

SERVER-BASED DATA PUSH ARCHITECTURE FOR DATA ACCESS PERFORMANCE OPTIMIZATION

Surendra Byna, Ph.D.

Illinois Institute of Technology, December 2006

Adviser: Dr. Xian-He Sun

In the past several years, High-End Computing (HEC) has seen enormous growth in peak performance, and development of Peta-flop supercomputer is in the near horizon. Despite these advances, data access delay has been a major reason for poor sustained system performance (SSP) on HEC machines. Multiple levels of memory hierarchy have been incorporated into computer architecture to take advantage of locality among data accesses to reduce the gap between peak and sustained performances. However, many applications lack locality, which make these advances inefficient. Researchers have proposed many optimization methods to improve locality and to prefetch data into these cache memories before CPU demands for it. However, there are limitations in applying these methods. First, locality is application dependent and choosing an efficient combination among all existing tuning methods at runtime remains elusive. Second, the current client-initiated prefetching strategies do not work well for applications with complex, non-contiguous data access patterns.

To bridge the performance gap, we introduce server-based data push architecture. In this architecture, a dedicated server named Data Push Server (DPS) initiates and proactively pushes data closer to the processing units in time. We addressed the issues of

monitoring data access history, making spatial and temporal access pattern predictions, architecture modifications to push the predicted data values close to processing cores, and modeling data access cost. We have quantified data access cost from communication and middleware latencies. We present analytical models for memory performance prediction based on data access patterns that are useful to choose effective optimization and prefetching strategies with low overhead. We have applied these models to improve the performance of Message Passing Interface (MPI) derived datatypes. We have studied the server-push architecture by enhancing SimpleScalar simulator with a dedicated processing unit that pushes data for another processor. The simulation results show significant performance gains. Our DPS architecture is extendable to various levels of memory hierarchy, and has a broader impact on high-end computing to improve productivity.