Scalable Computing Software Laboratory

Department of Computer Science

Illinois Institute of Technology

A Taxonomy of Data Prefetching Mechanisms

Surendra Byna **Yong Chen** **Xian-He Sun**

sbyna@iit.edu chenyon1@iit.edu sun@iit.edu

May 2007

Technical Report № IIT/CS-SCS07-01

http://www.cs.iit.edu

10 West 31st Street, Chicago, IL 60616

A Taxonomy of Data Prefetching Mechanisms[1]

Surendra Byna† **Yong Chen†** **Xian-He Sun†‡**

†Computer Science Department
Illinois Institute of Technology, Chicago, IL 60616

‡Computing Division
Fermi National Accelerator Laboratory, Batavia, IL 60510
{sbyna, chenyon1, sun}@iit.edu

Abstract

Data prefetching has been considered as an effective way to mask data access penalty caused by cache misses and to bridge the performance gap between processor and memory. With hardware and/or software support, data prefetching brings data closer to processor before it is actually needed. Many prefetching techniques have been developed to reduce data access latency for single-core processor. While some of these techniques are directly applicable to multi-core processors, many novel strategies have been proposed in the last few years to take advantage of multi-core architectures. In this paper, we propose a taxonomy that characterizes and classifies various design concerns in developing a prefetching strategy. We also compare pros and cons of various prefetching strategies, and discuss several representative prefetching projects.

Keywords: taxonomy of prefetching strategies, multi-core processors

[1] This research is supported in part by National Science Foundation under NSF CCF-0621435.