

Subtler tests urged for supercomputers

Jim Giles, London

The supercomputer chart will soon have a new number one: data just released by IBM make its Blue Gene/L machine the world's fastest computer.

But some supercomputer specialists, while applauding IBM's technical prowess, question the significance of these rankings. They point out that the technique used to compare supercomputer performance is badly out of date. And they worry that the chart, which guarantees widespread publicity for whichever machine hits the top spot, is skewing US priorities in computer science funding.

IBM says that Blue Gene/L, which is currently on company premises in Rochester, Minnesota, prior to shipment to the Lawrence Livermore National Laboratory in California, can perform 36,000 billion calculations a second. The result has already been submitted to the organizers of the TOP500 list of supercomputers, who will publish their next chart in November. Blue Gene/L is expected to run about ten times faster next year, when all its 130,000 processors are installed.

Alan Gara, the machine's chief architect, says that US government officials are likely to be as pleased with the milestone as IBM is. After the Earth Simulator, built by NEC in Japan, hit the top spot in 2002 (see *Nature* 416, 579–580; 2002), Congress began to pay renewed attention to supercomputer funding. Politicians “want to show that the United States is competitive”, says Gara.



Gene genie: IBM's prototype can perform up to 36,000 billion calculations per second.

Yet the researchers who create the chart say that its contents should not be taken too seriously. Machines are ranked by the time it takes them to solve a set of linear equations, a test known as the Linpack benchmark. The test is a good measure of the speed of a computer's processors, says Erich Strohmaier, a computer scientist at Lawrence Berkeley

National Laboratory in California, who helps to compile the TOP500 list. But the test is less sensitive when it comes to assessing the speed at which the processors can communicate with each other — a crucial factor in actual performance, he says.

Using more than one benchmark could help, suggests Dale Nielsen, a theoretical physicist at Lawrence Livermore. A test that involves solving another type of equation, known as a fast Fourier transform, would be a useful addition, adds Strohmaier.

But even if the chart provided a perfect measure of hardware power, it still wouldn't tell the full story. Computers are useless without software, and some scientists feel that the focus on building ever faster machines distracts from the real bottleneck for heavy-duty computer users: the shortage of code that can fully exploit supercomputers' power.

Rick Stevens, a computer scientist at Argonne National Laboratory in Illinois, estimates that some \$500 million a year is spent on funding scientific software development in the United States. But this cash is spread around every scientific discipline, he says, and some areas still lack adequate resources.

At present, Stevens says, researchers often find that code they want to use for a particular application is not optimized for the machine to which they have access. He adds that there has been little investment in adapting code to run efficiently on the fastest supercomputers. ■

Tardy earthquake excites California geophysicists

David Cyranoski

For a group of scientists in the little town of Parkfield, California, the Earth finally moved — when an earthquake of magnitude 6 hit last week. After lying in wait for a major rupture in this part of the San Andreas fault for more than 20 years, scientists in the area pounced on their huge array of monitors to examine the data.

The earthquake, which caused little damage and no deaths, will provide a wealth of information for researchers. But it has also thrown up questions about whether major earthquakes come in predictable periods or with predictable characteristics.

Scientists studying Parkfield in the late 1970s hypothesized that an earthquake would hit the area once every 22 years. The current earthquake, striking after a 38-year gap, might make them recalculate. It does seem to add to the evidence that this earthquake zone may be “quasi-periodic”, says David Jackson, an earthquake researcher at the University of California, Los Angeles. But he generally

doubts the idea that certain places have earthquakes with specific characteristics.

Previous Parkfield earthquakes have led scientists to expect foreshocks of up to magnitude 5 striking the area about 17 minutes before the main earthquake. Instead they got aftershocks of magnitude 5. The earthquake also moved north to south along the fault, the opposite of what normally occurs in this region. Showing up late, moving backwards and speaking out of order, the delinquent earthquake is posing a few problems. “These are interesting phenomena, but we don't yet have any explanation,” says geophysicist Steve Hickman of the US Geological Survey.

Because of the regularity expected of Parkfield's earthquakes, the area has been intensely monitored for decades. Instruments are used to measure details of structural changes deep below the surface and the movement of the fault. “The data will be copious and sweet,” says Jackson.

The earthquake did not come at a great

time for researchers working on the San Andreas Fault Observatory at Depth (SAFOD), also in the Parkfield area. This project aims to drill into the heart of an earthquake-generating zone about 5 kilometres away from last week's rupture to study the fault itself during an earthquake. The project reached a drill depth of 3 kilometres two weeks ago, but seismometers, which would have given unique data, have not yet been installed in the hole. And the project will not drill through the fault itself until next year.

A second, pilot hole a few metres away did have some seismometers in place. But this summer, workers accidentally snipped off 25 seismic stations from the bottom of a string of 32. They have not yet been replaced.

Last week's earthquake has got SAFOD project leaders excited, however. Their part of the fault line is expected to produce smaller earthquakes every two to three years, so they have plenty of time to see more earthquake activity. ■