Record Speedups for Parallel Processing

In computer circles, the full potential of parallel processing has always seemed out of reach. In theory, 1,000 processors working simultaneously – with each taking on a small portion of the total computation – could tackle a problem 1,000 times faster than one processor plodding through the problem step by step as most computers do now. In computer lingo, the 1,000-processor machine would then have a "speedup" of 1,000. This would be akin to constructing a building with 1,000 workers rather than just one.

But most computer scientists have been slow to harvest the fruits of this division of labor, because they've believed that such ideal speedups – equal to the number of processors used – were unattainable. This week researchers at Sandia National Laboratories in Albuquerque, N.M., shattered the skepticism by announcing that they have achieved record speedups of more than 1,000 on a 1,024-processor computer. By showing that it is possible to run a parallel machine at near-100 percent efficiency, they have ensured a prominent place for parallel processing in the future of computing.

This future had been in doubt because of the realization that what is gained by dividing up the computation is partly offset by the time it takes for processors to talk to one another and to perform "serial" tasks that simply can't be divvied up. The long-held belief has been that these chores will bog down a parallel machine so much that even with a billion processors, speedups of only 100 are the best anyone could hope for.

As a result, "there's been a psychological barrier to working with thousands of processors on a single problem," says Sandia computer scientist Robert E. Benner.

Now Benner and Sandia colleagues John L. Gustafson and Gary R. Montry have broken through this barrier by carefully crafting computer algorithms designed to get the most out of a parallel processing machine. They ran these programs on a recently acquired, state-of-the-art "hypercube" computer made by NCUBE Corp. of Beaverton, Ore. With the hypercube architecture, processors are connected as if they were sitting at the corners of cubes that fit inside of one another.

The researchers first considered problems involving a few thousand equations – the largest size that could be handled by a single processor – and then divided the problem among increasing numbers of processing units. By judiciously arranging their programs to maximize parallel tasks and minimize serial ones, they achieved speedups of 502 to 637 on the 1,024-processor hypercube. This means that a problem that took 30 hours to run on one processor was completed in about 3 1/2 minutes with the hypercube.

But Benner and his colleagues also realized that they could use the machine most efficiently if they expanded the problem size and complexity while they added more processors. In problems involving wave propagation, mechanics and fluid flow, they achieved speedups of 1,020, 1,019 and 1,011 respectively. These problems account for about half the kinds of scientific and engineering problems that normally concern Sandia researchers. One of the scientists' next projects is to see if other types of problems yield to their parallel processing approach as well.

"These guys did an outstanding job," says Alan Karp, a physicist based in Palo Alto, Calif. "They've shown that you can get almost all the speed that's [available to the machine]."

In so doing, Benner's group was the first to meet a challenge issued by Karp in 1985, who says that at the time there had been a lot of talk about building computers with 1,000 or 10,000 processors, but that no one had shown that these machines would be able to do anything useful. To spur the development of multiprocessors and their software, he challenged computer scientists to demonstrate a speedup of at least 200 on a general-purpose computer. "I didn't think anybody would [meet the challenge] so soon," he says.

With their speedups of 502 to 637 from the fixed-sized problems, Benner and his colleagues were also the recipients earlier this month of the first Gordon Bell Award, which was established to acknowledge important contributions to parallel processing applied to real problems. According to Karp, who was a judge in the Bell competition, second place went to a research consortium that achieved speedups of 458 on a 512-processor machine and 39 with 127 processors. The speedups of all remaining entries were 16 or less.

For the near term, the Sandia work shows that multiprocessors can solve problems as fast as current Cray supercomputers and other supercomputers that contain only a few processors at most – but at about one-tenth the cost, making supercomputer power accessible to more people. And for the distant future, it paves the way for succeeding generations of parallel computers that may contain hundreds of thousands of processors. "By developing more of these massively parallel applications in the future," says Benner, "we're preparing for the day when we'll have a truly awesome machine to run them." *– S. Weisburd*

Fermat's last theorem: A promising approach

The end of a centuries-long search for a proof of Fermat's last theorem, one of the most famous unsolved problems in mathematics, may at last be in sight. A Japanese mathematician, Yoichi Miyaoka of the Tokyo Metropolitan University, has proposed a proof for a key link in a chain of reasoning that establishes the theorem's truth. If Miyaoka's proof survives the mathematical community's intense scrutiny, then Fermat's conjecture (as it ought to be called until a proof is firmly established) can truly be called a theorem.

Miyaoka's method builds on work done by several Russian mathematicians and links important ideas in three mathematical fields: number theory, algebra and geometry. Though highly technical, his argument fills fewer than a dozen manuscript pages – short for such a significant mathematical proof. Miyaoka recently presented a sketch of his ideas at a seminar at the Max Planck Institute for Mathematics in Bonn, West Germany.

"It looks very nice," mathematician Don B. Zagier of the Max Planck Institute told SCIENCE NEWS. "There are many nice ideas, but it's very subtle, and there could easily be a mistake. It'll certainly take days, if not weeks, until the proof's completely checked."

Fermat's conjecture is related to a statement by the ancient Greek mathematician Diophantus, who observed that there are positive integers, x, y and z, that satisfy the equation $x^2 + y^2 = z^2$. For example, if $x = 3$ and $y = 4$, then $z = 5$. In fact, this equation has an infinite number of such solutions.

In the 17th century, French amateur mathematician Pierre de Fermat, while reading a book by Diophantus, scribbled a note in a margin proposing that there are no positive-integer solutions to the equation $x^n + y^n = z^n$, when n is greater than 2. In other words, when $n = 3$, no set of positive integers satisfies the equation $x^3 + y^3 = z^3$, and so on. Then, in a tantalizing sentence that was to haunt mathematicians for centuries to come, Fermat added that although he had a wonderful proof for the theorem, he didn't have enough room to write it out.

Later mathematicians found proofs for a number of special cases, and a com-