VLAB WORKS™

astc

Current Computer Architecture, Power Management (and ASTC VLAB™)

Jakob Engblom, PhD
Global Technical Marketing Manager, ASTC VLAB

2025-05-19

Jakob Engblom

- Global Technical Marketing Manager, at ASTC

- Educational background:
  - MSc, Computer Science, and PhD, Real-Time Systems, Uppsala
- Experience: virtual platforms, simulation, embedded systems
  - Product management, product marketing, technical sales, technical marketing, business development, training development, demos, ...
  - IAR Systems, Virtutech, Wind River, Intel, ASTC
- My own blog, since 2007:
  - https://jakob.engbloms.se

What is in a Computer?

What’s in a “Computer”?

- (Main) Processor cores
  - Run user-visible OS and applications
- Main memory (“RAM”)
- Compute acceleration
  - AI/ML, compression, networking, …
- Graphics and display
- Video and audio processing
- Storage (disks)
  - NVMe, SSD, HDD, SD, Flash, …
- Networking
  - Ethernet, WiFi, Bluetooth, …
- Local peripherals -
  - USB, Thunderbolt, Serial, Bluetooth, ...

Once Upon a Time...

- The "processor" was the essential part of a system
- It measured the goodness of the machine, with measures like:
  - Megahertz (not GHz yet)
  - Instructions per cycle
  - Cache size (from the 1990s on)
- Getting a better computer meant a better processor (mostly)
- The supporting chipset and I/O was quite basic

Year 2000: Intel Pentium 4

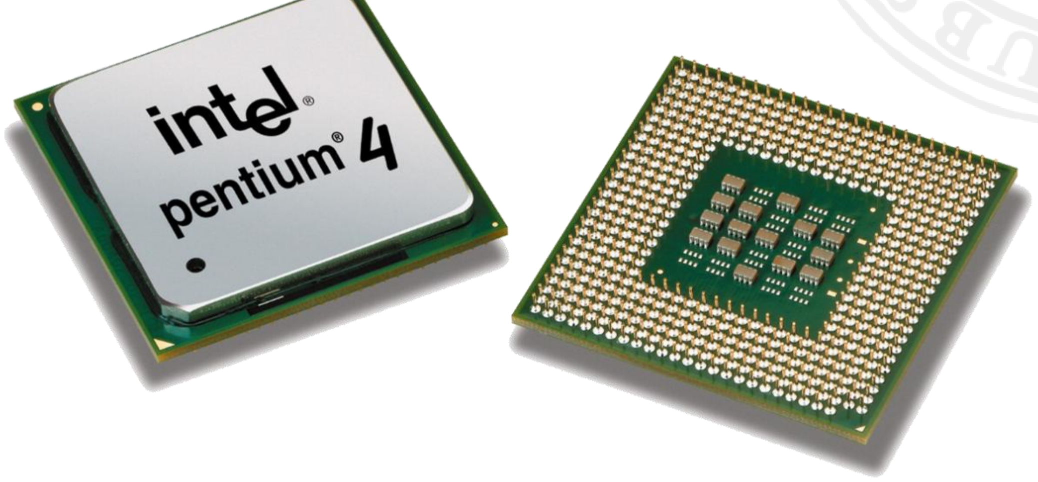

https://www.youtube.com/watch?v=wEiPDzDjbjU

2009: Intel® Core™ i7 Processor: Still a Processor

- Intel® Core™ i7-960 Processor (2009)
  - 4 Processor cores
    - (Multicore was the big trend of the early 2000s)
  - Cache
  - Memory controller
    - Moved on-chip in this generation
    - Used to be on an external chip
  - PCIe 2.0 for graphics
  - Intel QuickPath Interconnect (QPI)
    - Link to the chipset (rest of system)

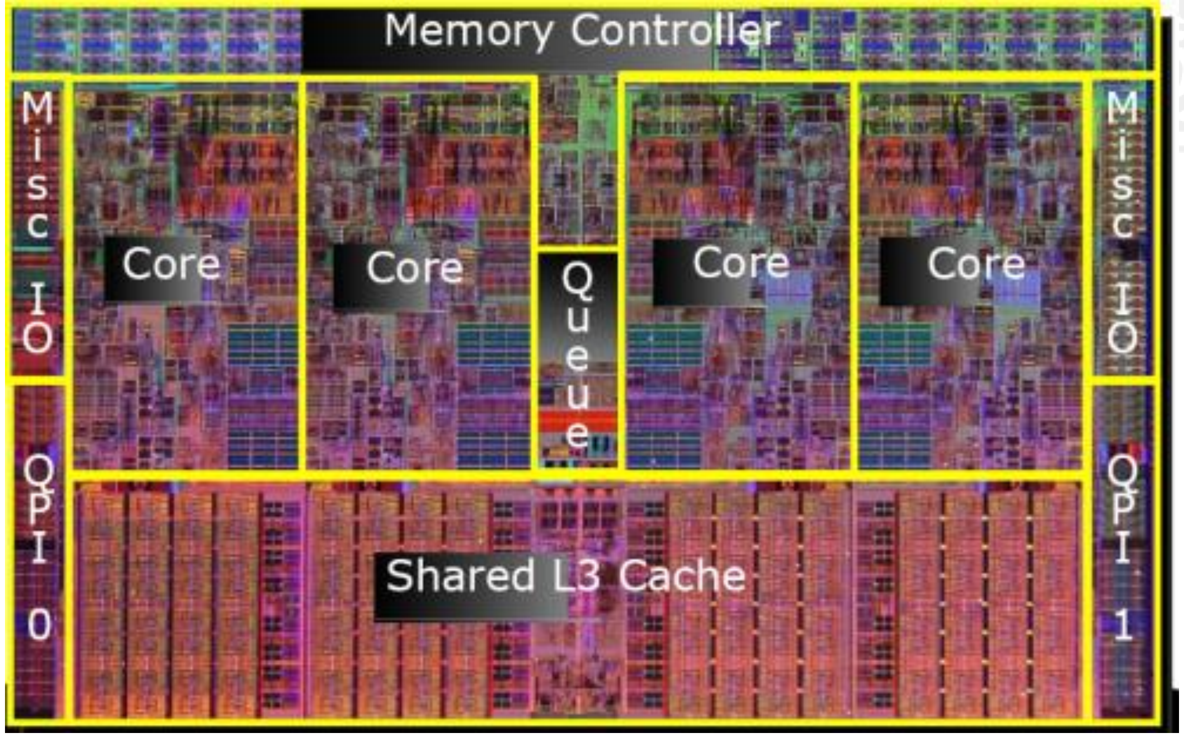

2009: Intel® X58 Express Chipset

By 2023: 1 channel of DDR5-5600 ≈ 45GB/s.
5.5x faster in 14 years

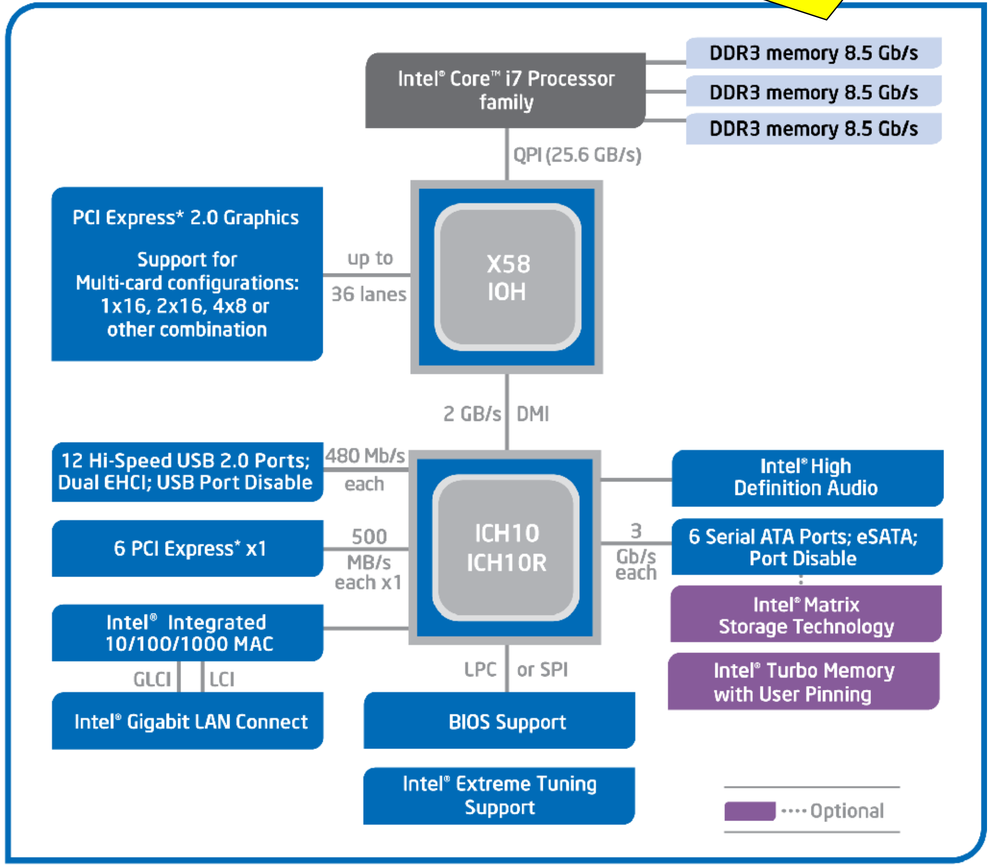

Intel® X58 Express Chipset Block Diagram

- Two chips + the processor
  - Today, all of this is integrated as a single unit
- I/O Hub (X58 IOH)
  - QPI - Fast link to the processor
  - Graphics cards and other high-bandwidth PCIe devices
- I/O Controller Hub (ICH10)
  - Linked to the IOH over a slow link
  - Main IO chip for slow IO
  - SATA, Audio, USB, PCIe, Ethernet

Chip

Note: “Chips”

Leading-edge fab companies
- TSMC
- Intel
- Samsung

Many more at trailing-edge nodes

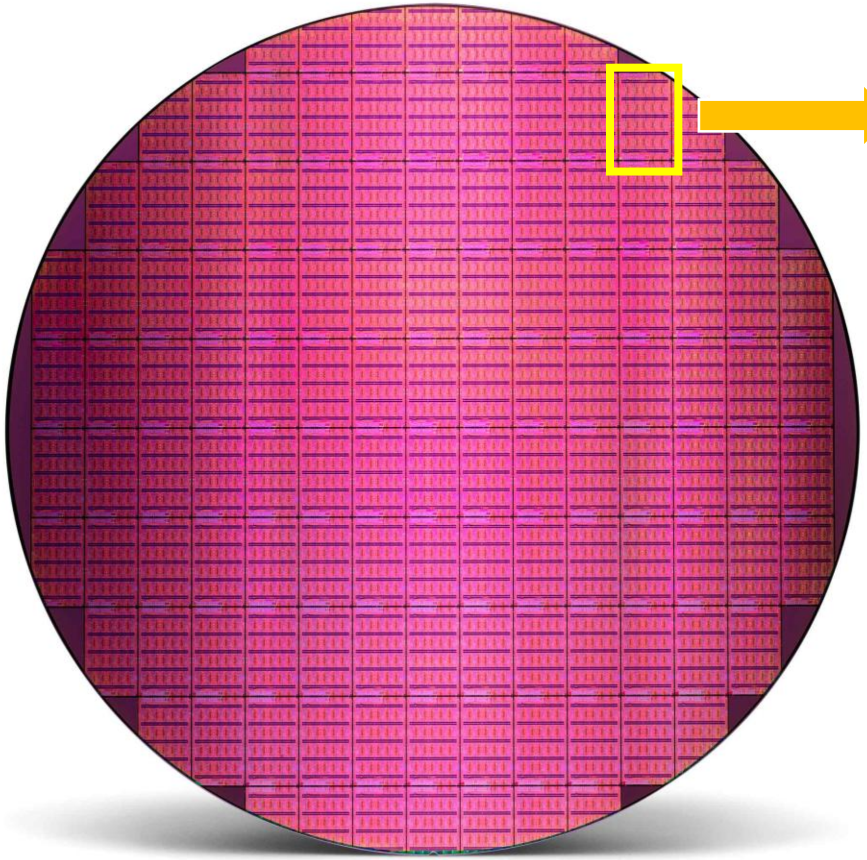

Wafer = what gets manufactured in a fab

Each wafer contains **dies**

“**The max reticle size**” is the biggest die that can be printed. Currently around 26x33mm.

One or more dies are put onto a **package**

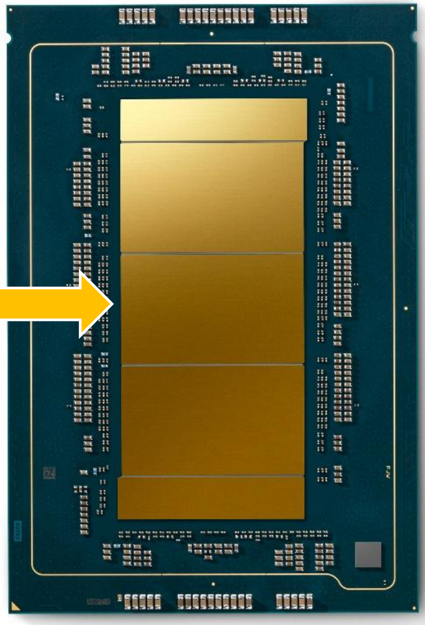

A package is what most people call a “chip”

The package is covered by a **lid**

The back side has **bumps** (pins) for signals and power

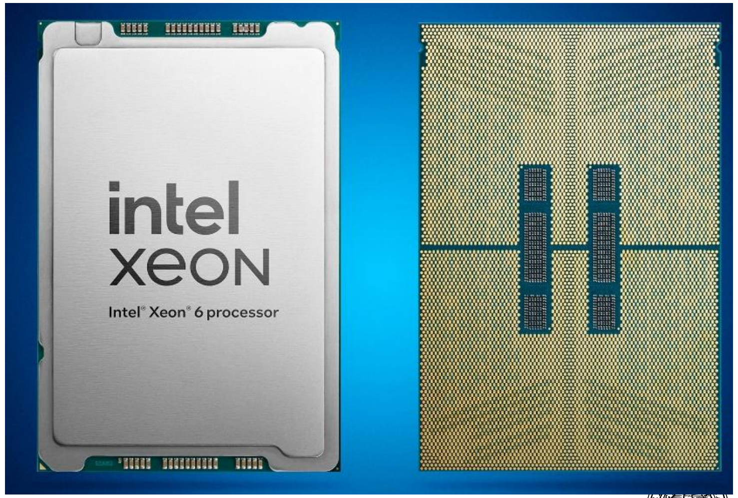

Computer Architecture Today & Trends

Trend

Rich I/O – Home Office Example

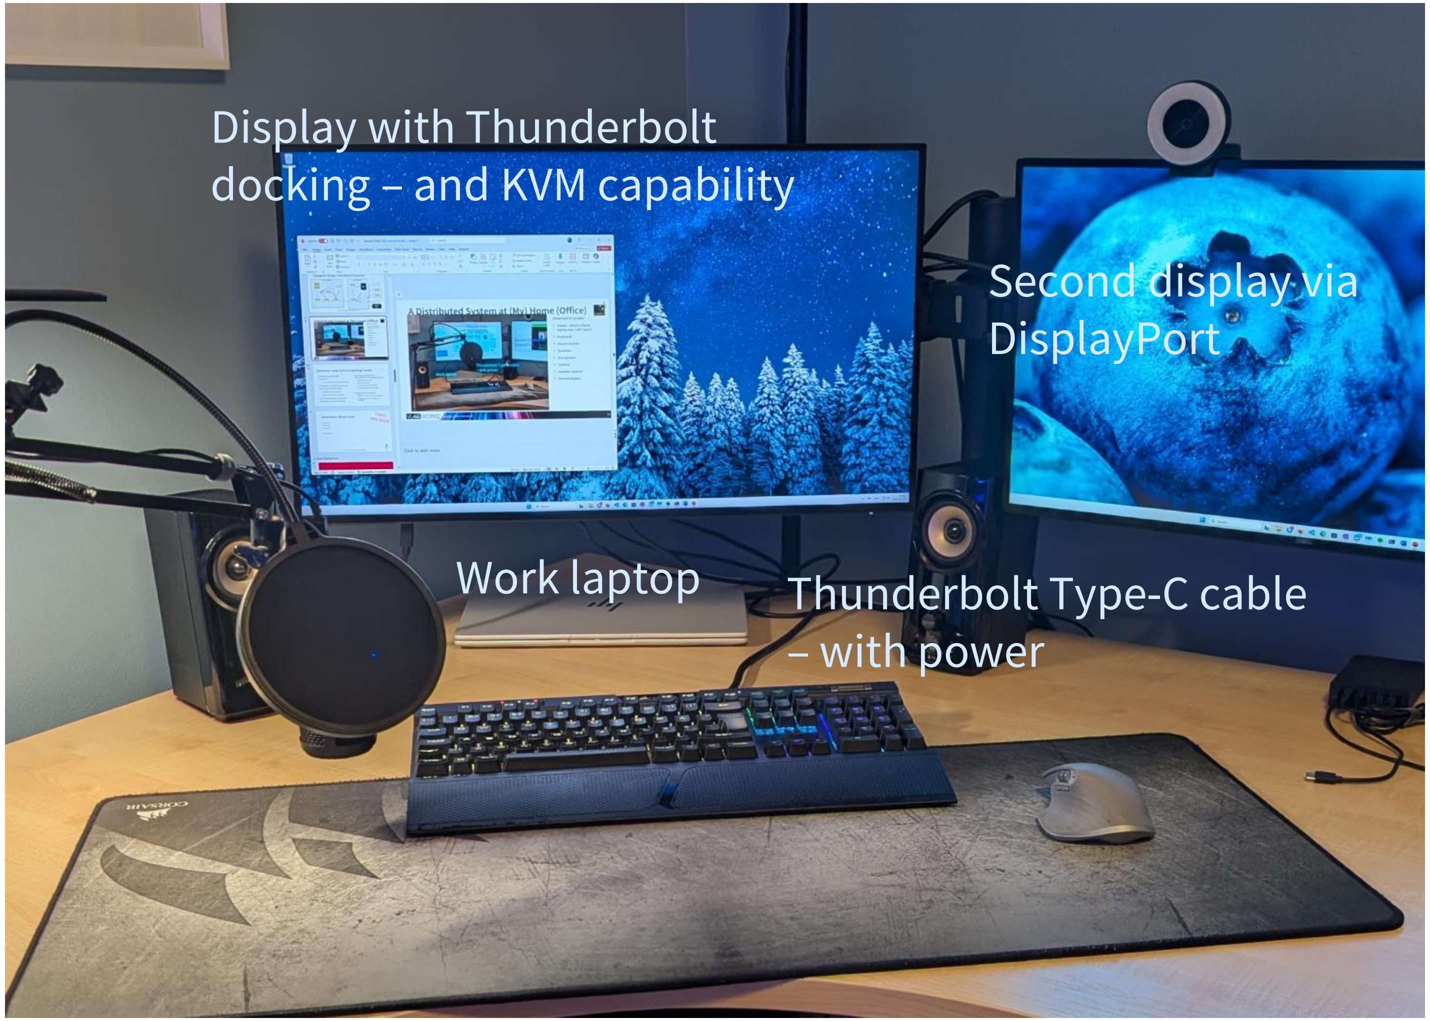

- Attached to primary screen:
  - Power – which is fed to laptop over USB Type-C
  - Keyboard (USB)
  - Mouse receiver (USB)
  - Speakers (USB)
  - Microphone (USB)
  - Camera (USB)
  - Second display (DP)
- Primary screen works as a Keyboard-Video-Mouse switch
  - To a second computer under the desk
    - Display port
    - Second USB-C input to screen

There is Software in that Screen

Misguided picture-in-picture?

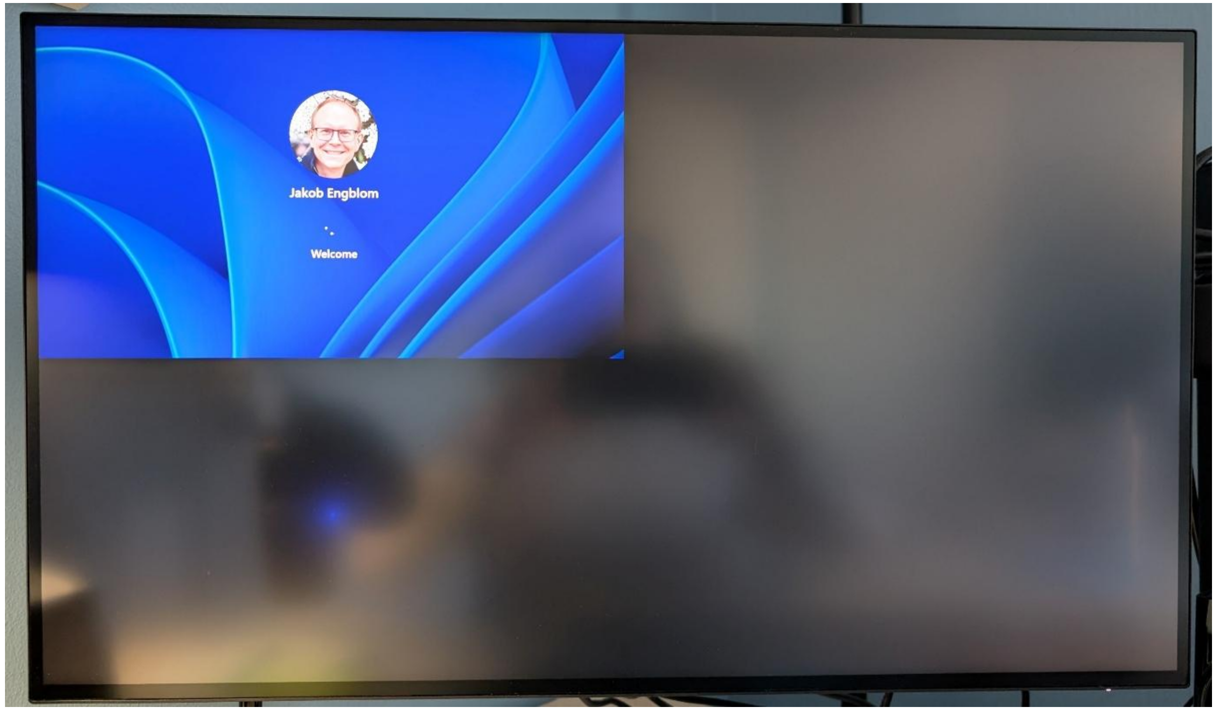

Strange skewing and splitting

Software Everywhere in the System

FW = Firmware, built-in software

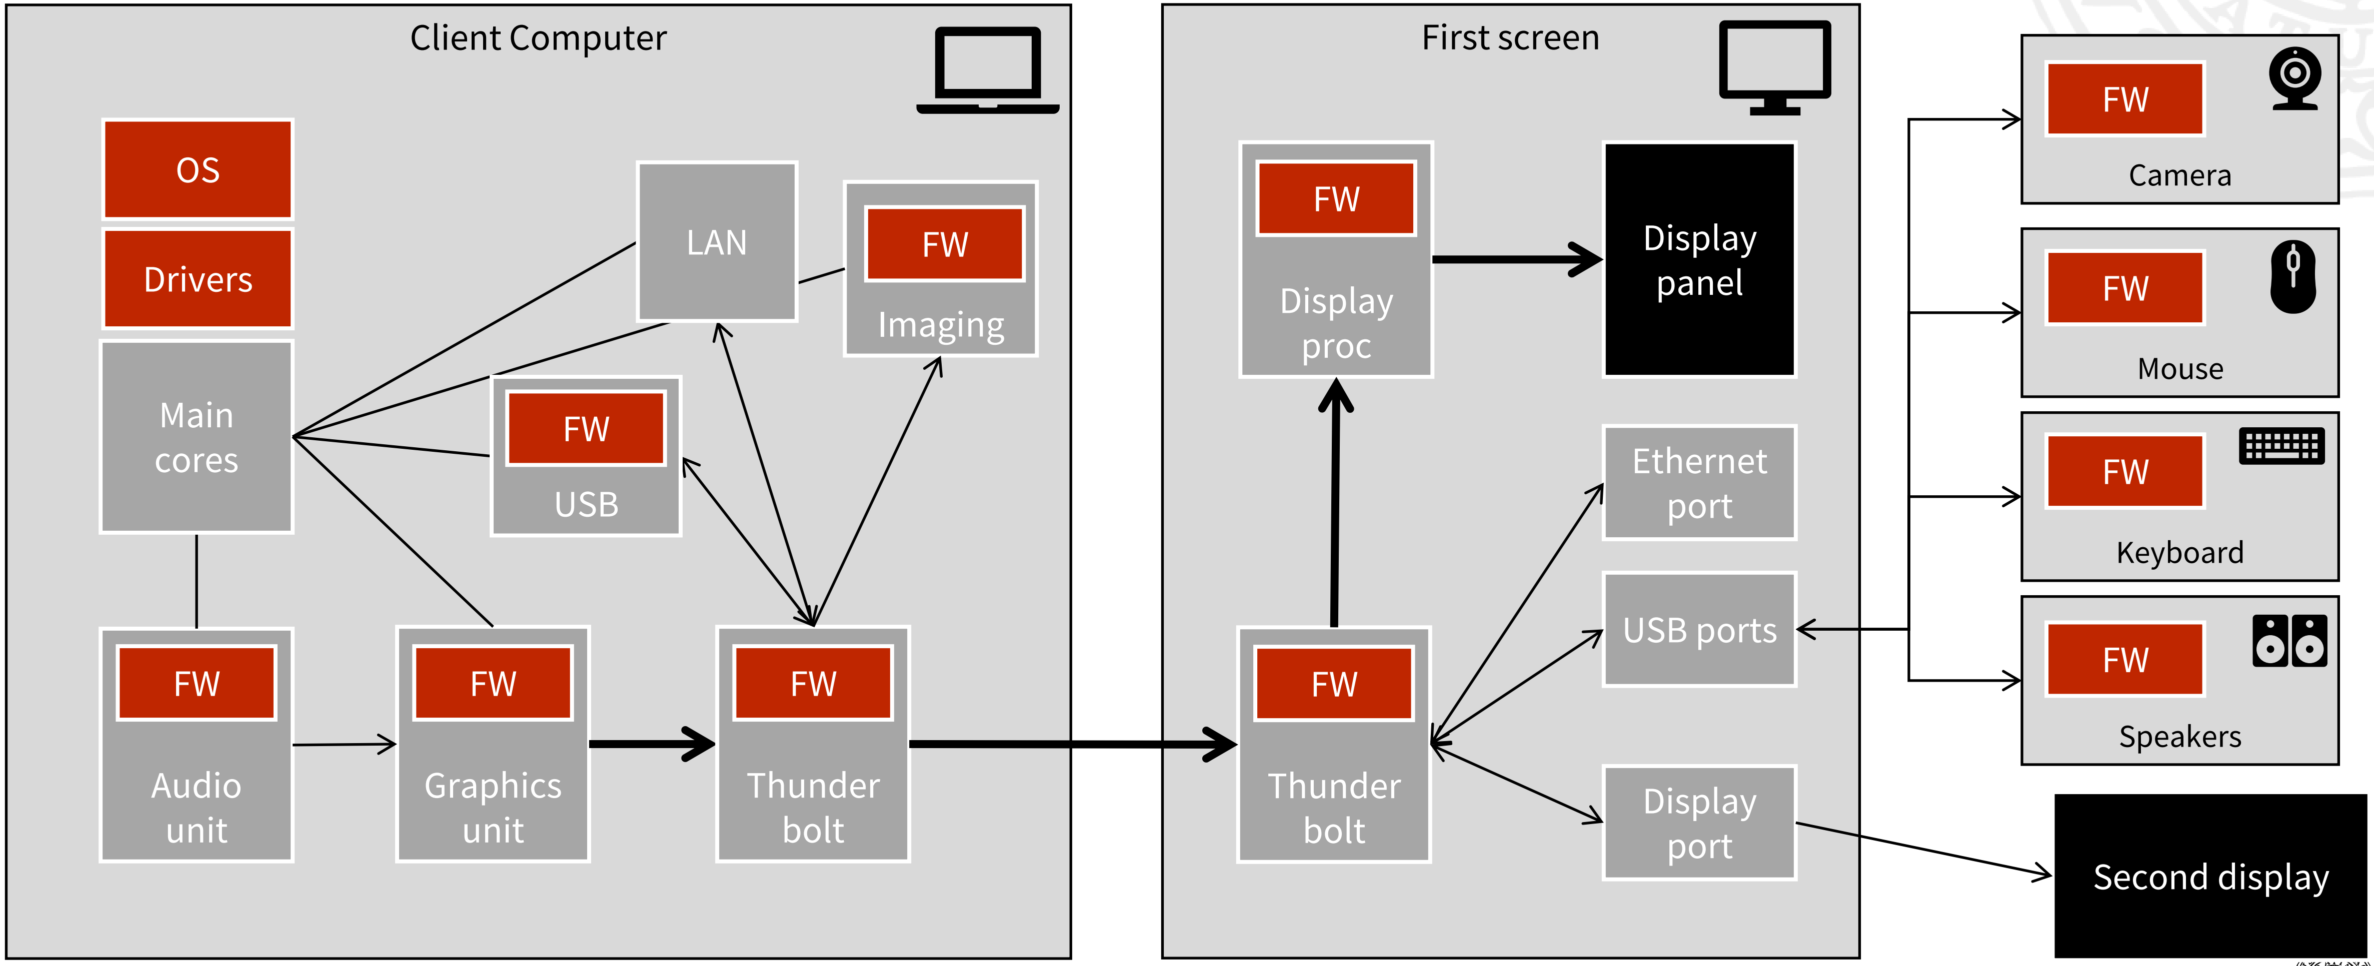

Trend

Big and Small (and Intermediate) Cores

Processor core performance/cost tradeoff

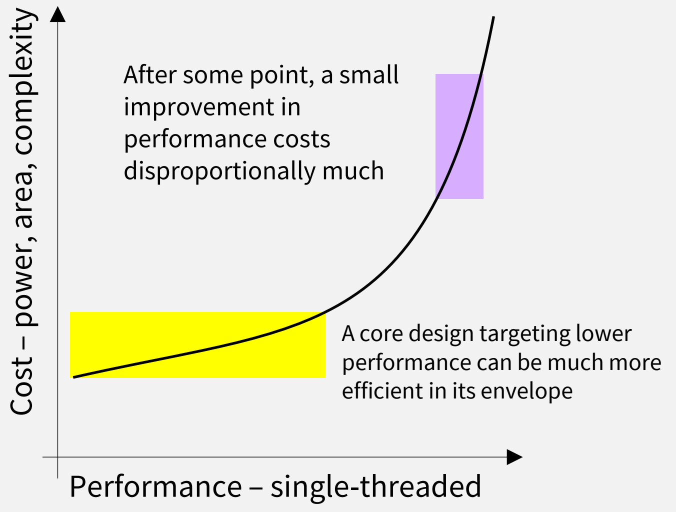

Graph just for illustration purposes – not based on any real numbers

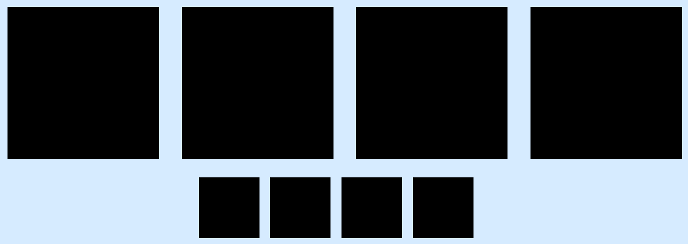

Mobile phone: Light background tasks on small cores, decent performance for foreground from big cores

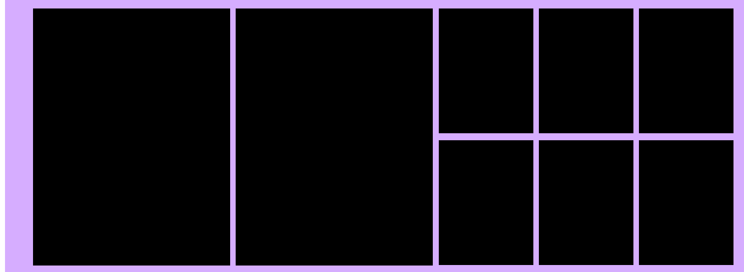

Laptop: Maximum performance big cores, small*er* cores can do significant processing at lower area and lower power – scale up multithreading

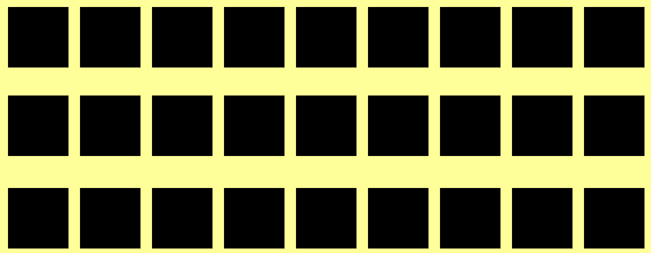

Throughput server or budget phone: Only small cores, trade single-threaded performance for multi-threaded throughput (or just lower cost)

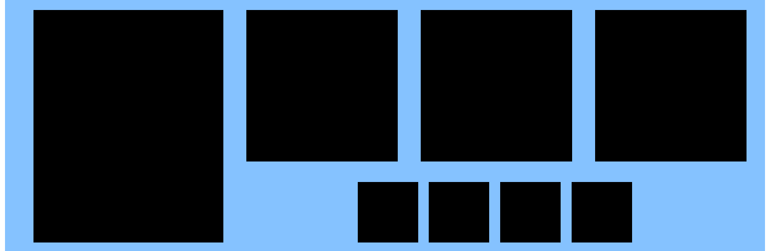

Mobile phone 2: One extra-big core for critical single-threaded tasks, plus “big”, plus “small”

More reading:
https://chipsandcheese.com/2021/12/21/gracemont-revenge-of-the-atom-cores/

https://www.tomshardware.com/reviews/intel-core-i9-12900k-and-core-i5-12600k-review-retaking-the-gaming-crown/6

https://www.anandtech.com/show/17102/snapdragon-8-gen-1-performance-preview-sizing-up-cortex-x2

2023: Intel Raptor Lake – Desktop Processor

- "Intel® Core™ i-13xxx"
  - On "old" Intel 7 process
- Mostly processor cores
  - 8 x Big/Performance "Raptor Cove"
  - 16 x Small/Efficiency cores "Gracemont"
- Small graphics/media block
  - Good enough for office use, no more
  - Most of it is video/media engines, not 3D graphics
- System agent
  - PCIe
  - DDR4/DDR5 memory controller
  - Display controller – can drive 5 displays!

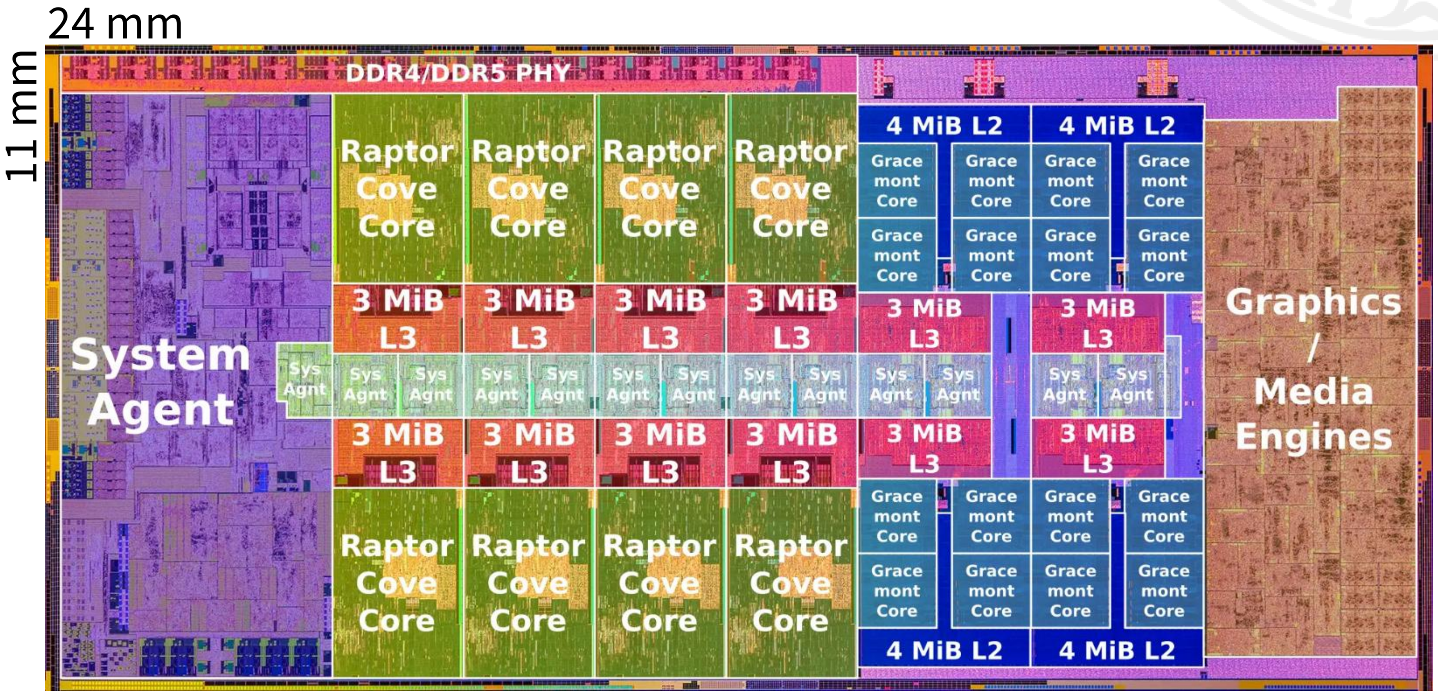

Source: Intel Raptor Lake annotated die shot from Wikichip
https://en.wikichip.org/wiki/File:intel_raptor_lake_die_%288%2B16%29_%28annotated%29.png

2023: Intel® Z790 Chipset – for Raptor Lake

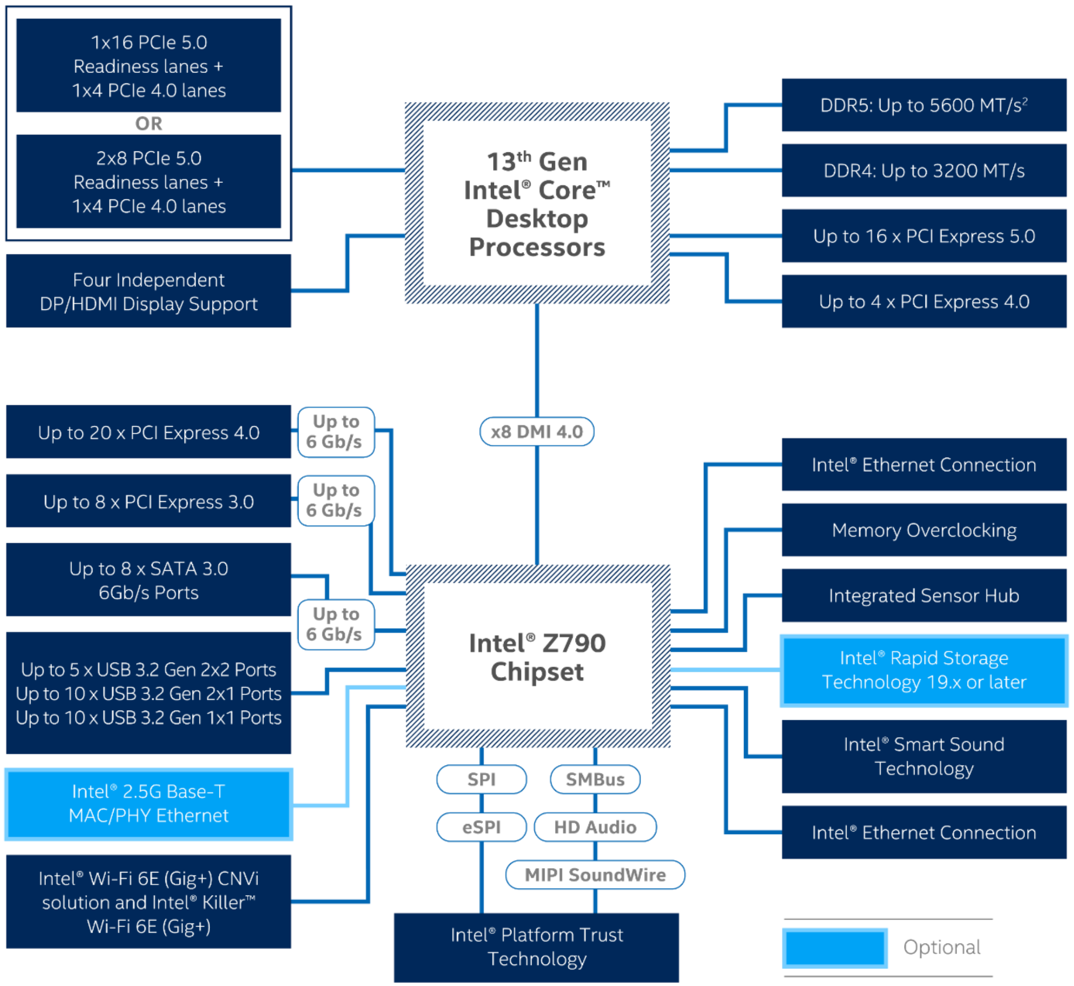

- Processor:
  - External connections to memory and PCIe
- Chipset: Platform Controller Hub (PCH)
- Massive I/O capabilities:
  - PCIe: 16 + 4 and 20 + 8 lanes
  - USB: 15-20 (including Type-C)
  - WiFi + wired Ethernet (require external PHYs)
  - Displays, sound
  - Thunderbolt – add as external chip
- Designed for user-installable processors
  - Desktop is not space-constrained = external chips OK
  - Motherboard manufacturers can differentiate with additional features around the chipset

Trend

Accelerators: Specialized Compute

Examples

Graphics rendering (GPU)	Audio processing (DSP)	Camera image processing (ISP)	Neural networks/AI (NPU)	Cryptography
Video input processing	Video encoding/ decoding	Data compression/ decompression	Network packet processing	Data movement

Work shifts to specialized subsystems for higher performance and lower power – provided the workload happens often enough to warrant the investment

Why Accelerators? Why not just Processors?

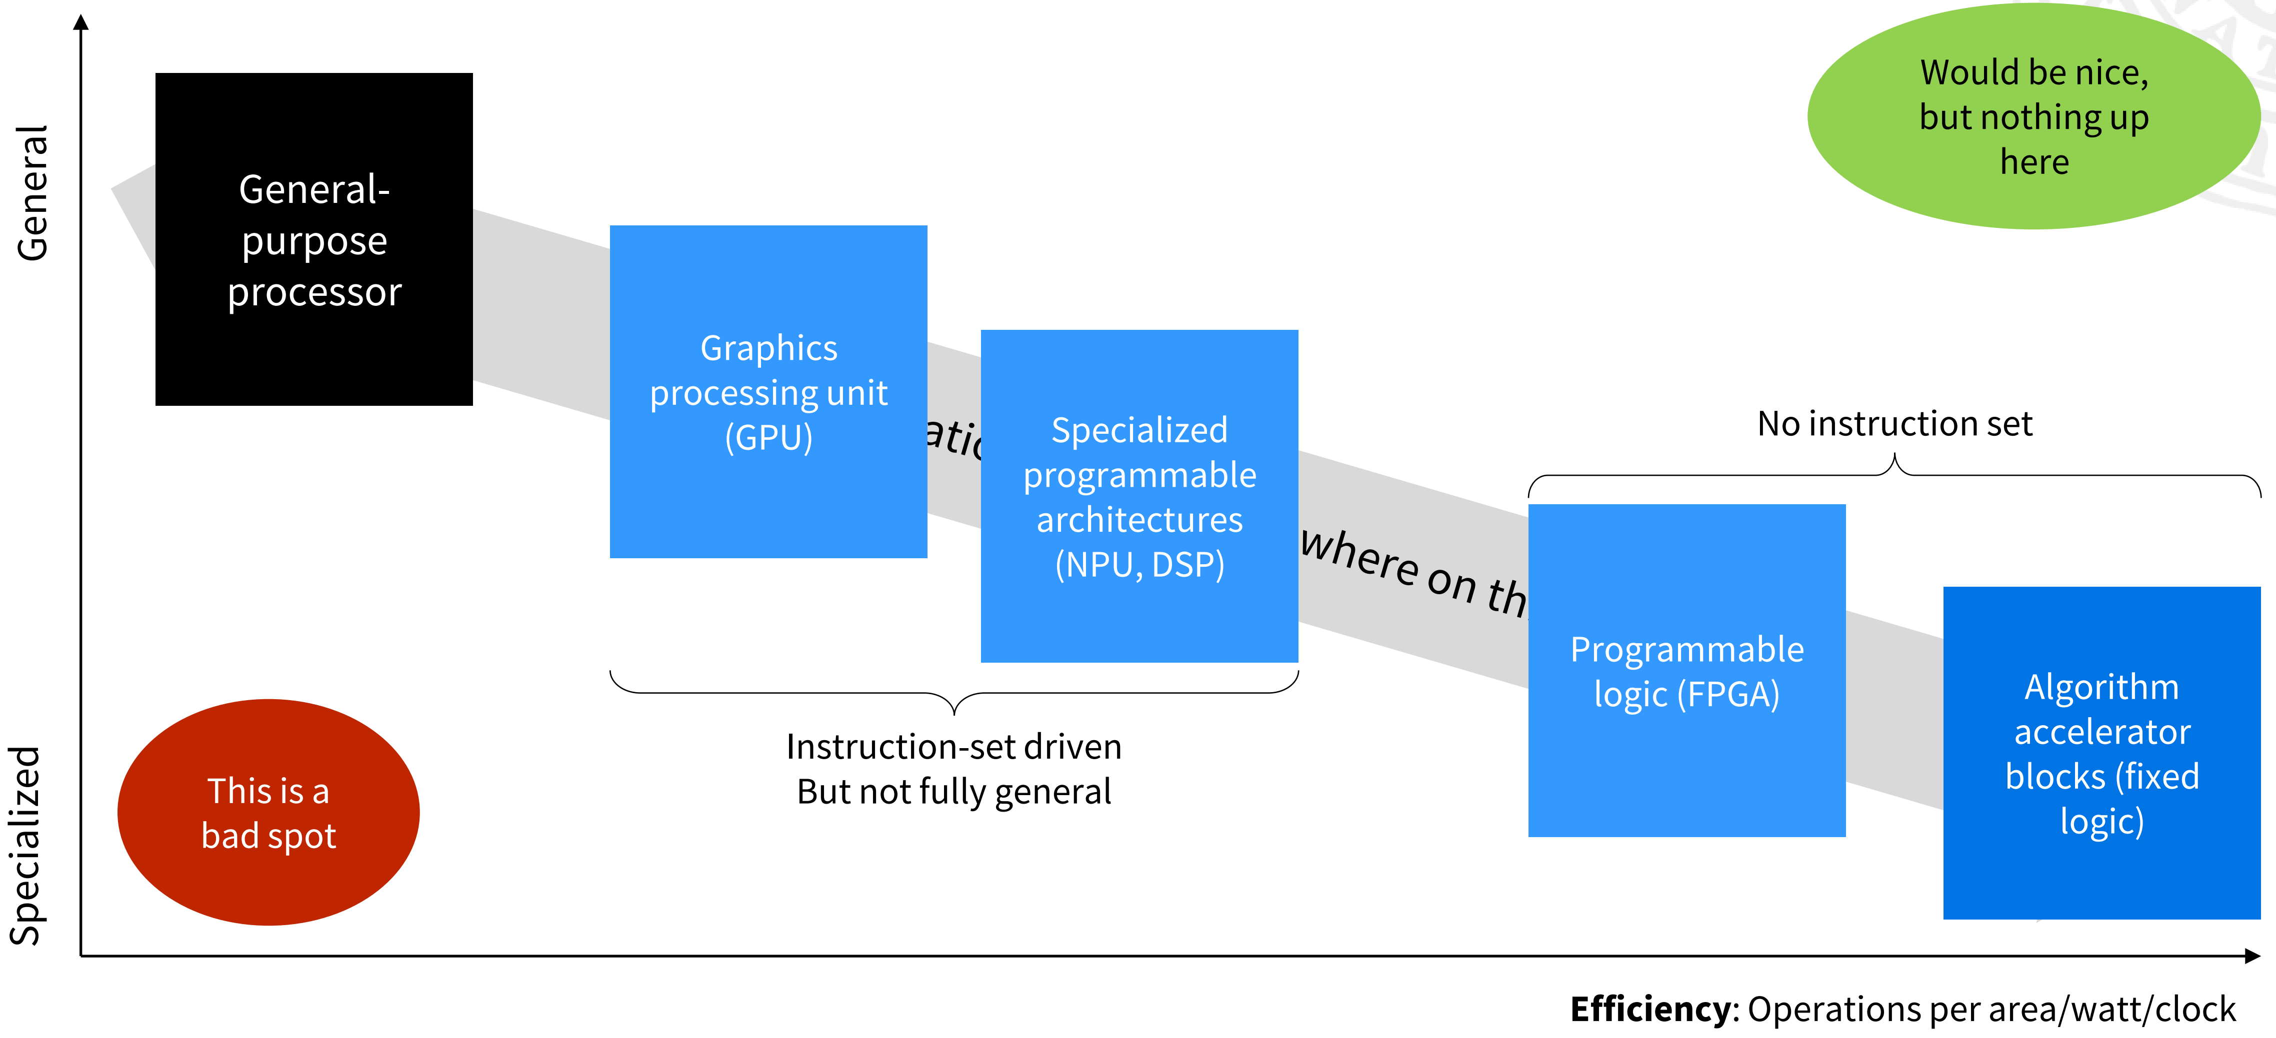

Trend

Current Hype: Neural Processing Units (for AI)

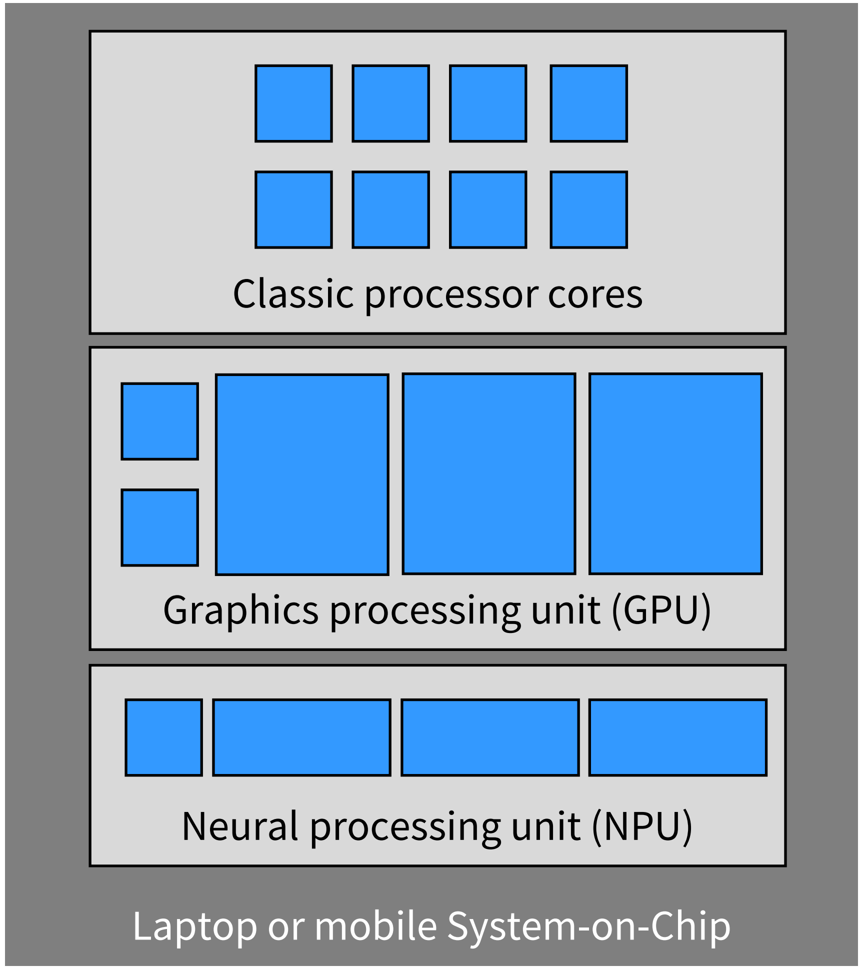

- **Neural Processing Units**, NPU
  - Added to *client* processors
  - Purpose: AI processing at *lower power*
  - User programmable
    - APIs similar to graphics units
- Rapidly increasing in power (and silicon size)
- Actual designs:
  - New accelerator design
    - Like Apple's Neural Engine
  - Repurpose existing processors
    - Qualcomm* Hexagon* DSP
    - Intel® Movidius video processing unit

More reading: https://chipsandcheese.com/2024/04/22/intel-meteor-lakes-npu

More reading: https://chipsandcheese.com/2023/10/04/qualcomms-hexagon-dsp-and-now-npu/

2024: AMD Strix Point Laptop Chip

- “AI PC” **laptop** chip
- No external chipset
- Compute resources on chip:
    - 25% - processors cores
        - 4 x Zen 5 Performance cores
        - 8 x Zen 5c Efficiency cores
    - 20% - GPU
    - 10% - NPU
- System functions
    - DDR5 memory controllers
    - PCIe controllers
    - Media engines and display
    - USB-C
    - Note: Thunderbolt not on chip!

19 mm

12 mm

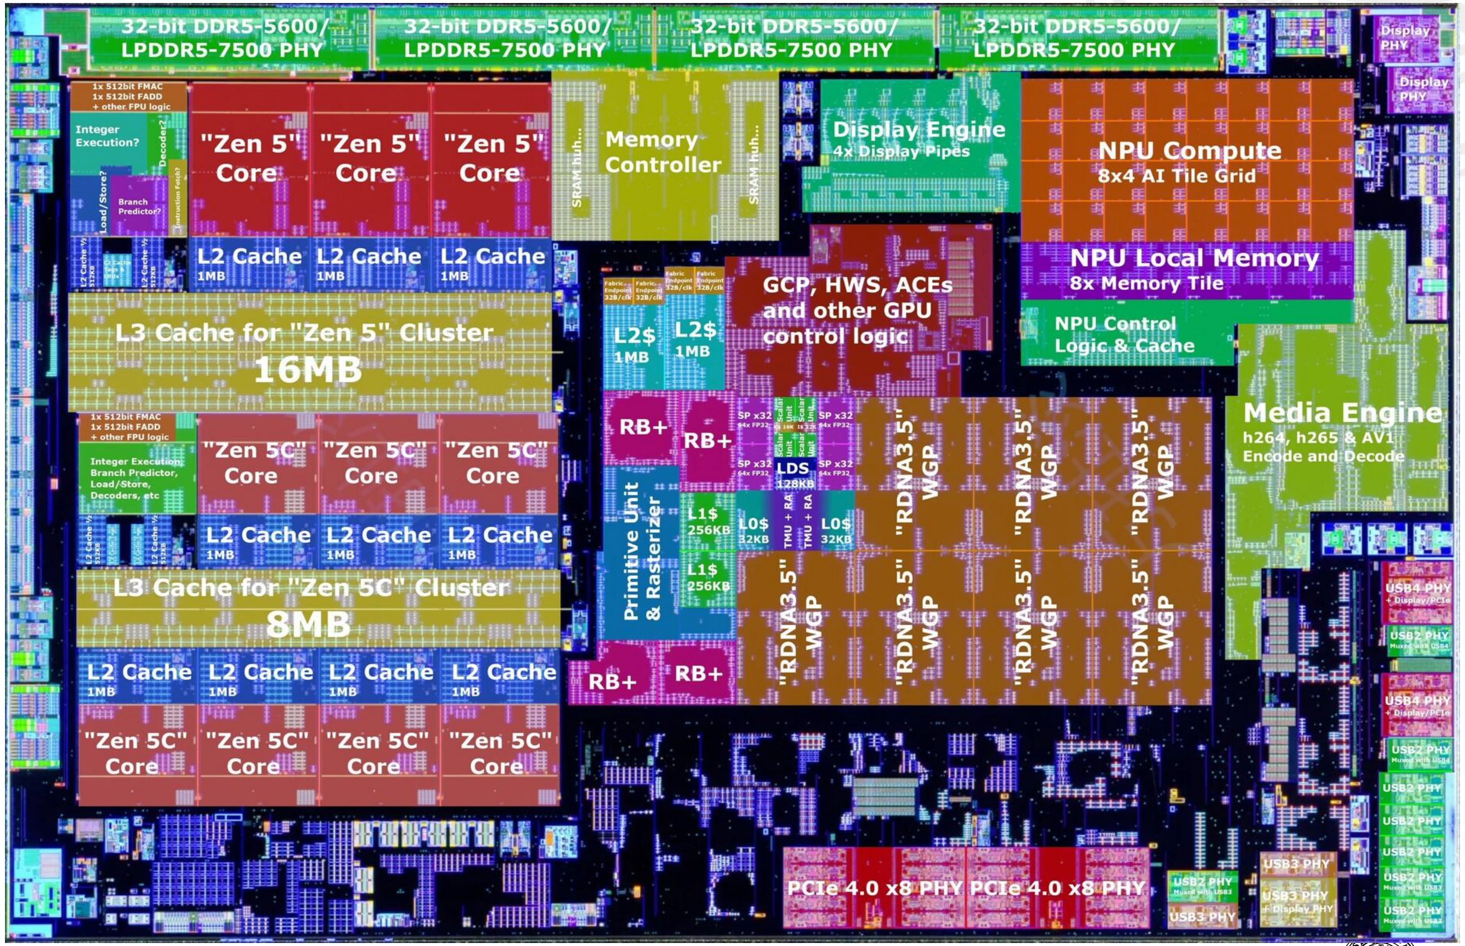

https://www.techpowerup.com/325035/amd-strix-point-silicon-pictured-and-annotated

Trend

Chiplets/Tiles/Dies

- Chiplets is the new standard way to build large “chips”
- **Problems addressed** by chiplets

Total chip size	Yield	Production cost	Die properties	Design cost	Variation
• Use multiple dies to exceed maximum lithography size	• Smaller dies = better yield	• Larger dies are disproportionally expensive • Use the cheapest node that meets requirements for each particular die	• Use technically optimal process for each subsystem (speed, density, power, analog properties, ….)	• Replicate an identical unit multiple times • Fewer chip variants to validate	• Vary the number of chiplets to scale system • Vary one subsystem in isolation

2024: AMD Granite Ridge (Desktop Processor)

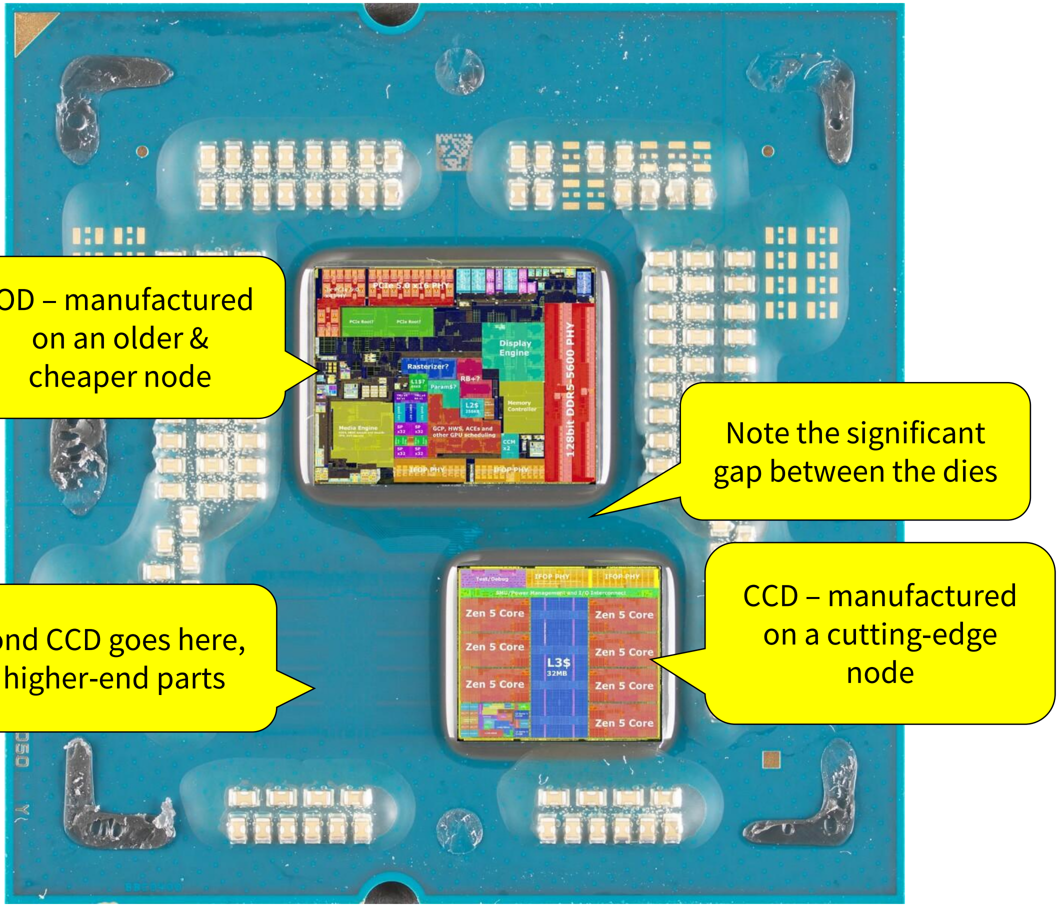

- AMD Granite Ridge, "AMD Ryzen 9000"
  - Desktop processor = plenty of space c
  - Attaches to an "AM5" chipset
- Chiplets:
  - "CCD", Core Compute Die
    - 8 Big "Zen 5" cores + 32MB Cache
    - *TSMC N4P process*
  - "IOD", Input-Output Die
    - Basic GPU
    - Memory controller (dual channel)
    - PCIe controller
    - Display & Media engine
    - USB
    - *TSMC N6 process*
  - Connected using "Infinity Fabric"
- Economical to build, some scalability

Image source: https://www.techpowerup.com/327388/amd-granite-ridge-zen-5-processor-annotated#g327388-1

UPPSALA UNIVERSITET

2023: Intel Meteor Lake (Laptop Processor)

- Intel "Meteor Lake", Core Ultra 100
  - Laptop = maximize integration
- Tiles/Dies (Tile=Intel term for die)
  - **Base tile** underneath to connect
  - Compute tile
    - 6 P-cores + 8 E-cores
    - *Intel 4 process*
  - Graphics processor unit (GPU)
    - *TSMC N5 process*
  - Memory and PCIe controller tile
    - *TSMC N6 process*
  - System tile
    - Low-power island with 2 E-cores
    - AI accelerator (NPU)
    - USB and Thunderbolt
    - Audio and video processing
    - Display – can drive 4 x 4k60 HDR displays
    - *TSMC N6 process*

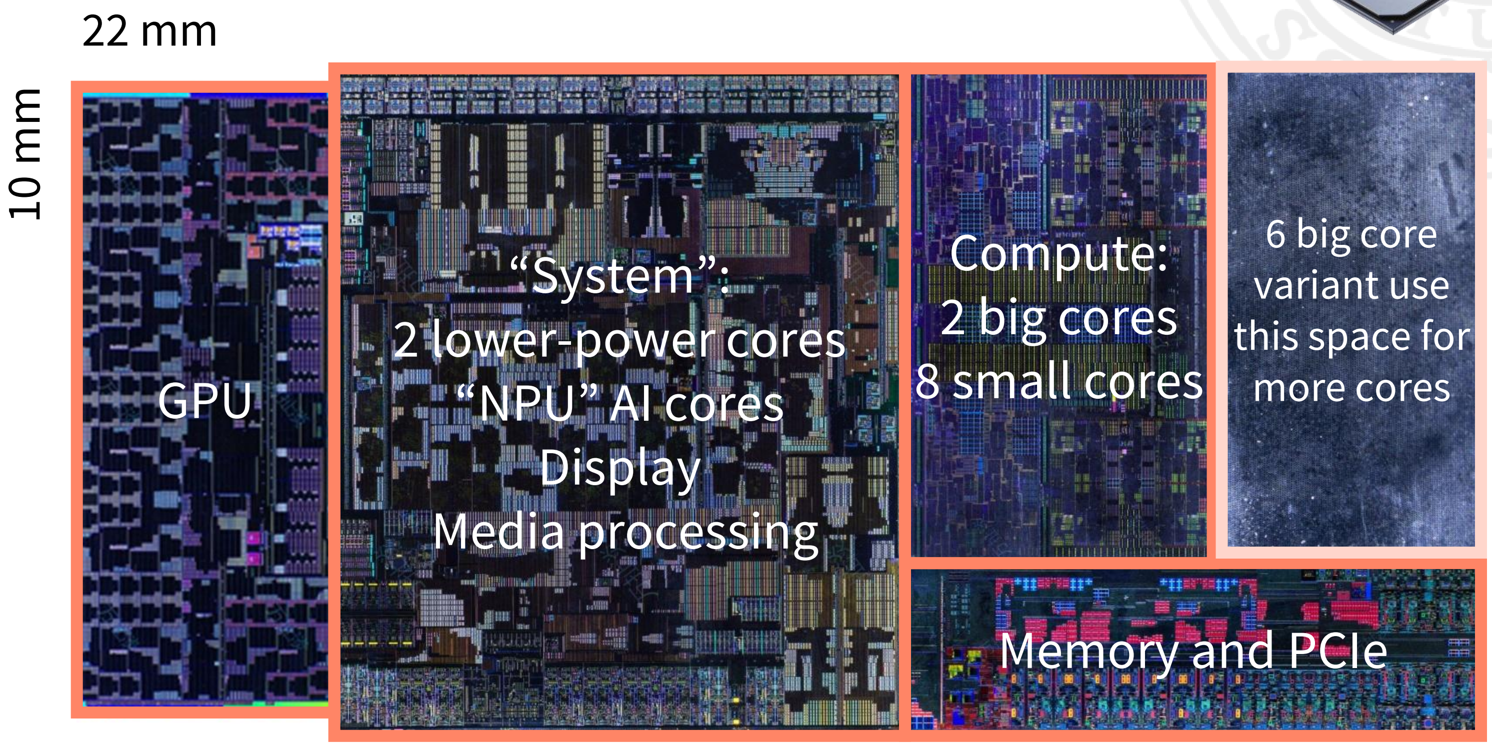

https://wccftech.com/intel-core-ultra-meteor-lake-cpu-die-shots-closer-look-at-various-cpu-gpu-io-chiplets/

https://www.tomshardware.com/pc-components/cpus/meteor-lake-die-shots-show-off-inner-workings-of-soc-and-gpu-tiles

YouTube Tip:

- https://www.youtube.com/watch?v=wusyYscQi0o
- Fantastic walk through an Intel Arrow Lake – similar tiles to the Meteor Lake

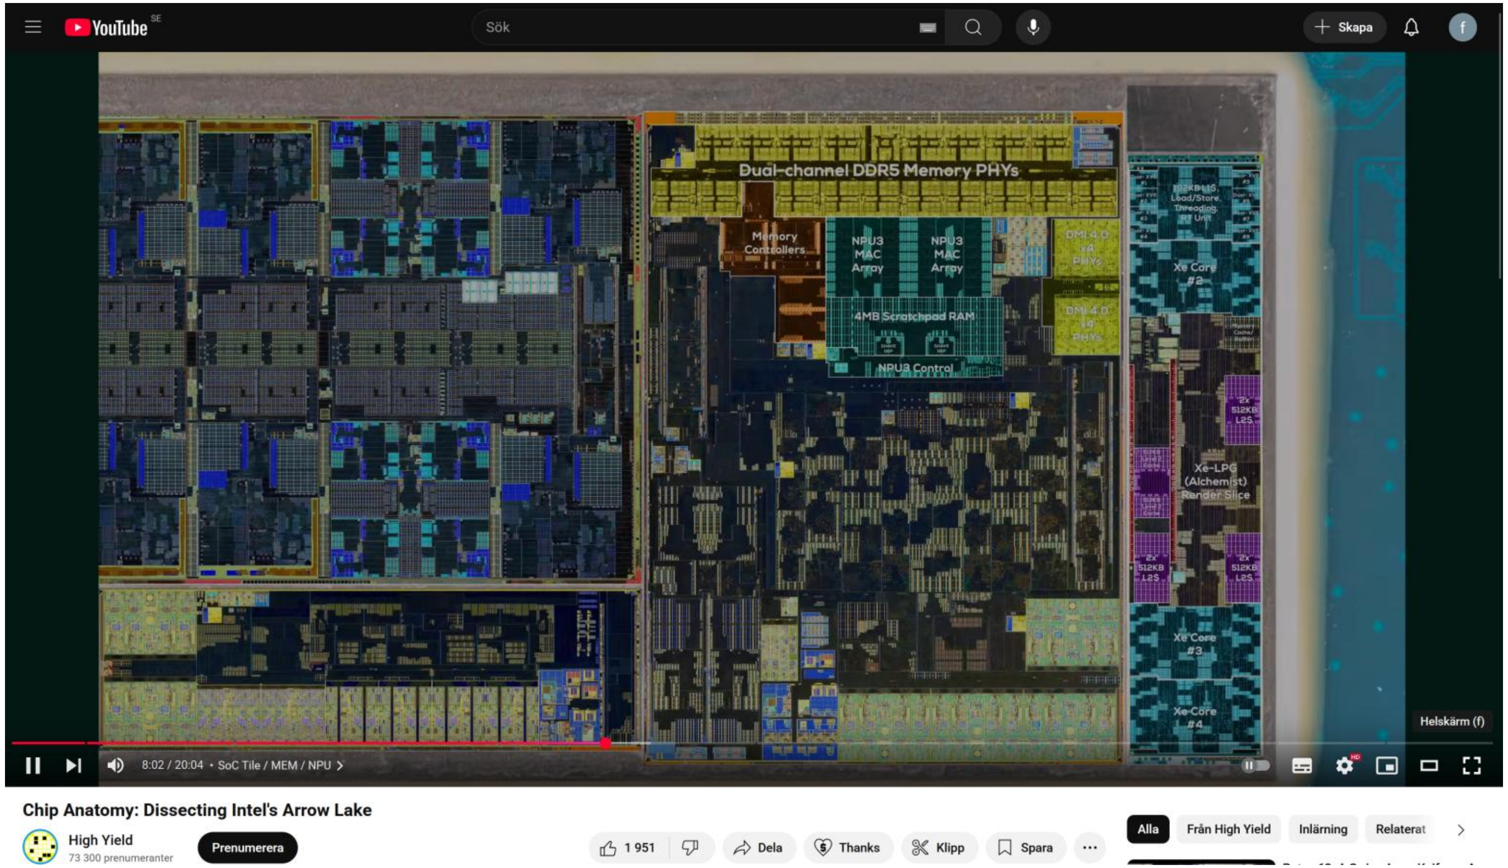

2024: Intel Granite Rapids Servers

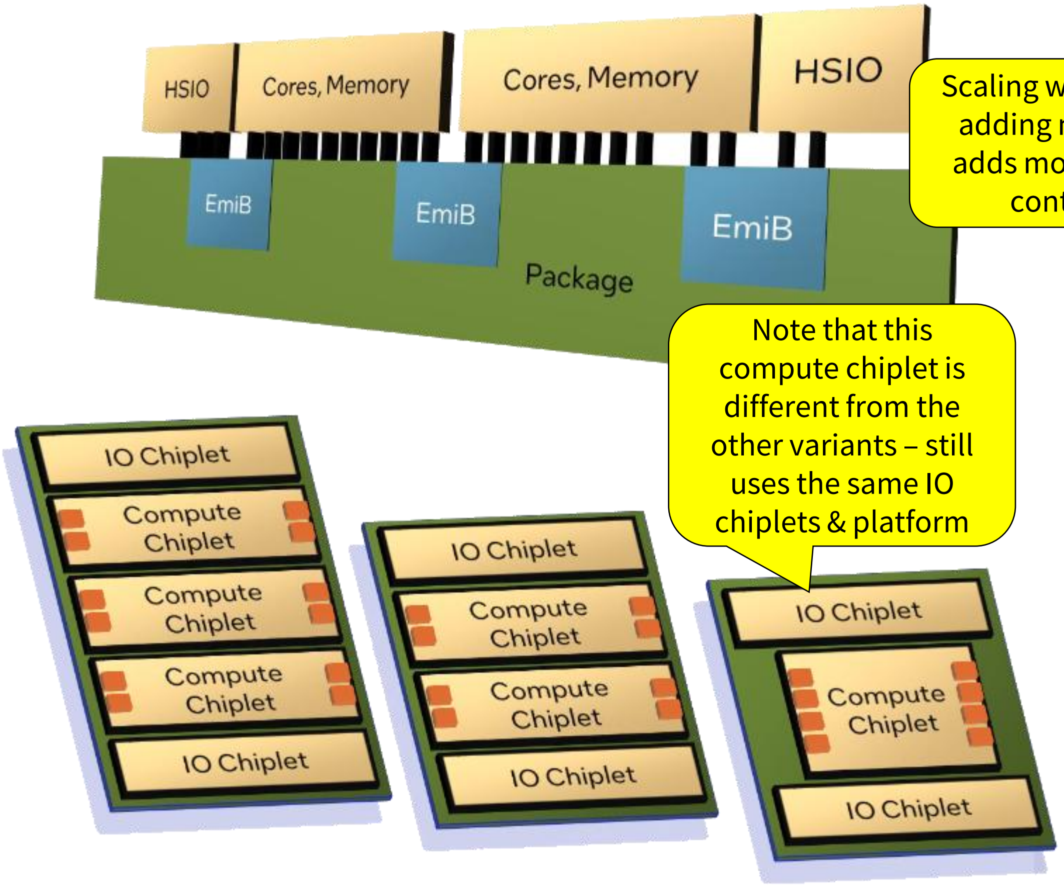

Scaling with chiplets: adding more cores adds more memory controllers

Note that this compute chiplet is different from the other variants – still uses the same IO chiplets & platform

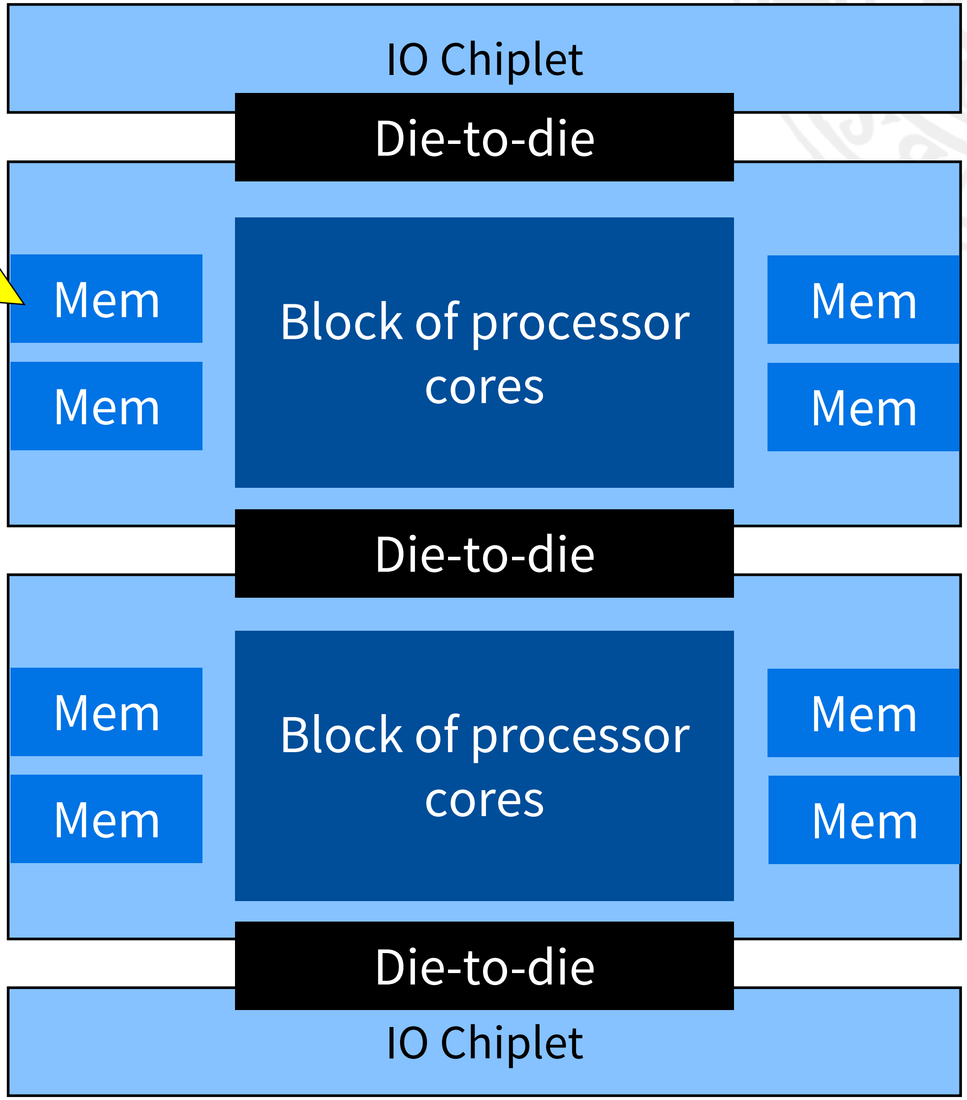

Source: https://www.intel.com/content/www/us/en/content-details/787386/hot-chips-2023-bhs-and-granite-rapids-xeon-architecture-presentation.html

Chiplets/Tiles/Dies – Tradeoffs

- Nothing is free…
- Problems **created** by chiplets

Power consumption	Performance	Packaging costs
• Communication across die interfaces use more power than inside a chip (or die)	• Added latencies between chiplets	• More complex packaging • Typically requires a base die

- In general: chiplets are good for scaling and flexibility, monolithic chips are better for power efficiency

UPPSALA UNIVERSITET

Company Choices: Chiplet vs Monolithic

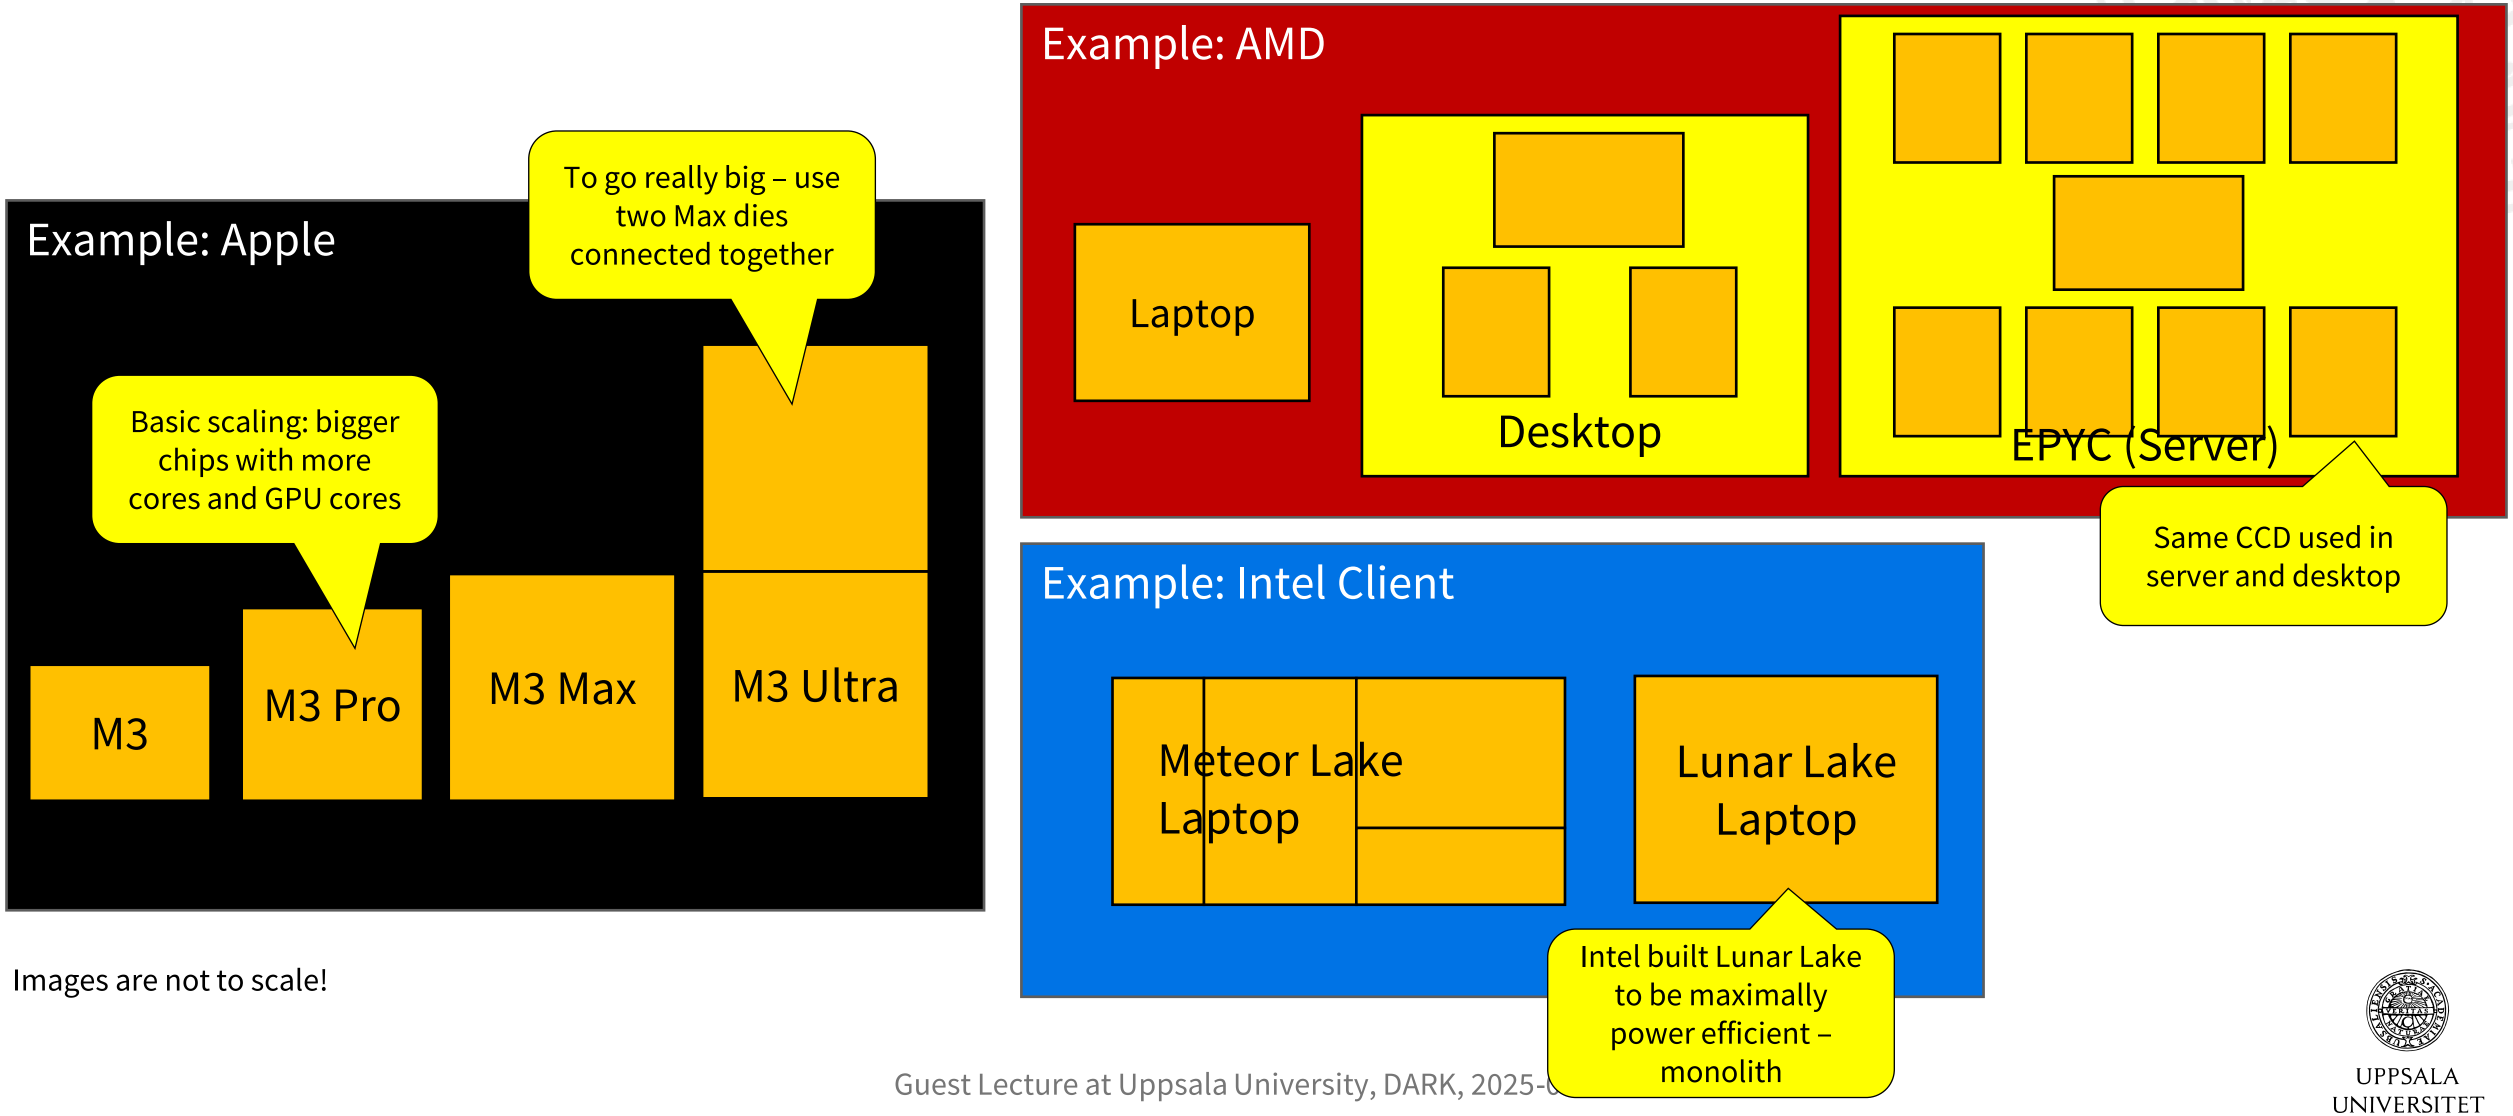

Images are not to scale!

Memory Variants – Flexibility, Performance, Power

Slotted. **Desktop**, **server**, some **laptops**. Prioritize configurability, repairability, expandability. 2 to 16 channels, 64 bits per channel.

Soldered. Standard **laptops**. Configuration on purchase. Prioritize size and cost. Lower power than slotted. 1 to 4 channels.

On-package. Premium **laptops, HPC, mobile**. Prioritize size, performance, power consumption. Expensive to produce, weak configurability. 2 to 16 channels.

HBM – on-package high-bandwidth memory. High-performance **compute**, **GPUs**, **AI** accelerators. Maximum performance at maximum cost & power consumption. Limited capacity, used with regular DRAM.

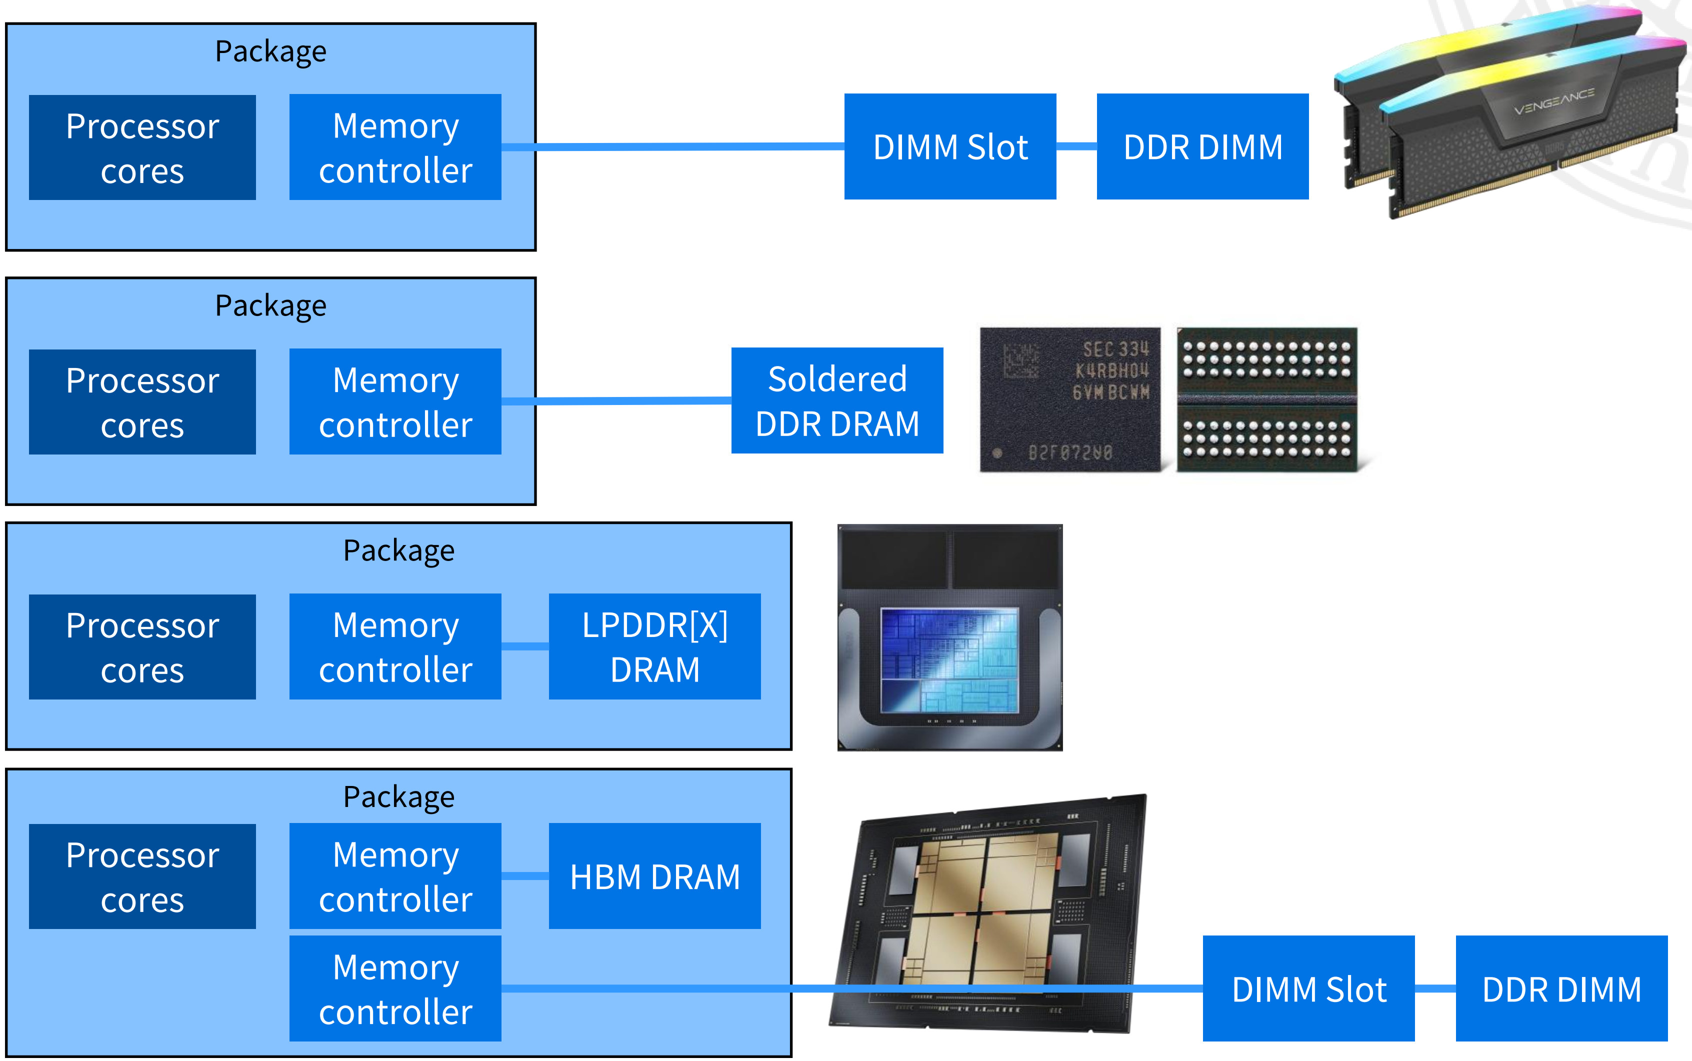

UPPSALA UNIVERSITET

Automotive Computing

Zonal Architecture

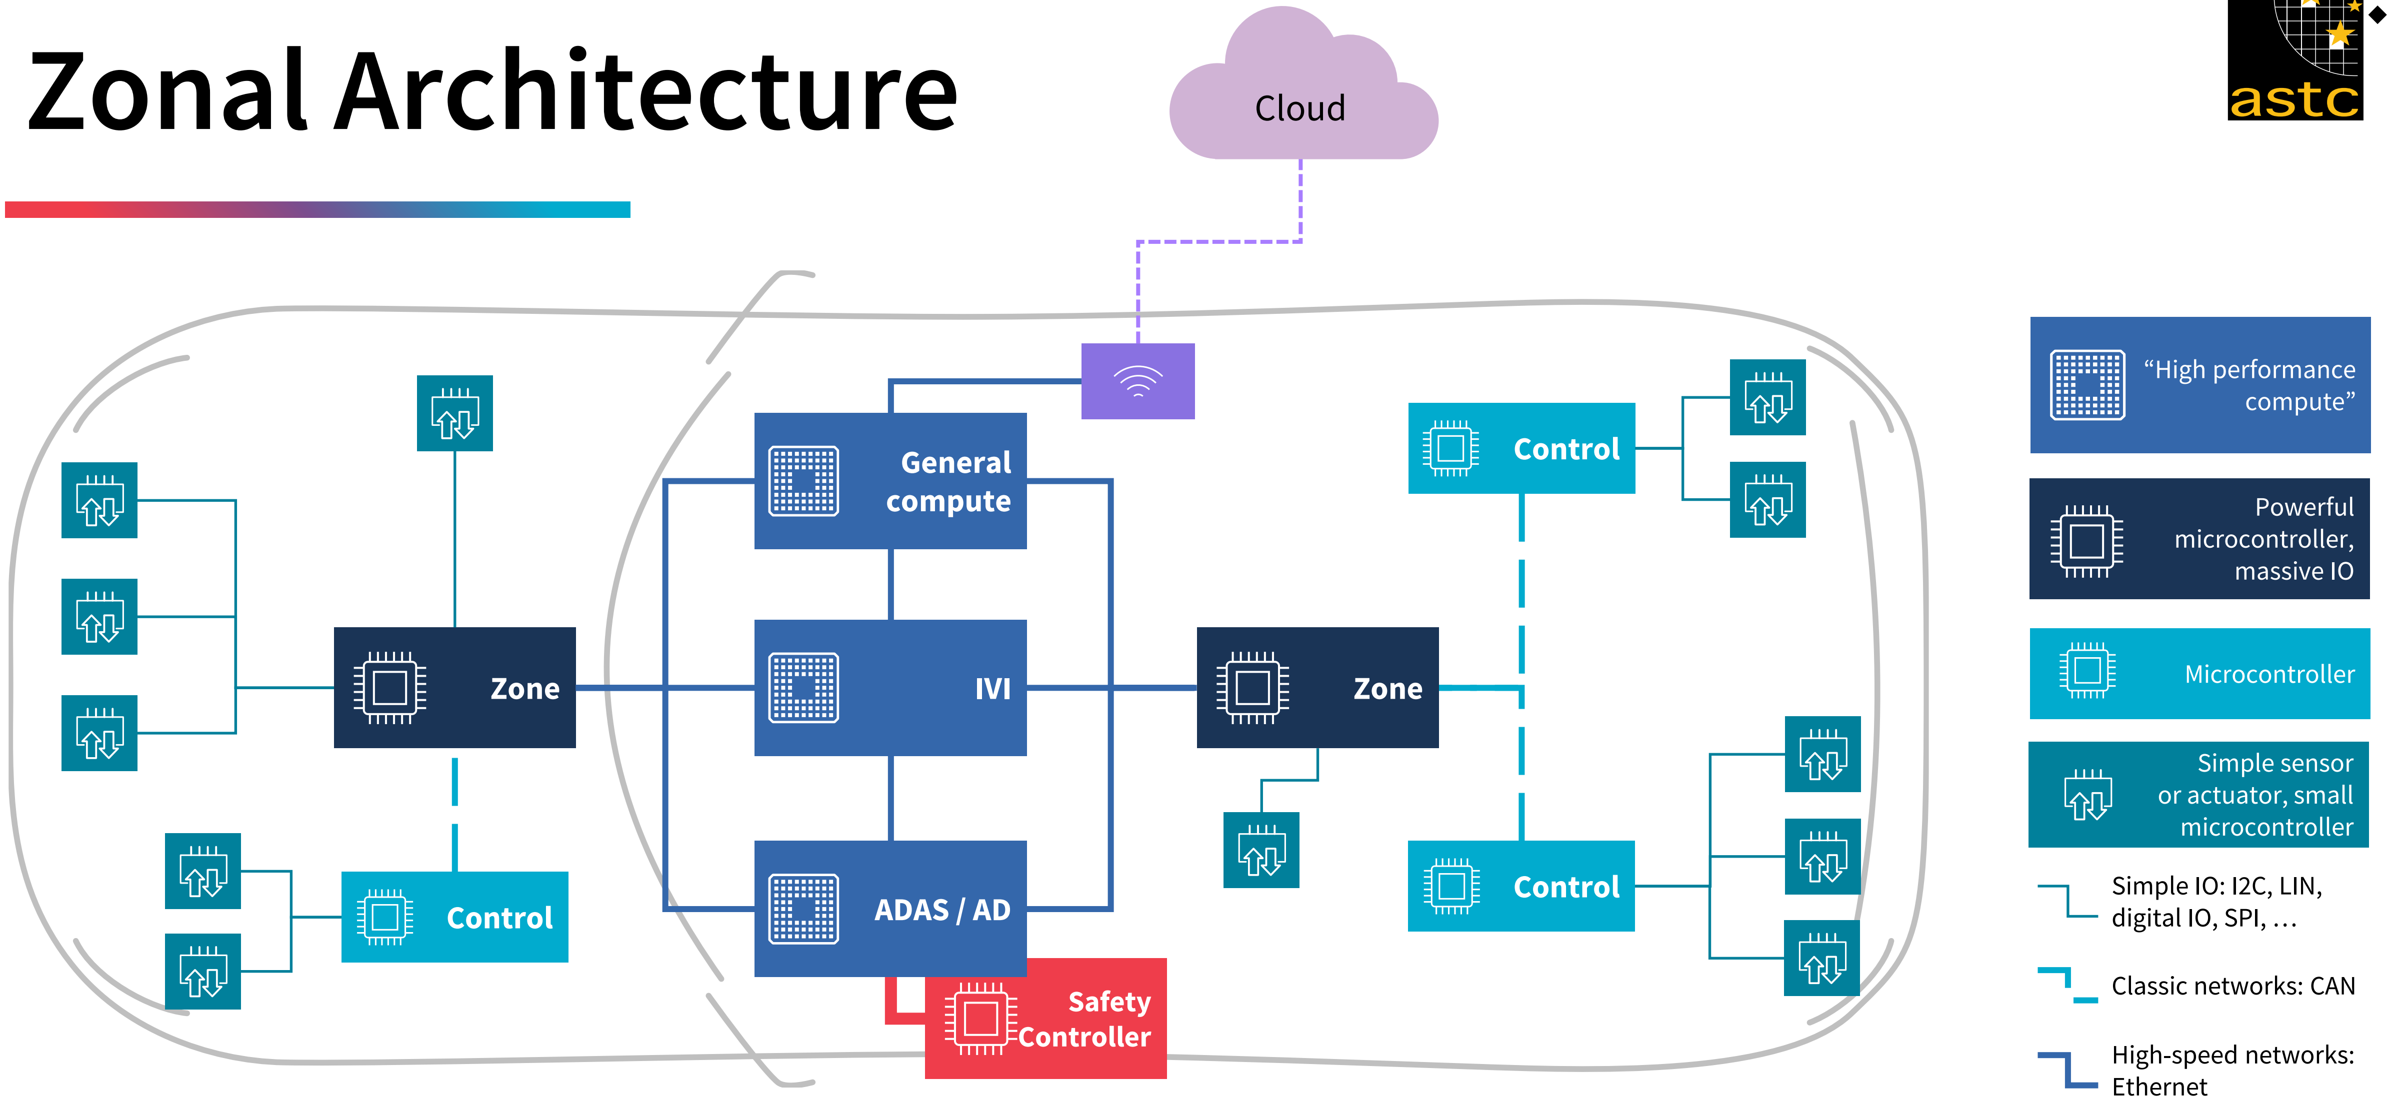

Control: engine, battery, chassis, …

ADAS=Advanced Driver Assistance Systems
AD = Autonomous Driving
IVI=In-Vehicle Infotainment

Zones: front, back, left, right, …

Automotive Controller: Infineon Aurix TC4x

- For Zonal and safety controllers, radar processing, object recognition, …
- Diverse processing
  - Main cores: 6 x **TriCore**, 32-bit RISC, 500 MHz
  - Vector engine ("PPU") (Synopsys ARC)
  - Programmable timer: GTM
  - Radar input processing
- On-chip memory
  - 25 MB FLASH
  - Closely coupled SRAM in multiple places
- Automotive networking
  - CAN, LIN, Ethernet
- Direct IO
  - PCIe (gen 3)
  - Analog/Digital converters
  - GPIO (digital lines)
  - QSPI & I2C (serial protocol to external hardware)
  - Audio inputs (for ADAS)
- Manufactured on a *28nm process*
- Security
  - CSRM security engine: ensure secure software, network traffic

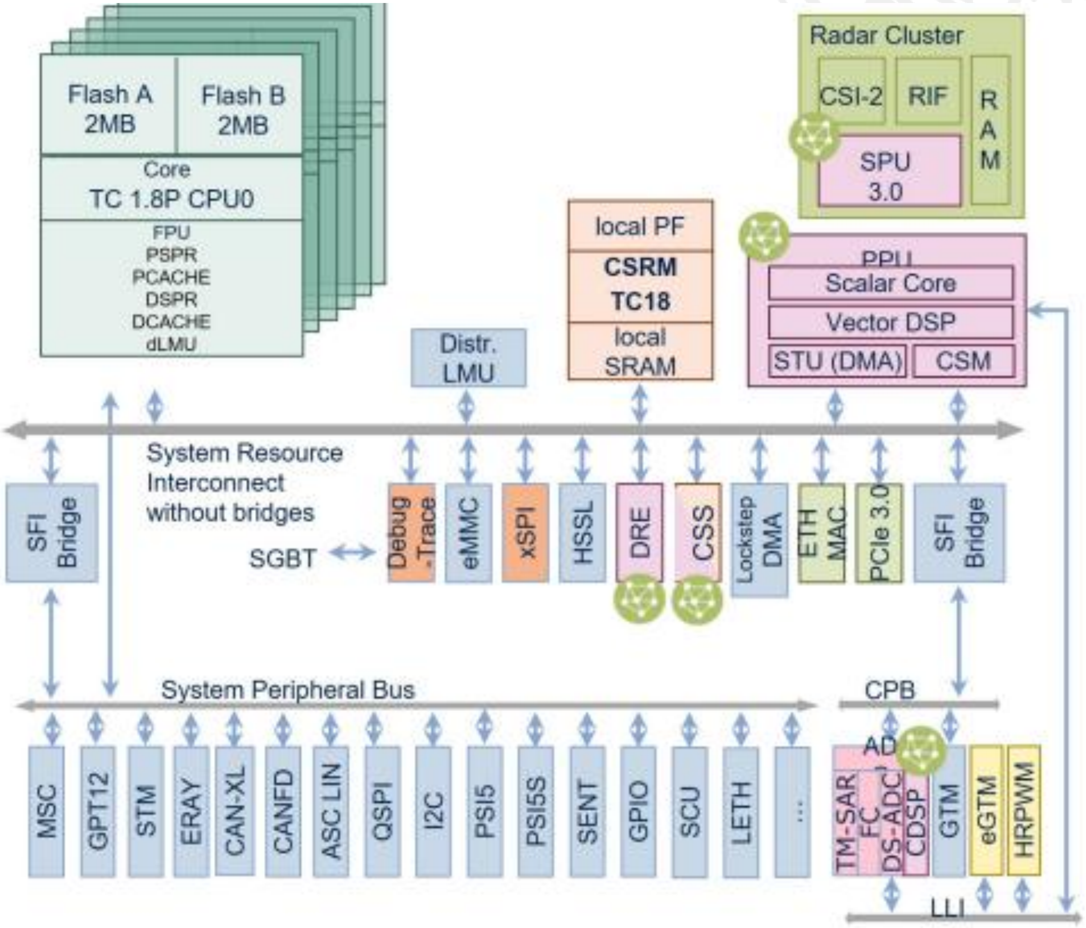

https://www.infineon.com/cms/en/product/microcontroller/32-bit-tricore-microcontroller/32-bit-tricore-aurix-tc4x/

Automotive Controller: Renesas R-Car V4H

- HPC: For Advanced driver assistance (ADAS) and autonomous driving (AD)
  - Much more compute power required
  - Manufactured on a 7 nm process

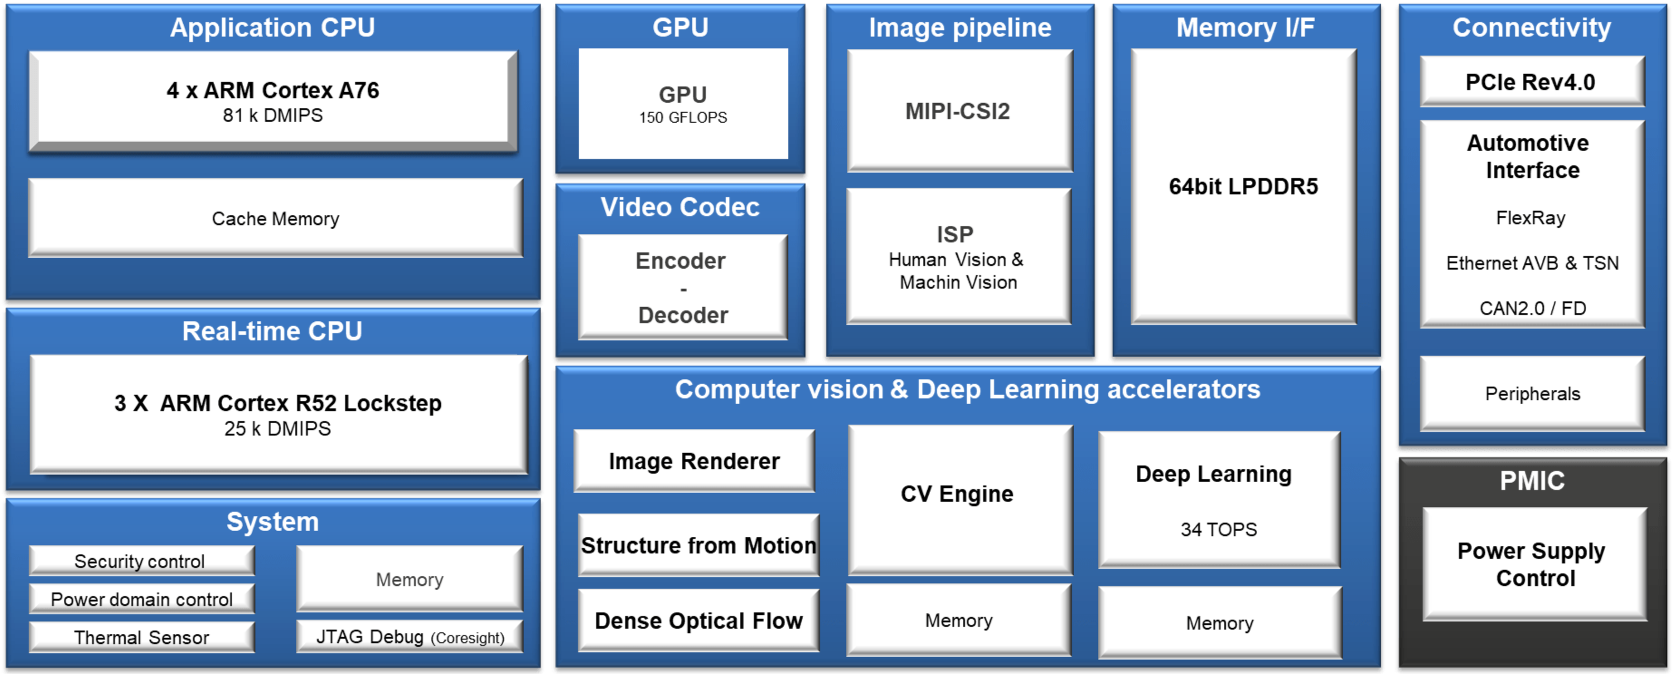

- Heterogenous compute engine
  - 4 x Arm Cortex-A78 cores, 1.8 GHz
  - 3 x Arm Cortex- R52 cores, 1.4 GHz
  - GPU (Imagination AXM-8, 150 GFlops)
  - Video in/out
  - Computer vision accelerator
  - Image renderer
    - Fisheye conversion
  - AI accelerator
    - (34 TOPS, same as CoPilot+ PCs)

https://www.renesas.com/en/products/automotive-products/automotive-system-chips-socs/r-car-v4h-best-class-deep-learning-very-low-power-system-chip-automated-driving-level-2level-3

Power Optimization & Power Management

Improving Power Efficiency

- **Higher performance** with **lower power** consumption

- For portables this means:
  - ... Lower weight
  - ... With longer battery life
  - ... A less fan noise

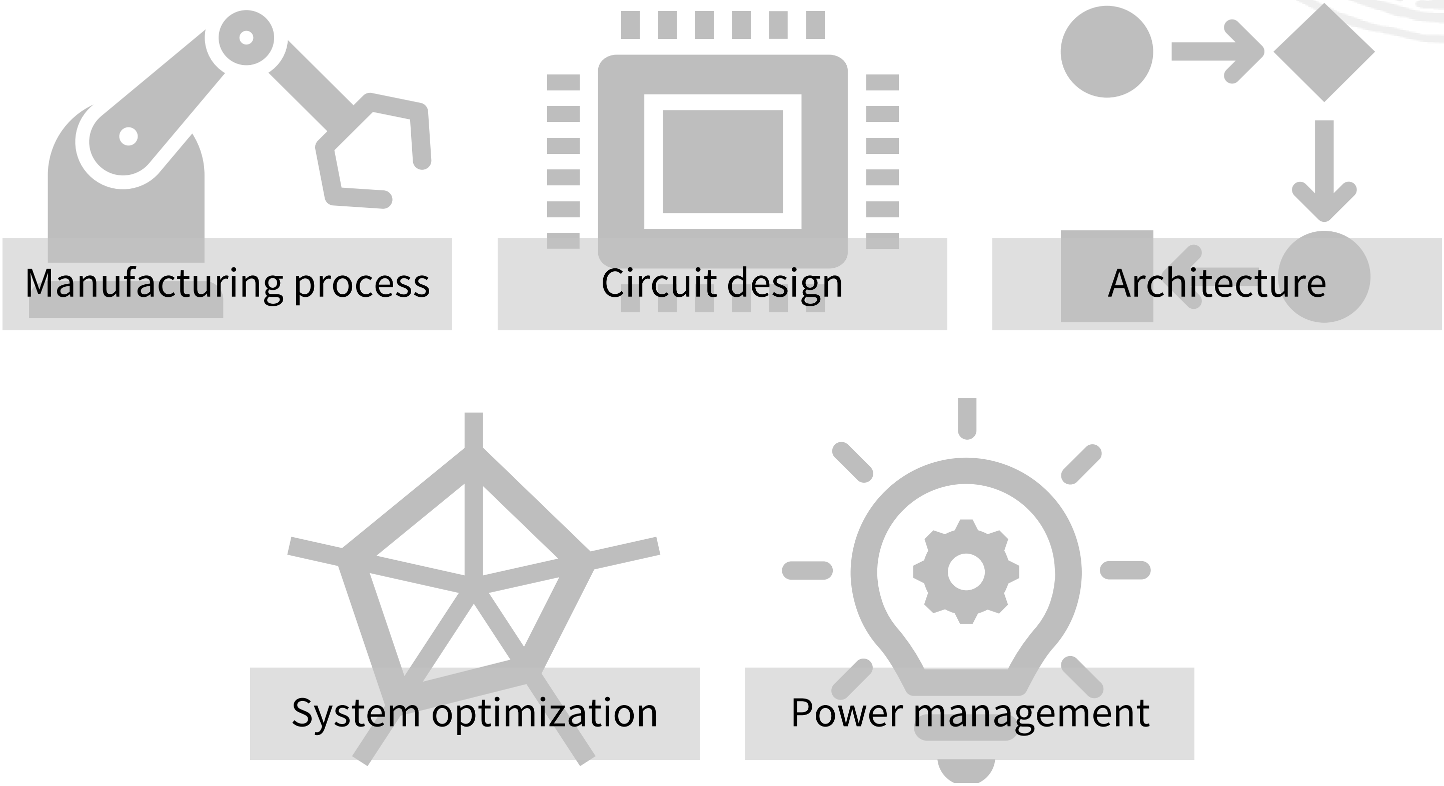

Fundamentals: Where Does the Power Go?

$$P = P_{dynamic} + P_{leakage}$$

- Dynamic power
  - Comes from transistors switching between states
  - Note: more transistors=more switching
- Leakage power
  - From just being powered-on

$$P_{dynamic} = CV^2 f$$

- Dynamic power:
  - Basic capacitance
  - × Voltage squared
  - × Frequency
  - Note that V affects f_{max}
    - Raising the frequency requires raising the voltage
- Old rule of thumb: doubling the clock frequency doubles the power

Better Process and Circuit Design

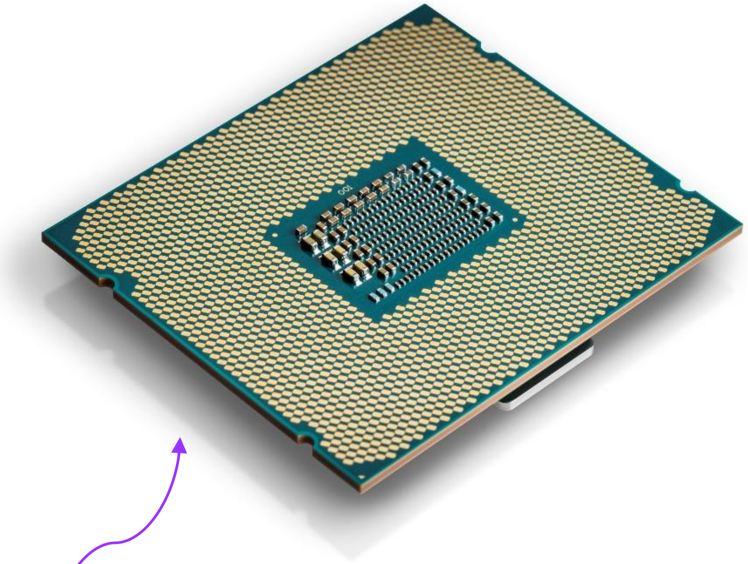

- Process technology and circuit design advances
  - Lower leakage power
  - Transistors that use less power individually
  - Lower drive voltages
    - Processors used to run on 5V, then 3.3V, now down to < 1V
      - Interesting side-effect: to power a 100W processor, we have to feed more than 100A
      - Approximately half of all "pins" on a chip package are for power distribution
  - Faster switching between different power modes and operating frequencies

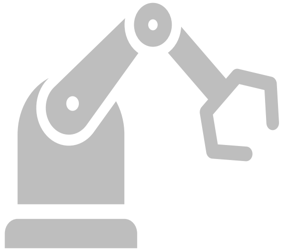

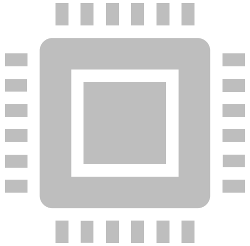

- All things equal, the same design on a better process gets
  - Lower power at same frequency
  - Higher frequency at the same power

Better Computer Architecture

- Reduce "wasted power" in the chip
  - Clock gating – shut off *clock* to subsystem
    - Removes dynamic power
  - Power gating – shut off *power*
    - Removes static power (leakage)
  - Over time, gating gets applied to smaller parts
- Adding more power states:
  - Granular control over activity
    - Settings for frequency, voltage, on/off, …
  - Subsystems are set to lowest possible state
  - Increasing the number of controllable units and the number of steps
- Processor core and pipeline design – trade max performance vs power
  - Different core designs hit different trade-offs
  - Many slow cores, a few fast cores, or a mix?
  - *(As discussed earlier)*
- Cache and memory system
  - Cache hit = lower power than memory access
  - Faster (external) RAM costs more power
  - Bigger caches reduce overall power consumption – but increase the cost of the chip
- Use special-purpose accelerators
  - Can do the same work at orders of magnitude less energy

System Optimization and End-Device Design

- Overall system design choices
  - Display size, technology, resolution, update frequency, brightness, …
  - Battery size
  - Choice of processor variant & tuning
    - Different manufacturers tune differently
  - Memory choice
    - LPDDR (Low-Power DDR) vs regular DDR vs on-chip memory, speed rating
  - Speed of wireless functions
  - Cooling efficiency
  - How many ports, how fast?

- Not easy for an end-user to grasp all the details

Certifications like Intel evo: provide consumers with an indicator to look for, and provide computer makers with advice to design better systems

Power Management (Software)

- Power efficiency does not only depend on static hardware factors
- Power management in software:
  - **Control feedback loop** implemented in hardware, firmware, and software – driving power states

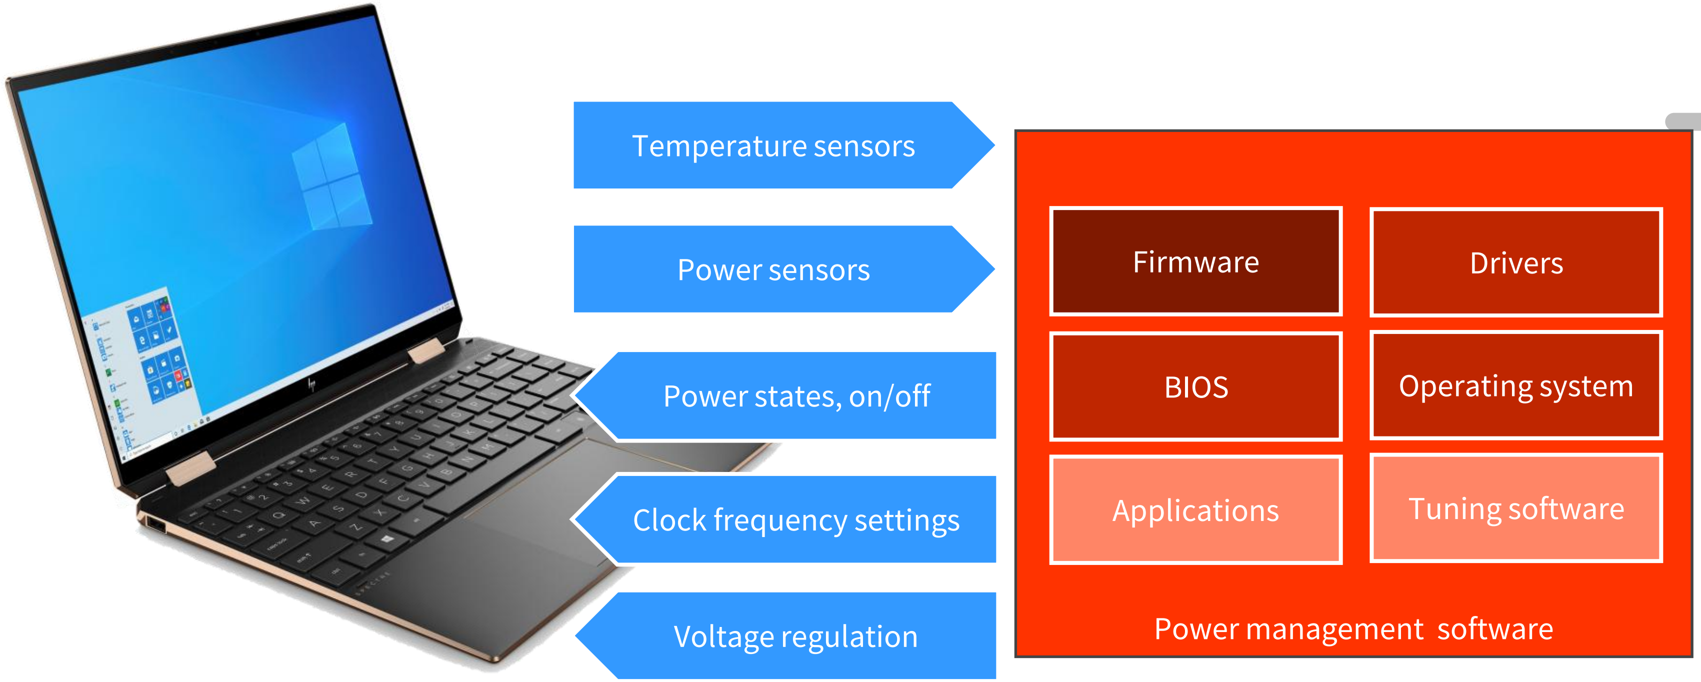

UPPSALA
UNIVERSITET

Power Management: Hardware Control Points

- Hardware provides proliferating control points:
  - *Per-core* voltage and clock-frequency adjustments (used to be per chip)
  - More power states, in more devices
  - Faster changes between power states
  - Faster frequency scaling:
    - Note that going to another power state is not free – it takes time to effect
- Sensors multiply across the chips and system
  - Power levels
  - Thermal sensors – very important limiter for performance
- All of which come together in the power management stack
  - Typically, one central control unit + many small control units spread across the chip

Power Management: Software and Firmware

- Optimize settings for current load
  - Profile current load and usage
  - Balance power draw vs user experience
  - Set power/performance operating points
  - Allow higher performance if temperature and power availability allows it
  - Select the right core to run a workload on for optimum results

- Sleep & wake-up
  - Put system into deeper sleep
  - Wake up when user comes back

- Avoid disaster
  - Throttle to avoid drawing too much power from the platform
    - Each chip has a design limit the rest of the computer expects it to adhere to
  - Throttle to avoid overheating the chip

Example: Overheating (Almost) with a Dead Fan

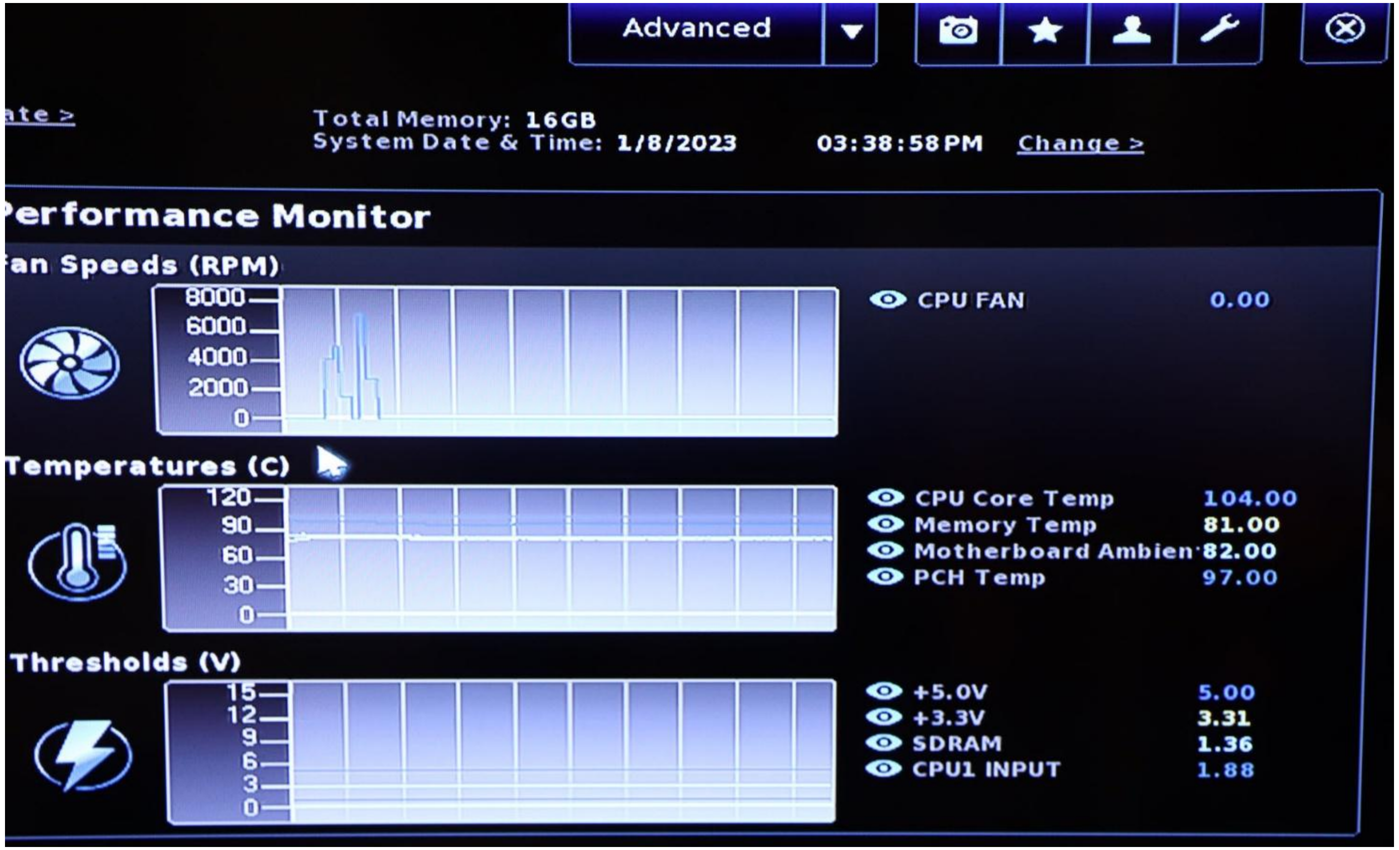

Photo of the BIOS performance monitoring screen

Fan speed spikes: from blowing compressed air into the system – speed sensor still worked

- Fan broke on my old NUC
  - The machine had to be replaced
  - https://jakob.engbloms.se/archives/3711
- Interestingly:
  - The processor still ran!
  - Chip temperature at **104° C**
  - Power management:
    - Throttled the processor down to <1GHz clock

UPPSALA UNIVERSITET

Power Management: Max is not Sum of all Max

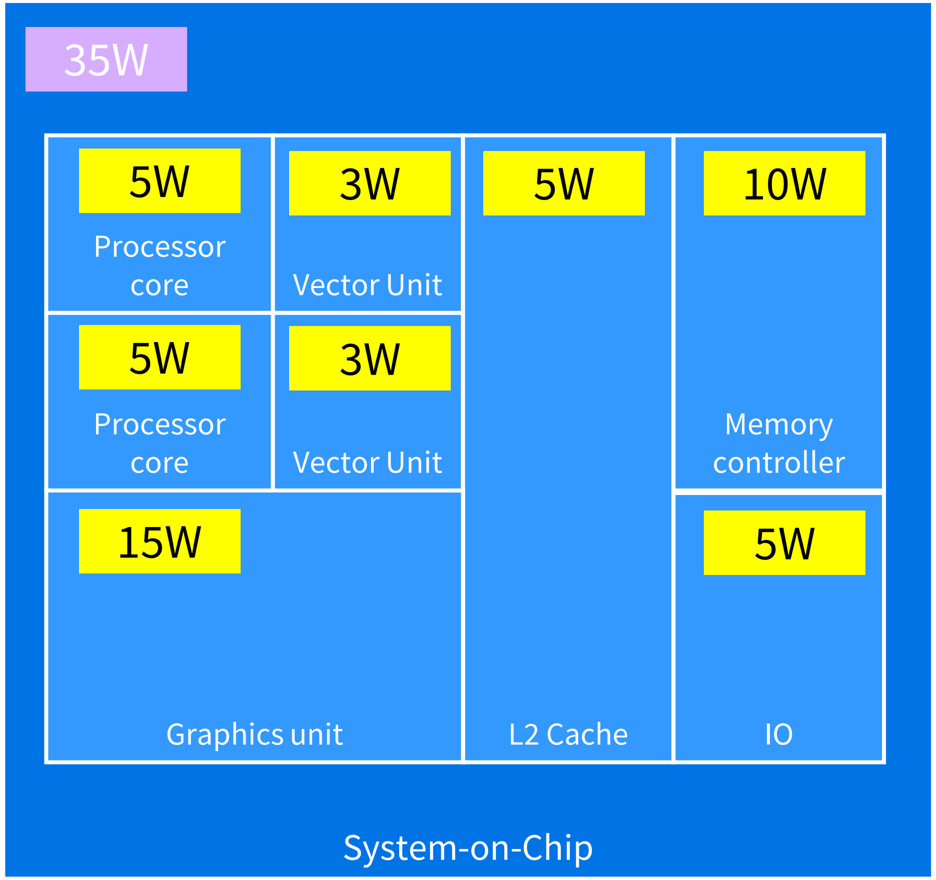

Hypothetical chip, rather simplified

- *Fictional example for illustration*
- Total chip power allowed = 35W
  - Dictated by heat sink, power supply, and market segmentation
- Total max power = 51W
  - Throttle one part of the chip to allow others to run at full speed
- Power management needs to keep the power inside allowed bounds

Power Management: Set According to Workload

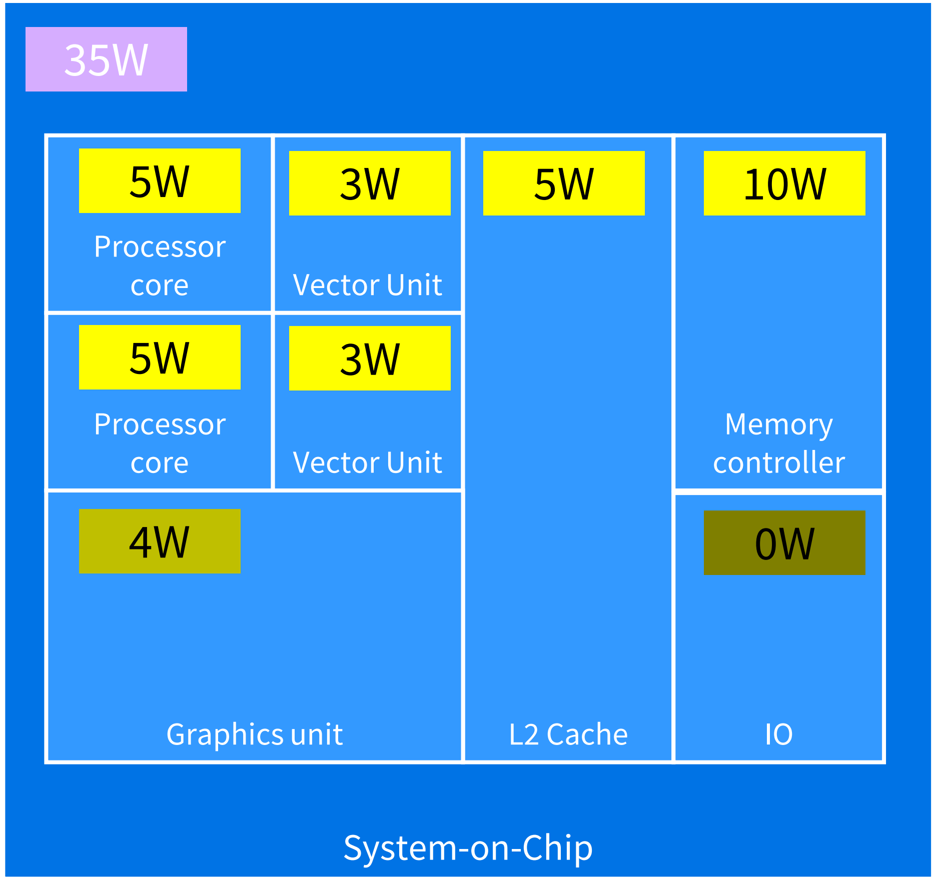

Hypothetical chip, rather simplified

- Compute-focus:
  - Power up cores, memory, and vector units
  - Throttle graphics to make room
  - Turn off IO, we assume we run from memory

UPPSALA UNIVERSITET

Power Management: Set According to Workload

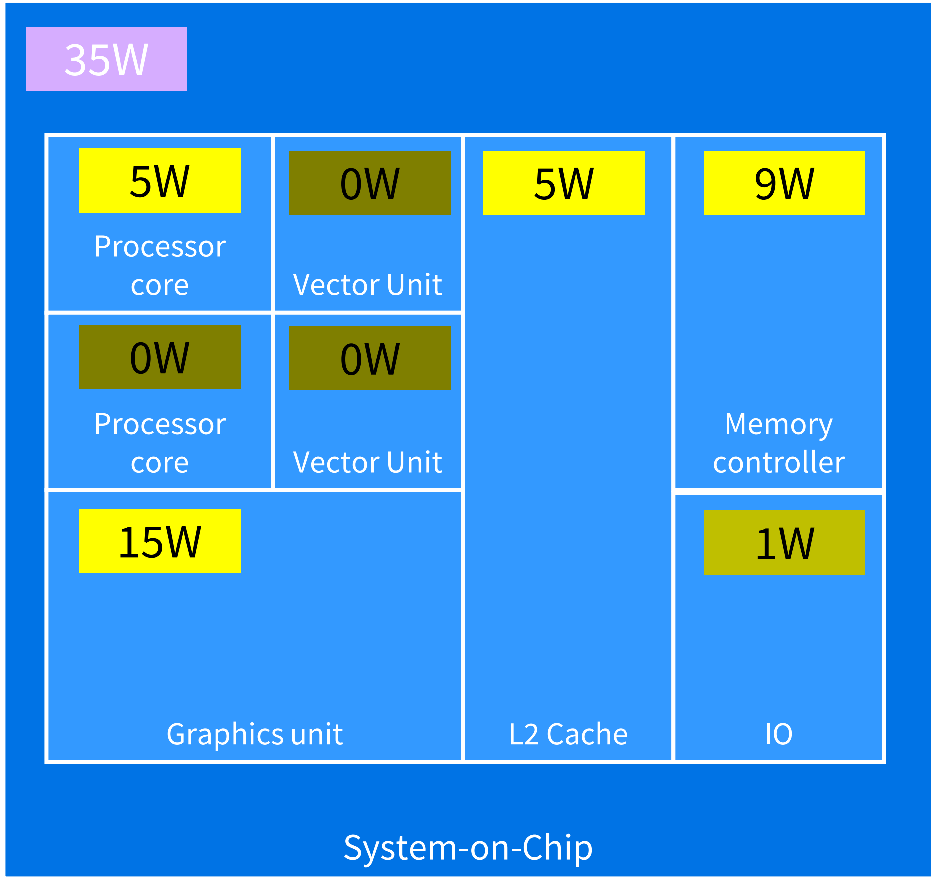

Hypothetical chip, rather simplified

- Gaming:
  - Graphic processing take priority
  - Run one processor core at full speed – latency matters more than throughput
  - Disable vector units – such work is now on the graphics unit
  - A bit of IO needed for sound and chat
  - Memory controller busy - but cannot be given full power since that would exceed the global limit
- Setting the trade-offs right is tricky
- Actual performance can be very different from theoretical peak performance

Heat: Limiting Factor

Noctua* DH15S cooler, image from
https://noctua.at/en/press-images/NH-D15S

Practical tip: If a computer does not seem to run optimally: check and clean fans and airflow

- Heat removal is often limiting practical performance
  - Power can be fed, but the resulting heat cannot be removed at the same rate
- Example:
  - Nvidia A100 SXM4 die could generate up to 1100W is all transistors switched at the same time
  - Limited to 400W in designs due to aircooling
  - =65% of the silicon is unused at any point in time
  - https://www.ri.se/en/news/blog/generative-ai-must-run-using-liquid-cooling
- **Cooling** a crucial part of system design
  - Heat sinks
  - Fans and airflow
  - Liquid cooling
  - Rule of thumb: cool 1W from 1mm^2
- Note: RAM and SSDs use power and need cooling too

Power Consumption Examples

- Maximize power/performance
  - Add specific accelerators
  - Design for efficiency
  - Stay within given envelope
- Maximize performance (at all costs)
  - Add more cores
  - Add more dies with more cores
  - Use higher clocks
  - = increase power to increase performance
  - Increasingly common in datacenter AI

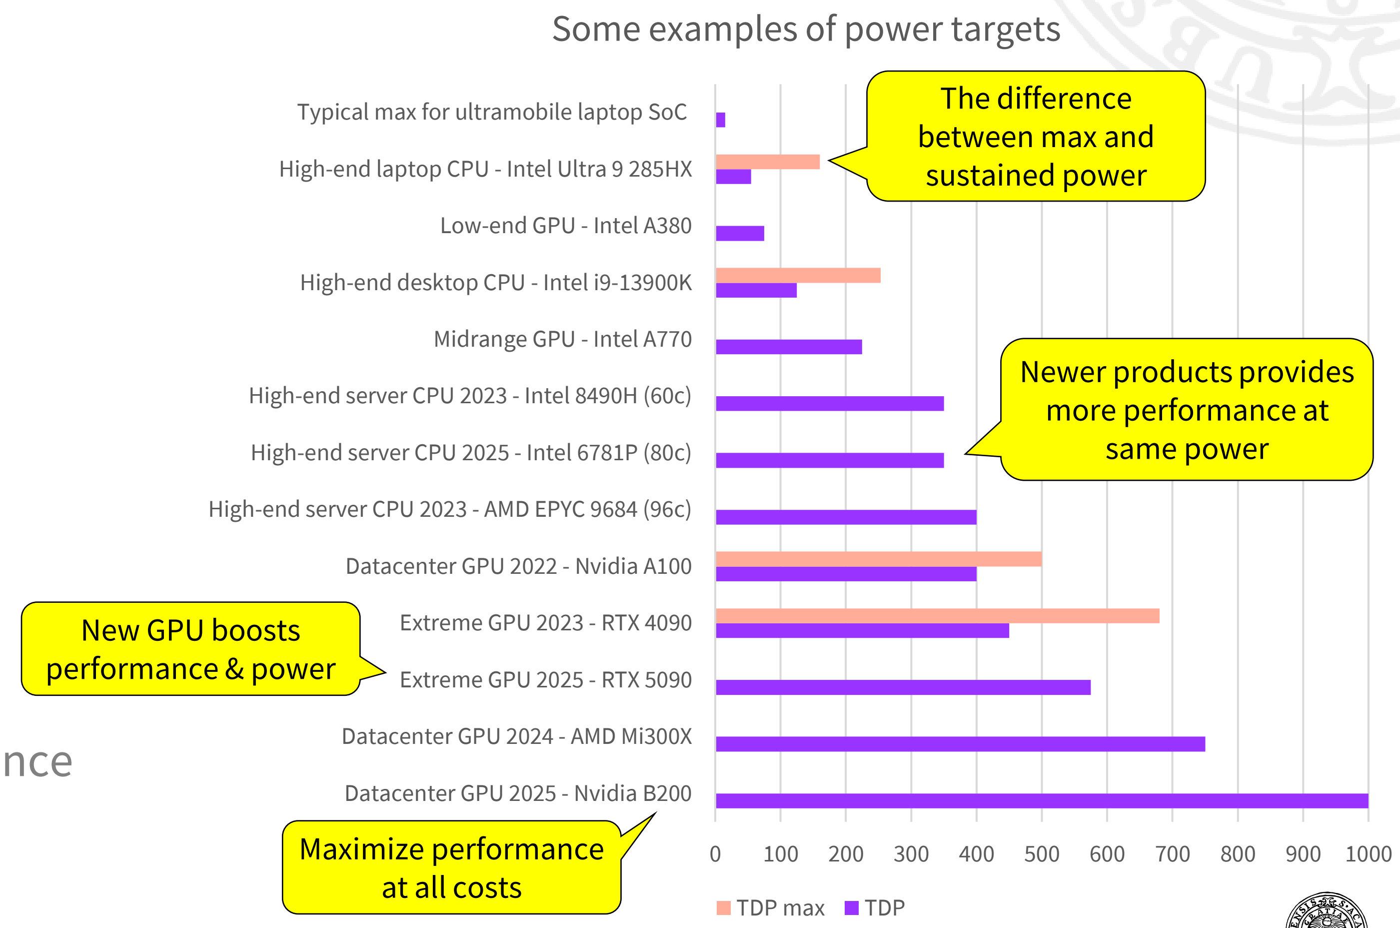

VLAB WORKS™

astc

The ASTC VLAB™ Simulator

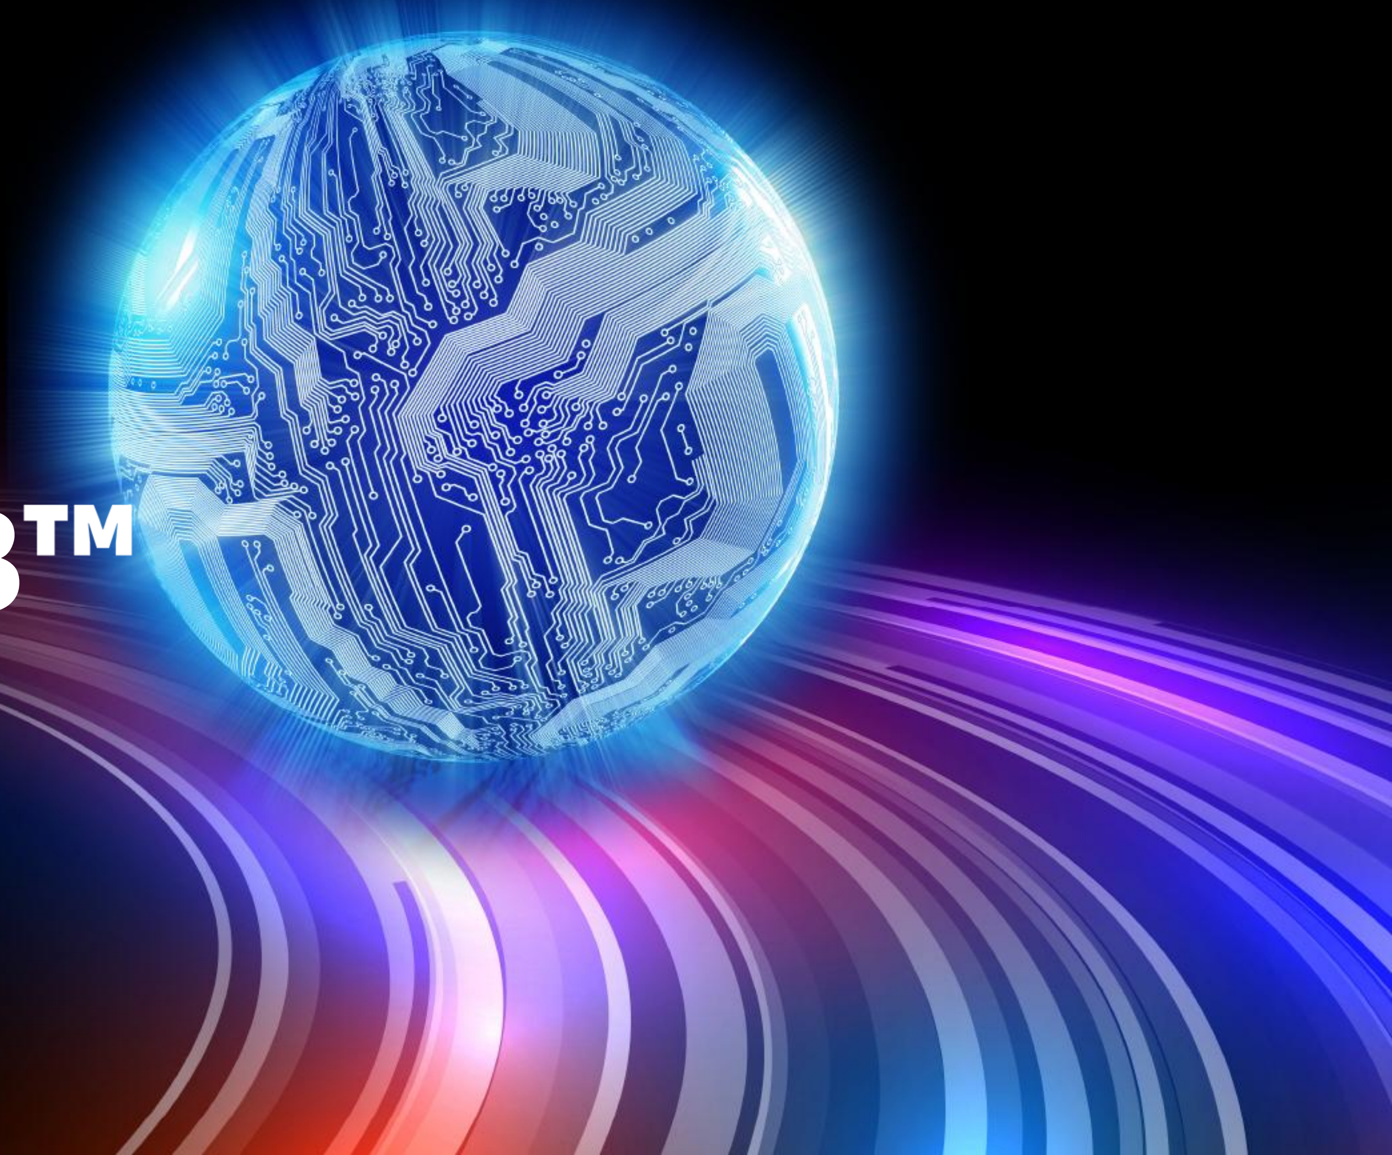

Also known as my day job

ASTC Globally

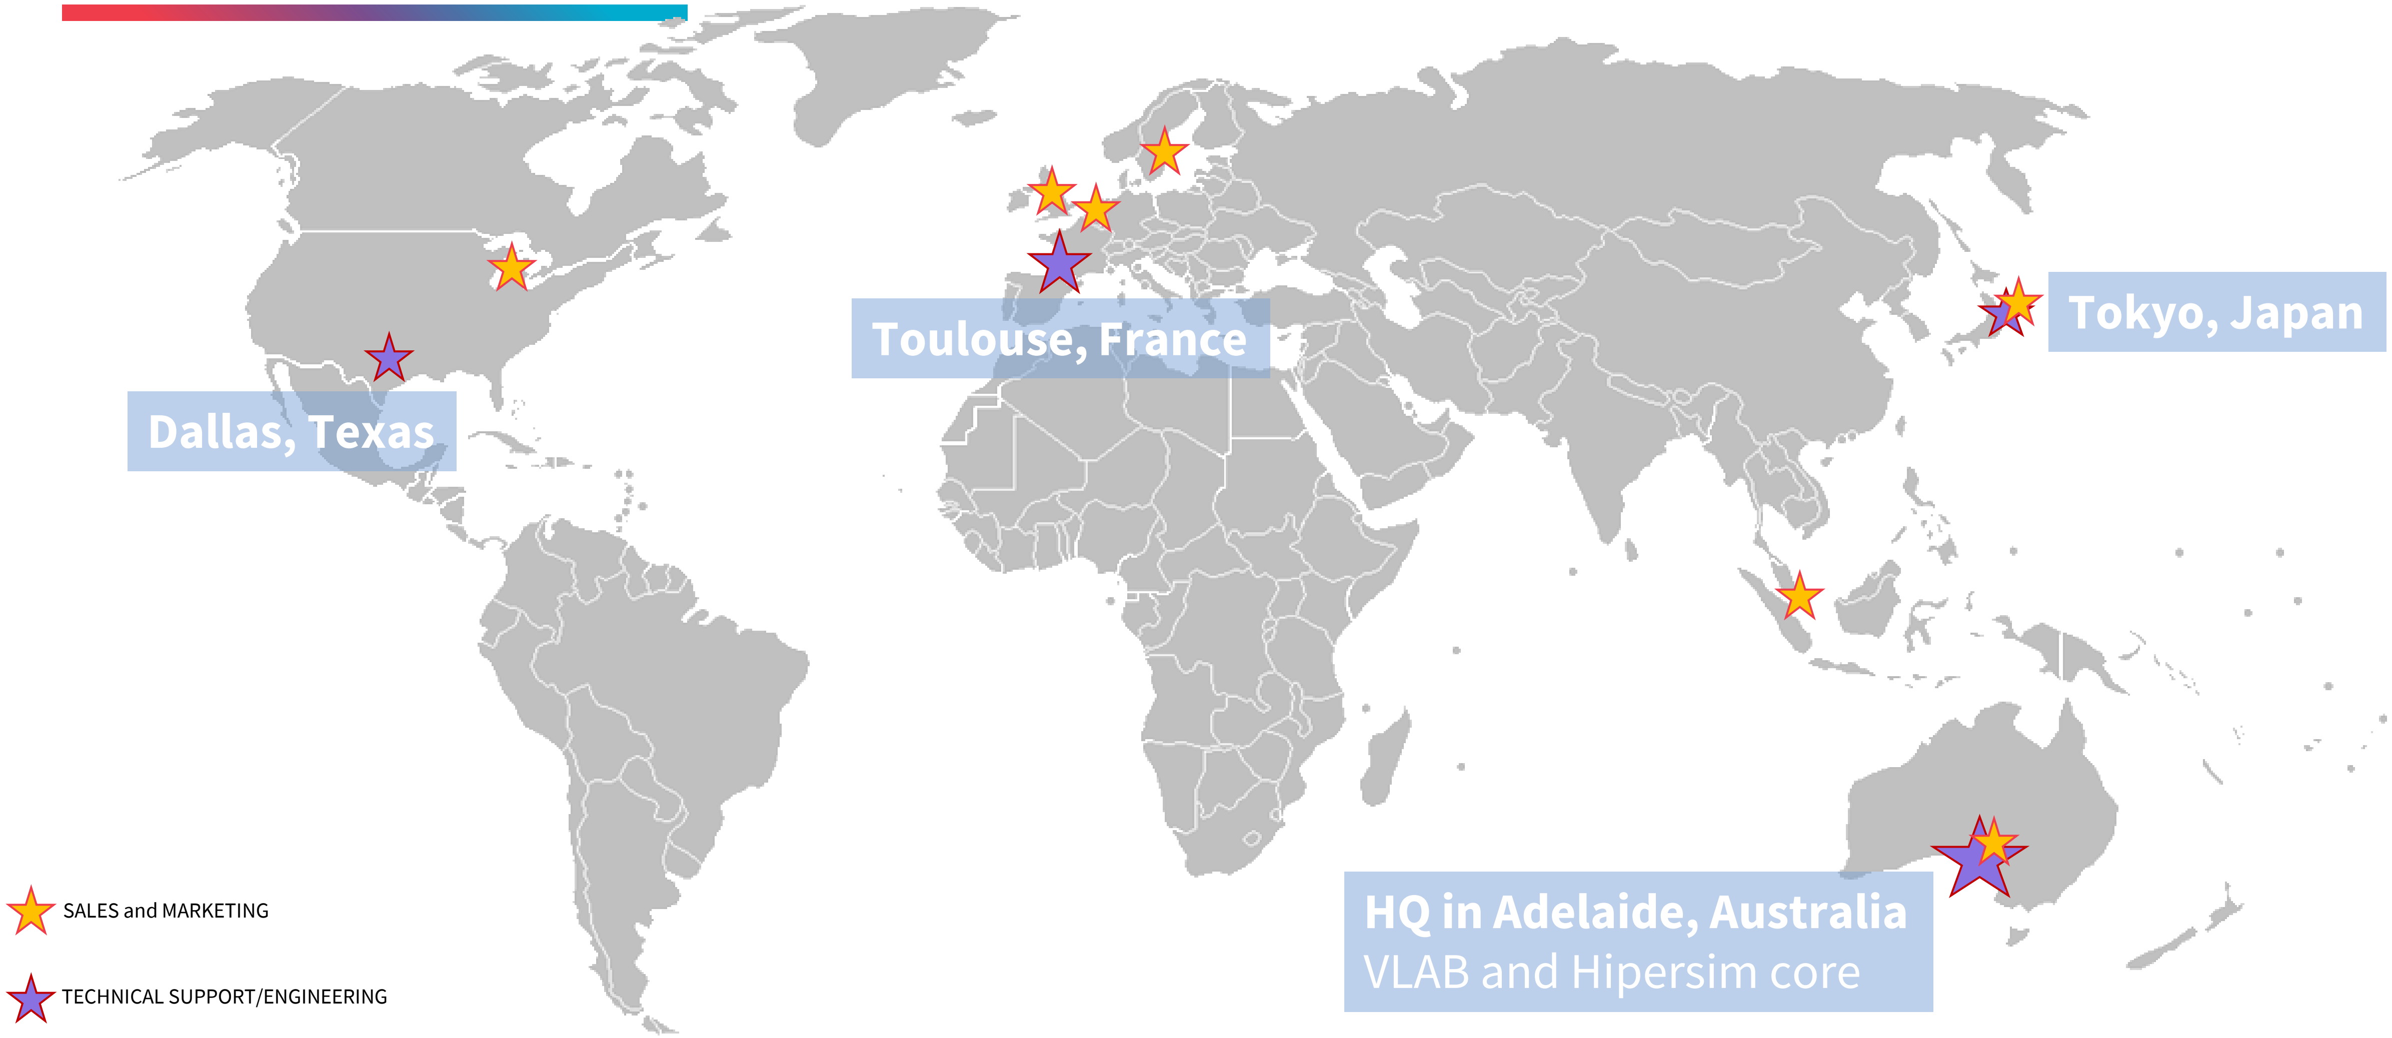

What is a VLAB™ VDM?

Virtual Development Machines (VDMs) are virtualizations of embedded hardware devices, running on a laptop, desktop, server, or in the cloud.

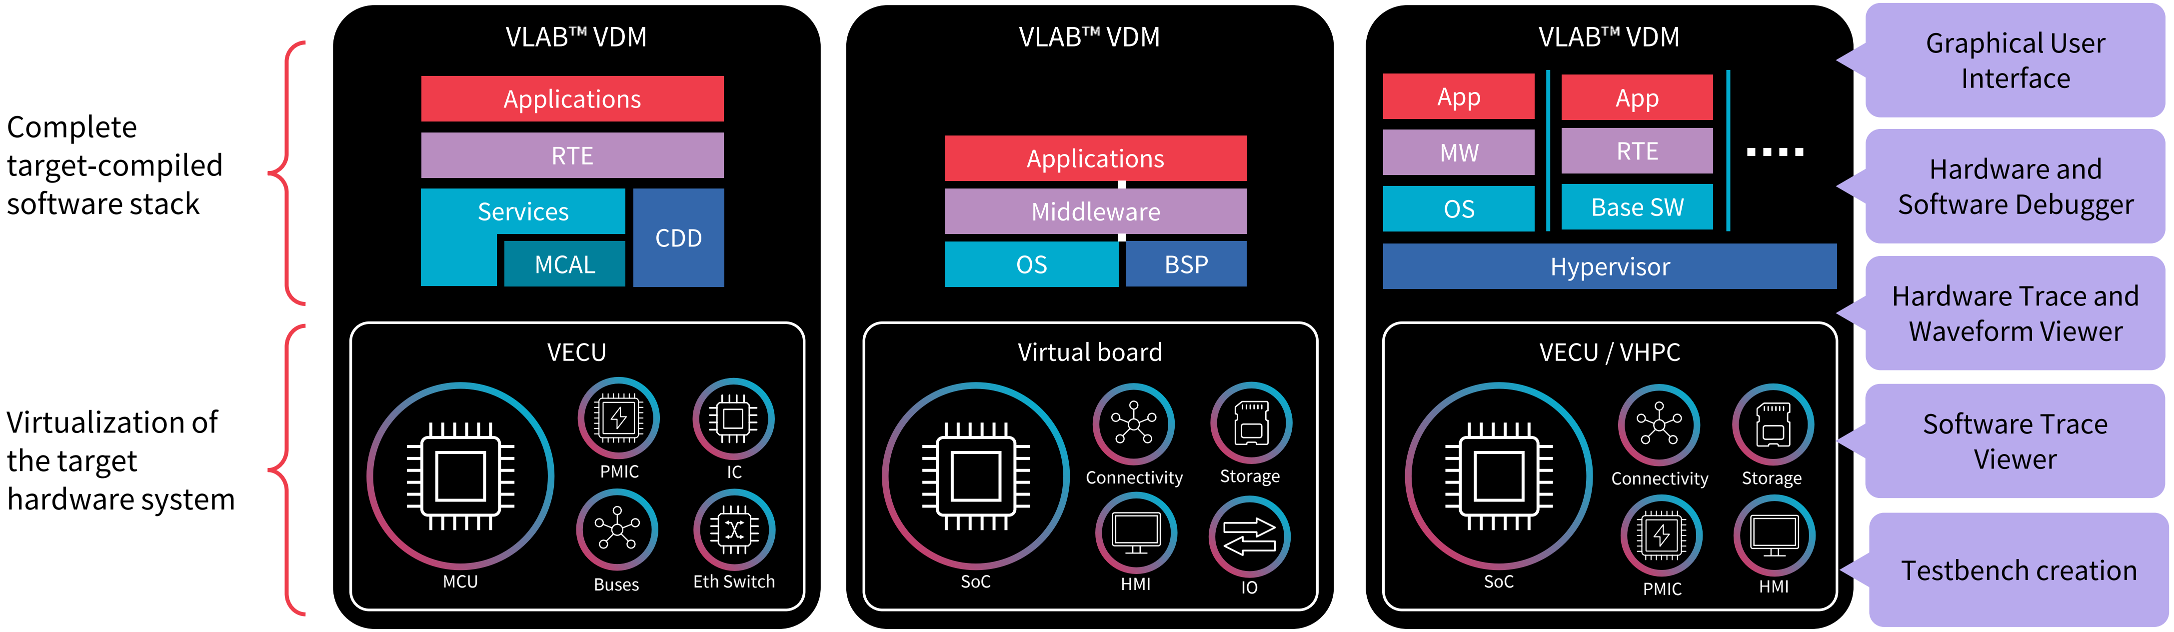

- Industry leading, high-performance simulation of embedded hardware
- Scalable, on-demand capacity for CD/CI/CT flows
- Access for local and global development teams

Why use simulation for software development?

Simple: Embedded Hardware is Difficult

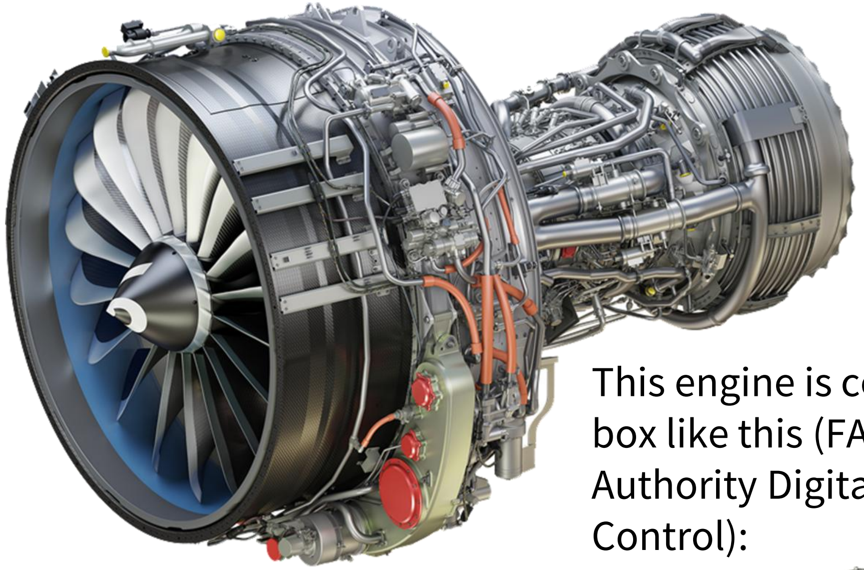

This engine is controlled by a box like this (FADEC, Full Authority Digital Electronic Control):

- Physical test systems:
  - Expensive
  - Limited in number
  - Difficult to use – built to run, not to develop
- The operation depends on the environment
  - Control systems need something to control
  - Testing "in the real world" can be interesting…
- Specialized computers
  - = cannot use a standard computer as a stand-in
- Limited access
  - Software load = updating FLASH
  - Limited memory and storage for logs & trace
  - Debug ports rare or non-existent
    - And requires special interface hardware
  - No user interface
    - Maybe not even a serial port or network login

https://www.mtu.de/engines/commercial-aircraft-engines/narrowbody-and-regional-jets/leap-1a/-1b/
https://www.safran-group.com/products-services/fadec-leap-pp20-full-authority-digital-engine-control-unit-leap-and-passport-20-engines

Embedded Hardware is Difficult

https://jaycarlson.net/2021/09/18/juggle-embedded-projects-home-office-workspace-tour/

https://jaycarlson.net/2021/09/18/juggle-embedded-projects-home-office-workspace-tour/

Free Developers from Hardware Limitations

Unlimited access to targets for testing - use any PC, server, or cloud VM to run tests

Apply **any test configuration on-demand** – not limited by availability of physical test rigs

Configure the target beyond what available hardware allows

Construct **arbitrary network topologies**, without physical limitations

Load software to the target system instantly – no complex flashing or download flows

Automate any action in the target—the simulator has perfect insight and control

Inject faults and provoke **boundary** conditions for better test coverage

Script target actions, irrespective of target user interfaces

Make Debugging Easier

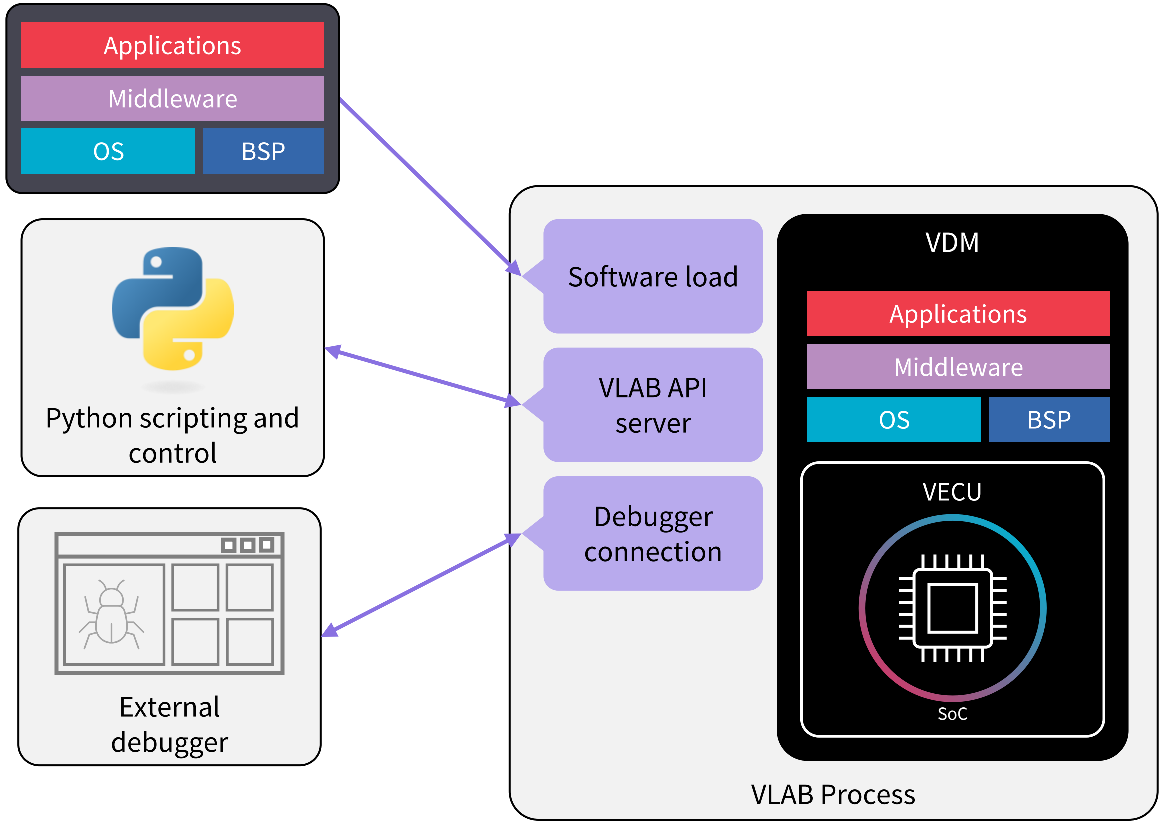

- Unlimited “hardware” breakpoints
- Breakpoints on hardware events
- Complete system inspection
- Unintrusive trace and debug
  - No timing disturbance
  - No software instrumentation/changes
- Deterministic re-execution
- Global system stop
- Script debug
  - From debugger
  - From Python (for simulator-specific actions)
- Quickly iterate new software builds

Virtual X-in-the-Loop Testing

- Test without the hardware
  - Mirror real-hardware setups
  - Simulated input/output
  - Real software stack
- Virtual system X-in-the-loop
  - Software-in-the-loop
  - Simulation-in-the-loop
  - Hardware-in-the-loop
    - Connect to the real world
  - Real-network-in-the-loop
- Attach test and calibration tools to the virtual platform

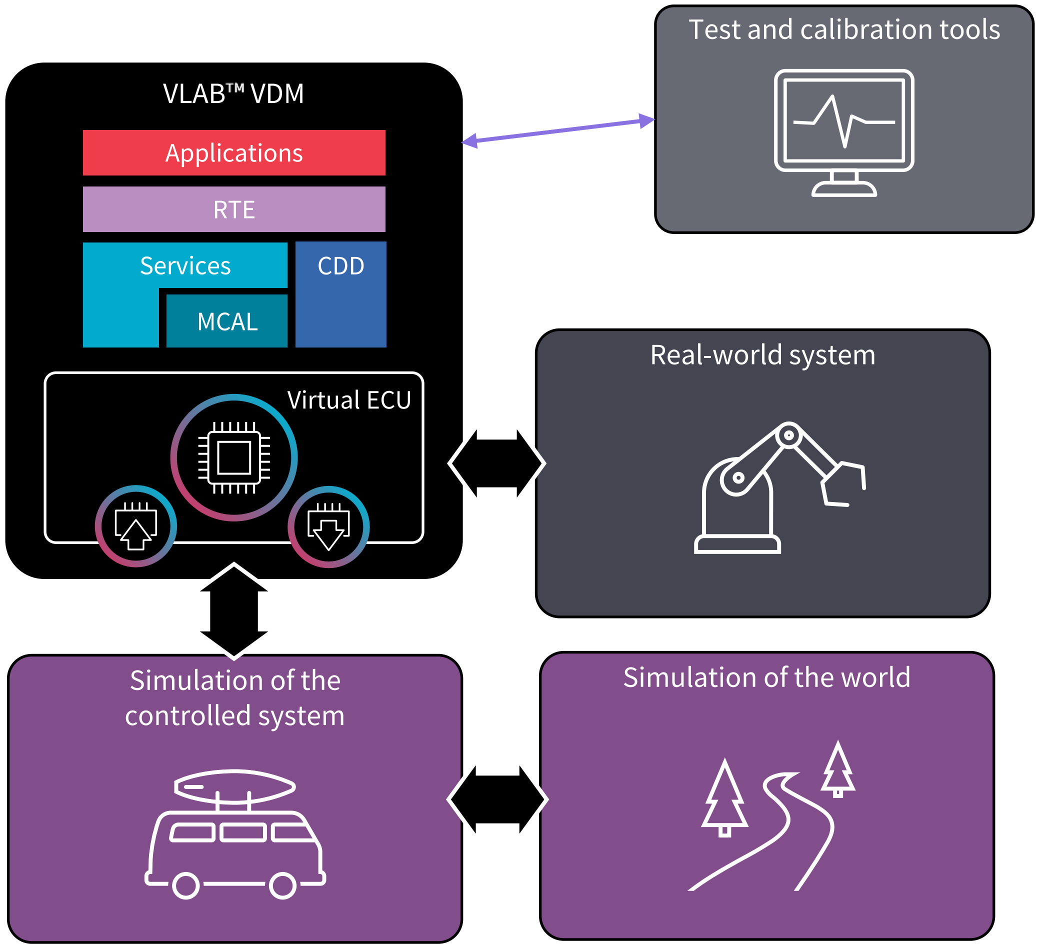

System-Level Virtualization: Digital Twin

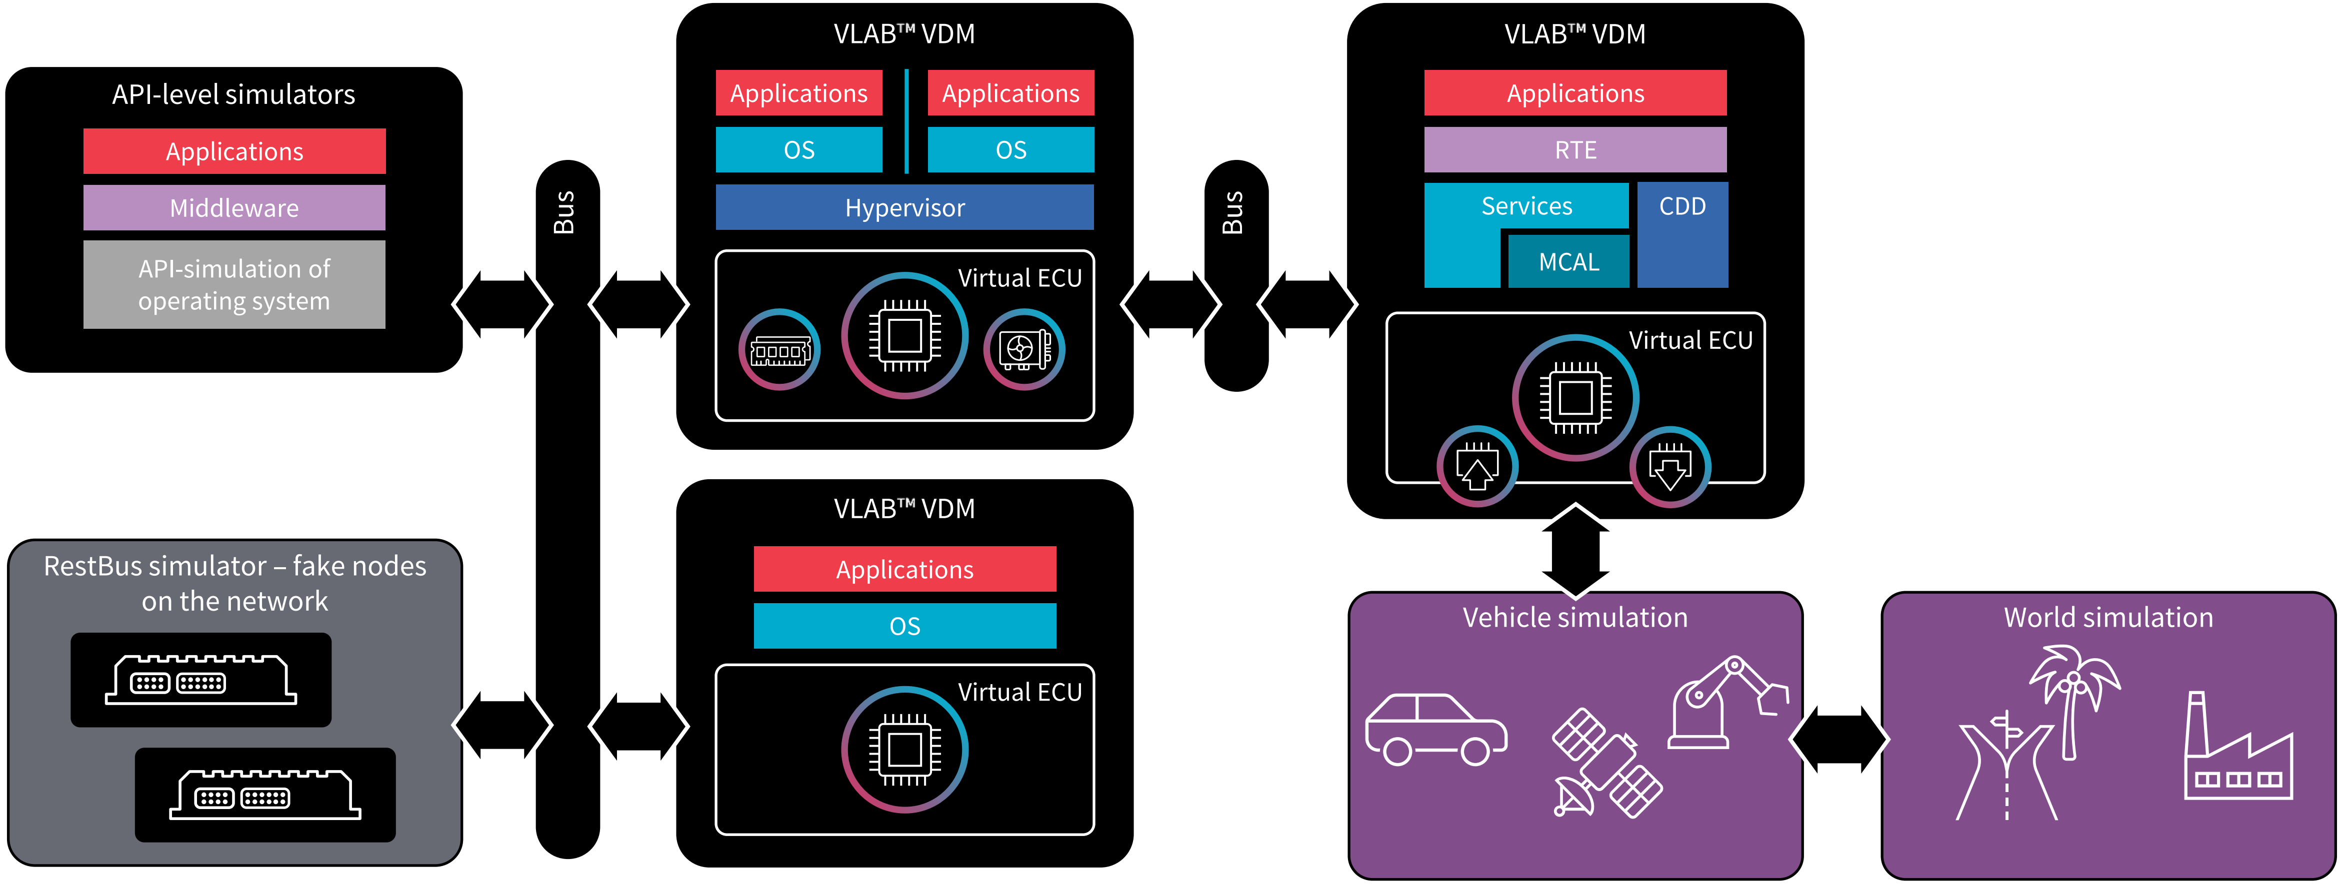

Continuous Integration and Deployment

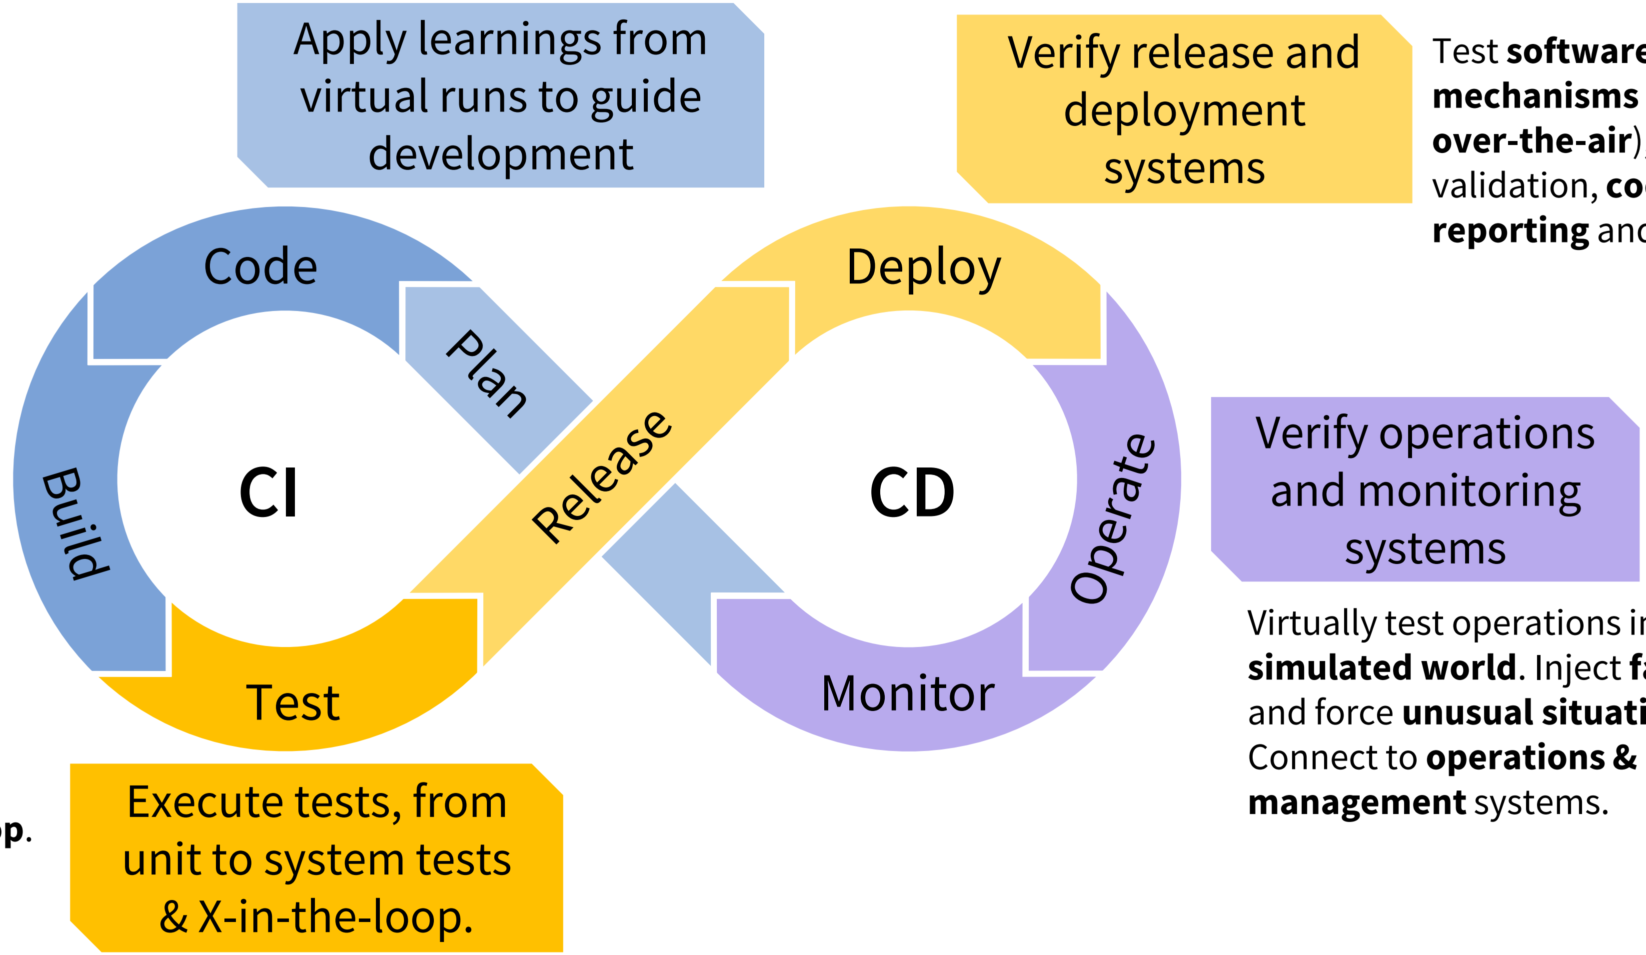

Test **software update mechanisms** (including **over-the-air**), code validation, **code loading**, **reporting** and **attestation**

Virtually test operations in a **simulated world**. Inject **faults** and force **unusual situations**. Connect to **operations & management** systems.

Scale testing **without hardware** bottlenecks. Provide **agility** and quick feedback loops. Virtualized **x-in-the-loop**.

"Shift Left"

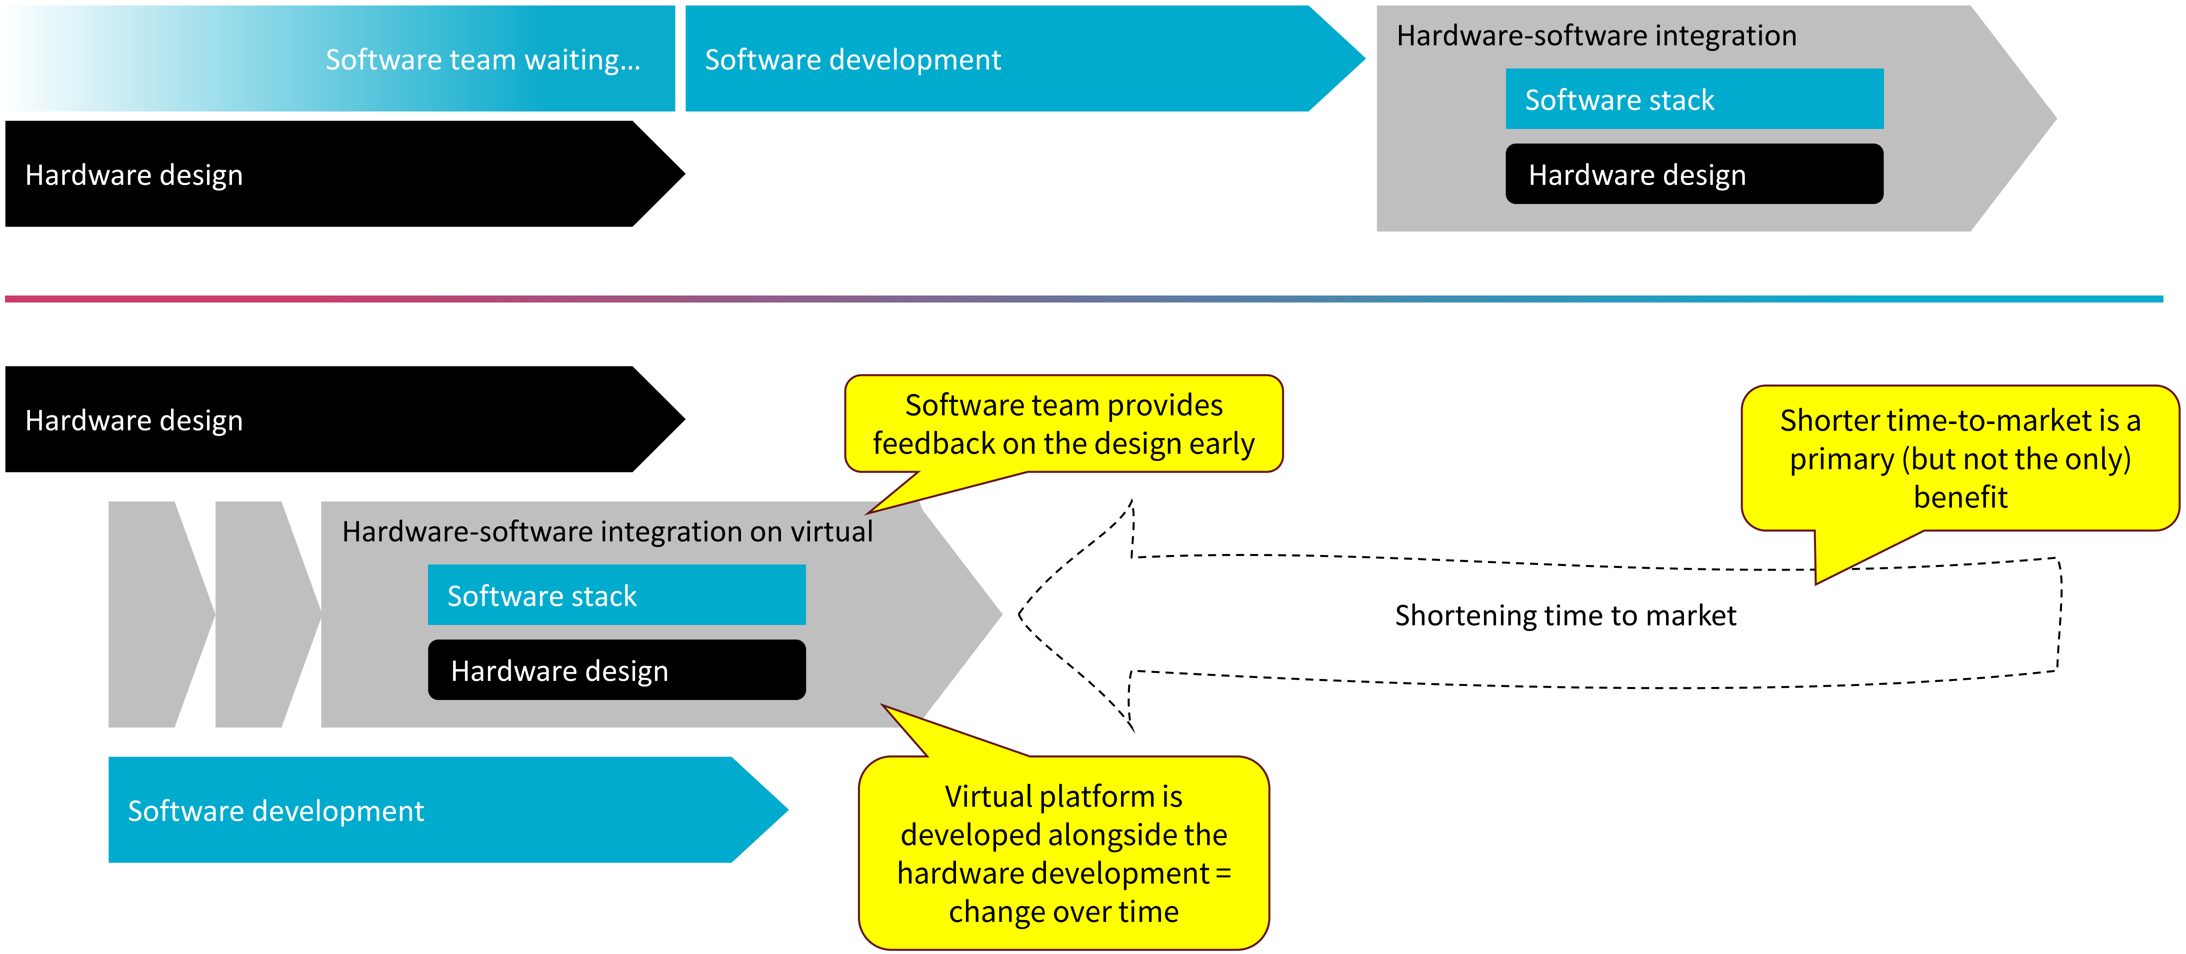

VLAB WORKS™

How does it Work?

Virtualization for Embedded Systems

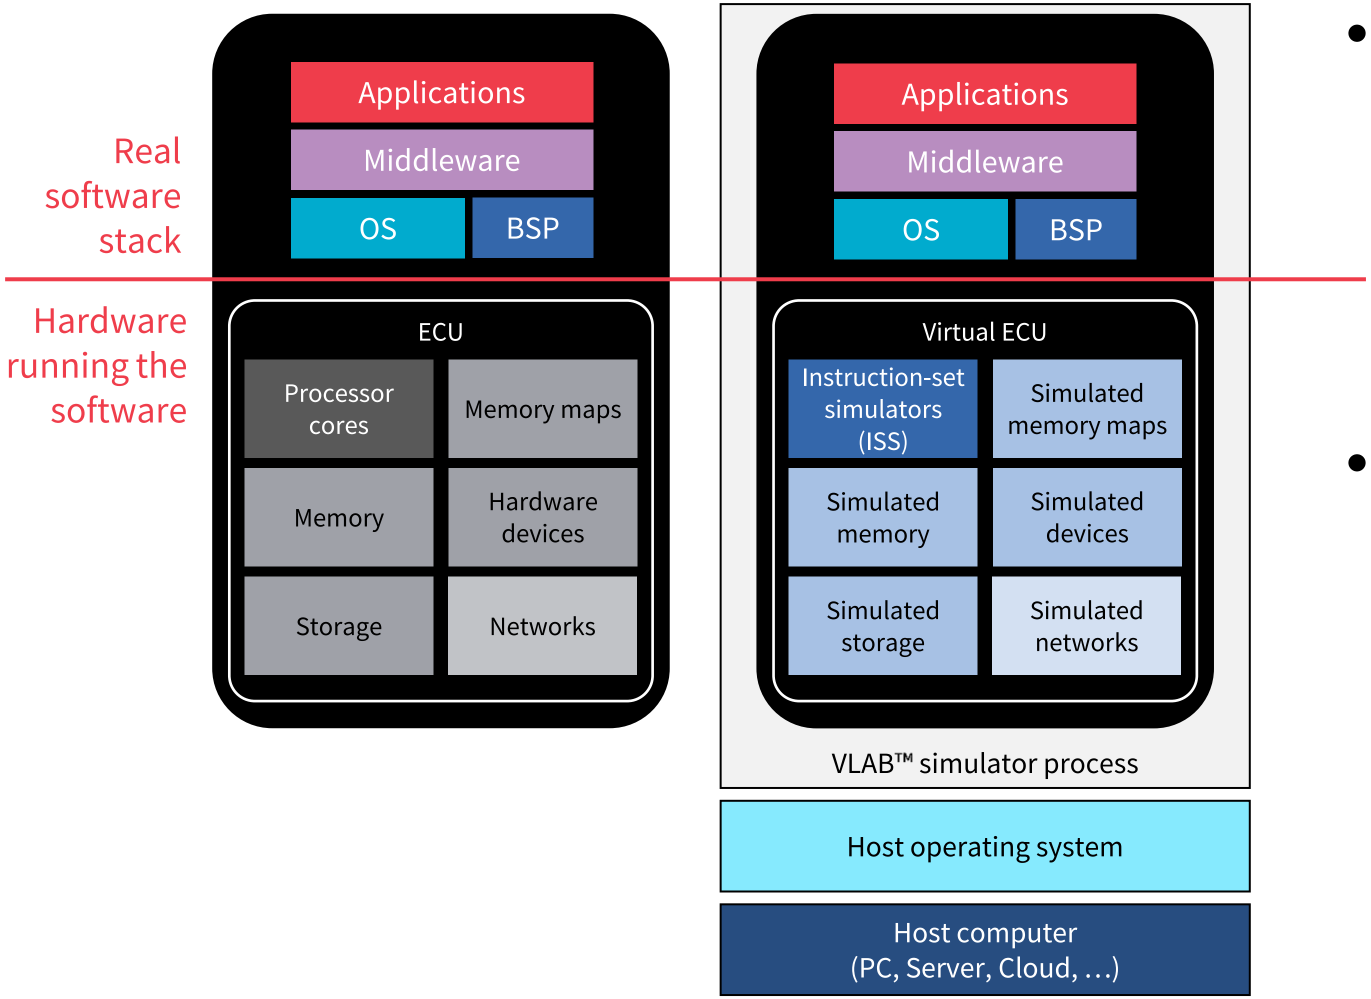

- VDM (Virtual Development Machine):
  - **Software simulation** of the hardware
  - Runs the **real software** – without the hardware
  - **Fast** – interactively useful
- Use cases:
  - Effective software debug
  - Software analysis and test
  - Scaling CI/CD
  - Fault injection
  - Regression testing
  - Early hardware access
  - …

Inside The Virtual Platform

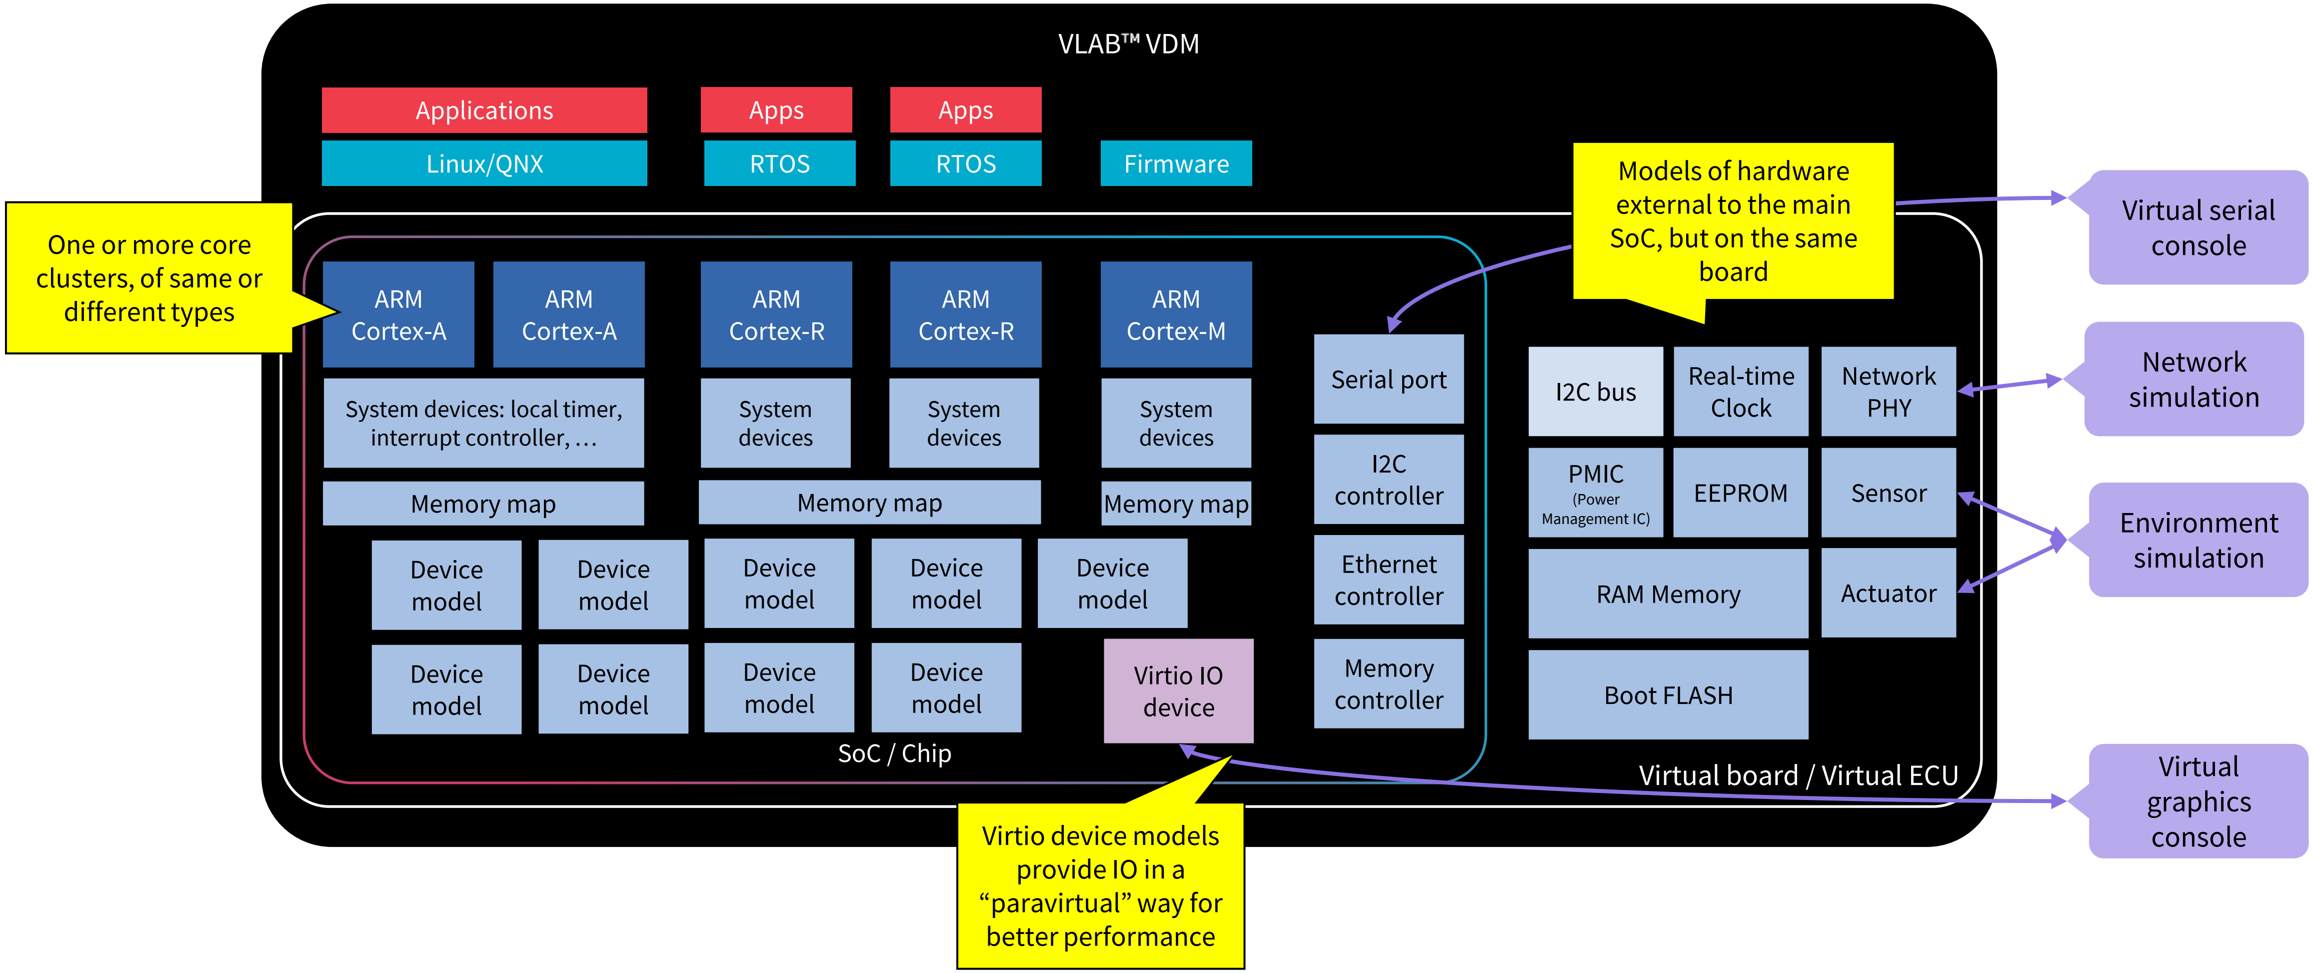

Variants of Virtual

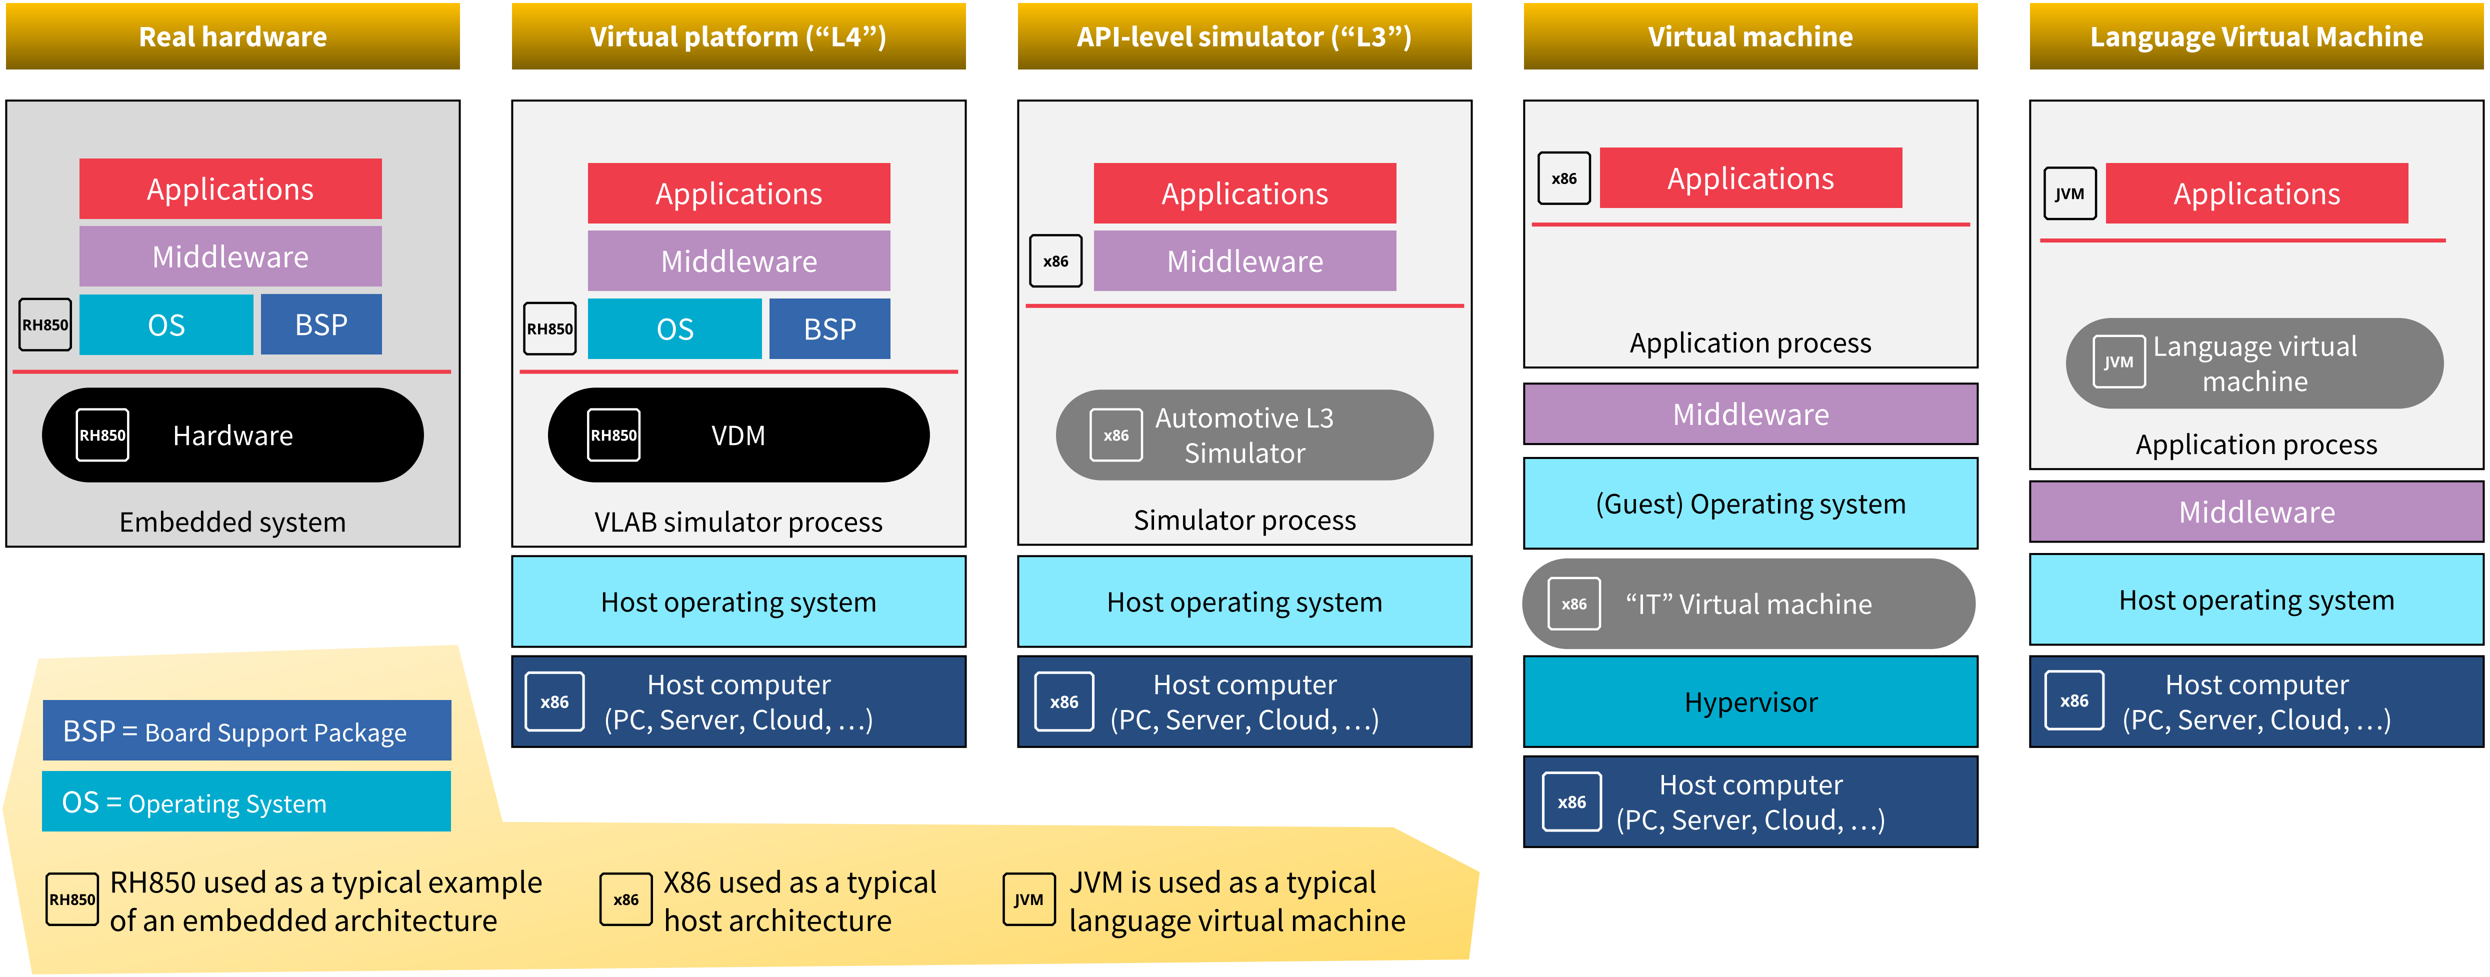

Computer Architecture Companies in Sweden

“Computer Architecture” Companies in Sweden

Company	Website	Processor cores	Hardware blocks	Chips	Software	Simulators and tools
Ericsson (Kista)	https://www.ericsson.com	✷	✷	✷	✷	✷
Intel (Kista)	https://www.intel.com				✷	✷
Arm (Lund)	https://www.arm.com/	✷	✷		✷	
Shortlink (Karlstad)	https://shortlink.se/		✷	✷	✷	
ZeroPoint (Göteborg)	https://www.zeropoint-tech.com/		✷		✷	
Gaisler (Göteborg)	https://www.gaisler.com/	✷	✷	✷	✷	✷
Fingerprint Cards (Göteborg)	https://www.fingerprints.com		✷	✷	✷	
Nordic Semiconductor (Stockholm)	https://www.nordicsemi.com/		?	?	?	?
MediaTek (Linköping)	https://www.mediatek.com	✷	?		?	?
Axis Communications (Lund)	https://www.axis.com	?	✷	✷	✷	✷
SAAB (Linköping, Stockholm, …)	https://www.saab.com/sv/		✷	✷	✷	✷
Qamcom (Göteborg, Stockholm)	https://www.qamcom.com/		✷	✷	✷	
Adtran (Stockholm)	https://www.adtran.com/	✷	✷	✷	✷	

Not exhaustive by any means

VLAB WORKS™

astc

Thank You!

www.vlabworks.com

Virtual Platforms for Embedded Systems

www.astc-design.com

Advancing the Design of Electronic Systems