

Computing 2010: from black holes to biology

Declan Butler

By 2010, a click on the PC on your desktop will suffice to call up instantly all the computing power you need from what by then will be the world's largest supercomputer, the Internet itself. Supercomputing for the masses will trigger a revolution in the complexity of problems that are tackled, whole disciplines will go digital and, rather than spending time collecting their own data, scientists will organize themselves around shared data sets.

"PAL, I'm ready; set up that conference call." The teraflop PC in George's Chicago laboratory obeys. Detectors stir and projectors purr. Ceiling, walls and floor shimmer sapphire. The virtual-reality cavern ready, PAL connects to the Internet at ten gigabytes per second, beaming in three-dimensional avatars of Juliette in Paris and Günther in Munich. The three shake hands. The experiment begins.

Elsewhere on the Internet, PAL is busy interpreting and assembling data sets from research centres across the globe, while simultaneously negotiating CPU time from a supercomputer in New Mexico, and petaflop clusters of PCs in London and Prague. Operation accomplished. A kaleidoscope cloud of millions of glowing pixels suddenly coalesces in the centre of the laboratory, sculpting two floating black holes that swirl slowly around one another.

Koichi, the principal investigator, takes control of the input parameters, steering a collision of the black holes using a handheld screen from his bedroom in Tokyo. One of the most violent events in the Universe unfolds in fast-forward in front of the scientists, as the black holes whirlpool in the final plunge, merge and mushroom.

Within minutes, PAL has crunched the terabytes of output. "Back in the 1990s, it would have taken months," muses George to himself. The results scroll down the high-resolution

wall. An unequivocal first; a strong burst of gravitational waves in the last milliseconds of the collision, exactly as predicted by Einstein's theory of relativity.

The excitement builds. PAL relays even better news — that the bursts also have a specific waveform — to its counterpart machine at VIRGO, the European gravitational-wave detector in Cascina, Italy. VIRGO had come on-line in 2008, and its kilometres-long laser interferometers had amassed a haystack of data; but finding the needle — tell-tale vibrations a tiny fraction of an atomic nucleus in magnitude — had proved elusive. VIRGO now had a vital clue as to what to look for.

The scientists pace for what seems an eternity, to be interrupted by on-line gatecrashers from VIRGO, popping champagne. PAL's waveform template matches blips buried in the noise of VIRGO's massive data sets. The signal-to-noise ratio adjusted, the blips even yield a first guess at the mass of a pair of real black holes. PAL generates a virtual draft manuscript, which is passed around, quickly edited and passed back. The computer rapidly ranks the citation and media impact of the publication in the EMBIH global literature repository as well as the handful of journals that now exist, and — showing signs of intelligence — decides to e-mail the paper to Nature. All in a day's work in 2010.

Journalistic licence? Certainly. Fiction? Probably not. The core science of the black-hole collision, and much of the Internet technology, was demonstrated in June this year, when a team led by Ed Seidel of the Albert Einstein Institute in Potsdam, and Wai-Mo Suen of Washington University in St Louis, broke supercomputing records with a 140,000-CPU-hour run that produced almost a terabyte of data, using the 256-processor 'Origin2000' at the US National Center for Supercomputing Applications (NCSA).

Prototype systems emerging from the world's most advanced computing and networking test beds suggest that such high-end

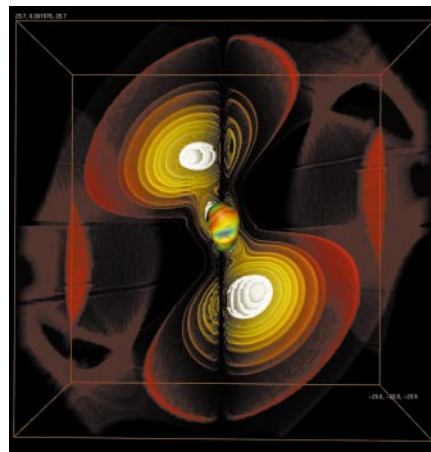
virtual reality, modelling and visualization will be routine over tomorrow's Internet. A proprietary virtual-reality theatre, CAVE, under development at the Electronic Visualization Laboratory of the University of Illinois at Chicago, is one of several systems for remote collaboration — or 'tele-immersion' — that are up and running. Nor will using such technologies require booking scarce run-time on supercomputers in advance; a click on the PC on your desktop will suffice to call up instantly all the computing power you need, from what by then will be the world's largest supercomputer, the Internet itself.

By 2010, today's top-end computing promises to be no longer the exotic pastime

of a privileged few, but the routine of the masses. A vision is emerging of a global computing grid, where computers of any shape and size are linked across the Internet to assemble giant distributed supercomputers. Scientists in every discipline will, from their desktops, have access to raw computing power that dwarfs today's most powerful supercomputers.

Seidel admits that simulations of black-hole collisions are still far from being sophisticated enough to provide realistic help to the planned European VIRGO (Variability of Solar Irradiance and Gravity Oscillations) and US LIGO (Laser Interferometer Gravity-wave Observatory) gravitational-wave observatories. But he predicts that they soon will be, as advances in computing power and algorithms enable such complex simulations to model reality for the first time.

The wide availability of massive comput-



Black-hole simulation showing a just-formed larger black hole that resulted from the collision of two smaller black holes. The surface of the newly formed hole is shown as a coloured surface in the centre of the simulation. This surface is semi-transparent, allowing one to see the two original holes inside. A burst of gravitational waves results from this collision, and these waves are shown as the red-yellow wisps emanating from the central region.

AEI/NCSA/UNIV. WASHINGTON


ing power to help tackle complex systems will prompt a profound change in science itself, predicts Ruzena Bajcsy, assistant director of the US National Science Foundation (NSF)'s directorate for computer and information science and engineering. She believes it will close the chapter of Cartesian reductionism in research, and usher in a new era where researchers will increasingly take on the understanding of complex dynamic systems, such as whole cells and the Earth. To advance the global grid, NSF is funding prototypes for scientific use — through the National Partnership for Advanced Computational Infrastructure, a consortium led by the San Diego Supercomputer Center — and for the population at large — through the National Computational Science Alliance, an NCSA-led consortium of over 50 research agencies and universities.

Conventional supercomputers will without doubt become major nodes on the grid at least at the beginning, and linking even a few of these would greatly increase the resources that could be brought to bear on individual problems. Petaflop supercomputers that are roughly 1,000 times faster than today's top teraflop (10^{12} floating-point operations per second) machines are expected soon after 2010 — and perhaps before if current research on technologies such as 'processor in memory' and the use of superconducting parts bears fruit.

But supercomputers are expensive. The world's most powerful supercomputer, Asci Red, at Sandia National Laboratories in New Mexico, which distributes tasks across an array of around 10,000 Pentium Pro chips, cost almost US\$60 million to build. A decade ago, Tom Sterling, who works at Caltech and NASA's Jet Propulsion Laboratory, thought that a cheaper way forward might be to simply string together lots of PCs. Sterling built his first 'Beowulf' cluster supercomputer by taking a then-barely-heard-of open-source operating system, Linux, and adding network drivers to it. The project worked, cost a tenth of the price of a conventional supercomputer and was also highly flexible — if more power was needed then more PCs were added. 'Avalon', a Beowulf at the Los Alamos National Laboratory, ties together 140 PCs and carries out 50 billion floating-point calculations per second for a construction cost of a mere \$300,000. The European Southern Observatory at La Silla, Chile, and LIGO have also chosen Beowulfs to meet their supercomputing needs in data processing

The 64-bit question

The Beowulf concept now looks set to take off in an even bigger way: on the Internet. Intel's latest chip, the IA64, or Merced — which is due to be shipped in 2000 — is more than just another faster Pentium. What makes Merced so special is that it is the first commodity PC chip to have a 64-bit architecture.



CAVE virtual-reality theatre.

Many scientific applications and high-resolution graphics require 64-bit floating-point accuracy — today's 32-bit PCs round numbers off too imprecisely for many tasks. The arrival of the 64-bit chip blurs further the differences between Beowulfs and conventional supercomputers. Being able to link up cheap 64-bit PCs for scientific purposes will open up cluster supercomputing to many more laboratories, while the movement will also be encouraged by the arrival of gigabit-per-second Ethernet local area networks.

The designers of tomorrow's Internet think they might be able to take the idea further. If you can combine the computing power of PCs over a local network, then why not also over the high-speed network that will be tomorrow's Internet? "Asci Red has 10,000 Pentium Pros; link up PCs on the Internet and you could link up tens of millions," says Larry Smarr, director of the National Computational Science Alliance. "We will have the accuracy we need for these large complex simulations but at the price of PCs".

Indeed, the characteristics of 64-bit chips, in principle, make them good candidates for such a feat. Although today's 32-bit PC chips can address 2^{32} other items, which in terms of memory represents 2 gigabytes, 64-bit chips can address 2^{64} , which in distributed computing terms means that one processor could in principle address the memory of every other computer on the planet. Link such PCs over a high-speed Internet and it would in effect become one

large computer, says George Lake, a scientist at the University of Washington and NASA project scientist for high-performance computing in Earth and space science. (The 128-bit chips on the horizon would allow the user to address a number equivalent to every atom in the Universe with space left over.)

"Think ahead ten years," says Smarr, "and you can see the supercomputer sort of dissolving into this vast planetary fabric of computing because all [computer technology] will be built on the same commodity processors, operating systems and interconnects."

Proof of principle for the use of PCs over the Internet comes from the Search for ExtraTerrestrial Intelligence screensaver SETI@home, a downloadable screensaver which has become the largest distributed computing project in history. Over 1 million users (including this author) have donated spare CPU time on their PCs and Macs to search for signals from extraterrestrials in the 35 gigabytes of data produced daily from the Arecibo radio telescope in Puerto Rico. The result is the planet's largest virtual supercomputer, running at 7 teraflops.

David Anderson, a computer scientist at the University of California, Berkeley, who leads the project, says that he now intends to use the same approach to tackle other supercomputing tasks including rational drug design, protein folding and high-resolution graphics. A web-based group, 'distributed.net', has also made a name for itself by using distributed computing over PCs to solve cryptography keys. But Sterling

remains cautious about the prospect of commodity PC chips becoming a major source of supercomputing power over the Internet, despite its extraordinary potential. He points out that the speed at which the memory is addressed — or latency — is crucial for many applications. At present, the fast interconnects needed are found only in dedicated supercomputers. The connections between clusters of PC operating over local area networks are still too slow or unreliable for applications requiring tight coupling, let alone clusters operating over the network.

Faster interconnects produced by Gigaset Inc. have produced tighter coupling in a 256-processor cluster unveiled in October by the Advanced Cluster Computing Consortium, an organization based at Cornell University which includes Microsoft, Dell and Intel. The project, based not on Linux but on Windows 2000, is a commercial bid to standardize reliable cluster computing systems for supercomputing in research and business.

Latency is a wider problem in computing. The time required to fetch something from the memory of say a Cray supercomputer — 300 nanoseconds — is not much less than half what it was 20 years ago. Ultimately, this means that significantly greater speeds will not necessarily follow from adding more transistors on a chip. Putting the central processing unit in the memory is the only feasible option, according to many computer scientists.

What seems clear is that as soon as the high-speed Internet becomes available, distributed computing in some form is going to come of age, probably involving a mix of supercomputers, clusters of PCs and workstations, and individual PCs. The backbone of the US research network currently runs at 2.5 gigabytes per second. But Douglas E. Van Houweling, president of the University Corporation for Advanced Internet Development (UCAID), which runs the US Internet 2 project, predicts that the network will run at terabyte levels by 2010, delivering to gigabit desktop connections.

Power hungry

Brute-force computing is badly needed. “With 100,000 sequences of 600 bases pouring out every day and needing to be screened against all others, we could easily use ten times the computing power right now just for sequencing,” says Craig Venter, chief executive officer of Celera Genomics Corporation. “If you do not have this kind of [high-end] computing capacity in biology, tomorrow you are going to get left behind.”

Demand for supercomputing power at the NSF is growing exponentially and will pass 10 teraflops by 2005, with most coming from new users. The National Computational Science Alliance has 1,536 processors for scientific supercomputing, but Smarr points out that with around 300 users a month, that

works out at a handful of processors per user. “The peer-review system to select projects is a polite term for ‘rationing,’” quips Smarr. Over half the top projects get rejected but Smarr is optimistic that over the next decade we will move from “an era of scarcity to an era of plenty.”

Attention is now turning to how best to make such distributed resources available to users. The possibilities are enormous. Imagine sitting at your screen in front of a simple Yahoo-type interface, customized to your personal research requirements. One click and databases around the world are scoured for the data you want, say spectral lines of several classes of stars. Little matter what format the image bank databases are in; the interface uses a standard format and automatically converts data having other formats. With the data collected and merged, another click and the interface negotiates with computers around the world to book in real time all the CPU time you need. Want a Fast Fourier Transform? One keystroke does it. If you wish to view the results in three dimensions, virtual reality or rotate them, click on a menu offering advanced visualization techniques.



Larry Smarr foresees supercomputers being superseded by a vast planetary fabric of computing.

And all translated into your own native language.

Such interfaces to a suite of tools written specifically to exploit the distributed database of the Internet and invisible to the user — or ‘middleware’ — are now the focus of development within many advanced Internet engineering projects such as Internet 2. By 2010 they will be a routine part of the infrastructure of working scientists, according to Rick Stevens, head of the mathematics and computer science division at Argonne National Laboratory.

Plug and play supercomputing

A big step in this direction is the recent release of Cactus, a novel software tool kit developed by Joan Massó and Paul Walker, at the Max Planck Institute for Gravitational Physics in Potsdam. It allows scientists to combine supercomputers over a high-speed network without having to know anything about the necessary advanced computing techniques. It is plug and play in that users simply take their own application (or ‘thorn’) — be it written in Fortran or C or whatever — and plug it into the flesh of the Cactus code. The Cactus automatically ‘parallelizes’ their program to run on virtually any computer system, from a portable PC to

distributed clusters of PCs or supercomputers. Cactus’s modules also allow booking of CPU time at remote centres, handle the output of terabytes of data, and also offer a suite of state-of-the-art three-dimensional visualization tools. The NSF has committed \$2.2 million to developing Cactus in the United States. The code is accessible to the public at <http://www.cactuscode.org>.

That is a sea change. In the past, supercomputing algorithms have generally been developed for specific problems, and therefore could not be used by others. Cactus will allow biologists, for example, to draw on decades of experience of their more mathematically inclined colleagues, physicists. “As the network and network software evolves I think the division between the PC on your desk and the assets the networks will offer will become steadily less apparent,” predicts Van Houweling. “Today, scientists mainly use the web to find information. In the future, accessing large-scale computing resources over the net will be the major use.”

Such ‘middleware’ is beginning to appear in biology. Even the most basic genomic analyses often require users to run a plethora of different software and data formats. NCSA’s ‘Biology Workbench’ and Baylor College of Medicine’s ‘Search Launcher’ are two prototype systems designed to replace these with user-friendly packages, along the lines of Microsoft Office. Behind the interface lies software that converts requests into the formats used by the various existing genome analysis tools, allowing users to screen multiple databases on the web as if they were one. But Andy Baxevanis, director of Computational Genomics at the National Institutes of Health’s National Human Genome Research Institute, which is itself about to release a similar system, warns against users becoming too dependent on interfaces as a “black box, where they stick sequences in, get the results out, and don’t know what the underlying method actually is”. Baxevanis argues that “there is no substitute for actually understanding the individual methods, whether you use them individually or in the context of a workbench”.

Collaborating in cyberspace

“By 2010, the relationship of scientists with their computers will change; they will see it as an information and computing fabric that they have customized for the work they are doing, and a window to their colleagues,” says Van Houweling. Greater collaboration will emerge naturally as the telephone, video and virtual reality become commonplace on the Internet. Smarr foresees “a persistent cyber café where you walk into a room with your world-wide colleagues.”

But there is more to collaboration than sharing a virtual tea-break with colleagues

NCSA, UNIV. ILLINOIS AT URBANA-CHAMPAIGN



Computing will be the biologist's number one tool in the future, predicts Craig Venter.

on the other side of the world. A decade ago, scientists studying the upper atmosphere tended to build their research agendas around their own particular instrument, be it a satellite or radar. This hardly made much sense given that the community was studying the same zone. Dan Atkins at the University of Michigan has led an Internet-based effort called the

Upper Atmosphere Research Collaboratory, which has united the community's panoply of instruments and data around common research programmes.

The nature of science is changing, says Lake. "Until now, astronomy has been about getting telescope time, looking at a point in the sky, getting the data, doing something with those data, keeping it proprietary and publishing a small result; but soon all the information is going to be available digitally." Disciplines, he reckons, will increasingly organize large-scale computational resources which are automatically called into play from your PC or workstation. The concept of scientists standing on each others' shoulders is in for a renaissance. Rather than spending most of their time collecting their own data, this will be shared, coded and made accessible to the whole community.

Going for grand challenges

If scientists have access to supercomputing power from their desktops through the Internet by 2010, it will allow them to take on much more ambitious research challenges. "I think we will have the kind of computing power and tools to build models of whole cells and organisms," says Van Houweling, "but these new tools that we are talking about are becoming available only now to ordinary scientists."

Most biologists at present are happy to look for the homologue of the particular gene or protein sequences they are interested in. But by using more powerful computers, it will be possible to start looking at the integration of information across the genome. "If we hope to understand biology, instead of looking at one little protein at a time, which is not how biology works, we will need to understand the integration of thousands of proteins in a dynamically changing environment," says Venter. This is the direction Celera will take next year once it has completed the sequencing of the human and mouse genomes. To do so, the company is building

what will be the world's largest supercomputer for biology, a 1,200-processor machine. "A computer will be the biologist's number one tool," predicts Venter. "The data sets are beyond the capacity of the human brain."

Gene Myers, who leads Celera's computational biology group, has cut the time it takes to piece together the *Haemophilus influenzae* genome from days to just 5 minutes. He is now turning to the task of developing sophisticated algorithms to analyse the human genome, beginning with whole-genome comparisons. He then intends to overlay and integrate gene expression data.

A pointer as to where computational biology is likely to go is a software package called E-cell, developed by Masaru Tomita from Keio University in Japan. The package, which can be downloaded from his website at <http://www.e-cell.org>, simulates basic cellular processes. Tomita has just completed a model of human erythrocytes and is building other models of human mitochondria, signal transduction for chemotaxis in the bacterium *Escherichia coli*, and gene expression networks in this bacterium's lactose operon (a collection of genes that are switched on when the bacterium is forced to feed on lactose sugar). Nourished properly, the 'Tamagotchi' erythrocyte reaches a steady state where metabolite concentrations compare well with those reported in real mammalian erythrocytes. Tomita is now inhibiting enzymes from the glycolytic pathway *in silico*, such as hexokinase, glucose-6-phosphate dehydrogenase, phosphofructokinase and pyruvate kinase, in a bid to throw light on the cellular metabolism of people suffering from hereditary anaemia, which is caused by deficiencies of these enzymes. Many are predicting that the fruitfly *Drosophila* will be the main target for complex computer modelling in developmental biology. Its genome will soon be available, and many hands make light work; there are some 6,000 *Drosophila* researchers out there.

There is no shortage of petaflop-scale problems in biology, from brain research to the modelling of organs. The principal limiting step in the development of advanced computing in biology is the lack of biologists who know how to do it. While physicists and astronomers have decades of experience of expressing mathematically complex problems and writing the software to carry them out, computational biology is in its infancy. So physicists will be able to translate more quickly improvements in hardware performance at the petaflop level into the scale or realism of grand challenges. One such challenge is to design a virtual sustained thermonuclear fusion reactor on a computer without having to build real prototypes, a problem estimated at 1 to 10 petaflops.

But even physicists are now running up against the difficulties caused by the sheer complexity of the software needed to model

such systems. US vice-president Al Gore has recently launched a plan for federal agencies to carry out a 'Digital Earth' survey, the stated goal of which is to create "a virtual representation of our planet that enables a person to explore and interact with the vast amounts of natural and cultural information gathered about the Earth". Modelling parts of the Earth such as the ocean or atmosphere is difficult enough; the thought of coupling them all into a working model of Earth is sending software engineers back to their drawing boards. "It makes the Sloan Digital Sky survey look easy," says one scientist (see *Nature* 399, 520; 1999).

Scientists like writing code, but software in the future is likely to be built using higher levels of abstraction. Much current inspiration is coming from *A Pattern Language*, a book written by architect Christopher Alexander in 1977, in which he proposes tackling large problems, like building a city, by building up to them. In this approach, all of the smaller-scale models, such as houses and streets, are solved first. Lake predicts that this concept of looking at the sorts of 'services and interfaces' between components of complex systems may be a key to better algorithm design.

Perhaps the most appealing aspect of distributed power is the prospect of an end to proprietary software, thanks to the input, not least, of researchers who should no doubt be concentrating on their science. The success of the Linux open-source operating system has shown that, when it comes to debugging codes, companies seem no match for the collective IQ of thousands of developers on the web.

But the Internet and computers are easier to invent than they are to predict. The biggest uncertainty is perhaps quantum computing. In 1999, Yasunobu Nakamura's team at the NEC Fundamental Research Laboratory in Tsukuba, Japan, demonstrated electrical control of a quantum data bit (or qubit), a possible first step on the road to the eventual creation of a solid-state quantum computer. Whereas today's computers store information as ones or zeros, quantum computers could use non-classical memory states that incorporate both values simultaneously, resulting in an exponential speed-up in computing time. One small quantum computer would easily outstrip all the computing power in the world combined. Although Dmitri Averin, an expert in quantum computing at the State University of New York, Stony Brook, believes that quantum computing will not become a practical tool within the next ten years, he predicts nonetheless that it will be a much bigger research area than today, and will yield exciting new discoveries in quantum mechanics.

Declan Butler is European correspondent at *Nature*, 3 rue de l'Arrivée, B.P. 264, 75749 Paris Cedex 15, France.