High Performance Computing for Large Eddy Simulations: Today and Tomorrow

N00014-90-J-4120

Justin Rattner
Intel Fellow
Intel Corporation
Supercomputer Systems Division
December 1990

intel®

High Performance Computing for Large Eddy Simulations: Today and Tomorrow

Section **Page**

Figures

High Performance Computing for Large Eddy Simulations: Today and Tomorrow

Justin Rattner
Intel Fellow
Intel Corporation
Supercomputer Systems Division

Introduction: Performance Needs for LES

Turbulence, said to be the last unsolved problem in classical physics, is among the most important areas of research facing scientists and engineers today. The study of turbulence is essential to many of the Grand Challenges of Science identified by the Federal High Performance Computing Program. Grand Challenge problems such as understanding global environmental change, designing efficient combustion systems and developing controlled nuclear fusion all require the ability to simulate turbulence realistically. The economic impact of turbulence studies affects industries ranging from aerospace and automotive engineering to oil, medicine and electronics.

Turbulence studies, including large eddy simulations, are also among the most computationally demanding research areas. The study of vehicle signature calls for sustained performance in the 10 GigaFLOPS range, and problems such as vehicle dynamics, ocean circulation, viscous fluid dynamics and climate modeling will require sustained performance of a TeraFLOPS (one trillion floating point operations/second). Lacking these levels of performance, LES researchers are forced to parameterize their simulations, reducing a simulation's complexity to fit the available computational power.

As higher levels of computing power become available, LES researchers will be able to increase the complexity of their models, running simulations with higher spatial and temporal resolution and more accurately representing the physical processes involved. Increased computing power will also make it possible to reduce the reliance on subgrid scale models and to move the domain of subgrid scale models to smaller scales, where confidence about predicted behavior is greater. In short, more powerful computers will enable scientists and engineers to solve current problems faster and more cost-effectively and to begin addressing problems that are beyond the scope of current technology. (See Figure 1. Computing performance requirements for Grand Challenge problems.) In both cases, the results will serve to validate the underlying assumptions of LES research and to extend the utility of LES methods to a broader range of applications.

Parallel supercomputers are the most promising means of obtaining the high levels of computing performance needed for large eddy simulations and other computational approaches to studying turbulence. While conventional supercomputers face limits of physics and thermodynamics that make performance improvements increasingly difficult to attain, microprocessor performance will double every two years through the end of this century, according to Intel studies undertaken for MIT's Technology 2000 project. By exploiting this rapidly rising performance curve, massively parallel, microprocessor-based computers promise

TeraFLOPS performance well before the close of this decade. (Figure 2. Parallel supercomputer performance.Today's parallel supercomputers, such as Intel's iPSC/860, provide a cost-effective means for researchers to obtain exceptionally high computing performance while ensuring compatibility with future performance gains.

Parallel Supercomputing from Intel

Intel Corp., which invented the microprocessor, is also the leading developer and manufacturer of massively parallel supercomputers. More than 250 iPSC /860 and iPSC/2 systems, developed by the company's Oregon-based Supercomputer Systems Division (iSC), have been installed worldwide. Intel supercomputers are used for LES and CFD applications at a wide range of sites, including MIT, Princeton University, NASA AMES, NASA Lewis and the University of Wisconsin.

iSC is also the developer of the world's fastest computer: the Touchstone DELTA System, which will be installed in Spring, 1991 at the California Institute of Technology. The DELTA System, which provides peak performance of 32 GFLOPS, will e a resource of the Concurrent Supercomputing Consortium, which comprises Caltech, DARPA , NASA, the Argonne, Lawrence Livermore, Los Alamos, Oak Ridge and Sandia National Laboratories, and six other organizations. The system will be used for turbulence studies and a number of other problems.

The iPSC family's maturity, high performance, flexible architecture and rich software environment, along with Intel's ability to offer extensive consultation and support, make the systems well suited for large eddy simulations and many other computationally intensive applications.

Scalable Performance. The latest iPSC computer, the iPSC/860 is built around Intel's i860R floating-point microprocessor, whose million transistors, 64-bit bus, 40 MHZ operation and superscalar architecture leapfrogged other RISC processors. The iPSC/860 computer scales from 8 to 128 i860-based processors allowing users to make a relatively modest initial investment and add processors and memory as their needs expand. Since each i860 processor delivers peak performance of 60 MFLOPS (double precision), peak system performance scales from 480 MFLOPS to 7.6 GFLOPS.

Concurrent I/O. Computationally intensive applications such as LES generally involve large data sets and require extensive data storage capability. The iPSC/860 offers scalable I/O with 1 to 127 80386-based I/O nodes. Intel's Concurrent File System increases ease of use by providing transparent I/O concurrency: users can create a single file as large as the total system disk capacity, or hundreds of files across many disks, all without concern for the physical location of any file, block or directory, and all accessible from any processor or node.

Message-Passing Internal Communications. In a multiprocessor computer, communication among nodes is a critical determinant of system performance. The iPSC/860 uses Intel's proprietary Direct-Connect m routing, a high- performance, message-passing system that dynamically creates hardware communication circuits between communicating nodes. Messages are transmitted at a bi-directional rate of 5.6 Mbytes/second, and total communication bandwidth scales with the number of computing nodes. In addition, because message-passing is a straightforward and relatively intuitive communication paradigm, the iPSC/860 simplifies the task of software development.

Multicomputer Architecture. The distributed memory, Multiple Instruction Multiple Data (MIMD) multicomputer architecture used in Intel's parallel computers is inherently amenable to the domain decomposition approach that characterizes compressible turbulence simulations and many other LES programs. Multicomputers also offer higher level parallelism, increased flexibility and ease of programming compared Single Instruction Multiple Data (SIMD) data parallel machines.

Flexible Networking. Intel supercomputers are designed for integration into a networked computing environment and can connect to IBM and VAX computers. The iPSC/860 provides industry standards such as Ethernet networking, TCP/IP software and the X Window System.

Software and Applications. System software includes UNIX System V, Release 3.2 on the System Resource Manager and an efficient kernel operating system, NX/2, on each node. iPSC computers also provide a broad set of tools for developing and porting parallel software, including libraries of linear algebra subroutines, matrix solvers, parallel CASE tools, optimizing C and FORTRAN compilers, an interactive parallel debugger and parallel performance analysis tools. A number of computational fluids packages, such as FL087 and Nekton, are supported.

Robustness. Traditional supercomputers typically are environmentally sensitive and need special electrical wiring and plumbing (for system cooling). Intel supercomputers are air-cooled and conventionally powered, which significantly lowers the systems' life-cycle costs. The iPSC/860 averages a highly reliable 96 to 98 percent up-time.

Intel Support. As with any still-maturing technology, supercomputer customers are best served by large, stable vendors who can make the large-scale, long-term commitment of resources needed to solve the complex challenges of high performance computing. Intel, founded in 1968, is a leading supplier of microcomputers; microcomputer components, modules and systems; and parallel supercomputers. The company had third quarter 1990 revenues of $1 billion. Intel's Supercomputer systems Division was formed as Intel Scientific Computers in 1984 and shipped the world's first parallel supercomputer, the iPSC/l, in 1985. The division offers its customers comprehensive, world-wide technical support, including applications consulting from the PhD-level scientists and mathematicians of Intel's Computational Sciences Group, systems integration support, porting services, benchmarking assistance and end-user training.

The Future of Parallel Supercomputing: The Intel/DARPA Touchstone Program

To advance the state of the art in scalable multicomputer systems, Intel and DARPA/ISTO[1] initiated the Touchstone Program, a comprehensive, three-year research and development project

[1] DARPA/ISTO -- The Defense Advanced Research Projects Agency/Information Science and Technology Office.

designed to achieve order-of-magnitude improvements in key aspects of multicomputer performance. Touchstone is funded by a $7.6M DARPA grant and $19.9M from Intel Corp.

In Sept. 1990, Intel demonstrated the Touchstone Delta System, the third of four major Touchstone prototype systems and the system that will be installed for the Concurrent Supercomputing Consortium. In contrast to the hypercube interconnect system of earlier Touchstone systems, Delta nodes are arranged in a two-dimensional mesh. Delta uses a mesh router chip designed at Caltech and a backplane-routing plane arrangement designed by Intel to scale to 512 nodes and provide peak performance of 30 GFLOPS. The interconnect system has a bisection bandwidth approaching 1 Gbyte/second.

By the end of 1991, Intel will demonstrate Sigma, the fourth and final Touchstone prototype. Sigma will scale to atleast 2,048 processors, 64 Gbytes of main memory and half a Terabyte of online storage. Peak aggregate performance will exceed 150 GFLOPS and 100,000 MIPS. Sigma will use high-density packaging that quadruples the system's packaging density over previous Touchstone prototypes Sigma will also incorporate scalable visualization facilities, multiple processors sharing memory within a node, an integrated development-tool environment and other advances.

Intel retains the rights to commercialize technology developed as a result of the Touchstone Program. Intel's iPSC/860 computer, Concurrent File System and Concurrent I/O facility are examples of technology transfer from Touchstone to Intel products.

Conclusion

The scientific and economic importance of turbulence studies as well as their enormous computational demands make it inevitable that LES researchers will continue to be on the forefront of high performance computing. As parallel supercomputing takes us to TeraFLOPS performance levels and beyond and as the amount of parallel-based software continues to grow, we can expect the decade to bring new insights and understanding -- and new research questions --for physicists, oceanographers, geophysicists, meteorologists, aeronautical engineers and others involved in the study of turbulence.

Figure 1.

The Grand Challenges – Performance

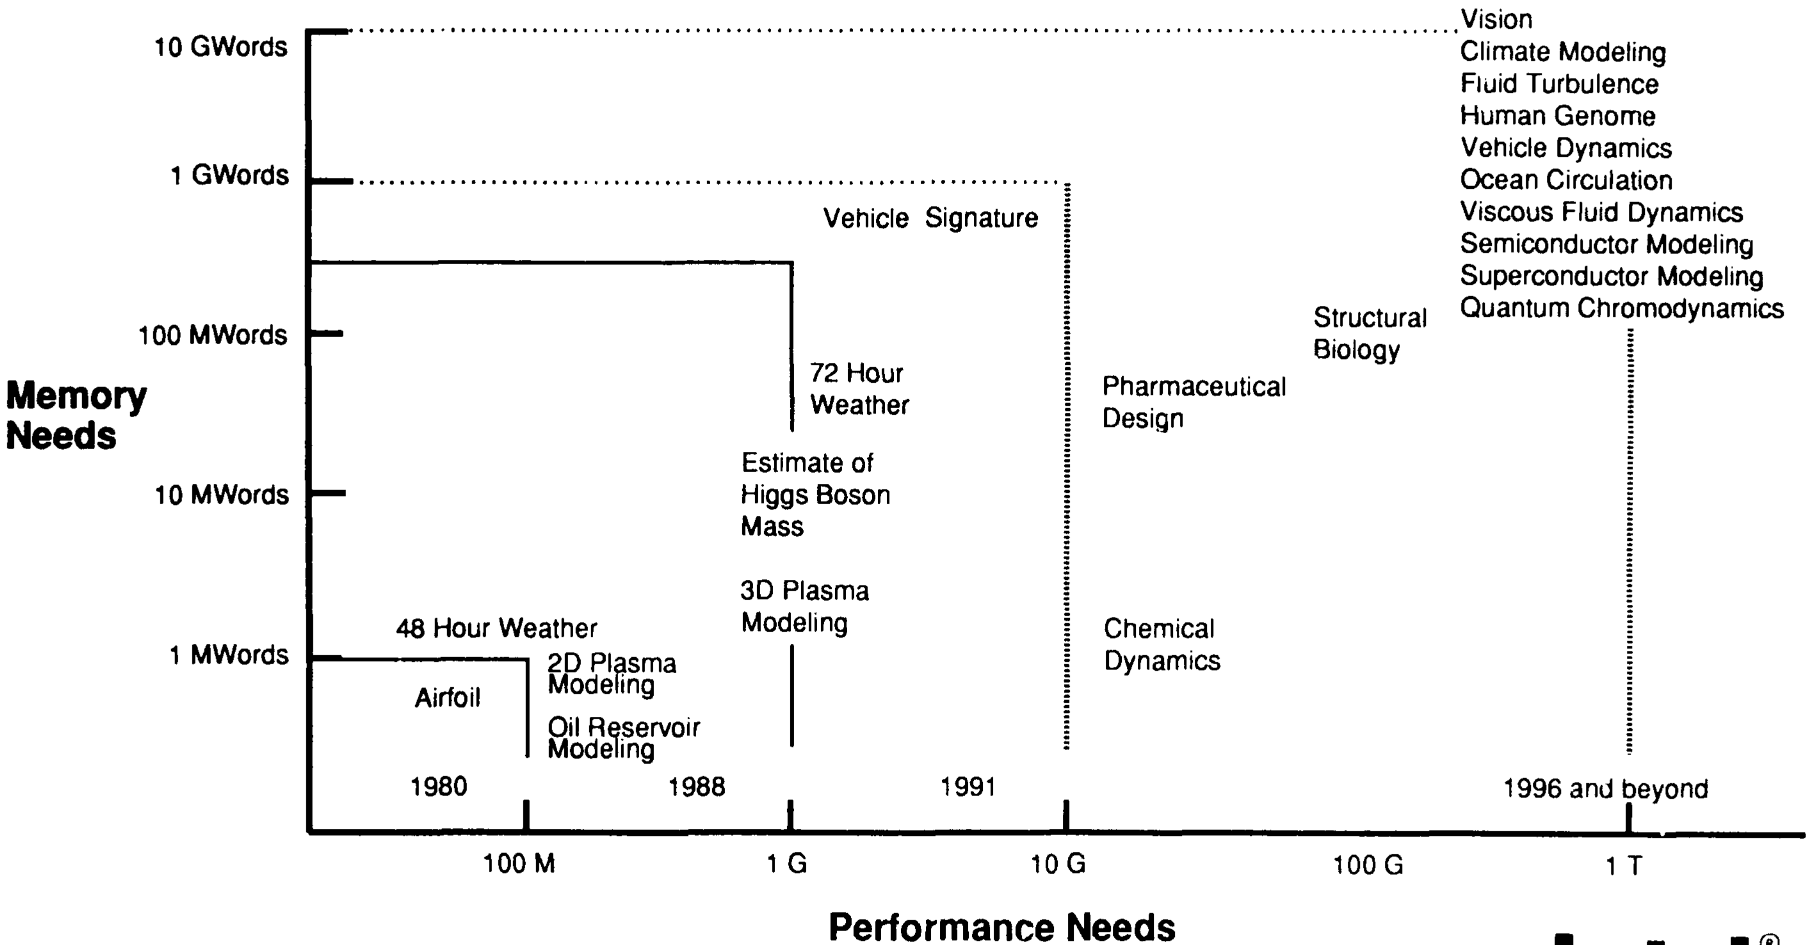

Supercomputer Systems Division
October 1990

intel®

Figure 2.

Parallel Supercomputer Performance

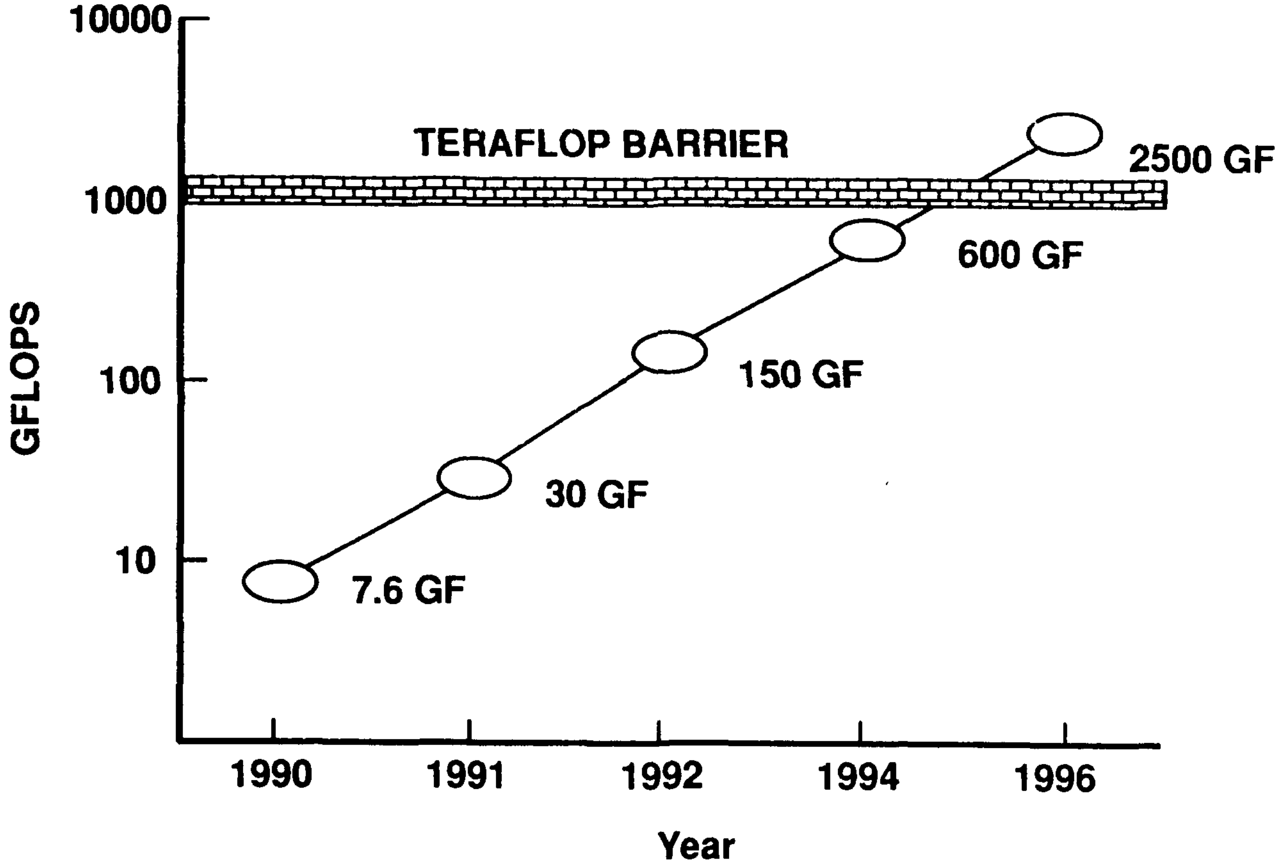

Supercomputer Systems Division
October 1990