7 Conclusion

7.1 Benchmark Results

The Matrix-Matrix Multiply (assembly coded DGEMM) benchmark showed that with the current DAISy system configuration individual node performance will peak at 22% of the theoretical maximum of 60 MFLOPS. Various other benchmark results (LINPACK and NPB 2.0) showed maximal performance rates of 56% of peak.

The lmbench suite provided an extensive test of the system latency and bandwidth among the CPU and (a) memory, (b) file system and, (c) network,. In comparison to other commodity type PCs the DAISy systems are beginning to show their age. Even though, at the time of purchase (October 1994) the DAISy systems were leading the field as far as commercial PCs were concerned, the performance growth of microcomputers has again hit a new height, with advances in cache memory and main memory. Increased performance of these critical components exceeds 50%.

In contrast, the results from the main memory benchmarks (LINPACK, SPEC, lmbench suite, and STREAM) main memory bandwidth of the 90MHz Pentium is well within that of many low-end workstations offered by traditional workstation vendors.

Disk I/O performance results from the bonnie benchmark show that the DAISy cluster systems have disk performance equivalent or superior to that of the measured HEAT systems. The NFS server performance shows the disadvantage of the current Intel motherboards inability to support more than one bus-mastering network interface card.

The Netperf suite provided additional network performance analysis verifying results obtained from the lmbench suite. The 100Mb DAISy network lagged significantly behind the performance achieved from the FDDI network connected to each HEAT cluster node. The cause of this performance lag is due to the memory bandwidth of motherboards used in the DAISy cluster. Using Triton PCI chipset based motherboards that support Pipelined Burst SRAM, current Pentium 100 MHz systems exceed 72 Mbit/s TCP throughput over Fast Ethernet as measured by *Netperf*. The DAISy nodes configured with P5-90 CPUs are not powerful enough to drive the networking hardware to the theoretical maximum.

The evaluation of the performance of the DAISy cluster on the three NPB application benchmarks LU, BT, and SP is complicated by the source and intentions of the data which is being examined. First, the results reported in [19] were obtained under NPB 1.0 rules. However, the data reported for DAISy was obtained under NPB 2.0 rules. The difference in the two sets of rules is essentially that NPB 1.0 rules allow intensive optimization of codes in order to assess the absolute maximum performance obtainable on the algorithm from a particular architecture, while NPB 2.0 rules are intended to determine the performance of a parallel architecture on a portable parallel code using MPI as the message passing standard. The codes run on DAISy were modified only to the extent needed to run; i.e., no algorithmic optimizations were made In view of these differences in the source of data, it is impressive that NPB SP, BT, and LU implementations run on DAISy have price/performance with effectively unmodified, portable MPI message passing codes that exceed that of highly optimized codes on several architectures, and is competitive with many, including shared memory architectures. This comparison between NPB 1.0 results for other systems and NPB 2.0 results from DAISy is necessitated by the lack of NPB 2.0 data at the time this analysis was performed. It is expected that DAISy will exhibit much better comparative performance against most other systems when NPB 2.0 data becomes available for them.

The Parallel Seismic Inverse Problem was the first demonstration of how **DAISy** can be used as a parallel processing system. The project goal was to demonstrate a parallel seismic inverse code that runs scalably on inexpensive IBM compatible platforms.

7.2 Overall

In the past year the DAISy (**D**istributed **A**rray of **I**nexpensive **Sy**stems) cluster has proven to be a cost effective method for demonstrating the feasibility of solving grand challenge type computing problems on a cluster of inexpensive commodity parts, but only due to the tremendous increase in the performance growth of microcomputers has practical parallel computing on these type of clusters been possible.

The present DAISy cluster distributed system is defined as a collection of personal computers working together on a local network to achieve a goal. The goal in this case is to perceive DAISy as a parallel processing system, even though DAISy is considered a distributed system. The key characteristics that define DAISy as a distributed system are: (1) resource sharing, (2) openness, (3) concurrency, (4) scalability, (5) fault tolerance, and (6) transparency.

To perceive DAISy as a parallel processing system, discussed were various viewpoints associated with MPPs and how they related to DAISy: (1) generalized parallel processor, (2) control vs. data parallelism, and (3) the single program, multiple data taxonomy (SPMD).

The technology available to make the DAISy project possible consists of the P54C Pentium 90Mhz Processor, the PCI (peripheral control interface) bus, frame switched Fast Ethernet, and a UNIX compatible operating system that can be run on architecture based on the 80x86 microprocessor. Workstation clusters were originally developed as a way to leverage the better cost basis of UNIX workstations to perform computations previously handled only by relatively more expensive Supercomputers. Commodity workstation clusters take this evolutionary process one step further by replacing equivalent proprietary workstation functionality with less expensive PC technology. As PC technology encroaches on proprietary UNIX workstation vendor markets, these vendors will see a declining share of the overall market.

7 References

[1] [1996]. “SPEC list - 1996/01/04”. WWW Tech. Doc., Web Page (January 4), URL: *http://hpwww.epfl.ch/bench/SPEC.html*.

[2] IDEAS International, [1996]. “Benchmark Results, Benchmarks -SPEC”. WWW Tech. Doc., Web Page (January), *http://www.ideas.com.au/bench/spec/spec.htm*.

[3] L. McVoy, C. Staelin [1996]. “lmbench: Portable Tools for Performance Analysis”, Proceedings of the USENIX Annual Technical Conference, San Diago, CA.

[4] Subhash Saini and David H. Bailey [1995]. “NAS Parallel Benchmark Results, 12-95”, Report NAS-95-021, Dec. 1995, NASA Ames Research Center, Moffett Field, CA.

[5] M. J. Quinn [1994]. “Parallel Computing, Theory and Practice, Second Edition”, McGraw-Hill, Inc., McGraw -Hill Series in Computer Science, New York.

[6] G. S. Almasi, A. Gottlieb [1994]. “Highly Parallel Computing, Second Edition”, The Benjamin/Cummings Publishing Company, Inc., Redwood City, California.

[7] D. Bailey, T. Harris, W. Saphir, R. van der Wijngaart, A. Woo, and M. Yarrow, [1995]. “The NAS Parallel Bencmarks 2.0”, Tech. Rep. (December), NASA Ames Research Center, Moffett Field, California.

[8] [1994]. “MPI: A Message-Passing Interface Standard”, Message Passing Interface Forum (May 5)

[9] G. Coulouris, J. Dollimore, T. Kindberg [1994]. “Distributed Systems, Concepts and Designs, Second Edition”, Addison-Wesley Publishing Company, Wokingham, England.

[10] V. S. Sunderam, [1996]. “Internet Parallel Computing Archieve: /parallel/environments/pvm3/emory-vss/”. WWW Tech. Doc. Web Page, URL: *http://www.hensa.ac.uk/parallel/environments/pvm3/emory-vss.*

[11] W.R. Stevens [1990]. “UNIX Network Programming”, Prentice-Hall, Inc., A Simon & Schuster Company, Englewood Cliffs, Jew Jersey.

[12] A. Geist, A. Beguelin, J. Dongarra, R. Manchek, V. Sunderam, [1994]. “PVM: Parallel Virtual Machine, A Users’ Guide and Tutorial for Network Parallel Computing”, MIT Press, Cambridge, Massachusettes.

[13] M. J. Bach [1986]. “The Design of the UNIX Operating System”, Prentice-Hall, Inc., A Simon & Schuster Company, Englewood Cliffs, Jew Jersey.

[14] E. Solari and G. Willse [1994]. “PCI Hardware and Software”, Annabooks, San Diego, CA.

[15] The Linux Documentation Project [1994]. “Dr. Linux”, Linux System Labs, Chesterfield, MI.

[16] Bay Networks, Inc. [1994]. “LattisSwitch Model 28115, Fast Ethernet Switching Hub”, Product Update (March 14), Bay Networks, Inc., Santa Clara, California.

[17] [1996]. “FreeBSD”. WWW Tech Doc. and Software Archieve, Web Page (January), URL: *http://www.freebsd.org*.

[18] S. S. Fried [1995]. “Pentium Optimization and Numeric Performance”, Dr.Dobb’s Journal, pp.18-29, (January).

[19] M. L. Schmit [1995]. “Pentium Processor Optimization Tools”, Academic Press, Cambridge MA.

[20] J. J. Dongarra, J. Du Croz, I. S. Duff, and S. Hammarling [1990] A Set of 3 Basic Linear Algebra Subprograms, ACM Trans. Math. Soft., 16 , pp. 1-17.

[21] J. J. Dongarra [1995]. “Performance of Various Computers Using Standard Linear Equations Software”, University of Tennessee Computer Science Technical Report CS-90-85, {February 24).

[22] J. D. McCalpin [1995]. “Memory Bandwidth and Machine Balance in Current High Performance Computers.” Invited for submission to IEEE Technical Committee on Computer Architecture newsletter. (December).

[23] R. Weicker, J. Reilly [1995]. "SPEC Frequently Asked Questions (FAQ) / SPEC Primer". WWW Tech. Doc., Web Page (December 15), Siemens Nixdorf, Paderborn/Germany, weicker.pad@sni.de, and (Intel) and R. Weicker, jwreilly@mipos2.intel.com. URL: *http://hpwww.epfl.ch/bench/SPEC.FAQ.html*.

[24] T. Bray, Bonnie source code, 1990.

[25] Information Networks Division, Hewlett-Packard Company [1995], “Netperf: A Network Performance Benchmark, Revision 2.0”, Tech. Rep. (February 15), Hewlett-Packard Company, URL http://www.cup.hp.com/netperf/NetperfPage.html.

[26] [1995]. “Netlib”, WWW.Tech. Doc. and Software Archieve, Web Page (January), URL: *http://www.netlib.org*

[27] S. White, A. Alund, and V.S. Sunderam, [199x]. “Perfromance of the NAS Parallel Benchmarks on PVM Based Networks”, Tech. Rep., Department of Mathematics and Computer Science, Emory University, Atlanta,Georgia.

[28] Rob Armstrong, [1994]. “Parallel Seismic Invers Model”. Tech. Demo. Handout (November), Sandia National