

How do the "minisupers" stack up?

by Jack J. Dongarra
Argonne National Laboratory*

In the last few years there has been an explosion in the number of computers that are based on advanced architecture design and have the capability for high performance. These machines cover a price range from \$50,000 to \$20,000,000. Clearly there is a wide range in their performance as well. Billed as "minisupercomputers," "near-supercomputers," "affordable supercomputers," or "Crayettes," a subset of these innovative machines provides close to the performance of the leading-edge supercomputers at a fraction of their multimillion-dollar price tag.

There are a number of reasons for the proliferation of these high-performance machines:

- use of off-the-shelf processors,
- standard bus interfaces,
- custom gate arrays, and
- availability of venture capital.

The performance gains on these systems come from a number of sources. Vector processing and pipelined-instruction execution were among the first ways performance improvements were made in supercomputers. A number of manufacturers of today's high-performance/low-cost systems are using similar techniques to achieve this goal.

In addition, some of the machines have multiple-function units in one CPU, allowing for independent operations to proceed in parallel within a single processor.

Another approach to getting high performance at low cost involves the use of multiple processors. Here two or more CPUs are connected in a network, cooperating to solve a given problem. The network may range from a standard bus to a multidimensional hypercube.

Along with parallel processing comes the question of where the memory is located. Implementations range from global memory accessible to all processors, to semiprivate memory available directly to a subset of processors, to local memory connected to a single processor requiring communication to transfer information.

Many of today's machines are really a hybrid design. For example, the CRAY X-MP has up to four processors (MIMD), but each processor uses pipelining (SIMD) for vectorization. Moreover, where there are multiple processors there may or may not be caches and virtual memory systems, and the interconnections can

*Work supported in part by the Applied Mathematical Sciences subprogram of the Office of Energy Research, U. S. Department of Energy, under Contract W-31-109-Eng-38.

be by crossbar switches, multiple bus-connected systems, timeshared bus systems, etc.

Here we focus on three minisupercomputer architectures: the Alliant FX/8, Convex C-1, and FPS-164 with MAX boards. Each of these machines has the capability of high performance (greater than 1

Mflops) and low cost (under one million dollars). Each machine has an architecture that requires the user to aid to various degrees to gain the maximum performance. Some software support for improved productivity is also already available.

Continued on p. 100

Table 1.
Solving a system of linear equations with LINPACK in full precision for a matrix of order 100.

| Computer | Compiler | Ratio | Mflops | Time (sec) | Unit (μ sec) |
|------------------------|---------------|-------|--------|------------|-------------------|
| NEC SX-2 | Fortran 77/SX | 0.26 | 46 | 0.015 | 0.043 |
| Cray X-MP/2 (1 proc.) | CFT 1.13 | 0.50 | 24 | 0.028 | 0.082 |
| CDC Cyber 205 (2-pipe) | FTN | 0.70 | 17 | 0.039 | 0.11 |
| Fujitsu VP-200 | Fortran 77 | 0.72 | 17 | 0.040 | 0.12 |
| Hitachi S-810/20 | FORT77/HAP | 0.74 | 17 | 0.042 | 0.12 |
| Cray-1S | CFT | 1 | 12 | 0.056 | 0.16 |
| IBM 3090 Model 200/VF | VS Fortran V2 | 1.0 | 12 | 0.057 | 0.17 |
| Alliant FX/8 (8 CEs) | Comp Dir | 2.0 | 6.2 | 0.11 | 0.320 |
| Convex C-1 | Fortran 1.6 | 4.2 | 2.9 | 0.235 | 0.69 |
| Alliant FX/8 (8 CEs) | FX Fortran | 4.9 | 2.5 | 0.274 | 0.80 |
| FPS-164 | F.01 D, opt=3 | 9.1 | 1.4 | 0.508 | 1.48 |
| Alliant FX/1 (1 CE) | FX Fortran | 9.6 | 1.3 | 0.540 | 1.57 |

Table 2.
Solving a system of linear equations using the vector unrolling technique for a matrix of order 300.

| Computer | Compiler | Mflops | Time (sec) | Unit (μ sec) |
|-----------------------|-----------------------------------|--------|------------|-------------------|
| Cray X-MP/4 | CFT (coded MV routines) | 480 | 0.038 | 0.0042 |
| Fujitsu VP-200 | Fortran 77 | 183 | 0.099 | 0.011 |
| Cray X-MP/2 | CFT | 161 | 0.113 | 0.012 |
| Hitachi S-180/20 | FORT77/HAP | 158 | 0.115 | 0.013 |
| Cray X-MP/1 | CFT | 106 | 0.172 | 0.019 |
| Cray 1-S | CFT | 66 | 0.273 | 0.030 |
| CDC Cyber 205 | ftn 200 opt=1 (coded MV routines) | 31 | 0.59 | 0.065 |
| IBM 3090 Model 200/VF | VS Fortran V2 (coded MV routines) | 27 | 0.673 | 0.074 |
| FPS-164 + 4 MAX | E, F77 (coded MV routines) | 26 | 0.70 | 0.078 |
| FPS-164 + 3 MAX | E, F77 (coded MV routines) | 24 | 0.77 | 0.085 |
| FPS-164 + 2 MAX | E, F77 (coded MV routines) | 20 | 0.89 | 0.098 |
| IBM 3090 Model 200/VF | VS Fortran V2 | 18 | 1.03 | 0.114 |
| FPS-164 + 1 MAX | E, F77 (coded MV routines) | 15 | 1.8 | 0.130 |
| Alliant FX/8 (8 CEs) | FX Fortran (coded MV routines) | 14 | 1.3 | 0.140 |
| Convex C-1 | Fortran 1.6 (coded MV routines) | 14 | 1.3 | 0.140 |
| Convex C-1 | Fortran 1.6 | 8.7 | 2.1 | 0.230 |
| FPS 164 | E, opt=3 (coded MV routines) | 8.7 | 2.1 | 0.231 |
| Alliant FX/8 (8 CEs) | FX Fortran | 7.3 | 2.5 | 0.275 |
| FPS 164 | E, opt=3 | 5.1 | 3.6 | 0.395 |
| VAX 11/780 FPA | Unix xl77 | 0.11 | 177.0 | 19.5 |

Notes

Full Precision implies the use of (approximately) 64-bit arithmetic, e.g., CDC single precision or IBM double precision.

Ratio indicates the speed of each particular machine configuration relative to the CRAY-1S using a Fortran coding for the BLAS in full precision (ratio = 1). Higher numbers indicate slower execution.

Mflops (millions of floating-point operations per second) measures rate of execution. For solving a system of n equations, approximately $2/3n^3 + 2n^2$ operations are performed (we count both additions and multiplications).

Coded MV routines refers to an assembly-language version of the matrix-vector subroutine.

the X.25 standards, transport class 4 and connectionless internet. The X.25 tests are used to test conformance with Federal Information Processing Standard 100 (FIPS 100), which is a subset of the international standard. FIPS 100 and the conformance tests are used by the Defense Data Network. The transport class 4 tests have been applied to over 25 vendor implementations and were used to support the 1984 NCC demonstration and the 1985 Autofact demonstration. The internet tests were developed by the NBS with assistance from engineers loaned by Honeywell, Intel, and NCR. These tests were applied to vendor implementations to support the 1985 Autofact demonstration. Currently, the NBS is developing conformance tests for IEEE 802.4 (token bus), with assistance from over six guest engineers working in the NBS laboratory. These tests will be completed in calendar year 1986. Priority has been given to the development of tests for File Transfer Access and Management (FTAM) and the Message Handling Protocol (CCITT X.400). This will be followed by tests for the Virtual Terminal Service basic class. The primary users of these tests have been

the vendors participating in the two demonstrations. The tests are applied during the development of the implementations. The tests have also been used for multi-vendor testing at the laboratories at the NBS, General Motors, and Boeing.

Performance Testing

Another road block to the implementation of open systems interconnection is the debate over the performance of individual OSI protocols or, of more significance, a seven layer stack of OSI protocols. The NBS is addressing this problem in its laboratories through the development of prototype implementations, modelling and simulation, performance metrics and testing systems, and experimental testing of the protocols. The protocols are tested over local networks, including CSMA/CD and token bus, X.25 public data networks, and Department of Defense datagram-based networks. The testing is also done against different application requirements ranging from bulk file transfers to real-time applications.

What comes next?

A major issue that must be addressed by both the standardization community and the implementors is network management. The system management functions for accounting, configuration control, fault management, security, performance, directory services, and naming and addressing must be addressed in order to achieve real distributed computing. The objective is to develop OSI-based computer networks, as opposed to networks of computers. This goal pushes at the edges of the state of the art and will not be done quickly. To address the current problems of developing compatible implementations and to address the future problems of network management, the NBS has proposed the formation of OSINET. The purpose of OSINET is to facilitate the cooperative development of OSI test methods, to test vendor implementations as a means of verifying those test methods, and to support multivendor testing so that vendors can test their implementations against each other. In addition, OSINET provides an excellent vehicle for trying out the OSI and CCITT concepts for network systems management. So far, in addition to the NBS, 15 corporations have volunteered to assist in the development of OSINET.

No single organization, no single industry, no single country has the significant critical mass to undertake these developments alone. Open systems interconnection has been a worldwide effort which will lead to the kind of compatibility that users want and provide the framework for vendors to develop distributed applications.

Robert P. Blanc is director of the Center for Computer Systems Engineering at the National Bureau of Standards. He directs research and engineering in computer networking, network protocols, local-area networks, computer security, parallel processing, and storage media. He is also actively involved in the NBS Workshop for Implementors of Open System Interconnection.

Before joining NBS, Blanc was the assistant director of the Computer Center at the University of Pittsburgh. He was a team member of the President's Reorganization Project on ADP and has taught at Catholic University, the American Association for the Advancement of Science, and the IEEE. He was awarded the William A. Jump Award for distinguished public service and the Department of Commerce Gold Medal for initiation and direction of the NBS Networking Program.

Questions about this article can be addressed to the author at the National Bureau of Standards, Gaithersburg, MD 20760; blanc@nbs-vms.arpa.

How do the "minisupers" stack up?

Continued from p. 93

Alliant is the first computer that provides for automatic parallelization of a user's Fortran program. The compiler detects opportunities for parallel execution and generates the appropriate code to allow the computation to be carried out across multiple processors.

The Convex Fortran vectorizing compiler is state of the art, providing vectorization potentially superior to that of the CRAY CFT compiler. The compiler is competitive with the current Japanese vectorizing compilers, such as those from Fujitsu, Hitachi, and NEC.

To compare the floating point performance of the Alliant FX/8, Convex C-1, and FPS-164, we ran the LINPACK benchmark on each machine. Table 1 reports timing results using LINPACK to solve a system of linear equations of order 100. The execution speeds, particularly for vector computers, may not have reached their asymptotic rates.

LINPACK routines SGEFA and SGESL were used for single precision, and routines DGEFA and DGESL were used for double precision. These routines perform standard LU decomposition with partial pivoting and back-substitution.

The LINPACK routines do not, however, reflect the true performance of "advanced scientific computers." A different implementation of the solution of linear equations better describes the per-

formance on such machines (Table 2). That algorithm is based on matrix-vector operations rather than just vector operations. It produces a program that has a high level of modularity or larger granularity, with the potential for better performance across a wide range of machines especially on high-performance computers. The number of floating-point operations required and the roundoff errors produced by both algorithms are the same; only the way in which the matrix elements are accessed is different. As before, a Fortran program was run, and the time to complete the solution of equations for a matrix of order 300 is reported.

Note that these numbers are for a problem of order 300 and all runs are for full precision.

Tables 1 and 2 were compiled over a period of time. Subsequent software and hardware changes to a computer system may alter the timing to some extent.

For further reading

J. J. Dongarra, "Performance of Various Computers Using Standard Linear Equations Software in a Fortran Environment," Argonne National Laboratory Report MCS-TM-23, January, 1986.

J. J. Dongarra and I. S. Duff, "Advanced Computer Architectures," Argonne National Laboratory Report MCS-TM-57, October, 1985.