

A long road to Exascale

Posted on [Sunday, 11 August 2013](#) by [Marko](#) — [No Comments](#) ↓



Curie supercomputer
© GENCI/CEA

There has been a lot of buzz these days about future supercomputing systems. Articles in the press, news feeds, conference sessions, they all talk about the need to drive forward in order to support growing computing demands. A new supercomputer has been introduced in China, taking the first spot at world's [TOP500](#) list [1], the UK announced a big contract with Cray to build next generation, national supercomputer, big systems operate both in USA and Japan, and in Europe [PRACE](#) lead the way with two supercomputers in the top 10.

Performance of modern supercomputers is measured in [FLOPS](#). The petaflops barrier was reached a while ago, so both scientific community and wider public have set a new landmark – a supercomputer of exascale power. But, how are we going to achieve this? It remains to be seen. At the moment, world's fastest supercomputer, China's Tianhe-2 system, has performance of 33.86 petaflop/s on the Linpack benchmark, while the second one, USA's Titan, achieved 17.59 petaflop/s. When I came here to UK, and started to work with HECToR, I thought it was a really fast machine, and it achieves “only” around 360 teraflop/s.

So, we have to build a supercomputer that is only about 1000 times faster than the current ones. Can we do that with the present technology? Theoretically, yes, we could that with state-of-the-art x86 processors. But, as Bill Dally, NVIDIA's chief scientist, pointed out at ISC'13 conference in Leipzig [2], that computer would require about 2GW of power, which is more or less the whole capacity of the Danube's biggest hydroelectric power station, [Iron Gate I](#). The same computer built as a combination of traditional x86 processors and accelerators like NVIDIA Kepler K20 or Intel Xeon Phi would consume “only” 150 MW of power. It's not only about feasibility, but also about the environmental impact. As someone said, supercomputers should simulate the climate, not change it.

- [Caoimhín](#)(2)
- [Elisavet](#)(1)
- [Evguenia](#)(3)
- [Ezhilmathi](#)(3)
- [Hannes](#)(3)
- [Jasper](#)(2)
- [Lukáš](#)(4)
- [Luna](#)(3)
- [Maria](#)(2)
- [Marko](#)(3)
- [Matteo](#)(3)
- [Mordechai](#)(6)
- [Nathan](#)(2)
- [Nicolas](#)(4)
- [Niki](#)(2)
- [Onur](#)(3)
- [Peter](#)(2)
- [Simone](#)(3)
- [SoHPC Team](#)(5)
- [Stamatia](#)(2)
- [Vito](#)(2)
- [Vojtech](#)(2)

Participants Profiles

- [Participants](#)

Many point out that developing exascale supercomputers would be an interesting challenge, and that it will require efforts to both reduce power, and improve programming models, operating systems, and processor architecture. It is widely believed that existing hardware and software will not scale to exascale level, so it is up to academia and industry to bring fresh, innovative ideas.

In order to achieve better energy efficiency it has already been clear that such a computer would combine CPUs together with accelerators. But, could we potentially save more energy if we learn some lessons from the mobile market? Researchers at the Barcelona Supercomputing Center launched the [Mont Blanc](#) project [3] which had success in building an HPC system with ARM processors. Scientists at Argonne National Laboratory, USA, are working on new, hierarchical, exascale OS, called Argo [4], while new programming models like MPI/OmpSs and PGAS (Partitioned global address space) are becoming more and more important.

As Daniel Reed from ACM labelled in his article [5], we need a “catastrophe” – a completely new approach to current architectural and programming challenges that will bring us new power by the end of the decade.

–

[1] TOP500 list, <http://www.top500.org>

[2] Improving Power and Programming: Keys to Exascale Kingdom, <http://blogs.nvidia.com/blog/2013/06/21/improving-power-and-programming-keys-to-the-exascale-kingdom/>

[3] Mont Blanc project, <https://www.montblanc-project.eu/home>

[4] A Sneak Peek at the Next-Gen Exascale Operating System, http://www.hpcwire.com/hpcwire/2013-07-31/an_early_peek_at_an_exascale_operating_system.html

[5] Leaping the Exascale Chasm, <http://cacm.acm.org/blogs/blog-cacm/166121-leaping-the-exascale-chasm/fulltext>

◀ My story with Johannes

Magic behind the Most of the CFD solvers for HPC ▶

Tagged with: computer architecture, exascale, power efficiency, PRACE, programming models, supercomputers

Leave a Reply

Your email address will not be published. Required fields are marked *

Name *

E-mail *

Website

Comment

Post Comment

