

NEWS

The petaflop challenge

IBM's announcement on 26 June that it was about to smash the 'petaflop' speed barrier could herald a new era in computing. But unless the research and computing communities get their programming act together, they risk having few scientific applications that can take advantage of this huge increase in power, say experts.

Called Blue Gene/P, the first of the new high-powered computers should be operational next year, and IBM has already lined up potential customers in the United States, Germany and Britain. The largest planned configuration for the machines would run continually at 1,000 trillion floating point operations per second, or teraflops, and be capable of peak speeds of up to 3,000 teraflops (3 petaflops). That would make it between three and ten times faster than the machine that tops the latest TOP500 list of supercomputers, the IBM Blue Gene/L at Lawrence Livermore National Laboratory in California, which peaks at 360 teraflops.

"I'm very excited," says Ray Bair, head of the Argonne Leadership Computing Facility at the Argonne National Laboratory in Illinois, which has worked with IBM on developing the Blue Gene/P, and which will host a version of the hardware. "We're crossing a threshold," Bair adds, claiming that the increased power will at last allow researchers to build and run models in the way they have always wanted.

At prices of around \$100 a gigaflop, petaflop computers will start off in the \$100-million range. A great deal of the added speed that money will buy simply comes from more processors. Supercomputer designers began taking parallel processing seriously in the 1990s, but few machines have been designed to work with more than 10,000 processors. In 2005, Blue Gene/L, which is almost three times as fast as its nearest competition, marked a significant step up with 131,072 processors. A 3-petaflop Blue Gene/P will boast 884,736 — a multitude that brings with it problems as well as promise.

The rapid recent growth in supercomputing power — Blue Gene/L is as powerful as the whole of 2002's TOP500 list combined — has come mostly from increases in the performance of the computer's component processors. But a few years ago that process "came to a grinding halt", says Jim Sexton, head of Blue Gene applications at IBM's Thomas J. Watson Research Center in Yorktown Heights, New York. As the processors got faster and faster, they began suffering from disproportionate increases

in power consumption and heat output.

To combat this, chipmakers made a historic switch. Since 2004, they have concentrated on increasing the number of processors on a chip, allowing the speed at which individual processors operate to plateau. So the dual and quad chips now common in laptops, offering two or four processors, may well be upgraded to 128 or 256 cores by 2015. This means that the cheap Linux supercomputing clusters common in universities would have hundreds of thousands of processors, and dedicated supercomputers might have hundreds of millions. The processors will not necessarily be blazingly fast — Blue Gene/P's 850-MHz chips are little faster than a Pentium III from 1999. But with their numerical advantage they won't have to be.

This new reliance on parallel processing for increased performance means that companies from Microsoft to Nintendo will have to rethink their software — and so will scientists. A few scientific applications fall into the favoured subset called 'embarrassingly parallel problems'. Genome analysis using BLAST software to compare sequences and mass spectroscopy for proteomics are generally fairly easy to parallelize, says Leroy Hood, president of the Institute for Systems Biology in Seattle, Washington. Each processor can take on a specific task without much reference to what all the others are doing. But other sorts of problem, in which many of the calculations depend on other calculations being done elsewhere, are not so tractable.

"The scientific community is not very good at software development."

For the moderate levels of parallelism seen to date, it is possible to get by with the current practice of writing code, and designing models, in terms of linear sequences of instructions and then parallelizing once satisfied. To get the most out of massive parallel clusters and machines, that will no longer be an option. "Coding models running across as many as a million processors is a new challenge we have to meet," says Tim Palmer of the European



IBM

Centre for Medium-Range Weather Forecasts in Reading, UK, who is interested in petaflop machines for climate modelling. "We have no choice but to follow the hardware trends."

Scientists need to shift to thinking in parallel from the outset, designing hypotheses and code accordingly, says Horst Simon, associate lab director for computing at Lawrence Berkeley National Laboratory, California. He describes what is needed as nothing short of a "revolution in scientific programming". Such a revolution has been brewing for decades, but there hasn't been much storming of the barricades. "The high-performance computing community has been working on the parallel-programming problem for over 25 years. And frankly, we don't have much to show for it," wrote Intel researcher Timothy Mattson on his company's research blog last week. "On the software front, it's chaos."

**MIND-ALTERING URINE**

Alpha-male pheromones cause females to make brain cells.

www.nature.com/news



Parallel lines: the Blue Gene/P computer (left) will be up to ten times faster than the current record-holder Blue Gene/L.

Such challenges, say experts, will expose serious weaknesses in the capacities of the scientific computation community. “The scientific community is not very good at software development,” says Simon. He reckons his 20-year-old son, who writes gaming software for fun, is way ahead of most scientists in addressing the challenge of parallel programming.

“I’m amazed at what he can do just using open-source libraries,” he says. Although there are exceptions, such as high-energy physics and bioinformatics, many labs keep their software development close to their chests, for fear that their competitors will put it to better use and get the credit for the academic application of the program. There is little incentive to get the software out there, says Simon, and such attitudes plague development.

One underlying reason is that in academia, software professionals are not given due recognition. “Why would you work at a low salary at a university, where the academic hotshot gets all the credit,” asks Simon, “when you could work at a gaming company, be the hotshot and get stock options?”

Declan Butler

UK science reshuffled

In the wide-ranging shake-up of government that accompanied Gordon Brown’s arrival as Britain’s new prime minister, science got both a new home and a new minister in charge. Although the move had not been announced at the time *Nature* went to press, the new minister will be Ian Pearson, previously Minister of State for Climate Change and Environment.

In a complex game of departmental division and recombination, Pearson will take office not in the Department of Trade and Industry (DTI), which used to be home to the science brief, but in the newly formed Department for Innovation, Universities and Skills (DIUS). This department marks a fusion between parts of the old DTI and part of the old Department for Education, and will also contain a new Office of the Chief Scientific Adviser. The trade part of the DTI’s brief will now be shared by the Department for International Development and a new Department for Enterprise, Business and Regulatory Reform.

Pearson comes to the science brief having worked recently on launching a consultation process on the introduction of a cap-and-trade carbon-reduction scheme for up to 5,000 UK businesses. Welcoming the appointment, a spokesperson for the Royal Society said: “The government’s track record in climate change has been good recently and Ian Pearson has obviously played his part in that. If he brings that track record to science we’d be very much encouraged.” A strong emphasis on the environment at high levels in government seems assured by the appointment of David Miliband as foreign secretary; as Secretary of State for the Environment, Miliband was a strong voice on climate change.

Yet Phil Willis, chairman of the parliamentary science and technology select committee, expressed

concerns that environmental science is about to hijack science as a whole. “A danger of the environment becoming so dominant,” he told *Nature*, “is that other branches of science will be playing second fiddle.” To do justice to science as a whole, Willis said, Pearson will have to “get to grips with a portfolio that is absolutely massive”.

The creation of the DIUS, which in institutional terms will treat science as more closely linked to education than to enterprise, has received mixed reviews. “It makes sense to put science and higher education together,” says Colin Blakemore, shortly to step down as chief executive of the Medical Research Council.

But David Brown, chief executive of the Institution of Chemical Engineers, warns that “Innovation and science should be located in the same ministry as enterprise and skills. I fear we will continue to be held back by the traditional interdepartmental boundaries.”

Katharine Sanderson



Climate for change: Ian Pearson will take on the science brief.

M. LIMA/AFP/GETTY