

HPCC Software for UNIX Clusters

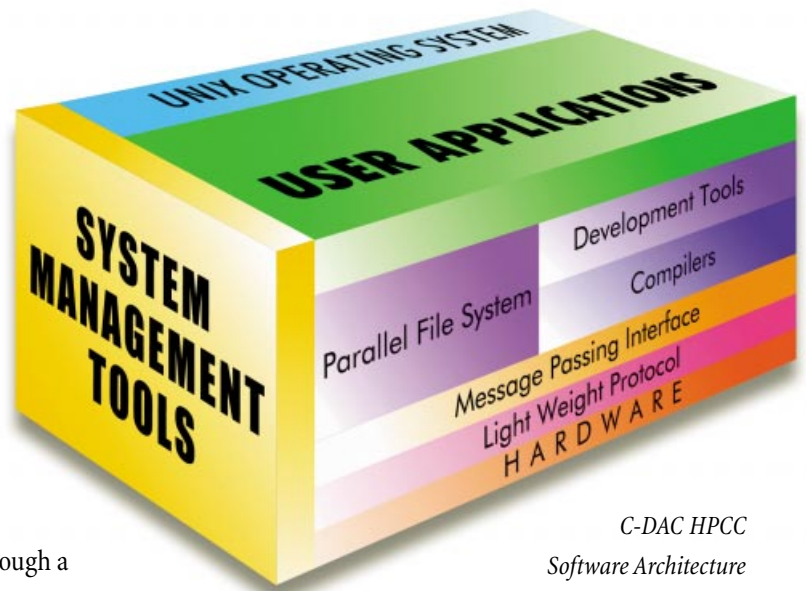
Flexible Parallel and Distributed Software Environment

INTRODUCTION

Workstation clusters are today's cost-effective solution for high performance computing. Programming the clusters to run efficiently and to provide a look and feel of conventional computers to maximise ease-of-use are the major software challenges. The C-DAC HPCC* (High Performance Computing and Communication) software effectively addresses these performance and usability challenges through a high performance flexible software environment, which adheres to established and emerging standards in parallel and distributed computing.

C-DAC HPCC software supports development and execution of sequential, message passing and data parallel programs and thus allows the ensemble of workstations to be viewed as independent workstations, cluster of workstations, or as a massively parallel processor system.

The HPCC software provides a complete solution for creating and executing parallel and distributed programs on UNIX clusters through high performance communication protocols and a rich set of program development, system management and software engineering tools.



C-DAC HPCC
Software Architecture

DESCRIPTION

The HPCC software suite of products includes high performance compilers, parallel debuggers, data visualisers and performance profilers. High performance communication protocols and efficient MPI implementation provide the right framework to extract maximum performance. Single system image at the system management level ensures ease of administration and monitoring of large UNIX clusters.

HIGHLIGHTS

- Exploits hardware features for communication, providing a low latency and high bandwidth communication substrate.
- Supports full implementation of the MPI standard. Collective algorithms tuned to a Cluster of Multi Processors (CLUMPS) architecture.
- Full function development environment for parallel and distributed applications with true support for parallel program debugging and profiling.
- Integrated Fortran 90 Development Environment.
- Unified system management tool for cluster monitoring and administration.
- Adheres to established and emerging standards in parallel and distributed computing.

The HPC software offering consists of:

- ❑ Base Software, which includes a lightweight communication substrate, optimized MPI, and a parallel file system with MPI-IO interface.
- ❑ Program Development environment, providing a rich set of FORTRAN 77 / Fortran 90 program development tools, and an integrated debugging environment for parallel programs.
- ❑ System Management tools for monitoring the effective utilization of the cluster resources and for administration of large UNIX clusters.
- ❑ Software engineering tools for developing high quality programs.

Details on these components of HPC software for UNIX clusters follow :

HPCC - BASE SOFTWARE

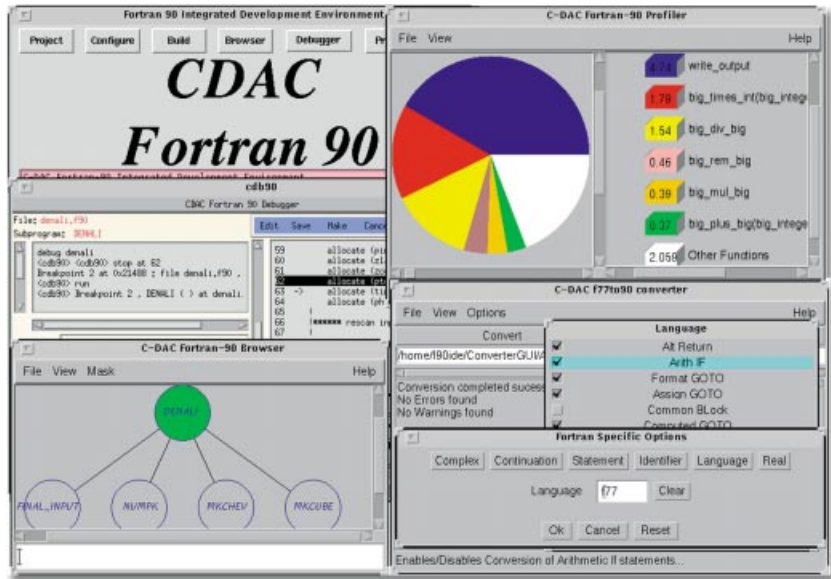
For parallel applications to scale on large clusters, the HPC base software

provides low overhead communication, optimized MPI and a parallel file system.

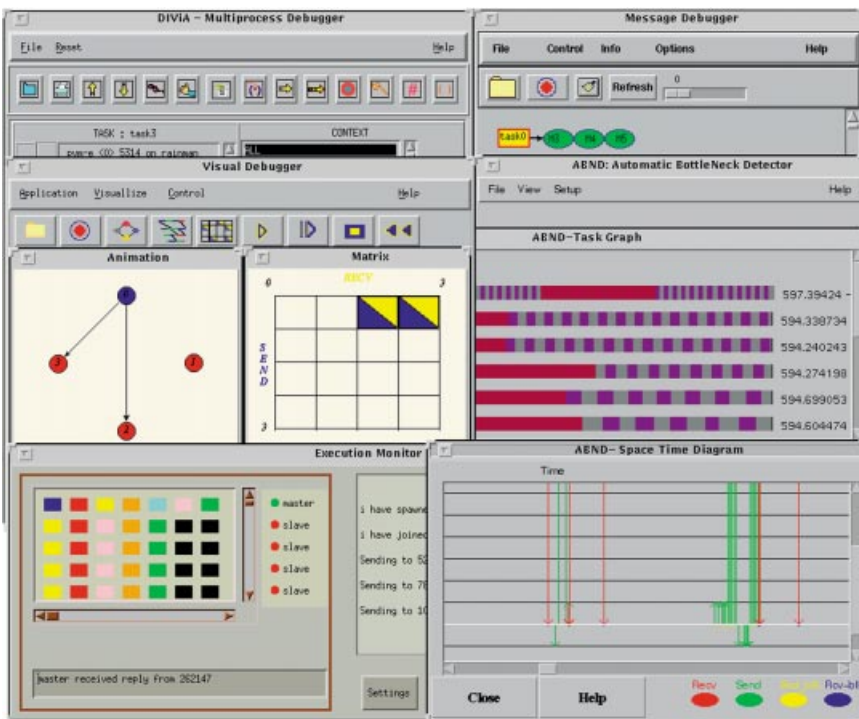
Legacy protocols like TCP/IP have large overheads, do not exploit hardware features nor allow experiments with new protocols and flow control mechanisms. HPC software's KSHIPRA* communication substrate

provides lightweight communication primitives on the ParamNet* and Myrinet* conforming to the Active Messages II specifications of the University of California, Berkeley. The implementation leverages on a mechanism, which decouples communication path from the operating system. This eliminates operating system overheads in frequent operations like send/receive. It provides a communication abstraction which allows the applications to exploit to the fullest the low latency and high bandwidth of the underlying high performance network. On top of this low level communication substrate, HPC software provides Message Passing Interface (MPI) application programming interface for parallel computing and fast sockets interface for distributed computing.

The MPI standard was jointly defined by a forum of parallel programming industry leaders including representatives from national laboratories, universities and key MPP vendors. To facilitate simultaneous communication among many parallel processes, MPI provides collective communication functions. The



F90 Integrated Development Environment



DIVIA Integrated Parallel Program Debugging Environment

performance of these functions can be enhanced by using appropriate algorithms for the architecture. C-MPI has tuned MPI collective algorithms to a cluster of SMP nodes. Also C-MPI achieves enhanced performance by layering MPI over lightweight communication protocols.

Parallel applications, to scale also require an efficient file system. HPCC software addresses this issue by providing C-PFS, an efficient parallel file system exporting MPI-IO interface.

HPCC - PROGRAM DEVELOPMENT

HPCC software provides an integrated Fortran program development environment, which includes the Fortran 90 compiler, debugger, source browser, profiler, project manager and FORTRAN 77 to Fortran 90 converter. CAF90* is a highly optimising Fortran-90 and FORTRAN-77 compiler developed jointly by C-DAC and Apogee Inc. It is fully ANSI X3.9:1978 and ISO/IEC 1539:1991 compliant. C-F77to90* is an extremely powerful tool for migrating FORTRAN 77 programs to Fortran 90.

For correctness and performance debugging of parallel applications the HPCC software provides DIViA* integrated debugging environment. DIViA (Debugger with Integrated Visualizer and Analyzer) works under both PVM and MPI environments and debugs C or Fortran language parallel programs.

Correctness Debugging Components

- ❑ Multiprocess Debugger - allows control of the execution of loosely synchronised, concurrently

running multiple tasks as a related group of tasks with source level debugging facility.

- ❑ Message Debugger - allows user to inspect, modify, block and in general control the messages between the various tasks of a parallel application.
- ❑ Visual Debugger - provides visual feedback on the execution of a parallel program, filtering out all unnecessary details and portraying sequence of communication and status of participant tasks.

Performance Debugging Components

- ❑ Automatic Bottleneck Detector - helps in detecting the communication and computational bottlenecks and their sources automatically, to help the user fine-tune the parallel application.

HPCC-SYSTEM MANAGEMENT

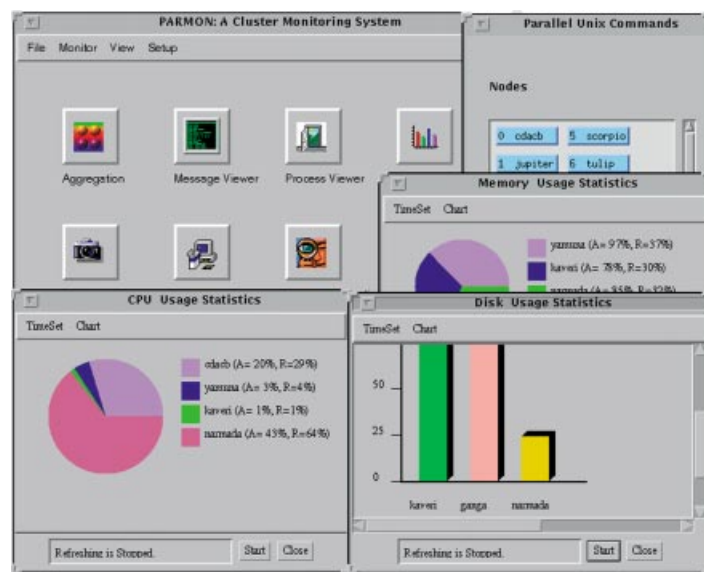
A large cluster requires tools to manage the system effectively. The system management tool -PARMON* - allows the cluster user or administrator to monitor activities and resource

utilization of various components of the cluster. It monitors the machine both at node level and at the entire system level exhibiting a single system image. The PARMON, monitors system activities like process activities, system log activities, kernel activities; controls devices; and generates and analyses events. It also provides physical and logical views of the components of the system.

The Parallel Unix Commands - PARCOM* - is a tool for executing UNIX commands in a cluster environment where there is a frequent need to execute a given UNIX command on several processors in the cluster at the same time.

HPCC- SOFTWARE ENGINEERING

HPCC software also provides programmers with tools for producing and maintaining quality code. The MetricAdvisor* calculates a variety of source code metrics, which is essential for generating quality code. It evaluates the Halstead, McCabe and Complexity Density metrics. In addition it includes a static analyser that generates the call graph of C language source files.



PARMON: Unified Cluster Monitoring System

STANDARDS ADHERENCE

C-DAC HPCC software follows established standards to reduce learning time and maximize ease-of-use. Most commands use familiar UNIX syntax. HPCC software graphical user interfaces comply with Motif or Java style standards. HPCC's

implementation of the MPI standard closely follows the published standard and allows programmers to write portable parallel applications which run on the PARAM OpenFrame* system and other cluster systems supporting the standard.

HPCC SOFTWARE : PRODUCT COMPONENTS

KSHIPRA	: Scalable Communication Substrate. Layered over ParamNet and Myrinet for UNIX Clusters.
C-MPI	: Optimized implementation of MPI for Cluster of Multi Processors (CLUMPS). Both point-to-point and collective calls have been optimized. Effectively uses both shared and distributed memory of CLUMPS.
C-PFS	: Parallel File System. Provides MPI-IO file system interface to parallel applications.
CAF90	: Fortran 90 compiler. Jointly developed by C-DAC and APOGEE, USA. Integrated development environment with debugging, profiling, browsing and project management support.
C-F77to90	: A converter to Fortran 90 for FORTRAN-77 programs. An essential product for programmers who have large FORTRAN 77 code to maintain.
DIViA	: Parallel program correctness and performance debugger. Detects communication bottlenecks, and supports message debugging. Provides architecture and communication layer neutrality.
PARMON	: Cluster monitoring tool. Monitors the cluster as a unified resource. Helps in detecting performance bottlenecks and also interfaces to diagnostic tools.
PARCOM	: Parallel Unix Commands. Provides parallel extensions to traditional UNIX commands.
MetricAdvisor	: Software engineering tool for metrics. Evaluates Halstead, McCabe, Complexity Density, Fan-in and Fan-out metrics.

AVAILABILITY

Supported Hardware	:	Workstation Clusters
Supported Operating System	:	Solaris 2.5 and above
User Interfaces	:	GUI
Supported Languages	:	FORTRAN 77/90, C
Prerequisite Software	:	Java, X/Motif
Prerequisite Hardware for Base Software	:	ParamNet, Myrinet



A Scientific Society of the
Department of Electronics
Government of India

Additional Information

For more information on CDAC HPCC software, contact your CDAC marketing representative, access the CDAC Home Page on the internet World-Wide Web (www.cdac.org.in), or send an e-mail over the internet to : ssg@cdacb.ernet.in

C-DAC reserves the right to change or modify any of the product or service specifications or features described herein without notice. The product summary is for information only. C-DAC makes no express or implied representations or warranties in this summary

*All trademarks and brand names are owned by their respective owners.

Headquarters

University of Poona Campus,
Ganesh Khind, Pune - 411 007, INDIA
Tel : 352461 Fax : 91-212-357551
Tlx : 0145-7615 CDAC IN
email : business@cdac.ernet.in

Business Division

Ramanashree Plaza, 2/1 Brunton Road,
Bangalore - 560 025, INDIA
Tel : 5584271 Fax : 91-80-5584893
Tlx : 0845-8413 CDAC IN,
email : bdm@cdacb.ernet.in

Delhi Centre

E-13, 2nd Floor, Hauz Khas,
New Delhi - 110 016, INDIA
Tel : 6863428 Fax : 91-11-6863428