CC PUB #268

January 1999

Random RAIDs with Selective Exploitation of Redundancy for High Performance, Agile Stream Servers

Yitzhak Birk and Noam Bloch

Department of Electrical Engineering
Technion - Israel Institute of Technology
Haifa 32000, Israel

Abstract

This paper jointly addresses the issues of load balancing, fault tolerance, responsiveness, agility, streaming capacity and cost-effectiveness of high-performance storage servers for data-streaming applications such as video-on-demand. Stripping the data of each movie across disks in a "random" order balances the load while breaking any correlation between user requests and the access pattern to disks. Parity groups are of fixed-size, comprising consecutive blocks of a movie and a derived parity block, and resulting in "random" disk-members of any given group. Consequently, the load of masking a faulty disk is shared by all disk drives, minimizing the degradation in streaming capacity. By using the redundant information to avoid accessing an overloaded disk drive, the occasional transient imbalance in disk load due to the randomization is partly prevented and, when occurring, can be circumvented. Finally and most important, making a distinction between data blocks and redundant blocks and using redundant blocks only when necessary is shown to substantially reduce required buffer sizes without giving up the benefits. The result is a simple, flexible and robust video-server architecture. The paper includes rationale and insights, as well as simulation results and analytical support for the observed behavior.