La Recherche
June 2009

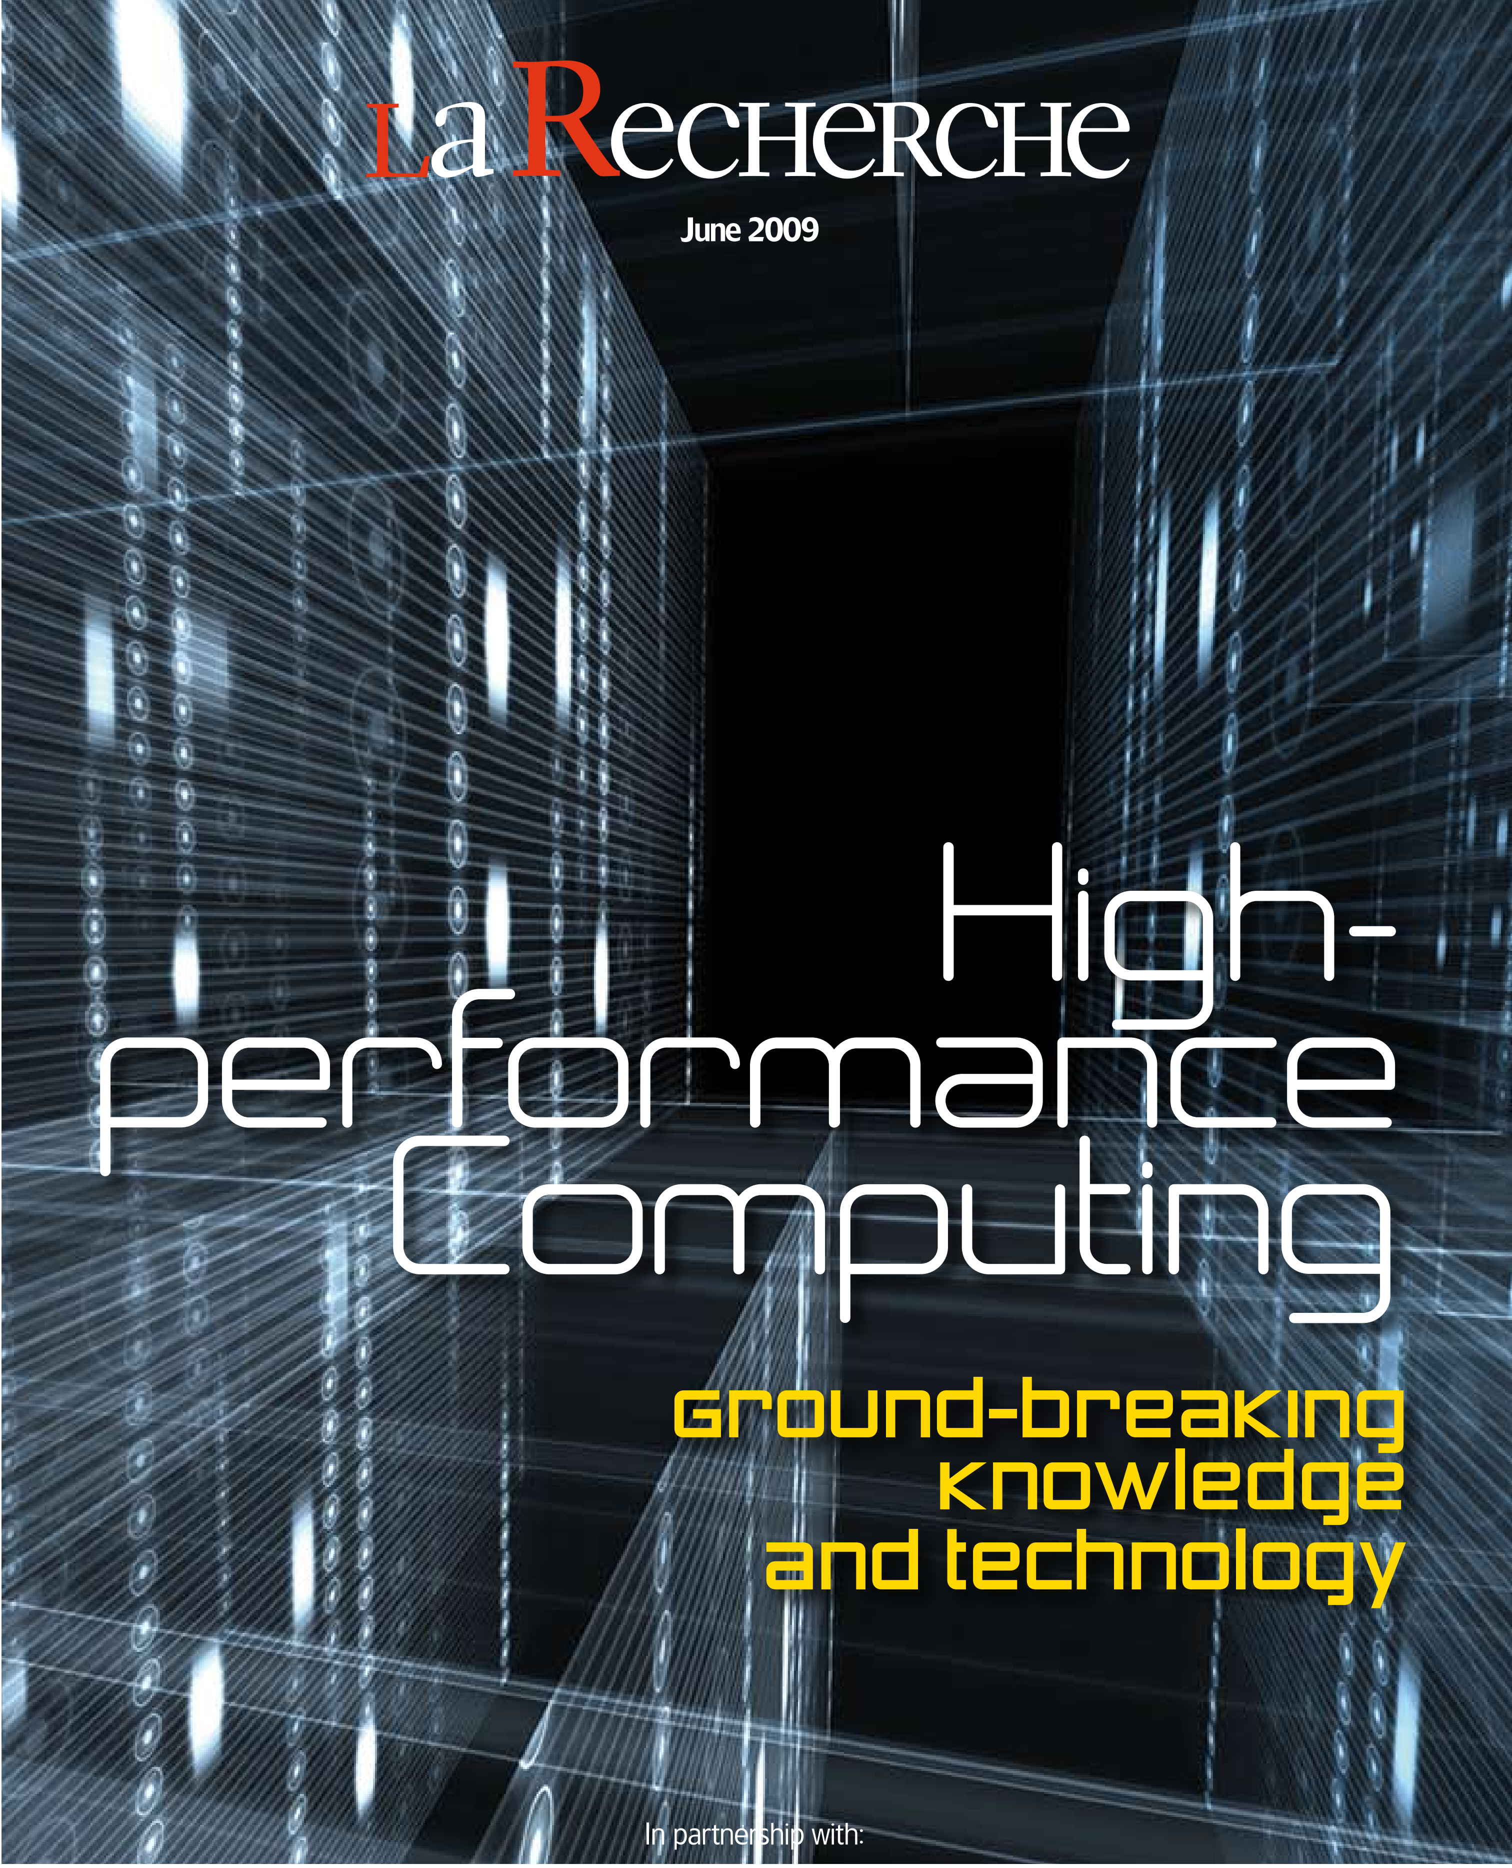

High-performance Computing

ground-breaking knowledge and technology

In partnership with:

TOTAL, cea, PRACE, GENCI, INRIA, Bull (Architect of an Open World™), intel

HIGH-PERFORMANCE computing is going to save research

Today's researchers are 'super-specialists', far removed from the philosophers and polymaths of the Age of the Enlightenment! This acute level of specialisation has contributed to the highly technical nature of research, the staggering increase in the number of researchers since 1950 and the exponential rise in the number of publications.

The same goes for the research industry: it too has experienced exceptional growth, to the point where it is now highly strategic for a good number of sectors.

You'd think this was a sign of good health; however this extreme specialisation has brought about a situation where exchange of information between researchers is problematic when they don't belong to the same sub-sub-sub-discipline. Research is now progressing at two different paces. Steadily, with gradual developments, and rapidly, with technological leaps, often carrying over from one field to another. You can't reinvent the light bulb by doing research on candles.

High-performance computer simulation – this brand new methodology which involves reproducing complex three-dimensional, time-dependent phenomena using computers – is a saving grace for our community. Not only does it produce results which would have seemed totally inaccessible thirty years back (some of which are discussed in this special issue) but it has managed to get specialists from various different fields around the same table, reconciling industrial research and university research. A kind of cross-fertilisation is at last taking place between high-performance computing – which hinges on the most recently available ground-breaking fundamental research – and industrial research which is pushing requirements to the extreme limits. Simulation calls for defining, by means of equation (so-called 'modelling'), the phenomenon under investigation in all its complexity and therefore across all disciplines.

Modelling a patient's radiotherapy treatment, for example, calls upon nearly all scientific disciplines. Over and above the benefit to patients, it also supports research itself.

Numerical simulation executed on CEA[(1)]'s Tera 10 supercomputer. Seismic wave propagation.

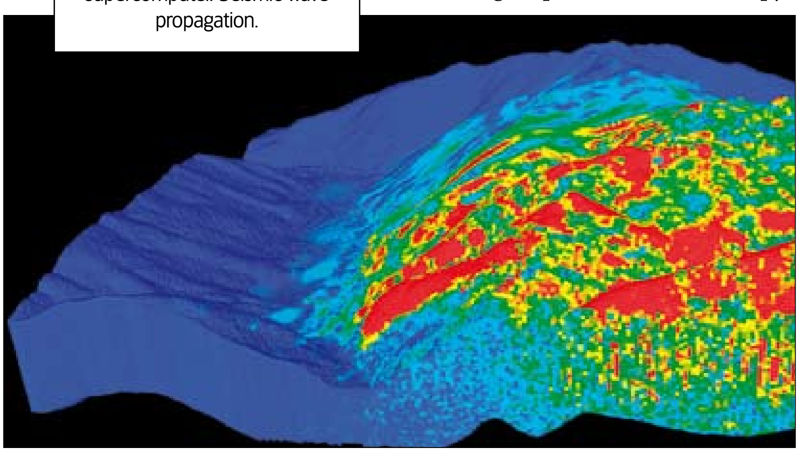

© CEA

By **Jean-Michel Ghidaglia,**
Professor at the Cachan Ecole Normale Supérieure, Scientific Director of La Recherche

(1) CEA - French Atomic Energy Commission

La Recherche

Supplement 2 of La Recherche cannot be sold separately from supplement 1 (LR No.432).
This 'La Recherche' magazine supplement has been produced in association with Bull, CEA, Genci, Intel, Inria, Prace and Total.
Sophia Publications
74, avenue du Maine - 75014 Paris - FRANCE
Tel.: +33 (0)1 44 10 10 10
Editorial office email: courrier@larecherche.fr
To contact a member of the editorial team by phone directly, dial +33 (0)1 44 10, followed by the four digits next to the name.
SUBSCRIPTIONS / BACK ISSUES / BINDERS
Email: abo.recherche@groupe-gli.com
La Recherche subscription department
22, rue René Boulanger, 75 472 Paris cedex 10 - France
Price in France: 1 year, 11 issues, €59.00
1 year, 11 issues + 4 supplements, €79.50
International pricing: please contact us.

Switzerland: Edigroup, (0041) 22 860 84 01
Belgium: Edigroup (0032) 70 233 304
Canada: Express Mag, 8155, rue Larrey, Anjou Québec HJ2L5
Editor-in-Chief: Aline Richard
Scientific Director: Jean-Michel Ghidaglia
Deputy Editor-in-Chief for supplement 2: Isabelle Bellin
Managing Editor for supplement 2: Sylvie Richardin (!Infokom)
Art Director for supplement 2: Pascal Brachet (!Infokom)
Layout Artist for supplement 2: Guylaine Vautier (!Infokom)
Production: Christophe Perrusson (13 78)
Partnership Project Manager: Stéphanie Jullien (54 55)
Sales assistant: Antoine Faure (54 53)

Administration and finance director: Dounia Ammor
Direct marketing and subscriptions Director: Virginie Marliac (54 49)
Marketing Manager: Estelle Castillo (54 51)
Distribution (distributors/newsagents): Evelyne Miont (13 80)
Management Coordinator: Isabelle Parez (13 60)
Accounts: Marie-Françoise Chotard (13 43)
Webmaster: Jean-Brice Ouvrier (54 52)
ADVERTISING: Le Point Communication
Head of advertising: Marie Amiel (12 57)
Sales and cultural sector assistant: Françoise Hullot (f.hullot@interdeco.fr)

La Recherche is published by Sophia Publications, a Financière Tallandier subsidiary.
Chairman and Managing Director: Philippe Clerget

Supplement 2 of La Recherche
Joint commission:
0909 K 85863 ISSN 0029-5671
Printed in Italy by Imprimerie Canale, Via Liguria 24, 10071 Borgaro, Torino. Copyright deposit

Contents

High-performance computing

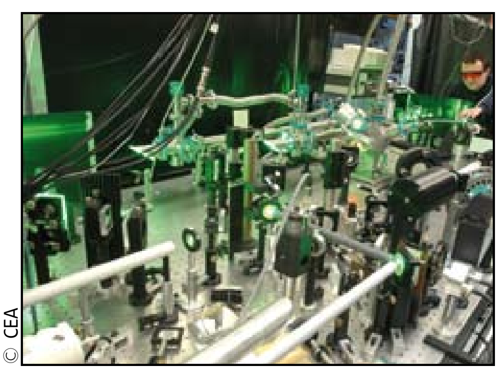
© CEA

TS

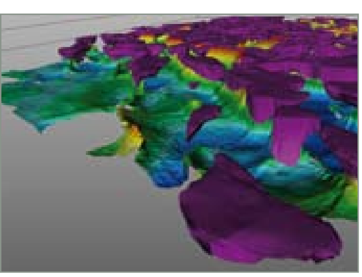
© TOTAL

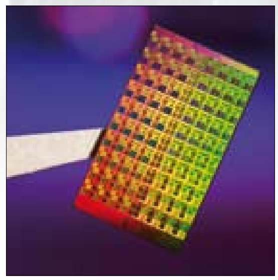
© Intel


© Intel

Supercomputing technology for E

In this interview, Bernard Bigot, chairman of the French Atomic Energy Commission (CEA), a major player in France for research, development and innovation, and Achim Bachem, chairman of

© CEA

Bernard Bigot, General administrator of the French Atomic Energy Commission (CEA)

Teraflops, petaflops, exaflops(2)… The capability of supercomputers is growing incessantly. High-performance computing (HPC) looks like an endless worldwide challenge. Why do we need so much simulation? Are we forced to go further?

Achim Bachem: It is not just a race for capability, it is a chance to solve the most difficult and complex problems of our society: climate change, energy, materials science, biology, drug design... Besides theory and experiments, simulation is the third pillar of science. Thanks to the current and the next generations of supercomputers, we are at the beginning of a new era. Simulation is often the only solution in numerous areas where physical experiments are impossible, such as climate change or astronomy. In other areas like aviation crashes or combustion tests, it is very difficult to do experiments. Simulation gives a realistic answer to complex problems. But complexity is always increasing. Think about climate change: even leading edge supercomputers have to run for years to forecast important parameters such as the possible increase in the earth's temperature in the next 100 years. But we need an answer right now. Only next generation machines, 100 to 1,000 times faster than the current ones, can help.

Bernard Bigot: Indeed, simulation can give us an accurate representation. It is the result of a long process, taking advantage of the combination of a very deep understanding of fundamental physical phenomena – from the nanometric to the universe-wide scale – and mathematical algorithms that allow these physical models to run on computers. As Professor Bachem says, for some issues, experiments are impossible. I would add that for other fields such as drug design, materials science, chemistry… experiments are possible but there are too many possible combinations of conditions. And we already know that the complexity of the phenomena we are dealing with will require more and more intensive computing.

> "With the present and the next generations of supercomputers, we are at the beginning of a new era"

To what extent is industry involved?

B.B.: Many areas of industry are involved to a great extent. Let's consider nuclear energy: building a new nuclear power plant requires complex simulation of heat transfer, of neutron flux, of the behavior of materials. For wind energy, simulation is fundamental to manage the behavior of a network of thousands of wind turbines properly. Issues of safety, economic performance, the cost of raw materials... mean that more and more simulation is needed for industry.

A.B. : Another example: in the aerospace industry, our European company EADS, needs leading edge technologies to stay competitive. Their latest aircraft, the A380, would never have been built in such a short time without the HPC expertise that exists in Europe. There is no sense in carrying out such developments on a national level.

Do you mean that HPC can only be achieved on at the European level?

A.B.: Yes. Above all, for European based industries such as aeronautics and aerospace, but also for other industries. We have to leverage our own key technologies, bundle our competencies, gather the knowledge in

(1) Partnership for Advanced Computing in Europe.

(2) A teraflops, a petaflops and an exaflops correspond respectively to processing capacities of 10^{12}, 10^{15} and 10^{18} operations per second.

is a key urope

the Forschunsgzentrum Jülich and coordinator of the PRACE(1) project, reveal their shared view of high-performance computing. They map out the best future of European policy.

Interview by
Isabelle Bellin

© Forschungszentrum Jülich

Achim Bachem, Chairman of Forschungszentrum Jülich and coordinator of the PRACE project.

Europe and give our manufacturers access to this computing technology. These machines are too expensive for a single country.

B.B.: I guess there are three levels of capabilities. The first one is at the European level: we need a few "world-class machines" with the most advanced technology, which our big companies and researchers can test at prototype stage and utilise to their full extent as soon as they are launched on the market . Building such machines requires close collaboration between the hardware and software research community and industry. At a lower level, numerous, but nonetheless powerful machines, based on proven technologies, should be run on a national scale. There is no real added value in managing them at the European level. At the third level: an individual body or a company only needs standard production machines, that are less expensive and with less capability.

How should these "world-class machines" be designed?

B.B.: Using the kinds of methods being developed through the PRACE program, which aims to create a few huge Petaflops infrastructures with highly concentrated expertise (see article on PRACE, p. 56). The whole European organization is currently debating how we could cooperate to build these world-class machines, so we have the same capabilities as the US, Japan and soon China.

A.B.: It is strategic for industry: these supercomputers funded by PRACE – like the prototype we are jointly proposing with the CEA, provided by the French server manufacturer Bull with the new Intel chip (Xeon Nehalem) – give you a big technology advantage against standard machines you can buy off the shelf: it's a huge worldwide advantage.

Europe must have independent access to the key HPC technology, in order to compete. Last year, the CEA and Jülich signed a cooperation contract to jointly develop fundamental IT-technologies used for HPC. It is very important that the big national labs in Europe, such as ours, come closer together and show the strength of Europe. But besides technology, we also have to foster education in the science of computer simulation.

"There are several cards to play"

How is the shared Franco-German vision of HPC manifesting itself when it comes to simulation?

A.B. : Fusion science is certainly one of the best examples. The CEA and FZJ are working very closely together on the technologies for the first European fusion computer, currently being installed at Jülich. The good news is that this machine is being jointly developed in Europe with the French company Bull, Intel, the German company ParTec, and the Israeli firm Melanox. Proof positive that we have the necessary skills.

B.B.: Jülich has the capability to host this new very large computer. We are sharing our own experience, on our Bull Tera 10 computer at the CCRT(3) (Bruyères-le-Châtel), with Jülich's project leaders. Following in-depth discussions they made their choice for the best capability, capacity and cost ratio, on their own. Besides fusion, we also cooperate in many areas: imaging technologies, nanoscience, new energy technologies like fuel cells, energy storage, materials science, quantum technology, biophysics, earth and atmospheric sciences… But no

(3) **Centre for Research and Technology Computing.**

matter what area you are researching, it is very important to have leading-edge computers. The researchers themselves define the most appropriate tools for each application. With PRACE's infrastructure, they will easily work together.

From your point of view, what is the best architecture?

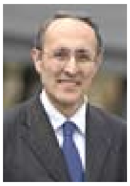

B.B. : No single technology is right for all applications. Some need a huge amount of memory, others, very high data transmission speed, or lots of computing power. A variety of strategies are needed. We know that parallelism can be extremely valuable, but there are many ways to organize it. Europe could certainly contribute, by promoting new combinations of the basic components to provide machines that are better suited to specific problems. In particular, we could influence these choices to reflect our intense concerns about saving energy.

A.B. : Reducing energy consumption is the biggest technological issue that needs to be resolved, otherwise every future Exaflops-scale machine will need its own power station! We also have to develop new algorithms and new software for massively parallel computers, for all different disciplines: a huge challenge. That's why within PRACE, we have decided to offer different technologies in different locations in Europe for different applications.

Last year, President Sarkozy suggested to Chancellor Merkel that France and Germany should join forces in supercomputer development. Is the discussion progressing?

A.B.: We are all convinced it is very important for Europe to foster HPC capabilities. It is being discussed intensively in Germany and France. We have to build up a joint taskforce to promote our activities in software and hardware. We already have Bull in France and SAP in Germany, one of the top three sellers of business software solutions worldwide. Also smaller countries like Finland, Sweden and some new European member states are committed to high-end computing. We also want to convince worldwide vendors to build research and development labs in Europe.

B.B. : For such a strategic technology, we must not trust these developments into the hands of others and buy these machines from outside Europe as we do with any other computer. We need several companies in Europe able to provide all the expertise and competence in hardware and software. We must convince companies to produce and invest in Europe in the conception, the production, and the maintenance of these types of systems. The French President asked me to look at how we could implement this successfully. I presented my report a few months ago and I am sure we will soon have an answer. Germany and France can put forward initiatives, but we must involve other countries more widely, including Spain, the UK, the Netherlands, Italy… We are working on this, and I am confident that these discussions will bear fruit. It is a major strategic decision, that has to be analyzed in depth to find the best conditions.

The USA has given huge financial support to its HPC industry since the 1990s. Japan does the same. What about Europe? Shouldn't we support more strongly our industry?

A.B. : We have European and national financial support in many countries, in France, in Germany, and through PRACE to help industry build these installations and develop the next generations of supercomputers. We made very good progress in the last two to three years. The race goes on, Europe has to catch up. I am sure we can be one of the big world players in supercomputers and simulation science in the future.

B.B. : But our governments must be more involved in HPC, as they are in the US and Japan. The industry cannot finance the development of these leading technologies on their own. As the US does, we have to support public sector research institutions so that they can cooperate with the industry and then have the means to buy equipment which is still not on the market. Maybe, we will need some adjustments but this will not break the commercial rules. I think that's essential, given the strategic role of these technologies. I am sure Europe could do as well as the US with a little bit less money. The aim is not to have as much money as the others: it is to have as much money as we need. We expect the next generation of petaflops-scale computers to be available by 2012. We are preparing it, but it makes no sense to have just a single shot. We must continue the process, to take advantage of our investment, in order to be able to be ready for the following generation and on the ones after that.

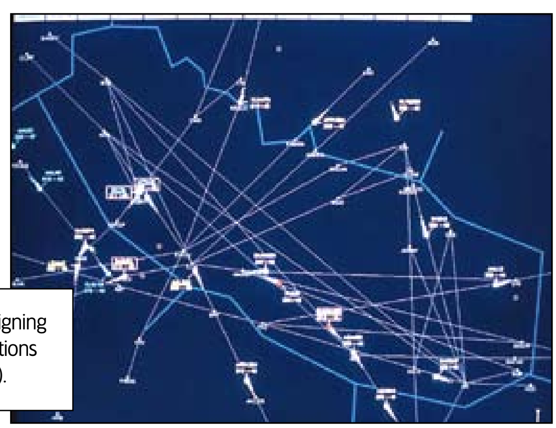

Air traffic simulation for designing future air traffic control stations (Lamih – Valenciennes).

Europe must finance ´public-private´ sector partnerships

Interview with Jean Gonnord, head of the digital simulation and computing project at the CEA.

Interview by
Isabelle Bellin

© DR/CEA

Jean Gonnord, head of the numerical simulation and computing project at the CEA.

You say that 2005 marked the recognition in France and the rest of Europe of the importance of high-performance computing, outside the defence sector. What's been happening in the field since then?

Jean Gonnord: Massive headway has been made in France and the rest of Europe, which has allowed us to bring our high-performance computing resources up to a 'regular' level, compatible with our industrial and research capacities. Europe is now in the running, accounting for 37% of the world's available computing power according to the latest list of the world's Top 500 largest computers (November 2008). The CEA has been heavily involved in this: with the Tera 10 supercomputer which held 5th place in the world ranking in 2006, then with the supercomputer installed in 2007 at the CCRT[1] in Bruyères-le-Châtel. More recently, a whole new ball game has started with the implementation of Genci[2] in France – coordinating supercomputer investments – and the PRACE[3] project in Europe which has been working on the installation of three to five world-class computers scheduled for 2010 (see article p. 56). The political world is now according high-power computing its rightful place. This is excellent news. But it's not enough.

What else would you like to see?

At this point in time, nobody would question the benefits and strategic nature of high-power computing. For the defence sector and nuclear dissuasion, of course. For scientific research, too, where simulation technology has become a crucial component. But, and this is all too often forgotten, high-power computing is also strategic for the competitiveness of our companies and generally speaking for society as a whole. It is the only way to reduce the design-production cycle, whether this be the design of an industrial product or a pharmaceutical drug. In fact, power is directly related to the speed of development, that's the crux of the matter. Large computers are arriving on the market two or three years after the first prototype and those who can get their hands on them as soon as they are designed will be ahead of the game.

Are you suggesting that, to stay in the running, we need to develop computers that match the needs of industries in Europe more closely?

Absolutely. And more and more of us are convinced of this in Europe. Without delay, we must set up a sustainable European model to fund R&D work for high-power computing to support our own companies in this area and to share the benefits of the latest computer technology with our industries. This is no more, and no less, than what the Americans, the Japanese and now the Chinese are doing. The American state, for example, provides IBM, Cray and other companies with funding, paying 300 million dollars for R&D each year, and guaranteeing the purchase of their prototypes for over a billion dollars per year.

How would you like to see this funding set up in Europe?

It would be enough for Europe to put 50 million euros funding per year into R&D and to invest 100 million euros per year in a prototype machine. To do this, we suggest setting up public-private sector partnerships on a 50/50 basis. This is what we did with Bull for the R&D for our future computer, Tera 100 (see article on page 53). We set up a shared laboratory with over 60 people, on the Ter@tec site near the CEA in Bruyères-le-Châtel, the biggest European HPC science park. In a year's time, it should bring together 200 people working on the architecture of future machines. Following this same model, with Genci and the University of Versailles-Saint-Quentin, we have signed an agreement with Intel to set up a shared laboratory for 60 engineers on the same site, by the end of the year. Called Ex@tec, its task will be to design software adapted to hardware architectures enabling Exaflops[4]-scale processing. This demonstrates Intel's acknowledgement, as the world's largest manufacturer of processors, of the expertise of French research teams in the HPC field. By 2010, we hope to bring together over 1,000 engineers from both the public and private sector working at Ter@tec.

(1) Centre for Research and Technology Computing.

(2) Major National Infrastructure for High-Performance Computing.

(3) Partnership for Advanced Computing in Europe.

(4) One exaflops is equivalent to a processing capacity of 10^{18} operations per second.

Europe is back in the run

The never-ending race

Twice a year, in June and November, the world's top 500 most powerful computers are awarded their rankings, based on the execution of a benchmark programme, named Linpack, by researchers from the University of Mannheim (Germany), the Lawrence Berkeley National Laboratory and the University of Tennessee (United States). At the 2008 awards, two machines broke the petaflops record (10^{15} operations per second): Roadrunner (IBM) and Jaguar (Cray). Like five other machines selected for the top 10, they are installed in National laboratories in America, supported by the Department of Energy (DOE).

The yellow 'curve' opposite shows the development of the world's most powerful machines at the time of their installation. Some stayed number 1 in the charts for several months, following in the footsteps of Earth Simulator and Blue Gene L. The United States are always the clear forerunners in the rankings. The blue 'curve' shows the growth of processing power installed in Europe.

Some European machines rank in the Top 10 to 15 at the time of their installation. Even more powerful machines will be released by 2010; we could be reaching exaflops by 2020, if IBM, Cray or Bull are to be believed.

Computing power installed by country

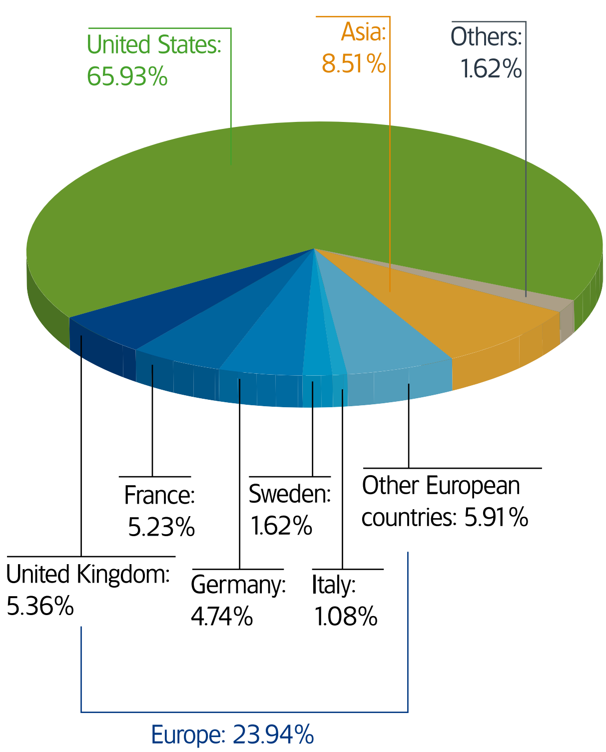

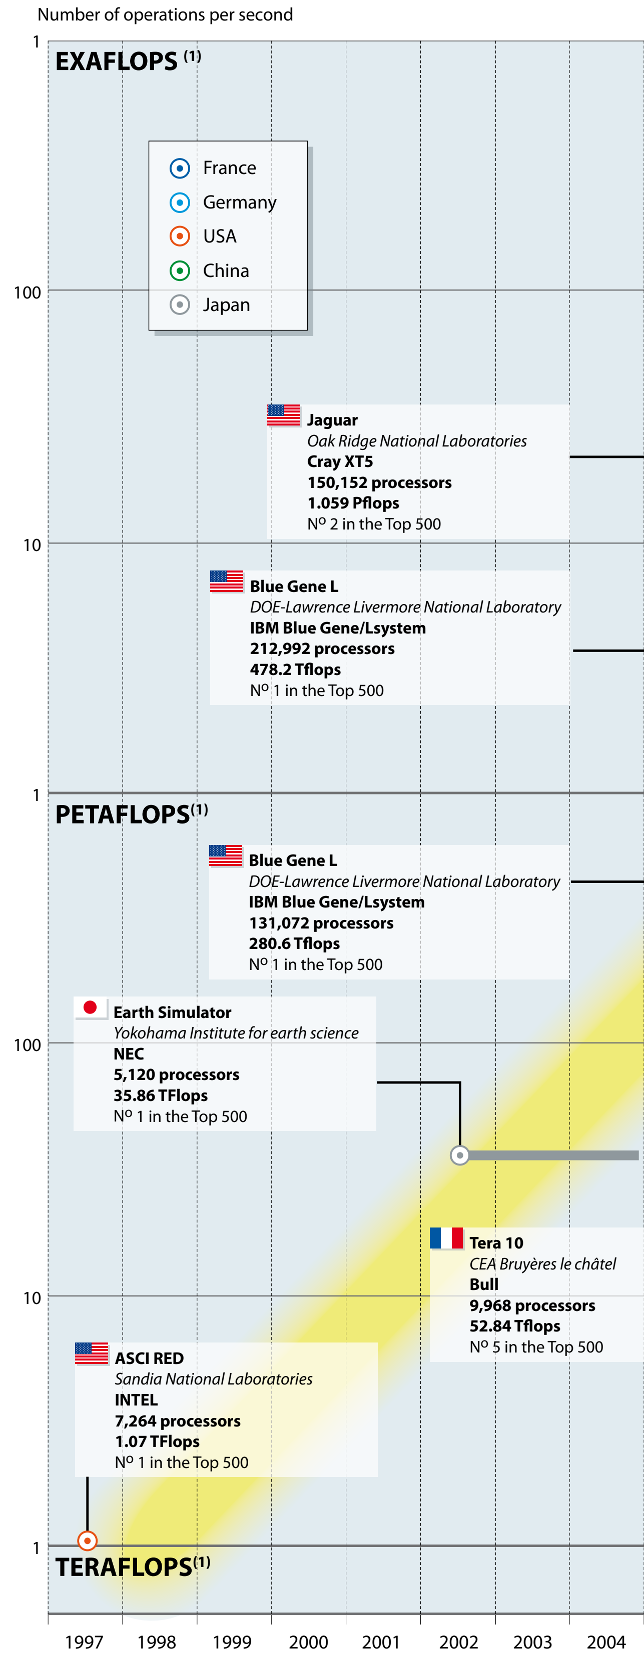

ning

Isabelle Bellin,
Scientific journalist

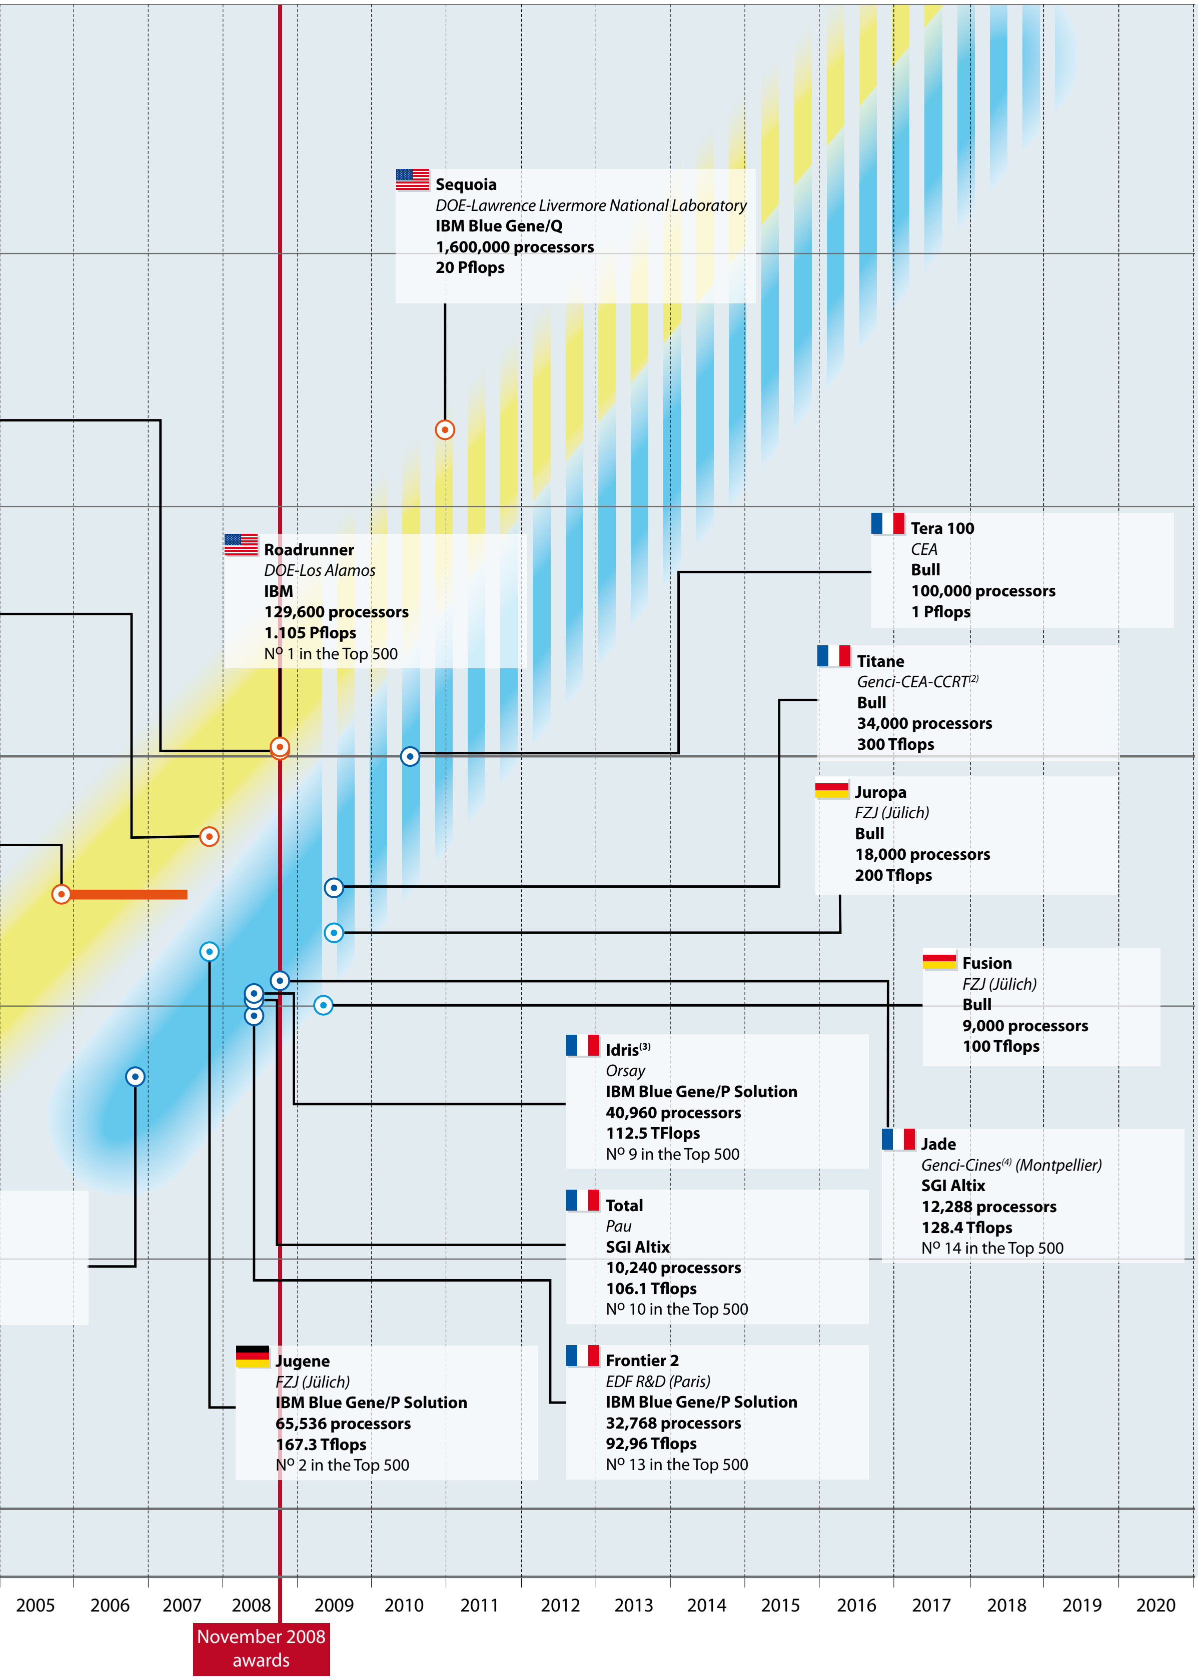

(1) A teraflops, a petaflops and an exaflops correspond respectively to processing capacities of 10^{12}, 10^{15} and 10^{18} operations per second.

(2) Centre for Research and Technology Computing.

(3) Institute for Development and Resources in Intensive Scientific computing.

(4) National computing centre for higher education.

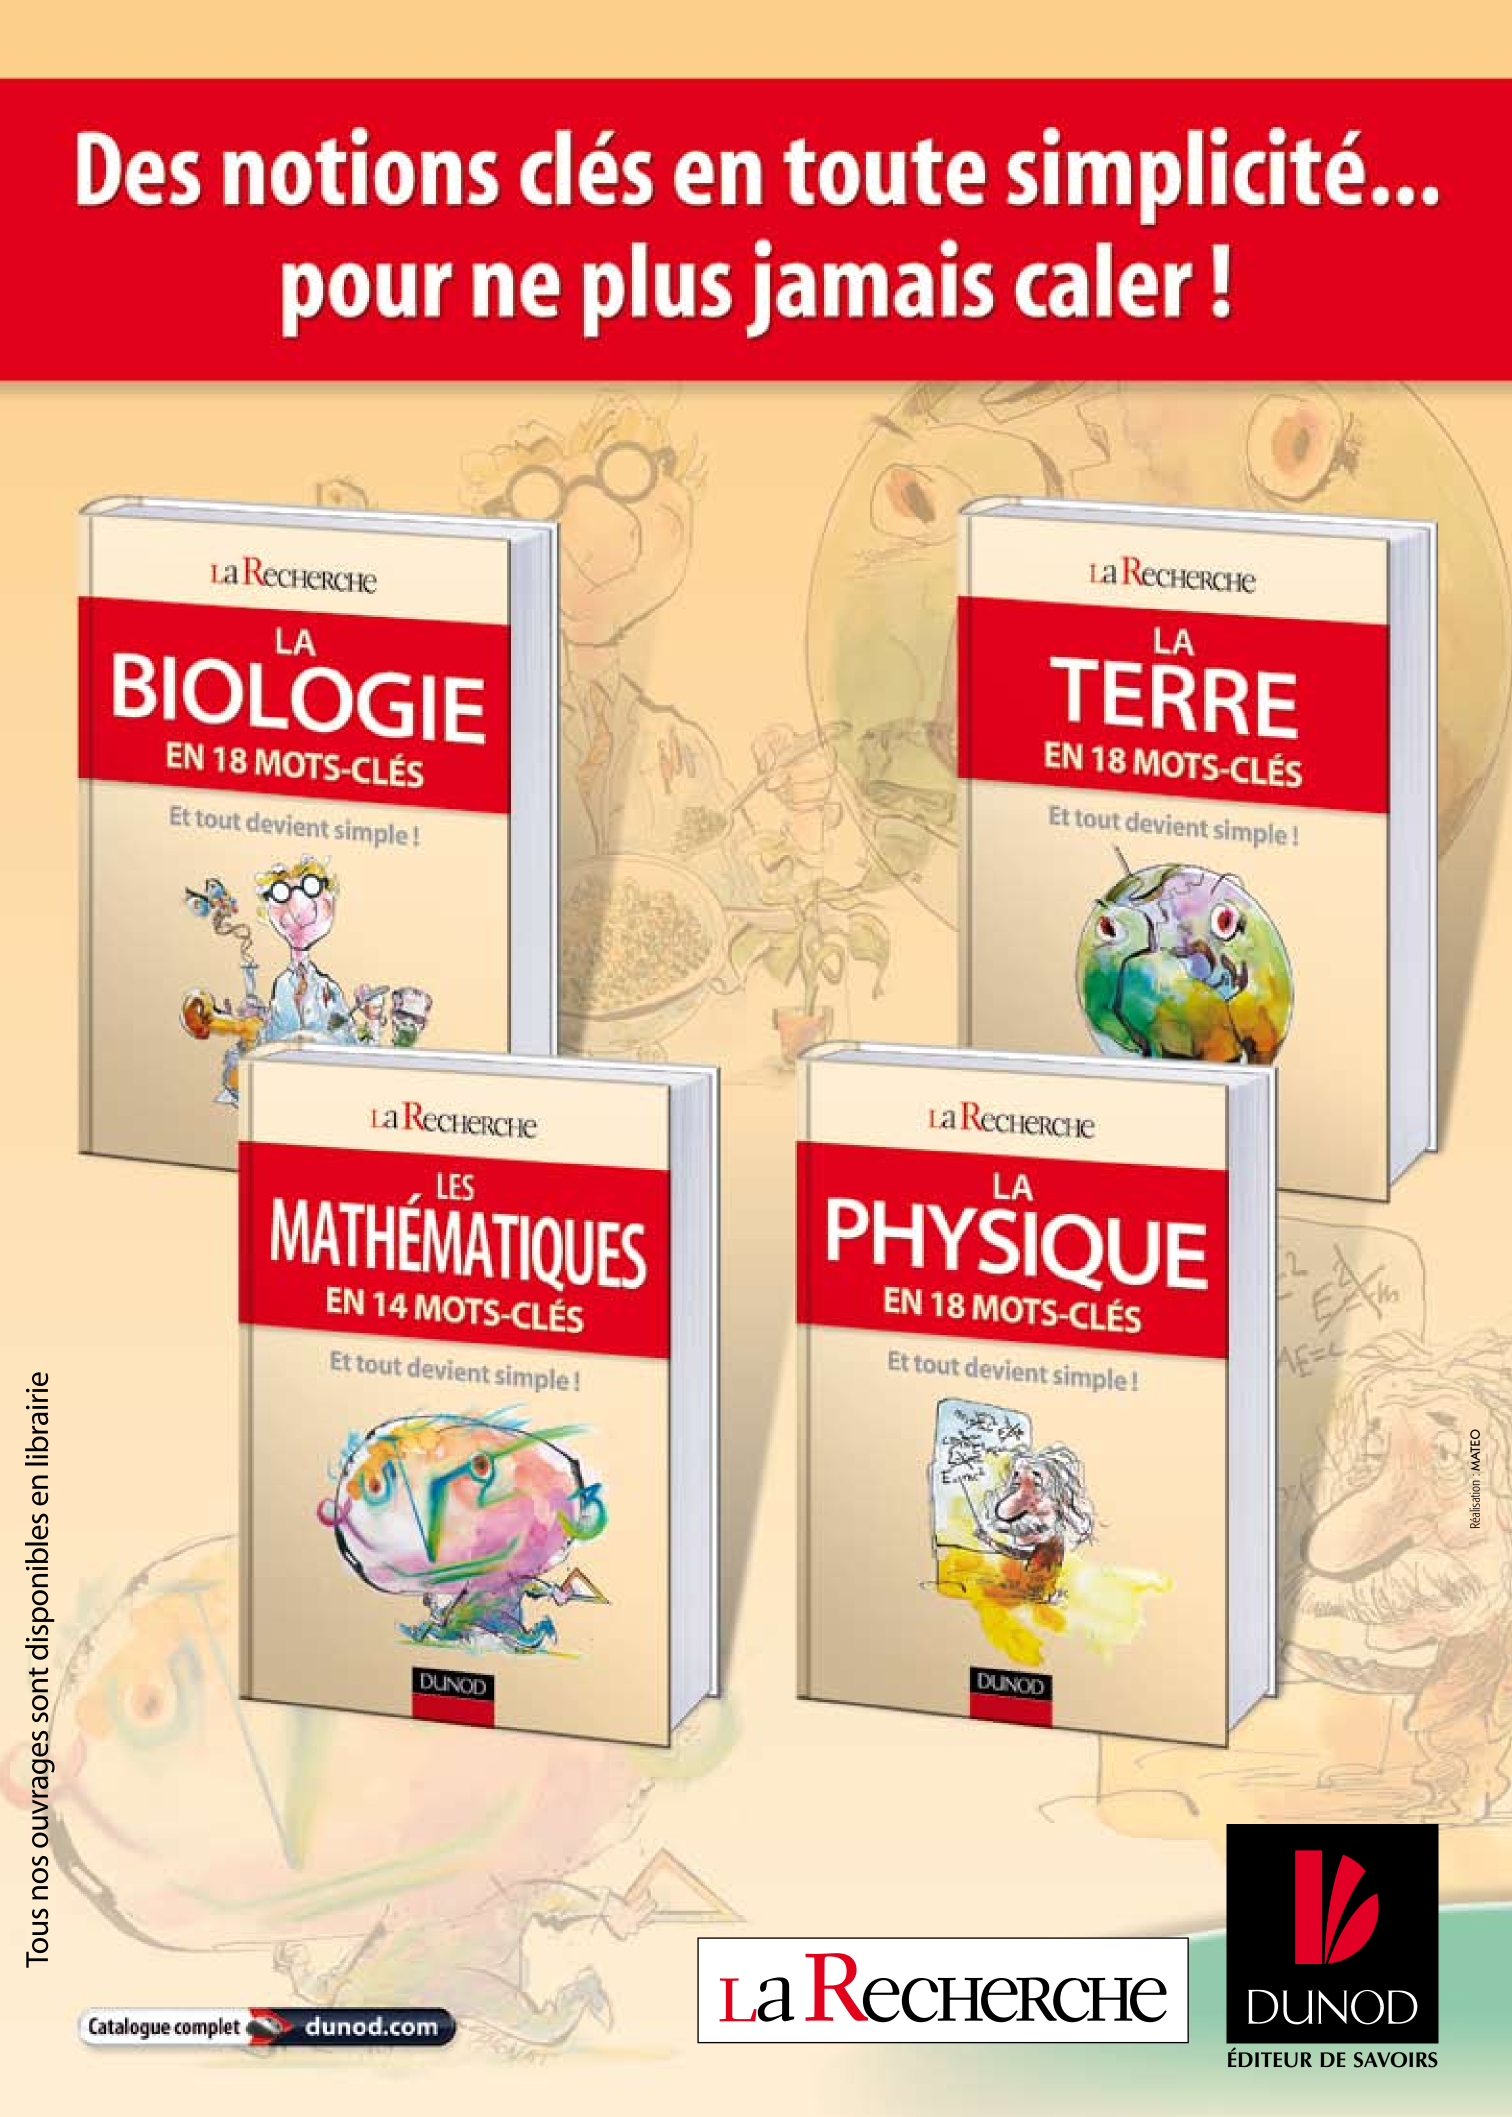
Des notions clés en toute simplicité...
pour ne plus jamais caler !
La Recherche
LA BIOLOGIE EN 18 MOTS-CLÉS
Et tout devient simple !
La Recherche
LA TERRE EN 18 MOTS-CLÉS
Et tout devient simple !
La Recherche
LES MATHÉMATIQUES EN 14 MOTS-CLÉS
Et tout devient simple !
DUNOD
La Recherche
LA PHYSIQUE EN 18 MOTS-CLÉS
Et tout devient simple !
DUNOD
Tous nos ouvrages sont disponibles en librairie
Réalisation : MATEO
Catalogue complet dunod.com
La Recherche
DUNOD
ÉDITEUR DE SAVOIRS

High-performance computing in the fight against cancer

Scientists have long pondered the seemingly impossible dream of using images of the patient's body to make accurate calculations, during radiotherapy sessions, of doses of ionising radiation to be 'deposited'. With high-performance computing, this dream is finally becoming reality.

Renaud Persiaux,
scientific journalist

Half of cancer patients undergo radiotherapy treatment. How can calculations of these radiotherapy doses be optimised to treat tumours more successfully, whilst sparing the healthy surrounding tissues? This calls upon highly sophisticated techniques, involving the development of a complex simulation of the position of the photon or electron beams in question, converging towards the tumour cells, using the patient's scanner images.

Dose calculation in less than 10 minutes.

Calculating such doses is one of Aurélie Isambert's daily tasks at the medical physics division of Gustave-Roussy Institute's radiotherapy department, where she is a medical physicist. *"The properties of the beams used (their shape and particle type) depend both on the size and location of the targeted tumour. The analytical methods currently used for such calculations are fast, but sometimes less satisfactory in certain configurations, for example, in the presence of heterogeneous environments such as lungs or bones. Another more accurate method of calculation, based on the statistical 'Monte-Carlo' method, has been around for years but has as yet been unusable on a daily basis because it takes up far too much computing time."* It consists of simulating random selections to study interactions between the particles and the regions of the human body that the beams cross.

Improved accuracy in dosimetric prediction

"The more we simulate particles, the more accurate the dose calculation becomes, but this also entails higher demand in terms of computing power: with a 'standard' computer, it can take several days to calculate the dose for a series of irradiation beams. Whereas we need a very fast system (less than ten minutes) to keep bottlenecks in the hospitals to a minimum", explains Aurélie Isambert.

The target for the Télédos project was to develop both routine and remote use of the Monte-Carlo method, to obtain statistical uncertainty of just 2% on the doses, within ten minutes maximum. This project was coordinated by Bull and brought together specialists in the field of radiation and computing. The project, which was funded by the French National Research Agency (ANR), was successfully completed last February.

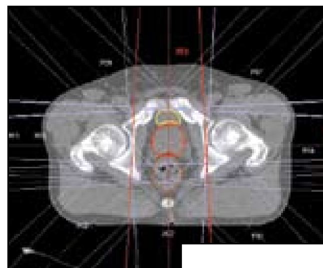

Using the ISOgray™ software (Dosisoft) to simulate the position of the beams as they will be emitted during therapy.

A world premier

"This is the first project of its type in the world, after a few attempts – for the most part American - which never got past the experimental phase due to the size and cost of equipment", states Mehdi Benkebil, R&D director at Dosisoft (a spin-off of the Gustave-Roussy and Curies Institutes), the project's industrial partner. Specialising in software for radiotherapy, it hosts the dedicated supercomputer built by Bull.

Its modestly-sized architecture consists of a large unit, 60cm wide by 90cm high, housing a large number of standard microprocessors, up to 99 networked processor cores. All this for 20,000 euros. *"We parallelised the computing code, in collaboration with the Saclay CEA (French Atomic Energy Commission), to use it on several processors and speed up dose calculation without compromising efficiency"*, adds Mehdi Benkebil. The next step is to contact anti-cancer centres likely to be interested in this new telemedicine service.

Radiation machine: Oncor linear accelerator (Siemens OCS).

Does simulation hold

Transforming an instrument used in fundamental physics into a therapeutic tool. In short, designing lasers capable of producing protons for treating cancerous tumours - this is the goal of the Saphir project.

Renaud Persiaux,
scientific journalist

Proton therapy is a technique for treating cancerous tumours which is more accurate and safer than traditional X-ray radiotherapy: it uses protons (hydrogen nuclei) which are much more difficult to produce than X-rays but they are very accurate from a 'ballistic' point of view. It is particularly suited, especially in children, to highly localised cancers in sensitive areas such as the eye or brain, and it limits the risk of secondary tumours because it does not touch the healthy tissues. Despite these advantages, it is rarely used (in France, only in Orsay and Nice), due to size constraints: the production of protons requires large particle accelerator machines.

Proton acceleration involves complex physics

For a few years now, several research teams across the world have been attempting to produce these protons by means of small lasers using ultra-short pulses. Among them, Erik Lefebvre, a physician at the CEA, Victor Malka and Jérôme Faure, of the Applied Optics Laboratory (ENSTA – Grand Ecole of engineering), all three awarded the Research Prize in 2008.

The idea of using small lasers came about by accident, in 2000, in the Lawrence Livermore laboratory (California) when Tom Cowan's team succeeded in accelerating protons using lasers emitting ultra-short pulses, in the order of several hundred femtoseconds (10-15 seconds): unfortunately such fundamental physics research tools are far from being of any therapeutic use.

In fact, to destroy the tumour and nothing else, the protons must be accelerated at a very accurate energy level: 70 MeV (million electronvolts) for the eyes and 200 MeV for the brain, for example. Whereas, *"for the moment, even with 'colossal' light power, ultra-short pulse lasers accelerate the protons to much weaker levels which are too variable, between 1 and 10 MeV"*, explains Erik Lefebvre.

What changes should be made to the laboratory lasers to save on energy and improve beam precision without electrical power consumption shooting off the scale? We are still a long way from the solution, since proton acceleration involves complex physics.

First, the lasers accelerate electrons, with energy levels on a scale that triggers numerous physical phenomena, which could potentially be harmful or cause experimental fluctuations. These electrons then form plasma, a kind of negatively charged loose cloud, in a state of equilibrium. This plasma surrounds a target that is coated with a layer which is positively charged with protons. The plasma then accelerates these protons using a very

Simulation of proton acceleration

A fine layer of protons is applied to the centre of the right-hand side of a solid, thin target. A short and intense laser pulse then hits the target from the other side, accelerating the target's electrons with a high energy beam (image 1). Part of this pulse is reflected by the target (image 2). The electrical field created by the energised electrons on the other side to the laser's impact accelerates the protons on it, which gradually detach from the target (images 3 and 4).

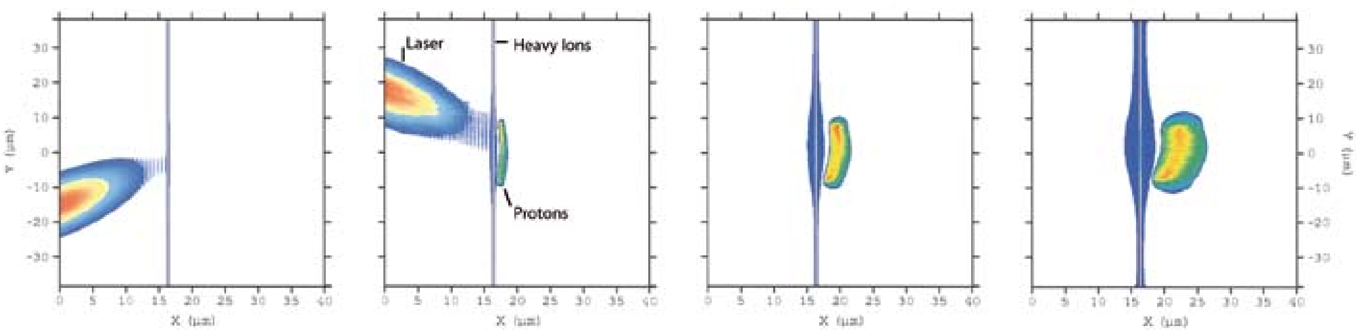

© CEA

the key?

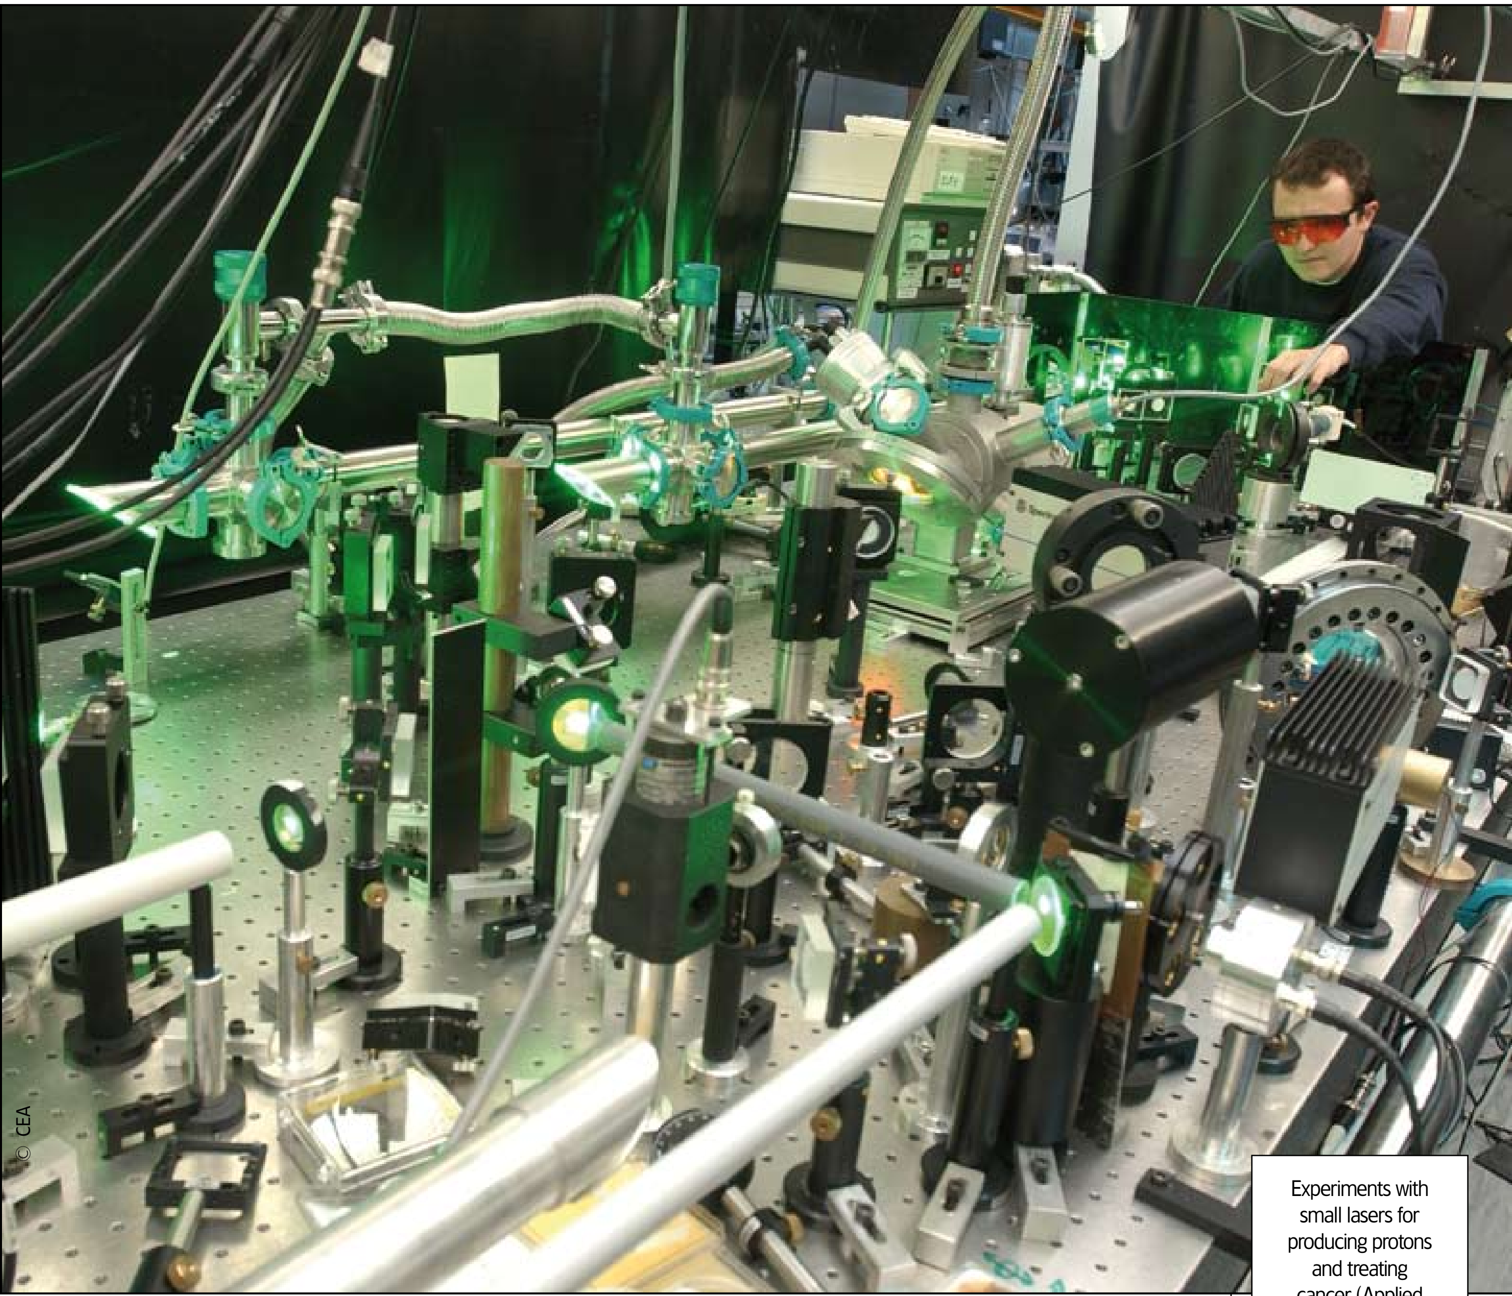

© CEA

Experiments with small lasers for producing protons and treating cancer (Applied Optics Laboratory – Ensta – Ecole Polytechnique).

narrow beam. This is more or less all that we know. This is where high-performance computing comes in. Highly realistic simulations of laser-matter interactions are needed to guide researchers in their work.

However, *"these phenomena do not involve linear description"*, explains Erik Lefebvre. *"To solve these equations, we quickly reach a point where we are simulating billions of particles moving around in a volume of hundreds of millions of mesh cells, during thousands of time steps. A colossal quantity of operations, only within the capacity of large, massively parallel scientific computers."*

To achieve this modelling process will be one of the main stages in the Saphir project, which aims to experimentally demonstrate the acceleration of protons up to energies of 150 MeV. Saphir brings together the CEA, the Applied Optics Laboratory, the Curie Institute, the Gustave-Roussy Institute, as well as Amplitude Technologie, a small business that supplies lasers to research laboratories, and Imagine Optic and Dosisoft, specialists in laser and radiotherapy applications respectively. Results are scheduled for 2015, before initial clinical trials can even be considered.

Brain: from structure to function

Performance in brain imaging combined with computer imaging has paved the way for a real revolution in neuroscience.

Dominique Chouchan, scientific journalist

Which zone of the brain is used for reading or arithmetic? Can we diagnose pathologies such as Alzheimer's or Parkinson's disease at an early stage? Such questions assume that we can describe the links between structure and function of each cerebral region. To do this, functional magnetic resonance imagery (MRIf)[1] tends to be the tool of choice. High magnetic fields such as those produced on the NeuroSpin platform[2] at the CEA in Saclay offer new perspectives in terms of space-time image resolution. But from observing the images to understanding them, the most important task lies in processing and interpreting them.

MRIf was conceived during the 1990's. This non-invasive technique is based on the observation of variations in oxygenation of the blood. Even if its time resolution (a few seconds) is slower than that obtained using magnetoencephalography (MEG), its spatial resolution is currently accurate to the millimetre. With the future NeuroSpin MRI system of 11.7 Teslas it will be accurate to 300 micrometres. *"MRIf is now the best possible way of producing functional mapping"* explains Bertrand Thirion, team manager for Parietal at Inria.

"Functional MRI is currently the best method available"

A truly functional atlas of the brain

One of the challenges is to equip neuroscientists with geometric referencing systems upon which they can base their studies on brain activity. But what kind of reference points should we choose? Standard brain mapping techniques are inadequate because the shape of the brain varies from one individual to another. The idea implemented by the Parietal team at Inria is to refer to a few dozen sulci separating the cerebral regions, which according to neurobiologists would correspond to the boundaries between the functional zones. The problem therefore lies in identifying these sulci from an image. This has been successfully tested using stochastic (probabilistic) methods: *"We are one of the few teams in the world to have connected the geometric information provided by the sulci to the functional information provided by the activity maps of the brain surface"*, points out Bertrand Thirion. Eventually, this type of work could contribute to the creation of a 'functional atlas' of the brain.

Another question is how does the brain respond to a given stimulus? This is an aspect which the Parietal team are also examining, together with the Unicog (Inserm-CEA team) laboratory directed by Stanislas Dehaene. For example, if an individual is presented with a mass of points, their brain will react differently according to the quantity of points shown. *"By simply observing the activated zone, we have succeeded in predicting the quantity actually perceived by the brain"*, adds Bertrand Thirion. Beyond the fundamental aspect, this kind of knowledge gives us better insight into certain disorders, particularly in areas of arithmetic or language and could eventually lead to establishing criteria for diagnosis. This assumes that classification methods capable of distinguishing between normal and pathological functions can be perfected: *"Finding the zone that differentiates two populations is a bit like searching for a needle in a haystack!"* comments the researcher.

(1) Magnetic resonance imaging makes use of phenomenon of magnetic resonance occurring in atom nuclei such as hydrogen atoms, within an intense magnetic field.

(2) Inaugurated at the end of 2006, the NeuroSpin platform, designed for studying the human brain, comprises MRI systems at 3 and 7 Teslas and soon, 11.7 Teslas.

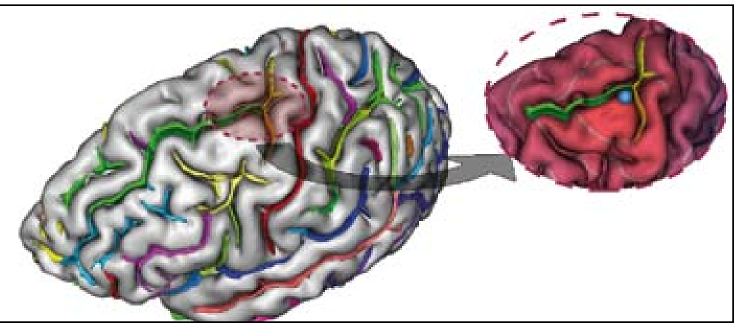

On the left of this diagram, an anatomical map of the sulci and on the right, the location of an activity zone: the blue dot near the intersection of the three sulci. This system of local coordinates (position in relation to the sulci) enables more accurate location of standard coordinates.

© B. Thirion-A. Tucholka - Parietal

Modelling the whole human body

Spotlight on the virtual physiological human (VPH). A European project is working on a global and unified framework for researchers and doctors, with the added advantage of many potential applications.

Renaud Persiaux,
scientific journalist

Modelling the entire human body, or to be precise, all of its physiological mechanisms – science-fiction? No, this is the ambitious, but very serious aim of a European framework project, named Virtual Physiological Human (VPH). To achieve this aim, in June 2008 a network of excellence involving many universities, research organisations and hospitals across Europe and New Zealand was set up and is set to continue for four and a half years.

Brain: working towards higher precision surgery

The Genius project (Grid-enabled neurosurgical imaging simulation) aims specifically to optimise the medical decision-making process using 3D images of a patient's cerebral vascular system. Based on programs designed by Peter Coveney and Marco Mazzeo (University College London), the system calculates a representation of key parameters. Generated just before surgery, these allow the doctors to test their results before performing the procedure.

It's an ambitious project which aims to offer a global framework enabling all researchers (academic, clinical and industrial) the world over to work together simultaneously on a unique and integrated model of the human body.

Its sights are set on achieving a better description and understanding of our pathophysiology, to make predictive hypotheses and simulations and to develop and test new therapies.

Customised treatment

From cancers to the AIDS virus, from drug design to precision surgery, potential applications will encompass the whole spectrum of human health. This can extend to customised therapy solutions based on individual patient data; reduced need for animal experimentation and more in silico tests; new global approaches to medicine, no longer perceiving the body as a collection of organs to be treated individually, but as a whole made up of several organs; and lastly development in prevention methods and quality of life.

© Fotolia

Representation of the human anatomy in 3D. In red, the colon.

"Building this virtual physiological human is a bit like trying to put together a unique work of art using the pieces from several different puzzles", explains Peter Coveney, director of the Centre of Computational Science (University College London) and one of the two project managers. Most of the modelling involves high-performance computing, but integrating it will be a complex task. *"One of the challenges will lie in integrating the various models and allowing them to communicate through a common computing language (probably XML)".* Another critical issue will be the use of very specific algorithms, allowing various computing resources to be rapidly allocated to medical emergencies. This will be the first step towards an 'on demand' medical computing method where doctors will be able to borrow computing time from supercomputer structures to optimise their decisions.

AIDS virus: new advances in resistance patterns

Currently, to find the antiretroviral agents best suited to the resistance pattern of the virus and to the patient's genetic profiling, doctors must often proceed by trial and error and use complicated and time-consuming genotyping tests. In 2008, using simulations running on supercomputers, Peter Coveney and his team (University College London) managed to predict the level at which a given treatment can inhibit the viral enzymes displaying resistance.

Seismic simulation: a n

Predicting the impact of an earthquake – dream or a reality? With the simulation software developed by CEA's Environmental Assessment and Monitoring Department, a new modelling method is now available for estimating the propagation of seismic waves and their effects on buildings. Feedback from an experiment is paving the way for potential new prediction tools.

Marie Schal,
scientific journalist

5.5 on the Richter scale, the earthquake spread fast through the small town in the foothills of an Alpine valley. Although built on a sedimentary basin, whose loose sub-soil trapped the energy and strongly amplified the seismic waves, the buildings withstood the force well. The tremor only triggered one boulder slide, on one of the mountain's slopes.

This event is fictional: it is the result of a groundbreaking simulation carried out at the end of 2007 on the CEA's Tera 10 supercomputer. The operation took 500 processors 40 hours, totalling 20,000 computing hours, making it possible to model the earthquake over a 11 x 11 km zone, 2 kilometres deep, for a time lapse of one minute, long enough to estimate the impact of the tremor.

"What is new is that we have now integrated the propagation of seismic waves on the ground and their effects on the structure of buildings, within one and the same calculation", explains Christian Mariotti, of the CEA's Environmental Assessment and Monitoring Department (DASE). It took ten years to develop the simulation software, which was christened Mka3D. Its originality lies in the method used to construct the model. Traditionally, earthquakes have been modelled using the finite element method, which consists of discretizing the ground into virtual cells and then simulating wave propagation through the geologic layers represented by this mesh. However it is difficult to apply such a method to processes dealing with fracturing and subsidence, when there is loss of contact between the environmental constituents, which become separated. CEA researchers, who began by looking into the simulation of explosions, preferred to opt for a so-called 'particle' method, developed in the 1970s, by Peter Cundall at the Imperial College of London, but which until now remained largely sidelined.

Rather than discretizing structures into virtual cells, the idea is to discretize them into solid and rigid particles, which interact with their neighbouring particles according to Newtonian Mechanics[1].

(1) C. Mariotti, Geophys. J. Int., 171, 857, 2007.

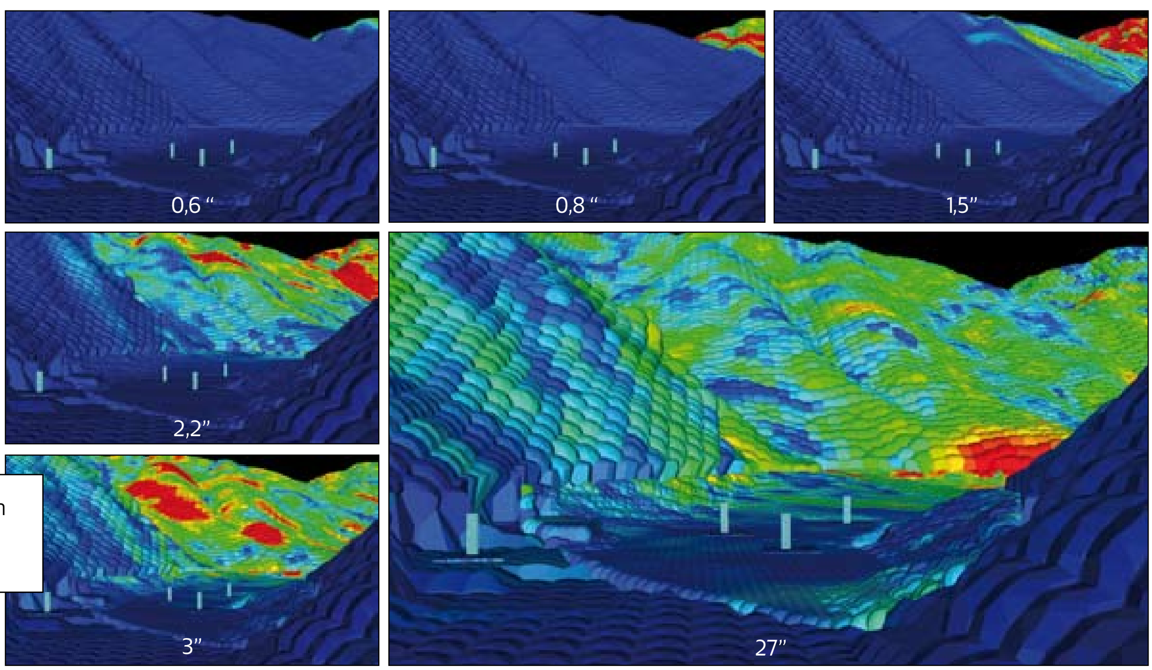

Propagation of seismic waves.

© CEA

ew approach

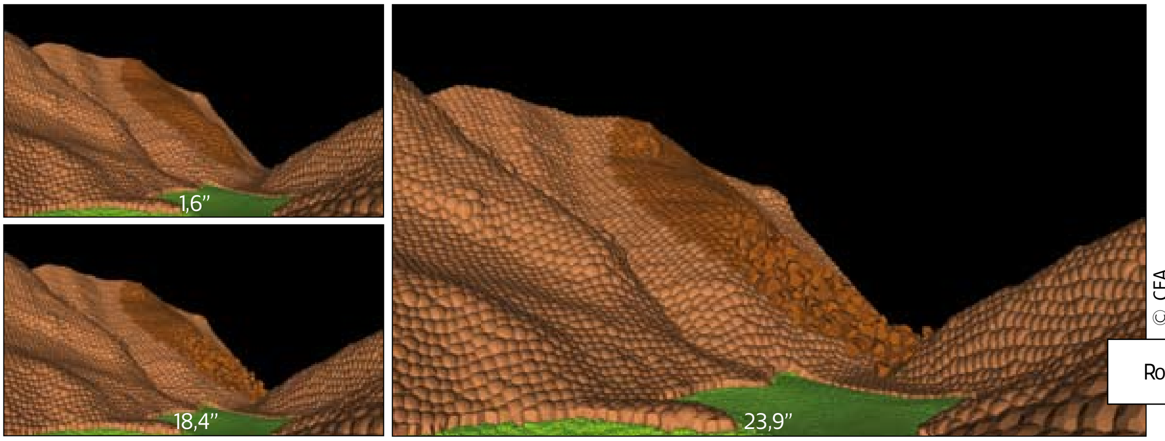

Rock slide.

"We are thus very close to the physics of nature", specifies Christian Mariotti, *"the links between particles are represented by small beams when the environment is still cohesive, then by contact force and friction when it is fragmented."* The load-bearing structure of buildings – floors, shells, pillars – is divided up in the same manner, even if the particles are 100 times smaller here (1 metre in diameter for buildings, compared with 100 m for the ground).

The method is demanding in terms of computing performance: to simulate the alpine earthquake 3 million polyhedral particles need to be accounted for, each equipped with 6 degrees of freedom (translations and rotations) and relating to ten neighbouring particles. Around sixty equations need to be solved per particle, at each stage of the calculation, not forgetting that each stage has had to be iterated 100,000 times to simulate a one-minute tremor. And this doesn't even account for the specific algorithms applied to processing contacts when subsidence occurs. *"Ten years ago, carrying out such simulations was unthinkable, we were struggling with just a few particles on our computers" recounts the researcher. These numerical methods were compromised by their high demand in computing power and we were only able to process very small-sized cases. Today, technological progress has reduced computing times and we are already preparing for migration to the CEA's future Tera 100 supercomputer, which by 2010 will represent a factor-10 gain in speed."*

Calculation iterated 100,000 times for a one-minute tremor

From fiction to reality

The aim behind the simulation launched in 2007 was to demonstrate the method's feasibility, with the ultimate plan in mind of applying it to a real town, to predict the impact of a potential earthquake and to present it as a tool for testing town-planning or building design methods. *"Generally, structural codes used in traditional simulations enable us to predict the fracture of a building, but that's as far as they go: will the building totally collapse or simply fissure? The upshot in terms of human life and public safety will of course be different"*, emphasises the researcher. A further benefit is the possibility of analysing seismic movements in the immediate vicinity of fault lines, where the non-linear nature of the phenomena in play (with accelerations greater than gravity) defy traditional modelling techniques. Outside the field of earthquakes, another area of research focuses on modelling the impact of accidental explosions, of a chemical product factory or warehouse for example, in view of improving designs for surrounding safety installations.

International test bench

The next step consists in validating results produced by the Mka3D software, by comparing them with those obtained from past tried and tested methods of seismic modelling. In addition to comparisons with data from real earthquakes, the similarity of results produced by different simulations is indeed a key argument in favour of their reliability. CEA's team is currently participating in the Euroseistest international numerical test bench or 'benchmark', coordinated by the CEA, the Grenoble Laboratory of Internal Geophysics and Tectonophysics and the Aristotle University of Thessaloniki in Greece. The objective is to simulate a fictional earthquake in the lake Volvi region in Macedonia. Fourteen French, Italian, German, Slovakian, American and Japanese laboratories are participating in the exercise, which was launched in September 2008. Each participant conducts blind simulations, using the same input data (relief, geological stratum, source and earthquake properties), then must reproduce the seismic signals at precise points of the model. The results are scheduled for end of 2009.

Model cars

Simulating all or part of a vehicle is now the norm for most manufacturers, up to microscopic scale for certain parts.

Marie Schal,
scientific journalist

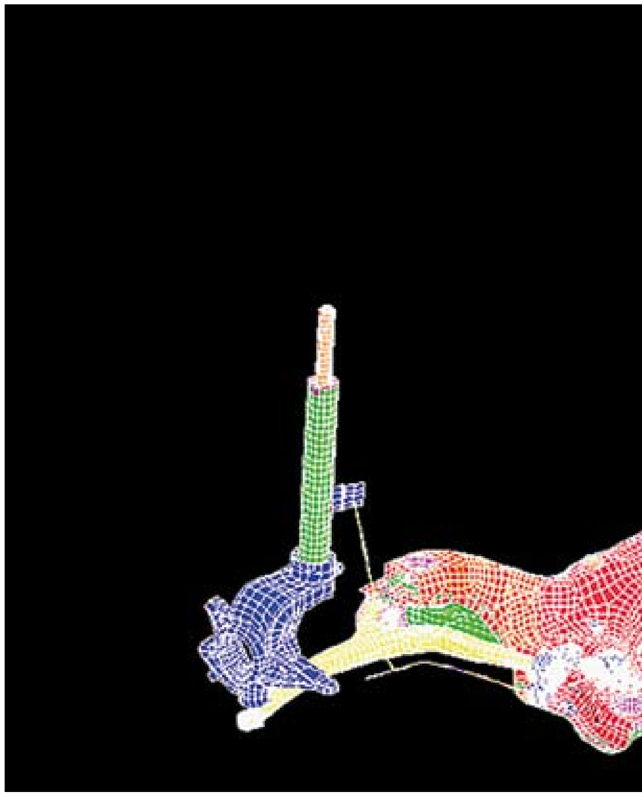

The association of the words 'simulation' and 'automobile' or 'aeronautics' often conjures images of virtual driving software and aerodynamic studies. There is however another sector of transport which has become an avid consumer of high-performance computing: that of elementary vehicle components, ranging from mechanical parts through to door joints or engine mounts. Its recurring theme: economise on weight, down to the nearest gram, to minimise fuel costs and consumption. *"Modelling has become an essential process. It plays a key role in optimising our products by considerably shortening our design lead-times"*, stresses Daniel Benoualid, director of the Hutchinson research centre, subsidiary of the Total group, specialists in rubber processing and suppliers of automobile and aeronautical parts.

The equipment manufacturer, which in 1997 brought in parallel computing techniques developed in collaboration with Onera, has just increased its computing power by integrating a cluster of around 5 teraflops(1). The simulations involve different disciplines of physics: structural mechanics for the analysis of strain and stress, fluid mechanics coupled with thermal analysis for designing injection or part extrusion tools, vibro-acoustics for predicting noise levels in passenger compartments, or even electromagnetism to account for the rapid growth in mecatronics, in other words the integration of electrical components in mechanical parts.

Today, the issue is no longer simply one of modelling an individual component, but of simulating its performance in the full vehicle, or at least in its immediate environment. This entails working on complex systems such as a full suspension system or a door locking system to analyse the performance of its joints. Multi-scale models also need to be developed. For the past twenty years, detailed research into materials such as plastics, elastomers and metals have indeed demonstrated their heterogeneity: they are formed from aggregates and these microstructures have an effect on fissure damage or initiation. Hence the importance of factoring in this material at a microscopic scale in the models. *"We have just acquired a nanotomography system for analysing our materials in a non-destructive way on the hundred nanometer scale. The idea is to model these microstructures to better predict the service life of parts"*, states Daniel Benoualid. This will give them a reliable alternative to real tests on fatigue testing machines. *"To do so, we will need to install new computing tools"*, continued the researcher. *Of course, initially, simulations will be time-consuming, like the first three-dimensional modelling 20 years back. But this is the future."* In the same line of thought, the Optiforge national project, with the participation of PSA Peugeot Citroën, aims to integrate the microstructure into the modelling of manufacturing processes for industrial forgings, to better predict their

"Improving prediction of the service life of parts"

(1) **A teraflops is the equivalent of a processing capacity of 10^{12} operations per second.**

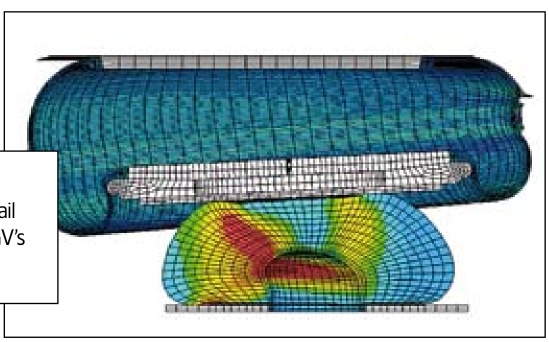

The same type of simulation is used in rail transport, here on a TGV's bolster spring.

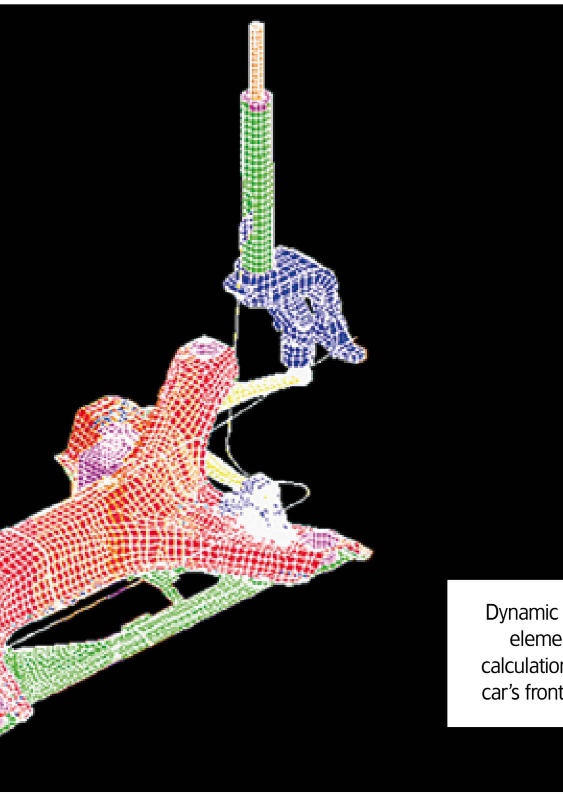

© Hutchinson

Dynamic finite element calculation of a car's front axle.

mechanical resistance to wear. This three-year project was launched in 2006 and was funded by the ANR (French National Research Agency), bringing together the Paris CEMEF (Mines ParisTech Centre for Material Forming), Lyon INSA (National Institute of Applied Sciences) and Angers ENSAM (National Institute of Industrial Engineering), as well as metallurgy manufacturers.

Improved calculation of the performance of materials

The increase in computing power has at last made it possible to tackle another essential question: that of part geometry and materials performance. One sheet-metal piece will never be exactly identical to the next off the vehicle production line; industries work with manufacturing tolerances, even if these are minimal. The same goes for the deformability of a joint, for example. Factoring in these statistical fluctuations for each of the vehicle's constituents increases the reliability of prediction of the performance of a part within the system. *"We have just started doing this, using specific algorithms which obviously require far more calculations'* continues Daniel Benoualid, *"rather than producing unique values, the models will then produce response ranges that are much more realistic. Such methods should, in the medium-term, become widespread."*

"Three questions for"

Philippe Thierry, High-performance computing (HPC) specialist, Intel.

Interview by **Isabelle Bellin**

Intel works in close collaboration with publishers of scientific software programs used in the automobile and aeronautics sector to use supercomputers to their best advantage.

What does this entail?
We help them to develop their codes and to diagnose their application in areas such as crash, fluid mechanics and vehicle structure simulation. Parallelism in HPC computers (see article on page 48) is a hot topic, but don't forget that calculations can only be executed in parallel if the code of the application in question is fully distributed between the different cores (or computing units) of each processor, which is rarely the case.

Does the same go for automobile applications?
No, due to the very nature of the physics and numerical schemes involved, these applications can barely use more than 256 to 512 cores, some can only use a few dozen. The main objective is therefore to optimise the sequential code process, which rarely uses more than 20% of the theoretical capacity of processors. The key is to increase the number of operations per clock cycle, to minimise time where processors are waiting to execute instructions by, for example, intervening on the memory access. Depending on the initial state of the code, acceleration factors from 1.5 to 5 are attainable.

Can the parallelism be optimised too?
Yes, we try to reorganise the code first to increase its share of parallel computing. We then work on minimising waiting times between calculations by improving core synchronisation, whilst balancing out their computing loads as much as possible. We also work on minimising message sizes and optimising the organisation of communications according to the architecture of the computer's communication network. Lastly, with current architectures, calculations can be distributed simultaneously across different computing nodes, generally containing eight cores. This two-tier (or more) parallelism also needs to be used to advantage.

Modelling where

From lab to lecture hall

Dominique Chouchan, scientific journalist

Since 2008, chemistry students from the University of Reims Champagne-Ardenne and even secondary school students have been benefiting from teaching content that is taken directly from recent results in molecular modelling. *"As modelling experts, we regularly publish our work, however we are only able to allow students to use our findings on rare occasions"*, regrets Eric Hénon, teacher-researcher at the University of Reims Champagne-Ardenne. He created his project with this in mind, stemming from his desire to use methods and numerical three-dimensional tools for teaching chemical reactions. The idea is to break away from the purely symbolic limitations of poor-quality molecular imaging (in two dimensions). This project, which began in 2007 and is based on the supercomputer ROMEO 2's resources[(1)], is just starting to bear its fruit. *"In chemistry, three-dimensional understanding of objects is difficult but essential, especially in the field of stereochemistry, which studies the spatial arrangement of atoms within a molecule"*, adds Eric Hénon. To bridge the gap between knowledge gathered from research and educational content, all we were lacking was skilful image engineering. The Region's co-funding of the project has allowed us to recruit the staff we needed.

Students can now work on models created by researchers

Secondary schools at the heart of the experiment

Result: presentations developed based on scientific publications[(2)]. At university level, these can be used as teaching material in lessons or consulted freely by students, in particular via the establishment's 'virtual office', a numerical platform available to lecturers and students. In secondary education, the tests were performed at the Libergier Lycée. *"Partnerships with secondary education establishments were planned right from the start, if only because a percentage of our secondary school students will be our future undergraduates"*, underlines Eric Hénon. The project has met with growing enthusiasm from chemistry teachers. Proof if any were required that the complementary nature of educational research projects led by university lecturers is not a pipe dream.

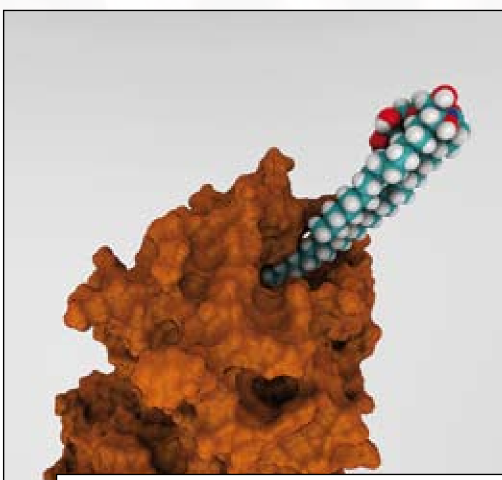

These images are extracted from animations of reactions between molecules. left: reaction between a protein and a ligand (a drug); right: catalysis process.

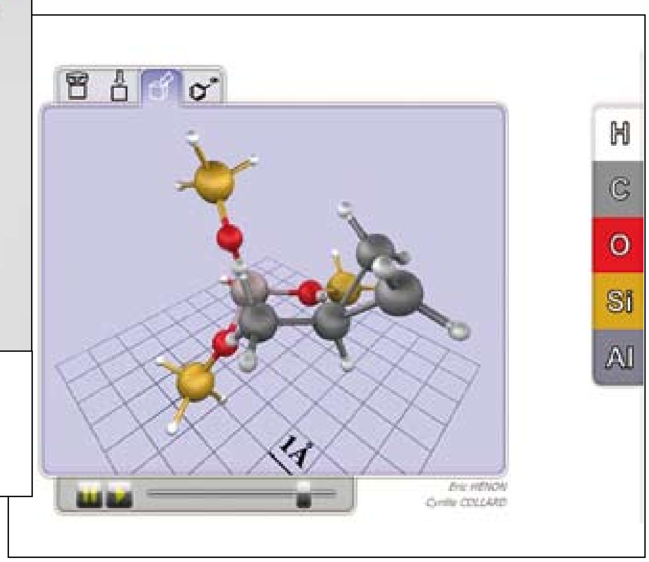

(1) The regional supercomputer in Champagne-Ardenne (ROMEO 2, Bull architecture), equipped with 104 cores, with a power of 600 Gflops (billions of operations per second).

(2) See the website http://sead.univ-reims.fr/courses/3DCHEMISTRY.

you least expect it

The game of 'Go'

© Université de Tainan (Taiwan)

In 2009, the professional Taiwanese Go player, Zhou Junxun, lost with handicap 7 against the program MoGo; this is the first time this has ever happened.

Computers have already beaten champion chess players, are they as good at playing Go? Not yet, for the time being, the humans still have the winning edge, but how long for? *"At the start of this year in Taiwan, it beat a champion player at a simplified form of the game"*, outlines Olivier Teytaud, researcher for the TAO team at Inria: on a 19 x 19 goban[3] with a handicap of 7 (7 stones, or pawns, placed on the goban at the start of the game) and on a 9 x 9 goban with zero handicap. The winner's name? The MoGo program, developed by French researchers (the majority working at Inria and the CNRS). To produce a better result, they would have to make the algorithms more 'intelligent' and increase computing power. One of the difficulties lies in the astronomic number of possible permutations for placing each stone: *"The total number of different games is approximately 10,600 and, each turn, the player has around 360 different possible moves, this number obviously decreasing over the course of the game"*, explains Olivier Teytaud. The algorithms created consist of simulating a very high number of random games, each game also taking into account the previous simulation. After 30 seconds, or once several million supercomputer calculations have been made, the pawn's next move is determined based on the game that offers the highest chance of success. Improvements are constantly being made, mainly based on learning techniques and the use of historical databases (archives of all past matches). But beyond the game, the methods that have been developed could also be applied to other problems requiring vast computing power: stock management, non-linear optimisation, energy management, and so forth.

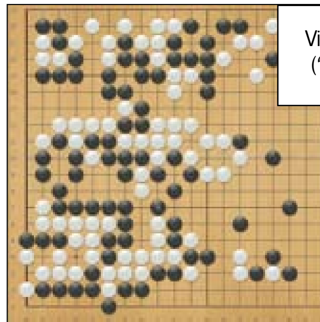

Virtual goban ('go' board).

© Université de Tainan (Taiwan)

(3) The goban is the chequered playing board on which the stones are placed: it is made up of 19 x 19 squares. Beginners play on a smaller board (9 x 9).

Champagne bubbles

It all started around a few glasses of champagne among friends: "The project to model champagne came about during professional and social meetings between Gérard Liger-Belair and myself", explains Nicholas Floquet, CNRS researcher at the Max Mousseron Biomolecule Institute (Montpellier Pharmacy Faculty). Nicholas Floquet works on the interfaces between physics, chemistry, biology and modelling and Gérard Liger-Belair is an experimental physician and teacher at the Reims Champagne-Ardennes University, who is also passionate about the region's famous bubbly[4]. Combining their respective expertise, they embarked upon a project to model champagne bubbles.

"I used a molecular dynamics model, similar to those that I use for studying biomolecules", explains Nicolas Floquet. These simulations of the behaviour of water-ethanol-CO_2 (the main constituents of champagne) sometimes take several weeks on a supercomputer with parallel architecture, such as ROMEO 2 belonging to the Reims Champagne-Ardenne region. Initial findings showed that the diffusion coefficients of CO_2 calculated appear to be consistent with those observed experimentally. The next step will consist in simplifying these models, in order to reduce computing time and simulate larger systems, up to the size of a bubble. The idea will then be to characterise the proteins that give champagne its distinctive taste, as well as certain physico-chemical properties, which could, in the future be used to optimise the quality of champagne. Well, why not?

(4) Gérard Liger-Belair, *Effervescence! La science du champagne*, Odile Jacob, 2006.

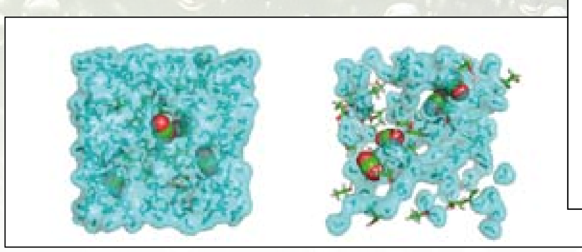

© Nicolas Floquet - Inria

Simulation of migration of carbon dioxide molecules (CO_2) in a 'simplified champagne' - a mixture of water (in blue), ethanol (the bars) and dissolved CO_2 (carbon in green, oxygen in red).

The megajoule laser, a colossus with feet of

Initially, it was feared that the optical components would prove to be the Achilles heel for this future giant built in the Gironde, France. During its development, computer scientists at the CEA and Inria were bent over the cradle like fairy godmothers in their attempts to undo this evil spell.

Xavier Muller,
scientific journalist

Any Superman fan will tell you: every superhero has their weak spot. This monster of light and energy is a keystone of the French nuclear deterrence programme which is scheduled for construction in the Gironde in 2014. Its one weak spot looked likely to be its optical components.

The nature of the beast: strange cracks appearing over time in the windows and lenses of all the powerful lasers and requiring regular and very expensive repairs. Whilst the CEA engineers now have empirical understanding of how to protect the MJL from these fissures, they had to investigate their underlying cause. In 2004, the CEA set up a partnership with Inria (French National Institute for Research in Computing Science and Control) to explore the secret world of glass at the very moment that it tears. Feedback from an example of applied high-performance computing that unveiled the mysteries behind the fissures.

Unlocking the mystery behind mysterious fissures

Five years ago, when Gilles Zérah, from the military application division (DAM) at the CEA, attempted to reproduce the progress of fissures, he rapidly realised the limited nature of these simulations that were based on the 'molecular dynamics' method. Ideal for observing the behaviour of the material on a small scale, the method is too rigorous to describe a phenomenon that occurs on the scale of a fissure. The result of this was a full month's work for each round of simulation and the ten million atoms it needed to deal with. To accelerate these calculations, the researcher turned to two high-performance computing specialists: Jean Roman and Olivier Coulaud, from the Scalapplix team at Inria Bordeaux. The first is an expert in algorithmics and parallel computing, the second has a solid background in applied mathematics and molecular dynamics.

The solution that these three scientists opted for was put together in a thesis, and consisted of conjointly using a rough model and a more precise one. *"We worked with a multi-scale model"*, explains Guillaume Anciaux, the former PhD student of the laboratory. In other words, two models of the material co-habit in the same simulation: the immediate environment of the fissure continues to be described with the molecular dynamics method, whereas the more distant 'outskirts' use a more global approach method, known as continuum mechanics. Used in many other domains such as fluid or structure mechanics, this technique is less demanding of computing power.

Athough the idea seems easy on paper, it is actually quite complex to apply. *"A large part of the thesis work involved was improving the consistent flow of information between the dimensions described and these very different scales"*, recalls Olivier Coulaud. When run consecutively, neither model allowed high frequency waves emitted by the tip of the fissure to pass through, despite their vital importance in the phenomenon of fissuring: unlike seismic waves, they cause the wave to propagate by moving the tensions around inside the material. To correct the problem, the researchers introduced a buffer zone between the immediate environment of the fissure and its 'outskirts': inside this region, both models – molecular dynamics and continuum mechanics – coexist. This new development is beginning to bear fruit: complexities linked to high frequency waves have virtually been eliminated.

Modelling and parallelisation

But after the pure modelling stage, researchers still needed to confront the question of efficient 'parallelisation' of the calculations, a vital step to ensuring that power-hungry simulation benefits

The energy monster

The megajoule laser (MJL), built in the village of Barp, in the Gironde, will concentrate the light emitted by 240 laser beams onto a ball filled with dueterium and tritium (two hydrogen isotopes). Its energy output, 1.8 megajoules, will be so great that it will heat up the isotopes, causing them to fuse together. With this tool, CEA engineers will have a much better understanding of the reactions that come into play during a nuclear explosion. This instrument is available to the scientific community and will also be used for astronomy and planetology applications.

glass

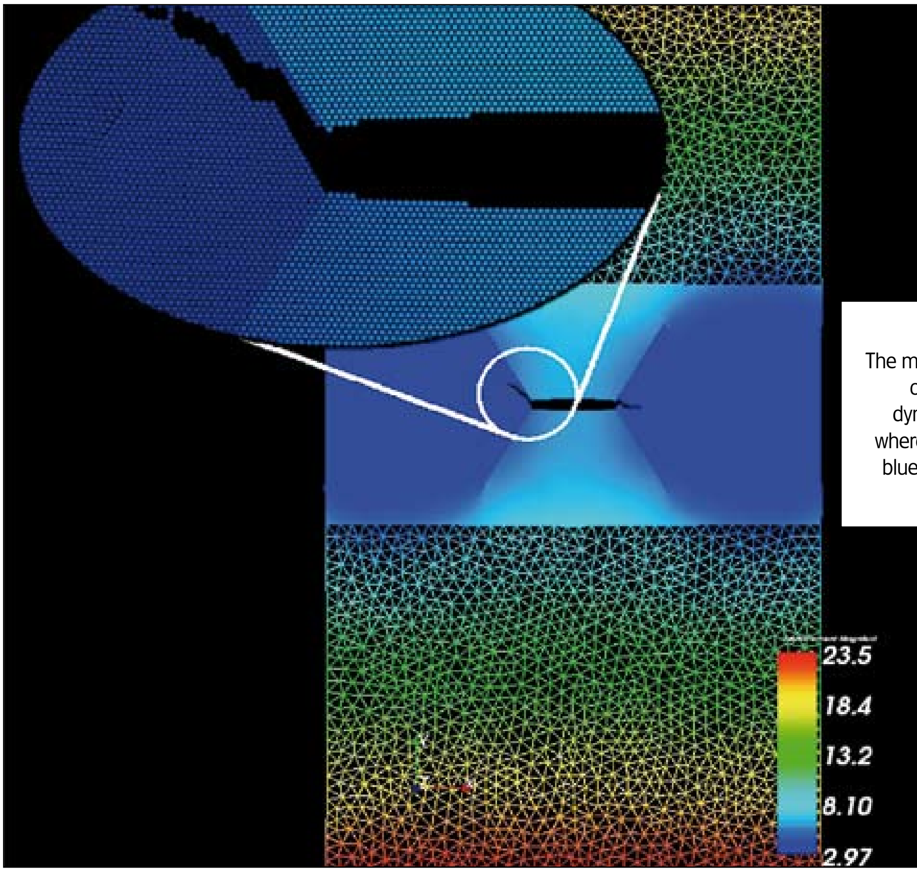

© CEA/Inria

How does glass tear?
The modelling of fissures in the optical components of the megajoule laser using a molecular dynamics model near to the fissure, the zone where the matter is processed microscopically (in blue). Elsewhere, the glass is examined with the 'continuum mechanics' method.

from the full power of a parallel supercomputer. Jean Roman and his colleagues opted for the use of two computing codes (one per model type), which, taken individually, are already 'parallelised'. Their task consisted therefore of combining these codes, in particular ensuring that the computing power available was evenly distributed between both codes. They naturally prioritised molecular dynamics, which is much slower than continuum mechanics, since it is computing on a much more precise scale. 2008 was a successful year: the new simulation was four times faster than the one based solely on molecular dynamics. Thanks to these performances, the CEA teams were then able to make a more exhaustive study of fissure propagation. This done, they succeeded in qualitatively reproducing the experimental observations for the first time: the progress in 'leaps' of the fissure, its speed of progression and above all the presence of dislocation lines (a disordered array in an atomic arrangement) along the length of the fissure. The aim of this exploratory research - to improve our understanding of the mechanics of fissuring - has now been achieved. The exact origin of the fissures, however, remains a mystery. The only hypothesis put forward is that miniscule defects are caused by polishing the glass'surfaces.

So what about the code developed by Inria? The end of the partnership between the CEA and Inria in 2008 did not mean that the project was consigned to a filing cabinet. *"We developed a generic 'coupler' (the program that combined the models), because we were intending to use it later"*, explains Olivier Coulaud. During his post-doctorate in Switzerland, Guillaume Anciaux used it to study contacts between two surfaces on a nanometric scale. To share the fruits of their work with the scientific community, the Inria researchers have now made the 'coupler' freely available as an open source download. Teams of computer scientists and physicians from around the world have already downloaded it. Although the mysterious fissures are restricted to optical components, the offshoots from this work continue to spread.

Air quality under surveillance

Hundreds of pollutants, gases and aerosols react with one another and with the atmosphere: forecasting their concentrations on a daily basis is a real headache.

Dominique Chouchan, scientific journalist

Every day, the French and European Air quality forecasting system (PREV'AIR[1]) draws up forecasting maps for the major pollutants in the atmosphere: ozone, nitrogen dioxide and fine particles amongst others. Due to the sheer number of parameters and mechanisms to be taken into account, both chemical and meteorological, but also the non-linear nature of certain processes, these simulations are sometimes beset by large degrees of uncertainty. To reduce these uncertainties, the strategy adopted by the Clime team at Inria consists of basing their work on a clever combination of models.

Simulations on a European scale

Due to the differences between the life cycles of pollutants, including the diversity of the spatial dispersion scale in question (global, continental, local), simulation of their concentration levels must be carried out at least on a European scale, even if we are only focussing on France. Furthermore, as Vivien Mallet, an Inria researcher working with the Clime team explains: *"We have to model air quality based on the whole atmospheric boundary layer[2], because pollution near ground level (a few metres) depends largely upon what is happening in the first few kilometres of the atmosphere."* The aim is to produce reliable forecasts with a range of 2 to 3 days. The challenge is therefore to develop models which combine meteorological, chemical and physico-chemical phenomena. So what are the origins of these uncertainties? They have many different sources. The physico-chemical processes are still only partially understood. We do not really know, for example, how aerosols are created (fine liquid or solid particles). The turbulent phenomena are difficult to factor in. Furthermore, numerical models have to be simplified to a certain extent, in order to limit computing time. The available data, whether relating to meteorological phenomena, pollutant emissions or their residues on the ground, is largely inaccurate. *"When you consider this, it is actually astonishing that we manage to produce accurate forecasts!"*, notes one researcher.

This was the idea behind using a large quantity of simulations, each distinguished by the way in which the different mechanisms are configured: *"From one model to another, different descriptions of the ozone cycle and/or different representations of turbulence can be used, or even a series of emissions from different chemical species"*, he explains. Ultimately, a hundred or so models are generated this way, from the same numerical platform. The next stage consists of weighing up each one's level of importance. A new model is then calculated based on data from the previous hundred models. The skill lies in selecting the appropriate weights. They are assessed using statistical learning algorithms based on forecasts produced in the past and their observations, all this on a daily basis. The result: *""Not only have we achieved better forecasts than with the best models, they are also better than those produced with an optimum combination of linear models, their weights are also constant over time"*, explains Vivien Mallet. A prototype of the system which calculates a weighted, linear combination of a few models, is now already in place on the PREVE'AIR platform.

Our sights are set on providing reliable forecasts within a two to three-day range.

(1) The PREV'AIR system, created in 2003 by the Ministry of Ecology, was set up by the National Institute for the Environment and Risk.

(2) The atmospheric boundary layer is the zone of the atmosphere directly influenced by the Earth's surface (land and ocean). Its thickness varies between a few hundred metres and a few kilometres.

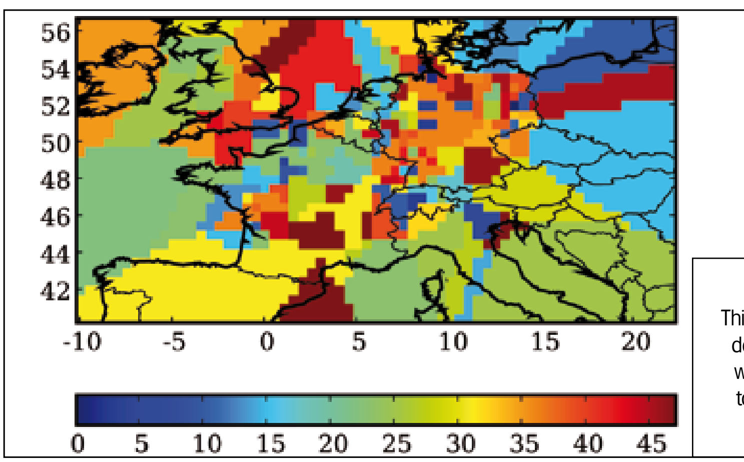

What's the best model?

This map shows the difference between prediction and reality for several dozen models. Each colour represents the best model, in comparison with measurements taken on the 7th May 2001 at the station closest to each zone in question. Model 42 (in red) emerges as the best in a section covering North Europe and England.

Simulating nuclear fusion

One of the biggest challenges in nuclear fusion is to succeed in maintaining this reaction, thereby controlling the mechanisms at work in the reactor's core.

Dominique Chouchan,
scientific journalist

The energy created by the fusion between atom nuclei is vast. On paper, this source of energy would appear to be one of the potential future alternatives to fossil fuels. However, its feasibility remains to be proven. The energy needed to maintain the reaction of nuclear fusion, even of light atom nuclei such as deuterium or tritium, two hydrogen isotopes, is still enormous: the fuel must be brought to temperatures of several tens of millions of degrees and, of course, energy loss must be avoided. Numerical simulation is essential for preparing and using experiments and, more generally, for optimizing the way in which the reactors operate.

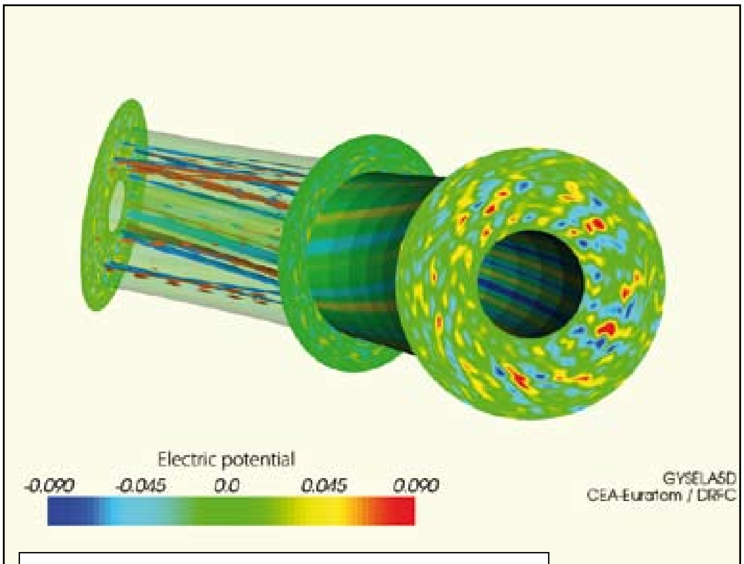

© CEA

Turbulence in tokamak plasma. Cross sections of these surfaces (as in foreground) showing the vortex

Complementary approaches

To produce a maximum number of thermonuclear reactions, the plasma[1] must be confined and kept away from the confinement chamber walls. For a reactor such as ITER, under construction in Cadarache (France), this confinement is provided by intense magnetic fields. Among the mechanisms to be modelled are plasma equilibrium, stability, turbulence effects, plasma heating and interactions between plasma, walls, and materials.

The modelling is based on two complementary approaches, explains Xavier Garbet, research director at the Magnetic Fusion Research Institute (IRFM) at the CEA in Cadarache, the so-called 'first principles' approach and the integrated approach: *"The first consists of solving physics equations describing the phenomena studied. But the simulation of an experiment in its entirety using this method would be far too expensive, it would cost in the region of year's worth of computing on a supercomputer! The second uses simplified models, and the 'first principles' simulations to constrain them."*

Studying turbulence in plasma: one of the greatest challenges to nuclear physics

So where do these energy losses originate? They are linked in particular to turbulence, which affects the way that heat particles are transported. We aim to avoid this in the hottest zones of the plasma, adds Xavier Garbet: *"Closer to the walls, on the other hand, it enables heat to be distributed more evenly."* The models have greatly improved over the past twenty years. However, although the thermal transport coefficients are now estimated with an accuracy of 20-30%, scientists are aiming for 10%, regardless of the plasma conditions.

Further instabilities include those of magnetohydrodynamic origin. On a larger scale than turbulence (tens of centimetres as opposed to a few millimetres), they are likely to cause the reactor to grind to a halt. *"They result particularly from pressure differences and currents running from the centres to the walls, explains Sibylle Günter, director of the Max Planck plasma physics Institute (Garching, Germany); modelling them requires enormous computing power, very efficient algorithms and calculation parallelization over a large number of processors, up to a hundred thousand."*

It is easy to see why the supercomputer installed in May 2009 at the Jülich computing centre (Germany) is useful. Designed by Bull, the HPC-FF (standing for High-Performance Computing For Fusion) is designed to validate numerical simulations as part of the European Union Fusion project; its power is equivalent to 100 Teraflops, or the capacity to perform 100 billion operations per second!

(1) Plasma is the mainly neutral mix of atom constituents (nucleus, electrons) which separate out under the effect of very high temperatures.

Reducing the uncertainties

Have we already reached the point of no return? High-performance computing will play a vital role in finding this out. This is because our current forecasts are limited by the inaccuracies in our description of man's impact on the climate. The next IPCC(1) report will include results produced by the most recent supercomputers. A crucial issue in terms of making the most suitable political decisions.

Renaud Persiaux,
scientific journalist

"In the field of climate modelling, we are still working at the very limits of supercomputer power", states Sylvie Joussaume, coordinator of the European ENES[2] and CNRS (French National Centre for Scientific Research) network and researcher at the IPSL (Institut Pierre Simon Laplace) in Paris, whose 80-person team has developed one of the 23 models used in the IPCC reports.
"We have been working on climate modelling for over 30 years, she continues. Initially with models that are limited to the atmosphere, because of low computing capacities, then ones with increasingly high performance. Developments in computers have pushed back the boundaries of potential areas for research – the models used for the 4th IPCC report were therefore all atmosphere-ocean coupled models."
Climate modelling is one of the most heavily consuming domains in terms of computing time. Because of this, accuracy of results relating to global warming is for the most part limited by the available computing resources. Doubts will remain until we increase the realism and complexity of models.
This accuracy comes at a price that translates first and foremost into computing time.
In 2006, a report[3] identified four challenges requiring high-performance computing in the field of climate change: achieving a very high resolution for quantifying extreme weather events forecasts and assessing social and economic impacts on a regional level, moving towards more comprehensive models of the Earth System, quantifying uncertainties, and lastly, investigating the possibility of unexpected climate-related events.
These needs relate to four factors, leading to an exponential number of computations: *"the model's complexity, its resolution, the time-span of the phenomenon being modelled and the number of simulations necessary"* according to Sylvie Joussaume.

Doubts will remain until we increase the realism and complexity of models

(1) IPCC - Intergovernmental Panel on Climate Change.

(2) European Network for Earth System Modelling is a European network including the main climate modelling centres. The IS-ENES (ENES Infrastructure) project aims to strengthen the ENES community's organisation at a European level.

(3) Scientific Case for European Petascale Computing (2006), p 35-36.

(4) Intensive computing for the climate and the environment (conducted from January 2008 to June 2009).

Adapting to the massively parallel

For several years, the climate community has been conducting major research on adapting programming codes to massively parallel machines. *"Our current models are essentially designed to run on standard parallel machines"*, explains Olivier Marti, researcher at the IPSL (Institut Pierre-Simone Laplace). *"Thanks to the CICLE[4] project that he directed, IPSL's 'Earth System' model now exists in a massively parallel version. Major work had to be done on the coupler, which was running on a single processor and which was causing bottlenecks"*. Once we have overcome these problem areas, climatology will then be capable of using any high-performance machines and of adapting to the profound evolutions in hardware that we expect to see.

Ranking first in the Top 500 from 2001 to 2004, the Earth Simulator is a massively parallel vector machine that is partially dedicated to climate modelling. Although it is now obsolete, this machine pioneered the world's first climate simulation at an unprecedented resolution of 3 km.

# of climate change

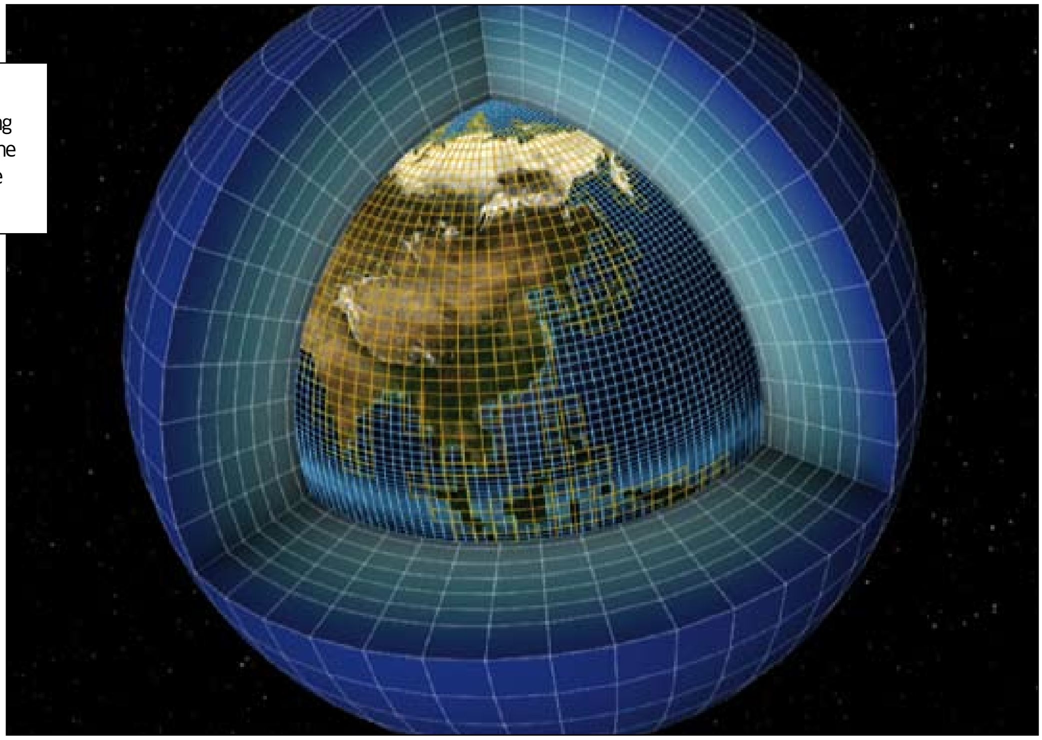

Cutaway view of the climate model, revealing the atmospheric grid, the ocean model and the ground model.

The first factor to be taken into account here is the model's complexity, i.e. the number of processes that it integrates. Although the models used for the 4th IPCC report were all atmosphere-ocean coupled models, researchers agree that it is time to move towards global models of the 'Earth System'. Such models would encompass all known processes affecting the climate. For the past 30 years, models have gradually begun to take the following factors into account: wind evolution, temperatures, radiation, clouds, topography, soil humidity, oceans and marine currents, sea ice, biogeochemical cycles and in particular vegetation, atmospheric pollutants and many others, including changes brought about by human activities. When considered as a whole, this gives us a vast quantity of calculations to be solved in the models grid. For example, integrating biochemical cycles increases the number of calculations by a factor of between 5 and 10.

The representation of the process and interactions between sub-systems must also be refined (such as water and heat flows at the point where the atmosphere meets the ocean). In fact, the way in which they are modelled influences the model sensitivity, a phenomenon which researchers have coined the 'climate sensitivity factor': the various models, when subjected to double the amounts of CO_2, simulate a global warming scenario that varies between 1.5° and 4.5°. These different responses mainly stem from cloud modelling", adds Sylvie Joussaume. Moreover, we now know that an increase in temperature changes the absorption capacity of the carbon sinks - the vegetation and oceans. In simple terms, the hotter it gets, the less CO_2 is likely to be absorbed. This mechanism is at the root of snowballing phenomena which are somewhat complex to model, and whose significance still needs to be assessed. The response from vegetation in particular is still poorly understood and therefore very rarely taken into account by climate simulation models.

> **"On an international level, teams equipped with the best computing resources will be working at 100 km in the atmosphere!"**

Optimising spatial resolution

The second factor is resolution. The terrestrial globe is discretized into a grid, with cells of varying sizes. At each point in the grid and at each time step, from tens to hundreds of equations need to be solved, depending on the model's complexity. *"For the last IPCC report, the atmospheric grid was made up of over 130,000 points(5), for a mesh of around 300 km. Today, the grid has almost 400,000 points(6), for a mesh of around 200 km, thereby increasing the number of calculations by roughly a factor of 6. On an international level, teams*

(5) Precisely 131,328 points: 96 points in longitude x 72 points in latitude x 19 altitude levels.

(6) Precisely 393,984 points: 144 x 144 x 19.

equipped with the best computing resources will be working at 100 km in the atmosphere! We are also examining the benefits of increasing the number of vertical levels", adds Olivier Marti, researcher at the IPSL. It is crucial that we increase spatial resolution further still to improve process representations, some of which take place at a very small scale, such as soil humidity, clouds, vegetation, atmosphere pollutant concentrations, as has been demonstrated in regional models in the order of 10 to 20 km (see "three questions for José Baldasano", on page 31). *"Increasing the resolution to less than a kilometre for regional models and less than 20 km for models on a planetary scale would increase the computing power needed by a factor of a thousand, compared with today's needs"*, explains Olivier Marti.

Accuracy for 'Earth System'

The third and fourth factors are the length and number of simulations. *"Since the calculations are made with a three-minute time step, this produces over 175,000 time steps for a year, with simulations spanning several centuries"*, explains Sylvie Joussaume. But a single simulation experiment is not enough because of the dependency on initial conditions or parameter values entered into the model, which influence simulation results: in particular on potential extreme weather events (see box below) or the risks of ruptures causing real climate disasters, such as a brutal change of balance (for example, ceasing of the Gulf Stream's Atlantic flow) or exceeding a point of no return beyond which an irreversible snowballing effect would occur. Furthermore, the different potential scenarios need to be simulated (around forty in the IPCC, depending on emission reduction policies, which are more or less voluntary). It's fair to say that there are many unknowns concerning greenhouse gases and aerosols, possibly even more so concerning the scale of global warming and its impacts.

"175,000 time steps for a year, with simulations spanning several centuries…"

When combined, all these factors bring about a combinatorial explosion in the order of 106 at least. Each power increase brings us forward on one of the three factors. *"For the 4th IPCC report, a total of 80,000 compu-*

(7) *Geophysical Research letters*, **vol. 35, 2008.**

(8) **Series of Simulations of Extreme Weather Events under Nonlinear Climate Change.**

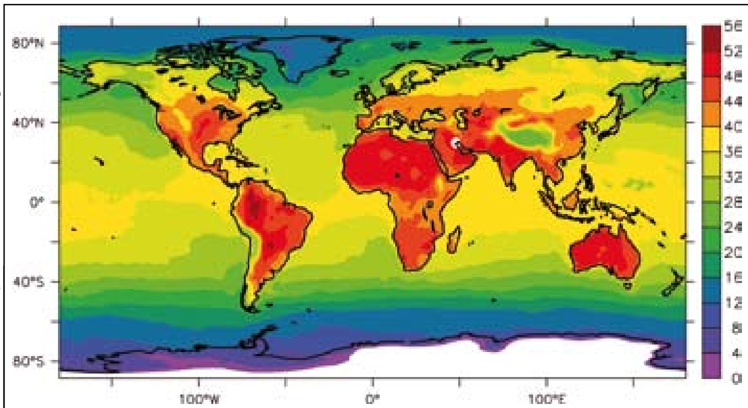

Simulation of extreme temperatures (in °C) forecast by the Henk Dijkstra team (University of Utrecht) for the decade 2090-2100.

Extreme weather events: multiplying models to for more accurate forecasts

Temperature peaks will increase two-fold faster than average temperatures and will cause increasingly frequent heat waves. With, at the end of the century, extreme temperatures exceeding 40 °C and even 50 °C in a very high number of regions. These findings were published last year[7] and have hit like a bomb shell. To estimate these extreme events, Henk Dijkstra (University of Utrecht, Netherlands) and his colleagues have run 17 models of the ESSENCE project[8], developed by the Max Planck Institute and used for the 4th IPCC report.

"To check the robustness of findings, we introduced disturbances into the initial conditions, and then took the average. With 17 different models, each as accurate as the next, we are more likely to see extreme weather events occurring", explains Henk Dijkstra. But 17 models mean 17 simulations and therefore 17 times more computing time. This is why very powerful supercomputers would be highly beneficial. *"Without high-performance computing, we wouldn't have been able to do any of this"*, he asserts.

© IPSL & Animea, Frédéric Durillon

A more or less clear atmosphere
Clouds simulated for two resolutions of the IPSL climate model. On the left: 360 points in longitude x 180 points in latitude x 55 vertical levels, and on the right, 96 x 72 x 19. The highest resolution is still difficult to use routinely with the computers available.

ting hours were needed (the equivalent of just over nine years if we had used a single processor), performed between 2003 and 2004, producing 40 Terabytes (10^{12} bytes) of data. At the time, six NEC SX6s processors were used", recalls Sylvie Joussaume. For the 5th IPCC report, scheduled for publication in 2013 and for which the simulations have to be finished in 2010, climatologists will have 48 SX9 processors, six times faster, so fifty times more powerful than when the first IPCC report was published. With increasing computing capacities available, the increasing resolution and simulation of the Earth System will be integrated into the 5th IPCC report.

Currently, the coupled IPSL model, for example, includes five aspects of the climate integrated in the Earth System (atmosphere, ocean and oceanic biogeochemistry, sea ice, continental surfaces and vegetation, atmospheric chemistry) interacting by means of a 'coupler'. *"Available in different configurations and different resolutions depending on the level of complexity in hand, it is perpetually evolving, thereby reflecting the international state of play. Our teams are conducting the last tests before launching the simulations destined to the IPCC"*, concludes Sylvie Joussaume.

So we can see that answering society's questions on global warming is wholly dependent on the accessibility of high-performance computing resources. To reduce doubt, quantify the probability of extreme weather events, count the carbon sinks and their possible evolution following the climate change, to study impacts on ecosystems, and so forth - all this will require an increase in computing power of a 10 – 10^6 factor compared with current resources. The world climate community is already preparing for the change in scale and is adapting its models, no longer only to Tera-scale computing[9], but to Peta-scale and even Exa-scale computing.

"Three questions for"

José Baldasano, director of the Earth Sciences Department at the Barcelona Supercomputing Centre

What exactly are you working on?

Our work involves examining the links between the climate, the circulation of atmospheric pollutants and attempts to control them by means of regulations.
The atmosphere is a complex reactor. Its many composites react dynamically to physical conditions. However the regulations have not yet taken stock of the long-term impact of climate change on air quality. The problem lies in the fact that the low resolution of current global simulations between the climate and chemical products have not yet authorised regional estimations.

What method have you been using for this?

We used very high-resolution regional models of the Mediterranean zone (5,000 by 2,500 km), with a 20 km resolution, and a vertical resolution of 31 layers of the troposphere, installed on our supercomputer Mare Nostrum. We examined the responses of ozone and particle concentrations to climate change, by drawing a comparison between the months of August 1960, 1980 and 2000, and two conditions for 2030, using scenarios where emission levels are maintained as well as scenarios with reduced emission levels.

What are the main findings?

The complex topography of the Mediterranean region engenders distinct regional features and differentiated pollutant behaviours across the western, central and eastern basins. Evolutions in ozone concentration are highly variable depending on locations, ranging from −20 to +70 $\mu g/m^3$. If we reduce emissions, the peaks will drop for the whole sector, particularly around Sicily where they were very high in 2000. If we don't clean up our act, the pollutant concentration levels will increase in the majority of regions, due to self-magnification phenomena.

(9) **Referring to supercomputer processing capacities of 10^{12}, 10^{15} and 10^{18} operations per second respectively.**

Treating water using computers

One fifth of the human race still doesn't have access to drinking water. A team of Finnish researchers uses high-performance computing to optimise the coagulation of pollutants in water and to facilitate their elimination.

Renaud Persiaux,
scientific journalist

Access to drinking water will be one of the greatest challenges of the 21st century. Numerous heavy metals - such as cadmium, mercury and lead - and highly-toxic pollutants are frequently used by, or produced by, the petrochemical, textiles and paper industries. Their processes require vast quantities of water, and current methods used for treating this water before it is discharged into the environment are still often too burdensome, too expensive and energy-consuming.

"It is vital that we have a means of purifying water, in the most energy-efficient manner possible", stresses Prof. Kari Laasonen from the University of Oulu Department of Chemistry. He has long been convinced that this can only be possible by modelling molecular events occurring in the water, notably to better exploit its extraordinary solubility properties. *"We use numerous chemical composites empirically to eliminate pollutants that have dissolved in water, but we know virtually nothing of their structure at an atomic level, or of the mechanisms at work."* What would be the optimum solutions for pollution removal? "*This question can't be answered with simple test samples*", he explains, *"given the scale, the complexity and the speed of mechanisms, as well as the number of molecules to be tested!"* He works *in silico* with his team, introducing known properties of existing and potential de-polluting chemical composites (notably using the Cambridge Crystallographic Database) into highly-realistic models of the physico-chemical properties of water (COSMO - Conductor-like screening model). *"To determine the precise relations between the properties of water, the composition of de-polluting or potentially de-polluting materials and their conditions of use, we first need to produce analyses, at multiple scales, and link them all together. Our simulations are highly demanding in terms of computing"*, continues Kari Laasonen, whose team is one of the biggest users of CSC's (Finnish IT Center for Science) supercomputers.

"Purifying water with optimum energy-efficiency"

After ten years of research work, his team - working in collaboration with the multinational Kemira, international leader of the coagulants sector - is at the cutting edge. *"With modelling techniques, we are now able to study almost all metal-ligand[(1)] complexes that are relevant to our work. We can therefore test prototypes of molecules before their production. This helps developers to better understand their products and points them in the right direction for making new developments."*

(1) A ligand is a molecule capable of binding to atoms or ions.

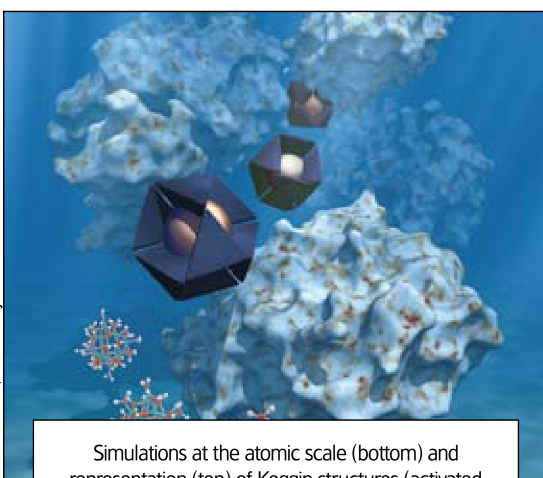

Simulations at the atomic scale (bottom) and representation (top) of Keggin structures (activated alumina), amidst dirty particles (in light blue).

Favouring pollution control

Kari Laasonen's team (University of Oulu, Finland) has been examining the behaviour of activated alumina, a compound used for removing pollutants from water. One of the forms that it can take is the 'Keggin' structure: a highly symmetrical positively-charged spherical assembly, with the property of attracting negatively-charged impurities and thus forming aggregates that are easy to filter. This is but one of its many possible conformations in water (see photo): *"the spectrometers reveal the existence of dozens, or even hundreds of relatively-stable alumina structures"*, explains Kari Laasonen. The team uses high-performance computing to identify all structures that are relevant to pollutant removal, and to understand and foster their formation. For the time being, it is still impossible to model them all, but he hopes that *"with lots of work and by increasing computing power, we might just achieve this"*.

Looking into the depths of ocean circulation

Climatic forecasting demands a minute knowledge of ocean processes. At equal computing power, clever algorithms can be used to display invisible details with a large-scale model.

Dominique Chouchan,
scientific journalist

How can the behaviour of the global ocean, and more specifically its response to interactions with the atmosphere (atmospheric forcing), be simulated with sufficient precision? This is the challenge that general circulation models, such as the European model NEMO[(1)] are faced with. This challenge demands multi-scale modelling, since the circulation is highly influenced by turbulent processes of local nature that are small in size (in the order of tenths of kilometres, or even less) but highly energetic. A collaboration between oceanographers and mathematicians has brought about the design of techniques by which given zones can be 'zoomed' in on without actually having to model the whole ocean on a minute scale.

Amongst the processes which impact the general circulation are inter-basin exchanges, explains Bernard Barnier, head of the ocean modelling team for the Laboratory of Geophysical and Industrial Fluid Flows (LEGI, Grenoble – France): *"This is typically the case for exchanges between the Arctic ocean and the North Atlantic via the Denmark strait. If we don't model the dynamics minutely, we won't be able to produce a high-resolution representation of thermohaline circulation[(2)] nor its potential sensitivity to climate changes."* Other local processes are occurring in zones where surface waters meet with deep waters, for example, in the Greenland sea or in the Labrador sea (see figure).

Such situations have prompted work on solving dynamic equations on a very small scale, by adapting the configuration of models. This is where the mathematician steps in: *"The power of computers is insufficient for modelling a whole ocean at the level of detail required in certain zones"*, explains Eric Blayo, in charge of an Inria applied mathematics for environmental science project (MOISE, Grenoble). However, sectioning off an area is out of the question, since each one is wholly dependent on its environment: *"We need to be capable of 'plugging' a very high-resolution local model into a model of the global ocean"*, he added. This methodology and the associated software have been developed by his team. What is it exactly? The strategy is based on local multigrid methods, also known as multiresolution algorithms. They can be used to solve equations relating to oceanic dynamics on increasingly fine grids (several levels of discretization) ensuring the continuity of flows at interfaces between grids and minimising uncertainties. *"The small-scale results then allow us to determine the configurations essential for the lowest-resolution models"*, highlights Bernard Barnier. Used in the majority of major oceanographic models, this software is now even an integral part of NEMO. However its potential fields of application reach far beyond this single framework, due to its generic character. One obvious possible application would be for atmosphere analysis. *"Astrophysicists are also interested in this"*, adds Eric Blayo, *"for increasing their model resolutions in certain areas of space"*.

Computing power is only exploited where it is absolutely necessary

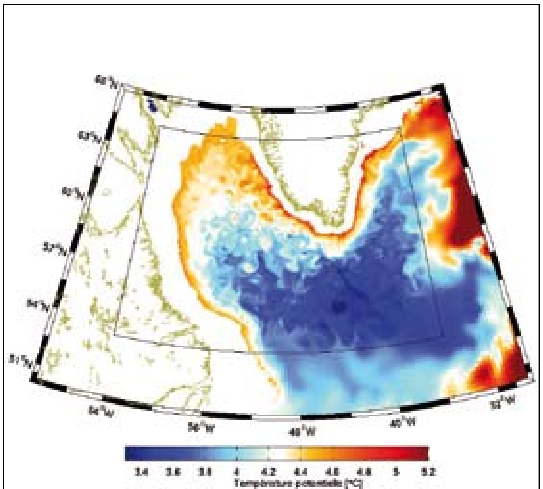

© Jérôme Chanut/Mercator Océan

On this simulation of the instant oceanic temperature in the Labrador sea, at a depth of 182 m, the local resolution (4 km) is five times more accurate than in the rest of the Atlantic for a model such as NEMO.

(1) NEMO, *Nucleus for European Modelling of the Ocean*, provides a framework for numerically modelling the ocean. It enables simulation of ocean dynamics, sea ice and biogeochemical processes. Managed by the INSU (National Institute of Universal Sciences), its development now takes place within a consortium bringing together different European research centres.

(2) Thermohaline circulation describes the large-scale circulation of oceanic water masses under the effect of temperature variations and salinity.

An essential springboard for tomorrow's growth

by **Didier Lamouche,**
Chief Executive Officer
of the Bull Group

© DR/Bull

Today's fault lines – globalization, revolutions in digital technology, nanotechnologies, biotechnologies and the environmental challenge –will become the main driving forces in the world that emerges from this recession. Innovation will be the key to success in this new world order. High-performance computing (HPC) enables the most complex equations to be solved and highly sophisticated models to be analyzed; opening up new perspectives for companies, reaching out to all sectors, from healthcare to energy, agronomy to finance, transport to construction. HPC is proving to be an essential tool… provided that the three major obstacles to its development are overcome.

To begin with, this is a technological and industrial challenge. In the world of HPC, the appetite for computing power cannot be met by ever-larger architectures. Major companies' demands are evolving, and technology suppliers are constantly running to catch up.

Creating a new European ecosystem

It's quite simple – the current performance of supercomputers is still not enough, despite the fact that demand for computing power is growing constantly, in line with the market. HPC is strategic when it comes to growth, innovation, even State sovereignty; and the Europe 'arm' of the global system cannot just rely on external know-how, it needs to actively assert its own technological expertise and come up with step-change solutions. Now is the right time to create a truly European ecosystem, bringing together suppliers of computing technologies, customers and users. The high-performance computing revolution needs to be a shared challenge, even more so in a 'systematized' global economy.

Then there is the financial and political challenge. One of the major differences between the USA and Europe is that although European funding may be sufficient to equip Data Centers, it is still not enough to fund technological developments themselves. The European approach differs here from the American approach, in that the European Union's programs do not aim to fund research and development for industry at a European level, but simply to equip researchers with the computing power that they need. This is a critical situation, which could eventually destabilize the competitive equilibrium.

> **"Three major obstacles to overcome"**

The final challenge is one of accessibility. The extraordinary potential of supercomputers may be available to major corporations, but it also needs to be accessible to smaller businesses and industries, with pre-integrated, high-performance turnkey solutions. But this technological revolution has to involve the SMEs themselves, in order to reach out to all sectors of the economy.

After all, HPC is an extraordinary tool when it comes to anticipating the future, offering many spectacular applications that will have a tangible impact on the future lives of our fellow citizens. By supporting the democratization and widespread use of HPC, the public powers in Europe could create a collective momentum in research and innovation. We are poised to take on the challenge; to enable France and the rest of Europe to cross a major technological divide, which will guarantee our industrial competitive position.

Petroleum, supercom

Xavier Muller,
scientific journalist

What's more risky than the casino? Oil prospecting. When they drill into a seabed that is likely to contain petroleum, Total engineers might discover a deposit of several million barrels of crude… or nothing at all. When you consider that exploratory drilling at sea costs between 60 and 80 euros, you can understand that when it comes to prospecting, the petroleum giant prefers not to leave things to chance. In the mid-90's, only 40% of deposits fulfilled their promises. Numerical simulations that analyse data obtained by seismic echography have radically changed the playing field. Armed with the third most powerful French supercomputer after the one at Idris (Institute for Development and Resources in Intensive Scientific Computing) at Orsay and Cines (National Computing Centre for Higher Education) in Montpellier, the Total engineers are now hitting the bull's eye in 60-70% of cases.
But during the long journey that petroleum travels between its extraction and its transformation to a finished product, prospecting is just the first of many tasks that require binary punching power. Refinery, the manufacture of bitumen and rubber seals are others. For these operations, the need for high-performance computing is largely due to the coexistence of phenomena taking place on different scales and mutually influencing one another. In recent years, Total has made massive investments in high-power computing. Proof, were any required, of the important place that supercomputers now occupy in petroleum engineering.

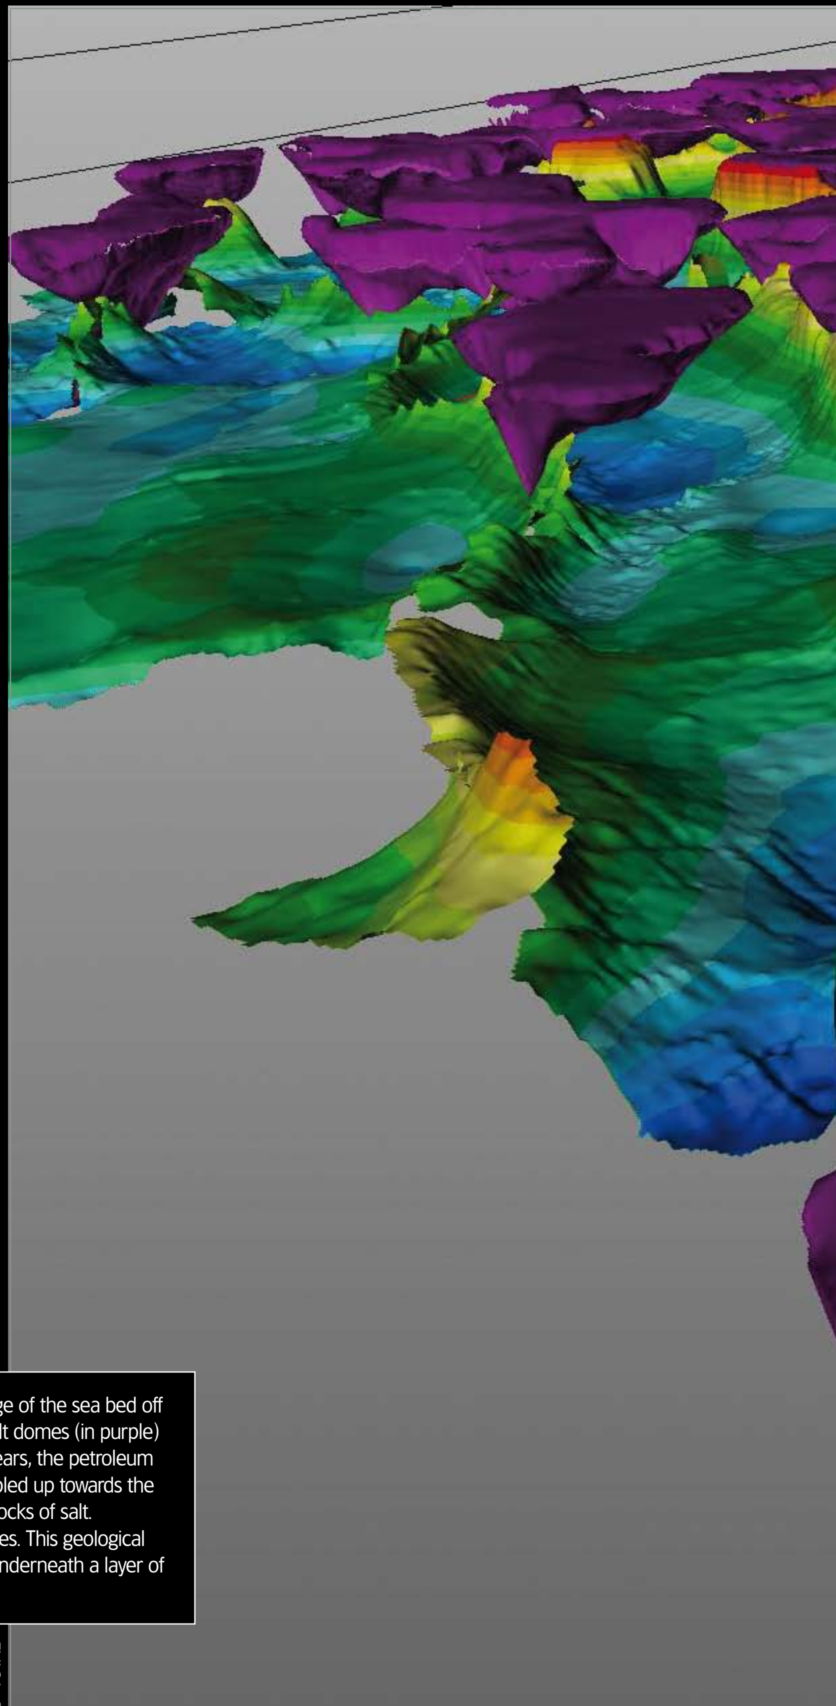

Using their supercomputer, Total engineers reconstructed this image of the sea bed off the coast of West Africa. The zone, covering 4,000 km², shows salt domes (in purple) covering sediment layers (from blue to red). Over thousands of years, the petroleum – which is lighter than water – contained in these sediments, bubbled up towards the surface and found itself trapped beneath impenetrable blocks of salt. Recovery wells were then bored down to the base of these domes. This geological landscape, which stretches to a height of six kilometres, is buried underneath a layer of 200 m of sediment, not shown in the image.

puting at all levels

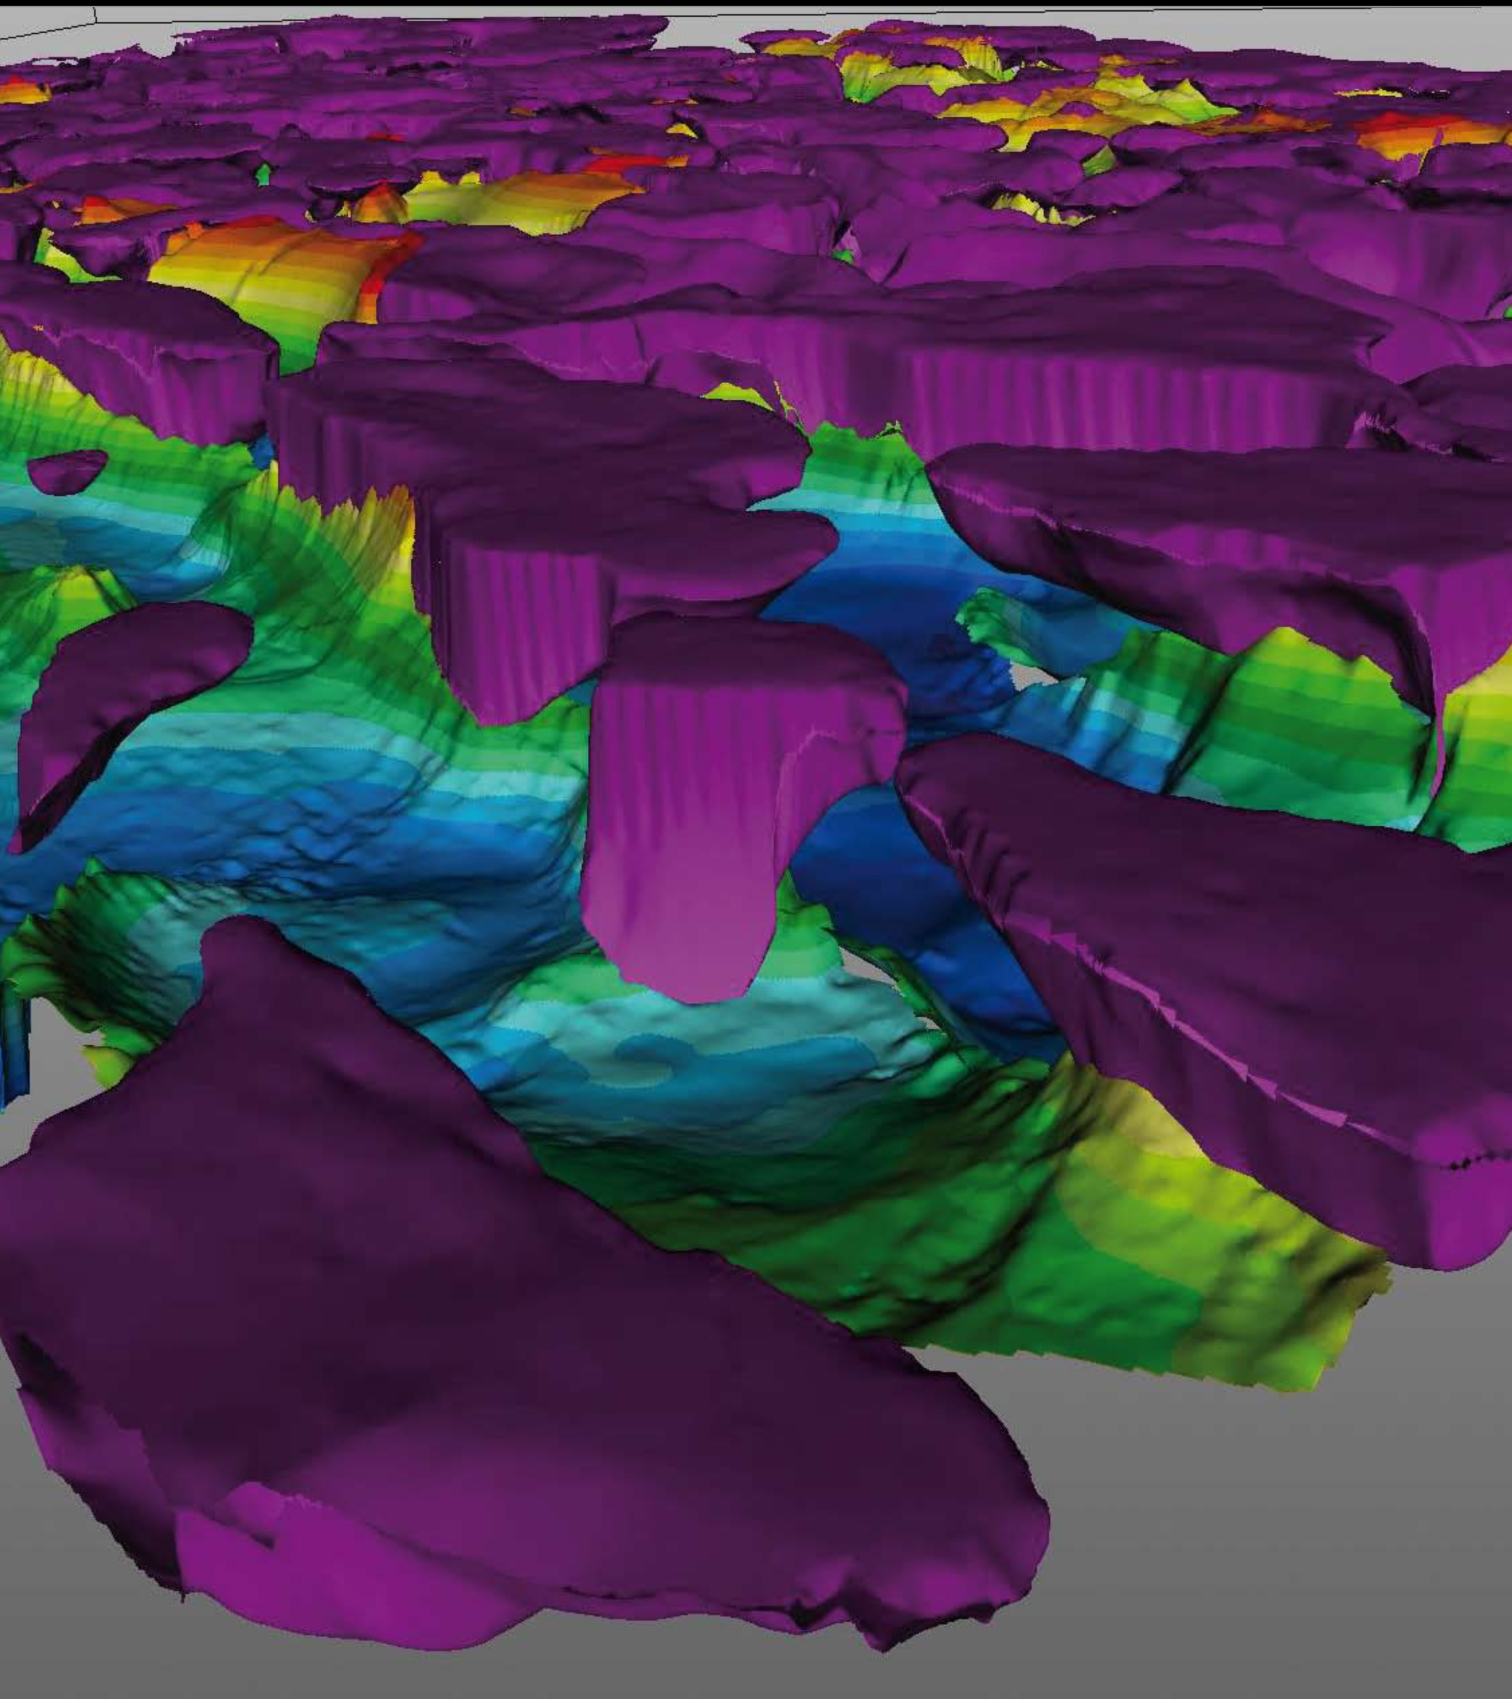

The quest for black gold "Seismic echography at sea off the coast of West Africa"

© TOTAL - Direction de la communication

1

Seismic echography, first stage of petroleum prospecting, is unlike its medical cousin: bursts of sound waves are sent, from a ship, towards the ocean floor. Next, thousands of sensors fixed along cables stretching out behind the ship record the echoes produced by the waves being reflected on the geological layers underneath. The cables, which reach lengths of up to 8 km, are linked to the boat with floaters equipped with small fins which ensure that the cables move correctly through the wash.

© TOTAL

2

Seismic recordings are then processed to obtain rough profiles of the sea-bed. Although the computer is the main actor in this phase, the human eye remains irreplaceable. Only geologists, specializing in the region of the globe concerned (here west Africa), know how to identify the boundaries between the rocky compartments (in orange and red) within this greyish layer-cake. Analysing these, solely from images and their overall knowledge of the region, they will use this as a basis for computer simulations, enabling them to go back and refine the profiles.

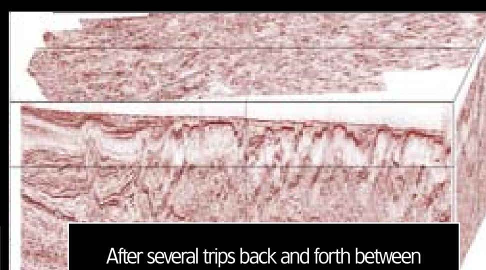

© TOTAL

3

After several trips back and forth between numerical computing specialists and geologists, a final image is produced. The scale of nuances shows the reflection of seismic waves: the darker zones correspond to a strong reflection.
The whitish undulating zone near to the surface is made up of salt. It lies over a pile of furrowed sediment, where oil is likely to be formed. Thanks to high-power computing and following campaigns that lasted from a few weeks to several months, seismic echography now enables scientists to view regions covering up to 7,000 km² over 10km depth, as is the case here, off the coast of Angola.

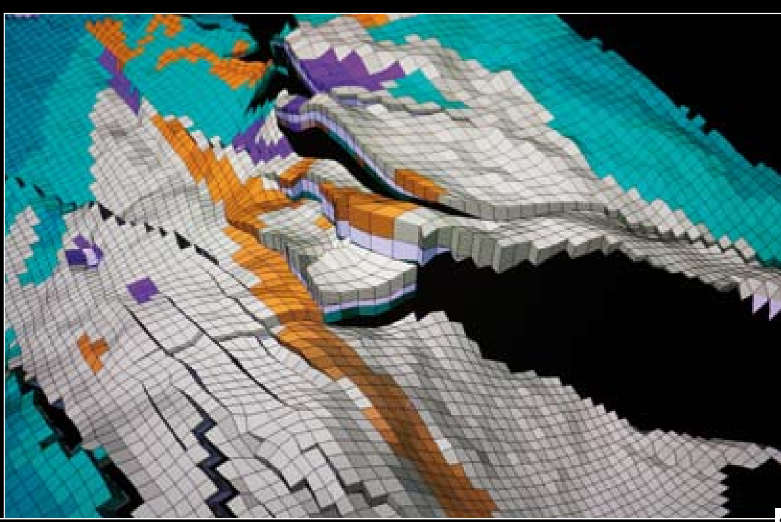

© TOTAL

4

Detail of previous image on a petroleum deposit (the white zone). A deposit occupies an area of a few hundred metres squared. Depending on the case, one or more borehole wells are installed to exploit it. In order to decide where exactly, engineers use numerical simulations reproducing the pumping of the pool of black gold through the rock's pores. They then adjust these simulations according to the viscous properties of the petroleum, established through coring techniques performed from a ship. By using simulations, the quality of the oil extracted increases by several percent compared with a random pumping method in the deposit.

Making seismic echography more reliable

Mathematically, seismic echography boils down to the resolution of an inverse problem of extreme difficulty: the nature and type of geological structures that reproduce as reliably as possible hundreds of thousands of seismic recordings. A mammoth task that Total has performed on its 123 teraflops supercomputer based in Pau.

1

© TOTAL - Direction de la communication

2

Seismic profiles of a region of the Gulf of Mexico. The top image, taken in 2003 using a computer with 64 processors, shows the bottom right-hand side of a structure shaped like a bowler hat, typical of a petroleum zone. Based on this image, Total engineers were ready to install boring equipment on this site. Fresh data analysis using a supercomputer revealed that the structure was an artefact.

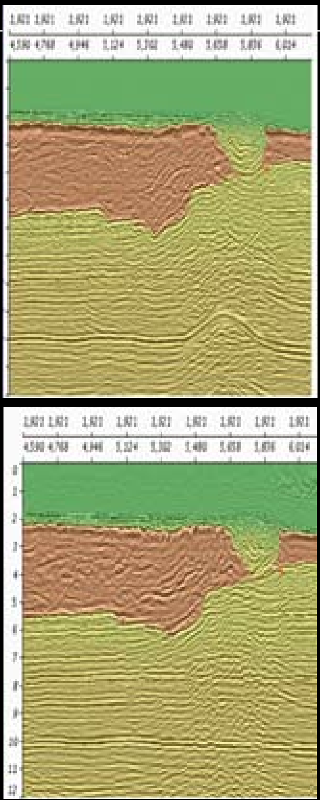

© TOTAL

3

Seismic echography on land requires more computing than at sea, because the surface from which the wave bursts are sent is irregular. Total has joined forces with Inria to develop specific numerical methods to take into account site topography. Validation tests are carried out on false data: researchers manufacture a geological structure from scratch, formed of an assembly of geometric shapes representing rocky structures, with variable speeds of propagation of the sound waves. They calculate the seismic recordings that they would have obtained by placing a string of detectors at ground level, and from these data they attempt to recover the initial geological structure (in red, the ground surface, in blue and green, geological compartments). This concordance between computed and initial structures is proof of the quality of the simulation.

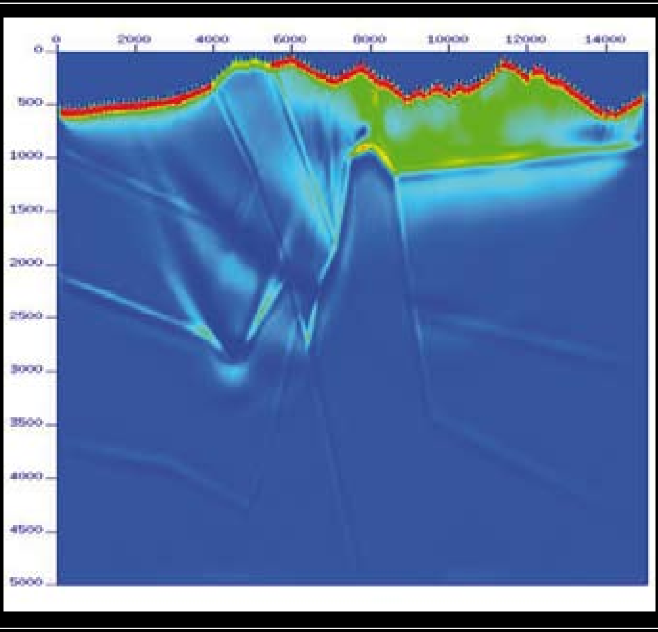

© TOTAL/INRIA - Hélène Barucq

Improving refinery

The Anvers refinery, in Belgium. The gleaming "rocket' which appears poised for launch is the reactor of a catalytic cracker. By way of a catalyser, it 'breaks' heavy hydrocarbons up into lighter elements. The nature of these elements (liquefied gas, diesel or petrol) depends in theory only on thermodynamic conditions (temperature and pressure) in the reactor. Using these simulations, engineers at Total can therefore plan which changes to make to these conditions to make the reactor produce another fuel. This is how they adapt production in refineries to local market changes, especially in the wake of the increased demand for diesel in Europe.

Simulation of the top part of the catalytic cracker. After transforming the light products in the bottom of the cracker, the hydrocarbons, still mixed with the catalyser grains, describe a spiral trajectory. The hydrocarbons leave the cracker through the top, whereas the catalyser grains exit via a pipe (on the left) before being recycled. This simulation was obtained on a supercomputer with a processing power of several tens of teraflops.

A catalytic cracker utilises phenomena which take place on several scales. This is why simulating the way it works requires powerful computers. On a millimetre scale, heavy hydrocarbons hit the catalysers, chemically reacting with them and forming turbulence at the back of the grains (image 1). On a centimetre scale, the catalysers form a suspension in the reactor's atmosphere (image 2). Finally, within the reactor, general mixing of the fluids can be observed (image 3).

image 1

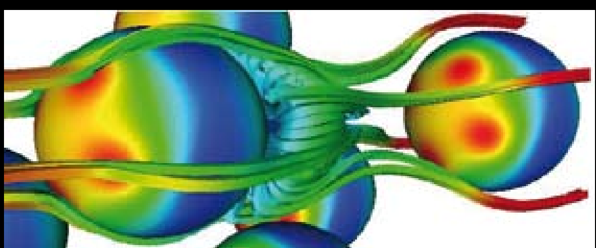

image 2

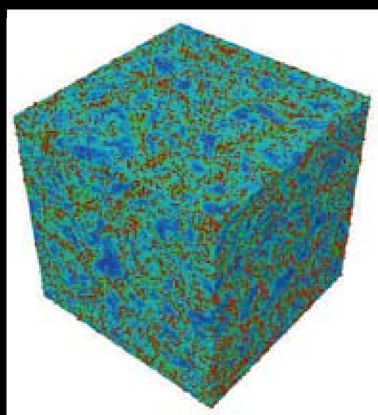

image 3

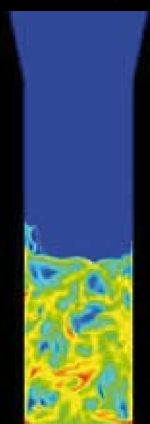

Transforming petroleum

© Hutchinson

Via its subsidiary Hutchinson, Total dedicates a large portion of its computing capacities to developing elastomers, artificial rubbers manufactured using polymers. Numerical simulations are used, for example, to perform virtual fatigue tests on car door joints or dust seals on the Paris metro (on the image of the door, left, the colours represent the different meshes that were used for the numerical representation; on the image of the dust seal, right, the red indicates the very stretched parts of the rubber, the dark blue, the least stretched.) The simulations of the door joints can be carried out on a 7 teraflops supercomputer.

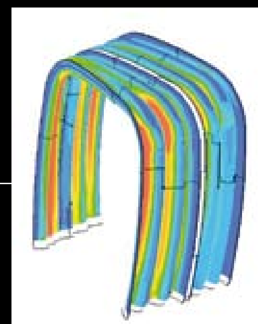

© Hutchinson

On the banks of the Mississippi river in Louisiana, the Total factory at Carville is the world's largest producer of polystyrene. The polystyrene is produced inside a reactor containing 150 tons of fluids, by polymerizing styrene (a benzene hydrocarbon) in the presence of catalysers. Can the reaction's efficiency and the quality of the product be improved through better design of the blades that mix the fluids? Simulations attempt to provide an answer (in green, the high speeds reached by the fluids in contact with the blades).

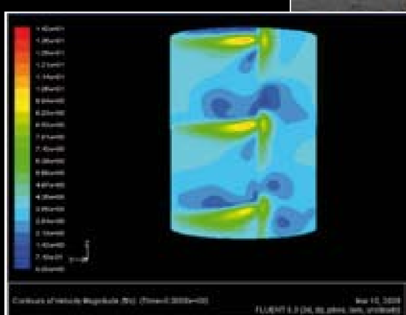

© TOTAL

© TOTAL – Direction de la communication

© TOTAL – Direction de la communication

Porosity test surfacing, a mixture of gravel and bitumen used in road surfacing. By injecting polymers into the bitumen, a road surfacing mix is created on which rain will form fewer puddles: the material becomes porous, which enables the water to run through it. Total uses simulations to gain a better understanding of how polymers, reacting with the carbon grains contained in the bitumen, produce this permeability.

Scientific computing go

Supercomputers are ravenous energy monsters. At this rate, they will have to be located within arm's reach of an electrical power plant. The time has come to curb their bad habits.

Léo Gerat,
science journalist

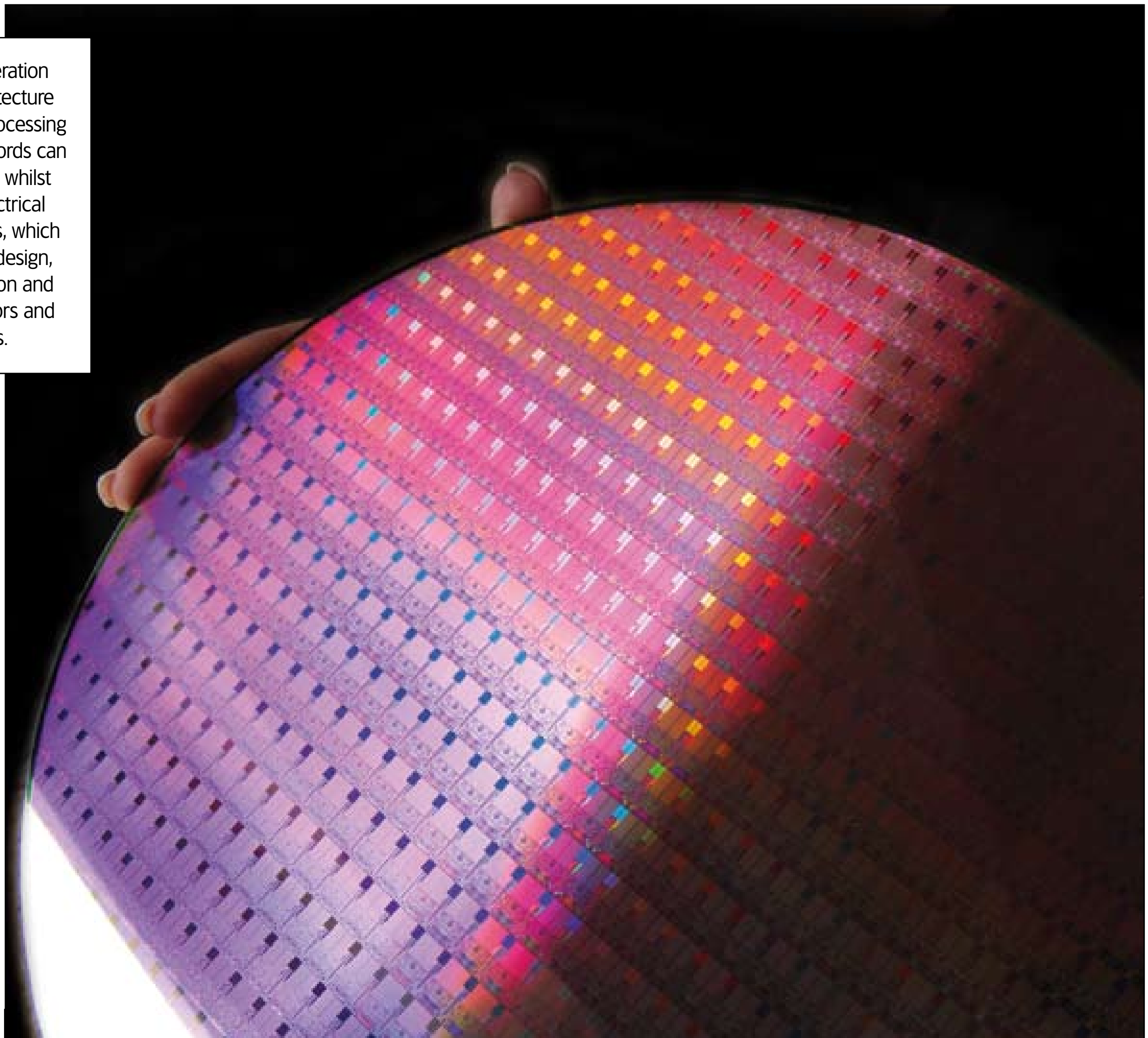

With new-generation Intel microarchitecture (Nehalem), all processing performance records can now be broken, whilst minimising electrical loss in transistors, which may affect the design, size, consumption and cost of processors and computers.

© Intel

Two years back, global computing was diagnosed with bulimia. It was energy-hungry and insatiable. The first alarming figures were published in 2007 by Jonathan Koomey, a professor at Stanford, showing that the planet's computer servers had consumed 0.8% of electricity production in 2005, twice the amount consumed in 2000.

In the same year, the Gartner design and engineering consultancy showed that global computing accounted for around 2% of the total CO_2 emissions produced by humans. That's as much as civil aviation! This is the reason behind the computing industry's growing consciousness over the past two years and the creation of structures such as *Green Grid.* Administrative organizations are starting to react and the European Commission launched its *Code of Conduct for Data Centers* on 30th October 2008. This text is fairly unrestricting and has no legal value, but we can only hope that governmental organisations will start to include these Codes of Conduct in their tenders.

And how about energy issues in the context of high-performance computing (HPC)? It's quite simple – at this pace, the largest supercomputers will be consuming as much as a town and those set to attain a quin-

es green

tillion operations per second (exaflops) around 2020 will need to be located next to electrical power plants. We're talking 50 megawatts (MW), 100 MW, maybe even more. Electricity bills may become prohibitive.
The profligate appetite of computing equipment is a multi-stage rocket. The processors are greedy. Memories and peripheral devices are too. They are often under-used and their electrical power supply systems are energy-inefficient. The energy consumed is essentially transformed into heat, which then needs to be evacuated to avoid the occurrence of detrimental temperature rises, using air-conditioning equipment, another major consumer.
The cure for cutting down will revolve around easy solutions focussed on eliminating identified wastage. Power supplies were showing poor efficiency levels, often transforming a third or a quarter of electricity into heat. Manufacturers have already begun to redress this balance and we can but hope that power supplies for computers yielding less than 90% of what they consume will soon be an exception.

Solutions for uninterrupted computer operations

One aspect of this issue is linked to the need to guarantee uninterrupted operation for numerous computer installations. We have got into the habit of applying easy rules, which are often somewhat simplistic.
It is therefore high time to install UPS (*Uninterruptible Power Supplies*) in critical computing centres. These basically consist of batteries and uninterruptible power sources that are capable of producing an alternative current similar to that of the mains supply, to take over when the mains supply fails.
Google, which owns the planet's largest computer park (perhaps a million servers) has just unveiled some of the methods it uses on its gigantic server 'farms'. One of these consists in equipping each server with its own battery. This distributed solution would then make it possible to come close to 100% efficiency, compared with the 92% efficiency of a centralised UPS. Another example can be found in the HPC solution now available from Bull, which has no need for UPS systems, because of supercondensors implemented within the computer enclosure. This type of component is far more robust than a battery and is capable of storing enough energy to counter power line disturbance on the mains supply. Whereas in the majority of HPC applications, power line disturbance is the only real concern, or at least in countries that benefit from a reliable network.
Let's start from scratch: data, handled by minute transistors engraved on silicium chips, consumes electricity. That's where it all starts. This consumption increases at a very high speed with the frequency at which we work, and ever since the invention of computers, it hasn't stopped climbing. Around 2002, the industry ran into an obstacle: beyond 25 watts per square centimetre (W/cm²), a chip becomes red-hot. This is why the industry did an about-turn at this time and started to exploit its capacity in a different way, to engrave an increasing number of transistors, which were increasingly small in size. In 2001, IBM launched the first 'dual-core' microchip (the Power 4) by grouping two central units together on one single microchip, an essential method for suppressing computer energy consumption (see article on page 48).

Reducing the consumption of microchips

The clock frequency of processors seems to have ground to a halt at between 2 and 5 gigahertz. IBM has even decided to backtrack with its BlueGene supercomputer range, by opting for a microchip from its Power range, with a very moderate frequency of 700 Megahertz (followed by 850). This is the reason behind its very low heat output and therefore unprecedented density: the first BlueGene/P, which was the world's most powerful machine in 2005, contained 1,024 of these dual-core microchips per enclosure. Furthermore, manufacturers are endlessly improving their techniques in an attempt to reduce the consumption of their microchips. Initially, this is achieved at a purely electronic level, involving methods that counter 'leakage current', 'stray capacity' and other miasma caused by miniaturisation. A well-known example is Intel's recent replacement of silicium dioxide with hafnium dioxide, which has a particularly high dielectric constant, in the transistor gates in its latest line of microchips engraved in 45 nanometres.
Another type of improvement consists of controlling the circuit's consumption with an increasing level of subtlety. By modulating the supply voltage of momentarily inactive modules, for example, or by scaling back the clock speed on all or part of the microchip, depending on workload and other criteria.

Replacing hard drives

More radical avenues are being explored for the near future. A laconic press-release from Intel recently referred to a prototype for a 'P-channel transistor' made on a silicium substrate, but using 'III-V' materials (in reference to their position in the Mendeleïev table). This creation, associated with a similar transistor - an 'N-channel' transistor, that has already been finalised, could pave the way to new integrated circuit technologies, supplied with a voltage half the current standard and consuming ten times less. Mass memories are also power-hungry. The progressive replacement of hard drives by units (SSD: *Solid State Drive*) that store data on 'flash memory' microchips is in the planning stages. This technology is used widely in handheld media players and other mobile devices. It is still more costly than disk technology, but it is making headway in leaps and bounds.

© Bull

Gate with high energy efficiency designed by Bull for its supercomputers. Air output is at an equal temperature to the ambient air and the energy consumption is reduced by 75% compared with traditional air conditioning.

Computers transform electrical power into heat, rather like a radiator. Power-loss occurs when they rise above a given temperature - the calories simply have to be evacuated. This is where the cumbersome 'Crac' (Computer Room Air Conditioning) systems, as they are called in the industry, enter the game. Smaller and larger fans are also used to dissipate heat right down to the level of the microchips themselves. These are high consumers of electrical power.

A small gain in efficiency has been achieved in recent years with 'blade chassis' systems which generally group together 6 to 8 servers, stacked vertically in racks, with a shared power supply and fan system.

Air flow is now being studied in depth and heat flow simulations are being used to find optimal integration solutions at all levels, from board components through to equipment in a room. Specialists are also endeavouring to find ways of regulating air-conditioning more locally and preventatively, using distributed software and thermal sensors.

When air isn't enough, liquid cooling technologies can be used. 'Cooling doors' are now creeping into the market. These doors are installed at the back of the computer enclosure and are refrigerated by means of water circulation. By forcing hot air out of the system with fans, via the blades of a radiator in which cooling water is circulating, the room is restored to initial temperature. This means that more power (computing and therefore electrical) can be concentrated in an enclosure (30 kilowatts and more) without drawing on the main air-conditioning system. Computer rooms increasingly use under-floor cooling water distribution networks.

This is a 'back to the future' scenario. In 1976, Cray-1, the mother of all supercomputers, was already using Freon liquid (used in most refrigerators at the time) for cooling purposes. Cray pursued this course (obviously using 'greener' liquids). Last year, the manufacturer launched ECOphlex, a new cooling technology using the phase exchange (gas-liquid) of a very effective cooling liquid, to transfer heat between an enclosure and an external air-conditioning installation.

Another idea is creeping in: how about if we just stopped setting the air-conditioning at 20 °C? We know why computers had to be kept cool initially, because of temperature-related failures. Intel wanted to kick this habit and conducted an experiment in New-Mexico, in two rooms filled with 450 servers each. One was permanently conditioned at 20 °C, while the other had an economical device which simply evacuated hot air, replacing it with outdoor air when it was at a temperature below 32°C. The findings of this experiment showed an electricity bill three times lower; and a higher, but acceptable, rate of failure. Food for thought.

Improved air-conditioning management

Generally, we talk about 'free cooling', i.e. air-conditioning as much as possible without using electricity. This is obviously much easier in cooler regions. In fact, Google has just announced plans to set up its next major European operation in Finland. Why ever would they choose this remote country with 5 million inhabitants? Its climate could well be the explanation. Not forgetting that computers can still be cooled using river water.

Since computer hardware produces heat, an alternative would consist in trying to put this heat to good use. This is what IBM did last year, for a client in the suburbs of Zurich. They installed a cooling system which transferred the heat output to the local swimming pool. A concept worthy of further investigation…

Since this all involves energy and waste, we might be tempted to lay the blame on computer hardware. However this is not the case. Firstly because a very efficient way of throwing kilowatt-hours out of the window is by running poorly-designed software on a computer. This does happen. But the HPC world has a different approach. We are more often inclined to optimise and re-optimise the codes to get the maximum possible out of the supercomputers that we have worked so hard to develop.

Lastly, the software itself can be of great assistance. Software solutions are now available for improving energy management at the level of each server, in a giant server farm for example, by optimising workload distribution, switching off unused equipment, and so forth.

Lastly, we should bear in mind that although computing has become a real energy drain, it will equip us with essential tools for a potentially greener lifestyle. Who would ever invent the car of the future, or clean energy generation systems, without using a computer?

(1) A megaflops corresponds to a processing capacity of one million operations per second.

(2) A petaflops corresponds to a processing capacity of 10^{15} operations per second.

Green 500, the energy-efficiency awards for supercomputers

The 'green' supercomputers rankings were launched in 2005, organised by two American researchers from Virginia Tech – Wu-chun Feng and Kirk Cameron. Green500.org uses the list of the most powerful machines, compiled by the famous Top500.org site and divides each one's power measured on the Linpack 'test bench' by its electrical consumption, to obtain an efficiency factor expressed in megaflops (1) per watt (Mflops/W). IBM dominated the top 20 places, followed by two Silicon Graphics and one Cray. The seven most efficient supercomputers are IBM's BladeCenter, with its famous 9-core Cell chip (designed for Sony's PlayStation). The winner, Nautilus, is installed in the Warsaw University and has an efficiency factor of 536 Mflops/W.
The n° 7 (445 Mflops/W) is no less than the Roadrunner, currently the world's most powerful computer (1.1 Pflops(2)), according to the Top 500.
IBM holds 8th to 17th position with ten of its famous BlueGenes (360-370 Mflops/W), equipped with a 'slow', and therefore cold, Power chip. We then drop down to the 260 Mflops/W range, with another IBM, holding 19th and 20th place with a third technological approach revolving around Intel's Xeon microchips.
We might add here that the manufacturer achieved this enviable ranking partly at the expense of innovations which are not equally beneficial to all application types (BlueGene is currently the most highly parallel computer) and which complicate programming, due to the heterogeneity of cores (up to three different types in the case of Roadrunner).
Cray's famous Jaguar (1.06 Pflops), number two in the Top500, is relegated to 79th position in the Green500 charts, consuming three times more electricity (7 MW) than Roadrunner (2.5 MW).

© IBM

Nautilus, number one in the Green 500 rankings, the world's most energy-efficient supercomputer. Installed at Warsaw University.

Grid: computers team

The purpose of grids is to get dozens of supercomputers, thousands of servers or even millions of PCs to work together. The idea is catching on.

Léo Gerat,
scientific journalist

The idea goes back over a decade. The oldest known grid, distributed.net (which still exists), was born in 1997. It now groups together some 160,000 PCs combining their energy to resolve cryptography problems. However Seti@home was the one to release the concept to the public in 1999. At the time this involved analysing chunks of signals received by a radio-telescope searching for possible evidence of radio transmissions from extraterrestrial intelligence. Three million people now dedicate their computer's spare time to this project.
In these examples, the grid unites the power of a very high number of personal computers. This type of machine community, which has very weak coupling (initially, the majority of PCs recruited were still dialling up a net connection via a modem), is suitable for a category of problems that are qualified as 'embarrassingly parallel'. Seti@home is a good illustration of this case: thousands of hours of signals recorded from the stars are parsed into small chunks which are then submitted to an algorithm which analyses them one by one. The volume of independent data renders the problem very parallelizable. But right from the outset, we envisaged applying the grid concept to other contexts. Particularly the concept of getting larger-scale machines to collaborate, linked together via broadband, to process parallelizable problems. The first 'tool box' for implementing a grid, Globus Toolkit, was created in 1997 by a team headed by Ian Foster (Argonne National Laboratory) and Carl Kesselmann (University of Southern California). The grid concept looks set for a dazzling career…
The initial idea is therefore to split heavy tasks into wafers, which are then distributed across a large number of scattered computers, which are theoretically heterogeneous and which are also used for other tasks. Thanks to an ad hoc software layer – we're talking about middleware here – the user sees all these computers as one single supercomputer. It's basically just another form of what is known as 'distributed computing' or 'grid computing'. This has become common currency at a time where all supercomputers are created by interconnecting a tribe of computers. The essential difference resides in the fact that these clusters are made up of identical nodes, gathered together in the same place and communicating via a high-speed 'interconnect'.
Ten years later, what's the latest with the grid concept? According to Thierry Priol, head of the Paris team at Inria (French national institute for research in computer science and control), *"It's true that the concept was somewhat oversold initially, something that happens to many new ideas, but it is now blooming and looks pretty good. We have a better idea of the limits and a better understanding of the challenges ahead."*

> **"Dedicated grids work better than those of a more general-purpose nature"**

A wide diversity of grids

The initial idea gave rise to a whole fauna of different species. We can now distinguish between calculation grids, data grids, and even mixed specimens. The superlative research instrument, the LHC (Large Hadron Collider) particle accelerator belonging to CERN (European Organization for Nuclear Research), in Geneva produces 15 Pb (petabytes) every year, i.e. 15 quadrillion bytes. These will be analysed by research teams scattered across the globe via the WLCG (Worldwide LHC Computing Grid) data grid. Another example is the Birn (Biomedical Informatics Research Network) grid in the United-States, which groups together laboratories specialising in neuroimaging.
Furthermore, we need to distinguish between the production grids and research grids. Egee, for example, which brings together dedicated machines, is an operational European Grid, while Grid'5000, in France, is quite clearly an experimental grid, used for research and gathering together 5,000 nodes distributed across 9 sites in France. Lastly, some grids are 'coarse-grained', such as the Deisa grid, which federates major computing centres, or on the contrary 'fine-grained', the best examples being Seti@home and its counterparts (around 40 projects of this type), which bring standard and private PCs together to work collectively via Internet.
"It should also be noted that a grid can function in two different modes", adds Thierry Priol. *"In 'Push' mode, where*

up

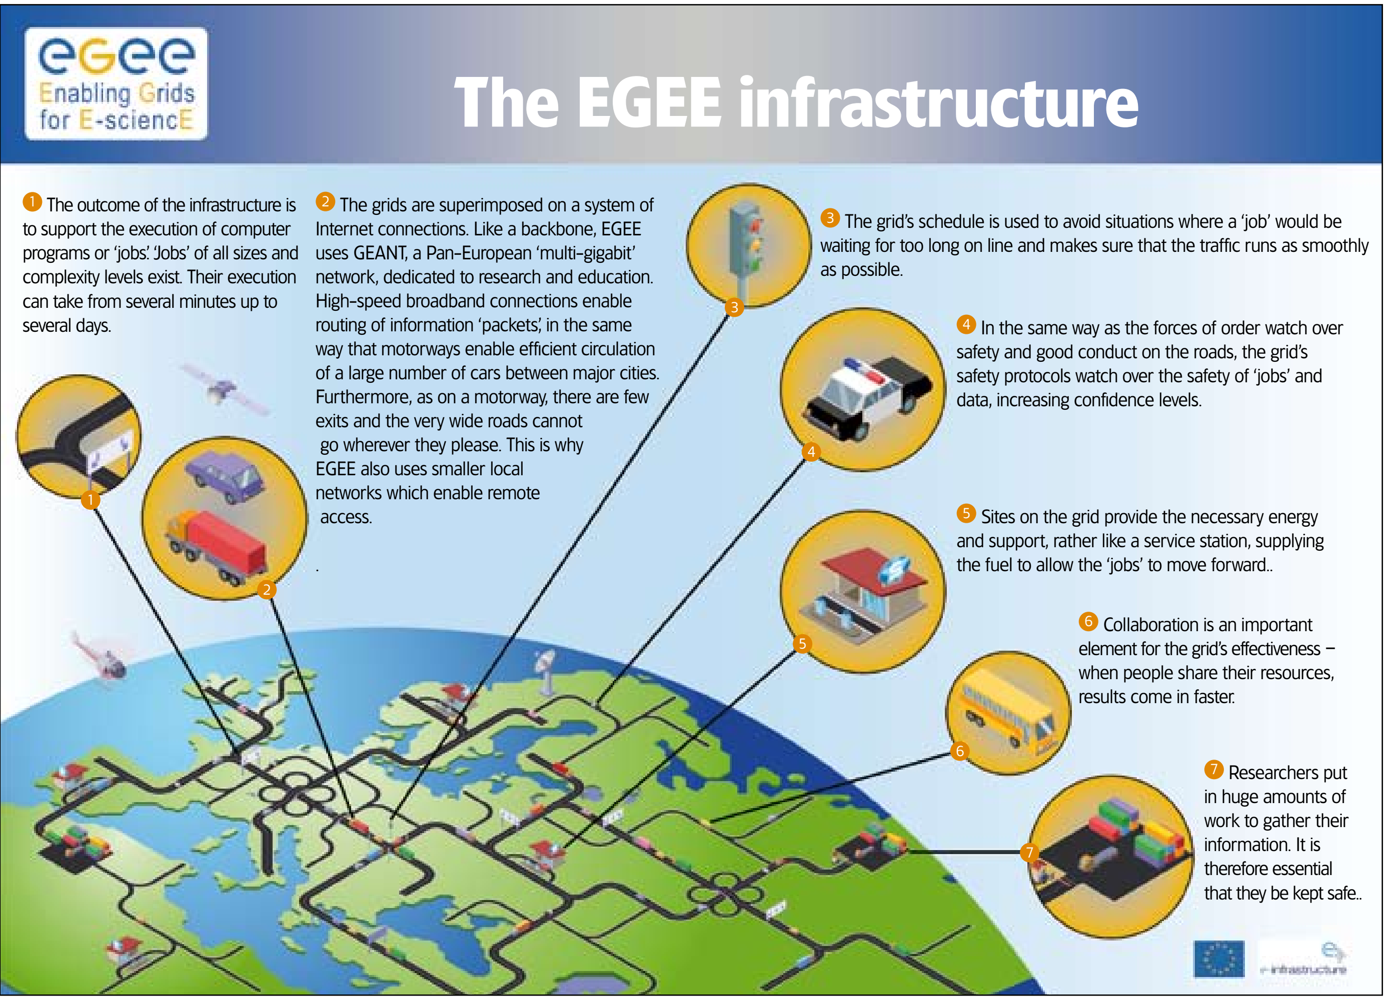

© EGEE

the grid administration software distributes the work and submits tasks to the affiliated computers, or in 'Pull' mode where members take the initiative and notify as and when they are available and request the work. PC grids of the Seti@home type work like this."

From *'Grids'* to *'Groud computing'*

"One thing clearer to us now is that dedicated grids work better than those of a more general-purpose nature," claims Franck Cappello, manager of Inria's Grand Large project, the originator of the Grid'5000 initiative, *"The European Deisa grid, which brings together a dozen major scientific computing centres, for example, or even Egee, which groups together some 300 sites and initially targets management of mass data, have proven their worth. On the other hand, the multi-task grid, Teragrid, in the United-States, is still trying to find its feet."*
Recently, a closer yet distinct paradigm is treading on the toes of 'grid computing': 'cloud computing'. This concept has emerged in a context far removed from scientific computing, with Web players such as Google, Yahoo, Amazon, Salesforce or IBM offering clients the opportunity to use computing resources offered by suppliers with vast 'server farms' via Internet, in an asymmetrical manner. Such resources can take relatively 'raw' forms such as processing power, or storage capacity. Inversely, more sophisticated forms such as plug-and-play applications are comprehensive services that free the user from any computing problems.
Is HPC really concerned by this 'cloud computing'? *"Not in the immediate future"*, adds Franck Cappello. *"It is lacking several critical elements, such as a public network of sufficient speed and internal networks that are powerful enough to satisfy intensive computing requirements."* We do however believe that these two concepts may draw closer together in the future, in one way or another, if not only by cohabiting on common installations. Ian Foster, the founder of the Grid, has already coined a name for the product of this union: 'Groud computing'.

Microchip cores: the rac

From now on, parallelism will burrow right into the microchip, microprocessors with several 'cores', becoming multiprocessors. Two, four, eight, sixteen cores... what next? Ideas are flowing in thick and fast.

Léo Gerat,
scientific journalist

Scientific computing has now entered the era of massive parallelism. From now on, any self-respecting supercomputer pulls out all the stops, right down to its microprocessors, in which several cores work side by side, sharing data. Two cores per microchip is already the minimum, four is commonplace, eight is standard and we are now waiting with baited breath for microchips with 16 cores, 32 cores and more.

Were these chips developed specifically for high-performance computing (HPC)? Not in the least. We find them in supercomputers because manufacturers no longer have the means of developing the microchips of their dreams, so they have become accustomed to using what they find off the shelf, whereas computing in general has moved on to multi-core technology. All, or virtually all, recent computers are equipped with a dual-core microprocessor. Does Jo Bloggs really need to use parallelism like nobody's business? The answer to this is "no". Nobody ever asked for this new form of parallelism, no more than they asked for any other. The simple truth is that if a manufacturer were to announce a new supercomputer tomorrow offering a power of several teraflops[1] without parallelism technology, running at a speed of just a few Terahertz, for example, we wouldn't really blame them. Above all, parallelism means programming problems. Multi-core is not the exception to the rule. Industry isn't doing this just for the fun of it. *"Multi-core is the educated choice of an almost overwhelming majority"*, states Olivier Temam, who heads the Alchemy project at Inria Saclay.

"Multi-core is an educated choice"

The computer industry took the multi-core avenue at the start of the century. A short history lesson: Initially, industry began to regularly reduce the size of transistors, which meant that even more could be fitted onto one and the same chip, simultaneously increasing operating speed. In theory. But the faster a transistor commutes, the more it consumes electricity and outputs heat. In addition, miniaturisation reduced the size of some of the insulating components of the circuit so much that we witnessed an increase in leakage current, the chip therefore started to consume significantly, even when it was inactive. Worse still, over the following years, reducing the size of transistors would not lead to any significant reduction in their consumption. The consequence of this was that they could no longer run simultaneously and we could no longer use all the cores of a multi-core, rendering the multi-core option useless.

At the start of the new millennium, it was evident that it had become increasingly challenging to increase speed. Even more so, given that we had become accustomed to upholding this parameter as an indicator of power for each new chip. If we were no longer capable of producing the new generation with double the speed, whatever were we meant to tell our clients? We made up our minds; we would exploit developments in miniaturisation by engraving two processors side by side on the same chip, with that little added extra to interconnect them, of course. In 2001, IBM launched the first dual-core microprocessor, the Power 4.

Intel and AMD followed suit in 2005. The clock then stopped at a handful of gigahertz (GHz) and now we're counting our cores.

The race is back on, with two, four or eight cores per chip, and plans for 16 then 32. We have started

NEHALEM, Intel's latest microchip.

© Intel

(1) A teraflops corresponds to a processing capacity of 10^{12} operations per second.

e is on!

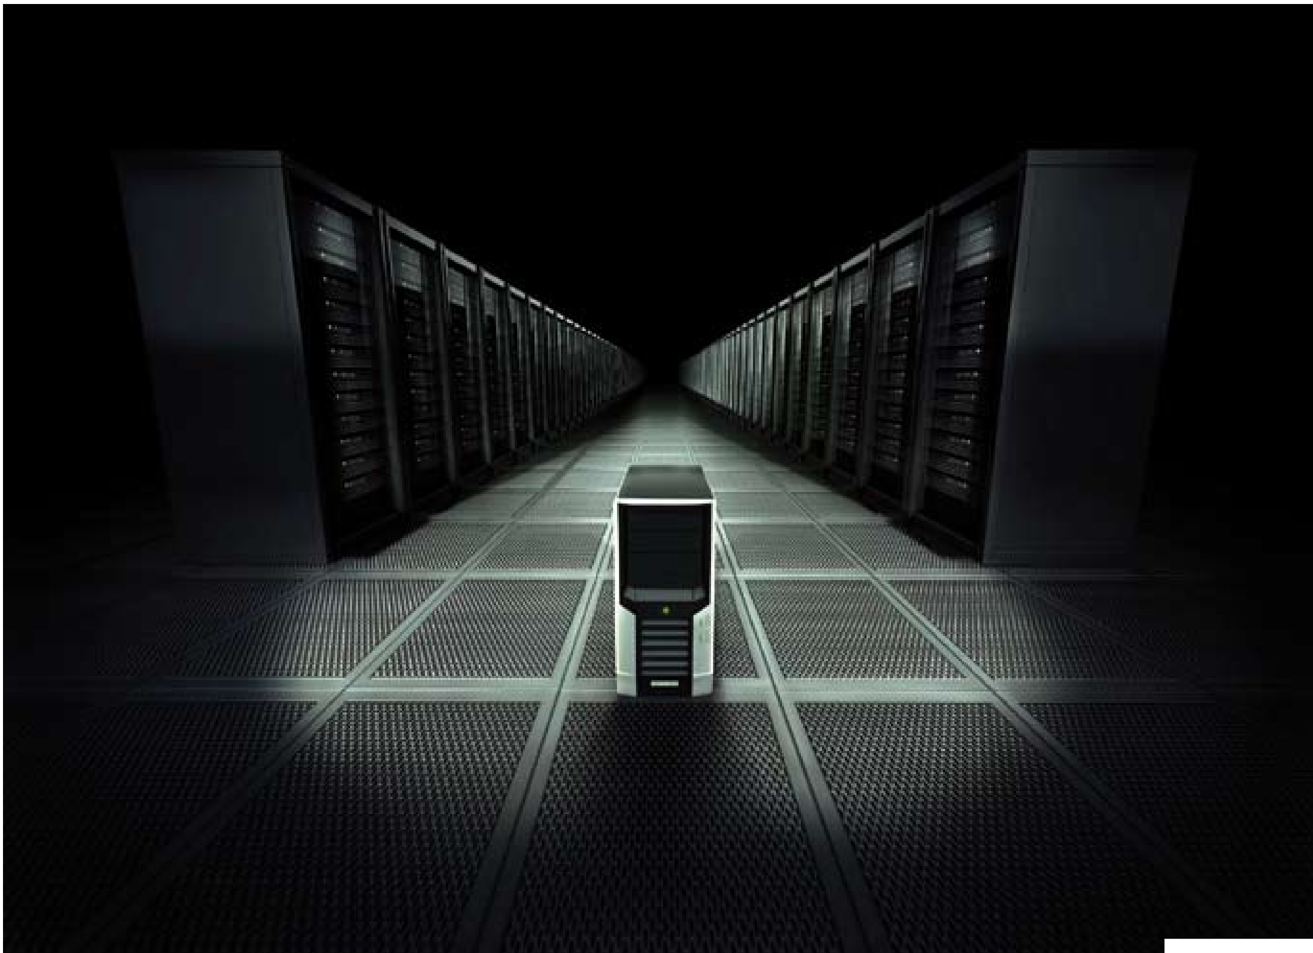

© Nvidia

The NVIDIA® Tesla™ personal supercomputer. Based on GPU computing architecture, it runs with 960 parallel cores.

to increase cores rather than GHz. But is it really that simple? To start with, juxtaposing these cores simply isn't enough for us. *"The essential question lies in the sophistication of core interconnection"*, explains Jean-Laurent Philippe, European technical director for Intel. The more there are of them, the more they need to be able to exchange efficiently, to share the same rapid access memory (RAM), thereby consuming silicium too.

However, it is quite clear that the number of cores looks set to rise exponentially. In December 2007 Intel presented an 80-core prototype named Polaris at the International Solid-State Circuits Conference in San Francisco.

But, hold on, do we really have to keep making more copies of the same core? And what if we were to add other types of processors, capable of doing things faster and better, to the standard, general-purpose central processing units (CPU)? This is standard practice in any personal computer: even the least-sophisticated laptop has a general-purpose microprocessor supporting all standard software programs, as well as a GPU (Graphics Processing Unit) on which all those computations producing the astounding images animating our screens are performed.

"A graphics processor", explains André Seznec, in charge of the CAPS project at the Rennes Inria institute, *"is no more than a large vector computer, a type of Cray."* In other words, a processor specialised in executing operations between vectors, where image processing, amongst other operations, is voracious.

A GPU can be found in any computer, but it is also often multi-core itself, and not just moderately so. nVidia's Tesla T10 chip has no less than 240 cores.

It should also be added that these latter perform a small set of very specific operations and that, due to this factor, they are less complex and occupy a smaller surface area of silicium than a general-purpose core, as would be found on the Nehalem chip, Intel's latest development.

It's one thing that these microchips were designed for graphic processing, but certain HPC consumers showed huge interest in them, although they're not actually using them for images. Over time, the

manufacturers widened the spectrum of potential for utilising GPU microchips, to account for demand, which is particularly prominent in the video gaming industry. Today's video games are going a lot further than retouching images. A quick look at them will show that they apply simulation of varying physical phenomena, for example. Furthermore, the idea has been brewing for some time now to create supercomputers, which are indeed rather specialised, using GPU microchips. A manufacturer such as nVidia, has lately been offering HPC solutions using its GPU microchips.
GPGPU (General-Purpose Processing on Graphics Processing Units) is the term used to describe the method of performing scientific computing on graphics microchips. Some are really pushing this to the limits. For example, a team at the Tokyo Institute of Technology created a hybrid machine, known by the name of Tsubame, which contains, amongst others, 680 nVidia T10 microchips, each equipped with 240 cores, representing a total of 163,200 graphic cores. In France, an exotic prototype of this type is currently being produced by Bull, commissioned by Genci (Major National Infrastructure for High-Performance Computing). This hybrid machine is set to integrate 192 nVidia Tesla T10 microchips.

© BULL

One of the most powerful hybrid systems in the world, recently installed at the CCRT (Centre for Research and Technology Computing) in Bruyères-le-Châtel.

GPU and CPU, soon on one and the same microchip!

From now on, if virtually all computers need a graphics processor in addition to their general-purpose microprocessor, we might ask ourselves why these two wouldn't be integrated on one and the same chip? According to André Seznec, *"Technically, it would have been logical for the GPU to have migrated to the CPU's microchip a long time ago. Today we could have microchips combining CPU and GPU, rather than dual-core CPUs. But market forces have decided to apply their own logic."*
One possible development of multi-core technology would then see graphics cores joining general-purpose cores on the same piece of silicium. *"This is one possibility"*, admits Intel's Jean-Laurent Philippe,

Moving towards variable geometry cores?

The ultimate dream, for the end user, would be to have access to a computer that is perfectly tailored to their problems, equipped with microchips designed to run their software as efficiently as possible. However, the computing industry does the exact opposite, investing colossal sums of money in the mass production of a small number of microchips targeting vast markets. Of late, supercomputers have been manufactured using the same chips as ordinary computers, sold in their millions An old idea aims to resolve this contradiction between the efficient application-specific circuit and the economical generic circuit. *"This is the theory underlying reconfigurable circuits"*, explains Olivier Temam of Inria. The idea consists in producing circuits made up of modules with an interconnection that can be programmed at a given point in time, thereby creating, momentarily, on demand, an operator capable of performing a given calculation. The most widely-known form is the FPGA (Field-Programmable Gate Array), a chip comprising a network of logic gates that are interconnectable on demand, controlled by a software program. This makes it possible to quite literally 'program' a logic circuit. *"A fascinating concept on paper, but which in practice would doubtless be somewhat restricting to begin with"*, adds Olivier Temam.
FPGA solutions are available on the market and are used to meet some specific needs. But there is nothing to say that a future generation of multi-core processors couldn't one day integrate circuits that are 'slightly' re-configurable. *"Potentially, this is an interesting idea, but is still on the drawing board,"* admits Intel's Laurent Philippe.

IBM's 9-core Cell microchip.

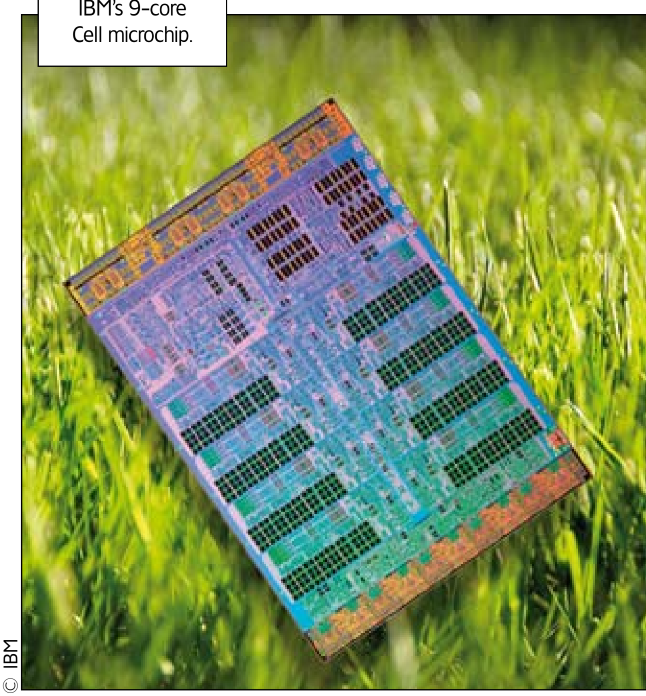

© IBM

"One thing is sure - the GPU, which is now often a far cry from the CPU (on a 'graphics card'), will find common ground and may eventually find itself on the same chip."

This movement will be even more desirable if, at the same time, the graphics processors were to gradually become less specialised. Intel's description of its future graphics microchip Larrabee – with the first 24-core version expected for release by the end of the year, is interesting from this point of view. Larrabee should offer both stunning performance as a graphics processor, whilst also lending itself to the GPGPU's game. We can probably expect to see it again on HPC computers.

Another microchip goes even further still. IBM's famous Cell designed in partnership with Sony comprises both a traditional core, from the Power range, with eight cores specifically designed for high-performance scientific computing. It shows the promising future of the heterogeneous multi-core microchip, comprising both general-purpose cores and more specialised cores.

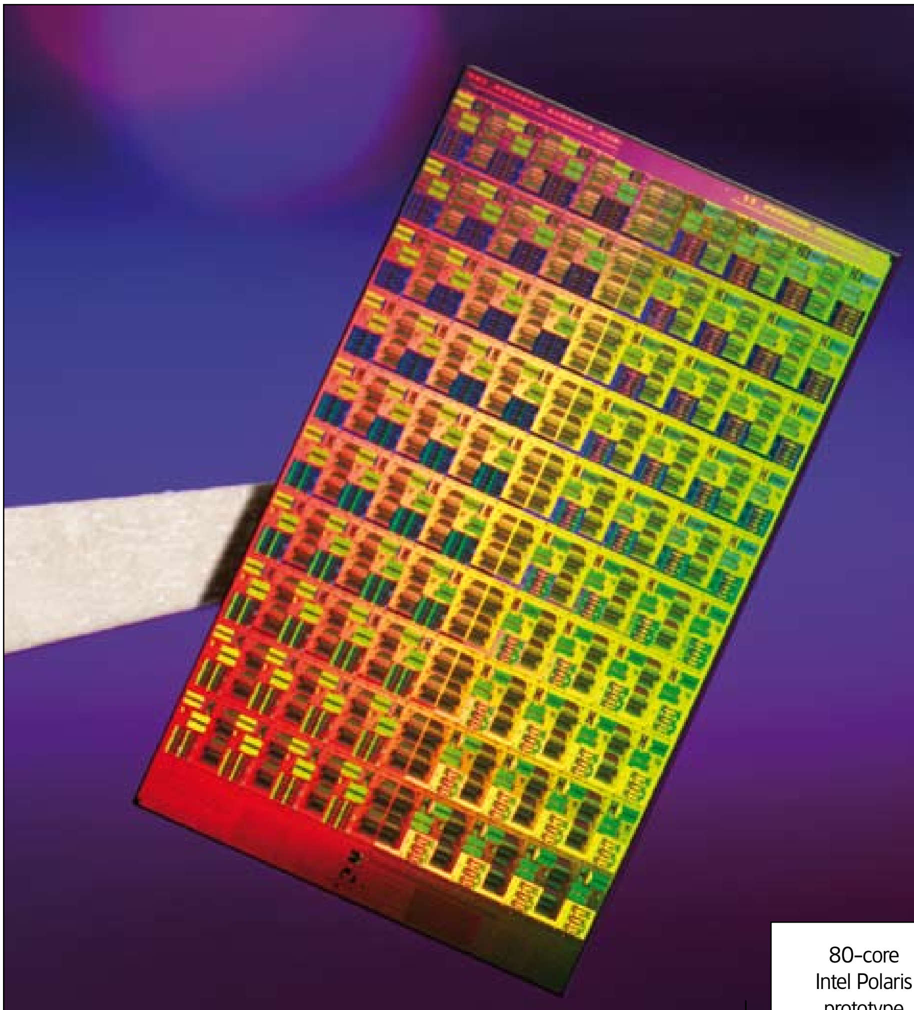

© Intel

80-core Intel Polaris prototype microchip.

We are even thinking about making the game more complex. According to Olivier Temam *"the multi-core is not the only way of utilising the increase in the number of available transistors. We can also use it to add all kinds of accelerators."* If in the near future, we have enough space on the silicium to install dozens of cores, why would we not keep a portion of them aside to engrave specialised operators on them, capable of performing given operations highly effectively?

Is it really worth the effort? Does the speed at which this highly specialised silicium is useful actually justify the space that it requires? Such choices will depend on robust market research. Ultimately, the answer simply lies in using the 'silicium budget' available to the circuit's designer in current technology, more efficiently.

In this quest for the perfect balance between efficient application-specific functions that are rarely used and generic functions, we would like to be able to draw on both possibilities. *"One idea consists in making sure that overly specialised functions are not engraved"*, adds Olivier Temam, *"but rather algorithm sets that can be used in several different contexts. We are attempting to define functions that are application–specific while being flexible at the same time."*

It is therefore impossible to draw a conclusion on the role of the multi-core in high-performance computing without taking into account the software aspect. *"We mustn't be fooled into the trap, it's not easy for an application to exploit the full potential of the multi-core"*, adds Thierry Priol, head of the PARIS team at Rennes Inria. The optimum use of this new way of producing scientific computers will require major work on the software side.

© NVIDIA

Tesla NVIDIA GPU card.

High-performance computing: child's play?

In the world's top 500 biggest supercomputers, 'hybrid' architectures are up and coming. Combining general-purpose computing processors and graphic cards such as those found in video gaming, they promise computation time accelerated by factor 10, provided that the software in question is 'hybridised'.

Isabelle Bellin,
scientific journalist

Ever since the advent of parallel architecture (see article on page 48), algorithms for scientific applications have been adapted to these new computing environments. For the past two years, the graphics processors on gaming consoles, a hundred times faster than general-purpose processors, have taken the multi-core route. This new architecture, dubbed hybrid, promises to boost high-performance computing. But reaping the benefit of this potential would mean adapting their programming technique - this is a complex area of research still in its infancy - whereas these new computing titans are already escalating. Last April, 300 teraflops (Tflops[1]), achieved in Europe for the first time and funded by Genci[2], were installed by Bull at the CEA at Bruyères-le-Châtel.

A major application: nanosimulations

At the front of the queue of applications interested in the technology, we have nanosimulations. These electronic-scale simulations are essential for manufacturing nanomaterials at a reasonable cost. *"100,000 hours of computing are needed with a standard supercomputer architecture to simulate the behaviour of 100 atoms for a picosecond with the Abinit software, one of the worldwide leaders in the field"* adds Thierry Deutsch, researcher at the atomistic simulation laboratory of the CEA in Grenoble. What needs to be done to ensure that such computation code benefits from the promising capacities of hybrid architecture? *"We need to know how to estimate which codes can be transferred, which size architecture they need, which computations should be allocated to the general-purpose or graphic processors and how to limit the exchange of data between them"*, answers Jean-François Méhaut, lecturer at the Joseph-Fourier university in Grenoble and researcher for the Mescal team at Inria. A tall order! The aim is ultimately to automate this process whilst maintaining the same graphic card performances from one generation to another, from one hybrid architecture to another. These teams at Inria and the CEA have been working on this together for two years. *"For the time being, we are gradually localising the most time-consuming portions of the computations and those most easily parallelised in order to attribute them to the graphics cards"*, explains Jean-François Méhaut. They are then accelerated by factor 30 or even 50. But the 'slowness' of other portions is still slowing down performances. *"On a small hybrid 20 Tflops supercomputer, we tested one of our nanosimulation (the wavelets method) codes, which we know to be massively parallel"*, adds Thierry Deutsch. *"We gained a factor 7 in computing speed without optimising the distribution of the computations between graphics and general-purpose processors."*

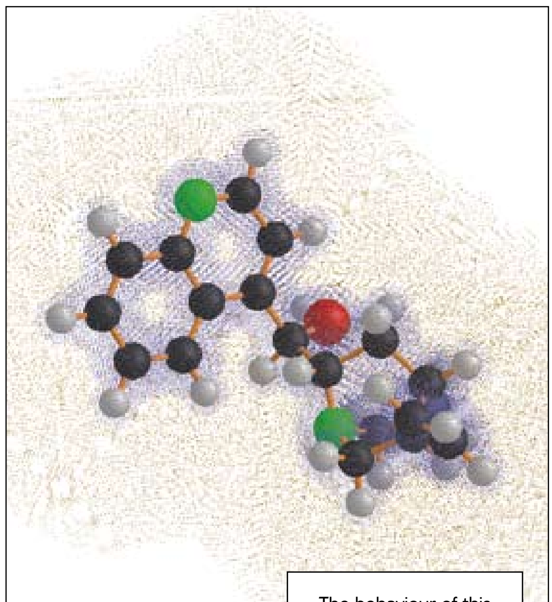

The behaviour of this type of molecule can be calculated very quickly with graphic processors. We can also see the adaptive mesh associated with it.

© CEA

Researchers are counting on a factor-10 gain. This is one of the goals of the ProHMPT[3] project, launched in January, coordinated by Inria and supported by the National Research Agency. *"We are also counting on the next world standards in programming dedicated to this architecture"*, continues Jean-François Méhaut. *"One of them could be developed by CAPS Entreprise (Inria-mentored spin-off), one of the partners of the ProHMPT project."*

(1) **One teraflops corresponds to a processing capacity of 10^{12} operations per second.**

(2) **Major National Infrastructure For High-Performance Computing (GENCI).**

(3) **ProHMPT involves researchers from Inria, the CEA, Bull, CAPS company and the University of Versailles-Saint-Quentin-en-Yvelines.**

Shattering the Petaflops record: a public-private partnership unique in Europe

Due to go live in 2010, Tera 100, developed by Bull in cooperation with the CEA, will be the first supercomputer designed in Europe to break the symbolic Petaflops record, or a computing capacity of a quadrillion operations per second.

By **Sophie Houssiaux,** director of the Tera 100 project at Bull, and **Pierre Leca,** head of the simulation and computing sciences department of the CEA military applications division

For the past year, teams of engineers at Bull and the CEA military applications division have been working together on the design of the Tera 100 supercomputer. This public-private sector team brings together a development capacity that is unique in Europe with an expert knowledge of application needs, integration constraints and computing architecture. The task is to group together several thousands of processing nodes (the basic elements in a supercomputer) using a network of dense interconnections. This will connect between 3,000 and 5,000 nodes, in this case multiprocessors using around thirty 'cores', the elementary computing units.

Tera 100 will have an 'island' architecture linking clusters of nodes which will have rapid access to data storage media: it will look like a dense network of large cities interconnected by motorways, rather than a network of medium-sized towns linked by minor roads.

Having taken into account both the total cost of ownership) and increasing ecological concerns, the Bull/CEA team decided to focus on energy-saving. A whole battery of measures was taken to limit electrical consumption, one of the main ones being a software optimisation system that enables the frequency of the cores to be lowered if they are not being used at full power, or even to shut down a core whilst controlling its impact on running computations. With an estimated 5 megawatt consumption, Tera 100 is throwing down the gauntlet in terms of the problem of heat dissipation. A water cooling system will lead to higher energy efficiency and will also be more compact. The supercomputer will therefore take up the same floor space as its predecessor Tera 10, or 600 square metres, for twenty times more computing capacity. Another specific feature: the multi-core processors are capable of executing 50 to 100 billion operations per second, and the physical network of interconnections between nodes is standard. As for the software, this is based on free software optimised to reach the levels of performance required. This deliberate building of the supercomputer from non-proprietary (and therefore more easily adaptable) elements was one of the foundations for the success of Tera 10. For Tera 100, the emphasis will be placed on a modular, general-purpose architecture, compatible with the environment of workstations running the Linux operating system, fault-tolerant and free of data processing bottlenecks. The whole research community will therefore benefit from the advances made during the development of Tera 100. This is the main advantage of such a co-operation between a public organisation and an industrial player, still unique in Europe, which in the future needs to be developed if we are to catch up with the major countries at the cutting edge of high-performance computing.

The whole research community will therefore benefit from the advances made during the development of Tera 100.

The whole research community will therefore benefit from advances made during the development of Tera 100.

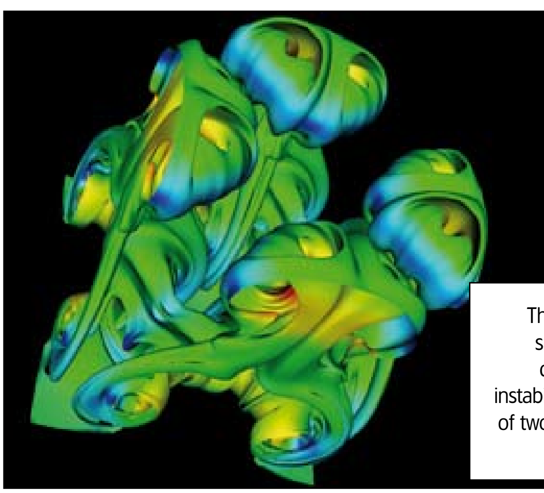

© CEA

Three-dimensional simulation of the development of instabilities at the interface of two fluids with different densities.

Ter@tec, Europe's first technical park dedicated to high-performance simulation.

© www.imaconcept.tv

On the road to

After the petaflops, the next challenge will be exaflops (performing a quintillion operations per second). There will be many obstacles to overcome.

Léo Gerat,
scientific journalist

The petaflops threshold (a quadrillion operations per second) was reached in 2008 by the IBM Roadrunner, eleven years after the teraflops threshold (1997: ASCI Red, Intel). Now scientists are setting their sights on a quintillion operations per second. Demand for this has already been expressed by governmental administrations such as the DOE (Department of Energy), in the United States, as well as certain scientific sectors such as climate modelling.

Exaflops within the next ten years? This might as well be infinity. How can this factor of a thousand be achieved? There are few voices offering an answer at this stage, but many more asking the questions. A report commissioned by Darpa[1], a team headed by Peter Kogge (University of Notre-Dame, Indiana), entitled Exascale Computing Study, set out four major 'challenges'. The first is the question of energy. Everybody knows it; this is the fundamental problem.

Roadrunner consumes 2.5 megawatts (MW), which already seems an achievement when you consider that Jaguar (Cray) is hot on its heels in terms of processing power, but consumes an almighty 7 MW. If the energy bill were to be multiplied by a thousand to achieve the exaflops level, we would be looking at gigawatts, which is roughly the output of a large power station. This is preposterous!

We might, at a push, consider a one-off computing installation consuming a tenth of a gigawatt. But 125 MW for an exaflops for example, would represent an increase in energy efficiency of a factor of 20 compared with Roadrunner, which is not to be sniffed at.

We really need to make serious headway in this area. This means first and foremost looking at the microchips that would supply the 'flops'. However, the cost of developing a new generation of microprocessors is now so great that supercomputer manufacturers can no longer envisage specially commissioning them for their own use.

According to William Jalby, team manager of the Prism laboratory's Arpa team at the University of Versailles-Saint-Quentin-en-Yvelines, *"this is a problem we are going to have to live with. Only mass-produced 'retail' components offer a satisfactory power-price ratio."*

Today's supercomputers use microprocessors designed for other uses. Not long ago, these were microchips designed for large computers, or servers. But more recently, we have started to look into more modest and cheaper microchips.

Enter the IBM BlueGene - number one in the Top 500 in 2004 - which was already using a modest microchip originally designed for embedded systems. The same was true for 2008's champion, Roadrunner, another IBM jewel, which this time uses a microchip designed for a games console, the Sony Playstation.

"The world of high-performance scientific computing has become accustomed to borrowing microchips which were not designed for it", states Jean-François Lavignon, R&D partnerships director at Bull. *"We saw a good example of this in France with this hybrid prototype that Bull built for Cines (National Computing Centre for Higher Education), that uses nVidia's Tesla T10 240-core microchip, designed for graphics cards."*

This is why it is often said that the next generation of supercomputers could be built using microchips designed for mobile consumer devices. At Santa Clara in California, a team of researchers at the Lawrence Livermore Laboratory and Stanford University found the

(1) Defense Advanced Research Projects Agency (American).

exaflops

© www.imaconcept.tv

The Ter@tec Campus at Bruyères-le-Châtel will have 1,000 people working in 15,000m² of laboratories and offices by 2011.

technology that should enable them to build an exaflopic computer, capable of simulating the earth's climate by fragmenting the atmosphere into 20 billion cubes.
Xtensa, a mobile microchip designed by Tensilica, has 32 cores and consumes only 0.09 W. According to these researcher's calculations, the processing power per watt obtained is four times what is offered by the microchip in IBM's BlueGene and 100 times higher than an Intel Core 2 microchip, designed for laptop computers.
The second challenge identified by Darpa is memory and storage. Feeding millions of starving cores will be no easy feat. New memory technology would be more than welcome. Steve Scott, technology director at Cray, is optimistic on this subject.

We're going to have to literally re-invent a new way of programming

As for long-term storage, flash memory (used in mobile devices) is expensive but very modest, and will probably play a major role. *"We can expect machines with a layer of flash memory built-in between the central memory and the disc memory in the future"*, claims Steve Scott. This opinion is shared by Steve Pawlowski, technology director at Intel, a major player in the flash memory market.
The third gamble identified by the Darpa is entitled "concurrence and location". *"We are going to have to learn how to distribute the application over a billion cores"*, says Willian Jalby. *"This is not too much of a problem for those algorithms that work frequently with local data. But what about the others….?"*
Franck Cappello, project manager of the Grand Large project at Inria Saclay, recalls that *"ten years ago, nobody would have predicted that multi-core technology would become a vital building block for petaflops. Now, we do not know what the building blocks for exaflops will be, or whether multi-core technology will play a positive role. It may not even be all that easy to exploit."* Because of this, we think that we will have to literally re-invent a new way of programming. We may even see some strange practices emerging. As William Jalby explains *"for example, it is sometimes better to re-compute the data that we need immediately, than to wait until it reaches a distant microchip."*
Energy management itself would also have a great influence on programming. Serge Petiton, head of the MAP team at the Lille Fundamental Computing Laboratory explains that "while the results from two parallel computations are needed in order to do something else, it is sometimes more logical to slow down the faster of the two, to save watts."
The fourth and final challenge of exascale identified by Darpa is reliability. According to Jean-François Lavignon *"this is a key question, one which will force us to re-invent the way in which we design computers." "The complexity of these machines is such that the average time between two failures will probably be hours, if not worse"*, claims Jean-Pierre Panziera, engineer at Silicon Graphics (SGI).
Franck Cappello feels that "failures will become normal events. We must therefore invent technologies which play them down, allowing us to continue running each software program and therefore applications, while we carry out repairs locally."

Ter@tec, architect's drawing.

© www.imaconcept.tv

European research at

The European PRACE (Partnership for Advanced Computing in Europe) project's main aim is to provide European researchers with the world's most cutting-edge simulation and modelling installations. This partnership, based on trans-national collaboration and interdisciplinary exchange, intends to equip Europe with the means to match world standard.

Jane Nicholson, Director of the PRACE initiative and research manager at the EPSRC (Engineering and Physical Sciences Research Council), the British governmental agency funding the research and training of physical sciences engineers.

Simulation and modelling is an increasingly important approach for understanding the world around us. Being able to model hypothesises and explore the details that simulation and modelling can provide can open doors to new scientific insight, enable studies to take place that could not be carried out in experiment, test new theories and provide the key pointers for the development of experimental work programmes helping to save time and money. This contributes not only to furthering our scientific understanding but also contributes to our ability to exploit science for both economic and societal gain.
Many models used to date have been based on simple systems with extrapolations made to try to use the information gained to understand larger scale problems. But real world scenarios do not often build on simple models but are highly complex incorporating many features which call out for increasingly greater modelling capabilities and the computational resource to tackle them.

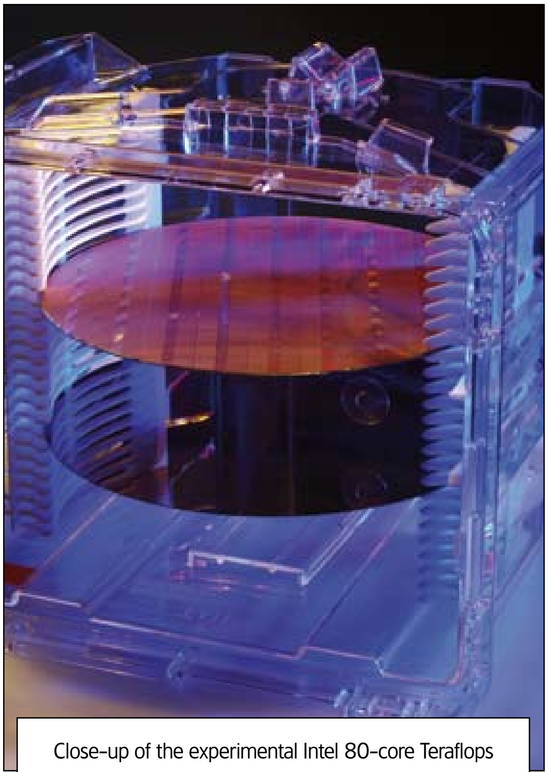

Close-up of the experimental Intel 80-core Teraflops processor.

Adequate computing resources

Challenges across diverse fields from fuel efficient aerospace engineering, understanding and predicting climate change, modelling the heart to fundamental science such as particle physics and attosecond science all benefit from access to leading edge computational resources. As can be noted these areas not only provide exciting science but also contribute to addressing real problems we face as a global society.
The value placed on ensuring that research communities have access to leading edge computational resources can be seen by the investments that many European countries have made in the last few years in their own national facilities. However, as we look to the future, while the need for access to computational resources will continue and the ingenuity of the hardware manufacturers continues to offer new machines that can enable increasingly complex problems to be tackled, the ability to be able to afford to provide a wide number of machines across Europe is likely to be a challenge from both the perspective of the costs involved and the facilities required to host these increasingly demanding machines.

"A challenge in terms of cost and expertise"

Thus the aspirations of the 18 partners in Partnership for Advanced Computing in Europe (PRACE) to develop a shared provision of leading computational resources is very timely. A collaborative approach across Europe to tackle the most complex problems that can exploit the thousands to the hundred thousand to a million cores that will be available in future supercomputers is the one which will enable

the forefront

© Tensilica

The Xtensa microchip from Tensilica, 0.09 W for 32 cores.

European research to maintain and develop it strong position on the world stage.

Way back in Eighteenth century Johann Wolfgang von Goethe said that "Science and the arts belong to the whole world. The barriers of nationality vanish before them" and this remains true today. Approaches that enable the best researchers from different nations to work together on shared resources will provide opportunities for outcomes that individuals working alone may not be able to achieve.

A collective movement to respond to future challenges

As the current chair of the PRACE Initiative, a collaboration of 18 countries who have all signed a Memorandum of Understanding, I have seen at first hand the enthusiasm and interest in developing a shared approach to providing European researchers with access to future facilities for computational simulation and modelling at the highest level. The Initiative has focused its efforts in first 18 to 24 months of operation on developing a project which will ensure we have the best hardware for the challenges to come. Called the PRACE project, this first project from the Initiative has been established to create a persistent pan-European High Performance Computing service for research. With funding from the partners and the European Commission the preparatory phase, which will run until the end of 2009, will establish the basis of transnational organisational structure for scientific supercomputing in Europe. By bringing together the know-how and resources of the partners, PRACE could provide European researchers with access to supercomputing capacities at a world-class level, transgressing those affordable at the national level. This includes a coordinated approach to hardware procurement and potentially a European platform for the development of hardware and software jointly with industry. Close cooperation with national and regional computer centres and scientific organisations will ease access to computing resources at all levels for scientists and engineers from academia and industry.

Every one who works in support of the use of high end computing knows that as much as having the

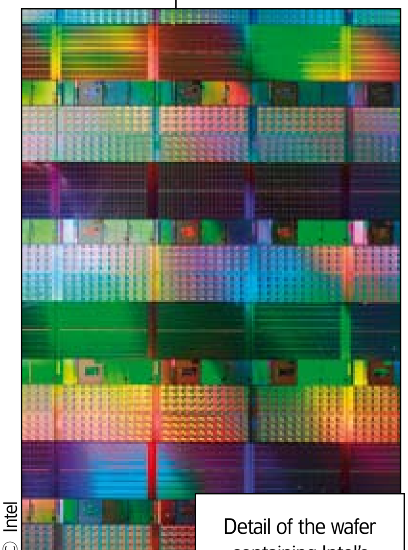

© Intel

Detail of the wafer containing Intel's first experimental Teraflops programmable processor capable of executing more than a trillion operations per second.

PRACE project objectives and participating countries

The objectives of the PRACE project are to:

- Prepare the creation of a persistent, sustain4able pan-European HPC service
- Prepare the deployment of three to five Peta4flop/s systems at different European sites
- Define and establish a legal and organisa4tional structure involving HPC centres, na4tional funding agencies, and scientific user communities
- Develop funding and usage models and es4tablish a peer review process
- Provide training for European scientists and create a permanent education programme

18 countries signed the PRACE initiative

- France
- Germany
- The Netherlands
- Spain
- UK
- Austria
- Finland
- Greece
- Italy
- Norway
- Poland
- Portugal
- Sweden
- Switzerland
- Cyprus
- Ireland
- Serbia
- Turkey

© CEA

Numerical simulation performed on a Tera 10 supercomputer. Here, a point-in-time snapshot of a building just as it is crossed by a seismic wave; note the deformation of the tower.

best computing hardware is essential, having the technology alone does not allow you to make progress. Users of supercomputers need to have access to software that can exploit the hardware to the full and to develop both of these we need a highly skilled group of people which includes those who know how to manage the hardware, programme the machines and prepare the software code to produce the results that science and industry needs. While PRACE to date has focused on the first of these three key ingredients, the members of the initiative have benefited from the connections and shared understanding they have formed over the last two years and are now beginning to look at how, by working together, we can tackle these other key issues. Collaborations are being formed that will look at the future areas to focus on for software development and PRACE has already run two successful training courses for young researchers with opportunity to do more in this space in the future.

A powerful knowledge database

From a UK perspective we have already been looking to ensure that future plans encompass not only the provision of hardware but also build knowledge and expertise in the research community in software development and to increase the linkages across disciplines bring together mathematicians and computational scientists so that future software developments for science can benefit from advances in numerical analysis and computer science.

It is proving to be a busy time for those involved in support of the computational requirements of research – the technological opportunities opened up by the move to multi and many core processors raise new challenges in how to programme them and exploit all that they can deliver, the potential scale of new systems provide challenges in securing the funds in these financial constrained times but there are rewards to be had in the development of scientific strengths within Europe and with partners from across the globe, if we tackle these challenges together.

(1) A teraflops corresponds to a processing capacity of 10^{12} operations per second.

(2) Centre for research and technology computing.

(3) Institute for Development and Resources in Intensive Scientific Computing.

(4) National Computing Centre for Higher Education.

(5) High-performance simulation for small businesses.

"Three questions for"

Catherine Rivière, Director of Genci (Major National Infrastructure for High-Performance Computing)

Interview by **Isabelle Bellin**

What's new at Genci, this company created in 2007 and jointly run by the state, the CEA, the CNRS and universities, and designed to fund and organize French computing resources?

Our annual budget of 25 million euros has enabled us to double the resources available in France for high-performance computing. The computing capacity available increased from 21 to 470 teraflops (1) at the start of 2009, thanks to the installation of machinery in the CCRT(2) at the CEA, IDRIS(3) at the CNRS and the Cines(4) for the universities. The computing time is now attributed in a transparent way and shared between the three centres by scientific peer communities.

Genci represents France as its main partner in the PRACE project. What does this job consist in?

This basically consists of our aim to make our national policy of high-performance computing part of a European perspective. This involves a financial commitment, of 25 million euros per year, from 2010 to 2015, and the installation after 2010 of a supercomputer in France on the site of the Very Large Computing Centre (TGCC) in Bruyères-le-Châtel, as part of the Jacques-Louis Lions Institute shared by the CNRS and the CEA.

What do you think will be the major knock-on effects of this in France?

Researchers will obviously have access to an exceptional level of processing power and a variety of supercomputing architectures. We are also opening up these resources to industrial partners and small businesses. We are organizing a seminar on this subject in Toulouse on the 8th and 9th September with our German colleagues at PRACE. With Inria in France, we are putting together a program(5) designed for small businesses, in partnership with competitive clusters, to facilitate access to high-performance calculation. This could concern around a hundred small businesses. This initiative will serve as a valuable example for PRACE.

bullx

Designed without compromise for UNLIMITED INNOVATION.

Intel Xeon inside
Powerful. Intelligent.

bullx extreme computing

The first supercomputer specifically designed for Extreme Computing by Europe's largest team of experts. Equipped with Intel® Xeon® 5500 processors.

Find out more about bullx, instruments for innovation on
www.bull.com/extremecomputing

Bull
Architect of an Open World™