AMD's Zen-based processor lines (Family 17h)

Dezső Sima

January 2019

(Ver. 2.2)

AMD's processor lines belonging to the Zen Family (Family 17h)

In the Lectures only Sections 1 - 4, 6.1 and 8.2 will be discussed (except Sections 2.1.3 and 2.1.4).

1. Introduction

1. Introduction

Overview of AMD's processor lines

AMD's in-house designed x86 families

32-bit x86 families	The Hammer family	Intermediate families	The Bulldozer family	The Cat family	The Zen family
K5/K6/K7 families	**K8/K10/K10.5 families (08h/10h/10.5h)**	**Families 11h/12h**	**Family 15h**	**Families 14h/16h**	**Family 17h**
(32-bit Mobile/DT)	(64-bit x86 family)	(Mobile/DT oriented)	(High-performance oriented)	(Low-power oriented)	(Modular design)
1996-2003)	(2003-2009)	(2008-2011)	(2011-2016)	(2011-2015)	(2017-)

K7: Athlon (1999-2003)

Remark

Before the in-house designed K5, AMD licensed and manufactured Intel designed processors.

1. Introduction

AMD's unit shares on the world market [51]

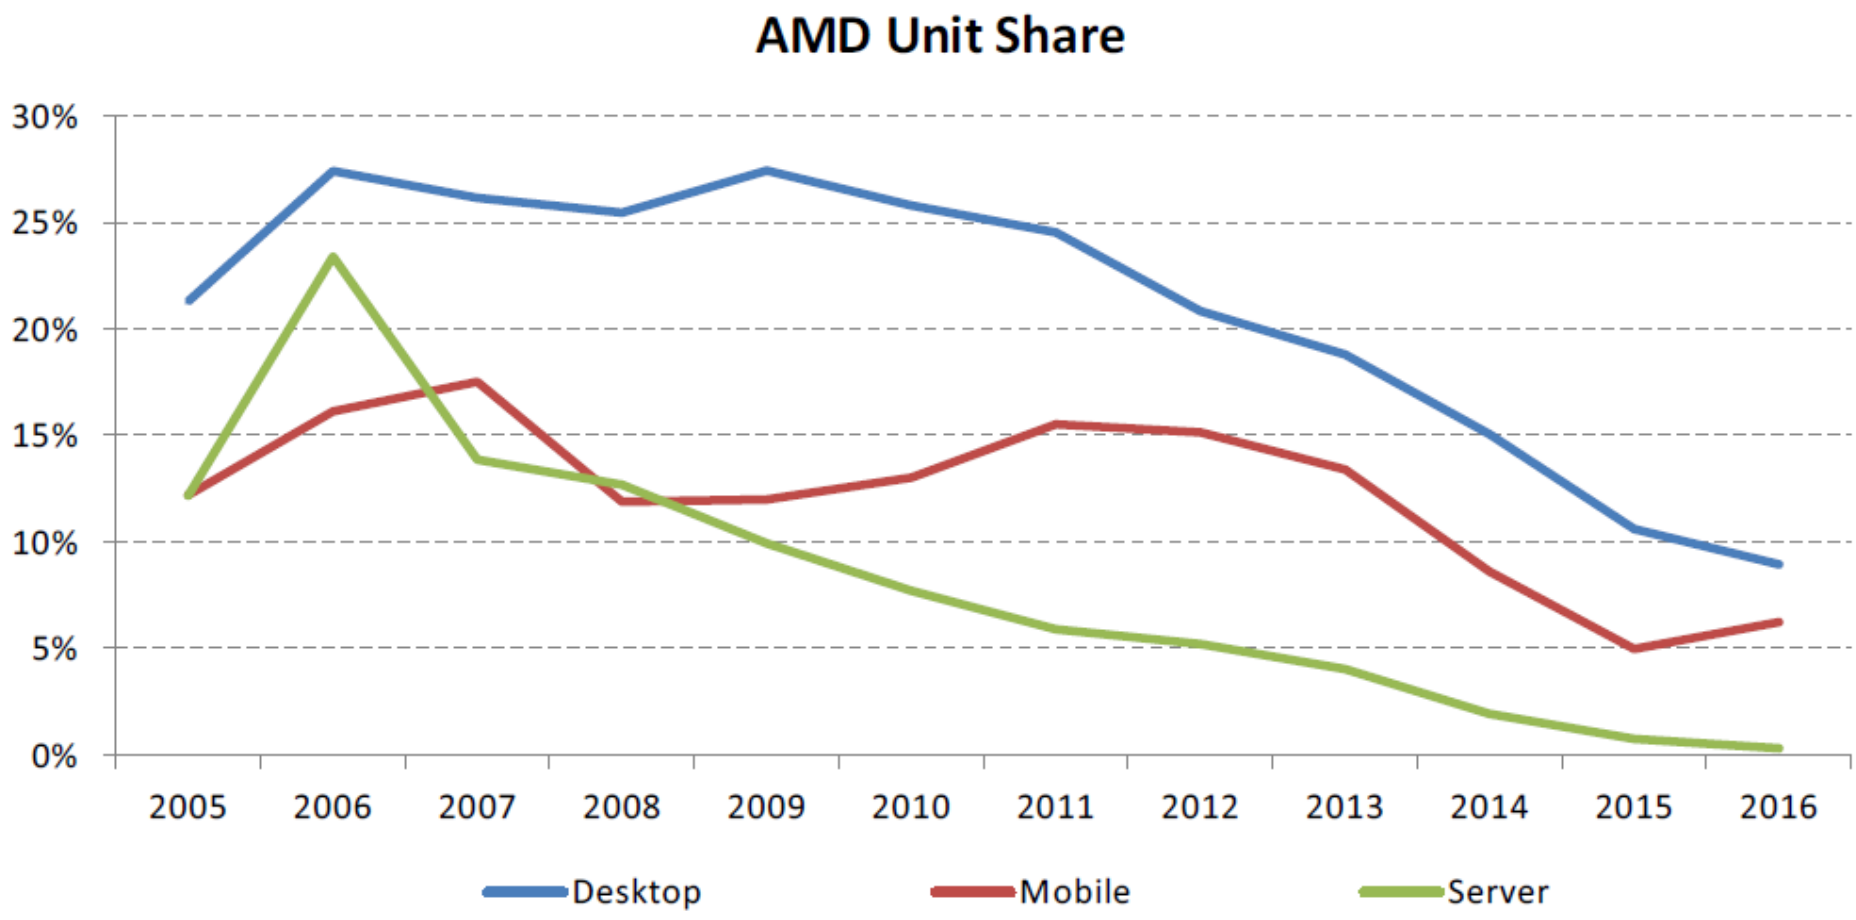

Source: Mercury, Morgan Stanley Research

Efficiency of Intel and AMD CPUs - 2004 to Ryzen [89]

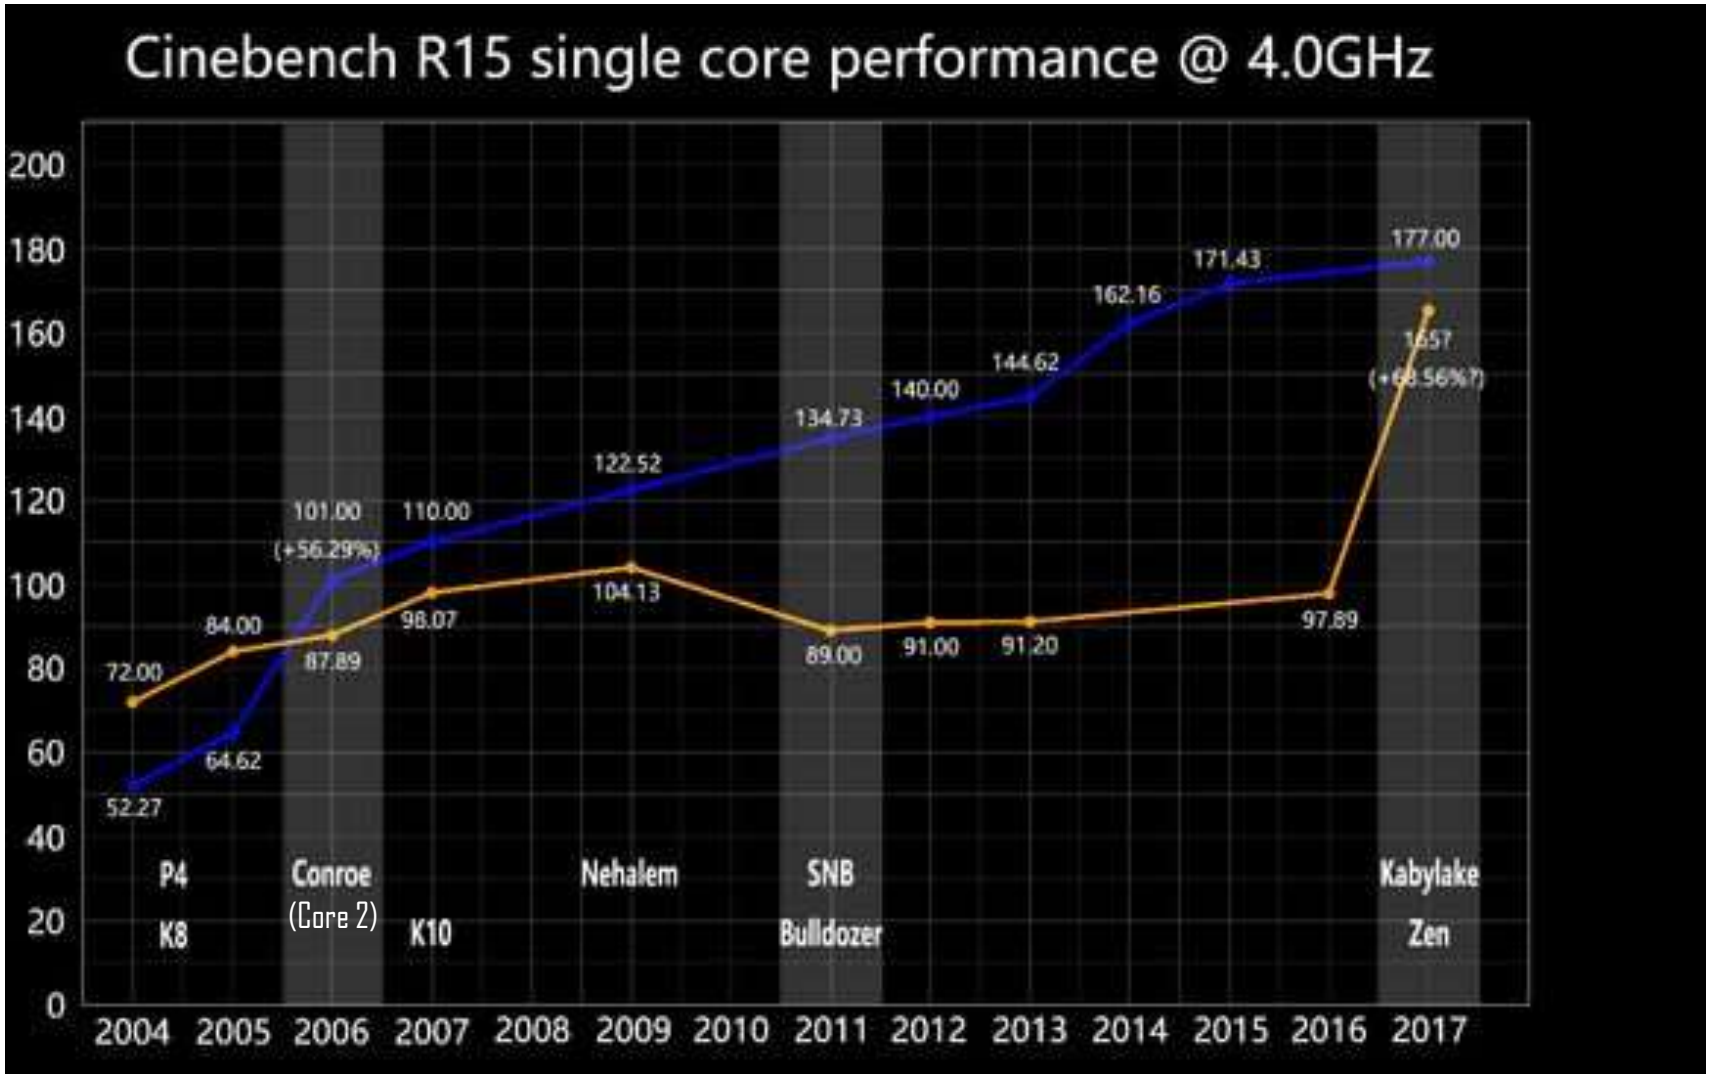

Announcing details of the Zen architecture at Computex 6/2016 [1]

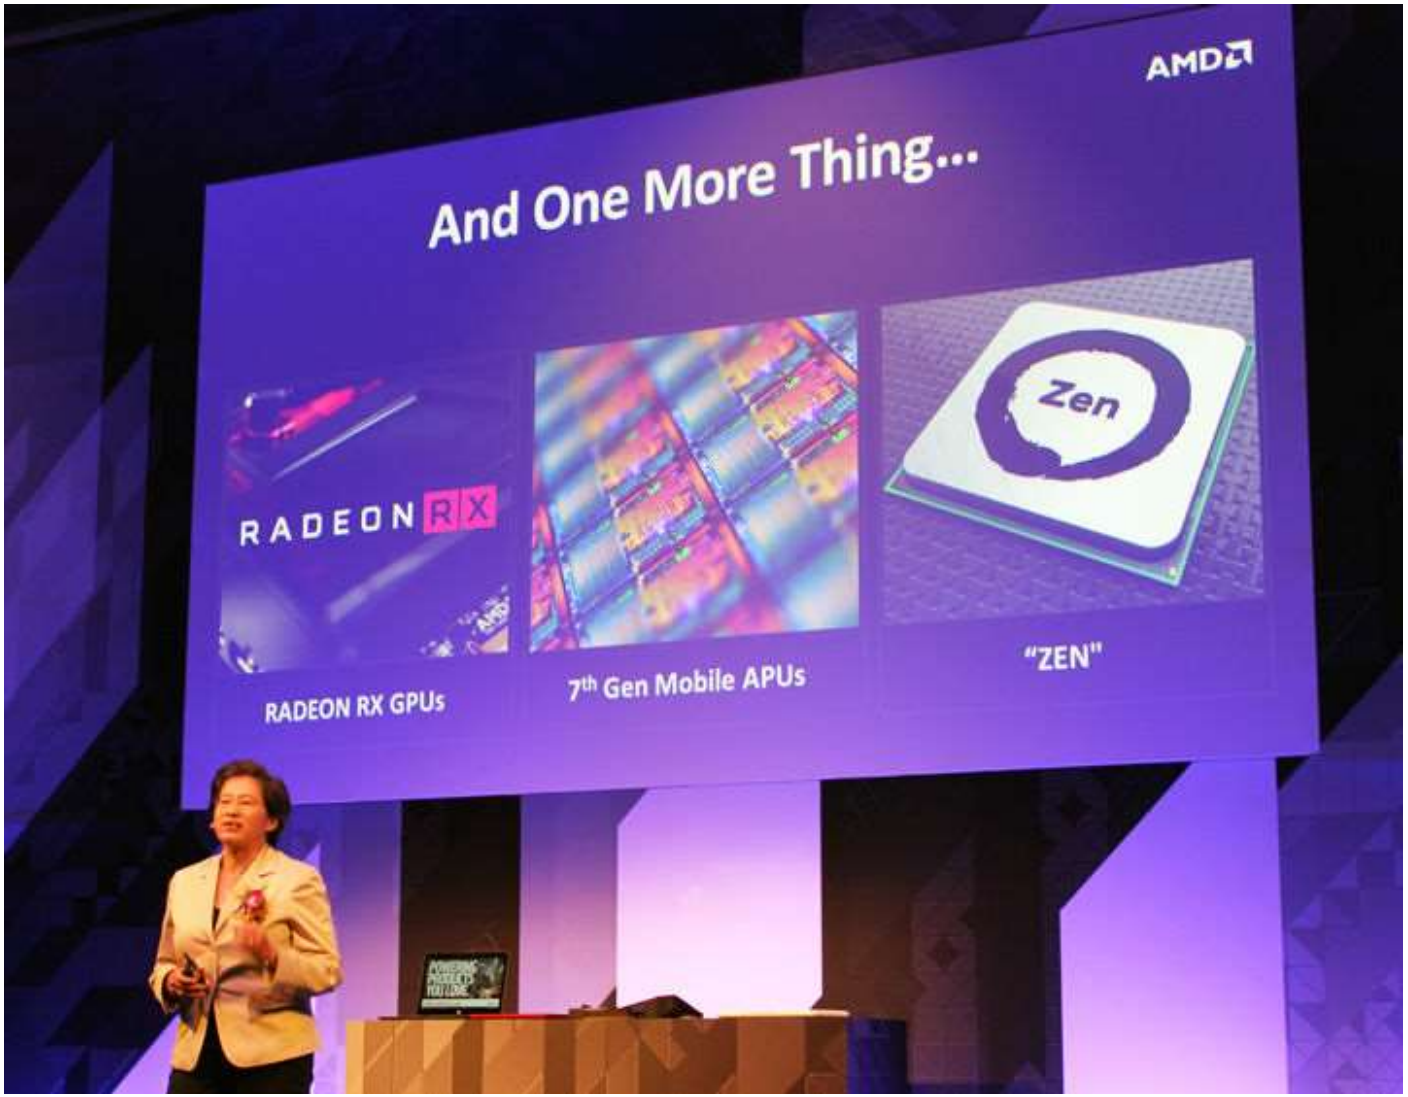

Brand names of AMD's Zen-based (Family 17h-based) processor lines

	Launched in	2017-2018	2018	2019
		Family 17h (00h-0Fh) (**Zen**)	**Family 17h** (00h-0Fh) (**Zen+**)	**Family 17h** (xxh-xxh) (**Zen+**)
		14 nm	12 nm	12 nm
Servers	**4P servers**			
	2P servers	Naples (*Epyc 7xx1)*		
	1P servers	Naples (*Epyc 7xx1P)*		
	(85-140 W)			
Desktops	**HED** (~95-125 W)	Whitehaven *(ThreadRipper (TR) 1xxxX)*	Pinnacle Ridge *(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)			Picasso (APU) Ryzen 7/5 3xx0H)
	Ultra portable (~10-15 W)	Raven Ridge (APU) *(Ryzen 7/5/3 2x00U)*		Picasso (APU) Ryzen 7/5/3 3xx0U)
	Tablet (~5 W)			

AMD Zen-based processor lines introduced in 2017 (except Pro series)

AMD's Zen-based processor lines introduced in 2017

	Ryzen Mobile APU (Raven Ridge)	**Ryzen DT** (Summit Ridge)	**ThreadRipper** (Whitehaven)	**Epyc** (Naples)
Market segment	Mobile	Desktop platform	HED	1S/2S server
µarch./Technology	Zen 14 nm	Zen, 14 nm	Zen, 14 nm	Zen, 14 nm
Launched models	Ryzen 7 2700U Ryzen 5 2500U (10/2017)	Ryzen 7 (3/2017) Ryzen 5 (4/2017) Ryzen 3 (7/2017)	1950X/1920X/1900X (8/2017)	Series 7000 (6/2017)
Layout	CCX + Vega 8/10	Zeppelin die with 2x CCX	MCM (2x Zeppelin die)	MCM (4x Zeppelin die)
Integrated GPU	Yes	No	No	No
Core count	4	4/6/8	8/12/16	8/16/24/32
SMT	SMT	SMT (except Ryzen 3)	SMT	SMT
Mem. channels/rate	2xDDR4-2400	2xDDR4-2666	4xDDR4-2666	8xDDR4-2666
PCIe 3.0 lanes	??	16xPCIe 3.0	60xPCIe 3.0	128 for 1S servers 64 for 2S servers
TDP	15 W	65/95 W	180 W	120/170/180 W
Socket	AM4 (1331)	AM4 (1331)	TR4 (SP3r2) (4094)	SP3 (4094)
Chipset	SoC	300-series	X399	No chipset, SOC

Processor logos of the Zen series [8]

AMD Zen-based processor brands						
Logo	Family	General Description	Features			
			Cores	Unlocked	AVX2	SMT
	Ryzen Mobile	Mobile processors with VEGA graphics	Quad	✓	✓	✓
	Ryzen 3	Entry level Performance	Quad	✓	✓	✗
	Ryzen 5	Mid-range Performance	Quad	✓	✓	✓
			Hexa	✓	✓	✓
	Ryzen 7	High-end Performance	Octa	✓	✓	✓
	Ryzen Threadripper	Enthusiasts	8-16	✓	✓	✓
	EPYC	High-performance Server Processor	8-32	✓	✓	✓

AMD Zen/Zen+ based processor lines introduced in 2018 (except Pro series)

AMD's Zen/Zen+ based processor lines introduced in 2018

	Ryzen Mobil APU (Raven Ridge)	**Ryzen DT APU** (Raven Ridge)	**2.G Ryzen DT** (Pinnacle Ridge)	**2.G ThreadRipper** (Pinnacle Ridge)
µarch./tech.	Zen/14 nm	Zen/14 nm	Zen+/12 nm	Zen+/12 nm
Launched models	Ryzen 3 2300U Ryzen 3 2200U (1/2018)	Ryzen 5 2400G/GE Ryzen 3 2200G/GE (G: 2/2018) (GE: 42018)	Ryzen 7 2700X Ryzen 7 2700 Ryzen 5 2600X Ryzen 5 2600 (4/2018)	2990WX (8/2018) 2950X (8/2018) 2970WX (10/2018) 2920X (10/2018)
Layout	CCX + Vega 6/3	CCX + Vega 11/8	Zeppelin die (2x CCX)	MCM (up to 4 Zeppelin dies)
Integrated GPU	Yes	Yes	No	No
Core count	2/4	4	2/4	Up to 32
SMT (Multithreaded)	SMT only for Ryzen 3 2200U	SMT only for Ryzen 5	SMT	SMT
Mem. channels/ data rate	2xDDR4-2400	2xDDR4-2667 (G) 2xDDR4-2933(GE)	2xDDR4-2933	4xDDR4-2933
PCIe lanes	??	??	16x PCIe 3.0	60x PCIe 3.0
TDP	15 W	35/65 W	65/95/105 W	180/250 W
Socket	FP5 (na.)	AM4 (1331)	AM4 (1331)	TR4 (SP3r2) (4094)
Chipset	SoC	300-series	300/400-series	X399

AMD Zen/Zen+ based processor lines introduced in 2019

AMD's Zen/Zen+ based processor lines introduced in 2019

	2.G Ryzen Mobil APU (Picasso)			
µarch./tech.	Zen+/12 nm			
Launched models	Ryzen 7 3750H/3700U Ryzen 5 3550H/3500H Ryzen 3 3300U/3200U (1/2019)			
Layout	CCX + Vega 10/8/3			
Integrated GPU	Yes			
Core count	2/4			
SMT (Multithreaded)	Yes			
Mem. channels/ data rate	2xDDR4-2400			
PCIe lanes	??			
TDP	H:35W U:15 W			
Socket	FP5			
Chipset	SoC			

AMD's effort to design the Zen core [10]

M. Papermaster, CTO of AMD: Up to 300 engineers were working on the Zen core spending over two million working hours.

Aim: optimizing the power/frequency curve.

CTO: Chief Technology Officer

Innovations introduced within the Zen family [10]

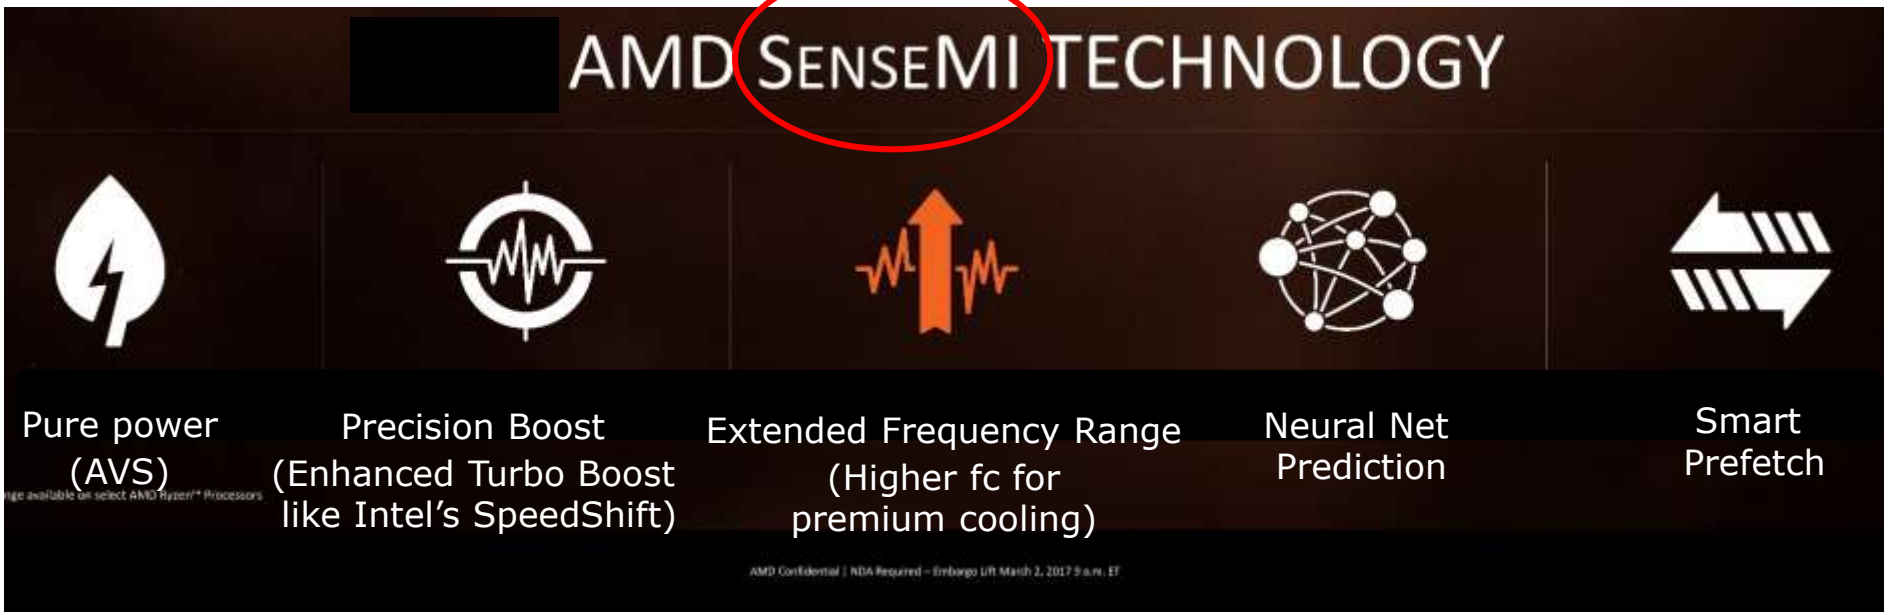

To reduce dissipation

See Section 2.1.3.1

To raise fc

See Sections 2.1.3.2/2.1.3.3

To raise IPC

See Section 2.1.2

Use of advanced technologies in AMD's mobile and desktop series [75] -1

	Ryzen DT	2nd gen. Ryzen DT	Ryzen DT APU with Radeon Vega graphics	Ryzen mobile APU with Radeon Vega graphics
Extended Frequency Range (XFR)	Yes (XFR)	Yes (XFR 2)	Yes (XFR 2)	Yes (Mobile XFR)[4]
Precision Boost (Turbo Boost)	Yes Precision Boost)	Yes (Precision Boost 2)	Yes (Precision Boost 2)	Yes (Precision Boost 2)
Neural Net Prediction (for branch prediction)	Yes	Yes	Yes	Yes
Pure Power (AVS)	Yes	Yes	Yes	Yes
Smart Prefetch (Advanced data prefetch)	Yes	Yes	Yes	Yes

Use of advanced technologies in AMD's HED and 2S server series [75] -2

	ThreaRripper	2nd gen Ryzen ThredRipper	Epyc
Extended Frequency Range (XFR)	Yes (XFR)	Yes (XFR 2)	No
Precision Boost (Turbo Boost)	Yes Precision Boost)	Yes (Precision Boost 2)	Yes (Precision Boost)
Neural Net Prediction (for branch prediction)	Yes	Yes	Yes
Pure Power (AVS)	Yes	Yes	Yes
Smart Prefetch (advanced data prefetch)	Yes	Yes	Yes

AMD's Zen or Zen+ based processor series

Series	CPU core	µarch.	Techn.	Launched	Models
Ryzen DT	Summit Ridge	Zen	14 nm	3/2017	Ryzen 1xxx Ryzen 1xxxX
2. gen. Ryzen DT	Pinnacle Ridge	Zen+	12 nm	4/2018	Ryzen 2xxx Ryzen 2xxxX
3. gen. Ryzen DT	Matis	Zen 2	7 nm	1/2019	n.a.
Ryzen DT APU	Raven Ridge	Zen	14 nm	2/2018	Ryzen 2xxxG Ryzen 2xxxGE
Ryzen mobile APU	Raven Ridge	Zen	14 nm	10/2017	Ryzen 2xxxU
2. gen. Ryzen mobile APU	Picasso	Zen+	12 nm	1/2019	Ryzen 3xxxU/H
ThreadRipper HED	Whitehaven	Zen	14 nm	8/2017	TR 1xxx
2. gen. Threadripper HED	Pinnacle Ridge	Zen+	12 nm	8/2018	TR 2xxxX TR2xxxWX
EPYC 1S/2S server	Naples	Zen	14 nm	6/2017	EPYC 7xxx

AMD's 2017-2019 desktop/notebook roadmap [64]

	2017	2018	2019
Desktop CPU (No GPU)	Summit Ridge Zen/14 nm ▲ Up to 16 Zen Threads ▲ Socket AM4	Pinnacle Ridge Zen+/12 nm ▲ Summit Ridge architecture ▲ Performance uplift ▲ Socket AM4	Matisse Zen 2/7 nm ▲ Zen 2 Cores ▲ Socket AM4
Desktop/Notebook APU	Bristol Ridge Excavator/28 nm ▲ Excavator CPU ▲ Polaris GPU ▲ Socket AM4 desktop ▲ Socket FP4 notebook	Raven Ridge Zen/14 nm ▲ Up to 8 Zen Threads ▲ Up to 11 Vega CU's ▲ Socket AM4 desktop ▲ Socket FP5 notebook	Picasso Zen+/12 nm ▲ Raven Ridge architecture ▲ Power/Performance uplift ▲ Socket AM4 desktop ▲ Socket FP5 notebook

Design paradigms for segmenting multicore processors

Design paradigms for segmenting multicore processors

Monolithic implementation

All cores are implemented on the same die

Example

Intel's Skylake-SP (2017) *[2]*

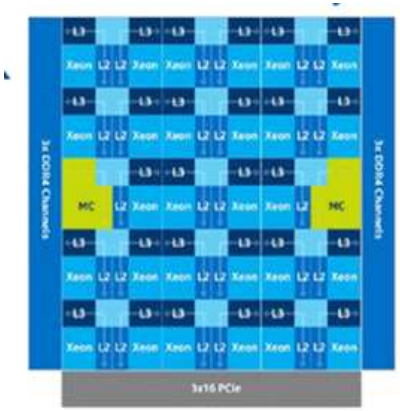

Up to 28 cores

Multi-Chip-Module (MCM)

Cores are implemented on a number of dies, they are properly interconnected and mounted into the same package

AMD's Epyc (2017) *[2], [3]*

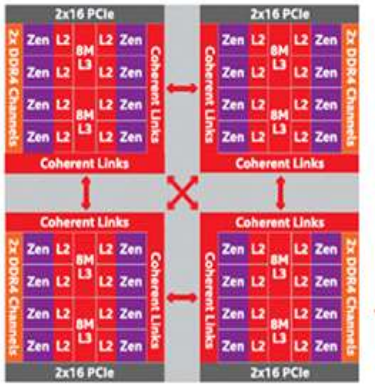

Up to 4x8 cores

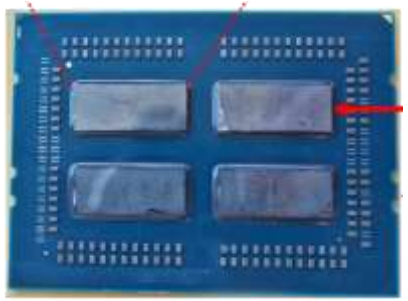

Single package

*

Example of Intel's early MCM design (Dual core DT Pentium D-820) (5/2005)) [91], [92]

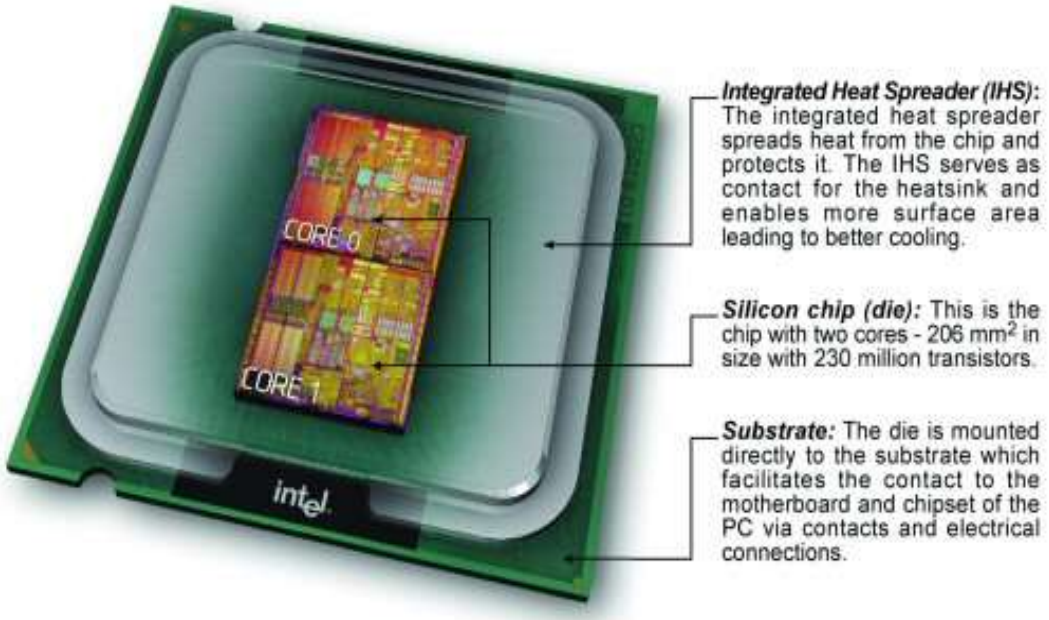

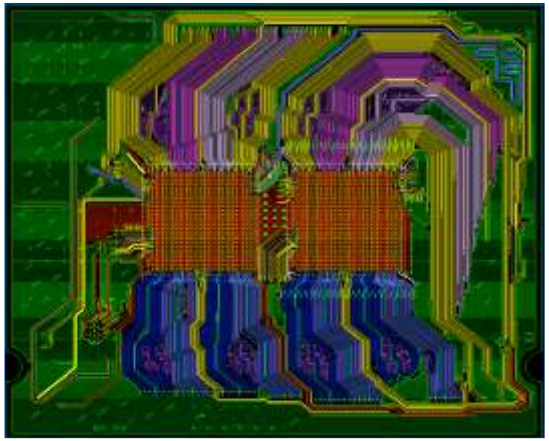

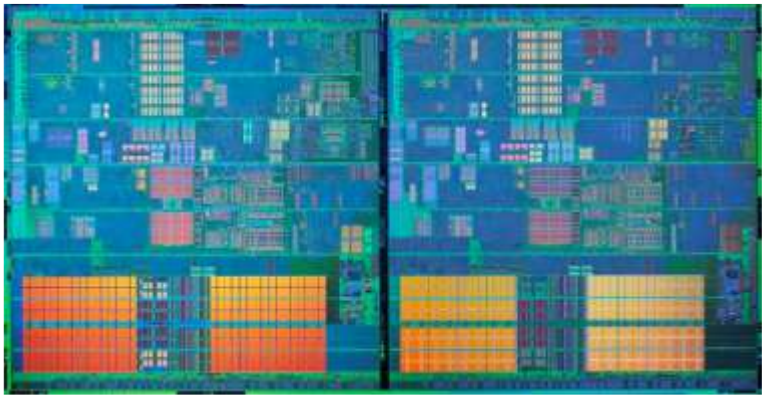

Further examples for MCM designs of server processors

Intel

Paxville DP 2.8 (2x1C), 2xDP-enhanced Pentium 4 Prescott (10/2005)
Xeon 5000 (Dempsey), (2x1C), ~ the 65 nm shrink of the 90 nm Paxville DP 2.8 (5/2006)

Xeon 5300 (Clowertown) [4], [5]

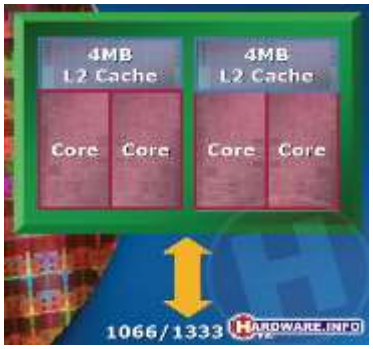

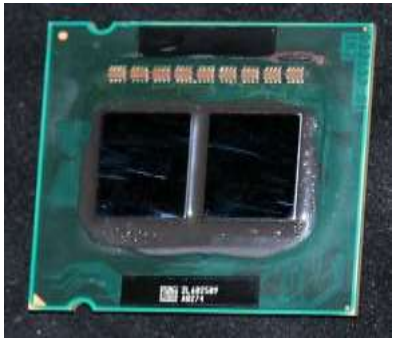

Xeon 5300 (Clowertown) (2x2C) 2xCore 2-based 65 nm Xeon 5100 (Woodcrest), 11/2006
Xeon 5400 Harpertown 2x2C, 2x Core 2-based 45 nm 2C Harpertown (11/2007)

AMD

Opteron 6100 (Magny-Course) 2x6C, 2xBulldozer-based 45 nm Istambul die, (3/2010)
Opteron 6200 (Interlagos), 2x8C, 2xBulldozer-based 32 nm Orochi die, (11/2011)
Opteron 6300 (Abu Dhabi), 2x8C, 2xPiledriver-based 32 nm Abu Dhabi die, (11/2012)
Opteron 6400 (Warsaw), 2x8C, 2xPiledriver-based, 32-nm Abu-dhabi die, (2014)

IBM

POWER4, 2C on 4 chips, 180 nm, (12/2001)
z10, 10/2008) 5x4C,(up to 17C), 65 nm (10/2008)

z10 [6]
5 proc. dies
2 mem. control dies

NVIDIA's possible future use of MCMs, revealed in a research paper [32]

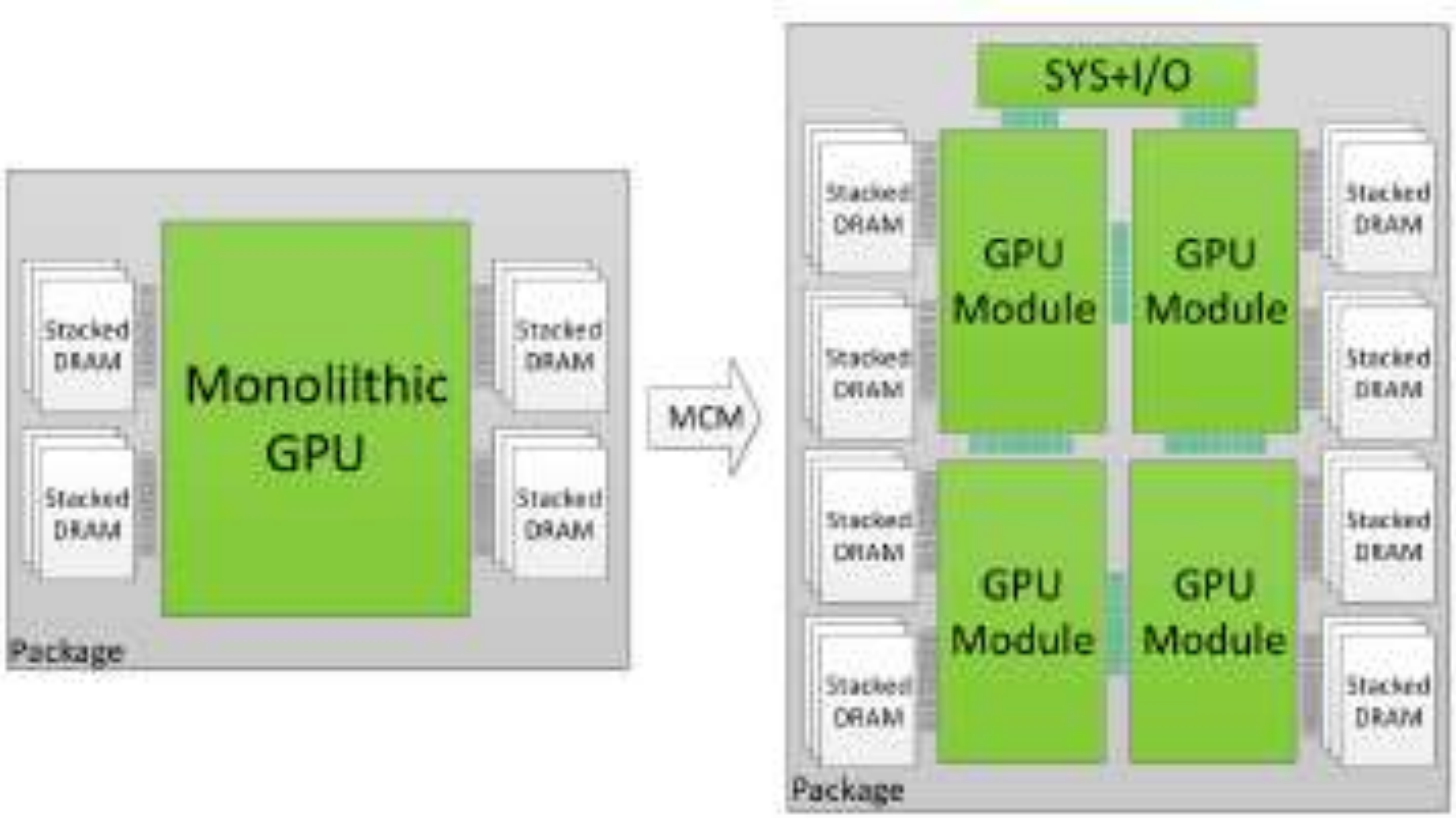

AMD's cost argumentation for the MCM layout over the monolithic one [7]

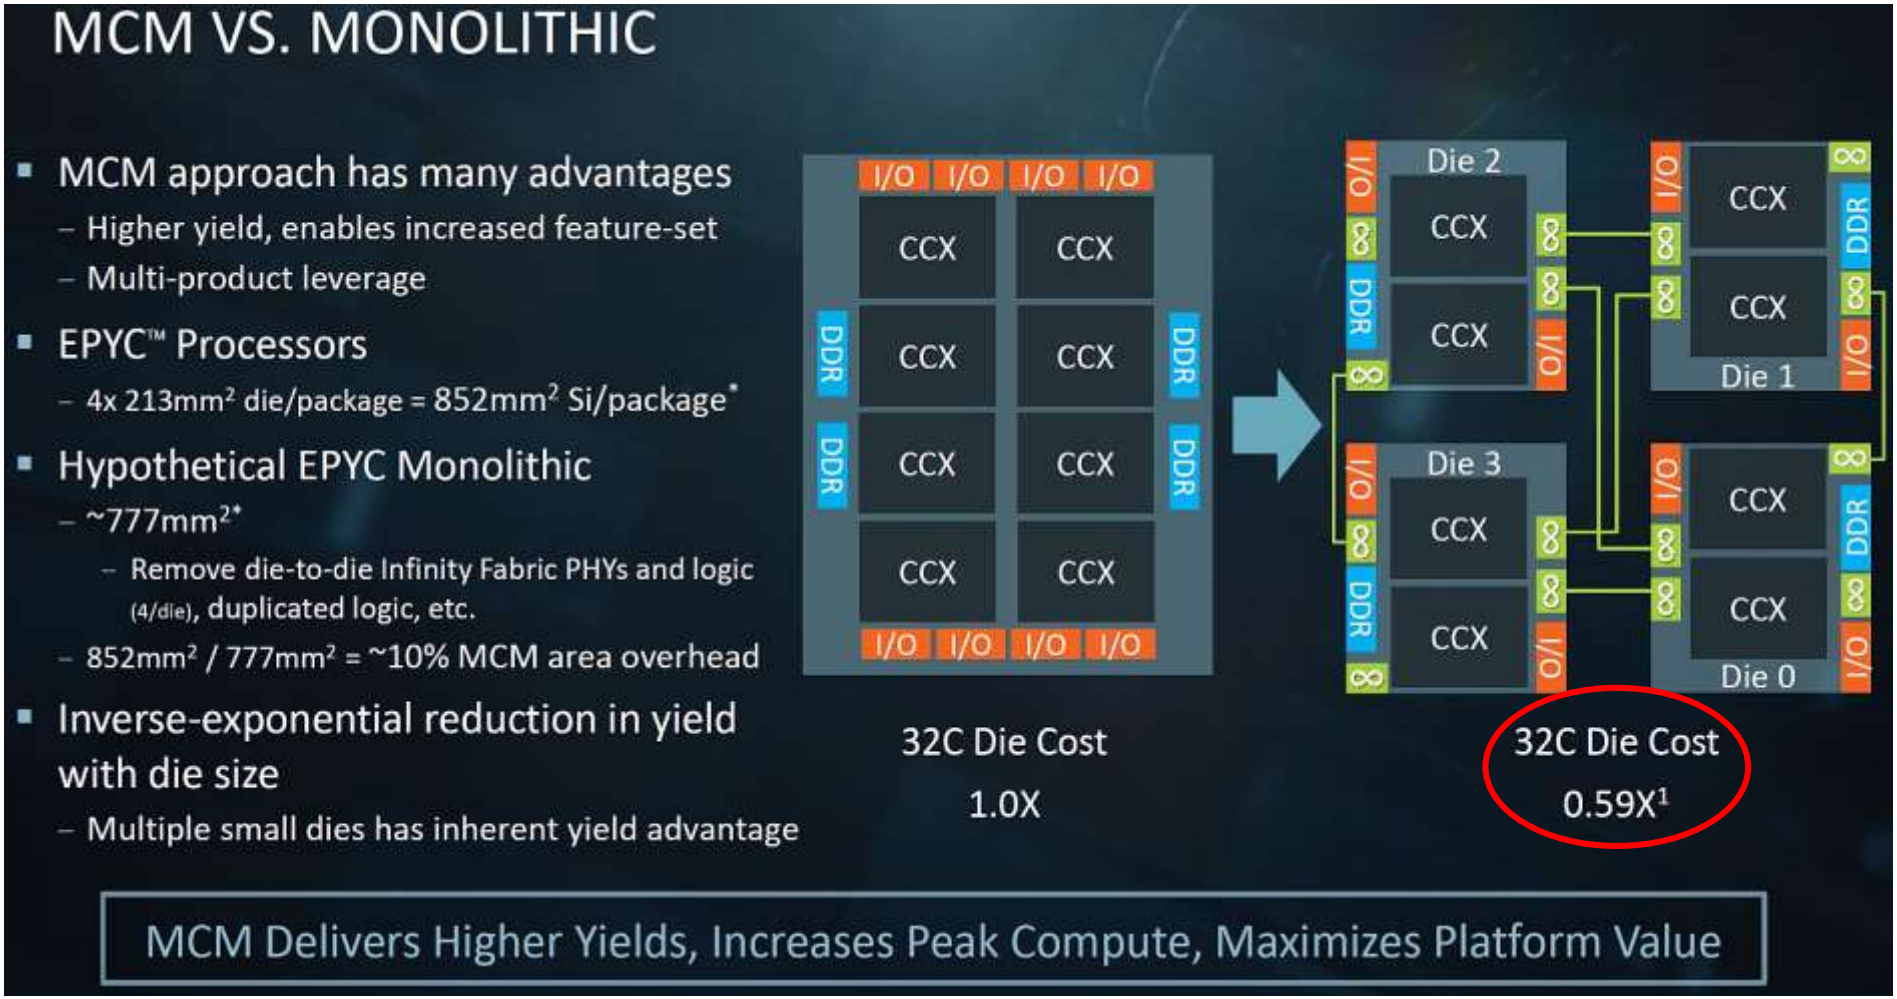

Presented by AMD at the Hot Chips Conference in Aug. 2017

*

Pros and cons of a modular processor design [3]

Pros

- It is more economical to manufacture large core count, and therefore large dies if segmented into a number of smaller dies
- Memory channels and I/O become linearly scaled with the die count (by proper design)
- It becomes possible to design processors for different market segments
- by including different numbers of the same die

Cons

- High die-to-die transfer latencies that degrade performance

*

Basic building blocks of AMD's Zen-based processors

Basic building blocks of AMD's Zen-based processors

Zen core [8]

4-core CCX [9]
(Core CompleX)

8-core Zeppelin module [10]
2x CCX

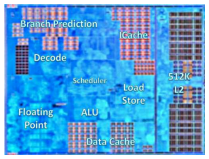

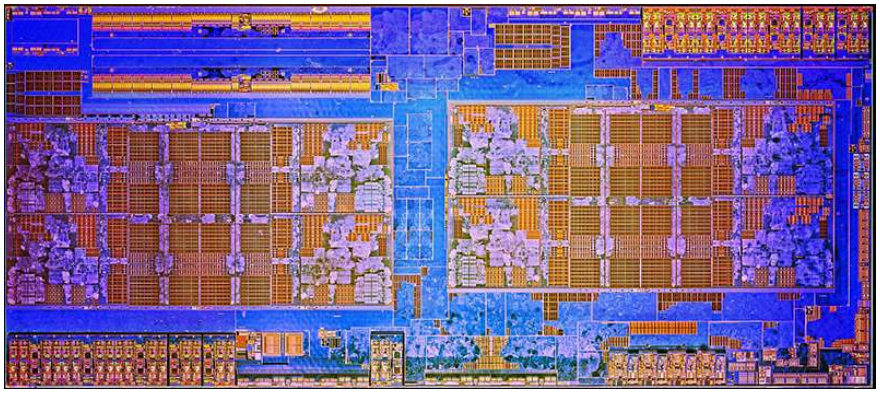

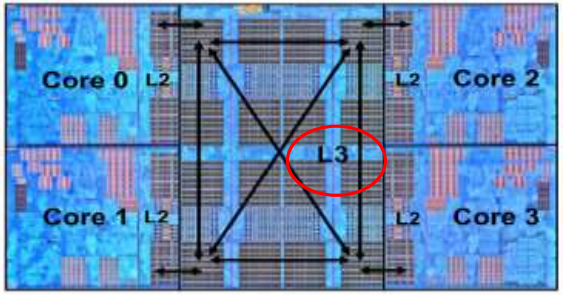

Overview of AMD's Zen-based processor lines

AMD's Zen-based processor lines

Ryzen Mobile (Mobil)	**Ryzen** (DT)	**ThreadRipper** (HED)	**Epyc** (1S/2S server)
Single CCX + Vega GPU [54]	Zeppelin chip (2x CCX)	1. gen: 2 Zeppelin chips interconnected by the IF, implemented as an MCM 2. gen: 4 Zeppelin chips	4 Zeppelin chips interconnected by the IF, implemented as an MCM

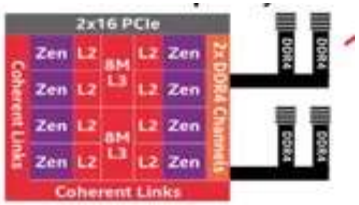

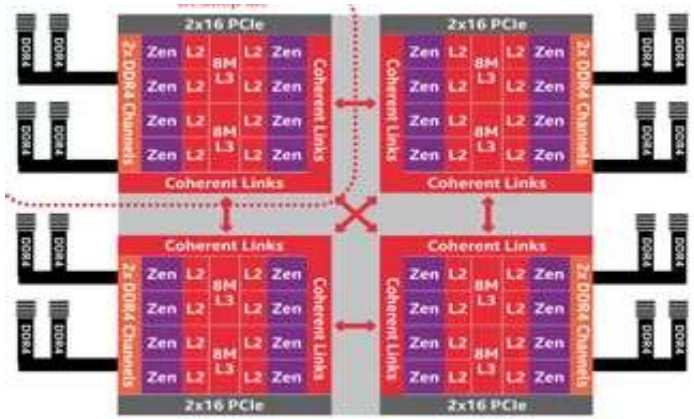

1. gen.

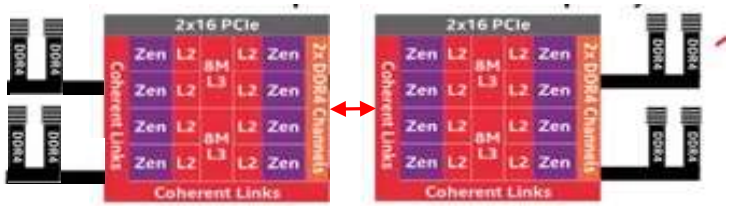

Drawings based on [11]

*

The cache coherent interconnect, called Infinity Fabric (IF)

- It is used to interconnect
  - CCX modules on a die
  - dies to dies in an MCM (Multi-Chip-Module) or
  - two sockets in a 2S server,

  as indicated in the next Figures, and discussed in more detail in Section 7.

*

Interconnecting two CCX core complexes by means of the Infinity Fabric (IF) in a Zeppelin die [88]

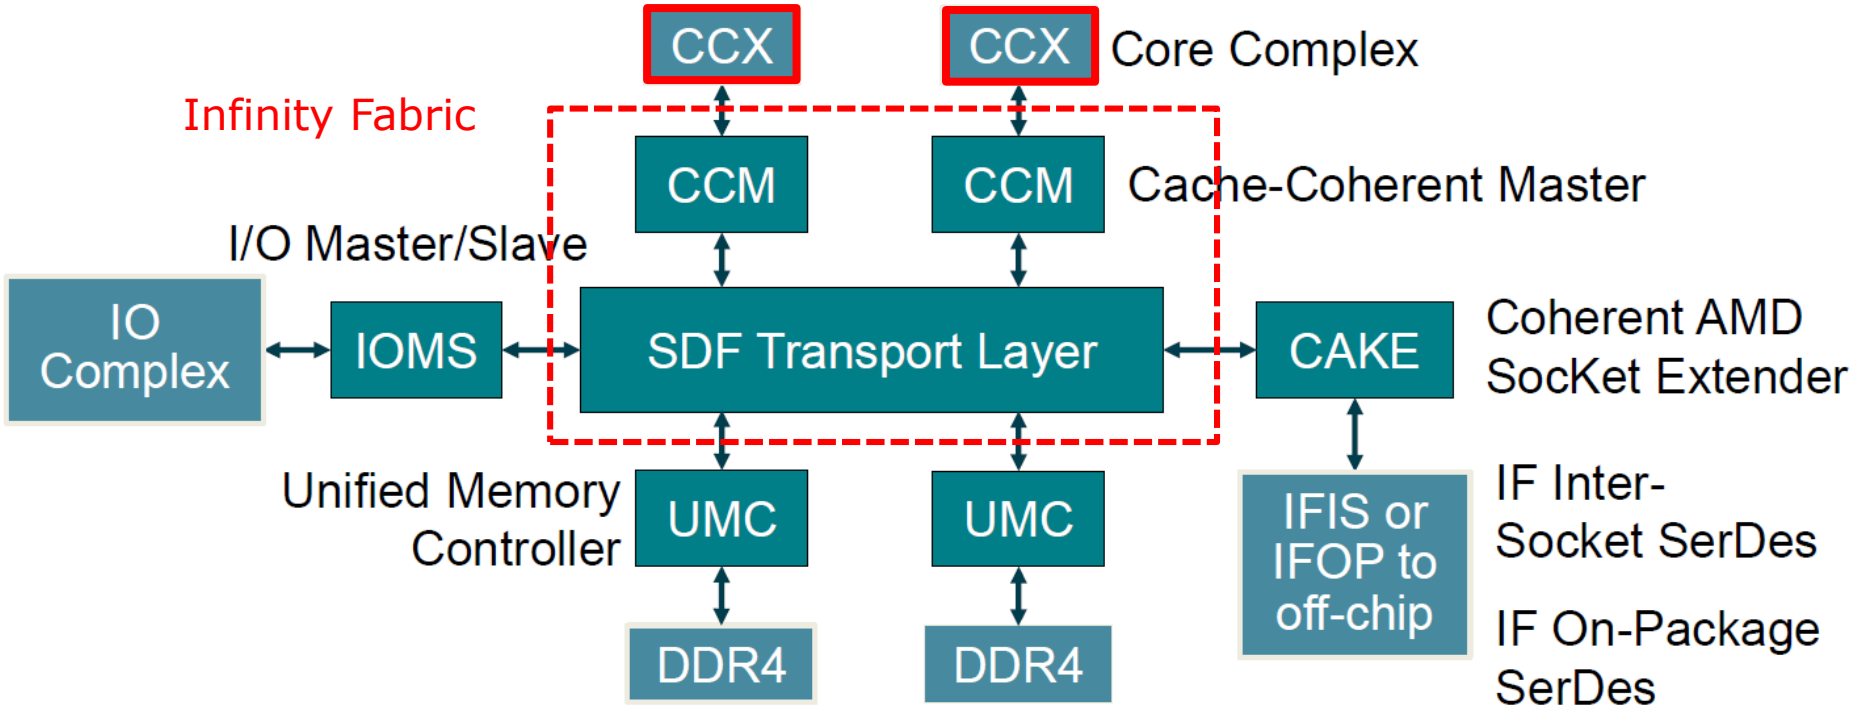

IFOP: IF On-Package interconnect links
IFIS: IF Inter-Socket interconnect links

SDF implements a crossbar-type on-die interconnect.

SerDes: (Serializer/Deserializer) interface to another chip within the system

Interconnecting dies within a socket and two sockets of a server by means of the infinity Fabric (IF) [3]

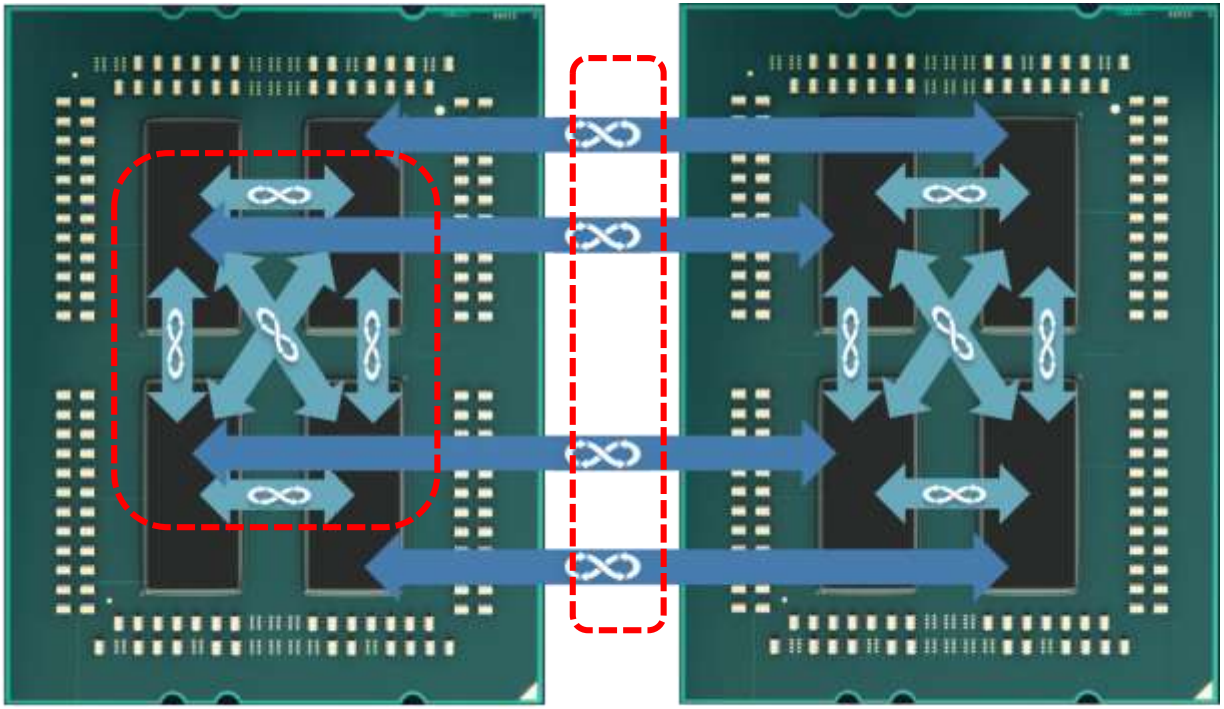

Source: AMD

Key features of the cache coherent interconnect, called Infinity Fabric

- It is an enhanced HyperTransport interconnect with
  - lower latency
  - higher bandwidth
  - lower power consumption and
  - the option to scale it up or down as needed.

*

AMD's new Zen 2-based Epyc server (Rome) for 64 cores (11/2018) -1

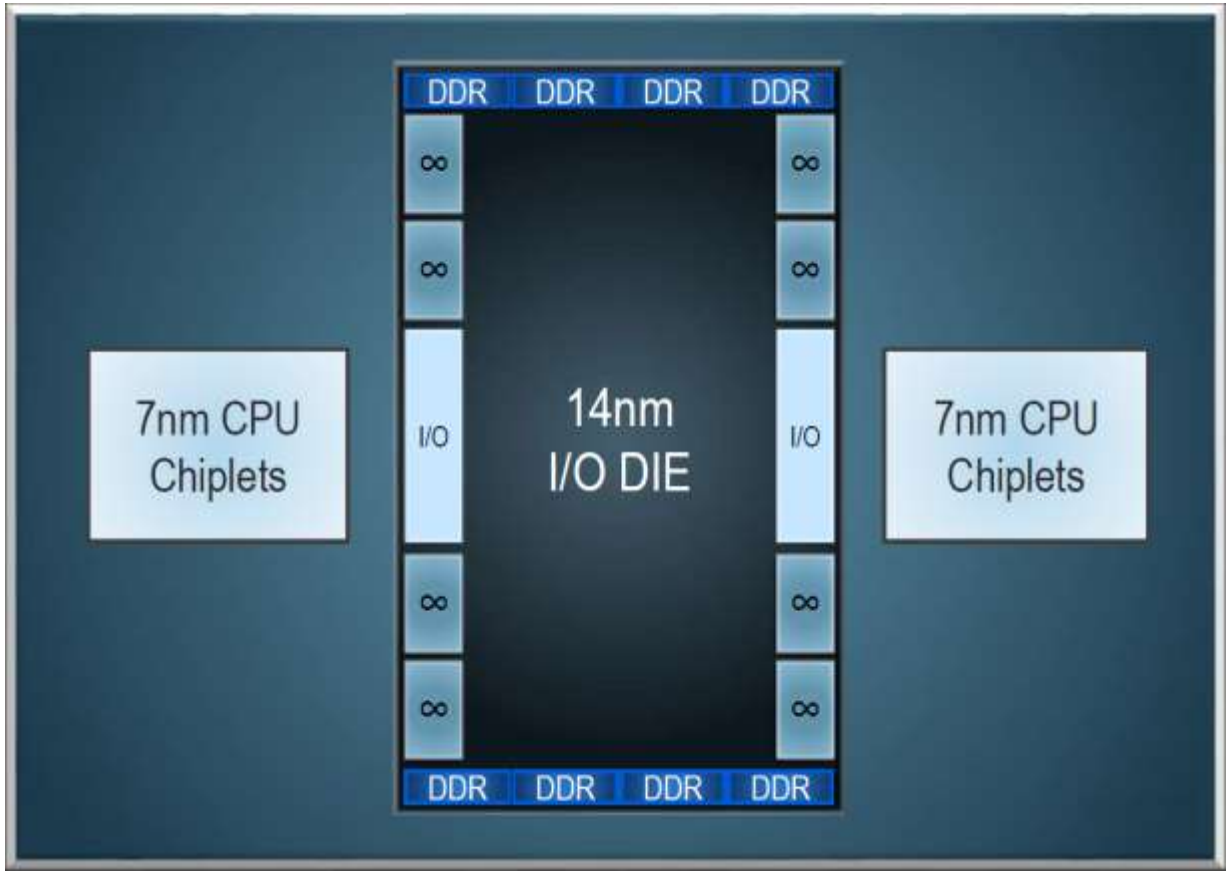

- Eight channel memory controller with equal access latency.
- Up to 4 TB DDR4 memory and 128 PCIe 4.0 lanes per socket (first PCIe 4.0!).

Similarity of AMD's MCM system concept to ARM's monolithic system concept

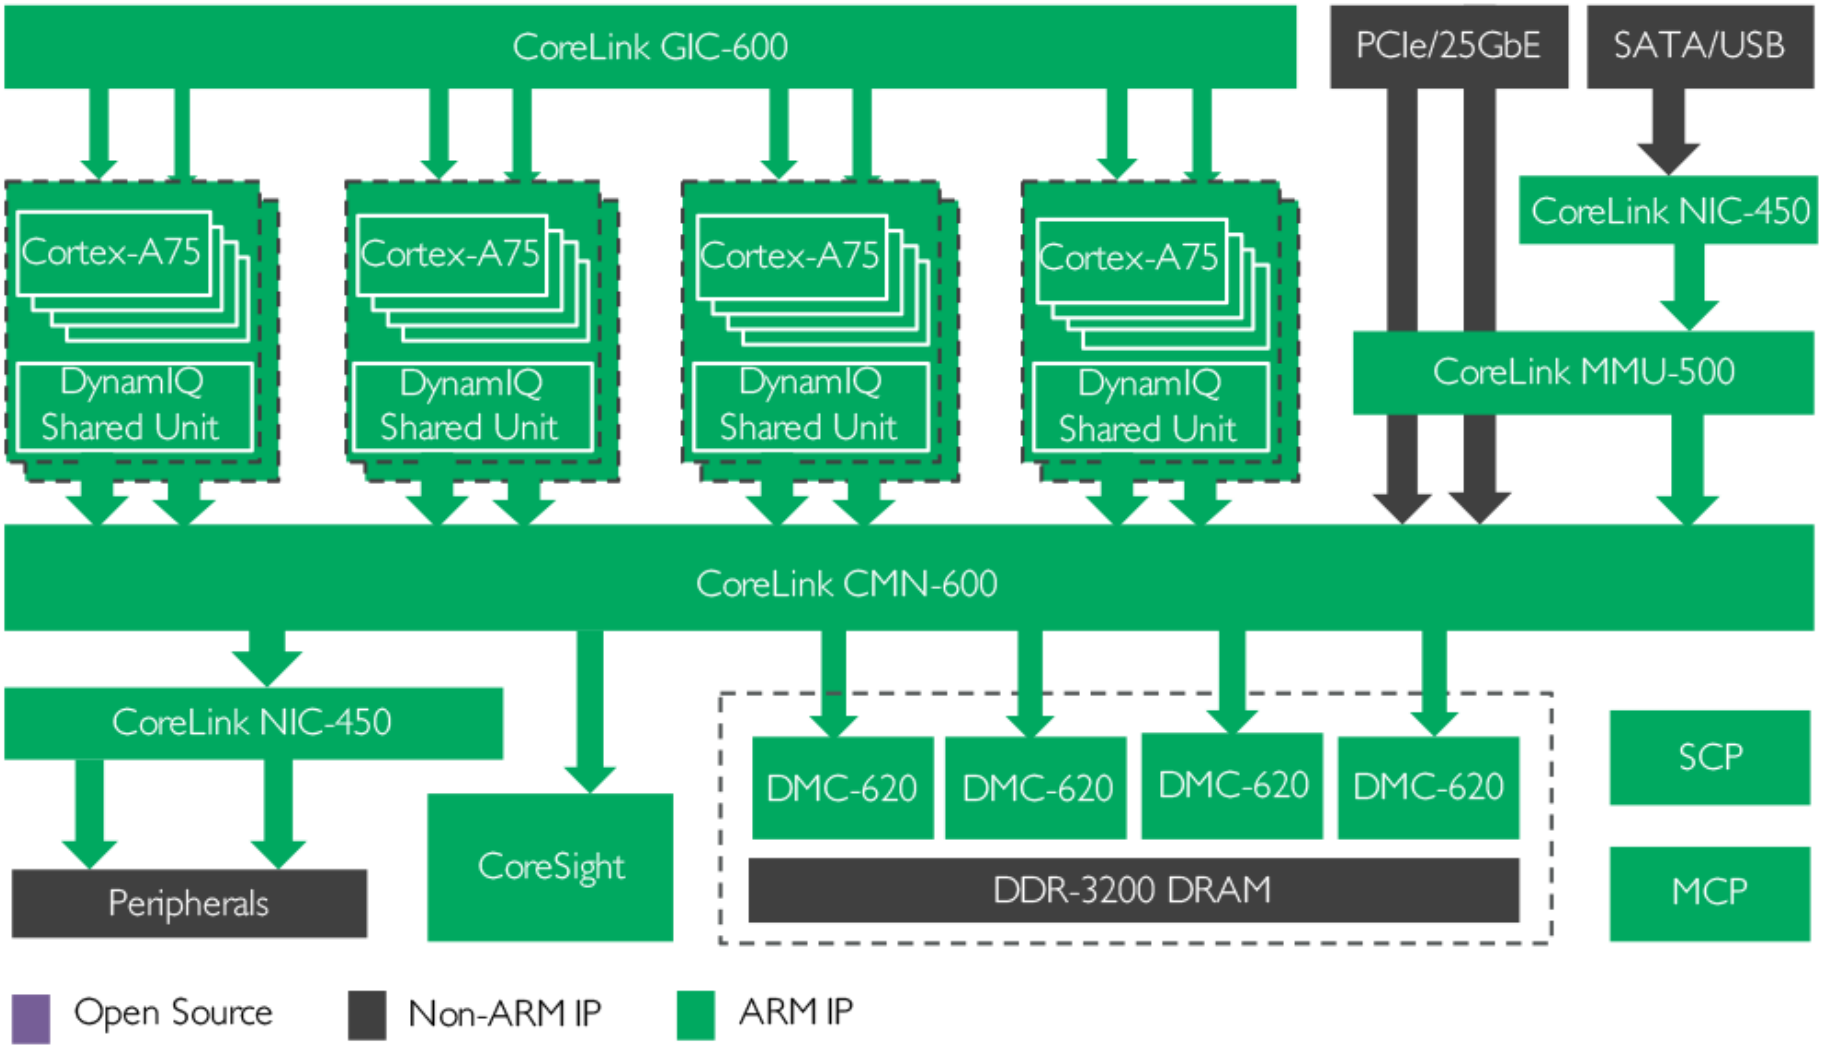

Figure: ARM's mesh interconnect (CMN-600) to build a server [97]

*

AMD’s new Zen 2-based Epyc server (Rome) for 64 cores (11/2018) -2

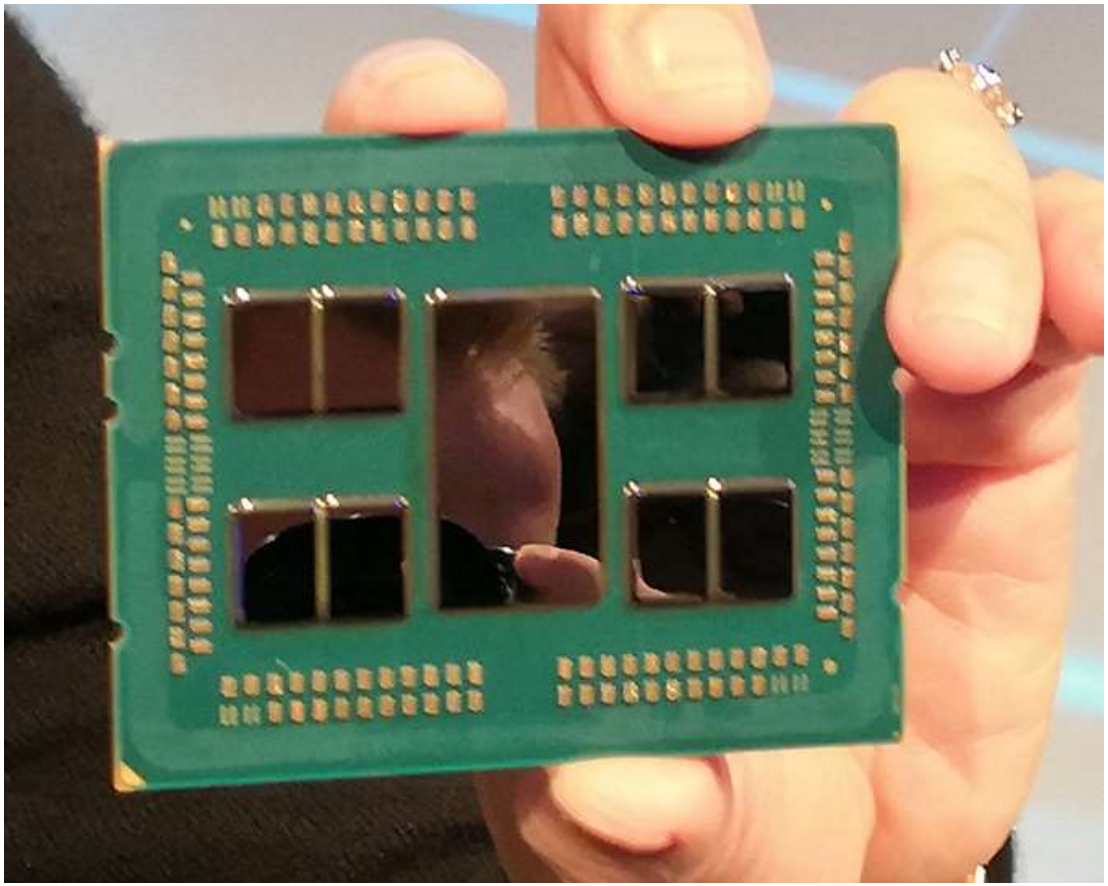

- Eight CPU chiplets with 64 cores and a 14 nm I/O chip.
- AMD has already started sampling EPYC 'Rome' processors to enterprise customers.

2. The Zen cores

- 2.1 The Zen core
- 2.2 The Zen+ core

Roadmap of the Zen core as published in 2/2018 [68]

*

2.1 The Zen core

- 2.1.1 Introduction to the Zen core
- 2.1.2 Innovations and enhancements introduced to raise IPC
- 2.1.3 Power management in Zen cores
- 2.1.4 Further details relating the Zen core

Only Sections 2.1.1 and 2.1.2 will be discussed

2.1.1 Introduction to the Zen core

2.1.1 Introduction to the Zen core [13]

- First announced at AMD's Financial Analyst Day in 5/2016.
- First details of the Zen architecture were published at the Computex Conference in 6/2016.
- According to M. Papermaster, CTO of AMD up to 300 engineers were working on the core design.
- Key design objective was raising IPC rather than power gain.
- To achieve this AMD developed the SenseMI technology package, to be discussed in Sections 2.1.2 and 2.1.3.
- Key benefits of the Zen core vs. the previous Excavator architecture
  - 14 nm process vs. 28 nm
  - 52 % performance improvement for single threaded workloads
  - Up to 3.7x improvement in performance/Watt
  - SMT support (2-way)

*

2.1 The Zen core (2)

Remark [14]

With Zen AMD departed from the Bulldozer design style cramming two integer parts and an FP part into a module and declaring the module as being dual cores and followed the regular way of building up cores.

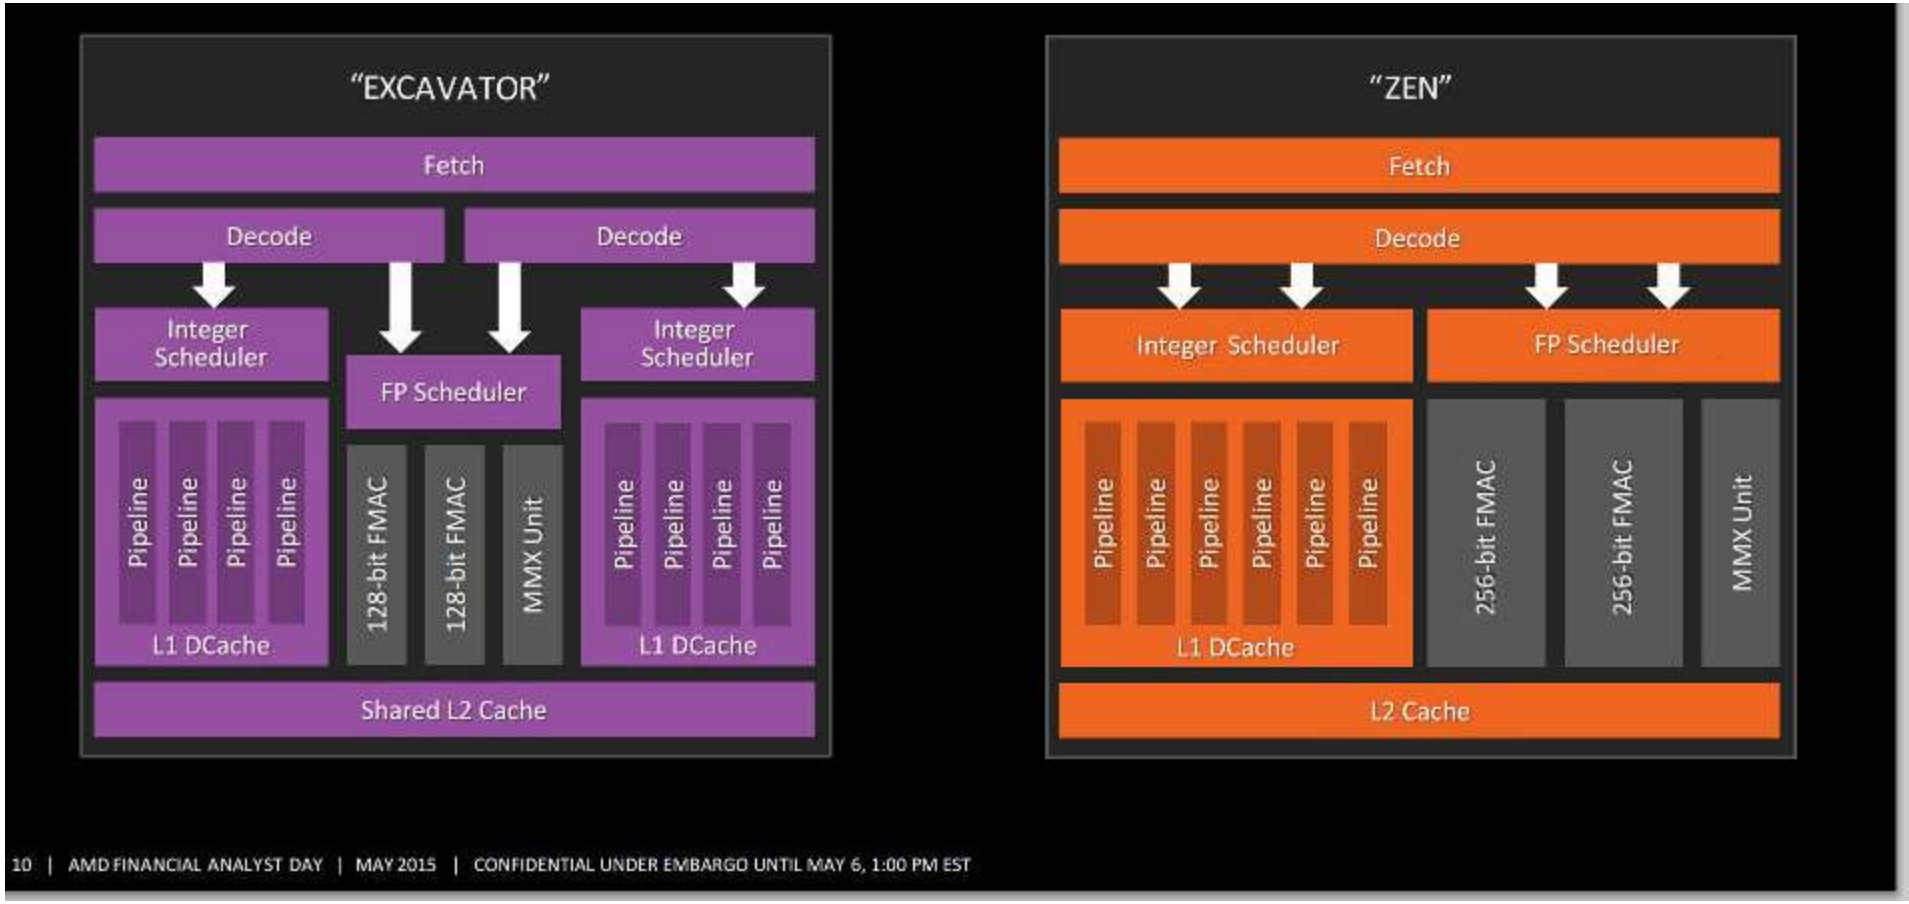

*

IPC improvements of the Bulldozer Family and the Zen architecture [15]

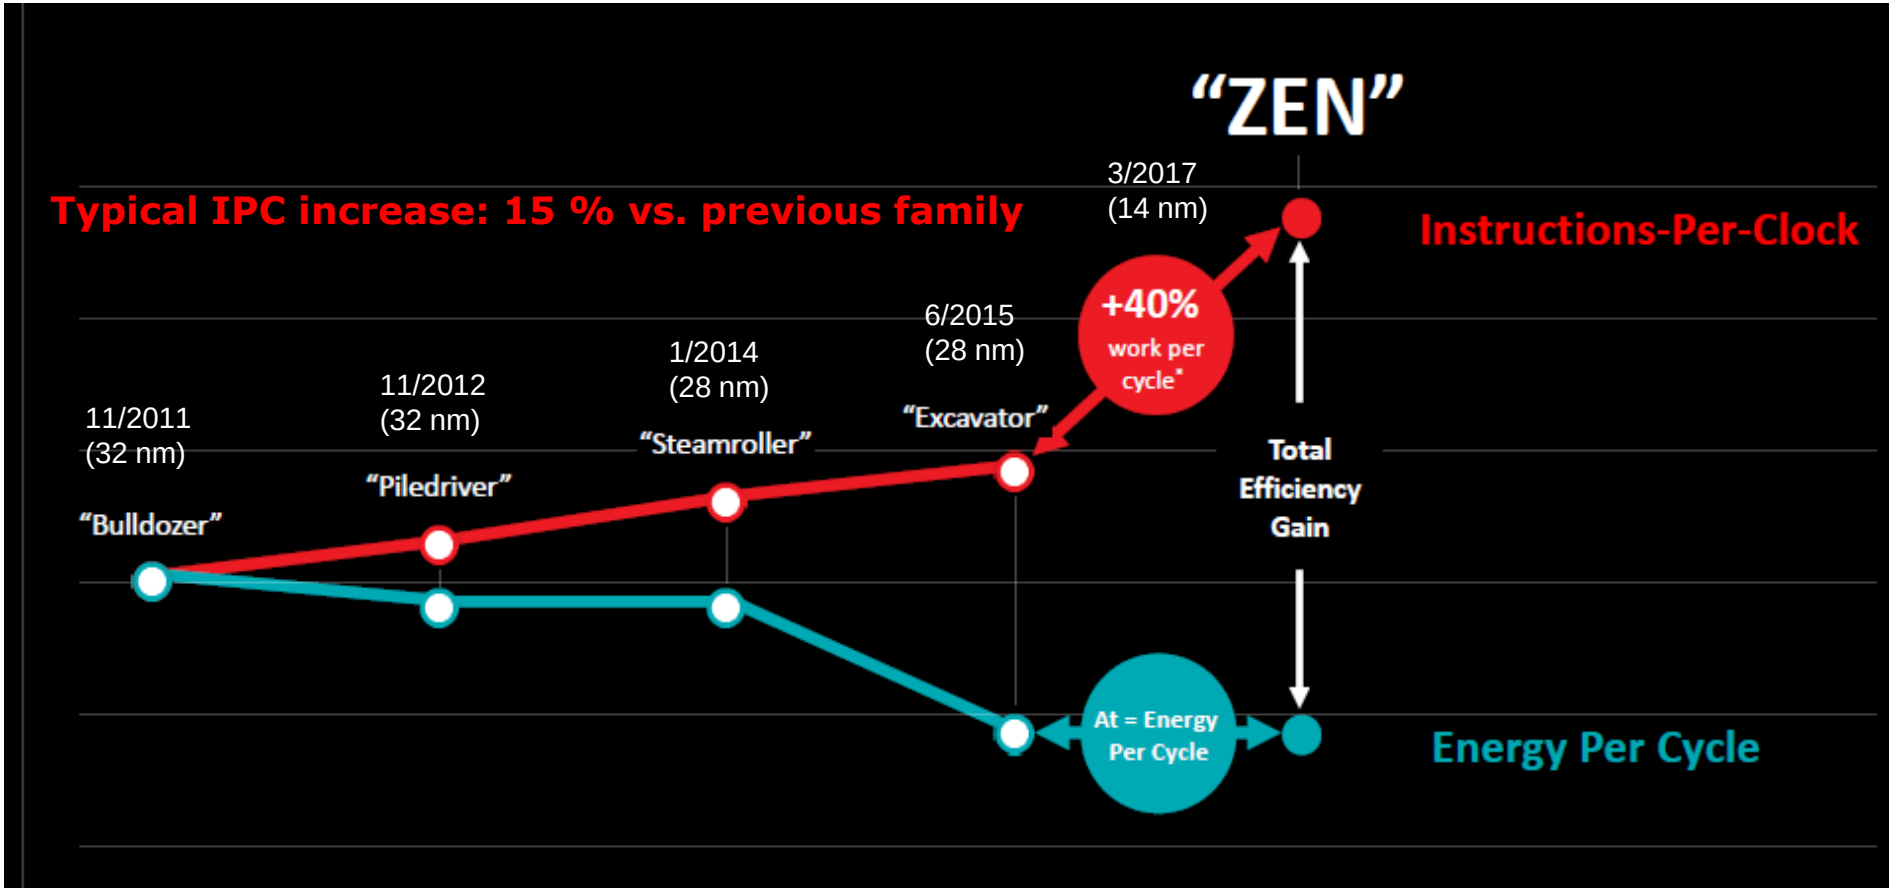

2.1.2 Innovations and enhancements introduced to raise IPC

2.1.2 Innovations and enhancements introduced to raise IPC

Innovations

a) Neural net improved branch predictor (Part of SenseMI)
b) Smart prefetch (Part of SenseMI)
c) Large Micro-Op cache (2K instructions)

Enhancements

d) Wider µop dispatch
- Six-issue FX execution (up from 4)
- Quad-issue FP execution (up from 3)

e) Further improvements of the microarchitecture

*

a) Neural-net improved branch predictor (Part of Sense MI) [10]

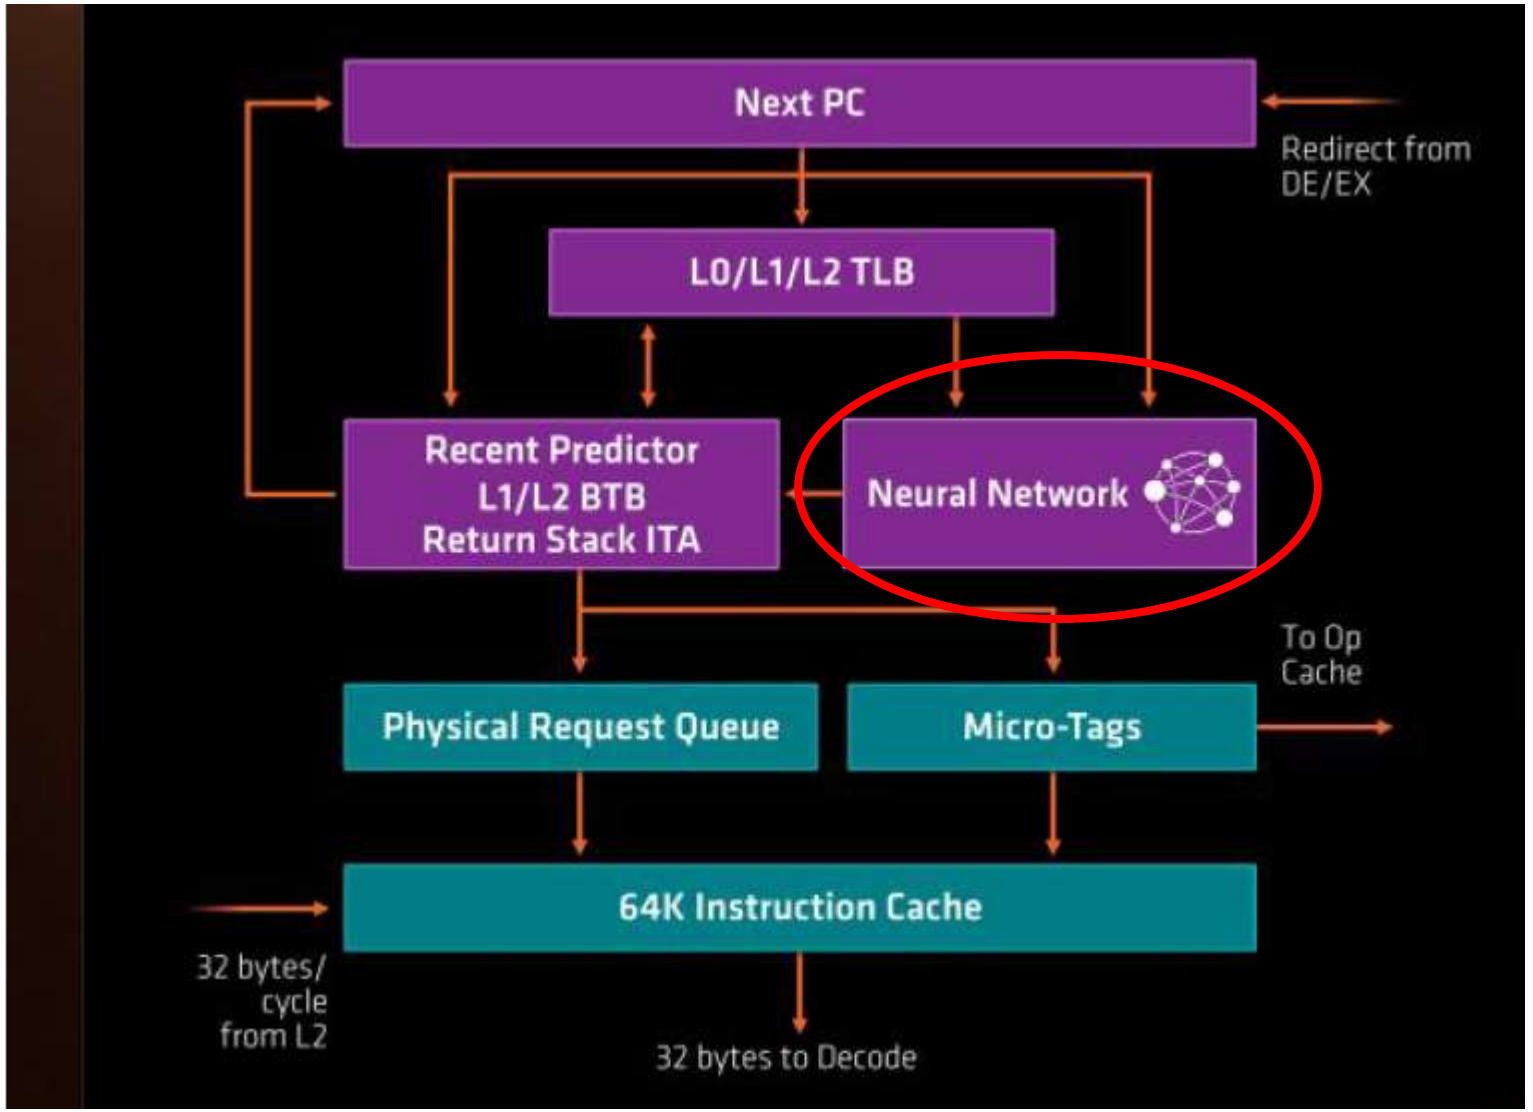

b) Smart prefetch (Part of Sense MI) [10]

Smart Prefetch

The Right Data At The Right Time

- Anticipates the location of future data accesses by application code
- Sophisticated learning algorithms model and learn application data access patterns
- Prefetches vital data into local cache so it's ready for immediate use

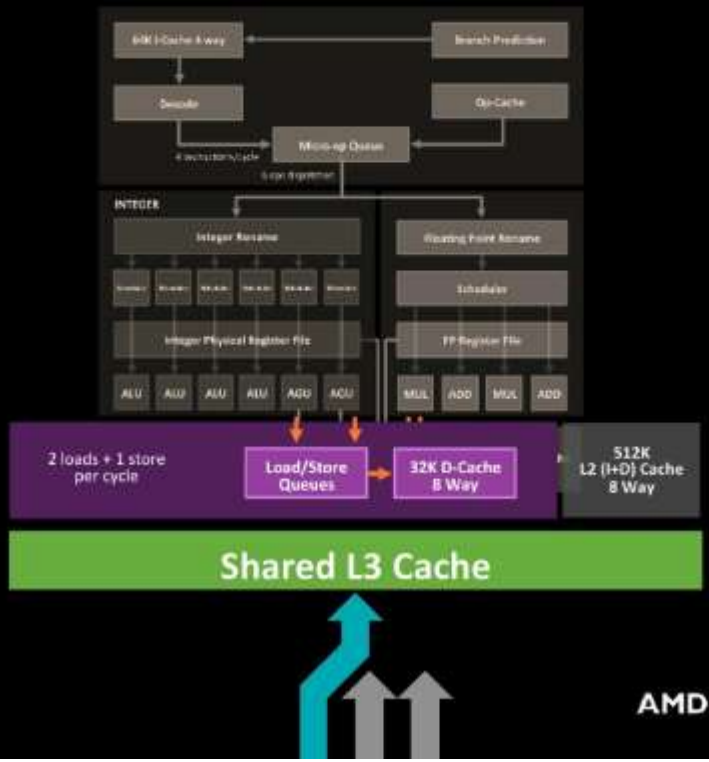

AMD | ZEN

c) Large Micro-Op cache (2K instructions) [15]

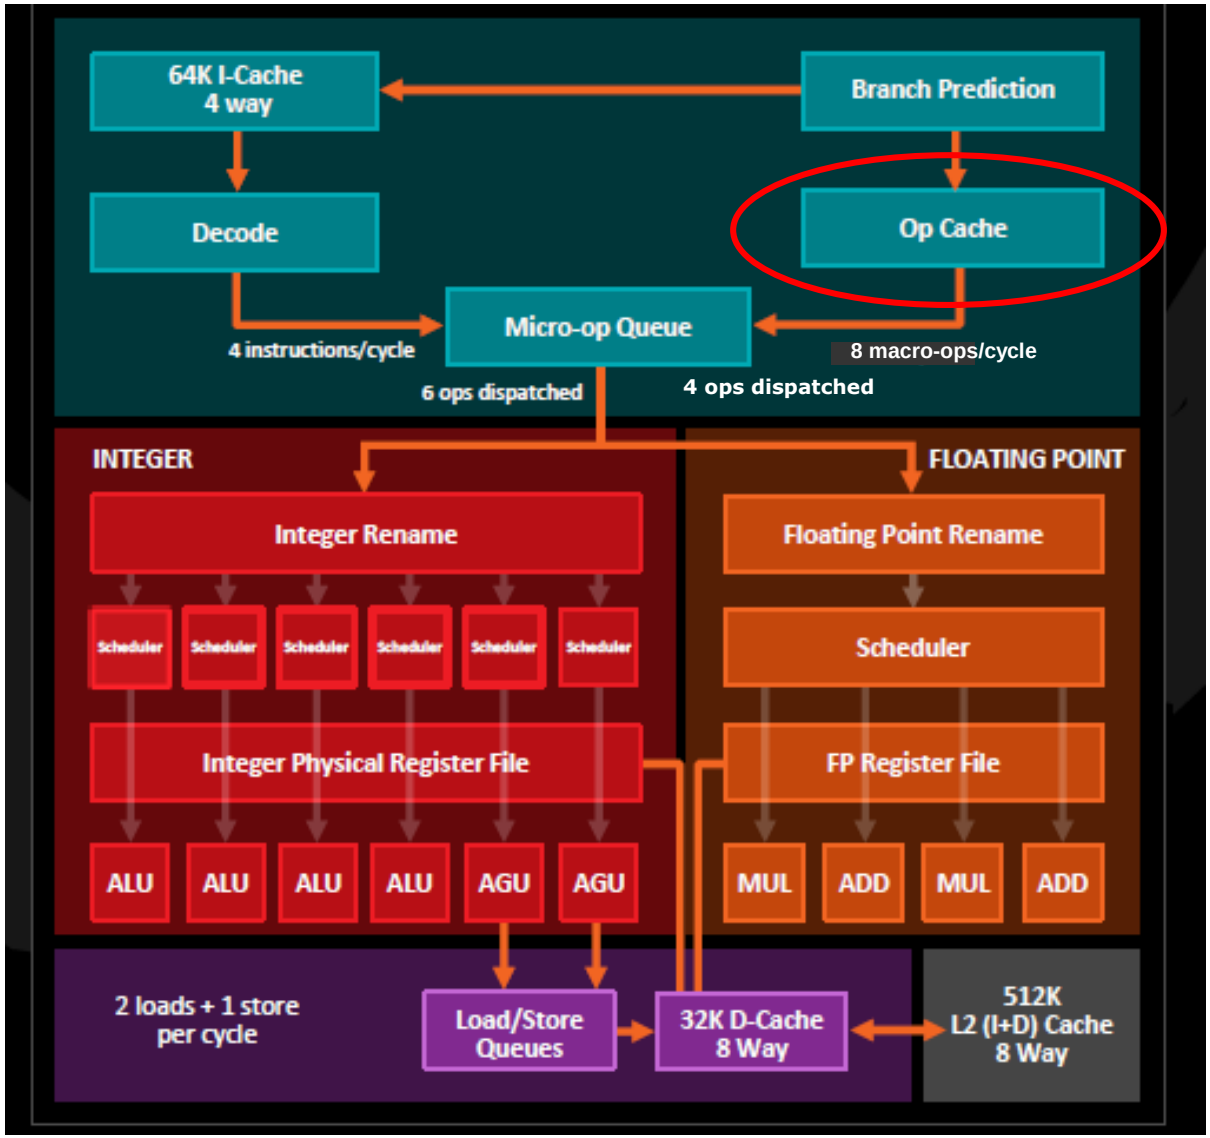

By contrast: The micro-op cache introduced in Intel's Sandy Bridge [90]

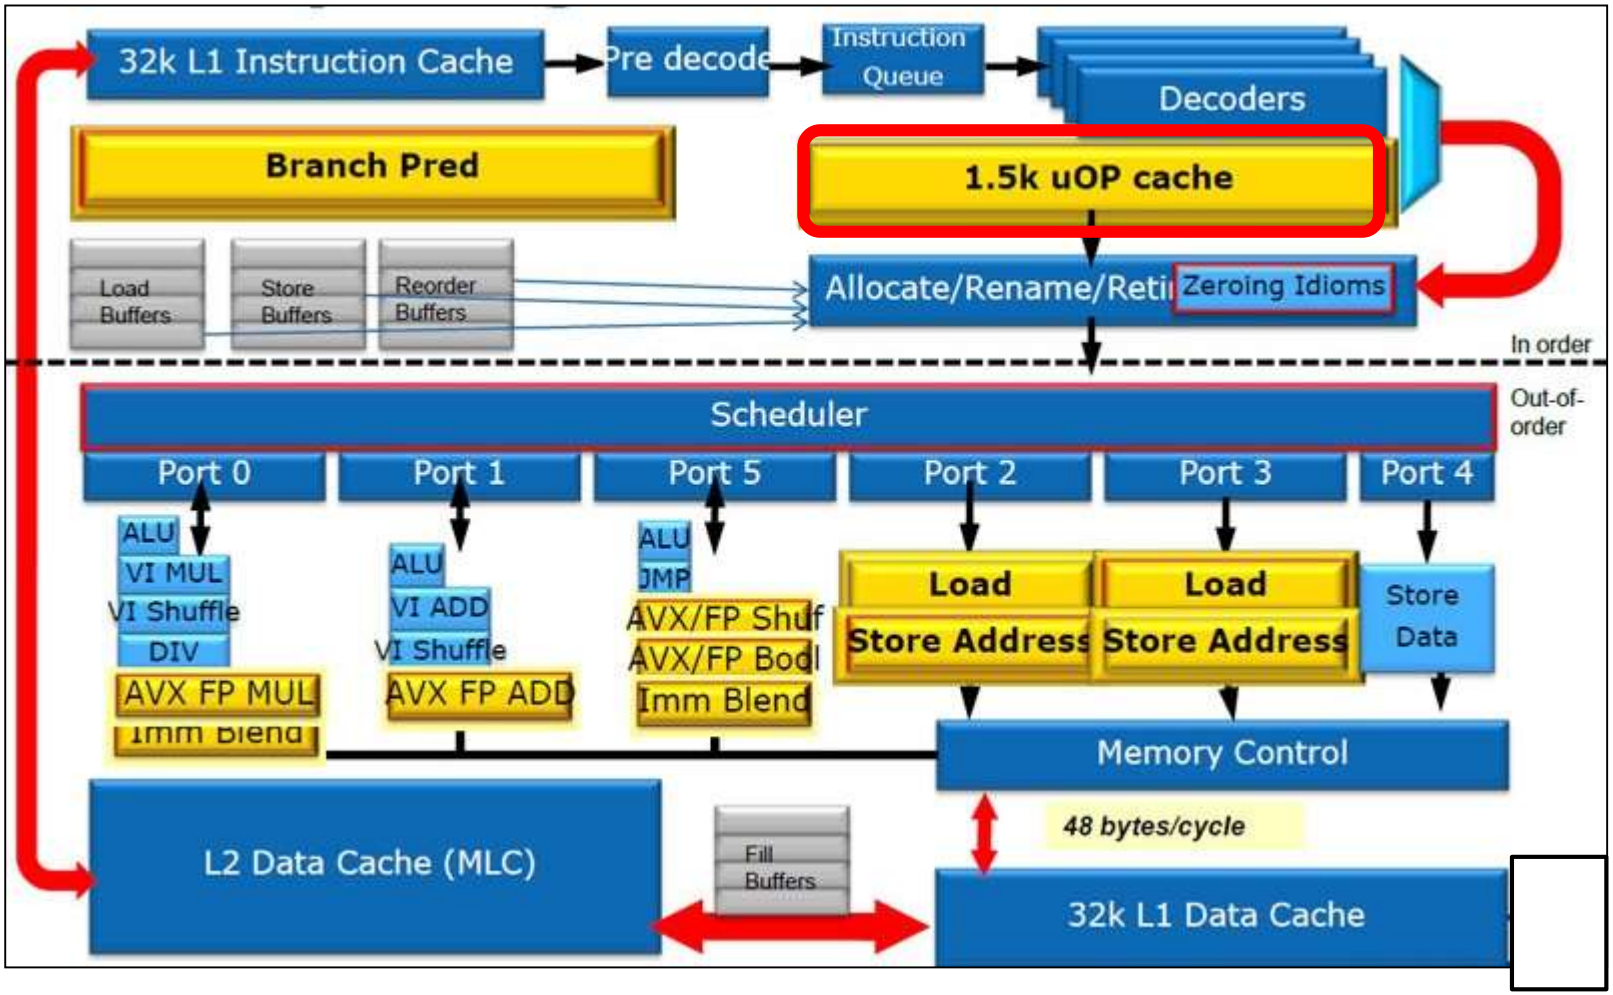

2.1.2 Innovations and enhancements introduced to raise IPC (6)

d) Wider μop dispatch

- Six-issue FX execution (up from 4) - Quad-issue FP execution (up from 3)

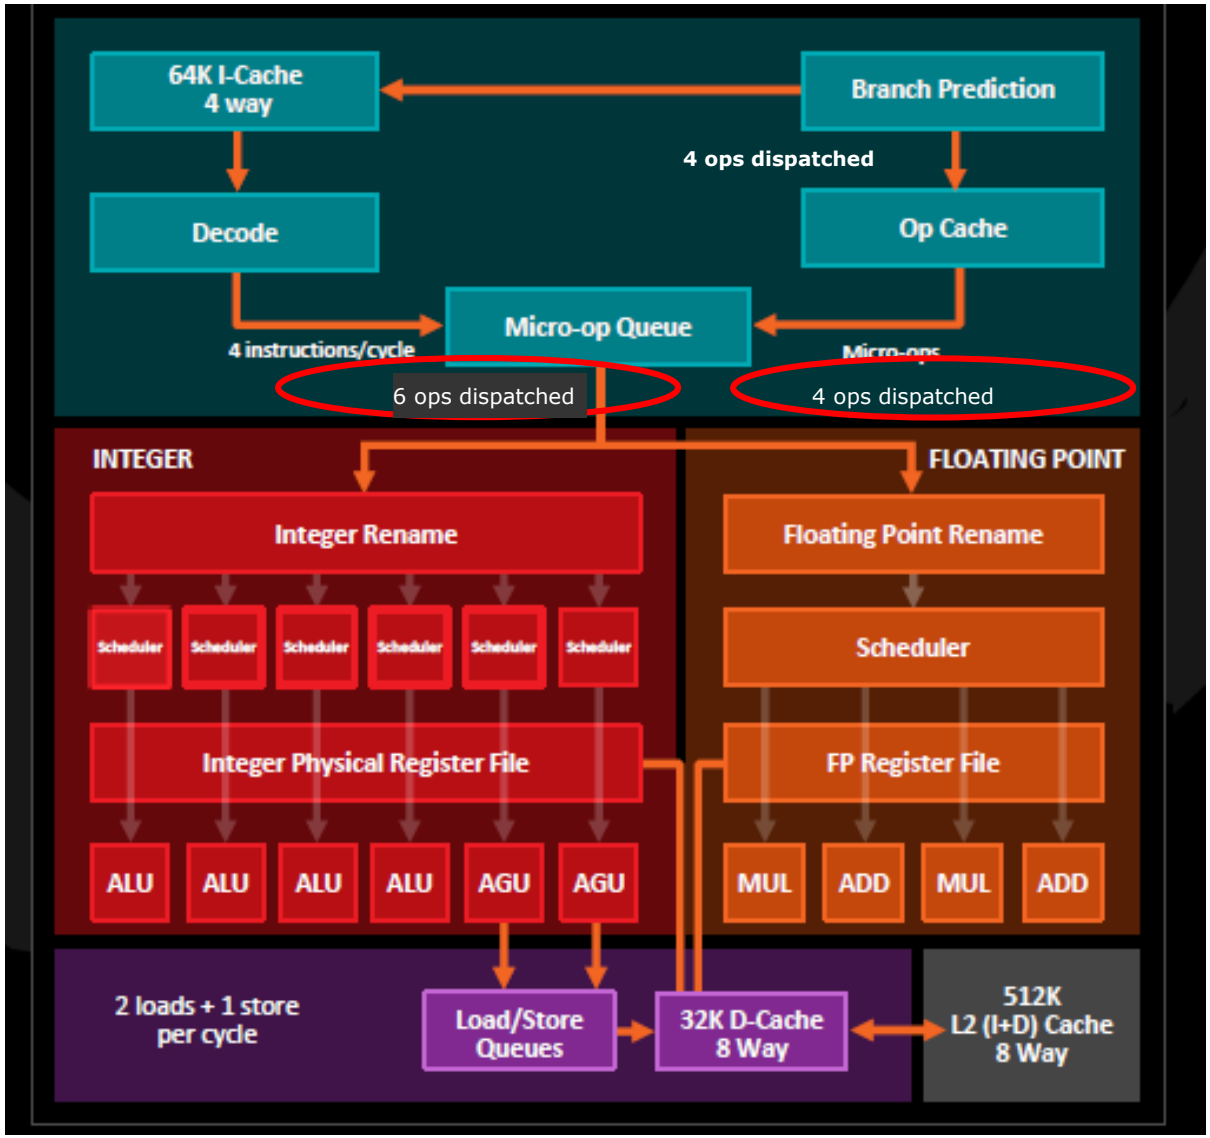

Comparison: Decoding in Intel's Skylake microarchitecture

5-wide decoding, as shown below.

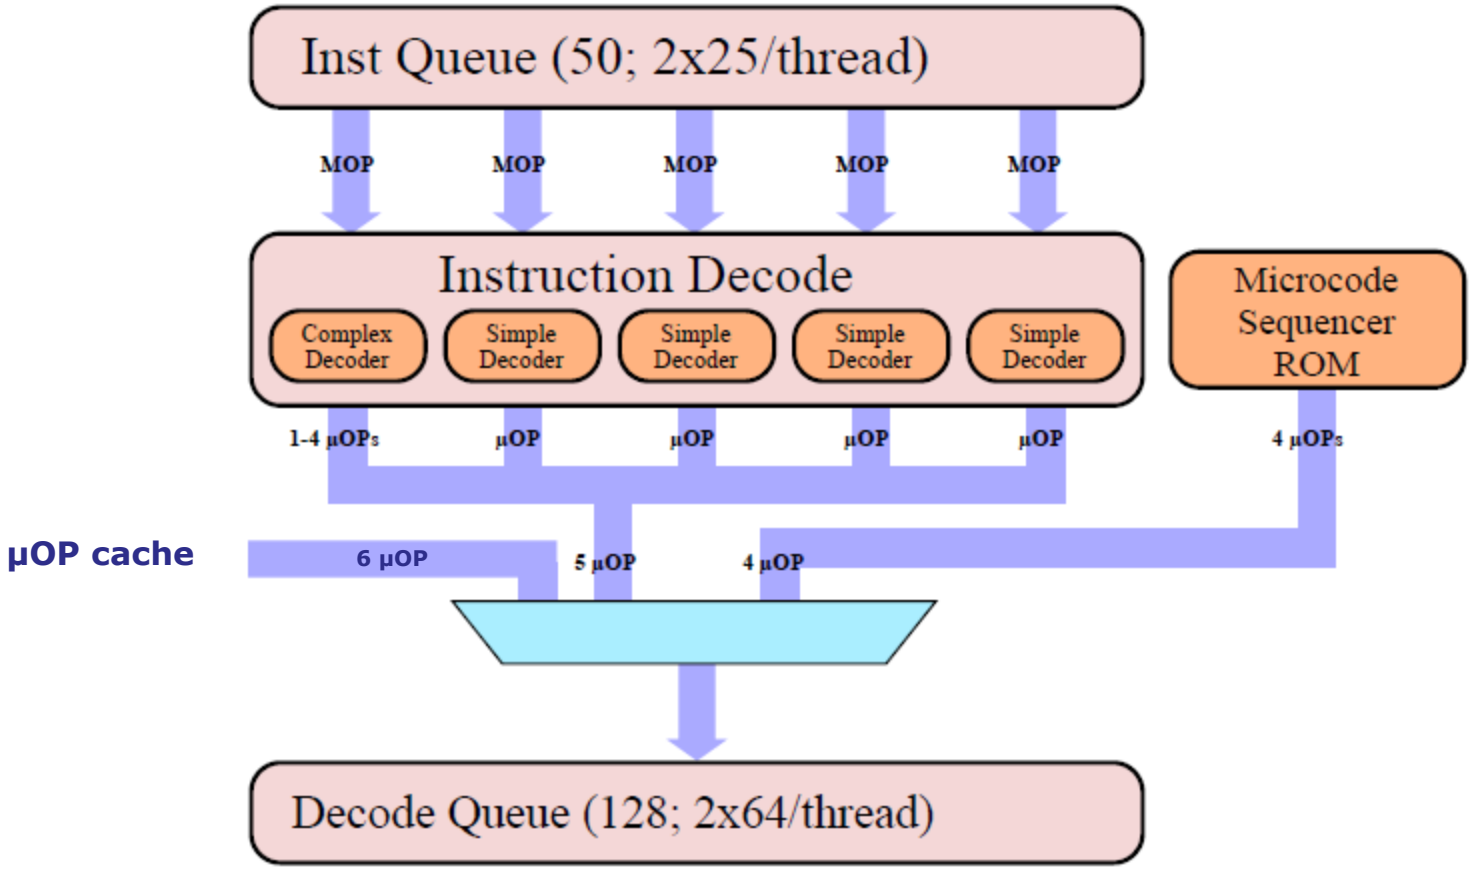

Figure: Decoding in Intel's Skylake's core [35]

e) Further improvements of the microarchitecture of Zen

Feature	Zen	Excavator
ROB size (duplicated for each thread)	192	128
ROB throughput	8 instr./cyle	4 instr./cycle
Integer scheduler size	84	48
FP scheduler size	96	60
Load Queue	72	44
Store Queue	44	32

2.1.3 Power management in Zen cores

It will be omitted

2.1.3 Power management in Zen cores

As part of the SenseMI technology suite AMD introduced the following three new technologies aiming at the power management:

- Pure Power technology (actually AVS)
- Precision Boost technology (actually Turbo Boost similar to Intel's SpeedShift)
- XFR (eXtreme Frequency Range) technology (actually configurable TDP)

to be discussed next.

*

2.1.3.1 Pure Power technology

- It is basically an AVS (Adaptive Voltage Scaling) technology.
- AVS is a kind of AVFS (Adaptive Voltage and Frequency Scaling) technology.
- AVFS is an enhancement of the DVFS (Dynamic Voltage and Frequency Scaling).

*

DVFS (Dynamic Voltage and Frequency Scaling)

- Scaling means here controlling core voltage (Vcc) and frequency (fc) according to the actual workload.
- With DVFS the OS typically assesses the actual workload and scales down fc as far as feasible without noticeable lengthening the runtime, then it selects the Vcc value that is needed to maintain the chosen fc from a look up table.

  Vcc is determined a priory at chip testing while a guard band is also provided.

Figure:
Principle of DVFS

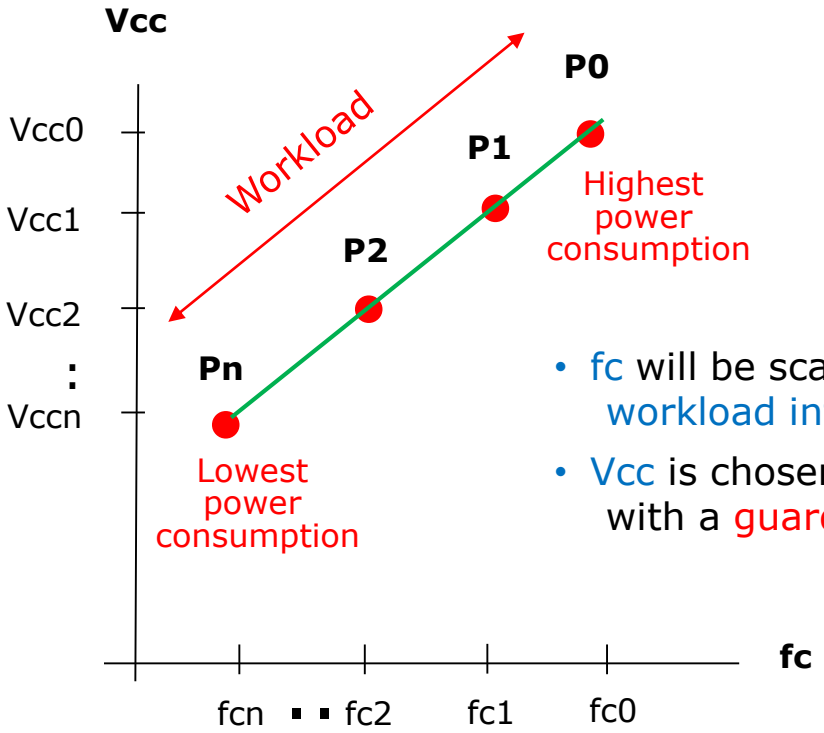

- fc will be scaled according to the workload intensity.
- Vcc is chosen to maintain fc with a guard band.

*

2.1.3 Power management in Zen cores (4)

Example: Minimal (measured) and specified core voltage (Vcc) needed for sustaining a given core frequency (fc) [61]

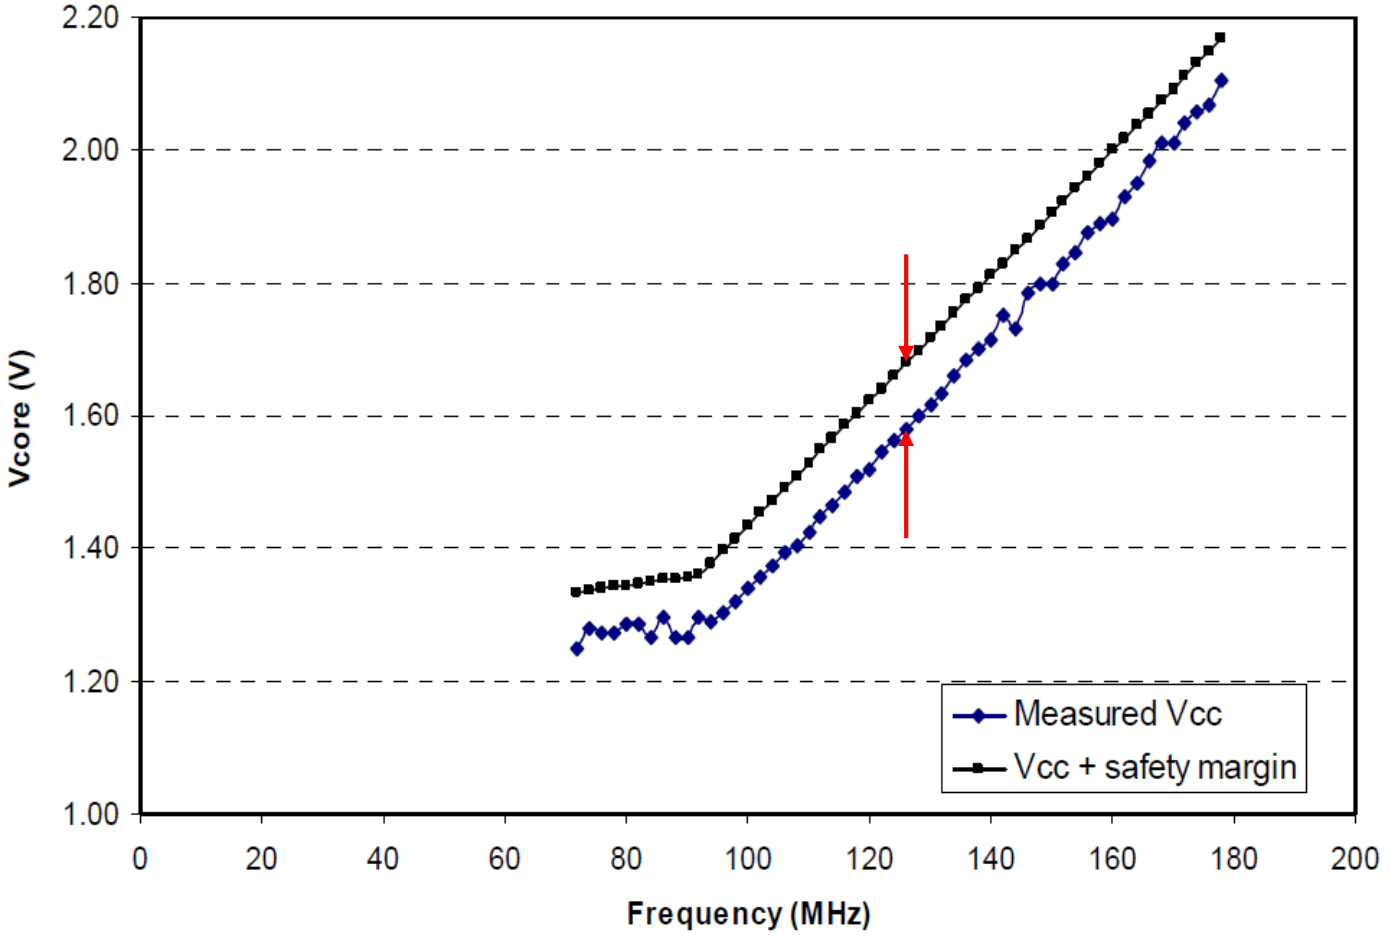

AVFS vs. DVFS

DVFS

(Dynamic Voltage and Frequency Scaling)

DVFS is performed in an open loop.

It means that the OS scales down
core frequency (fc)
according to the actual workload
and determines core voltage (Vcc)
statically (from a look-up table).

Vcc was determined previously
at chip testing while
a certain guard band
was taken into account
to compensate for PVT variations.

AVFS

(Adaptive Voltage and Frequency Scaling)

Aim: to reduce guard band and thus dissipation

AVFS is performed in a closed loop.

It means that first the OS scales down
core frequency (fc) and voltage (Vcc)
according to the actual workload
similarly to DVFS, but then the processor
strongly reduces the guard band
based on monitoring the actual parameters, like
speed, dissipation, die temperature etc..

PVT: Process, Voltage, Temperature

Illustration of AVS and AFS

Use cases of AVFS to exploit guard band reduction

AVS
(Adaptive Voltage Scaling)

It is also called **Undervolting**
(Called **Poor Power** by AMD (Zen)

AVS is used to reduce Vcc
at a given fc (P-state)
to a min. allowed value,
to reduce power consumption.

AFS
(Adaptive Frequency Scaling)

It is also called **Overclocking**
(Called **Precision Boosting** by AMD (Zen)

AFS is used to increase fc
at a given Vcc (P-state)
to the max. allowed value
to raise performance.

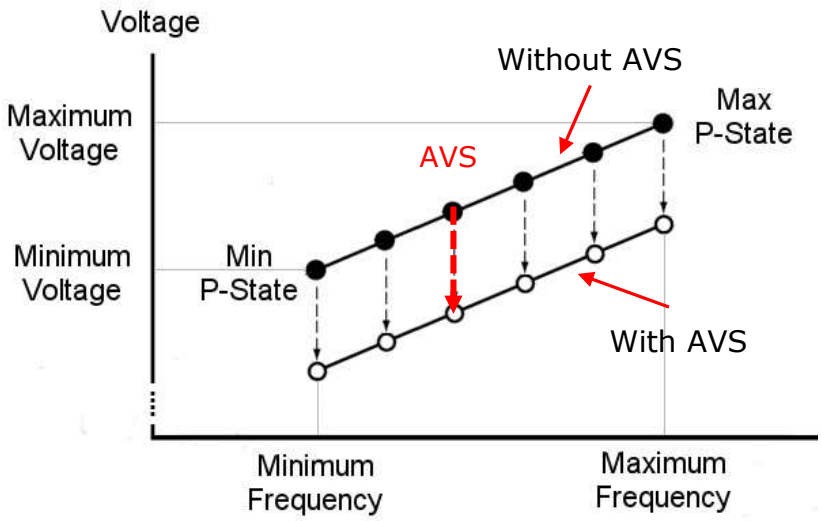

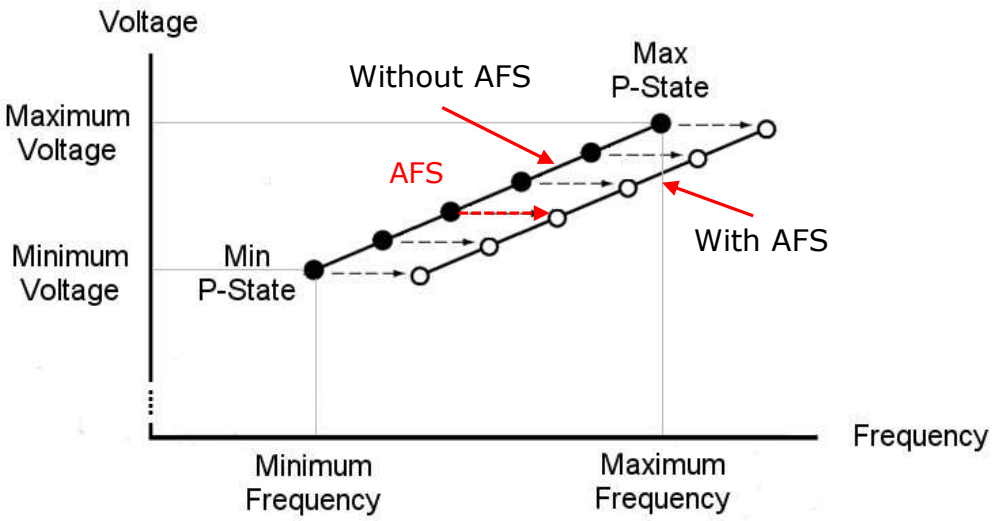

*

Overview of introduced AVFS implementations

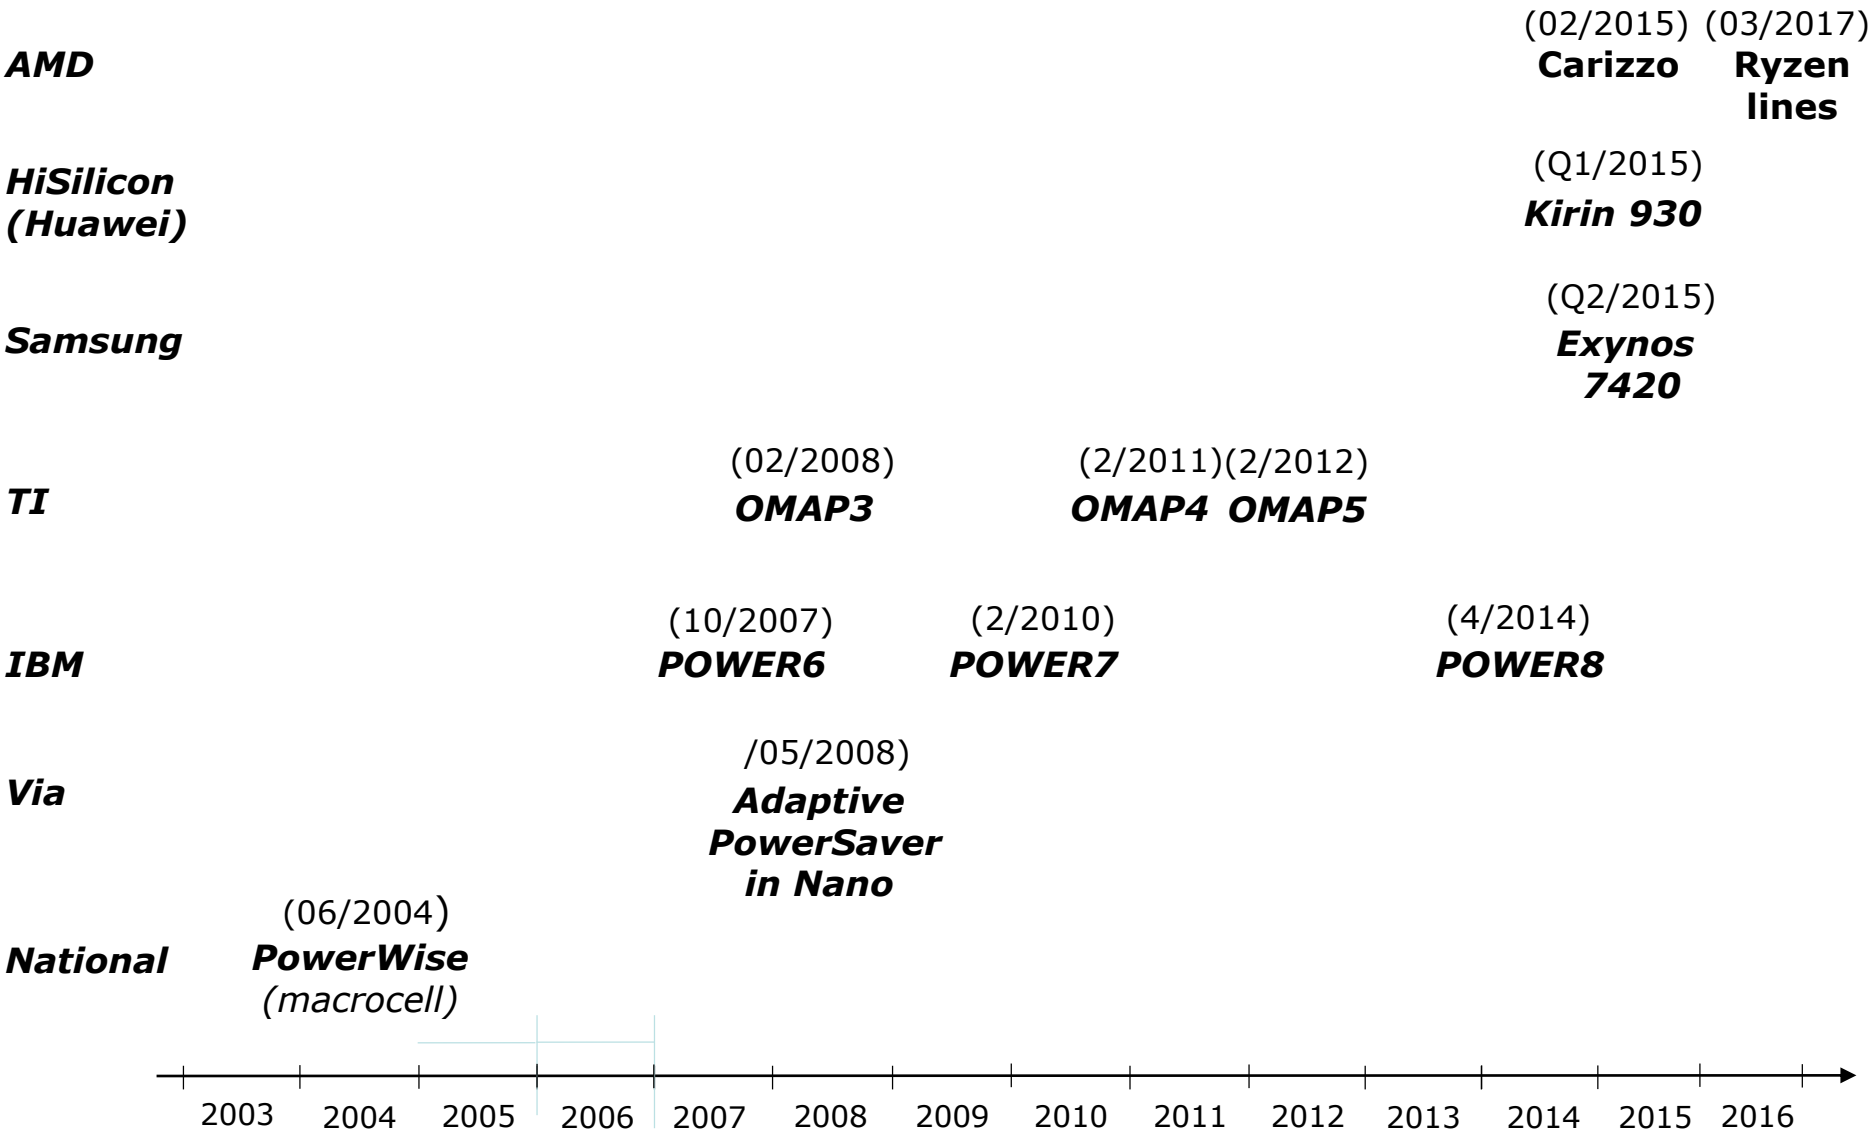

Principle of AMD's Pure Power technology that implements AVS [10]

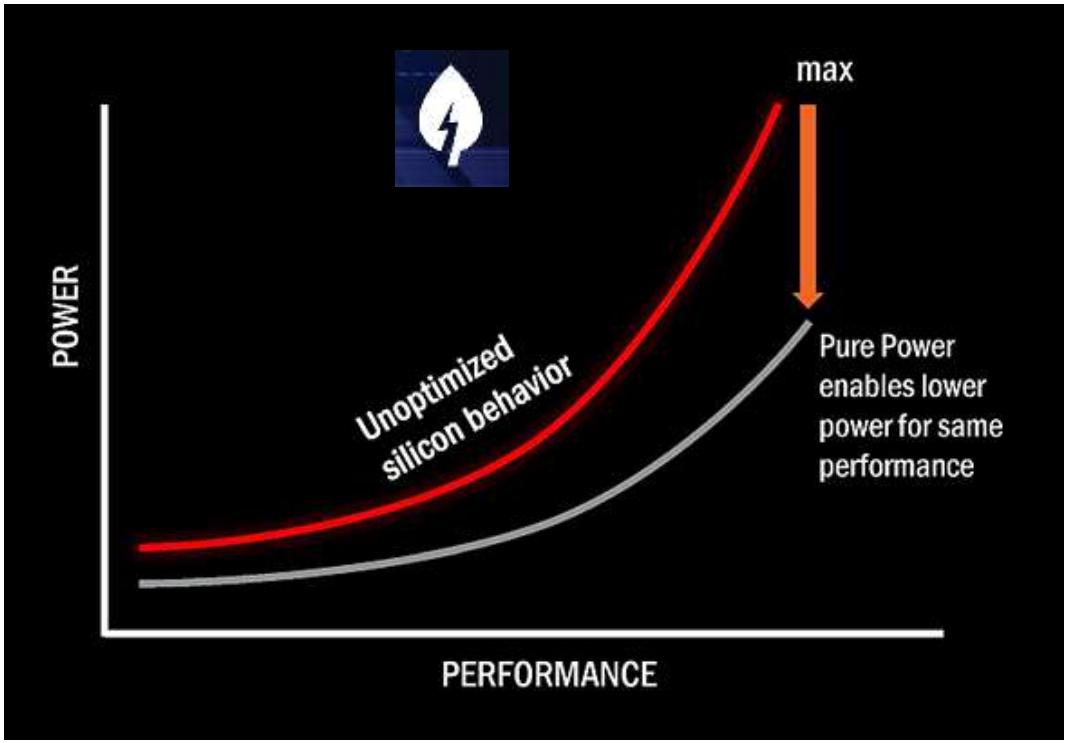

Using Energy Wisely

- Monitors temperature, speed and voltage
- Adaptive control manages real time for lower power usage
- Closed loop control with Infinity Fabric

Key points of the implementation of AMD's Pure Power (actually AVS)

a) Real time monitoring of temperature, speed and voltage by sensors,
b) Adaptive closed-loop control
c) Per-core scaling of core frequency and voltage

*

2.1.3 Power management in Zen cores (10)

a) Real time monitoring of temperature, speed and voltage by sensors

1000+ embedded sensors per CCX monitor temperature, speed and voltage per ms.

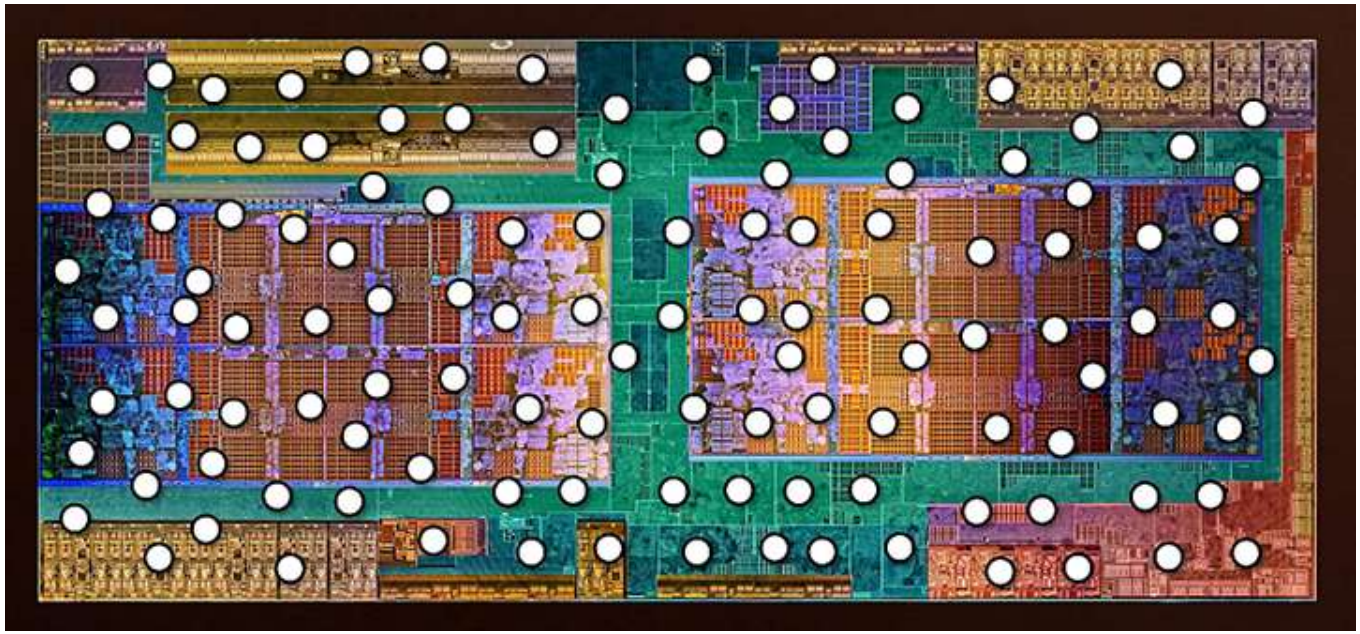

Figure: Distributed embedded sensors on a Ryzen die (including two CCXs) [10]

Sensors implemented on a CCX die to support AVS [9]

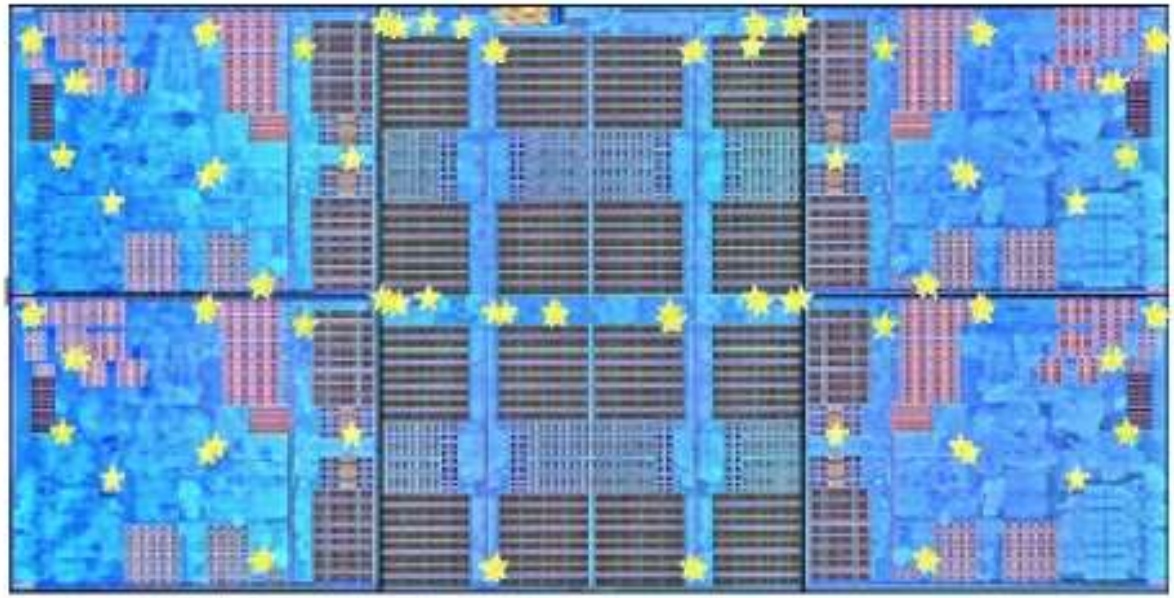

- >1300 critical path monitors
- 48 on-die high speed power supply monitors
- 20 thermal diodes
- 9 high speed droop detectors

Critical path monitors overlook processor speed (clock frequency), they allow to reduce core voltage in a controlled way as far as possible whereby a given clock frequency remains maintained.

Droops are rapid spikes in the voltage supply, they need to be addressed to avoid erroneous operations (not detailed here).

Critical path monitors: kritikus áramköri szakaszok sebességét ellenörző monitorok

b) Adaptive closed-loop control -1

The Infinity System's Scalable Control Fabric provides the required infrastructure for AVS, as seen below.

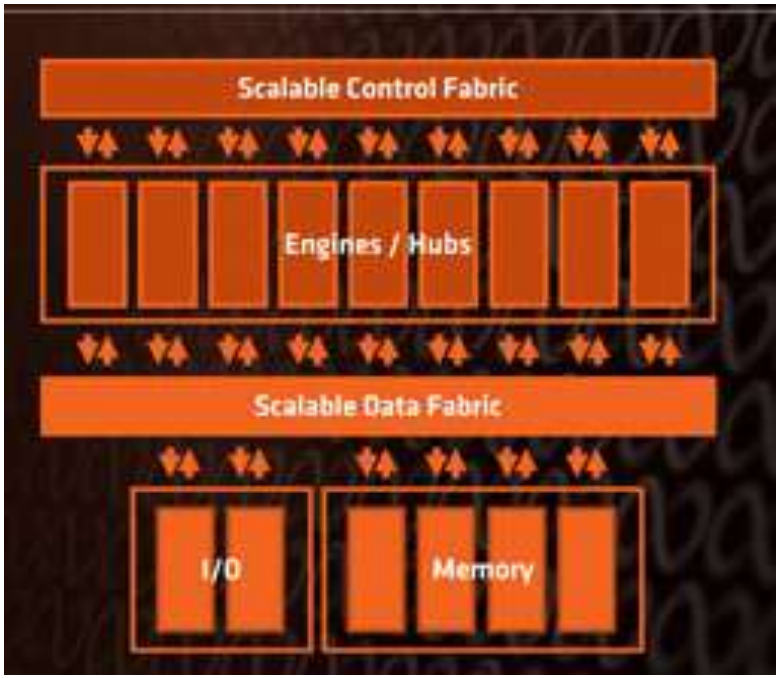

Figure: The Infinity Fabric [38]

b) Adaptive closed-loop control -2

There is closed loop control for implementing AVS, as indicated below.

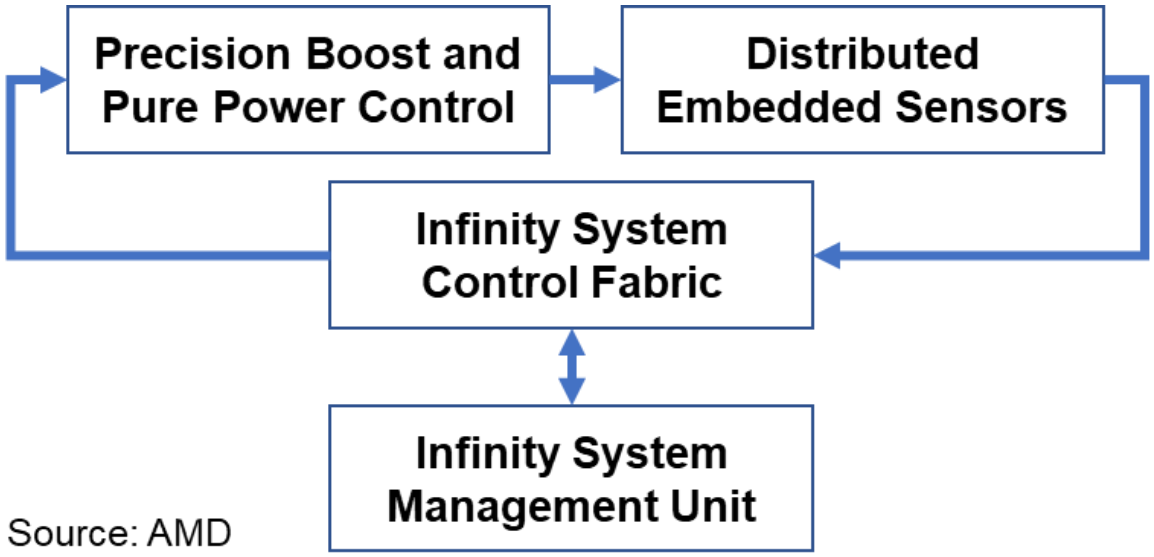

Figure: The closed loop control mechanism for Pure Power (i.e. AVS) and Precision Boost (i.e. Turbo Boost) [3]

b) Adaptive closed-loop control -3

- An Infinity System Management Unit (based on the SMU microcontroller) is used for system settings according to user requirements, as the Figure below shows.
- The SMU (System Management Unit) microcontroller allows to chose the proper balance of performance and power consumption.

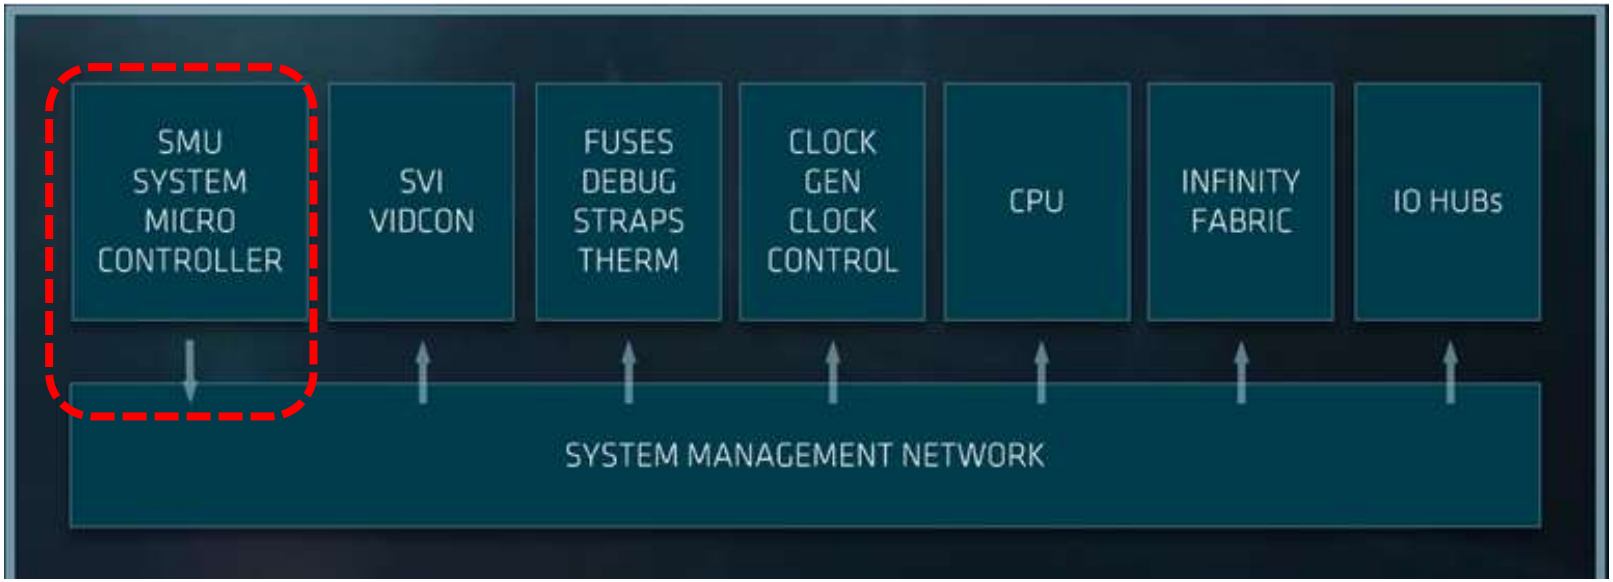

Figure: Infinity System Management [3]

SVI: Serial VID Interface

VID: Voltage Identifier (transferred to the VR (Voltage Regulator)

*

c) Per core scaling of core frequency and voltage [16]

- Each core has
  - a digital Low Drop-Out (LDO) Voltage Regulator (VR) and
  - a digital frequency synthesizer (DFS)

  to independently vary core frequency and voltage for reducing power consumption.
- The voltage regulators operate also as power gates.

Remark

LDOs are one type of linear regulators, e.g. as seen below [39].

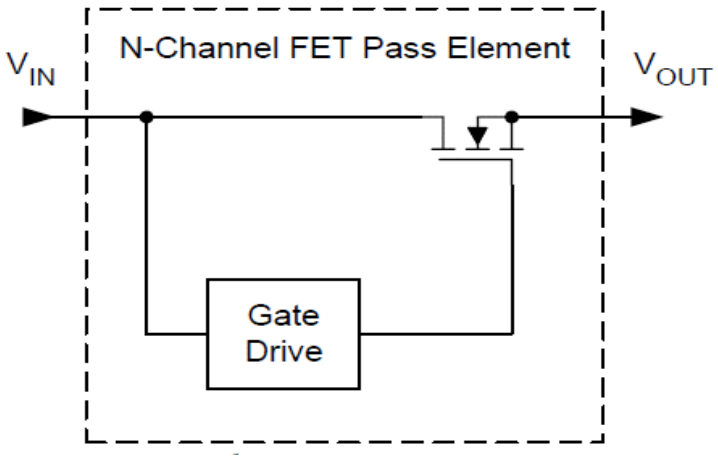

FET: Field Effect Transistor (térvezérlésű tranzisztor)

*

Implementing separate VRs for each core [50]

- Two stage voltage regulation (similar to Intel's two stage FIVR implementation):
  - 1. stage: on the mainboard, it reduces voltage from 12 V to 0.9 V,
  - 2. stage: separate on-die VRs for each core.
- Example for per-core voltage regulation in a Ryzen desktop processor:

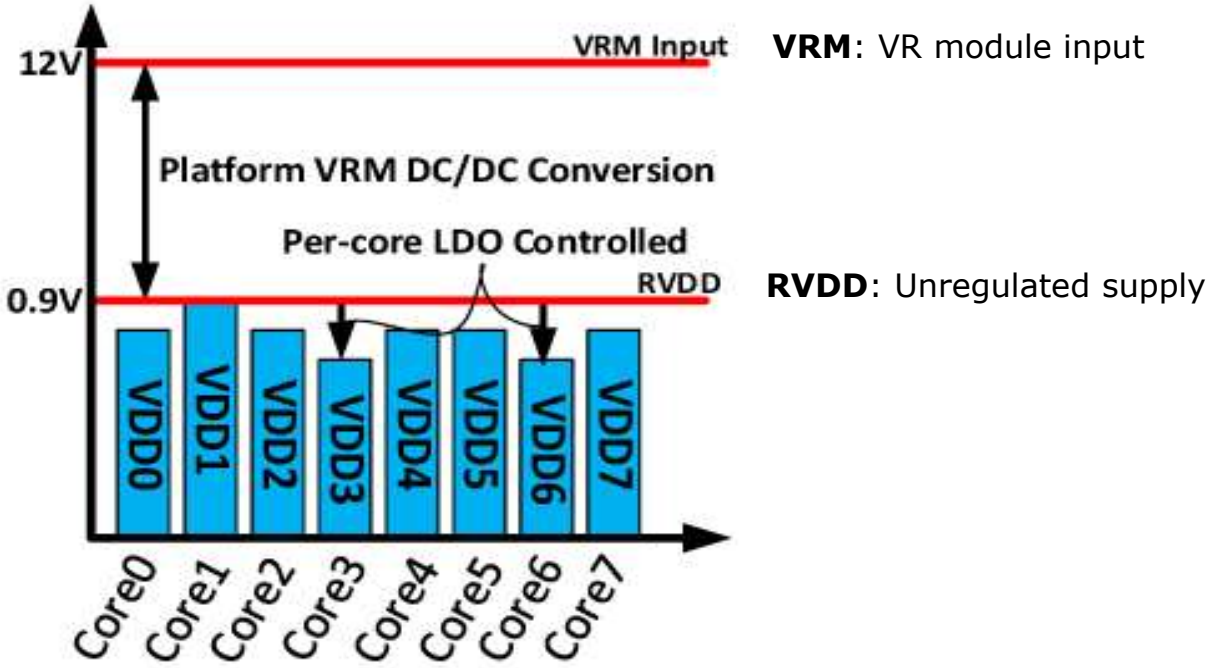

Figure: Per-core voltage regulation in a Ryzen desktop processor [50]

*

2.1.3.2 Precision Boost -1

- Basically, it is AMD's Turbo Boost implementation.
- It makes use of the same closed loop control mechanism as AMD's Pure Power (that is AVS):

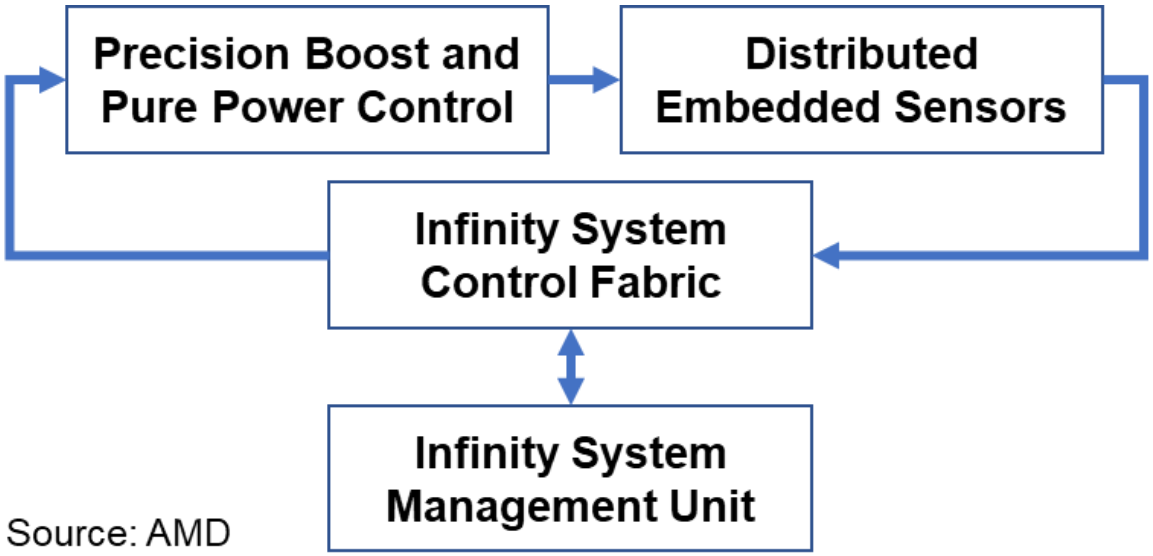

Figure: The closed loop control mechanism for Pure Power (i.e. AVS) and Precision Boost (i.e. Turbo Boost) [3]

2) Precision Boost -2

Specific features of the Precision Boost technology:

a) Specifying the Turbo Boost frequencies

b) Fine grained tuning of the clock frequency (in 25 MHz steps)

*

a) Specifying the Turbo Boost frequency [38]

Precision Boost makes use of two Turbo Boost frequencies,

- one higher value for the case if up to 2 cores are active (in DTs), and
- a lower value for more active cores,

as indicated in the Figure below.

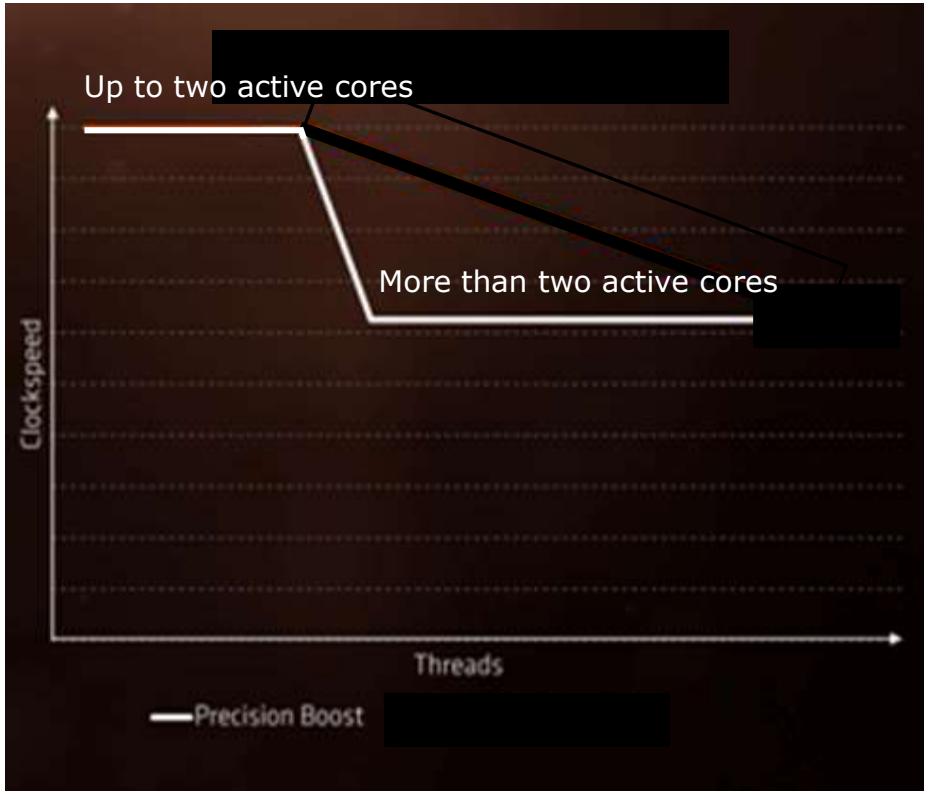

Figure: Boost behavior in the Ryzen line [38]

*

b) Fine grained tuning of the clock frequency (in 25 MHz steps) [10]

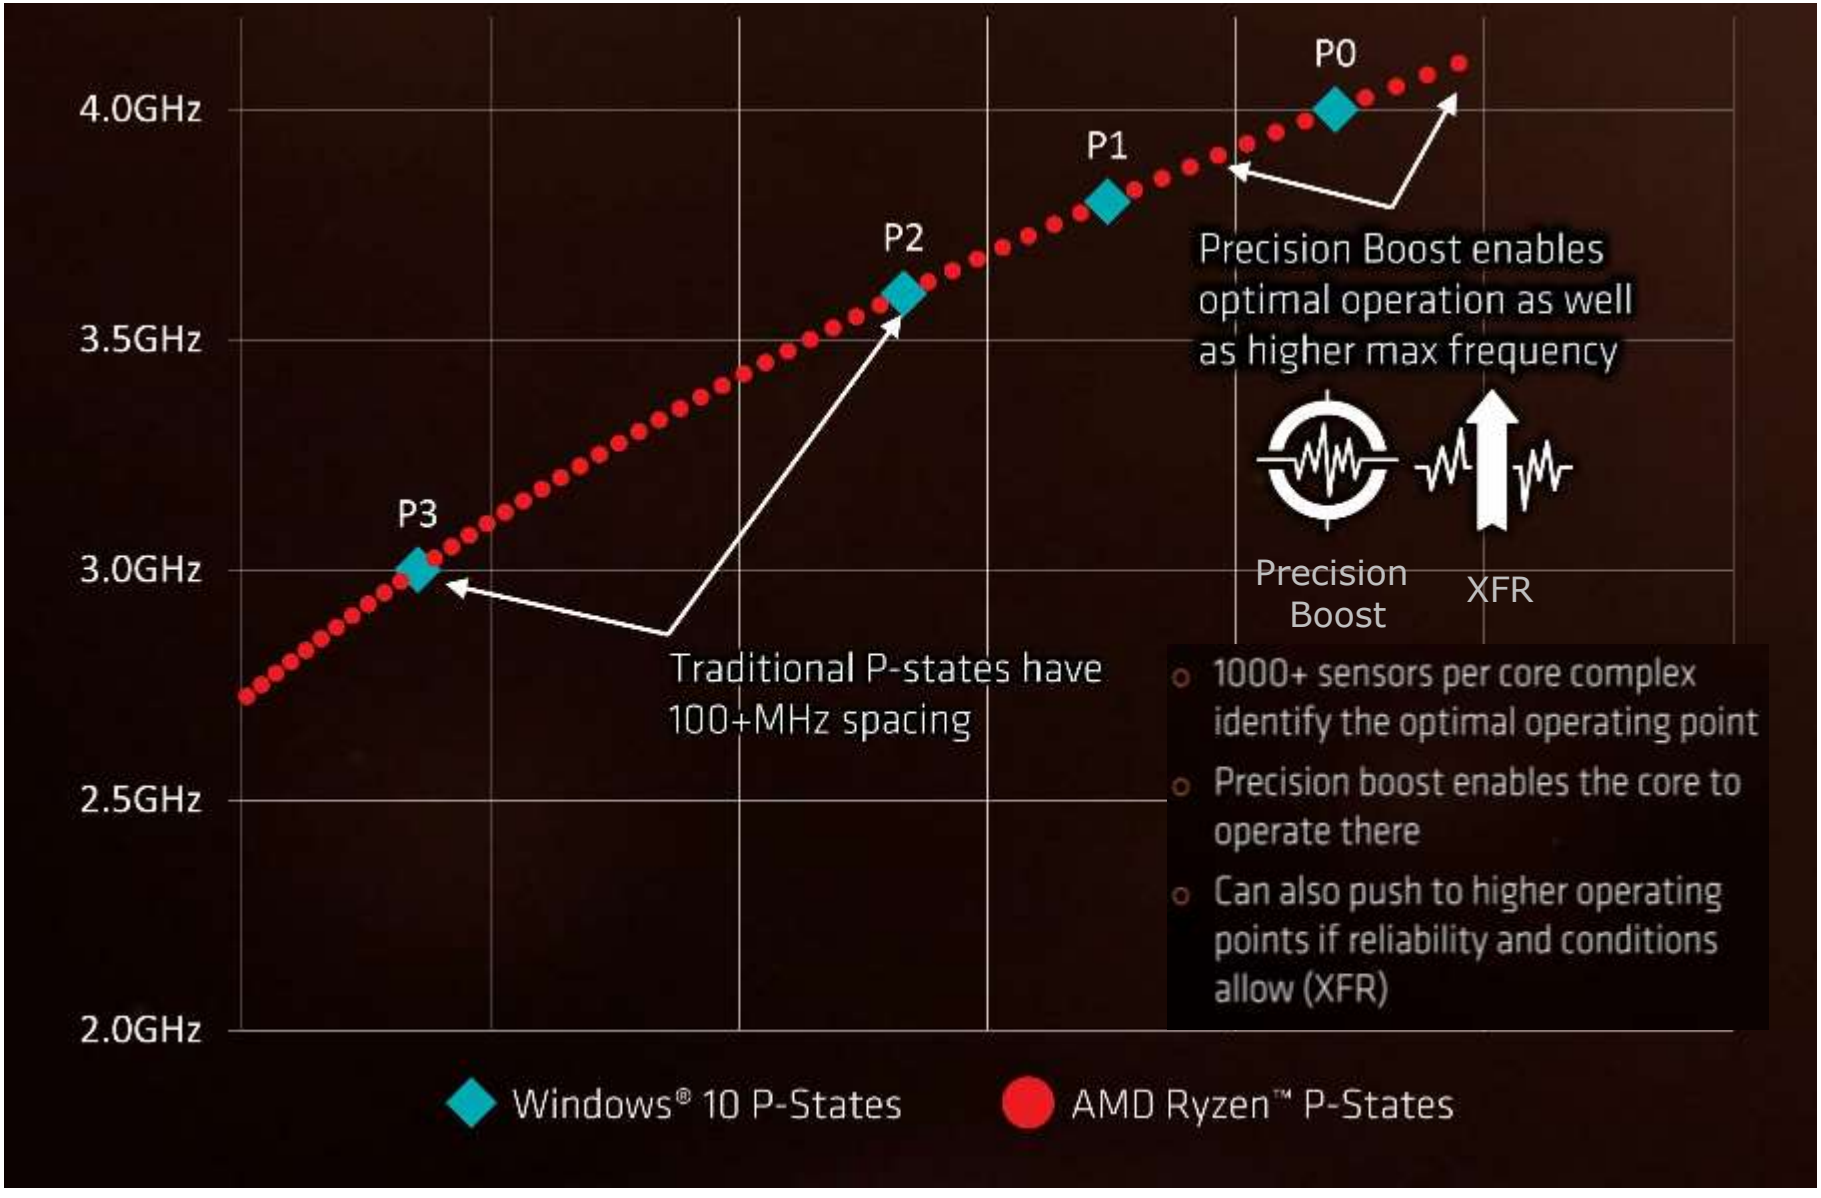

*

Assumed principle of fine grained tuning

- While using Precision Boost all active cores operate at the same clock frequency and core voltage.
- Fine grained tuning will be performed by the SMU microprocessor that raises the core voltage and the clock frequency (in 25 MHz steps) whereas checks whether the limiting factors, including the temperature, SoC power and VRM current (and with it the socket power consumption) remain in the factory boost range.

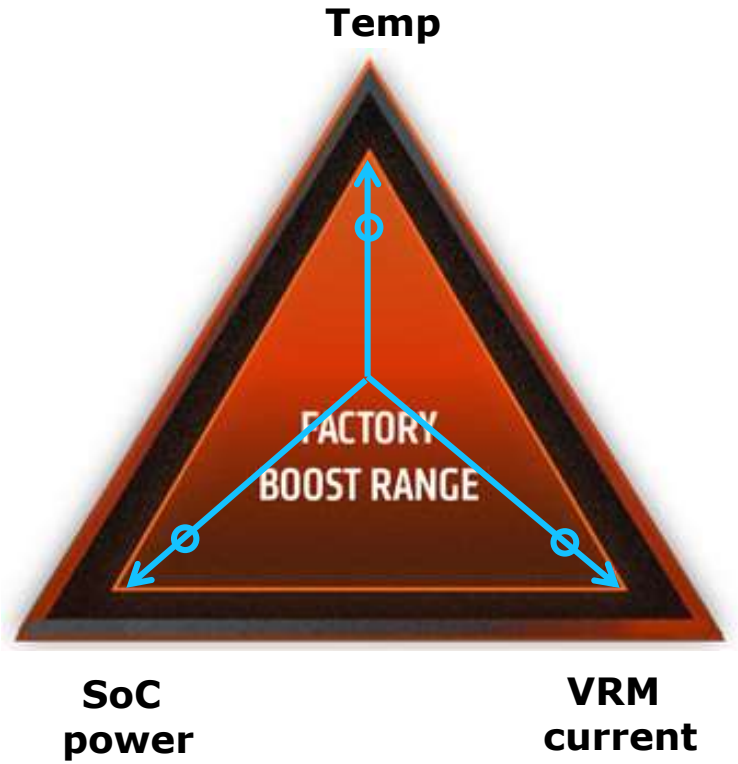

Figure: Conceptual view of extending the limiters of Precision Boost with PBO [82]

2.1.3.3 Raising the clock frequency over the Turbo boost frequency - XFR

The XFR (eXtended Frequency Range) technology presumes a high performance (premium) cooler system to keep die temperature within the given limit [10].

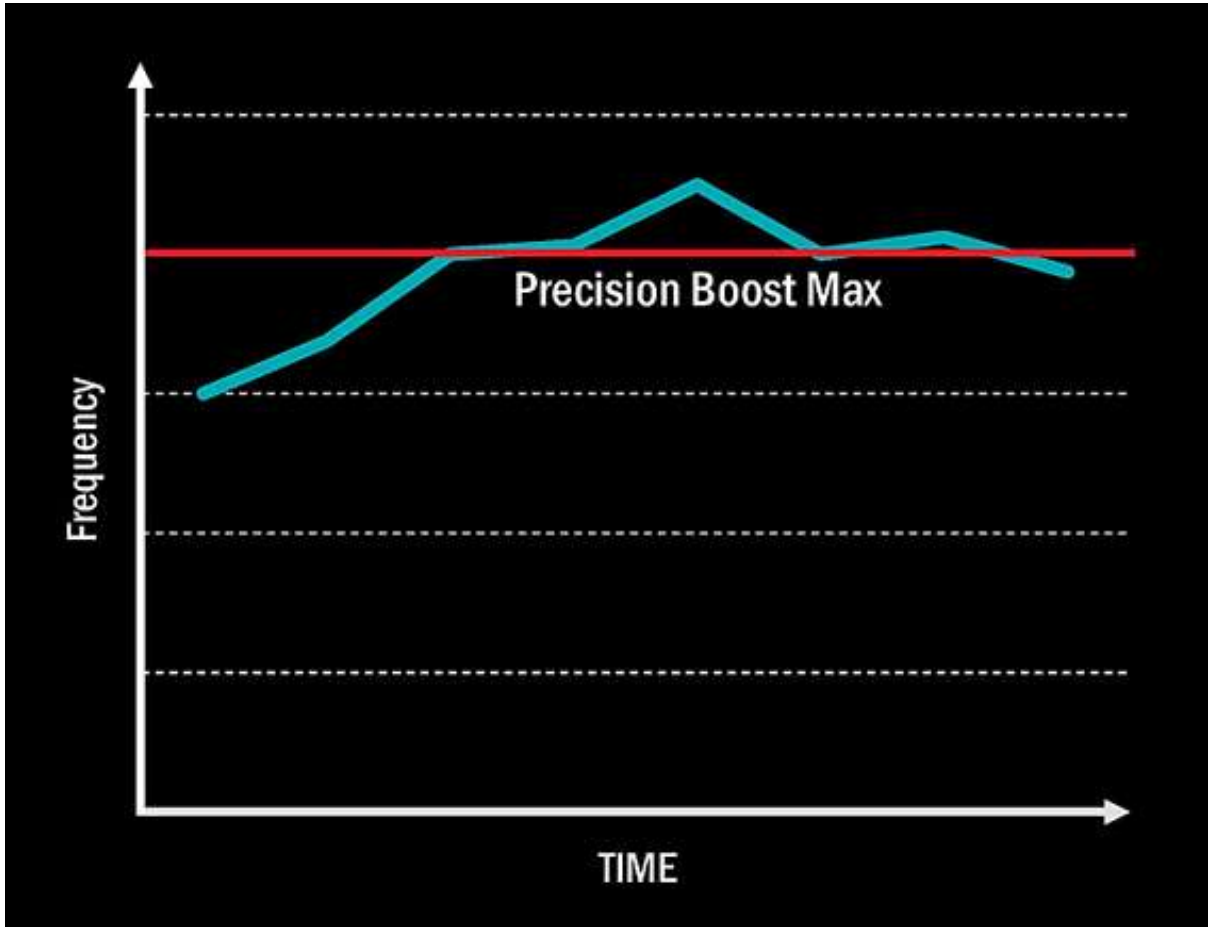

XFR provides an automated overclocking over the max. Turbo Boost frequency if there is a premium cooler system and the operating conditions of the processor (e.g. temperature data) allow it.

With XFR enabled the processor may boost clock frequency by 50 to 200 MHz above the advertised max. Turbo Boost frequency.

Figure: Principle of operation of XFR [10]

*

Illustration of the operation of the XFR technology [10]

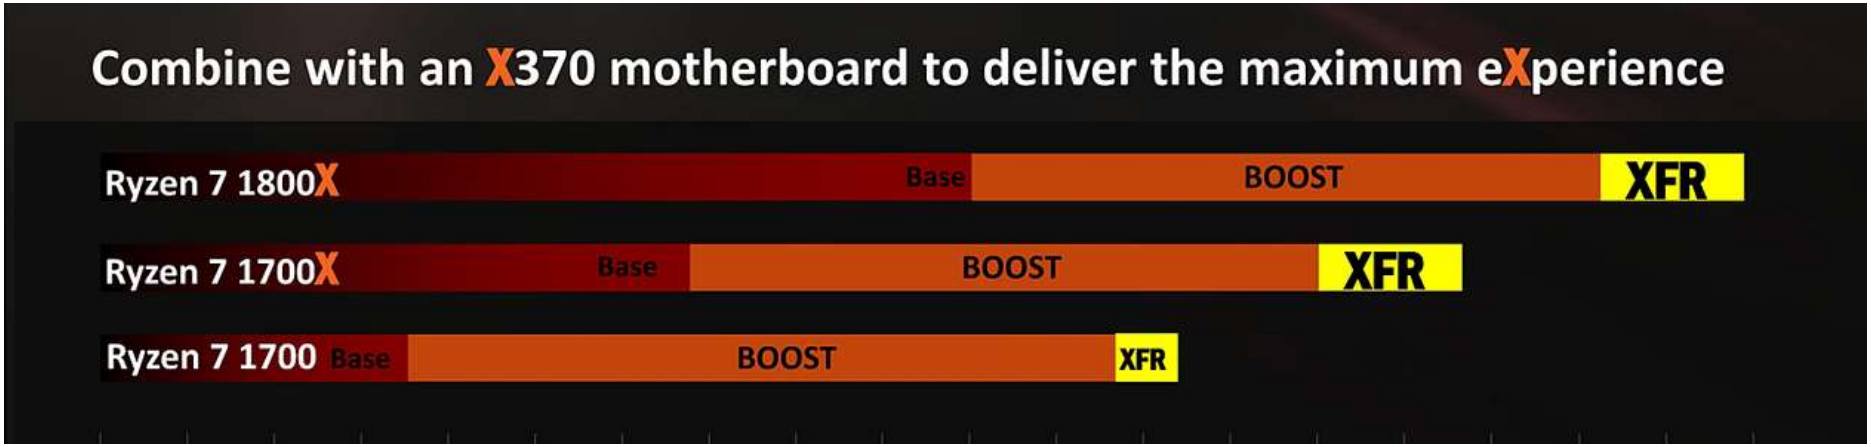

Example frequency scaling in the Ryzen 7 1800X [10]

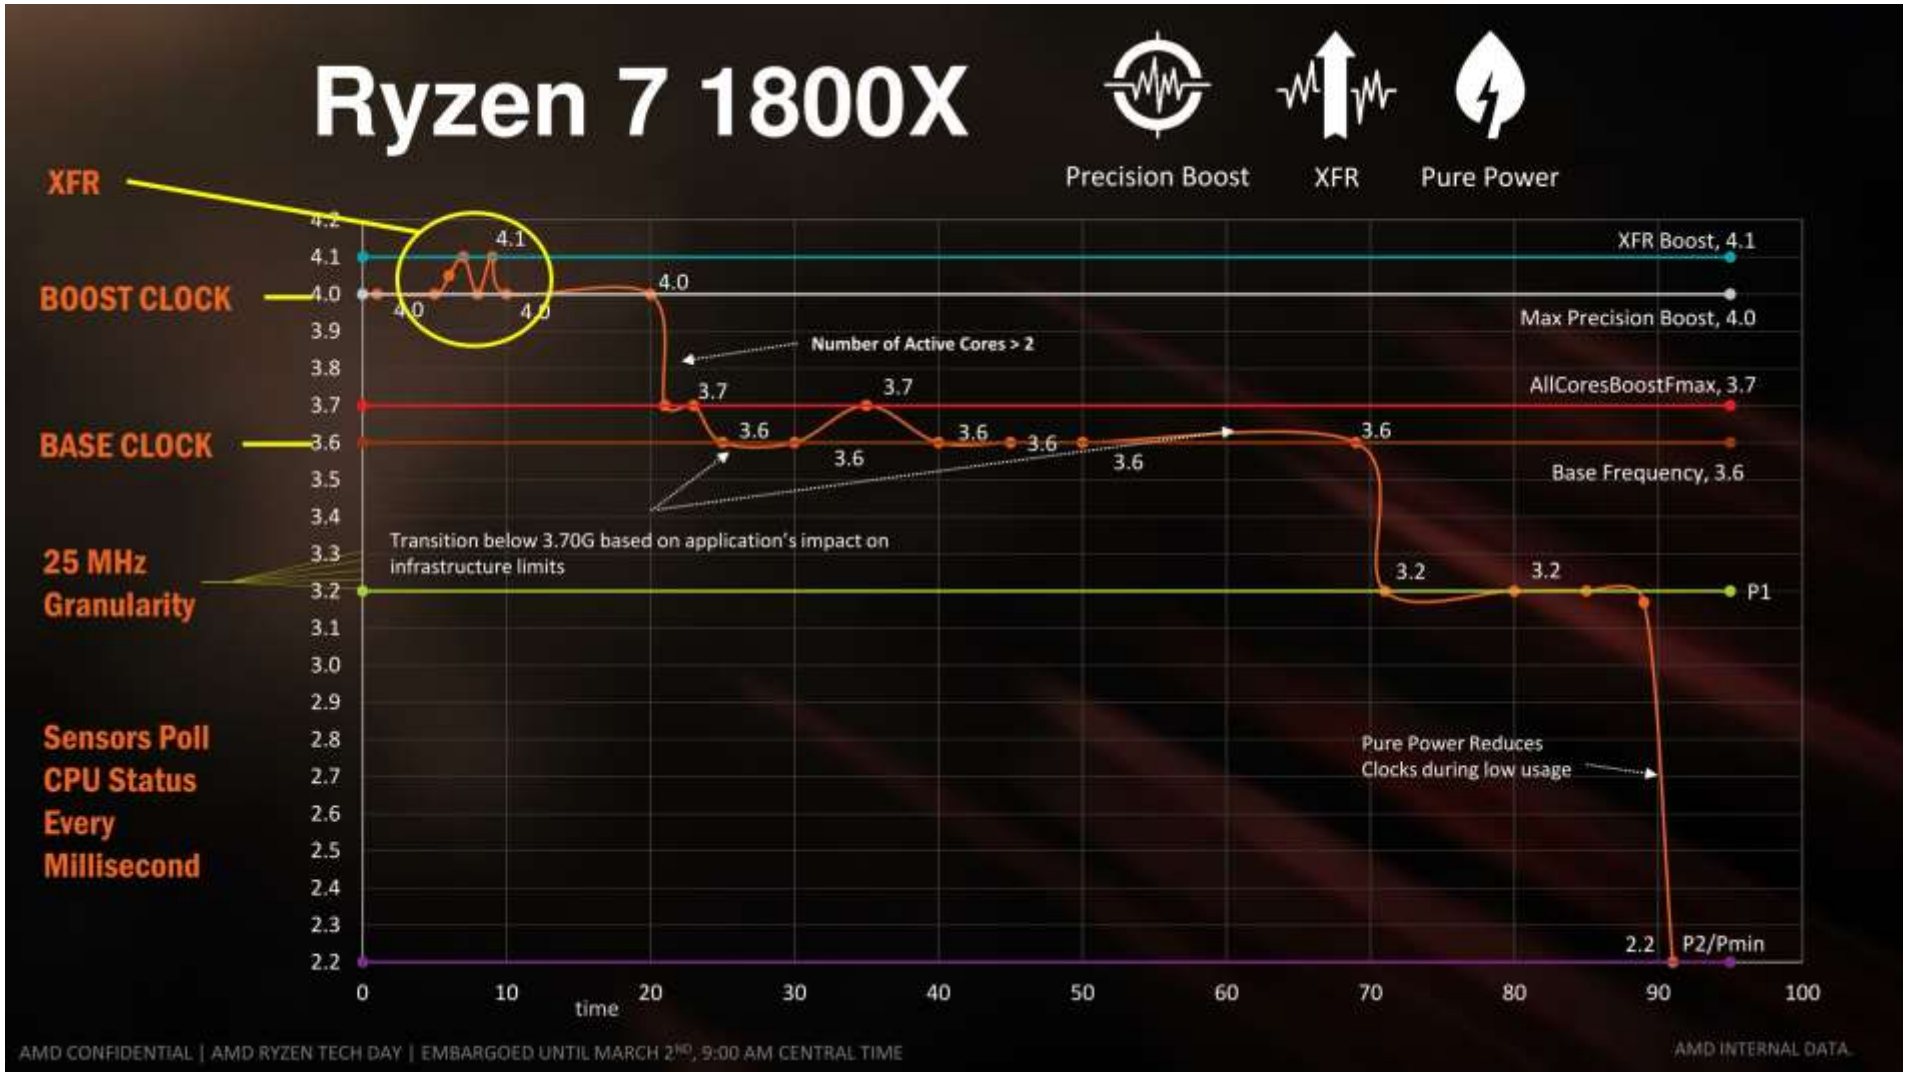

2.1.4 Further details relating the Zen core

It will be omitted

2.1.4 Further details relating the Zen core

Block diagram of the FX-execution of Zen [15]

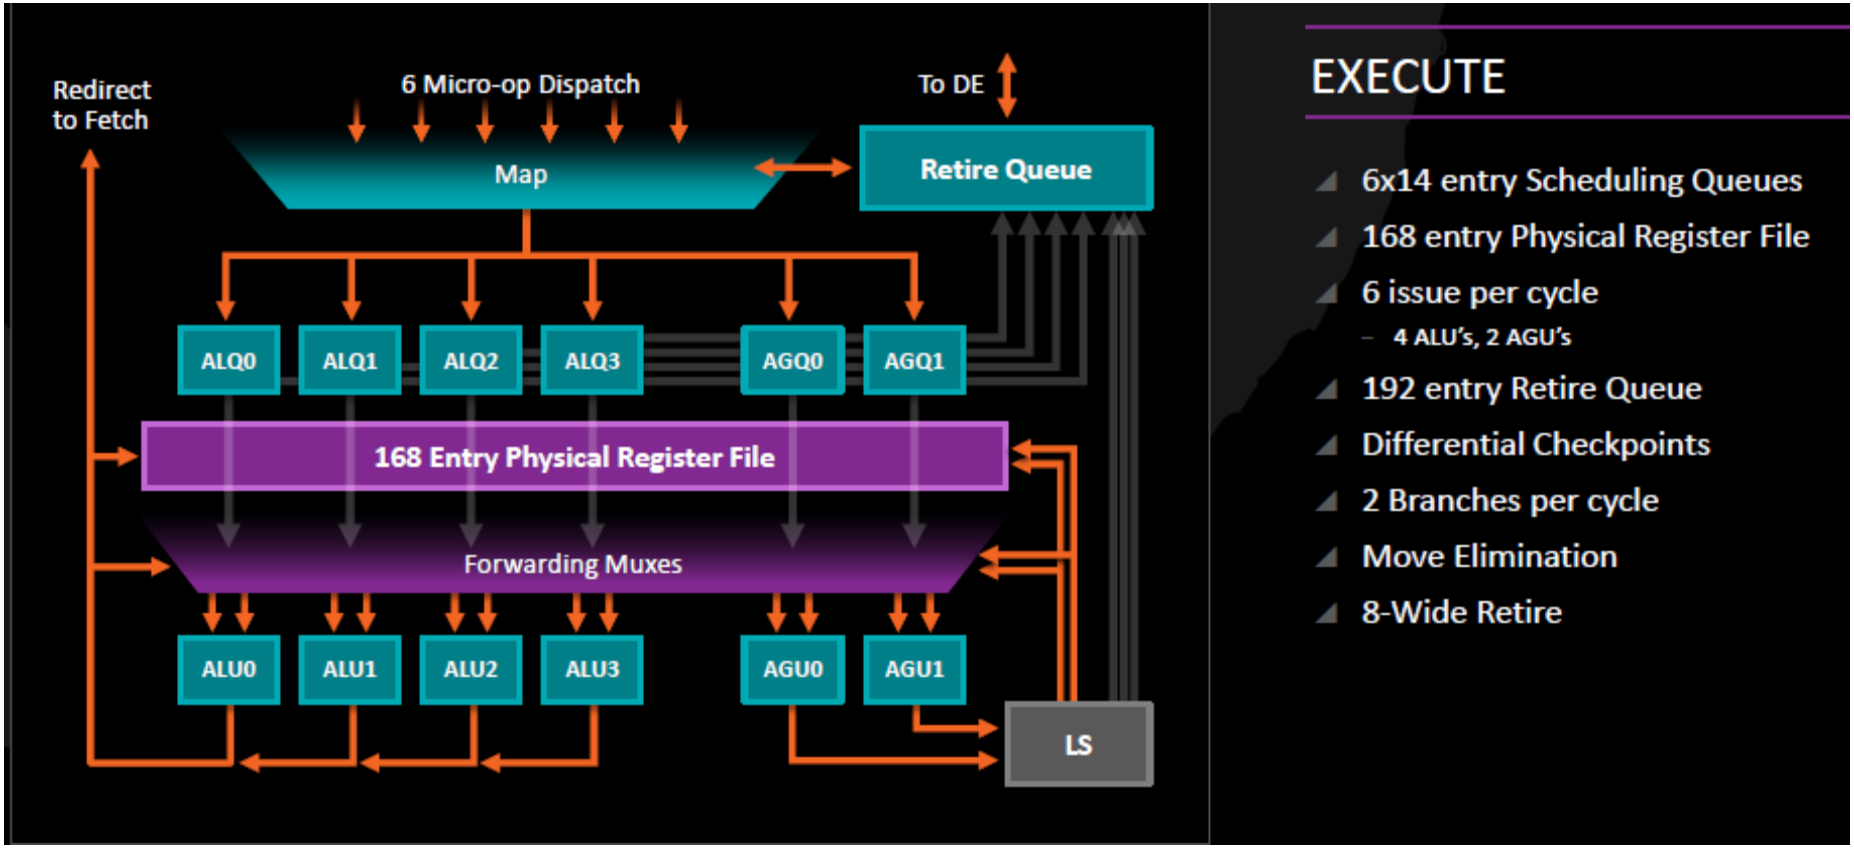

Block diagram of the FP-execution of Zen [15]

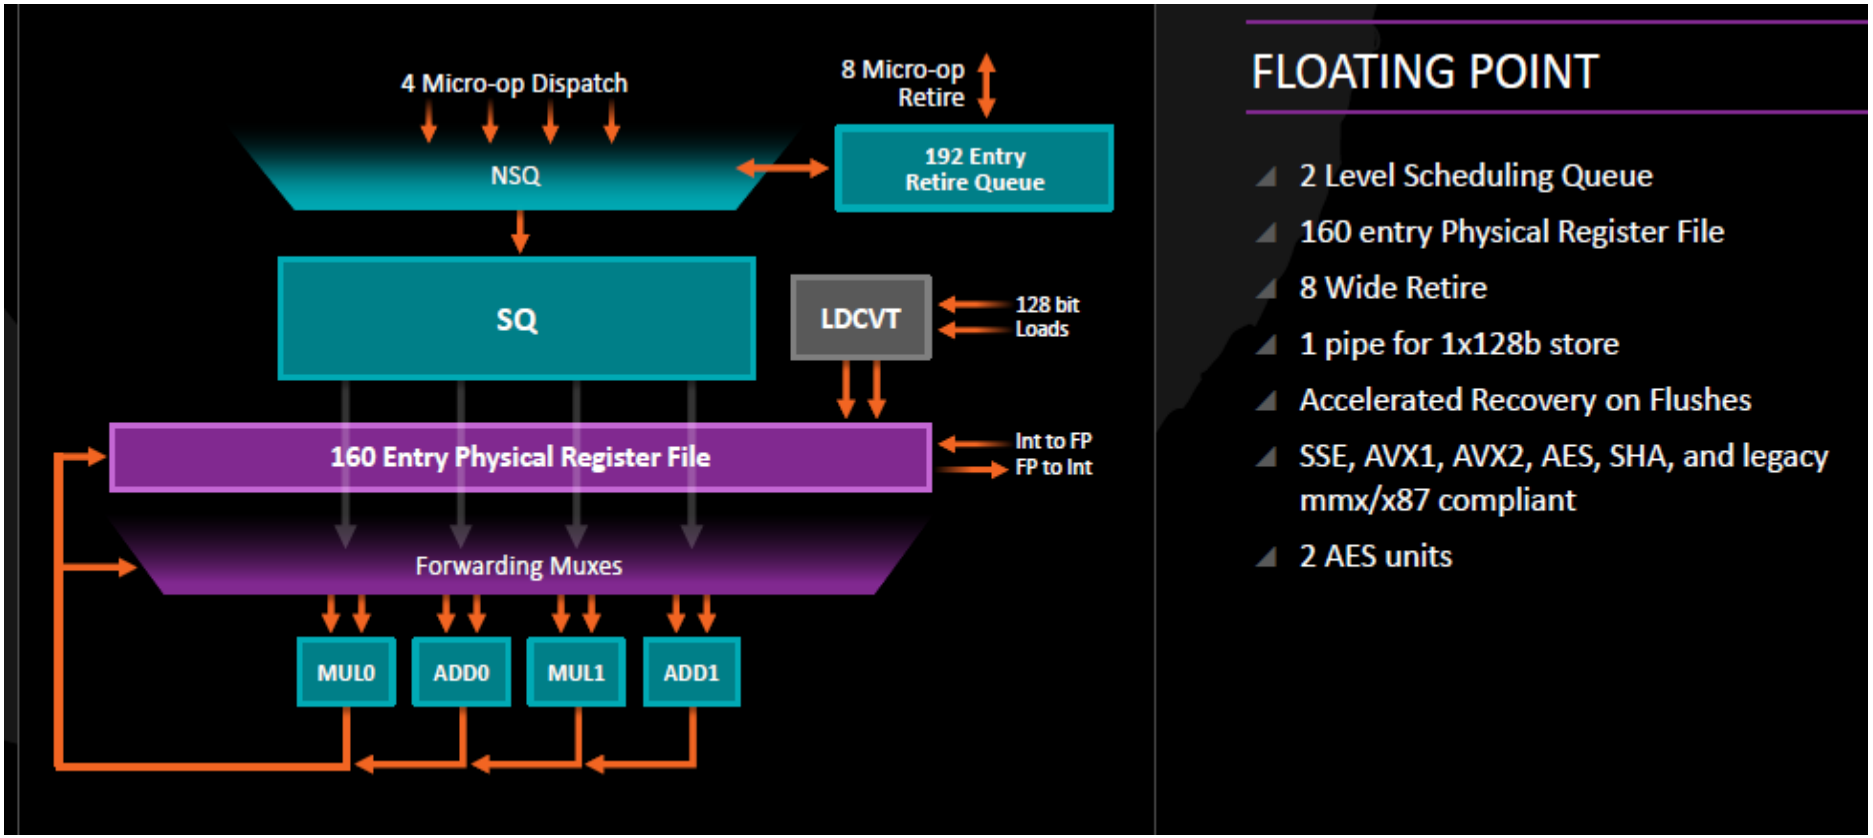

128-bit datapaths, 128-bit FP-units.

The cache architecture of Zen [15]

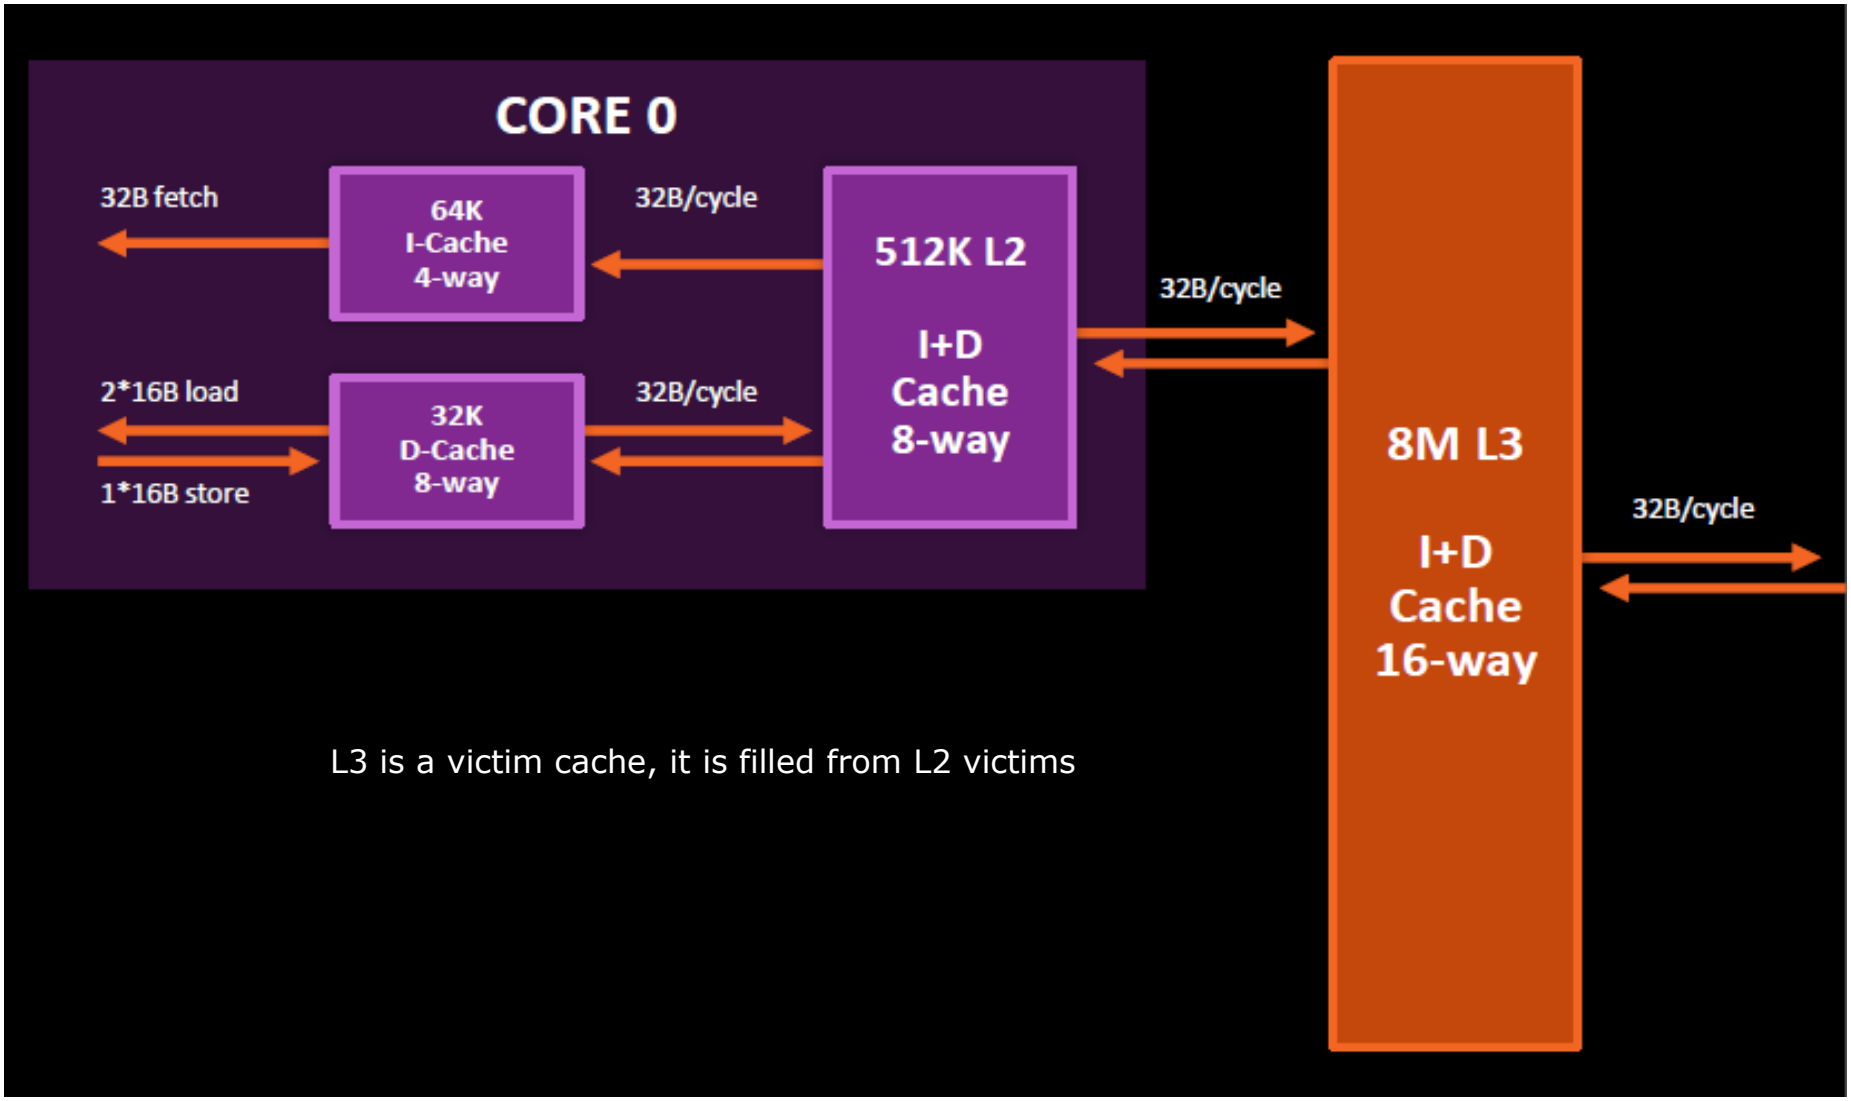

Implementation of SMT in Zen [10]

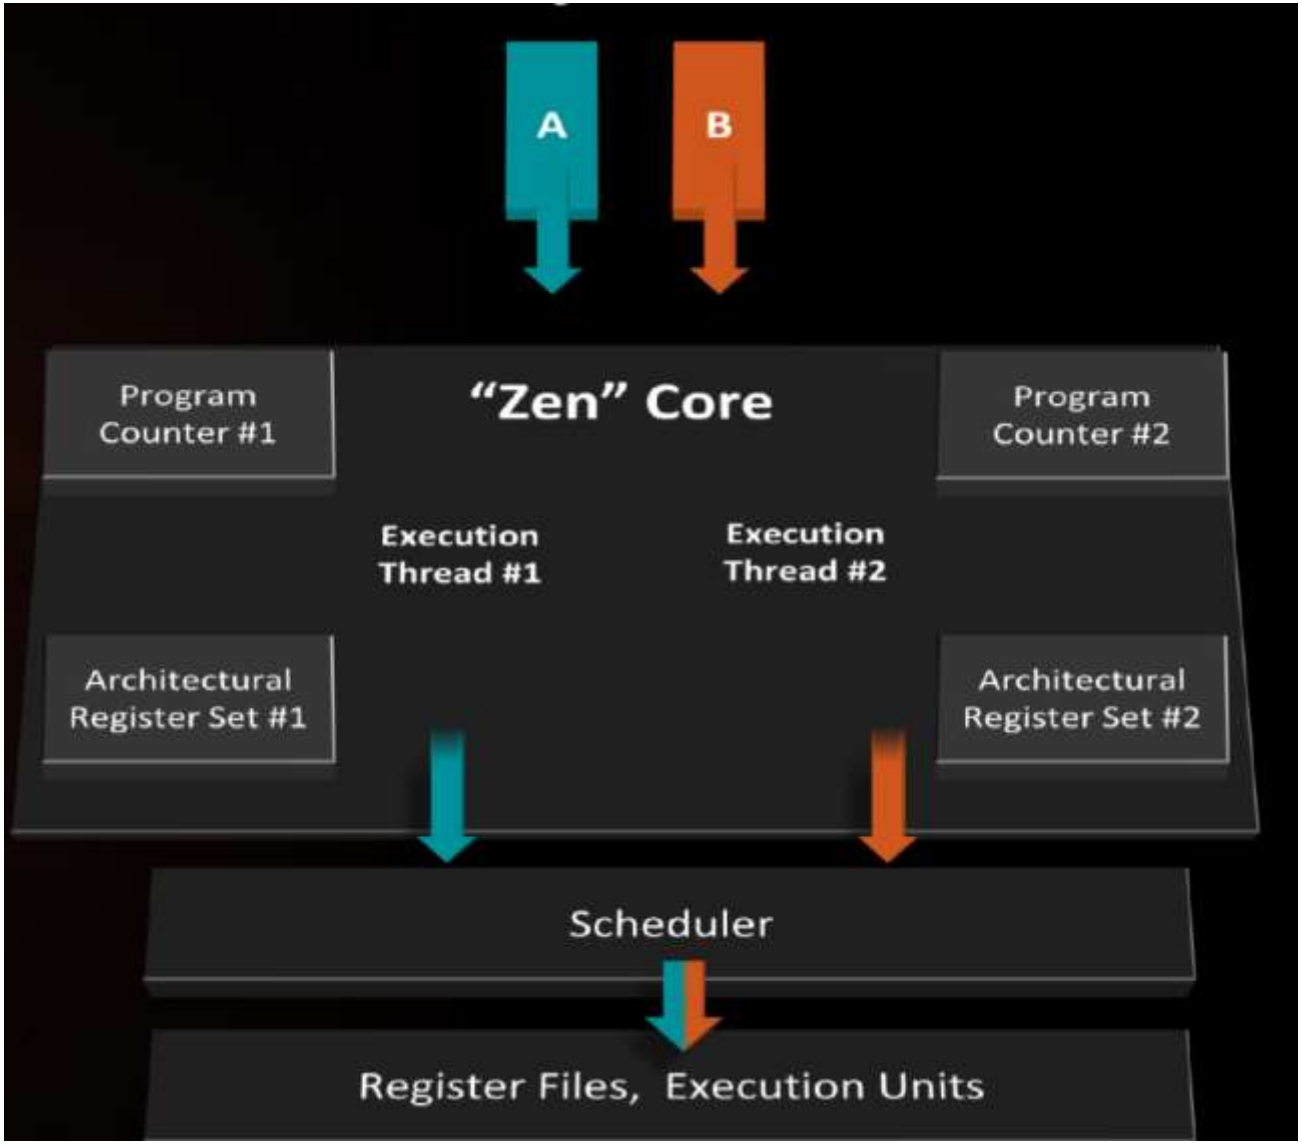

The Zen die [8]

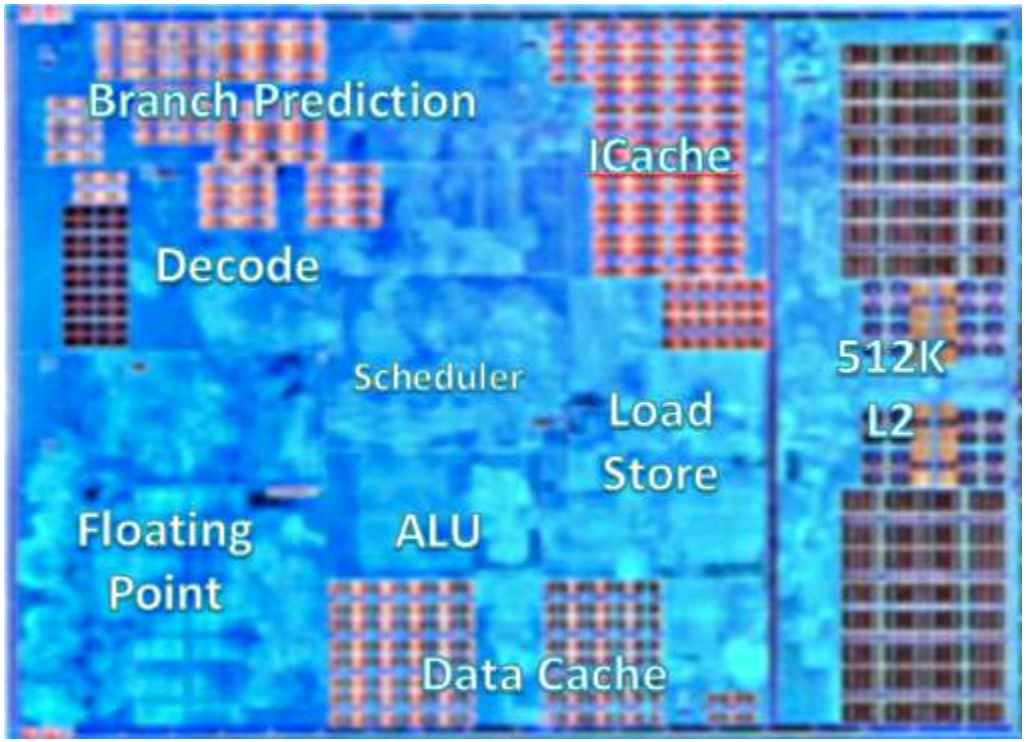
Branch Prediction
ICache
Decode
Scheduler
512K
L2
Load
Store
Floating
Point
ALU
Data Cache

Implementation of functional units on the Zen die [10]

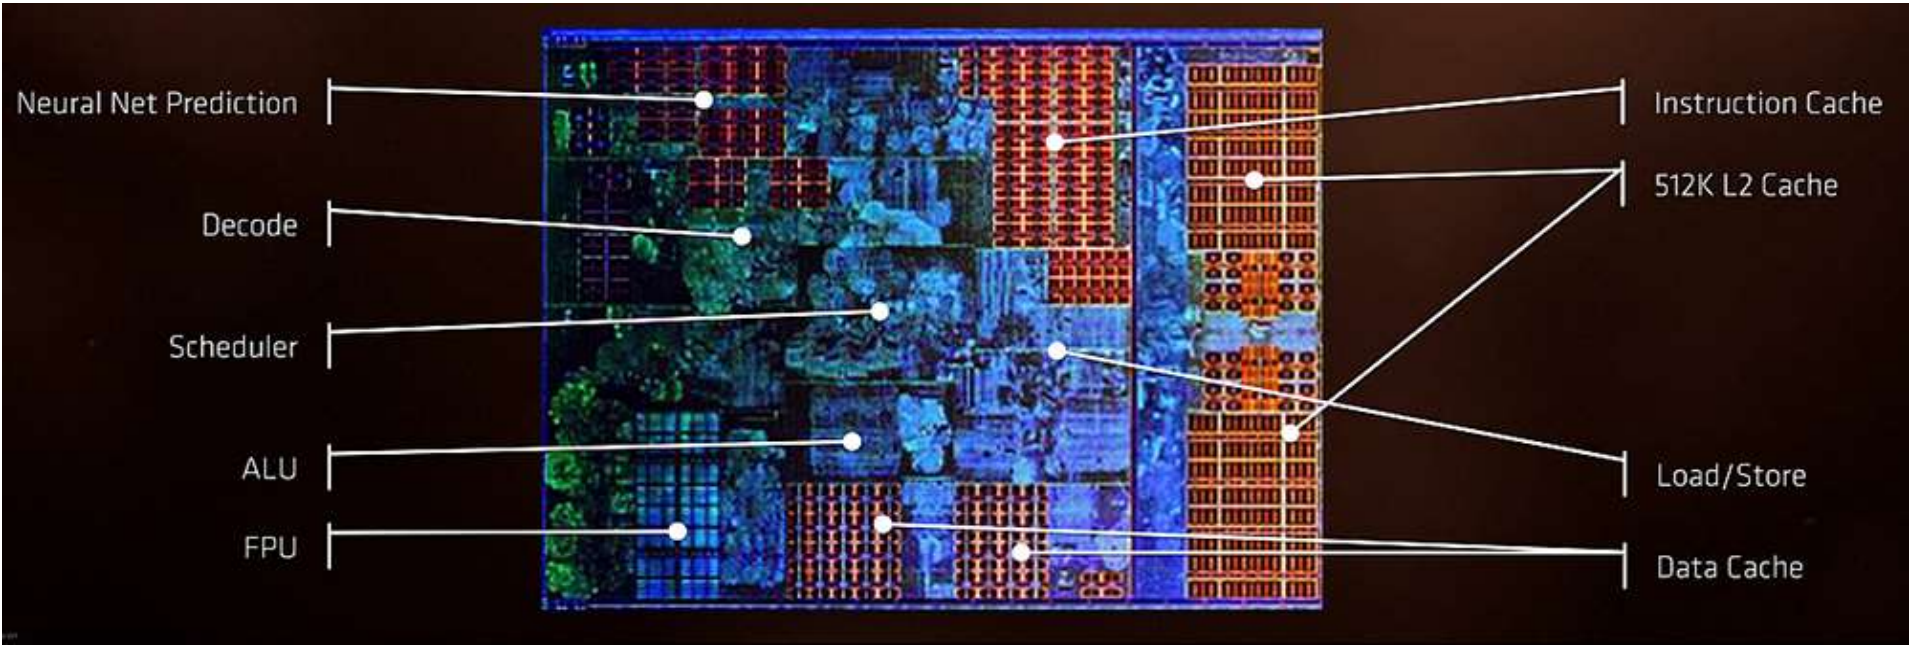

2.2 The Zen+ core

2.2 The Zen+ core

Improvements yielding from the 12 nm technology [74]

AMD's use of the 12LP (Low Power) technology (shown as an example while implementing a part of a datapath) [74]

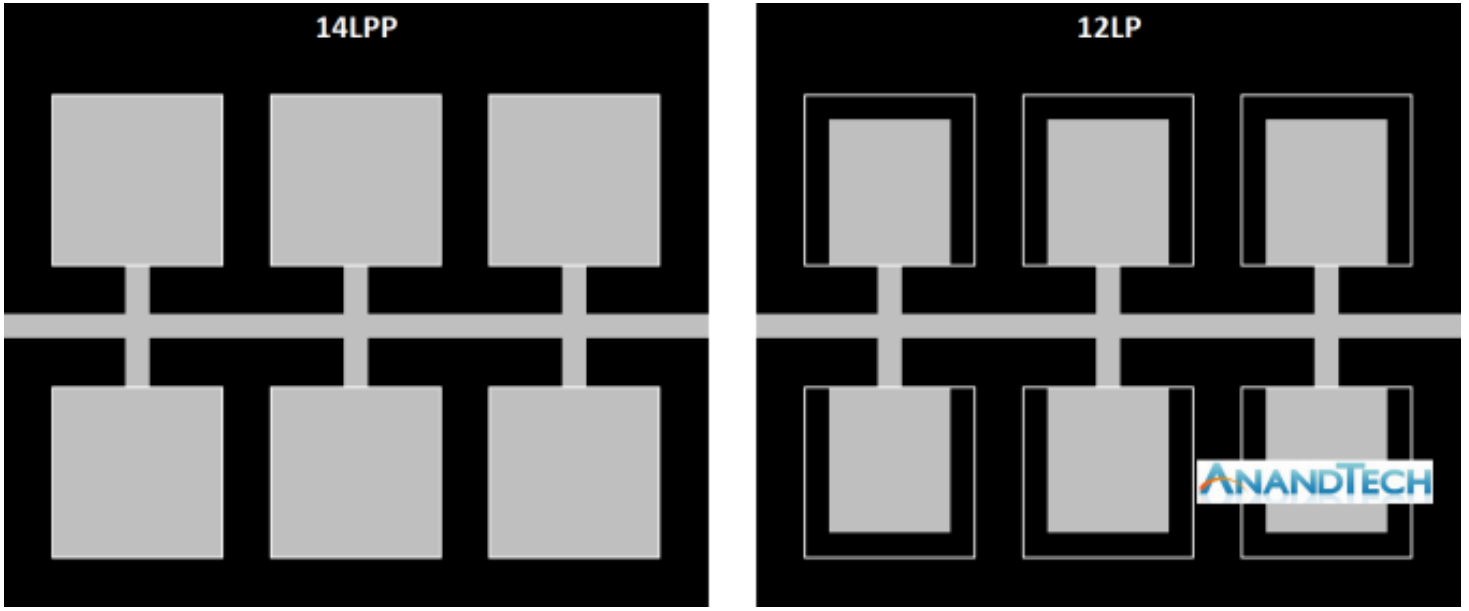

- While employing the 12LP technology AMD implemented the refreshed Ryzen DT line on the same die size (213 mm^2) and made use of the same number of transistors (4.8 billion) as in their original design.
- As a consequence, with reducing the feature size the dark silicon (unused silicon) area has grown.
- Since dark silicon acts as a thermal buffer between high-consumption parts of the die it improves the thermal behavior of the die. *

Benefits of the 12LP technology according to AMD [74]

- The new Ryzen 2000 series processors draw about 11 % less power than the Ryzen 1000 series models at the same clock frequency.
- It results about 15 % more performance at the same power.

*

Benefits of the 12LP technology vs 14 nm technology in cache latencies as given by AMD [74]

- 13% lower L1 latency (1.10ns vs. 0.95ns)
- 34% lower L2 latency (4.6ns vs. 3.0ns)
- 16% lower L3 latency (11.0ns vs. 9.2ns)
- 11% lower memory latency (74ns vs 66ns at DDR4-3200)
- Increased DRAM rate (DDR4-2666 vs. DDR4-2933)

Contrasting cache latencies in AMD's 1000 and 2000 series Ryzen DT processors [74]

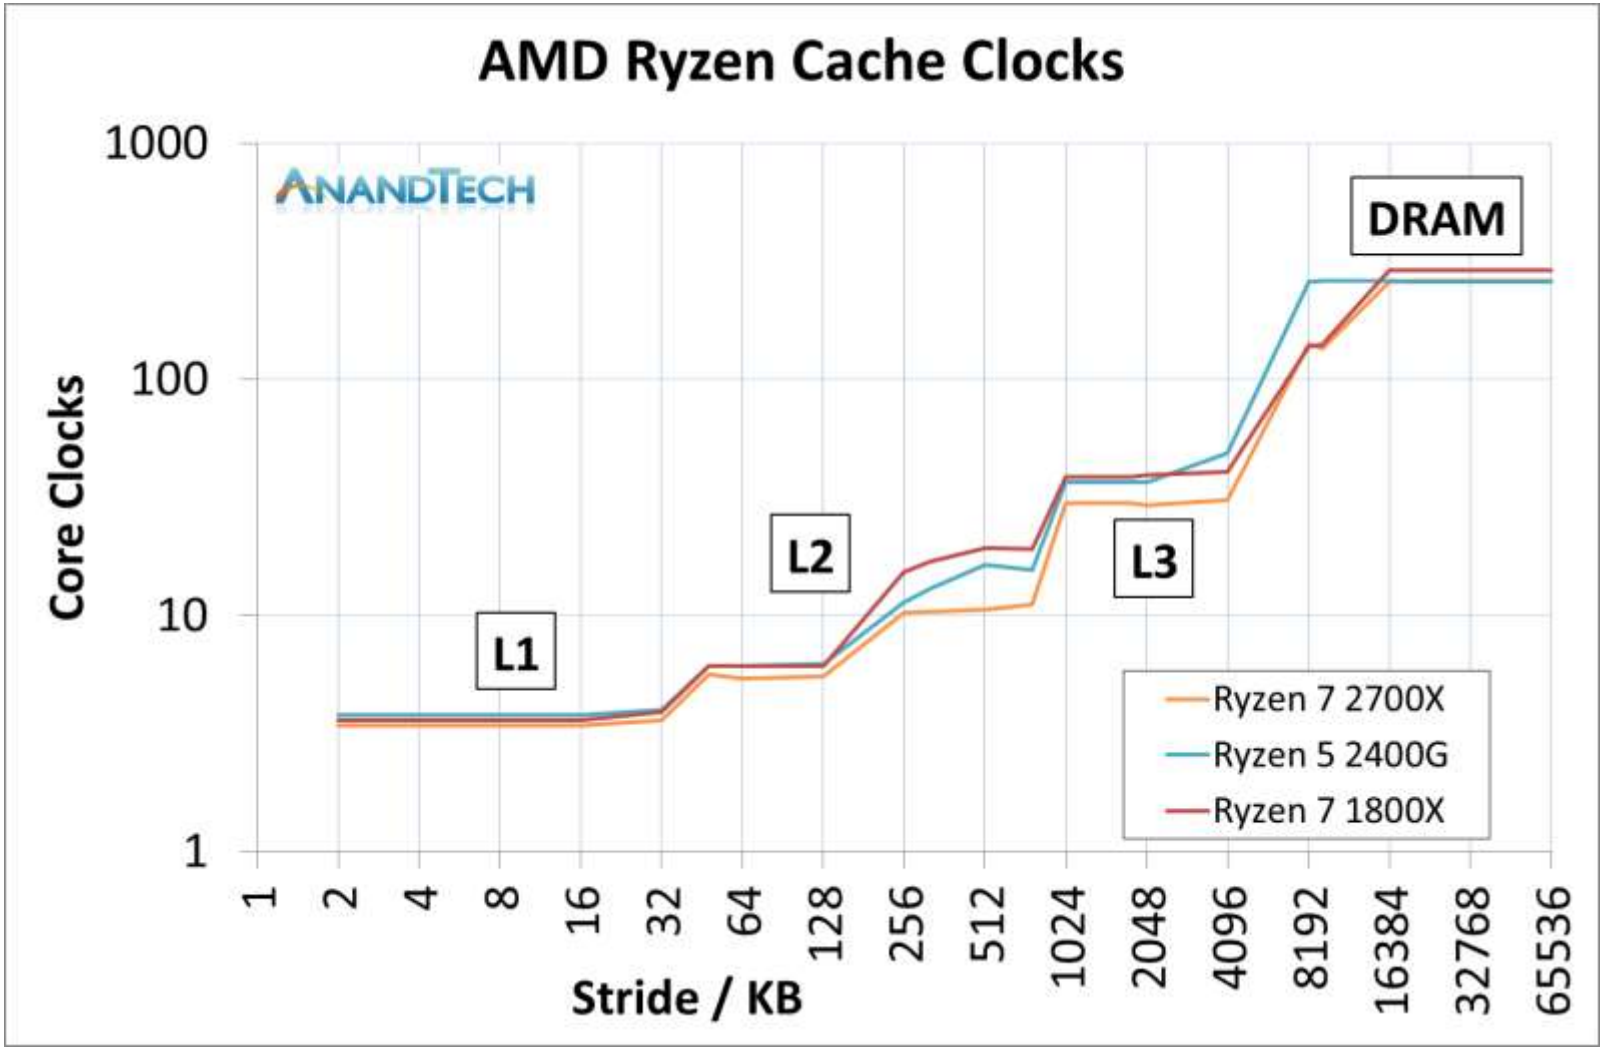

Contrasting cache latencies of AMD's and competing Intel's DT models [74]

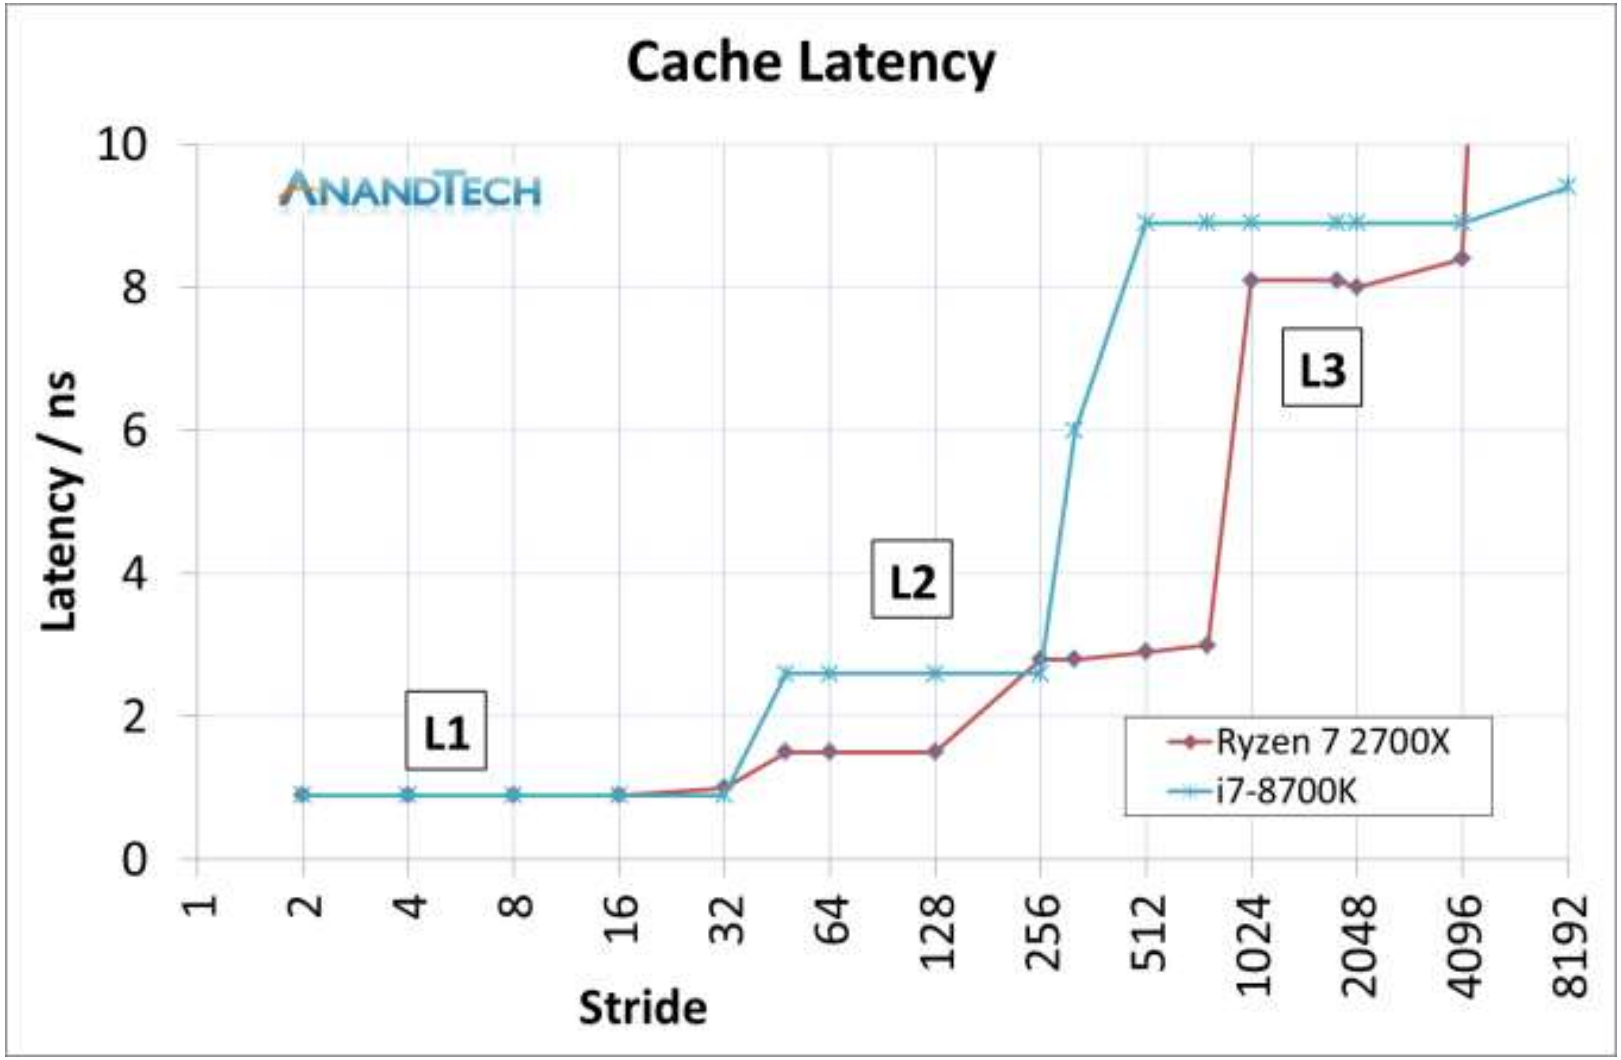

Contrasting IPC figures of subsequent DT models [74] -1

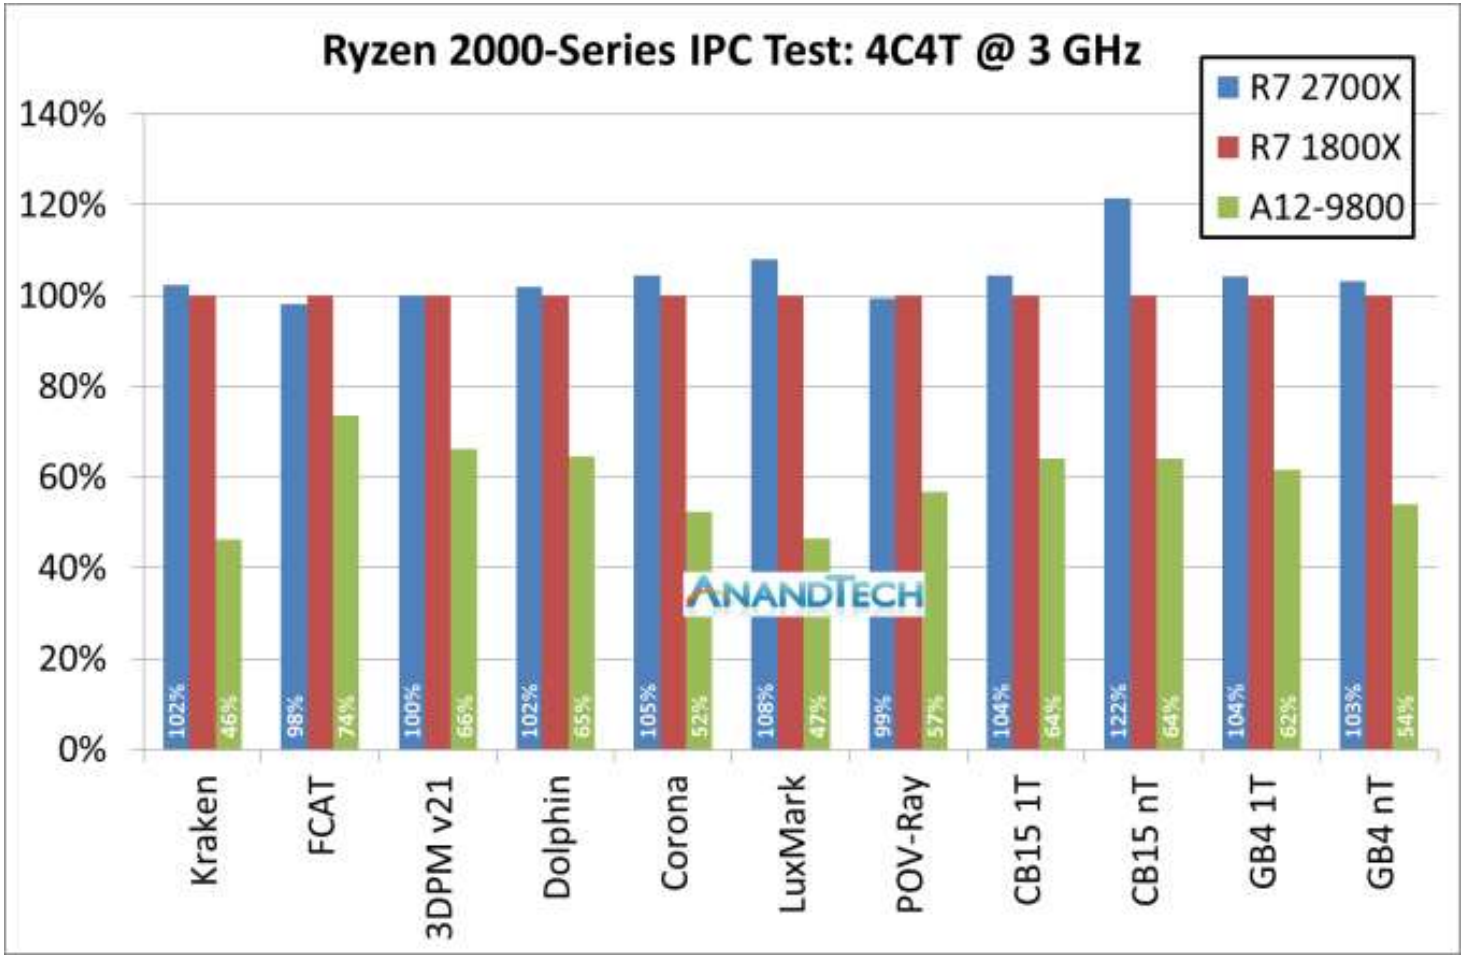

(12 nm)
(14 nm)

A12-9800:
Excavator-based
(28 nm)

As the Figure shows the R7 2700X provides on average only a few % (about 3 %) higher IPC than the preceding R7 1800X.

Contrasting IPC figures of subsequent DT models [74] -2

This about 3 % IPC boost is increased by

- about 6 % average clock frequency gain and
- by a few % gain due to Precision Boost 2 and XFR2

and results in an average performance gain of about 10 % over the previous DT generation.

Performance gain in subsequent Zen-based lines expected by AMD [74]

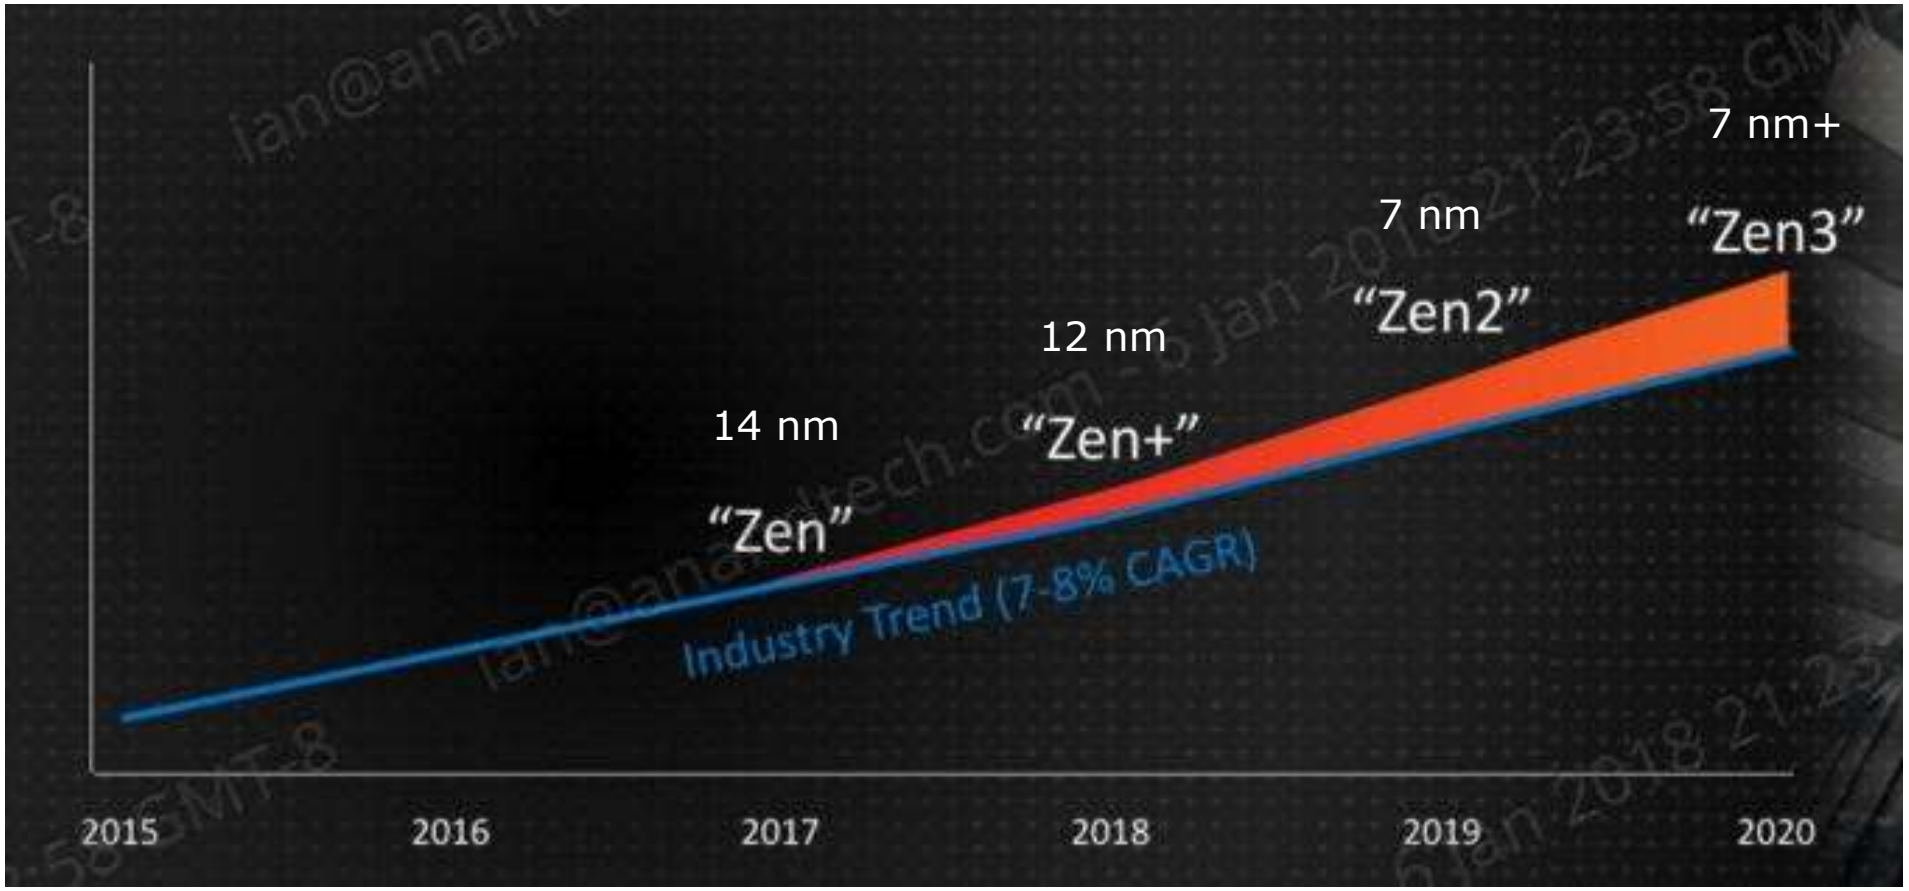

CAGR: Compounded Annual Rate of Growth

(Compounded: year/previous year)

2.3 The Zen 2 core

2.3 The Zen 2 core

- Introduced in 11/2018.
- Based on 7 nm technology from TSMC.
- First processors with the Zen 2 core inside are the Rome servers, AMD already started sampling EPYC 'Rome' processors to enterprise customers in 11/2018.
- Key features
  - Doubled FP and load/store performance through 256-bit execution units and datapaths,
  - half the energy per operation through 7 nm feature size vs. 12 nm.

Front-end enhancements [93]

Larger OP-cache

Load/store- and FP execution enhancements [93]

The Zen 2 core has 256-bit Load/Store and FP execution units and datapaths that is twice as wide than implemented in the Zen core.

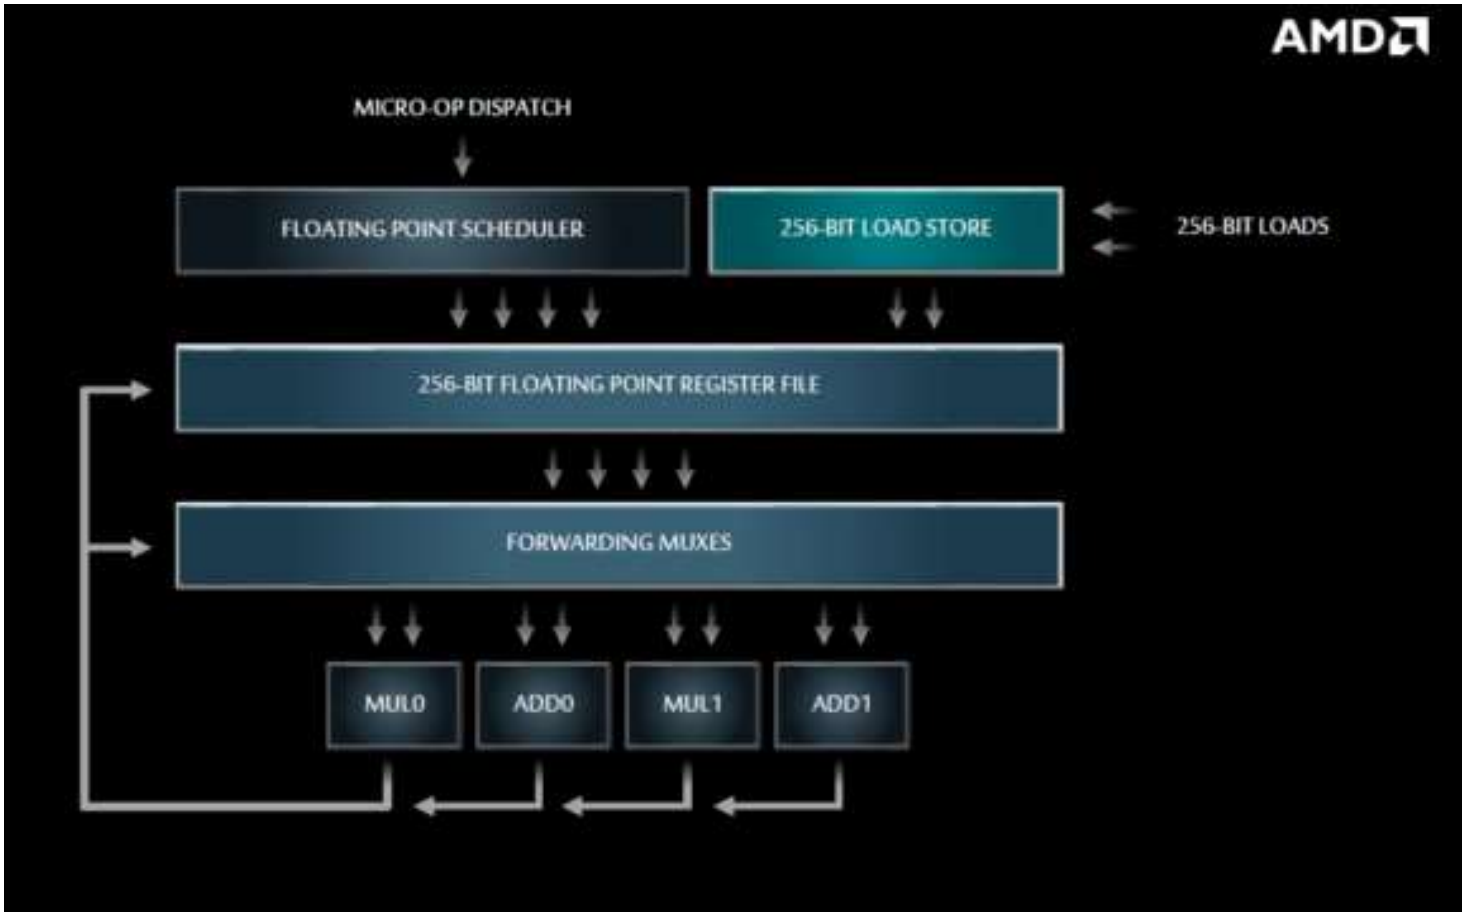

Figure: Load/Store and FP execution in the Zen 2 core [93]

For comparison: Block diagram of Load/Store- and FP-execution in Zen [15]

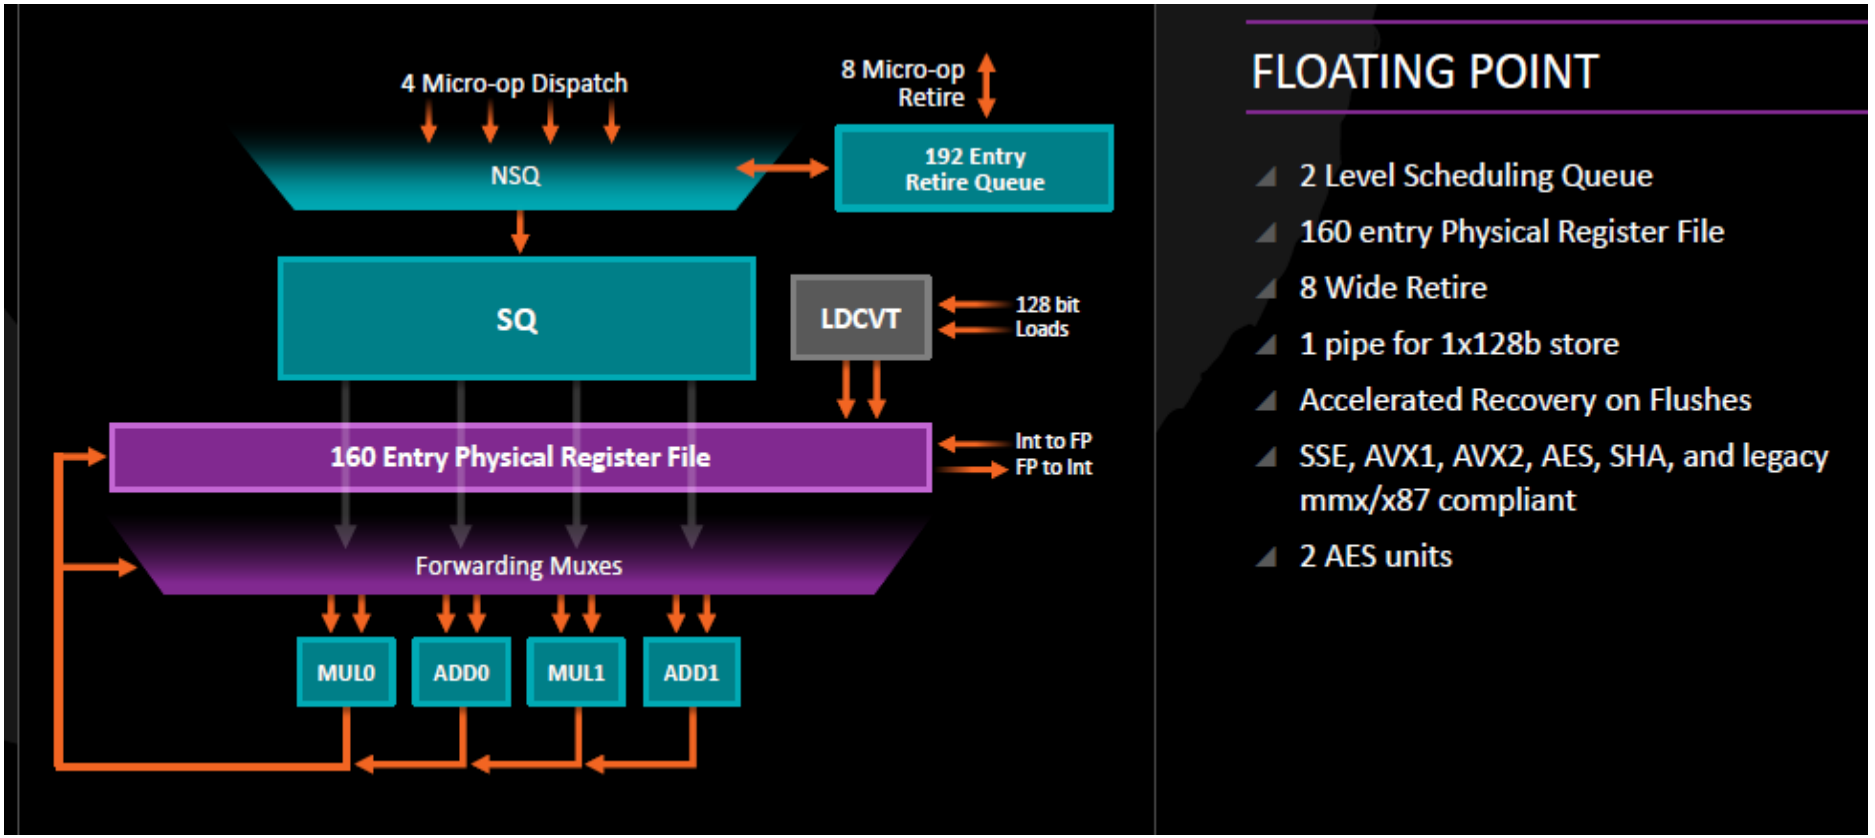

Note: There are 128-bit Load/Store and FP-execution unit as well as datapaths.

3. The 4-core CCX building block

3. The 4-core CCX building block -1

Basic building blocks of AMD's Zen-based processors

Zen core [8]

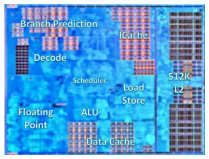

4-core CCX [9]
(Core CompleX)

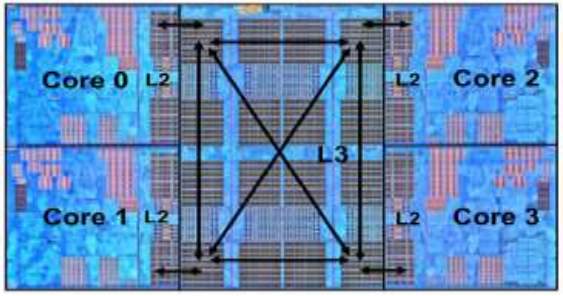

8-core Zeppelin module [10]
2x CCX

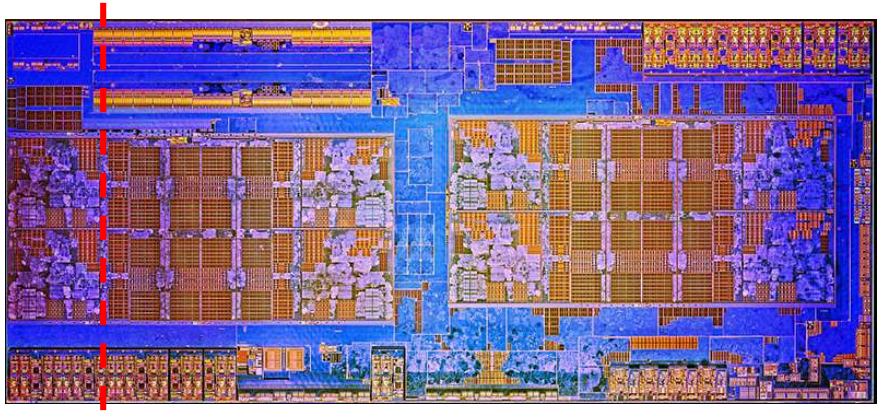

The 4-core CCX building block -2 [16]

- The 4-bit CCX (CPU-Complex) is the basic building block of Zen-based processor lines.
- Each core has a private 512 kB L2 cache whereas all 4 cores share an 8 MB L3 cache that is split into 4 slices (see Figure) by low-order address interleave.
- It includes 1.4 billion transistors implemented on a die area of 44 $mm^{2.}$

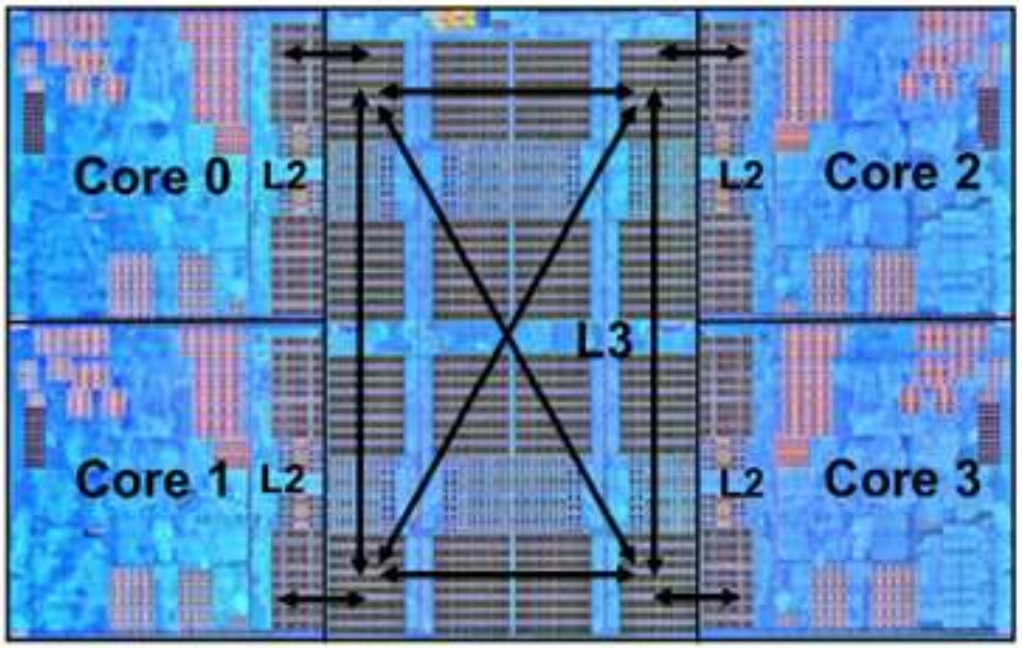

Figure: The CCX building block [16]

The 4-core CCX building block -3 [15]

- The L3 cache is mostly exclusive to the L2 cache, it is actually a victim cache for the L2 cache.
- Due to the internal crossbar interconnection (see previous Figure), each core can access each L3 cache segment with the same average latency.

*

3. The 4-core CCX building block (4)

Block diagram and die photograph of the CCX Core CompleX [15], [59]

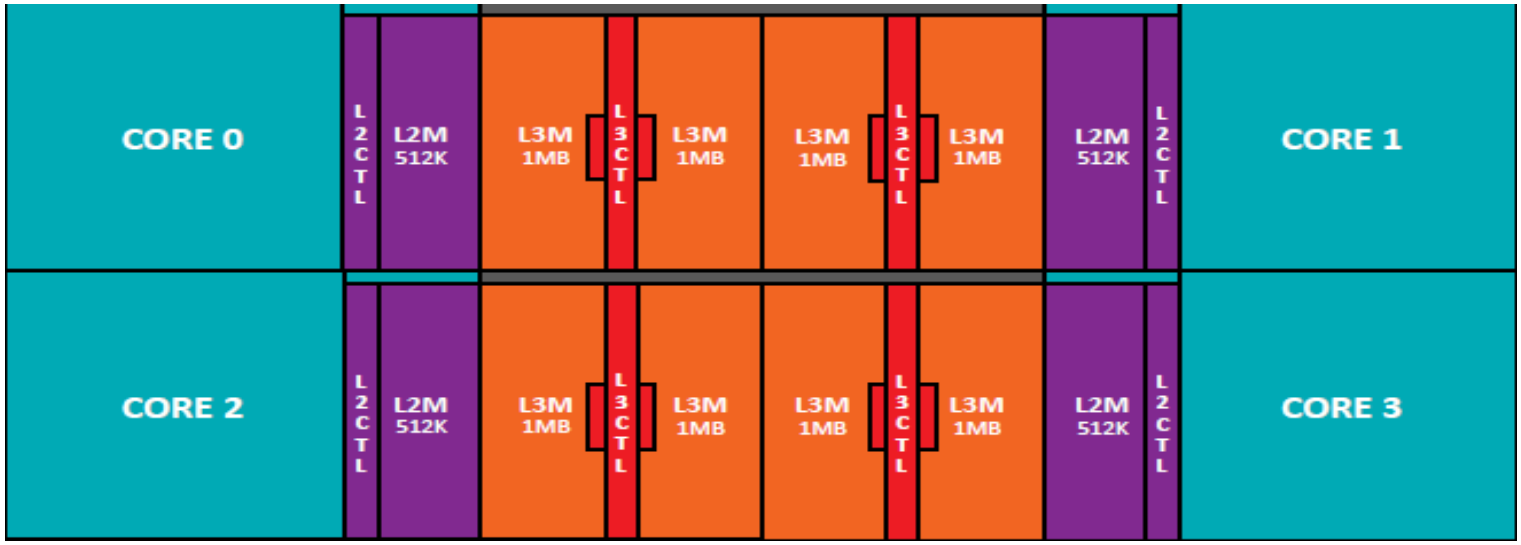

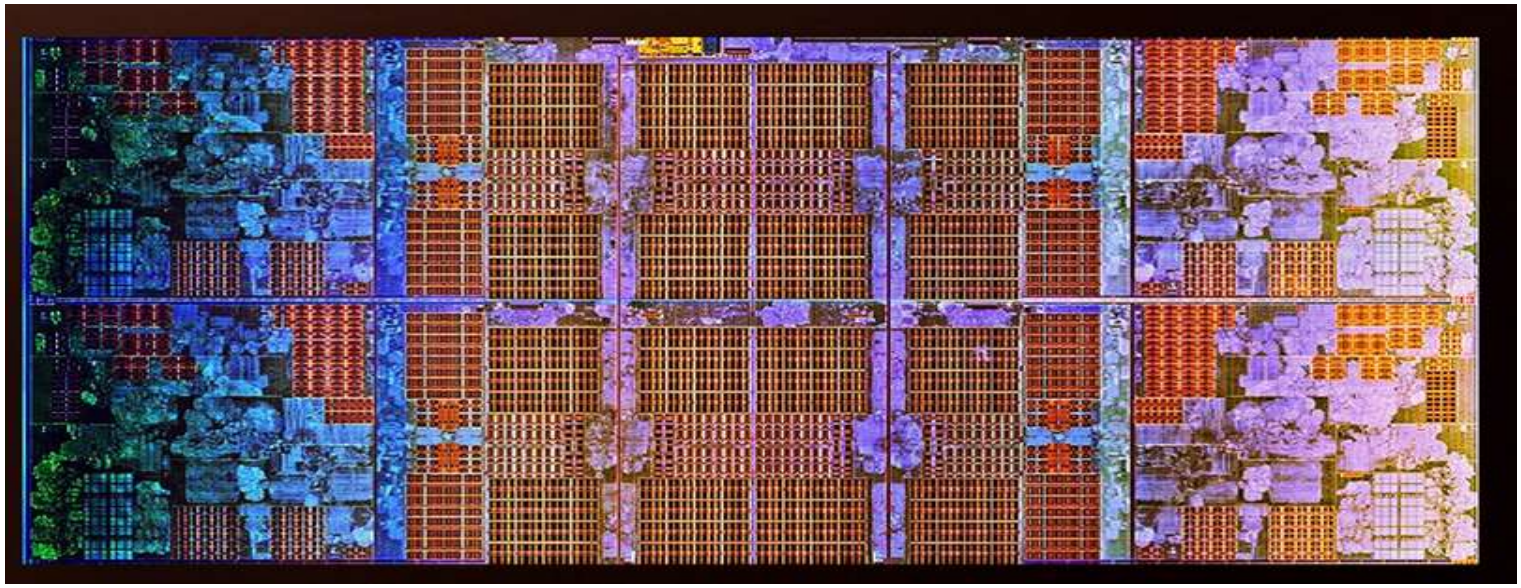

Comparing AMD's CCX (a) and ARM's MPCore (actually Cortex-A 73) (b)

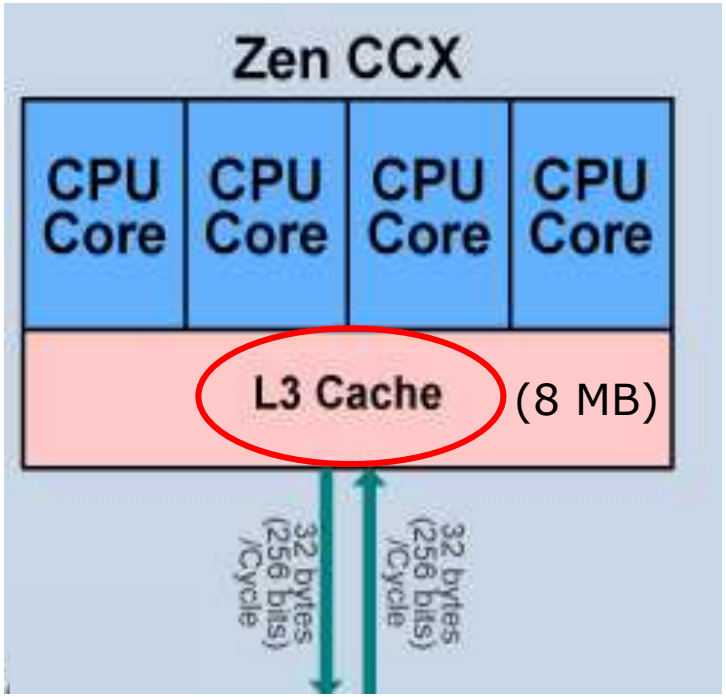

a) (2017)

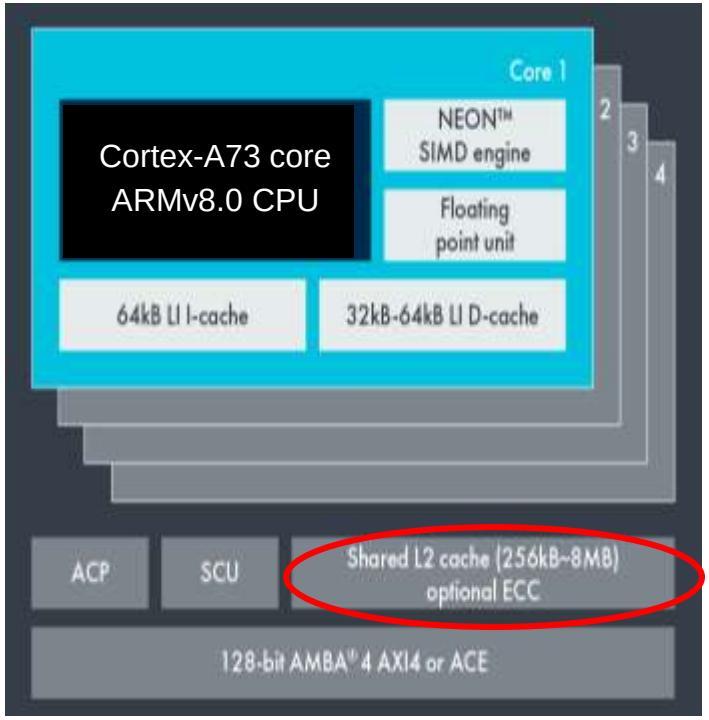

b) 2016) [55]

ACP: **A**ccelerator **C**oherence **P**ort
(to connect non-cached coherent data sources)

SCU: **S**noop **C**ontrol **U**nit
(Provides coherence within the the processor)

ECC: **E**rror **C**orrecting **C**ode

AMBA: **A**dvanced **M**icrocontroller **B**us **A**rchitecture
(On-chip bus standard for SoC) designs)

ACE: **A**XI **C**oherency **E**xtensions
AMBA4 cache coherent interface
(Used in big.LITTLE systems for smartphones, tablets, etc.)

AXI4: Advanced eXtensible Interface
(Not cache coherent interface) *

Comparing ARM's MPCore (Cortex-A73) (a) and DynamIQ core clusters (b)

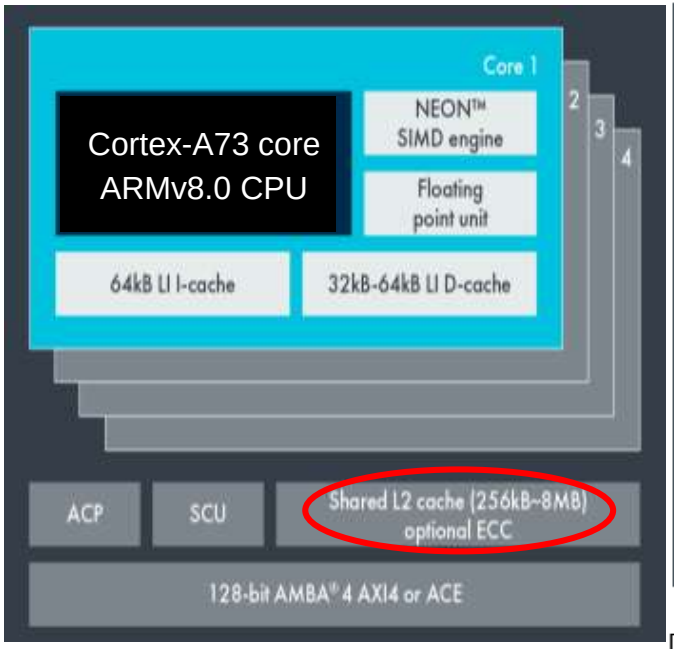

(a) (2016) [55]

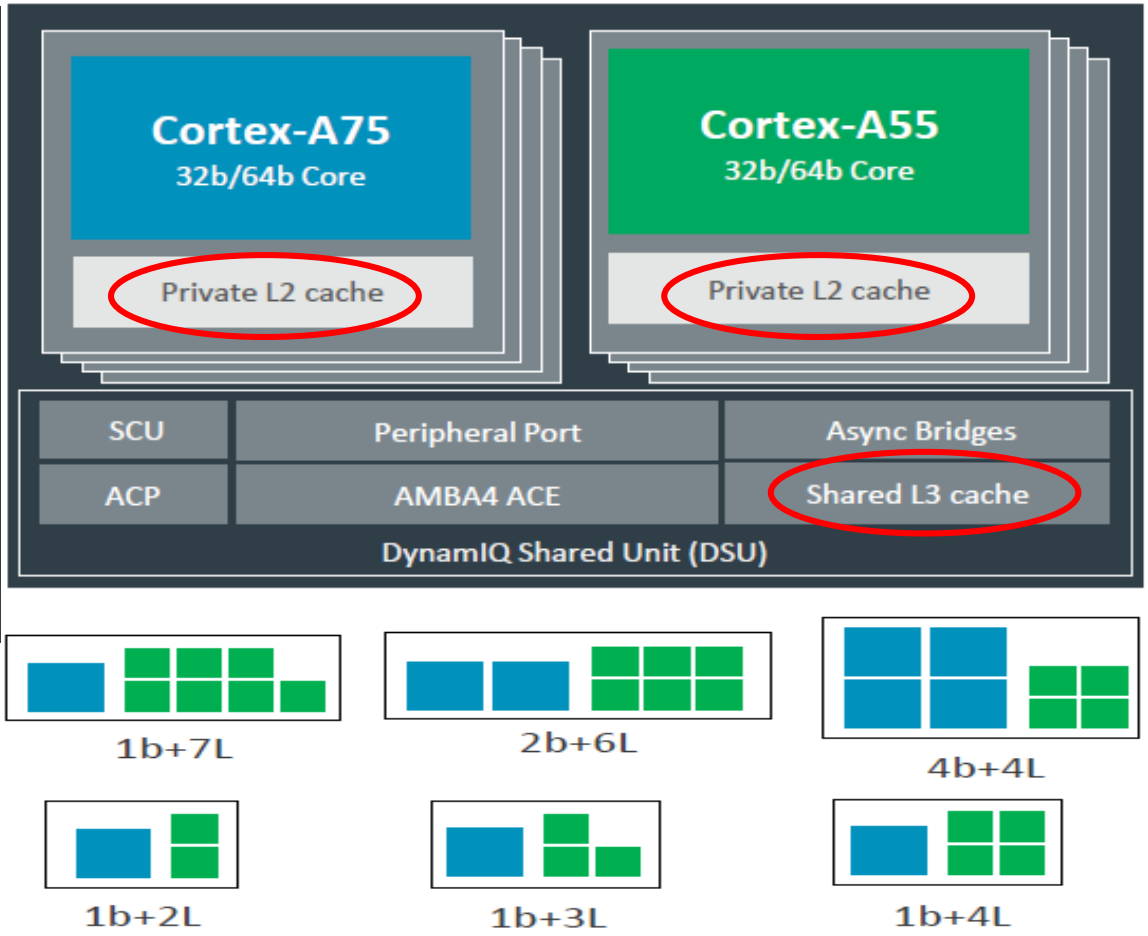

(b) (2017) [56]

- ARMv8.1 support, e.g. Cortex-A53/ A57/A72/A73, up to 4 cores
- Shared L2 cache

- ARMv8.2 support, (Cortex-A55/75), up to 8 cores
- Shared L3 cache, partial L3 cache power down,
- Private L2 caches

*

4. The Zeppelin module

4. The Zeppelin module

Basic building blocks of AMD's Zen-based processors

Zen core [8]

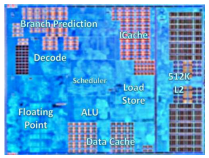

4-core CCX [9]
(Core CompleX)

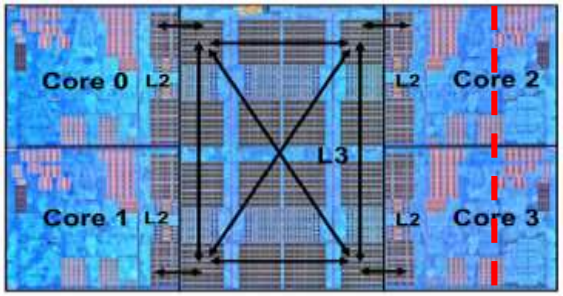

8-core Zeppelin module [10]
2x CCX

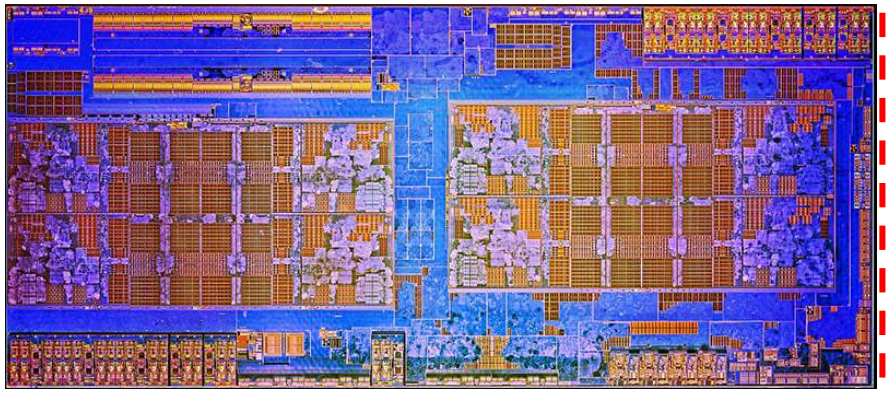

The Zeppelin module

It is based on two CCX blocks connected together by the Infinity Fabric, as seen below.

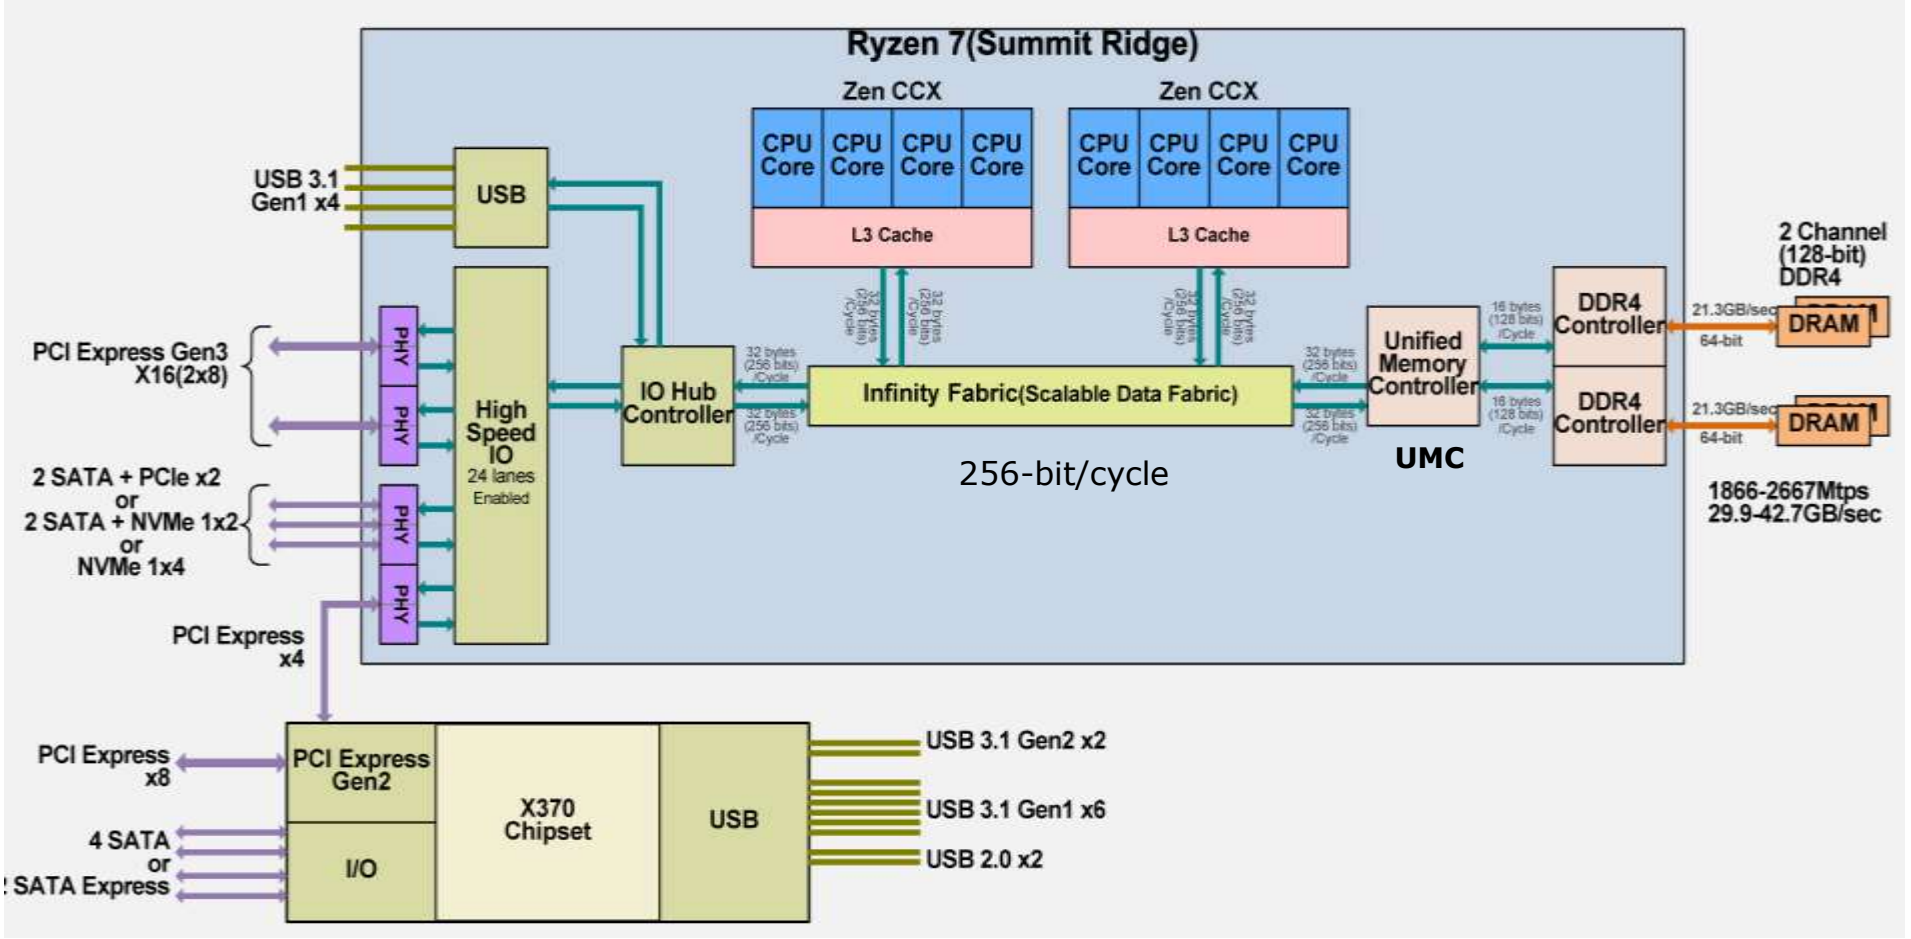

Figure: Block diagram of the Zeppelin module based Ryzen processor [17]

*

. The Zeppelin module (2b)

Embedding the Zeppelin module into a Threadripper or Epyc processor [88]

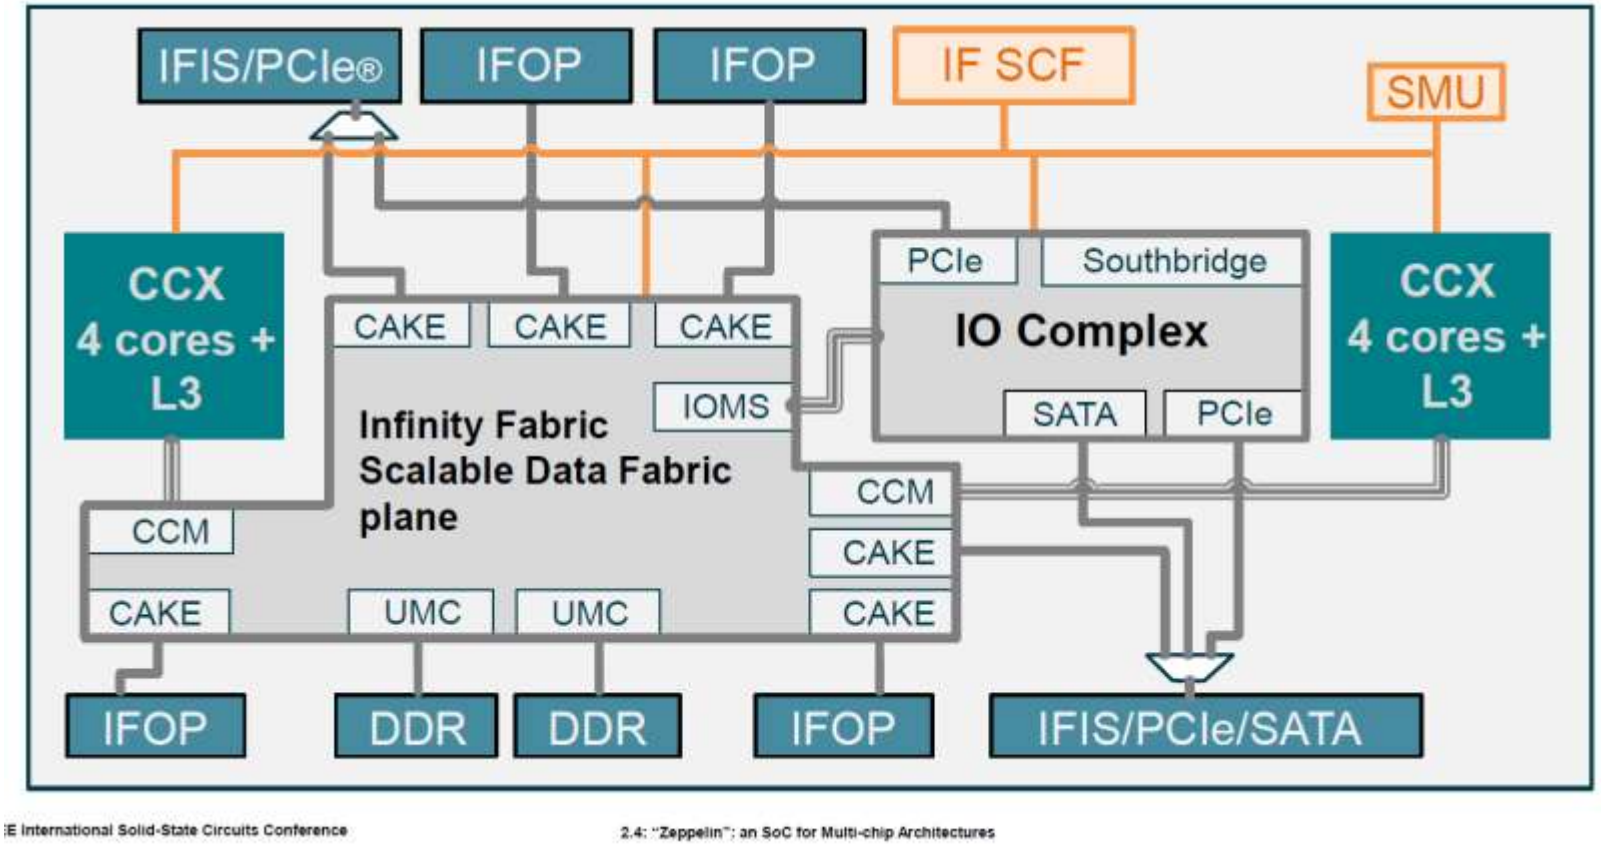

CAKE. The Coherent AMD socKet Extender
IFSCF: IS Scalable Control Fabric
SMU: System Management Unit (microcontroller)

UMC: Unified Memory Controller
IFOP: IF On-Package interconnect links
IFIS: IF Inter-Socket interconnect links

The server first nature of the Zeppelin module [18]

AMD designed the Zeppelin module for the server market, but they utilize it also as a modular building block for their consumer (Ryzen) and HED (Threadripper) lines, as indicated in one of the preceding Figures.

4. The Zeppelin module (4)

Die photograph of the 14 nm Zeppelin die (Ryzen die) [78]

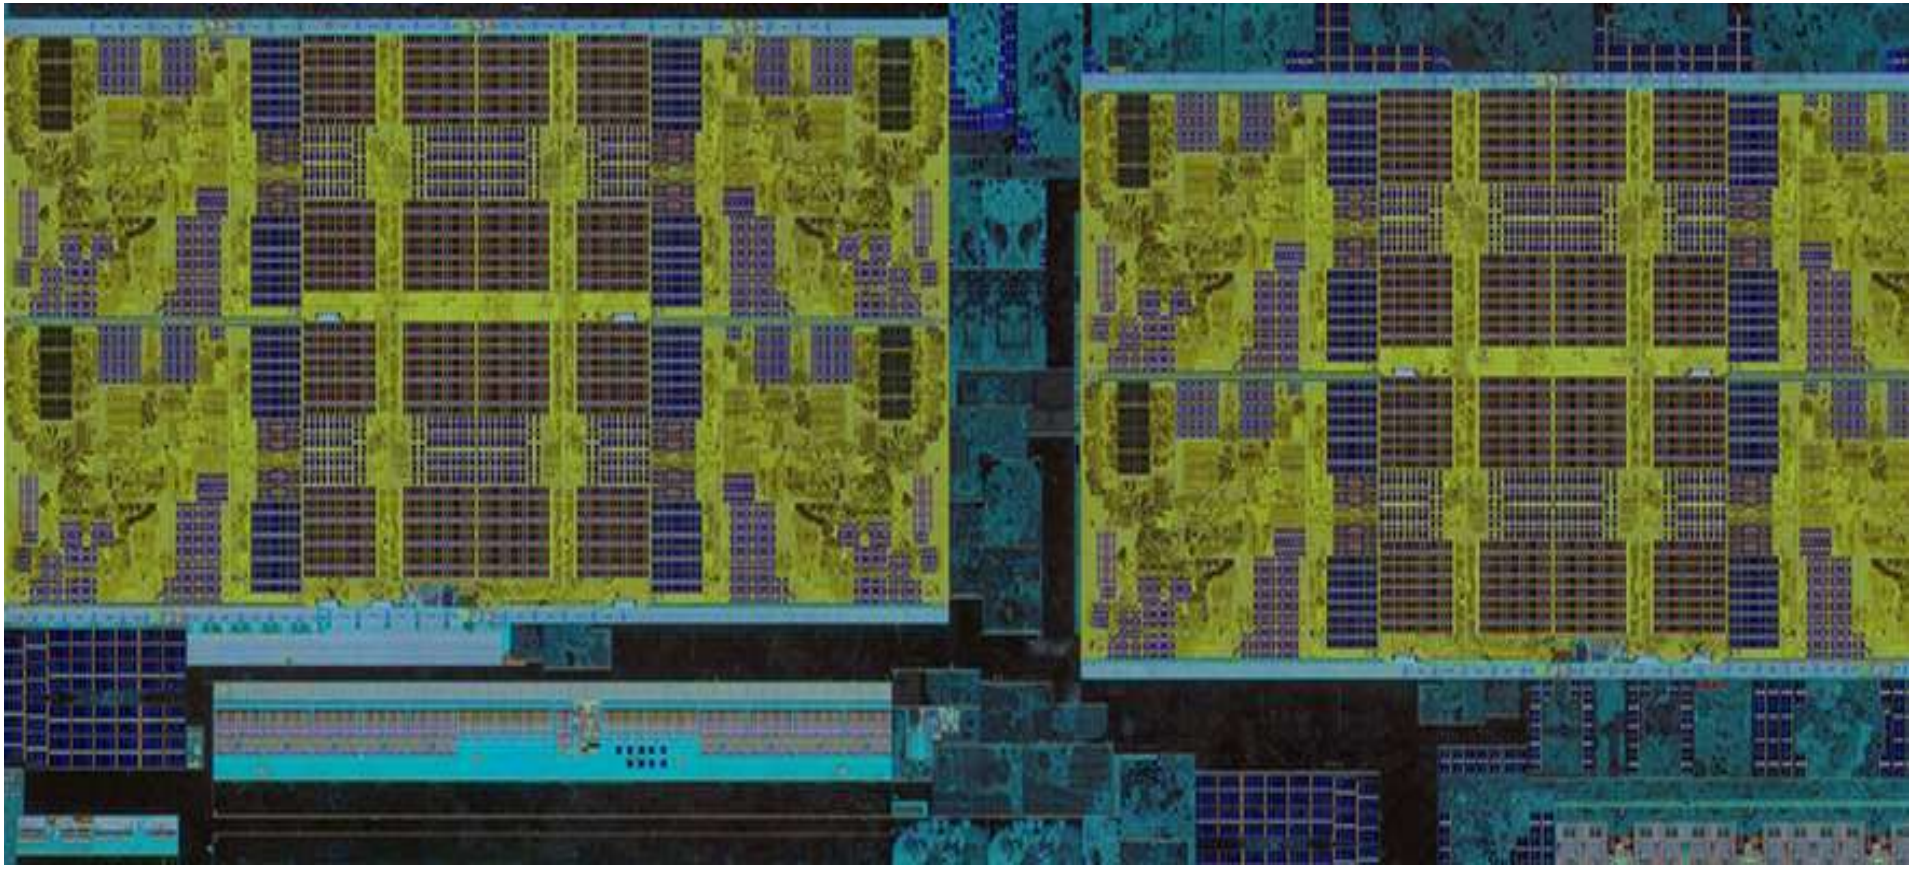

Floorplan of the 14 nm Zeppelin die (Ryzen die) [77] -1

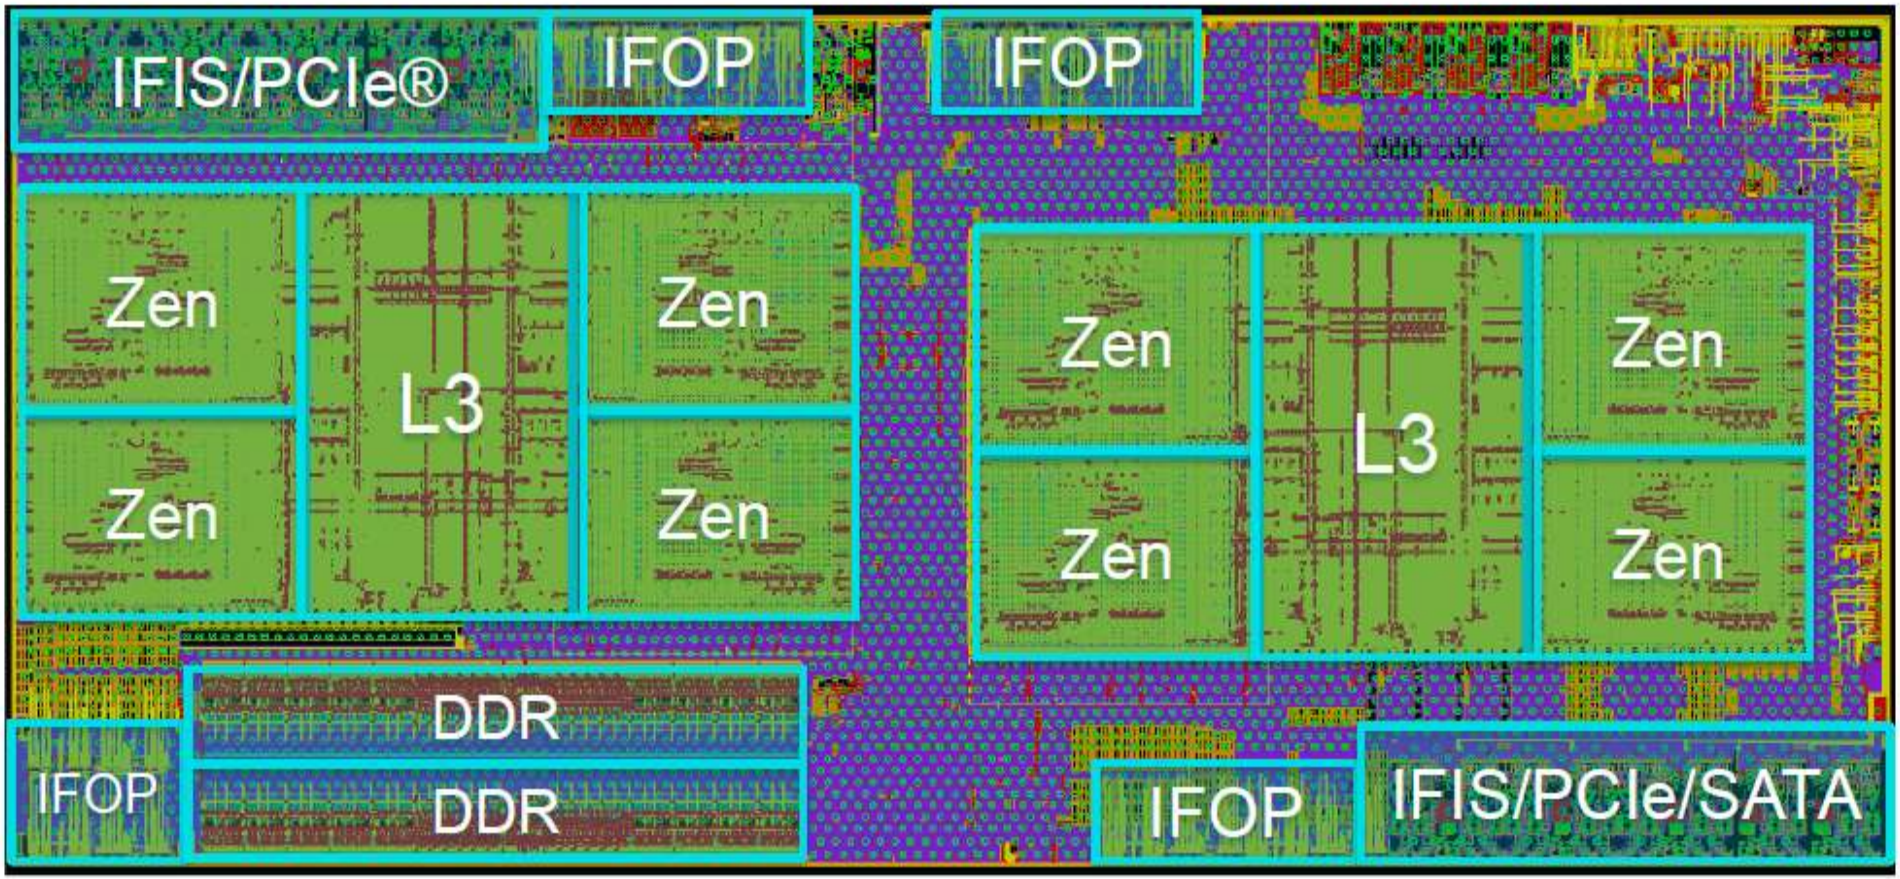

IFIS: Infinity Fabric InterSocket/PCIe Gen 3 x16
IFOP: Infinity Fabric On-Package x4

Die photograph of the 12 nm Zeppelin die (Ryzen die) [63]

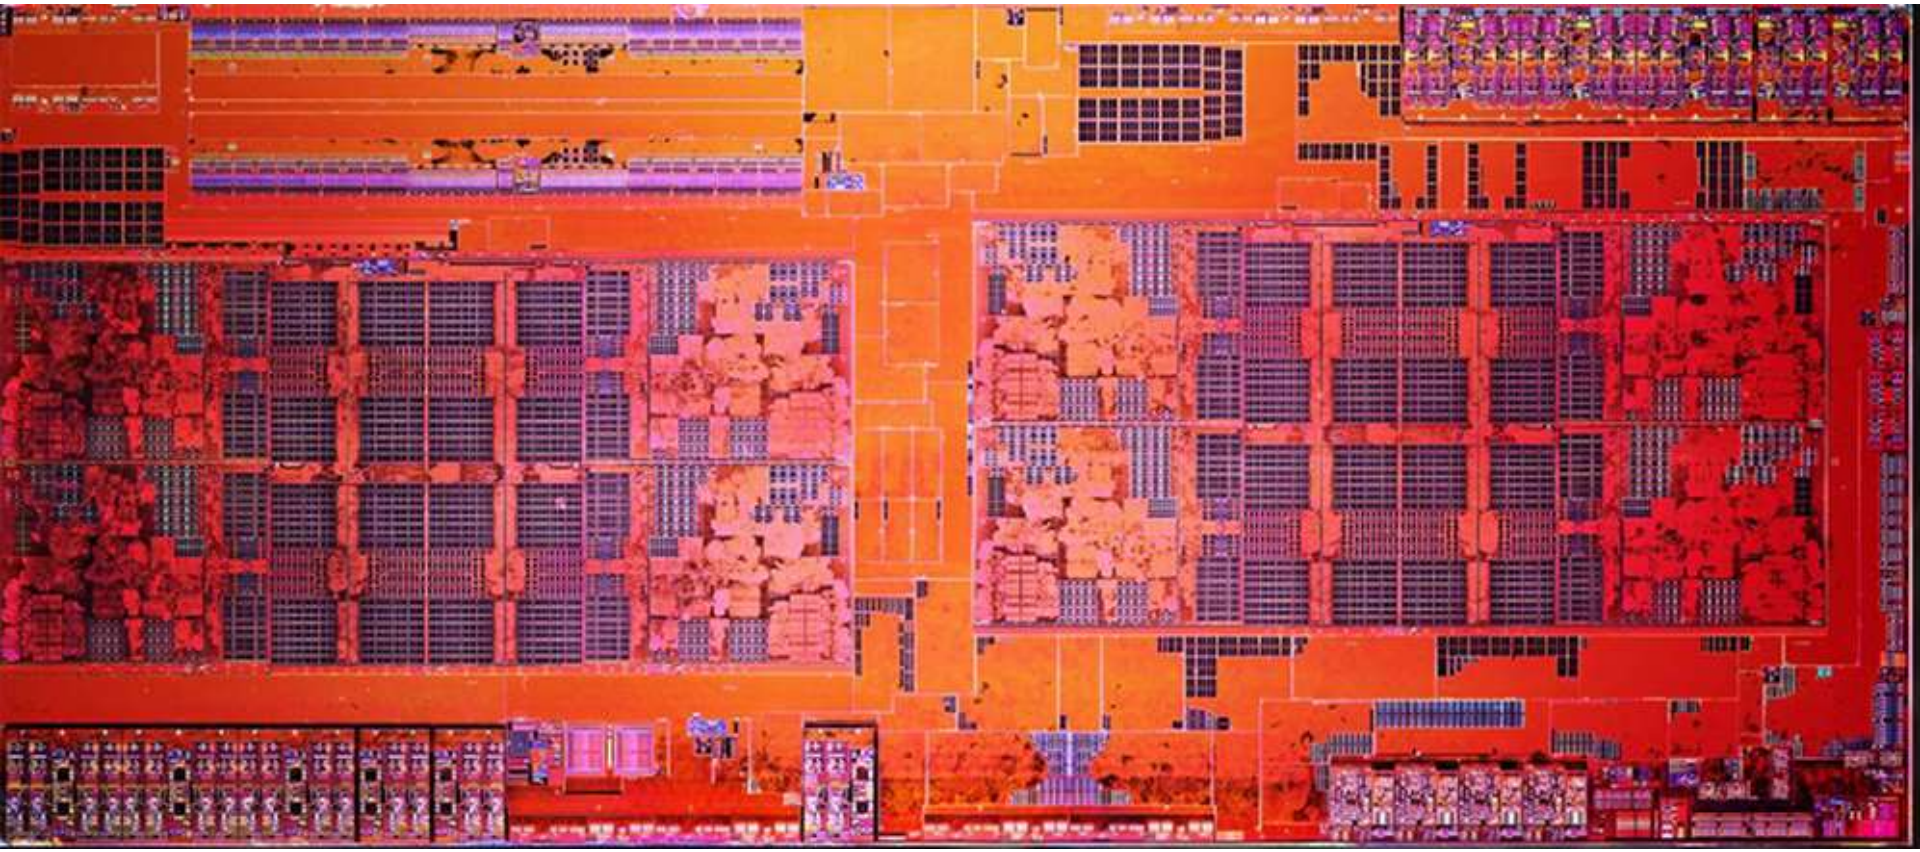

5. The Infinity Fabric

Will not be discussed

5. The Infinity Fabric (IF)

- The IF is AMD's recent interconnection fabric used for linking system units in a cache coherent way and provides also system wide control functions, such as power management or system security.
- It is the evolution of the HyperTransport bus, it is similar to Intel's QPI or UPI interconnection technology or ARM's CCN (Cache Coherent Network).
- The Infinity Fabric consists of two key parts, a Scalable Control Fabric, and a Scalable Data Fabric, as seen in the next Figure.

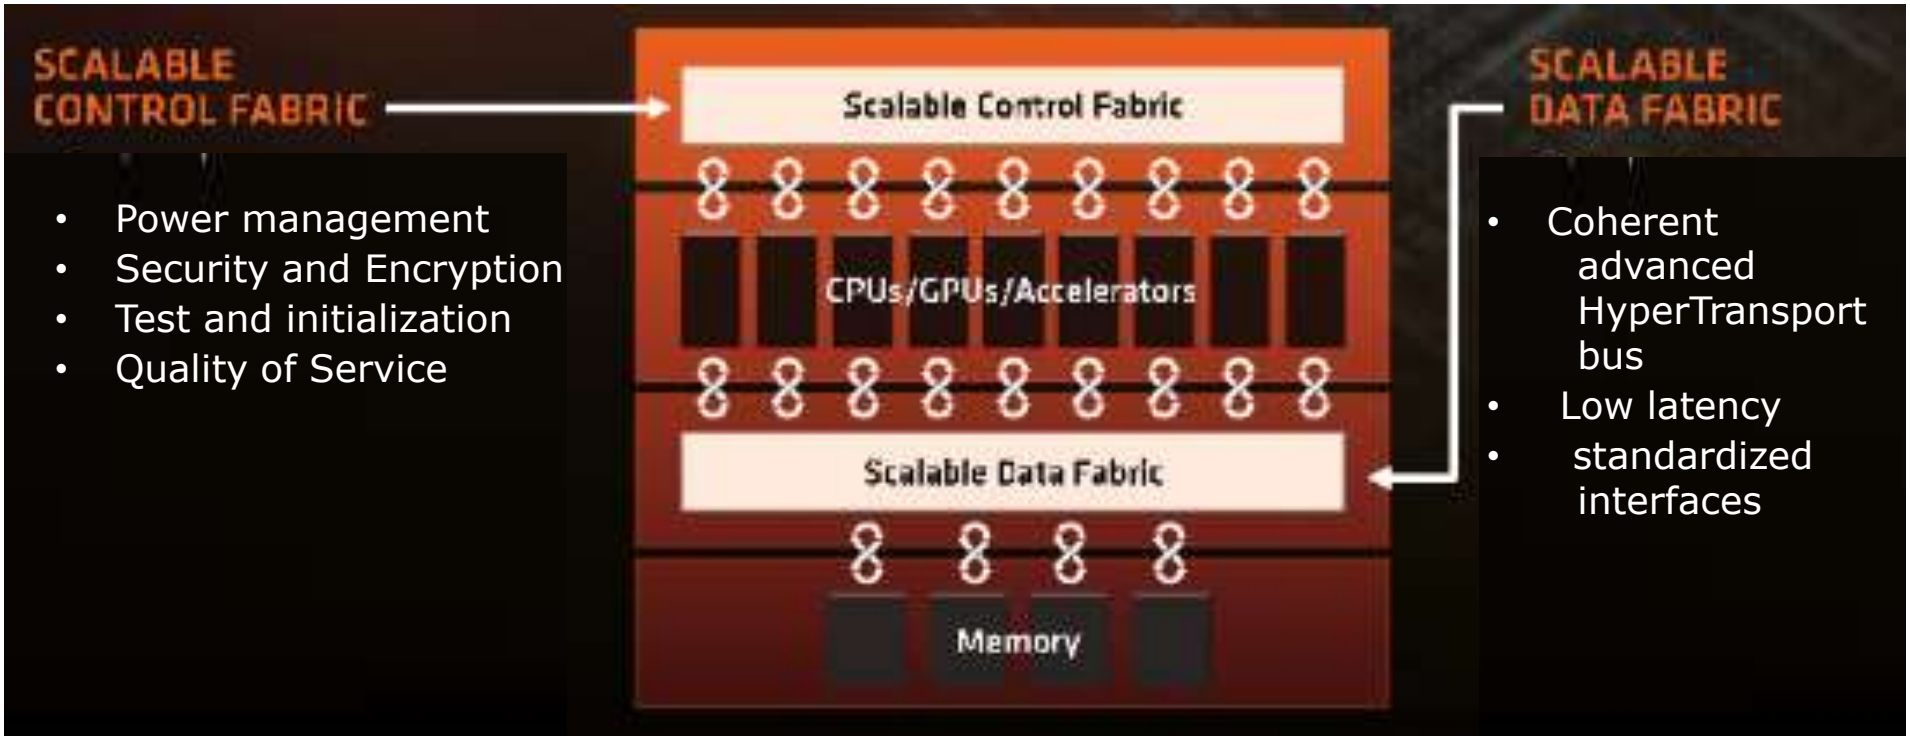

Figure: AMD's Infinity Fabric [19]

*

AMD's Infinity Fabric in the Ryzen chip in more details [41]

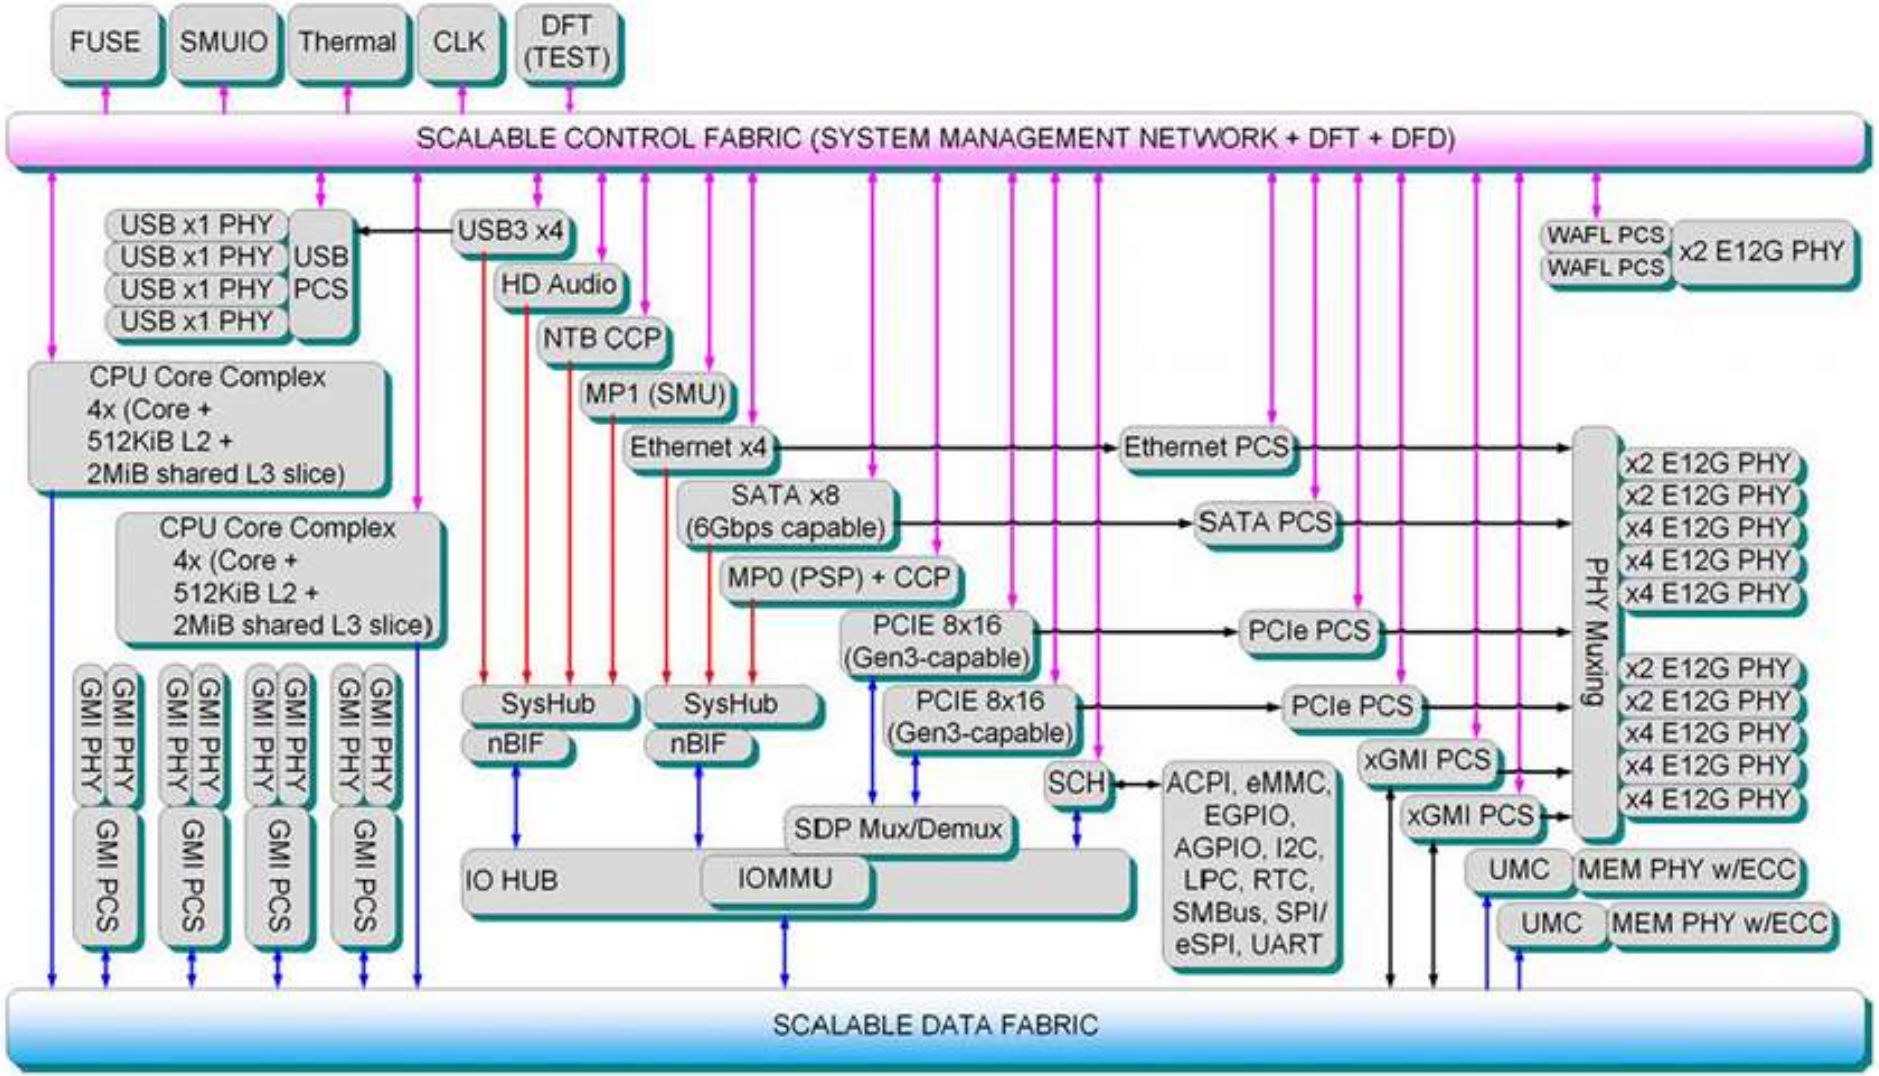

SMUIO: System Management Unit I/O registers

DFT/DFD: Design for Test and Debug

The Scalable Control Fabric -1

The Scalable Control Fabric is responsible for the system wide

- Power management
- Security and Encryption
- Test and initialization
- Quality of Service,

as indicated below.

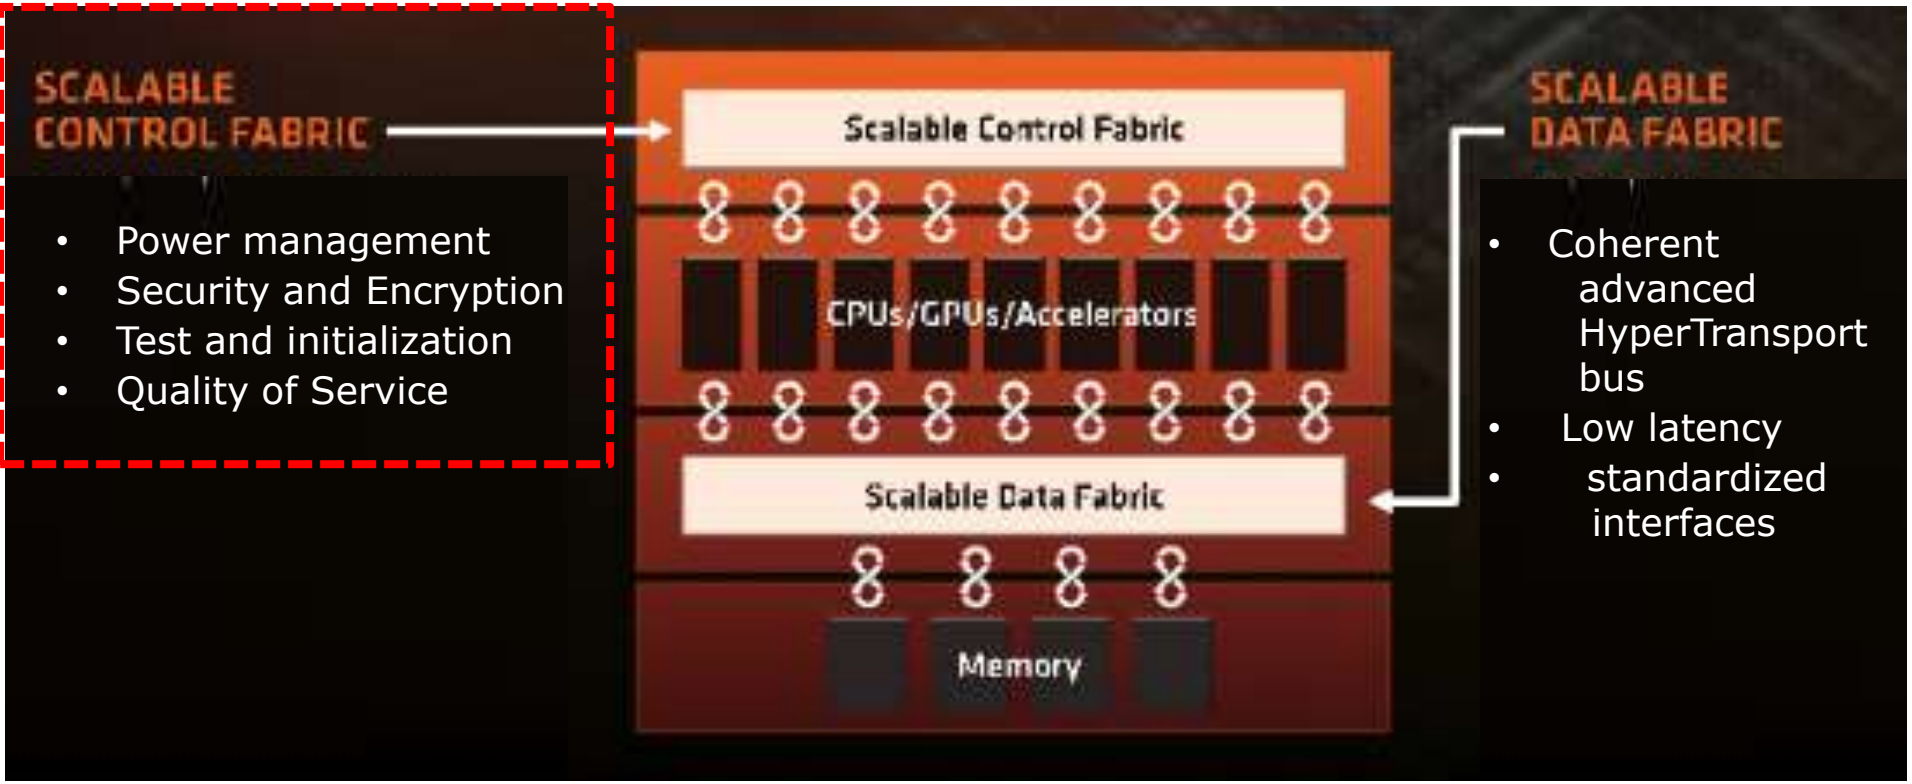

Figure: AMD's Infinity Fabric [19]

*

The Scalable Control Fabric -2

- The Scalable Control Fabric incorporates central control elements, and remote embedded elements dispersed in all blocks of the SoC.
- The remote elements are mainly sensors monitoring temperature, speed and voltage, their data is fed into the central control element and allows to optimize power management in a closed control loop, as seen in the next Figure and detailed in Section 7.

*

Principle of the operation of the Infinity System Control Fabric [3]

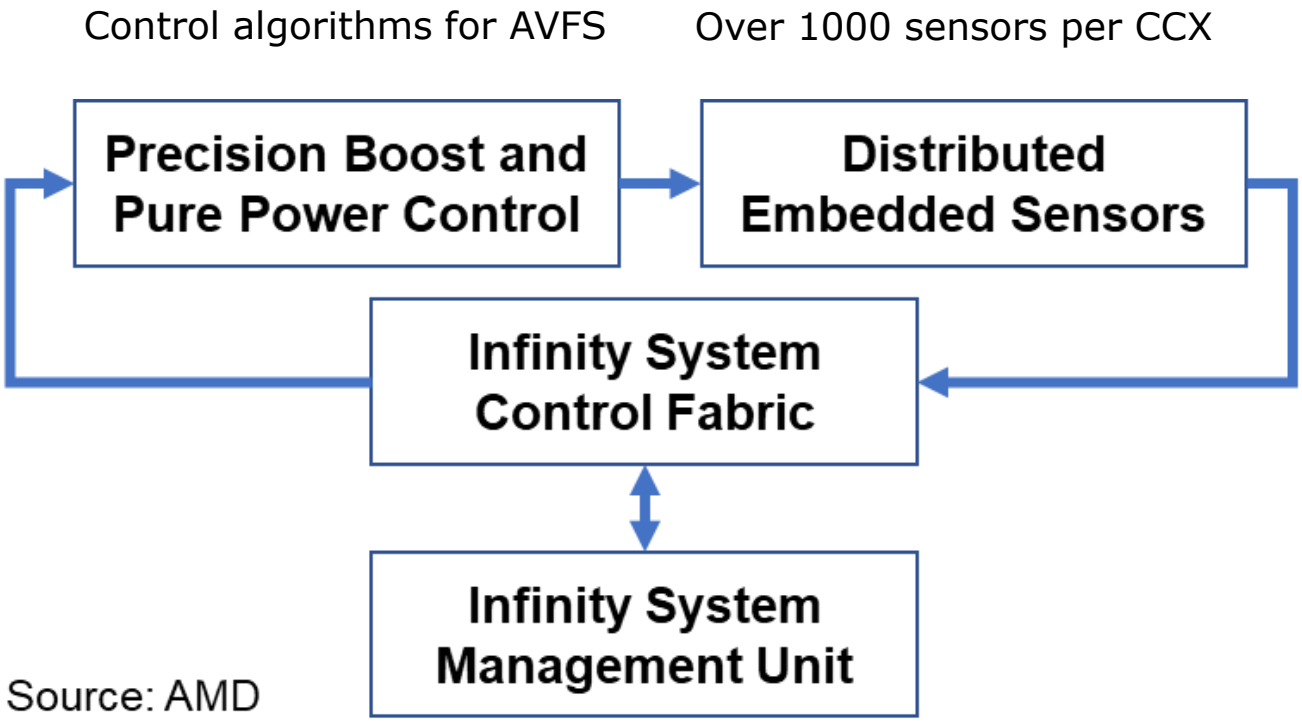

The system may be customized
by choosing the proper balance between
performance and power consumption

For details see Section 7.

The Scalable Data Fabric (SDF) -1

- The Scalable Data Fabric provides a high performance cache coherent data pathway for interconnecting all system blocks.
- It is an evolution of the HyperTransport bus.

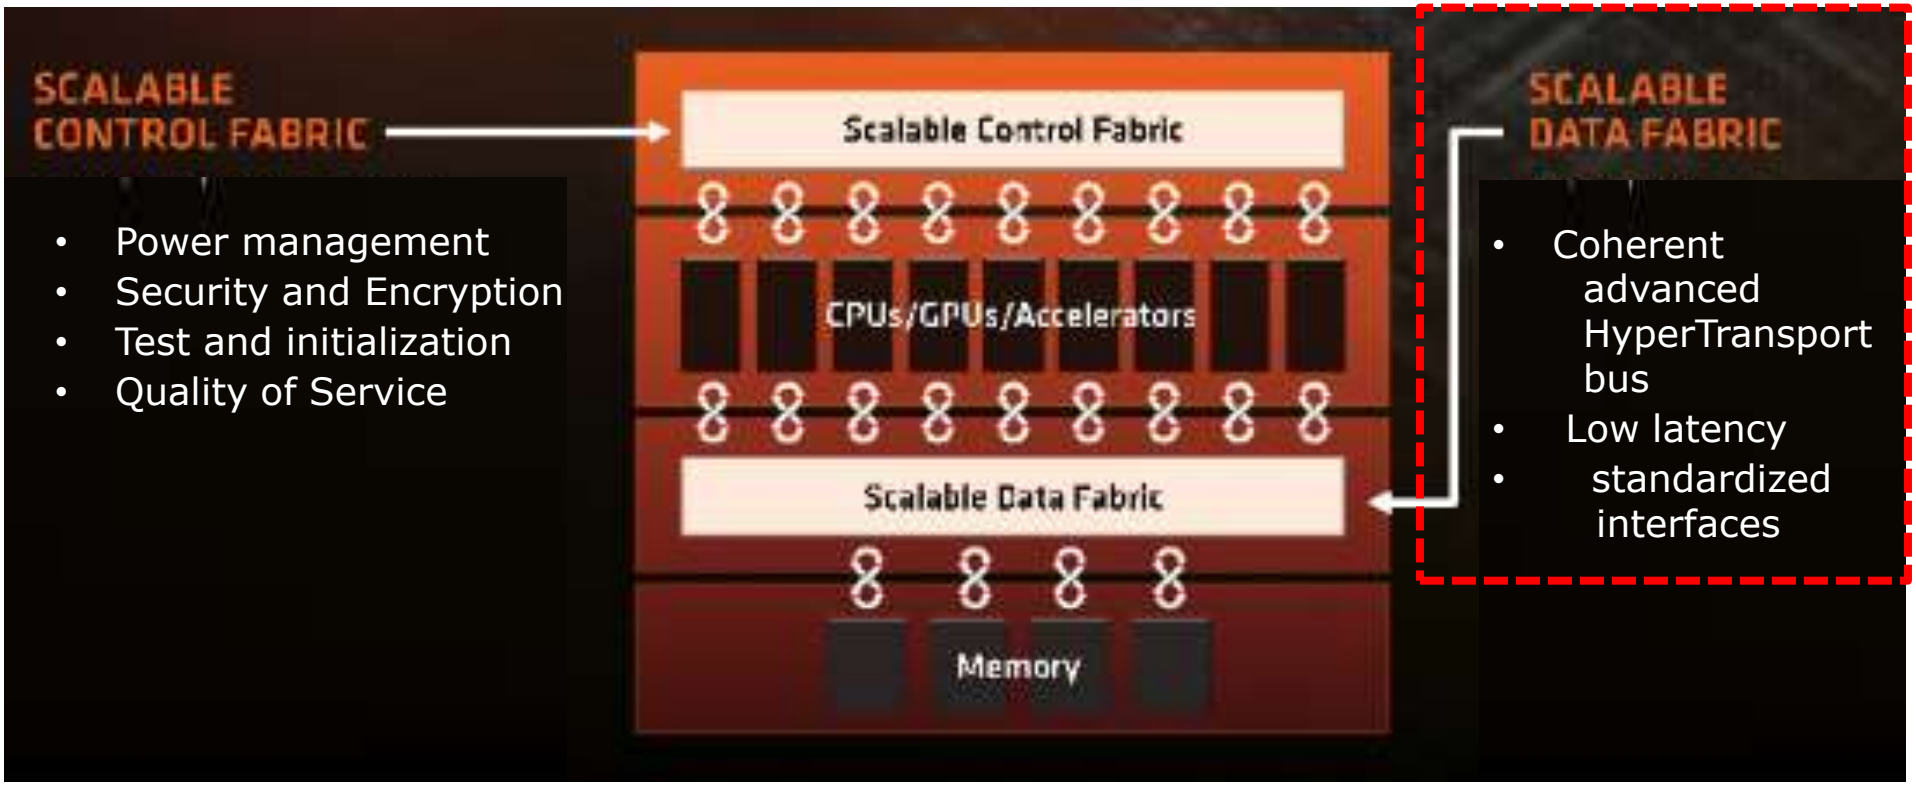

Figure: AMD's Infinity Fabric [19]

*

The Scalable Data Fabric (SDF) -2

It interconnects
- CCX building blocks within a die (not seen in the Figure)
- different dies within a socket (designated as CPUs in the Figure),
- both sockets of a 2S server or even further sockets.

as detailed subsequently.

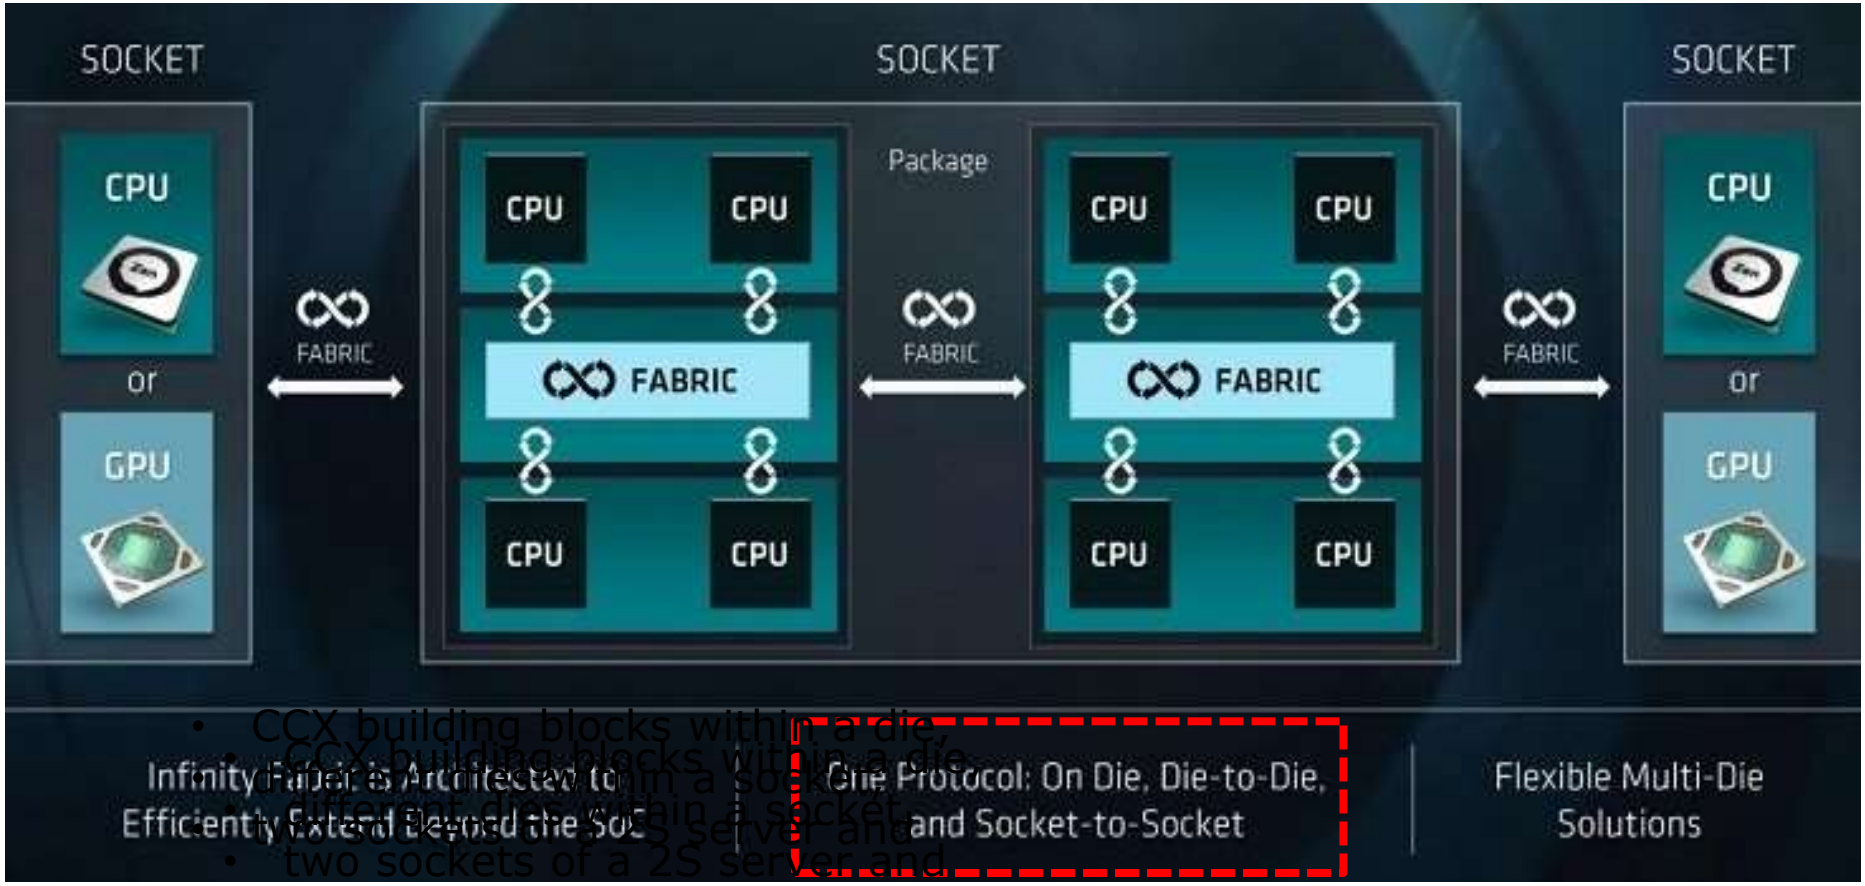

- CCX building blocks within a die,
- CCX building blocks within a die,
- different dies within a socket,
- two sockets of a 2S server and

Figure: Concept of the SDF [19]

*

5. The Infinity Fabric (8)

a) Interconnecting two CCX building blocks within a Ryzen chip [20]

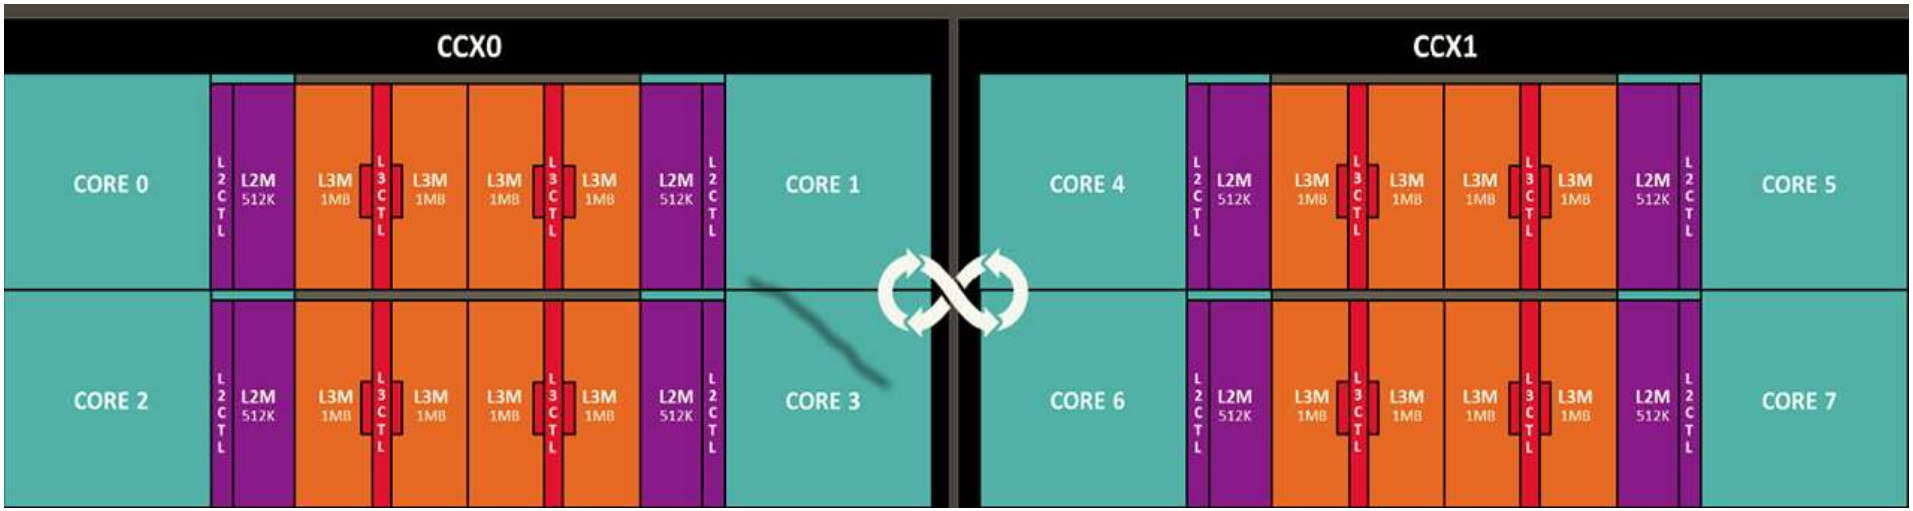

Interconnecting CCX building blocks and other units within a Ryzen chip [17]

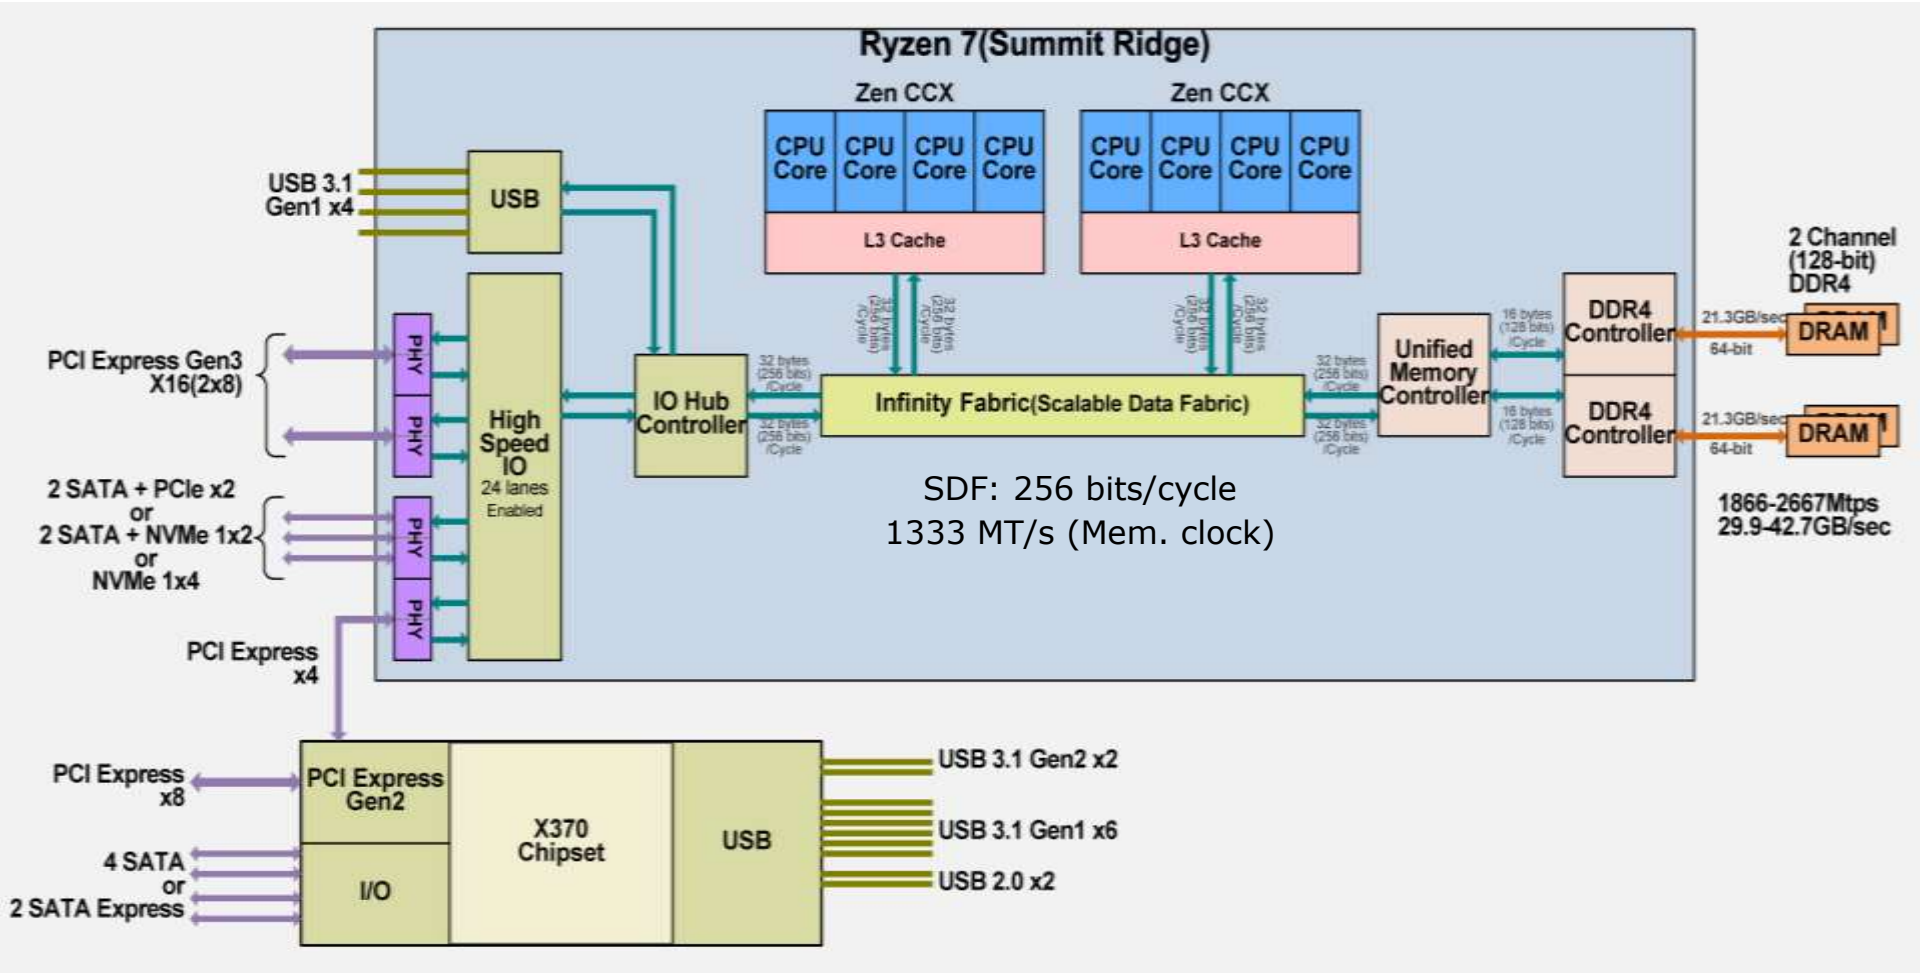

Note: Bandwidth of 2 DDR4-2666 memory channels: 2x 8x 2666 MT/s = 42.667 GB/s
Bandwidth of the SDF: 256 bit x 1333 MT/s = 42.667 GB/s (bidirectional)

Conceptual layout of the Infinity Fabric (IF) in the Zeppelin die [88]

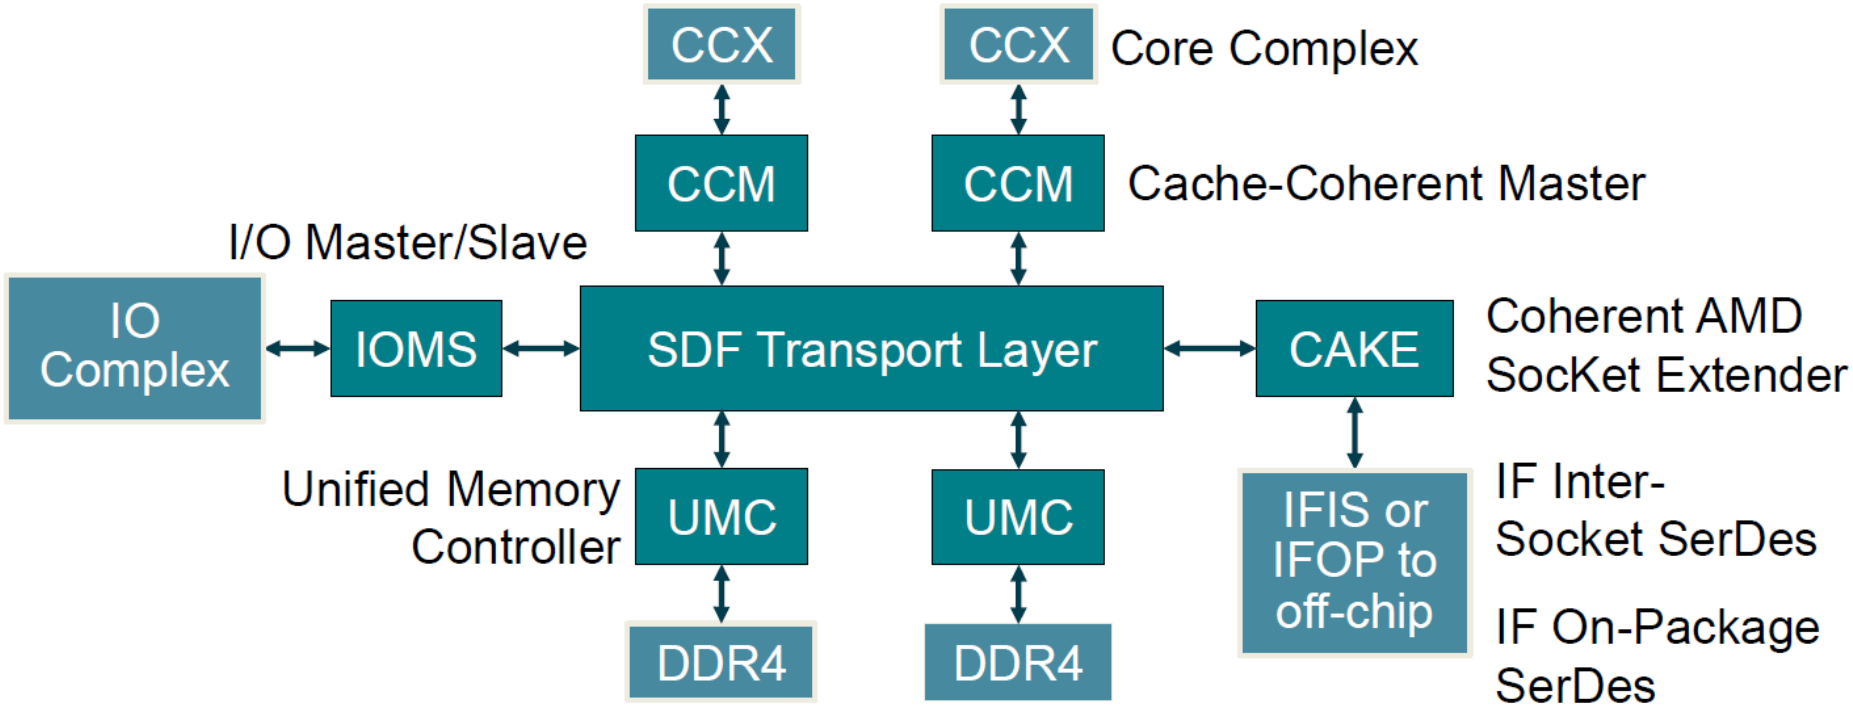

IFOP: IF On-Package interconnect links
IFIS: IF Inter-Socket interconnect links

SDF implements a crossbar-type on-die interconnect.

SerDes: (Serializer/Deserializer) interface to another chip within the system

SDF local memory access [88]

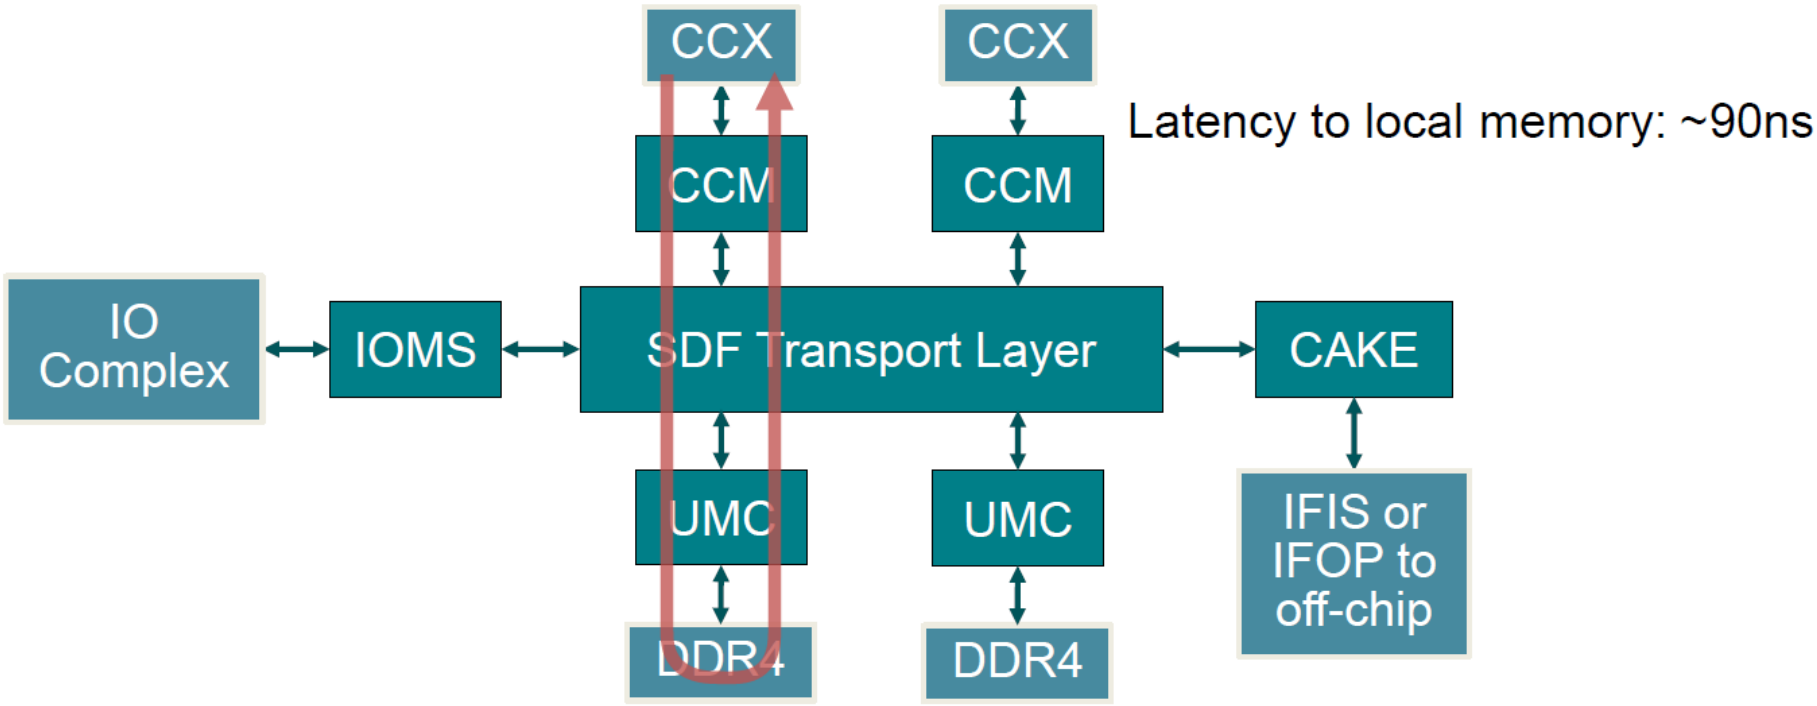

SDF remote memory access to on-die or off-die accessible memory [88]

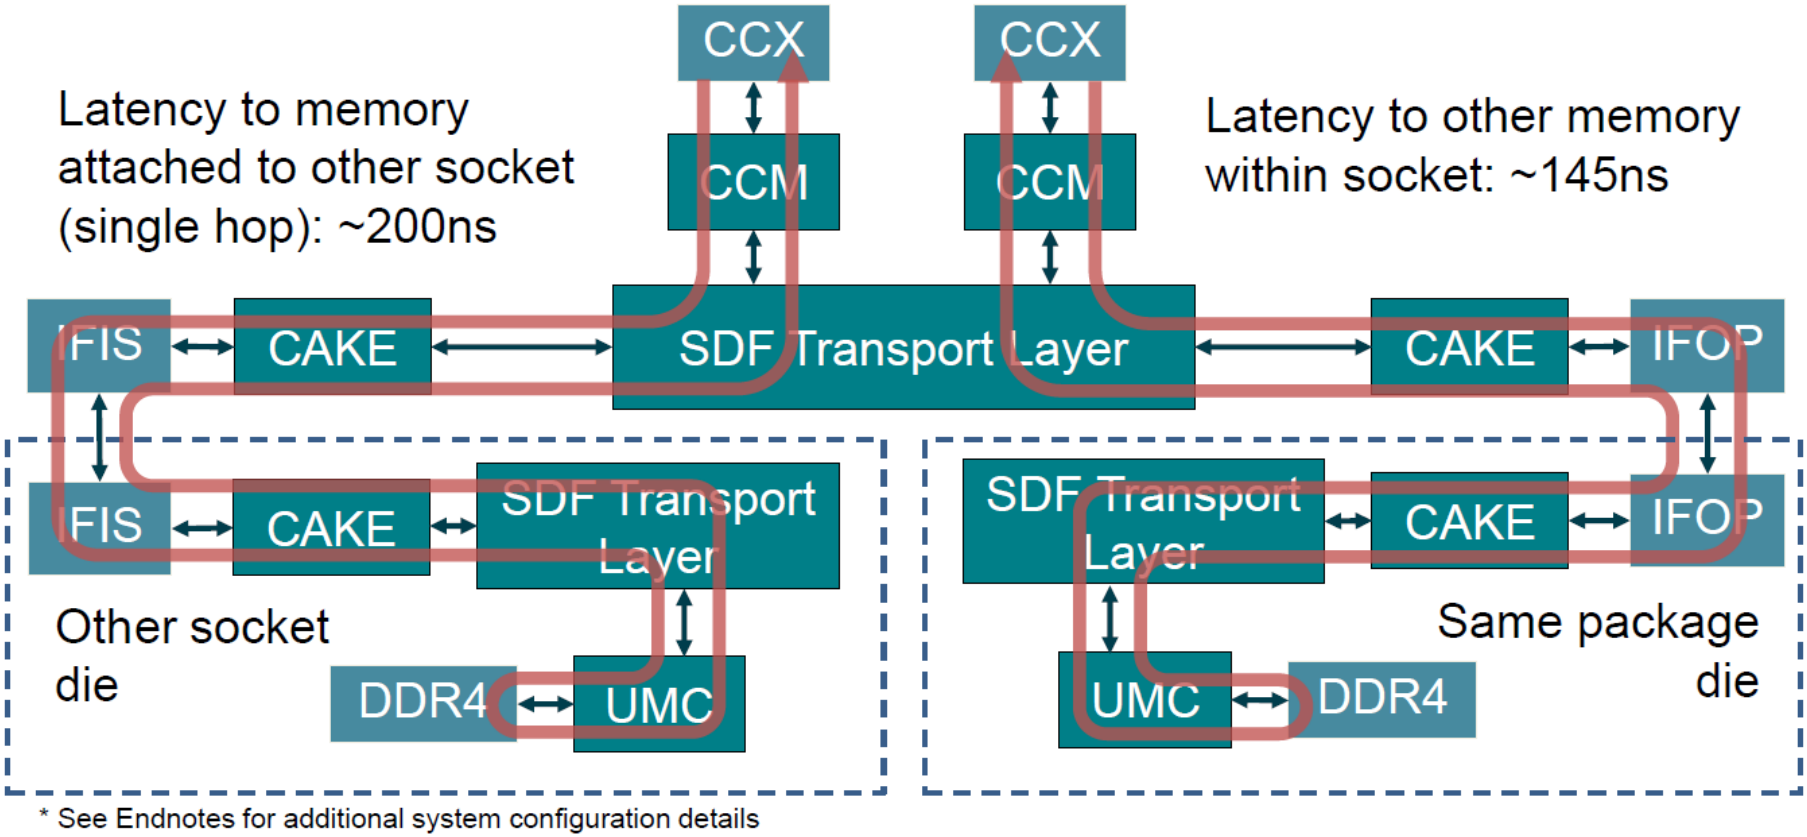

* See Endnotes for additional system configuration details

Comparison: Intel's and ARM's on-die interconnects -1

Topology of on-die interconnects

Shared bus based interconnects

Higher performance for higher nc

Crossbar-based interconnects

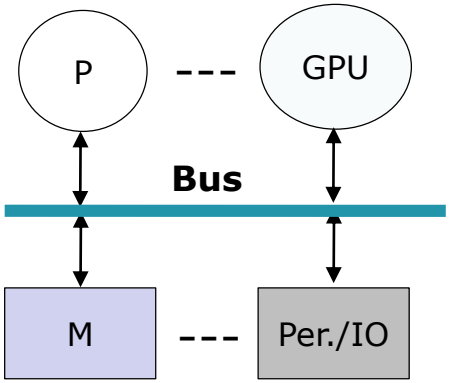

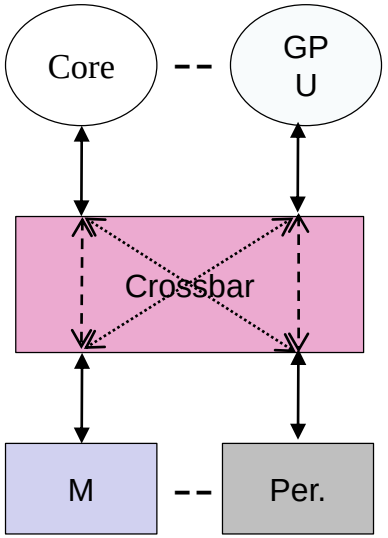

Examples

Intel		*Nehalem (CC, on QPI, 2008)*
ARM	*AHB based SoCs (non CC, ~2001) [8]*	*PL-300 interconnect (non CC, on AXI3, 2004)*
		through the CCI 550 interconnect (CC, on ACE, 2015)

CC: Cache Coherent

Comparison: Intel's and ARM's on-die interconnects -2

Topology of on-die interconnects

Less cost for higher nc

Ring bus-based interconnect topology

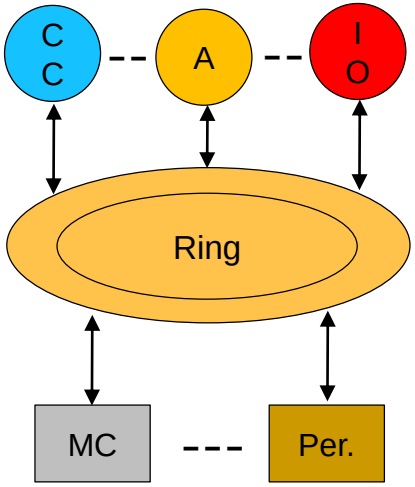

Shorter access time for higher nc

2D mesh-based interconnect topology

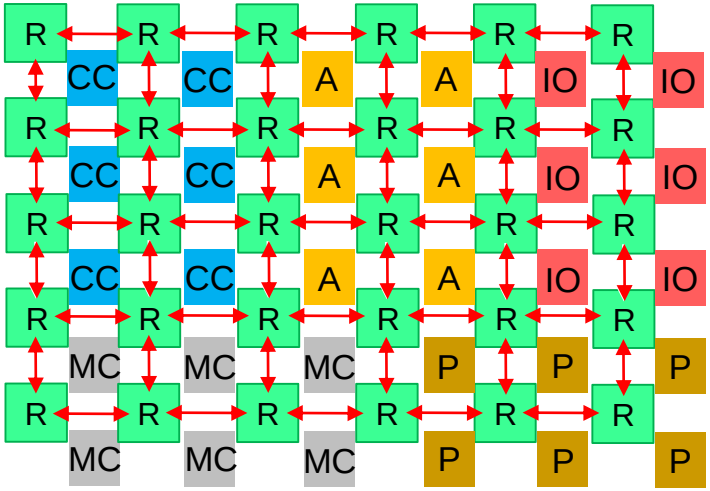

CC: Core Cluster
A: Accelerator (e.g. GPU, NPU)
nc: Core count

Examples

Intel	*Sandy Bridge (CC, on QPI, 2011)*	*Skylake-SP for servers (CC, UPI, 2017)* *Knights Landing for AI, (2016)*
ARM	*CCN-5xx interconnect for servers (CC, on CHI, 2012)*	*CMN-6xx Interconnect for servers (CC, on CHI, 2017)*

**

Example: ARM's CCI-400 crossbar interconnect fabric [46]

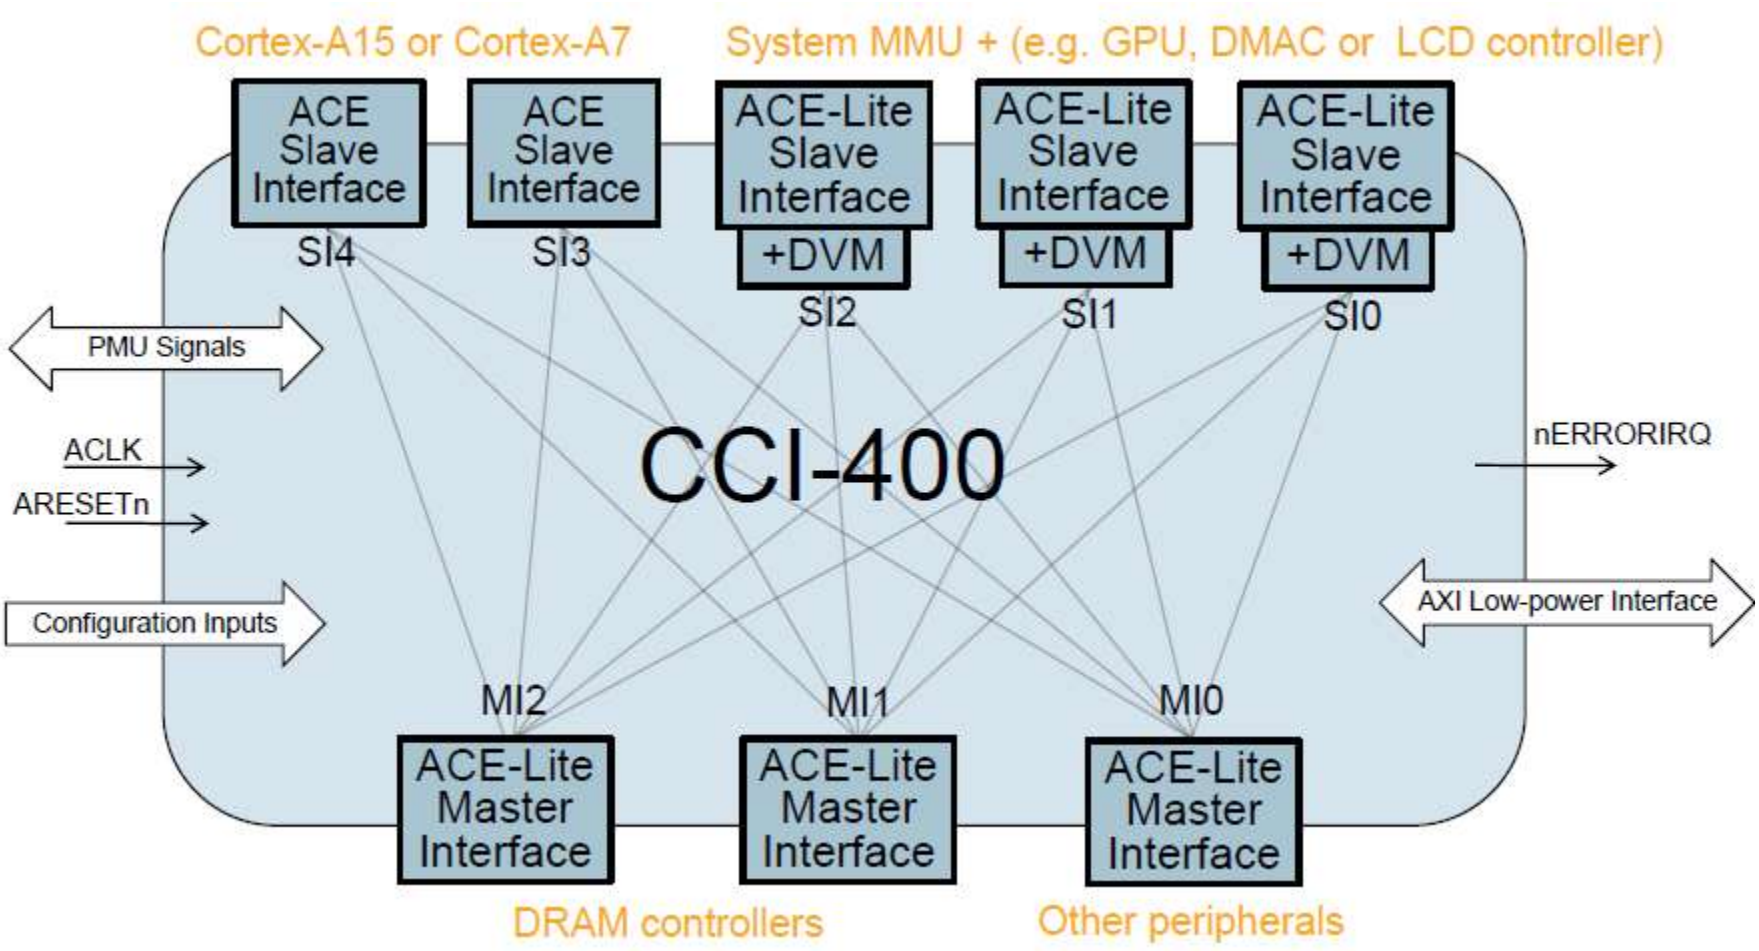

ACE: (AXI Coherency Extensions). ARM's coherent interface specification developed for core clusters, introduced in 2011

Example: ARM's CCN-504 ring interconnect fabric (dubbed Dickens) [47]

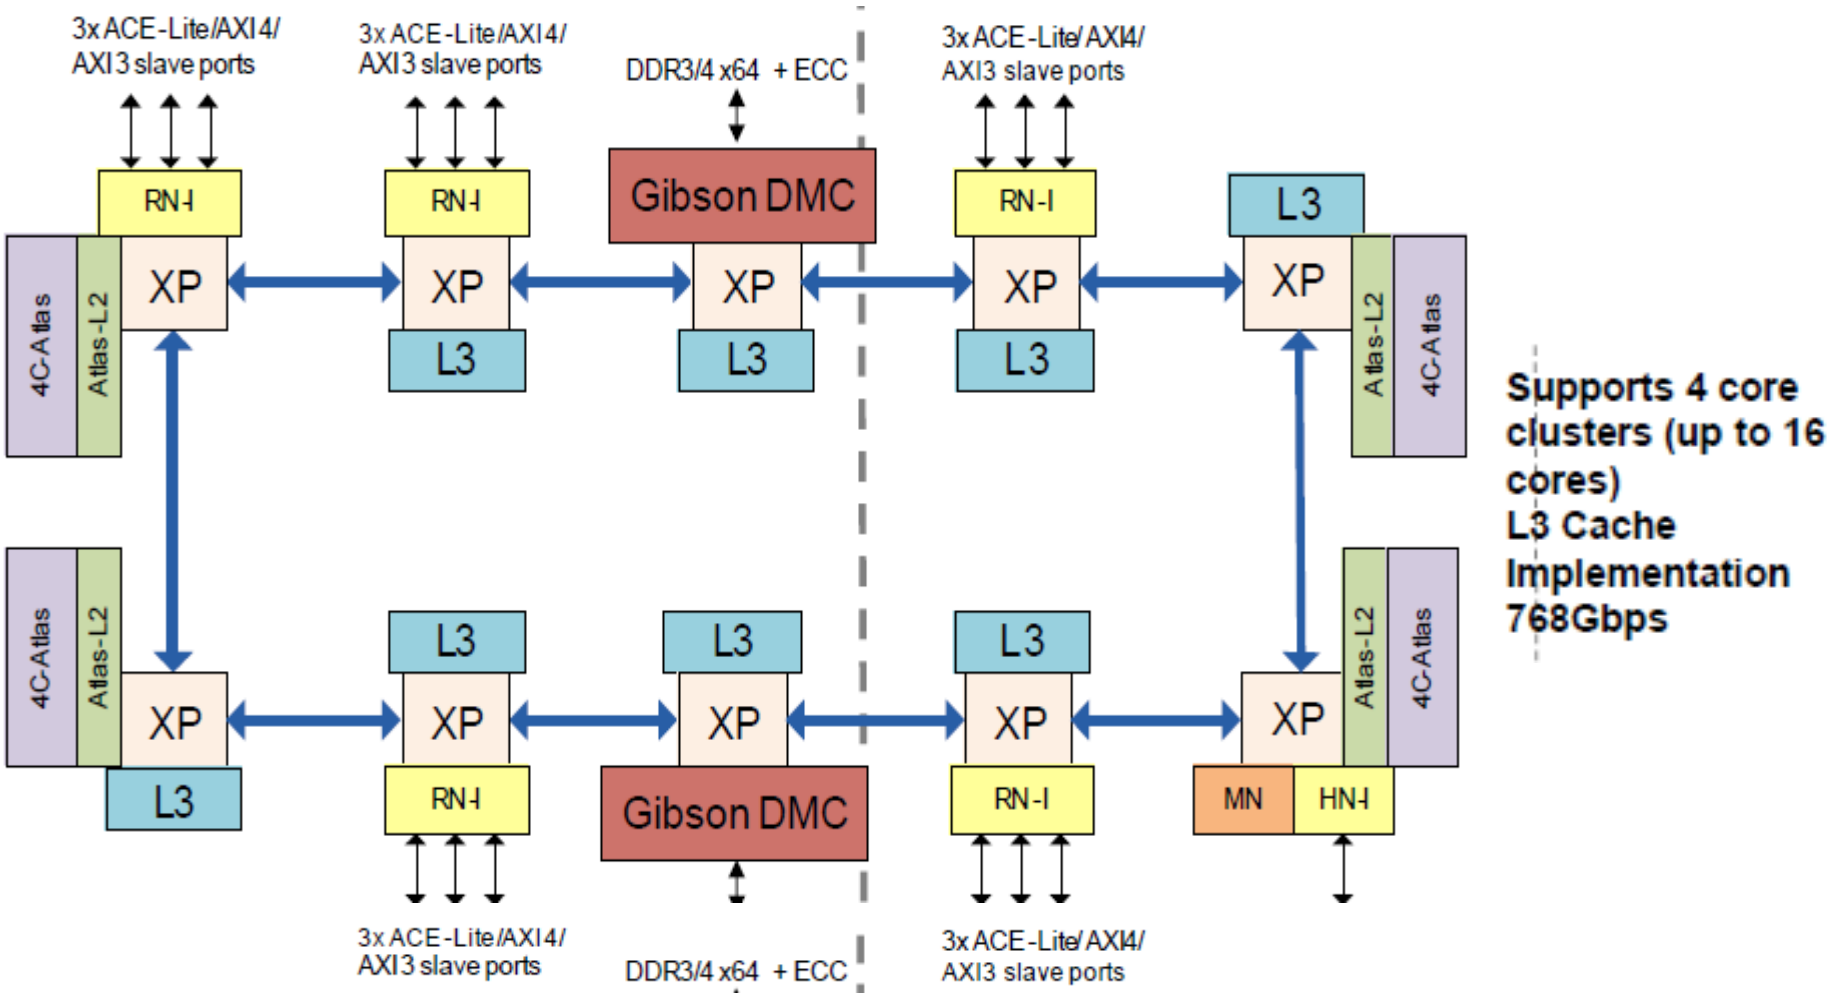

Remark: The Figure indicates only 15 ACE-Lite slave ports and 1 master port whereas ARM's specifications show 18 ACE-Lite slave ports and 2 master ports.

Example: Use of ARM's cache coherent CCN-504 interconnect in a server system [49]

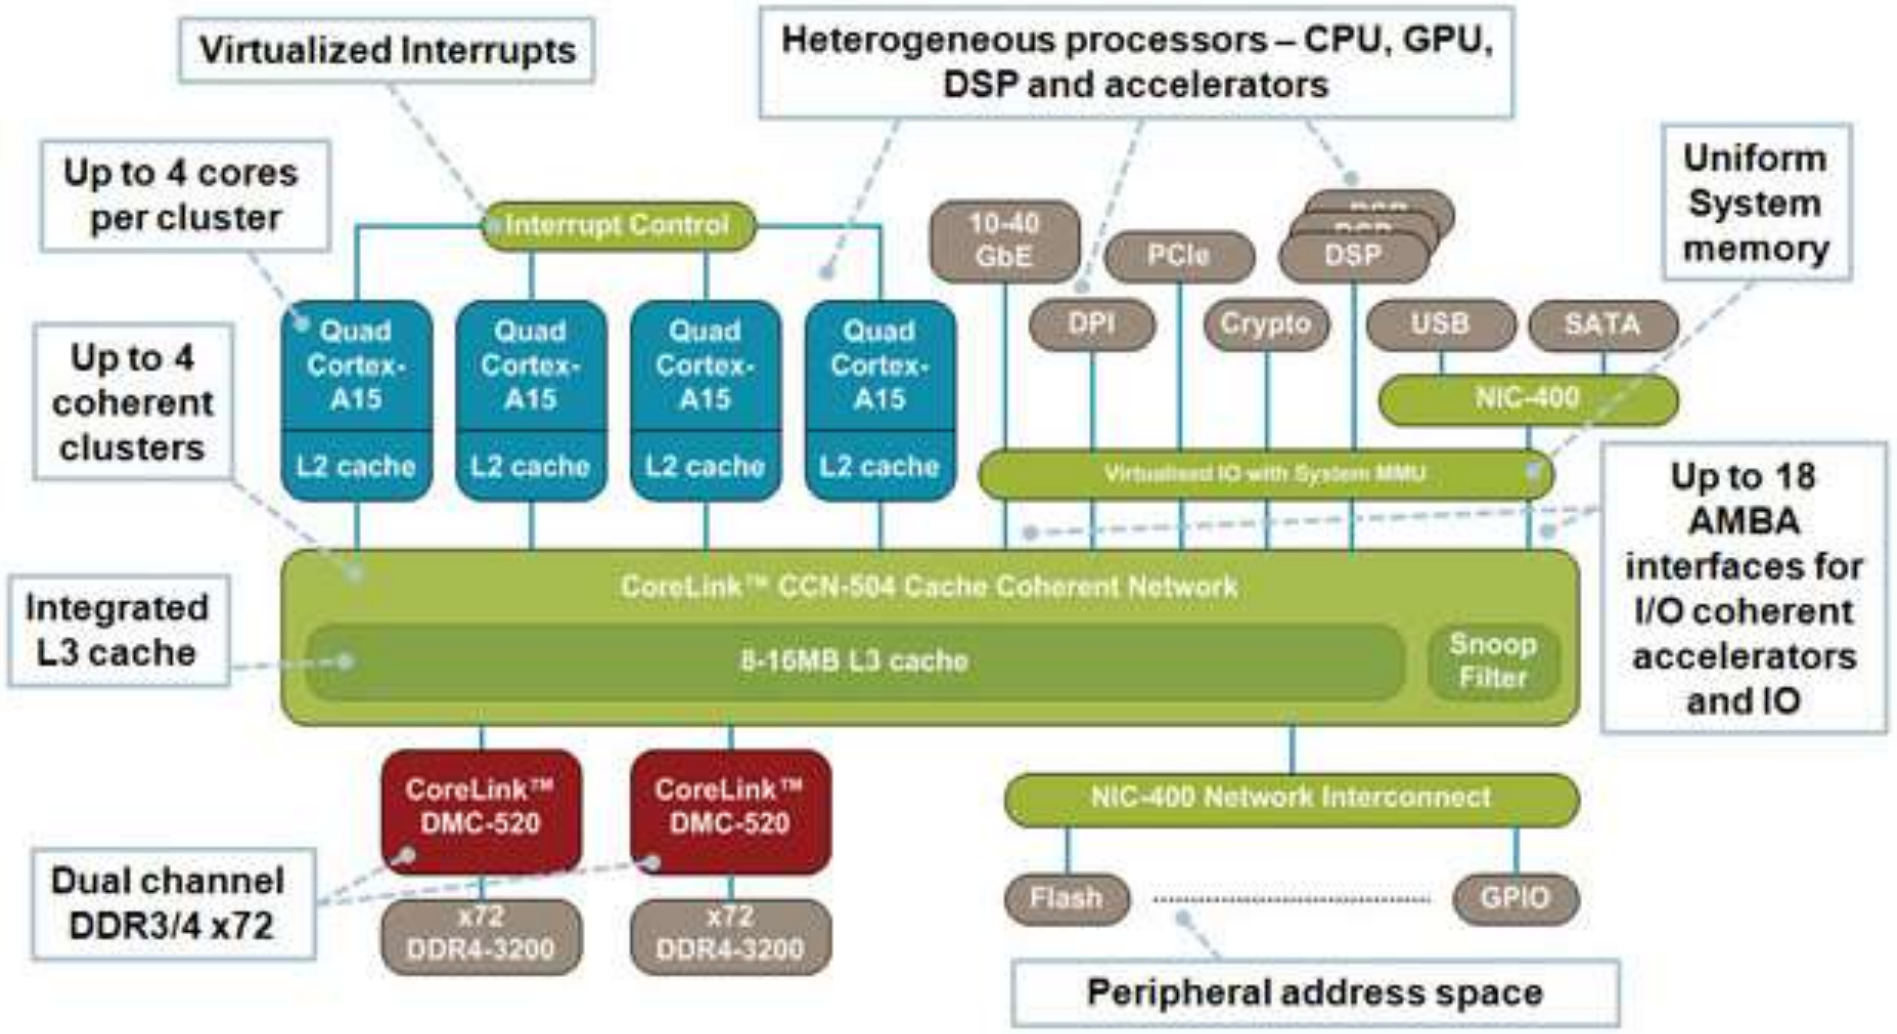

b) Die-to-die interconnect in an Epyc server processor [21] (Dies are Zeppelin modules)

INFINITY FABRIC: DIE-TO-DIE INTERCONNECT

- Fully connected Coherent Infinity Fabric within socket
- Optimized low power, low latency MCM links between die
- 42GB/sec bi-dir BW per link, ~2pJ/bit TDP
- Single-ended, low power zero transmission
- Infinity Control Fabric connected between dies

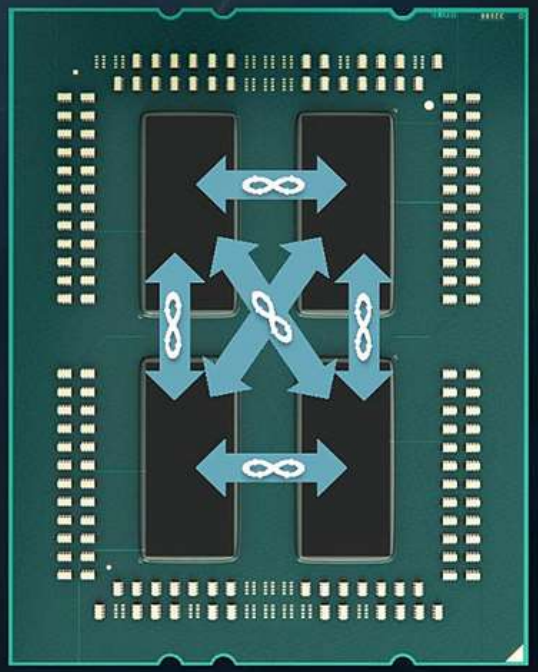

Purpose-built MCM links optimized for power, bandwidth, and latency

Revealed in AMD's Epyc Launch event presentation, June 20 2017

Die-to-die interconnect links, called IFOP (IF On-Package) interconnect links
e.g. in an Epyc server processor - detailed [7]

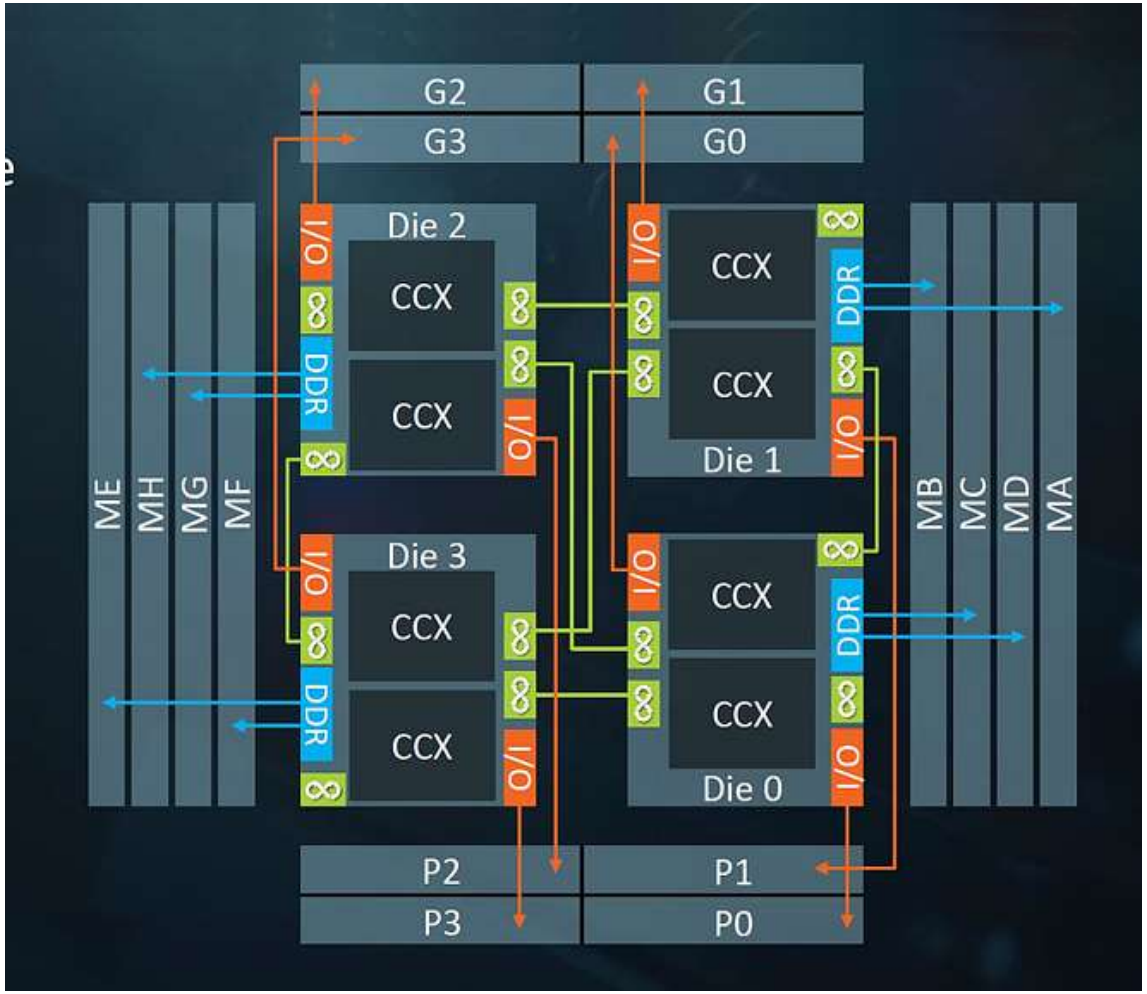

G0 - G3: 4 x 16 SERDES lanes
(in each direction)
Used for socket-to-socket
communication or PCIe 3.0 lanes

MA-MH: 8 x 64+8-bit DDR4
DRAM channels

∞: Die-to-die IF links (IFOP links)
32-bit per link per direction
interconnecting the chips
in a crossbar fashion
by single-ended signals
Differential clocking
(2 pJ/bit)
4 transfers/mem. clock

P0 - P3: 4 x 16 SERDES lanes
(in each direction)
Used to implement PCIe 3.0 lanes

Bi-dir. BW/link for DDR4-2666: 2x 4 x 4 B x 1333 MT/s = 42.67 GB/s

Die-to-die interconnect through the IFOP (IF On-Package) interconnect link within a package [43]

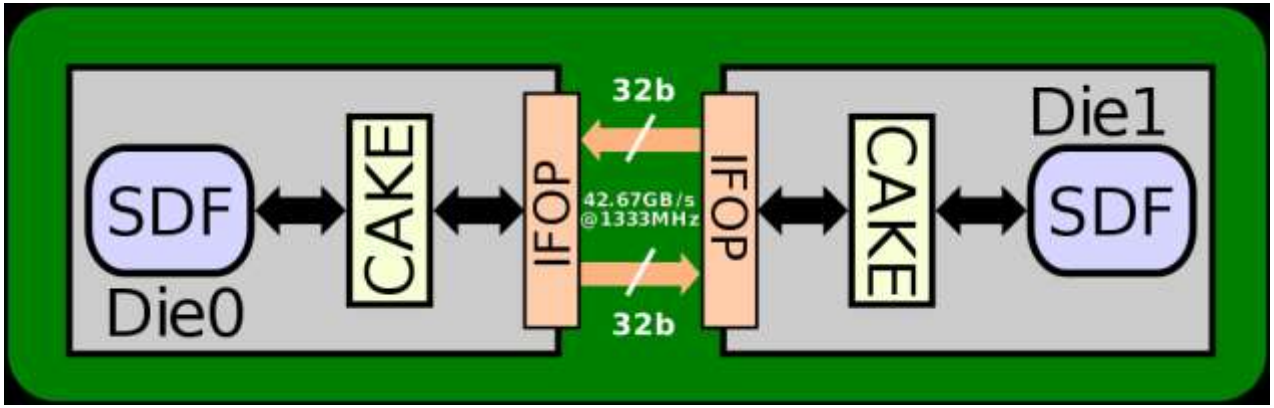

Die-to-die IF links (IFOP links)

- 32-bit per link per direction interconnecting the chips in a crossbar fashion by single-ended signals (2 pJ/bit)
- Differential clocking
- 4 transfers/mem. clock

Implementation of the IFOP (IF On-Package) interconnect links [88]

- 2pJ/bit power consumption
- Low-swing, single-ended data for ~50% of power of an equivalent differential driver
- Zero power driver state during logic 0 transmit
  - Transmit/receive impedance termination to ground while driver pullup is disabled
  - Also applied during link idle
- Data bit inversion encoding saving 10% average power per bit

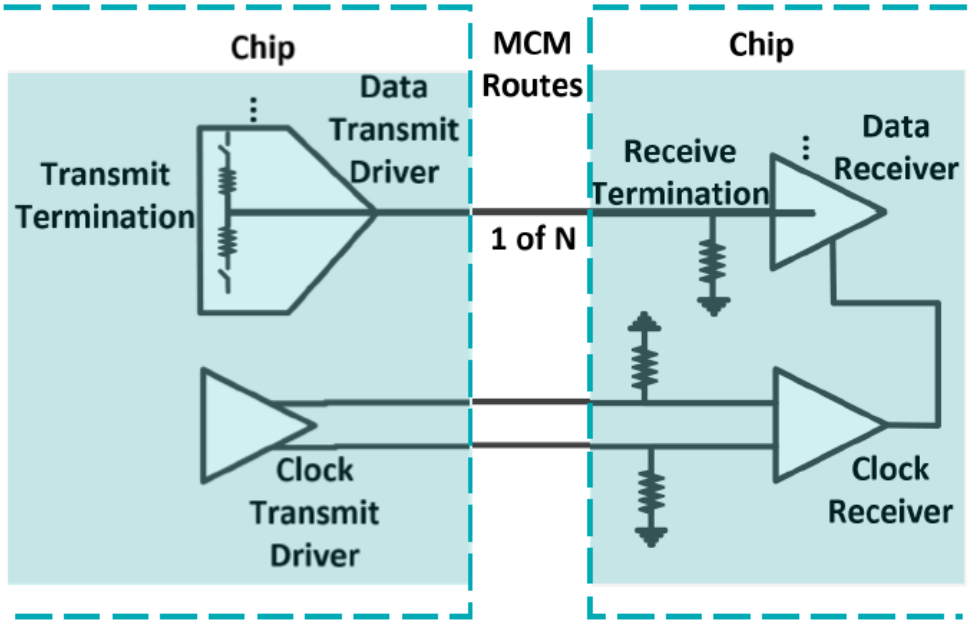

Implementation of the IFOP (IF On-Package) interconnect links for both directions [87]

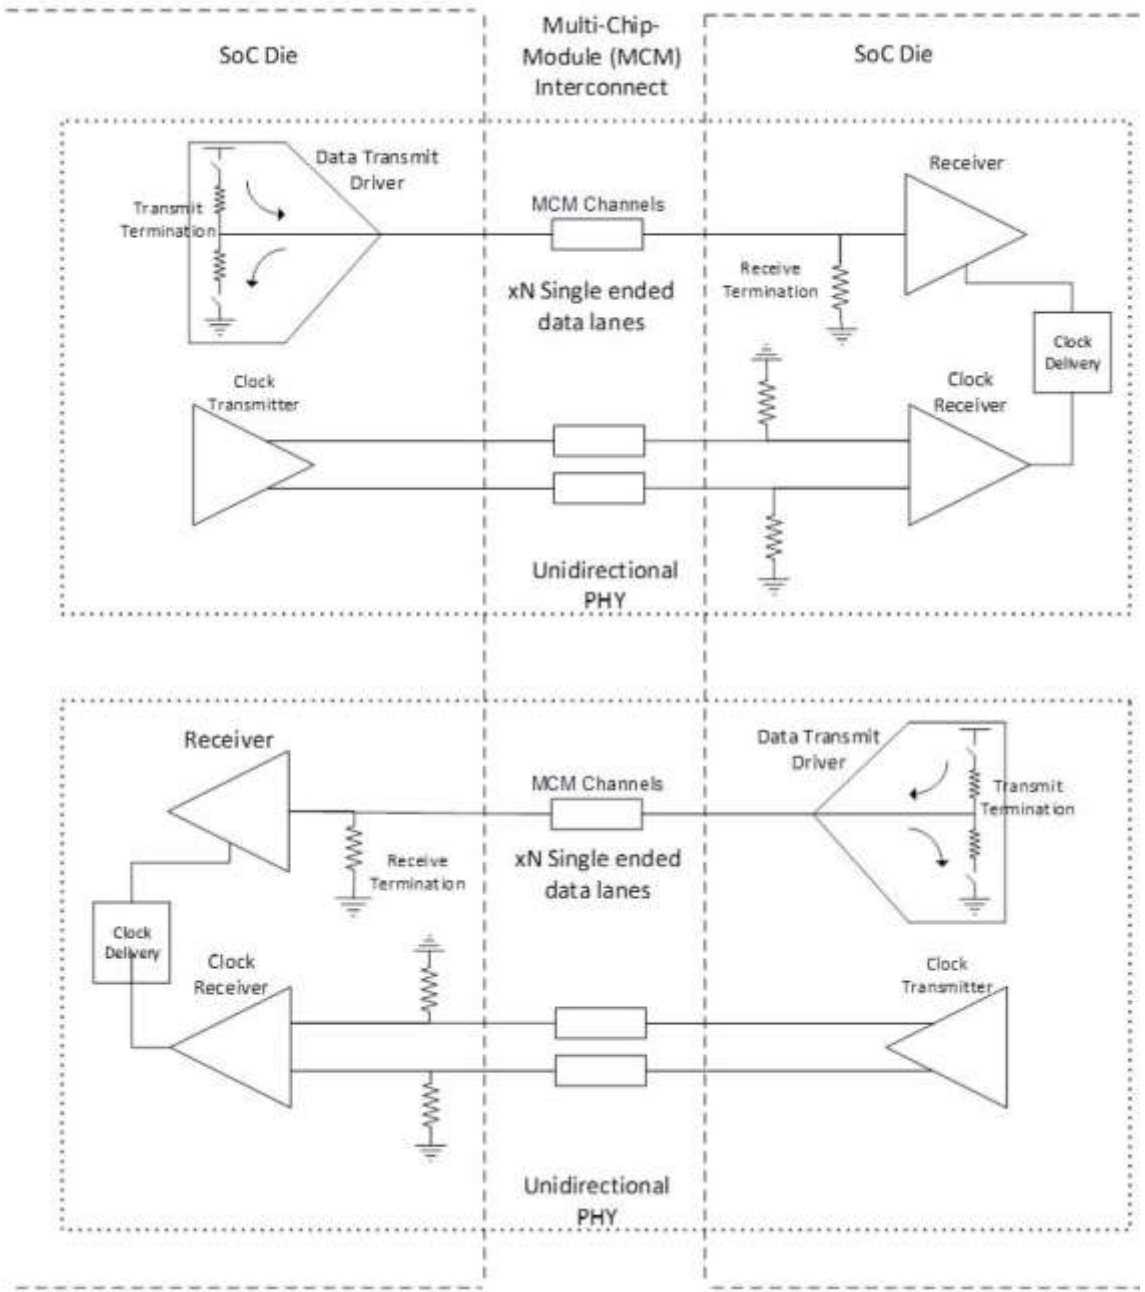

c) Socket-to-socket interconnect links, called IFIS (IF Inter-Socket) interconnect links in a 2S Epyc server [21]

INFINITY FABRIC: SOCKET-TO-SOCKET INTERCONNECT

- 4 links between sockets in 2P
  - Each CPU die connected to peer die on second socket
  - 2-hop maximum system diameter
- 38GB/s bi-dir BW per link, 152GB/s between sockets, ~9pJ/bit TDP
- Infinity Control Fabric connected between sockets

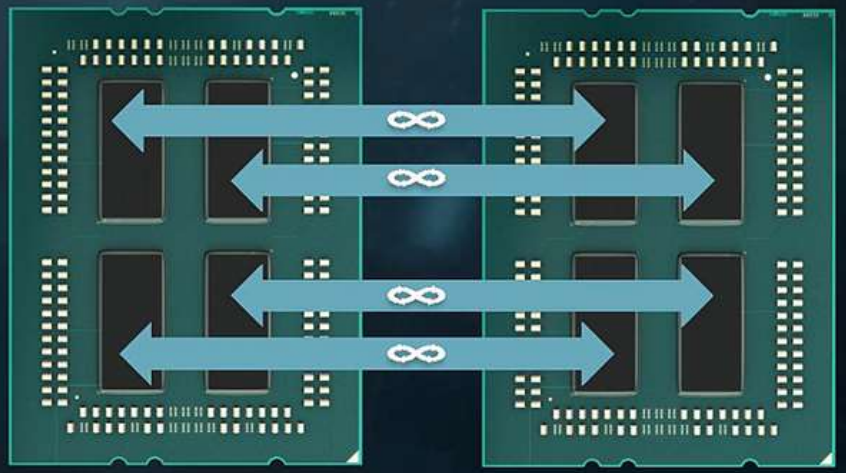

LOW POWER, HIGH BW
2x Maximum socket-to-socket links vs E5-2699v4*

Revealed in AMD's Epyc Launch event presentation, June 20 2017

Socket-toSocket interconnect through the IFIS (IF InterSocket) links between two packages [43]

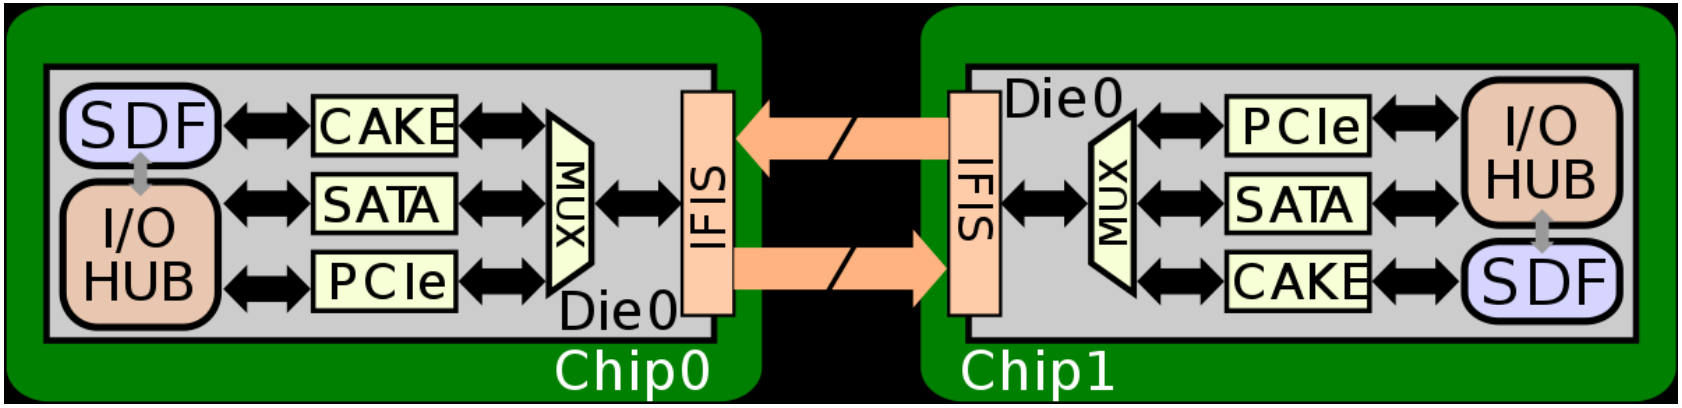

Socket-to-Socket IF links (IFIS links)

- 4x16 lanes per direction, 8 transfers/memory clock
- Differential signaling, (9 pJ/bit according to [43] or 11 pJ/bit according to [87])
- In-band CRC, 8/9 bandwidth
- BW for DDR4-2666: 8/9x 8x 2B x 1333 MT/s 2 (bi-dir.) = 37.92 GB/s

Socket-to-socket interconnect links, called IFIS (IF Inter-Socket) interconnect links in a 2S Epyc server – detailed [43]

4x 16 lanes per direction, 8 transfers/memory clock
Differential signaling, (9 pJ/bit according to [43] or 11 pJ/bit according to [87])
In-band CRC, 8/9 bandwidth
BW for DDR4-2666: 8/9x 8x 2B x 1333 MT/s 2 (bi-dir.) = 37.92 GB/s

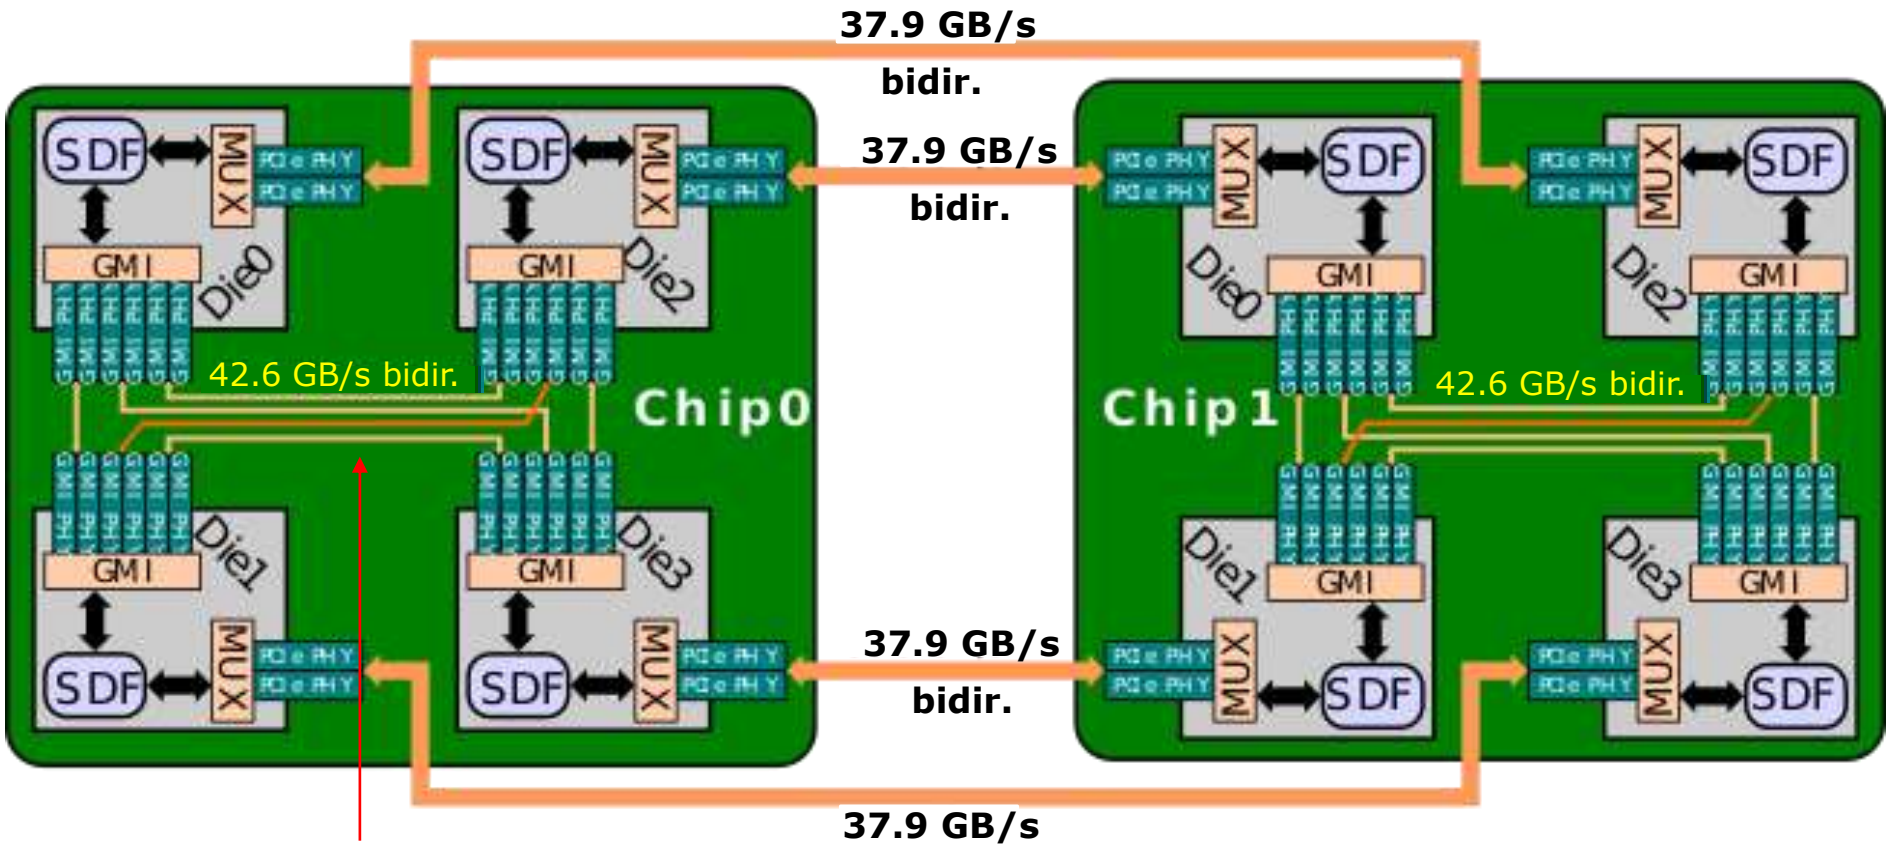

Balancing bandwidths in a 2S Epyc server chip with DDR4-2667 [21]

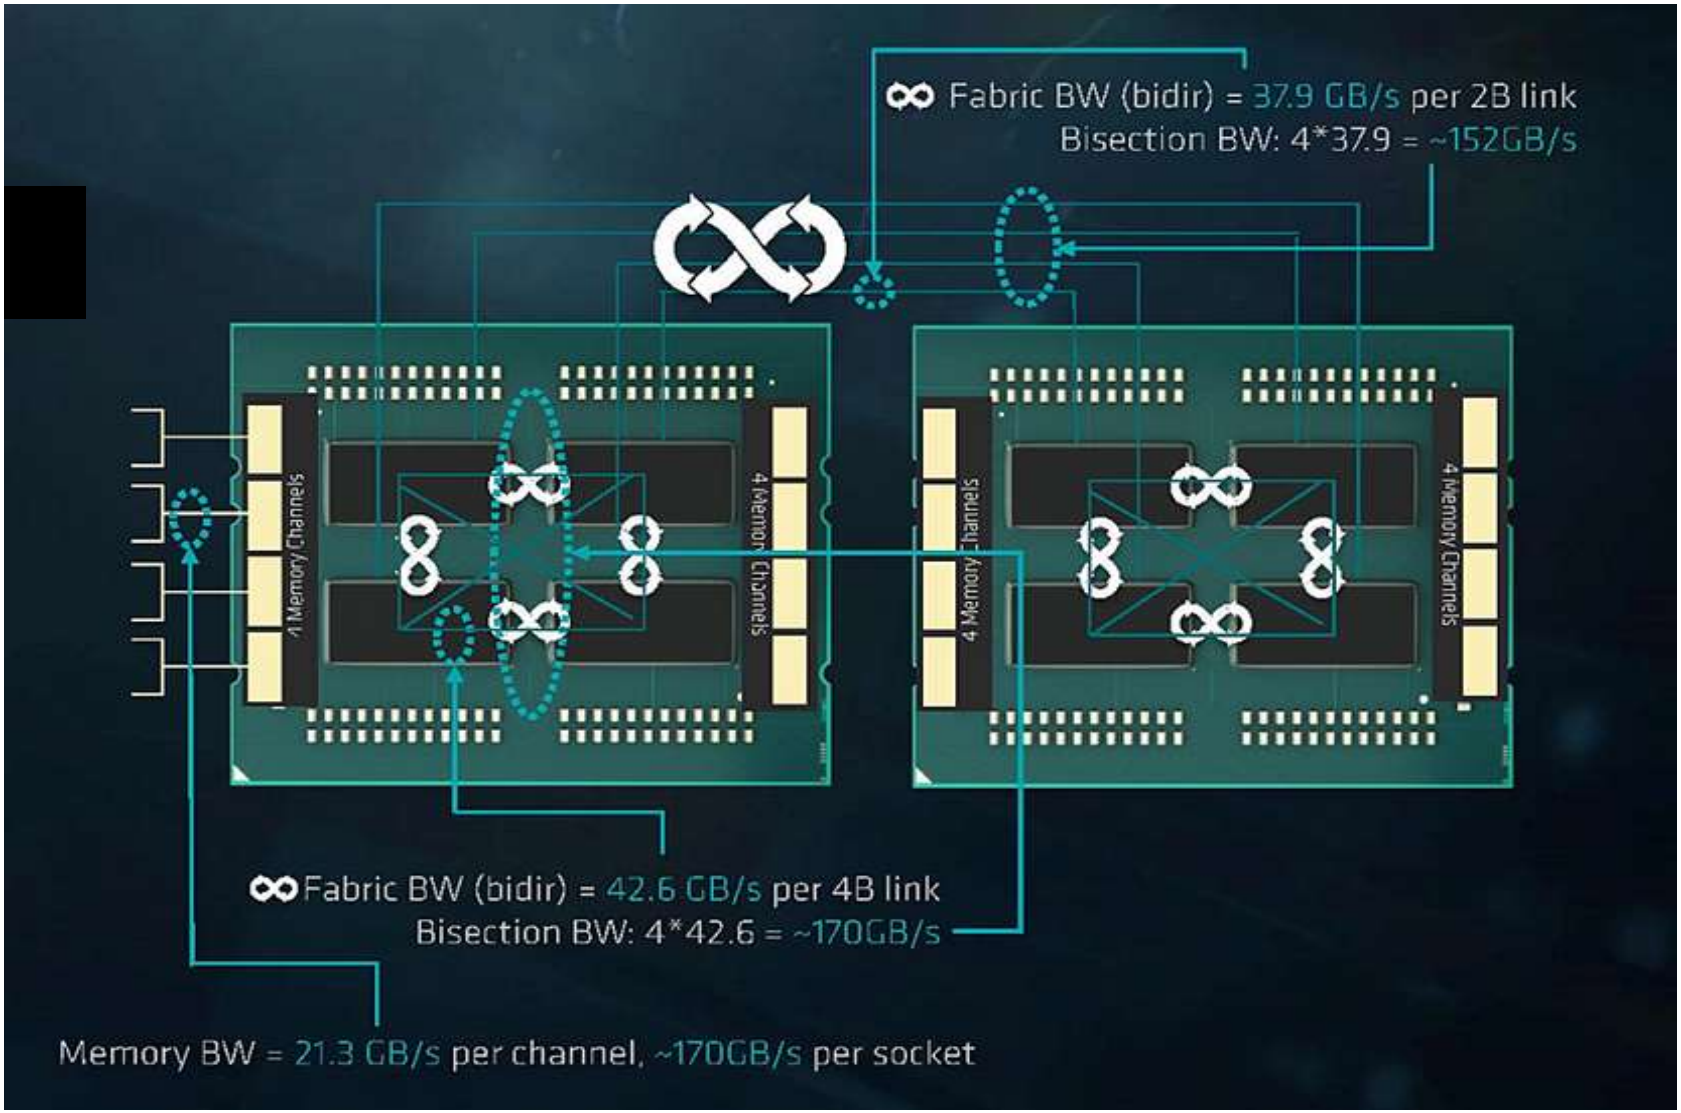

Revealed in AMD's Epyc Launch event presentation June 20 2017

I/O (PCIe 3.0) links provided by a 1S server [21]

- 8 x16 links available in 1 and 2 socket systems
  - Link bifurcation support; max of 8 PCIe® devices per x16
- 32GB/s bi-dir bandwidth per link, 256GB/s per socket
- Integrated SATA support

Architected for Massive I/O Systems
with Leading-edge Featureset

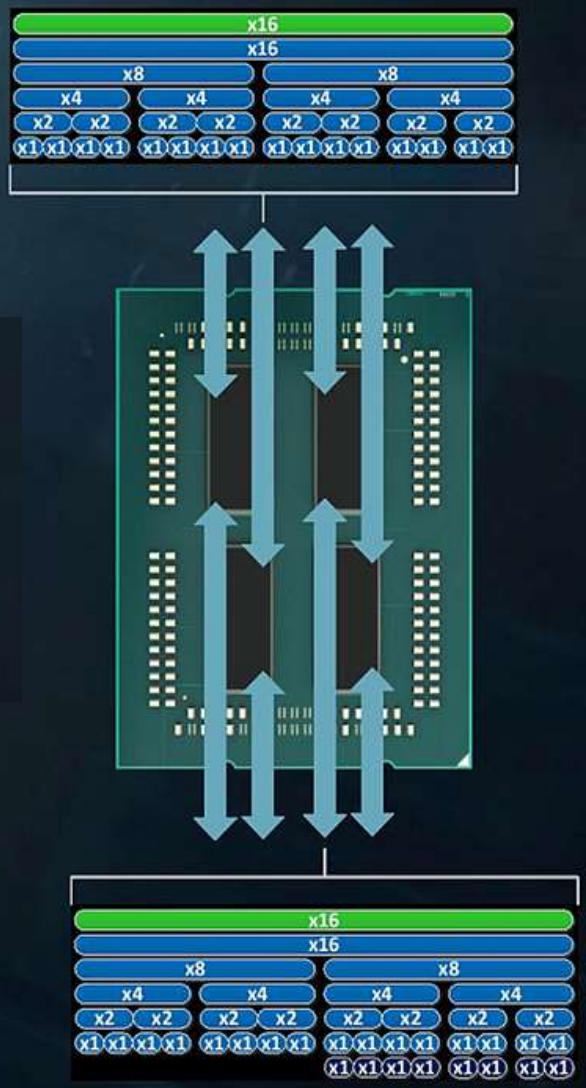

Latency and bandwith of IF links in EPYC processors [21]

EPYC (Naples) Thread Ping Connections		
	Latency	**Bandwidth**
Within A Core	26 ns	-
Core-to-Core, Same CCX	42 ns	-
Core-to-Core, Different CCX, Same Die	142 ns	-
Die-to-Die, Same Package	?	42.6 GB/s
Die-to-Die, Different Socket, One Hop	?	37.9 GB/s
Die-to-Die, Different Socket, Two Hops	?	37.9 GB/s
Core to DRAM, Same Die	90 nm	42.6 GB/s
Core to DRAM, Different Die	145 ns	42.6 GB/s
Core to DRAM, Different Socket, One Hop	200 ns	37.9 GB/s
Core to DRAM, Different Socket, Two Hops	?	37.9 GB/s

Bandwidth: bidirectional

Latencies and bandwidth in Intel's Skylake and Kaby Lake processors [22]

Processor	Line	Core count	Intra-Core Latency	Core-To-Core Latency	Core-To-Core Average Latency	Average Transfer Bandwidth
Core i9-7900X	Skylake-X	10	14.5 - 16ns	69.3 - 82.3ns	75.56ns	83.21 GB/s
Core i9-7900X @ 3200 MT/s	Skylake-X	10	16 - 16.1ns	76.8 - 91.3ns	83.93ns	87.31 GB/s
Core i7-6950X	Broadwell-E	10	13.5 - 15.4ns	54.5 - 70.3ns	64.64ns	65.67 GB/s
Core i7-7700K	Kaby Lake-S	4	14.7 - 14.9ns	36.8 - 45.1ns	42.63ns	35.84 GB/s
Core i7-6700K	Skylake-S	4	16 - 16.4ns	41.7 - 51.4ns	46.71ns	32.38 GB/s

Intra-core latency quantifies latency between threads that are resident on the same physical core, while the core-to-core figures reflect thread-to-thread latency between two physical cores.

Core i9-7900K is most comparable to the 10-core Core i7-6950X, but the four-core models are also included as a reference point.

6xxx models belong to the Skylake line, 7xxx models to the Kaby Lake line.

DDR + IFOP (IF On-Package) routing [88]

- Vertical and Horizontal IFOP: 2 layers each
- Diagonal IFOP: 1 layer each
- DDR channel: 1 layer each

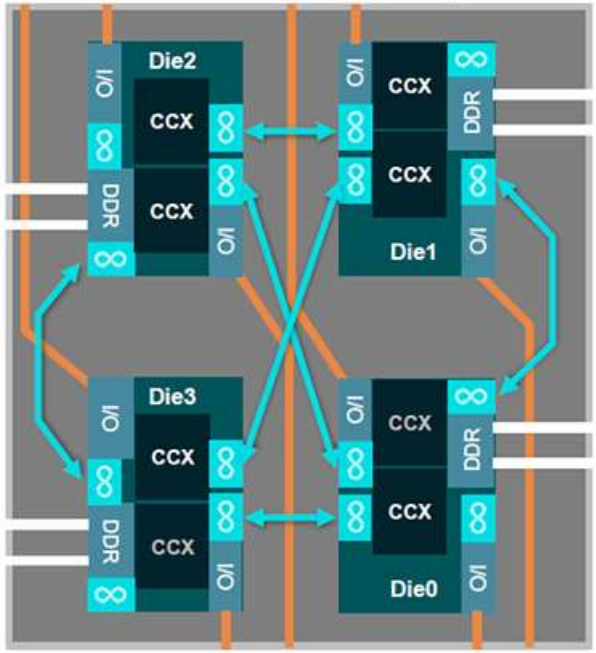

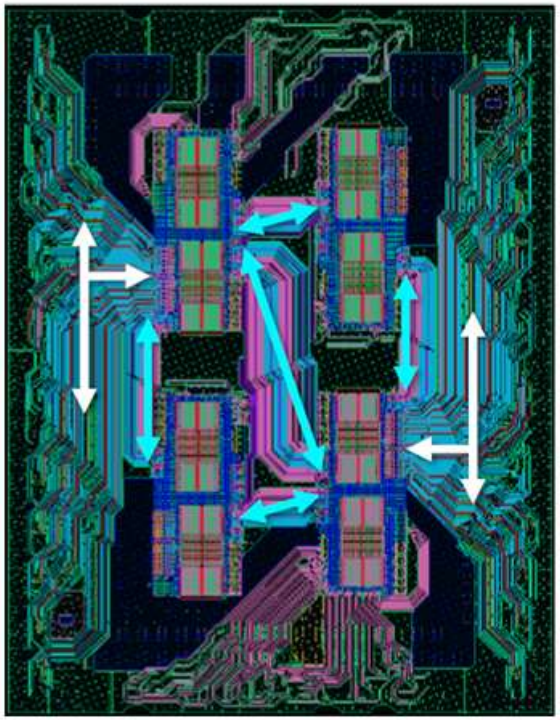

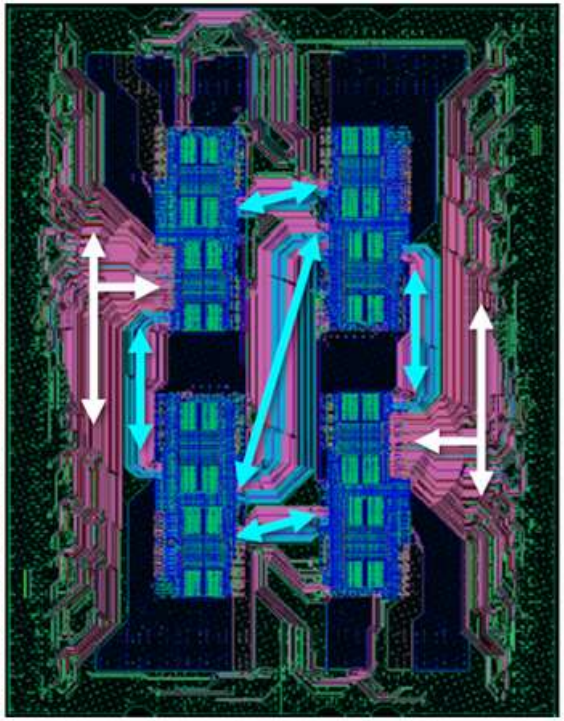

DDR + IFIS (IF Inter-Socket) package routing [88]

- DDR channel: 1 layer each
- IFIS links: 2 layers each

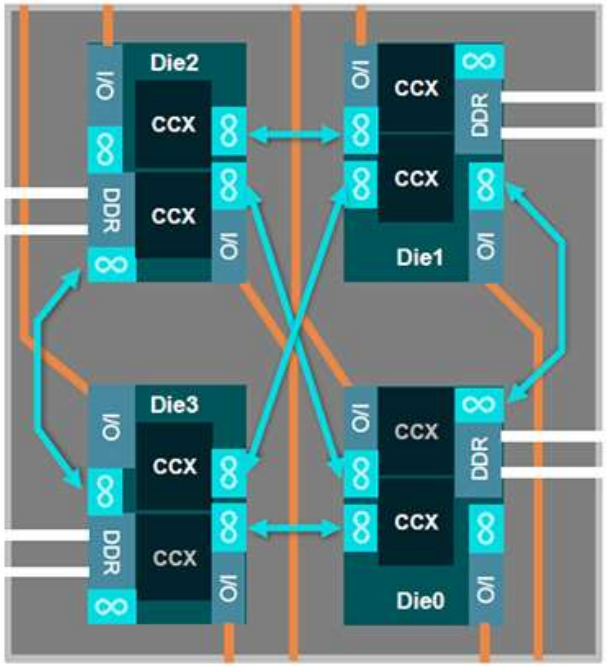

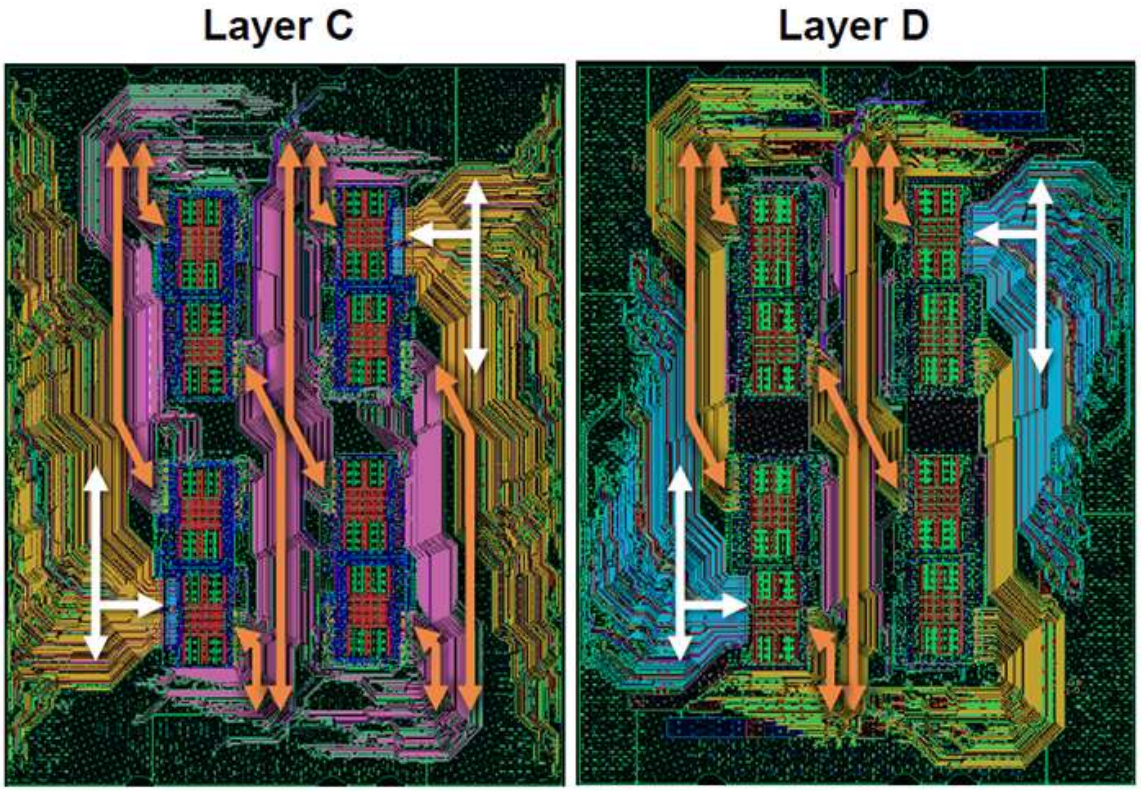

6. The Ryzen desktop lines

- 6.1 Ryzen desktop lines without integrated GPU
- 6.2 Ryzen desktop lines with integrated GPU

Only Section 6.1 will be discussed

The Ryzen desktop lines

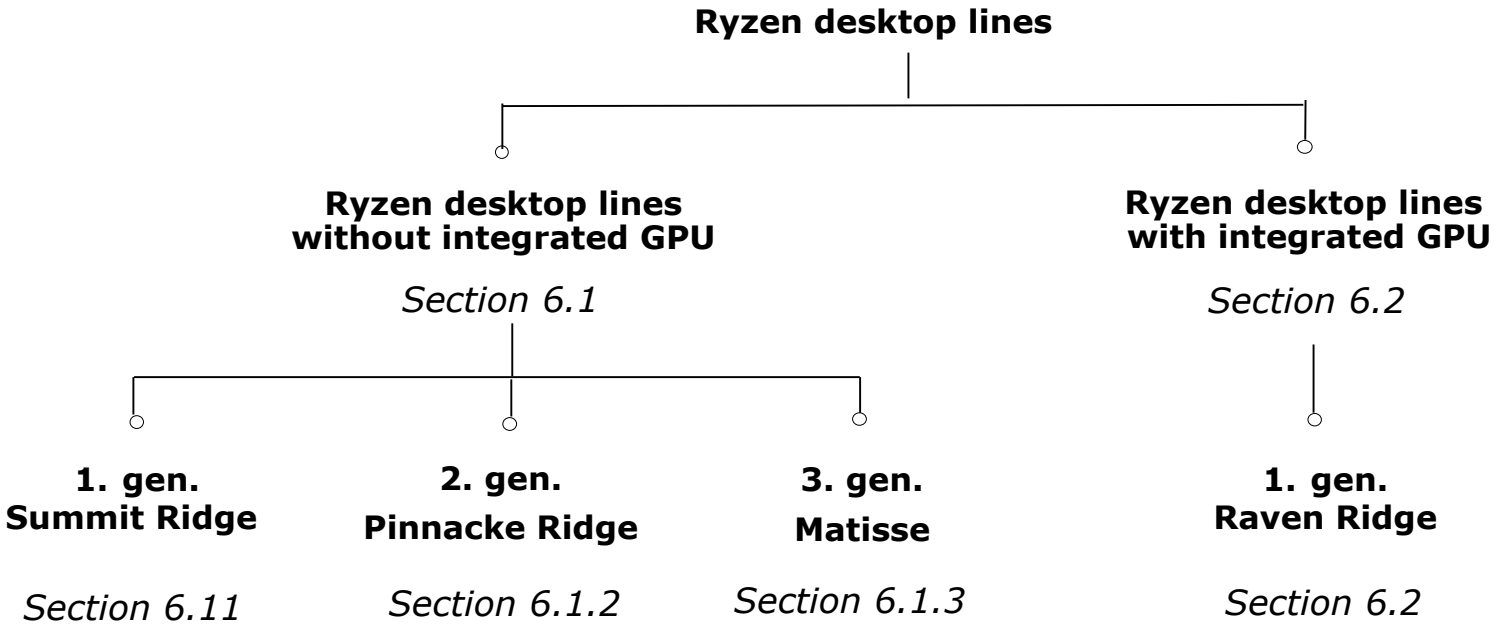

6.1 Ryzen desktop lines without integrated GPU

- 6.1.1 The 1. gen. Ryzen desktop line without GPU (Summit Ridge)
- 6.1.2 The 2. gen. Ryzen desktop line without GPU (Pinnacle Ridge)
- 6.1.3 The 3. gen. Ryzen desktop line without GPU (Matisse)

Only Section 6.1 will be discussed

6.1 Ryzen desktop lines without integrated GPU [64]

6.1.1 The 1. gen. Ryzen desktop line without GPU (Summit Ridge)

6.1.1 The 1. gen. Ryzen desktop line without GPU (Summit Ridge)

- The 1. gen. Ryzen desktop line without GPU is the Summit Ridge line.
- Launched in 3/2017 (Ryzen 7)
  4/2017 (Ryzen 5)
  7/2017 (Ryzen 3)
- AMD introduced the Ryzen processors at their AMD Tech Day in March 2 2017 by saying proudly:

"We are AMD, creators of Athlon, Radeon and other famous microprocessors. We also power the Xbox One and PS4. Today we want to talk RYZEN, our new high-speed CPU five years in the making" [60].

- Ryzen processors include a single Zeppelin die, called also the Ryzen die, which is built up basically of two CCX complexes interconnected by the IF (Infinity Fabric), as seen in the next slides.
- Models of the Ryzen desktop line are designated as Ryzen 7/5/3 1xxx.
- They are unlocked.
- They have AM4 socket (1331).

*

The increased maximum core count of the Ryzen desktops

Until introducing the Ryzen desktop line desktops provided up to 4 cores, as the Figure below indicates.

Figure: Increasing the maximum core count along with the introduction of the Ryzen desktop line [66]

- The Ryzen desktop line provides up to 8 cores.
- Subsequently, also Intel raised the max. core count to 6 in their 7. gen. Kaby Lake Refresh desktop line in Q3 2017 to cope with AMD's Ryzen line for MT workloads.

Positioning AMD's Zen-based Ryzen desktop lines

AMD's Zen-based processor lines

Ryzen Mobile
(Mobil)

Single CCX +
Vega GPU

Ryzen
(DT)

Zeppelin chip
(2x CCX, no GPU)

ThreadRipper
(HED)

2 Zeppelin chips
interconnected by the IF,
implemented as an MCM

Epyc
(1S/2S server)

4 Zeppelin chips
interconnected by the IF,
implemented as an MCM

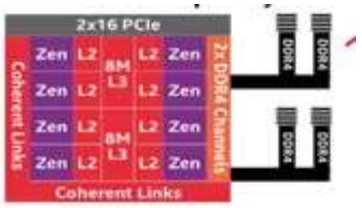

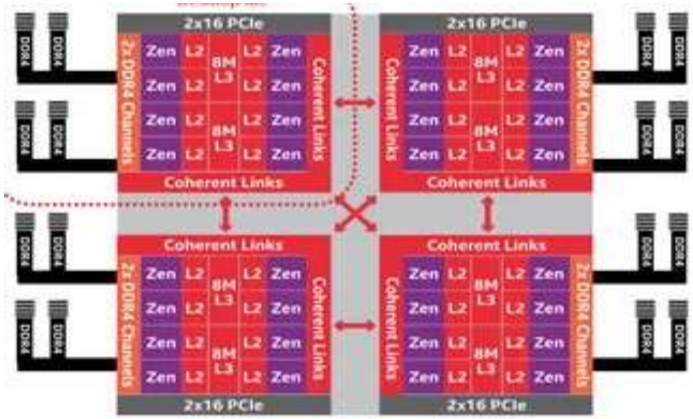

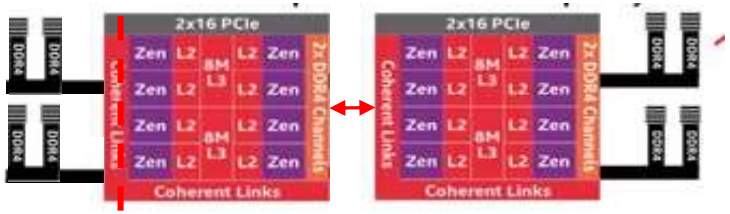

Drawings based on [11]

Basic layout of AMD's Zen-based Zeppelin chip [20], [10]

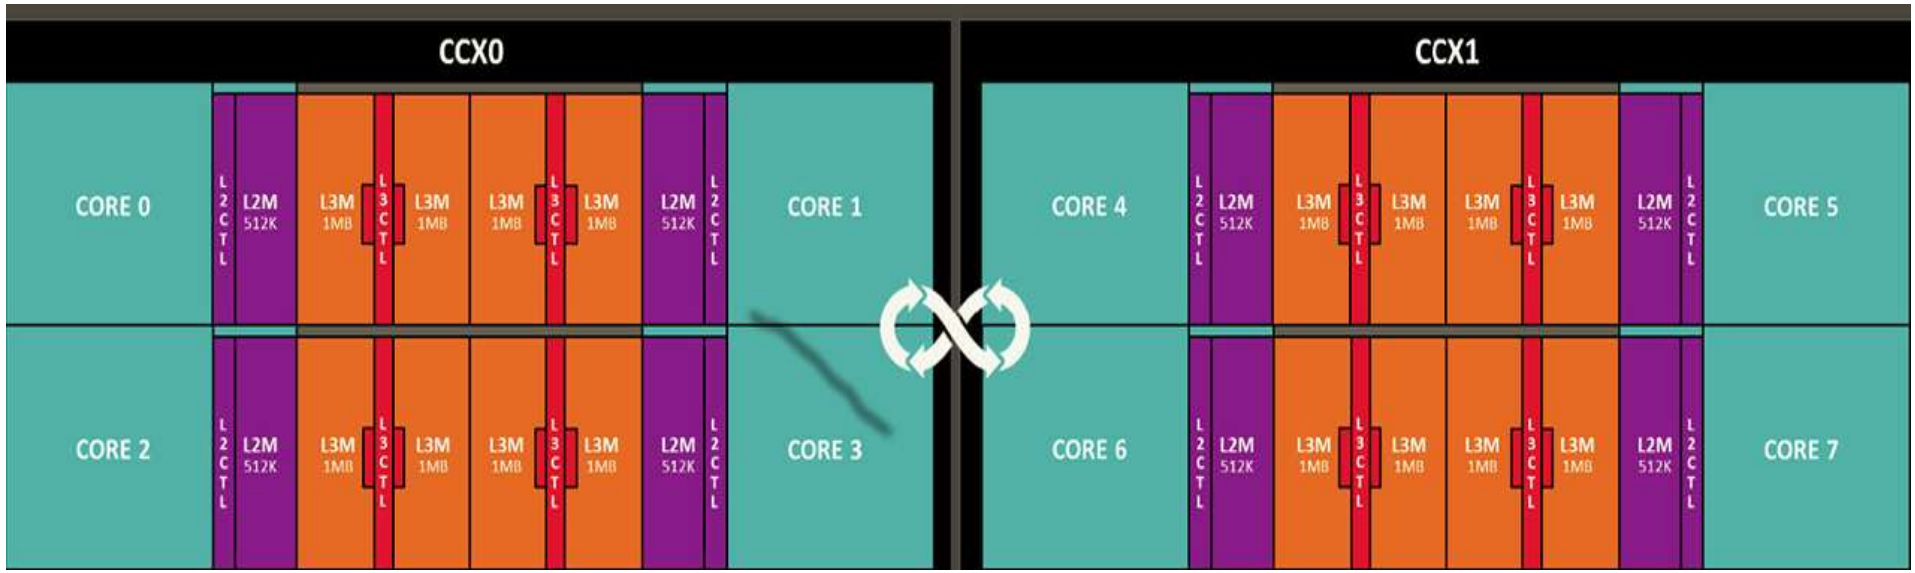

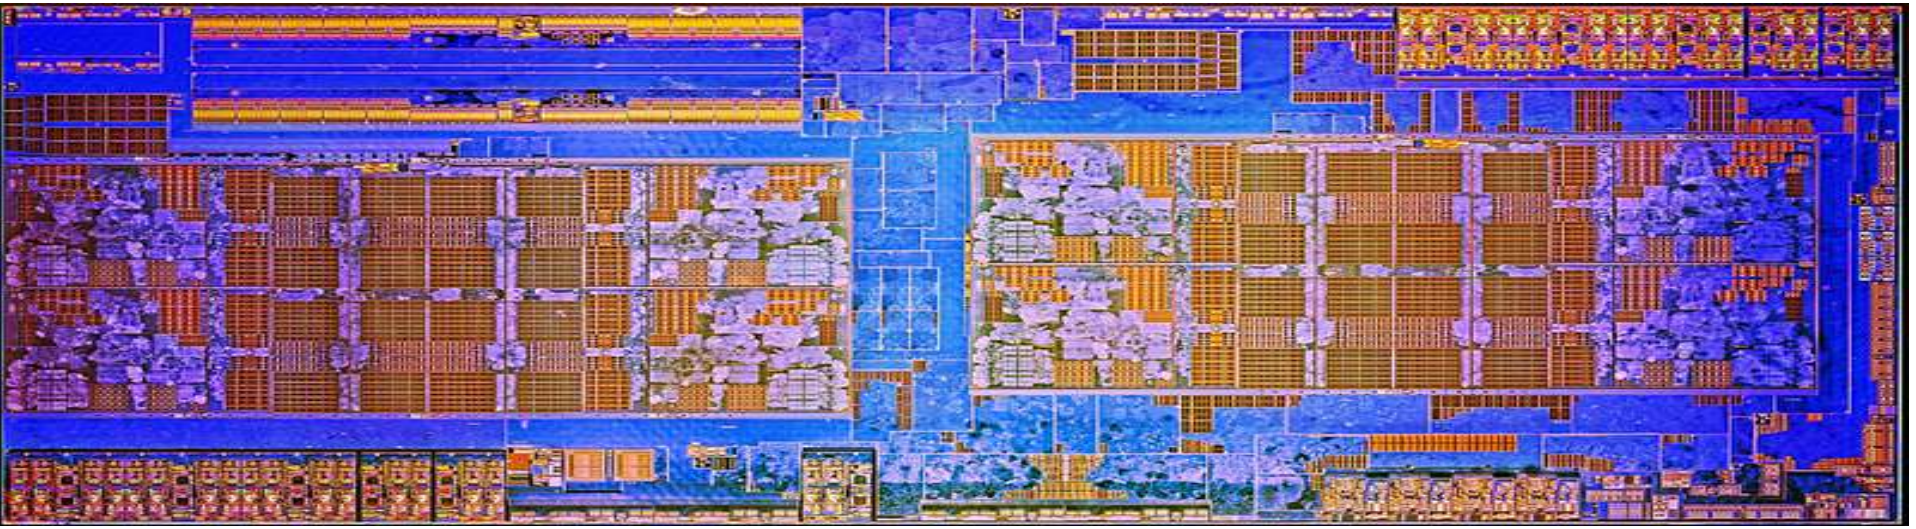

4.8 billion transistors, 192mm^2 die

Block diagram of AMD's Zeppelin chip [17]

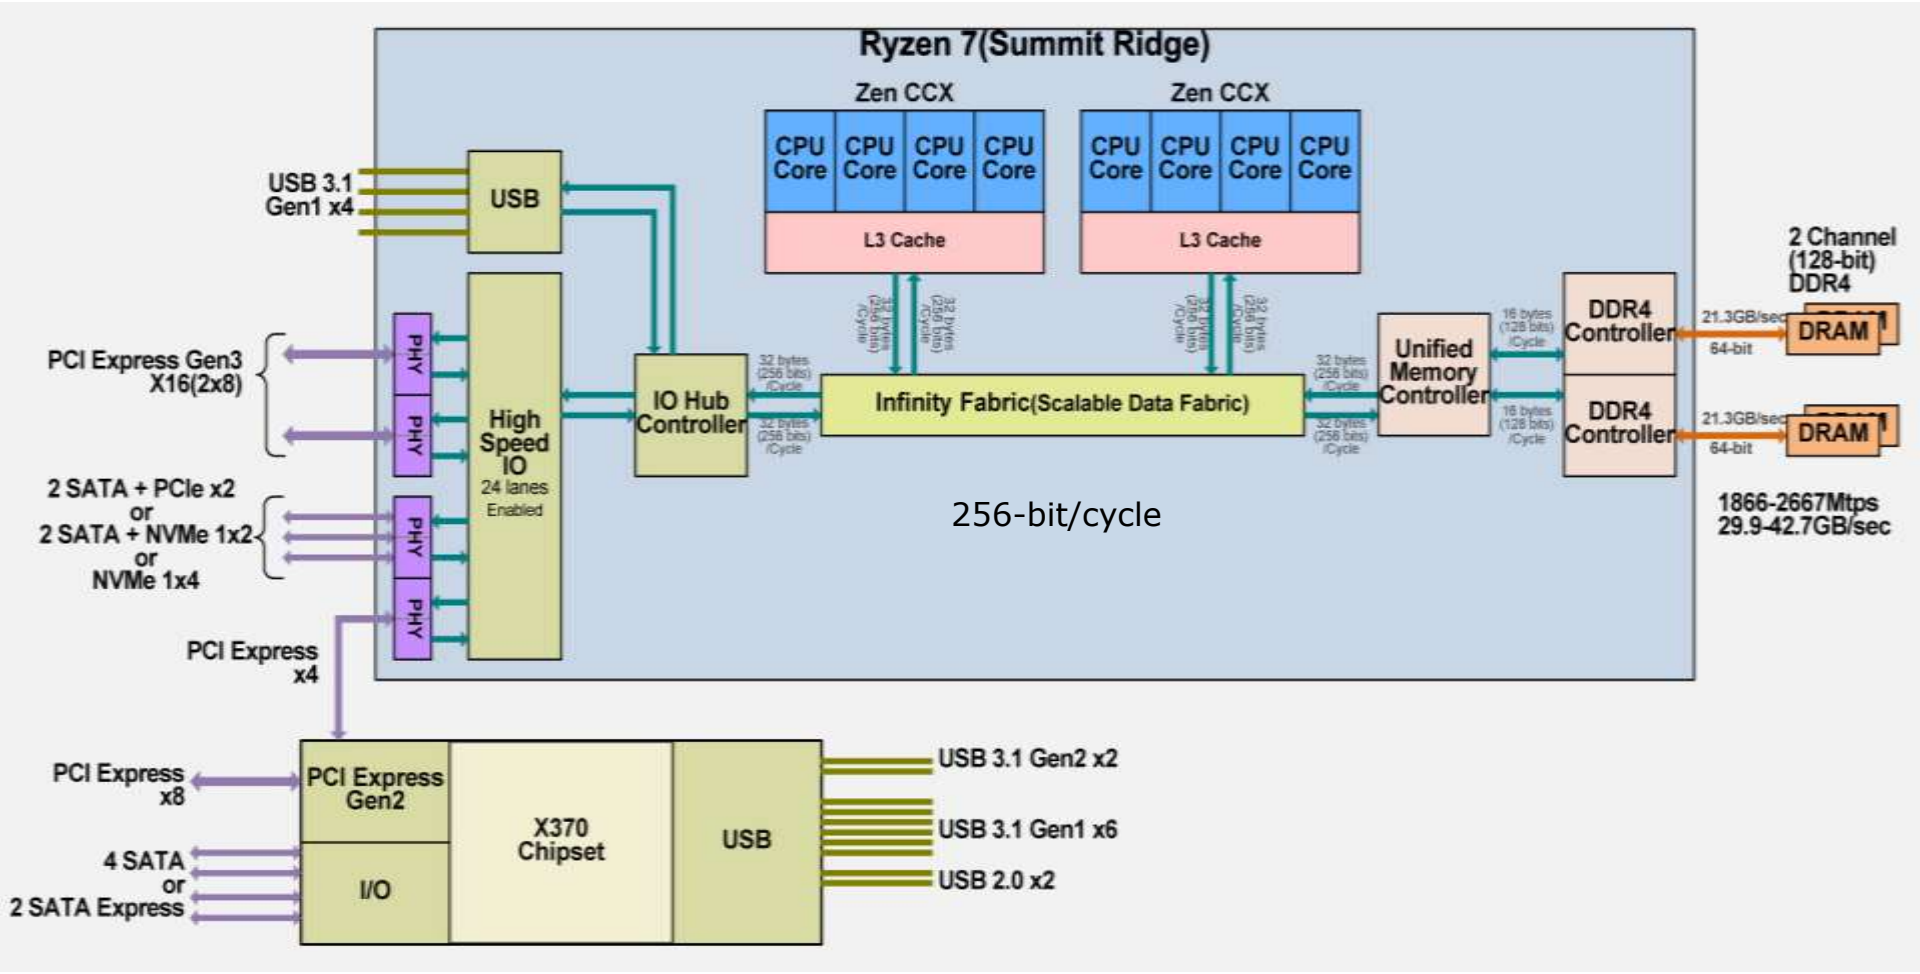

Main features of AMD's Ryzen desktop line introduced in 2017

AMD's Zen-based processor lines introduced in 2017

	Ryzen Mobile APU (Raven Ridge)	**Ryzen DT** (Summit Ridge)	**ThreadRipper** (Whitehaven)	**Epyc** (Naples)
Market segment	Mobile	Desktop platform	HED	1S/2S server
µarch./Technology	Zen 14 nm	Zen, 14 nm	Zen, 14 nm	Zen, 14 nm
Launched models	Ryzen 7 2700U Ryzen 5 2500U (10/2017)	Ryzen 7 (3/2017) Ryzen 5 (4/2017) Ryzen 3 (7/2017)	1950X/1920X/1900X (8/2017)	Series 7000 (6/2017)
Layout	CCX + Vega 8/10	Zeppelin die with 2x CCX	MCM (2x Zeppelin die)	MCM (4x Zeppelin die)
Integrated GPU	Yes	No	No	No
Core count	4	4/6/8	8/12/16	8/16/24/32
SMT	SMT	SMT (except Ryzen 3)	SMT	SMT
Mem. channels/rate	2xDDR4-2400	2xDDR4-2666	4xDDR4-2666	8xDDR4-2666
PCIe 3.0 lanes	??	16xPCIe 3.0	60xPCIe 3.0	128 for 1S servers 64 for 2S servers
TDP	15 W	65/95 W	180 W	120/170/180 W
Socket	AM4 (1331)	AM4 (1331)	TR4 (SP3r2) (4094)	SP3 (4094)
Chipset	SoC	300-series	X399	No chipset, SOC

Main features of the Ryzen and ThreadRipper lines [23]

AMD Ryzen SKUs								
	Cores/ Threads	Base/ Turbo	XFR	L3	DRAM 1DPC	PCIe	TDP	SRP
TR 1950X	**16/32**	**3.4/4.0**	**+200**	**32 MB**	**4x2666**	**60**	**180W**	**$999**
TR 1920X	**12/24**	**3.5/4.0**	**+200**	**32 MB**	**4x2666**	**60**	**180W**	**$799**
TR 1920**	12/24	3.2/3.8	?	32 MB	4-Ch?	60	140W	?
TR 1900X	8/16	3.8/4.0	+200	16 MB*	4x2666*	60	180W*	$549
Ryzen 7 1800X	8/16	3.6/4.0	+100	16 MB	2x2666	16	95 W	$499
Ryzen 7 1700X	8/16	3.4/3.8	+100	16 MB	2x2666	16	95 W	$399
Ryzen 7 1700	8/16	3.0/3.7	+50	16 MB	2x2666	16	65 W	$329
Ryzen 5 1600X	6/12	3.6/4.0	+100	16 MB	2x2666	16	95 W	$249
Ryzen 5 1600	6/12	3.2/3.6	+100	16 MB	2x2666	16	65 W	$219
Ryzen 5 1500X	4/8	3.5/3.7	+200	16 MB	2x2666	16	65 W	$189
Ryzen 5 1400	4/8	3.2/3.4	+50	8 MB	2x2666	16	65 W	$169
Ryzen 3 1300X	4/4	3.5/3.7	+200	8 MB	2x2666	16	65 W	$129
Ryzen 3 1200	4/4	3.1/3.4	+50	8 MB	2x2666	16	65 W	$109

Product specifications subject to change.

Infinity Fabric thread latencies in AMD's Ryzen line [24]

Processor	Intra-Core Latency	Intra-CCX Core-to-Core Latency (same die)	Cross-CCX Core-to-Core Latency (different dies)	Cross-CCX Average Latency	Die-to-Die Latency	Die-To-Die Average Latency	Average Transfer Bandwidth
Ryzen 7 1800X	14.8ns	40.5 - 82.8ns	120.9 - 126.2ns	122.96ns	**X**	**X**	48.1 GB/s
Ryzen 5 1600X	14.7 - 14.8ns	40.6 - 82.8ns	121.5 - 128.2ns	123.48ns	**X**	**X**	43.88 GB/s

Interpretation of the Infinity Fabric latencies in the Table above [24]

- Intra-core latency reflects the communication between two logical threads resident on the same physical core, it is unaffected by memory speed.
- Intra-CCX latencies quantify latency between threads that are on the same CCX but not resident on the same core, it is also largely unaffected by memory speed.
- Cross-CCX latency denotes latency between threads located on two separate CCX building blocks.
- Die-to-Die latency relates to latency between threads that are on different dies.
- Ryzen line processors include only a single die so no Die-to-Die latencies can occur.

Ryzen based system architecture [25]

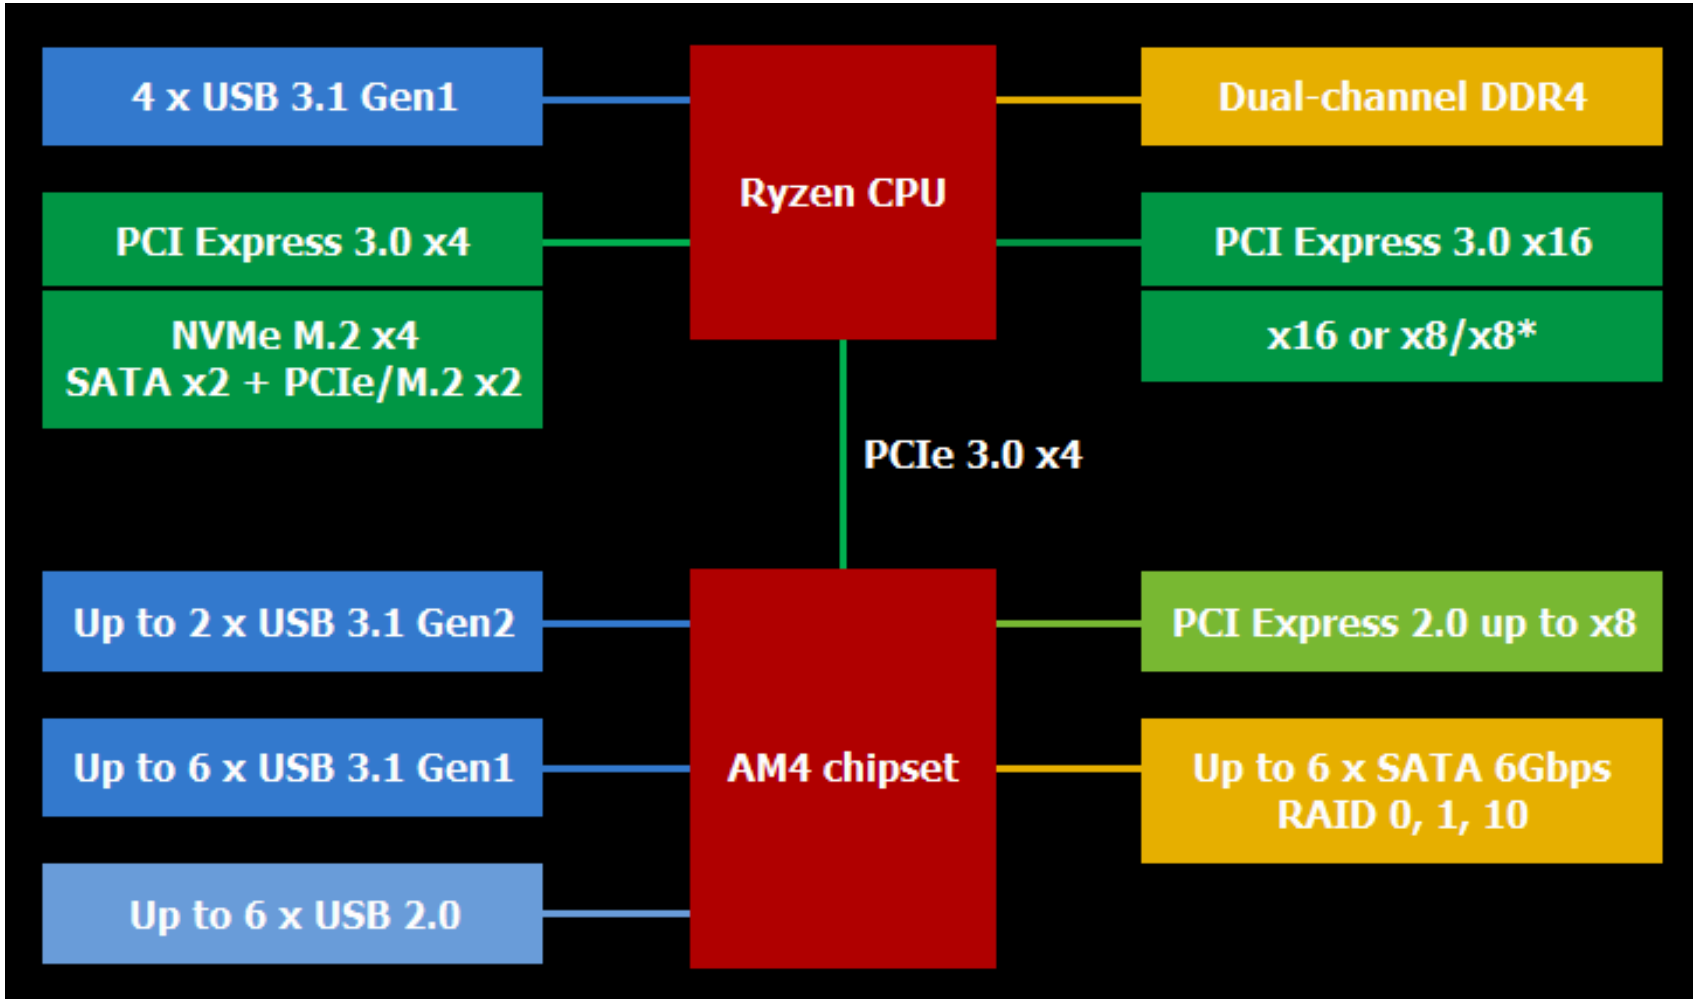

Single-threaded performance per Dollar on CPU benchmarks [10]

(Only Ryzen 7 models are shown)

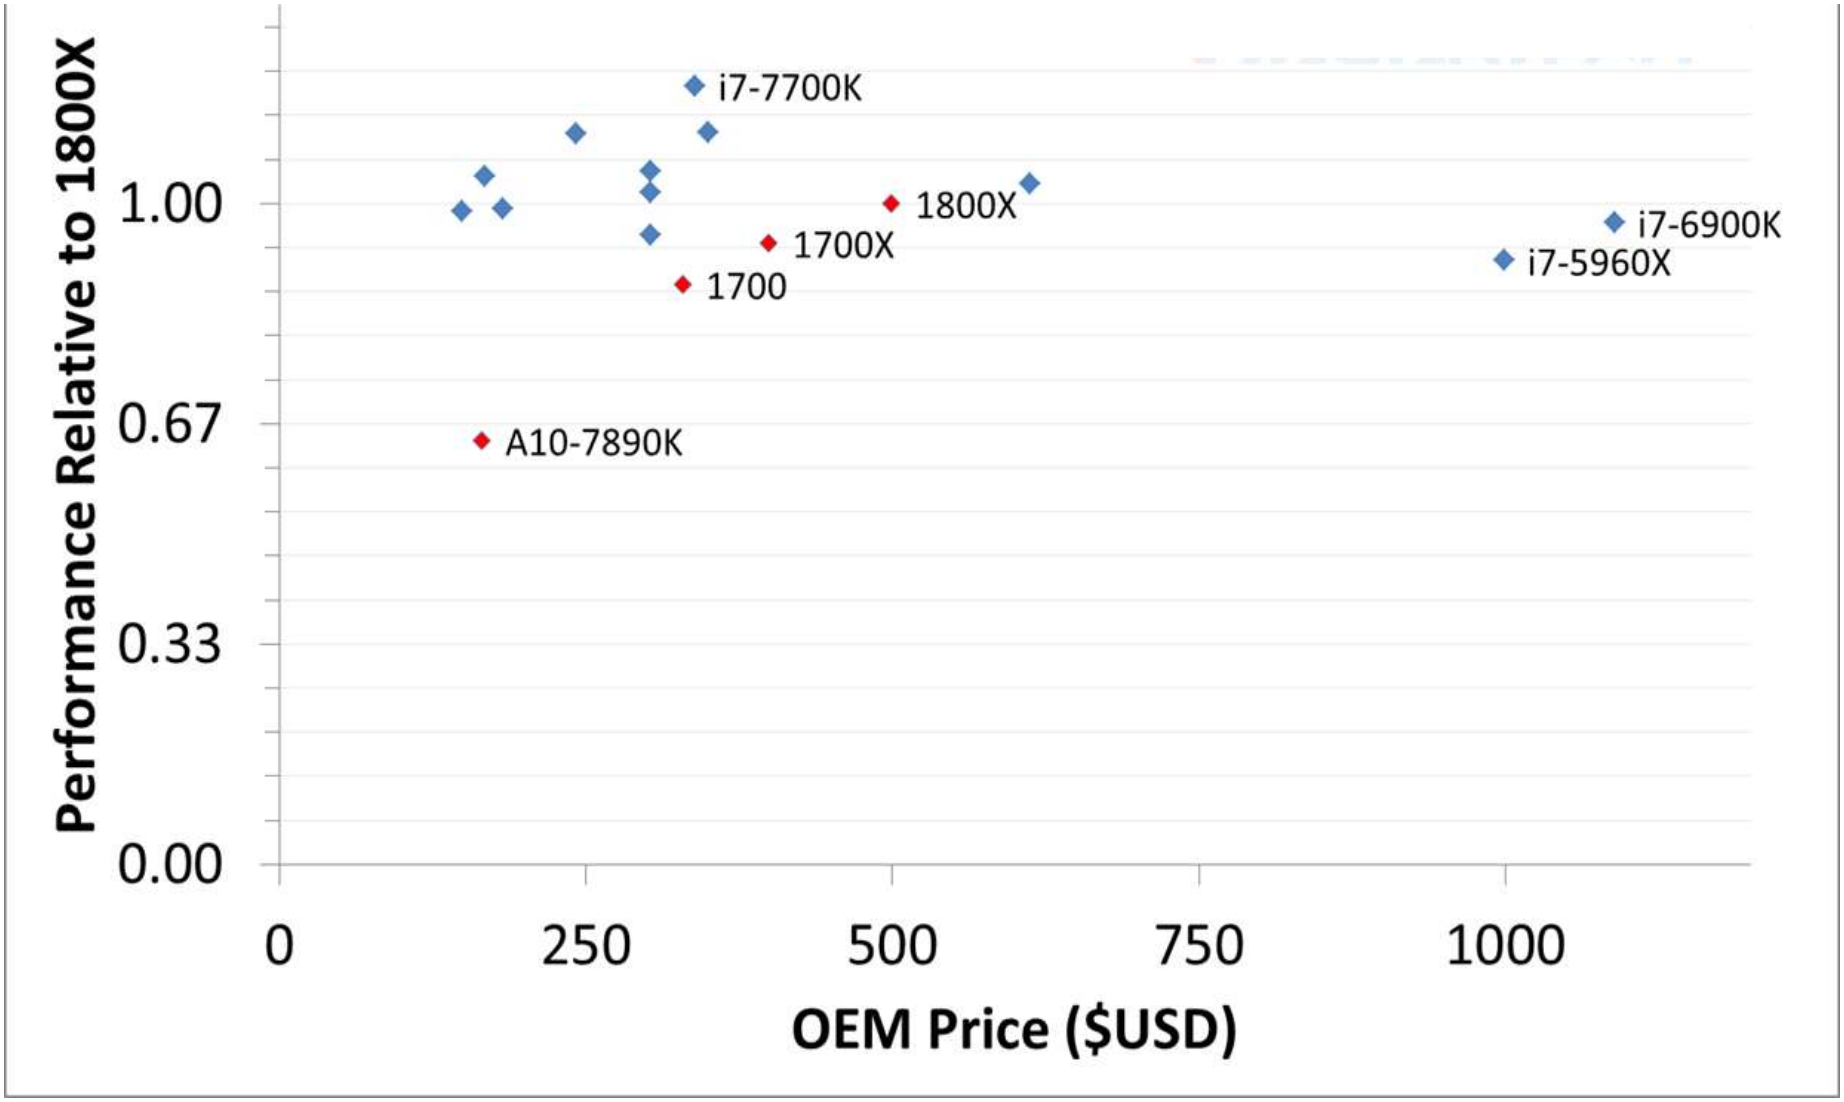

Multi-threaded performance per Dollar on CPU benchmarks [10]

(Only Ryzen 7 models are shown)

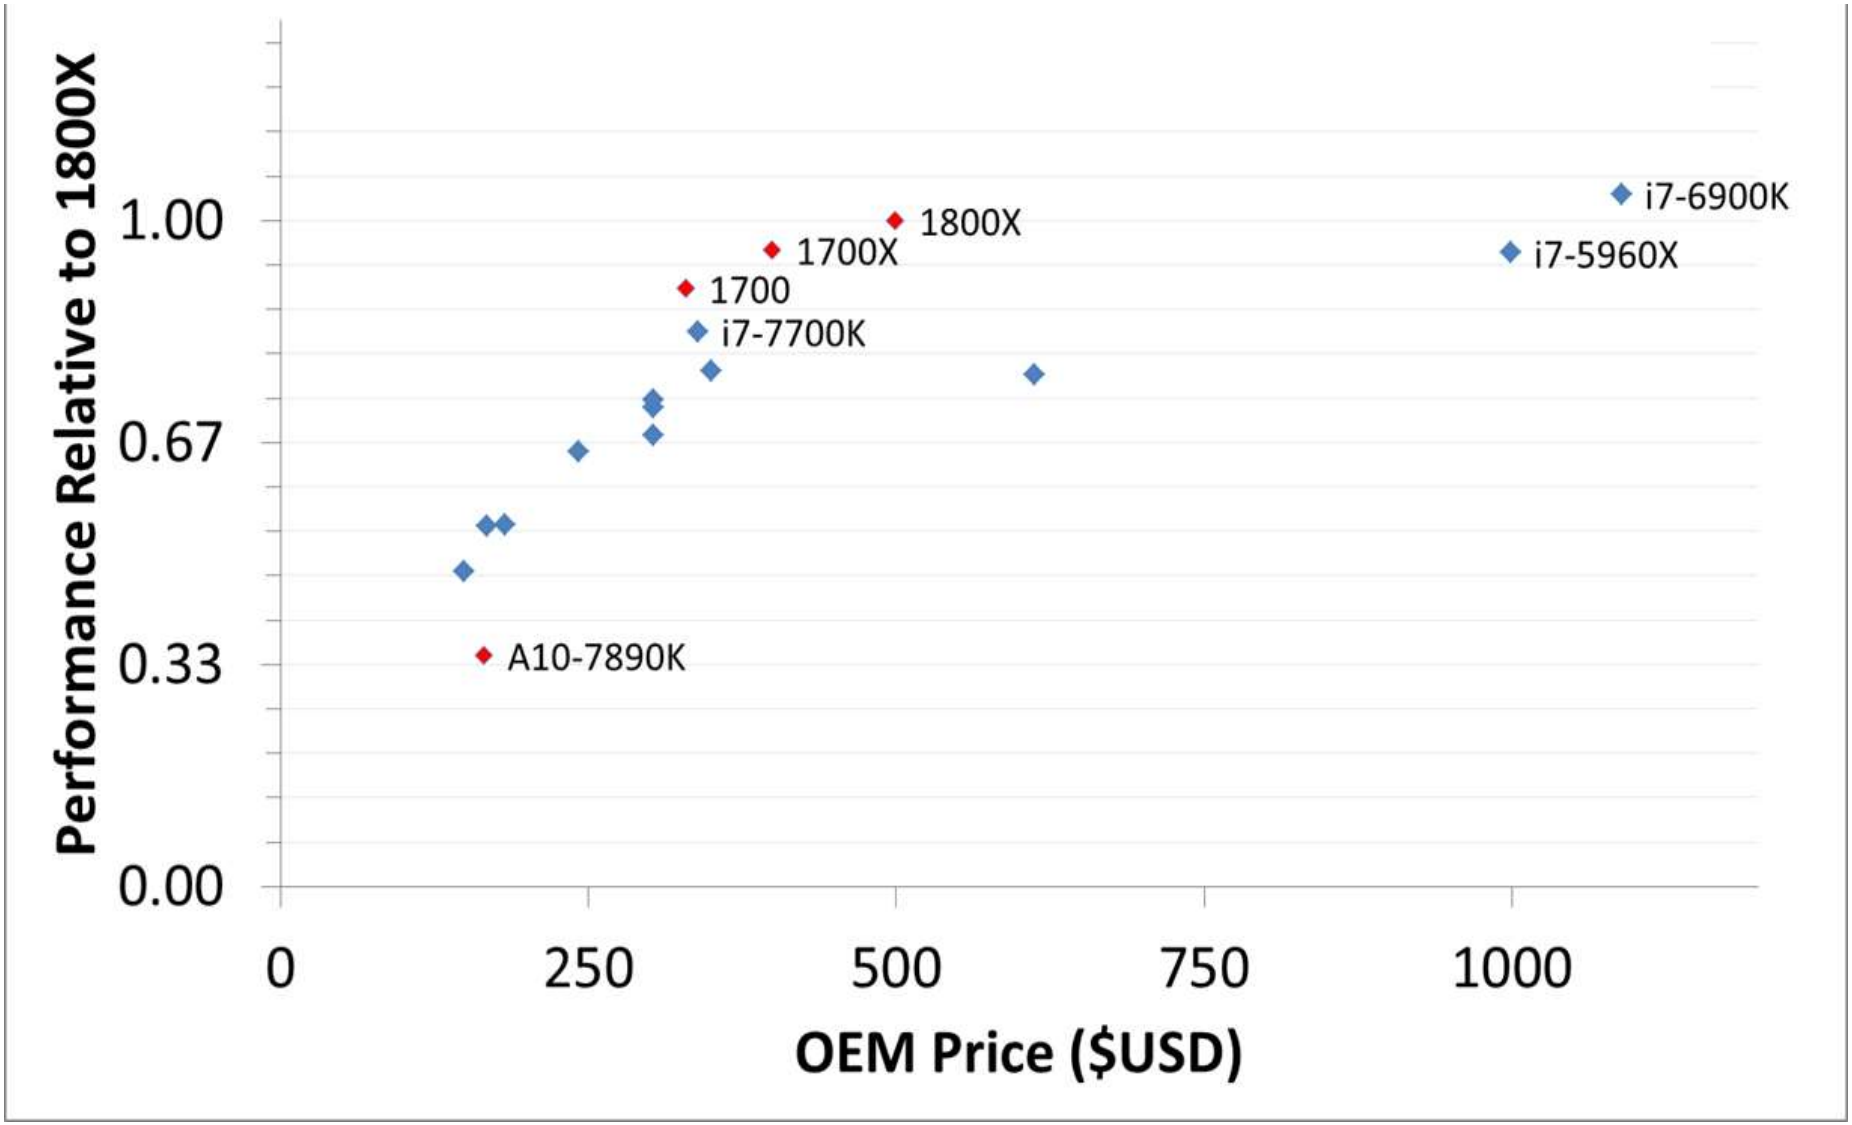

6.1.2 The 2. gen. Ryzen desktop line without GPU (Pinnacle Ridge)

6.1.2 The 2. gen. Ryzen desktop line without GPU (Pinnacle Ridge)

- Code-named as the Pinnacle Ridge line.
- Launched in 4/2018 as the Ryzen 7 2700X/2700 and
  Ryzen 5 2600X/2600
  models.
- The line is manufactured on the 12 nm technology (by GlobalFoundries).
- These processors are unlocked.
- They have AM4 socket.
- The line is accompanied by new motherboards (400-series) but they remained compatible with the previous 300-series motherboards.
- The processors come bundled with AMD's Wraith Spire cooler.

*

Main features of the 2. gen. Ryzen desktop line

AMD's Zen/Zen+ based processor lines introduced in 2018

	Ryzen Mobil APU (Raven Ridge)	**Ryzen DT APU** (Raven Ridge)	**2.G Ryzen DT** (Pinnacle Ridge)	**2.G ThreadRipper** (Pinnacle Ridge)
µarch./tech.	Zen/14 nm	Zen/14 nm	Zen+/12 nm	Zen+/12 nm
Launched models	Ryzen 3 2300U Ryzen 3 2200U (1/2018)	Ryzen 5 2400G/GE Ryzen 3 2200G/GE (G: 2/2018) (GE: 42018)	Ryzen 7 2700X Ryzen 7 2700 Ryzen 5 2600X Ryzen 5 2600 (4/2018)	2990WX (8/2018) 2950X (8/2018) 2970WX (10/2018) 2920X (10/2018)
Layout	CCX + Vega 6/3	CCX + Vega 11/8	Zeppelin die (2x CCX)	MCM (up to 4 Zeppelin dies)
Integrated GPU	Yes	Yes	No	No
Core count	2/4	4	2/4	Up to 32
SMT (Multithreaded)	SMT only for Ryzen 3 2200U	SMT only for Ryzen 5	SMT	SMT
Mem. channels/ data rate	2xDDR4-2400	2xDDR4-2667 (G) 2xDDR4-2933(GE)	2xDDR4-2933	4xDDR4-2933
PCIe lanes	??	??	16x PCIe 3.0	60x PCIe 3.0
TDP	15 W	35/65 W	65/95/105 W	180/250 W
Socket	FP5 (na.)	AM4 (1331)	AM4 (1331)	TR4 (SP3r2) (4094)
Chipset	SoC	300-series	300/400-series	X399

Advantages of the 12 nm technology [65]

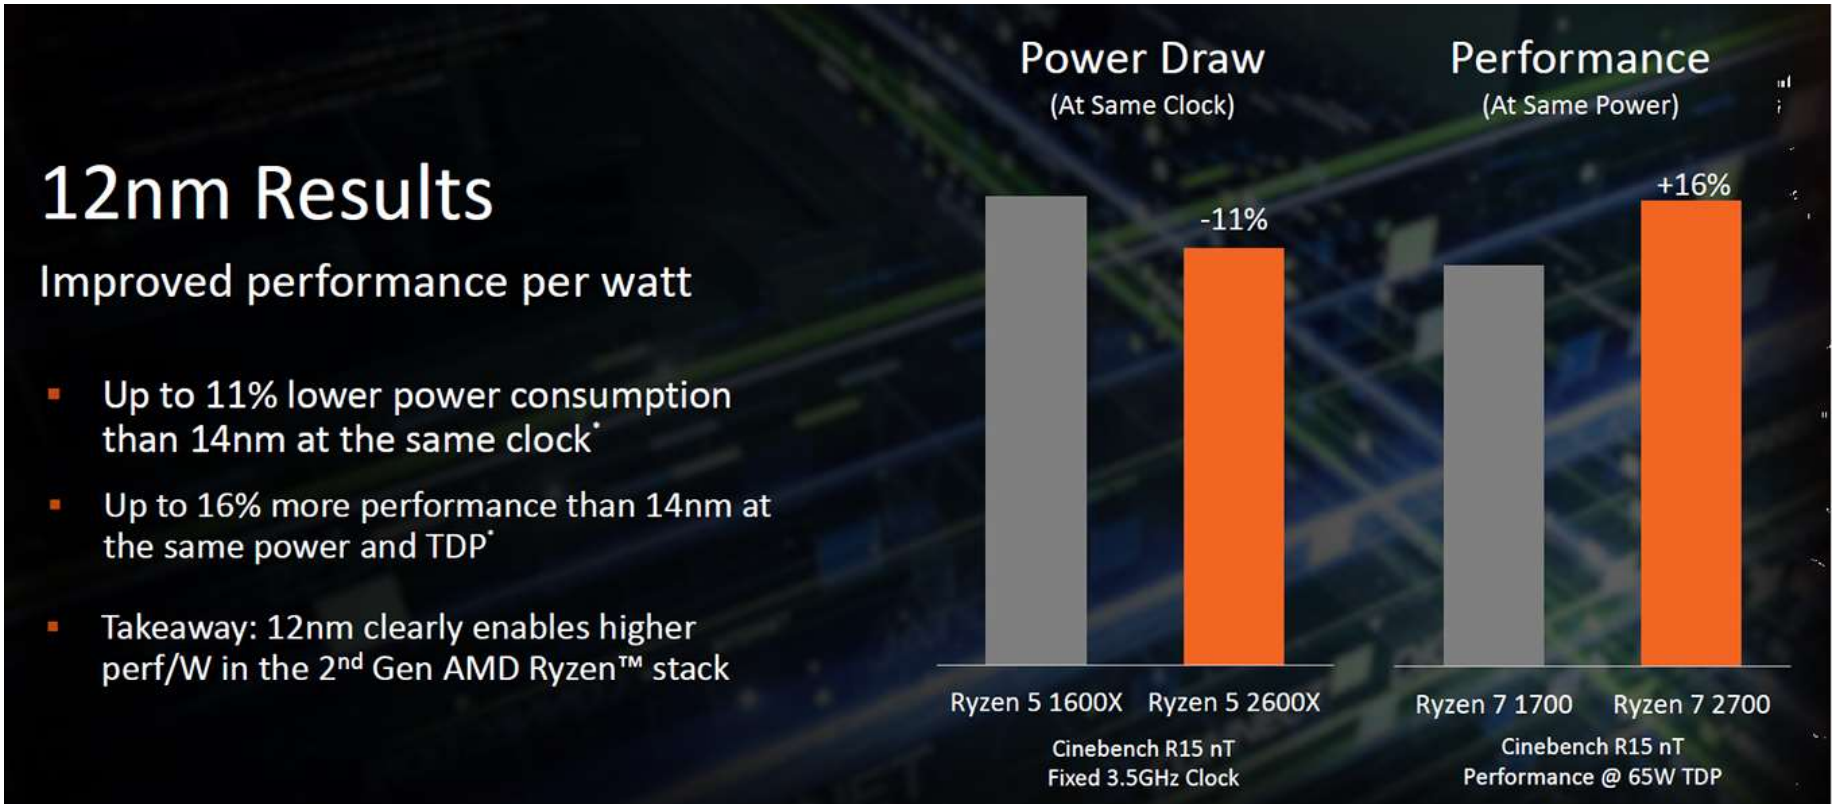

The Ryzen+ die [63]

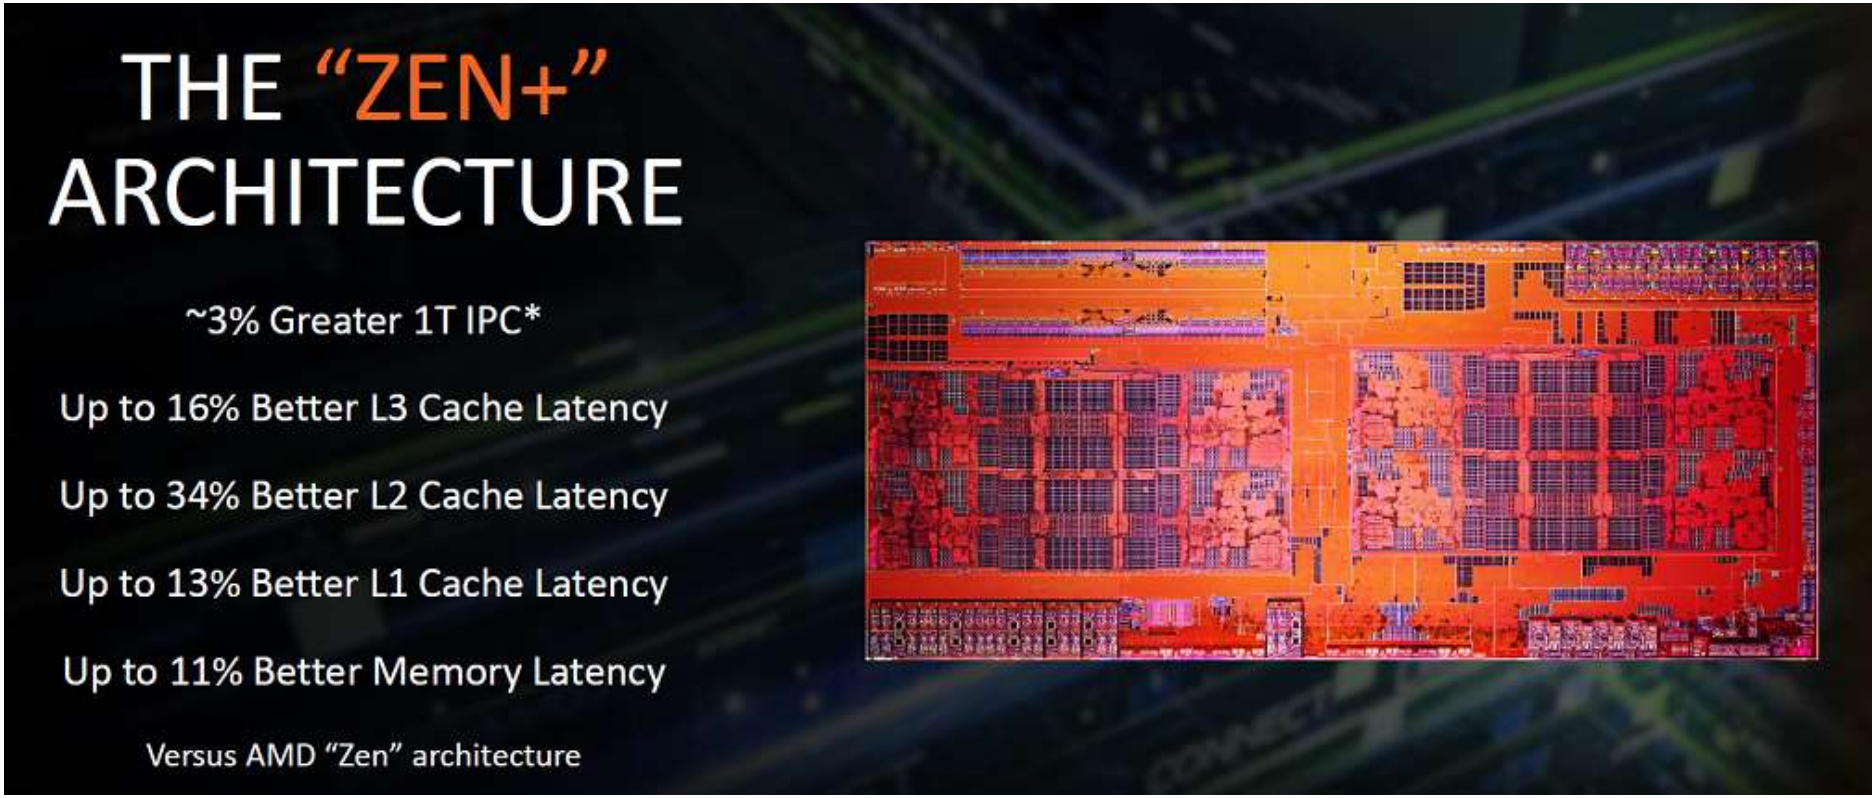

Die size: 213 mm^2, core count: 4.8 billion transistors.
This is the same data as for the previous Ryzen die.

Comparing the 14 nm and the 12 nm technologies [63]

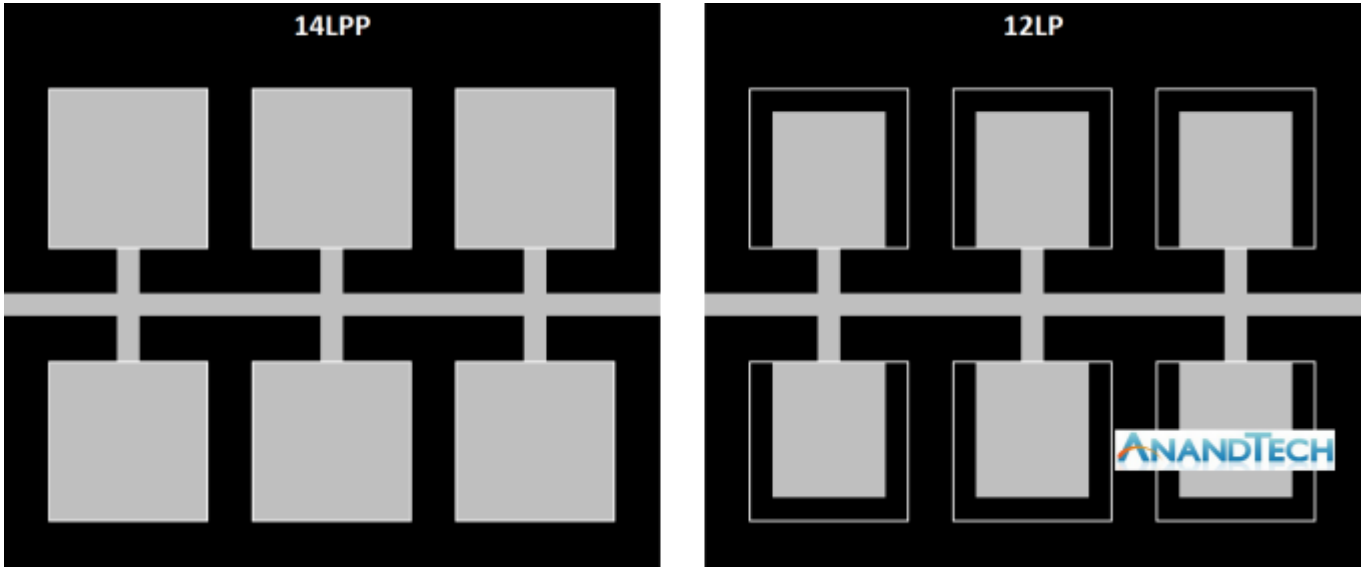

The 12 nm part has more dark silicon, that is beneficial for the thermal behavior.

Main features of the Ryzen+ (Pinnacle Ridge) desktop models [63]

AMD Ryzen 2000-Series CPUs (Pinnacle Ridge)				
	Ryzen 7 2700X	Ryzen 7 2700	Ryzen 5 2600X	Ryzen 5 2600
CPU Cores/Threads	8 / 16	8 / 16	6 / 12	6 / 12
Base CPU Frequency	3.7 GHz	3.2 GHz	3.6 GHz	3.4 GHz
Turbo CPU Frequency	4.3 GHz	4.1 GHz	4.2 GHz	3.9 GHz
TDP @ Base Frequency	105 W	65 W	95 W	65 W
L1 Cache	I: 64K. D: 32K	I: 64K. D: 32K	I: 64K. D: 32K	I: 64K. D: 32K
L2 Cache	512 KB/core	512 KB/core	512 KB/core	512 KB/core
L3 Cache	16 MB	16 MB	16 MB	16 MB
DRAM Support	DDR4-2933 Dual Channel	DDR4-2933 Dual Channel	DDR4-2933 Dual Channel	DDR4-2933 Dual Channel
PCIe Lanes (CPU)	16 Free + 4 NVMe	16 Free + 4 NVMe	16 Free + 4 NVMe	16 Free + 4 NVMe
Price	$329	$299	$229	$199
Bundled Cooler	AMD Prism RGB	AMD Spire RGB	AMD Spire	AMD Stealth

Enhancements of the Ryzen+ desktop line vs. the previous Ryzen line

a) Precision Boost 2 technology
b) Precision Boost Overdrive (PBO)
c) XFR2 (Extended Frequency Range) technology
d) StoreMI technology

6.1.2 The 2. gen. Ryzen desktop line without GPU (Pinnacle Ridge) (8)

a) Precision Boost 2 technology

Precision Boost 2 has been introduced in a number of series as follows:

- Ryzen mobile APU series (Zen (Raven Ridge)-based, 14 nm, 10/2017)
- Ryzen DT APU series (Zen+ (Pinnacle Ridge)-based, 12 nm 4/2018)
- 2. gen. Ryzen DT series (Zen+ (Pinnacle Ridge)-based, 12 nm, 4/2018) and
- 2. gen. ThreadRipper series, (Zen+ (Pinnacle Ridge)-based, 12 nm, 8/2018).

Boost behavior in the Precision Boost technology

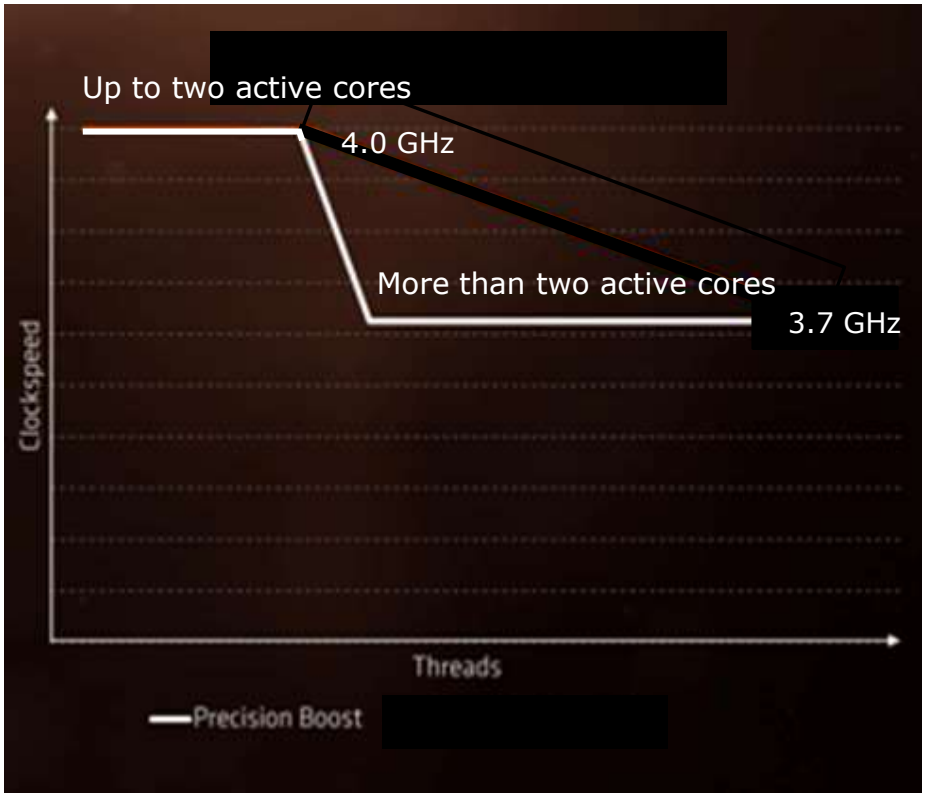

Figure: Boost behavior of the Ryzen DT and mobile lines [38]

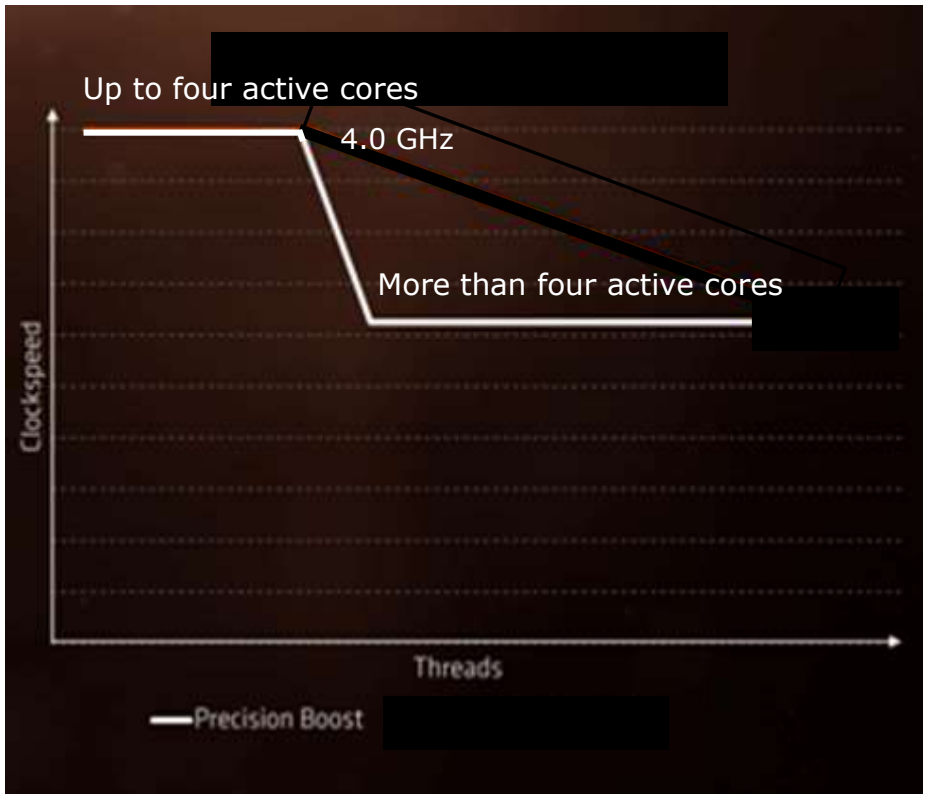

Figure: Boost behavior of the ThreadRipper line (based on [38])

Deficiency of the boost behavior of the implemented Precision Boost technology [76]

- The step-like reduction of the boost frequency while entering more active threads can unfavourable impact performance, especially in games.

  In this case several additional threads could often be spawned in addition to the one or two main threads and while the additional threads frequently do not contribute as much to the performance they can cause the boost frequency significantly drop.

- This could partly explain the moderate performance of AMD's Zen-based processors in certain gain titles and other benchmarks compared to Intel.
- In order to mend this deficiency AMD implemented a gradual decrease of the boost frequency along with raising the number of the active threads in their Precision Boost 2. technology, as discussed next.

Concept of the Precision Boost 2 technology -1 [76]

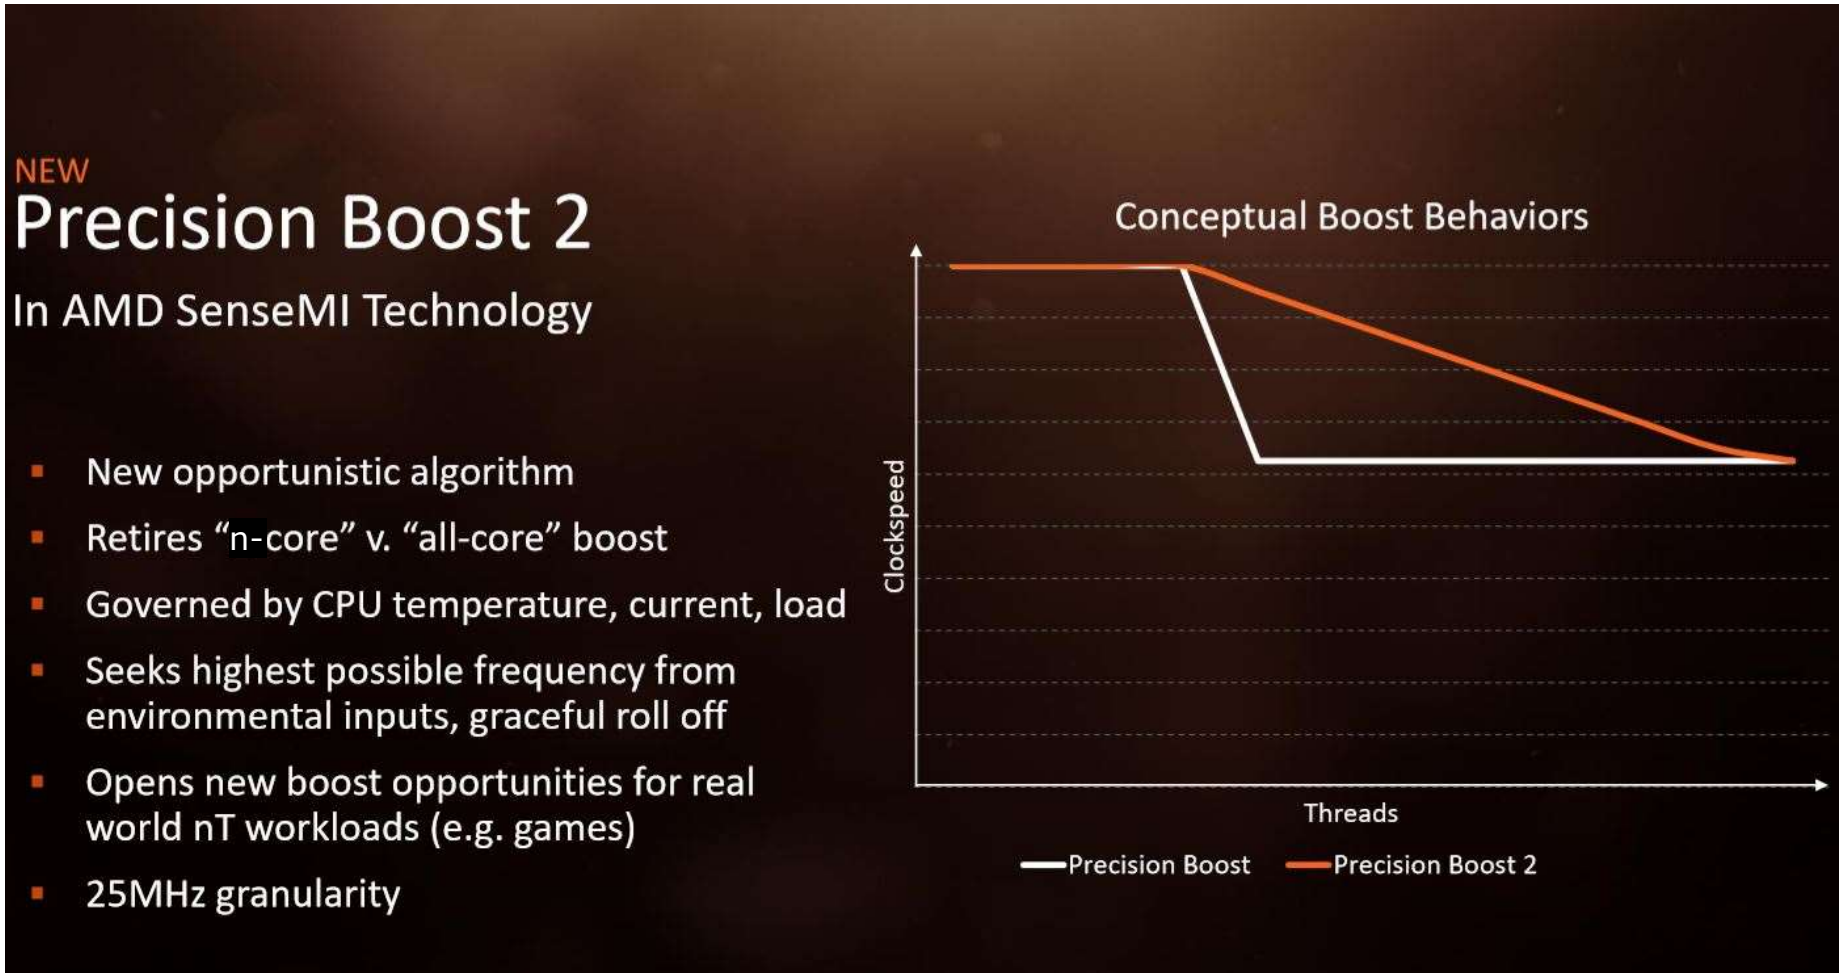

https://www.bit-tech.net/blogs/cpu-frequency-boosting-could-be-the-hot-topic-of-2018/1/

Principle of operation of Precision Boost 2 [83]

- Precision Boost 2 technology raises the clock frequency for any active core count as much as given limits (defined by the factory boost settings) allow it.
- The limiting factors are
  - SoC power ("PPT (Total Platform Limit"): the maximum amount of power the CPU can draw before boost levels off (measured in watts)
  - VRM current ("TDC Limit"): the maximum amount of current the motherboard is allowed to deliver to the CPU after warming up to a steady-state temperature before boost levels off
  - EDC current Limit - Electrical Design Current, maximum peak current that can be delivered by the motherboard' voltage regulator (VRM) for a short time, as a spike
  - Temp (°C): measured in degrees Celsius, the maximum temperature the CPU die can reach before boost levels off.

Conceptual view of implementing Precision Boost 2 [82]

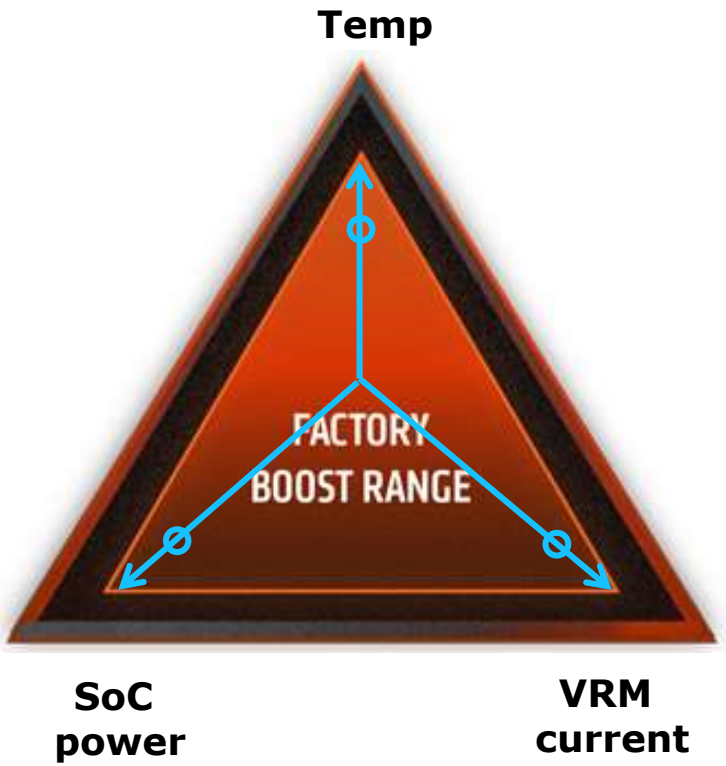

- The inner triangle illustrates the factory boost range settings, the blue circles represent actual operating points of the key limiting factors.
- Clock frequency and core voltage of all active cores will be increased in 25 MHz steps as long as one of the limiting factors reach the factory setting.

Precision Boost 2 technology as implemented in the 2. gen Ryzen DT line [63]

Boost behavior with Precision Boost 2
in the 2. gen. Ryzen Desktop line

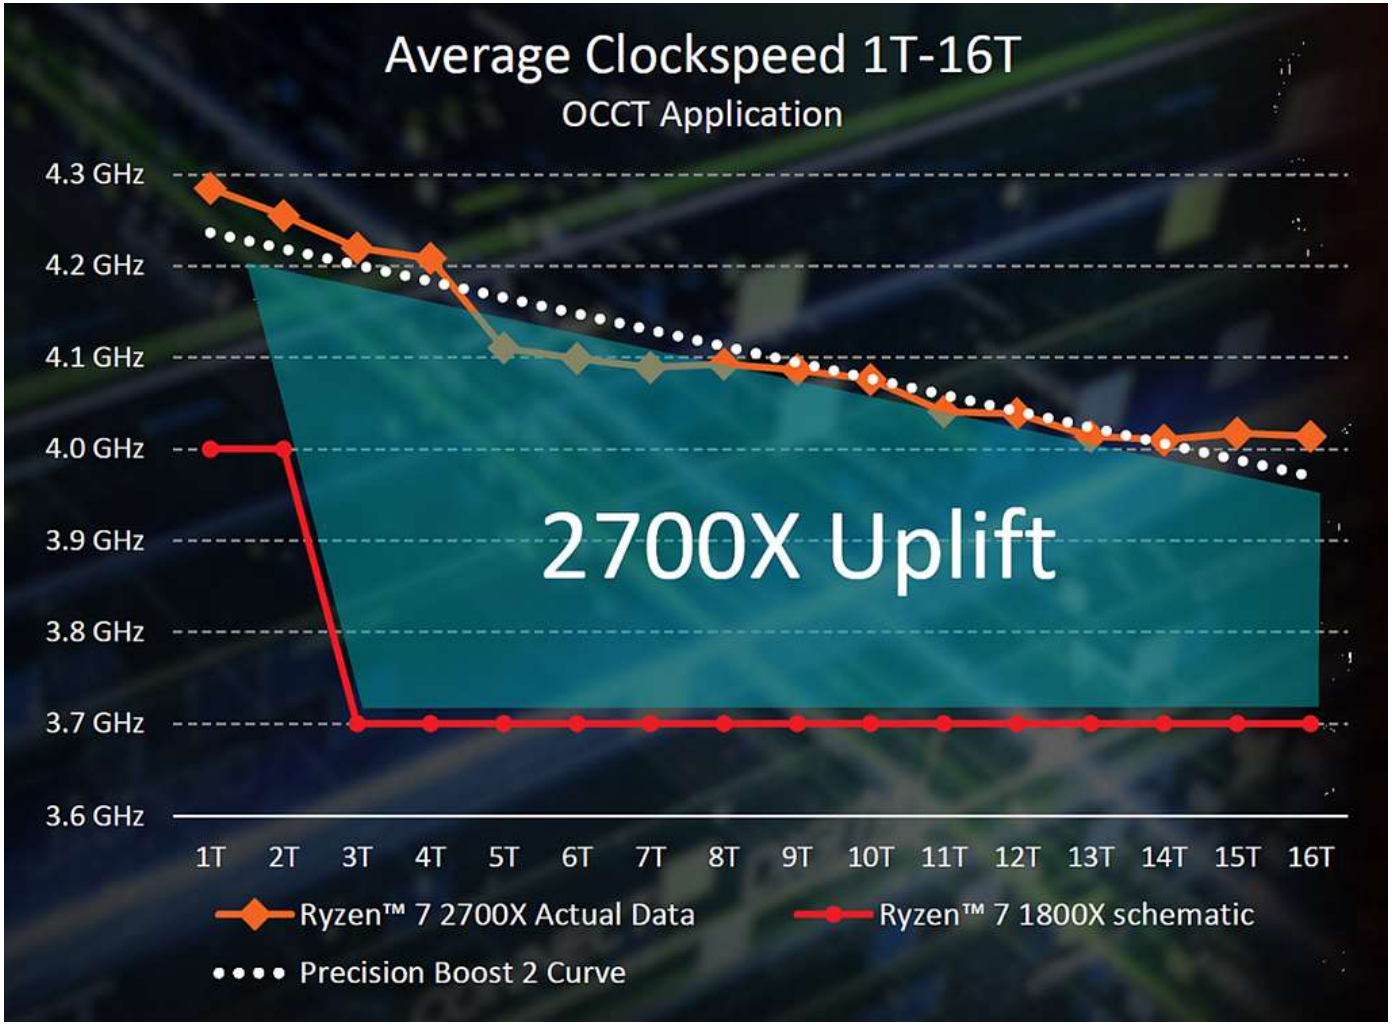

Power consumption of the cores while raising the number of active cores in the Ryzen 7 2700X 2. gen. desktop [63]

ANANDTECH	Ryzen 7 2700X (4+4) Loading (W), 105W TDP							
Core	2T	4T	6T	8T	10T	12T	14T	16T
0	23.16	20.71	19.22	16.92	14.53	12.23	10.78	11.06
1	0.09	20.67	18.94	16.88	14.47	12.45	10.99	11.20
2	0.08	0.13	19.31	16.90	14.56	12.28	10.77	10.79
3	0.05	0.11	0.12	16.73	14.40	12.34	10.87	10.44
4	0.49	0.05	0.05	0.27	0.49	0.12	1.35	10.35
5	0.04	0.36	0.14	0.86	0.44	2.53	12.34	10.82
6	0.04	0.52	0.06	0.54	1.07	14.58	12.34	10.65
7	0.09	0.11	0.11	0.26	17.29	14.75	12.22	10.75
Total	24.04	42.66	57.95	69.36	77.25	81.28	81.66	86.06

There are up to two active CCX units with up to 4 active cores/unit.
Green figures: not loaded cores - Yellow to red figures: loaded cores

Red figures appearing for few active cores indicate that in this case presumably die temperature will constrain turbo clock frequency and determine the related power consumption, whereas for yellow figures, seeing at many active cores, the total power consumption limits clock frequency.

b) Precision Boost Overdrive (PBO) [82]

It is a new feature of 2. gen. Ryzen DT and 2. gen. ThreadRipper processors, aiming at increasing multihreaded performance.

It requires however the use of specific chipsets, as follows

- in case of the 2. gen. DT series the series 400 chipsets (X470/B450) or
- in case of the 2. gen. ThreadRipper series the X399 chipset

and the Ryzen Master overdrive utility release 1.3 or higher.

Unlike traditional overlocking while PBO is activated both Precision Boost 2 and XFR2 (eXtended Frequency Range) remains enabled.

Relative merit of PBO [81]

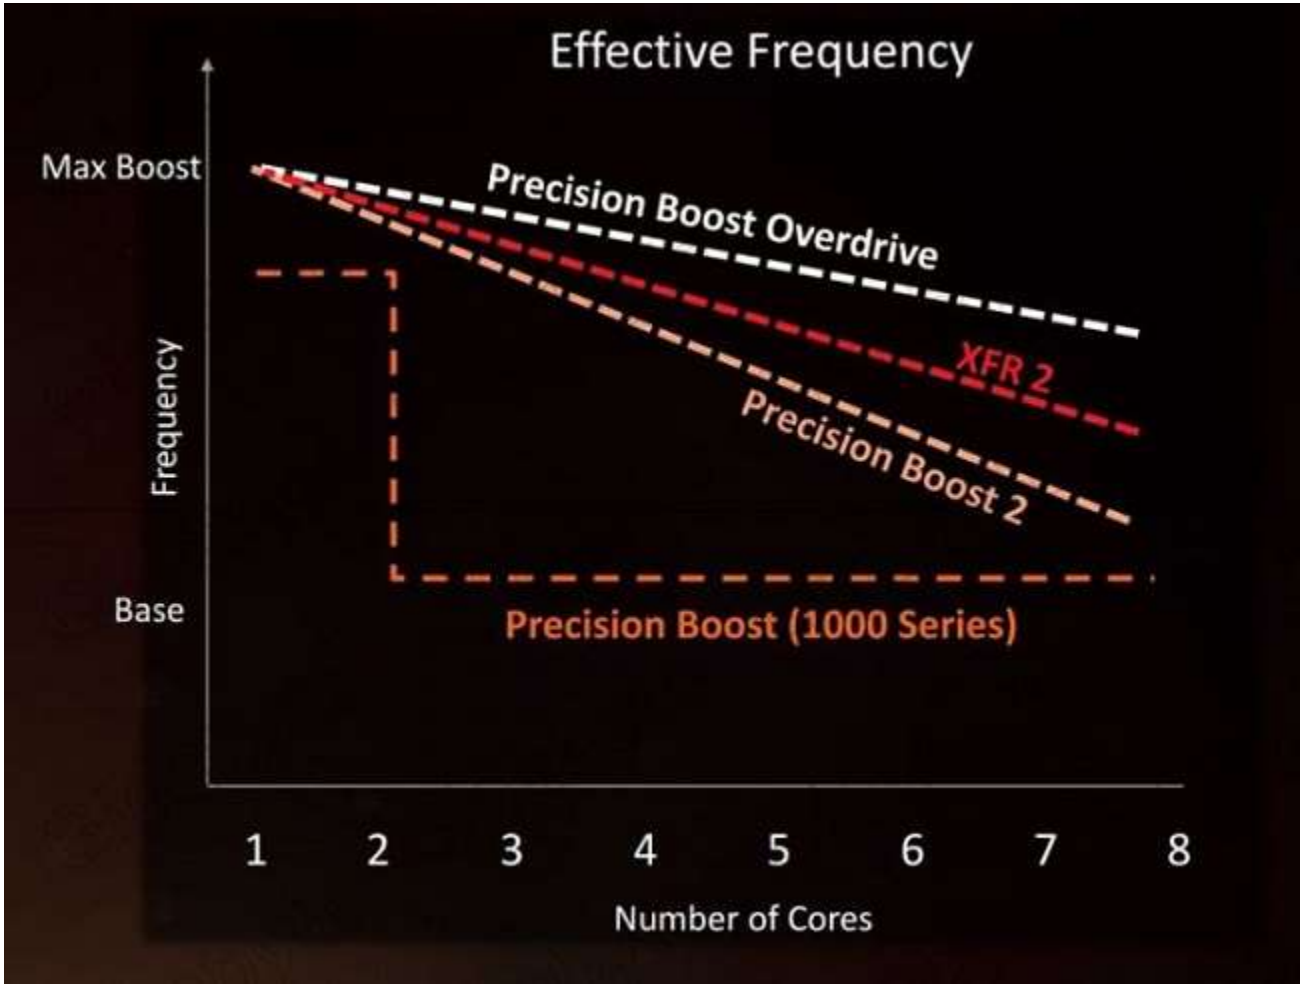

Principle of operation of PBO [83]

- When PBO is activated by means of the Ryzen Master overlocking utility the factory settings that determine the operation limits of Precision Boost 2 will be overwritten and the processor will operate outside the factory settings.

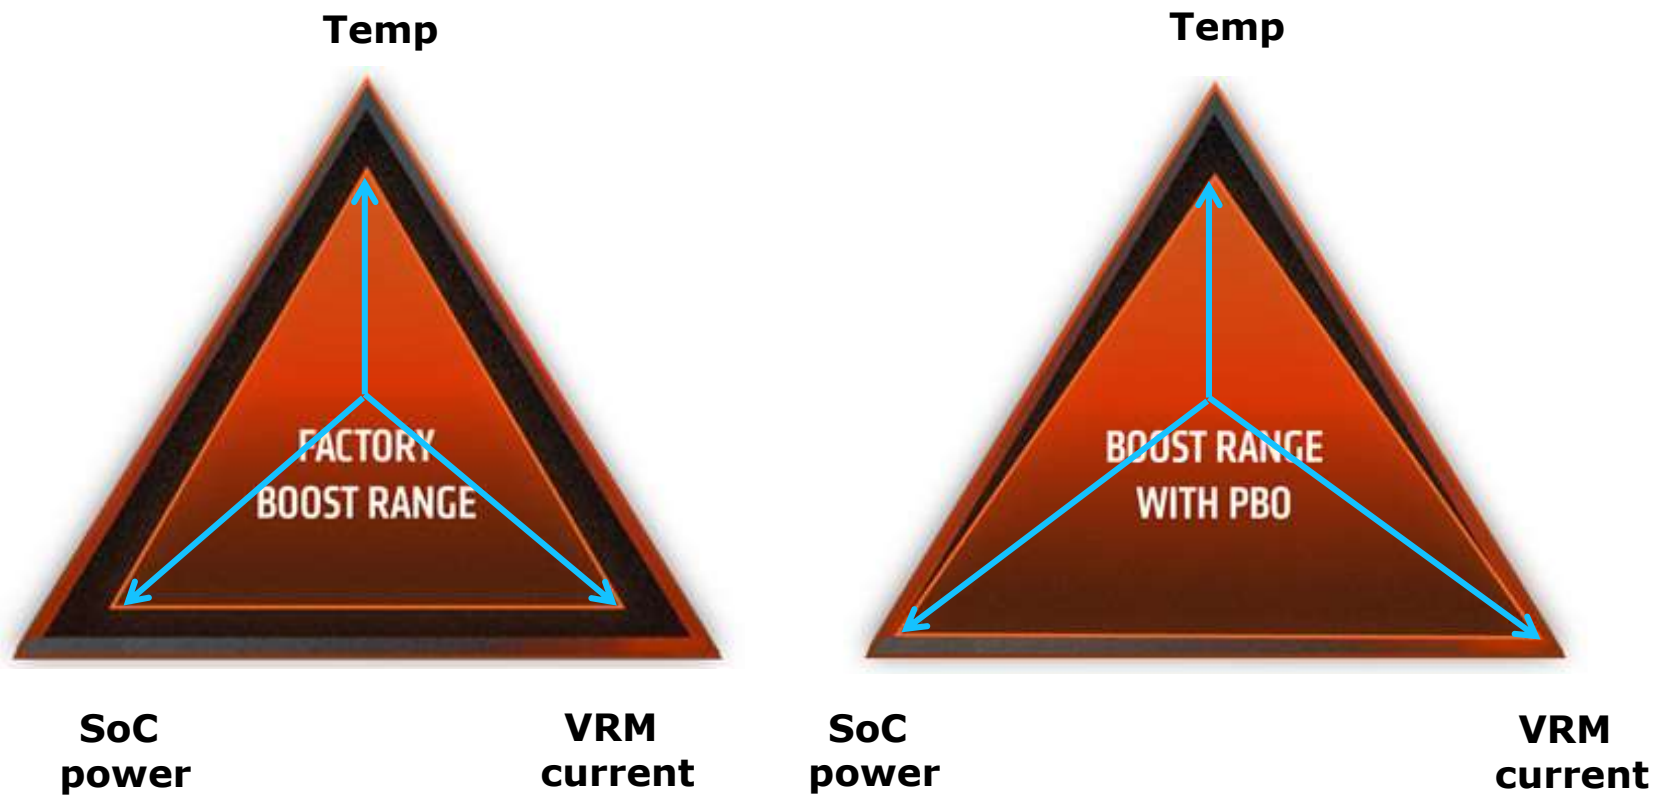

Figure: Conceptual view of extending the limiters of Precision Boost with PBO [82]

- Note that the temperature limit with PBO remains unchanged.

Example for extending the limiters of Precision Boost with PBO [83]

- For the 2. gen. Ryzen 2700X DT with 105 W TDP the default limits of Precision Boost 2 are as follows:
  - PPT (Total Platform Limit): 141.75 W
  - TDC (Thermal Design Current): 95 A
  - EDC Electrical Design Current: 140 A and
  - Tjmax (Die temperature): 85^0 C.

  When Precision Boost Overdrive mode is enabled the operation limiters will be raised as given below:

  - PPT (Total Platform Limit): 1000 W
  - TDC (Thermal Design Current): 114 A
  - EDC Electrical Design Current: 168 A and
  - Tjmax (Die temperature): 85^0 C.
- Nevertheless, with enabling the PBO mode invalidates AMD's warranty.

Contrasting implications of using overclocking and PBO [73]

Experience	Out-of-box	Overclock	PBO
nT Frequency	Stock	User Controlled	PBO Controlled
Power Limits	Stock	User Controlled	Increased
Idle Downclocking	Yes	No	Yes
Warranty	Yes	No	No
Precision Boost 2	On	Off	On
Performance	Stock	Increased	Increased

Example for performance improvement achieved by employing PBO in a ThreadRipper 2990WX for a multithreaded benchmark [82]

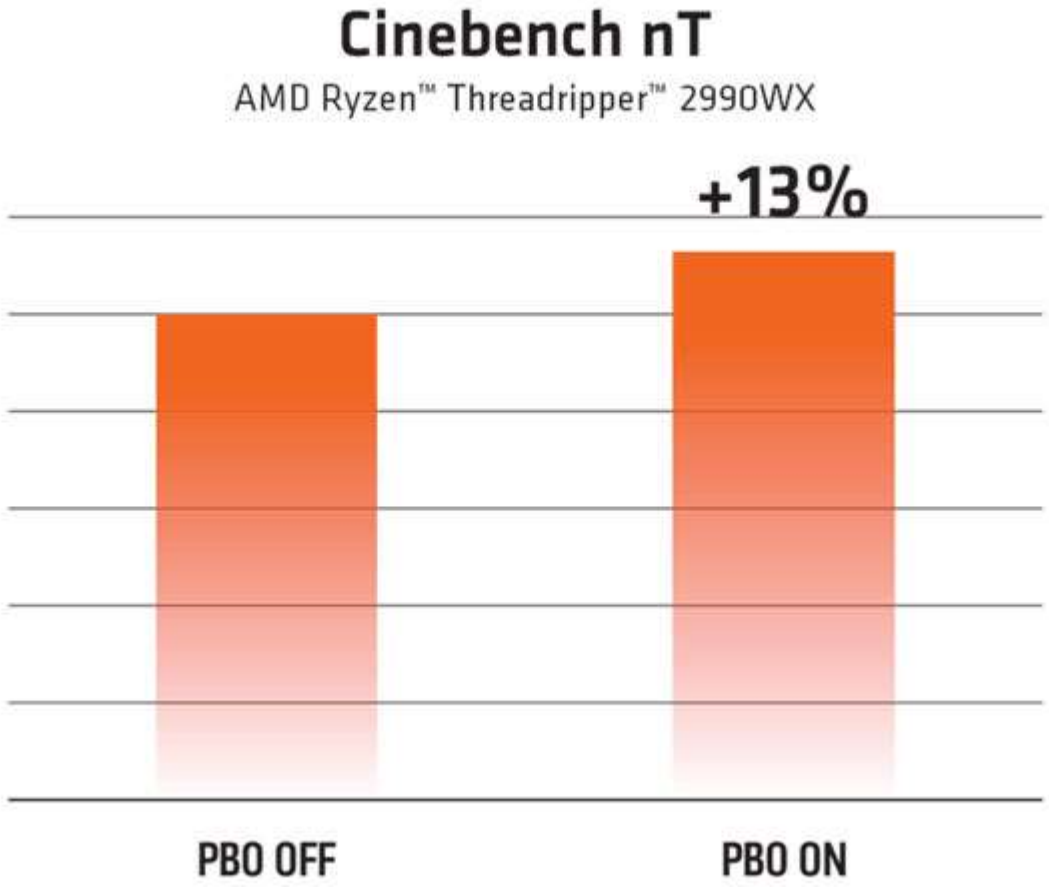

c) The XFR2 (eXtended Frequency Range) technology [63]

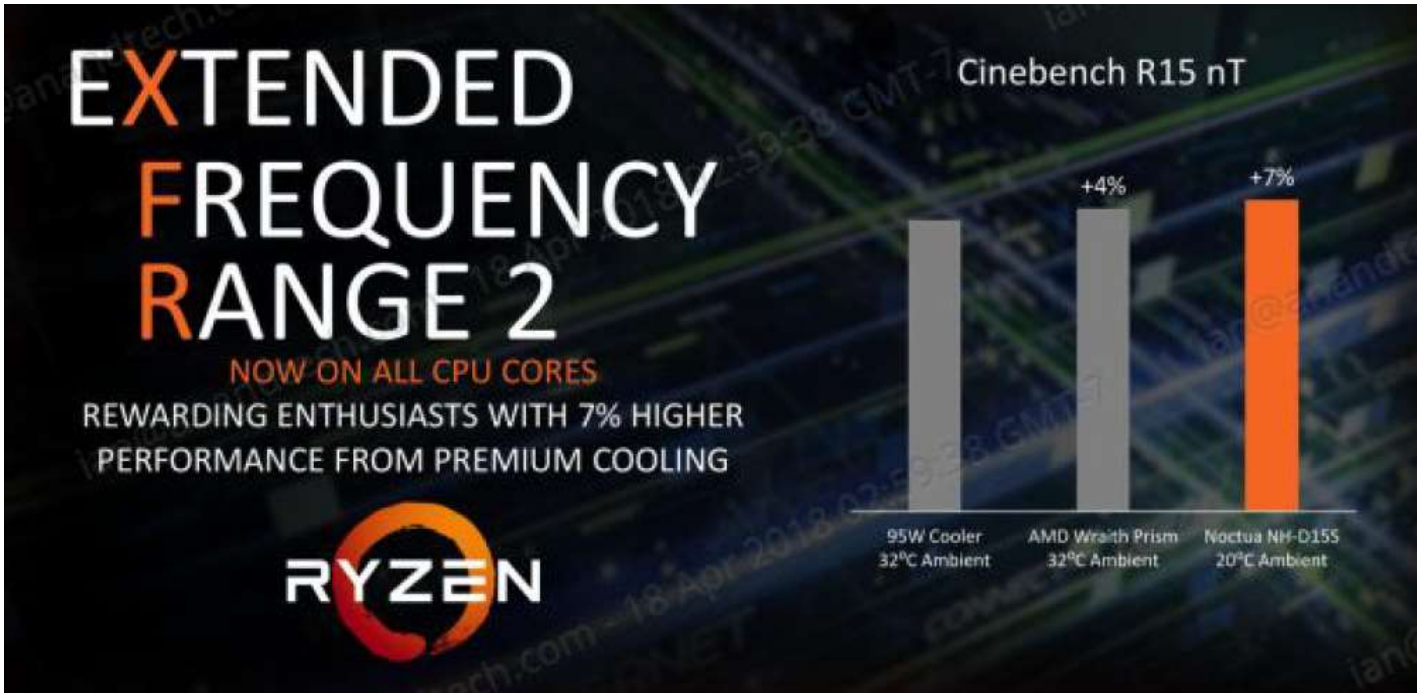

- In the previous generation certain processor models could raise clock frequency over the turbo boost frequency by 50 to 200 MHz if a premium cooling system was used.
- XFR2 allows a higher clock frequency boost assuming a premium cooling system and a CPU temperature below 60 C^0.

d) The StoreMI technology [79] -1

The StoreMI technology combines fast SSD (up to 256 GB NVMe or SATA) and large capacity HDD storage into a single drive (as indicated below) and automatically moves the data accessed the most to the SSD.

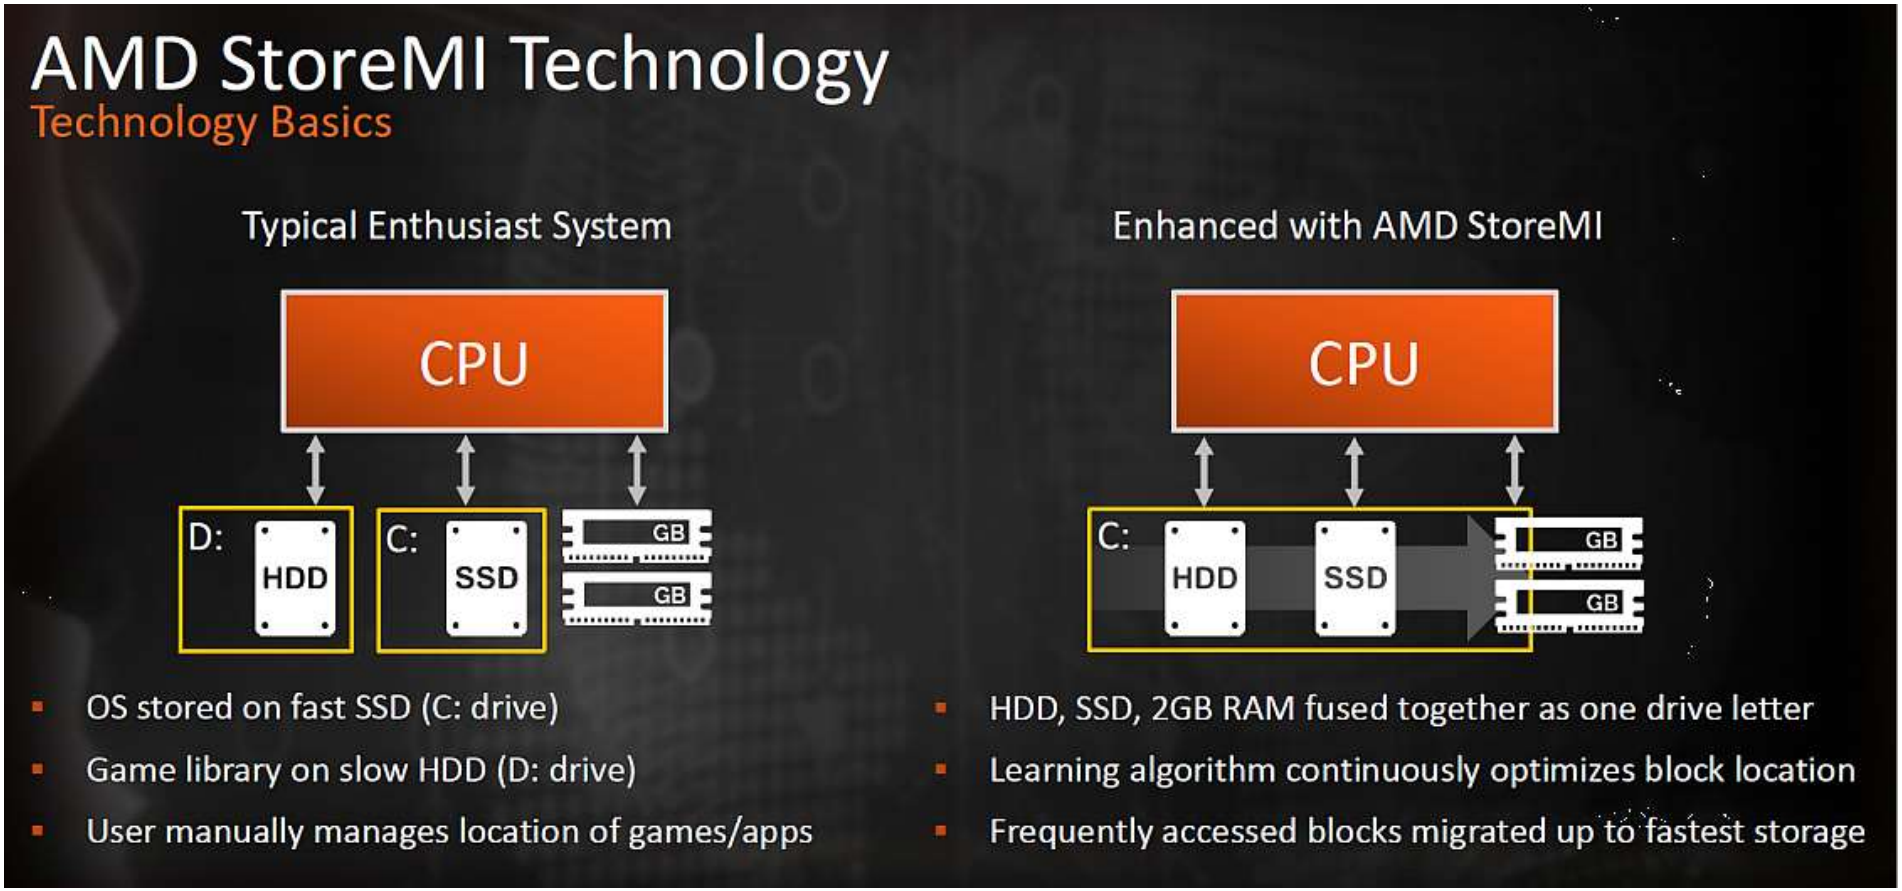

Figure: The StoreMI technology providing a unified, fast, large capacity drive [79]

d) The StoreMI technology [79] -2

- The unified storage has the responsiveness of SSD, and capacity of low priced mechanical HDD.
- In addition up to 2GB of DRAM can be included to have a last-level cache for ultra-fast data access.
- The StoreMI technology requires for 2. gen. Ryzen 2000-series DT processors an AMD 400-series or for 2. gen. ThreadRipper processors an AMD X399 chipset.
- The storage hierarchy is controlled by the AMD's StoreMI software utility that initially was sold for $20 but later became free of charge.

Benefit of AMD's StoreMI technology in speeding up starting applications [66]

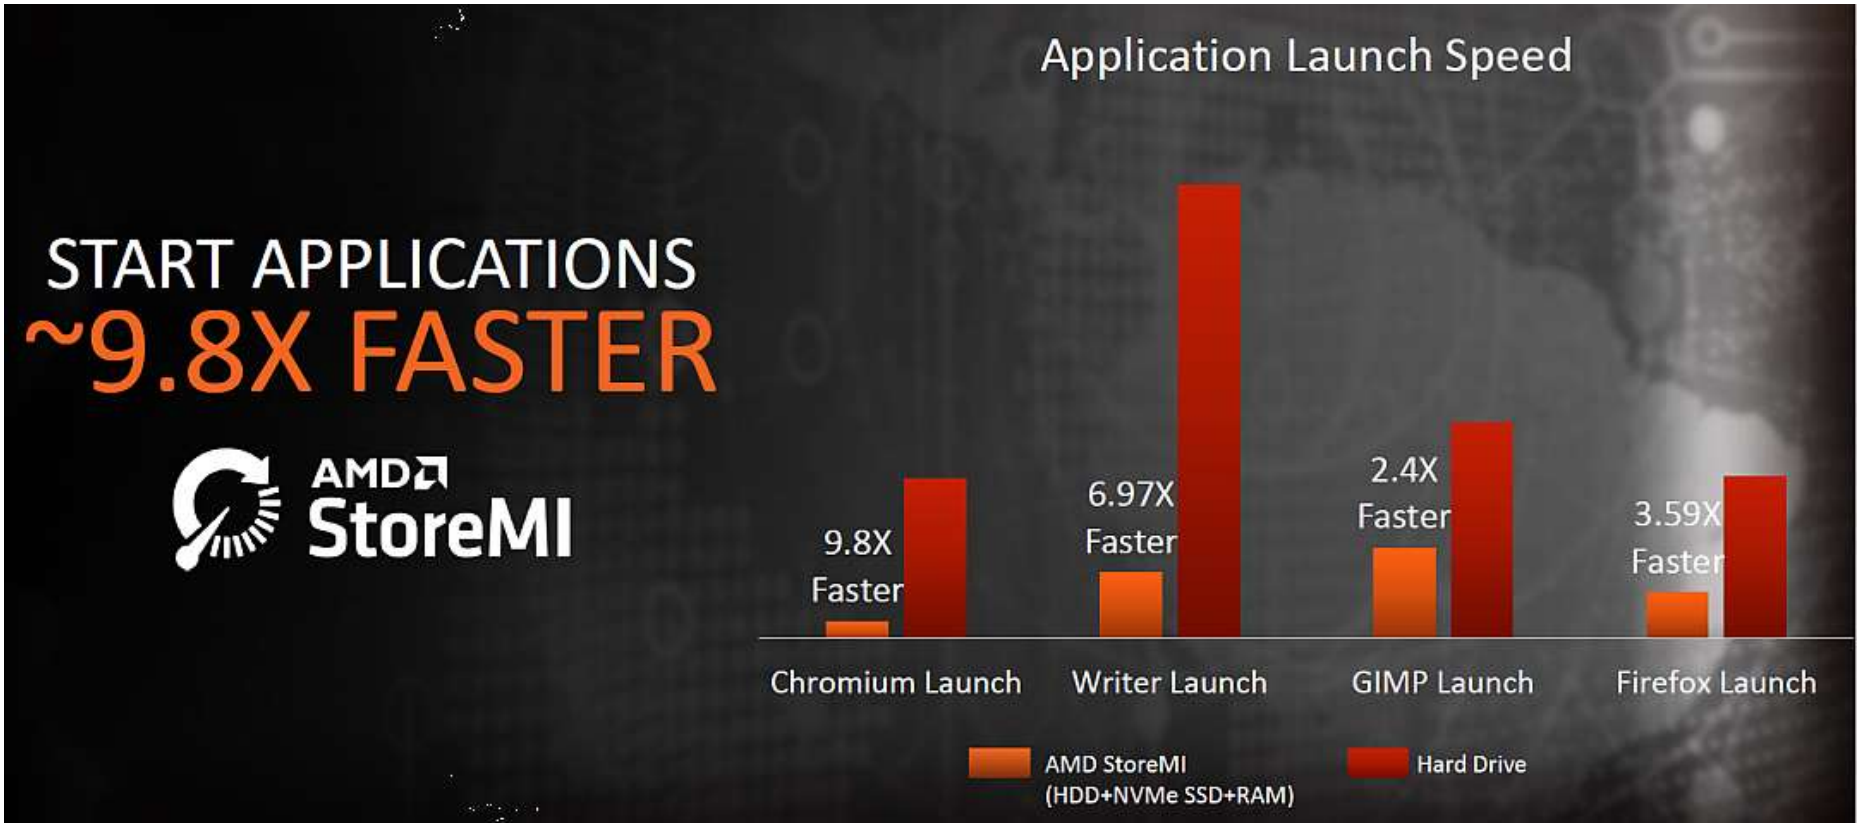

Intel's competing solution: their Octane technology

Along with the 7th generation Core line, called the Kaby Lake line, Intel introduced in 2017 their Octane memory technology and enhanced it in their subsequent 8th generation Coffee Lake line in 2018.

In order to contrast it with AMD's StoreMI technology next we give a brief account of Intel's Octane memory enhancement solution.

Basic facts of Intel's Octane memory

- The Optane memory is nonvolatile memory typically used as a cache of a HDD drive.
- It is based on the 3D XPoint memory technology, announced by Intel and Micron
- in 2016.
- It has a typical size of 16 or 32 GB.
- It is mounted on an M.2 card that is attached via 2 to 4 PCIe lanes.
- Its use needs the Rapid Storage Technology driver (appropriate release no.).

*

Example: Optane memory installed on a mainboard [85]

Enhancements of the 2. gen. Optane memory

- In the 1. gen. Optane memory, introduced along with the Kaby Lake line in 2017, a single HD Boot drive was assumed and Octane memory served as a cache of it.
- By contrast, its 2. gen., introduced along with the Coffee Lake line in 2018, assumes a fast SSD Boot drive and a large HD data drive.

  In this case the Optane memory serves as a cache for the data drive (see below).

	1. generation	2. generation
Storage system assumed	HD Boot drive and Optane memory	Fast SSD Boot drive, a large HD data drive and Optane memory
Principle of operation	Optane memory serves as a cache of the Boot drive (HD)	Optane memory serves as cache of the large HD data drive
Processor support	7. Generation (Kaby Lake)	8. Generation (Coffee Lake)
PCH support	200-series PCH or later	300-series PCH
Intro	2017	2018
Driver	RST 15.5 or later	RST 16.02 or later
OS	Windows 10 64-bit	Windows 10 64-bit

Table: Key features of the 1. and 2. generation Optane memory

*

Benefits of using Intel's 2. generation Optane memory [80]

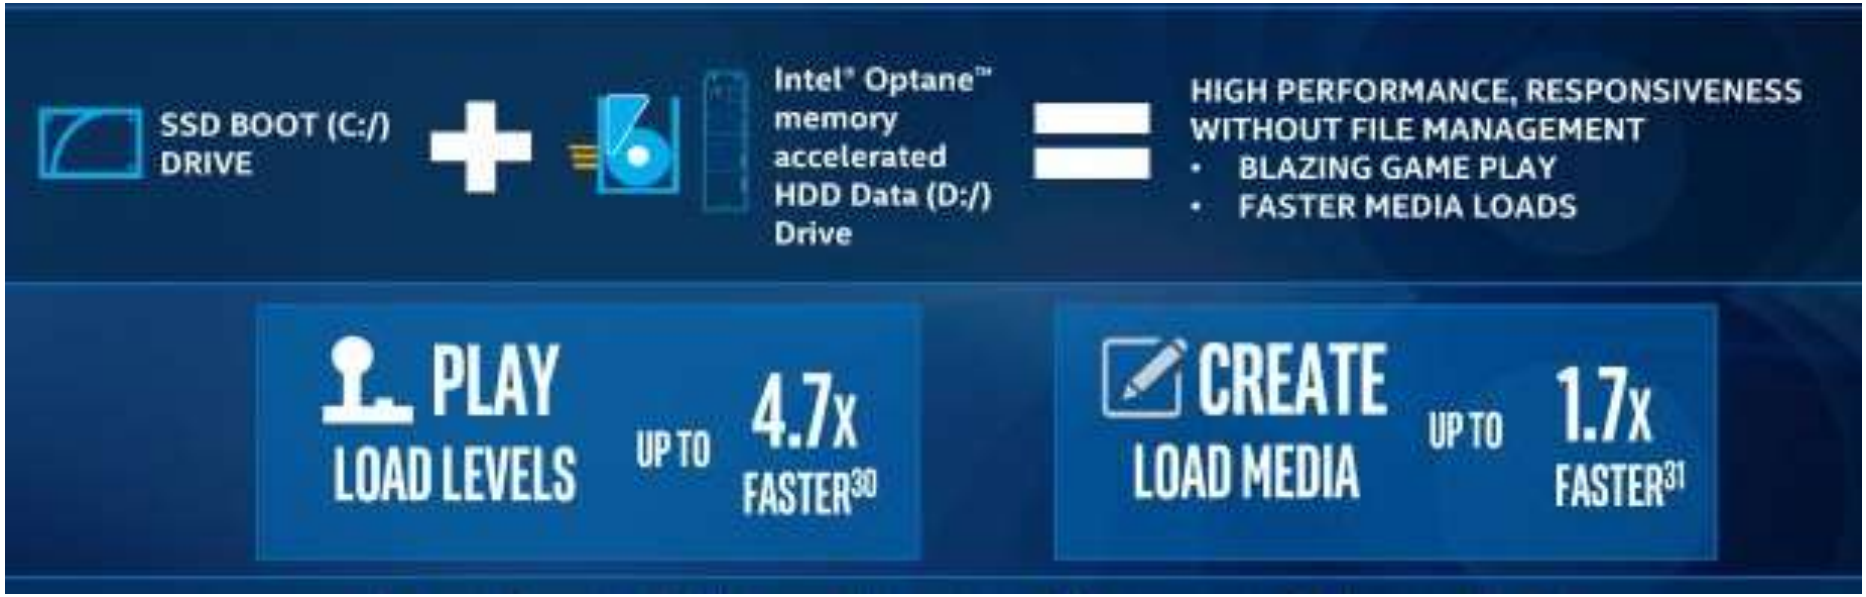

*

Increasing IPC in AMD's subsequent desktop series [63]

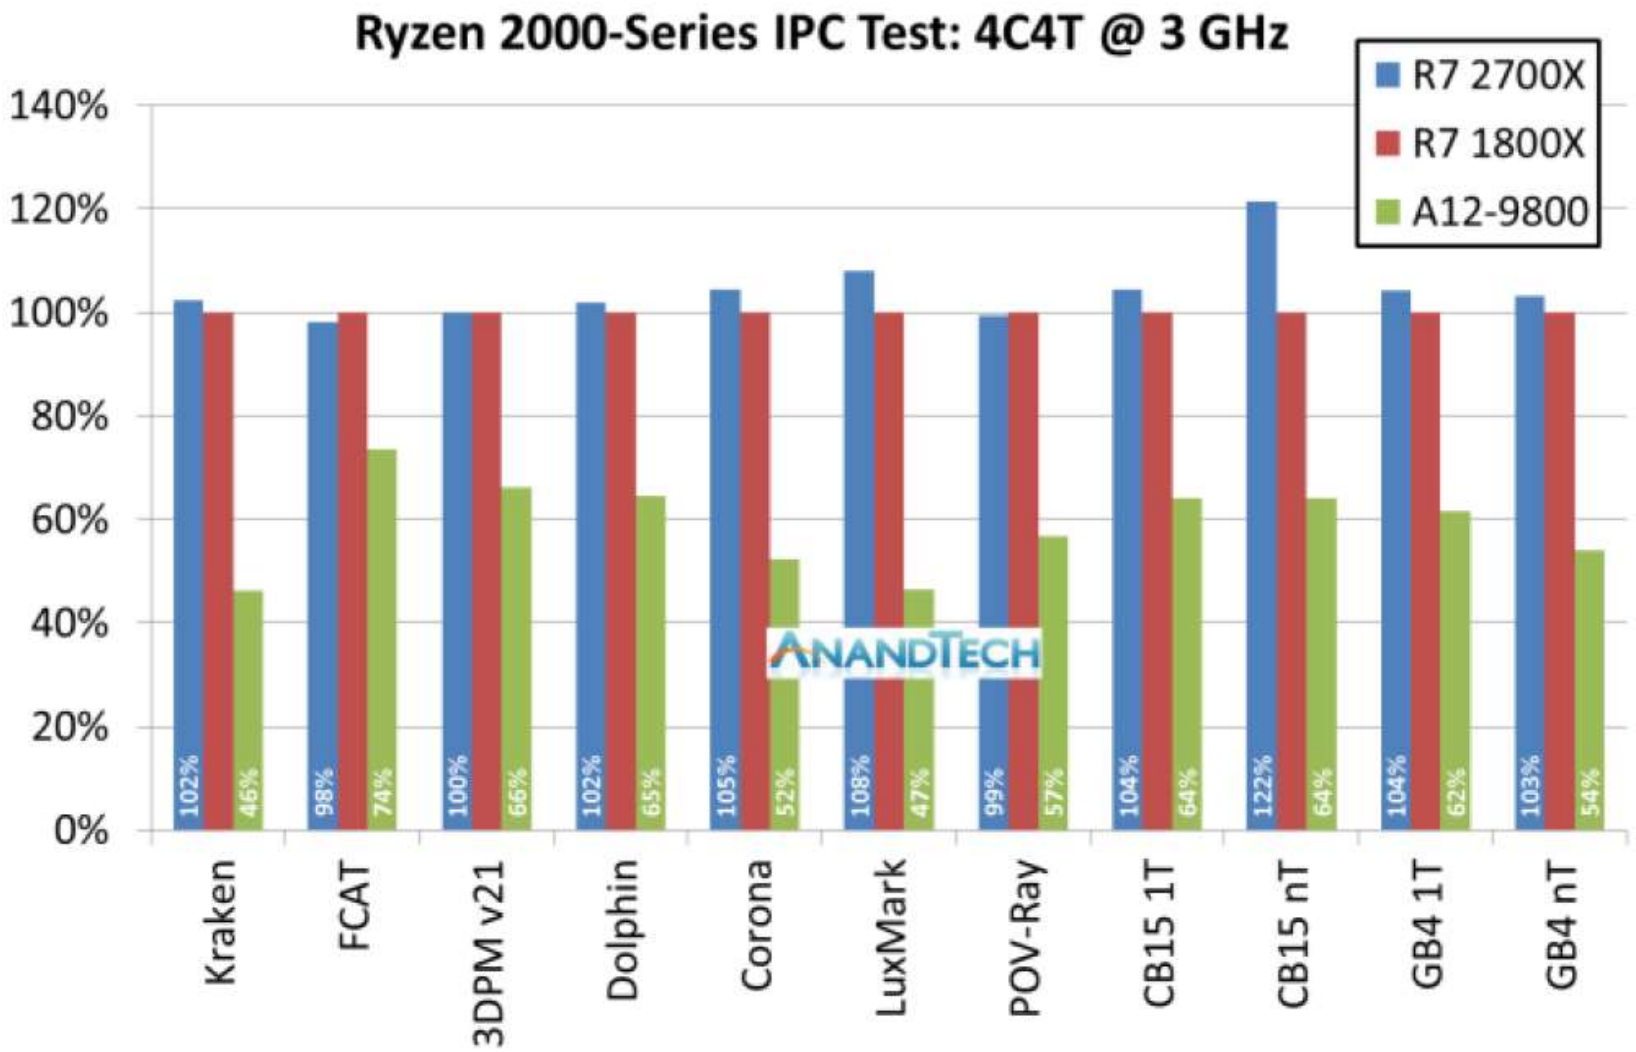

Comparing single-threaded performance of competing desktop processors [63

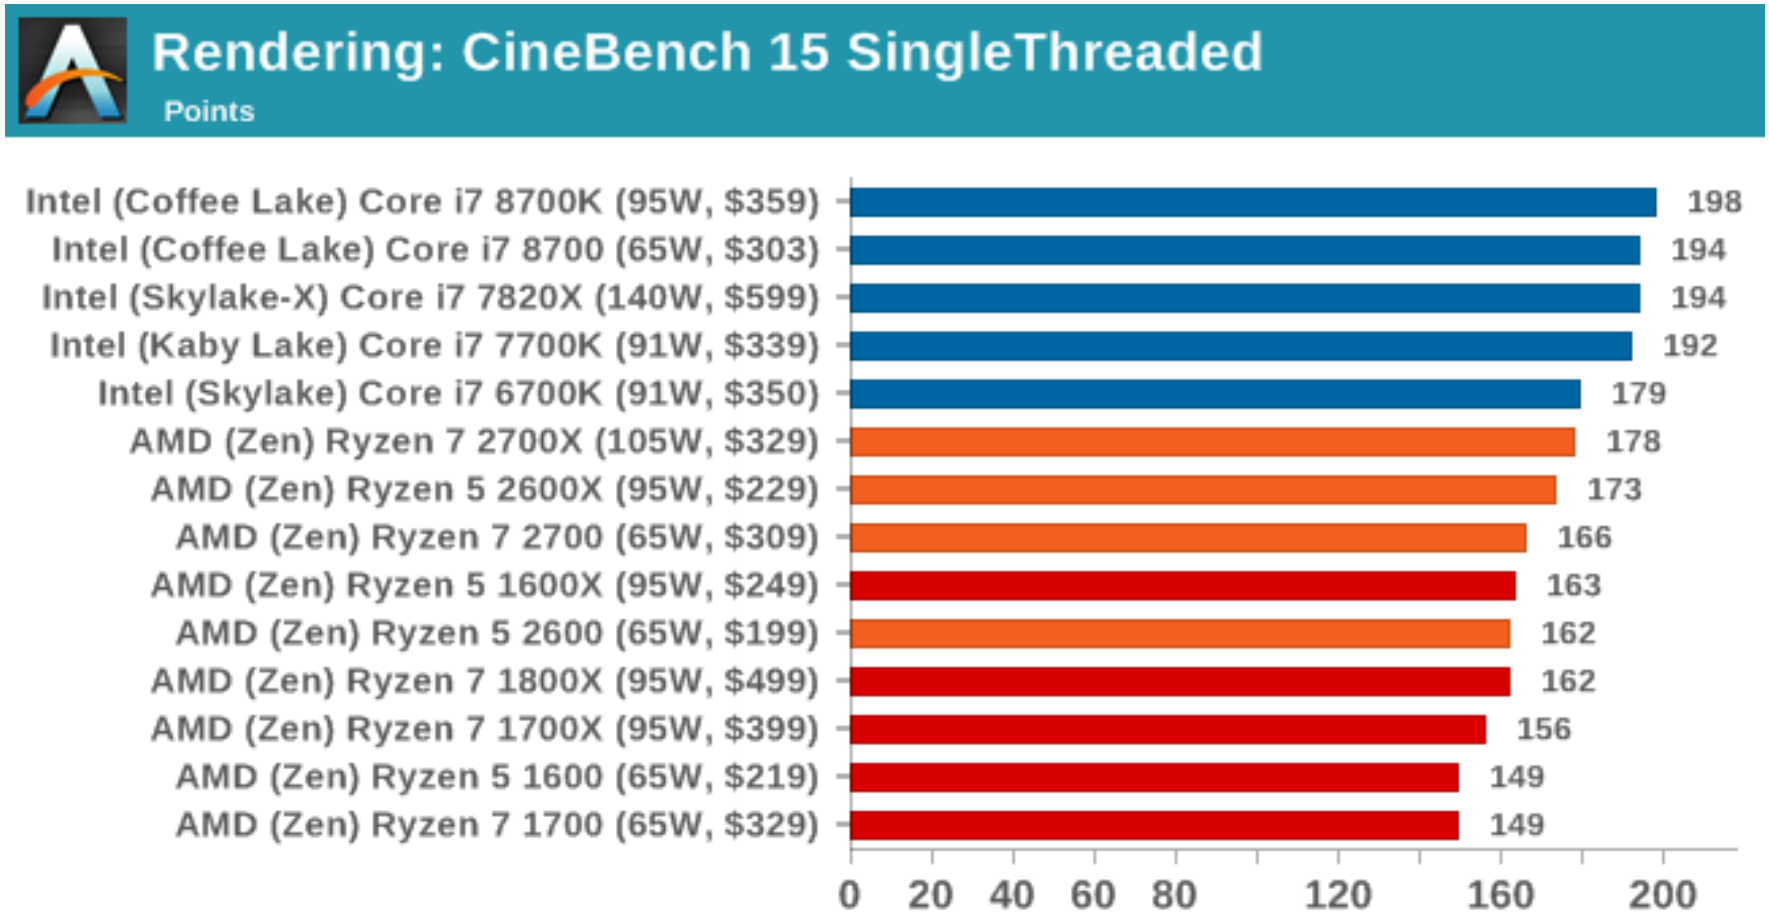

As the Figure indicates Intel's 7. and 8. gen. desktops provide the highest single thread performance for the benchmark used, nevertheless the performance gap of AMD's Ryzen+ processors became smaller.

Comparing multi-threaded performance of competing desktop processors [63]

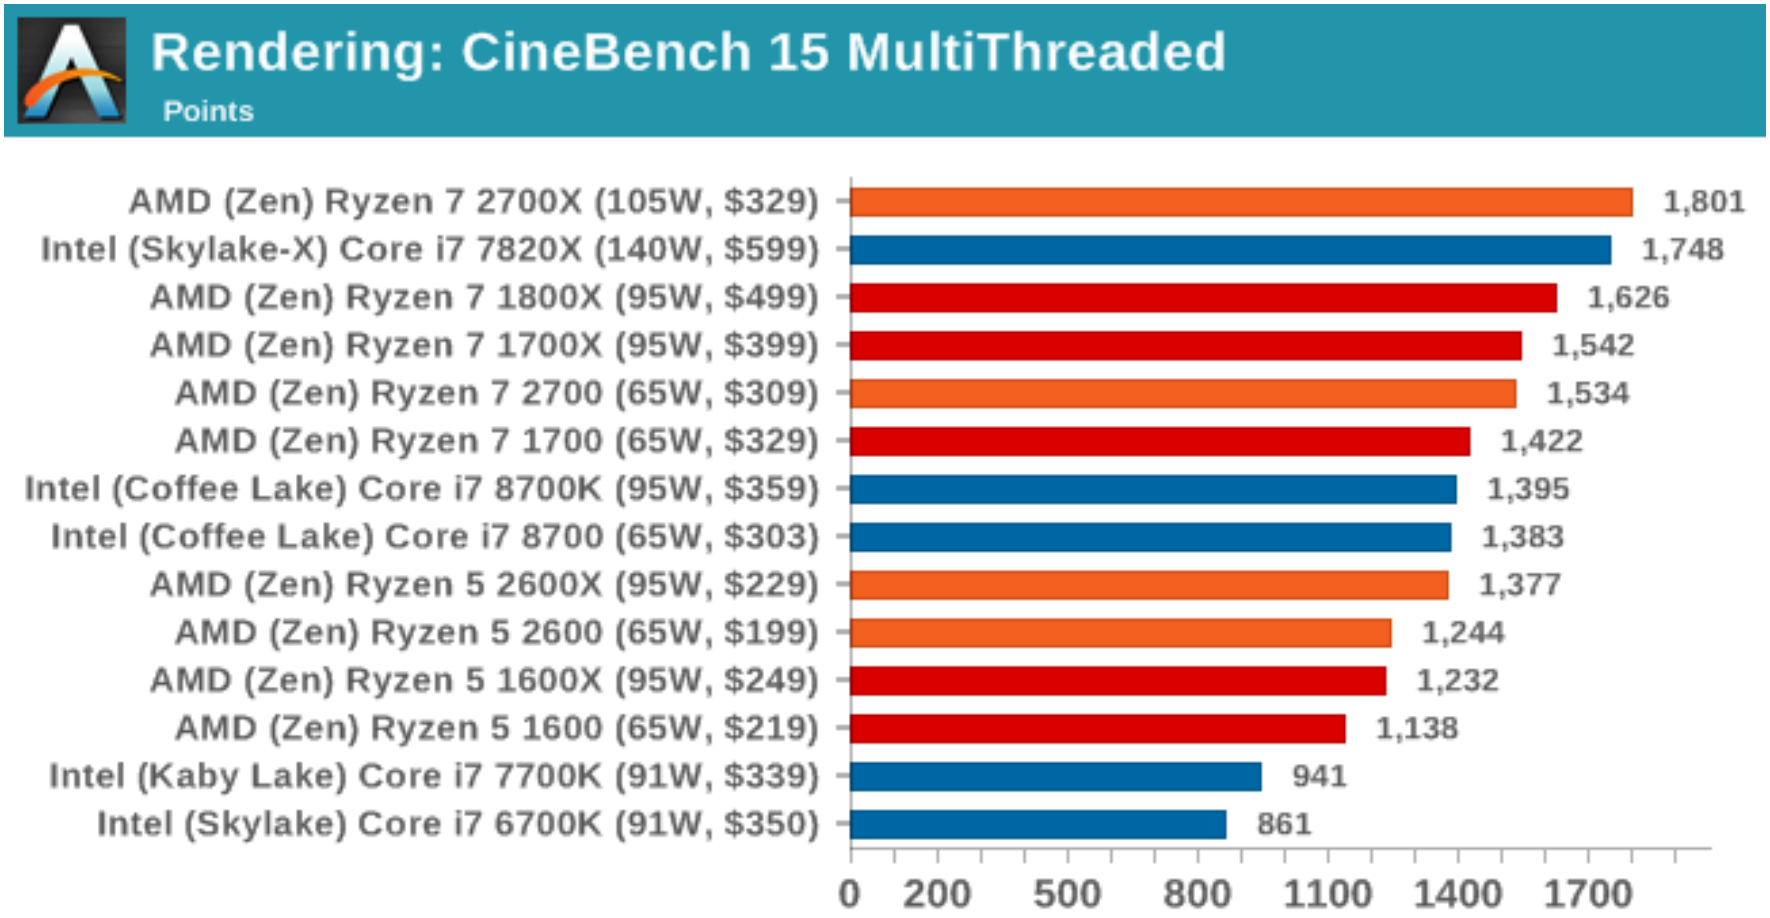

The Figure indicates that AMD's Ryzen+ based 8-core desktops have performance and price advantages over Intel's 4 or 6-core 7. and 8. gen. desktop processors.

6.1.3 The 3. gen. Ryzen desktop line without GPU (Matisse)

6.1.3 The 3. gen. Ryzen desktop line without GPU (Matisse) [64]

6.1.3 The 3. gen. Ryzen desktop line without GPU (Matisse) [102] -1

- Code-named as the Matisse line.
- Announced in 1/2019, to be launched in Q2/2019 or Q3/2019.
- Implemented as an MCM package, consisting of two dies:
  - an eight-core 7nm CPU chiplet, made at TSMC, and
  - a 14nm input/output chiplet with dual memory controllers and PCIe 4.0 x16 lanes, made at GlobalFoundries.

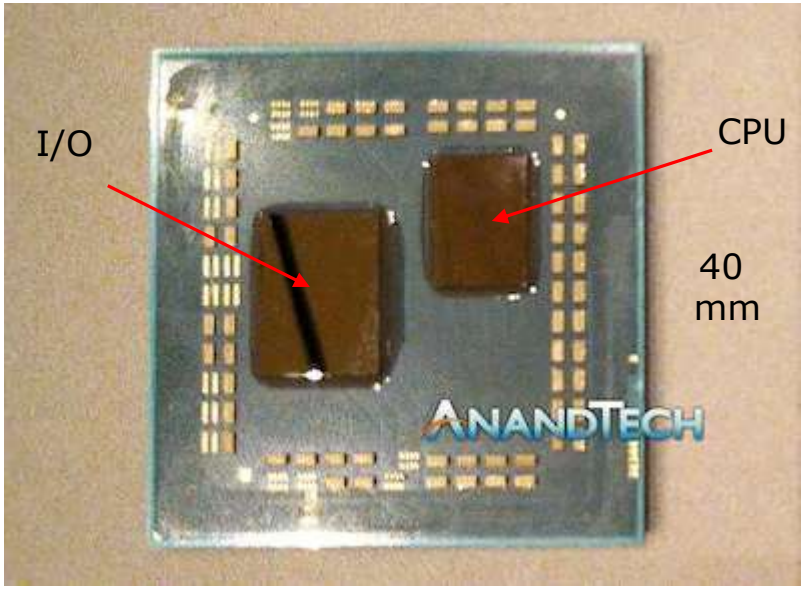

Figure: Layout of a 3. gen. Ryzen desktop processor [102]

https://www.anandtech.com/show/13829/amd-ryzen-3rd-generation-zen-2-pcie-4-eight-core

6.1.3 The 3. gen. Ryzen desktop line without GPU (Matisse) [102] -2

- It will be the world's first 7nm gaming CPU and the world's first mainstream CPU to support PCIe 4.0 x16.
- They will have AM4 socket.
- It will work in 300 or 400 series chipsets.

Performance expectations [102]

- First Cinebench 15 results show about 15 % performance increase vs. the the 2nd Generation Ryzen 7 2700X , as seen below.
- This indicates about the same performance level as Intel's flagship Core i9-9900K.

AMD Benchmarks at CES 2019						
	System Power	Idle Power*	Chip Power	CB 15 MT Score (pre-brief)	CB 15 MT Score (on-stage)	All-Core Frequency
AMD Zen 2	130W	55W	75W	2023	2057	?
Intel i9-9900K	180W	55W	125W	2042	2040	4.7 GHz

Table: First Cinebench 15 scores of AMD's 3. gen. Ryzen desktop processors [102]

- Note that AMD's 3. gen. Ryzen desktop processors consume much less power than Intel's i9-9900 (75 W vs. 125 W).

Speculations over possible extensions in the MCM [102]

As the Figure below shows, the MCM has place for a second die as well, probably for a second CPU die to raise the core count to 16 or for a GPU.

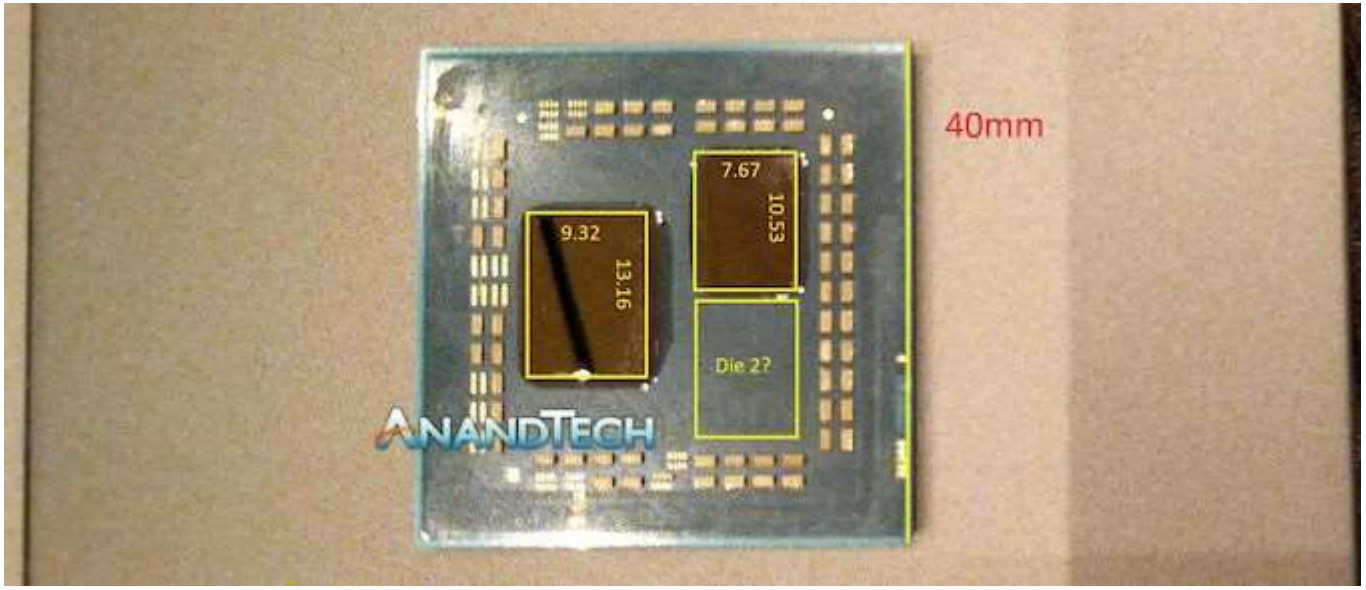

Figure: Possible extension of the MCM of the 3. gen. Ryzen desktop processor [102]

6.2 The Ryzen desktop APU line

6.2 The Ryzen desktop APU line

- Code-name of the CPU: the Raven Ridge line.
- The line includes the models
  - Ryzen 5 2400G with Radeon RX Vega 11 and
    Ryzen 3 2200G with Radeon RX Vega 8
    graphics, launched in 02/2018 and the
    Ryzen 5 2400GE with Radeon RX Vega 11 and
    Ryzen 3 2200GE with Radeon RX Vega 8
    graphics, launched in 04/2018.
- The Ryzen desktop APU line is unlocked.
- The line is manufactured on the 14 nm technology (by GlobalFoundries).
- They have AM4 socket.
- Unlike the 'G' models, however, the 'GE' processors are not bundled with a cooler.

Main features of AMD's Ryzen desktop APU models [68]

Raven Ridge (Zen-based) vs. Bristol Ridge (Excavator-based)						
	Ryzen 5 2400G	Ryzen 5 2400GE	A12-9800	Ryzen 3 2200G	Ryzen 3 2200GE	A10-9700
Core uArch	Zen	Zen	Excavator	Zen	Zen	Excavator
Launched	02/2018	04/2018	09/2016	02/2018	04/2018	09/2016
Cores/Threads	4 / 8	4 / 8	2 / 4	4 / 4	4 / 4	2 / 4
Base CPU Frequency	3.6 GHz	3.2 GHz	3.8 GHz	3.5 GHz	3.2 GHz	3.5 GHz
Turbo CPU Frequency	3.9 GHz	3.8 GHz	4.2 GHz	3.7 GHz	3.6 GHz	3.8 GHz
TDP	65 W	35 W	65 W	65 W	35 W	65 W
cTDP	46-65 W	na	45-65W	46-65 W	na	45-65W
L2 Cache	512 KB/core	512 KB/core	1 MB/core	512 KB/core	512 KB/core	1 MB/core
L3 Cache	4 MB	4 MB	-	4 MB	4 MB	-
Graphics	Vega 11	Vega 11	GCN 3 Gen	Vega 8	Vega 8	GCN 3 Gen
Compute Units	11 CUs	11 CUs	8 CUs	8 CUs	8 CUs	6 CUs
Streaming Processors	704 SPs	704 SPs	512 SPs	512 SPs	512 SPs	384 SPs
Base GPU Frequency	1250 MHz	1250 MHz	1108 MHz	1100 MHz	1100 MHz	1029 MHz
DRAM Support	DDR4-2667	DDR4-2933	DDR4-2400	DDR4-2667	DDR4-2933	DDR4-2400
Price	$169	na.	$99	$99	na.	$79

7. The Ryzen Mobile line

- 7.1 The 1. gen. Ryzen Mobile line (Raven Ridge)
- 7.2 The 2. gen. Ryzen mobile line (Picasso)

Will not be discussed

7. The Ryzen Mobile line -1

AMD's mobile processors since 2011 [38]

Roadmap of AMD's notebook lines 2017 - 2019 [64]

7.1 The 1. gen. Ryzen Mobile line (Raven Ridge)

Will not be discussed

7.1 The 1. gen. Ryzen Mobile line (Raven Ridge) -2

- Launched in 10/2017
- Designated as Raven Ridge
- Main components of the mobile models:
  - a 4-core CCX complex and
  - a Vega GPU
  - interconnected by AMD's Infinity Fabric,

  as indicated in the next Figure.

*

High-level block diagram of the Ryzen Mobile processors [38]

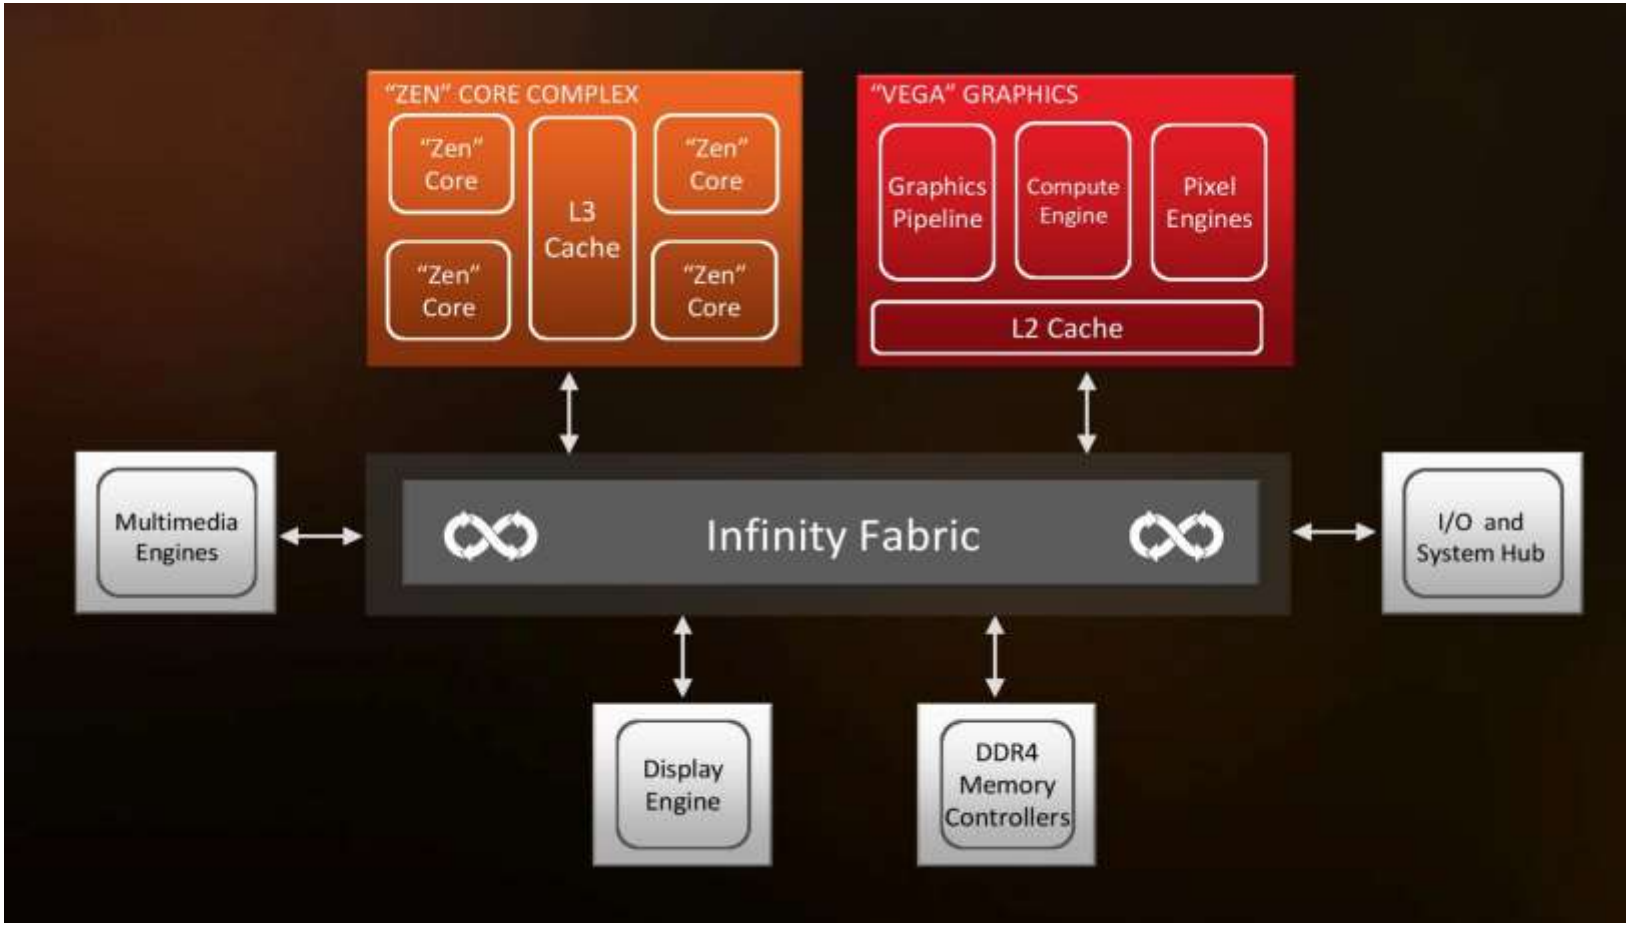

- One coherent control and data fabric integrates and manages the full SoC.
- The Infinity Fabric services 6 clients in the SoC.
- Development work needed 4 years.

*

7.1 The 1. gen. Ryzen Mobile line (Raven Ridge) (3)

Comparison: ARM's CCN-504 Cache-Coherent Interconnect based system architecture

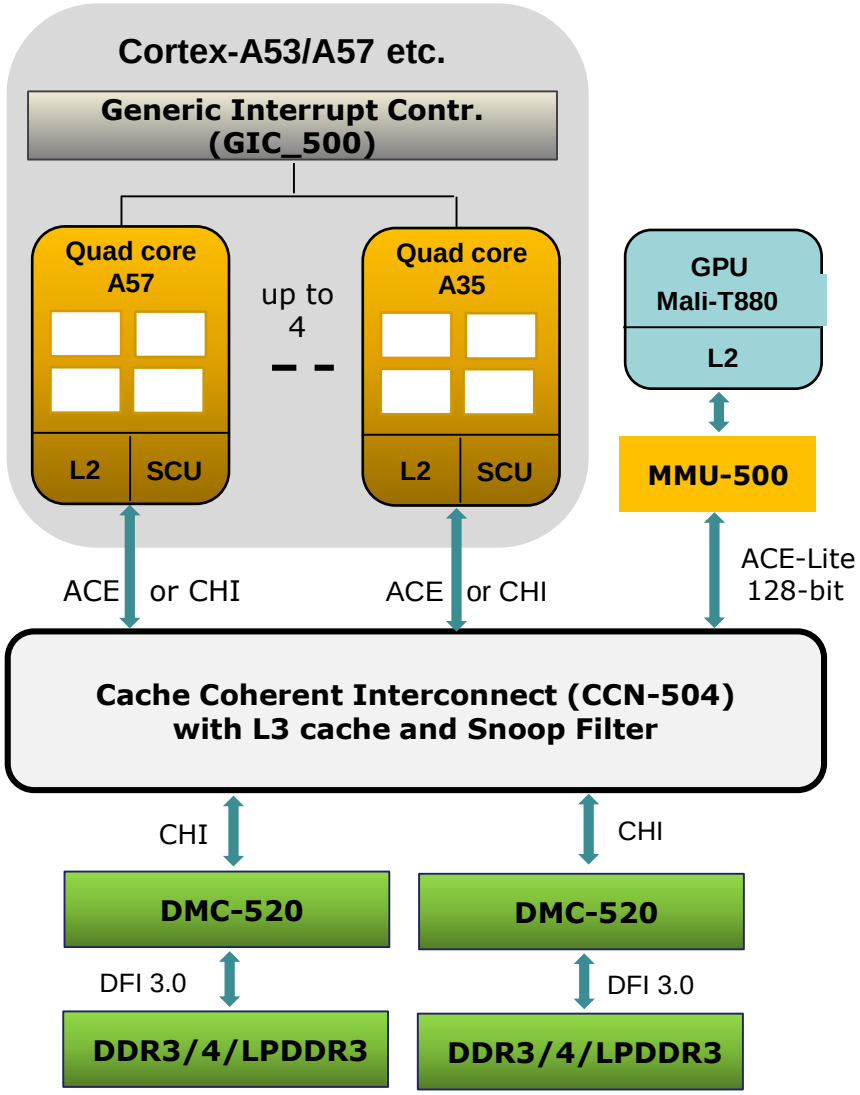

More detailed block diagram of the Ryzen Mobile processors [98]

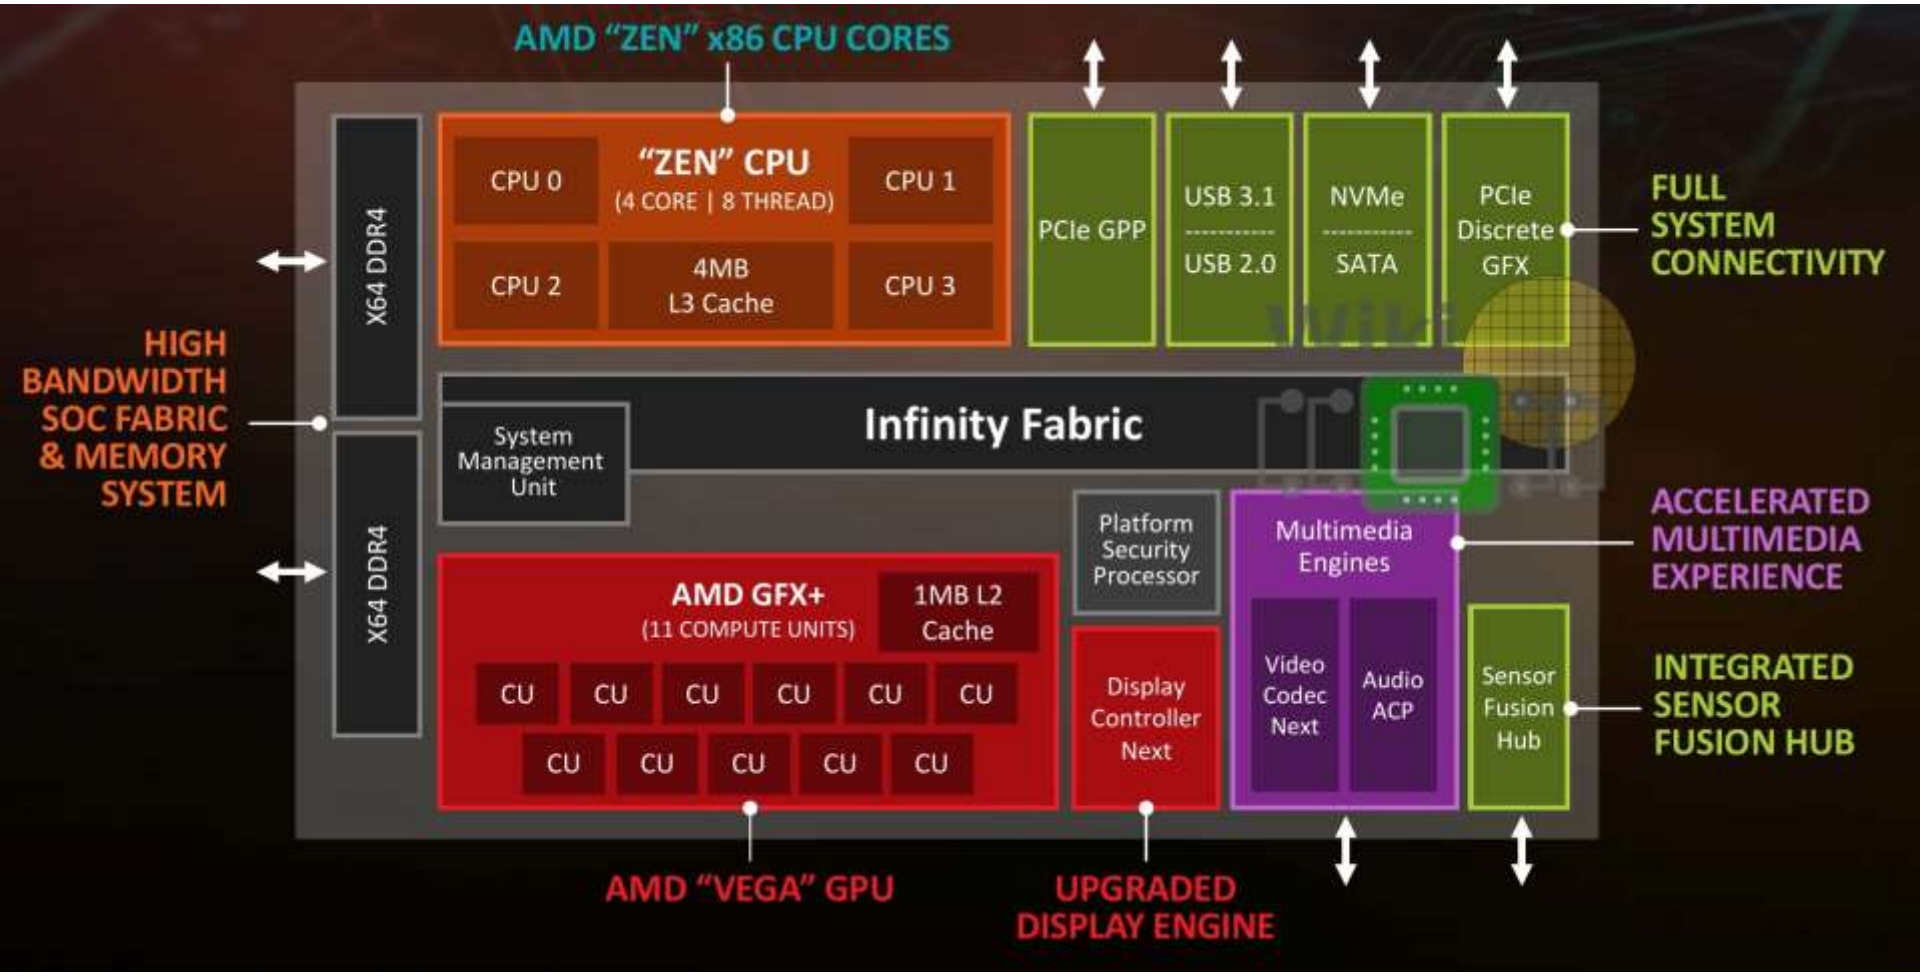

Implementation of the Infinity Fabric [98]

"Raven Ridge" Optimizations

- 32 Byte internal datapath width
- Up to 1.6GHz for bandwidth exceeding 50GB/s
- Up to 5 transfers/clock per switch
- Improved CPU latency under load, while maintaining DRAM efficiency
- Structured for multi-region power gating
- Floorplan-aware, optimized display to memory routing

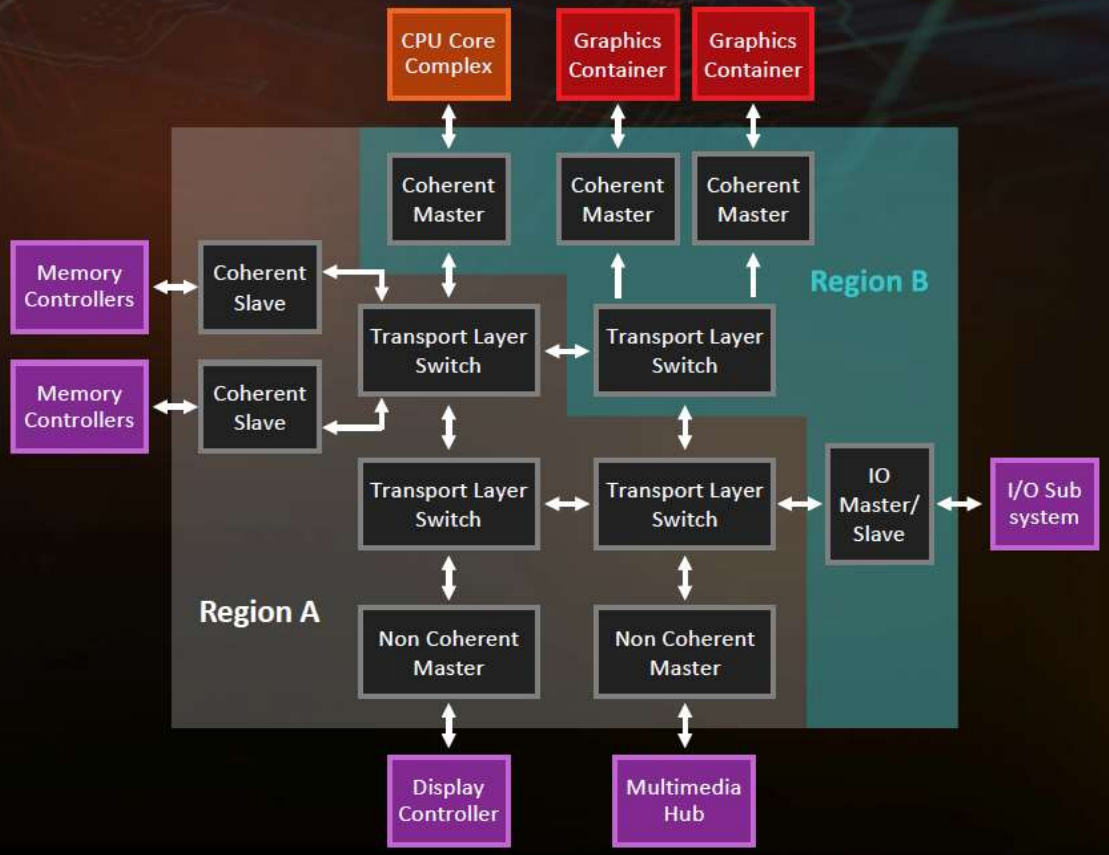

Delivering A New Level Of Visual Performance In An Soc Amd "Raven Ridge" Apu
Dan Bouvier, Jim Gibney, Alex Branover, Sonuarora

Die photograph of a Ryzen Mobile die [37]

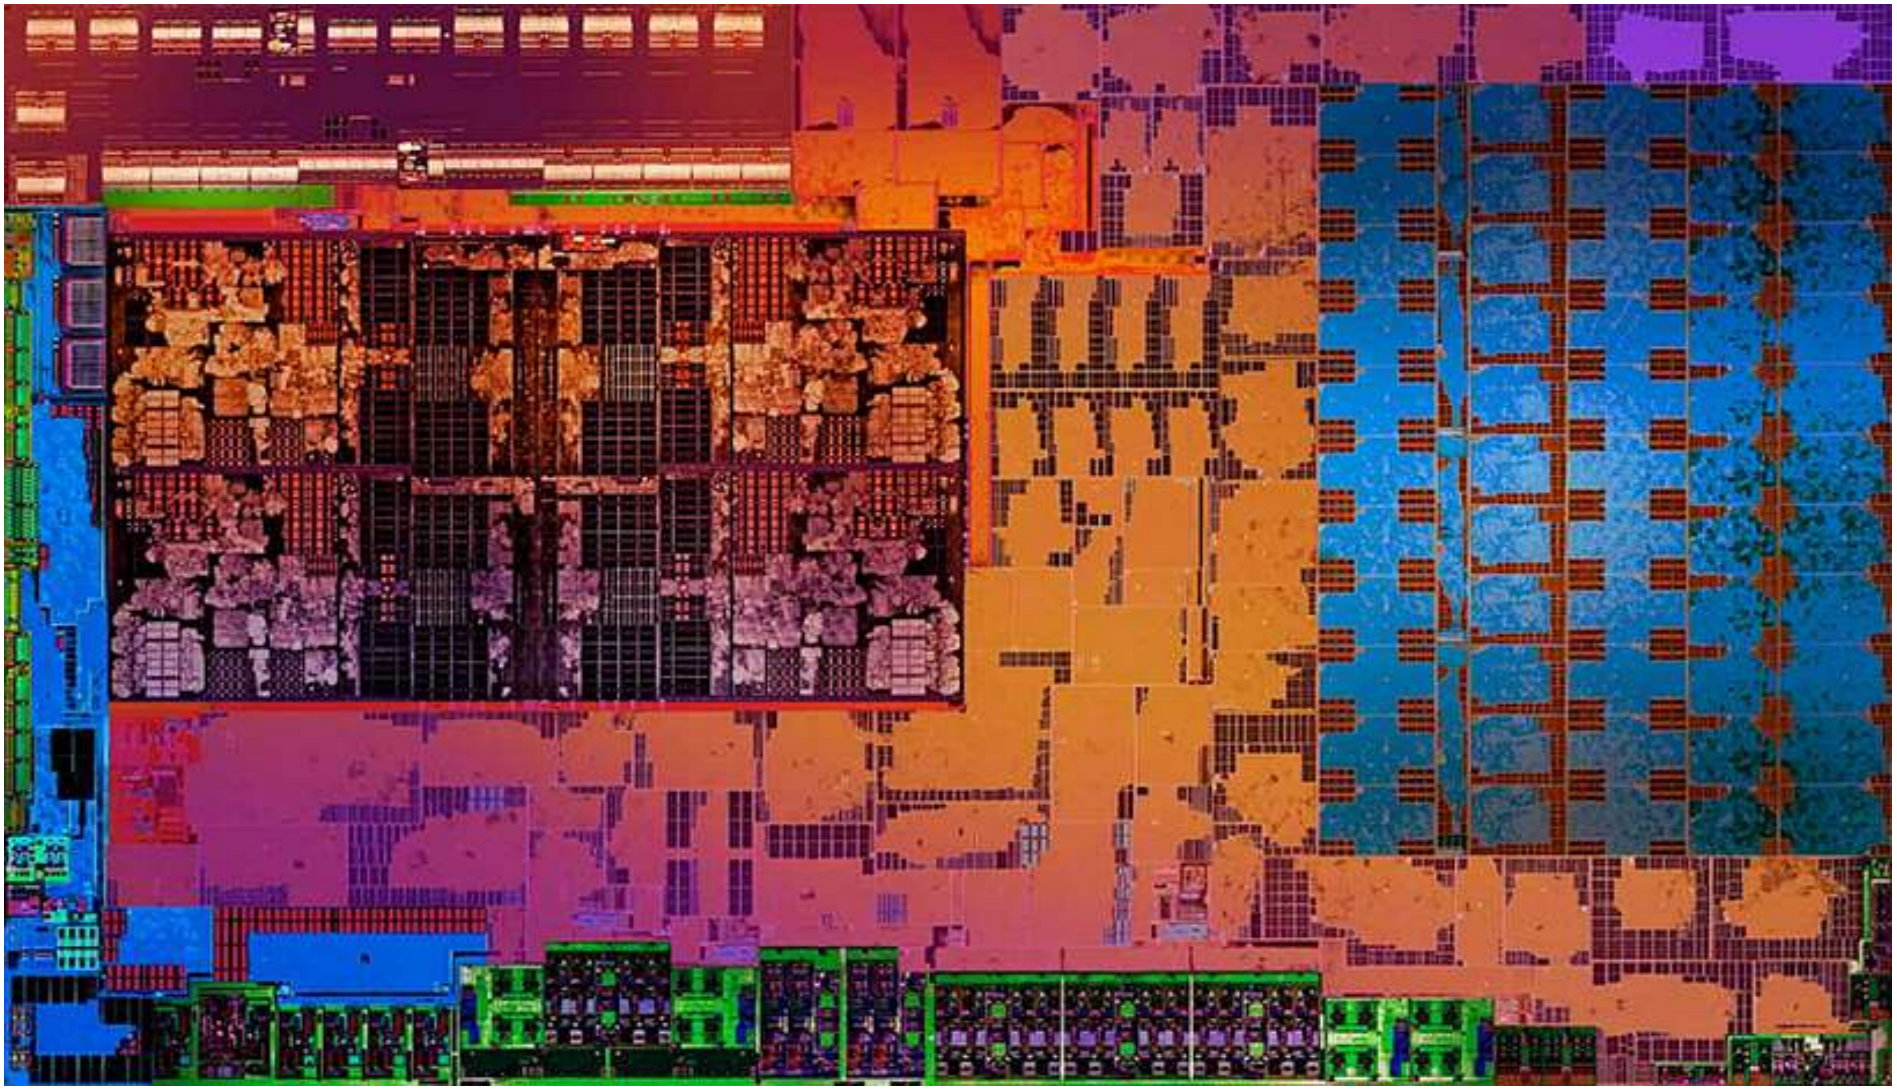

4.94 billion transistors, 210 mm^2

Designation of AMD's Ryzen and Ryzen Mobile models [48]

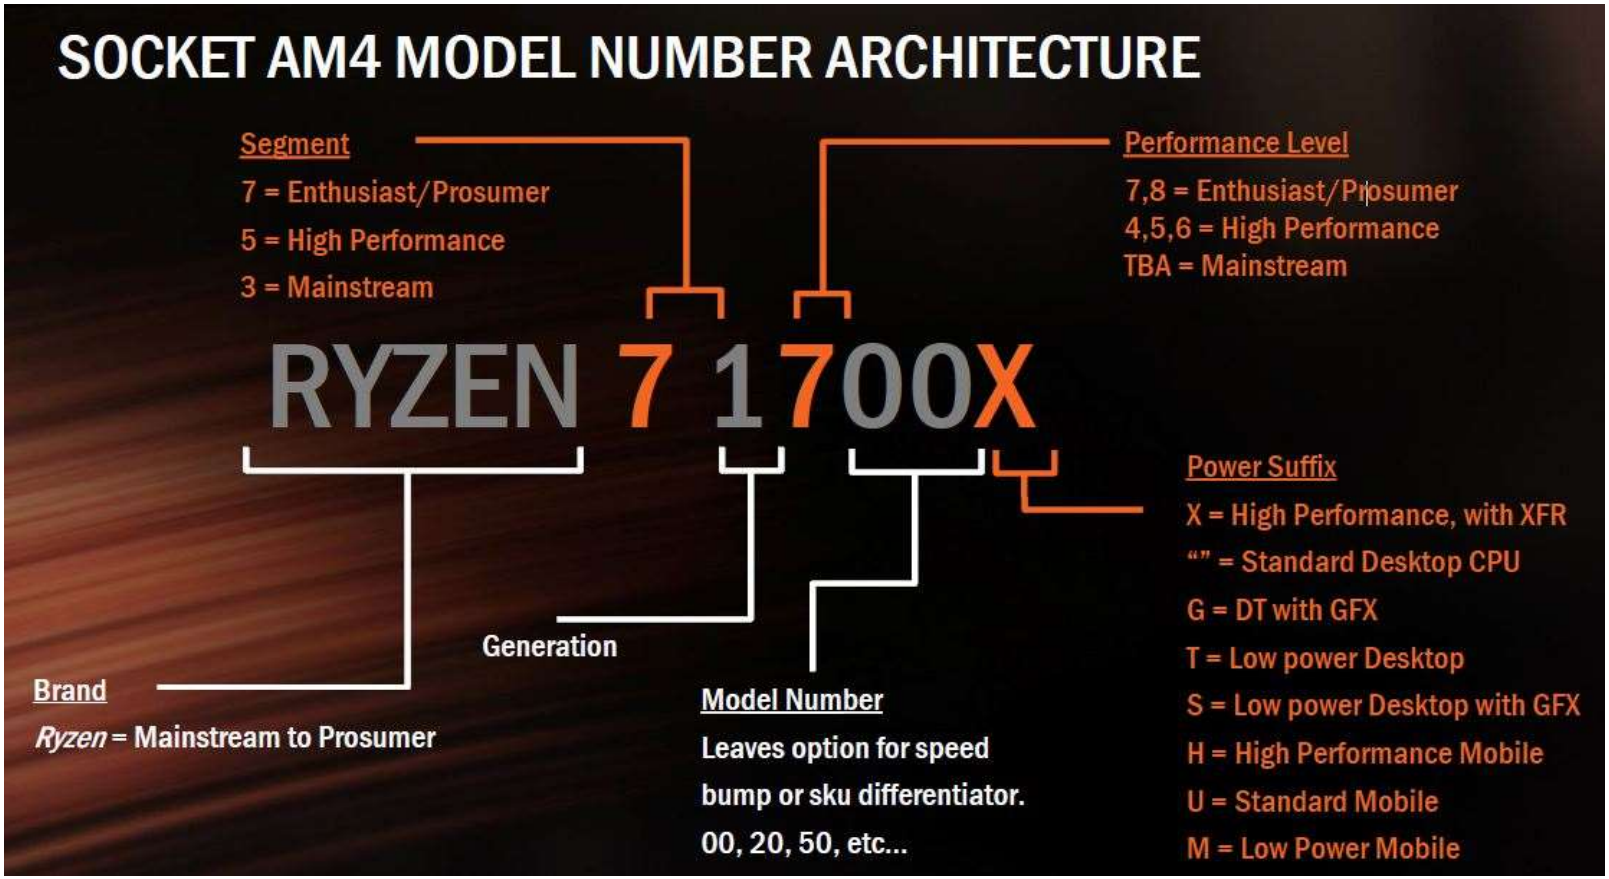

- AMD designates their mobile line as 2. generation models, model numbering begins with 2 rather than 1) .
- Mobile models provide improvements first of all in the power management, as discussed in this Section.

Prosumer: Between professional and consumer

SKU: Stock Keeping Unit

Main features of AMD's Ryzen Mobile line introduced in 2017-1

AMD's Zen-based processor lines introduced in 2017

	Ryzen Mobile APU (Raven Ridge)	**Ryzen DT** (Summit Ridge)	**ThreadRipper** (Whitehaven)	**Epyc** (Naples)
Market segment	Mobile	Desktop platform	HED	1S/2S server
µarch./Technology	Zen 14 nm	Zen, 14 nm	Zen, 14 nm	Zen, 14 nm
Launched models	Ryzen 7 2700U Ryzen 5 2500U (10/2017)	Ryzen 7 (3/2017) Ryzen 5 (4/2017) Ryzen 3 (7/2017)	1950X/1920X/1900X (8/2017)	Series 7000 (6/2017)
Layout	CCX + Vega 8/10	Zeppelin die with 2x CCX	MCM (2x Zeppelin die)	MCM (4x Zeppelin die)
Integrated GPU	Yes	No	No	No
Core count	4	4/6/8	8/12/16	8/16/24/32
SMT	SMT	SMT (except Ryzen 3)	SMT	SMT
Mem. channels/rate	2xDDR4-2400	2xDDR4-2666	4xDDR4-2666	8xDDR4-2666
PCIe 3.0 lanes	??	16xPCIe 3.0	60xPCIe 3.0	128 for 1S servers 64 for 2S servers
TDP	15 W	65/95 W	180 W	120/170/180 W
Socket	AM4 (1331)	AM4 (1331)	TR4 (SP3r2) (4094)	SP3 (4094)
Chipset	SoC	300-series	X399	No chipset, SOC

Main features of AMD's first introduced Ryzen Mobile models -2 [38]

AMD Ryzen Mobile APUs (Raven Ridge)			
	Ryzen 7 2700U with Vega 10	Ryzen 5 2500U with Vega 8	FX-9800P (Bristol Ridge)
CPU	Quad-Core with SMT 2.2 GHz Base 3.8 GHz Turbo Zen Cores, 14nm	Quad-Core with SMT 2.0 GHz Base 3.8 GHz Turbo Zen Cores, 14nm	Dual Module 2.7 GHz Base 3.6 GHz Turbo Excavator, 28nm
GPU	Vega 10 10 CUs (640 SPs) Up to 1300 MHz	Vega 8 8 CUs (512 SPs) Up to 1100 MHz	GCN 1.2 8 CUs (512 SPs) Up to 758 MHz
TDP	15W	15W	15W
DRAM	Up to DDR4-2400	Up to DDR4-2400	Up to DDR4-1866
L2 Cache	512 KB/core	512 KB/core	1 MB/module
L3 Cache	1 MB/core	1 MB/core	-
PCIe Lanes	?	?	8 x PCIe 3.0
Die Size	209.78 mm2	209.78 mm2	250.4 mm2
Transistors	4.95 billion	4.95 billion	3.1 billion
Launch	October 2017	October 2017	May 2016

Power management enhancements vs. the Ryzen line

a) Integrated voltage and frequency management
b) Precision Boost 2 technology (actually an Enhanced Turbo Boost technology)
c) Skin Temperature Aware Power Management (STAPM)
d) Mobile XFR (eXtended Frequency Range) (actually Configurable TDP).

*

a) Integrated voltage and frequency management

- This means integrated, fine-grained voltage and frequency management for the CPU cores and the GPU (i.e. power is allocated to the units that demand it).
- Obviously, it needs per core and per GPU frequency and voltage regulation.

Implementing per core and per GPU voltage regulation [38]

Two stage voltage regulation (VR):

- 1. stage: on the mainboard,
- 2. stage: separate on-die VRs
  - for each CPU core and
  - the graphics.
- VRs are implemented as LDOs (Digital Low-Dropout Regulators)
- The LDOs serve also as power gates.
- Power reduction: ~ 35%.

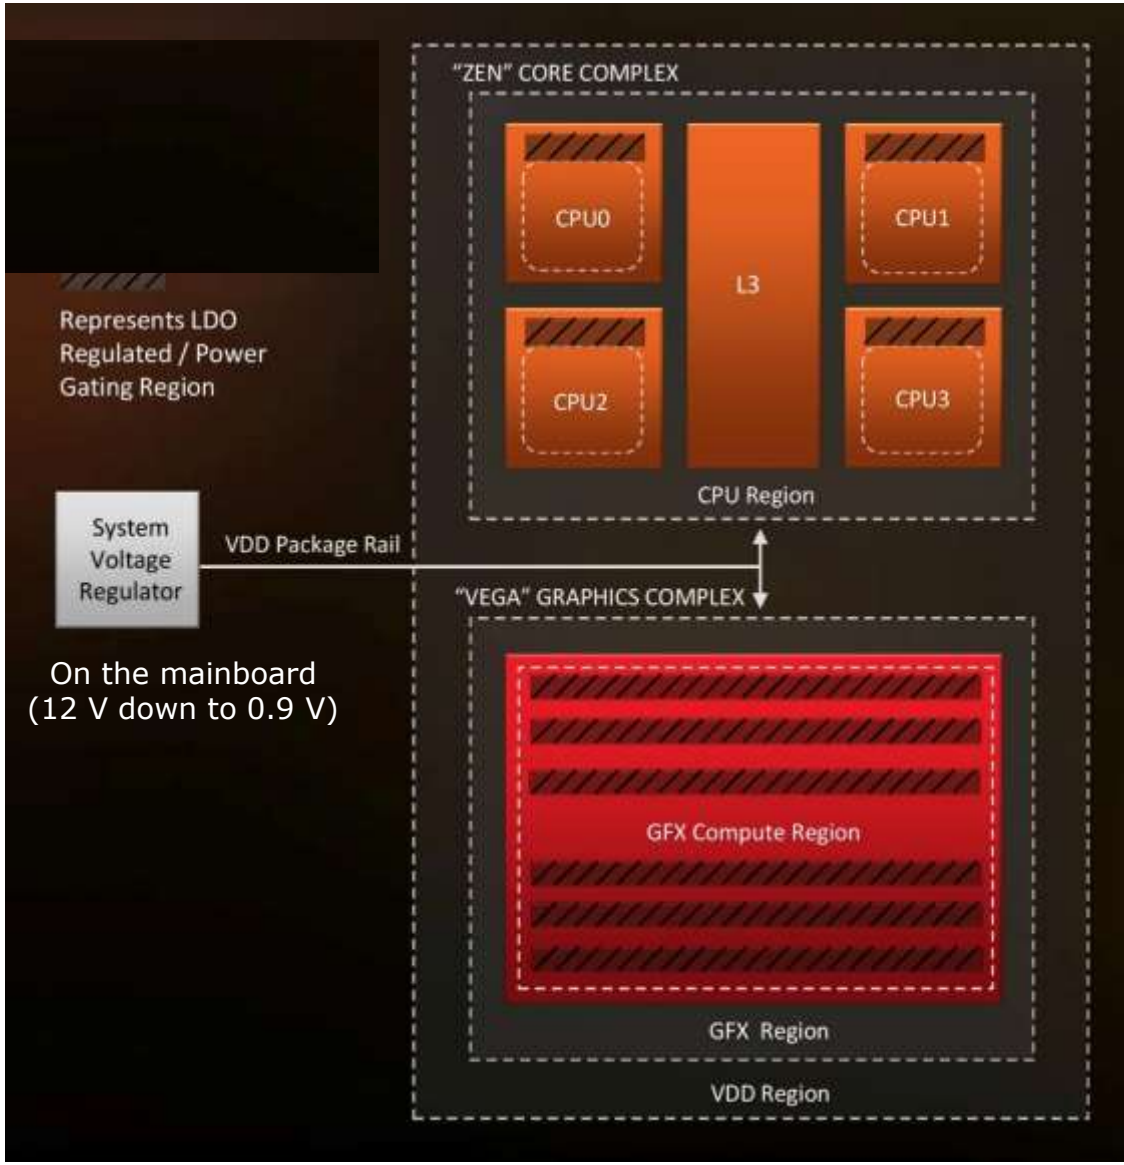

Principle of implementing integrated voltage and frequency management [38]

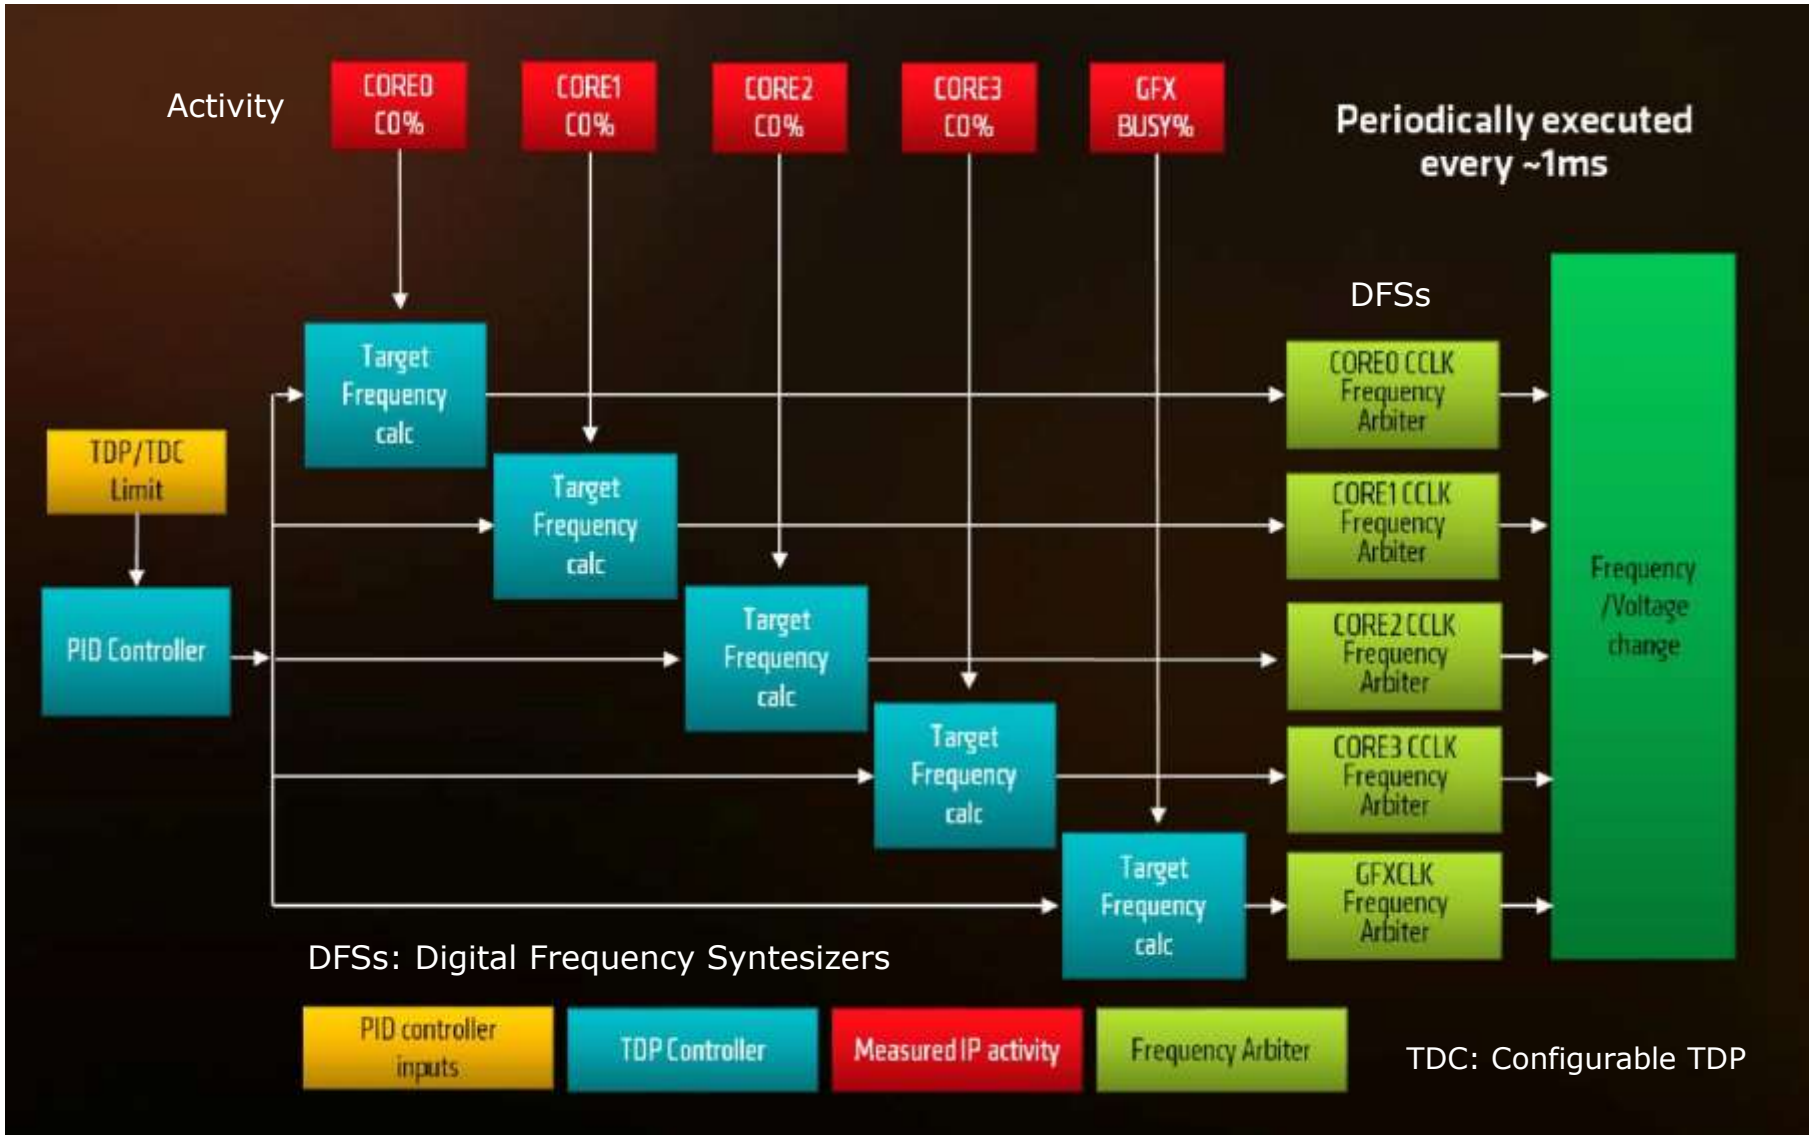

PID Controller: A proportional–integral–derivative *controller*

*

Example for integrated per-core frequency and voltage management [38]

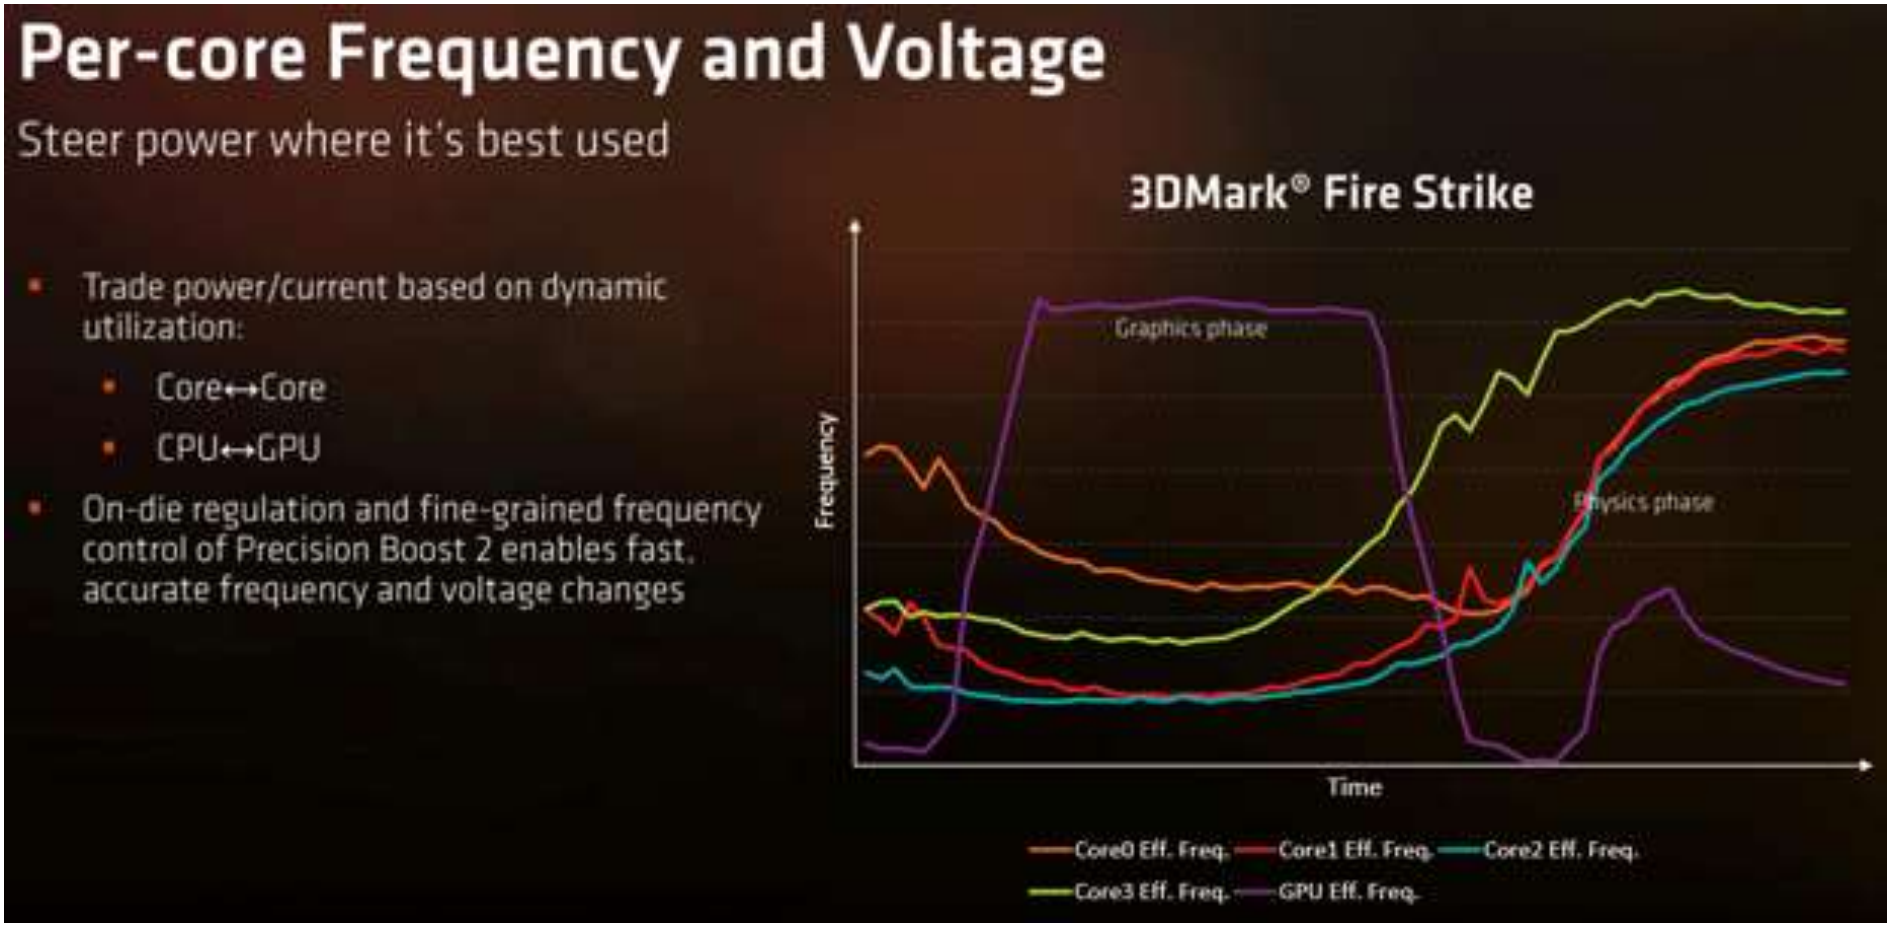

b) Enhanced Turbo Boost technology (termed Precision Boost 2 technology) [38]

Boost behavior in the Ryzen line

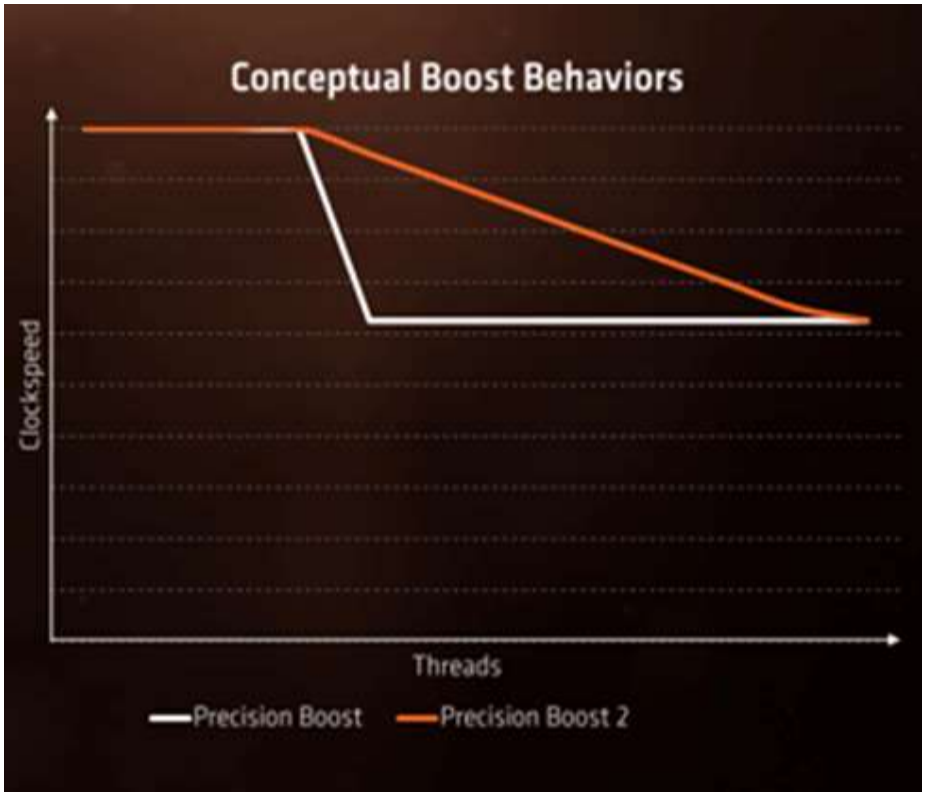

- There exist two Turbo Boost frequencies, one if up to 2 cores are active the other if more cores.
- 25 MHz granularity.

Boost behavior in the Ryzen Mobile line

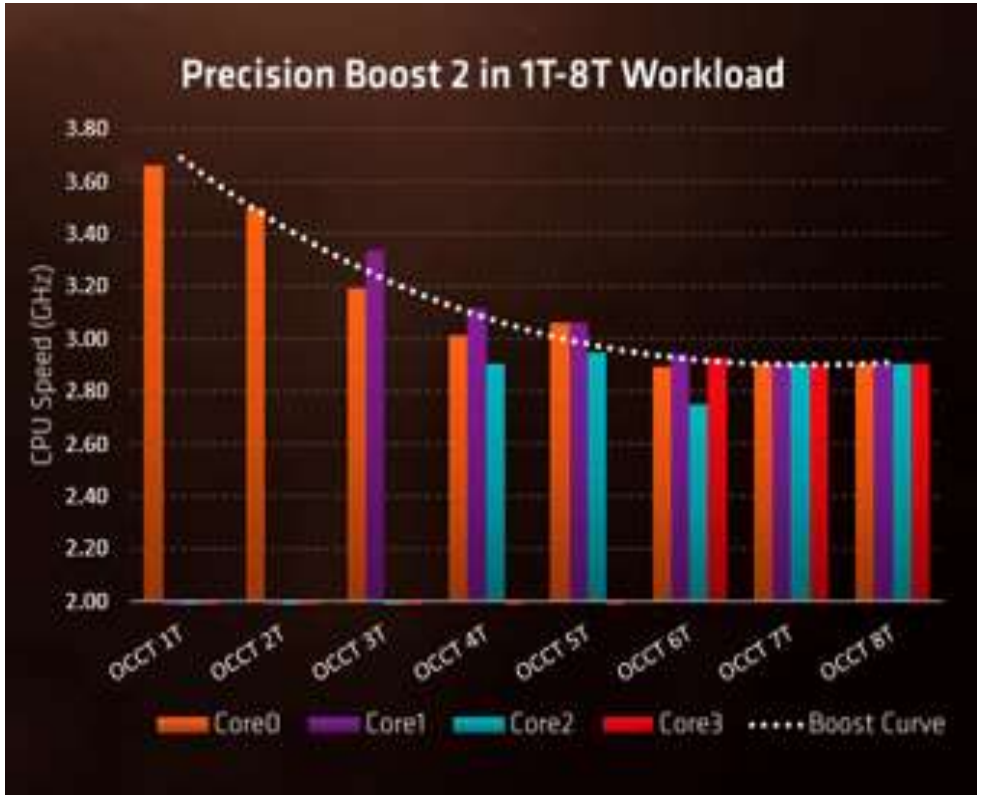

- The Turbo Boost frequency gracefully rolls out depending on the number of active threads, CPU temperature and currents.
- 25 MHz granularity.

OCCT: OverClock Checking Tool

*

c) Skin Temperature Aware Power Management (STAPM)

Before STAPM:

APU guard-banded to Tj~60C to meet Tskin requirements

After STAPM:

Delta between ambient and Tskin calculated based on the power/activity system components

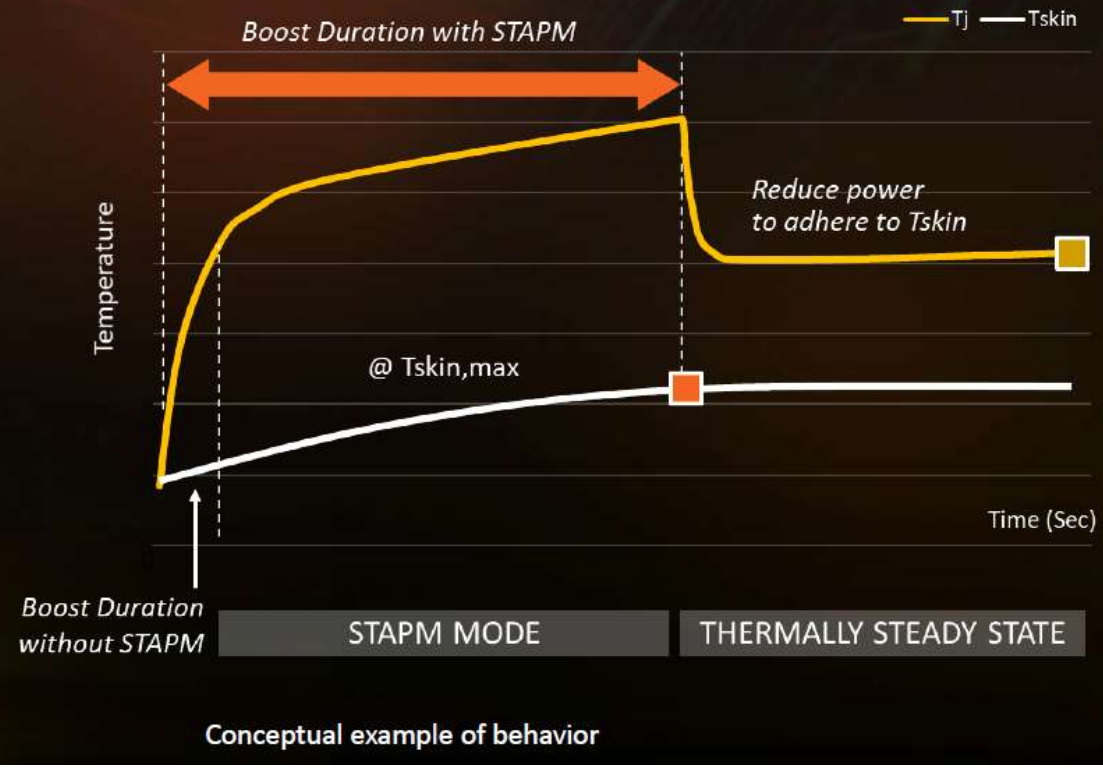

Conceptual example of behavior

d) Mobile XFR (eXtended Frequency Range) [38]

- Assuming an enhanced cooling system, the sustained clock frequency can be raised by the Mobile XFR (eXtended Frequency Range) technology within the configurable TDP (TDC) limits.
- This results in higher performance, as indicated in the Figure below.

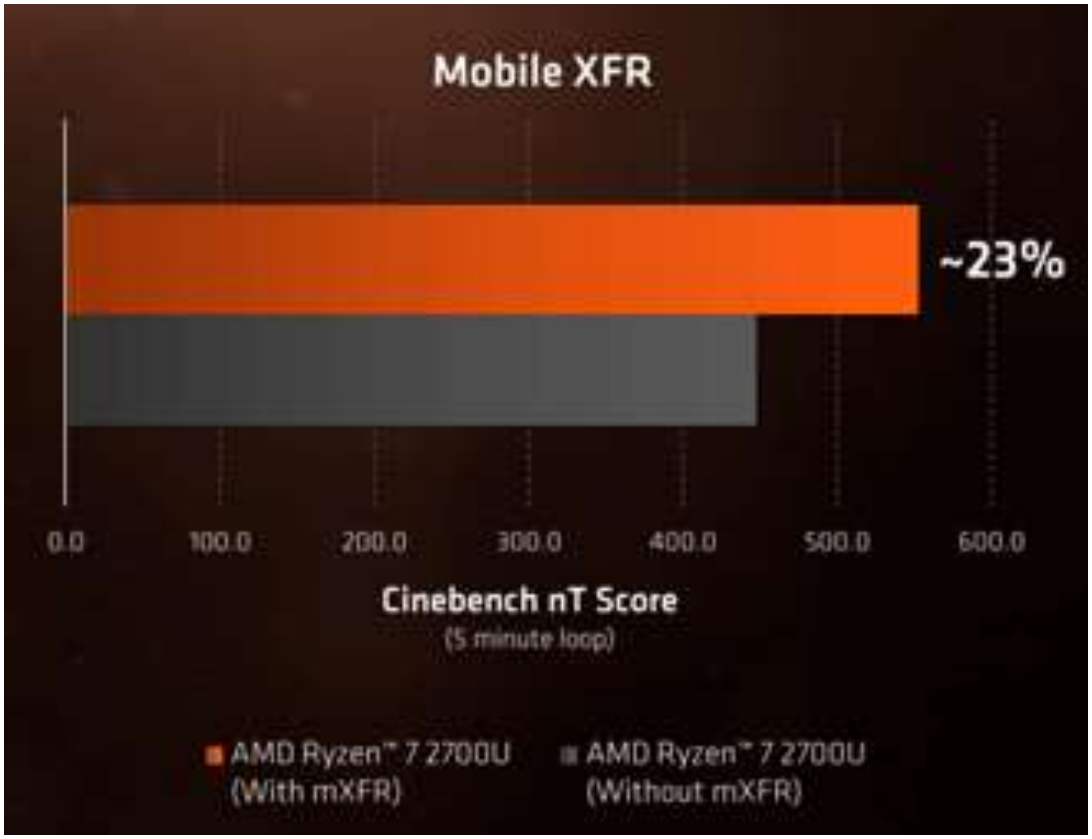

Figure: Raising performance by the Mobile XFR technology [38]

*

1T/nT performance comparison with Intel processors [72]

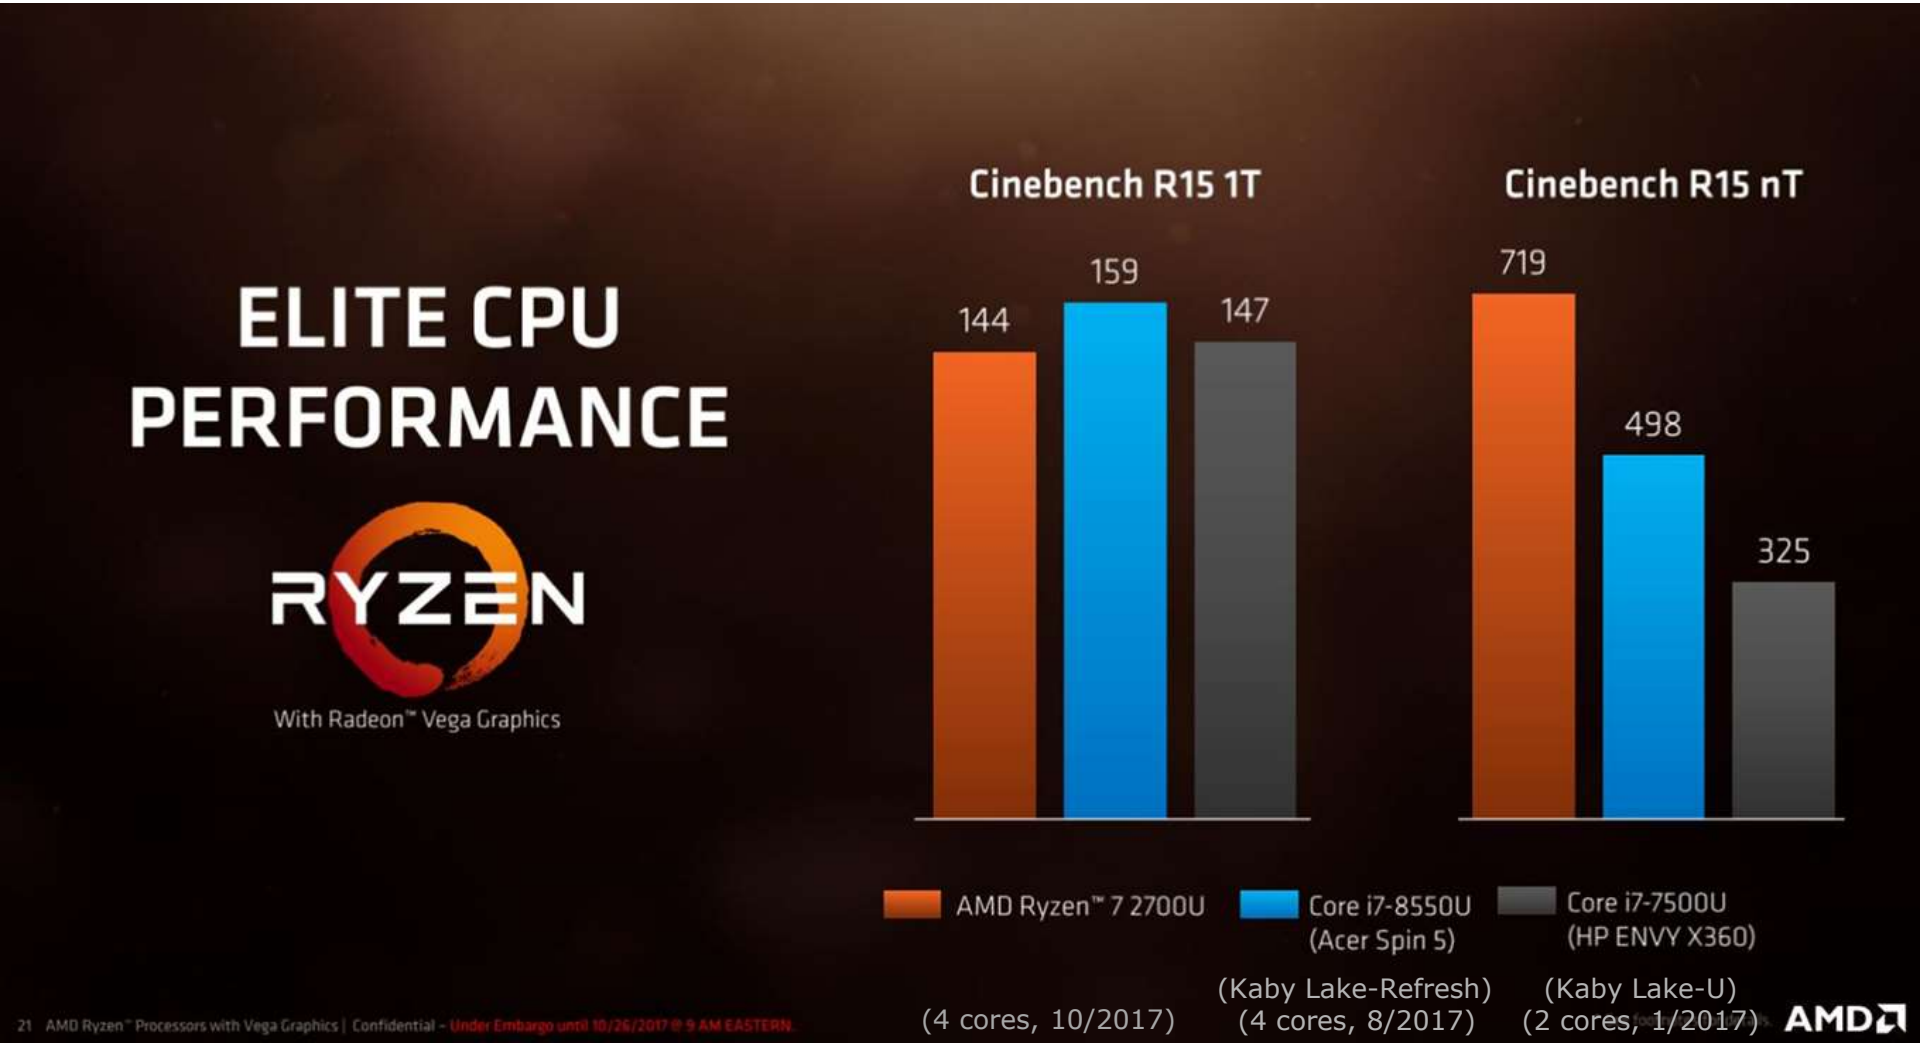

Cinebench: Cross platform CPU and graphics benchmark.

Remark

- On 06. Nov. 2017 Intel announced that it has been developing an 8. generation H-series processor that makes use of AMD's Radeon Vega GPU along with a high-bandwidth HBM2 stacked memory integrated in the same processor package, as indicated in the Figure below.
- The GPU and the HBM2 memory is interconnected by means of Intel's EMIB (Embedded Multi-die Interconnect Bridge) technology (see the next Figure).

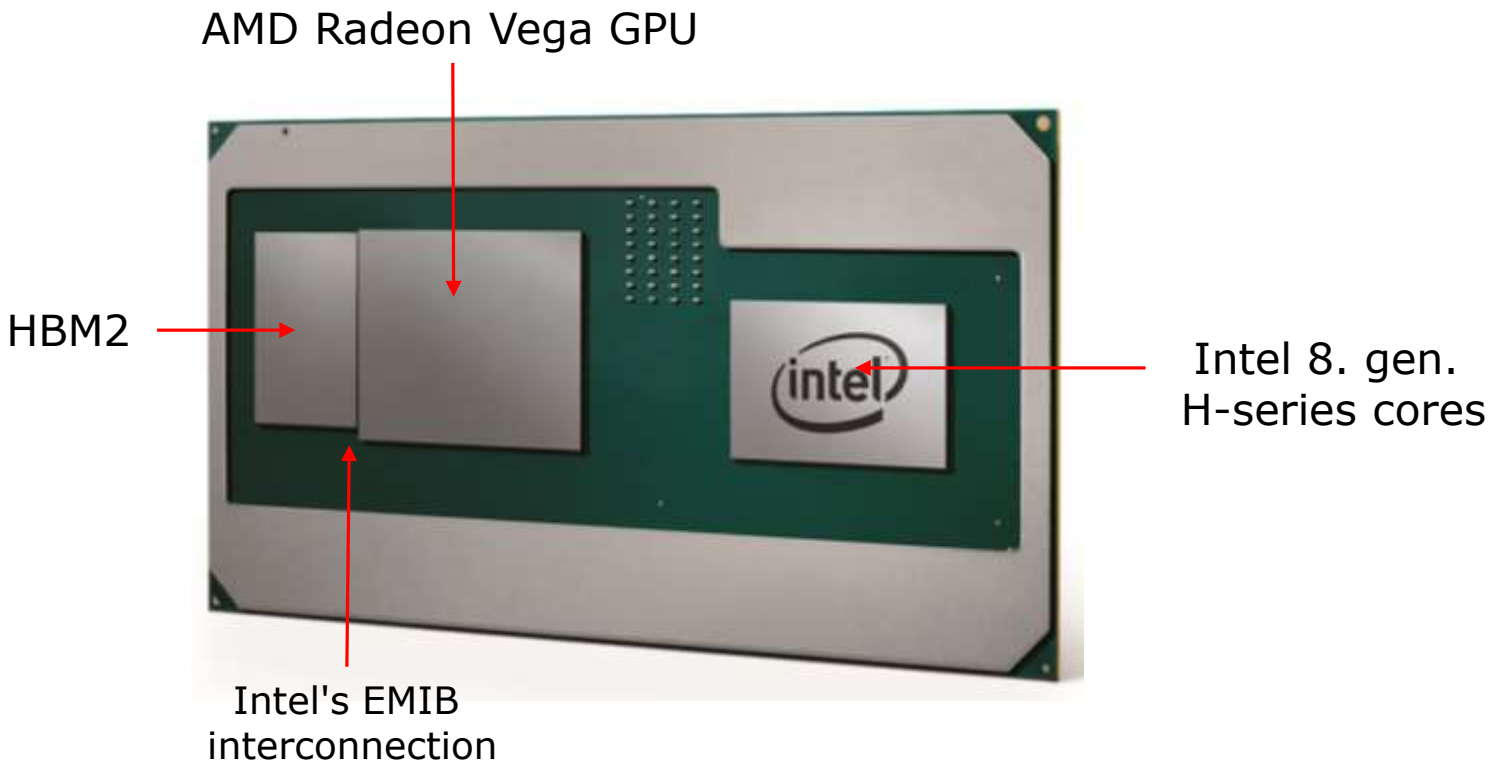

Figure: Intel's planned 8. gen. H-series processor [52]

Heterogeneous integration options [69]

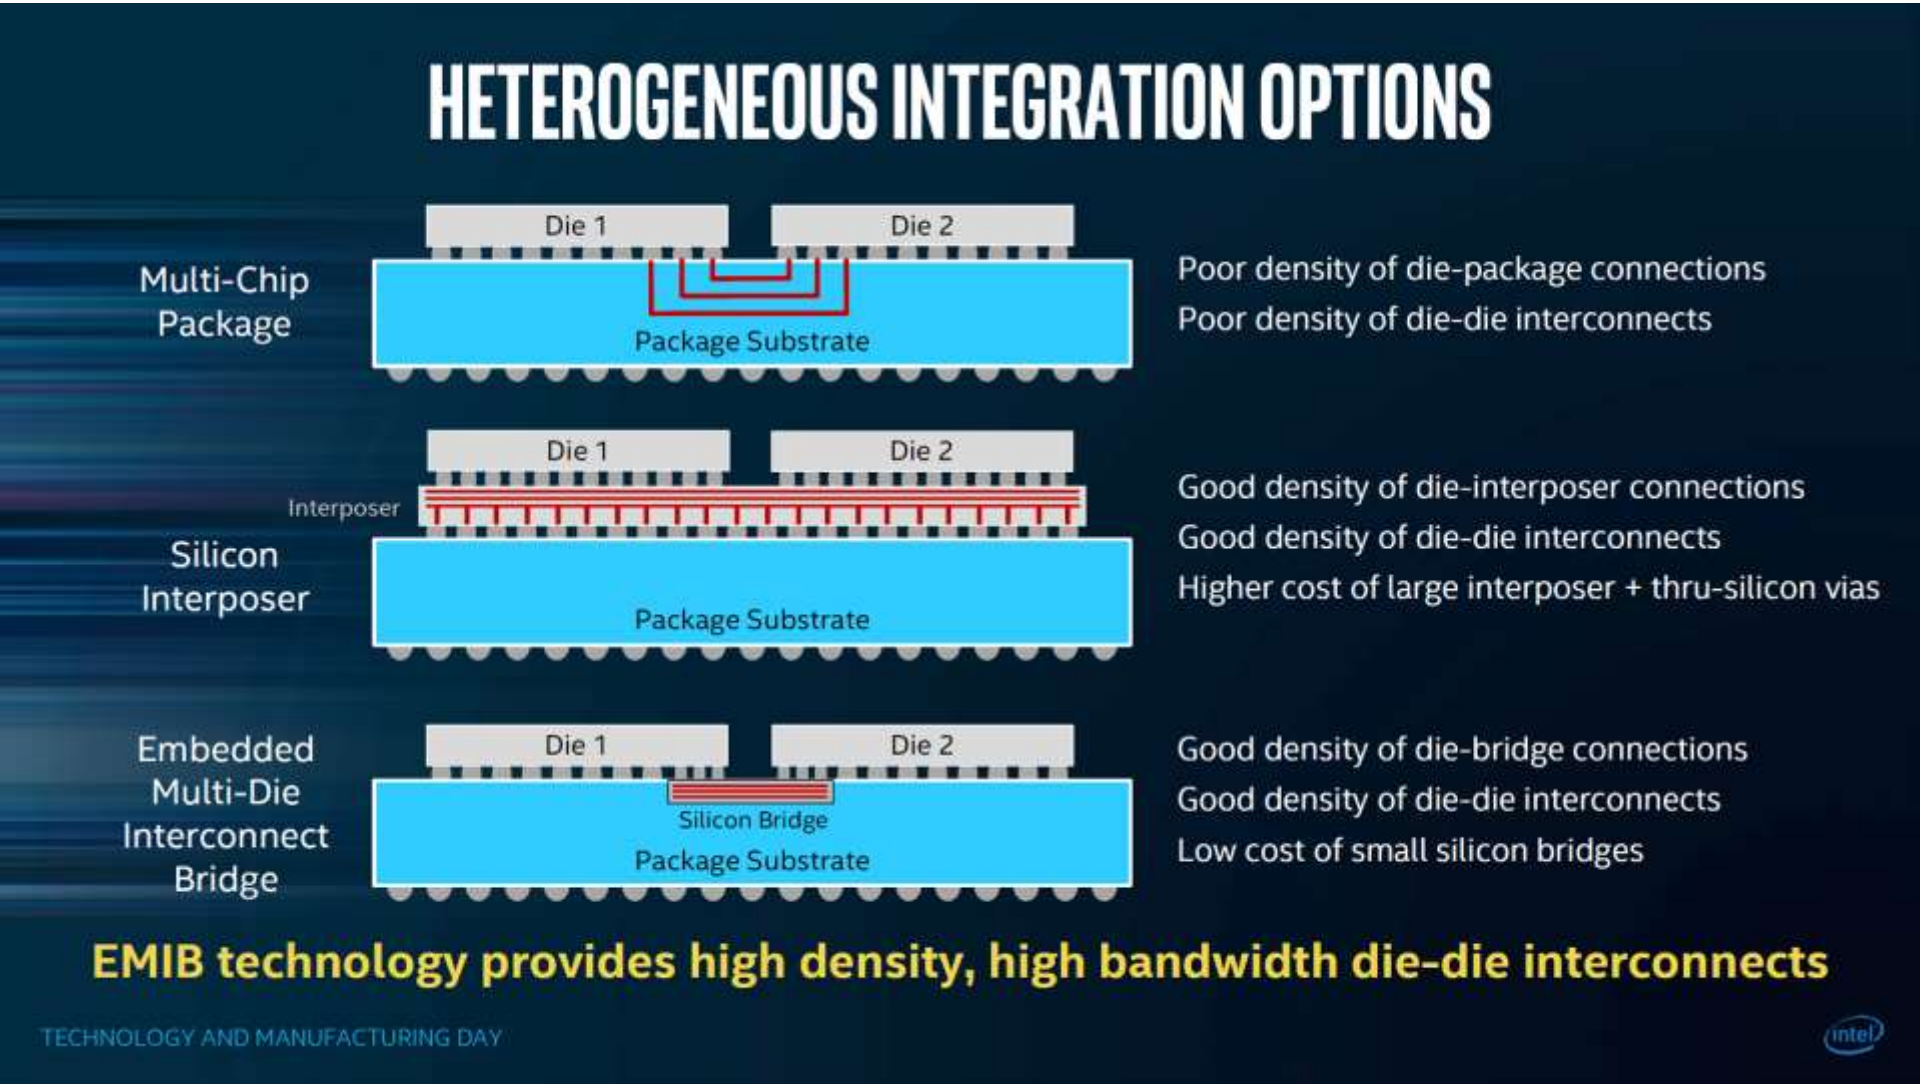

Implementation of EMIB [62]

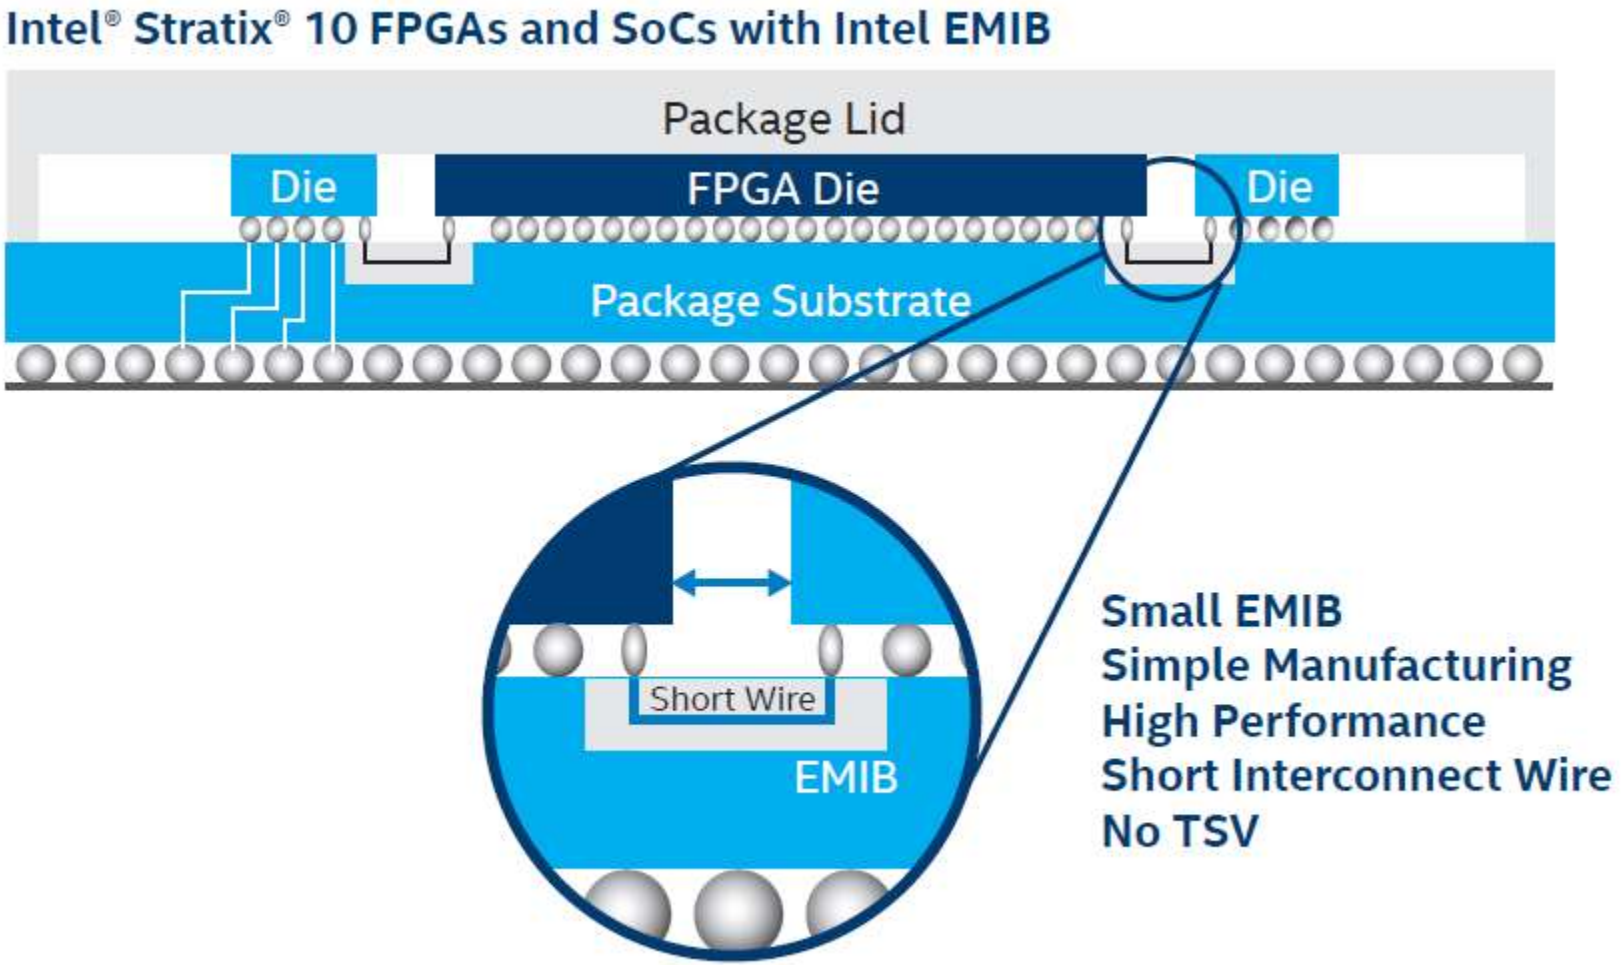

Intel's concept for interconnecting multiple dies through EMIB [57]

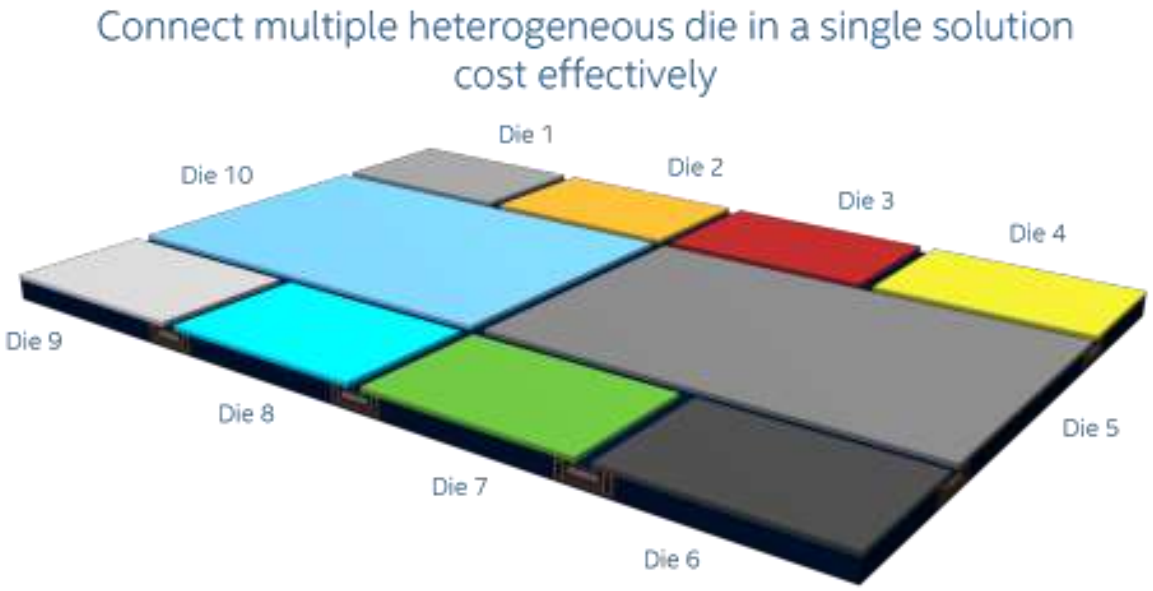

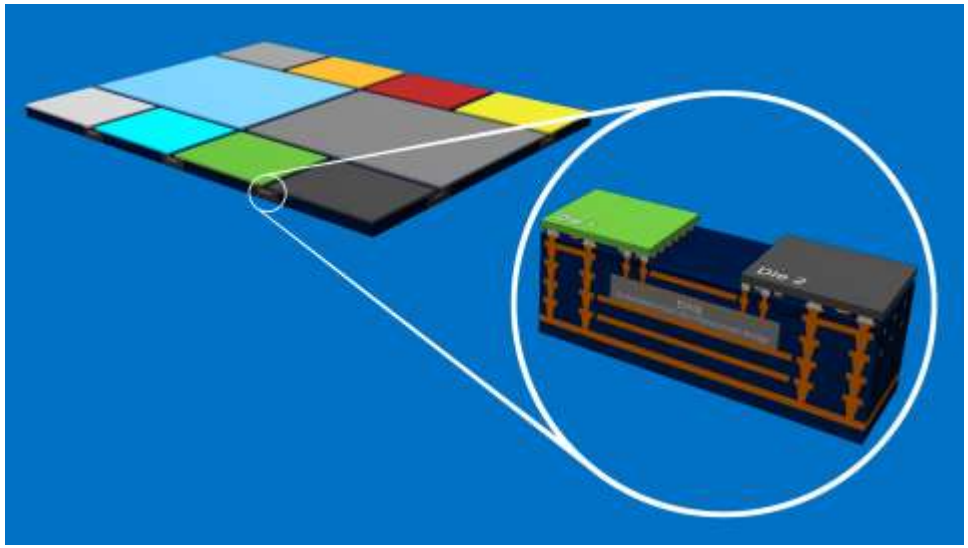

Remark

- In 01/2018 AMD extended its Ryzen mobile line with Ryzen 3 models, as follows:
  - Ryzen 3 2300U and
  - Ryzen 3 2200U
- Main features of these models are shown in the next Table.

Main features of AMD's 2. wave mobile processors [84]

Model	Launched	CPU Cores	Threads	Max Clock (GHz)	Graphics Compute Units[iii]	Max GPU clock (MHz)	L2/L3 cache (MB)	cTDP (Watts)
Ryzen™ 7 2700U with Radeon™ RX Vega 10 Graphics	10/2017	4	8	3.8	10	1300	2/4	15W Nominal
Ryzen™ 5 2500U with Radeon™ Vega 8 Graphics	10/2017	4	8	3.6	8	1100	2/4	15W Nominal
Ryzen™ 3 2300U with Radeon™ Vega 6 Graphics	1/2018	4	4	3.4	6	1100	2/4	15W Nominal
Ryzen™ 3 2200U with Radeon™ Vega 3 Graphics	1/2018	2	4	3.4	3	1000	1/4	15W Nominal

7.2 The 2. gen. Ryzen mobile line (Picasso)

Will not be discussed

7.2 The 2. gen. Ryzen mobile line (Picasso) [64] -1

7.2 The 2. gen. Ryzen mobile line (Picasso) [64] -2

- Launched in 1/2019
- Designated as the Picasso line.
- Introduced models:
  - Ryzen 7/5/3 3x00U models (15 W TDP)
  - Ryzen 7/5 3x50H models (35 W TDP)
- The new models are upgraded versions of the first generation, 14 nm parts, they are based on the Zen+ architecture and are fabricated on 12 nm technology by GlobalFoundries.

7.2 The 2. gen. Ryzen mobile line (Picasso) (3)

Performance and power advantages of the 12 nm LP scaling vs. the 14 nm LP [99]

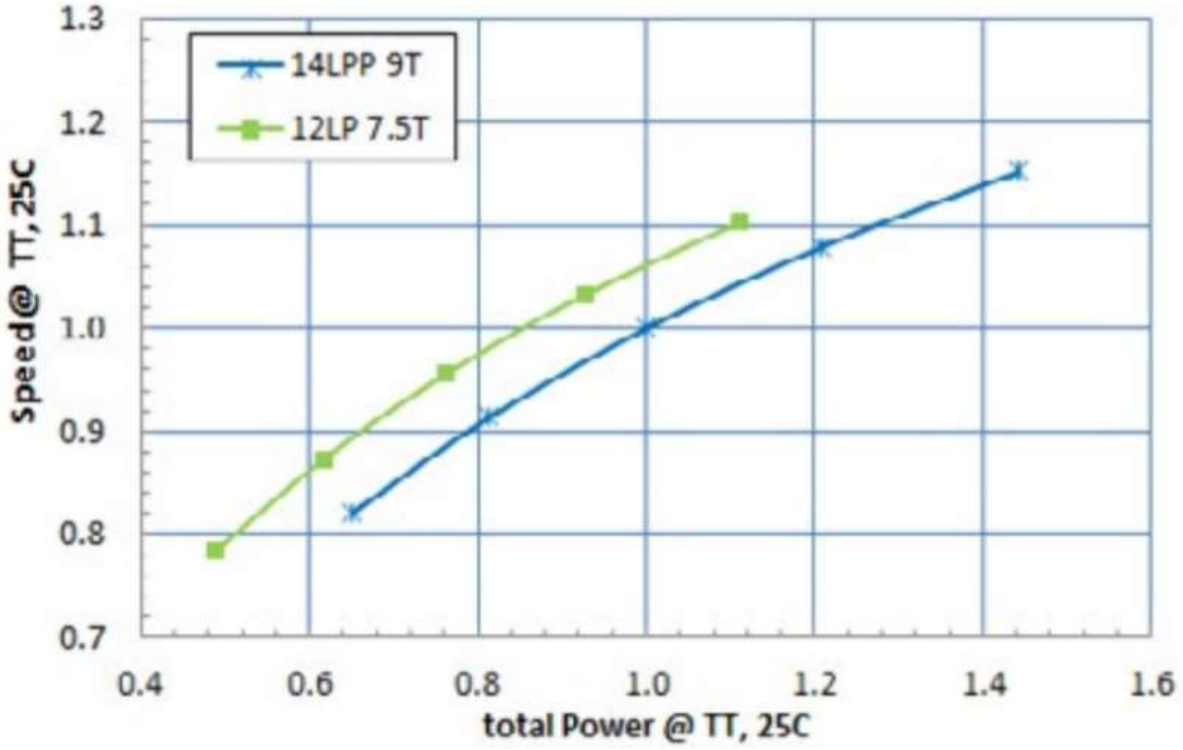

12 nm LP technology provides vs. the 14 nm LP process about

6% performance improvement at iso power

15% power reduction at iso frequency

7.2 The 2. gen. Ryzen mobile line (Picasso) (4)

Main features of AMD's 2. gen. Ryzen mobile line vs. the 1. gen. [100]

Model	Cores/ Threads	TDP	Base / Boost Freq. (GHz)	Graphics	Graphics Base / Boost (MHz)	L3 Cache	Memory	Node (GloFo)	Core
Ryzen 7 3750H	**4 / 8**	**35W**	**2.3 / 4.0**	**Radeon Vega 10**	**Up to 1,400**	**6MB**	**DDR4-2400**	**12nm**	**Zen+**
Ryzen 7 3700U	**4 / 8**	**15W**	**2.3 / 4.0**	**Radeon Vega 10**	**Up to 1,400**	**6MB**	**DDR4-2400**	**12nm**	**Zen+**
Ryzen 7 2700U	4 / 8	12 - 25W (15W)	2.2 / 3.8	Radeon Vega 10	Up to 1,300	4MB	DDR4-2400	14nm	Zen
Ryzen 5 3550H	**4 / 8**	**35W**	**2.1 / 3.7**	**Radeon Vega 8**	**Up to 1,200**	**6MB**	**DDR4-2400**	**12nm**	**Zen+**
Ryzen 5 3500U	**4 / 8**	**15W**	**2.1 / 3.7**	**Radeon Vega 8**	**Up to 1,200**	**6MB**	**DDR-2400**	**12nm**	**Zen+**
Ryzen 5 2500U	4 / 8	12 - 25W (15W)	2.0 / 3.6	Radeon Vega 8	Up to 1,100	4MB	DDR4-2400	14nm	Zen
Ryzen 3 3300U	**4 / 4**	**15W**	**2.1 / 3.5**	**Radeon Vega 8**	**Up to 1,200**	**6MB**	**DDR4-2400**	**12nm**	**Zen+**
Ryzen 3 2300U	4 / 4	15W	2 / 3.4	Radeon Vega 6	Up to 1,100	4MB	DDR4-2400	14nm	Zen
Ryzen 3 3200U	**2 / 4**	**15W**	**2.6 / 3.5**	**Radeon Vega 3**	**Up to 1,200**	**5MB**	**DDR4-2400**	**12nm**	**Zen+**
Ryzen 3 2200U	2 / 4	15W	2.5 / 3.4	Radeon Vega 3	Up to 1,100 MHz	4MB	DDR4-2400	14nm	Zen
AMD Athlon 300U	**2 / 4**	**15W**	**2.4 / 3.3**	**Radeon Vega 3**	**Up to 1,000**	**5MB**	**DDR4-2400**	**12nm**	**Zen+**

7.2 The 2. gen. Ryzen mobile line (Picasso) (5)

Main improvements of the released models compared to previous related ones [100]

a) +100 MHz base frequency and +100 or +200 MHz Turbo frequency boost.

b) Emerging two 35 W TDP parts designated as H-series models (Ryzen 7 3750H and Ryzen 5 3550H).

Aim of the H-series Picasso models with a TDP value of 35 W

- H-series models have the same technical parameters, except their TDP value.
- Higher TDP values allows the processor to boost more frequently and for longer periods of time.
- OEMs are expected to design more mainstream gaming laptops with discrete graphics cards that are based on H-series processors with aTDP of 35 W, since then the power delivery system can allocate more of the available power budget to the CPU cores instead of the Vega graphics to achieve higher performance.

7.2 The 2. gen. Ryzen mobile line (Picasso) (6)

Remark []

- With it's A-series processors AMD entered the Chromebook market with two new, A-series low power processors.
- These 28 nm processors are built on the Excavator based Carrizo processor and provide up to ten hour battery life.
- These processors challenge Intel's Pentium and Celeron processors that dominate the growing Chromebook market.

Models	Cores / Threads	TDP (typ.)	Base / Boost GHz	Radeon Graphics	GPU Cores	GPU Frequency
AMD A6-9200C	2 / 2	6W	1.8 / 2.7	R5 Series	3 / 192 (GCN 1.2)	720 MHz
AMD A4-9120C	2 / 2	6W	1.6 / 2.4	R4 Series	3 / 192 (GCN 1.2)	600 MHz

Table: AMD's Carrizo based A-series processors introduced for the Chromebook market in 1/2019 [100]

2. gen. Ryzen mobile line (Picasso) (7)

[101]

cost laptops or tablets primarily used to run Google's
h most application data are kept in the cloud rather than

rating under the Linux based Chrome OS, but all models
017 can also run Android apps.

entered the market in 2011, since then their market share rose continuously.

- In March 2018 Chromebooks achieved a market share of up to 60 % in computers purchased by schools in the USA.

8. The ThreadRipper HED line

- 8.1 The 1. gen. ThreadRipper HED line
- 8.2 The 2. gen. Threadripper HED line

Only Section 8.2 will be discussed

8. The ThreadRipper HED line

AMD launched until now two generations of ThreadRipper processors:

- the 1. generation in 2017 (called Whitehaven) and
- the 2. generation in 2018 (called Pinnacle Ridge)

These two generations will be discussed subsequently.

8.1 The 1. gen. ThreadRipper HED line

8.1 The 1. gen. ThreadRipper HED line

- Launched in 8/2017, designated as Whitehaven.
- HED (or HEDT) meaning High-End Desktop.
- 2 Zeppelin chips interconnected by the IF, implemented as an MCM,
  as indicated in the next slides.

Positioning AMD's Zen-based Threadripper HED line

AMD's Zen-based processor lines

Ryzen mobile (Mobil)	**Ryzen (**DT)	**ThreadRipper** (HED)	**Epyc** (1S/2S server)
Single CCX + Vega GPU	Zeppelin chip (2x CCX, no GPU)	2 Zeppelin chips interconnected by the IF, implemented as an MCM	4 Zeppelin chips interconnected by the IF, implemented as an MCM

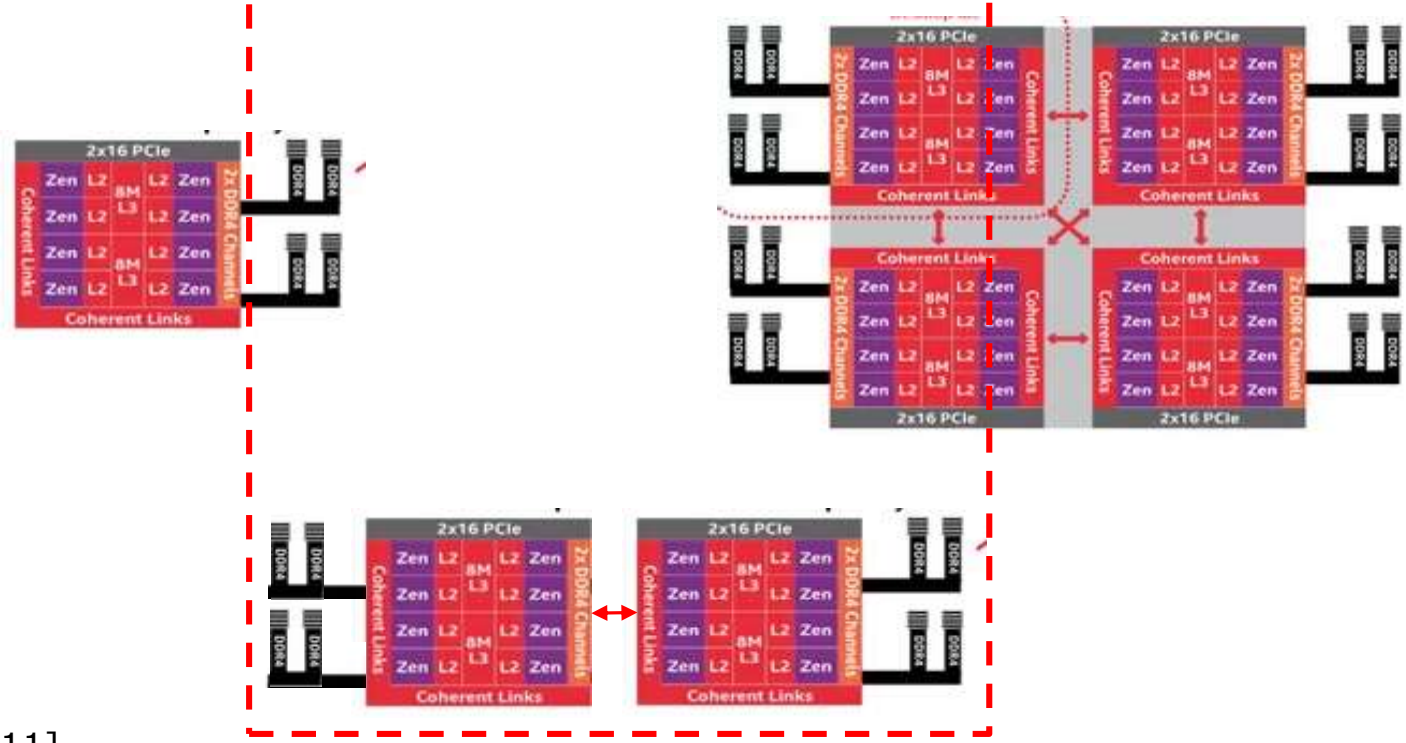

Drawings based on [11]

Block diagram of 1. gen. ThreadRipper processors [36]

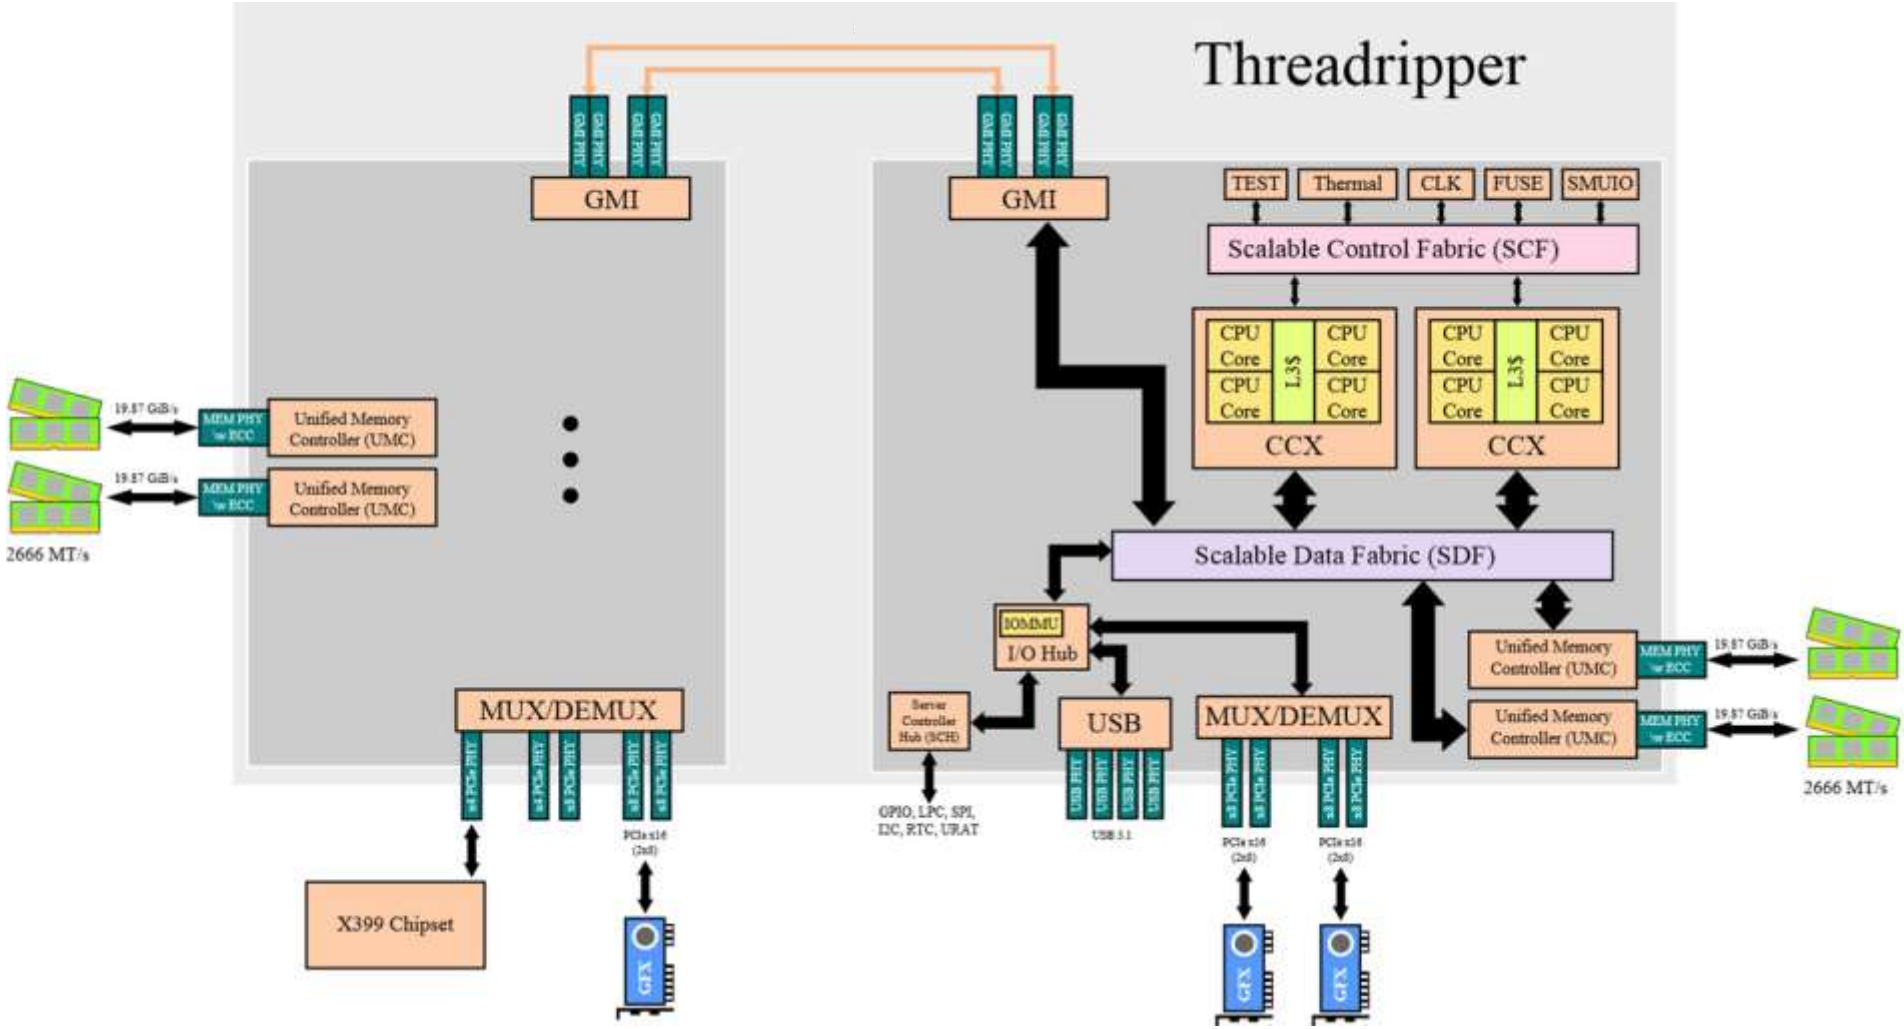

Main features of the 1. gen. ThreadRipper line introduced in 2017

AMD's Zen-based processor lines introduced in 2017

	Ryzen Mobile APU (Raven Ridge)	**Ryzen DT** (Summit Ridge)	**ThreadRipper** (Whitehaven)	**Epyc** (Naples)
Market segment	Mobile	Desktop platform	HED	1S/2S server
µarch./Technology	Zen 14 nm	Zen, 14 nm	Zen, 14 nm	Zen, 14 nm
Launched models	Ryzen 7 2700U Ryzen 5 2500U (10/2017)	Ryzen 7 (3/2017) Ryzen 5 (4/2017) Ryzen 3 (7/2017)	1950X/1920X/1900X (8/2017)	Series 7000 (6/2017)
Layout	CCX + Vega 8/10	Zeppelin die with 2x CCX	MCM (2x Zeppelin die)	MCM (4x Zeppelin die)
Integrated GPU	Yes	No	No	No
Core count	4	4/6/8	8/12/16	8/16/24/32
SMT	SMT	SMT (except Ryzen 3)	SMT	SMT
Mem. channels/rate	2xDDR4-2400	2xDDR4-2666	4xDDR4-2666	8xDDR4-2666
PCIe 3.0 lanes	??	16xPCIe 3.0	60xPCIe 3.0	128 for 1S servers 64 for 2S servers
TDP	15 W	65/95 W	180 W	120/170/180 W
Socket	AM4 (1331)	AM4 (1331)	TR4 (SP3r2) (4094)	SP3 (4094)
Chipset	SoC	300-series	X399	No chipset, SOC

Memory access times and bandwidth while using the Infinity Fabric in AMD's ThreadRripper processors [23]

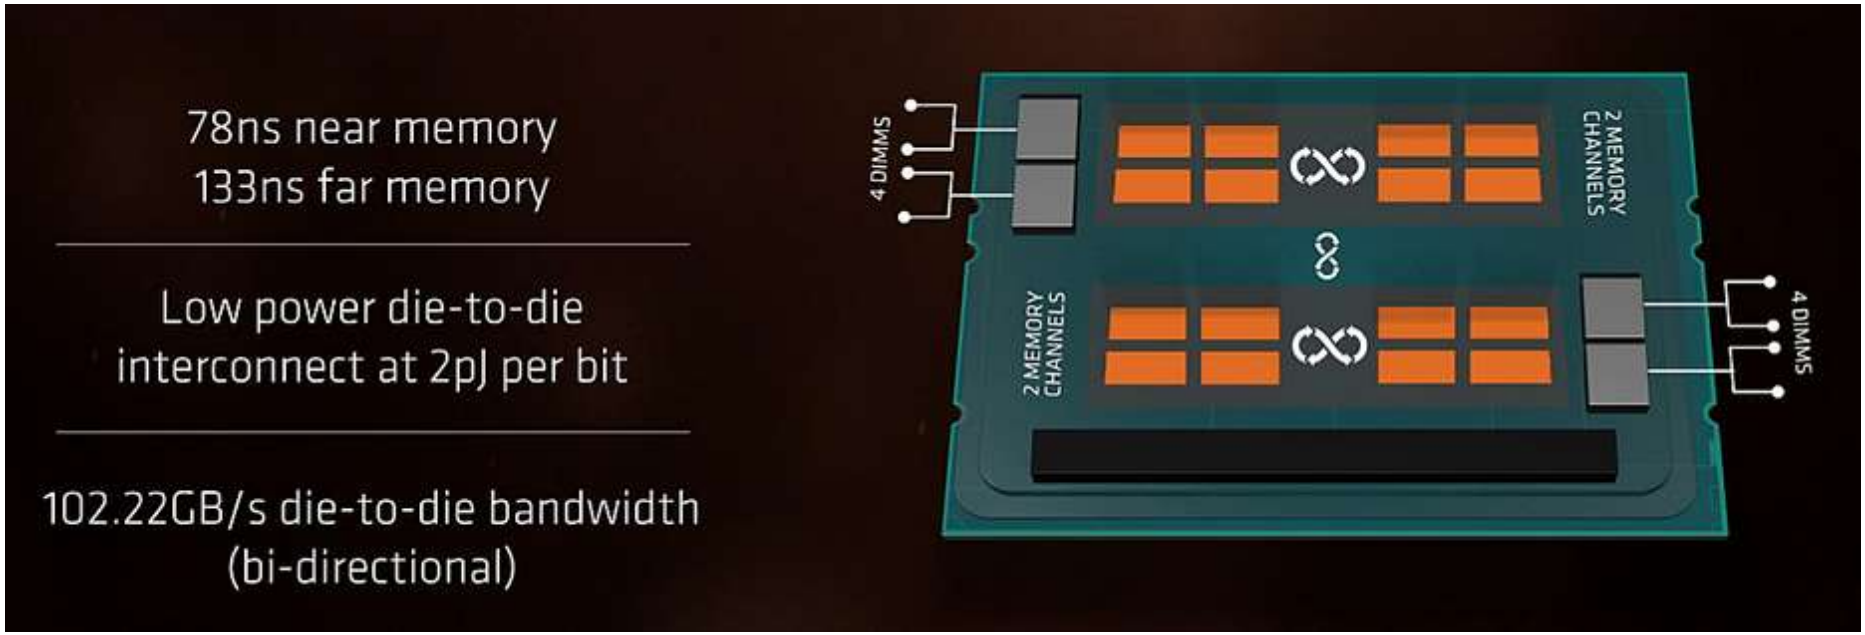

Infinity Fabric latencies in AMD's ThreadRipper and Ryzen lines [24]

Processor	Intra-Core Latency	Intra-CCX Core-to-Core Latency	Cross-CCX Core-to-Core Latency	Cross-CCX Average Latency	Die-to-Die Latency	Die-To-Die Average Latency	Average Transfer Bandwidth
TR 1950X Creator Mode DDR-2666	13.7 - 14.1	39.4 - 43.2ns	157.6 - 171.3	168ns	**180.6 - 256.7ns**	**238.47ns**	90.26 GB/s
TR 1950X Creator Mode DDR4-3200	13.8 - 14.9	39.2 - 45.4ns	144.9 - 167.2ns	160.1ns	**213.1 - 227.8ns**	**216.9ns**	91.67 GB/s
TR 1950X Game Mode DDR4-2666	13.9 - 14.2ns	39.5 - 42.3ns	149.2 - 164.1ns	159.66ns	**X**	**X**	46.58 GB/s
TR 1950X Game Mode DDR4-3200	14.3 - 14.9ns	41.2 - 46.2ns	123 - 150.6ns	145.44ns	**X**	**X**	45.52 GB/s
Ryzen 7 1800X	14.8ns	40.5 - 82.8ns	120.9 - 126.2ns	122.96ns	**X**	**X**	48.1 GB/s
Ryzen 5 1600X	14.7 - 14.8ns	40.6 - 82.8ns	121.5 - 128.2ns	123.48ns	**X**	**X**	43.88 GB/s

Interpretation of the Infinity Fabric latencies in the Table above [24]

The intra-core latency reflects the communication between two logical threads resident on the same physical core, it is unaffected by memory speed.

Intra-CCX latencies quantify latency between threads that are on the same CCX but not resident on the same core**, it is** also largely unaffected by memory speed.

Cross-CCX latency denotes latency between threads located on two separate CCX building blocks.

Die-to-Die latency relates to latency between threads that are on different dies.

In Creator mode of the TR 1950X both dies are active but in the Game mode one die is disabled so no Die-to-Die latency will occur.

Similarly, Ryzen line processors include only a single die, so no Die-to-Die latency can occur.

Block diagram of a ThreadRipper platform [23]

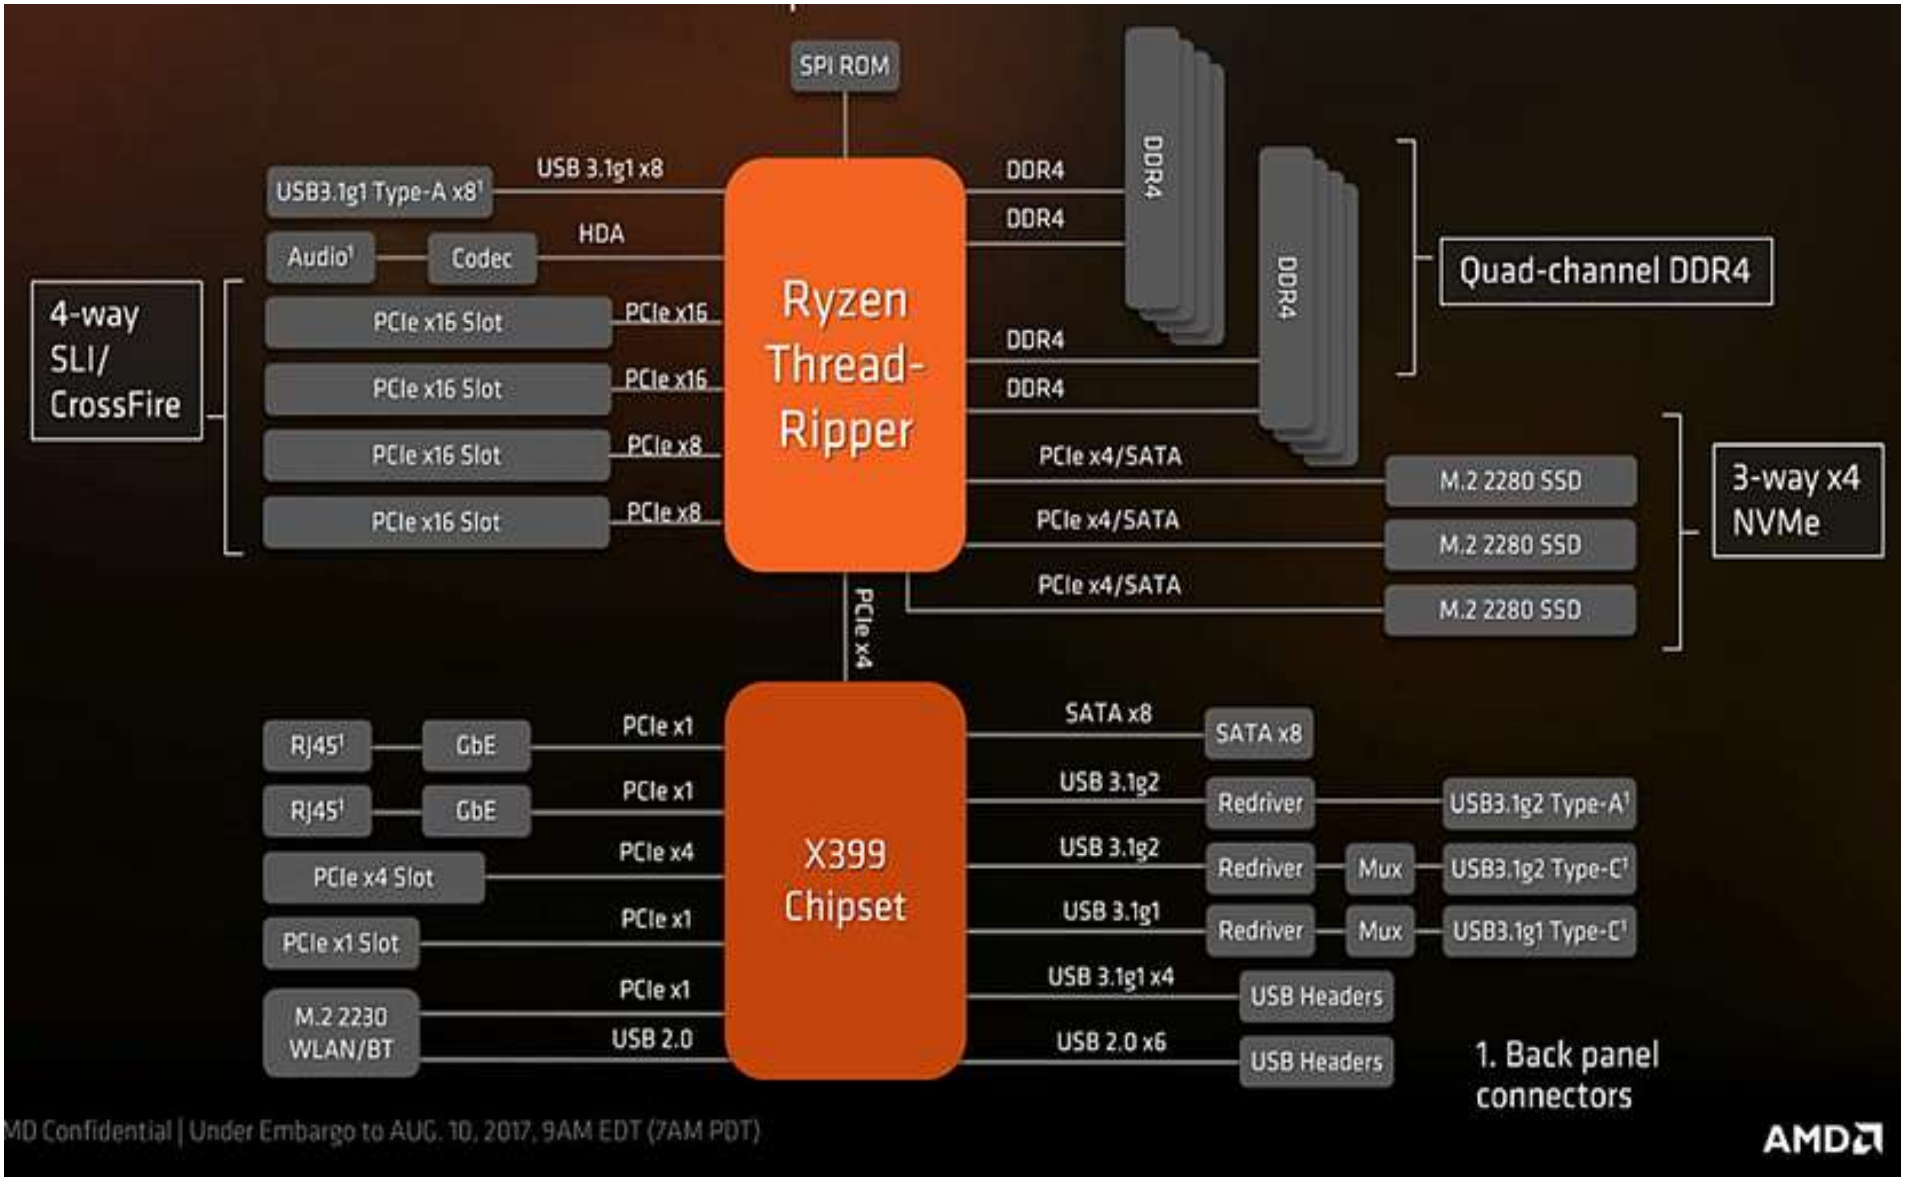

MCM implementation of the ThreadRipper processor line

AMD reuses the packaging of the EPYC line including the cooling solution, it means

- they put 4 chips into the package whereas two chips are actually dummies, as seen below and
- they use the TR4 socket with 4094 contacts.

Figure: Cracked open Threadripper package [26]

Main features of the Ryzen and ThreadRipper models [23]

AMD Ryzen SKUs								
	Cores/ Threads	Base/ Turbo	XFR	L3	DRAM 1DPC	PCIe	TDP	SRP
TR 1950X	**16/32**	**3.4/4.0**	**+200**	**32 MB**	**4x2666**	**60**	**180W**	**$999**
TR 1920X	**12/24**	**3.5/4.0**	**+200**	**32 MB**	**4x2666**	**60**	**180W**	**$799**
TR 1920**	12/24	3.2/3.8	?	32 MB	4-Ch?	60	140W	?
TR 1900X	8/16	3.8/4.0	+200	16 MB*	4x2666*	60	180W*	$549
Ryzen 7 1800X	8/16	3.6/4.0	+100	16 MB	2x2666	16	95 W	$499
Ryzen 7 1700X	8/16	3.4/3.8	+100	16 MB	2x2666	16	95 W	$399
Ryzen 7 1700	8/16	3.0/3.7	+50	16 MB	2x2666	16	65 W	$329
Ryzen 5 1600X	6/12	3.6/4.0	+100	16 MB	2x2666	16	95 W	$249
Ryzen 5 1600	6/12	3.2/3.6	+100	16 MB	2x2666	16	65 W	$219
Ryzen 5 1500X	4/8	3.5/3.7	+200	16 MB	2x2666	16	65 W	$189
Ryzen 5 1400	4/8	3.2/3.4	+50	8 MB	2x2666	16	65 W	$169
Ryzen 3 1300X	4/4	3.5/3.7	+200	8 MB	2x2666	16	65 W	$129
Ryzen 3 1200	4/4	3.1/3.4	+50	8 MB	2x2666	16	65 W	$109

Advanced technologies implemented in the ThreadRipper line

1) Pure Power technology (actually AVS)
2) Precision Boost technology (actually Turbo Boost)
3) XFR (eXtreme Frequency Range) technology (actually configurable TDP)

These technologies has already been discussed in Section 6.1, here we will briefly touch upon some peculiarities of the implementation of the Precision Boost technology in the ThreadRipper line.

Boost behavior of the original Precision Boost technology -1

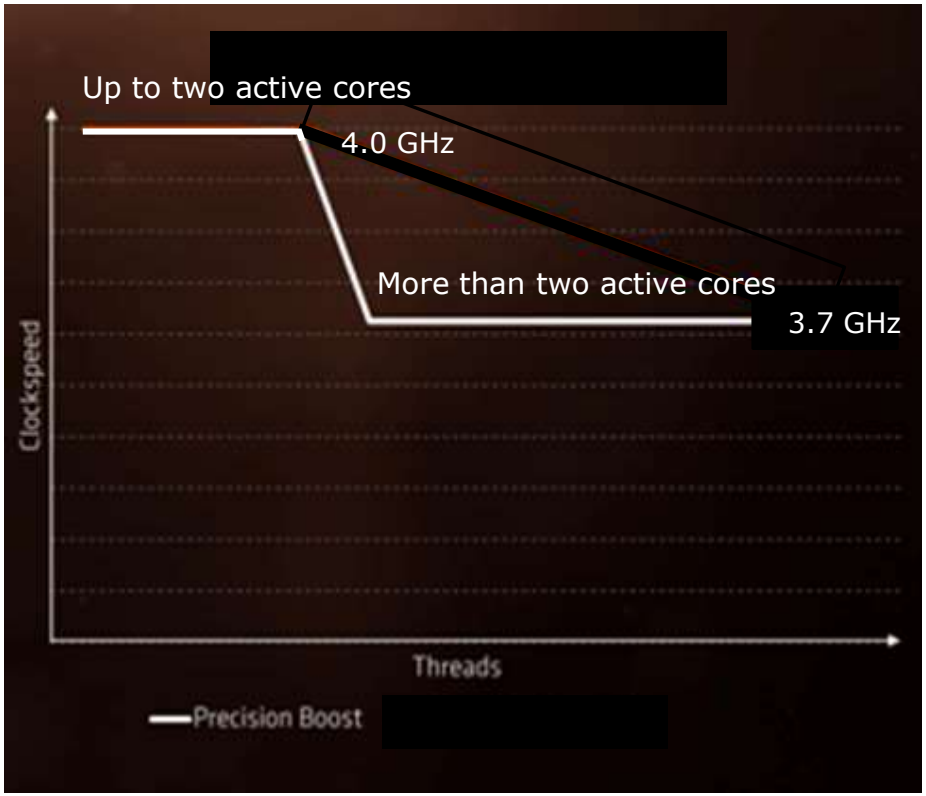

Figure: Boost behavior of the Ryzen DT and mobile lines [38]

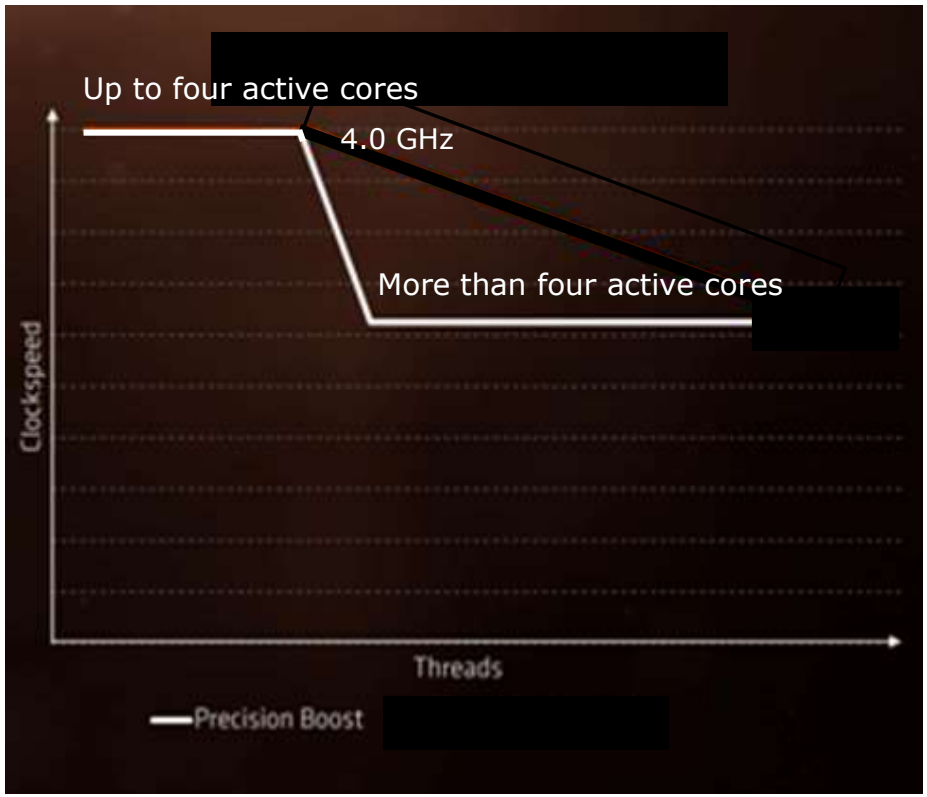

Figure: Boost behavior of the ThreadRipper line (based on [38])

Main parameters of the implementation of Precision Boost and XFR in the ThreadRipper line [86]

Model	Base clock	Boost all	Boost up to 4 cores	XFR
TR1900X	3.8 GHz	3.7 GHz	4.0 GHz	4.2 GHz
TR1900X	3.5 GHz			
TR1900X	3.4 GHz			

Content creation performance

Remarks on the CineBench benchmark [27]

- It is a real-world cross-platform test suite that evaluates CPU and graphics performance for 3D content creation.
- It is based on MAXON's award-winning animation software Cinema 4D.

CineBench R15.38 Single-Threaded performance [20]

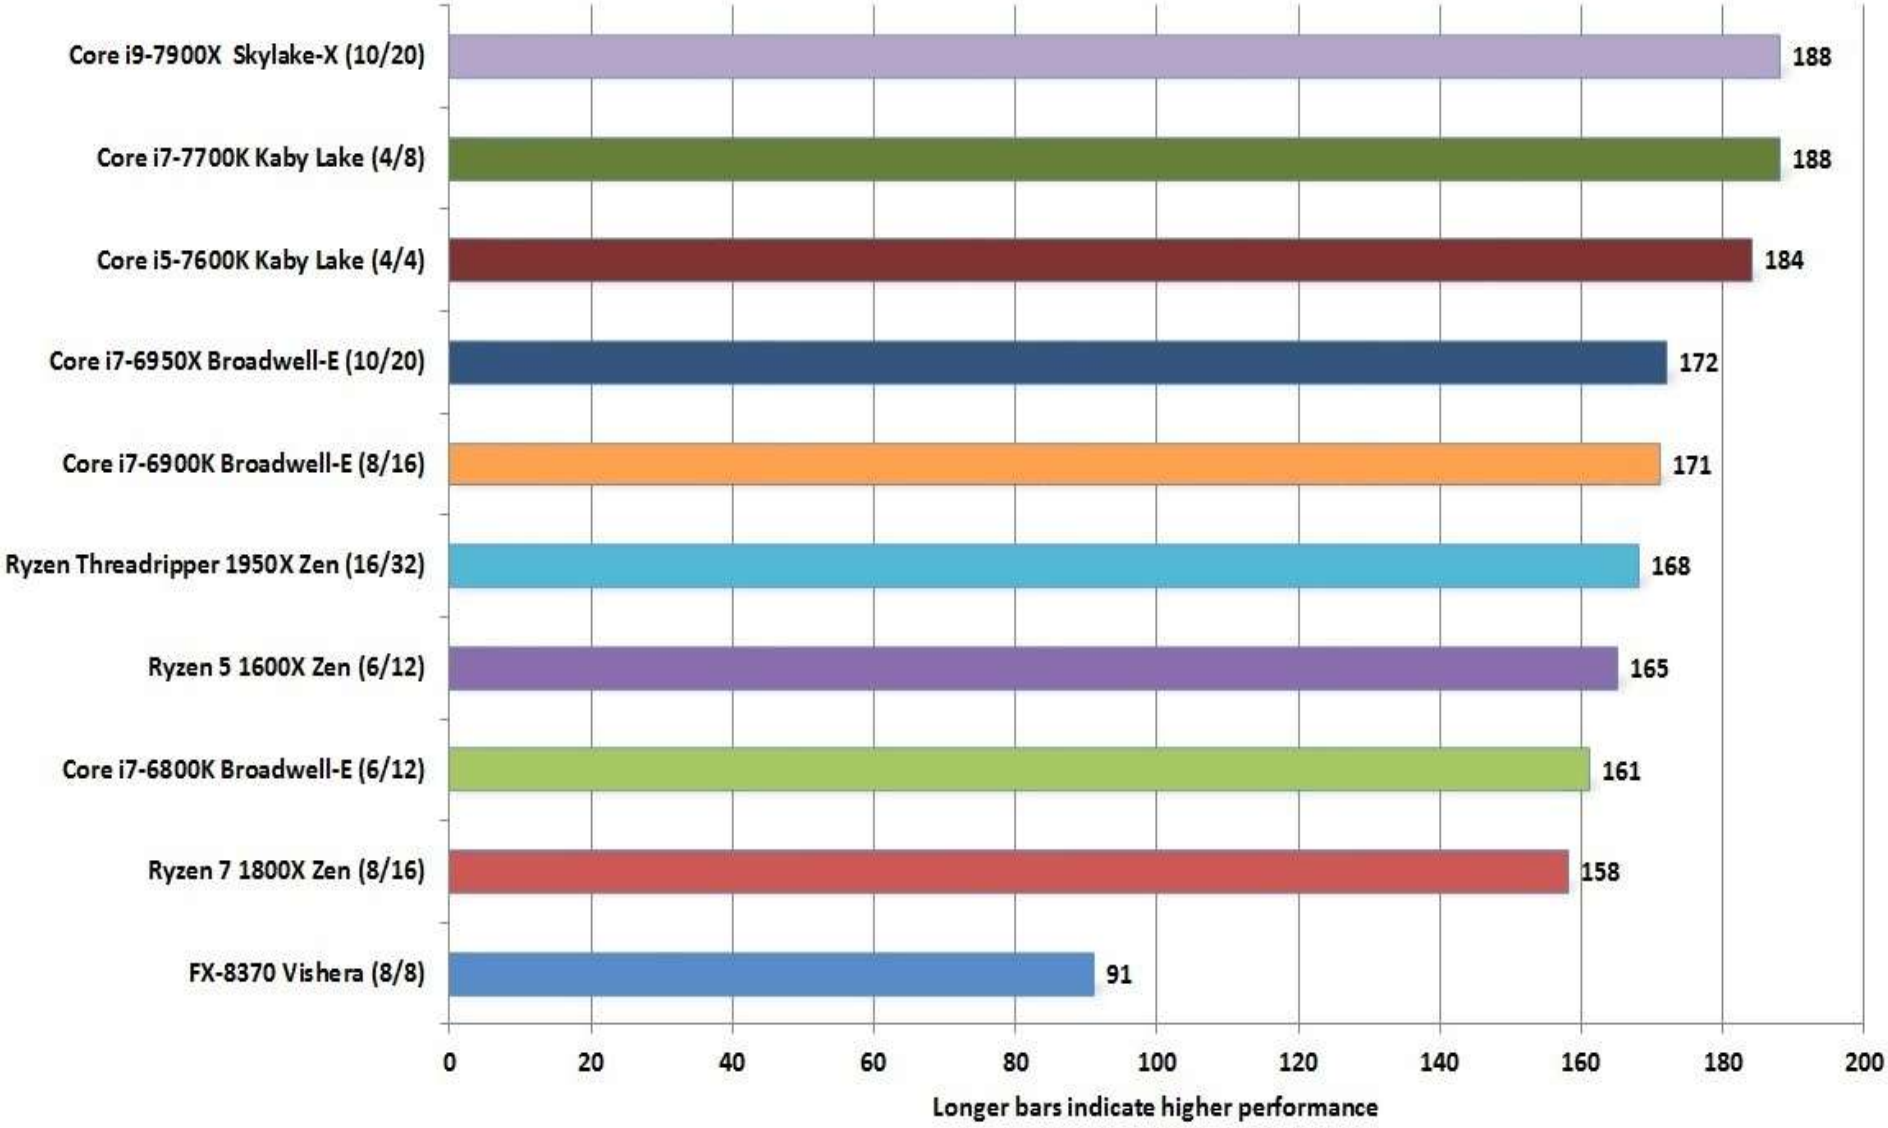

CineBench R15.38 Multi-Threaded performance [20]

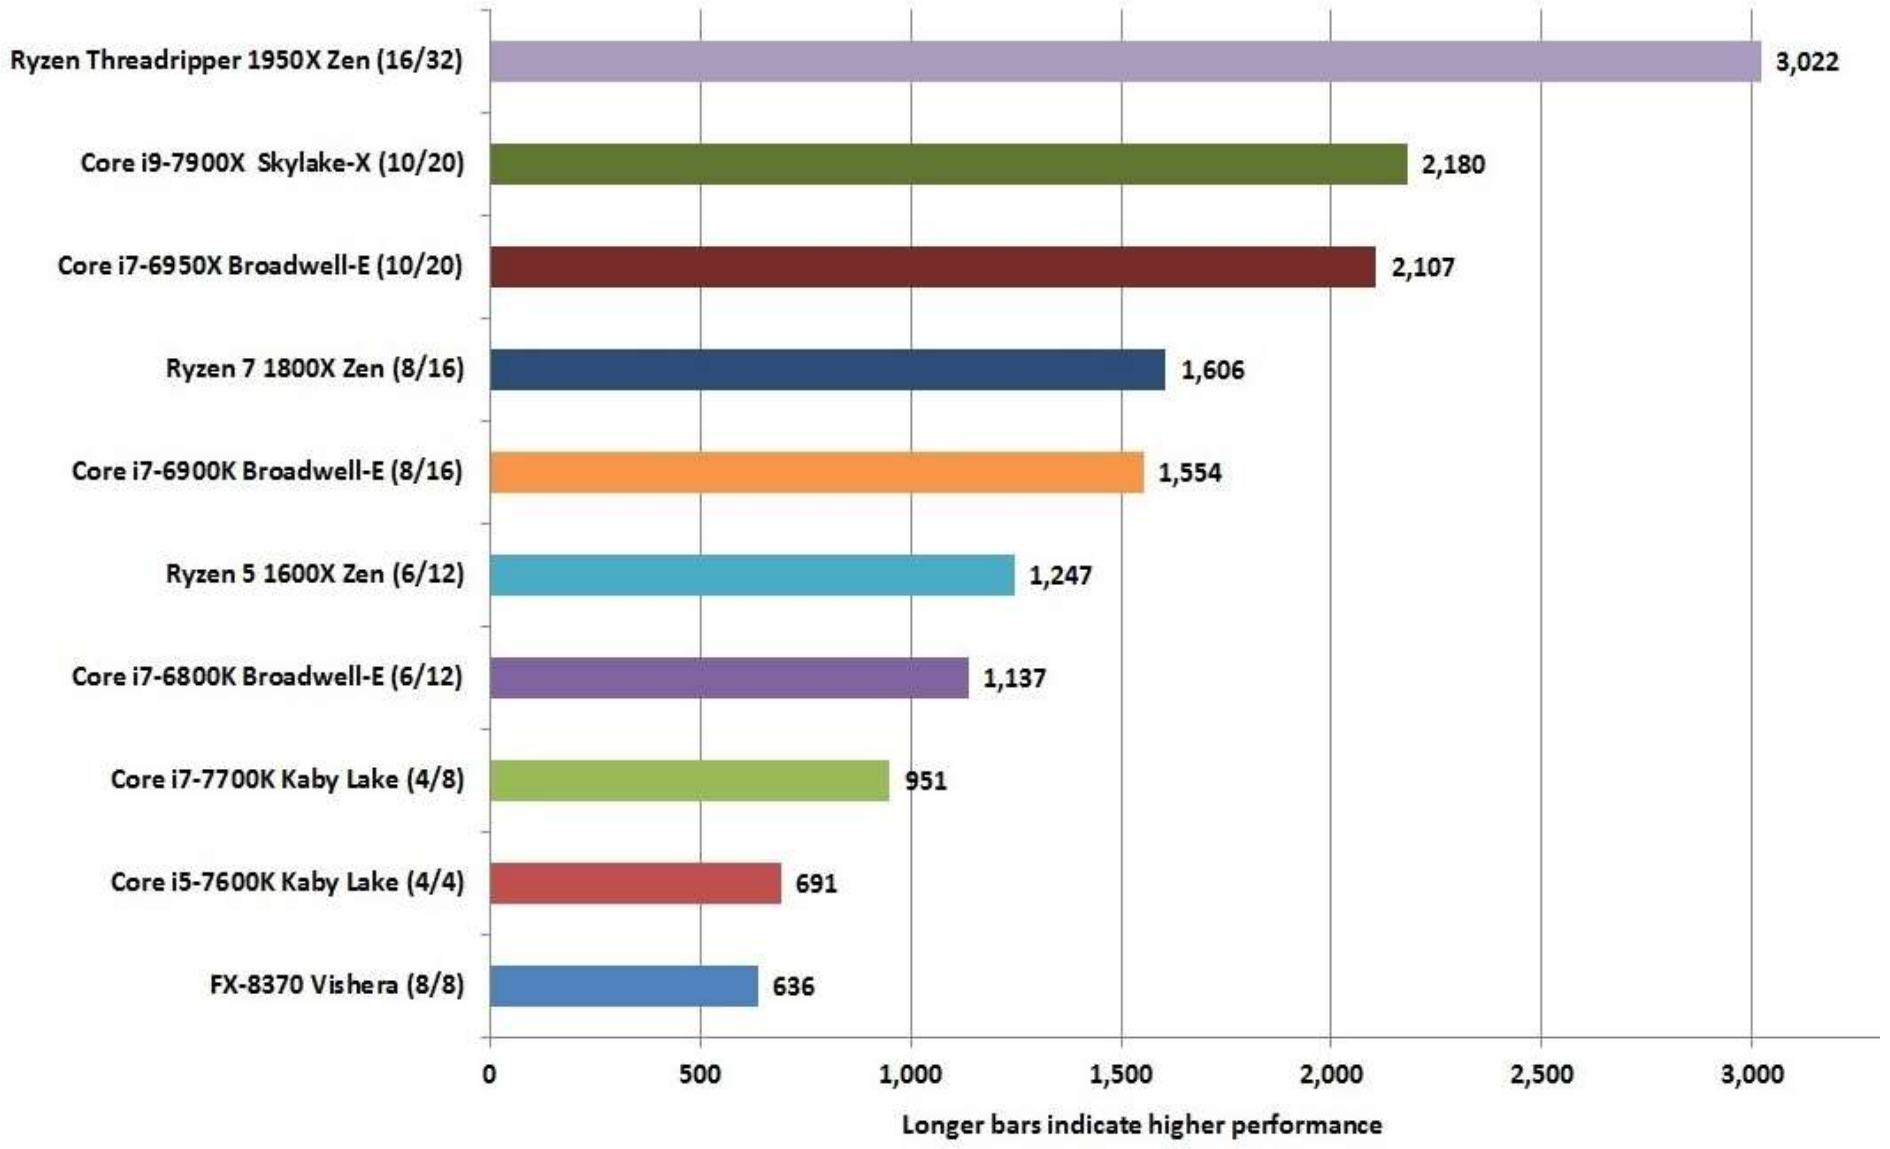

Content creation performance summary

As far as single thread performance matters Intel's high-end HED processors provide a better performance (IPC), nevertheless for multithread applications AMD's ThreadRipper 1950X is the clear winner.

Gaming performance

Creator mode - Game mode of the ThreadRipper line

- The Creator mode is the natural configuration of ThreadRipper processors with all cores enabled.
- In Gaming mode only one die is enabled, so only half of the cores are active.

  This can be beneficial while running older games that are not prepared to use more than 8 cores, in this case long Die-to-Die latencies could significantly reduce performance.

Gaming performance (Rainbow Six Siege 19x10 medium) average FPS [20]

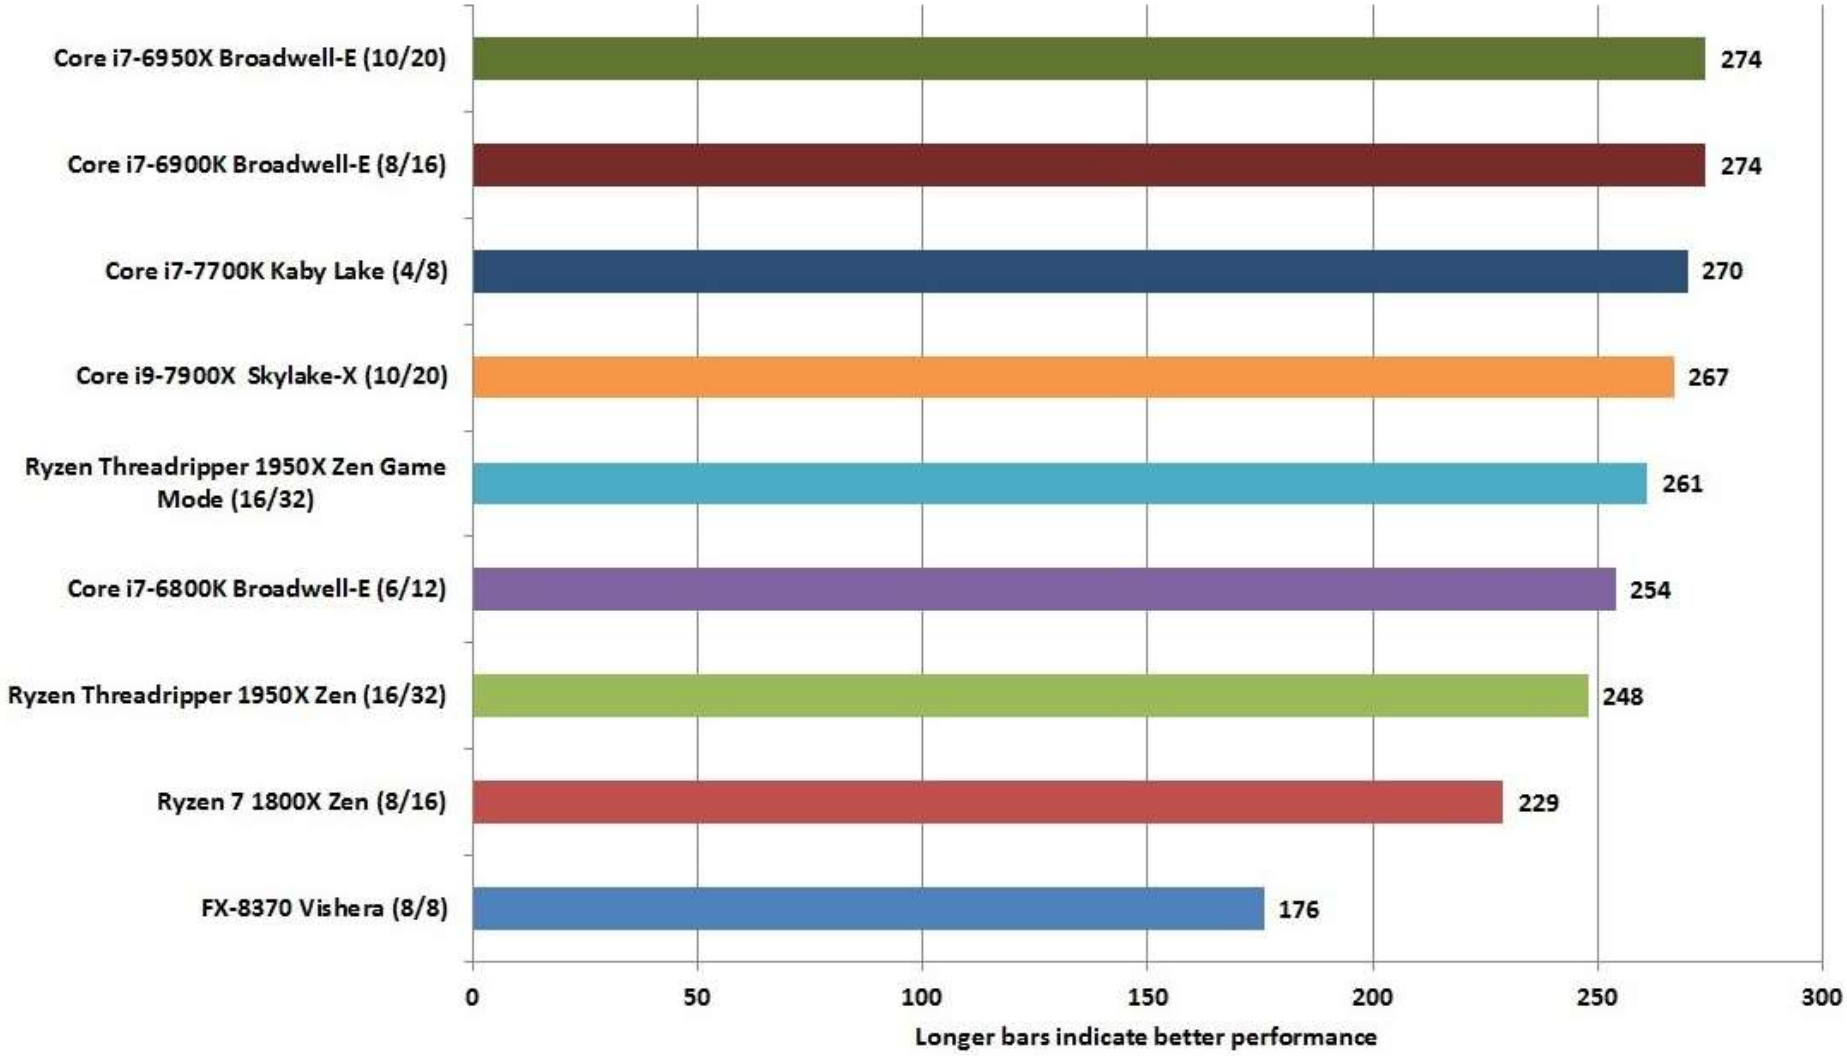

Gaming performance summary

In gaming Intel's HED processor provide 5 to 20 % better performance than AMD's ThreadRipper processors, nevertheless "no one should buy a $1,000, 16-core CPU just to play conventional gaming or run lightly threaded applications„ [28].

Contrasting AMD's and Intel's competing HED processor models [23]

Intel's vs. AMD's competing processors									
		Cores/ Threads	Base/ Turbo	XFR	L3	DRAM 1DPC	PCIe	TDP	Cost (8/10)
AMD	**TR 1950X**	**16/32**	**3.4/4.0**	**+200**	**32 MB**	**4x2666**	**60**	**180W**	**$999**
Intel	**i9-7900X**	10/20	3.3/4.3	+200	13.75	4x2666	44	140W	$980
Intel	**i7-6950X**	10/20	3.0/3.5	+500	25 MB	4x2400	40	140W	$1499
AMD	**TR 1920X**	**12/24**	**3.5/4.0**	**+200**	**32 MB**	**4x2666**	**60**	**180W**	**$799**
Intel	**i7-7820X**	8/16	3.6/4.3	+200	11 MB	4x2666	28	140W	$593
AMD	**TR 1900X**	8/16	3.8/4.0	+200	16 MB	4x2666	60	180W	$549
AMD	**R7 1800X**	8/16	3.6/4.0	+100	16 MB	2x2666	16	95 W	$419
AMD	**R7 1700X**	8/16	3.4/3.8	+100	16 MB	2x2666	16	95 W	$350
AMD	**R7 1700**	8/16	3.0/3.7	+50	16 MB	2x2666	16	65 W	$291

Sales figures Intel's Skylake/Kaby Lake HEDs vs. AMD's Ryzen HEDs March - Aug. 2017 by Mindfactory Germany (on-line retailer company) [34]

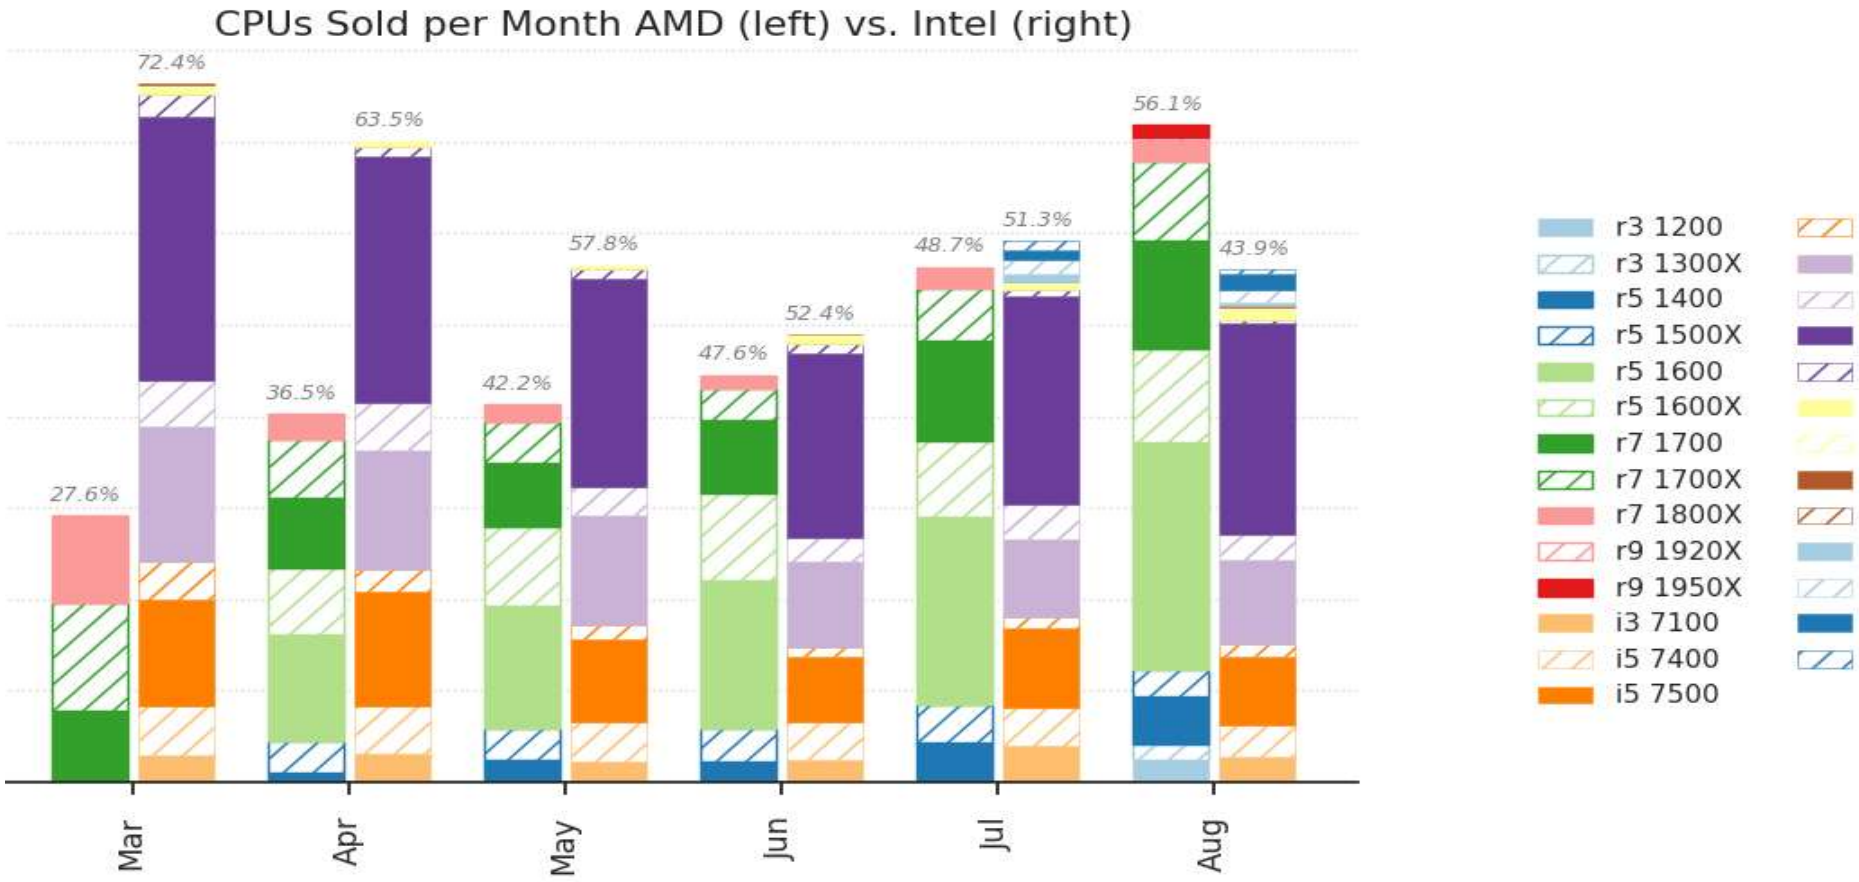

8.2 The 2. gen. ThreadRipper HED line

8.2 The 2. gen. ThreadRipper HED line -1

- Launched in 8/2018, designated as Pinnacle Ridge.
- It is based on the Zen+ μarchitecture, implemented on 12nm technology.

The 2. gen. ThreadRipper HED line -2

AMD's Zen/Zen+ based processor lines introduced in 2018

	Ryzen Mobil APU (Raven Ridge)	**Ryzen DT APU** (Raven Ridge)	**2.G Ryzen DT** (Pinnacle Ridge)	**2.G ThreadRipper** (Pinnacle Ridge)
µarch./tech.	Zen/14 nm	Zen/14 nm	Zen+/12 nm	Zen+/12 nm
Launched models	Ryzen 3 2300U Ryzen 3 2200U (1/2018)	Ryzen 5 2400G/GE Ryzen 3 2200G/GE (G: 2/2018) (GE: 42018)	Ryzen 7 2700X Ryzen 7 2700 Ryzen 5 2600X Ryzen 5 2600 (4/2018)	2990WX (8/2018) 2950X (8/2018) 2970WX (10/2018) 2920X (10/2018)
Layout	CCX + Vega 6/3	CCX + Vega 11/8	Zeppelin die (2x CCX)	MCM (up to 4 Zeppelin dies)
Integrated GPU	Yes	Yes	No	No
Core count	2/4	4	2/4	Up to 32
SMT (Multithreaded)	SMT only for Ryzen 3 2200U	SMT only for Ryzen 5	SMT	SMT
Mem. channels/ data rate	2xDDR4-2400	2xDDR4-2667 (G) 2xDDR4-2933(GE)	2xDDR4-2933	4xDDR4-2933
PCIe lanes	??	??	16x PCIe 3.0	60x PCIe 3.0
TDP	15 W	35/65 W	65/95/105 W	180/250 W
Socket	FP5 (na.)	AM4 (1331)	AM4 (1331)	TR4 (SP3r2) (4094)
Chipset	SoC	300-series	300/400-series	X399

Positioning the 2. gen. ThreadRipper HED models [73]

1. Gen ThreadRipper (2017)				2. gen. ThreadRipper (2018)		
Model	Cores/ Threads	Price		Price	Cores/ Threads	Model
-				$1799	32/64	**TR 2990WX**
-				$1299	24/48	**TR 2970WX**
TR 1950X	16/32	$999		$899	16/32	**TR 2950X**
TR 1920X	12/24	$799		$649	12/24	**TR 2920X**
TR 1900X	8/16	$549				

As seen in the Table, two of the new models are direct replacements of previous 1. gen. models (TR 2951X and the TR 2920X) whereas two new models (TR 2990WX and TR2970WX) extend available core counts to 32 and 24 cores, respectively.

Main features of the 2. gen. ThreadRipper processors [73]

2. gen. ThreadRipper processors								
	Launched	Cores/ Threads	Base/ Turbo	L3	DRAM 1DPC	PCIe	TDP	SRP
TR 2990WX	8/2018	32/64	3.0/4.2	64 MB	4x2933	60	250 W	$1799
TR 2970WX	10/2018	24/48	3.0/4.2	64 MB	4x2933	60	250 W	$1299
TR 2950X	8/2018	16/32	3.5/4.4	32 MB	4x2933	60	180 W	$899
TR 2920X	10/2018	12/24	3.5/4.3	32 MB	4x2933	60	180 W	$649
Ryzen 7 2700X (DT in contrast)	*4/2018*	*8/16*	*3.7/4.3*	*16 MB*	*2x2933*	*16*	*105 W*	*$329*

Architectural layout of the 2. gen. ThreadRipper models [73] -1

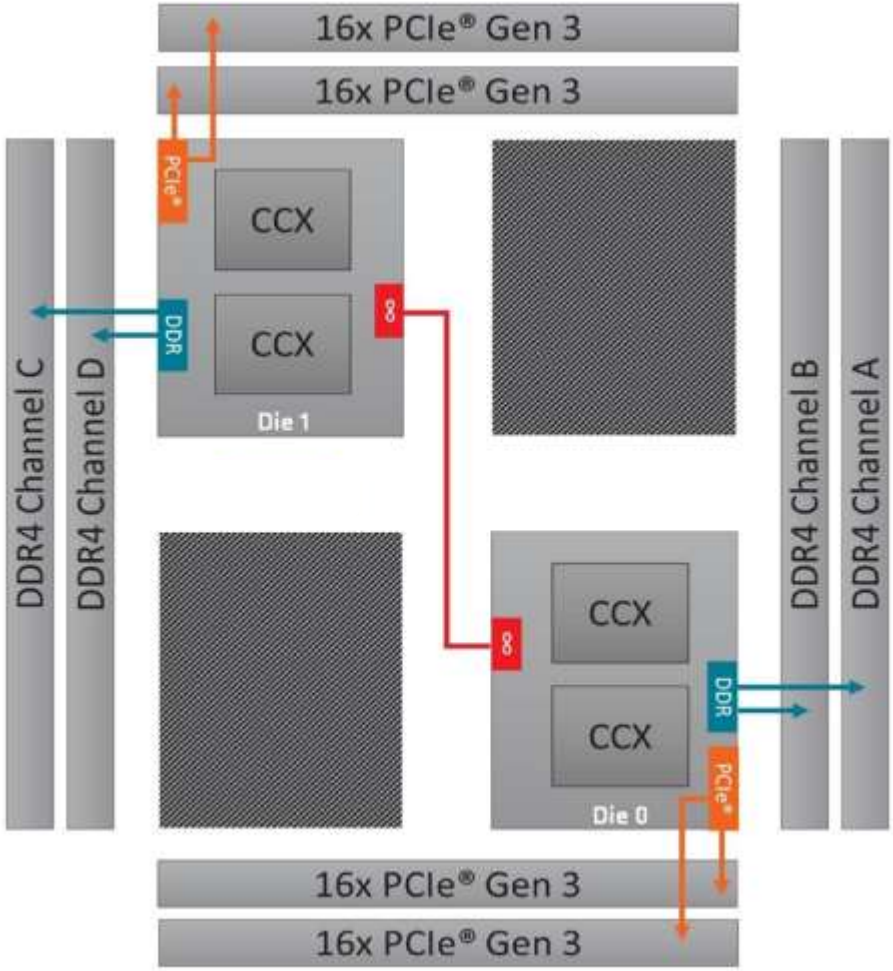

TR 2950: 8+8 cores
TR 2920: 6+6 cores

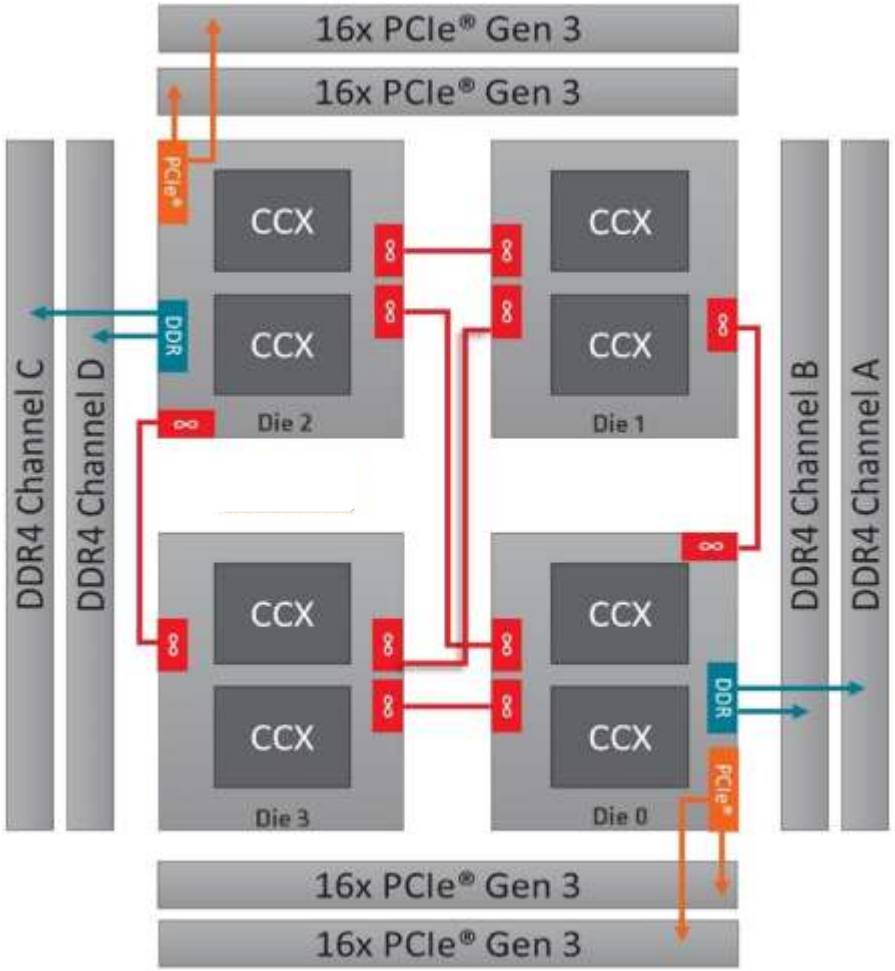

TR 2990WX: 8+8+8+8 cores
TR 2970WX: 6+6+6+6 cores

Architectural layout of the 2. gen. ThreadRipper models [73] -2

What the above Figure indicates is that

- the low core count models have two active dies with 8 or 6 cores each, with both dies having direct access to memory and I/O, similar to the 1. gen. models, whereas
- the high core count models incorporate four active dies with 8 or 6 cores each.
- In the latter case
  - two dies have direct access to memory and I/O whereas
  - the other two dies have no direct access to memory or I/O, only via the direct-access dies, obviously with higher latency.
- The new high core count models have more compute resources and the same memory and I/O resources, as the low core count models, obviously they target users with compute bound applications, like simulations or physics calculations.
- The above asymmetric kind of resource usage can be designated as being bi-modal.

*

Implications of the bi-modal layout of the high-core ThreadRipper models for the scheduler

- For obvious reasons the scheduler will assign tasks first to cores that are directly attached to the memory and I/O before loading other cores.
- Nevertheless, to prevent overheating, the scheduler will likely avoid loading all cores of the direct connected dies before allocating tasks to not directly connected dies, e.g. after loading 12 or 14 cores out of 16 it begins already loading cores of not directly connected dies.

*

Contrasting the layouts of the high core count ThreadRipper and the EPIC models [73] -1

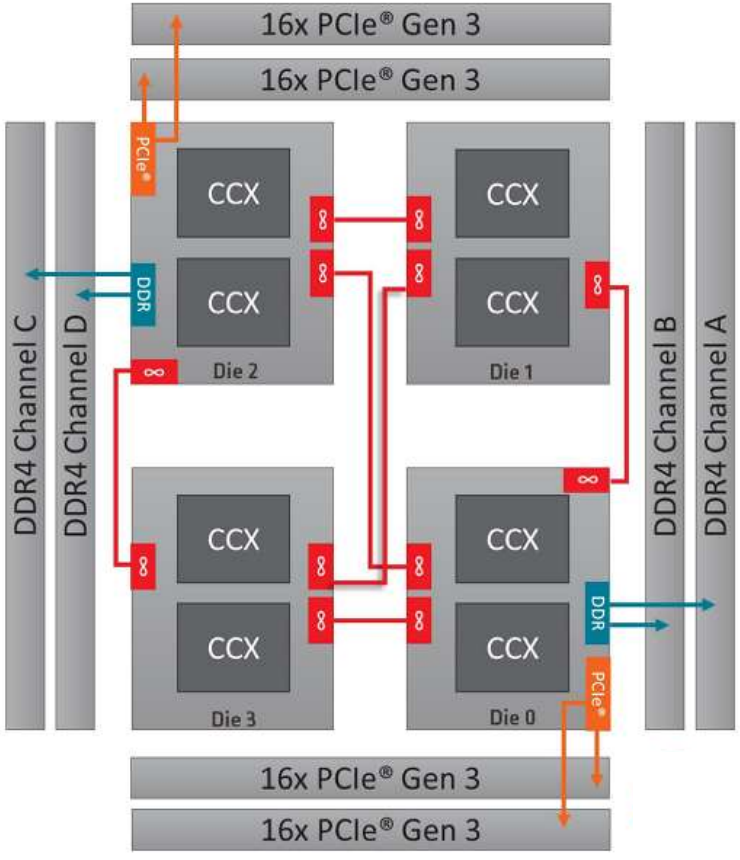

ThreadRipper 2990WX/2970WX
(TR4 (SP3r2) socket, 4094 pins))

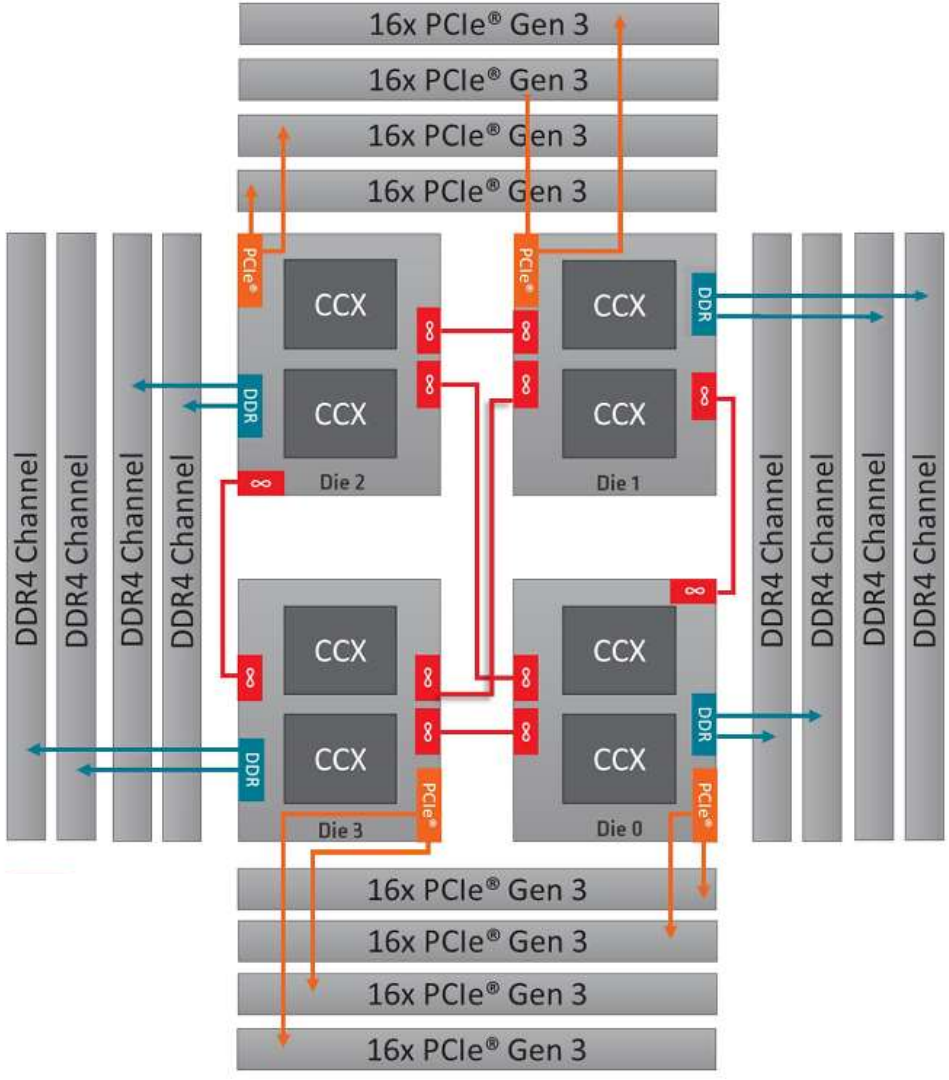

EPIC
(SP3 socket, 4094 pins, not compatible)

Contrasting the layouts of the high core count ThreadRipper and the EPIC models [73] -2

- In the 2. gen. 2990WX/2970WX design (seen on the left of the Figure), the two inactive dies are now enabled, and there is a crossbar die-to-die interconnect between all four cores.

  Cores on the now enabled two dies do not have direct access to memory or I/O thus for these cores each memory access requires an extra hop through the die-to-die interconnect using AMD's interconnect fabric which consumes extra power (as will be discussed subsequently).

- By contrast, each die of an Epyc server has direct access to memory and I/O, resulting in doubling of the available memory bandwidth and number of PCIe lanes.

  This reduces the average access time of the memory but increases the power consumption needed for data transfers.

Recap: The Infinity Fabric (IF) [73]

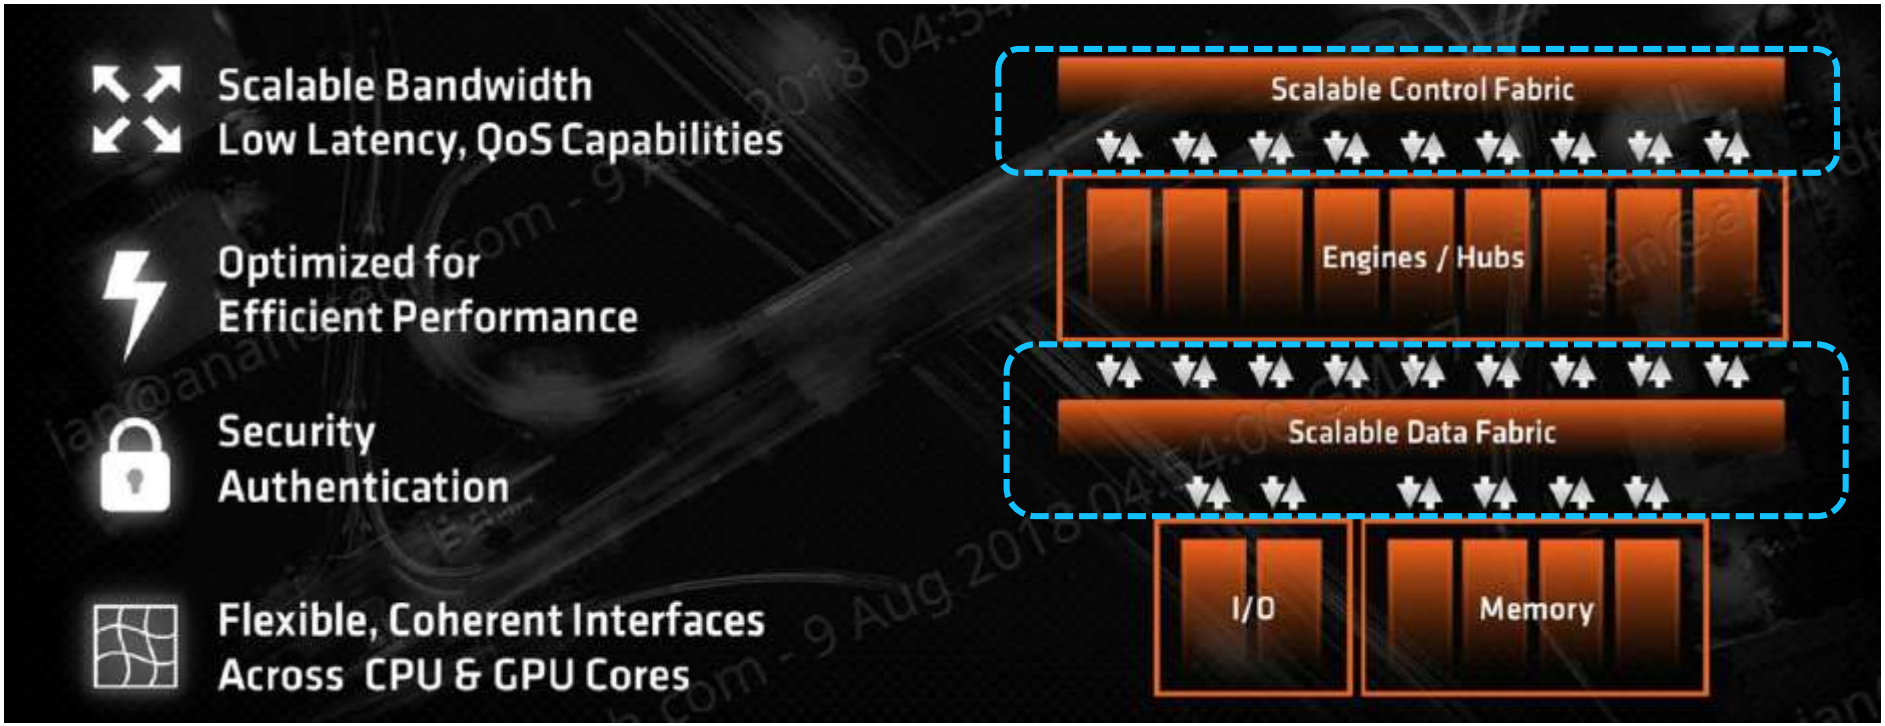

It interconnects computing resources (like the cores, the GPU), memory and I/O.

8.2 The 2. gen. ThreadRipper HED line (10)

Power consumption of the cores while raising the number of active cores in the 2. gen. 16-core ThreadRipper 2950X processor (TDP: 180 W) [63]

ANANDTECH — Ryzen Threadripper 2950X (8+0+8+0) Loading (W)

Core	2T	4T	6T	8T	10T	12T	14T	16T	18T	20T	22T	24T	26T	28T	30T	32T
0	22.97	22.97	18.73	17.32	16.83	16.11	15.72	14.80	14.22	12.57	11.45	10.59	9.39	7.90	9.04	7.68
1	0.06	22.88	18.42	17.37	16.72	16.26	15.78	14.65	14.35	12.62	11.38	10.80	9.43	7.97	9.13	9.01
2	0.05	0.08	18.75	17.42	16.80	15.99	15.72	14.72	14.26	12.65	11.36	10.60	9.36	7.93	7.09	7.77
3	0.05	0.06	0.04	17.38	16.74	16.08	15.78	14.82	14.35	12.64	11.35	10.64	9.39	10.29	9.22	8.68
4	0.04	0.04	0.08	0.11	18.23	17.04	16.88	14.89	14.49	12.72	11.59	10.76	9.40	9.30	7.46	8.12
5	0.05	0.06	0.08	0.10	0.09	17.40	16.98	14.81	14.20	12.68	11.35	10.91	9.50	9.29	9.51	9.80
6	0.04	0.04	0.06	0.06	0.10	0.07	16.74	14.69	14.19	12.61	11.35	10.60	9.35	7.34	7.63	7.97
7	0.05	0.06	0.07	0.10	0.08	0.09	0.15	14.76	14.23	12.72	11.48	10.62	9.42	9.33	9.66	8.25
8	0.04	0.04	0.03	0.10	0.03	0.15	0.05	0.09	17.49	15.15	13.22	11.26	9.77	9.33	8.81	8.84
9	0.05	0.04	0.04	0.08	0.03	0.04	0.04	0.08	0.75	15.20	13.22	11.44	9.84	9.24	8.85	8.68
10	0.04	0.04	0.04	0.04	0.03	0.08	0.04	0.40	0.07	0.12	13.34	11.26	9.80	9.32	8.89	7.40
11	0.04	0.04	0.03	0.04	0.07	0.07	0.04	0.04	0.05	0.10	1.07	11.19	9.74	9.35	8.89	8.88
12	0.13	0.10	0.05	0.05	0.04	0.04	0.04	0.05	0.04	0.05	1.84	0.07	12.73	11.44	10.32	9.43
13	0.09	0.10	0.08	0.08	0.11	0.12	0.08	0.12	0.08	0.08	0.08	1.74	0.12	11.54	10.47	9.22
14	0.06	0.06	0.04	0.05	0.03	0.04	0.04	0.04	0.04	0.03	0.03	0.04	0.71	3.75	7.97	7.88
15	0.07	0.08	0.06	0.06	0.06	0.06	0.27	0.06	0.06	0.07	0.05	0.06	0.08	0.11	0.17	7.15
Cores	23.83	46.69	56.60	70.36	85.99	99.64	114.35	119.02	132.87	132.01	134.16	132.58	128.03	133.43	133.11	134.76
Package	58.17	83.98	95.34	111.09	127.99	142.99	159.35	163.81	176.04	174.22	174.73	173.77	175.73	175.68	176.09	177.88

There are up to two active dies with up to 8 active cores/unit.
Green figures: not loaded cores - Yellow to red figures: loaded cores

Red figures appearing for few active cores indicate that in this case presumably die temperature will constrain turbo clock frequency and determine the related power consumption, whereas for yellow figures, seeing at many active cores, clearly the total power consumption limits clock frequency.

*

Breakdown of power consumption between the cores and the IF in the 16 cores (2 active dies) TR 2950X depending on loading [73]

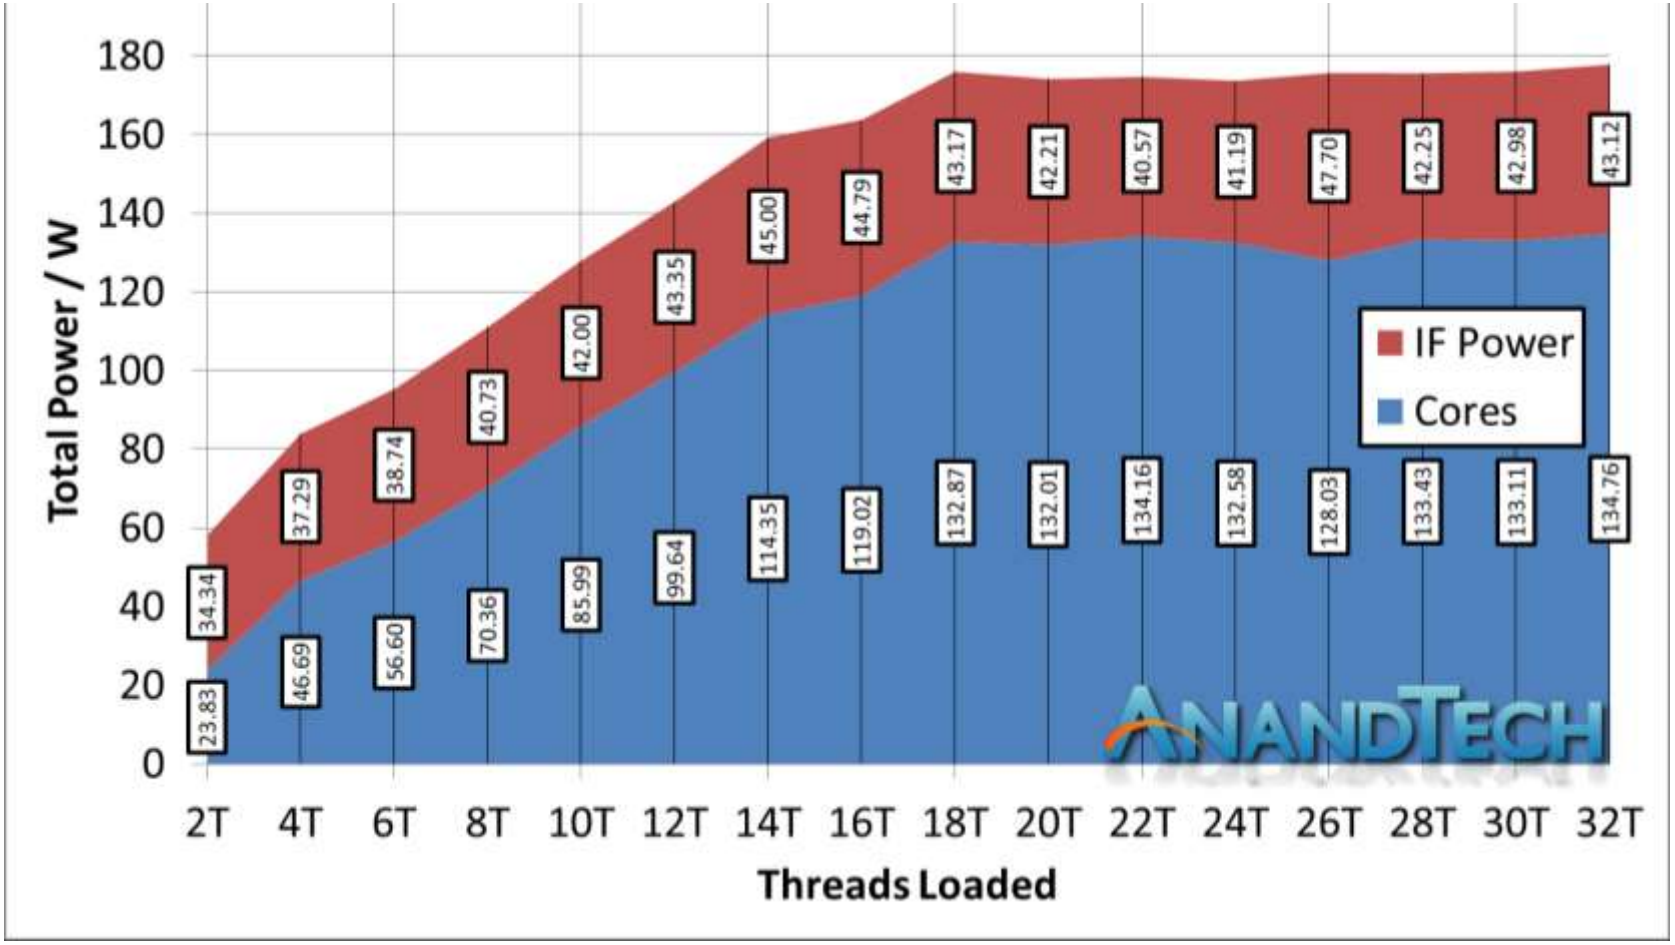

As seen, the IF of the 16-core (2 active dies) TR 2950X consumes about 25 % of the total power.

*

Breakdown of power consumption between the cores and the IF in the 32-core, 4 active dies) TR 2990WX depending on loading [73]

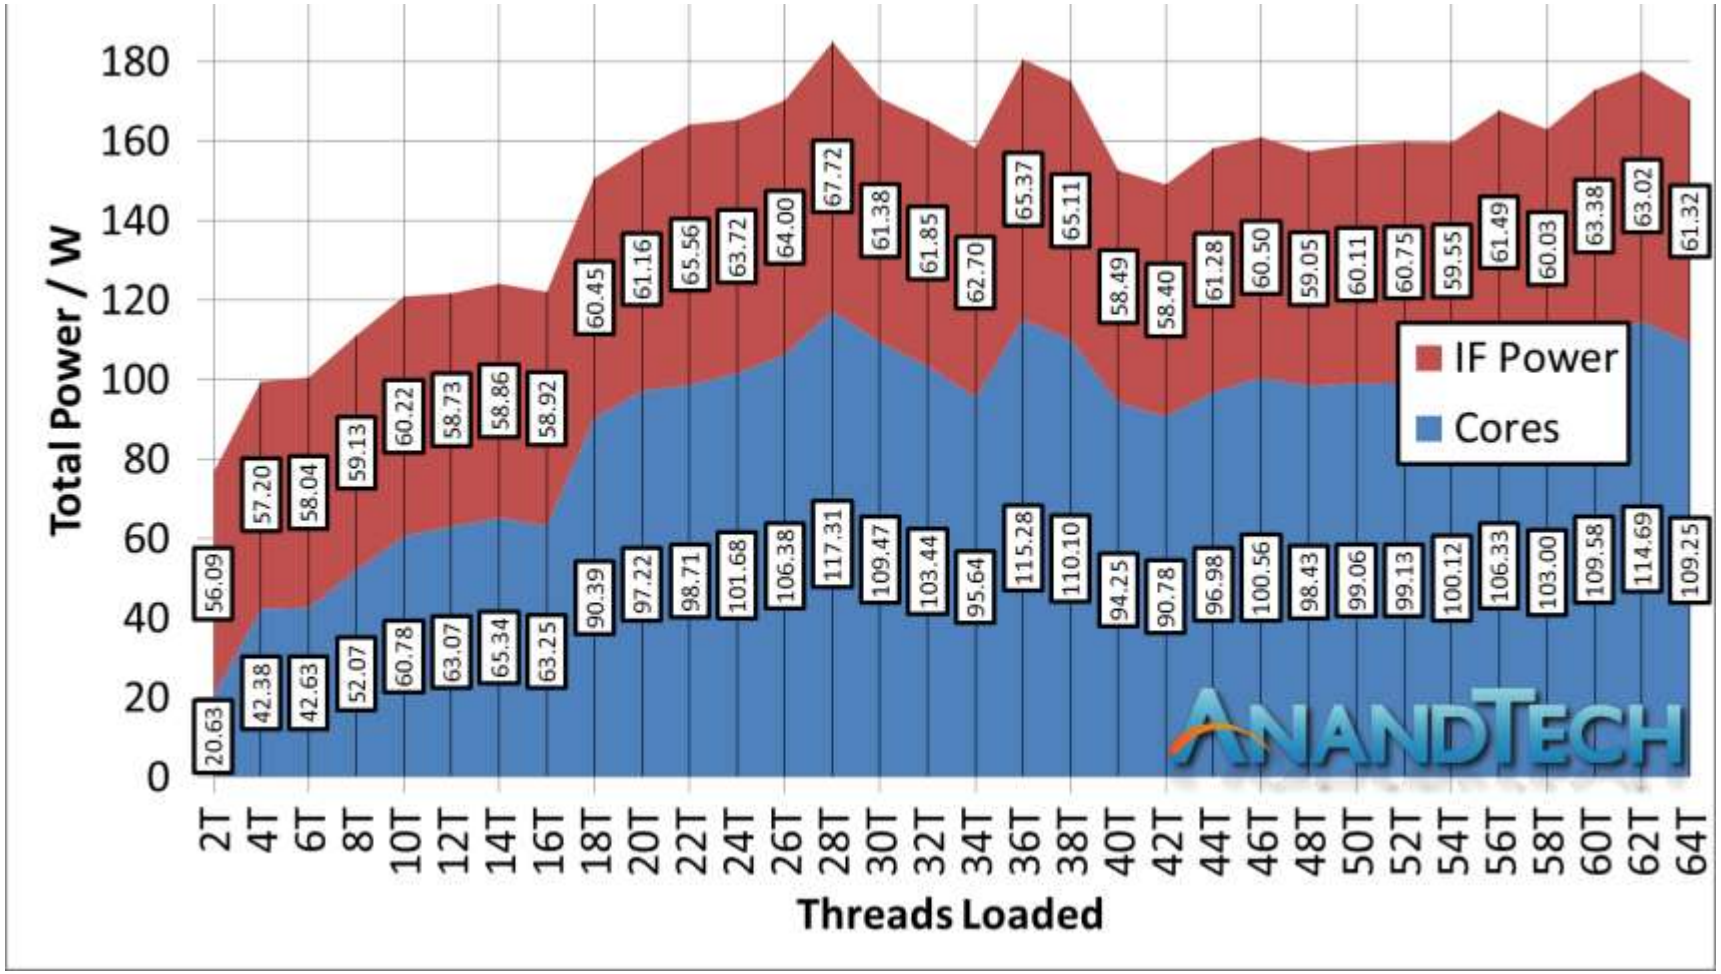

As seen, the IF of the 32-core (4 active dies) TR 2990WX consumes already about 35 % of the total power due to the much more complex die-to-die interconnects. *

In contrast: Breakdown of power consumption between the cores and the IF in the 32-core 4 active dies) EPIC 7610 depending on loading [73]

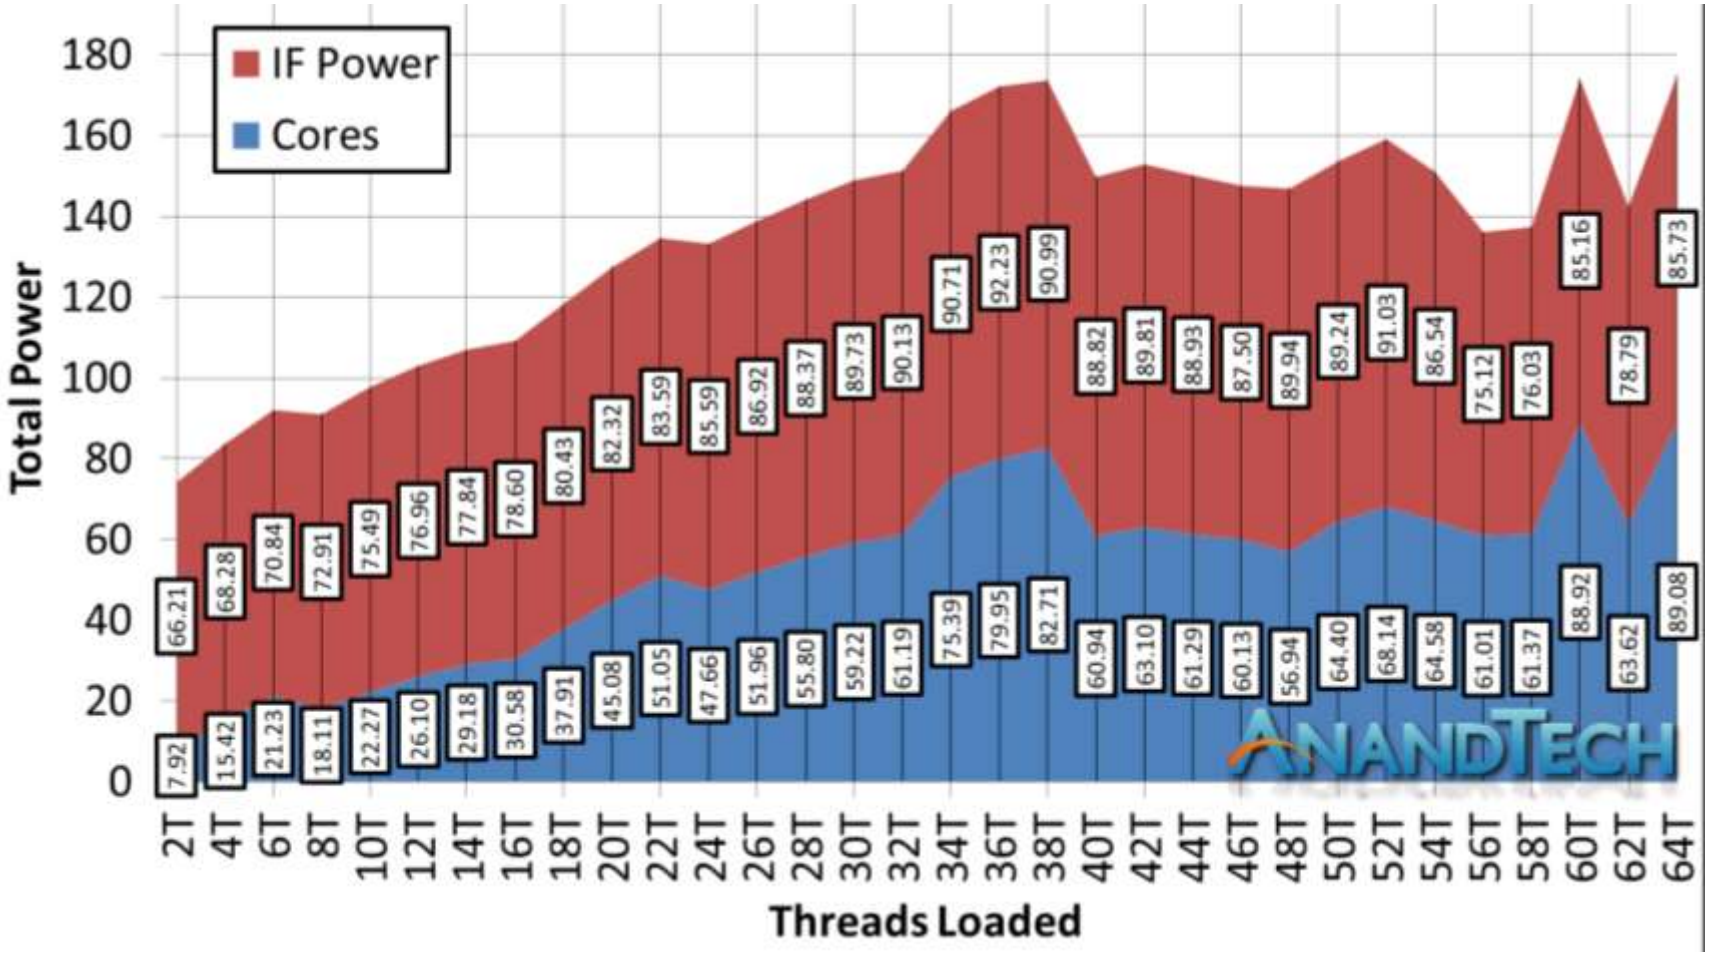

As seen, the IF of the 32-core (4 active dies) EPIC 7601 consumes already up to 50 % of the total power due to the much more complex die-to-die interconnect. *

Note [73]

Preceding Figures let’s conclude that with increasing core count - beyond providing enough memory bandwidth - the implementation of low power, scalable and high performance interconnects will be a great challenge for processor designers.

Enhancements of the 2. gen. ThreadRipper line vs. the 1. gen. line

a) Higher core counts (24 and 32 core models) vs. up to 16-core models
b) Precision Boost 2 (decrease of clock frequencies for higher number of active cores in Turbo mode)
a) Precision Boost Overdrive (overriding factory settings)
b) StoreMe (unified storage)

- The point of higher core counts has already been previously discussed, whereas the enhancements b) to d) has already been detailed in Section 6.2.
- Here subsequently we will only emphasize some peculiarities of the implementation of the Precision Boost 2 technology in the 2. gen. ThreadRipper line.

In the lecture we will not detail the enhancements of the 2. gen. ThreadRipper line.

Implementation of Precision Boost 2 in the 2. gen. ThreadRipper line

Precision Boost 2 has been introduced in a number of series as follows:

- Ryzen mobile APU series (Zen (Raven Ridge)-based, 14 nm, 10/2017)
- Ryzen DT APU series (Zen+ (Pinnacle Ridge)-based, 12 nm 4/2018)
- 2. gen. Ryzen DT series (Zen+ (Pinnacle Ridge)-based, 12 nm, 4/2018) and
- 2. gen. ThreadRipper series, (Zen+ (Pinnacle Ridge)-based, 12 nm, 8/2018).

Boost behavior in the original Precision Boost technology -1

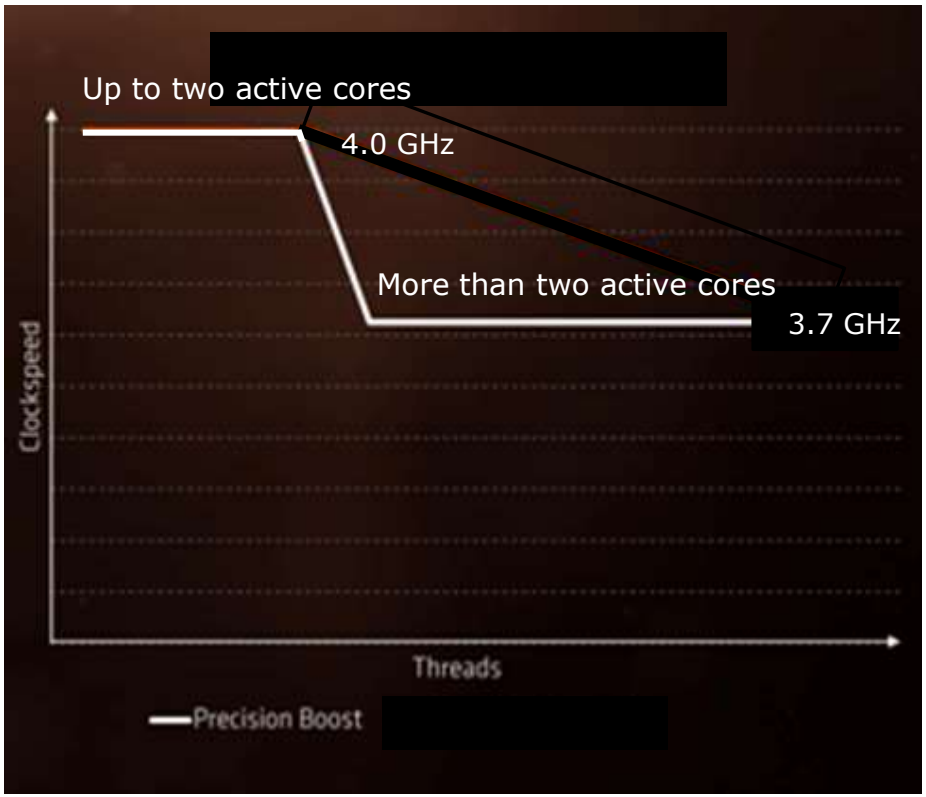

Figure: Boost behavior of the Ryzen DT and mobile lines [38]

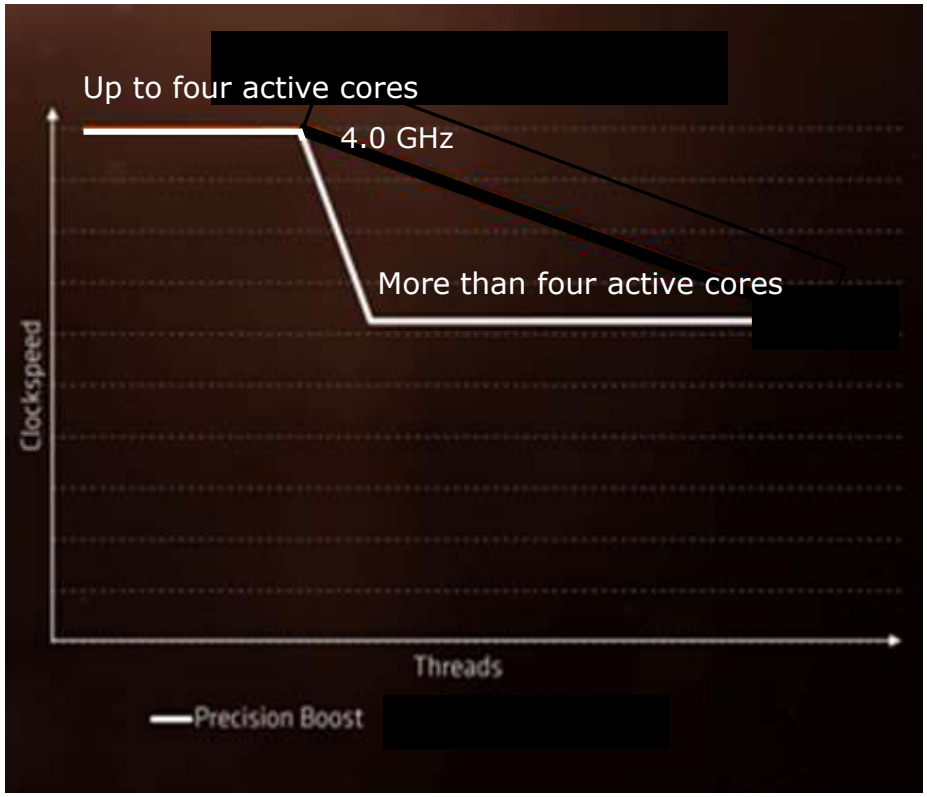

Figure: Boost behavior of the ThreadRipper line (based on [38])

*

Precision Boost 2 as implemented in the 2. gen. TR 2990WX [73]

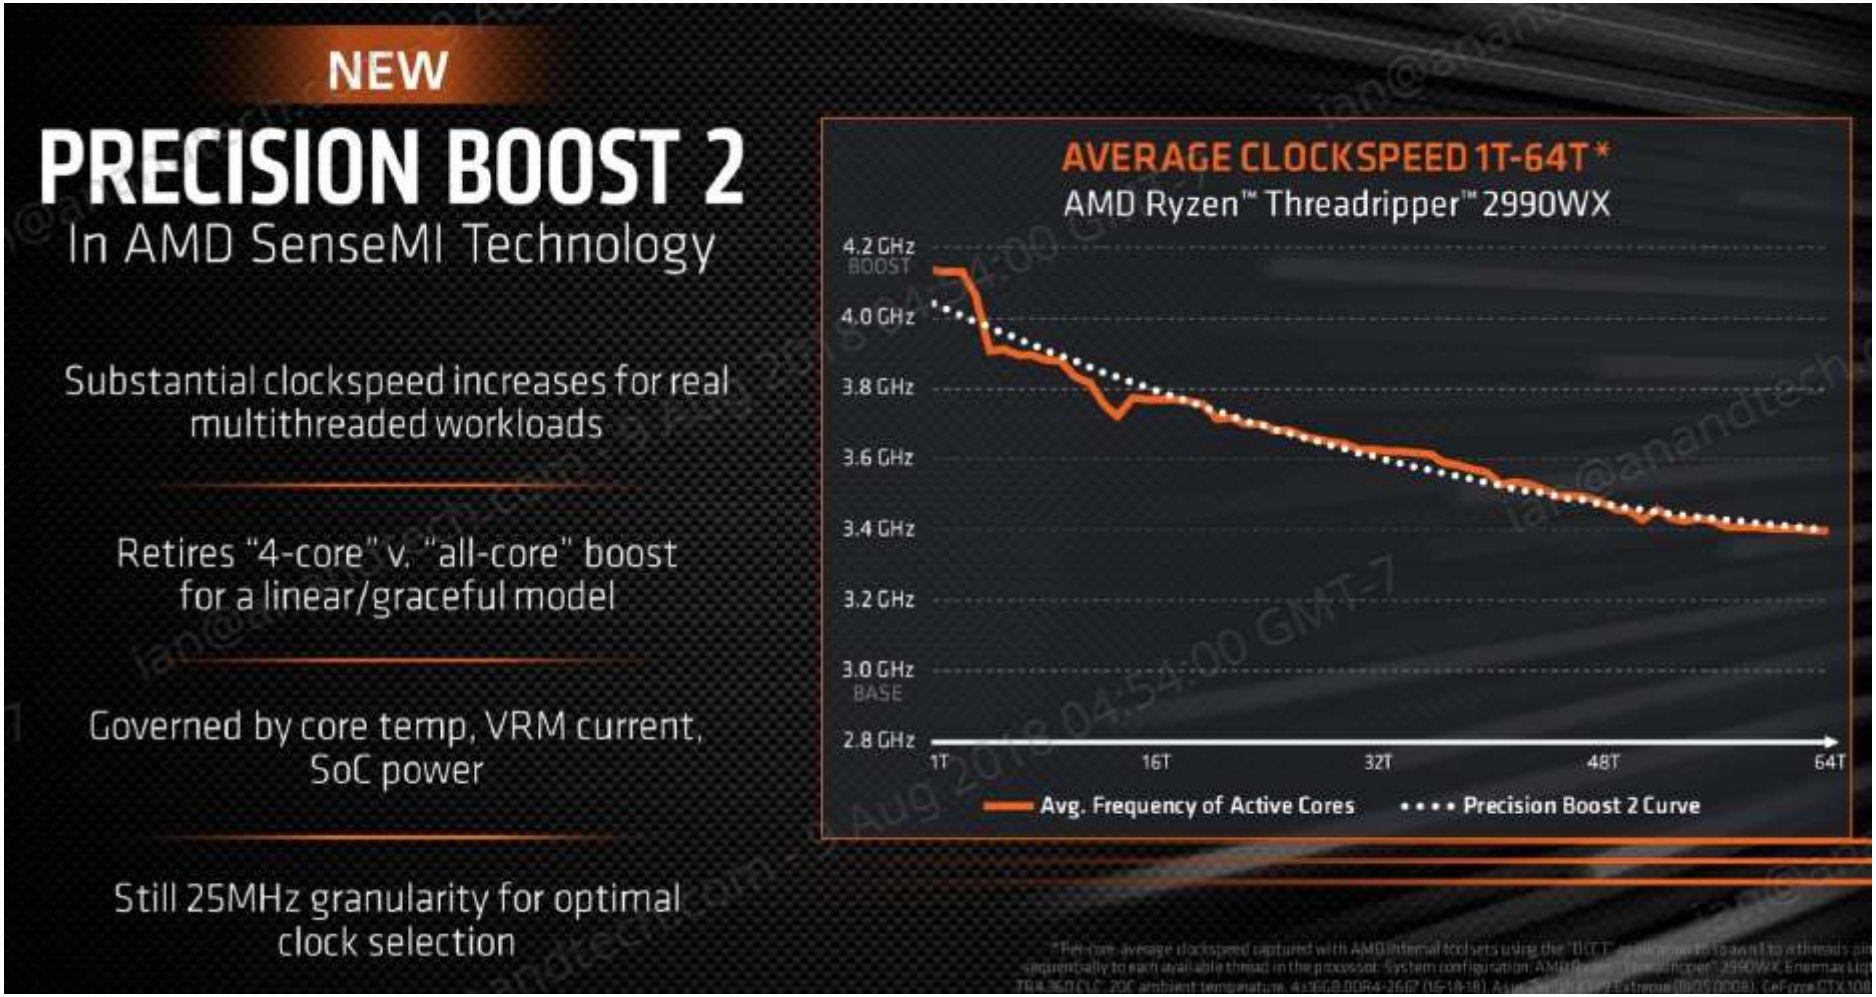

Benchmark results for the PCMark10 Extended Score [73]

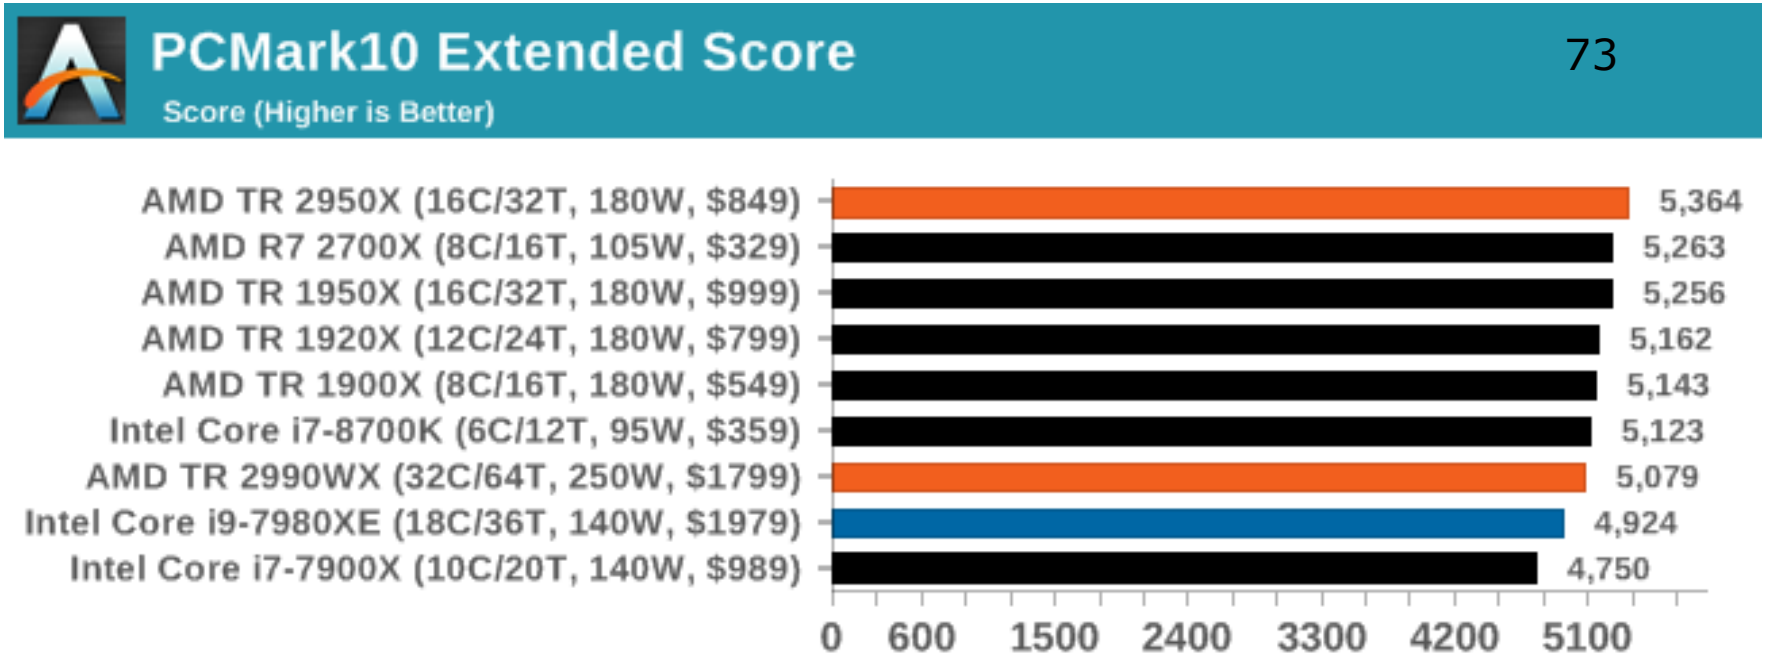

Remark

PCMark covers about 14 different areas, including application startup, web, spreadsheets, photo editing, rendering, video conferencing, and physics.

Benchmark results for the single thread Geekbench 4 [73]

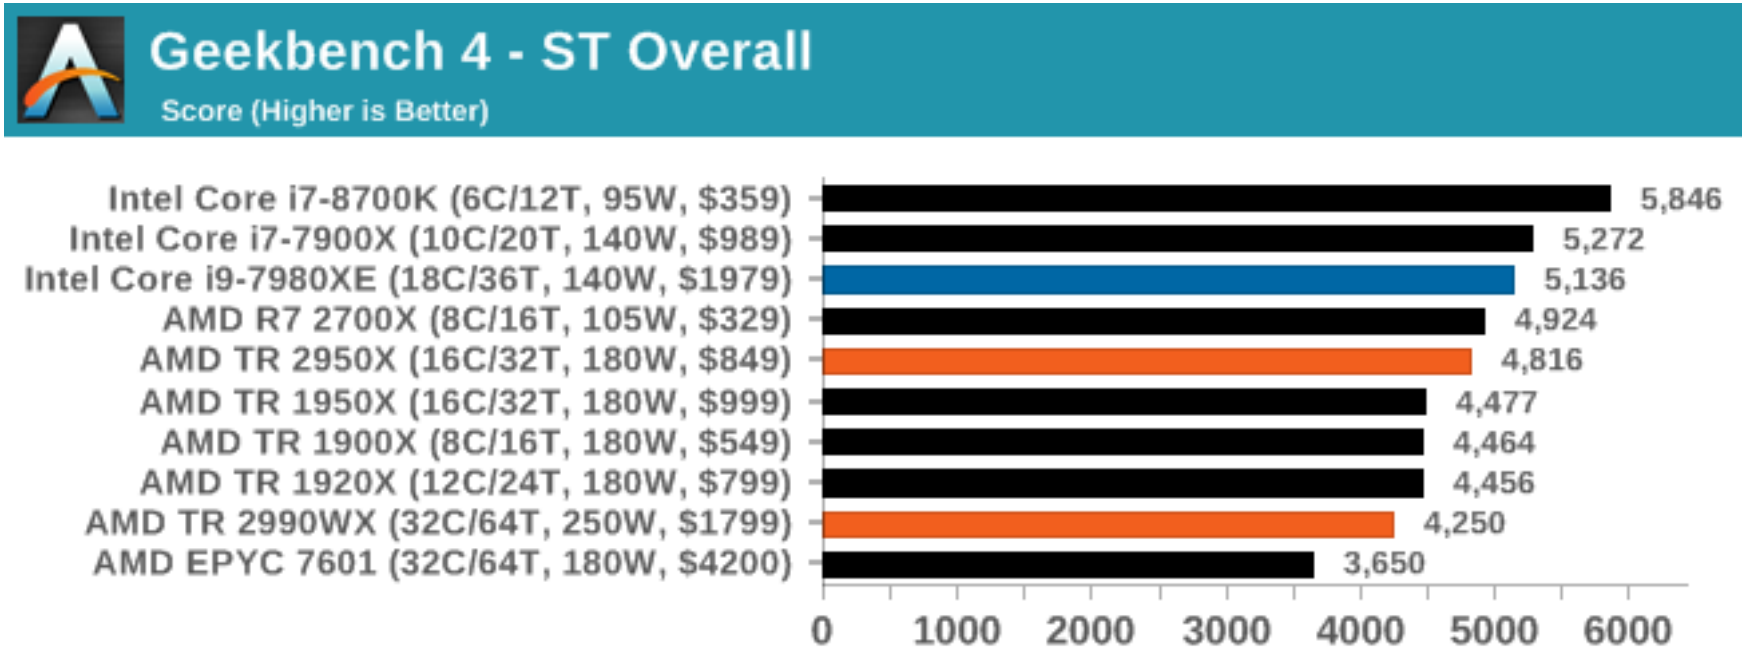

Remark

Geekbench 4 is a common tool for cross-platform testing between mobile, PC, and Mac.

It is an advanced tool in synthetic testing across a range of algorithms assessing peak throughput.

Tests include encryption, compression, fast Fourier transform, memory operations, n-body physics, matrix operations, histogram manipulation, and HTML parsing.

Benchmark results for the multi thread Geekbench 4 [73]

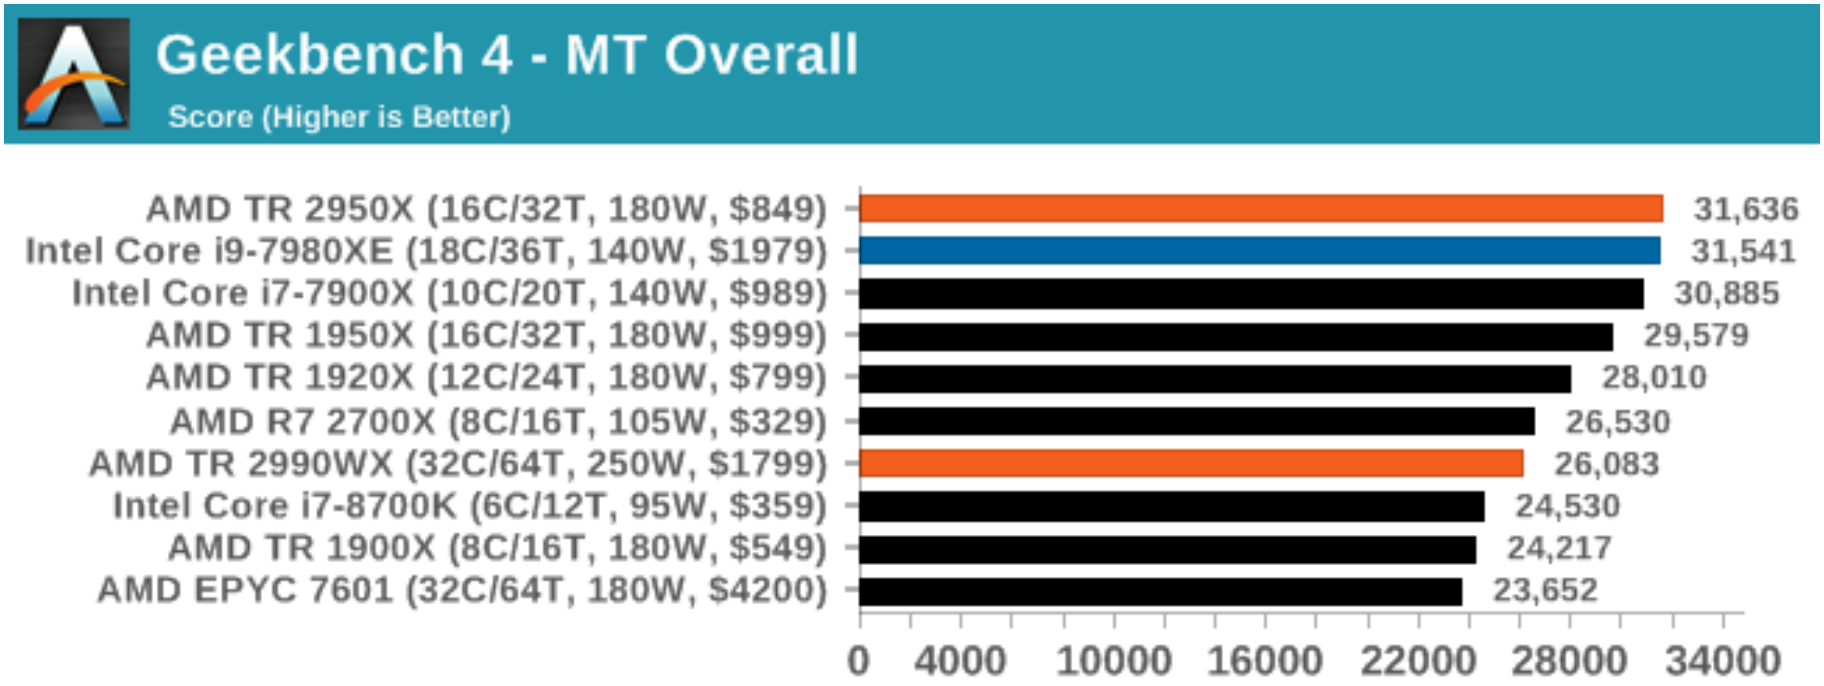

As the benchmark results show Intel retains further on their single thread advantage, whereas high-end 2. gen. ThreadRipper models provide the highest performance in a large number of benchmarks.

Benchmark results for the multi-thread Cinebench R15 [89]

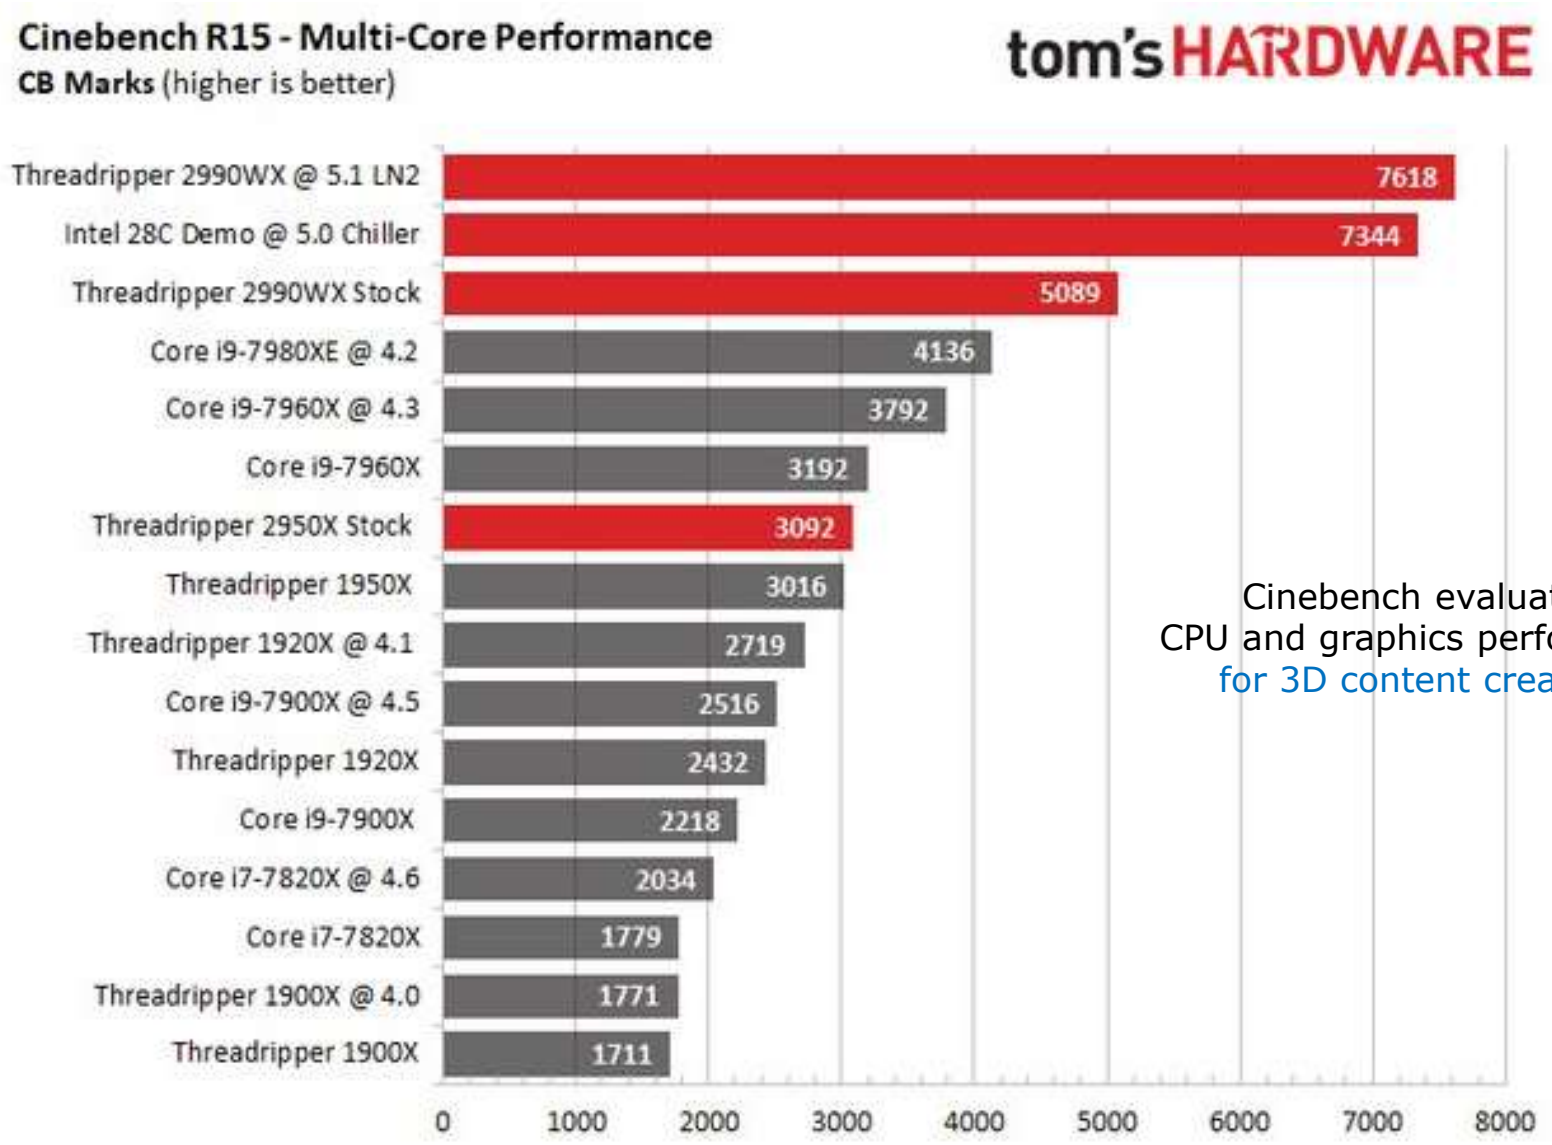

Cinebench evaluates
CPU and graphics performance
for 3D content creation.

Intel's response to the ThreadRipper threat – the Xeon W-3175X [96]

PROCESSOR NUMBER[1]	BASE CLOCK SPEED (GHZ)	INTEL® TURBO BOOST TECHNOLOGY 2.0 MAXIMUM SINGLE CORE TURBO FREQUENCY (GHZ)	CORES/ THREADS	TDP	INTEL® SMART CACHE	UNLOCKED[4]	PLATFORM PCIE LANES	MEMORY SUPPORT	STANDARD RAS SUPPORT	ECC SUPPORT
Intel® Xeon® W-3175X	3.1	4.3	28/56	255W	38.5 MB	✓	Up to 68	Six channels DDR4-2666	✓	✓

(TR 2990WX 3.0 4.2)

It is probably Intel 28C Demo processor shown @ 5 GHz overclocked, with a water chiller, in Computex in 06/2018

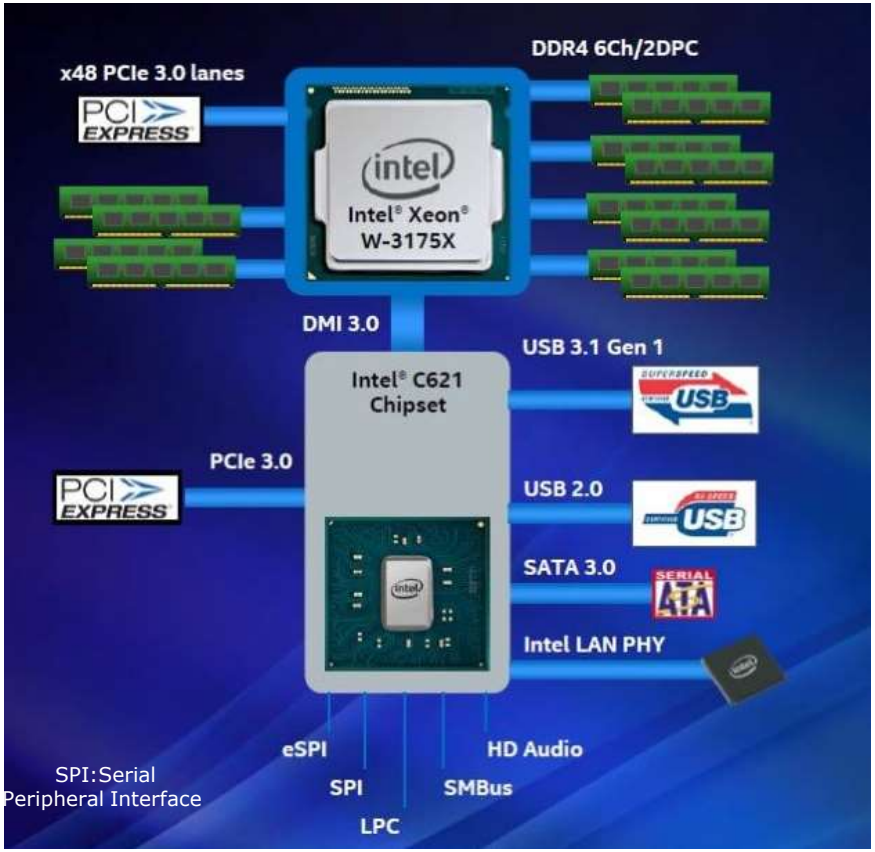

Announced in 10/2018

To be shipped in 12/2018

9. The Epyc server line

- 9.1 The 1. gen. Naples Epyc server line
- 9.2 The 2. gen. Rome Epyc server line

Will not be discussed

9. The Epyc server lines

AMD unveiled until now two generations of Epyc processor lines:

- the 1. generation, called Naples line, launched in 06/2017 and
- the 2. generation, called Rome line, introduced in 11/2018.

These two generations will be discussed subsequently.

9.1 The 1. gen. Naples Epyc server line

9.1 The 1. gen. Naples Epyc server line -1

Positioning AMD's 1. gen. Epyc server line in AMD's 1.gen. Zen-based lines

AMD's Zen-based processor lines

Ryzen Mobile (mobil)	**Ryzen** (DT)	**ThreadRipper** (HED)	**Epyc** (1S/2S server)
Single CCX + Vega GPU	Zeppelin chip (2x CCX, no GPU)	2 Zeppelin chips interconnected by the IF, implemented as an MCM	4 Zeppelin chips interconnected by the IF, implemented as an MCM

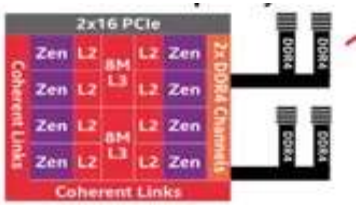

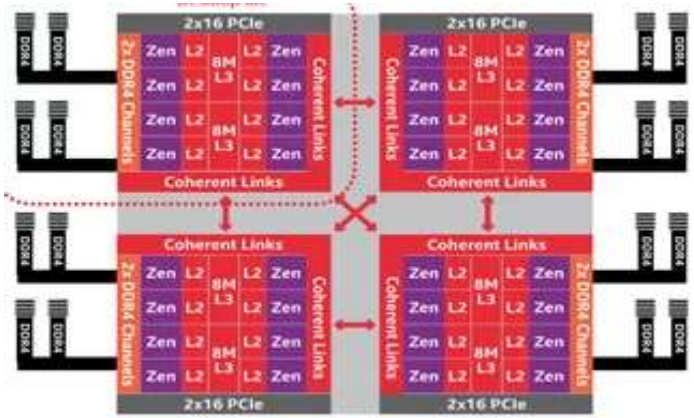

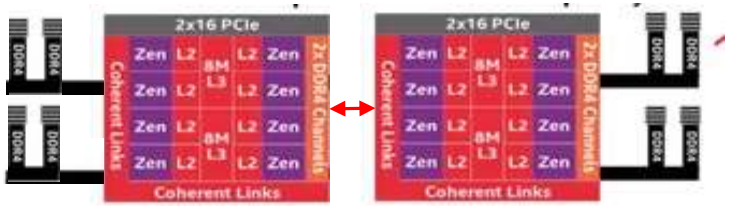

Drawings based on [11]

The 1. gen. Naples Epyc server line -2

- Launched in 6/2017 based on the Zen core, manufactured on 14 nm technology.
- It covers 1S and 2S servers.

The reason why AMD focuses on up to 2S servers is that these servers cover recently over 90 % of the server market, as indicated below.

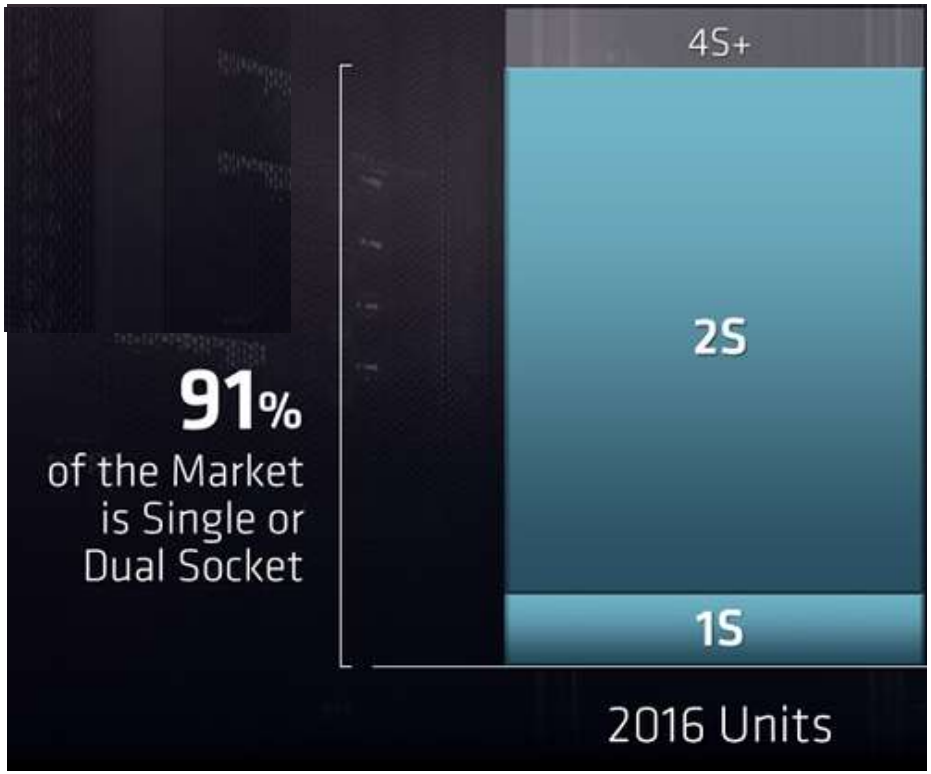

Figure : Recent market share of 1S and 2S servers [29]

The EPIC package [21]

Package size: 58 x 75 mm, LGA 4094 socket (SP3).

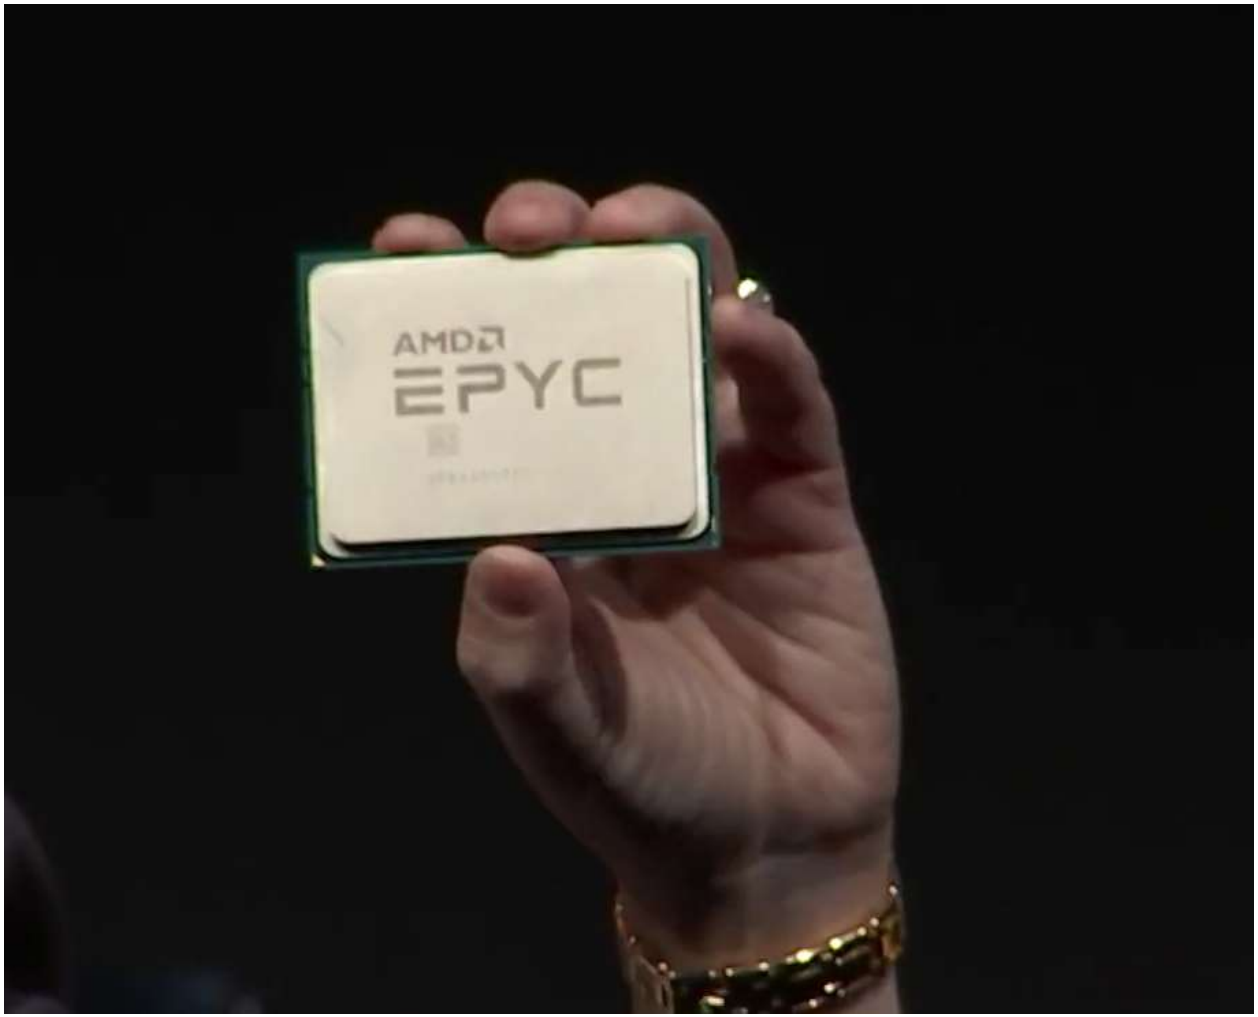

Main features of AMD's Epyc line introduced in 2017

AMD's Zen-based processor lines introduced in 2017

	Ryzen Mobile APU (Raven Ridge)	**Ryzen DT** (Summit Ridge)	**ThreadRipper** (Whitehaven)	**Epyc** (Naples)
Market segment	Mobile	Desktop platform	HED	1S/2S server
µarch./Technology	Zen 14 nm	Zen, 14 nm	Zen, 14 nm	Zen, 14 nm
Launched models	Ryzen 7 2700U Ryzen 5 2500U (10/2017)	Ryzen 7 (3/2017) Ryzen 5 (4/2017) Ryzen 3 (7/2017)	1950X/1920X/1900X (8/2017)	Series 7000 (6/2017)
Layout	CCX + Vega 8/10	Zeppelin die with 2x CCX	MCM (2x Zeppelin die)	MCM (4x Zeppelin die)
Integrated GPU	Yes	No	No	No
Core count	4	4/6/8	8/12/16	8/16/24/32
SMT	SMT	SMT (except Ryzen 3)	SMT	SMT
Mem. channels/rate	2xDDR4-2400	2xDDR4-2666	4xDDR4-2666	8xDDR4-2666
PCIe 3.0 lanes	??	16xPCIe 3.0	60xPCIe 3.0	128 for 1S servers 64 for 2S servers
TDP	15 W	65/95 W	180 W	120/170/180 W
Socket	AM4 (1331)	AM4 (1331)	TR4 (SP3r2) (4094)	SP3 (4094)
Chipset	SoC	300-series	X399	No chipset, SOC

Die-to-die interconnections in a 1S EPYC server processor

A 1S EPYC server processor is built up of 4 Zeppelin dies interconnected by a crossbar Infinity Fabric switch and implemented as an MCM (Multi-chip_Module), as shown in the Figure below.

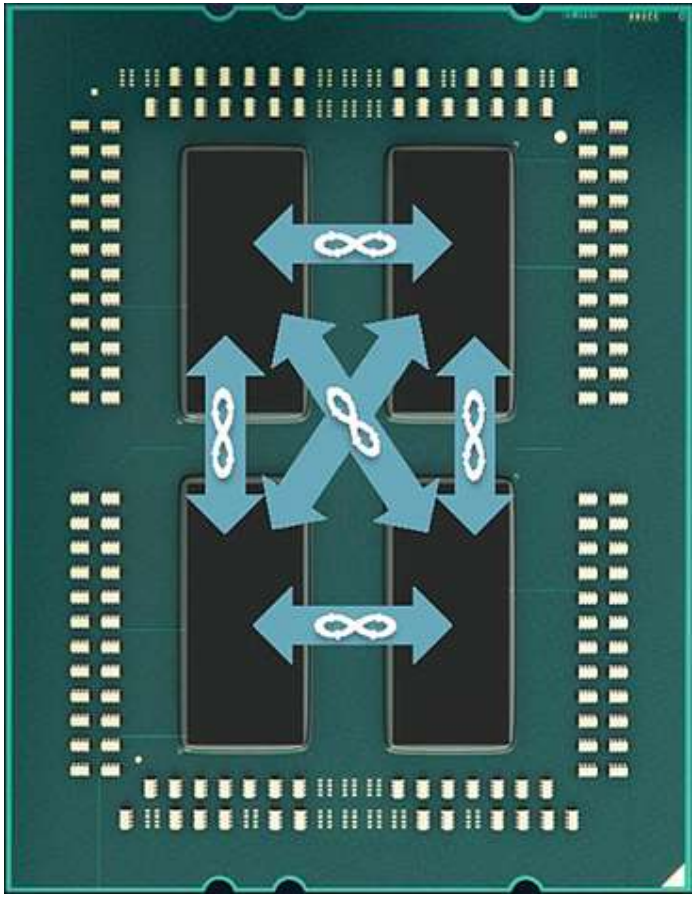

32-bit per link per direction
interconnecting the chips
in a crossbar fashion
by single-ended signals
with a bandwidth of 42.8 GB/s
per link in each direction.

Figure: Die interconnections in an 1S EPYC server [21]

A more detailed presentation of a 1S Epyc server processor [7]

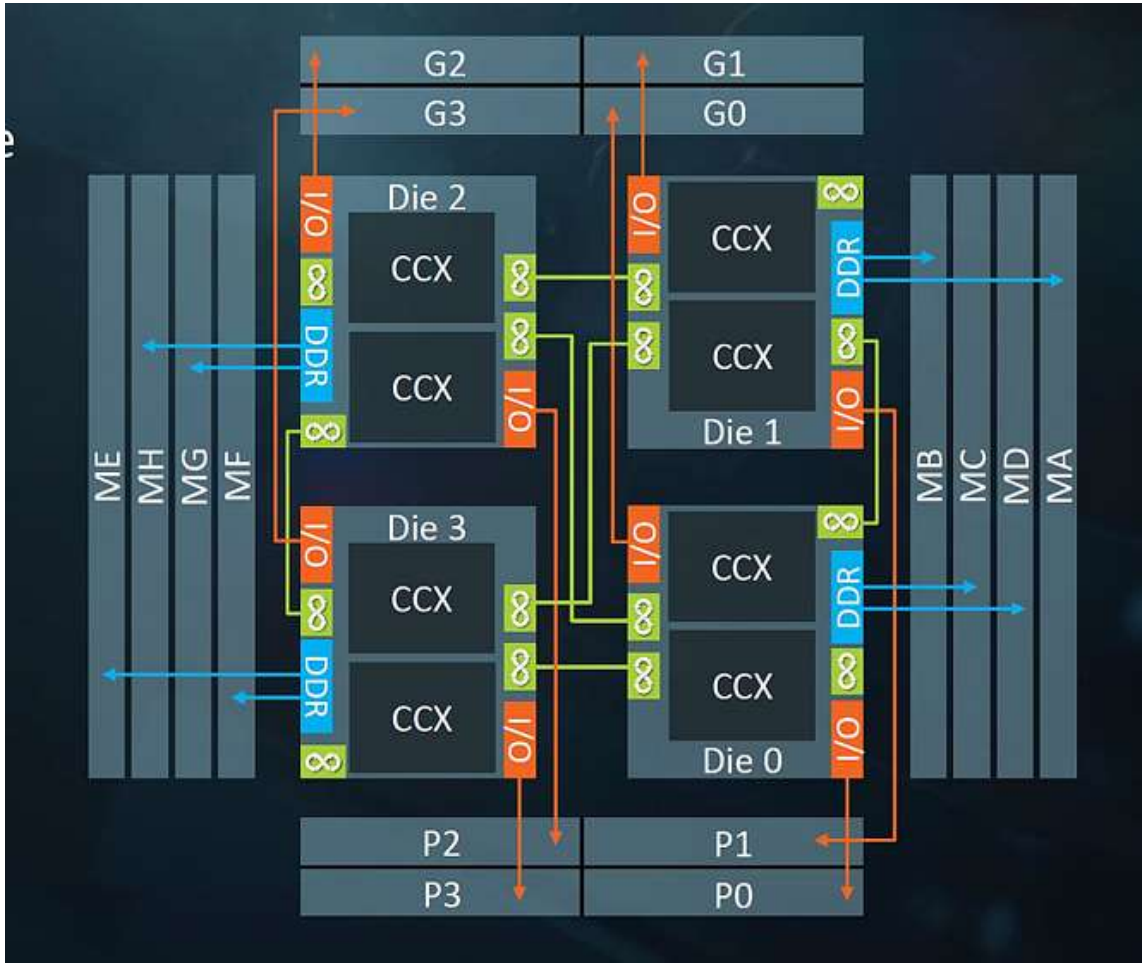

G0 - G3: 4 x 16 SERDES lanes
(in each direction)
Used for socket-to-socket
communication

MA-MH: 8 x 72-bit DDR4
DRAM channels
Up to 2 TB DDR4 memory

∞: Die-to-die IF links
16-bit per link per direction
interconnecting the chips
in a crossbar fashion
by single-ended signals
(2 pJ/bit TDP)
with a bandwidth of 42.6 GB/s
per link in each direction.

P0 - P3: 4 x 16 SERDES lanes
(in each direction)
Used to implement PCIe 3.0 lanes

Dies and die interconnections within a 1S EPYC server [30]

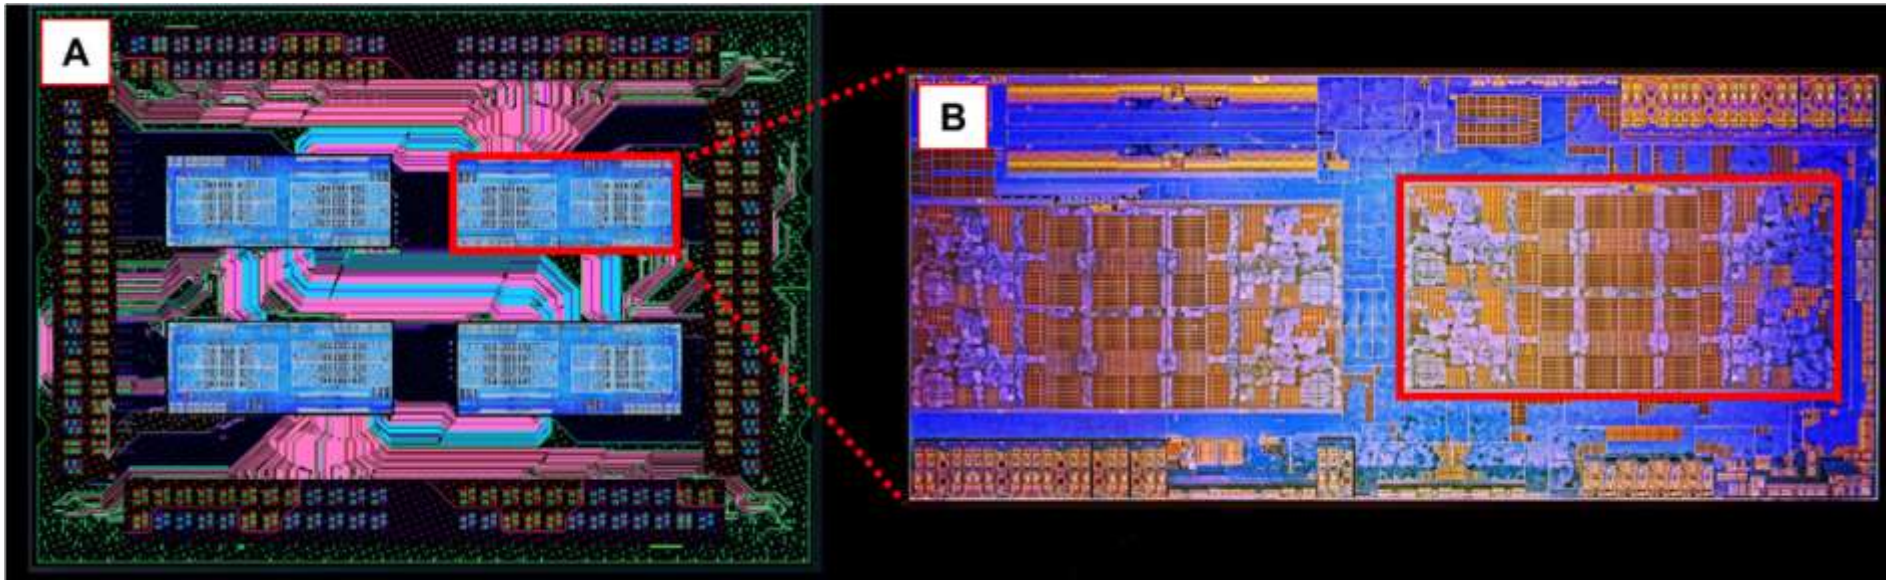

A: 4 Zeppelin dies/socket B: A single Zeppelin die

MCM-to-MCM interconnections in a 2S EPYC server processor

A 2S EPYC server processor is built up 2 EPYC processors that are interconnected by the Infinity Fabric (IF), as shown in the Figure.

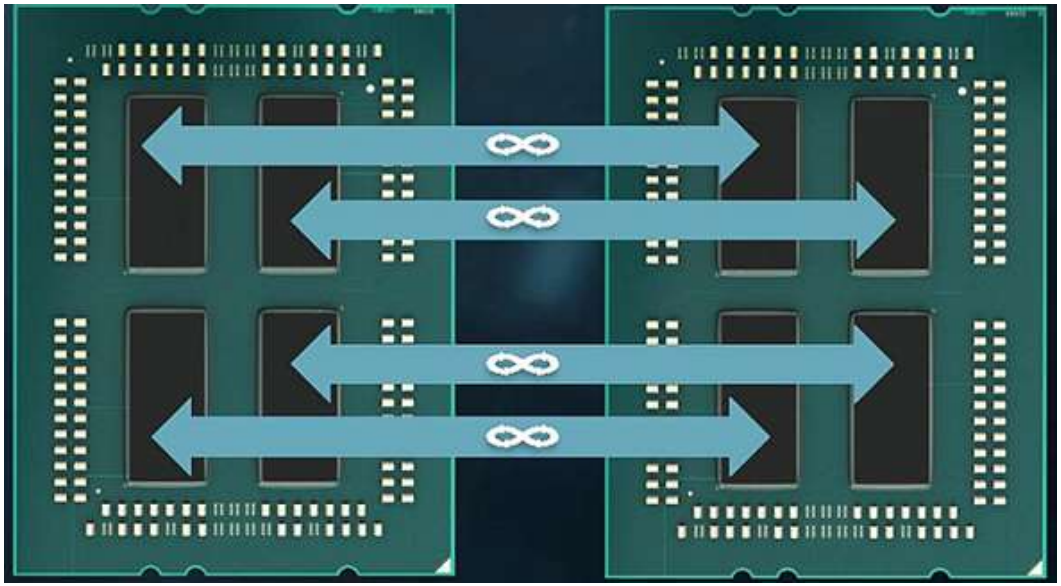

Figure: Built up of a 2S EPYC server [21]

Up to 2x 32 cores,

2x 4 channel DDR4 memory,

up to 2x 2 TB memory,

2x 64 PCIe 3 lanes

4x die-to-die IF links
16 differential lanes per direction
(9 pJ/bit TDP)
interconnecting related dies
on the two sockets,
37.9 GB/s bandwidth
bidirectional.

(Up to two hops are needed
for die-to-die transfers).

Layout of a dual socket Epyc system [88]

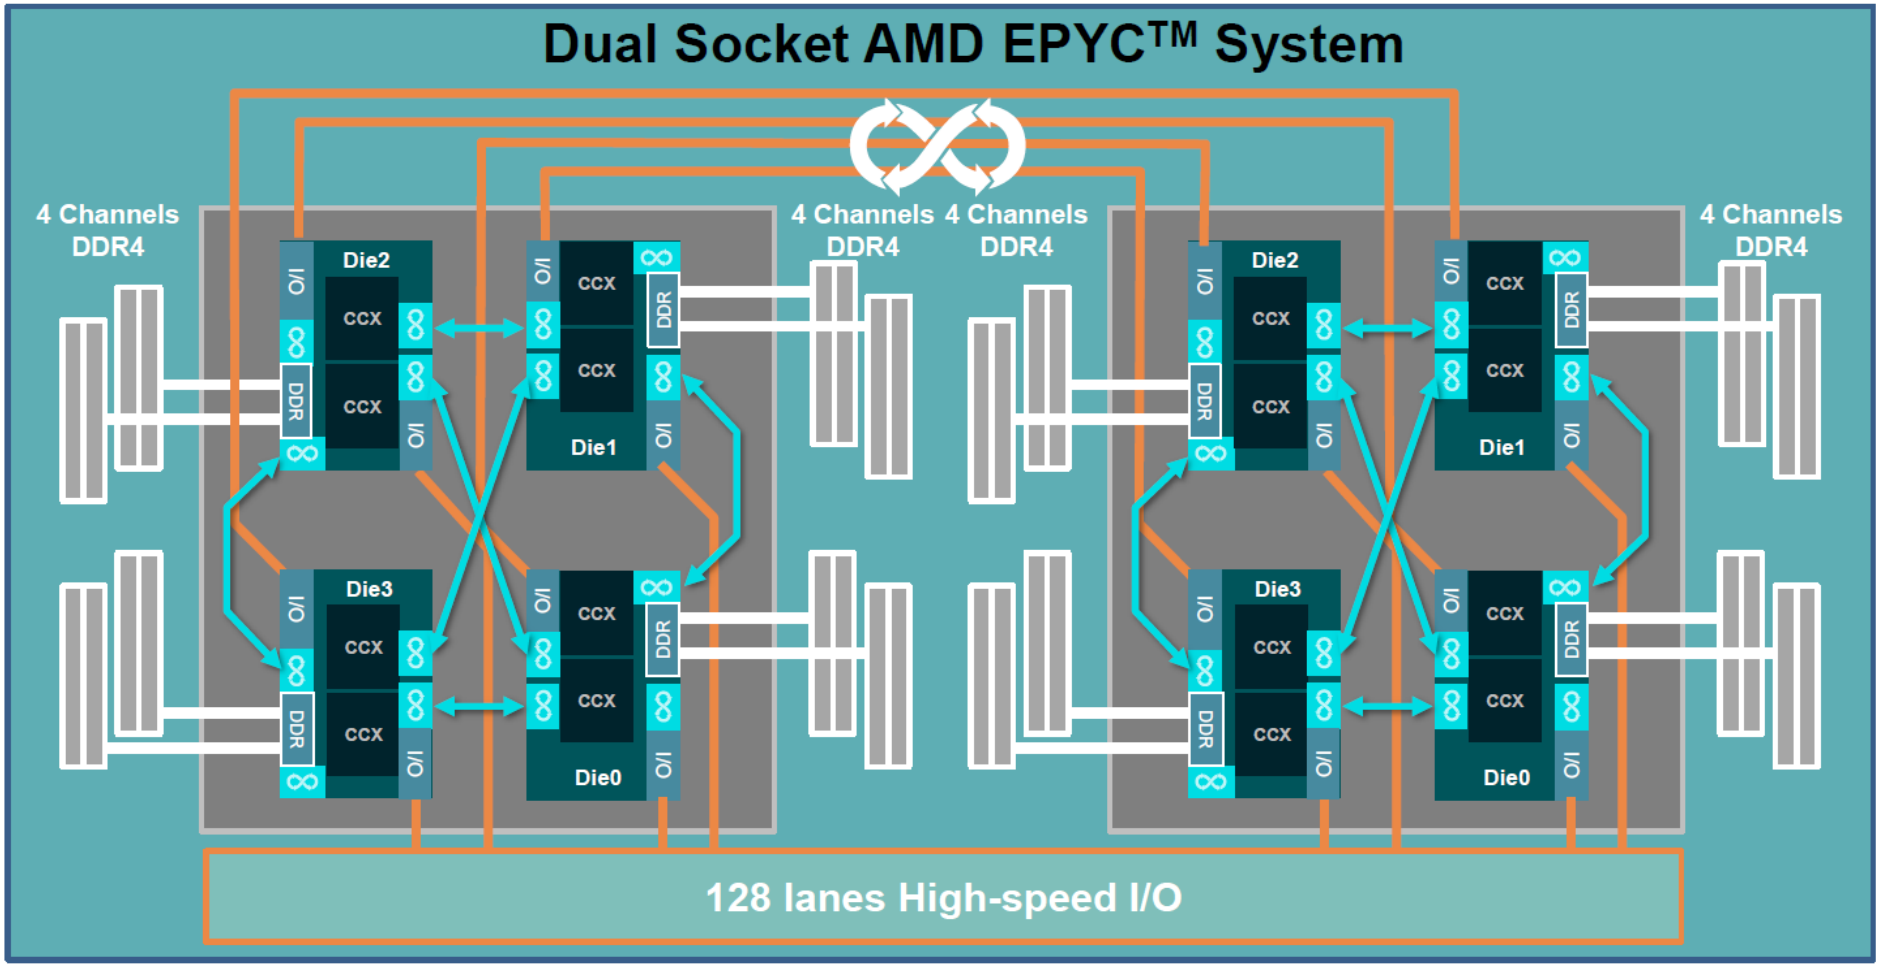

MCM-to MCM IF interconnections in a 2S Epyc server processor - detailed [43]

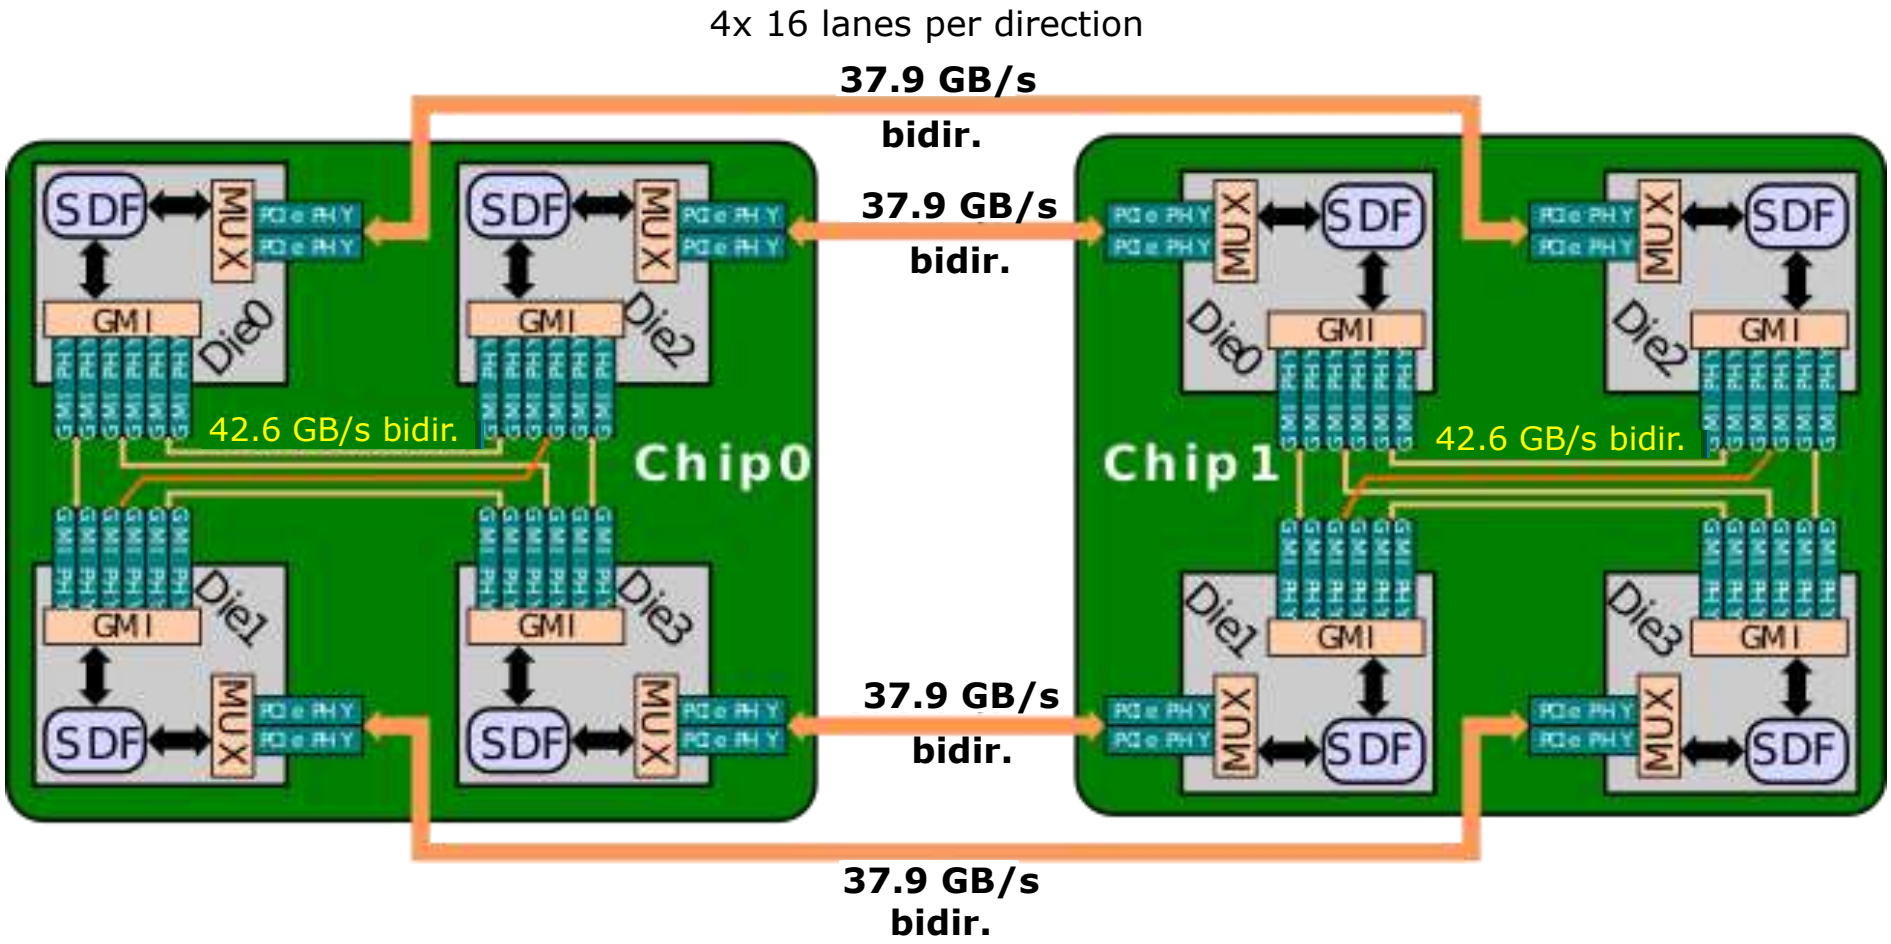

GMI: Global Memory Interconnect
SDF: Scalable Data Fabric

Comparing memory latencies in Intel's Xeon and AMD's Epyc [70]

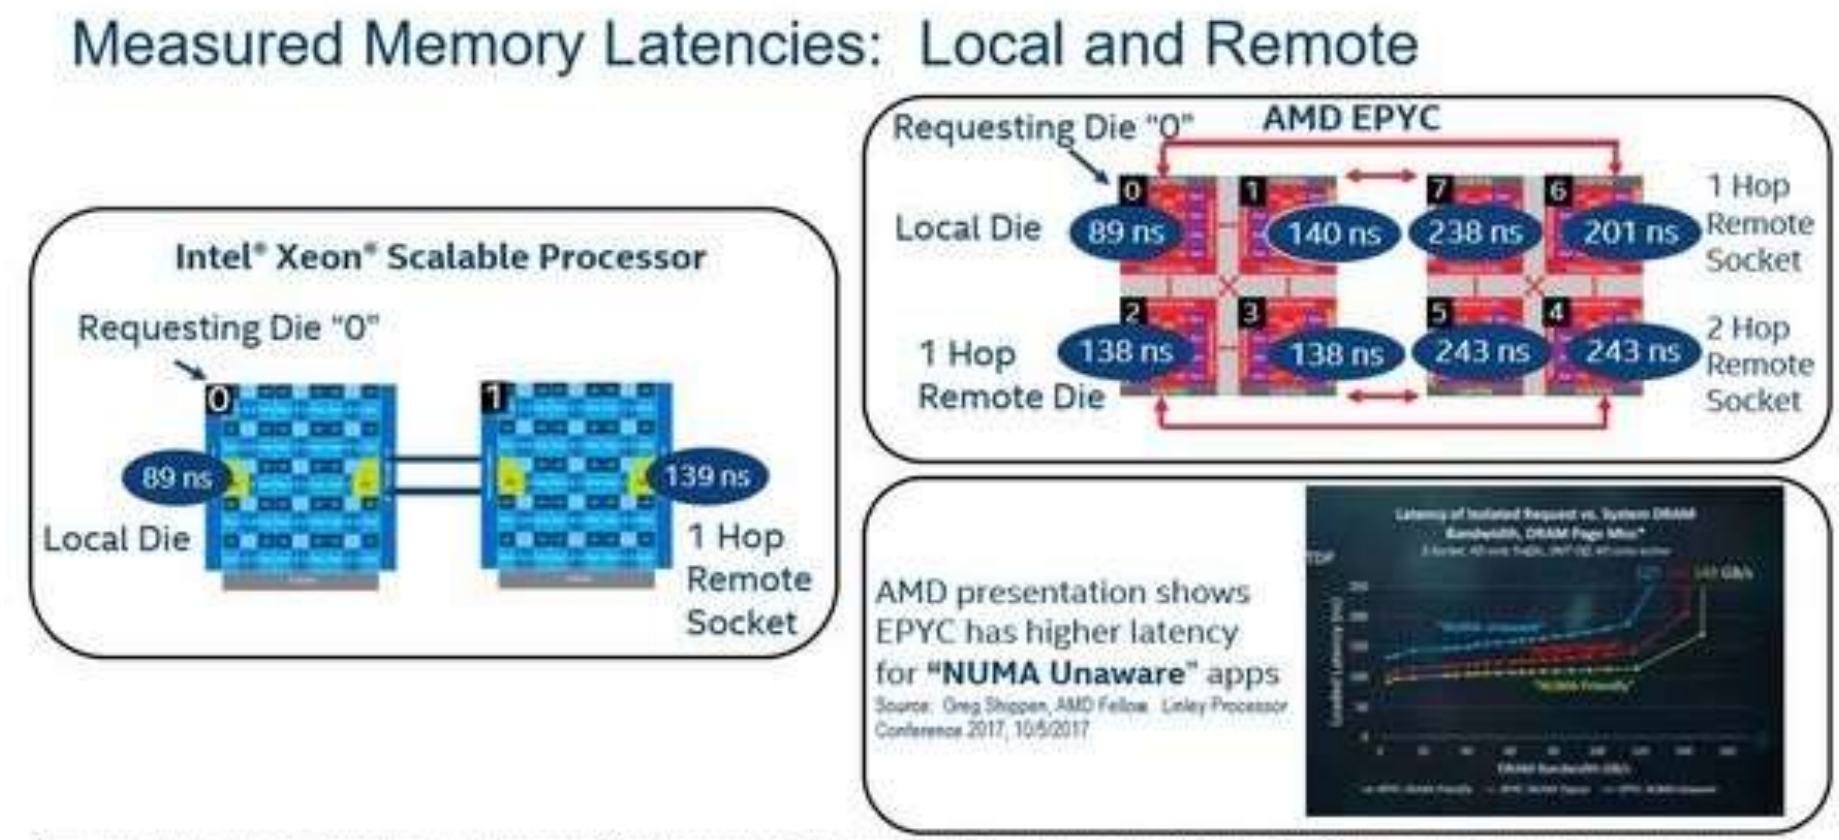

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. Intel does not control or audit the design or implementation of third party benchmark data or Web sites referenced in this document. Intel encourages all of its customers to visit the referenced Web sites or others where similar performance benchmark data are reported and confirm whether the referenced benchmark data are accurate and reflect performance of systems available for purchase.
Refer to Configuration, Configuration: Memory Latency Local and Remote, slide 19

Restrictions on the number of PCIe 3.0 lanes per socket -1

An Epic processor has 8 x 16 serialization-deserialization circuits (called SERDES) that can be used either for implementing the interconnection lanes between two sockets of a 2S server processor or PCIe 3.0 lanes, as indicated below.

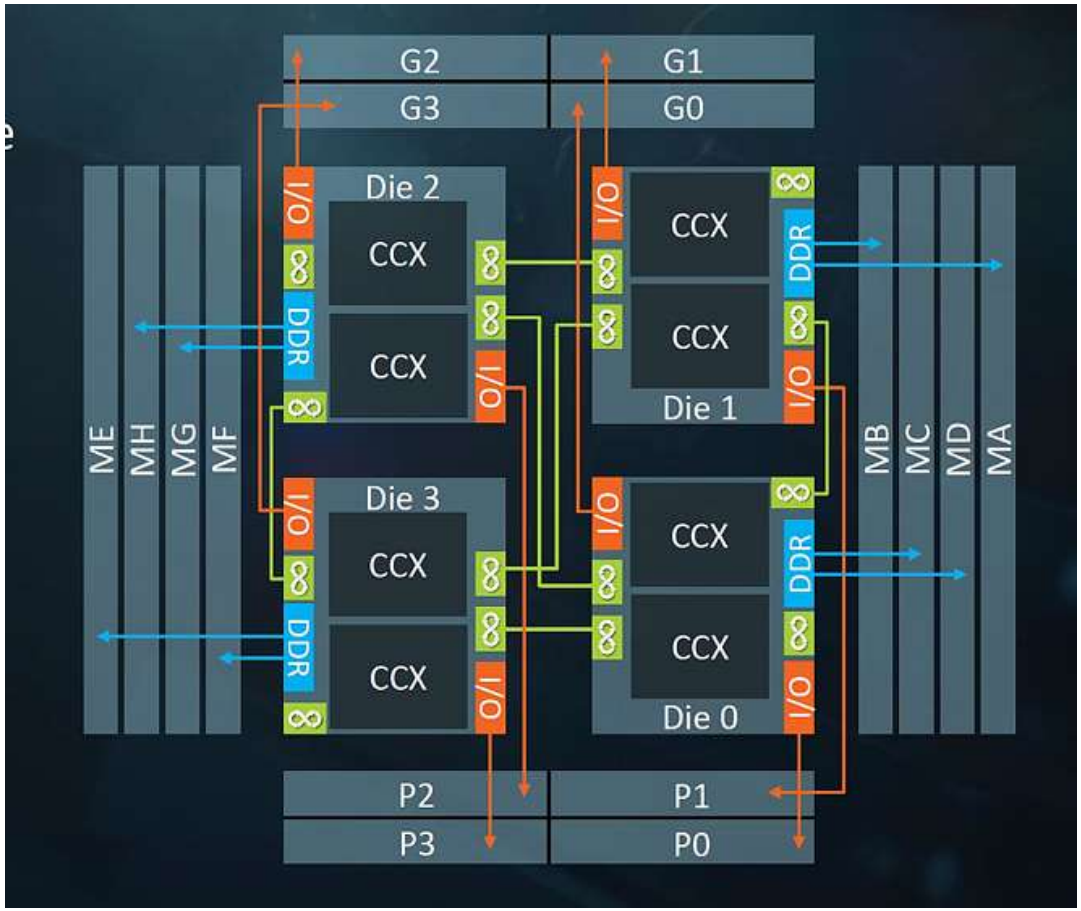

G0 - G3: 4 x 16 SERDES circuits
(in each direction)
Used for PCIe 3.0 lanes in 1S
and for MCM-to-MCM
interconnection lanes
in a 2S configuration

P0 - P3: 4 x 16 SERDES lanes
(in each direction)
Used for PCIe 3.0 lanes

Figure: High-level block diagram of a Ryzen processor

Restrictions on the number of PCIe 3.0 lanes per socket -2

- In a 1S Epyc processor all SERDES circuits are available for PCIe 3.0 lanes, thus 1S Epyc server processors can provide 6x16 = 128 PCIe 3.0 lanes.
- By contrast, in 2S Epyc processors 4x16 lanes are needed to interconnect both sockets, then only 4x 16 = 64 PCIe 3.0 lanes can be implemented per socket.
- This design approach implies however, a serious restriction on the feasible server configurations since a 4S configuration would require already 2x 64 lanes, i.e. all 128 available SERDES circuits for interconnecting 4 sockets, consequently, for a 4S configuration no PCIe 3 lanes could be provided.

  In other words, the chosen IF solution allows only to implement 1S or 2S configurations.
- AMD is full aware of this restriction and did not announce any plans to develop 4S or 8S servers while employing IF.

  This is understandable due to the low market share of 4S and 8S servers vs. 1S and 2S servers.

Width and bit rate of Infinity Fabric links in an Epyc processor [3]

	Cross-Sectional Links	Lanes per Link	Total Lanes	One-Way Lanes	Bit Rate Per Lane (Mb/sec)	Byte Rate Per Lane (MB/sec)	Byte Rate Per Link (GB/sec)	Cross Sectional (GB/sec)
Die-to-Die	4	32	128	16	5325	666	10.65	42.6
MCM-to-MCM	4	16	64	8	9475	1184	9.48	37.9
Configurable			64					
Total per Die			256					

Remarks

1) From the 4 available Cross-Sectional Links per die only 3 are needed to interconnect the actual die with the further three dies of the processor.

2) The Die-to-Die interconnects are actually not differential lanes but single-ended lines to reduce power consumption (resulting in 2 pJ/bit vs. 9 pJ/bit TDP.

Additional features of EPYC-7000-series processors

a) Workload aware power management

b) Per core frequency and voltage scaling

c) Integrated secure processor

a) Workload aware power management of the EPYC line [21]

WORKLOAD AWARE POWER MANAGEMENT

- Many workloads utilize all the cores to do sporadic tasks which are not overly latency sensitive
- Default behavior has cores running at a high frequency during the active phases, wasting power
- The AMD workload aware efficiency algorithm recognizes these cases and dynamically optimizes operation for up to 10% performance/Watt gains*

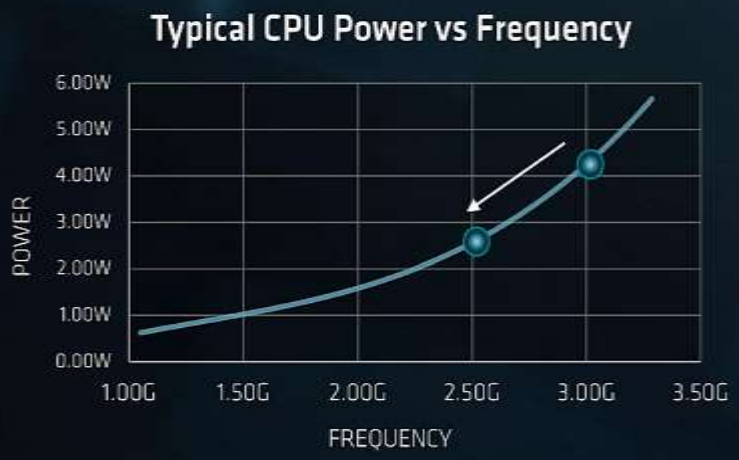

*See endnotes

AMD

b) Per-core frequency and voltage scaling [21]

PER-CORE LINEAR REGULATION

- With 32 cores across 4 different die, running all cores at the voltage required by the slowest would waste power
- Our per-core regulator capabilities enable AVFS-informed adaptation of each core's voltage to its particular characteristics and operating environment
- The result is significant core power savings and variation reduction

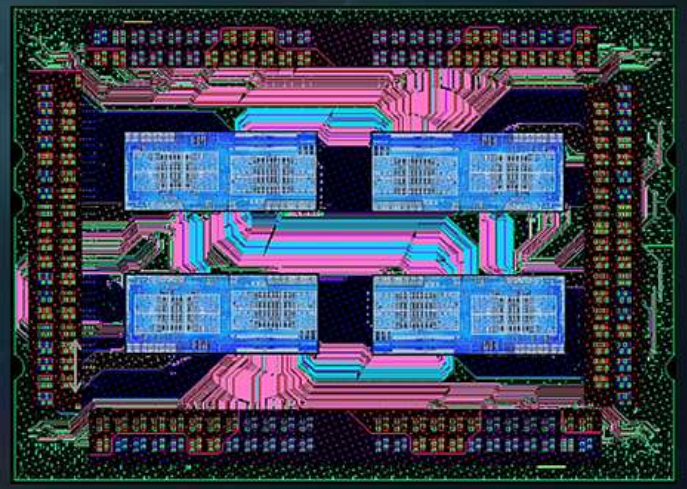

LDO: Low Drop Out Voltage Regulator

c) Integrated secure processor [31]

Main features of AMD's 1S EPYC-7000 series processors [21]

	Cores Threads	Frequency (GHz)			L3	DRAM	PCIe	TDP	Price
		Base	All	Max					
EPYC 7551P	32 / 64	2.0	2.6	3.0	64 MB	8 Ch DDR4 up to 2666 MT/s	8 x16 PCIe	180W	$2100
EPYC 7401P	24 / 48	2.0	2.8	3.0				155W/170W	$1075
EPYC 7351P	16 / 32	2.4	2.9					155W/170W	$750

Note that the P tag means 1S configuration.

The last numeric character in the model designation, i.e. 1 indicates generation 1.

Main features of AMD's 2S EPYC-7000 series processors [21]

	Cores Threads	Frequency (GHz)			L3	DRAM	PCIe	TDP	Price
		Base	All	Max					
EPYC 7601	32 / 64	2.20	2.70	3.2	64 MB	8 Ch DDR4 up to 2666 MT/s	8 x16 PCIe	180W	$4200
EPYC 7551	32 / 64	2.00	2.55	3.0				180W	>$3400
EPYC 7501	32 / 64	2.00	2.60	3.0				155W/170W	$3400
EPYC 7451	24 / 48	2.30	2.90	3.2				180W	>$2400
EPYC 7401	24 / 48	2.00	2.80	3.0				155W/170W	$1850
EPYC 7351	16 / 32	2.40	2.9					155W/170W	>$1100
EPYC 7301	16 / 32	2.20	2.7					155W/170W	>$800
EPYC 7281	16 / 32	2.10	2.7		32 MB			155W/170W	$650
EPYC 7251	8 / 16	2.10	2.9					120W	$475

Layout of a 2S EPYC server [29]

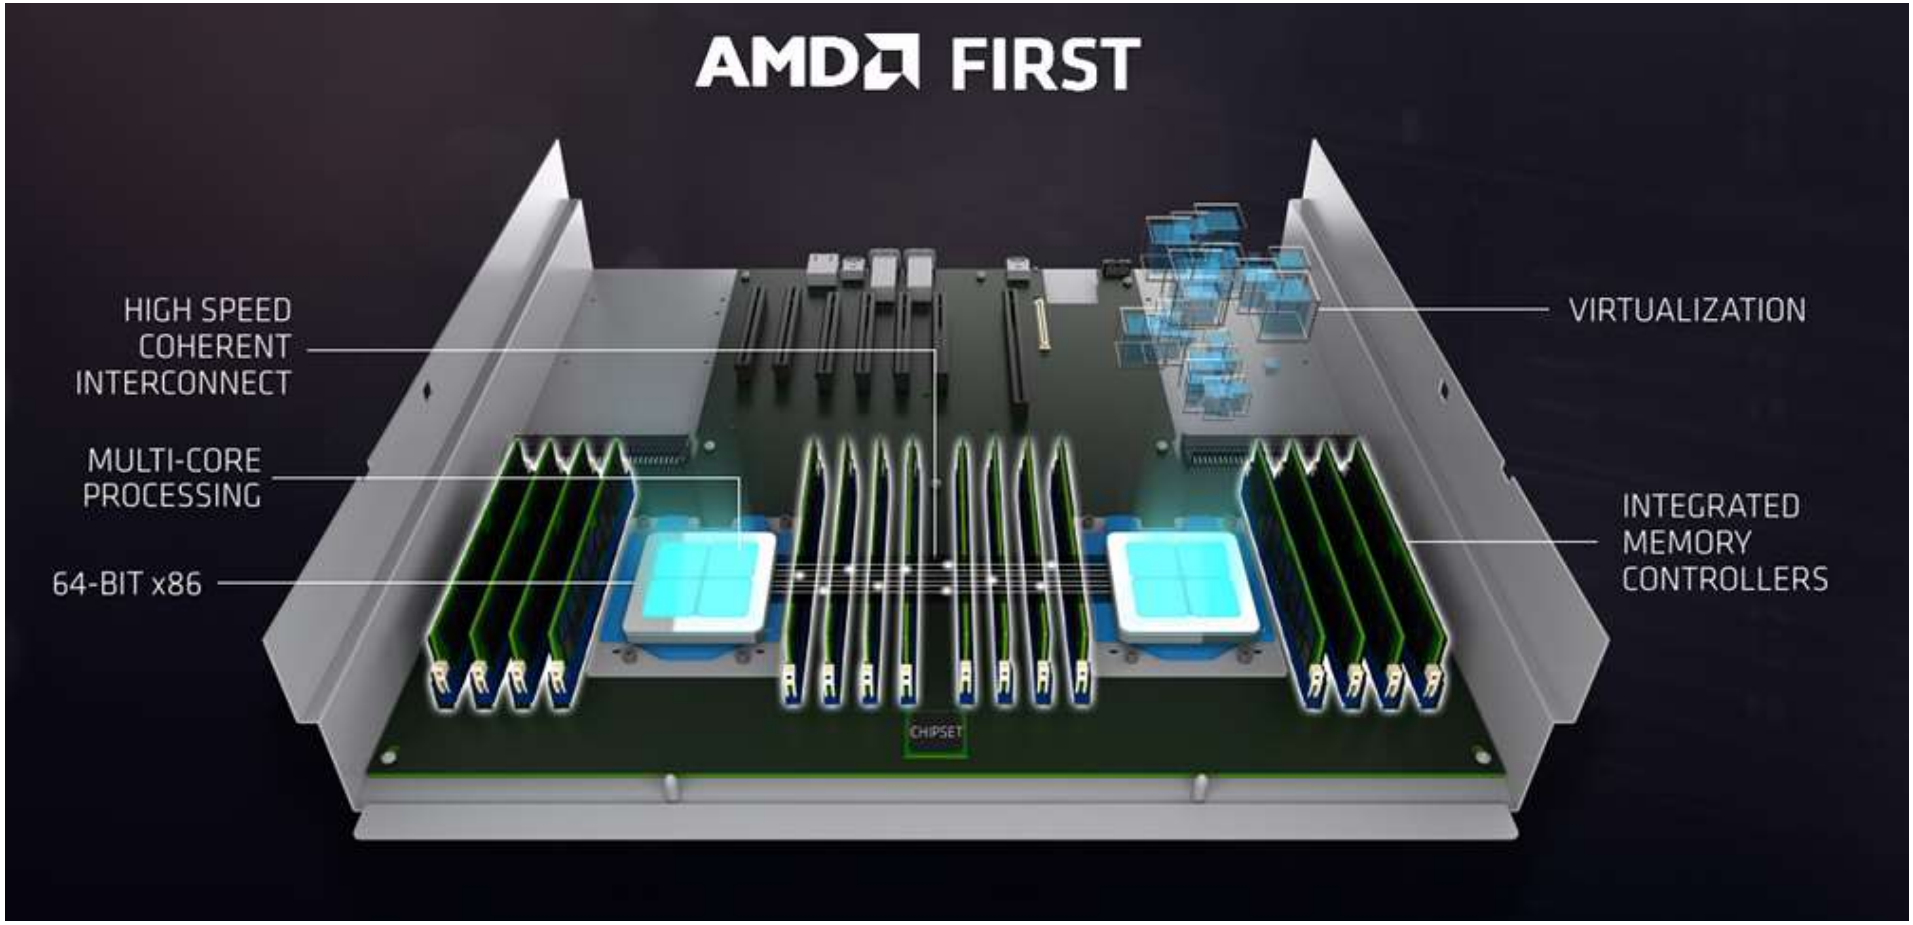

Note that it does not require a PCH.

Example EPYC 2S motherboard (SuperMicro's H11DSi) -1 [33]

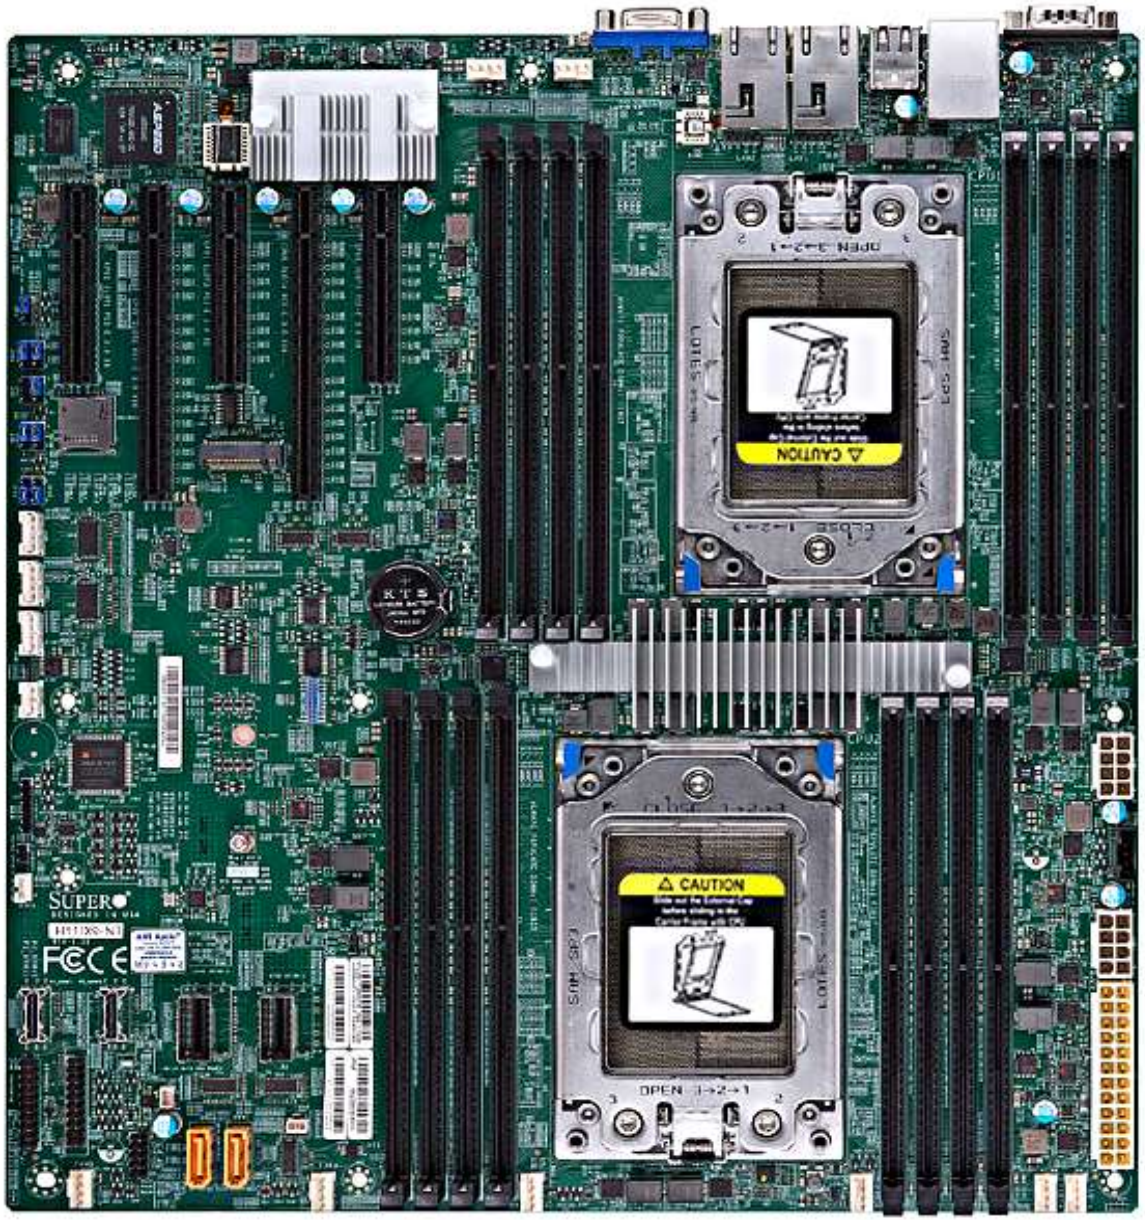

1. Dual AMD EPYC™ 7000-Series Processors
2. 2TB Registered ECC DDR4 2666MHz SDRAM in 16 DIMMs
3. Expansion slots:
   2 PCI-E 3.0 x16
   3 PCI-E 3.0 x8

4. 10 SATA3, 1 M.2, 2 SATA DOM
5. Dual Gigabit Ethernet LAN Ports
6. ASPEED AST2500 BMC graphics
7. Up to 2 USB 3.0 ports
   Up to 4 USB 2.0 ports
8. 8 4-pin PWM Fan & Speed control

Example block diagram of a 2S Epyc server [71]

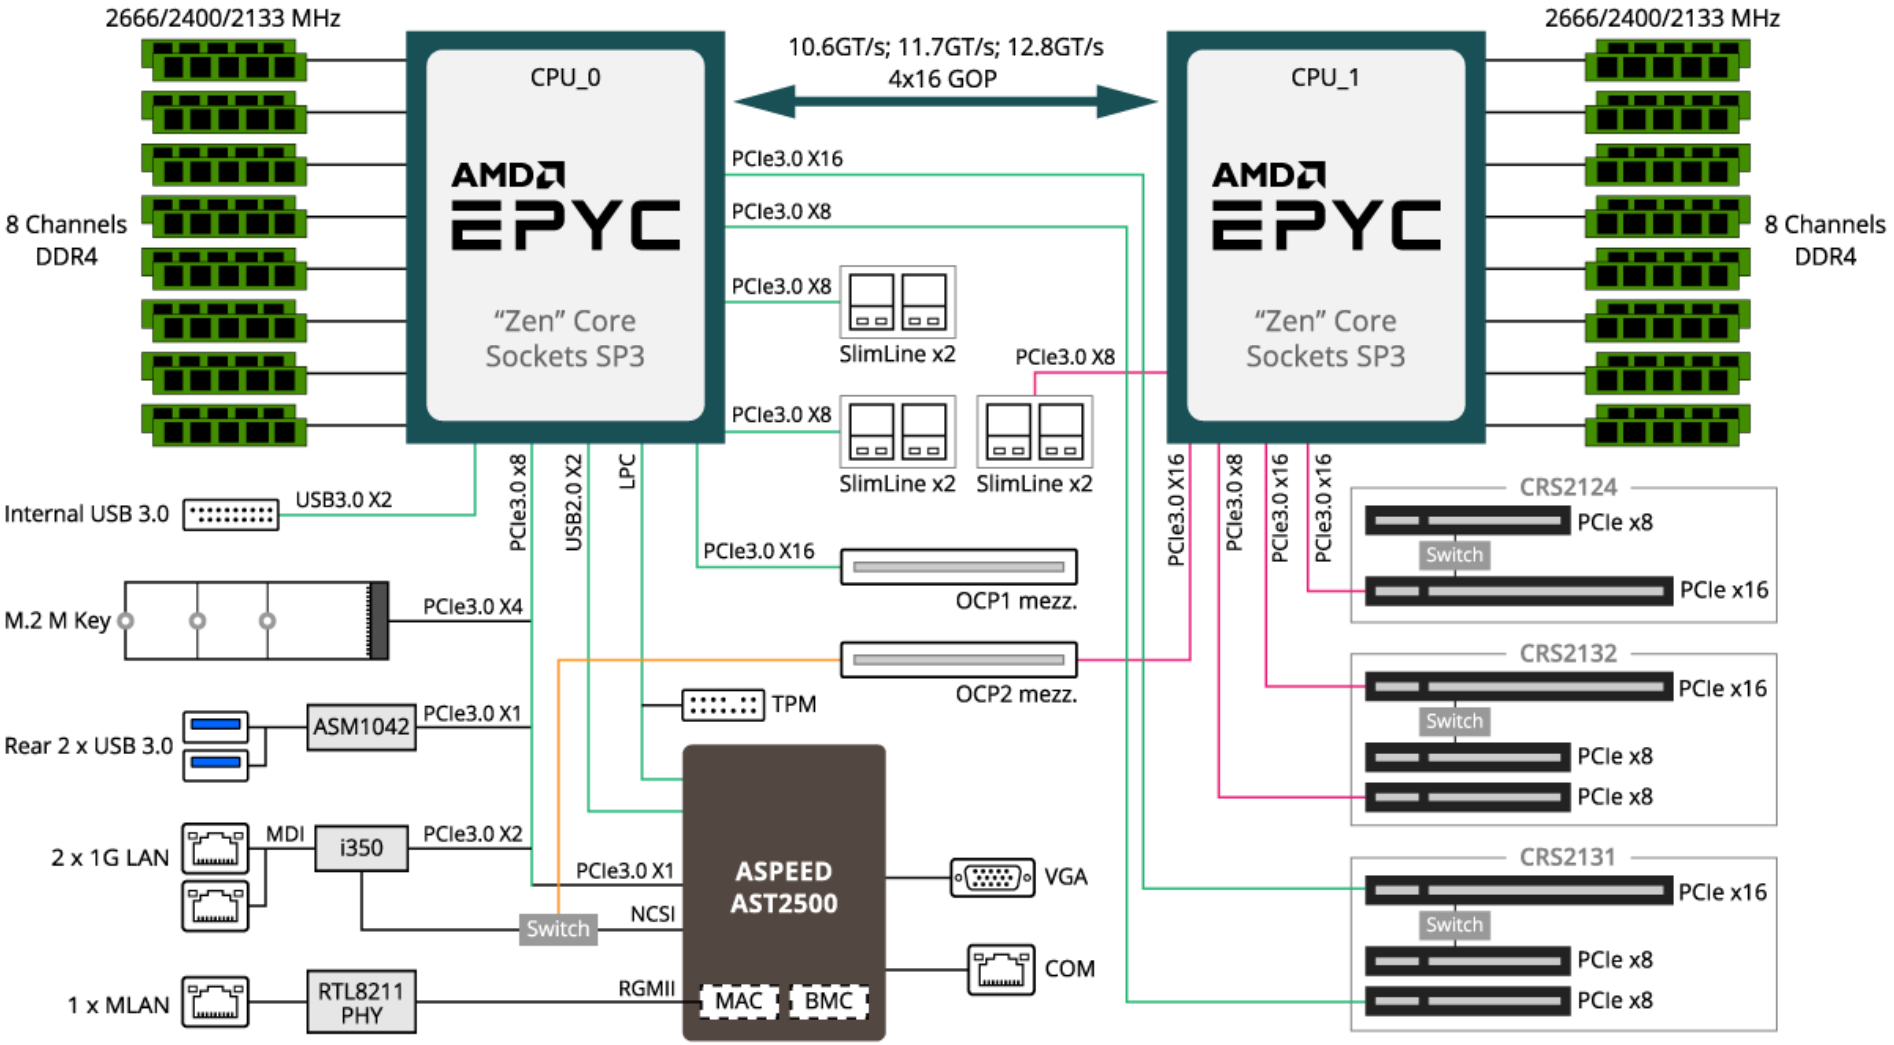

Performance comparison AMD's EPIC vs. Intel's Broadwell server line models while running the SPECint_rate_base2006 benchmark [21]

9.2 The 2. gen. Rome Epyc server line

The 2. gen. Rome EPYC server line

- Introduced in 11/2018, sampled in Q4 2018 and to be launched commercially
- in Q1 2018.
- It is based on the Zen 2 core and is fabricated on 7 nm technology (TSMC).
- It covers 1S and 2S servers.

Key features of the 2. gen. Rome Epyc server line -1

- Novel layout based on an I/O die (manufactured on 14 nm) and 8 core chiplet, each with 8 Zen 2 cores, as seen below.

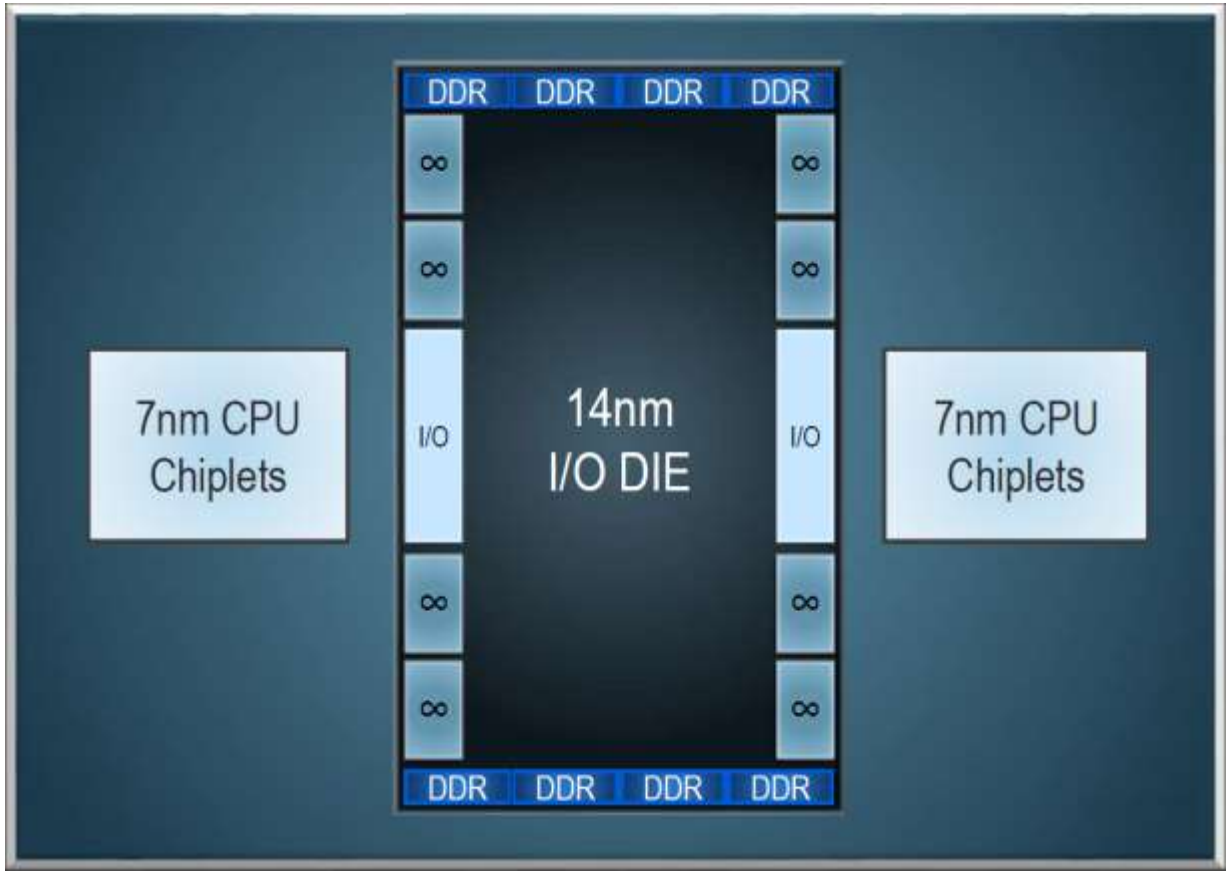

Figure: Layout of the 2. gen. Rome Epyc line [93]

Contrasting the layouts of the 1. and 2. gen. Epyc server lines [94]

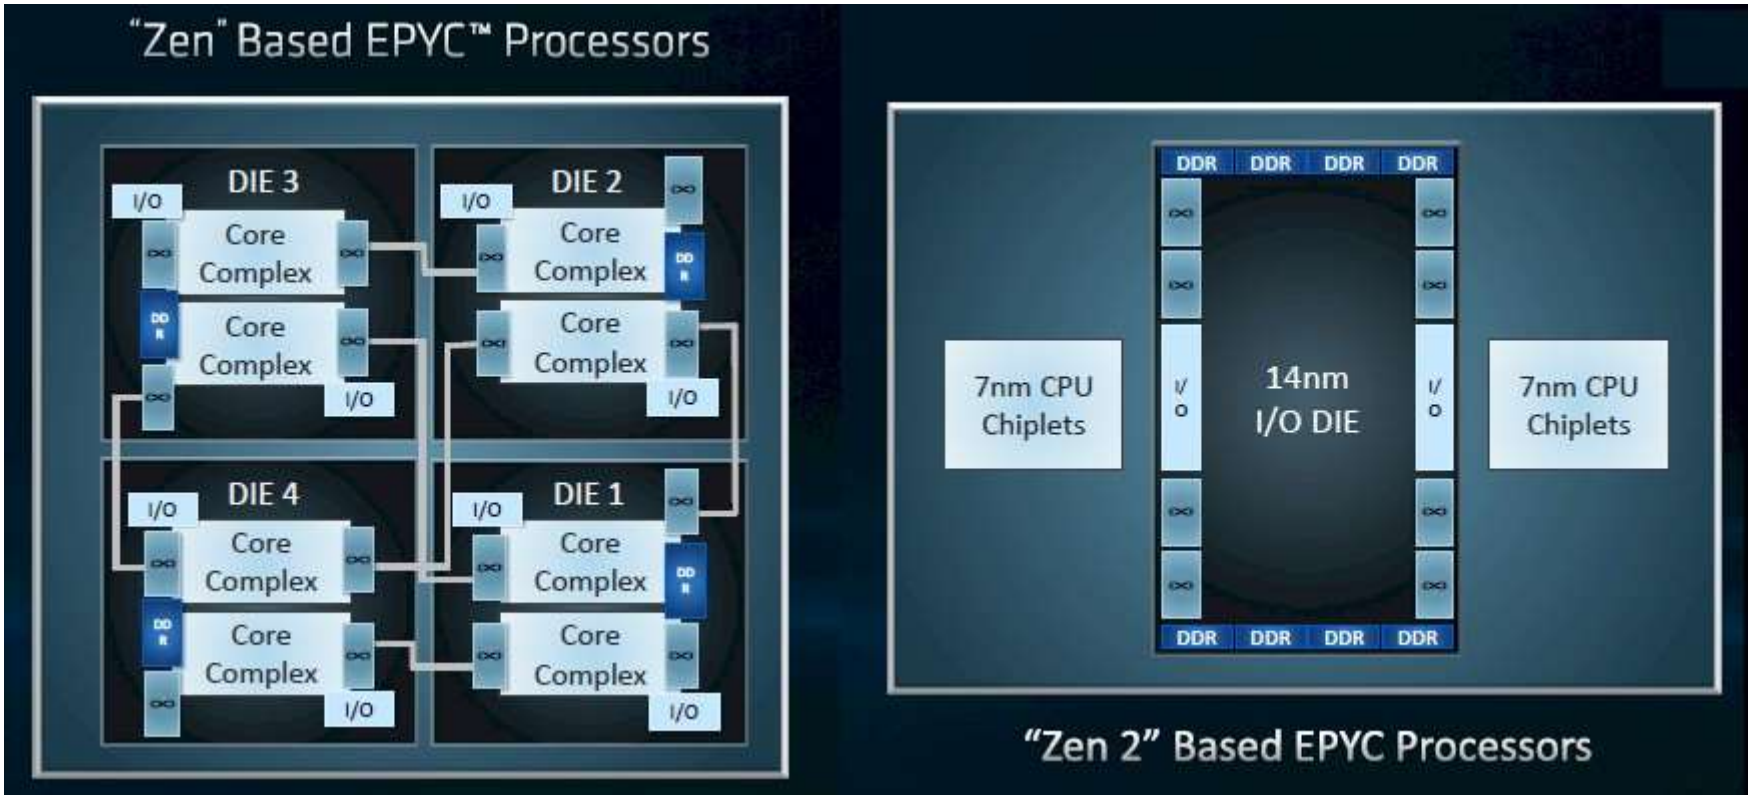

Key features of the 2. gen. Rome Epyc server line -2

- Eight channel DDR4 memory controller that provides equal access latency to all chiplets.
- Up to 4 TB DDR4 memory.
- 128 PCIe 4.0 lanes per socket, which is the first PCIe 4.0 implementation in processors.
- Greatly improved Infinity Fabric speed.
- Ability to connect GPUs via the I/O chip and Infinity Fabric protocol without using PCIe lanes.
- Socket compatibility with Naples processors.

Picture of the 2. gen. Rome Epyc server (11/2018) -2

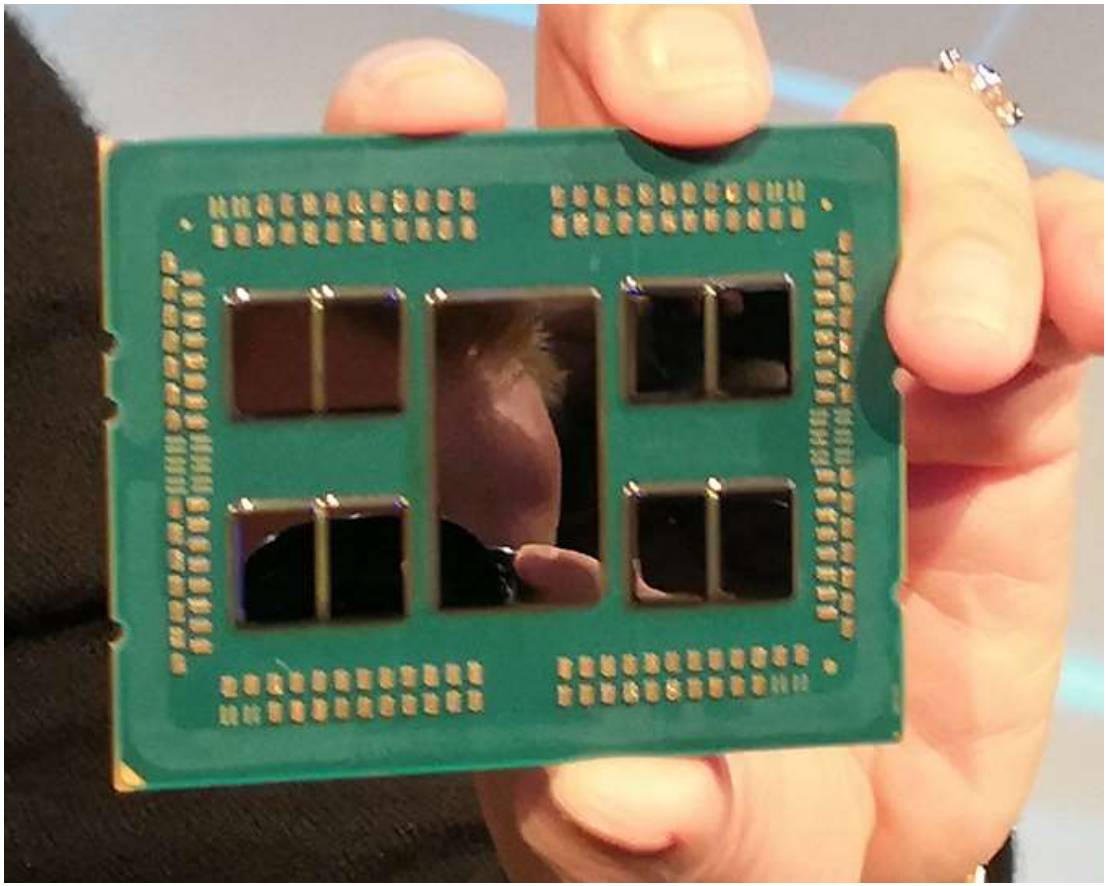

- AMD has already started sampling EPYC 'Rome' processors in about Q3/2018.

Intel's response to AMD's Rome, announced in 08/2018 at Hot Chips 30 [95]

Note that it is implemented as a Multi-Chip_Package with 48 cores, to be launched in H1 2019 on advanced 14 nm technology.

10. References

10. References (1)

[1]: Morris J., Computex 2016: AMD inches closer to Zen, announces Polaris graphics and new APUs, ZD Net, June 1 2016, http://www.zdnet.com/article/computex-2016-amd-inches-closer-to-zen-announces-polaris-graphics-and-new-apus/

[2]: Chiappetta M., Intel Xeon Scalable Debuts: Dual Xeon Platinum 8176 With 112 Threads Tested, Hot Hardware, July 11 2017, https://hothardware.com/reviews/intel-xeon-scalable-processor-family-review

[3]: Teich P., The Heart Of AMD's Epyc Comeback Is Infinity Fabric, The Next Platform, July 12 2017, https://www.nextplatform.com/2017/07/12/heart-amds-epyc-comeback-infinity-fabric/

[4]: Intel quadcore Xeon 5300 review, Nov. 13 2006, Hardware.Info, http://www.hardware.info/en-US/articles/amdnY2ppZGWa/Intel_quadcore_Xeon_5300_review

[5]: Wasson S., Intel's Woodcrest processor previewed, The Bensley server platform debuts, May 23, 2006, The Tech Report, http://techreport.com/articles.x/10021/1

[6]: Wikipedia, IBM z10, https://en.wikipedia.org/wiki/IBM_z10

[7]: Kennedy P., AMD EPYC Infinity Fabric Update and MCM Cost Savings, STH, Aug. 22 2017, https://www.servethehome.com/amd-epyc-infinity-fabric-update-mcm-cost-savings/

[8]: WikiChip, Zen - Microarchitectures – AMD, https://en.wikichip.org/wiki/amd/microarchitectures/zen

[9]: Amjad T., AMD Zen CPU Core Implementation Details Shared At ISSCC, SegmentNext, Febr. 9 2017, https://segmentnext.com/2017/02/09/amd-zen-cpu/

10. References (2)

[10]: Cutress I., The AMD Zen and Ryzen 7 Review: A Deep Dive on 1800X, 1700X and 1700 AnandTech, March 2 2017, http://www.anandtech.com/show/11170/the-amd-zen-and-ryzen-7-review-a-deep-dive-on-1800x-1700x-and-1700

[11]: Alcorn P., Intel Plays Defense: Inside Its EPYC Slide Deck, Tom's Hardware, July 17 2017, http://www.tomshardware.com/reviews/intel-amd-die-fabric-slides,5125-2.html

[12]: Wan S., AMD Releases EPYC Enterprise Processors, eTeknix, 2017, https://www.eteknix.com/amd-releases-epyc-enterprise-processors/

[13]: Cutress I., AMD Gives More Zen Details: Ryzen, 3.4 GHz+, NVMe, Neural Net Prediction, & 25 MHz Boost Steps, AnandTech, Dec. 13 2016, http://www.anandtech.com/show/10907/amd-gives-more-zen-details-ryzen-34-ghz-nvme-neural-net-prediction-25-mhz-boost-steps

[14]: Hruska J., New leak hints at AMD Zen's architecture, organization, Extreme Tech, April 29 2015, https://www.extremetech.com/gaming/204523-new-leak-hints-at-amd-zens-architecture-organization

[15]: Clark M., A New x86 Core Architecture for the Next Generation of Computing, Hot Chips 28, Aug. 23 2016, http://www.hotchips.org/wp-content/uploads/hc_archives/hc28/HC28.23-Tuesday-Epub/HC28.23.90-High-Perform-Epub/HC28.23.930-X86-core-MikeClark-AMD-final_v2-28.pdf

[16]: Singh T., Zen: A Next-Generation High-Performance x86 Core, Proc. ISSCC 2017, Session 3, Digital Processors

[17]: Goto H., AMD Infinity Fabric, PC Watch, April 6 2017, http://pc.watch.impress.co.jp/docs/column/kaigai/1053318.html

10. References (3)

[18]: Morgan T. P., AMD Winds Up One-Two Compute Punch For Servers, The Next Platform, June 19 2017, https://www.nextplatform.com/2017/06/19/amd-winds-one-two-compute-punch-servers/

[19]: Koduri R., AMD's 2017 Financial Analyst Day, May 16 2017

[20]: Ung G. M., Ryzen Threadripper review: AMD's monster 1950X stomps on other CPUs, PC World, Aug. 10 2017, https://www.pcworld.com/article/3214635/components-processors/ryzen-threadripper-review-we-test-amds-monster-cpu.html

[21]: Cutress I., AMD's Future in Servers: New 7000-Series CPUs Launched and EPYC Analysis, AnandTech, June 20 2017, http://www.anandtech.com/show/11551/amds-future-in-servers-new-7000-series-cpus-launched-and-epyc-analysis/2

[22]: Alcorn P., Intel Core i9-7900X Review: Meet Skylake-X, Tom's Hardware, June 19 2017, http://www.tomshardware.com/reviews/intel-core-i9-7900x-skylake-x,5092-2.html

[23]: Cutress I., The AMD Ryzen Threadripper 1950X and 1920X Review: CPUs on Steroids, AnandTech, Aug. 10 2017, http://www.anandtech.com/show/11697/the-amd-ryzen-threadripper-1950x-and-1920x-review

[24]: Alcorn P., AMD Ryzen Threadripper 1950X Review, Tom's Hardware, Aug. 10 2017, http://www.tomshardware.com/reviews/amd-ryzen-threadripper-1950x-cpu,5167-2.html

[25]: Gasior G., Your guide to the Ryzen AM4 platform and its X370, B350, and A320 chipsets, March 2 2017, https://rog.asus.com/articles/technologies/your-guide-to-the-ryzen-am4-platform-and-its-x370-b350-and-a320-chipsets/

10. References (4)

[26]: Ung G. M., Why Ryzen Threadripper has two extra chips, PC World, July 27 2017, https://www.pcworld.com/article/3211409/computers/why-ryzen-threadripper-has-two-mysterious-chips.html

[27]: Cinebench, Maxon, https://www.maxon.net/en/products/cinebench/

[28]: AMD: Threadripper Owns Intel, Seeking Alpha, Aug. 11 2017, https://seekingalpha.com/article/4097976-amd-threadripper-owns-intel?auth_param=djq9a:1corga9:7e836bbeb489dc69af8486d4cc569a2e

[29]: A New Day for the Datacenter, Presentation at AMD Financial Analyst Day, May 16 2017, http://ir.amd.com/phoenix.zhtml?c=74093&p=irol-analystday

[30]: Teich P., The New Server Economies Of Scale For AMD, The Next Platform, July 13 2017, https://www.nextplatform.com/2017/07/13/new-server-economies-scale-amd/

[31]: Kingsley-Hughes A., AMD EPYC: What you need to know, ZD Net, June 21 2017, http://www.zdnet.com/article/amd-epyc-what-you-need-to-know/

[32]: MCM-GPU: Multi-Chip-Module GPUs for Continued Performance Scalability, Nvidia, http://research.nvidia.com/publication/2017-06_MCM-GPU%3A-Multi-Chip-Module-GPUs

[33]: H11DSI-NT Users's Manual, Supermicro, 2017

[34]: Hruska J., Report: AMD Stealing Significant Market Share, Revenue from Intel, Extremetech Sept. 6, 2017, https://www.extremetech.com/computing/255122-european-retailer-shows-amd-stealing-significant-market-share-revenue-intel

10. References (5)

[35]: Skylake (client) - Microarchitectures - Intel, https://en.wikichip.org/wiki/intel/microarchitectures/skylake_(client)

[36]: Zen - Microarchitectures - AMD, WikiChip, https://en.wikichip.org/wiki/amd/microarchitectures/zen

[37]: Kampman J., AMD's Ryzen 7 2700U and Ryzen 5 2500U APUs revealed, TechReport, Oct. 26, 2017 http://techreport.com/review/32743/amd-ryzen-7-2700u-and-ryzen-5-2500u-apus-revealed

[38]: Cutress I., Ryzen Mobile is Launched: AMD APUs for Laptops, with Vega and Updated Zen, AnandTech, Oct. 26, 2017, https://www.anandtech.com/show/11964/ryzen-mobile-is-launched-amd-apus-for-laptops-with-vega-and-updated-zen

[39]: Day M., Understanding Low Drop Out (LDO) Regulators, Texas Instruments, 2006 http://www.ti.com/lit/ml/slup239/slup239.pdf

[40]: Bright P., CPU competition at last: AMD Ryzen brings 8 cores from just $329, ArsTechnica, 2/22/2017, https://arstechnica.com/information-technology/2017/02/amd-ryzen-arrives-march-2-8-cores-16-threads-from-just-329

[41]:Processor Programing Reference (PPR) for AMD Family 17h Model 01h, Revision B1 Processors, 54945 Rev. 1.14, April 15, 2017 ,

[42]: Kennedy P., Second AMD Naples Server Pictures and Platform Block Diagram 112 PCIe Lanes Exposed!, ServeTheHome, April 6, 2017, https://www.servethehome.com/wp-content/uploads/2017/04/AMD-Naples-Server-Block-Diagram.jpg

10. References (6)

[43]: Infinity Fabric (IF) - AMD, https://en.wikichip.org/wiki/amd/infinity_fabric

[44]: Frumusanu A., ARM Reveals Cotex-A72 Architecture Details, AnandTech, Apr. 23, 2015, http://www.anandtech.com/show/9184/arm-reveals-cortex-a72-architecture-details

[45]: Lee H-J., Shin Y., Bae S., Kim M., Kim K., 20nm High-K Metal Gate Heterogeneous 64-bit Quad-core CPUs and Hexa-core GPU for High performance and Energy-efficient Mobile Application Processor, ISOCC, 2015

[46]: Multi-core and System Coherence Design Challenges, ARM, http://www.ece.cmu.edu/~ece742/lib/exe/fetch.php?media=arm_multicore_and_system_coherence_-_cmu.pdf

[47]: Cheng M., NXP QorIQ Product Family Roadmap, Freescale, APF-NET-T0795, April 2013, http://www.freescale.com/files/training/doc/dwf/DWF13_APF_NET_T0795.pdf

[48]: Kennedy P., AMD Ryzen 5 1600X and 1500X (6-core and 4-core) Announced, STH, March 2, 2017, https://www.servethehome.com/amd-ryzen-5-1600x-and-1500x-6-core-and-4-core-announced/

[49]: ARM Cortex-A57 — So Big is Relative but How Relative is Your Big?, ARM, https://community.arm.com/soc/b/blog/posts/arm-cortex-a57-so-big-is-relative-but-how-relative-is-your-big

[50]: Moammer K., AMD Ryzen Architectural Deep-Dive – Ending The Intel Monopoly, WCCFtech, Feb 15, 2017, https://wccftech.com/amd-ryzen-architecture-detailed

10. References (7)

[51]: Cho A., Examining AMD Bear Logic, Seeking Alpha, Nov. 6. 2017, https://seekingalpha.com/article/4121329-examining-amd-bear-logic

[52]: Bright P., Intel to build discrete GPUs, hires AMD's top graphics guy to do it, ArsTechnica, 11/9 2017, https://arstechnica.com/gadgets/2017/11/intel-poaches-amds-top-gpu-architect-to-build-its-own-discrete-graphics-chips

[53]: Alcorn P., Hot Chips 2017: Intel Deep Dives Into EMIB, Hot Chips 2017, August 25, 2017, http://www.tomshardware.com/news/intel-emib-interconnect-fpga-chiplet,35316.html

[54]: Mardi S., AMD Announces New Ryzen Mobile Processors, TeckKnow, Oct. 26, 2017 https://www.teckknow.com/amd-announces-new-ryzen-mobile-processors/

[55]: Cortex-A73, ARM Developer, https://developer.arm.com/products/processors/cortex-a/cortex-a73

[56]: ARM Unveils Cortex-A75, A55 Processors And Mali-G72 GPU, ARM Developer, https://www.xda-developers.com/arm-unveils-cortex-a75-a55-processors-and-mali-g72-gpu/

[57]: Embedded Multi-die Interconnect Bridge, Intel, 2017, https://www.intel.com/content/www/us/en/foundry/emib.html

[58]: Mujtaba H., Intel Kaby Lake-G Processors To Pack Fast Discrete GPU and HBM2 Memory – Will Utilize Multi-Die Package Design, WCCFTech, Apr. 3, 2017 https://wccftech.com/intel-kaby-lake-g-hbm2-gpu-multi-die/

10. References (8)

[59]: Clark M., The "Zen" Architecture, ARM Ryzen Tech Day presentation, 2. March, 2017

[60]: AMD Tech Day, March 2, 2017, AMD Marketing
https://www.reddit.com/r/Amd/comments/5x4hxu/we_are_amd_creators_of_athlon_radeon_and_other

[61]: Intelligent Energy Management: An SoC design based on ARM926EJ-S, Hot Chips 2003,
https://www.hotchips.org/wp-content/uploads/hc_archives/hc15/2_Mon/5.arm.pdf

[62]: Deo M., Enabling Next-Generation Platforms Using Intel's 3D System, White Paper, Altera, 2017, https://www.altera.com/content/dam/altera-www/global/en_US/pdfs/literature/wp/wp-01251-enabling-nextgen-with-3d-system-in-package.pdf

[63]: Cutress I., The AMD 2nd Gen Ryzen Deep Dive: The 2700X, 2700, 2600X, and 2600 Tested, AnandTech, April 19 2018, https://www.anandtech.com/print/12625/amd-second-generation-ryzen-7-2700x-2700-ryzen-5-2600x-2600

[64]: AMDs 2019er CPU-Generationen heißen "Matisse" (CPU) und "Picasso" (APU), 3D Center, Sept. 26 2017, http://www.3dcenter.org/news/amds-2019er-cpu-generationen-heissen-matisse-cpu-und-picasso-apu

[65]: Bassiri S., AMD Ryzen 7 2700X and Ryzen 5 2600X Review, Tweak Town, April 19 2018, https://www.tweaktown.com/reviews/8602/amd-ryzen-7-2700x-5-2600x-review/index2.htm

[66]: Kingsley-Hughes A., AMD's 2nd-generation Ryzen - The ultimate desktop processor, ZD Net, April 19 2018, https://www.zdnet.com/article/amds-2nd-generation-ryzen-the-ultimate-desktop-processor/

10. References (9)

[67]: Kennedy P., Intel Optane Memory – 3D XPoint for Client Workloads Launched, STH, March 27 2017, https://www.servethehome.com/intel-optane-memory-3d-xpoint-client-workloads-launched/

[68]: Cutress I., AMD Tech Day at CES: 2018 Roadmap Revealed, with Ryzen APUs, Zen+ on 12nm, Vega on 7nm, AnandTech, Febr. 1 2018, https://www.anandtech.com/print/12233/amd-tech-day-at-ces-2018-roadmap-revealed-with-ryzen-apus-zen-on-12nm-vega-on-7nm

[69]: Garrou P., IFTLE 329 3D Integration Leaders – Europe; Trolls; More on Intel EMIB, Solid State Technology, April 6 2017, http://electroiq.com/insights-from-leading-edge/2017/04/iftle-329-3d-integration-leaders-europe-trolls-more-on-intel-emib/

[70]: Alcorn P., Intel Presents Its AMD EPYC Server Test Results, Tom's Hardware, Nov. 27 2017, https://www.tomshardware.com/news/amd-intel-epyc-xeon-benchmarks,35993.html

[71]: R281-Z91, https://b2b.gigabyte.com/Rack-Server/R281-Z91-rev-100#ov

[72]: AMD Introduces New Ryzen Mobile Processors, TechPowerUp, Oct. 26 2017, https://www.techpowerup.com/forums/threads/amd-introduces-new-ryzen-mobile-processors.238193/

[73]: Cutress I., The AMD Threadripper 2990WX 32-Core and 2950X 16-Core Review, AnandTech, August 13, 2018, https://www.anandtech.com/show/13124/the-amd-threadripper-2990wx-and-2950x-review

[74]: Cutress I., The AMD 2nd Gen Ryzen Deep Dive: The 2700X, 2700, 2600X, and 2600 Tested, AnandTech, Apr. 19, 2018, https://www.anandtech.com/show/12625/amd-second-generation-ryzen-7-2700x-2700-ryzen-5-2600x-2600

10. References (10)

[75]: AMD SenseMI Technology, AMD, 2018, https://www.amd.com/en/technologies/sense-mi

[76]: Leather A., Frequency boosting could be the hot CPU topic of 2018, bit-tech, April 5, 2018, https://www.bit-tech.net/blogs/cpu-frequency-boosting-could-be-the-hot-topic-of-2018/1

[77]: Beck N. C al., "Zeppelin": an SOC for Multi-chip Architectures, Presentation at ISSCC 2018, Febr. 12, 2018, https://www.slideshare.net/AMD/isscc-2018-zeppelin-an-soc-for-multichip-architectures

[78]: Schor D., ISSCC 2018: AMD's Zeppelin; Multi-chip routing and packaging, WikiChip Fuse, March 24, 2018, https://fuse.wikichip.org/news/1064/isscc-2018-amds-zeppelin-multi-chip-routing-and-packaging/

[79]: AMD StoreMI Technology, AMD, 2018, https://www.amd.com/en/technologies/store-mi

[80]: Cutress I., Intel Expands 8th Gen Core: Core i9 on Mobile, Iris Plus, Desktop, Chipsets, and vPro, AnandTech, April 3 2018, https://www.anandtech.com/show/12607/intel-expands-8th-gen-core-core-i9-on-mobile-iris-plus-desktop-chipsets-and-vpro

[81]: Raettig N., AMD Ryzen 2000 - Spiele-Benchmarks, Taktraten und Preise aufgetaucht, GameStar, 07-03-2018, https://www.gamestar.de/artikel/amd-ryzen-2000-benchmarks-taktraten-und-preise, 3326979.html

[82]: Hallock R., Understanding Precision Boost Overdrive in Three Easy Steps, AMD Communities, 13. Aug. 2018, https://community.amd.com/community/gaming/blog/2018/08/13/understanding-precision-boost-overdrive-in-three-easy-steps

10. References (11)

[83]: Gupsterg, Crosshair VII Hero Essential Info Thread, Republic of Gamers, 04-25_2018, https://rog.asus.com/forum/showthread.php?101617-Crosshair-VII-Hero-Essential-Info-Thread

[84]: AMD Redefines High-Performance Computing with New Processor and Graphics Products Preview at CES 2018, AMD, 2018 CES 1/7/2018, https://www.amd.com/en-us/press-releases/Pages/ces-2018-2018jan07.aspx

[85]: Bott E., Intel Optane: Memory acceleration kicks PCs into warp drive on April 24, ZD Net, March 27 2017, http://www.zdnet.com/article/intel-optane-memory-acceleration-modules-to-go-on-sale-april-24/

[86]: Johndms, AMD Extended Frequency Range (XFR) – Explained, Linustechtips, Oct.23, 2017, https://linustechtips.com/main/topic/850358-amd-extended-frequency-range-xfr-explained/

[87]: Beck, N. & al., "Zeppelin": An SoC for Multichip Architectures, ISSCC 2018, Febr. 12, 2018, Digest of Technical Papers, pp. 40-42

[88]: Beck, N. & al., "Zeppelin": An SoC for Multichip Architectures, ISSCC 2018, Febr. 12, 2018, Presentation slides, https://www.slideshare.net/AMD/isscc-2018-zeppelin-an-soc-for-multichip-architectures

[89]: O Donnell D., OC'd Threadripper 2990WX beat 5.0GHz and topped Cinebench scores...with caveats, Notebookcheck, Aug. 10, 2018, https://www.notebookcheck.net/OC-d-Threadripper-2990WX-beat-5-0GHz-and-topped-Cinebench-scores-with-caveats.321206.0.html

[90]: Gwennap : L., Sandy Bridge Spans Generations, Microprocessor, 9/27/10, http://www.cs.nyu.edu/courses/fall13/CSCI-GA.3033-008/Microprocessor-Report-Sandy-Bridge-Spans-Generations-243901.pdf

10. References (12)

[91]: Intel Pentium D 820 Review Overclock Benchmark –Overclocking D820, Xtreview, 2006-05-26, http://xtreview.com/review88.htm

[92]: Douglas J., Intel 8xx series and Paxville Xeon-MP Microprocessors, Hot Chips 17, 2005, https://www.hotchips.org/wp-content/uploads/hc_archives/hc17/3_Tue/HC17.S8/HC17.S8T1.pdf

[93]: Mujtaba H., AMD Officially Talks Zen 2 CPU Architecture: Significant Performance Uplift With 2X Throughput, Doubled Core Density With Up To 64 Cores, PCIe Gen 4.0 Support, Aiming Server EPYC Rome 7nm CPUs First, WCCFTech, No. 6, 2018, https://wccftech.com/amd-zen-2-7nm-cpu-architecture-epyc-rome-and-ryzen-official

[94]: Prickett T., AMD's Long Road From Naples To Milan Centers On Rome, TheNextPlatform, Nov. 6, 2018, https://www.nextplatform.com/2018/11/06/amds-long-road-from-naples-to-milan-centers-on-rome/

[95]: Prickett T., Intel To Challenge AMD With 48 Core "Cascade Lake" Xeon App, TheNextPlatform, Nov. 5, 2018, https://www.nextplatform.com/2018/11/05/intel-to-challenge-amd-with-48-core-cascade-lake-xeon-ap/

[96]: Schiesser T., Our take on Intel's newly announced Skylake-X Refresh and Xeon W-3175X 28-core, TechSpot, Oct. 10, 2018, https://www.techspot.com/news/76855-our-take-intel-newly-announced-skylake-x-refresh.html

[97]: System Guidance for Infrastructure, ARM Developer, https://developer.arm.com/products/system-design/system-guidance/system-guidance-for-infrastructure

10. References (13)

[98]: Bouvier D., Delivering A New Level of Visual Performance in an SOC AMD “Raven Ridge” APU, Hot Chips 30, 2018, http://www.hotchips.org/archives/2010s/hc30/

[99]: 12LP 12nm FinFET Technology, Product Brief, PB12LP-1.1, GlobalFoundries, https://www.globalfoundries.com/sites/default/files/product-briefs/pb-12lp-11-web.pdf

[100]: Alcorn P., AMD Launches 3000-series Picasso APUs, new H-Series and A-Series Processors, TomsHardware, Jan. 6, 2019, https://www.tomshardware.com/news/amd-3000-series-picasso-apu-ryzen,38290.html

[101]: Chromebook, 2019, https://en.wikipedia.org/wiki/Chromebook

[102]: Cutress I., AMD Ryzen 3rd Gen 'Matisse' Coming Mid 2019: Eight Core Zen 2 with PCIe 4.0 on Desktop, AnandTech, Jan. 9, 2019, https://www.anandtech.com/show/13829/amd-ryzen-3rd-generation-zen-2-pcie-4-eight-core