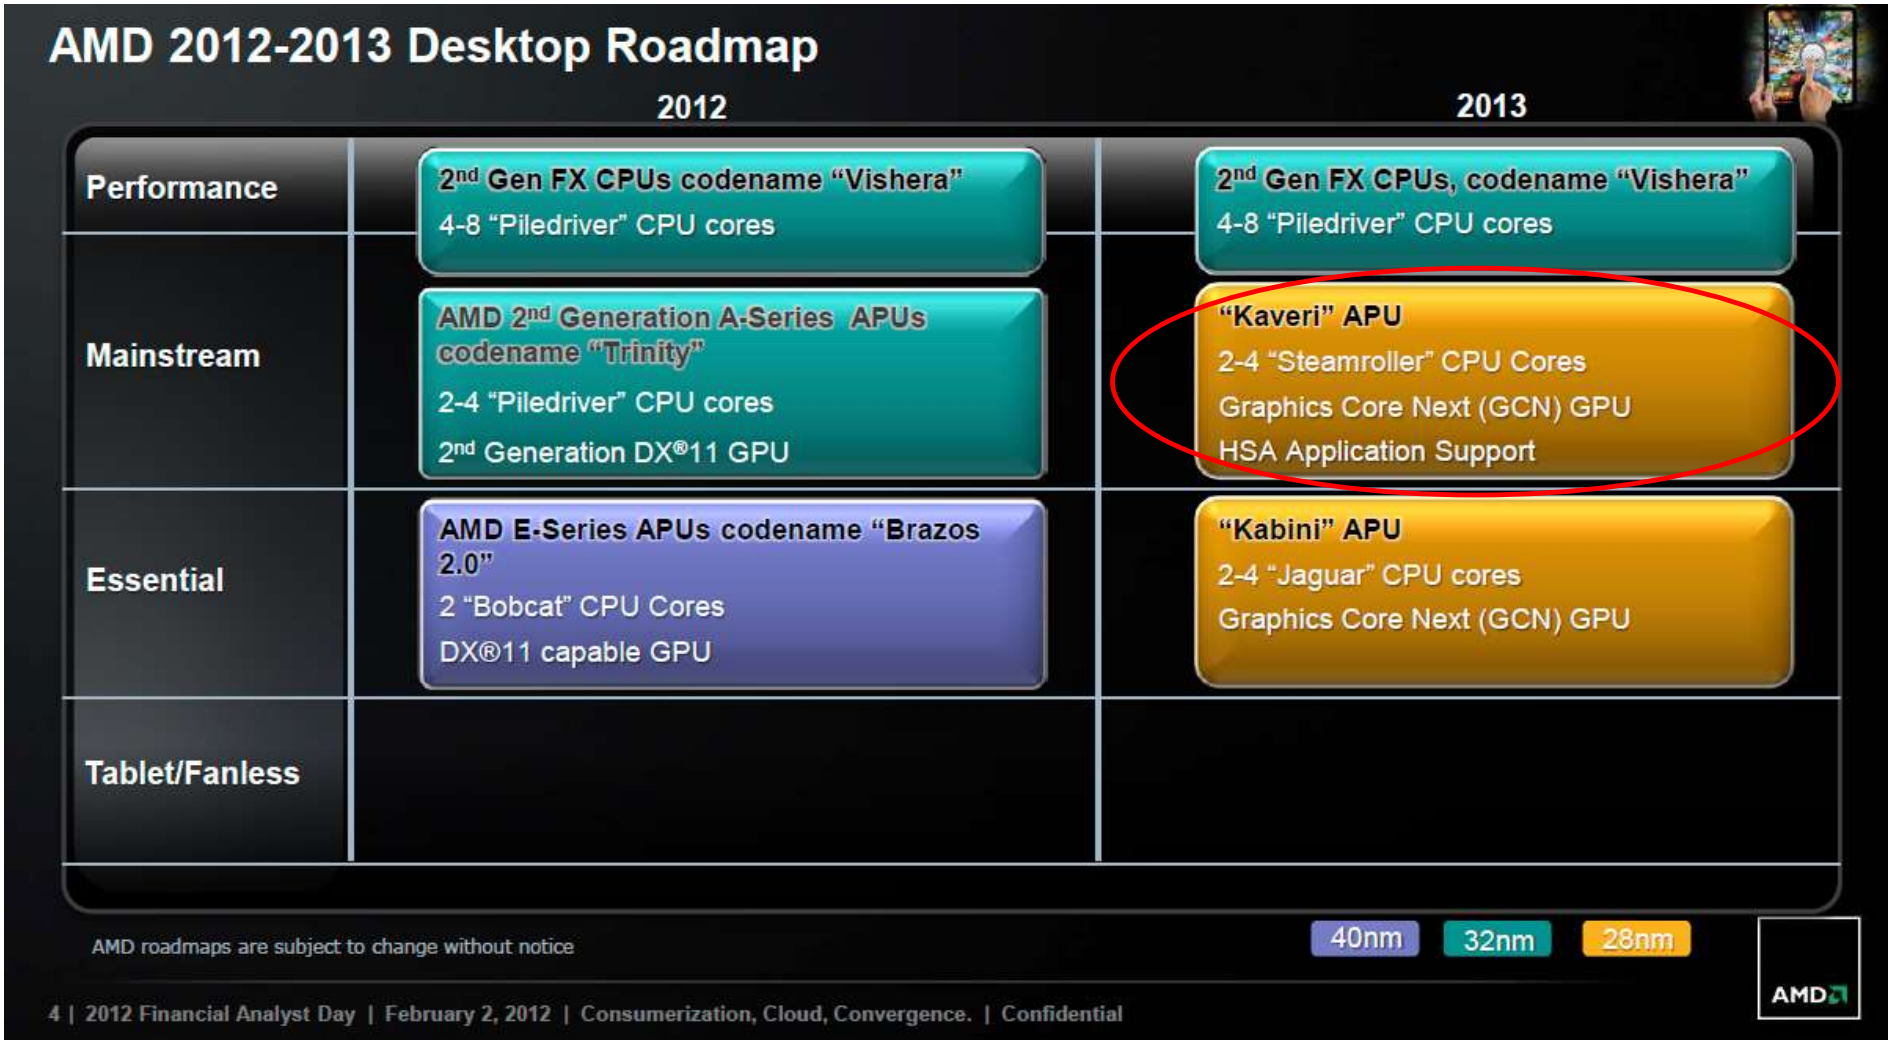

Positioning the Kaveri APU line as performance/mainstream mobile line [51]

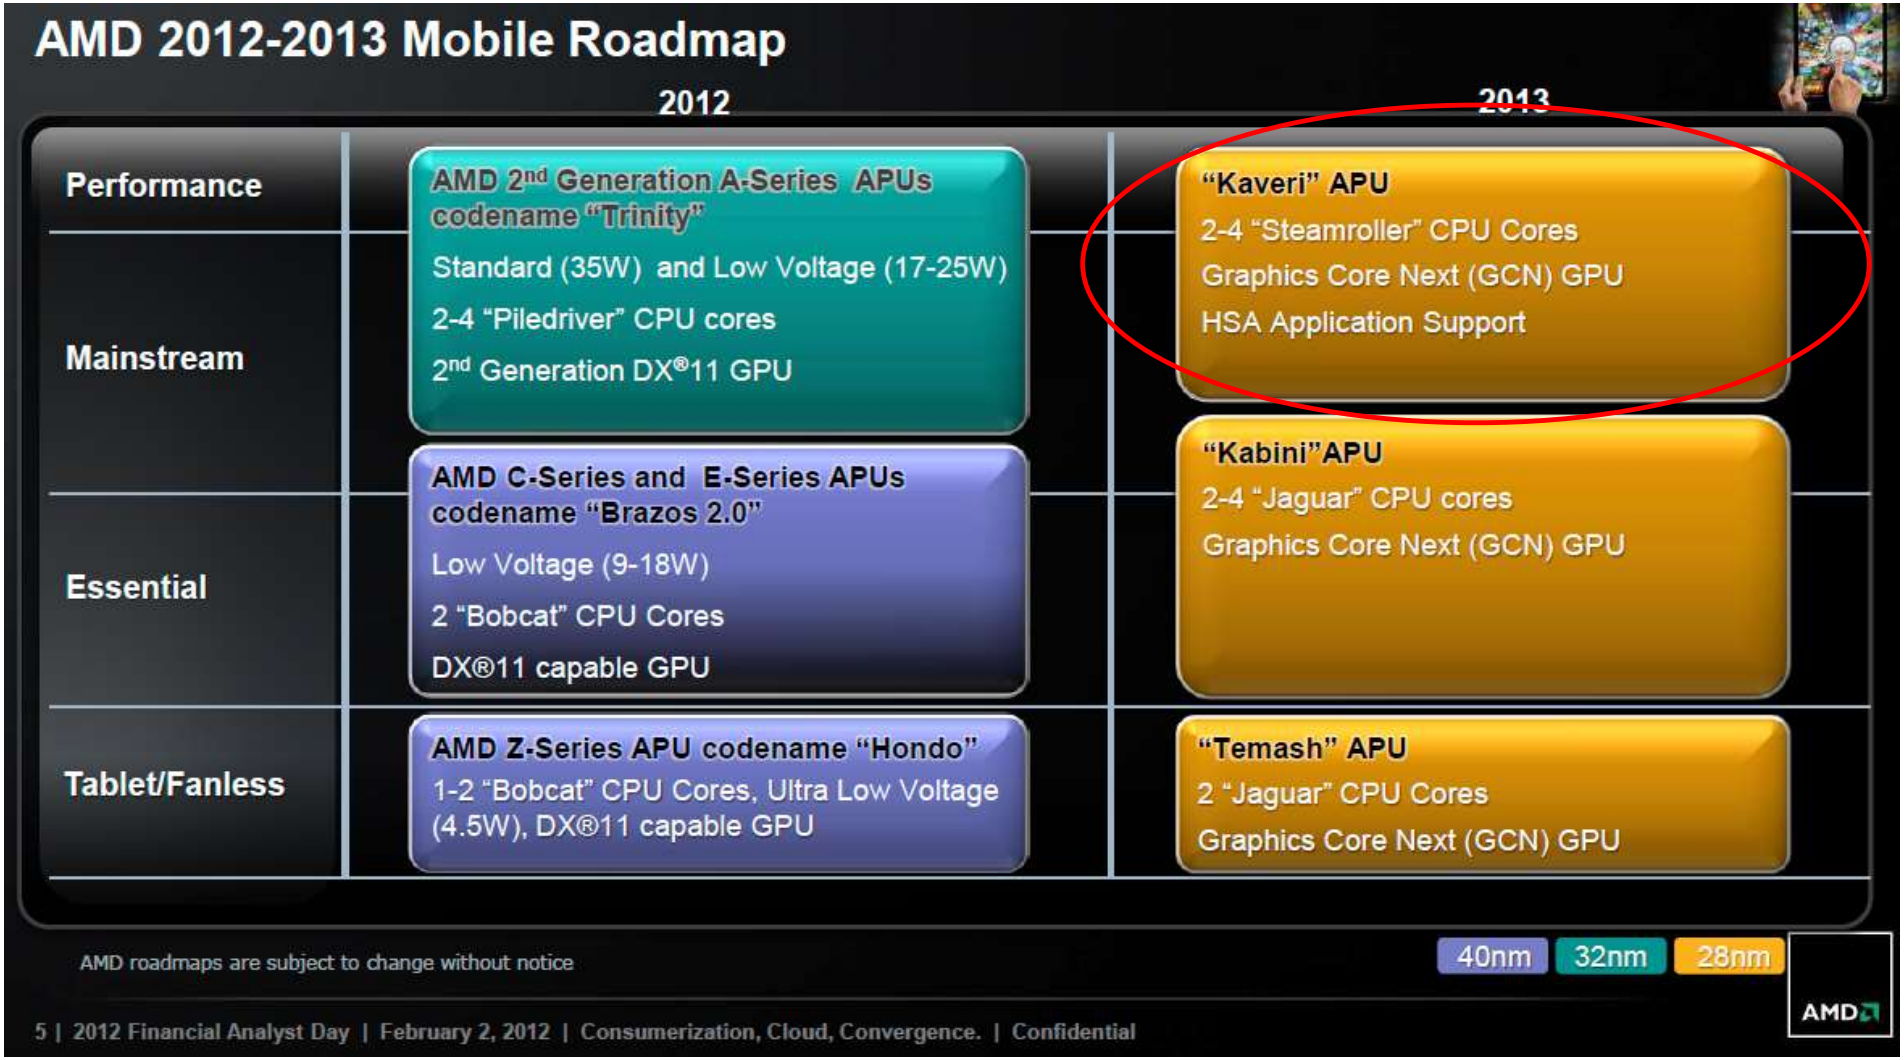

Revised positioning of the the Kaveri APU line [52]

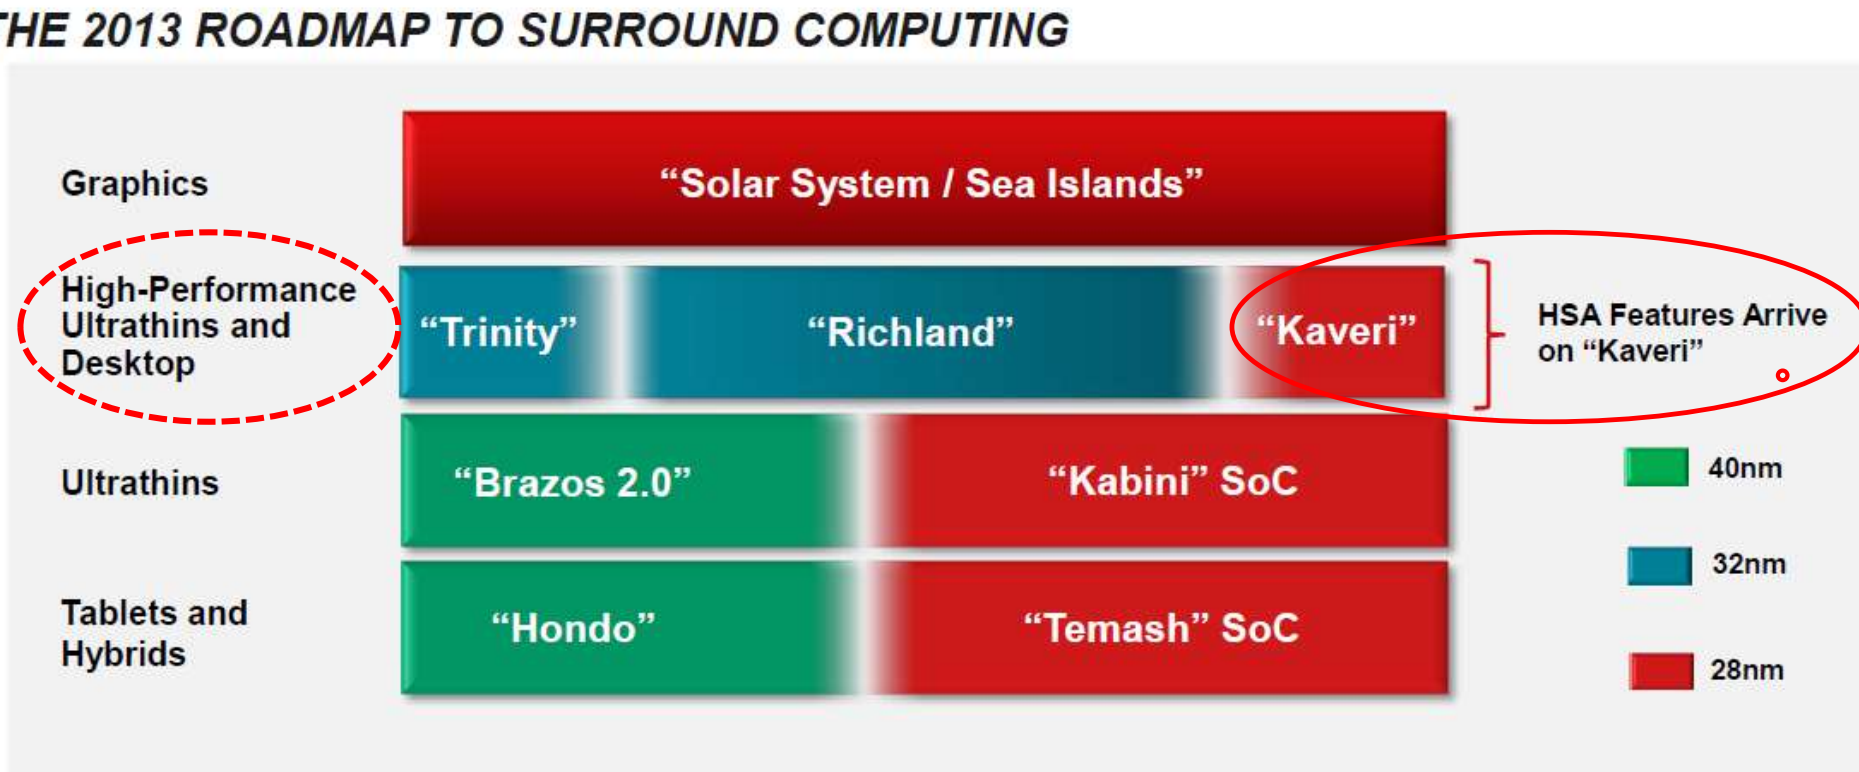

Key features of the Kaveri die [96]

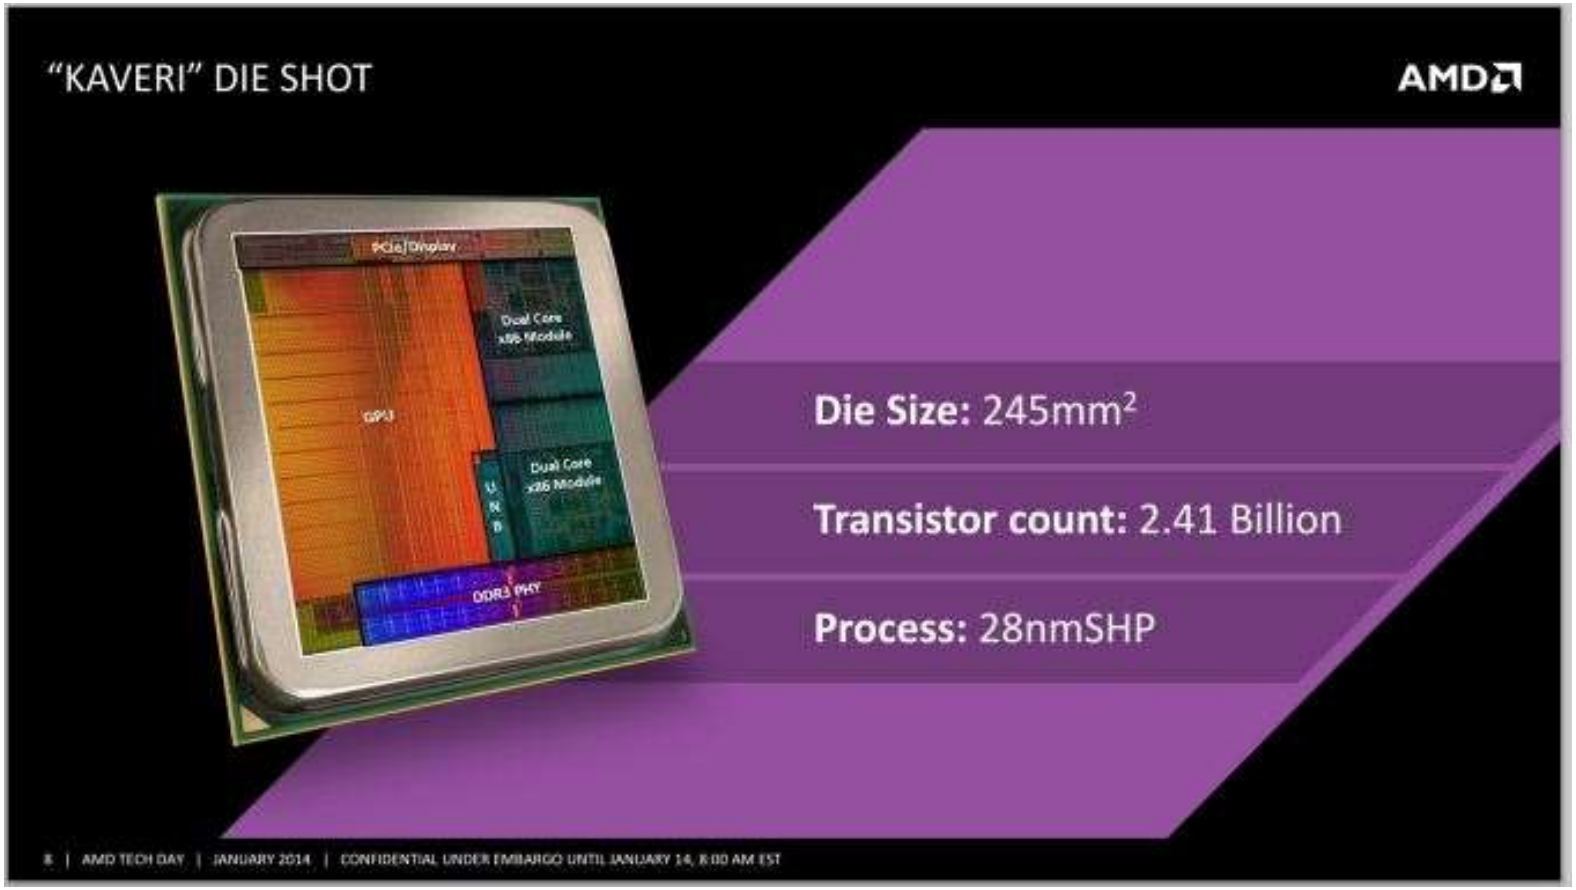

Die shot of a dual-module Kaveri [96]

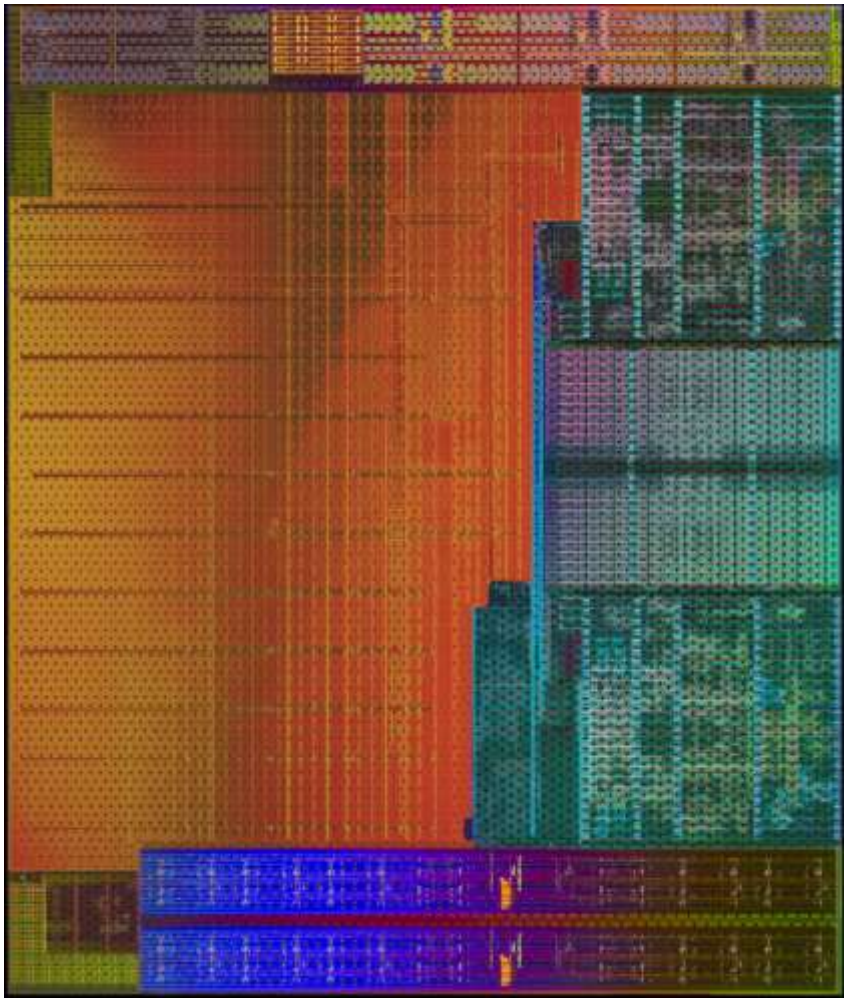

Overview of the Steamroller-based processor lines

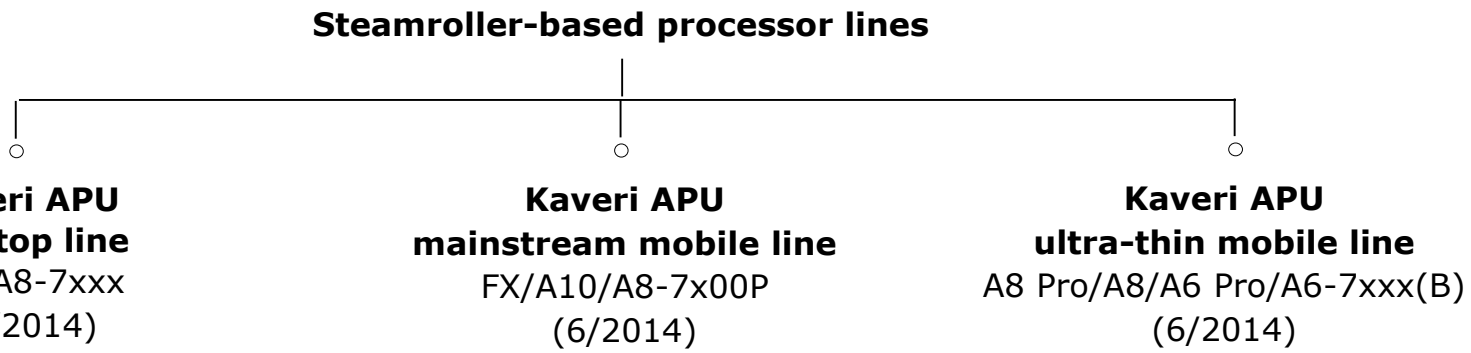

Product positioning of first introduced (01/2014) Kaveri models [96]

4.3 The Kaveri desktop and mobile APU line (9)

Main components of Kaveri APUs

- 1-2 Steamroller Compute Modules (2-4 Steamroller CPU cores)
- GCN (Graphics Core Next) GPU (Presumably models of the Sea Islands (HD8000) family)

Main features of the Kaveri DT processors [97]

Model	Stepping	CPU				GPU			Memory support	TDP	Release date
		Modules (cores)	Clock	Turbo	L2 cache	Model	Cores	Clock			
A8-7600	KV-A1	2(4)	3.1 GHz	3.8 GHz	2× 2 MB	R7	6	720 MHz	DDR3-2133	65 W	Jul 31, 2014
A8 PRO-7600B											
A8-7650K			3.3 GHz							95 W	Jan 7, 2015
A8-7670K	GV-A1		3.6 GHz	3.9 GHz				757 MHz			Jul 20, 2015
A8 PRO-8650B			3.2 GHz	3.9 GHz						65 W	Sep 29, 2015
A10-7700K	KV-A1		3.4 GHz	3.8 GHz				720 MHz		95 W	Jan 14, 2014
A10-7800			3.5 GHz	3.9 GHz			8	720 MHz		65 W	Jul 31, 2014
A10 PRO-7800B											
A10-7850K			3.7 GHz	4.0 GHz						95 W	Jan 14, 2014
A10 PRO-7850B											Jul 31, 2014
A10-7860K	GV-A1		3.6 GHz					757 MHz		65 W	Feb 2, 2016
A10-7870K			3.9 GHz	4.1 GHz				866 MHz		95 W	May 28, 2015
A10-7890K			4.1 GHz	4.3 GHz							Mar 1, 2016
A10 PRO-8750B			3.6 GHz	4.0 GHz				757 MHz		65 W	Sep 29, 2015
A10 PRO-8850B			3.9 GHz	4.1 GHz				800 MHz		95 W	

4.3 The Kaveri desktop and mobile APU line (10b)

Main features of the Kaveri mainstream and ultra-thin mobile processors [98]

Model	Cores / Threads	Frequency	Turbo frequency	L2 cache	TDP
AMD A8-Series for Notebooks family, BGA (FP3)					
A8-7100	4 / 4	1.8 GHz	3 GHz	4 MB	19W
A8 Pro-7150B	4 / 4	1.9 GHz	3.2 GHz	4 MB	19W
A8-7200P	4 / 4	2.4 GHz	3.3 GHz	4 MB	35W
Other families, Steamroller micro-architecture, BGA (FP3)					
A6 Pro-7050B	2 / 2	2.2 GHz	3 GHz	1 MB	17W
A6-7000	2 / 2	2.2 GHz	3 GHz	1 MB	17W
A10-7300	4 / 4	1.9 GHz	3.2 GHz	4 MB	19W
A10 Pro-7350B	4 / 4	2.1 GHz	3.3 GHz	4 MB	19W
FX-7500	4 / 4	2.1 GHz	3.3 GHz	4 MB	19W
A10-7400P	4 / 4	2.5 GHz	3.4 GHz	4 MB	35W
FX-7600P	4 / 4	2.7 GHz	3.6 GHz	4 MB	35W

Mainstream mobile: 35 W Ultra-portable mobile: 17/19 W

Main innovations introduced along with the Kaveri line

a) Support of HSA (Heterogeneous System Architecture)
b) AMD's dual graphics (called also hybrid graphics)
c) Use of GPU cores as compute cores

a) Support of HSA (Heterogeneous System Architecture)

- Aim: Architectural integration of the CPU and the GPU cores in the Kaveri APU line
- Main features
  - Unified address space,
  - The GPU uses pageable system memory via CPU pointers and
  - there exist a fully consistent memory between CPU and GPU, as indicated in the next Figure.

Remarks

HSA standards are maintained by the non-profit standardization body, called HSA Foundation, established in 2012 by leading processor vendors, like AMD, ARM, Imagination Technologies, MediaTech, Qualcomm, Samsung (nevertheless, Intel does not take part in it).

Aim of the HAS standards is to dramatically easier to program heterogeneous computing devices.

HSA develops royalty-free standards and open-source software.

The HSA standards comprise three specifications:

- HSA Platform System Architecture Specification
- HSA Programmer Reference Manual Specification
- HSA Runtime Specification

At the time being there are three versions of the specifications, versions 1.0, 1.1 and 1.2.

State of the art of supporting shared memory for CPU and GPU before HAS [96]

“The terms "shared memory" or "unified memory" are actually thrown about quite frequently in the industry and can mean different things in different contexts. We examine the current state of art across platform distributors:
NVIDIA has introduced "unified memory" in CUDA. However, on current chips, it is a software-based solution that is more of a convenience for software developers and hidden behind APIs for ease of use. The price of data transfer still needs to be paid in terms of performance, and NVIDIA's tools merely hide some of the software complexity. However, NVIDIA is expected to offer true shared memory in the Maxwell generation, which will likely be integrated into the successor of Tegra K1 in 2015 or 2016.
AMD: AMD touts "zero copy" on Llano and Trinity for OpenCL programs. However, in most cases, this only provides a fast way to copy data from CPU to GPU and the ability to read data back from GPU in some limited cases. In practice, the zero copy feature has limited uses due to various constraints such as high initialization cost. For most use cases, you will end up copying data between CPU and GPU.
Intel: Intel provides some support for shared memory today on the Gen7 graphics in Ivy Bridge and Haswell exposed through OpenCL and DirectX. Intel's CPU/GPU integration is actually more impressive than Llano or Trinity from the perspective of memory sharing. However, sharing is still limited to some simple cases as it is missing pointer sharing, demand-based paging and true coherence offered in HSA and thus the integration is far behind Kaveri. I am expecting better support in Broadwell and Skylake. Intel's socketed Knights Landing (future Xeon Phi) product may also enable heterogeneous systems where both CPU and accelerator access the same memory, which might be the way forward for discrete GPUs as well (if possible).
Others: Companies like ARM, Imagination Technologies, Samsung and Qualcomm are also HSA Foundation members and probably working on similar solutions. Mali T600 and T700 GPUs expose some ability of sharing GPU buffers between CPU and GPU through OpenCL 1.1. However, I don't think we will see a full HSA stack from vendors other than AMD in the near future.
As of today, HSA model implemented in Kaveri is the most advanced CPU-GPU integration yet and offers the most complete solution of the bunch.”

Evolution of HSA in AMD's subsequent mobile APU lines [48]

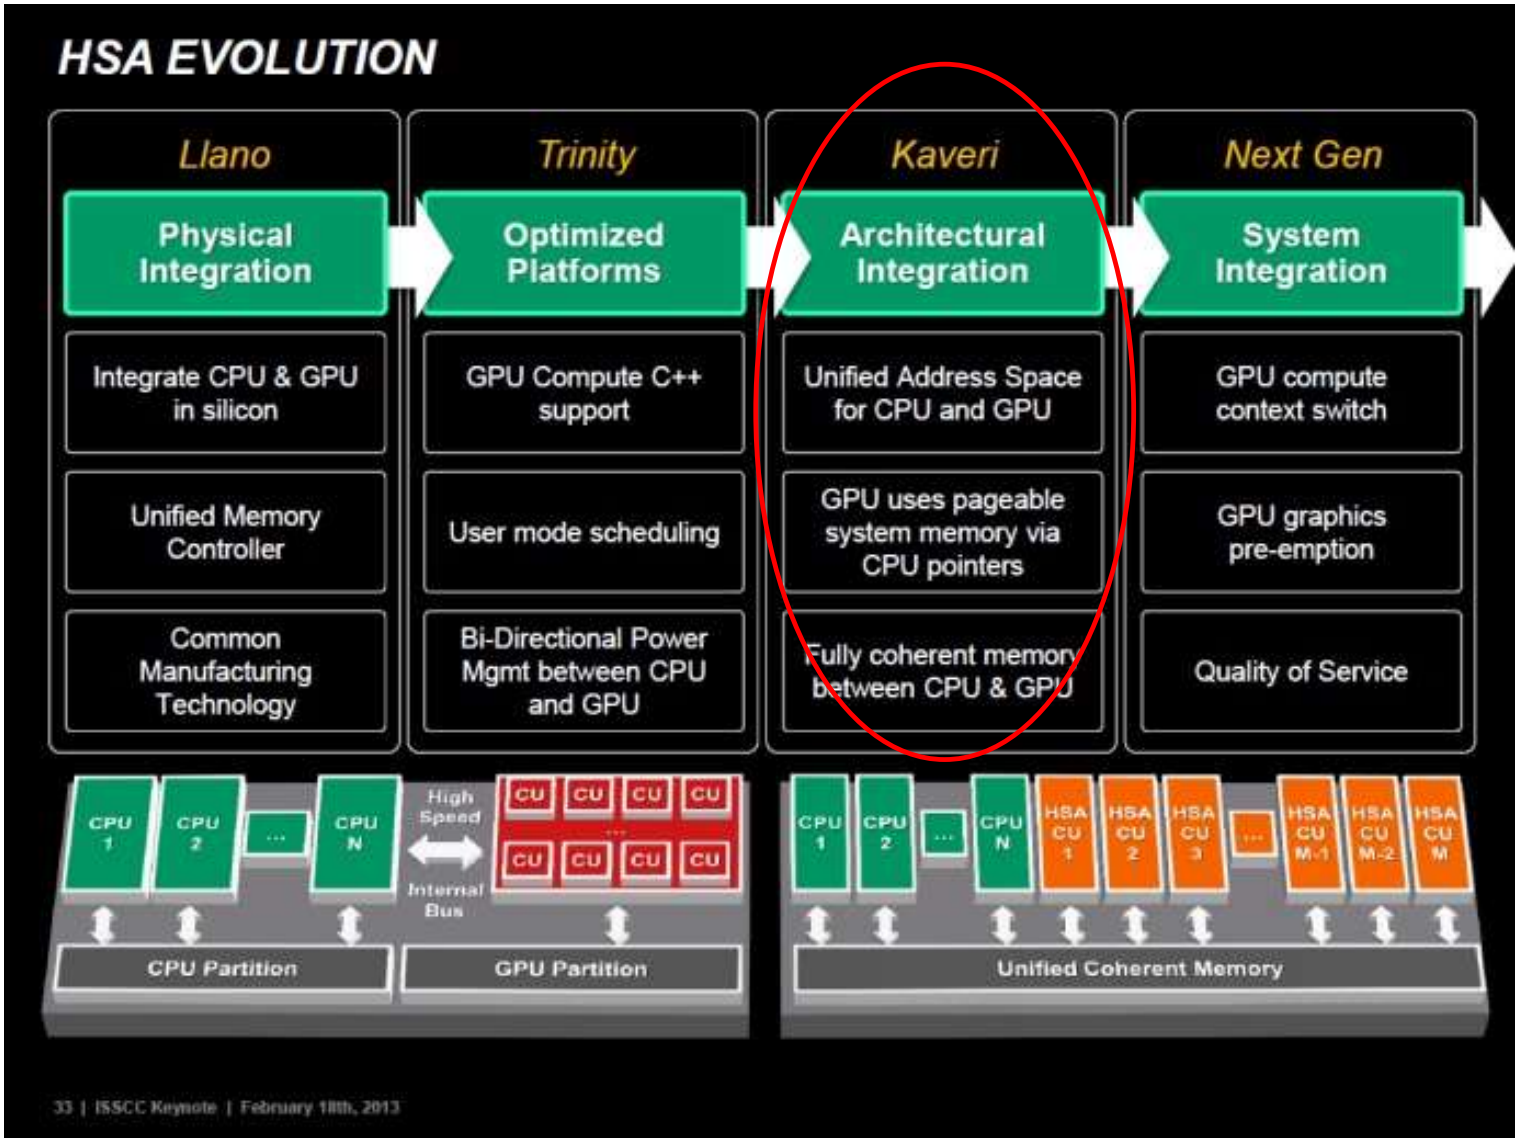

Main components of the Heterogeneous Unified Memory Architecture (hUMA) [96] -1

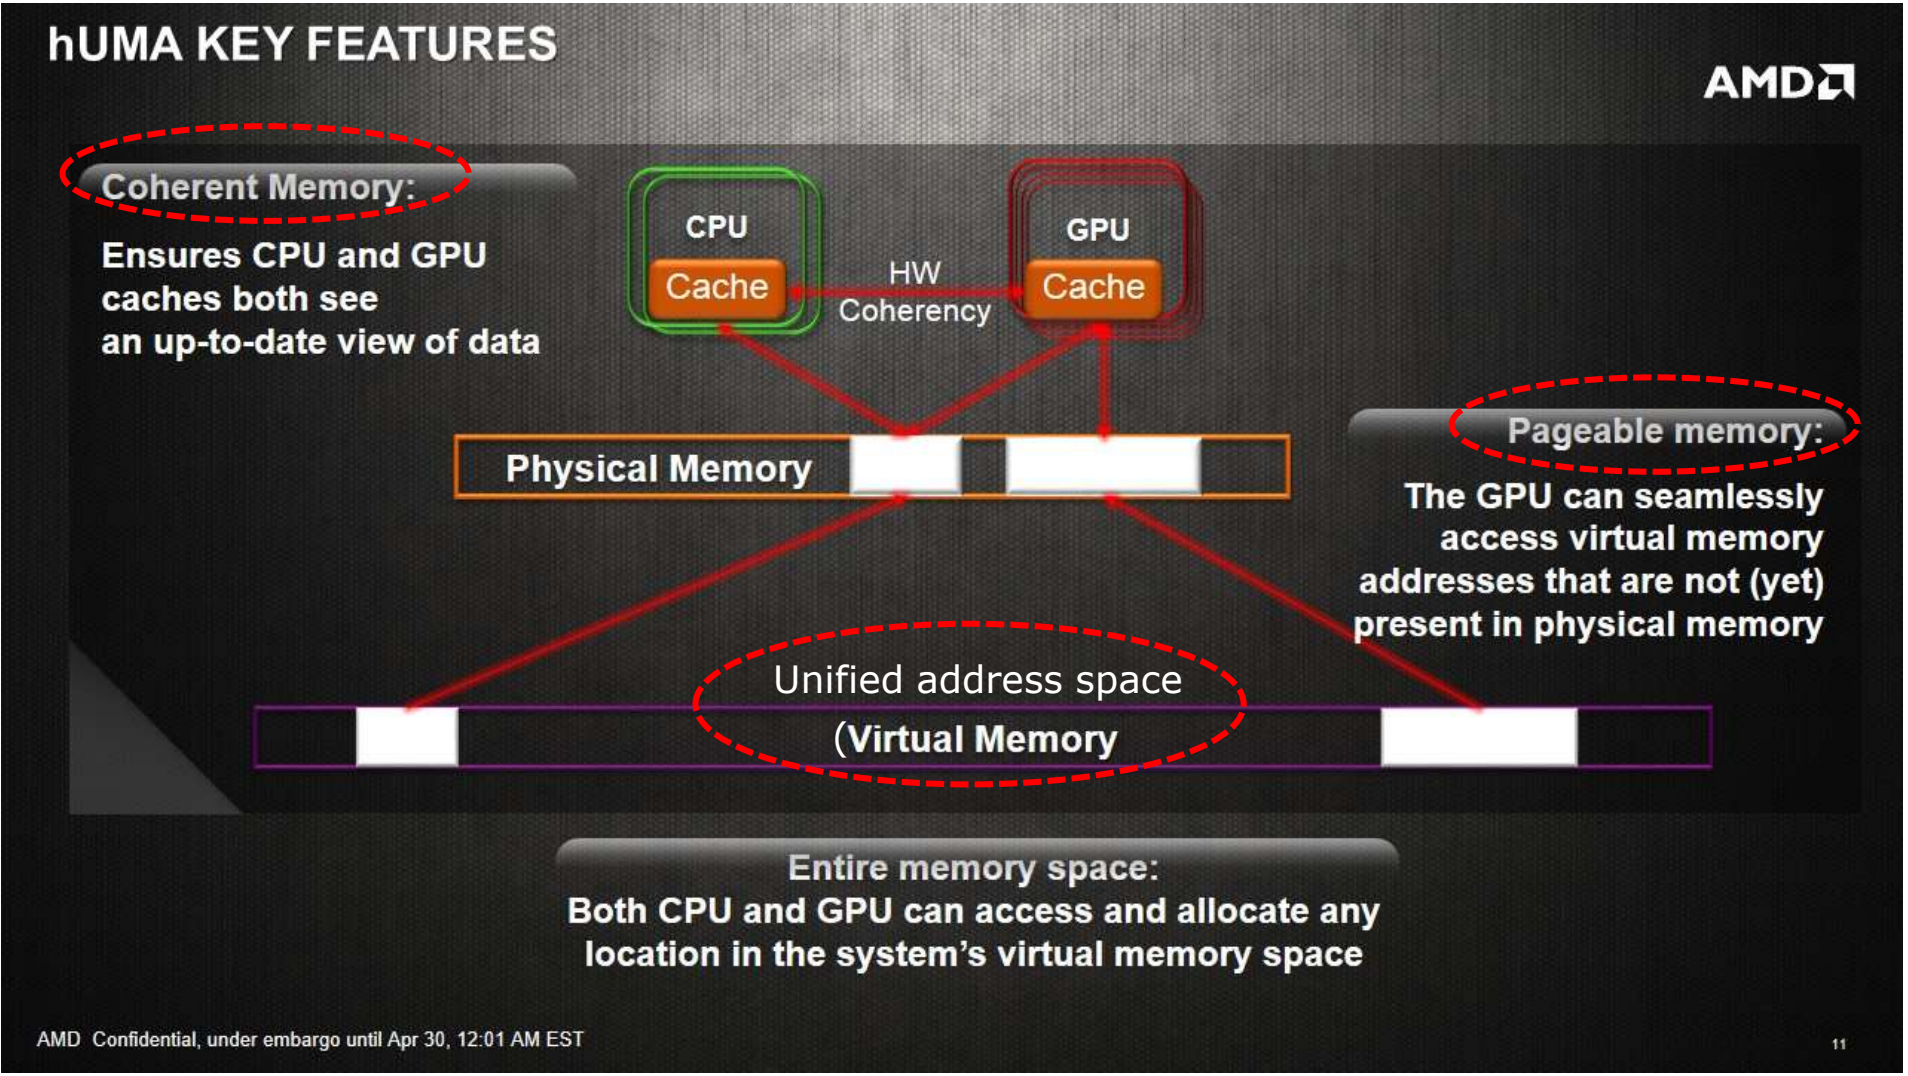

Benefits of the Heterogeneous Unified Memory Architecture (hUMA) [96] -2

"Eliminating CPU-GPU data copies: GPU can now access the entire CPU address space without any copies. In an HSA system, the copy of input data to GPU and copy of results back to CPU can be eliminated.

Access to entire address space: In addition to the performance benefit of eliminating copies, the GPU is also no longer limited to the onboard RAM as is usually the case with discrete GPUs. Even top-end discrete cards top out at about 12GB of onboard RAM currently while a CPU had the advantage of having access to potentially a much larger pool of memory. In many cases, such as scientific simulations, this would mean that the GPU can now work on much larger datasets without any special effort on the part of the programmer to somehow fit the data into GPU's limited address space. Kaveri will have access up to 32GB DDR3 memory, whereby the limiting factor is more the lack of 16GB unregistered non-ECC memory sticks on the market. The latency between the APU and the DRAM still exists however, meaning that a large L3 or eDRAM in the future might improve the scenario, especially in memory bandwidth limited scenarios and pre-empting data fetching.

Unified addressing in hardware: This is the big new piece in Kaveri and HSA that is not offered by any other system currently. Application programs allocate memory in a virtual CPU memory space and the OS maintains a mapping between virtual and physical addresses. When the CPU encounters a load instruction, it converts the virtual address to physical address and may need the assistance of the OS. The GPU also has its own virtual address space and previously did not understand anything about the CPU's address space. In the previous generation of unified memory systems like Ivy Bridge, the application had to ask the GPU driver to allocate a GPU page table for a given range of CPU virtual addresses. This worked for simple data structures like arrays, but did not work for more complicated structures. Initialization of the GPU page table also created some additional performance overhead".

GPU co-processing without HAS i.e. without pointers and data sharing [91]

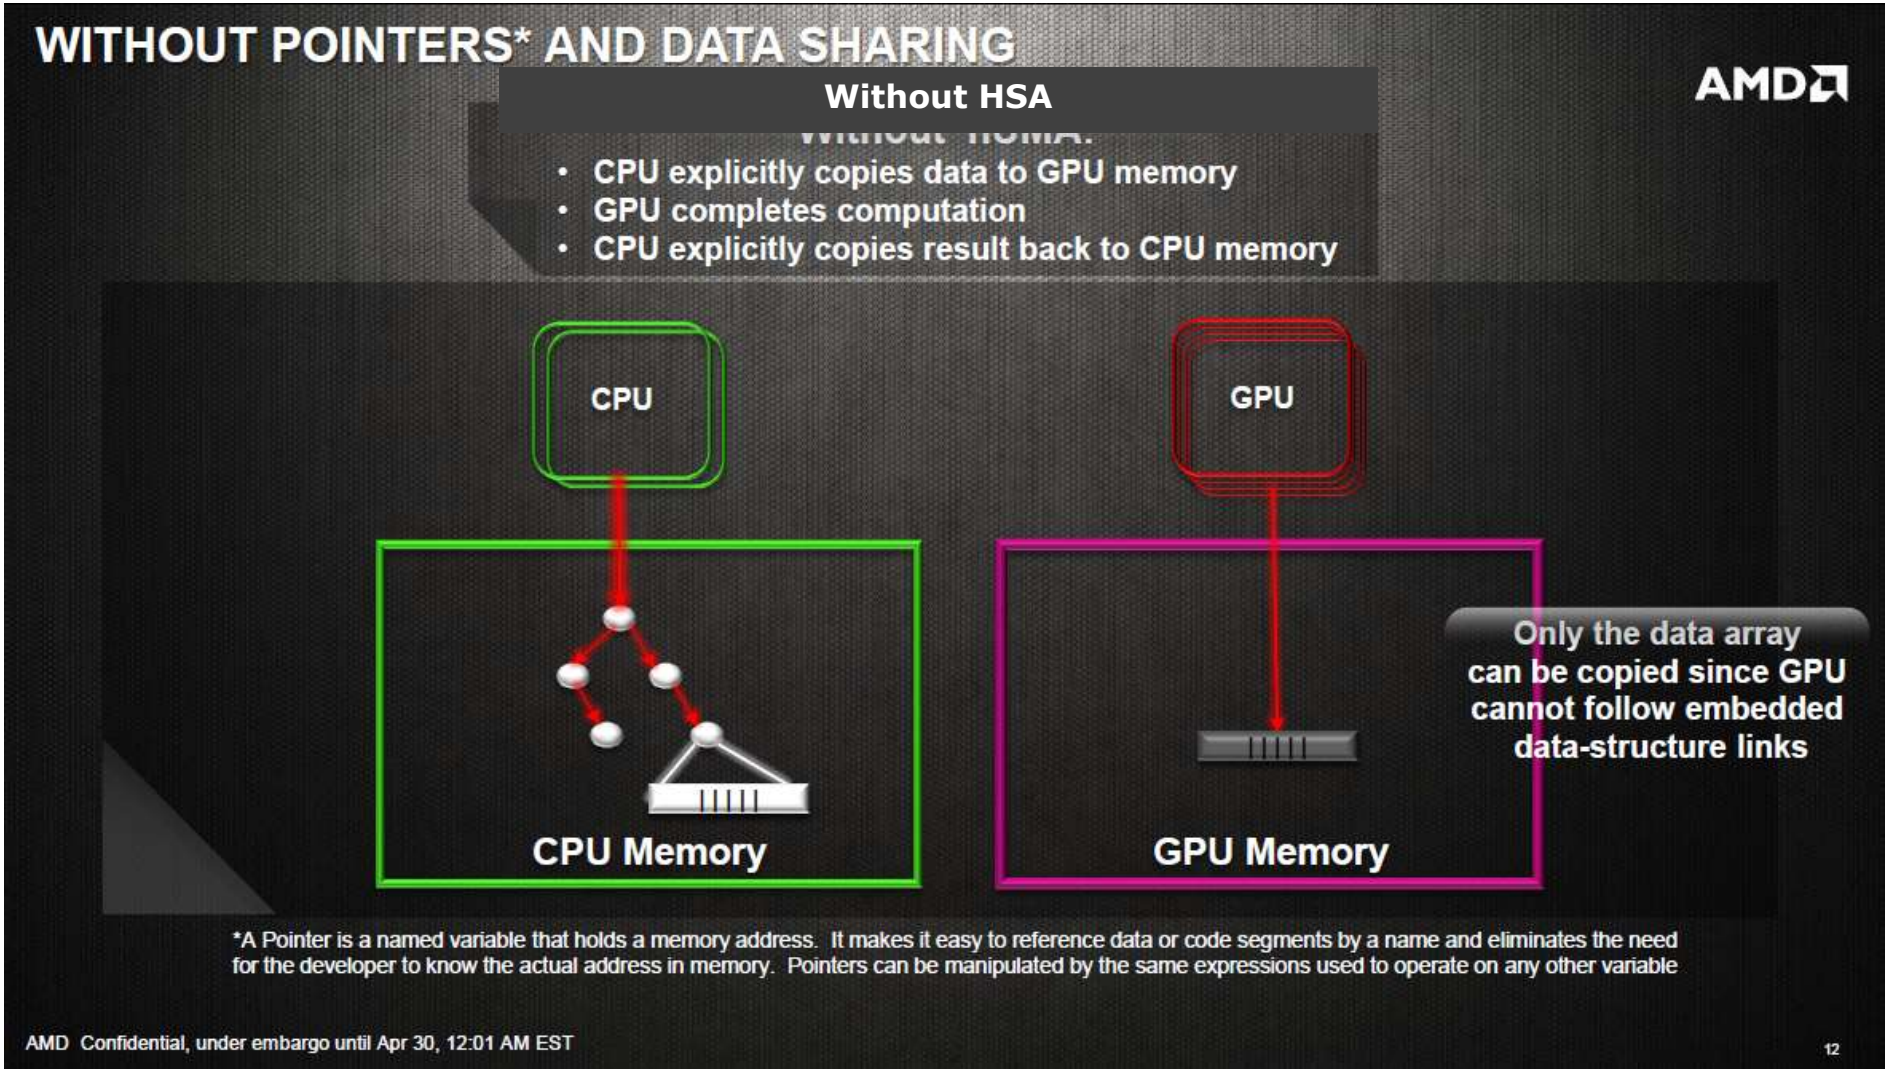

GPU co-processing with HSA i.e. by by using pointers and data sharing [91]

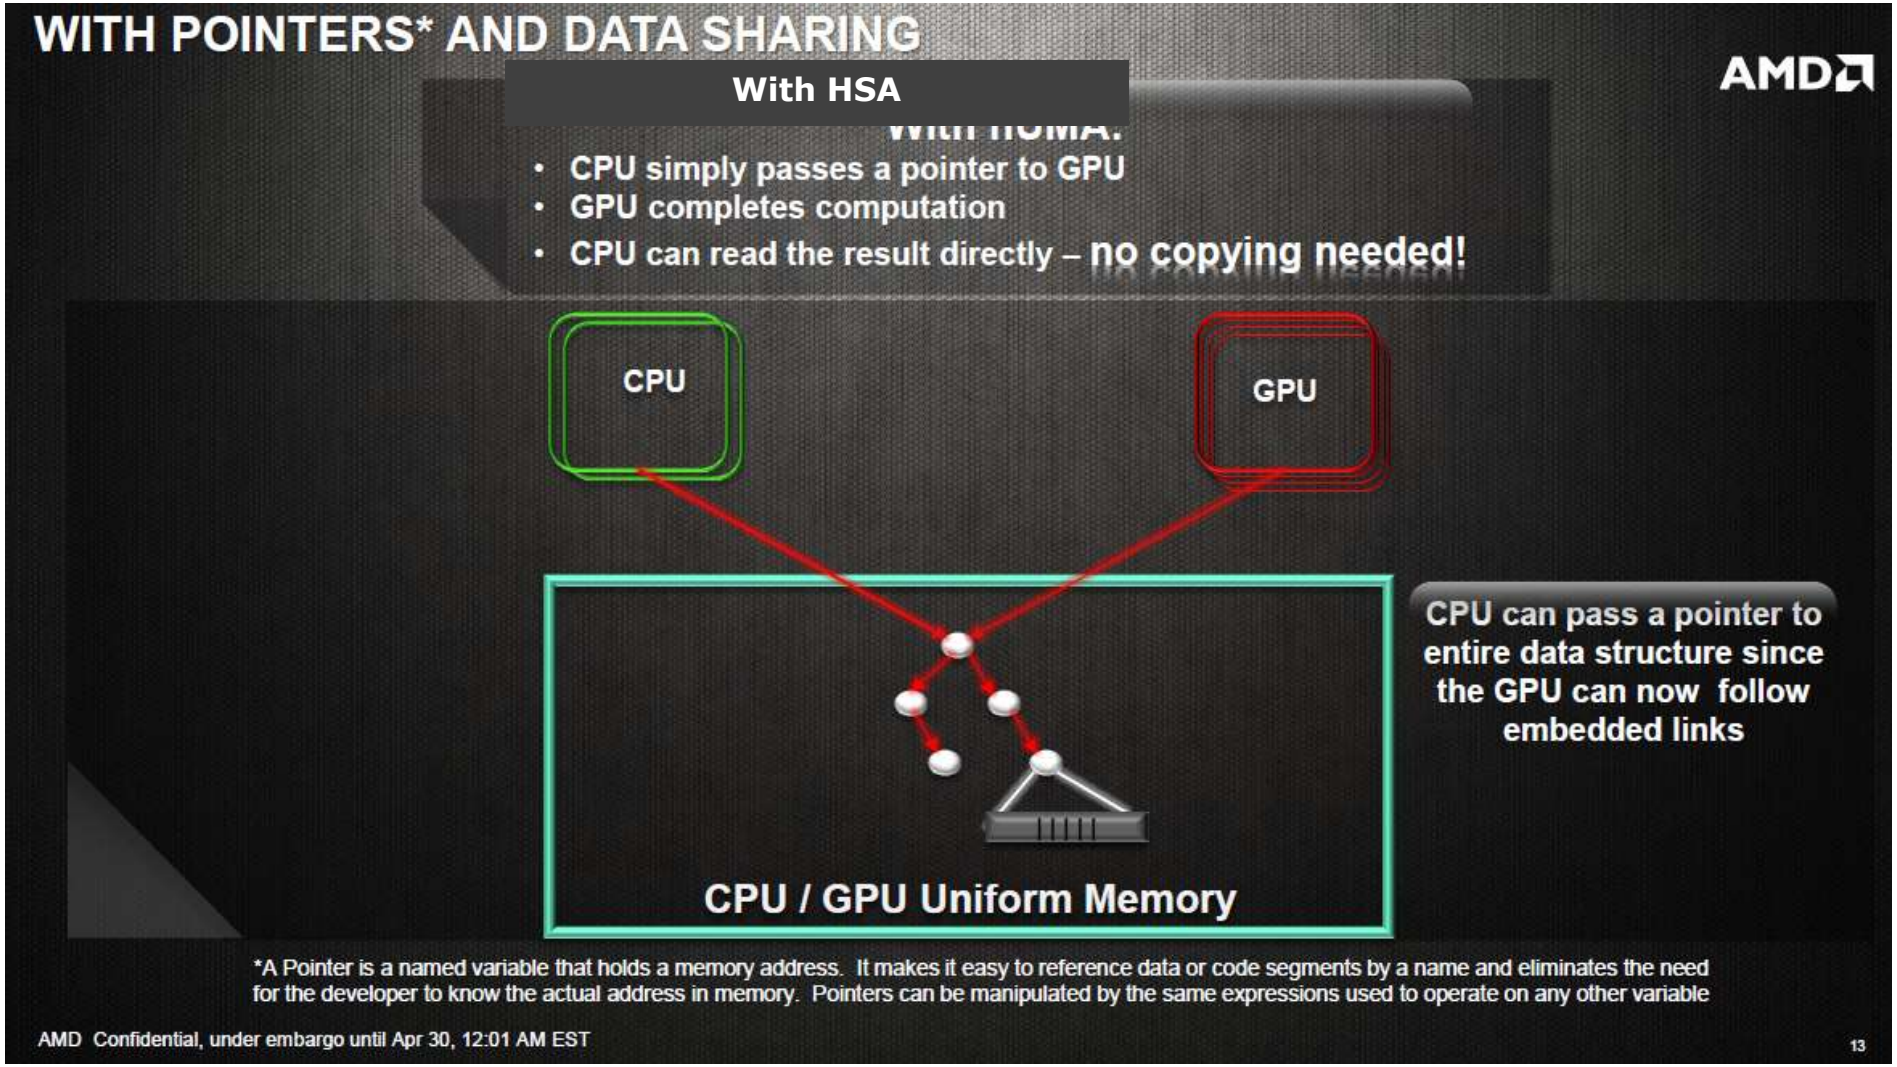

Remark [48]

In order to increase the efficiency of the HSA APU the width of the internal interface, that connects

- the GPU to the coherent system memory space and
- the CPU to the Frame Buffer part of the memory

(termed also as the Fusion Control Link (FCL) or Onion interface, shown in the next Figure) has been widened from 128-bit to 256-bit in both directions.

This enhancement increases the data transfer bandwidth between the CPU and the GPU significantly.

For comparison: Data transfers in the memory hierarchy of the Llano APU (called Fusion Memory Hierarchy) [53]

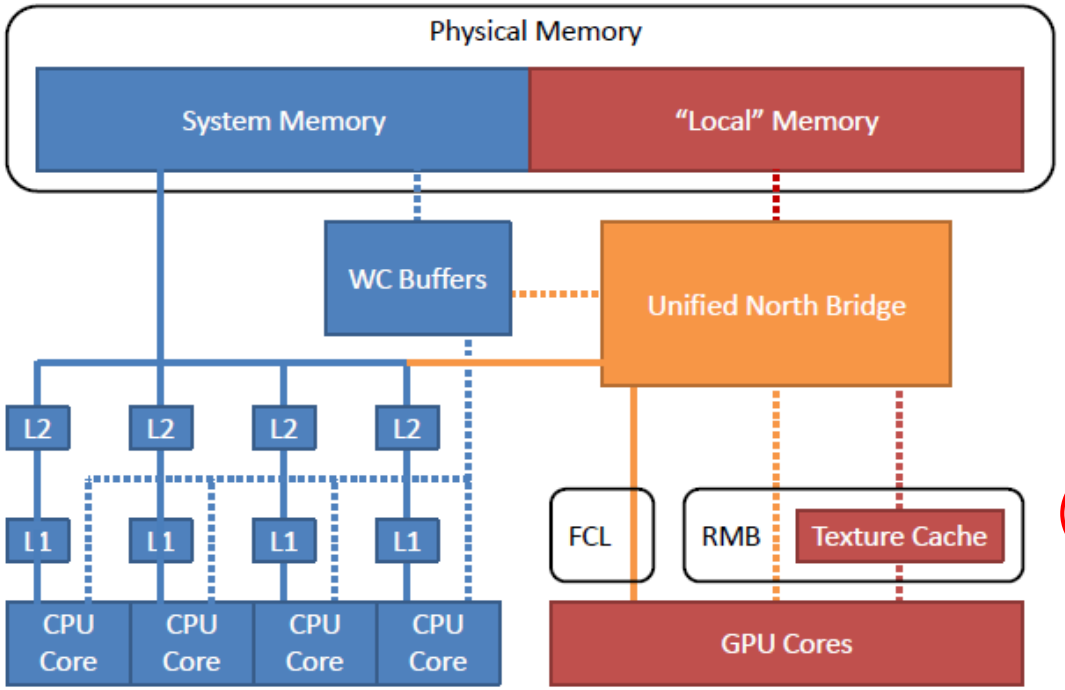

"Frame Buffer" in the physical memory

FCL: Bi-directional, 2x128-bit wide (Onion interface)

RMB: Bi-directional, 2x256-bit wide

The Fusion Memory Hierarchy. The solid lines in this figure indicate cache coherent connections, and the dashed lines show lack of coherence. Blue indicates components of a traditional CPU memory hierarchy and red shows components of a traditional GPU hierarchy. For example, the CPU usually accesses System Memory through the L2 cache and the write-combining buffers. Orange indicates novel features and paths in Fusion. The familiar cache hierarchy of the CPUs is connected to the GPU cores by the FCL. The RMB preserves high bandwidth access from the GPU cores to the "Local" memory (optionally storing data in the texture cache). The CPU cores can access this same "Local" memory via the write-combining buffers through the Unified North Bridge.

HSAIL: Portable Pseudo-ISA for Heterogeneous Compute [96]

“The HSA Foundation wants that the same heterogeneous compute applications run on all HSA-enabled systems. Thus, they needed to standardize the software interface supported by any HSA-enabled system. HSA foundation wanted a low-level API to the hardware that can be targeted by compilers of different languages. Typically compilers target the instruction-set of a processor. However, given the diversity of hardware being targeted by HSA (CPUs, GPUs, DSPs and more), standardizing on an instruction-set was not possible. Instead, HSA Foundation has standardized on a pseudo-instruction set called HSAIL. HSAIL stands for HSA Intermediate Language. The idea is that the compiler for a high-level language (like OpenCL, C++ AMP or Java) will generate HSAIL and the HSA driver will generate the actual binary code using just-in-time compilation. The idea of a pseudo-ISA has been used in many previous portable technologies such as Java bytecode and the Direct3D bytecode. HSAIL is low-level enough to expose many details of the hardware and has been carefully designed such that the conversion from HSAIL to binary code can be very fast.
In terms of competition, Nvidia provides PTX which has similar goals to HSAIL in terms of providing a pseudo instruction set target for compilers. PTX is only meant for Nvidia systems, though some research projects do provide alternate backends such as x86 CPUs. HSAIL will be portable to any GPU, CPU or DSP that implements HSA APIs”

b) AMD's dual graphics (called also hybrid graphics)

- With Kaveri AMD allows the use of hybrid graphics, designated as dual graphics, that is using both a discrete graphics card and the integrated graphics at the same time in order to boost graphics performance.
- Precondition is the use of an R7-based APU and GDDR3 based R7 GPU.
- Dual graphics needs a new driver.

Benefit of dual graphics [96]

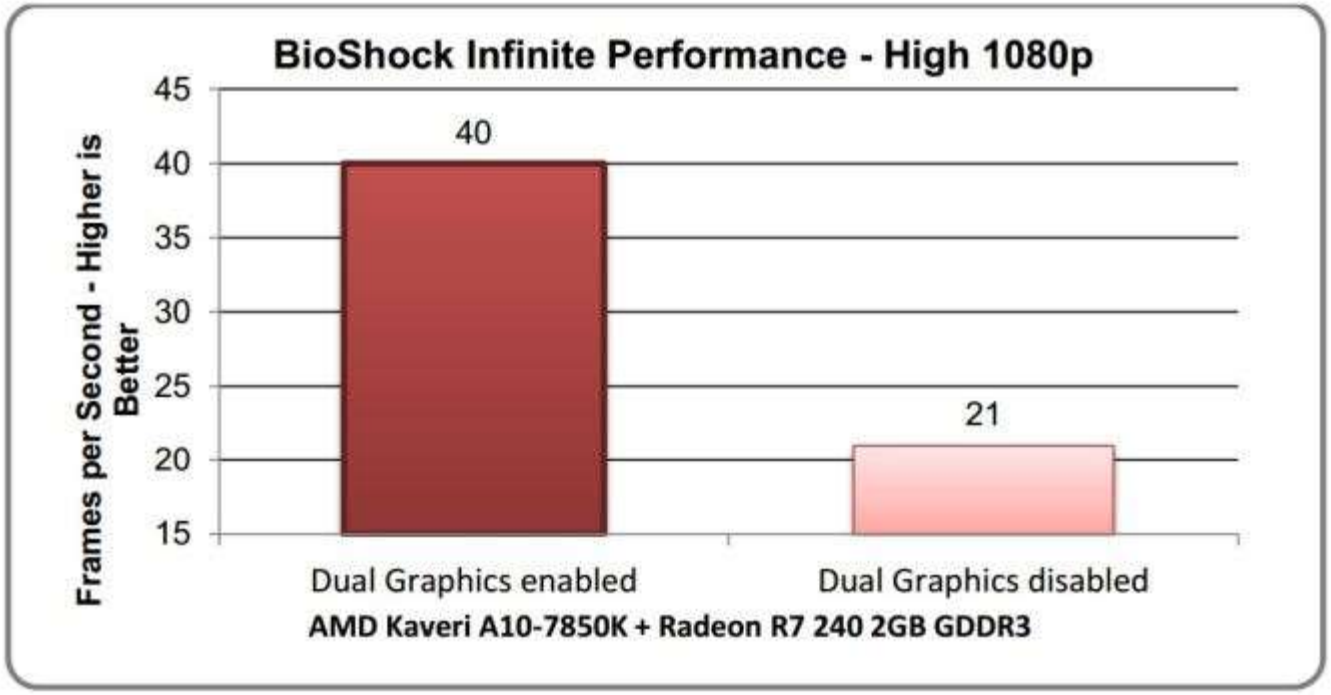

c) Use of GPU cores as compute cores [96] -1

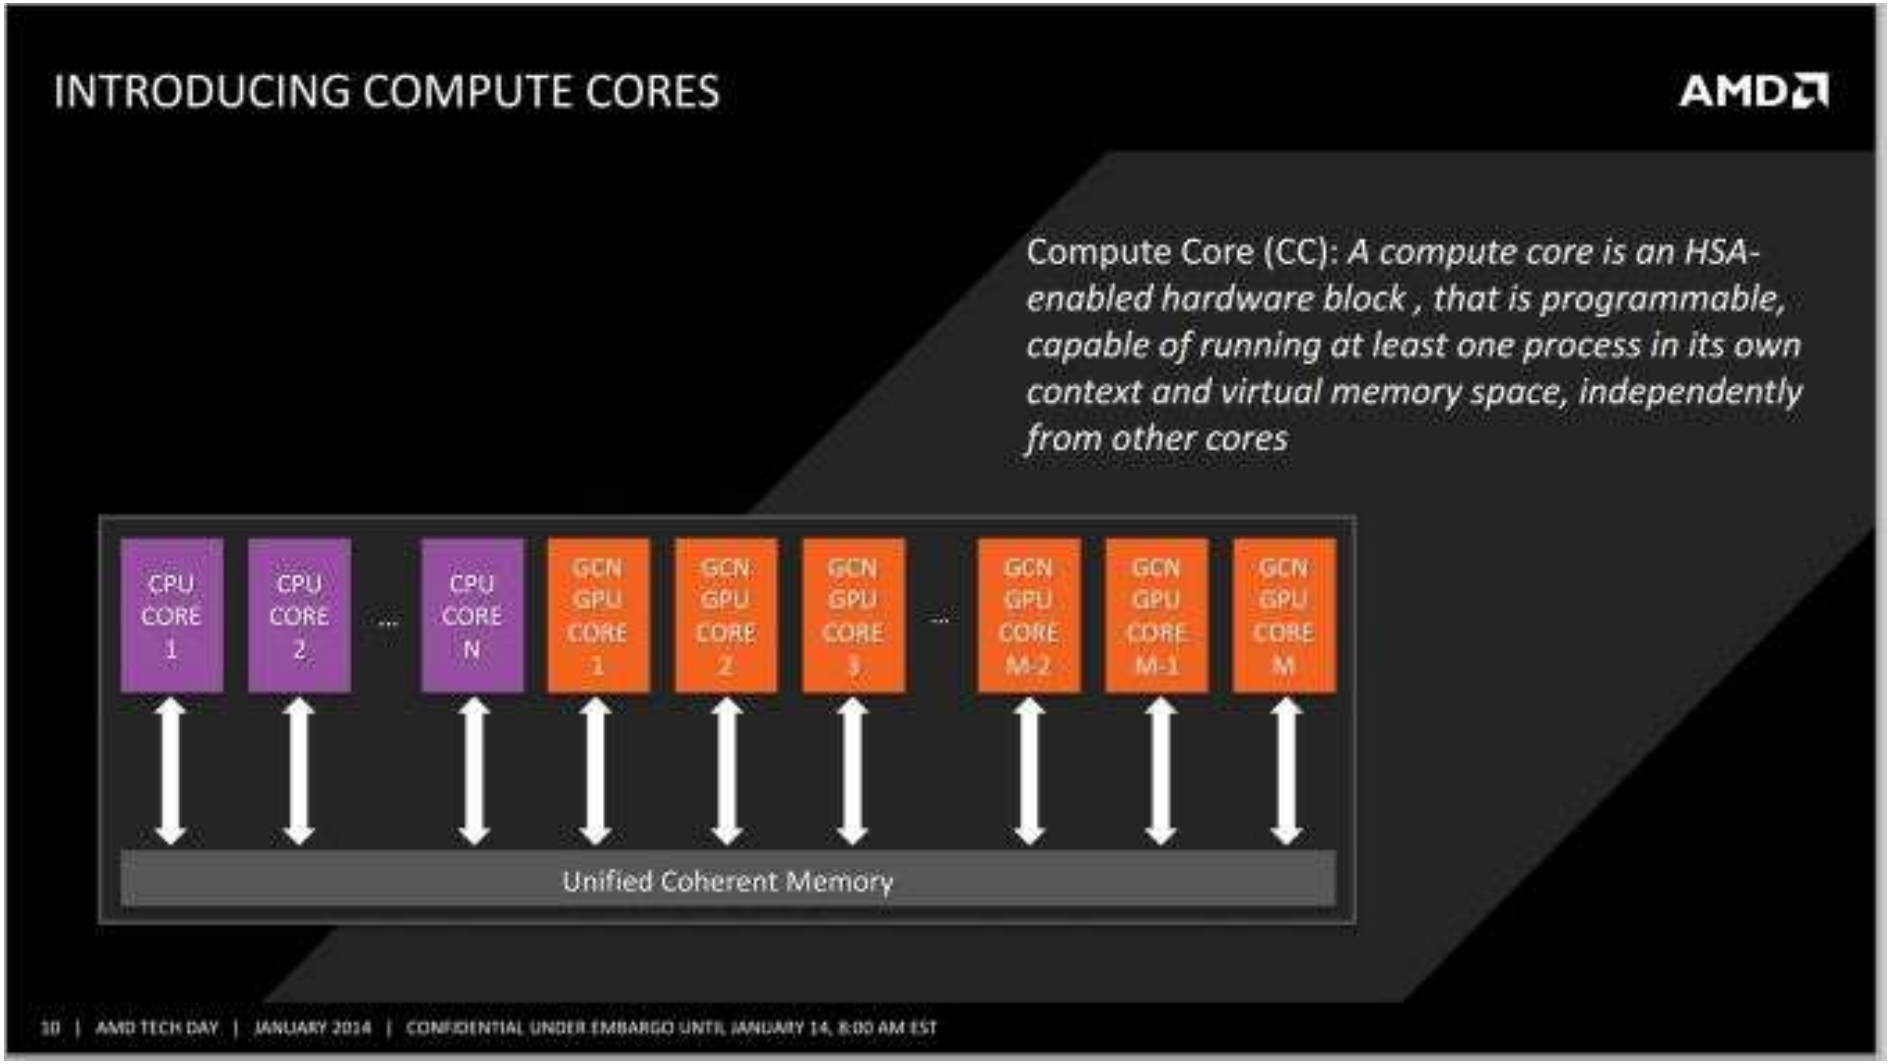

c) Use of GPU cores as compute cores [96] -2

- Each GPU core of the integrated graphics and each CPU core can execute separate code.
- Further on, the GCN architecture allows to spawn as many kernels as compute units.

  Before however, the GPU was restricted to run a single compute kernel at once.
- The 12 compute units are not equivalent, the CPU cores and the GPU cores need different code.

Compatibility between the FM2 and FM2+ sockets [96]

Socket Compatibility Chart		
	Will Work in FM2	**Will Work in FM2+**
Richland	Yes	Yes
Kaveri	No	Yes

- Kaveri is launched with the FM2+ socket.
- This socket has two extra pins that are unused in the FM2.

Benchmark results of DT models for assessing the gaming performance [96]

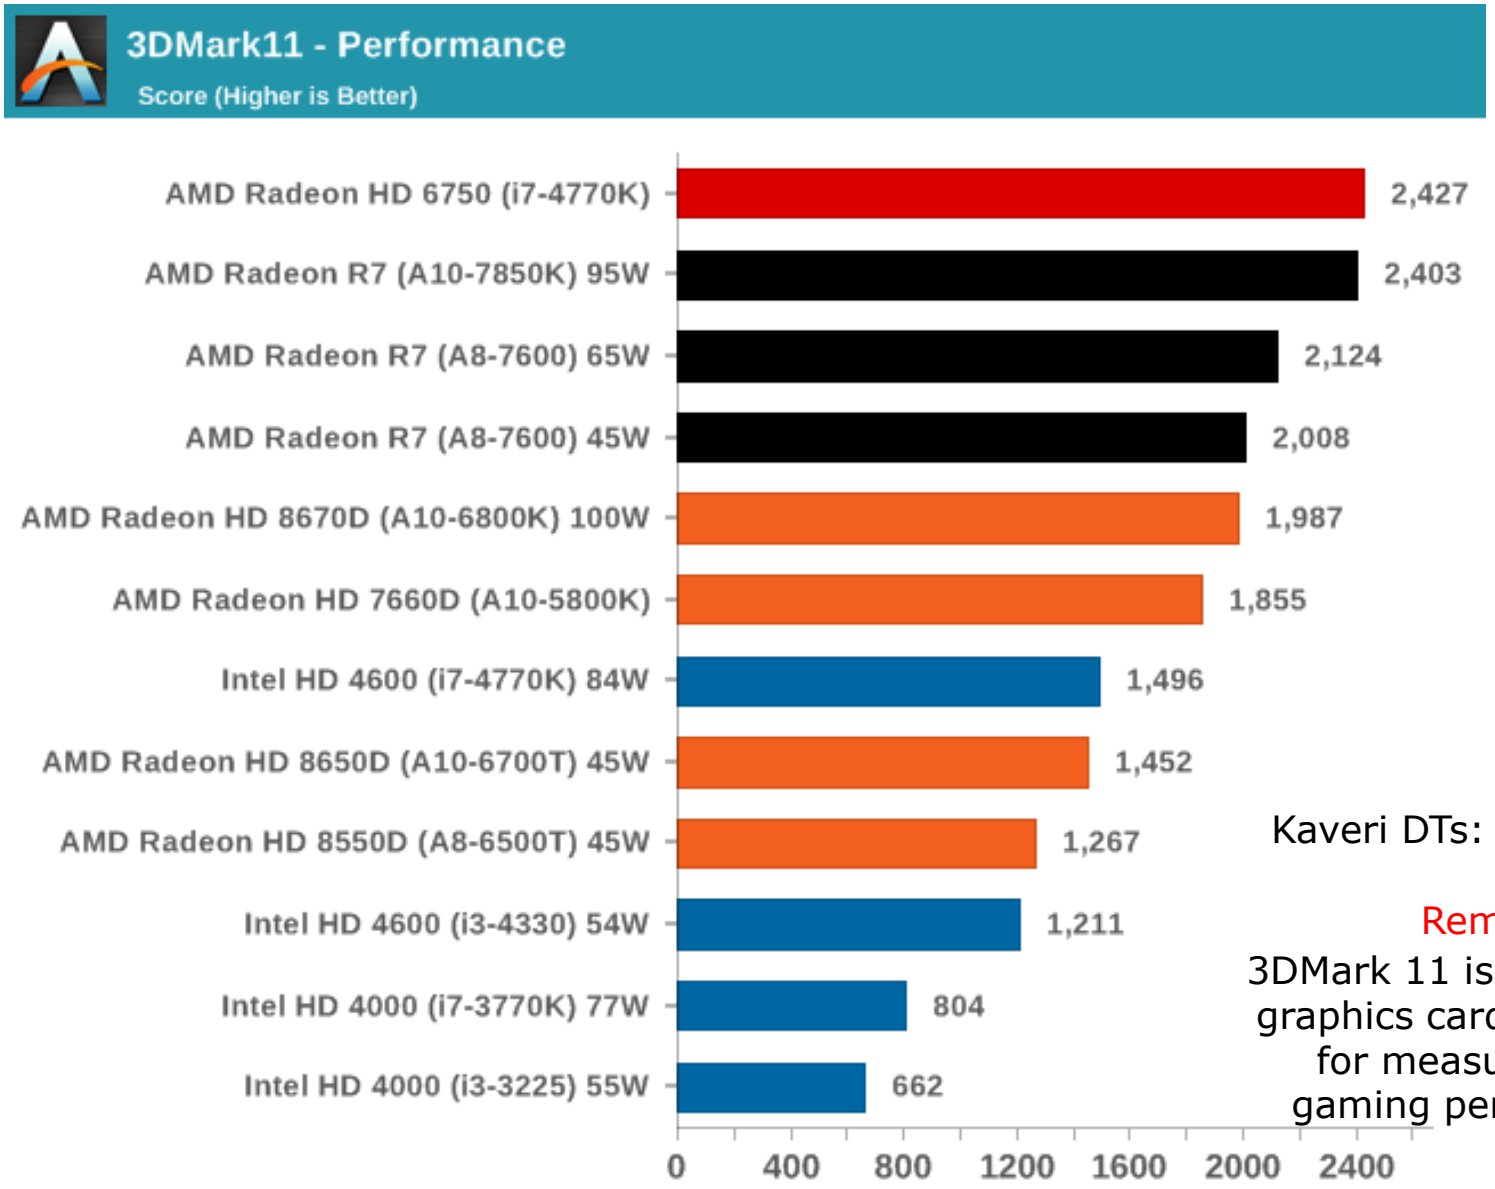

Kaveri DTs: A10/A8 7xxx

Remark

3DMark 11 is a DirectX 11 graphics card benchmark for measuring PC's gaming performance.

5. Fourth generation Excavator-based (Family 15h Models 60h-6Fh and 70h-7Fh) processor lines

- 5.1 Overview of the Excavator-based processor lines
- 5.2 The Excavator (version 1) Compute Module
- 5.3 The The 1. gen. Excavator-based Carrizo mobile and DT line
- 5.4 The 2. generation (version 2) Compute Module
- 5.5 2. generation Excavator-based desktop and mobile APU lines

5.1 Overview of the Excavator-based processor lines

5.1 Overview of the Excavator-based processor lines -1

	Launched in	2011	2012	2013	2013	2015	2016
		Family 15h (00h-0Fh) (**Bulldozer**)	**Family 15h** (10h-1Fh) (**Piledriver**)	**Family 15h** (10h-1Fh) **(Piledriver v.2)**	**Family 15h** (30h-3Fh) (**Steamroller**)	**Family 15h** (60h-6Fh) (**Excavator v.1**)	**Family 15h** (77h-3Fh) (**Excavator v.2**)
Servers	**4P servers** (85-140 W)	Interlagos	Abu Dhabi				
	2P servers (85-140 W)	Valencia	Seoul				
	1P servers (85-140 W)	Zurich	Delhi				
Desktops	**High perf.** (~95-125 W)	Zambezi FX-Series	Vishera FX-Series				
	Mainstream (~65-95 W)		Trinity A10-A4	Richland A10/A8/A6/A4	Kaveri A10/A8		
Notebooks	**Mainstream** (~25-35 W)		Trinity A10/A8/A6M	Richland A10/A8/A6M	Kaveri FX/A10/A8P		Bristol Ridge FX/A12/A10P
	Ultra-thin (~10 – 15 W)		Trinity A10/A6M	Richland A10/A8/A6/A4M	A8 Pro/A8(B) A6 Pro/A6(B)	Carrizo FX/A10/A8P	Bristol Ridge FX/A12/A10P Stoney Ridge A9/A6
Tablets (~5 W)							

5.1 Overview of the Excavator-based processor lines -2

AMD's Family 15h Excavator-based processor lines

Excavator v1-based ultra-thin mobile APU line

Family 15h Models 60h-6Fh

Carrizo (6/2015)
Ultra-thin
mobile APU line

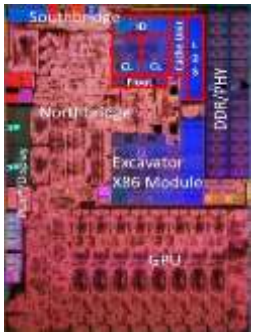

FX8800P/A10/A8 8xxxP
2 modules
R8/R7/R6
2x DDR3-2133
BGA (FP4)
28 nm, 3.1 btrs, 245 mm2

Excavator v2-based mobile APU lines

Family 15h Models 70h-7Fh

Bristol Ridge (5/2016)
Standard and ultra-thin
mobile APU line

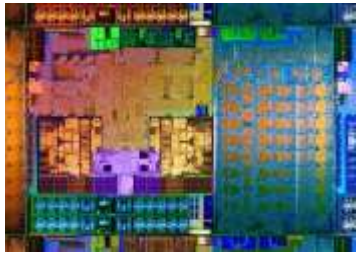

FX98xxP/A12/10/8/6 9xxxP
2 modules
R7/R5 graphics
2x up to DDR4-2400
BGA (FP4)
28 nm, 3.1 btrs, 250 mm2

Stoney Ridge (5/2016)
Ultra-thin
mobile mobile line

A9/A6 9xxx
1 module
R5/R4 graphics
1x up to DDR4-2133
BGA (FP4)
28 nm, 1.2 btrs, 124 mm2

AMD's client roadmap from 03/2015 [99]

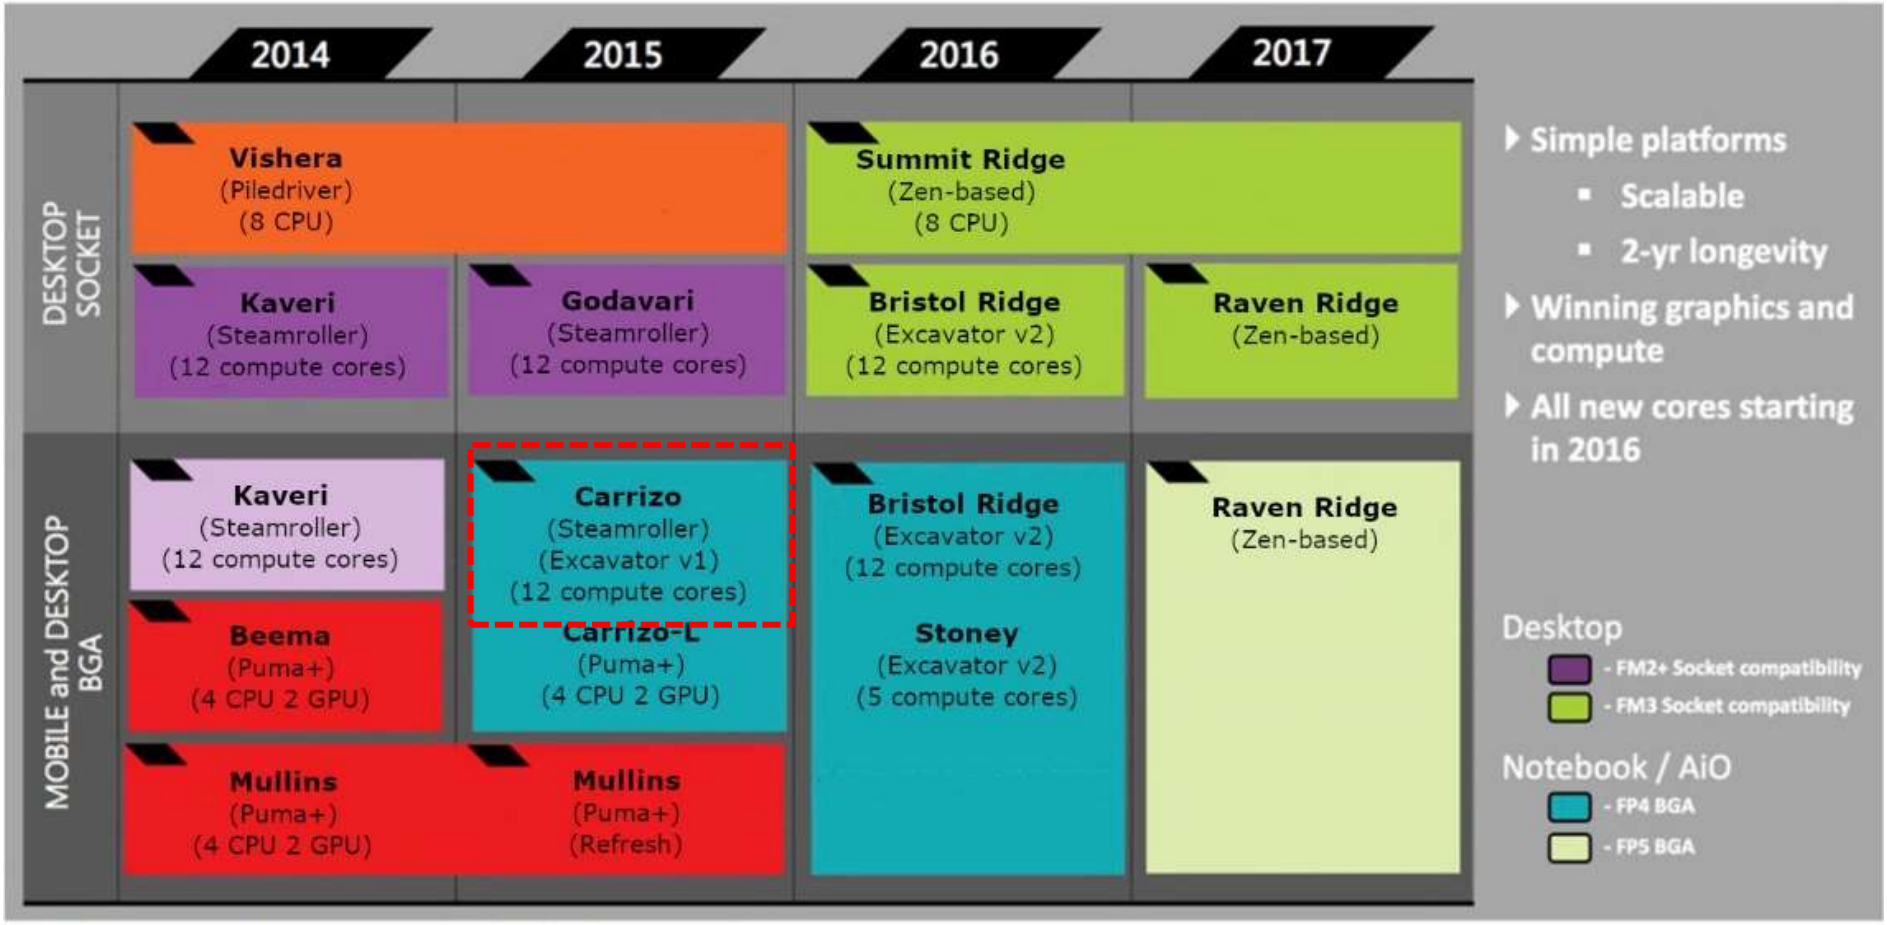

5.2 The Excavator version 1 Compute Module

5.2 The Excavator version 1 Compute Module

Main enhancements of the Excavator (v. 1) Compute module to raise performance [100] **-1**

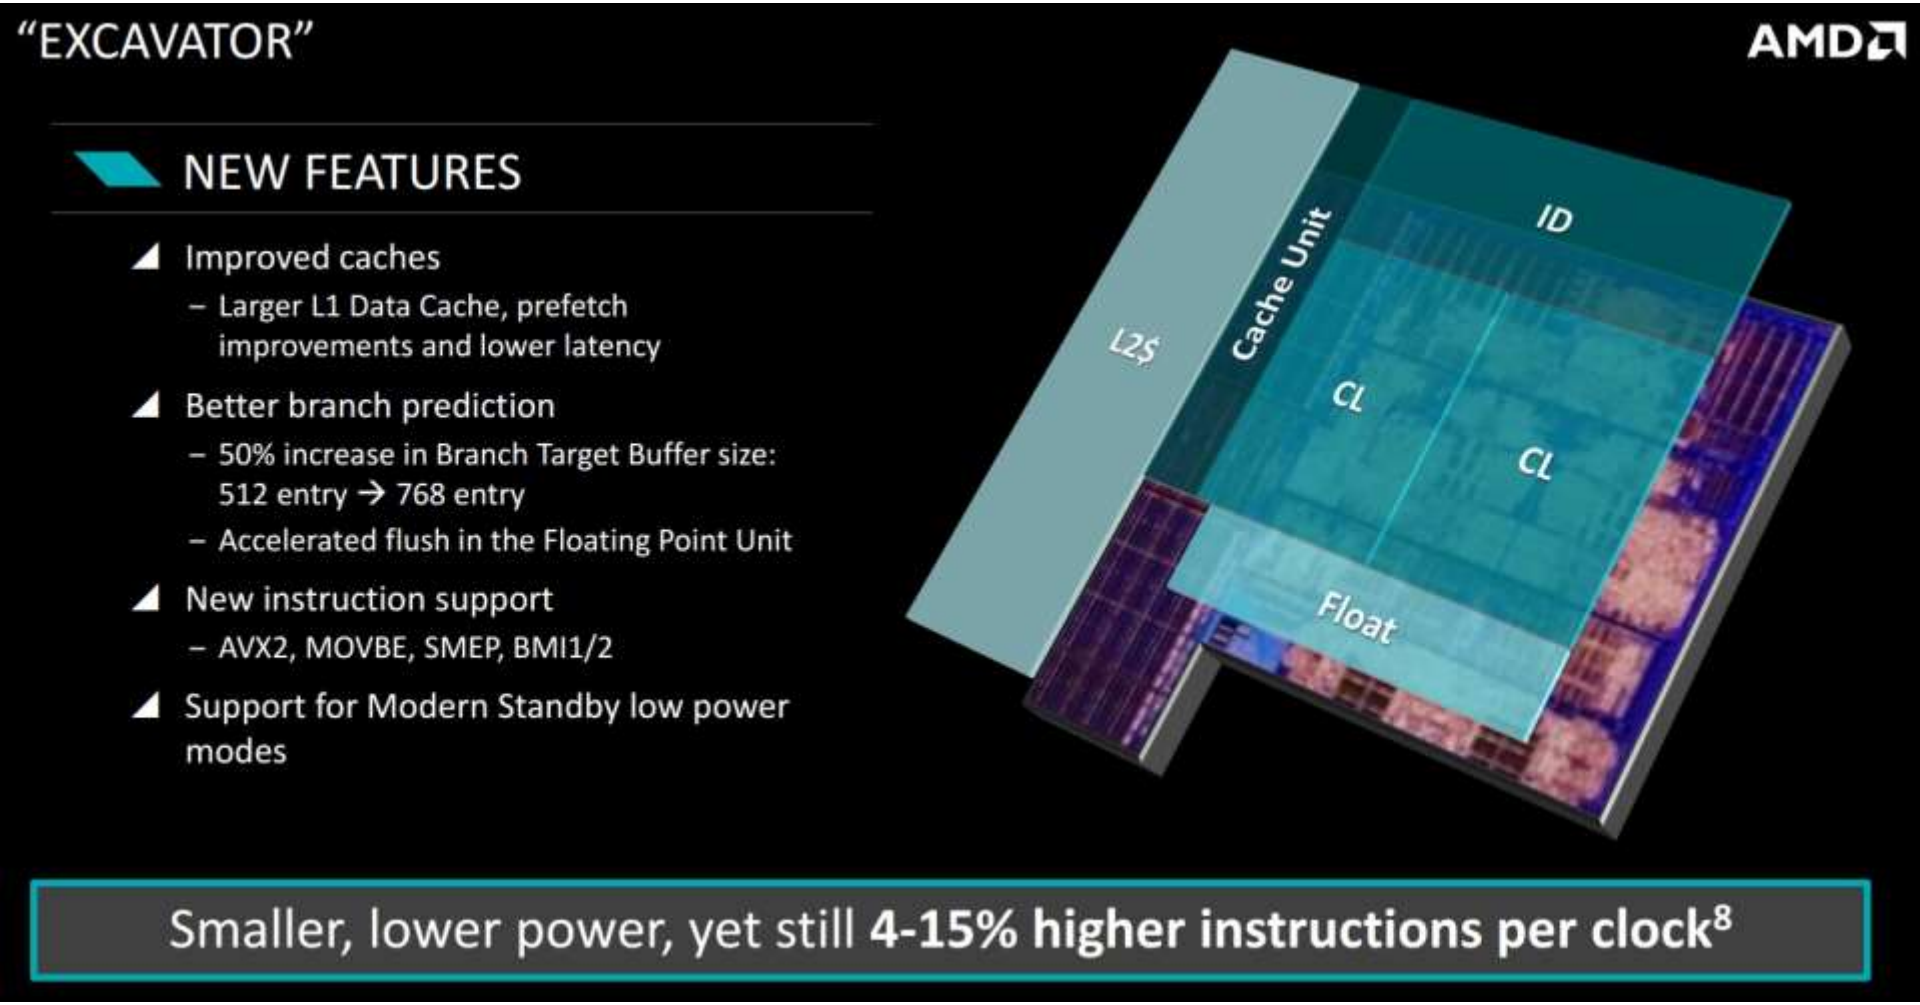

5.2 The Excavator version 1 Compute Module (2)

Main enhancements of the Excavator (v. 1) Compute module to raise performance [100] **-2**

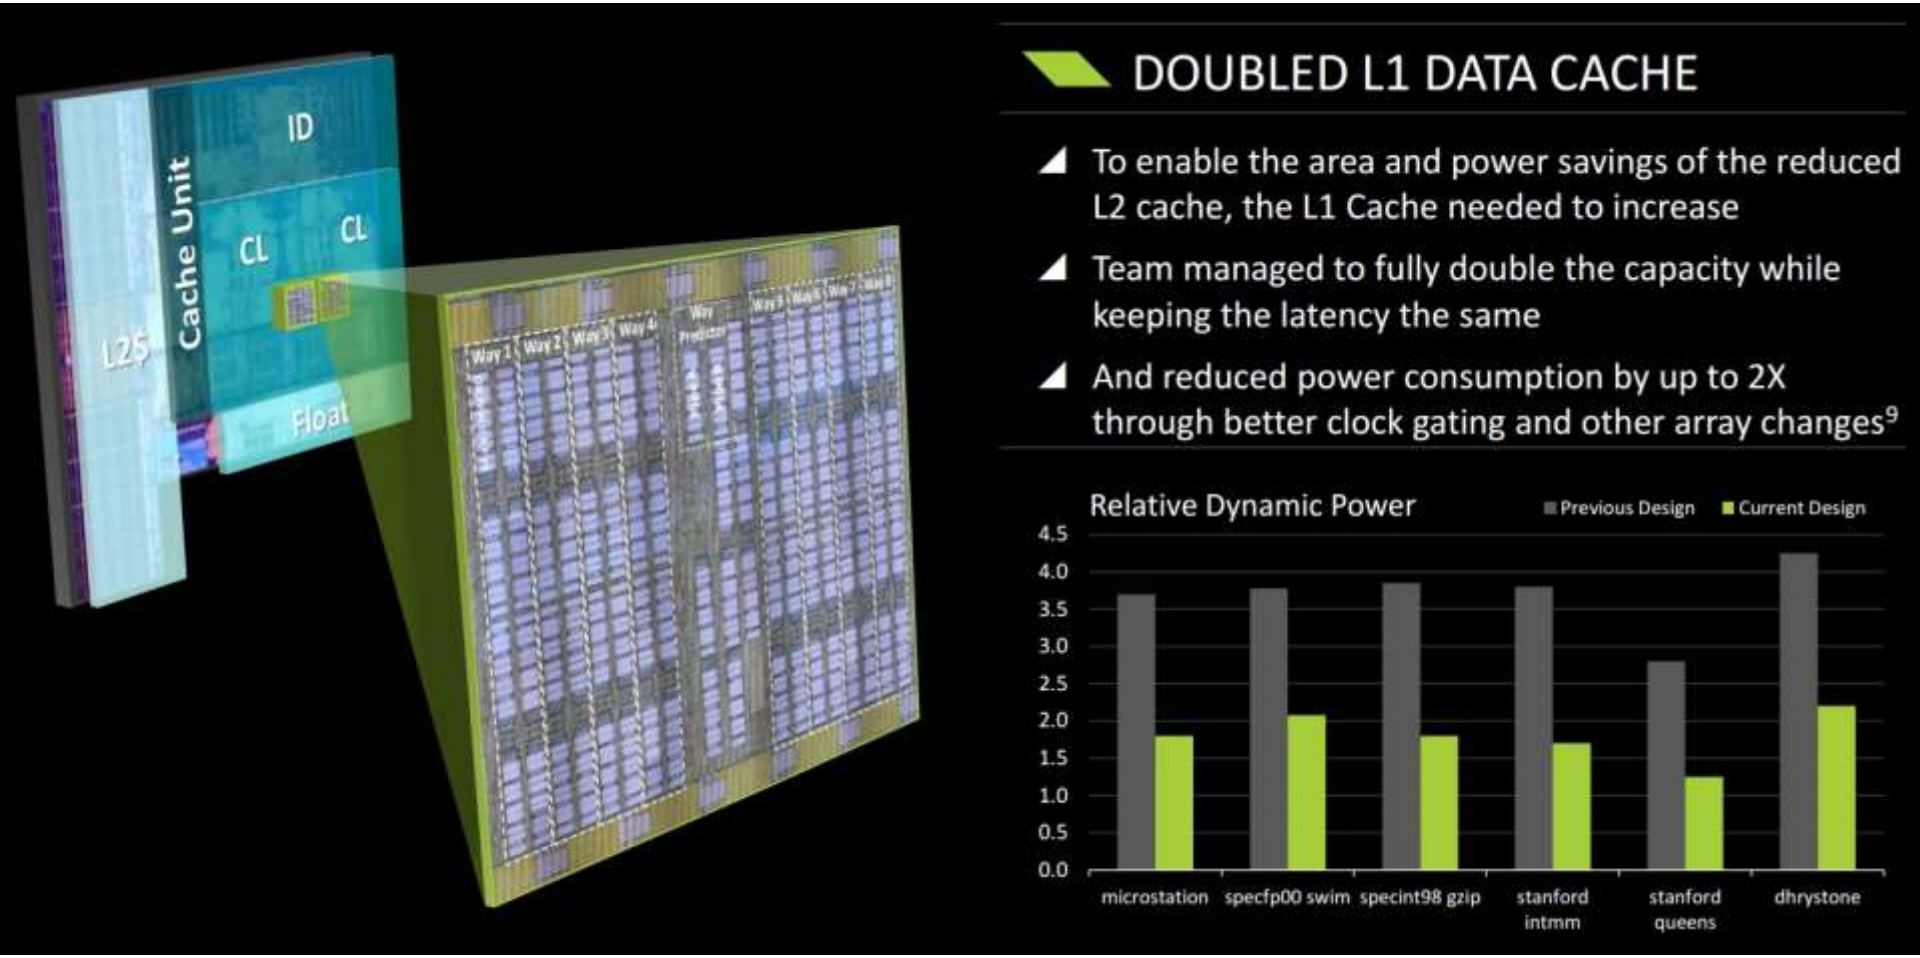

L2 cache size is reduced from 2 MB/module to 1 MB/module
L1D cache size is increased from 16 KB/core to 32 KB/core,
as shown in the next Figure.

5.2 The Excavator version 1 Compute Module (3)

Evolution of the L1/L2 cache architecture of the subsequent compute modules of the Bulldozer family [95], [45], [101]

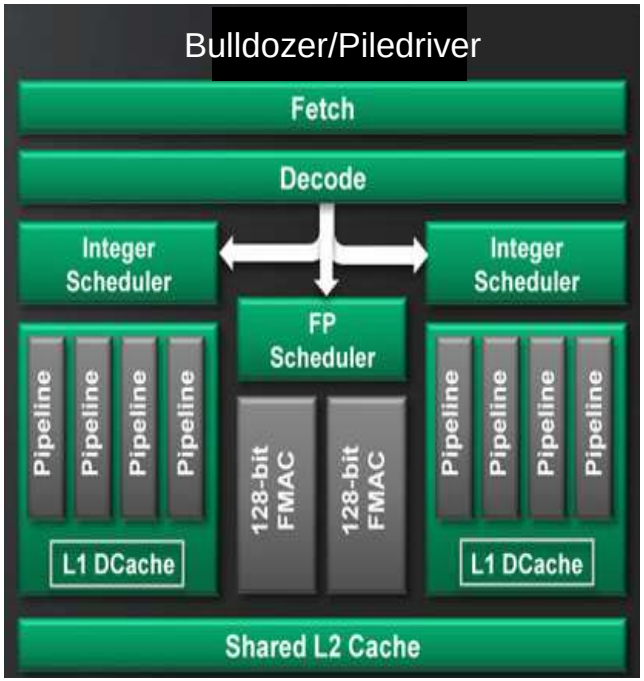

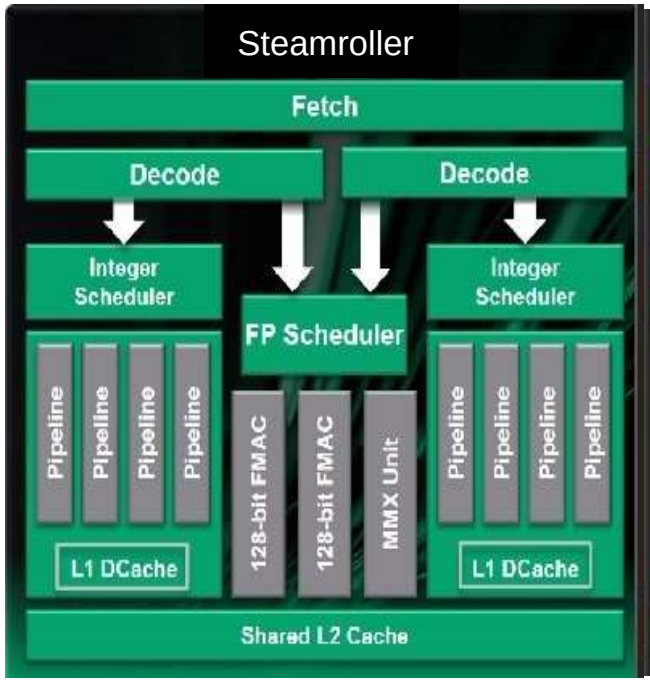

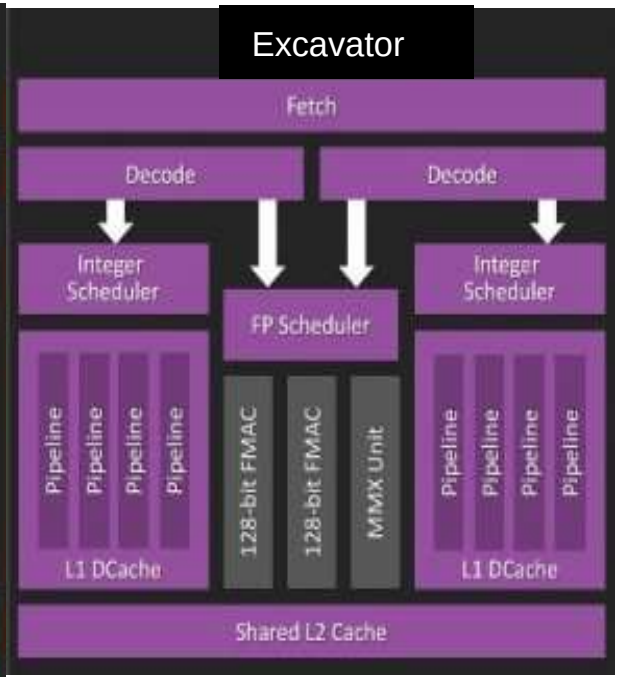

64 KB L1I/module
16 KB L1D/core
2 MB L2/module

96 KB L1I/module
16 KB L1D/core
2 MB L2/module

96 KB L1I/module
32 KB L1D/core
1 MB L2/module

Note

Bulldozer-based servers and high-end DTs, such as the Bulldozer-based Zambezi and the Piledriver-based Vishera DT line, have L3 caches of the size of 2 MB/module.

This is in sharp contrast to Intel's DT and mobile lines that have beginning with the Nehalem microarchitecture L3 caches.

The L1/L2 cache architecture of Zen-based processors [101]

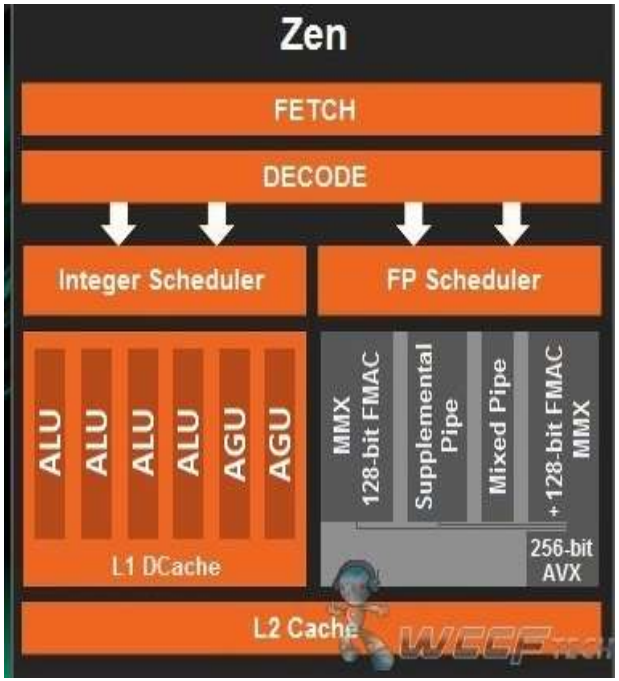

64 KB L1I/core
32 KB L1D/core
1/2 MB L2/core

Zen-based processors are organized into 4-core modules with an L3 cache of 8 MB.

Main enhancements of the Excavator (v. 1) Compute module to raise performance [100] **-1**

PUTTING IT ALL TOGETHER

AMD

"EXCAVATOR" PERFORMANCE

- "Excavator" is optimized for 15W design point
  - Enables increased frequency for up to 39% more performance[10]
- Significant IPC enhancements contribute an additional 9-13% performance[10]
  - Without increasing power consumption
- Total increase of up to 55% in key industry benchmarks such as Cinebench over previous generation[10]

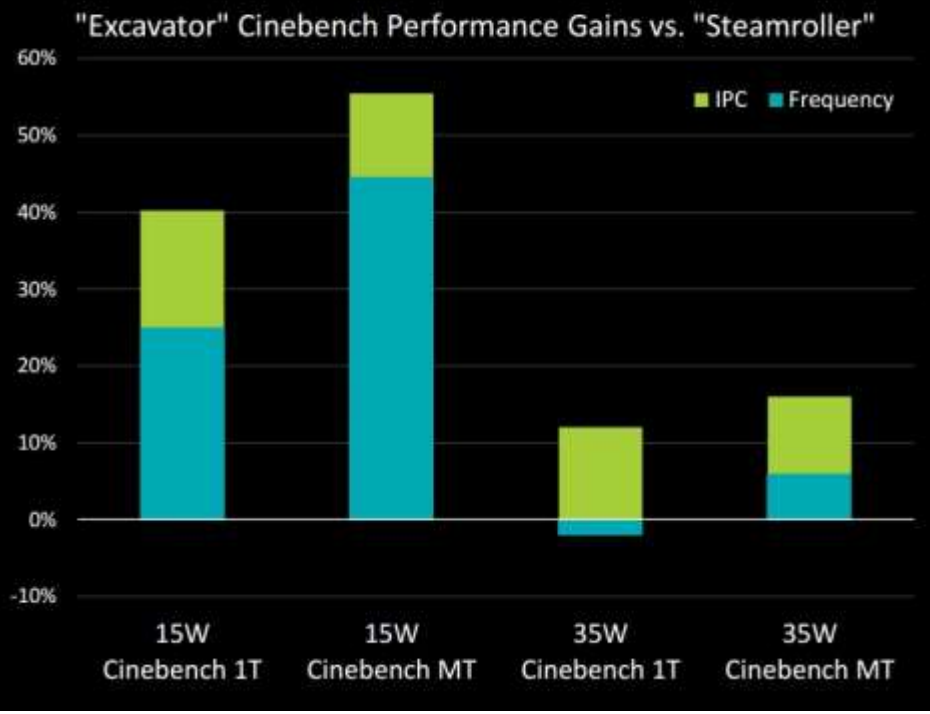

5.2 The Excavator version 1 Compute Module (6)

Main enhancements of the Excavator (v. 1) Compute module to reduce power [100] **-1**

a) Voltage adaptive operation to encounter short voltage drops
b) AVFS (Adaptive Voltage and Frequency Scaling)

a) Voltage adaptive operation to encounter short voltage drops [102]

VOLTAGE ADAPTIVE OPERATION

AMD

- Delivering low noise voltage to high performance CPUs, GPUs and APUs has always been a challenge for the industry
- The variations that happen are typically about 10% of the nominal value – that means at least 20% power is wasted covering these voltage variations (power goes as the square of voltage)
- AMD's unique voltage adaptation feature recovers much of that wasted power by operating at the average voltage and quickly reducing frequency for the brief periods when the voltage reduces

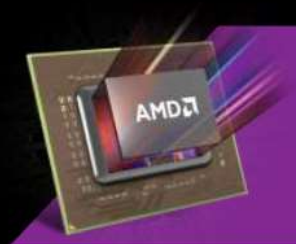

Voltage adaptive feature applied to both CPU and GPU in "Carrizo" results in 19% and 10% power savings respectively

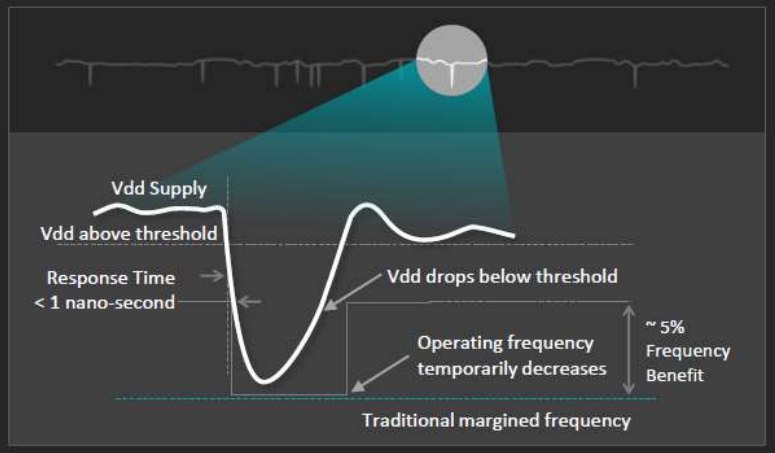

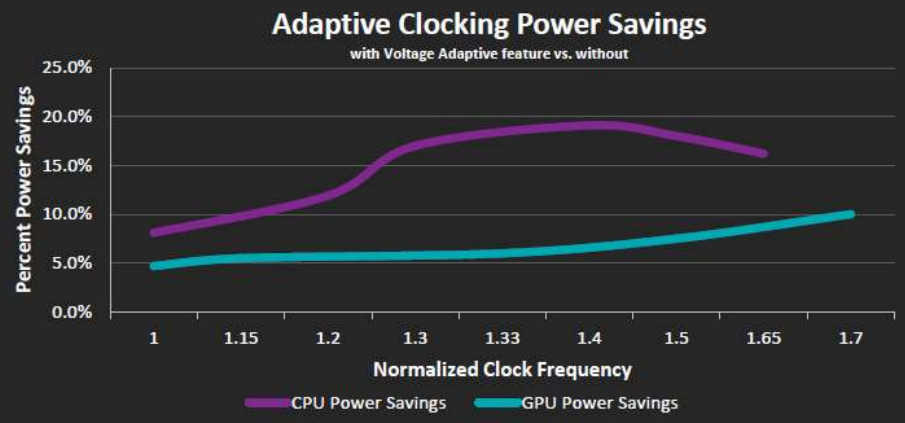

Principle of Voltage adaptive operation [102]

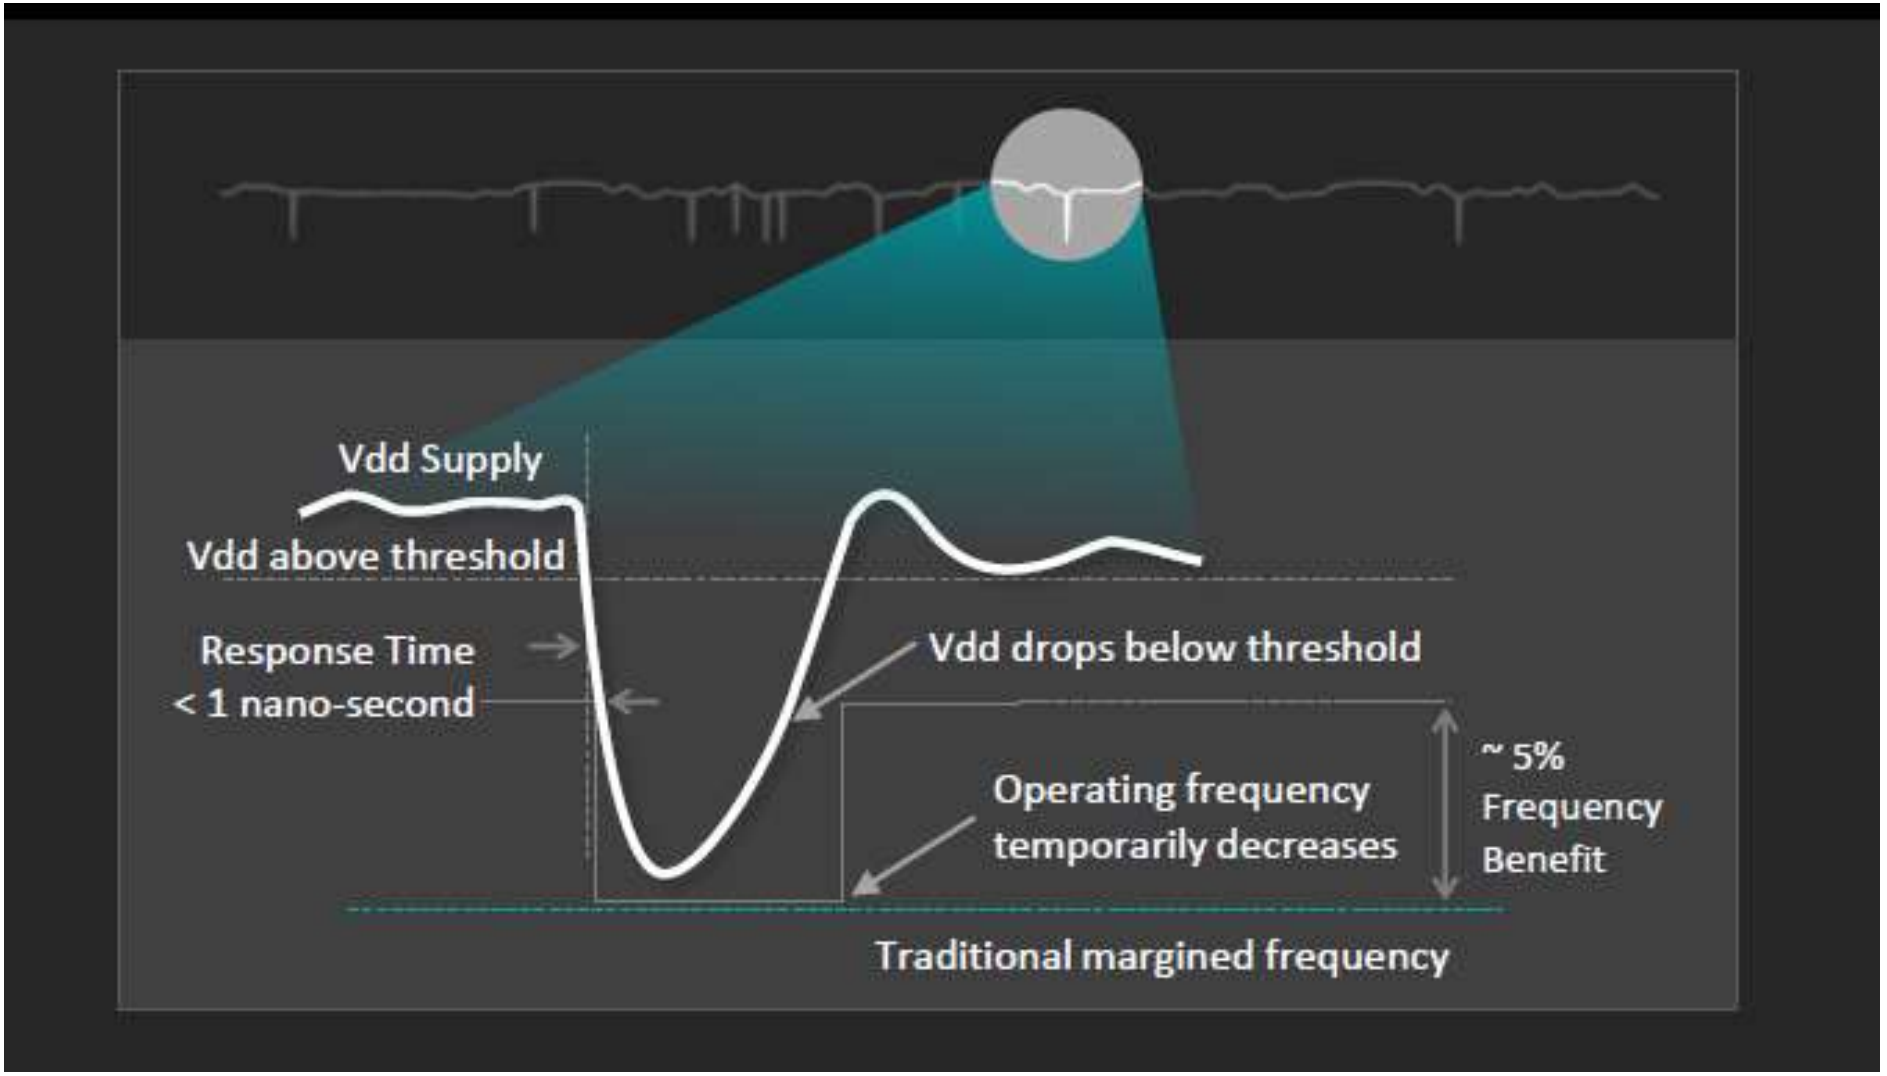

5.2 The Excavator version 1 Compute Module (9)

Remark

AMD introduced Voltage Adaptive Operation already in 4/2014 in their Puma+ core based Beema and Mullins APUs targeting in the first line tablets and notebooks.

5.2 The Excavator version 1 Compute Module (10)

b) AVFS (Adaptive Voltage and Frequency Scaling)

- AMD implemented AVFS first in their Excavator core and then in the Zen-based CCX module.
- AMD's first AVFS implementation is based on the patent US 9,575,553 B2, filed on 19. 12. 2014 [103].
- Next, we will describe the principle of AMD's AVFS implementation based on the cited patent.

Brief description of AMD's AVFS implementation in the Excavator compute unit (simplified) [102]

- AMD implements AVFS by self calibrating the supply voltage needed for a given clock frequency in the functional units, like the CPU or GPU.
- Self calibration makes use of replica paths, that are critical circuit paths those propagation delay limit the max. clock frequency at a given supply voltage or vice versa govern the minimal supply voltage needed for a given clock frequency.
- There are about 500 replica paths on the Excavator die, about 300 are gate dominated, 100 wire dominated and 100 cache dominated.
- Using the replica paths the implementation allows to collect a statistical sample of the propagation delays, and this statistics is utilized to choose the required supply voltage, as detailed subsequently.
- In order to collect statistics the replica paths are connected to 10 Critical Path Accumulators (CPAs), (see the next Figure), as described next.

Critical Path Accumulators on the Excavator die (called AVFS modules in the Figure [102]

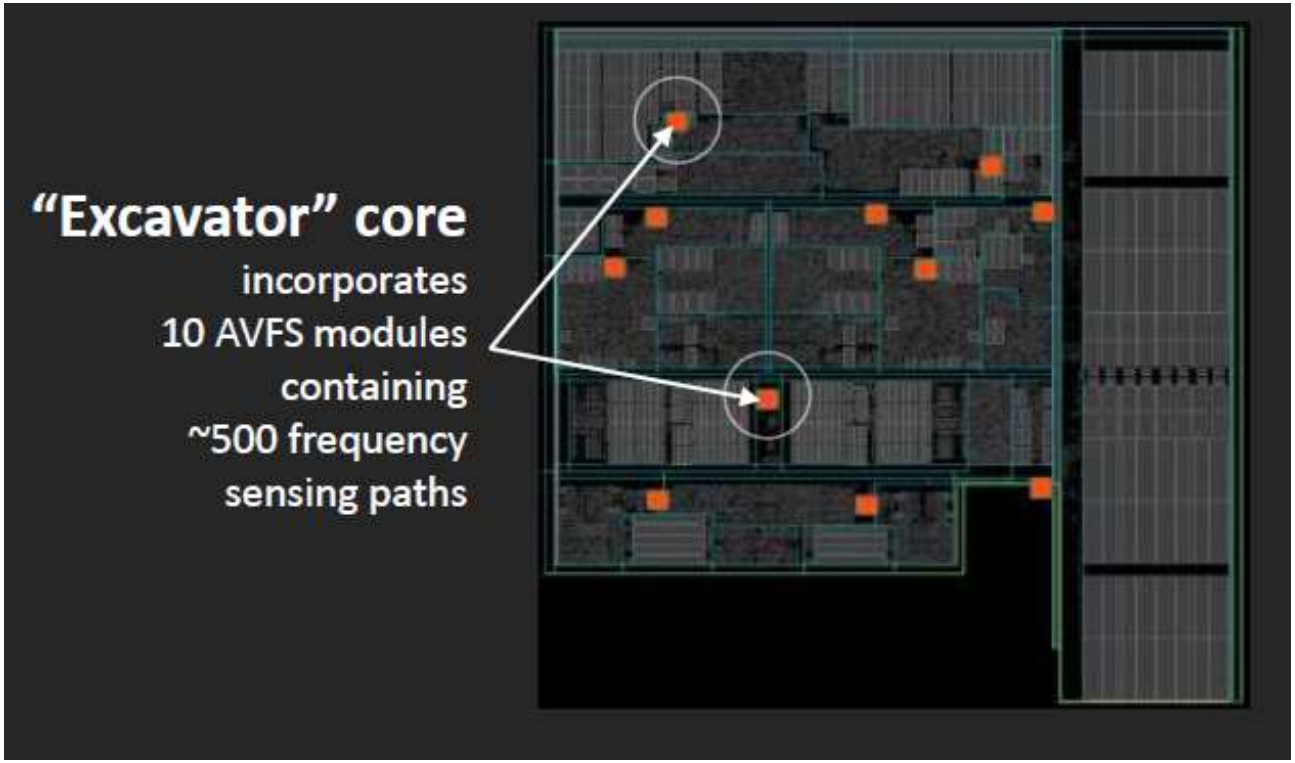

Block diagram of a CPA [103]

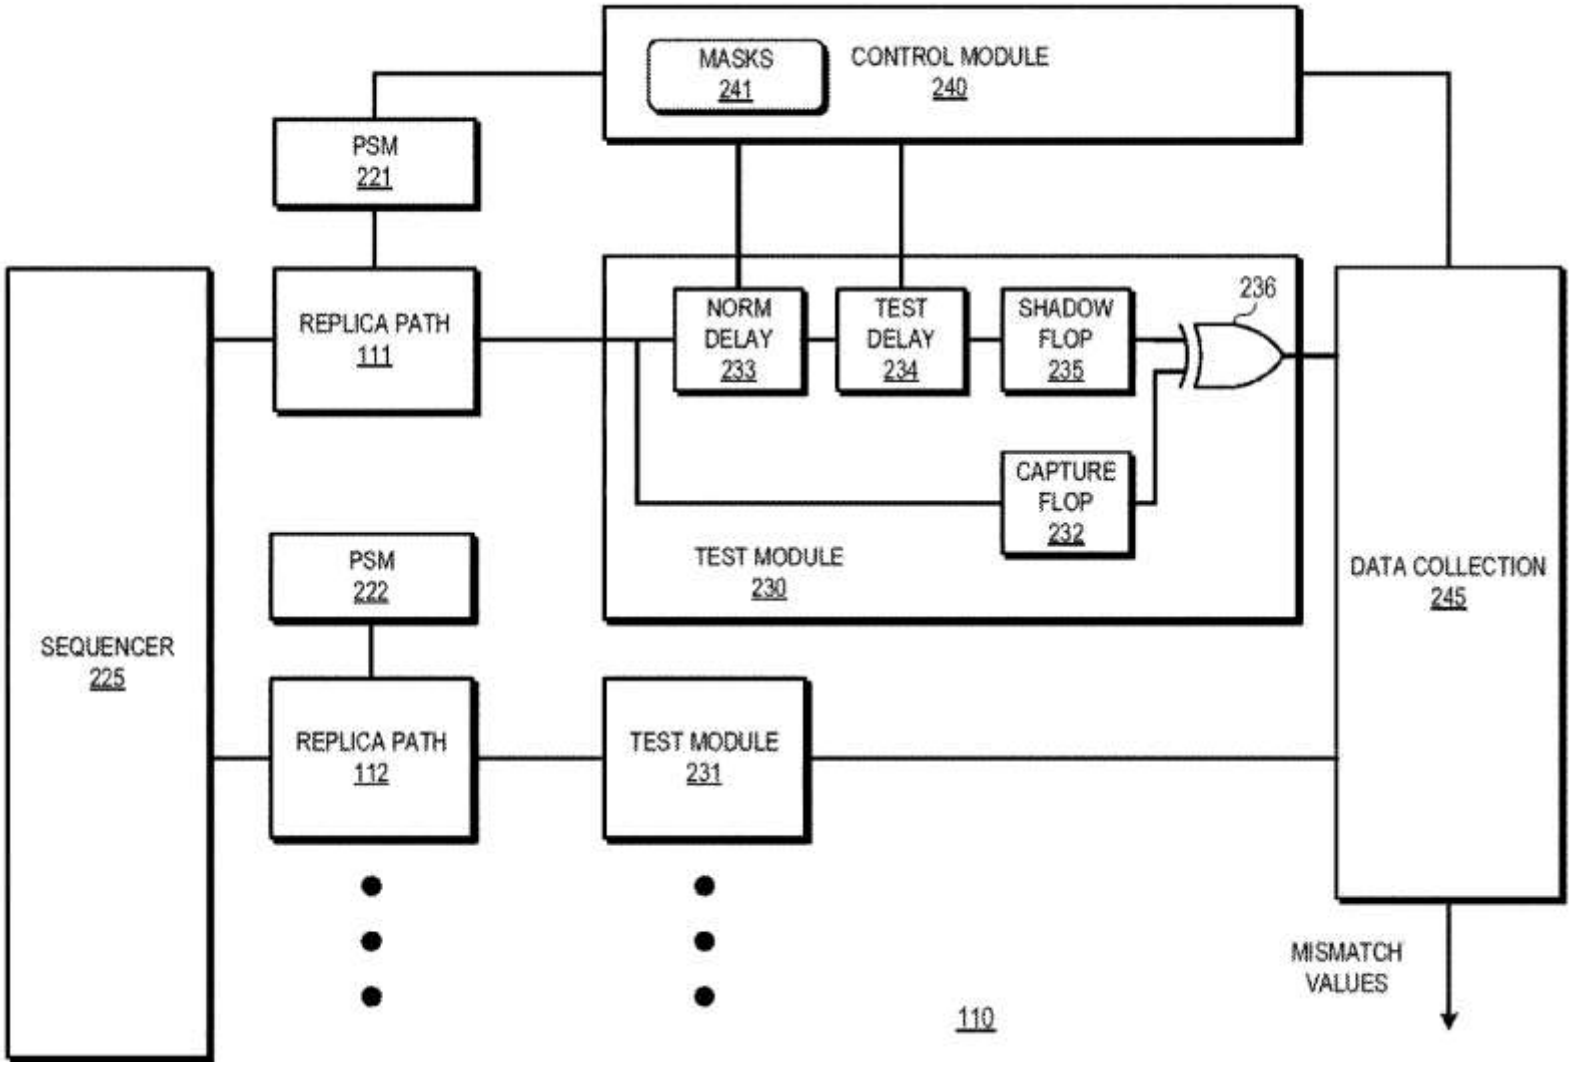

Principle of operation (simplified) -1

- As the above Figure shows, each replica path has both a normalization delay element and a test delay element.
- The normalization delay elements are used only once for a given design (during testing) to set the same delay for each replica path (within a specified tolerance).

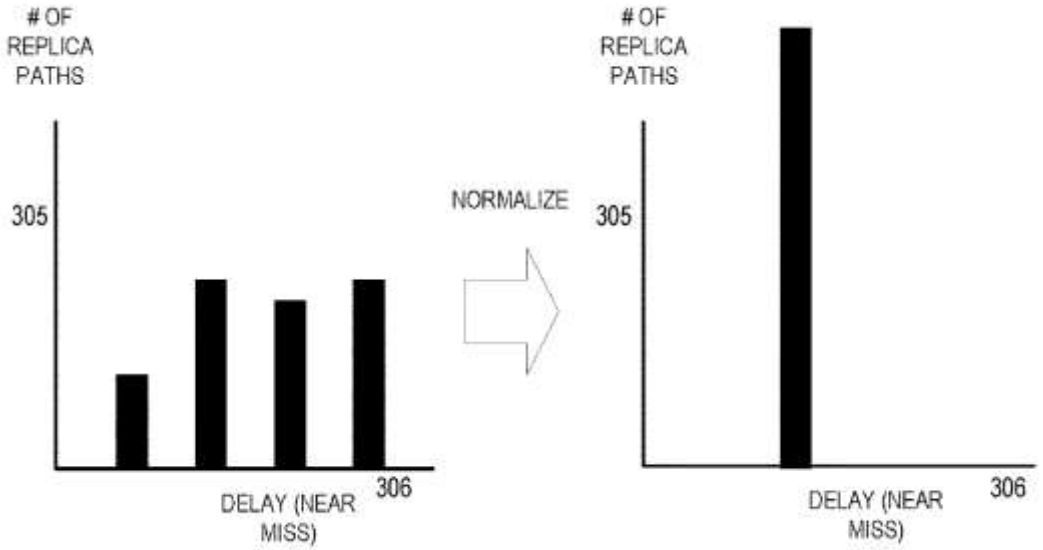

Figure: Distribution of replica path delays before and after normalization [103]

Principle of operation (simplified) -2

- During self calibration the CPA steps over the replica paths connected to it and determines the available timing margin for each path, as follows.
- The Control module (see Figure) increases the delay of each replica path by means of the test delay element until the delay becomes too large, this is noticed when the Shadow flop gates data too late in respect to the Capture flop and the output of the XOR gate becomes "1".
- The delay value that results in a mismatch is referred to a the mismatch value, it indicates the timing margin for the corresponding replica path.
- The CPA forwards a distribution of the mismatch values of the replica paths to the AVS control unit that is in fact a microprocessor responsible for power management.
- The AVS determines an average and standard distribution deviation of the timing margins of the replica paths, as indicated for example in the next Figure.

Figure: Near miss statistics and estimation of the median value [104]

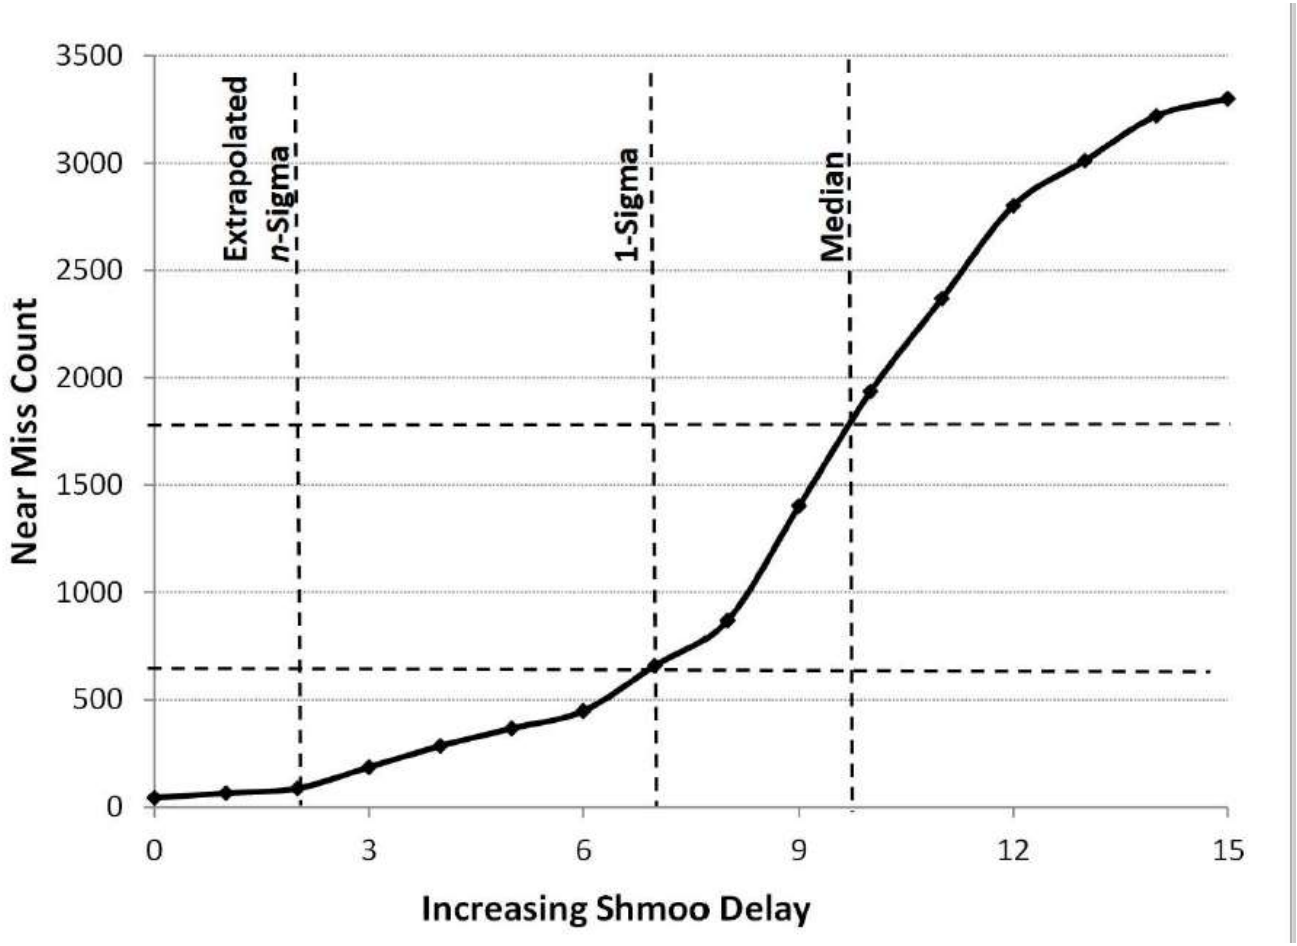

Principle of operation (simplified) -3

- The AVS control module sets the minimum operating voltage for a given operating frequency based on the average and standard deviation of the mismatch values.
- In addition there are Power supply monitors (PSMs) that monitor variations in the VDD, e.g. by determining an average value for the VDD.

  If VDD, as applied to a given replica path, varies from the average VDD value by more than a threshold, the CPA can adjust the distribution of mismatch values based on the variation in VDD.
- Compared to the previous implementation (Steamroller) the Excavator module provides about 40 % power saving due to using AVFS [104].

Use of the high density library [105]

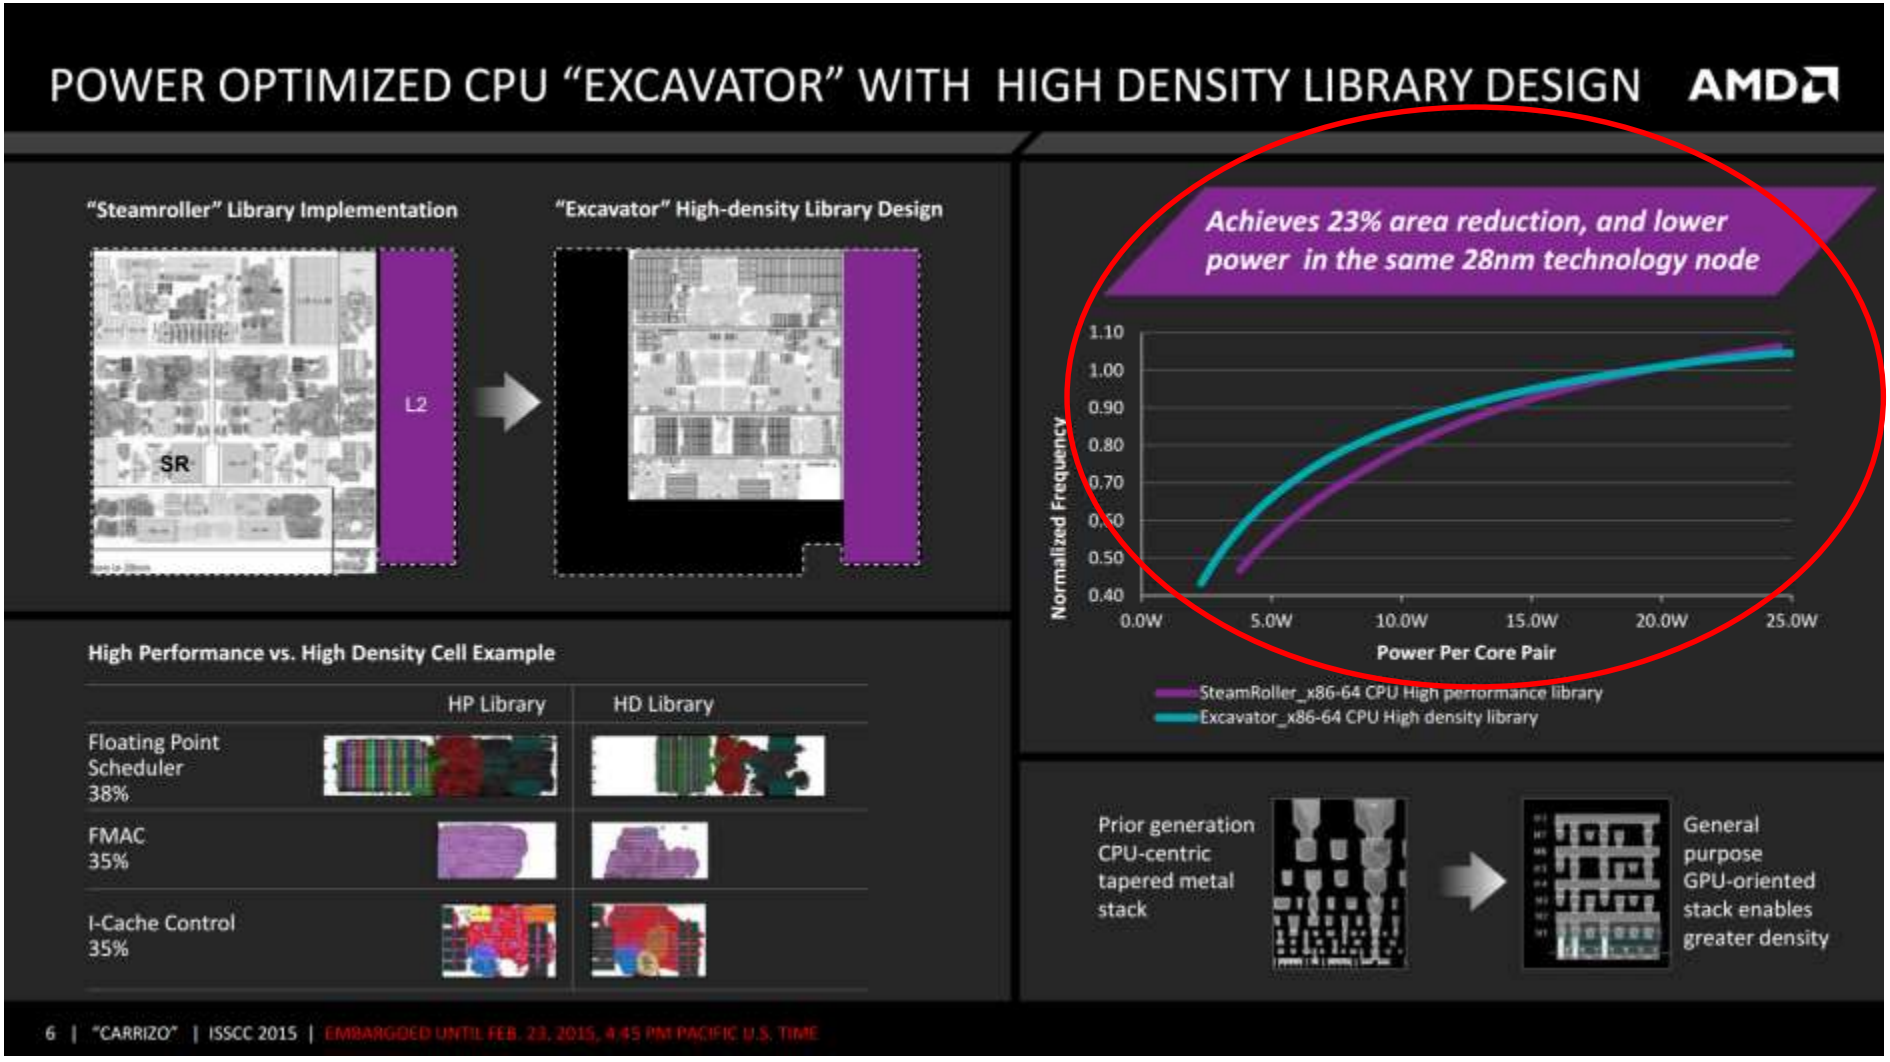

Achieved power reduction at a given performance level [102]

AVFS TO OPTIMIZE PERFORMANCE PER WATT

AMD

- Reliably extract the true silicon speed capability of CPU
  - Includes effects of part-to-part processing, temperature and power delivery
  - Add both a voltage and frequency sensor to existing power and temperature sensors
- Enables accurate setting of the optimal operating point for a given power or performance level across process, voltage and temperature ranges
  - Improved energy efficiency across the entire voltage/temperature operating range

"Excavator" core
incorporates
10 AVFS modules
containing
~500 frequency
sensing paths

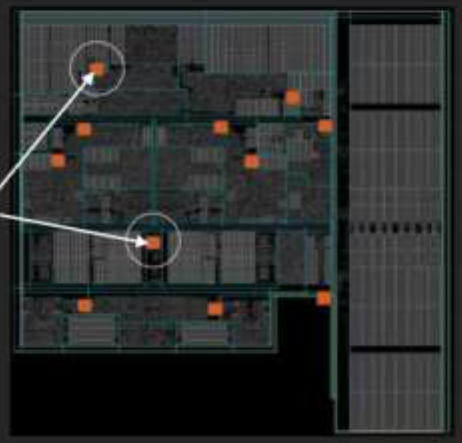

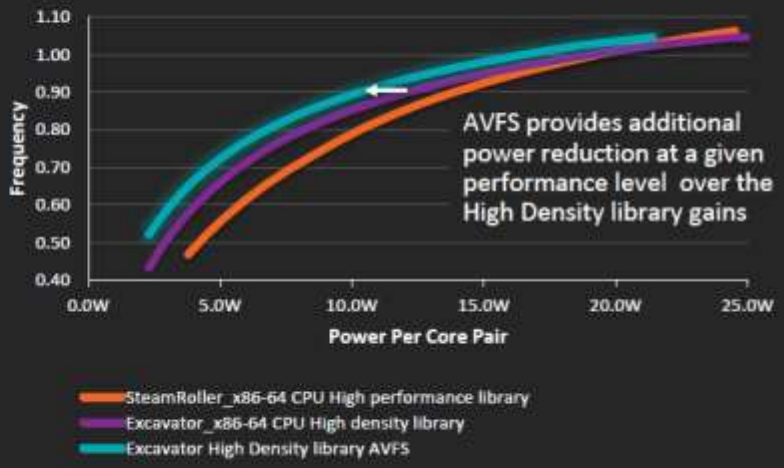

5.3 The Carrizo mobile and desktop line

5.3 The Carrizo mobile line -1

AMD's Family 15h Excavator-based processor lines

Excavator v1-based ultra-thin mobile APU line

Family 15h Models 60h-6Fh

Carrizo (6/2015)
Ultra-thin
mobile APU line

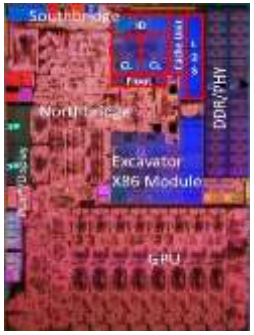

FX8800P/A10/A8 8xxxP
2 modules
R8/R7/R6
2x DDR3-2133
BGA (FP4)
28 nm, 3.1 btrs, 245 mm^2

Excavator v2-based mobile APU lines

Family 15h Models 70h-7Fh

Bristol Ridge (5/2016)
Standard and ultra-thin
mobile APU line

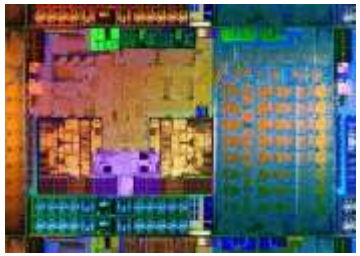

FX98xxP/A12/A10 9xxxP
2 modules
R7/R5 graphics
2x up to DDR4-2400
BGA (FP4)
28 nm, 3.1 btrs, 250 mm^2

Stoney Ridge (5/2016)
Ultra-thin
mobile mobile line

A9/A6 9xxx
1 module
R5/R4 graphics
1x up to DDR4-2133
BGA (FP4)
28 nm, 1.2 btrs, 124 mm^2

The Carrizo mobile line -2

- Introduced in 06/2015
- Manufactured on 28 nm technology, 3.1 btrs, 245 mm^2 die size
- It belongs to the 6. generation APUs.

Key features of the Carrizo mobile APU [105] -1

NEW PERFORMANCE MOBILE APU – "CARRIZO"

AMD

NEXT GENERATION PERFORMANCE APUs
WITH FULL HSA CAPABILITY

- Single, scalable infrastructure shared with "Carrizo-L"
- New "Excavator" core optimized for low power notebook/convertible form factors
- Next Generation AMD Radeon™ Graphics Core Next architecture with support for Mantle, DirectX® 12, and Dual Graphics
- Single-chip integration of the APU and the Southbridge onto a single die
- Significant performance and battery life improvements. First processor in the world with full HSA 1.0 support
- AMD Secure Processor, leveraging ARM® TrustZone technology for Enterprise-class security

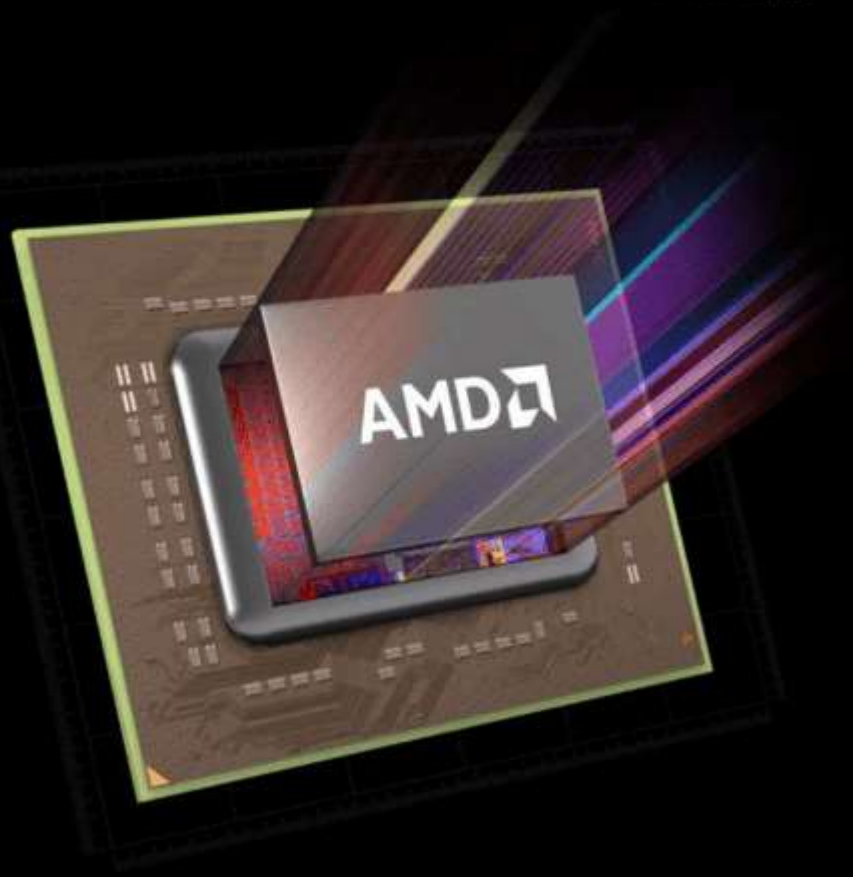

Key features of the Carrizo mobile APU [105] **-2**

NEW PERFORMANCE MOBILE APU – "CARRIZO"

ISSCC 2015 DISCLOSURES

AMD

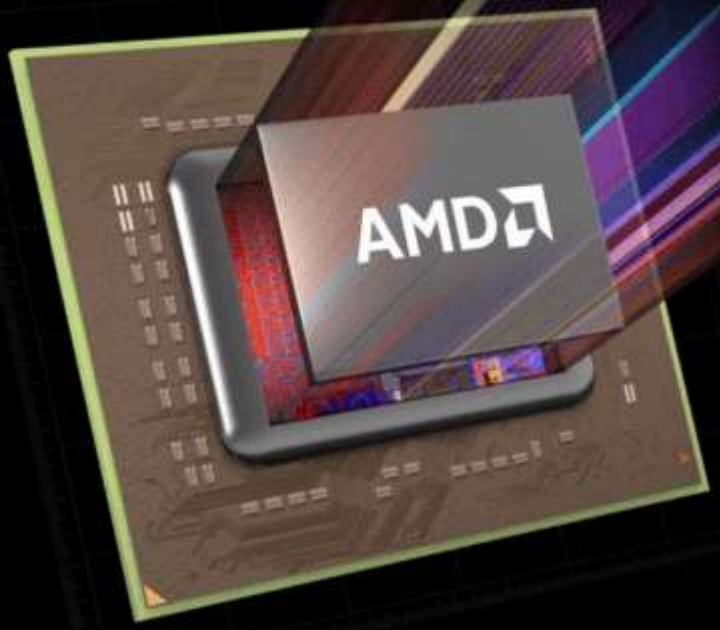

A 28NM X86 APU OPTIMIZED FOR POWER AND AREA EFFICIENCY – SESSION 4.8*

- High density design library resulting in 29% more transistors than "Kaveri" in approximately the same die area
  - 3.1 billion transistors
- Excavator cores: 5% more IPC at 40% less power and 23% less area
- H.265 support and > 3.5x transcode performance of "Kaveri"
- Device selection and implementation tuning enable the eight AMD Radeon™ cores to reduce power 20% from "Kaveri"
- Double digit increases in performance and battery life

*A 28nm x86 APU Optimized for Power and Area Efficiency , presented by Kathryn Wilcox. Session 4.8, *Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*. 2015 ISSCC Conference, February, 2015.

Key features of the Carrizo mobile APU [106] **-3**

Process	28nm	
Package	FP4	
CPU	4 "Excavator" cores / 2MB L2 Cache	
GPU	3rd Gen GCN, 8 Graphics CUs, 2RBs	
Memory	DDR3 Dual-Channel up to 2133	
HSA	Designed to meet Full HSA 1.0 spec	
Integrated Southbridge	Yes	
Display	3 Display Engines, 3 DDI ports	DCE11
Audio	TrueAudio support, Azalia HD Audio or I2S	
Multimedia	UVD6, VCE 3.1 with dual VCE engines, ACP2.1	
PCIE	x8 PCIe® G3/G2	
UART/I2C	UART 2 links; I2C 4 links	
Core Power Supplies	3 Rails – VDD, VDDNB, VDDGraphics	2SVI2 interfaces
Security	AMD Secure Processor / TPM 2.0, crypto acceleration and secure boot	
Software	Windows® 10 and DirectX® 12 Ready	
Streaming Media	HEVC / H.265 Decoding	

Highly enhanced energy efficiency [106] **-1**

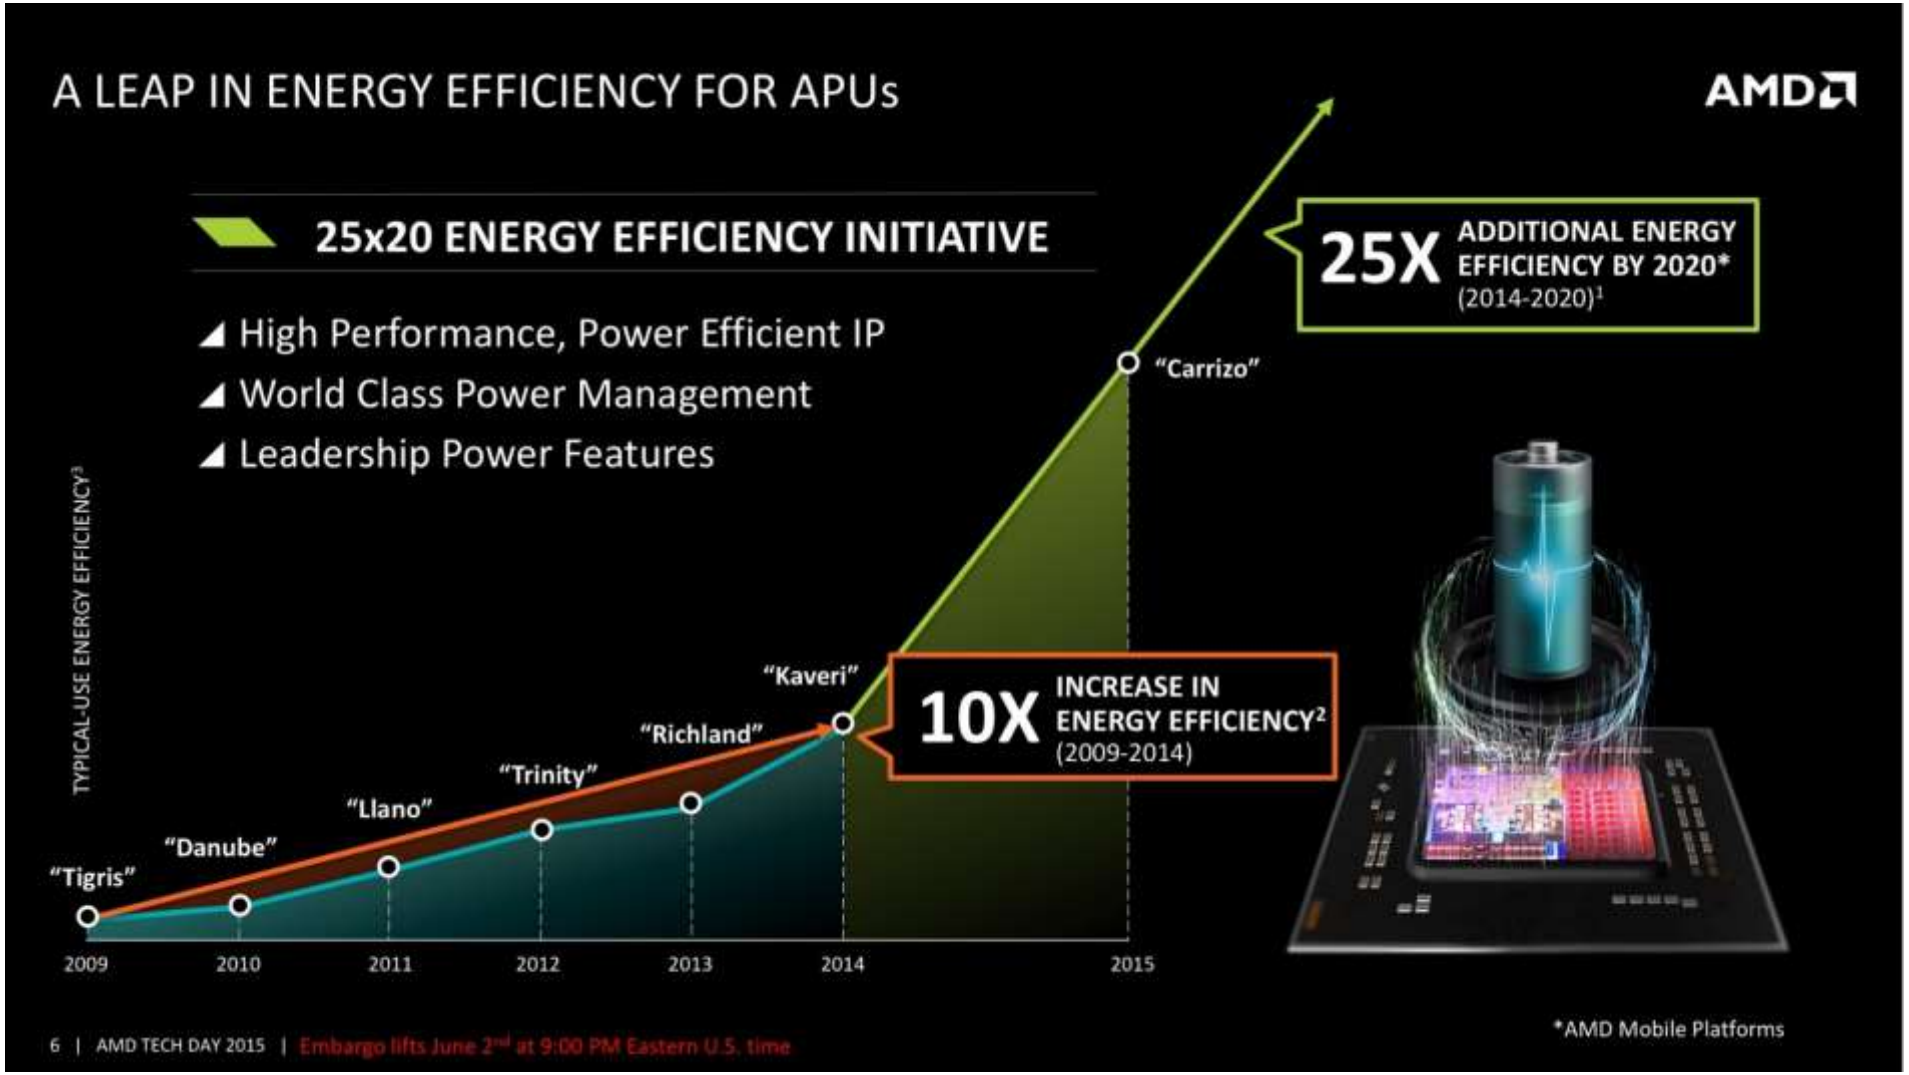

Highly enhanced energy efficiency [106] **-2**

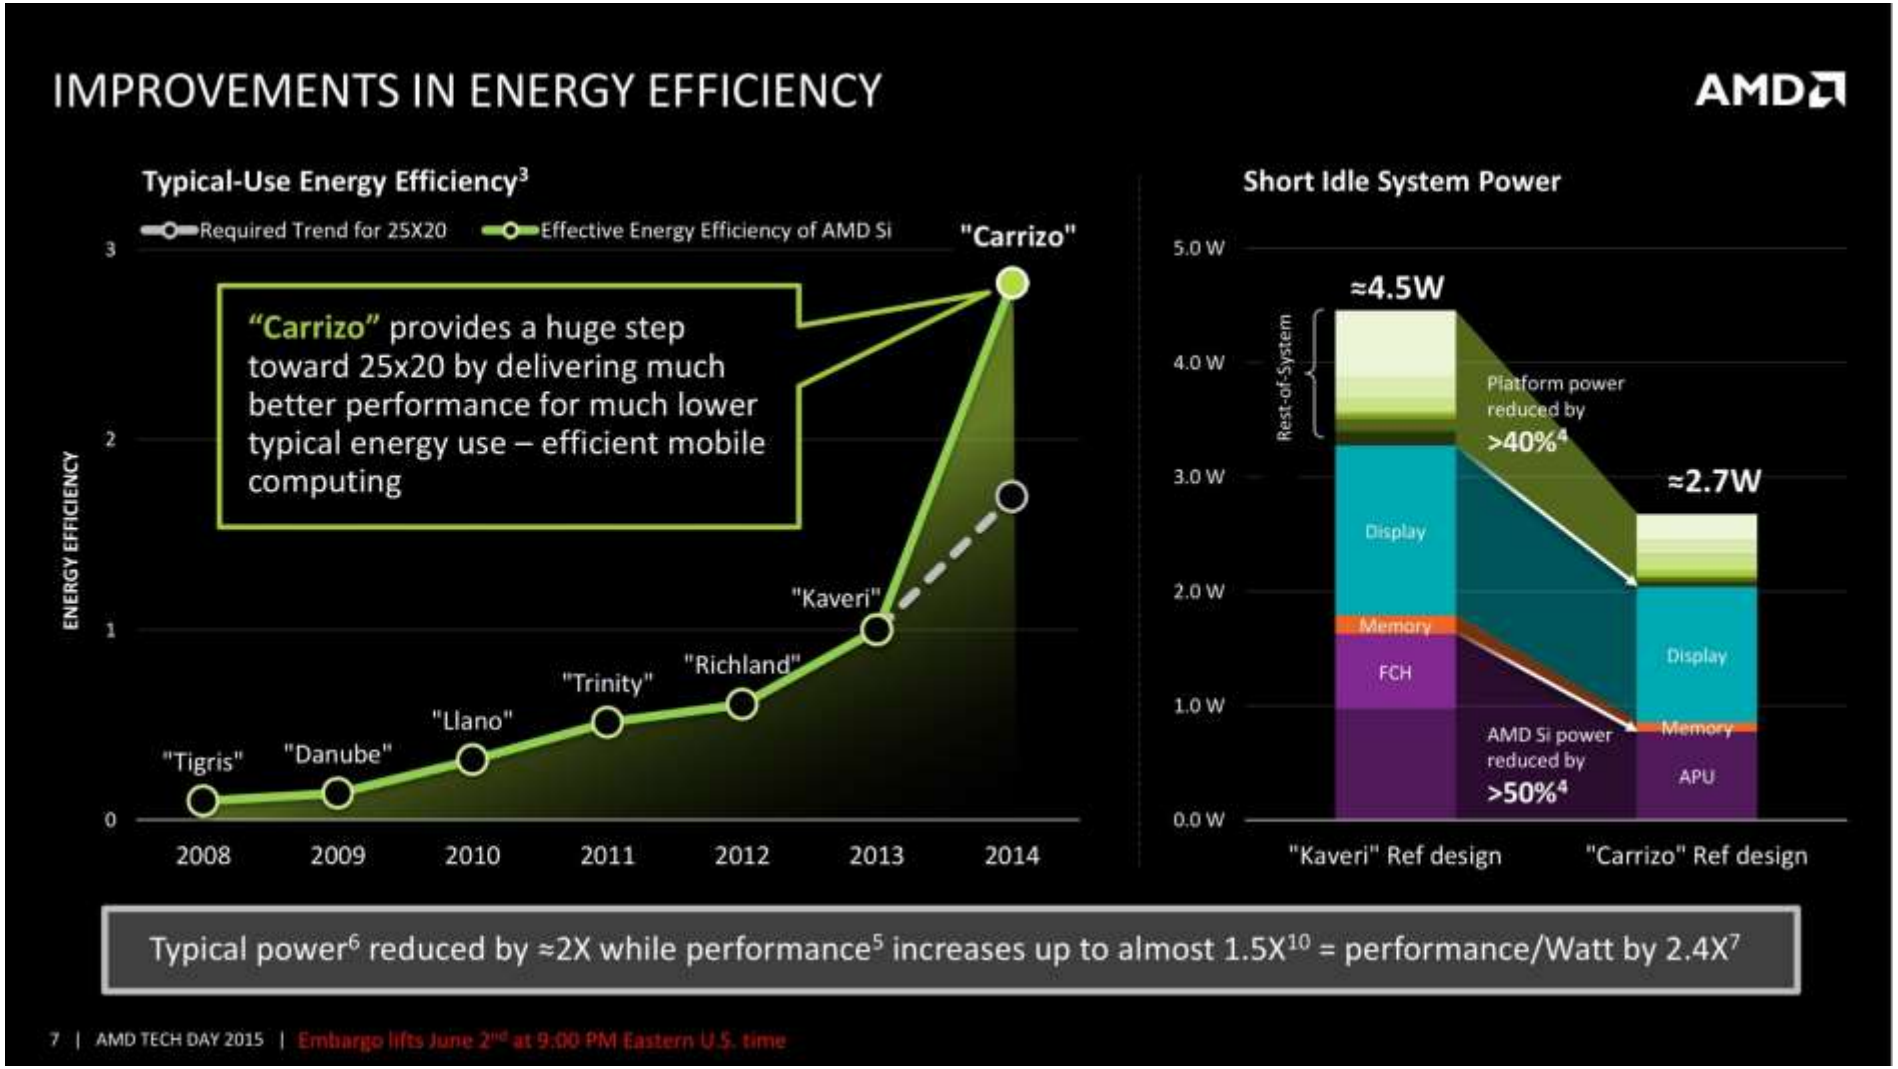

5.3 The Carrizo mobile and desktop line (8)

Main features of the models of the Carrizo mobile line

6th Generation AMD A-Series Processors						
Model	**Radeon™ Brand**	**TDP**	**CPU Cores**	**Compute Cores**	**CPU Clock (Max/Base)**	**L2 Cache**
FX-8800P	R7	12-35W	4	12 (4C+8G)*	Up to 3.4GHz	2MB
A10-8780P Extreme	R8	15W	4	12 (4C+8G)*	Up to 3.3GHz	2MB
A10-8700P	R6	12-35W	4	10 (4C+6G)*	Up to 3.2GHz	2MB
A8-8600P	R6	12-35W	4	10 (4C+6G)*	Up to 3.0GHz	2MB

5.3 The Carrizo mobile and desktop line (9)

Innovations introduced by the Excavator v1. based Carrizo mobile line

a) Skin temperature aware power management (STAM)
b) Full HSA 1.0 support
c) Low-power optimized graphics
d) Support for ARM TrustZone via integrated Cortex-A5 processor

a) Skin Temperature Aware Power Management (STAPM)

This is in fact chassis temperature aware turbo boost technology.

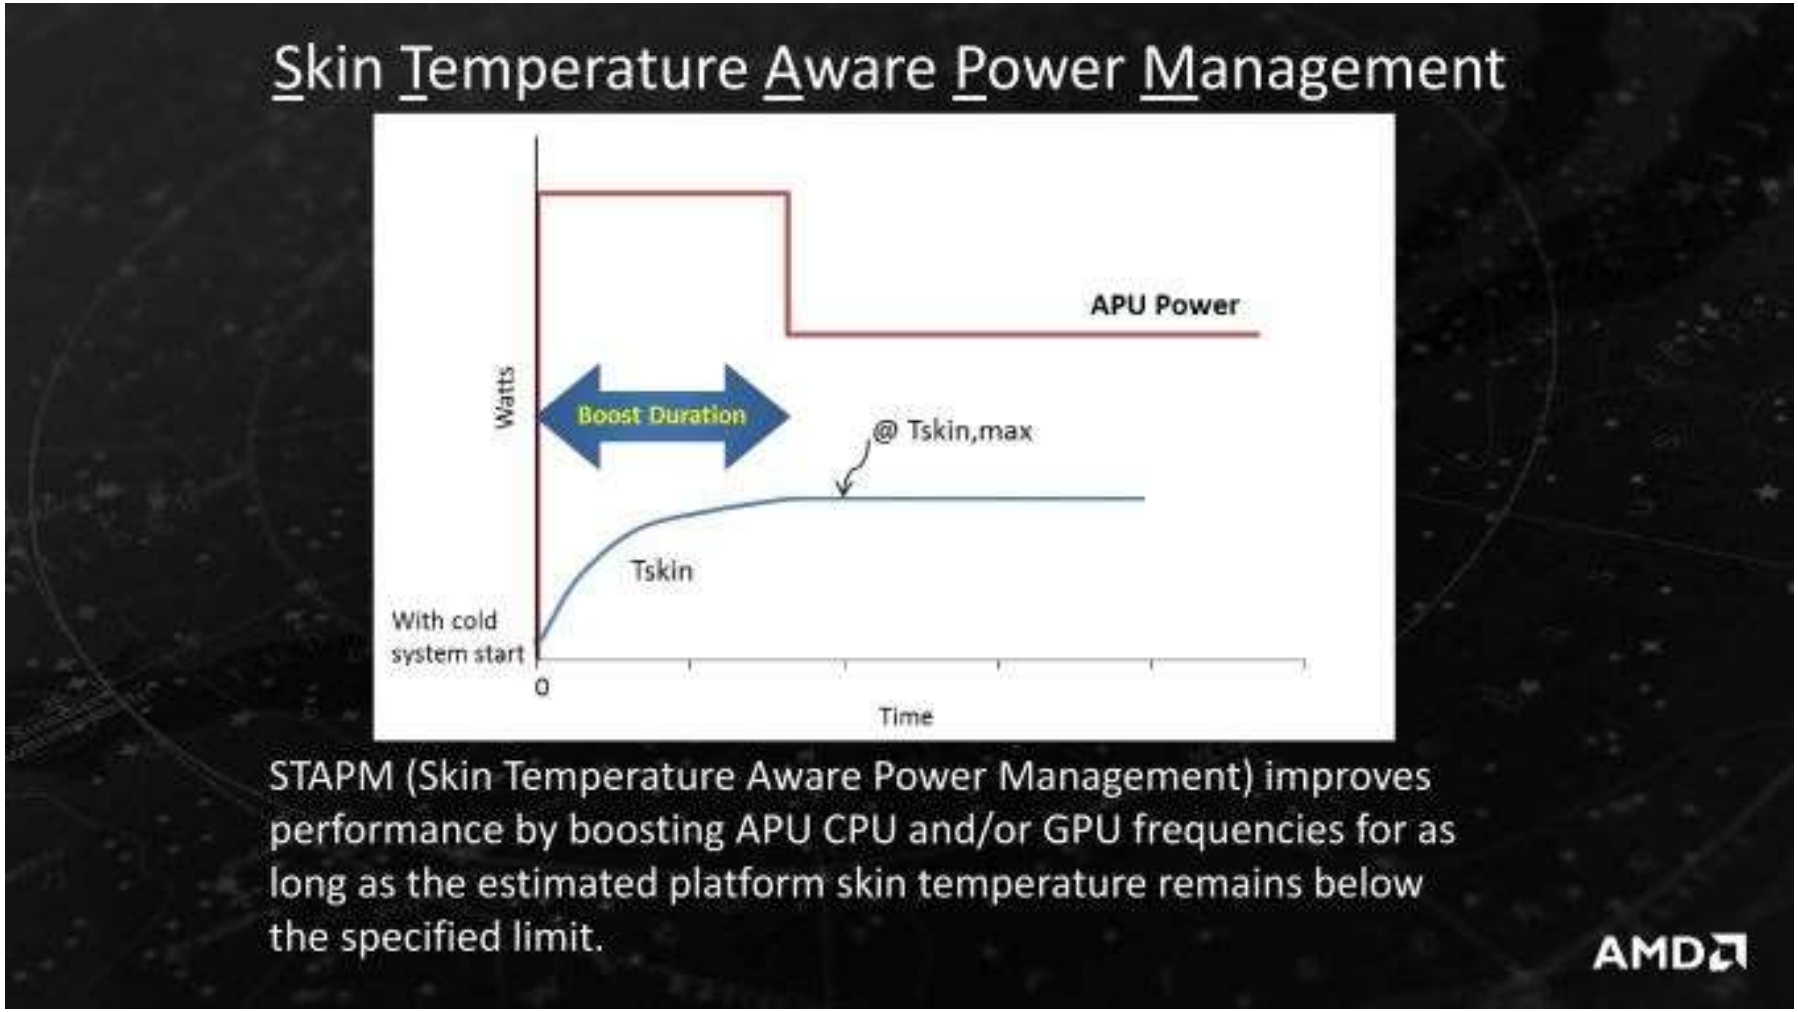

Most use cases for mobile devices are short in duration, this result in many cases in higher performance [107]

5.3 The Carrizo mobile and desktop line (11)

Remark

AMD introduced STAPM already along with the Puma+ based APUs (Beema/Mullins) processor lines (4/2014)

b) Full HSA 1.0 support [106]

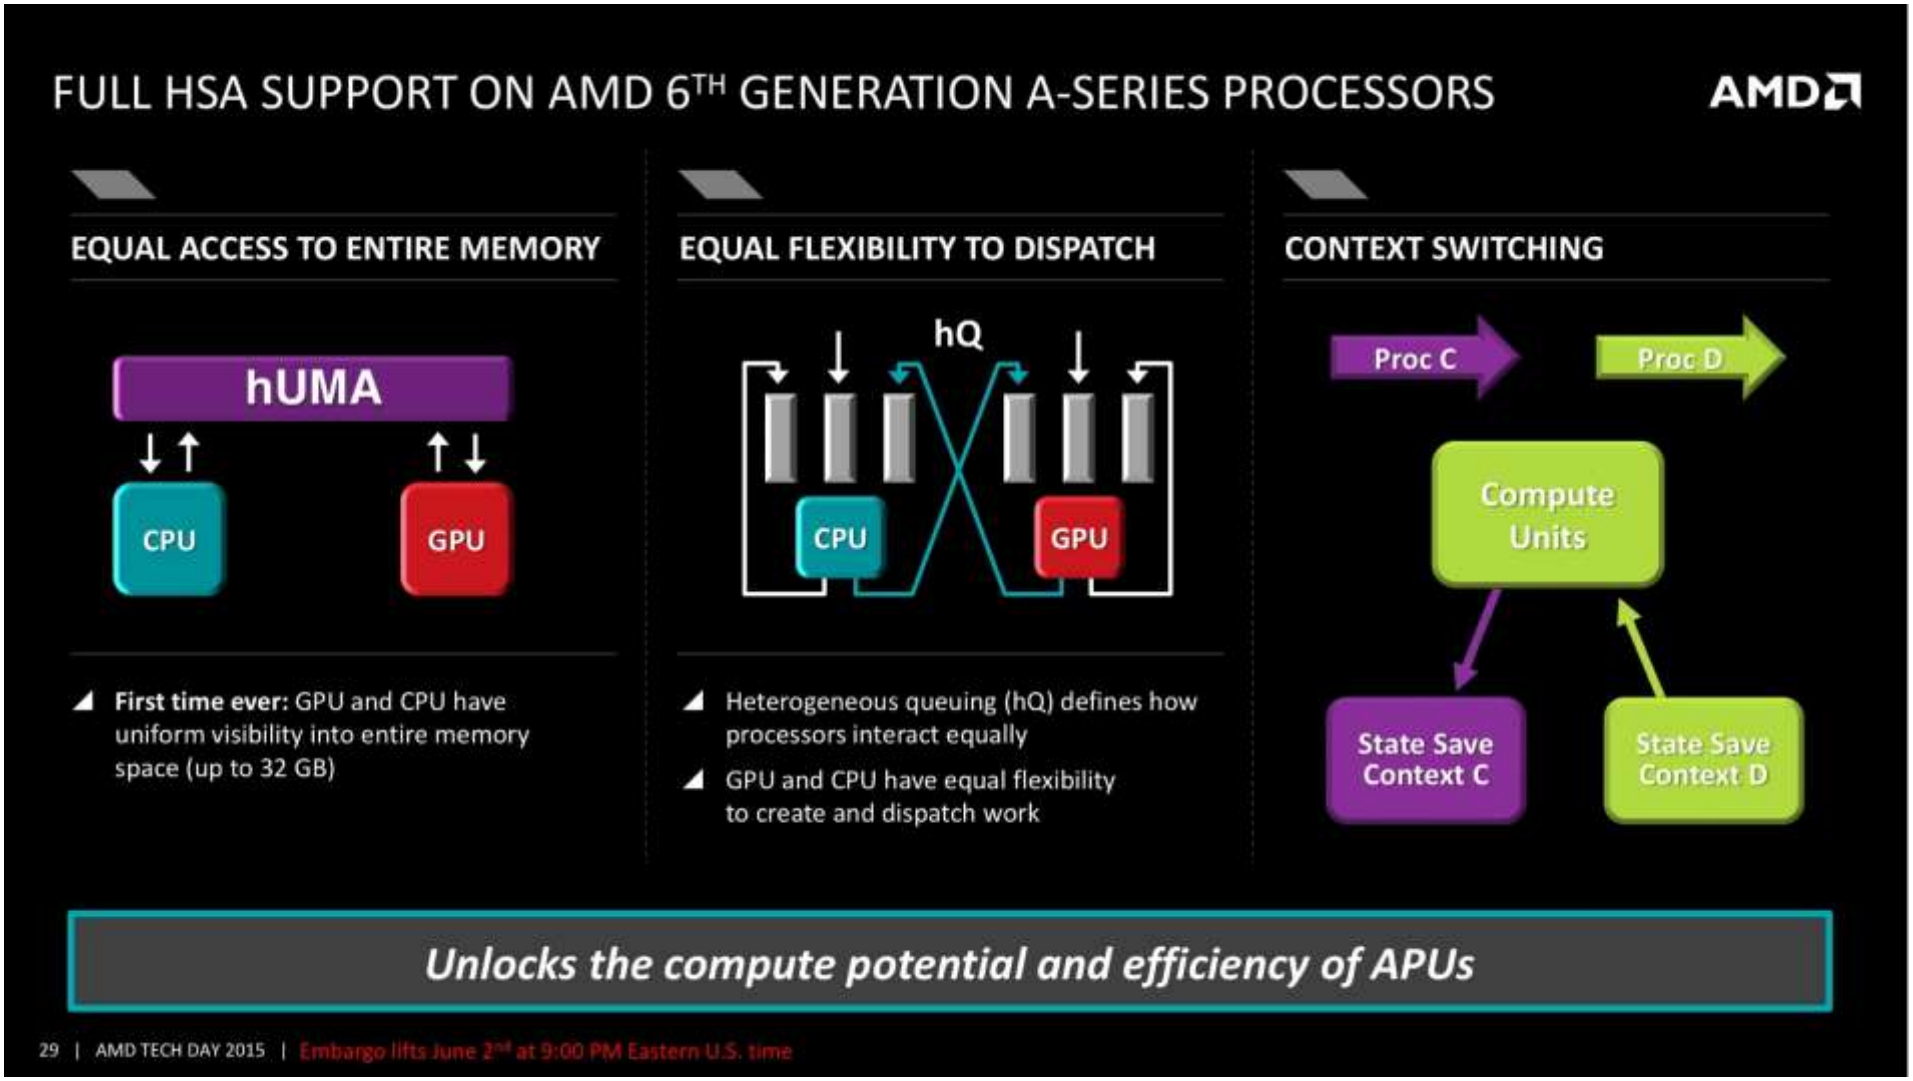

Use of HSA for more energy efficient computing in many workloads [105]

AMD APU ENERGY EFFICIENCY WITH HSA

"CARRIZO" IS THE FIRST FULLY HSA COMPLIANT SOC

AMD

WHAT DOES THIS MEAN FOR POWER?

- Many workloads execute more efficiently using GPU compute resources rather than CPU only
  - E.g. video indexing, natural human interfaces, pattern recognition
- For the same power, much better performance: lower energy per operation → greater efficiency

COMPUTE CAPACITY TREND IN PCs

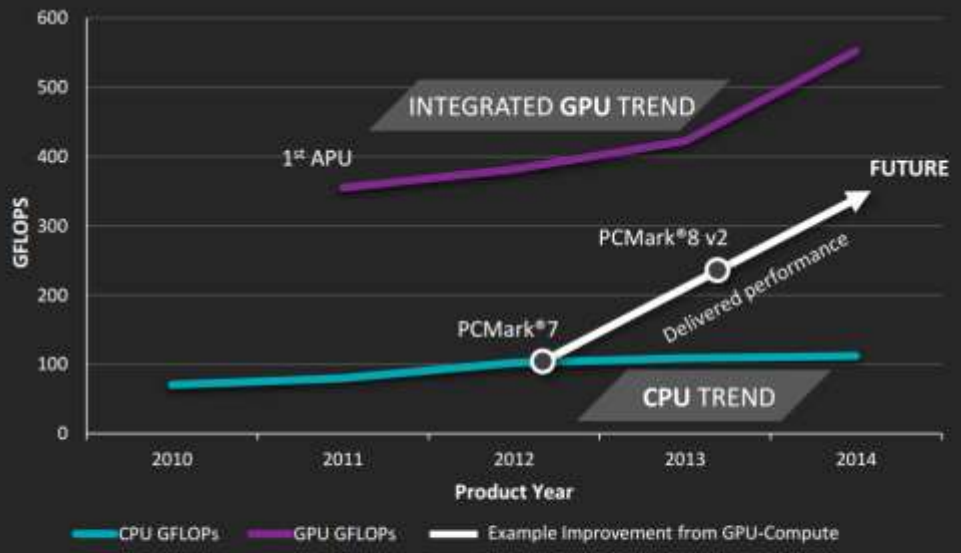

Source: AMD Internal data

c) Low-power optimized graphics [108]

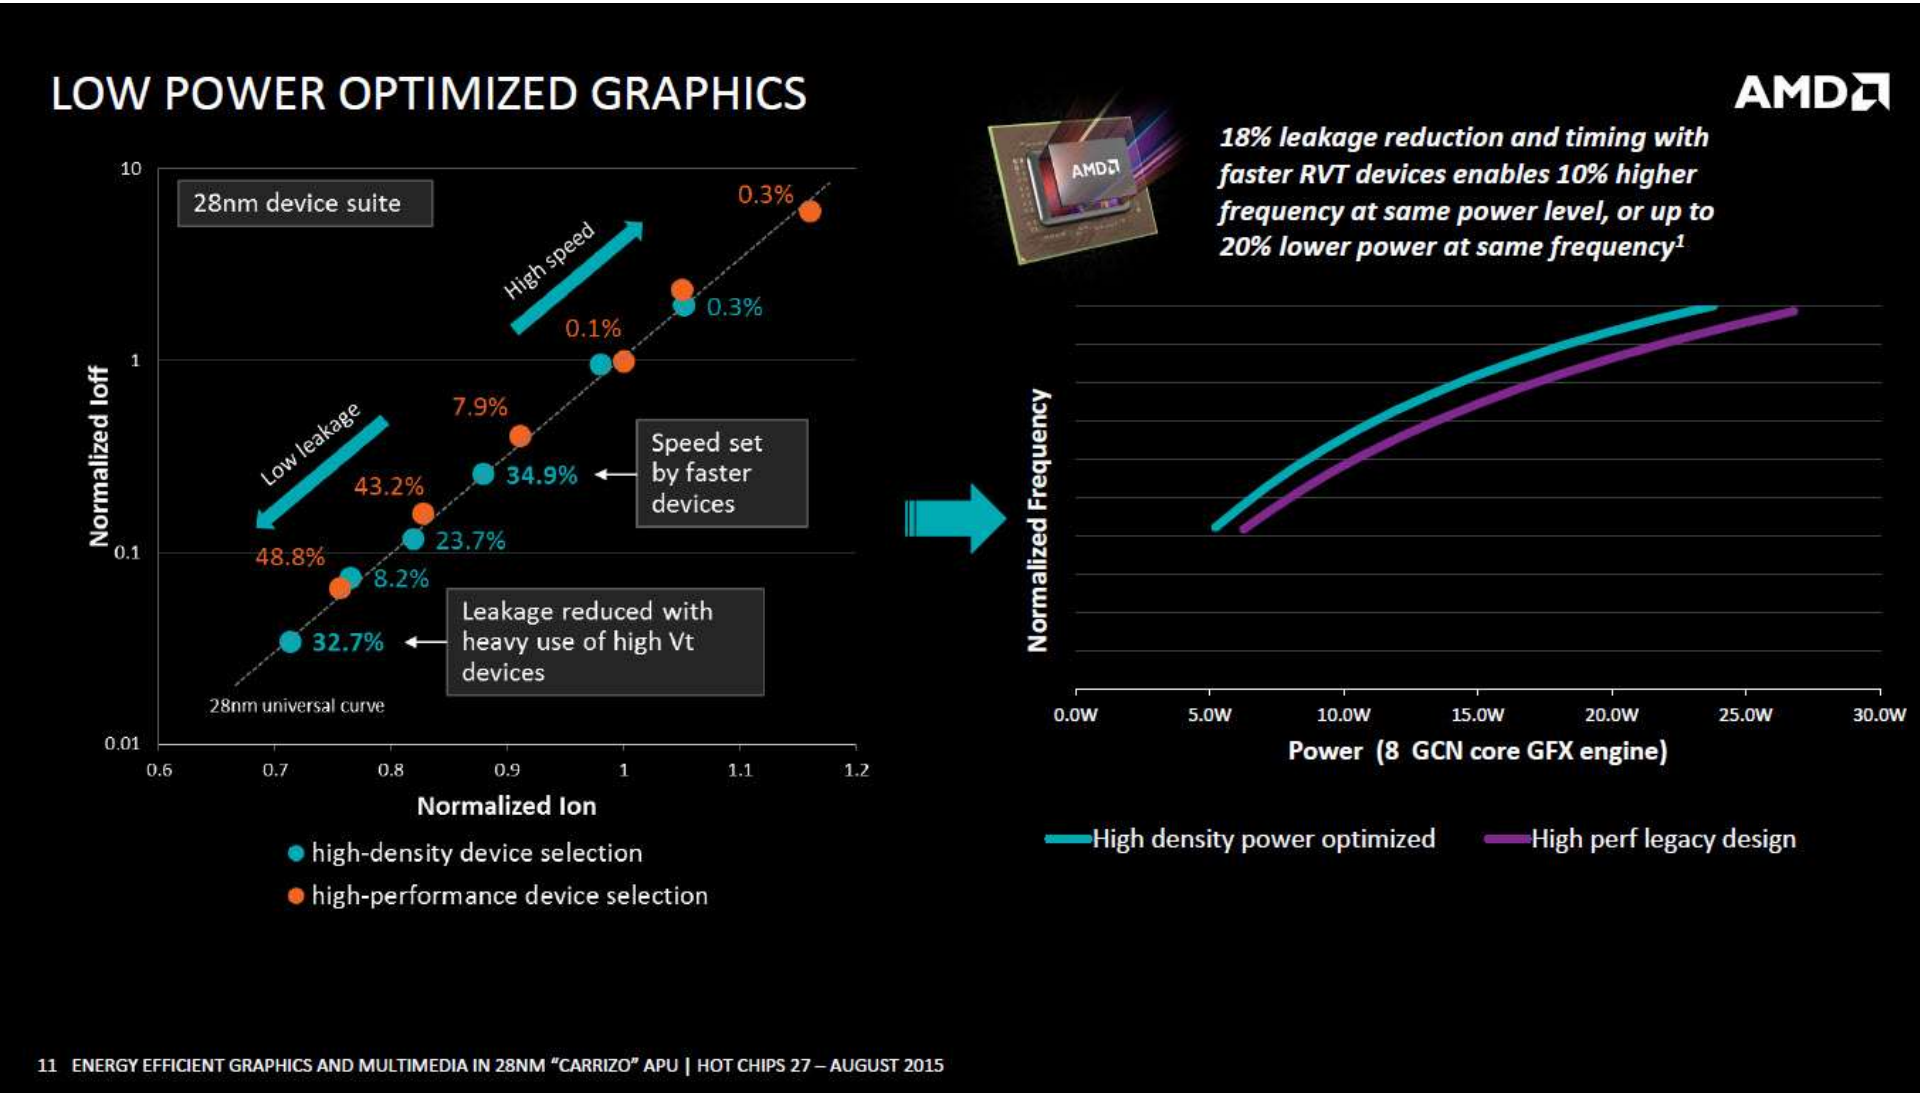

d) Support for ARM TrustZone via integrated Cortex-A5 processor [93] **-1**

- This technology has already been introduced in AMD's Puma+ based Mullins and Beema mobile lines in 04/2014, as documented in the BIOS and Kernel Developer's Guide for AMD Family 16h Models 30h-3Fh .
- It is analogous to Intel's Trusted Computing technology.

FIRST TIME ON PERFORMANCE A-SERIES: AMD SECURE PROCESSOR

AMD

DEDICATED SECURITY SUBSYSTEM INTEGRATED WITHIN AMD 6TH GENERATION A-SERIES PROCESSORS

PSP HARDWARE INCLUDES:

- ~~Dedicated 32-bit microcontroller (ARM® Cortex®-a5)~~
- Isolated on-chip ROM & SRAM
- Access to system memory / resources
- OTP for platform-unique key material
- HW logic for secure control of x86 core boot
- Cryptographic co-processor
  - RSA (up to 16384-bit)
  - Sha (sha1, sha-224, sha-256, sha-512)
  - ECC (basic mathematical computations, up to 384-bit)
  - AES engine (ECB, CBC, CFB, CFB8, OFB, CTR, GCM, CMAC, GMAC, IAPM, XTS-AES128)
  - Zlib (decompression)
  - TRNG (basis for RDRAND)

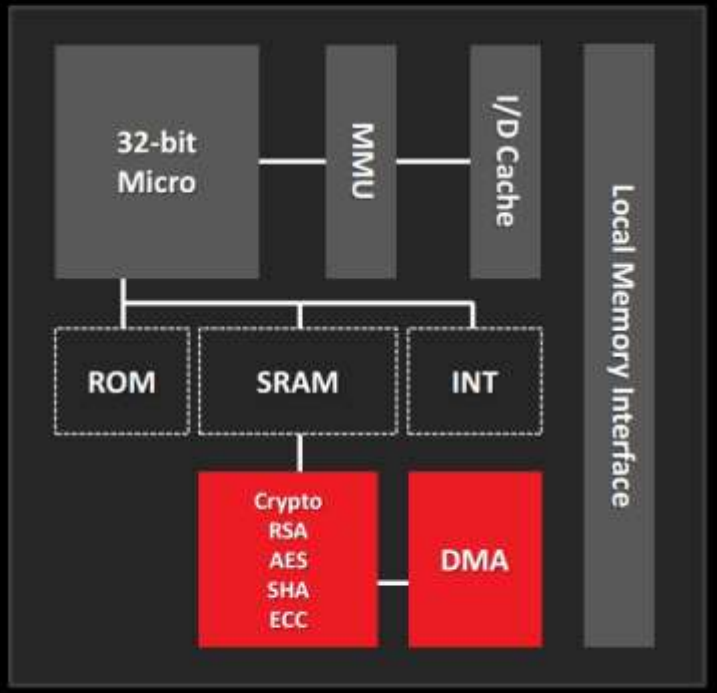

Support for ARM TrustZone via integrated Cortex-A5 processor [109] -2

- Provides a Trusted Execution Environment (TEE)
  - Protects against software attack from open/rich OS side of system
  - Provides scalable environment for secure applications like user authentication, anti-malware, content management, online payments, etc.
- Delivers two separate domains, normal and **secure**
  - Extends across entire system
  - Beyond simply the processor/SOC
  - Can deliver secure
    - Processing data path
    - On/off-chip memory
    - I/O and display

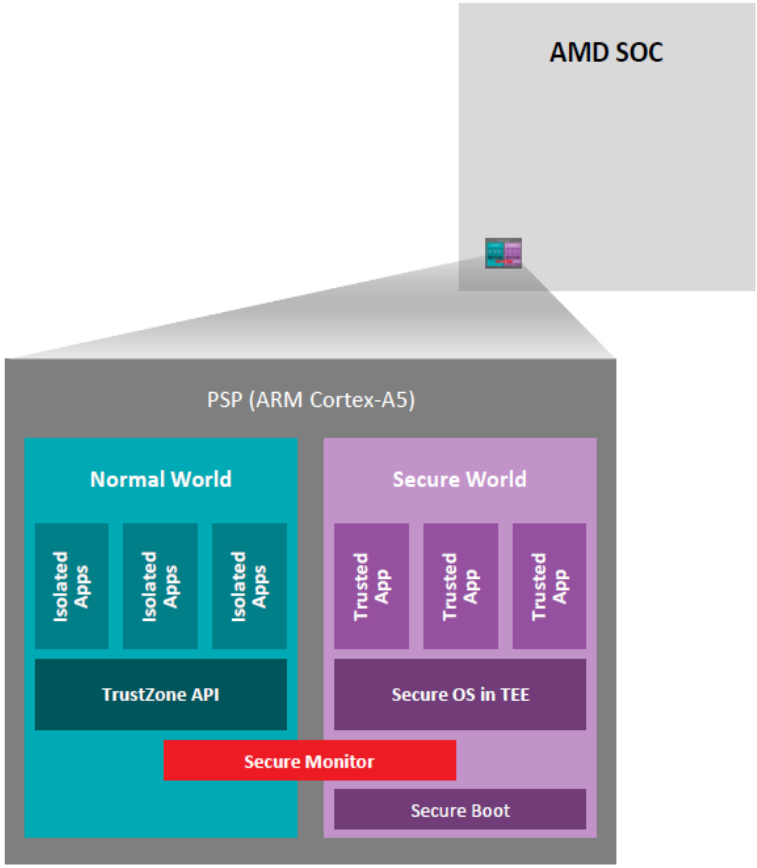

5.4 The Excavator version 2 Compute Module

5.4 The Excavator version 2 Compute Module

Process technology improvements [93]

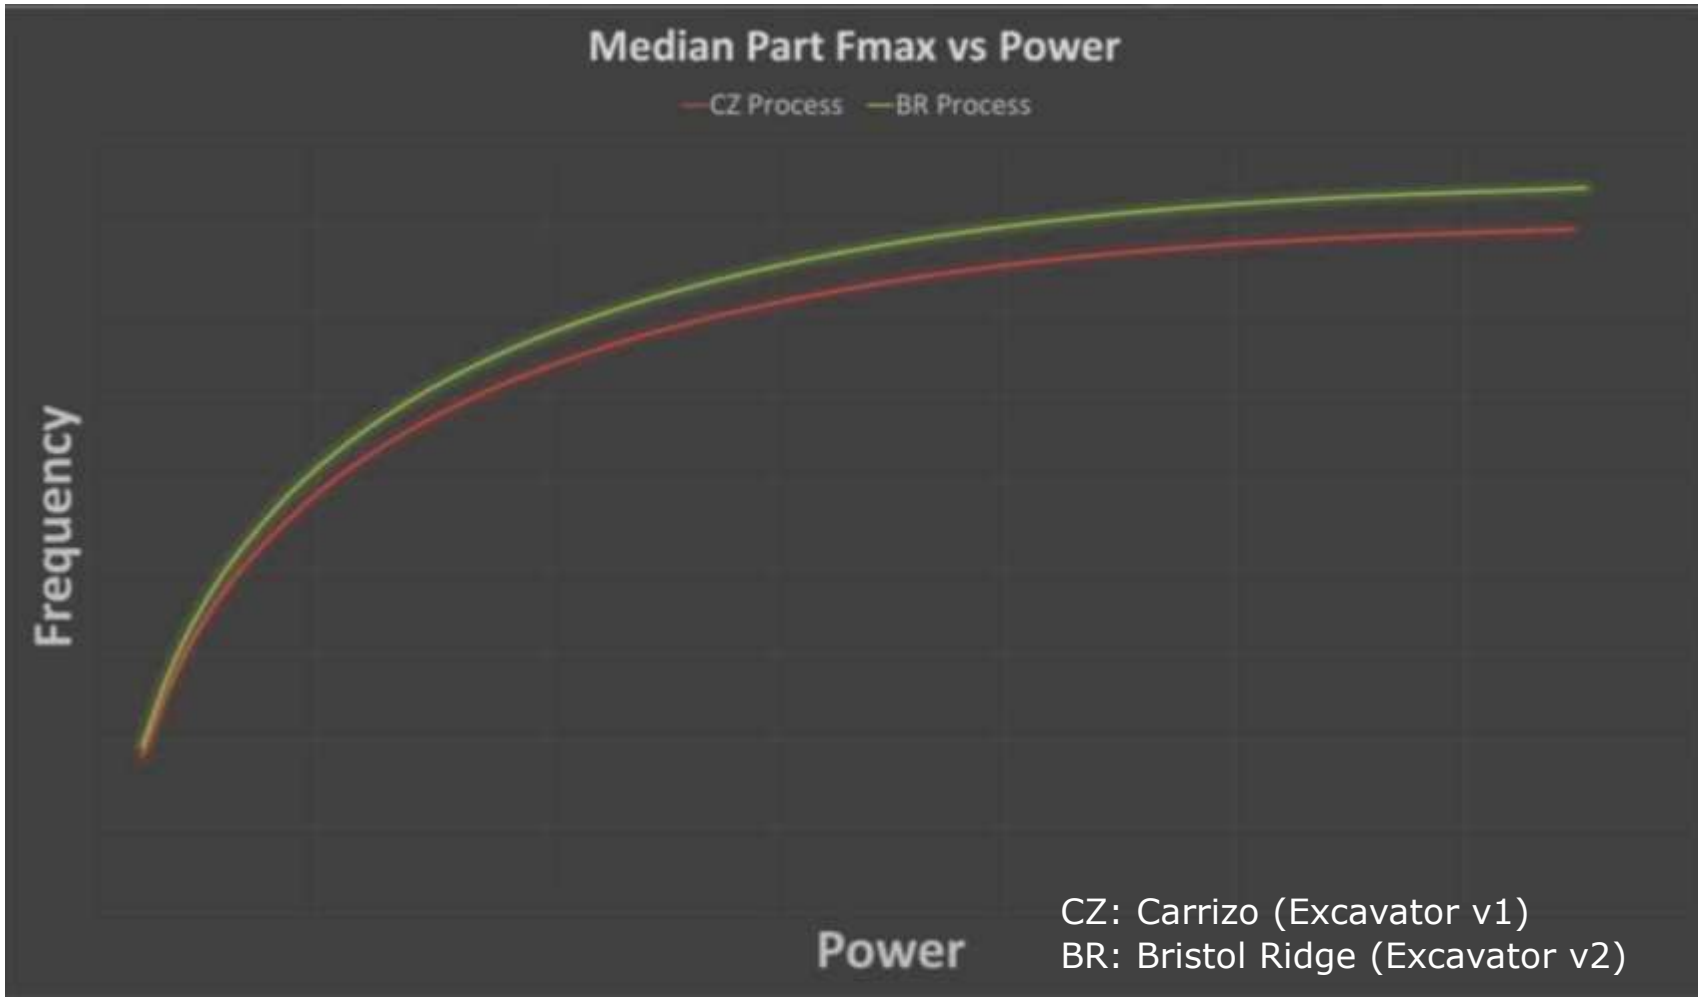

5.5 The Excavator version 2 based processor lines

5.5.1 Overview of the Excavator v2-based processor lines

5.5.1 Overview of the Excavator v2-based processor lines

AMD's Family 15h Excavator-based processor lines

Excavator v1-based ultra-thin mobile APU line

Family 15h Models 60h-6Fh

Carrizo (6/2015)
Ultra-thin
mobile APU line

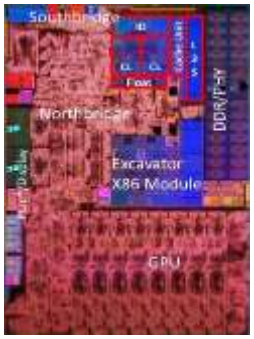

FX8800P/A10/A8 8xxxP
2 modules
R8/R7/R6
2x DDR3-2133
BGA (FP4)
28 nm, 3.1 btrs, 245 mm^2

Excavator v2-based mobile APU lines

Family 15h Models 70h-7Fh

Bristol Ridge (5/2016)
Standard and ultra-thin
mobile APU line

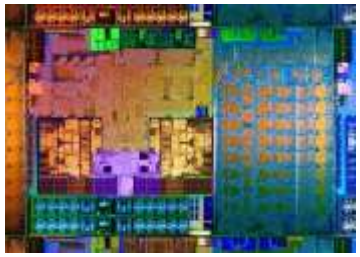

FX98xxP/A12/A10 9xxxP
2 modules
R7/R5 graphics
2x up to DDR4-2400
BGA (FP4)
28 nm, 3.1 btrs, 250 mm^2

Stoney Ridge (5/2016)
Ultra-thin
mobile mobile line

A9/A6 9xxx
1 module
R5/R4 graphics
1x up to DDR4-2133
BGA (FP4)
28 nm, 1.2 btrs, 124 mm^2

5.5.2 The Bristol Ridge mobile and DT line

5.5.2 The Bristol Ridge mobile APU line -1

AMD's Family 15h Excavator-based processor lines

Excavator v1-based ultra-thin mobile APU line

Family 15h Models 60h-6Fh

Carrizo (6/2015)
Ultra-thin mobile APU line

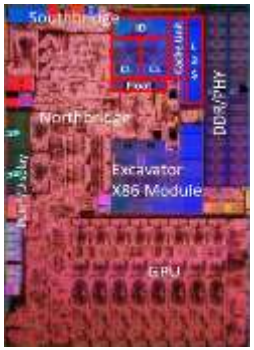

FX8800P/A10/A8 8xxxP
2 modules
R8/R7/R6
2x DDR3-2133
BGA (FP4)
28 nm, 3.1 btrs, 245 mm^2

Excavator v2-based mobile APU lines

Family 15h Models 70h-7Fh

Bristol Ridge (5/2016)
Standard and ultra-thin mobile APU line

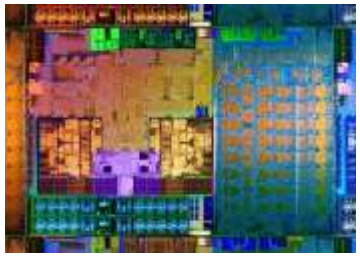

FX98xxP/A12/A10 9xxxP
2 modules
R7/R5 graphics
2x up to DDR4-2400
BGA (FP4)
28 nm, 3.1 btrs, 250 mm^2

Stoney Ridge
(5/2016)
Ultra-thin
mobile APU line

A9/A6 9xxx
1 module
R5/R4 graphics
1x up to DDR4-2133
BGA (FP4)
28 nm, 1.2 btrs, 124 mm^2

5.5.2 The Bristol Ridge mobile and DT line (2)

The Bristol Ridge mobile APU line -2

- Introduced in 05/2016 first only to OEMs, then in 07/2017 also for retail
- Manufactured on 28 nm technology, 3.1 btrs, 250 mm2 die size
- It belongs to the 7. generation APUs.
- It provides ~ 20% boost in CPU performance and 37 % increase in GPU performance over the predecessor Carrizo processor.

5.5.2 The Bristol Ridge mobile and DT line (3)

Main features of the Bristol Ridge mobile APU line [110]

7th-Generation Bristol Ridge A-Series	FX 9830P	FX 9800P	A12-9730P	A12-9700P	A10-9630P	A10-9600P
Launched	Q2/2016	Q2/2016	Q2/2016	Q2/2016	Q2/2016	Q2/2016
Radeon Graphics	R7 Graphics	R7 Graphics	R7 Graphics	R7 Graphics	R5 Graphics	R5 Graphics
CPU Cores	4	4	4	4	4	4
Max/Base CPU Frequency (GHz)	3.7 / 3.0	3.6 / 2.7	3.5 / 2.8	3.4 / 2.5	3.3 / 2.6	3.3 / 2.4
Graphics Cores	8	8	6	6	6	6
Process	28nm	28nm	28nm	28nm	28nm	28nm
DDR4 Dual Channel Memory Support	DDR4-2400	DDR4-1866	DDR4-2400	DDR4-1866	2400 MHz	DDR!-1866
TDP	35W	15W	35W	15W	35W	15W
Configurable TDP Range	25-45W	12-15W	25-45W	12-15W	25-45WE	12-15W

Chipset features supporting the Excavator v2-based Bristol Ridge APU processor [111]

7TH GEN AMD APU & SOCKET AM4 CHIPSET I/O

AMD

PROVIDING THE I/O YOU WANT – NATIVE USB 3.1 GEN2 SUPPORT

Segment	7th Gen APU Processor Features					Chipset Features				
	AM4 CPU	PCI Express® Gen3	DDR4	USB 3.1 G2 + 3.1 G1 + 2.0	Storage & GPP PCIe G3	Chipset	USB 3.1 G2 + 3.1 G1 + 2.0	SATA + SATA Express	PCI Express® Gen 2 General Purpose	SATA RAID
Mainstream	7th Gen AMD APU	x8 Gen3	2CH	0+4+0	2 SATA + x2 NVMe or 2 SATA + x2 PCIe®	B350	2+2+6	2+1	6 Lanes Gen2	0,1,10
						A320	1+2+6	2+1	4 Lanes Gen2	0,1,10
SFF Options	SoC Capabilities as described above					X/B/A300	---	---	---	0,1

Notes: Features are preliminary and subject to change without notice. Customer should always consult the latest technical documentation for design and product specifications

Generational performance improvements [112]

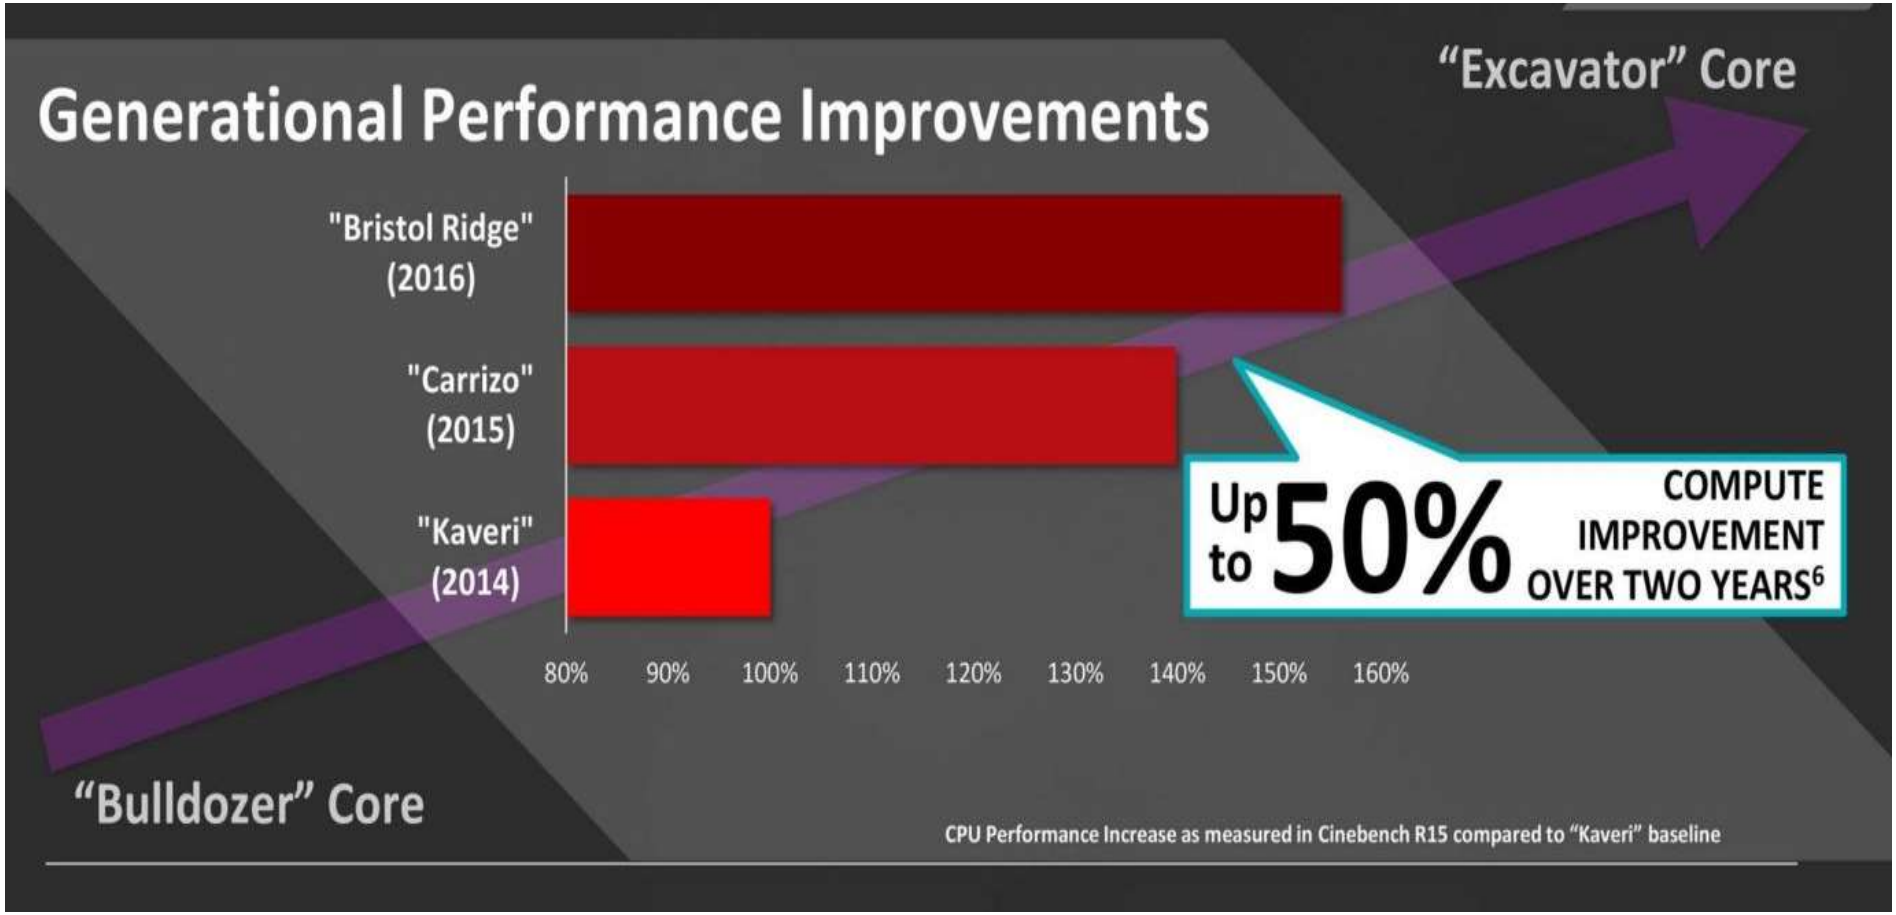

5.5.2 The Bristol Ridge mobile and DT line (6)

Cinebench R11.5 benchmark results for the A10-9600P [113]

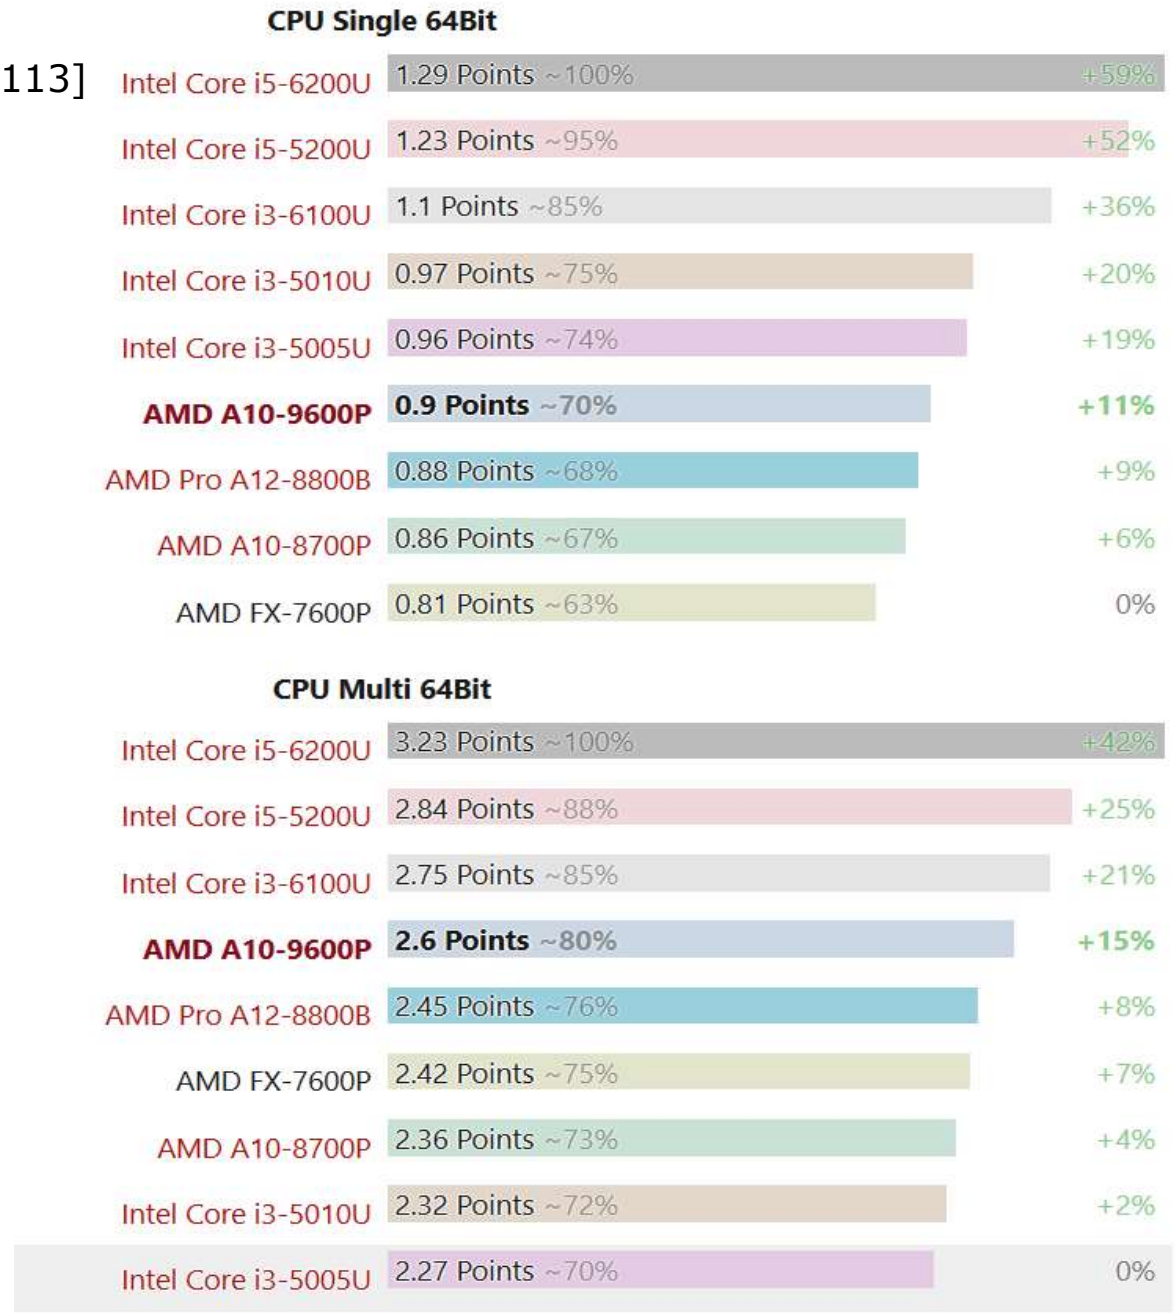

5.5.3 The Stoney Ridge mobile line

5.5.3 The Stoney Ridge mobile line -1

AMD's Family 15h Excavator-based processor lines

Excavator v1-based ultra-thin mobile APU line

Family 15h Models 60h-6Fh

Excavator v2-based mobile APU lines

Family 15h Models 70h-7Fh

Carrizo (6/2015)
Ultra-thin
mobile APU line

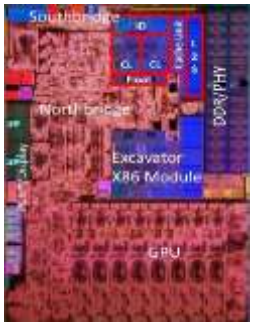

FX8800P/A10/A8 8xxxP
2 modules
R8/R7/R6
2x DDR3-2133
BGA (FP4)
28 nm, 3.1 btrs, 245 mm^2

Bristol Ridge (5/2016)
Standard and ultra-thin
mobile APU line

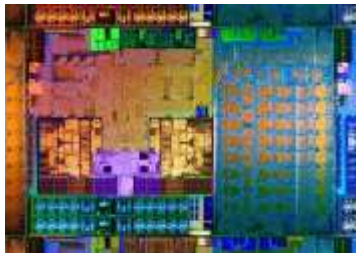

FX98xxP/A12/A10 9xxxP
2 modules
R7/R5 graphics
2x up to DDR4-2400
BGA (FP4)
28 nm, 3.1 btrs, 250 mm^2

Stoney Ridge (5/2016)
Ultra-thin
mobile mobile line

A9/A6 9xxx
1 module
R5/R4 graphics
1x up to DDR4-2133
BGA (FP4)
28 nm, 1.2 btrs, 124 mm^2

The Stoney Ridge mobile line -2

- Introduced in 05/2016
- Manufactured on 28 nm technology, 1.2 btrs, 124 mm2 die size
- It belongs to the 7. generation APUs.
- It includes only a single compute module with 1 MB L2 cache and a single memory channel (up to DDR4-2133)

5.5.3 The Stoney Ridge mobile line (3)

Main features of the models of the Stoney Ridge line launched in Q2/2016 [110]

6th-Generation Carrizo A-Series	A9-9410	A6-9210	E2-9010
Radeon Graphics	R5 Graphics	R4 Graphics	R2 Graphics
CPU Cores	2	2	2
Max/Base CPU Frequency (GHz)	3.5 / 2.9	2.8 / 2.4	2.2 / 2.0
Graphics Cores	3	3	2
Process	28nm	28nm	28nm
DDR4 Single Channel Memory Support	2133 MHz	2133 MHz	2133 MHz
TDP	15W	15W	15W
Configurable TDP Range	10-25	10-15	10-15

5.5.3 The Stoney Ridge mobile line (4)

Main features of the models of the Stoney Ridge line launched in Q2/2017 [113]

	A4-Series for Notebooks	A6-Series for Notebooks	A9-Series for Notebooks
Manufacturing process	0.028 micron		
Cores	2		
Frequency (MHz)	2200	2000 - 2500	2400 - 3000
Boost Frequency (MHz)	Up to 2500	Up to 2900	Up to 3600
Fastest processor	A4-9120	A6-9220	A9-9420
L2 cache size	1024 KB		
L3 cache size	None		
No. of GPU cores	2	3	3
Thermal Design Power (Watt)	15	10 - 15	25
Package	micro-BGA		
Socket	BGA	BGA BGA (FP4)	

5.5.3 The Stoney Ridge mobile line (5)

Integer performance results for the GeekBench 3 benchmark [115]

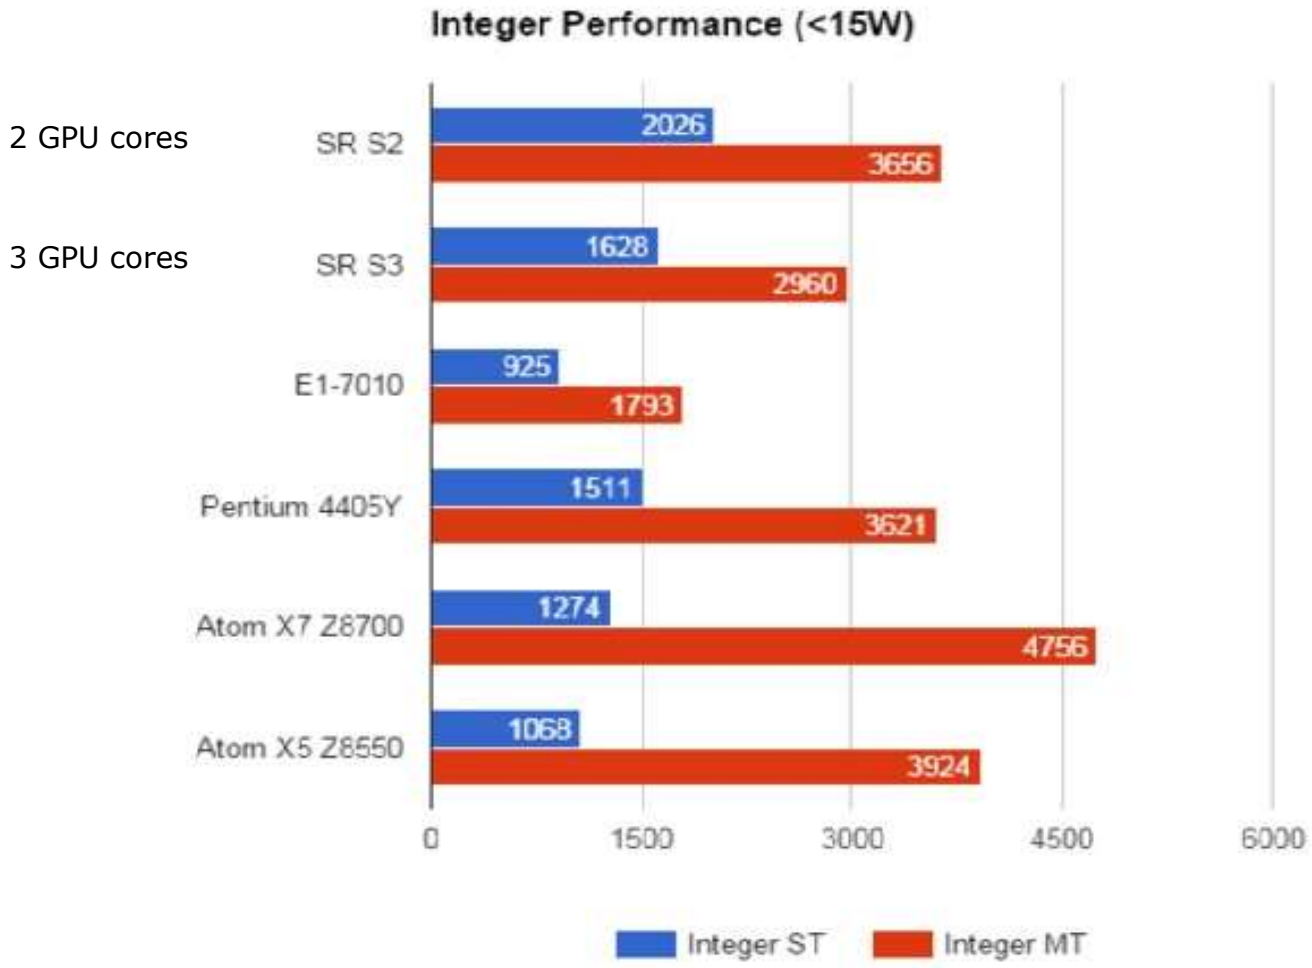

5.5.3 The Stoney Ridge mobile line (6)

Floating point performance results for the GeekBench 3 benchmark [115]

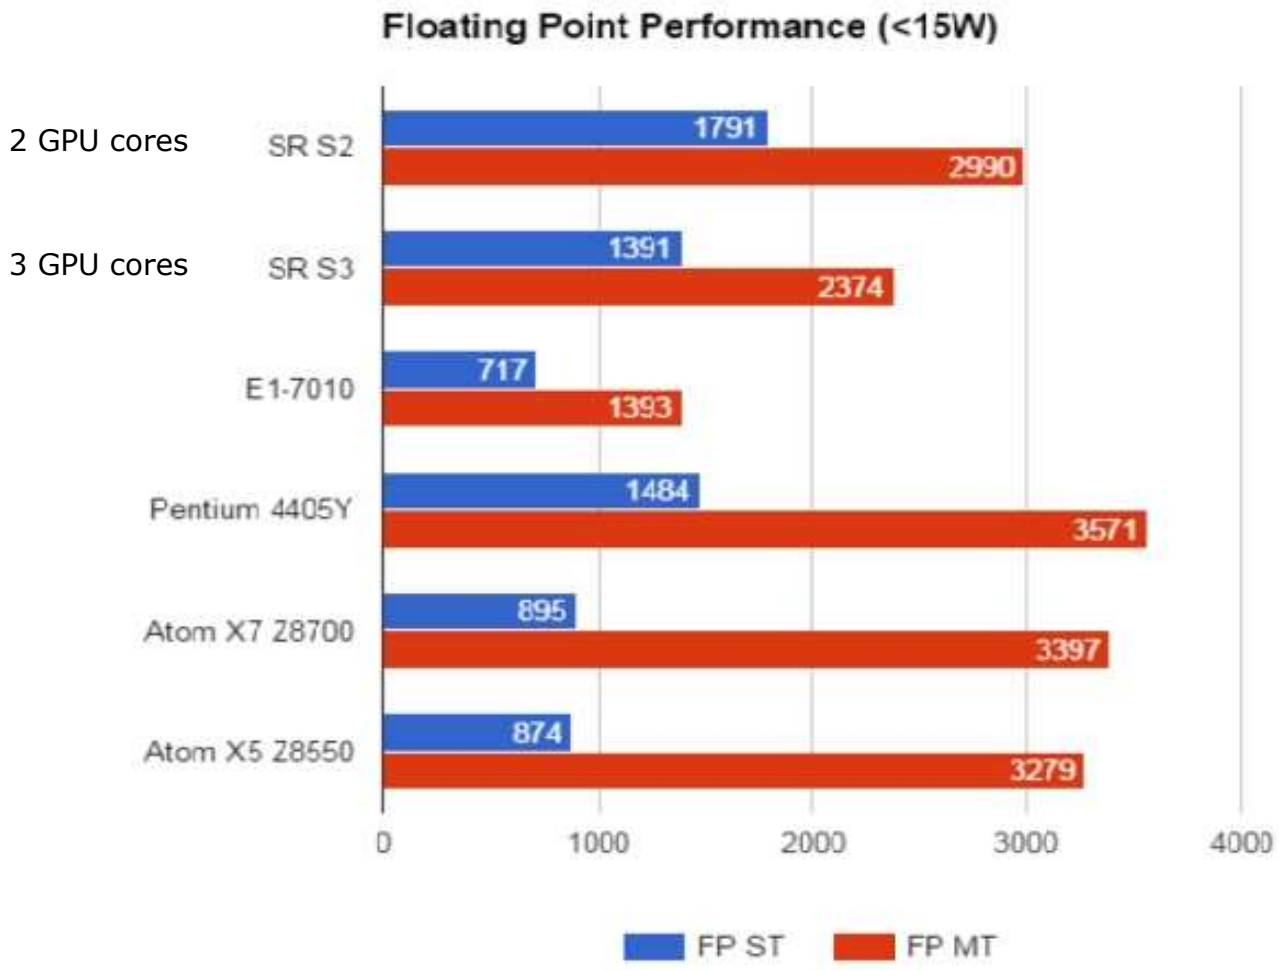

6. References

6. References (1)

[1]: Három referenciatabletet demonstrált az AMD az MWC-n, Prohardver, Febr. 27 2013, http://prohardver.hu/hir/harom_referenciatabletet_demonstralt_amd_mwc.html

[2]: Su L., Consumerization, Cloud, Convergence, AMD 2012 Financial Analyst Day, Febr. 2 2012

[3]: Bright P., Can AMD survive Bulldozer's disappointing debut?, Ars Technica, Oct. 20 2011, http://arstechnica.com/gadgets/news/2011/10/can-amd-survive-bulldozers-disappointing-debut.ars/1

[4]: Heidekrüger A., CPU / GPU Technologies Now and Future, 2010, http://www.hpcadvisorycouncil.com/events/2011/switzerland_workshop/pdf/Presentations/Day%202/10_AMD_CPU.pdf

[5]: A Nagy AMD Llano APU Megateszt, Pro Hardver, Aug. 1 2011, http://prohardver.hu/teszt/amd_llano_apu_megateszt/hammertol_huskyig.html

[6]: White S., High-Performance Power-Efficient X86-64 Server and Desktop Processors, Using the core codenamed „Bulldozer”, Aug. 19 2011, http://hotchips.org/uploads/hc23/HC23.19.9-Desktop-CPUs/HC23.19.940-Bulldozer-White-AMD.pdf

[7]: AMD Opteron Platform Overview and Product Strategy, April 2011, http://www.hp-sp.ch/events/techcircle/pastEvents/server_storage_juni2011/images/hp_techcircle_bern_amd_part.pdf

[8]: Kahn O., Valentine B., Microarchitecture Codename Sandy Bridge: New Processor Innovations, Presentation ARCS001, IDF San Francisco Sept. 2010

[9]: Shimpi A.L., The AMD FX (Bulldozer) Scheduling Hotfixes Tested, Jan. 27 2012, http://www.anandtech.com/show/5448/the-bulldozer-scheduling-patch-tested

References (2)

[10]: Kanter D., AMD's Bulldozer Microarchitecture, Real World Technologies, Aug. 26 2010, http://www.realworldtech.com/page.cfm?ArticleID=RWT082610181333&p=10

[11]: De Gelas J., Intel Core versus AMD's K8 architecture, AnandTech, May 1 2006, http://www.anandtech.com/cpuchipsets/intel/showdoc.aspx?i=2748&p=1

[12]: Expected Ivy Bridge performance, AnandTech Forums, Febr. 4 2012, http://forums.anandtech.com/showthread.php?p=32952446

[13]: Wikipedia, List of AMD Opteron microprocessors, http://en.wikipedia.org/wiki/List_of_AMD_Opteron_microprocessors#Opteron_4200-series_.22Valencia.22_.2832_nm.29

[14]: AMD Phenom™ II Processor Model Number and Feature Comparisons, http://www.amd.com/us/products/desktop/processors/phenom-ii/Pages/phenom-ii-model-number-comparison.aspx

[15]: CAS2K11 / UCAR AMD Opteron Platform Overview and Product Strategy, 2011, http://www.cisl.ucar.edu/dir/CAS2K11/Presentations/laurie/CAS2K11-AMD-September-2011-web.pdf

[16]: Wesner S., HERMIT – Petaflop/s Performance for Engineering Applications, May 26 2011, http://www.t-systems-sfr.com/e/downloads/2011/vortraege/08Wesner.pdf

[17]: Kanter D., Intel's Sandy Bridge Microarchitecture, Real World Technologies, Sept. 25 2010, http://www.realworldtech.com/page.cfm?ArticleID=RWT091810191937&p=10

[18]: Bergman R., AMD Financial Analyst Day, Nov. 11 2009, http://www.slideshare.net/AMDUnprocessed/amd-financial-analyst-day

References (3)

[19]: Schilling A., Bulldozer-Nachfolger kommen im Jahresrhythmus, Hardware Luxx, Oct. 12 2011, http://www.hardwareluxx.de/index.php/news/hardware/prozessoren/20155-bulldozer-nachfolger-kommen-im-jahresrhythmus.html

[20]: Shilov A., Ex-AMD Engineer Explains Bulldozer Fiasco: Lack of Fine Tuning, Xbit Labs, Oct. 13 2011, http://www.xbitlabs.com/news/cpu/display/20111013232215_Ex_AMD_Engineer_Explains_Bulldozer_Fiasco.html

[21]: Angelini C., Meet AMD Zambezi, Valencia, And Interlagos, Tom's Hardware, Oct. 12 2011, http://www.tomshardware.com/reviews/fx-8150-zambezi-bulldozer-990fx,3043-10.html

[22]: Goto H., AMD has pulled back the veil of Bulldozer chip 8-core version finally, 2011 http://pc.watch.impress.co.jp/docs/column/kaigai/20110830_473823.html

[23]: AMD Opteron 4200 Series Processor, http://www.siliconmechanics.com/files/BulldozerValencialInfo.pdf

[24]: BIOS and Kernel Developer's Guide (BKDG) for AMD Family 15h Models 00h-0Fh Processors, 42301 Rev. 3.08, March 12 2012, http://support.amd.com/us/Processor_TechDocs/42301_15h_Mod_00h-0Fh_BKDG.pdf

[25]: McIntyre H., Arekapudi S., Busta E., Fischer T., Golden M., Horiuchi A., Meneghini T., Naffziger S., Vinh J., Design of the Two-Core x86-64 AMD „Bulldozer" Module in 32 nm SOI CMOS, IEEE Vol. 47 No. 1, Jan. 2012, http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=06060836

[26]: De Gelas J., Bulldozer for Servers: Testing AMD's "Interlagos" Opteron 6200 Series, AnandTech, Nov. 15 2011, http://www.anandtech.com/show/5058/amds-opteron-interlagos-6200/5

References (4)

[27]: Shimpi A. L., AMD's 2012 - 2013 Server Roadmap: Abu Dhabi, Seoul & Delhi CPUs, AnandTech, Febr. 2 2012, http://www.anandtech.com/show/5488/amds-2012-2013-server-roadmap-abu-dhabi-seoul-delhi-cpus

[28]: White J., The AMD FX-8150 "Bulldozer CPU and Scorpius (FX) Platform Reviewed – Part One, Future Looks, Oct. 14 2011, http://www.futurelooks.com/the-amd-fx-8150-bulldozer-cpu-and-scorpius-fx-platform-reviewed-part-one/1/

[29]: AMD FX-8150 review, Channel Pro, Jan. 25 2012, http://www.channelpro.co.uk/reviews/6503/amd-fx-8150-review

[30]: Négy csúcs 990FX-es AM3+ alaplap a porondon, Pro Hardver, Dec. 19 2011, http://prohardver.hu/teszt/negy_csucs_990fx-es_am3_alaplap_a_porondon/a_scorpius_platform.html

[31]: Angelini C., Benchmark Results: PCMark 7, Tom's Hardware, Oct. 12 2011, http://www.tomshardware.com/reviews/fx-8150-zambezi-bulldozer-990fx,3043-12.html

[32]: AMD FX-8150 - Bulldozer im ausführlichen Test, HT4U.net, Oct. 12 2011, http://ht4u.net/reviews/2011/amd_bulldozer_fx_prozessoren/

[33]: Woligroski D., Best Gaming CPUs For The Money: April 2012, Tom's Hardware, April 2 2012, http://www.tomshardware.com/reviews/gaming-cpu-review-overclock,3106.html

[34]: Rotem E., Power Management Architecture of the 2nd Generation Intel Core Microarchitecture, Formerly Codename Sandy Bridge, Hot Chips, Aug. 2011, http://www.hotchips.org/wp-content/uploads/hc_archives/hc23/HC23.19.9-Desktop-CPUs/HC23.19.921.SandyBridge_Power_10-Rotem-Intel.pdf

References (5)

[35]: AMD | Bulldozer, Fusion, AM3+, FM1, and What's To Come, NeoGaf Belive, http://www.neogaf.com/forum/showthread.php?p=31582137

[36]: Shimpi A. L., The Bulldozer Review: AMD FX-8150 Tested, AnandTech, Oct. 12 2011, http://www.anandtech.com/show/4955/the-bulldozer-review-amd-fx8150-tested/4

[37]: Prior J., AMD Bulldozer - FX 8150 Performance Review, Rage 3D, Oct. 11 2011, http://www.rage3d.com/reviews/cpu/amd_fx_8150/index.php?p=2

[38]: Angelini C., Enabling Turbo Core, Tom's Hardware, Oct. 12 2011, http://www.tomshardware.com/reviews/fx-8150-zambezi-bulldozer-990fx,3043-8.html

[39]: George V., 45nm Next Generation Intel Core Microarchitecture (Penryn), Hot Chips 19, 2007, http://www.hotchips.org/archives/hc19/3_Tues/HC19.08/HC19.08.01.pdf

[40]: Gelsinger P., "Invent the new reality," IDF Fall 2008, San Francisco http://download.intel.com/pressroom/kits/events/idffall_2008/PatGelsinger_day1.pdf

[41]: Glaskowsky P., Explaining Intel's Turbo Boost technology, cnet News, Sept. 28. 2009, http://news.cnet.com/8301-13512_3-10362882-23.html

[42]: McGrath D., Former IBM, Lenovo exec takes the helm at AMD, EE Times, Aug. 25 2011, http://www.eetimes.com/electronics-news/4219307/AMD-appoints-former-Lenovo-exec-CEO

[43]: Patel N., AMD lays off 1400 employees, including some senior execs, The Verge, Nov. 3 2011, http://www.theverge.com/2011/11/3/2536299/amd-lays-off-1400-employees-including-some-senior-execs

References (6)

[44]: AMD High-Performance Core Roadmap 2011-2014, 3D Center, Oct. 13 2011, http://www.3dcenter.org/abbildung/amd-high-performance-core-roadmap-2011-2014

[45]: Papermaster M., The Surround Computing Era, Hot Chips Symposium, Aug. 28 2012, http://www.hotchips.org/wp-content/uploads/2012/08/HC24.28.key1-SurroundComputingEra-Papermaster-AMD.pdf

[46]: Butler M., Barnes L., Sarma D.D., Gelinas B., Bulldozer: An Approach to Multithreaded Compute Performance, IEEE Micro, Vol. 31, Issue 2, March-April 2011

[47]: Angelini C., The Piledriver Architecture: Improving On Bulldozer, Tom's Hardware, Oct. 23 2012, http://www.tomshardware.com/reviews/fx-8350-vishera-review,3328-3.html

[48]: Pollice M., Analysis: AMD Kaveri APU and Steamroller Core Architectural Enhancements Unveiled, BSN, March 6 2013, http://www.brightsideofnews.com/news/2013/3/6/analysis-amd-kaveri-apu-and-steamroller-core-architectural-enhancements-unveiled.aspx

[49]: Shimpi A.L., The Vishera Review: AMD FX-8350, FX-8320, FX-6300 and FX-4300 Tested, AnandTech, Oct. 23 2012, http://www.anandtech.com/show/6396/the-vishera-review-amd-fx8350-fx8320-fx6300-and-fx4300-tested

[50]: AMD Q1 2013 Investor Presentation, March 2013

[51]: Consumerization, Cloud, Convergence, AMD 2012 Financial Analyst Day, Febr. 2 2012, AMD Product and Technology Roadmaps, http://ir.amd.com/phoenix.zhtml?c=74093&p=irol-2012analystday

References (7)

[52]: Su L., Consumers and the World of Surround Computing, CES 2013 Press Conference, Jan. 7 2013, http://www.slideshare.net/AMD/amd-ces-2013-press-conference

[53]: Spafford K., Meredith J.S., Lee S., Li D., Roth P.C., Vetter J.S., The Tradeoffs of Fused Memory Hierarchies in Heterogeneous Computing Architectures, May 15-17 2012, http://ft.ornl.gov/~dol/papers/cf12_llano.pdf

[54]: Bennett K., AMD FX-8350 Piledriver Processor - IPC and Overclocking, Hard OCP, Oct. 22 2012, http://www.hardocp.com/article/2012/10/22/amd_fx8350_piledriver_processor_ipc_overclocking/#.UcLRYthVbps

[55]: Walton J., The AMD Trinity Review (A10-4600M): A New Hope, AnandTech, May 15 2012, http://www.anandtech.com/show/5831/amd-trinity-review-a10-4600m-a-new-hope

[56]: Angelini C., AMD FX-8350 Review: Does Piledriver Fix Bulldozer's Flaws?, Tom’s Hardware, Oct. 22 2012, http://www.tomshardware.com/reviews/fx-8350-vishera-review,3328-3.html

[57]: Payne D., Clock Design for SOCs with Lower Power and Better Specs, SemiWiki, Dec. 15 2011, http://www.semiwiki.com/forum/content/917-clock-design-socs-lower-power-better-specs.html

[58]: Clock Distribution, Acsel-lab.com, July 28 2004, http://www.acsel-lab.com/Projects/clocking/clock_distribution.htm

[59]: Restle P.J., A Clock Distribution Network for Microprocessors, IEEE Journal of Solid-State Circuits, Vol. 36, No. 5, May 2001, http://weble.upc.es/ifsin/Block5/00918917.pdf

References (8)

[60]: AMD FX-8350: Vishera, a lánctalpas cölöpverő, Prohardver, Oct. 23 2012, http://prohardver.hu/teszt/amd_fx-8350_vishera_piledriver_teszt/piledriver_v2_bulldozer_kipofozva.html

[61]: Chan S.C., A Resonant Global Clock Distribution for the Cell Broadband Engine Processor, IEEE Journal of Solid-State Circuits, Vol. 44, No. 1, Jan. 2009, http://www.ece.ncsu.edu/asic/ece733/2011/docs/ResonantClock.pdf

[62]: Ishii A.T., A Resonant-Clock 200MHz ARM926EJ-S™ Microcontroller, ESSCIRC, 2009

[63]: Sathe V., Arekapudi S., Ishii A., Ouyang C., Papaefthymiou M., Naffziger S., Resonant Clock Design for a Power-efficient, High-volume x86-64 Microprocessor, http://www.eecs.umich.edu/eecs/about/articles/2012/ISSCC_2012_Piledriver_final_submission.pdf

[64]: Courtland R., Power-Saving Clock Scheme in New PCs, IEEE Spectrum, June 28 2012, http://spectrum.ieee.org/semiconductors/processors/powersaving-clock-scheme-in-new-pcs

[65]: Shilov A., AMD Quietly Starts to Sell Two New Six-Core and Quad-Core FX Processors, Xbit Labs, March 11 2013, http://www.xbitlabs.com/news/cpu/display/20130311070551_AMD_Quietly_Starts_to_Sell_Two_New_Six_Core_and_Quad_Core_FX_Processors.html

[66]: Shilov A., GlobalFoundries Teams Up with Cyclos to Speed Up ARM Cortex-A15 Designs, Xbit Labs, Febr. 6 2013, http://www.xbitlabs.com/news/cpu/display/20130206061859_GlobalFoundries_Teams_Up_with_Cyclos_to_Speed_Up_ARM_Cortex_A15_Designs.html

[67]: AMD Opteron 6300 Series Processor, Codenamed „Abu Dhabi”, Sales-in Presentation, July 2012, http://www.abacus.cz/web/Konfig/AMD/Prodejn%C3%AD%20argumenty.pdf

References (9)

[68]: De Gelas J., AMD Launches Opteron 6300 series with "Piledriver" cores, AnandTech, Nov. 5 2012, http://www.anandtech.com/show/6430/amd-launches-opteron-6300-series-with-piledriver-cores

[69]: Hruska J., AMD Launches New Piledriver-Based Opteron 6300 Family, Hot Hardware, Nov. 5 2012, http://hothardware.com/News/AMD-Launches-New-PiledriverBased-Opteron-6300-Family/

[70]: Kozak A., AMD Desktop Platforms, 2012 AMD FX, Oct. 2012, http://www.xbitlabs.com/hot-gallery/20

[71]: Shimpi A.L., AMD A10-5800K & A8-5600K Review: Trinity on the Desktop, Part 1, AnandTech, Sept. 27 2012, http://www.anandtech.com/show/6332/amd-trinity-a10-5800k-a8-5600k-review-part-1

[72]: Woligroski D., AMD A10-4600M Review: Mobile Trinity Gets Tested, Tom's Hardware, May 15 2012, http://www.tomshardware.com/reviews/a10-4600m-trinity-piledriver,3202-4.htm

[73]: Wasson S., AMD's A10-4600M 'Trinity' APU reviewed, Tech Report, May 16 2012, http://techreport.com/review/22932/amd-a10-4600m-trinity-apu-reviewed

[74]: Naffziger S., US Patent 8010824, Aug. 30 2011

[75]: Foley D., AMD's „LLANO" Fusion APU, Hot Chips 23, Aug. 19 2011, http://www.hotchips.org/archives/hc23/HC23-papers/HC23.19.9-Desktop-CPUs/HC23.19.930-Llano-Fusion-Foley-AMD.pdf

References (10)

[76]: Angelini C., AMD Trinity On The Desktop: A10, A8, And A6 Get Benchmarked!, Tom's Hardware, Sept. 26 2012, http://www.tomshardware.com/reviews/a10-5800k-a8-5600k-a6-5400k,3224-3.html

[77]: Két 65 wattos Trinity: A10-5700 és A8-5500, Prohardver, Febr. 27 2013, http://prohardver.hu/teszt/ket_65_wattos_trinity_a10-5700_es_a8-5500/az_a10-5700_es_a8-5500.html

[78]: Kozak A., AMD Desktop Platforms, 2012 AMD A-Series, Sept. 25 2012, http://enfasys.net/ar/news/imagenes/pdf/review_amd.pdf

[79]: Valich T., AMD Virgo Uncovered: Trinity Gives You Wings?, BSN, Sept. 27 2012, http://www.brightsideofnews.com/news/2012/9/27/amd-virgo-uncovered-trinity-gives-you-wings.aspx

[80]: Altavilla D., AMD Trinity A10-4600M Processor Review, Hot Hardware, May 15 2012, http://hothardware.com/Reviews/AMD-Trinity-A104600M-Processor-Review/

[81]: Richland Die Shot AMD APU, Tom's Hardware, http://www.tomshardware.com/gallery/DieShot,0101-375851-0-2-3-1-png-.html

[82]: Valich T., AMD Launches "Elite APU" with Richland, Successor to Trinity, BSN, March 12 2013, http://www.brightsideofnews.com/news/2013/3/12/amd-launches-e2809celite-apue2809d-with-richland2c-successor-to-trinity.aspx

[83]: Sakr S., AMD Richland chips will arrive in notebooks next month, promise better graphics, battery life and a few extras, Engadget, March 12 2013, http://www.engadget.com/2013/03/12/amd-richland-details/

References (11)

[84]: Hruska J., AMD's new Richland APU boosts clocks and adds features, but is ultimately just a minor Trinity refresh, Extreme Tech, March 12 2013, http://www.extremetech.com/computing/150451-amds-new-richland-apu-boosts-clocks-and-adds-features-but-its-a-just-modest-refresh

[85]: Broekhuijsen N., New Details Revealed on AMD's Upcoming Richland Chips, Tom's Hardware, March 12 2013, http://www.tomshardware.com/news/Richland-APU-AMD,21318.html

[86]: AMD Richland APU Preview: Trinity Gets a Facelift, Hardware Canucks, March 10 2013, http://www.hardwarecanucks.com/forum/hardware-canucks-reviews/60112-amd-richland-apu-preview-trinity-gets-facelift.html

[87]: Richland, Kaveri, Kabini & Temash; AMD's 2013 APU Lineup Examined, Hardware Canucks, Jan. 9 2013, http://www.hardwarecanucks.com/forum/hardware-canucks-reviews/59053-richland-kaveri-kabini-temash-amd-s-2013-apu-lineup-examined.html

[88]: New AMD A-Series APU Offers Mobile PC Users Innovative Experiences, Elite Graphics Performance and Longer Battery Life, March 12 2013, http://www.amd.com/us/press-releases/Pages/new-amd-a-series-2013mar12.aspx

[89]: Hachman M., AMD fires back at Intel's Haswell with its A-Series desktop processors, PC World, June 4, 2013, https://www.pcworld.com/article/2040763/amd-fires-back-at-intels-haswell-with-its-a-series-desktop-processors.html

[90]: Hinum K., Intel Core i7-3520M, Notebook Check, May 3 2012, http://www.notebookcheck.net/Intel-Core-i7-3520M-Notebook-Processor.74446.0.html

[91]: Rogers P., Macri J., Marinkovic S., AMD Heterogeneous Uniform Memory Access, Apr. 30 2013, http://events.csdn.net/AMD/130410%20-%20hUMA_v6.6_FINAL.PDF

References (12)

[92]: Wikipedia, Unified Video Decoder, https://en.wikipedia.org/wiki/Unified_Video_Decoder

[93]: Cutress I., AMD Announces the 7th Generation APU: Excavator mk2 in Bristol Ridge and Stoney Ridge for Notebooks, AnandTech, May 31 2016, https://www.anandtech.com/show/10362/amd-7th-generation-apu-bristol-ridge-stoney-ridge-for-notebooks

[94]: AMD Opteron 6300 Series Processors, https://www.amd.com/en-us/products/server/opteron/6000/6300

[95]: AMD Details Bulldozer Processor Architecture, TechPowerUp, Aug. 24 2010, https://www.techpowerup.com/129392/amd-details-bulldozer-processor-architecture

[96]: Cutress I., Garg R., AMD Kaveri Review: A8-7600 and A10-7850K Tested, AnandTech, Jan. 14 2014, https://www.anandtech.com/show/7677/amd-kaveri-review-a8-7600-a10-7850

[97]: Wikipedia, List of AMD accelerated processing unit microprocessors, https://en.wikipedia.org/wiki/List_of_AMD_accelerated_processing_unit_microprocessors#%22Kaveri%22_(2014)

[98]: AMD A8-7100 specifications, CPU-World, http://www.cpu-world.com/CPUs/Bulldozer/AMD-A8-Series%20A8-7100.html

[99]: Shilov A., AMD 'Raven Ridge': Mainstream APU with 'Zen' cores due in 2017, KITGURU, June 24 2015, https://www.kitguru.net/components/cpu/anton-shilov/amd-raven-ridge-mainstream-apu-with-zen-cores-due-in-2017/

References (13)

[100]: Cutress I., AMD Launches Carrizo: The Laptop Leap of Efficiency and Architecture Updates, AnandTech, June 2 2015, https://www.anandtech.com/show/9319/amd-launches-carrizo-the-laptop-leap-of-efficiency-and-architecture-updates/4

[101]: Hruska J., New leak hints at AMD Zen's architecture, organization, Extreme Tech, April 29 2015, https://www.extremetech.com/gaming/204523-new-leak-hints-at-amd-zens-architecture-organization

[102]: Cutress I., AMD at ISSCC 2015: Carrizo and Excavator Details, AnandTech, Febr. 23 2015, https://www.anandtech.com/show/8995/amd-at-isscc-2015-carrizo-and-excavator-details

[103]: Toh S.O., McLellan E.J. et al., Replica path timing adjustment and normalization for adaptive voltage and frequency scaling, Patent US9575553B2, Febr. 21 2017

[104]: Wilcox K., Akeson D. et al., A 28nm x86 APU optimized for power and area efficiency, ISSCC 2015

[105]: AMD Carrizo (A10-8700P) review – will the new series of APU chips turn AMD around?, Laptopmedia, March 15 2016, https://laptopmedia.com/review/amd-carrizo-a10-8700p-review-will-the-new-series-of-apu-chips-turn-amd-around/

[106]: Walton M., Sixth time lucky: AMD details the Carrizo APU, Ars Technica, June 3 2015, https://arstechnica.com/information-technology/2015/06/sixth-time-lucky-amd-details-the-carrizo-apu/

[107]: Woligroski D., Mullins And Beema APUs: AMD Gets Serious About Tablet SoCs, Tom's Hardware, April 28 2014, https://www.tomshardware.com/reviews/amd-tablet-processor,3813.html

References (14)

[108]: Mujtaba H., AMD Details Carrizo APUs Energy Efficient Design at Hot Chips 2015 – 28nm Bulk High Density Design With 3.1 Billion Transistors, 250mm2 Die, WCCF Tech, Aug. 26 2015, https://wccftech.com/amd-carrizo-apu-architecture-hot-chips/

[109]: AMD 2014 Mobility APU Lineup Announcement, Slideshare, Nov. 15 2013, https://www.slideshare.net/AMD/amd-mobility-apu-lineup-announcement?from_action=sav

[110]: Alcorn P., AMD Details Bristol Ridge And Stony Ridge A-Series APUs At Computex 2016, Tom's Hardware, June 1 2016, https://www.tomshardware.co.uk/amd-bristol-ridge-stony-apu,news-53130.html

[111]: Wee W.C., AMD officially launches its new 7th generation A-series "Bristol Ridge" APUs, Hardware Zone, Sept. 6 2016, https://www.hardwarezone.com.sg/tech-news-amd-officially-launches-its-new-7th-generation-series-bristol-ridge-apus

[112]: Pirzada U., AMD 7th Generation APU Lineup And Specifications Leaked – Flagship SKU Features 4 x86 Excavator Based Cores and 8 GCN 3.0 Cus, WCCF Tech, May 21 2016, https://wccftech.com/amd-7th-generation-apu-lineup-specifications-leaked-flagship-sku-features-4-x86-excavator-based-cores-8-gcn-30-cus/

[113]: Schönborn T., Bristol Ridge in Review: AMDs A10-9600P Against the Competition, Notebook Check, July 5 2016, https://www.notebookcheck.net/Bristol-Ridge-in-Review-AMDs-A10-9600P-Against-the-Competition.168477.0.html

[114]: AMD Stoney Ridge core, CPU-World, http://www.cpu-world.com/Cores/Stoney_Ridge.html

[115]: AMD's Stoney Ridge Performance And Market Positioning Detailed, WCCF Tech, March 5 201 https://wccftech.com/amds-stoney-ridge-performance-market-positioning-detailed/

References (15)

[116]: Mujtaba H., AMD Richland Mobile "Elite Performance APU Platform Launched – Powered by Piledriver Cores, WCCFTech, May 23, 2013, https://wccftech.com/amd-richland-mobile-elite-performance-apu-platform-launched-powered-piledriver-cores