

Challenging, exciting, promising times

Giampietro Tecchioli expects the techniques of high-performance computing to become more pervasive but worries about shortage of skilled manpower

Computational science is now commonly considered a third mode of science and engineering, complementing and adding to field/experimentation/observation and theory. Computer models and computer simulations have become an important part of the research repertoire, supplementing (and in some cases replacing) experimentation.

Scientific computing spans different disciplines such as computer science, applied mathematics and science/engineering. It is a complex and diverse area, requiring domain expertise, mathematical modelling, numerical analysis, algorithm development, software implementation, program execution, analysis, validation, and visualisation of results.

The fundamental technological enabler of scientific computing is high performance computing (HPC). The linkage between scientific computing and HPC exists at different levels: some scientific and engineering problems need to scale over large computational resources. Others are suited to thousands of simulations running in parallel. A third group deals with large amounts of data.

These three areas share the same need for more computing. The key point is exactly in those two words: 'more computing', because there lies the very essence of the challenges HPC is facing – challenges that will shape the future of HPC and scientific computing.

The first, and very considerable, challenge is exascale, which in concrete terms is the ability to install and run a production supercomputer with a peak performance of at least 1000 PFlop/s and a power consumption limited to 20 MW.

Another challenge is to take the technologies for exascale to much smaller dimensions so that they become available to all users, bringing benefits such as fast



Eurotech's Tigon board shows the company's cooling technology

and energy-efficient computation to many companies and institutions in a variety of markets.

A third challenge is skills. With HPC systems becoming larger and more pervasive and with the application of scientific computing in many additional areas, there is and there will continue to be a lack of specific competences, at least until proper education becomes more widespread and starts producing positive effects.

The future of scientific computing will be very much conditioned by exascale and its derivations. The current largest Top 500 supercomputers are petascale installations that already seem large enough in terms of energy consumption, operational cost, and space occupancy. But, they are not large enough for the challenges science is facing in many fields: from high energy physics to astronomy, climate modelling, chemistry, material science, biology, and others.

Exascale supercomputers can enable progress in 'old' science such as climate modelling, molecular dynamics, aerodynamic design, and cosmology. In other words, there is a part of science and engineering for which it is already known that more computing will lead to faster and better results. But, exascale could also open possibilities of 'new' science and engineering, so new applications that will lead to discoveries that we don't yet imagine.

Take, for example, a project like the

'Human Brain', in which Eurotech is taking part. With current computational resources, we can map and simulate 5-10 per cent of the human brain but not in real time. Extrapolation demonstrates that a real-time complete human brain simulation will need from 1 to 10 exaflop/s and 4 PB of memory.

While theoretical projections suggest the possibility of having an exascale system by 2018, reality tells us that a usable supercomputer of that size will require at least few years into the next decade. Simply adopting the current approach – more of the same, but bigger and faster – will not work, due to constraints in power availability, power cost, reliability, and scalability of applications.

Exascale poses new constraints that will force the entire HPC community to think differently. Power (cost and availability) is the largest of these constraints. In the last decade, we realised a one-time gain in power efficiency by switching to accelerators/manycore. This is not a sustainable trend, without any future technological discontinuity in combination with a new approach. New technologies such as silicon photonics, as well as improvements on existing technologies, such as low power processors, accelerators, liquid cooling and 3D Torus, will be combined with a new way of using systems and new programming models (so, in essence, a new scientific computing) that will not focus only on performance but also on energy efficiency. For instance, since the majority of energy consumed by today's supercomputers is used to move data around the system, a lot of attention will be given to concurrency and locality.

As a manufacturer, Eurotech believes it can contribute to the exascale effort by working with research institutions (as we are doing in projects like QPACE2, DEEP, and Human Brain) and using technology advances in systems that are usable and affordable.

What we envision is that the combination of novel, extremely energy-efficient architectures and liquid cooling should provide the ground on which to build exascale systems. We recently presented a

new HPC architecture, based on X-Gene, the Applied Micro ARM 64 bit CPU, with the support for 4 Nvidia Tesla K40 GPUs. We think that steps like these are important to reach exascale, although we are perfectly aware we will need to take a lot of these steps and at a rather fast pace.

Another aspect of future systems will be liquid cooling. Seven years ago we bet on direct-contact hot-water cooling and we will continue to develop this technology through stages of improvement. We have developed the second generation of Aurora direct liquid cooling – lighter, more compact and effective – allowing extraordinary densities.

It is highly likely that future exascale systems will be heterogeneous, including in one system different computation and storage components, such as processors, accelerators, FPGAs, NVMs, etc The Eurotech Aurora Bricks architecture goes into the direction of modular heterogeneity.

Also, exascale systems will use so many components, that it will be almost impossible for the whole system to operate without faults. Resiliency – the ability to recover from faults – will be paramount. As a manufacturer, Eurotech aims to make systems as reliable as possible and also to provide to applications all information they

need to prevent and manage faults. However, there are limits to what we can attain with hardware reliability alone. In the future, it is likely there will be the adoption of an integrated approach to fault management across the system stack, with a balance of



PEOPLE WHO HAVE SPECIALISED IN HPC ARE CURRENTLY SCARCE, ESPECIALLY IN EUROPE

GIAMPIETRO TECCHIOLLI

proactive prevention and reactive resilience.

The future of scientific computing is not only exascale and extreme scale systems. Exascale technologies will be used in different applications, including mid-range systems and solutions that are not conventional HPC. This will match the expansion of HPC usage driven by a demand for acceleration coming from SMEs and new HPC application areas over and beyond scientific computing. OSINT (open source

intelligence), cybersecurity, computational finance, media and rendering and, real time situational awareness, high performance embedded computing will likely become areas in which HPC techniques and technologies will be more in demand.

People who have specialised in scientific computing and enabling technologies such as HPC are currently scarce, especially in Europe. This will surely trigger a reaction and the growth of training programmes so one can expect a more structured education in HPC in the future, even if it will take a while to fill the current gap. Organisations or countries which are better equipped with those skills will be in a position of competitive advantage.

These have triggered a wealth of attention and investment worldwide, laying the groundwork for technological discontinuities. At the same time, simulation is being adopted by more industries and an increasing amount of data is produced within science and beyond. All of this will determine an extraordinary evolution of scientific computing, at a time that is turbulent, exciting and promising.

Giampietro Tecchiolli is CTO of Eurotech

The best is yet to come

Although advances in HPC have been stellar, in **Mark Seager's** view there is even more still to come

The single most important truth about high-performance computing (HPC) over the next decade is that it will have a more profound societal impact with each passing year. The issues that HPC systems address are among the most important facing humanity: disease research and medical treatment; climate modelling; energy discovery; nutrition; new product design; and national security. In short, the pace of change and of enhancements in HPC performance – and its positive impact on our lives – will only grow.

This phenomenon stems from what can generally be called 'predictive scientific simulation,' which has revolutionised the scientific method itself. Since Galileo invented the telescope 405 years ago and observed the moons of Jupiter, scientists have moved research

forward through theory and experiment. Now the pace of scientific discovery has been radically accelerated as a result of theory and experiment being augmented by predictive scientific simulations using parallel supercomputers.

Scientific simulations inform theory and experiment in three ways. The 'hero' simulations, typically use the entire system on a single record-setting simulation, and produce higher-fidelity results. The 'ensemble' simulations are groups of simulations run in throughput mode, typically utilising major fractions of the system, and which provide sensitivity information of the 'hero' results on inputs and model parameters that allow one to estimate confidence in the 'hero' results (e.g. error bars). Lastly, the rapid analysis of large experimental data sets increases the usefulness

of these results. These procedures enable significant savings through informed decisions and actions in the real world.

The effect of scientific simulations on the scientific method has been enabled by the astonishing increase in the computational power of supercomputer systems. Over the past 20 years – which is a heartbeat compared to long span of the scientific method (and even shorter when compared to the evolution of biological systems or to the timescales of geological change) – supercomputers have been transformed in many ways.

Here's a quick analysis of the Top500 list of supercomputers, comparing the Thinking Machine's CM-5/1024 of 1993 to China's Tianhe-2 of 2013:

● **Performance:** Systems have exploded