

Recent Trends in Cluster Computing

Anne C. Elster*, Otto Anshus**, Amund Tveit* and Cyril Banino*

*Norwegian Univ. of Science & Technology (NTNU), Norway

**Univ. of Tromsø, Norway

This paper gives an overview of recent cluster computing technology trends with a focus on clusters that are used for high-performance computing (HPC). Here, a network of workstations and/or PCs is used rather than a traditional supercomputer. Although the network speed is still generally slow compared to available processing speed, and implementing distributed memory programs is generally harder than shared-memory programs, the price and flexibility of such systems cannot be underestimated. Tools like TSpaces that emulate the shared-memory paradigm on clusters, as well as standards like MPI and others, also make these systems worth a serious look in the coming years.

We discuss various hardware configurations from systems based on COTS (commercial off-the-shelf parts) such as PCs and the new Sony Playstation2-based cluster at University of Illinois Urbana-Champaign to the recent high-end SuperCluster (Altrix 3000) from SGI. Software packages, including performance tools, MPI and others are also discussed. We also describe recent benchmarking programs and algorithmic techniques that adjust to a given configuration.

Newer HPC areas such as using clusters for Massively Multiplayer Online Games (MMOGs) and parallel visualization are also introduced. Recent developments in clusters for GRID computing as well as fault-tolerant cluster computing are also included.

So far only a few very large clusters from purely commodity components have yet been introduced as general purpose supercomputers. Lately such clusters have been promoted as silver bullets in the price-performance war of HPC. Unfortunately, this has led to many users getting unrealistic expectations to the performance of commodity clusters. However, the tools and techniques are advancing and flexibility and proliferation of smaller to medium-sized clusters should not be underestimated.

There are many possible levels between classic supercomputers and clusters of commodity PCs. The component that is most often replaced with a non-commodity part is the network since significant improvement can be achieved by replacing Fast-Ethernet by a high-performance cluster-interconnect, i.e., Myrinet or Quadrics. Another often seen improvement is to replace standard PCs by more costly server-machines including multi-CPU systems. Another approach again is to keep the hardware completely commodity but replace the operating system or parts thereof with custom middleware.

The new MMOGs (Massively Multiplayer Online Games) should also open exciting opportunities and challenges. These systems involve parallel simulation environments where the domain may be distributed on several thousand processors each handling 100 or more users affecting the distributed domain. Together with the demand for GRID computing they indicate that the world of cluster computing is just beginning!