



Physics teaching
UK physicists bring the school syllabus up to date
p708



Laser target
The US Nova laser target chamber goes to France on loan
p709



Mars blow
The loss of the Mars Polar Lander gives NASA a headache
p711



X-ray success
European space scientists relieved at XMM launch
p713

IBM promises scientists 500-fold leap in supercomputing power...

Paris

Biologists are hailing IBM's US\$100 million project to build a 'petaflop' computer as likely to revolutionize our understanding of cellular biology. The computer, nicknamed 'Blue Gene', would be around 500 times faster than today's most powerful supercomputer.

Computer scientists say that the planned machine, details of which were revealed last week, is the first large leap in computer architecture in decades.

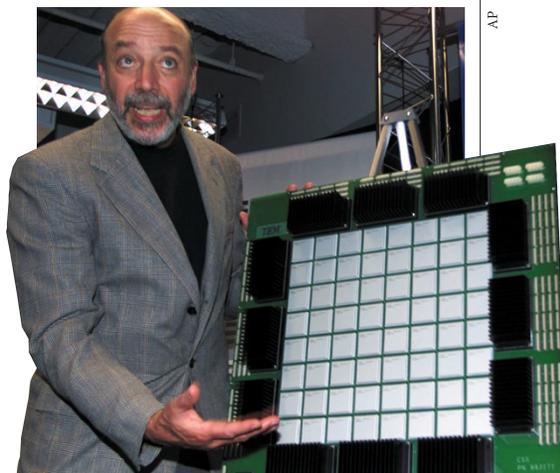
IBM will build the programme around the challenge of modelling protein folding (see below), with much of the research costs going on designing software. It will involve 50 scientists from IBM Research's Deep Computing Institute and Computational Biology Group, and unnamed outside academics.

But Blue Gene's hardware will not be customized to the problem and, if IBM's blueprint works, it will offer all scientific disciplines petaflop computers. These will be

capable of more than one quadrillion floating point operations ('flops') per second — around two million times more powerful than today's top desktops. Most experts have predicted that fundamental technological difficulties would prevent a petaflop computer being built before around 2015 (see *Nature* 402 supp., C67).

"It is fantastic that IBM is doing this," says George Lake, a scientist at the University of Washington and NASA project scientist for high-performance computing in Earth and space science. IBM is showing leadership by ushering in a new generation of supercomputers, he says.

The biggest technological constraints to building a petaflop machine have been latency — increasing the speed with which a chip addresses the memory — and reducing power consumption. A petaflop computer built using conventional chips would consume almost one billion watts of power. IBM reckons Blue Gene will use just one million.



On the board: IBM senior vice-president Paul Horn with a mock-up of a Blue Gene array.

Although processor speeds have increased exponentially, the time to fetch data from the memory of a supercomputer, 300 nanoseconds, is only slightly less than half what it was 20 years ago. Putting more and more transistors on a chip is therefore unlikely to lead to much greater speed.

"We set out from scratch, completely ignoring history, and thought how can we get the highest performance out of silicon," says Monty Denneau, a scientist at IBM's Thomas J. Watson Research Center in Yorktown Heights, New York, who is assistant architect of Blue Gene.

Arvind, a professor of computer science at MIT who is considered one of the top authorities on computer architecture, applauds IBM's approach. "It has made very big steps in rethinking computer architecture to try to do without the components that consume power; it has taken all these research ideas and pulled them together."

One innovation is IBM's use of a prototype technology that combines processor and memory on the same chip. This radically reduces the access time and bandwidth at which the processor can address the memory, as well as reducing power consumption.

Denneau claims that Blue Gene's chips get latency down to just ten nanoseconds. Blue Gene's bandwidth is such that it could

... and a chance to tackle protein structure

Paris

IBM's choice of protein folding as the focus of its project to develop the supercomputer Blue Gene is part of a strategy to position itself as a leader in computational biology, as the post-sequencing phase of the Human Genome Project approaches (see above).

The company has regrouped its life-science activities into a single unit focused on genomics and drug design. Sharon Nunes, senior research manager at IBM's Computational Biology Center at the company's Deep Computing Institute in New York, says that predicting protein structure will be a key use of genome data. "Using today's computers it would take a thousand years to fold a protein," she says.

"Blue Gene is tremendously

exciting," says Christopher Dobson, who works on protein folding at the University of Oxford. "Folding is how the cell converts information from DNA into biological function, how the cell targets proteins to particular parts of the cell, controls and localizes them, and is key to turning activity on and off," he says.

Some are sceptical. Computing power alone is not enough, says Geoff Barton, a scientist at the European Bioinformatics Institute in Cambridge. He argues that solving folding will require enormous progress in understanding the interactions between the atoms in a protein.

"There is no guarantee that if you find the minimum energy conformation it will actually be the

native protein conformation, because the potentials used to model the pairwise interactions are imperfect. There are also difficulties with modelling the true solvent environment of the protein."

Ajay Royyuru, a researcher at IBM's Computational Biology Center, says this will be a major thrust of the Blue Gene project. Dobson says the timing is right in that "many of the fundamental principles of folding are beginning to emerge".

Whatever the reservations about protein folding, few biologists doubt that, if it succeeds, the project's contribution to handling the mountains of data emerging from genome sequencing programmes will be revolutionary. **D.B.**

download the entire Internet — around 100 terabytes — in less than one second.

Arvind points out that, although a substantial improvement, this is not enough to allow uninterrupted traffic between the processor and the memory, and bottlenecks in processing would still occur. The conventional solution is to put a special high-speed memory, or cache, in front of the processor, to supply it with the most frequently requested instructions and data. These can then be accessed much faster than instructions and data in the main memory.

The problem facing designers of petaflop computers is that caches consume vast amounts of power. “IBM’s real innovation is to have done away with the cache,” says Arvind. To do so, it has turned to another prototype technology: multi-threading.

Each of Blue Gene’s one million processors can do eight tasks, or threads, simultaneously. If one set of threads is busy, the next instantly takes up the relay. “No-one has attempted multi-threading at this level,” says Arvind. “Current supercomputers just don’t have this feature.”

When built, Blue Gene will have one million processors — each capable of one billion operations per second (one gigaflop) — of which 32 will be placed on each chip. A board containing 64 of these chips would be capable of two teraflops, equivalent to Asci Red, the world’s most powerful supercomputer, at Sandia National Laboratory in New Mexico.



Chip ahoy: IBM vice-president Ambuj Goyal.

Asci Red distributes tasks across an array of about 10,000 Pentium Pro chips, and cost almost \$60 million to build. Eight of these boards will be placed in two-metre-high racks, and 64 racks will combine to give one-petaflop performance.

Denneau admits that supercomputers are facing stiff competition from clusters of PCs and workstations, which can deliver supercomputing power at a fraction of the cost (op. cit.). But he says that these are still too slow or unreliable for applications requiring tight coupling and low latency.

The cost of building Blue Gene will be a fraction of that of Asci Red, adds Denneau, thanks to a secret proprietary technology used to etch the silicon in the bulky chips. Most of the costs will be in devising parallel computing models and programming methodologies.

Blue Gene’s main drawback, says one scientist, is that it may be relatively short of memory. Arvind says the biggest unknown will be reliability. “To put together a petaflop computer in five years is a very big deal,” he says. “It may simply not be possible.”

Declan Butler

Germany drags its feet over demand for genome funds

Munich

A widely heralded strategy paper proposing significantly higher public investment in genomics research in Germany appears to have been put on hold. German scientists are concerned that the move could threaten efforts to catch up with other countries.

The paper was drawn up by the research ministry, and had been expected to be made public at the annual meeting of the German Human Genome Project in Munich earlier this month. But Wolf-Michael Catenhusen, secretary of state for research, says the paper will now be presented to the government in early summer next year.

The German research ministry has been giving mixed messages about the cause of the delay, which dashes hopes for more cash for genomics next year. Hans Lehrach of the Max Planck Institute for Molecular Genetics in Berlin, a spokesman for the German Human Genome Project, says the delay could be “disastrous” for Germany’s hopes of catching up in genome research after a late start.

The research ministry’s strategy paper was drawn up after a rare period of lobbying by the scientific community. This began a year ago at a meeting called by Ernst-Ludwig Winnacker, president of the Deutsche Forschungsgemeinschaft, Germany’s grant-giving agency for research in universities.

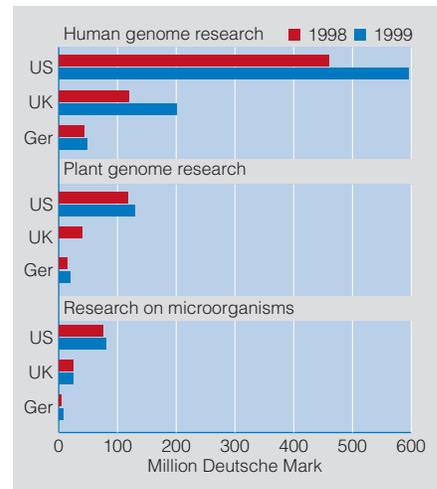
Ministry representatives at that meeting were dismayed by the lack of agreement between scientists on the best way forward for genome research. Various research organizations then prepared individual reports and reached a general agreement that funding should be increased by an order of magnitude if Germany was to compete internationally.

Proposed strategies for spending an enlarged budget were also convergent, focusing primarily on analysis of the function of human genes and development of model organisms for this, but also including some economically, environmentally or medically relevant plant and microorganism genomics.

These aims are reflected in an early draft of the ministry’s strategy document which recommends the expansion of the main genome research centres in Munich, Berlin and Heidelberg, and the creation of more focused scientific networks.

According to the draft, such a strategy would require “a new quality of state financing” — it refers to the large differences in the funding of genomic research compared with the United States, United Kingdom and France (see figure). And it should be complemented by industry, and by a reorientation of research institutes’ core funding.

Catenhusen, who had been confident in



Public funding of genome research. (Source: German research ministry.)

the summer that there would be no delay, agrees that the scientific community is now “speaking with one voice”, and says there is no dissent in the ministry about the strategy paper’s content.

But he now says the paper needs to wait until other decisions are made, such as how the Berlin-based resource centre should be restructured, and how national research centres plan to expand their own genome research through internal reorganization (see *Nature* 402, 450; 1999).

He remains bullish about the strategy paper’s big aims. It would be relatively simple to continue modest increases in genome research funding “to a level of say DM90 million [US\$47 million] per year, which would be a 40 per cent increase [since 1996]”, he says. “But this is not our strategy’s message; genome research funding should be DM250–400 million per year. We need time to prepare ourselves carefully for a round of persuasive discussion with the government.”

But many scientists are dismayed by what they see as delaying tactics. Lehrach says: “We lost a lot of time in the last ten years when other countries were moving slowly in genomics, but now genomics is moving very fast, so any months we lose in failing to bolster funds will be a disaster.”

While some scientists worry that the delays could reflect splits in the Social Democrat/Green coalition government over a genomics strategy that includes plant genomics, a more likely explanation is Germany’s economic squeeze. “If politicians are holding back at the moment, it is because of the unclear financial situation,” says one ministry official. “But the will to expand genomics research is very strong.”

Alison Abbott