

Preface

[Lin Chao](#)

Editor

Intel Technology Journal

This Q1'98 issue of the *Intel Technology Journal* focuses on Intel's tera-scale supercomputer and research on multithreading software libraries for applications.

On June 11, 1997, the Intel supercomputer, containing over 9,200 Pentium® Pro processors, set the world's fastest computing record. Using the industry standard Linpack measurement method, the system calculated 1.34 trillion operations per second (teraflops) making it faster than a speeding bullet. By the time a bullet travels one foot, the computer will have completed 667 million calculations. The supercomputer is used for important scientific simulations such as the effect of a kilometer-wide comet striking the Atlantic Ocean. It will also be used to ensure the safety, reliability and effectiveness of the U.S. nuclear stockpile through computer simulation instead of nuclear testing.

The contract to build this supercomputer was awarded to Intel under the U.S. government's Accelerated Strategic Computing Initiative. It is a joint development of the Department of Energy, Sandia National Labs in New Mexico, and Intel. Besides its 9,200 Pentium Pro processors, it has 573 gigabytes of system memory and 2.25 terabytes of disk storage. It weighs about 44 tons and has 86 cabinets taking up 1,728 square feet.

The Intel TFLOPS supercomputer itself is the subject of four out of the five papers in this Q1'98 issue. The first paper gives an overview that includes a look at system architecture and how 4,536 compute nodes are connected into a single massively parallel supercomputer. The second paper looks at how high performance is achieved with coding and system parallelism. The third paper describes managing and optimizing large-scale parallelism from the point of view of an operating system, and the fourth one looks at the design of the system management environment.

Finally, the fifth paper in this issue describes research in a parallel processing run-time library for applications supporting loop-level parallelism, task-level parallelism and nested parallel threads.