# Development of a Beowulf-Class High Performance Computing System for Computational Science Applications

**Rafael Saldaña, Jerrold Garcia, Felix Muga II, and William Yu**
Ateneo High Performance Computing Group
Ateneo de Manila University, Quezon City, Philippines
E-mail: raf@mathsci.math.admu.edu.ph

## ABSTRACT

Using Beowulf cluster computing technology, the Ateneo High Performance Computing Group has developed a high performance computing system consisting of eight compute nodes. Called the AGILA HPCS this Beowulf cluster computer is designed for computational science applications. In this paper, we present the motivation for the AGILA HPCS and some results on its performance evaluation.

*Key words:* High-performance computing system, cluster computing, Beowulf cluster, performance evaluation, computational science applications

## INTRODUCTION

In the Philippines today, computing power in the range of gigaflops is not generally available for use in research and development. Conventional supercomputers or high performance computing systems are very expensive and are beyond the budgets of most university research groups especially in developing countries such as the Philippines. A lower cost option available to such research groups is to establish a cluster of high performance workstations designed to work in parallel (Saldaña, 1999; Tan, 1998).

The Beowulf parallel workstation developed by the National Aeronautic Space Agency (Beowulf Project, 2001; Ewing, 1997; Sterling, 1995) employs 16 PC-based processing modules integrated with multiple Ethernet networks. Large disk capacity and high disk to memory bandwidth is achieved through the use of hard disk and controller for each processing module supporting up to 16 way concurrent accesses. A series of experiments that measures the scaling characteristics of Beowulf clusters in terms of communication bandwidth, file transfer rates, and processing performance was conducted by a team of NASA scientists (Sterling, 1995). Their evaluation includes a computational fluid dynamics code of an N-body gravitational program. Their results show that the Beowulf architecture provides a new operating point in performance to cost for high performance workstations, especially for file transfers under favorable condition (Sterling, 1995).

To use cluster for parallel processing, a software layer must be installed to make a virtual parallel machine from a group of nodes. This software usually takes care of parallel task creation/deletion, passing data among tasks, and synchronizes tasks execution. Examples of systems available publicly are PVM (from the university of Tennessee and Oakridge National Laboratory), MPI (from MPI Forum), and BSP (from Oxford University).

The typical steps employed in building a Beowulf cluster are as follows (Uthayopas, 1998):

- Install Linux operating system on the front-end.
- Install each compute node and link them together to form a cluster.
- Set-up a single system view.
- Install parallel computing systems.

## THE AGILA HPCS

The High Performance Computing Group of the Ateneo de Manila University has developed **AGILA**, which stands for the *Ateneo Gigaflops-Range Performance, Linux Operating System, Athlon Processors High Performance Computing System.* It consists of eight (8) compute nodes connected by a 100Mbps Fast Ethernet and supports parallel programming using message passing software such as LAM-MPI and PVM. (See Fig. 1 for the AGILA HPCS hardware design). As a scalable high performance computing system, the AGILA HPCS can be increased to 64 compute nodes.

The AGILA HPCS is an interdisciplinary project (mathematics, physics, engineering, and computer science) aimed at supporting the computational science and engineering research at the Ateneo de Manila University. It is also intended to support the parallel computing courses offered by the Ateneo de Manila University, particularly those computational subjects offered in the applied mathematics/computational science program of the Mathematics Department.

Some of the identified applications for the AGILA HPCS are as follows:

- Modeling vehicular traffic dynamics
- Modeling air pollution dynamics
- Climate change modeling
- Fluid dynamics modeling for biophysics and medical applications
- Financial systems modeling
- Artificial neural networks modeling
- Network Cluster Computing
- Coding theory and cryptography
- Ecosystems modeling, and
- Parallel computing education

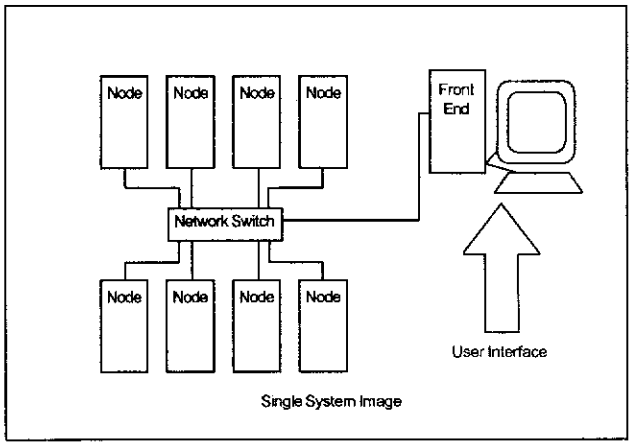


Figure 1. AGILA Hardware Design

## PERFORMANCE EVALUATION

Table 1 lists some preliminary results of the performance evaluation of the AGILA 8-node Athlon Cluster.

The benchmarking software used is the High Performance Linpack (HPL), a freely available implementation of the High Performance Computing Linpack benchmark for distributed memory computers. It is a software package that solves a (random) dense linear system of equations in double precision (64 bits) arithmetic. HPL provides a testing and timing program to quantify the accuracy of the obtained solution as well as the time it took to compute it.

Based on these results, the AGILA HPCS has achieved a peak performance of 3 Gigaflops.

It should be noted that the performance measure for different systems may vary with different factors. The most important factors are the processor, memory, and network connection. HPL assumes that the power of

Table 1. AGILA HPCS performance evaluation results

| N | NB | P | Q | Time | Gflops |
|---|---|---|---|---|---|
| 10000 | 32 | 2 | 4 | 239.59 | 2.783 |
| 10000 | 64 | 2 | 4 | 221.21 | 3.014 |
| 10000 | 80 | 2 | 4 | 224.37 | 2.972 |
| 10000 | 80 | 1 | 8 | 257.23 | 2.592 |
| 10000 | 80 | 8 | 1 | 277.40 | 2.404 |
| 5000 | 80 | 8 | 1 | 56.12 | 1.486 |
| 8000 | 80 | 8 | 1 | 162.10 | 2.107 |
| 2000 | 80 | 8 | 1 | 6.45 | 0.796 |

every node in the cluster is the least powerful of all the nodes.

## CONCLUSION

Using Beowulf cluster computing technology, the Ateneo High Performance Computing Group has developed a low-cost high performance computing system consisting of eight compute nodes. This system is called the AGILA HPCS and it is designed for computational science applications. The initial performance evaluation of the AGILA HPCS shows that its peak performance is in the order of 3 gigaflops.

## ACKNOWLEDGMENTS

This project was funded by the CHED Center of Excellence in Science and Mathematics Grant of the Ateneo de Manila University.

## REFERENCES

Beowulf Project at the Center of Excellence in Space Data and Information Sciences. http://www.beowulf.org.

Ewing, M., 1997. Beowulf at SC97. http://www.yale.edu/secf/etc/sc97/beowulf97.htm

Reschke, C., T. Sterling, D. Ridge, D. Savarese, D. Becker, & P. Merkay, 1996. A design study of alternative network topologies for the Beowulf Parallel Workstation. Proceedings of the Fifth IEEE International Symposium on High Performance Distributed Computing.

Saldaña, R., 1999. A report on the workshop and conference on modeling, simulation, and scientific computing (MODEL '99). Meeting of the Steering Committee of the High-Performance Computing in Asia-Pacific (HPC-Asia), Maui High Performance Computing Center, Hawaii.

Sterling, T., D. Becker, & D. Savarese, 1995. Beowulf: A parallel workstation for scientific computation. Proceedings of the Fourth 1995 International Conference on Parallel Processing (ICPP).1:111-114.

Sterling, T., D. Becker, D. Savarese, M. Berry, & C. Reschke, 1996. Achieving a balanced low-cost architecture for mass storage management through multiple fast ethernet channels on the Beowulf Parallel Workstation. Proceedings of the International Conference on Parallel Processing (ICPP).

Tan, K., L. Leslie, & R. Morison. 1998. A high performance workstation cluster for numerical weather prediction. Proceedings of the High Performance Computing Conference in Asia-Pacific (HPC-Asia), Singapore, 23-25 September (1998).

Uthayopas, P., T. Aungsakul, & J. Maneesilp, 1998. Building a parallel computer from cheap PCs: SMILE cluster experience. Proceedings of the Second Annual National Symposium on Computational Science and Engineering. Bangkok, Thailand: National Science and Technology Development Agency.