

CS-1988-23

Prefetching in File Systems for MIMD Multiprocessors

Carla Schlatter Ellis
David Kotz

Department of Computer Science
Duke University
Durham, North Carolina 27708-0129

November 1988

Prefetching in File Systems for MIMD Multiprocessors

Carla Schlatter Ellis
David Kotz

July 1988

Abstract

The problem of providing file I/O to parallel programs has been largely neglected in the development of multiprocessor systems. There are two essential elements of any file system design intended for a highly parallel environment: parallel I/O and effective caching schemes. This paper concentrates on the second aspect of file system design and specifically, on the question of whether prefetching blocks of the file into the block cache can effectively reduce overall execution time of a parallel computation. MIMD multiprocessor architectures have a profound impact on the nature of the workloads they support. In particular, it is the collective behavior of the processes in a parallel computation that often determines the performance. The assumptions about file access patterns that underlie much of the work in uniprocessor file management are called into question. Results from experiments performed on the Butterfly Plus multiprocessor are presented showing the benefits that can be derived from prefetching (e.g. significant improvements in the cache miss ratio and the average time to perform an I/O operation). We explore why it is not trivial to translate these gains into much better overall performance.

This technical report, although still available on hardcopy from Duke University, has been superseded by [EK89] and later published more thoroughly as [KE90]. For yet another presentation and many more results, see also [Kot91, KE91b, KE92b, KE91a, KE92a, Kot92].

For more information, write to

Technical Reports
Department of Computer Science
Duke University
Durham, North Carolina 27708-0129
techreports@cs.duke.edu

References

- [EK89] Carla Schlatter Ellis and David Kotz. Prefetching in file systems for MIMD multiprocessors. In *Proceedings of the 1989 International Conference on Parallel Processing*, pages I:306–314, August 1989.
- [KE90] David Kotz and Carla Schlatter Ellis. Prefetching in file systems for MIMD multiprocessors. *IEEE Transactions on Parallel and Distributed Systems*, 1(2):218–230, April 1990.
- [KE91a] David Kotz and Carla Schlatter Ellis. Caching and writeback policies in parallel file systems. In *1991 IEEE Symposium on Parallel and Distributed Processing*, pages 60–67, December 1991. To appear in the *Journal of Parallel and Distributed Computing*.

- [KE91b] David Kotz and Carla Schlatter Ellis. Practical prefetching techniques for parallel file systems. In *Proceedings of the First International Conference on Parallel and Distributed Information Systems*, pages 182–189, December 1991. To appear in *Distributed and Parallel Databases*.
- [KE92a] David Kotz and Carla Schlatter Ellis. Caching and writeback policies in parallel file systems. *Journal of Parallel and Distributed Computing*, January 1992. Submitted.
- [KE92b] David Kotz and Carla Schlatter Ellis. Practical prefetching techniques for multiprocessor file systems. *Distributed and Parallel Databases*, 1992. To appear.
- [Kot91] David Kotz. *Prefetching and Caching Techniques in File Systems for MIMD Multiprocessors*. PhD thesis, Duke University, April 1991. Available as technical report CS-1991-016.
- [Kot92] David Kotz. Multiprocessor file system interfaces. Technical Report PCS-TR92-179, Dept. of Math and Computer Science, Dartmouth College, March 1992. Abstract appeared in 1992 Usenix Workshop on File Systems.