Vittoria Rezzonico
EPFL/CSIN
Information Systems Coordination

Vincent Keller
CADMOS
Center for Advanced Modeling Science

EDITOR'S LETTER

In 2008, the PetaFLOPs (10^{15} Floating Point Operations per second) barrier was broken, and we expect to reach the ExaFLOPs milestone sometime around 2020. As for ZettaScale, we won't get there before Zetta's 20th anniversary, if Gordon Moore is right.

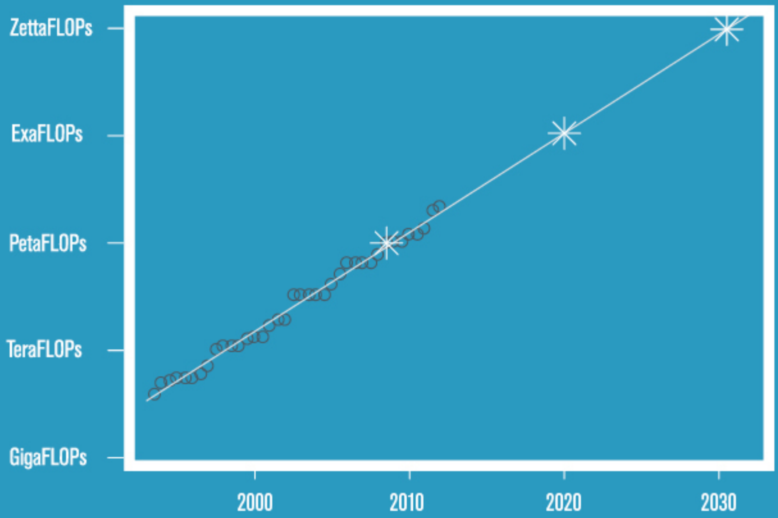

Evolution of the computational power of the most powerful machine in the world, the latest milestone reached and the next two. An iPhone has a computational power of the order of the GigaFLOPs. A standard desktop computer today reaches 100 GigaFLOPs and would have been the most powerful computer in the world in 1993.

If we admit that (1) all of the problems which arise with a smaller machine have been solved and (2) science needs such a machine, then the most obvious problems we have to solve in order to achieve the next milestones are:

— the increasing number of cores per machine, leading to shorter Mean Time Between Failures (of the order of a couple of hours: not even enough to run a benchmark);
— the power consumption of one machine;
— the knowledge it takes to handle an increasing number of cores.

The last issue is already being tackled: we are educating students to think differently in order to use those machines. Parallel thinking, math knowledge, engineering and necessary programming techniques will be some of the students' capacities.

As for the power consumption, we have to admit that most of today's supercomputers are composed of commodity parts you can buy in bulk off the Internet. Maybe that's where the problem resides: the applications have to adapt to hardware that, originally, was not designed to handle the specific requirements. Chip founders and HPC system assemblers are aware of this problem. It has lead to a new discipline — namely co-design — which enforces the latter to take into account applications' requirements in their development.

Suppose we have solved both the resiliency problem and the energy consumption issue. In order for either an Exa- or ZettaScale machine to be useful, we must be able to take advantage of millions of cores. Not only does this require advanced programming techniques, but also superscalable algorithms. Here, we conscientiously omit the unknown solutions to handle the PetaBytes of produced data by the simulations of such a machine: will we need two ZettaScale machines in the future, one to run the simulation, the other to post-process its results?

But let's come back to reality: we're a long way from ZettaScale. Most of our problems do not need such power, and computational science and engineering are not about big machines; it's all about different disciplines communicating.

This first issue of Zetta features articles about science and engineering, thanks to computational resources that are described at the end of the magazine. In each issue, we would like to have a look-back section, to remind us that the problems involving Computational Science and Engineering (CSE) have never been trivial.