



Arcosu Cluster Status Report, January 2002

Covering the period from January 2001 to December 2001

**Alan Louis Scheinine
Carlo Podda**

CRS4

**Centro di Ricerca, Sviluppo e Studi Superiori in Sardegna
Sesta Strada, Ovest
Zona Industriale Macchiareddu
09010 Uta (Cagliari) Italy**

E-mail: scheinin@crs4.it carlo@crs4.it

Arcosu Cluster Status Report January 2002

Covering the period from January 2001 to December 2001

Alan Louis Scheinine

Carlo Podda

CRS4

Centro di Ricerca, Sviluppo e Studi Superiori in Sardegna

Sesta Strada, Ovest

Zona Industriale Macchiareddu

09010 Uta (Cagliari) Italy

E-mail: scheinin@crs4.it carlo@crs4.it

¹*Cover page: a photo-montage of four nodes of the Arcosu cluster with the park Monte Arcosu in the background.*

Contents

1	Hardware	1
2	Software	1
3	Documentation	2

1 Hardware

The cluster began with four Hewlett-Packard 'Kayak' PCs with dual-processor main boards that use the Pentium III processor at 500 MHz. Four additional PCs arrived near the end of 2000, assembled locally by AC Computer, for which installation was completed at the beginning of 2001. These PCs have dual-processor main boards that use the Pentium III processor. The CPUs chosen have a clock frequency of 800 MHz, which was relatively fast at the time the computers were purchased. The second group of four PCs have a 64-bit, 66 MHz PCI bus, whereas a typical home computer has a 32-bit, 33 MHz PCI bus. Both the older and the newer computers have one gigabyte of memory.

Also in 2001, a Myrinet switch and interface cards were purchased in order to have a fast interconnection. Each connection has a peak speed of 1.9 gigabits/sec. Myrinet, moreover, has a rather low latency, about 5 to 7 microseconds when used with MPI – much lower than the latency of Gigabit Ethernet. The Myrinet interface cards can use a PCI bus of either 33 or 66 MHz, and either 32 or 64 bit slots. A PCI slot at 64 bits and 66 MHz gives the highest Myrinet bandwidth, which was the primary motivation for the selection of a PC with the high PCI bus bandwidth.

For the second group of PCs, each has two SCSI hard disks of 9 gigabytes and an IDE hard disk of 20 gigabytes. The IDE disk is used for scratch space.

Important file systems such as installed software (described in the next section) are now under automatic backup to a "juke box" of rewritable optical disks. In addition, copies of installed software are made to a second hard disk on a computer different from the one that usually exports the files with NFS. If the main NSF server PC should fail, it should be a rapid procedure to switch to using the second copy on the other PC.

2 Software

The software described in last year's report has also been installed on the four new PCs. This includes the Portland Group Fortran 90 and HPF compilers, Message Passing Interface (version MPI), FFTW, an up-to-date Java, IBM Open Data Explorer, Perl with many modules from CPAN, and Geomview. The cluster uses the Linux operating system, distribution Red Hat, version 6.2.

With regard to MPI, Myrinet provides a version of MPICH tailored for use with

the Myrinet interconnection. In addition a version of MPICH was compiled for use with debugging.

In 2001 we purchased the Kuck Associates (KAI) C++ compiler. In addition, the parallel solver for sparse linear equations, Aztec, was installed. Also, the graph partitioning and sparse matrix ordering software Metis was installed. Arpack and Parpak were also installed. Aztec was compile in two ways, one version for Fast Ethernet and one version for Myrinet. In addition, Metis and Arpack were compiled in two versions, one for Ethernet and one for Myrinet.

The commercial debugger for parallel programs, TotalView, was purchased and installed.

The commercial software StarCD was installed, and also Cubit, a mesh generation tool.

Big Brother “health” monitoring system was installed, which notifies the system administrators of anomalies. We added to Big Brother a script to monitor the license manager daemons for Portland Group software, TotalView and the KAI compiler. Also, the network time protocol (NTP) daemon was activated on the clusters and synchronized with three external time sources.

The NFS structure was changed to simplify software installation, nearly all installed software is available using NFS mounts rather than installing the software locally.

3 Documentation

The High Performance Computing Services (HPCS) group installed the http server Apache for WWW access to documentation. All of the documentation associated with the installed software on the cluster can be accessed from this server. In addition, the HPCS group wrote some “first facts” documentation for using the cluster and for Myrinet, as well as writing a Frequently Asked Questions (FAQ) document.

The HPCS group also has an internal newsgroup (crs4.hpc) for announcements and for discussions. At the present time, the HPCS WWW pages are not visible to the outside world because a firewall has been installed at CRS4.

Acknowledgments

This work has been partially supported by the grant “progetto Murst 1.2, Laboratorio Avanzato per la Progettazione e la Simulazione Assistita al Calcolatore” and by the Sardinian Regional Authorities.