Effectiveness of Aggessive Out of Order Scheduling in the IBM PowerPC Processors

Lizy K. John and Juan Rubio
The Laboratory for Computer Architecture
Department of Electrical and Computer Engineering (UT-ECE)
The University of Texas at Austin

In Collaboration With: Alex Mericas, IBM Austin Server Group

Abstract

The past decade has seen the development of several out of order, dynamically scheduled superscalar microprocessors. While they have been commercially successful, there has been some concerns on the effectiveness of out of order execution for commercial workloads. Commercial workloads including web werver, E-commerce and data base workloads are extremely important for future IBM server microprocessors. The primary objective of the proposed project is to determine the usefulness of aggressive out-of-order scheduling in the IBM PowerPC (PPC) processors. The study will utilize the IBM RS64-II, PPC 604e and POWER3 processors. The performance counters in the respective processors will be used to conduct the measurements. The three aforementioned processors cover the spectrum of PPC machines from in-order microarchitectures to aggressively out of order and speculative microarchitectures. The primary objective of our study will be to measure and identify components of processor performance, and study the effects of speculative and out-of-order execution on resource usage and overall performance. Another objective of the research would be to comment on the performance monitoring facility of the PowerPC processors. Overall, the study would result in suggesting useful design directions for future IBM server microprocessors, particularly those aimed at web server, data base and E-commerce workloads.

OBJECTIVES

The past decade has seen the development of several out of order, dynamically scheduled superscalar microprocessors. While they have been commercially successful, there has been some concerns on the effectiveness of out of order execution for commercial workloads. In 1998, Keeton et al [Ke98] conducted a study on On-Line Transaction Processing (OLTP) workloads on the Intel PentiumPro processor. They examined the effectiveness of out of order execution, branch prediction, superscalar issue and retire, caching and multiprocessor scaling, using the performance monitoring counters on the PentiumPro. They concluded that out of order execution is only moderately effective for OLTP workloads. It was also seen that PentiumPro's branch prediction scheme is not nearly as effective for database workloads as it is for SPEC workloads. At about the same time, Radhakrishnan and Rawson [Ra98] conducted a study of three web server workloads on two x86 microarchitectures, the Pentium and the PentiumPro. The Pentium is an in order two-issue processor while the PentiumPro uses

speculative, out of order execution. The PentiumPro was seen to lower the program execution time and cycles per instruction, illustrating exploitation of speculation and dynamic scheduling.

Thus past research is not very conclusive about the usefulness of out of order execution techniques on commercial workloads. It is extremely important to examine these results on other platforms and confirm or contradict the previous inferences. Three of IBM's recent PowerPC processors, the RS64-II, the PPC 604e and the POWER3 cover a nice spectrum of PPC machines from in-order microarchitectures to aggressively out of order and speculative microarchitectures. All these machines are equipped with scores of on-chip performance monitoring counters making them nice platforms to pursue this question on the effectiveness of out of order execution for commercial workloads. Hence in this project, we undertake the task of examining the effectiveness of branch prediction and out of order execution on three IBM PowerPC microarchitectures.

The primary objective of our study will be to measure and identify components of processor performance (CPI), and study the effects of speculative and out-of-order execution on resource usage and overall performance. Using performance counter data, an analysis will be done to identify the performance bottlenecks.

Another objective of the research would be to suggest design directions for future IBM server microprocessors. The data obtained in the various experiments will be analyzed to draw conclusions on memory access patterns, effectiveness of caches and dat prefetching, characteristics of data sharing, non-cacheability, I/O access patterns, SMP scaling characteristics, etc. Keeton's study [Ke98] also made a few other observations. Caches were found to be effective at reducing processor memory traffic. Only 0.33% of database accesses and 0.47% of OS accesses reached memory. They also found that exclusive state of the four-state MESI cache coherence policy is underutilized and a three state (MSI) cache coherence protocol may be sufficient. We will investigate various such design tradeoffs for future microprocessors for E-commerce and web server workloads.

A third objective of the research would be to comment on the performance monitoring facility of the PowerPC processors. Controllability and observability of desired performance evaluation points will be analyzed. Other useful additions to the performance monitors will be suggested. Ofcourse, it will be extremely important to keep the invasiveness of the counters to be minimal. We will also evaluate whether some of the features that exist in other processors including Compaq Alpha's ProfileMe facility [De97] would have made a big difference to the easiness of studies such as this.

EXPERIMENT METHODOLOGY

The study will utilize the performance counters in the IBM PowerPC processors. IBM has developed an API that can be used to program and collect data from any PowerPC performance monitor. Calls from this API will be used to interface with the performance monitor.

Platforms

Three different PowerPC processors [Pa98, Po98, Bo99] with varying amounts of aggressiveness in code scheduling will be used for ths study. The **PPC 604e** is a superscalar, out-of-order , speculative execution, in-order completion machine. It can dispatch up to four and retire up to four instructions per cycle. For branch prediction, it uses a 64 entry BTAC and a 512 entry, two-bits-per-entry branch history table. The RS64-II, interestingly, is a superscalar, in-order, speculative execution machine which can fetch, dispatch, and retire up to 4 instructions per cycle. The **RS64-II**, does not employ any sophisticated branch prediction, but simply predicts every branch to be not taken. Once the branch is resolved in the dispatch stage, it either continues fetching from the current instruction stream with no penalty or flushes the instructions after the branch and starts fetching from the branch target buffer, with a penalty of at most one and often zero cycles The **POWER3**, the most complex of the three processors, is a superscalar, out-of-order, speculative execution (speculative loads but not stores), in-order completion machine. It can fetch eight instructions simultaneously and dispatch up to eight in the same cycle. It has a 256 entry BTAC and a 2048 entry (bits per entry not available) branch history table for use in branch prediction. The three microarchitectures are very different in their aggressiveness and form interesting platforms for a study on effectiveness of out of order scheduling and branch prediction.

Benchmarks

The web server workloads used in Radhakrishnan et. al. [Ra98] [Ra99] will be used as the starting point for our study. These include static and dynamic web server applications. SPECweb99 workloads will be the next target. We are also setting up a TPC-C workload environment. Effort to obtain real e-commerce workloads will also be made in the long run.

STATUS

The characterization of web server benchmarks is progressing as described in the companion paper [Pa00]. SPECweb99 has been aquired and will be the next set of benchmarks. A TPC-C like workload on a DB-II system is being set up. Results from experiments from the IBM groups with TPC workloads and web server workloads will be correlated with the results obtained from this research activity.

FUTURE WORK

Effort will be made to obtain more interesting commercial workloads. Help from various IBM groups is solicited in this matter. Based on this study, we hope to suggest many useful design directions for future IBM server microprocessors, particularly those aimed at web server, data base and E-commerce workloads.

The Laboratory of Computer Architecture (LCA) [lca] has performed several characterization studies on multimedia, web server, and Java applications. The expertise of the laboratory members will be utilized to provide a meaningful evaluation of the aforementioned architectural issue on the IBM processors.

REFERENCES

[Bo99] J. Borkenhagen, and S. Storino, "4th Generation 64-bit PowerPC-Compatible Commercial Processor Design", White Paper, IBM Corporation, http://www.rs6000.ibm.com/resource/technology/nstar.html.

[De97] J. Dean, J. E. Hicks, C. A. Waldspurger, W. E. Weihl, and G. Chrysos, "ProfileMe: Hardware Support for Instruction-Level Profiling on Out-of-Order Processors", Proceedings of MICRO-30, 1997, pp. 292-302.

[Ke98] Kimberly Keeton, David A. Patterson, Yong Quiang He, Roger C. Raphael, and Walter E. Baker, "Performance Characterization of a Quad PentiumPro SMP Using OLTP Workloads", Proceedings of the International Symposium on Computer Architecture, Barcelona, Spain, 1998, pp. 15-26.

[lca] The Laboratory for Computer Architecture, Electrical and Computer Engineering Department, The University of Texas at Austin, http://www.ece.utexas.edu/projects/ece/lca/ Publications available at http://www.ece.utexas.edu/projects/ece/lca/publications.html

[Pa00] Pattabi Seshadri and Lizy K. John, "Characterization of Web Workloads for Three Generations of IBM Power PC Microarchitectures", Digest IBM Center for Advanced Studies Conference (this report), July 18, 2000.

[Pa98] M. Papermaster, R. Dinkjian, M. Mayfield, P. Lenk, B. Ciarfella, F. O'Connell, and R. DuPont, "POWER3: Next Generation 64-bit PowerPC Processor Design", White Paper, IBM Corporation, 1998.

[Po98] "PowerPC 604e Microprocessor User's Manual", IBM Corporation and Motorola Corporation, 1998.

[Ra98] Ramesh Radhakrishnan and Freeman L. Rawson, III, "Characterizing the Behavior of Windows NT Web Server Workloads Using Processor Performance Counters", Presented at the First Workshop on Workload Characterization, November 1998, Dallas, Texas, Also appears in Workload Characterization: Methodology and Case Studies, edited by Lizy Kurian John and Ann Marie Maynard, IEEE Computer Society Press, 1999, pp. 76-85.

[Ra99] Radhakrishnan, R. and John, L. K., "A Performance Study of Modern Web Applications," Proceedings of the European Parallel Processing Conference (Euro-Par 99), Lecture Notes in Computer Science, Springer, pp. 239-247.