

DEPARTMENT: NOVEL ARCHITECTURES

The European Factor: From ARM to Atos

Anne C. Elster , Norwegian University of Science and Technology (NTNU), Trondheim, Norway

It has been both a strange and unsettling year with a lot of uncertainties related to COVID-19. Both the premier supercomputing conferences ISC'20 and SC20 went virtual, preventing users from exploring the show floors with their examples of novel architectures and seeing the demos of their related tools. Companies have also slowed their procurement of high-end systems during 2020, as was discussed during the presentation of top500.org statistics at the Top-500 Birds-of-a-Feather session at SC20. The Top-500 list is a measurement of the world's largest computers using the High-Performance Linpack benchmark (see <http://www.top500.org>). The systems must solve large dense systems of linear equations without using the operation-saving Strassen algorithm or a mix of lower precision, but performing the calculations conforming to LU factorization with partial pivoting using $2/3n^3 + O(n^2)$ double-precision (64-b) floating point numbers.

Despite less turnover of the world's largest computer systems, technologies continue to advance, and the trends are more interesting and diverse than ever.

The US dominated the computing industry—including the supercomputer industry—for years, until Japan gave the US a jolt with the introduction of the Earth Simulator back in 2002. This put supercomputing back in the spotlight, especially not only in the US but also in Europe. For the past decade, the US, Japan, and China have dominated at the top of the Top-500 list. As can be seen from Table 1, these systems were built primarily on US-designed chips from IBM and Intel, coupled with accelerators from NVIDIA. The AMD-based *Titan* Cray/Hewlett Packard Enterprise (HPE) system and the *Sunway* system in China, as well as the new ARM-based Fujitsu *Fugaku* supercomputer, are notable exceptions. In this article, we will discuss the history of

ARM and its current impact on HPC as well as some other European trends.

ARM—FROM RISC AND EMBEDDED SYSTEMS TO SUPERCOMPUTERS

ARM is a 30-year-old U.K.-based semiconductor Intellectual Property (IP) company that started as Advanced RISC Machines Ltd., a joint venture that included former U.K. company Acorn Computers, Apple Computers (now Apple Inc.),¹ and VLSI Technology (a US company later bought by Phillips for USD 1 billion and now part of the Philips spin-off NXP Semiconductors). ARM was bought in 2016 by the Japanese conglomerate and investing company Softbank Group, which also has a main stake in US Sprint and also has significant T-Mobile stock via the merger of Sprint and T-Mobile that was completed in April 2020. In September 2020, NVIDIA announced that it is buying ARM,² but it may take time to complete since it is running into some licensing issues regarding its subsidiary ARM-China.

Unlike Intel, AMD, and Freescale, ARM does not produce any computer chips, but licenses its architecture designs and product ecosystem to other vendors. It has particularly focused on low-power CPUs and GPUs for mobile and embedded devices, and has over the last several years acquired dozens of smaller companies, such as Falanx, a Norwegian startup that is now known as ARM Norway and is behind their Mali GPU processor.

ARM now offers a wide range of IP design products, including CPUs, GPUs, and microcontrollers. ARM CPUs and neural processing units include Cortex-A, Cortex-M, Cortex-R, Neoverse, Ethos, and SecureCore. Mobile companies such as Apple and Samsung often license the designs from ARM to manufacture and integrate them into their own system-on-chip designs with other components such as GPUs (sometimes ARM's Mali) or modem/radio baseband components (for mobile phones). In fact, due to the proliferation of mobile and handheld devices such as tablets, ARM was claimed to have 75% of the world's CPU market in 2005. As the demand for processing power of these devices has increased so has the

TABLE 1. Top500 no. 1 systems 2010–2020.

Name	Dates no. 1	Where	Architecture	Cores
Tianhe-1A	Nov. 2010	China	NUDT THMPP, Intel Xeon X5670 @ 2.93 GHz, Infiniband FDR FT-1000, Nvidia 2050 GPUs	186 368
K-Computer	June 2011–Nov. 2011	Japan	Fujitsu SPARC64 VIIIfx @ 2.0GHz with Tofu interconn.	705 024
Sequoia	June 2012	US	IBM BlueGene/Q, IBM Power BQC 16C @ 1.6GHz, Custom interconnect	1 572 864
Titan	Nov. 2012	US	Cray/HPE Cray XK7, AMD Opteron 6274 16C @ 2.2GHz w/ Cray Gemini interconn., Nvidia K20x GPUs	560 640
Tianhe-2A	June 2013–Nov. 2015	China	Intel Xeon, TH Express-2, Matrix 2000	4 981 760
Sunway TaihuLight	June 2016–Nov. 2017	China	Sunway MPP , Sunway SW26010 260C, NRCPC	10 649 600
Summit	June 2018–Nov. 2019	US	IBM AC922 Power9 + Nvidia GV100 GPUs, Mellanox	2 424 492
Fugaku SC	June 2020–Nov. 2020	Japan	Fujitsu/RIKEN ARM-based A64FX 48C @ 2.2 GHz w/ Tofu interconnects	7 299 072 (Jun.) 7 630 848 (Nov.)

demand for their CPUs; since 2018, ARM has also offered server-class CPU designs. If NVIDIA's purchase of ARM goes through, it will be very interesting to see how it impacts future ARM designs that are already influenced by AI workloads.

ARM-BASED ASTRA

The first ARM-based supercomputer to appear on the Top-500 list was *Astra* at Sandia National Labs, announced at SC18. The HPE-built system ranked at number 204 with 1.529 petaFLOPS, and made number 36 on the HP Conjugate Gradient benchmark, a benchmark more targeted to its use. *Astra* was the first system that was part of the US DOE National Nuclear Security Administration Vanguard program that looks at prototype systems for advanced architectures.

One of the *Astra* nodes is about 100 times faster than the ARM-based CPU chips found in cell phones. The *Astra* system has 5184 ARM-based Cavium ThunderX2 28-core processors. Cavium Inc. grew from designing networking processors to itself being bought in 2018 by Marvell Technology group, a Bermuda-registered conglomerate.

Michael Aguilar, the main systems administrator of *Astra*, gave a nice introduction to the system in 2019 at Stanford.³ He describes how they have ported many of the familiar packages ranging from OpenMPI to UCX, where one of the interesting challenges they faced was

that, unlike Intel cores, the ARM cores did not appear next to each other, so they behaved in a very nonuniform fashion. Socket-direct, a new feature on the architecture, seems to alleviate many performance issues associated with core locality, but requires a new awareness at this level. They have had good results with the system so far, achieving from 1.6X speedup on Monte Carlo to 1.87 on linear solver versus the Intel Haswell-based *Trinity* ASC platform they used as a baseline.

However, as far as power consumption, they did not see it particularly more efficient than the Intel ones, but have not been focusing on this either.

FUGAKU AND ARM

In Summer 2020, the new Japanese *Fugaku* Supercomputer at RIKEN gained the top spot on the Top-500 list. Named "High Peak" after Mt. Fuji, its ARM-based A64FX 48C chips are jointly developed by RIKEN and Fujitsu and uses Fujitsu's Tofu (torus fusion) interconnects. The A64 chip is a Taiwan Semiconductor Manufacturing Company (TSMC) 7-nm FinFET with 8.8 billion transistors and 594 signal pins. It also has SVE 512 × 2 vector extensions.⁴

Note that one of the biggest differences between x86 chips from Intel and AMD, from a vector operation point of view, is that modern Intel chips have one 512-b AVX2 vector unit, whereas AMD chips have two to three 128-b units.

Fugaku maintained its number-1 spot in November 2020,⁵ but had, unlike the three next systems (*Summit*, *Sierra*, and *Sunway*), added over 300,000 cores (from 7,299,072 cores to 7,630,848 cores). Its A64 chip has in addition to the 48-core compute unit, two or four assist units that run the operating system.

The *Fugaku* system is also very energy efficient, ranking number 10 on the Green500 list, unlike *Summit*, *Sierra*, and *Sunway*. *Selene*, the Nvidia DGX A100 system at NVIDIA with 555,520 cores, ranked in November 2020 at number 5 both on the Top-500 and Green500 lists.

The Fujitsu A64fx is also featured in Stony Brook's HPE Apollo 80/Cray computer *Ookami* ("wolf" in Japanese), announced in November 2020. (Cray was formally acquired by HPE in September 2019.) The *Ookami* system is the first system announced outside Japan that will feature the A64fx.

ATOS AND AMD EPYC

Another interesting entry on the November 2020 Top-500 list is the number-7 *JUWELS Booster Module*, an Atos Bull Sequena XH200 based on the AMD EPYC chips installed at Forschungszentrum Jülich in Germany.

Atos is a French company that emerged from the merger of two French companies and a Dutch IT company. It has since grown with the merger/acquisition of several companies, including Siemens IT (2011), Bull, and Xerox ITO (2014), the latter a well-known HPC vendor in Europe. It has also strengthened its US presence through the acquisition of Xerox ITO (2014) and Syntel (2018).

As a historical note, Heide^{6,7} describes how both Bull and IBM have their roots dating back to punch-card machines. Bull is in fact named after the Norwegian pioneer, Fredrik Rosing Bull who in 1918 started to build punched-card machines. Bull died in 1924, but his work was continued by others in Norway, and later (1931) it was relocated to Paris, France, where it, after several reorganizations, emerged as Bull, a subsidiary of Atos Technologies.

In May 2020, Atos announced that it was adding the new NVIDIA A100 GPUs to the Jülich system,⁸ which clearly contributed to its number 7 November 2020 ranking on the Top-500.

THE NANOMETER RACE: INTEL VERSUS TSMC

Currently, just three chip manufacturers in the world are using less than 14-nm node technologies: Intel, TSMC, and Samsung. China's SMIC has yet to do smaller than 14-nm processes. The nanometer denomination previously referred to gate length, but has in the last several years been tied to "technology node" size and become more vendor dependent. For instance, Intel's

10-nm nodes are assumed to be compatible with foundries' 7-nm nodes.

So How Is This Affecting Supercomputing?

Fugaku and AMD's recent EPYC 7-nm chip are both manufactured at TSMC, which also produces the new 5-nm M1 chip for Apple. However, TSMC reports it is struggling to meet the demand for the Apple chips, so Apple may have to again also use Samsung.

Intel announced through its earnings statement in July 2020 that it is having issues with its 7-nm manufacturing, thus causing an expected 1-year or more delay of its GPU Ponte Vecchio chips. This may again delay the *Aurora* system planned at Argonne National Labs,^{9,10} one of the two new exascale systems planned in the US in 2021. The AMD-powered *Frontier* system at Oak Ridge National Laboratory is supposedly still on schedule, and will likely become the first exascale system in the US.

Of course, if we allow for low precision or mixed precision and look at the HPL-AI benchmarks—differing from the Linpack benchmarks used for the Top-500 rankings that only allow for full 64-b precisions—we are already in the exascale era. *Fugaku* also ranked number 1 on this benchmark in November 2020 peaking at 2.0 exaFLOPS.¹¹

EUROPEAN HPC JOINT UNDERTAKING (EUROHPC JU) AND THE NEXT WAVE OF EUROPEAN SUPERCOMPUTERS

Europe have had some systems in the top ten of the Top-500 list, including the current Atos systems, but is now pushing hard for their new EuroHPC JU with a roadmap that includes the following systems.

Pre-Exascale Class:

- Large Unified Modern Infrastructure (LUMI) at CSC's center in Kajaani, Finland.^{12,13}
- *Marenostrum 5* at Barcelona Supercomputing Centre, Spain.
- *Leonardo* at CINECA, Italy.¹⁴

Petascale Supercomputers:

- LuxProvide, Luxembourg.
- IZUM, Slovenia.
- IT4Innovations National Supercomputing Center, Czech Republic.
- Sofiatech, Bulgaria.
- Minho Advanced Computing Centre, Portugal.

LUMI will be an HPE Cray EX supercomputer with AMD 64-core EPYC CPUs and feature the new AMD

Instinct GPUs, and is projected to be faster than the current number 1 *Fugaku* supercomputer. And *Leonardo* is now announced to be a 240-petaFLOPS Atos' BullSequana XH2000 with ParTec's ParaStation Modulo software. *Marenostrum 5's* design has not been announced as of this article's writing.

Field-Programmable Gate Array (FPGA)—Finally Maturing for HPC?

FPGAs that let you encode algorithms and functions in hardware offer great potential speeds for user application codes, but have been, and still are, fairly hard to program for scientists. However, two of the world's largest CPU makers have bought FPGA companies (AMD announced in October 2020 that it will be buying the FPGA company Xilinx, Inc.). Although Xilinx is headquartered in San Jose, it has a strong presence in Europe and Asia, and has also started to offer FPGA-based accelerator cards targeting AI and cloud workloads. The deal is expected to go through by the end of 2021. Xilinx's new data center accelerator card, the Alveo FPGA U250, targets AI/ML, graphics computing, and HPC computations.¹⁵

FPGAs and other emerging accelerators will be discussed in a future department feature.

REFERENCES

1. Accessed: Jan. 15, 2021. [Online]. Available: <https://www.latimes.com/archives/la-xpm-1990-11-28-fi-4993-story.html>
2. Accessed: Jan. 15, 2021. [Online]. Available: <https://www.globenewswire.com/news-release/2020/09/13/2092654/0/en/NVIDIA-to-Acquire-Arm-for-40-Billion-Creating-World-s-Premier-Computing-Company-for-the-Age-of-AI.html>
3. Accessed: Jan. 15, 2021. [Online]. Available: <https://insidehpc.com/2019/02/astra-a-large-scale-arm64-hpc-deployment/>
4. Accessed: Jan. 15, 2021. [Online]. Available: <https://www.fujitsu.com/global/about/innovation/fugaku/>
5. Accessed: Jan. 15, 2021. [Online]. Available: <https://www.japantimes.co.jp/news/2020/11/17/national/japan-fugaku-fastest-supercomputer/>
6. L. Heide, "From invention to production: The development of punched-card machines by F. R. Bull and K. A. Knutsen 1918–1930," *Ann. Hist. Comput.*, vol. 13, no. 3, pp. 261–272, 1991. doi: [10.1109/mahc.1991.10024](https://doi.org/10.1109/mahc.1991.10024).
7. L. Heide, "Punched cards for professional European offices: Revisiting the dynamics of information technology diffusion from the United States to Europe 1889–1918," *Hist. Technol.*, vol. 24, no. 4, pp. 307–320, 2008. doi: [10.1080/07341510802044686](https://doi.org/10.1080/07341510802044686).
8. Accessed: Jan. 15, 2021. [Online]. Available: https://atos.net/en/2020/press-release/general-press-releases_2020_05_14/atos-launches-first-supercomputer-equipped-with-nvidia-a100-gpu
9. Accessed: Jan. 15, 2021. [Online]. Available: <https://www.datacenterdynamics.com/en/news/intel-delays-7nm-considers-using-rival-foundries/>
10. Accessed: Jan. 15, 2021. [Online]. Available: <https://phonemantra.com/amd-powered-frontier-may-become-first-us-exascale-supercomputer-due-to-intels-7nm-process-issues/>
11. Accessed: Jan. 15, 2021. [Online]. Available: <https://spectrum.ieee.org/tech-talk/computing/hardware/japans-fugaku-supercomputer-is-first-in-the-world-to-simultaneously-top-all-high-performance-benchmarks>
12. Accessed: Jan. 15, 2021. [Online]. Available: <https://www.lumi-supercomputer.eu/lumi-one-of-the-worlds-mightiest-supercomputers/value>
13. Accessed: Jan. 15, 2021. [Online]. Available: <https://www.hpcwire.com/2020/11/19/eurohpc-exec-dir-talks-procurement-epi-and-europes-efforts-to-control-its-hpc-destiny/value>
14. Accessed: Jan. 15, 2021. [Online]. Available: <https://insidehpc.com/2020/11/cineca-partners-with-atos-for-240pflops-leonardo-supercomputer/>
15. Accessed: Jan. 15, 2021. [Online]. Available: <https://www.xilinx.com/products/boards-and-kits/alveo/u250.html>

ANNE C. ELSTER (Senior Member, IEEE) was born near the Arctic Circle in Mo i Rana, Norway. Since 2001, she has been a faculty member in computer science with the Norwegian University of Science and Technology (NTNU), Trondheim, Norway, where she is currently a Full Professor and has held various leadership positions. She was with Schlumberger ASC, Austin, TX, USA, where she also held appointments at The University of Texas at Austin and started her own company before returning to Norway in 2001. She has also authored several book chapter and articles, as well as advised more than 90 graduate students. She received the bachelor's degree in computing systems engineering from The University of Massachusetts at Amherst, MA, USA, in 1985, and the master's and Ph.D. degrees in electrical engineering from Cornell University, NY, USA, in 1988 and 1994, respectively. She is an IEEE CS Distinguished Visitor Program Speaker (2019–2022) for both the North America and Europe, Middle East and Africa Regions. Dr. Elster is a Life Member of the AGU, and a member of IEEE, IEEE CS, SIAM, Tekna (Norwegian Engineering Society), among others. Contact her at anne.elster@gmail.com.