1 Introduction

Due to the cost of Supercomputers and massively parallel processing (MPP) systems there has been a rise in cluster based parallel processing systems. Grand challenge type computational problems can now be performed economically using workstation cluster technology. Workstation clusters are networked RISC UNIX systems that have been commonly provided by such vendors as IBM, Hewlett Packard, SUN, Silicon Graphics, and Digital Equipment Corporation. Each of these vendors is the predominant or proprietary supplier of both the hardware and UNIX operating systems. The applications that run on proprietary workstation clusters use UNIX features that enable high performance floating point, fast disk drive transfer rates, and robust networking performance in a multi-user, multi-system environment. Achieving supercomputing efficiencies on many computational intensive tasks using UNIX workstation clusters is possible. Additionally, due to the nature of clusters, individual cluster systems can be and are used as desktop UNIX workstations

The size of the personal computing (PC) market is about nine times greater than the proprietary UNIX workstation market, with Intel Pentium based systems as the performance leader. Only until recently (late 1994), the gap that proprietary workstation hardware and software greatly exceeded in (i.e.: sophistication, capability, and performance) has lessened to that of commodity PCs. Recent PC technology advances have dramatically improved the performance of the control processing unit (CPU), main memory, and cache memory. Although, some high end models offered by proprietary workstation vendors still maintain a significant advantage in peak floating point performance. However, the development and widespread availability of high performance peripheral control interface (PCI) bus network, video, and disk controllers for PCs has erased the traditional hardware I/O performance advantage of proprietary UNIX workstations. A similar evolution has occurred in PC system software. Today there are several flavors of fully compatible UNIX OSs available for the PC market. The main target for these UNIX OSs have been systems designed around the Intel 80x86 microprocessor. Due to the size of the commodity PC market the lowest possible cost for hardware is ensured. As an alternative to Supercomputers, MPPs, and a cluster of UNIX workstations, the Distributed Computing Research Department at Sandia National Laboratories, CA (SNL,CA) has constructed the 16 node PC cluster, DAISy, running a full UNIX compatible OS based upon the BSD 4.4 UNIX OS.

The DAISy (**D**istributed **A**rray of **I**nexpensive **Sy**stems) Cluster is a research prototype and testbed used for demonstrating the feasibility of solving grand challenge type computing problems on a cluster of inexpensive commodity parts. The cluster, will in the course of testing, be used to solve important problems in the fields of science and technology at the laboratory. Recent development of hardware and software have allowed for MPP type computation to be performed on PCs.

The problems currently facing researchers at SNL are the cost of Supercomputers. Because of their great cost and limitations, these systems are being retired from the nation's research institutions. Rising in their place are the MPP systems and cluster based distributed systems. While MPPs have a place in current scientific and technological efforts, not every organization can afford such systems. As an alternative to these expensive MPPs, affordable cluster based distributed systems are coming forth.

The focus of this paper will be on the hardware involved and system performance analysis of the DAISy cluster. In analyzing the DAISy performance another cluster available at SNL,CA, HEAT (**H**eterogeneous **E**nvironment **A**nd **T**estbed), will be reviewed.

The layout of this paper will be to discuss some of the history behind parallel computing, the hardware/software configuration of the DAISy cluster, the results from the benchmarks run on DAISy, and a conclusion of the overall performance and use of DAISy as a parallel processing system.

The history of parallel computing section will discuss the performance growth of microprocessors, the definition of a distributed system, and the key characteristics that make up a distributed system. Also discussed will be MPP viewpoints that are a virtual overlap when classifying cluster based distributed systems as parallel processing systems, such as, defining a generalized parallel processor, control vs. data parallelism, and the single program, multiple data (SPMD) taxonomy. Communication models and the BSD sockets interprocess communication method will be reviewed along with a layout for analyzing system performance.

The hardware/software configuration sections discuss the hardware/software associated with the DAISy cluster that help elevate PC technology to new heights: the P54C Pentium ^{TM} 90 MHz Processor, the peripheral control interface (PCI) bus, framed switched and 100Mb/s fast ethernet, the freely

redistributable FreeBSD operating system, and message passing software. Also discussed is the hardware configuration of the HEAT cluster.

The benchmark results section gives detailed information regarding DAISy's system performance, including single node, network, and overall parallel processing performance analysis. Included are not only results form both the DAISy and HEAT clusters, but various published results from SPEC [1], [2] (Standard Performance Evaluation Corporation), the lmbench suite [3] technical document and the NAS parallel benchmarks 12/95 results report [4].

The conclusion section wraps up the paper with an analysis of DAISy's system performance and an overall discussion addressing the viability of using distributed PC system based clusters as parallel processing systems.