

Source sites	2
Topics	3
Information to collect	3
Introduction	4
General Purpose Microprocessors	5
x86 processors	5
Intel	5
Xeon Scalable	5
Nervana	6
Foveros and hybrid CPUs	6
AMD	6
Power processors	7
ARM	7
ThunderX2	7
eMAG	8
Graviton	8
Other	8
Graphics Processors (GPU/GPGPU)	8
Nvidia GPUs	8
AMD GPUs	8
GPGPU Software ecosystems	8
Machine Learning Processors	8
Google TPU	9
Intel Nervana	9
FPGA	9
Other Accelerator/CoProcessors	9
Embedded Microprocessors	9
Supporting Technology	9
Memory Technology	9
Interconnect Technology	10
PCI-e	10
OpenCAPI	10
NVLink	10

CCIX	10
GenZ	10
Packaging Technology	10
2D Packaging	10
2.5D Packaging	11
3D Packaging	11

Source sites

From Bernd's presentations:

<https://www.extremetech.com/category/computing>

https://www.theregister.co.uk/data_centre/servers/

<https://www.tomshardware.com/>

<https://www.nextplatform.com/>

Andrea S.:

<https://www.anandtech.com/>

<https://www.techspot.com/>

<https://www.techarp.com/guides/workstation-server-cpu-comparison/>

<https://www.servethehome.com/>

<https://hothardware.com/>

<https://www.phoronix.com/>

<https://www.networkworld.com/category/data-center/>

<https://www.computerworld.com/category/data-center/>

<https://datacenterfrontier.com/>

Shigeki M:

<https://www.realworldtech.com>

<https://en.wikichip.org/wiki/WikiChip>

<https://www.hotchips.org/>

Olof B:

<https://www.smartbrief.com/industry/tech> in particular the JEDEC channel

<https://www.datacenterdynamics.com/news/>

Tristan S.:

<https://www.semiaccurate.com/>

Harvey N.:

<https://www.cpubenchmark.net/>

Servesh.M:

<https://ark.intel.com/>

<https://www.agner.org/optimize/> (Specifically the CPU Blog)

<https://semiengineering.com/>
<https://www.cpubenchmark.net>

Topics

Topic	People
x86	Mattieu Puel (for AMD) , Luca Atzori, Tristan Suerink (for AMD and Intel) , Andrea Chierici, Michele Michelotto, Andrea Sciabà
ARM	Tristan, Niko , Pepe
POWER	Tristan , Niko
RISC-V	Fons Rademakers, Tristan
GPUs	Servesh Muralidharan, Felice Pantaleo
FPGA	Niko Neufeld, Servesh

Information to collect

- CPUs
 - Architectures
 - List, shortly describe and compare the main processor architectures and implementations from different vendors
 - X86, ARM, POWER, RISC-V
 - Roadmaps for development
 - Relevance for HEP
 - Manufacturing and sales
 - Process technologies and evolution
 - Market shares and units shipped
 - Pricing
- GPUs
 - Architectures
 - List and compare the main GPU architectures and implementations from different vendors
 - Roadmaps for development
 - Relevance for HEP
 - Manufacturing and sales
 - Process technologies and evolution

- Market shares and units shipped
 - Pricing
 - Effect of cryptomining?
 - Licensing for data centers?
- FPGA
 - Architectures
 - List and compare the main FPGA architectures and implementations from different vendors
 - Roadmaps for development
 - Relevance for HEP
 - Manufacturing and sales
 - Process technologies and evolution
 - Market shares and units shipped
 - Pricing

HEPiX Techwatch WG : CPU and Accelerators

Introduction

“It was the best of times, it was the worst of times, it was the age of wisdom, it was the age of foolishness, ... “, *A Tale of Two Cities*, Charles Dickens.

“[It's the ecosystem, stupid!](#)”, from [EE Times](#).

The central processing unit (CPU) or microprocessor, are the “brains” of the computing and storage systems that are critical in the collection and analysis of data obtained from HEP experiments. The past decade and a half brought a period of relative stability, or stagnation depending on your perspective, in microprocessors. During this period, the x86 architecture, for all intents and purposes, had relegated other CPU architectures to “niche” markets. This and the lackluster performance of AMD’s “Bulldozer” in 2011, effectively left Intel as the only game in town. After a series of “tick-tock” improvements to their x86 architecture processor, a confluence of events, both inside and outside of Intel, has triggered a minor renaissance in CPU’s. Whether viewed as a positive development or an act of desperation, the new developments in the computing field promise exciting times ahead.

Recent events influencing the direction of CPU technology includes:


1. Problems with Intel’s 10nm process, causing disruption in the Intel product pipeline.
2. Resurgence of AMD with their new Zen microarchitecture.
3. Renewed attempts by IBM to create a broader ecosystem for their Power instruction set architecture (ISA) with their OpenPower consortium.

4. Invasion of the server market by ARM and ARM licensees
5. Consolidation of the foundry business
6. Transition to <10nm processes and EUV lithography
7. Rise of GPGPU computing, accelerated by Nvidia's CUDA framework.
8. Interest in Machine Learning, leading to interest in application specific processors

Overview

Since th

Scaling trends for general purpose microprocessors, from [Karl Rupp's "42 Years of Microprocessor Trend Data"](#) web page.



This document focuses on CPU's, application specific accelerators, CPU interconnect technology, and packaging technology.

General Purpose Microprocessors

The general purpose microprocessor arena is inhabited by three main instruction set architectures, x86, Power, and ARM. These are the remaining survivors from the plethora of CPU architectures that existed in the 80's and 90's. x86 is the proverbial 800 lb gorilla that decimated ranks of the RISC/UNIX CPUs. The x86 architecture benefited from economies of scale (technical, manufacturing, and financial) and a large software and hardware ecosystem. At the sole surviving RISC/UNIX CPU, the Power ISA benefited from the technical and financial strength of IBM. Finally, ARM has achieved market dominance in the IoT and cell phone market, and is attempting to become a player in the server market.

x86 ISA processors

Intel (Andrea, Tristan)

Xeon Scalable

The current server CPU lineup is the Xeon Scalable introduced in 2017, which comprises several families (Bronze, Silver, Gold and Platinum), with a core number up to 28 for 56 threads. The Skylake architecture was initially introduced in 2015 for the desktop CPU Core lineup, which went through some minor refinements (Kaby Lake, Coffee Lake, Ice Lake). The upcoming successor is the Cascade Lake architecture, still built on the 14nm process and foreseen for 2019. The main new features setting it apart from its predecessor are:

- Core count up to 28
- Hardware mitigations for Spectre, Meltdown and L1TF
- Support for Optane DIMMs
- AVX-512 Vector Neural Network Instructions
- Better power efficiency (14nm++ process)

As for Skylake, Cascade Lake will use Ultra Path Interconnect (UPI), which allows multiple processors to share the same address space. The first Cascade Lake CPUs belong to the Xeon Gold and Platinum families and were launched in December 2018.

Cascade Lake-AP will be a special implementation consisting of two Cascade Lake CPUs in the same package for a total of 48 cores and it will be relevant only for massive multicore workloads.

Another architecture upgrade, Cooper Lake, is foreseen later in 2019, featuring relatively minor improvements, namely support for bfloat16 and eight memory channels; it will be adopted by all Xeon Scalable families and it is considered as the answer to the AMD EPYC Rome CPU, in particular if it will also include support for PCIe Gen4.

A much more significant improvement is expected with the introduction of the Sunny Cove microarchitecture, due for mid 2019, manufactured using the 10nm+ process and to be adopted for the Ice Lake CPUs. In this case, the core architecture will undergo several changes, among which:

- Two additional ports (for a total of 10), one for storing data and one for memory access
- 5-wide allocation
- 50% bigger L1 caches, larger micro-op cache, second-level TLB and L2 caches
- Two SIMD shuffles and four LEA units

- Several additional AVX-512 extensions

A double-digit gain in IPC is to be expected also for general purpose, unoptimized code, while performance for specialized operations (compression, cryptography, security, machine learning) will particularly benefit from many of the planned enhancements.

Additional microarchitecture improvements are planned for Willow Cove, which should undergo a cache redesign, a transistor optimization and more security features, probably in 2020, and Golden Cove, with better single threaded and AI performance, supposedly from 2021. More information can be found in reports from Intel 2018 Architecture Day.

Nervana

As part of its strategy to focus on AI applications, Intel introduced a new line of CPUs called Neural Network Processors (NNP). The first-generation CPU, Lake Crest, was delivered to a small number of partners in 2018. The second-generation CPU, the L-1000 (code name Spring Crest), should be 3-4 times faster and will become broadly available in 2019. Among the main improvements there is the support for bfloat16, an increasingly popular numerical format for neural networks. The Nervana CPUs should directly compete with Nvidia GPUs for accelerating ML applications.

Foveros and hybrid CPUs

Intel has recently presented Foveros, a new 3D chip stacking technology that is supposed to allow for very flexible CPU designs, for example with heterogeneous compute nodes. Apart from a demo, however, no product has been announced yet and it is not clear if and when Foveros will be adopted for server class CPUs.

AMD

Following a period of famine, resulting from the underwhelming performance of the Bulldozer follow on to the successful K10 processor, AMD is on the verge of a remarkable resurgence, based on their new Zen microarchitecture.

The Zen microarchitecture, the basis for the AMD Ryzen desktop and Epyc server processors, is a “clean slate” microarchitecture that aimed to dramatically boost instructions per cycle (IPC) per core over Excavator (the final iteration of the Bulldozer microarchitecture) to be more competitive with Intel processors. A significant differentiating characteristic of the Zen and follow on processors is the use of multi-chip modules (MCM) and “chiplets” to implement a scalable processor family.

The use of multi-chip modules and chiplets is fairly radical move for AMD, but is understandable in light of its financial and technical resources.

Zen

First generation implementation of Zen micro architecture, consisting of a “Zeppelin” SoC building block. Each Zeppelin SoC die consists of :

- 14nm engraving fineness
- Two CPU Complexes (CCX) per die
- Four Cores per CCX
- Two DDR4 memory channels per die, up to 2600 MHz
- 32 Infinity Fabric/PCI-e Gen 3 lanes per die
- Four Infinity Fabric “on package” lanes per die
- Memory and I/O controllers
- 32 PCI gen3 lanes per processor (16 on dual socket configurations)

Four Zeppelin dies are aggregated with a multi-chip module (MCM) design to build an EPYC (codename “Naples”) processor :

- SP3 socket
- TDP range from 120W to 180W
- Available in 8 (bi-socket configuration), 16, 24 and 32 cores

EPYC Naples processors are available since June 2017. They have similar computing power compared to Intel Skylake processors (HS06 benchmarks on close frequencies and CPU core counts) with cutoff prices up to 49% (AMD claim). They are mostly compatible with Intel x86, which spares for user code modifications.

Zen+

Zen+ is a little optimization of Zen architecture, bringing 12nm, higher clock speeds and lower power consumption.

Zen 2

Zen2 based EPYC processors are codenamed “Rome”. Those processors embed nine dies : one shared IO die (14nm, GlobalFoundries) supporting I/O and memory accesses and eight 7nm chiplets (CPU dies processed by TSMC). AMD’s strategy regarding IO die seems to look for cost reduction as IO does not scale as much as CPU. The possible frequency increase (thanks to 7nm engraving fineness) is expected to reach 300~400Mhz for low core count processors and equivalent frequency for 64 cores processors compared to Naples. GlobalFoundries giving up 7nm graving.

Main specs:

- 9 dies per chip : a 7nm single IO/memory die and 8 CPU 7nm chiplets
- avoids dedicated chipset
- range of frequencies are not disclosed as of this writing
- 8 DDR4 memory channels, up to 3200 MHz
- up to 64 cores (128 threads) per processor
- up to 128 PCI gen3/4 lanes per processor
- SP3 / LGA-4094 sockets

- TDP range: 120W-225W (max 180W for SP3 compatibility)

Rome design ended up on Q1 2018. The processors are expected in production as of Q2 2019. The ZEN2 core architecture is followed by a ZEN3 architecture, processors codenamed "Milan".

Power ISA processors

1. Power 9
2. Power 10

Power CPUs will be the first on the market with PCIe Gen5 support.

Power 9 is equipped with Coherent Accelerator Processor Interface (CAPI) 2.0 I/O interfaces to enable the following:

1. Coherent user-level access to accelerators and I/O devices
2. Access to advanced memories via read/write or user-level DMA semantics

CAPI is designed to reduce latency and increase bandwidth to accelerators and I/O devices. Power 9 is also equipped with NVLink interfaces to increase bandwidth to NVidia GPUs.

ARM ISA processors

1. Marvell (Cavium) ThunderX2
2. Ampere eMAG
3. Amazon (Annapurna Labs) Graviton
4. Huawei Kunpeng 920
5. Qualcomm Centriq
6. Fujitsu A64FX
7. Mellanox Bluefield

ThunderX2

Information on the Marvell ThunderX2 processor can be found at the WikiChips [Vulcan](#) and [ThunderX2](#) web pages.

eMAG

Information on the Ampere eMAG ARM processor can be found at the WikiChips [Skylark](#) and [eMag](#) web pages.

Graviton

Other

Graphics Processors (GPU/GPGPU)

1. Nvidia
2. AMD

Nvidia GPUs

NVLink/OpenCAPI

AMD GPUs

GPGPU Software ecosystems

1. CUDA - Nvidia
2. OpenCL - Khronos Group

Machine Learning Processors

1. Google Tensor Processing Unit (TPU)
2. Intel Nervana NNP
3. Amazon Inferentia

Google TPU

Intel Nervana

FPGA

1. Altera
2. Xilinx

Other Accelerator/CoProcessors

Embedded Microprocessors

1. ARM
2. RISC-V
3. MIPS

ARM

RISC-V

MIPS

Supporting Technology

Memory Technology

1. DDR
2. GDDR
3. HBM
4. NVDIMM

Discussion on memory technology that impacts CPUs are discussion in the reports generated by the HEPiX [Techwatch Memory](#) WG (DDR/GDDR/HBM) and Storage WG (NVDIMM).

Interconnect Technology

1. PCI-e
2. Open Coherent Accelerator Processor Interface (OpenCAPI)
3. NVLink
4. Cache Coherent Interconnect for Accelerators (CCIX)
5. GenZ
6. Infinity Fabric
7. Ultra Path Interconnect (UPI)

PCI-e

1. PCI-e Gen 4
2. PCI-e Gen 5

OpenCAPI

NVLink

CCIX

GenZ

PCI-e is developed by the [PCI-Sig](#).

[OpenCAPI Consortium](#) manages the Open CAPI standard

NVLink is developed by Nvidia

[CCIX Consortium](#) is the organization developing the CCIX interconnect standard

[Gen-Z Consortium](#) is responsible for the Gen-Z interconnect standard

Packaging Technology

CPUs are built on 300 mm silicon wafers, with typically ~150 CPUs or “die” per wafer. Individual CPU die are separated from the wafer via wafer dicing. The bare die needs to be package to protect the die and to provide a way to connect the CPU to a circuit board. Although the majority of existing CPUs on the market consist of one die per package, a non trivial fraction of CPU’s are built with multiple die per package. In the latter case, the die in each package are typically not the same. For example, Intel’s Kaby Lake-G consists of an Intel CPU die, an AMD Radeon GPU die and High Bandwidth Memory (HBM) die in a single package.

There are many types of packaging, but they basically fall into three categories, as follows:

1. 2D Packaging (MCM)
2. 2.5D Packaging
3. 3D Packaging

These packaging types differ in how the die are arranged and interconnected.

2D Packaging

2D packaging refers to die packaging that consists of a multiple die mounted on a substrate. The 2D designation comes from the fact that the die all sit on a common plane. Feature sizes on the substrate are typically much larger than the die, resulting in limited connectivity and power and performance penalties.

2.5D Packaging

2.5D packaging refers to the use of an interposer between the die and the substrate. The interposer connect the multiple die together, as well as connecting the die to the substrate. The interposer allows for substantially higher interconnectivity between die at higher speed, and lower power than 2D packaging.

3D Packaging

3D packaging refers to the interconnection of die that are physical stacked, like floors in a skyscraper building. Communication between die is accomplished via “through silicon vias”. An example of 3D packaging are stacked DRAM die in High Bandwith Memory.