



Completed

Processor Trends

HEPiX Techwatch Working Group
July 10, 2024

Trends Affecting CPUs

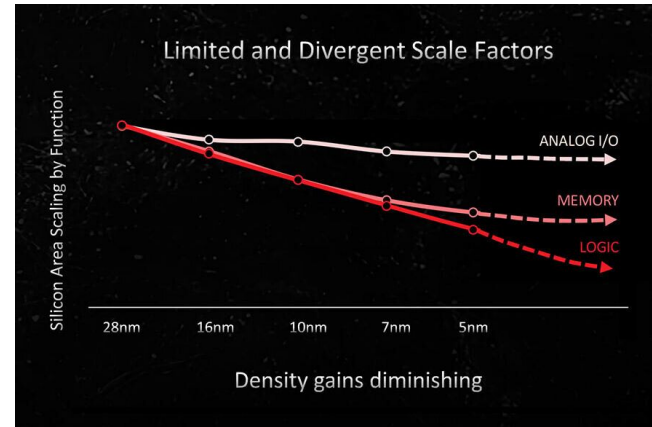
- End of Dennard scaling
- Scaling imbalance
 - Static RAM cell size vs logic
 - I/O vs logic
- Reduced lithography reticle limits
- Advances in packaging
- Explosive growth in the use of AI/ML technologies
- Changing dynamics at semiconductor foundries
- Increased competition in the CPU market
- Changing relationship between CPU producers and consumers

End of Dennard Scaling

- Dennard scaling
 - Transistor feature size is scaled by $1/K$
 - Transistor area decreases by $1/K^2$
 - Delay decreases by $1/K$ - Max frequency increase by K
 - Transistor power consumption decreases by $1/K^2$
 - Power consumption per unit area remains the same
- Power density no longer constant as logic density increases
 - For a given die size, power consumption increases (roughly $\propto K^2$)
 - CPU frequency has effectively stalled at $\sim 4\text{GHz}$

SRAM Scaling

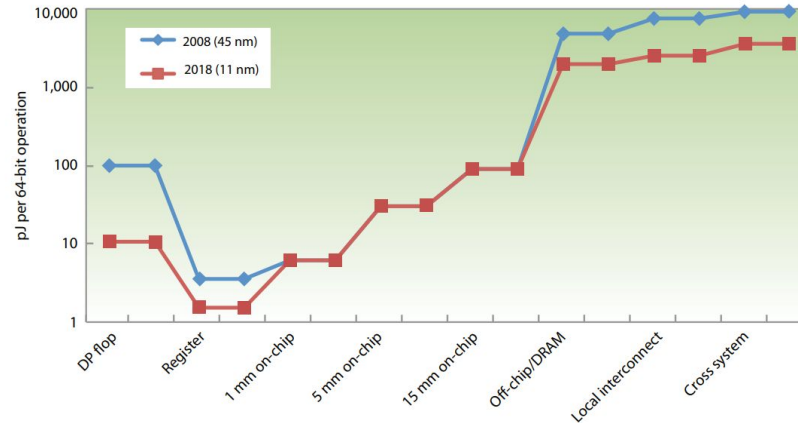
- For a fixed SRAM capacity and logic gate count, SRAM die area remains roughly constant as logic shrinks in newer processes
 - Trade off between core count and on die SRAM cache capacity per core when moving to more advanced nodes.
 - SRAM costs potentially higher with more advanced processes



TechPowerup.com, "AMD Explains the Economics Behind Chiplets for GPUs", Nov 14, 2022, <https://www.techpowerup.com/301071/amd-explains-the-economics-behind-chiplets-for-gpus>

I/O Scaling

- Scaling of communication circuitry has not matched logic scaling
 - I/O energy consumption per bit roughly constant
 - I/O energy per bit increases with distance
 - Size of I/O circuitry remains roughly constant
 - I/O circuitry does not benefit from process feature size shrink



Choices in Power/Performance/Area (PPA) Optimization

- Net Affect
 - Switch to multi-core processors to compensate for slower increases in single core performance
 - “Dark Silicon” workarounds
 - Dynamic power management (e.g. Turbo mode)
 - On die application specific hardware with different power/performance profiles (e.g., “big” core, “little” core, GPU, crypto HW, AI HW)
 - HPC vs HTC CPUs - Smaller # faster cores vs larger # slower cores

Drive to Chiplets


- I/O Power/Area consumption
 - Shorter signal distances, higher bandwidth, lower power consumption
 - HBM die stacking
 - LPDDR5/GDDR7 MCM
- Process optimized fabrication
 - Older process for I/O and SRAM
 - Newer processes for logic
- Reduced Reticle Area - Limits size of monolithic die.
 - i193 (DUV) and EUV limit ~ 853mm²
 - High-NA EUV limit ~ 450mm²
- Yield increases with smaller die sizes
 - Reduces on die device variation, increasing yield at higher performance
 - Loss due to defects is reduced
- Enables modular construction CPU designs with a wider array of customizations

Advanced Die Packaging

- Low power, low latency, high bandwidth die to die interconnect a key enabler of chiplet based CPU implementations
 - HBM memory utilizes 3D die stacking internally and 2.5D interconnect for connectivity to CPU or GPU die.
 - 2D, 2.5D, and 3D interconnect increases signal density, reduces power and increase bandwidth.
- Universal Chiplet Interconnect Express (UCIe)
 - Proposed chiplet interconnect standard to enable an interoperable ecosystem of chiplets from multiple vendors

I/O Scaling

- Scaling of communication performance and power consumption has not kept up with advances in logic
 - Use of LPDDR5x and HBM3e/4 as a replacement or a cache in front of DDR5 are mitigation strategies to increase memory performance and power consumption.
 - PCI-e Gen6 moves to PAM4, doubling bandwidth (Zero)
- HBM3e utilizes 3D die stacking in order to increase bandwidth and reduce power connectivity to CPU or GPU die.
- 2D, 2.5D, and 3D die to die interconnects to reduce power and increase bandwidth.



Technology Changes Affecting CPUs

- Failure of Dennard Scaling
 - Power density no longer constant as device (transistor) sizes shrinks
 - Effective limit on CPU frequency
- Static RAM cell size not scaling with logic
- EUV and High-NA EUV reticle limits
 - EUV and High-NA EUV limits CPU die sizes to $\sim 850\text{mm}^2$ and $\sim 450\text{mm}^2$ respectively.
- Bit transfer “performance” not scaling with logic
 - I/O rates lagging logic performance
 - On chip and off chip I/O power consumption not scaling with logic
- Advances semiconductor packaging
 - 2D/2.5D/3D interconnect

Changing Landscape in CPU Market

- Semiconductor foundries
 - Intel late with multiple generations of new semiconductor processes over the past decade
 - Intel late to transition to EUV lithography
 - TSMC capitalized on problems at Intel and successfully moved EUV lithography into production.
- Increased competition in the CPU market
 - Resurgence of AMD with the introduction of Zen
 - Substantial penetration of ARM ISA CPUs (Apple M series) in the CPU market
 - ARM Neoverse and Neoverse CSS
- Explosive growth in AI/ML applications and AI/ML accelerators

DRAFT

Explosive Growth in AI/ML Applications

- AI/ML accelerators imbedded in CPU's
- Driving need for higher memory bandwidth
- Driving need for tighter coupling between CPU's and external AI/ML/GPU accelerators
 - CXL 1.1
 - NVLink



DRAFT

Intel Product Portfolio

- Intel Xeon Max
- Intel 6th Generation Xeon
 - Granite Rapids
 - P Cores SMT?
 - Sierra Forest
 - E cores no SMT

AMD Product Portfolio

- Genoa
 - Zen 4 CCD - Chiplet with performance optimized, SMT enabled, Zen 4 cores + L1/L2/L3 cache
 - Separate I/O die
- Genoa X
 - Zen 4 CCD stacked with SRAM chiplet
- Bergamo
 - Zen4c CCD - Area optimized, SMT enabled, Zen 4 cores (Zen 4c) with L1, L2, and reduce capacity L3 cache

Chiplet design optimizes overall yield in addition to frequency binned yield. CCD die size well with EUV/High-NA EUV reticle limit. Separate I/O die allows for use of older processes better suited to I/O

Chiplet design allows for “mix and match” to create a broader product line. Heavily dependent on availability, performance, and maturity of die interconnect technologies.



Completed

ARM and the Changing Producer/Consumer Relationship

- ARM is a supplier of CPU IP that entered the server market in 2018
 - ARM does not sell CPU chips
 - ARM sells designs for the major components of a complete CPU; cores, MMU, interconnect fabric, memory controller, etc.
 - Designs are either “soft” or “hard” IP i.e., logical designs (RTL models) or physical implementations from foundry partners (e.g. TSMC)
- Three generations of ARM Neoverse cores, E_n , N_n , and V_n , where n =generation (1,2,3). Core types target different environments
 - E_n - Low power (energy efficiency)
 - N_n - “Balanced” power and performance
 - V_n - High performance
- ARM IP significantly reduces the expertise and effort required to develop a complete CPU
 - Costs well within the budget of the large public cloud providers



Completed

ARM Based Data Center CPUs


- Neoverse V2 derived CPUs
 - Amazon Graviton4
 - Nvidia Grace
 - Google Axion
- Neoverse CSS N2 derived CPUs
 - Microsoft Cobalt 100
- Neoverse N1 derived CPUs
 - Ampere Altra / Altra Max
- Custom (non Neoverse) derived CPUs
 - Ampere AmpereOne



Completed

ARM IP Market Disruption

| IP Development | Compute Subsystem | Top-Level SoC (Arm owns) | BackEnd (Arm owns) (Reference) | Software (Partner owns) (Reference) |
|------------------------------------|--|---|--------------------------------------|--|
| Arch, CPU, CMN, System, POP/RFM | Arch, IP Config, Perf, RTL, Verify/SBSA, FPGA Image | SoC Arch, 3PIP Config, 3PIP Perf, 3PIPVerify | Impl pkg, TO | FVP, FW, OS |



1. 3rd Party IP (3PIP) includes Memory, PCIe/CXL, Die-to-Die PHY components
2. Engineering Year (EY). 80 EY savings is estimated from a typical CSS customer.

Neoverse CSS

- Neoverse CSS IP
 - Preconfigured, mostly complete SoC with tunables (e.g., #cores, cache size)
 - Chiplet support via UCIe or customer proprietary
 - External accelerator support via PCIe-5/CXL1.1
 - Significantly reduces effort and expertise required to design from components IP
- Open question for non-captive ARM developers:
 - What is the value proposition?
 - Is there enough demand to support a custom core or Neoverse derived ARM CPU in the open market?



DRAFT

Conclusions

