AMD's early processor lines, up to the Hammer Family (Families K8 - K10.5h)

Dezső Sima

October 2018

(Ver. 1.1)

AMD's early processor lines, up to the Hammer Family (Families K8 - K10.5h)

1. Introduction to AMD’s processor families

1. Introduction to AMD's processor families

AMD's early x86 processor history [1]

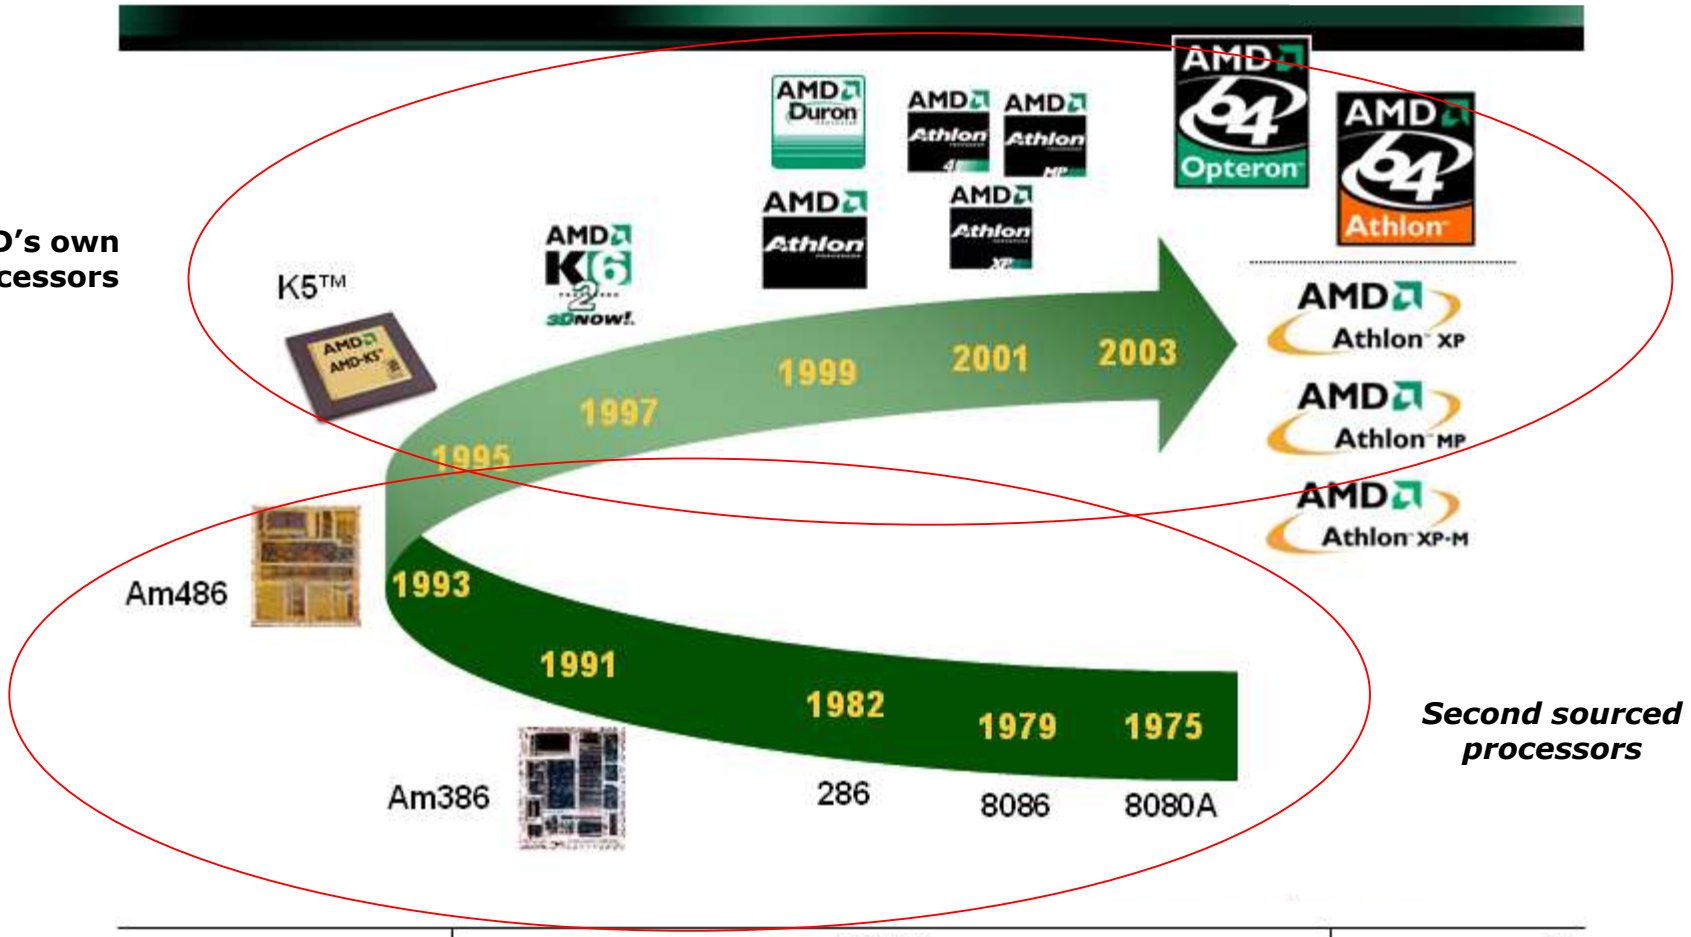

1. Introduction to AMD's processor families (2)

Evolution of AMD's early processors [2]

16	→ 32 →	64 Bits
Integer	→	FPU, SIMD Vector
1/3 Issue	→	9 Issue
Trivial Cache	→	1MB Cache
CPU	→	SOC

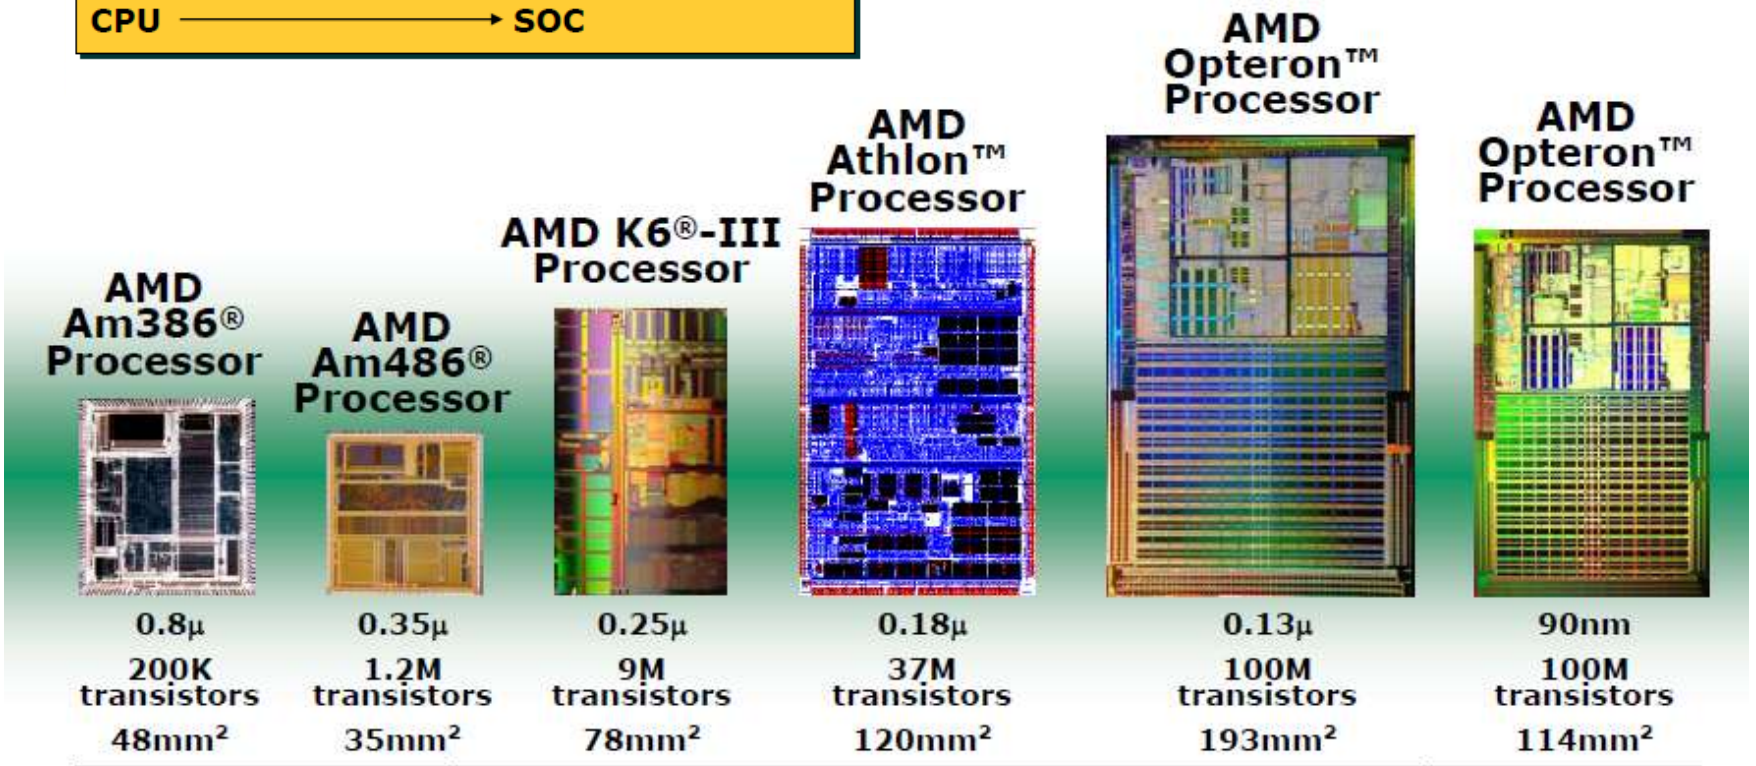

1. Introduction to AMD's processor families (3)

Historical remarks

1) Beyond x86 processors AMD also designed and marketed two embedded processor families;
   - the 2900 family of bipolar, 4-bit slice microprocessors (1975-?) used in a number of processors, such as particular DEC 11 family models, and
   - the 29000 family (29K family) of CMOS, 32-bit embedded microcontrollers (1987-95).

   In late 1995 AMD cancelled their 29K family development and transferred the related design team to the firm's K5 effort, in order to focus on x86 processors [3].
2) Initially, AMD designed the Am386/486 processors that were clones of Intel's processors.
3) Then the K5 was AMD's first in-house designed processor.
4) The K6 was originally developed by NexGen, a firm that AMD purchased in 1995.

   This processor was pin-compatible with Intel's Pentium.

   Subsequent K6 models became competitive with Intel's Pentium II/III.

1. Introduction to AMD's processor families (4)

The K and Family xxh nomenclature of AMD's processor families

- The K designation used previously by AMD is a counterpart of Intel's P designation for their processor families.
- It was inspired by comic books, since Kryptonite was the only substance that could bring Superman to knees. Obviously, Superman stays for Intel [4].
- Presumably, a similar inspiration is behind AMD's core names such as Sledgehammer, Clawhammer, or even Bulldozer and its successors including Piledriver, Steamroller etc.
- In Nov. 2004 AMD abandoned using the K moniker for their basic architectures in order to signalize their move to enter a wide variety of markets and started to use their own in-house Family xxh designations.

  Nevertheless, outside AMD the K designation was used further on for simplicity and clarity more or less until arriving the Family 11h (Griffin-based) lines.

  In this chapter we will also make use of the K nomenclature including the K10.5 family and will change to the Family 1xh designation beginning with the Family 11h (Griffin-based).

1. Introduction to AMD's processor families (5)

AMD's x86-64 family designations and related main features

FamilyE	Intro.	Core	Techn. (nm)	Used typically in	Core contr.	Market segment
F0h (K8)	2003	Sledghammer	130	Sledgehammer	1	S, DT, M
	2004	Athens	90	Athens	1	S, DT, M
	2005	Egypt	90	Egzpt	2	S, DT, M
0Fh NPT (K8)	2006	Hammer	90	Santa Rose	2	S, DT, M
10h	2007	Greyhound	65	Barcelona	4	S, DT
	2008	Greyhound+	45	Shanghai	4	S, DT, M
	2009			Istambul	6	S, DT
	2010			Magny Course	2*6	S
11h	2008	Griffin	65	Lion	2	DT, M
12h	2011	Husky (Liano)	32	Fusion A/E2	4+ GPU	DT, M
14h	2011	Bobcat	40	Fusion C/E/G/Z	2 + GPU	M
15h	2011	Bulldozer/Piledriver/ Steamroller/ Excavator	32	Interlagos (S) Zambezi (DT)	2*8 8	S, DT
16h	2012	Jaguar/Puma+	28	Fusion C/E/G/Z	2 + GPU	M
17h	2017	Zen		EPYC/ThreadRipper/ Ryzen	8/16/32	S, DT, M

S: Server DT: Desktop M: Mobile

1. Introduction to AMD's processor families (6)

Brand names of AMD's processor lines

AMD typically, assigns unique brand names to their processor lines that indicate

- the processor famiy
- market segment and
- relative performance

the particular processor line supports.

1. Introduction to AMD's processor families (7)

Main market segments

- AMD strives to cover main market segments.
- Salient market segments vary however, temporally.
- Concerning this, AMD designed and marketed a low-power oriented processor family, targeting tablets and smart phones, the Cat family, in the first halve of the 2010's (between 2011 and 2015).

  Nevertheless, the Cat family was not successful and AMD cancelled the Cat line in 2015.

  As an example for the main market segments, the next Figure shows AMD's market segments favored in their Hammer family (K8 to K10,5), that is
  - servers
  - desktops and
  - mobiles.
- Obviously, AMD's different processor families emphasize different market segments.

1. Introduction to AMD's processor families (8)

Example for main market segments in AMD's K8 – K10.5h processor lines

		2003-2007	2007-2008	2008-2011	2009	2009
		K8 (Hammer)	**K10 (Barcelona)**	**K10.5 (Shanghai)**	**K10.5 (Istanbul)**	**K10.5 (Magny- Course)**
Servers	**4P servers**	**See Section 4**	**Barcelona (834x-836x))**	**Shanghai (837x-839x)**	**Istambul (8410-8430)**	**Magny-Course (6100)**
	2P servers		**Barcelona (234x-236x)**	**Shanghai (237x-239x)**	**Istambul (241x-243x)**	**Lisbon (4100)**
	1P servers		**Budapest (135x-136x)**	**Suzuka (138x-139x)**		
Desktops	**High perf.** (~80-120W)		Phenom X4-X2	Phenom II X4-X2	Phenom II X6-X4	
	Mainstream (~60-90W)	Athlon 64 Athlon 64 X2	Athlon X2	Athlon II X4-X2		
	Value (~40-60W)	Sempron		Sempron		
MobIles	**High perf.** (~30-40W)	Turion 64 X2 (TL 6/5) Turion 64 (ML/MT)		Phenom II (N/P 9xx-6xx) Turion II Ultra (M6xx) Turion II (M/N/P 5xx)		
	Mainstream (~20-30W)	Athlon 64 X2 (TK-5x/4x) Athlon 64 (2xxx+-4xxx+)		Athlon II (M/N/P 3xx) Sempron (M1xx)		
	Ultraportable (~10-20W)	Mobile Sempron (2xxx+-4xxx+) Sempron 2100 fanless		Turion II Neo (K6xx) Athlon II Neo (K1xx) V-series (V1xx)		
Embedded (~10-20W)				Turion II Neo X2 Athlon II Neo X2 Athlon II Neo		

Performance classes within market segments

Typically, in each market segment processor lines are broken down into performance classes, indicating the relative performance of the processor lines within a particular market segment, like

- high performance desktops
- mainstream desktops and
- value desktops,

as shown below again for the Hammer family (K8 – K10.5) for the desktop lines.

1. Introduction to AMD's processor families (10)

Example for performance classes within the desktop segment in AMD's K8 – K10.5h lines

		2003-2007	2007-2008	2008-2011	2009	2009
		K8 (Hammer)	**K10 (Barcelona)**	**K10.5 (Shanghai)**	**K10.5 (Istanbul)**	**K10.5 (Magny- Course)**
Servers	**4P servers**	**See Section 4**	**Barcelona (834x-836x))**	**Shanghai (837x-839x)**	**Istambul (8410-8430)**	**Magny-Course (6100)**
	2P servers		**Barcelona (234x-236x)**	**Shanghai (237x-239x)**	**Istambul (241x-243x)**	**Lisbon (4100)**
	1P servers		**Budapest (135x-136x)**	**Suzuka (138x-139x)**		
Desktops	**High perf.** (~80-120W)		Phenom X4-X2	Phenom II X4-X2	Phenom II X6-X4	
	Mainstream (~60-90W)	Athlon 64 Athlon 64 X2	Athlon X2	Athlon II X4-X2		
	Value (~40-60W)	Sempron		Sempron		
MobIles	**High perf.** (~30-40W)	Turion 64 X2 (TL 6/5) Turion 64 (ML/MT)		Phenom II (N/P 9xx-6xx) Turion II Ultra (M6xx) Turion II (M/N/P 5xx)		
	Mainstream (~20-30W)	Athlon 64 X2 (TK-5x/4x) Athlon 64 (2xxx+-4xxx+) Mobile Sempron (2xxx+-4xxx+)		Athlon II (M/N/P 3xx) Sempron (M1xx)		
	Ultraportable (~10-20W)	Sempron 2100 fanless		Turion II Neo (K6xx) Athlon II Neo (K1xx) V-series (V1xx)		
Embedded (~10-20W)				Turion II Neo X2 Athlon II Neo X2 Athlon II Neo		

Brand names of AMD's processor lines

As an example, the next Figure shows brand names of K10.5 Shanghai based desktop processor lines of different performance potentials, such as:

- Phenom II
- Athlon II and
- Sempron.

1. Introduction to AMD's processor families (12)

Brand names of AMD's 64-bit K8 – Family 10.5h processor lines

		2003-2007	2007-2008	2008-2011	2009	2009
		K8 (Hammer)	**K10 (Barcelona)**	**K10.5 (Shanghai)**	**K10.5 (Istanbul)**	**K10.5 (Magny- Course)**
Servers	**4P servers**	**See Section 4**	**Barcelona (834x-836x))**	**Shanghai (837x-839x)**	**Istambul (8410-8430)**	**Magny-Course (6100)**
	2P servers		**Barcelona (234x-236x)**	**Shanghai (237x-239x)**	**Istambul (241x-243x)**	**Lisbon (4100)**
	1P servers		**Budapest (135x-136x)**	**Suzuka (138x-139x)**		
Desktops	**High perf.** (~80-120W)		Phenom X4-X2	Phenom II X4-X2	Phenom II X6-X4	
	Mainstream (~60-90W)	Athlon 64 Athlon 64 X2	Athlon X2	Athlon II X4-X2		
	Value (~40-60W)	Sempron		Sempron		
Mobiles	**High perf.** (~30-40W)	Turion 64 X2 (TL 6/5) Turion 64 (ML/MT)		Phenom II (N/P 9xx-6xx) Turion II Ultra (M6xx) Turion II (M/N/P 5xx)		
	Mainstream (~20-30W)	Athlon 64 X2 (TK-5x/4x) Athlon 64 (2xxx+-4xxx+)		Athlon II (M/N/P 3xx) Sempron (M1xx)		
	Ultraportable (~10-20W)	Mobile Sempron (2xxx+-4xxx+) Sempron 2100 fanless		Turion II Neo (K6xx) Athlon II Neo (K1xx) V-series (V1xx)		
Embedded (~10-20W)				Turion II Neo X2 Athlon II Neo X2 Athlon II Neo		

Model designations within a processor line

In addition to the brand names model designations differentiate particular models of a processor line.

Model designations

- may include a tag, such as X2 or X4 that indicates the number of cores (e.g. X2 meaning dual cores etc.) and
- a model number that specifies the features of the processor, such as the clock frequency, L2 or L3 cache size, wattage (dissipation etc.) as shown in an example given for the K10.5 Shanghai based Phenom II X3 7xx line in the next table.

1. Introduction to AMD's processor families (14)

Example: Processor model designations of the Phenom II X3 line (Desktop line based on the K10.5 Shanghai derived Deneb core) [5]

Model Number	Step.	Freq.	L2 Cache	L3 Cache	HT	Multi[1]	Voltage	TDP	Socket	Release Date
Phenom II X3 700e	C2	2.4 GHz	3x 512 KB	6 MB	2 GHz	12x	0.825 - 1.25	65 W	AM3	June 2, 2009
Phenom II X3 705e	C2	2.5 GHz	3x 512 KB	6 MB	2 GHz	12.5x	0.800 - 1.25	65 W	AM3	June 2, 2009
Phenom II X3 710	C2	2.6 GHz	3x 512 KB	6 MB	2 GHz	13x	0.875 - 1.425	95 W	AM3	Febr. 9, 2009
Phenom II X3 715 *Black Edition*[2]	C2	2.8 GHz	3x 512 KB	6 MB	1.8 GHz	14x	0.875 - 1.425	95 W	AM2+	???
Phenom II X3 720	C2	2.8 GHz	3x 512 KB	6 MB	2 GHz	14x	0.875 - 1.425	95 W	AM3	???
Phenom II X3 720 *Black Edition*[2]	C2	2.8 GHz	3x 512 KB	6 MB	2 GHz	14x	0.850 - 1.425	95 W	AM3	Febr. 9, 2009
Phenom II X3 740 *Black Edition*[2]	C2	3.0 GHz	3x 512 KB	6 MB	2 GHz	15x	0.850 - 1.425	95 W	AM3	Sept. 2009

1. Introduction to AMD's processor families (15)

Implementation of cores in a processor family

- Each processor family, like the K10.5 Shanghai family, etc. often is based on a number of processor lines with each line based on one or more different cores having different features, like the number of CPU cores or the size of the L2 cache.

  Obviously, different processor cores target different market segments (like servers, desktopsr or mobiles) and performance levels, like high-performance, mainstream or low cost processors.

- As an example, the next slide depicts all the cores of the K10.5 (Shanghai) family.

1. Introduction to AMD's processor families (16)

AMD's K10.5 Shanghai based processor lines [based on xx]

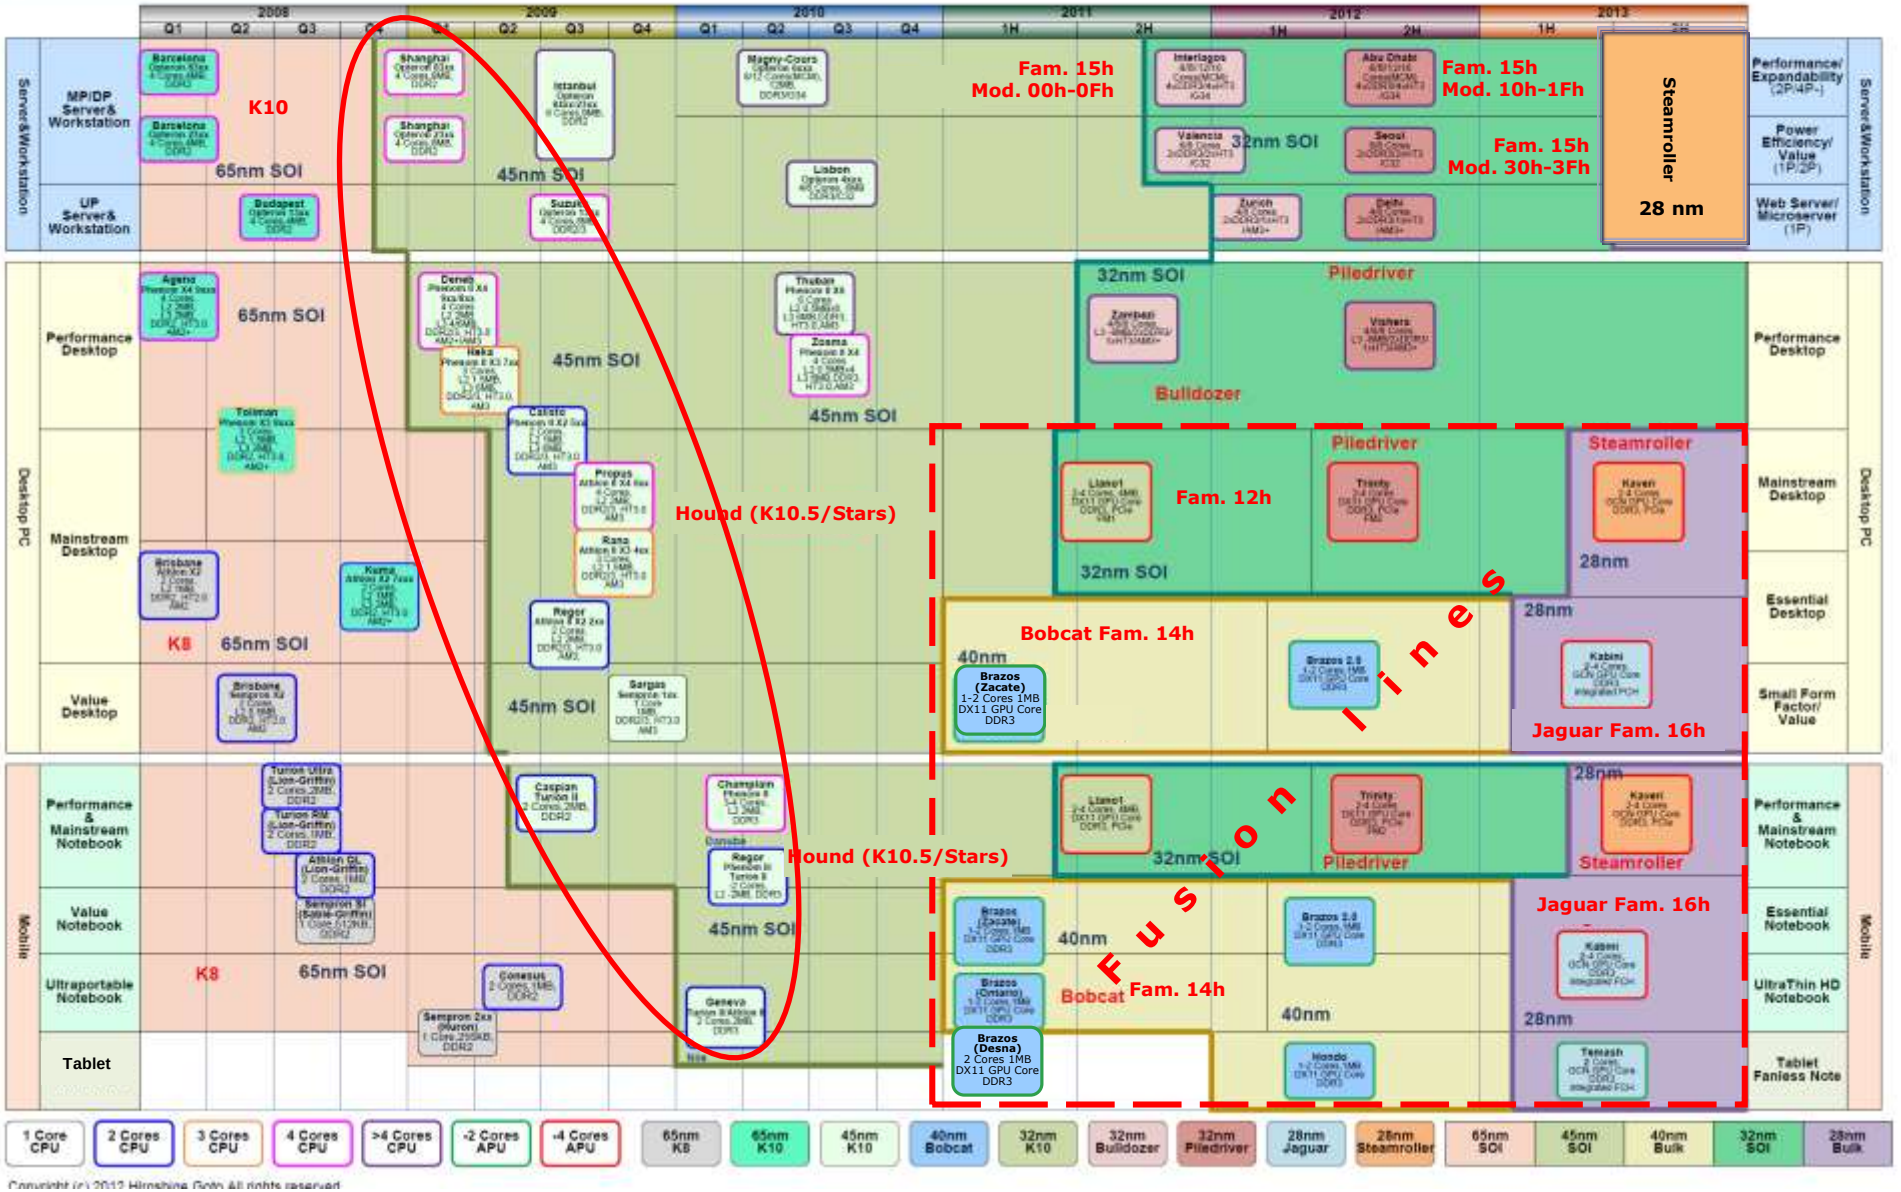

Set of processor cores targeting a particular market segment

Typically, cores of a processor family, like the K10.5 Shanghai family, may be subdivided into a certain core sets, with each set aiming at different market segments, such as the server, desktop or mobile segment, as the next example for the K10.5 Shanghai family shows.

1. Introduction to AMD's processor families (18)

Sets of processor cores in the K10.5 Shanghai family [based on xx]

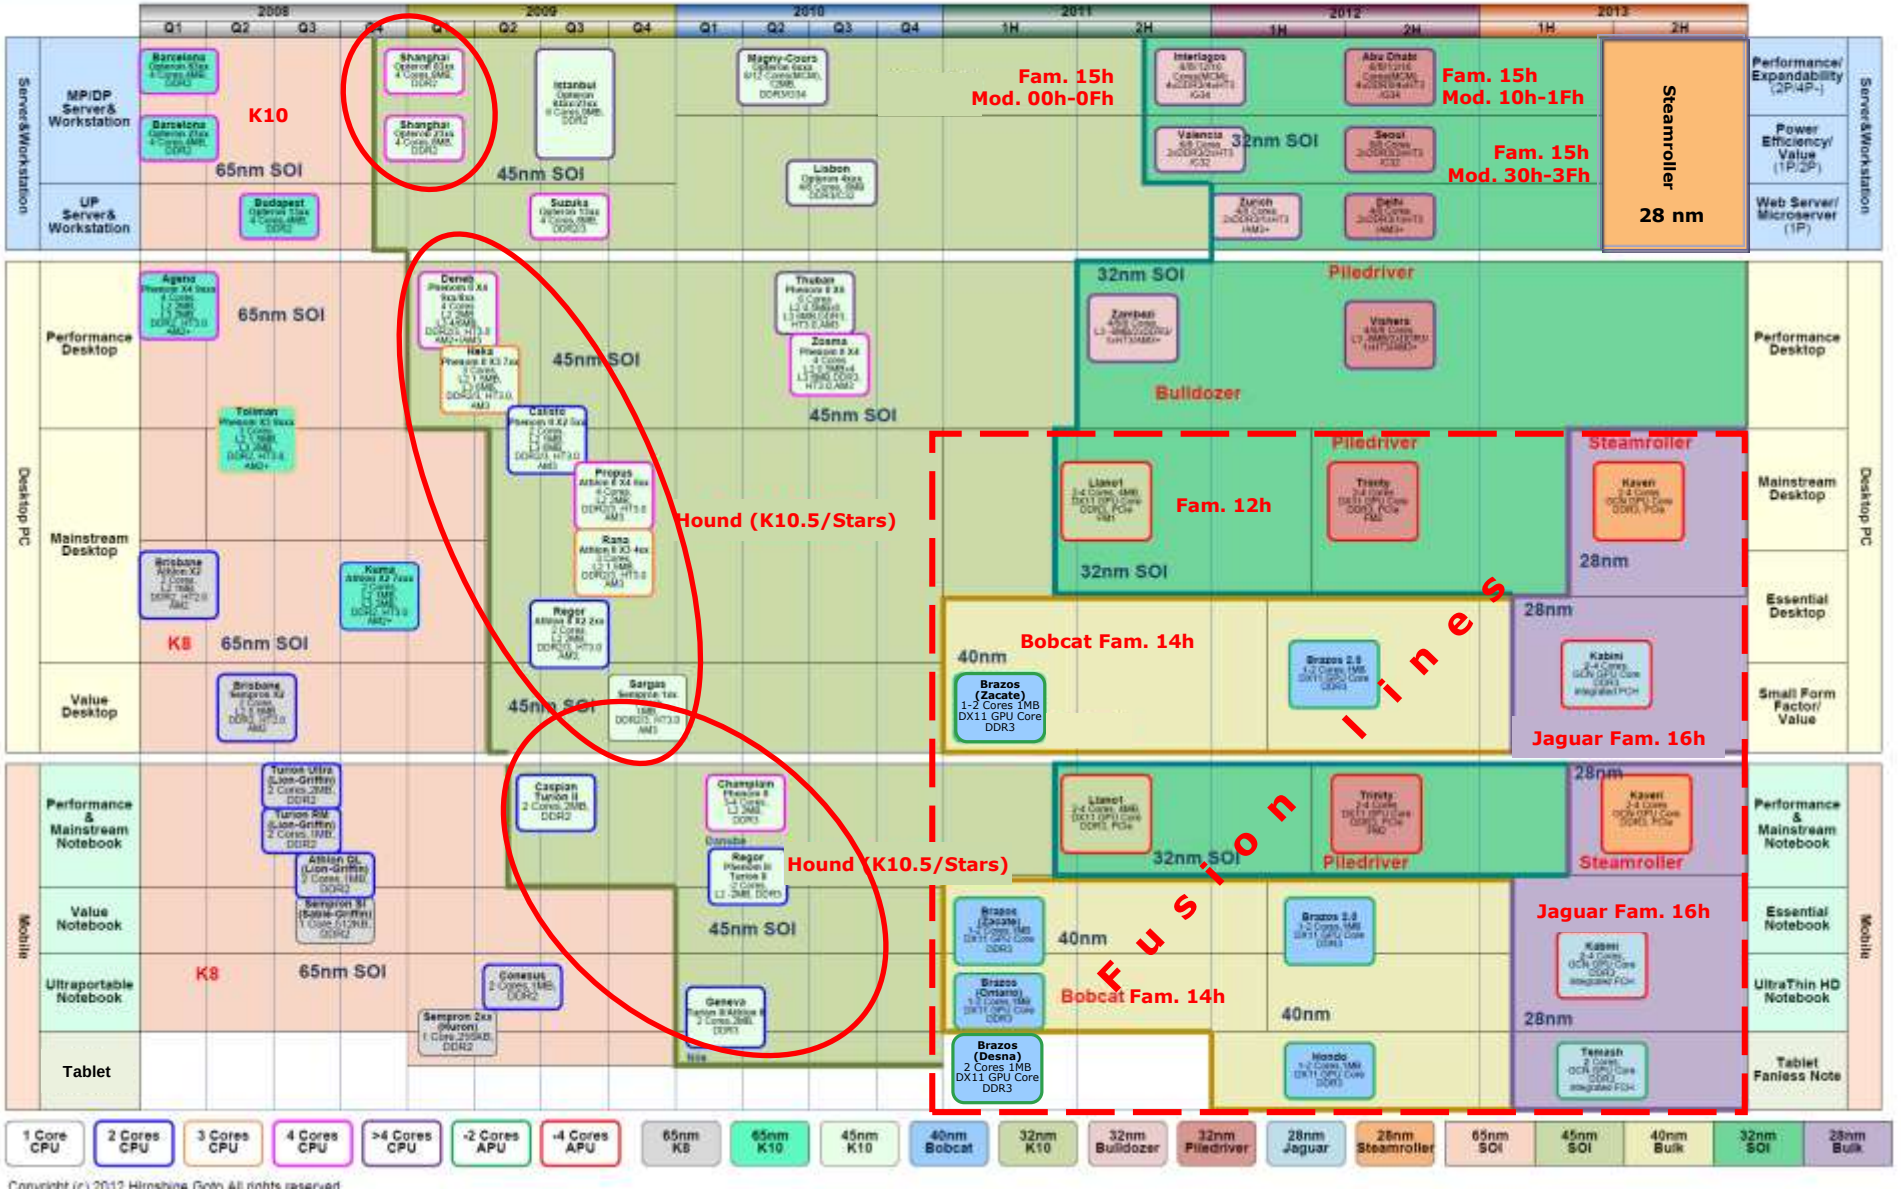

Subsets of processor cores within a set of cores

In some cases even a given set of processor cores that targets a particular market segment consists of a few subsets, each targeting different performance levels, such as high-performance, mainstream or low-cost processors, like the desktop cores of the K10.5 Shanghai famiiy, as the next figure shows.

1. Introduction to AMD's processor families (20)

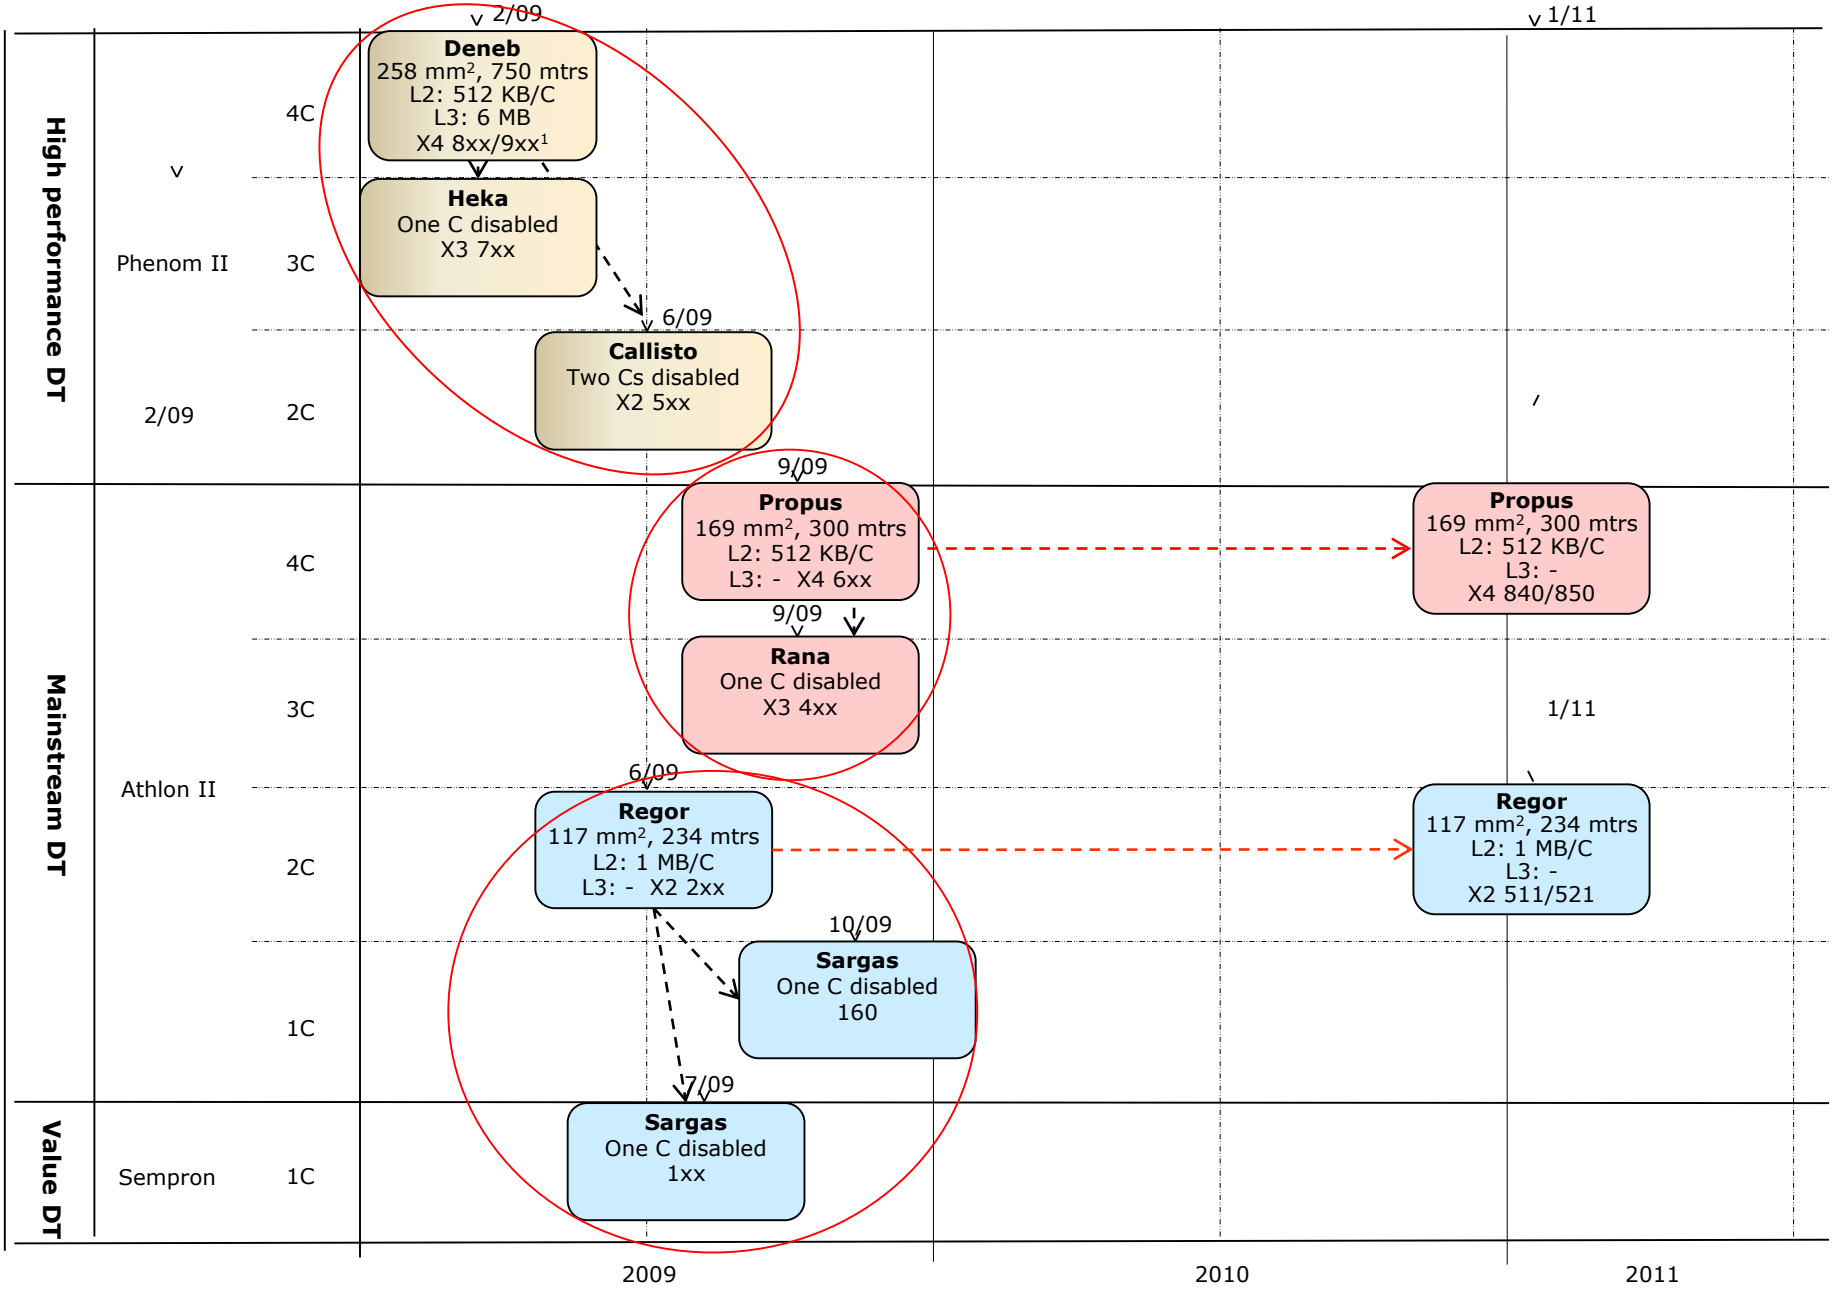

[1] For X4 8xx; L3=4MB

Figure: Subsets of AMD's Shanghai based desktop cores (45 nm)

Native designs and partly disabled cores

Sets of cores that can not be broken down into subsets or particular subsets typically include a native design and a few partly disabled cores.

As an example, the desktop cores of the K10.5 Shanghai family are based on three native core designs, called the

- Deneb
- Propus and
- Regor,

cores, as indicated in the next figure.

1. Introduction to AMD's processor families (22)

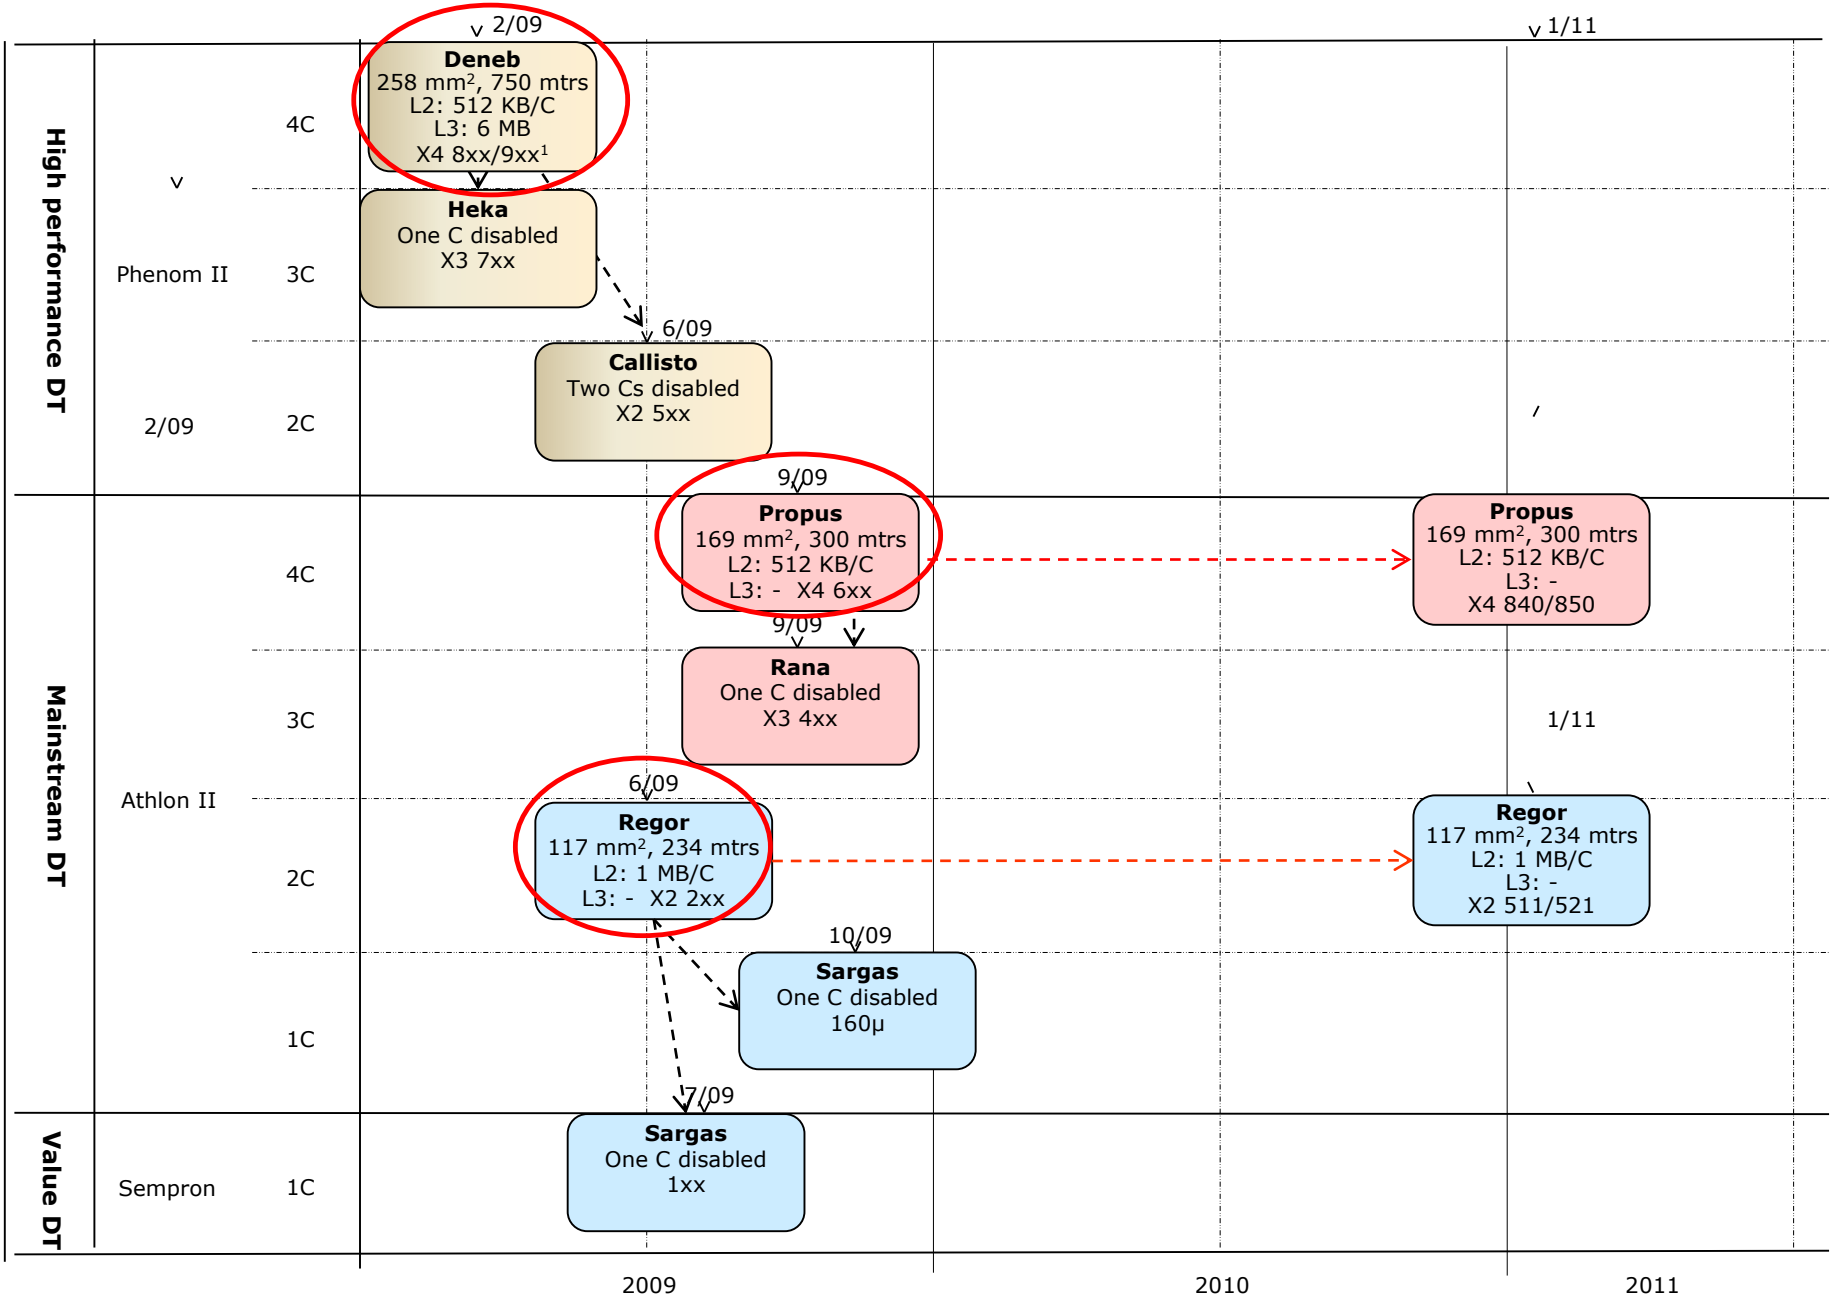

¹ For X4 8xx; L3=4MB

Figure: Subsets of AMD's Shanghai based desktop cores (45 nm)

1. Introduction to AMD's processor families (23)

In the above Figure arrows indicate partly disabled cores that are derived from the native designs;

E.g. In the Phenom II lines of the desktop cores of the K10.5 Shanghai family the 4-core Deneb core represents the native design whereas

- the Heka core is in fact a Deneb core with a single core disabled and
- the Callisto core is a Deneb core with two cores disabled,

as the next Figure shows.

1. Introduction to AMD's processor families (24)

Subsets of AMD's Shanghai based desktop cores (45 nm)

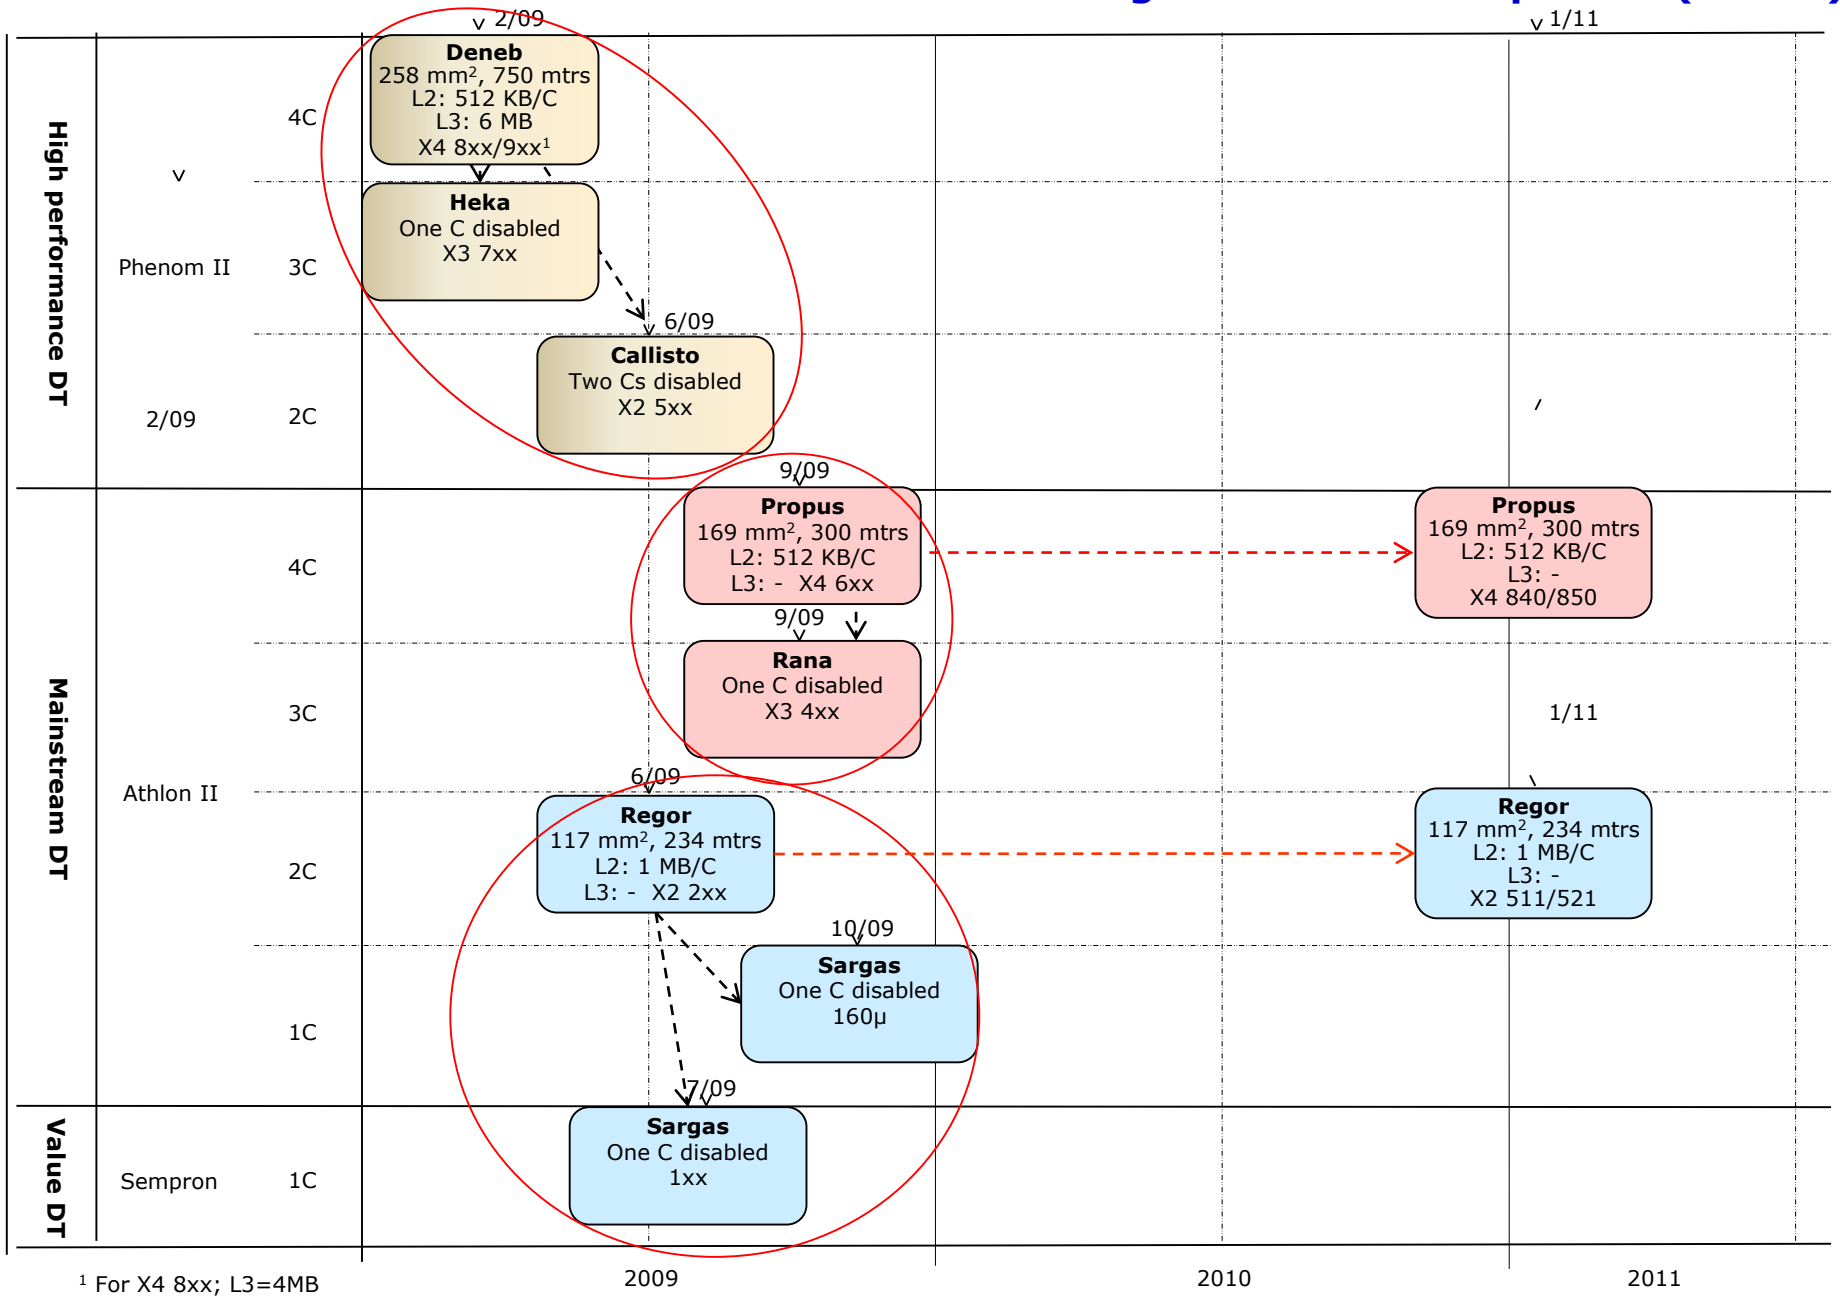

[1] For X4 8xx; L3=4MB

Possible aims of disabling features of native designs

Disabling defective units

Aim: To sell processors with defective features, such as less cores, smaller L2, no SSE3 or AMD-V at lower price

E.g. models of a particular line with smaller L2, e.g. 512 K L2 instead of 1 MB L2

Disabling available features to reduce functionality

Aim: To avoid separate designs for lower priced lines but to maintain enough difference in functionality to higher priced lines

E.g. Sempron lines are typically native 512 KB L2 Athlon 64 designs with reduced functionality, such as 128/256 KB L2 and disabled SSE3 or AMD-V

Remark

In this chapter we focus on AMD's processor designs rather than marketing issues, so typically we disregard partly disabled native designs.

Remark

Core designations are often irritating and serve marketing purposes as for example the previous figure demonstrates.

Overview of AMD's processor lines

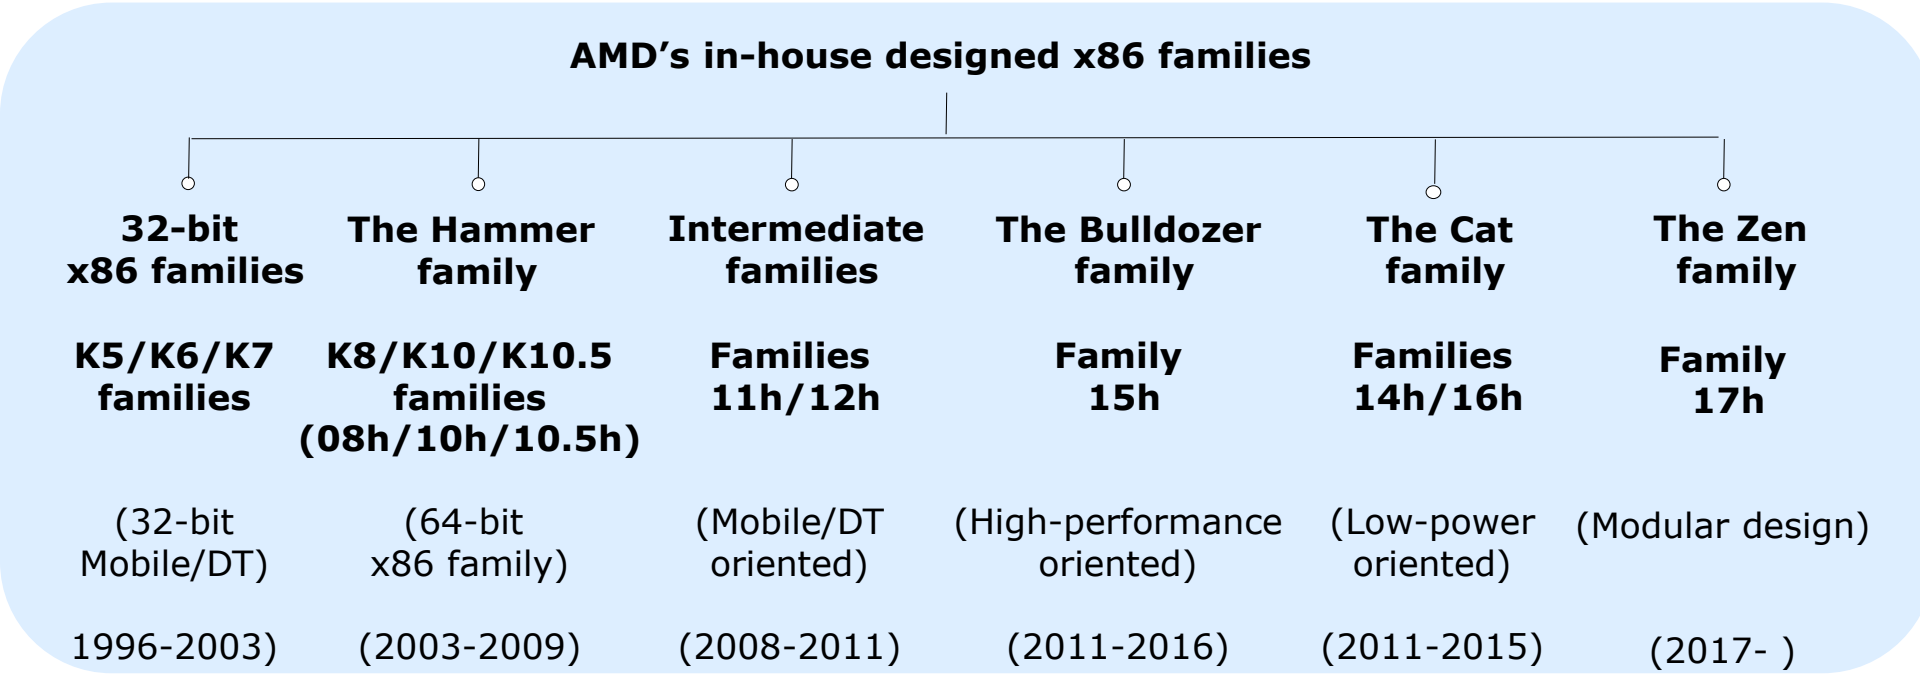

Remark

Before the K5 AMD manufactured (licensed) Intel designed processors rather than own designs

AMD's 32-bit x86 families

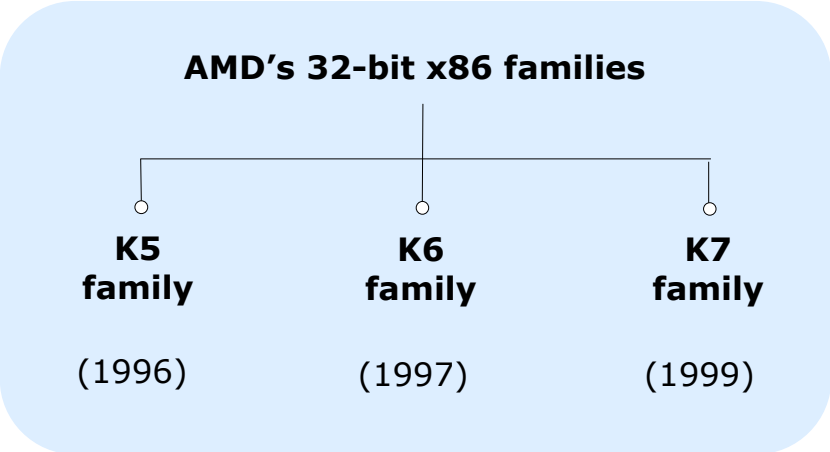

Overview and major innovations of AMD's K5/K6 families

CPU Family	Intro.	CPU core	Brand name	Techn. (µm)	New key feature	Typ. Application
K5	1996		K5	0.5/0.35	2. gen. superscalar (32-bit), Pentium competitor	DT
K6	1997		K6	0.35/0.25	2.5 gen. superscalar, MMX (NexGen design)	DT/M
K6-2	1998	Chomper	K6	0.25	3. gen. superscalar, 3DNow!	DT/M
K6III **K6-2+** **K6III+**	1999 2000 2000	Sharptooth na. na.	K6	0.25 0.18 0.18	On-die L2 PowerNow![1] PowerNow![1]	DT/M

[1] PowerNow! was introduced in the Mobile K6-2+ and Mobile K-III+ processors in 2000

[2] Athlon: Attained performance lead over Intel's Pentium III

1. Introduction to AMD's processor families (30)

Overview and major innovations in AMD's K7 (Athlon) families

<table>
<tr><th colspan="2">Base arch./ stepping</th><th>Intro.</th><th>Core</th><th>Brand name</th><th>Techn. (µm)</th><th colspan="2">L2 cache</th><th>FSB</th><th>ISA</th><th>PowerNow!</th><th>Typ. Appl</th></tr>
<tr><td rowspan="16">K7</td><td>Mod.1</td><td>6/1999</td><td>Argon</td><td rowspan="2">Athlon</td><td>0.25</td><td colspan="2" rowspan="2">In-package integrated 512 KB</td><td rowspan="16">DDR FSB</td><td rowspan="5">Enh. 3DNow!</td><td rowspan="5">-</td><td rowspan="2">DT</td></tr>
<tr><td>Mod.2</td><td>11/2000</td><td>Pluto/Orion</td><td rowspan="7">0.18</td></tr>
<tr><td rowspan="2">Mod.3</td><td>6/2000</td><td rowspan="2">Spitfire</td><td>Duron</td><td rowspan="2">64 KB</td><td rowspan="15">On-die (exclusive)</td></tr>
<tr><td>1/2001</td><td>Mobile Duron</td><td>M</td></tr>
<tr><td>Mod.4</td><td>6/2000</td><td>Thunderbird</td><td>Athlon</td><td rowspan="3">256 KB</td><td>DT</td></tr>
<tr><td rowspan="3">Mod.6</td><td>7/2001</td><td rowspan="3">Palomino</td><td>Mobile Athlon4</td><td rowspan="14">3DNow! Prof./SSE</td><td>PowerNow!</td><td>M</td></tr>
<tr><td>10/2001</td><td>Athlon XP</td><td rowspan="3">-</td><td>DT</td></tr>
<tr><td>4/2001</td><td>Athlon MP</td><td>S</td></tr>
<tr><td rowspan="2">Mod.7</td><td>8/2001</td><td>Morgan</td><td>Duron</td><td rowspan="2">64 KB</td><td>DT</td></tr>
<tr><td>1/2001</td><td>Camaro</td><td>Mobile Duron</td><td>PowerNow!</td><td>M</td></tr>
<tr><td rowspan="4">Mod.8</td><td>11/2002</td><td rowspan="2">Thoroughbred</td><td>Duron</td><td rowspan="8">0.13</td><td rowspan="4">-</td><td>DT</td></tr>
<tr><td>3/2003</td><td>Athlon XP</td><td rowspan="3">256 KB</td><td>S</td></tr>
<tr><td>8/2003</td><td>Applebread</td><td>Athlon MP</td><td rowspan="2">M</td></tr>
<tr><td>4/2002</td><td>Thoroughbred</td><td>Mobile Athlon XP-M</td></tr>
<tr><td rowspan="4">Mod.10</td><td>9/2003</td><td>Thorton</td><td rowspan="2">Athlon XP</td><td>DT</td></tr>
<tr><td>9/2003</td><td rowspan="3">Barton</td><td rowspan="3">512 KB</td></tr>
<tr><td>5/2003</td><td>Athlon MP</td><td>S</td></tr>
<tr><td>3/2003</td><td>Mobile Athlon XP-M</td><td>PowerNow!</td><td>M</td></tr>
</table>

The Hammer family

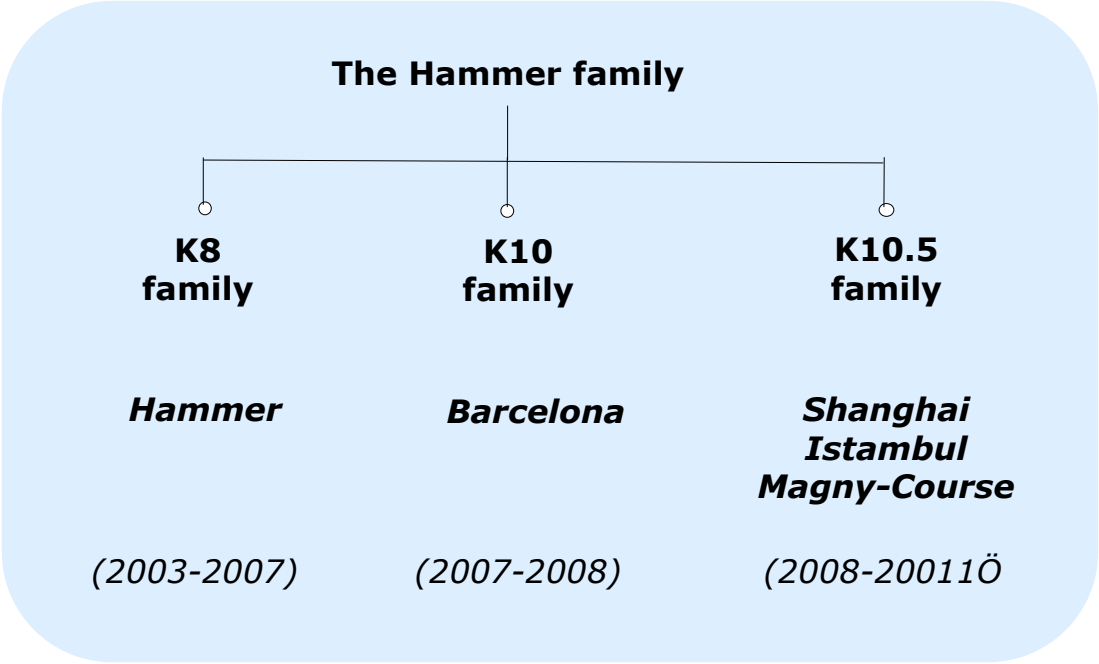

1. Introduction to AMD's processor families (32)

Brand names of AMD's 64-bit K8 – Family 10.5h processor lines

		2003-2007	2007-2008	2008-2011	2009	2009
		K8 (Hammer)	**K10 (Barcelona)**	**K10.5 (Shanghai)**	**K10.5 (Istanbul)**	**K10.5 (Magny- Course)**
Servers	**4P servers**	**See Section 4**	**Barcelona (834x-836x))**	**Shanghai (837x-839x)**	**Istambul (8410-8430)**	**Magny-Course (6100)**
	2P servers		**Barcelona (234x-236x)**	**Shanghai (237x-239x)**	**Istambul (241x-243x)**	**Lisbon (4100)**
	1P servers		**Budapest (135x-136x)**	**Suzuka (138x-139x)**		
Desktops	**High perf.** (~80-120W)		Phenom X4-X2	Phenom II X4-X2	Phenom II X6-X4	
	Mainstream (~60-90W)	Athlon 64 Athlon 64 X2	Athlon X2	Athlon II X4-X2		
	Value (~40-60W)	Sempron		Sempron		
Mobiles	**High perf.** (~30-40W)	Turion 64 X2 (TL 6/5) Turion 64 (ML/MT)		Phenom II (N/P 9xx-6xx) Turion II Ultra (M6xx) Turion II (M/N/P 5xx)		
	Mainstream (~20-30W)	Athlon 64 X2 (TK-5x/4x) Athlon 64 (2xxx+-4xxx+)		Athlon II (M/N/P 3xx) Sempron (M1xx)		
	Ultraportable (~10-20W)	Mobile Sempron (2xxx+-4xxx+) Sempron 2100 fanless		Turion II Neo (K6xx) Athlon II Neo (K1xx) V-series (V1xx)		
Embedded (~10-20W)				Turion II Neo X2 Athlon II Neo X2 Athlon II Neo		

Overview of AMD's K8-based processor lines

The 130 nm – 90 nm K8-based lines [13]

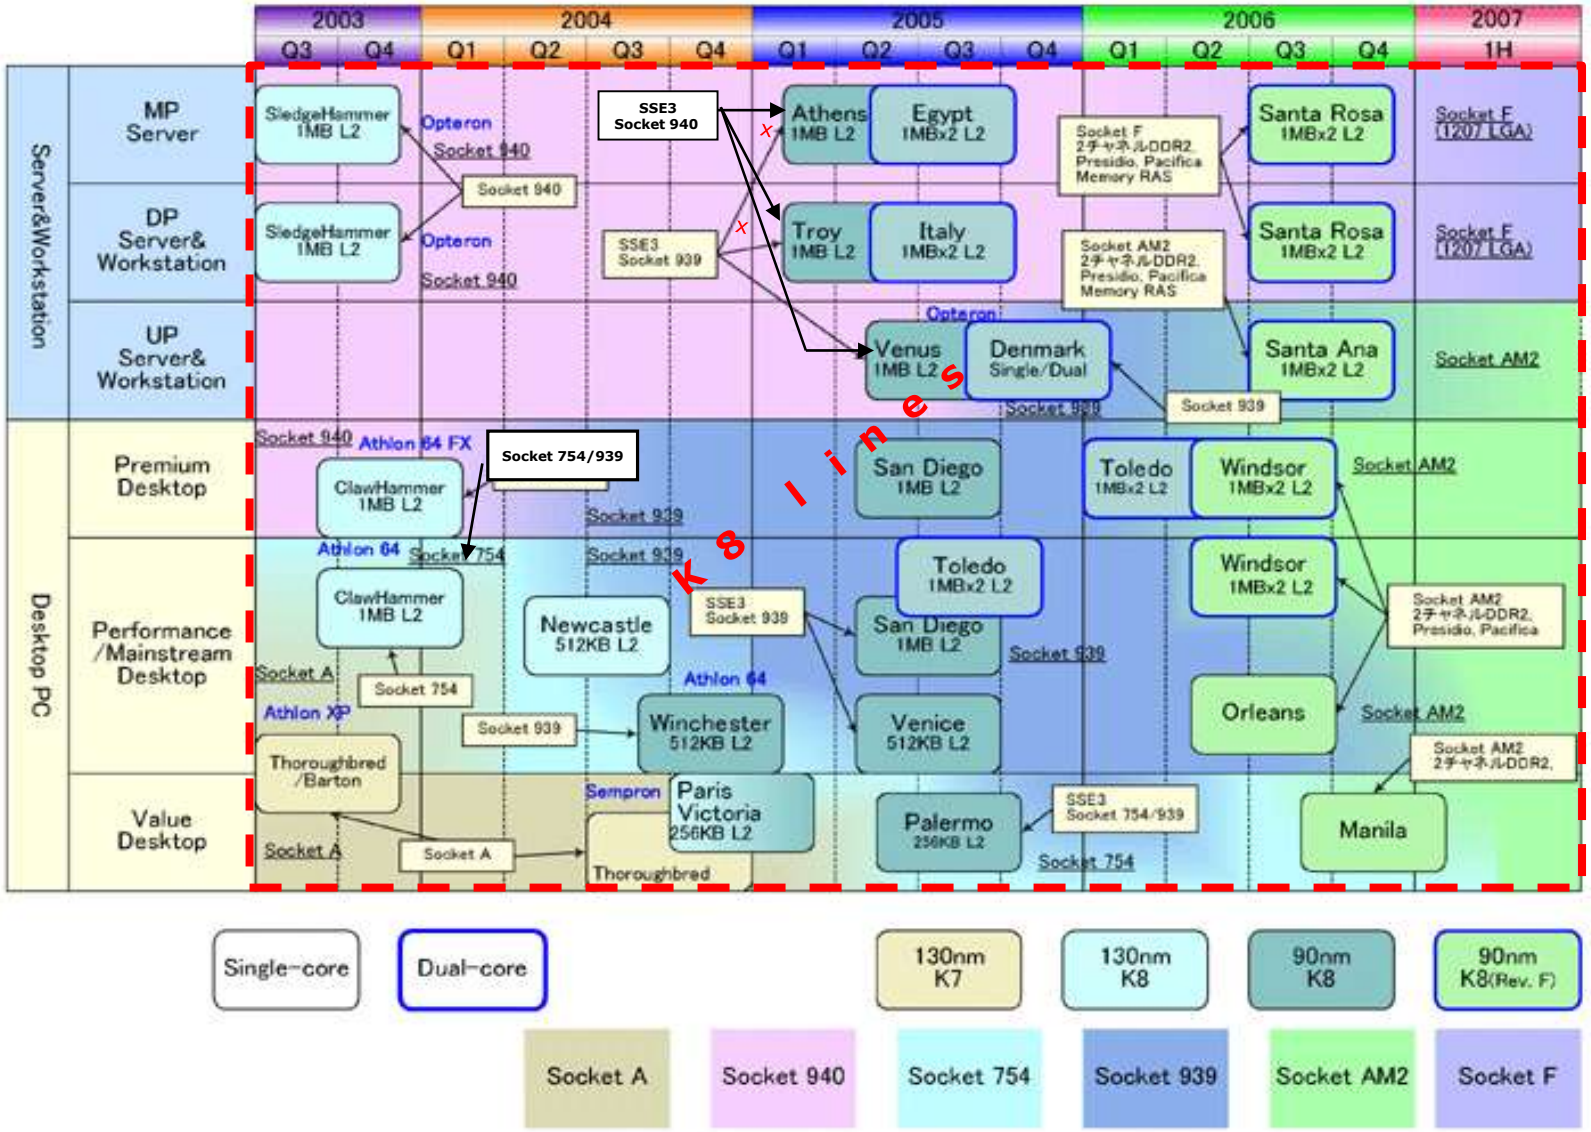

1. Introduction to AMD's processor families (34)

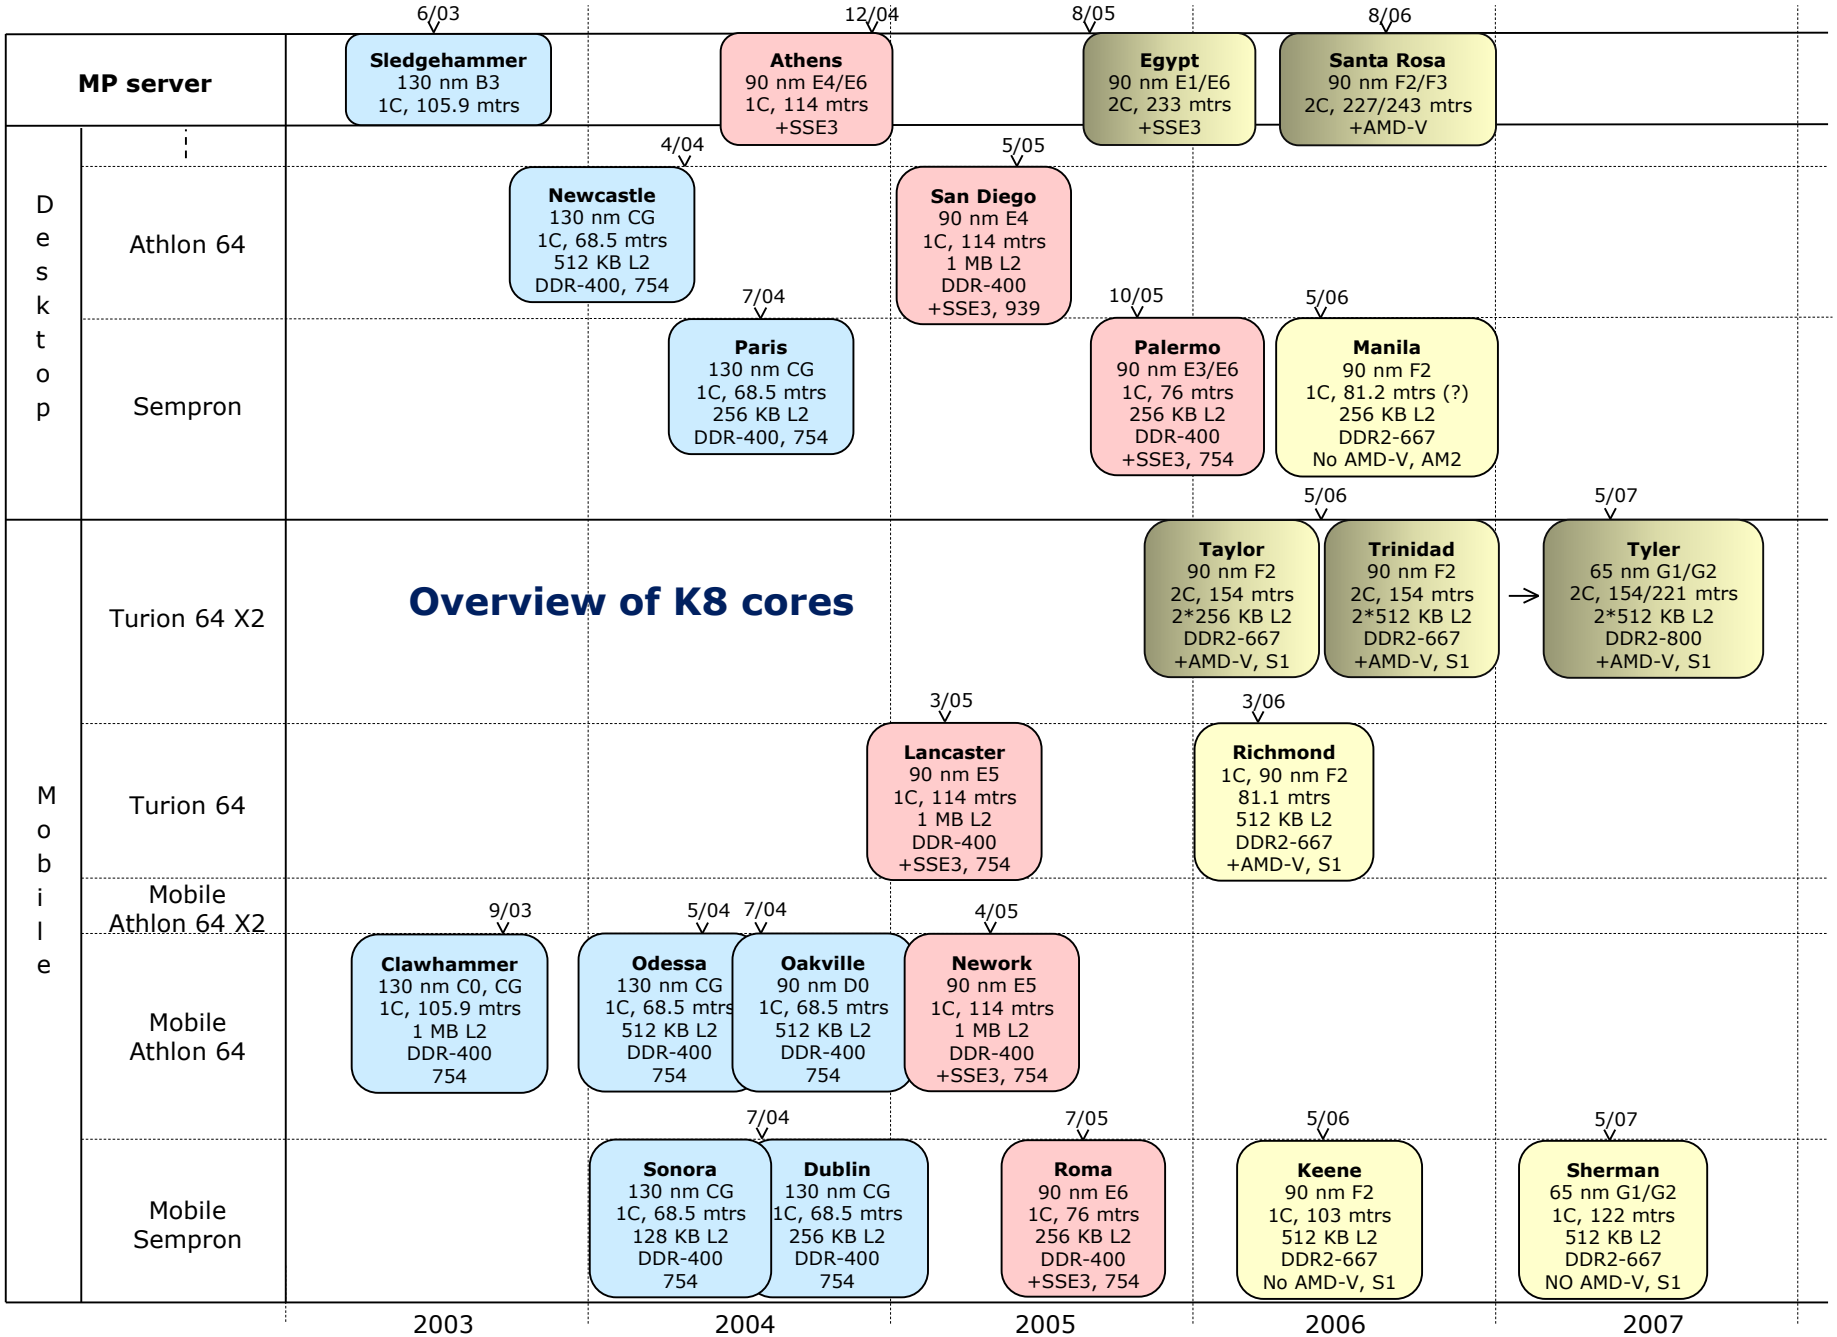

1. Introduction to AMD's processor families (35)

Overview of AMD's K10/K10.5-based processor lines [14]

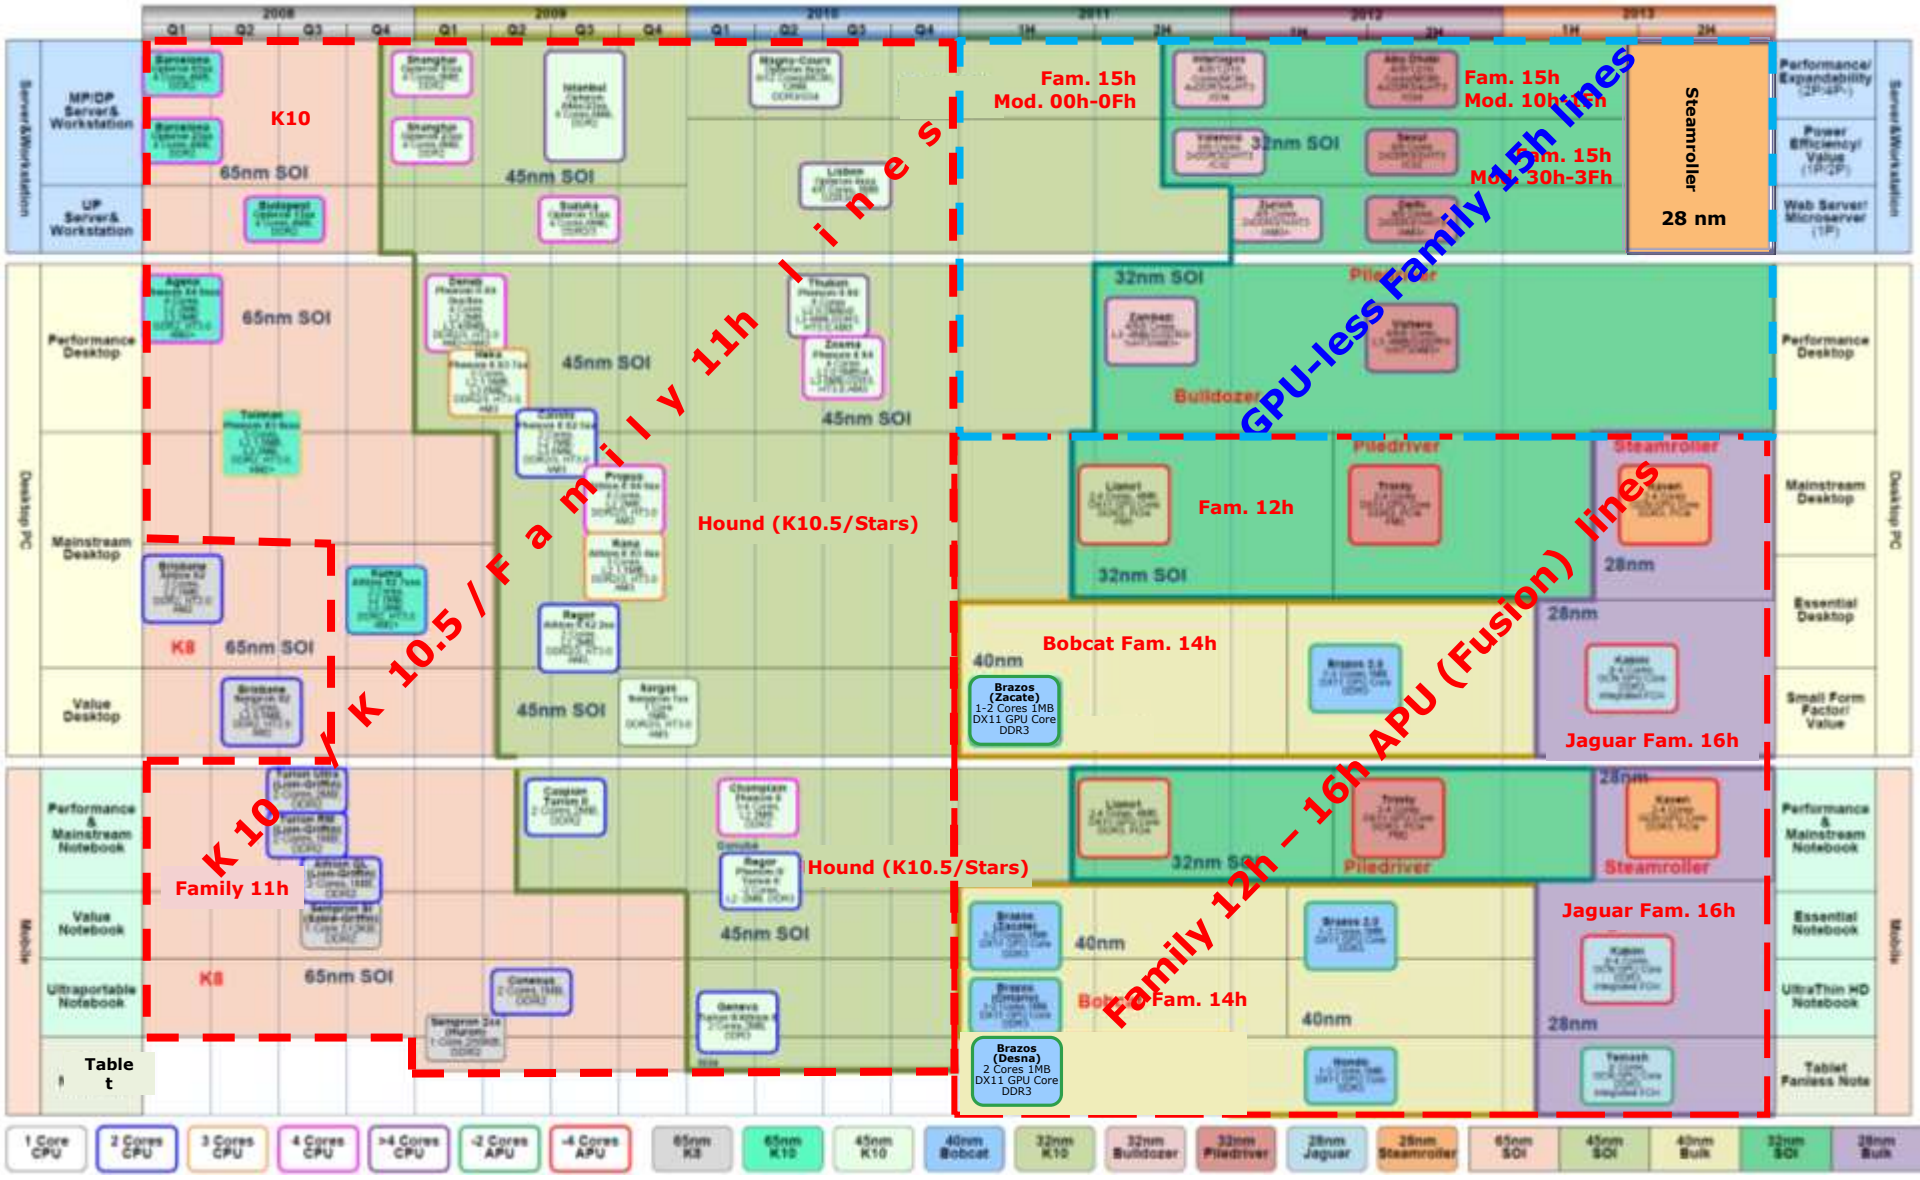

AMD's intermediate families

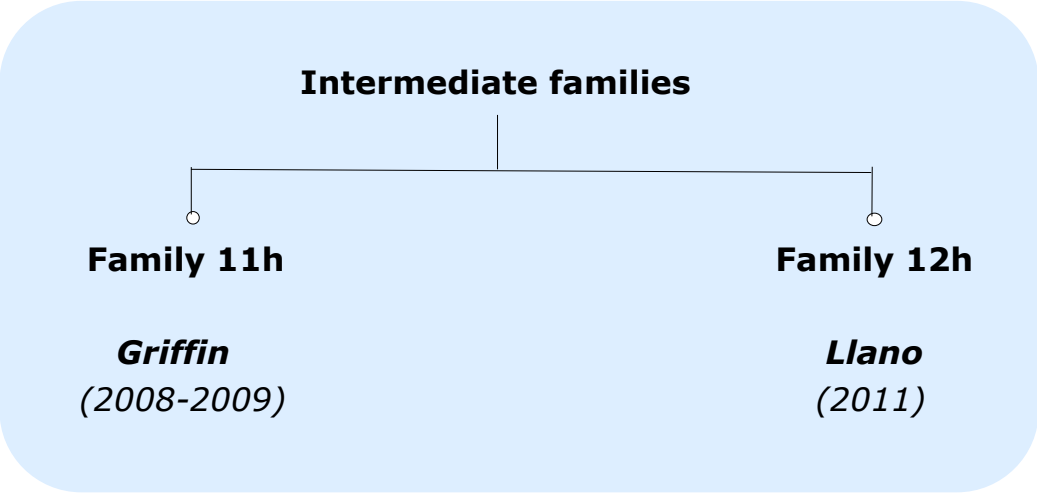

1. Introduction to AMD's processor families (37)

Brand names of AMD's Intermediate (Family 11h – Family 12h) processor lines

	Launched in	2008-2009	2011
		Family 11h (**Griffin**)	**Family 12h** (**Llano**)
Servers	**4P servers**		
	2P servers		
	1P servers		
	(85-140 W)		
Desktops	**High perf.** (~95-125 W)		
	Mainstream (~65-100 W)		Llano A8/A6/A4/E2 Sempron X2
	Entry level (40-60 W)		
Notebooks	**High perf.** (~30-60 W)	Turion X2 Ultra (ZM-xx) Turion X2 (RM-xx)	Llano A8 M
	Mainstream/Entry (~20-30 W)	Athlon X2 (QL-xx) Sempron (SI-xx)	Llano A6/A4/E2 M
	Ultra portable (~10-15 W)	Turion Neo X2 (L6xx) Turion X2 (RM-xx) Athlon Neo X2 (L3xx) Sempron (200U/210U)	
Tablet (~5 W)			
Embedded (~10 – 20 W)		Turion Neo X2 (L6xx) Athlon Neo X2 (L3xx) Sempron (200U/210U)	

The Bulldozer family

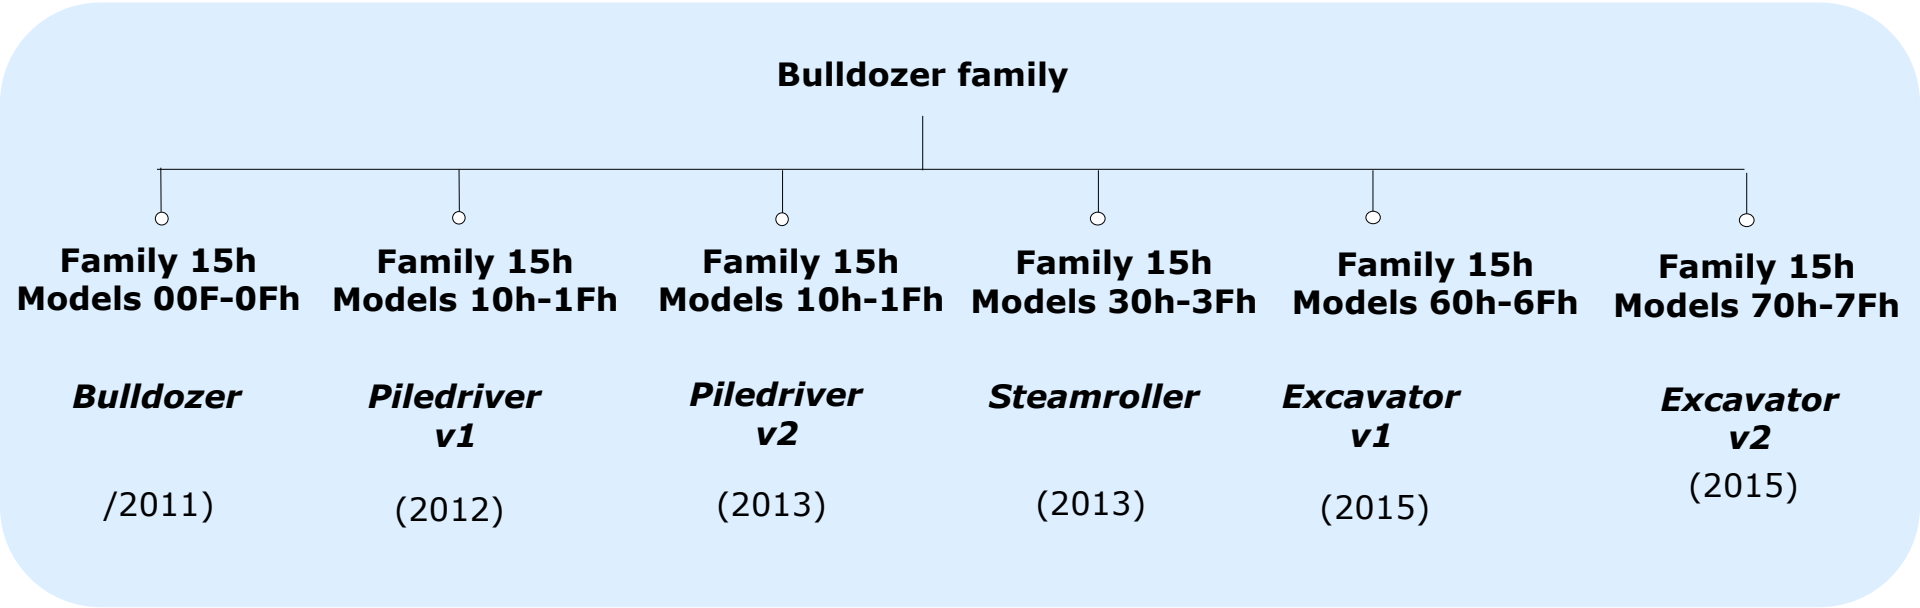

Brand names of AMD's Bulldozer-based processor lines

	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 (~65-100 W)		Trinity A10-A4 Series	Richland A10-A4 Series	Kaveri A10-A8		
	Entry level (~40-60W)						
Tablets/Notebooks	**High perf.** (~30-40 W)		Trinity A10 M-A6 M	Richland A10 M-A4 M	Kaveri FX/A10/A8		
	Mainstream/ Entry level (~20-30 W)		Trinity A10 M-A6 M	Richland A10 M-A4 M	Kaveri FX/A10/A8/A6	Carrizo FX/A12-A6	Bristol Ridge FX/A12/A10
	Ultra portable (~10 – 15 W)						Stoney Ridge A9/A6/A4/E2
	Tablet (~5 W)						

The Cat family

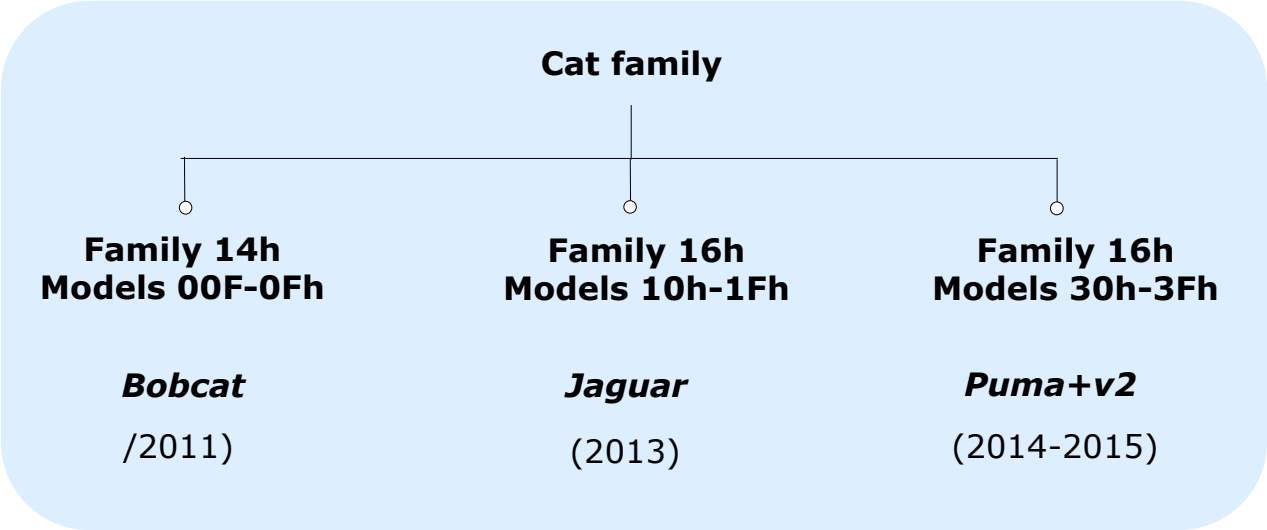

Brand names of AMD's Family 12h, 14h and 16h-based processor lines

	Launched in	2011	2012	2013	2014	2015
		Family 14h (00h-0Fh) (**Bobcat**)	**Family 14h** (00h-0Fh) (**Bobcat**)	**Family 16h** (00h-0Fh) (**Jaguar**)	**Family 16h** (30h-3Fh) (**Puma+**)	**Family 16h** (30h-3Fh) (**Puma+**
Servers	**4P servers**					
	2P servers					
	1P servers					
	(85-140 W)					
Desktops	**High perf.** (~95-125 W)					
	Mainstream (~65-100 W)					
	Entry level (~30-60 W)					
Notebooks	**High performance/ mainstream/entry** (~30-60 W)			Kabini A6		
	Ultra portable (~10-15 W)	Zacate E-Series Ontario C-Series	Zacate E1/E2	Kabini A/E-Series	Beema A/E-Series	Carrizo-L A/L-Series
Tablet (~5 W)		Desna Z-Series		Temash A Series	Mullins A Series/E1	

The Zen family

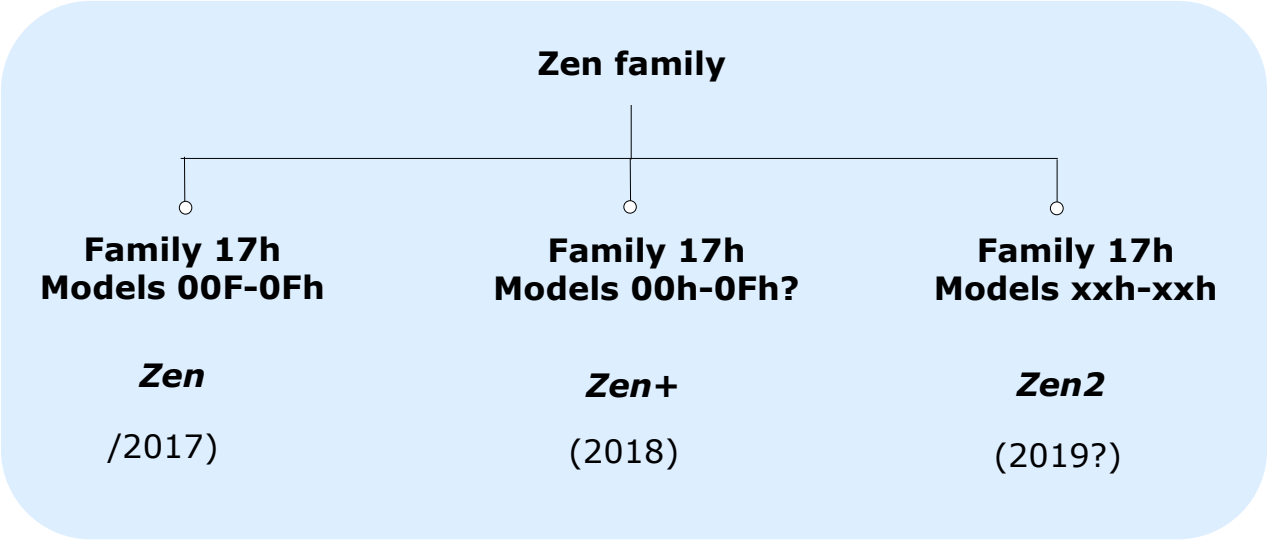

Brand names of AMD's Family 17h-based processor lines

	Launched in	2017-2018	2018	2019?
		Family 17h (00h-0Fh) (**Zen**)	**Family 17h** (00h-0Fh) (**Zen+**)	**Family 17h** (xxh-xxh) (**Zen2**)
Servers	**4P servers**			
	2P servers	Epyc 7xx1		
	1P servers	Epyc 7xx1P		
	(85-140 W)			
Desktops	**High perf.** (~95-125 W)	ThreadRipper (TR 1xxxX)	ThreadRipper (TR 2xxxX/WX)	
	Mainstream/ Entry level (30-95 W)	Summit Ridge (Ryzen 7/5/3 1xxx/1xxxX) Raven Ridge (APU) (Ryzen 7/5/3 2000G/GE)	Pinnacle Ridge (Ryzen 7/5 2xxx/2xxxX)	
Notebooks	**High perf.** (~30-60 W)			
	Mainstream/Entry (~20-30 W)			
	Ultra portable (~10-15 W)	Raven Ridge (APU) (Ryzen 7/5/3 2x00U)		
Tablet (~5 W)				

1. Introduction to AMD's processor families (44)

Main features of AMD's 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)	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 (NPT)	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. 17h Mod. 00h-0Fh		6/2017	Epyc (2S!!) (4 dies/proc.)	7000	14 nm	4x(2x4) (32C)	½ MB/C	2 MB/C	DDR4-2666	IFIS 75.8 GB/s	SP3

Evolution of main features of AMD's DP/MP servers

Base arch./ stepping		Intro.	Core	Techn. (nm)	Server family name	Cores	New key features						Use
							L3	Mem.	On-die MC	HT	ISA extension	NX	
K8	B3/CG	4/2003	Sledge-hammer	130	Sledgehammer	1C	-	DDR	+On-die MC	3xHT 1.0	+SSE2	+NX -bit	S/DT/M
	E4	12/2004		90	Athens					3xHT 2.0	+SSE3		
	E1/E6	4/2005			Egypt	2C							
	F2/F3	8/2006			Santa Rosa			DDR2					
	G1/G2	12/2006		65	DT: Brisbane								DT
K10	B2/B3	9/2007	Greyhound	65	Barcelona[5]	4C	6 MB				+SSE4a		S/DT
K10.5	C2/C3	11/2008	Greyhound +	45	Shanghai					3xHT 3.0			S/DT/M
	CE	6/2009			Istambul	6C				4xHT 3.1			S/DT
	D1	3/2010			Magny Course (2xIstambul)	2x6C	2x 6 MB	DDR3					S
Fam. 15h	Mod. 0xh	11/2011	Bulldozer	32	Interlagos (2xOrochi)	2x8C	2x 8 MB				+SSE4.1/4.2, AES, AVX, XOP, FMA4,CMUL		S/DT
	Mod. 1xh	11/2012	Piledriver		Abu Dhabi (Dual dies)						+FMA3, CVT16, BMI, TBM		S/DT/M
Fam. 17h	Mod. 0xh	6/2017	Zen	14	DT: Epic (4 dies)	4x8C	2 MB/ core	DDR4		IFIS	na.		S/HED

[1] x4UMI: 4x PCIe 2.0

[2] ISA enh.: +AES, +AVX, +FMA4, +XOP, +PCLMULQDQ

[3] PCIe 1.0/2.0

[4] 3DNow! Prof. dropped

[5] The Barcelona die supports already 4xHT 3.0 and DDR3 but Socket F used for DP/MP servers restricts supported features to 3xHT 2.0 and DDR2

Overview of subsequent K10/K10.5 implementations (as used in MP/DP servers) [88]

MP/DP Platforms – 8000 and 2000 Series

65 nm
45nm

Platform Segment	2008	2009	2010
CPU	**"Barcelona"** 4-Core • 2M L3 • RDDR-2 • 3x HT-1 • AMD-V™ • 65nm	**"Shanghai"** 4-Core • 6M L3 • RDDR-2 • cHT-3 • AMD-V • 45nm	**"Istanbul"** 6-Core 2H09 • 6M L3 • RDDR-2 • cHT-3 • AMD-V • 45nm
			"Magny-Cours" 1H10 12-Core • 12M L3 • Probe Filter • 4x HT-3 • HTC • U/R DDR-3 • APML • 45nm • AMD-V
			"Sao Paulo" 1H10 6-Core • 6M L3 • Probe Filter • 4x HT-3 • HTC • U/R DDR-3 • APML • 45nm • AMD-V
Chipset	Nvidia nForce 3600/3050 Broadcom HT-2100/1000		AMD RD890S w/IOMMU AMD RD870S w/IOMMU AMD SB700S
Platform	**Socket F (1207)** • 3x HT-1 (moving to cHT-3) • DDR-2 (Dual Channel)		**"Maranello"** • 4x HT-3 • DDR-3

Key parameters of subsequent versions of the HyperTransport standard [58]

HT version	Year	Max. HT frequency	Max. link width	Max. bandwidth at 16-bit unidirectional
1.0	2001	800 MHz	32-bit	3.2 GB/s
1.1	2002	800 MHz	32-bit	3.2 GB/s
2.0	2004	1.4 GHz	32-bit	5.6 GB/s
3.0	2006	2.6 GHz	32-bit	10.4 GB/s
3.1	2008	3.2 GHz	32-bit	12.8 GB/s

1. Introduction to AMD's processor families (48)

Main features of AMD's high-performance desktop lines (except Bulldozer-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. 17h (Zen)		3/2017	Summit Ridge	Ryzen 7	14 nm	8	8x1/2 MB	16 MB	DDR4-2993	-	AM4
Fam. 17 (Zen+)		4/2018	Pinnacle Ridge	Ryzen 7	12 nm	8	8x1/2 MB	16 MB	DDR4-2933	-	AM4

Main features of 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. 15h based lines		-	-	-	-	-	-	-	-	-	-

Main features of Hammer (K8 – K 10.5)-based high performance mobile lines

Base arch./ stepping		Intro	High perf. mobile family name	Series	Techn.	Core count (up to)	L2 (up to)	L3	Memory (up to)	HT/ dir. (up to)	Sock et
K8	C0, CG	9/2003	Clawhammer	Mobile Athlon 64	130 nm	1	512 KB	-	DDR-400	HT 1.0: 3.2 GB/s	754
	E5	3/2005	Lancaster	Turion 64	90 nm	1	1 MB	-	DDR-400	HT 1.0: 3.2 GB/s	754
	F2	5/2006	Trinidad	Turion 64 X2	90 nm	2	2*512 KB	-	DDR2-667	HT 1.0: 3.2 GB/s	S1
K10	-	-	-	-	-	-	-	-	-	-	-
K10.5	DA-C2	9/2009	Caspian	Turion II	45 nm	2	2*512 KB/ 2*1 MB[1]	-	DDR2-800	HT 3.0: 7.2 GB/s	S1g3
	DA-C3	5/2010	Champlain	Turion X4	45 nm	4	4*512 KB	-	DDR3-1066	HT 3.0: 7.2 GB/s	S1g4

[1]: 2*512 KB for Turion II, 2*1 MB for Turion II Ultra

Main features of the Intermediate-based high-performance mobile lines

Base arch./ stepping	Intro	High perf. mobile family name	Series	Techn.	Core count (up to)	L2 (up to)	L3	Memory (up to)	HT/ dir. (up to)	Sock et
Family 11h (K11) (Griffin)	6/2008	Lion (no APU) (not SoC)	Turion X2 Ultra	65 nm	2	2x512 KB/ 2*1 MB[2]	-	DDR2-800	HT 3.0: 10.4 GB/s	S1g2
Family 12h (K12) (Llano)	6/2011	Llano (APU) (not SoC)	Fusion A8 M	32 nm	4	4x1 MB	-	DDR3-1600	-	FM1

APU: Accelerated Processing Unit (CPU +GPU) CCX: Core CompleX

[2]: 2*512 KB for Turion X2, 2*1 MB for Turion X2 Ultra

UMI: Universal Media Interface

1. Introduction to AMD's processor families (51b)

Main features of Cat and Zen-based ultra-portable mobile lines

Base arch./ stepping	Intro	Ultra-portable mobile family	Series	Techn.	Core count (up to)	L2 (up to)	L3	GPU (APU)	Memory (up to)	TDP [W]	Socket
Family14h (00h-0Fh) (Bobcat)	1/2011	Zacate (not SoC)	E Series	40 nm	2	512 KB/ core Private	-	Yes	DDR3L-1333	18	FT1 (BGA)
	6/2012	Zacate (not SoC)	E1/E2 Models	40 nm	2	512 KB/ core Private	-	Yes	DDR3L-1333	18	FT1 (BGA)
	1/2011	Ontario	C Series	40 nm	2	512 KB/ core private	-	Yes	DDR3-1066	9	FT1 (BGA)
Family 16h (10H-1fH) (Jaguar)	5/2013	Kabini (SoC)	A Series	28 nm	4 cores with a shared L2 cache	2 MB shared	-	Yes	DDR3L-1866	9/ 15	FT3
Family 16h (30H-3fH) (Puma+)	4/2014	Beema (SoC)	A Series	28 nm		2 MB shared	-	Yes	DDR3L-1866	15	FT3b
	5/2015	Carrizo-L (SoC)	A Series	28 nm		2 MB shared	-	Yes	DDR3L-1866	10/ 15	FP4
Family 17h (00H-0fH) (Zen)	10/2017	Raven Ridge (SoC)	Ryzen 7/5/3	14 nm	4-core CCX, private L2 and shared L3 cache(s)	½ MB/ core	1 MB/ core	Yes	DDR4-2400	15	AM4

APU: Accelerated Processing Unit (CPU +GPU) CCX: Core CompleX

[2]: 2*512 KB for Turion X2, 2*1 MB for Turion X2 Ultra

UMI: Universal Media Interface

Remark

AMD's chipsets

- AMD started offering own chipsets to their processors with their 640/645 chipset (yet licensed from VIA) in 1997 to support their K6, Cyrix 6x6 and Pentium processors.
- It was followed by the in-house developed 750 chipset intended for the Athlon Model 1 (1999).
- Since then AMD usually provides own chipsets for their processors.

Introducing the platform concept by AMD-1

- For years AMD disdained the platform approach arguing that unlike Intel's Centrino platform they want to give OEMs the choice to select components from a wide range of suppliers.

  But OEMs prefer platforms since its components are already tested and their integration is aready validated by the manufacturer [6].

  Also manufacturers benefit from the platform concept as it motivates OEMs to buy all key components of a computer system from the same manufacturer.

- For the reasons mentioned, two to three years after Intel, also AMD introduced the platform approach.

- First, AMD announced their platform concept in the mobile segment (like Intel) with their Kite platform in 2006.

  (This platform supported the K8-based dual core Turion 64 X2 and sigle core Turion 64 and Mobile Sempron processors, as shown in the next Figure).

AMD's 2006 mobile roadmap showing their first platform, the Kite mobile platform [130]

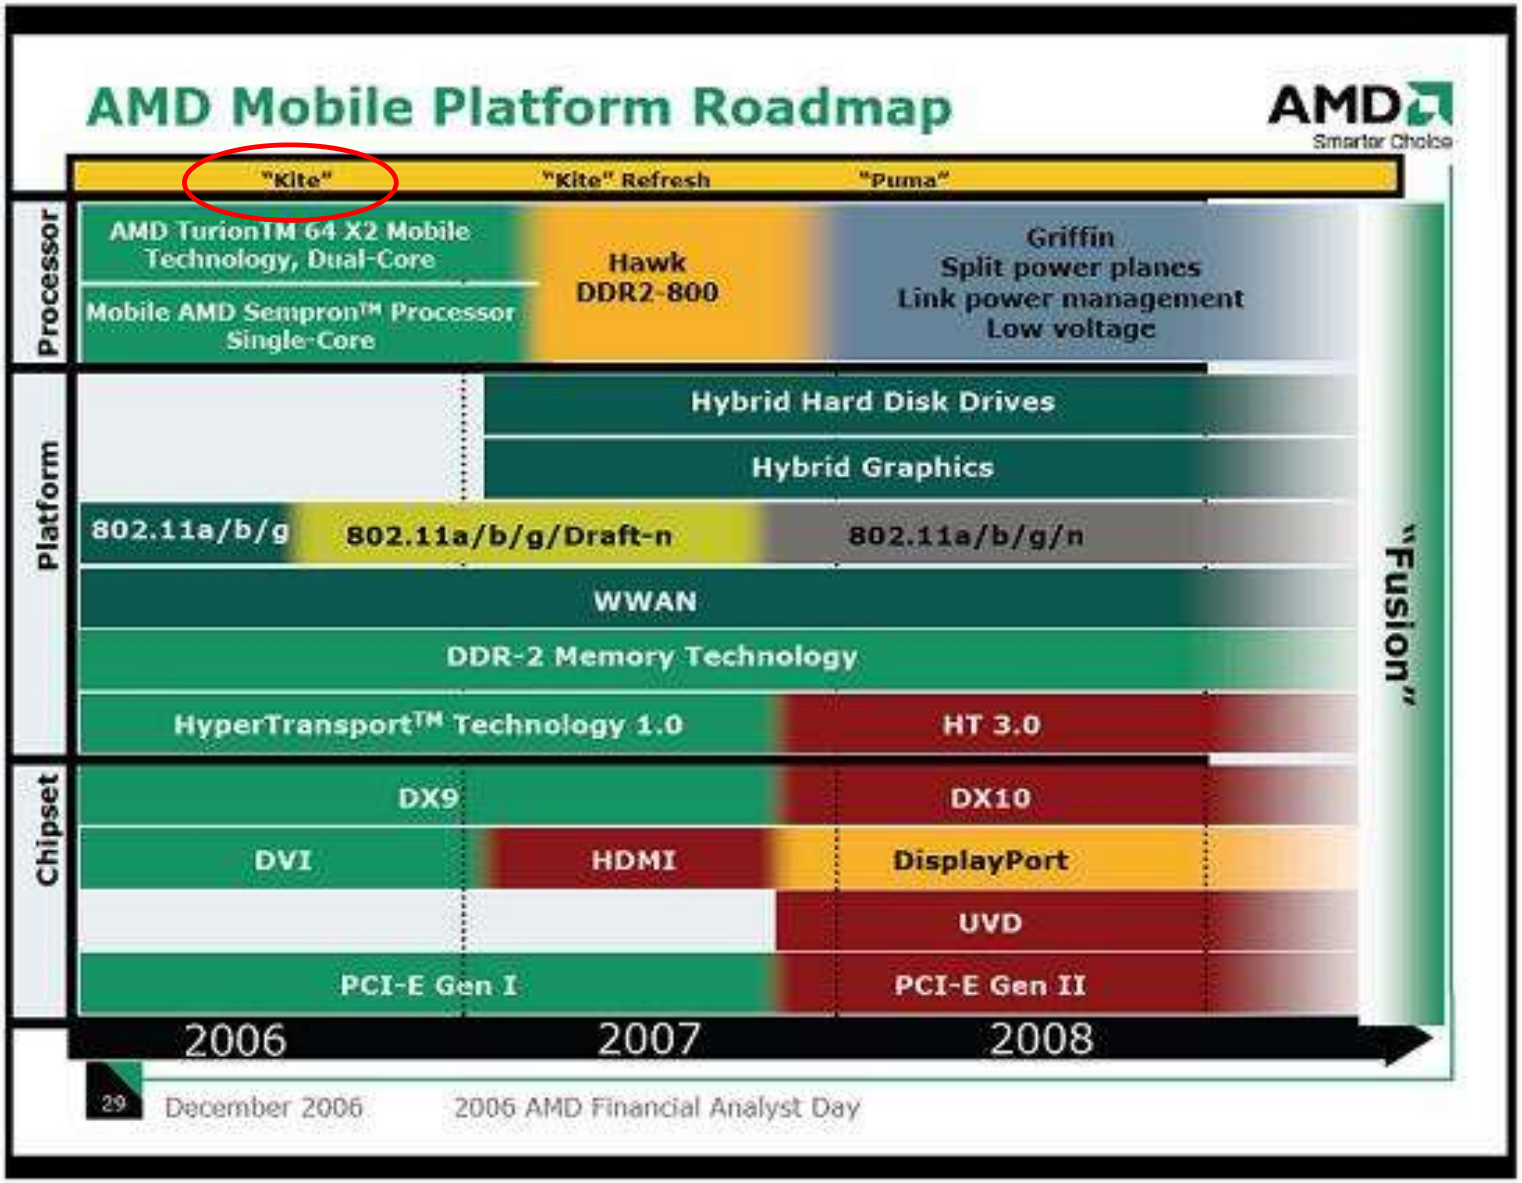

Main features of the Kite platform [7]

Introduced in 2006

AMD mobile	Kite platform
Mobile processor	Processors - Socket S1 • Mobile Sempron single-core 64-bit processor (codenamed *Keene*), or • Turion 64 single-core 64-bit processor (codenamed *Richmond*), or • Turion 64 X2 dual-core 64-bit processor (codenamed *Taylor*, *Trinidad*)
Mobile chipset	• DVI and HyperTransport 1.0 • DDR2-667 SO-DIMM
Mobile support	• Wireless IEEE 802.11 b/g mini-PCIe WiFi adapter

Introducing the platform concept by AMD-2

- In the desktop segment AMD introduced their platform concept in 2007 with their Spider platform.

The Spider platform supported the K10 Barcelon-based Phenom X2/X4 processors, as indicated in the next Figure.

The Phenom X2/X4 processors were built on the K10-based Agena and Kuma desktop cores and targeted gamers).

AMD's first desktop platform: the K10 Barcelona based Spider platform (2007) [8]

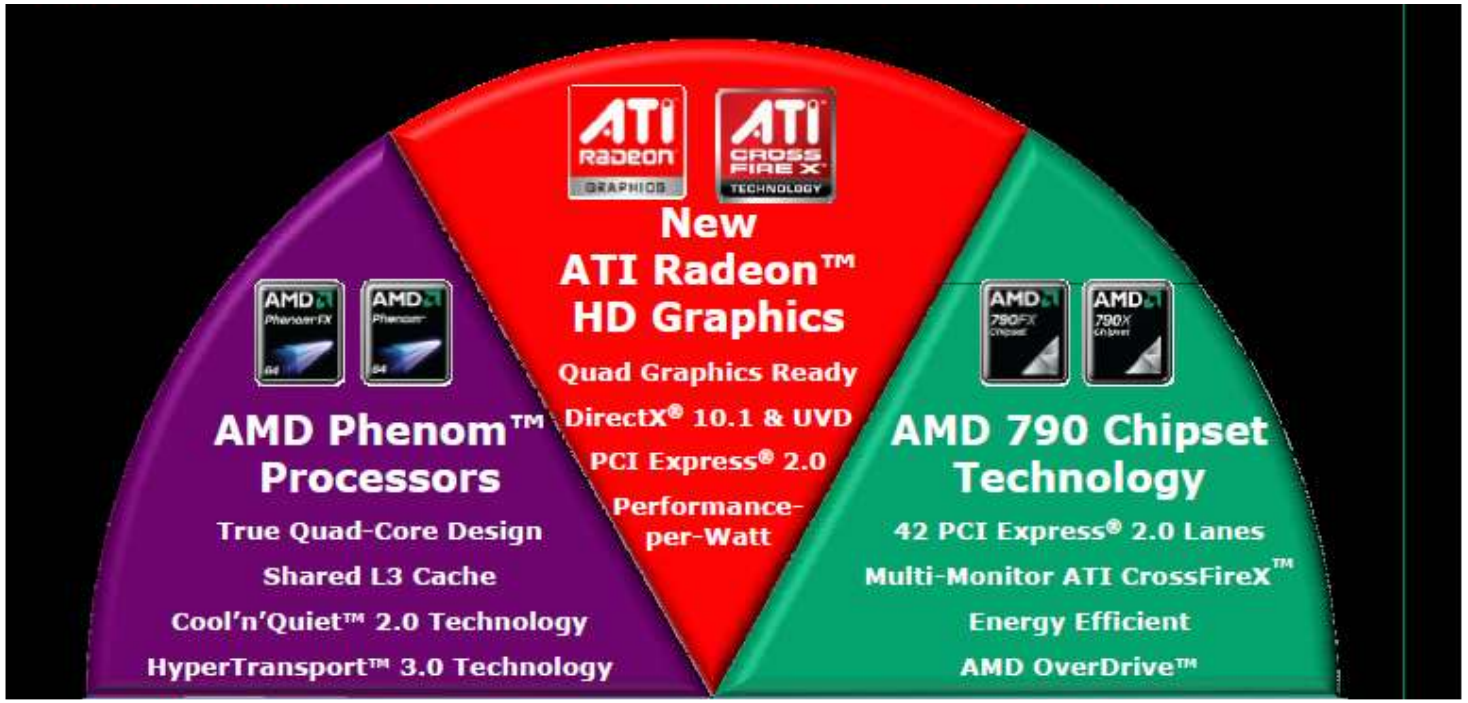

Introducing the platform concept by AMD-3

- AMD's platform concept has a peculiarity, as AMD's desktop platforms cover - beyond the usual main components processor and chipset - also the graphics cards (see next slide).

This reflects AMD's strategy to support aggressively graphics, leading to the acquisition of one of the major graphics card supplier ATI in 2006.

Remark

AMD revealed already in 2006 their desktop roadmap, nevertheless without publishing actual platform designations, as indicated below [131].

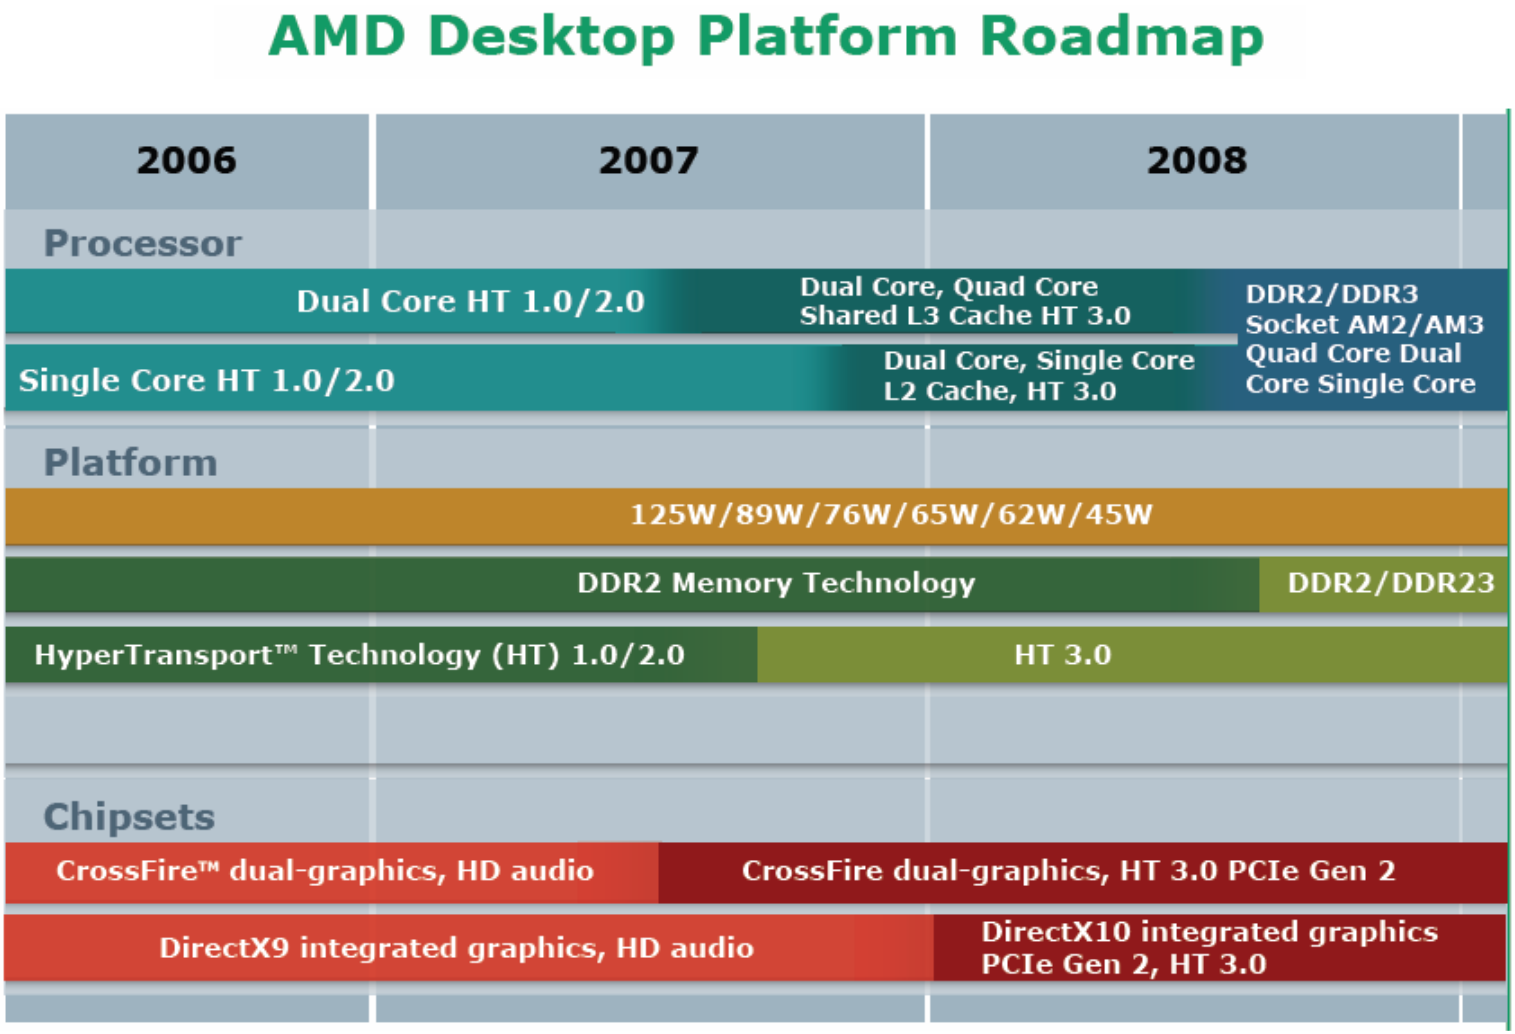

AMD's first desktop roadmap with platform names (published in 2007) [128]

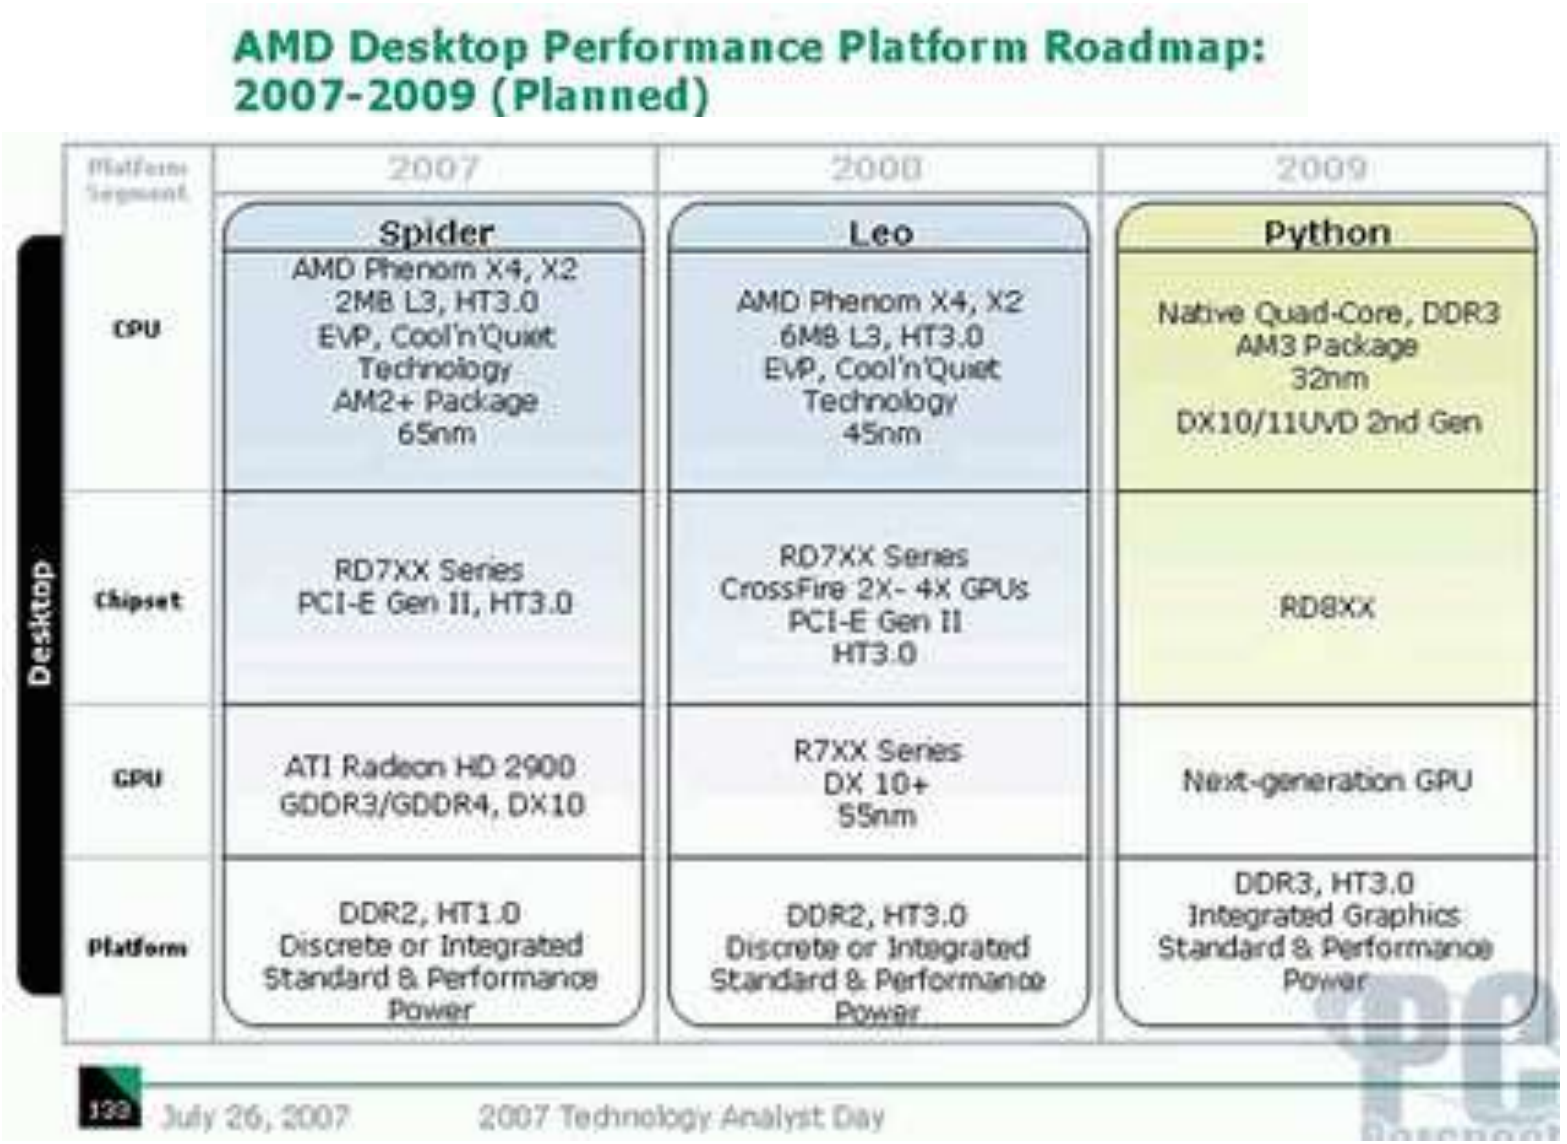

AMD's power management techniques K8 – Family 15h (Bulldozer) (based on [53])

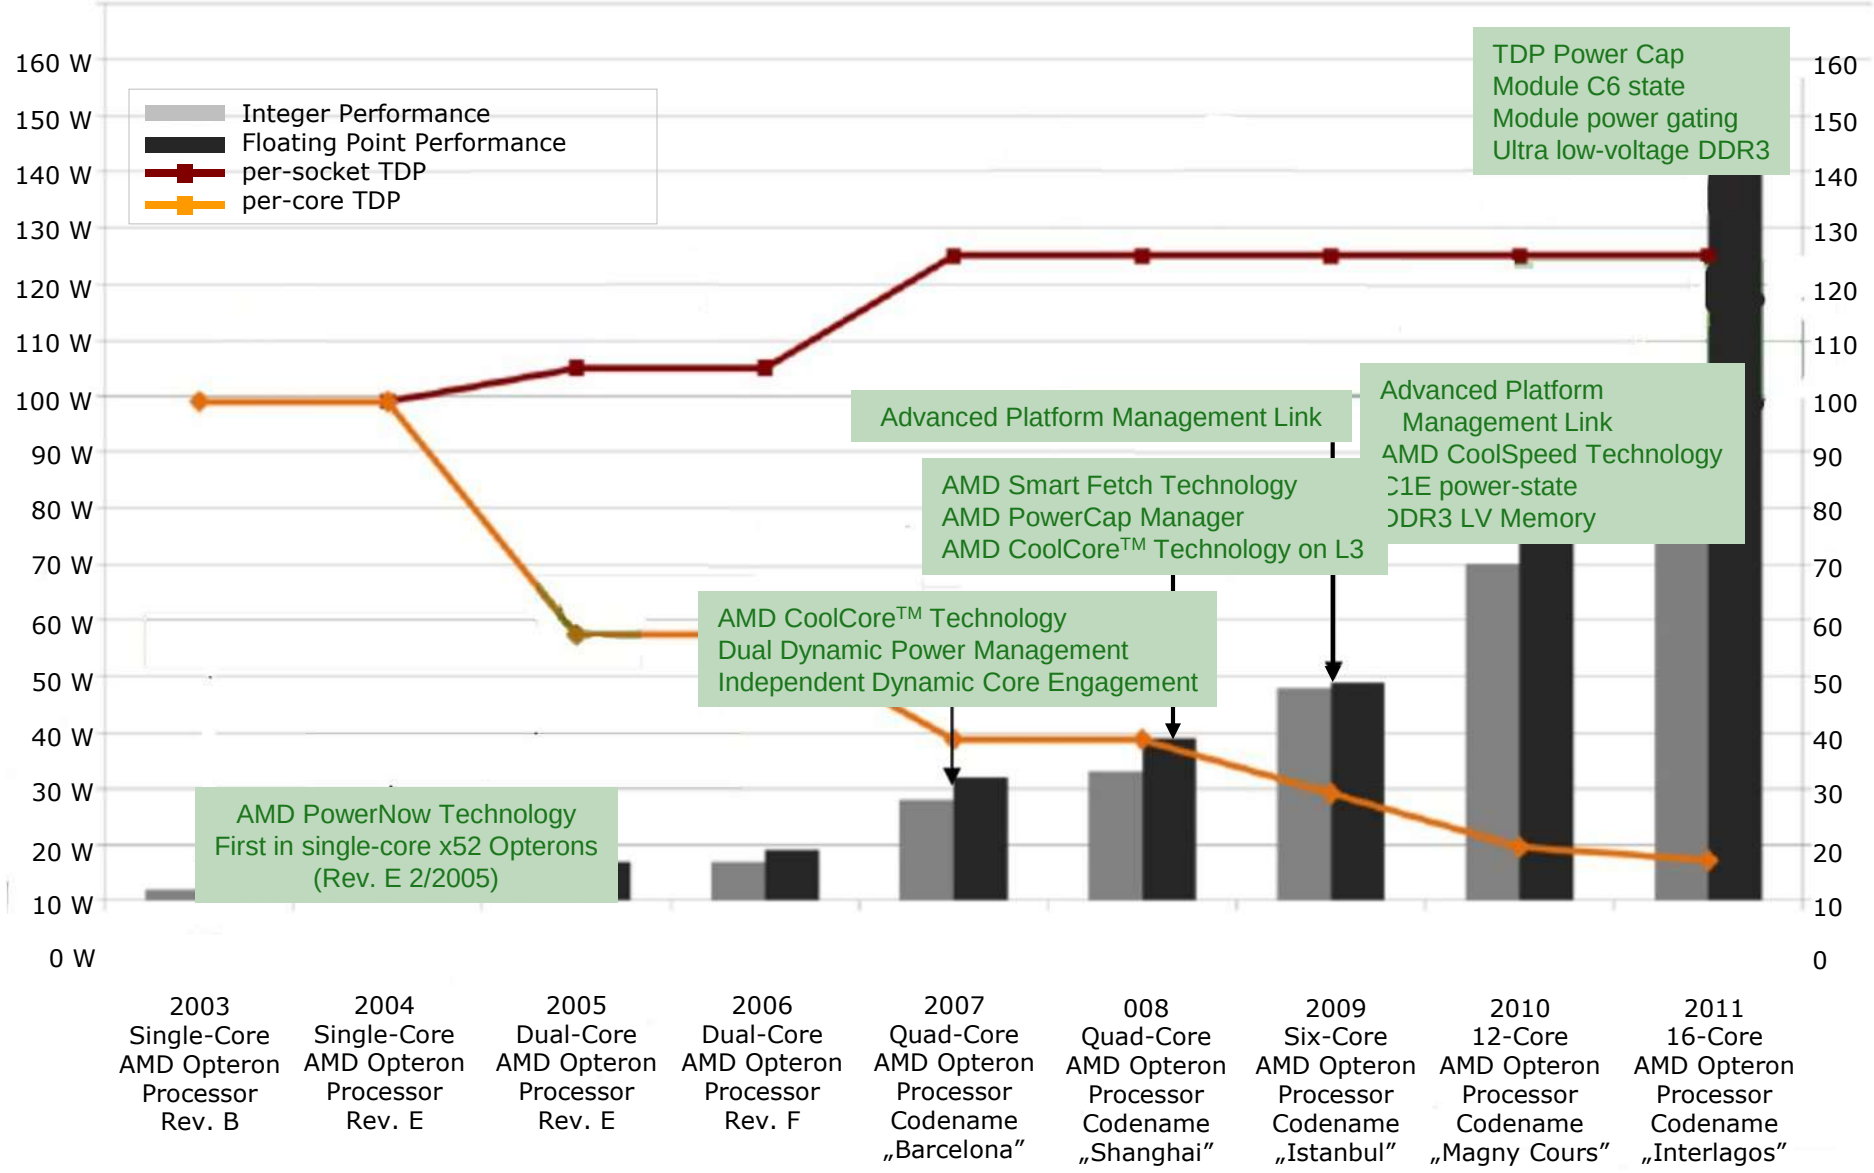

1. Introduction to AMD's processor families (62)

Sockets of AMD's x86-64 processors [based on 42]

Supports DDR3

Socket **G34** (1944 pins, MP servers K10.5 „Magny-Course", Fam. 15h)
Socket **C32** (1207 pins, DP servers K10.5 „Magny –Course, Fam. 15h)
Socket **AM3** (941 pins, K10.5 Shanghai UP servers, desktops, except first released desktops, K10.5 Istambul desktops)
Socket **AM3+** (942 pins, UP servers Fam. 15h, desktops Fam. 15h)
Socket **FM1** (905 balls BGA, desktops, mobiles Fam. 12h)
Socket **S1g4** (638 pins, K10.5 mobile)
Socket **FT1** (413 balls BGA, desktops, mobiles Fam. 14h)
Socket **FT3** (769 balls BGA, desktops, mobiles Fam. 16h)

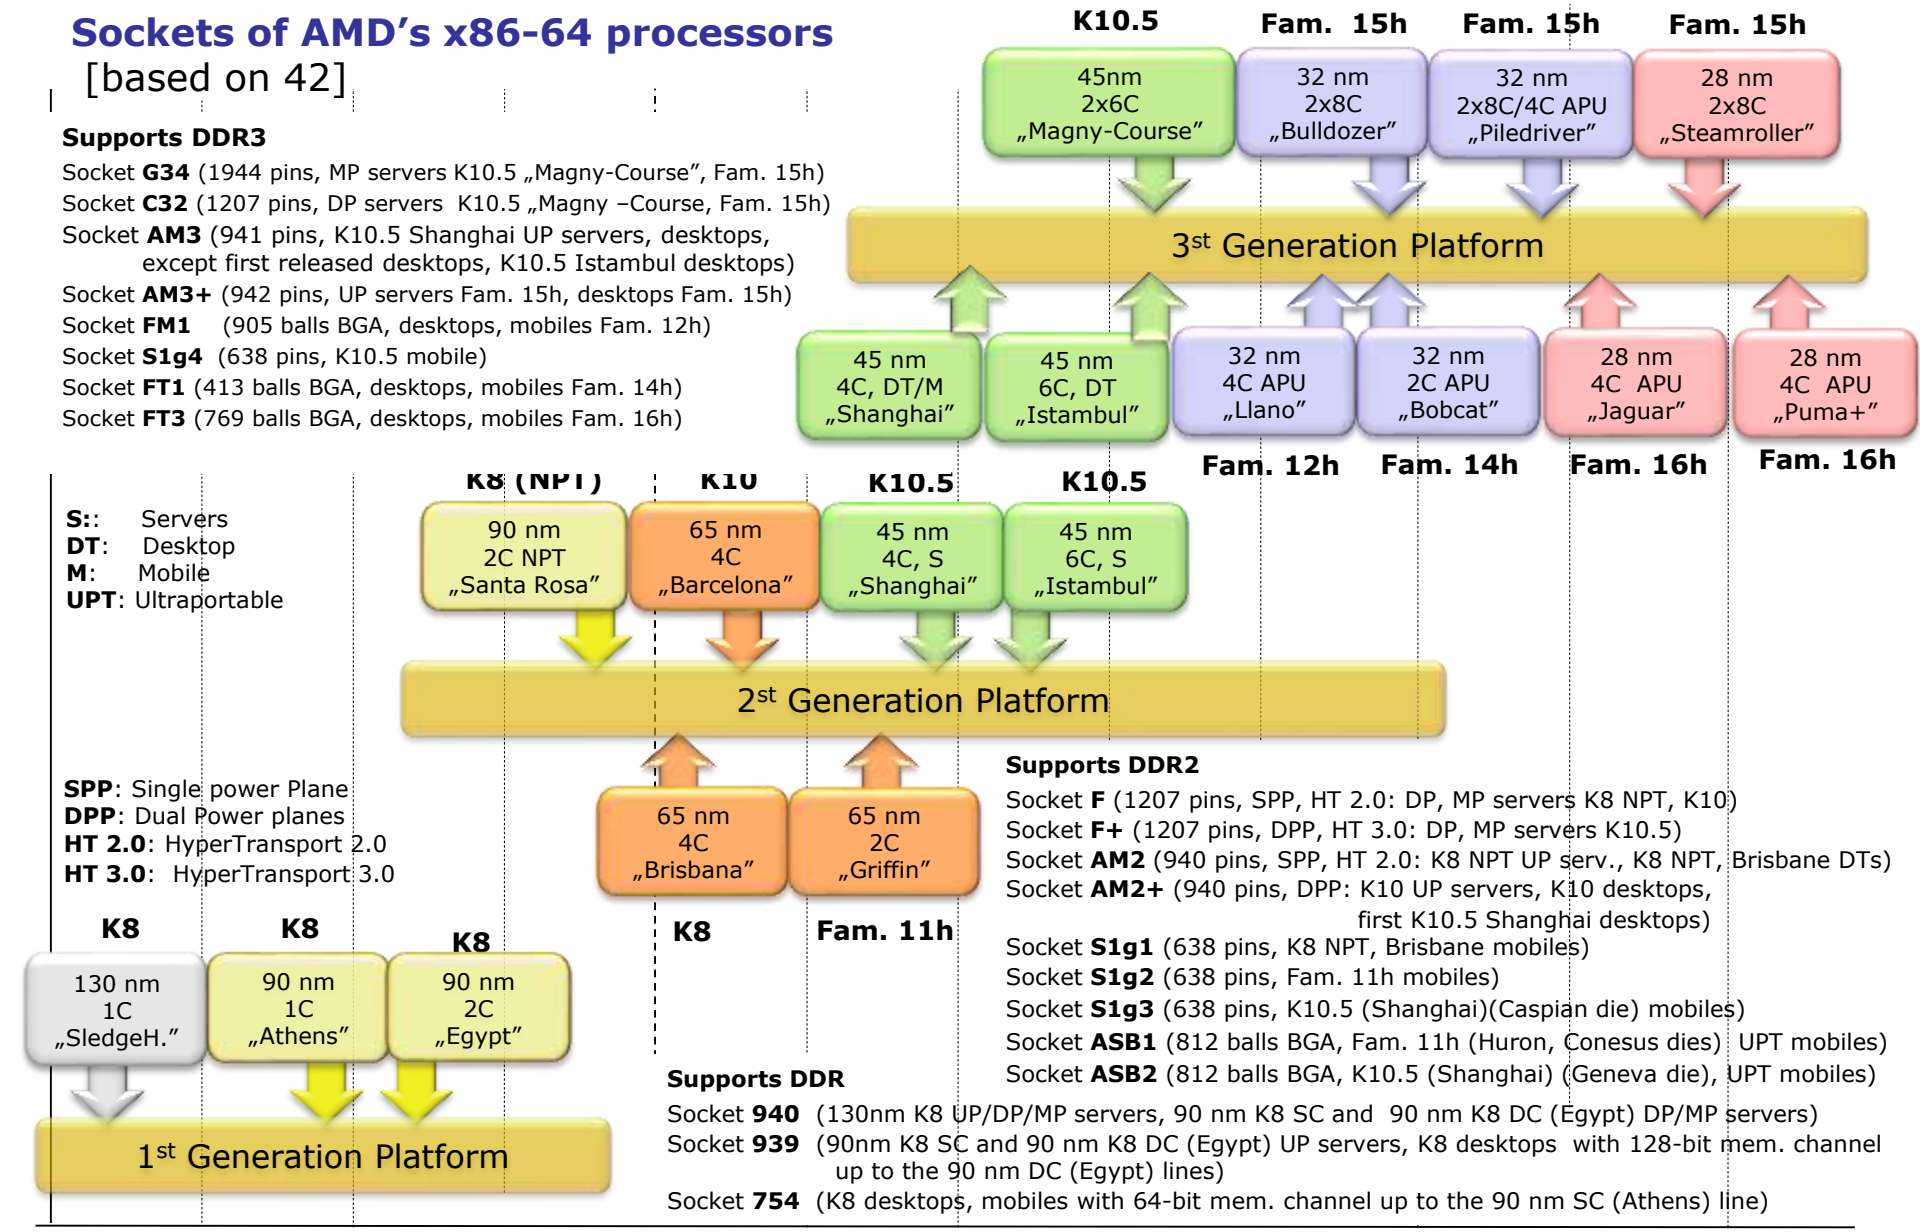

Supports DDR2

Socket **F** (1207 pins, SPP, HT 2.0: DP, MP servers K8 NPT, K10)
Socket **F+** (1207 pins, DPP, HT 3.0: DP, MP servers K10.5)
Socket **AM2** (940 pins, SPP, HT 2.0: K8 NPT UP serv., K8 NPT, Brisbane DTs)
Socket **AM2+** (940 pins, DPP: K10 UP servers, K10 desktops, first K10.5 Shanghai desktops)
Socket **S1g1** (638 pins, K8 NPT, Brisbane mobiles)
Socket **S1g2** (638 pins, Fam. 11h mobiles)
Socket **S1g3** (638 pins, K10.5 (Shanghai)(Caspian die) mobiles)
Socket **ASB1** (812 balls BGA, Fam. 11h (Huron, Conesus dies) UPT mobiles)
Socket **ASB2** (812 balls BGA, K10.5 (Shanghai) (Geneva die), UPT mobiles)

Supports DDR

Socket **940** (130nm K8 UP/DP/MP servers, 90 nm K8 SC and 90 nm K8 DC (Egypt) DP/MP servers)
Socket **939** (90nm K8 SC and 90 nm K8 DC (Egypt) UP servers, K8 desktops with 128-bit mem. channel up to the 90 nm DC (Egypt) lines)
Socket **754** (K8 desktops, mobiles with 64-bit mem. channel up to the 90 nm SC (Athens) line)

2. AMD’s 32-bit x86 families

2. AMD’s 32-bit x86 families

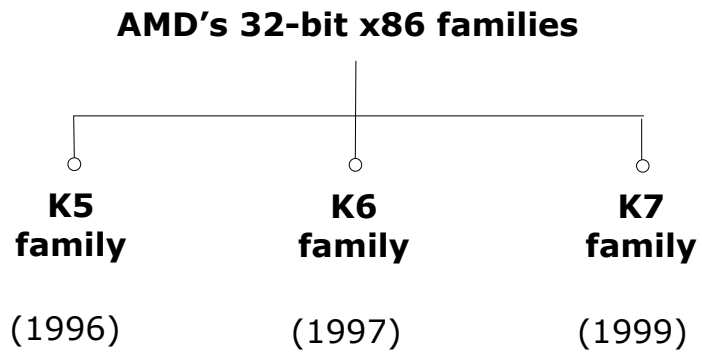

2. AMD's 32-bit x86 families (2)

Overview and major innovations in AMD's K5/K6 families

CPU Family	Intro.	CPU core	Brand name	Techn. (µm)	New key feature	Typ. Application
K5	1996		K5	0.5/0.35	2. gen. superscalar (32-bit), Pentium competitor	DT
K6	1997		K6	0.35/0.25	2.5 gen. superscalar, MMX (NexGen design)	DT/M
K6-2	1998	Chomper	K6	0.25	3. gen. superscalar, 3DNow!	DT/M
K6III **K6-2+** **K6III+**	1999 2000 2000	Sharptooth na. na.	K6	0.25 0.18 0.18	On-die L2 PowerNow![1] PowerNow![1]	DT/M

[1] PowerNow! was introduced in the Mobile K6-2+ and Mobile K-III+ processors in 2000
[2] Athlon: Attained performance lead over Intel's Pentium III

2. AMD's 32-bit x86 families (3)

Microarchitecture of AMD's 2. generation superscalar K5 [9]

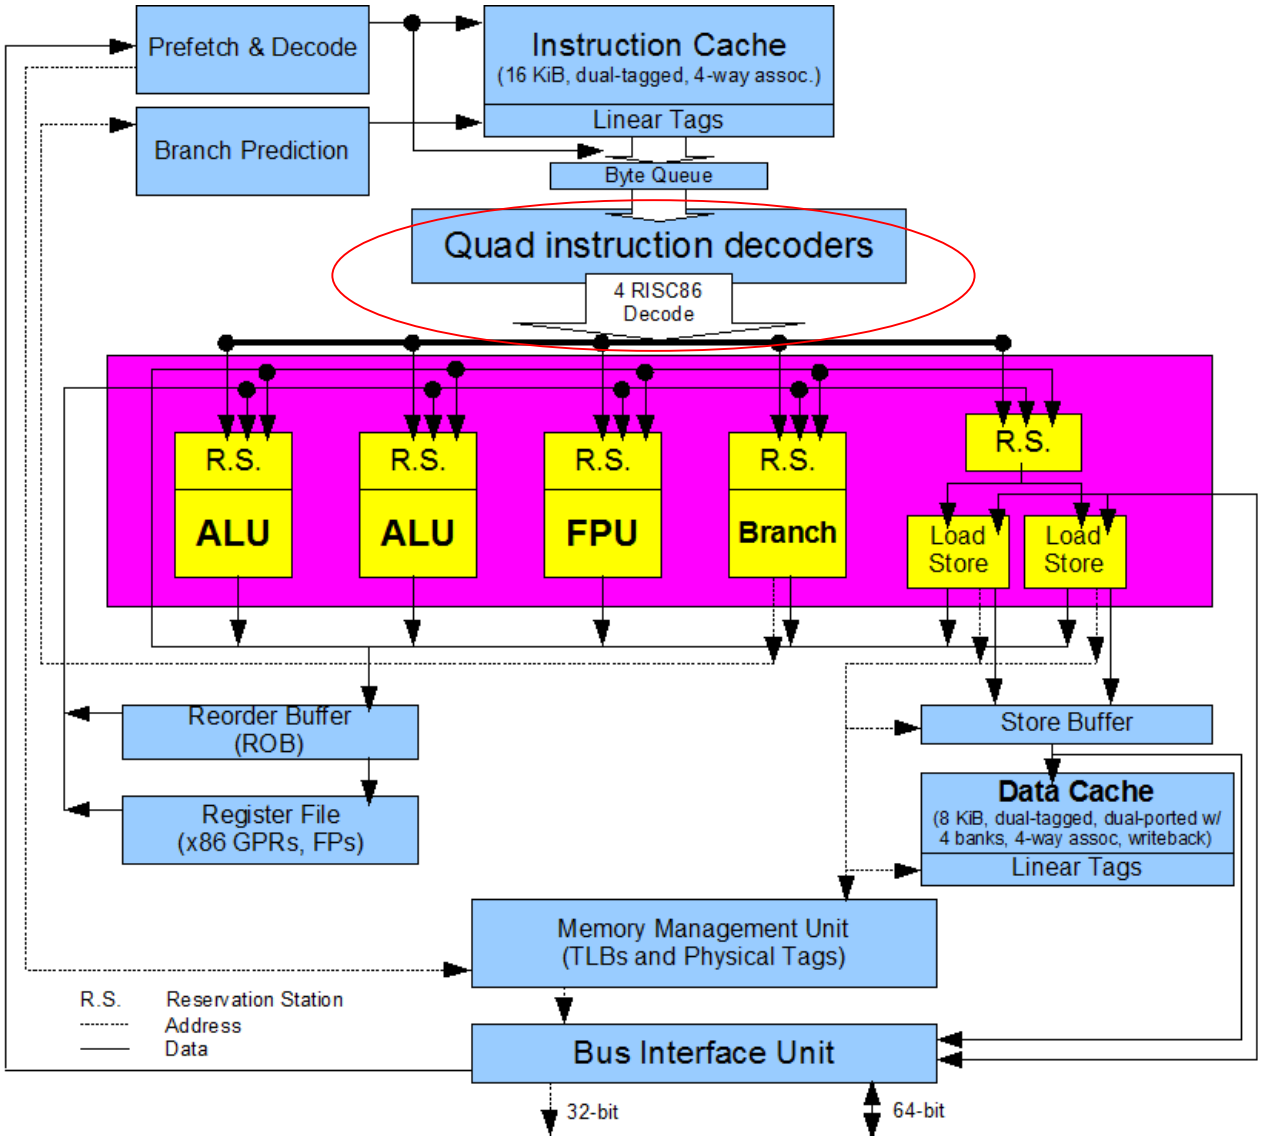

2. AMD's 32-bit x86 families (4)

Microarchitecture of AMD's 2.5 generation superscalar K6 (NexGen design) [10]

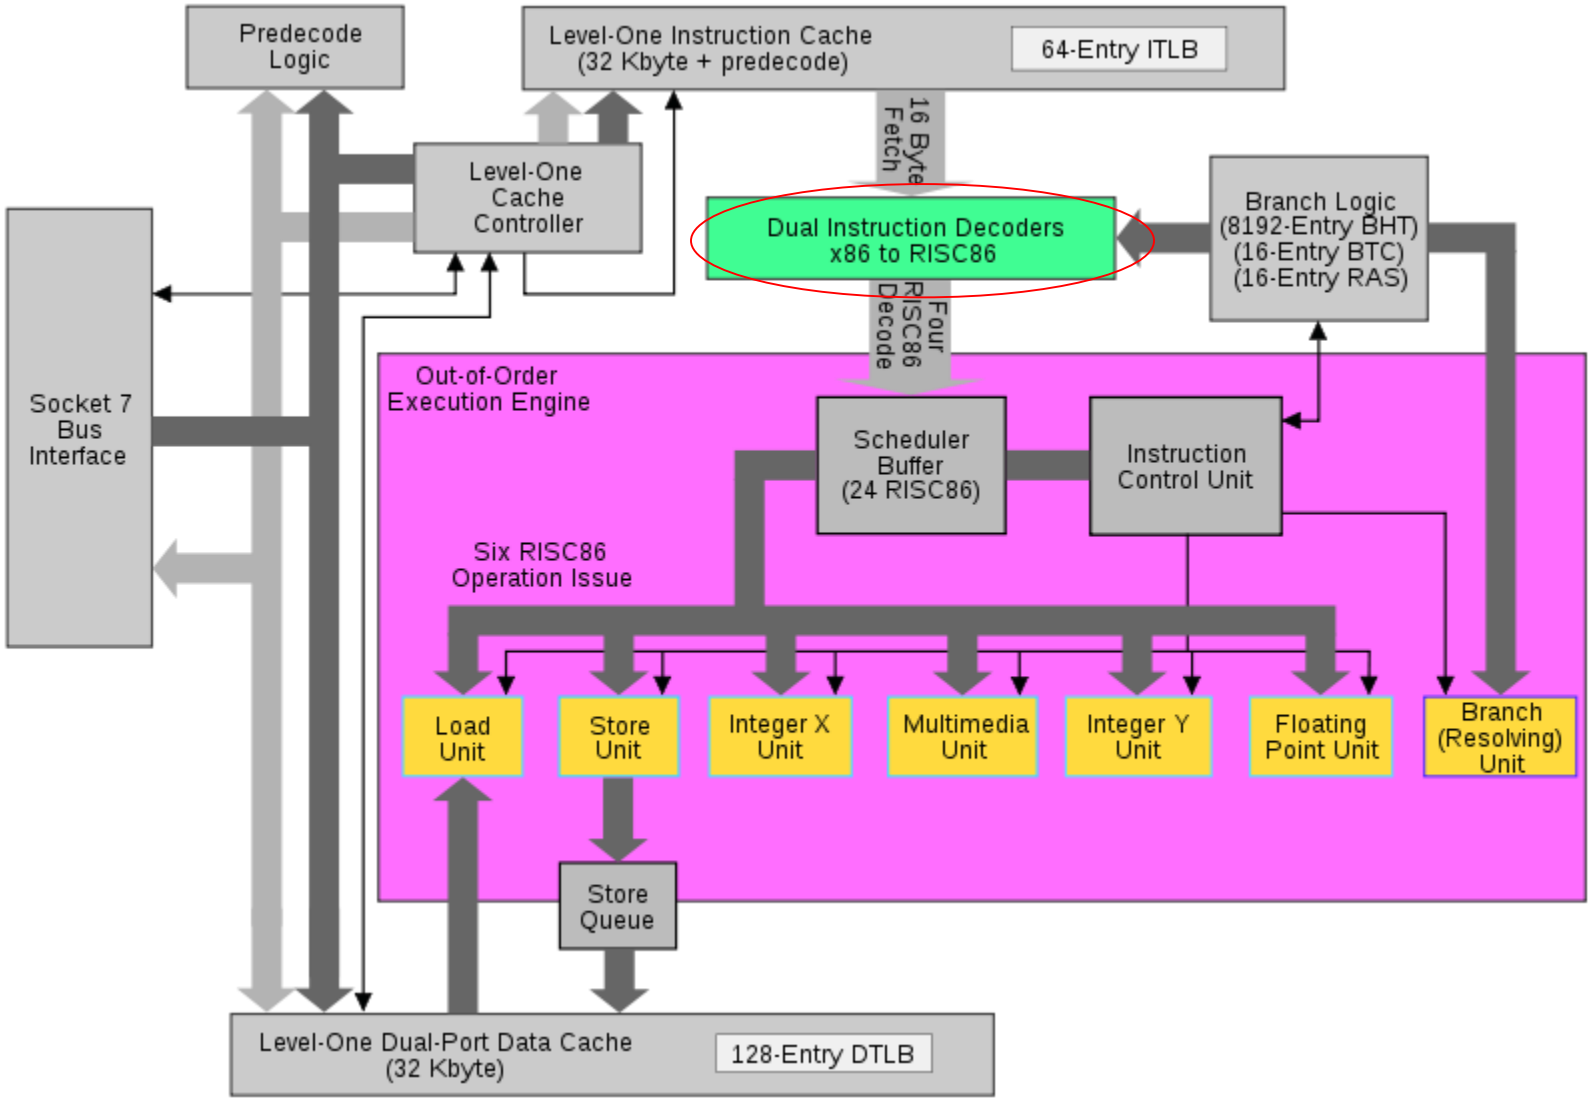

2. AMD's 32-bit x86 families (5)

Microarchitecture of AMD's 3. generation superscalar K7 (Athlon) [12]

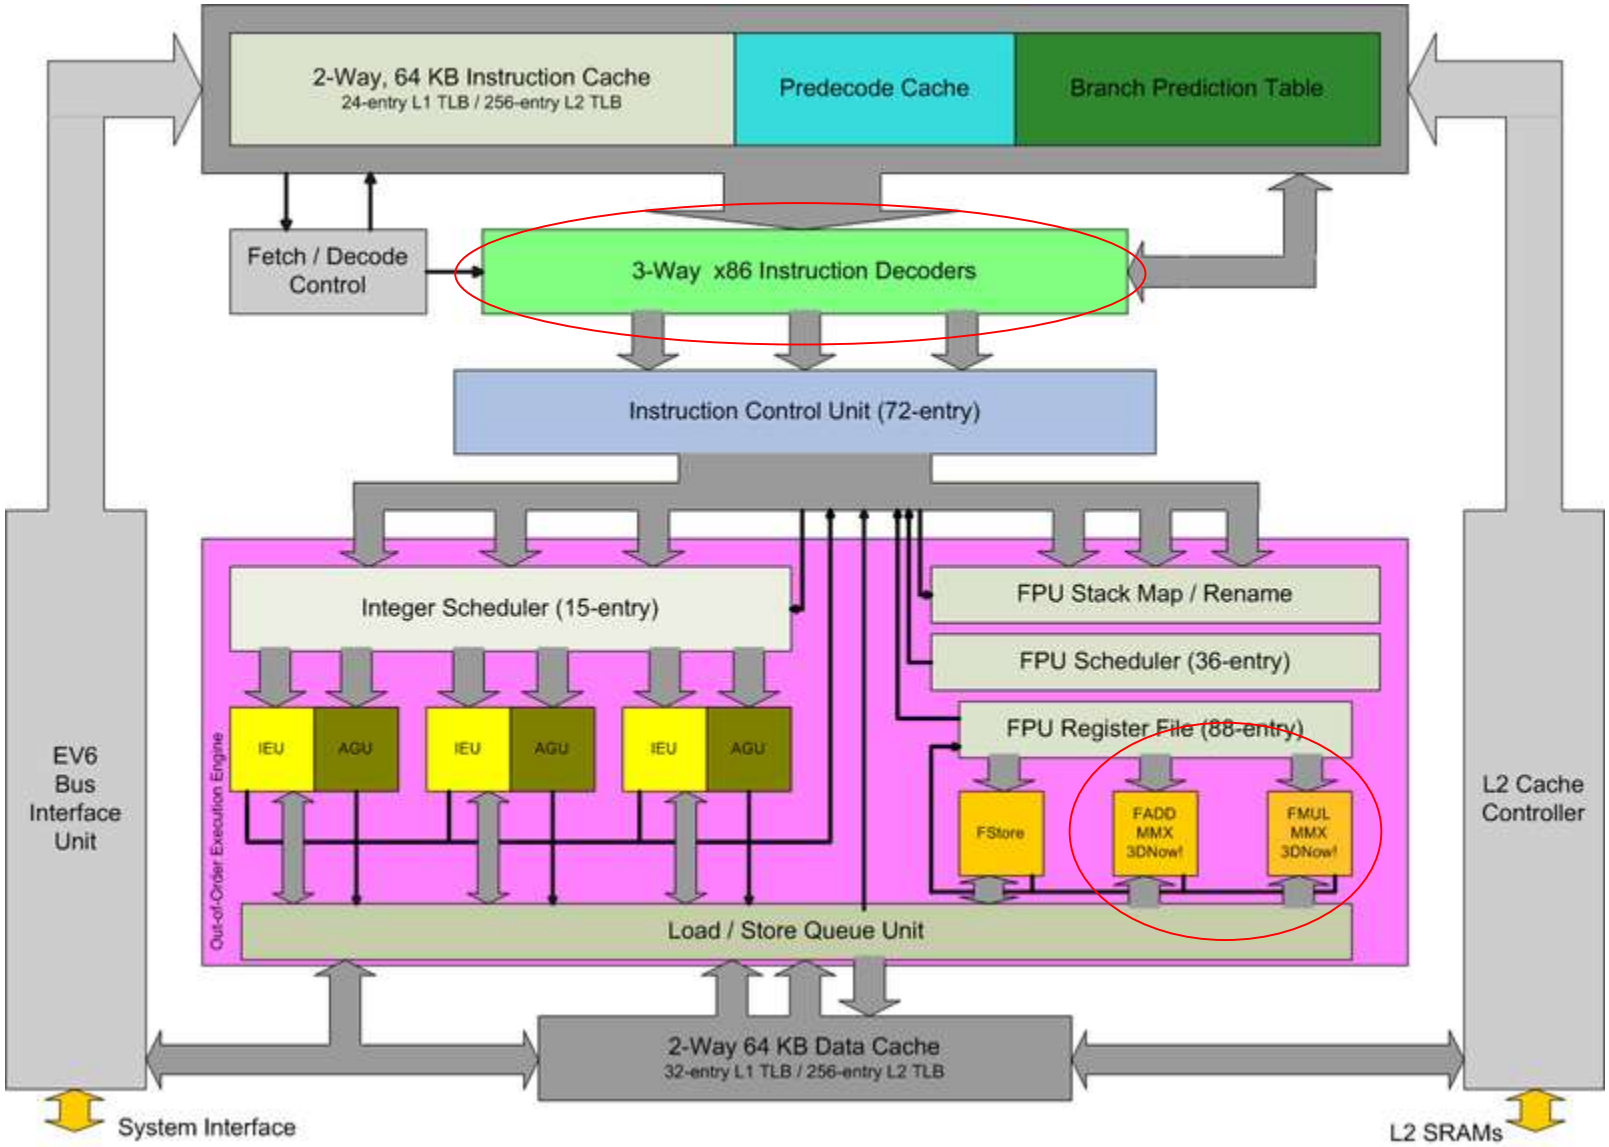

2. AMD's 32-bit x86 families (6)

Overview and major innovations in AMD's K7 (Athlon) families

Base arch./ stepping		Intro.	Core	Brand name	Techn. (µm)	L2 cache		FSB	ISA	PowerNow!	Typ. Appl
K7	Mod.1	6/1999	Argon	Athlon	0.25	In-package integrated 512 KB		DDR FSB	Enh. 3DNow!	-	DT
	Mod.2	11/2000	Pluto/Orion		0.18						
	Mod.3	6/2000	Spitfire	Duron		64 KB	On-die (exclusive)				DT
		1/2001		Mobile Duron							M
	Mod.4	6/2000	Thunderbird	Athlon		256 KB					DT
	Mod.6	7/2001	Palomino	Mobile Athlon4					3DNow! Prof./SSE	PowerNow!	M
		10/2001		Athlon XP						-	DT
		4/2001		Athlon MP							S
	Mod.7	8/2001	Morgan	Duron		64 KB					DT
		1/2001	Camaro	Mobile Duron						PowerNow!	M
	Mod.8	11/2002	Thoroughbred	Duron	0.13	256 KB				-	DT
		3/2003		Athlon XP							S
		8/2003	Applebread	Athlon MP							M
		4/2002	Thoroughbred	Mobile Athlon XP-M							
	Mod.10	9/2003	Thorton	Athlon XP							DT
		9/2003	Barton			512 KB					
		5/2003		Athlon MP							S
		3/2003		Mobile Athlon XP-M						PowerNow!	M

2. AMD's 32-bit x86 families (7)

Remark

Naming schemes of Intel's and AMD's processors

Traditionally, Intel named their processors by a character string, like Pentium, extended with a number reflecting the clock speed, like an early Pentium 133, meaning a Pentium processor with 133 Hz clock rate etc.

For a long time Intel's processors were designed for raw speed, achieved in the first line by using long pipelines (up to about 30 stages in the 3. core of Pentium 4 termed Prescott.

By contrast, AMD followed a different design philosophy, preferring efficiency (IPC) vs clock speed.

As a consequence, in those times Intel had a more favorable market position than AMD as costumers looked for high clock speed while buying computers.

In 2001 AMD tried to amend this drawback by introducing the PR (Performance Rating) scheme into the naming of their processors.

Usually, AMD's PR figures were interpreted as providing comparable or better performance than Intel's processors with the clock speed given in the PR rating.

E.g. an Athlon XP 1800+ was interpreted as having the same or higher performance than an Intel Pentium 4 1800 processor, despite the fact that its clock speed was actually only 1.53 GHz [11].

AMD employed the PR naming scheme also in their K7 Palomino based server (Athlon MP) and mobile (Mobile Athlon4) lines about the end of 2001.

After Intel met the thermal wall with their Pentium 4 Prescott core in 2004 and clock speeds became leveled off, Intel abandoned their clock speed based naming scheme and introduced a new naming scheme along with their Core2 family in 2006.

Also AMD abandoned the PR rating scheme when they introduced their quad-core K10 (Barcelona) line in 2007.

2. AMD's 32-bit x86 families (8)

Example: AMD's PR rating figures and related clock frequencies of the Athlon XP line [11]

CPU	FSB Frequency	Multiplier	Actual Core Frequency
Athlon XP 1800+	133MHz	11.5x	1.53 GHz
Athlon XP 1700+	133MHz	11.0x	1.47 GHz
Athlon XP 1600+	133MHz	10.5x	1.40 GHz
Athlon XP 1500+	133MHz	10.0x	1.33 GHz

3. Migration of 32-bit ISAs and microarchitectures to 64-bit

3. Migration of 32-bit ISAs and microarchitectures to 64-bit

Motivations for increasing the word length of processors and their underlying ISAs

There are two key motivations for increasing the word length of processors and their underlying ISAs, as follows:

- The demand for using larger data spaces

  In the course of the evolution of computing, applications need larger and larger data spaces, so their addressing requires more and more address bits.

- The demand for increasing performance

  Architectures with a longer word length can process more data per clock cycle, e.g. as long as 32-bit architectures can perform e.g. two 16-bit operations or a single 32-bit operation per clock cycle, 64-bit architectures are able to perform four 16-bit operations, two 32-bit or a single 64-bit operation per cycle.

→ As computing evolves there is a continuous motivation to increase the word length of processors, as demonstrated in the next Figure.

3. Migration of 32-bit ISAs and microarchitectures to 64-bit (2)

The evolution of the word length of x86 processors and their underlying ISA until the middle of the 1990's

Year of intro.	Processor	Data word length	Addressing
1978	8086	16-bit	16-bit/20 bit (segmented)
1982	80286	16-bit	16-bit/20 bit (segmented)
1985	80386	32-bit	32-bit
1995	Pentium Pro	32-bit	36-bit

→ The word length of x86 processors and their underlying ISA became 32-bit as early as in 1985 but get stuck at this figure for nearly 20 years.

3. Migration of 32-bit ISAs and microarchitectures to 64-bit (3)

Migration of 32-bit RISC ISAs to 64-bit

In contrast to CISCs, RISC ISAs and processors made the leap to 64-bit already in the first half of the 90's

- either while firms enhanced legacy 32-bit RISC ISAs and processors to 64-bit, as shown below

Year of extension	RISC ISA
1991	MIPS III
1993	SPARC V9
1995	PowerPC-AS
1996	PA-RISC 2.0

- or when firms developed new superscalar 64-bit RISC processors and underlying ISAs, from the scratch, like

  DEC (later COMPAC then HP) the Alpha ISA and the related Alpha line of processors.

3. Migration of 32-bit ISAs and microarchitectures to 64-bit (4)

Remark

Wider processors need obviously wider system buses to avoid bottlenecks in data transfers.
The related evolution of the system buses of early x86 processors is given in the next Table.

Main features of the system-bus of x86 processors

Width of the	8086	8088	80286	80386	80486	Pentium	Pentium Pro PII, PIII	P4
address bus (bit)	20[1]	20[1]	24	32[2]	32[2]	32[3]	36	36
data bus (bit)	16[1]	8[1]	16	32	32	64	64+8[4]	64+8[4]

[1] Multiplexed
[2] Bits 0,1 not implemented (Doubleword aligned)
[3] Bits 0-2 not implemented (Quadword aligned)
[4] For error protection

Table 3.1: Main features of the system bus

Referring to the above Table we point out that it was the 32-bit Pentium processor whose data bus became widened to 64-bit in order to increase its memory bandwidth (as the memory transfer is carried over the system bus).

The resulting 64-bit data bus width evoked the emergence of 64-bit memory DIMMs.

3. Migration of 32-bit ISAs and microarchitectures to 64-bit (5)

The demand for upgrading 32-bit x86 ISAs and microarchitectures to 64-bit

At the latest when most RISC processors already migrated to 64-bit in the middle of the 1990's it became obvious that there is an urgent need to upgrade the 32-bit x86 ISAs and processors to 64-bit as well.

Expected key benefits of upgrading to 64-bit are [18]

- The capability to directly address more than 4 GB, since 32-bit addressing limits the direct addressability of the physical memory to 4 GB.

  High performance servers or applications with large databases will benefit from available larger data spaces.
- The capability to perform twice as much data operations at the same time.
- The possibility to extend the number of available programmable registers, as for computer intensive applications the small number of registers available in legacy x86 architectures limits performance.

3. Migration of 32-bit ISAs and microarchitectures to 64-bit (6)

Intel's and HP's approach to introduce a 64-bit ISA

- 6/1994 Intel and HP announced that they formed an alliance to develop a new 64-bit ISA which would become the basis for a line of Intel microprocessors [19].
- 10/1997 Intel and HP disclosed main features of the 64-bit IA-64 ISA, revealing that it is based on the EPIC execution model (renewed and strongly enhanced VLIW model).

  Intel and HP also announced that the first IA-64 processor, termed as Merced, is slated for 1999 [19].

  The new IA-64 ISA and the related processor family aimed at workstations and servers while providing appropriate means to run software written for Intel's existing 32 bit IA-32 and HP's PA-RISC processors (whose word length was already extended to 64-bit along with the PA-RISC 2.0 ISA in 1996 [20].

Chosen techniques to run the existing PA-RISC and x86 (32-bit) code bases on IA-64 processors:

- PA-RISC code is automatically converted to the native IA-64 code by a dynamic object code translator [21].
- The compatibility with the existing x86 code base was implemented by an additional x32 decoding unit and shared resources [22], as shown below.

Block diagram of the first Intel Itanium processor [23]

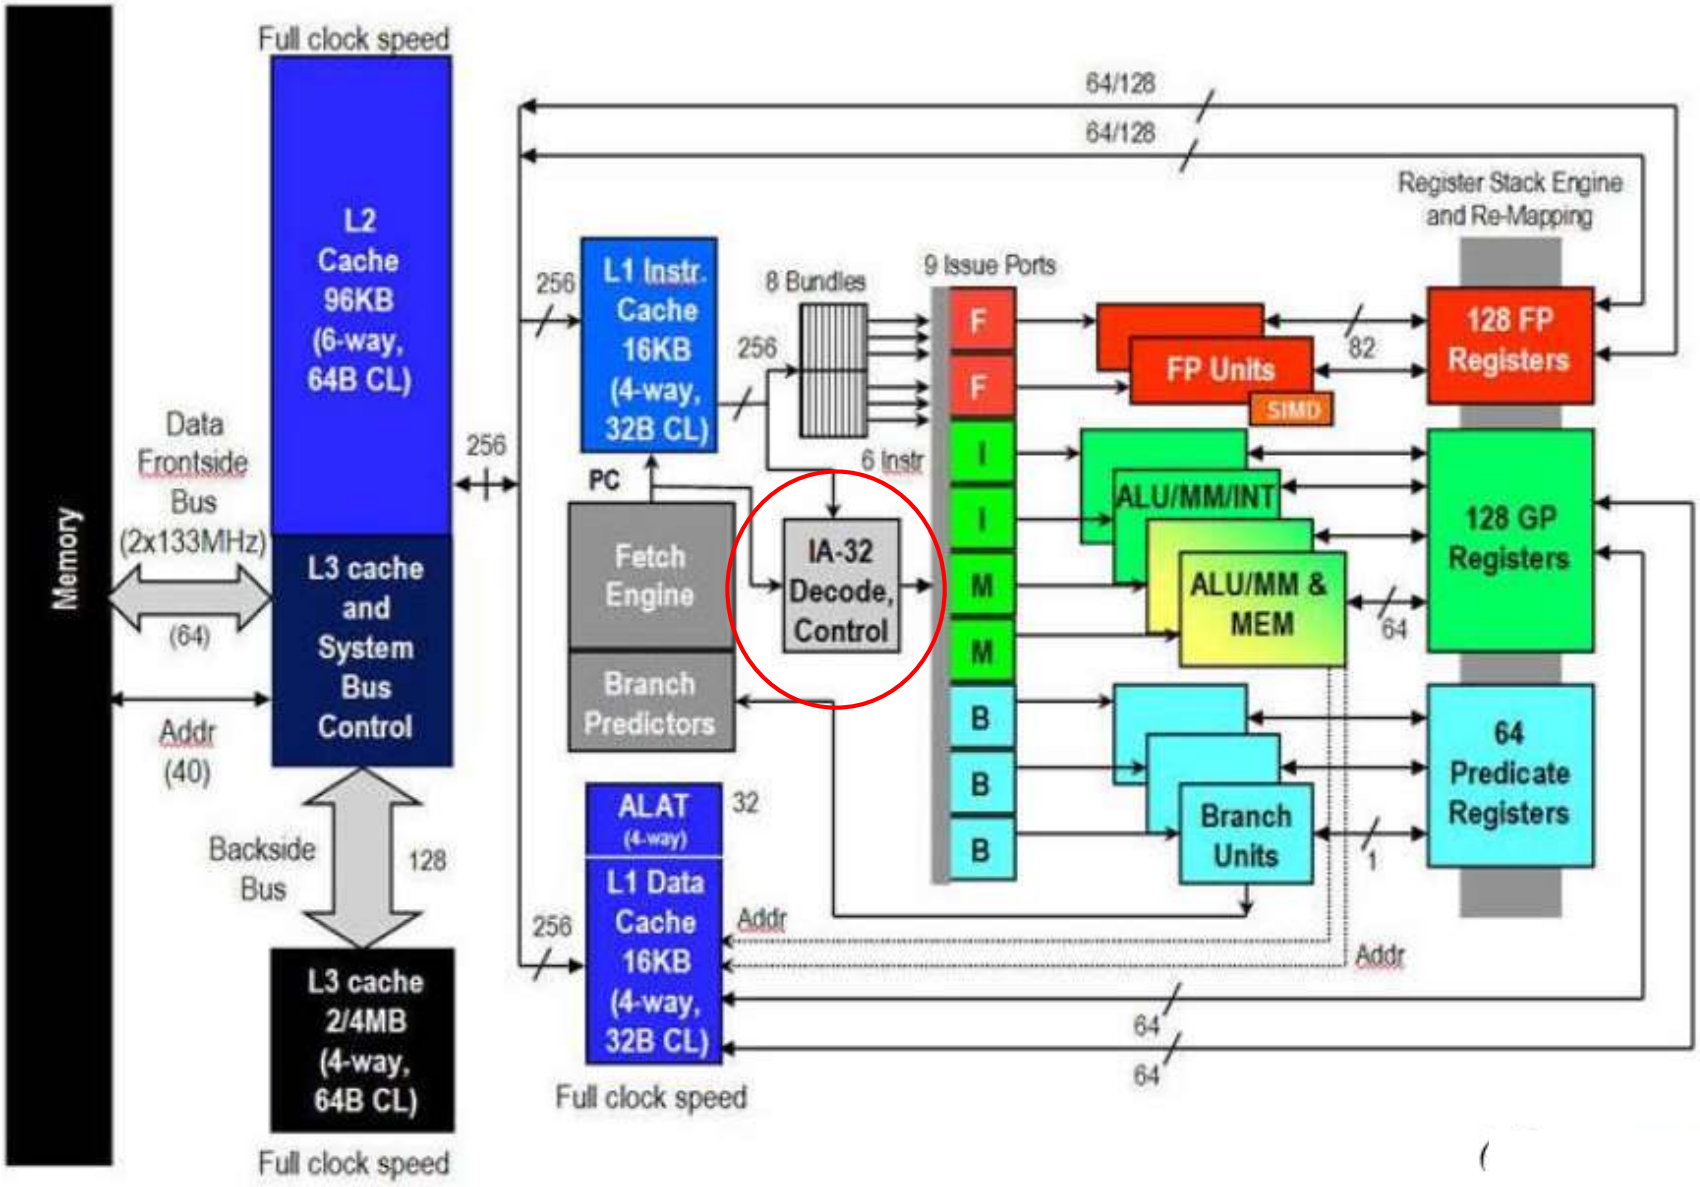

3. Migration of 32-bit ISAs and microarchitectures to 64-bit (8)

Intel's milestones in introducing the IA64 ISA and the related Itanium line-1

- 5/1999 Intel and HP revealed the IA-64 ISA [24].
- In 1999 Intel named their new IA64 processor line as the Itanium line and its first processor Merced.
- Merced was scheduled to appear in 1999, but its launch was delayed about two years until 6/2001.

3. Migration of 32-bit ISAs and microarchitectures to 64-bit (9)

Market reactions to Intel's IA-64 based processor implementation

- At its delayed introduction in 5/2001 Merced showed considerable lower performance than expected.
- Due to this and due to the fact that the IA64 ISA means a radical departure from the x86 software environment the Itanium line gained a much slower and lower market penetration than expected [29], as shown below.

Remark

Intel announced the official name of the processor, *Itanium*, on October 4, 1999. [29]

3. Migration of 32-bit ISAs and microarchitectures to 64-bit (10)

Itanium's sales forecasts and actual sales figures [29]

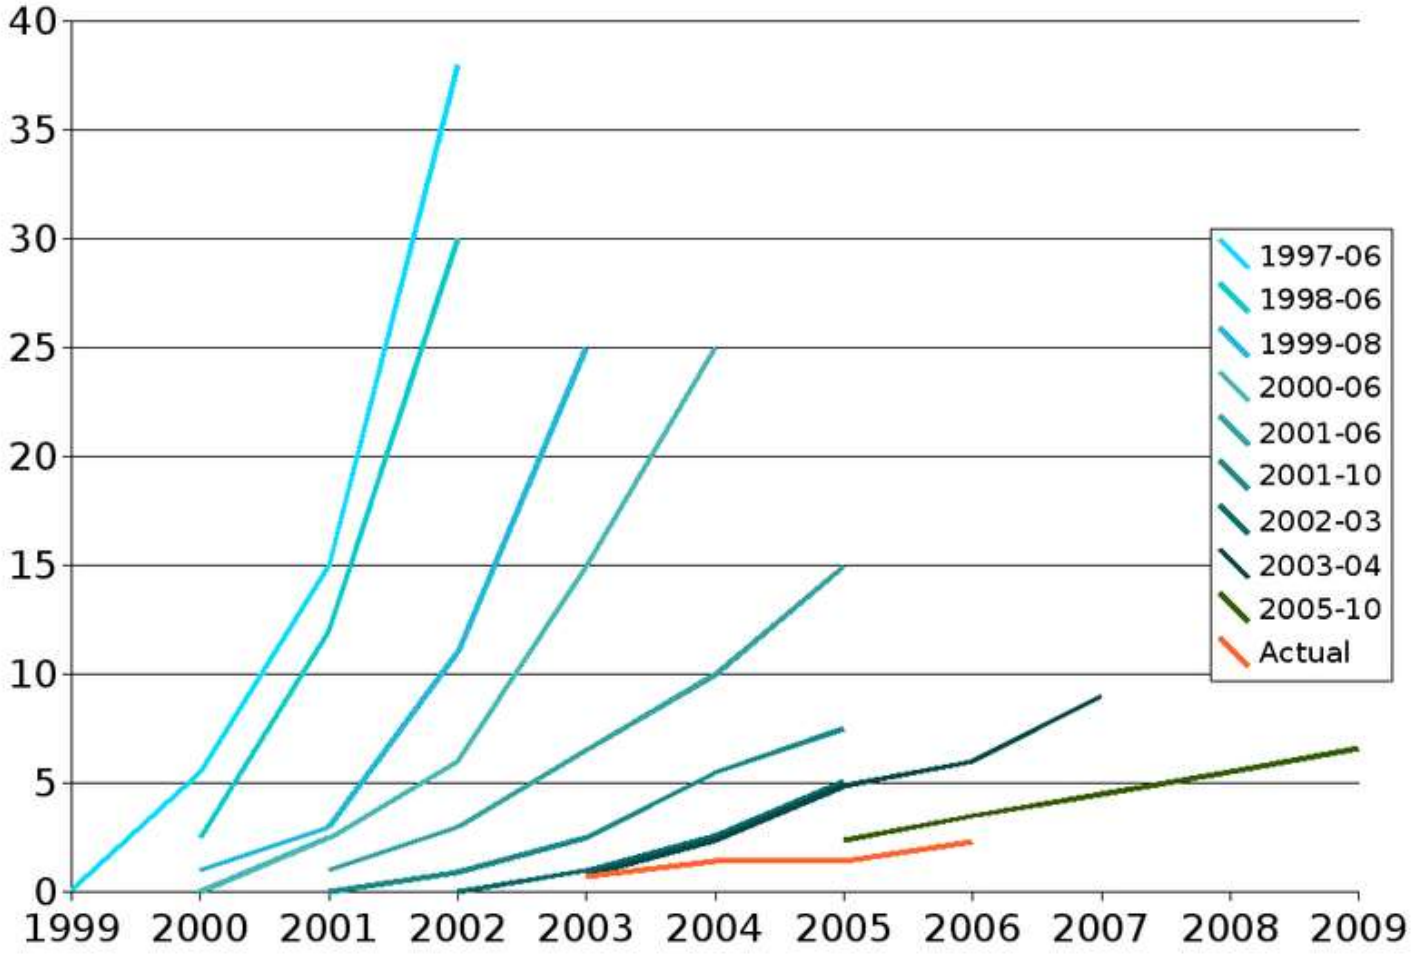

AMD's milestones in introducing the x86-64 ISA and related processor lines-1

- 10/1999 AMD disclosed their plan
  - to make a compatible extension of the x86 ISA, designated as the x86-64 ISA
  - to implement it as their eights generation (K8) processor family, code named as Sledgehammer and
  - to use the serial two byte wide Lighting Data Transport bus (renamed in 2001 to HyperTransport bus) as a chip-to-chip interconnect bus to provide enough I/O bandwidth.
- 8/2000 Release of the x86-64 architecture specification to encourage the software community to begin incorporating x86-64™ technology into operating systems, applications, drivers and development tools [25] .

AMD's competing approach to introduce 64-bit computing

AMD's goal was clearly formulated at revealing the x86-64 ISA specification in 8/2000 [25]:

- "Ultimately this technology is designed to help preserve the enterprise community's enormous financial investment in 32-bit operating systems, applications, development tools and support infrastructure while providing a seamless path to deploy future 64-bit technology."
- "Perhaps the most noteworthy feature of AMD's approach to 64-bit computing is that it is an extension to the 32-bit environment prevailing in the industry today rather than a radical departure."

The programmer's model of the available register set in the x86-64 ISA [28]

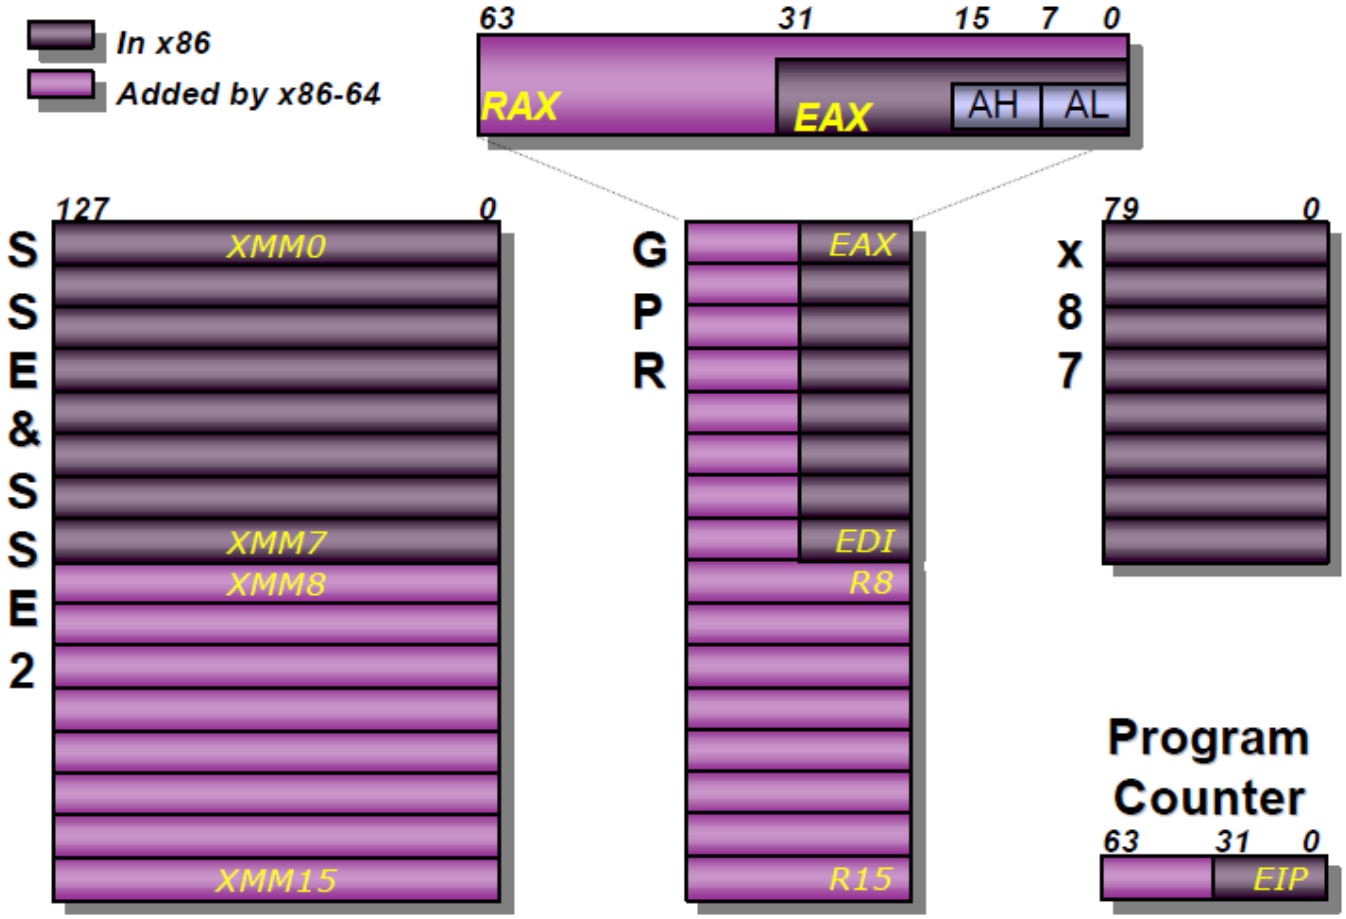

Intel's reaction to AMD's disclosure to develop a smooth migration from x86 to x86-64 computing-1

After AMD disclosed their plan to migrate to 64-bit computing on a smooth evolutionary path unlike Intel's revolutionary approach, Intel recognized that they have to react to it by a plan B for the case when the Itanium project should fail [31].

Intel had two options:

- either to define from the scratch on a new x86-64 ISA, preannounce it and convince Microsoft and the rest of the software industry to develop a related new operating system and a full set of software tools, like compilers, debuggers and the like,
- or to accept AMD's x86-ISA and utilize subsequently the whole software environment to be developed for it.

Intel chose the second option as plan B, as described in more details in [31] and is cited below.

Intel's reaction to AMD's disclosure to develop a smooth migration from x86 to x86-64 computing-2

"One option was to preannounce a competing ISA with a RISC-like 64-bit extension to x86. This would have been risky: Microsoft and other vendors were unlikely to develop softeware for another, non-compatible x86 extension without a major performance win. Furthermore, Intel did not want to damage the IA64 project, and disclosure of an alternative 64-bit plan so far in advance of the Itanium release would hurt the project and the HP relationship. Nor could Intel work on this alternative secretly: if Intel's x86 extensions were not compatible with AMD64, Intel would have to disclose the plan to vendors to enable them to develop compilers and operating systems for the platform.

Eventually, Bhandarkar was responsible for proposing the effort that began in June of 2000 to release an AMD64-compatible Intel ISA that was variously called Yamhill, Clackamas and finally EM64T during its secretive development cycle (it was eventually renamed Intel64). Intel knew that vendors would be able to release Intel64-compatible software quickly on the heels of their AMD64 development efforts. In fact, Intel monitored the Windows source for AMD64-releated changes, ran their own builds, and tested those builds on pre-silicon simulators for validation before sharing the plan with Microsoft so as to keep the possibility of leaks to a minimum. In January of 2002 they began to disclose their plans to partners, followed by testing on prototype systems in 2003, and production systems in early 2004."

Intel's reaction to AMD's disclosure to develop a smooth migration from x86 to x86-64 computing-3

Microsoft's role in Intel's decision not to develop a competing x86-64 ISA

As reported by various sources about 4/2002, top Microsoft decision makers viewed the x86-64 ISA as the clearly superior solution over IA64 [132].

Accordingly, Microsoft has pressured Intel into supporting AMD's x86-64 ISA else allegedly they may drop supporting Intel's IA64 ISA [18].

AMD's milestones in introducing the x86-64 ISA and related processor lines-2

- 10/2001 Announcing the planned "Hammer" architecture, renamed later (in 2004) to Direct Connect Architecture again at the Microprocessor Forum in form of a detailed presentation [26].
- 4/2002 AMD announces Microsoft support for their x86-64 lines [27].
- 4/2002 Disclosure of the Opteron designation for the x86-64 server line.
- 4/2003 First shipment of K8 processors.

  Both the introduced server and desktop processors were superior to Intel's existing Pentium 4 based processors [123].

Intel's way to develop their x86-64 processors-1

- As the B plan, Intel started silently to develop an AMD compatible 64-bit enhancement to be included into their third Pentium 4 core, called the Prescott core about 6/2000 [31].

  This secretive development effort was named differently first as Yamhill, later Clackamas and finally EM64T [31].

  For years however, Intel officially denied that they do develop an x86-64 extension (see e.g. [32], [34]), for not to undermine their own 64-bit IA-64 (Itanium) project.

  The existence of a 64-bit extension was however, obvious for insiders as Prescott had a transistor count of 125 million, more than twice as much as their previous Northwood core that incorporated only 55 million transistors [33].

  On the other hand, Intel informed already in 1/2002 their key partners about supporting AMD's x86-64 ISA in the upcoming 3. core of the Pentium 4.

- In 2/2004 then, Intel introduced the third core of their Pentium 4 line (termed as the Prescott core) with the x86-64 extension included but not disclosed.

3. Migration of 32-bit ISAs and microarchitectures to 64-bit (19)

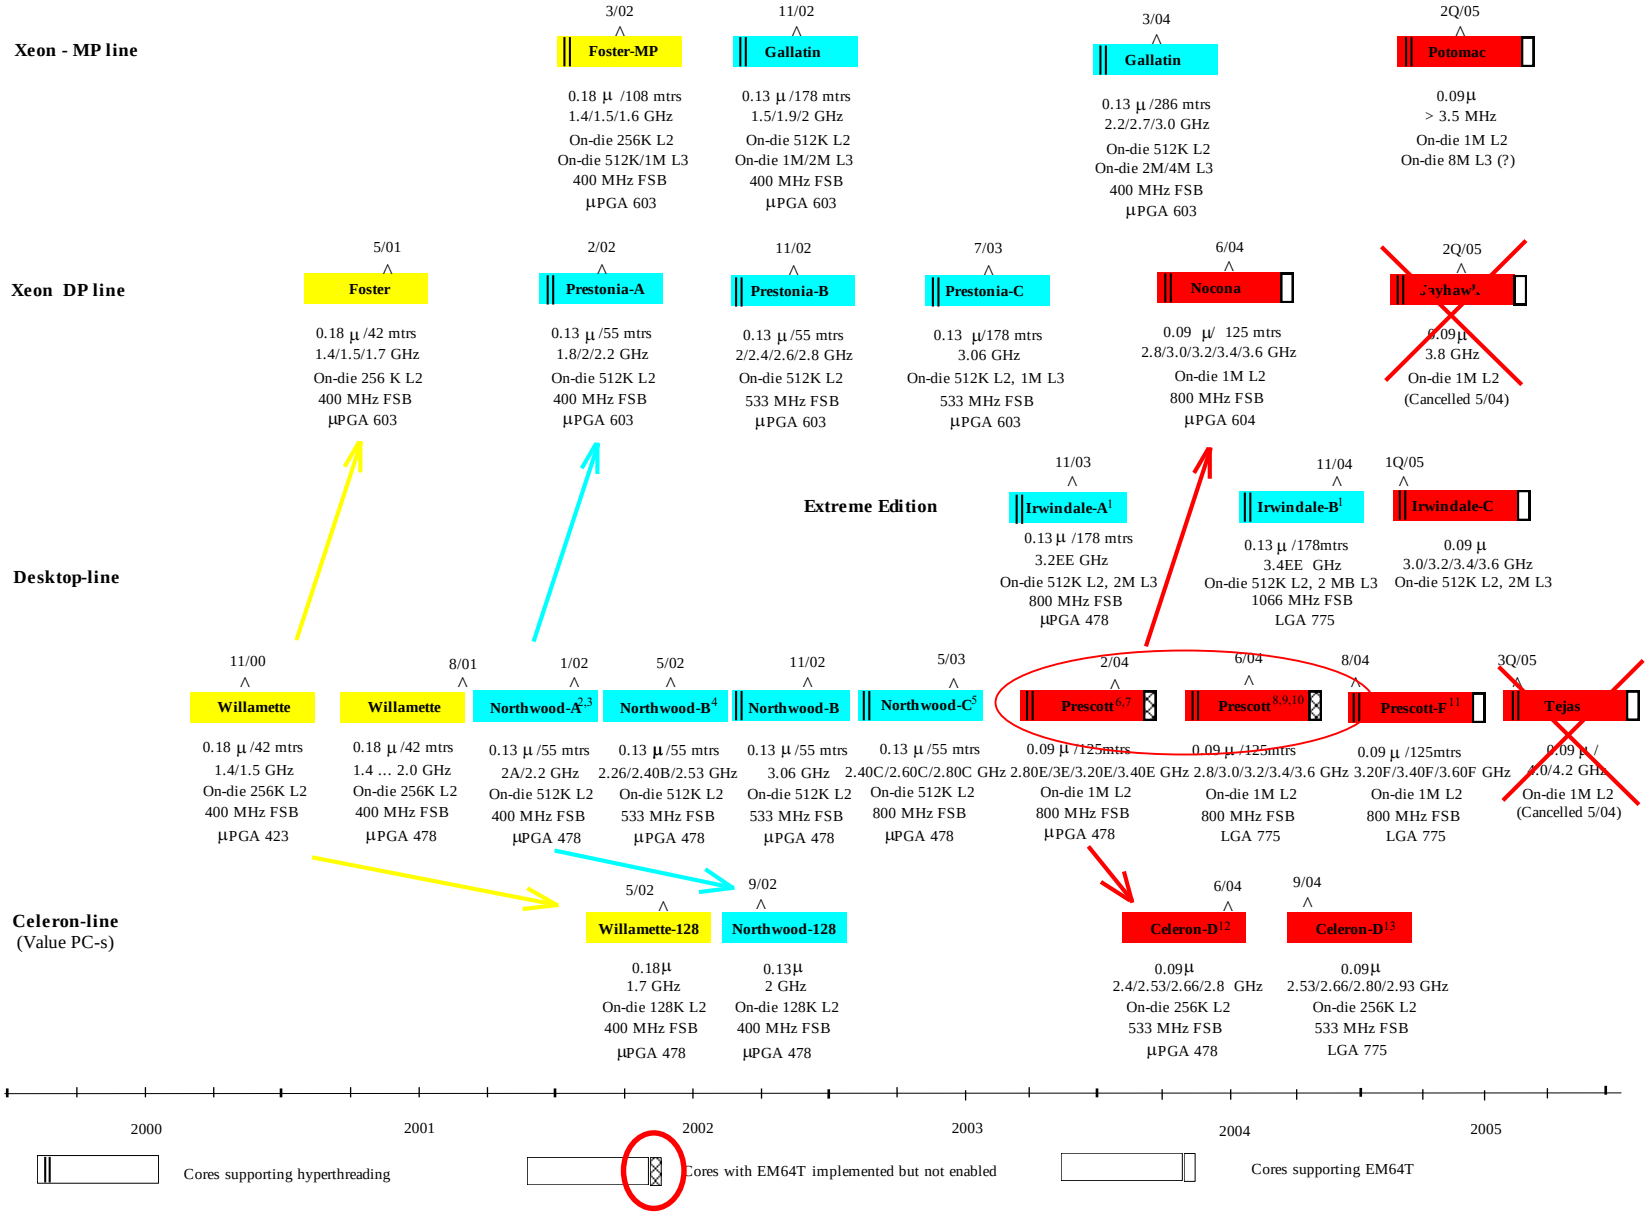

Figure: Intel' P4 processor family (Implementing the Netburst microarchitecture)

Intel's way to develop their x86-64 processors-2

Finally, in summer 2004 Intel revealed without much PR their x86-64 extension first in their server lines and subsequently also in their desktops, designated as the EM64T (Extended Memory 64 technology).

3. Migration of 32-bit ISAs and microarchitectures to 64-bit (21)

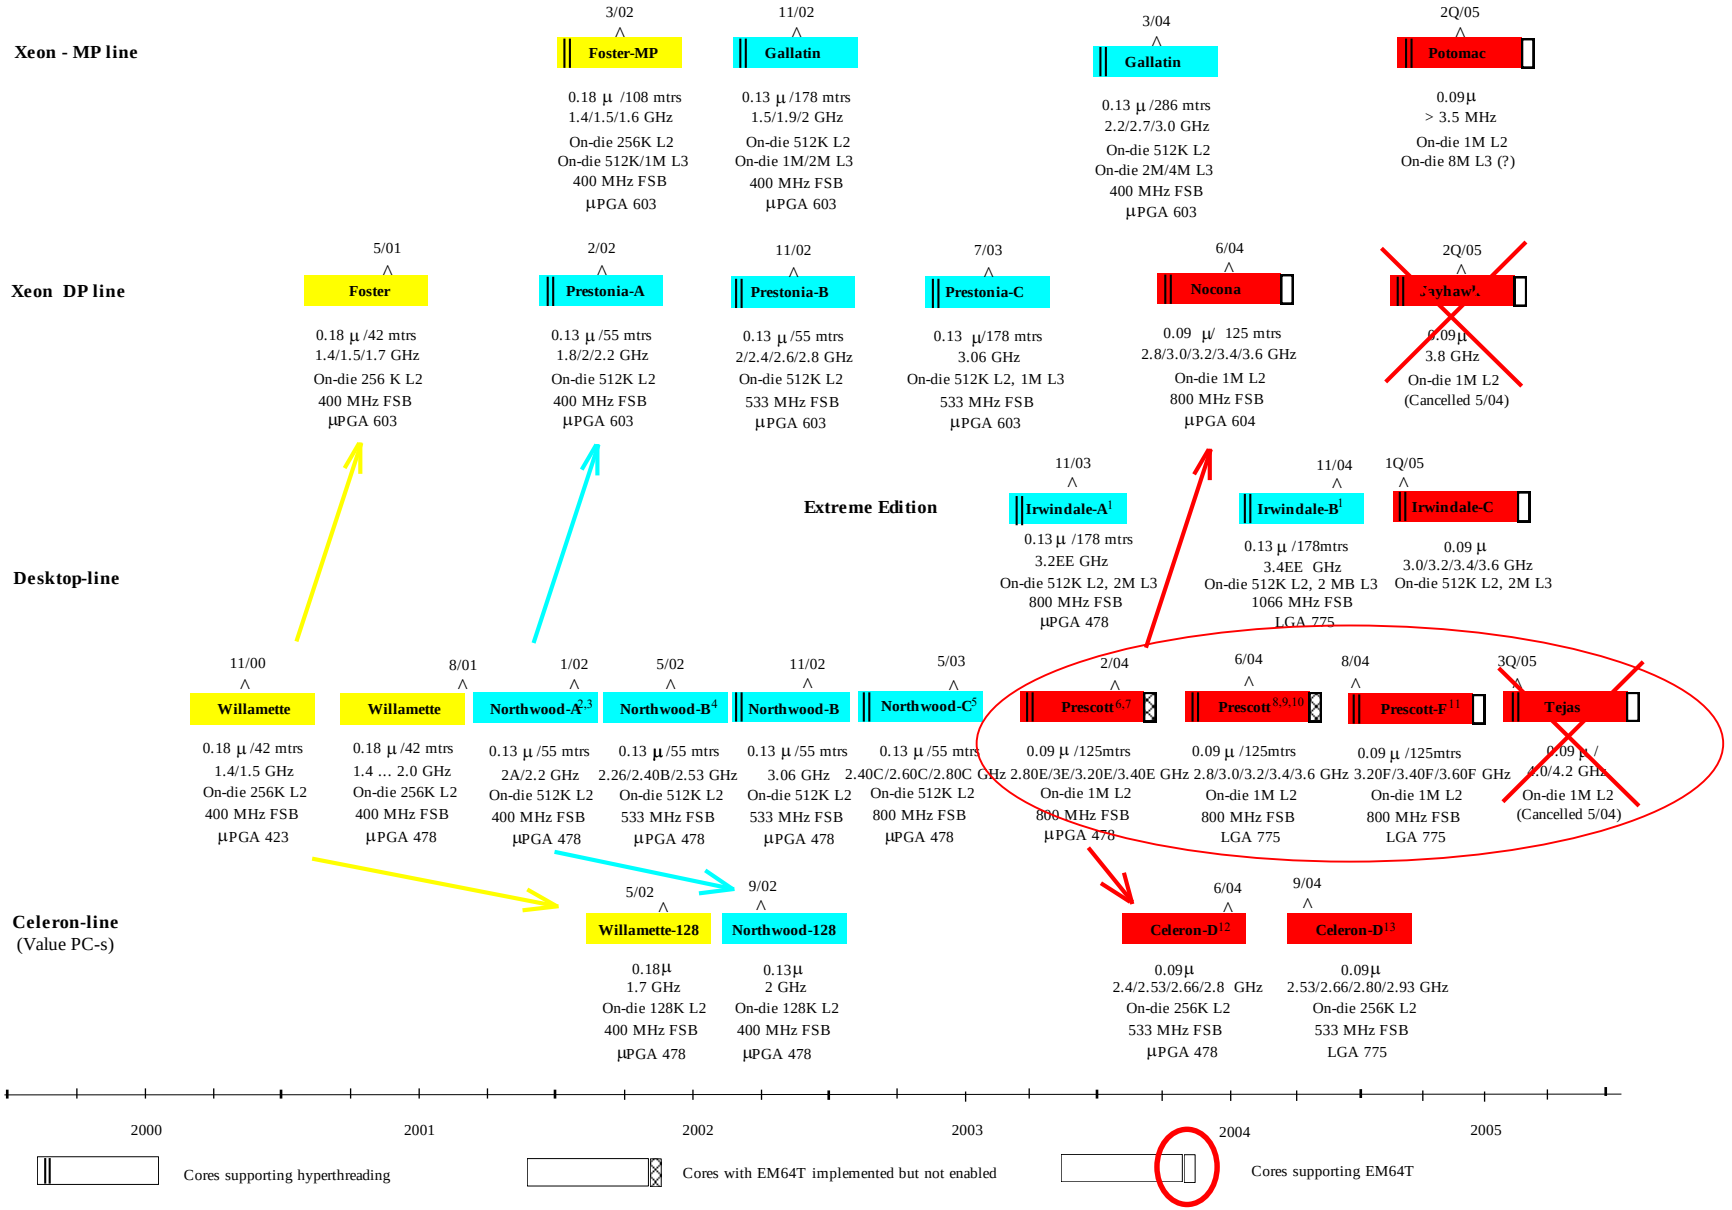

Figure: Intel' P4 processor family (Implementing the Netburst microarchitecture)

Intel's transition from x86 (IA-32) to x86-64 (called first as EM64T) [35]

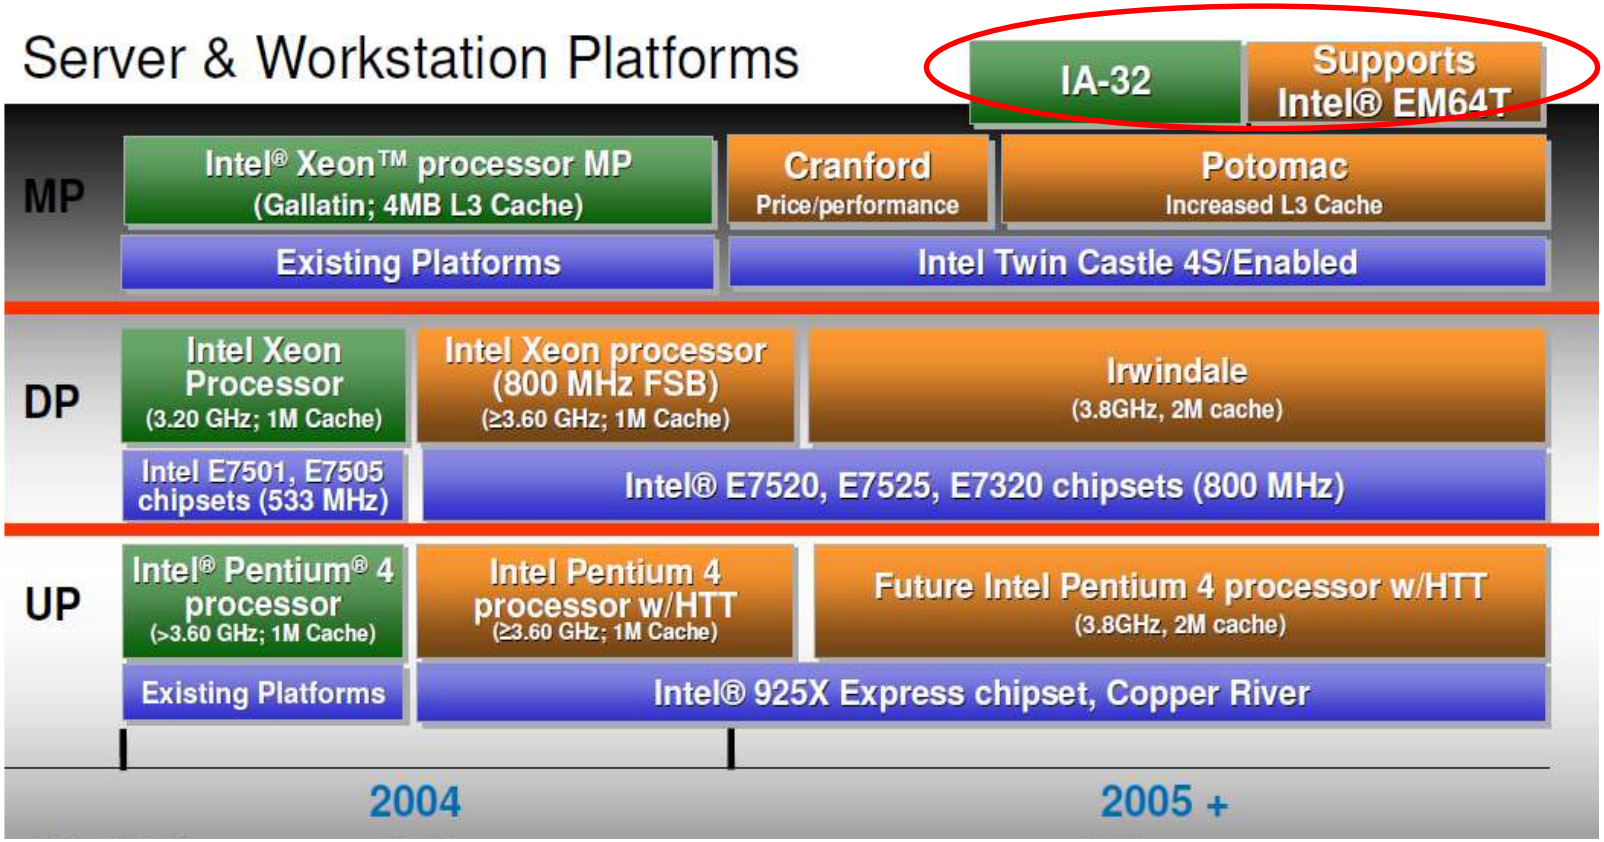

Note

a) EM64T: Extended Memory 64 Technology, a designation that was used only for a few years.
b) The x86-64 transition requires new chipsets, as shown.

3. Migration of 32-bit ISAs and microarchitectures to 64-bit (23)

Remark

A fascinating, very detailed description of the migration from x86 to x86-64 can be found in [31]

The fate of the IA-64 architecture

Due to the rapid evolution of the mainline multicore x86-64 processors the target market segment of the IA-64 line, that is large, high performance mission critical servers, became more and more smaller.

→ 4/2010 Microsoft and about one year later (3/2011) also Oracle announced that they will discontinue the support the IA-64 (Itanium) line.

The last supported versions are:

- the Windows Server 2008 R2,
- the SQL Server 2008 R2 database management tool and the
- Visual Studio 2010 developer tools [36].

4. Overview of AMD’s K8 – K10.5 (Hammer-based) families

4. Overview of AMD's K8 K10.5 (Hammer-based) families

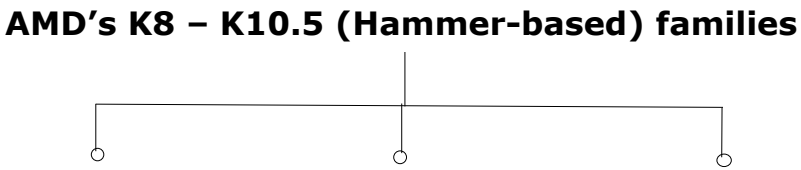

	K8 family	K10 family	K10.5 family
Typical family designations	*Hammer*	*Barcelona*	*Shanghai* *Istambul* *Magny-Course*
	Section 5	*Section 6*	*Sections 7-9*

4. Overview of AMD's K8 – K10.5 (Hammer-based) families (2)

Overview AMD's 64-bit K8-based lines [13]

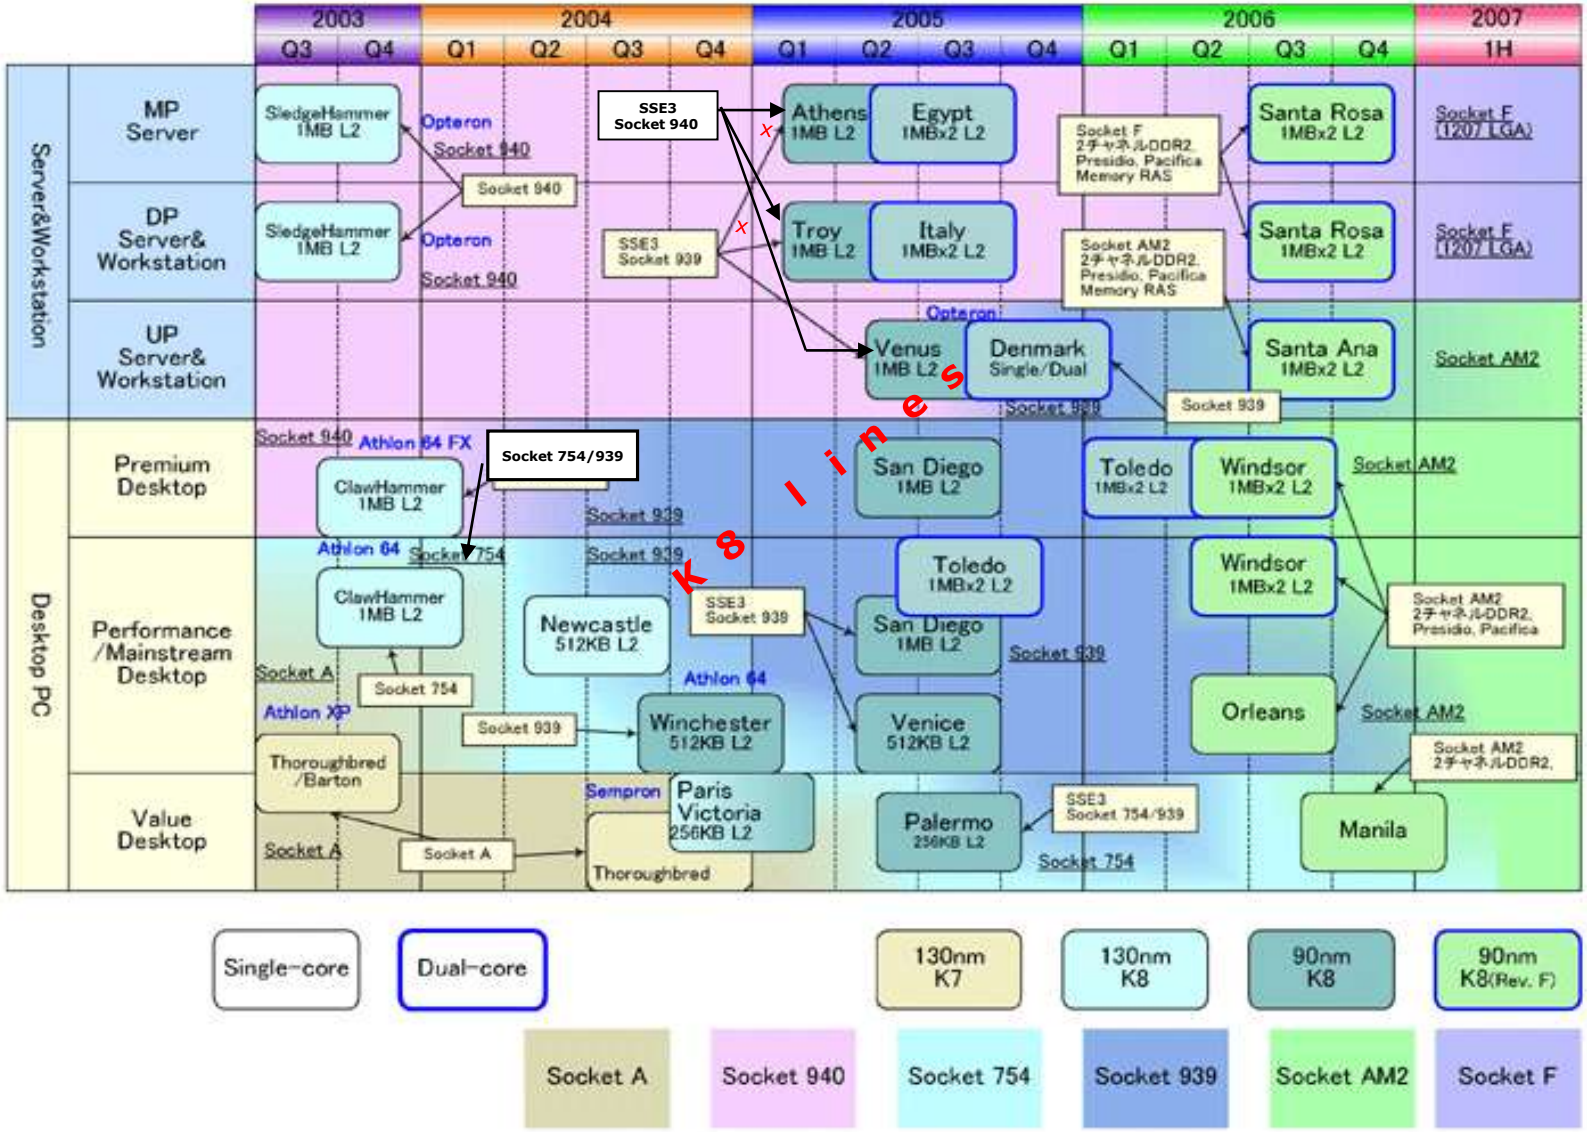

4. Overview of AMD's K8 – K10.5 (Hammer-based) families (3)

Overview of AMD's K10 and subsequent x86-64 families [14]

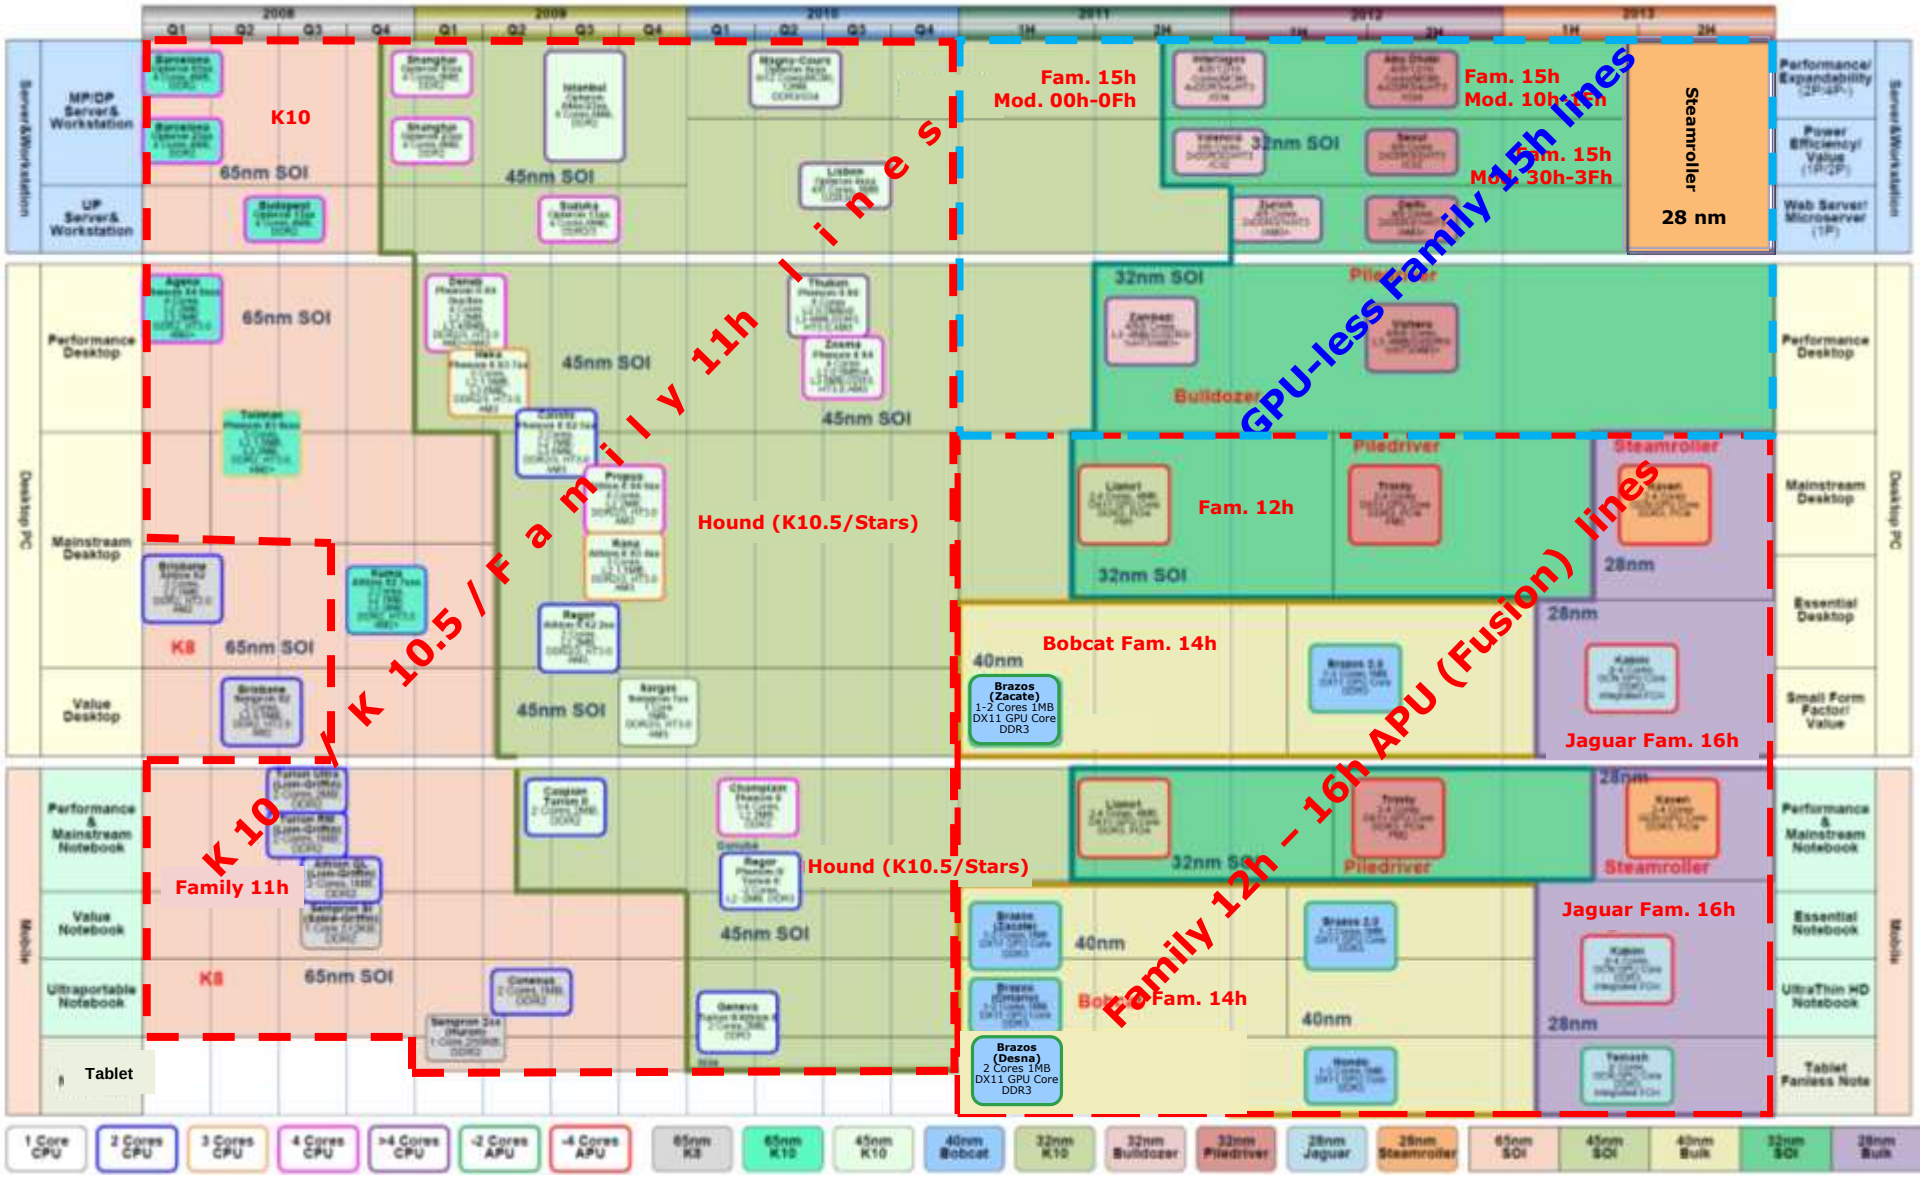

4. Overview of AMD's K8 – K10.5 (Hammer-based) families (4)

Brand names of AMD's K8 – Family 10.5h (Hammer)-based processor lines

		2003-2007	2007-2008	2008-2011	2009	2009
		K8 (Hammer)	**K10 (Barcelona)**	**K10.5 (Shanghai)**	**K10.5 (Istanbul)**	**K10.5 (Magny- Course)**
Servers	**4P servers**	**See Section 4**	**Barcelona (834x-836x))**	**Shanghai (837x-839x)**	**Istambul (8410-8430)**	**Magny-Course (6100)**
	2P servers		**Barcelona (234x-236x)**	**Shanghai (237x-239x)**	**Istambul (241x-243x)**	**Lisbon (4100)**
	1P servers		**Budapest (135x-136x)**	**Suzuka (138x-139x)**		
Desktops	**High perf.** (~80-120W)		Phenom X4-X2	Phenom II X4-X2	Phenom II X6-X4	
	Mainstream (~60-90W)	Athlon 64 Athlon 64 X2	Athlon X2	Athlon II X4-X2		
	Value (~40-60W)	Sempron		Sempron		
Mobiles	**High perf.** (~30-40W)	Turion 64 X2 (TL 6/5) Turion 64 (ML/MT)		Phenom II (N/P 9xx-6xx) Turion II Ultra (M6xx) Turion II (M/N/P 5xx)		
	Mainstream (~20-30W)	Athlon 64 X2 (TK-5x/4x) Athlon 64 (2xxx+-4xxx+)		Athlon II (M/N/P 3xx) Sempron (M1xx)		
	Ultraportable (~10-20W)	Mobile Sempron (2xxx+-4xxx+) Sempron 2100 fanless		Turion II Neo (K6xx) Athlon II Neo (K1xx) V-series (V1xx)		
Embedded (~10-20W)				Turion II Neo X2 Athlon II Neo X2 Athlon II Neo		

4. Overview of AMD's K8 – K10.5 (Hammer-based) families (5)

Main features of AMD's K8 – K10.5-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)	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 (NPT)	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. 17h Mod. 00h-0Fh		6/2017	Epyc (2S!!) (4 dies/proc.)	7000	14 nm	4x(2x4) (32C)	½ MB/C	2 MB/C	DDR4-2666	IFIS 75.8 GB/s	SP3

4. Overview of AMD's K8 – K10.5 (Hammer-based) families (6)

Main features of AMD's high-performance K8- K10.5-based desktop 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. 15h based lines		-	-	-	-	-	-	-	-	-	-

4. Overview of AMD's K8 – K10.5 (Hammer-based) families (7)

Main features of AMD's K8 – K 10.5-based high performance mobile lines

Base arch./ stepping		Intro	High perf. mobile family name	Series	Techn.	Core count (up to)	L2 (up to)	L3	Memory (up to)	HT/ dir. (up to)	Sock et
K8	C0, CG	9/2003	Clawhammer	Mobile Athlon 64	130 nm	1	512 KB	-	DDR-400	HT 1.0: 3.2 GB/s	754
	E5	3/2005	Lancaster	Turion 64	90 nm	1	1 MB	-	DDR-400	HT 1.0: 3.2 GB/s	754
	F2	5/2006	Trinidad	Turion 64 X2	90 nm	2	2*512 KB	-	DDR2-667	HT 1.0: 3.2 GB/s	S1
K10	-	-	-	-	-	-	-	-	-	-	-
K10.5	DA-C2	9/2009	Caspian	Turion II	45 nm	2	2*512 KB/ 2*1 MB[1]	-	DDR2-800	HT 3.0: 7.2 GB/s	S1g3
	DA-C3	5/2010	Champlain	Turion X4	45 nm	4	4*512 KB	-	DDR3-1066	HT 3.0: 7.2 GB/s	S1g4

[1] 2*512 KB for Turion II, 2*1 MB for Turion II Ultra

4. Overview of AMD's K8 – K10.5 (Hammer-based) families (8)

Brand names of AMD's K8 – K10.5-based processor lines

		2003-2007	2007-2008	2008-2011	2009	2009
		K8 (Hammer)	**K10 (Barcelona)**	**K10.5 (Shanghai)**	**K10.5 (Istanbul)**	**K10.5 (Magny- Course)**
Servers	**4P servers**	**See Section 4**	**Barcelona (834x-836x))**	**Shanghai (837x-839x)**	**Istambul (8410-8430)**	**Magny-Course (6100)**
	2P servers		**Barcelona (234x-236x)**	**Shanghai (237x-239x)**	**Istambul (241x-243x)**	**Lisbon (4100)**
	1P servers		**Budapest (135x-136x)**	**Suzuka (138x-139x)**		
Desktops	**High perf.** (~80-120W)		Phenom X4-X2	Phenom II X4-X2	Phenom II X6-X4	
	Mainstream (~60-90W)	Athlon 64 Athlon 64 X2	Athlon X2	Athlon II X4-X2		
	Value (~40-60W)	Sempron		Sempron		
Mobiles	**High perf.** (~30-40W)	Turion 64 X2 (TL 6/5) Turion 64 (ML/MT)		Phenom II (N/P 9xx-6xx) Turion II Ultra (M6xx) Turion II (M/N/P 5xx)		
	Mainstream (~20-30W)	Athlon 64 X2 (TK-5x/4x) Athlon 64 (2xxx+-4xxx+)		Athlon II (M/N/P 3xx) Sempron (M1xx)		
	Ultraportable (~10-20W)	Mobile Sempron (2xxx+-4xxx+) Sempron 2100 fanless		Turion II Neo (K6xx) Athlon II Neo (K1xx) V-series (V1xx)		
Embedded (~10-20W)				Turion II Neo X2 Athlon II Neo X2 Athlon II Neo		

4. Overview of AMD's K8 – K10.5 (Hammer-based) families (9)

Remark

AMD's naming conventions in their K8- K14 lines [15]

- K10/K10.5/K15 servers: GP racing places
- K8 desktops/mobiles: cities (except the first K8 implementation (ClawHammer)
- K10/K10.5 desktops/mobiles: stars
- Server sockets (platforms): Ferrari facilities and race tracks
- K10 and K14 mobile platforms: rivers
- K10 based high performance desktop platforms: animals

Examples for naming of K10-K10.5 servers

65nm Quad-Core: Barcelona [Circut de Cataluña, Spanish GP],
45nm Quad-Core: Shanghai [Shanghai Guoji Saichechang, Chinese GP]
45nm Sexa-Core: Istanbul [Istanbul Park, Turkish GP]
45nm Quad-Core: Lisbon [Autodromo do Estoril, former Portugese GP]
45nm Octa- to Duodec-Core [8 to 12 cores]: Magny Cours [Circuit de Nevers Magny-Cours, former French GP]
32nm Sexa- to Octa-Core Single-die [6 to 8 cores]: Valencia [Valencia Street Circuit, European GP]
32nm Duodec- to Sedec-Core [12 to 16 cores]: Sao-Paolo or Interlagos [Autodromo Jose Carlos Pace, Brasil GP]

4. Overview of AMD's K8 – K10.5 (Hammer-based) families (10)

Remark (cont.)

The reason for choosing car racing related processor designations lies in the connection of AMD with motor racing, as

- FIA (Federation Internationale de Automobile) is using an AMD-based supercomputer, and
- AMD has good connections to Ferrari through personal interests of a former top manager [15].

4. Overview of AMD's K8 – K10.5 (Hammer-based) families (11)

Evolution of AMD's x86-64 server lines K8 – Family 10.5 – Key features [16]

	2003	2005	2007	2008	2009	2010
	AMD Opteron™	AMD Opteron™	"Barcelona"	"Shanghai"	"Istanbul"	"Magny-Cours"
Mfg. Process	130 nm ~~90nm~~ SOI	90nm SOI	65nm SOI	45nm SOI	45nm SOI	45nm SOI
CPU Core	K8	K8	Greyhound	Greyhound+	Greyhound+	Greyhound+
L2/L3	1MB/0	1MB/0	512kB/2MB	512kB/6MB	512kB/6MB	512kB/12MB
Hyper Transport™ Technology	3x 1.6GT/.s	3x 1.6GT/.s	3x 2GT/s	3x 4.0GT/s	3x 4.8GT/s	4x 6.4GT/s
Memory	2x DDR1 ~~300~~ 333	2x DDR1 400	2x DDR2 667	2x DDR2 800	2x DDR2 1066	4x DDR3 1333

Max Power Budget Remains Consistent

Evolution of AMD's x86-64 server lines K8 – Family 1.5 – Main features [17]

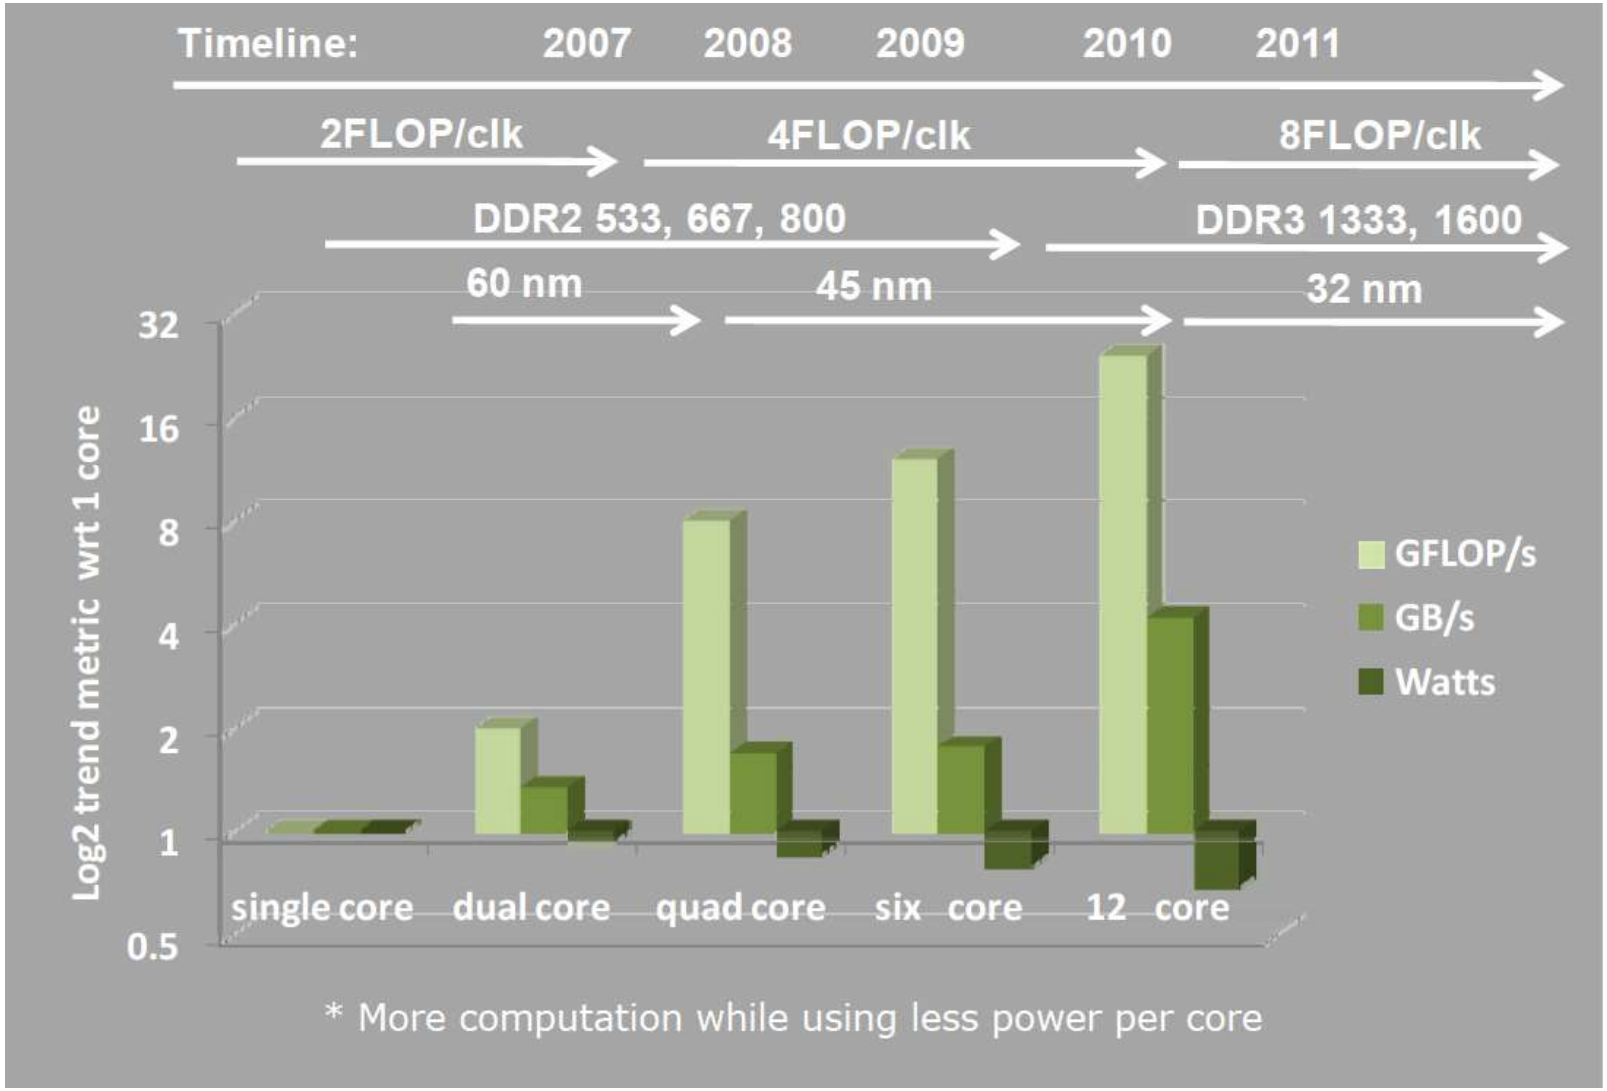

4. Overview of AMD's K8 – K10.5 (Hammer-based) families (13)

Evolution of main features of AMD's DP/MP servers K8 – Family 17h

Base arch./ stepping		Intro.	Core	Techn. (nm)	Server family name	Cores	New key features						Use
							L3	Mem.	On-die MC	HT	ISA extension	NX	
K8	B3/CG	4/2003	Sledge-hammer	130	Sledgehammer	1C	-	DDR	+On-die MC	3xHT 1.0	+SSE2	+NX -bit	S/DT/M
	E4	12/2004		90	Athens					3xHT 2.0	+SSE3		
	E1/E6	4/2005			Egypt	2C							
	F2/F3	8/2006			Santa Rosa			DDR2					
	G1/G2	12/2006		65	DT: Brisbane								DT
K10	B2/B3	9/2007	Greyhound	65	Barcelona[5]	4C	6 MB				+SSE4a		S/DT
K10.5	C2/C3	11/2008	Greyhound +	45	Shanghai					3xHT 3.0			S/DT/M
	CE	6/2009			Istambul	6C				4xHT 3.1			S/DT
	D1	3/2010			Magny Course (2xIstambul)	2x6C	2x 6 MB	DDR3					S
Fam. 15h	Mod. 0xh	11/2011	Bulldozer	32	Interlagos (2xOrochi)	2x8C	2x 8 MB				+SSE4.1/4.2, AES, AVX, XOP, FMA4,CMUL		S/DT
	Mod. 1xh	11/2012	Piledriver		Abu Dhabi (Dual dies)						+FMA3, CVT16, BMI, TBM		S/DT/M
Fam. 17h	Mod. 0xh	6/2017	Zen	14	DT: Epic (4 dies)	4x8C	2 MB/ core	DDR4		IFIS	na.		S/HED

[1] x4UMI: 4x PCIe 2.0

[2] ISA enh.: +AES, +AVX, +FMA4, +XOP, +PCLMULQDQ

[3] PCIe 1.0/2.0

[4] 3DNow! Prof. dropped

[5] The Barcelona die supports already 4xHT 3.0 and DDR3 but Socket F used for DP/MP servers restricts supported features to 3xHT 2.0 and DDR2

5. The K8 (Hammer) family

- 5.1 Overview of the K8 family
- 5.2 AMD's Direct Connect Architecture (DCA) multiprocessor concept
- 5.3 The microarchitecture of the K8 (Sledgehammer) core
- 5.4 Overview of AMD's native K8 designs
- 5.5 K8 server lines
- 5.6 K8 desktop lines
- 5.7 K8 mobile lines
- 5.8 Evolution of the high performance K8-based Athlon 64 and Athlon 64 X2 desktop lines as reflected by die shots

5.1 Overview of the K8 family

5.1 Overview of the K8 family

(The introductory part is recited from Section 3 to give a more complete account of the evolution of the K8 family.)

5.1 Milestones of the introduction of AMD's K8 family-1

- A few month after introducing their highly successful K7 (Athlon) family in 10/1999 AMD disclosed their plan
  - to make a compatible extension of the x86 ISA, designated as the x86-64 ISA
  - to implement it as their eights generation (K8) processor family, code named as Sledgehammer and
  - to use the serial two byte wide Lighting Data Transport bus (renamed in 2001 to HyperTransport bus) as a chip-to-chip interconnect bus to provide enough I/O bandwidth.
- 8/2000 Release of the x86-64 architecture specification (i.e. x86-64 ISA) to encourage the software community to begin incorporating x86-64™ technology into operating systems, applications, drivers and development tools [25] .

5.1 Overview of the K8 family (2)

AMD's approach to introduce 64-bit computing

AMD's goal was clearly formulated at revealing the x86-64 ISA specification in 8/2000 [25]:

- "Ultimately this technology is designed to help preserve the enterprise community's enormous financial investment in 32-bit operating systems, applications, development tools and support infrastructure while providing a seamless path to deploy future 64-bit technology."
- "Perhaps the most noteworthy feature of AMD's approach to 64-bit computing is that it is an extension to the 32-bit environment prevailing in the industry today rather than a radical departure."

5.1 Overview of the K8 family (3)

The programmer's model of the available register set in the x86-64 ISA [28]

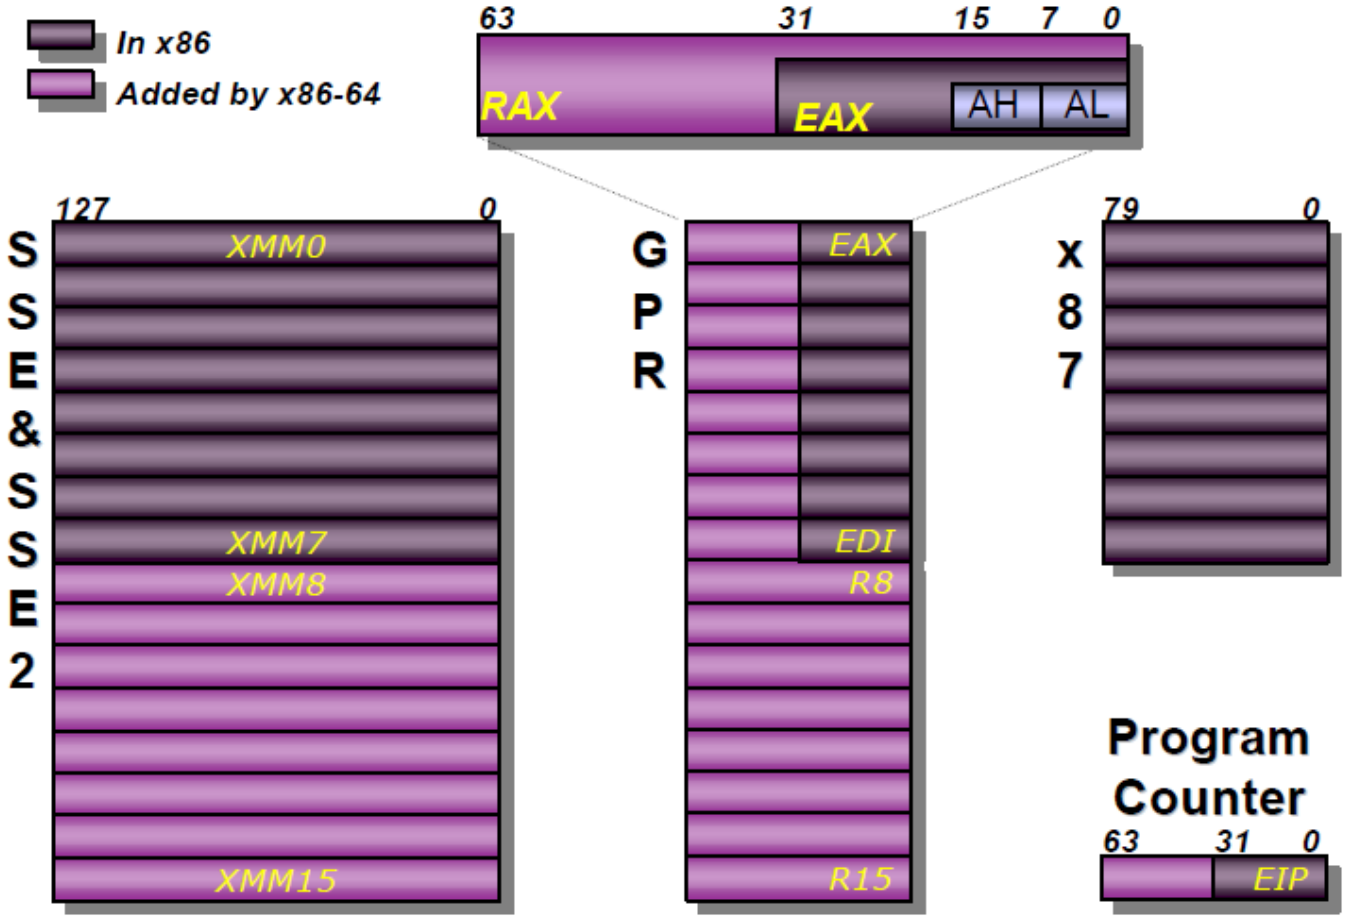

5.1 Overview of the K8 family (4)

Milestones of the introduction of AMD's K8 family-2

- In 10/2001 AMD reveals details of the planned "Hammer" multiprocessor architecture, that is described subsequently.

  The Hammer architecture became renamed in 2004 to Direct Connect Architecture [26].
- 4/2002 AMD announces Microsoft support for their x86-64 family [27].
- 4/2002 Disclosure of the Opteron designation for the x86-64 server line.
- 4/2003 First shipment of K8 processors (the Sledgehammer line of Opteron servers).

  Both the introduced server and subsequent desktop processors were superior to Intel's existing Pentium 4 based processors [123].

5.2 AMD’s Direct Connect Architecture (DCA) multiprocessor concept

5.2 AMD's Direct Connect Architecture (DCA) multiprocessor concept

It is AMD's multiprocessor architecture concept.
It originates from about 1999 and was designated as the Hammer Architecture until 2004.

AMD's Direct Connect Architecture (DCA) concept provides an efficient framework for multiprocessors and is characterized by

- an Integrated Memory Controller
  to eliminate the memory bottleneck
- up to 3 serial HyperTransport links
  for chip-to-chip-communication
  to eliminate the FSB bottleneck [28].

  These links interconnect processors
  in DP (2P) and MP (4P) multiprocessors
  as well as processors with the chipset.

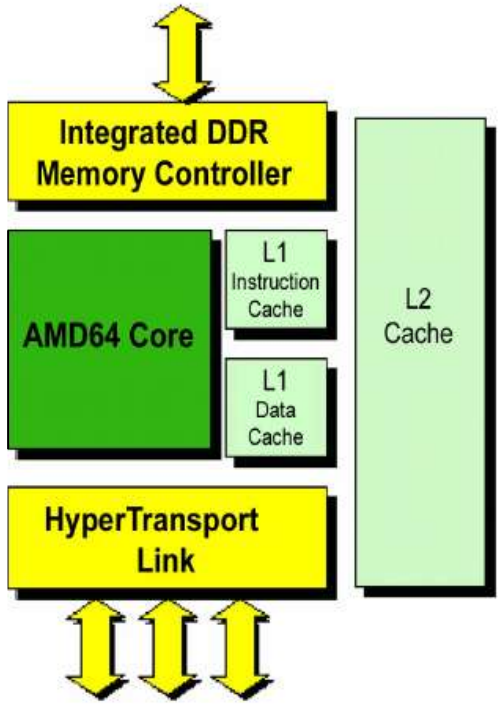

Figure: AMD's Direct Connect Architecture [37]

5.2 AMD's Direct Connect Architecture (DCA) multiprocessor concept (2)

The HyperTransport links

There are 2-Byte wide standard serial chip-to-chip interconnect buses (called before 2001 Lighting Data Transport buses) [26].

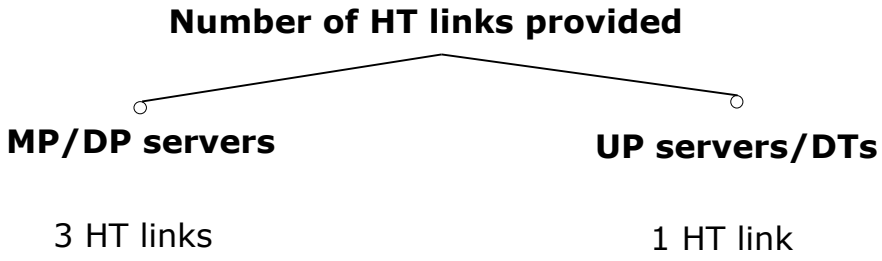

5.2 AMD's Direct Connect Architecture (DCA) multiprocessor concept (3)

Use of HT links in MP servers [28]

Two links are used to interconnect processors and a 3. HT link is available to connect I/O. Note that the DCA implemented only a partial mesh for connecting the processors.

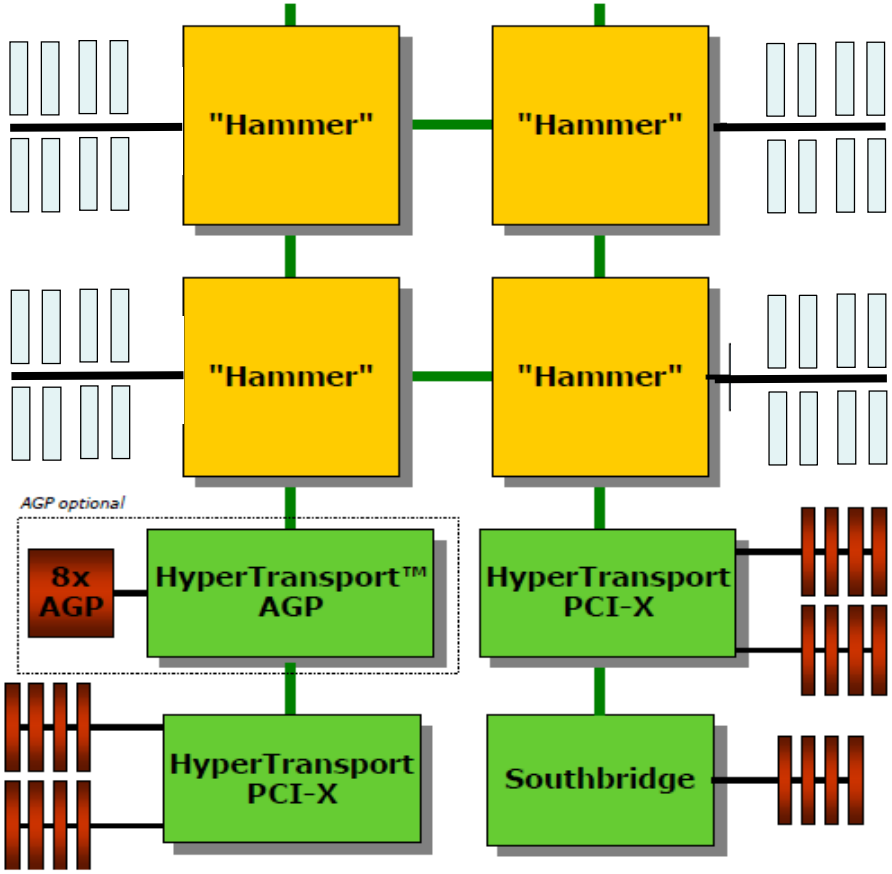

5.2 AMD's Direct Connect Architecture (DCA) multiprocessor concept (4)

Use of HT links in DP servers [28]

One link is used to interconnect both processors and a two HT links are available to connect I/O

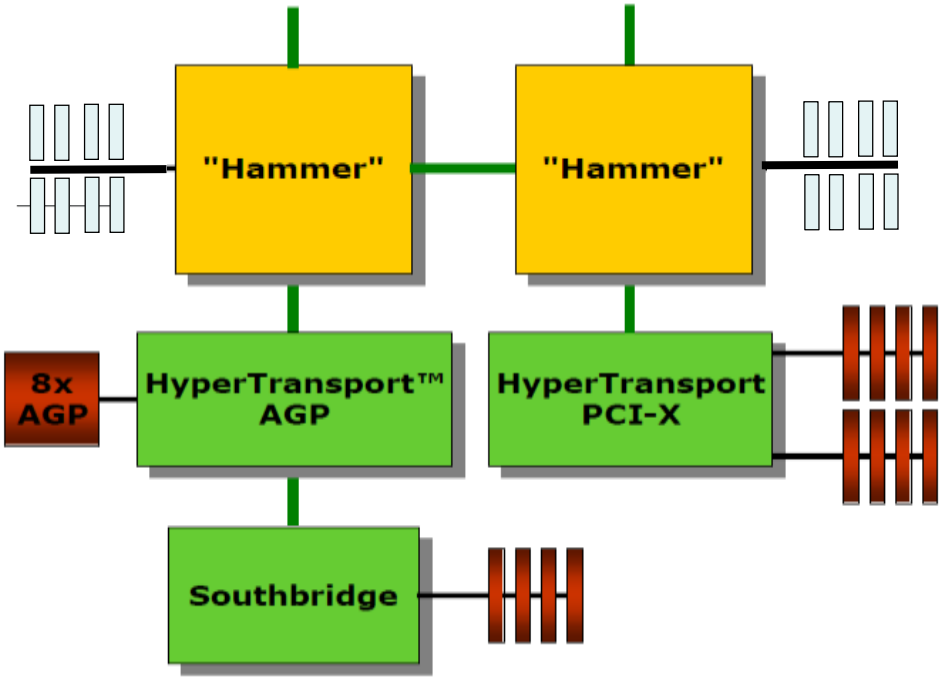

5.2 AMD's Direct Connect Architecture (DCA) multiprocessor concept (5)

Use of HT links in UP servers and DTs [28]

The available HT link is used to connect I/O

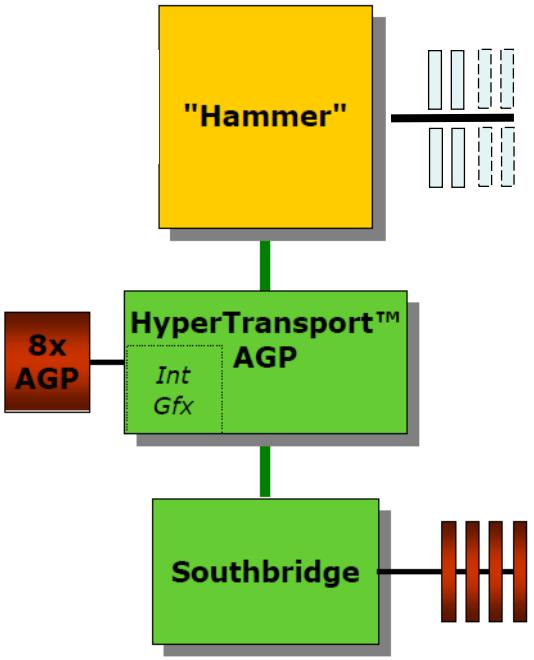

Typical memory

UP servers: 2x4 RDIMMs
DT servers: 2x2 UDIMMs

5.2 AMD's Direct Connect Architecture (DCA) multiprocessor concept (6)

Remarks

1) When AMD announced the „Hammer" architecture in 2001 and later unveiled the Opteron and Athlon64 processors in 2003 they did not yet use the term "Direct Connect Architecture".

This term appeared in the literature about in the beginning of 2004, i.e. roughly one year after shipping the first x86-64 processors.

5.2 AMD's Direct Connect Architecture (DCA) multiprocessor concept (7)

Remarks - cont.

2) Intel introduced a similar microarchitecture that provides integrated memory controllers and 4 serial links, called the Nehalem microarchitecture, in 2008.

3) Also AMD revamped their DCA in two steps;

- first in their K10 based chips AMD implemented already a 4. HT link (in 2007). Nevertheless, using the 4. link would require to switch to a new socket, but AMD stuck to the old one and does not make use of the 4. link in K10-based servers.
- Subsequently, along with their K10.5 based Magny-Course processor AMD enhanced their Direct Connect Architecture (DC) to Direct Connect Architecture 2.0 (DC2.0) and at the same time they also switched to a new socket (Socket G34) that allowed to use DDR3 memory and utilize all 4 HT links.

5.3 The microarchitecture of the K8 (Sledgehammer) core

5.3 The microarchitecture of the K8 (Sledgehammer) core [38]

- It supports the Direct Connect Architecture multiprocessor concept.
- It is a three-wide 3. gen. superscalar.

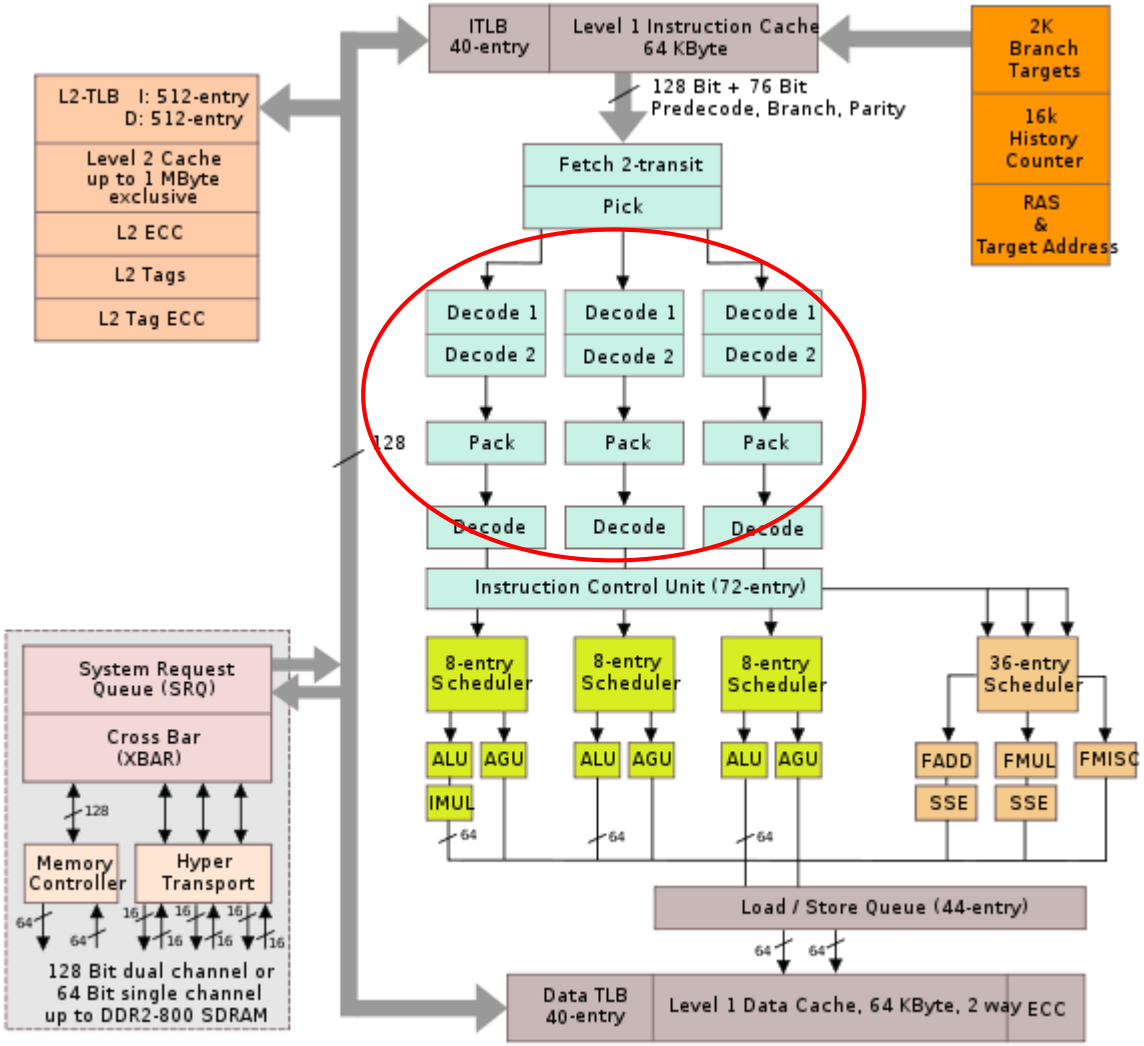

5.3 The microarchitecture of the K8 (Sledgehammer) core (2)

Die plot of the K8 (Sledgehammer) core [39]

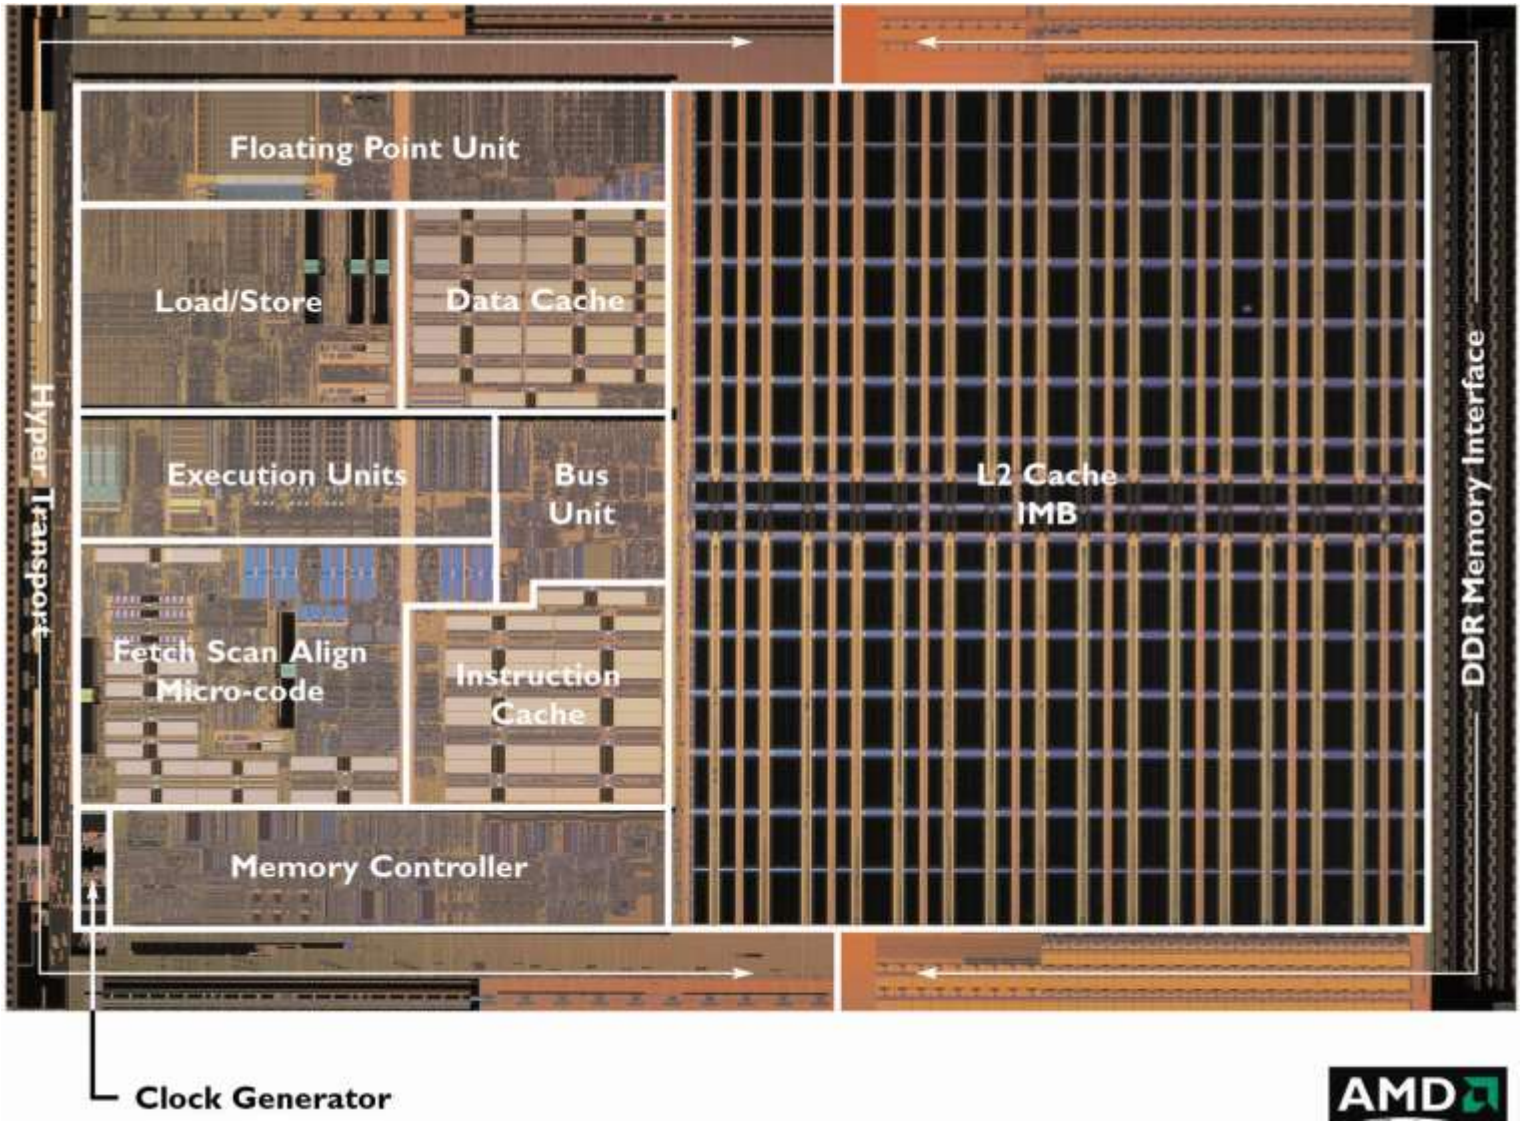

5.3 The microarchitecture of the K8 (Sledgehammer) core (3)

The Integrated Memory Controller of K8-based processors (based on [28])

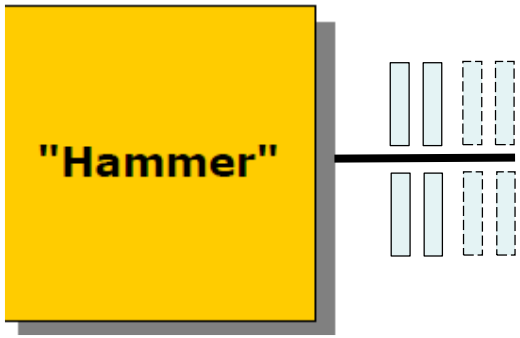

- Single 144 bit wide DDR memory channel (128 bit data and 16 bit ECC)
- Ganged memory controller (DIMMs are used in pairs)
- Up to 4x2 RDIMMs for MP and DP servers
- Up to 2x2 UDIMMs for UP servers and desktops with dual memory channels
- Up to 2 UDIMMs for desktops with a single memory channel

5.3 The microarchitecture of the K8 (Sledgehammer) core (4)

Ganged memory controller [40]

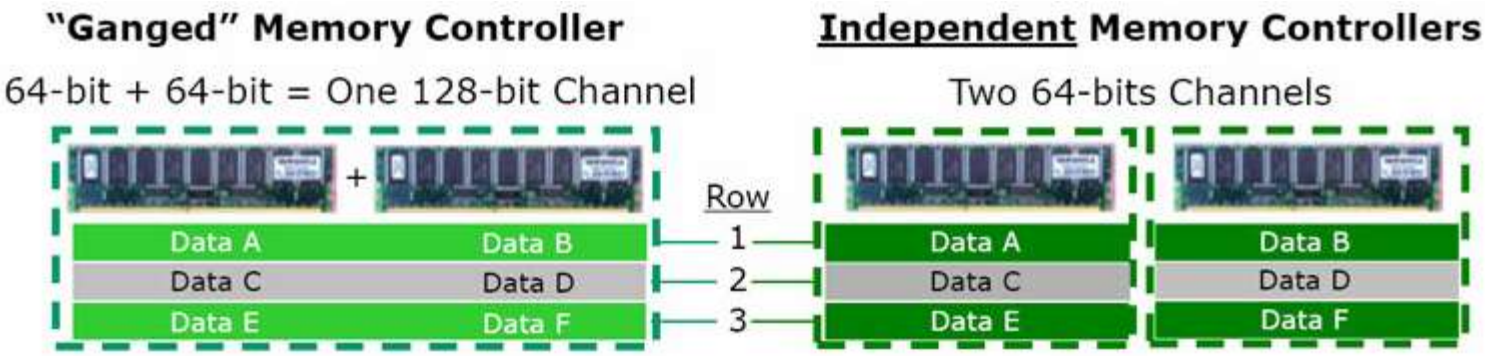

Both have the same max bandwidth, but independent is more efficient. Why?

Let's say we need to multiply Data A times Data F …

Ganged

- Requires two data fetches

- Half of this bandwidth is wasted

Used in K8-based lines

Independent

- Can access two rows at once
- Requires one data fetch

- No wasted bandwidth

Introduced in K10-based lines

5.3 The microarchitecture of the K8 (Sledgehammer) core (5)

Remark

Below we give a more detailed description of the "ganged" memory controller of the K8 microarchitecture.

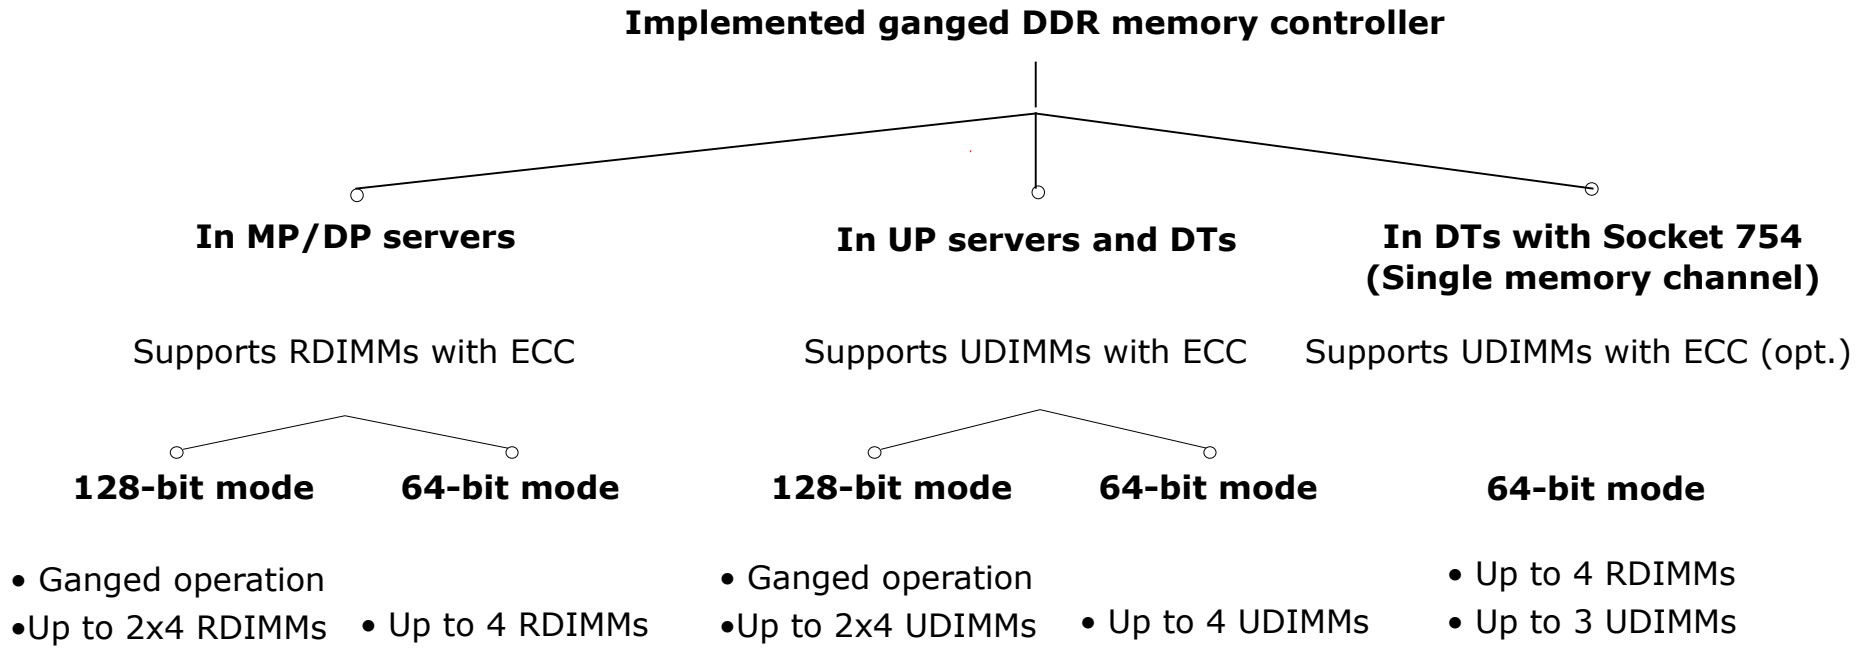

5.4 Overview of AMD’s native K8 designs

5.4 Overview of AMD's native K8 designs

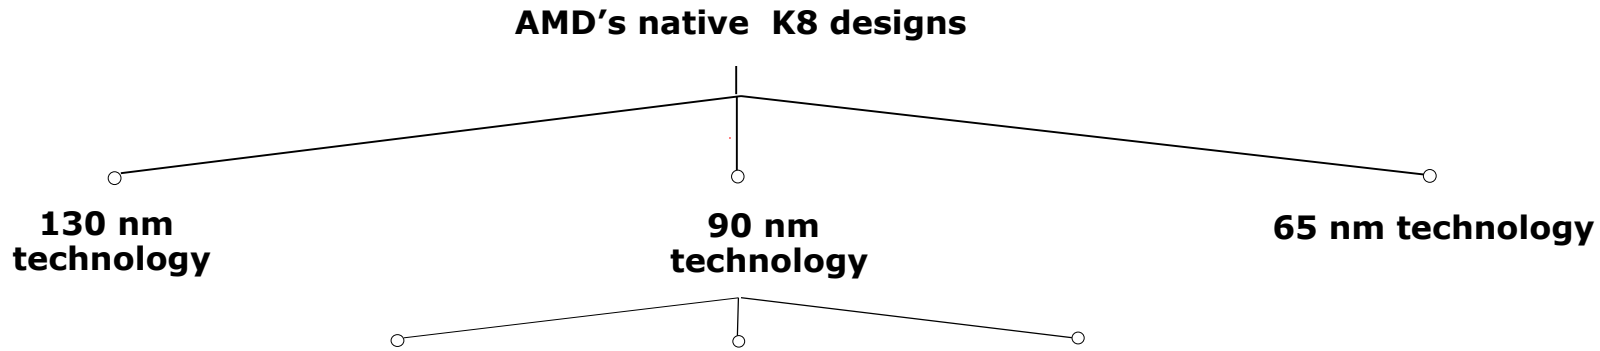

Termed as	***Sledgehammer lines***	***Athens lines***	***Egypt lines***	***Santa Rosa lines***	***Brisbane lines***
Steppings	*B3/C0/CG*	*D4/E4*	*E4/E6*	*F2/F3*	*G1/G2*

5.4 Overview of AMD's native K8 designs (2)

Overview AMD's 64-bit K8-based server and desktop lines (130-90 nm) [13]

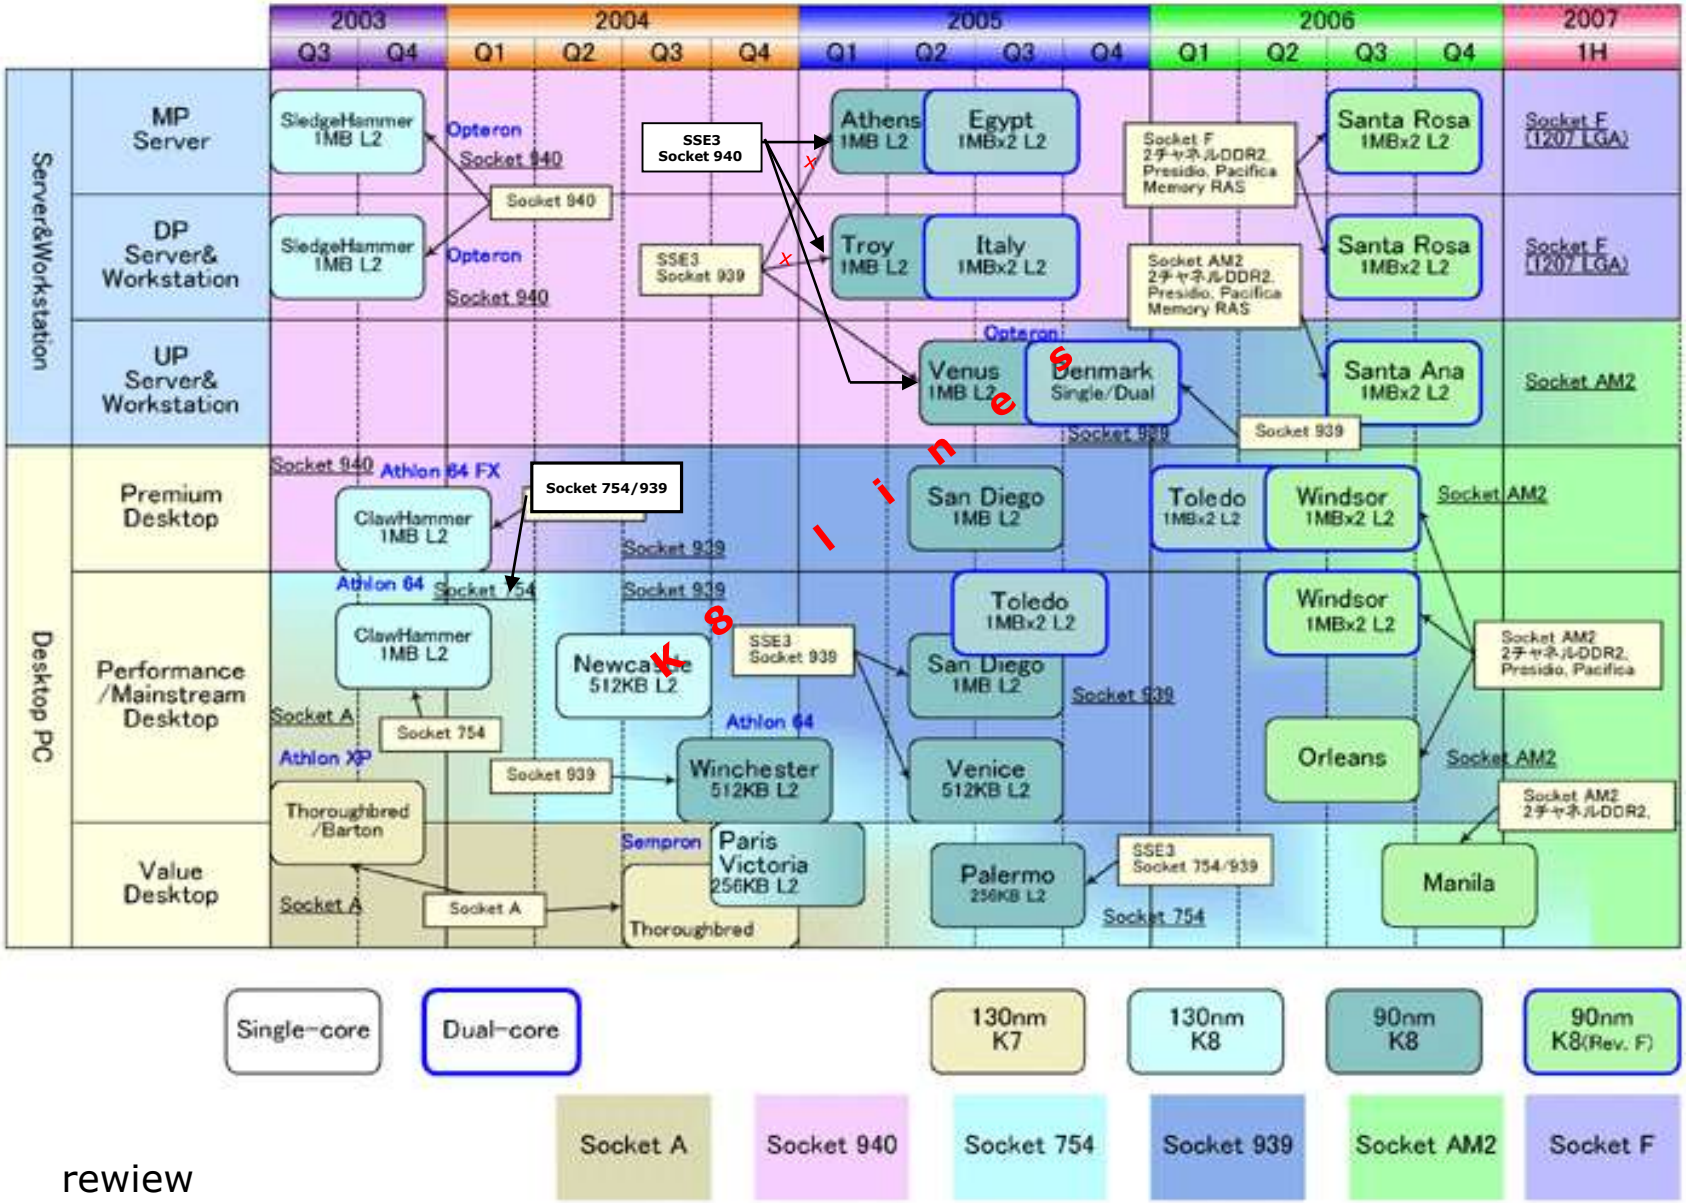

rewiew

5.4 Overview of AMD's native K8 designs (3)

Overview of AMD's 65 nm K8 Brisbane lines [14]

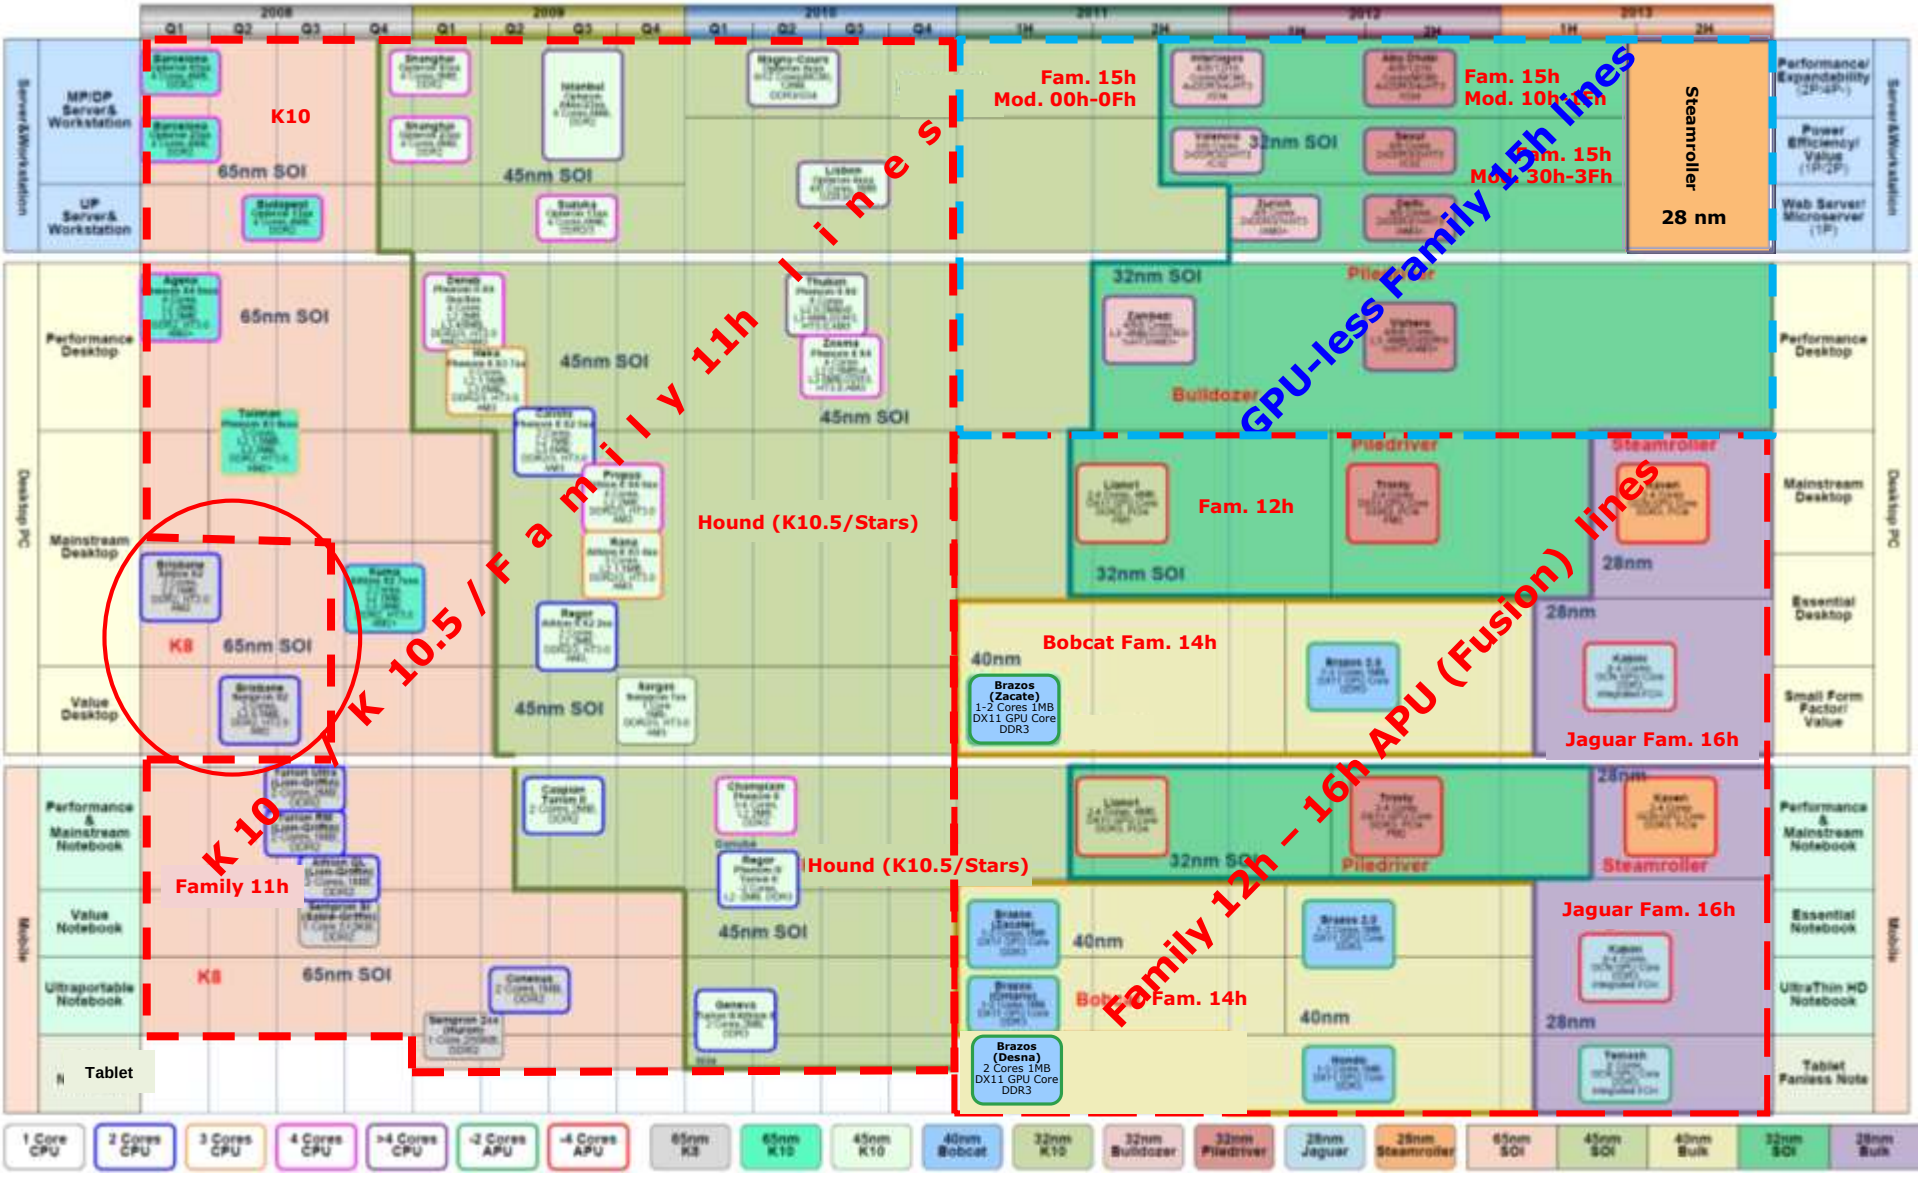

5.4 Overview of AMD's native K8 designs (4)

Evolution of main features of AMD's K8 –based DP/MP servers

Base arch./ stepping		Intro.	Core	Techn. (nm)	Server family name	Cores	New key features						Use
							L3	Mem.	On-die MC	HT	ISA extension	NX	
K8	B3/CG	4/2003	Sledge-hammer	130	Sledgehammer	1C	-	DDR		3xHT 1.0	+SSE2		S/DT/M
	E4	12/2004		90	Athens					3xHT 2.0	+SSE3		
	E1/E6	4/2005			Egypt	2C							
	F2/F3	8/2006			Santa Rosa								
	G1/G2	12/2006		65	DT: Brisbane								DT
K10	B2/B3	9/2007	Greyhound	65	Barcelona[5]	4C	6 MB	DDR2	+On-die MC		+SSE4a	+NX -bit	S/DT
K10.5	C2/C3	11/2008	Greyhound +	45	Shanghai					3xHT 3.0			S/DT/M
	CE	6/2009			Istambul	6C							S/DT
	D1	3/2010			Magny Course (2xIstambul)	2x6C	2x 6 MB	DDR3		4xHT 3.1			S
Fam. 15h	Mod. 0xh	11/2011	Bulldozer	32	Interlagos (2xOrochi)	2x8C	2x 8 MB				+SSE4.1/4.2, AES, AVX, XOP, FMA4,CMUL		S/DT
	Mod. 1xh	11/2012	Piledriver		Abu Dhabi (Dual dies)						+FMA3, CVT16, BMI, TBM		S/DT/M
Fam. 17h	Mod. 0xh	6/2017	Zen	14	DT: Epic (4 dies)	4x8C	2 MB/ core	DDR4		IFIS	na.		S/HED

[1] x4UMI: 4x PCIe 2.0

2 ISA enh.: +AES, +AVX, +FMA4, +XOP, +PCLMULQDQ

[3] PCIe 1.0/2.0

[4] 3DNow! Prof. dropped

[5] The Barcelona die supports already 4xHT 3.0 and DDR3 but Socket F used for DP/MP servers restricts supported features to 3xHT 2.0 and DDR2

5.4 Overview of AMD's native K8 designs (5)

AMD's K8 family

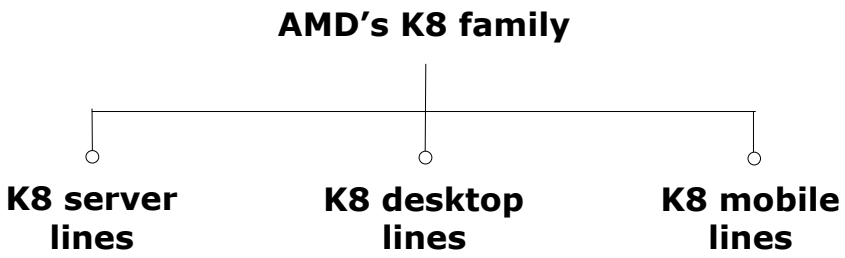

5.5 K8 server lines

5.5 K8 server lines
Overview of AMD's K8 server designs

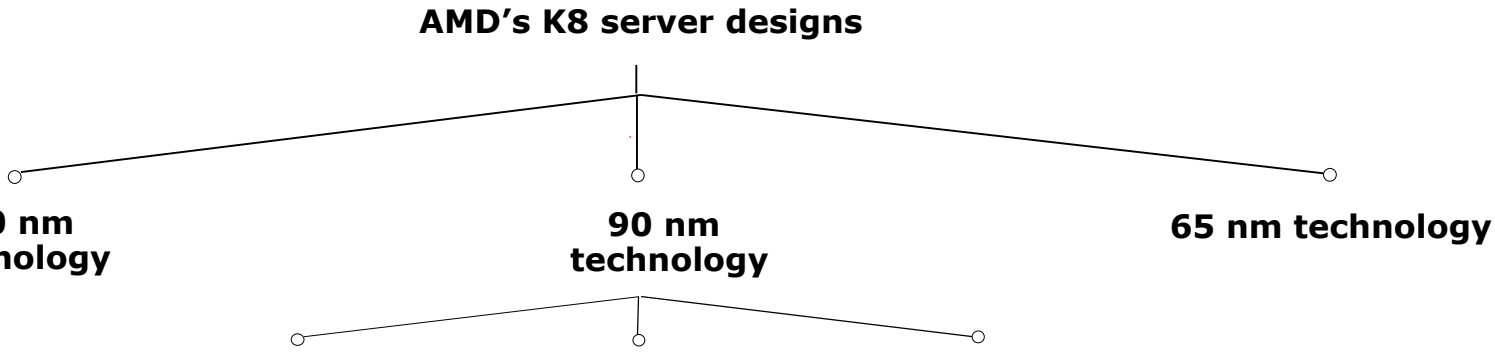

	Sledgehammer lines	***Athens lines***	***Egypt lines***	***Santa Rosa lines***	***Brisbane lines***
Steppings	*B3/C0/CG*	*D4/E4*	*E4/E6*	*F2/F3*	
MP	*Sledgehammer*	*Athens*	*Egypt*	*Santa Rosa*	
DP	*Sledgehammer*	*Troy*	*Italy*	*Santa Rosa*	
UP		*Venus*	*Denmark*	*Santa Ana*	

5.5 K8 server lines (2)

Overview AMD's K8 server lines [13]

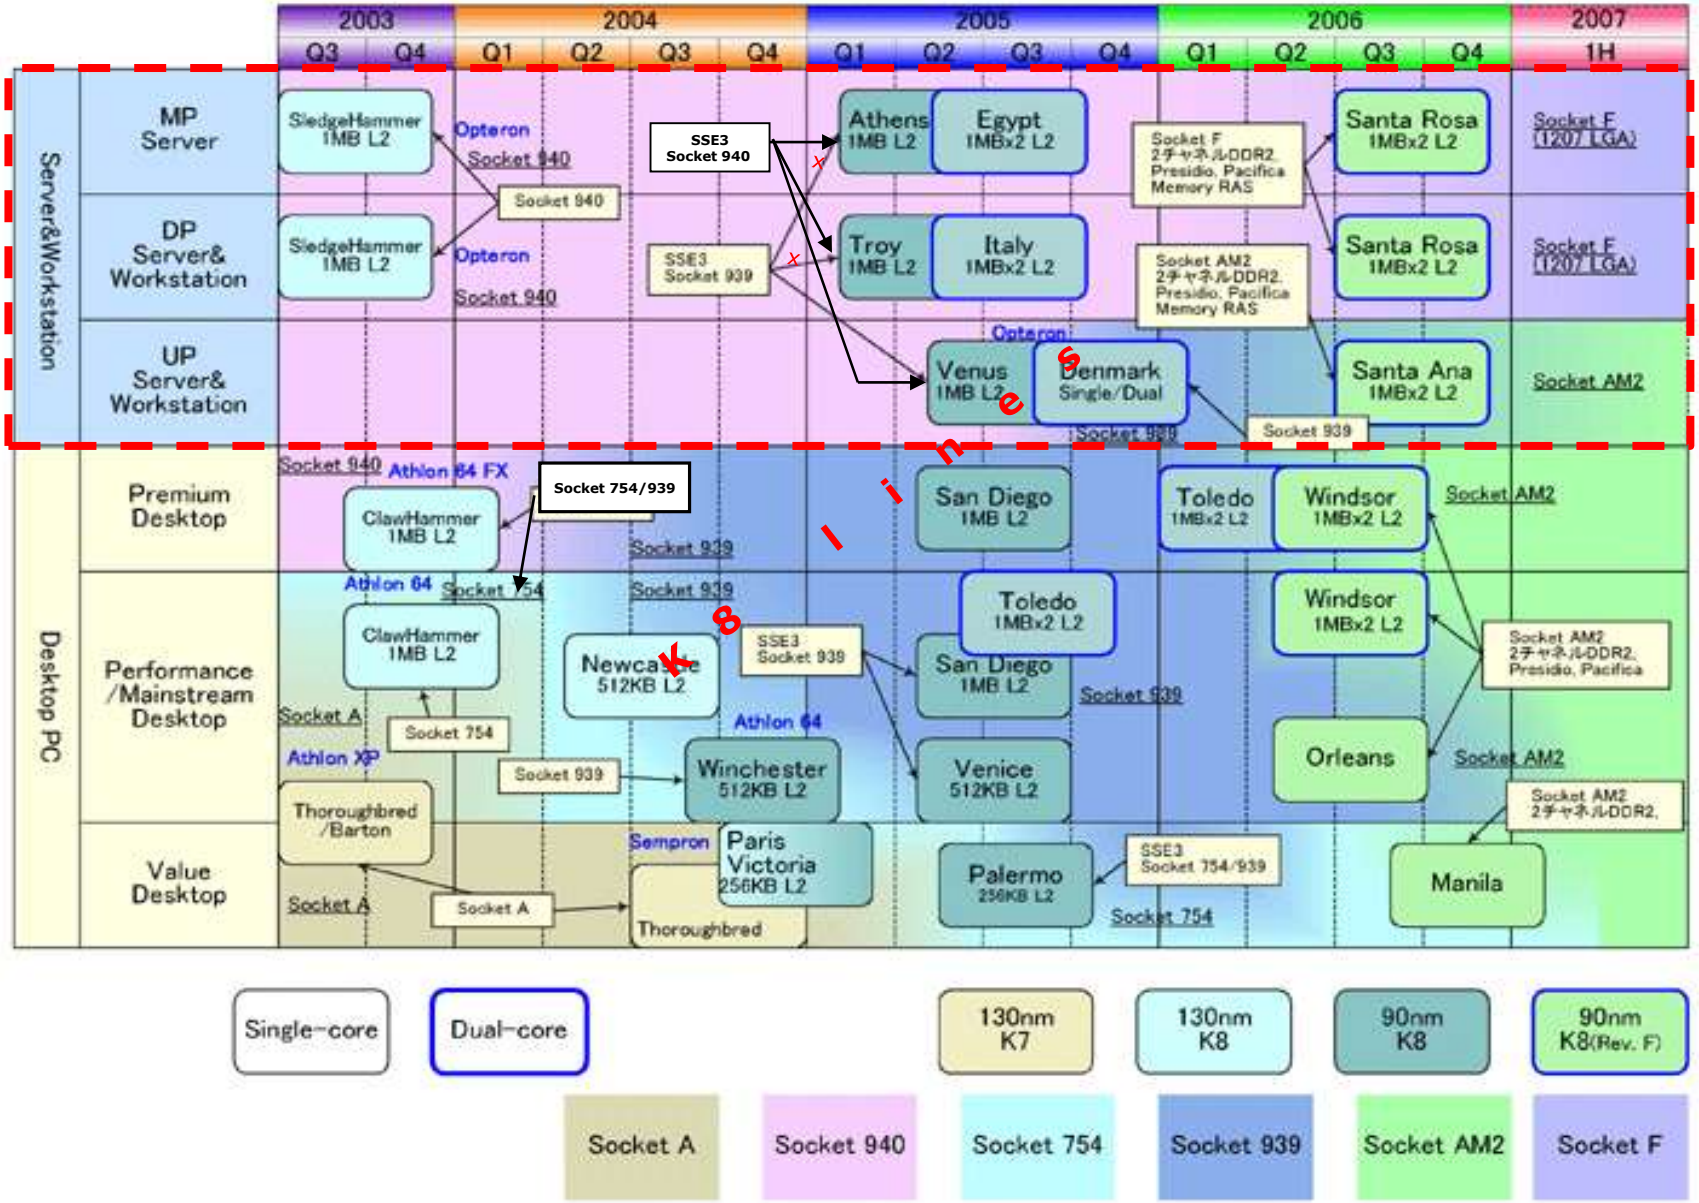

5.5 K8 server lines (3)

Overview of AMD's K8 native K8 server cores (Data are based on [133])

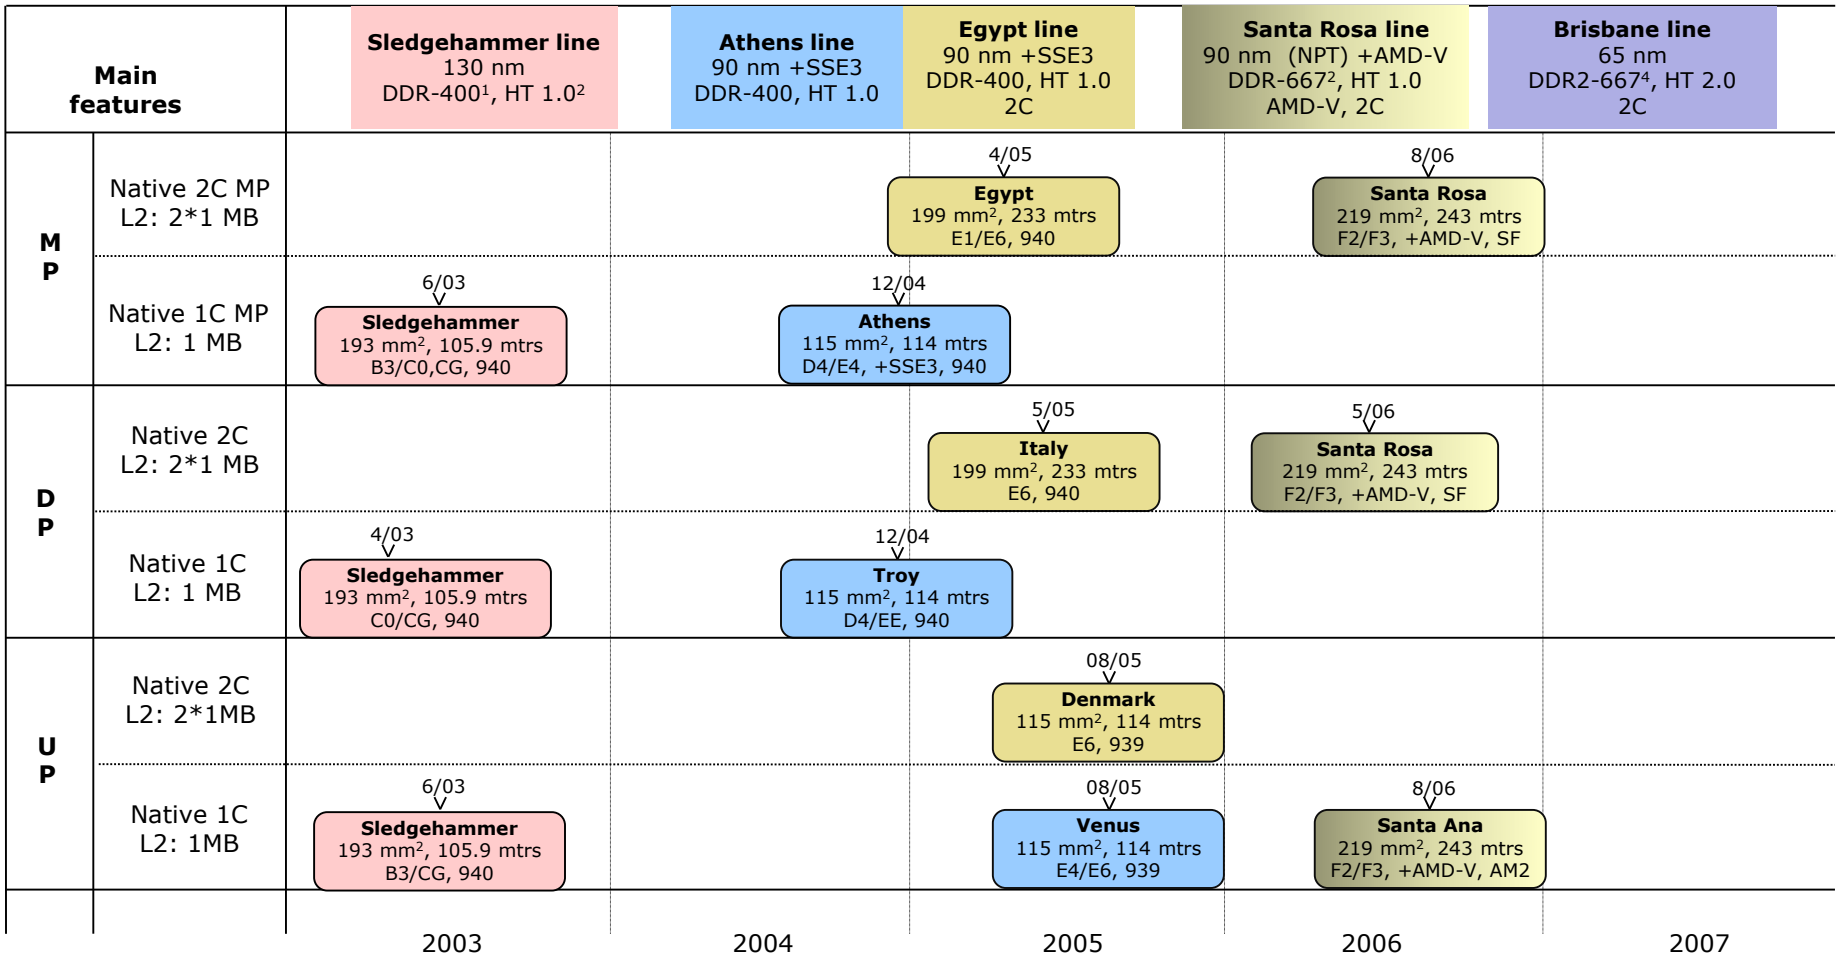

5.5 K8 server lines (4)

Sockets of K8 server processors

[based on 42]

Supports DDR3

Socket **G34** (1944 pins, MP servers K10.5 „Magny-Course", Fam. 15h)
Socket **C32** (1207 pins, DP servers K10.5 „Magny –Course, Fam. 15h)
Socket **AM3** (941 pins, K10.5 Shanghai UP servers, desktops, except first released desktops, K10.5 Istambul desktops)
Socket **AM3+** (942 pins, UP servers Fam. 15h, desktops Fam. 15h)
Socket **FM1** (905 balls BGA, desktops, mobiles Fam. 12h)
Socket **S1g4** (638 pins, K10.5 mobile)
Socket **FT1** (413 balls BGA, desktops, mobiles Fam. 14h)
Socket **FT3** (769 balls BGA, desktops, mobiles Fam. 16h)

S:: Servers
DT: Desktop
M: Mobile
UPT: Ultraportable

SPP: Single power Plane
DPP: Dual Power planes
HT 2.0: HyperTransport 2.0
HT 3.0: HyperTransport 3.0

Supports DDR2

Socket **F** (1207 pins, SPP, HT 2.0: DP, MP servers K8 NPT, K10)
Socket **F+** (1207 pins, DPP, HT 3.0: DP, MP servers K10.5)
Socket **AM2** (940 pins, SPP, HT 2.0: K8 NPT UP serv., K8 NPT, Brisbane DTs)
Socket **AM2+** (940 pins, DPP: K10 UP servers, K10 desktops, first K10.5 Shanghai desktops)
Socket **S1g1** (638 pins, K8 NPT, Brisbane mobiles)
Socket **S1g2** (638 pins, Fam. 11h mobiles)
Socket **S1g3** (638 pins, K10.5 Shanghai (Caspian die) mobiles)
Socket **ASB1** (812 balls BGA, Fam. 11h (Huron, Conesus dies) UPT mobiles)
Socket **ASB2** (812 balls BGA, K10.5 (Shanghai) (Geneva die), UPT mobiles)

Supports DDR

Socket **940** (130nm K8 UP/DP/MP servers, 90 nm K8 SC and 90 nm K8 DC (Egypt) DP/MP servers)
Socket **939** (90nm K8 SC and 90 nm K8 DC (Egypt) UP servers, K8 desktops with 128-bit mem. channel up to the 90 nm DC (Egypt) lines)
Socket **754** (K8 desktops, mobiles with 64-bit mem. channel up to the 90 nm SC (Athens) line)

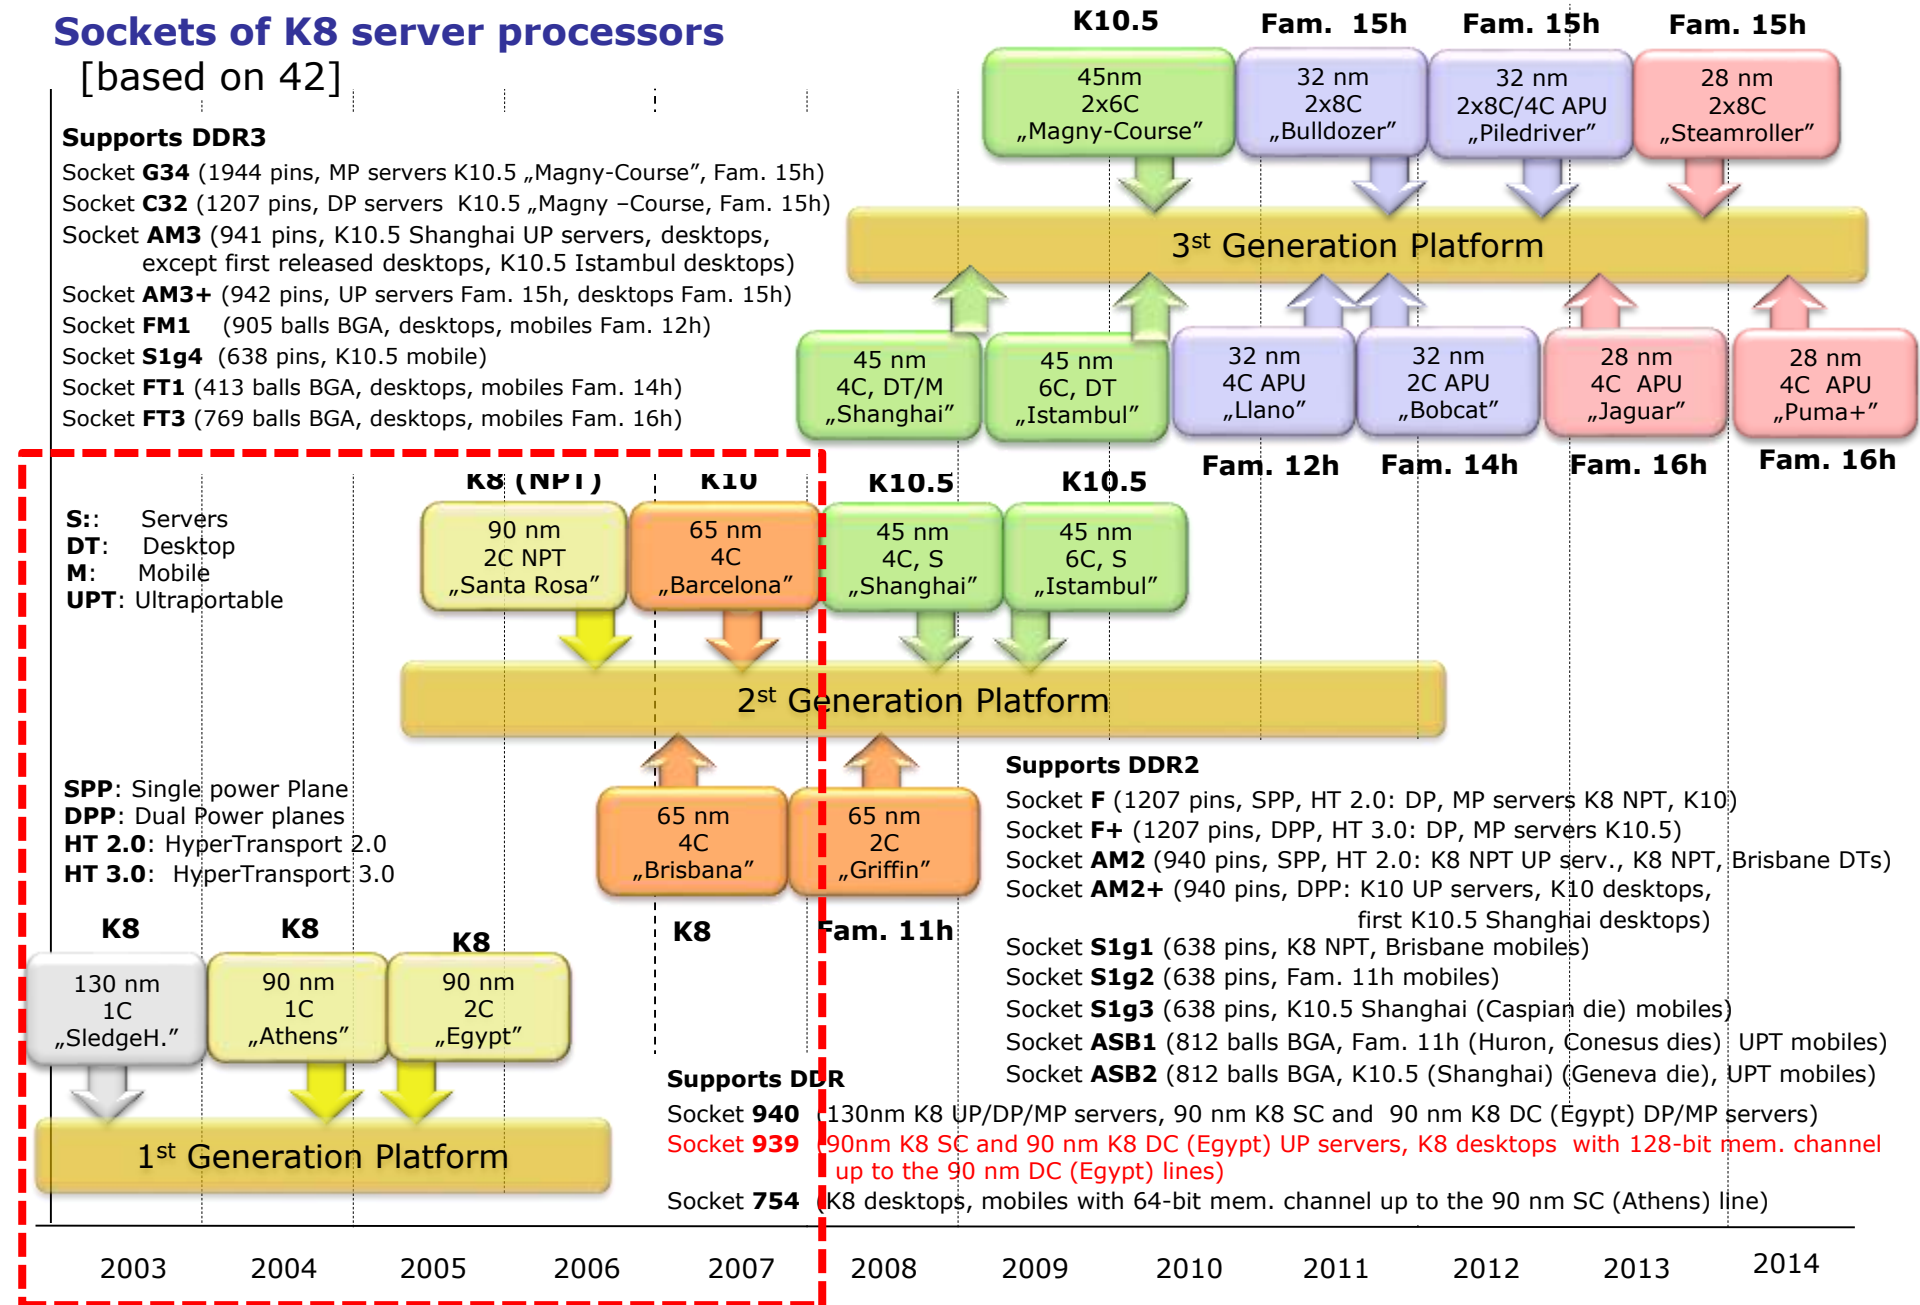

5.5 K8 server lines (5)

Main features of AMD's K8-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)	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 (NPT)	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. 17h Mod. 00h-0Fh		6/2017	Epyc (2S!!) (4 dies/proc.)	7000	14 nm	4x(2x4) (32C)	½ MB/C	2 MB/C	DDR4-2666	IFIS 75.8 GB/s	SP3

5.5 K8 server lines (6)

Main features of AMD's 130 nm SC K8-based Opteron lines (without HE/SE/EE lines)

	4/03 (DP); 6/03 (UP/MP)	9/03	11/03	5/04
MP	**Opteron**	**Opteron**	**Opteron**	**Opteron**
DP	**Sledgehammer**	**Sledgehammer**	**Sledgehammer**	**Sledgehammer**
UP	**Sledgehammer**	**Sledgehammer**	**Sledgehammer**	**Sledgehammer**
Based on	K8	K8	K8	K8
Revision	B3	C0	C0	CG
No of cores	1	1	1	1
Technology	130 nm	130 nm	130 nm	130 nm
Transistors	105.9 mtrs	105.9 mtrs	105.9 mtrs	105.9 mtrs
MP models	844 - 840	846 - 840	848	850 - 840
fc	1.8 - 1.4 GHz	2.0 - 1.4 GHz	2.2 GHz	2.4 - 1.4 GHz
DP models	244 - 240	246 - 240	248	250 - 240
fc	1.8 - 1.4 GHz	2.0 - 1.4 GHz	2.2 GHz	2.4 - 1.4 GHz
UP models	144 - 140	146 - 140	148	150 - 140
fc	1.8 - 1.4 GHz	2.0 - 1.4 GHz	2.2 GHz	2.4 - 1.4 GHz
L2	1 M	1 M	1 M	1 M
L3	–	–	–	–
Memory (reg.)	DDR-333 144 bit	DDR-333 144 bit	DDR-400 144 bit	DDR-400 144 bit
HT	800 MHz	800 MHz	800 MHz	800 MHz
PM	–	–	–	–
TDP	84.7 W	82.1-89.0 W	89.0 W	82.1-89.0 W
NX	NX	NX	NX	NX
Virt. techn.	–	–	–	–
Vcc	1.55 V	1.50 V	1.50 V	1.50 V
Socket	940	940	940	940
C-states	C1, C2, C3	→	→	→
CIE				
S-states	S1, S3-S5	→	→	→
HW throttling	HW Clock throttling	→	→	→

5.5 K8 server lines (7)

Main features of AMD's 90 nm SC K8-based Opteron lines (without HE/SE/EE lines)

	12/04	2/05	8/05	4/06
MP DP UP	**Opteron Athens Troy Venus**	**Opteron Athens Troy Venus**	**Opteron Athens Troy Venus**	**Opteron Athens Troy Venus**
Based on	K8	K8	K8	K8
Revision	E4	E4	E4	E4
No of cores	1	1	1	1
Technology	90 nm	90 nm	90 nm	90 nm
Transistors	114 mtrs	114 mtrs	114 mtrs	114 mtrs
MP models	850 - 842	852	854	856
fc	2.4- 1.6 GHz	2.6 GHz	2.8 GHz	2.6 GHz
DP models	250 - 242	252	254	256
fc	2.4- 1.6 GHz	2.6 GHz	2.8 GHz	3.0 GHz
UP models	150 - 142	152	154	156
fc	2.4- 1.6 GHz	2.6 GHz	2.8 GHz	3.0 GHz
L2	1 M	1 M	1 M	1 M
L3	–	–	–	–
Memory (reg.)	DDR-400 144 bit	DDR-400 144 bit	DDR-400 144 bit	DDR-400 144 bit
HT	1000 MHz	1000 MHz	1000 MHz	1000 MHz
PM	-	PowerNow!	PowerNow!	PowerNow!
TDP	85.3 W	92.6 W	92.6 W	104 W
NX	NX	NX	NX	NX
Virt. techn.	–	–	–	–
Vcc	1.40/1.35 V	1.4 / 1.35 V	1.4 / 1.35 V	1.4 / 1.35 V
Socket	940	940	940	940
C-states CIE	C1, C2, C3 →			
S-states	S1, S3-S5 →			
HW throttling	HW Clock throttling →			

5.5 K8 server lines (8)

AMD's power management techniques K8 – Family 15h (Bulldozer) (based on [53])

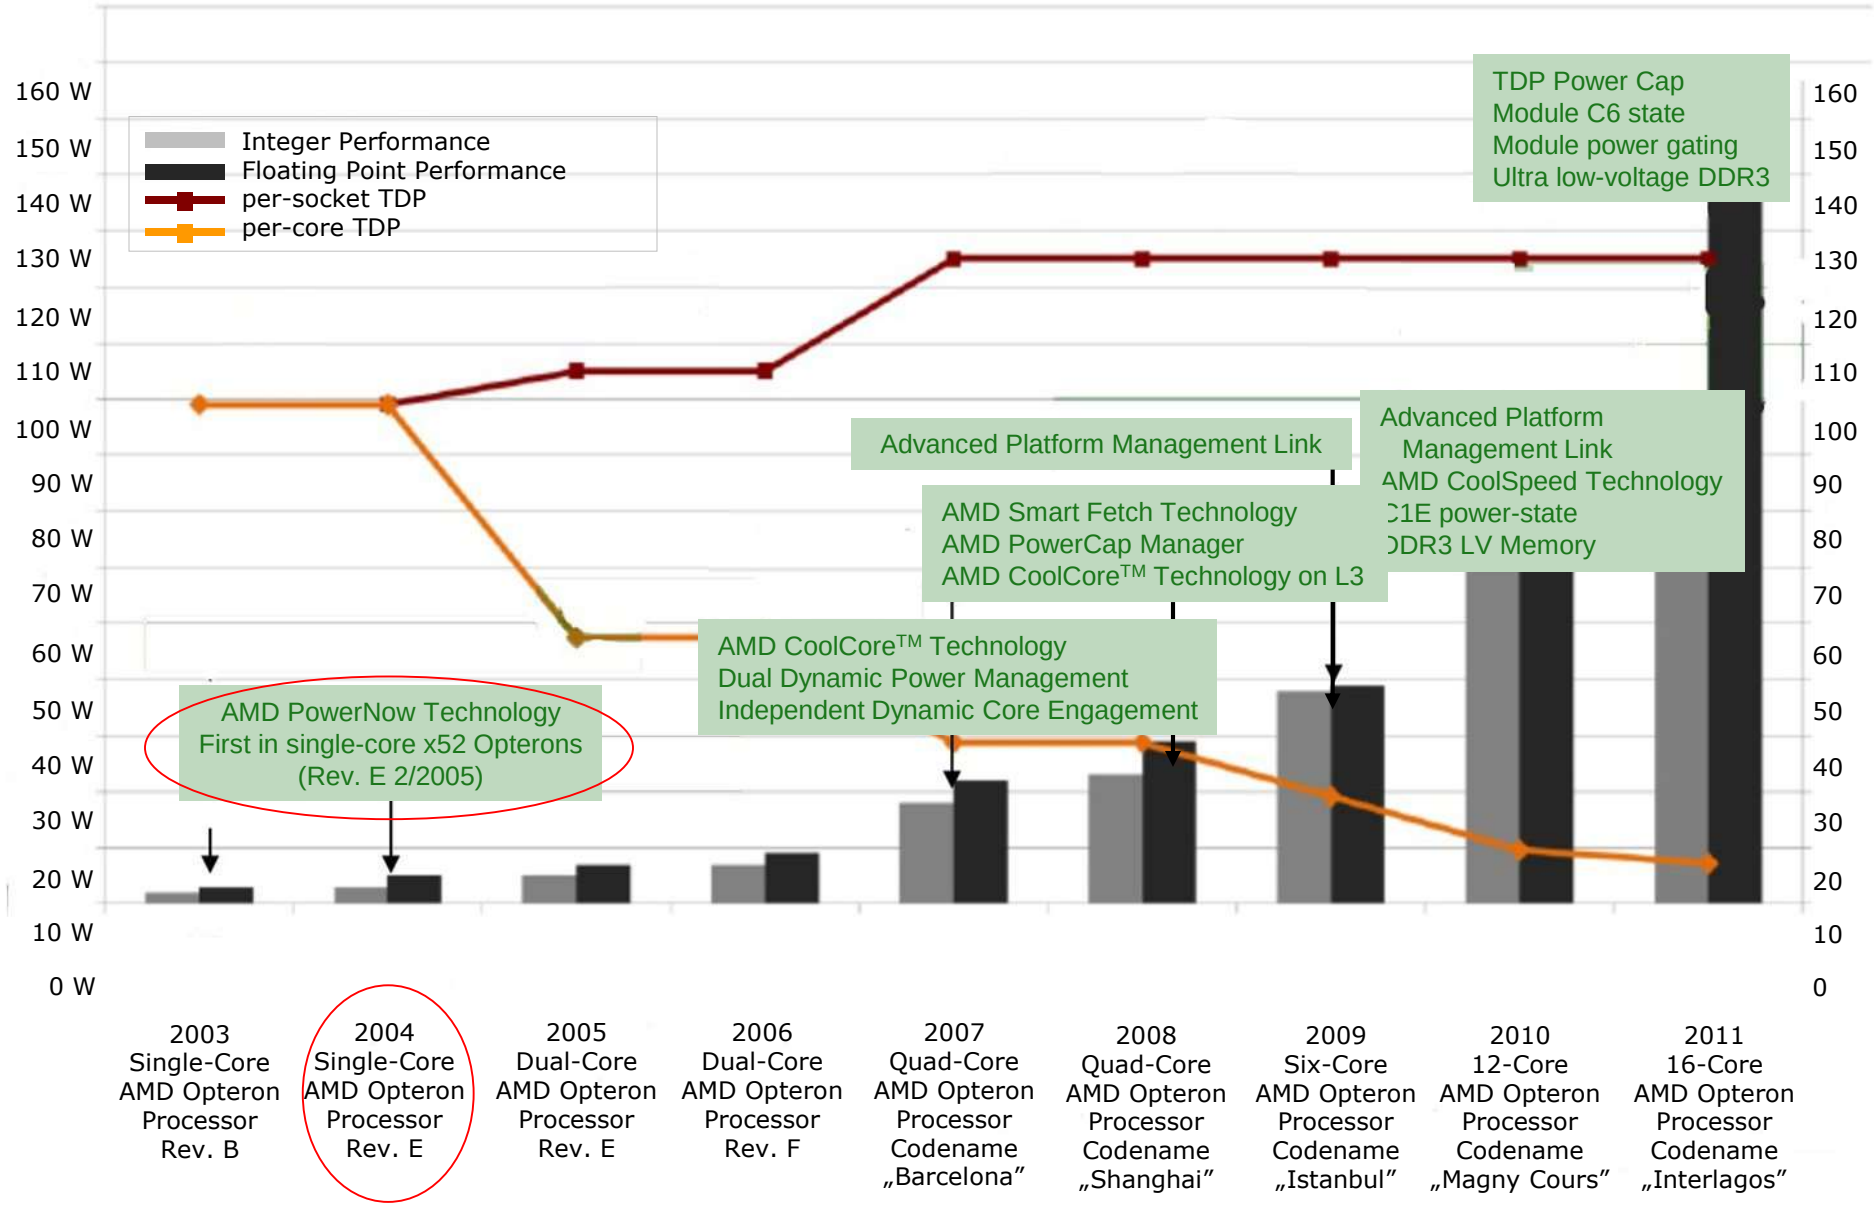

5.5 K8 server lines (9)

Main features of AMD's 90 nm DC K8-based Opteron lines (without HE/SE/EE lines)

	4/05	5/05- 9/05	3/06	2/07
MP DP UP	**Opteron Egypt**	**Opteron Egypt Italy Denmark**	**Opteron Egypt Italy Denmark**	**Opteron Egypt Italy**
Based on	K8	K8	K8	K8
Revision	E1	E6	E6	E6
No of cores	2	2	2	2
Technology	90 nm	90 nm	90 nm	90 nm
Transistors	233 mtrs	233 mtrs	233 mtrs	233 mtrs
MP models	875 /870/ 865	880 - 865	885	890
fc	2.2 / 2.0/1.8 GHz	2.4 - 1.8 GHz	2.6 GHz	2.8 GHz
DP models		280 - 265	285	290
fc		2.4 - 1.8 GHz	2.6 GHz	2.8 GHz
UP models		180 - 165	185	
fc		2.4 - 1.8 GHz	2.6 GHz	
L2	2*1 M	2*1 M	2*1 M	2*1 M
L3	–	–	–	–
Memory (reg.)	DDR-400 144 bit	DDR-400 144 bit	DDR-400 144 bit	DDR-400 144 bit
HT	1000 MHz	1000 MHz	1000 MHz	1000 MHz
PM	PowerNow!	PowerNow!	PowerNow!	PowerNow!
TDP	95 W	95 / 110 W	95 / 110 W	95 / 110 W
NX	NX	NX	NX	NX
Virt. techn.	–	–	–	–
Vcc	1.35 / 1.3 V	1.35 / 1.3 V	1.35 / 1.3 V	1.35 / 1.3 V
Socket	940	UP: 939, DP/MP: 940	UP: 939, DP/MP: 940	DP/MP: 940
C-states CIE	C1, C2, C3	→	→	→
S-states	S1, S3-S5	→	→	→
HW throttling	HW Clock throttling	→	→	→

5.5 K8 server lines (10)

Main features of AMD's 90 nm DC K8-based (NPT) Opteron lines (without HE/SE/EE lines)

	8/06 ^	2/07 ^	8/07 ^
MP DP UP	**Opteron Santa Rosa Santa Rosa Santa Ana**	**Opteron Santa Rosa Santa Rosa Santa Ana**	**Opteron Santa Rosa Santa Rosa Santa Ana**
Based on	K8	K8	K8
Revision	F2/F3	F3	F3
No of cores	2	2	2
Technology	90 nm	90 nm	90 nm
Transistors	114 mtrs	114 mtrs	114 mtrs
MP models	8218 - 8212	8220	8222
fc	2.6- 2.0 GHz	2.8 GHz	3.0 GHz
DP models	2218 - 2210	2220	2222
fc	2.6- 1.8 GHz	2.8 GHz	3.0 GHz
UP models	1218 - 1210	1220	1222
fc	2.6- 1.8 GHz	2.8 GHz	3.0 GHz
L2	2*1 M	2*1 M	2*1 M
L3	–	–	–
Memory (reg.)	UP: DDR2-800 144 bit DP/MP: DDR2-677 144 bit	UP: DDR2-800 144 bit DP/MP: DDR2-677 144 bit	UP: DDR2-800 144 bit DP/MP: DDR2-677 144 bit
HT	1000 MHz	1000 MHz	1000 MHz
PM	PowerNow!	PowerNow!	PowerNow!
TDP	95 / 103 W	95 / 103 W	95 / 103 W
NX	NX	NX	NX
Virt. techn.	AMD-V	AMD-V	AMD-V
Vcc	1.35/1.3 V	1.35 / 1.3 V	1.35 / 1.3 V
Socket	UP: AM2, DP/MP: F	UP: AM2, DP/MP: F	UP: AM2, DP/MP: F
C-states	C1, C2, C3 →		
CIE			
S-states	S1, S3-S5 →		
HW throttling	HW Clock throttling →		

Example: Basic architecture of dual core Opteron processors [41]

UP: Opteron 100/1000

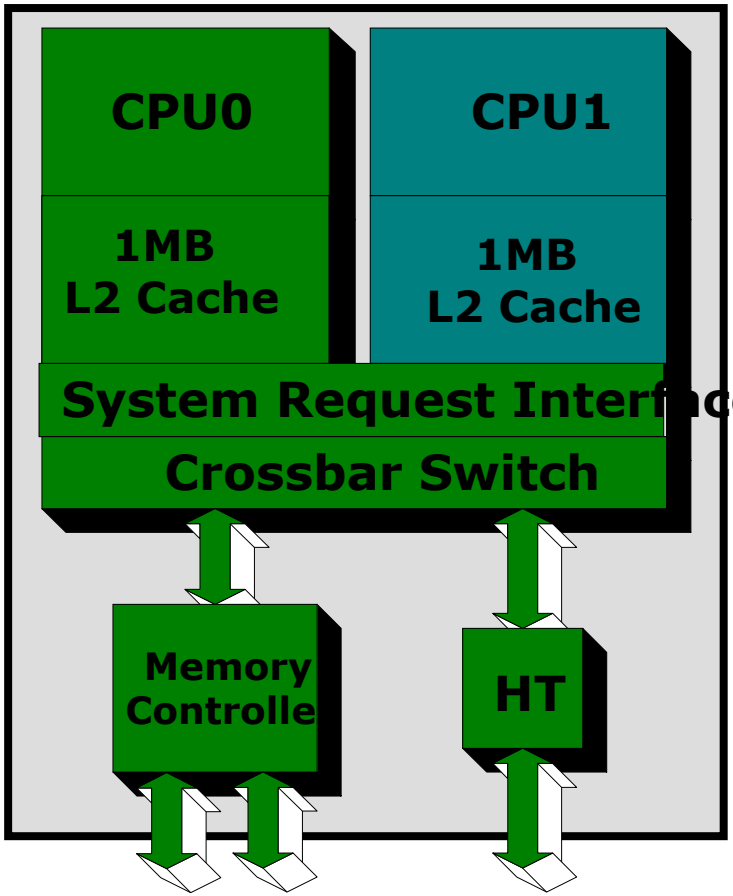

2 x 72 bit

DP: Opteron 200/2000, MP: 800/8000

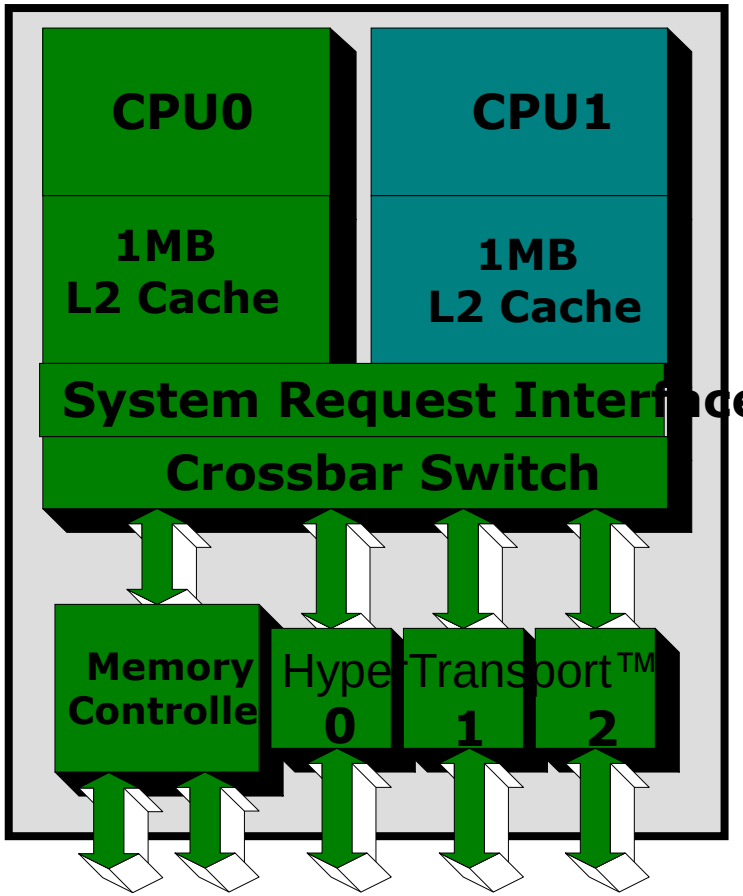

2 x 72 bit

800/8000: 3 coherent links
200/2000: 1 coherent link

5.6 K8 desktop lines

5.6 K8 desktop lines

Overview of AMD's K8 desktop lines

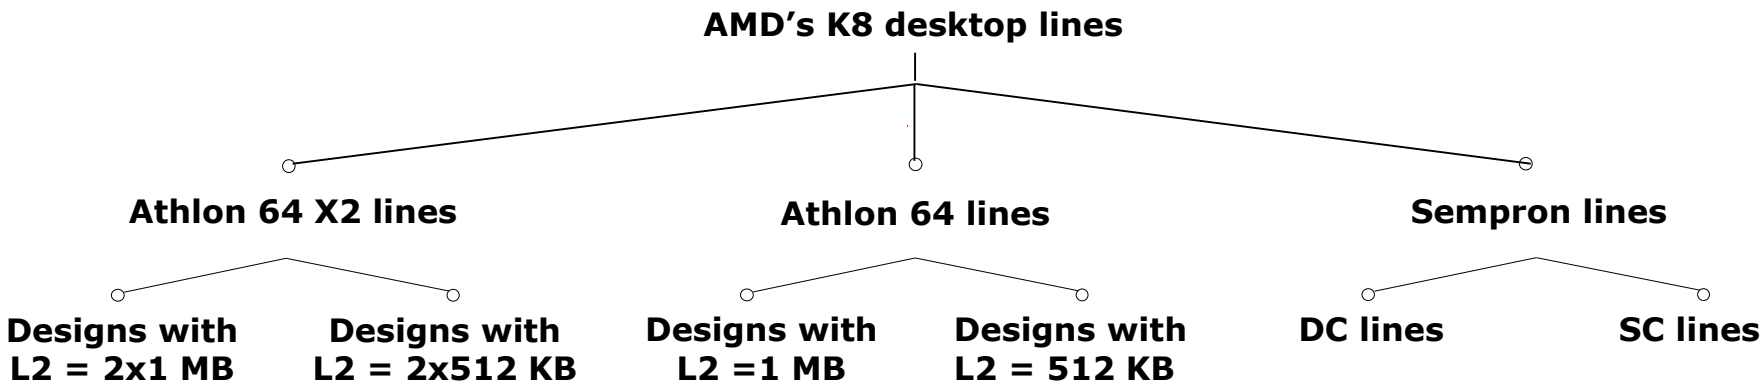

No native designs.
Processor designs are the same as the 2x512 KB L2 Athlon 64 X2 and 512 Kb L2 Athlon 64 designs and are partly disabled.

Remark

The subsequent overview of AMD's K8-based native desktop lines does not reflect models with disabled functionality.

5.6 K8 desktop lines (2)

Overview of AMD's K8-based desktop lines (except the Brisbane line) [13]

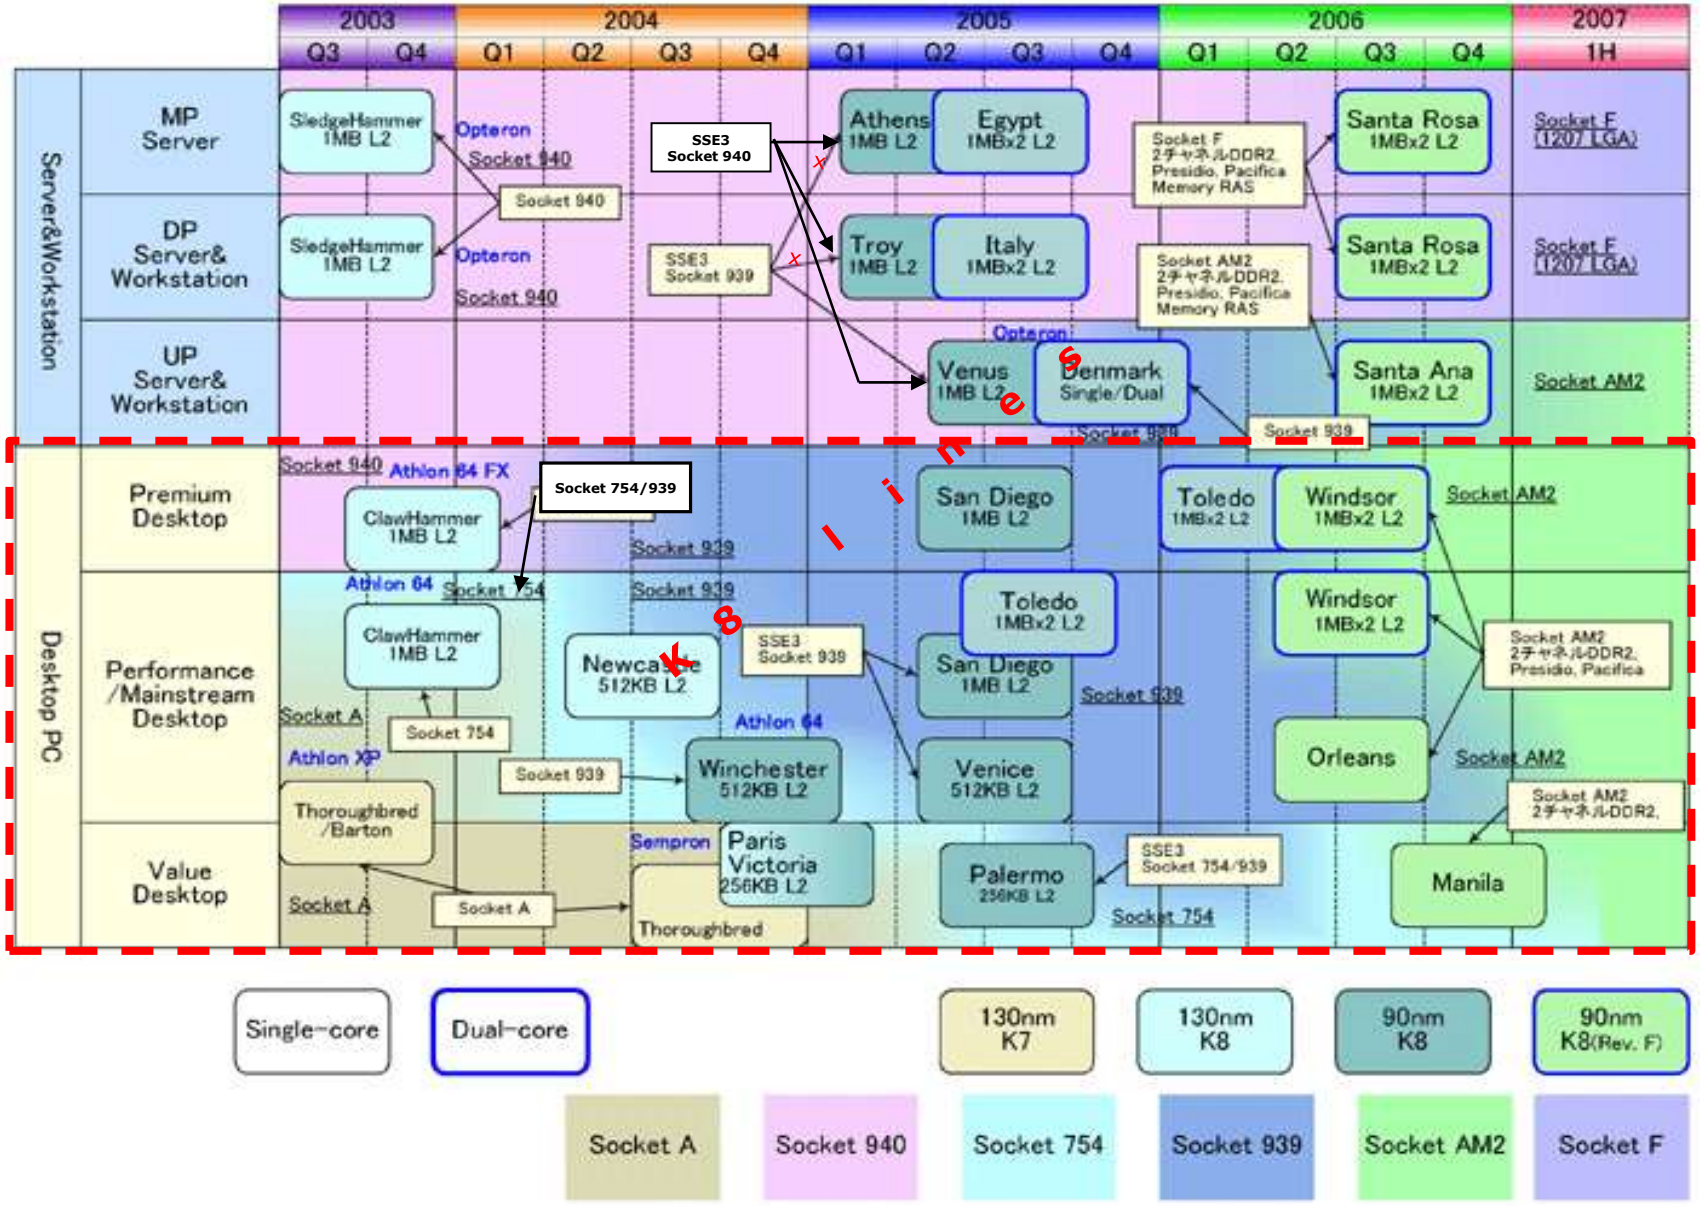

5.6 K8 desktop lines (3)

Overview of AMD's K8 desktop lines – the Brisbane line [14]

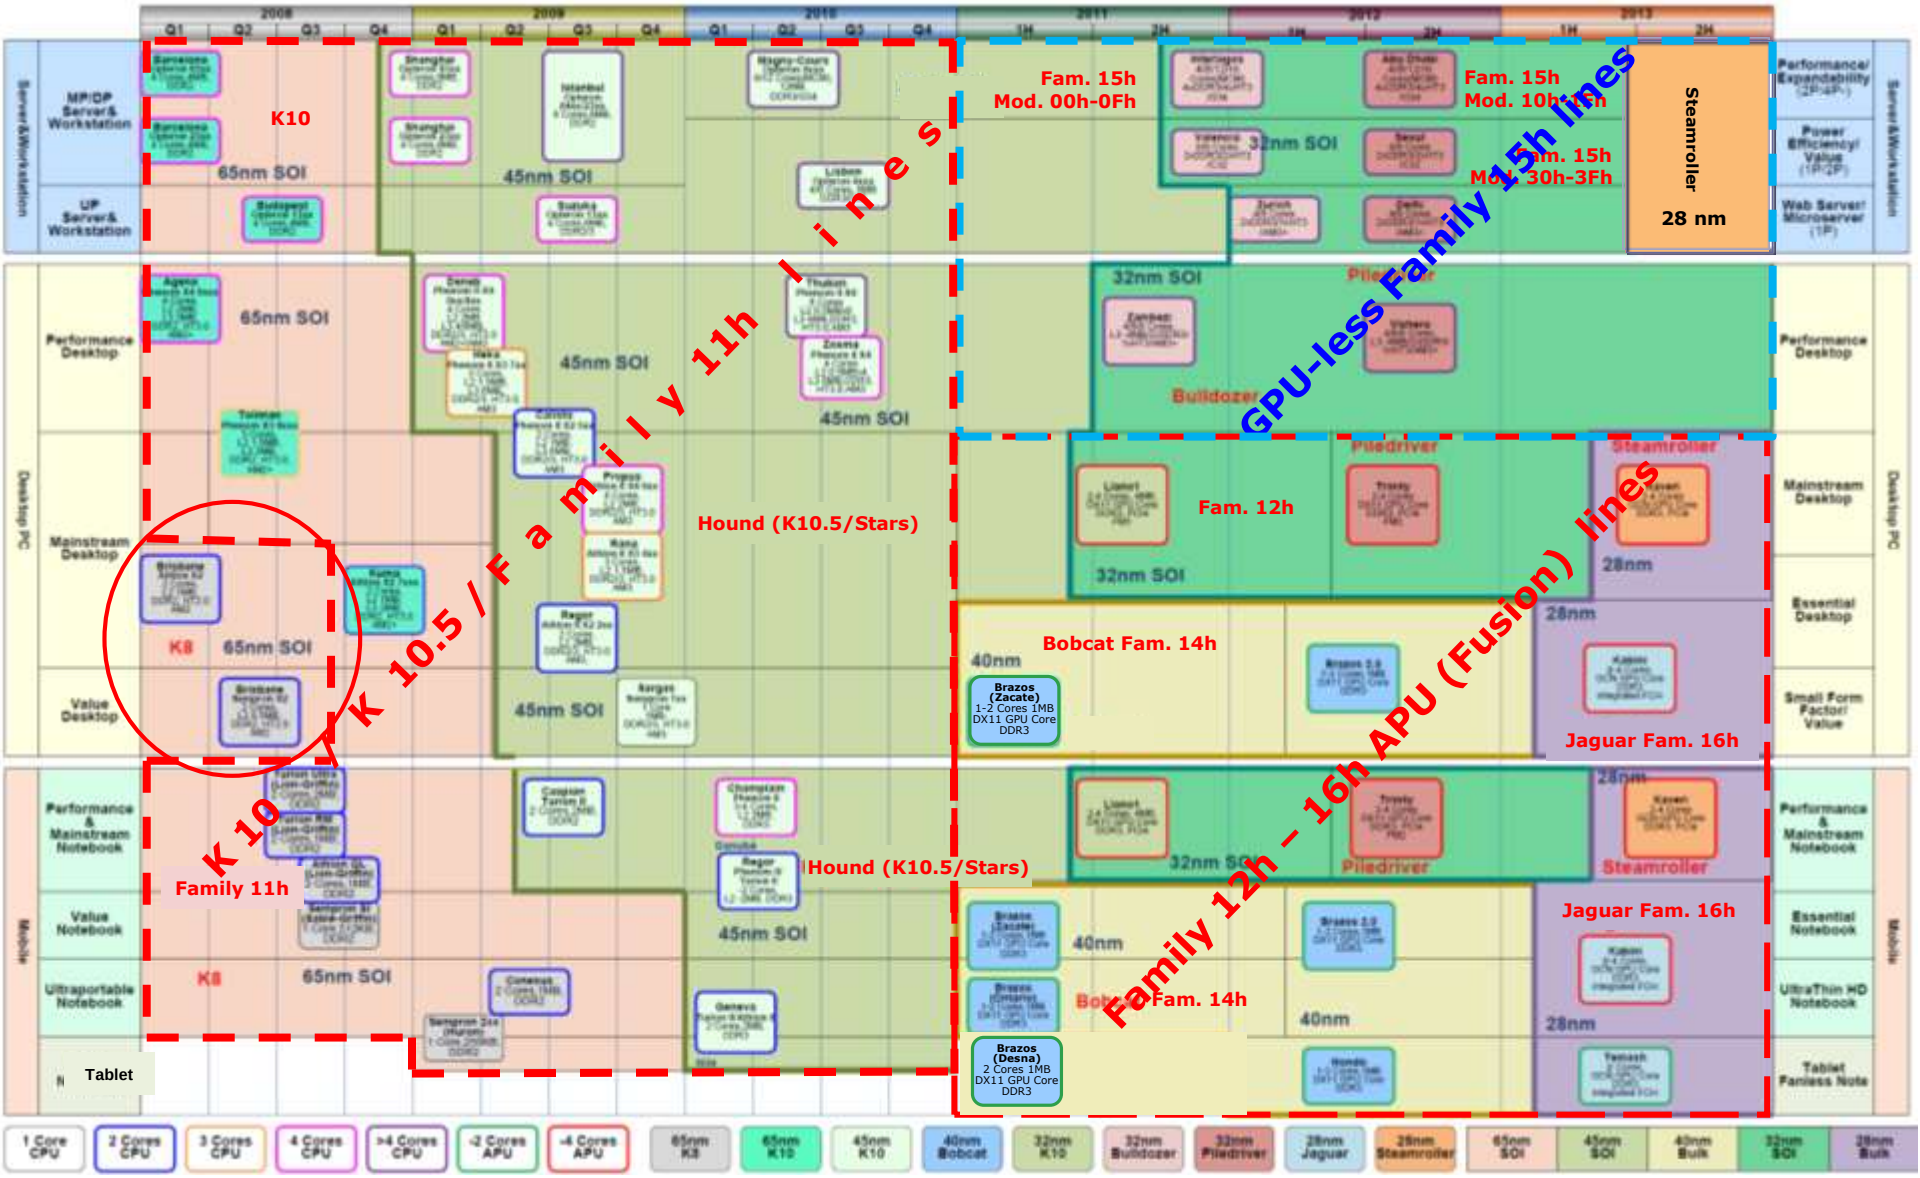

5.6 K8 desktop lines (4)

Overview of AMD's native K8 desktop cores (Data based on [134])

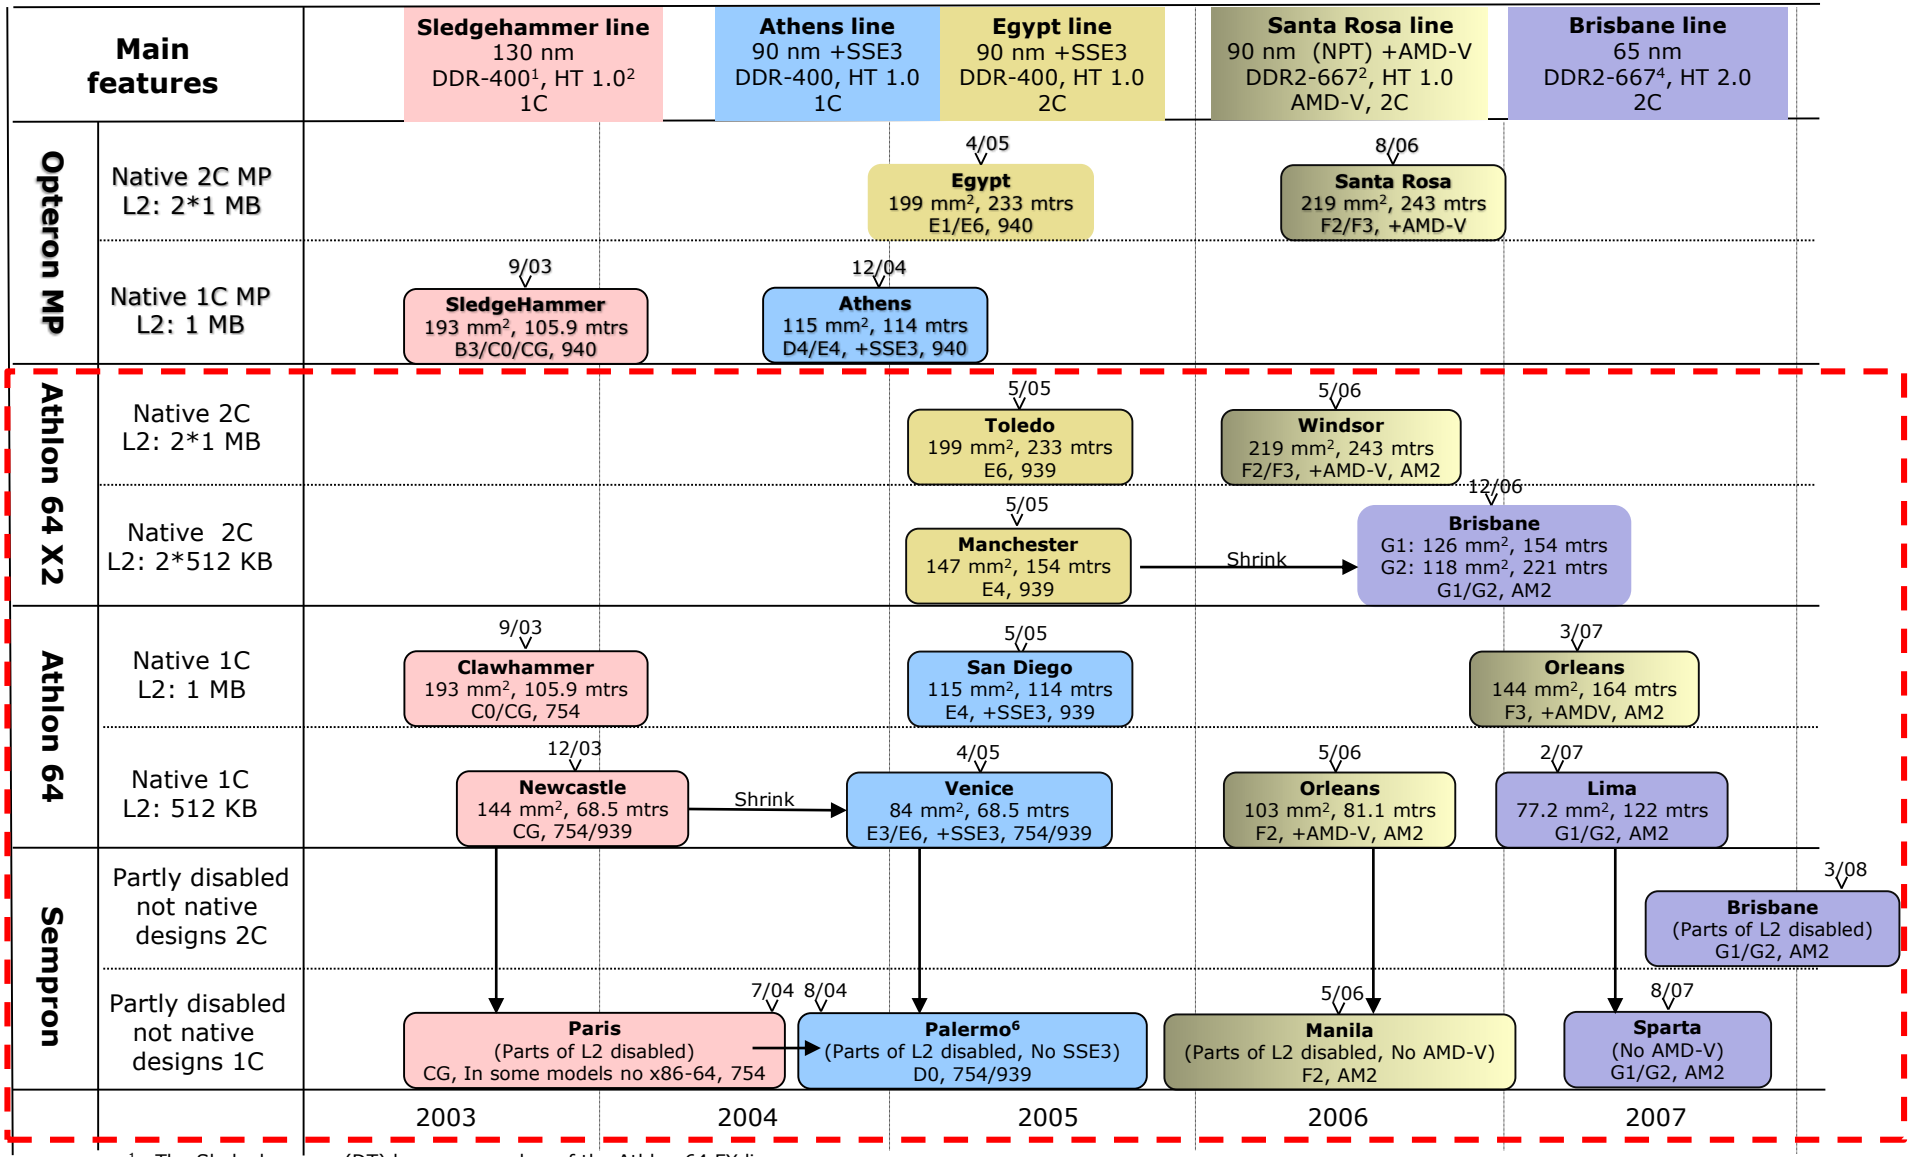

[1] : The Sledgehammer (DT) became member of the Athlon 64 FX line
Subsequent members of the Athlon FX line are the same processors
as the highest frequency Athlon 64/Athlon 64 FX lines

[2] : G1: 219 mm², 243 mtrs; G2: 118 mm², 221 mtrs

5.6 K8 desktop lines (5)

Sockets of K8 desktop processors [based on 42]

Supports DDR3

Socket **G34** (1944 pins, MP servers K10.5 „Magny-Course", Fam. 15h)
Socket **C32** (1207 pins, DP servers K10.5 „Magny –Course, Fam. 15h)
Socket **AM3** (941 pins, K10.5 Shanghai UP servers, desktops, except first released desktops, K10.5 Istambul desktops)
Socket **AM3+** (942 pins, UP servers Fam. 15h, desktops Fam. 15h)
Socket **FM1** (905 balls BGA, desktops, mobiles Fam. 12h)
Socket **S1g4** (638 pins, K10.5 mobile)
Socket **FT1** (413 balls BGA, desktops, mobiles Fam. 14h)
Socket **FT3** (769 balls BGA, desktops, mobiles Fam. 16h)

S:: Servers
DT: Desktop
M: Mobile
UPT: Ultraportable

SPP: Single power Plane
DPP: Dual Power planes
HT 2.0: HyperTransport 2.0
HT 3.0: HyperTransport 3.0

K10.5 | Fam. 15h | Fam. 15h | Fam. 15h

45nm 2x6C „Magny-Course" | 32 nm 2x8C „Bulldozer" | 32 nm 2x8C/4C APU „Piledriver" | 28 nm 2x8C „Steamroller"

3st Generation Platform

45 nm 4C, DT/M „Shanghai" | 45 nm 6C, DT „Istambul" | 32 nm 4C APU „Llano" | 32 nm 2C APU „Bobcat" | 28 nm 4C APU „Jaguar" | 28 nm 4C APU „Puma+"

Fam. 12h | Fam. 14h | Fam. 16h | Fam. 16h

K8 (NPT) | K10 | K10.5 | K10.5

90 nm 2C NPT „Santa Rosa" | 65 nm 4C „Barcelona" | 45 nm 4C, S „Shanghai" | 45 nm 6C, S „Istambul"

2st Generation Platform

65 nm 4C „Brisbana" | 65 nm 2C „Griffin"

K8 | Fam. 11h

Supports DDR2

Socket **F** (1207 pins, SPP, HT 2.0: DP, MP servers K8 NPT, K10)
Socket **F+** (1207 pins, DPP, HT 3.0: DP, MP servers K10.5)
Socket **AM2** (940 pins, SPP, HT 2.0: K8 NPT UP serv., K8 NPT, Brisbane DTs)
Socket **AM2+** (940 pins, DPP: K10 UP servers, K10 desktops, first K10.5 Shanghai desktops)
Socket **S1g1** (638 pins, K8 NPT, Brisbane mobiles)
Socket **S1g2** (638 pins, Fam. 11h mobiles)
Socket **S1g3** (638 pins, K10.5 Shanghai (Caspian die) mobiles)
Socket **ASB1** (812 balls BGA, Fam. 11h (Huron, Conesus dies) UPT mobiles)
Socket **ASB2** (812 balls BGA, K10.5 (Shanghai) (Geneva die), UPT mobiles)

K8 | K8 | K8

130 nm 1C „SledgeH." | 90 nm 1C „Athens" | 90 nm 2C „Egypt"

1st Generation Platform

Supports DDR

Socket **940** (130nm K8 UP/DP/MP servers, 90 nm K8 SC and 90 nm K8 DC (Egypt) DP/MP servers)
Socket **939** (90nm K8 SC and 90 nm K8 DC (Egypt) UP servers, K8 desktops with 128-bit mem. channel up to the 90 nm DC (Egypt) lines)
Socket **754** (K8 desktops, mobiles with 64-bit mem. channel up to the 90 nm SC (Athens) line)

2003 2004 2005 2006 2007 2008 2009 2010 2011 2012 2013 2014

5.6 K8 desktop lines (6)

Main features of AMD's high-performance desktop 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. 15h based lines		-	-	-	-	-	-	-	-	-	-

5.6 K8 desktop lines (7)

Block diagram of AMD's 90 nm Athlon 64 X2 DC (Egypt) line supporting DDR memory [147]

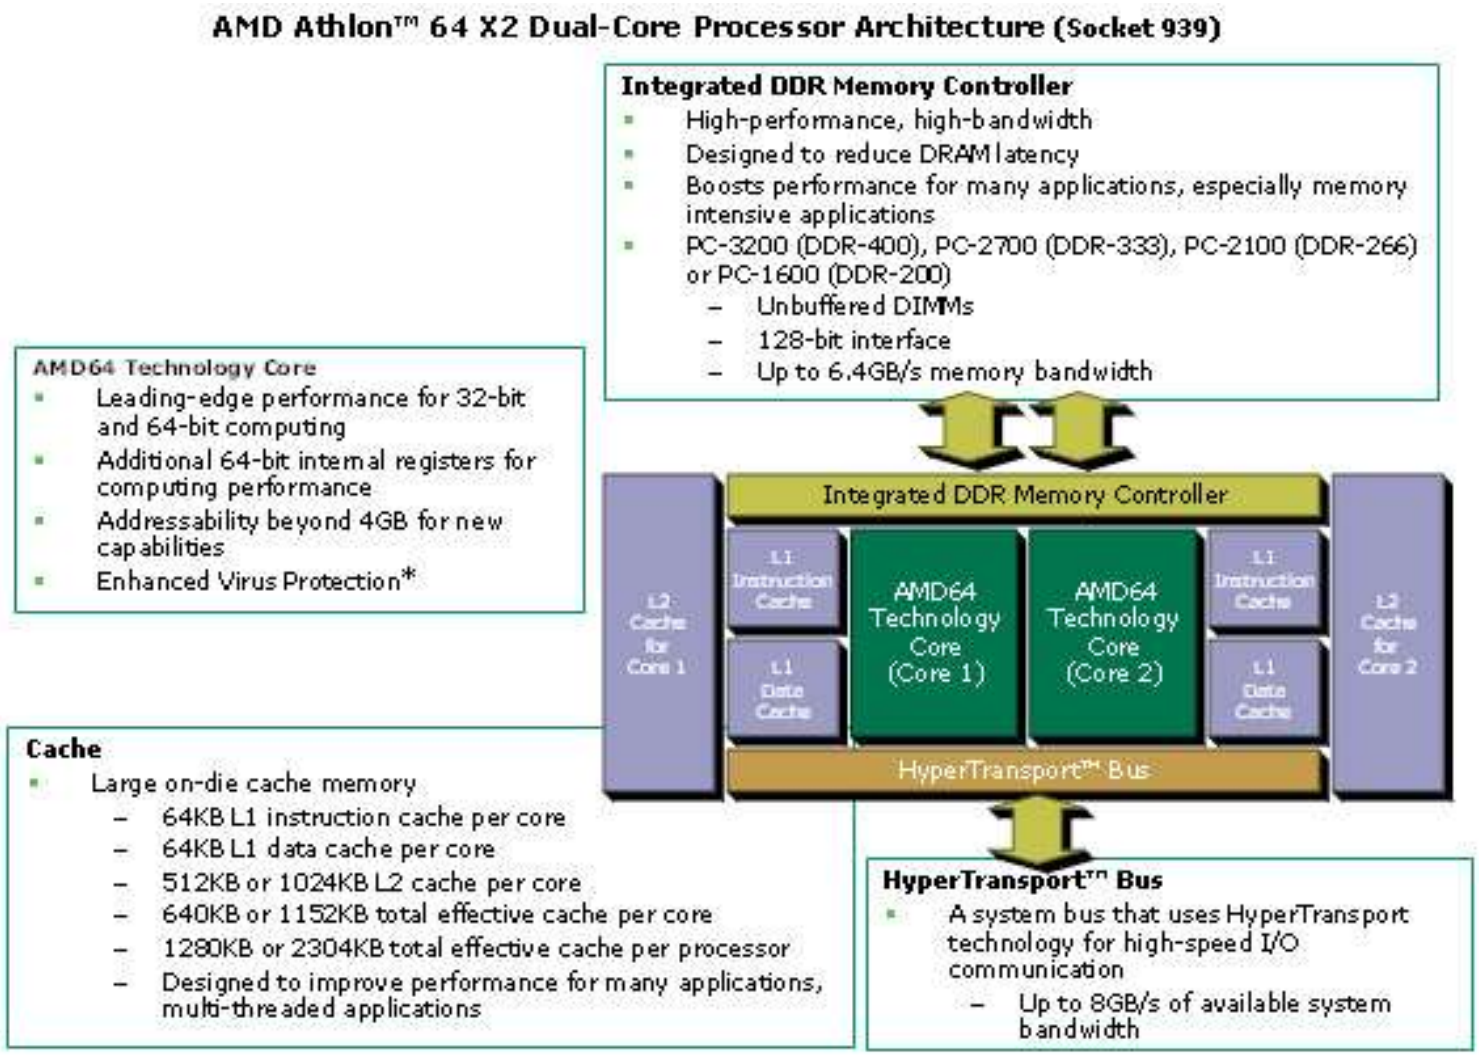

5.6 K8 desktop lines (8)

Block diagram of AMD's 90 nm Athlon 64 X2 DC (NPT) line supporting DDR2 memory [147]

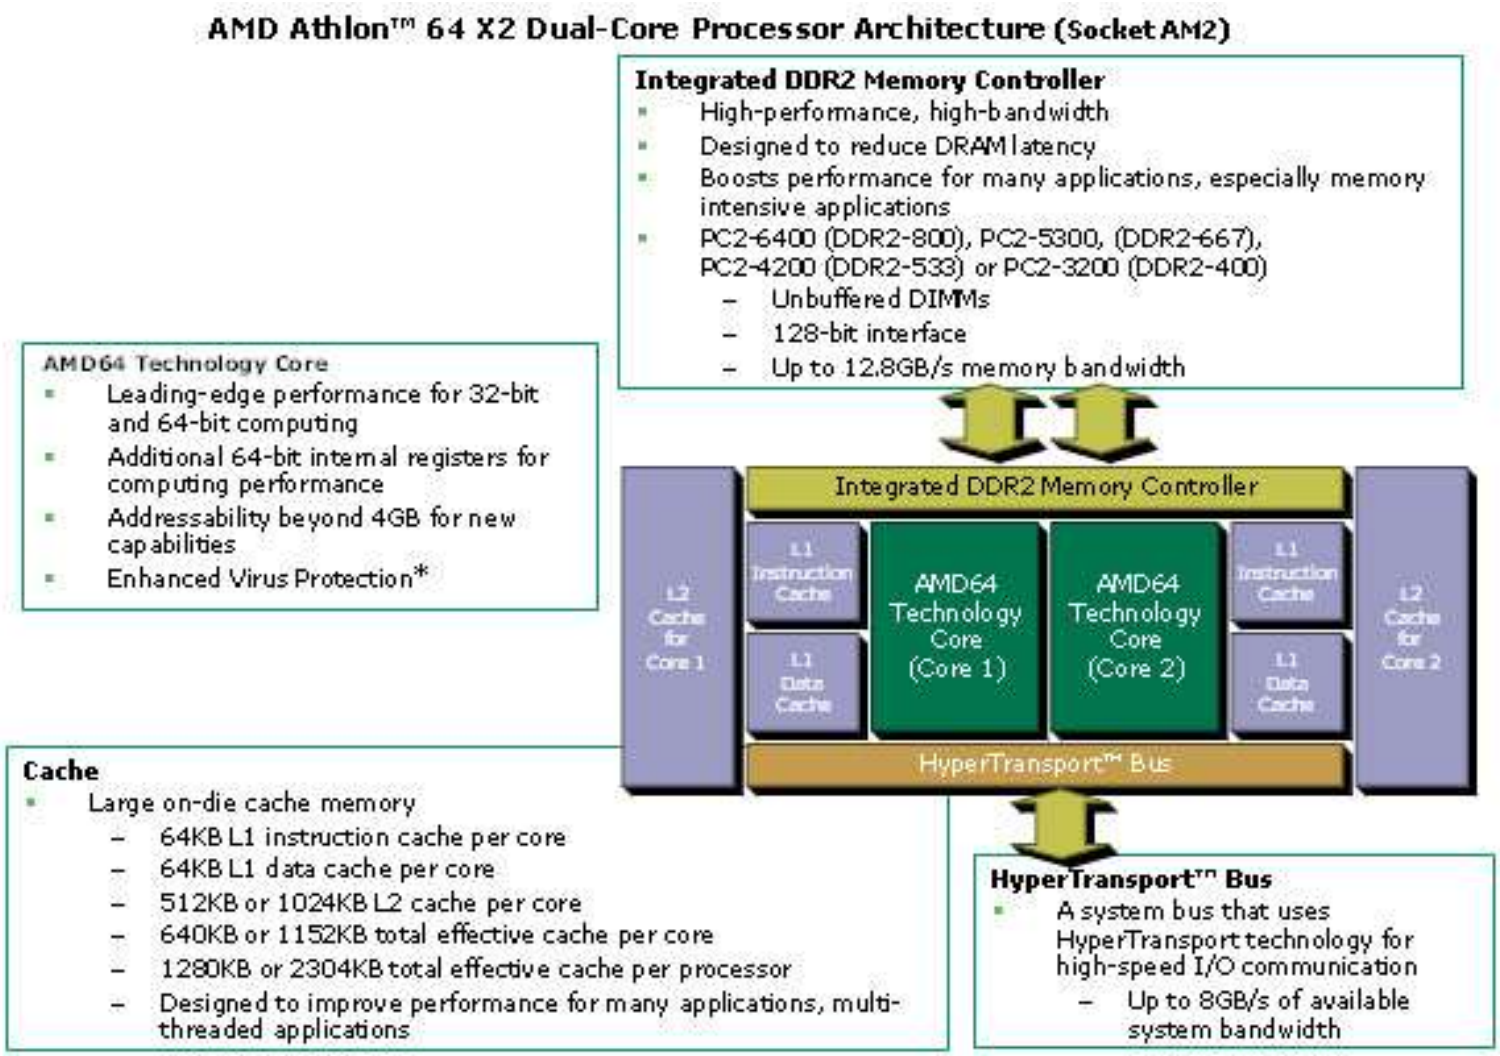

5.7 K8 mobile lines

5.7 K8 mobile lines

Overview of AMD's K8 mobile lines

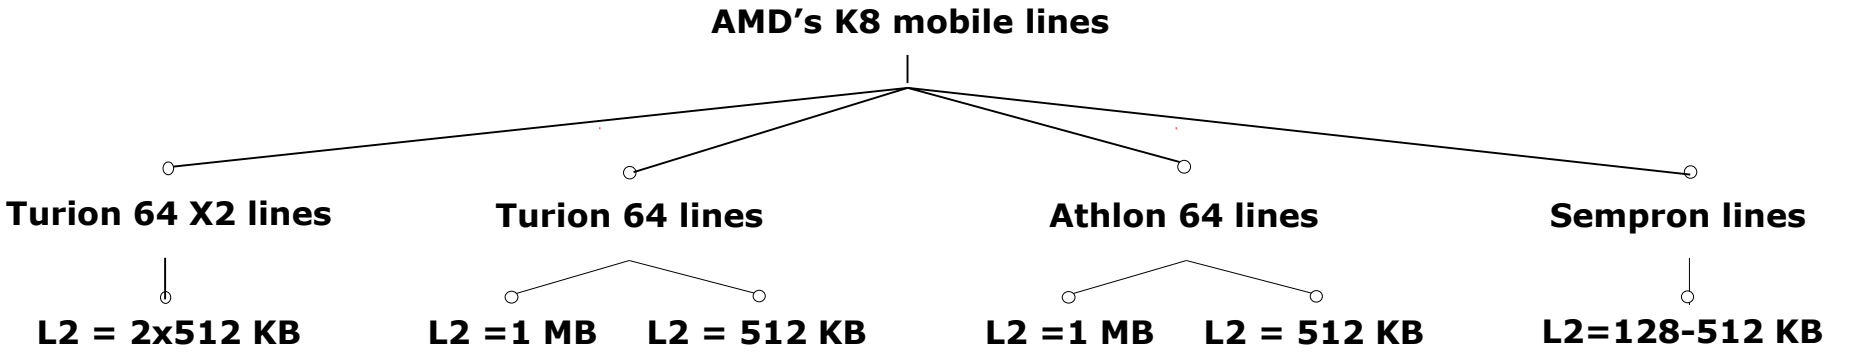

No native designs but Turion 64 or Athlon 64 designs with partly disabled L2

Remark

The subsequent overview of AMD's K8-based native mobile lines does not reflect models with disabled functionality, such as the Athlon 64 X2 (TK-4x/5x) lines with 2x512 or 2x 256 KB L2.

5.7 K8 mobile lines (2)

Overview of AMD's K8-based mobile lines (Data based on [7])

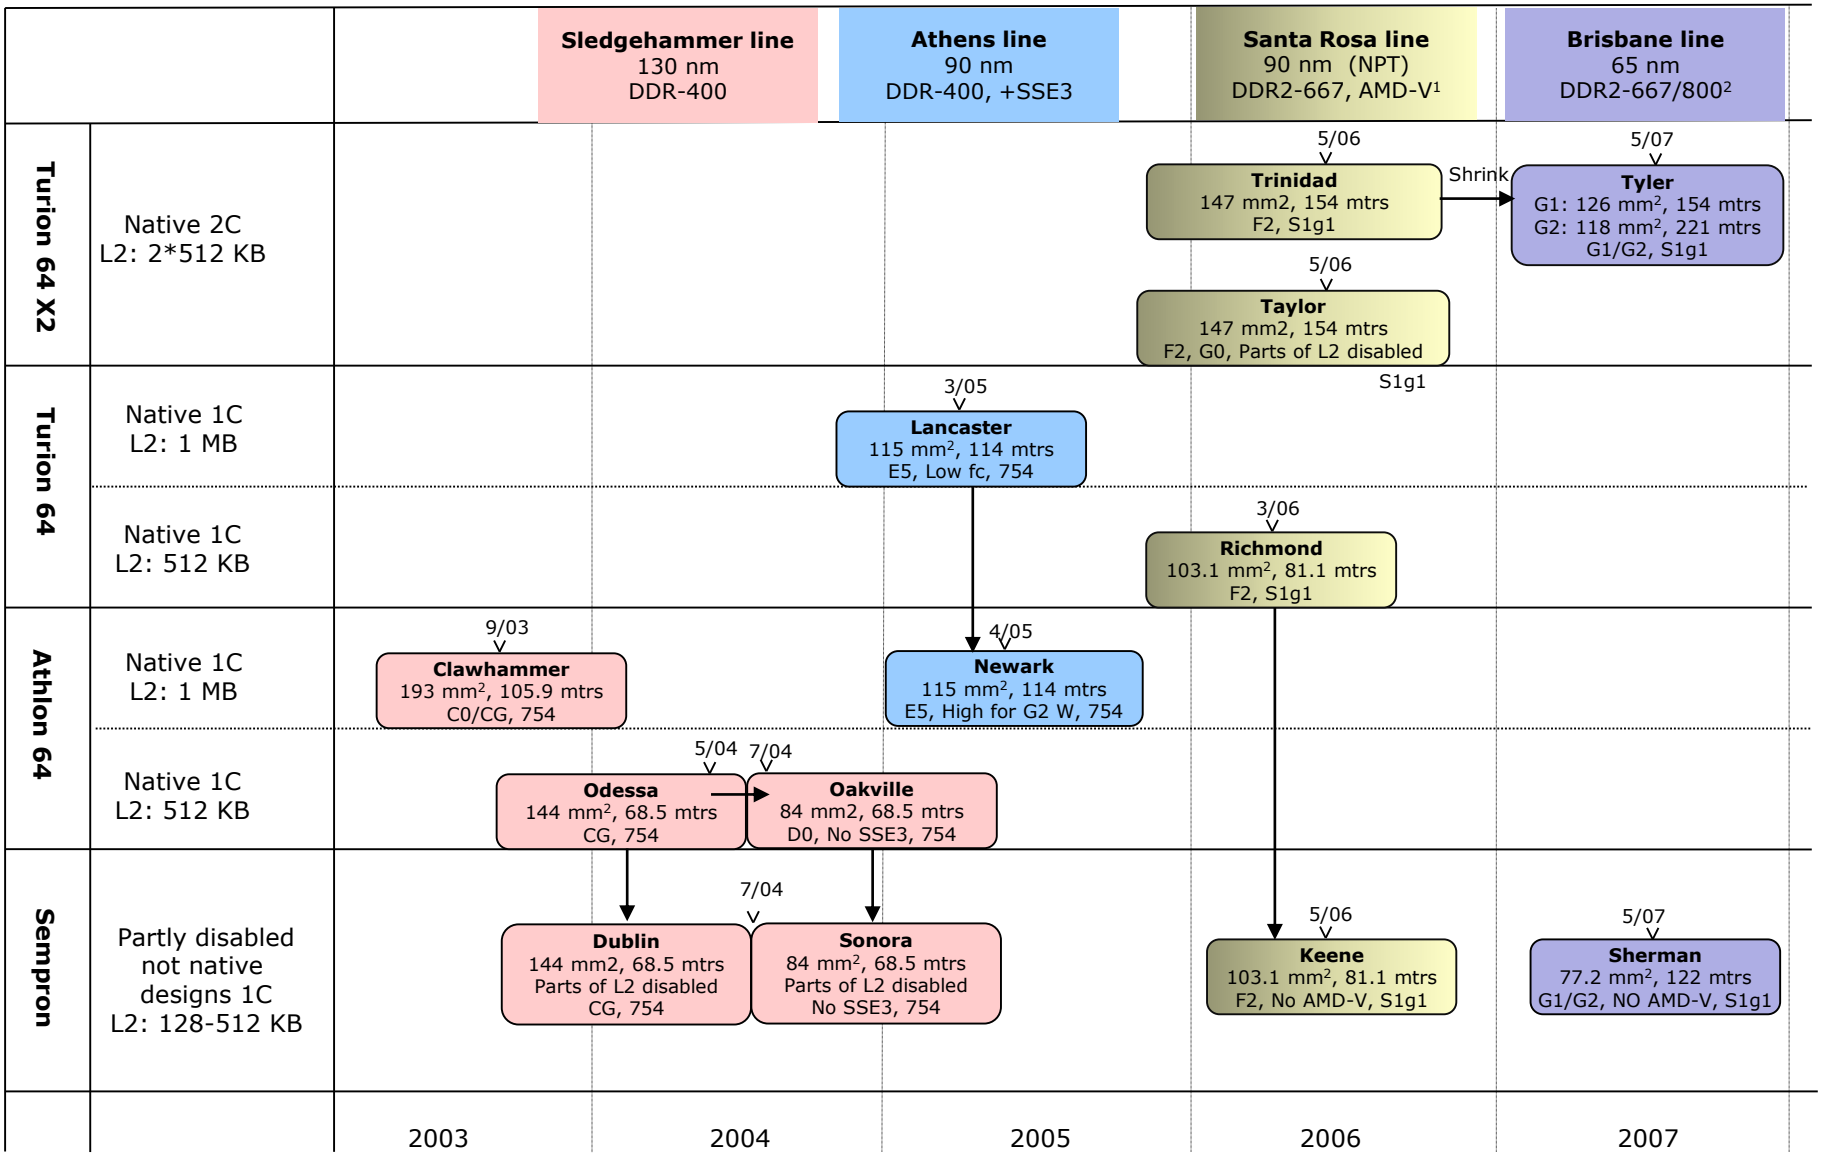

5.7 K8 mobile lines (3)

Main features of AMD's high performance K8 (Hammer)-based mobile lines

Base arch./ stepping		Intro	High perf. mobile family name	Series	Techn.	Core count (up to)	L2 (up to)	L3	Memory (up to)	HT/ dir. (up to)	Sock et
K8	C0, CG	9/2003	Clawhammer	Mobile Athlon 64	130 nm	1	512 KB	-	DDR-400	HT 1.0: 3.2 GB/s	754
	E5	3/2005	Lancaster	Turion 64	90 nm	1	1 MB	-	DDR-400	HT 1.0: 3.2 GB/s	754
	F2	5/2006	Trinidad	Turion 64 X2	90 nm	2	2*512 KB	-	DDR2-667	HT 1.0: 3.2 GB/s	S1
K10	-	-	-	-	-	-	-	-	-	-	-
K10.5	DA-C2	9/2009	Caspian	Turion II	45 nm	2	2*512 KB/ 2*1 MB[1]	-	DDR2-800	HT 3.0: 7.2 GB/s	S1g3
	DA-C3	5/2010	Champlain	Turion X4	45 nm	4	4*512 KB	-	DDR3-1066	HT 3.0: 7.2 GB/s	S1g4

[1]: 2*512 KB for Turion II, 2*1 MB for Turion II Ultra

5.7 K8 mobile lines (4)

Sockets of AMD's x86-64 processors [42]

Note

Sockets (termed by AMD as platforms).

They determine the memory type supported, as indicated in the next Figure.

5.7 K8 mobile lines (5)

Sockets of K8 mobile processors [based on 42]

Supports DDR3

Socket **G34** (1944 pins, MP servers K10.5 „Magny-Course", Fam. 15h)
Socket **C32** (1207 pins, DP servers K10.5 „Magny –Course, Fam. 15h)
Socket **AM3** (941 pins, K10.5 Shanghai UP servers, desktops, except first released desktops, K10.5 Istambul desktops)
Socket **AM3+** (942 pins, UP servers Fam. 15h, desktops Fam. 15h)
Socket **FM1** (905 balls BGA, desktops, mobiles Fam. 12h)
Socket **S1g4** (638 pins, K10.5 mobile)
Socket **FT1** (413 balls BGA, desktops, mobiles Fam. 14h)
Socket **FT3** (769 balls BGA, desktops, mobiles Fam. 16h)

S:: Servers
DT: Desktop
M: Mobile
UPT: Ultraportable

SPP: Single power Plane
DPP: Dual Power planes
HT 2.0: HyperTransport 2.0
HT 3.0: HyperTransport 3.0

K10.5 | Fam. 15h | Fam. 15h | Fam. 15h

45nm 2x6C „Magny-Course" | 32 nm 2x8C „Bulldozer" | 32 nm 2x8C/4C APU „Piledriver" | 28 nm 2x8C „Steamroller"

3st Generation Platform

45 nm 4C, DT/M „Shanghai" | 45 nm 6C, DT „Istambul" | 32 nm 4C APU „Llano" | 32 nm 2C APU „Bobcat" | 28 nm 4C APU „Jaguar" | 28 nm 4C APU „Puma+"

Fam. 12h | Fam. 14h | Fam. 16h | Fam. 16h

K8 (NPT) | K10 | K10.5 | K10.5

90 nm 2C NPT „Santa Rosa" | 65 nm 4C „Barcelona" | 45 nm 4C, S „Shanghai" | 45 nm 6C, S „Istambul"

2st Generation Platform

65 nm 4C „Brisbana" | 65 nm 2C „Griffin"

K8 | Fam. 11h

Supports DDR2

Socket **F** (1207 pins, SPP, HT 2.0: DP, MP servers K8 NPT, K10)
Socket **F+** (1207 pins, DPP, HT 3.0: DP, MP servers K10.5)
Socket **AM2** (940 pins, SPP, HT 2.0: K8 NPT UP serv., K8 NPT, Brisbane DTs)
Socket **AM2+** (940 pins, DPP: K10 UP servers, K10 desktops, first K10.5 Shanghai desktops)
Socket **S1g1** (638 pins, K8 NPT, Brisbane mobiles)
Socket **S1g2** (638 pins, Fam. 11h mobiles)
Socket **S1g3** (638 pins, K10.5 Shanghai (Caspian die) mobiles)
Socket **ASB1** (812 balls BGA, Fam. 11h (Huron, Conesus dies) UPT mobiles)
Socket **ASB2** (812 balls BGA, K10.5 (Shanghai) (Geneva die), UPT mobiles)

K8 | K8 | K8

130 nm 1C „SledgeH." | 90 nm 1C „Athens" | 90 nm 2C „Egypt"

1st Generation Platform

Supports DDR

Socket **940** (130nm K8 UP/DP/MP servers, 90 nm K8 SC and 90 nm K8 DC (Egypt) DP/MP servers)
Socket **939** (90nm K8 SC and 90 nm K8 DC (Egypt) UP servers, K8 desktops with 128-bit mem. channel up to the 90 nm DC (Egypt) lines)
Socket **754** (K8 desktops, mobiles with 64-bit mem. channel up to the 90 nm SC (Athens) line)

2003 | 2004 | 2005 | 2006 | 2007 | 2008 | 2009 | 2010 | 2011 | 2012 | 2013 | 2014

5.8 Evolution path of AMD’s K8-based Athlon 64/64 X2 desktop cores

5.8 Evolution path of AMD's K8-based Athlon 64/64 X2 desktop cores (1)

Evolution path of AMD's K8 based Athlon 64/64 X2 desktop cores (Data based on [134])

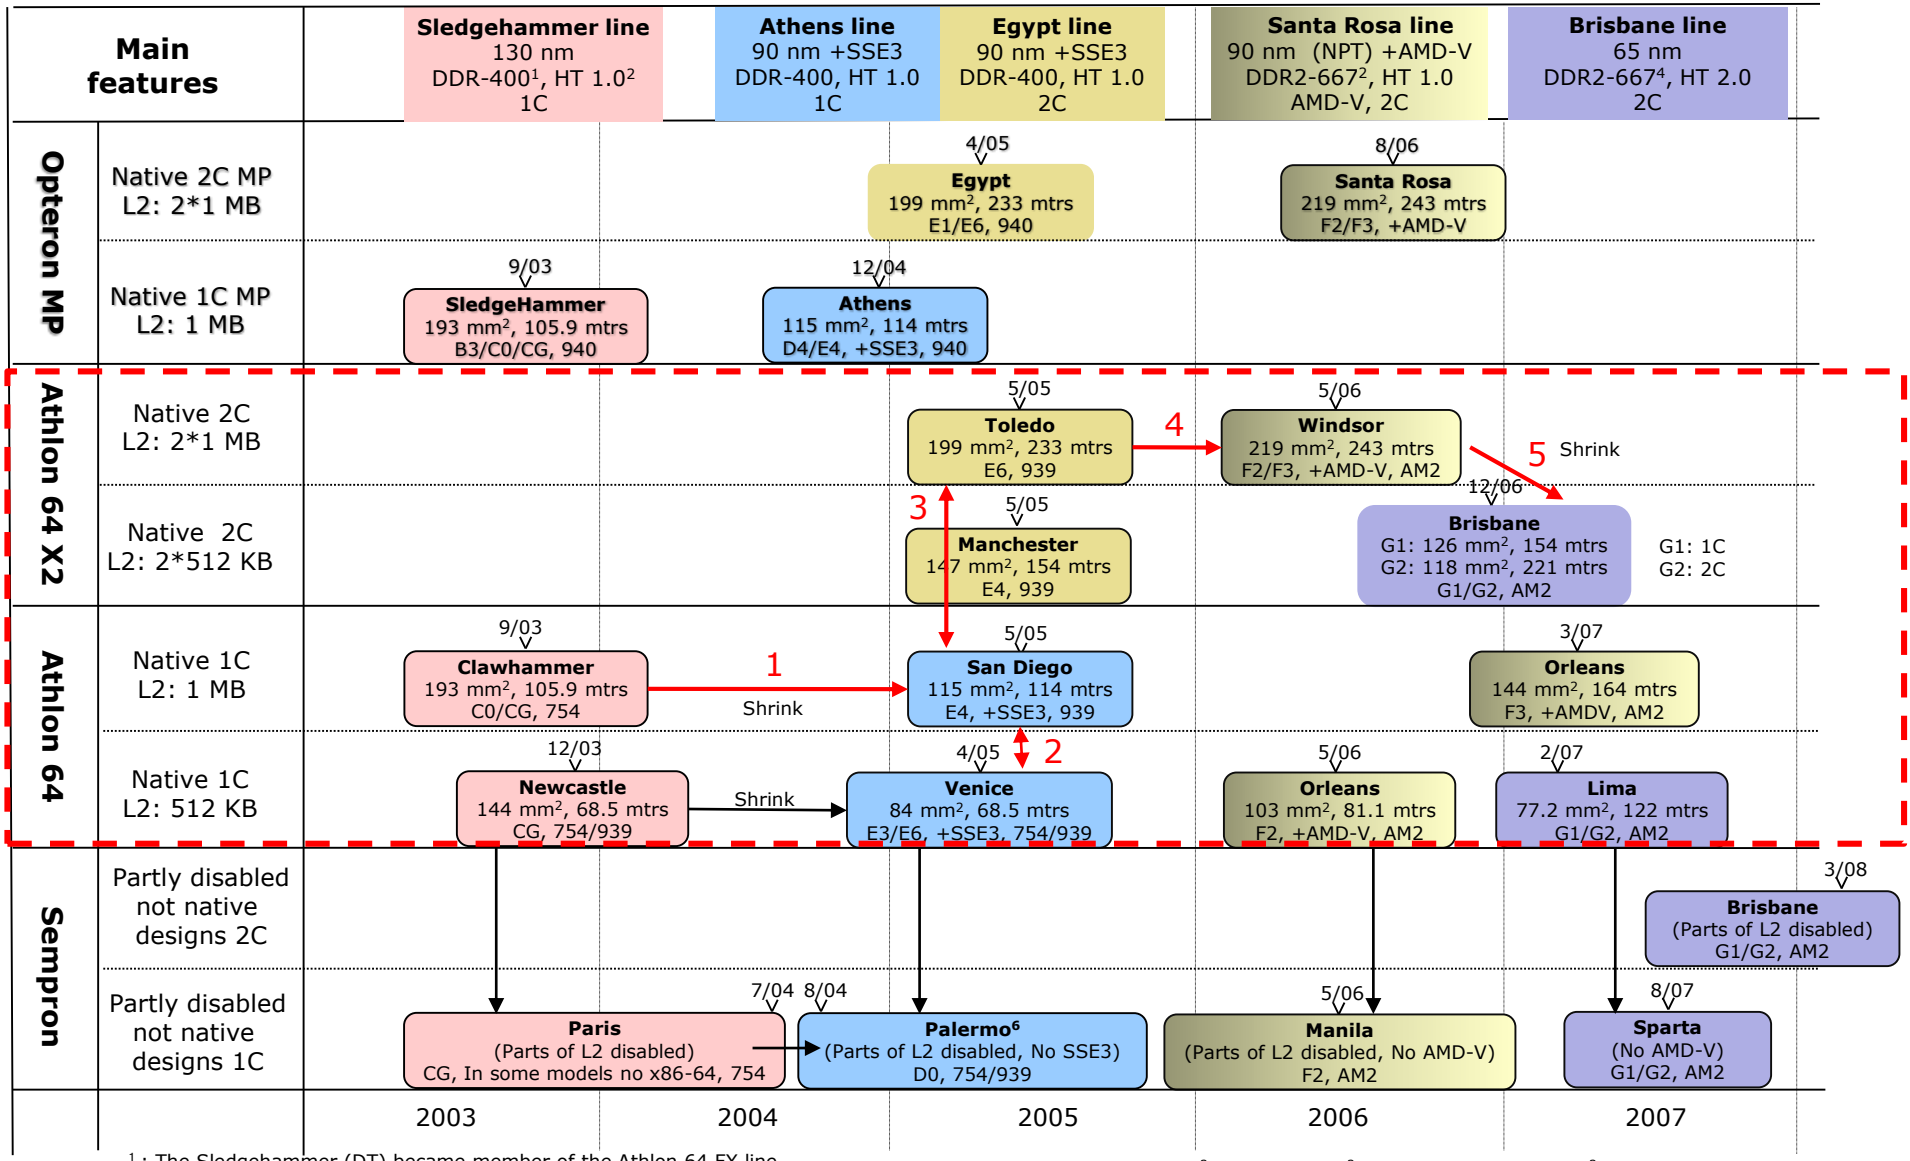

[1] : The Sledgehammer (DT) became member of the Athlon 64 FX line
Subsequent members of the Athlon FX line are the same processors
as the highest frequency Athlon 64/Athlon 64 FX lines

[2] : G1: 219 mm², 243 mtrs; G2: 118 mm², 221 mtrs

5.8 Evolution path of AMD's K8-based Athlon 64/64 X2 desktop cores (2)

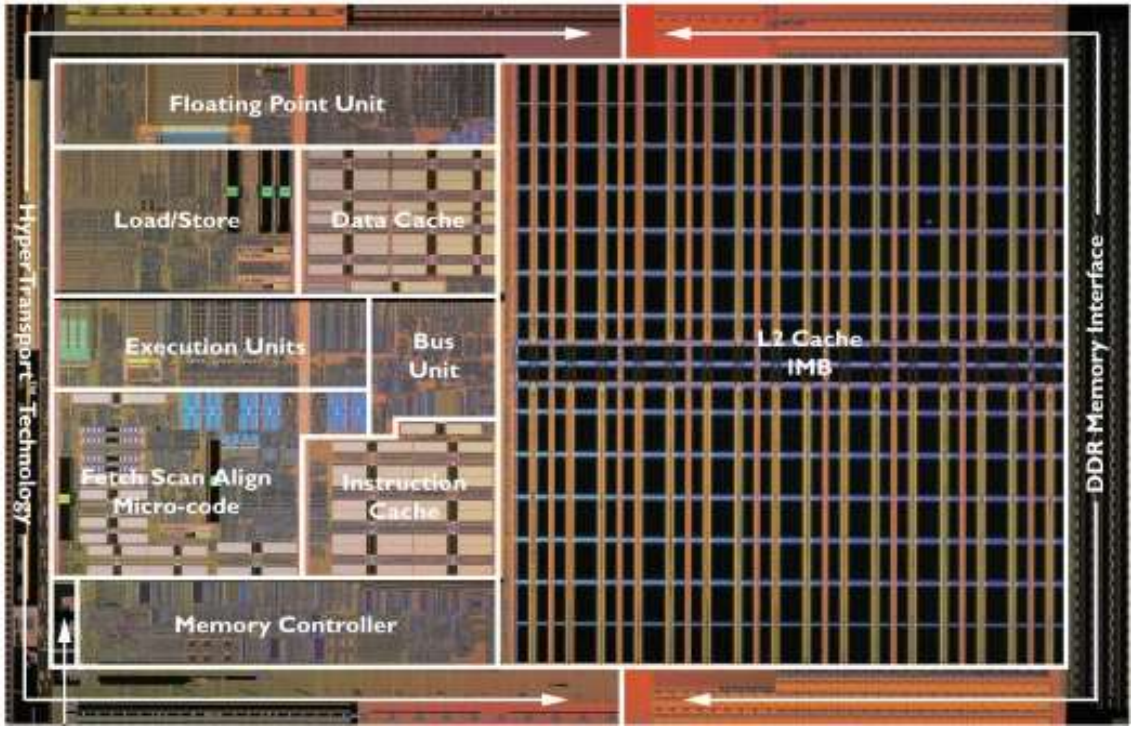

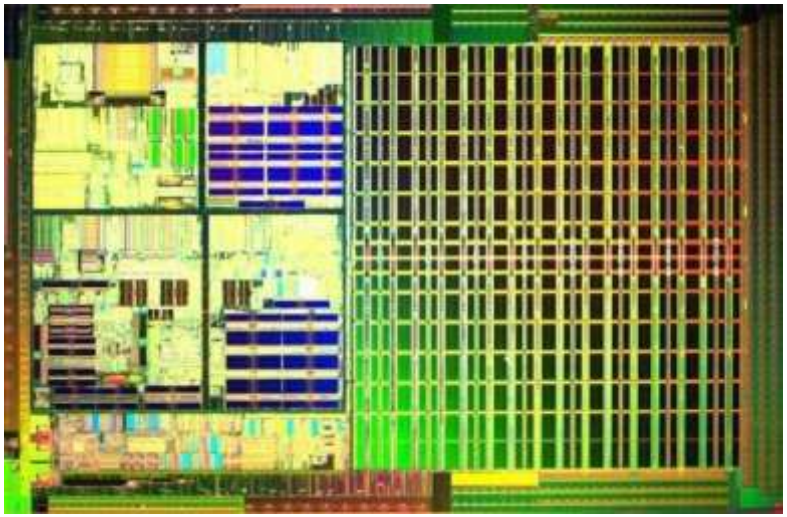

Evolution of the native SC designs with 1 MB L2

130 nm K8-based Athlon 64

ClawHammer (9/2003) [39]
Stepping C0

130 nm
193 mm^2
105.9 mtrs
L2: 1 MB

Shrink

90 nm K8-based Athlon 64

San Diego (5/2005) [43]
Stepping E4

90 nm
115 mm^2
114 mtrs
L2: 1 MB

5.8 Evolution path of AMD's K8-based Athlon 64/64 X2 desktop cores (3)

Contrasting 90 nm SC 512 KB and 1 MB Athlon 64 designs

Venice [43]

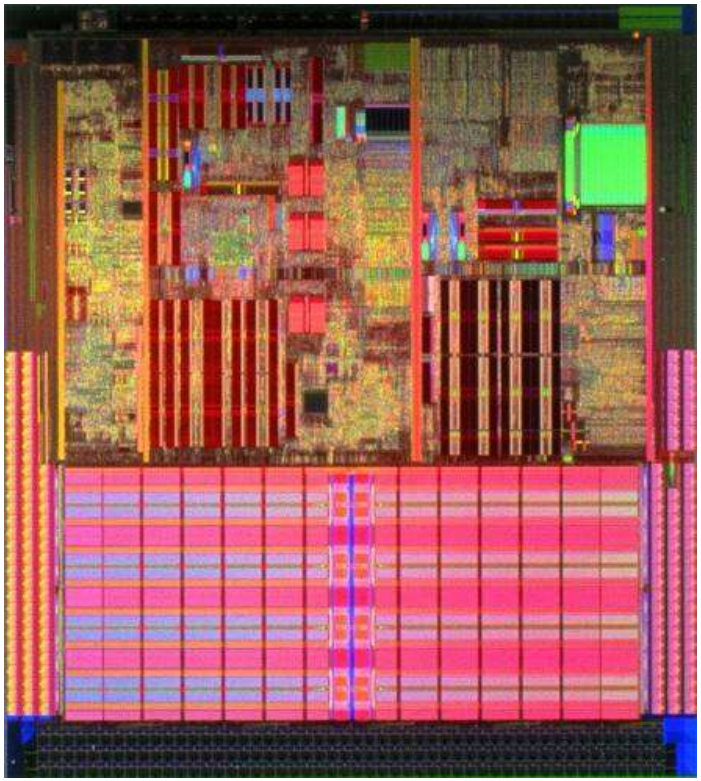

Stepping E3/E6, 90 nm, L2: 512 KB
84 mm^2, 77 mtrs

Venice: Only 512 KB L2

San Diego [43]

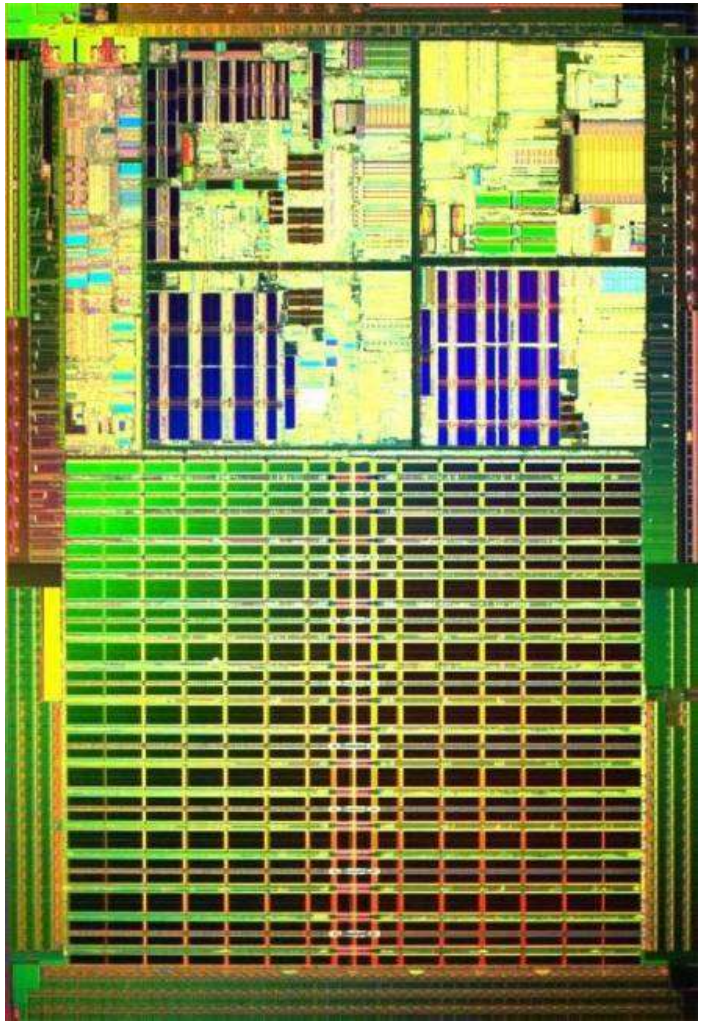

Stepping E4/E6, 90 nm, L2: 1 MB
115 mm^2, 114 mtrs

5.8 Evolution path of AMD's K8-based Athlon 64/64 X2 desktop cores (4)

Contrasting 90 nm SC 1 MB L2 and DC 2x1 MB L2 Athlon 64 designs

San Diego (5/2005) [43]

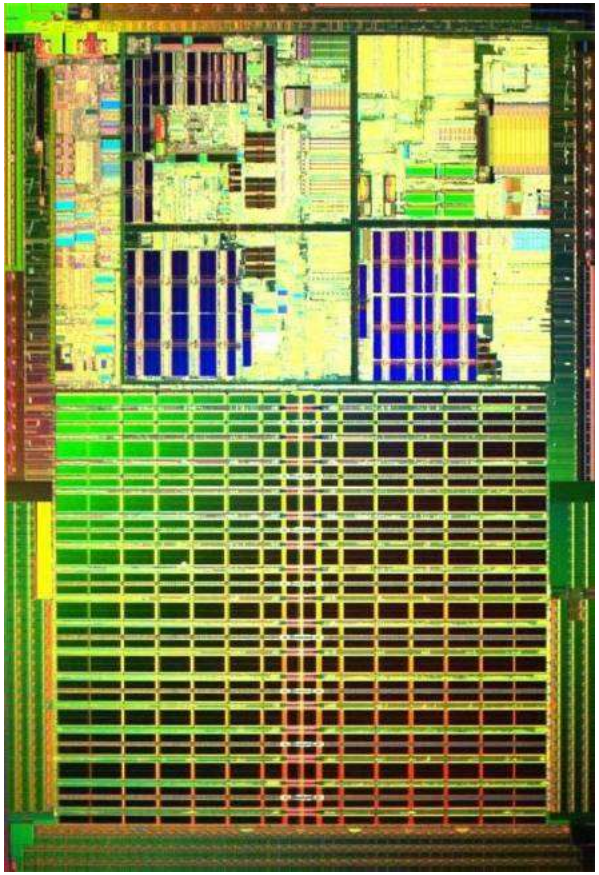

115 mm2, 114 mtrs
Stepping E4 L2: 1MB

Toledo (5/2005) [44]

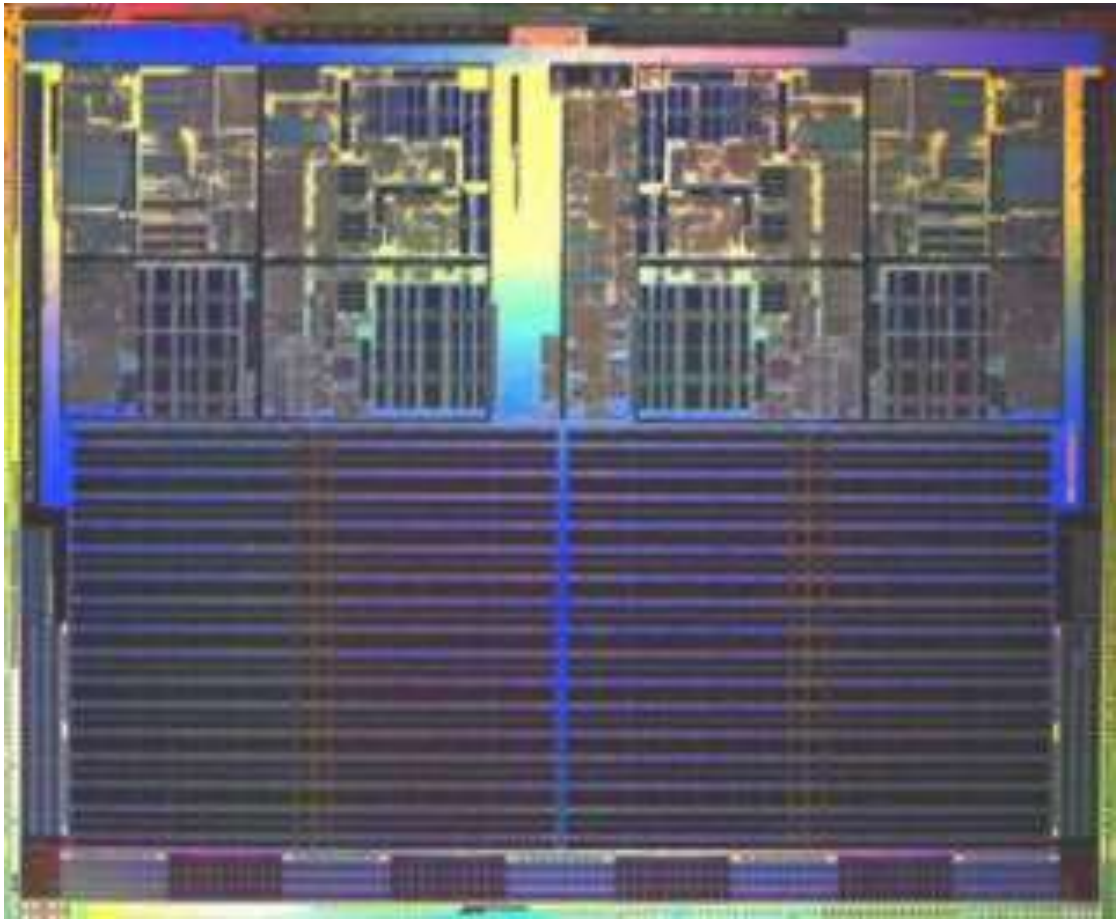

199 mm^2, 233 mtrs
Stepping E6, L2:2x1 MB

The dual-core Toledo core includes basically two San Diego cores (90 nm, L2: 1 MB)

5.8 Evolution path of AMD's K8-based Athlon 64/64 X2 desktop cores (5)

Contrasting the first 90 nm DC and the 90 nm NPT DC 2x1 MB L2 Athlon 64 designs

Toledo (5/2005) [44]

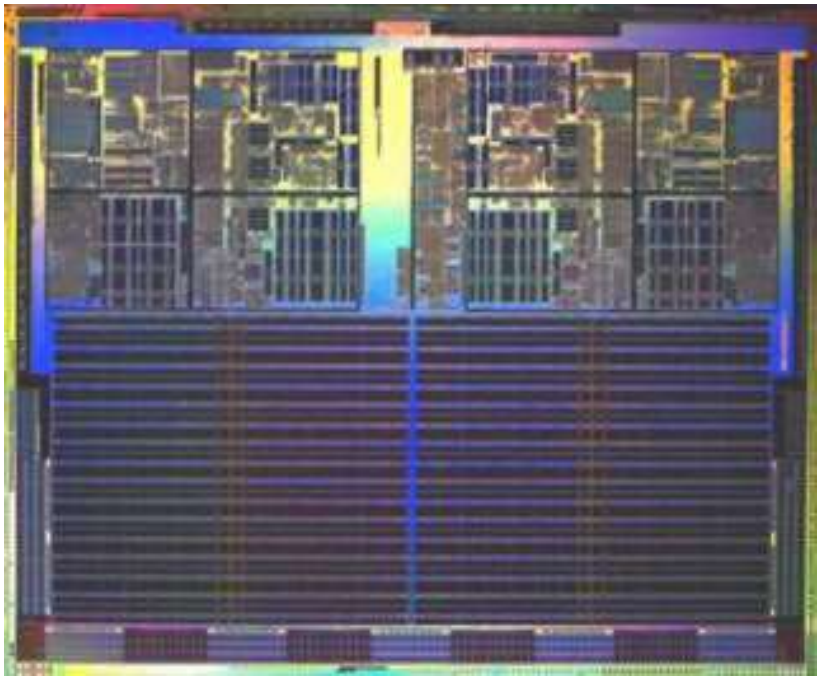

(Stepping E6, 90 nm/199mm^2/233 mtrs.)

L2: 2x1MB (Socket 939)

Windsor (5/2006) [44]

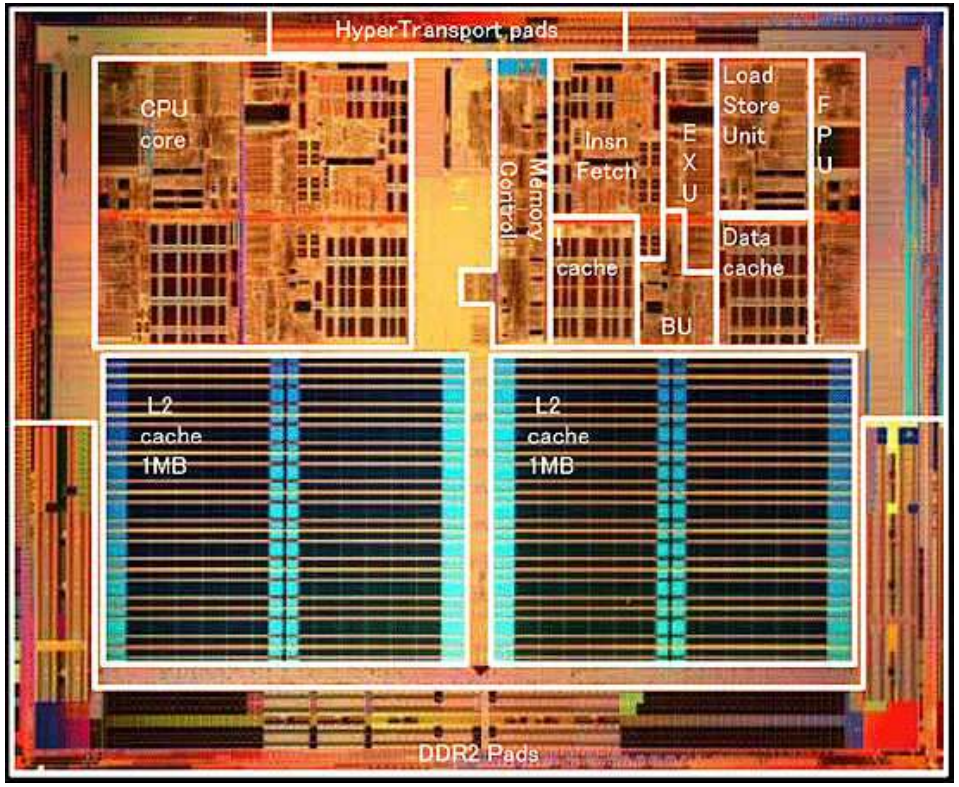

(Stepping F2/3, 90 nm/230 mm^2/227 mtrs.)

L2: 2x1 MB)(Socket AM2)

Stepping F supports DDR2 and the Pacifica virtualization technology

5.8 Evolution path of AMD's K8-based Athlon 64/64 X2 desktop cores (6)

Contrasting the 90 nm DC 2x1 MB Windsor (NPT) and the 65 nm SC 512 KB Brisbane

Windsor (5/2006) [44]

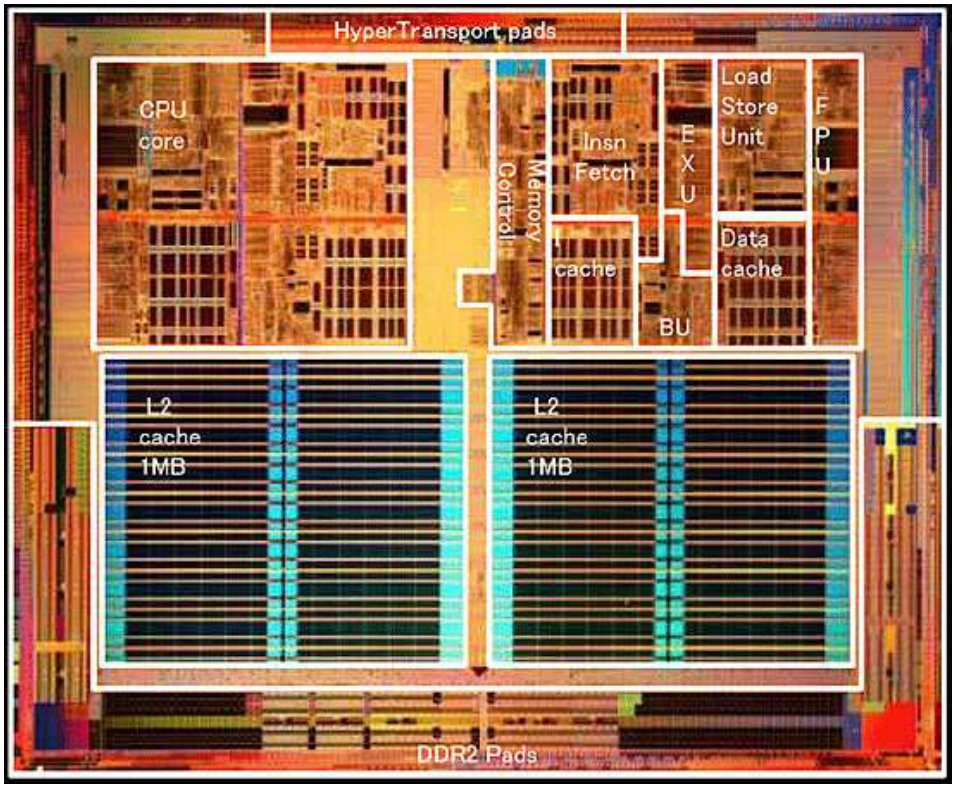

(Stepping F2/F3, 90 nm, L2: 2 x 1 MB)

Brisbane (12/2006) [45]
(Athlon 64 X2)

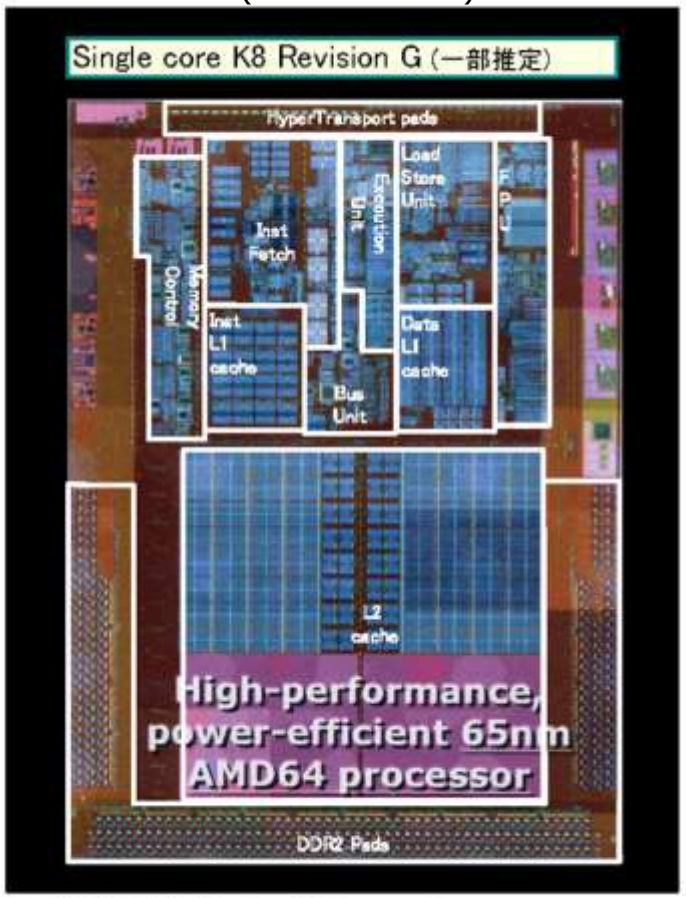

(Stepping G1, 65 nm, 512 KB, enlarged)

Brisbane is the 65 nm shrink of Windsor

5.8 Evolution path of AMD's K8-based Athlon 64/64 X2 desktop cores (7)

Contrasting die plots of the Brisbane and Barcelona cores

Brisbane (12/2006) [44]

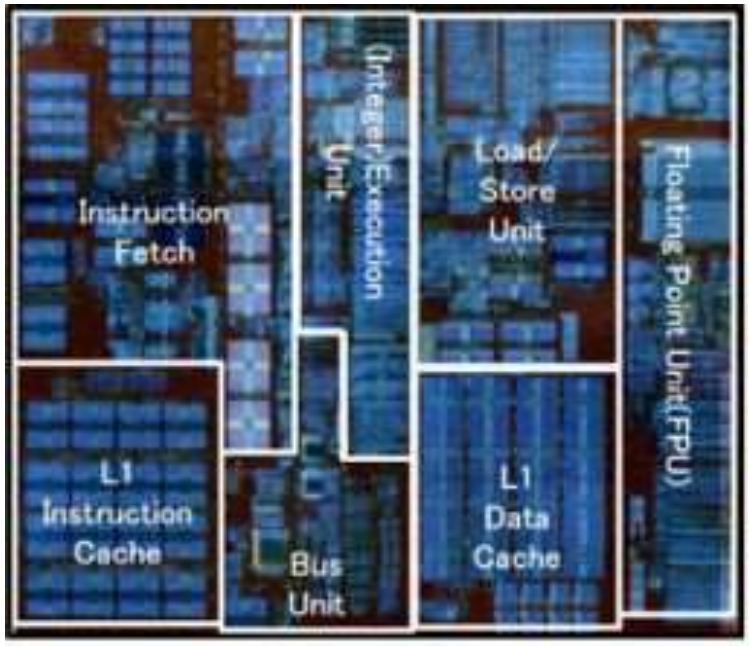

4. gen. K8-based
Athlon 64 X2 scaled down to 65 nm

(65 nm)
126 mm^2, 154 mtrs
K8 based, Stepping G1

Barcelona (8/2007) [44]

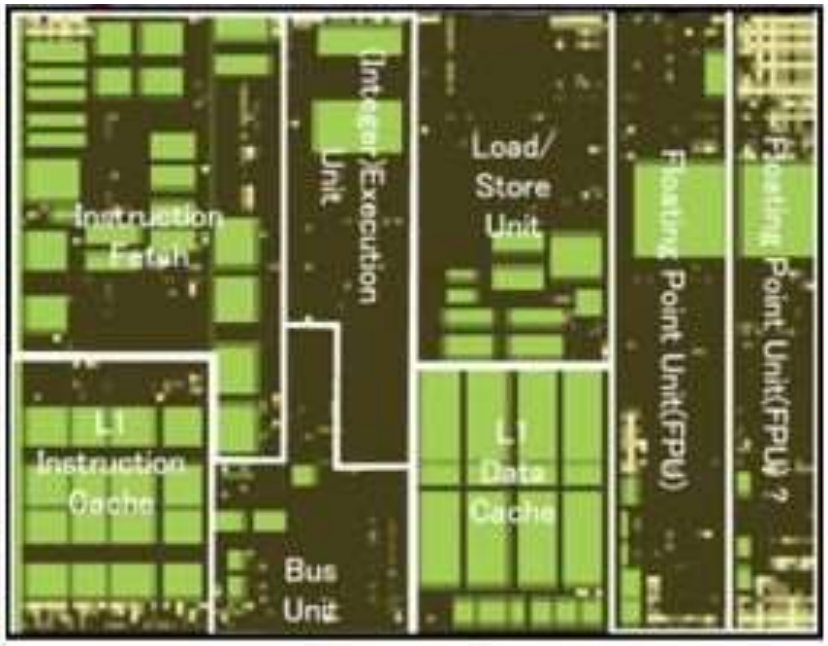

K10-based server

(65 nm)
K10 based
Stepping B1/B2

5.8 Evolution path of AMD's K8-based Athlon 64/64 X2 desktop cores (8)

Main features of AMD's Cool'n'Quiet' technology in K8-based desktops

- It provides dynamic voltage and frequency scaling
  (supported first in the Athlon 64 line (2003).
- Dynamic voltage scaling is performed in the Working state.
  It is implemented by voltage stepping.
- Dynamic frequency scaling is performed in the Stop Grant state
  (while both the processor's clock and the PLL is halted).
  It is implemented by reconfiguring the PLL in the Stop Grant state and waiting for locking the PLL (this needs about 10 μs).
- OS support
  - either by legacy OSs (not supporting the Cool'n'Quiet technology)
    augmented by an AMD provided software package that implements a non ACPI compliant P-state management,
  - or by Windows OSs with native Cool'n'Quiet technology supporting K8-based desktops
    (such as Windows Vista (2006)) which is an ACPI 2.0 compliant solution.

5.8 Evolution path of AMD's K8-based Athlon 64/64 X2 desktop cores (9)

Main features of the implementation of AMD's Cool'n'Quiet technology in K8-based desktops[1] [148]

- Separation of the VID change and FID change phases of P-state transitions in order to be prepared to reduce out time during transitions.
- Introduction of incremental voltage stepping instead of a robust one step voltage ramp in order to reduce noise and allow the processor to resume instruction execution during the voltage change phase as well.

Note

Earlier AMD implementations transitioned voltage in a single step causing significant noise that inhibited resuming instruction execution. For this reason, earlier implementations performed P-state transitions usually in the Deep Sleep state.

[1] AMD's K8-based servers make use of the same technology, as well, presumable also AMD's mobiles, but no related reference could be found for the latter.

6. The K10 Barcelona family

- 6.1 Overview of AMD's K10 processor lines
- 6.2 Main innovations and enhancements of K10 servers
- 6.3 Contrasting utilized and implemented features of the microarchitecture of K10 servers
- 6.4 K10 server lines
- 6.5 Main innovations and enhancements of K10 desktops
- 6.6 K10 desktop lines

6.1 Overview of AMD’s K10 processor lines

6.1 Overview of AMD's K10 processor lines (1)

6.1 Overview of AMD's K10 processor lines-1 [14]

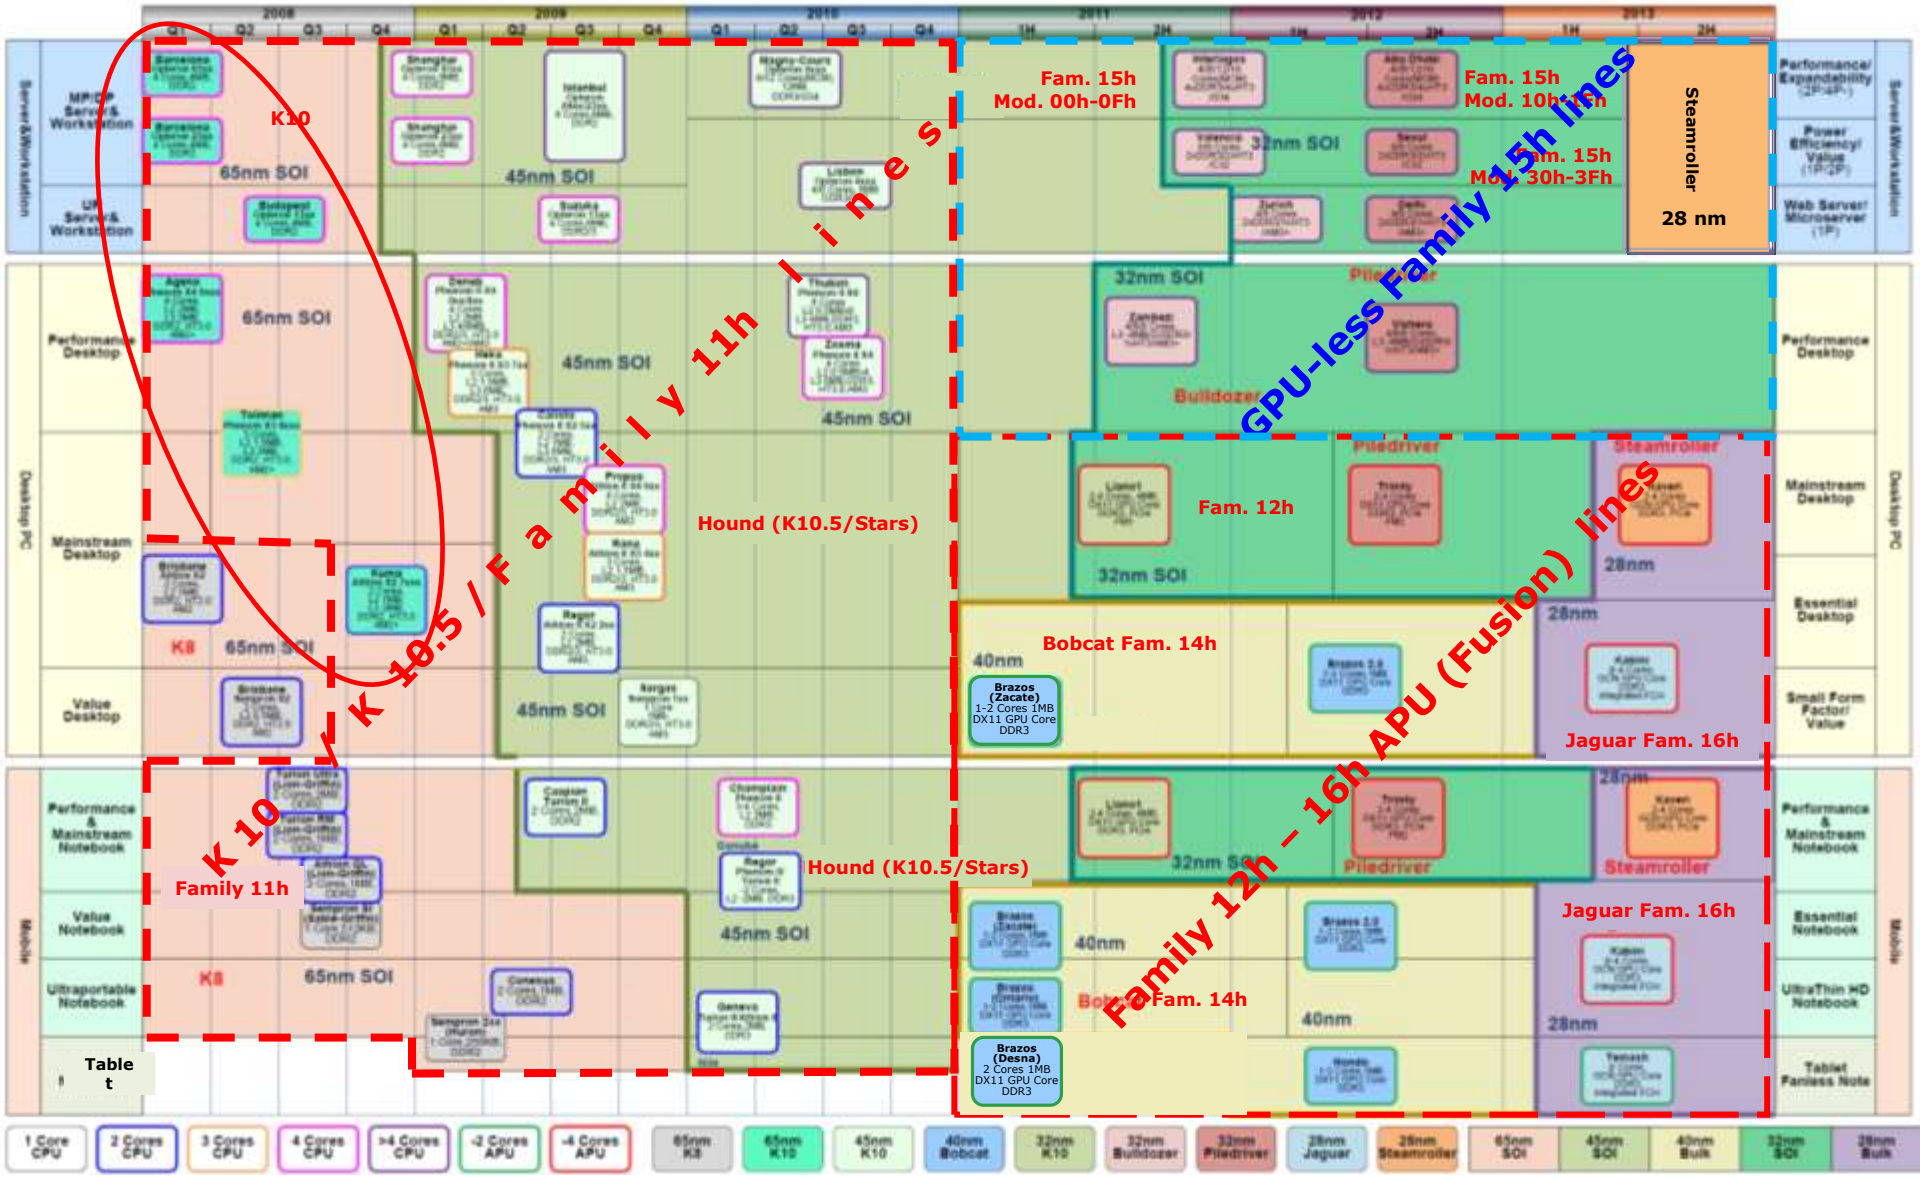

6.1 Overview of AMD's K10 processor lines (2)

Brand names of AMD's K10 (Barcelona)-based processor line

		2003-2007	2007-2008	2008-2011	2009	2009
		K8 (Hammer)	**K10 (Barcelona)**	**K10.5 (Shanghai)**	**K10.5 (Istanbul)**	**K10.5 (Magny- Course)**
Servers	**4P servers**	**See Section 4**	**Barcelona (834x-836x))**	**Shanghai (837x-839x)**	**Istambul (8410-8430)**	**Magny-Course (6100)**
	2P servers		**Barcelona (234x-236x)**	**Shanghai (237x-239x)**	**Istambul (241x-243x)**	**Lisbon (4100)**
	1P servers		**Budapest (135x-136x)**	**Suzuka (138x-139x)**		
Desktops	**High perf.** (~80-120W)		Phenom X4-X2	Phenom II X4-X2	Phenom II X6-X4	
	Mainstream (~60-90W)	Athlon 64 Athlon 64 X2	Athlon X2	Athlon II X4-X2		
	Value (~40-60W)	Sempron		Sempron		
Mobiles	**High perf.** (~30-40W)	Turion 64 X2 (TL 6/5) Turion 64 (ML/MT)		Phenom II (N/P 9xx-6xx) Turion II Ultra (M6xx) Turion II (M/N/P 5xx)		
	Mainstream (~20-30W)	Athlon 64 X2 (TK-5x/4x) Athlon 64 (2xxx+-4xxx+)		Athlon II (M/N/P 3xx) Sempron (M1xx)		
	Ultraportable (~10-20W)	Mobile Sempron (2xxx+-4xxx+) Sempron 2100 fanless		Turion II Neo (K6xx) Athlon II Neo (K1xx) V-series (V1xx)		
Embedded (~10-20W)				Turion II Neo X2 Athlon II Neo X2 Athlon II Neo		

6.1 Overview of AMD's K10 processor lines (3)

Overview of AMD's K10 processor lines –The underlying microarchitectures

- AMD designed two distinct K10 processor lines, one for servers and another one for desktops.
- These lines are based on different microarchitectures, as indicated below.

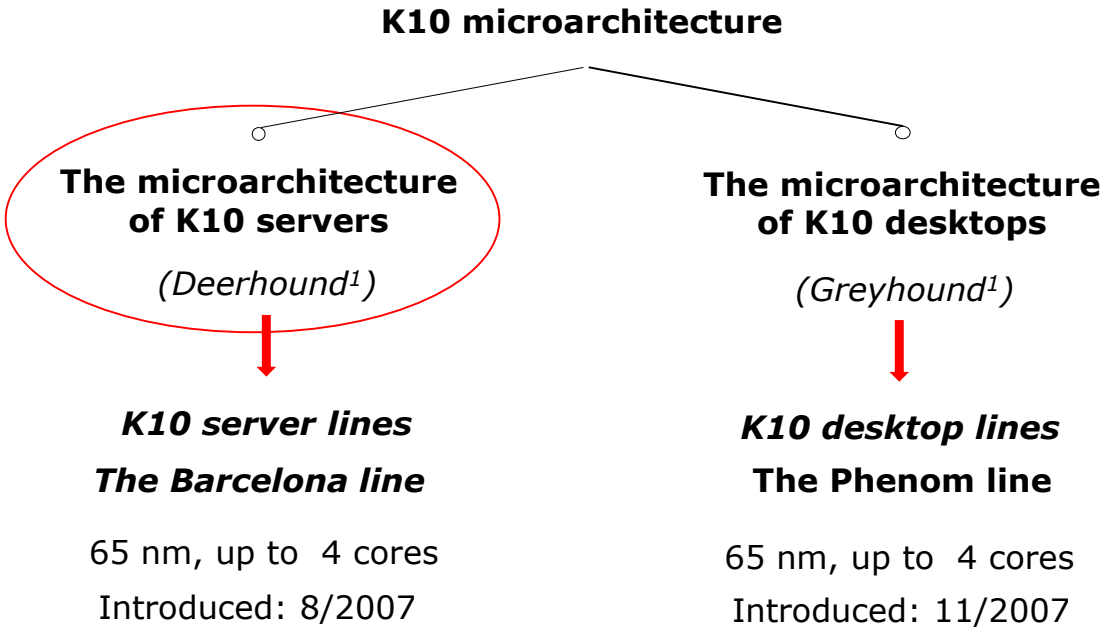

[1] The Deerhound/Greyhound designations are in the available literature not consistently used. Some places use these designations as given above e.g. [46] [124], others adversely, the name Deerhound is used for desktops and Greyhound for servers.

Nevertheless, we assume that the given name allocation is correct since later publications, e.g. [47] use the Greyhound name for subsequent K10.5-based processors that support HT3.0 and DDR2/3 i.e. features that are provided by the Greyhound processor according to the above interpretation.

6.1 Overview of AMD's K10 processor lines (4)

Positioning of the K10 (Barcelona)-based MP server line within AMD's MP 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)	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 (NPT)	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. 17h Mod. 00h-0Fh		6/2017	Epyc (2S!!) (4 dies/proc.)	7000	14 nm	4x(2x4) (32C)	½ MB/C	2 MB/C	DDR4-2666	IFIS 75.8 GB/s	SP3

6.1 Overview of AMD's K10 processor lines (5)

Positioning of the high performance K10 DT line within AMD's high performance DT 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. 17h (Zen)		3/2017	Summit Ridge	Ryzen 7	14 nm	8	8x1/2 MB	16 MB	DDR4-2993	-	AM4
Fam. 17 (Zen+)		4/2018	Pinnacle Ridge	Ryzen 7	12 nm	8	8x1/2 MB	16 MB	DDR4-2933	-	AM4

6.1 Overview of AMD's K10 processor lines (6)

Sockets of K10 DP, MP servers and desktops [based on 42]

Supports DDR3

Socket **G34** (1944 pins, MP servers K10.5 „Magny-Course", Fam. 15h)
Socket **C32** (1207 pins, DP servers K10.5 „Magny –Course, Fam. 15h)
Socket **AM3** (941 pins, K10.5 Shanghai UP servers, desktops, except first released desktops, K10.5 Istambul desktops)
Socket **AM3+** (942 pins, UP servers Fam. 15h, desktops Fam. 15h)
Socket **FM1** (905 balls BGA, desktops, mobiles Fam. 12h)
Socket **S1g4** (638 pins, K10.5 mobile)
Socket **FT1** (413 balls BGA, desktops, mobiles Fam. 14h)
Socket **FT3** (769 balls BGA, desktops, mobiles Fam. 16h)

S:: Servers
DT: Desktop
M: Mobile
UPT: Ultraportable

SPP: Single power Plane
DPP: Dual Power planes
HT 2.0: HyperTransport 2.0
HT 3.0: HyperTransport 3.0

K10.5 | **Fam. 15h** | **Fam. 15h** | **Fam. 15h**

45nm 2x6C „Magny-Course" | 32 nm 2x8C „Bulldozer" | 32 nm 2x8C/4C APU „Piledriver" | 28 nm 2x8C „Steamroller"

3st Generation Platform

45 nm 4C, DT/M „Shanghai" | 45 nm 6C, DT „Istambul" | 32 nm 4C APU „Llano" | 32 nm 2C APU „Bobcat" | 28 nm 4C APU „Jaguar" | 28 nm 4C APU „Puma+"

Fam. 12h | **Fam. 14h** | **Fam. 16h** | **Fam. 16h**

K8 (NPT) | **K10** | **K10.5** | **K10.5**

90 nm 2C NPT „Santa Rosa" | 65 nm 4C „Barcelona" | 45 nm 4C, S „Shanghai" | 45 nm 6C, S „Istambul"

2st Generation Platform

65 nm 4C „Brisbana" | 65 nm 2C „Griffin"

K8 | **Fam. 11h**

K8 | **K8** | **K8**

130 nm 1C „SledgeH." | 90 nm 1C „Athens" | 90 nm 2C „Egypt"

1st Generation Platform

Supports DDR2

Socket **F** (1207 pins, SPP, HT 2.0: DP, MP servers K8 NPT, K10)
Socket **F+** (1207 pins, DPP, HT 3.0: DP, MP servers K10.5)
Socket **AM2** (940 pins, SPP, HT 2.0: K8 NPT UP serv., K8 NPT, Brisbane DTs)
Socket **AM2+** (940 pins, DPP: K10 UP servers, K10 desktops, first K10.5 Shanghai desktops)
Socket **S1g1** (638 pins, K8 NPT, Brisbane mobiles)
Socket **S1g2** (638 pins, Fam. 11h mobiles)
Socket **S1g3** (638 pins, K10.5 Shanghai (Caspian die) mobiles)
Socket **ASB1** (812 balls BGA, Fam. 11h (Huron, Conesus dies) UPT mobiles)
Socket **ASB2** (812 balls BGA, K10.5 (Shanghai) (Geneva die), UPT mobiles)

Supports DDR

Socket **940** (130nm K8 UP/DP/MP servers, 90 nm K8 SC and 90 nm K8 DC (Egypt) DP/MP servers)
Socket **939** (90nm K8 SC and 90 nm K8 DC (Egypt) UP servers, K8 desktops with 128-bit mem. channel up to the 90 nm DC (Egypt) lines)
Socket **754** (K8 desktops, mobiles with 64-bit mem. channel up to the 90 nm SC (Athens) line)

2003 2004 2005 2006 2007 2008 2009 2010 2011 2012 2013 2014

6.2 Main enhancements of K10 servers

6.2 Main enhancements of K10 servers [48]

Comprehensive Upgrades for SSE128

Can quadruple floating-point capabilities

New Highly Efficient Cache Structure with Shared L3 Cache

Balance of dedicated and shared cache for optimum quad-core performance

CPU Core Enhancements

To benefit applications by improving overall efficiency and performance of cores

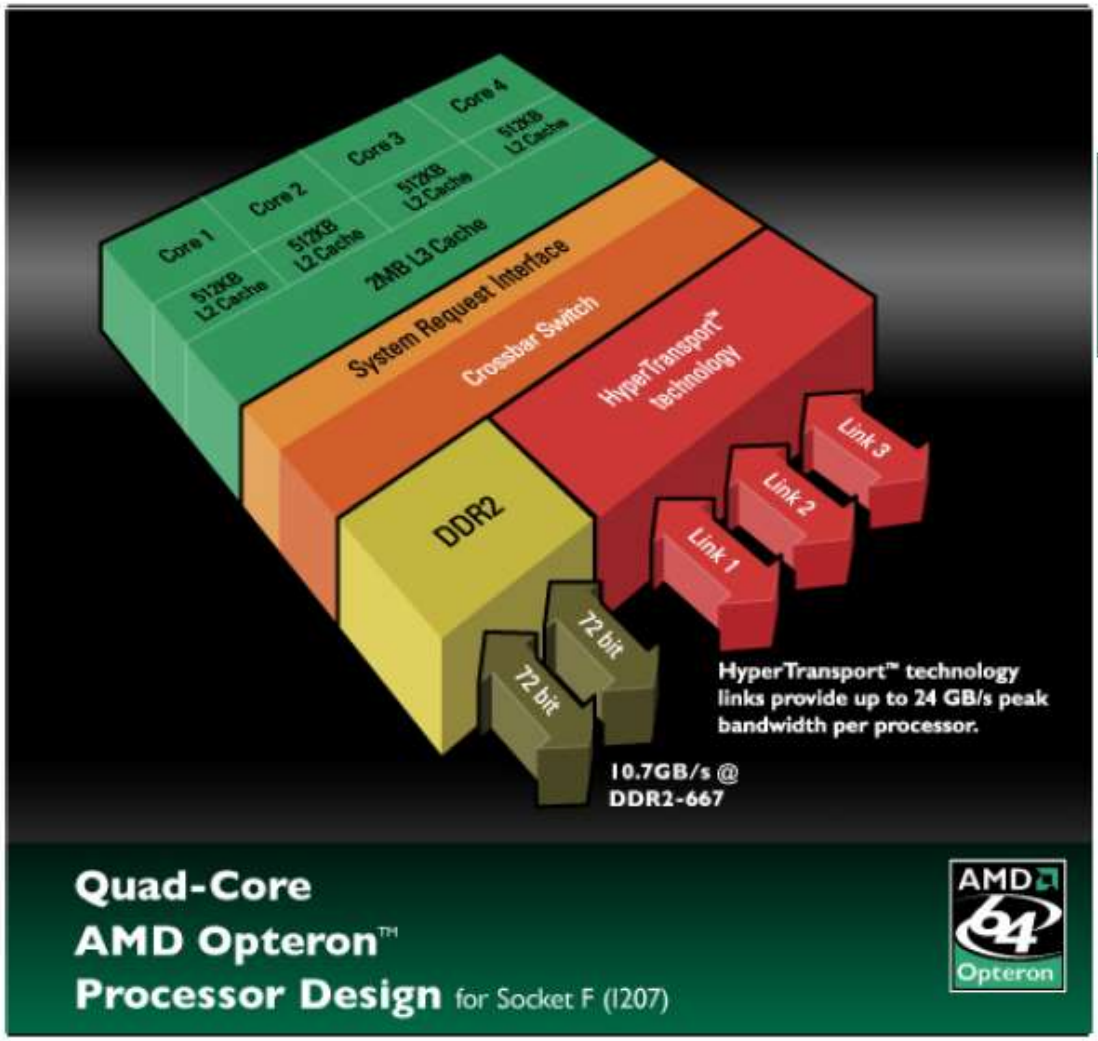

Virtualization Enhancements

New "Nested Paging" feature designed for near native performance on virtualization applications

Advanced Power Management

Provides granular power management resulting in improved power efficiency

DRAM Controller Enhancements

To improve overall memory performance with native quad-core processing

6.2 Main enhancements of K10 servers (2)

a) 128-bit SSE and FP units as well as 128-bit loads [49]

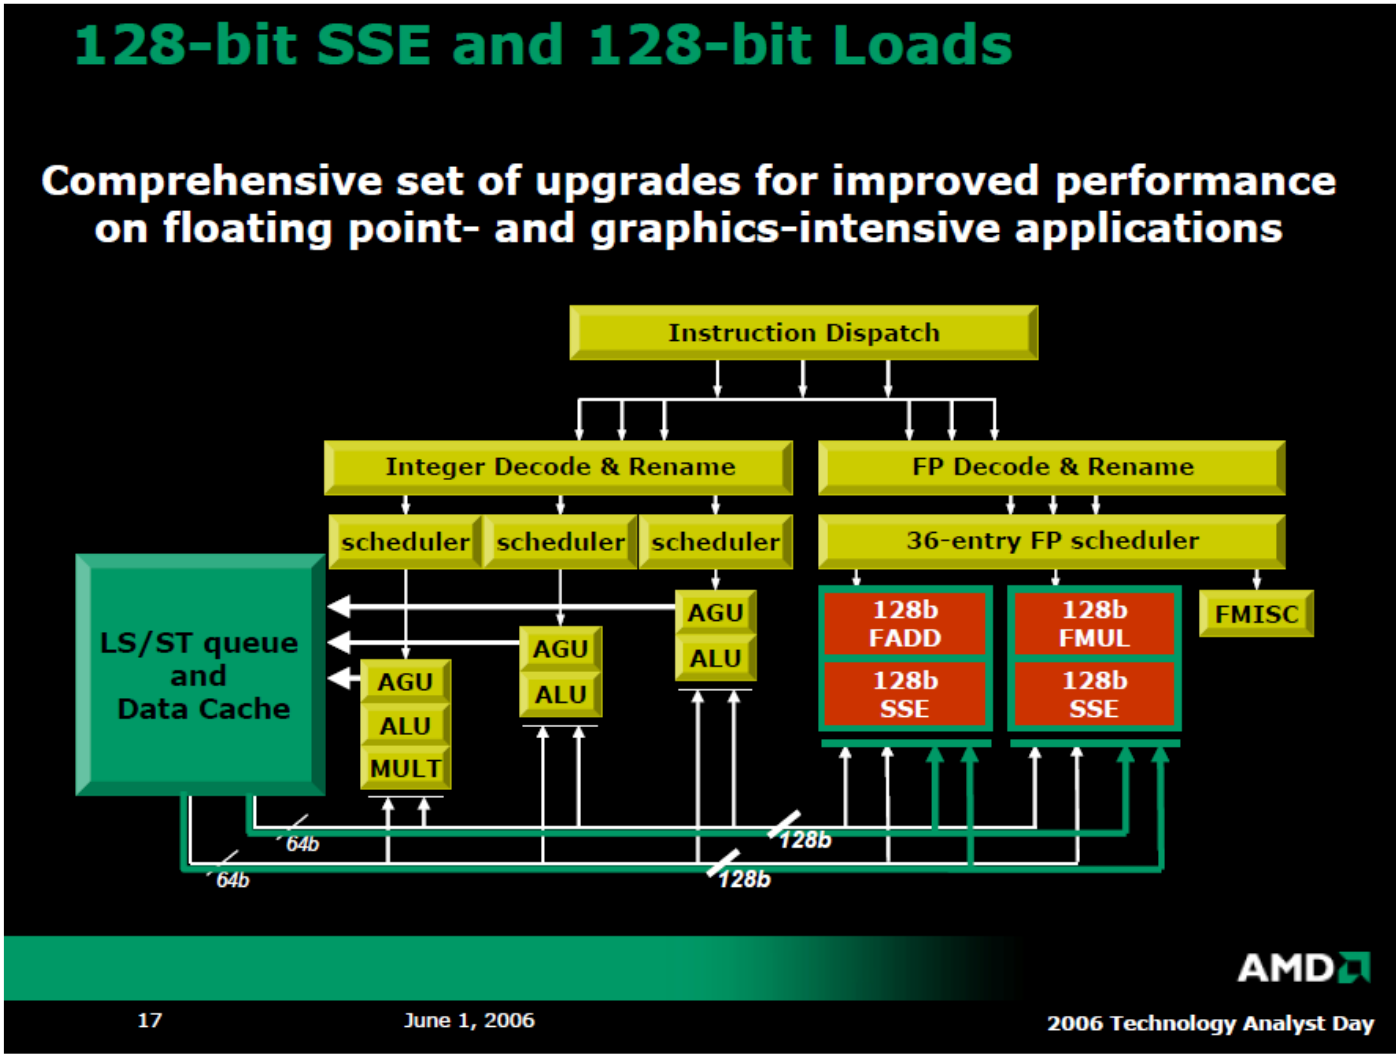

6.2 Main enhancements of K10 servers (3)

b) 3-level cache system with a shared exclusive L3 cache [63]

Dedicated L1
- Locality keeps most critical data in the L1 cache
- Lowest latency
- 2 loads per cycle

Dedicated L2
- Sized to accommodate the majority of working sets today
- Dedicated to eliminate conflicts common in shared caches
  - Better for Virtualization

Shared L3 – NEW
- Victim-cache architecture maximizes efficiency of cache hierarchy
- Fills from L3 leave likely shared lines in the L3
- Sharing-aware replacement policy
- Ready for expansion at the right time for customers

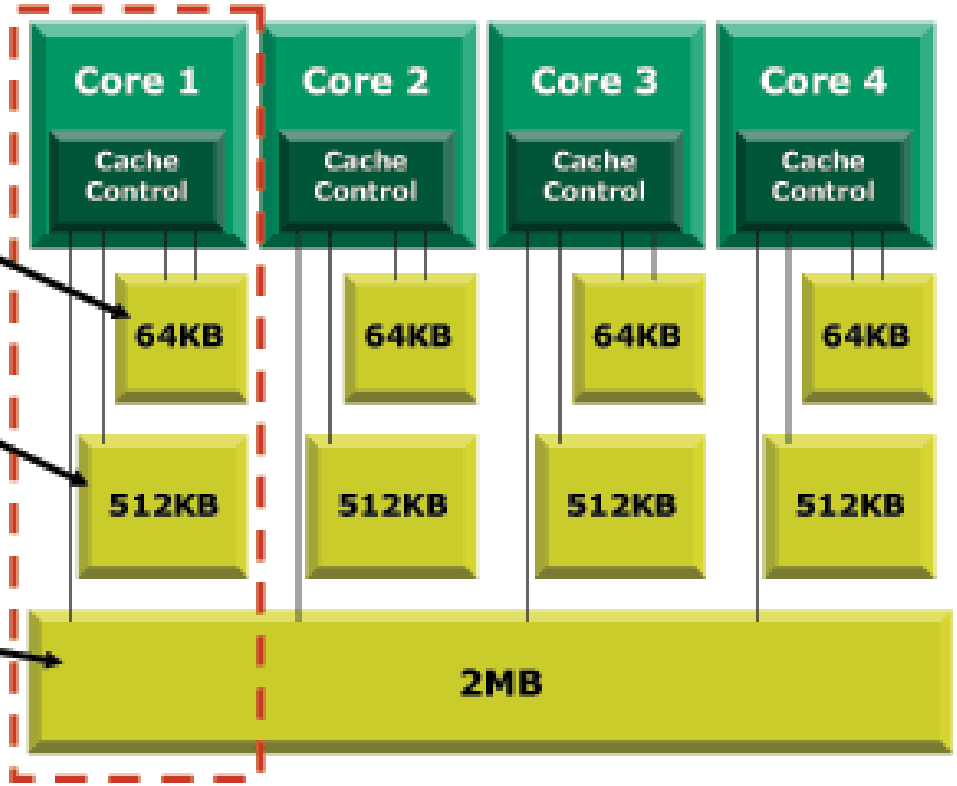

Figure: Cache architecture of the QC Barcelona [63]

6.2 Main enhancements of K10 servers (4)

Exclusive L2 and L3 caches [40]

Both the L2 and L3 caches are exclusive.
This is in contrast to Intel's cache hierarchy design, as Intel prefers inclusive caches.

L2:512 K

L3: 2 M

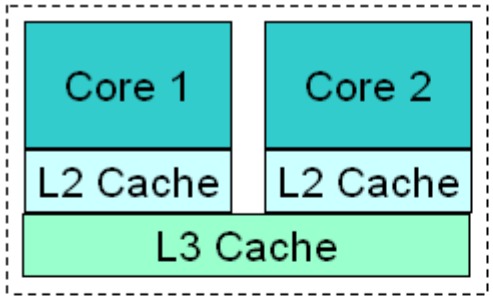

K10-based dual-core CPU

L2:512 K

L3: 2 M

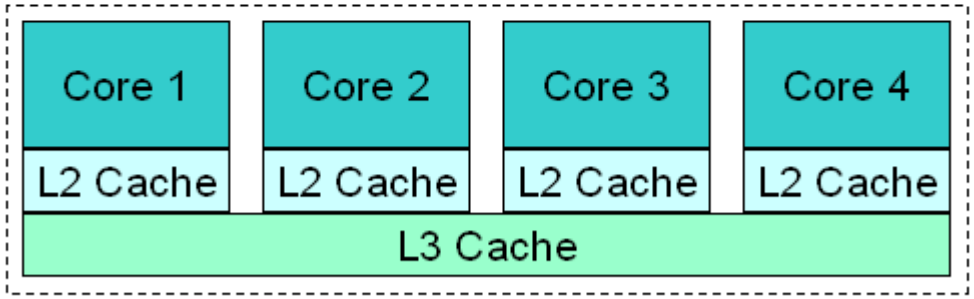

K10-based quad-core CPU

Remark

Particular K10-based desktops and mobile models may have smaller L2 and L3 caches.

6.2 Main enhancements of K10 servers (5)

Modified exclusive L3 cache policy [50]

(Termed also as "mostly exclusive caching")

If addressed data are missing in the entire cache system the referenced line will be loaded from the memory directly into the L1 cache, as in case of exclusive caches.

When later this line becomes evicted first from L1 and subsequently also from L2, it will be brought into L3.

If data kept in the L3 cache will be referenced anew, AMD's L3 cache behaves differently than "ordinary" exclusive caches.

In case of "ordinary" exclusive caches referenced data kept in the L3 cache will be evicted from the L3 cache and brought into the L1 cache in order to free space for victim lines from L2.

In "mostly exclusive caching" employed by AMD for their L3 caches, if it is likely that accessed data from the L3 cache will be used by multiple cores, data will not be removed from the L3 cache but will be let in it.

By contrast, if it is likely, that data accessed in the L3 cache will only be used by a single core, data will be removed from the L3 cache as in case of "ordinary" exclusive caches.

Another feature of the cache policy is that in case of evicting L3 cache lines non-shared lines will be preferred over shared lines.

6.2 Main enhancements of K10 servers (6)

Contrasting exclusive and inclusive caches

Exclusive caches allow a more efficient use of a given transistor budget, as they avoid data duplication in subsequent cache levels but they generate more coherency traffic.

For higher core counts and for larger caches available with raising transistor counts at hand, however the disadvantage of higher coherency traffic presumably outweights the advantage of a more efficient cache implementation, so in this case inclusive caches promise more benefits.

c) Independent memory controllers [40]

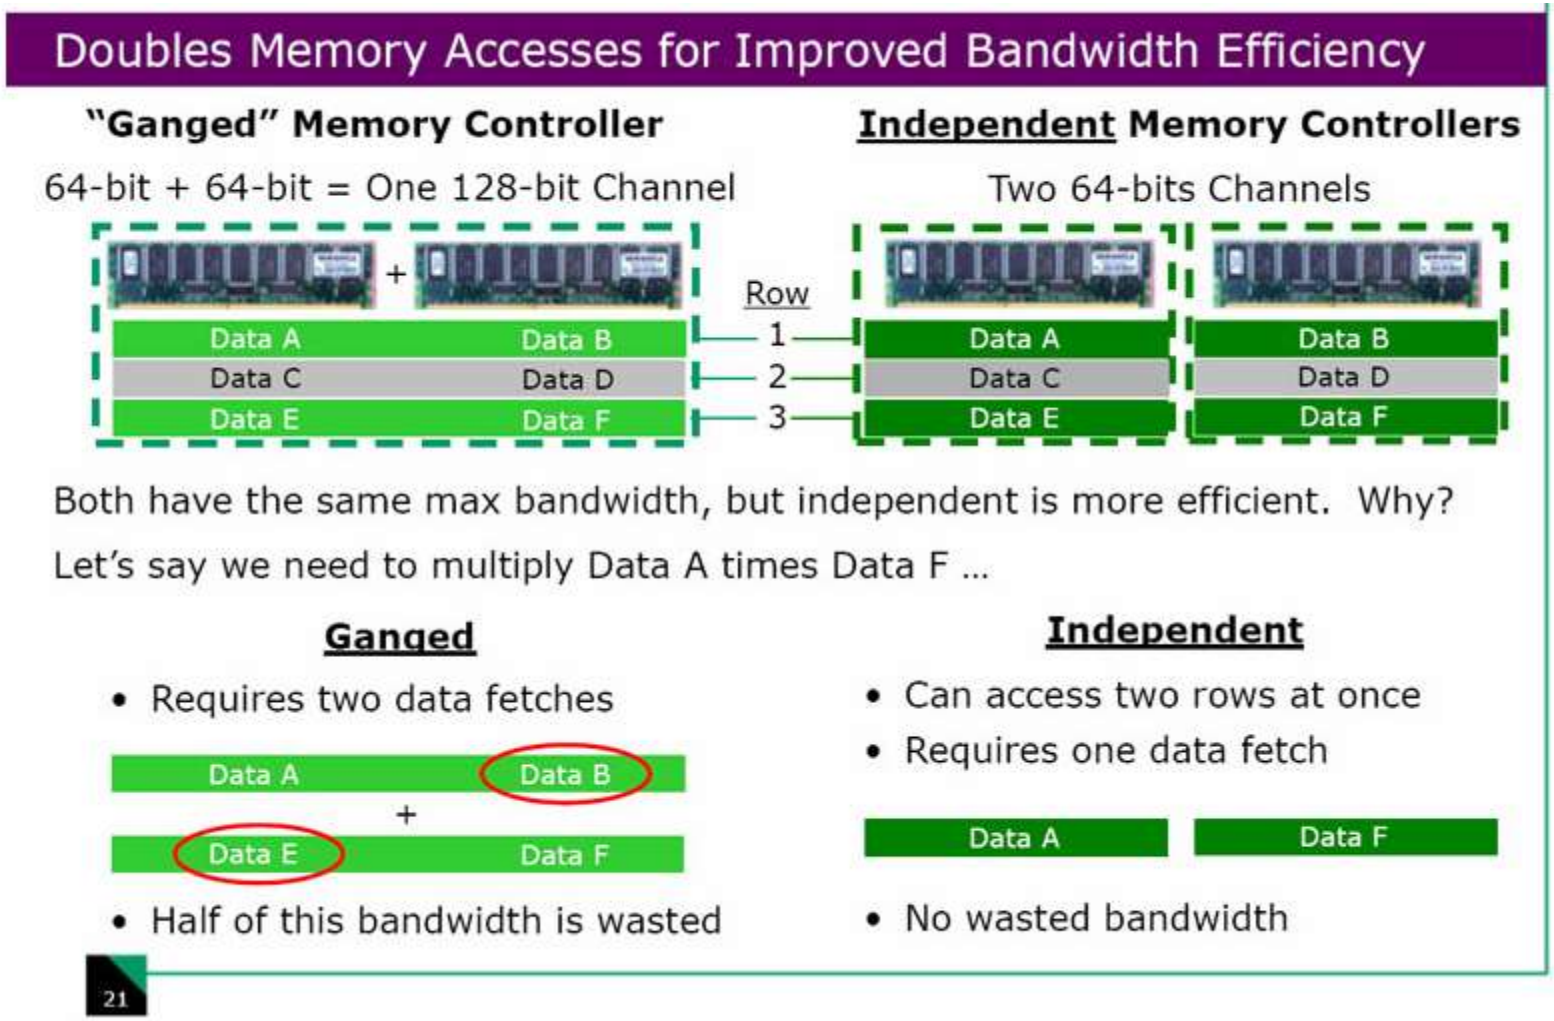

6.2 Main enhancements of K10 servers (8)

d) Supporting HT 3.0 in K10 UP servers-1

In contrast to DP and MP servers, K10 UP servers (designated as Budapest that were released about 8 months later than the K10 DP/MP (Barcelona servers) were already equipped with the Socket AM2+.

Socket AM2+ supported already HT 3.0 and dual power planes, needed for Dual Dynamic Power management.

6.2 Main enhancements of K10 servers (9)

Sockets of K10 UP servers [based on 42]

Supports DDR3

Socket **G34** (1944 pins, MP servers K10.5 „Magny-Course", Fam. 15h)
Socket **C32** (1207 pins, DP servers K10.5 „Magny –Course, Fam. 15h)
Socket **AM3** (941 pins, K10.5 Shanghai UP servers, desktops, except first released desktops, K10.5 Istambul desktops)
Socket **AM3+** (942 pins, UP servers Fam. 15h, desktops Fam. 15h)
Socket **FM1** (905 balls BGA, desktops, mobiles Fam. 12h)
Socket **S1g4** (638 pins, K10.5 mobile)
Socket **FT1** (413 balls BGA, desktops, mobiles Fam. 14h)
Socket **FT3** (769 balls BGA, desktops, mobiles Fam. 16h)

S:: Servers
DT: Desktop
M: Mobile
UPT: Ultraportable

SPP: Single power Plane
DPP: Dual Power planes
HT 2.0: HyperTransport 2.0
HT 3.0: HyperTransport 3.0

K10.5 | Fam. 15h | Fam. 15h | Fam. 15h
45nm 2x6C „Magny-Course" | 32 nm 2x8C „Bulldozer" | 32 nm 2x8C/4C APU „Piledriver" | 28 nm 2x8C „Steamroller"

3st Generation Platform

45 nm 4C, DT/M „Shanghai" | 45 nm 6C, DT „Istambul" | 32 nm 4C APU „Llano" | 32 nm 2C APU „Bobcat" | 28 nm 4C APU „Jaguar" | 28 nm 4C APU „Puma+"
Fam. 12h | Fam. 14h | Fam. 16h | Fam. 16h

K8 (NPT) | K10 | K10.5 | K10.5
90 nm 2C NPT „Santa Rosa" | 65 nm 4C „Barcelona" | 45 nm 4C, S „Shanghai" | 45 nm 6C, S „Istambul"

2st Generation Platform

65 nm 4C „Brisbana" | 65 nm 2C „Griffin"
K8 | Fam. 11h

K8 | K8 | K8
130 nm 1C „SledgeH." | 90 nm 1C „Athens" | 90 nm 2C „Egypt"

1st Generation Platform

Supports DDR2

Socket **F** (1207 pins, SPP, HT 2.0: DP, MP servers K8 NPT, K10)
Socket **F+** (1207 pins, DPP, HT 3.0: DP, MP servers K10.5)
Socket **AM2** (940 pins, SPP, HT 2.0: K8 NPT UP serv., K8 NPT, Brisbane DTs)
Socket **AM2+** (940 pins, DPP: K10 UP servers, K10 desktops, first K10.5 Shanghai desktops)
Socket **S1g1** (638 pins, K8 NPT, Brisbane mobiles)
Socket **S1g2** (638 pins, Fam. 11h mobiles)
Socket **S1g3** (638 pins, K10.5 Shanghai (Caspian die) mobiles)
Socket **ASB1** (812 balls BGA, Fam. 11h (Huron, Conesus dies) UPT mobiles)
Socket **ASB2** (812 balls BGA, K10.5 (Shanghai) (Geneva die), UPT mobiles)

Supports DDR

Socket **940** (130nm K8 UP/DP/MP servers, 90 nm K8 SC and 90 nm K8 DC (Egypt) DP/MP servers)
Socket **939** (90nm K8 SC and 90 nm K8 DC (Egypt) UP servers, K8 desktops with 128-bit mem. channel up to the 90 nm DC (Egypt) lines)
Socket **754** (K8 desktops, mobiles with 64-bit mem. channel up to the 90 nm SC (Athens) line)

2003 2004 2005 2006 2007 2008 2009 2010 2011 2012 2013 2014

6.2 Main enhancements of K10 servers (10)

e) Virtualization enhancements

It will not discussed here.

6.2 Main enhancements of K10 servers (11)

f) Related enhancements of the K10 microarchitecture vs. the K8 microarchitecture [51]

Remark

In 3-level caches 512 KB cache size is typically chosen as an expected optimum when considering both the access time and cache misses.

Red marks the differences between the K8 and the k10 microarchitectures

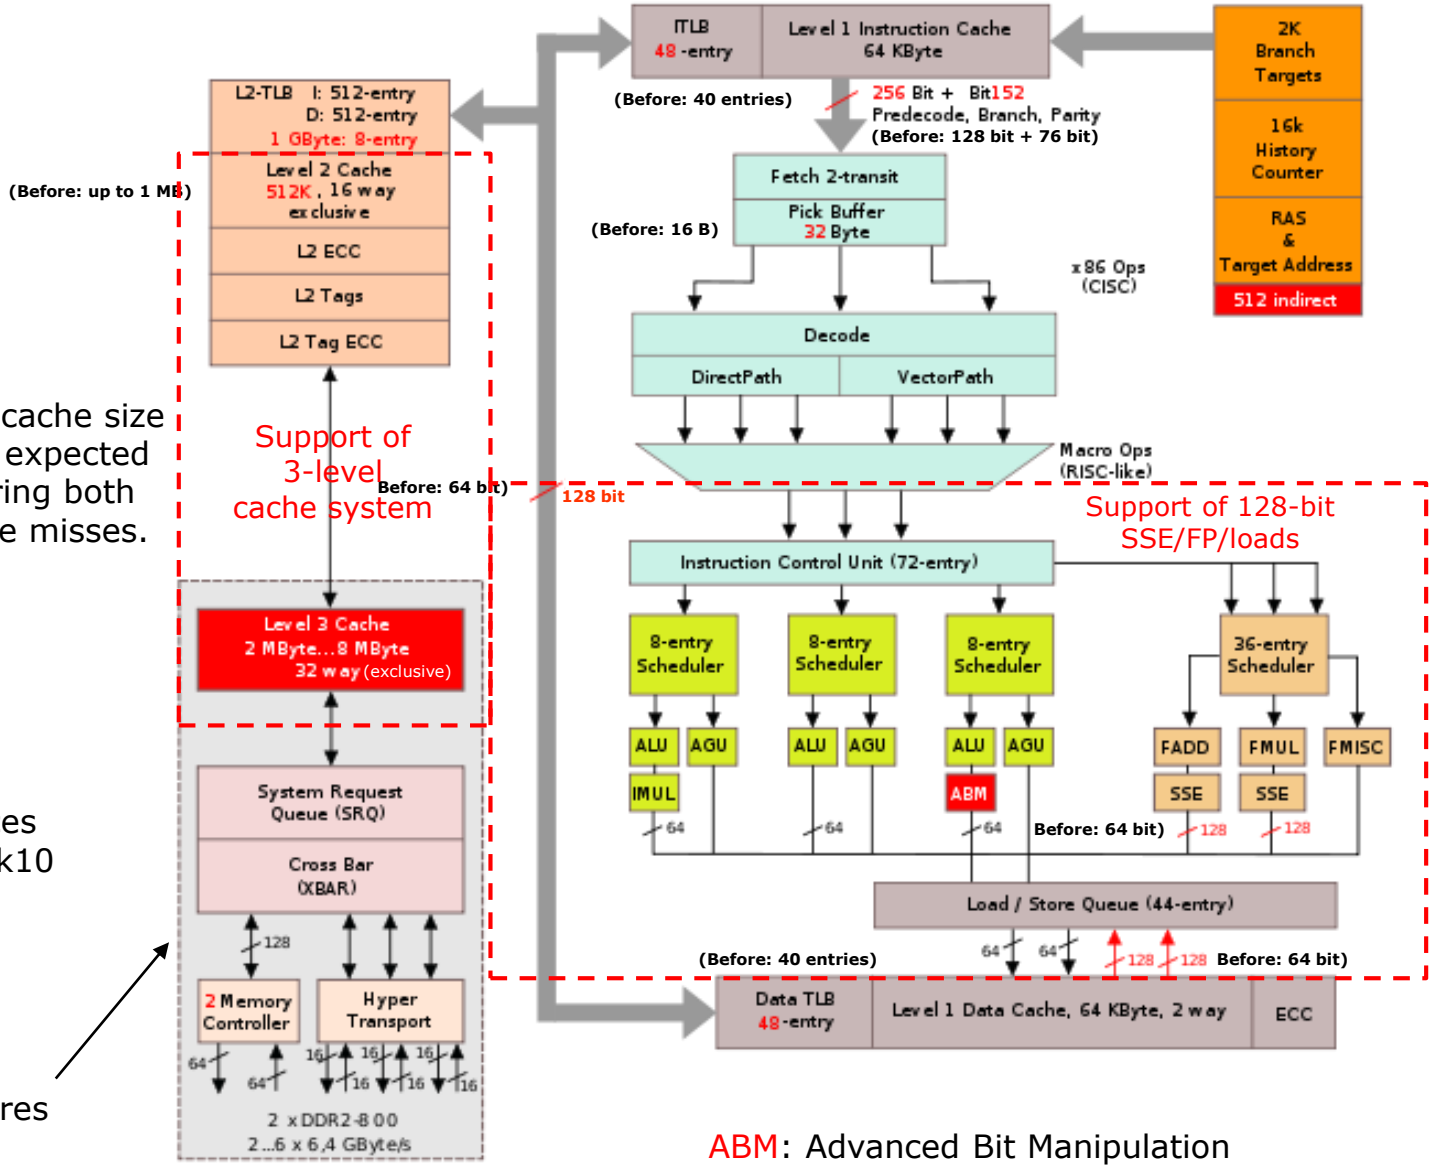

Shared by all four cores

ABM: Advanced Bit Manipulation

6.2 Main enhancements of K10 servers (12)

AMD's power management techniques K8 – Family 15h (Bulldozer) (based on [53])

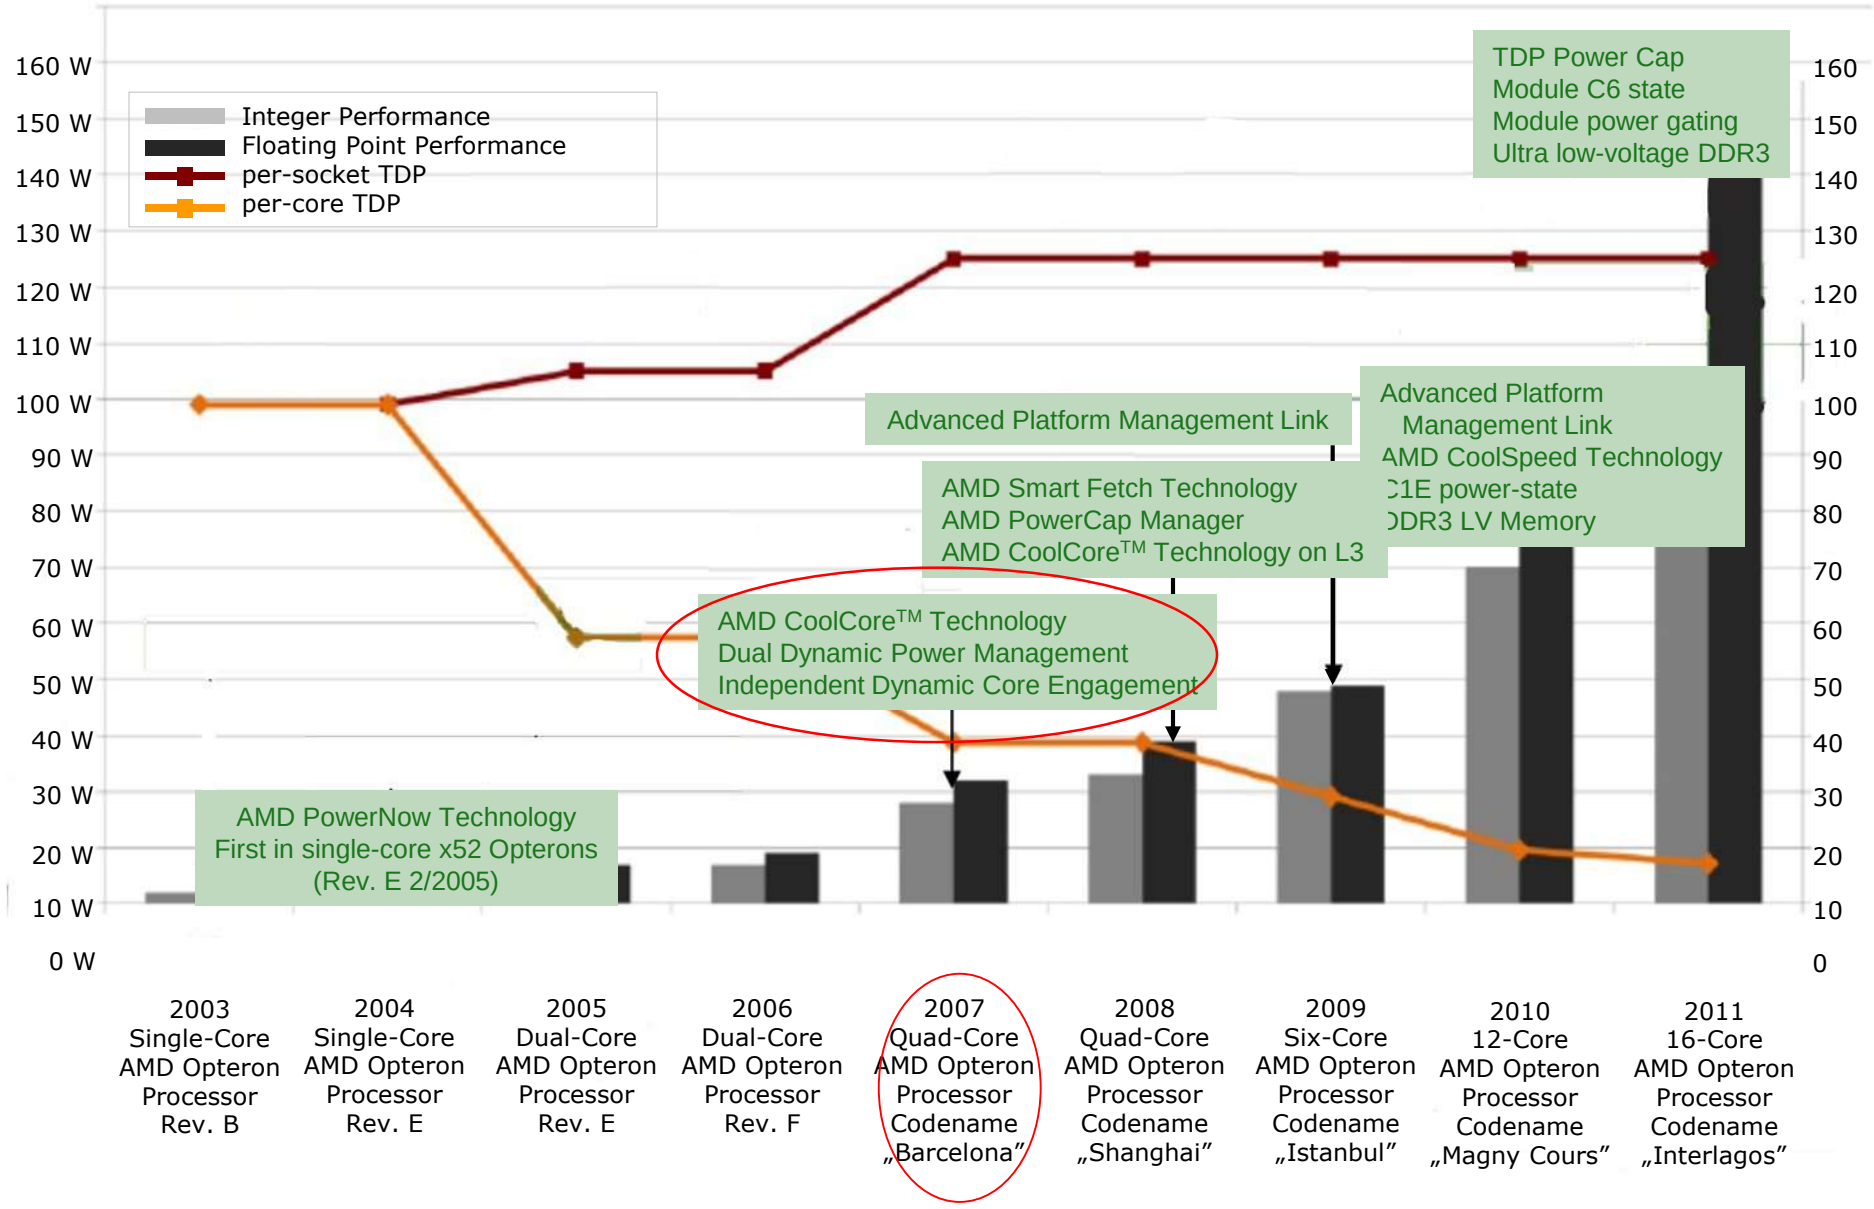

g) Advanced power management techniques of K10-based servers-2 [54]

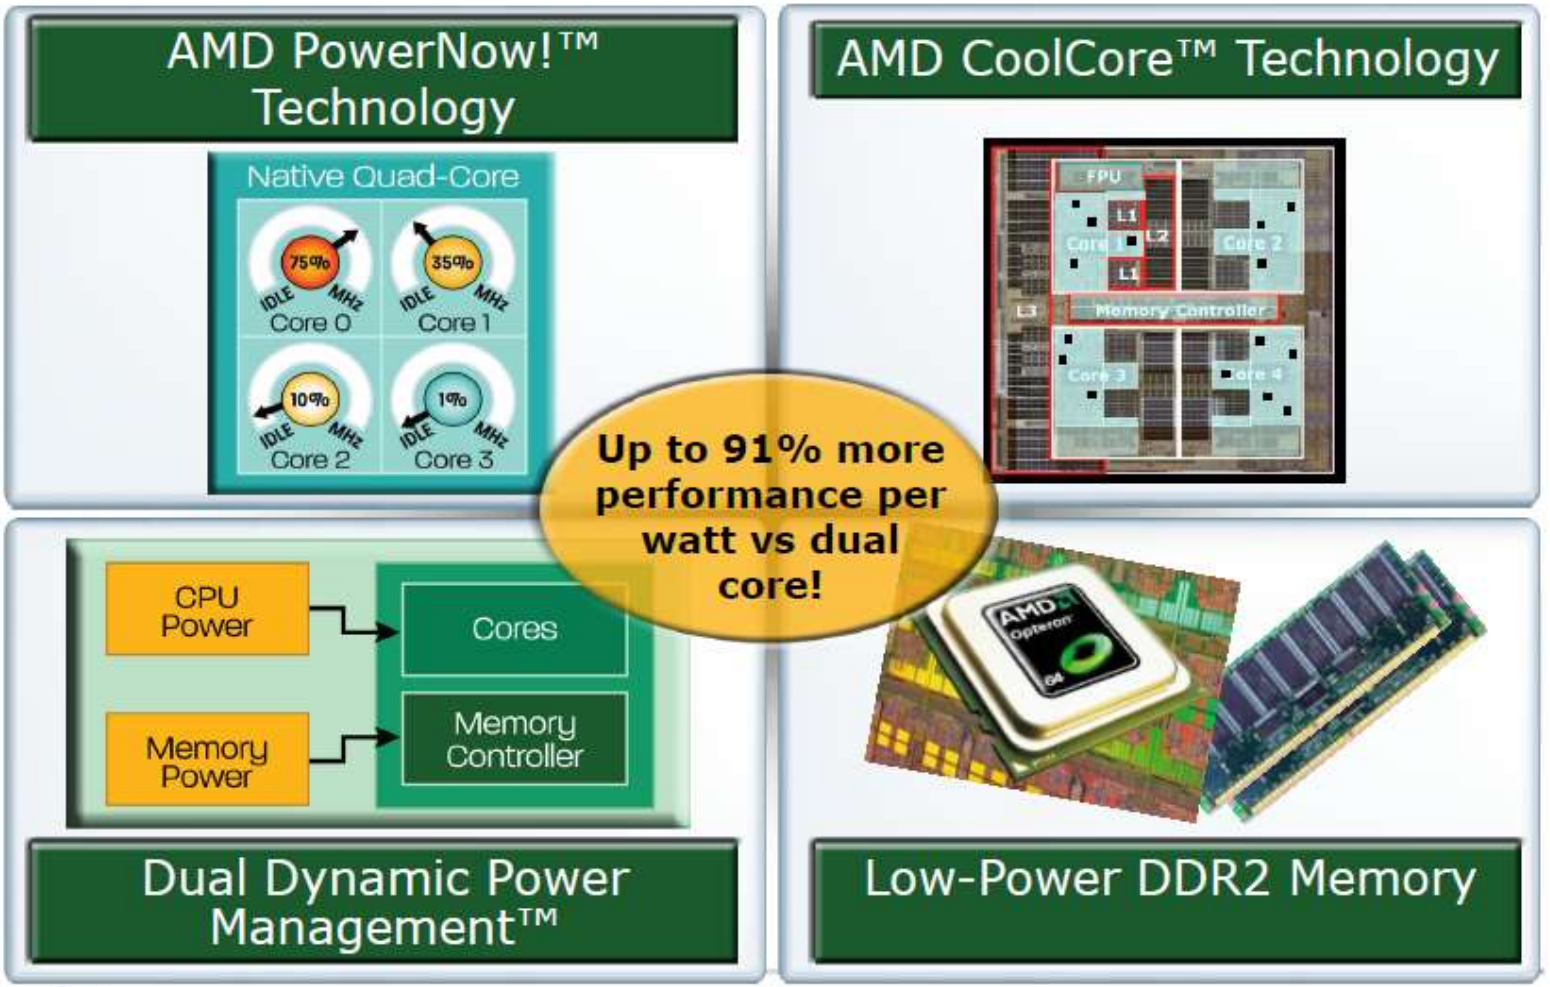

AMD The future is fusion

g1) AMD CoolCore Technology [40]

It turns off not used functional units

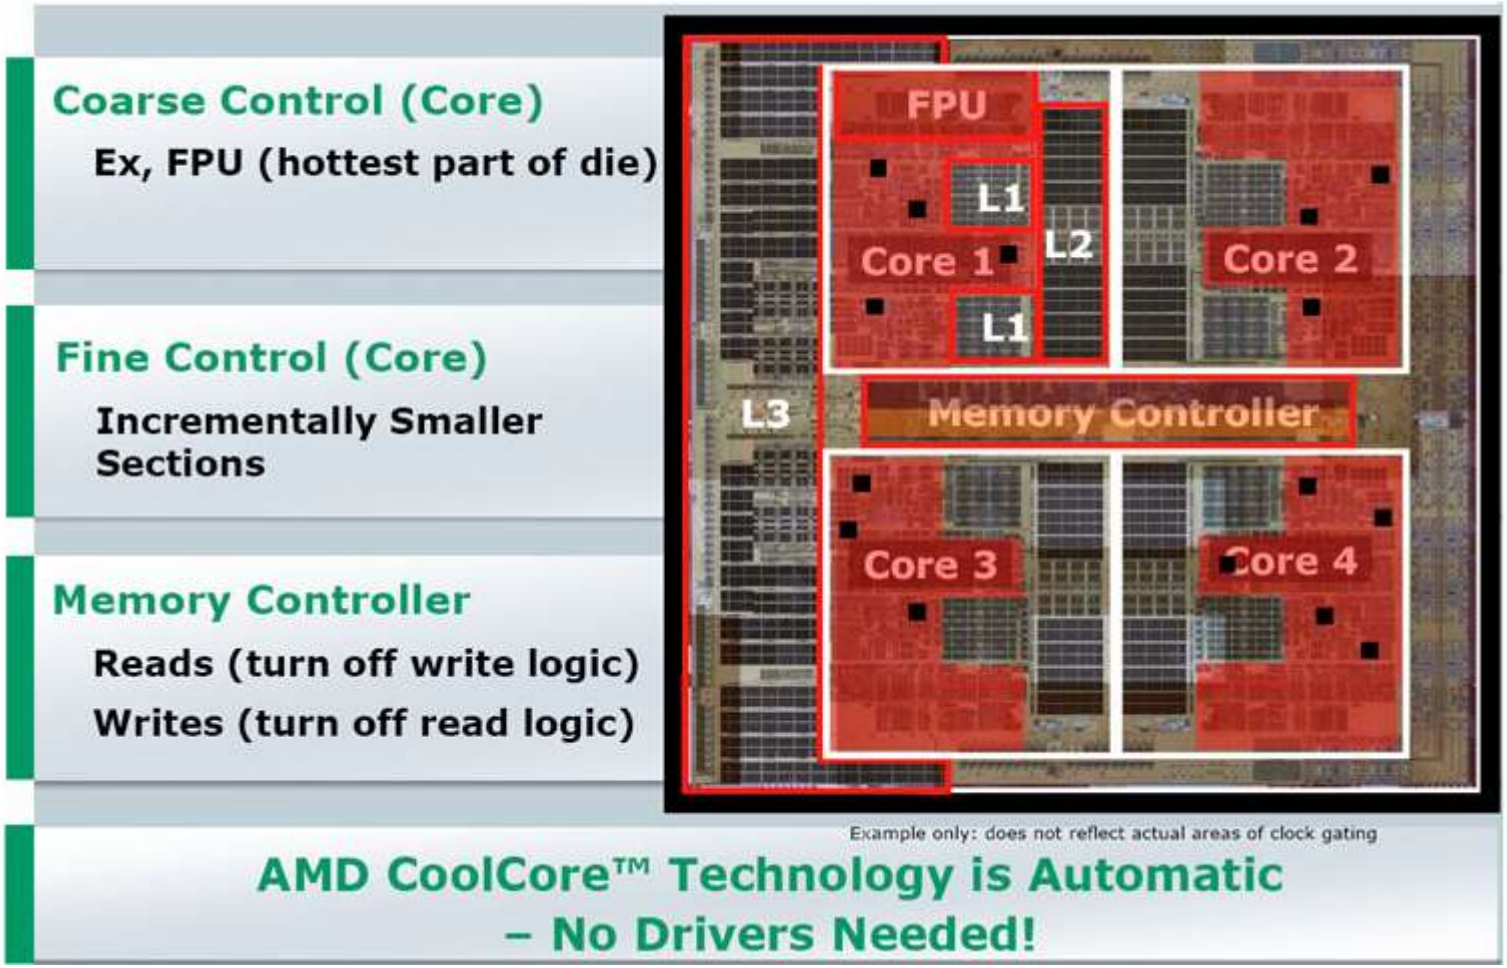

6.2 Main enhancements of K10 servers (15)

g2) Dual Dynamic Power Management

- Separate power planes (split planes) for cores and the memory controller,
- Enables cores to operate at reduced power consumption level while memory controller continues to run at full speed,
- Memory controller can operate at higher frequency for increased bandwidth and performance (see next slide).

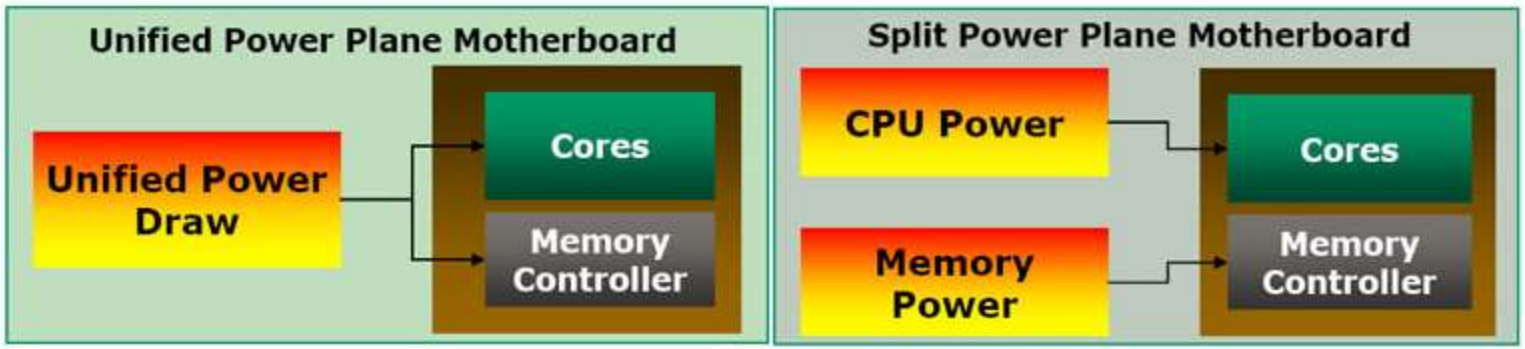

Figure: Contrasting single (unified) and split power plane motherboards [40]

Remark

Dual Dynamic Power Management requires the use of Socket F revision 2 (Socket Fr2) or higher (the Socket of K10.5 DP/MP servers) and motherboard support dual (isolated voltage supplies).

6.2 Main enhancements of K10 servers (16)

Example: Increased memory clock achieved by Dual Dynamic Power Management [40]

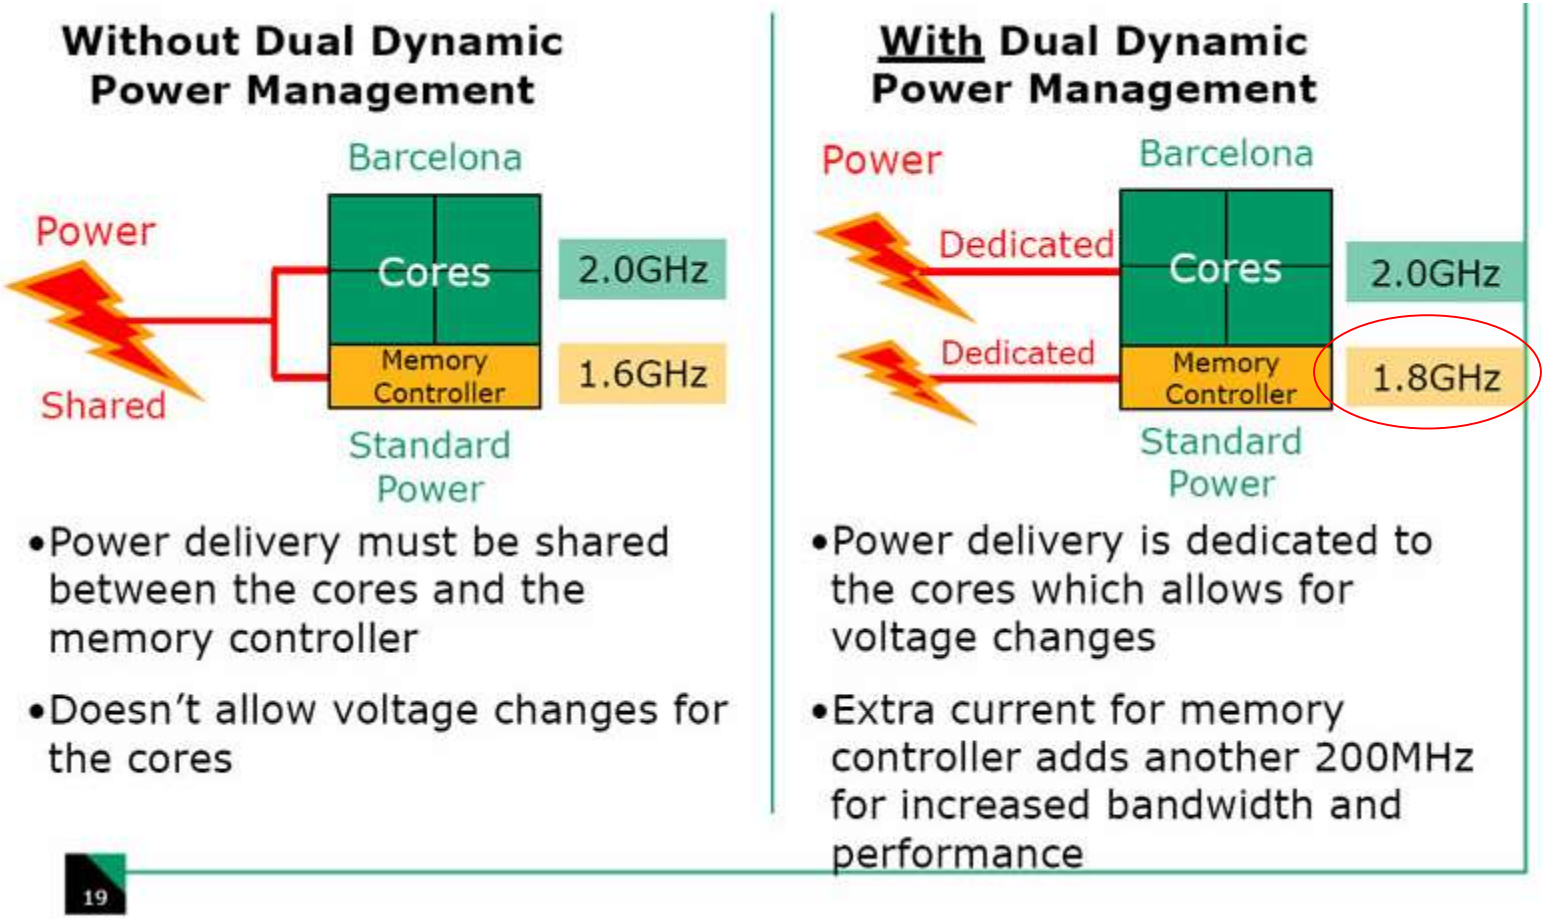

6.2 Main enhancements of K10 servers (17)

g3) Independent Dynamic Core Technology [127]

- Core frequencies are separately adjusted according to their level of utilization.
  This saves power on less active cores.
- All cores share the same power plane, the shared voltage is determined according to the highest P-state of the cores.

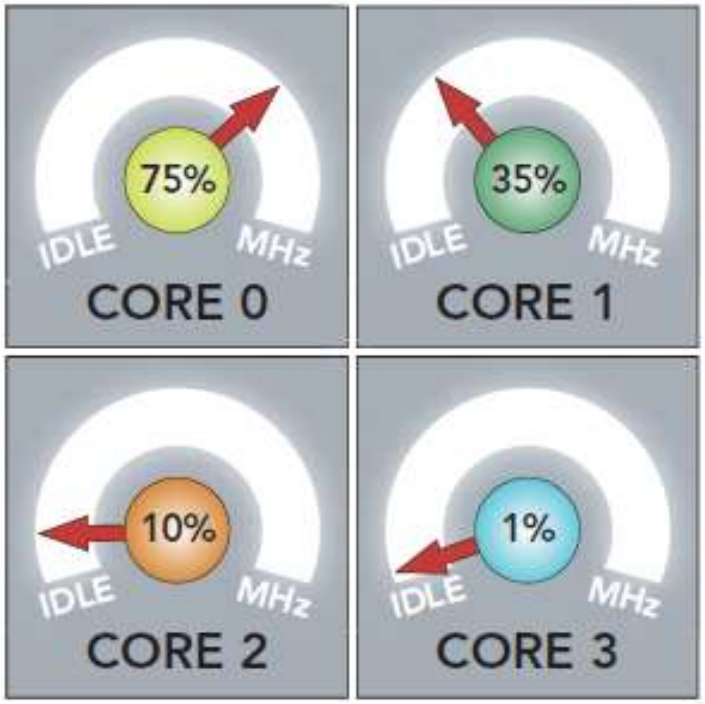

Independent Dynamic Core Technology is part of AMD's power management technology implemented for K10 servers and desktops, as outlined next.

6.2 Main enhancements of K10 servers (18)

Overview of AMD's P-state management in K10-based desktops and servers [135] **(1)**

- It provides dynamic voltage and frequency scaling (DVFS)
  (supported first in the Athlon 64 line (2003).
- Dynamic voltage scaling is performed in the Working state.
  It is implemented
  - either by a single step voltage ramping (called voltage slamming)
    if the voltage regulator provides built-in output-voltage slew rate control,
  - else by voltage stepping [135].
- Dynamic frequency scaling is performed also in the Working state.

6.2 Main enhancements of K10 servers (19)

Accomplishing dynamic frequency scaling (frequency transitions) in K10 processors

Up to K8-based processors frequency transitions in AMD processors were performed in the Halt state, as indicated in the next Figure.

6.2 Main enhancements of K10 servers (20)

Accomplishing frequency transitions in processors

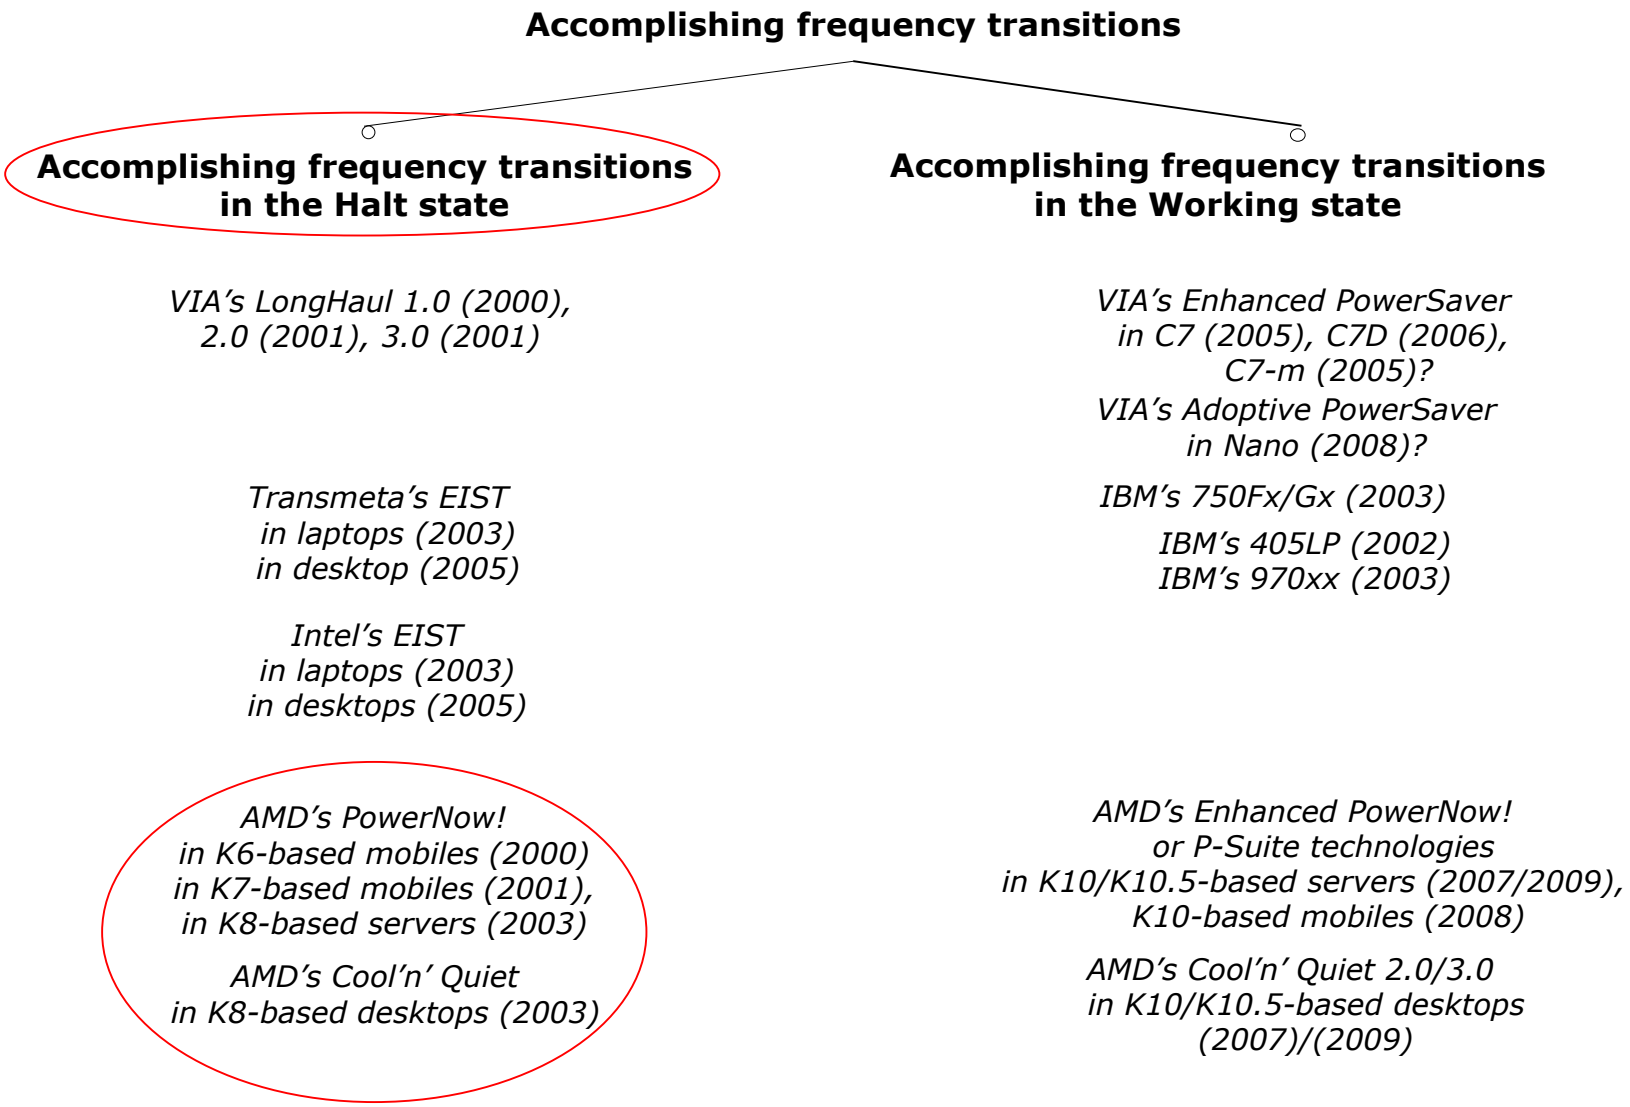

6.2 Main enhancements of K10 servers (21)

The principle of accomplishing frequency transitions in the Halt state

- before modifying the clock frequency the clocking was stopped and the PLL was shut down
- then the PLL was set to the required value and
- after the settling time of the PLL (about nx10 μs) clocking could be restarted.

This mechanism causes an efficiency degradation during each P-state transition.

For this reason, starting with K10 based servers and desktops the clock generator was redesigned and frequency transitions became performed in the working state without halting the PLL.

6.2 Main enhancements of K10 servers (22)

Accomplishing frequency transitions in processos

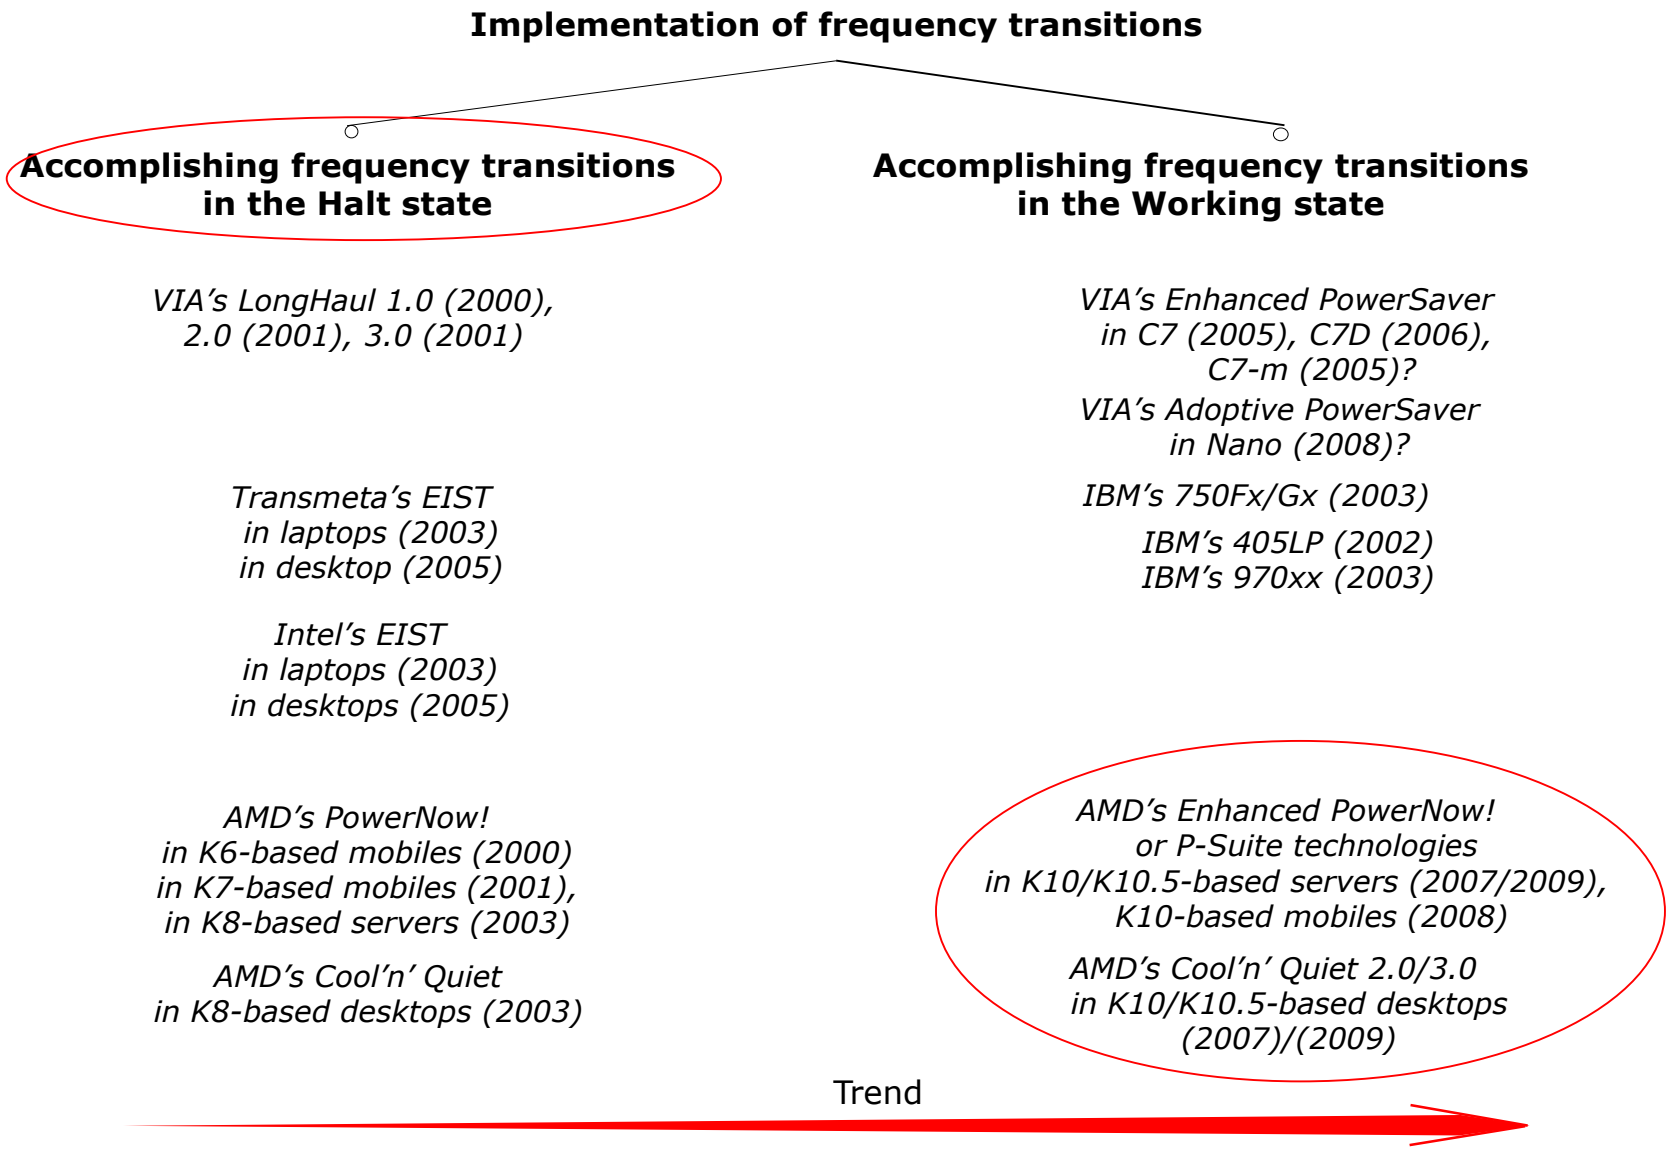

6.2 Main enhancements of K10 servers (23)

Basic layout of the clock generator in K10 servers

- Separate PLLs for each core.
- Frequency transitions achieved by modifying the clock multiplier ni.

⟹ This causes a PLL lock time of about 16 µs [135].

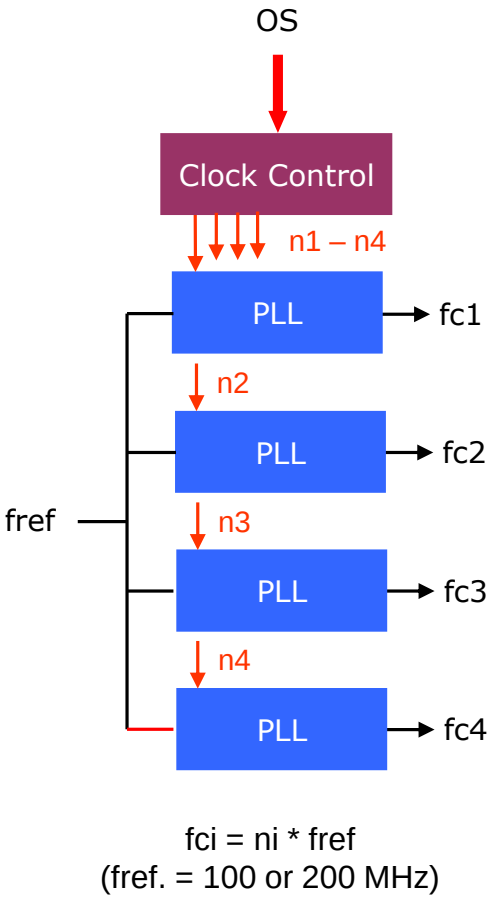

Figure: Basic layout of the clock generator in K10 servers

6.2 Main enhancements of K10 servers (24)

Coordination of P-state transitions of the cores

- Core frequencies can be controlled independently from each other.
- To the contrary, core voltages are the same for all cores, where the common core voltage is the highest one needed for the highest core frequency.
  Dedicated hardware of the processor performs the required P-state coordination.

6.3 Contrasting utilized and implemented features of the microarchitecture of K10 servers

6.3 Contrasting utilized and implemented features of the microarchitecture of K10 servers

In fact, the underlying microarchitecture of K10 servers has much more functional capabilities as utilized in them.

The reason for this is AMD's decision that DP/MP servers should use the same sockets as previously introduced for the K8 NPT servers, as shown in the next Figure.

6.3 Contrasting utilized and implemented features of the microarchitecture of K10 servers (2)

Sockets of K10 DP, MP servers

[based on 42]ery first re

Supports DDR3

Socket **G34** (1944 pins, MP servers K10.5 „Magny-Course", Fam. 15h)
Socket **C32** (1207 pins, DP servers K10.5 „Magny –Course, Fam. 15h)
Socket **AM3** (941 pins, K10.5 Shanghai UP servers, desktops, except first released desktops, K10.5 Istambul desktops)
Socket **AM3+** (942 pins, UP servers Fam. 15h, desktops Fam. 15h)
Socket **FM1** (905 balls BGA, desktops, mobiles Fam. 12h)
Socket **S1g4** (638 pins, K10.5 mobile)
Socket **FT1** (413 balls BGA, desktops, mobiles Fam. 14h)
Socket **FT3** (769 balls BGA, desktops, mobiles Fam. 16h)

S:: Servers
DT: Desktop
M: Mobile
UPT: Ultraportable

SPP: Single power Plane
DPP: Dual Power planes
HT 2.0: HyperTransport 2.0
HT 3.0: HyperTransport 3.0

K10.5 | **Fam. 15h** | **Fam. 15h** | **Fam. 15h**

45nm 2x6C „Magny-Course" | 32 nm 2x8C „Bulldozer" | 32 nm 2x8C/4C APU „Piledriver" | 28 nm 2x8C „Steamroller"

3st Generation Platform

45 nm 4C, DT/M „Shanghai" | 45 nm 6C, DT „Istambul" | 32 nm 4C APU „Llano" | 32 nm 2C APU „Bobcat" | 28 nm 4C APU „Jaguar" | 28 nm 4C APU „Puma+"

Fam. 12h | **Fam. 14h** | **Fam. 16h** | **Fam. 16h**

K8 (NPT) | **K10** | **K10.5** | **K10.5**

90 nm 2C NPT „Santa Rosa" | 65 nm 4C „Barcelona" | 45 nm 4C, S „Shanghai" | 45 nm 6C, S „Istambul"

2st Generation Platform

65 nm 4C „Brisbana" | 65 nm 2C „Griffin"

K8 | **Fam. 11h**

K8 | **K8** | **K8**

130 nm 1C „SledgeH." | 90 nm 1C „Athens" | 90 nm 2C „Egypt"

1st Generation Platform

Supports DDR2

Socket **F** (1207 pins, SPP, HT 2.0: DP, MP servers K8 NPT, K10)
Socket **F+** (1207 pins, DPP, HT 3.0: DP, MP servers K10.5)
Socket **AM2** (940 pins, SPP, HT 2.0: K8 NPT UP serv., K8 NPT, Brisbane DTs)
Socket **AM2+** (940 pins, DPP: K10 UP servers, K10 desktops, first K10.5 Shanghai desktops)
Socket **S1g1** (638 pins, K8 NPT, Brisbane mobiles)
Socket **S1g2** (638 pins, Fam. 11h mobiles)
Socket **S1g3** (638 pins, K10.5 Shanghai (Caspian die) mobiles)
Socket **ASB1** (812 balls BGA, Fam. 11h (Huron, Conesus dies) UPT mobiles)
Socket **ASB2** (812 balls BGA, K10.5 (Shanghai) (Geneva die), UPT mobiles)

Supports DDR

Socket **940** (130nm K8 UP/DP/MP servers, 90 nm K8 SC and 90 nm K8 DC (Egypt) DP/MP servers)
Socket **939** (90nm K8 SC and 90 nm K8 DC (Egypt) UP servers, K8 desktops with 128-bit mem. channel up to the 90 nm DC (Egypt) lines)
Socket **754** (K8 desktops, mobiles with 64-bit mem. channel up to the 90 nm SC (Athens) line)

2003 2004 2005 2006 2007 2008 2009 2010 2011 2012 2013 2014

6.3 Contrasting utilized and implemented features of the microarchitecture of K10 servers (3)

Enhanced features implemented on the K10 Barcelona server die but not utilized in K10 Barcelona server lines

Relating memory and HT links AMD implemented three innovative features on the K10 Barcelona server die, as follows:

- supporting DDR3 memory,
- providing HyperTransport 3.0 links and
- implementing four HypertTransport links [55].

However, along with the introduction of K10 DP and MP servers AMD made further on use of the Socket F introduced previously for the K8 NPT line in order to remain compatible.

Socket F however, supported only

- DDR2 memory,
- HyperTransport 2.0 links,
- three HypertTransport links and
- a single power plane (restricting the use of Dual Dynamic Power Management).

As a consequence, K10 DP/MP servers equipped with the Socket F support only the restricted features, listed above.

Releasing the restrictions of K10 DP/MP servers caused by sticking to Socket F in subsequent K10.5 DP/MP lines-1

The K10.5 Shanghai and Istambul lines replaced Socket F by Socket F+ (Socket Fr2).

This Socket allowed already the use of HT 3.0 but neither DDR3 memory nor four HT 3.0 were supported.

6.3 Contrasting utilized and implemented features of the microarchitecture of K10 servers (5)

Sockets of AMD's x86-64 processors [based on 42]

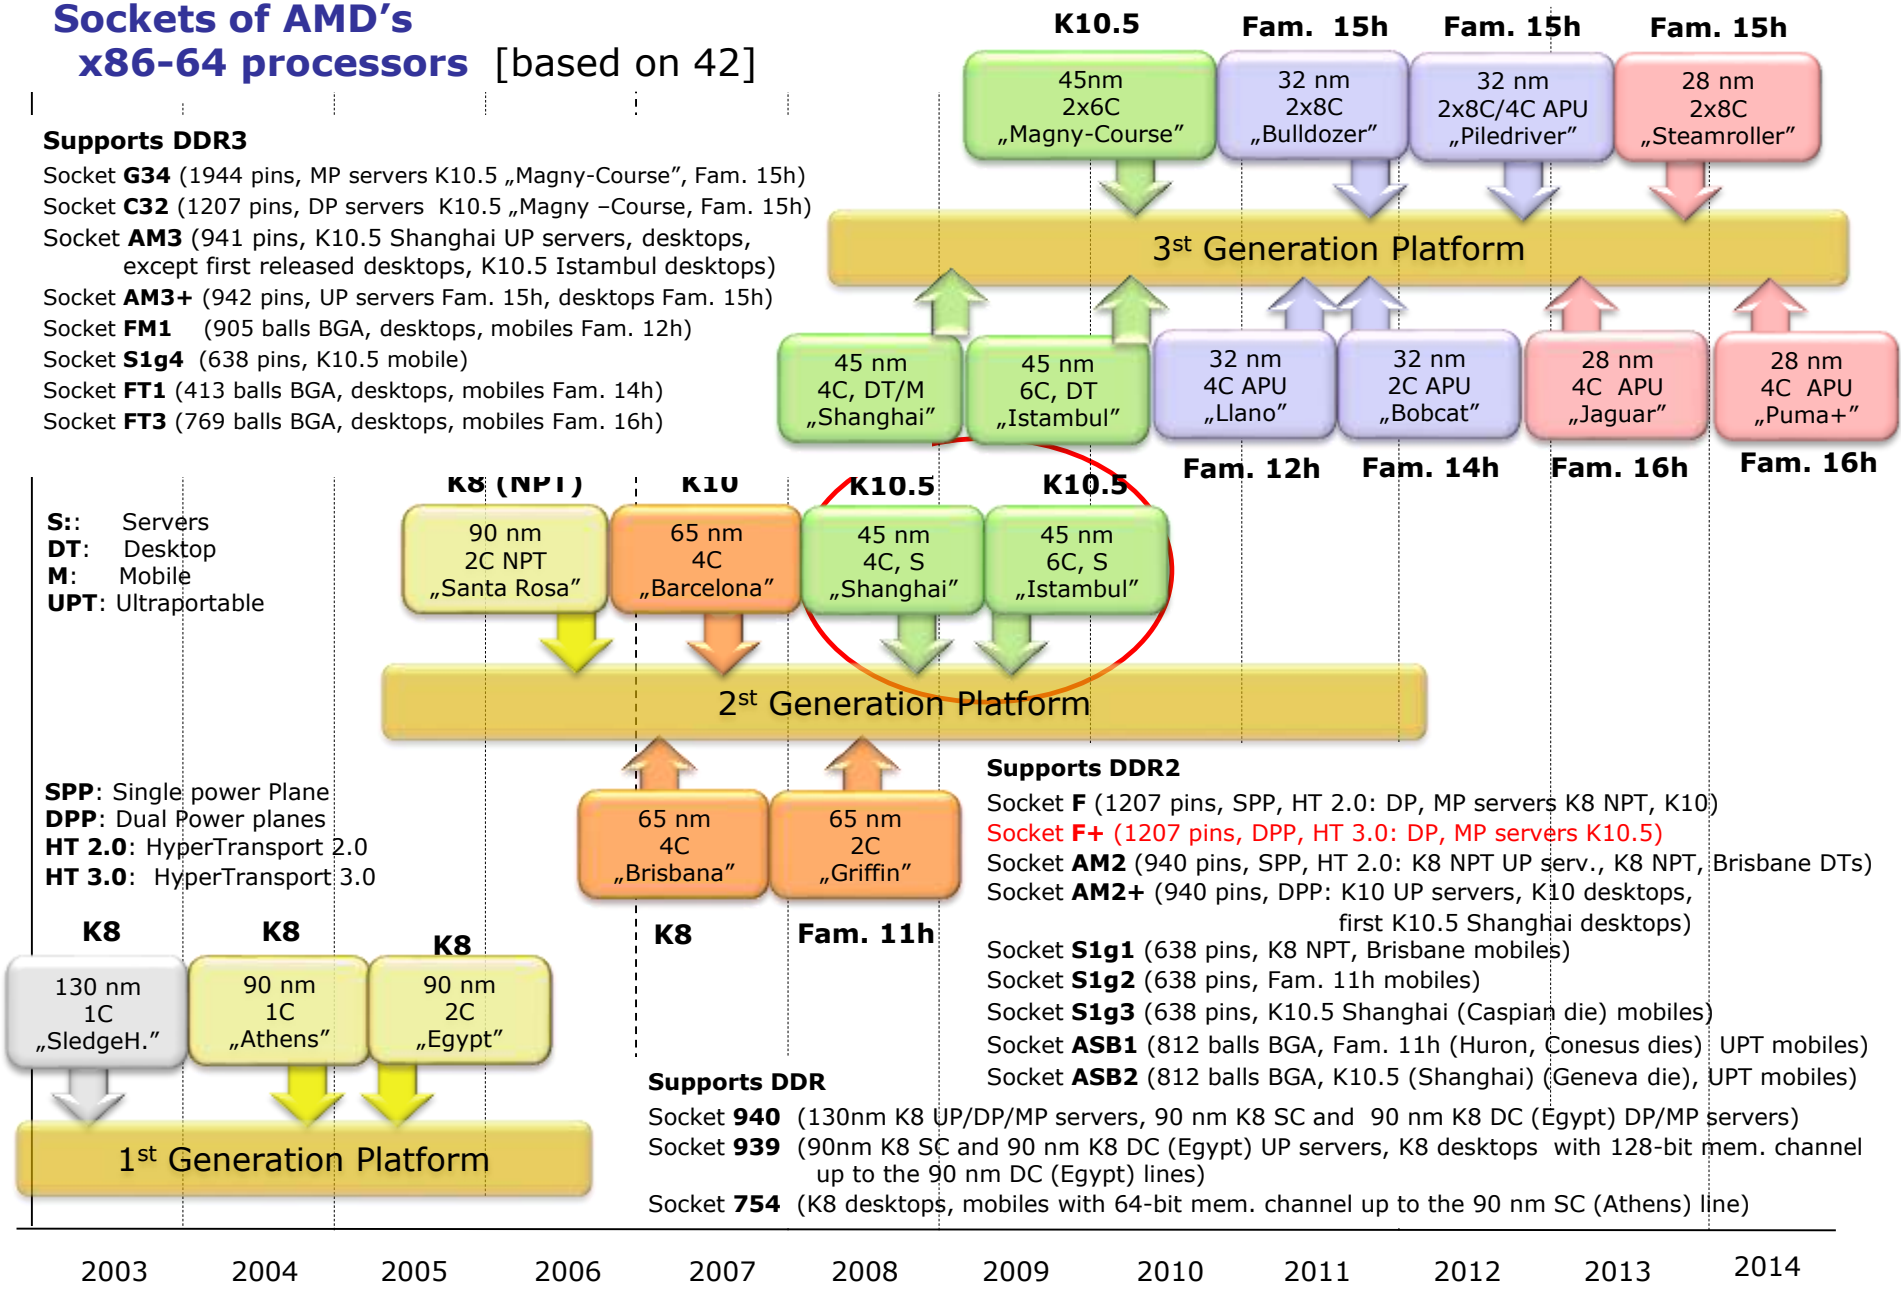

Releasing the restrictions of K10 DP/MP servers caused by sticking to Socket F in subsequent K10.5-based DP/MP lines-2

Finally, the K10.5 based dual-chip Magny-Course MP server switched to the G34 Socket.

With this Socket all previous restrictions were lifted and the Magny-Course MP server provided support for DDR3 memory and four HT 3.0 links.

6.3 Contrasting utilized and implemented features of the microarchitecture of K10 servers (7)

ckets of AMD's 86-64 processors [based on 42]

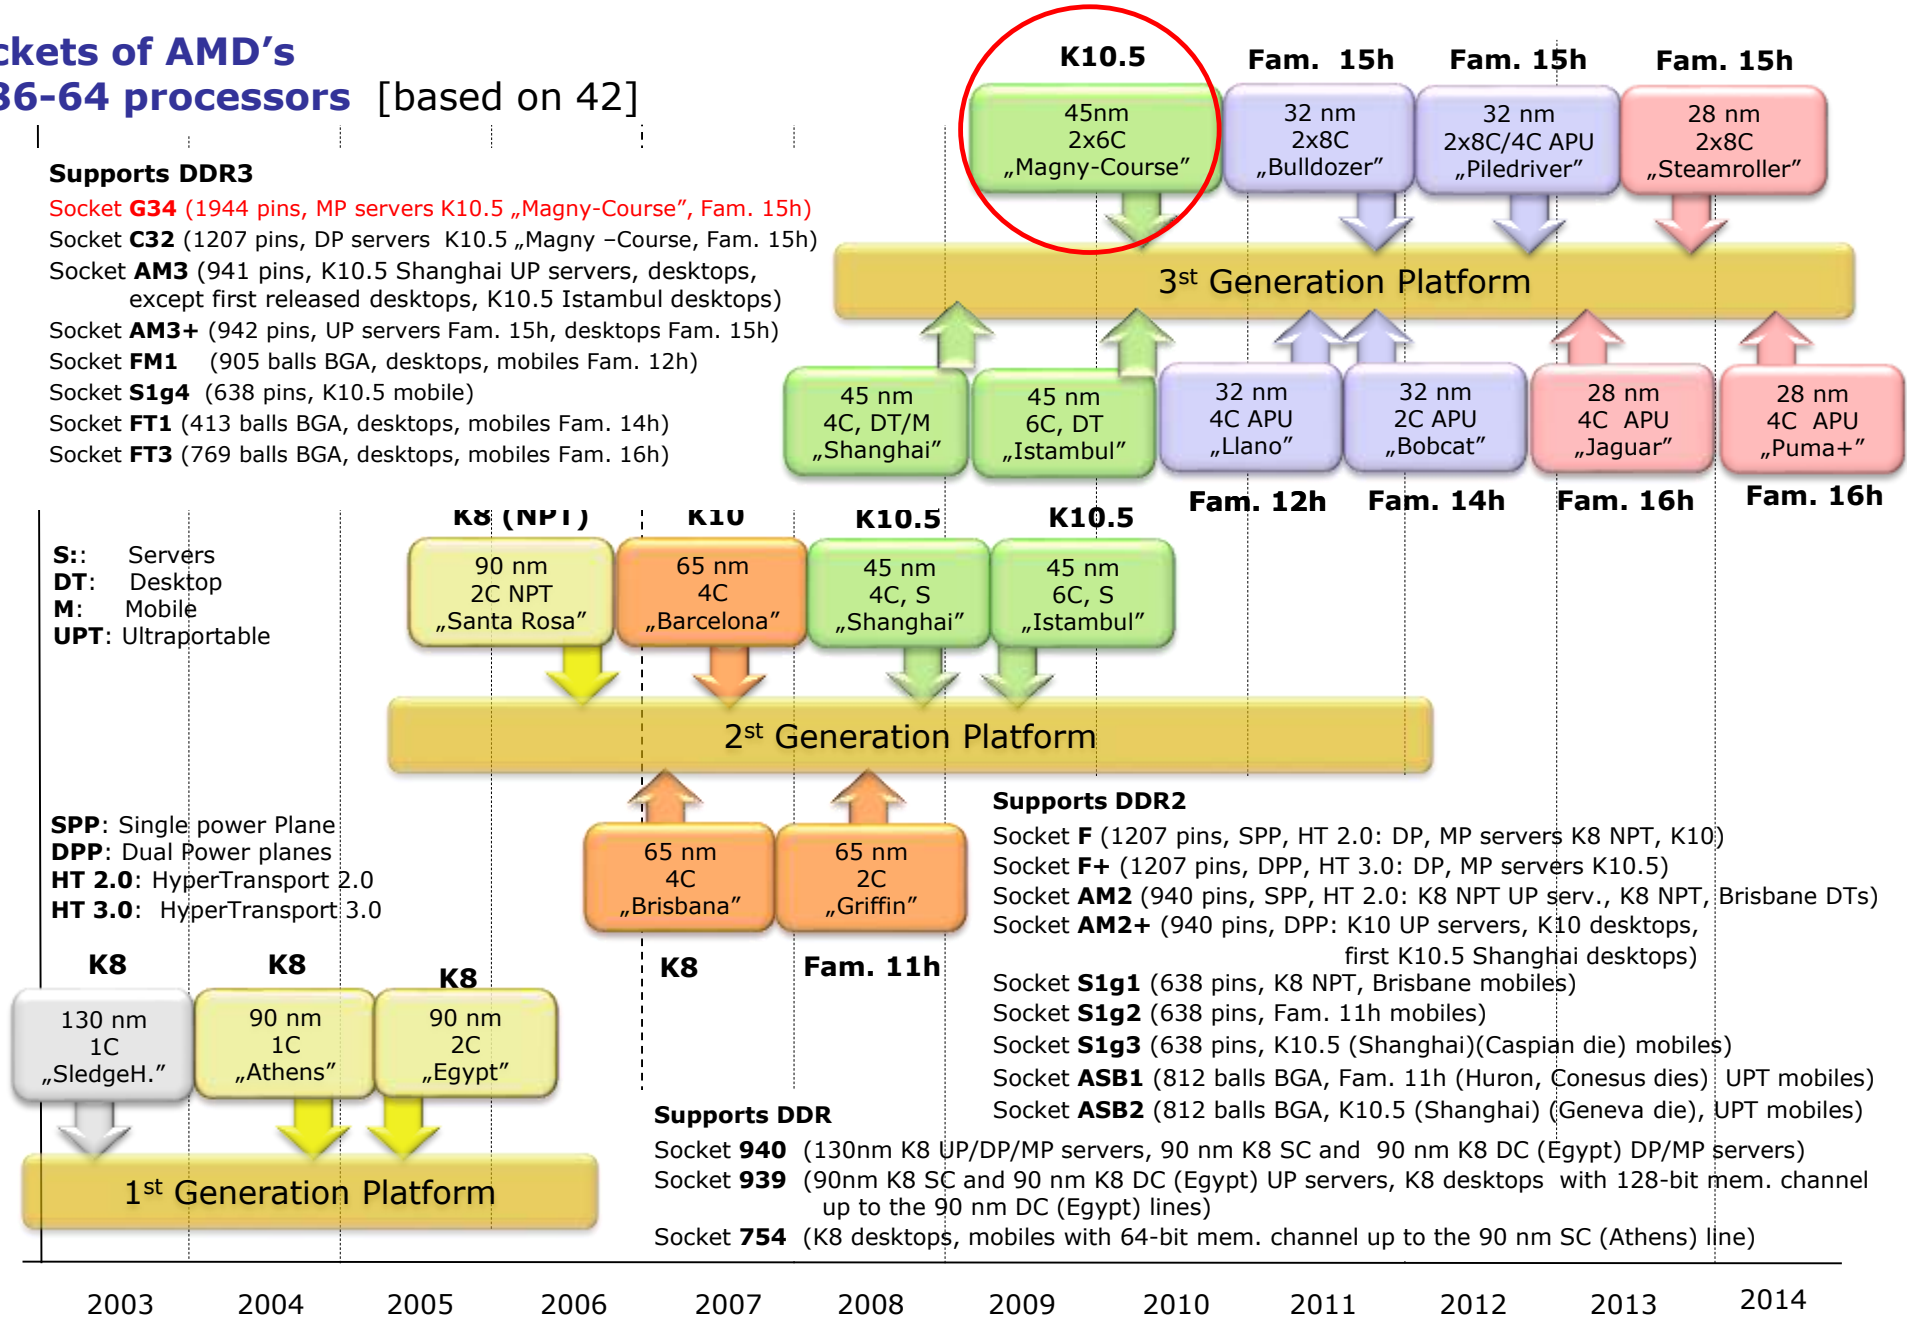

Remark to the implementation of the HyperTransport unit [55]

The chosen implementation includes dual HyperTransport paths, one for HT 1 and HT2 and another one for HT 3.0, as indicated for the HT receiver unit on the next Figure.

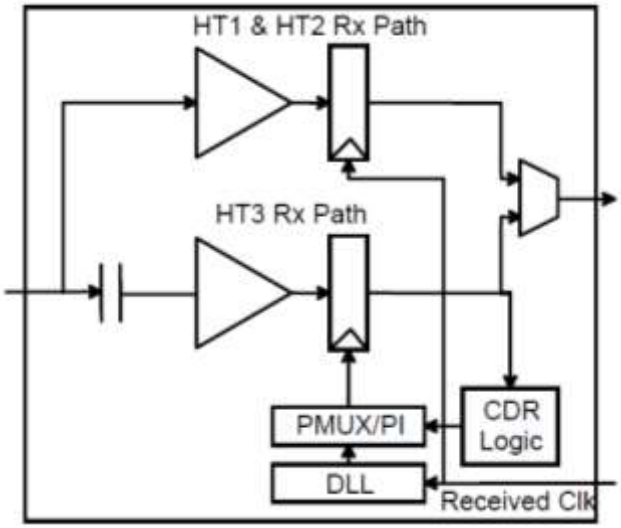

Figure: The receiver part of the HyperTransport transceiver unit [55]

The HT unit operates such that in legacy mode the related legacy path and in HT 3.0 mode the associated HT 3.0 path is activated.

In legacy mode the processor supports HT clock frequencies up to 1.0 GHz whereas in HT 3.0 mode up to 1.3 GHz.

In K10 server chips the legacy mode is activated whereas in subsequent K10.5 parts (Shanghai lines etc.) and in K10 desktop chips from the beginning on the HT 3.0 mode.

Supporting HT 3.0 in K10 UP servers

In contrast to DP and MP servers, K10 UP servers (designated as Budapest that were released about 8 months later than the K10 DP/MP (Barcelona servers) were already equipped with the Socket AM2+.

Socket AM2+ supported already HT 3.0 and dual power planes, needed for Dual Dynamic Power management.

Subsequently, K10.5 UP servers (of the Shanghai family, designated as Suzuka) and K10.5 desktops (both of the Shanghai and the Istambul family) got already the Socket AM3 and thus DDR3 support.

Sockets of AMD's x86-64 processors [based on 42]

Supports DDR3

Socket **G34** (1944 pins, MP servers K10.5 „Magny-Course", Fam. 15h)
Socket **C32** (1207 pins, DP servers K10.5 „Magny –Course, Fam. 15h)
Socket **AM3** (941 pins, K10.5 Shanghai UP servers, desktops, except first released desktops, K10.5 Istambul desktops)
Socket **AM3+** (942 pins, UP servers Fam. 15h, desktops Fam. 15h)
Socket **FM1** (905 balls BGA, desktops, mobiles Fam. 12h)
Socket **S1g4** (638 pins, K10.5 mobile)
Socket **FT1** (413 balls BGA, desktops, mobiles Fam. 14h)
Socket **FT3** (769 balls BGA, desktops, mobiles Fam. 16h)

S:: Servers
DT: Desktop
M: Mobile
UPT: Ultraportable

SPP: Single power Plane
DPP: Dual Power planes
HT 2.0: HyperTransport 2.0
HT 3.0: HyperTransport 3.0

K10.5 | Fam. 15h | Fam. 15h | Fam. 15h
45nm 2x6C „Magny-Course" | 32 nm 2x8C „Bulldozer" | 32 nm 2x8C/4C APU „Piledriver" | 28 nm 2x8C „Steamroller"

3st Generation Platform

45 nm 4C, DT/M „Shanghai" | 45 nm 6C, DT „Istambul" | 32 nm 4C APU „Llano" | 32 nm 2C APU „Bobcat" | 28 nm 4C APU „Jaguar" | 28 nm 4C APU „Puma+"
Fam. 12h | Fam. 14h | Fam. 16h | Fam. 16h

K8 (NPT) | K10 | K10.5 | K10.5
90 nm 2C NPT „Santa Rosa" | 65 nm 4C „Barcelona" | 45 nm 4C, S „Shanghai" | 45 nm 6C, S „Istambul"

2st Generation Platform

65 nm 4C „Brisbana" | 65 nm 2C „Griffin"
K8 | Fam. 11h

K8 | K8 | K8
130 nm 1C „SledgeH." | 90 nm 1C „Athens" | 90 nm 2C „Egypt"

1st Generation Platform

Supports DDR2

Socket **F** (1207 pins, SPP, HT 2.0: DP, MP servers K8 NPT, K10)
Socket **F+** (1207 pins, DPP, HT 3.0: DP, MP servers K10.5)
Socket **AM2** (940 pins, SPP, HT 2.0: K8 NPT UP serv., K8 NPT, Brisbane DTs)
Socket **AM2+** (940 pins, DPP: K10 UP servers, K10 desktops, first K10.5 Shanghai desktops)
Socket **S1g1** (638 pins, K8 NPT, Brisbane mobiles)
Socket **S1g2** (638 pins, Fam. 11h mobiles)
Socket **S1g3** (638 pins, K10.5 (Shanghai)(Caspian die) mobiles)
Socket **ASB1** (812 balls BGA, Fam. 11h (Huron, Conesus dies) UPT mobiles)
Socket **ASB2** (812 balls BGA, K10.5 (Shanghai) (Geneva die), UPT mobiles)

Supports DDR

Socket **940** (130nm K8 UP/DP/MP servers, 90 nm K8 SC and 90 nm K8 DC (Egypt) DP/MP servers)
Socket **939** (90nm K8 SC and 90 nm K8 DC (Egypt) UP servers, K8 desktops with 128-bit mem. channel up to the 90 nm DC (Egypt) lines)
Socket **754** (K8 desktops, mobiles with 64-bit mem. channel up to the 90 nm SC (Athens) line)

2003 2004 2005 2006 2007 2008 2009 2010 2011 2012 2013 2014

Die plot of the K10 Barcelona server chip [55]

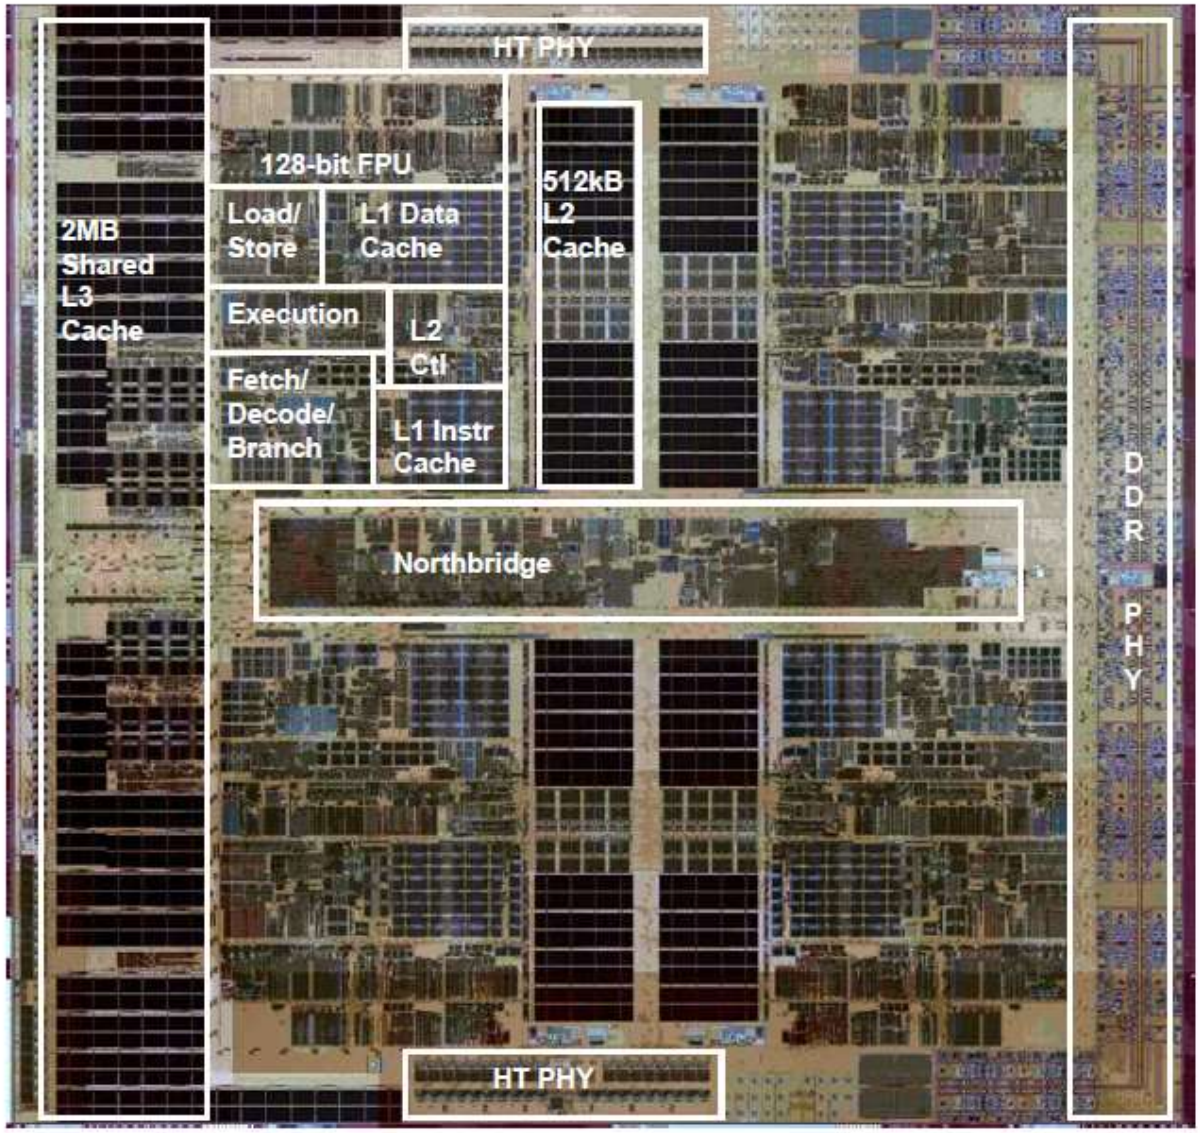

6.4 K10 (Barcelona) server lines

6.4 K10 (Barcelona) server lines

- They were introduced in 8/2007.
- They include models for UP, DP and MP servers, as indicated below.

Model numbers of the K10 Barcelona server lines	
Server type	Model numbers
UP servers	Opteron 135x - 036x
DP servers	Opteron 234x – 236x
MP servers	Opteron 834x – 836x

6.4 K10 (Barcelona) server lines (2)

Brand names of AMD's K10-based server lines

		2003-2007	2007-2008	2008-2011	2009	2009
		K8 (Hammer)	**K10 (Barcelona)**	**K10.5 (Shanghai)**	**K10.5 (Istanbul)**	**K10.5 (Magny- Course)**
Servers	**4P servers**	**See Section 4**	**Barcelona (834x-836x))**	**Shanghai (837x-839x)**	**Istambul (8410-8430)**	**Magny-Course (6100)**
	2P servers		**Barcelona (234x-236x)**	**Shanghai (237x-239x)**	**Istambul (241x-243x)**	**Lisbon (4100)**
	1P servers		**Budapest (135x-136x)**	**Suzuka (138x-139x)**		
Desktops	**High perf.** (~80-120W)		Phenom X4-X2	Phenom II X4-X2	Phenom II X6-X4	
	Mainstream (~60-90W)	Athlon 64 Athlon 64 X2	Athlon X2	Athlon II X4-X2		
	Value (~40-60W)	Sempron		Sempron		
Mobiles	**High perf.** (~30-40W)	Turion 64 X2 (TL 6/5) Turion 64 (ML/MT)		Phenom II (N/P 9xx-6xx) Turion II Ultra (M6xx) Turion II (M/N/P 5xx)		
	Mainstream (~20-30W)	Athlon 64 X2 (TK-5x/4x) Athlon 64 (2xxx+-4xxx+)		Athlon II (M/N/P 3xx) Sempron (M1xx)		
	Ultraportable (~10-20W)	Mobile Sempron (2xxx+-4xxx+) Sempron 2100 fanless		Turion II Neo (K6xx) Athlon II Neo (K1xx) V-series (V1xx)		
Embedded (~10-20W)				Turion II Neo X2 Athlon II Neo X2 Athlon II Neo		

6.4 K10 (Barcelona) server lines (3)

Main features of the K10 server lines [149]

MODEL NUMBER	CORE FREQUENCY	I/O BUS FREQUENCY*	MAX I/O BANDWIDTH	SOCKET	CMOS TECH	L2 CACHE	L3 CACHE	ACP***
8360 SE**	2.5 GHz	1000 MHz	24 GB/s	F (1207)	65 nm SOI	512 KB/core	2MB	105W
2360 SE**	2.5 GHz	1000 MHz	24 GB/s	F (1207)	65 nm SOI	512 KB/core	2MB	105W
8358 SE**	2.4 GHz	1000 MHz	24 GB/s	F (1207)	65 nm SOI	512 KB/core	2 MB	105W
2358 SE**	2.4 GHz	1000 MHz	24 GB/s	F (1207)	65 nm SOI	512 KB/core	2 MB	105W
8356	2.3 GHz	1000 MHz	24 GB/s	F (1207)	65 nm SOI	512 KB/core	2 MB	75W
2356	2.3 GHz	1000 MHz	24 GB/s	F (1207)	65 nm SOI	512 KB/core	2 MB	75W
1356**	2.3 GHz	2000 MHz	16 GB/s	AM2	65 nm SOI	512 KB/core	2 MB	75W
8354	2.2 GHz	1000 MHz	24 GB/s	F (1207)	65 nm SOI	512 KB/core	2 MB	75W
2354	2.2 GHz	1000 MHz	24 GB/s	F (1207)	65 nm SOI	512 KB/core	2 MB	75W
1354**	2.2 GHz	1800 MHz	14.4 GB/s	AM2	65 nm SOI	512 KB/core	2 MB	75W
2352	2.1 GHz	1000 MHz	24 GB/s	F (1207)	65 nm SOI	512 KB/core	2 MB	75W
1352**	2.1 GHz	1800 MHz	14.4 GB/s	AM2	65 nm SOI	512 KB/core	2 MB	75W
8350	2.0 GHz	1000 MHz	24 GB/s	F (1207)	65 nm SOI	512 KB/core	2 MB	75W
2350	2.0 GHz	1000 MHz	24 GB/s	F (1207)	65 nm SOI	512 KB/core	2 MB	75W
8347 HE**	1.9 GHz	1000 MHz	24 GB/s	F (1207)	65 nm SOI	512 KB/core	2 MB	55W
2347 HE**	1.9 GHz	1000 MHz	24 GB/s	F (1207)	65 nm SOI	512 KB/core	2 MB	55W
8346 HE**	1.8 GHz	1000 MHz	24 GB/s	F (1207)	65 nm SOI	512 KB/core	2 MB	55W
2346 HE**	1.8 GHz	1000 MHz	24 GB/s	F (1207)	65 nm SOI	512 KB/core	2 MB	55W

6.4 K10 (Barcelona) server lines (4)

Note

We point out that in the above table AMD specified power consumption by a new measure, called the ACP value.

AMD introduced ACP along with the K10 Barcelona server line.

They understand ACP as the Average CPU Power that is the power consumption when a processor runs a suite of server workloads that represents the breadth of typical server application, like TPC-C, SPECcpu2006, SPECjbb2005 and STREAM.

The geometric mean of measurements, taken during running these workloads, is the ACP [136].

The ACP value may be relevant for data centers to specify power supply and cooling .

By contrast, the TDP (Thermal Design Power) characterizes the power consumption while running power intensive applications.

It should be emphasized that TDP is not the worst case power consumption of a processor that can be drawn e.g. while running a "power virus" but it is the power design point for the cooling system.

TDP is the key input parameter for designing the cooling system of a computer.

The cooling system has to be designed such that for a power consumption of TDP the junction temperature of the processor chip should remain below a given value.

6.4 K10 (Barcelona) server lines (5)

We note that ACP is typically lower than TDP.

A comparison of ACP and TDP values is given in [137].

Accordingly, comparable ACP and TDP values are e.g. as follows:

40	Watt ACP	=	60 Watt TDP
55	Watt ACP	=	79 Watt TDP
75	Watt ACP	=	115 Watt TDP
105	Watt ACP	=	137 Watt TDP

6.5 Main enhancements of K10 (Barcelona) desktops

6.5 Main enhancements of K10 (Barcelona) desktops

K10 desktops

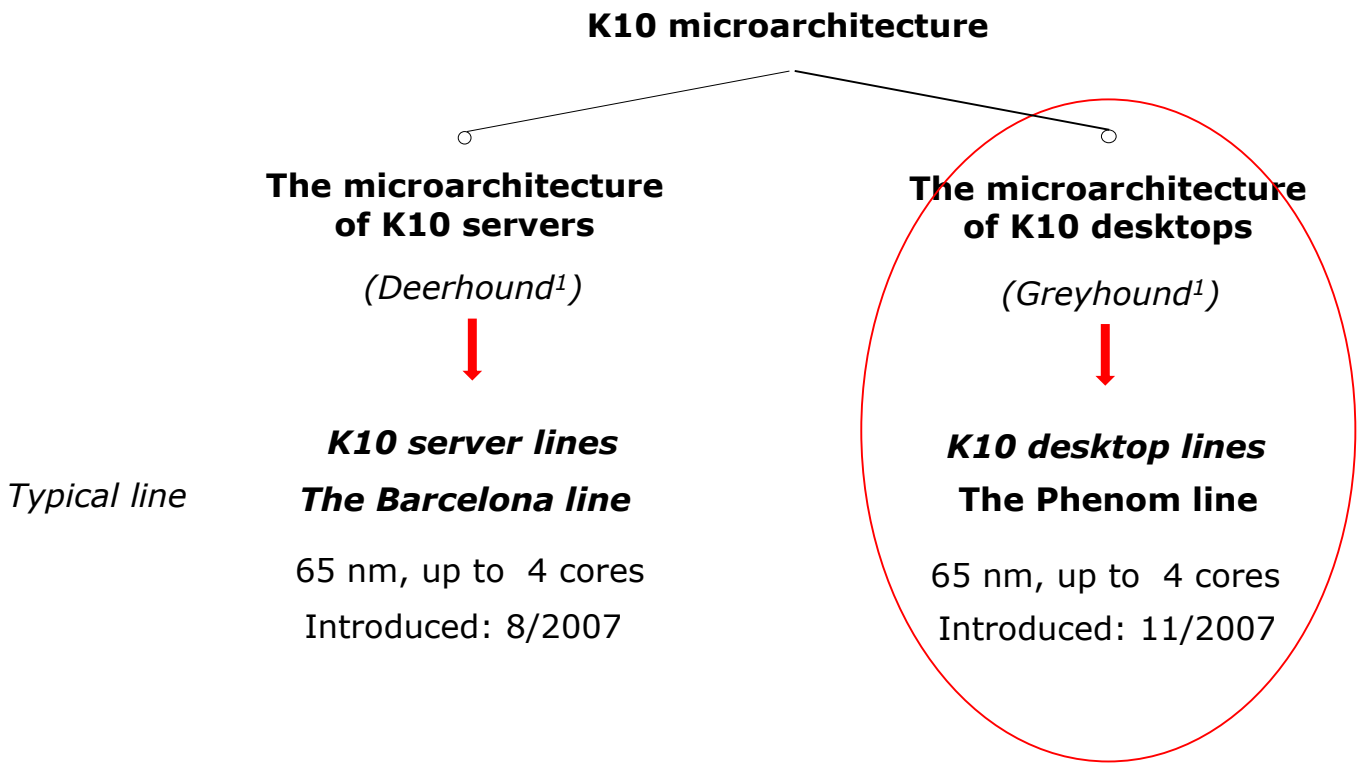

6.5 Main enhancements of K10 (Barcelona) desktops (2)

K10 desktops – coarse block diagram [56]

(HT 3.0 link not shown)

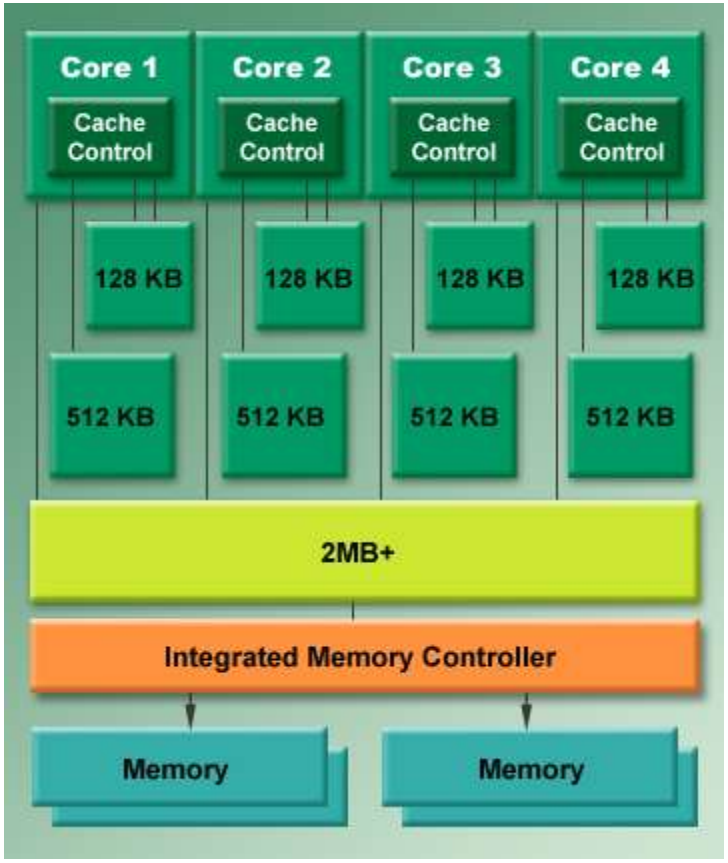

6.5 Main enhancements of K10 (Barcelona) desktops (3)

Main innovations and enhancements of K10 desktops [8]

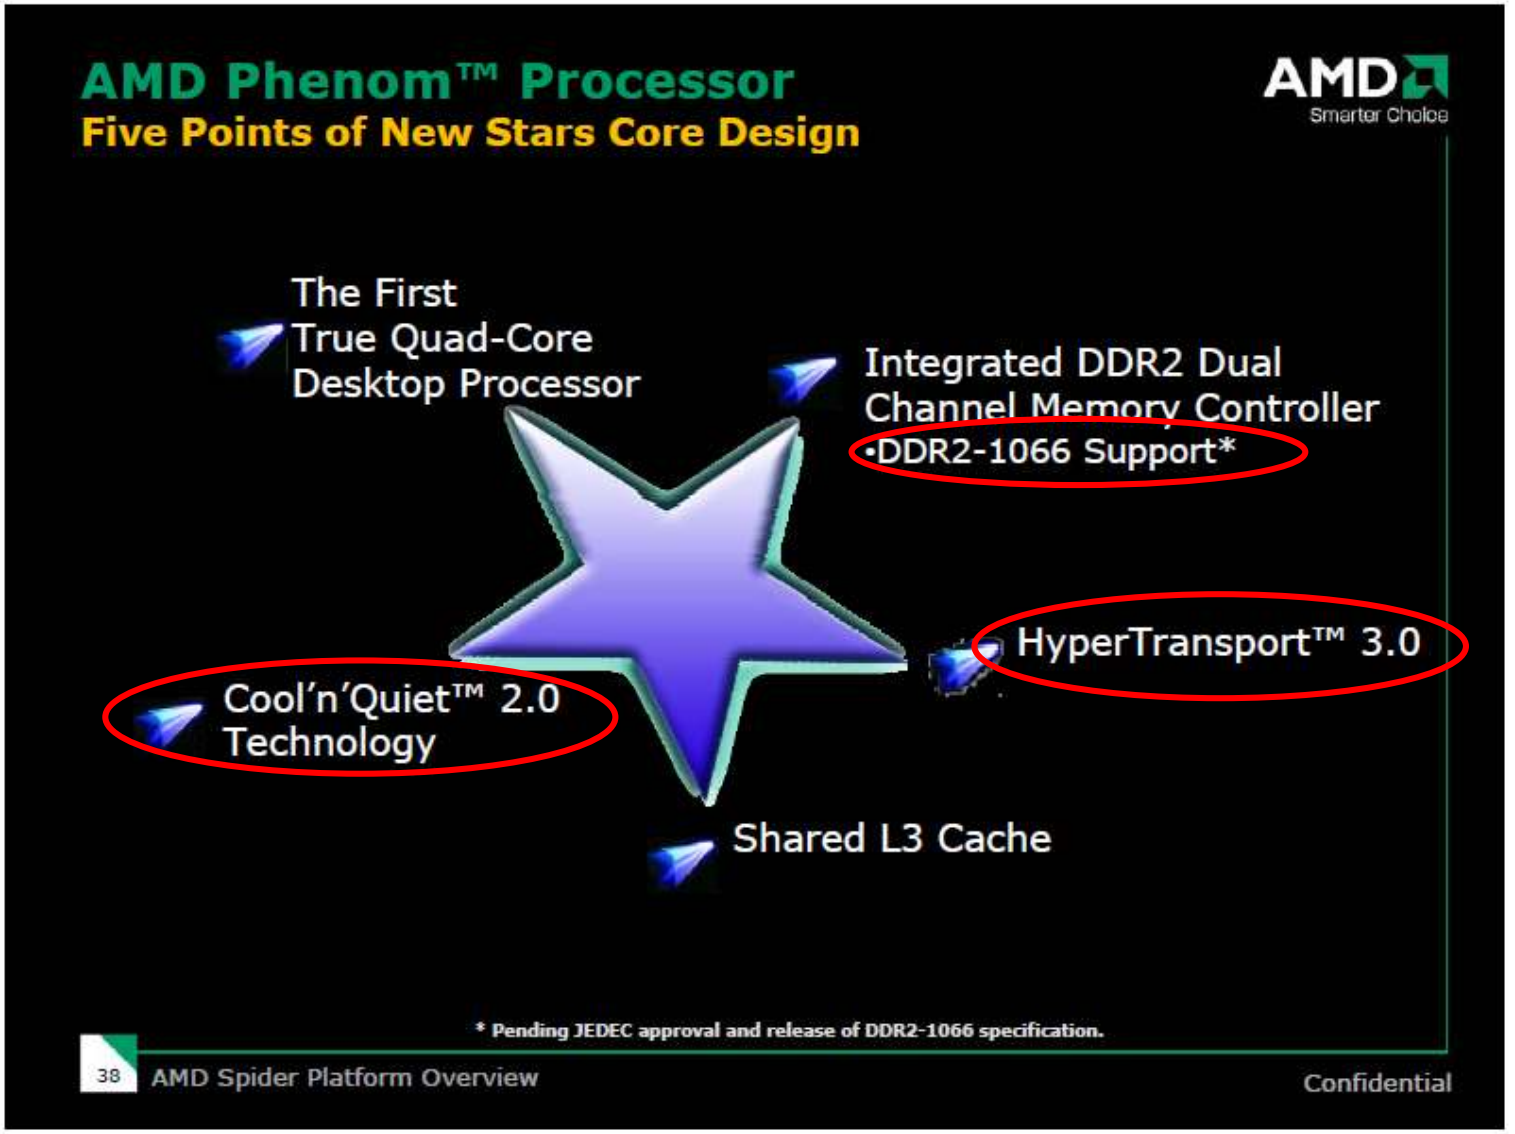

6.5 Main enhancements of K10 (Barcelona) desktops (4)

a) HyperTransport 3.0 support in K10 desktops

- Originally, AMD introduced the 130 nm K8 microarchitecture with HT 1.0 links operating at 0.8 GHz in 2003.
- Soon, with the introduction of 90 nm K8 cores AMD begun to use enhanced (HT 2.0) links at 1.0 GHz link speed in 2004.
- In K10-based desktops AMD introduced already the AM2+ socket.
- Socket AM2+ supports HT 3.0 links with increased link speeds of 1.6 to 2.0 GHz and dual power planes needed for Dual Dynamic Power Management.

  Nevertheless, Socket AM2+ supports only DDR2 memory.

  (DDR3 support was introduced then with the AM3 socket in K10.5 (Shanghai)-based desktops).

6.5 Main enhancements of K10 (Barcelona) desktops (5)

Key parameters of subsequent versions of the HyperTransport standard [138]

HT version	Year	Max. HT frequency	Max. link width	Typical link width	Max. bandwidth at 16-bit unidirectional
1.0	2001	800 MHz	32-bit	16-bit	3.2 GB/s
1.1	2002	800 MHz			3.2 GB/s
2.0	2004	1.4 GHz			5.6 GB/s
3.0	2006	2.6 GHz			10.4 GB/s
3.1	2008	3.2 GHz			12.8 GB/s

6.5 Main enhancements of K10 (Barcelona) desktops (6)

b) Increased DDR2 speed

K10 desktops support up to DDR2-1067 memory whereas K10-based servers only DDR2-667 memory.

6.5 Main enhancements of K10 (Barcelona) desktops (7)

c) Cool'n'Quiet 2.0

It includes

- PowerNow 2.0
- Dual Dynamic Power Management
- CoolCoore Technology,

as discussed before for the K10 Barcelona servers

but additionally also

- Wideband Frequency Control
- Multi-Point Thermal Control
- C1E state.

6.5 Main enhancements of K10 (Barcelona) desktops (8)

c1) Wideband Frequency Control

It refers to a new scheme of frequency transitions based on the redesign of the clock generator.

Whereas on the K10 server die each core contains its own PLL [55], on the K10 desktop die a single PLL with per core clock divider supplies the clock frequency, as indicated in the next Figures.

6.5 Main enhancements of K10 (Barcelona) desktops (9)

Basic layout of the clock generator in K10 servers

- Separate PLLs for each core.
- Frequency transitions achieved by modifying the clock multiplier ni.

⟹ This causes a PLL lock time of about 16 µs [135].

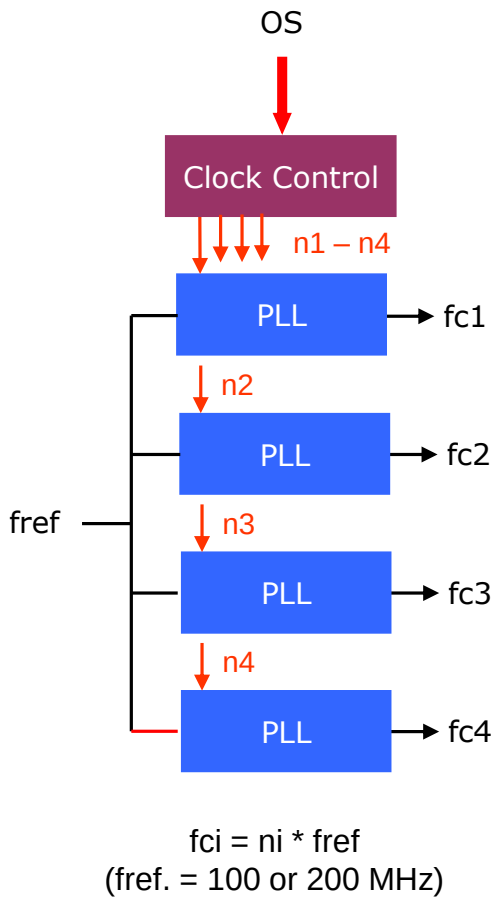

Figure: Basic layout of the clock generator in K10 servers

6.5 Main enhancements of K10 (Barcelona) desktops (10)

Basic layout of the clock generator in K10 desktops, called Wideband Frequency Control

- Use of a single PLL for all cores with individual (per core) clock dividers.
  The clock multiplier (n) is set statically (by BIOS) and remains constant during frequency transitions.
- Frequency transitions accomplished simply by modifying the clock dividers mi.

⟹ No PLL lock time occurs during frequency transitions.

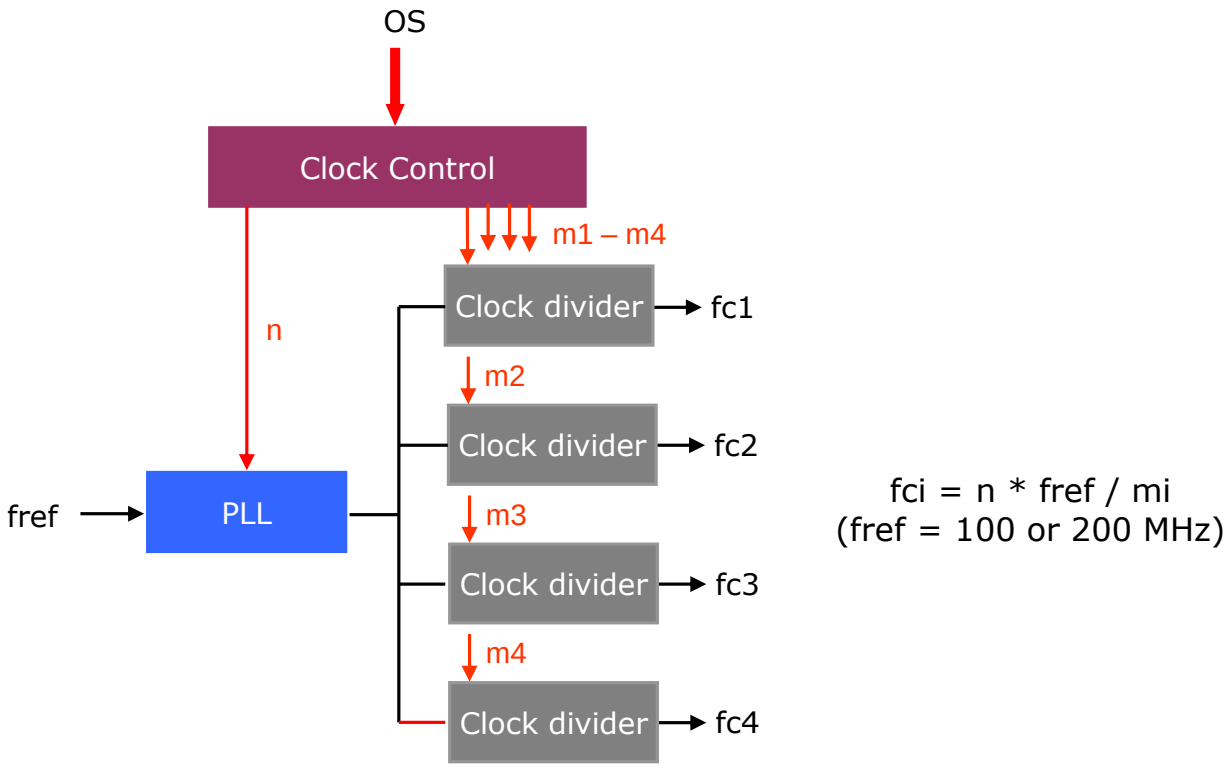

Figure: Basic layout of the clock generator in K10 desktops

6.5 Main enhancements of K10 (Barcelona) desktops (11)

c2) Multi-Point Thermal Control

It is a temperature sensors initiated overheating protection of the processor.

Principle of operation [59]

- The processor includes multiple thermal sensors typically at hot spots.
- The measured values will be continuously scanned.
- When the temperature exceeds a pre-set limit the P-state will be automatically reduced.

No detailed description could be found of AMD's Multi-Point Thermal Control mechanism.

This technique is similar to Intel's overheating protection mechanisms introduced in the Pentium 4 (2000) and its enhanced version in the mobile Core Duo (2006) processor, to be described briefly in the Remark next.

6.5 Main enhancements of K10 (Barcelona) desktops (12)

Remark 1

Principle of Intel's overheating protection as implemented in the Core Duo processor [62], [125]

- There are digital temperature sensors (first used in Core Duo, designated by red circles) beyond an analog sensor (indicated by a diode symbol) on the die.

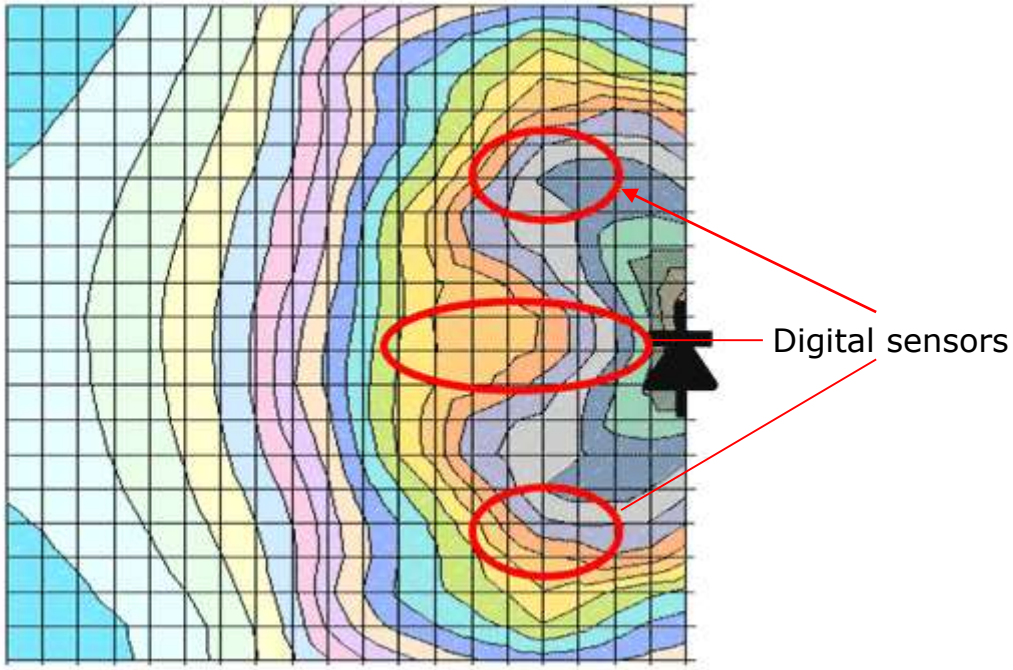

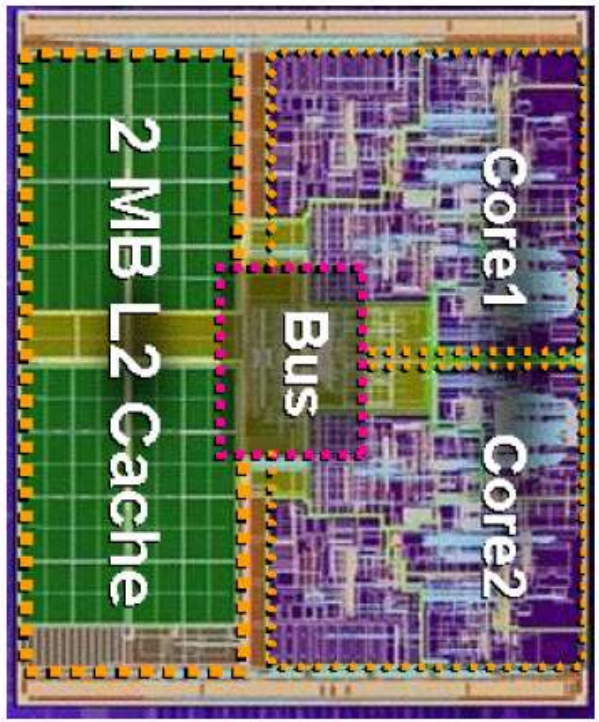

6.5 Main enhancements of K10 (Barcelona) desktops (13)

- Temperature values are continuously scanned and evaluated.
- When the temperature exceeds predefined values the control unit initiates appropriate actions to avoid overheating [62].

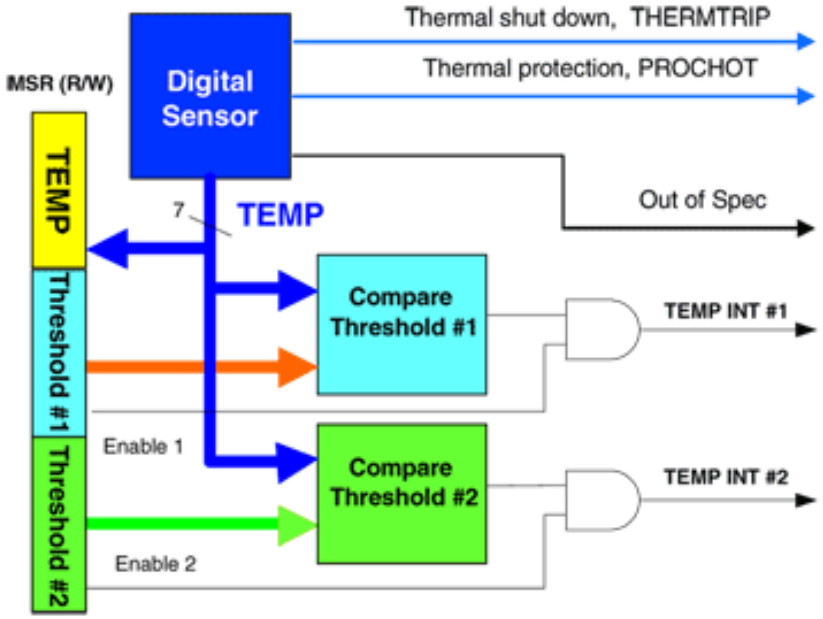

Figure: Principle of overheating control in Intel's Core Duo processor [62].

6.5 Main enhancements of K10 (Barcelona) desktops (14)

AMD's Multi-Point Thermal Control mechanism was subsequently used also in

- K10.5 Shanghai-based desktops (branded as Phenom II/Athlon II/Sempron),
- K10.5 Istambul-based desktops (branded as Phenom II) and
- K11 (Griffin)-based mobiles (branded as Turion X2).

It can be assumed that AMD replaced this mechanism by performance monitor based temperature and performance control while introducing turbo mode in Family12h-based Llano and subsequent Family 14h/15h/16h processors.

6.5 Main enhancements of K10 (Barcelona) desktops (15)

Remark

The Family 11h-based Griffin also implements Cool'n'Quiet 2.0 including also multiple on-die thermal sensors through an integrated SMBUS (SB-TSI) interface (that replaces the thermal monitor circuit chip and the SMBUS in its predecessors) [60].

An additional MEMHOT signal sent from the embedded controller to the processor can reduce memory temperature.

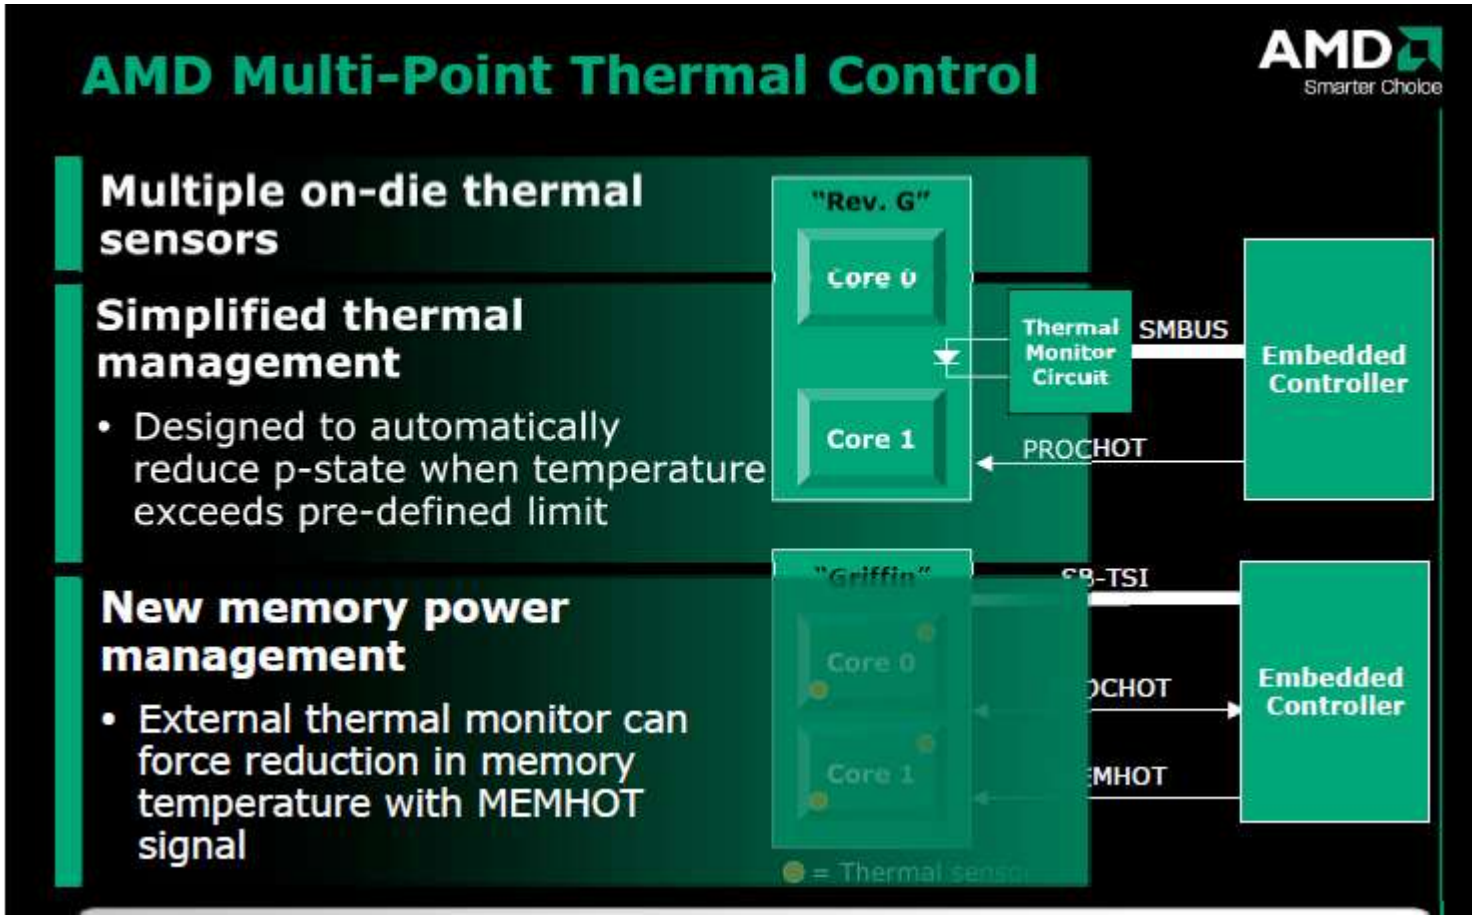

6.5 Main enhancements of K10 (Barcelona) desktops (16)

c3) C1E state

Aim of the C1E state

to reduce power consumption when all cores become idle through reducing supply voltage [61].

Principle of operation

- If all four cores become idle (i.e. enter the C1 state) the processor enters the C1E state.
- In the C1E state
  - the HT link will be deactivated,
  - the system memory will be placed into a low power state,
  - the internal clock generator will be shut down, and
  - a lower alternative voltage may be applied to the CPU cores and the NB.
    Separate voltages may be applied to the cores and the NB (in split power-plane mode).
- If the graphics card requests data from the system memory while the processor is in the C1E state, the memory interface will be waked up from its power saving mode, it sends the requested data and goes back-to-sleep without waking the cores from their C1E state.

6.5 Main enhancements of K10 (Barcelona) desktops (17)

Contrasting the Agena desktop and the Barcelona server die

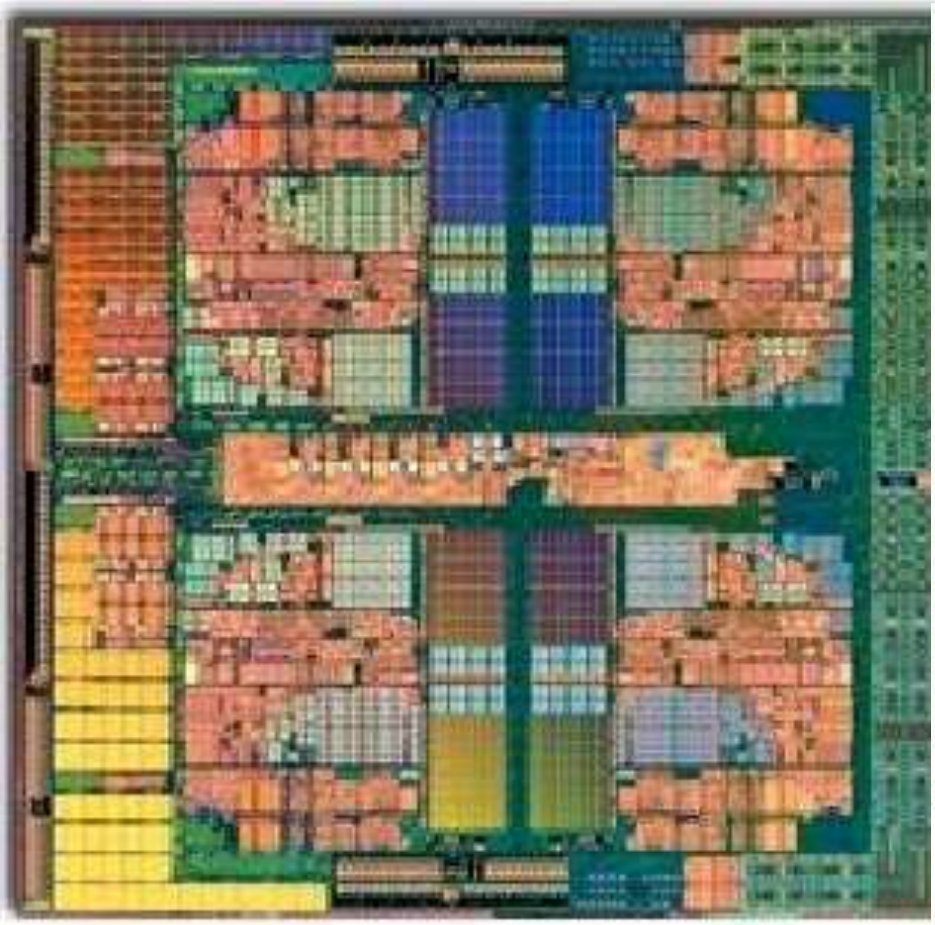

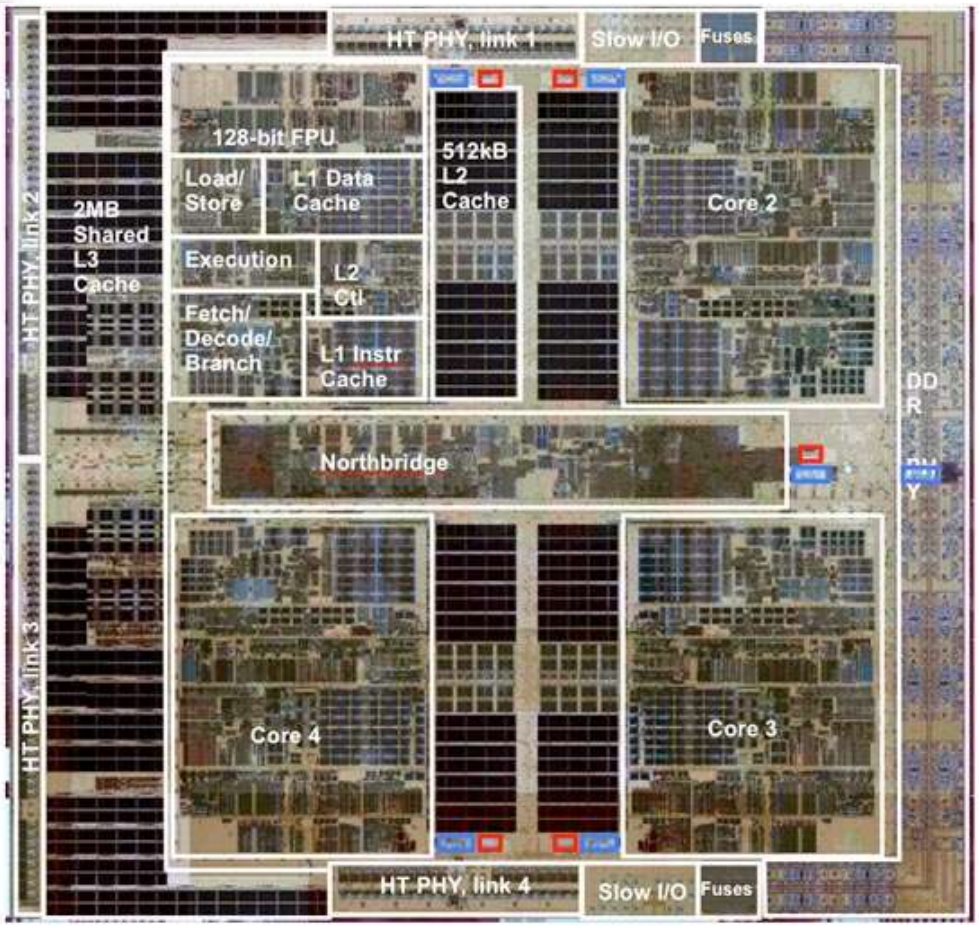

Agena desktop die [57]
285 mm², 463 mtrs

Barcelona server die [55]
285 mm², 463 mtrs

6.6 K10 (Barcelona) desktop lines

6.6 K10 (Barcelona) desktop lines (1)

Brand names of AMD's K10 desktop lines

		2003-2007	2007-2008	2008-2011	2009	2009
		K8 (Hammer)	**K10 (Barcelona)**	**K10.5 (Shanghai)**	**K10.5 (Istanbul)**	**K10.5 (Magny- Course)**
Servers	**4P servers**	**See Section 4**	**Barcelona (834x-836x))**	**Shanghai (837x-839x)**	**Istambul (8410-8430)**	**Magny-Course (6100)**
	2P servers		**Barcelona (234x-236x)**	**Shanghai (237x-239x)**	**Istambul (241x-243x)**	**Lisbon (4100)**
	1P servers		**Budapest (135x-136x)**	**Suzuka (138x-139x)**		
Desktops	**High perf.** (~80-120W)		Phenom X4-X2	Phenom II X4-X2	Phenom II X6-X4	
	Mainstream (~60-90W)	Athlon 64 Athlon 64 X2	Athlon X2	Athlon II X4-X2		
	Value (~40-60W)	Sempron		Sempron		
Mobiles	**High perf.** (~30-40W)	Turion 64 X2 (TL 6/5) Turion 64 (ML/MT)		Phenom II (N/P 9xx-6xx) Turion II Ultra (M6xx) Turion II (M/N/P 5xx)		
	Mainstream (~20-30W)	Athlon 64 X2 (TK-5x/4x) Athlon 64 (2xxx+-4xxx+)		Athlon II (M/N/P 3xx) Sempron (M1xx)		
	Ultraportable (~10-20W)	Mobile Sempron (2xxx+-4xxx+) Sempron 2100 fanless		Turion II Neo (K6xx) Athlon II Neo (K1xx) V-series (V1xx)		
Embedded (~10-20W)				Turion II Neo X2 Athlon II Neo X2 Athlon II Neo		

6.6 K10 (Barcelona) desktop lines (2)

K10 desktop lines - The cores [14]

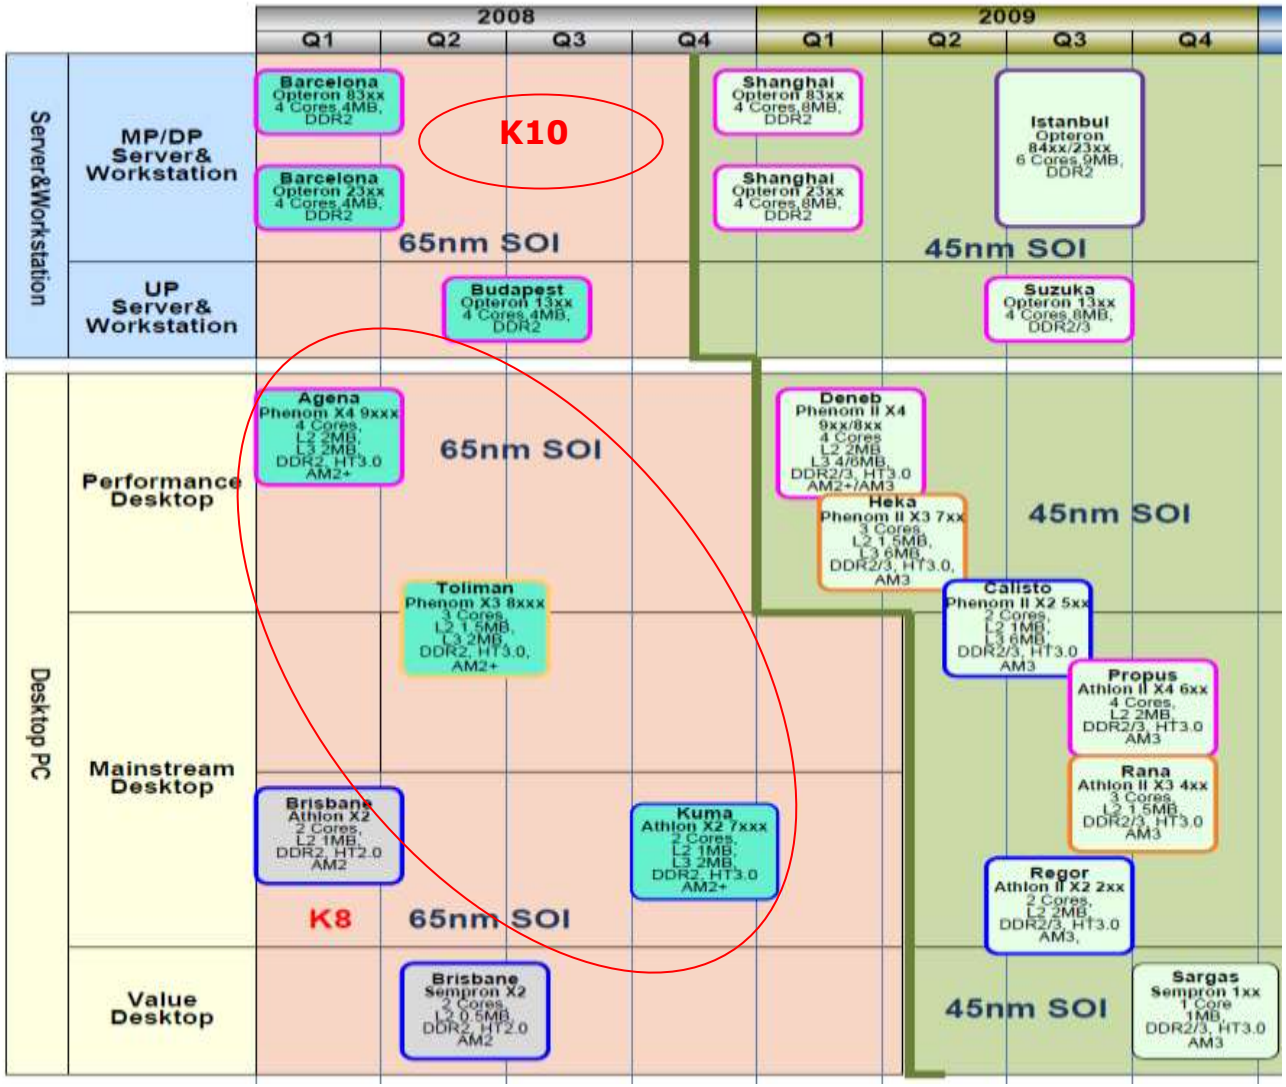

6.6 K10 (Barcelona) desktop lines (3)

Overview and main features of K10 desktop lines and cores

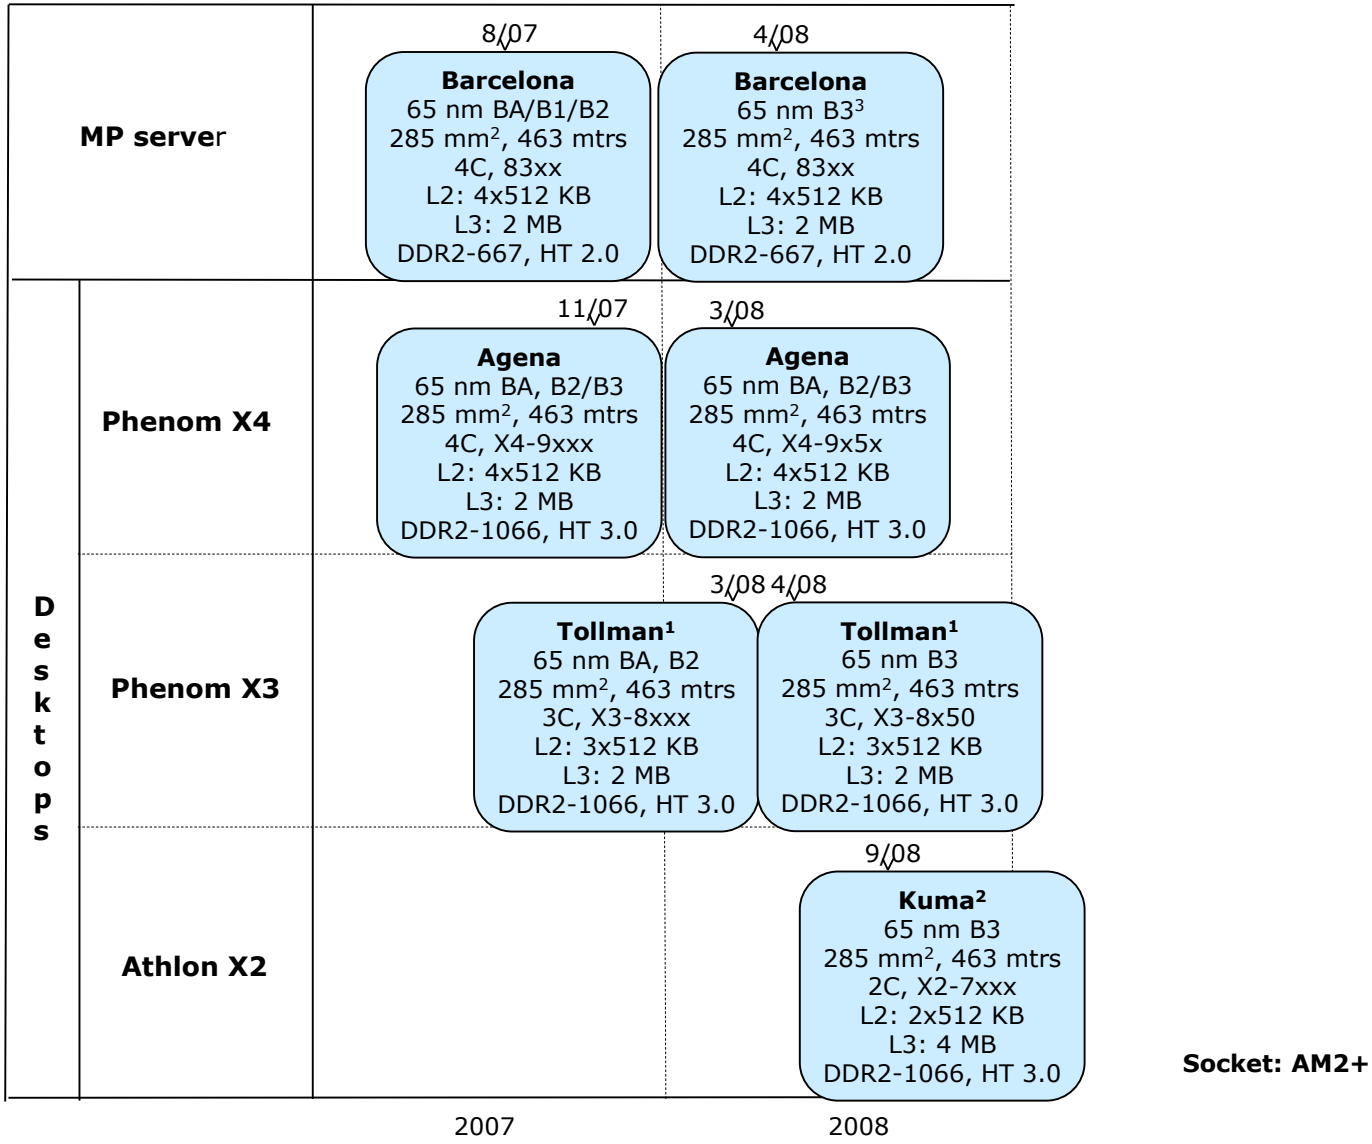

[1] 1C disabled [2] 2 C disabled 3 In stepping BA/B1/B2 ther was a TLB bug that becaume repaired in stepping B3

6.6 K10 (Barcelona) desktop lines (4)

TLB bug in AMD's K10 based desktops

K10-based desktops (named as Phenom X4/X3 or Athlon X2 lines) included a TLB bug up to the stepping B2 that became fixed in the new stepping B3.

6.6 K10 (Barcelona) desktop lines (5)

Overview of K10 desktop lines

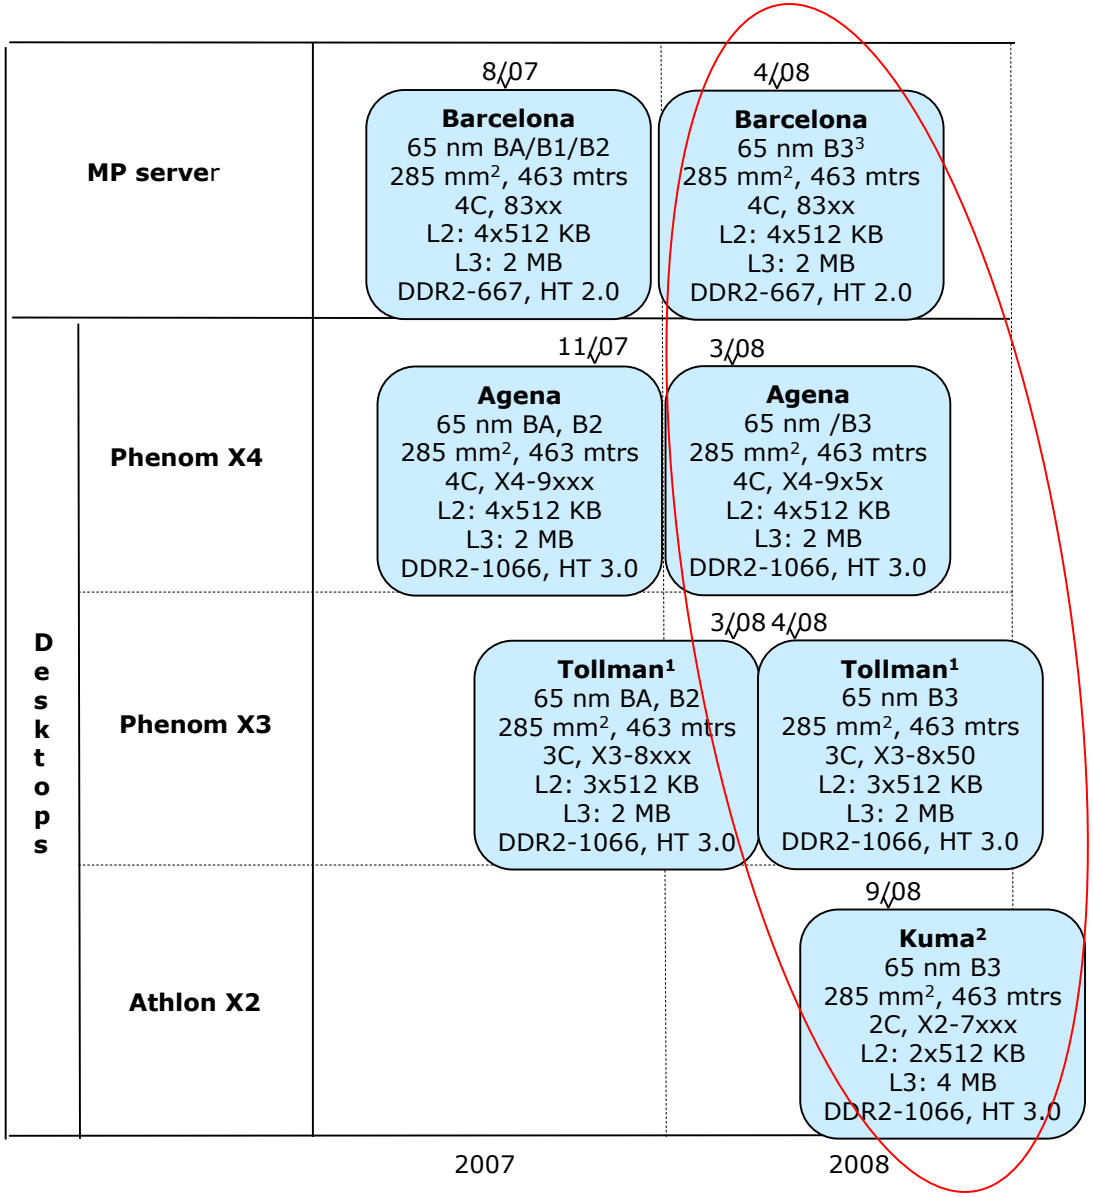

TLB bug free
B3 stepping

Socket: AM2+

[1] 1C disabled [2] 2 C disabled 3 In stepping BA/B1/B2 ther was a TLB bug that becaume repaired in stepping B3

6.6 K10 (Barcelona) desktop lines (6)

Main features of AMD's high-performance K10-based desktop 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. 15h based lines		-	-	-	-	-	-	-	-	-	-

6.6 K10 (Barcelona) desktop lines (7)

Main features of K10 desktop lines [139]

Specifications							Steppings	
Model / Part	# Cores	# Threads	Frequency	L2 Cache	L3 Cache	TDP	B2	B3
Phenom X4 9100e	4	4	1.8GHz	2MB	2MB	65W	+	
Phenom X4 9150e	4	4	1.8GHz	2MB	2MB	65W		+
Phenom X4 9350e	4	4	2GHz	2MB	2MB	65W		+
Phenom X4 9450e	4	4	2.1GHz	2MB	2MB	65W		+
Phenom X4 9500	4	4	2.2GHz	2MB	2MB	95W	+	
Phenom X4 9550	4	4	2.2GHz	2MB	2MB	95W		+
Phenom X4 9600	4	4	2.3GHz	2MB	2MB	95W	+	
Phenom X4 9600	4	4	2.3GHz	2MB	2MB	95W	+	
Phenom X4 9600B	4	4	2.3GHz	2MB	2MB	95W	+	
Phenom X4 9600B	4	4	2.3GHz	2MB	2MB	95W		+
Phenom X4 9650	4	4	2.3GHz	2MB	2MB	95W		+
Phenom X4 9700	4	4	2.4GHz	2MB	2MB	125W	+	
Phenom X4 9750	4	4	2.4GHz	2MB	2MB	95W		+
Phenom X4 9750	4	4	2.4GHz	2MB	2MB	125W		+
Phenom X4 9750B	4	4	2.4GHz	2MB	2MB	95W		+
Phenom X4 9850	4	4	2.5GHz	2MB	2MB	95W		+
Phenom X4 9850	4	4	2.5GHz	2MB	2MB	125W		+
Phenom X4 9850	4	4	2.5GHz	2MB	2MB	125W		+
Phenom X4 9850B	4	4	2.5GHz	2MB	2MB	95W		+
Phenom X4 9950	4	4	2.6GHz	2MB	2MB	125W		+
Phenom X4 9950	4	4	2.6GHz	2MB	2MB	140W		+

6.6 K10 (Barcelona) desktop lines (8)

Remarks to the designation of K10 desktops

1) K10/K10.5-based lines altogether are often designated as the Hound lines due to the names of their underlying microarchitectures (Deerhound, Greyhound).

2) K10/K10.5-based desktop lines are also designated as Stars lines,

- like the Phenom line (that is the desktop version of the K10 Barcelona server line) or
- the Pnenom II line (that is the desktop version of the K10.5 Shanghai server line),

as individual processors of the lines mentioned, like Agena, Toliman of the Phenom line, or Deneb, Heka, Calisto etc. of the Phenom II line, are designated after names of stars.

7. The K10.5 Shanghai family

- 7.1 Overview of the K10.5 Shanghai family
- 7.2 Key enhancements of the K10.5 Shanghai family
- 7.3 K10.5 Shanghai-based server lines
- 7.4 K10.5 Shanghai-based desktop lines
- 7.5 K10.5 Shanghai-based mobile lines
- 7.6 K10.5 Shanghai-based embedded lines

7.1 Overview of the K10.5 Shanghai family

7.1 Overview of the K10.5 Shanghai family

- Introduced in 11/2008
- 45 nm technology vs. Barcelona's 65 nm feature size.
- Improved IC technology to reduce leakage.
- Many improvements or enhancements to increase performance or reduce power consumption.

7.1 Overview of the K10.5 Shanghai family (2)

Overview of AMD's K10.5 Shanghai family [14]

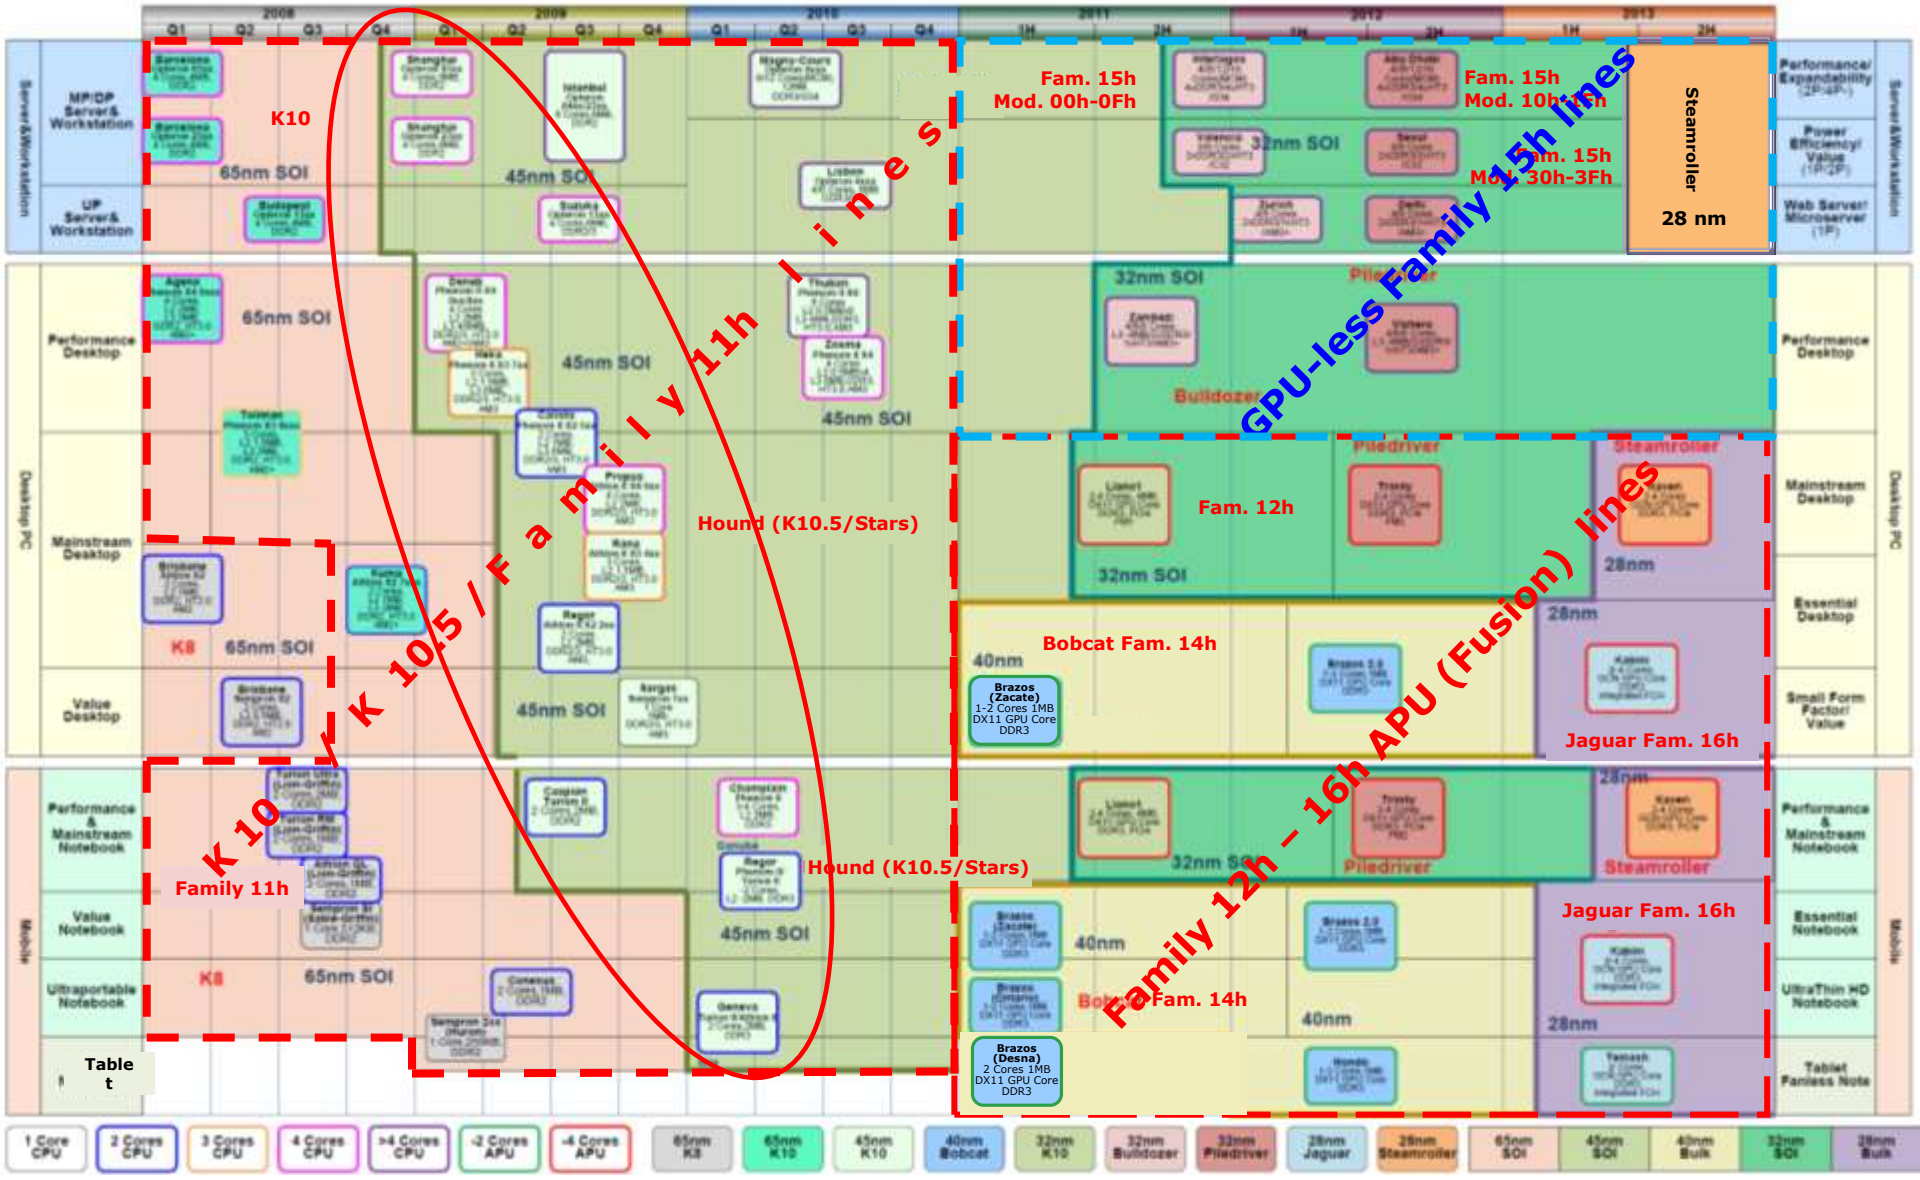

7.1 Overview of the K10.5 Shanghai family (3)

Brand names of AMD's K10.5h Shanghai-based processor lines

		2003-2007	2007-2008	2008-2011	2009	2009
		K8 (Hammer)	**K10 (Barcelona)**	**K10.5 (Shanghai)**	**K10.5 (Istanbul)**	**K10.5 (Magny- Course)**
Servers	**4P servers**	**See Section 4**	**Barcelona (834x-836x))**	**Shanghai (837x-839x)**	**Istambul (8410-8430)**	**Magny-Course (6100)**
	2P servers		**Barcelona (234x-236x)**	**Shanghai (237x-239x)**	**Istambul (241x-243x)**	**Lisbon (4100)**
	1P servers		**Budapest (135x-136x)**	**Suzuka (138x-139x)**		
Desktops	**High perf.** (~80-120W)		Phenom X4-X2	Phenom II X4-X2	Phenom II X6-X4	
	Mainstream (~60-90W)	Athlon 64 Athlon 64 X2	Athlon X2	Athlon II X4-X2		
	Value (~40-60W)	Sempron		Sempron		
Mobiles	**High perf.** (~30-40W)	Turion 64 X2 (TL 6/5) Turion 64 (ML/MT)		Phenom II (N/P 9xx-6xx) Turion II Ultra (M6xx) Turion II (M/N/P 5xx)		
	Mainstream (~20-30W)	Athlon 64 X2 (TK-5x/4x) Athlon 64 (2xxx+-4xxx+)		Athlon II (M/N/P 3xx) Sempron (M1xx)		
	Ultraportable (~10-20W)	Mobile Sempron (2xxx+-4xxx+) Sempron 2100 fanless		Turion II Neo (K6xx) Athlon II Neo (K1xx) V-series (V1xx)		
Embedded (~10-20W)				Turion II Neo X2 Athlon II Neo X2 Athlon II Neo		

7.2 Key enhancements of the K10.5 Shanghai family

7.2 Key enhancements of the K10.5 Shanghai family

7.2.1 Overview [129]

Performance

45nm with Immersion Lithography

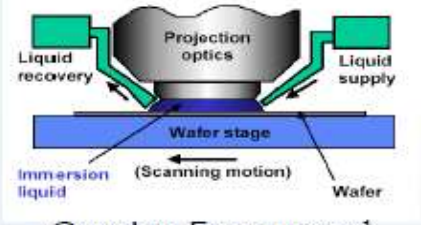

Greater Frequency[1]

AMD Memory Optimizer Technology

Enhanced Pre-fetching

2x Core Probe Bandwidth[1]

AMD Balanced Smart Cache

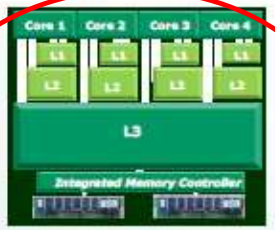

8MB: 2x total cache[1]

HyperTransport ™ technology 3.0[2]

Up to 17.6GB/s per link

Power Efficiency

45nm with Immersion Lithography

~35% Lower CPU Idle Power[1]

AMD Smart Fetch

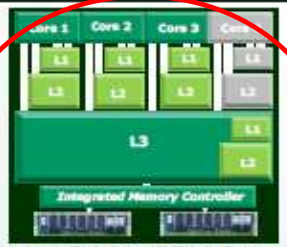

Up to 21% CPU Power Savings

Virtualization

Enhanced RVI

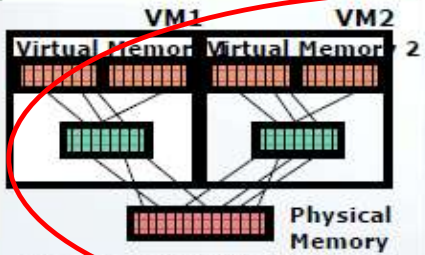

Increased Performance [1]

RAS

L3 Cache Index Disable[3]

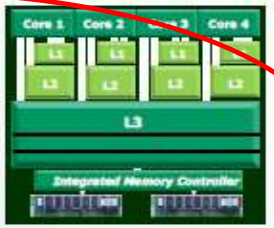

Superior Data Integrity

7.2 Key enhancements of the K10.5 Shanghai family (2)

7.2.2 Key enhancements of the K10.5 Shanghai family targeting higher performance

a) Increased L3 cache size

Increasing L3 from 2 MB (Barcelona) to 6 MB [63], [64]

Barcelona (65 nm)

Dedicated L1
- Locality keeps most critical data in the L1 cache
- Lowest latency
- 2 loads per cycle

Dedicated L2
- Sized to accommodate the majority of working sets today
- Dedicated to eliminate conflicts common in shared caches
  - Better for Virtualization

Shared L3 – NEW
- Victim-cache architecture maximizes efficiency of cache hierarchy
- Fills from L3 leave likely shared lines in the L3
- Sharing-aware replacement policy
- Ready for expansion at the right time for customers

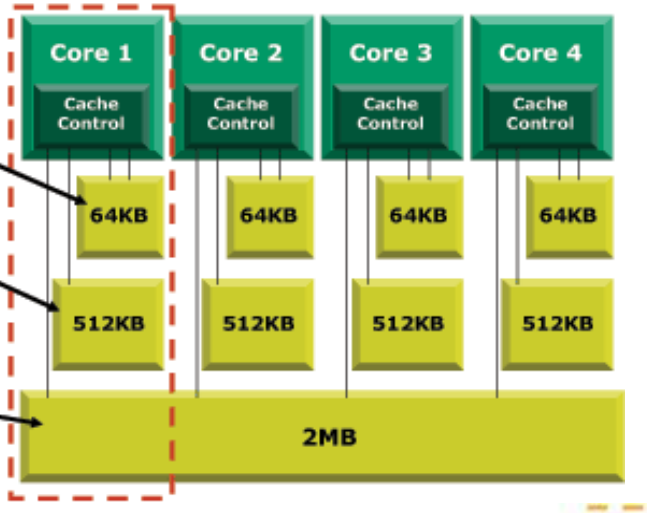

Shanghai (45 nm)

Shared L3 Cache
- Victim-cache architecture maximizes efficiency of cache hierarchy
- Fills from L3 leave likely shared lines in the L3
- Sharing-aware replacement policy

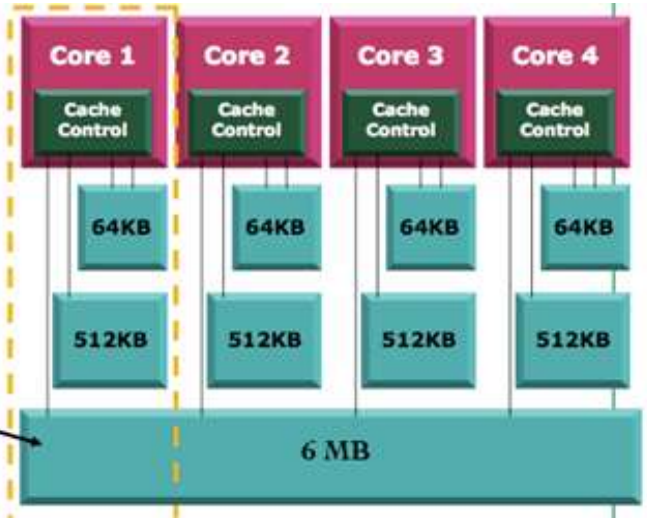

7.2 Key enhancements of the K10.5 Shanghai family (3)

b) Increased memory speed

Implemented both in servers, desktops and mobiles, as indicated next.

7.2 Key enhancements of the K10.5 Shanghai family (4)

Main features of AMD's K10.5 Shanghai-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)	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 (NPT)	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. 17h Mod. 00h-0Fh		6/2017	Epyc (2S!!) (4 dies/proc.)	7000	14 nm	4x(2x4) (32C)	½ MB/C	2 MB/C	DDR4-2666	IFIS 75.8 GB/s	SP3

7.2 Key enhancements of the K10.5 Shanghai family (5)

Main features of AMD's high-performance K10.5 Shanghai-based desktop 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. 15h based lines		-	-	-	-	-	-	-	-	-	-

7.2 Key enhancements of the K10.5 Shanghai family (6)

Main features of AMD's high performance K10.5 Shanghai-based mobile lines

Base arch./ stepping		Intro	High perf. mobile family name	Series	Techn.	Core count (up to)	L2 (up to)	L3	Memory (up to)	HT/ dir. (up to)	Sock et
K8	C0, CG	9/2003	Clawhammer	Mobile Athlon 64	130 nm	1	512 KB	-	DDR-400	HT 1.0: 3.2 GB/s	754
	E5	3/2005	Lancaster	Turion 64	90 nm	1	1 MB	-	DDR-400	HT 1.0: 3.2 GB/s	754
	F2	5/2006	Trinidad	Turion 64 X2	90 nm	2	2*512 KB	-	DDR2-667	HT 1.0: 3.2 GB/s	S1
K10	-	-	-	-	-	-	-	-	-	-	-
K10.5	DA-C2	9/2009	Caspian	Turion II	45 nm	2	2*512 KB/ 2*1 MB[1]	-	DDR2-800	HT 3.0: 7.2 GB/s	S1g3
	DA-C3	5/2010	Champlain	Turion X4	45 nm	4	4*512 KB	-	DDR3-1066	HT 3.0: 7.2 GB/s	S1g4

[1]: 2*512 KB for Turion II, 2*1 MB for Turion II Ultra

7.2 Key enhancements of the K10.5 Shanghai family (7)

Spreading of K10.5 Shanghai cores in AMD's desktop and mobile lines – Overview [14]

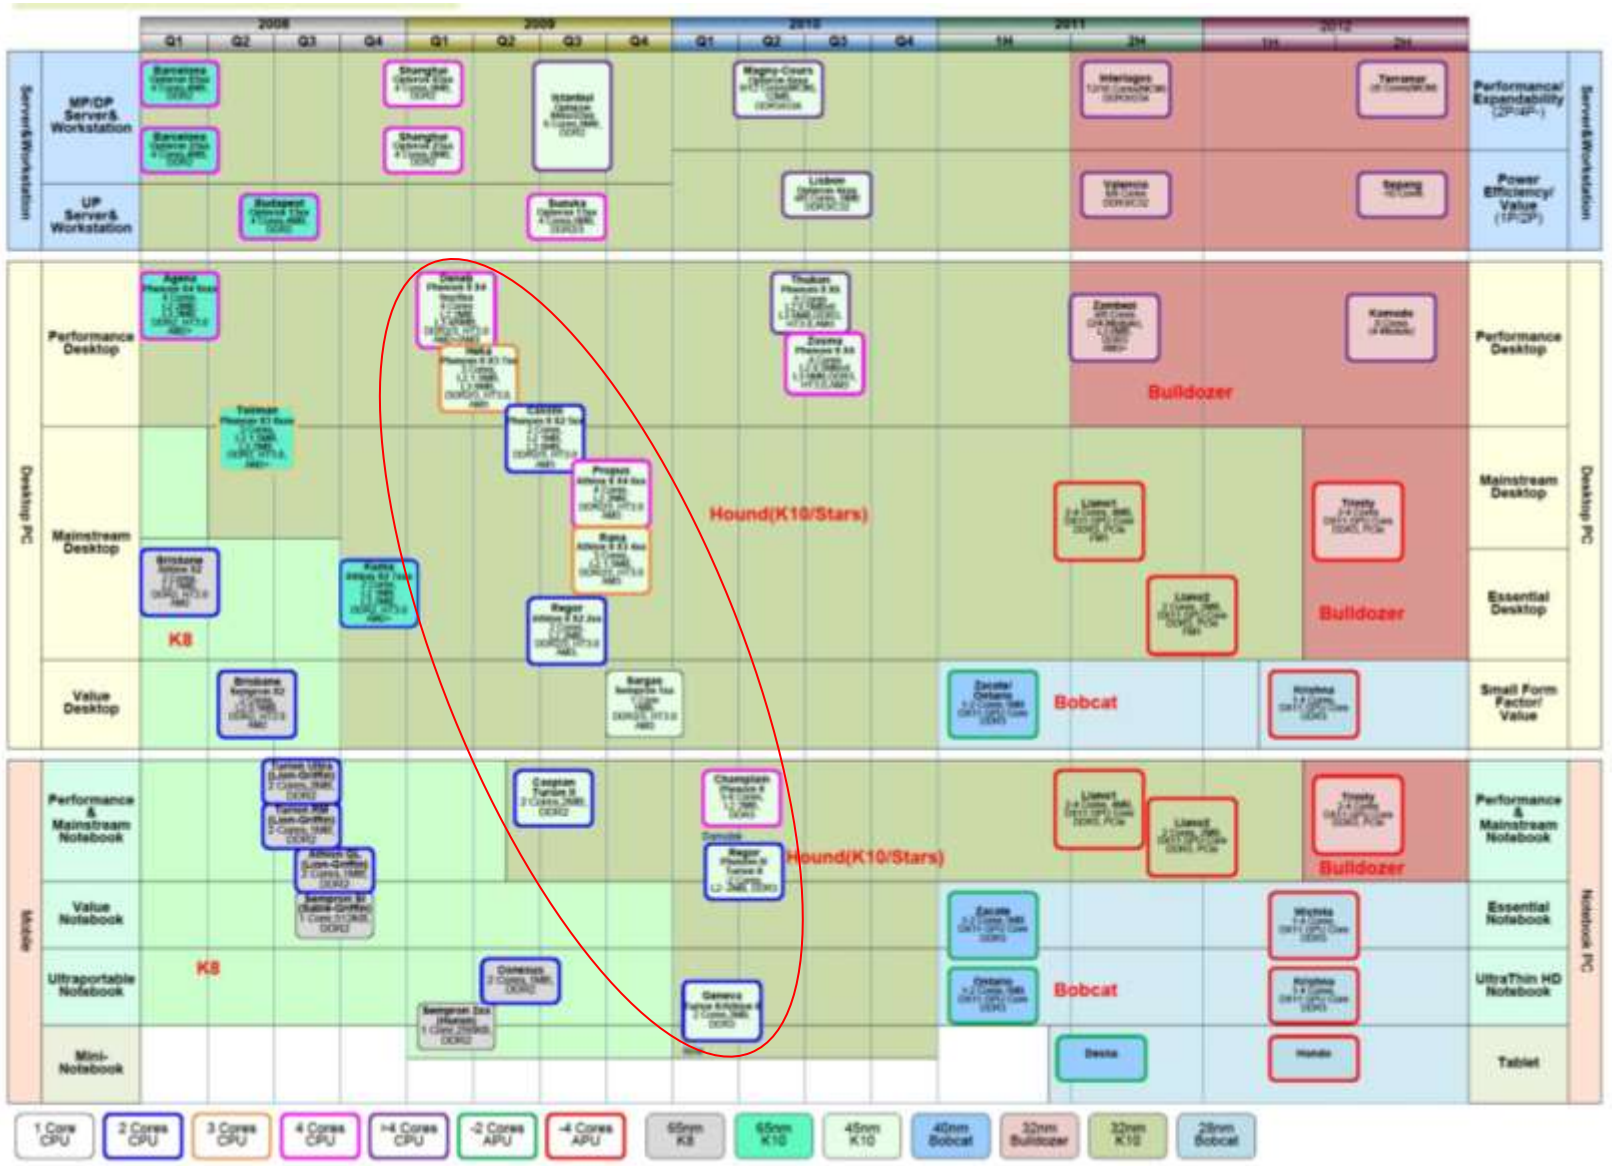

7.2 Key enhancements of the K10.5 Shanghai family (8)

c) HT 3.0 links for DP and MP servers

AMD implemented HT 3.0 links as early as on the K10 Barcelona server die, nevertheless, K10 DP and MP servers could not make use of it, since AMD insisted on the Socket F just introduced for the previous K8 NPT line for reasons of compatibility.

So HT 3.0, implemented on the die, became available only for K10 UP servers and desktops along with a new socket (AM2+).

With their K10.5 Shanghai family AMD initially replaced the former Socket F (Fr1) to the the Socket F+ (more precisely Fr2) to allow the utilization of Dual Dynamic Power Management (i.e. dual power planes).

Socket Fr2 however, supported only HT 2.0.

Only the second major wave of K10.5 Shanghai DP/MP servers, launched about half a year after the initial release, gave support for HT 3.0 links, as indicated for K10.5 Shanghai MP servers in the next table.

These processor need however, an enhanced Socket F termed as Socket Fr5.

We note that Socket Fr5 has the same pin count as previous Socket F versions (1207 pins).

7.2 Key enhancements of the K10.5 Shanghai family (9)

Main features of K10.5 Shanghai MP servers [133]

Model Number	Step.	Freq.	L2-Cache	L3-Cache	HT	Mult	Voltage	ACP	TDP	Socket	Release Date
C2, Quad Core											
Opteron 8378	C2	2.4 GHz	4x 512 KB	6 MB	1 GHz[7]	12x	1.35	75 W	115 W	Socket Fr2	November 13, 2008
Opteron 8380	C2	2.5 GHz	4x 512 KB	6 MB	1 GHz[7]	12.5x	1.35	75 W	115 W	Socket Fr2	November 13, 2008
Opteron 8382	C2	2.6 GHz	4x 512 KB	6 MB	1 GHz[7]	13x	1.35	75 W	115 W	Socket Fr2	November 13, 2008
Opteron 8384	C2	2.7 GHz	4x 512 KB	6 MB	1 GHz[7]	13.5x	1.35	75 W	115 W	Socket Fr2	November 13, 2008
Opteron 8386 SE	C2	2.8 GHz	4x 512 KB	6 MB	1 GHz[7]	14x	1.325	105 W	137 W	Socket F Fr2	January 26, 2009
Opteron 8387	C2	2.8 GHz	4x 512 KB	6 MB	2.2 GHz	14x	1.325	75 W	115 W	Socket Fr5	April 22, 2009
Opteron 8389	C2	2.9 GHz	4x 512 KB	6 MB	2.2 GHz	14.5x	1.325	75 W	115 W	Socket Fr5	April 22, 2009
Opteron 8393 SE	C2	3.1 GHz	4x 512 KB	6 MB	2.2 GHz	15.5x	1.325	105 W	137 W	Socket Fr5	April 22, 2009

HT speed of 1.0 GHz: HT 2.0
HT speed of 2.2 GHz: HT 3.0

7.2 Key enhancements of the K10.5 Shanghai family (10)

Nevertheless, from the 4 HT 3.0 links implemented on the die (since the K10 Barcelona implementation), further on only 3 HT links are supported in K10.5 MP servers, as utilizing all four links requires a substantially new platform.

This happened subsequently for the K10.5 Magny-Course MP server. with the Socket G34.

7.2 Key enhancements of the K10.5 Shanghai family (11)

Achieved performance increase - Shanghai vs. Barcelona [140]

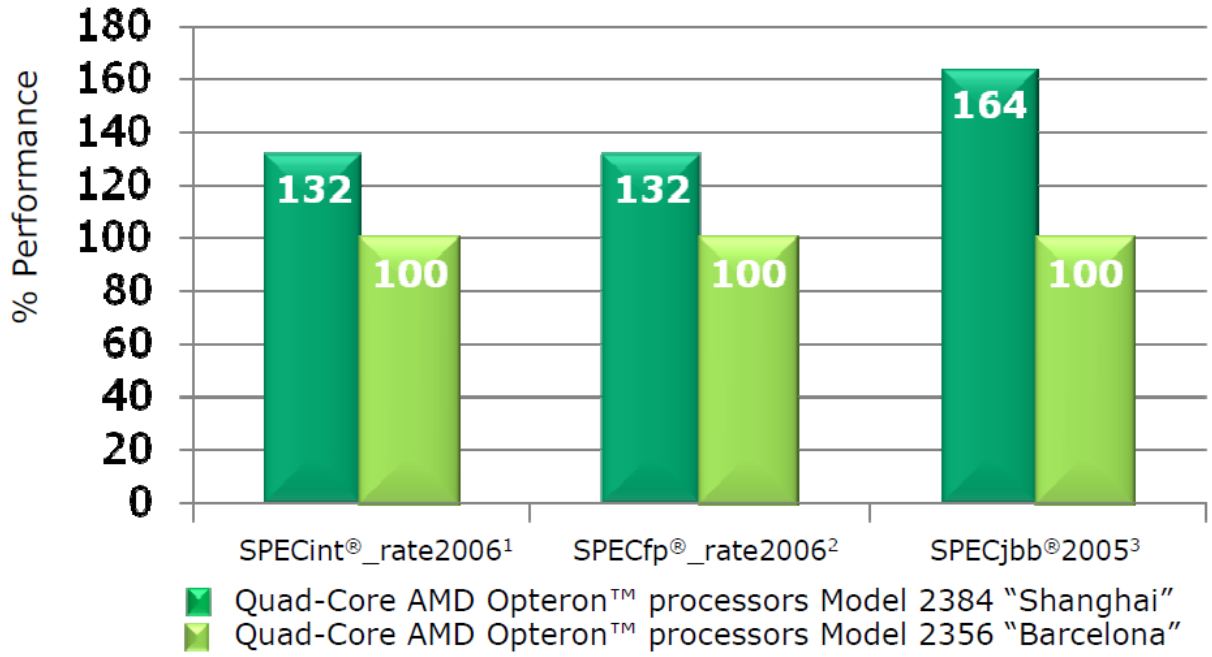

7.2 Key enhancements of the K10.5 Shanghai family (12)

7.7.3 AMD's power management techniques K8 – Family 15h (Bulldozer) (based on [53])

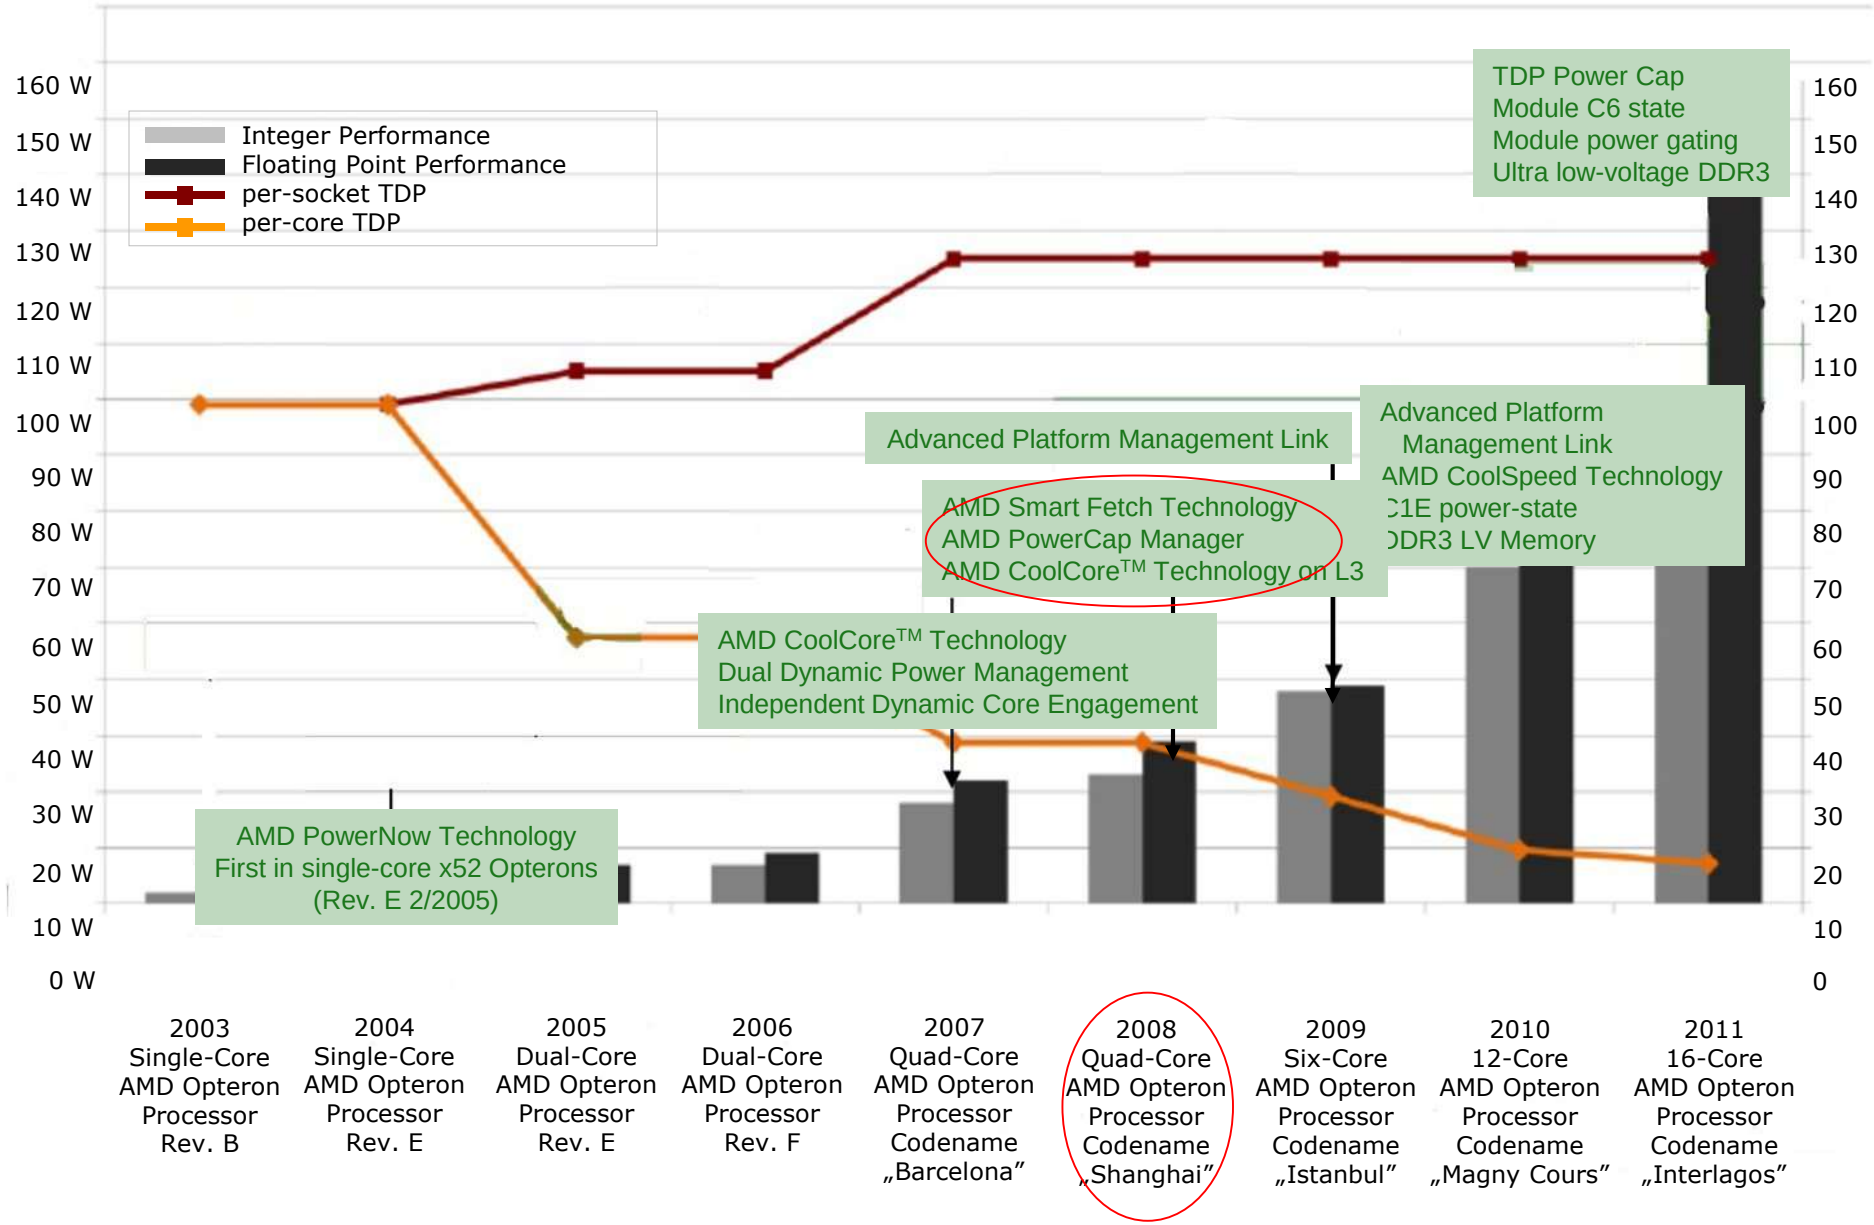

7.2 Key enhancements of the K10.5 Shanghai family (13)

a) AMD Smart Fetch Technology [65], [66]

(Clock gating at the core level).

Aim

Reducing the power consumption of idle cores.

The problem to be amended arises from snoops initiated by active cores for L1 or L2 data of idle cores (exclusive cache!), as these snoops might wake up a sleeping core causing higher power consumption.

E.g.

Quad-Core AMD Opteron™ Processor (“Barcelona”)

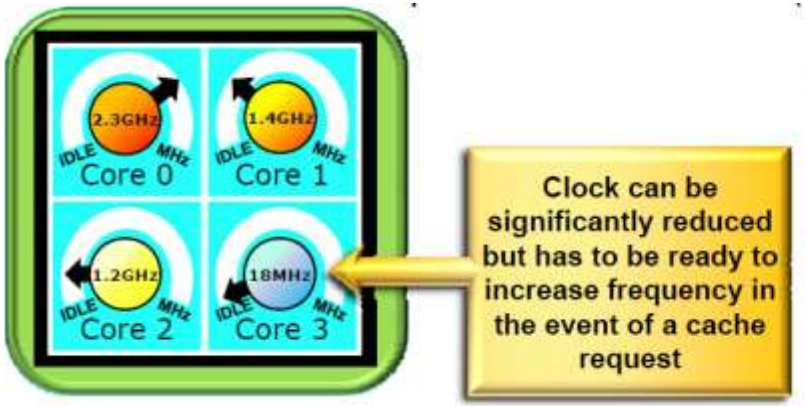

- Cores can’t shut all the way down
- Cores need to remain clocked in case other cores need data from L1 or L2

7.2 Key enhancements of the K10.5 Shanghai family (14)

Principle of the Smart Fetch Technology

Smart Fetch Technology avoids waking up an idle core by

- dumping the contents of L1 and L2 caches of the idle core into the L3 cache and then
- let enter the idle core to the clock gated C1 state (with clocking shut down).

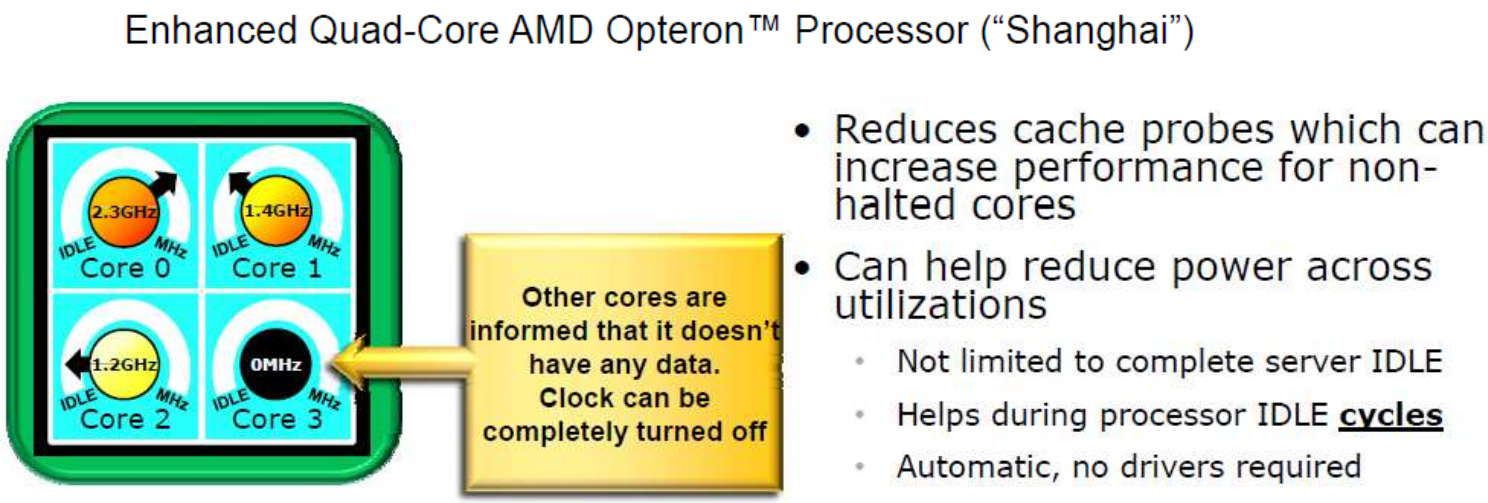

This results in a power saving of ~ 5 W per core [66] or
to a power reduction of an idle processor from ~ 25 W (Barcelona) to ~ 10 W (Shanghai).

7.2 Key enhancements of the K10.5 Shanghai family (15)

Remark [66]

Smart Fetch Technology could not be introduced in Barcelona as 2 MB L3 cache size was too small to dump the L1 and L2 caches of all 3 cores (3x(64+64) KB +3x512 KB) into L3 while leaving enough room for operation.

A further improvement of this technique is when an idle core goes into the hardware controlled C1E state with lower supply voltage (as in the Magny-Course) or even with power gating while shutting down the supply voltage (as in the Bulldozer processor, termed as the C6 state) rather than putting it into the merely clock gated C1 state.

7.2 Key enhancements of the K10.5 Shanghai family (16)

b) AMD PowerCap Manager-1 [67]

- It is a BIOS extension to limit maximum power consumption during operation.
- The BIOS option caps clock frequency and voltage with four options:
  - Disabled full-range voltage and clock frequency, no effective power savings
  - 1 – Caps power consumption to 70% of normal, effectively saving 30% power
  - 2 – Caps power consumption to 60% of normal, effectively saving 40% power
  - 3 – Caps power consumption to 40% of normal, effectively saving 60% power
- This feature is targeted at data center operations when power consumption is more critical than absolute performance (such as in case of cloud computing workloads).

b) AMD PowerCap Manager-2 [65]

What is AMD PowerCap management?

- BIOS selectable options that set maximum MHz/voltage limits
  - Choose a setting and the processor(s) can operate up to that set limit

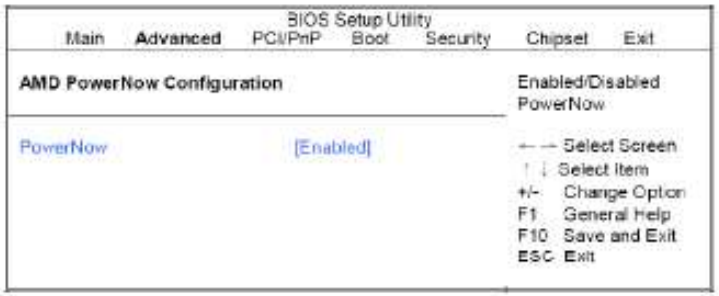

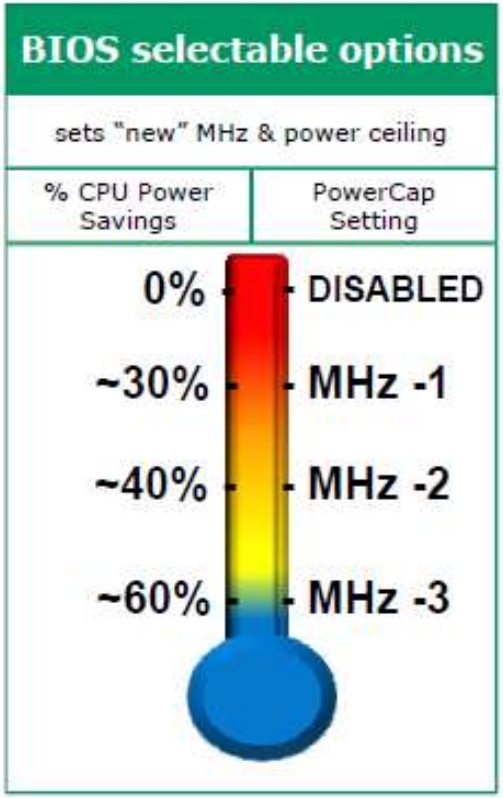

AMD PowerCap manager requires no special drivers, operating systems or hypervisor support.

7.2 Key enhancements of the K10.5 Shanghai family (18)

b) AMD PowerCap Manager-3 [65]

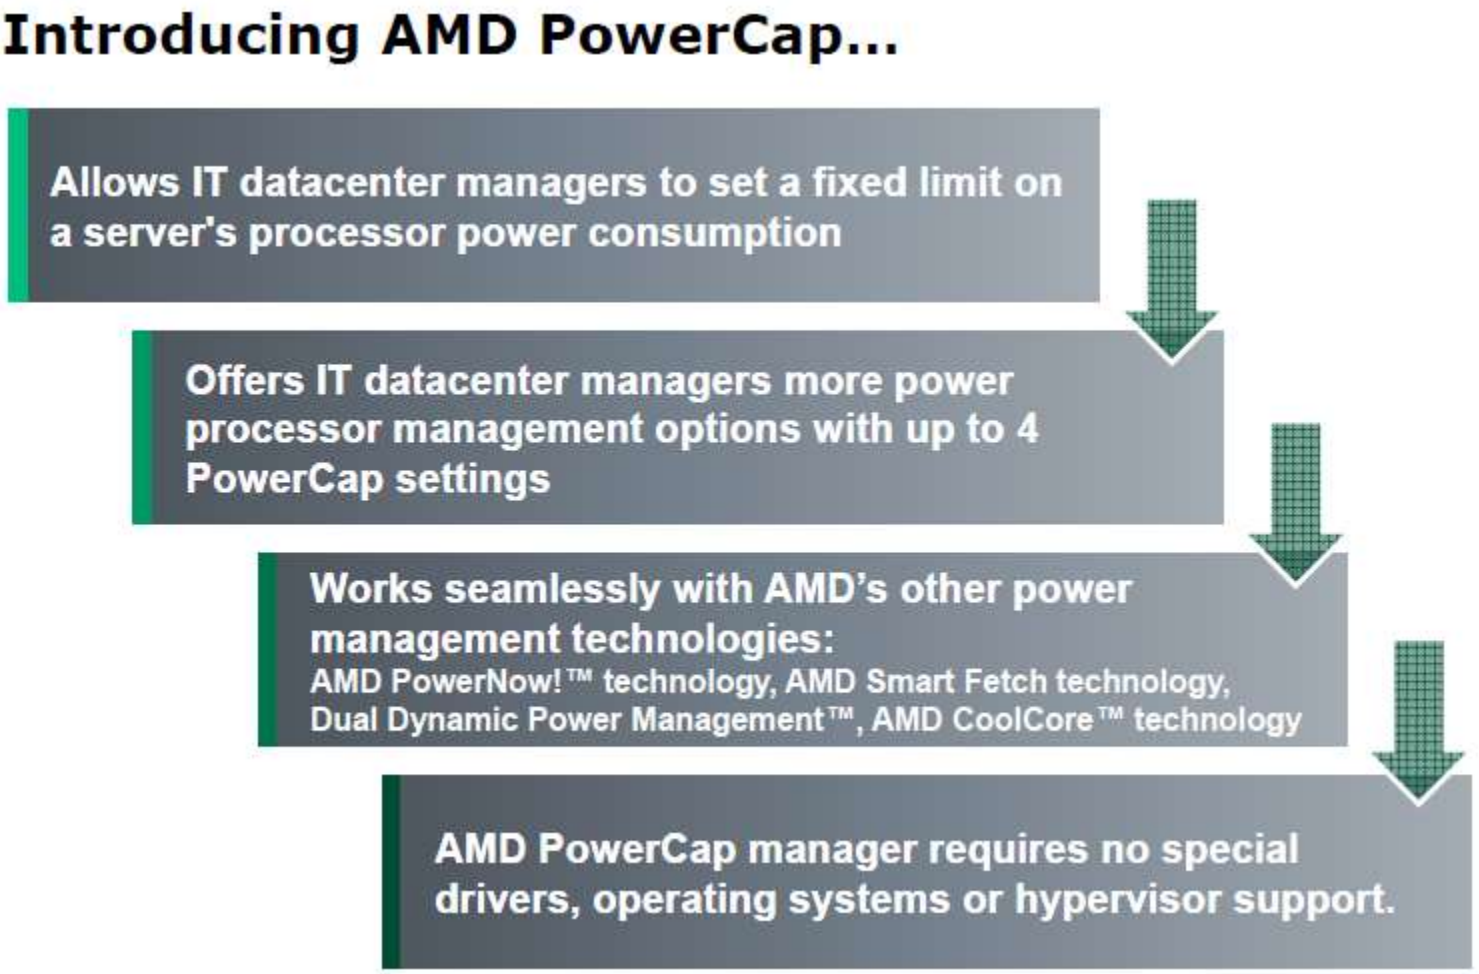

7.2 Key enhancements of the K10.5 Shanghai family (19)

c) Extension of the CoolCore Technology to the L3 cache [65]

- If sections of the L3 cache are not used, they can be turned off to save power.
- L3 cache can be gated in 4 sections
  - 2x1MB blocks
  - 2x2MB blocks.

AMD CoolCore™ Technology

Can reduce processor energy consumption by dynamically turning off sections of the processor when inactive.

No software requirements, part of CPU design logic

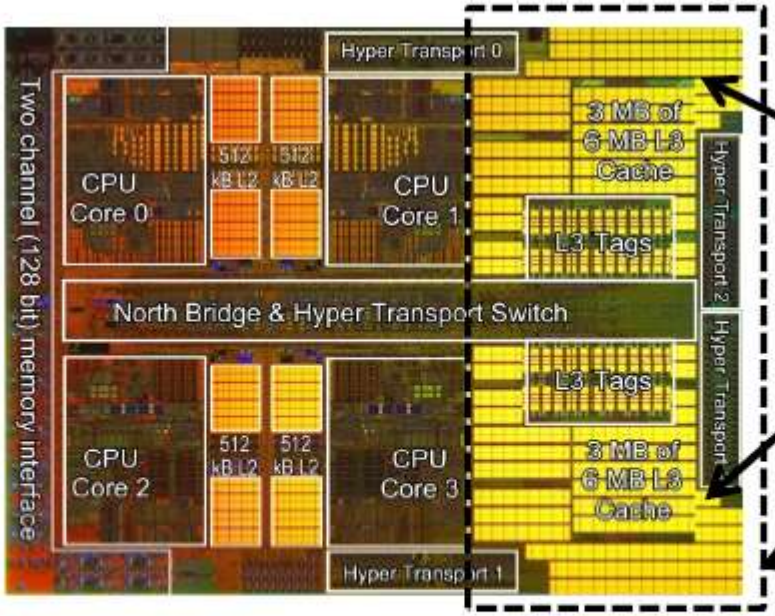

NEW! Now extended to the L3 cache

L3 cache logic considerable portion of die

Dynamically turns off sections dependant on workload

- L3 cache can be gated in 4 sections
  - 2x 1MB blocks
  - 2x 2MB blocks

7.2 Key enhancements of the K10.5 Shanghai family (20)

Performance per power ratios Shanghai vs. Barcelona [65]

Results obtained while comparing performance per power ratios at 2.3 GHz processor frequencies.

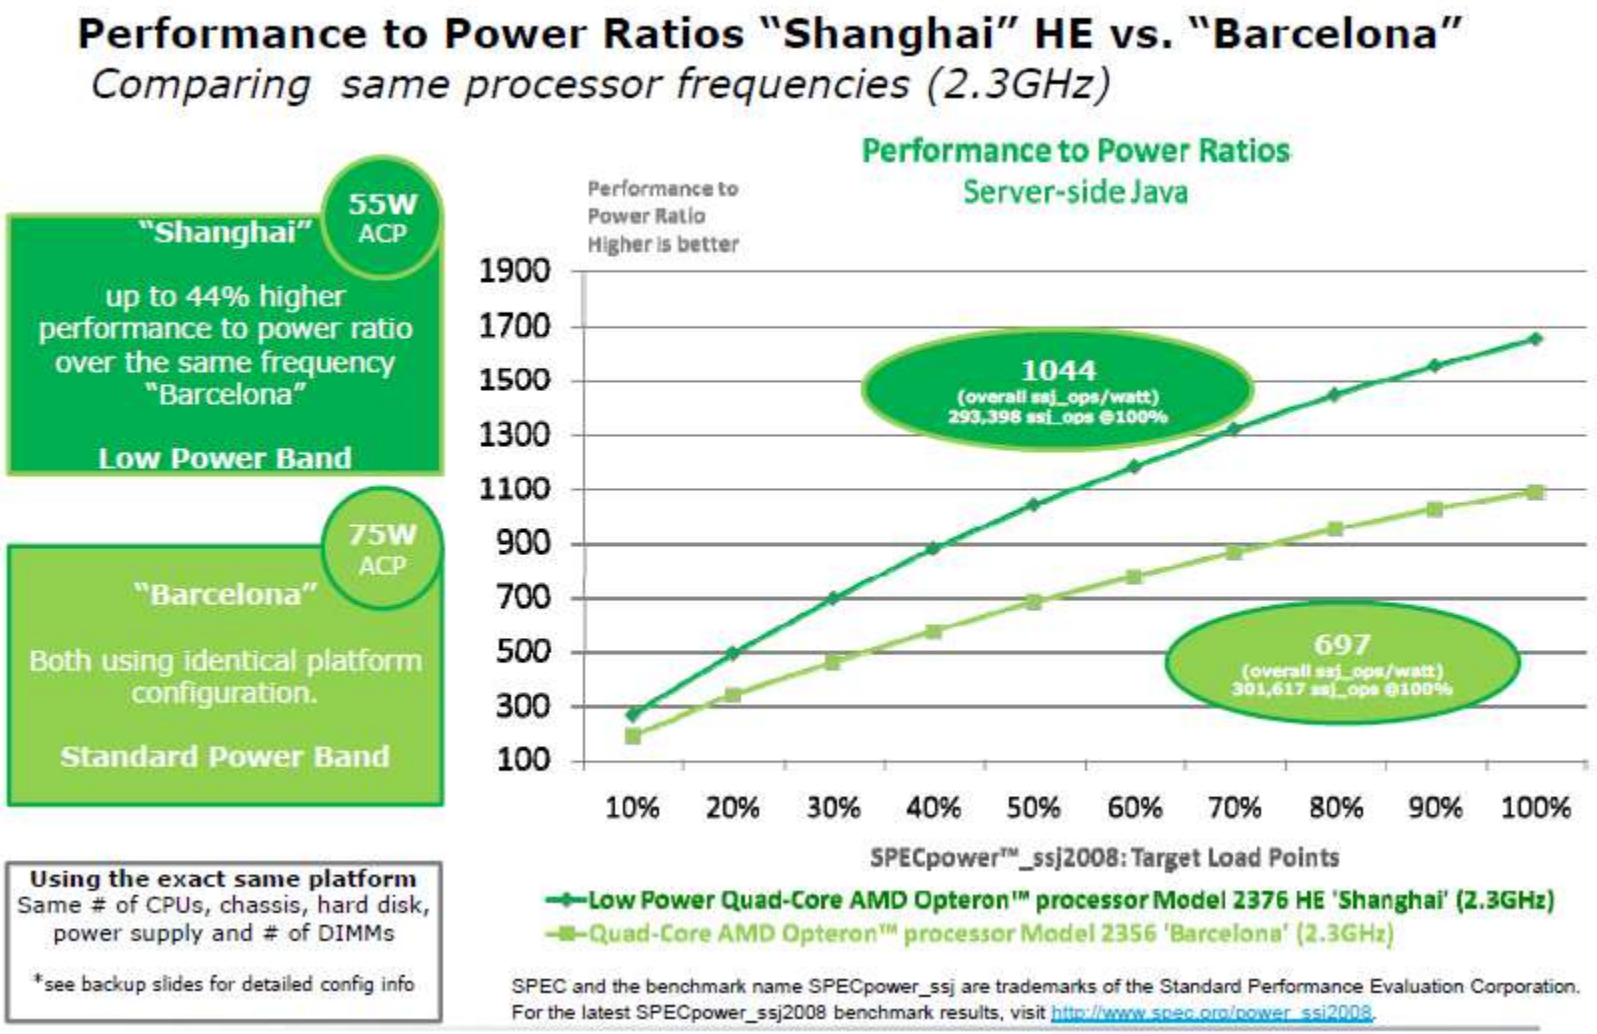

7.7.4 Key enhancements of the K10.5 Shanghai family targeting increased efficiency of virtualization and RAS

These topics will not be discussed here.

7.2 Key enhancements of the K10.5 Shanghai family (22)

Contrasting particular features of AMD's and Intel's processors [66]

Architecture	AMD "Barcelona"	AMD "Shanghai" & "Istanbul"	Intel "Harpertown"	Intel "Nehalem"
Example Server CPU series	Opteron 235x	Opteron 238x and Opteron 24xx	Xeon 54xx	Xeon 55xx, 35xx, 34xx
Dynamic frequency regulation	per core	per core	per CPU	Per CPU, turbo mode per core
Dynamic Voltage regulation	per CPU	per CPU	per CPU	Per CPU
lowest power state one Core idling	C1	C1	p-state of most loaded core	C6-state
effect of Core idling	lower frequency	clock gated: O Hz	possibly lower frequency	power gated
Cache sizing	no	no	yes	yes

7.2 Key enhancements of the K10.5 Shanghai family (23)

Die shot and floor plan of the Shanghai die [72]

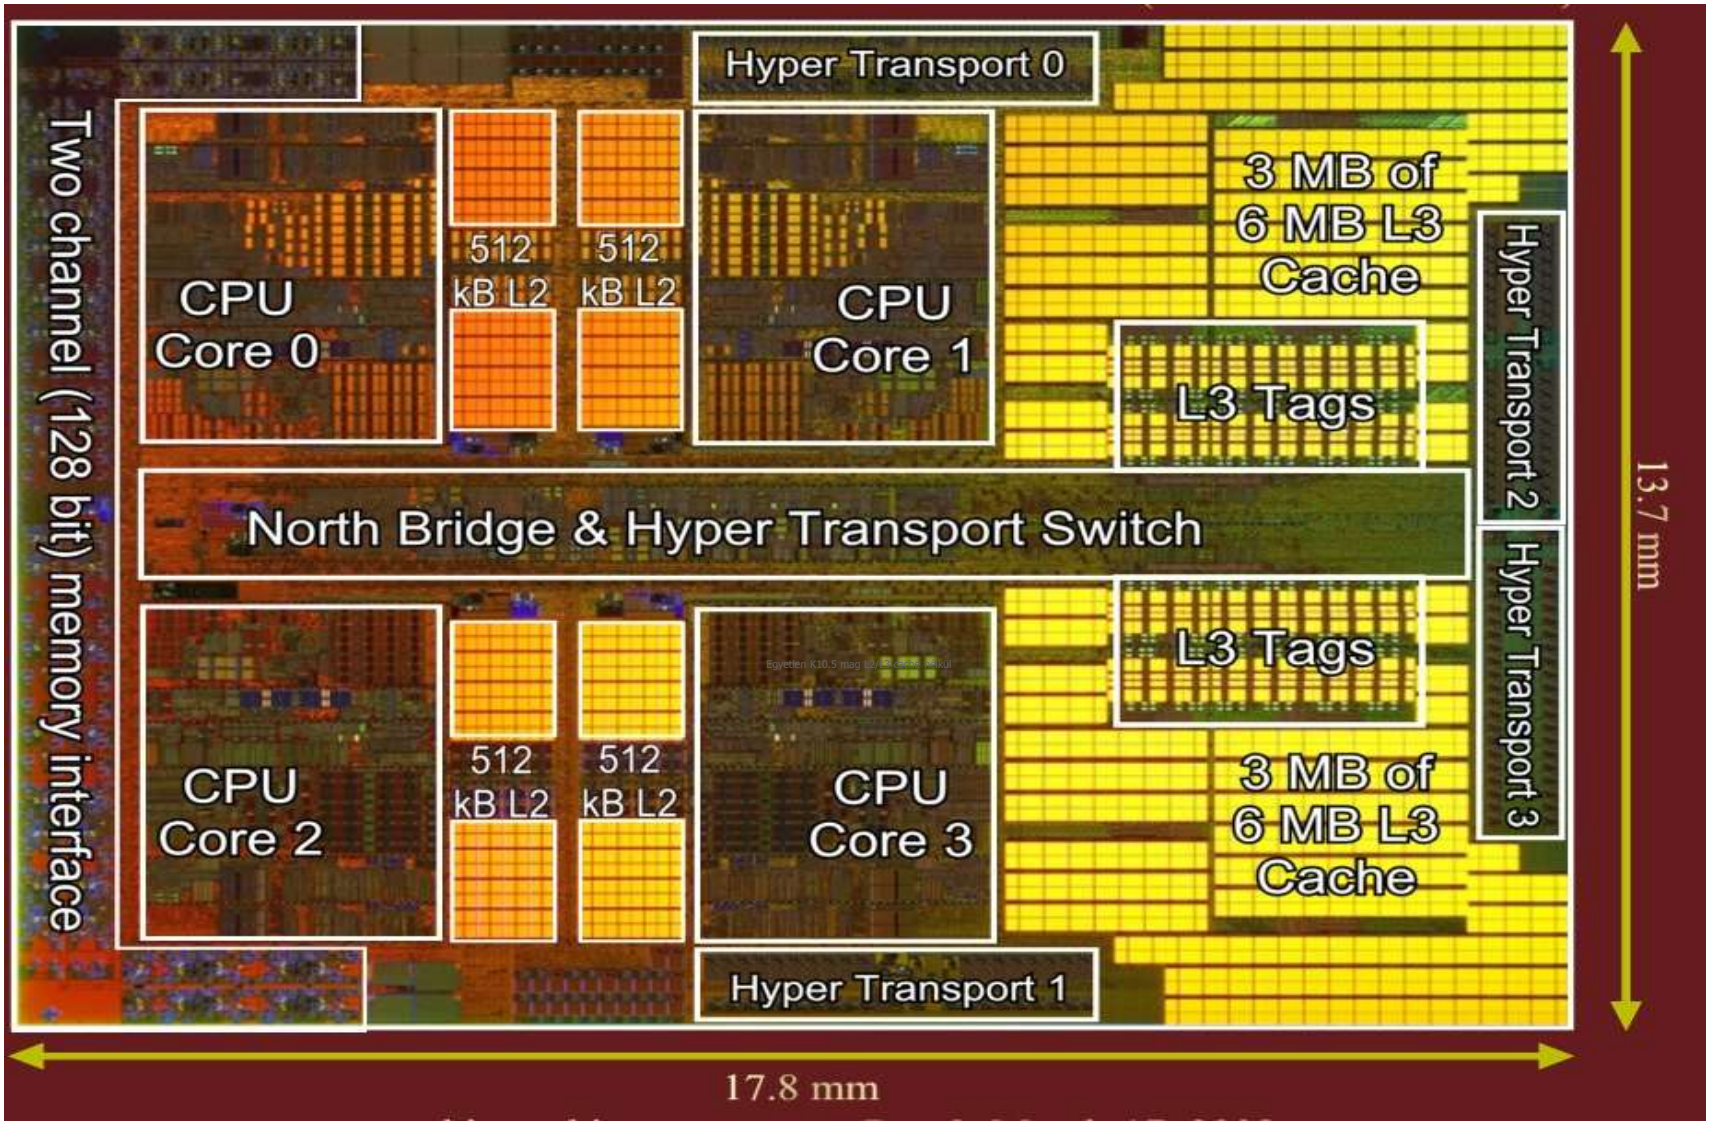

7.3 K10.5 Shanghai-based server lines

7.3 K10.5 Shanghai –based server lines

Overview of subsequent K10/K10.5 DP/MP server implementations [88]

65 nm
45nm

Platform Segment	2008	2009	2010
CPU	"Barcelona" 4-Core • 2M L3 • RDDR-2 • 3x HT-1 • AMD-V™ • 65nm	"Shanghai" 4-Core • 6M L3 • RDDR-2 • cHT-3 • AMD-V • 45nm	"Istanbul" 2H09 6-Core • 6M L3 • RDDR-2 • cHT-3 • AMD-V • 45nm
			"Magny-Cours" 1H10 12-Core • 12M L3 • APML • Probe Filter • 45nm • 4x HT-3 • AMD-V • HTC • U/R DDR-3
			"Sao Paulo" 1H10 6-Core • 6M L3 • APML • Probe Filter • 45nm • 4x HT-3 • AMD-V • HTC • U/R DDR-3
Chipset	Nvidia nForce 3600/3050 Broadcom HT-2100/1000		AMD RD890S w/IOMMU AMD RD870S w/IOMMU AMD SB700S
Platform	Socket F (1207) • 3x HT-1 (moving to cHT-3) • DDR-2 (Dual Channel)		"Maranello" • 4x HT-3 • DDR-3

7.3 K10.5 Shanghai-based server lines (2)

First introduced K10.5 Shanghai Opteron DP and MP models [141]

AMD Shanghai Overview		
Model	**CPU Clock**	**MC Clock**
Opteron 2384	2.7GHz	2.2GHz
Opteron 2382	2.6GHz	2.2GHz
Opteron 2380	2.5GHz	2.0GHz
Opteron 2378	2.4GHz	2.0GHz
Opteron 2376	2.3GHz	2.0GHz
Opteron 8384	2.7GHz	2.2GHz
Opteron 8382	2.6GHz	2.2GHz
Opteron 8380	2.5GHz	2.0GHz
Opteron 8378	2.4GHz	2.0GHz

7.4 K10.5 Shanghai-based desktop lines

7.4 K10.5 Shanghai-based desktop lines (1)

AMD's K10.5 Shanghai-based desktop lines - Overview [14]

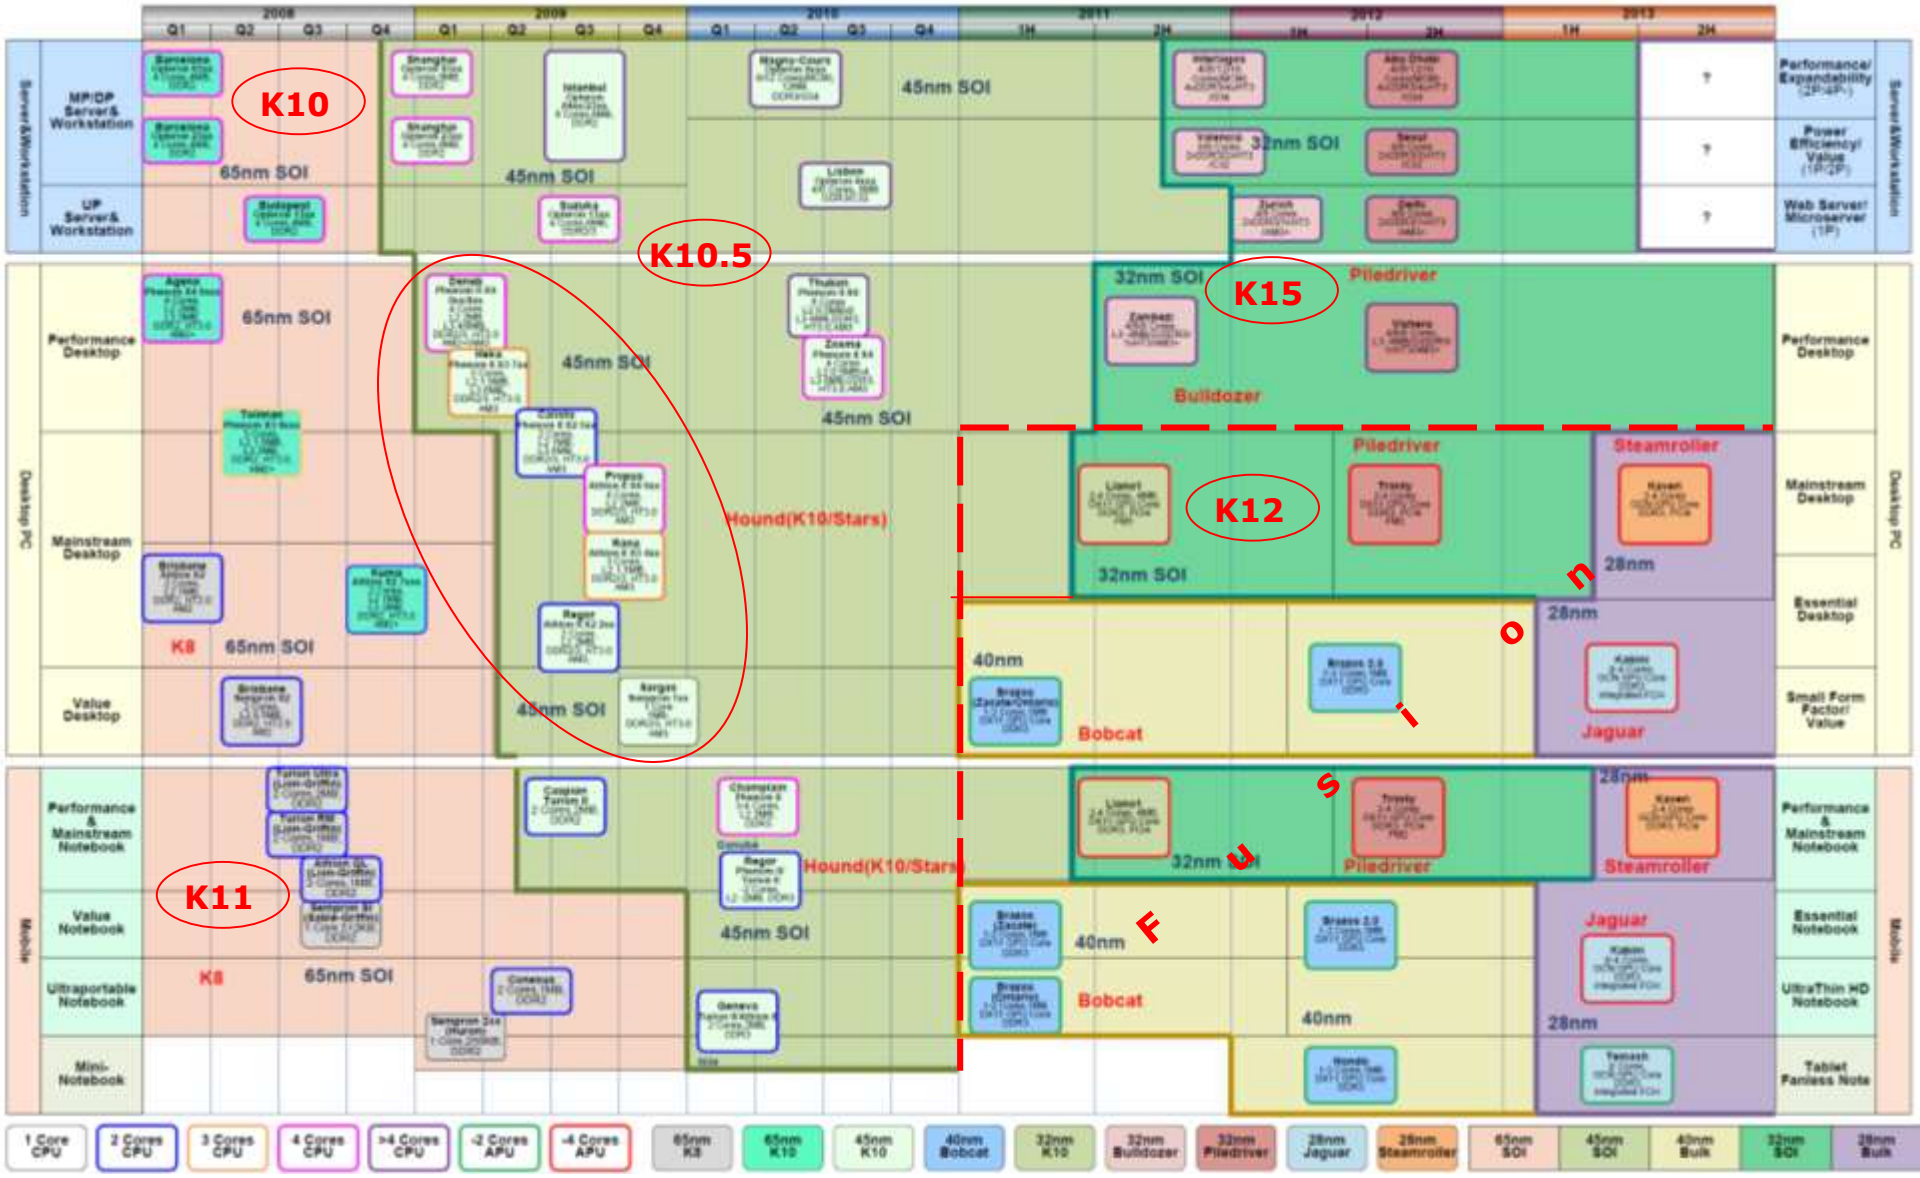

7.4 K10.5 Shanghai desktop lines

Brand names [68]

7.4 K10.5 Shanghai-based desktop lines (3)

Brand names of AMD's 10.5h Shanghai-based processor lines

		2003-2007	2007-2008	2008-2011	2009	2009
		K8 (Hammer)	**K10 (Barcelona)**	**K10.5 (Shanghai)**	**K10.5 (Istanbul)**	**K10.5 (Magny- Course)**
Servers	**4P servers**	**See Section 4**	**Barcelona (834x-836x))**	**Shanghai (837x-839x)**	**Istambul (8410-8430)**	**Magny-Course (6100)**
	2P servers		**Barcelona (234x-236x)**	**Shanghai (237x-239x)**	**Istambul (241x-243x)**	**Lisbon (4100)**
	1P servers		**Budapest (135x-136x)**	**Suzuka (138x-139x)**		
Desktops	**High perf.** (~80-120W)		Phenom X4-X2	Phenom II X4-X2	Phenom II X6-X4	
	Mainstream (~60-90W)	Athlon 64 Athlon 64 X2	Athlon X2	Athlon II X4-X2		
	Value (~40-60W)	Sempron		Sempron		
Mobiles	**High perf.** (~30-40W)	Turion 64 X2 (TL 6/5) Turion 64 (ML/MT)		Phenom II (N/P 9xx-6xx) Turion II Ultra (M6xx) Turion II (M/N/P 5xx)		
	Mainstream (~20-30W)	Athlon 64 X2 (TK-5x/4x) Athlon 64 (2xxx+-4xxx+)		Athlon II (M/N/P 3xx) Sempron (M1xx)		
	Ultraportable (~10-20W)	Mobile Sempron (2xxx+-4xxx+) Sempron 2100 fanless		Turion II Neo (K6xx) Athlon II Neo (K1xx) V-series (V1xx)		
Embedded (~10-20W)				Turion II Neo X2 Athlon II Neo X2 Athlon II Neo		

7.4 K10.5 Shanghai-based desktop lines (4)

AMD-s K10.5 Shanghai-based native desktop dies

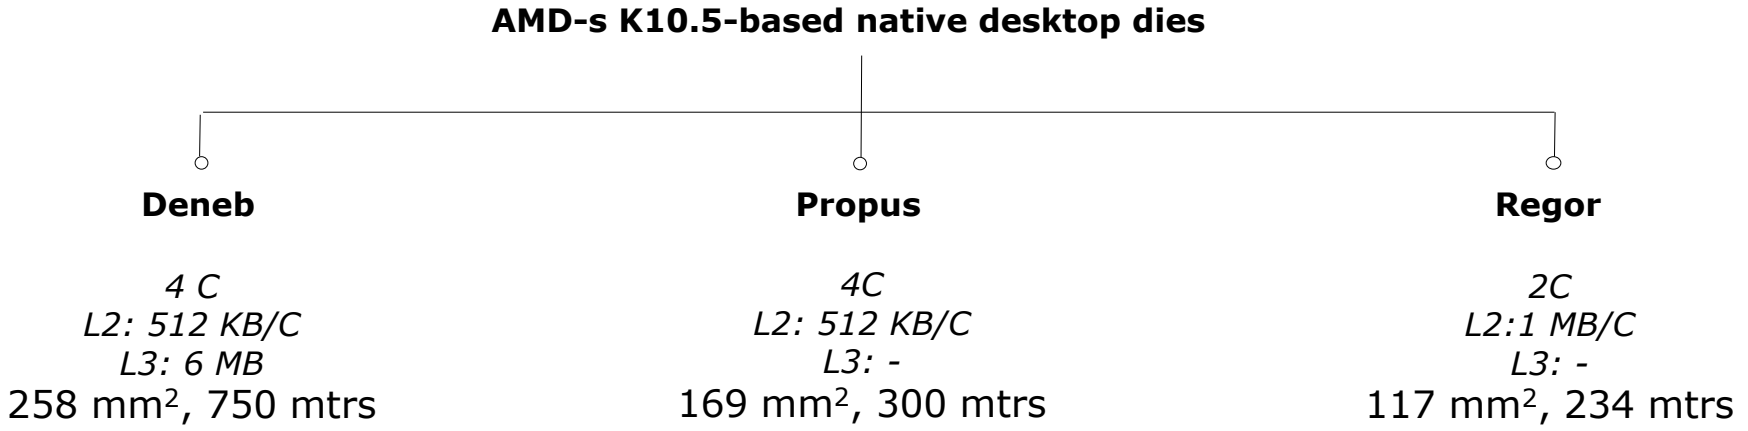

7.4 K10.5 Shanghai-based desktop lines (5)

AMD's K10.5 Shanghai-based desktop dies and desktop lines

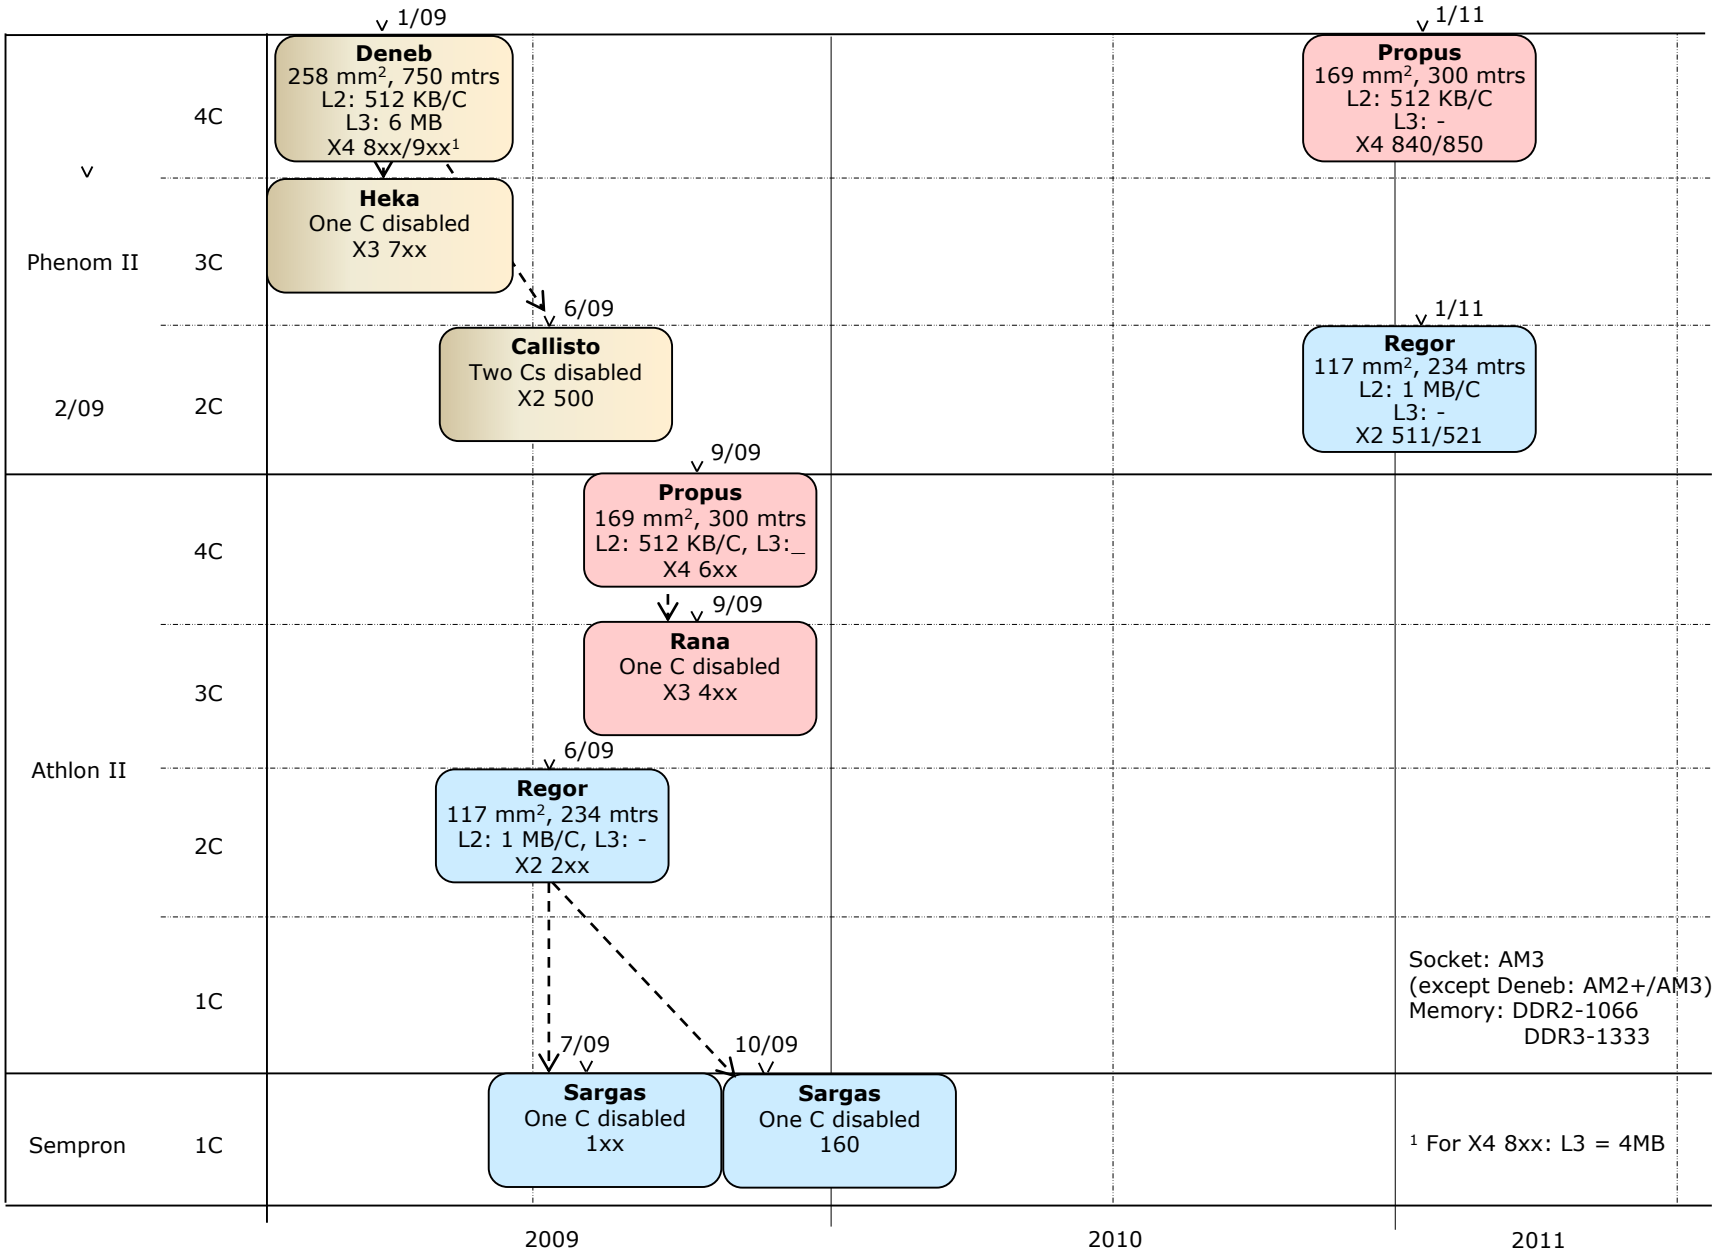

7.4 K10.5 Shanghai-based desktop lines (6)

Main features of AMD's high-performance K10.5 Deneb-based desktop line

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. 15h based lines		-	-	-	-	-	-	-	-	-	-

7.4 K10.5 Shanghai-based desktop lines (7)

AMD-s K10.5 Shanghai-based native desktop dies

Deneb

4 C
L2: 512 KB/C
L3: 6 MB
258 mm^2, 758 mtrs

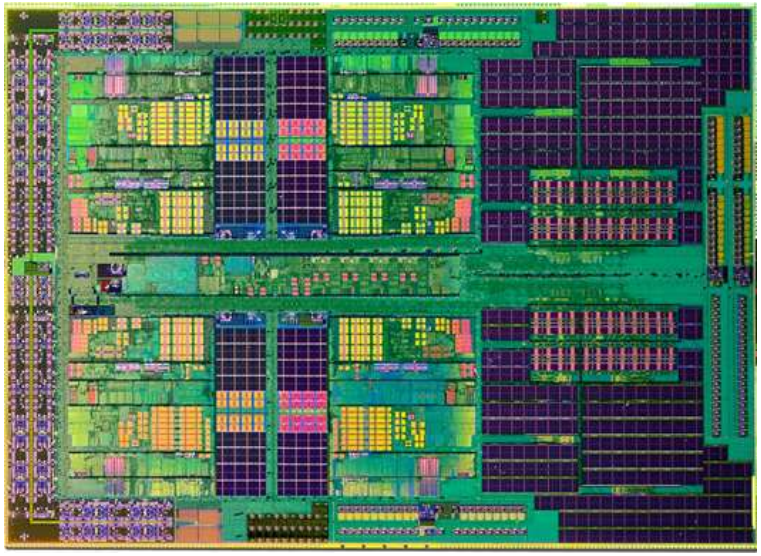

[69]

Propus

4C
L2: 512 KB/C
L3: -
169 mm^2, 300 mtrs

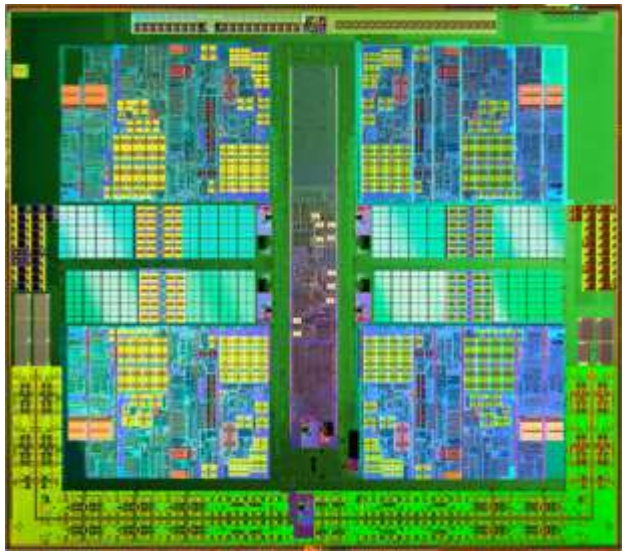

[70]

Regor

2C
L2:1 MB/C
L3: -
117 mm^2, 234 mtrs

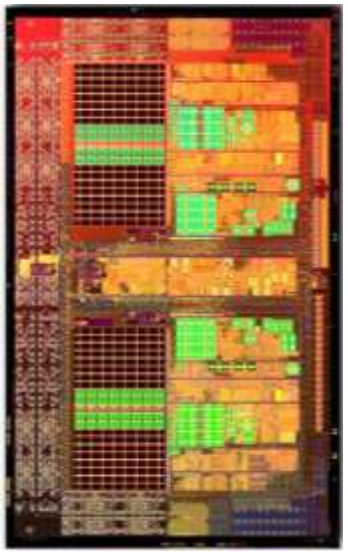

[71]

7.4 K10.5 Shanghai-based desktop lines (8)

Contrasting the K10.5 Shanghai-based server and Deneb desktop dies

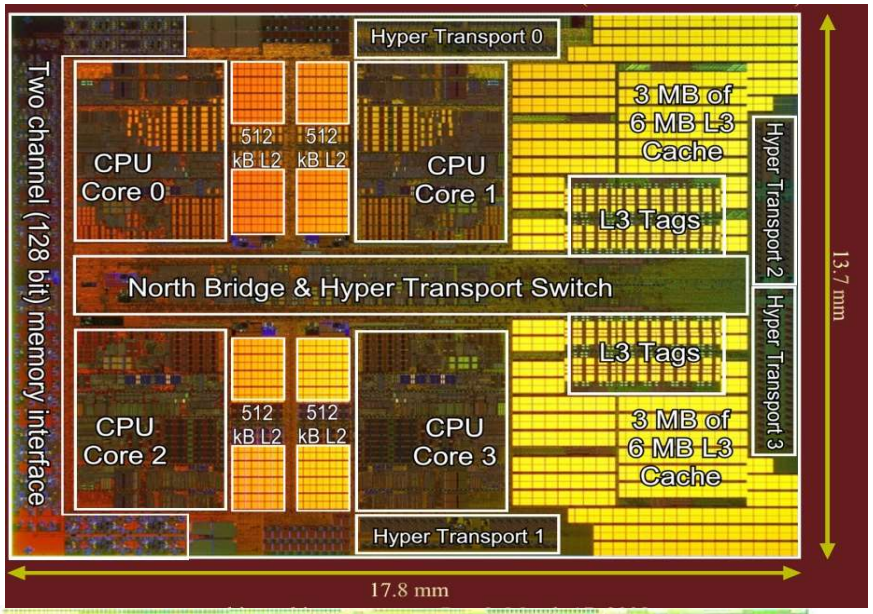

Shanghai core [72]

4 C
L2: 512 KB/C
L3: 6 MB
258 mm^2, 758 mtrs

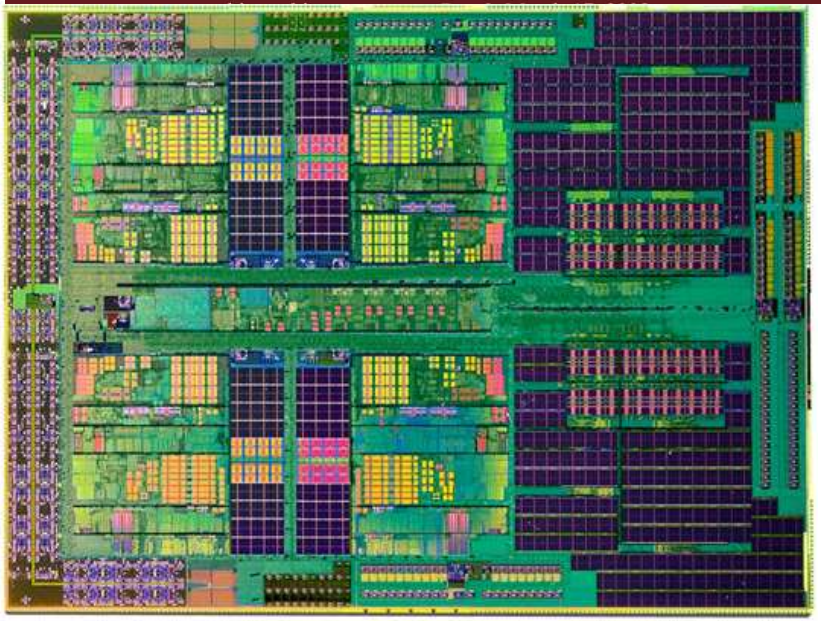

Deneb core [69]

4 C
L2: 512 KB/C
L3: 6 MB
258 mm^2, 758 mtrs

7.4 K10.5 Shanghai-based desktop lines (9)

Contrasting the K10.5-based Deneb and Propus desktop dies

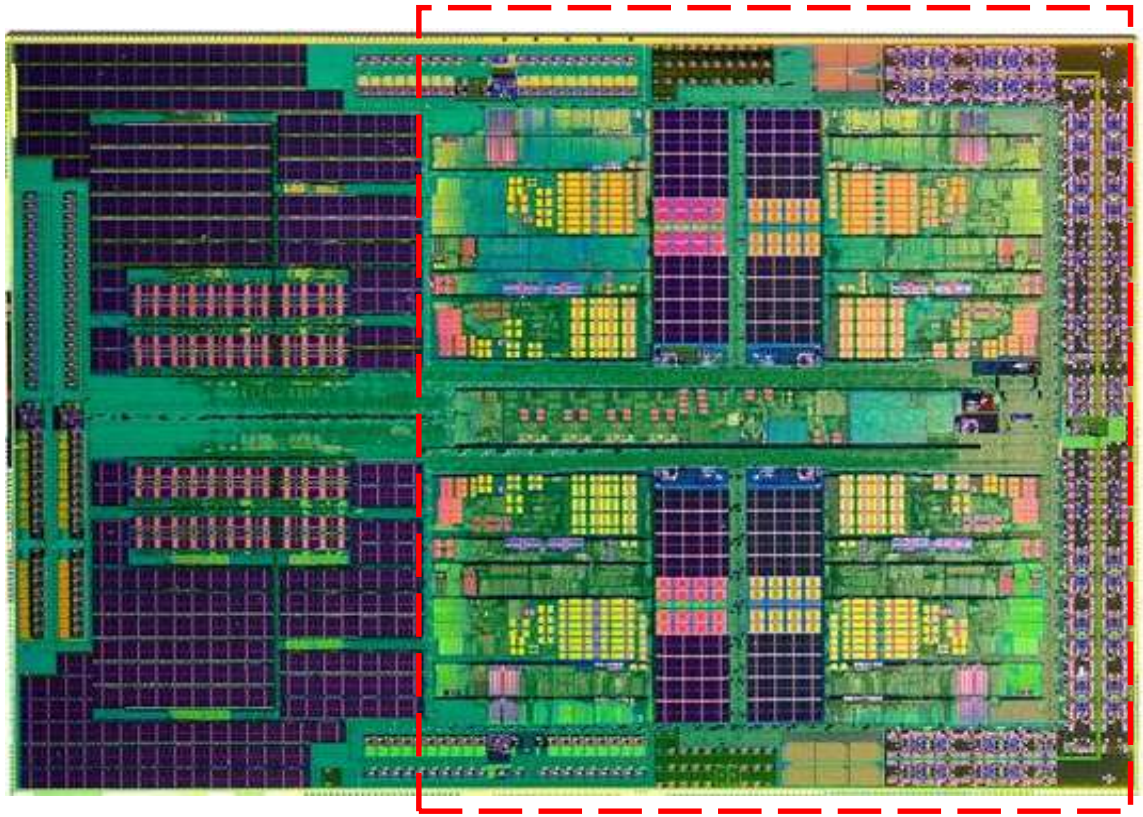

Deneb 45 nm [69]

Performance desktop
L3: 6 MB
258 mm^2, 758 mtrs

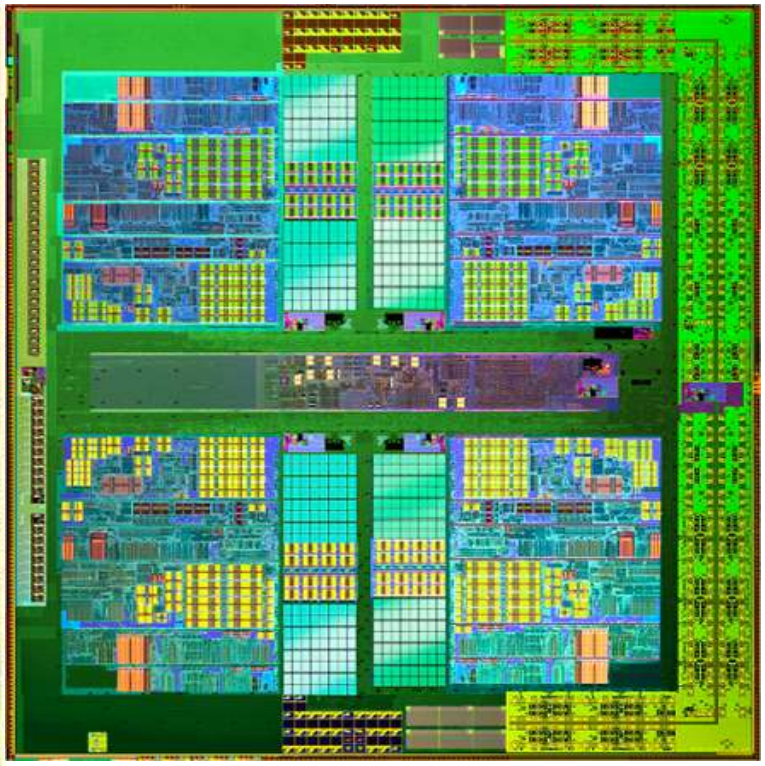

Propus 45 nm [70]

Mainstream desktop
no L3
169 mm^2, 300 mtrs

7.4 K10.5 Shanghai-based desktop lines (10)

Contrasting the K10.5-based Deneb and Regor desktop dies

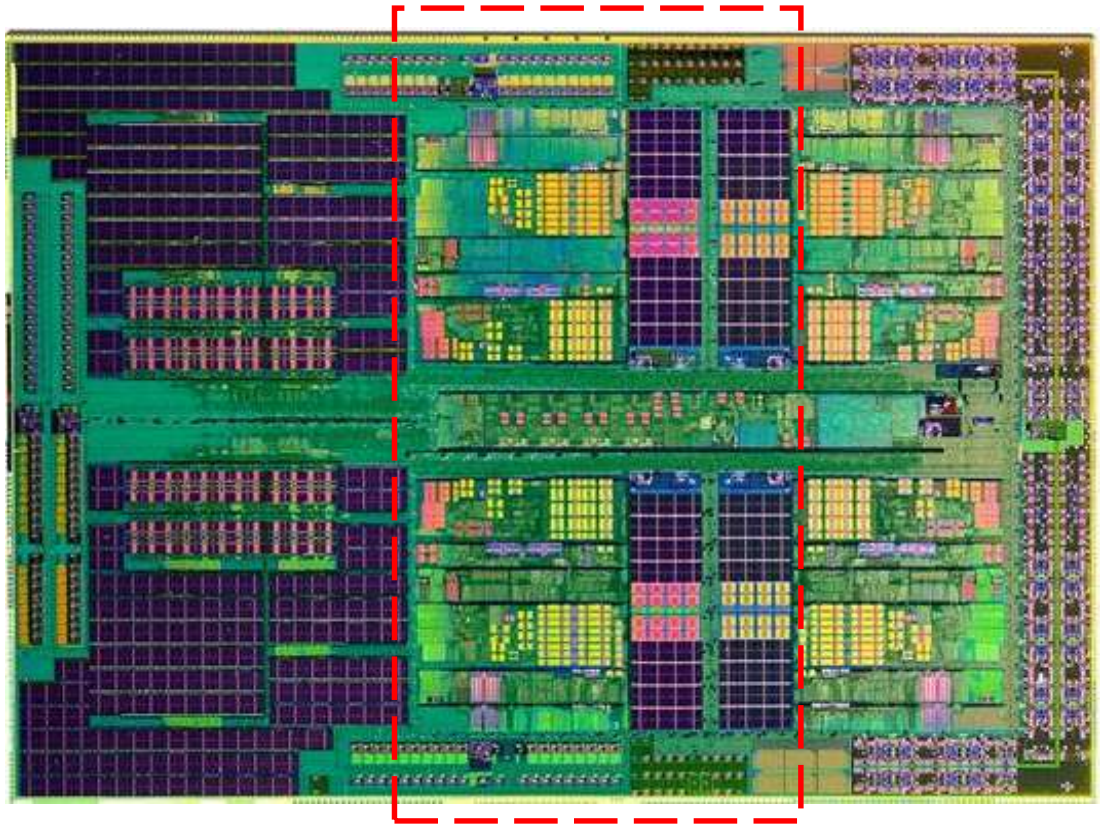

Deneb 45 nm [69]
Performance desktop
L3: 6 MB
258 mm^2, 758 mtrs

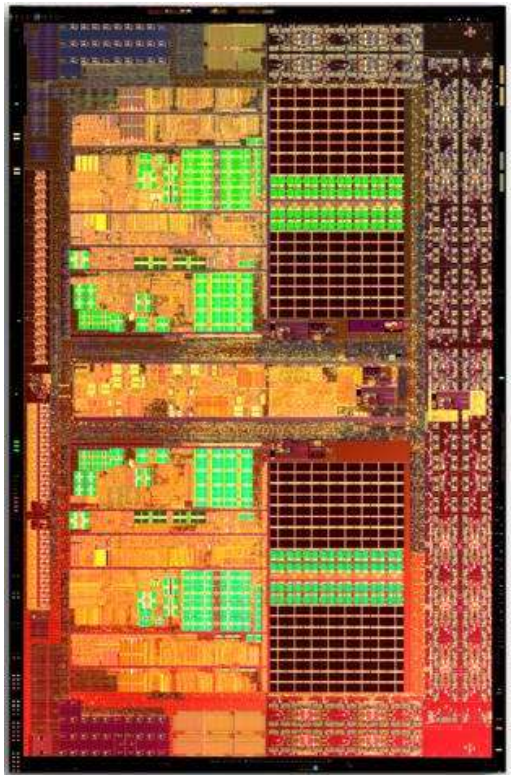

Regor 45 nm [71]
Mainstream desktop
no L3
117 mm^2, 234 mtrs

7.4 K10.5 Shanghai-based desktop lines (11)

Deriving the Deneb X4 800/Heka/Calisto lines from the Deneb X4 900 line [73]

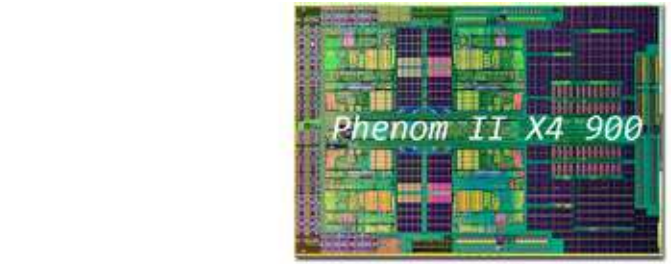

758 mtrs 258 mm2

Deneb

One core disabled

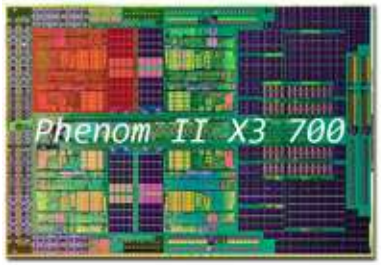

Heka

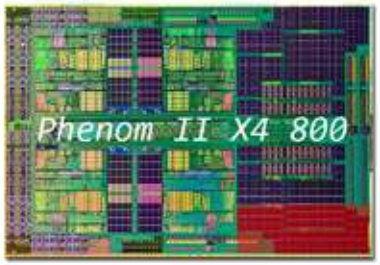

Deneb

2 MB L3 disabled

Two cores disabled

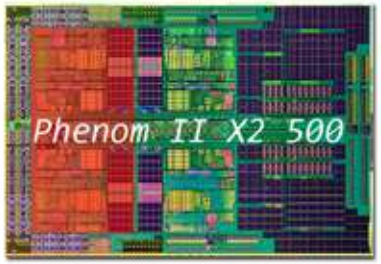

Callisto

7.4 K10.5 Shanghai-based desktop lines (12)

Remarks to the implementation of DVFS in K10.5 Shanghai-based (Phenom II) processors
[74], [75]

- In the K10 Barcelona-based Phenom line all four cores could run at different clock speed, i.e. all cores had separate clock domains.

  But under Windows Vista separate clock domains cause an efficiency problem for single threaded applications as the OS scheduler migrated the single thread to all cores in a round-robin fashion.

  Switching threads from one core to another, however reduces efficiency.
- To avoid this problem, in the Phenom II line AMD switched back to their previous core coordination policy and use BIOS code to lock clock frequencies of all four cores together.
- Nevertheless, in Windows 7 (introduced in 10/2009) Microsoft changed the scheduler policy of the OS and for single threaded applications lets run a single core without interruption, so all idle cores can enter a low power state (C1E).
- Accordingly, in their Phenom II X6 Istambul line (introduced in 4/2010) AMD implemented again separate clock planes for the cores (but a common voltage plane for all six cores).

7.5 K10.5 Shanghai-based mobile lines

7.5 K10.5 Shanghai-based mobile lines (1)

7.5 K10.5 Shanghai-based mobile lines - Overview [14]

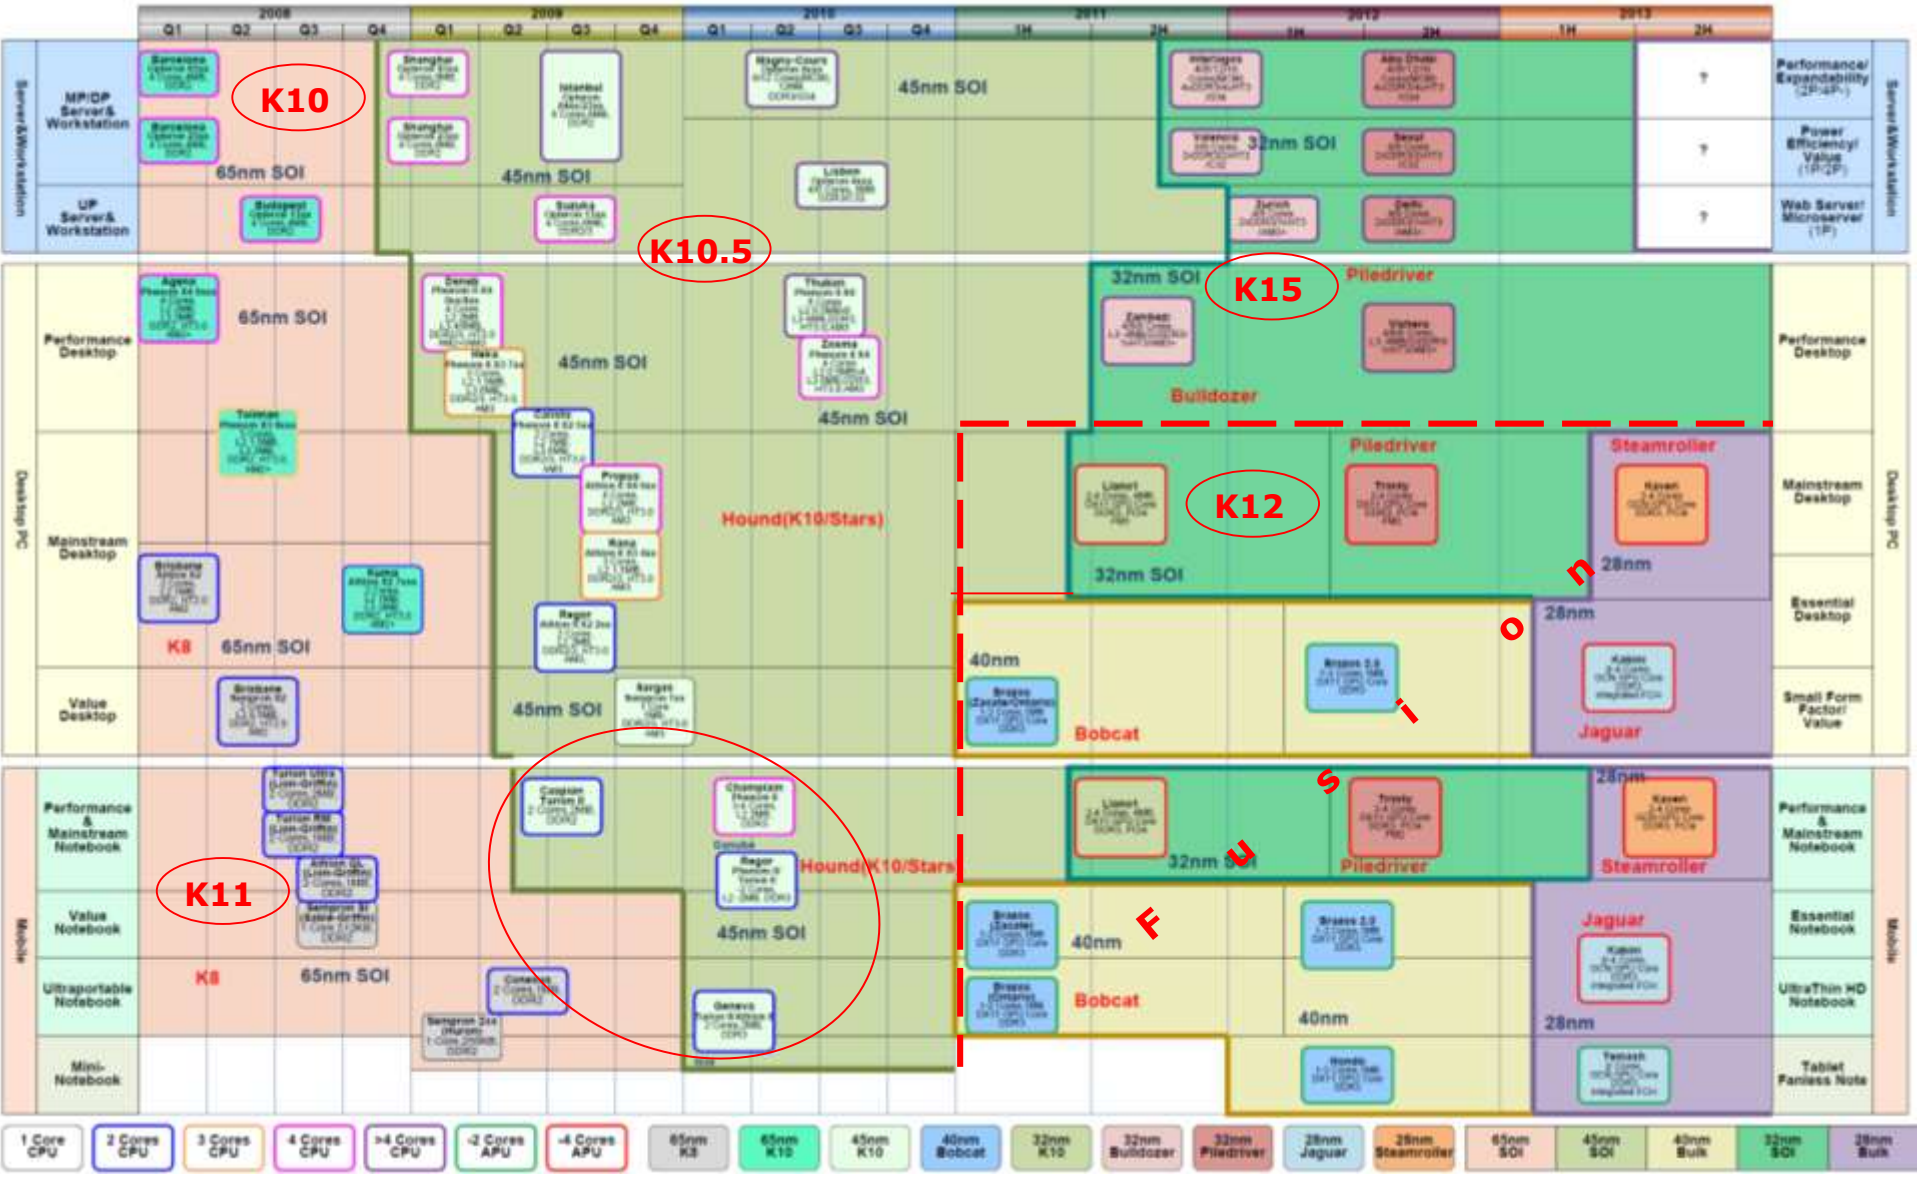

7.5 K10.5 Shanghai-based mobile lines (2)

Brand names of AMD's K10.5h Shanghai-based mobile lines

		2003-2007	2007-2008	2008-2011	2009	2009
		K8 (Hammer)	**K10 (Barcelona)**	**K10.5 (Shanghai)**	**K10.5 (Istanbul)**	**K10.5 (Magny- Course)**
Servers	**4P servers**	**See Section 4**	**Barcelona (834x-836x))**	**Shanghai (837x-839x)**	**Istambul (8410-8430)**	**Magny-Course (6100)**
	2P servers		**Barcelona (234x-236x)**	**Shanghai (237x-239x)**	**Istambul (241x-243x)**	**Lisbon (4100)**
	1P servers		**Budapest (135x-136x)**	**Suzuka (138x-139x)**		
Desktops	**High perf.** (~80-120W)		Phenom X4-X2	Phenom II X4-X2	Phenom II X6-X4	
	Mainstream (~60-90W)	Athlon 64 Athlon 64 X2	Athlon X2	Athlon II X4-X2		
	Value (~40-60W)	Sempron		Sempron		
Mobiles	**High perf.** (~30-40W)	Turion 64 X2 (TL 6/5) Turion 64 (ML/MT)		Phenom II (N/P 9xx-6xx) Turion II Ultra (M6xx) Turion II (M/N/P 5xx)		
	Mainstream (~20-30W)	Athlon 64 X2 (TK-5x/4x) Athlon 64 (2xxx+-4xxx+)		Athlon II (M/N/P 3xx) Sempron (M1xx)		
	Ultraportable (~10-20W)	Mobile Sempron (2xxx+-4xxx+) Sempron 2100 fanless		Turion II Neo (K6xx) Athlon II Neo (K1xx) V-series (V1xx)		
Embedded (~10-20W)				Turion II Neo X2 Athlon II Neo X2 Athlon II Neo		

7.5 K10.5 Shanghai-based mobile lines (3)

AMD's K10.5 Shanghai-based native mobile dies

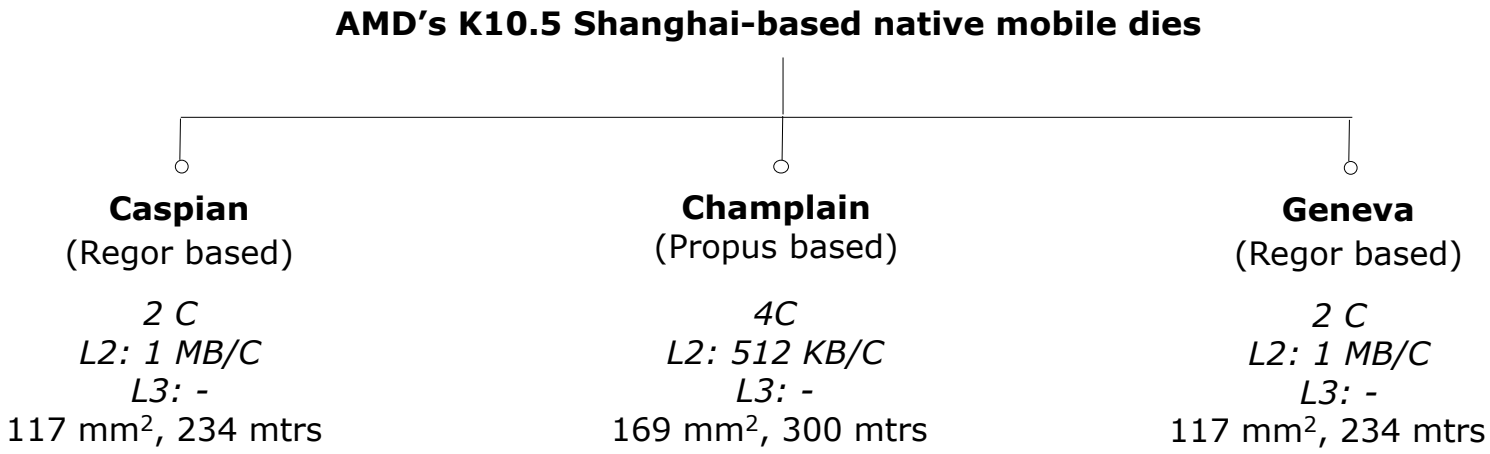

7.5 K10.5 Shanghai-based mobile lines (4)

AMD's K10.5 Shanghai-based native mobile dies and mobile lines-1

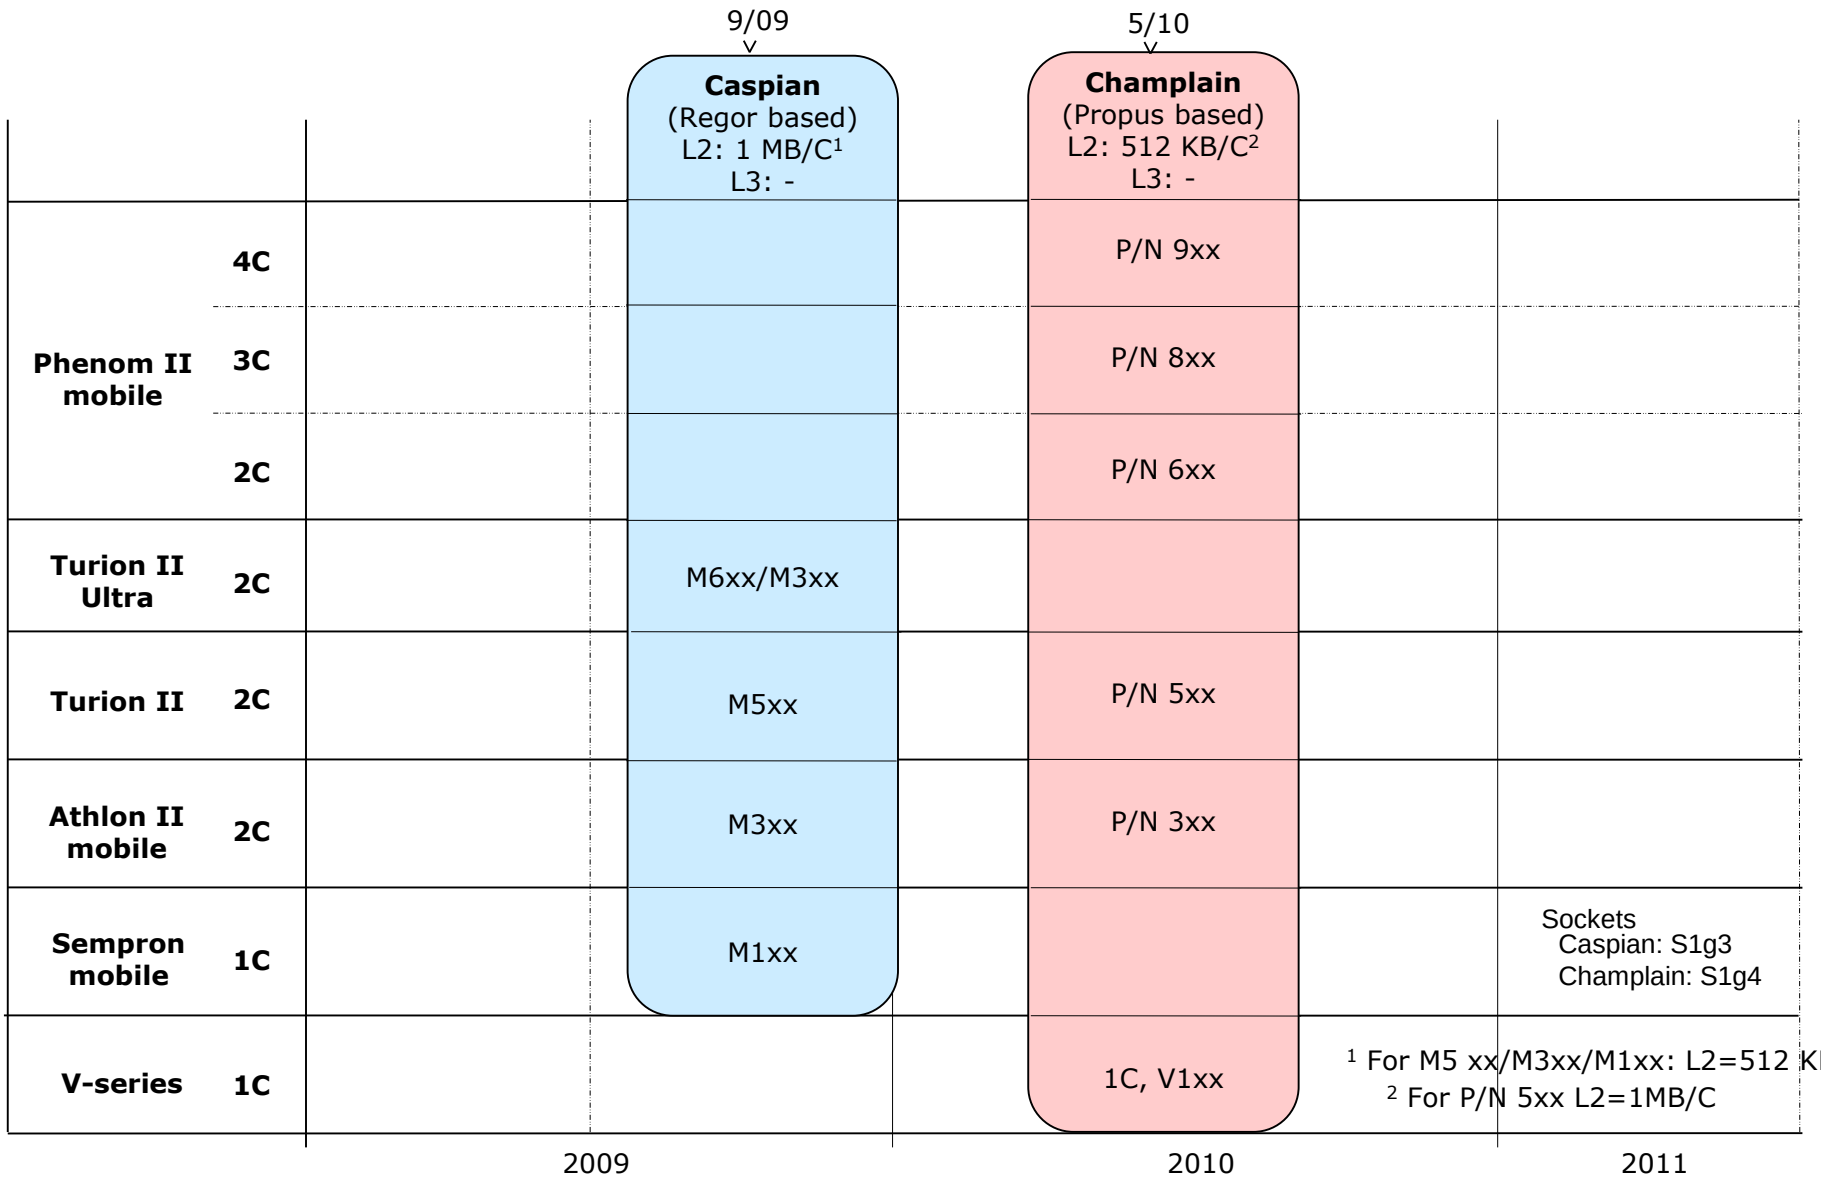

7.5 K10.5 Shanghai-based mobile lines (5)

Main features of AMD's high performance K10.5 Shanghai-based mobile lines

Base arch./ stepping		Intro	High perf. mobile family name	Series	Techn.	Core count (up to)	L2 (up to)	L3	Memory (up to)	HT/ dir. (up to)	Sock et
K8	C0, CG	9/2003	Clawhammer	Mobile Athlon 64	130 nm	1	512 KB	-	DDR-400	HT 1.0: 3.2 GB/s	754
	E5	3/2005	Lancaster	Turion 64	90 nm	1	1 MB	-	DDR-400	HT 1.0: 3.2 GB/s	754
	F2	5/2006	Trinidad	Turion 64 X2	90 nm	2	2*512 KB	-	DDR2-667	HT 1.0: 3.2 GB/s	S1
K10	-	-	-	-	-	-	-	-	-	-	-
K10.5	DA-C2	9/2009	Caspian	Turion II	45 nm	2	2*512 KB/ 2*1 MB[1]	-	DDR2-800	HT 3.0: 7.2 GB/s	S1g3
	DA-C3	5/2010	Champlain	Turion X4	45 nm	4	4*512 KB	-	DDR3-1066	HT 3.0: 7.2 GB/s	S1g4

[1]: 2*512 KB for Turion II, 2*1 MB for Turion II Ultra

7.5 K10.5 Shanghai-based mobile lines (6)

AMD-s K10.5 Shanghai-based native mobile dies

Caspian/Geneva)
(Regor based)

2 C
L2: 1 MB/C
L3: -
117 mm^2, 234 mtrs

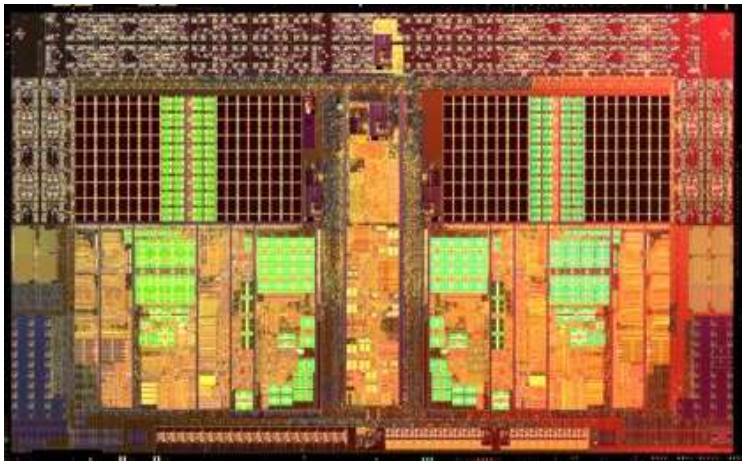

[77]

Champlain
(Propus based)

4C
L2: 512 KB/C
L3: -
169 mm^2, 300 mtrs

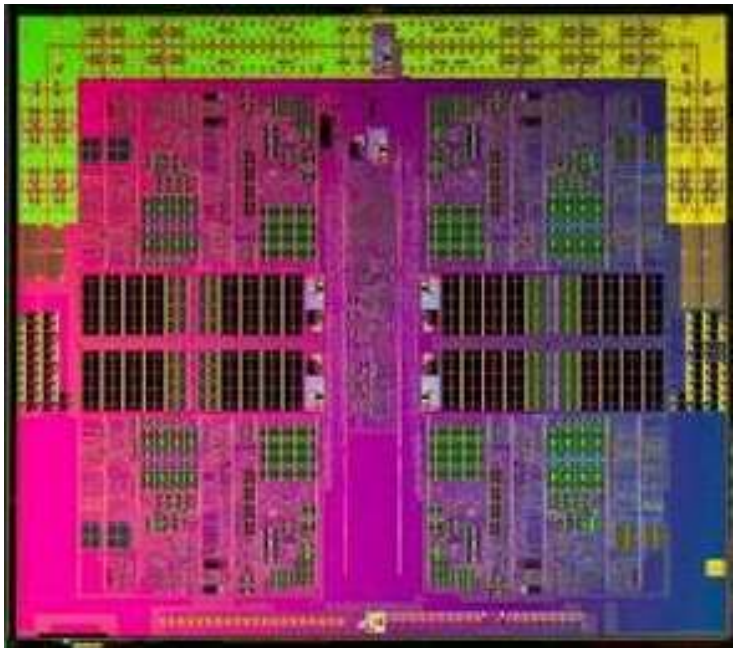

[78]

7.5 K10.5 Shanghai-based mobile lines (7)

Contrasting the Caspian [77] and the Regor [71] dies

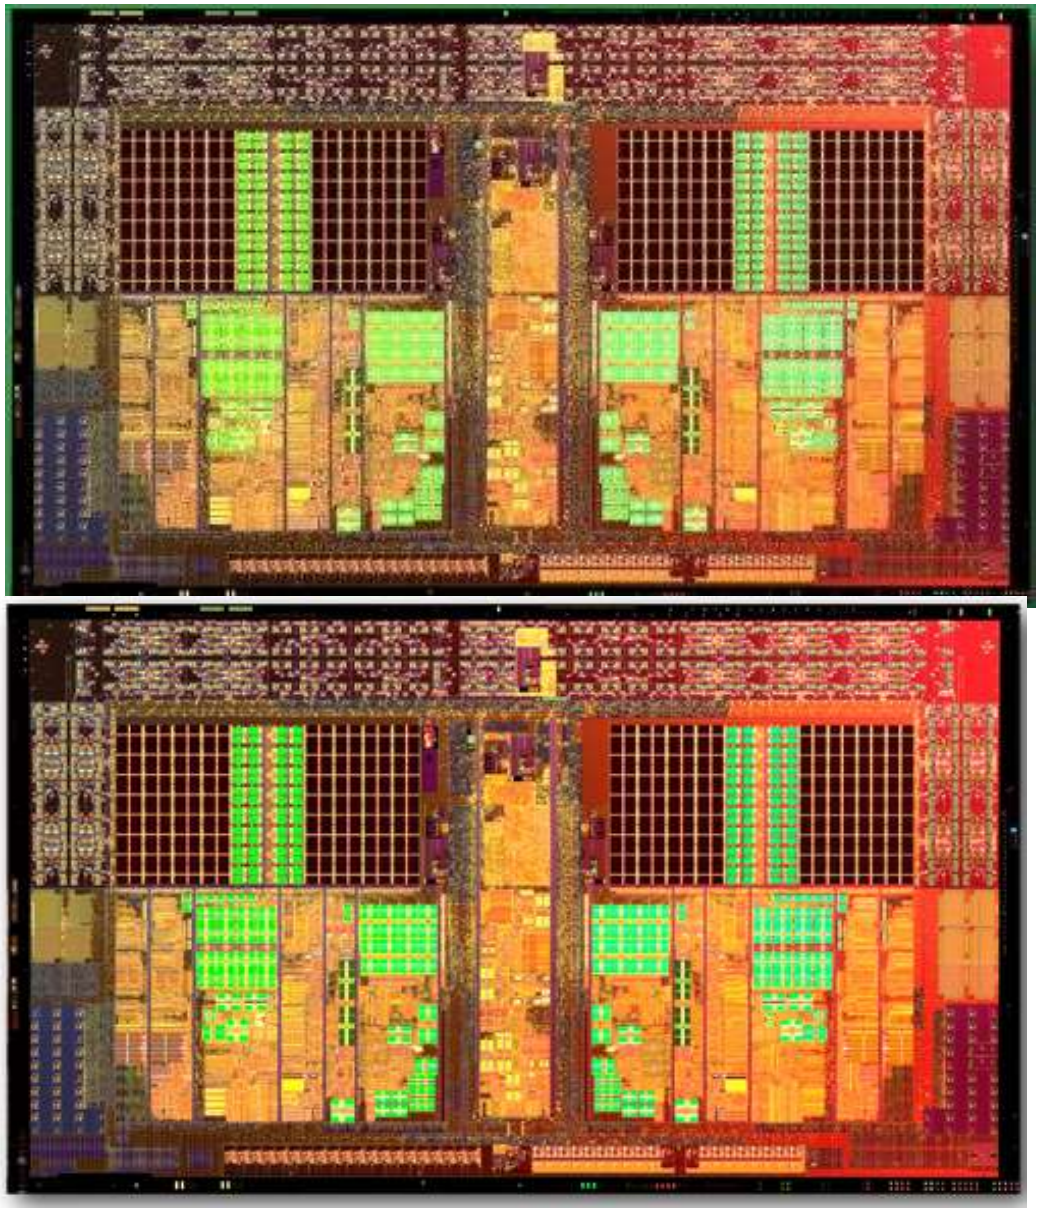

Caspian
45 nm
mainstream
notebook
die
117 mm^2
234 mtrs

Regor
45 nm
desktop
die

117 mm2
234 mtrs

7.5 K10.5 Shanghai-based mobile lines (8)

AMD's K10.5 Shanghai-based mobile lines-2 (45 nm)

Ultraportable mobiles

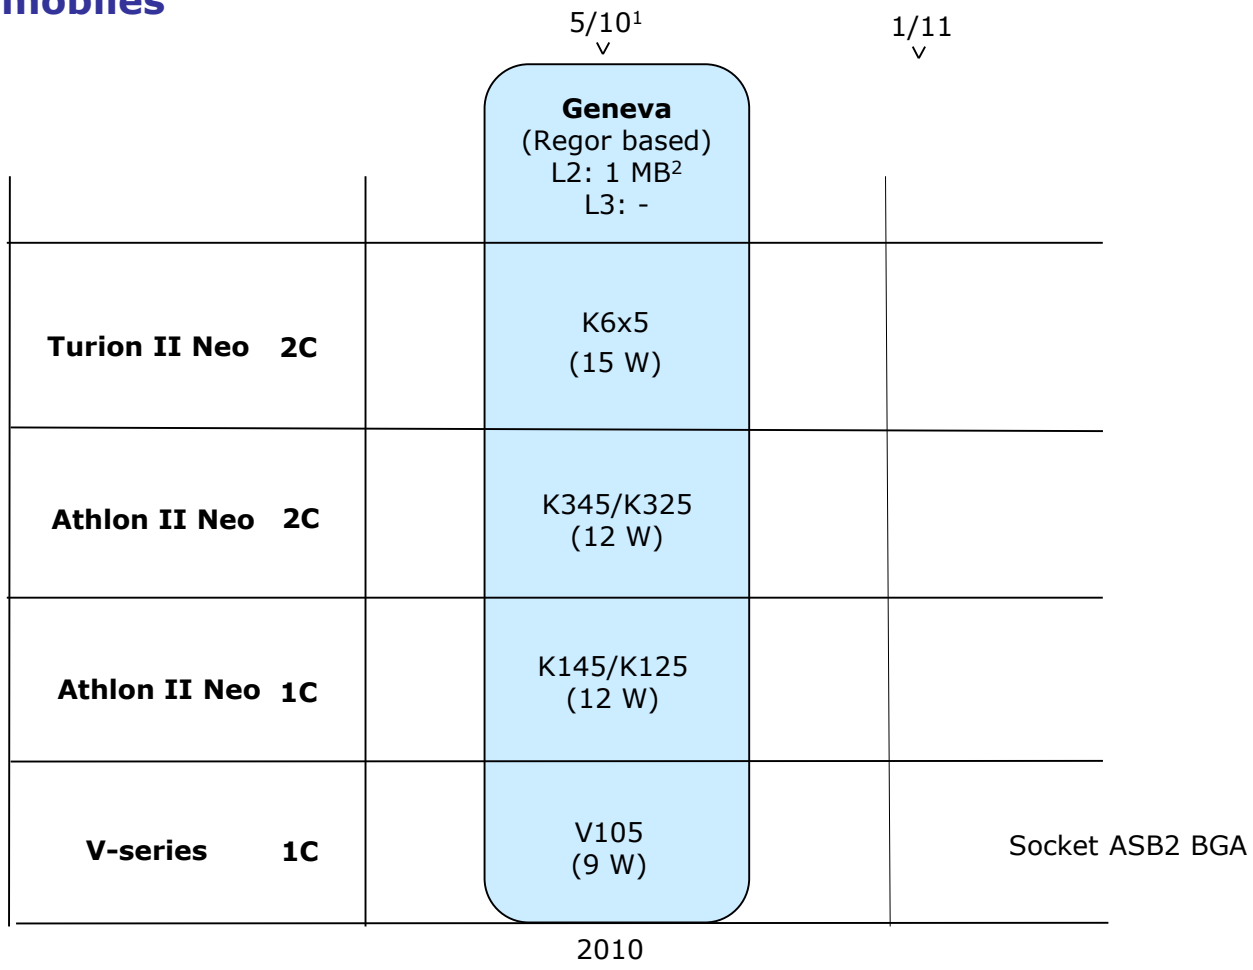

[1] Some of the models indicated were introduced only in 1/2011

[2] In V105: L2=512 KB

7.5 K10.5 Shanghai-based mobile lines (9)

The ASB2 BGA Socket

It is an update of the ASB1 BGA Socket introduced one year before.

The ASB1 BGA Socket [102]

The ASB1 family of processors are soldered directly to the board and reduce vertical height requirements from up to 8.6mm to 2.03mm. This enables small form factor and rugged designs to be created.

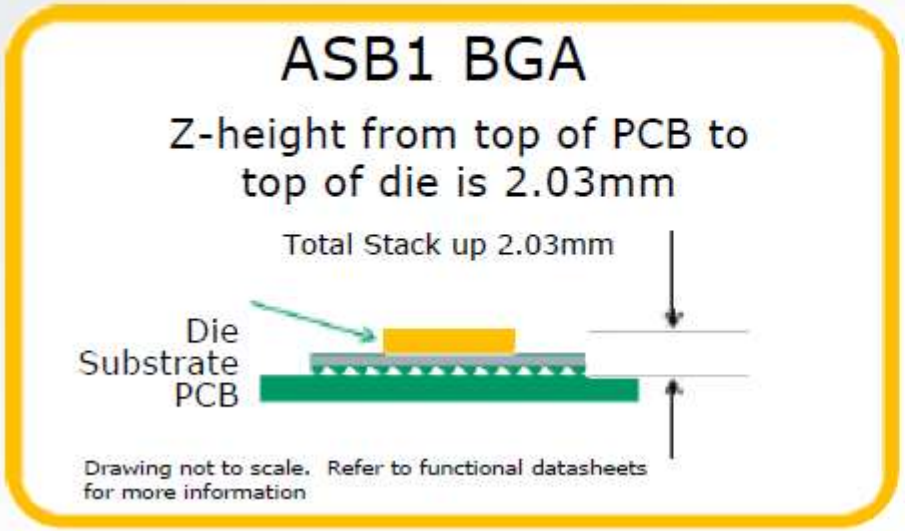

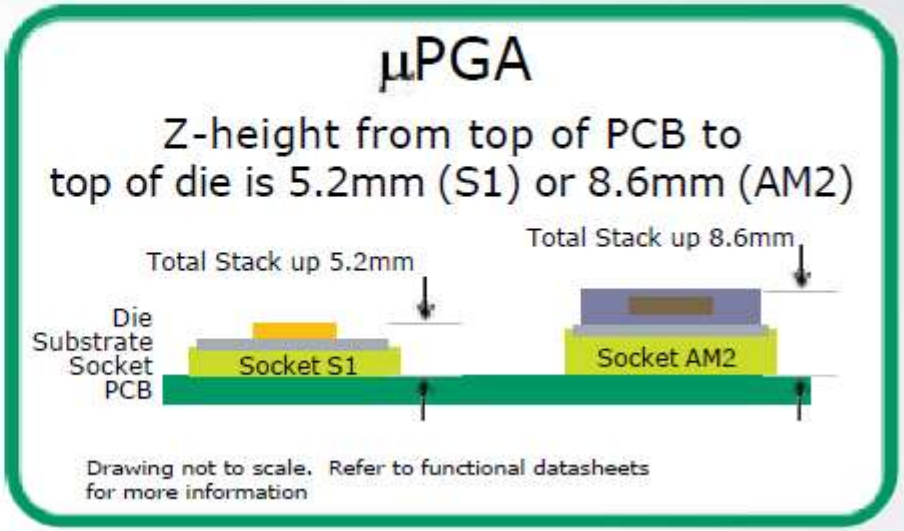

7.5 K10.5 Shanghai-based mobile lines (10)

Remark

The Geneva-based ultraportable processor family is part of AMD's Nile mobile platform and is actually AMD's third ultraportable family.

The first two ultraportable families were as follows [150]:

1. Ultraportable Yukon platform:
   Based on the 65 nm single core Huron die (belonging to the Family 11h Griffin lines).
   Introduced in 1/2009, available in 4/2009.
2. Ultraportable Congo platform:
   Based on the 65 nm dual core Conesus die (belonging to the Family 11h Griffin lines).
   Introduced in in 8/2009.

Both processor lines had ASB1 BGA sockets to allow direct soldering onto the mainboard.

AMD's ultraportable processor families can be considered as competing parts to Intel's 45 nm low power Atom family announced in 3/2008 with availability in 4/2009.

7.5 K10.5 Shanghai-based mobile lines (11)

Main features of K10.5 Shanghai-based ultraportable mobiles [7]

Turion II Neo

Model Number	Frequency	L2-Cache	FPU width[8]	HT	Multiplier[1]	TDP	Socket	Release date
Turion II K625	1500 MHz	2 × 1 MB	128-bit	1600 MHz	7.5x	15 W	Socket ASB2	May 12, 2010
Turion II K645	1600 MHz	2 × 1 MB	128-bit	1600 MHz	8x	15 W	Socket ASB2	January 4, 2011
Turion II K665	1700 MHz	2 × 1 MB	128-bit	1600 MHz	8.5x	15 W	Socket ASB2	May 12, 2010
Turion II K685	1800 MHz	2 × 1 MB	128-bit	1600 MHz	9x	15 W	Socket ASB2	January 4, 2011

Athlon II Neo

Model Number	Frequency	L2-Cache	FPU width[8]	HT	Multiplier[1]	TDP	Socket	Release date
Athlon II K325	1300 MHz	2 × 1 MB	64-bit	1000 MHz	6.5x	12 W	Socket ASB2	May 12, 2010
Athlon II K345	1400 MHz	2 × 1 MB	64-bit	1000 MHz	7x	12 W	Socket ASB2	January 4, 2011

Athlon II Neo

Model Number	Frequency	L2-Cache	FPU width[8]	HT	Multiplier[1]	TDP	Socket	Release date
Athlon II K125	1700 MHz	1 MB	64-bit	1000 MHz	8.5x	12 W	Socket ASB2	May 12, 2010
Athlon II K145	1800 MHz	1 MB	64-bit	1000 MHz	9x	12 W	Socket ASB2	January 4, 2011

V series processors

Model Number	Frequency	L2-Cache	FPU width[8]	HT	Multiplier[1]	TDP	Socket	Release date
V 105	1200 MHz	512 KB	64-bit	1000 MHz	6x	9 W	Socket ASB2	May 12, 2010

7.6 K10.5 Shanghai-based embedded lines

7.6 K10.5 Shanghai-based embedded lines (1)

Brand names of AMD's K10.5 Shanghai-based embedded lines

		2003-2007	2007-2008	2008-2011	2009	2009
		K8 (Hammer)	**K10 (Barcelona)**	**K10.5 (Shanghai)**	**K10.5 (Istanbul)**	**K10.5 (Magny- Course)**
Servers	**4P servers**	**See Section 4**	**Barcelona (834x-836x))**	**Shanghai (837x-839x)**	**Istambul (8410-8430)**	**Magny-Course (6100)**
	2P servers		**Barcelona (234x-236x)**	**Shanghai (237x-239x)**	**Istambul (241x-243x)**	**Lisbon (4100)**
	1P servers		**Budapest (135x-136x)**	**Suzuka (138x-139x)**		
Desktops	**High perf.** (~80-120W)		Phenom X4-X2	Phenom II X4-X2	Phenom II X6-X4	
	Mainstream (~60-90W)	Athlon 64 Athlon 64 X2	Athlon X2	Athlon II X4-X2		
	Value (~40-60W)	Sempron		Sempron		
Mobiles	**High perf.** (~30-40W)	Turion 64 X2 (TL 6/5) Turion 64 (ML/MT)		Phenom II (N/P 9xx-6xx) Turion II Ultra (M6xx) Turion II (M/N/P 5xx)		
	Mainstream (~20-30W)	Athlon 64 X2 (TK-5x/4x) Athlon 64 (2xxx+-4xxx+)		Athlon II (M/N/P 3xx) Sempron (M1xx)		
	Ultraportable (~10-20W)	Mobile Sempron (2xxx+-4xxx+) Sempron 2100 fanless		Turion II Neo (K6xx) Athlon II Neo (K1xx) V-series (V1xx)		
Embedded (~10-20W)				Turion II Neo X2 Athlon II Neo X2 Athlon II Neo		

7.6 K10.5 Shanghai-based embedded lines (2)

AMD's K10.5 Shanghai-based embedded lines

These lines are actually the same as the K10.5 Shanghai-based ultraportable mobile lines, as indicated below.

Ultraportable mobiles

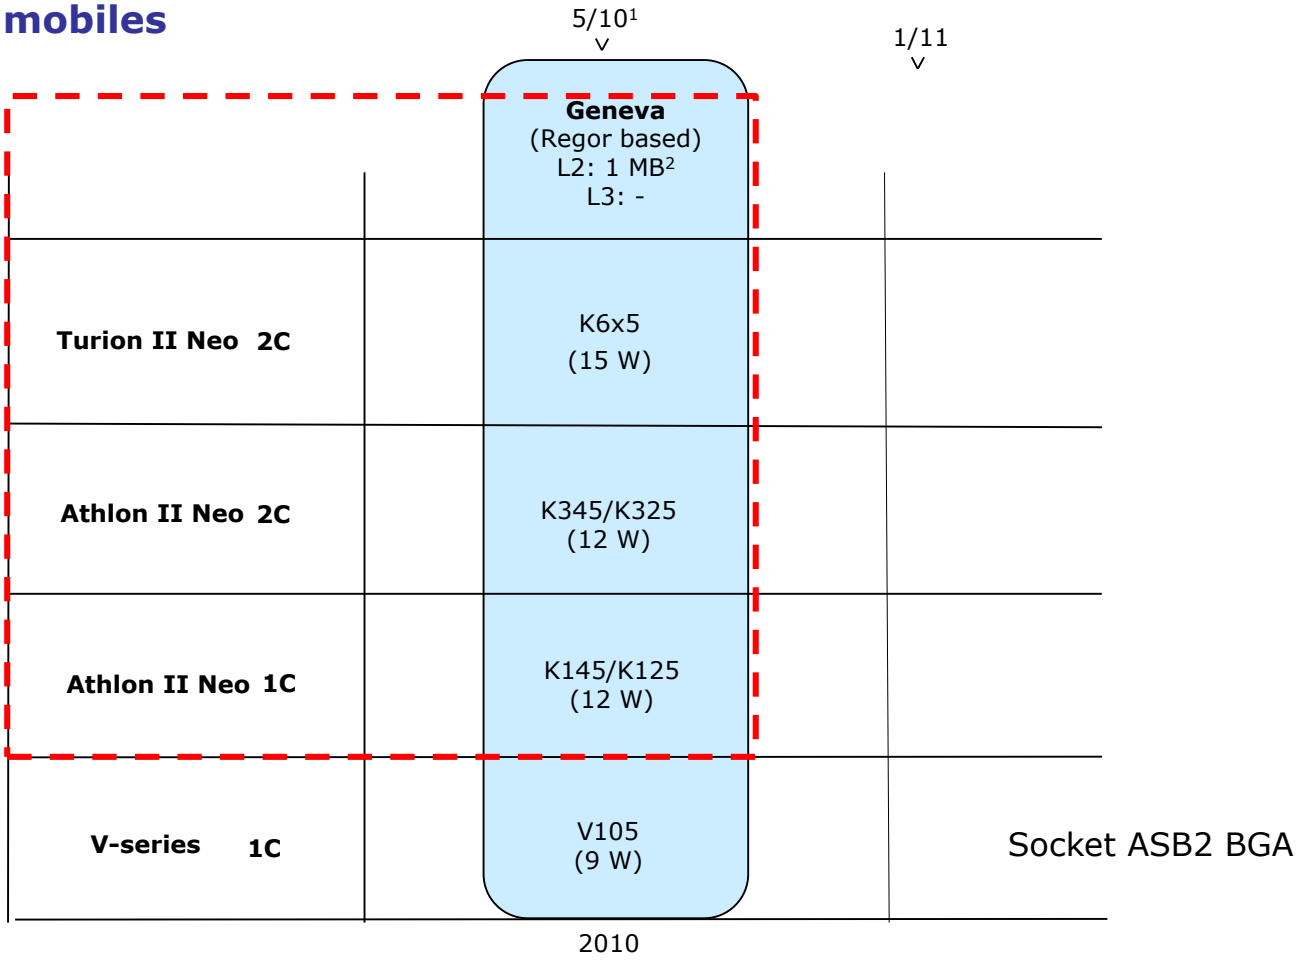

[1] Some of the models indicated were introduced only in 1/2011

[2] In V105: L2=512 KB

7.6 K10.5 Shanghai-based embedded lines (3)

The K10.5 Shanghai-based embedded lines – The related platform [142]

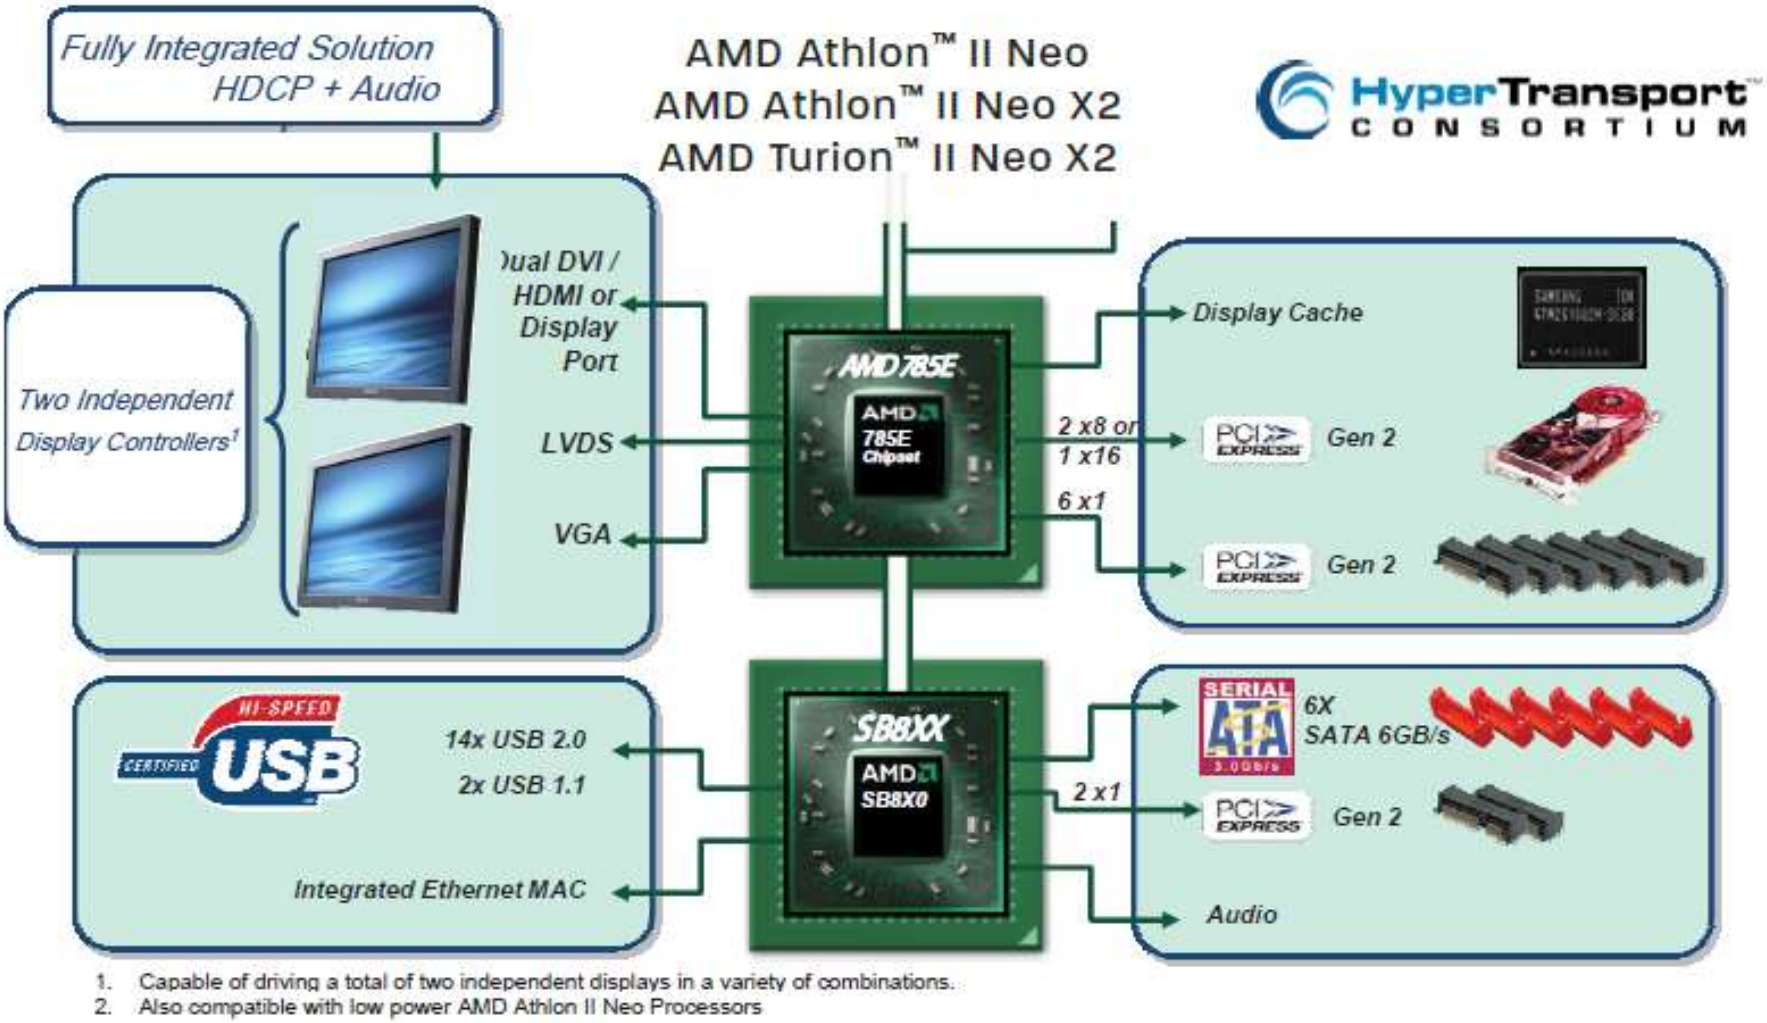

1. Capable of driving a total of two independent displays in a variety of combinations.
2. Also compatible with low power AMD Athlon II Neo Processors

7.6 K10.5 Shanghai-based embedded lines (4)

The ASB2 BGA Socket

It is an update of the ASB1 BGA Socket introduced one year before.

The ASB1 BGA Socket [102]

The ASB1 family of processors are soldered directly to the board and reduce vertical height requirements from up to 8.6mm to 2.03mm. This enables small form factor and rugged designs to be created.

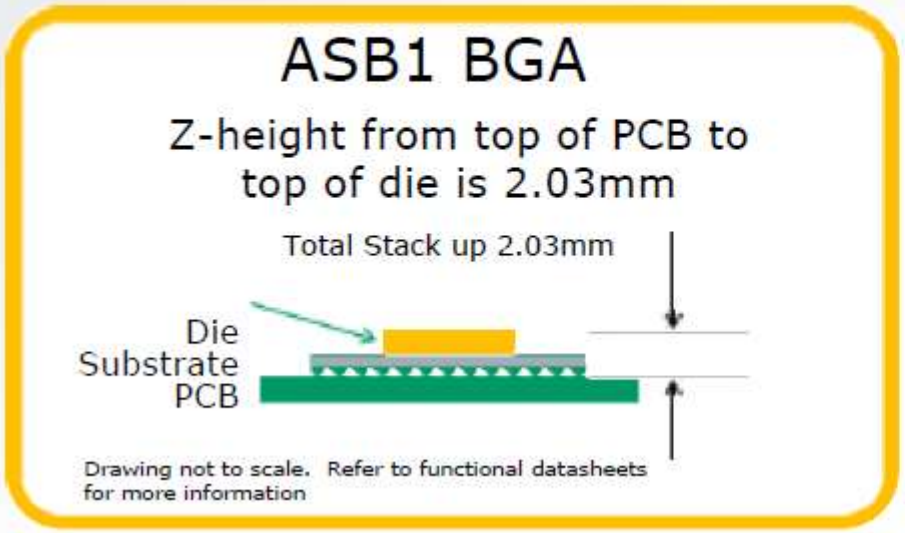

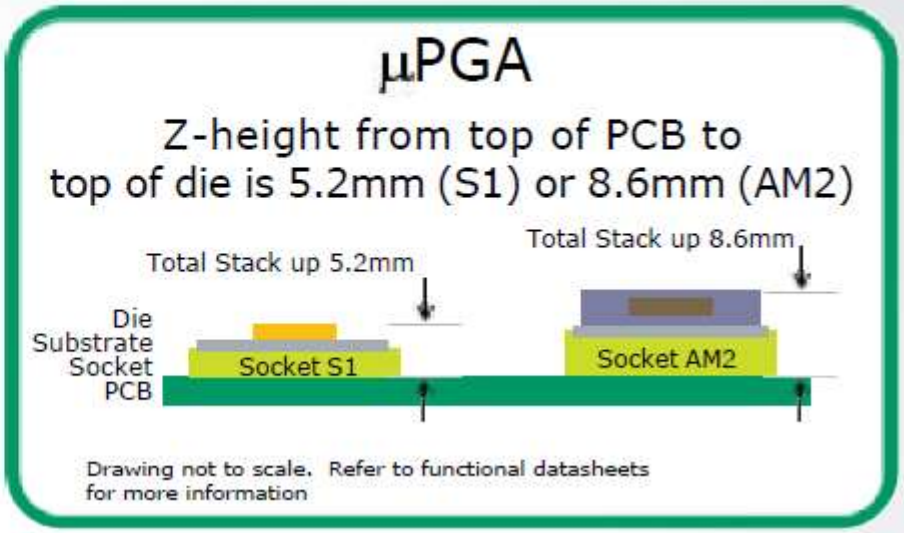

8. The K10.5 Istambul family

- 8.1 Overview of the K10.5 Istambul family
- 8.2 K10.5 Istambul-based server lines
- 8.3 K10.5 Istambul-based desktop lines

8.1 Overview of the K10.5 Istambul family

8.1 Overview of the K10.5 Istambul family

AMD designed first a server die and then an upgraded desktop die for their K10.5 Istambul-based servers and desktops, as indicated below.

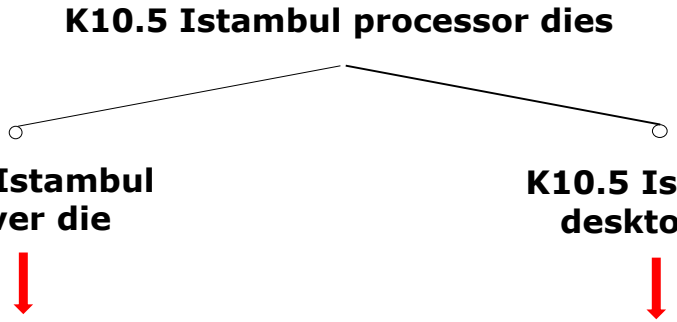

	K10.5 Istambul server lines	K10.5 Istambul desktop lines
Typical line	***The Istambul MP line***	***The Phenom II X6 line***
	45 nm up to 6 cores	45 nm with high-k dielectric [143] up to 6 cores
Introduced:	6/2009	4/2010

8.1 Overview of the K10.5 Istambul family (2)

Overview of subsequent K10/K10.5 DP/MP server implementations [88]

65 nm
45nm

Platform Segment	2008	2009	2010
CPU	"Barcelona" 4-Core • 2M L3 • RDDR-2 • 3x HT-1 • AMD-V™ • 65nm	"Shanghai" 4-Core • 6M L3 • RDDR-2 • cHT-3 • AMD-V • 45nm	"Istanbul" 2H09 6-Core • 6M L3 • RDDR-2 • cHT-3 • AMD-V • 45nm
			"Magny-Cours" 1H10 12-Core • 12M L3 • APML • Probe Filter • 45nm • 4x HT-3 • AMD-V • HTC • U/R DDR-3
			"Sao Paulo" 1H10 6-Core • 6M L3 • APML • Probe Filter • 45nm • 4x HT-3 • AMD-V • HTC • U/R DDR-3
Chipset	Nvidia nForce 3600/3050 Broadcom HT-2100/1000		AMD RD890S w/IOMMU AMD RD870S w/IOMMU AMD SB700S
Platform	Socket F (1207) • 3x HT-1 (moving to cHT-3) • DDR-2 (Dual Channel)		"Maranello" • 4x HT-3 • DDR-3

8.1 Overview of the K10.5 Istambul family (3)

AMD's K10.5 Istambul-based processor lines – Overview [14]

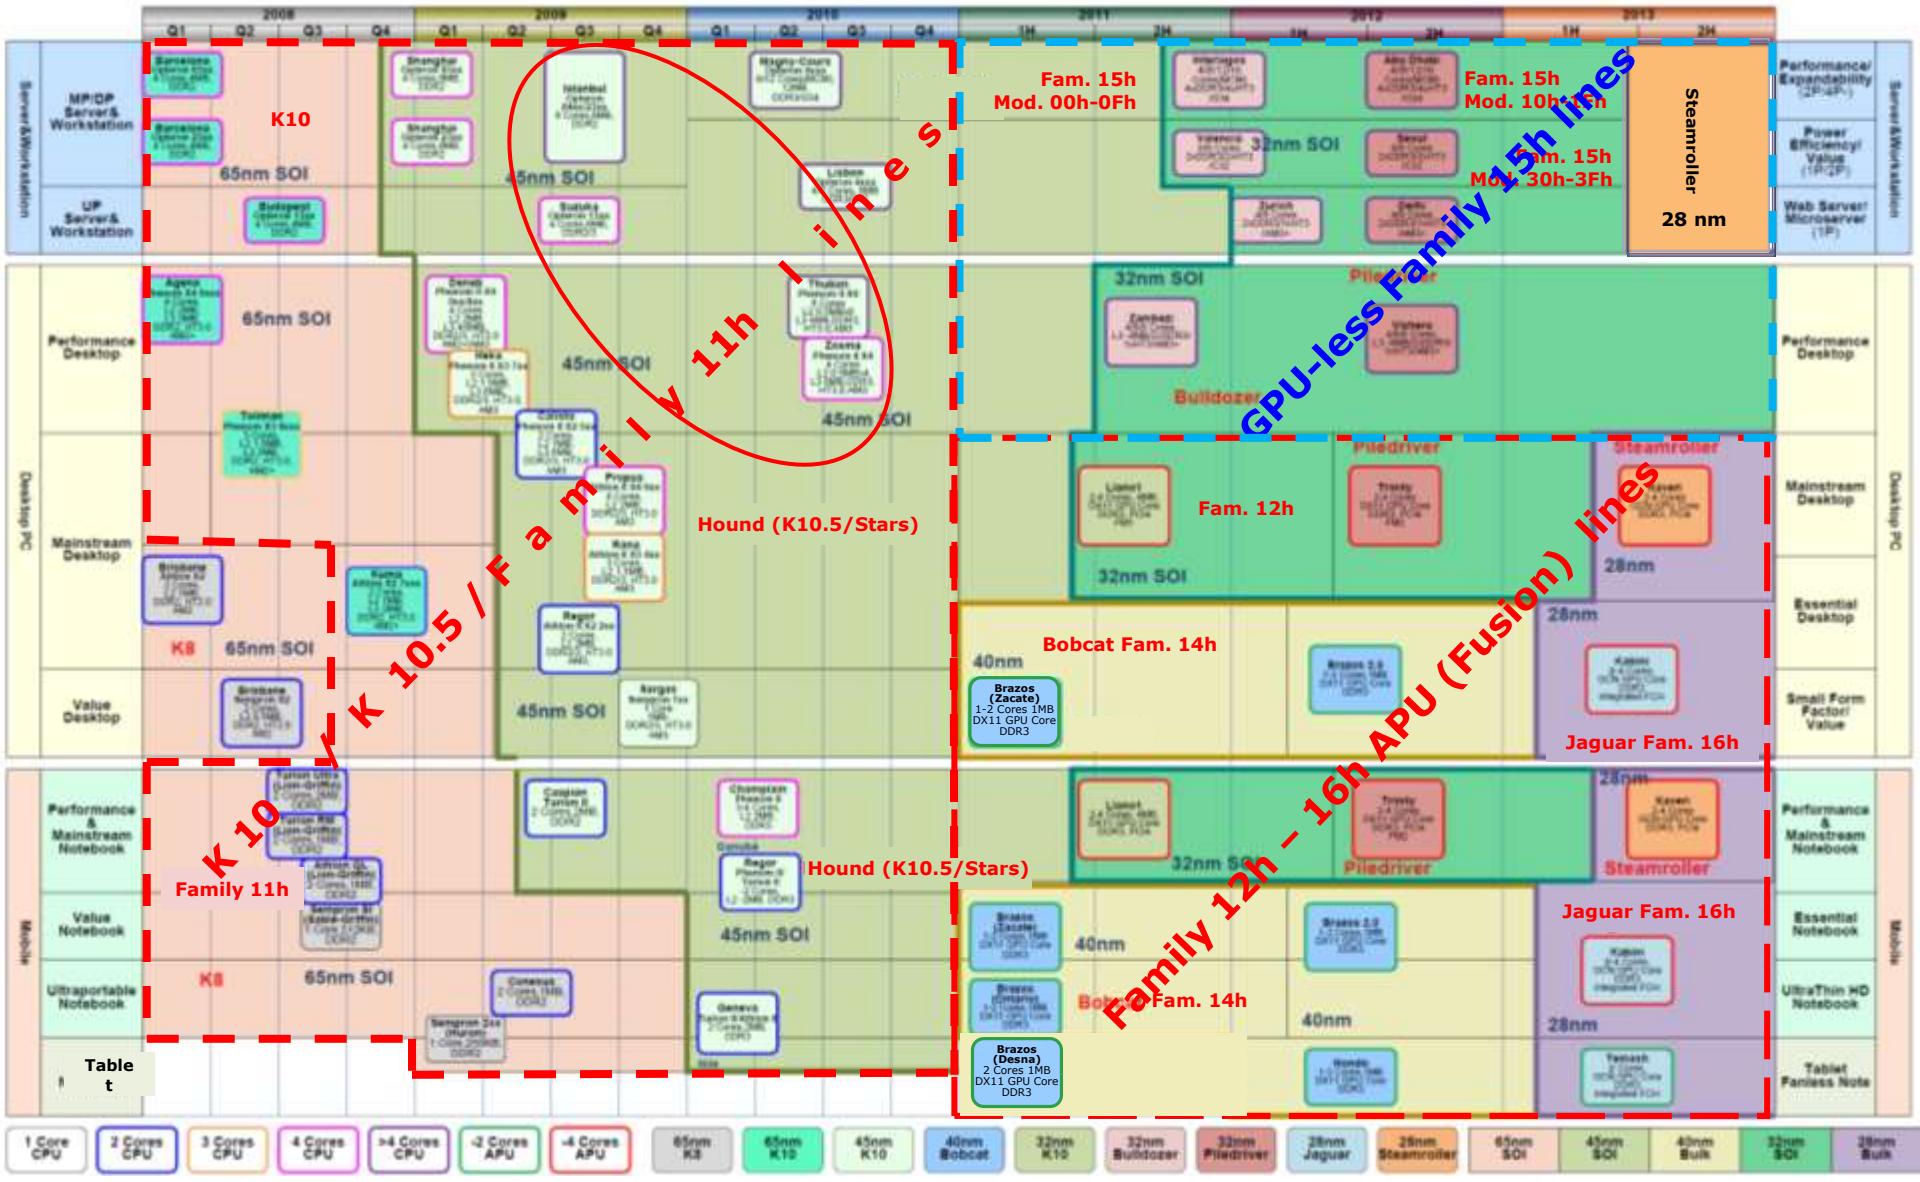

8.1 Overview of the K10.5 Istambul family (4)

Brand names of AMD's K10.5 Istambul-based processor lines

		2003-2007	2007-2008	2008-2011	2009	2009
		K8 (Hammer)	**K10 (Barcelona)**	**K10.5 (Shanghai)**	**K10.5 (Istanbul)**	**K10.5 (Magny- Course)**
Servers	**4P servers**	**See Section 4**	**Barcelona (834x-836x))**	**Shanghai (837x-839x)**	**Istambul (8410-8430)**	**Magny-Course (6100)**
	2P servers		**Barcelona (234x-236x)**	**Shanghai (237x-239x)**	**Istambul (241x-243x)**	**Lisbon (4100)**
	1P servers		**Budapest (135x-136x)**	**Suzuka (138x-139x)**		
Desktops	**High perf.** (~80-120W)		Phenom X4-X2	Phenom II X4-X2	Phenom II X6-X4	
	Mainstream (~60-90W)	Athlon 64 Athlon 64 X2	Athlon X2	Athlon II X4-X2		
	Value (~40-60W)	Sempron		Sempron		
Mobiles	**High perf.** (~30-40W)	Turion 64 X2 (TL 6/5) Turion 64 (ML/MT)		Phenom II (N/P 9xx-6xx) Turion II Ultra (M6xx) Turion II (M/N/P 5xx)		
	Mainstream (~20-30W)	Athlon 64 X2 (TK-5x/4x) Athlon 64 (2xxx+-4xxx+)		Athlon II (M/N/P 3xx) Sempron (M1xx)		
	Ultraportable (~10-20W)	Mobile Sempron (2xxx+-4xxx+) Sempron 2100 fanless		Turion II Neo (K6xx) Athlon II Neo (K1xx) V-series (V1xx)		
Embedded (~10-20W)				Turion II Neo X2 Athlon II Neo X2 Athlon II Neo		

8.2 K10.5 Istambul-based server lines

8.2 K10.5 Istambul-based server lines

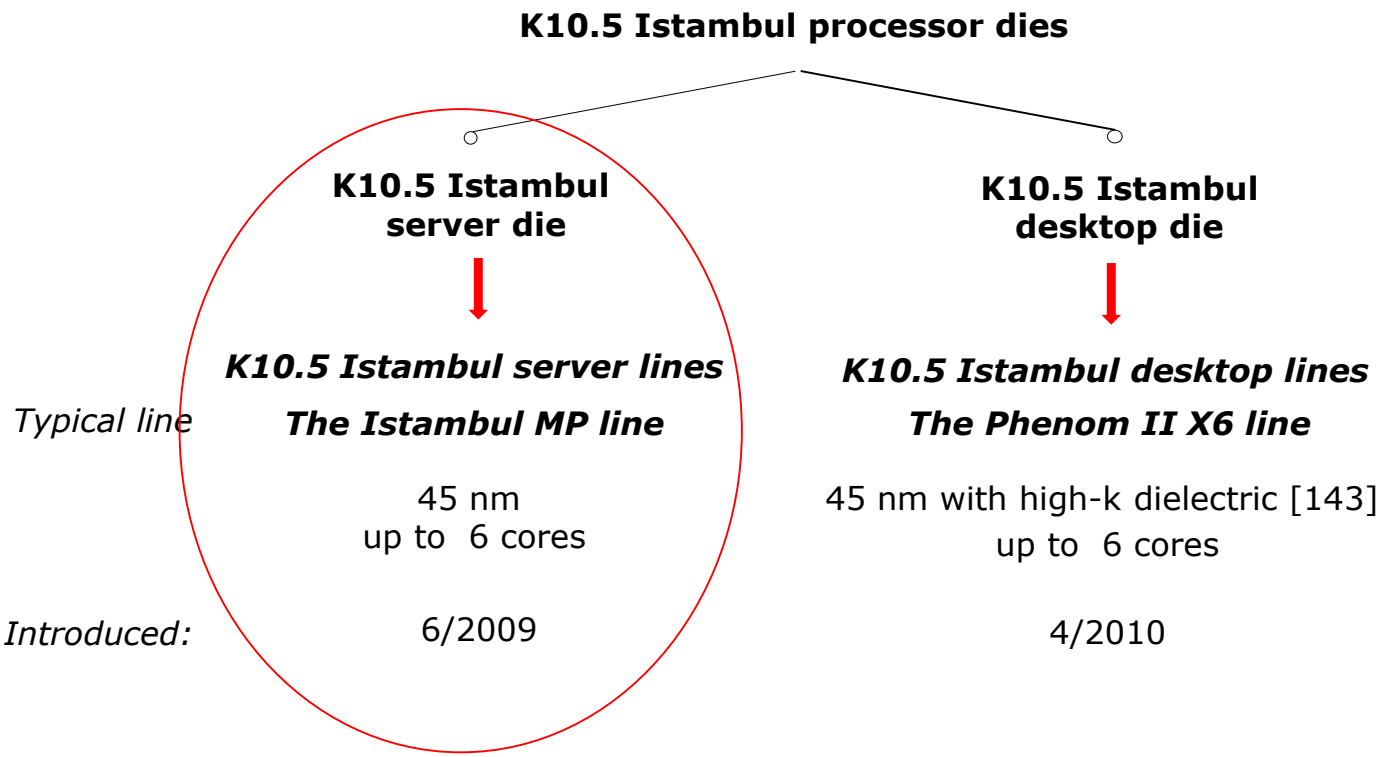

8.2 K10.5 Istambul-based server lines (2)

K10.5 Istambul-based server lines

Announced 4/2009, first delivered 6/2009

They are socket compatible with previous server generations.

8.2 K10.5 Istambul-based server lines (3)

Overview of subsequent K10/K10.5 DP/MP server implementations [88]

65 nm
45nm

Platform Segment	2008	2009	2010
CPU	**"Barcelona"** 4-Core • 2M L3 • RDDR-2 • 3x HT-1 • AMD-V™ • 65nm	**"Shanghai"** 4-Core • 6M L3 • RDDR-2 • cHT-3 • AMD-V • 45nm	**"Istanbul"** 6-Core 2H09 • 6M L3 • RDDR-2 • cHT-3 • AMD-V • 45nm
			"Magny-Cours" 1H10 12-Core • 12M L3 • Probe Filter • 4x HT-3 • HTC • U/R DDR-3 • APML • 45nm • AMD-V
			"Sao Paulo" 1H10 6-Core • 6M L3 • Probe Filter • 4x HT-3 • HTC • U/R DDR-3 • APML • 45nm • AMD-V
Chipset	Nvidia nForce 3600/3050 Broadcom HT-2100/1000		AMD RD890S w/IOMMU AMD RD870S w/IOMMU AMD SB700S
Platform	**Socket F (1207)** • 3x HT-1 (moving to cHT-3) • DDR-2 (Dual Channel)		**"Maranello"** • 4x HT-3 • DDR-3

8.2 K10.5 Istambul-based server lines (4)

Basic structure of the Istambul MP server [79]

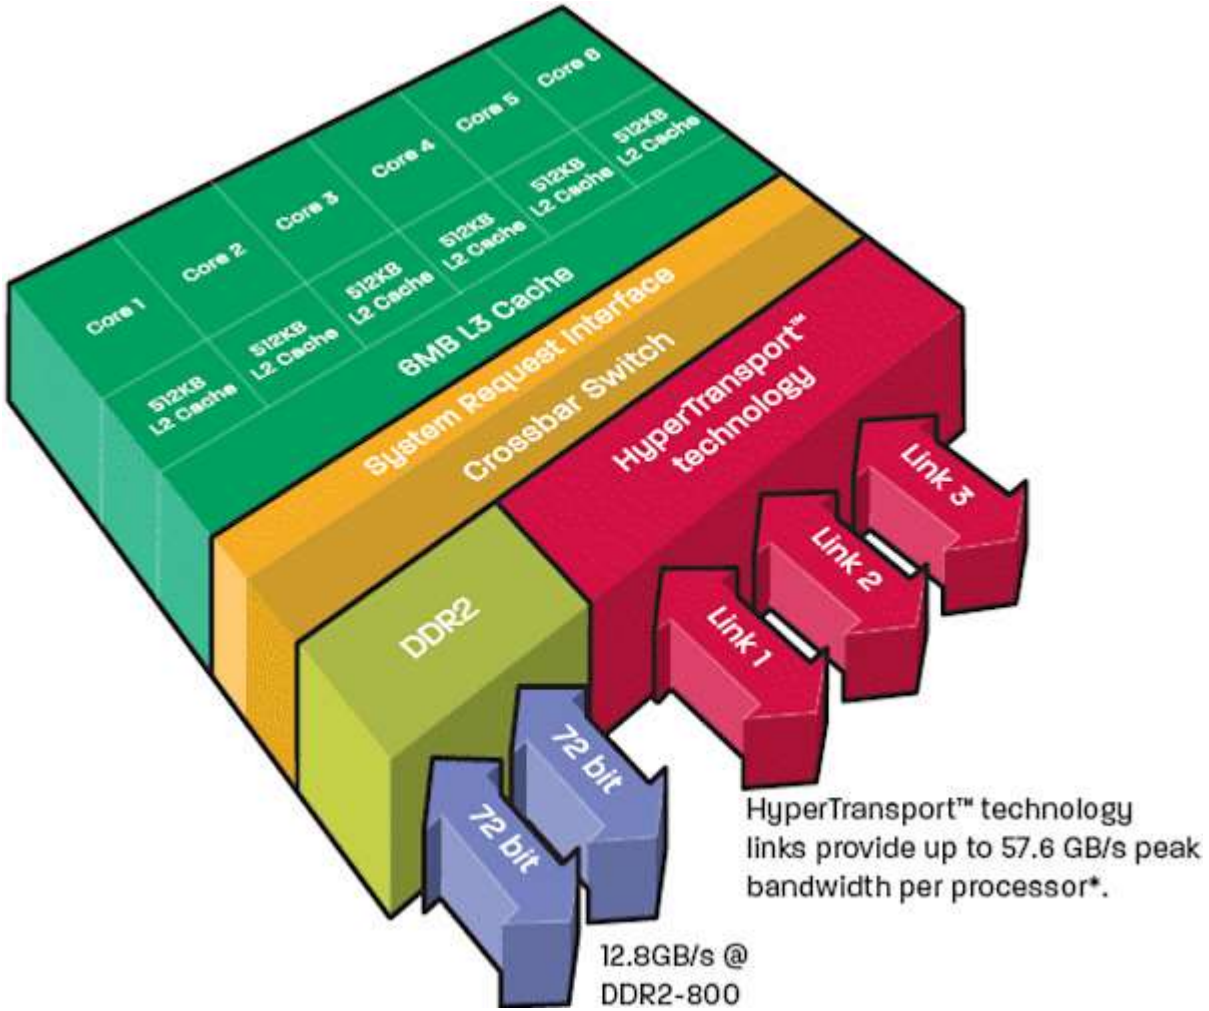

8.2 K10.5 Istambul-based server lines (5)

Istambul's per core microarchitecture [17]

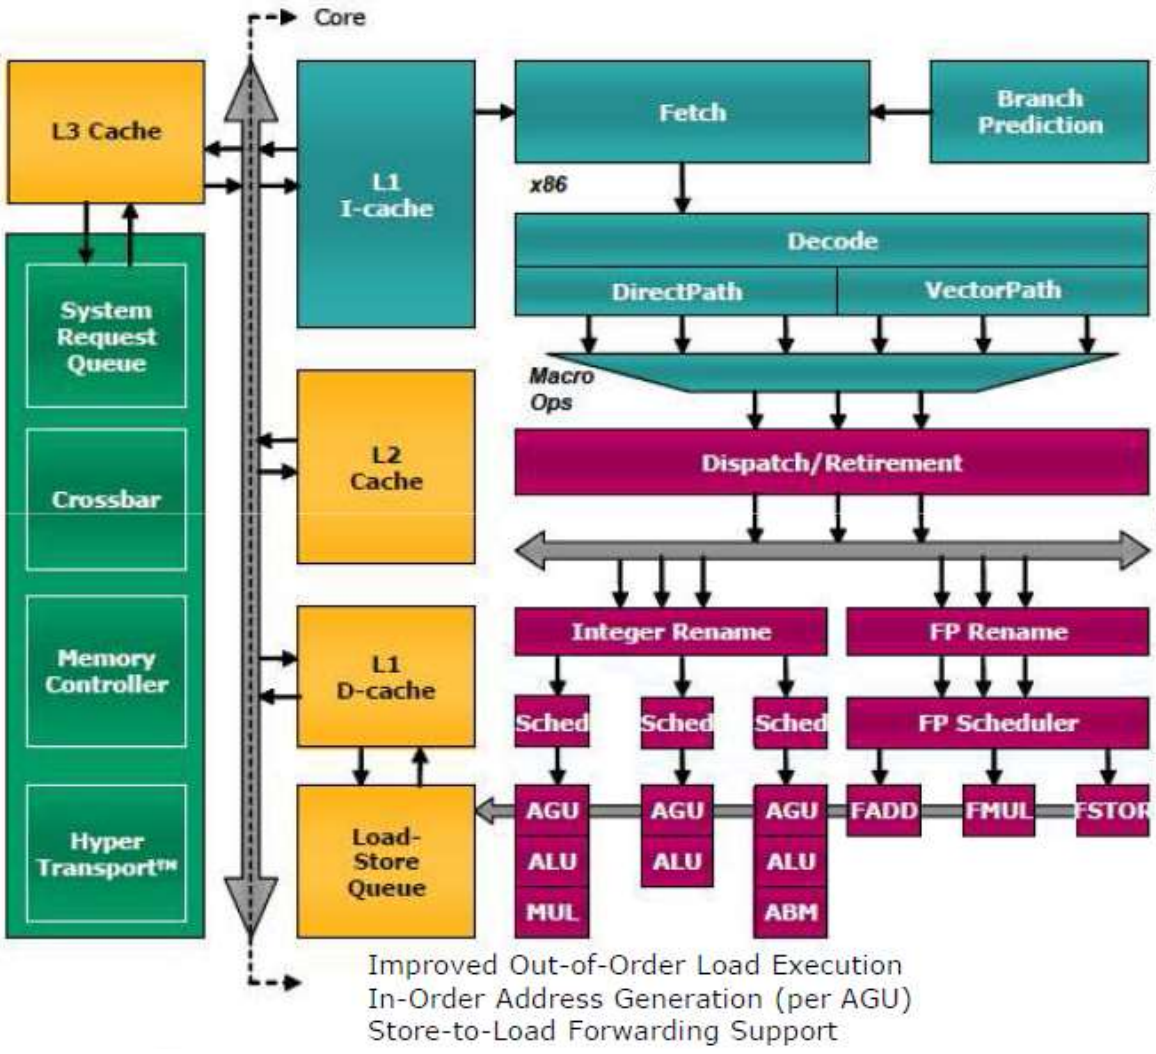

Notes / Considerations

Avoid having more than 2 or 3 branches per 16B of instructions.

Three Decode Categories (FastPath also called DirectPath)
- DirectPath Single - *best*
- DirectPath Double - *better*
- VectorPath (microcode) - *good*

Macro-Ops tracked ReOrder Buffer (ROB).
- 72 entries (3 wide x 24 deep)
- In-order dispatch, retirement

Micro-Ops issue from Sched to Execution Units
- "Sched" aka "Reserv. Station"
- Out-of-order issue
- FP scheduler shared across units
- INT Schedulers are "per unit"

8.2 K10.5 Istambul-based server lines (6)

Main features of AMD's K10.5 Istambul-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)	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 (NPT)	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. 17h Mod. 00h-0Fh		6/2017	Epyc (2S!!) (4 dies/proc.)	7000	14 nm	4x(2x4) (32C)	½ MB/C	2 MB/C	DDR4-2666	IFIS 75.8 GB/s	SP3

8.2 K10.5 Istambul-based server lines (7)

Main enhancements of the K10.5 Istambul-based server family [80]

Feature	Description	Benefit
Six Cores per Socket	*Six Core support for F (1207) Socket infrastructure*	*Improves Performance*
HT Assist	*Reduces probe traffic and resolves probes more quickly in multi-socket systems*	*Increases HT bus efficiency*
Higher HyperTransport™ 3.0 Technology Speeds	*Support for up to 4.8GT/s per link*	*Overall System Performance*
APML Remote Power Management Interface (RPMI)	*Remote monitor and control of P-state limits*	*Processor Power Savings*
x8 ECC	Correction for x4 and x8 device failures	*Superior Reliability*
Continued Drop-in Upgradeability for F (1207) Platforms	*Six Cores within same power bands*	*Investment Protection and Time to Market*

8.2 K10.5 Istambul-based server lines (8)

a) HT Assist (HyperTransport Assist)

It is a probe or snoop filter that reduces so called coherency traffic between the cores needed to maintain cache coherency.

It supports only quad-socket systems (MP systems)

The cache coherency problem

- Lets consider a 4 socket multiprocessor system with each processor having caches [81].

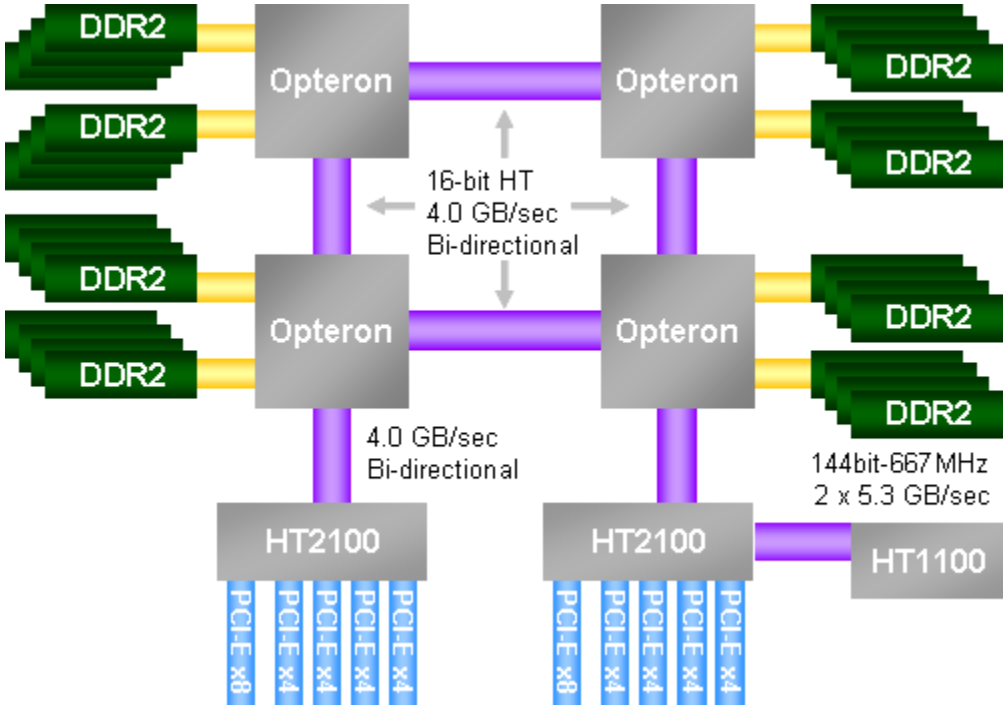

8.2 K10.5 Istambul-based server lines (9)

- For simplicity lets take for granted that each processor has an inclusive last level cache,
  i.e. the last level cache keeps all data incorporated in lower level (e.g. L1, L2) caches.
- Obviously, each processor operates independently on its own.

  Then at a given time different copies of the same data (more precisely of data belonging to the same address) may exist in different cache states in the individual last level caches of the processors.
- Based on the state information kept in each last level cache a cache coherency protocol is used to assure that processors access always the most recent copy of the referenced data.
- There are different schemes and different cache coherency protocols to maintain cache coherency (not discussed here).
- As an example, a possible scheme to maintain cache coherency is shown below for an MP system (Opteron's Istambul system without using HT Assist) [81].

8.2 K10.5 Istambul-based server lines (10)

Example for accessing cached data without HT assist [81]

- Let's consider an MP system, as indicated below.

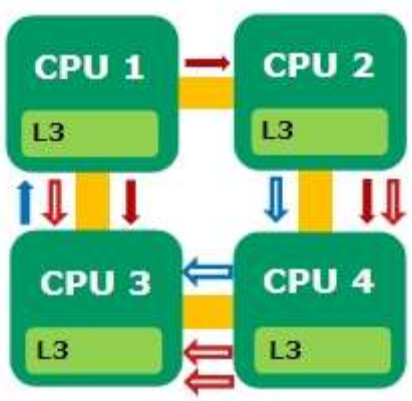

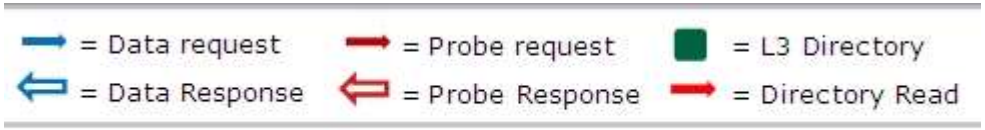

- Lets take for granted that CPU 3 requests data from the data space maintained by CPU 1 by sending a Data request to it.
- CPU 1 snoops then all processors for the most recent value of the referenced data by sending Probe requests to all other processors.
- CPU 3 idles waiting for the requested data.
- Assuming that CPU 2 has the most recent value (revealed by the cache line state) CPU 2 sends the requested data to CPU 3 (through CPU 4).

In this case 9 transactions are needed in total to get the referenced data.

8.2 K10.5 Istambul-based server lines (11)

HT Assist

It maintains a directory of Probe Filter entries in a portion of each processor's L3 cache. (Its size is configurable, typical size is 1 MB [16]).

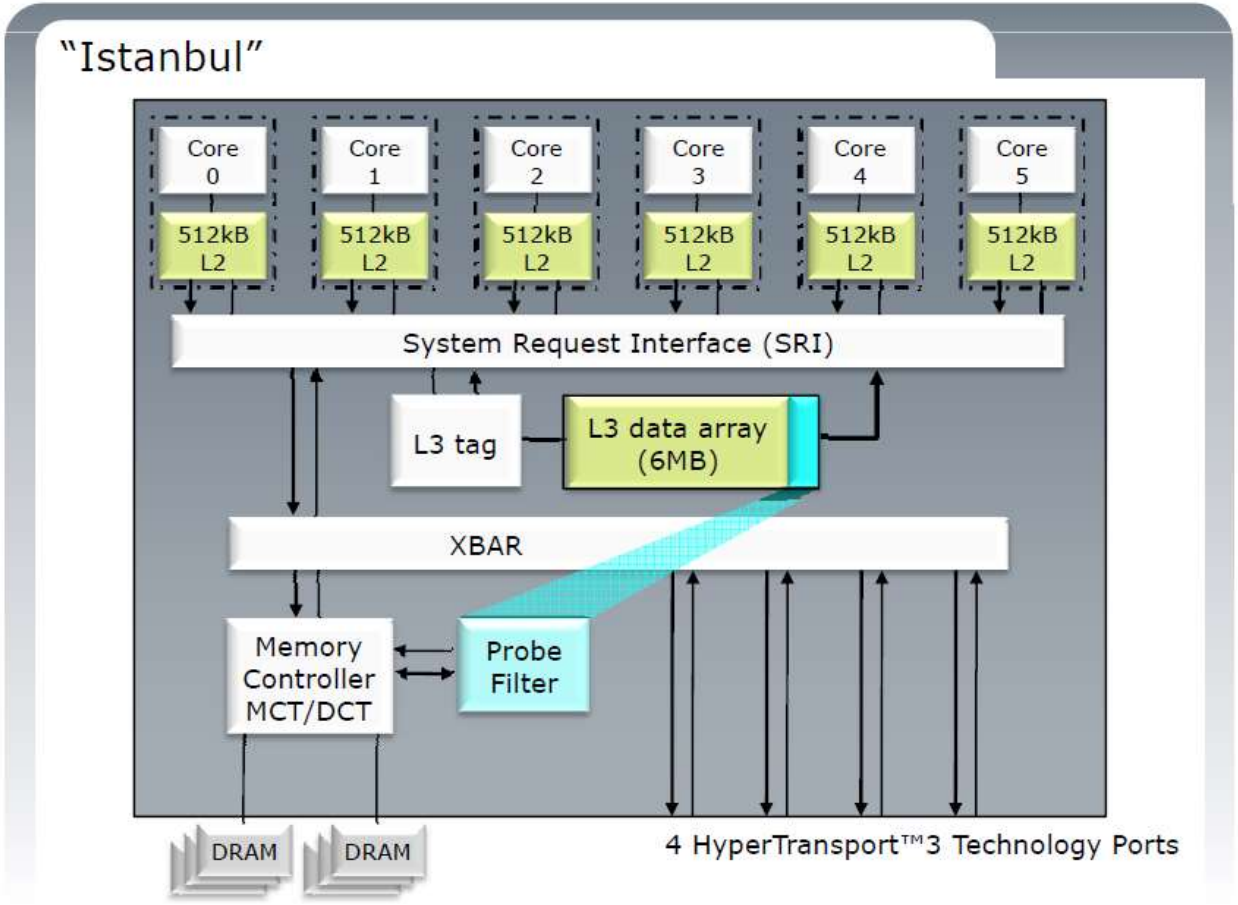

Figure: Probe filter of the Istambul processor [16]

8.2 K10.5 Istambul-based server lines (12)

The Probe Filter entries hold data about the state and the owner of the most recent copy of the associated cache lines.

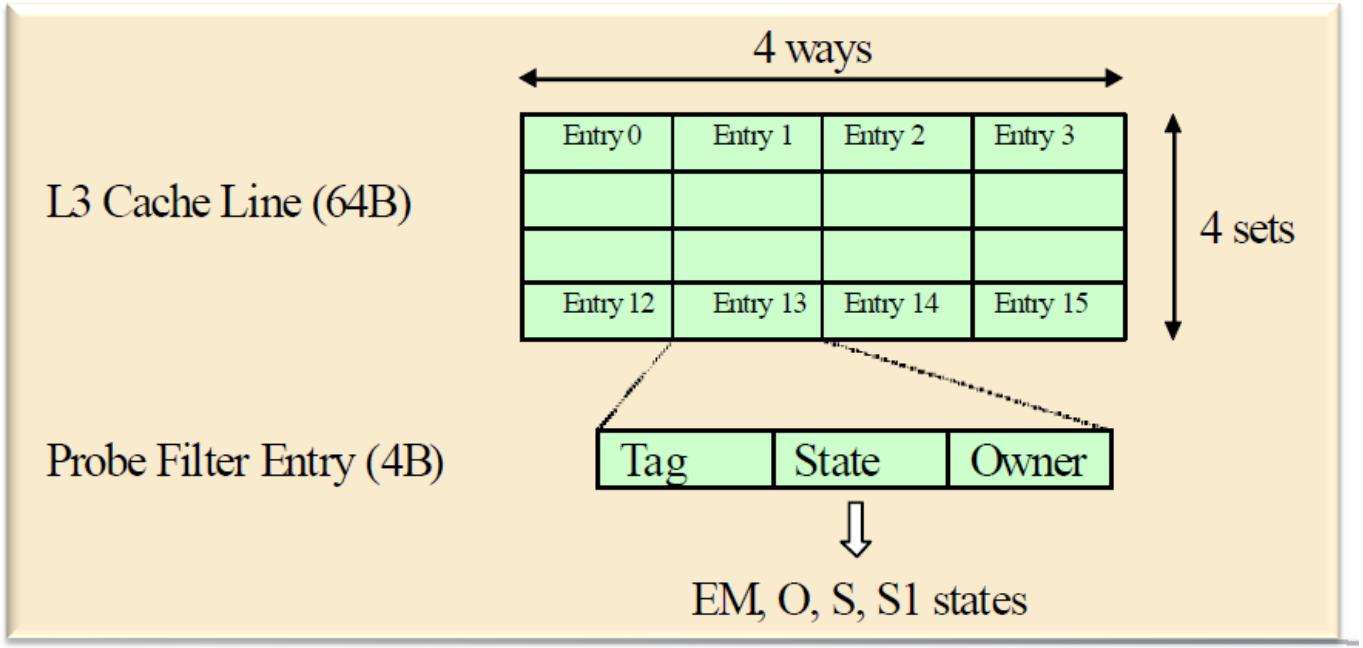

Figure: Format of a Probe Filter Entry [16]

(For a straightforward discussion we do not want to go into details of the possible cache line states and the related cache coherency protocol).

If there exists no probe filter entry to the referenced data, referenced data is not cached.

8.2 K10.5 Istambul-based server lines (13)

Example for accessing cached data with HT Assist [81]

Case 1: The most recent copy of the referenced cache line is held in CPU 2.

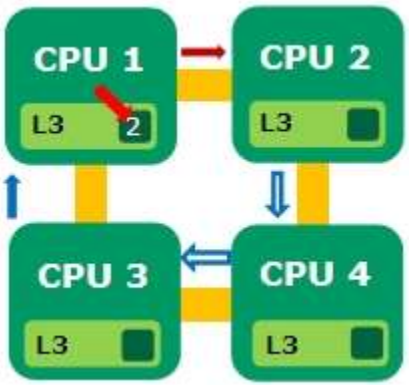

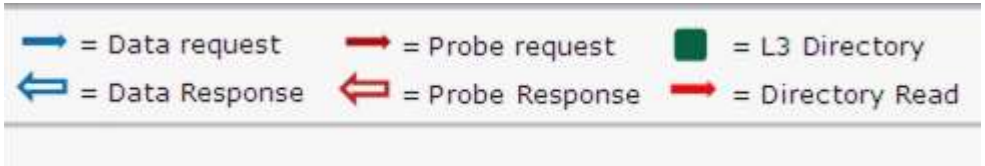

- Lets assume again that CPU 3 requests data from the data space maintained by CPU 1 by sending a Data request to it.
- Then CPU 1 checks its Probe Filter to locate the most recent value of the requested data.
- CPU 1 finds that CPU 2 has the most recent data value and sends a Probe request to CPU 2.
- CPU 2 sends the requested data to CPU 3 (through CPU 4).

In this case 4 transactions are needed in total to get the referenced data.

8.2 K10.5 Istambul-based server lines (14)

Case 2: The most recent coy of the referenced cache line is held in CPU 1 [81]

In this case accessing referenced data will be further simplified, as shown below.

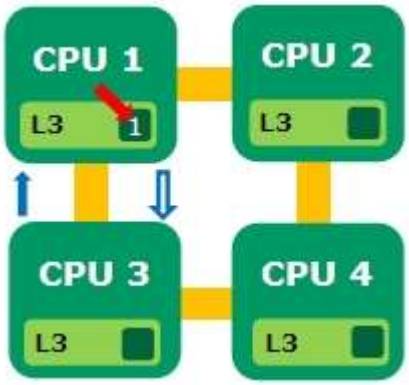

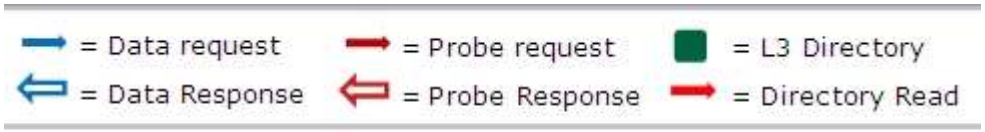

- Lets assume again that CPU 3 requests data from CPU 1 by sending a Probe request to it.
- CPU 1 checks its Probe Filter directory to locate the most recent value of the requested data.
- CPU 1 finds now that it has the most recent value of the requested data and sends it to the requester (CPU 3) directly.

In that case there are only two transactions needed.

8.2 K10.5 Istambul-based server lines (15)

Remarks

1) As AMD make use of exclusive L2 and "most exclusive" L3 caches the last level cache (L3 cache) does not incorporate all data included in lower level caches (i.e. in the per core available L1, L2 caches).

   So AMD's HT assist implementation had to cover this situation.

   Nevertheless, no details were found about AMD's solution.

2) When considering the hardware support of the cache coherency problem by probe filters it can presumably stated that efficiency benefits of exclusive caches vanish and inclusive caches (used by Intel) become more beneficial.

3) Intel introduced their first snoop filter (probe filter) in their B5000 platform, including the Blackford north bridge, supporting the Intel 5100 dual-core and 5300 quad-core processors in 2006.

Benefits of using HT Assist

HT Assist reduces the number of probe requests that are sent out by the referenced processors.

⟹ This leaves additional bandwidth for other requests and reduces the average memory latency, as indicated in the next Figure.

8.2 K10.5 Istambul-based server lines (17)

Memory bandwidth and latency improvements due to HT Assist [82]

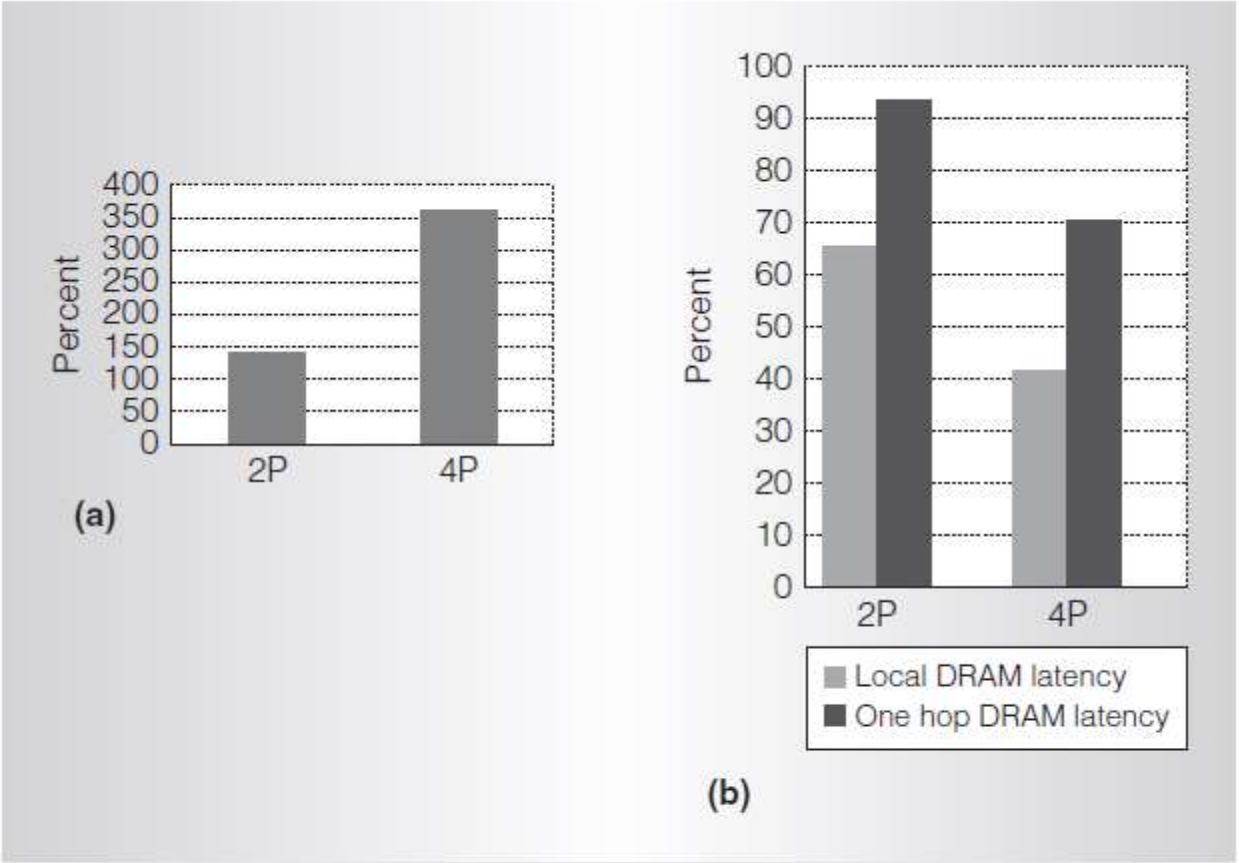

Figure 8. Memory bandwidth improvement (a) and relative memory latency (b) with and without HT Assist. Latency is normalized to HT Assist disabled as 100 percent (lower is better).

8.2 K10.5 Istambul-based server lines (18)

Performance improvements due to HT Assist for different commercial workloads (simulation results) [82]

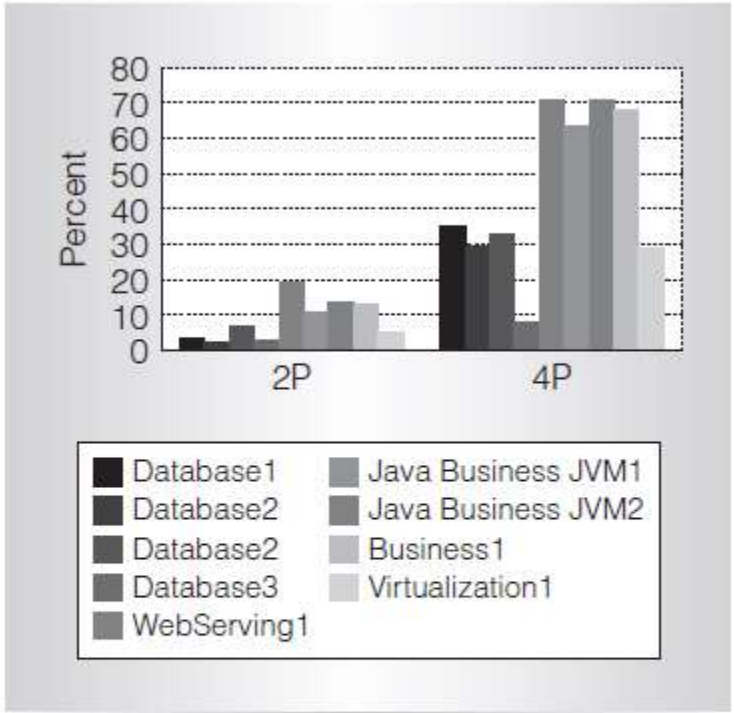

Figure 10. Simulation-based studies of performance improvement with HT Assist across a diverse set of commercial workloads.

8.2 K10.5 Istambul-based server lines (19)

b) Enhanced HyperTransport 3.0 speed

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 (NPT)	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. 17h Mod. 00h-0Fh		6/2017	Epyc (2S!!) (4 dies/proc.)	7000	14 nm	4x(2x4) (32C)	½ MB/C	2 MB/C	DDR4-2666	IFIS 75.8 GB/s	SP3

8.2 K10.5 Istambul-based server lines (20)

c) Advanced Platform Management Link (APML) (based on [53])

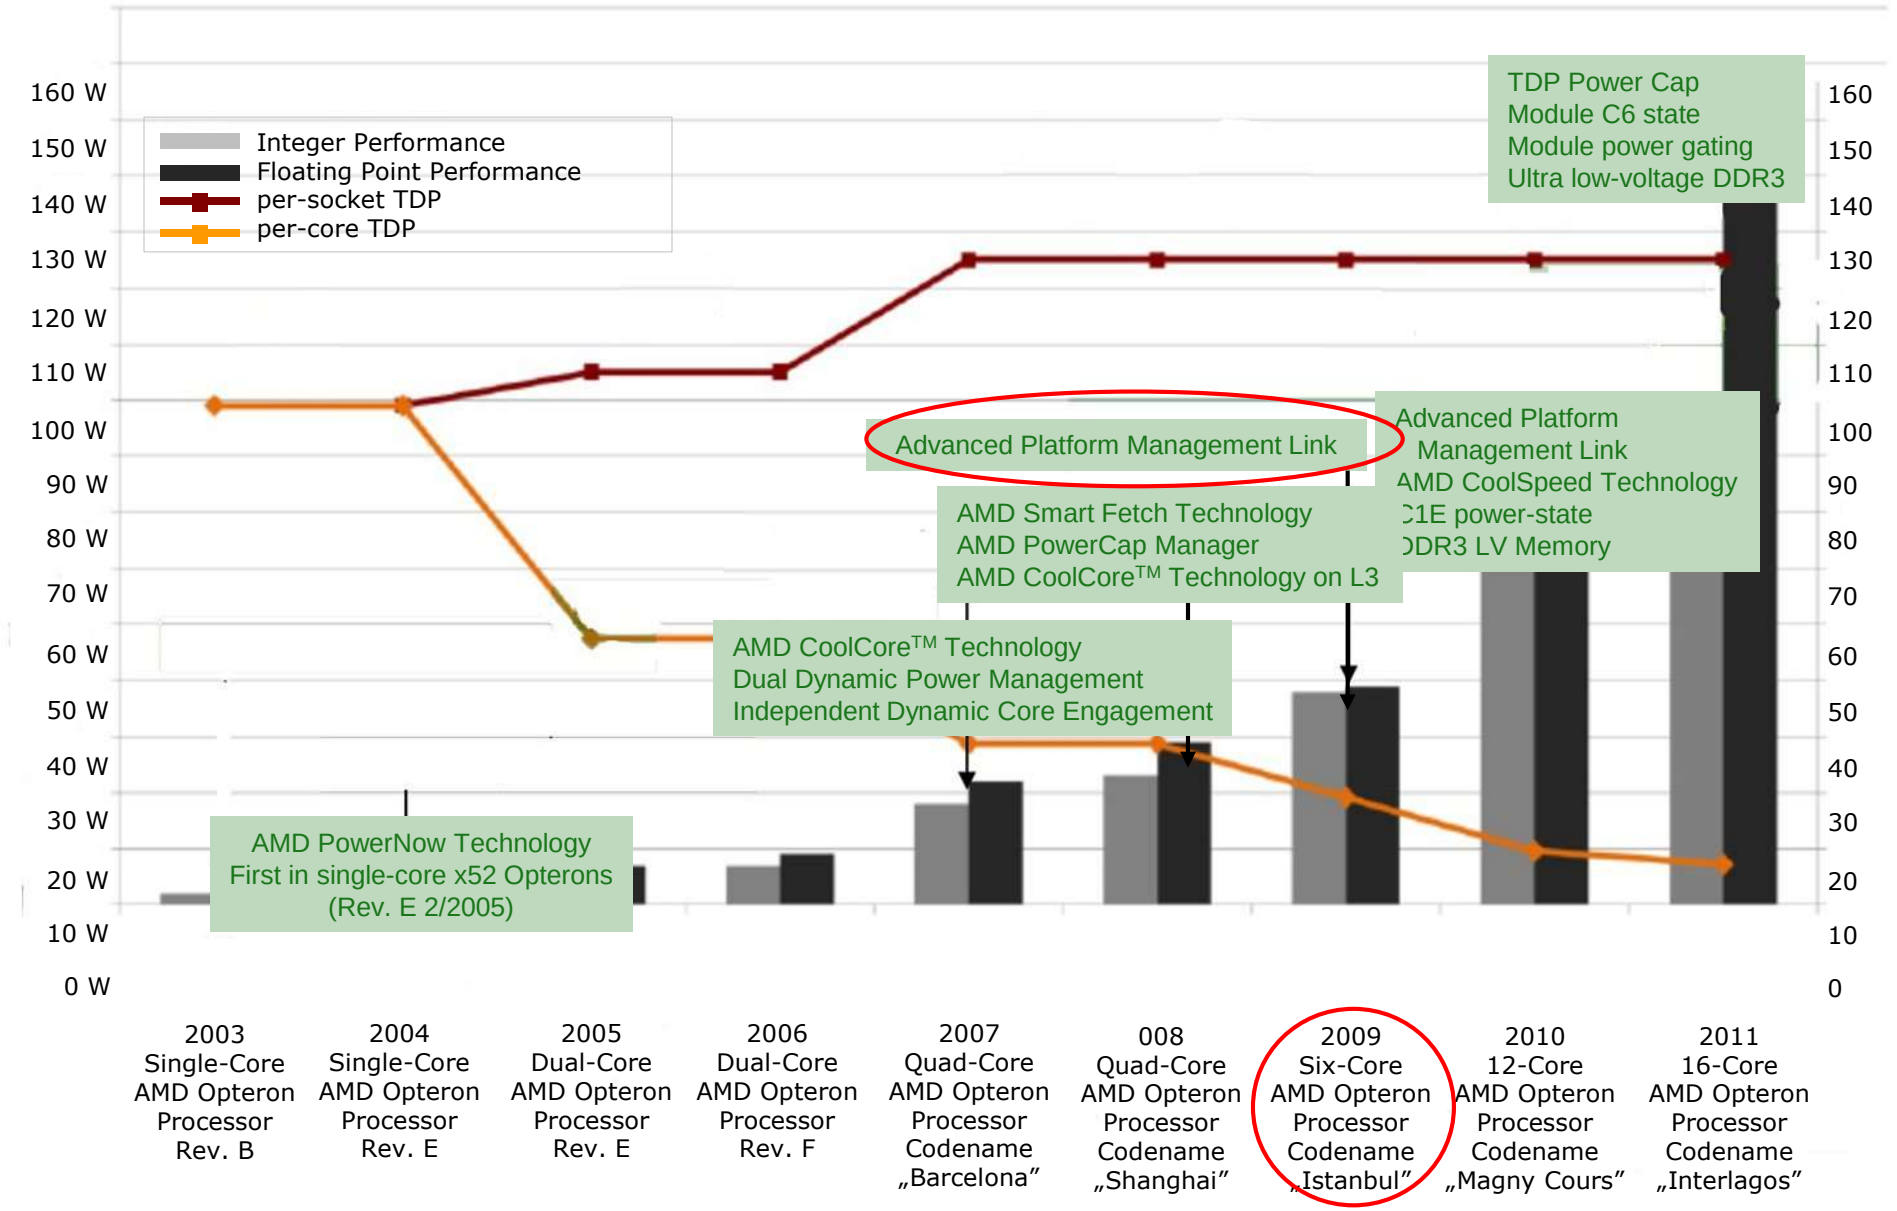

8.2 K10.5 Istambul-based server lines (21)

APML (Advanced Platform Management Link)-1 [83], [84], [85]

(also referred as the Sideband Interface (SBI))

- APML is physically a 2-wire bus that follows the SMBus v2.0 specification with a few exceptions.
- Its use requires a bus master, called the Management Controller or Service Processor.
- APML allows system administrators to remotely monitor and control particular system settings actually by reading and writing limited processor state through predefined interfaces via the Service Processor.

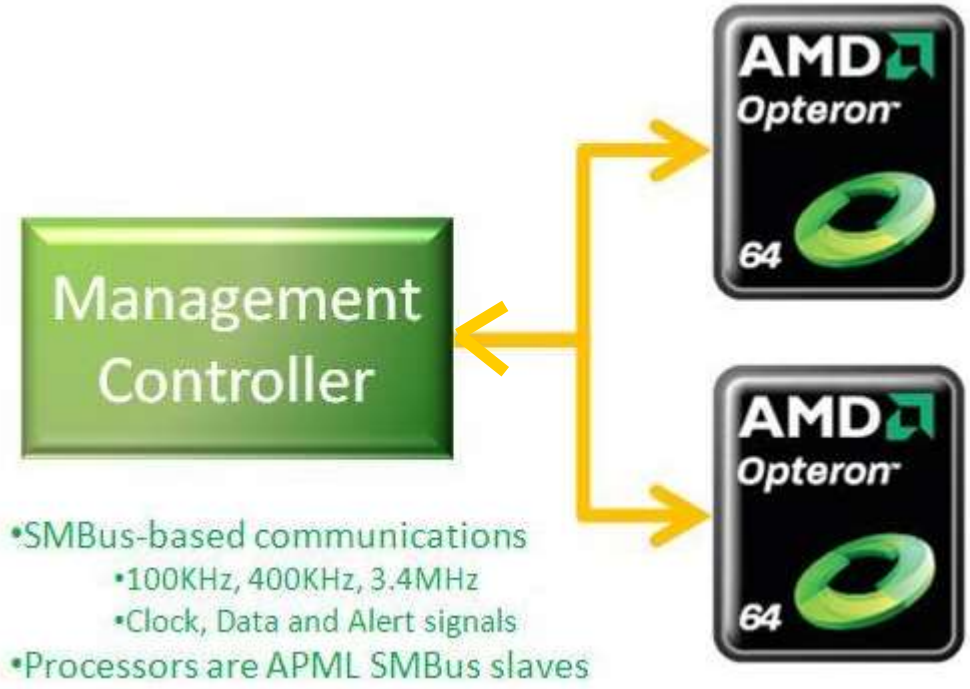

8.2 K10.5 Istambul-based server lines (22)

APML (Advanced Platform Management Link)-2 [83], [84], [85]

Aim of the APML link

In the Istambul line of servers APML can be used

- to remotely monitor and cap platform power consumption by imposing P-state limits directly for a particular processor via the SBI Remote Management Interface (SB-RMI), and
- to remotely monitor the internal temperature sensors and to specify temperature thresholds for thermal protection through the Temperature Sensor Interface (SB-TSI).

via the Service Processor.

Remark

Capping power is useful in datacenters to maintain limited power and cooling capabilities.

8.2 K10.5 Istambul-based server lines (23)

APML (Advanced Platform Management Link)-3 [83], [84], [85]

The Remote Management and the Temperature Sensor interfaces

The Remote Management Interface (SB-RMI) and the Temperature Sensor Interface (SB-TSI) define communication protocols to read and write particular internal processor registers, such as the P-state Status Register, the P-State Current Limit Register or the Hardware Thermal Control Register through the APML interface via the service processor.

8.2 K10.5 Istambul-based server lines (24)

Remarks

Basic layout of the communication protocol over the APML link [83]

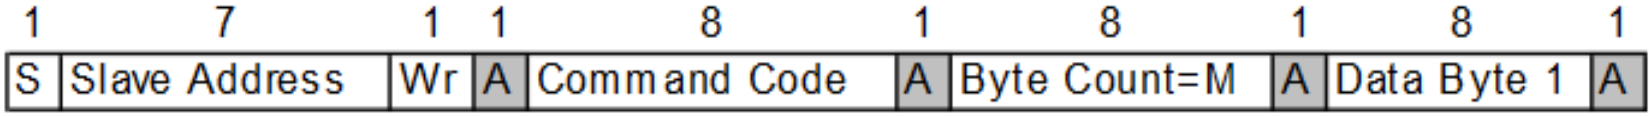

Master-to-Slave

Slave-to-Master

S Start Condition

A Acknowledge (this bit position may be '0' for an ACK or '1' for a NACK)

8.2 K10.5 Istambul-based server lines (25)

Example for the commands (Functions) related to the SBI-RMI interface [83]

Function	Description	Core Specific[1]
CPUID	Access to CPUID using read CPUID command. General purpose registers are not altered unlike a processor CPUID instruction. See the *BIOS Kernel and Developer's Guide* of the processor family for more information about CPUID functions.	Y
HTC	Register read or write command to register address C001_003Eh to access the Hardware Thermal Control (HTC) Register (F3x64).	N
Current P-state	Register read command to register address C001_0063h to access the P-State Status Register (MSRC001_0063).	Y
Set P-state limit	Register read or write command to register address C001_0072h to access the SBI P-state Limit Register (F3xC4).	N
Current P-state limit	Read command to register address C001_0061h to access the P-State Current Limit Register (MSRC001_0061).	N
MCA Registers	Register read or write command using the MSR address as the register address to access MSR0000_0179, MSR0000_017A, MSR0000_017B, MSR0000_0400 through MSR0000_0417, and MSRC000_04[0A:08].	Y
1. Functions that are not core specific must use SB-RMI 41[CoreNum] as the core in the command.		

8.2 K10.5 Istambul-based server lines (26)

APML (Advanced Platform Management Link)-4 [83], [84], [85]

- APML supports 100 KHz, 400 KHz and 3.4 MHz clock speeds.
  - In mobile and desktop environment ACPI supports 400 kHz operation, whereas
  - in server environments 3.4 MHz operation.
- In addition to an APML compatible service processor using APML requires also OS support.

8.2 K10.5 Istambul-based server lines (27)

Remarks

1) Principle of informing the OS about setting a new P-state (simplified) [86]

- When the Service processor changes the P-state limit by issuing an APML command also the OS needs to be informed about the new limit.

  This can be done when the platform invokes a System Control Interrupt (SCI).
- The SCI handler calls then the ACPI machine language routine (AML) that interrogates the Service Processor for the new P-state limit value.
- The ACPI machine language routine (AML) updates the related ACPI _PPC object and generates a message to the OS, indicating that the _PPC object should be re-evaluated.

  (The _PPC (Performance Present Capabilities) object is actually an ACPI method that indicates to the Power Management routine of the OS the highest P-state that the OS can use at a given time) [87].

2) At the introduction of the Istambul server line (6/2009) both APML hardware and software were in development and became available a few month later in 8/2009 [85].

8.2 K10.5 Istambul-based server lines (28)

3) In fact, APML became introduced already with the Barcelona line of servers in 2007, however only via the Temperature Sensor Interface (SB-TSI) in order to access the internal temperature sensors and to specify temperature thresholds via an external service processor.

The extended use of APML for remotely monitoring and capping platform power consumption by imposing P-state limits over the Remote Management Interface (SB-RMI) was introduced only first in the Istambul line of server processors, a few month later than AMD launched this server line (6/2009).

4) APML is an alternative to IPMI (Intelligent Platform Management Interface, introduced in 1998, promoted by Intel, Dell, hp and NEC, and adopted by a large number of companies [88].

8.2 K10.5 Istambul-based server lines (29)

Die shot of the Istambul server chip [17]

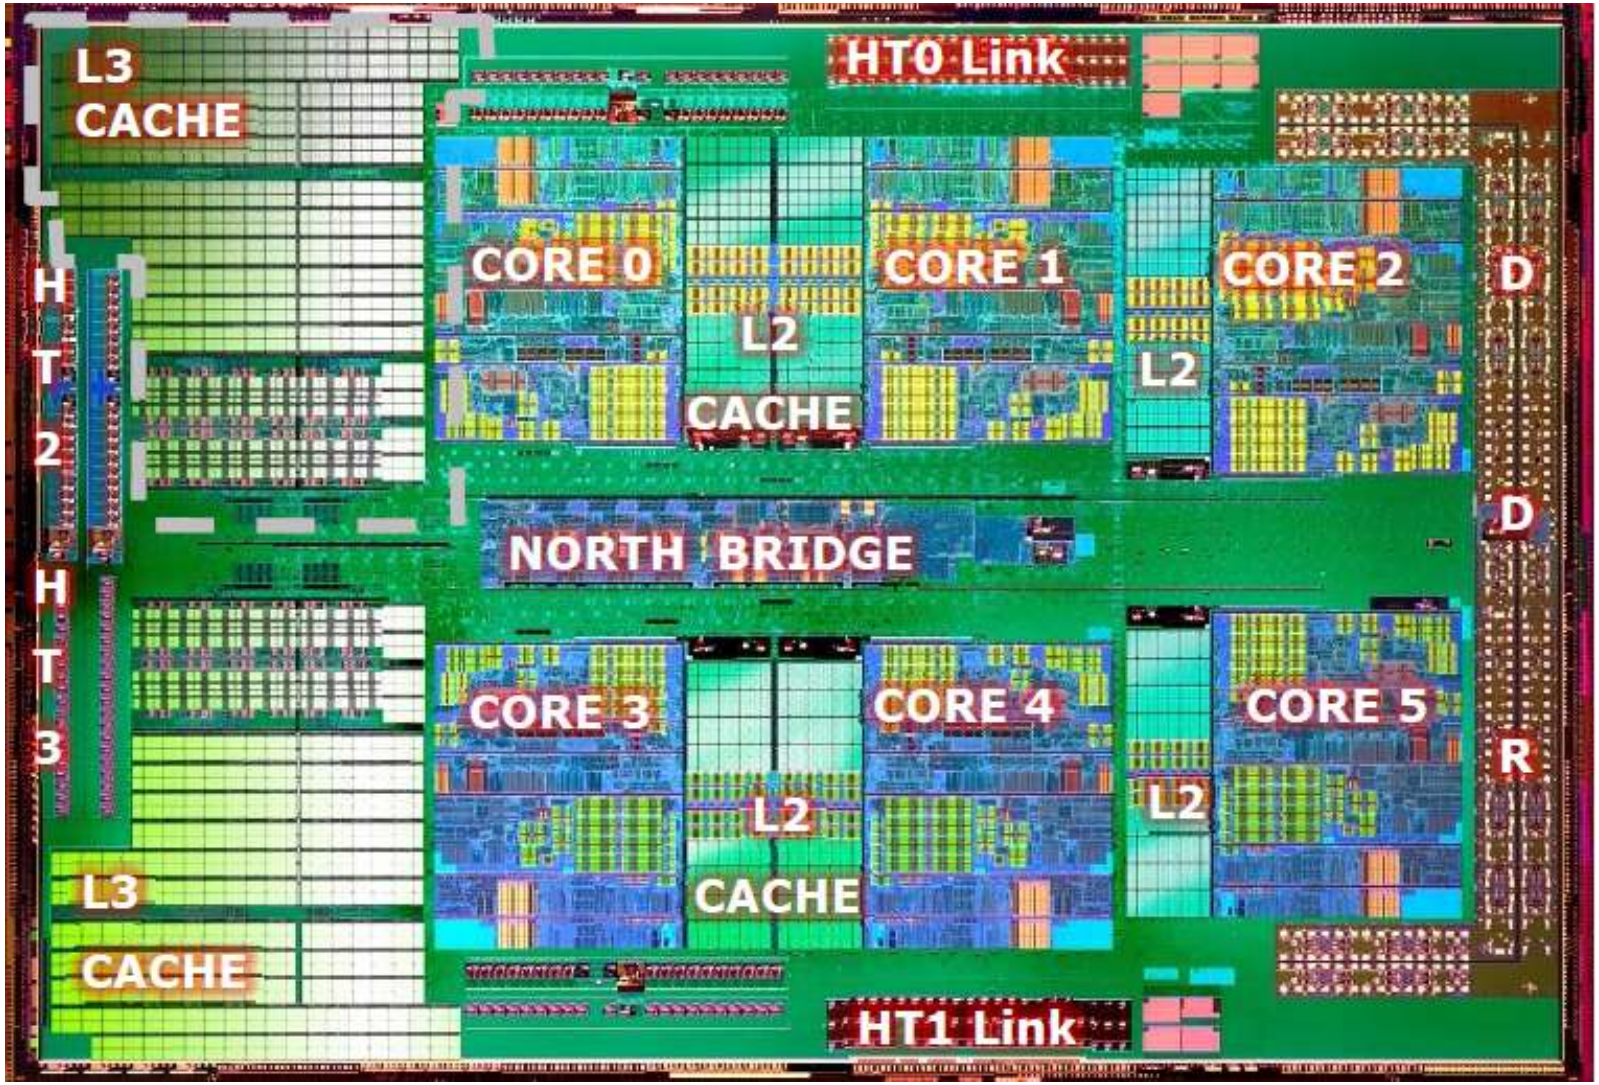

8.2 K10.5 Istambul-based server lines (30)

Main parameters of AMD's K10.5 Istambul-based DP server models [79]

Model	Cores	Clock speed	North bridge/ L3 cache speed	HyperTransport speed	ACP
Opteron 2435	6	2.6GHz	2.2GHz	2.4GHz	75W
Opteron 2431	6	2.4GHz	2.2GHz	2.4GHz	75W
Opteron 2427	6	2.2GHz	2.2GHz	2.4GHz	75W
Opteron 2389	4	2.9GHz	2.2GHz	2.2GHz	75W
Opteron 2387	4	2.8GHz	2.2GHz	2.2GHz	75W
Opteron 2384	4	2.7GHz	2.2GHz	2.2GHz	75W
Opteron 2382	4	2.6GHz	2.2GHz	2.2GHz	75W
Opteron 2380	4	2.5GHz	2.0GHz	2.0GHz	75W
Opteron 2378	4	2.4GHz	2.0GHz	2.0GHz	75W
Opteron 2376	4	2.3GHz	2.0GHz	2.0GHz	75W

8.2 K10.5 Istambul-based server lines (31)

Main parameters of AMD's K10.5 Istambul-based MP server models [79]

Model	Cores	Clock speed	North bridge/ L3 cache speed	HyperTransport speed	ACP
Opteron 8435	6	2.6GHz	2.2GHz	2.4GHz	75W
Opteron 8431	6	2.4GHz	2.2GHz	2.4GHz	75W
Opteron 8389	4	2.9GHz	2.2GHz	2.2GHz	75W
Opteron 8387	4	2.8GHz	2.2GHz	2.2GHz	75W
Opteron 8384	4	2.7GHz	2.2GHz	2.2GHz	75W
Opteron 8382	4	2.6GHz	2.2GHz	2.2GHz	75W
Opteron 8380	4	2.5GHz	2.0GHz	2.0GHz	75W
Opteron 8378	4	2.4GHz	2.0GHz	2.0GHz	75W

8.2 K10.5 Istambul-based server lines (32)

Performance gain of using K10.5-based 6 core Istambul servers vs. K10.5-based 4 core Shanghai servers [144]

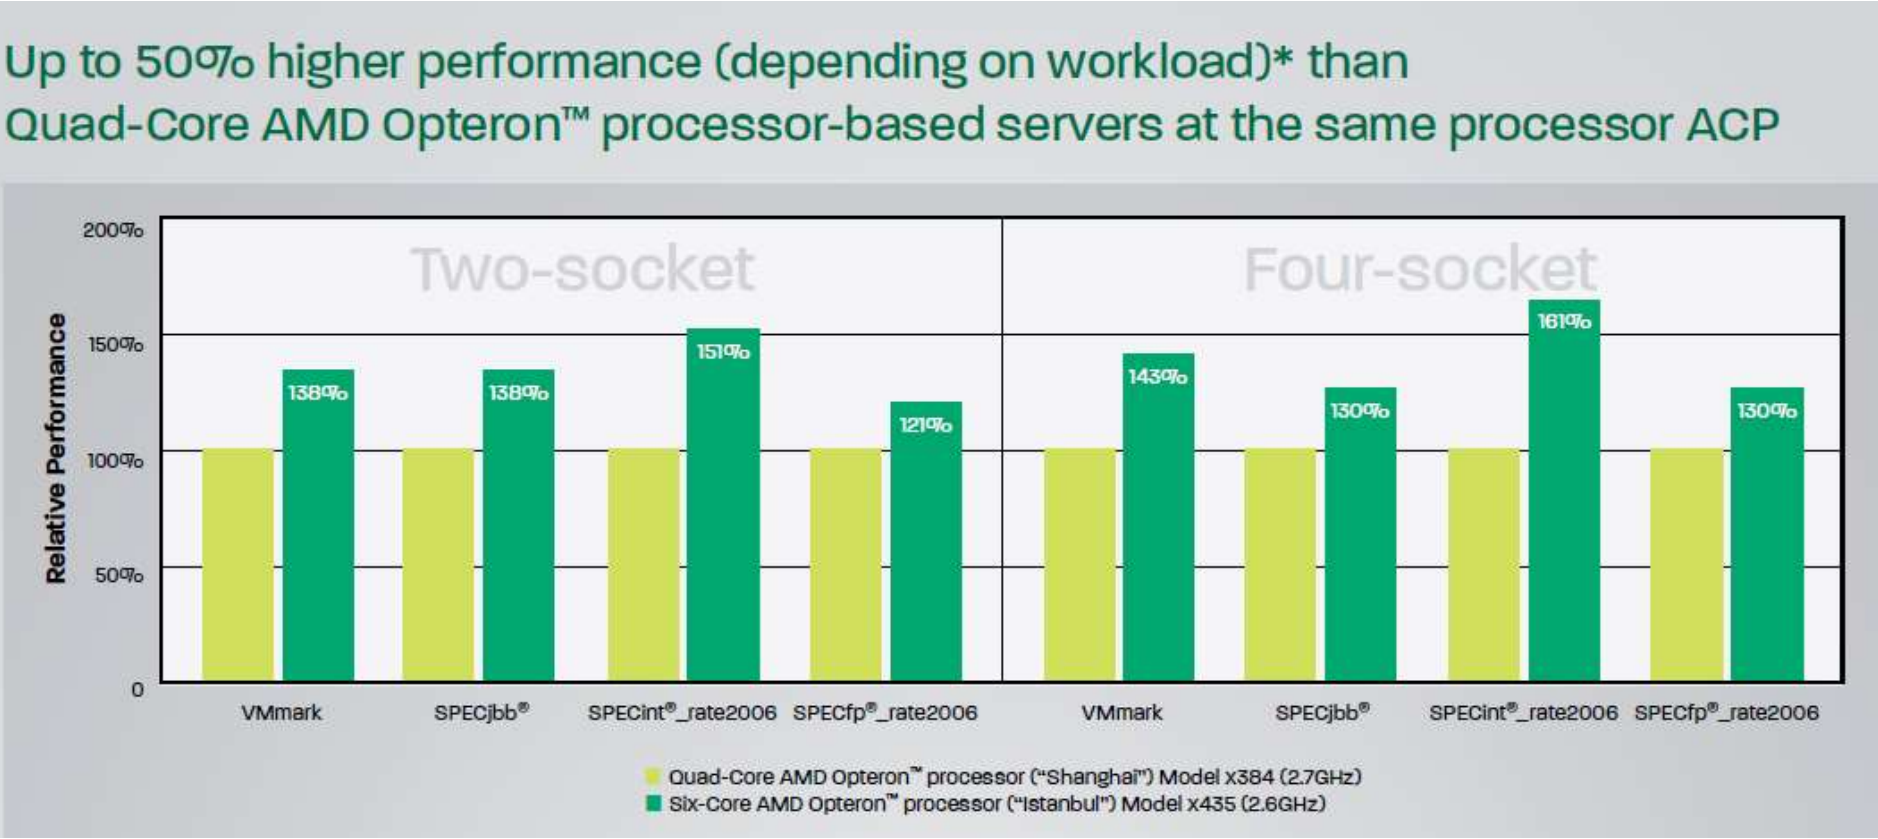

8.2 K10.5 Istambul-based server lines (33)

Relative performance of AMD's servers related to the original Opteron server [16]

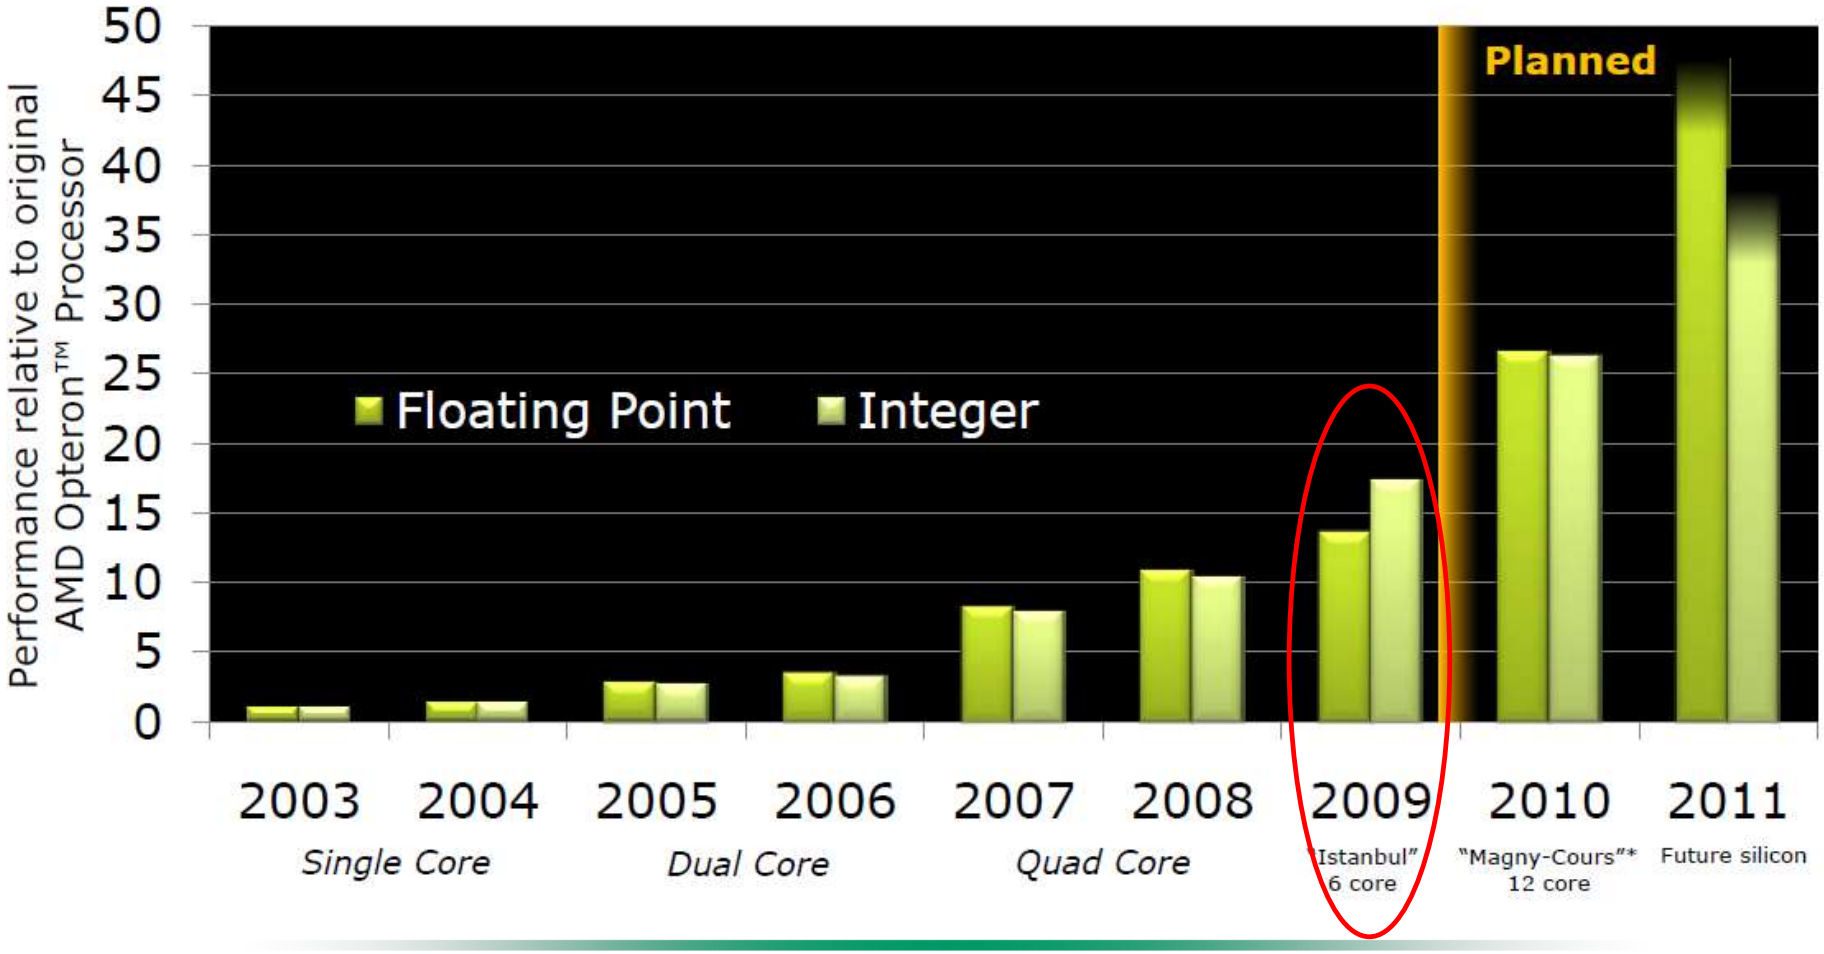

"Shanghai" to "Istanbul" delivers 34% more performance in the same power envelope

8.3 K10.5 Istambul-based desktop lines

8.3 K10.5 Istambul-based desktop lines

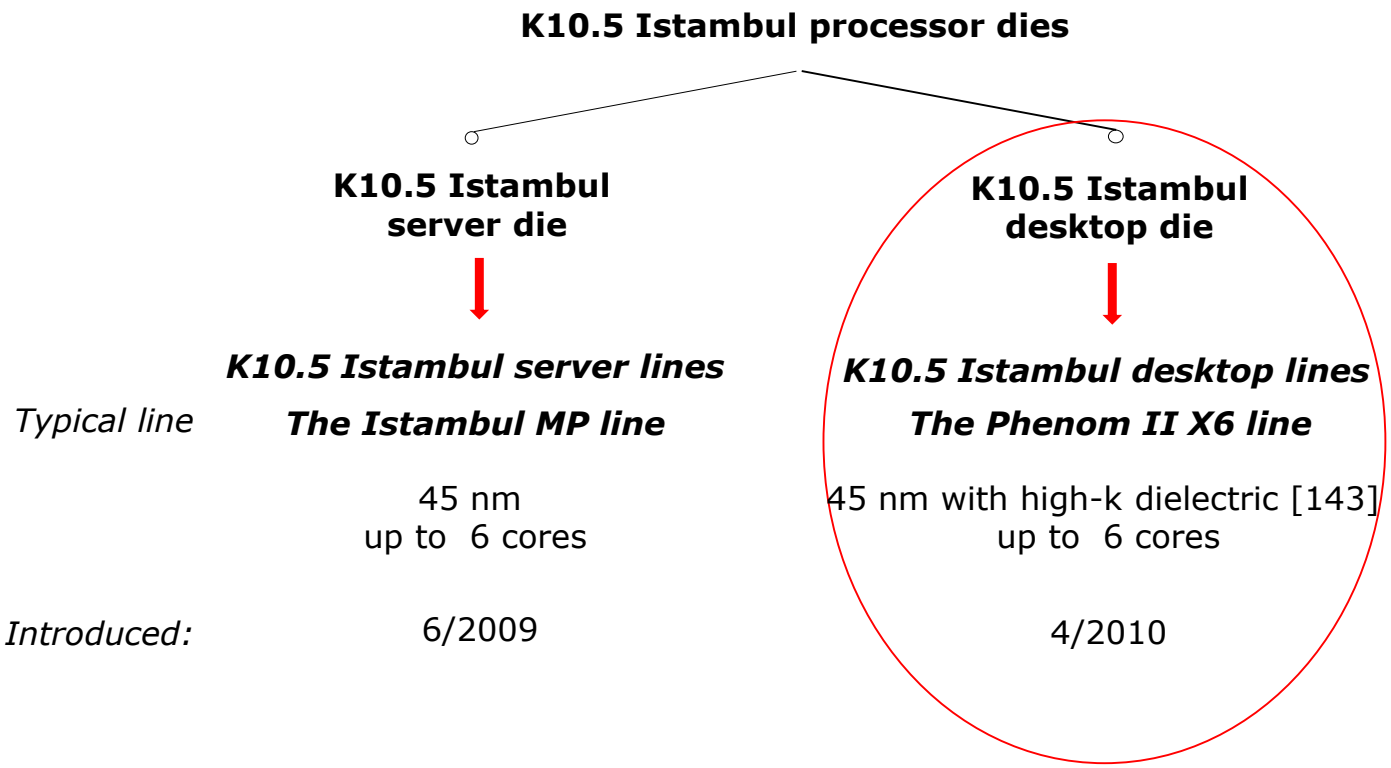

For manufacturing this chip, GlobalFoundries has added a low-k dielectric to its high-performance 45-nm SOI fabrication process, in order to reduce leakage power.

8.3 K10.5 Istambul-based desktop lines (2)

Positioning of AMD's K10.5 Phenom II X6 desktop line

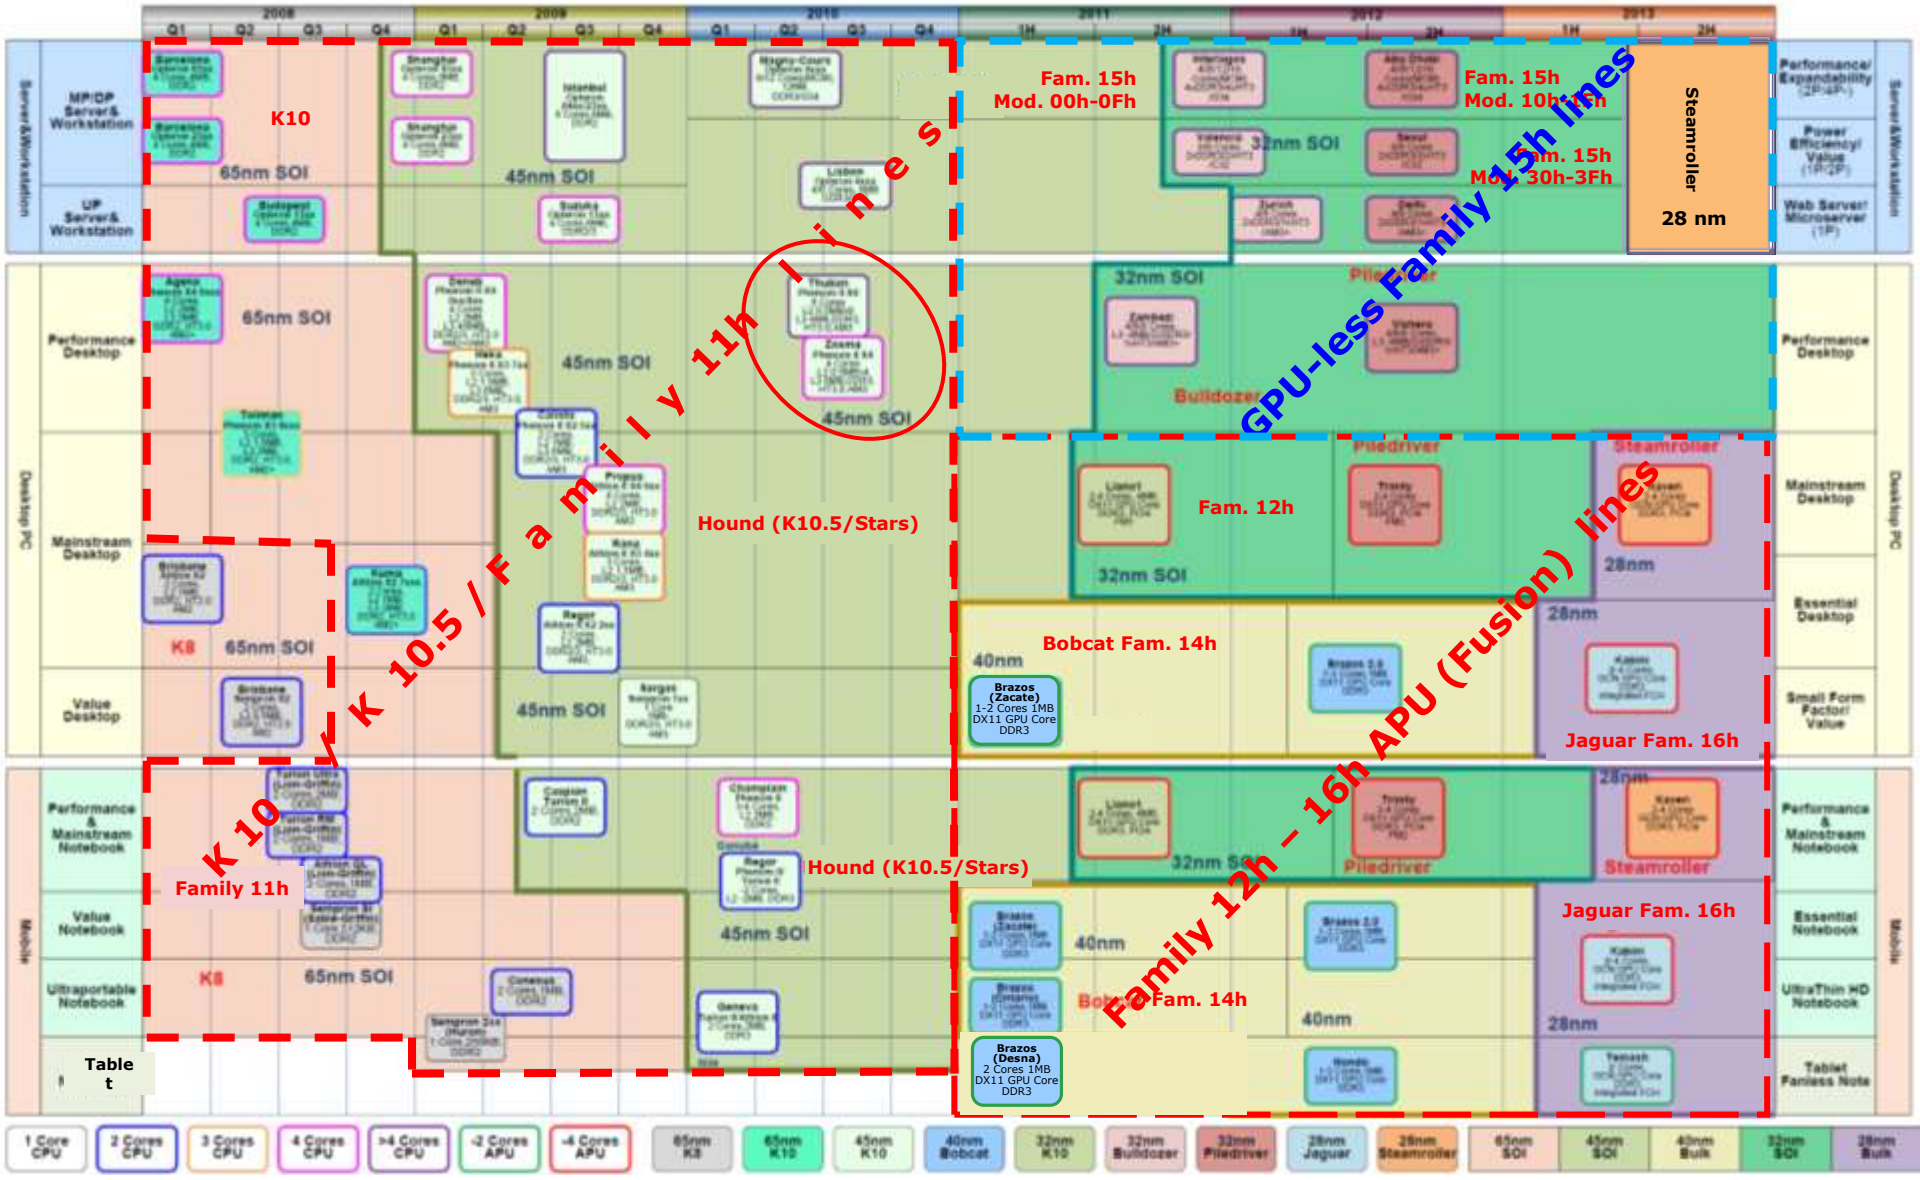

8.3 K10.5 Istambul-based desktop lines (3)

Main features of AMD's high performance K10.5 Phenom II X6 DT line

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. 15h based lines		-	-	-	-	-	-	-	-	-	-

8.3 K10.5 Istambul-based desktop lines (4)

AMD's K10.5 Istambul-based desktop lines – Overview

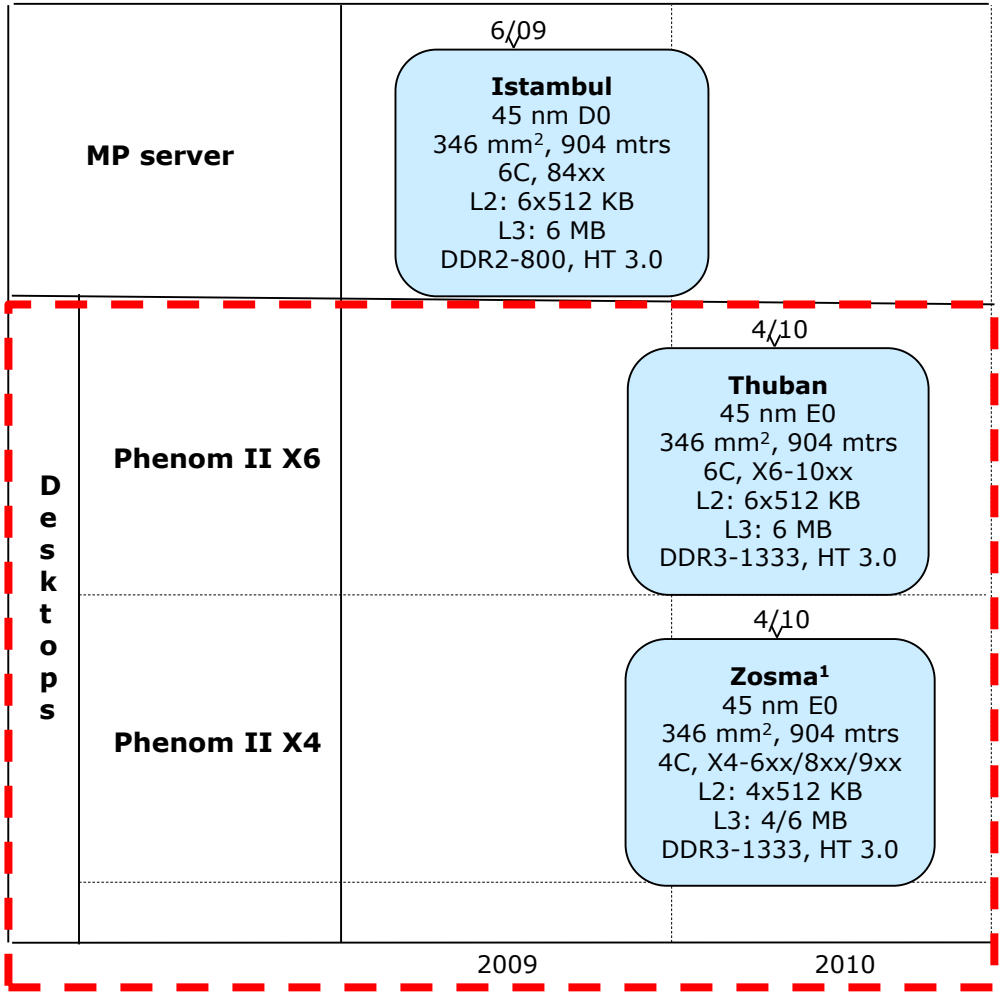

[1] 2 cores disabled

8.3 K10.5 Istambul-based desktop lines (5)

Contrasting the desktop and server aimed 6-core K10.5 Istambul dies [79], [91]

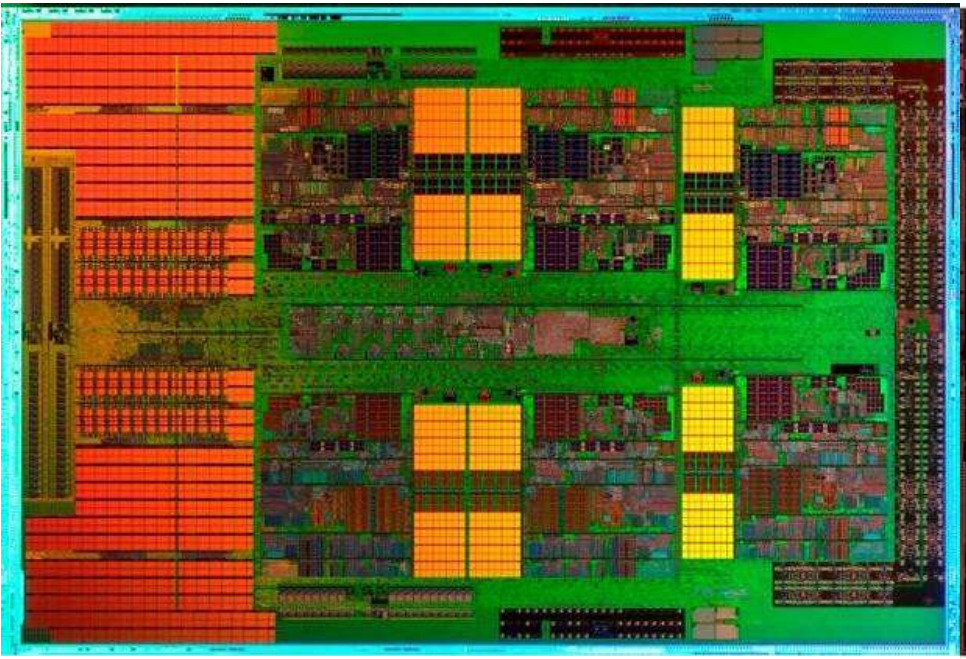

Server aimed **Istambul** die [79]

346 mm^2, 904 mtrs
6C, 6 MB L3, ½ MB/C L2

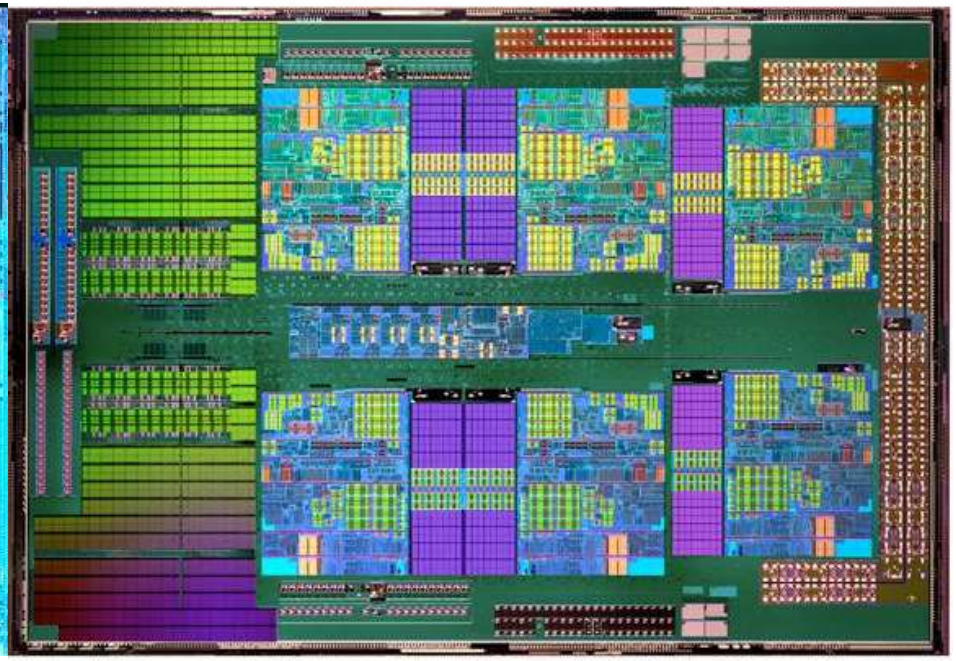

Desktop aimed **Thuban (Phenom II X6)** die [91]

346 mm2, 904 mtrs
6C, 6 MB L3, ½ MB/C L2

8.3 K10.5 Istambul-based desktop lines (6)

Major innovations introduced with the K10.5 Istambul-based desktop line (Phenom II X6 line)

a) Unlocking cores during P-state switches

b) Turbo core technology in select models

8.3 K10.5 Istambul-based desktop lines (7)

a) Unlocking cores during P-state switches

In their Shanghai-based desktop lines AMD linked together the P-states of all four cores by BIOS to avoid performance degradation experienced due to the scheduler policy of the Microsoft's Windows Vista operating system [75].

(In this operating system the scheduler loaded all cores one after the another in a round-robin fashion even in single thread applications that caused performance degradation.)

With Windows 7 (introduced in 10/2009) Microsoft changed their scheduler policy by loading a single core for single threaded applications, consequently also AMD unlocked the P-states of the cores in their Phenom II X6 line allowing independent P-state control for the individual cores [145].

Nevertheless, all cores are supplied by the same voltage further on, that is determined by the core running at the highest clock frequency.

8.3 K10.5 Istambul-based desktop lines (8)

b) Turbo CORE technology in select Thuban (Pnenom II X6) models

(Supported by models ending with the letter T)

The introduction of the Turbo CORE technology became feasible in Istambul-based desktops as for these dies Global Foundries already made use of the 45 nm high-k fabrication process.

While using this process technology leakage becomes reduced and a higher power headroom became available when there are idle cores in the processor.

Principle of operation

When

- three or more cores enter the idle state (Boost eligible state) and
- the active cores are in the P0 state,
- the idle cores (cores in the Boost eligible state) will be placed into the low power, low clock frequency state.

  In this state cores are clocked only at 800 MHz and their supply voltage becomes decreased to a lower value.
- Then up to three active cores will be switched into the PO Boost state.

  In this state both their supply voltage and clock speed will be increased.

  The frequency boost is up to 500 MHz.

The Turbo CORE mode is beneficial for single threaded or light threaded applications as these applications can run at an up to 500 MHz higher clock speed than without Turbo CORE mode.

8.3 K10.5 Istambul-based desktop lines (9)

Comparing AMD's Turbo CORE technique as implemented in their K10.5 Istambul based desktops (Phenom II X6) with Intel's Turbo Boost technique as implemented in their Nehalem line (2008)

AMD's Turbo CORE mode as implemented in their Phenom II X6 desktop line is less efficient than Intel's Turbo Boost as implemented in their Nehalem lines in 2008 [92].

One of the key differences is that in their Nehalem lines Intel already make use of the C6 state with completely shutting down inactive cores through using power gates, so there is a larger temperature headroom that can be utilized by the active cores in the Turbo mode.

Missing the C6 mode and power gates can be the reason why AMD do not allows to switch more than 3 cores into the Turbo CORE mode.

Nevertheless, in their subsequent lines (K12-based Llano, K14-based Bobcat and K15-based Bulldozer) introduced in 2011 also AMD introduced the C6 state along with power gates and improved significantly the implementation of their Turbo CORE technology.

Sequence of P-state transitions in Turbo CORE technology [93]

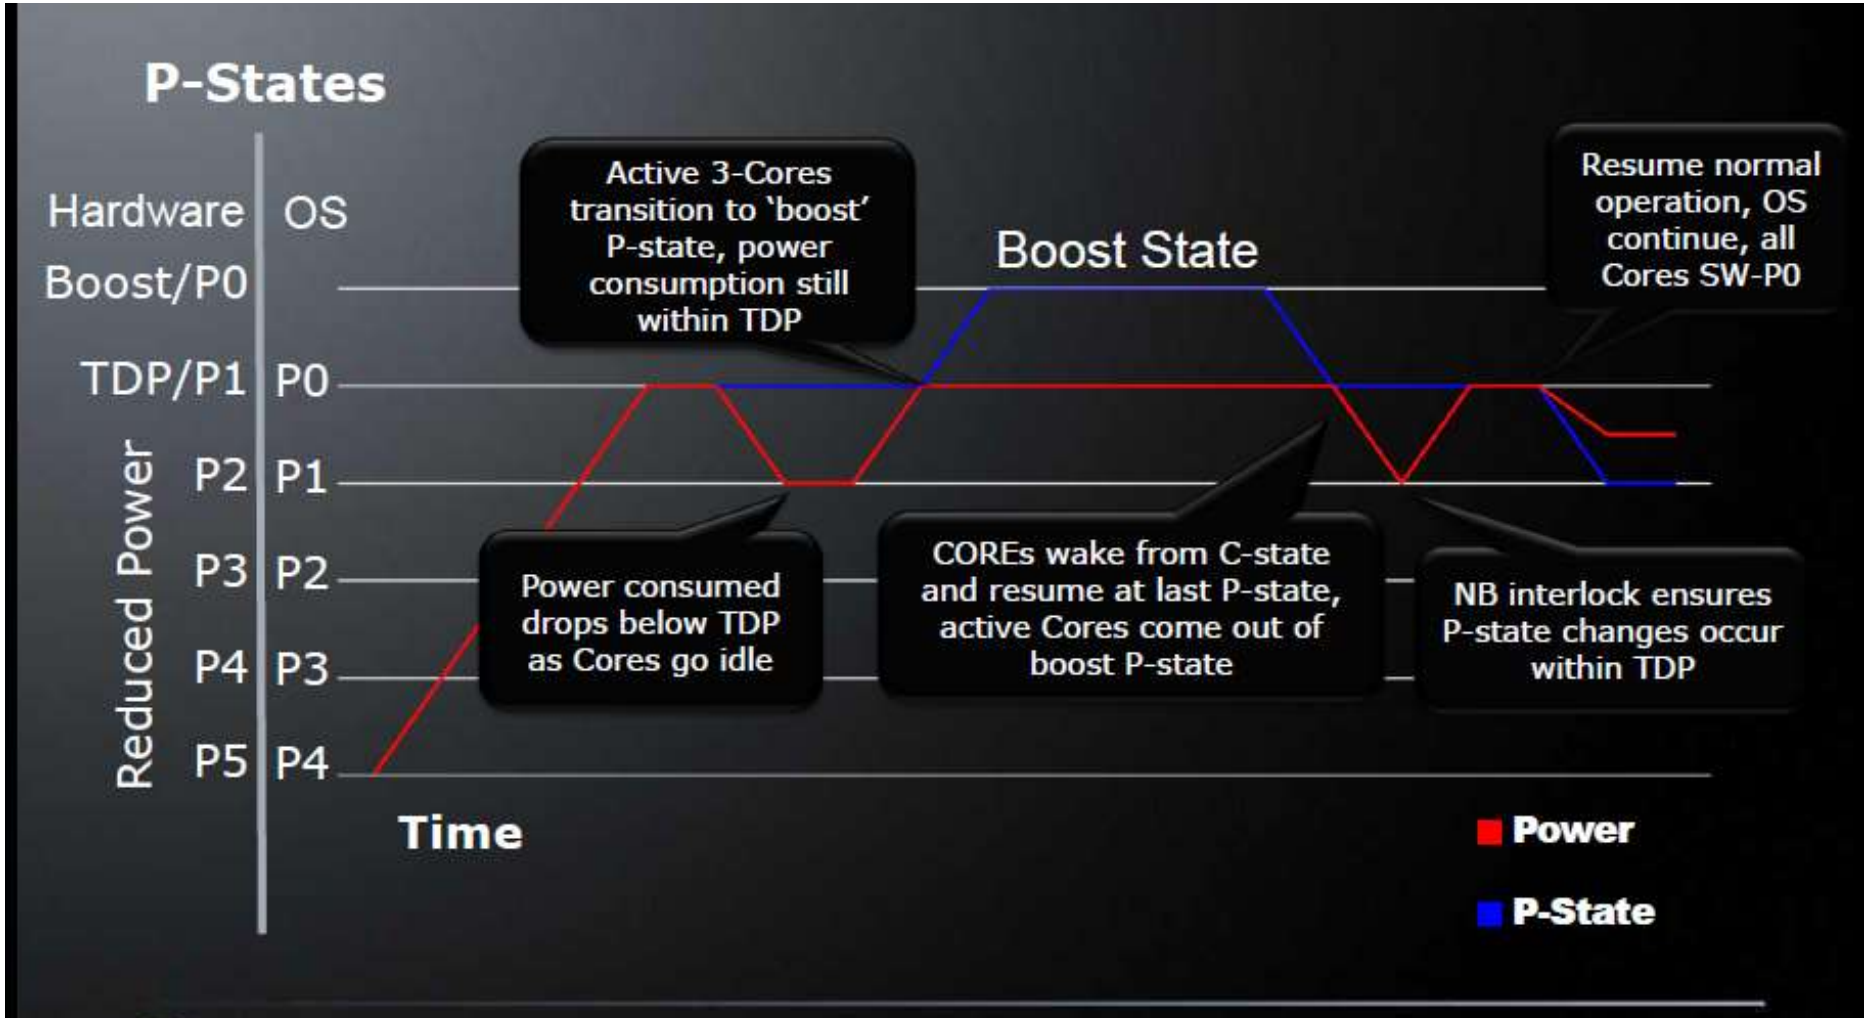

8.3 K10.5 Istambul-based desktop lines (11)

Resulting performance gain when using Turbo CORE technology [93]

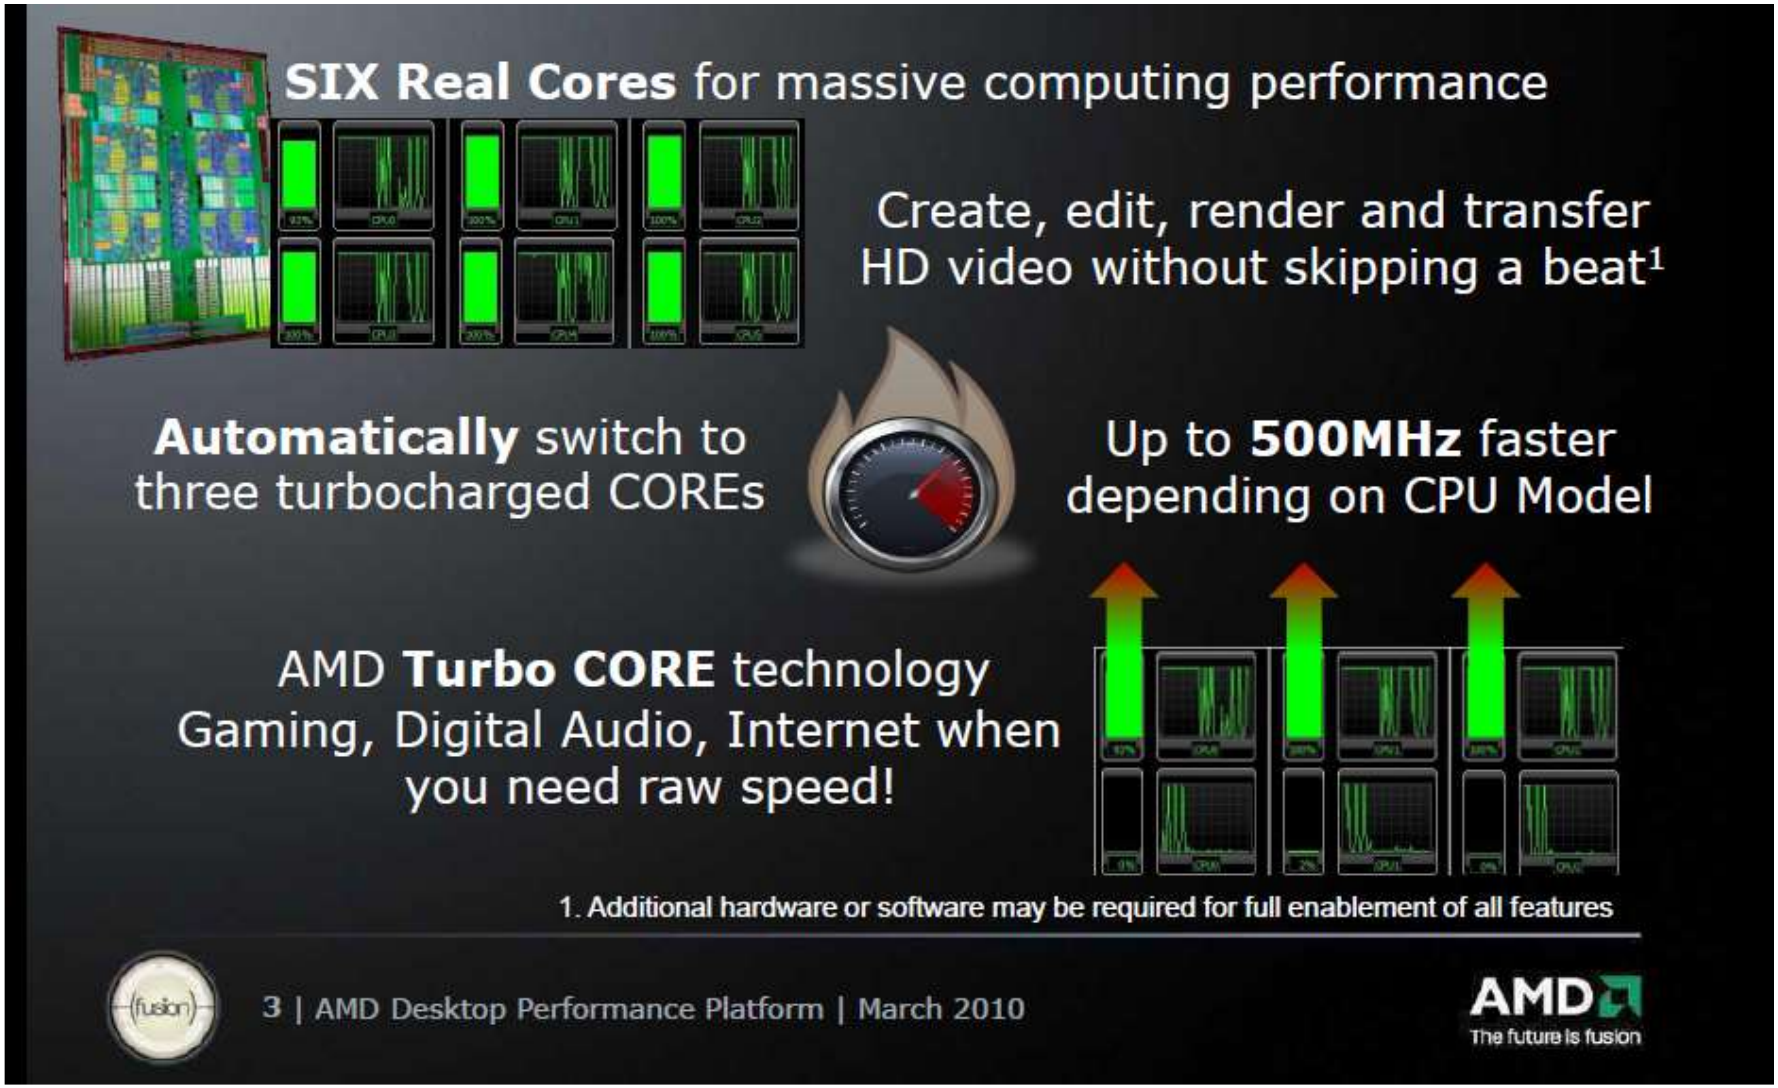

9. The K10.5-based Magny-Course/Lisbon family

- 9.1 Overview of the K10.5 Magny-Course/Lisbon family
- 9.2 Main enhancements of the K10.5 Magny-Course MP servers
- 9.3 K10.5 Magny-Course-based server lines

9.1 Overview of the K10.5 Magny-Course/Lisbon family

9.1 Overview of the K10.5 Magny-Course/Lisbon family

The Magny-Course processor [96]

- Released: 3/2010
- Third (and last) K10.5-based Opteron line, designated as the 6100 series processors
- 2 x 6 Istambul cores
- 692 mm^2 die

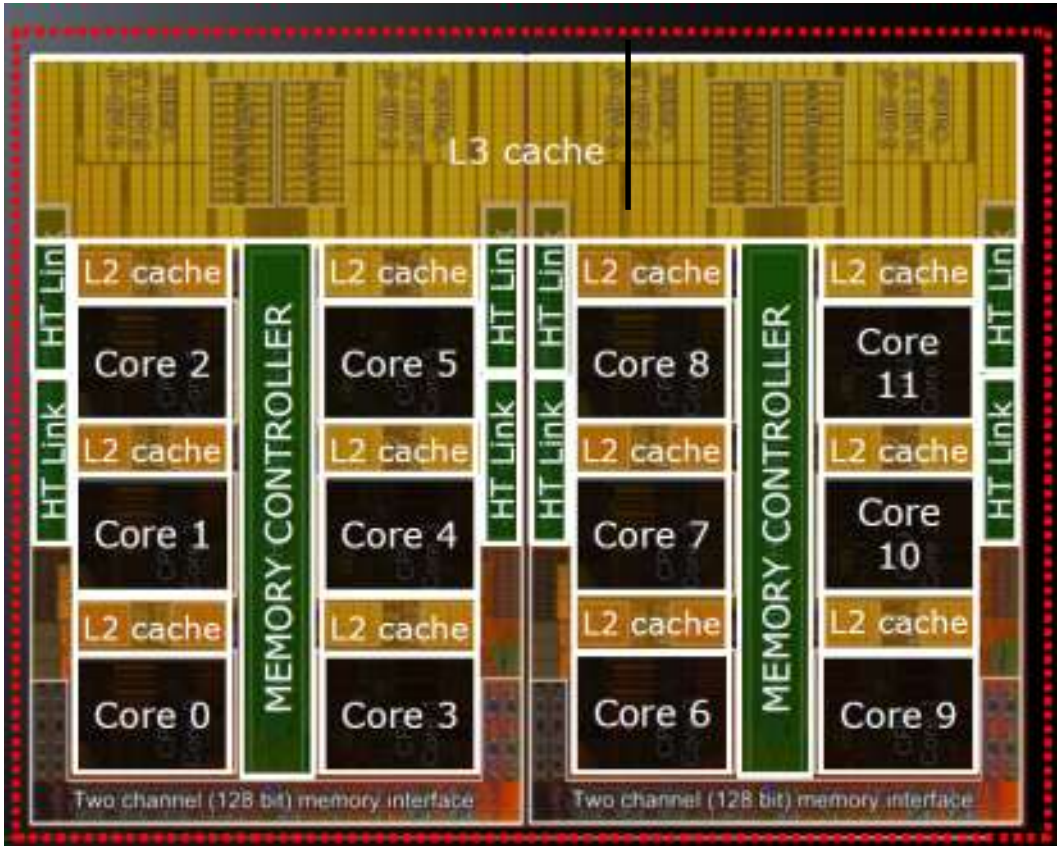

- 4 HT links
- 2 DDR3 memory channels

L1: 64 K I$/64 K D$
L2: 512 K/core
L3: 12 M (shared)

9.1 Overview of the K10.5 Magny-Course/Lisbon family (2)

AMD's server and platform roadmap (DP/MP servers) (based on [88])

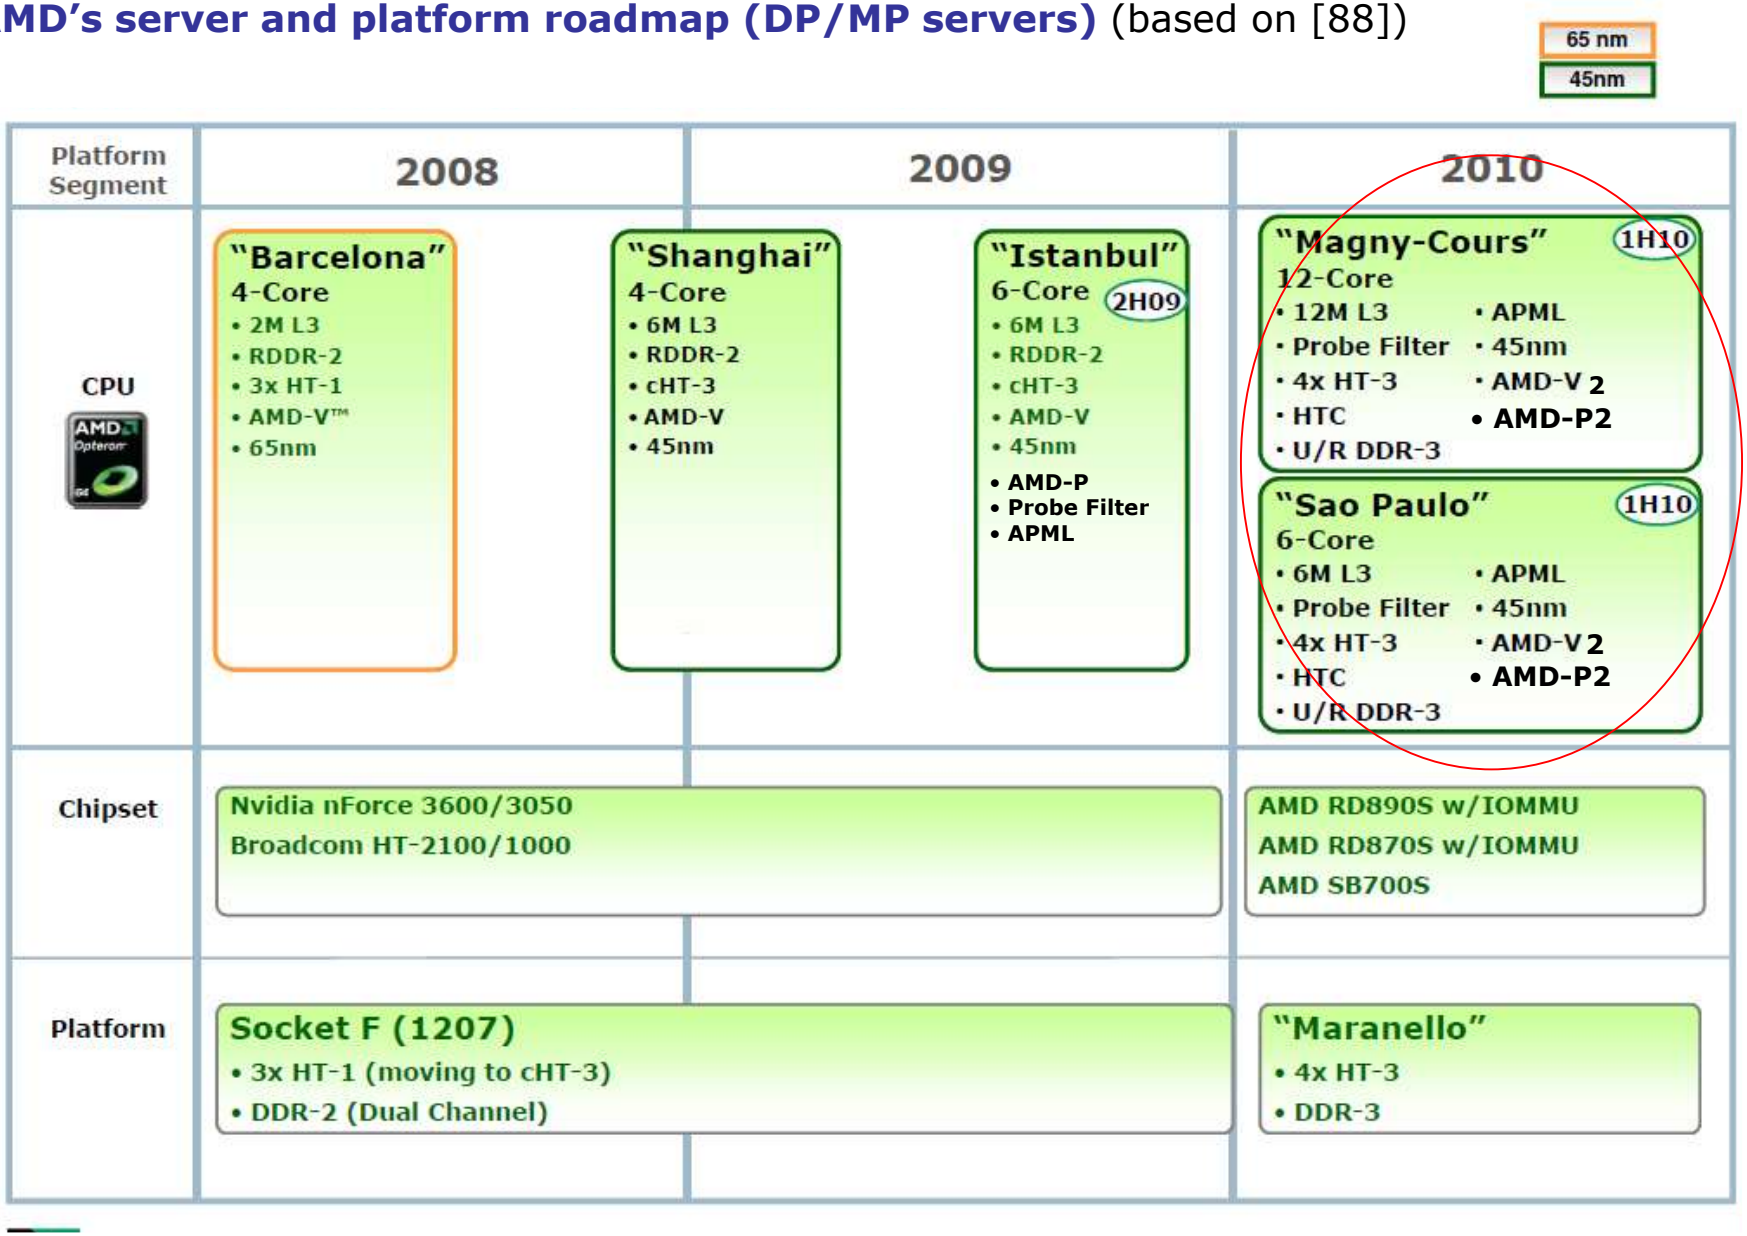

9.1 Overview of the K10.5 Magny-Course/Lisbon family (3)

Positioning of AMD's K10.5-based Magny-Course/Lisbon processor lines [14]

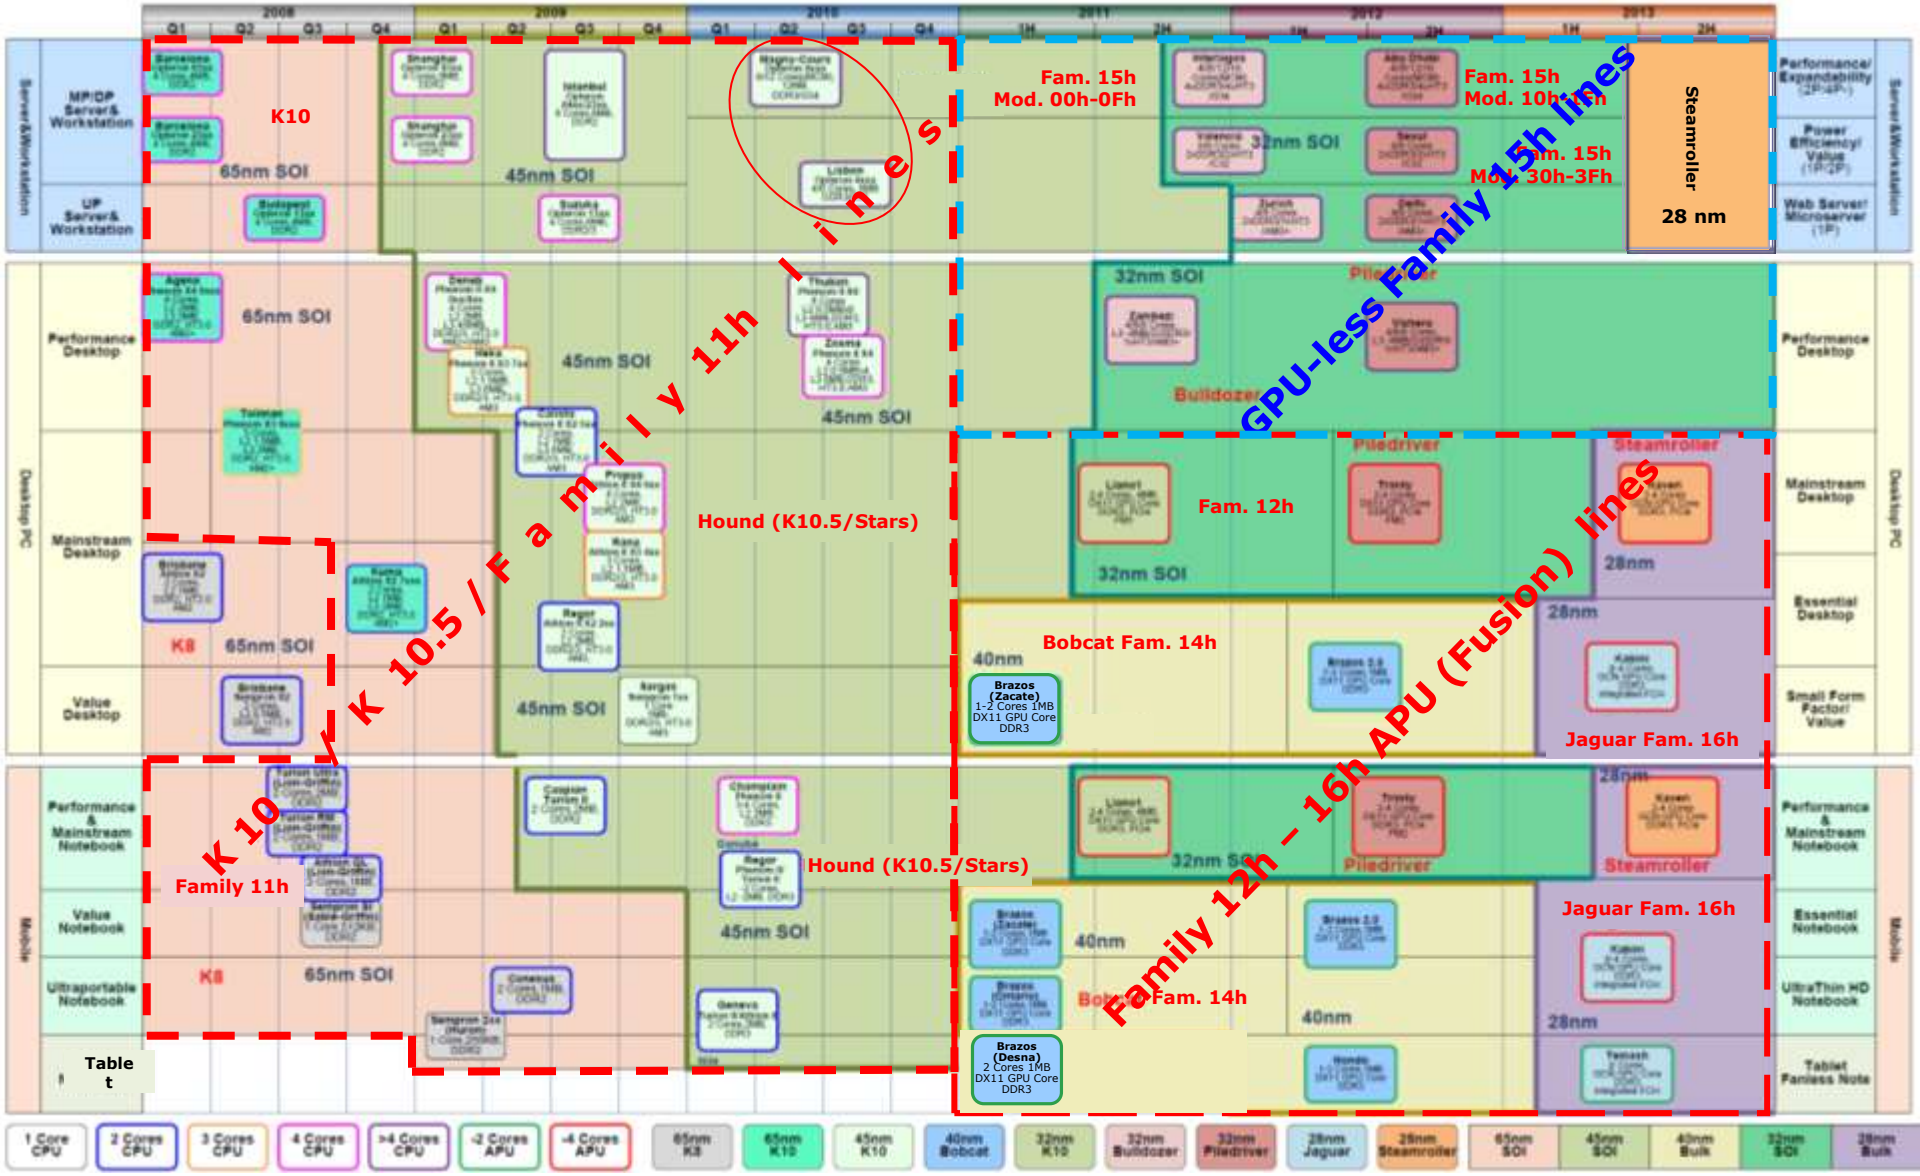

9.1 Overview of the K10.5 Magny-Course/Lisbon family (4)

Brand names of AMD's K10.5h Magny-Course-based server lines

		2003-2007	2007-2008	2008-2011	2009	2009
		K8 (Hammer)	**K10 (Barcelona)**	**K10.5 (Shanghai)**	**K10.5 (Istanbul)**	**K10.5 (Magny- Course)**
Servers	**4P servers**	**See Section 4**	**Barcelona (834x-836x))**	**Shanghai (837x-839x)**	**Istambul (8410-8430)**	**Magny-Course (6100)**
	2P servers		**Barcelona (234x-236x)**	**Shanghai (237x-239x)**	**Istambul (241x-243x)**	**Lisbon (4100)**
	1P servers		**Budapest (135x-136x)**	**Suzuka (138x-139x)**		
Desktops	**High perf.** (~80-120W)		Phenom X4-X2	Phenom II X4-X2	Phenom II X6-X4	
	Mainstream (~60-90W)	Athlon 64 Athlon 64 X2	Athlon X2	Athlon II X4-X2		
	Value (~40-60W)	Sempron		Sempron		
Mobiles	**High perf.** (~30-40W)	Turion 64 X2 (TL 6/5) Turion 64 (ML/MT)		Phenom II (N/P 9xx-6xx) Turion II Ultra (M6xx) Turion II (M/N/P 5xx)		
	Mainstream (~20-30W)	Athlon 64 X2 (TK-5x/4x) Athlon 64 (2xxx+-4xxx+)		Athlon II (M/N/P 3xx) Sempron (M1xx)		
	Ultraportable (~10-20W)	Mobile Sempron (2xxx+-4xxx+) Sempron 2100 fanless		Turion II Neo (K6xx) Athlon II Neo (K1xx) V-series (V1xx)		
Embedded (~10-20W)				Turion II Neo X2 Athlon II Neo X2 Athlon II Neo		

9.1 Overview of the K10.5 Magny-Course/Lisbon family (5)

The Magny Course die [94]

Magny Course die
692 mm^2
(2 x Istambul die)
(2x346 mm^2)

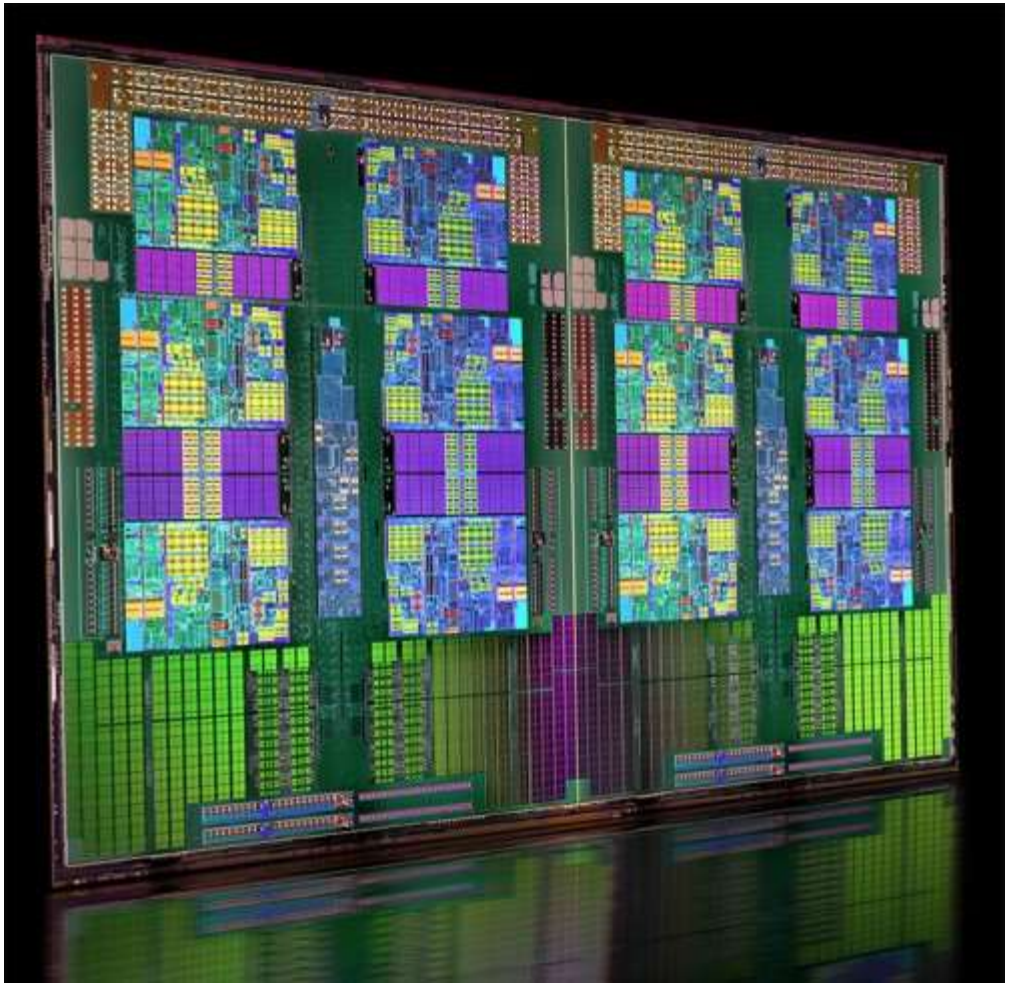

9.1 Overview of the K10.5 Magny-Course/Lisbon family (6)

½ of the Magny-Course die (actually the Istambul die) [17]

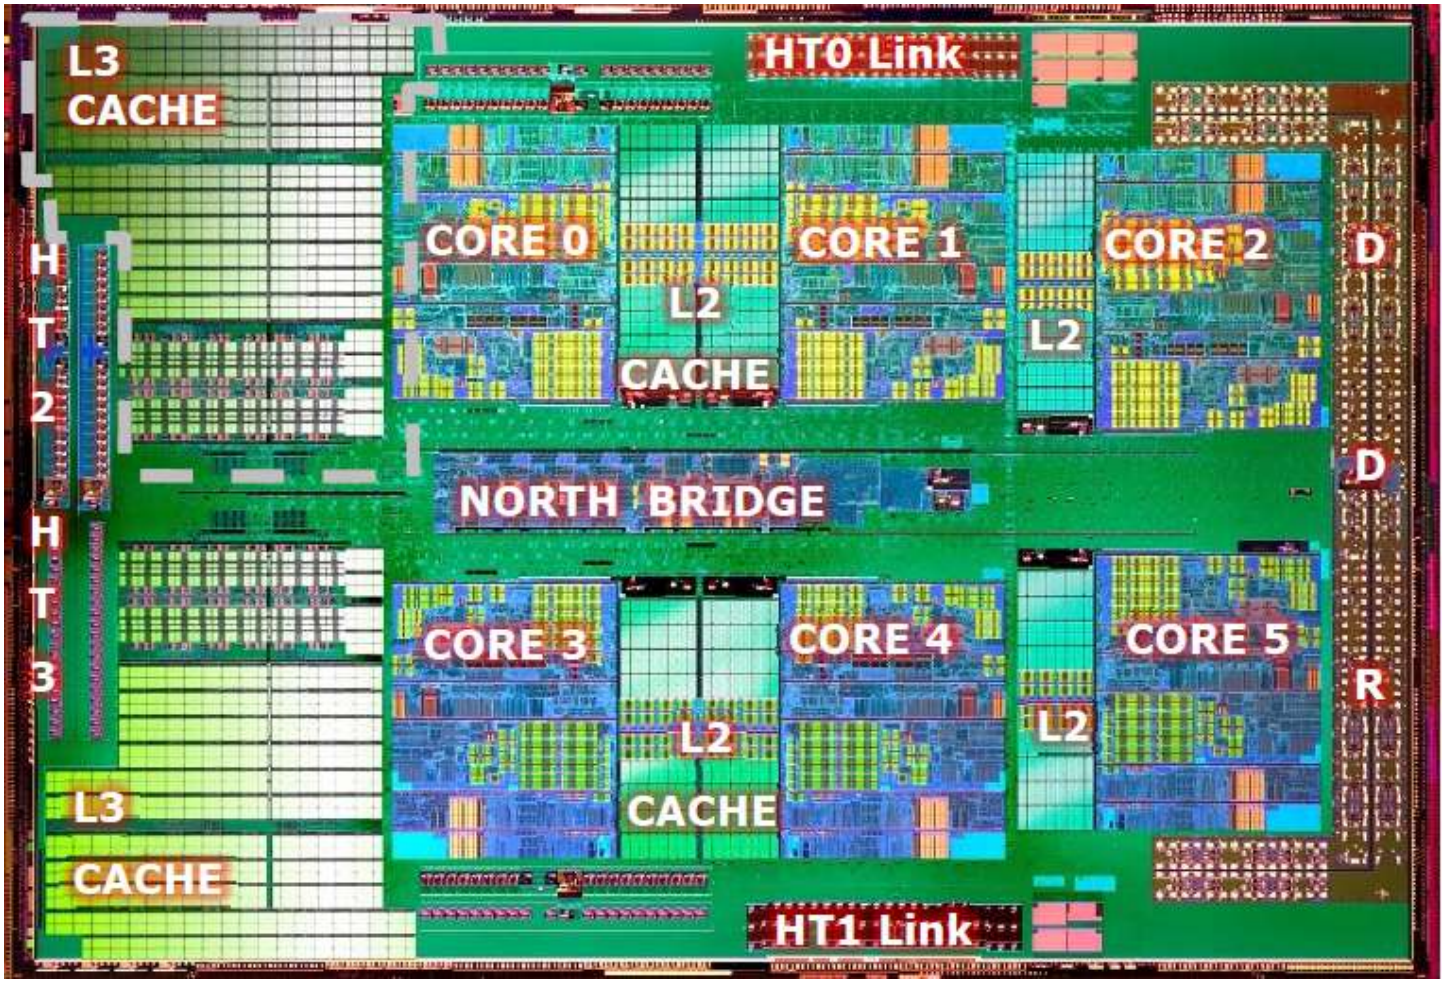

9.1 Overview of the K10.5 Magny-Course/Lisbon family (7)

Main features of AMD's K10.5 Magny-Course-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)	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 (NPT)	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. 17h Mod. 00h-0Fh		6/2017	Epyc (2S!!) (4 dies/proc.)	7000	14 nm	4x(2x4) (32C)	½ MB/C	2 MB/C	DDR4-2666	IFIS 75.8 GB/s	SP3

9.1 Overview of the K10.5 Magny-Course/Lisbon family (8)

Sockets of Magny-Course/Lisbon DP/MP servers [based on 42]

Supports DDR3

Socket **G34** (1944 pins, MP servers K10.5 „Magny-Course", Fam. 15h)
Socket **C32** (1207 pins, DP servers K10.5 „Magny –Course, Fam. 15h)
Socket **AM3** (941 pins, K10.5 Shanghai UP servers, desktops, except first released desktops, K10.5 Istambul desktops)
Socket **AM3+** (942 pins, UP servers Fam. 15h, desktops Fam. 15h)
Socket **FM1** (905 balls BGA, desktops, mobiles Fam. 12h)
Socket **S1g4** (638 pins, K10.5 mobile)
Socket **FT1** (413 balls BGA, desktops, mobiles Fam. 14h)
Socket **FT3** (769 balls BGA, desktops, mobiles Fam. 16h)

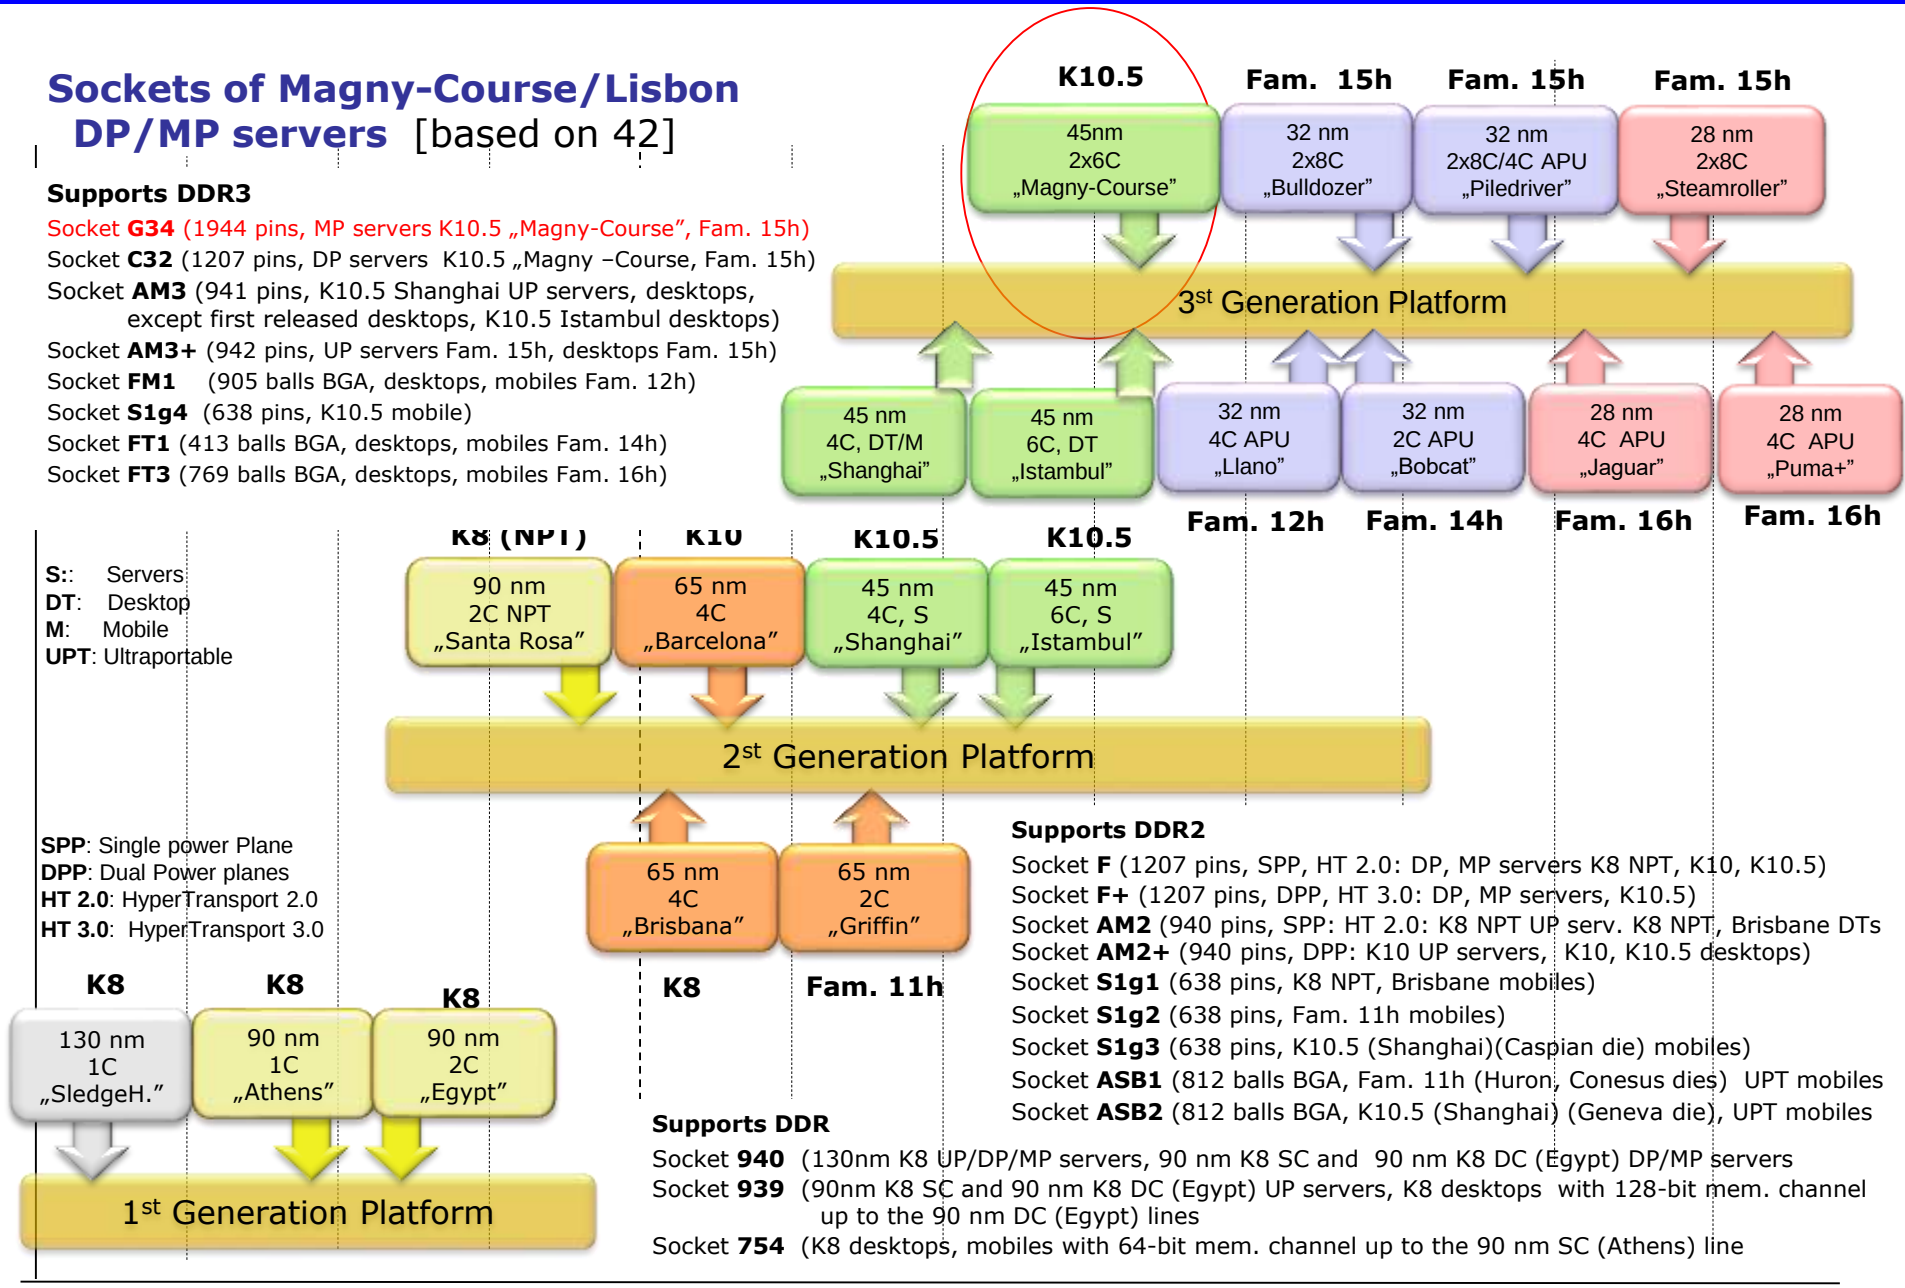

Supports DDR2

Socket **F** (1207 pins, SPP, HT 2.0: DP, MP servers K8 NPT, K10, K10.5)
Socket **F+** (1207 pins, DPP, HT 3.0: DP, MP servers, K10.5)
Socket **AM2** (940 pins, SPP: HT 2.0: K8 NPT UP serv. K8 NPT, Brisbane DTs
Socket **AM2+** (940 pins, DPP: K10 UP servers, K10, K10.5 desktops)
Socket **S1g1** (638 pins, K8 NPT, Brisbane mobiles)
Socket **S1g2** (638 pins, Fam. 11h mobiles)
Socket **S1g3** (638 pins, K10.5 (Shanghai)(Caspian die) mobiles)
Socket **ASB1** (812 balls BGA, Fam. 11h (Huron, Conesus dies) UPT mobiles
Socket **ASB2** (812 balls BGA, K10.5 (Shanghai) (Geneva die), UPT mobiles

Supports DDR

Socket **940** (130nm K8 UP/DP/MP servers, 90 nm K8 SC and 90 nm K8 DC (Egypt) DP/MP servers
Socket **939** (90nm K8 SC and 90 nm K8 DC (Egypt) UP servers, K8 desktops with 128-bit mem. channel up to the 90 nm DC (Egypt) lines
Socket **754** (K8 desktops, mobiles with 64-bit mem. channel up to the 90 nm SC (Athens) line

2003 2004 2005 2006 2007 2008 2009 2010 2011 2012 2013 2014

9.2 Main enhancements of the K10.5 Magny-Course servers

9.2 Main enhancements of the K10.5 Magny-Course MP servers [89]

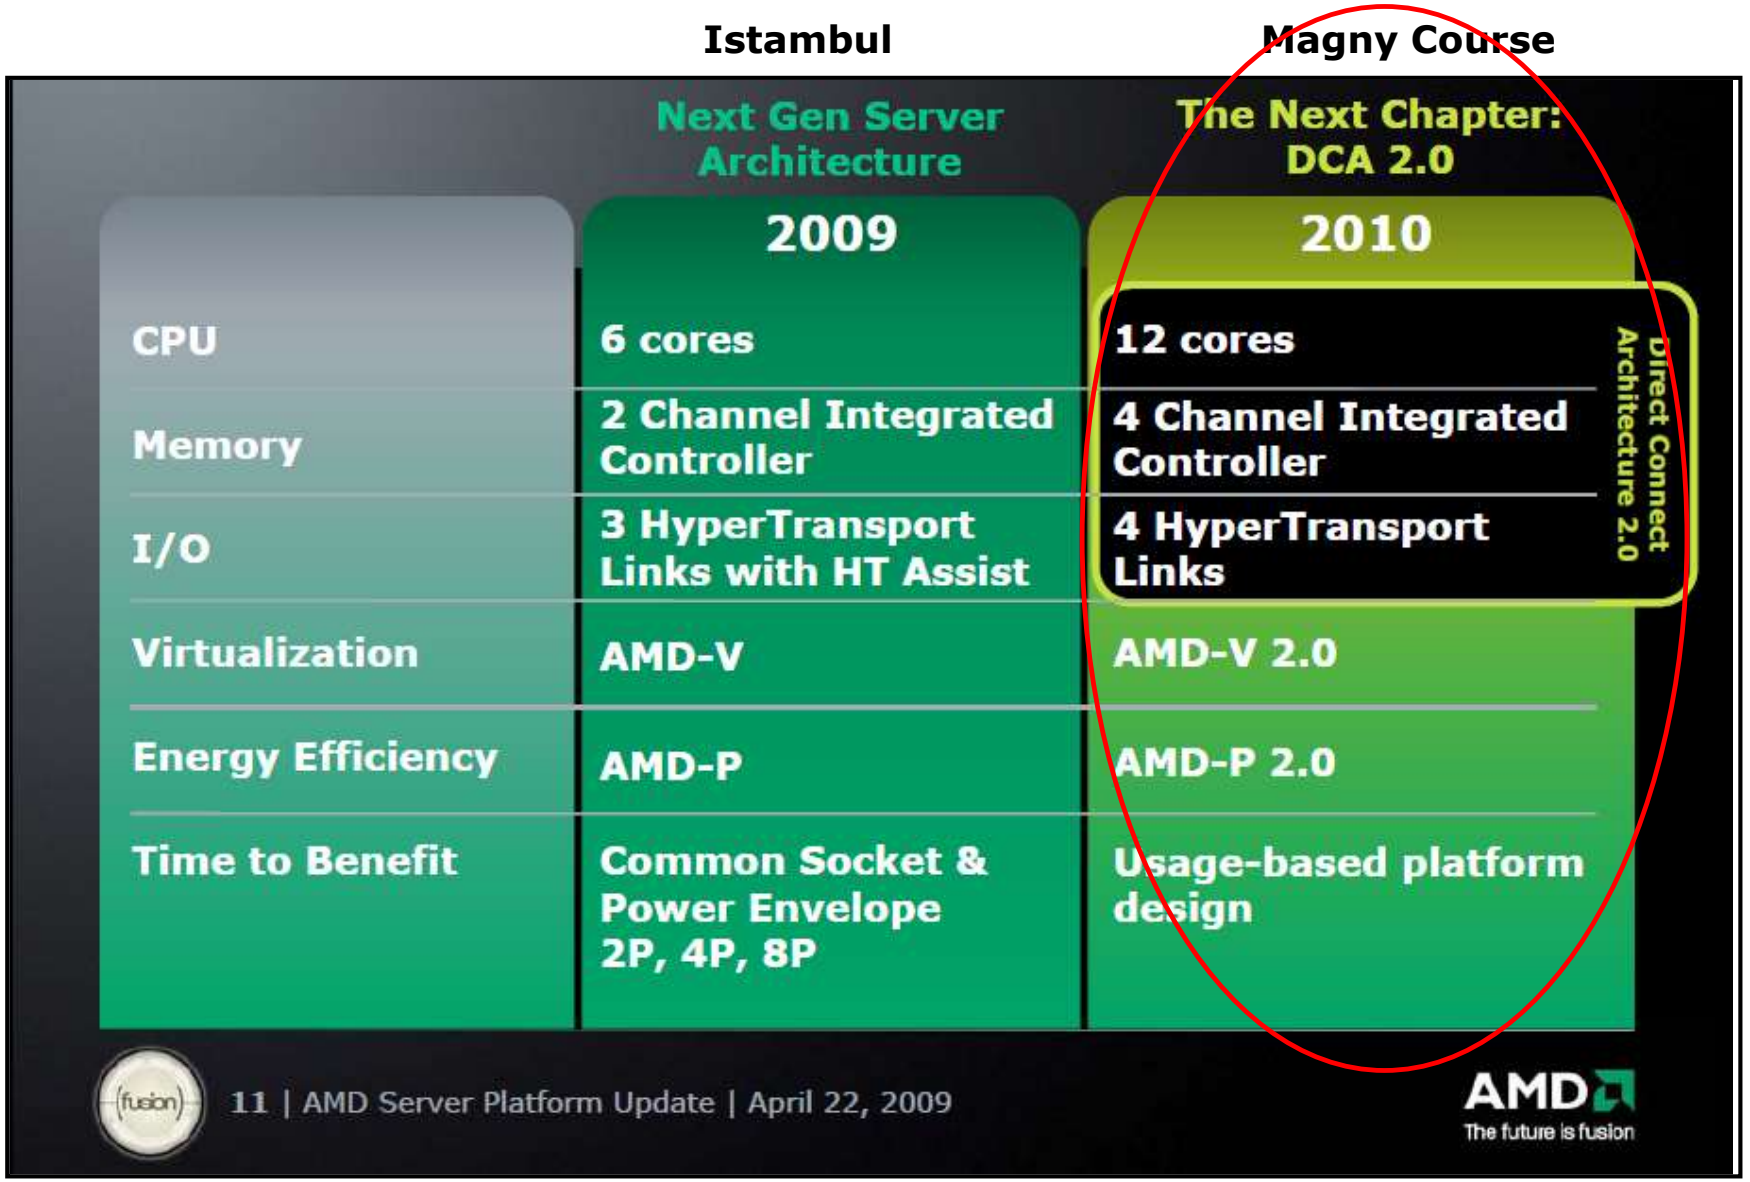

9.2 Main enhancements of the K10.5 Magny-Course servers (2)

a) Direct Connect Architecture 2.0

The previous implementation

Direct Connect Architecture 1.0, introduced with K8 MP servers in 2003 [95].

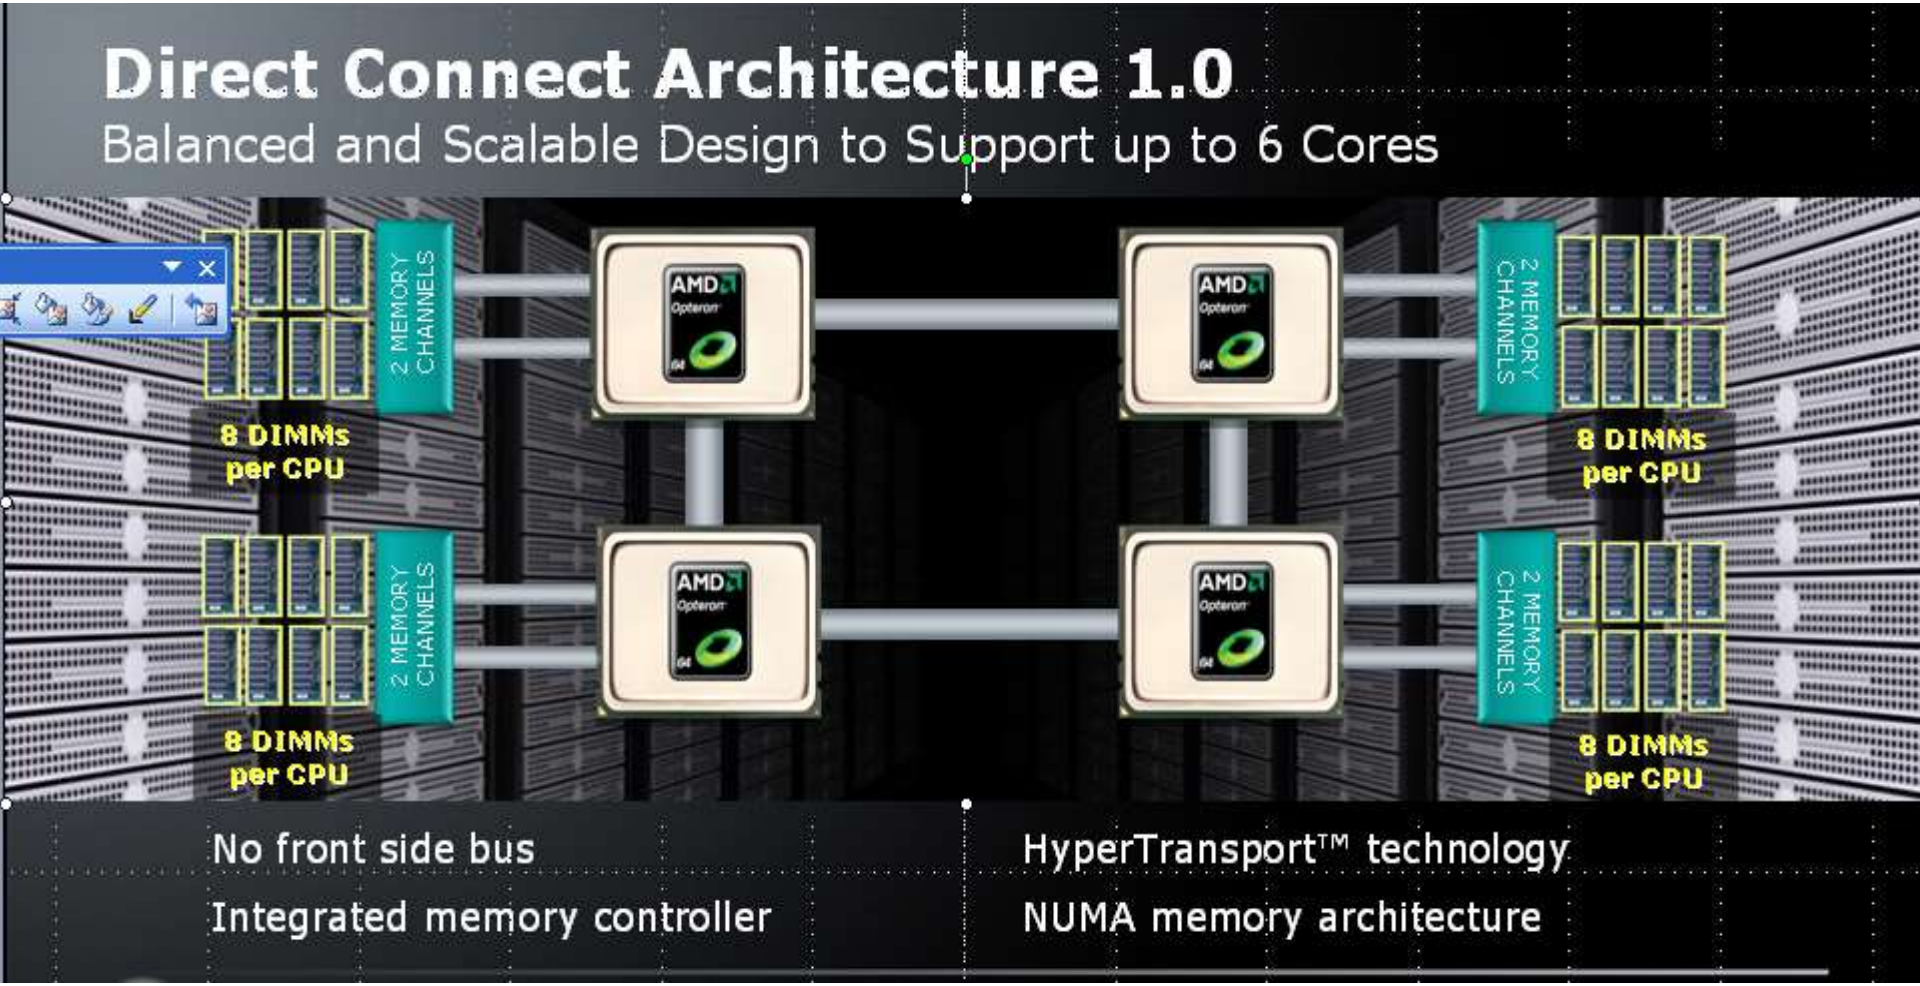

9.2 Main enhancements of the K10.5 Magny-Course servers (3)

Direct Connect Architecture 2.0 [95]

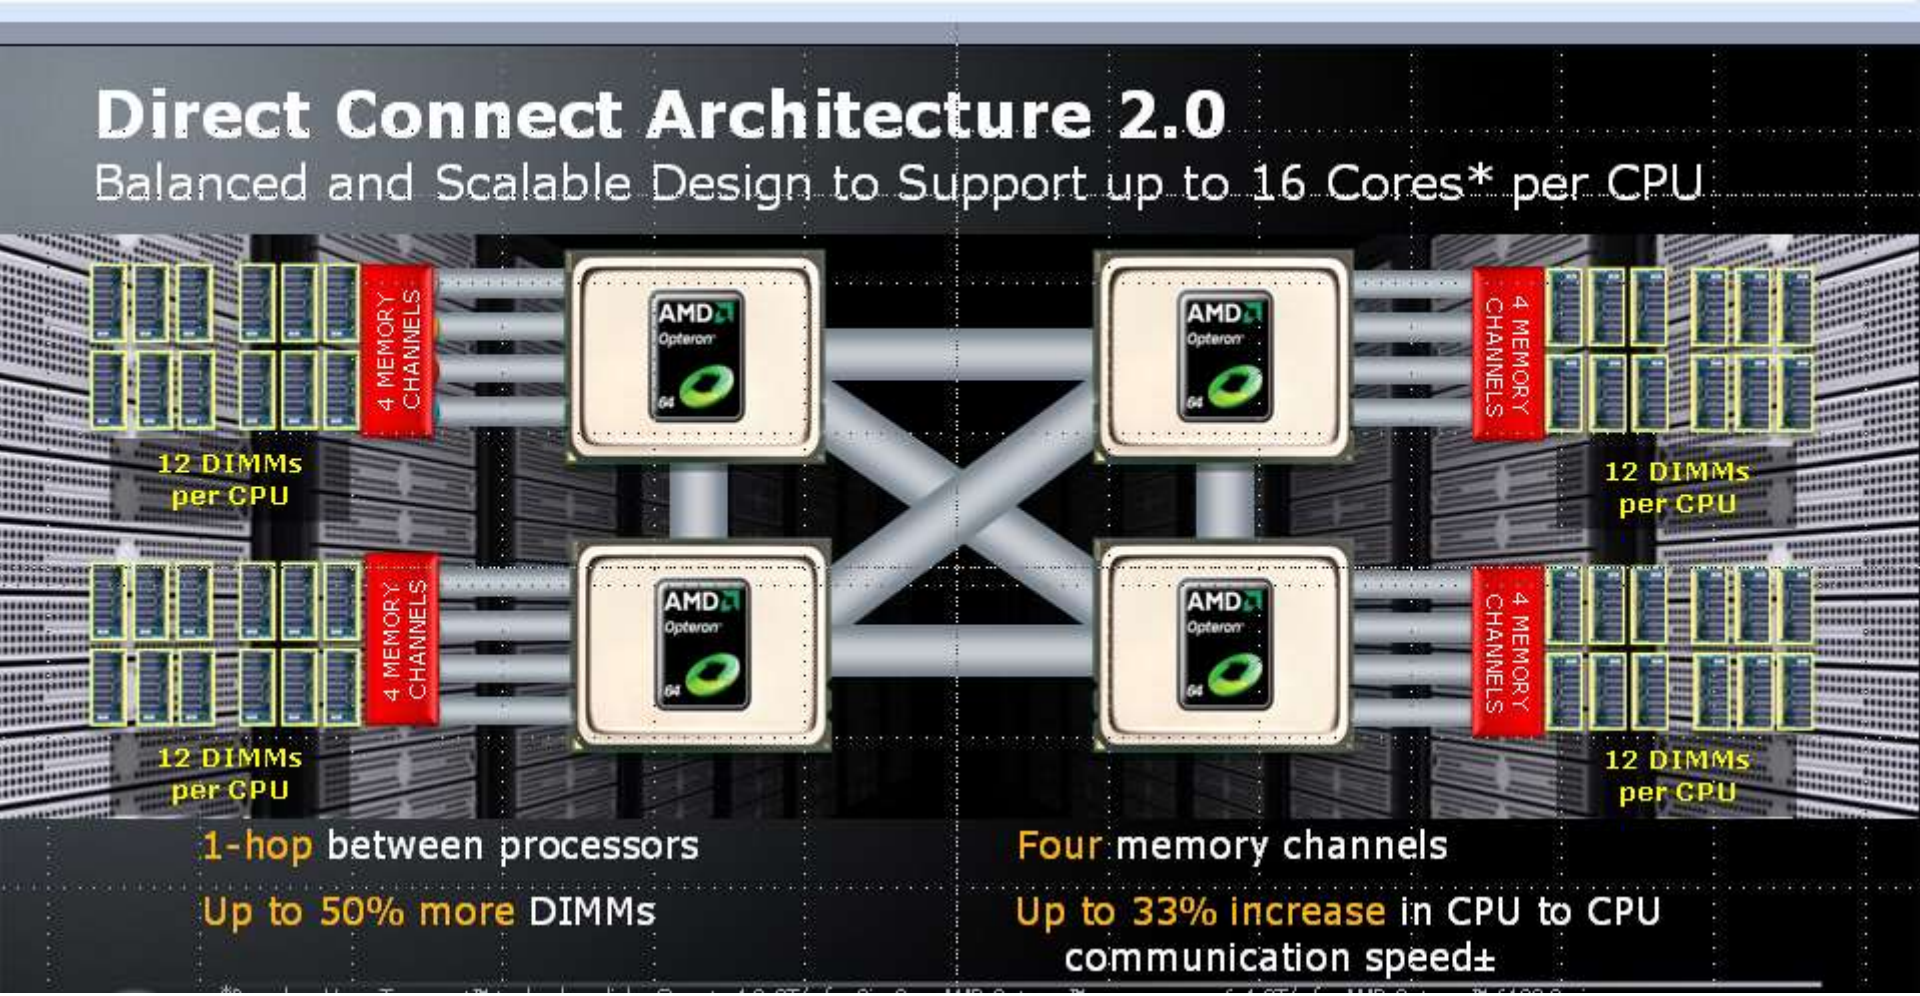

9.2 Main enhancements of the K10.5 Magny-Course servers (4)

Available interconnections of the chip (two Istambul dies) [96]

- 4 DDR3 memory channels (DDR3-1333)
- 2 external 16-bit cache coherent HT 3.0 links or
- 2 external 8-bit cache coherent HT 3.0 links (= 1 external 16-bit cache coherent HT 3.0 link)
- 1 external 16-bit non cache coherent HT 3.0 link

Non cc
HT 3.0 link
(e.g. to a GPU)

- Both dies are internally interconnected by a 16 + 8 bit HT link.

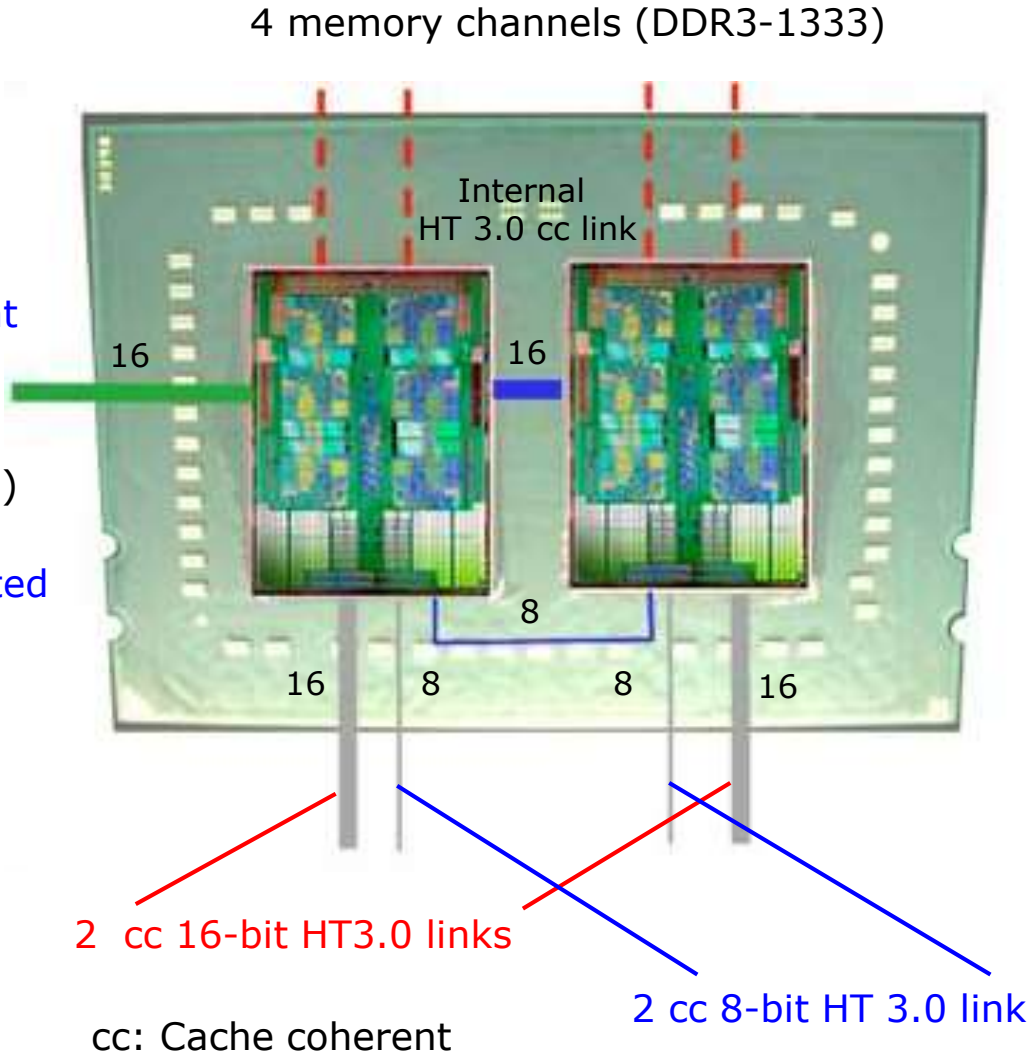

9.2 Main enhancements of the K10.5 Magny-Course servers (5)

The possibility for splitting HT 3.0 16-bit links into two 8-bit wide links

The HT 3.0 protocol allows to split each 16-bit link to two 8-bit wide links. This features can be utilized in MP servers to build fully connected 8P systems with 8-bit wide links, as shown in the next Figure.

9.2 Main enhancements of the K10.5 Magny-Course servers (6)

Possible uses of four HT 3.0 links in 4P and 8P servers [52]

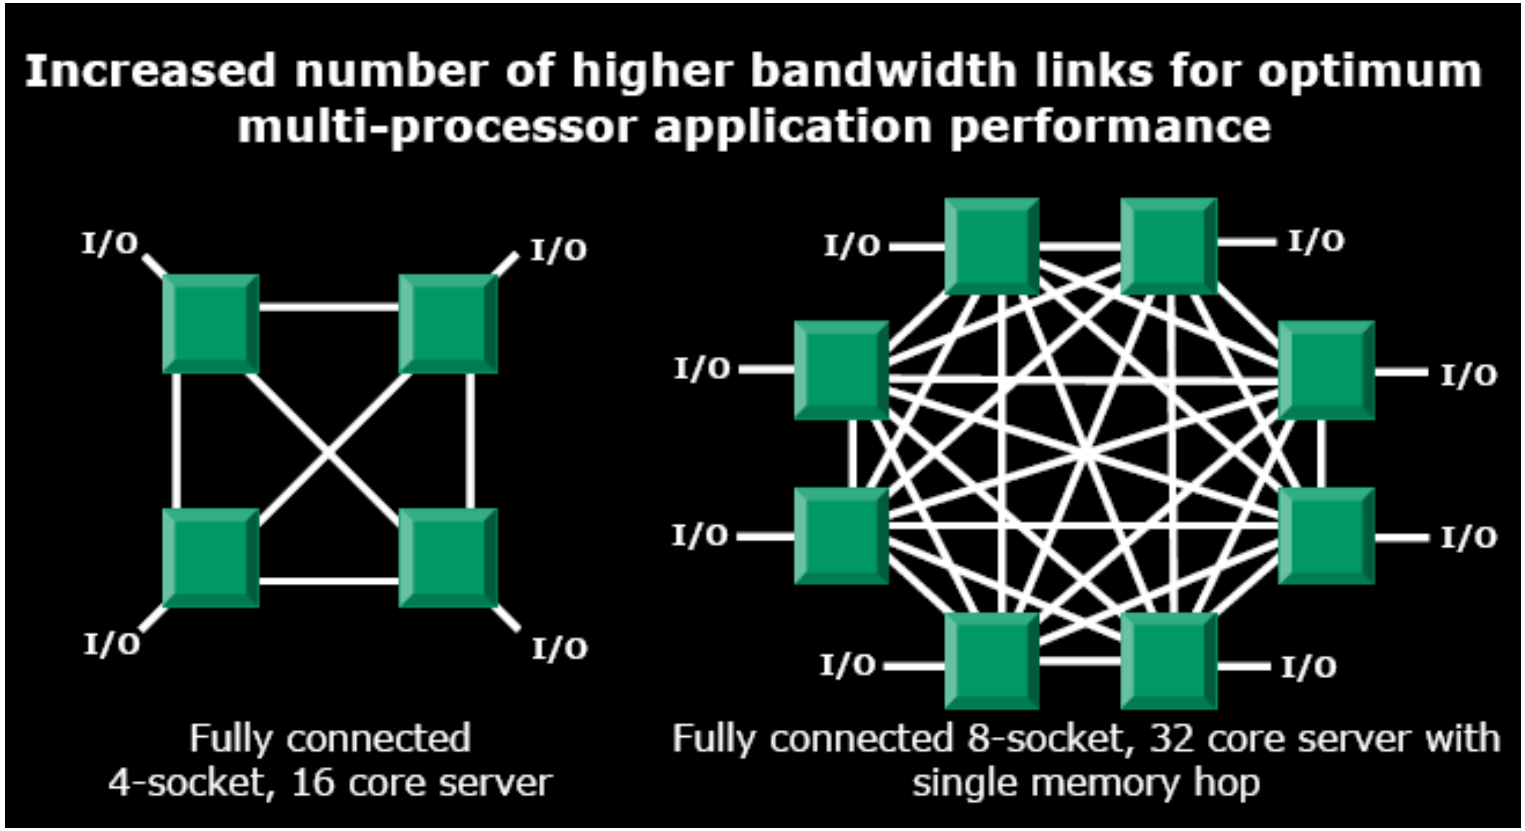

9.2 Main enhancements of the K10.5 Magny-Course servers (7)

b) AMD V 2.0 [89]

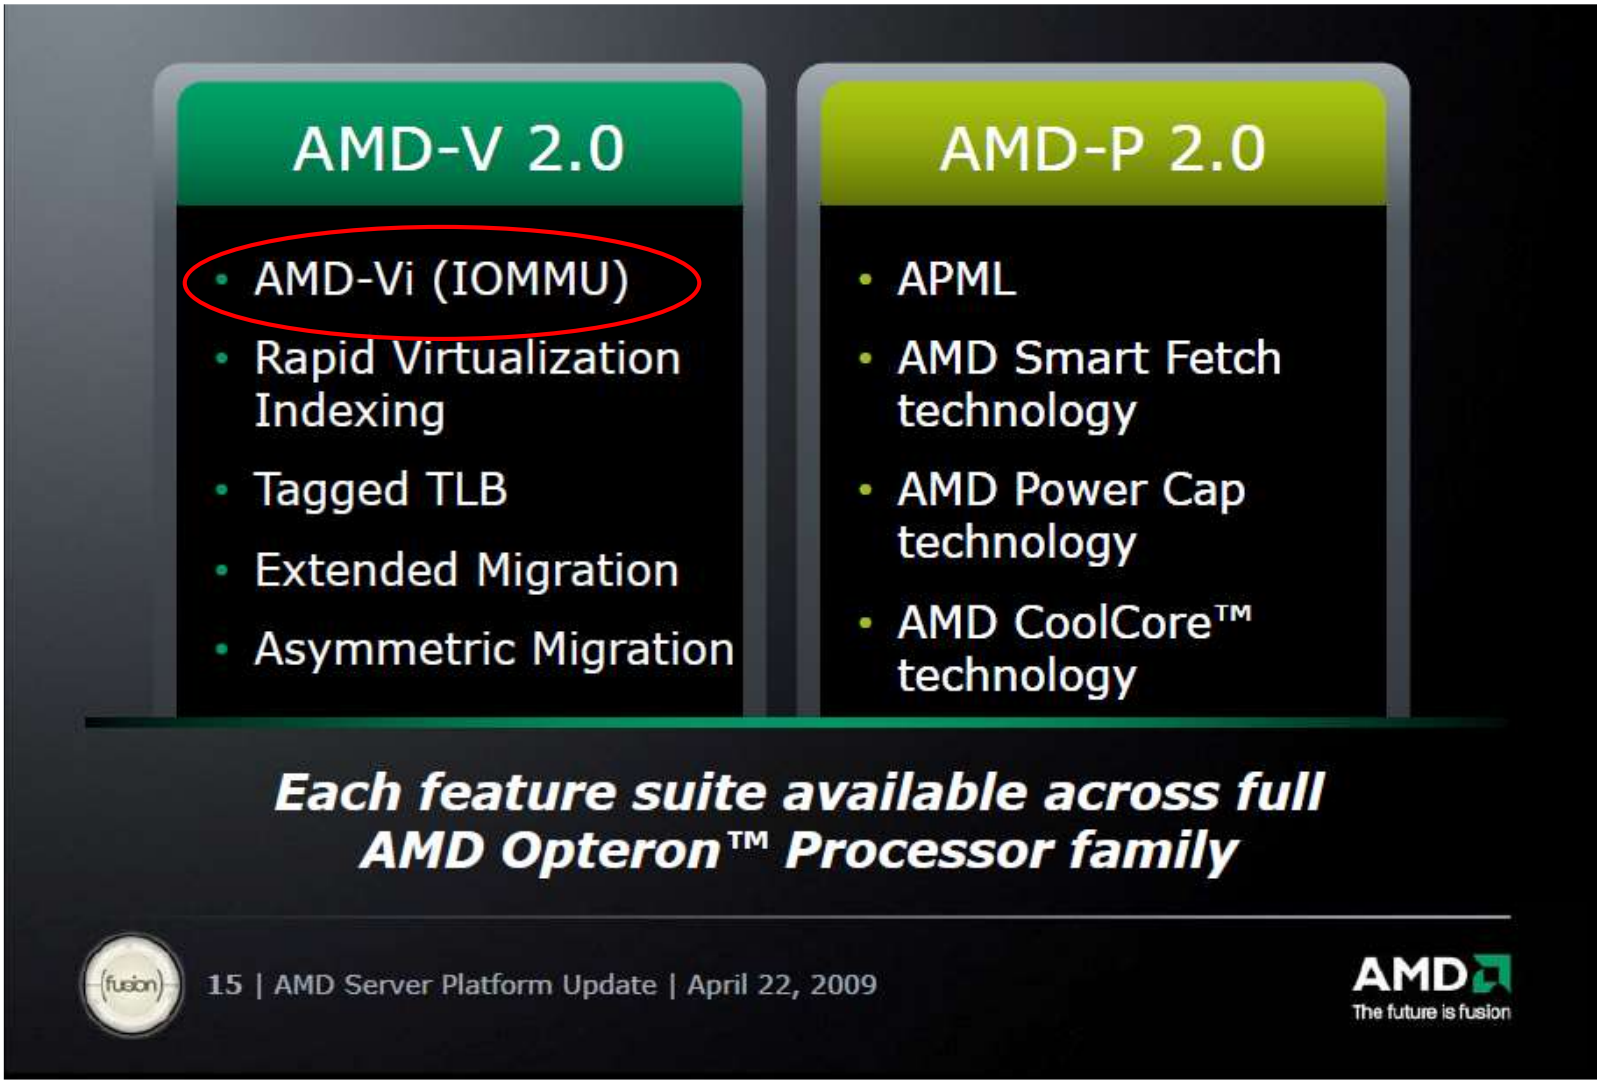

9.2 Main enhancements of the K10.5 Magny-Course servers (8)

AMD-V2.0 major innovation: I/O Virtualization (IOMMU) [151]

Chipset feature providing speedup for performing I/O in virtual environments

Without IOMMU

A VM (Virtual Machine) guest is not allowed to directly access an I/O device
to prevent guests from programming a device such that other guest's memory will be corrupted.

$\Longrightarrow$ Hypervisors must emulate devices, which causes significant performance loss.

With IOMMU

I/O devices can be assigned directly to a given VM guest by limiting memory accesses of the device to addresses that belong to this guest.

9.2 Main enhancements of the K10.5 Magny-Course servers (9)

c) Power management techniques of K10.5 Magny-Course servers (based on [53])

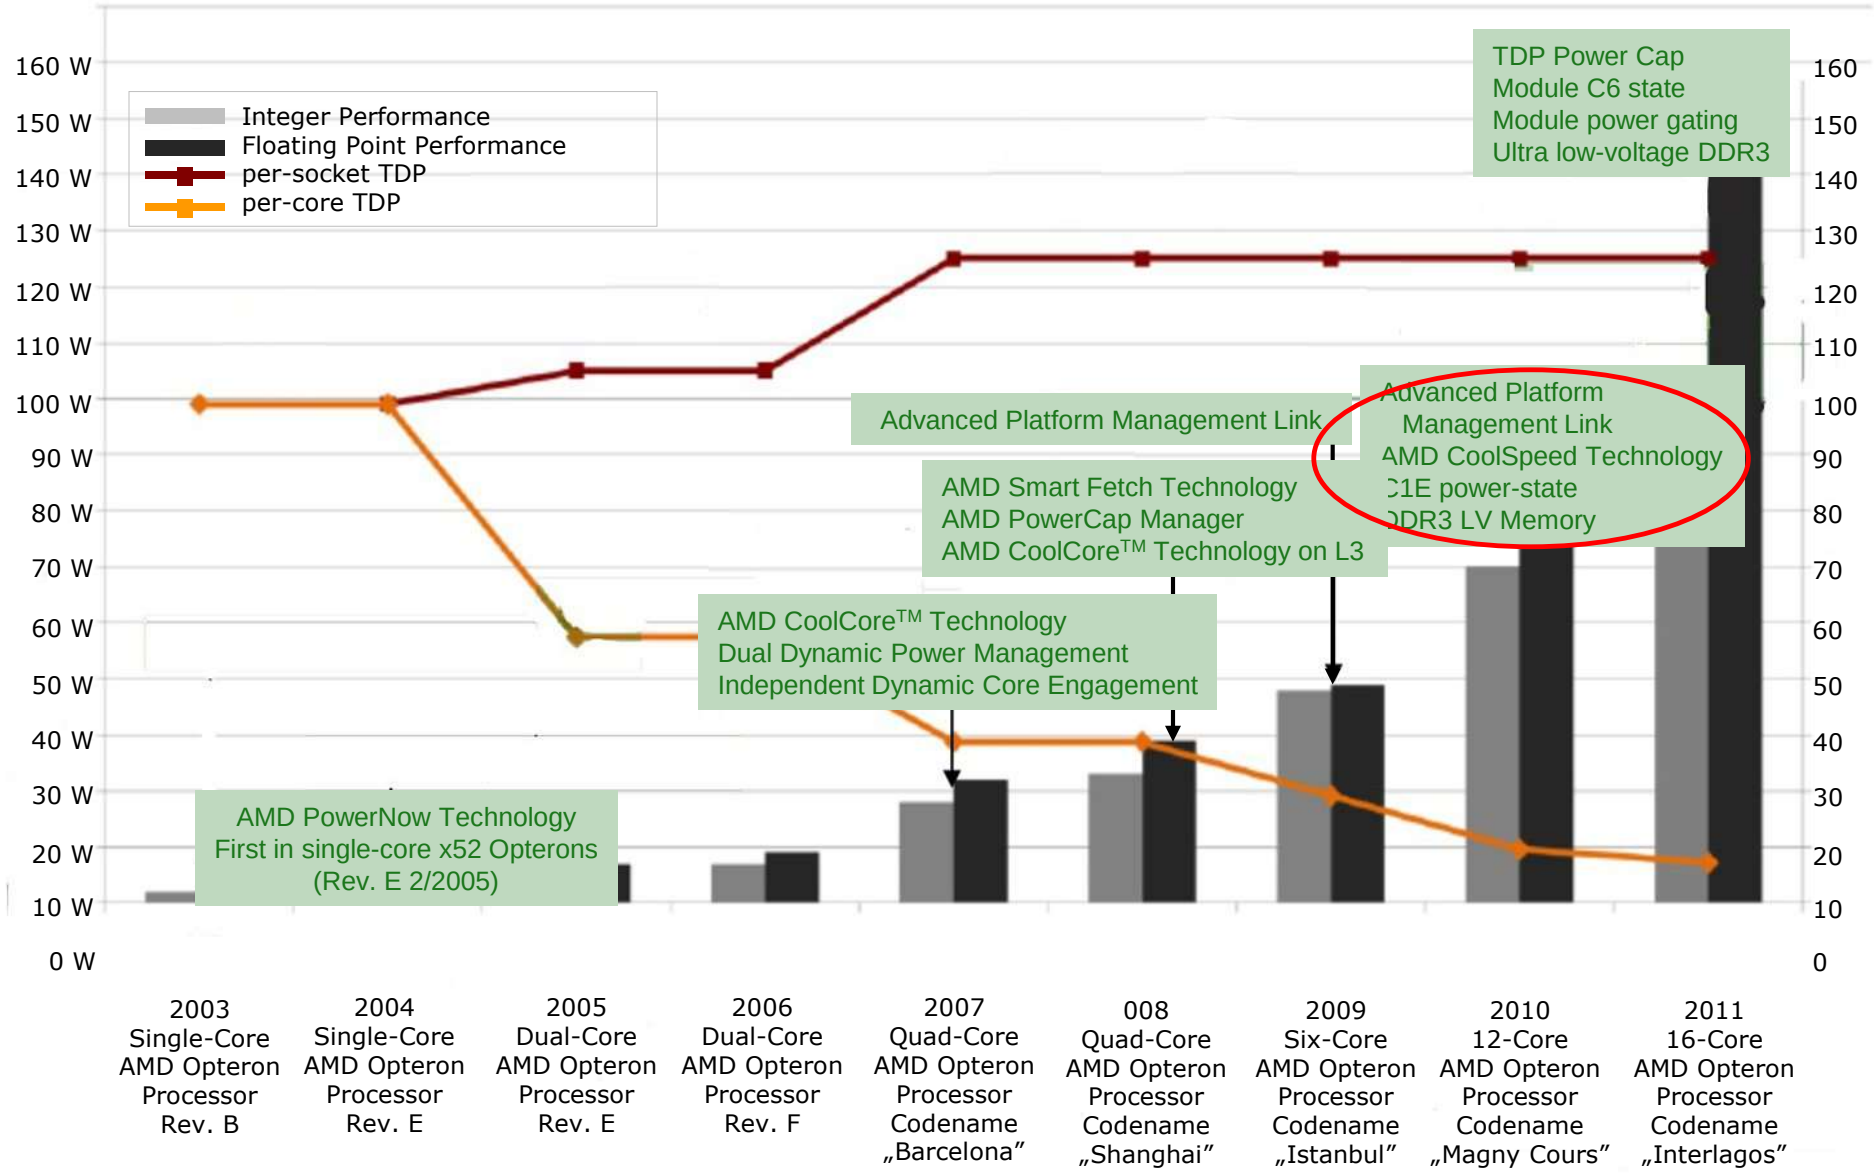

9.2 Main enhancements of the K10.5 Magny-Course servers (10)

Remark

The designation P2.0 power saving technology suit is merely a marketing term that covers a number of power saving technologies but does not identify the actual technologies introduced into Magny-Course servers, so we present these technologies based on the above Figure.

Overview of the AMD-P power saving technologies [97]

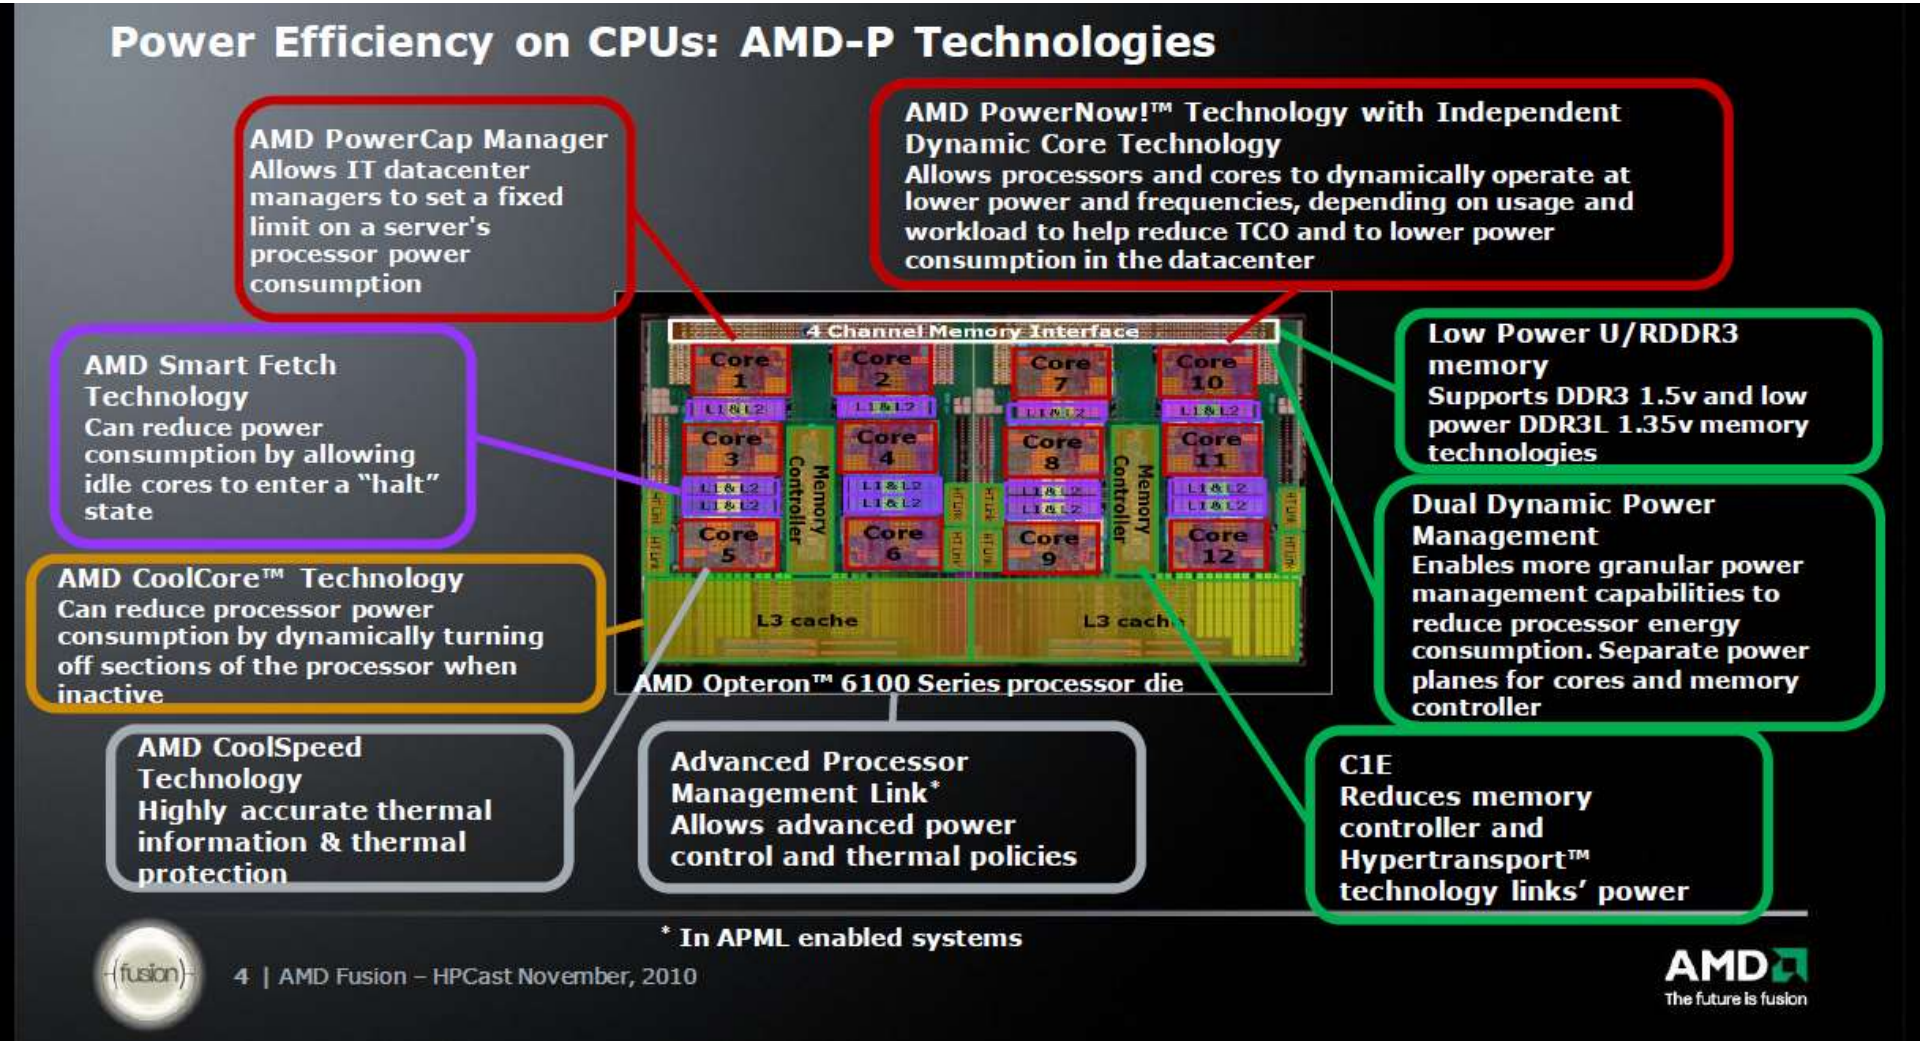

9.2 Main enhancements of the K10.5 Magny-Course servers (12)

New power saving features of AMD-P2 technology vs. AMD-P [96], [97]

c1) AMD CoolSpeed Technology

It provides highly accurate thermal information and thermal protection.

It reduces P-states when a temperature limit is reached to allow a server to operate.

if the processors thermal environment exceeds safe operational limits.

9.2 Main enhancements of the K10.5 Magny-Course servers (13)

c2) C1E state

- AMD introduced the C1E state first in their K10 Barcelona-based desktops, termed as the Phenom line in order to save power in times when all cores of a processor are inactive.
- In their servers AMD introduced the C1E state along with their Magny-Course line.
  This C1E state differs however, from the previously introduced C1E state.

Principle of operation [61]

- If all processors of a DP or MP server become idle for a longer period of inactivity the processor enters the C1E state.
- In the C1E state
  - all cores flush their L1 and L2 caches into the L3 cache,
  - the clocking of all cores will be shut down,
  - the HT link will be put into a lower power state (LS2),
  - the system memory will be placed into a low power state,
  - the L3 cache, the north bridge and the memory controller will be clocked at a low rate and to save power a lower alternative voltage (Altvid) may be applied to the CPU cores and the NB since then the static power will be lower.

    Separate voltages may be applied to the cores and the NB (in split power-plane mode).
- DMA events will wake up the processor from the C1E state.

9.2 Main enhancements of the K10.5 Magny-Course servers (14)

Areas covered by different power saving technologies [96]

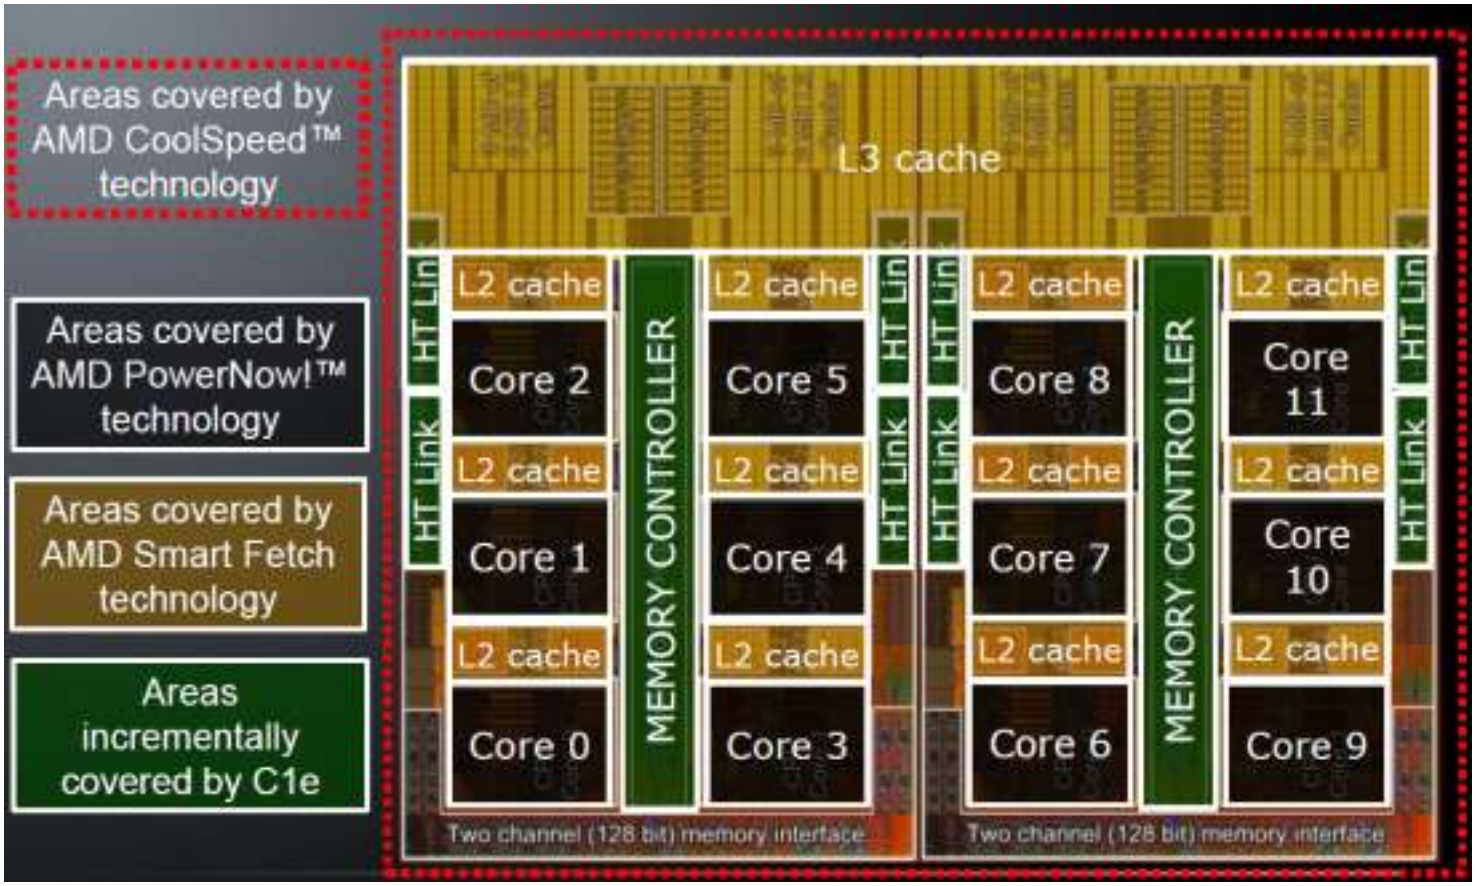

9.2 Main enhancements of the K10.5 Magny-Course servers (15)

c3) LV-DDR3 support [96]

It allows using 1.35 V LV-DDR3 DIMMs instead of 1.5 V regular DDR3 DIMMs.
Both unregistered and registered DIMMs are supported.

9.3 K10.5 Magny-Course-based server lines

9.3 K10.5 Magny-Course-based server lines

Main features of the Magny-Course MP server line [95]

Model Number	Core Count	Core Speed	ACP*	North Bridge†	1KU Pricing at intro.
6176 SE	12	2.3GHz	105W	1.8GHz	$1386
6174	12	2.2GHz	80W	1.8GHz	$1165
6172	12	2.1GHz	80W	1.8GHz	$989
6168	12	1.9GHz	80W	1.8GHz	$744
6136	8	2.4GHz	80W	1.8GHz	$744
6134	8	2.3GHz	80W	1.8GHz	$523
6128	8	2.0Ghz	80W	1.8GHz	$266
6164 HE	12	1.7GHz	65W	1.8GHz	$744
6128 HE	8	2.0GHz	65W	1.8GHz	$523
6124 HE	8	1.8GHz	65W	1.8GHz	$455

9.3 K10.5 Magny-Course-based server lines (2)

Performance increase of AMD's DP servers [146]

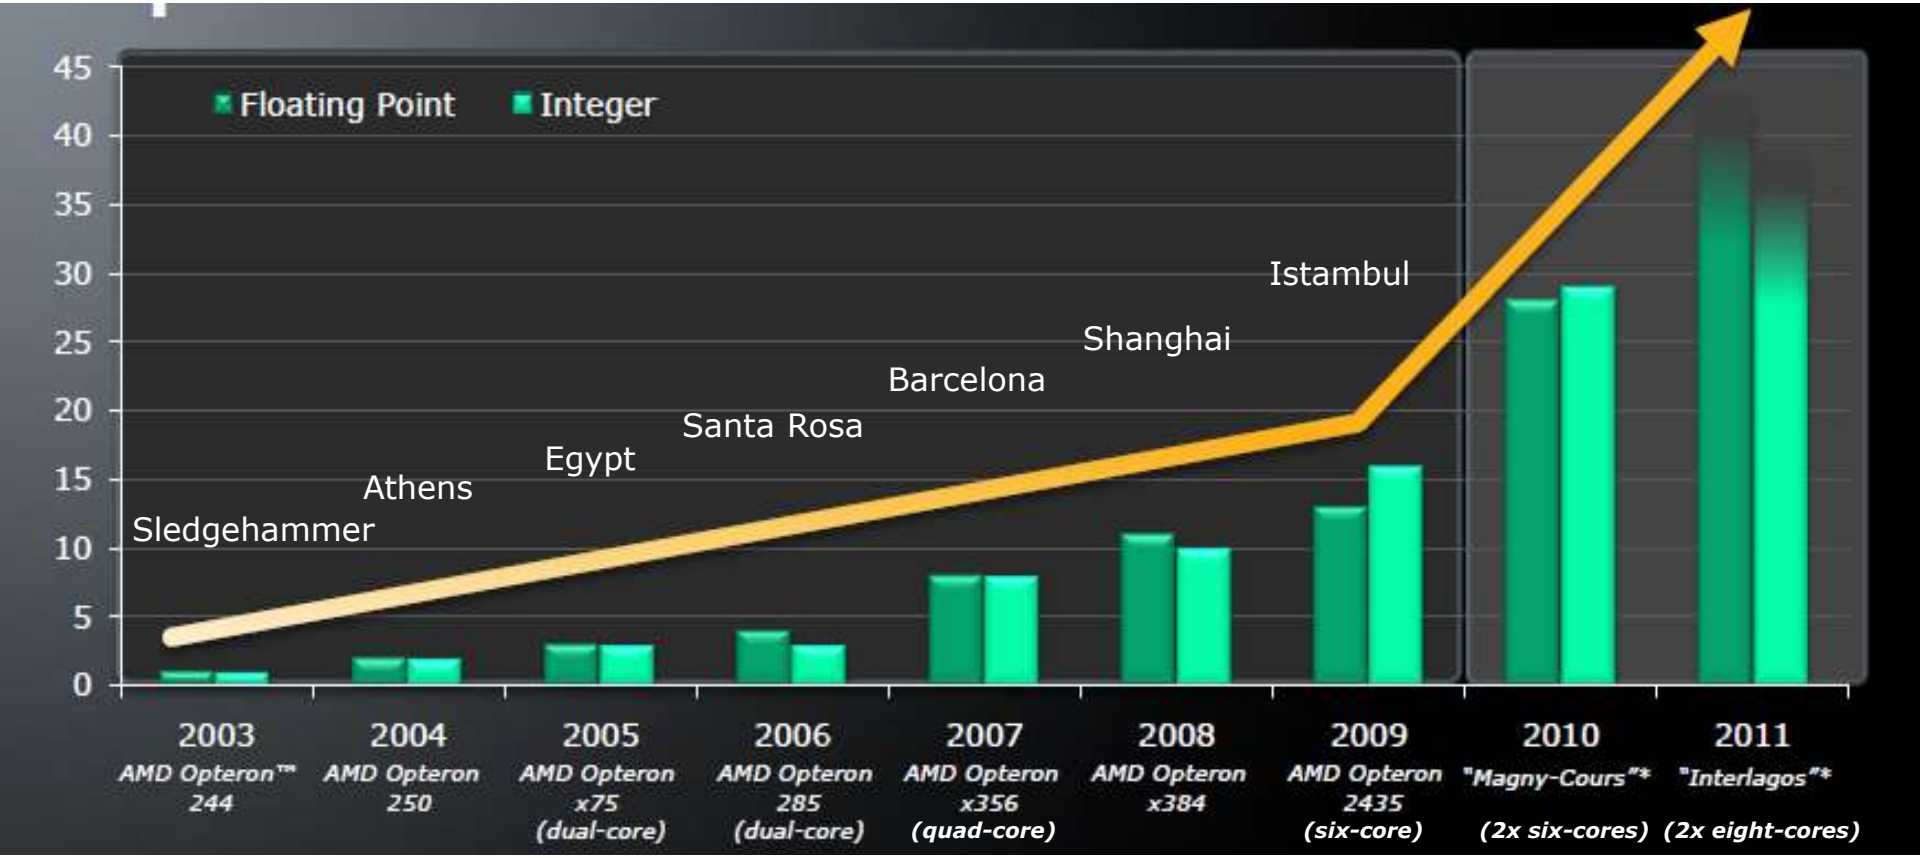

9.3 K10.5 Magny-Course-based server lines (3)

Note

If we relate the performance increase of the processors considered to the last single core Opteron (termed as Athens) instead of the first x86-64 server (Sledghammer), as done in the Figure, we can state that in the multi-core era

processor performance increases roughly linearly with the core count, as expected.

This means that in the multi-core era both the clock frequency and also the efficiency of the microarchitecture (IPC) remained roughly constant.

9.3 K10.5 Magny-Course-based server lines (4)

FP-performance, memory bandwidth and power consumption trends in AMD's Opteron family [17]

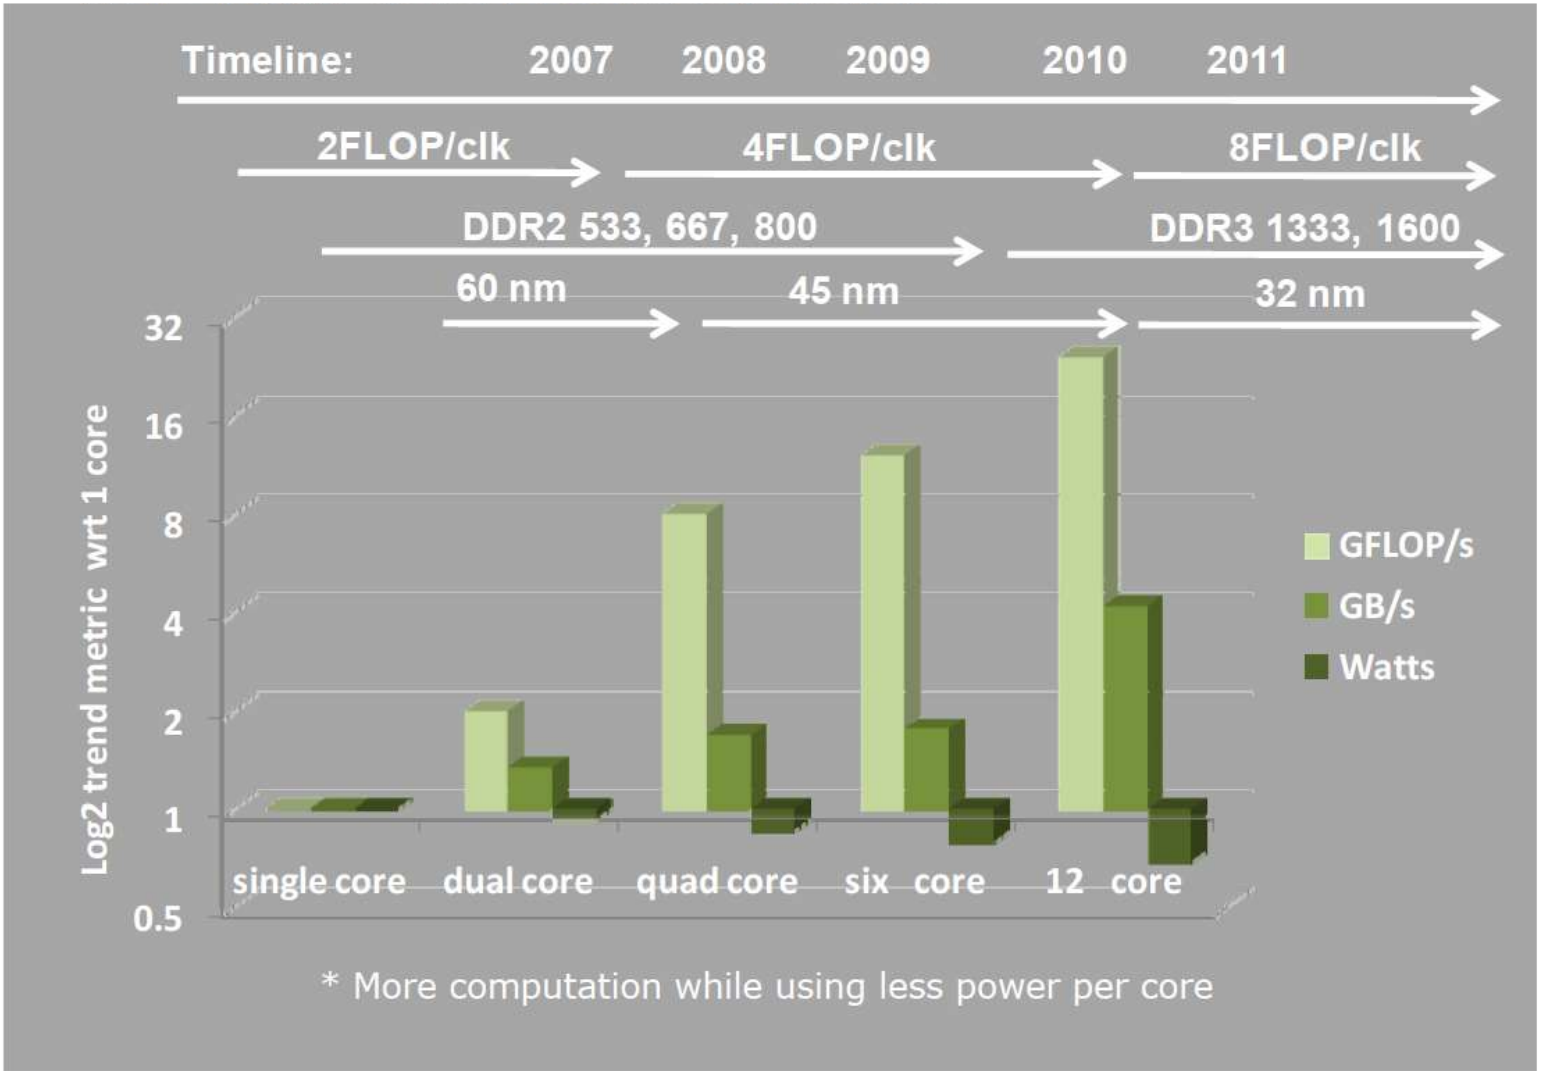

9.3 K10.5 Magny-Course-based server lines (5)

Changing the interpretation of "value" in AMD's processor lines [95]

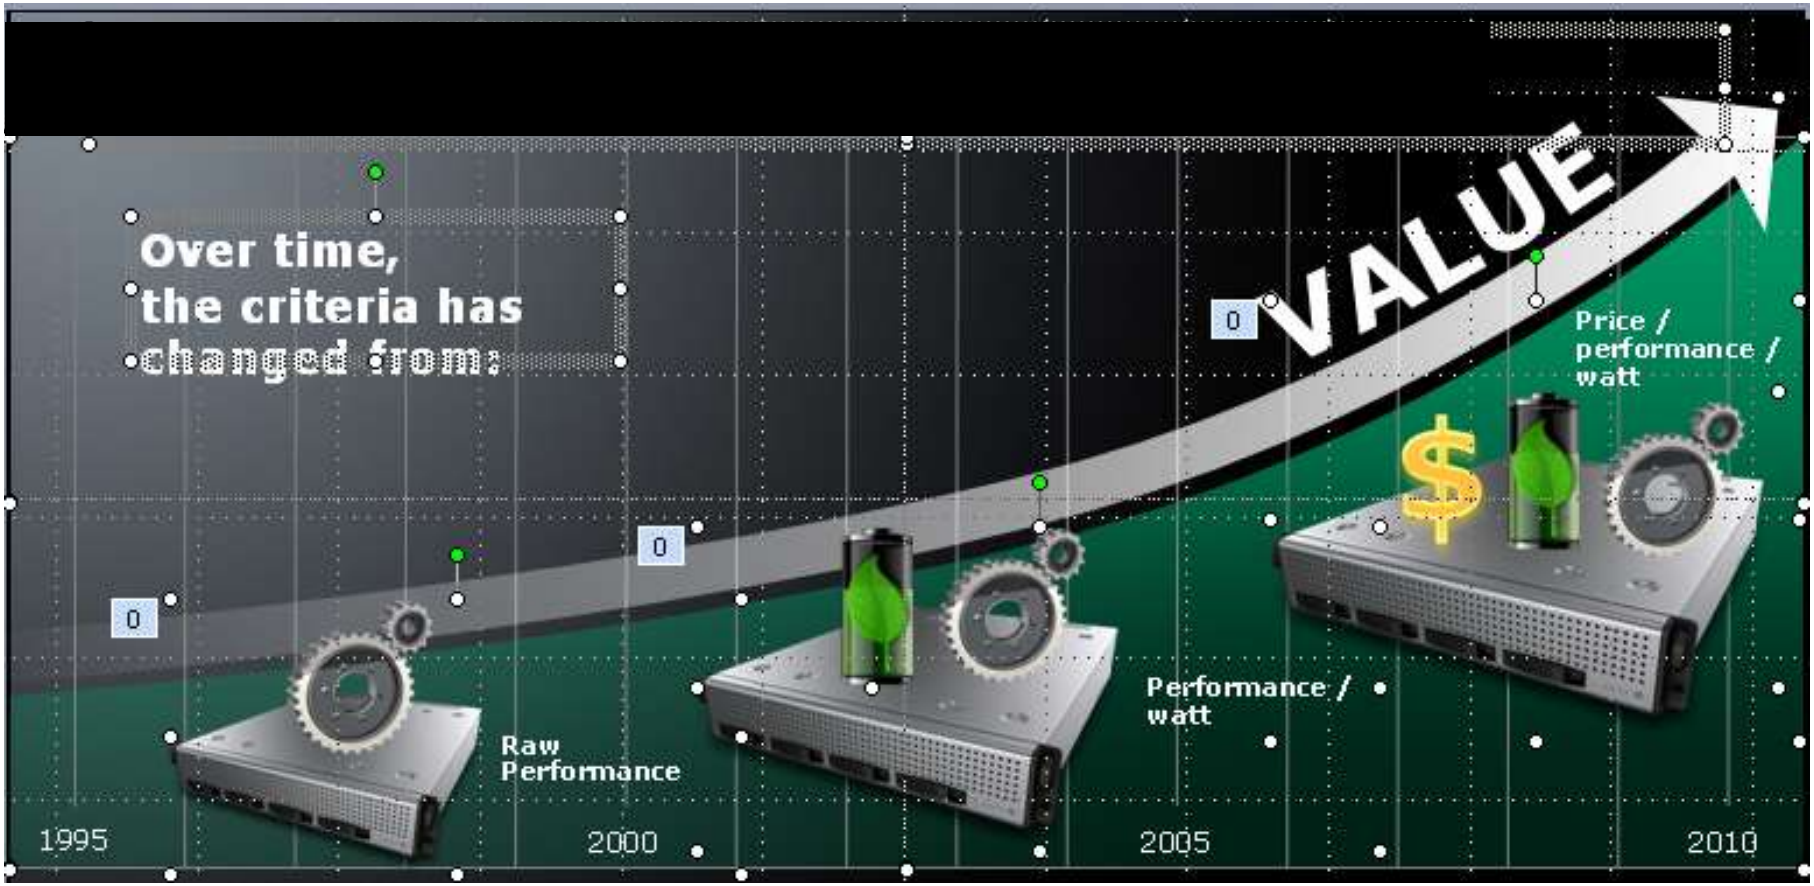

10. References

References (1)

[1]: Dognini R., SISSA Trieste AMD and SUN, Trieste, Nov. 22 2004

[2]: Underhill J., AMD64 with Direct Connect Architecture, A Solid Foundation for Dense Compute Environments, Aug. 13 2004, http://www.cfroundtable.org/hdcc/081104/DENSE%20COMPUTE%20ARCHITECTURES%20by%20Underhill.pdf

[3]: Wikipedia, AMD Am29000, http://en.wikipedia.org/wiki/AMD_Am29000

[4]: Hesseldahl A., Why Cool Chip Code Names Die, Forbes.com, July 6 2000 http://www.forbes.com/2000/07/06/mu2.html

[5]: Wikipedia, AMD Phenom, http://en.wikipedia.org/wiki/AMD_Phenom

[6]: Krazit T., AMD hopes for desktop PC boost with Spider, CNET.com, Nov. 18 2007, http://news.cnet.com/8301-13579_3-9819123-37.html#ixzz1mbCYk1zy

[7]: Wikipedia, List of AMD mobile microprocessors, http://en.wikipedia.org/wiki/List_of_AMD_mobile_microprocessors

[8]: Introducing AMD „Spider" Platform, Media Presentation, Nov. 19 2007, http://download.amd.com/Corporate/SpiderPlatformPresentationv3.pdf

[9]: Wikipedia, AMD K5, http://en.wikipedia.org/wiki/AMD_K5

[10]: Wikipedia, AMD K6, http://en.wikipedia.org/wiki/AMD_K6

[11]: Gavrichenkov I., AMD Athlon XP Processor Family Review, Xbitlabs.com, Nov. 7 2001, http://www.xbitlabs.com/articles/cpu/display/athlon-xp-1800.html

References (2)

[12]: Wikipedia, Athlon, http://en.wikipedia.org/wiki/Athlon

[13]: Goto H., AMD CPU, 2006, http://pc.watch.impress.co.jp/docs/2006/0313/kaigai05l.gif

[14]: Goto H., AMD CPU Transition, 2011, http://pc.watch.impress.co.jp/video/pcw/docs/473/823/p7.pdf

[15]: Valich T., AMD Opteron, Phenom codenames explained, Bright side of news, May 1 2009, http://www.brightsideofnews.com/news/2009/5/1/amd-opteron2c-phenom-codenames-explained.aspx

[16]: Conway P., Kalyanasundharam N., Donley G., Lepak K., Hughes B., Blade computing with the AMD Opteron Processor („Magny-Cours"), Aug. 2009, http://hotchips.org/uploads/hc21/2_mon/HC21.24.100.ServerSystemsI-Epub/HC21.24.110.Conway-AMD-Magny-Cours.pdf

[17]: Waldecker B., Conway P., AMD Opteron Multicore Processors, Febr. 1 2009, http://www.nersc.gov/assets/Uploads/AMDMultiCoreCrayNersc020110.pdf

[18]: The AMD x86-64 Architecture Programmers Overview, Aug. 2000, http://www.weblearn.hs-bremen.de/risse/RST/docs/AMD/x86amd64.pdf

[19]: Crawford J., Introducing the Itanium Processors, IEEE, Sept.-Oct. 2000, http://www.cs.virginia.edu/~gjp5j/cs854/m5009.pdf

[20]: Gwennap L., Intel, HP Make EPIC Disclosure, Microprocessor Report, Vol. 11. No. 14., Oct. 27 1997, http://www.ele.uva.es/~jesman/BigSeti/ftp/Cajon_Desastre/MPR/epic.pdf

[21]: Mulder H., Huck J., Yu A., Announcing the IA-64 Architecture, http://www.weblearn.hs-bremen.de/risse/RST/docs/Intel/techpres.pdf

References (3)

[22]: Petrtylová B., Kubes T., Intel Itanium IA64 – Merced, April 24 2003, http://www.tomaskubes.net/CVUT/download/36aps_ia64_merced_aj.ppt

[23]: Cornelius H., Intel Itanium Architecture, Jan. 28 2003, http://www.rrze.de/dienste/arbeiten-rechnen/hpc/vortraege/IntelCornelius.pdf

[24]: Intel, HP Reveal IA-64 Instruction Set Architecture, Business Wire, May 26 1999, http://www.thefreelibrary.com/Intel,+HP+Reveal+IA-64+Instruction+Set+Architecture.-a054726825

[25]: AMD Releases x86-64™ Architectural Specification; Enables Market Driven Migration to 64-Bit Computing, Aug. 10 2000 hhttp://www.amd.com/us/press-releases/Pages/Press_Release_715.aspx

[26]: AMD Discloses New Technologies At Microporcessor Forum, Oct. 5 1999 http://www.amd.com/us/press-releases/Pages/Press_Release_751.aspx

[27]: AMD and Microsoft Collaborate to further 64-Bit Computing, April 24 2002 http://www.amd.com/us/press-releases/Pages/Press_Release_19906.aspx

[28]: Weber F., AMD's Next Generation Microprocessor Architecture, Oct. 2001, http://www.datasheetcatalog.org/datasheet/AdvancedMicroDevices/mXsutyv.pdf

[29]: Wikipedia, Itanium, http://en.wikipedia.org/wiki/Itanium

[30]: x86 64-bit support for Windows, Ars Technica Forum http://arstechnica.com/civis/viewtopic.php?f=2&t=1027

[31]: Kerner M., Padgett N., A History of Modern 64-bit Computing, Febr. 2007, http://www.cs.washington.edu/education/courses/csep590/06au/projects/history-64-bit.pd

References (4)

[32]: Magee M., Intel burns AMD-clone Yamhill idea, The Inquirer, Sept. 30 2002, http://www.theinquirer.net/inquirer/news/1047441/intel-burns-amd-clone-yamhill-idea

[33]: de Vries H., Looking at Intel's Prescott die, Chip Architect, April 20 2003, http://chip-architect.com/news/2003_04_20_Looking_at_Intels_Prescott_part2.html# Yamhill shines out of the "blue"

[34]: Intel confirms existence of X86-64 Yamhill chip, The Inquirer, Dec. 24 2003, http://www.theinquirer.net/inquirer/news/1032095/intel-confirms-existence-of-x86-64-yamhill-chip

[35]: Perich D., Intel Volume platforms Technology Leadership, Presentation at HP World 2004, http://98.190.245.141:8080/Proceed/HPW04CD/papers/4194.pdf

[36]: Fried I., Microsoft ending Itanium support, CNET.com, April 5 2010, http://news.cnet.com/8301-13860_3-20001746-56.html

[37]: Builder's Guide for AMD Opteron Processor-Based Servers and Workstations, Febr. 2004 http://support.amd.com/us/Processor_TechDocs/30925.pdf

[38]: Wikipedia, Fichier:AMD A64 Opteron arch.svg, http://fr.wikipedia.org/wiki/Fichier:AMD_A64_Opteron_arch.svg

[39]: Cook H., Sackey K., Weatherton A., The AMD Opteron

[40]: Torres G., Inside AMD K10 Architecture, Hardware Sectrets, Sept. 3 2007, http://www.hardwaresecrets.com/article/Inside-AMD-K10-Architecture/480/3

[41]: Sander B., AMD Microprocessor Technologies, 2006, http://www.ewh.ieee.org/r4/chicago/foxvalley/IEEE_AMD_Meeting.ppt

References (5)

[42]: Nordlund L., The Quad-Core AMD Opteron™ Processor, July 19 2006

[43]: Walrath J., AMD Athlon X2 3800+ and Athlon 3800+, Same Numbers, Different Results http://penstarsys.com/reviews/cpu/amd/x2_3800_939/x2_3800_3.htm

[44]: Romanchenko V., Sofronov D., Debut of AMD AM2: the long-awaited DDR2 on AMD Athlon X2, May 23 2006, http://www.digital-daily.com/cpu/amd_am2_4000/

[45]: PC Watch, May 2006, http://pc.watch.impress.co.jp/docs/2006/0503/kaigai267.htm

[46]: Shilov A., AMD Quad-Core Opteron, Athlon 64 Processors Details Leak, Xbit laboratories, May 3 2006, http://www.xbitlabs.com/news/cpu/display/20060503150902.html

[47]: Chevanne H., AMD Opteron processors scalability and Roadmap, March 17 2010, http://www.hpcadvisorycouncil.com/events/switzerland_workshop/pdf/Presentations/Day%203/9_AMD.pdf

[48]: AMD WW HPC 07, 2007, http://cisl.ucar.edu/dir/CAS2K7/Presentations/torricelli.pdf

[49]: Kubicki K., AMD Announces More K8L Details, Daily Tech, June 1 2006, http://www.dailytech.com/AMD+Announces+More+K8L+Details/article2637.htm

[50]: Larger L3 cache in Shanghai, Nov. 13 2008, http://blogs.amd.com/developer/2008/11/13/larger-l3-cache-in-shanghai-part-i/

[51]: Wikipedia, File:AMD K10 Arch.svg, http://en.wikipedia.org/wiki/File:AMD_K10_Arch.svg

References (6)

[52]: Kanter D., "AMD's K8L and 4x4 Preview, Real World Tech. June 02 2006, http://www.realworldtech.com/page.cfm?ArticleID=RWT060206035626&p=1

[53]: 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

[54]: Boggs J., Microsoft PDC 2008, Oct. 11 2008

[55]: Dorsey J., Searles S., Ciraula M., Johnson S., Bujanos N., Wu D., Bragaza M., Meyers S., Fang E., Kumar R., An Integrated Quad-Core Opteron Processor, ISSCC 2007

[56]: Schmid P., Phenom Models and Details, Continued, Tom's Hardware, Dec. 19 2007, http://www.tomshardware.com/reviews/amd-phenom-athlon-64-x2,1746-3.html

[57]: Henning W., AMD Phenom 9600 Review, Neoseeker, Nov. 29 2007, http://www.neoseeker.com/Articles/Hardware/Reviews/phenom_9600/

[58]: Wikipedia, Hyper Transport, http://en.wikipedia.org/wiki/HyperTransport

[59]: AMD Turion™ X2 Ultra Dual-Core Mobile Processors and AMD Turion™ X2 Dual-Core Mobile Processors Key Architecture Features, http://www.amd.com/us/infrastructure/processors/turion-x2/Pages/turion-x2-mobile-features.aspx

[60]: Owen J., Next-Generation Mobile Computing: Balancing Performance and Power Efficiency, Hot Chips 19, 2007, http://hotchips.org/uploads/hc19/3_Tues/HC19.08/HC19.08.02.pdf

[61]: Cool'n'Quiet 2.0 In Detail II, Tom's Hardware, Nov. 19 2007, http://www.tomshardware.com/reviews/spider-weaves-web,1728-13.html

References (7)

[62]: Naveh A., Rotem E., Mendelson A., Gochman S., Chabukswar R., Krishnan K., Kumar A., Power and thermal management in the Intel® Core™ Duo processor, May 15 2006, http://www.intel.com/technology/itj/2006/volume10issue02/art03_Power_and_Thermal_Management/p01_abstract.htm

[63]: Bailey A., Barcelona's Innovative Architecture Is Driven by a New Shared Cache, Aug. 14 2007, http://developer.amd.com/documentation/articles/pages/8142007173.aspx

[64]: Larger L3 cache in Shanghai, Nov. 13 2008, AMD, http://forums.amd.com/devblog/blogpost.cfm?threadid=103010&catid=271

[65]: Sahl H., Energieeffizienz aus sich des Prozessorhersteller, Jan. 2009, http://www.it-business.de/inhalte/whitepaper/downloads/10087/

[66]: De Gelas J., Dynamic Power Management: A Quantitative Approach, AnandTech, Jan. 18 2010, http://www.anandtech.com/show/2919/4

[67]: AMD Istanbul Launch: Shipping Today, Solori.net, June 1 2009, http://blog.solori.net/2009/06/01/amd-istanbul-launch-shipping-today/

[68]: Brandão R., Tecnologia – AMD Brasil, AMD Athlon II, Sept. 23 2009, http://www.slideshare.net/rfbrandao/web-seminario-athlon-ii

[69]: AMD Phenom II X4 975 BE 3.60 GHz, Tech Power Up, Jan. 5 2011, http://www.techpowerup.com/reviews/AMD/Phenom_II_X4_975/

[70]: AMD Phenom II X4 840 3.20 GHz, Tech Power Up, Jan. 5 2011, http://www.techpowerup.com/reviews/AMD/Phenom_II_X4_840/

References (8)

[71]: Gavrichenkov I., AMD Phenom II X2 550 and AMD Athlon II X2 250 Processors Review, Xbitlabs.com, June 1 2009, http://www.xbitlabs.com/articles/cpu/display/phenom-athlon-ii-x2_3.html

[72]: Enderle R., AMD Shanghai "We are back!", TGDaily, Nov. 13, 2008, http://www.tgdaily.com/content/view/40176/128/

[73]: Phenom II X2 és Athlon II X2, Pro Hardver, June 24 2009, http://prohardver.hu/teszt/phenom_ii_x2_es_athlon_ii_x2/vegre_igazi_ketmagos_az_amd-tol.html

[74]: Wasson S., AMD's FX-8150 'Bulldozer' processor, Tech Report, Oct. 12 2011, http://techreport.com/articles.x/21813/3

[75]: Angelini C., Power Management, Tom's Hardware, Oct. 12 2011, http://www.tomshardware.com/reviews/fx-8150-zambezi-bulldozer-990fx,3043-7.html

[76]: AMD Champlain core, CPU World, http://www.cpu-world.com/Cores/Champlain.html

[77]: Walrath J., AMD Introduces New Mainstream and Ultra-Portable Platforms, PC Perspective, Sept. 12 2009, http://www.pcper.com/reviews/Processors/AMD-Introduces-New-Mainstream-and-Ultra-Portable-Platforms/2009-Mainstream-Platfo

[78]: K10: Barcelona, Shanghai, Quad-Core Opteron, Phenom, AMD Zone, http://www.amdzone.com/phpbb3/viewtopic.php?f=52&t=137000

[79]: Wasson S., AMD's 'Istanbul' six-core Opteron processors, Tech Report, June 1 2009, http://techreport.com/articles.x/17005

References (9)

[80]: Introducing Six-Core AMD Opteron Processors, Codename „Istanbul”, June 1 2009, http://i.zdnet.com/blogs/istanbul_nprp_preso_legally_approved.pdf

[81]: AMD Opteron System Architecture http://www.qdpma.com/SystemArchitecture/SystemArchitecture_Opteron.html

[82]: Conway P., Kalyanasundharam N., Donley G., Lepak K., Hughes B., Cache Hierarchy and Memory Subsystem of the AMD Opteron Processor, IEEE Micro, March/April 2010, http://portal.nersc.gov/project/training/files/XE6-feb-2011/Architecture/Opteron-Memory-Cache.pdf

[83]: Advanced Platform Management Link (APML) Specification, Aug. 25 2009, http://support.amd.com/us/Processor_TechDocs/41918.pdf

[84]: Advanced Platform Management Link (APML) Tools, http://developer.amd.com/tools/apml/Pages/default.aspx

[85]: De Gelas J., AMD's Six-Core Opteron 2435, Anand Tech, June 1 2009, http://www.anandtech.com/show/2774/2

[86]: Using ACPI to Report APML P-State Limit Changes to Operating Systems and VMM’s, AMD, Aug. 7 2009, http://developer.amd.com/Assets/ACPI-APML-PState-rev12.pdf

[87]: Advanced Configuration and Power Interface (ACPI) Specification, Revision 5.0, Dec. 6 2011, http://www.acpi.info/spec.htm

[88]: Kowaliski C., AMD cooks up its own Opteron chipsets for 2009, Tech Report, Sept. 8 2008, http://techreport.com/discussions.x/15474

References (10)

[89]: AMD Server Platform Update, April 22 2009, http://issuu.com/lemmingoverlord/docs/amd_opteron_6th_anniversary_press_presentation

[90]: Six-Core AMD Opteron Processors, 2009, http://www.sgi.com/company_info/acceleratingresults/six_core_compare.pdf

[91]: Shimpi A. L., AMD's Six-Core Phenom II X6 1090T & 1055T Reviewed, Anand Tech, April 27 2010, http://www.anandtech.com/show/3674

[92]: Walrath J., AMD's Turbo Core Technology, PC Perspective, April 11 2010, http://www.pcper.com/reviews/Processors/AMDs-Turbo-Core-Technology

[93]: AMD Turbo CORE Technology, On Select AMD Phenom II Processors, March 2010, http://www.amd-news.com/assets/files/amd-cn/cn_2010-20_amd-turbo-core-technology.p

[94]: Elrajtoltak a 12 magos Opteronok, Pro Hardver, March 29 2010, http://prohardver.hu/hir/elrajtoltak_12_magos_opteronok_6100_amd.html

[95]: The AMD Opteron™ 6000 Series Platform: More Cores, More Memory, Better Value, Slide Share, March 26 2010, http://www.slideshare.net/AMDUnprocessed/amd-opteron-6000-series-platform-press-presentation-final-3564470

[96]: De Gelas J., AMD's 12-core "Magny-Cours" Opteron 6174 vs. Intel's 6-core Xeon, Anand Tech, March 29 2010, http://www.anandtech.com/show/2978/amd-s-12-core-magny-cours-opteron-6174-vs-intel-s-6-core-xeon/14

[97]: Nordlund L., Driving HPC Performance Efficiency with Heterogeneous Computing, June 19 2011, http://www.hpcadvisorycouncil.com/events/2011/european_workshop/pdf/7_amd_sponsor.pdf

References (11)

[98]: Wikipedia, Turion, http://en.wikipedia.org/wiki/Griffin_(processor)#Turion_X2_Ultra

[99]: AMD's CPU Roadmap, 2008-2011, Firing Squad, http://www.firingsquad.com/hardware/amd_cpu_roadmap_update_2008/

[100]: Sandhu T., AMD's Puma platform set to kick Centrino into touch?, Hexus, June 4 2008, http://hexus.net/tech/features/laptop/13529-amd039s-puma-platform-set-kick-centrino-touch/?page=3

[101]: AMD Sempron 200U and 210U Processors for Embedded Applications, 2008, http://www.amd.com/us/Documents/45626B_Sempron_BGA_brief.pdf

[102]: Introducing AMD Turion Neo X2 and AMD Athlon Neo X2 Dual-Core ASB1 (BGA) Processors for Embedded Applications, 2009, http://www.amd.com/us/Documents/47413A_Dual_Core_BGA_brief_PDF.pdf

[103]: AMD A-Series APU, EMEA Press Call, June 7 2011, http://img.zwame.pt/nemesis11/Amd_A_series/AMD.pdf

[104]: Wikipedia, AMD Vision, http://en.wikipedia.org/wiki/AMD_Vision

[105]: Wikipedia, List of AMD Fusion microprocessors, http://en.wikipedia.org/wiki/List_of_AMD_Fusion_microprocessors

[106]: 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 (12)

[107]: Shimpi A. L., The AMD A8-3850 Review: Llano on the Desktop, Anand Tech, June 30 2011, http://www.anandtech.com/show/4476/amd-a83850-review

[108]: Jotwani R., Sundaram S., Kosonocky S., Schaefer A., Andrade V. F., Novak A., Naffziger S., An x86-64 Core in 32 nm SOI CMOS, IEEE Xplore, 2010, http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=05624589

[109]: Altavilla D., AMD Fusion: A8-3500M A-Series Llano APU Review, Hot Hardware, June 14 2011, http://hothardware.com/Reviews/AMD-Fusion-A83500M-ASeries-Llano-APU-Review/?page=2

[110]: A Nagy AMD Llano APU Megateszt, Pro Hardver, Aug. 1 2011, http://prohardver.hu/teszt/amd_llano_apu_megateszt/hammertol_huskyig.html

[111]: Chiappetta M., AMD A8-3850 Llano APU and Lynx Platform Preview, Hot Hardware, June 30 2011, http://hothardware.com/Reviews/AMD-A83850-Llano-APU-and-Lynx-Platform-Preview/

[112]: Walton J., Shimpi A. L., The AMD Llano Notebook Review: Competing in the Mobile Market, Anand Tech, June 14 2011, http://www.anandtech.com/show/4444/amd-llano-notebook-review-a-series-fusion-apu-a8-3500m/4

[113]: Naffziger S. D., Sampling chip activity for real time power estimation, Patent Genius Aug. 30 2011, http://www.patentgenius.com/patent/8010824.html

[114]: Silcott G., AMD Talks-Up "Llano" x86 Innovation at ISSCC, Febr. 8 2010, http://blogs.amd.com/fusion/2010/02/08/amd-talks-llano-x86-innovation-isscc/

References (13)

[115]: Shimpi A. L., AMD Reveals More Llano Details at ISSCC: 32nm, Power Gating, 4-cores, Turbo?, Anand Tech, Febr. 8 2010, http://www.anandtech.com/show/2933

[116]: Kosonocky S., Practical Power Gating and Dynamic Voltage/Frequency Scaling, Aug. 17 2011, http://hotchips.org/uploads/hc23/HC23.17.1-tutorial1/HC23.17.111.Practical_PGandDV-Kosonocky-AMD.pdf

[117]: Branover A., Foley D., Steinman M., AMD Fusion APU: Llano, IEEE Micro, 2012, http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=06138843

[118]: Kosonocky S., Patent Application Publication, No. US 2011/0186930 A1, Aug. 4 2011

[119]: 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

[120]: AMD to Introduce Three New Bulldozer-based APUs in 2012, Softpedia, http://news.softpedia.com/newsImage/AMD-to-Introduce-Three-New-Bulldozer-based-APUs-in-2012-6.jpg/

[121]: Hugosson J., AMD's mobile roadmap up till 2013, early launch of Trinity confirmed, Nordic Hardware, Sept. 21 2011, http://www.nordichardware.com/news/69-cpu-chipset/44214-amds-mobile-roadmap-up-till-2013-early-launch-of-trinity-confirmed.html

[122]: 2010 Corporate Responsibility Report, Building Momentum, AMD, Aug. 9 2011, http://www.amd.com/us/Documents/2010_CRR.pdf

References (14)

[123]: AMD64 Technology Wins Best Of Show At 2004 Teched, June 1 2004, http://www.amd.com/us/press-releases/Pages/Press_Release_85830.aspx

[124]: Schafer S., Unleash the Hounds with AMD's Upcoming Quad-core Processors, June 28 2006, http://developer.amd.com/documentation/articles/pages/628200631.aspx

[125]: Gochman S., Mendelson A., Naveh A., Rotem E., Introduction to Intel® Core™ Duo processor architecture, May 15 2006, http://www.intel.com/technology/itj/2006/volume10issue02/art01_Intro_to_Core_Duo/p01_abstract.htm

[126]: Bennett K., AMD's Griffin Processor & Puma Mobile Platform, Hardocp, May 18 2007, http://www.hardocp.com/article/2007/05/18/amds_griffin_processor_puma_mobile_platform/

[127]: AMD Quad-core Press Briefing – First Quarter Update

[128]: AMD Desktop Mainstream and Performance Platform Road-Map, July 30 2007, http://xtreview.com/images-added.php?image=images/amd-future-platform-pc1.gif&id=3009

[129]: Quad-Core AMD Opteron Processor Code-named „Shanghai" Overview, Oct. 2008

[130]: Crothers B., AMD revisits Puma mobile technology, again, CNet, March 4 2008, http://news.cnet.com/8301-13924_3-9885550-64.html

[131]: Boggs J., AMD CPU Roadmap, July 2007, http://developer.amd.com/wordpress/media/2012/10/Develop_Brighton_Justin_Boggs-1.pdf

[132]: Microsoft is Committed to AMD?s x86-64, Extreme Tech, April 25 2002, http://discuss.extremetech.com/forums/thread/127612009.aspx

References (15)

[133]: Wikipedia, List of AMD Opteron microprocessors, http://en.wikipedia.org/wiki/List_of_AMD_Opteron_microprocessors

[134]: Wikipedia, List of AMD Athlon 64 microprocessors, http://en.wikipedia.org/wiki/List_of_AMD_Athlon_64_microprocessors

[135]: BIOS and Kernel Developer's Guide (BKDG) For AMD Family 10h Processors, Sept. 7 2007, http://wr0.wr.inf.h-brs.de/wr/hardware/nodes2/amd/BKDG.pdf

[136]: ACP – The Truth About Power Consumption Starts Here, AMD White Paper, 2010, http://www.amd.com/us/Documents/43761D-ACP_PowerConsumption.pdf

[137]: Wikipedia, Average CPU Power, http://en.wikipedia.org/wiki/Average_CPU_power

[138]: HyperTransport Link Specifications, HyperTransport Consortium, http://www.hypertransport.org/default.cfm?page=HyperTransportSpecifications

[139]: AMD Phenom X4 9500 - HD9500WCJ4BGD / HD9500WCGDBOX, CPU World, http://www.cpu-world.com/CPUs/K10/AMD-Phenom%20X4%209500%20-%20HD9500WCJ4BGD%20%28HD9500WCGDBOX%29.html

[140]: Introducing 45nm Quad-Core AMD Opteron Processors, Codenamed „Shanghai", Nov. 2008

[141]: Clark J. & Whitehead R., "AMD Shanghai Launch, Anandtech, Nov. 13 2008, http://www.anandtech.com/showdoc.aspx?i=3456

[142]: Exceptional Performance Per Watt and Flexibility AMD Turion II Neo Dual-Core Processor and AMD 785E Chipset, 2011, http://www.amd.com/us/Documents/48243_ASB2_Platform_Brief_web.pdf

References (16)

[143]: Phenom II X6: the six-core CPUs, HubPages, Oct. 6 2010, http://ancillotti.hubpages.com/hub/Phenom-II-X6-the-six-core-CPUs

[144]: Six-Core AMD Opteron Processors: Top-line performance that's bottom-line efficient, 2009, http://www.redhat.com/f/pdf/rhevonhp/SixCore_47469B.pdf

[145]: Wasson S., Kowaliski C., AMD's Phenom II X6 processors, Tech Report, April 27 2010, http://techreport.com/review/18799/amd-phenom-ii-x6-processors

[146]: AMD Financial Analyst Day, Nov. 11 2009

[147]: AMD Athlon 64 X2 Dual-Core Processor Key Architectural Features, http://www.amd.com/us/products/desktop/processors/athlon-x2/Pages/amd-athlon-x2-dual-core-processors-key-architectural-features.aspx

[148]: BIOS and Kernel Developer's Guide for AMD Athlon 64 and AMD Opteron Processors, Rev. 3.30, Febr. 2006, http://support.amd.com/us/Processor_TechDocs/26094.PDF

[149]: Quad-Core AMD Opteron Processors, Fast Facts, 2008

[150]: Wikipedia, AMD mobile platform, http://en.wikipedia.org/wiki/AMD_mobile_platform

[151]: Carver T., "Magny-Cours" and Direct Connect Architecture 2.0, AMD Developer Central, March 29 2010