Concurrent and Distributed Data Structures for Multikey Sorting on Computer Clusters

Extended Abstract

Abdelaziz Fellah
Dept. of Mathematics & Computer Science, University of Lethbridge
Lethbridge, AB, Canada, T1K 3M4
fellah@cs.uleth.ca

Parallel computing has long offered the promise of very high performance for a wide range of applications, but unfortunately it is hampered by several drawbacks. Among the drawbacks of exploiting parallelism is the cost of message-passing and communication overhead. Consequently, the trend in parallel and distributed computations is moving away from specialized platforms, such as the Cray/SGI T3E, Thinking Machine CM-5, and IBM SP2 to inexpensive, general purpose parallel machines consisting of clusters of workstations or PC's with Gigabit Ethernet and Myrinet networking. Over the next several years, computer clusters will likely be the predominant architecture for scalable high-performance computing; however, still there is much work to be done in several applications to exploit and support high-performance computing. Computer clusters have triggered the investigation of several high-performance algorithms for scientific applications in both research and industry, for example, pattern recognition, cryptology, combinatorial computations, sorting very large or massive amount of data. Sorting and searching multi-dimensional data on more than one key and in different order can be used in many important applications, such as data base management systems, image processing, geographical information systems and web-search engines.

This paper focuses on theoretical and practical aspects of the high-performance multikey sorting problem on computer clusters, with particular emphasis on the Alpha Maci Cluster, a world-class high-performance supercomputer that has many processors interconnected by a wide range of high-speed network connections. Even though the focus of this paper is on multikey sorting problems, developing new data structures and techniques for designing high-performance algorithms on computer clusters are of both theoretical and practical interest. We investigate strategies for developing, implementing, and refining high-performance algorithms for sorting multi-dimensional data on computer clusters. In addition, maximizing the performance of such distributed memory machines requires efficient data structures coupled with good load balancing. Similar to one-dimensional keys, multi-dimensional keys (multikey, for short) are subject to lexicographic order and must be suitably stored, and efficiently retrieved, sorted and maintained. We show that the multikey sorting algorithm is very amenable to performance improvement through three parameters: (1) efficient distributed and concurrent data structures, (2) exploitation of parallelism, and (3) high-performance capabilities of computer cluster architectures. We introduce five concurrent and distributed multikey structures that we refer to as α-*trees* (*alpha trees*), α-*parallel trees*, α-*search trees*, α-*lists*, and α-*matrices*. An $O(n)$ space algorithm is presented for constructing α-trees and α-matrices, where n is the number of multi-dimensional records to be sorted. We propose an $O(\frac{n}{\log n})$ space algorithm for constructing a hybrid of α-trees and α-matrices. We prove that the maximum depth of α-search tree structures is $\lceil \log n \rceil$. In addition, we show that these distributed data structures facilitate the development and execution of parallel programs and can be used for concurrent operations. We are currently implementing several of our algorithms and will report performance results in the near future.

References

[1] A. Fellah, M. Abaza, and A. Maamir. "Efficient Data and Workload Balancing for Multikey Sorting on Hypercube Multicomputers", *Proc. of the Int. Conf. on Applied Informatics,* Innsb., Austria, 2001, 355-358.

[2] A. Fellah, M. Abaza, and A. Maamir. "Distributed Data Structures for Multikey Sorting", *Int. J. of Parallel and Distributed Systems and Networks*, 1999, Vol. 2(2), 62-68.

[3] J.L. Bentley and R. Sedgewick. " Fast Algorithms for Sorting and Searching Strings", *Proc. of the 8th ACM-SIAM Symp. on Discrete Algorithms,* 1997, 360-369.

[4] D.R. Helman, D.A. Bader, and J. JáJá. "Sorting on Clusters of SMPs", Informatica, 1999, Vol. 23(1), 113-121.

Proceedings of the 16th Annual International Symposium on High Performance Computing Systems and Applications (HPCS'02)
0-7695-1626-2/02 $17.00 © 2002 IEEE