Peer-to-Peer Implementations

Alex Wild
Entropia Europe, Ltd
Alex.Wild@Entropia.com

Abstract

Distributed computing represents an extremely cost-effective way to gain supercomputer-scale power to run certain types of compute-intensive applications. Remarkably, the majority of a PC's time is spent doing nothing. The average PC is idle between 60 and 90%, even when it is being used.

Distributed computing platforms split large computational problems into many small tasks and distribute those tasks together with the algorithm to PCs connected within a corporate network or to the Internet. Applications integrated today include codes that run 'embarrassingly parallel', such as docking of small molecules to proteins or Fourier analysis of radio signals. Monte Carlo simulations and applications exploring multidimensional parameter spaces also fall into this category. More recently, 'divide and conquer' algorithms such as sequence alignment codes have been implemented. Eventually, it will become possible to distribute many more classes of applications using a variety of techniques.

Several large enterprises are currently deploying distributed computing technology. Very thorough standards in stability, security, manageability and scalability have to be met. The future will show a convergence and compatibility of different standards such as Globus and Entropia, leading to global computing grids with unprecedented computational capacity.

Proceedings of the First International Conference on Peer-to-Peer Computing (P2P'01)
0-7695-1503-7/02 $17.00 © 2002 IEEE