1 Personal Information

Author's Name: Katie A. Moor
University Affiliation: University of Notre Dame
Project Affiliation: Sandia National Laboratories/CA
U.S. Department of Energy
E-mail Address: kmoor@cse.nd.edu
Poster Title: Evaluation of Gigabit Ethernet for Cluster Computing
Category: Graduate

2 Problem and Motivation

Some may intuitively think the more compute nodes interconnected on a high-speed network, the quicker a problem could be solved. However, this is not true because at some point calculation completion times get worse due to communication overhead. My research aims to define the optimal amount of nodes for Gigabit Ethernet clusters given today's technology. Knowing the optimal sizes of compute clusters for Gigabit Ethernet will assist researchers when creating their own cost efficient, computational networks and programmers when they have to determine what size of network will be needed for a new parallel application.

3 Background and Related Work

People with an interest in hi-performance computing and networking would appreciate my research the most. From Stanford to Sun, there are many educational and corporate research groups studying cluster computing. The researchers who evaluate performance usually base their studies on benchmarks for varied latency, communication overhead, and throughput. My research is different because I evaluated performance by looking at how cluster sizes effect speed-up and calculation completion times with real parallel applications. The only study that comes close to it is "The MIT Alewife Machine: Architecture and Performance" presented at the International Society for Computers and their Applications (ISCA) Conference in 1995 (http://www.cag.lcs.mit.edu/alewife/papers). The Alewife study used 1 to 32 node clusters to evaluate speed-up, but was primarily concerned with physical topology and hardware than actual performance. I, on the other hand, used 1 to 256 node clusters with a two-dimensional mesh logical topology. Another key difference is that I eliminated the 32 and 128 node clusters

because my parallel algorithm called for square dimensional meshes. Comparing square meshes to rectangle meshes is not feasible because a square mesh would have more interior nodes, thus more communication overhead and higher completion times.

4 Approach and Uniqueness

To make my study as realistic as possible, the first thing established was the size of messages the cluster would pass. Message sizes ranging from 32-2048 bytes were used because these are the commonly used sizes for parallel applications. Clusters were defined to have 1 to 256 nodes for this study. As stated in the previous section, clusters composed of 32 and 128 nodes were not tested, unlike previous studies, because they are rectangle, not square hence they would have less interior nodes, communication overhead, and overall completion time. One and two node clusters were tested for simple error checking. One node gave the overall completion time which could be checked with a mathematical calculation. Two nodes checked the communication overhead time. Nodes had a two-dimensional mesh logical topology which means each node could only communicate with its respective x and y coordinate neighbor. OPNET Modeler, a Third Millennium Technologies' network simulation application was used to create and simulate the clustered nodes. OPNET is a dynamic application that allows researchers to create their own models and specify different amounts of traffic among other functions. After simulations were completed, the completion times were graphed. Speed-up was calculated to see the overall time taken by each processor, whereas efficiency was derived as a measurement of communication overhead.

5 Results and Contributions

An evaluation of Gigabit Ethernet for cluster computing contributes to the networking field because it shows how many nodes in a cluster would give optimal compute performance using today's technology. Programmers would be assisted by this research because they would know how many nodes they would need to run a particular application. The hi-performance cluster computing community would financially benefit from the results because they would only need to create clusters that would produce the maximum speed-up.

6 Succinct Abstract

At what point does increased cluster size no longer yield increased compute performance? This study evaluated the tradeoffs of cluster size versus speed-up over the message sizes commonly used in parallel computing applications. Gigabit Ethernet connected clusters ranging from 1 to 256 nodes were simulated on OPNET Modeler, a Third Millennium Technologies network simulation application, for each of the message sizes that were 32 to 2048 bytes long. The results not only exhibited which clusters were optimal for all message sizes, but also gave insight into how many compute nodes would be optimal to run a particular application.