


Supercomputers

Intel's hypercube architecture

Washington

INTEL, the computer-chip manufacturer, is coming out with the first of what it hopes will be a new generation of supercomputers, inexpensive machines that offer high computational power by hooking many small computers together in parallel operation. The Intel machines, which will initially be offered in configurations consisting of 32, 64 or 128 separate computers, or "nodes", are to be aimed specifically at scientific users, particularly those tackling large-scale



A hypercube connects $N = 2^n$ small computers, called nodes, through point-to-point communication channels in the Cosmic Cube. Shown here is a two-dimensional projection of a six-dimensional hypercube, or binary 6-cube, which corresponds to a 64-node machine.

modelling problems in areas such as fluid dynamics and circuit simulation.

The chief attraction of the Intel computers is certain to be their price, which will range from \$150,000 for the 32-node model to \$520,000 for the 128-node. The largest has a peak capacity of 10 million floating-point operations per second (MFLOPS), or about 25 times that of Digital Equipment Corporation's VAX 11/780, which costs about \$225,000. Perhaps the most familiar supercomputer, the Cray-1, costs \$5-10 million. Intel has announced that its computers will be available this summer.

Unlike the other parallel-processing supercomputers on the market, the Intel machines break totally with the conventional computer architecture based on executing a series of instructions in sequence. Each node consists of a separate central processor, floating-point processor and 512-kbyte random-access memory. Memory is not shared between the separate nodes. Each node can communicate with those adjacent to it in a hypercube layout (see diagram). In essence, the idea is to increase performance by stringing together many cheap processors, taking advantage of the recent advances in Very Large Scale Integration technology that have sent prices plummeting, rather than spending more and more on ultimately diminishing returns in attempting to improve performance of individual processors.

Intel's problem with this, its first entry into the marketing of complete computer systems, is that people do not know how to program highly-parallel processors. Although the company says it is encouraging third parties to develop applications software, it is offering only an operating system initially. Conventional supercomputers, such as the Cray, allow users to write their programs in the ordinary sequential manner; smart Fortran compilers (for example) are then employed to pick out the steps that can be executed simultaneously. The Cray-type parallel processor operates most efficiently when tackling, for example, vector arithmetic operations; its parallel processors can simultaneously perform the requested operation on each element of the vector. The inherent disadvantage of this approach, however, is that the overall execution is still sequential, and the high speed gained in vector operations inevitably runs into a "bottleneck" when single scalar operations are reached during the program.

To program the Intel-type machine successfully, the problem has to be split into independent pieces that each node can tackle simultaneously. In fluid dynamics problems, for example, each node can be assigned one sector of coordinate space; since the underlying physical laws involve only nearest-neighbour interactions of particles, each sector can be computed independently; the only communication between nodes that is necessary is the result at the edge of each node's sector.

In dealing with such problems, each node is used to its fullest, unlike the Cray-type machine which operates in practice only at 5-20 per cent of its peak capacity. (The Cray-1 has a peak capacity of about 160 MFLOPS.)

In justifying its somewhat daring move, Intel points to a potential market for such "near supercomputers" of \$100 million now, reaching ten times that in 1990. Supercomputers now account for \$300 million of the scientific computing market, estimated to be at least \$2,000 million. Intel clearly also has its eyes on the National Science Foundation's plans to support the establishment of supercomputer centres and to purchase time on existing machines in response to a considerable perceived demand for scientific supercomputing.

The Intel machines are based on the design developed by Charles Seitz, a researcher at California Institute of Technology. Intel is licensing the design from Caltech.

Stephen Budiansky

Correction

THE article "Corporate plans all the rage" (*Nature News* 31 January, p.341) incorrectly states that the Ordnance Survey is the responsibility of the Ministry of Defence rather than the Department of the Environment. □

Japan Prize

First awards go overseas

Tokyo

JAPAN'S answer to the Nobel prize was launched on 15 February with the announcement of the first two Japan Prize laureates. Honoured by the Science and Technology Foundation of Japan were Ephraim Katchalski-Katzir of the Weizmann Institute of Science, Israel, and John Robinson Pierce, professor emeritus of Stanford University in the United States.

The Science and Technology Foundation is itself the creation of its president, Konosuke Matsushita, founder of the vast Matsushita Electrical Industries, Japan's fifth biggest company. Matsushita is well known as a philanthropist and even philosopher; he established the spiritual and educational PHP (peace and happiness through prosperity) movement that has gained enormous popularity in Japan. Although Matsushita donated 3,000 million yen (\$12 million) to establish the Prize, he has stayed very much behind the scenes since — the foundation is officially a non-profit foundation endorsed by the Cabinet for the awarding of medals of honour. The foundation's funds have, however, made it possible to give a prize of ¥50 million (\$200,000) with the medal. The Japan Prize thus rivals the Nobel prize financially; the intention is that with the passage of time it will also rival it in prestige.

The Japan Prize's aim is a little different from that of the Nobel prizes: it is intended for those who have achieved breakthroughs that can clearly be put to work in new ways that will "contribute to world peace and prosperity", rather than at discoveries in basic science. Perhaps the prize will help in some way to overcome Japan's inferiority complex that the development of original new technology is in some sense a less creative endeavour than research in pure science.

For this first year, three target areas — biotechnology, medical engineering, and information and communication — were selected and nominations of scientists in these fields sought from well-known academics around the world. Some 432 recommendations for 291 candidates were received, roughly two-thirds of them from outside Japan. A 24-member committee, split into sub-groups for the three areas, then chose just two laureates.

Considerable surprise was expressed (by members of the Japanese press, at least) that no Japanese was chosen for the prize even though Japan has made a considerable contribution in the three areas. The simplest explanation, other than plain modesty, seems to be that the principal Japanese workers in these areas were all members of the selection committee.

In the biotechnology category, Ephraim Katchalski-Katzir was awarded the prize