An Evaluation of Parallel Computing in PC Clusters with Fast Ethernet

M. Acacio, O. Cánovas, J.M. García, and P.E. López-de-Teruel

Dpto. de Ingeniería y Tecnología de Computadores
Universidad de Murcia
Campus de Espinardo, s/n, 30080 Murcia (Spain)
{meacacio,jmgarcia,pedroe}@ditec.um.es
ocanovas@fcu.um.es

Abstract. Currently, clusters of Personal Computers (PCs) with sharing Fast Ethernet may offer a significant price/performance ratio in parallel processing. This paper presents a study of the promising possibilities of parallelism in this environment, using both a classical numerical method to solve a differential equation and an innovative algorithm to solve the statistical problem of density estimation. In addition, and to compare, some comparisons with an IBM SP2 are carried out.

1 Motivations

The evolution of CPU technology has brought high-end Personal Computers (PCs) to performance levels in the range of workstations at a very competitive cost. So, PCs are already a very cost effective alternative to workstations as processing nodes in NOW platforms [2, 3]. The major limitation of PC clusters is the high communication overhead when exchanging messages between nodes. Unlike a multiprocessor that has a custom low latency/high bandwidth network, nodes in a PC cluster are loosely connected by a LAN with a lower bandwidth.

Our work is situated in the context of an ambitious research project, called ELDORADO, whose aim is to show how using a high-performance comunication channel as the underlying network, clusters of powerful PCs can even overtake traditional parallel machines in a wide range of applications. We show the performance of two real problems implemented with MPI: a numerical algorithm (the parallel Jacobi method for solving differential equations) and a new parallel algorithm used for non-parametric density estimation from uncertain samples. Both algorithms have been successfully executed in a cluster of PCs with Fast Ethernet, and the results we have achieved are very good. A comparison with an IBM SP2 parallel computer has been performed for the first algorithm.

2 Preliminary Results, Conclusions, and Future Work

Jacobi relaxation method is an algorithm for solving a differential equation called Laplace's equation. We have used it in the iterative algorithm that calculates

the temperature of each particle in a body. The EDR algorithm [1], on the other side, is an extension of the Parzen method. It is used for nonparametric density estimation when working with uncertain input samples. We have developed and measured MPI parallel versions of these algorithms. Our tests have been carried out on a cluster of Intel Pentium 200 MHz processor (32 MB RAM) interconnected by a Fast Ethernet 3Com 905-network. The operating system was Linux 2.0.32 and the MPI implemetation MPICH v.1.0.13. Programs were compiled with the best optimization options. Figure 1 shows some illustrative results.

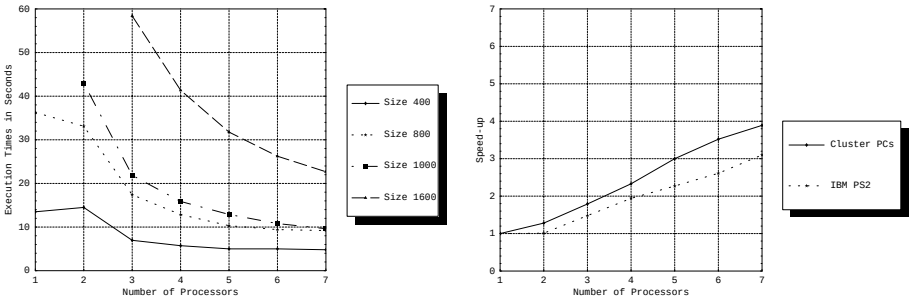

Fig. 1. Execution times in seconds for several sizes of the problem in the EDR algorithm, and speed-up comparison with IBM SP2 for the Jacobi method

As a conclusion, we can say that parallel computing on clusters of workstations and PCs has very high potential, since it takes advantage of existing hardware and software. Performance tests of the implementations show that they are superior to much existing parallel programming environments for some application problems.

Our research group is currently working in practical applications in different fields, such as real-time computer vision, and computationally expensive statistical and machine learning problems.

This work has been partially supported by the Spanish CICYT under grant TIC97-0897-C04-02.

References

1. López-de-Teruel, P.E., Acacio, M.E., Cánovas, O., García, J.M.: EDR: An Algorithm for Parallel Implementation of Parzen Density Estimation from Uncertain Observations. Submitted to the 2nd Merged Symposium IPPS/SPDP (1999)
2. Nupairoj, N., Ni, L.: Performance Evaluation of Some MPI Implementations on Workstations Clusters. Procceedings of the 1994 Scalable Parallel Libraries Conference, (1994) 98–105
3. Wang, X. Blum, E.K.: Parallel Execution of Iterative Computations on Workstations Clusters. Journal of Parallel and Distributed Computing, vol. 34 (1996) 218–226