Undergraduate Research Award Report:
Construction of Seattle University Beowulf Cluster (SUBC) for Use in Solving Computation Physics Problems

Seattle University SPS Chapter Seattle, W A

November 26, 2002

Abstract from A ward Proposal

Many members of the Seattle University SPS chapter are cUITently researching, under faculty direction, a variety of computational problems in nuclear physics, shock-wave physics, solid-state physics, fluid mechanics, and quantum mechanics. These student projects at Seattle University would be greatly facilitated by the building and running of a Linux-based Beowulf cluster, which speeds up Fortran and C++ physics computations through parallel processing. The Seattle University SPS chapter requests funds to build a small scale Beowulf cluster (5 computers) for which Seattle University is committed to matching any funds received from Sigma Pi Sigma.

History of the Project

During the end of Winter Quarter and into Spring Quarter of 2002, the Seattle University SPS chapter constructed a Beowulf cluster consisting of five computers ("nodes") and a Fast Ethernet network connecting them. The student most responsible for the construction of the Beowulf cluster is Jason Shoup, physics graduate in 2002.

Over the Summer of 2002, Mike Clement, a Seattle University Computer Science major and SPS chapter member, conducted research on the cluster, focusing on software configuration and programming on the cluster. He continued in the configuration of the cluster where Mr. Shoup ended as he graduated in June. By the end of the summer, the cluster was fully operational with all necessary software to run parallel applications. Mr. Clement's research culminated in a physics cluster resource web page available at:
http://www . seattleu.edu/ scieng/phys/ subclindex.html

This web page offers documentation of the cluster, information on CUITent research being conducted on the cluster, and a listing of resources for those who wish to further their knowledge about clusters and their applications in physics.

Cluster Design

Each machine was designed with the fastest processors available at the time, and equipped with amounts of processing power and data storage sufficient to execute physics computations in greatly reduced time. The processor chosen for each node was t4e Advanced Micro Devices

Athlon XP 1800+. It features a 1.5 Gigahertz clock speed, and the best floating-point arithmetic logic available at the time the cluster was built. The Athlon XP 1800+ was also the least expensive top of the line processor at the time, thus allowing the cluster to be both as fast and as economical as possible. Each node in the cluster also has 512 Megabytes of Double-Data-Rate RAM, giving the cluster the ability to store large amounts of data in memory, thus increasing the processing speed compared to the scenario in which data must be continuously accessed from hard disks.

All nodes within the cluster are connected via a Fast Ethernet network, with a transfer rate of 100 Megabits per second. The network is connected through a switch, which offers the advantage of providing a fulll 00 Megabits per second to each distinct connection, as opposed to low-end hubs where all connections share a limited total amount of network bandwidth. This also assists in the efficiency of the cluster, as program code and data transfers between machines quickly.

The cluster runs on Red Hat Linux version 7.1. Linux is a highly stable operating system that is specifically designed for network applications and highly-streamlined program code. The client, or processing, nodes access all data and program code from the server node via the Network File System, a file-sharing system built for Unix-like operating systems. All client nodes can be controlled from the server node, which uses Secure Shell (SSH) and Remote Shell (RSH) to start parallel applications on the client nodes.

Parallel applications run using the Message Passing Interface (MPI) standard. The chosen implementation of the MPI standard is LAM-MPI[l]. It is a highly-efficient, fullyfeatured implementation of MPI, complete with debugging and tracing tools, enabling researchers to expedite the process of developing program code for use on the cluster. It has already been tested with scenarios such as approximating pi to several digits using a Riemann Sum, and modeling three-dimensional environments using POV-Ray[2] and MPIPOV[3].

Future Projects

The cluster web page being newly released, the focus of the SPS chapter is currently to familiarize all members with the cluster, and to begin several research projects on the cluster. Mr. Clement is currently involved in research of parallel rendering of graphics and animation, useful in simulation of physical phenomena.

Additionally, several physics professors in the department have expressed interest in conducting research in their respective fields, alongside students, making use of the cluster for computations and simulations. Among the fields of interest are fluid dynamics, geophysics, and particle physics.

Footnotes

[1] See http://www.1am-mpi.org/ for details.
[2] See http://www.povray.org/ for details.
[3] See http://www.verrall.demon.co.uk/mpipov/ for details.