Tutorial II

High Performance Clusters: State of the Art and Challenges Ahead

David E. Culler
culler@cs.berkeley.edu
Computer Science Division
University of California, Berkeley

Clusters have become established as a viable, flexible, and exciting architecture for parallel computing. The complete system on every node, including full I/O capability, virtual memory, and scheduling, provides a great deal of power, allowing cluster architectures to be applied to a wide range of interesting, data intensive problems. They are inexpensive and easy to build. Recent innovations allow fast, protected messaging in a general purpose environment, and clusters are on the verge of a broad commercial transformation with the Intel/Microsoft/Compaq standardization on the Virtual Interface Architecture, as well as the arrival of gigabit Ethernet. Clusters are no longer just a personal supercomputer.

This tutorial will have three parts. The first focuses on the current state of the art in commercial and research clusters. We will discuss current System Area Network technology (including Myrinet, ServerNet, and Gigabit Ethernet), network interface support (including the Virtual Interface Architecture and Memory Channels), fast communication abstractions (including Active Messages and Virtual Memory Mapped Communication), cluster software layers, and what is out there and easy to build. The second part will address conceptual developments that make clusters powerful and interesting from pragmatic and theoretical viewpoints, including virtualization, implicit scheduling, I/O streaming, resource management, and failure independence. The final part will be a discussion on future challenges and open problems. The material will be self-contained and accessible to a broad audience of researchers.