

Scalable Computing Software Laboratory Technical Report

Department of Computer Science

Illinois Institute of Technology

Exploring Parallel I/O Concurrency with Speculative Prefetching

Yong Chen, Surendra Byna, Xian-He Sun, Rajeev Thakur, William Gropp

{chenyon1, sbyna, sun}@iit.edu, thakur@mcs.anl.gov, wgropp@uiuc.edu

November 2007

Technical Report № IIT/CS-2007-21

<http://www.cs.iit.edu>

10 West 31st Street, Chicago, IL 60616

LIMITED DISTRIBUTION NOTICE: This report has been submitted for publication outside of IIT-SCS and will probably be copyrighted if accepted for publication. It has been issued as a Technical Report for early dissemination of its contents. In view of the transfer of copyright to the outside publisher, its distribution outside of IIT-SCS prior to publication should be limited to peer communications and specific requests. After outside publication, requests should be filled only by reprints or legally obtained copies of the article (e.g. payment of royalties).

Exploring Parallel I/O Concurrency with Speculative Prefetching

Yong Chen¹ Surendra Byna^{1,2} Xian-He Sun¹ Rajeev Thakur² William Gropp³

¹ *Department of Computer Science, Illinois Institute of Technology, Chicago, IL*

{chenyon1, bynasur, sun} @iit.edu

² *Mathematics & Computer Science Division, Argonne National Laboratory, Argonne, IL*

thakur@mcs.anl.gov

³ *Department of Computer Science, University of Illinois Urbana-Champaign, Urbana, IL*

wgropp@uiuc.edu

Abstract

Parallel applications can benefit greatly from massive computational capability, but their performance usually suffers due to large latency in I/O accesses. Conventional I/O prefetching techniques are conservative and are limited by low accuracy and coverage. As the processor performance has been increasing rapidly and the computing power is virtually free, we introduce a novel speculative approach for comprehensive and aggressive parallel I/O prefetching in this study. We present the design of our approach, as well as challenges, solutions, and the implementation on ROMIO of MPICH2. The experiments have shown promising results in reducing I/O access latency.

Keywords: high performance computing, I/O performance, parallel application performance, speculative prefetching