An Efficient Way of Distributing Data on Parallel Memory Systems *

Sajal K. Das
Department of Computer Sciences
University of North Texas
Denton, TX 76203, USA

M. Cristina Pinotti
IEI, Consiglio Nazionale delle Ricerche
Via S. Maria, 46
56126 Pisa, ITALY

Abstract

A *perfect* memory system should supply immediately any datum requested by the CPU. However, since such an ideal memory is not practically feasible, this goal is commonly approximated by organizing the memory hierarchically and using a parallel memory system consisting of several modules (banks of main memory and/or I/O disks) at each level of the hierarchy. In this paper, we develop efficient data distribution schemes as a technique for improving the performance of a parallel memory system. The objective is to fetch in a single (or as few as possible) memory cycle specific subsets of nodes, called *templates*), of a given data structure, called a *host*. If a template to be retrieved is stored in a single memory module, its nodes can only be fetched sequentially, one after the other, thus incurring a very high access latency. This is independent of the number of memory modules available in the parallel memory system and hence, the potential system bandwidth. So, our goal is to distribute the host nodes evenly among the memory modules so as to minimize the number of nodes of a given template assigned to the same memory module. Interestingly, to retrieve a template in just one cycle (i.e., no conflicts), the absolute minimum number of memory modules equal to the template size is not always sufficient.

After defining the data access problem and its optimality criteria, we propose efficient schemes for conflict-free access to subtrees in k-ary trees and binomial trees. The proposed assignments, based on simple indexing, are optimal in terms of the number of memory modules required, and they balance the memory load among the memory modules. We also consider the problem of distributing n-dimensional vector spaces (hypercubes, Q_n) in such way that k-dimensional subspaces (subcubes) can be accessed efficiently. The strict correlation between this problem and the theory of error correcting codes is stated as follows: declustering the nodes of Q_n according to a linear code C of length n, size M and minimum distance d, any occurrence of a subcube $Q_s \subset Q_n$, for $0 \leq s \leq d-1$, can be accessed without conflicts using $\lceil \frac{2^n}{M} \rceil$ memory modules.

*This work is supported by Texas Advanced Technology Program grant under Award No. TATP-003594031. The authors can be reached at these e-mail addresses: das@cs.unt.edu, pinotti@iei.pi.cnr.it.