e science TALK

e-ScienceBriefings

Talking about e-science

Supercomputing: Empowering research

Computer simulations have evolved as an essential part of scientific research, complementing theory and experiment. Scientists and engineers use simulations when problems are complex, or experiments are too dangerous, expensive, or impossible to carry out. Today supercomputing has simulated rat brains, tested engineering structures and modelled global warming. Investment in these technologies can lead to real economic benefits while ensuring researchers remain internationally competitive.

What is a supercomputer?

Unlike distributed grids, where thousands of networked computers are linked together across many geographical locations, a supercomputer consists of thousands of integrated computers in a single room. These supercomputers are at the frontline of current processing capability, particularly speed of calculation.

A supercomputer's power is measured by how many floating point operations it can complete per second (flop/s). The world's fastest machines are listed in the biannual Top500 list. In November 2010, the Chinese Tianhe-1A system at the National Supercomputer Center in Tianjin topped the list, with a record peak performance of 2.57 petaflop/s - the equivalent of everybody on the planet doing 367,000 calculations a second.

Supercomputers are used for highly complex problems in which many intermediary results are dependent upon one another. Therefore supercomputers are often used for highly calculation-intensive tasks such as those found in astrophysics, climate modelling and biomedicine. The use of supercomputers is commonly referred to as 'high performance computing' or HPC.

European percentage of worldwide HPC revenue

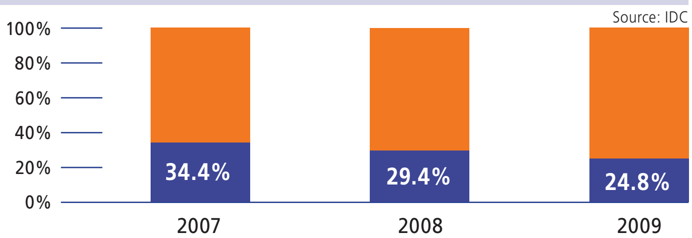

A study in the US has shown that that investment in HPC increases research competitiveness at academic institutions. In Europe there is a wealth of HPC-related experience and talent. However a 2010 IDC report for the European Commission found that recently Europe has fallen behind other regions. In order to fully benefit from HPC the report recommends that we need to both increase its HPC investments and find ways to apply HPC in a more productive and innovative manner.

Unravelling the mystery of Dark Energy

Friedrich Röpke's research group at the Max-Planck-Gesellschaft, Garching, Germany is using PRACE's JUGENE Tier-0 supercomputer to study type Ia supernovae explosions. They hope their work will contribute to a better understanding of mysterious dark energy in the universe.

Type Ia supernovae are exploding stars that are among the brightest objects in the universe. They are an important phenomenon, as scientists can use them to infer cosmic distances, by looking at their apparent brightness.

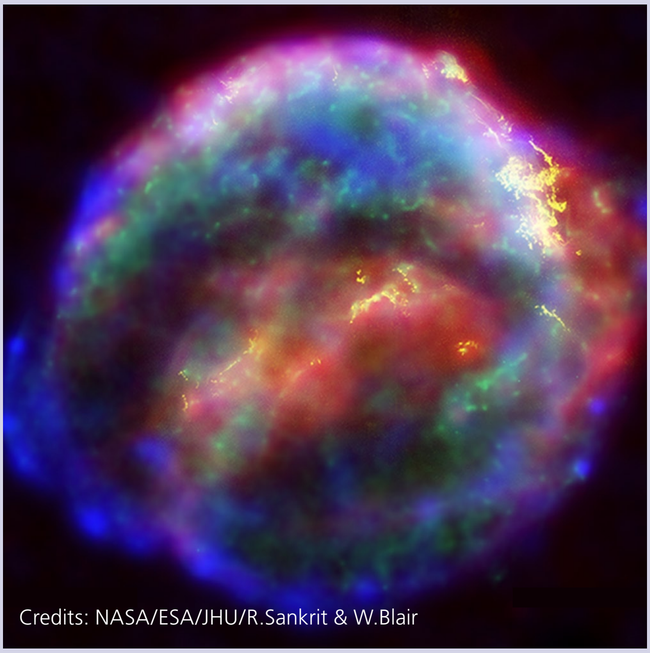

Credits: NASA/ESA/JHU/R.Sankrit & W.Blair

Distance determinations derived using type Ia supernovae show that the universe is expanding at an accelerated rate. However the reason for this effect is still unclear. Researchers attribute it to mysterious 'Dark Energy' which they believe forms the main constituent of the universe today.

Type Ia supernovae distance determination can contribute to a better understanding of this mysterious energy form. The necessary distance measurements, however, require high precision and great observational effort. Therefore we need to match this with a sound theoretical understanding of the supernova explosions.

Röpke's group aims to shed light on the physics of the explosion mechanism. But the explosion physics are complex and can only be studied in sophisticated numerical simulations which require the huge amounts of computational power provided by PRACE.

Thomas Eickermann, PRACE - *"European researchers and engineers need access to world-class supercomputer resources and services to remain internationally competitive. Therefore, Europe needs to continuously invest in this area to provide a leading-edge HPC infrastructure. Strategically, Europe also needs to strive for independent access to supercomputing, as it already has for aviation and space."*

Grids of supercomputers

In Europe, PRACE (Partnership for Advanced Computing in Europe) has created a high-end supercomputing infrastructure. Together with grid facilities, data resources, software and experimental facilities, the PRACE Research Infrastructure will meet the needs of scientific and industrial domains in Europe. Currently, PRACE has two world-class supercomputer (Tier-0) systems: JUGENE, at FZJ, Jülich, Germany and CURIE at CEA, Bruyères-le-Châtel, France. A third Tier-0 system, SuperMUC at Leibniz Supercomputing Centre, Munich, Germany is underway.

PRACE collaborates closely with DEISA, the Distributed European Infrastructure for Supercomputing Applications, which currently provides access to eleven European national supercomputing centres (Tier-1 systems). User access and coordinated operation of PRACE centres employs software and services that have been pioneered and deployed by the DEISA project. When DEISA ends in 2011, PRACE plans to integrate current DEISA services into its own infrastructure. PRACE also has close links with EGI, the European Grid Infrastructure, to ensure grid access and interoperability today and in the future.

Simulating blood flow

Federico Toschi and his team at the Eindhoven University of Technology are using PRACE supercomputers to learn more about blood flow in the body.

Studying this phenomenon is not easy as it requires blood flow simulations which have to take into account interactions between large numbers of cells in a variety of geometries. Previous models simulate blood as either a homogeneous fluid or as a relatively small number of red blood cells, but the latter are not able to reach relevant time or length scales.

The team in the Netherlands has developed a model which allows them to perform simulations of millions of blood cells over more realistic times and distances. They have been awarded 39 million compute hours on PRACE resources to run their model, and ultimately learn more about fatal and serious blood clots.

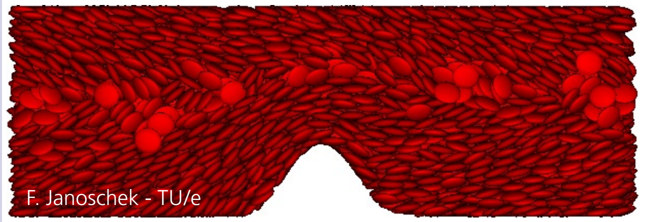

F. Janoschek - TU/e

In the US, The National Science Foundation's TeraGrid integrates high-performance computers, data resources and tools, with high-end experimental facilities around the country. Currently, TeraGrid resources include more than a petaflop of computing capability and more than 30 petabytes of online and archival data storage, with rapid access and retrieval over high-performance networks. Researchers can also access more than 100 discipline-specific databases.

John Towns, TeraGrid Forum Chair - *"The next phase of the US National Science Foundation's (NSF) investment in high-end computing will be launched in 2011: the eXtreme Digital, or XD, program which will succeed TeraGrid. Many of the existing resources will continue, and new ones will be added."*

Supercomputing challenges

Initiatives such as PRACE, DEISA and TeraGrid help to expand scientific knowledge and in turn deliver social and economic benefits. However, the use of supercomputers raises a number of issues:

- **Cost:** Supercomputers are expensive. The Chinese Tianhe-1A system for example is estimated to have cost 600 million yuan (68.7 million euro). But in such a fast moving field, these machines can rapidly become out-of-date. The number one computer in the Top500 list is rarely number one the following year. However such investment is necessary if we are to benefit from HPC science and industry research.
- **Access:** Supercomputers are powerful scientific instruments that can be used by researchers across the world. In the US, for example, supercomputers at the National Center for Supercomputing Applications (NCSA) are used remotely by more than 1,500 scientists and engineers. However, in practice, research time on these machines is limited. Supercomputing providers need to find fair ways to allocate time on their machines. PRACE, invites researchers to submit proposals which are subject to both technical and scientific peer review. TeraGrid also allocates resources via a quarterly review process, with successful proposals receiving allocations of high-performance computing, scientific visualisation, data storage, and advanced user support for 12 months at a time.

Elizabeth Cherry, Rochester Institute of Technology - *"Myself and colleagues from Cornell University and the Max Planck Institutes in Germany used TeraGrid resources to develop an effective and potentially less painful defibrillation method for treatment of atrial fibrillation, a heart disorder that affects millions of people. Cardiac modelling demands supercomputing because it requires a great range of scales - both spatially, from a single cell to the full-size heart, and over time, from microseconds to minutes. Our supercomputer simulations have shown the feasibility of our low-energy approach."*

- **Efficiency of use:** The TOP500 lists sites according to the Linpack benchmark - a 'horsepower number' which correlates to the machine's ability to perform calculations. However, harnessing that power to solve real world problems efficiently is still a major engineering challenge, requiring a change in algorithmic and programme design to fully utilise future multi- and many-core processors. PRACE experts continuously work on scientific software packages to make best use of these new computing engines as they are deployed in Europe.

Daniel Ahlin, PDC, KTH - *"I expect HPC power efficiency to be an extremely important field in the next few years since the environmental and power cost problems it tries to mitigate are critical issues, not just for HPC, but for society as a whole. A good, easy-to-reach and high-yield example is heat-reuse which has just begun to be explored more widely."*

- **Energy consumption:** As supercomputers become more powerful their energy consumption becomes a real issue. Energy costs make up a substantial proportion of supercomputing costs – cooling processors is a particular problem. PRACE is engaged with several green IT prototypes, including its new Tier-0 system, SuperMUC. SuperMUC will use water to cool its processors and memory, resulting in a 40% energy reduction compared to an air-cooled machine. The warm water can also be redirected to the heating systems of nearby buildings. As the trend towards more powerful computers continues the Green 500 list, which ranks computers according to their energy efficiency, is expected to become more important than the Top500 list.
- **Training:** In order to achieve the best results from these instruments, researchers need to be trained to use supercomputers. Most supercomputing providers offer this, enabling researchers to make the most out of these machines. TeraGrid offers an abundance of HPC training opportunities and PRACE has also developed an extensive training programme, including seasonal schools throughout Europe.

Pekka Manninen, PRACE training coordinator CSC - *"Using Tier-0 systems efficiently is highly non-trivial; in order to benefit from the competitive edge inherent in the infrastructure, our users must be skilled enough to use it properly. The PRACE training scheme builds a permanent network for training in the field of Tier-0 computational science. In the First Implementation Project (PRACE-1IP), the most visible actions are a top quality face-to-face training event curriculum and the establishment of an online training portal together with various education outreach activities."*

Reaching the exascale

These are not the only challenges. Petascale technologies have allowed us to simulate protein folding and model climate data. However the complexity of these scientific models is such that to truly understand them we require exascale computing – a thousand times more powerful than the current top of the line petascale computers.

Rosa M. Badia, Barcelona Supercomputing Center - *"Supercomputers are essential tools for today's research in physics, weather forecasting, molecular modeling and many other disciplines. However, the complexity of programming such computers prevents efficient use of them. Our research in programming models enable other scientists to boost their computations in MareNostrum (the supercomputer at the Barcelona Supercomputing Center) and promote the advancement of their research."*

The most beautiful supercomputer in the world?

When the Spanish and Catalan government needed a location for one of the fastest computers in Europe as well as an innovative research center, they came up with a novel solution.

Tasked with building a new supercomputer in just four months, the newly created Barcelona Supercomputing Center had a tough challenge on their hands. Surprisingly the best place was a chapel, Torre Girona, located at the North Campus of the university in Barcelona, which had enough interior space to fit the new computer, named MareNostrum.

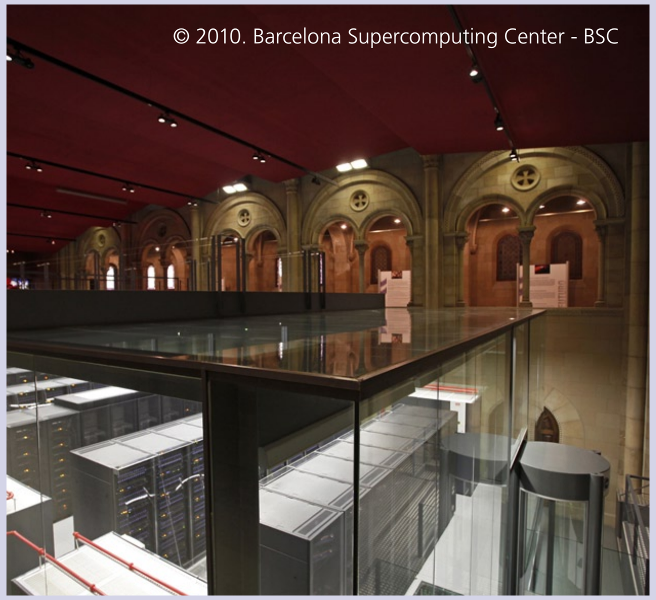

© 2010. Barcelona Supercomputing Center - BSC

Today, after an increase in capacity in November 2006, MareNostrum is the 118th most powerful computer in the world, as named in November 2010's Top500 list. It aids research in areas such as computer sciences, life sciences, earth sciences and computer applications in science and engineering.

Moving towards exascale computing will not only exacerbate issues of cost, access and training; it will raise entirely new problems. Whereas in the past there was scope to scale up, we are now reaching fundamental physical limits of technology. Current machines cannot simply be made bigger to achieve exascale processing – entirely new technologies are required. Exascale computers will also need an energy capacity which cannot be delivered by any computer centre today.

Engineers already acknowledge that exascale computers will have to rely on parallelism, along with compatible algorithms and software. Because memory bandwidth and capacity do not follow Moore's Law, developers will have to become better at managing the constraints of less memory per core and investigate novel communication-avoiding algorithms.

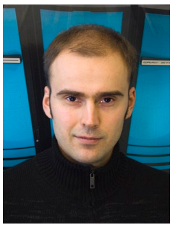

Sebastian von Alfthan, CSC – IT Center for Science - *"Exascale computing brings major challenges that need to be met on the hardware, software and programming model level; but also opportunities for groundbreaking research at an unprecedented scale. Addressing these challenges provides Europe a chance to take a greater role in shaping future HPC technologies. "*

An expected 100 million cores will be needed to deliver an exaflop but at such large scale, frequent hardware failures will be inevitable. Application developers must embrace this and factor it into algorithms and software. At the same time physically interconnecting 100 million cores will not be trivial. Network reliability, intrinsic latencies and message size will all need to be considered when designing the next generation of supercomputers. And even when solved, it is highly unlikely that one architecture will be suitable for all applications.

With so many hurdles in the way, exascale computing will need dedicated time, effort and money. Experts believe that we could see the first exascale systems by 2018 but warn that they will require an estimated $1 billion of investment. However this investment could lead to real economic benefits.

Glossary

Parellelism – a type of computing which carries out many small tasks concurrently.
Memory bandwidth – the rate at which data can be read to or from memory by a CPU.
Moore's Law – states that the number of transistors on a chip will double about every two years.
Core – Part of a processor, which reads and executes instructions independently of other cores on the same chip.
Latency - the delay experienced in transmitting signals from one part of a computer to another.

Scan this QR code into your smart phone for more on this e-ScienceBriefing

Cleaning up the spill

When the blowout of the BP Macondo well destroyed the Deepwater Horizon oil rig in the Gulf of Mexico on April 20, 2010, it sparked a massive environmental disaster. TeraGrid resources were ready to assist in the mitigation process.

Clint Dawson, head of the Computational Hydraulics Group at The University of Texas at Austin's Institute for Computational Engineering and Sciences (ICES) received an emergency three-million-hour allocation from the NSF to computationally model the spill. Using the Ranger supercomputer at the Texas Advanced Computing Center, Dawson's group created high-resolution forecasts showing how the spill might migrate over a 72-hour period and affect large parts of the coastline. These models helped response teams position booms and prepare for the possibility of a natural disaster – a hurricane - which could compound one of the nation's worst man-made disasters.

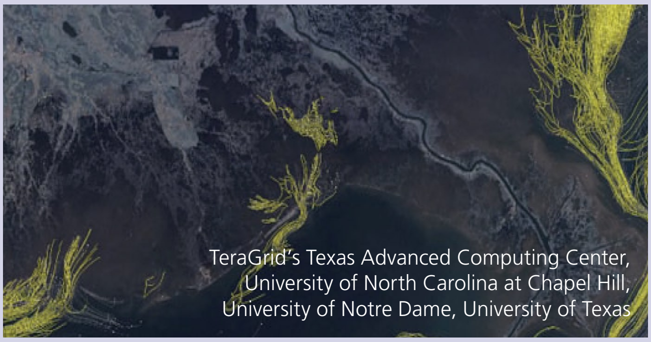

TeraGrid's Texas Advanced Computing Center, University of North Carolina at Chapel Hill, University of Notre Dame, University of Texas

Separately, the Louisiana Optical Network Initiative, another TeraGrid partner, provided the National Oceanographic and Atmospheric Administration's surveillance team with critical high-bandwidth network connections to quickly transfer large amounts of vital data between the spill site and Washington D.C.

For more information:

Barcelona Supercomputing Center: www.bsc.es

DEISA: www.deisa.eu

HPC in Europe (with link to IDC report): www.hpcuserforum.com/EU

PRACE: www.prace-ri.eu

Study into HPC and US research competitiveness: www.jiti.net/v10/jiti.v10n2.087-098.pdf

TeraGrid: www.teragrid.org

The Green500: www.green500.org

TOP500: www.top500.org

EGI: www.egi.eu

iSGTW: www.isgtw.org

e-ScienceTalk : www.e-sciencetalk.org

e-ScienceTalk is co-funded by the EC under FP7