

Runtime Data Declustering over SAN-Connected PC Cluster System

Masato Oguchi^{1,2} and Masaru Kitsuregawa²

¹ Research and Development Initiative, Chuo University
42-8 Ichigaya Honmura-cho, Shinjuku-ku Tokyo 162-8473, Japan

² Institute of Industrial Science, The University of Tokyo
oguchi@computer.org

Recently, personal computer/workstation (PC/WS) clusters have come to be studied intensively in the field of parallel and distributed computing. In the viewpoint of applications, data intensive applications including data mining and ad-hoc query processing in databases are considered very important for massively parallel processors, in addition to the conventional scientific calculation. Thus, investigating the feasibility of such applications on a PC cluster is meaningful. A PC cluster connected with Storage Area Network(SAN) is built as shown in Figure 1, and evaluated with data mining application.




Figure 1. An overview of SAN-connected PC cluster pilot system

In the case of SAN-connected cluster, each node can access all shared disks directly without using LAN; thus, SAN-connected clusters achieve much better performance than LAN-connected clusters for disk-to-disk copy operations. However, if a lot of nodes access the same shared disk simultaneously, application performance degrades due to I/O-bottleneck. A runtime data declustering method, in which data is declustered to several other disks dynamically during the execution of application, is proposed to resolve this problem.

An association rule mining is implemented and evaluated on the SAN-connected PC cluster pilot system. This application requires iterative scans of a shared disk, which degrade execution performance severely due to I/O-bottleneck. The runtime data declustering method is applied to the execution of application. In this method, portion of the data is copied to multiple disks which is used exclusively when the data is read in the first pass, then the copied data, instead of the original one, is accessed afterward.

The number of transaction is 20,000,000, the number of different items is 5,000, the minimum support is 0.7%, and the number of nodes used in this application is 16 in this experiment. The execution time of the proposed runtime data declustering method are shown in Figure 2. The proposed method is compared with the original one in which data is read only from the same shared disk repeatedly.

Application performance has improved since pass 3 because they are I/O-bound condition in the original method. According to the results of experiment, the proposed method prevents performance degradation caused by shared disk bottleneck in SAN-connected clusters.




Figure 2. Execution time of association rule mining (declustering from 1 to 16 disks)