COMP 122

Rev 8-1-24

Computer Org & ASSEMBLY Programming

Intel

CPU's

Dr Jeff Drobman

website → *drjeffsoftware.com/classroom.html*

email → *jeffrey.drobman@csun.edu*

Index

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP122

Tech Titan Timeline

Historical Perspective

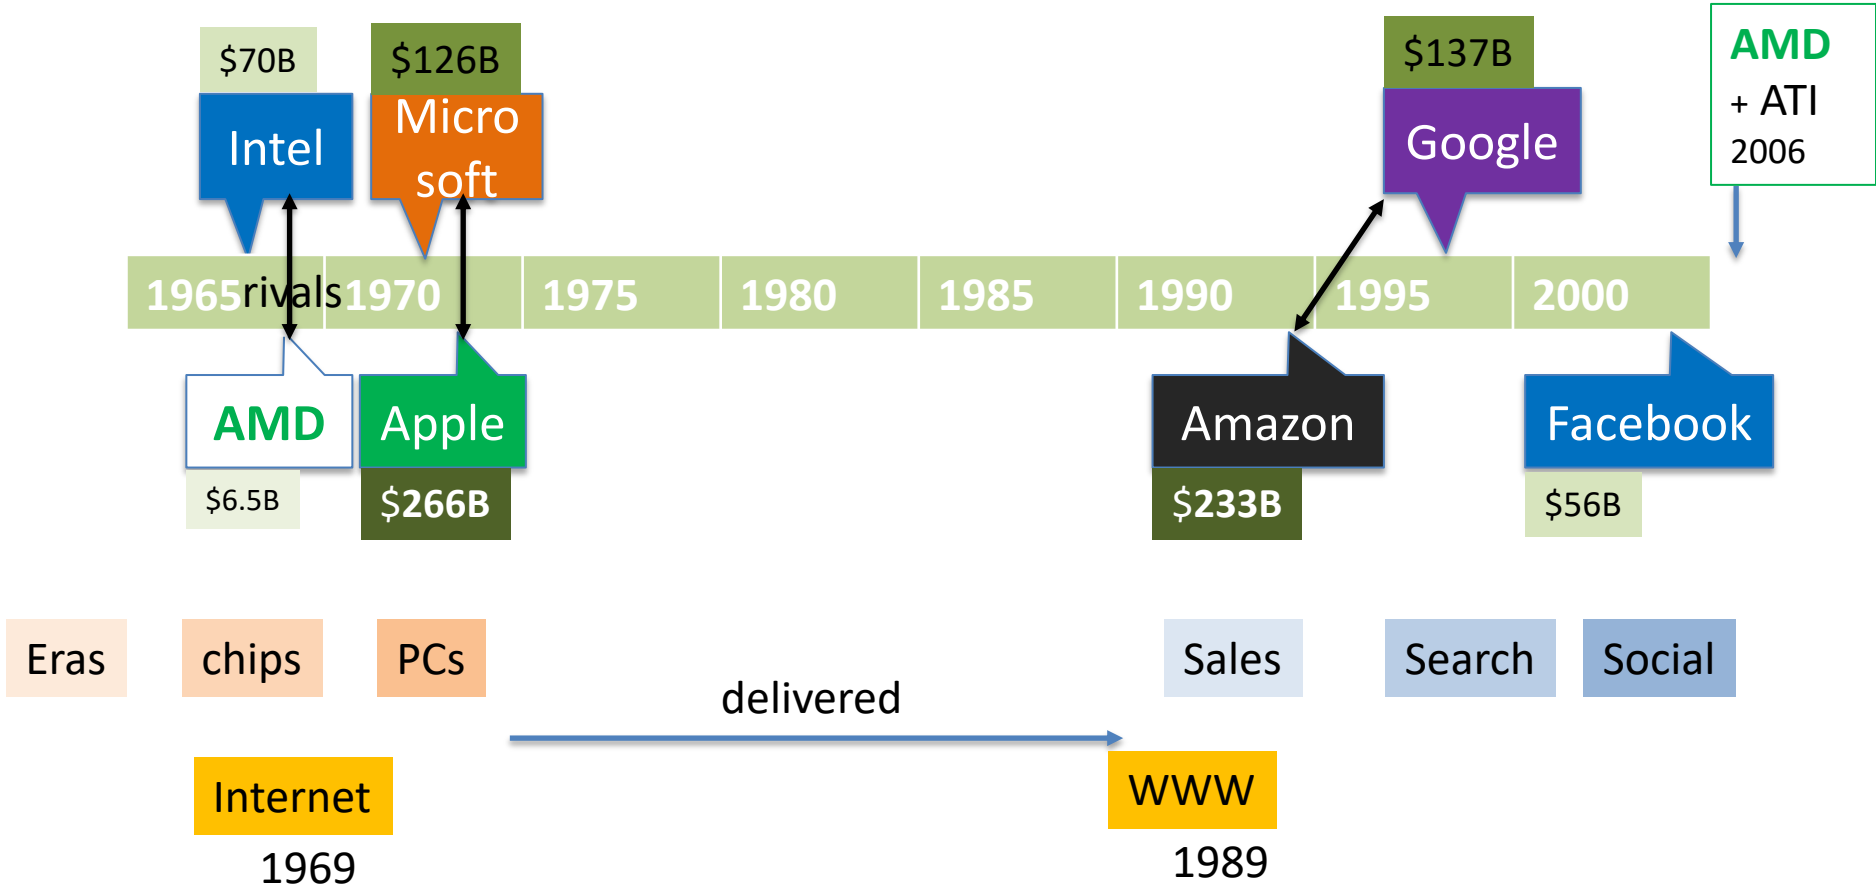

Intel

$70B

Intel's current logo, used since 2006

Intel's headquarters in Santa Clara, California

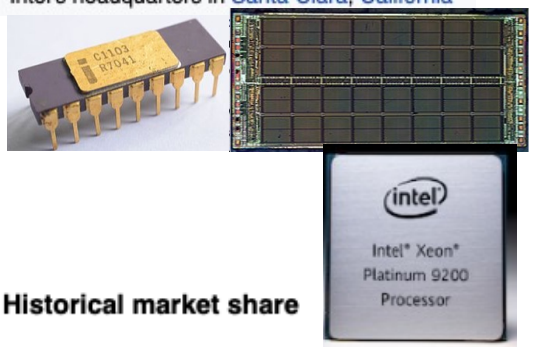

Intel logo used from 1968 to 2006

f. 1968 | Tech Titan | 2019

Founded	July 18, 1968; 51 years ago
Founders	Gordon Moore Robert Noyce
Revenue	▲ US$70.8 billion (2018)
Operating income	▲ US$23.3 billion (2018)
Net income	▲ US$21.0 billion (2018)
Number of employees	110,200 (2019)[2]
Subsidiaries	Mobileye McAfee (49%) Here (15%)

Andy Grove, Robert Noyce and Gordon Moore in 1978

Products

- Central processing units
- Microprocessors
- Integrated graphics processing units (iGPU)
- Systems-on-chip (SoCs)
- Motherboard chipsets
- Network interface controllers
- Modems
- Mobile phones
- Solid state drives
- Wi-Fi and Bluetooth Chipsets
- Flash memory
- Vehicle automation sensors

Historical market share

In the 1980s Intel was among the top ten sellers of semiconductors (10th in 1987) in the world. In 1992,[13] Intel became the biggest chip maker by revenue and has held the position ever since. Other top semiconductor companies include TSMC, Advanced Micro Devices, Samsung, Texas Instruments, Toshiba and STMicroelectronics.

Operating segments [edit]

- **Client Computing Group** – 55% of 2016 revenues – produces hardware components used in desktop and notebook computers.[10]
- **Data Center Group** – 29% of 2016 revenues – produces hardware components used in server, network, and storage platforms.[10]
- **Internet of Things Group** – 5% of 2016 revenues – offers platforms designed for retail, transportation, industrial, buildings and home use.[10]
- **Non-Volatile Memory Solutions Group** – 4% of 2016 revenues – manufactures NAND flash memory and 3D XPoint, branded as Optane, products primarily used in solid-state drives.[10]
- **Intel Security Group** – 4% of 2016 revenues – produces software, particularly security, and antivirus software.[10]
- **Programmable Solutions Group** – 3% of 2016 revenues – manufactures programmable semiconductors (primarily FPGAs).[10]

Founders HoF

Bell Labs

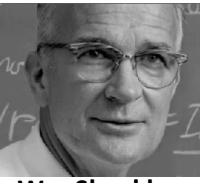

Wm Shockley

Fairchild founders (8)

Wilf Corrigan

Fairchild Chairman/CEO, LSI Logic founder

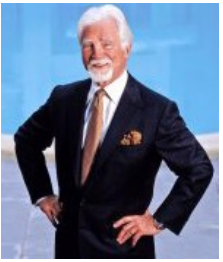

Jerry Sanders

CEO, AMD
1969–2002

Jack Gifford

AMD co-founder

In 1983, Gifford co-founded Maxim Integrated Products

Bob Noyce

Gordon Moore

Cypress Semi founder

Intel

1968

Intel Originals

L to R:
Andy Grove
Bob Noyce
Gordon Moore

Founders:
Bob Noyce
Gordon Moore

Intel CEO History

Intel CEOs

IC KNOWLEDGE LLC

CEO	Years	Background
Robert N. Noyce	1968-1975	Ph.D. in physics, co-inventor of the IC
Gordan E. Moore	1975-1987	Ph.D. in chemistry, creator of Moore's law
Andrew S. Grove	1987-1998	Ph.D. in chemical engineering, wrote Physics and Technology of Semiconductor Devices, co-discoverer of the Deal-Grove Oxidation Law
Craig R. Barrett	1998-2005	Ph.D. in materials science, professor at Stamford
Paul S. Otellini	2005-2013	MBA
Brian M. Krzanich	2013-2018	B.S. chemistry, manufacturing background
Robert H. Swan	2018-2021	MBA
Patrick P. Gelsinger	2021-	M.S. in electrical engineering, architect of the 80486

Until 2005 Intel was run by highly technical CEOs.

Semiconductors & IC's

MILESTONES

	1968	❖ **Intel** founded
	1969	❖ **AMD** founded
MPU	1971	❖ **Microprocessor** & **RAM** in **MOS** invented by Intel (i4004/i8008, i1101/3)
	1972	❖ Ion Implantation (replaces chemical diffusion) Mfg -- process
	1974	❖ Digital **Watch** IC invented
	1978	❖ **Wafer stepper** invented for fabs Mfg -- process
	1979	❖ **IDT** founded (**CMOS**) Mfg -- process
ASIC	1981	❖ **LSI Logic** founded (**ASIC**)
FPGA	1984	❖ **Xilinx** makes 1st **FPGA**, **MIPS** founded as early ***RISC*** pioneer (licenses LSI & IDT)
	1985	❖ **ARM** founded as Acorn RISC Machines
Flash	1987	❖ **Toshiba** intro's **Flash** EEPROM, **TSMC** founded (foundry) Mfg -- process
GPU	1998	❖ **Nvidia** founded (1993) – 1st **GPU's** (1998)
	2002	❖ Intel goes to **300mm** (12in) wafers Mfg -- process
	2009	❖ AMD spins off fabs to **Global Foundries** (owner Abu Dhabi) Mfg -- process
	2011	❖ Intel **FinFET** Mfg -- process
Dec	2019	❖ Intel intro's 1st **QC** chip ("Horse Ridge")
	2020	❖ ARM intro's "backside power" process Mfg -- process

First MPUs

4004
2300 transistors

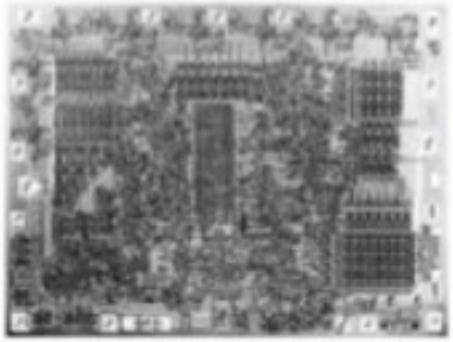

8008
3098 transistors

Intel MCS-4 and MCS-8 design team and CPU chips

Chips-Moore's Law

https://www.youtube.com/watch?v=ylgk3HEyZ_g

transistors doubles every 2 years

Plaque to Moore\'s Law at the technology plaza in Mountain View, beneath a model of the Silicon crystal

Dave Laws (2018)

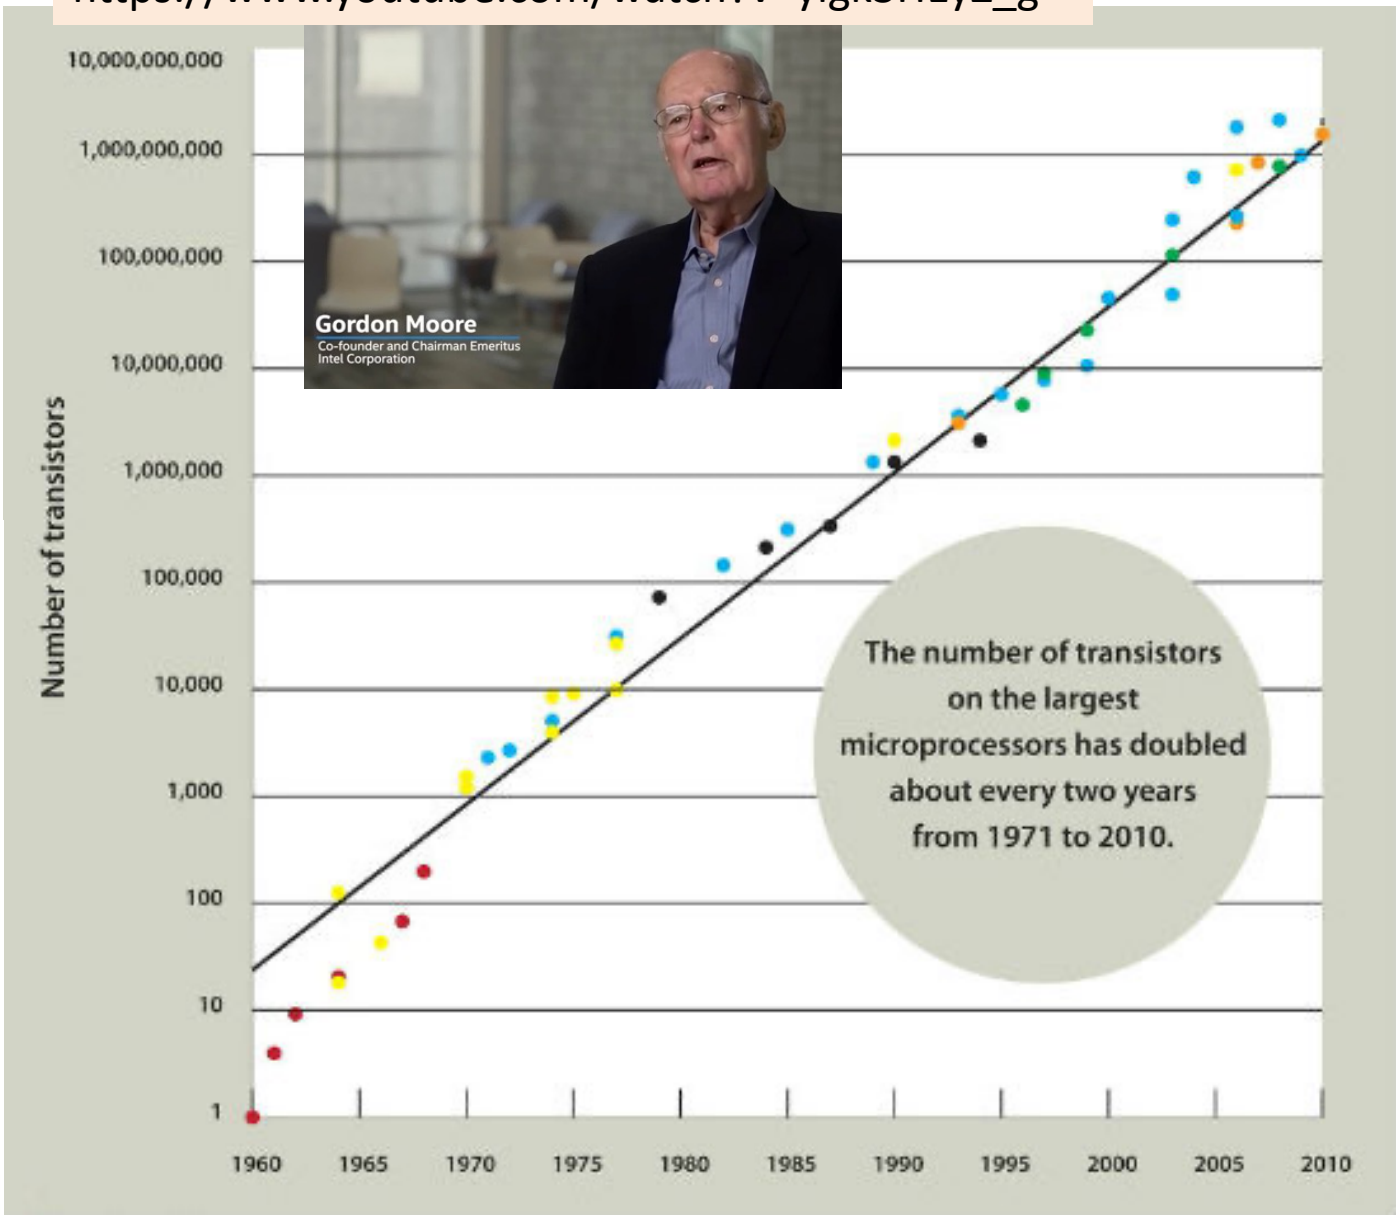

Section

i8086

See separate slide set ”x86”

8086 Reset Vector

Quora

Drazen Zoric · Follow
Lives in Cork, Ireland ·

8086 supports 1MB of external memory and reset vector is hardcoded to 0xFFFF0.
This gives BIOS 16 bytes at top for JMP to initialization routine somewhere. Well, BIOS was mostly on top of 1MB range, let's say top 64kB or more.

In later versions reset vector was moved to, again top, 0xFFFFFFF0.

Later things changed and today is more complicated cause of microcode updates, security, etc.

In original IBM PC lower 640kB was for RAM, followed by expansion slots and finally BIOS.

x86 Modes

In the x86-64 computer architecture, **long mode** is the mode where a 64-bit operating system can access 64-bit instructions and registers. 64-bit programs are run in a sub-mode called 64-bit mode, while 32-bit programs and 16-bit protected mode programs are executed in a sub-mode called compatibility mode. Real mode or virtual 8086 mode programs cannot be natively run in long mode.

An x86-64 processor acts identically as an IA-32 processor when running in real mode or protected mode, which are supported sub-modes when the processor is *not* in long mode.

"IA" = Intel Architecture

Long mode

From Wikipedia, the free encyclopedia

Part of a series on

Microprocessor modes for the x86 architecture

Real mode (Intel 8086)
8080 emulation mode (NEC V20/V30 only)
Protected mode (Intel 80286)
Unreal mode (Intel 80286)
Virtual 8086 mode (Intel 80386)
System Management Mode (Intel 386SL)
Long mode (AMD Athlon 64)
x86 virtualization (Intel Pentium 4, AMD Athlon 64)

First supported platform shown in parentheses

Section

i8088

Intel i8088

IBM PC 1981

Jeff Drobman · just now
Lecturer at California State University, Northridge (2016–present)

the i8008 and later i8080 were the world's first 8-bit microprocessors. then came the i8086 as the first 16-bit one, followed closely by the M68000 and the Z8000. the i8088 was a special design for IBM's PC: an i8086 with an 8-bit data bus (so IBM only had to build an 8-bit memory, but it was 16-bit addressable (64KiB).

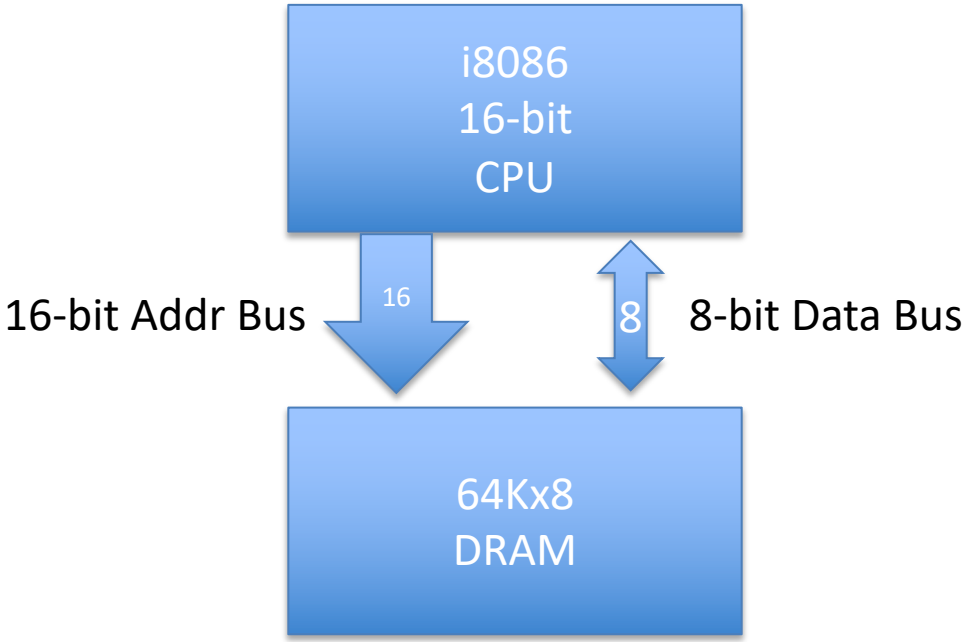

i8088 Muxed Buses

Quora

Joe Zbiciak · Follow
Processor and memory system architect for ~25 years ·

The 8086 and 8088 do not *require* a data bus transceiver.

The Intel® 8086 Family User's Manual shows the following configuration for an 8088 system as an example:

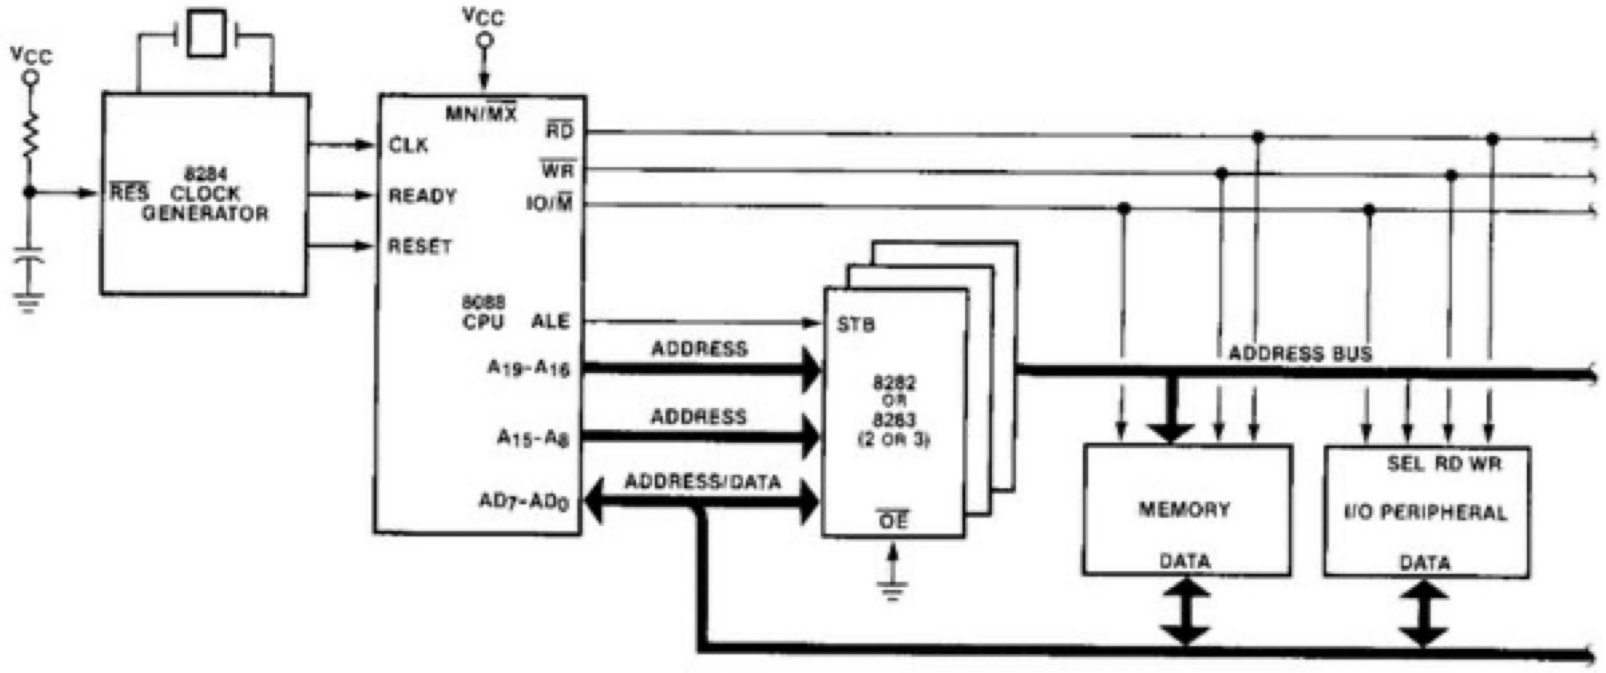

Figure 4-11. Minimum Mode 8088 Demultiplexed Address Bus

Old Intel 8088 Data Book

i8088

JULY 1981

iAPX 88 BOOK

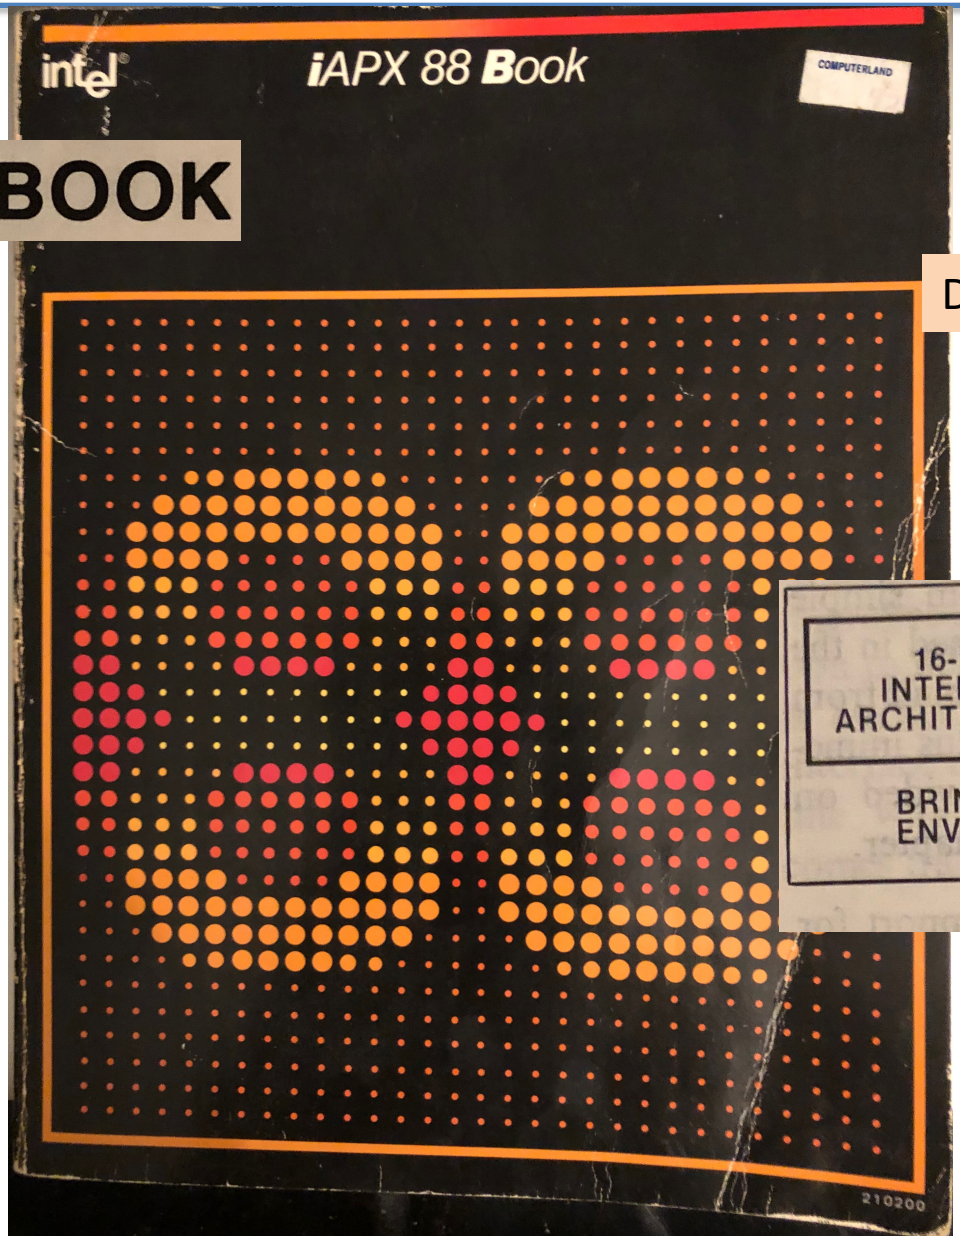

Designed for **IBM PC** (Aug 1981)

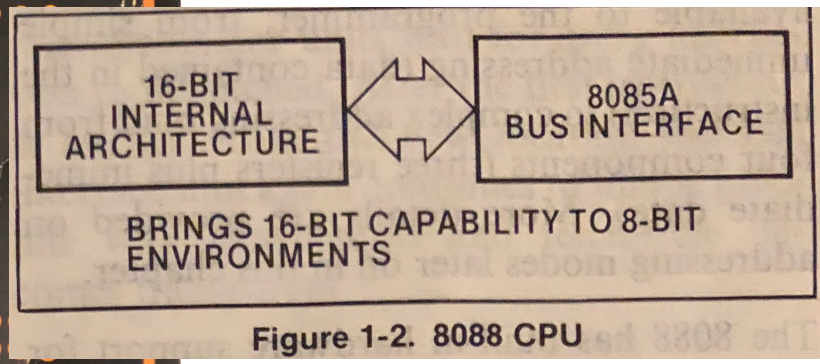

Figure 1-2. 8088 CPU

Old Intel 8088 Data Book

3 Simple Buses

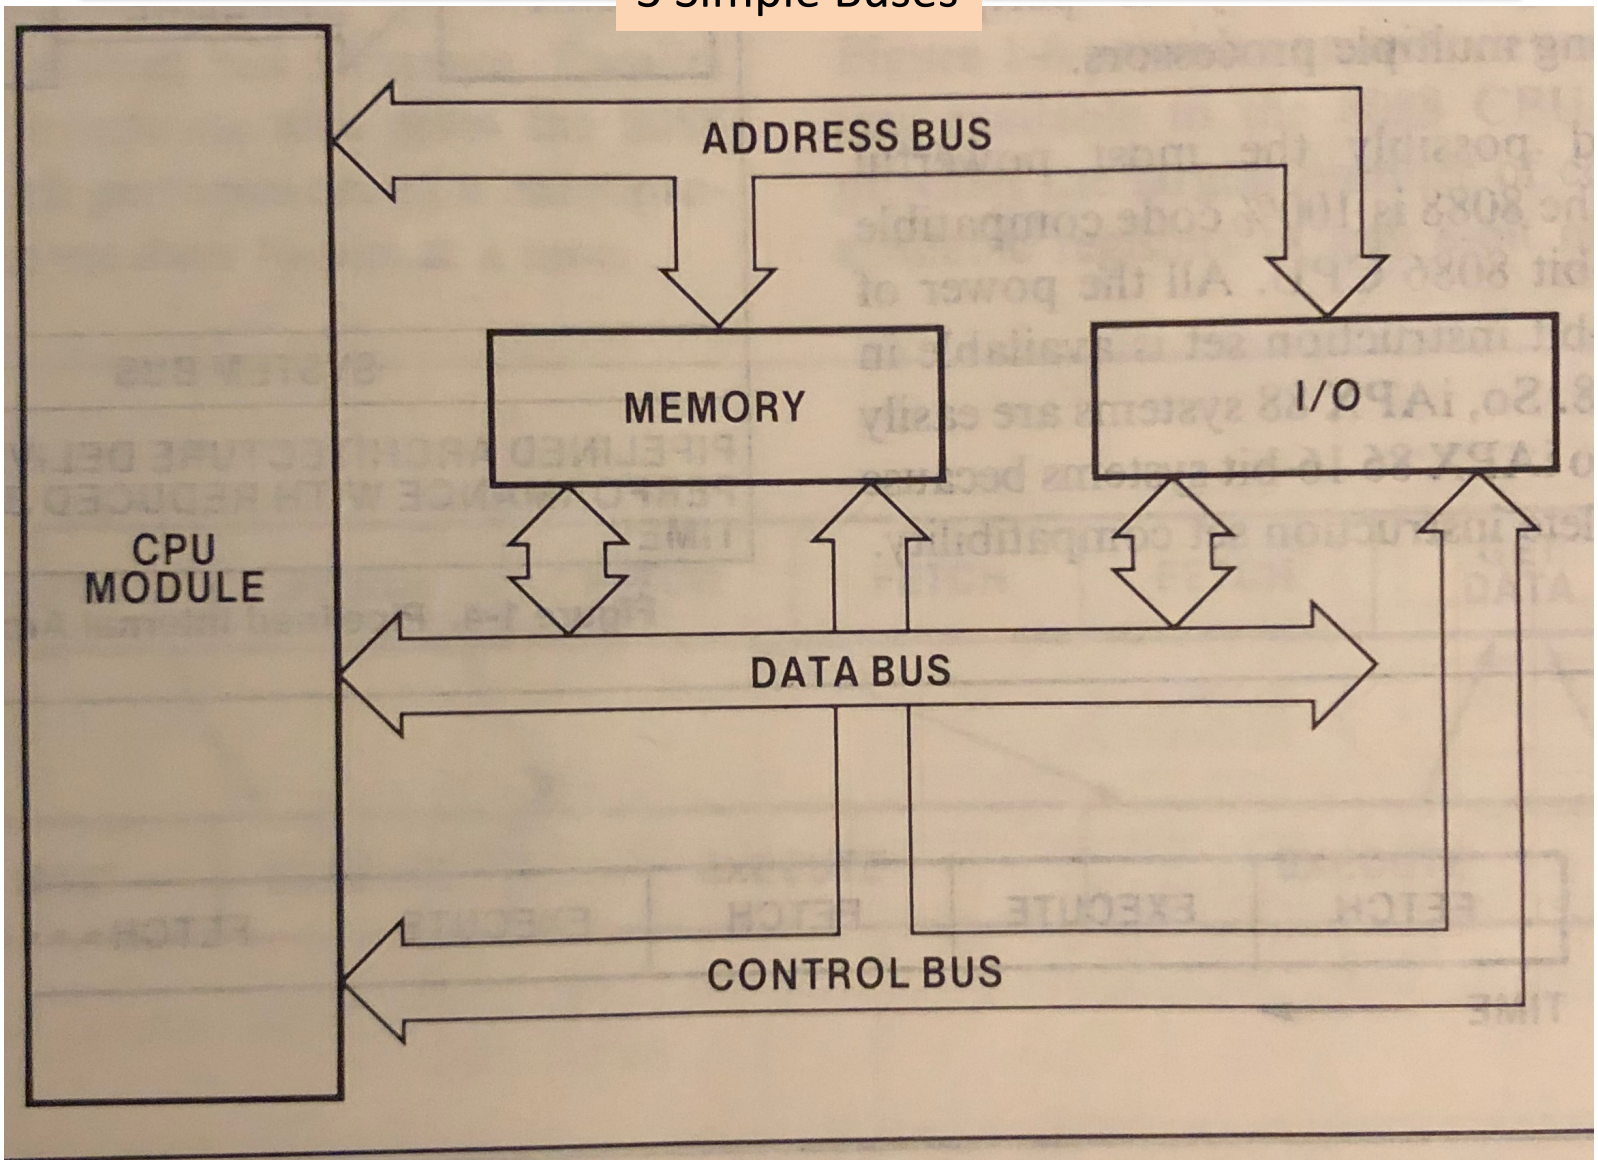

Figure 1-1. Microcomputer Block Diagram

Old Intel 8088 Data Book

i8088

FPU

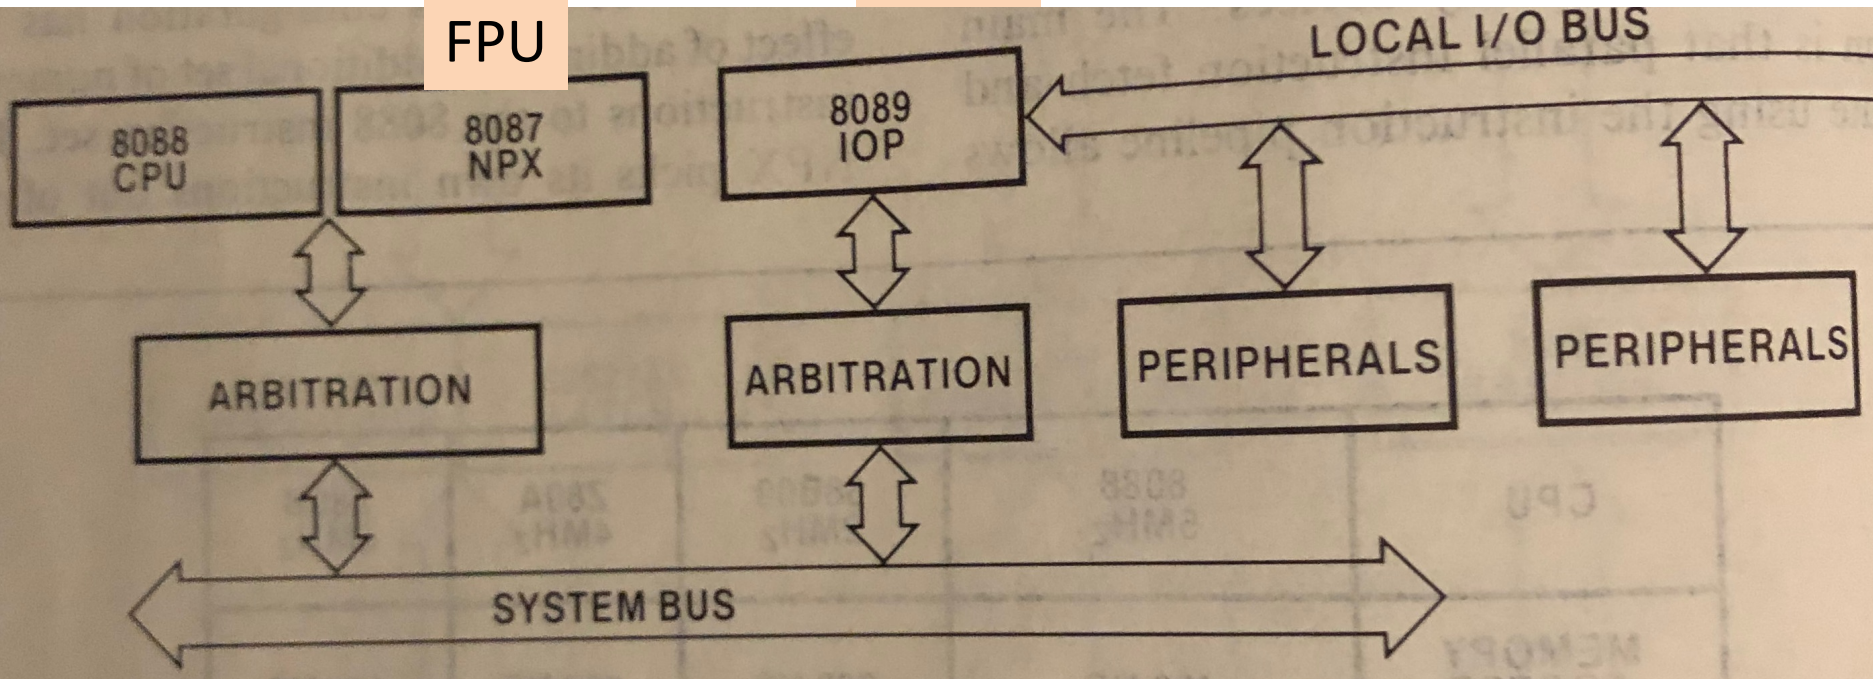

ARCHITECTURE EXTENDS FOR EVEN MORE PERFORMANCE

Figure 1-26. iAPX 88 Processor Extensions

Old Intel 8088 Registers

i8088

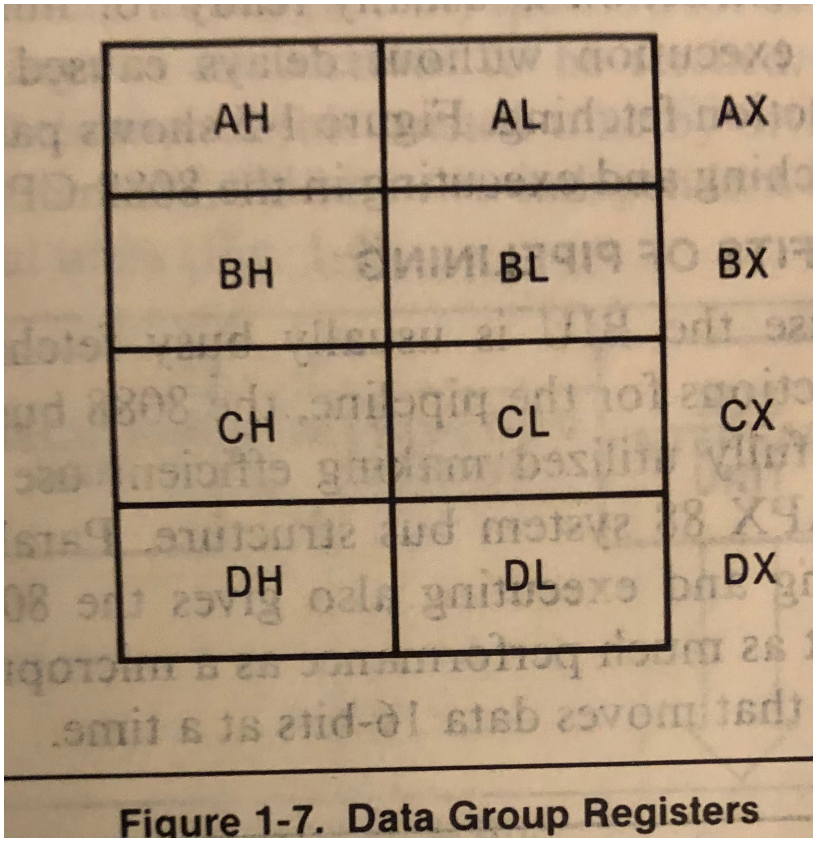

Figure 1-7. Data Group Registers

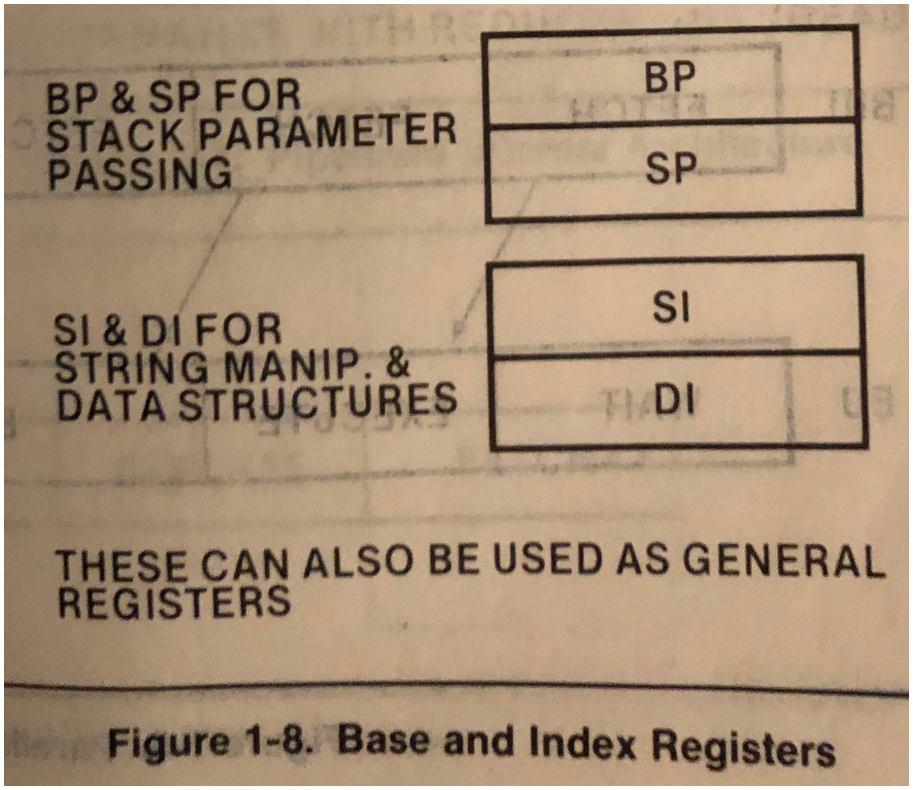

Figure 1-8. Base and Index Registers

Old Intel 8088 Registers

i8088

DATA REGISTERS

	7 0	7 0
AX	AH	AL
BX	BH	BL
CX	CH	CL
DX	DH	DL

POINTER AND INDEX REGISTERS

	15 0	
SP		STACK POINTER
BP		BASE POINTER
SI		SOURCE INDEX
DI		DESTINATION INDEX

Old Intel 8088 Registers

i8088

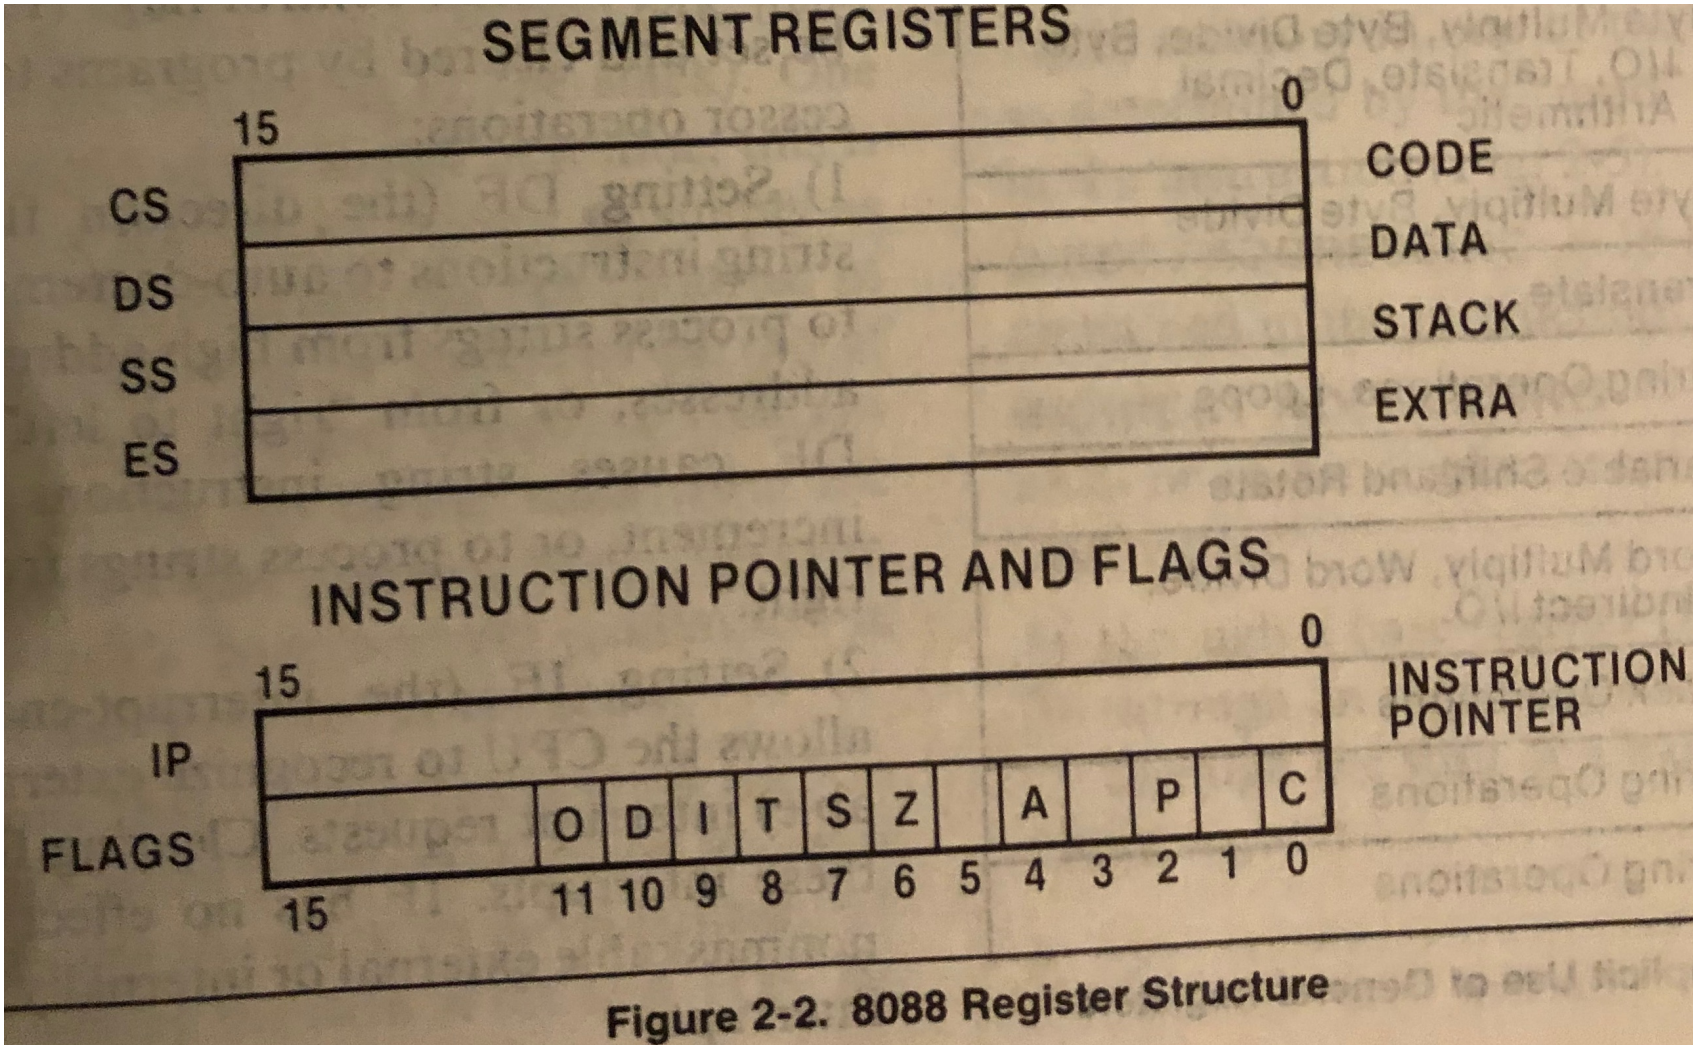

Figure 2-2. 8088 Register Structure

Old Intel 8088 Data Book

EA

FIRST OPERAND CHOICE DEPENDS ON ADDRESSING MODE:

FIRST OPERAND IN MEMORY			FIRST OPERAND IN REGISTER		
INDIRECT ADDRESSING		DIRECT ADDRESSING	MOD = 11		
MOD = 00* : DISP = 0 MOD = 01 : DISP = DISP-LO SIGN EXTENDED 10 : DISP = DISP-HI, DISP-LO		MOD = 00 AND R/M = 110 OPERAND EFFECTIVE ADDRESS = DISP-HI, DISP-LO	R/M:	REGISTER	
R/M:	OPERAND EFFECTIVE ADDRESS			8-BIT (W = 0)	16-BIT (W = 1)
000	(BX) + (SI) + DISP		000	AL	AX
001	(BX) + (DI) + DISP		001	CL	CX
010	(BP) + (SI) + DISP		010	DL	DX
011	(BP) + (DI) + DISP		011	BL	BX
100	(SI) + DISP		100	AH	SP
101	(DI) + DISP		101	CH	BP
110	(BP) + DISP		110	DH	SI
111	(BX) + DISP		111	BH	DI

Where ( ) means "contents of"
*Exception—direct addressing mode

Figure 2-5. Determining First Operand

Old Intel 8088 Data Book

EA

DATA STRUCTURE	DATA MEMORY		STACK
	WITHOUT BASE	WITH BASE	
SIMPLE VARIABLE	DIRECT	BX + OFFSET	BP + OFFSET
ARRAYS	SI DI	BX + SI BX + DI	BP + SI BP + DI
ARRAYS OF RECORDS	SI + OFFSET DI + OFFSET	BX + SI + OFFSET BX + DI + OFFSET	BP + SI + OFFSET BP + DI + OFFSET

Figure 2-6. Effective Addresses Used with Different Data Structures

Old Intel 8088 Data Book

"reg" Field Bit Assignments:

16-Bit (w = 1)	8-Bit (w = 0)	Segment
000 AX	000 AL	00 ES
001 CX	001 CL	01 CS
010 DX	010 DL	10 SS
011 BX	011 BL	11 DS
100 SP	100 AH	
101 BP	101 CH	
110 SI	110 DH	
111 DI	111 BH	

Old Intel 8088 Data Book

"r/m" Field Bit Assignments:

r/m	Operand Address
000	(BX) + (SI) + DISP
001	(BX) + (DI) + DISP
010	(BP) + (SI) + DISP
011	(BP) + (DI) + DISP
100	(SI) + DISP
101	(DI) + DISP
110	(BP) + DISP
111	(BX) + DISP

DISP = Offset

EA

DISP follows 2nd byte of instruction (before data if required).

*except if mod = 00 and r/m = 110 then EA = disp-high: disp-low

Old Intel 8088 Data Book

i8088

Encoding:

Memory or Register Operand with Register Operand:

000000dw	mod reg r/m

2 bytes

if d = 1 then LSRC = REG, RSRC = EA, DEST = REG
else LSRC = EA, RSRC = REG, DEST = EA

Immediate Operand to Memory or Register Operand:

100000sw	mod 000 r/m	data	data if s:w=01

4 bytes

LSRC = EA, RSRC = data, DEST = EA

Immediate Operand to Accumulator:

0000010w	data	data if w=1

3 bytes

if w = 0 then LSRC = AL, RSRC = data, DEST = AL
else LSRC = AX, RSRC = data, DEST = AX

Old Intel 8088 Clocks

EA

EA COMPONENTS		CLOCKS*
Displacement Only		6
Base or Index Only	(BX,BP,SI,DI)	5
Displacement + Base or Index	(BX,BP,SI,DI)	9
Base + Index	BP + DI, BX + SI	7
	BP + SI, BX + DI	8
Displacement + Base + Index	BP + DI + DISP BX + SI + DISP	11
	BP + SI + DISP BX + DI + DISP	12

*Add 2 clocks for segment override

Figure 2-10. Effective Address Calculation Time

Old Intel 8088 Clocks/Bytes

i8088

ADD Operands	Clocks*	Transfers	Bytes	ADD Coding Examples
register, register	3	—	2	ADD CX, DX
register, memory	9(13) + EA	1	2-4	ADD DI, [BX].ALPHA
memory, register	16(24) + EA	2	2-4	ADD TEMP, CL
register, immediate	4	—	3-4	ADD CL, 2
memory, immediate	17(25) + EA	2	3-6	ADD ALPHA, 2
accumulator, immediate	4	—	2-3	ADD AX, 200

*b(w): where b denotes the number of clock cycles for

Memory (vars)

Old Intel 8088 ISA

signed binary, signed binary

GENERAL PURPOSE	
MOV	Move byte or word
PUSH	Push word onto stack
POP	Pop word off stack
XCHG	Exchange byte or word
XLAT	Translate byte
INPUT/OUTPUT	
IN	Input byte or word
OUT	Output byte or word
ADDRESS OBJECT	
LEA	Load effective address
LDS	Load pointer using DS
LES	Load pointer using ES
FLAG TRANSFER	
LAHF	Load AH register from flags
SAHF	Store AH register in flags
PUSHF	Push flags onto stack
POPF	Pop flags off stack

Figure 1-16. Data Transfer Instructions

ADDITION	
ADD	Add byte or word
ADC	Add byte or word with carry
INC	Increment byte or word by 1
AAA	ASCII adjust for addition
DAA	Decimal adjust for addition
SUBTRACTION	
SUB	Subtract byte or word
SBB	Subtract byte or word with borrow
DEC	Decrement byte or word by 1
NEG	Negate byte or word
CMP	Compare byte or word
AAS	ASCII adjust for subtraction
DAS	Decimal adjust for subtraction
MULTIPLICATION	
MUL	Multiply byte or word unsigned
IMUL	Integer multiply byte or word
AAM	ASCII adjust for multiply
DIVISION	
DIV	Divide byte or word unsigned
IDIV	Integer divide byte or word
AAD	ASCII adjust for division
CBW	Convert byte to word
CWD	Convert word to doubleword

Figure 1-17. Arithmetic Instructions

Old Intel 8088 ISA

LOGICALS	
NOT	"Not" byte or word
AND	"And" byte or word
OR	"Inclusive or" byte or word
XOR	"Exclusive or" byte or word
TEST	"Test" byte or word
SHIFTS	
SHL/SAL	Shift logical/arithmetic left byte or word
SHR	Shift logical right byte or word
SAR	Shift arithmetic right byte or word
ROTATES	
ROL	Rotate left byte or word
ROR	Rotate right byte or word
RCL	Rotate through carry left byte or word
RCR	Rotate through carry right byte or word

Figure 1-18. Bit Manipulation Instructions

MOVS	Move byte or word string
MOVSB/MOVSW	Move byte or word string
CMPS	Compare byte or word string
SCAS	Scan byte or word string
LODS	Load byte or word string
STOS	Store byte or word string
REP	Repeat
REPE/REPZ	Repeat while equal/zero
REPNE/REPNZ	Repeat while not equal/not zero

Figure 1-19. String Instructions

Old Intel 8088 ISA

CONDITIONAL TRANSFERS	
JA/JNBE	Jump if above/not below nor equal
JAE/JNB	Jump if above or equal/not below
JB/JNAE	Jump if below/not above nor equal
JBE/JNA	Jump if below or equal/not above
JC	Jump if carry
JE/JZ	Jump if equal/zero
JG/JNLE	Jump if greater/not less nor equal
JGE/JNL	Jump if greater or equal/not less
JL/JNGE	Jump if less/not greater nor equal
JLE/JNG	Jump if less or equal/not greater
JNC	Jump if not carry
JNE/JNZ	Jump if not equal/not zero
JNO	Jump if not overflow
JNP/JPO	Jump if not parity/parity odd
JNS	Jump if not sign
JO	Jump if overflow
JP/JPE	Jump if parity/parity even
JS	Jump if sign

UNCONDITIONAL TRANSFERS	
CALL	Call procedure
RET	Return from procedure
JMP	Jump
ITERATION CONTROLS	
LOOP	Loop
LOOPE/LOOPZ	Loop if equal/zero
LOOPNE/LOOPNZ	Loop if not equal/not zero
JCXZ	Jump if register CX = 0
INTERRUPTS	
INT	Interrupt
INTO	Interrupt if overflow
IRET	Interrupt return

Figure 1-20. Program Transfer Instructions

Old Intel 8088 ISA

i8088

FLAG OPERATIONS	
STC	Set carry flag
CLC	Clear carry flag
CMC	Complement carry flag
STD	Set direction flag
CLD	Clear direction flag
STI	Set interrupt enable flag
CLI	Clear interrupt enable flag
EXTERNAL SYNCHRONIZATION	
HLT	Halt until interrupt or reset
WAIT	Wait for $\overline{\text{TEST}}$ pin active
ESC	Escape to external processor
LOCK	Lock bus during next instruction
NO OPERATION	
NOP	No operation

Figure 1-21. Processor Control Instructions

Old Intel 8088 Data Book

i8088 Segments

```
1.  MY_DATA     SEGMENT                                  ;data segment
2.  SUM         DB          ?                            ;reserve a byte for SUM
3.  MY_DATA     ENDS
4.  MY_CODE     SEGMENT                                  ;code segment
5.              ASSUME      CS:MY_CODE, DS:MY_DATA       ;contents of CS and DS

6.  PORT_VAL    EQU         3                            ;symbolic name for port number
7.  GO:         MOV         AX,MY_DATA                   ;initialize DS to MY_DATA
8.              MOV         DS,AX
9.              MOV         SUM,0                        ;clear sum
10. CYCLE:      CMP         SUM,100                      ;if SUM exceeds 100
11.             JNA         NOT_DONE
12.             MOV         AL,SUM                       ;...then output SUM to port 3
13.             OUT         PORT_VAL,AL
14.             HLT                                      ;...and stop execution
15. NOT_DONE:   IN          AL,PORT_VAL                  ;otherwise add next input
16.             ADD         SUM,AL
17.             JMP         CYCLE                        ;and repeat the test
18. MY_CODE     ENDS
19.             END         GO                           ;this is the end of the assembly
```

Intel 8088 Int/Exceptions

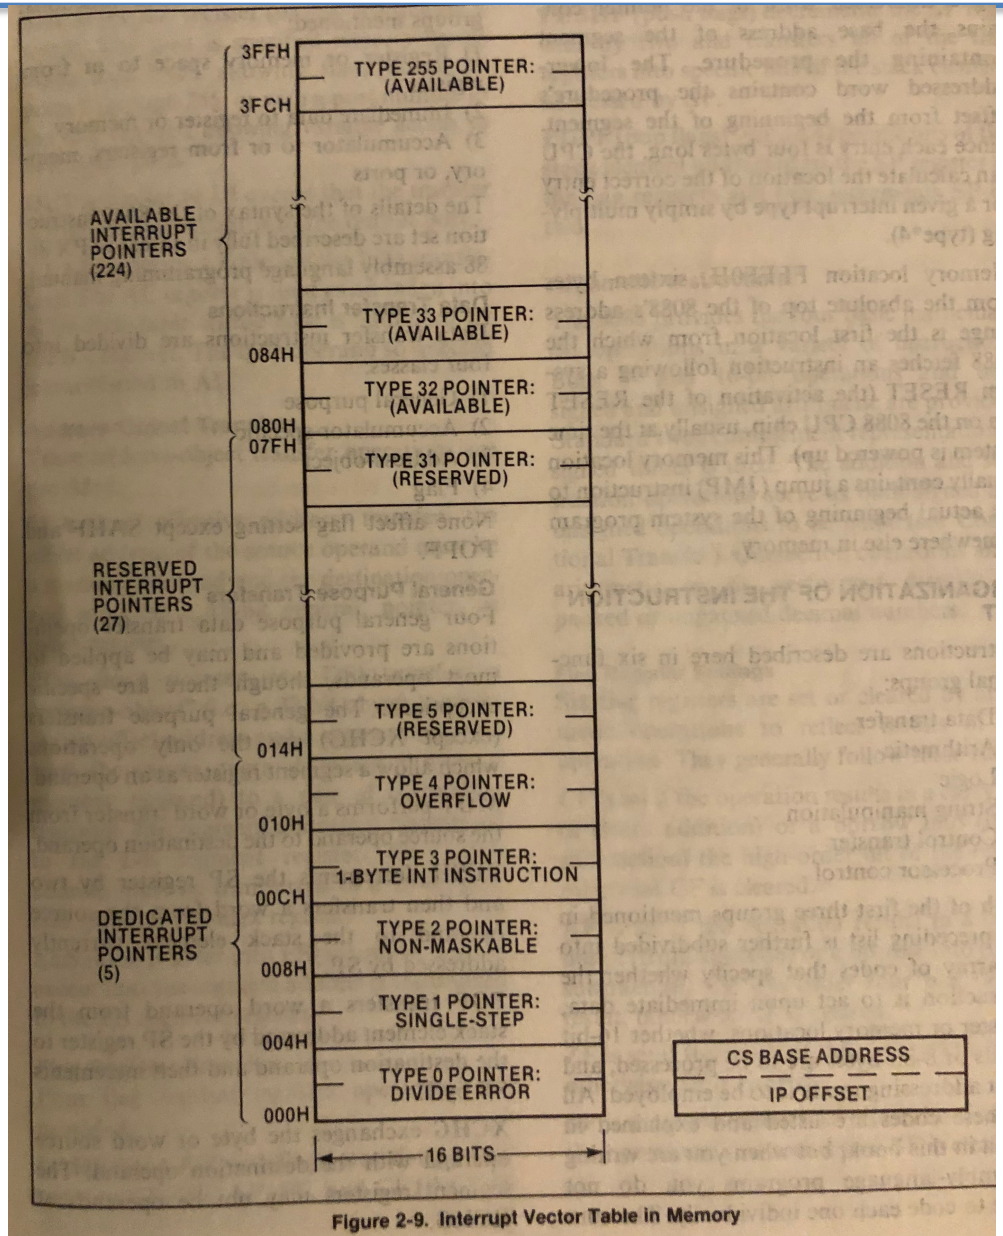

Figure 2-9. Interrupt Vector Table in Memory

Computer Architecture

Intel CPU's

Modern CPUs

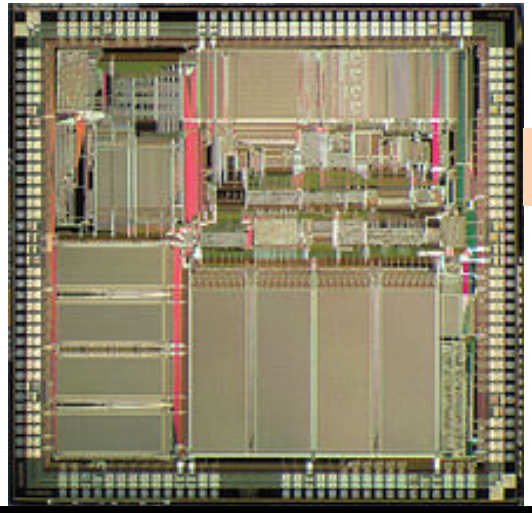

ARM 610

AMD “Mullins” APU

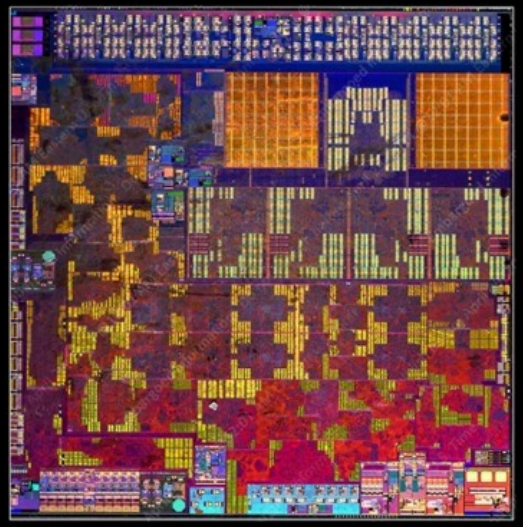

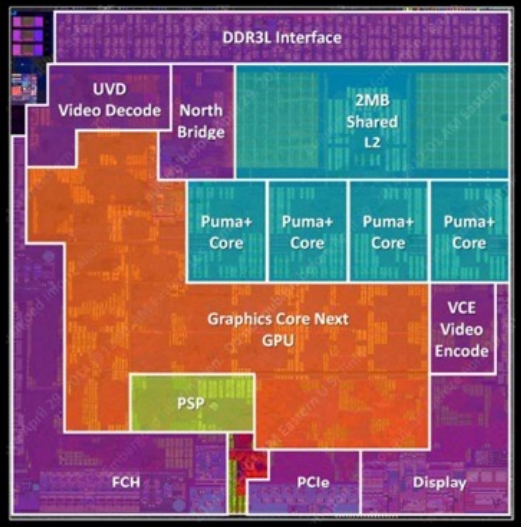

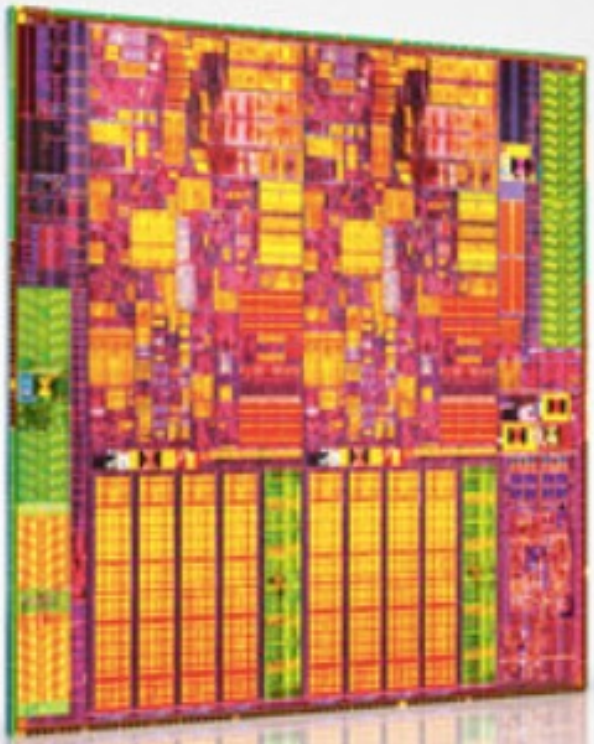

Intel “Westmere” CPU

The Westmere Die, a processor introduced by Intel in 2010. Intel

Intel Pentium "M"

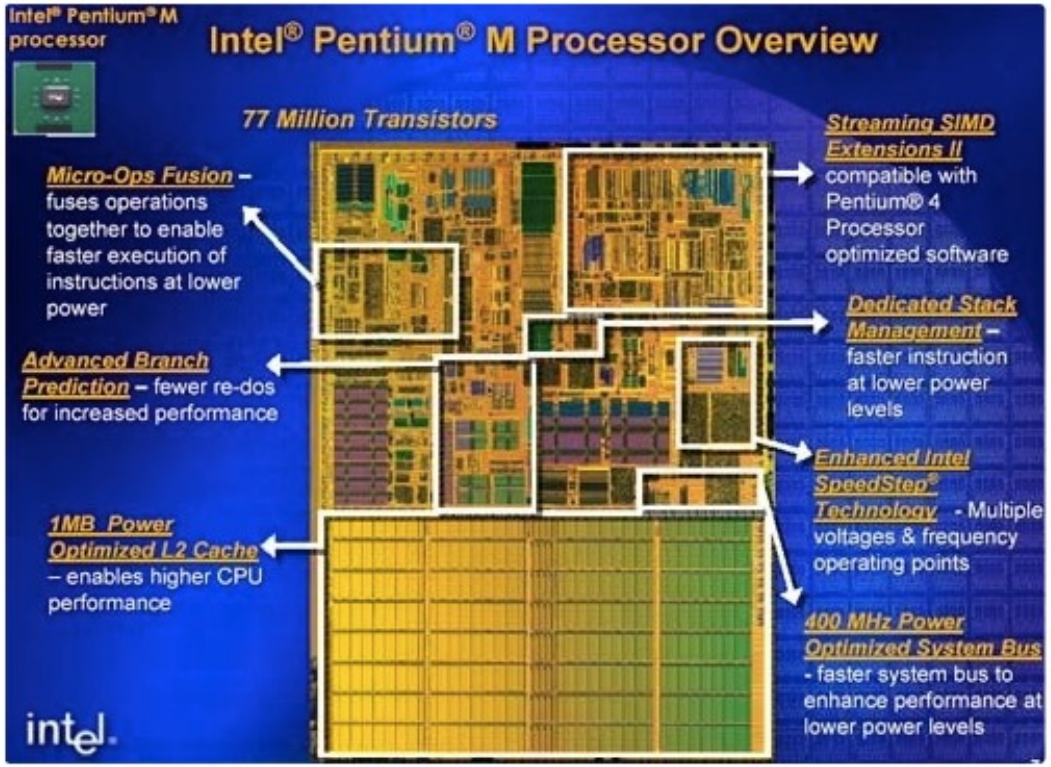

The arrival of the Pentium M on the market immediately gave Intel back the performance lead in the mobile segment. A second revision, the "Dothan," increased this lead. The Pentium M was so superior to even desktop Pentium 4's that many gamers bought them and ran them on desktop motherboards released for exactly that purpose. Of course, it wasn't long before Intel started contemplating the release of a real desktop version of the Pentium M. When the next-generation Netburst design, called "Tejas," proved to a disaster (engineering samples used 150W at 2.8 GHz), Intel cancelled it and decided to make the Pentium M design the basis of all future Intel releases.

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP122

Intel Itanium: IA-64

IA-64

From Wikipedia, the free encyclopedia

Not to be confused with x86-64.

IA-64 (**Intel Itanium architecture**) is the instruction set architecture (ISA) of the Itanium family of 64-bit Intel microprocessors. The basic ISA specification originated at Hewlett-Packard (HP), and was evolved and then implemented in a new processor microarchitecture by Intel with HP's continued partnership and expertise on the underlying EPIC design concepts. In order to establish what was their first new ISA in 20 years and bring an entirely new product line to market, Intel made a massive investment in product definition, design, software development tools, OS, software industry partnerships, and marketing. To support this effort Intel created the largest design team in their history and a new marketing and industry enabling team completely separate from x86. The first Itanium processor, codenamed *Merced*, was released in 2001.

The Itanium architecture is based on explicit instruction-level parallelism, in which the compiler decides which instructions to execute in parallel. This contrasts with superscalar architectures, which depend on the processor to manage instruction dependencies at runtime. In all Itanium models, up to and including *Tukwila*, cores execute up to six instructions per clock cycle.

In 2008, Itanium was the fourth-most deployed microprocessor architecture for enterprise-class systems, behind x86-64, Power ISA, and SPARC.[1]

Intel CPU Cores

You have to keep in mind that chips from the 7th, 8th, 9th and 11th gen series are all derivatives of the Skylake architecture.

This is the Core i7–7700K which was basically a re-spin of the Core i7–6700K with higher clock speeds and minor improvements:

With the 8th gen, Intel glued two more Skylake cores for a total of 6 (eg: Core i7–8700K). Later on, they added two additional cores for the 9th gen, bringing the core count to 8 (i9–9900K) and did the same thing to the 10th gen which maxed out at 10 cores.

Intel CPU Cores

Those Sunny Cove cores were originally designed for Intel's 10th gen Ice Lake laptop processors on the 10nm process but Intel backported it to 14nm which is 2.7X less dense and called it Cypress Cove. This implies that more silicon area is now required for each core.

- It uses a new Intel Xe GPU with 32EUs, this was again a backport but from their 11th gen Tiger Lake laptop processors.

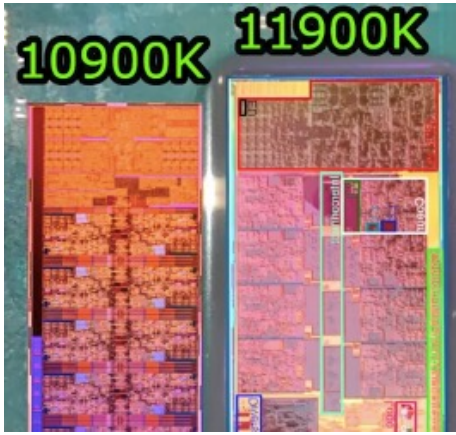

Moreover, Intel added the following:

- 28 PCIe Gen4 lanes, of which 20 are usable and 8 are reserved for DMI. This is up from the 20 Gen3 lanes (16 usable) of the previous generation
- Updated fixed-function media encoders, a new display interface with native HDMI 2.0b support (up from 1.4a)

All this means that the new chip is now 28% larger than the 10-core i9–10900K:

Intel CPU Cores

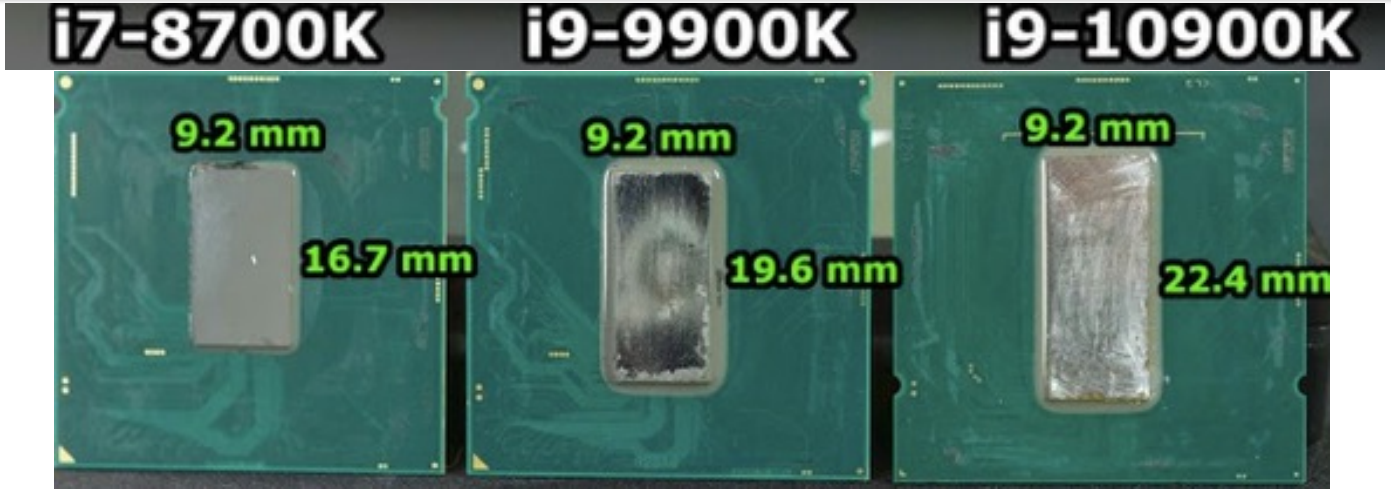

By that time, the die size was already becoming huge and the fact that Intel was reusing their 2015 Skylake architecture meant that desktop IPC had effectively stagnated despite the clock and memory speed increases.

Then came Rocket Lake which is effectively a stop gap and short lived product. But in order to remain competitive, Intel had to resort to backporting.

The 11th gen Rocket is thus a mix of several technologies:

- It uses Sunny Cove cores from 2018 which are larger. A single Sunny Cove core requires 300M transistors while the older Skylake ones require 217M. This is Intel's first new architecture on the desktop since Skylake which debuted in 2015 and it also includes AVX-512 which requires additional silicon

Intel CPU Cores

Intel could have added 2 more cores but this would have required an even larger die which would not make any sense given that performance would then be limited by thermal and cooling issues (Rocket Lake is already a power hog with AVX-512):

*That's the Core i9–11900K pulling almost **300W** when running AVX-512 workloads (source).*

Since each core now pulls nearly 50W of power alone, a 10 cores variant would consume nearly 350W which is above the limit for standard cooling.

Intel Itanium

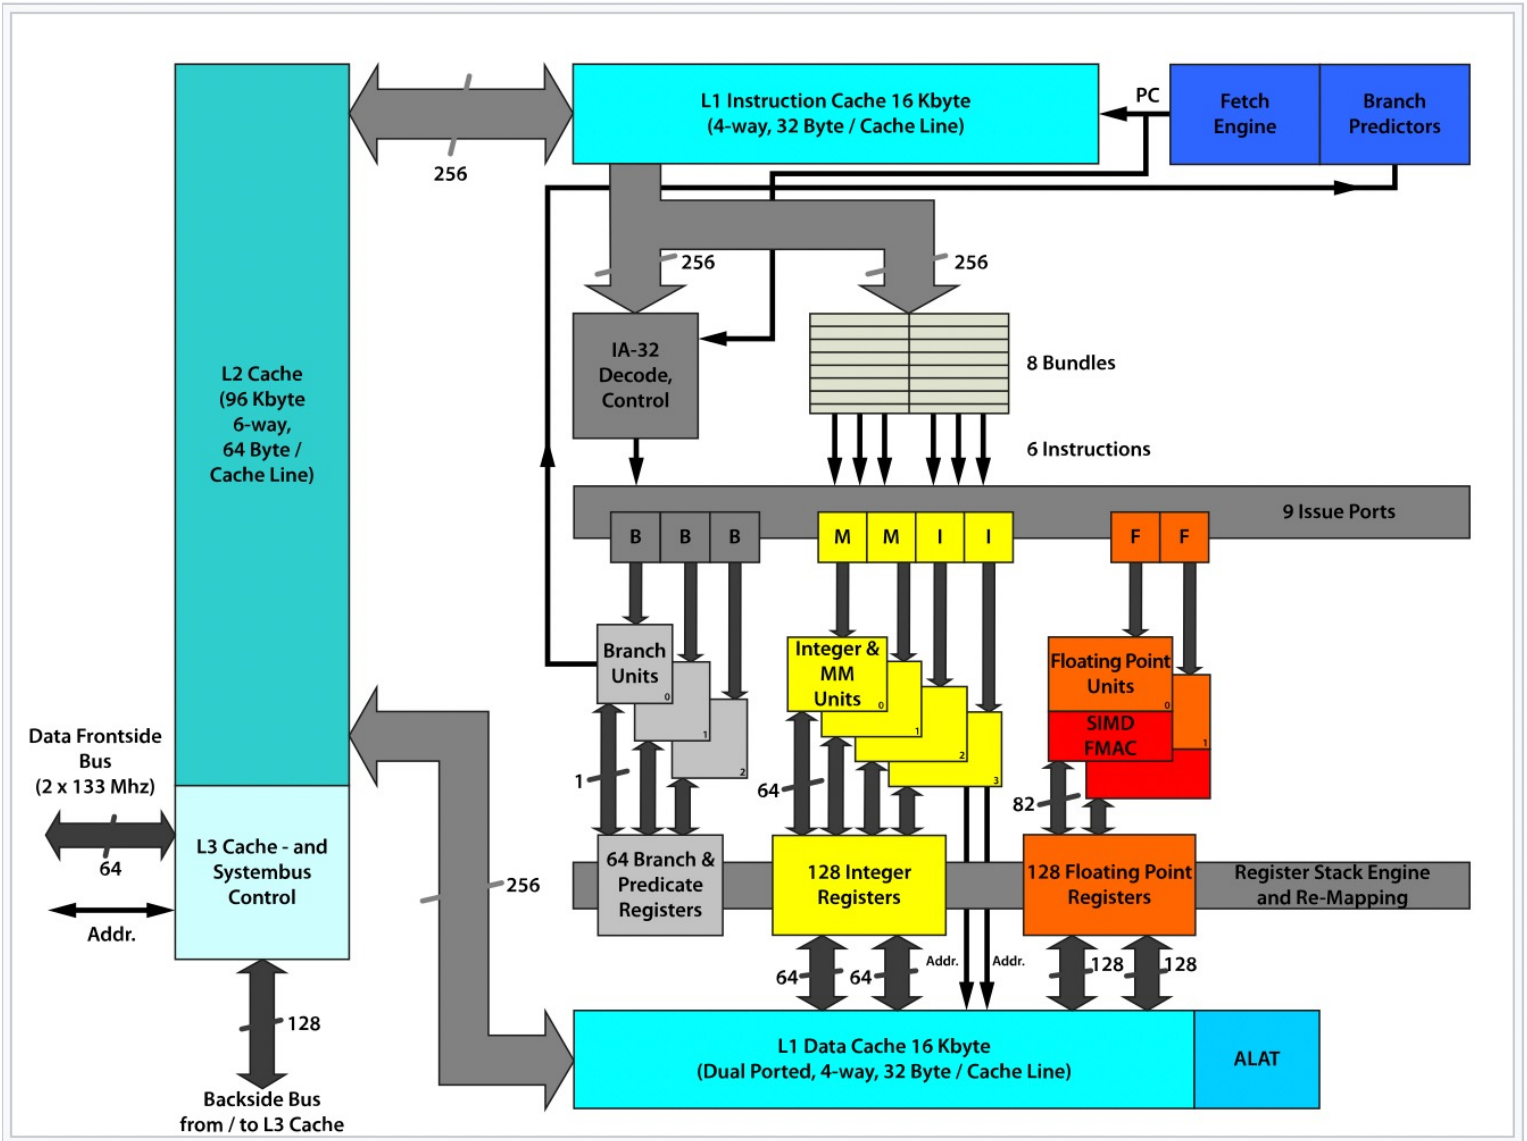

The Intel Itanium architecture

Intel CPU's

Xeon

back Intel® Xeon® Processors

3rd Generation Intel® Xeon® Scalable Processors	Intel® Xeon® Processor E5 v3 Family
2nd Generation Intel® Xeon® Scalable Processors	Intel® Xeon® Processor E5 v2 Family
Intel® Xeon® Scalable Processors	Intel® Xeon® Processor E5 Family
Intel® Xeon® E Processor	Intel® Xeon® Processor E3 v6 Family
Intel® Xeon® W Processor	Intel® Xeon® Processor E3 v5 Family
Intel® Xeon® D Processor	Intel® Xeon® Processor E3 v4 Family
Intel® Xeon® Processor E7 v4 Family	Intel® Xeon® Processor E3 v3 Family
Intel® Xeon® Processor E7 v3 Family	Intel® Xeon® Processor E3 v2 Family
Intel® Xeon® Processor E7 v2 Family	Intel® Xeon® Processor E3 Family
Intel® Xeon® Processor E7 Family	Legacy Intel® Xeon® Processors
Intel® Xeon® Processor E5 v4 Family	

Intel CPU's

back Processors

Intel® Core™ Processors	Intel® Pentium® Processor	Intel® Celeron® Processor
Intel® Xeon® Processors	Intel® Xeon Phi™ Processors	Intel® Itanium® Processor
Intel Atom® Processor	Intel® Quark™ SoC	

Intel CPU's

New Features

6TH GEN INTEL® CORE PROCESSOR™ FEATURES AT A GLANCE

FEATURES[1]	BENEFITS
Intel® Active Management Technology (Intel® AMT)[15]	Using built-in platform capabilities and popular third-party management and security applications, Intel AMT allows IT to discover, heal, and protect computing assets on wired and wireless networks. Intel AMT is supported on platforms that have Intel® vPro™.
Intel® Rapid Storage Technology (Intel® RST)[17]	Offers excellent levels of performance, responsiveness, and expandability. Take advantage of the enhanced performance and lower power consumption available with Intel® RST with one or more SATA or PCIe storage drives. With additional SATA drives, Intel® RST provides quicker access to digital photo, video, and data files with RAID 0, 5, and 10, and greater data protection against a storage disk drive failure with RAID 1, 5, and 10. Dynamic Storage Accelerator unleashes the maximum performance of Solid State Drives (SSD) when multitasking.[17]
Intel® Speed Shift Technology	Delivers dramatically quicker responsiveness with single-threaded, transient (short duration) workloads, such as web browsing, by allowing the processor to more quickly select its best operating frequency and voltage for optimal performance and power efficiency.

Intel CPU Prices

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP122

Intel BIOS

Intel® Server Board S2600WT BIOS and Firmware Update for EFI

Version: R01.01.0029 (Latest) Date: 7/6/2020

Available Downloads

S2600WT_EFI_BIOS_R01.01.0029_M

OS Independent
Language: English
Size: 23.59 MB
MD5: 5e8be6ce68a07a74168e20a4a8d3c9fb

Download

Other Versions

R01.01.0028

Detailed Description

uEFI System Update Package (SUP) Installation and Operation Instructions

Intel highly recommends reading this document in its entirety before performing the system update. Verifying that your system meets the documented requirements will ensure a successful update and provide the most reliable system
functionality after the update has completed.

Released - Jul. 1 2020

This update package includes the following production level system software updates and update utilities:

System BIOS - 01.01.0029
ME Firmware - 03.01.03.072
BMC Firmware - 01.60.12355
FRUSDR - 1.19

iFlash32.efi - Version 14.1 Build 15
fwpiaupd.efi - Version 14.1 Build 18
frusdr.efi - Version 14.1 Build 18

BIOS

You literally have everything you need right in the BIOS of your motherboard.

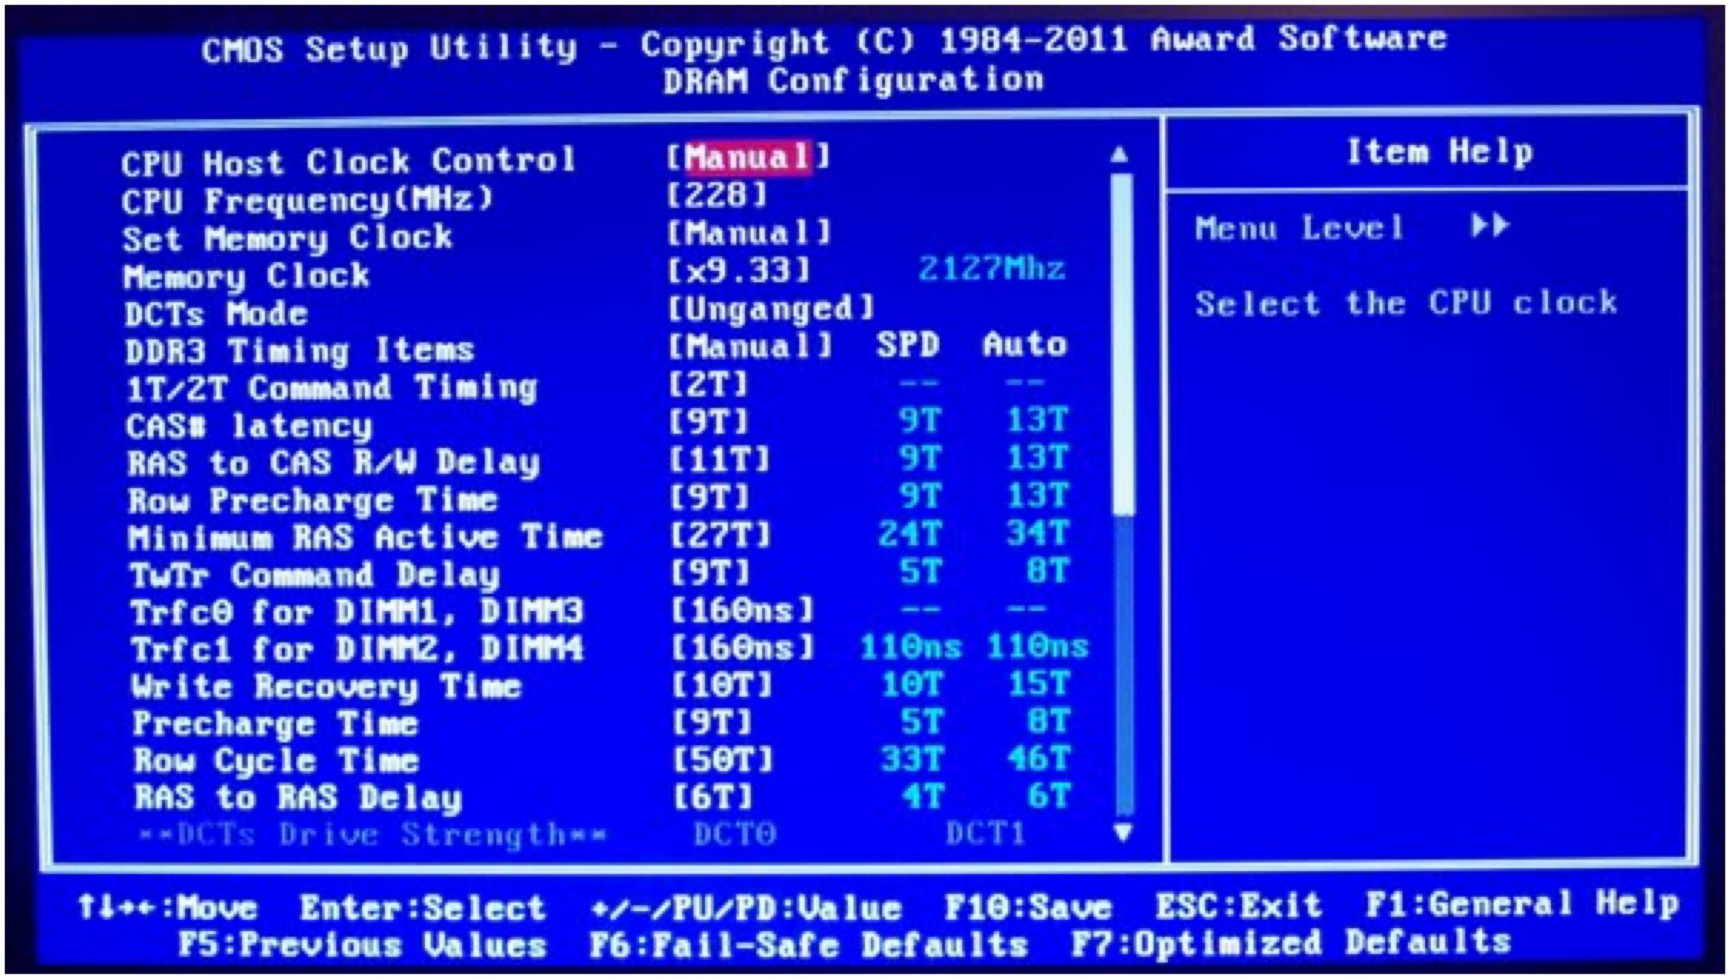

Intel CPU's

Xeon

3rd Generation Intel® Xeon® Scalable Processors

Filter: **View All** | Server

Product Name	Status	Launch Date	# of Cores	Max Turbo Frequency	Processor Base Frequency	Cache
Intel® Xeon® Platinum 8380HL Processor	Launched	Q2'20	28	4.30 GHz	2.90 GHz	38.5 MB
Intel® Xeon® Platinum 8380H Processor	Launched	Q2'20	28	4.30 GHz	2.90 GHz	38.5 MB
Intel® Xeon® Platinum 8376HL Processor	Launched	Q2'20	28	4.30 GHz	2.60 GHz	38.5 MB
Intel® Xeon® Platinum 8376H Processor	Launched	Q2'20	28	4.30 GHz	2.60 GHz	38.5 MB
Intel® Xeon® Platinum 8360HL Processor	Launched	Q3'20	24	4.20 GHz	3.00 GHz	33 MB
Intel® Xeon® Platinum 8360H Processor	Launched	Q3'20	24	4.20 GHz	3.00 GHz	33 MB
Intel® Xeon® Platinum 8356H Processor	Launched	Q3'20	8	4.40 GHz	3.90 GHz	35.75 MB
Intel® Xeon® Platinum 8354H Processor	Launched	Q2'20	18	4.30 GHz	3.10 GHz	24.75 MB
Intel® Xeon® Platinum 8353H Processor	Launched	Q2'20	18	3.80 GHz	2.50 GHz	24.75 MB

Intel CPU's

Intel® Quark™ Microcontroller D1000 Series

Filter: **View All** | Embedded

Product Name	Status	Launch Date	# of Cores	Processor Base Frequency	Cache	TDP
Intel® Quark™ Microcontroller D1000	Launched	Q3'15	1	33 MHz	0 KB	0.025 W

Intel® Quark™ Microcontroller D2000 Series

Filter: **View All** | Embedded

Product Name	Status	Launch Date	# of Cores	Processor Base Frequency	Cache
Intel® Quark™ Microcontroller D2000	Launched	Q3'15	1	32 MHz	0 KB

Product Name	Status	Launch Date	# of Cores	Processor Base Frequency	Cache
Intel® Quark™ SE C1000 Microcontroller	Launched	Q4'15	1	32 MHz	8 KB

Intel CPU's

Yowan Rajcoomar, Computer Technician (2008-present)

Answered 5h ago

The 11th gen Tiger Lake series get Xe graphics which can scale to 96EUs (Iris G7) but performance is still lacklustre due to the limited memory bandwidth and it can barely keep up with an MX150.

> Unoptimized drivers, a shared TDP pool, and lack of VRAM on the Iris Xe G7 may be attributing to some steep performance deficits when compared to the GeForce MX150, MX250, or MX350 series when running certain games. A select few titles, however, show the Intel chip to be on par with Nvidia.

Source: Tiger Lake Iris Xe G7 and GeForce MX350 are neck-to-neck in 3DMark — so why is the Intel chip still so much ... (more)

Intel i3/5/7 CPU Models

Jae Alexis Lee · April 19, 2017
I torture CPUs on a regular basis

Are all of Intel's Core i7 processors faster than the i5 series, and are the i5 series better than the i3 lineups?

You have to be *really clear* about what you're talking about when you speak about i3 / i5 / i7 characteristics relative to each other. Broad sweeping statements will blow you up every time.

We know, for example, just by looking at benchmarks that an Intel Core i7-920 @ 2.67GHz gets blown out of the water by an Intel Core i3-7350K @ 4.20GHz. That's grossly unfair though, we're comparing a first generation i7 to a *seventh generation* i3.

We also know that comparing desktop chips to laptop chips is pretty unfair. After all if you put an Intel Core i7-7500U @ 2.70GHz next to an Intel Core i5-7600K @ 3.80GHz, the desktop chip is going to run the laptop chip into the ground.

When you compare Intel chips understand that there are multiple *classes* (S-class, U-class, X-class) within each product *tier* i3/i5/i7.) There are also multiple product *lines* with in a product *class* (T-series, K-series, HQ-series, etc.)

Intel CPU Models

Jae Alexis Lee · April 19, 2017
I torture CPUs on a regular basis

Here's an example of an S-class product hierarchy, slowest to fastest:

i3–7100T -> i3–7100 -> i3–7300 T -> i3–7300 -> i3–7320 -> i3–7350K -> i5–7400T -i5–7400 -> i5–7500T -> i5–7500 -> i5–7600T -> i5–7600 -> i5–7600K -> i7–7700T -> i7–7700 -> i7–7700K

As long as you stay in the right *class and generation* you can generally say that i7s are faster than i5s which are faster than i3s. You can also understand that chips that end with *T* are lower power versions of the chip of a given model number and therefor a little slower, while chips that end with a *K* are unlocked chips that, even before you overclock them, have a higher stock clock speed than the same part number without the K. So in order of performance: i7–7700T (low power) -> i7–7700 -> i7–7700K (unlocked).

Oh, BTW, X-class chips like the i7–5960X on my desk or an i7–6900K... there are no i5 or i3 versions of these chips. X-class chips don't share silicon with S-class or U-class chips, they share silicon with *Xeons*, typically Xeon E5s. (Of course, all of this changes with Skylake-X and Kaby Lake-X but shhh... that's another answer.)

Intel CPU's

Heikki Kultala, Technical leader, SoC architecture at Nokia (2020-present)

Answered May 14

X86 instructions can be anything between 1–15 bytes long, and there are tens of different instruction formats. This makes decoding of x86 very complex.
One of the main ideas of RISC was to make instruction decoding simple. So, based on instruction encoding, x86 is very clearly CISC.

X86 instruction set is very unorthogonal, many instructions operate implicitly on some one register. This further complicates instruction decoding. One of the ideas of RISC was that any instruction can operate on any register. So, x86 is very un-RISCy here, clearly CISC.

X86 instructions can access memory even multiple times. This means that they are hard to trivially pipeline. To pipeline these instructions, they have to be splitted into multiple micro-operations. One of the main ideas of RISC is that all instructions are trivially pipelineable, so x86 is also far from RISCs here, clearly CISC.

x86 has things like "rep movsb" which performs arbitrary-length mempcy with one instruction. Microcode is the only way to implement this kind of instructions. One of the main points of RISC was to get rid of microcode. So again, very un-RISCy, very clearly CISC.

All these are very clearly CISCy features, X86 is very clearly CISC.

Intel CPU's

Christopher F Clark, Yacc++ author, Intel chip designer, ex-Googler

Answered May 14

To the user, x86 architecture is *very* CISC. Inside, a chip, it can be a different story.

A modern high end Intel chip takes those CISC instructions and translates (decodes) them down into an internal RISC-like instruction set and executes those in parallel and can often "execute" (retire is Intel's term for when an instruction completes execution) as many as 4 at a time (that was what they could do in 2015 on Haswell class machines).

The computer only has to maintain the fiction that it did what the original CISC instructions specified and only to the extent that such results are well-defined. So, if it takes a short-cut in the RISC machine that you cannot detect, it is allowed to.

Of course, some of those short-cuts don't work quite the way the chip designers hoped and you get things like Meltdown and Spectre which exploit the fact that you can detect certain edge cases where the RISC machine behavior leaks out inappropriately.

Intel CPU's

Paul Olaru, Been using computers since I was 10, fiddling since I was 14

Answered May 14

I'd say CISC and give you only one (of several) reasons.

```
push dword [eax + 2 * ebx + 4]
```

This is an extreme CISC-only thing. You have a "push" instruction (which is already CISC-y, although *I think* some impure RISCs have adopted it), and the pushed value is in memory, not in a register (already CISC all the way). And the complex address calculation (RISCs can do register + constant offset at most) is another trademark of CISCs.

Intel Power

Brett Bergan, Building PC's for 25 years

Answered 4h ago

Two basic reasons:

1. Intel is still stuck on 14nm lithography, which uses about double the power per unit quantity of transistors as the 7nm equivalents. 7nm actually has nearly 4X the areal density as 14nm, but since the AMD 7nm is more like Intel 10nm it does not directly compute to a purely quantifiable ratio.
2. Intel shoots for higher clock speeds. Wattage at clock speeds above the optimal threshold goes up exponentially.

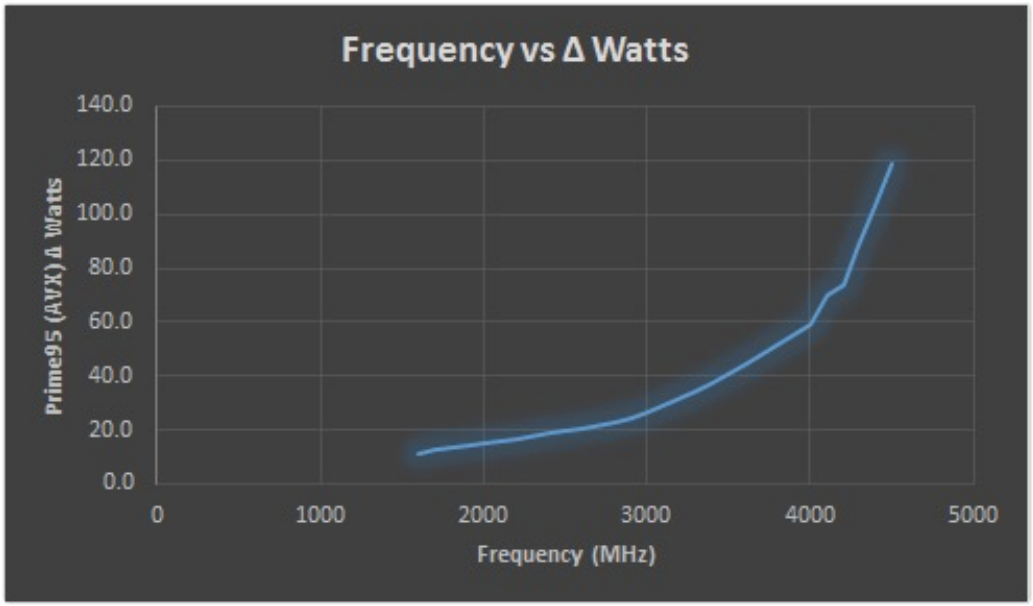

Once you hit 4GHz, power requirements start going through the roof. Keep in mind that some Intel CPU's for Gen 10 can surge to 5.3GHz (single-core) for brief moments.

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP122

Intel Core i3/5/7/9 + Xeon

Intel® Core™

- Intel® Core™ X-Series
- 9th Gen Intel® Core™ i9
- 10th Gen Intel® Core™ i7
- 10th Gen Intel® Core™ i5
- 10th Gen Intel® Core™ i3
- 8th Gen Intel® Core™ m3
- 9th Gen Intel® Core™ vPro™

1 or 2 **Threads**/Core

1.75-3B Transistors

Intel® Xeon®

- Intel® Xeon® Scalable
- Intel® Xeon® D
- Intel® Xeon® W
- Intel® Xeon® E

Intel Atom®

- Intel Atom® C
- Intel Atom® E
- Intel Atom® X

Pentium®

Celeron®

Itanium®

Intel® Quark™

Bottom of an LGA1151 CPU.

Top of an Intel Core i7-6700K (6th Gen).

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP122

Intel Core i3/5

2-4 Cores

10th Generation Intel® Core™ i3 Processors

Filter: View All | Mobile

Product Name	Status	Launch Date	# of Cores	Max Turbo Frequency	Processor Base Frequency	Cache
Intel® Core™ i3-10110U Processor	Launched	Q3'19	2	4.10 GHz	2.10 GHz	4 MB Intel® Smart Cache

Core i3	Clarkdale	2	32 nm	January 2010
	Sandy Bridge	2	32 nm	February 2011
	Ivy Bridge	2	22 nm	September 2012
	Haswell	2	22 nm	September 2013
	Skylake	2	14 nm	September 2015
	Kaby Lake	2	14 nm	January 2017
	Coffee Lake	4	14 nm	October 2017
	Coffee Lake	4	14 nm	Jan. & April 2019

Intel Core i3/5

2-6 Cores

10th Generation Intel® Core™ i5 Processors

Filter: View All | Mobile

Product Name	Status	Launch Date	# of Cores	Max Turbo Frequency	Processor Base Frequency	Cache
Intel® Core™ i5-1035G7 Processor	Launched	Q3'19	4	3.70 GHz	1.20 GHz	6 MB Intel® Smart Cache

Core i5	Lynnfield	4	45 nm	September 2009
	Clarkdale	2	32 nm	January 2010
	Sandy Bridge	4	32 nm	January 2011
	Sandy Bridge	2	32 nm	February 2011
	Ivy Bridge	2-4	22 nm	April 2012
	Haswell	2-4	22 nm	June 2013
	Broadwell	4	14 nm	June 2015
	Skylake	4	14 nm	September 2015
	Kaby Lake	4	14 nm	January 2017
	Coffee Lake	6	14 nm	October 2017
	Coffee Lake	6	14 nm	Oct. 2018 & Jan. 2019

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP122

Intel Core i7

DR JEFF SOFTWARE
INDIE APP DEVELOPER

4-8 Cores

10th Generation Intel® Core™ i7 Processors

Filter: View All | Mobile

Product Name	Status	Launch Date	# of Cores	Max Turbo Frequency	Processor Base Frequency	Cache
Intel® Core™ i7-10710U Processor	Launched	Q3'19	6	4.70 GHz	1.10 GHz	12 MB Intel® Smart Cache
Intel® Core™ i7-1065G7 Processor	Launched	Q3'19	4	3.90 GHz	1.30 GHz	8 MB Intel® Smart Cache
Intel® Core™ i7-7820X X-series Processor	Launched	Q2'17	8	4.30 GHz	3.60 GHz	11 MB L3 Cache
Intel® Core™ i7-7800X X-series Processor	Launched	Q2'17	6	4.00 GHz	3.50 GHz	8.25 MB L3 Cache

Core i7	Bloomfield	4	45 nm	November 2008
	Lynnfield	4	45 nm	September 2009
	Gulftown	6	32 nm	July 2010
	Sandy Bridge	4	32 nm	January 2011
	Sandy Bridge-E	6	32 nm	November 2011
	Sandy Bridge-E	4	32 nm	February 2012
	Ivy Bridge	4	22 nm	April 2012
	Haswell	4	22 nm	June 2013
	Ivy Bridge-E	4-6	22 nm	September 2013
	Broadwell	4	14 nm	June 2015
	Skylake	4	14 nm	August 2015
	Kaby Lake	4	14 nm	January 2017
	Coffee Lake	6	14 nm	October 2017
	Coffee Lake	8	14 nm	October 2018

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP122

Intel Core i9

6-18 Cores

Intel® Core™ X-series Processors

Filter: View All | Desktop | Mobile | Retail Box

Product Name	Status	Launch Date	# of Cores	Max Turbo Frequency	Processor Base Frequency	Cache
Intel® Core™ i9-10980XE Extreme Edition Processor	Launched	Q4'19	18	4.60 GHz	3.00 GHz	24.75 MB Intel® Smart Cache
Intel® Core™ i9-10940X X-series Processor	Launched	Q4'19	14	4.60 GHz	3.30 GHz	19.25 MB Intel® Smart Cache
Intel® Core™ i7-7820X X-series Processor	Launched	Q2'17	8	4.30 GHz	3.60 GHz	11 MB L3 Cache
Intel® Core™ i7-7800X X-series Processor	Launched	Q2'17	6	4.00 GHz	3.50 GHz	8.25 MB L3 Cache

Core i9	Skylake-X	10	14 nm	June 2017
	Skylake-X	12	14 nm	August 2017
	Skylake-X	14-18	14 nm	September 2017
	Coffee Lake	8	14 nm	October 2018

Intel Core 10th Gen

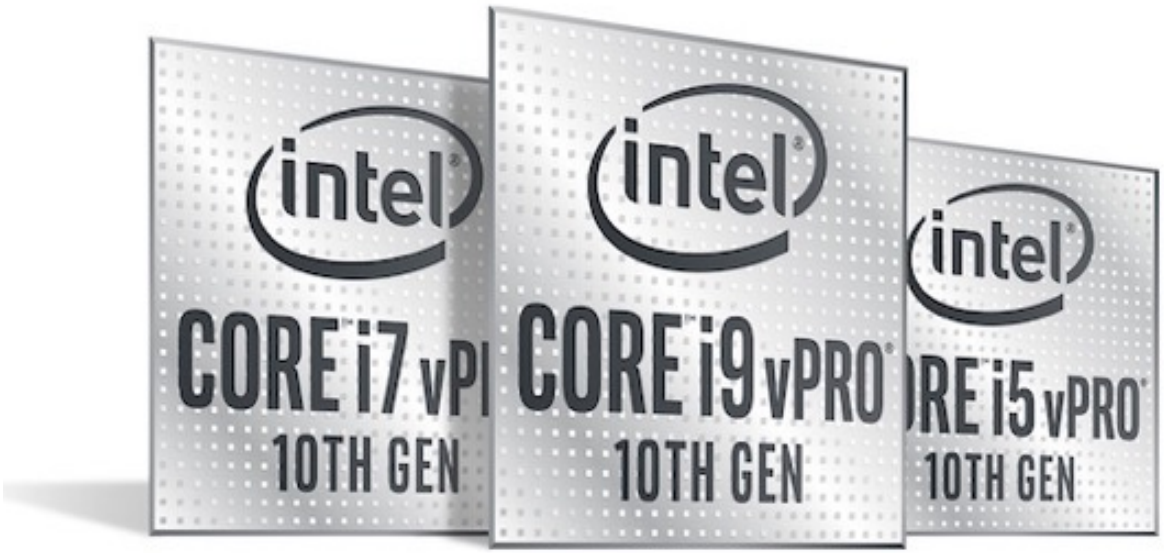

Source: Intel

New Intel vPro Platform Bolsters the Modern Workforce

Remote work may be our new normal, but that doesn't mean it's been easy. Enter: 10th-generation Intel vPro processors. Designed to deliver better connectivity, productivity, and remote manageability, the new Intel vPro platform offers hardware-based security features and the best Wi-Fi technology for video calls.

Intel Core 10th Gen

But the chances of Intel retaining that lead past November 5, 2020, aren't looking all that great. That's given its recent announcement that its coming 11th Generation Rocket Lake line of 14nm desktop CPUs won't be launching until sometime in the first quarter of 2021, well after the holiday season has already passed.

That's late, relative to Zen 3. But just as important: Will Intel's refined 14nm+++ process be able to continue leapfrogging AMD once Zen 3 is launched? Time (along with our benchmarking suite) will tell that story with more clarity. But from a distance and based on what we know from specs and AMD's claims today, don't expect a massive reversal of fortune. The next year of desktop CPU prospects does not look promising for Intel,

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP122

Intel CPU Cores

DR JEFF SOFTWARE
INDIE APP DEVELOPER

10TH GEN INTEL® CORE™ DESKTOP PROCESSORS

PROCESSOR NUMBER	BASE CLOCK SPEED (GHZ)	INTEL® TURBO BOOST TECHNOLOGY 2.0 MAXIMUM SINGLE CORE TURBO FREQUENCY (GHZ)	INTEL® TURBO BOOST MAX TECHNOLOGY 3.0 FREQUENCY (GHZ)	INTEL® THERMAL VELOCITY BOOST TECHNOLOGY SINGLE / ALL CORE TURBO FREQUENCY (GHZ)[4]	INTEL® ALL CORE TURBO FREQUENCY (GHZ)	CORES/ THREADS	THERMAL DESIGN POWER	UNLOCKED[2]	PLATFORM PCIE 3.0 LANES	MEMORY SUPPORT[3]	PROCESSOR GRAPHICS	INTEL® OPTANE™ MEMORY[9]	RCP PRICING (USD 1K)
i9-10900K	Up to 3.7	Up to 5.1	Up to 5.2	Up to 5.3 / 4.9	Up to 4.8	10/20	125	✔	Up to 40	Two Channels DDR4-2933	Intel® UHD Graphics 630	✔	$488
i9-10900KF	Up to 3.7	Up to 5.1	Up to 5.2	Up to 5.3 / 4.9	Up to 4.8	10/20	125	✔	Up to 40	Two Channels DDR4-2933		✔	$472
i9-10900	Up to 2.8	Up to 5.0	Up to 5.1	Up to 5.2 / 4.6	Up to 4.5	10/20	65		Up to 40	Two Channels DDR4-2933	Intel® UHD Graphics 630	✔	$439
i9-10900F	Up to 2.8	Up to 5.0	Up to 5.1	Up to 5.2 / 4.6	Up to 4.5	10/20	65		Up to 40	Two Channels DDR4-2933		✔	$422
i7-10700K	Up to 3.8	Up to 5.0	Up to 5.1	NA	Up to 4.7	8/16	125	✔	Up to 40	Two Channels DDR4-2933	Intel® UHD Graphics 630	✔	$374
i7-10700KF	Up to 3.8	Up to 5.0	Up to 5.1	NA	Up to 4.7	8/16	125	✔	Up to 40	Two Channels DDR4-2933		✔	$349
i7-10700	Up to 2.9	Up to 4.7	Up to 4.8	NA	Up to 4.6	8/16	65		Up to 40	Two Channels DDR4-2933	Intel® UHD Graphics 630	✔	$323
i7-10700F	Up to 2.9	Up to 4.7	Up to 4.8	NA	Up to 4.6	8/16	65		Up to 40	Two Channels DDR4-2933		✔	$298

Intel® processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. All processors are lead-free (per EU RoHS directive July 2006) and halogen free (residual amounts of halogens are below November 2007 proposed IPC/JEDEC J-STD-709 standards). All processors support Intel® Virtualization Technology (Intel® VT-x).

- DDR4 maximum speed support is 1 and 2 DPC for UDIMMs but only 1 DPC for SODIMMs. DDR4 2DPC UDIMM 2933 or 2666 is capable when same UDIMM part number are populated with in each channel.
- Intel® Optane™ memory requires specific hardware and software configuration. Visit www.intel.com/OptaneMemory for configuration requirements.
- Intel® Thermal Velocity Boost feature is opportunistic at a temperature of 70°C or lower and when turbo power budget is available. The frequency gain and duration is dependent on the workload (best for bursty workloads), capabilities of the individual processor, and the processor cooling solution. Frequencies may reduce over time and longer workloads may start at the max frequency but drop as processor temperature increases.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No product or component can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

Intel Cores

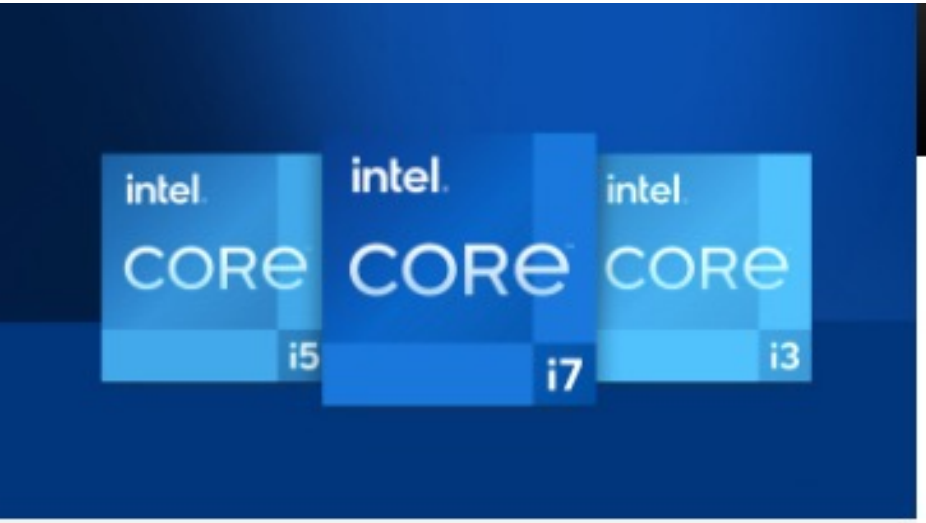

Processor

11th Gen Intel® Core™ processors: Break the boundaries of performance →

CPU

Graphics

Discover Intel® Iris® Xe graphics →

GPU

New Intel Core i7

UserBenchmark

11th gen

Intel Core i7-11700K $400

The i7-11700K is the second fastest CPU in Intel's 11th Gen Rocket Lake-S lineup. Rocket Lake brings higher IPC (early samples indicate +19%) and 50% stronger integrated graphics. There are several 500 series chipset improvements including: 20 PCIe4 CPU lanes (up from 16) and USB 3.2 Gen 2x2 (20 Gbps up from 10 Gbps). Rocket Lake's IPC uplift translates to around a 10% faster Effective Speed than both Intel's 10th Gen and AMD's 5000 series. Despite Intel's performance lead, AMD continues to outsell Intel. Given Intel's colossal R&D operation, it's bewildering that their marketing remains so neglected. Little effort is

IPC

Intel CPU Cores

You have to keep in mind that chips from the 7th, 8th, 9th and 11th gen series are all derivatives of the Skylake architecture.

This is the Core i7–7700K which was basically a re-spin of the Core i7–6700K with higher clock speeds and minor improvements:

With the 8th gen, Intel glued two more Skylake cores for a total of 6 (eg: Core i7–8700K). Later on, they added two additional cores for the 9th gen, bringing the core count to 8 (i9–9900K) and did the same thing to the 10th gen which maxed out at 10 cores.

Intel CPU Cores

Those Sunny Cove cores were originally designed for Intel's 10th gen Ice Lake laptop processors on the 10nm process but Intel backported it to 14nm which is 2.7X less dense and called it Cypress Cove. This implies that more silicon area is now required for each core.

- It uses a new Intel Xe GPU with 32EUs, this was again a backport but from their 11th gen Tiger Lake laptop processors.

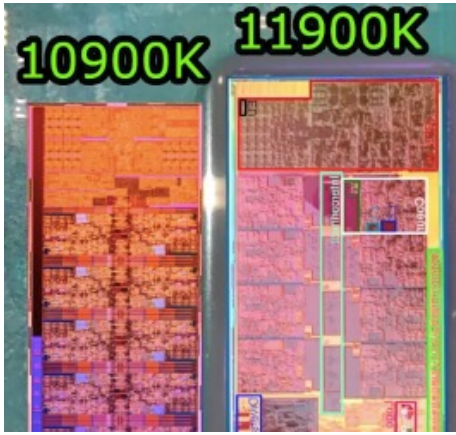

Moreover, Intel added the following:

- 28 PCIe Gen4 lanes, of which 20 are usable and 8 are reserved for DMI. This is up from the 20 Gen3 lanes (16 usable) of the previous generation
- Updated fixed-function media encoders, a new display interface with native HDMI 2.0b support (up from 1.4a)

All this means that the new chip is now 28% larger than the 10-core i9–10900K:

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE

COMP122

Intel CPU's

Core 11th Gen

11th Generation Intel® Core™ i3 Processors

Filter: **View All** | Embedded | Mobile

Product Name	Status	Launch Date	# of Cores	Max Turbo Frequency	Processor Base Frequency	Cache
Intel® Core™ i3-1125G4 Processor	Announced	Q3'20	4	3.70 GHz		8 MB Intel® Smart Cache
Intel® Core™ i3-1125G4 Processor	Announced	Q3'20	4	3.70 GHz		8 MB Intel® Smart Cache
Intel® Core™ i3-1120G4 Processor	Announced	Q3'20	4	3.50 GHz		8 MB Intel® Smart Cache
Intel® Core™ i3-1115GRE Processor	Launched	Q3'20	2	3.90 GHz	2.20 GHz	6 MB Intel® Smart Cache
Intel® Core™ i3-1115G4E Processor	Launched	Q3'20	2	3.90 GHz	2.20 GHz	6 MB Intel® Smart Cache
Intel® Core™ i3-1115G4 Processor	Launched	Q3'20	2	4.10 GHz		6 MB Intel® Smart Cache
Intel® Core™ i3-1115G4 Processor	Launched	Q3'20	2	4.10 GHz		6 MB Intel® Smart Cache
Intel® Core™ i3-1110G4 Processor	Launched	Q3'20	2	3.90 GHz		6 MB Intel® Smart Cache

Intel CPU's

Core 11th Gen

11th Generation Intel® Core™ i5 Processors

Filter: **View All** | Embedded | Mobile

Product Name	Status	Launch Date	# of Cores	Max Turbo Frequency	Processor Base Frequency	Cache
Intel® Core™ i5-1145GRE Processor	Launched	Q3'20	4	4.10 GHz	1.50 GHz	8 MB Intel® Smart Cache
Intel® Core™ i5-1145G7E Processor	Launched	Q3'20	4	4.10 GHz	1.50 GHz	8 MB Intel® Smart Cache
Intel® Core™ i5-1135G7 Processor	Launched	Q3'20	4	4.20 GHz		8 MB Intel® Smart Cache
Intel® Core™ i5-1135G7 Processor	Launched	Q3'20	4	4.20 GHz		8 MB Intel® Smart Cache
Intel® Core™ i5-1130G7 Processor	Launched	Q3'20	4	4.00 GHz		8 MB Intel® Smart Cache

Intel CPU's

Core 11th Gen

11th Generation Intel® Core™ i7 Processors

Filter: **View All** | Embedded | Mobile

Product Name	Status	Launch Date	# of Cores	Max Turbo Frequency	Processor Base Frequency	Cache
Intel® Core™ i7-1185GRE Processor	Launched	Q3'20	4	4.40 GHz	1.80 GHz	12 MB Intel® Smart Cache
Intel® Core™ i7-1185G7E Processor	Launched	Q3'20	4	4.40 GHz	1.80 GHz	12 MB Intel® Smart Cache
Intel® Core™ i7-1185G7 Processor	Launched	Q3'20	4	4.80 GHz		12 MB Intel® Smart Cache
Intel® Core™ i7-1165G7 Processor	Launched	Q3'20	4	4.70 GHz		12 MB Intel® Smart Cache
Intel® Core™ i7-1165G7 Processor	Launched	Q3'20	4	4.70 GHz		12 MB Intel® Smart Cache
Intel® Core™ i7-1160G7 Processor	Launched	Q3'20	4	4.40 GHz		12 MB Intel® Smart Cache

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE

COMP122

Intel 12th Gen

12th Gen Intel® Core™ Desktop Processors

12th Generation Intel® Core™ i7 Processors

Product brief: 12th Gen Intel® Core™ desktop processors →

Product brief: Intel® Z690 Chipset →

NOW!

2 Products COMPARE ALL

Compare	Product Name	Status	Launch Date	# of Cores	Max Turbo Frequency	Cache	Processor Graphics ‡
☐	Intel® Core™ i7-12700KF Processor (25M Cache, up to 5.00 GHz)	Launched	Q4'21	12	5.00 GHz	25 MB Intel® Smart Cache	
☐	Intel® Core™ i7-12700K Processor (25M Cache, up to 5.00 GHz)	Launched	Q4'21	12	5.00 GHz	25 MB Intel® Smart Cache	Intel® UHD Graphics 770

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE

COMP122

Intel 12th Gen

DR JEFF SOFTWARE
INDIE APP DEVELOPER

12th Gen Intel® Core™ Desktop Processors

12th Gen Intel® Core™ Desktop Processors: Features at a Glance

Feature	Benefit
Performance Hybrid Architecture	Performance hybrid architecture, combining Performance-cores (P-cores) and Efficient-cores (E-Cores) to deliver balanced single-thread and multi-threaded real-world performance.
Intel® Thread Director[1]	Optimizes workloads by helping the OS scheduler intelligently distribute workloads to the optimal cores.
PCIe 5.0 up to 16 Lanes	Offers readiness for up to 32 GT/s for fast access to peripheral devices and networking with up to 16 PCI Express 5.0 lanes.
PCIe 4.0 up to 4 Lanes	Offers up to 16 GT/s for fast access to peripheral devices and networking with up to 4 PCI Express 4.0 lanes.
Up to DDR5 4800 MT/s[2]	This industry first memory technology supports fast frequencies and high bandwidth and throughput leading to enhanced workflow and productivity.

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE

Intel 12th Gen

DR JEFF SOFTWARE
INDIE APP DEVELOPER

12th Gen Intel® Core™ Desktop Processors Comparisons

12th Gen Intel® Core™ Desktop Processors

	Intel® Core™ i9-12900K & i9-12900KF[4]	Intel® Core™ i7-12700K & i7-12700KF[4]	Intel® Core™ i5-12600K & i5-12600KF[4]
Max Turbo Frequency [GHz]	Up to 5.2	Up to 5.0	Up to 4.9
Intel® Turbo Boost Max Technology 3.0 Frequency [GHz]	Up to 5.2	Up to 5.0	n/a
Single P-core Turbo Frequency [GHz]	Up to 5.1	Up to 4.9	Up to 4.9
Single E-core Turbo Frequency [GHz]	Up to 3.9	Up to 3.8	Up to 3.6
P-core Base Frequency [GHz]	3.2	3.6	3.7

Intel 12th Gen

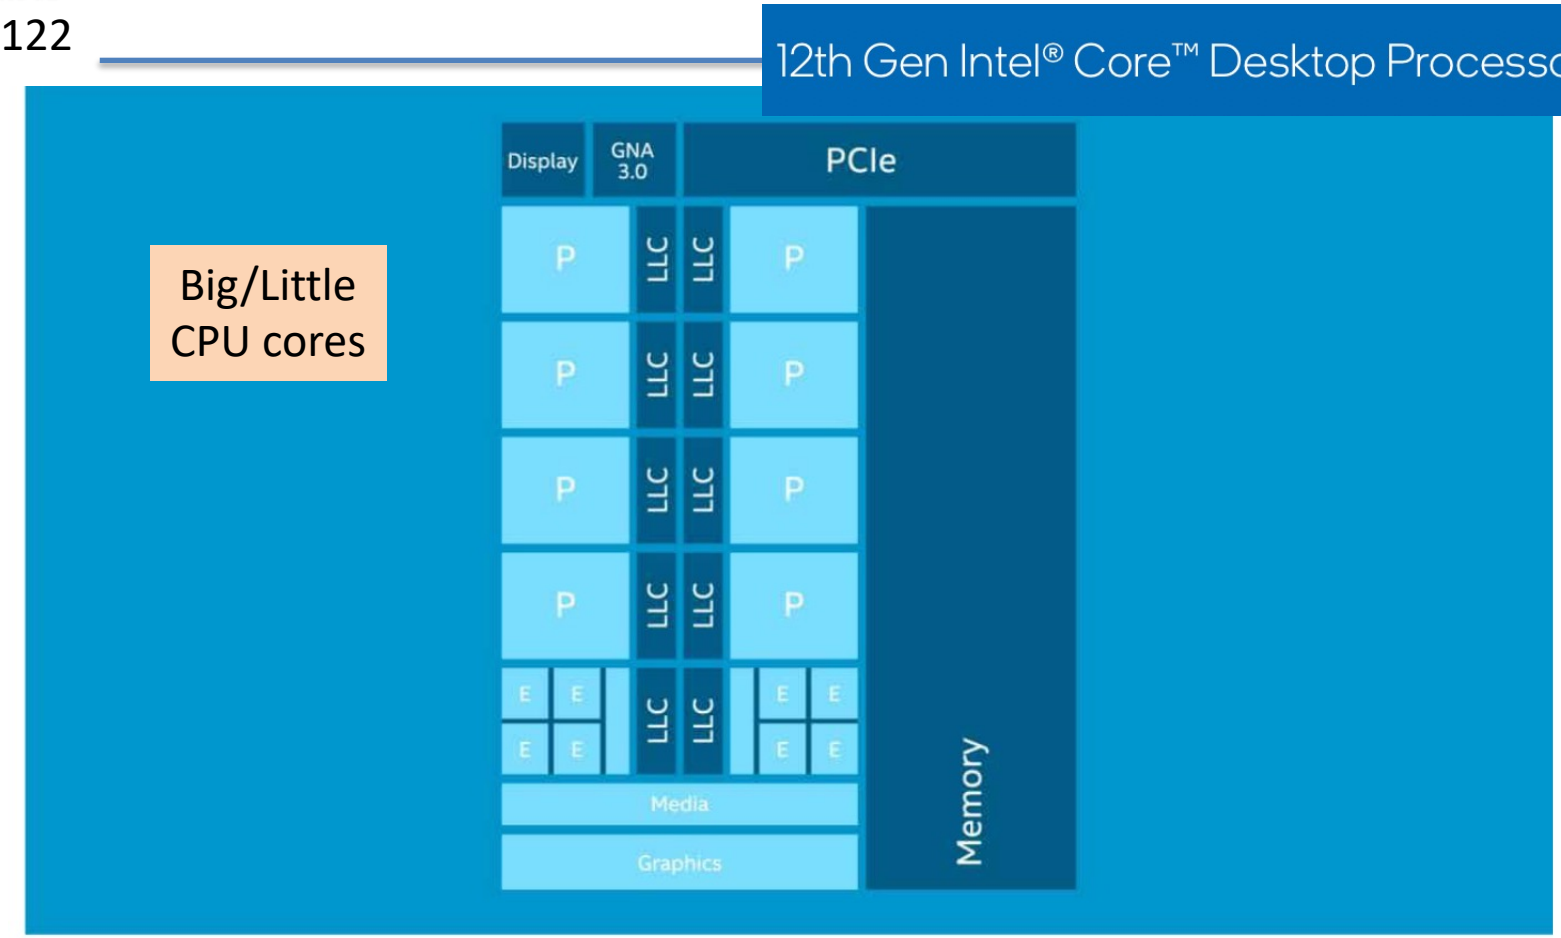

In a game-changing advance for core performance, 12th Gen Intel® Core™ desktop processors power a revolutionary approach to the x86 architecture. Its Performance-cores—or "P-cores"—are optimized for single & lightly-threaded performance, while its Efficient-cores—or "E-cores"—are optimized for scaling highly-threaded workloads. Intel® Thread Director helps to monitor and analyze performance data in real-time to seamlessly place the right application thread on the right core and optimize performance per watt.[1] That means gamers, creators, and

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP122

Intel 12th Gen

12th Gen Intel® Core™ Desktop Processors

12TH GEN INTEL® CORE™ DESKTOP PROCESSORS COMPARISONS

	Intel® Core™ i9-12900K & i9-12900KF[4]	Intel® Core™ i7-12700K & i7-12700KF[4]	Intel® Core™ i5-12600K & i5-12600KF[4]
Max Turbo Frequency [GHz]	Up to 5.2	Up to 5.0	Up to 4.9
Intel® Turbo Boost Max Technology 3.0 Frequency [GHz]	Up to 5.2	Up to 5.0	n/a
Single P-core Turbo Frequency [GHz]	Up to 5.1	Up to 4.9	Up to 4.9
Single E-core Turbo Frequency [GHz]	Up to 3.9	Up to 3.8	Up to 3.6
P-core Base Frequency [GHz]	3.2	3.6	3.7
E-core Base Frequency [GHz]	2.4	2.7	2.8
Processor Cores (P-cores + E-cores)	16 (8P + 8E)	12 (8P + 4E)	10 (6P + 4E)
Intel® Hyper-Threading Technology[5]	Yes	Yes	Yes
Total Processor Threads	24	20	16
Intel® Thread Director[1]	Yes	Yes	Yes
Intel® Smart Cache (L3) Size [MB]	30	25	20
Total L2 Cache Size [MB]	14	12	9.5

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE

COMP122

Intel 12th Gen

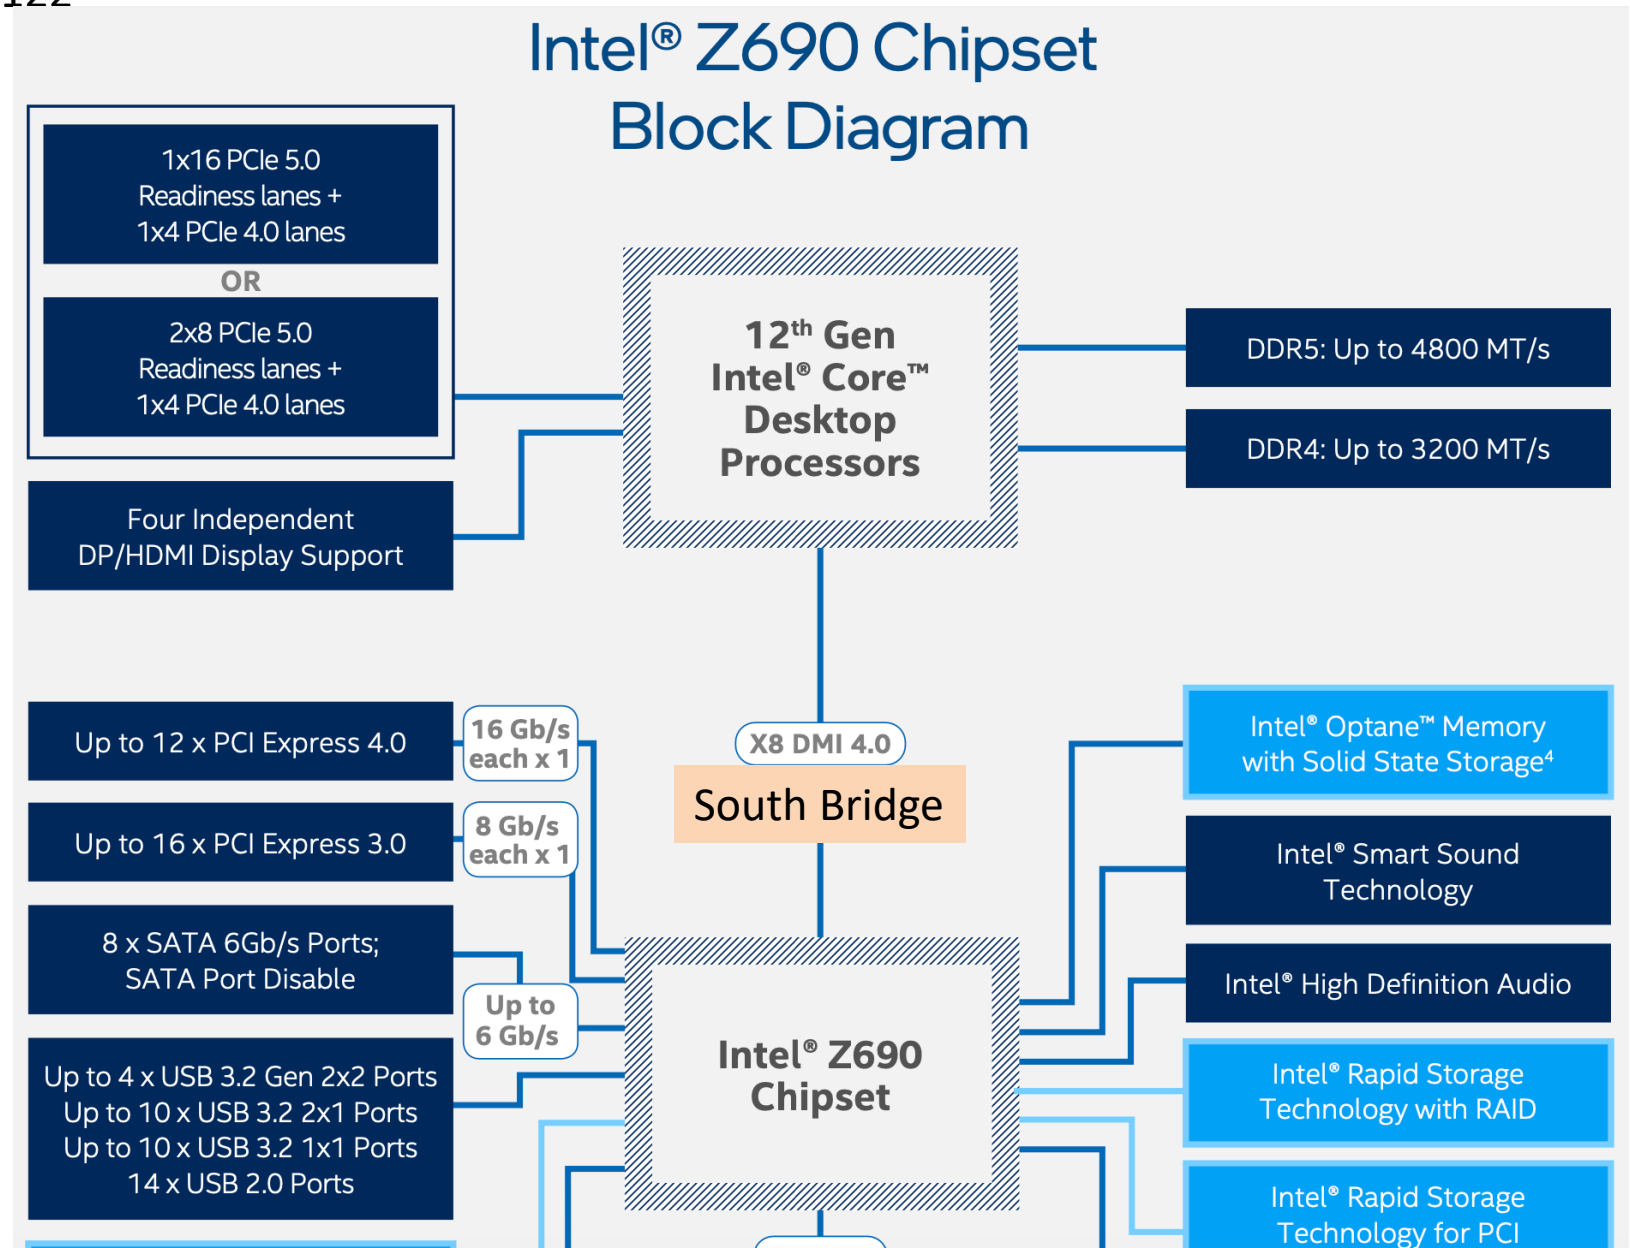

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP122

Intel 12th Gen Z690

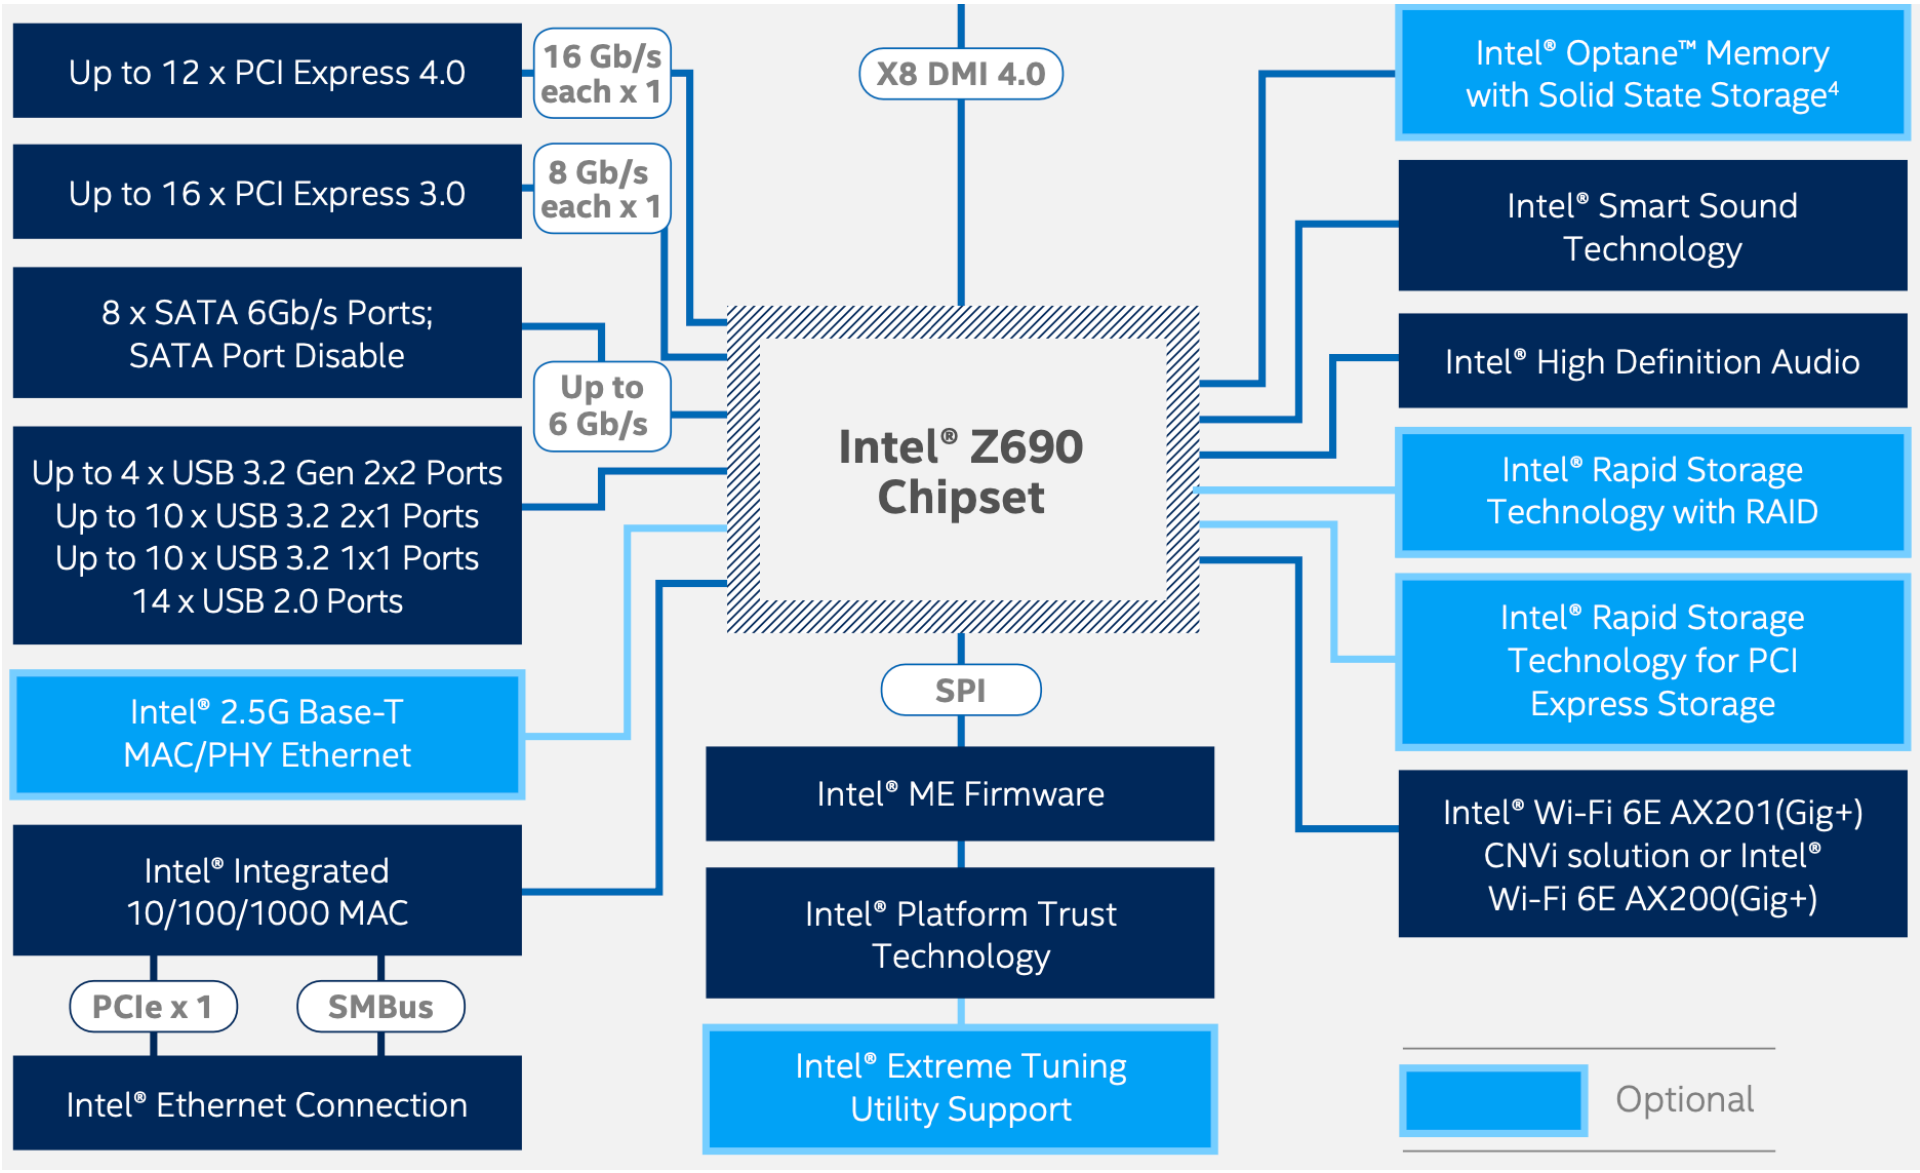

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE

COMP12?

Intel ‘s New 13th Gen

Intel Newsroom / **Intel Launches 13th Gen Intel Core Processors**

Oct 2022

Intel Launches 13th Gen Intel Core Processor Family Alongside New Intel Unison Solution

13th Gen Intel Core desktop processors deliver the world’s best gaming experience and unmatched overclocking capabilities.

Intel ‘s New 13th Gen

Oct 2022

Industry-Leading Features for Desktop Platforms

The 13th Gen Intel Core desktop processors empower users with leading-edge performance and experiences across gaming, content creation and work, with several new and improved features including:

- Intel® Adaptive Boost Technology and Thermal Velocity Boost opportunistically boosts processor clock frequencies based on power and thermal headroom during a given workload. Available in Intel Core i9 unlocked SKUs.
- More E-cores across Intel Core i5, i7, i9 power a big leap in multi-threaded performance and better multi-tasking/mega-tasking experience for users.
- PCIe Gen 5.0 support, with as many as 16 lanes off the processor.
- Increased memory support to DDR5-5600 and DDR5-5200, while maintaining DDR4 compatibility.
- Up to 2x the L2 cache and increased L3 cache.

Intel ‘s New 13th Gen

Oct 2022

Introducing the Intel 700 Series Chipset with Backward Compatibility

Alongside the 13th Gen Intel Core desktop processors, Intel is launching the new Intel 700 Series chipset with advanced features for increased reliability and performance. Eight additional PCIe Gen 4.0 lanes combined with PCIe Gen 3.0 provide 28 total lanes off the chipset, increased USB 3.2 Gen 2x2 (20Gbps) ports provide improved USB connectivity speed, and DMI Gen 4.0 increases the chipset-to-CPU throughput for fast access to peripheral devices and networking. Additionally, Intel is bringing forward and backward compatibility. Take advantage of 13th Gen Intel Core processor performance improvements with existing Intel 600 chipset-based motherboards.

Availability

13th Gen Intel Core desktop “K” processors and the Intel Z790 chipset will be available starting Oct. 20, 2022, including boxed processors, motherboards and desktop system sales.

Intel ‘s New 13th Gen

Oct 2022

13th Gen Intel® Core™ Unlock

Processor Number	Processor Cores (P+E)	Processor Threads	Intel® Smart Cache (L3)	Total L2 Cache	P-core Max Turbo Frequency (GHz)	E-core Max Turbo Frequency (GHz)	P-core Base Frequency (GHz)	E-core Base Frequency (GHz)
i9-13900K	24 (8+16)	32	36MB	32MB	Up to 5.8	Up to 4.3	3.0	2.2
i9-13900KF	24 (8+16)	32	36MB	32MB	Up to 5.8	Up to 4.3	3.0	2.2
i7-13700K	16 (8+8)	24	30MB	24MB	Up to 5.4	Up to 4.2	3.4	2.5
i7-13700KF	16 (8+8)	24	30MB	24MB	Up to 5.4	Up to 4.2	3.4	2.5
i5-13600K	14 (6+8)	20	24MB	20MB	Up to 5.1	Up to 3.9	3.5	2.6
i5-13600KF	14 (6+8)	20	24MB	20MB	Up to 5.1	Up to 3.9	3.5	2.6

Intel 's New 13th Gen

Oct 2022

Desktop Processors

Processor Graphics	Total CPU PCIe Lanes	Max Memory Speed (MT/S)	Memory Capacity	Processor Base Power (W)	Max Turbo Power (W)	RCP (USD)
Intel® UHD Graphics 770	20	DDR5 5600 DDR4 3200	128GB	125	253	$589
n/a	20	DDR5 5600 DDR4 3200	128GB	125	253	$564
Intel® UHD Graphics 770	20	DDR5 5600 DDR4 3200	128GB	125	253	$409
n/a	20	DDR5 5600 DDR4 3200	128GB	125	253	$384
Intel® UHD Graphics 770	20	DDR5 5600 DDR4 3200	128GB	125	181	$319
n/a	20	DDR5 5600 DDR4 3200	128GB	125	181	$294

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP122

Intel Xeon

>60 Cores

Xeon Phi 7200 Series	sSpec Number	Cores (Threads)	Clock (MHz)		L2 Cache	MCDRAM Memory		DDR4 Memory		Peak DP Compute	TDP (W)	Soc-ket	Release Date	Part Num
			Base	Turbo		Quantity	BW	Quantity	BW					
Xeon Phi 7210[81]	SR2ME (B0)	64 (256)	1300	1500	32 MB	16 GB	400+ GB/s	384 GB	102.4 Gbit/s	2662 GFLOPS	215	SVLCLGA3647	20 June, 2016	HJ8066702859300
	SR2X4 (B0)													
Xeon Phi 7210F[82]	SR2X5 (B0)										230			HJ8066702975000
Xeon Phi 7230[83]	SR2MF (B0)										215			HJ8066702859400
	SR2X3 (B0)													
Xeon Phi 7230F[84]	SR2X2 (B0)										230			HJ8066702269002
Xeon Phi 7250[85]	SR2MD (B0)	68 (272)	1400	1600	34 MB					3046 GFLOPS[86]	215			HJ8066702859200
	SR2X1 (B0)													
Xeon Phi 7250F[87]	SR2X0 (B0)										230			HJ8066702268900
Xeon Phi 7290[88]	SR2WY (B0)	72 (288)	1500	1700	36 MB					3456 GFLOPS	245			HJ8066702974700
Xeon Phi 7290F[89]	SR2WZ (B0)										260			HJ8066702975200

Intel Xeon Die

18 CPU Cores

Intel Haswell Xeon e7

(Note: display size on this page is limited to 1024px wide. Use URL below to retrieve raw file.)
URL:http://spectrum.ieee.org/image/MjgyMjYxNg.jpeg

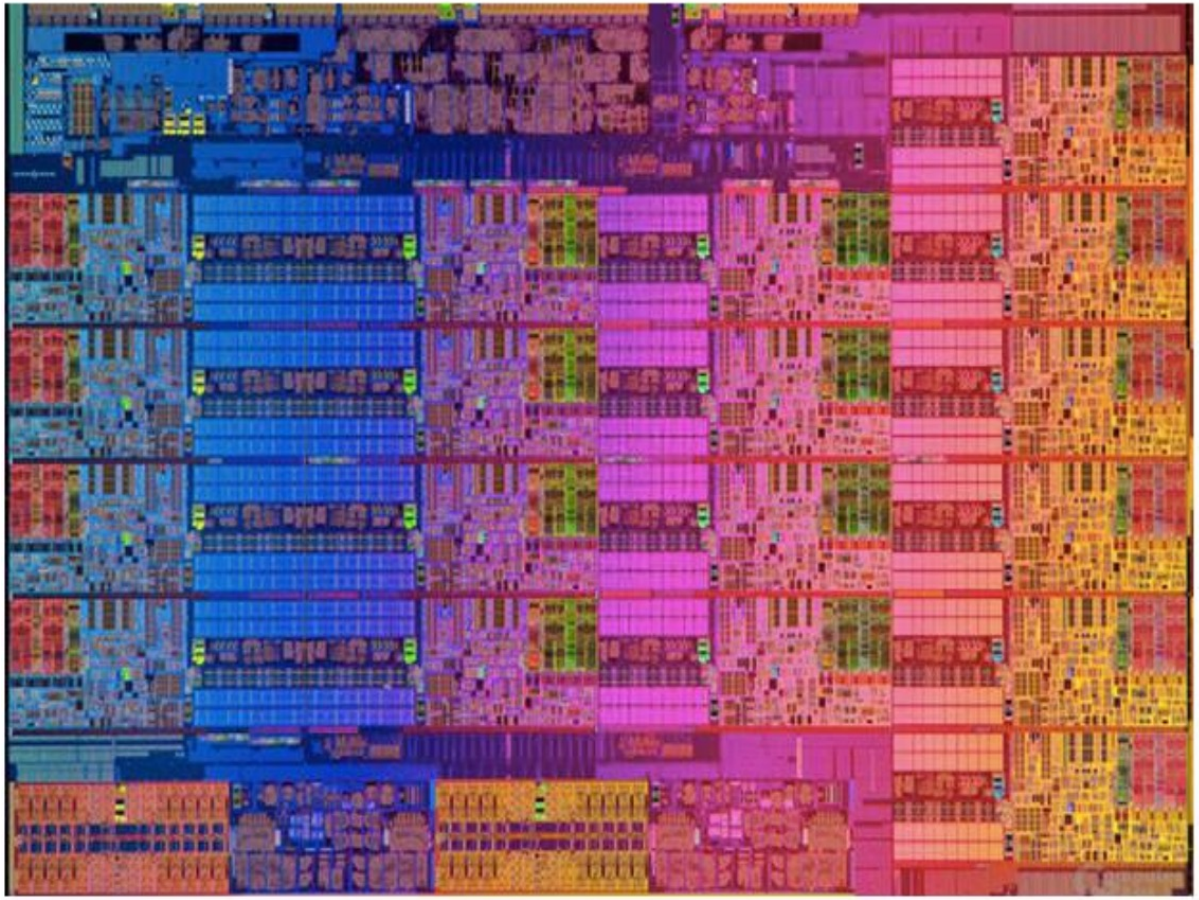

Asset id: 3808
Type: Image Mime:
Caption: This Intel Haswell-EX Xeon E7 V3 processor has 18 cores trying to work together with

Intel AVX, Xeon, FPGA

Intel® Xeon® Scalable Processors

Intel® FPGAs

Intel Xeon E5

Xeon 6

June 2024

Intel® Xeon®6 Processor

The Best Processors to Meet Your Diverse Performance & Efficiency Requirements

P-core

Performance

Industry-leading Performance-cores excel at the widest range of workloads

Best for compute-intensive and AI workloads

E-core

Efficiency

A new class of Efficient-cores deliver unmatched core density & performance per watt

Distinct advantages for cloud-native & hyperscale workloads

Xeon 6

June 2024

Product Name	Total Cores	Max Turbo Frequency	Processor Base Frequency	Cache
Intel® Xeon® 6740E Processor (96M Cache, 2.40 GHz)	96	3.2 GHz	2.4 GHz	96 MB
Intel® Xeon® 6756E Processor (96M Cache, 1.80 GHz)	128	2.6 GHz	1.8 GHz	96 MB
Intel® Xeon® 6780E Processor (108M Cache, 2.20 GHz)	144	3 GHz	2.2 GHz	108 MB
Intel® Xeon® 6710E Processor (96M Cache, 2.40 GHz)	64	3.2 GHz	2.4 GHz	96 MB
Intel® Xeon® 6766E Processor (108M Cache, 1.90 GHz)	144	2.7 GHz	1.9 GHz	108 MB
Intel® Xeon® 6731E Processor (96M Cache, 2.20 GHz)	96	3.1 GHz	2.2 GHz	96 MB
Intel® Xeon® 6746E Processor (96M Cache, 2.00 GHz)	112	2.7 GHz	2 GHz	96 MB

Intel vs AMD

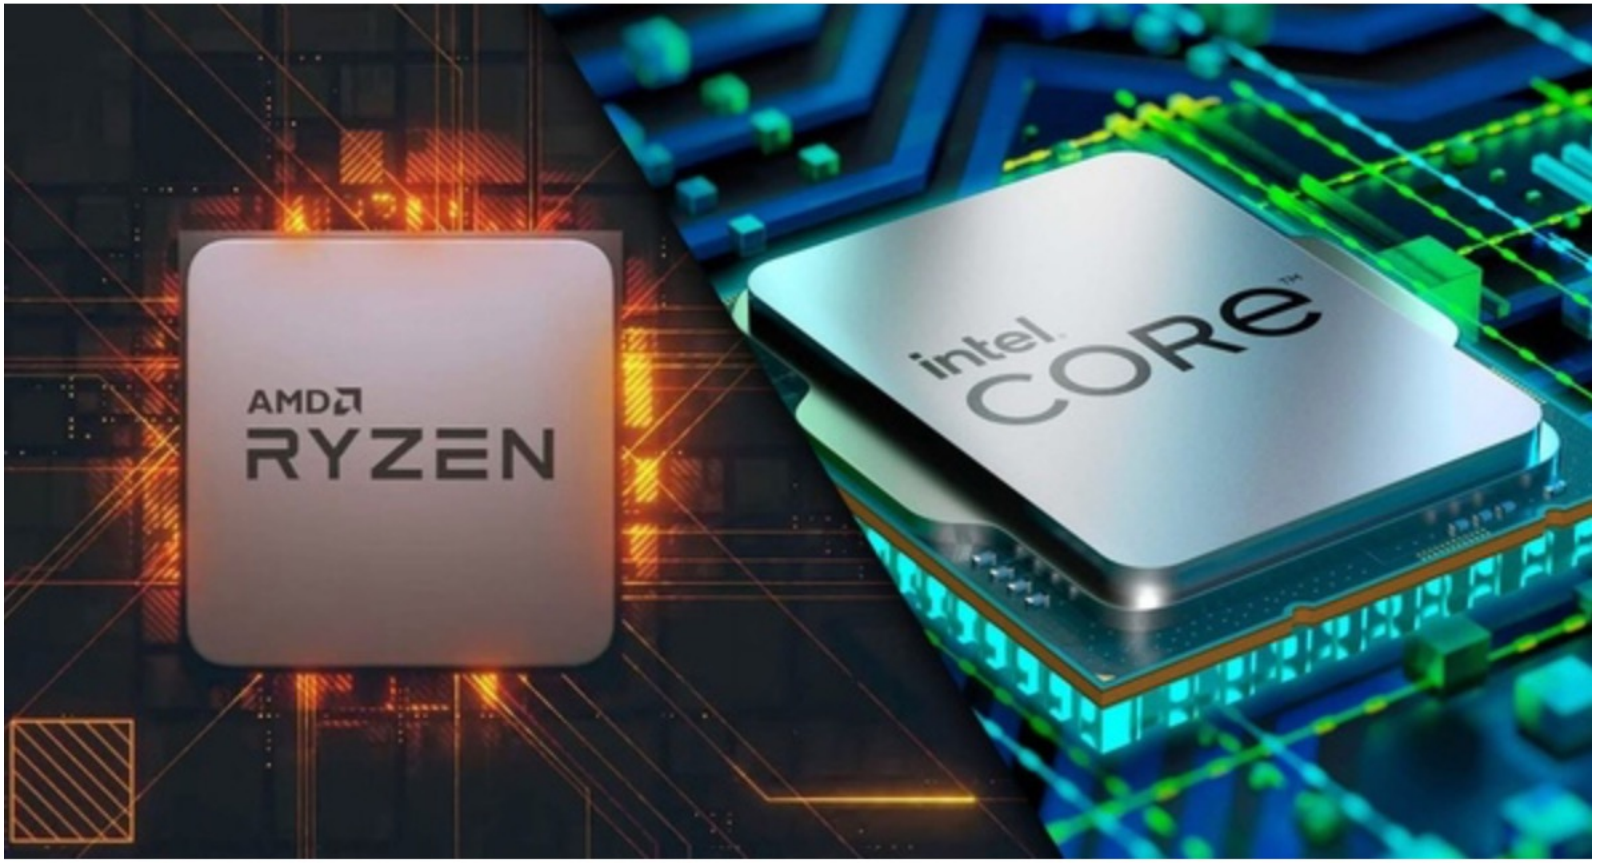

Intel vs AMD

Quora

What are the differences between Intel processors and AMD's Ryzen CPUs? Can they be compared, or are they designed for different purposes?

Jeff Drobman
Worked at Advanced Micro Devices (company) · Just now ·

yes they are comparable. AMD and Intel both have high performance CPU cores that run the x86–64 ISA. both also have models with and without integrated graphics (GPU). both are now using multi-chip modules they call chiplets and tiles, respectively. and their top chips are both manufactured at TSMC in the same 5nm and 3nm processes. the main difference is how many cores of each type per chip, and how many dice per module (Epyc or Xeon).

Intel x86-64 from AMD

X86-64 a 64-bit version of the x86 instruction set, first released in 1999. It introduced two new modes of operation, 64-bit mode and compatibility mode, along with a new 4-level paging mode.

AMD vs. Intel

Which chipmaker is growing faster?

Based on these facts, it's easy to see why AMD grew at a faster clip than Intel over the past three years.

Revenue Growth (YOY)	2018	2019	2020
Intel (NASDAQ:INTC)	13%	2%	8%
AMD (NASDAQ:AMD)	23%	4%	45%

DATA SOURCE: ANNUAL REPORTS. YOY = YEAR OVER YEAR.

AMD's growth decelerated in 2019, mainly due to slower sales of gaming consoles. However, its growth accelerated again in 2020 as it launched its new Ryzen CPUs and Radeon GPUs, while Intel's ongoing chip shortage continued to generate tailwinds for both its PC and data center businesses.

Intel's overall revenue growth initially seems stable, but its data center chip sales declined year over year in the second half of 2020 -- which partly offset its higher sales of PC CPUs during the pandemic.

Intel recently outsourced the production of some of its chips to TSMC to resolve its ongoing shortages, but it doesn't expect to launch its next-gen 7nm chips until 2023. AMD launched its first 7nm CPUs back in 2019, and it will likely launch its first 5nm chips in 2022.

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE

COMP122

AMD vs. Intel CPU's

The Motley Fool

Latest Stock Picks | Investing Basics | Stock Market | Retirement | Personal Financ

The key differences between Intel and AMD

Intel generated 56% of its revenue from PC-centric chips last year. It generated 36% of its revenue from data center chips, while the remaining 8% came from other types of chips -- including IoT (Internet of Things) chips, programmable chips, computer vision chips, and memory chips.

Intel agreed to sell most of its memory chip business to **SK Hynix** last October, and it's developing discrete GPUs to complement its CPUs and compete against NVIDIA and AMD. However, Intel's core business still mainly relies on its sales of x86 CPUs for PCs and data centers.

Intel manufactures its own chips with its internal foundry. But Intel's foundry fell behind **Taiwan Semiconductor Manufacturing** (NYSE:TSM), the world's largest third-party foundry, in the "process race" to create smaller chips in recent years. Those mistakes clogged up Intel's foundries and caused shortages of its latest CPUs.

AMD, which outsources the production of its CPUs and GPUs to TSMC instead of manufacturing them internally, didn't struggle with any shortages. As a result, AMD's share of the global x86 CPU market more than doubled from 18.1% to 39.4% between the first quarters of 2017 and 2021, according to PassMark Software, as Intel's market share plunged from 81.9% to 60.5%.

Last year, AMD generated nearly two-thirds of its revenue from its computing and graphics segment, which sells its Ryzen CPUs and Radeon GPUs. The rest of its revenue came from its EESC (enterprise, embedded, and semi-custom) business, which mainly sells custom chips for gaming consoles (including **Sony**'s PS5 and **Microsoft**'s Xbox Series consoles) and Epyc CPUs for data centers.

Computer Architecture

Future
Intel CPU's

Intel New Chips

Why doesn't Intel have as strong of integrated graphics on their CPUs, such as their Intel UHD 630 graphics, compared to AMD's Vega 11 integrated graphics?

Brett Bergan, Building PC's for 25 years

Answered 48m ago

Unfortunately for AMD fans, Vega 11 was a great product that found its way to ONE processor (actually two if you consider the 2400G and 3400G two different CPUs)

But that detail aside, Intel has been recycling the same 14nm "Skylake" HD 530/630 GPU for five years already. I have a sneaky suspicion that 10th gen Comet Lake CPU models consist of a lot of recycled Coffee Lake silicon that didn't get sold in 2018. The i3–10100 hyperthreaded quad is essentially a i7–7700 that has a locked multiplier set at 4.3GHz

Same CPU. Same GPU. Just three generations later.

10nm Ice Lake with its somewhat improve ... (more)

Intel New Chips

Yowan Rajcoomar, Computer Technician (2008-present)

Answered 1h ago

Intel desktop chips are set to receive a GPU upgrade with the 11th generation. The 24EU GPU will be replaced with an Xe GPU with 32EUs which will bring a 50% improvement in performance.

This still won't be enough to run triple-A games but you need to keep in mind that Intel considers integrated graphics on the desktop as a 'value added' solution while dedicating more die area to CPU cores. In contrast, AMD dedicates much area to the GPU (around half the entire die).

Intel was stuck with 24EUs since 2015 because they've been using the same Skylake architecture in 7-10th gen desktop chips. The HD530 was simply renamed as HD630 with the 7th gen and again renamed to UHD630 with the 8th gen. That exact same GPU had been in use for 5 year without any changes whatsoever.

Intel CPU's

Power Mgt

Sandy Bridge power mgmt ID card

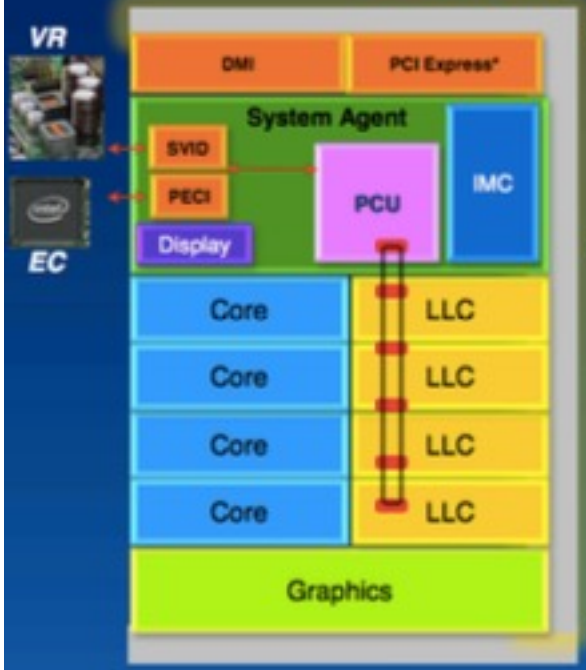

- **Sandy Bridge is:**
  - 1-4 CPU cores + PG
  - Integrated System Agent (SA)
  - Sliced LLC shared by all cores/PG
  - Ring interconnect + power management link
- **Package Control Unit (PCU) :**
  - On chip logic and embedded controller running power management firmware
  - Communicates internally with cores, ring and SA
  - Monitors physical conditions
    - Voltage, temperature, power consumption
  - Controls power states
    - CPU and PG voltage and frequency
    - Controls voltage regulators DDR and system
- **External power management interface**
  - Accepts external inputs
    - System power management requests and limits
    - Power and temperature reading
  - MSR, MMIO and PECI system bus

intel Leap ahead

Sandy Bridge - Hot Chips 2011

4

Intel CPU's

Power Mgt

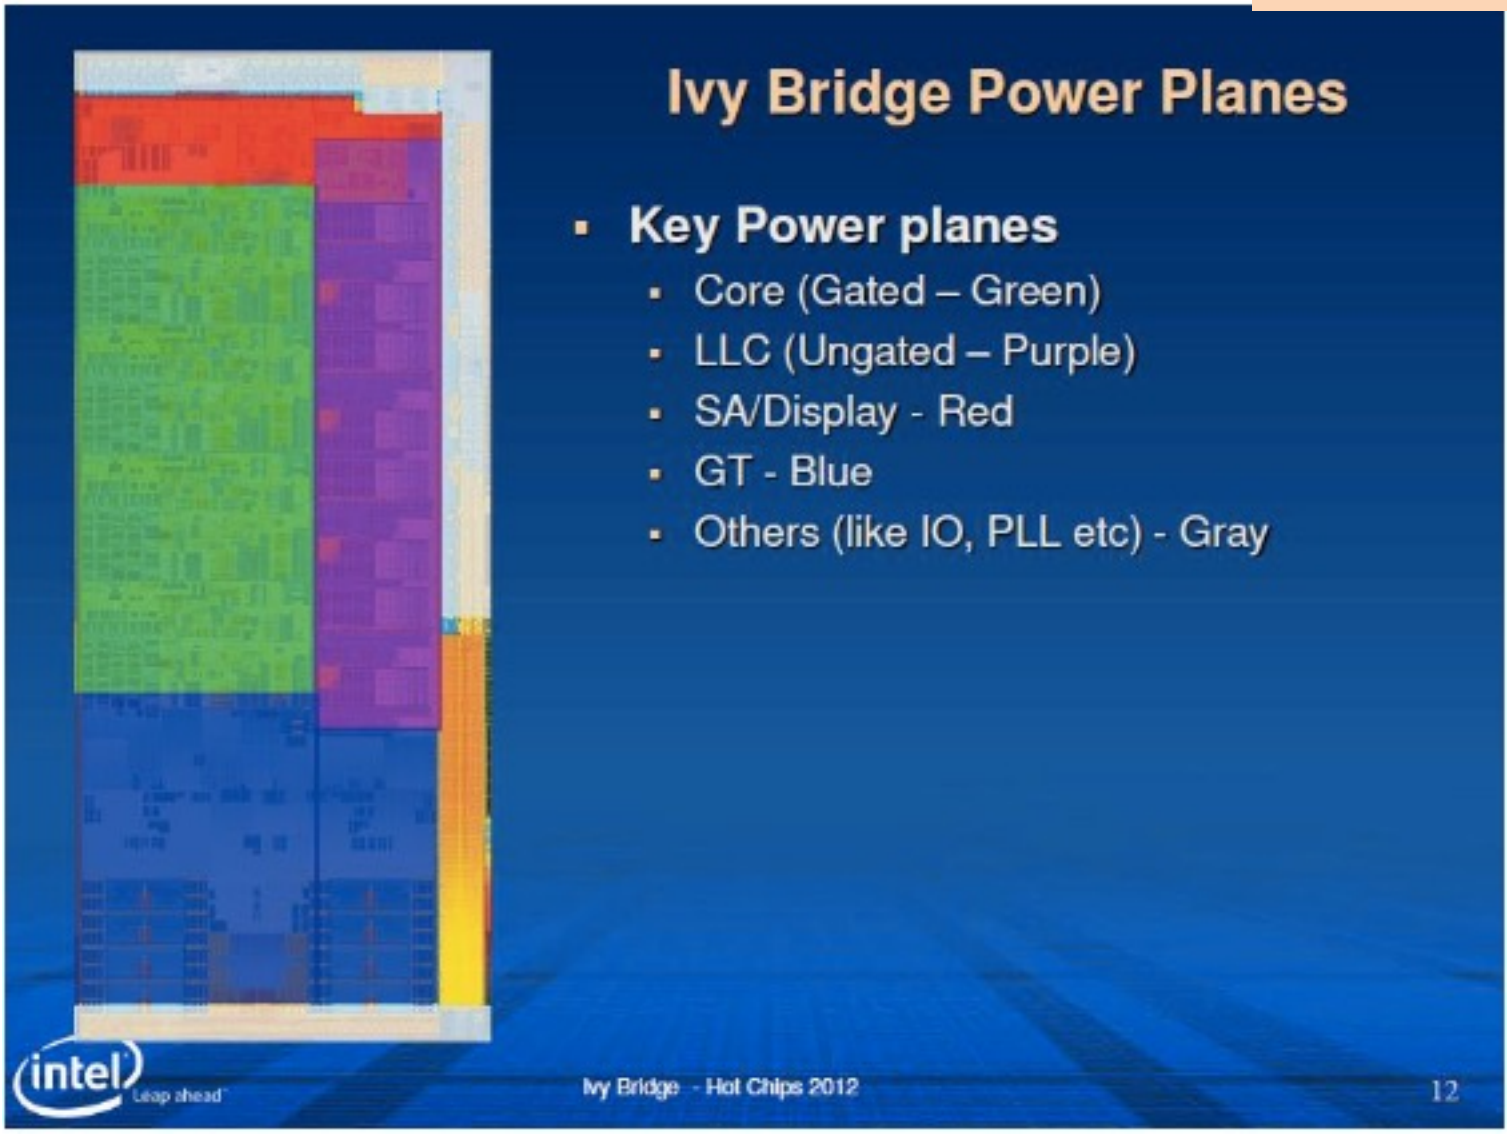

Intel CPU's

Power Mgt

Voltage and frequency domains

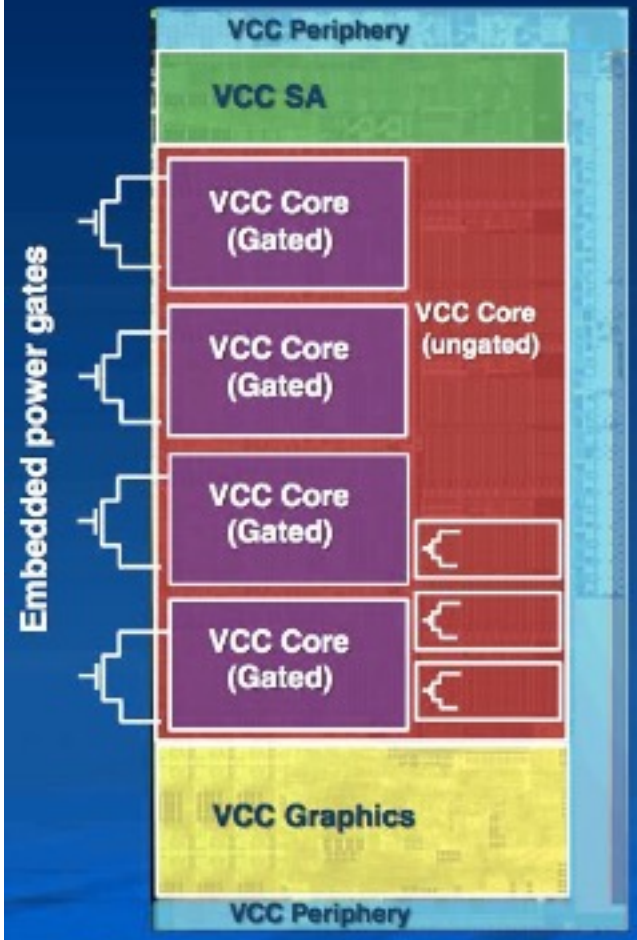

- **Two Independent Variable Power Planes:**
  - CPU cores, ring and LLC
    - Embedded power gates - Each core can be turned off individually
    - Cache power gating - Turn off portions or all cache at deeper sleep states
  - Graphics processor
    - Can be varied or turned off when not active
- **Shared frequency for all IA32 cores and ring**
- **Independent frequency for PG**
- **Fixed** Programmable **power plane for System Agent**
  - Optimize SA power consumption
  - System On Chip functionality and PCU logic
  - Periphery: DDR, PCIe, Display

intel Leap ahead

Sandy Bridge - Hot Chips 2011

5

Intel CPU's

Power Mgt

LLC - Dynamic Cache Shrink Feature

Power Down LLC

- **LLC organized in 16 ways.**
- **When PCU detects low activity workload**
  - Flushes 14 ways of the cache and puts ways to sleep
  - Shrinks active ways from 16 to 2 to improve VccMin
- **When PCU detects high activity**
  - Expands active ways back to 16 to improve cache hit rate.

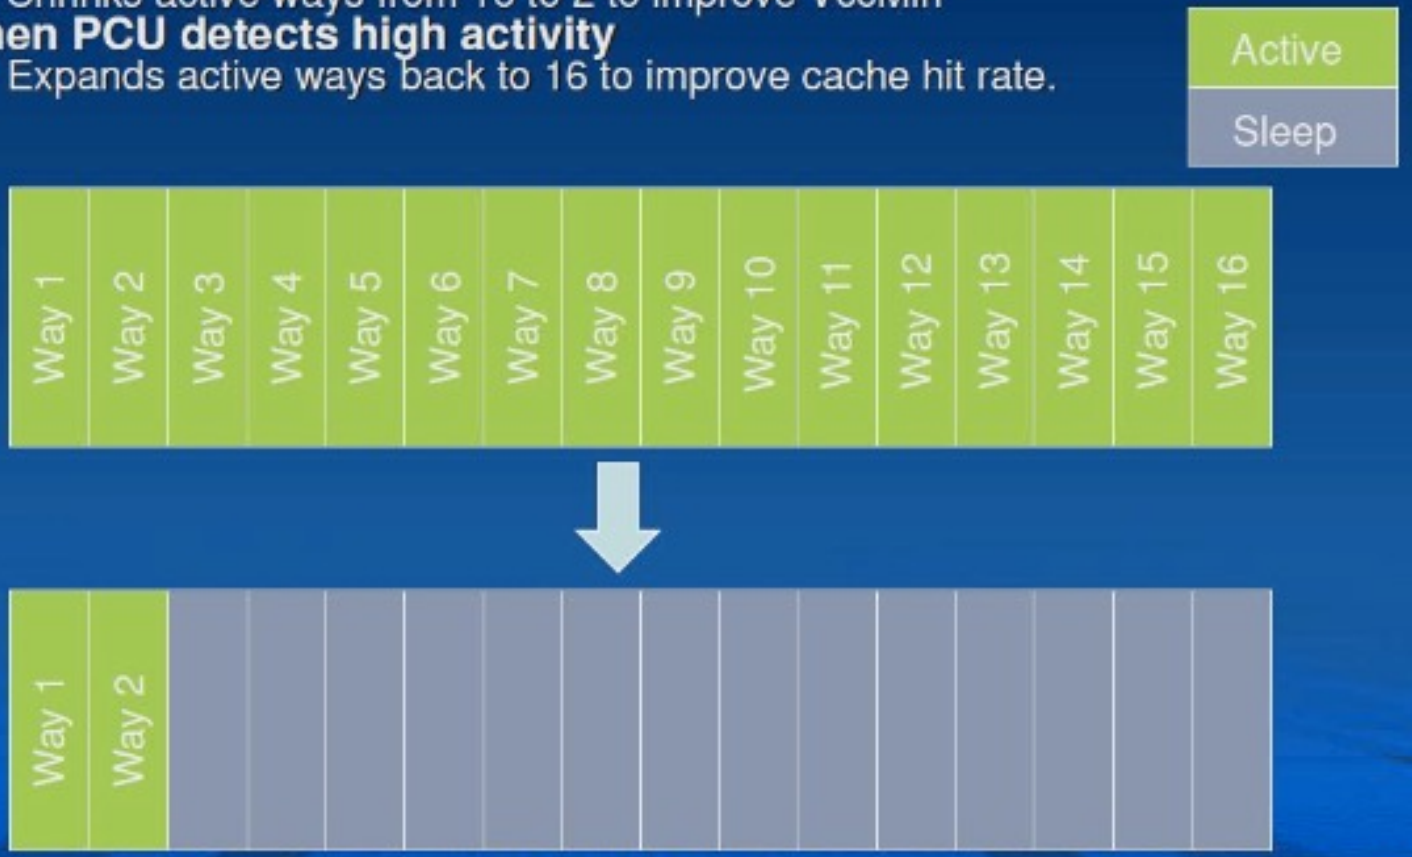

Intel CPU's

14.4 HARDWARE-CONTROLLED PERFORMANCE STATES (HWP)

Intel processors may contain support for Hardware-Controlled Performance States (HWP), which autonomously selects performance states while utilizing OS supplied performance guidance hints. The Enhanced Intel Speed-Step® Technology provides a means for the OS to control and monitor discrete frequency-based operating points via the IA32_PERF_CTL and IA32_PERF_STATUS MSRs.

In contrast, HWP is an implementation of the ACPI-defined Collaborative Processor Performance Control (CPPC), which specifies that the platform enumerates a continuous, abstract unit-less, performance value scale that is not tied to a specific performance state / frequency by definition. While the enumerated scale is roughly linear in terms of a delivered integer workload performance result, the OS is required to characterize the performance value range to comprehend the delivered performance for an applied workload.

When HWP is enabled, the processor autonomously selects performance states as deemed appropriate for the applied workload and with consideration of constraining hints that are programmed by the OS. These OS-provided hints include minimum and maximum performance limits, preference towards energy efficiency or performance, and the specification of a relevant workload history observation time window. The means for the OS to override HWP's autonomous selection of performance state with a specific desired performance target is also provided, however, the effective frequency delivered is subject to the result of energy efficiency and performance optimizations.

Intel CPU's

taking-overclocking-to-the-next-level.html at 24:

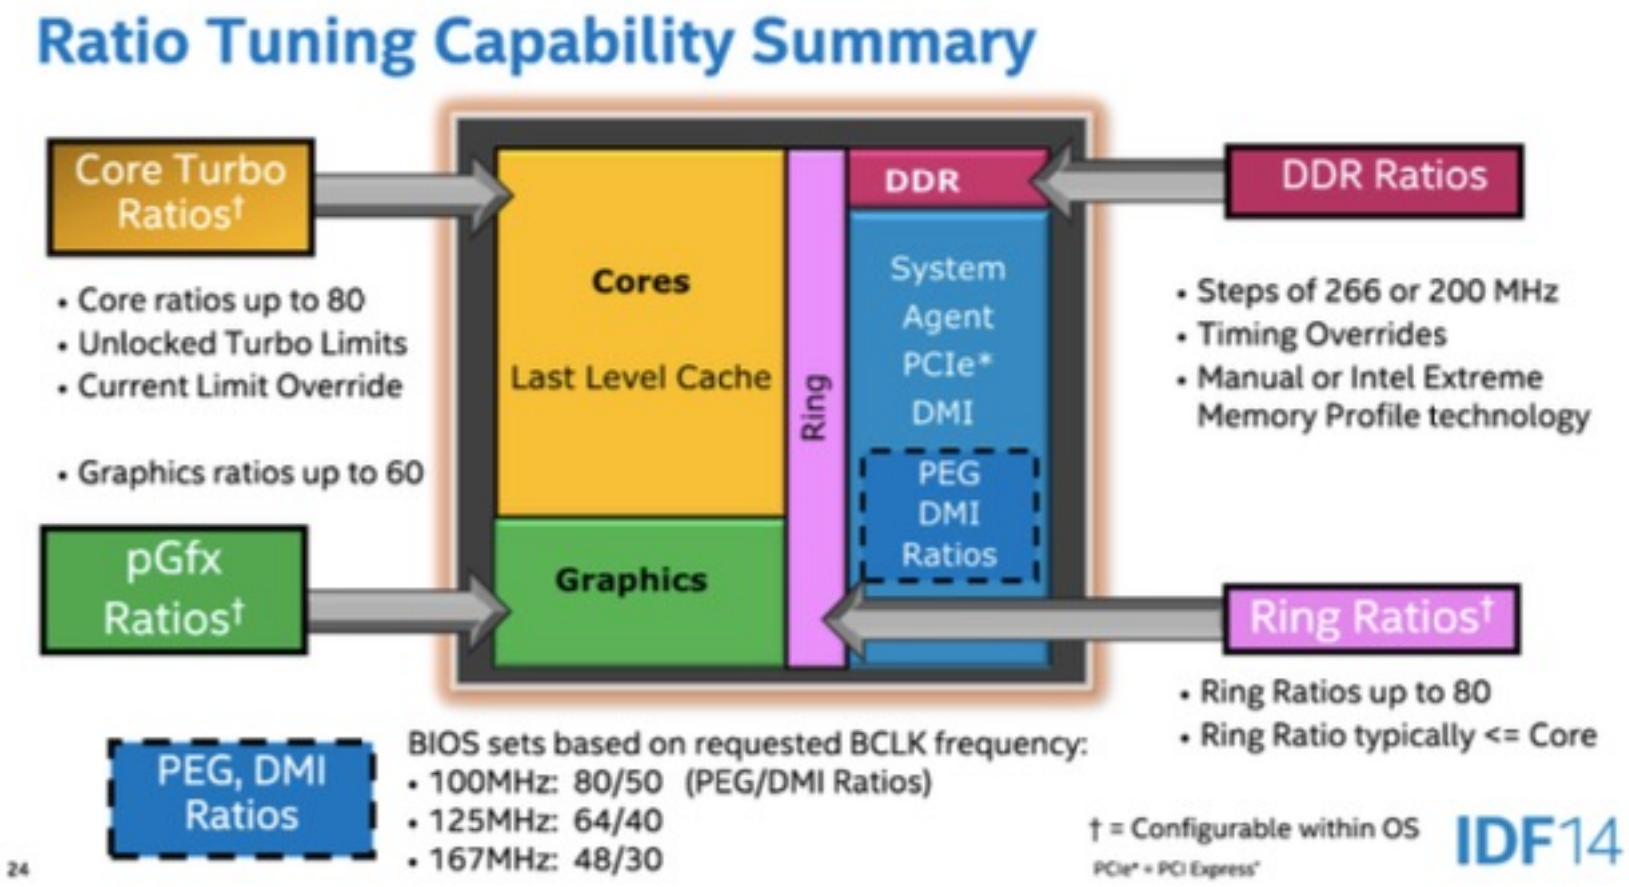

Intel Future CPUs

Xeon

Intel

Source: Intel

Next-Gen Intel Xeon Scalable Processors to Deliver Breakthrough Platform Performance

This week, Intel announced its next-generation Intel Xeon Scalable processor platform. Codenamed Cooper Lake, the latest addition to the processor family boasts up to 56 processor cores per socket, higher memory bandwidth, and built-in AI inference and training acceleration. Learn more about the high-core-count Cooper Lake processor platform, which will be available in the first half of 2020.

<= 56 cores in 1H 2020

Intel Future CPUs

Source: TechRadar

Intel Ice Lake Has Landed, and It's Smarter, Faster and More Efficient Than Ever

Built on 10nm technology, Intel's Ice Lake platform is the first step toward Project Athena—an initiative that aims to leverage next-generation tech for advanced mobile computing experiences. Ice Lake's deep-learning capabilities will deliver optimized performance while the inclusion of Gen11 graphics will make for more sophisticated gaming experiences. Here's what else you can expect from Ice Lake.

Intel Future CPUs

Intel

Source: Intel IT Peer Network

Exascale Computing Will Redefine Content Creation

In line with Intel's pursuit of exponential increases in computing capabilities, Intel's Architecture, Software, and Graphics group has been challenged to deliver a 1000x workflow improvement for content creators. Enter exascale computing. With the ability to process a quintillion (that's a billion billion, for those who can't count that high) calculations per second, exascale computing will enable content creators to push the limits for high-quality animation.

Exascale = 10^18 IPS

Intel Future CPUs

Hot Chips 2019 | Intel

AT HOT CHIPS, INTEL PUSHES 'AI EVERYWHERE'

What's New: At Hot Chips 2019, Intel revealed new details of upcoming high-performance artificial intelligence (AI) accelerators: Intel® Nervana™ neural network processors, with the NNP-T for training and the NNP-I for inference. Intel engineers also presented technical details on hybrid chip packaging technology, Intel® Optane™ DC persistent memory and chiplet technology for optical I/O.

Source: Intel

Intel Future CPUs

Hot Chips 2019 | Intel

What Intel Presented at Hot Chips 2019:

Intel Nervana NNP-T: Built from the ground up to train deep learning models at scale: Intel Nervana NNP-T (Neural Network Processor) pushes the boundaries of deep learning training. It is built to prioritize two key real-world considerations: training a network as fast as possible and doing it within a given power budget. This deep learning training processor is built with flexibility in mind, striking a balance among computing, communication and memory. While Intel® Xeon® Scalable processors bring AI-specific instructions and provide a great foundation for AI, the NNP-T is architected from scratch, building in features and requirements needed to solve for large models, without the overhead needed to support legacy technology. To account for future deep learning needs, the Intel Nervana NNP-T is built with flexibility and programmability so it can be tailored to accelerate a wide variety of workloads – both existing ones today and new ones that will emerge. View the presentation for additional technical detail into Intel Nervana NNP-T's (code-named Spring Crest) capabilities and architecture.

Intel Nervana NNP-I: High-performing deep learning inference for major data center workloads: Intel Nervana NNP-I is purpose-built specifically for inference and is designed to accelerate deep learning deployment at scale, introducing specialized leading-edge deep learning acceleration while leveraging Intel's 10nm process technology with Ice Lake cores to offer industry-leading performance per watt across all major datacenter workloads. Additionally, the Intel Nervana NNP-I offers a high degree of programmability without compromising performance or power efficiency. As AI becomes pervasive across every workload, having a dedicated inference accelerator that is easy to program, has short latencies, has fast code porting and includes support for all major deep learning frameworks allows companies to harness the full potential of their data as actionable insights. View the presentation for additional technical detail into Intel Nervana NNP-I's (code-named Spring Hill) design and architecture.

Intel Future CPUs

Hot Chips 2019 | Intel

TeraPHY: An in-package optical I/O chiplet for high-bandwidth, low-power communication: Intel and Ayar Labs demonstrated the industry's first integration of monolithic in-package optics (MIPO) with a high-performance system-on-chip (SOC). The Ayar Labs TeraPHY* optical I/O chiplet is co-packaged with the Intel Stratix 10 FPGA using Intel Embedded Multi-die Interconnect Bridge (EMIB) technology, offering high-bandwidth, low-power data communication from the chip package with determinant latency for distances up to 2 km. This collaboration will enable new approaches to architecting computing systems for the next phase of Moore's Law by removing the traditional performance, power and cost bottlenecks in moving data. View the presentation for additional technical detail and design decisions on creating processors with optical I/O.

Intel Optane DC persistent memory: Architecture and performance: Intel Optane DC persistent memory, now shipping in volume, is the first product in the memory/storage hierarchy's entirely new tier called persistent memory. Based on Intel® 3D XPoint™ technology and in a memory module form factor, it can deliver large capacity at near-memory speeds, latency in nanoseconds, while also natively delivering the persistence of storage. Details of the two operational modes (memory mode and app direct mode) as well performance examples show how this new tier can support a complete re-architecting of the data supply subsystem to enable faster and new workloads. View the presentation for additional architectural details, memory controller design, power fail implementation and performance results for Intel Optane DC persistent memory.

Intel Moore's Tick Tock

staggered it allows a once a year cadence of new product introduction. One year the new chip is a process shrink of an existing architecture. The next year the new chip uses the same process node but introduces a new architecture.

The beauty of this chip development model is that it is a business innovation that takes advantage of a classic pipeline that hides latency to achieve throughput.

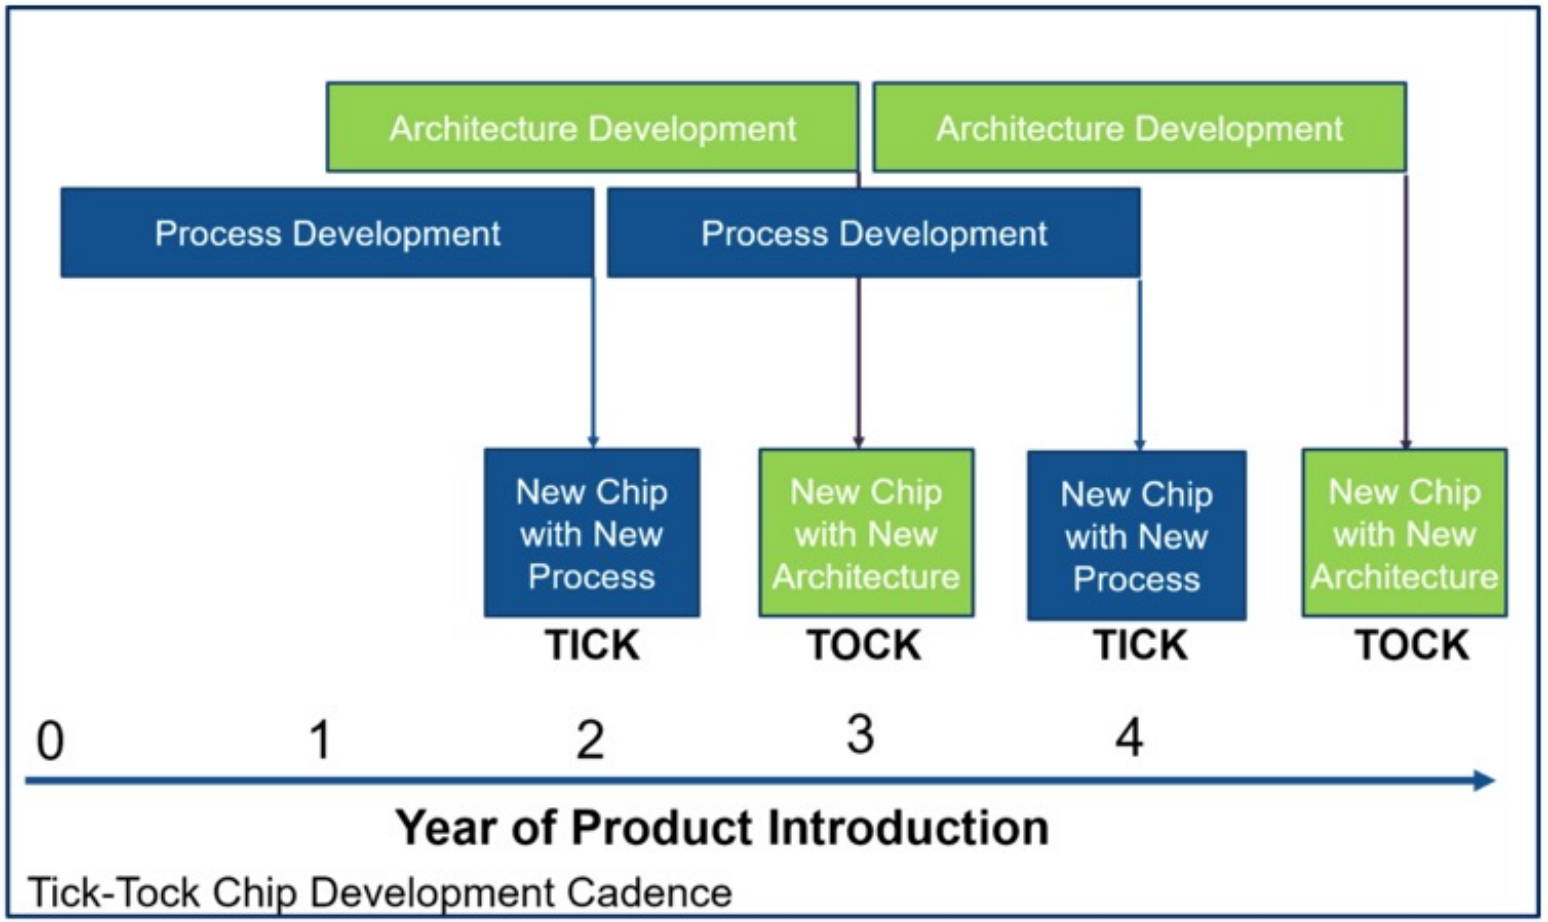

Tick-Tock Chip Development Cadence

Intel vs. Apple

The 4-core 11th gen Tiger Lake chips such as the Core i7–1165G7 has similar single threaded performance but higher multi-core performance despite having two less cores than the A13:

A13

Geekbench 5 Score

1637	4903
Single-Core Score	Multi-Core Score

Here is the A14 GB5 score:

A14

iPhone 12 Pro

Geekbench 5 Score

1599	4068
Single-Core Score	Multi-Core Score

Geekbench 5.2.5 for iOS AArch64

Don't forget that the mobile SoCs have limited sustained performance due to their low thermal envelope of under 5 watts while laptop chips from Intel have a base TDP ranging from 7 W to 28 W and can boost to 44W. Apple's SoCs is more energy efficient but Intel remains unchallenged when it comes to raw performance, especially when you consider that their chips support complex extensions such as VNNI and AVX-512.

MacBook Dual-Core ARM

2012

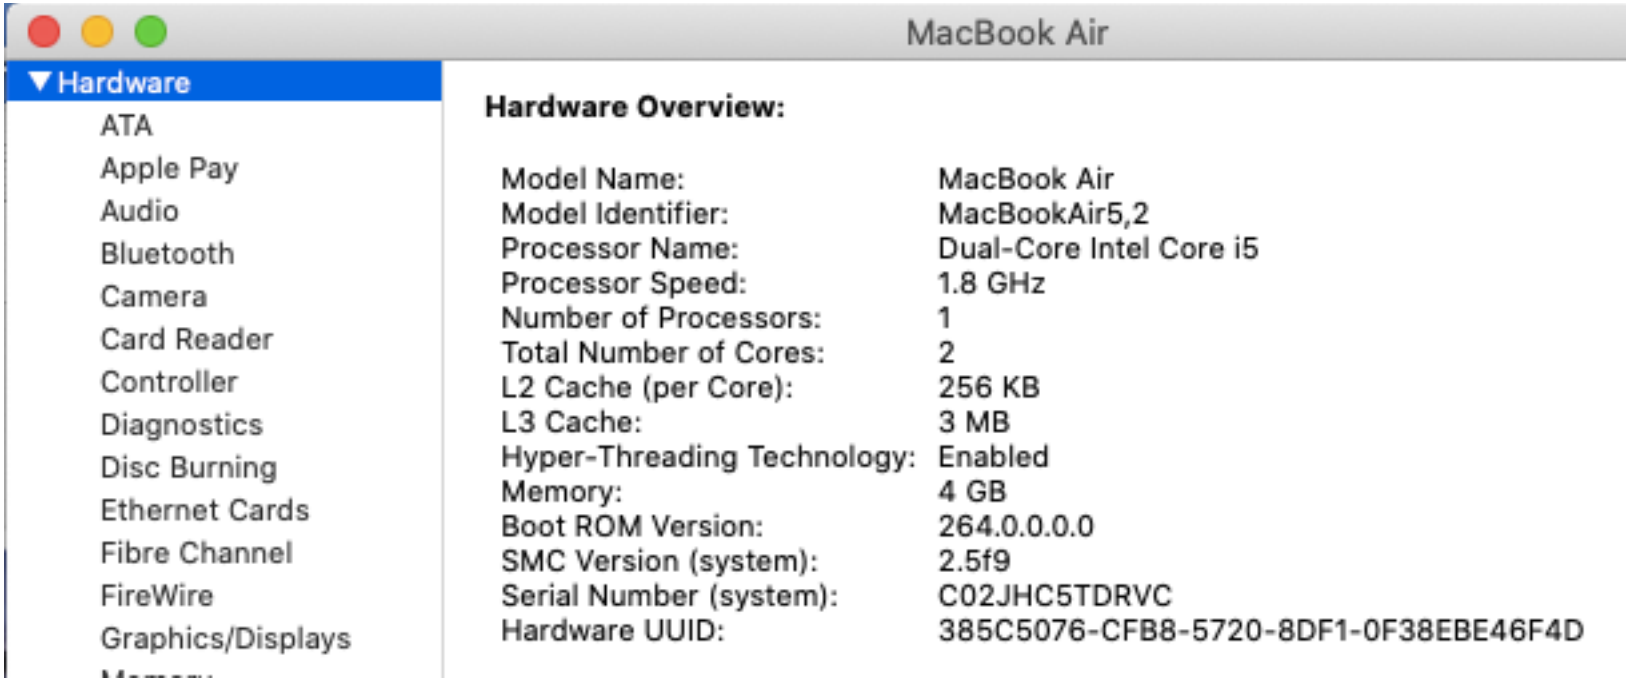

Section

Intel GPU's

See separate slide set "GPU"

Older Intel GPU Cores

Intel chips have had integrated GPU graphics for years. The Sandy/Ivy Bridge graphics were rather slow and you can see how small the GPU was in comparison to the actual CPU cores. This was 2500 series HD graphics.

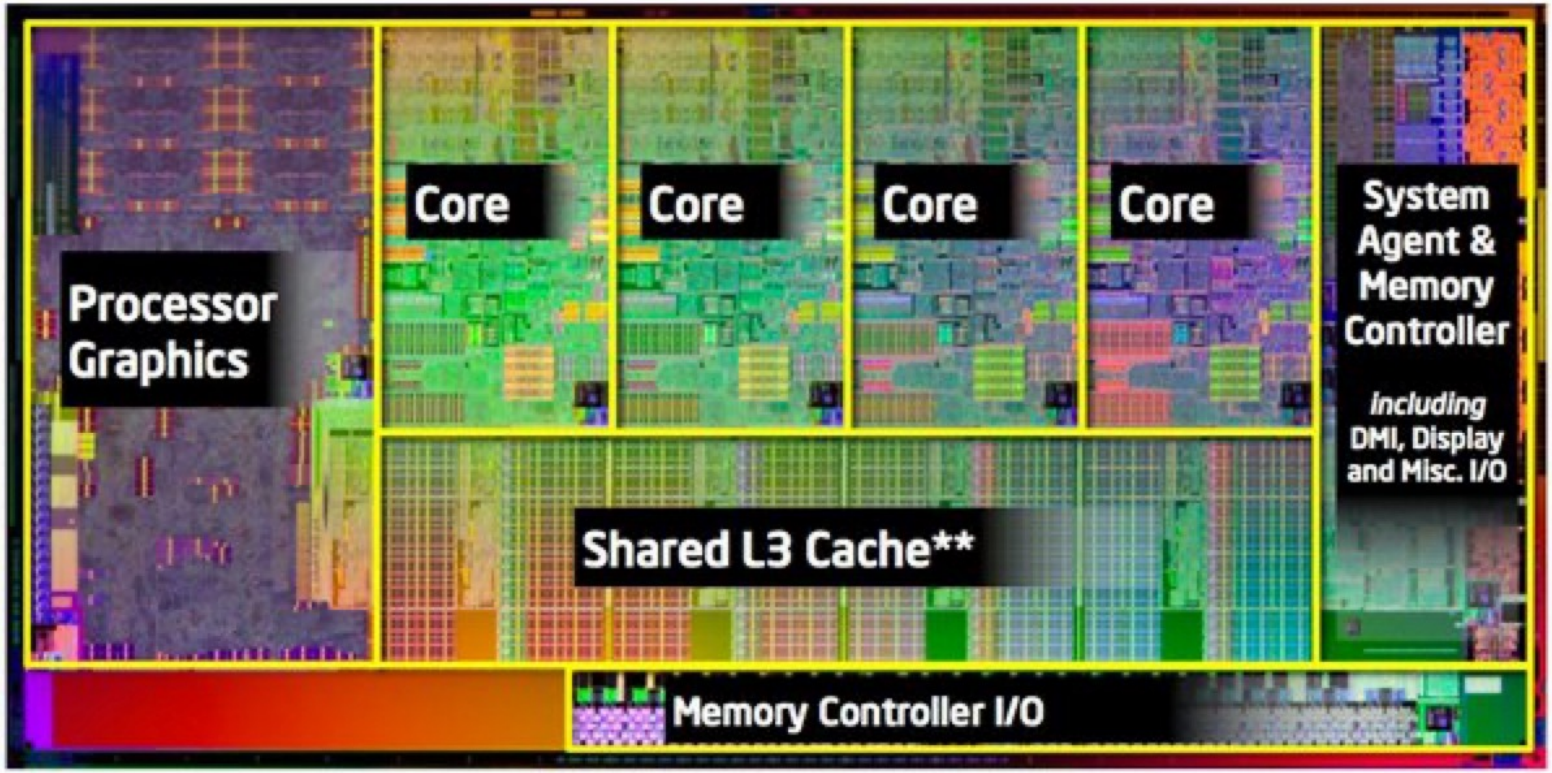

> *Six TMUs and **48 shader cores** in the HD2000/HD2500 GPU*

Newer Intel GPU Cores

Advancing to Broadwell/Haswell the proportion of GPU power is evident by an increase in the number of cores in the GPU compute architecture—shown here, essentially duplicated in mirror image and turned sideways.

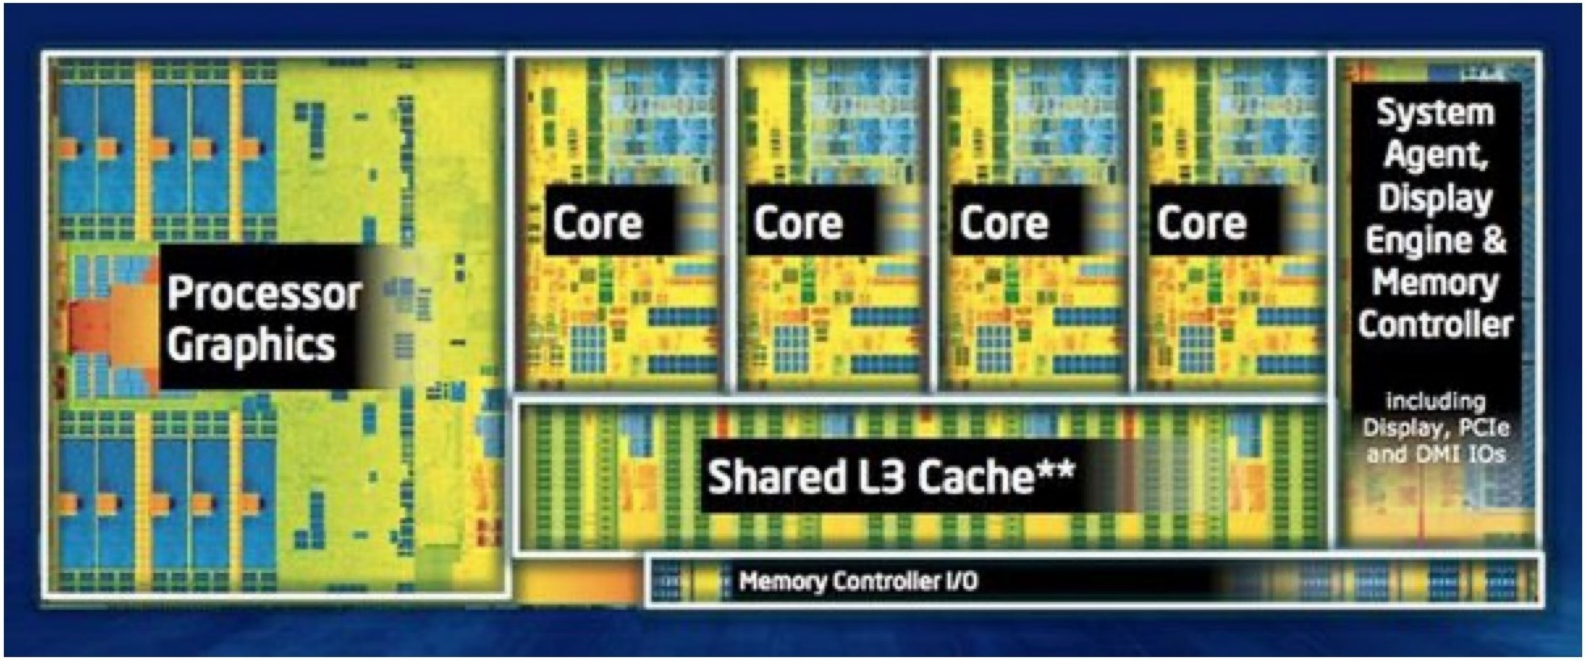

> *Sixteen TMUs and **128 shader cores** of Broadwell/Haswell HD4000 GPU off huge improvements over Ivy/Sandy Bridge.*

Intel i9: AVX CPU + Xe GPU

Quora

Brett Bergan, Building PC's for 25 years

Answered Sat

The 12900K Alder lake appears to be the 11900K with eight efficiency cores wedged into the mix. My understanding is that the so-called “Gracemont” cores are roughly akin to updated Atom-type cores running at a base clock of 3.2GHz. You can see from the die shot that four Gracemont cores take up less space than a single “Golden Cove” P-Core.

Besides the addition of the additional mini-cores each “P-Core” has 1.25MB L2 Cache which is greatly increased over the 512K L2 cache of previous generations. The 3MB cache per P-Core or E-Core cluster adds up to 30MB, which is at least on par withe the 32MB L3 cache of the Ryzen 5 5600G.

Intel i9: AVX CPU + Xe GPU

Intel i9: AVX CPU + Xe GPU

Quora

The 12900KF obviously has the iGPU disabled, which on this large of a chip is only a very small section of silicon, although the 256-core Xe GPU is a formidable resource in its own right. I would assume that the "Media Engine" is retained on a 12900KF but I don't know that for certain.

Because all cores are not created equal, Alder Lake must now have a hierarchical **Thread Director** to optimize tasks for delegation either the efficiency cores or to the performance cores. All Windows background tasks can now be done exclusively in the E-Cores, freeing the P-Cores for the heavy lifting.

Other features include **sixteen PCIe 5.0 lanes** in addition to **four native PCIe 4.0 lanes**. The Z690 chipset also includes an additional **twelve PCIe 4.0 lanes** and **sixteen PCIe 3.0 lanes** for a pretty insane amount of high speed interconnects. In addition to PCIe 5.0 the support for DDR5 is a major redesign that will have many benefits in the future even if its practicality is rather limited right now.

Peak turbo clock speeds are limited to 5.1GHz, which is significantly lower than the 5.3GHz clock speed reached by the i9–10900K.

Section

Intel Strategy

Intel Architectures

Diverse Architectures. Unprecedented Choice.

Computing applications have become increasingly varied, and different types of workloads require different types of computing architectures. Intel is unique positioned to deliver a diverse mix of scalar, vector, matrix, and spatial architectures deployed in CPU, GPU, accelerator, and FPGA sockets. As a result, our customers can use the most appropriate type of compute when and where it's needed.

Scalar

Versatile, high-performance, and power-efficient general-purpose compute with world-class Intel® CPUs.

Vector

Highly parallel processing with Intel® integrated graphics and discrete GPUs.

Matrix

Innovative new CPU instructions and purpose-built Intel® accelerators for AI performance.

Spatial

Reprogrammable Intel® FPGAs for customizable acceleration.

- FPGA (from Atmel)

x64/IA-64 ISA

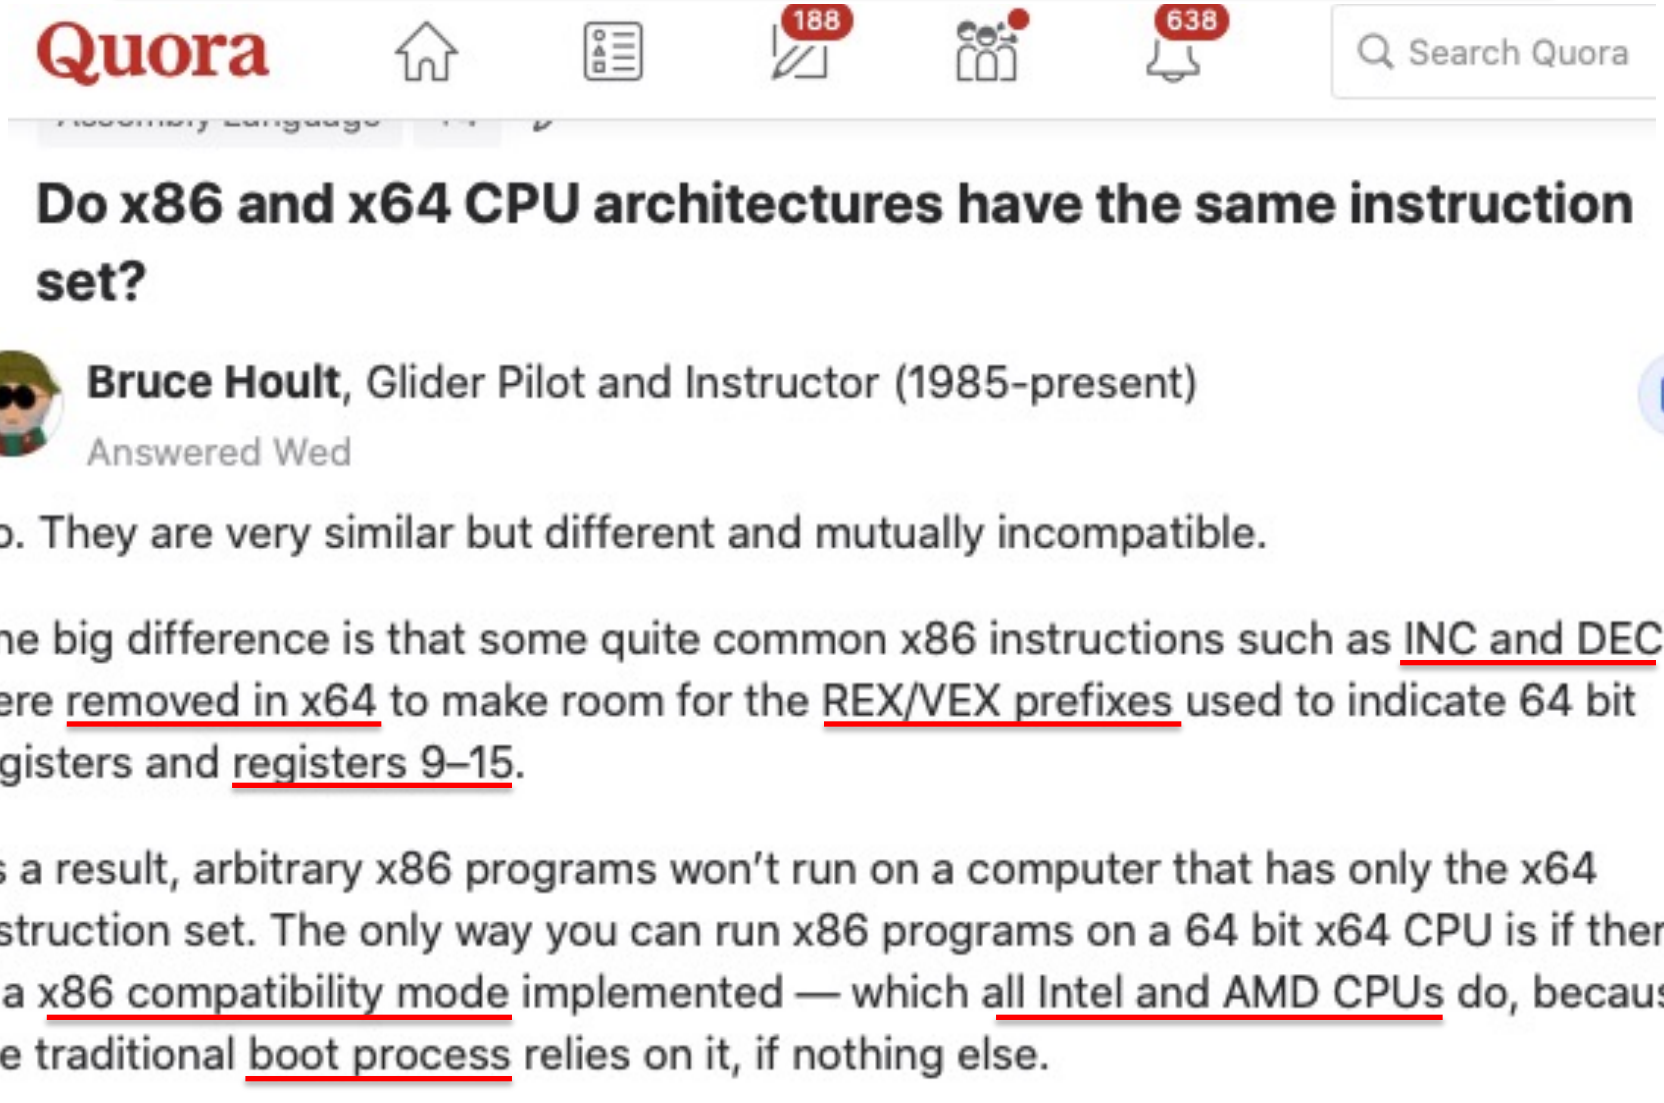

Quora 188 638 Search Quora

Do x86 and x64 CPU architectures have the same instruction set?

Bruce Hoult, Glider Pilot and Instructor (1985-present)
Answered Wed

No. They are very similar but different and mutually incompatible.

One big difference is that some quite common x86 instructions such as INC and DEC were removed in x64 to make room for the REX/VEX prefixes used to indicate 64 bit registers and registers 9–15.

As a result, arbitrary x86 programs won't run on a computer that has only the x64 instruction set. The only way you can run x86 programs on a 64 bit x64 CPU is if there is a x86 compatibility mode implemented — which all Intel and AMD CPUs do, because the traditional boot process relies on it, if nothing else.

Intel Strategy

Intel Is Reportedly Working On A Brand New Implementation of The x86 Architecture: One That Is Much Faster and Leaner

By Usman Pirzada

Dec 26, 2016 18:40 EST

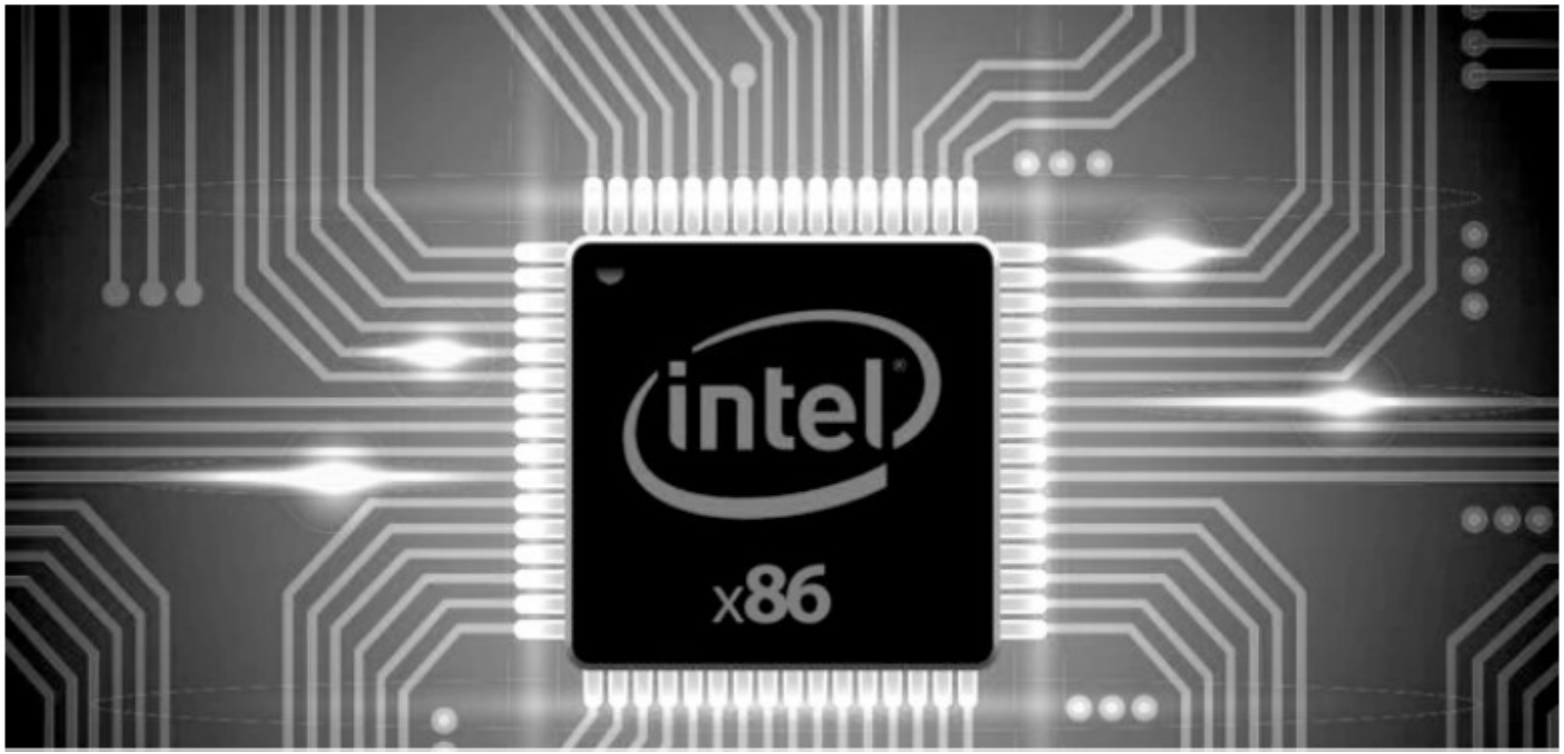

Intel Strategy

The Intel 'Core' legacy will come to an end with Tiger Lake in 2019 - To be succeeded by a lean and mean approach to x86

So when they say new architecture, what exactly does that mean. We talk about architectures a lot as far as GPUs go but the fact of the matter is that the architecture of a CPU works in a slightly different way. Where GPUs have historically freely shifted architectures on a very significant scale, Intel really hasn't shifted from its primary one in ages. As the inventor of the x86 architecture, it has a cross-licensing agreement with AMD to utilize the x86_64 extension of the same architecture (which is basically the 64-bit implementation of x86). It's a funny state of affairs since Intel cannot offer 64-bit compatibility without licensing x86_64 from AMD and the latter cannot ship 64-bit processors unless it has licensed the base x86 from Intel.

All this time, and through the generations, however, Intel processors have remained 100% backward compatible with all previous iterations. The same basic architecture expanded over and over with new features with time. Every new Intel architecture that we talk about on here (or every tock of Intel's PAO cadence) is essentially the same underlying x86 architecture expanded with new features with every iteration. For the first time, however, and if this rumor turns out to be true, things might actually change. Intel might introduce an x86 architecture that we can consider to be truly different from what we have seen so far. The reason? It will be a lean mean, x86-on a diet where backward compatibility is no longer assured. Intel's PAO cadence) is essentially the same underlying x86 architecture expanded with new features with every iteration.

Intel

Fabs/Foundry

Intel at TSMC: 3nm

July 2021

Apple and Intel become first to adopt TSMC's latest chip tech

Seeing Apple using TSMC's latest node is no surprise....but Intel?! Yes, Nikkei Asia has apparently learned that Apple and Intel have emerged as the first adopters of TSMC's 3nm process node. I would've thought AMD would have been behind Apple...but it looks like Intel must have paid a pretty penny to ensure itself in the front of the line.

- Intel beats AMD to 3nm at TSMC

Intel's Newest Fabs

$**20B** in Ohio

Online end of **2025**

SATURDAY, JANUARY 22, 2022 A9

Intel to build chip factories in Ohio

Company will invest $20 billion as a global shortage highlights the risks of reliance on manufacturers in Asia.

Samsung in Texas

Chipmakers are diversifying their manufacturing sites in response to the shortages. Samsung said in November that it planned to build a $17-billion factory outside Austin, Texas.

Micron Technology, based in Boise, Idaho, said it would invest $150 billion globally over the next decade in developing its line of memory chips, with a potential U.S. manufacturing expansion if tax credits can help make up for the higher costs of American manufacturing.

Micron globally

10,000 jobs in Ohio

Two chip factories on the 1,000-acre site in Licking County, just east of Columbus, are expected to create 3,000 company jobs and 7,000 construction jobs, and to support tens of thousands of additional jobs for suppliers and partners, the com-

CHIPS for America Act

Lawmakers have been urging House and Senate leaders to fully fund a law meant to address the semiconductor shortage. They want Congress to fully fund the $52-billion CHIPS for America Act, allowing for stateside investment in semiconductor factories.

More on Intel's New Fabs

Semi Wiki

$52B Chips for America is barely a rounding error

When you assume that the Chips for America act is a one time, one shot disbursement spread over a number of years and a number of companies it becomes clear that its not much against TSMC's spend.

It also does not compare to what China as a whole is spending on semiconductor technology.

Basically the US is being outclassed and outgunned by both China and Taiwan (probably part of China in the not too distant future).

Even if Intel got the whole $52B it still couldn't keep up as the spend would be over several years. Never mind that only $10B of the $52B is for fab projects with a $3B limit per project. Essentially the $52B will be spread so thin as to be ineffective versus the focused sharp spend of TSMC.

More on Intel's New Fabs

Semi Wiki

Can the US fabs being built make a difference?

Intel announced two fabs in Arizona at $10B each along with TSMC announcing a 5NM fab in Arizona which by the time its operational will be a drop in the bucket trailing edge fab perhaps meant to mollify the US.

Samsung has announced a $17B in Texas in addition to existing facilities there. It looks like Intel has chosen Ohio for its "megafab" project and Micron is eyeing North Carolina.

While details are scarce, it sounds like the Intel Ohio and Samsung Texas fabs are the most impactful on the US. Samsung would be somewhat less impactful as we assume that bleeding edge technology R&D will continue to be done in Korea making the Texas fab a "fast follower" much as the existing Samsung fab in Texas is today. That leaves Intel Ohio as the only trail blazing R&D facility in the US.

It also remains to be seen if the brain trust in Portland can either be moved or shared with Ohio or if Portland remains the R&D center with Ohio for production.

Intel's New Fab Equip.

EUV from ASML

Intel Places Order for ASML's Extreme Ultraviolet Technology

04:49 AM EST, 01/19/2022 (MT Newswires) -- Intel (INTC) and ASML Holding (ASML) said Wednesday that the US chip maker has placed its first purchase order for ASML's TWINSCAN EXE:5200 extreme ultraviolet high-volume production system. The purchase ... (MT Newswires 04:49 AM ET 01/19/2022)

Intel News: Italy Fab

Market Chatter: Intel, Italy Reportedly Pick Veneto Region as Location for Proposed Chip Factory

4:39 AM ET, 09/26/2022 - MT Newswires

04:39 AM EDT, 09/26/2022 (MT Newswires) -- Intel (INTC) and the Italian government have selected the town of Vigasio in the country's Veneto region as the location for a proposed chip factory, Reuters reported Sunday, citing anonymous sources familiar with the matter.

The factory, with an initial investment of some 4.5 billion euros ($4.36 billion), is part of the company's planned investment of 80 billion euros ($77.34 billion) to grow capacity in Europe, according to the report.

The new facility is expected to create 1,500 jobs and is slated to start operations between 2025 and 2027, Reuters reported.

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE

COMP122

Global Foundries

July 2021

Chip Fab

7 nm **TSMC**

What's a 5nm FinFET transistor look like?

Well... I didn't find a good pic for 5nm, but I found a pic for 7nm here:

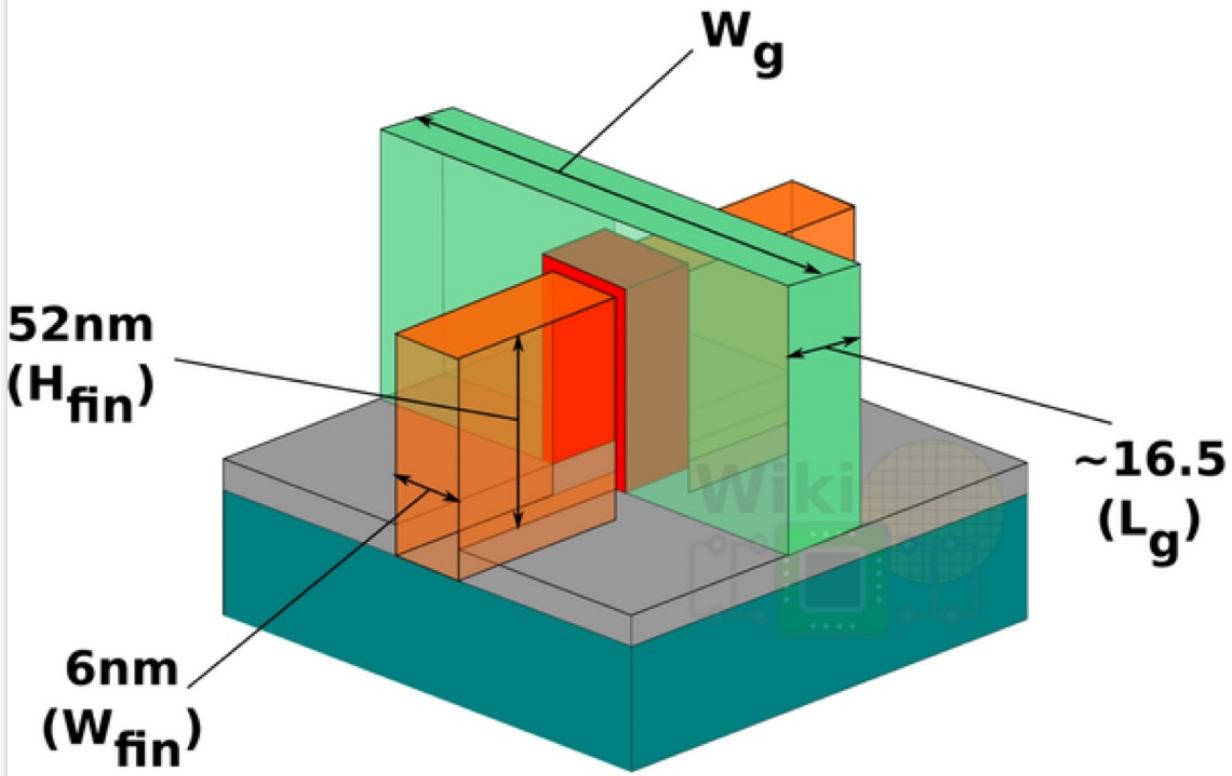

We can assume 5nm is slightly smaller on all axes.

The first thing to note is that the gate length (L_g) is 16.5nm. The width of the gate will vary, but I am going to go out on a limb and guess it is no smaller than $3W_{\text{fin}}$, or 18nm. And the overall structure is tall: 52nm plus another few nm

Intel's New Foundry Strategy

The central tenants of IDM 2.0 for Intel are:

1. Utilize the Intel internal factory network to build the majority of Intel's products internally.
2. Expand use of foundries so that all products have some level of foundry production.
3. Increasing engagement with TSMC, Samsung, GLOBALFOUNDRIES (GF) and UMC.
4. Plan to be a major foundry with US and European based manufacturing to balance the reliance on Asia.

There was slide that showed something like 80% of leading edge in Asia centered around Taiwan and South Korea, 15% in the US and 5% in Europe.

Intel IDM 2.0

Intel CEO Pat Gelsinger sent the following statement announcing creation of an internal foundry model for external customers and Intel product lines and the creation of the **IDM 2.0 Acceleration Office**

Internal Foundry Model Explained

Implementing an internal foundry model means establishing consistent processes, systems and guardrails between our business unit, design and manufacturing teams. This will allow us to identify and address structural inefficiencies that exist in our current model by driving accountability and costs back to decision-makers in real time. It will also put Intel's product groups on a similar footing as external Intel Foundry Services customers and vice versa.

For example, our business unit and design teams will be able to consider the potential impact on their margins if they want to run an additional product stepping, while the manufacturing team will be able to assess requests based on actual costs and impact on factory output. This will give teams the tools and transparency they need to find the most effective and cost-efficient solutions before implementation in silicon, ultimately helping us maximize factory output, reduce costs and shorten design cycles.

Intel Process Nodes

Slower Node Transitions Versus Foundries

IC KNOWLEDGE LLC

	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023
Intel	14nm					10nm				7nm
Samsung	14nm		10nm		7nm	5nm			3nm	
TSMC		16nm	10nm	7nm		5nm		3nm		2nm?

- Intel takes bigger density jumps but less often.
- TSMC and Samsung take smaller jumps more frequently, 5 nodes versus Intel's 3.

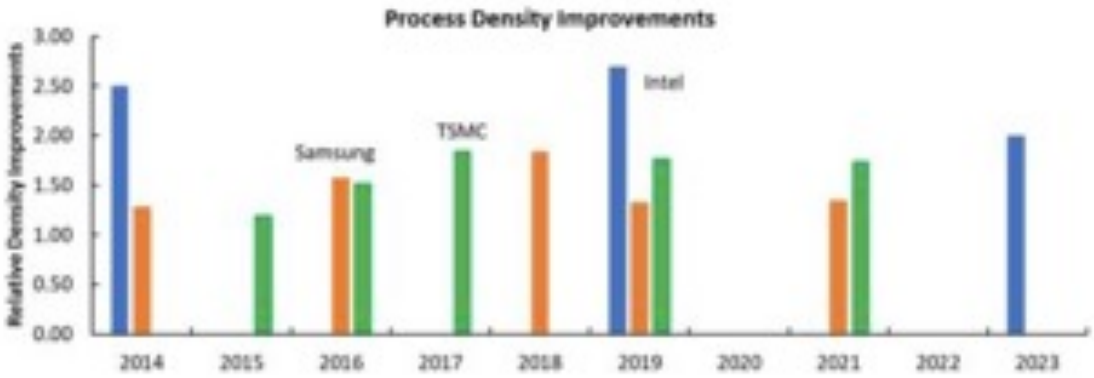

3/24/2021 4

Figure 4. Node Introductions.

Intel Process Nodes

Intel Nodes Versus Time

IC KNOWLEDGE LLC

	Original	Reset 1	Reset 2	Reset 3	Reset 4	Interval (year)	Comments
130nm	**2001**						
90nm	**2003**					2	Strain
65nm	**2005**					2	
45nm	**2007**					2	High-k Metal Gate
32nm	**2009**					2	
22nm	**2011**					2	FinFET
14nm	2013	**2014**				3	1 year late, slow yield ramp
10nm	2015	2017	**2019**			5	2+ years late, still has yield issues
7nm	2017	2021	2022	2023		4?	2 years late?
5nm	2019				2025?	2	Back to 2 years
3nm	2021				2027?	2	

Intel Process Nodes

Node Name Disconnect

IC KNOWLEDGE LLC

- Intel nodes names suggest larger process features relative to foundries.
- Using transistor density to convert Intel nodes to TSMC nodes yields ~4nm for Intel 7nm process and ~2.5nm for Intel 5nm.

	2017	2018	2019	2020	2021	2022	2023	2024	2025	2026	2027
TSMC											
- Node	7	7/7+	5	5	3	3	2	2	1.5	1.5	
- Status	Risk	Full	Risk	Full	Risk	Full	Risk	Full	Risk	Full	
Intel											
- Node	14	14	10	10SF	10SF	7	7	5	5	3	3
- EN	16.1	16.1	7.4	7.4	7.4	4.3	4.3	2.5	2.5	1.5	1.5
- Status	Full	Full	Full	Full	Full	Ramp	Full	Ramp	Full	Ramp	Full

- Risk = risk starts, Ramp = production ramp, Full = full production, EN = TSMC equivalent node.
- TSMC is assumed to stay on a new node every two years with shrink similar to the 5nm and 3nm announced shrinks.
- Intel 5nm and 3nm are assumed to be on two year intervals and to be 2x density improvements consistent with the announced 7nm density shrink.

3/24/2021 6

Figure 6. Node Name Disconnect.

Intel Process Nodes

Figure 3. illustrates Intel's hyper scaling.

Hyper Scaling

- Intel started hyper scaling at the same time that scaling was getting harder.
- 2.5x density improvement for 14nm was 1 year late.
- 2.7x density improvement for 10nm was 2+ years late and still has yield issues.

IC KNOWLEDGE LLC

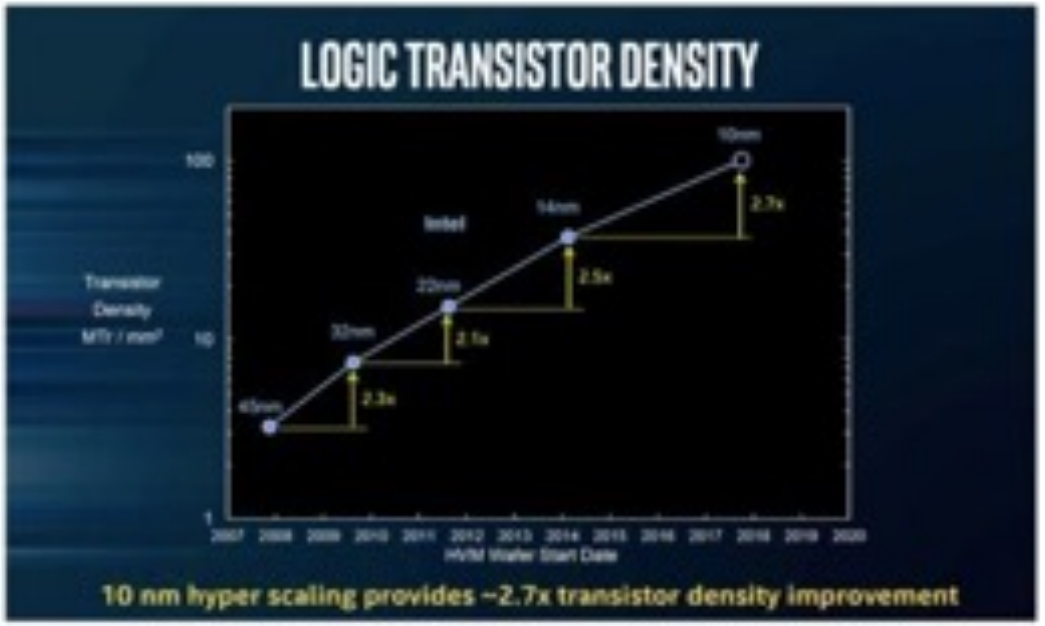

Moore's Law Leadership, Mark Bohr, Intel Manufacturing Day 2017

Intel New Processes

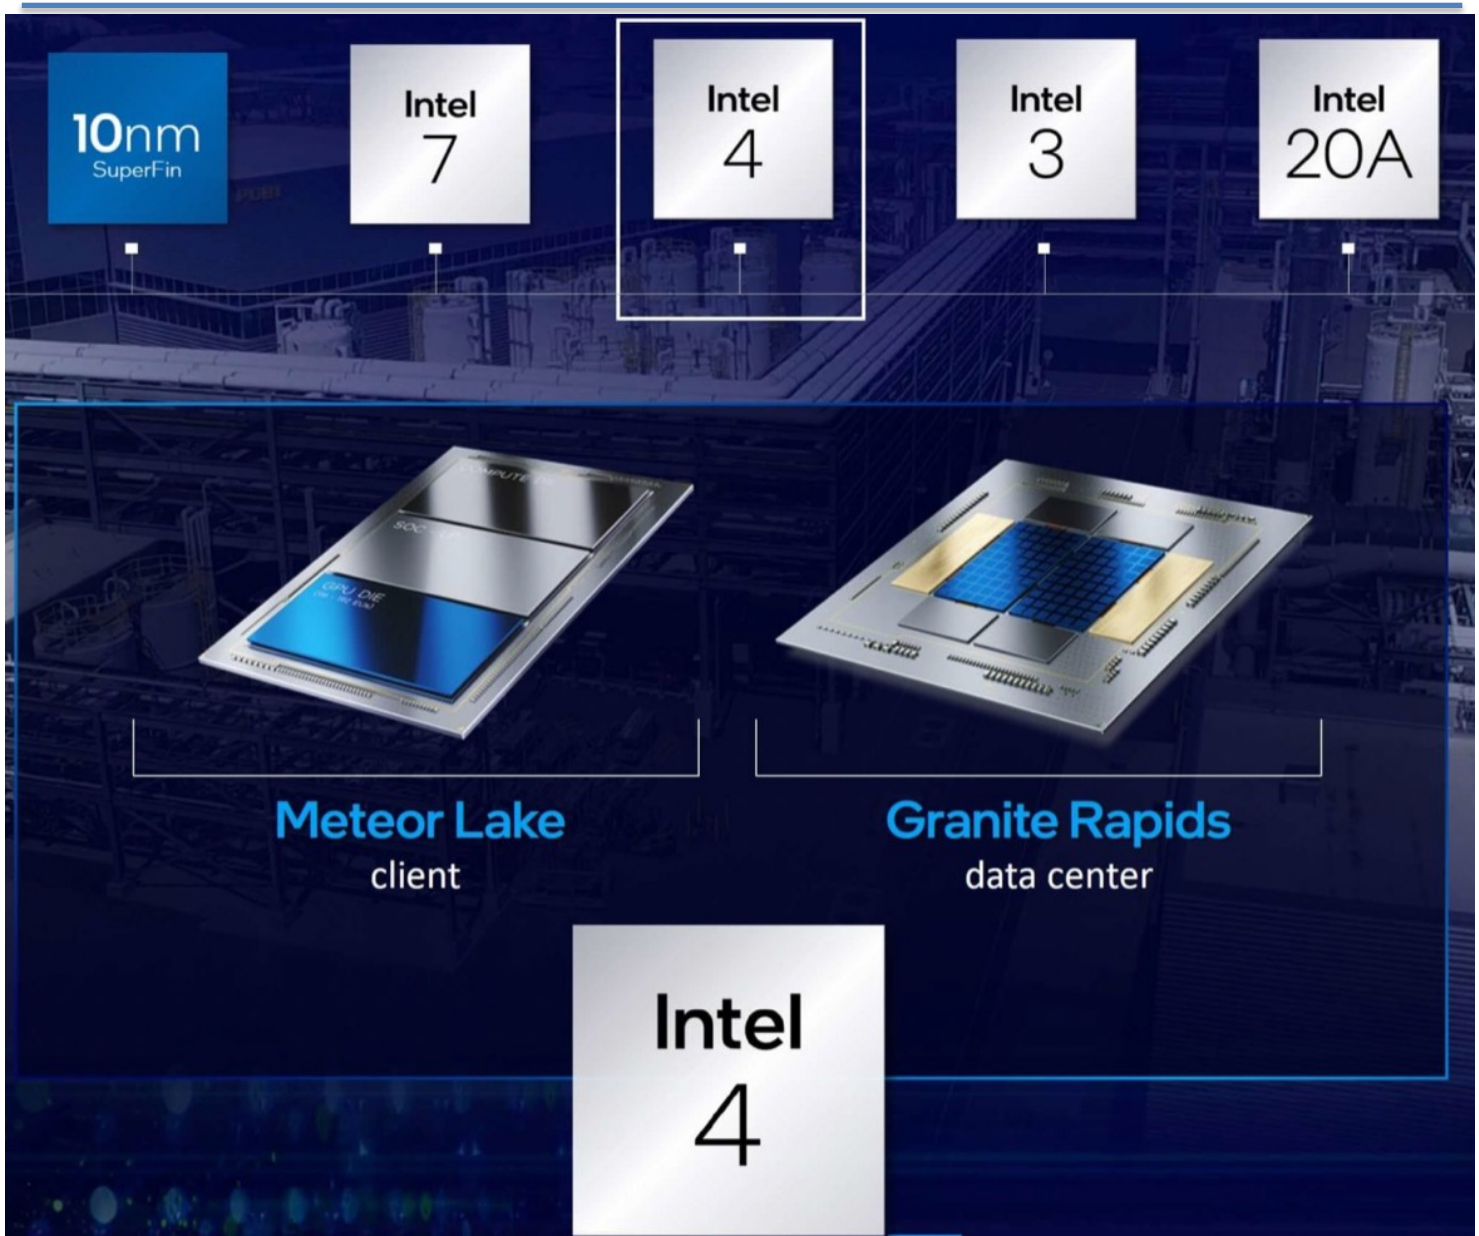

Intel Process Nodes

Intel has said that its own version of 3-D stacking technology, which will premiere on its Ponte Vecchio chip for graphics and A.I. processing, is expected to debut in the Argonne National Laboratory's Aurora supercomputer next year. Next, Intel will use the technology in its Meteor Lake line of desktop processors due in 2023.

Intel Process Roadmap

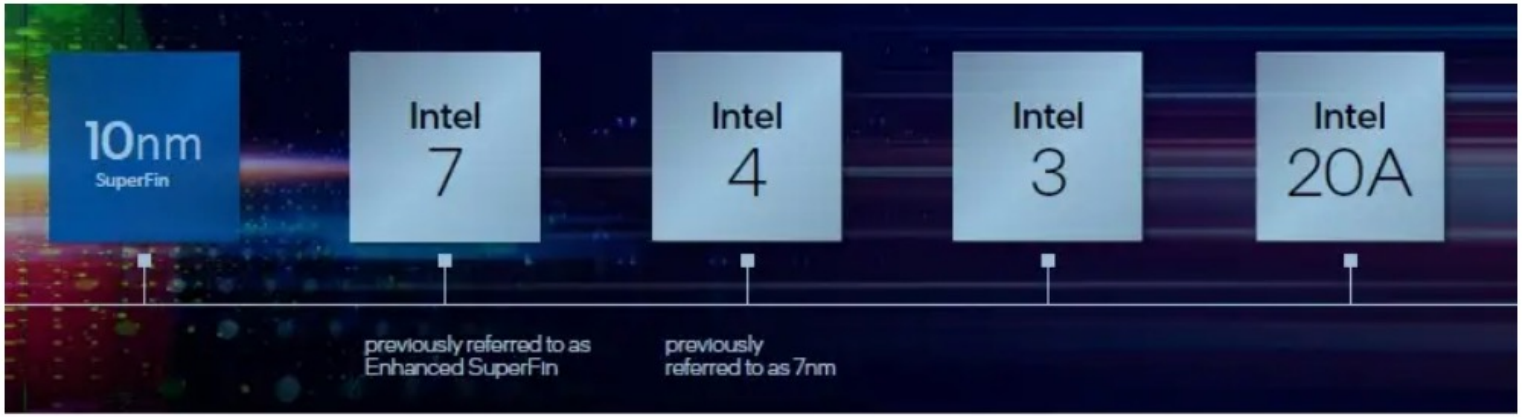

10 SuperFin (10SF)

- currently in volume production,
- "the most significant intra-node performance improvement (over 10+) in Intel's history"

I7 (was originally denoted as '10 Enhanced SuperFin")

Intel Process Roadmap

I7 (was originally denoted as '10 Enhanced SuperFin")

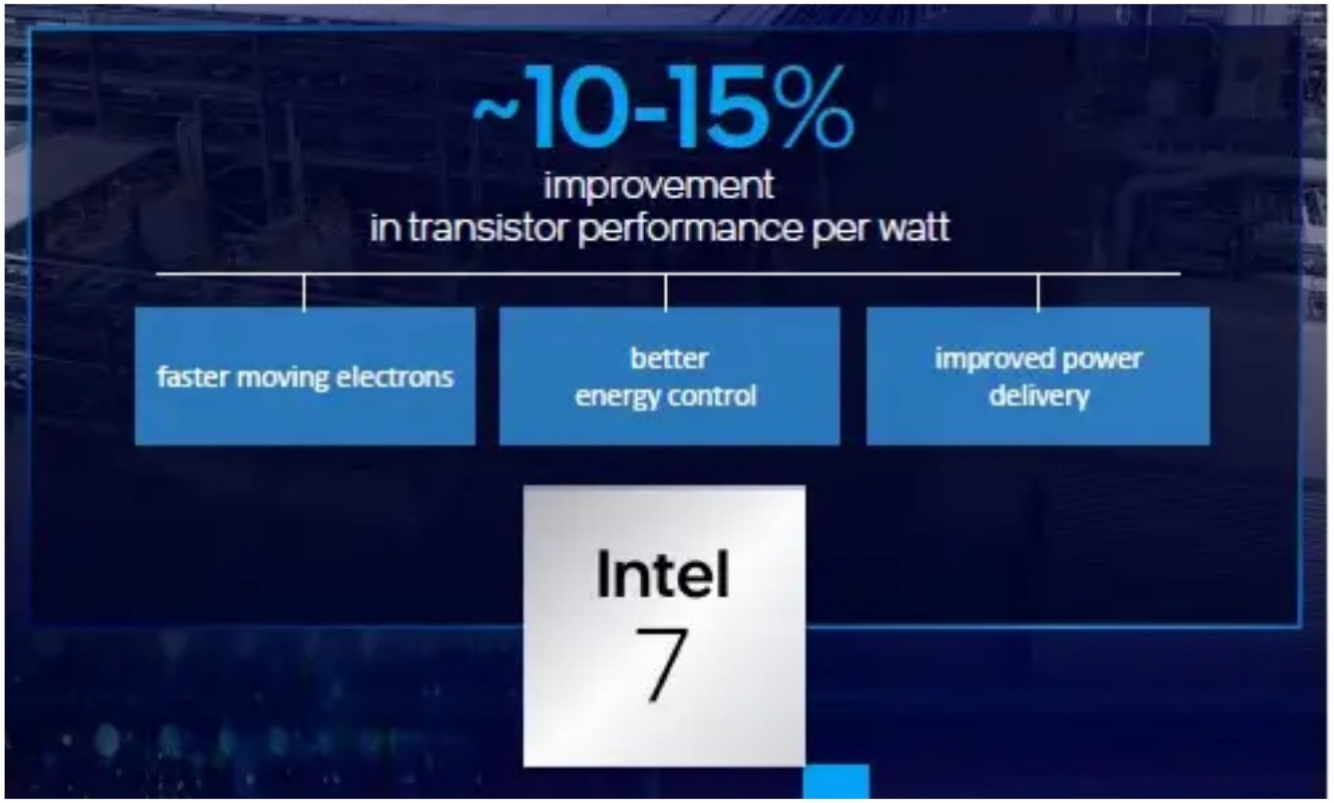

- +10-15% performance per watt gains over 10SF
- increased device channel strain
- Adler Lake (big/little X86 core mix) available late 2021, Sapphire Rapids available 1Q22

Intel Process Roadmap

I3

- +18% performance per watt compared to I4 (FinFET-based)
- new high-performance library
- improved transistor drive current
- improved BEOL characteristics (e.g., reduced via resistance)
- increased EUV deployment
- ramp in 2H2023, HVM in 2024

Intel Process Roadmap

I20A

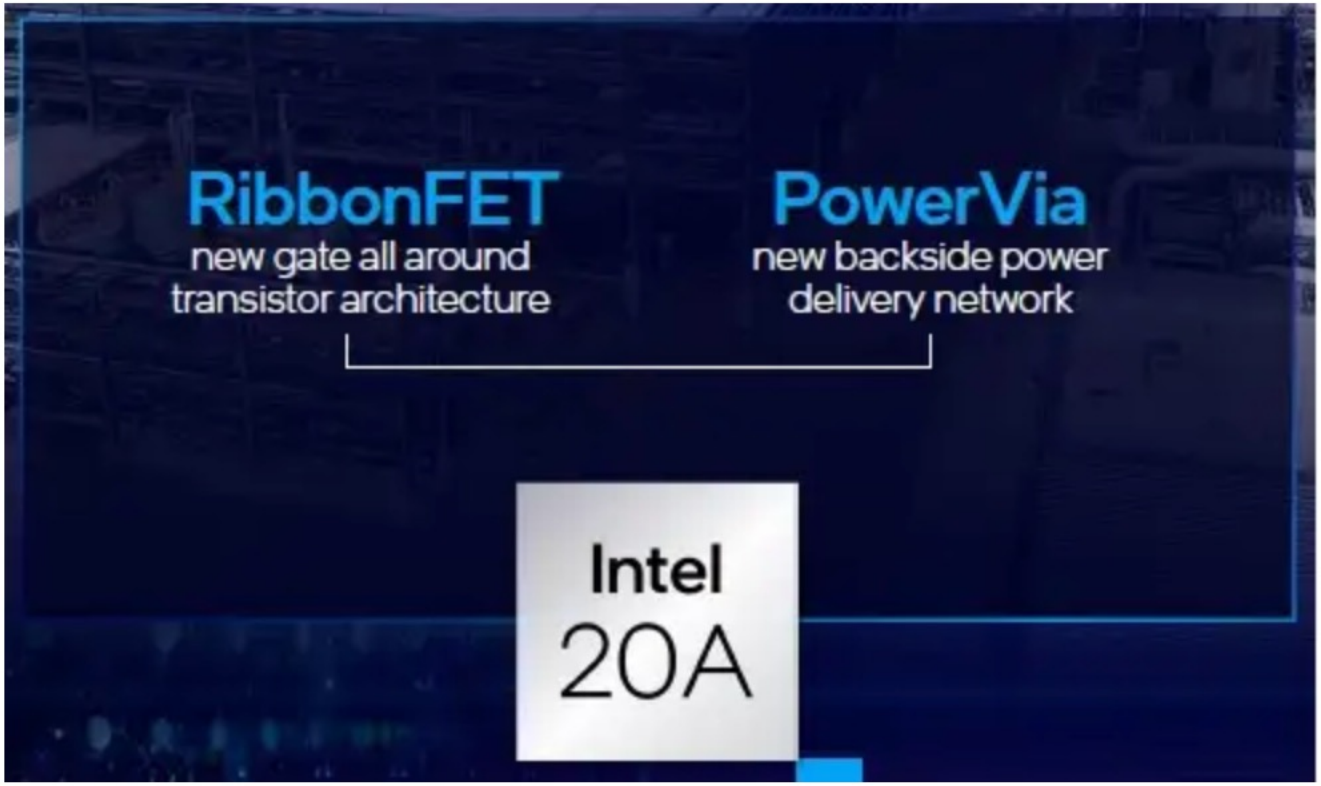

- ramp in 1H2024
- transition to “ribbon FET” (gate-all-around) device topology
- introduction of backside power delivery

Intel Process Roadmap

GAA

As mentioned above, node I20A incorporates a transition from FinFET devices to a gate-all-around configuration, denoted by Intel as "*ribbon FETs*", illustrated below.

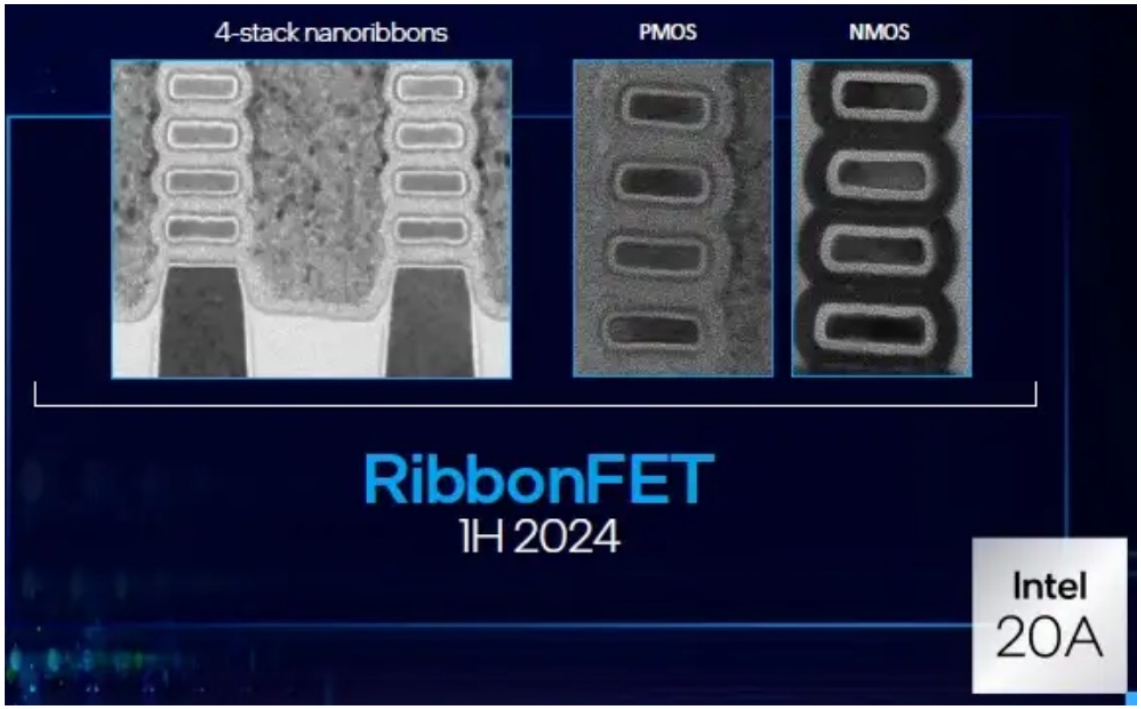

(A cynical engineer would observe that there are several different names being used for the GAA device – so much for nomenclature consolidation in the industry.)

The figure above shows 4 stacked device channels surrounded by the gate –

Intel Process Roadmap

Backside Power Delivery and "Power Vias"

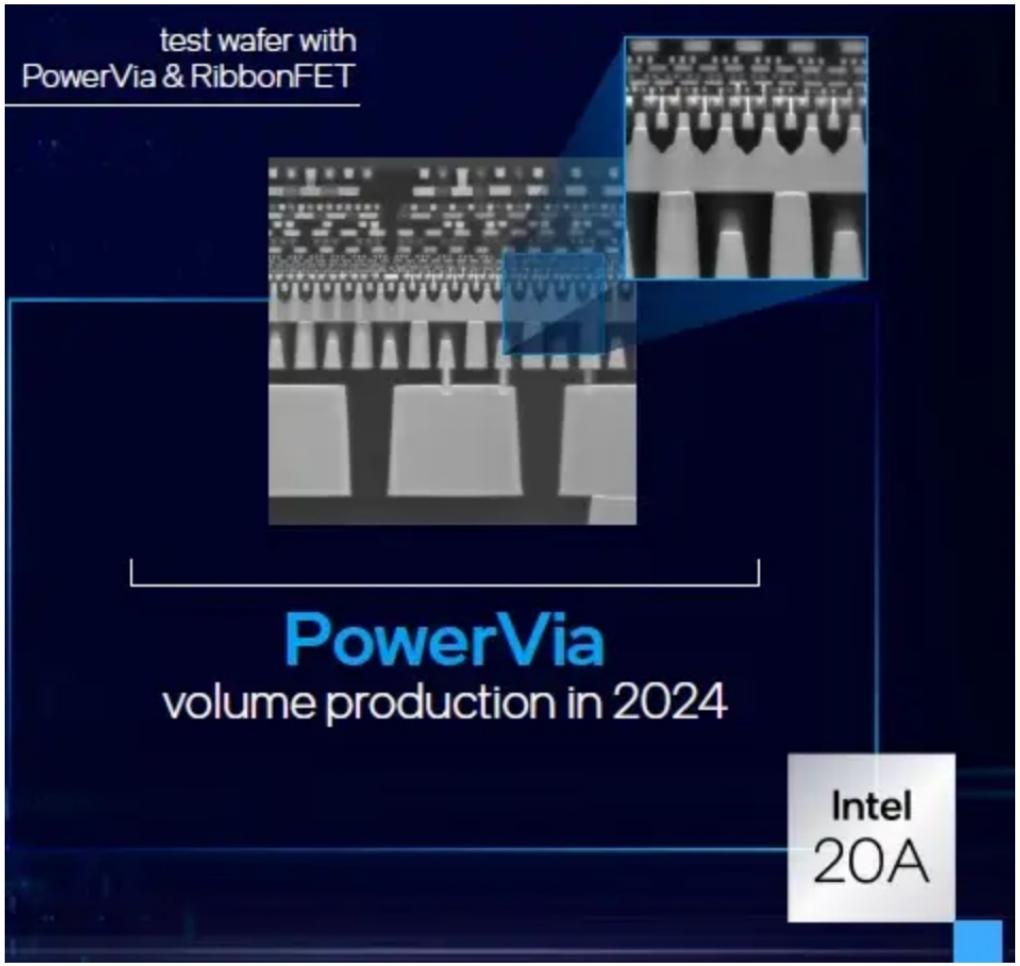

Intel showed die cross-sections of backside power delivery metallization, utilizing "power vias" – the target node is I20A, concurrent with the ribbon FET introduction.

Intel Process Roadmap

EMIB

The Embedded Multi-die Interconnect Bridge (EMIB) provides for 2.5D package interconnections between die, introduced by Intel in 2017. it utilizes a dense pitc microbump connection between the die edges and a silicon "bridge", as shown below. (also, a previous SemiWiki article link)

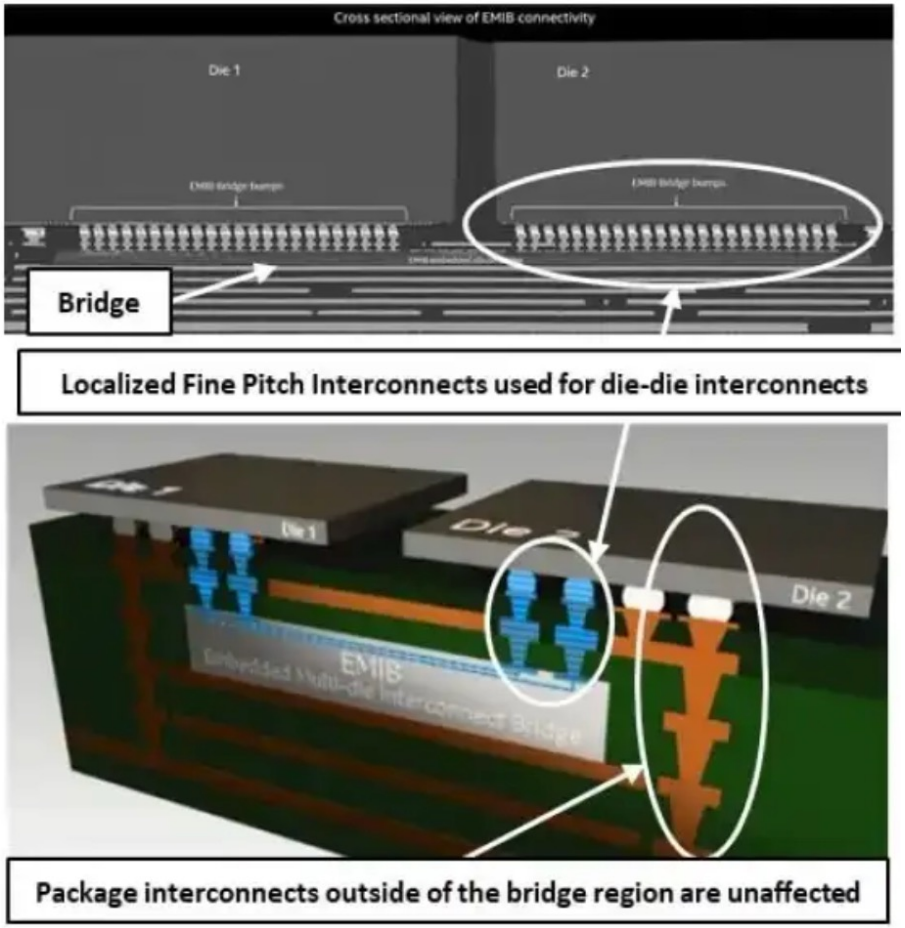

Intel Process Roadmap

Foveros

Intel's 3D stacked die technology is denoted as Foveros. As illustrated in the figures below, Foveros encompasses two configurations – a microbump-based die attach technology, and a direct hybrid bonded connection.

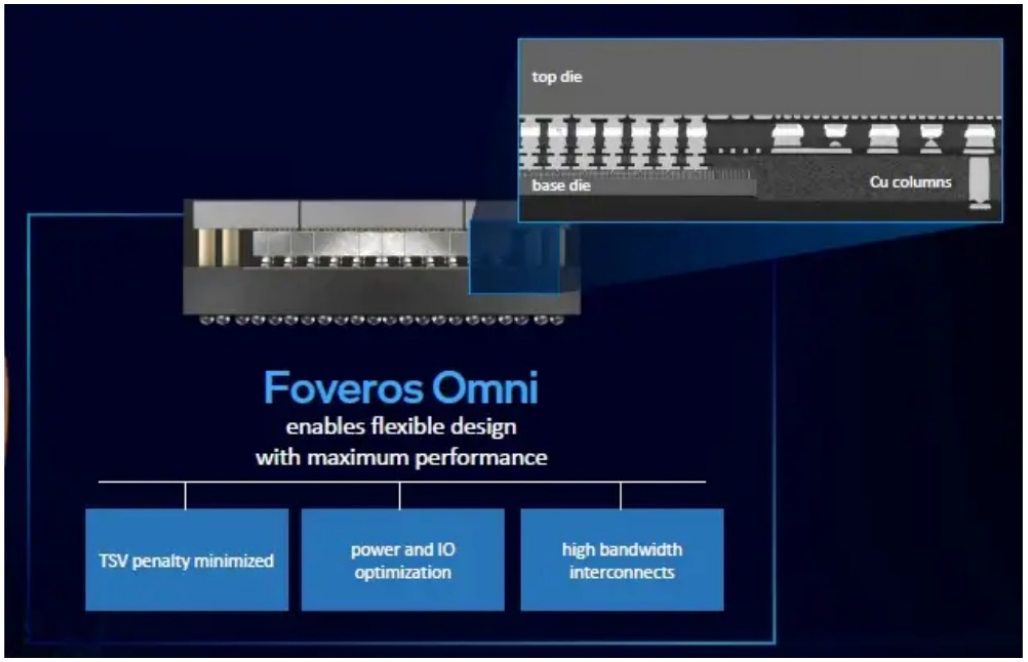

The second generation microbump Foveros technology, denoted as Foveros-Omni, introduces a 36um microbump pitch. The upcoming Meteor Lake product family announcement in process node I4 will showcase the Foveros-Omni offering.

The roadmap presented suggests a subsequent microbump pitch of 25um. Through silicon vias and through package copper columns complete the overall package interconnect offering.

Intel Process Roadmap

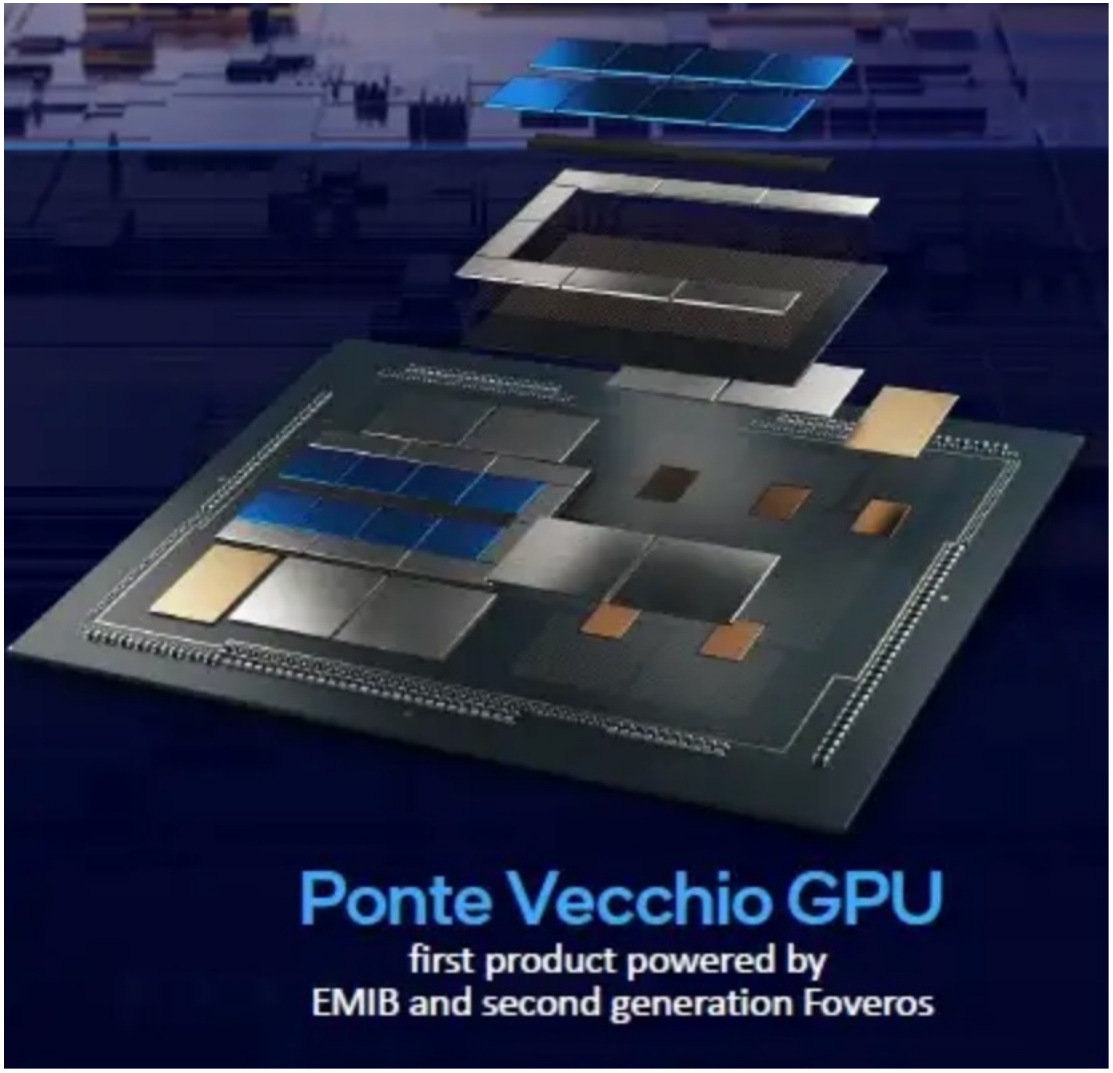

Section

Dev Con 2021

Intel Dev Con 2021

annual "Architecture Day", providing an extensive set of presentations on their product roadmap.

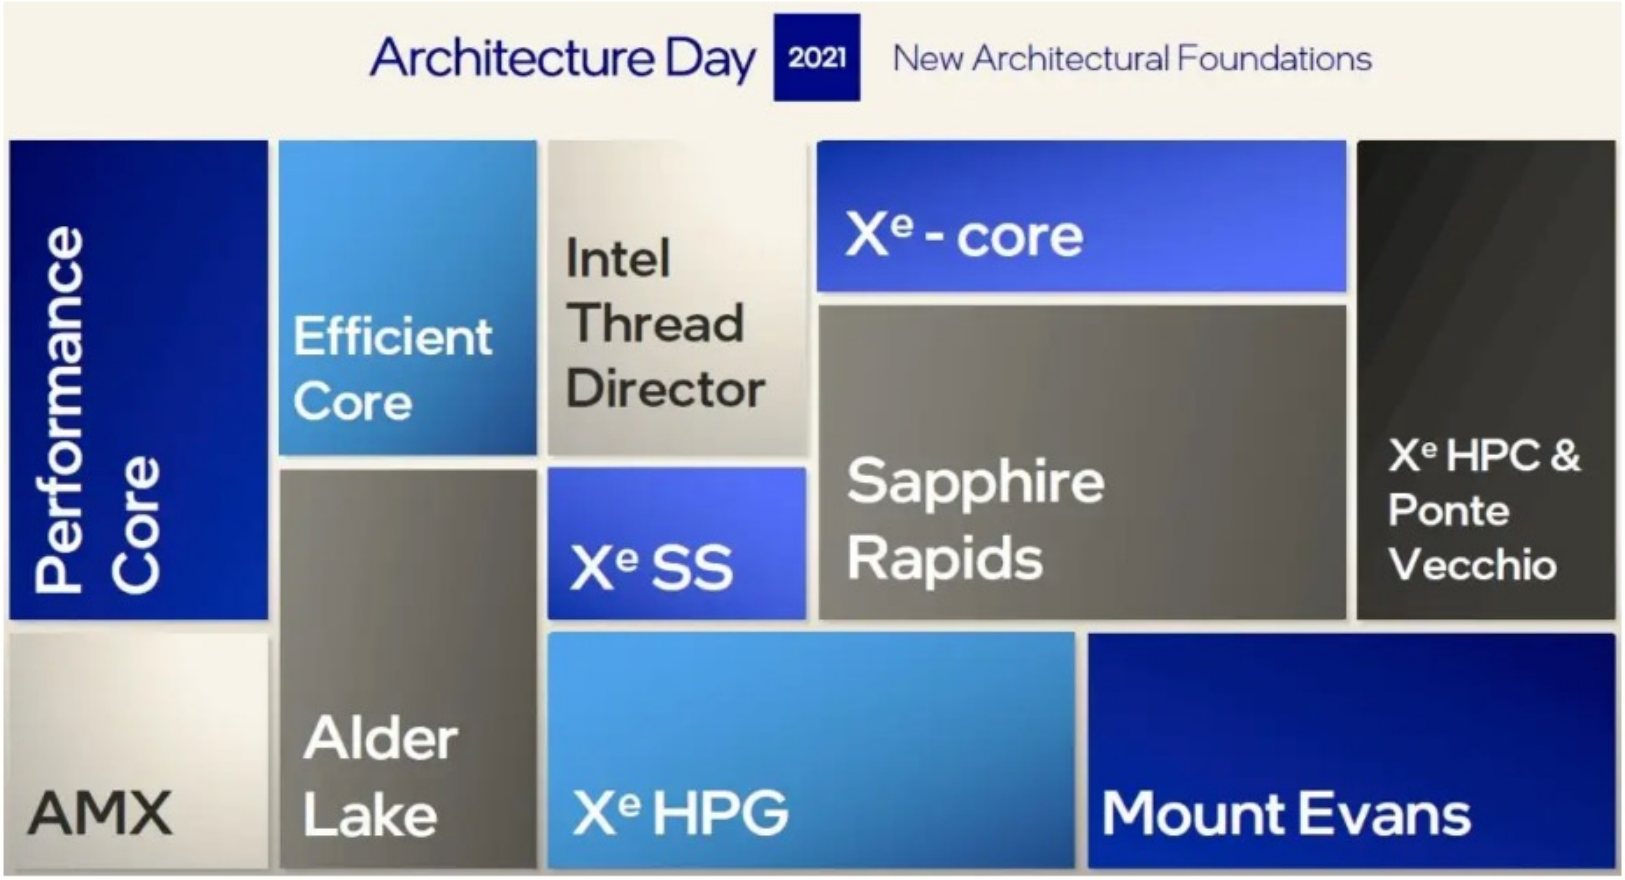

Intel Dev Con 2021

The breadth of topics was vast, encompassing:

- (client and data center) x86 CPUs
- (discrete and integrated) GPUs, from enthusiast and gaming support to high performance AI-centric workloads
- Interface Processing Units (IPUs), to optimize cloud service provider efficiency
- operating system features for managing computing threads in a multi-core complex
- open industry standards for software development application interfaces, focused on the integration of CPU and accelerator devices

Intel Dev Con 2021

An indication of the e-core performance measures is shown below, with comparisons to the previous generation "Skylake" core architecture – one core executing one thread on the left, and four e-cores running four threads on the right:

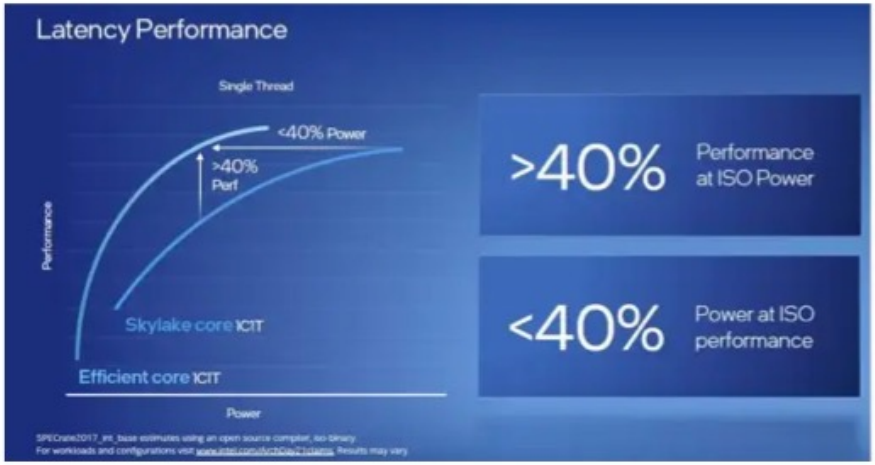

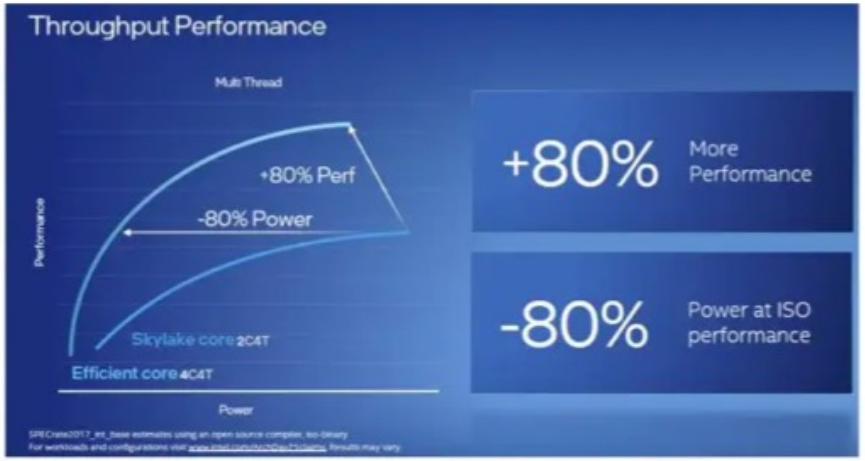

Intel Dev Con 2021

Whereas the efficient-core is a highly scalable microarchitecture focusing on multi-core performance per watt in a small footprint, the performance-core focuses on performance, low-latency and multi-threaded performance, with additional AI acceleration.

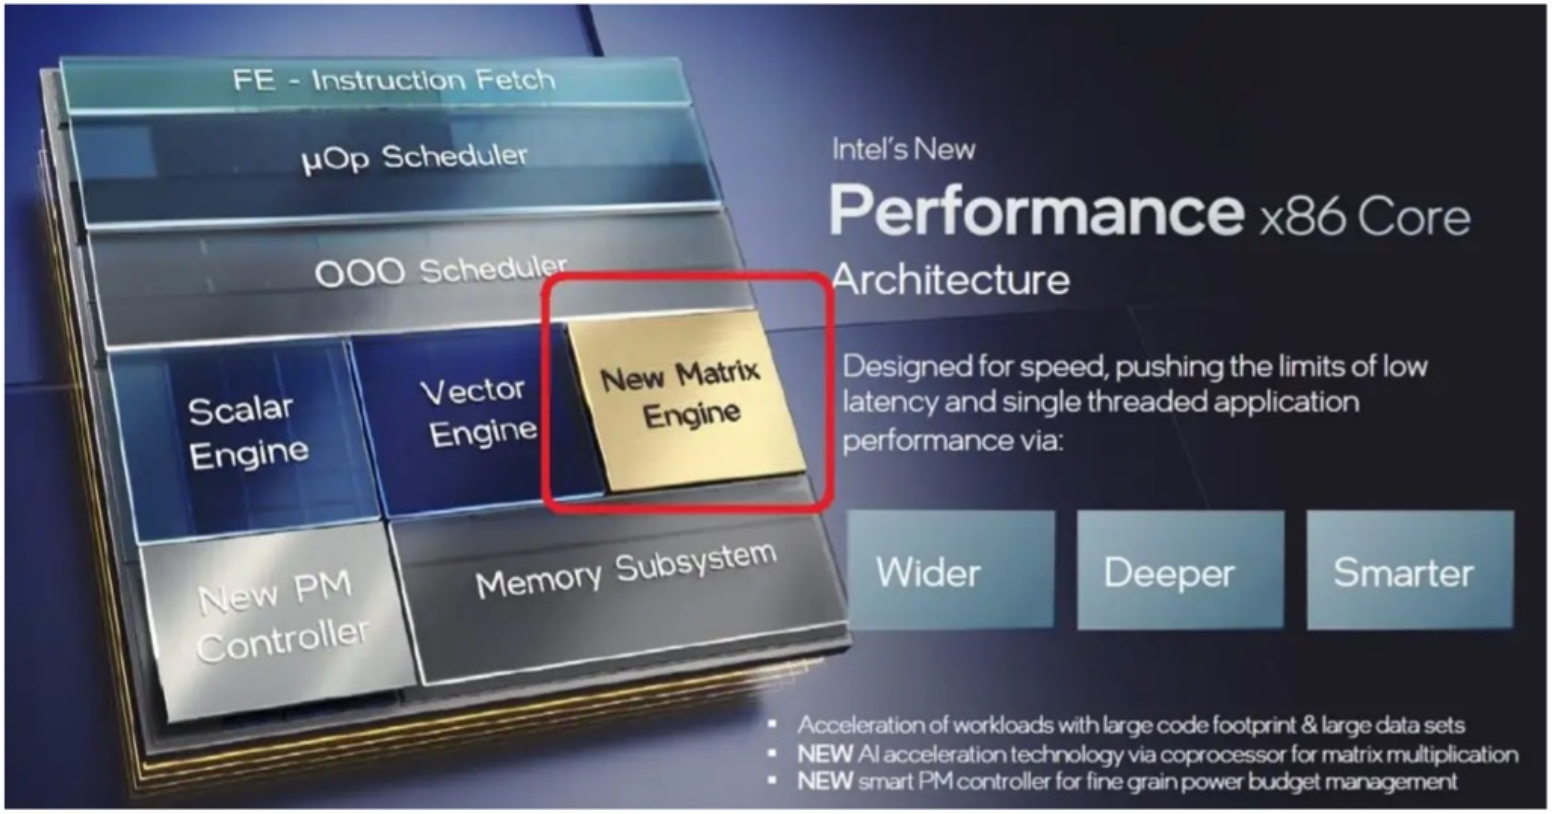

Intel Dev Con 2021

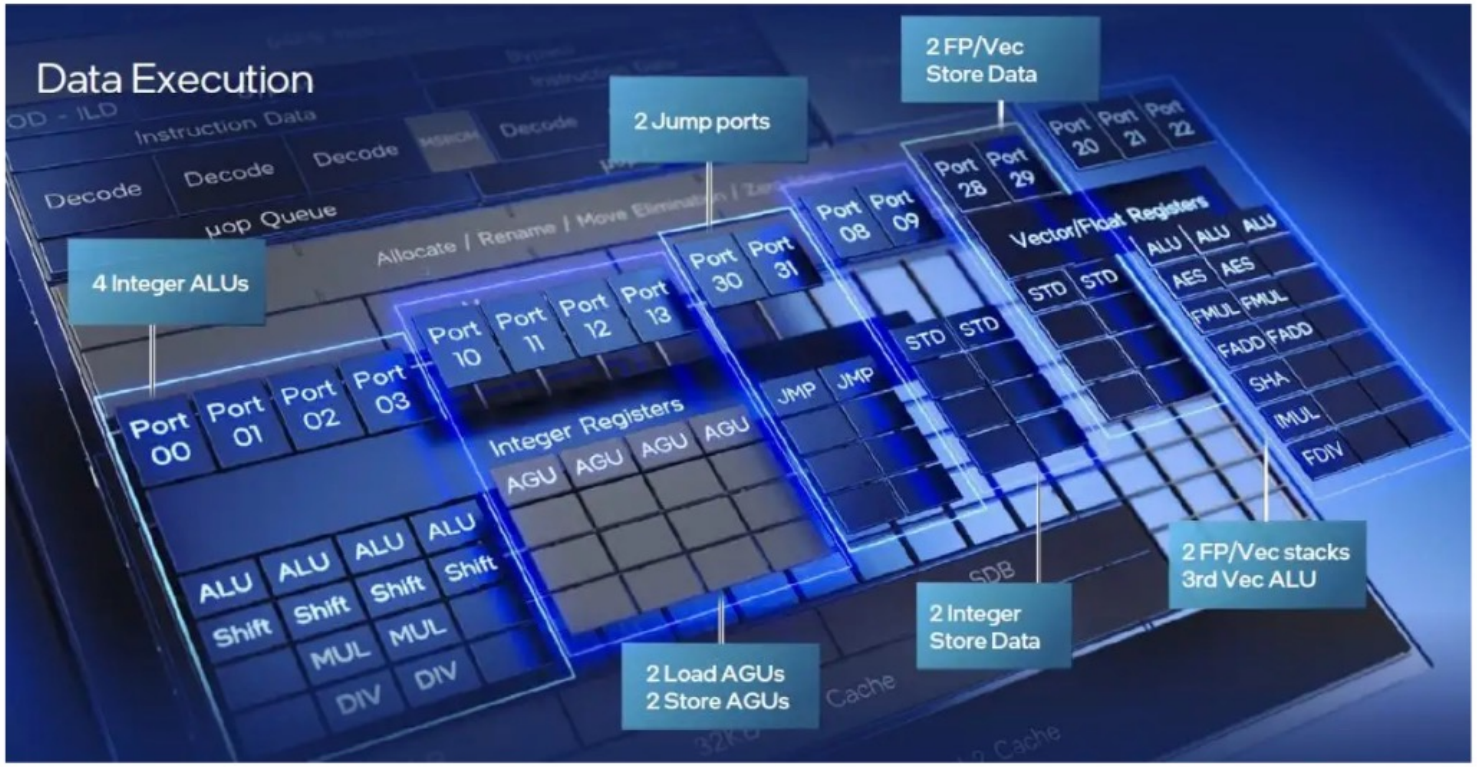

In the figure above, note the units associated with the x86 instructions using vector-based operands, to improve performance of the "dot-product plus accumulate" calculations inherent to deep learning software applications:

- Vector Neural Network Instructions (VNNI, providing int8 calculations)
- Advanced Vector Extensions (AVX-512, for fp16/fp32 calculations)

Intel Dev Con 2021

These instruction extensions accelerate neural network throughput. Active research is underway to determine the optimal data format(s) for neural network inference (with high accuracy), specifically the quantization of larger data types to smaller, more efficient operations – e.g., int4, int8, bfloat16. (The p-core adds another extension to further address machine learning application performance.)

An indication of the e-core performance measures is shown below, with comparisons to the previous generation "Skylake" core architecture – one core executing one thread on the left, and four e-cores running four threads on the right:

Intel Dev Con 2021

"Performance" and "Efficient" x86 Cores

Intel introduced two new x86 core implementations – an "efficient" (e-core) and a performance-centric (p-core) offering.

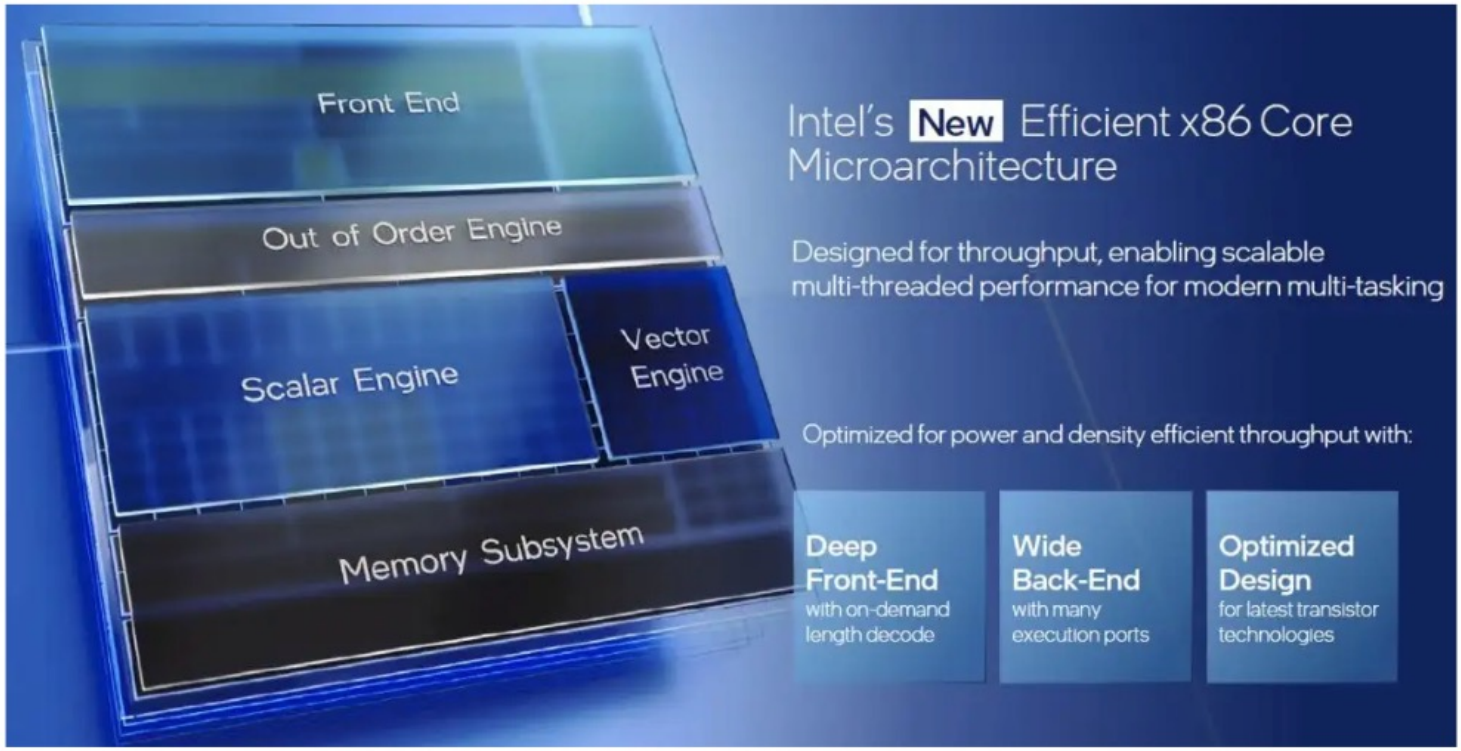

The design considerations for the e-core included:

Intel Dev Con 2021

The design considerations for the e-core included:

- cache pre-fetch strategy
- instruction cache size, and data cache size
- L2$ (shared memory) architecture across cores
- branch prediction efficiency, branch target buffer entries
- instruction prefetch bandwidth, instruction retire bandwidth
- x86 complex instruction micro-op decode and reuse strategy
- out-of-order instruction dependency management resources (e.g., allocate/rename register space)
- configuration of various execution units, and address generation load/store units

To maximize the power efficiency of the e-core, a wide (dynamic) supply voltage range is supported.

Intel Dev Con 2021

Perhaps the most noteworthy addition to the p-core is the Advanced Matrix Extension instruction set. Whereas vector-based data serve as operands for AVX instructions, the AMX operations work on two-dimensional datasets.

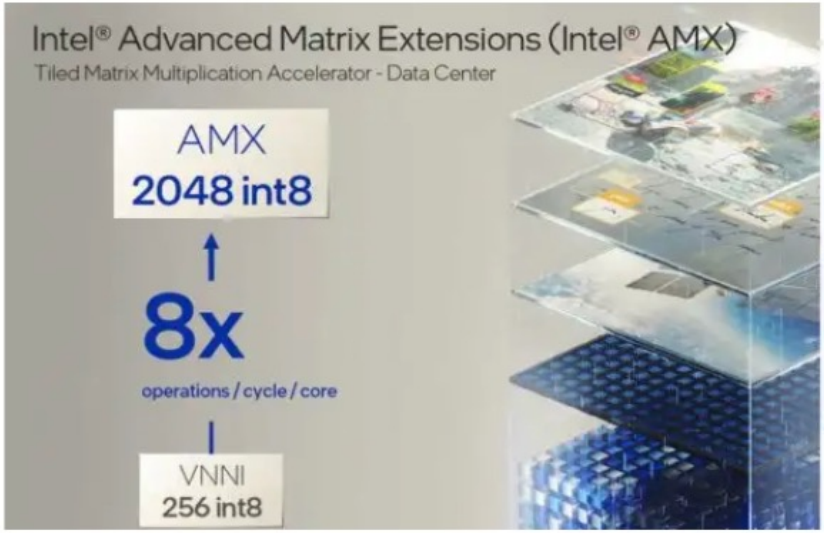

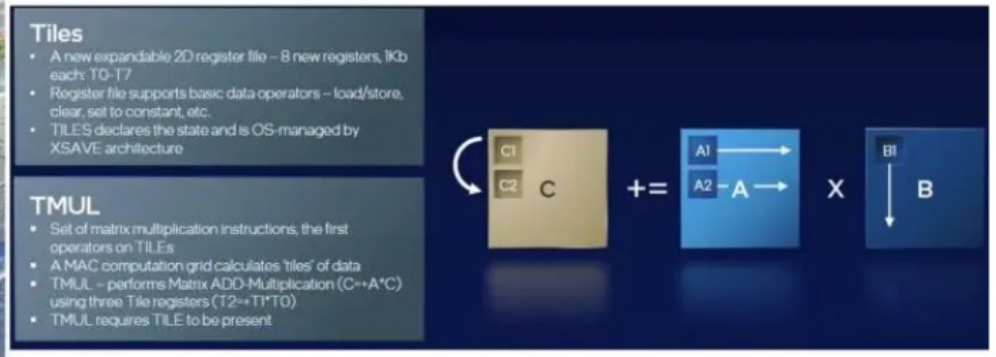

Intel Dev Con 2021

Silicon “tiles” representing 2D register files are integrated with “TMUL” engines providing the matrix operations, as illustrated above.

The addition of AMX functionality is an indication of the diversity of AI workloads. The largest of deep learning neural networks utilize GPU-based hardware for both training and (batch > 1) inference. Yet, there are many AI applications where a relatively shallow network (often with batch = 1) is utilized – and, as mentioned earlier, the use of smaller data types for inference may provide sufficient accuracy, with better power/performance efficiency. It will be very interesting to see how a general purpose CPU with AMX extensions competes with GPUs (or other specialized hardware accelerators) for ML applications.

Intel Dev Con 2021

As mentioned above, applications are selecting a more diverse set of data formats – the p-core also adds fp16 operation support.

Intel Dev Con 2021

Perhaps the most noteworthy addition to the p-core is the Advanced Matrix Extension instruction set. Whereas vector-based data serve as operands for AVX instructions, the AMX operations work on two-dimensional datasets.

Silicon "tiles" representing 2D register files are integrated with "TMUL" engines providing the matrix operations, as illustrated above.

The addition of AMX functionality is an indication of the diversity of AI workloads. The largest of deep learning neural networks utilize GPU-based

Intel Dev Con 2021

Additionally, based on thread priority, an executing thread could transition between a p-core and e-core. Also, threads may be “parked” or “unparked”.

Intel Dev Con 2021

Intel Dev Con 2021

Thread Director

A key performance optimization in any computer architecture is the scheduling of program execution threads by the operating system onto the processing resources available.

One specific tradeoff is the allocation of a new thread to a core currently executing an existing thread. "Hyperthread-enabled" cores present two logical (virtual) processors to the O/S scheduler. Dual architectural state is provided in the core, with a single execution pipeline. Register, code return stack buffers, etc. are duplicated to support the two threads, at a small cost in silicon area, while subsets of other resources are statically allocated to the threads. Caches are shared. If execution of one thread is stalled, the other is enabled. The cache memory offers some benefit to the two threads, as shared code libraries may be common between threads of the same process.

Intel Dev Con 2021

Thread Director

Another option is to distribute thread execution across separate (symmetric) cores on the CPU until all cores are busy, before invoking hyperthreading.

A combination of p-cores and e-cores in the same CPU (otherwise known as a "big/little" architecture) introduces asymmetry into the O/S scheduler algorithm. The simplest approach would be to distinguish threads based on foreground (performance) and background (efficiency) processes – e.g., using "static" rules for scheduling. For the upcoming CPUs with both p- and e-cores, Intel has integrated additional power/performance monitoring circuitry to provide the O/S scheduler with "hints" on the optimum core

Intel Dev Con 2021

Alder Lake

The first release of a client CPU with the new p- and c-cores will be Alder Lake, expected to be formally announced at the Intel InnovatiON event in October.

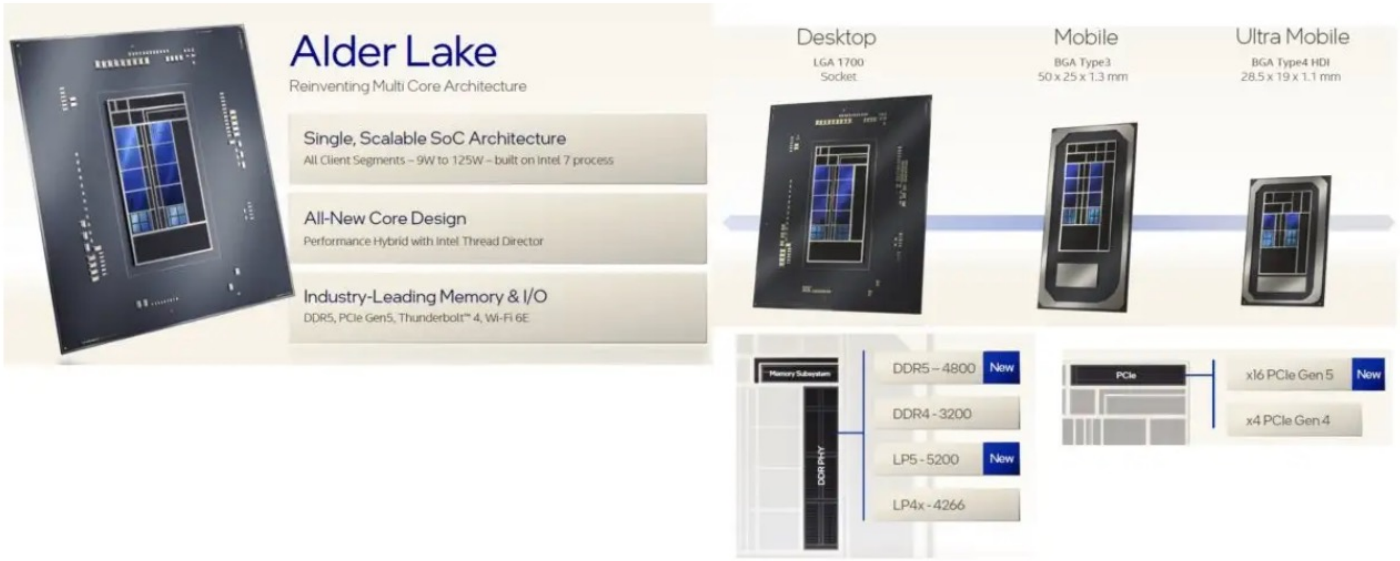

In addition to the new cores, Alder Lake incorporates PCIe Gen 5 and DDR5 memory interfaces. The Alder Lake product family will span a range of target markets, from desktop (125W) to mobile to ultra-mobile (9W), with an integrated GPU core (I7 node).

Intel Dev Con 2021

Sapphire Rapids

High Bandwidth Memory

Significantly Higher Memory Bandwidth
vs. baseline Xeon-SP with 8 channels of DDR 5

Increased capacity and Bandwidth
some usages can eliminate need for DDR entirely

2 Modes

HBM Flat Mode	HBM Caching Mode
Flat Mem Regions w/ HBM & DRAM	DRAM backed cache

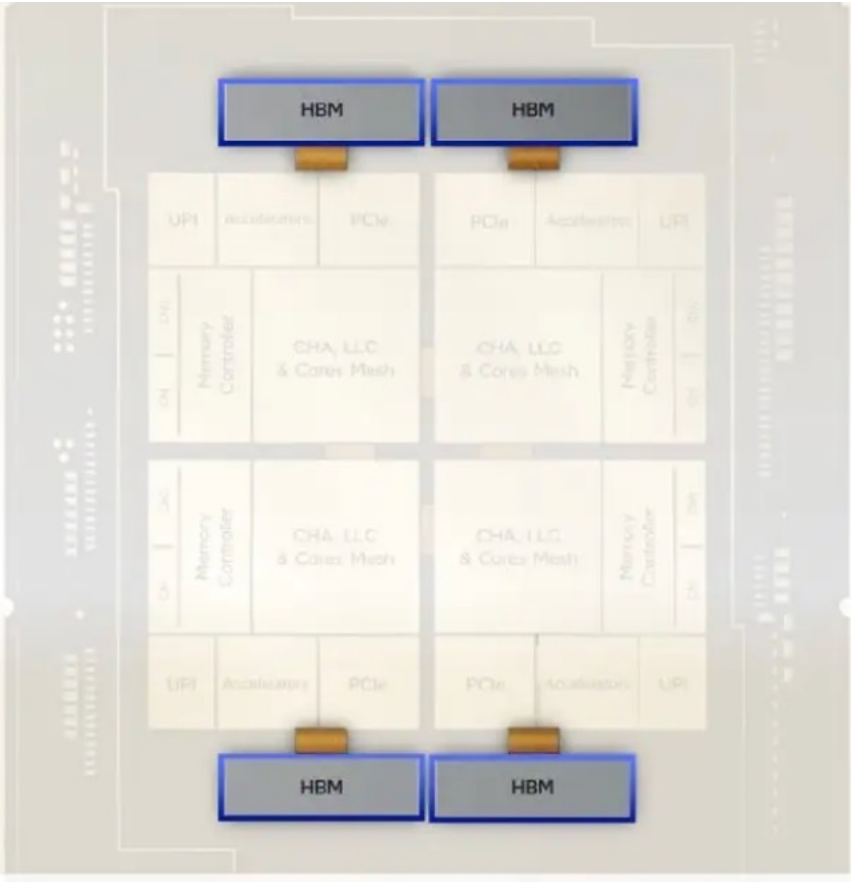

The intent is to provide memory “tiering” – the HBM2 stacks could serve as part of the “flat” memory space, or as caches to external system memory.

Intel Dev Con 2021

Sapphire Rapids

The first data center part with the new p-core will be Sapphire Rapids, to be formally announced in 1Q2022 (Intel7 node).

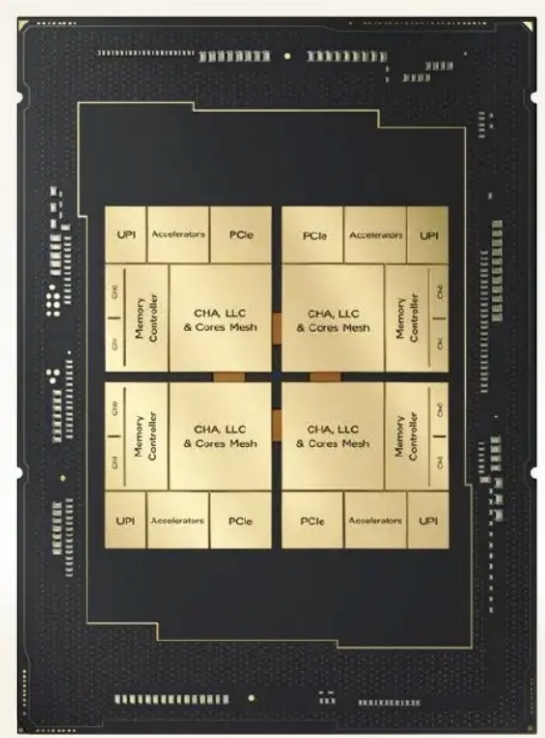

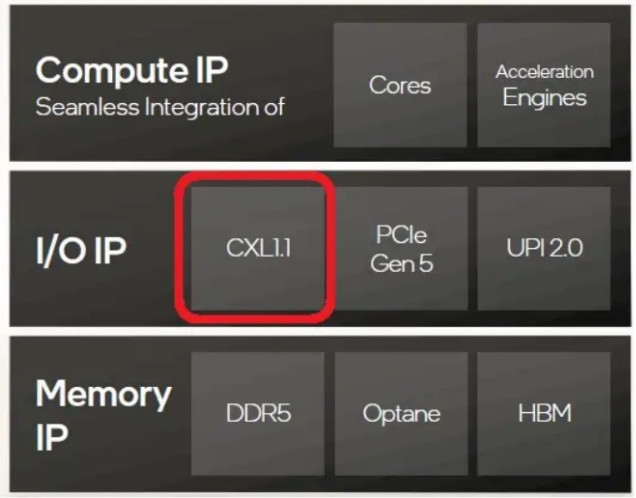

The physical implementation of Sapphire Rapids incorporates “tiles” (also known as “chiplets”), which utilize the unique Intel EMIB bridge silicon for adjacent tile-to-tile dense interconnectivity.

Intel Dev Con 2021

ULTRA formalISA®
Predictable RISC-V formal verification

Coverage analyzer example

Proving that specific instruction sequencing works and visualizing it

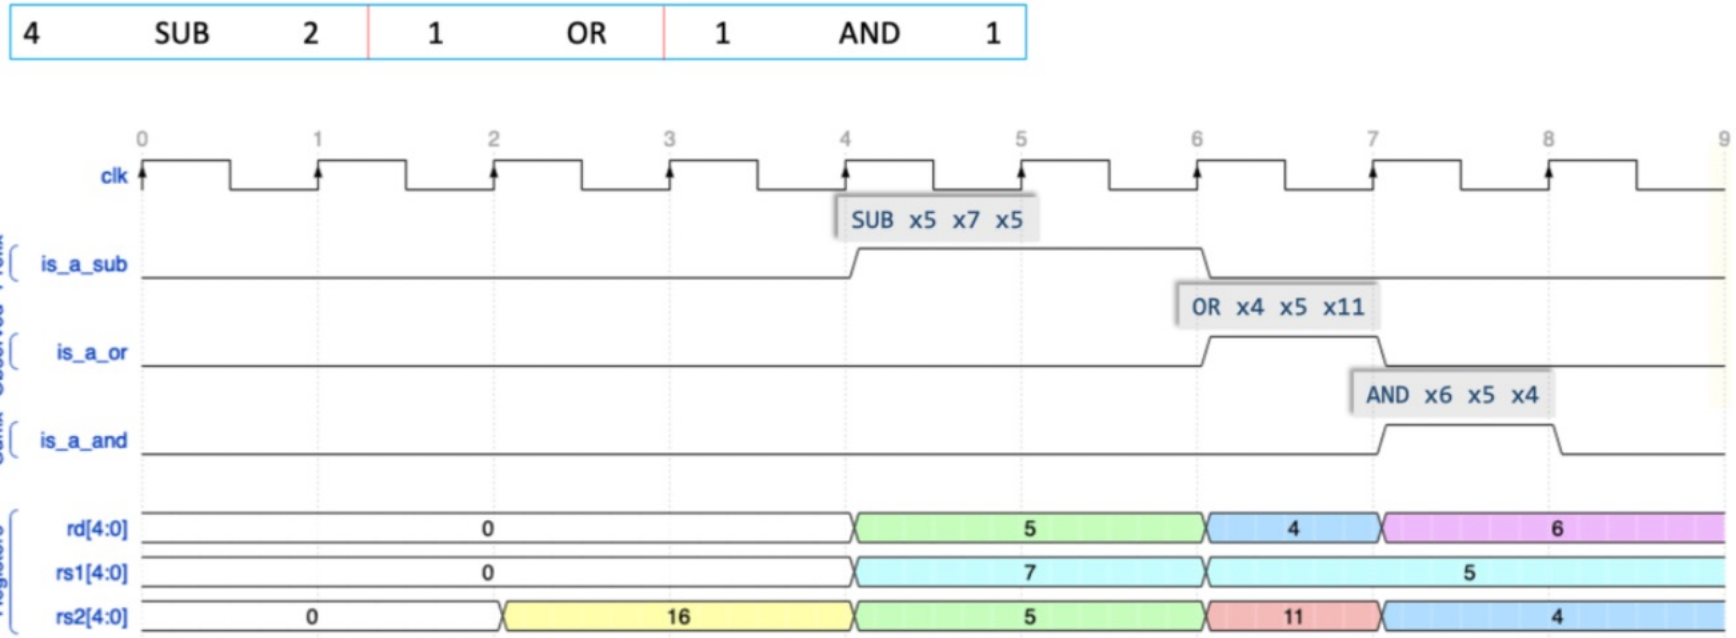

Fig 2. Scenario coverage example

Intel Dev Con 2021

Example of a specification bug caught by *formalISA*

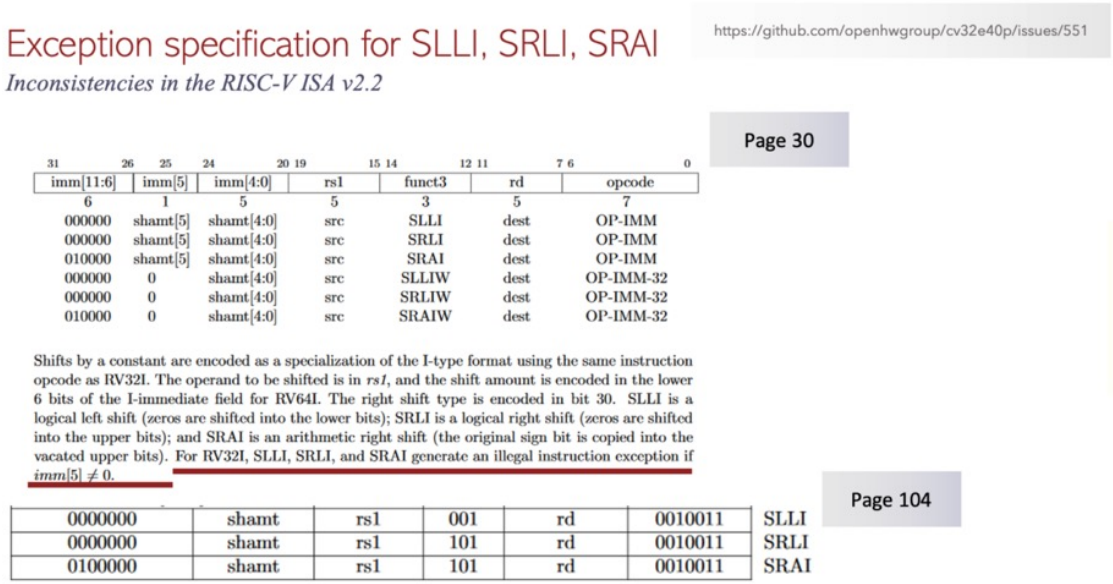

Exception specification for SLLI, SRLI, SRAI

Inconsistencies in the RISC-V ISA v2.2

https://github.com/openhwgroup/cv32e40p/issues/551

Page 30

31 26	25	24 20	19 15	14 12	11 7	6 0
imm[11:6]	imm[5]	imm[4:0]	rs1	funct3	rd	opcode
6	1	5	5	3	5	7
000000	shamt[5]	shamt[4:0]	src	SLLI	dest	OP-IMM
000000	shamt[5]	shamt[4:0]	src	SRLI	dest	OP-IMM
010000	shamt[5]	shamt[4:0]	src	SRAI	dest	OP-IMM
000000	0	shamt[4:0]	src	SLLIW	dest	OP-IMM-32
000000	0	shamt[4:0]	src	SRLIW	dest	OP-IMM-32
010000	0	shamt[4:0]	src	SRAIW	dest	OP-IMM-32

Shifts by a constant are encoded as a specialization of the I-type format using the same instruction opcode as RV32I. The operand to be shifted is in *rs1*, and the shift amount is encoded in the lower 6 bits of the I-immediate field for RV64I. The right shift type is encoded in bit 30. SLLI is a logical left shift (zeros are shifted into the lower bits); SRLI is a logical right shift (zeros are shifted into the upper bits); and SRAI is an arithmetic right shift (the original sign bit is copied into the vacated upper bits). For RV32I, SLLI, SRLI, and SRAI generate an illegal instruction exception if $imm[5] \neq 0$.

Page 104

0000000	shamt	rs1	001	rd	0010011	SLLI
0000000	shamt	rs1	101	rd	0010011	SRLI
0100000	shamt	rs1	101	rd	0010011	SRAI

Fig 1. Inconsistency bug caught by formalISA app.

In the published RISC-V ISA specification v2.2, on page 30 it requires the generation of illegal instruction exception for the shift instructions SLLI, SRLI and SRAI if imm[5] is not equal to 0. However, it is not possible for imm[5] to be non-zero if the instruction decoded is one of the SLLI, SRLI and SRAI as the specification provides the opcodes for these on page 104.

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP122

Intel Dev Con 2021

DR JEFF SOFTWARE
INDIE APP DEVELOPER

Part 2: GPUs, IPUs, XeSS, OpenAPI

INTEL

Intel Architecture Day – Part 2: GPUs, IPUs, XeSS, OpenAPI

by Tom Dillinger on 09-01-2021 at 10:00 am

Introduction

At the recent Intel Architecture Day presentations, a breadth of roadmap plans were provided – an earlier article focused on the x86 client and data center cores and products. This article focuses on the GPU and IPU announcements.

Intel Dev Con 2021

Part 2: GPUs, IPUs, XeSS, OpenAPI

X^e Graphics Core

The Intel GPU architecture for embedded, discrete, and data center acceleration is based on the X^e graphics core. The figure below illustrates the integration of multiple cores with other units to provide a "render slice" block in the overall GPU hierarchy.

Intel Dev Con 2021

Arc

Intel introduced the Arc brand, to refer to the discrete graphics card product roadmap. Arc also incorporates new "unified, re-factored" graphics software drivers.

The first Arc card is codenamed "Alchemist", with a X^e chip fabricated using TSMC's N6 process (1Q2022).

X^eSS Super Sampling

A unique software feature added to the GPU family is X^e super sampling. This image sampling method utilizes a combination of spatial and temporal data to upscale frame resolution, such as a 1080p to 4K video stream.

Intel Dev Con 2021

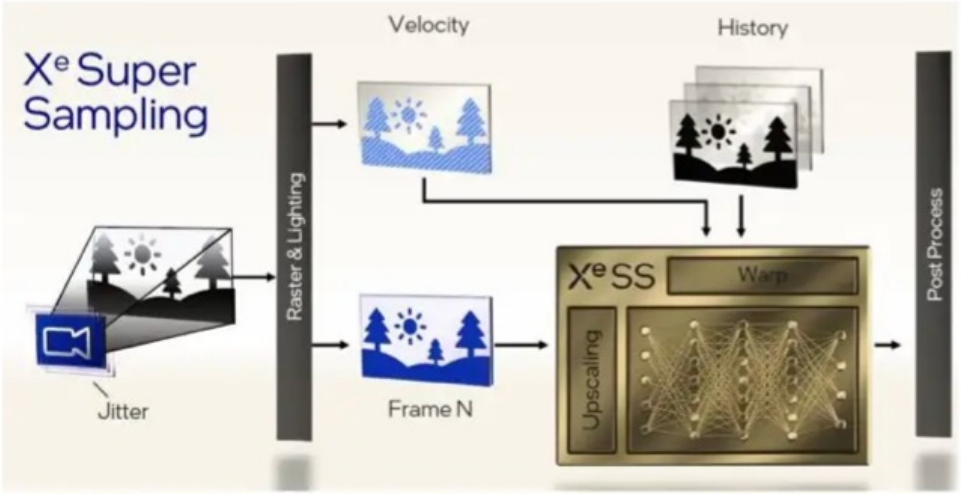

The graph on the right in the figure above illustrates the time to render an image for various methods (shorter is better). The XeSS algorithm combined with the XMX vector accelerator unit enables excellent 4K image throughput. Intel provided a demo of a 1080p video upscaled using XeSS to 4K resolution – the distinction between the upscaled and native 4K video (@60fps) was imperceptible, offering a unique power/performance

Intel Dev Con 2021

Infrastructure Processing Unit (IPU)

Intel provided a very compelling picture of the inefficiencies in current cloud data center services. The figure below shows that a "traditional" CPU plus SmartNIC cloud server architecture requires that the CPU spends considerable cycles performing infrastructure micro-services, such as storage management, security authentication, data encryption/decryption – a range of 31% to 83% overhead, as illustrated below.

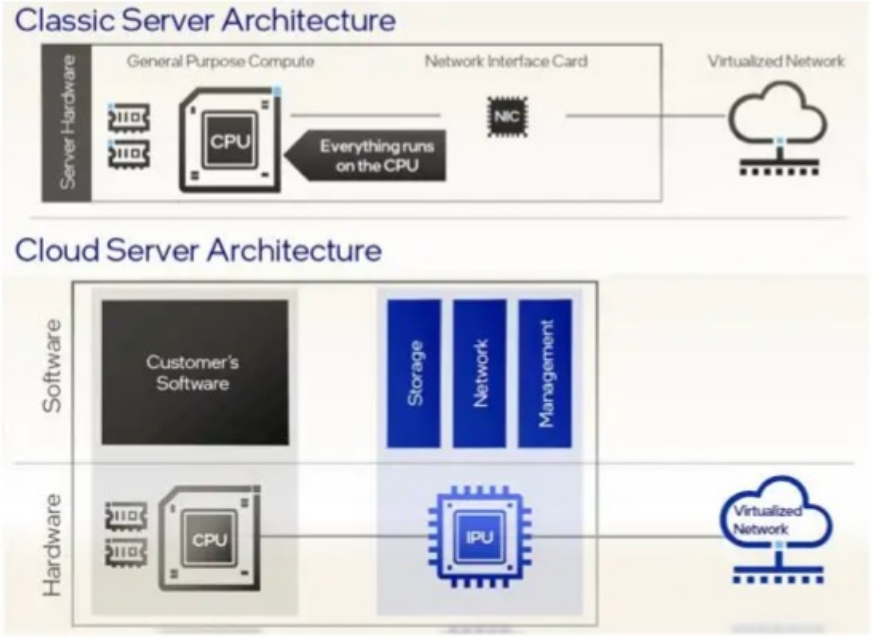

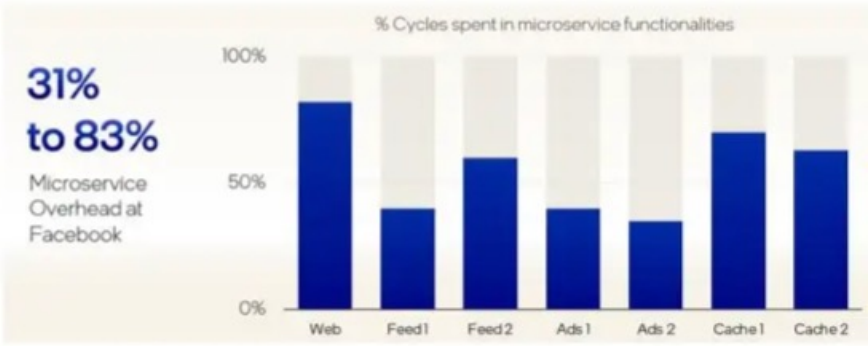

Intel Dev Con 2021

The X^e link tile (TSMC N7) enables direct connection of a variety of GPU topologies, as shown below.

Preliminary (A0 silicon) performance measurements indicated an extremely competitive positioning relative to the GPUs in prevalent use in today's data centers.

Intel Dev Con 2021

Ponte Vecchio High Performance Computing GPU

The most advanced illustration of Intel's packaging technology was provided as part of the Ponte Vecchio data center GPU presentation.

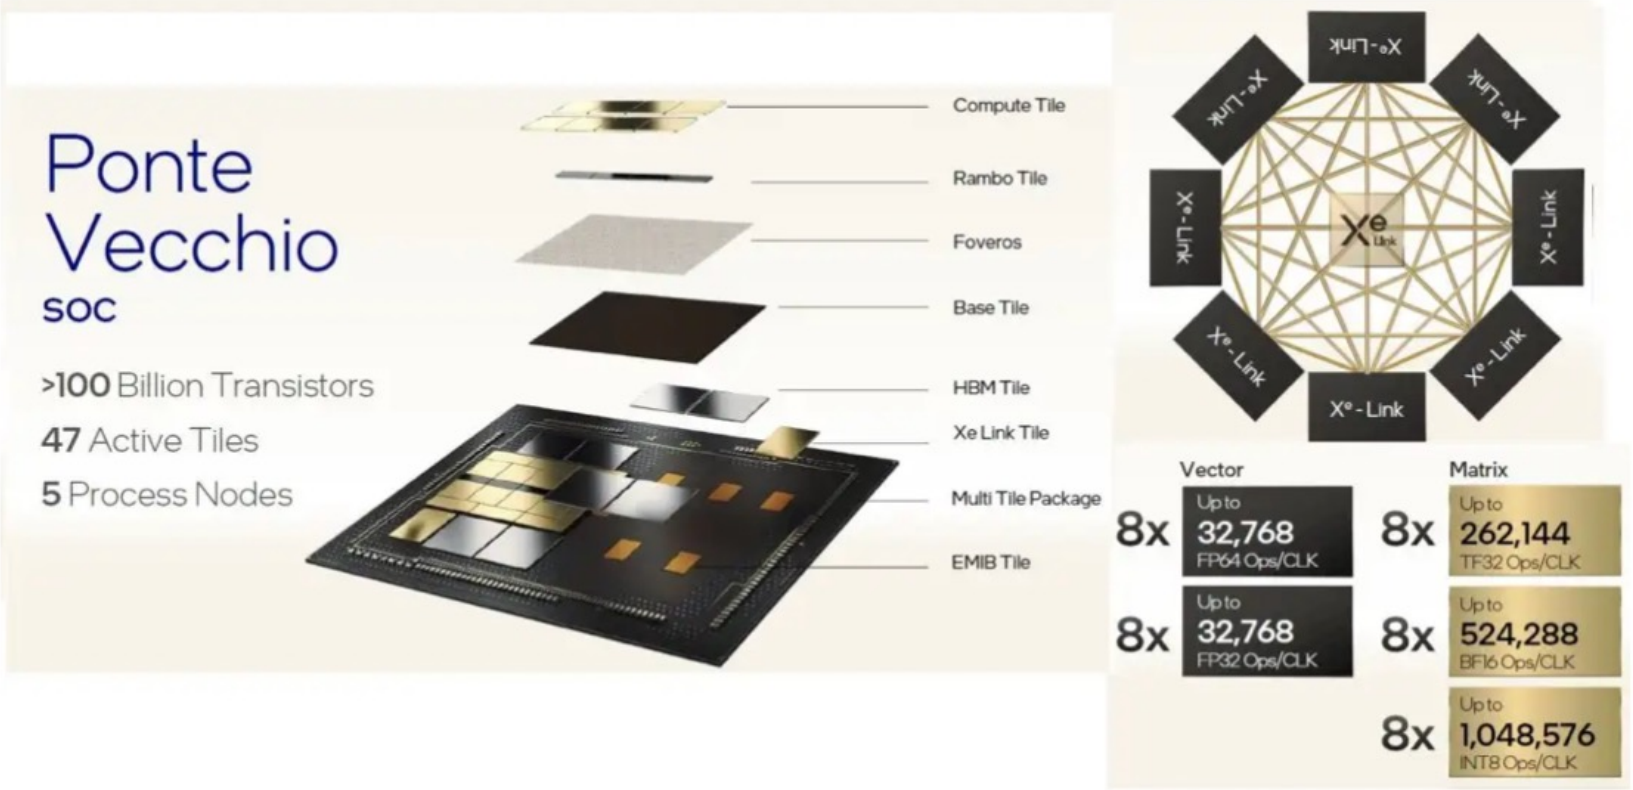

Intel Dev Con 2021

The (massive) package integrates various tiles, and utilizes both Intel's 2.5D EMIB interconnect bridges and 3D Foveros vertical stacked die. Of particular note is the constituent tiles are sourced from both TSMC (e.g., compute tile: TSMC N5) and Intel (e.g., base tile: Intel 7).

The X^e link tile (TSMC N7) enables direct connection of a variety of GPU topologies, as shown below.

Intel Dev Con 2021

Part 2: GPUs, IPUs, XeSS, OpenAPI

Key areas of focus are the:

- the definition of required hardware accelerator capabilities and services to interface with the software libraries
- acceleration of data de/compression
- optimization of the map-reduce framework (for faster database searches)
- optimization of the data storage footprint

For more information on OneAPI, please refer to: www.oneapi.com .

Summary

Although best known for their CPU offerings, Intel's breadth encompasses a much richer set of computing hardware and software products. At the recent Intel Architecture Day, they presented an aggressive roadmap for integrated, discrete, and (especially) data center GPUs, vying for leading performance across the full range of enthusiast/gamer and data center applications.

Intel Dev Con 2021

Part 2: GPUs, IPUs, XeSS, OpenAPI

A close collaboration with a major CSP promises to significantly upgrade the efficiency of cloud operations, replacing the SmartNIC with a richer set of functionality in the IPU.

The OneAPI initiative will no doubt lead to higher software development productivity across a myriad of CPU plus accelerator architectures.

The Ponte Vecchio GPU deserves special mention, as an example of the tradeoff decisions in building a complex GPU accelerator, integrating silicon tiles from both TSMC and Intel foundries with Intel's advanced packaging capabilities.

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP122

Intel Dev Con 2021

DR JEFF SOFTWARE
INDIE APP DEVELOPER

Part 2: GPUs, IPUs, XeSS, OpenAPI

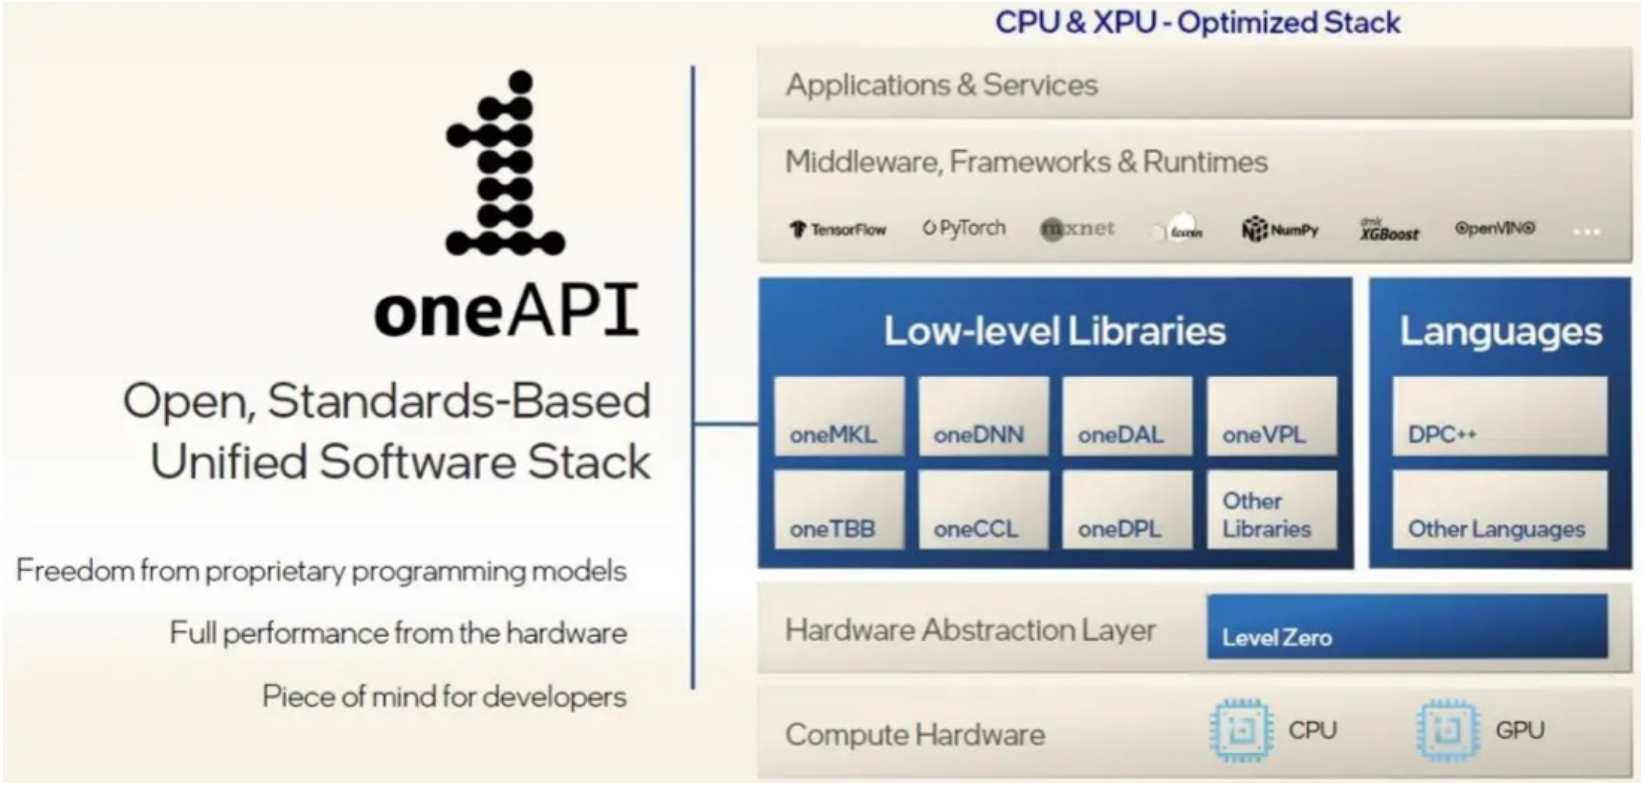

Intel Dev Con 2021

OneAPI

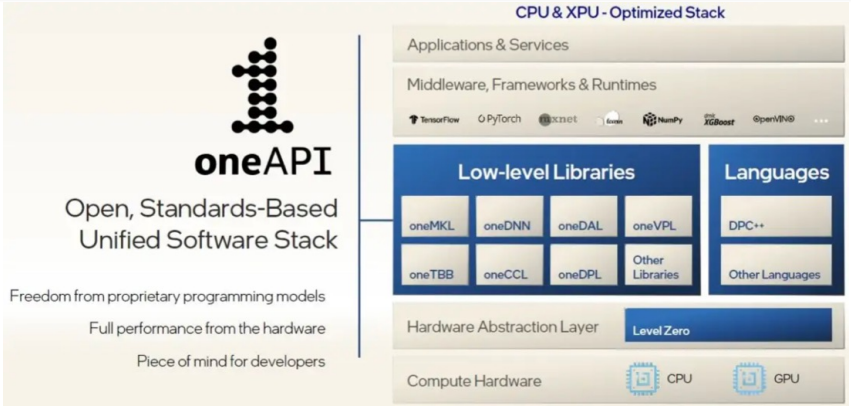

Briefly, Intel described their work on the industry-standard "OneAPI" software toolkit development, an effort to provide:

- a data-parallel software language (e.g., DPC++, especially for accelerators)
- an open S/W development stack for CPUs and XPUs (e.g., GPUs, accelerators)
- software library APIs for machine learning, video/media stream processing, matrix/vector math
- a full development toolkit (compiler, debugger, accelerator hardware interface models)

Intel Dev Con 2021

FPGA + IPU

Intel showed both an FPGA-based solution, and a new ASIC-based IPU named Mount Evans, as shown below. The cores in Mount Evans are based on the new Arm Neoverse (N1) architecture that is tightly coupled with the best-in-class packet processing pipeline and hardware accelerators.

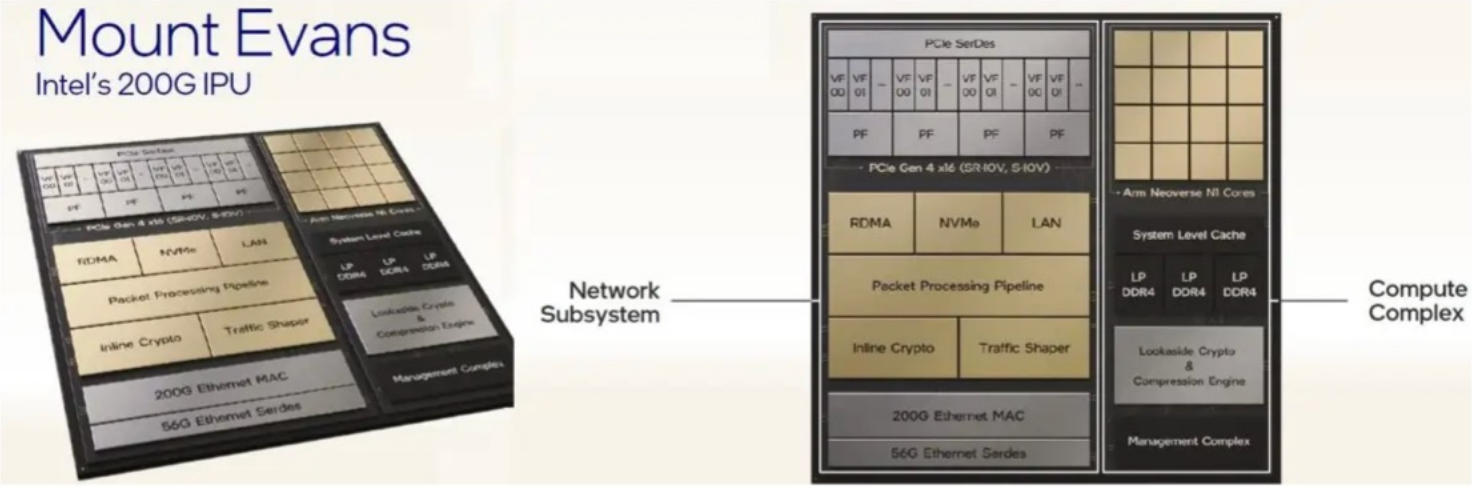

Intel Dev Con 2021

Summary

Intel recently described several new products based on p-core and e-core x86 architectural updates, to be shipped in the next few calendar quarters – Alder Lake (client) and Sapphire Rapids (data center). A new Advanced Matrix Extension (AMX) microarchitecture offers a unique opportunity to accelerate ML workloads, fitting a gap in the power/performance envelope between AVX instructions and dedicated GPU/accelerator hardware. The execution of multiples threads on asymmetric cores will benefit from the real-time interaction between the CPU and the (Windows-11) O/S scheduler.

These products also support new PCIe Gen5, DDR5, and the CXL1.1 protocol for unified memory space management across devices.

Section

Intel Architecture Day

Intel Architecture Day

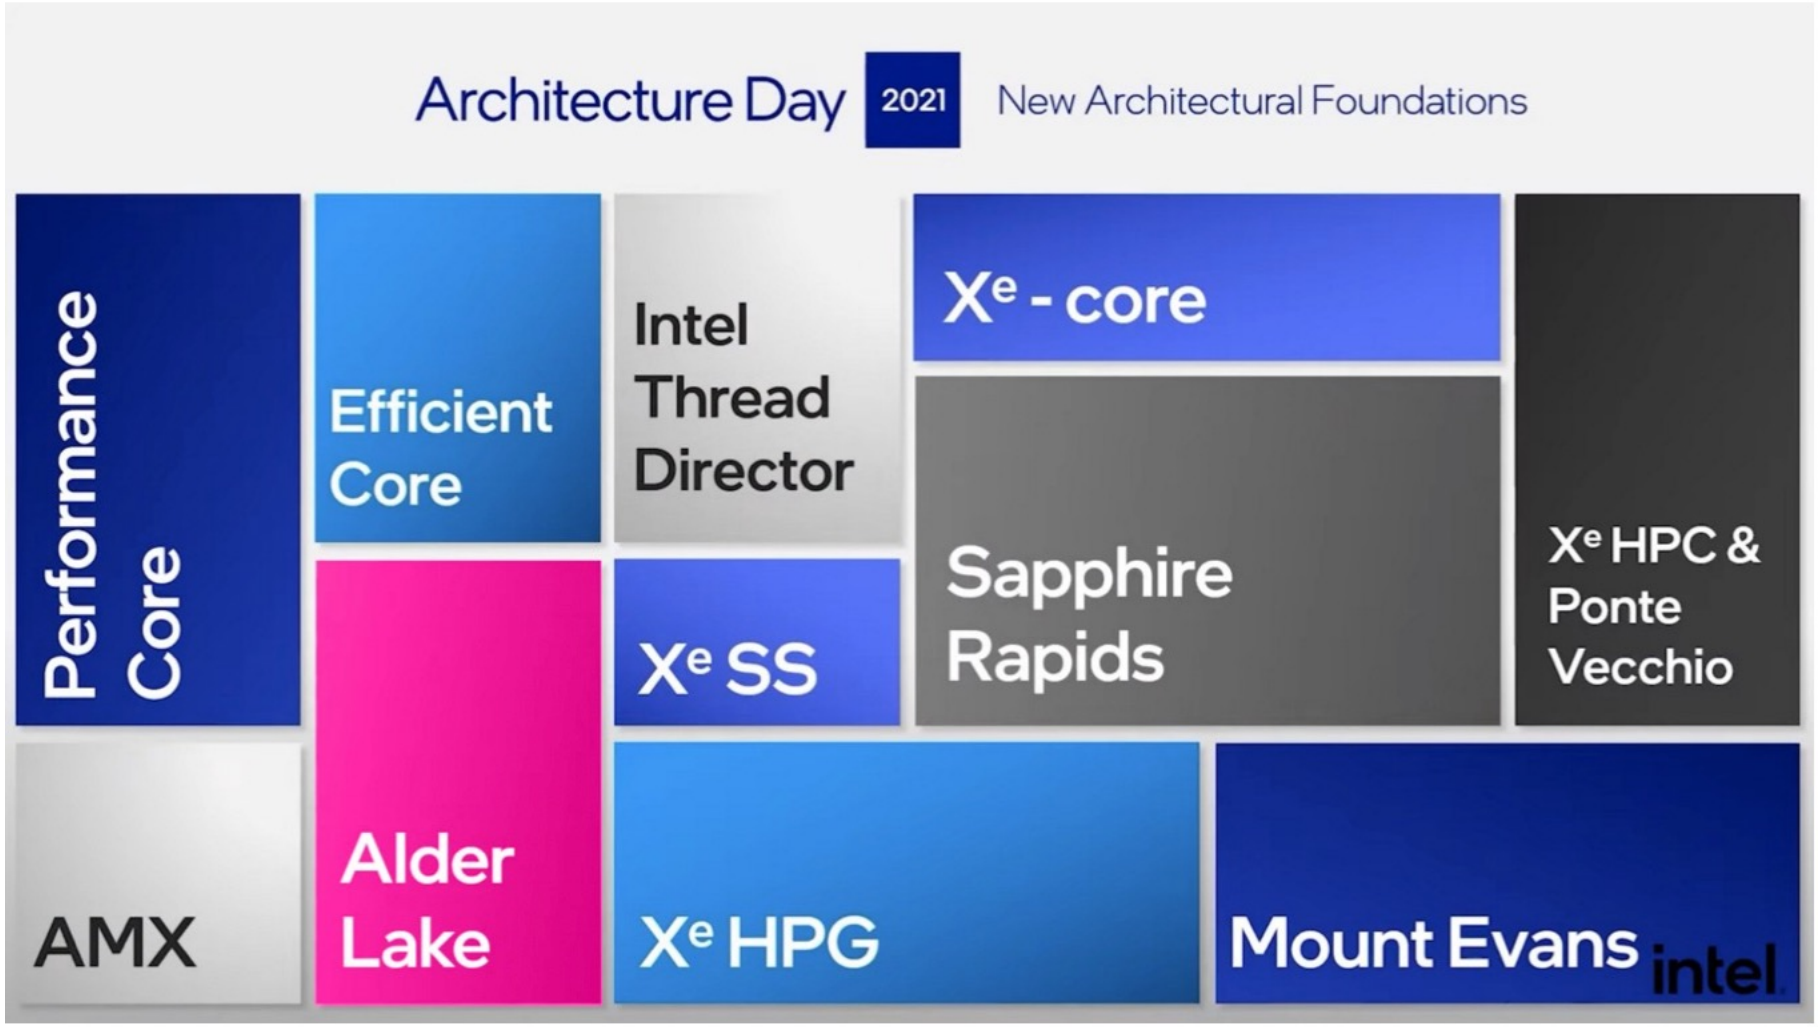

Source: Intel Architecture Day (highlight added by author)

Architecture Day 2021

After claiming that the Golden Cove represents Intel's largest architectural shift in over a decade, Yoaz went on to provide additional architecture details, including:

- To reduce latency the length decode was doubled to run at 32 bytes per cycle and 2 decoders were added (increasing from 4 to 6)
- Double the number of 4K pages stored in iTLB to better support software with a large code footprint
- A 2x+ branch target buffer that utilizes a machine learning algorithm for improved branch prediction and reduced jump mispredicts
- A wider, deeper, and smarter out-of-order engine
- Enhanced integer execution units with increased single-cycle operations
- Improved memory subsystem with reduced effective load latency and increased parallelism

Architecture Day 2021

Robinson went on to provide additional microarchitecture details, including:

- Improved branch prediction with a 5,000 entry branch target cache
- 64KB instruction cache
- Intel's first on-demand instruction length decoder, which enables bypassing the length decoder on subsequent execution, increasing performance while saving power
- Hardware-driven load balancing that extracts parallelism via a 256-entry out-of-order window and seventeen execution ports.
- A dual load + dual store memory subsystem

Architecture Day 2021

Section

News

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP122

Intel News

News Story

Print

Intel Partners with Microsoft against Cryptojacking

1:44 PM 4/26/2021 - MT Newswires

01:44 PM EDT, 04/26/2021 (MT Newswires) -- Intel (INTC) said Monday it has teamed up with Microsoft (MSFT) for the software company to expand its use of Intel technology to activate central processing unit based cryptomining machine learning detection.The technology equips endpoint detection and response solutions for advanced memory scanning, cryptojacking and ransomware detection.Price: 261.17, Change: +0.02, Percent Change: +0.01

What Is Intel® SGX?

There is tremendous opportunity for application and solution developers to take charge of their data security using new hardware-based controls for cloud and enterprise environments. Intel® Software Guard Extensions (Intel® SGX)[1 2] offers hardware-based memory encryption that isolates specific application code and data in memory. Intel® SGX allows user-level code to allocate private regions of memory, called enclaves, which are designed to be protected from processes running at higher privilege levels. Only Intel® SGX offers such a granular level of control and protection.

Intel News

Foundational Security

Ensuring a critical base of protection across the platform, focused on identity and integrity. Intel has a long history delivering technology to help ensure the platform comes up correctly and runs as expected. Our security engines have been used more than a billion times worldwide, and our processors feature enhanced cryptography to accelerate performance and help secure global commerce.

- Intel® Crypto Acceleration
- Firmware Update/Recovery
- Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI)
- Intel® BIOS Guard
- Intel® Boot Guard
- Intel® Converged Security and Management Engine (Intel® CSME)
- Intel® Platform Firmware Resilience (Intel® PFR)
- Intel® Platform Trust Technology (Intel® PTT)
- Intel® QuickAssist Technology (Intel® QAT)
- Intel® Runtime BIOS Resilience
- Intel® System Resources Defense
- Intel® System Security Report
- Intel® Total Memory Encryption (Intel® TME)
- Intel® Trusted Execution Technology (Intel® TXT)

Intel News

Semiconductors in Science and Industry

Semiconductor Industry News

Semiconductor company news:

- Intel news:
  - Intel agreed to sell its NAND flash memory-chip business that provides the gigabytes of storage used in smartphones, computers and other gadgets, to South Korea's SK Hynix for about $9B.
  - MarketWatch discussed how Intel lost "its Silicon Valley crown" and the company's recent mistakes, which have led to it falling behind "AMD in engineering process and Nvidia in market cap as a troubling decade takes its toll."
  - During the Next@Acer conference, Acer announced a number of new products, including the Acer Swift 3^X featuring 11^{th} Gen Intel Core processors and Iris X^e MAX discrete graphics. Press were excited to see how the new processors and graphics will perform and interested to see if the promised 17.5-hour battery life turns out to be true. Coverage appeared in AnandTech, CNET and others.
  - Google announced new solutions aimed to help organizations move toward more cloud-based technologies with Chrome OS, including Parallels Desktop for Chromebook Enterprise, which is only available on select Chromebooks powered by Intel Core Processors. This new technology provides users quick access to legacy and full-featured Windows applications, locally on Chrome OS, without needing to be connected to the internet. Coverage appeared in TechCrunch, The Verge and others.
  - AnandTech and others reported that Intel's DG1 GPU is now shipping. Intel also used its earnings to announce that its next GPU, the DG2—based on the upcoming Xe-HPG architecture—is currently powered on in Intel labs.
  - ASUS announced the VivoBook Flip 14, revealing that this will be the first laptop to feature Intel Xe Max discrete graphics.
  - Moore's Law Is Dead published a video on Intel's alleged Xe HPG GPU launching in 2021, noting efficiency is likely comparable to RDNA 2 and it will launch within 4-6 months of RDNA 3. The card is based on an enhanced version of Tiger Lake's GPU and will be one 512 EU Tile, verses multiple titles, as many initially predicted. Finally, it will utilize 16GB of GDDR6 VRAM over a 256-bit bus.

Intel News

Sep 20, 2020

- Following the embargo lift on reviews of reference design laptops containing Intel's 11th Gen Core i7-1185G7 mobile processor, sentiment was positive overall, with press highlighting big gains in single-threaded performance and positive gaming experiences with the new Intel Xe integrated graphics. Coverage appeared in AnandTech, Digital Trends, Engadget, Gizmodo, PCWorld and others.

8-core CPU

- A handful of outlets, including PCWorld, reported on a recent blog post by Intel's Corporate Vice President of Client Computing Group, Boyd Phelps, revealing that the Tiger Lake lineup will also feature an 8-core variant in addition to quad-core processors

960-EU GPU

- VideoCardz picked up rumors from an unidentified source that Intel will potentially release a 960 EU (Execution Units) chip with Xe-HPG architecture for Intel's DG2 series. The outlet noted it would be unlikely for Intel to mass produce the GPU at a 960 EU option, given Intel is already planning to offer 128, 384 and 512 EU's.
  - VideoCardz also shared Intel's DG2 series has been listed with 8GB GDDR6 memory within a spec photo of an undetermined device, per a Twitter user.
- Intel selected Austin, TX as the first U.S. site for the expansion of its Intel Ignite acceleration program for early stage startup companies.

Intel

Semiconductors in Science and Industry

Semiconductor Industry News

- Intel news:
  - ServeTheHome reported that Trenton Systems, a system integrator, posted photos and block diagrams of how they intend to use Intel's Ice Lake processors for their server platforms. Photos revealed a 3U platform "that is prioritizing size and PCIe expansion."
  - AnandTech covered new 10nm Intel Atom Embedded processors, codenamed "Elkhart Lake." The new processors are built with Tremont Atom cores and will come as three series of processors: Pentium, Celeron and Atom x6000E.
  - In a LinkedIn post, William Magro revealed he was leaving his position at Intel as fellow and chief technologist in HPC to assume the same role with Google Cloud.
  - Intel posted a blog by Technology Evangelist, Marcus Yam, announcing the company will be bringing 11th Generation Intel Core "Tiger Lake" processors with Iris Xe graphics to Chromebooks. Coverage appeared in Android Central, HotHardware, PCWorld, TechRadar, ZDNet and others.
- Nvidia news:
  - Wccftech, VideoCardz, eTeknix and others reported on crashes and black screens experienced by new RTX 3080 users, likely a result of overclocked boost speeds above 2.0GHz.
  - The Wall Street Journal's Christopher Mims wrote a piece titled, "Huang's Law Is the New Moore's Law, and Explains Why Nvidia Wants Arm." The newly coined Huang's Law "describes how the silicon chips that power artificial intelligence more than double in performance every two years." ExtremeTech responded saying, "there's no such thing" and explained why the definition of the rule does not work.