Notes

The Personal Supercomputing Revolution

Supercomputing is fast becoming a household term in many industries. This is due to a revolution in both architectures and usage models that is culminating in the emergence of an amazing new machine—the cluster workstation, also known as the world's first personal supercomputer. The cluster workstation is being pioneered by Orion Multisystems (http://www.orionmulti.com) of Santa Clara, California, USA.

Cluster workstations are continuing the evolution begun decades ago, an evolution in the direction of powerful and personal. Orion personal supercomputers fit the bill on both counts by providing the performance of a back-room cluster with the usage model of a personal computer.

Perhaps the most revolutionary part of this emergence is the expense, or lack thereof. Cluster workstations, and clusters in general, are based on open, inexpensive standards, such as commodity x86 (PC) hardware. Open-source software is also the rule, particularly the Linux operating system and GNU development tools. This lowers both entry and maintenance costs dramatically, and represents a strong departure from traditional supercomputing technology that relies on expensive custom multiprocessing environments.

Cluster Migration

Clusters are groups of individual computers networked together in such a way that they behave as one large computer for certain tasks. The kinds of tasks appropriate for this type of computer are those that can be easily performed in parallel. Jobs that are "parallelizable" include database searches across large collections of data, like genomes, chemical sequences, or weather data, as well as calculations on large information sets like weather, astronomical, or financial data, or graphics rendering. In general, any computational task that can be divided into smaller parts is appropriate for parallel processing. Those tasks whose parts do not depend on each other at all are known as "embarassingly parallel".

It is not by accident that these are the same types of problems approriate to larger supercomputers, as clusters were created to provide more manageable, less expensive solutions to the same problems.

X86-based clusters are fast becoming the de facto standard for high-performance computing. Sixty percent of the top 500 supercomputers are clusters, and a very

Notes

rapidly growing number of the top supercomputers are based on x86 standards, according to http://www.top500.org:

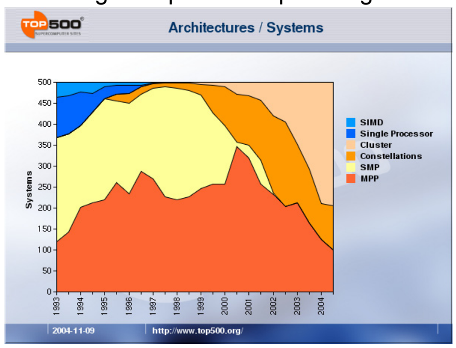

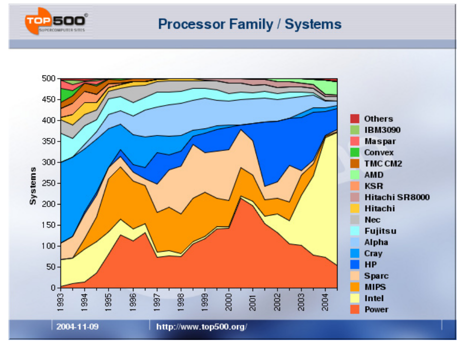

A standard cluster is managed by one node, the “head node”. The head node controls job distribution and monitoring, external network connectivity, disk access, and the user interface, while the other nodes do the work. Adding performance to the system is often no more difficult than adding inexpensive Linux computers to a network, although this still represents the “back-room” time-sharing usage model.

Linux, the freely available operating system, has revitalized high-performance computing similarly to its success in personal computing. Linux contains native support for clusters, making it the obvious choice for cash-strapped labs, universities, and burgeoning industries who need high-performance computing. Applications and supporting programs have exploded for this platform, with many of the most popular open-source and freely available, such as the MPI parallel programming interface, becoming standards themselves.

The Orion cluster workstation uses very low-power commodity x86 components to create a twelve-node cluster on a single board. This provides a very low-power, homogenous alternative to “back-room” clusters, which are often made up of very different individual machines individually powered. The most dramatic example is the DS-96 Deskside workstation, in which 8 Orion boards are combined to provide a 96-node cluster in a single case. The DS-96 fits unobtrusively under a desk, powered by a single plug in a standard wall socket with no special cooling considerations—and it draws less power than a hair dryer.

The Workstation Model

Traditional supercomputers required a separate dedicated room with special wiring and air conditioning, and often a fleet of maintenance personnel. This massive operational expensve yielded the obvious restriction that very few machines could be purchased, often only one. These systems operate on a “time-share” usage model. Those who need computing cycles had to wait in line—or compete with others—to get access. The time-share usage model was born.

Time-share systems enable organizations to centralize resources, but they naturally create conflict among groups who need access at the same time. Additionally, if one

Notes

project crashes the system, all jobs are lost. Current mid-range clusters are in a similar predicament, as they must also be relegated to a machine room for centralized maintenance, cooling, and wiring, and normally must operate on the time-share model.

In the middle 1980s, time-sharing minicomputers like the DEC VAX were popular options for the middle range of high-performance computing. A few years later, however, individual high-performance machines began to proliferate. These machines were called “workstations”. By themselves they were not as powerful as the hulking minicomputer down the hall, but they were able to be operated and maintained by one person—if they crashed, they took only one person’s job with them. Workstations rapidly became a standard by virtue of their vastly-improved usage model.

This same evolution is taking place in the cluster community. Orion has taken the standards agreed upon by the high-performance computing community—x86 clusters running Linux—and successfully applied them to the workstation usage model. Orion’s personal supercomputers are designed to be used by one person or small group—they have the performance of a back-room mid-range cluster, but a usage model similar to that of a modern PC.

Notes