Intel Supercomputer Users' Group
1992 ANNUAL USERS' CONFERENCE

October 4-7, 1992
Harvey Hotel
Dallas, Texas

Foreward

This proceedings document is being published to promote the exchange of information among Intel Supercomputer Users' Group Members. Contained herein are abstracts and/or copies of the presentations and papers presented by users at the 1992 Annual Users' Conference in Dallas, Texas. It is hoped that the information contained in these proceedings will be of use to you, as an Intel supercomputer user, and will encourage you to participate in and benefit from Users' Group membership. If you have questions about one of the presentations, or would like more information about the Users' Group, please call the Intel Supercomputer Users' Group Administrator, JoAnne Wold, at (503) 531-5322 *joanne@ssd.intel.com.*

Touchstone Delta Update
(extended abstract)

John K. Salmon
206-49 Caltech
Pasadena, CA 91125
email: johns@ccsf.caltech.edu

Paul Messina
158-79 Caltech
Pasadena, CA 91125
email: messina@ccsf.caltech.edu

1 What is the CSCC?

The Concurrent SuperComputing Consortium (CSCC) consists of vendors, federal agencies, research laboratories, and universities. Its purpose is to use leading-edge, scalable MIMD computers and integrate and adapt new hardware and software technologies for use on large-scale technical computing problems.

In addition, several teams of experts (from consortium partner institutions and industry) have been formed to address specific issues such as fast I/O, high-performance file systems, and programming environments.

An important aspect of the CSCC approach is the strategy of allocating large amounts of time to relatively few, very ambitious computational projects. This enables unprecedented computations to be undertaken and yields both new results and reveals strengths and weaknesses of the computing systems.

The CSCC demonstrated the ability to use massively parallel processors (MPPs) remotely. This is an issue because some people perceive software on MPPs to be so crude and shaky that productive remote use is not possible.

2 What is the Delta?

The Intel Touchstone Delta is the CSCC's first large machine. It was delivered to Caltech in May 1991. It consists (effectively) of 512 processing nodes (all Intel i860) and 64 auxiliary nodes which carry out specific functions like I/O control, ethernet management, handling user logins, and HiPPI interfacing. The auxiliary nodes have an assortment of Intel 80386 and i860 CPU's and various complements of external devices and memory. The Delta has 90Gb of online disk, as well as connections to the CSCC fileserver through HiPPI and ethernet. Users typically access the delta via rlogin or rsh.

The 512 processing nodes form the computational core of the machine. They are arranged in a 16x32 mesh (without toroidal connections). Each has 16Mb of memory and four communication channels nominally rated at 12Mb/sec. (see the talk by Rik Littlefield in this section). The machine has a nominal performance of 33Gflops/sec (double precision). In practice it is rare for performance to exceed 15Gflops/sec due to difficulties associated with keeping the parallel adder and multiplier pipelines on the i860 chips constantly full. Nevertheless, the Delta was the "worlds fastest computer" for well over a year.

3 Scientific advances using the Delta

The CSCC's goal is to advance the state of the art in parallel computing by making machines of unprecedented performance available to the scientists at member institutions. The number and significance of scientific results achieved using the Delta is a measure of how successful we have been in attaining this goal. Three of five Gordon Bell Prize finalists (Cwik & Patterson, Jones & Plassman, Warren & Salmon) in 1992 did their work on the Delta. In addition, the Delta is "saturated", i.e., there is very little idle time. This is noteworthy because many (but not all) previous parallel computers have been lightly used.

Over 300 users from over fifty institutions have worked on the delta. A small sample of the scientific achievements of the Delta in its first year include:

- processing of Venus Magellan image data (3-dimensional perspective rendering) at the rate of four frames per second, thus bringing within reach real-time image processing (Groom, Curkendall, Li; JPL)
- solution of an electromagnetics problem that involved 48,000 dense complex linear equations, using out-of-core methods and sustaining a computation rate of ten gigaflops per second (Cwik, Patterson; JPL)
- performed the largest direct numerical simulation of the time-dependent compressible Navier-Stokes equations to date. The Delta will permit three-dimensional compressible turbulence simulations to be conducted for turbulence Reynolds numbers a factor of five larger than what has been achieved on a CRAY 2 (Eidson; Langley, Erlebacher; ICASE)
- three-dimensional simulations of semiconductor devices for problems with up to five million grid points (which requires the solution of matrices with ten million variables) have been solved successfully and revealed that coarser resolutions would yield results for these problems that were wrong by over 40simulations will help advance the design and manufacture of electronic circuits (Dutton, Goosens, Wu; Stanford)
- computation of ab initio full configuration interaction (CI) electronic structure models, including the largest full-CI calculation yet completed (Harrison; Argonne National Laboratory, Stahlberg; Pacific Northwest Laboratory)
- derived a phylogenetic tree of 473 bacteria, a result that might yield insight into the origins of life and evolution of life forms. Previous results were only able to deal with approximately twenty bacteria (Overbeek, Hagstrom; Argonne, Olsen, Woese; University of Illinois at Urbana)
- computation of a highly-resolved flow over a three-dimensional aircraft configuration in under ten minutes. The computation involved a mesh with 804,056 points and 4.5 million tetrahedra (Das, Mavriplis, Saltz, Gupta, Ponnusamy; ICASE/NASA Langley)
- by far the largest N-body computations (17.2 million bodies) ever attempted, used in astrophysics research on the evolution of the universe (Salmon, Caltech; Warren, Los Alamos)