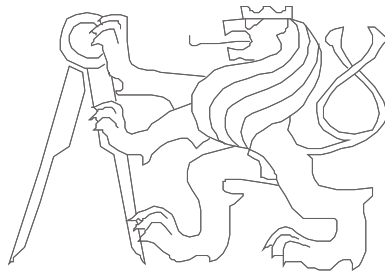


Advanced Computer Architectures


History and Future



Czech Technical University in Prague, Faculty of Electrical
Engineering

Slides authors: Michal Štepanovský, update Pavel Píša

ISA development history



1936 Alan Turing: On computable numbers, with an application to the Entscheidungsproblem

1937: Howard Aiken: Concept of **Automatic Sequence Controlled Calculator – ASCC.**

1945 **John von Neumann**: First Draft of a Report on the EDVAC. New idea: **Stored-program computer.** Previous computers required to physically modified to for given task. Remark: stored-program idea appeared even earlier in 1943 year – ENIAC development: J. P. Eckert a J. Mauchly

Computers of that era are equipped by accumulator (one register) for arithmetic and logic operations which is fixed destination and one of source operands.

1939 **Bombe**: designed to crack Enigma

1941 Konrad Zuse: **Z3** – the world first functional Turing complete computer, program controlled

1944 **Harvard Mark I**

1944 **Colossus**

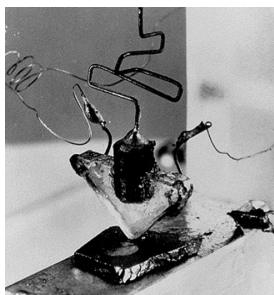

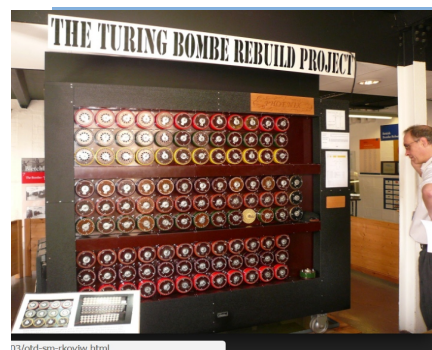
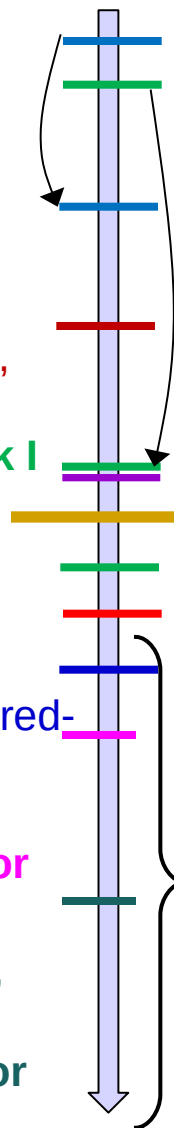
1946 **ENIAC**

1947 **Transistor**

1948 **Manchester Baby** – the first stored-program computer

1949 **EDSAC** – **single accumulator**

1953 **EDSAC**, **Manchester Mark I**, **IBM 700 series**: **single accumulator + index register**



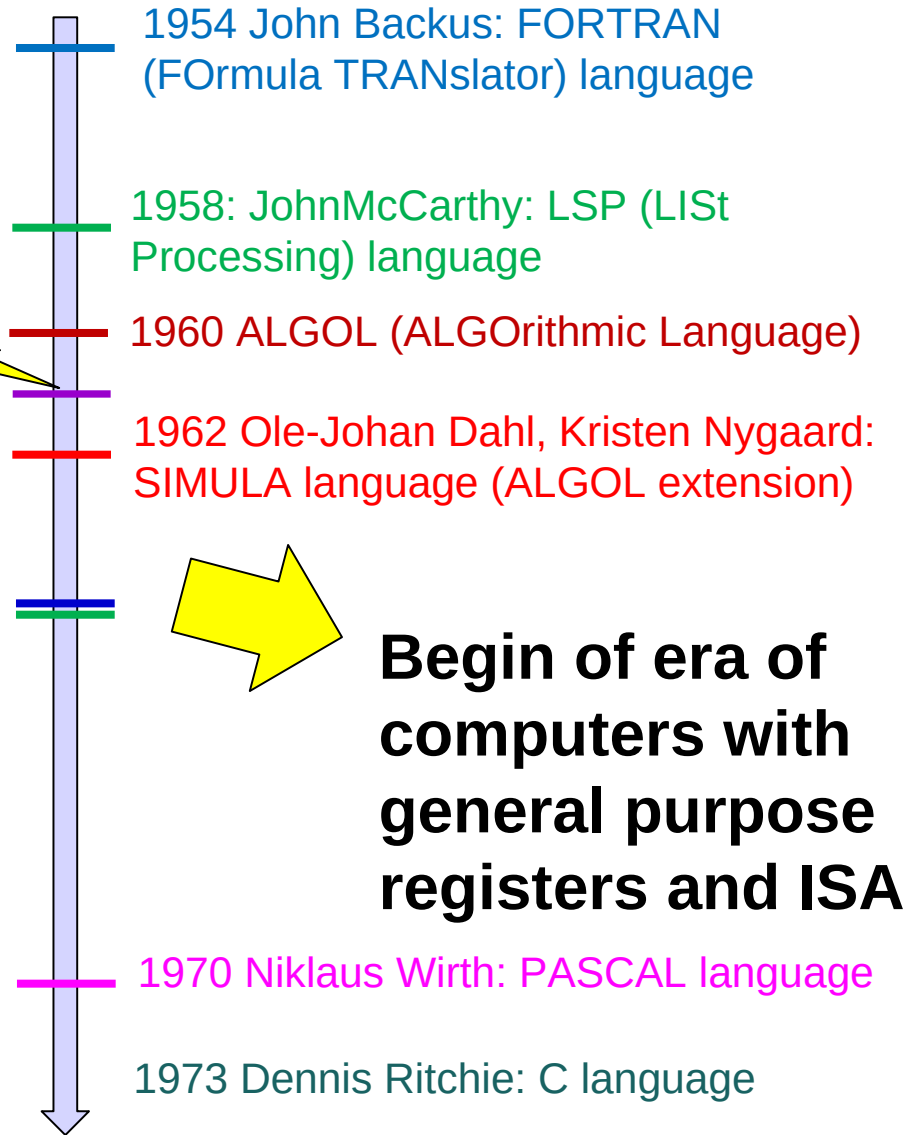
ISA development history

There is a significant separation of the programming model from implementation !!!

1961 B5000: Computer designed and optimized for ALGOL 60 processing
=> orientation to higher level languages


1964 IBM System/360 – under Gene Amdahl lead: strict separation of architecture from implementation – oriented to instructions/assembler

1964 CDC 6600 – the most powerful computer of its era



Begin of era of computers with general purpose registers and ISA

Control Data Corporation (CDC) 6600




1964 - Control Data Corporation (CDC) 6600

- Designed by Seymour Cray and Jim Thornton
- One CPU - 10 parallel functional units, each specialized (FP operations, Bool operations), 40MHz
- 60 registers, 60 bit each
- price \$8 millions (today equivalent of \$60 millions)
- peak performance of 3 MFLOPS - it is 10× more than IBM 7030 Stretch, which required 900m² area
- CDC required area of 4 cabinets
- Freon cooled
- 10 peripheral processors



1976 - the Cray 1 (Cray Research)



1976 - the Cray 1

- Designed by Seymour Cray
- Los Alamos National Laboratory
- The most successful supercomputer
- Price between \$5 and \$8 millions (today equivalent \$25 millions)
- It uses integrated circuits
- Word size 64 bits
- 136 MFLOPS
- Freon cooled
- Vector machine but multiple operations on block
 $a(1..1000000) = \text{addv } b(1..1000000),$
 $c(1..1000000)$

1985 - Cray 2



1985 - Cray 2

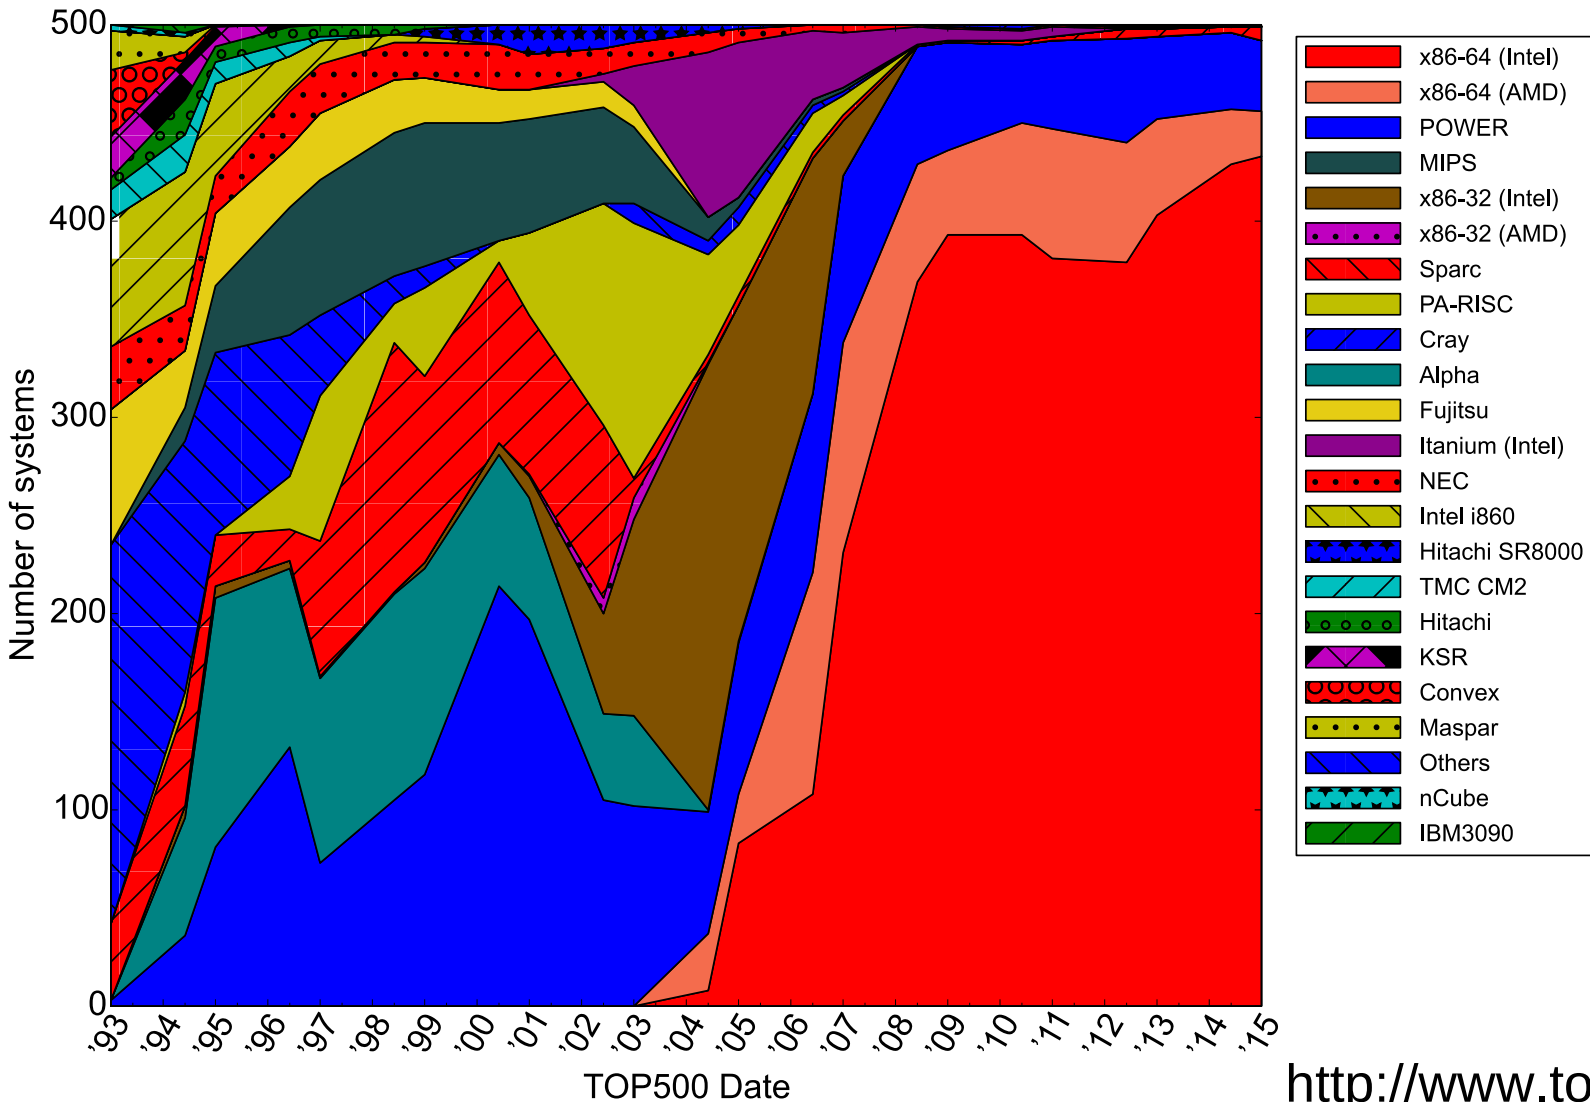
- Vector supercomputer
- 8 processors
- 1,9 GFLOPS (the fastest till 1990 year)
- UniCOS and Unix System V



From 1990 year

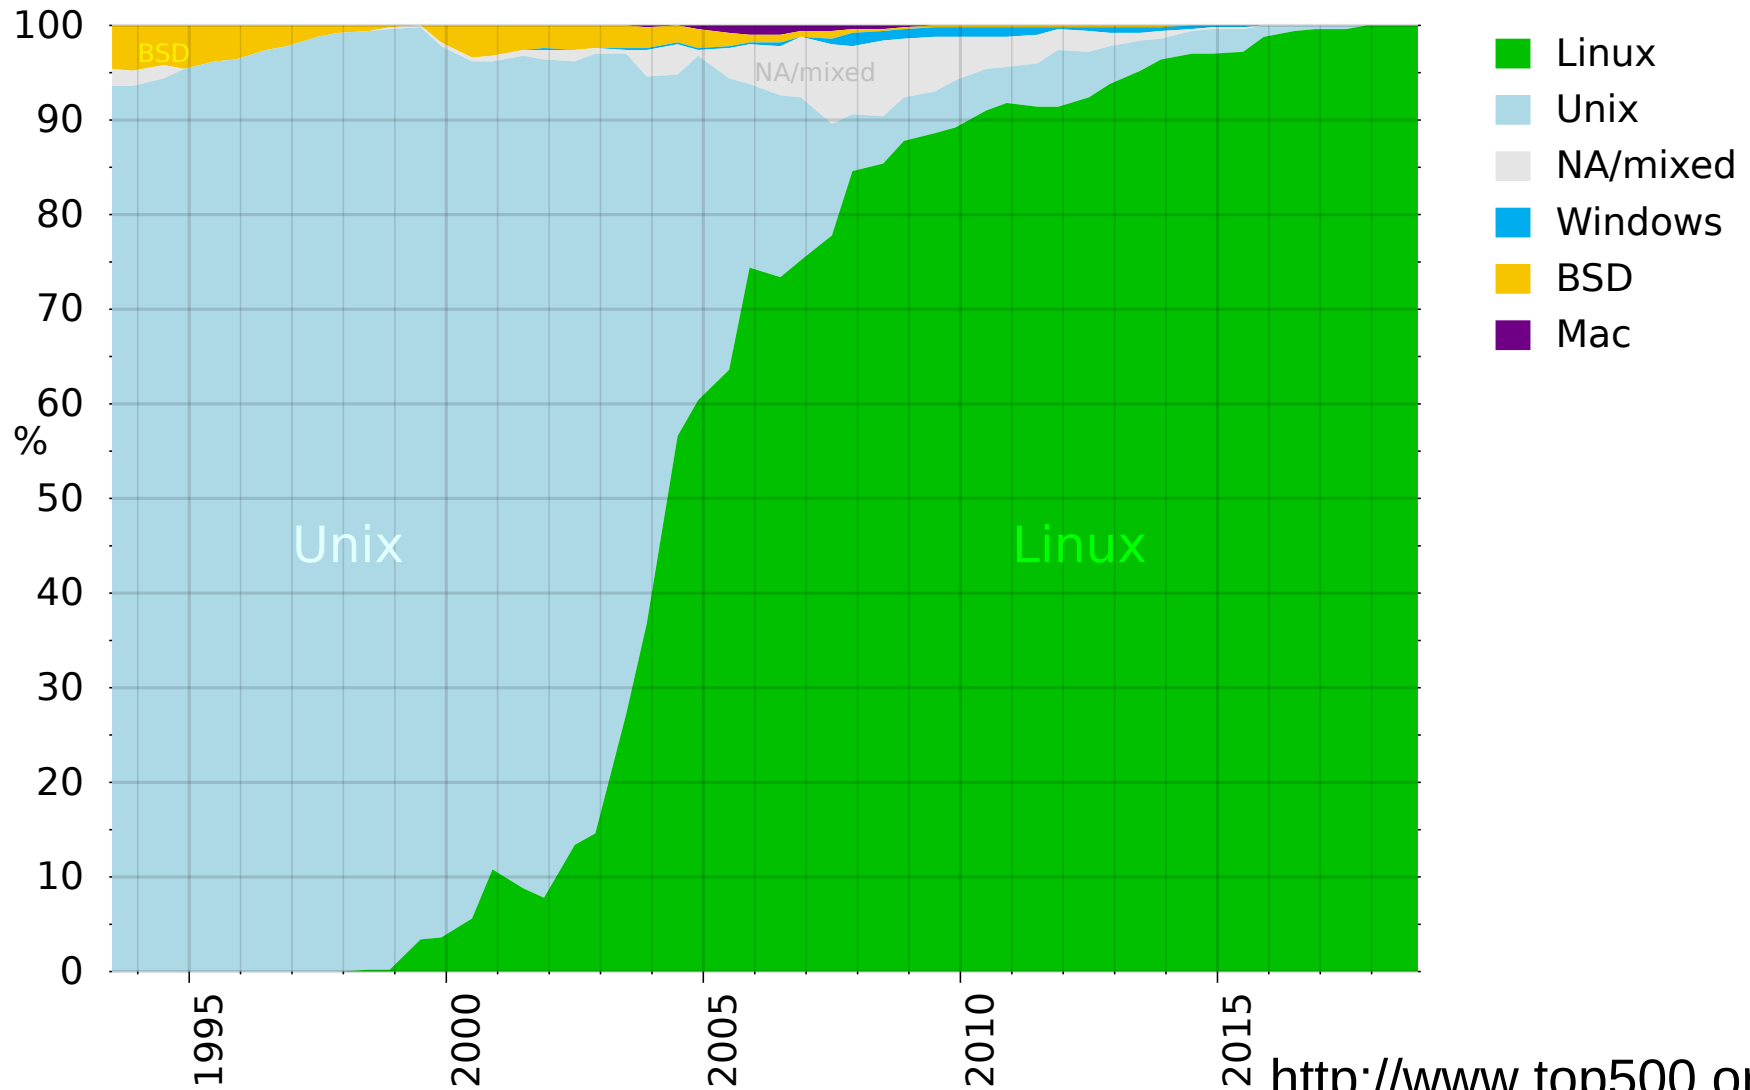
- 1993 Intel Paragon, max 4000× Intel i860 – 143 GFLOPS
- 1994 Fujitsu's Numerical Wind Tunnel, 166 vector processors, 1.7 GFLOPS/CPU – 170 GFLOPS
- 1996 Hitachi CP-PACS/2048 – 368 GFLOPS
- 1999 Intel ASCI Red/9632 – 2.3796 TFLOPS – Pentium II Xeon, 333 MHz.
- 2000 IBM ASCI White – 7.226 TFLOPS – IBM POWER
- 2002 NEC Earth Simulator – 35.86 TFLOPS – 5120× SX-6 (Cray license)
- 2004 – 2007 IBM Blue Gene/L – last version up-to 478 TFLOPS – QCDOC, 2× PowerPC 440
- 2008 IBM Roadrunner – 1.105 PFLOPS – 12,960 IBM PowerXCell 8i CPUs, 6,480 AMD Opteron dual-core processors, Infiniband
- 2009 Cray Jaguar – 1.759 PFLOPS – 224,256 AMD Opteron processors
- 2010 Tianhe-1A – 2.566 PFLOPS – 14,336 Xeon X5670 processors and 7,168 Nvidia Tesla
- 2011 Fujitsu K computer – 10.51 PFLOPS – 88,128 SPARC64 VIIIfx processors, Tofu interconnect (6D torus)
- 2012 IBM Sequoia – 16.32 PFLOPS – 98,304 compute nodes
- 2012 Cray Titan – 17.59 PFLOPS – 8,688 AMD Opteron 6274 16-core CPUs 18,688 Nvidia Tesla K20X
- 2013 NUDT Tianhe-2 – 33.86 PFLOPS – 32,000 Intel Xeon E5-2692 12C with 2.200 GHz 48,000 Xeon Phi 31S1P
- 2016 Sunway TaihuLight 93 PFLOPS – 40,960 SW26010 (Chinese) – total 10,649,600 cores

Development in top500 – used CPU architectures



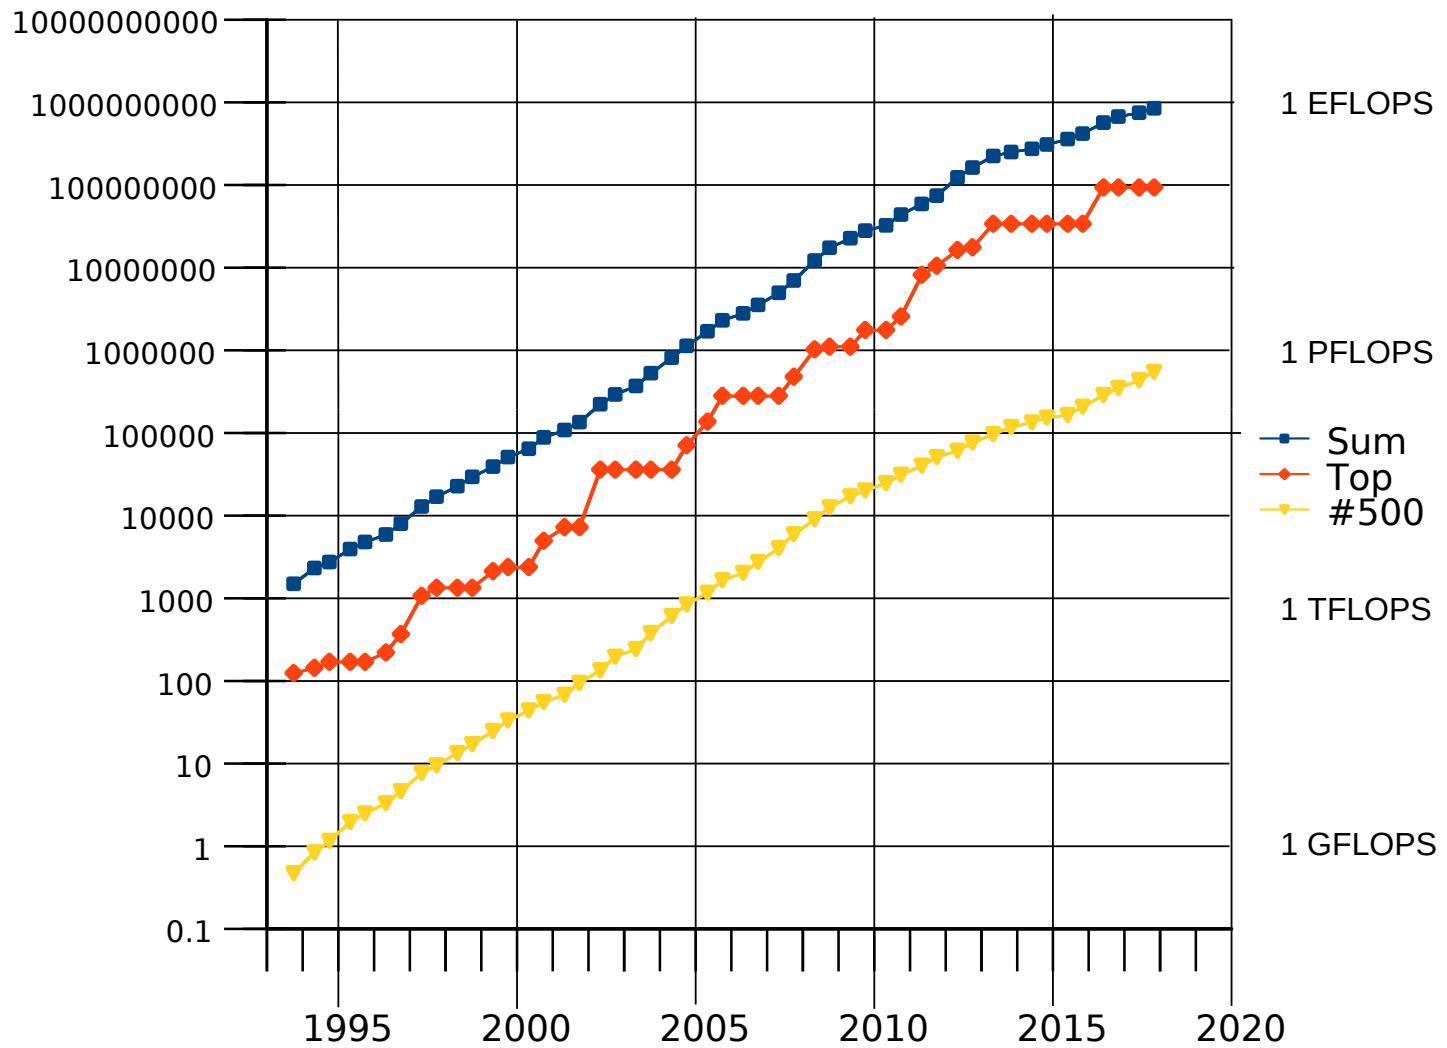
<http://www.top500.org/>

Development in top500 – used operating system



<http://www.top500.org/>


Processing power development GFLOPS



Source: https://en.wikipedia.org/wiki/History_of_supercomputing

Sunway TaihuLight, the Top500 #1 2016 – 2018

- 93 PFLOPS (LINPACK benchmark), peak 125 PFLOPS
- Interconnection 14 GB/s, Bisection 70 GB/s
- Memory 1.31 PB, Storage 20 PB
- 40,960 SW26010 (Chinese) – total 10,649,600 cores
- SW26010 256 processing cores + 4 management
- 64 KB of scratchpad memory for data (and 16 KB for instructions)
- Sunway RaiseOS 2.0.5
(Linux based)
- OpenACC
(for open accelerators)
programming standard
- Power Consumption
15 MW (LINPACK)



Summit supercomputer – IBM AC922

Plan 2018, US Oak Ridge National Laboratory (ORNL), 200 PetaFLOPS, 4600 “nodes”, 2× IBM Power9 CPU + 6× Nvidia Volta GV100

96 lanes of PCIe 4.0, 400Gb/s

NVLink 2.0, 100GB/s CPU-to-GPU,
GPU-to-GPU

2TB DDR4-2666 per node

1.6 TB NV RAM per node

250 PB storage

POWER9-SO, Global Foundrie 14nm FinFET, 8×10^9 tran., 17-layer, 24 cores, 96 threads (SMT4)

120MB L3 eDRAM (2 CPU 10MB), 256GB/s



Source: <http://www.tomshardware.com/>

Summit supercomputer – IBM AC922 – Volta

80 activated SMs (5,120 CUDA cores)

Volta die 815mm²,


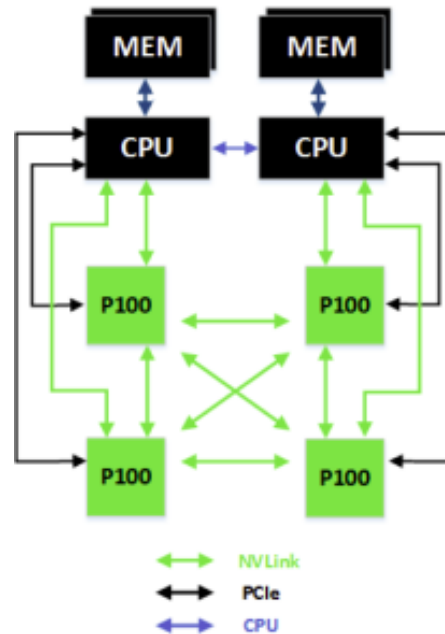
21×10⁹ tran. TSMC 12nm FFN

HBM2 (16GB total)

1.6 TB NV RAM per node

250 PB storage

300W, <1V ⇒ 300A



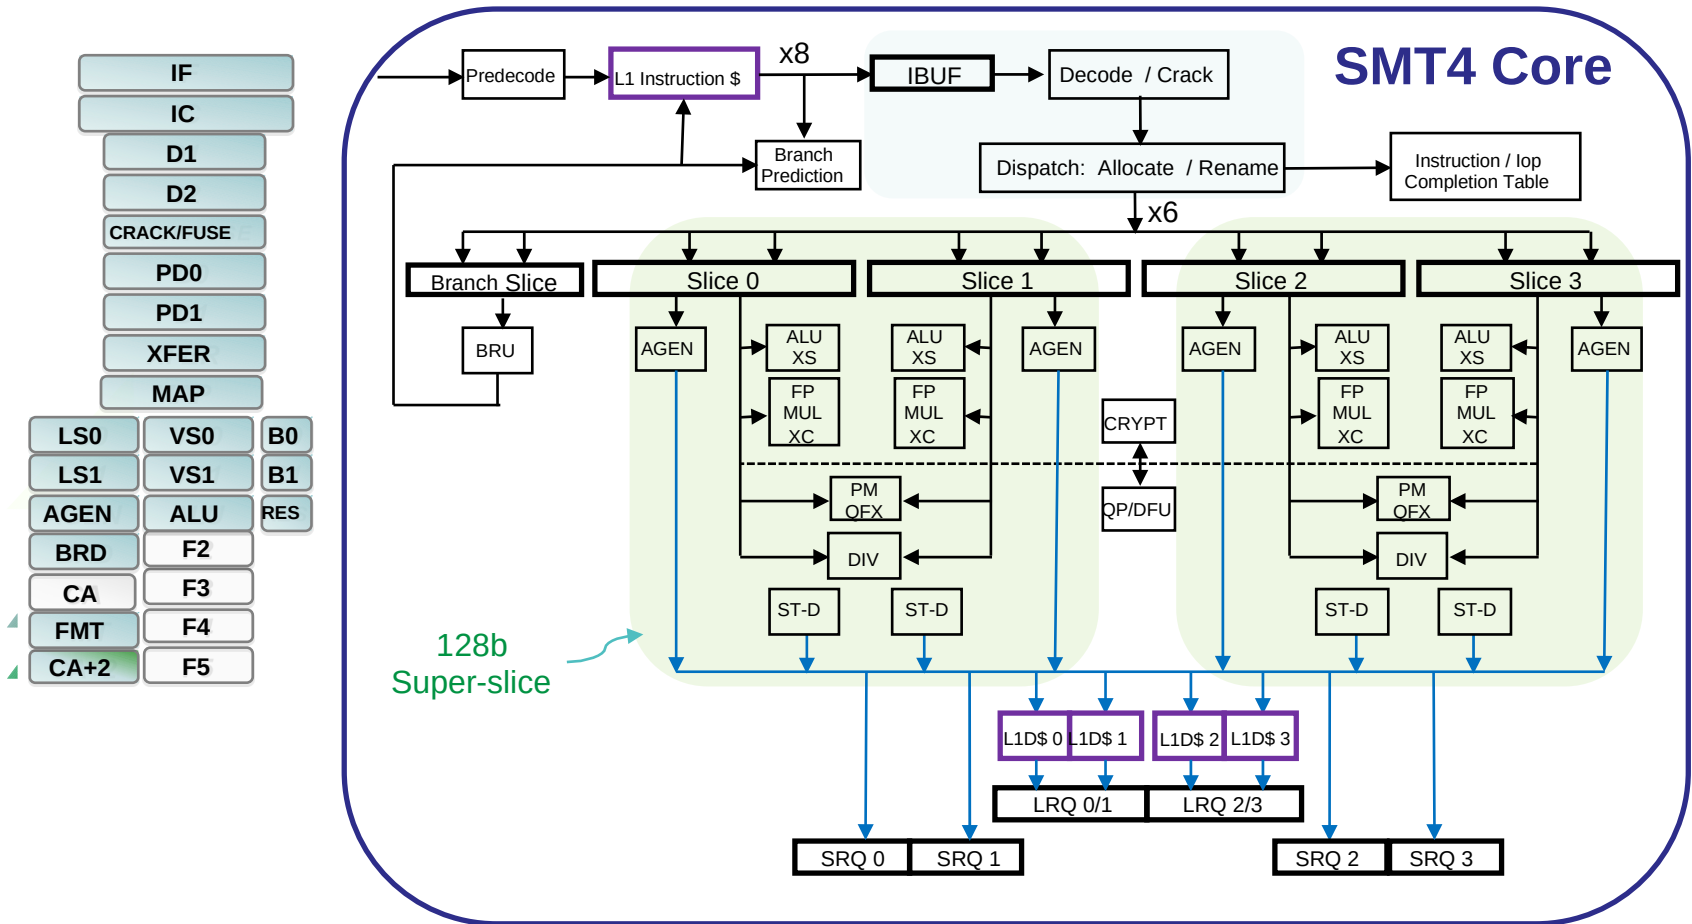
Source: <http://www.tomshardware.com/>

Power9 architecture

- L1I Cache, 32 KiB, 8-way, per SMT4 Core, line 128 (4× 32) CritSF
- L1D Cache, 32 KiB, 8-way, per SMT4 Core, line 128 (2× 64) CritSF
- L2 Cache, 258 KiB per SMT4 core
- L3 Cache, 120 MiB eDRAM, 12×10 MiB 20-way 7 TB/s
- Fetch/Branch – 8 fetch, 6 decode 1× branch execution
- Slices issue VSU & AGEN – 4× scalar-64b / 2× vector-128b
4× load/store AGEN
- VSU Pipe – 4× ALU, 4× FP + FX-MUL + Complex (64b),
2× Permute (128b), 2× Quad Fixed (128b),
2× Fixed Divide (64b), 1× Quad FP & Decimal FP,
1× Cryptography
- LSU Slices – 32 KiB L1D\$, Up to 4 DW Load or Store

Source: <https://en.wikichip.org/wiki/ibm/microarchitectures/power9>

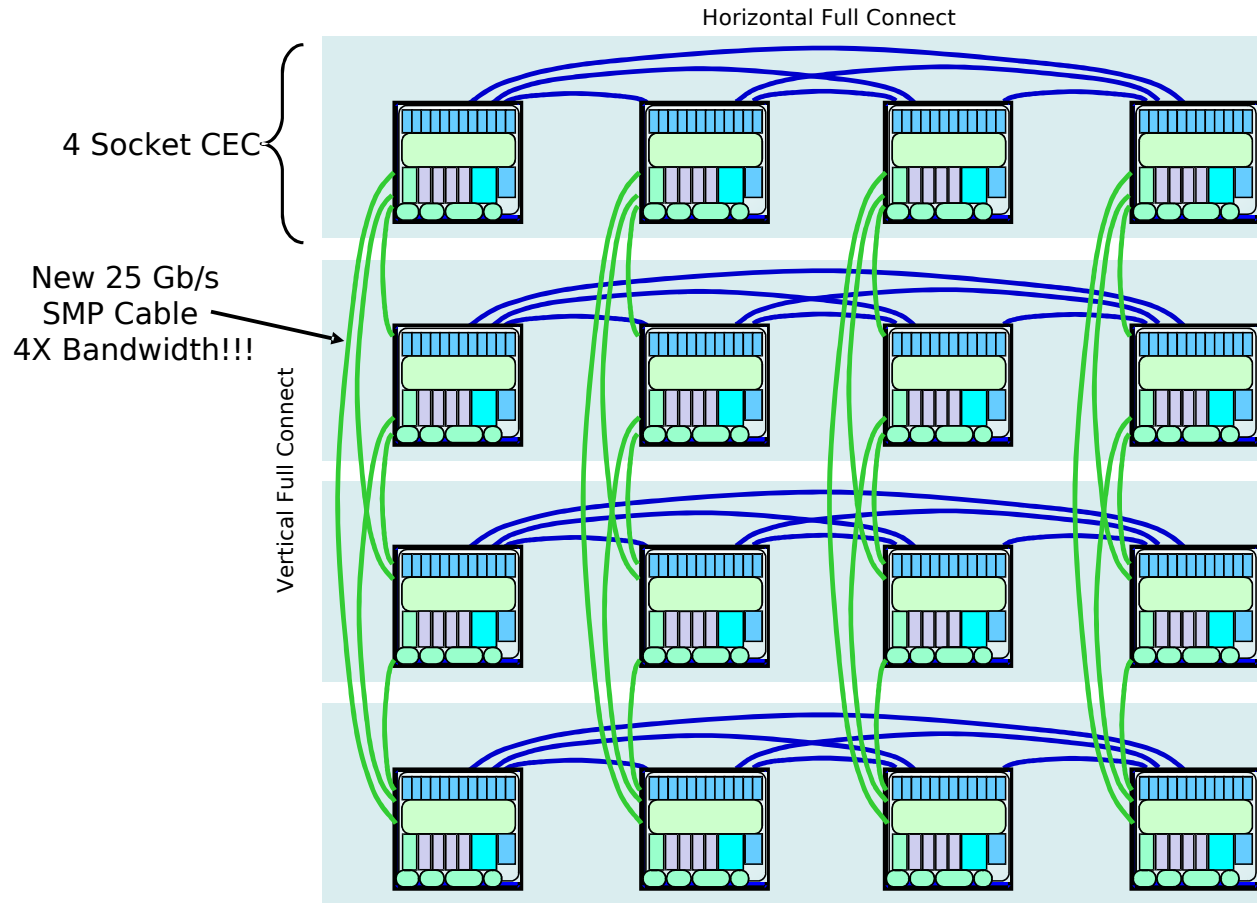
Power9 architecture – pipeline



Source: POWER8/9 Deep Dive, Jeff Stuecheli, POWER Systems, IBM Systems

Power9 architecture – Interconnect

16 Socket 2-Hop POWER9 Enterprise System Topology



Source: POWER9, Jeff Stuecheli, POWER Systems, IBM Systems

AMD Exascale Computing Technologies

Peak performance: 1.5 EFLOPS

Node: 4× AMD Radeon Instinct GPUs, 1× AMD EPYC CPU

AMD technologies

CPU-GPU Interconnect: AMD Infinity Fabric

Coherent memory across the node

System interconnect: Multiple Slingshot NICs providing 100 GB/s network bandwidth (4× Summit)

Storage: like Summit.


AMD open-source software

Cray PE, AMD ROCm, GCC

HIP C/C++, OpenMP (offload)

Delivery 2021 to the Oak Ridge

Source: <https://www.amd.com/en/products/frontier>



Big single systems image

Examples of the today biggest systems which memory is under control of single operating system kernel (single system image)

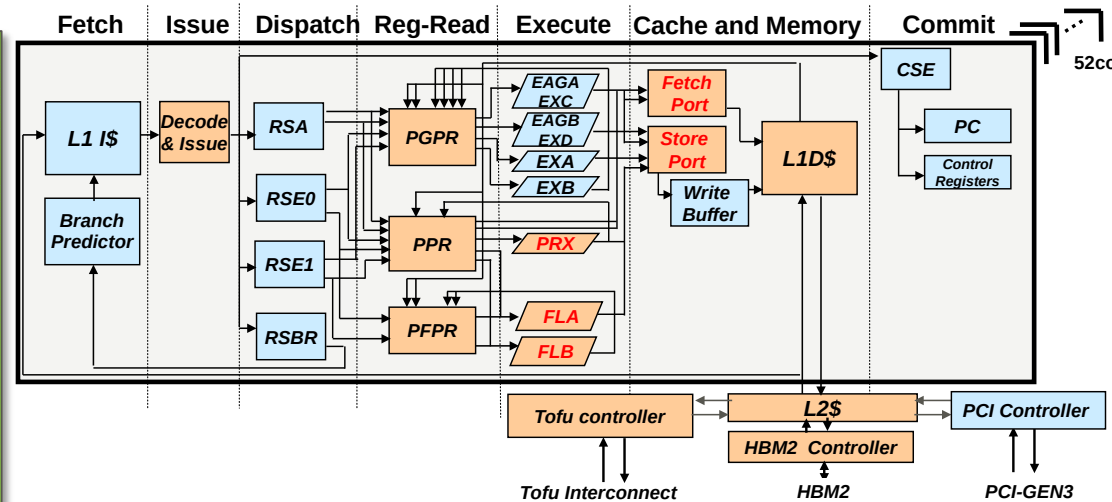
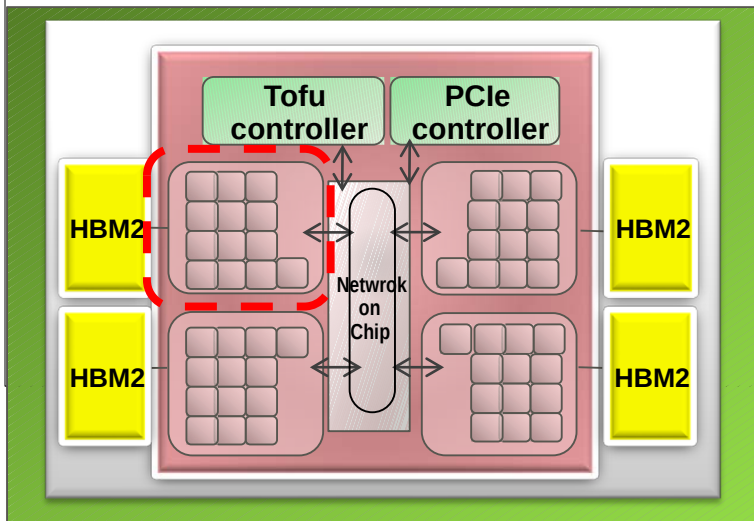
	SGI UV 2000	SGI UV 20
CPU Speed (Cores)	Intel® Xeon® processor E5-4600 product family 2.4GHz-3.3GHz	Intel® Xeon® processor E5-4600 product family 2.4GHz-3.3GHz
Min/Max Sockets	4/256	2/4
Min/Max Cores (Threads)	32/2048 (4096)	8/48
Max Memory	64TB	1.5TB
Interconnect	NUMALink® 6	Intel® Quickpath
Enclosure	10U rackmount	2U rackmount
Rack Size	Standard 19" Rack	Standard 19" Rack

Intel and AMD

	Intel Core i7-8700K	Intel Core i7-8700	Ryzen 7 1700X	Ryzen 7 1700
Socket	LGA 1151	LGA 1151	PGA 1311	PGA 1311
Cores/Threads	6 / 12	6 / 12	8 / 16	8 / 16
Base Frequency	3.7 GHz	3.2 GHz	3.4 GHz	3.0 GHz
Boost Frequency	4.7 GHz	4.6 GHz	3.8 GHz	3.7 GHz
Memory Speed	DDR4-2666	DDR4-2666	DDR4-1866 to DDR4-2667	DDR4-1866 to DDR4-2667
Memory Controller	Dual-Channel	Dual-Channel	Dual-Channel	Dual-Channel
Unlocked Multiplier	Yes	No	Yes	Yes
PCIe Lanes	x16 Gen3	x16 Gen3	x16 Gen3	x16 Gen3
Integrated Graphics	Intel UHD Graphics 630 (up to 1,200MHz)	Intel UHD Graphics 630 (up to 1,200MHz)	No	No
Cache (L2+L3)	13.5MB	13.5MB	20MB	20MB
Architecture	Coffee Lake	Coffee Lake	Zen	Zen
Process	14nm++	14nm++	14nm GloFo	14nm GloFo
TDP	95W	65W	95W	65W
Price (@1k)	\$359	\$303	\$399	\$329

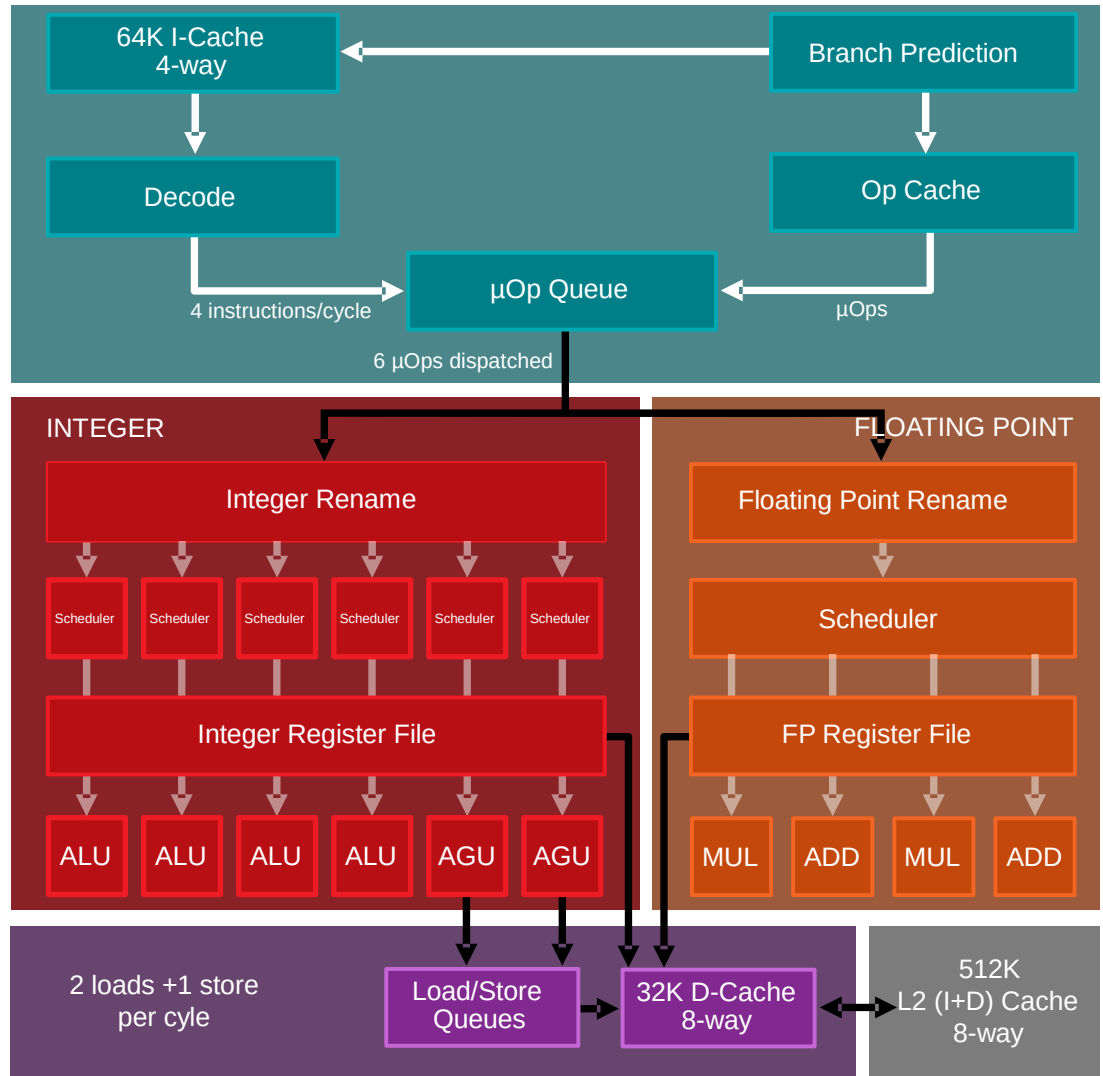
Fujitsu – Supercomputer Fugaku – A64FX, 2020 TOP500 #1

- Combine Armv8.2-A (AArch64 only) with Fujitsu supercomputer technology, SPARC64 V till now
- 48 computing cores + 4 assistant cores, SVE 512-bit wide SIMD
- HBM2 32GiB, 7nm FinFET, 8,786M transistors
- Tofu 6D Mesh/Torus, 28Gbps x 2 lanes x 10 ports, PCIe



AMD ZEN/Ryzen

- Ryzen 3 Mobile APUs: January 9th
- Ryzen Desktop APUs: February 12th
- Second Generation Ryzen Desktop Processors: April.
- Ryzen Pro Mobile APUs: Q2 2018
- Second Generation Threadripper Processors: 2H 2018
- Second Generation Ryzen Pro Desktop Processors: 2H 2018

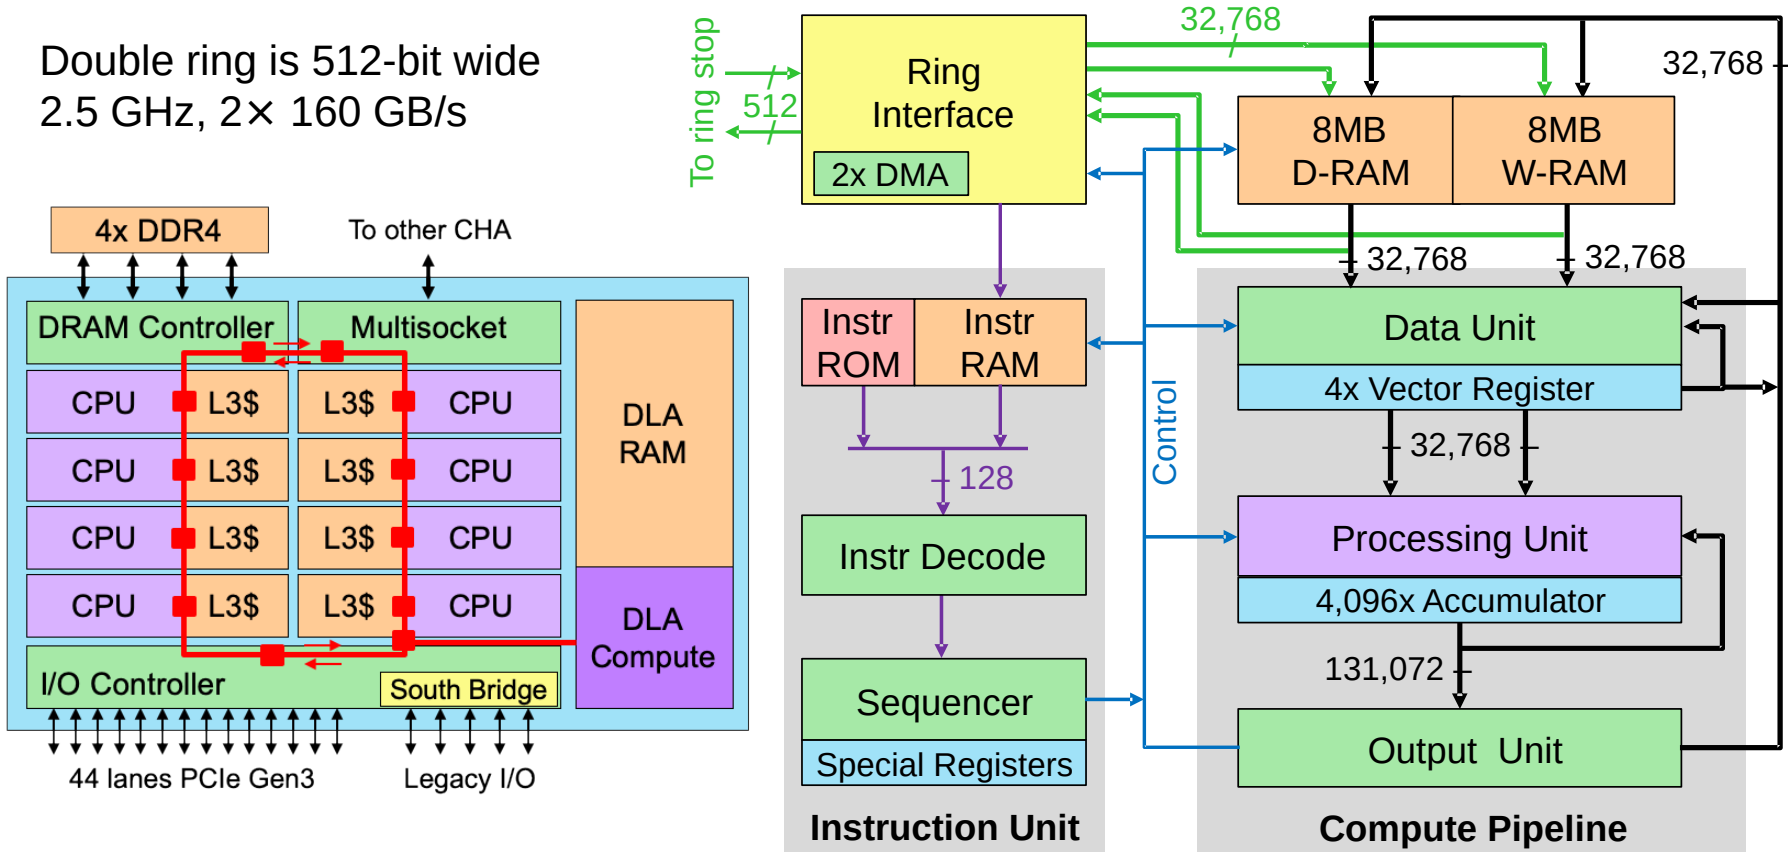


Centaur Technology, CHA, x86 for AI

16-nanometer x86 SoC

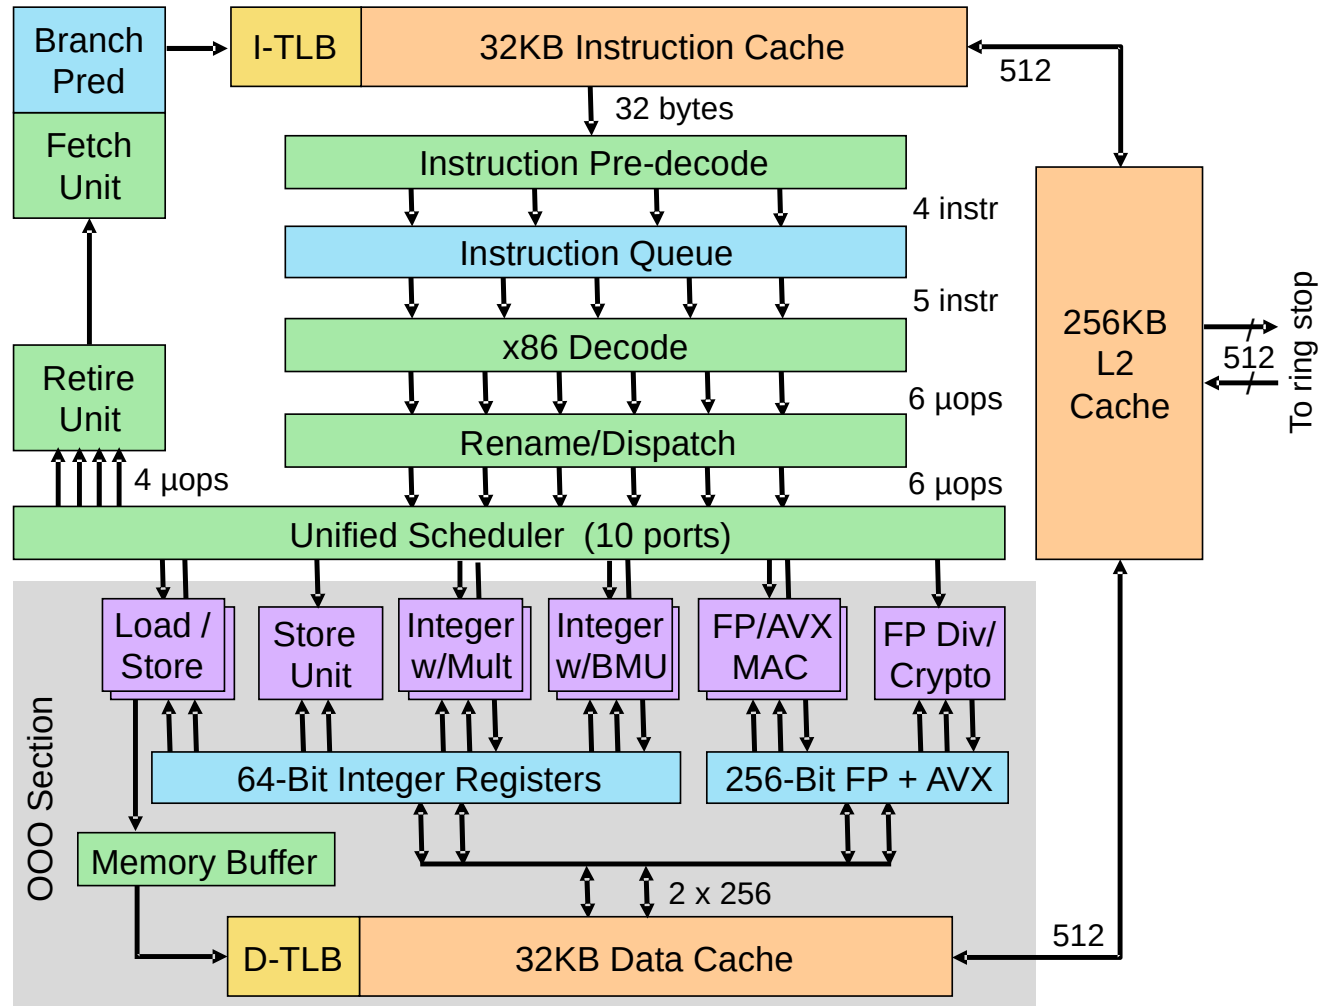
NCORE datapath 4,096-byte wide
Result accumulator 16 KiB

Double ring is 512-bit wide
2.5 GHz, 2 × 160 GB/s



source: https://www.centtech.com/wp-content/uploads/2020/10/MPR_Centaur_CHA_2020_10_12.pdf

Centaur CNS microarchitecture



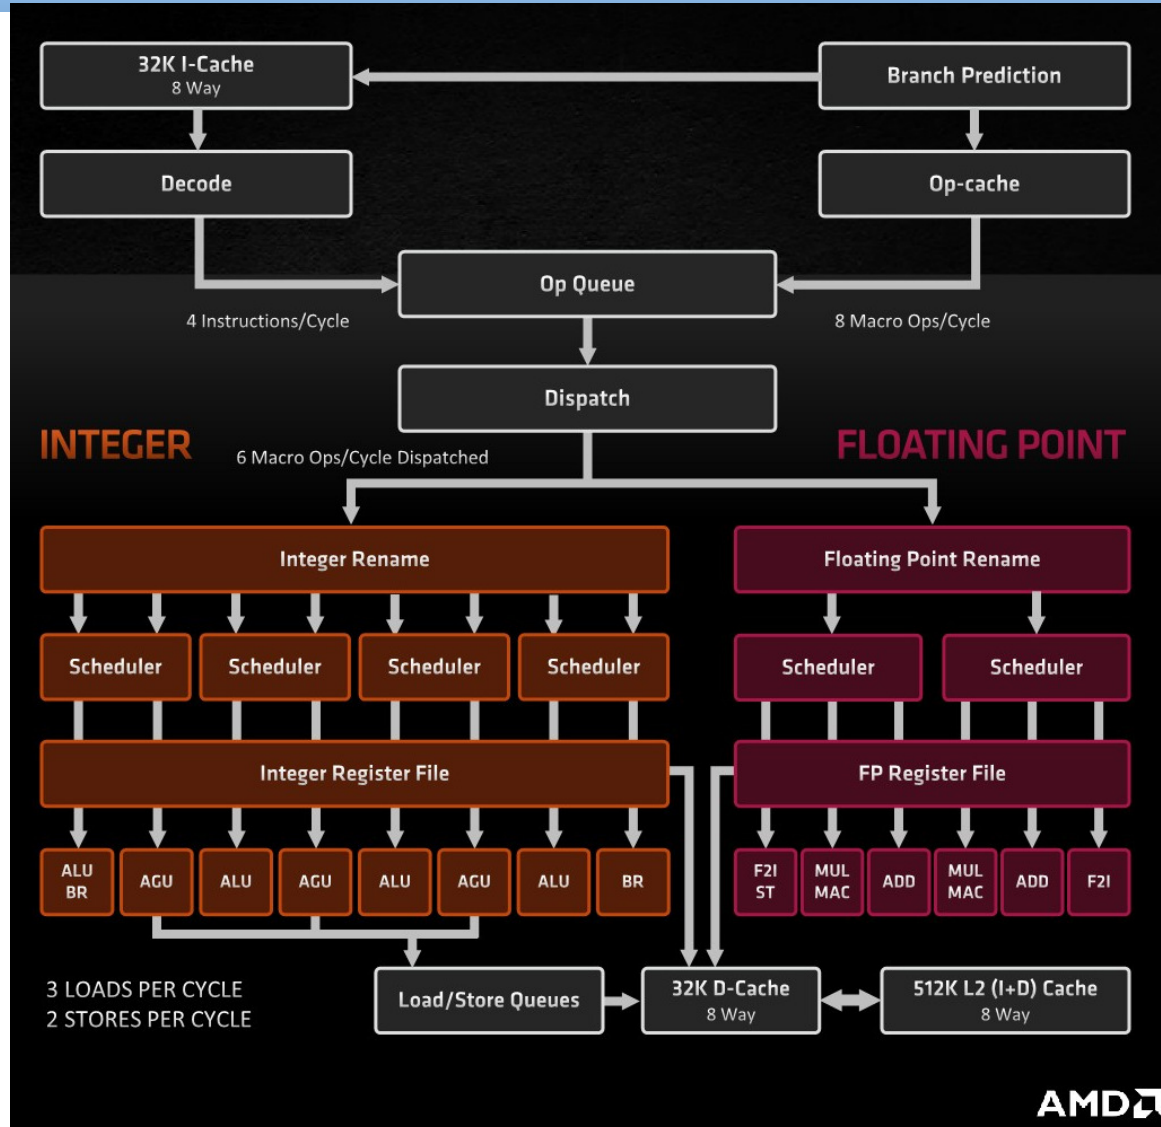
2019/20 x86 Comparison CenTauru CHA included

Company	Centaur CNS	AMD Zen	AMD Zen 2	Intel Cofee Lake	Intel Sunny Cove
L1I Capacity	32 KiB	64 KiB	32 KiB	32 KiB	32 KiB
L1I Org	8-w, 64 s	4-w, 256 s	8-w, 64 s	8-w, 64 s	8-w, 64 s
Renaming	4/cycle	6/cycle	6/cycle		5/cycle
Max in flight	192	192	224		352

AMD Zen 3 Microarchitecture, November 2020

- 8-core chiplet, shared 32 MiB L3, latency 46 cycles
- Integer register file 192 entries, 96 entry integer scheduler
- 256 entry reorder-buffer
- 6 μ OP dispatch width
- Issue width 10, 1 dedicated branch, 2 separated store data pathways
- 4 cycles for fused-multiply-add-ops
- div/mod - 8 bit \rightarrow 10 cycles, 16 \rightarrow 12, 32 \rightarrow 14, 64 \rightarrow 20
- BTB 1024 entries
- TLB L1 64 I&D all page sizes, TLB L2 512 I, 2K D no 1G
- +19% IPC compared to Zen 2


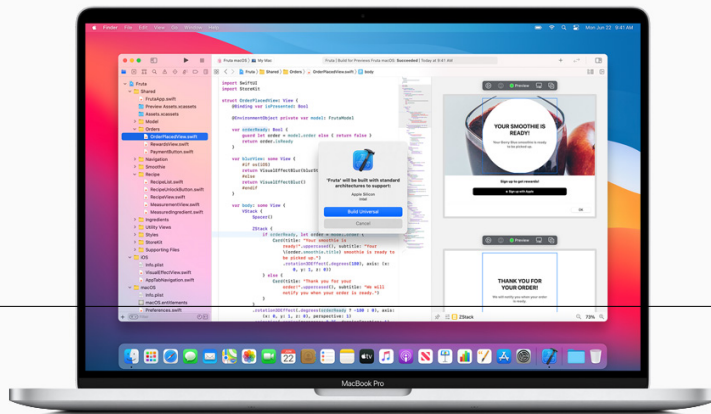
Zen 3 diagram



Source: <https://www.anandtech.com/print/16214/amd-zen-3-ryzen-deep-dive-review-5950x-5900x-5800x-and-5700x-tested>

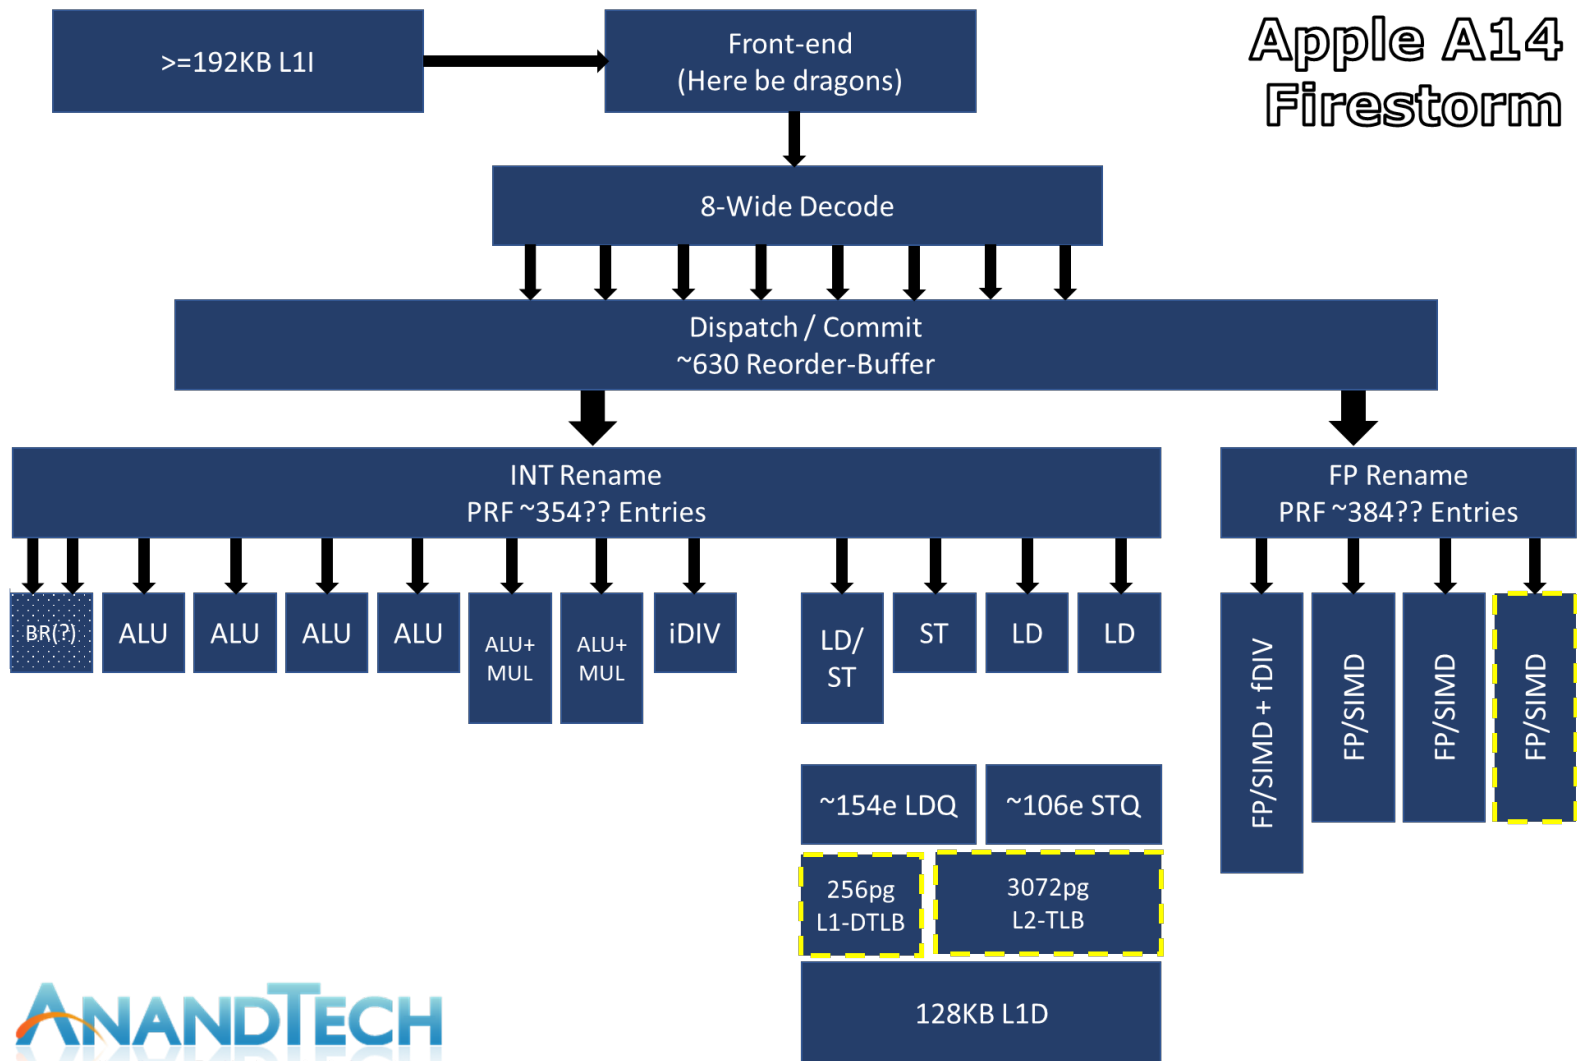
Apple A12Z Bionic – 64-bit ARM-based

- People who are really serious about software should make their own hardware. *Alan Kay*
- *Apple A12Z, 8 cores (ARM big.LITTLE: 4 "big" Vortex + 4 "little" Tempest), Max. 2.49 GHz, ARMv8.3-A*
- *Cache L1 128 KB instruction, 128 KB datam L2 8 MB*
- *GPU Apple designed 8-Core*



Apple M1, A14, 4 Firestorm, 4 Icestorm

Apple A14
Firestorm



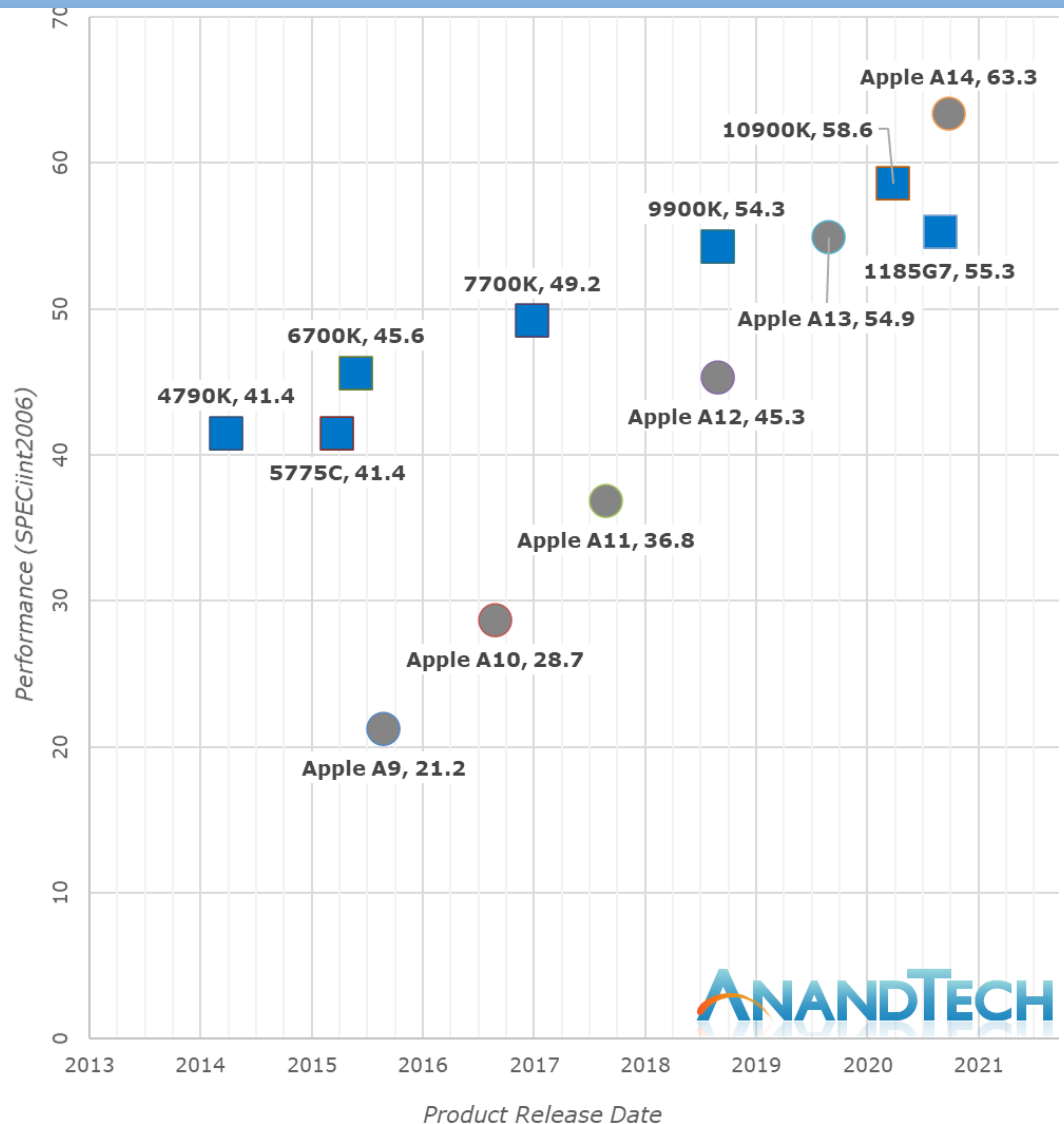
Source: <https://www.anandtech.com/show/16226/apple-silicon-m1-a14-deep-dive>

Intel versus Apple Top Single Thread Performance

Last 5 years

Intel +28%

Apple +198%
2.98x



Data storage, future directions

It is possible that current classical file systems, based on the concept of block devices and the transfer of part of data to the pagecache, will be completely inappropriate for future computer systems

New memory technologies that allow access by byte granularity, can be mapped directly into the physical address space of the CPU and are almost as fast as conventional SDRAM chips. Therefore, it is not necessary to copy data to access them from applications. However, use as a classical PFN is problematic because there is too much memory required for service structures in SDRAM and new media is used for service structures, it will also experience physical wear and tear. An interesting view of the issue is described, for example, in the article

XFS: There and back ... and there again?

<https://lwn.net/Articles/638546/>

But we also need to be thinking a little further ahead. Looking at the progression of capacities and access times for "spinning rust" shows 8GB, 7ms drives in the mid-1990s and 8TB, 15ms drives in the mid-2010s. That suggests that the mid-2030s will have 8PB (petabyte, 1000 terabytes) drives with 30ms access times.

The progression in solid-state drives (SSDs) shows slow, unreliable, and "damn expensive" 30GB drives in 2005. Those drives were roughly \$10/GB, but today's rack-mounted (3U) 512TB SSDs are less than \$1/GB and can achieve 7GB/second performance. That suggests to him that by 2025 we will have 3U SSDs with 8EB (exabyte, 1000 petabytes) capacity at \$0.1/GB.

persistent memory

NVDIMM battery-backed DIMMs 8 GB

400GB Memristors

New data storage concepts, 3D XPoint

- NVM Intel and Micron Technology 2015/2017
- based on a change of bulk resistance
- Intel Optane Memory 16 / 32 GB
- M.2 2280 - PCIe 3.0 x2 NVMe
- Read latency 6 μ s, Write Latency 16 μ s
Read seq/rand 1200 MB/s, Write seq/rand 280 MB/s (4 kB)
- Endurance
100 GB/day
- The future format
same as NV DIMM



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

Quantum computers

- IBM, 50-qubit quantum chip
- Intel, 49-qubit Tangle Lake, superconducting quantum chip
- 1,000-qubit in 5 to 7 years perspective
- real commercial usability probably when million qubit scale is reached
- superconducting qubit × qubits in silicon.