



# CPU Hardware Architecture and Performance Optimization

G. Amadio (CERN)

# Our Main Goals

- Understand the architecture of modern CPU hardware
  - Hardware evolution
  - Main features of modern hardware
- Understand how to analyze the performance of our code
  - How to identify performance bottlenecks
  - What to measure and how to measure it
- Combine architectural knowledge and performance analysis
  - How to interpret performance measurements
  - What changes to make to the software

# CPU Hardware Architecture and Evolution


# Early Computing Devices



2700 – 2300 BC

## Abacus


Used since ancient times, until Arabic numerals became the norm. Still in use as an educational tool.



1620– 1630

## Slide Rule


Uses logarithm scales to help with multiplications and also computing other functions. Extensively used by engineers in the last century, before computers became powerful.



1642

## Pascaline

Mechanical calculator invented by Blaise Pascal to help his father with tax calculations. Could add and subtract.



1820s

## Difference Engine


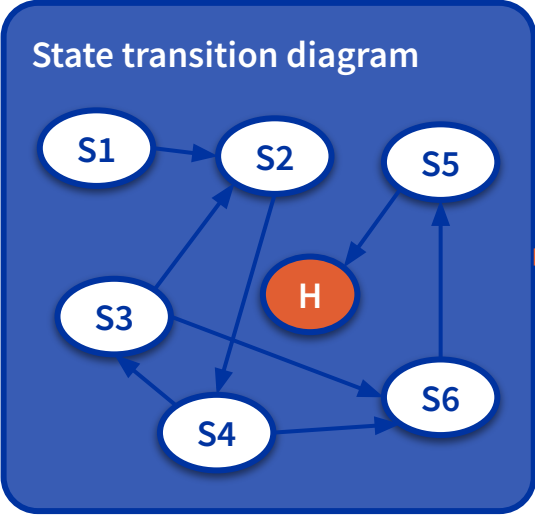
Automatic mechanical calculator designed to tabulate polynomial functions. Designed and first created by Charles Babbage.

# Ada Lovelace, the first computer programmer

**Augusta Ada King, Countess of Lovelace** (10 December 1815 – 27 November 1852) was an English mathematician and writer, known for her work on Charles Babbage's proposed mechanical general-purpose computer, the Analytical Engine. She was the first to recognize that the new machine had applications beyond simple calculations. She arguably wrote the first “computer program”. In her article entitled “note G” on the Analytical Engine, she described in detail an algorithm to compute a sequence of **Bernoulli numbers** using it.



# The Turing Machine: concept of first generic computer



A Turing machine is capable of computing any computable sequence.

HEAD: READ/WRITE/MOVE



Infinite tape

# From Turing Machine to Stored-Program Computer

## Turing Machine


Conceptually the first general computing machine.

## ABC, Colossus, ENIAC

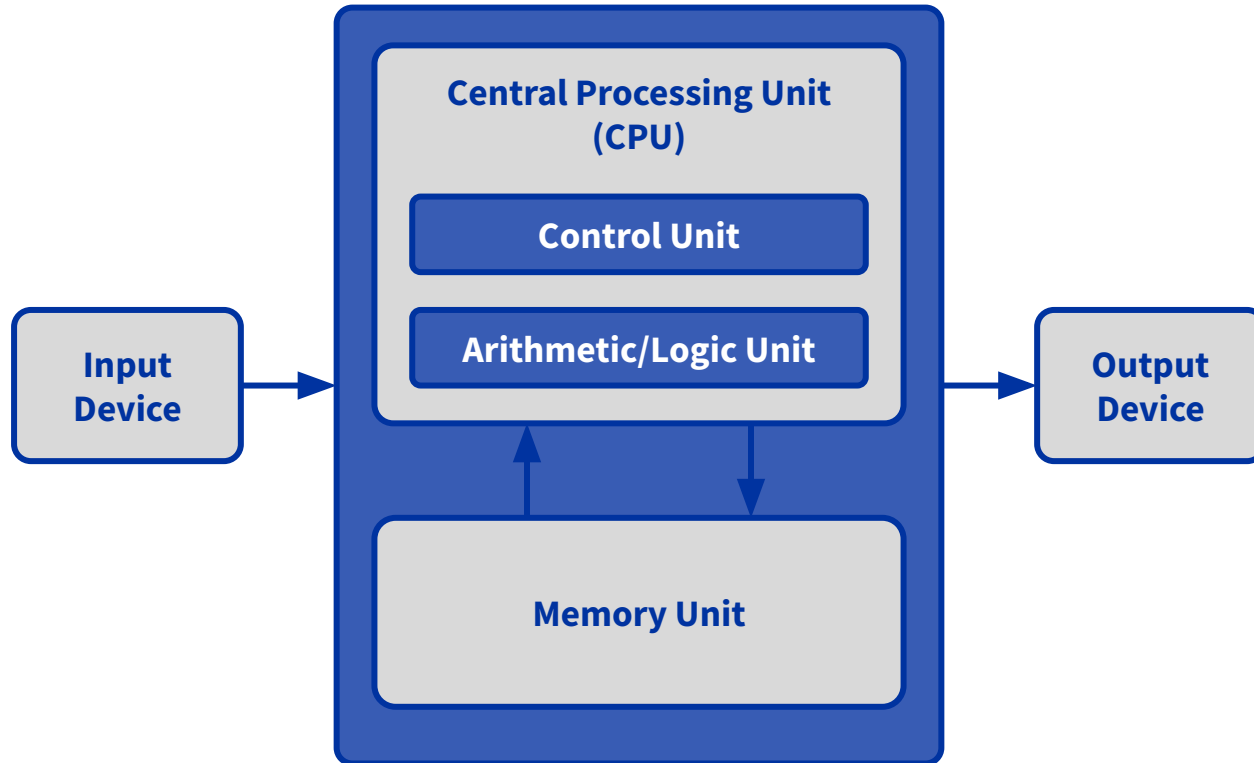
First truly digital computers, based on boolean logic and vacuum tubes.

## Assembly Language

Beginning of standardization of how to program computer with abstract instruction sets.




# John von Neumann Architecture





# Transistor count



Data source: Wikipedia ([wikipedia.org/wiki/Transistor\\_count](https://wikipedia.org/wiki/Transistor_count))


OurWorldinData.org – Research and data to make progress against the world's largest problems.

Licensed under CC-BY by the authors Hannah Ritchie and Max Roser.

# Integrated Circuit-Based Microprocessors

Intel 4004


First Intel microprocessor. 2,300 Transistors.



1972

Intel 8008

First 8 bit microprocessor. 3,500 Transistors.



MOS 6502

Powered many popular devices, such as the Apple II, Atari 2600, Commodore 64, and the NES. 3,510 Transistors.




1975

1976

Zilog Z80

8-bit microprocessor. Powered devices such as Sega Master System and Mega Drive, and Sinclair's ZX Spectrum.



1978

Intel 8086


First 16-bit microprocessor. 29,000 Transistors. Its successor, the Intel 8088, a slightly modified version, powered first IBM PC.



1985

Intel 80386

32-bit microprocessor. 275,000 Transistors, 33MHz. Cemented Intel's PC market dominance.



# Intel's 8086 Registers and Assembly

1 9 1 8 1 7 1 6 1 5 1 4 1 3 1 2 1 1 1 0 0 9 0 8 0 7 0 6 0 5 0 4 0 3 0 2 0 1 0 0 (bit position)

## Main registers

	AH	AL	<b>AX</b> (primary accumulator)
0000	BH	BL	<b>BX</b> (base, accumulator)
	CH	CL	<b>CX</b> (counter, accumulator)
	DH	DL	<b>DX</b> (accumulator, extended acc)

## Index registers

0000	SI	<b>Source Index</b>
0000	DI	<b>Destination Index</b>
0000	BP	<b>Base Pointer</b>
0000	SP	<b>Stack Pointer</b>

## Program counter

0000	IP	<b>Instruction Pointer</b>
------	----	----------------------------

## Segment registers

CS	0000	<b>Code Segment</b>
DS	0000	<b>Data Segment</b>
ES	0000	<b>Extra Segment</b>
SS	0000	<b>Stack Segment</b>

## Status register

- - - - O D I T S Z - A - P - C Flags

```

; _strtolower:
; Copy a null-terminated ASCII string, converting
; all alphabetic characters to lower case.
; ES=DS
; Entry stack parameters
; [SP+4] = src, Address of source string
; [SP+2] = dst, Address of target string
; [SP+0] = Return address
;
_strtolower proc
    push    bp                ;Set up the call frame
    mov     bp,sp
    push    si
    push    di
    mov     si,[bp+6]        ;Set SI = src (+2 due to push bp)
    mov     di,[bp+4]        ;Set DI = dst
    cld                    ;string direction ascending

loop:    lodsb                ;Load AL from [si], inc si
        cmp     al,'A'        ;If AL < 'A',
        jl     copy          ;Skip conversion
        cmp     al,'Z'        ;If AL > 'Z',
        jg     copy          ;Skip conversion
        add     al,'a'-'A'    ;Convert AL to lowercase
copy:    stosb                ;Store AL to [di], inc di
        or     al,al         ;If AL <= 0,
        jne    loop         ;Repeat the loop

done:    pop     di            ; restore di and si
        pop     si
        pop     bp            ;Restore the prev call frame
        ret                 ;Return to caller
end proc
    
```

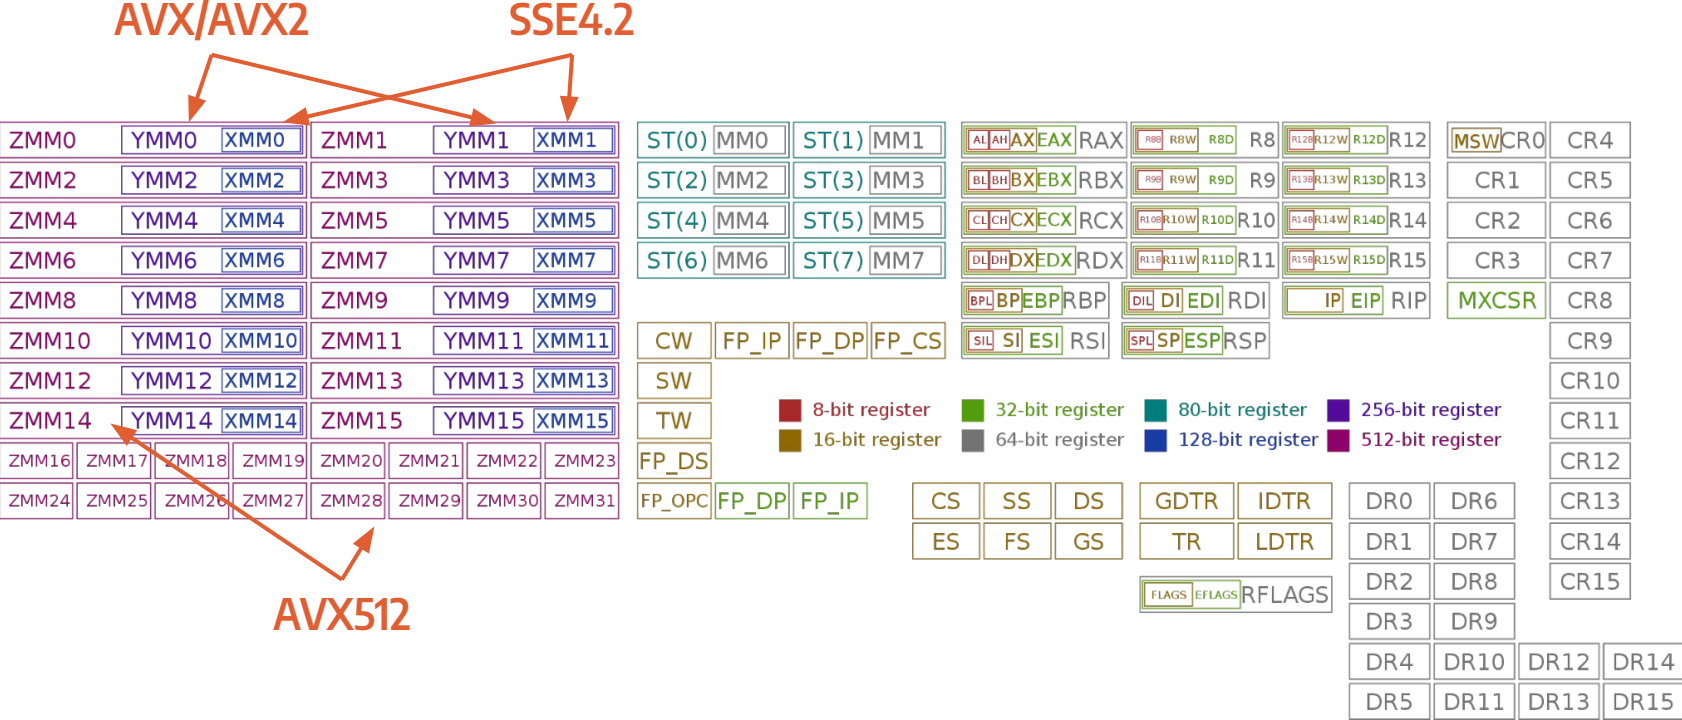
Source: Wikipedia

# Intel x86 Assembly

```
1  __attribute__((noinline))
2  int is_odd(unsigned long long n)
3  {
4      return n & 1;
5  }
6
7  unsigned int collatz_count(unsigned long long n)
8  {
9      unsigned int count = 0;
10
11     while (n != 1)
12     {
13         if (is_odd(n))
14             n = 3 * n + 1;
15         else
16             n = n / 2;
17
18         ++count;
19     }
20
21     return count;
22 }
```

```
1  is_odd:
2      mov     eax, edi
3      and     eax, 1
4      ret
5  collatz_count:
6      xor     edx, edx
7      cmp     rdi, 1
8      jne    .L7
9      jmp    .L3
10 .L10:
11     lea    rdi, [rdi+1+rdi*2]
12     add    edx, 1
13     cmp     rdi, 1
14     je     .L3
15 .L7:
16     call   is_odd
17     test   eax, eax
18     jne    .L10
19     shr    rdi
20     add    edx, 1
21     cmp     rdi, 1
22     jne    .L7
23 .L3:
24     mov     eax, edx
25     ret
```

# Registers available in the x86-64 instruction set



# Instruction Sets

- **CISC (Complex Instruction Set Computer)**
  - Intel x86 and AMD64
    - Most laptop and desktop PCs, Playstation 5, Xbox One
  - IBM System z (mainframe computers)
- **RISC (Reduced Instruction Set Computer)**
  - ARM
    - Amazon Graviton (AWS VMs)
    - Apple M1–M4 (iPhone, iPad, iMacs)
    - Ampere Altra, Fujitsu A64FX, etc
    - Qualcomm (mobile phones, tablets)
    - Nintendo Game Boy Advance, DS, 3DS and Switch, Raspberry Pi, etc
  - IBM's PowerPC
    - Apple Macintosh (1994–2005), Nintendo GameCube and Wii, Playstation 3, Xbox 360
  - DEC Alpha, MIPS, Motorola 68000, RISC-V, SPARC, SuperH
    - Apple II (M68k), Nintendo 64, PlayStation 1 and 2 (MIPS), Sega Saturn and Dreamcast (SuperH)

# Programming Language Evolution

1947  
Assembly

Low-level language.  
High correspondence between language and hardware instructions.

Code written in assembly is converted to machine code using an *assembler*, which was a big upgrade over previous forms of programming.

1954  
Fortran

One of the earliest high-level imperative programming languages.

Introduced procedural programming, double precision, and complex numbers. Still popular in HPC, including in GPU application programming.

1963  
BASIC

**B**eginner's **A**ll-purpose **S**ymbolic **I**nstruction **C**ode is a family of high-level languages.

BASIC became popular during the 8-bit era, but declined in popularity in the 90s, when more advanced languages like C were the norm.

1972  
C

Originally developed to implement many of the utilities for UNIX OSs. Still in wide use today.


C is a portable, imperative language with a static type system and which supports structured programming.

1979  
C++

C++ was designed with systems programming, embedded software, and efficiency in mind.

Although many think of C++ as a superset of C or C with classes, their latest versions are not fully compatible. Used extensively in HEP and HPC nowadays.


# 50 Years of Microprocessor Trend Data



source: <https://github.com/karlrupp/microprocessor-trend-data>



# 50 Years of Microprocessor Trend Data



source: <https://github.com/karlrupp/microprocessor-trend-data>

# Breakdown of Dennard's Scaling

- Power density per unit area stopped decreasing
- Frequency could no longer keep increasing after each die shrink
  - But the transistor numbers kept growing
- Single-thread performance gains continued, albeit at a slower pace
  - More complexity: pipelining, superscalar, out-of-order execution, SIMD
- AMD and Intel bring 64-bit CPUs to the mainstream market
  - Intel with IA-64, and AMD with amd64 (x86\_64), announced in 1999
- From symmetric multiprocessing (SMP) to multithreading (SMT)
  - In the '90s, dual socket high-end servers became popular
  - First SMT capable CPU was the Intel Pentium 4, released in 2002
- First dual core processors began to appear in mid 2000s
- The era of parallelism is born

# Instruction-Level Parallelism

## Pipelining



## Operation-Level Parallelism


## Superscalar Execution



Time →

# Symmetric multithreading (SMT)


Without SMT




Threads scheduled one at a time on each physical core



With SMT



Threads run simultaneously on two logical cores



# CPU Architecture

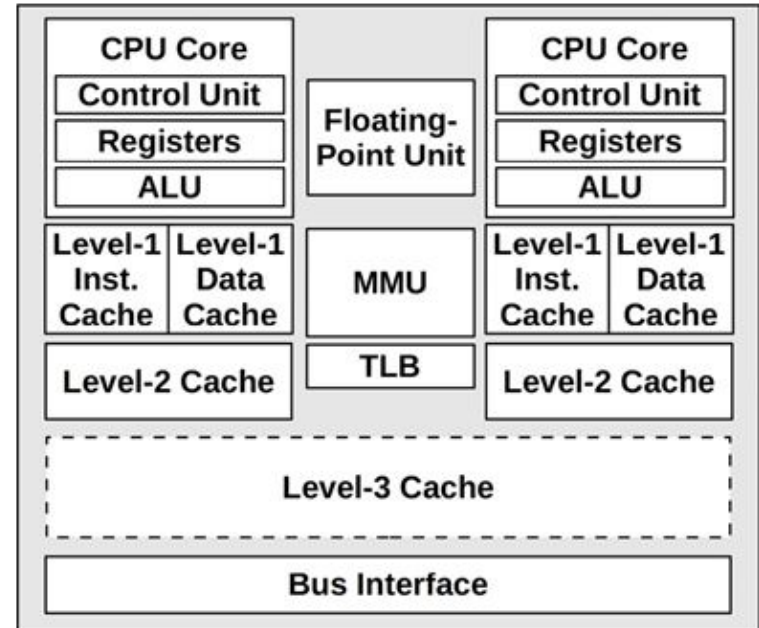
- Logical Components

- Control Unit
- Arithmetic Logic Units (ALUs)
- Floating Point Unit (FPU)
- Branch Predictor Unit (BPU)
- Memory Management Unit (MMU)
- Translation Lookaside Buffer (TLB)

- Memory Subsystem

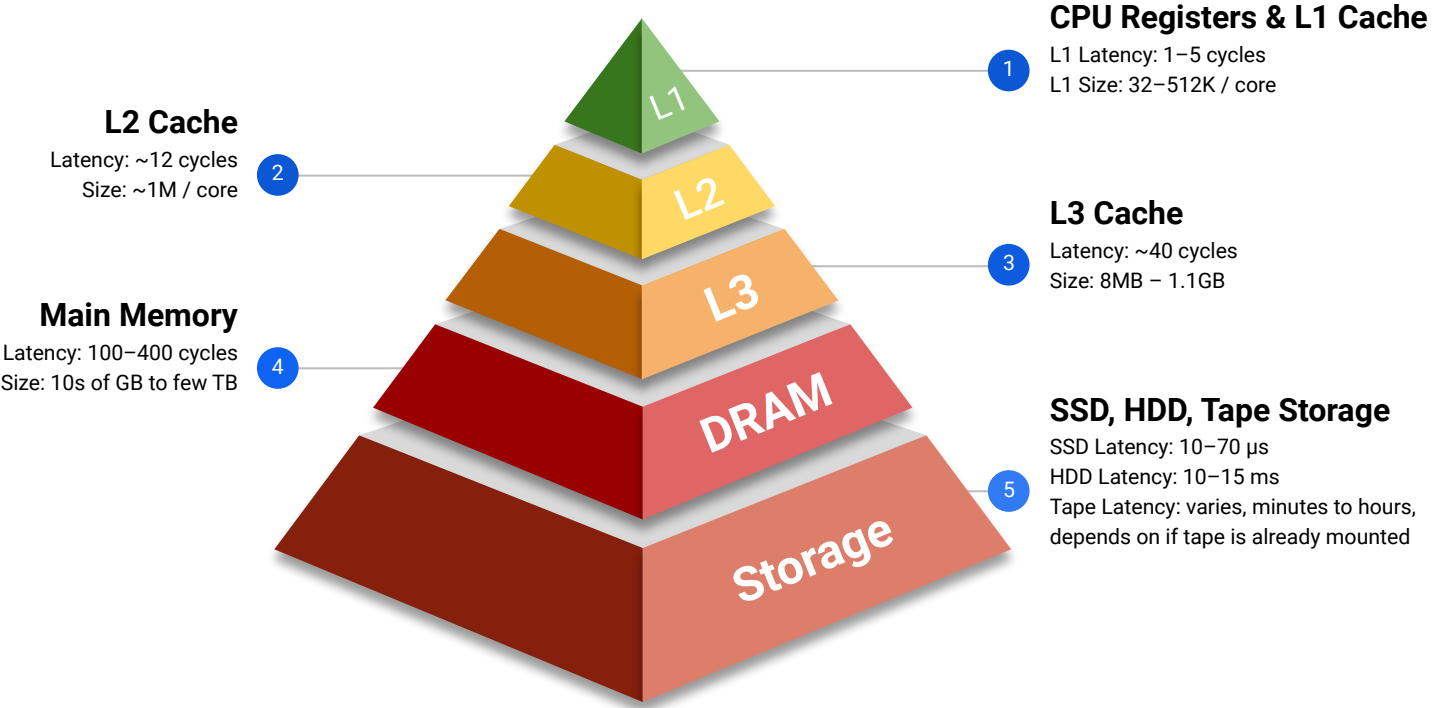
- L1 (~32–512KB per core)
  - L1 Instruction Cache
  - L1 Data Cache
- L2 (~1–8MB per core)
  - Instruction/Data Shared Cache
- L3 (up to ~8MB–1.1GB per socket)
  - Last level cache (LLC)

## Generic Dual Core CPU

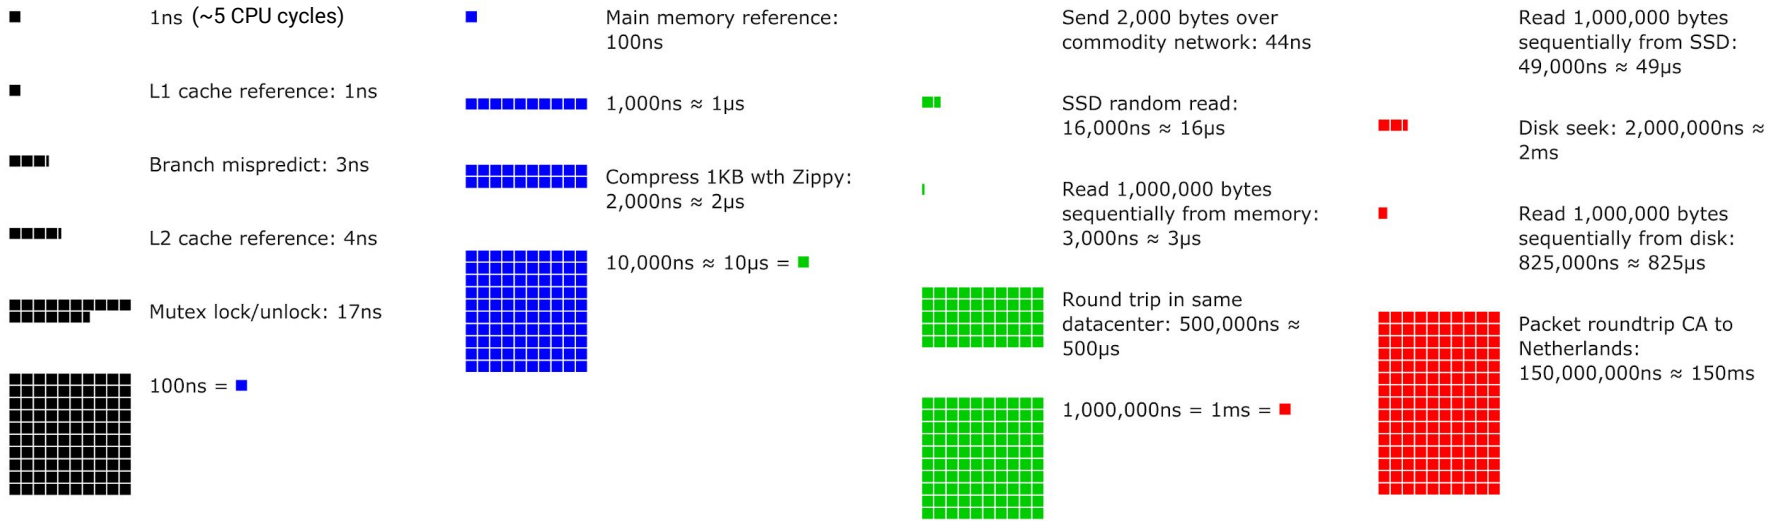


Source: *Systems Performance 2nd Edition*, Brendan Gregg

# Memory Hierarchy



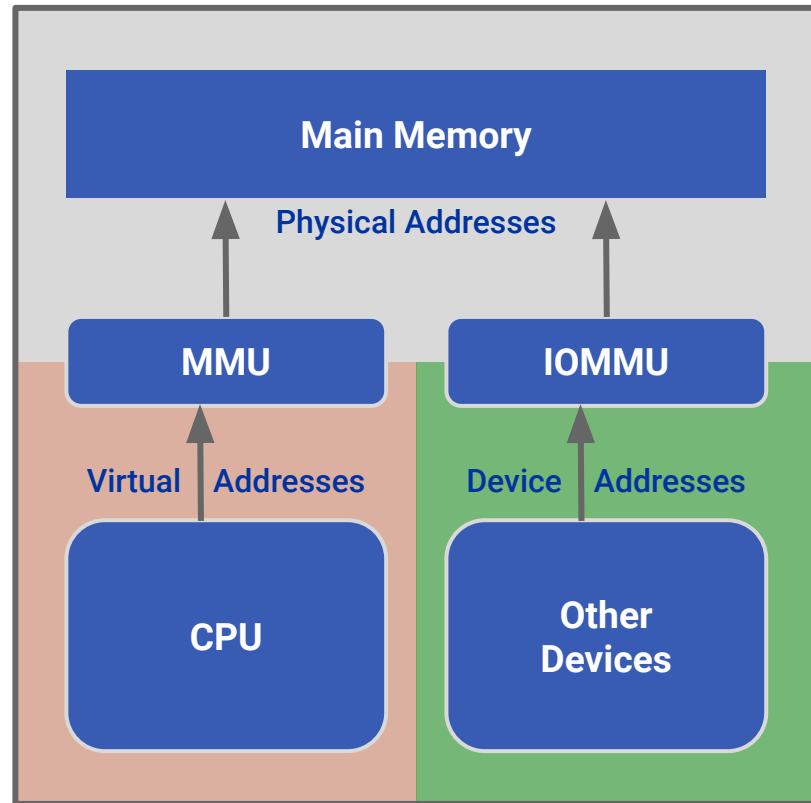
# Latency Numbers Every Programmer Should Know



Source: [https://colin-scott.github.io/personal\\_website/research/interactive\\_latency.html](https://colin-scott.github.io/personal_website/research/interactive_latency.html)

# Virtual Memory

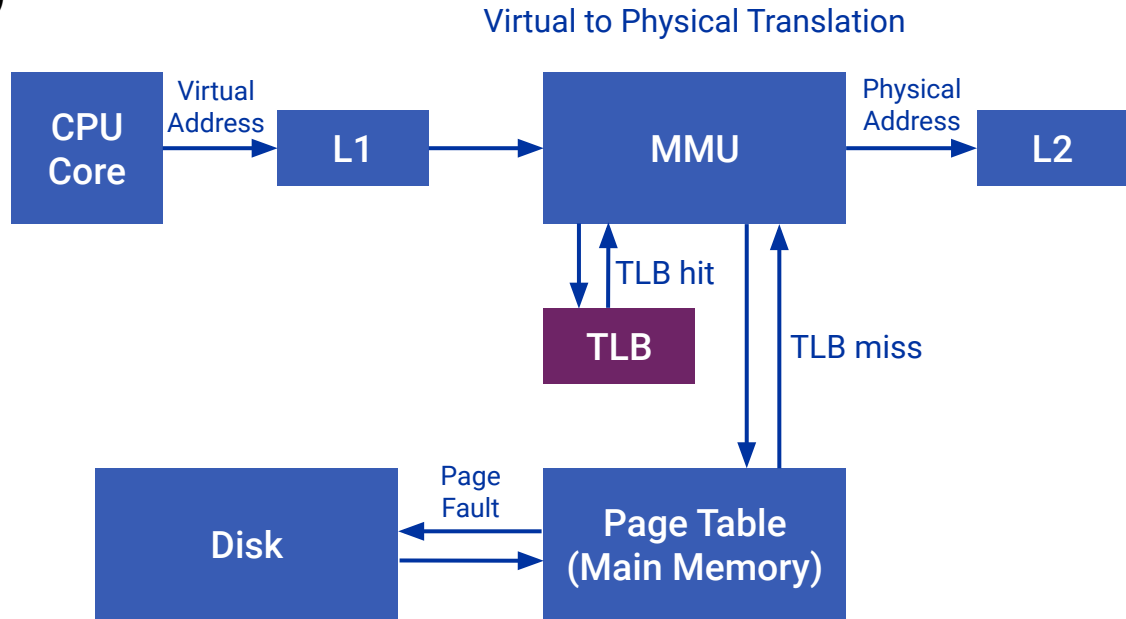
- First appeared in the [Atlas computer](#) in 1962
- Memory Management Unit (MMU)
- Memory managed in pages
  - Page sizes are usually 4K, 16K, 64K
  - May also support “huge pages” of 2MB, 1GB
- Hides fragmentation of physical memory
- Memory hierarchy managed by the kernel
- Makes application programming easier
  - Memory looks contiguous
  - No need to worry about fragmentation
  - Seems to own whole address space
  - Enabled timesharing features






# The Translation Lookaside Buffer

“A translation lookaside buffer (TLB) is a memory cache that is used to reduce the time taken to access a user memory location. It is a part of the chip's memory management unit (MMU). The TLB stores the recent translations of virtual memory to physical memory and can be called an address-translation cache.”




# The Parallel Era


- First dual core CPUs debut in 2004
  - Pentium D, based on Pentium 4
  - AMD Athlon X2
- Quickly evolved from 2 to 4 cores
  - Stagnated at 4 cores for several years
- Ryzen brought AMD back in the game
  - Offered more cores, forced Intel to do the same
- ARM finally begins move from phones to servers
  - Amazon Graviton, Fujitsu A64FX, Ampere Altra
- Innovations in packaging led to multi-chip CPUs
  - AMD EPYC, Intel Sapphire and Emerald Rapids



AMD EPYC Bergamo




Fujitsu A64FX



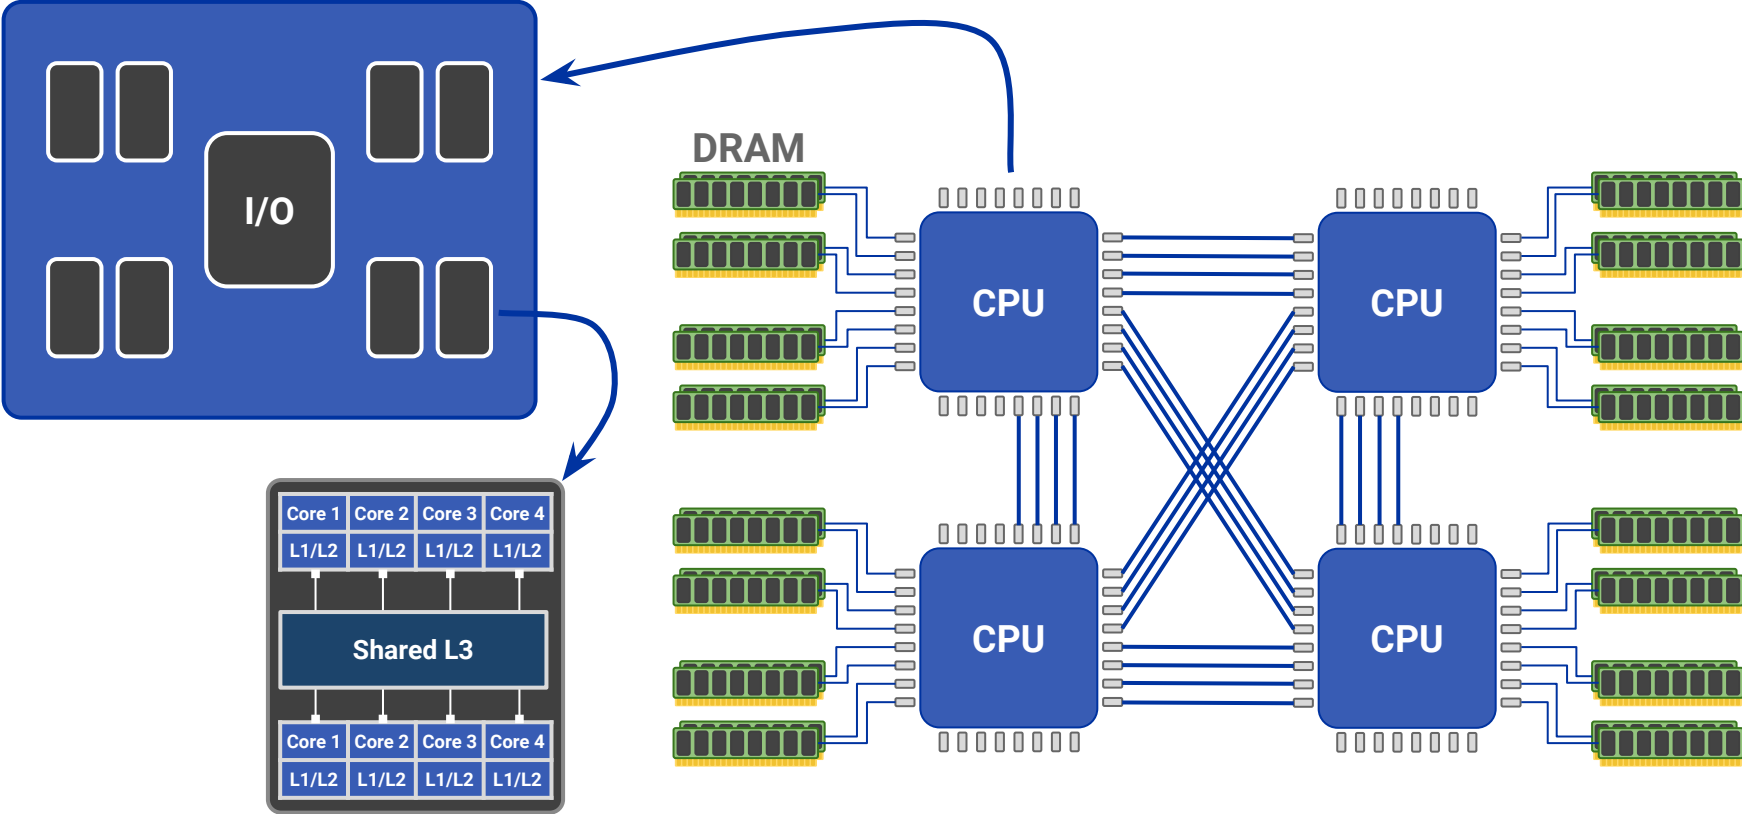
Intel Emerald Rapids

# Ampere Roadmap 2020 – 2026



Source: Ampere

# Non-Uniform Memory Architecture (NUMA)



	AMD EPYC 7001 'NAPLES'	AMD EPYC 7002 'ROME'	AMD EPYC 7003 'MILAN'	AMD EPYC 9004, 8004 'GENOA', 'SIENA'
				
<b>Core Architecture</b>	'Zen'	'Zen 2'	'Zen 3'	'Zen 4' and 'Zen 4c'
<b>Cores</b>	8 to 32	8 to 64	8 to 64	8 to 128
<b>IPC Improvement Over Prior Generation</b>	N/A	~24% <sup>ROM-236</sup>	~19% <sup>MLN-003</sup>	~14% <sup>EPYC-038</sup>
<b>Max L3 Cache</b>	Up to 64 MB	Up to 256 MB	Up to 256 MB	Up to 384 MB (EPYC 9004) Up to 128 MB (EPYC 8004)
<b>Max L3 Cache with 3D V-Cache™ technology</b>			768 MB	Up to 1152 MB
<b>PCIe® Lanes</b>	Up to 128 Gen 3	Up to 128 Gen 3	Up to 128 Gen 4	Up to 128 Gen 5 8 bonus lanes Gen 3
<b>CPU Process Technology</b>	14nm	7nm	7nm	5nm
<b>I/O Die Process Technology</b>	N/A	14nm	14nm	6nm
<b>Power (Configurable TDP [cTDP])</b>	120-200W	120-280W	155-280W	70-400W
<b>Max Memory Capacity</b>	2 TB DDR3-2400/2666	4 TB DDR4-3200	4 TB DDR4-3200	6 TB DDR5-4800

Source: AMD

# AMD Zen4 Architecture in detail

## COMPUTE

- AMD "Zen4" x86 cores (Up to 12 CCDs / 96 cores / 192 threads)
- 1MB L2/Core, 96MB L3/CCD / Total up to 1,152MB L3
- ISA updates: BFLOAT16, VNNI, AVX-512 (256b data path)
- Memory addressability with 57b/52b Virtual/Physical Address
- Updated IOD and internal AMD Gen3 Infinity Fabric™ architecture with increased die-to-die bandwidth

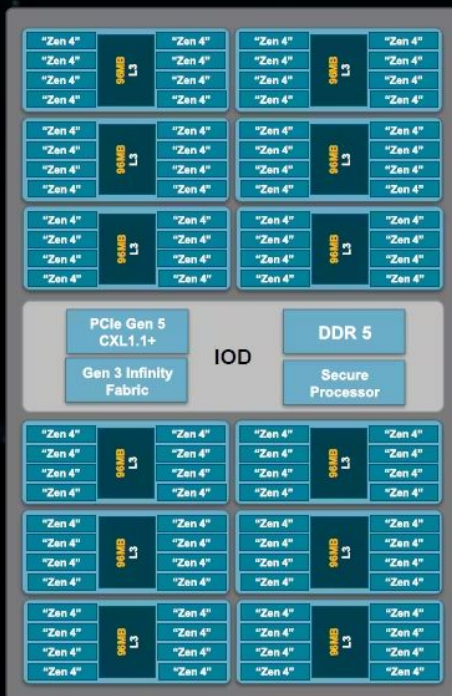
• Target TDP range: Up to 400W (cTDP)

• Updated RAS

## Memory

- 12 channel DDR5 with ECC up to 4800 MHz
- Option for 2, 4, 6, 8, 10, 12 channel memory interleaving<sup>1</sup>
- RDIMM, 3DS RDIMM
- Up to 2 DIMMs/channel capacity with up to 12TB in a 2 socket system (256GB 3DS RDIMMs)<sup>1</sup>

Source: AMD



ORANGE indicates difference from General Purpose

## SP5 Platform

- New socket, increased power delivery and VR
- Up to 4 links of Gen3 AMD Infinity Fabric™ with speeds of up to 32Gbps
- Flexible topology options
- Server Controller Hub (USB, UART, SPI, I2C, etc.)

## Integrated I/O – No Chipset

Up to 160 IO lanes (2P) of PCIe® Gen5

- Speeds up to 32Gbps, bifurcations supported down to x1
- Up to 12 bonus PCIe® Gen3 lanes in 2P config (8 lanes-1P)
- Up to 32 IO lanes for SATA

- 64 IO Lanes support for CXL1.1+ w/bifurcations supported down to x4

## Security Features

Dedicated Security Subsystem with enhancements

Secure Boot, Hardware Root-of-Trust

SME (Secure Memory Encryption)


SEV-ES (Secure Encrypted Virtualization & Register Encryption)

SEV-SNP (Secure Nested Paging), AES-256-XTS with more encrypted VMs




# Intel® Xeon® Die Package Enhancements

Scalable, Balanced Architecture



4th Gen Intel Xeon  
(Four-Tile Architecture)



5th Gen Intel Xeon  
(Two-Tile Architecture)

intel

# 5th Gen Intel® Xeon® Processors Turbo Frequencies

Introducing Improved 5 Turbo Ratio Levels

4th Gen Intel® Xeon® CPU

Instruction Class	Cdyn Class			
	0	1	2	3
SSE	128 Light	128 Heavy		
AVX2	256 Light	256 Moderate	256 Heavy	
AVX512	512 Ultra-Light	512 Light	512 Moderate	512 Heavy
AMX		AMX Light	AMX Moderate	AMX Heavy
Turbo Frequency	SSE	AVX2	AVX512	AMX

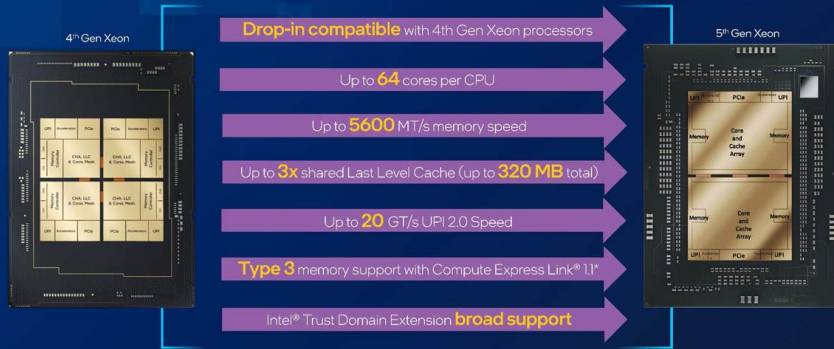
5th Gen Intel® Xeon® CPU

Instruction Class	Cdyn Class				
	0	1	2	3	4
SSE	128 Light	128 Heavy			
AVX2	256 Light	256 Moderate	256 Heavy		
AVX512	512 Ultra-Light	512 Light	512 Moderate	512 Heavy	
AMX	<b>ENHANCED</b>	AMX Ultra-Light	AMX Light	AMX Moderate	AMX Heavy
Turbo Frequency	SSE	AVX2	AVX512	AVX512 <b>NEW</b>	AMX

- Improves Turbo Frequencies for Intel® AVX heavy and Intel® AMX light workloads including HPC and AI
- ~2 bins Turbo Frequency Upside on Intel® AVX-512 Heavy usage
- ~9% performance improvement on low load (4T or 8T) LINPACK AVX512
- ~5% performance improvement on low instance (4 or 32) Resnet50 amx\_int8, amx\_bfloat16 and avx\_fp32
- Lowers the Turbo frequency penalty for using AVX512 or AMX, broadening usability of these instruction sets

# 5th Gen Intel® Xeon® Processors

Platform enhancements delivering significant gains at the same power envelope<sup>1</sup>




intel

<sup>1</sup> Type 3 CXL, targeted deployment. See backup for configuration details. Results may vary.

# Compute Express Link® 1.1 Enhancements

Type 3 memory support with 5th Gen Intel® Xeon® processors



## 2-Tier Memory Support

Type 3 Memory Expansion Devices:  
Capacity Expansion

- Tier 1 memory = native DDR, Tier 2 memory = CXL® attached memory
- Supports up to 4 channels of CXL memory across two CXL type 3 devices
- Supports CXL memory latency QoS distress signaling

Increased transactions per second for In-Memory databases (e.g. Redis)

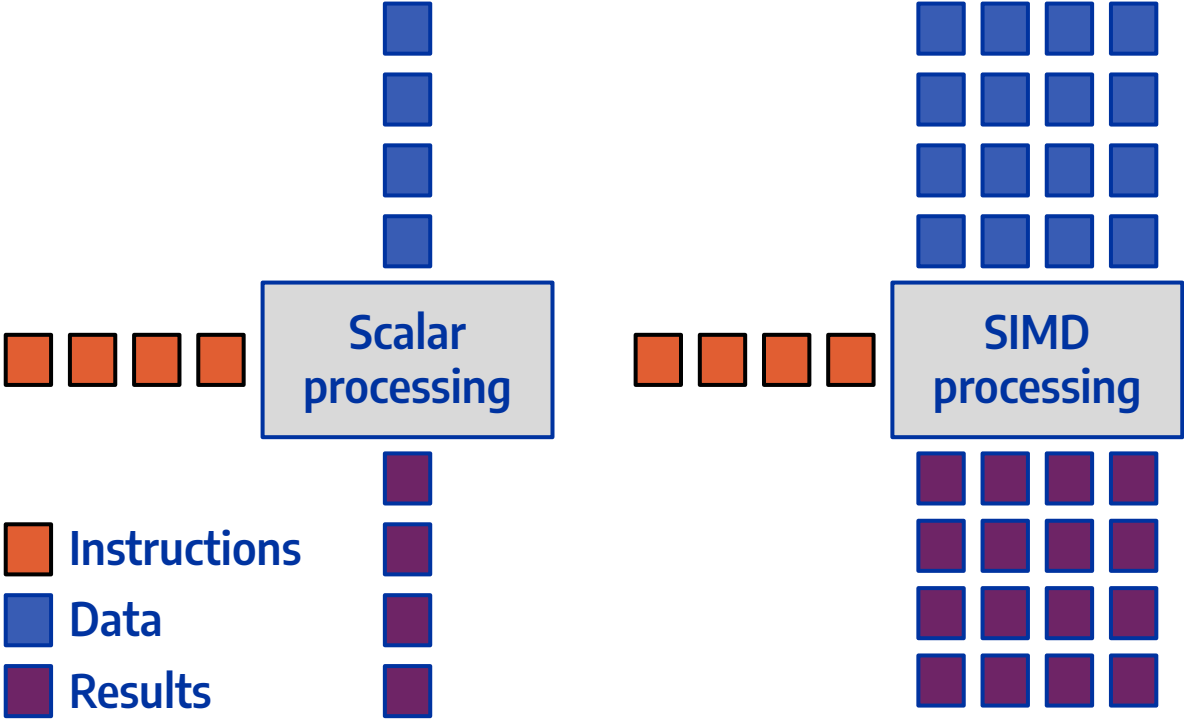
## Single Tier Memory Support

12 channel DDR+CXL interleaved memory

- Either for capacity or bandwidth expansion

intel


# SIMD Vectorization





# History of Intel® SIMD ISA Extensions

- Intel® Pentium Processor (1993)

 32bit


- Multimedia Extensions (MMX in 1997)

 64bit integer support only

- Streaming SIMD Extensions (SSE in 1999 to SSE4.2 in 2008)

 32bit/64bit integer and floating point, no masking

- Advanced Vector Extensions (AVX in 2011 and AVX2 in 2013)


 Fused multiply-add (FMA), HW gather support (AVX2)

- Many Integrated Core Architecture (Xeon Phi™ Knights Corner in 2013)

 HW gather/scatter, exponential

- AVX512 on Knights Landing, Skylake Xeon, and Core X-series (2016/2017)

 Conflict detection instructions

 = 32 bit word

# Evolution of Intel® SIMD ISA Extensions

- AVX 10
  - Supported on both P-cores and E-cores
  - Brings benefits of AVX512 to smaller registers
- Advanced Matrix Extensions (AMX)
  - Targeted at AI applications
  - SIMD for small matrix operations
  - Available on 4th and 5th generation Xeon
- Advanced Performance Extensions (APX)
  - Adds new features that improve general-purpose performance
  - Expands x86 instruction set with more general-purpose registers (from 16 to 32)
  - New REX2 prefix provides uniform access to the new registers
  - Adds conditional forms of load, store, and compare/test instructions
  - New prefix increase average instruction length, but there are less instructions overall

Intel® AVX	Intel® AVX2
128/256-bit FP	Float16
16 registers	128/256-bit FP FMA
NDS (and AVX128)	256-bit int
Improved blend	PERMD
MASKMOV	Gather
Implicit unaligned	

Intel® AVX-512
128/256/512-bit FP/int
32 vector registers
8 mask registers
512-bit embedded rounding
Embedded broadcast
Scalar/SSE/AVX "promotions"
Native media additions
HPC additions
Transcendental support
Gather/Scatter
Flag-based enumeration
Intel® Xeon P-core only

Intel® AVX10.1 (pre-enabling)
Optional 512-bit FP/int
128/256-bit FP/int
32 vector registers
8 mask registers
512-bit embedded rounding
Embedded broadcast
Scalar/SSE/AVX "promotions"
Native media additions
HPC additions
Transcendental support
Gather/Scatter
<b>Version-based</b> enumeration
Intel® Xeon P-core only

Intel® AVX10.2
<b>New data movement, transforms and type instructions</b>
<b>Optional 512-bit FP/int</b>
128/256-bit FP/int
32 vector registers
8 mask registers
<b>256/512-bit</b> embedded rounding
Embedded broadcast
Scalar/SSE/AVX "promotions"
Native media additions
HPC additions
Transcendental support
Gather/Scatter
<b>Version-based</b> enumeration
<b>Supported on P-cores, E-cores</b>

# Microarchitecture of a Modern Intel Core


- Front End

- Instruction Fetch and Decode
- Branch Predictor Unit (BPU)
- L1 Instruction Cache
- Instruction TLB

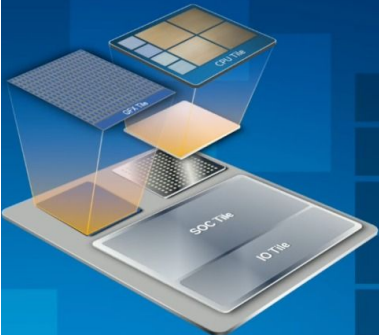
- Back End

- Execution Engine
  - Scheduler
  - Register File
  - Execution Units (EUs)
- Memory Subsystem
  - Load/Store Units (LSU)
  - L1 / L2 Data Cache
  - Data TLB

## Intel Golden Cove Core



# Meteor Lake Hybrid Architecture




**Meteor Lake**  
built on Intel 4 process node  
powering on in Q2'22, shipping in 2023

- New Flexible Tiled Architecture
- Hybrid Cores
- Lower Power
- Next Gen Graphics Engine: tGPU
- Integrated AI Acceleration


Ultra Mobile



Mobile

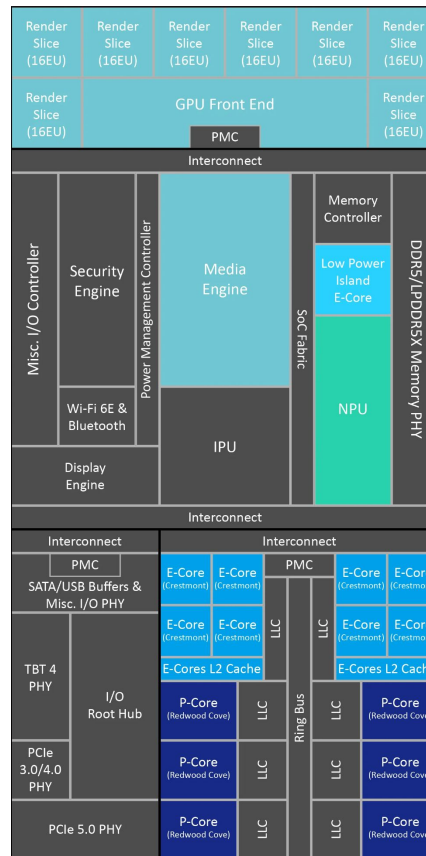


Desktop



Source: Intel

Meteor Lake Block Diagram



REDWOOD COVE

# New P-core

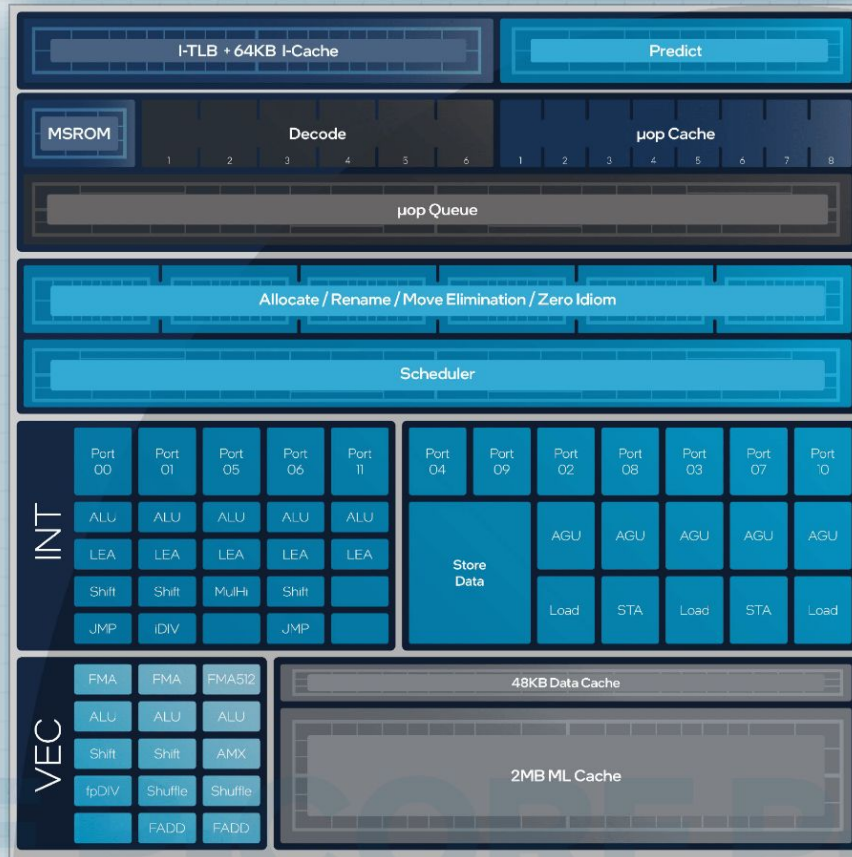
Targeted for efficient performance

Improved performance efficiency\*

Increased BW per core package\*

Improved Performance Monitoring Unit

Improved feedback Intel Thread Director



Source: Intel

\*Architectural simulation vs. Golden Cove architecture. Results may vary across workloads.



CRESTMONT

# New E-core

Significant improvements over prior E-core

IPC gains  
over prior E-  
cores\*

AI  
acceleration  
VNNI, ISA  
improvements\*

Enhanced  
branch  
prediction

Enhanced  
Feedback  
Intel Thread  
Director



\*Architectural simulation vs. Gracemont architecture across a broad set of workloads. VNNI improvements based on doubling the number of VNNI ports. Results may vary.

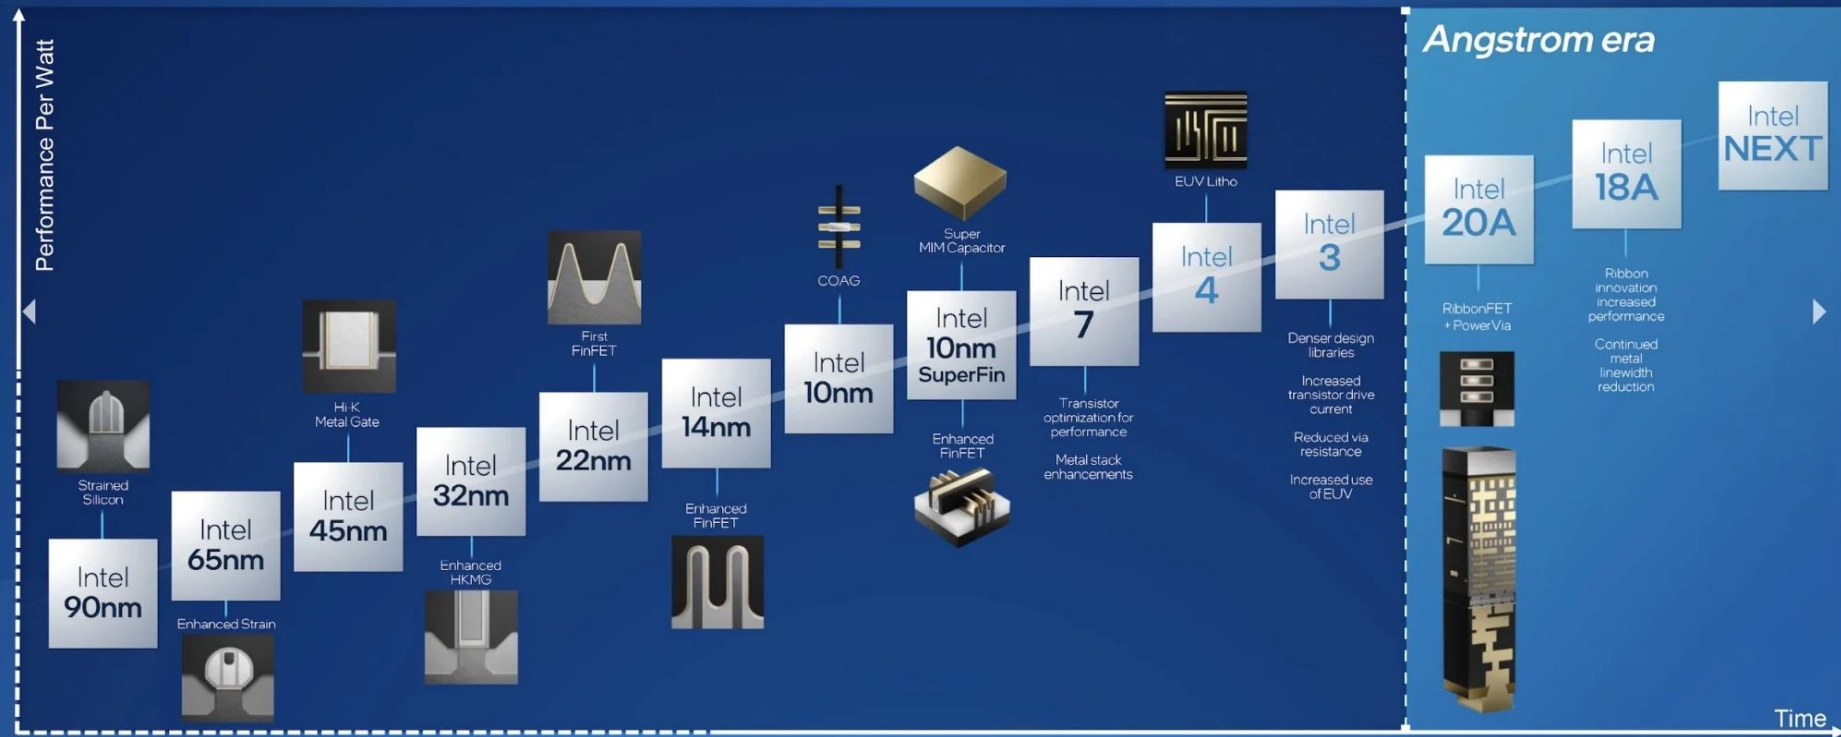
# DCAI Architecture Evolution



Source: Intel

Roadmap: 2023-2025

# Intel Process Technology

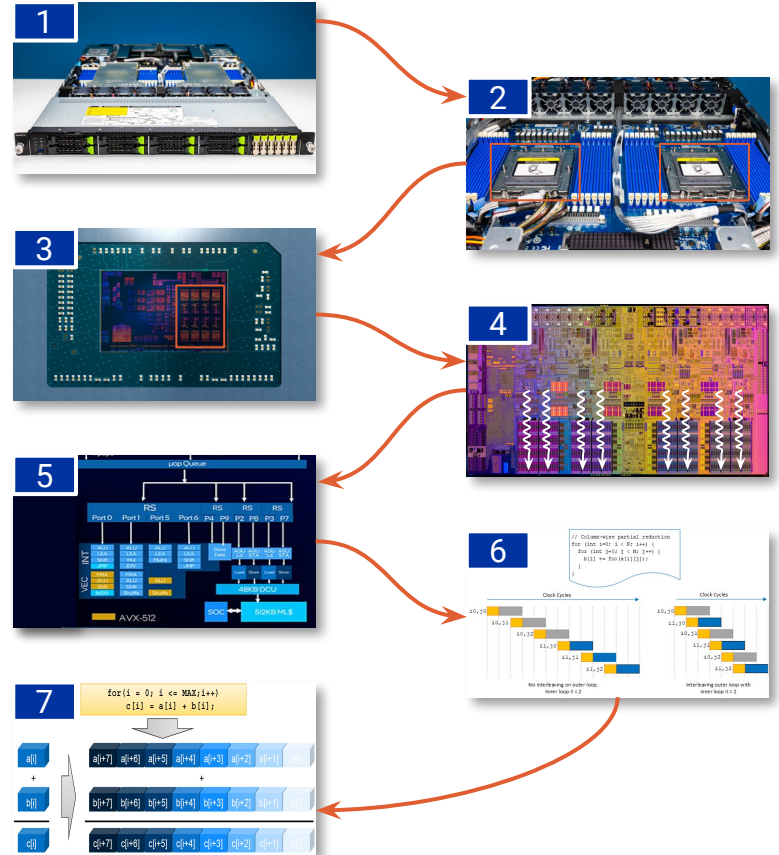


Source: Intel



# Modern Hardware

- 1 Systems**
  - Has always been there
  - Run copies of your code on each node
- 2 Sockets**
  - Modern machines have up to 8 sockets
  - Trend towards reduction back to 2 only
- 3 Cores**
  - Higher frequency no longer possible
  - 128 cores in each socket now common
- 4 Threads**
  - Hardware has N physical cores
  - OS sees 2N logical cores, shared exec.
- 5 Ports**
  - Superscalar execution
  - Multiple instructions executed at once
- 6 Pipelining**
  - Parallel instruction execution steps
  - Fetch / Decode / Execute / Write Back
- 7 SIMD Vectors**
  - SSE4.2, AVX2, AVX512, AVX10
  - PowerPC Altivec, ARM SVE, etc



# Summary

- We've come a long way, modern hardware is quite complex
  - NUMA Architecture (multi socket)
  - High parallelism (multicore, superscalar)
  - Advanced Packaging (chiplets)
  - Hybrid Architectures (performance/efficiency)
  - Variable CPU frequency scaling (turbo boost, thermal throttling)
  - Accelerators and Heterogeneity (GPUs, NPUs, FPGAs, ASICs)
- Performance does not come for free, we needed to adapt our software
  - Concurrency and Parallelism (processes, threads, SIMD)
  - Memory alignment, access patterns, fragmentation
  - Code layout, compiler optimizations, data structures, software design
  - Need the right tools to guide us: profilers, static analysis, etc
  - Need the right methodology: identify causes of bottlenecks, address the right issue