2000 Sixth International Symposium on High-Performance Computer Architecture

Proceedings

Sixth International Symposium on High-Performance Computer Architecture HPCA-6

January 8–12, 2000
Toulouse, France

Sponsored by
IEEE Computer Society
Technical Committee on Computer Architecture

With the support of
Région Midi-Pyrénees
Centre National de la Recherche Scientifique
Mairie de Toulouse
Silicon Graphics, Inc.
Sun Microsystems
Hewlett-Packard
IEEE France
U.S. National Science Foundation

IEEE COMPUTER SOCIETY

Los Alamitos, California
Washington • Brussels • Tokyo

IEEE Computer Society Order Number PR00550
ISBN 0-7695-0550-3
ISBN 0-7695-0552-X (microfiche)
IEEE Order Plan Catalog Number PR00550
Library of Congress Number: 99-069213

Additional copies may be ordered from:

IEEE Computer Society
Customer Service Center
10662 Los Vaqueros Circle
P.O. Box 3014
Los Alamitos, CA 90720-1314
Tel: + 1-714-821-8380
Fax: + 1-714-821-4641
E-mail: cs.books@computer.org

IEEE Service Center
445 Hoes Lane
P.O. Box 1331
Piscataway, NJ 08855-1331
Tel: + 1-732-981-0060
Fax: + 1-732-981-9667
mis.custserv@computer.org

IEEE Computer Society
Asia/Pacific Office
Watanabe Building,
1-4-2 Minami-Aoyama
Minato-ku, Tokyo 107-0062 JAPAN
Tel: + 81-3-3408-3118
Fax: + 81-3-3408-3553
tokyo.ofc@computer.org

Editorial production by A. Denise Williams

Cover art production by Alex Torres

Printed in the United States of America by The Printing House, Inc.

IEEE COMPUTER SOCIETY

Proceedings

Sixth International Symposium on High-Performance Computer Architecture HPCA-6

Message from the General Co-Chairs

Welcome to Toulouse and to the 6th International Symposium on High-Performance Computer Architecture (HPCA-6). We hope that your stay here is both technically rewarding and enjoyable.

HPCA is already established as an outstanding forum for the presentation of the latest research contributions in computer architecture. This year, we have given the conference a more international flavor by holding it outside of the United States for the first time. We hope that, as a result, the conference will benefit technically and the participants can find international links for joint research.

Many people deserve credit for making this conference a success. First and foremost, we acknowledge the efforts of the Program Chair, Kai Li, and his program committee. Their dedication has, once again, resulted in a program of the highest technical excellence. The authors of the accepted papers are to be congratulated.

In addition to the main conference, we have a set of interesting workshops and tutorials. For this, we must acknowledge the efforts of the organizers and of the Workshop and Tutorial Chair, Lawrence Rauchwerger.

Several people in the organizing committee have devoted many hours to ensure the success of this conference. In particular, we thank our tireless Local Arrangements Chair, Pascal Sainrat. He has taken care of a myriad of things and has been the real engine behind the organization of this conference. We also thank the efforts of the Registration Chair, Franck Cappello; the Publicity Chair, Ahmed Louri; our webmaster in Princeton, Sanjeev Kumar; and Mark Franklin, who provided experienced advice.

Generous financial support was provided by our sponsors, including IEEE, IEEE Computer Society, IEEE Computer Society Technical Committee on Computer Architecture, U.S. National Science Foundation, Centre National de la Recherche Scientifique, Mairie de Toulouse, Region Midi-Pyrenees, Silicon Graphics, Sun Microsystems, and Hewlett-Packard. We also thank IEEE France for providing awards for the best papers.

In addition to enjoying the technical program, we encourage you to explore Toulouse. The city is the aeronautics and space capital of Europe, as it has the Airbus assembly plant, the Centre National d'Etudes Spatiales, Matra Marconi Space and Alcatel Space Industries facilities, and the first European aeronautics and space museum. Add to this many universities, and the result is a highly-qualified workforce in Information Technology.

At the same time, Toulouse is rich in history and architecture. It is known as the Rose Town for the color of the brick walls in its palaces and monuments. By boat, you can reach the Mediterranean Sea on the Canal du Midi, built in the 16th century by Riquet, and now belonging to the Heritage of Humanity. Other attractions are the beautiful Pyrenees Mountains, which can be reached in one hour by car, and good food and wines, especially the typical "France Sud-Ouest'' products such as cassoulet and foie-gras.

Overall, we hope that you will enjoy your stay, and we thank you for making HPCA-6 a success!

Daniel Litaize and Josep Torrellas
General Co-Chairs

Message from the Program Chair

It is my great pleasure to present this collection of papers for the Sixth International Symposium on High-Performance Computer Architecture (HPCA-6).

HPCA continues to be an international conference that attracts high-caliber submissions from institutions around the world. We received 163 submissions, the largest number ever. A total of 426 authors are from 125 institutions in 26 countries. I would like to thank all the authors for submitting high-quality papers to this conference and for their willingness to submit papers in PDF format to ease the review process. In most cases, each paper was reviewed by three program committee members and one external reviewer. We received a total of 644 reviews (3.95 reviews per paper) from 184 reviewers whose names are listed in the proceedings. I would like to thank all of them for their detailed comments and suggestions.

We used a perl script package, which was used by several program committees in computer architecture, to process the reviews. I would like to thank Hank Levy for his initial development of the package, and Guri Sohi for giving me the latest version of the package. Sanjeev Kumar served as our perl script and web wizard. His efficiency, patience, and dedication greatly helped the entire review process.

I am fortunate to have worked with 25 knowledgeable, dedicated colleagues in the program committee. The program committee meeting was held at Princeton on September 17, 1999. Unfortunately, Hurricane Floyd came to the east coast the previous evening. Although all flights to the east coast were cancelled and several roads were flooded, 11 program committee members came to Princeton and 6 others called in (our meeting lasted ten hours that day). In order to make sure that every program committee member had a chance to participate in discussions of the papers they read, we arranged three conference calls during the week of September 20th, and each lasted more than two hours. The final decision was made on September 24th. All committee members were provided with a book containing all reviews and numerical scores for all papers without authors' names, allowing examination of reviews by the entire committee. In general, we processed the papers in the order of their numerical rankings, but we used only the rankings as a guide. Each paper was presented by a program committee member who had read it. The decision on each paper was made only after all program committee members who had read it expressed their opinions to the program committee and only after the program committee reached a consensus. Papers co-authored by program committee members were handled in special sessions in which the authors were not present. The reviews for these papers were compiled separately. The committee accepted 8 out of 14 papers authored by committee members.

In the end, the program committee accepted 35 papers for the conference. For 8 of the papers, a program committee member was asked to "shepherd" the revision process. The result is the excellent program you find in the proceedings.

I would like to thank all the program committee members for their dedication and their efforts in this unusually long paper selection process. I owe a great debt of gratitude to each of the program committee members. Without their extraordinary efforts and support, putting such an excellent program together would have been impossible. I would like to thank the general chairs, Daniel Litaize and Josep Torrellas, for the organization of this conference. I would also like to thank Pascal Sainrat and Lawrence Rauchwerger for their help in the publication matters.

Kai Li
Program Chair

Organizing Committee

General Chairmen

Daniel Litaize, *IRIT, Université Paul Sabatier,*

Josep Torrellas, *University of Illinois at Urbana-Champaign*

Program Chair

Kai Li, *Princeton University*

Workshops & Publication Chair

Lawrence Rauchwerger, *Texas A&M University*

Local Arrangements Chair

Pascal Sainrat, *IRIT, CNRS*

Registration Chair

Franck Cappello, *LRI, Université d'Orsay*

Publicity Chair

Ahmed Louri, *University of Arizona*

Program Committee

Chair: Kai Li, *Princeton University*

David Culler, *University of California at Berkeley*

Jose Duato, *Universidad Politecnica de Valencia*

Michel Dubois, *University of Southern California*

Hans Eberle, *Sun Microsystems*

Guang R. Gao, *University of Delaware*

Garth Gibson, *Carnegie Mellon University*

Allan Gottlieb, *NYU and NEC Research*

Dirk Grunwald, *University of Colorado at Boulder*

Wen-mei Hwu, *University of Illinois at Urbana-Champaign*

Norman P. Jouppi, *Compaq WRL*

Henry Levy, *University of Washington*

Matt Blumrich *IBM T.J. Watson Research Center*

Daniel Litaize, *Institut de Recherche en Informatique de Toulouse*

Geoffrey Lowney, *Compaq*

Margaret Martonosi, *Princeton University*

Sally McKee, *University of Utah*

Trevor Mudge, *University of Michigan*

Toshiyuki Nakata, *NEC Corporation*

André Seznec, *INRIA/IRISA*

Michael D. Smith, *Harvard University*

Guri Sohi, *University of Wisconsin at Madison*

Per Stenström, *Chalmers University of Technology*

Wen-Hann Wang, *Intel MRL*

Pen-Chung Yew, *University of Minnesota*

Willy Zwaenepoel, *Rice University*

List of Referees

Michael Adler
Gagan Agrawal
Jose Nelson Amaral
Mohit Aron
Andrea C. Arpaci-Dusseau
Remzi H. Arpaci-Dusseau
David August
Todd Austin
Jean-Loup Baer
Vasanth Bala
Ron Barnes
Phil Bernstein
Laxmi Bhuyan
Ricardo Bianchini
Angelos Bilas
Matthias Blumrich
Francois Bodin
Laszlo Boeszoermenyi
Gilad Bracha
Philip Buonadonna
John Carter
Pedro Celis
Jeff Chase
Peter Chen
Lynn Choi
Fred Chong
Douglas Clark
William Cohen
Robert Cohn
Dan Connors
Christophe Coustet
Alan Cox
Michael Cox
Kevin Crozier
David Culler
David E. Culler
William J. Dally
Chita R. Das
Brian T. Davis
Nathalie Drach
Steve Dropsho
Peter Druschel
Jose Duato
Michel Dubois
Sandhya Dwarkadas
Hans Eberle
Masato Edahiro
Avinoam Nomik Eden
Jan Edler
Susan Eggers
Joel Emer
Roger Espasa
Alexandre Farcy
Keith I. Farkas
Mamoun Filali-Amine
Tryggve Fossum
Daniel Franco-Puntes
Manoj Franklin
Vincent W. Freeh
Chao-ying Fu
Greg Ganger
Guang R. Gao
Alain Gefflaut
Thomas Geiger
Kourosh Gharachorloo
Garth Gibson
Maya Gokhale
Allan Gottlieb
Jim Griffioen
Dirk Grunwald
Steven Guccione
John C. Gyllenhaal
Erik Hagersten
Thierry Haquin
David Harper
Gerd Heber
Ariel Hendel
Mark D. Hill
Wilson Hsieh
Y. Charlie Hu
Norman C. Hutchinson
Wen-mei Hwu
Dongming Jiang
D. Joseph
Norman P. Jouppi
Toni Juan
Prasad Kakulavarapu
Vijay Karamcheti
Stefanos Kaxiras
Ashfaq A. Khokhar
Hong-Seok Kim
John Kubiatowicz
John Lafferty
Konrad Lai
William Lam
Bernard Lecussan

Whay S. Lee
Yong-fong Lee
Henri Levy
Zhiyuan Li
Daniel Litaize
Pedro Lopez
Geoff Lowney
Shih-Lien Lu
Stefan Ludwig
Chi-Keung Luk
Steven S. Lumetta
Alan Mainwaring
Bill Mangione-Smith
Andres Marquez
Jose F. Martinez
Margaret Martonosi
Mark D. Matson
Hiroshi Matsuoka
Sally McKee
Nick McKeown
Matthew C. Merten
Maged Michael
Christopher J. Morrone
Andreas Moshovos
Todd C. Mowry
Trevor Mudge
Shubu Mukherjee
Jean-Michel Muller
David Nagle
Walid A. Najjar
Toshiyuki Nakata
AshwiniNanda
David Nassimi
Nyles Nettleton
Erik Nystrom
Subbarao Palacharla
Dhabaleswar K. Panda
Derek Pappas
Jih-Kwon Peir
Timothy M. Pinkston
Matt Postiff
Sanjay Rajopadhye
Ravi Rajwar
Parthasarathy Ranganathan
Alastair Reid
Steve Reinhardt
Glenn Reinman
Christine Rochange
Erven Rohou
Amir Roth
Sean Ryan
Pascal Sainrat
Yiannakis Sazeides
Karsten Schwan
Andre Seznec
John W. Sias
Jaswinder Pal Singh
Kevin Skadron
J. E. Smith
Michael D. Smith
Guri Sohi
Vijayalakshmi Srinivasan
Carl Staelin
Darko Stefanovic
J. Gregory Steffan
Per Stenstrom
Xinan Tang
Olivier Temam
Chandu Thekkath
Kevin B. Theobald
Josep Torrellas
Andrew Trick
Jenn-Yuan Tsai
Dean Tullsen
Gary Tyson
Geoff Voelker
Steven Wallace
Randolph Wang
Wen-Hann Wang
Chris Wilkerson
Frederick Wong
David A. Wood
Zhaoyun Xing
Pen-Chung Yew
Cliff Young
Craig Zilles
Willy Zwaenepoel

Table of Contents

Sixth International Symposium on High-Performance Computer Architecture HPCA-6

Session 1:
System Architecture Tradeoffs

Impact of Chip-Level Integration on Performance of OLTP Workloads

Luiz André Barroso, Kourosh Gharachorloo, Andreas Nowatzyk and Ben Verghese

Western Research Laboratory
Compaq Computer Corporation
{barroso,kourosh,verghese}@pa.dec.com, agn@acm.org

Abstract

With increasing chip densities, future microprocessor designs have the opportunity to integrate many of the traditional system-level modules onto the same chip as the processor. Some current designs already integrate extremely large on-chip caches, and there are aggressive next-generation designs that attempt to also integrate the memory controller, coherence hardware, and network router all onto a single chip. The tight coupling of these modules will enable efficient memory systems with substantially better latency and bandwidth characteristics relative to current designs. Among the important application areas for high-performance servers, online transaction processing (OLTP) workloads are likely to benefit most from these trends due to their large instruction and data footprints and high communication miss rates.

This paper examines the design trade-offs that arise as more system functionality is integrated onto the processor chip, and identifies a number of important architectural choices that are influenced by chip-level integration. In addition, the paper presents a detailed study of the performance impact of chip-level integration in the context of OLTP workloads. Our results are based on full system simulations of the Oracle commercial database engine running on both in-order and out-of-order issue processors used in uniprocessor and multiprocessor configurations. The results show that chip-level integration can improve the performance of both configurations by about 1.4 to 1.5 times, though for different reasons. For uniprocessors, integration of the L2 cache and the resulting lower hit latency is the primary factor in performance improvement. For multiprocessors, the improvement comes from both the integration of the L2 cache (lower L2 hit latency) and the integration of the other memory system components (better dirty remote latency). Furthermore, we find that the higher associativity afforded by integrating the L2 cache plays a critical role in counteracting the loss of capacity relative to larger off-chip caches. Finally, we find that the relative gains from chip-level integration are virtually identical for in-order and out-of-order processors.

1 Introduction

Advances in semiconductor technology enable next-generation microprocessor designs with well over a hundred million transistors on a single die. At the same time, the increasing density and speed of transistors make off-chip communication relatively more expensive. These technology trends provide an incentive to integrate functional modules that traditionally appear at the system-level onto the same die as the processor. The current HP PA-8500 microprocessor already integrates a 1MB data cache and a 0.5MB instruction cache with the processor. Similarly, the next-generation Alpha 21364 plans to aggressively exploit such integration by including a 1GHz 21264 core, two levels of caches, memory controller, coherence hardware, and network router all on a single die. Figure 1 shows a block diagram of the Alpha 21364 along with how the processors can be connected to form a scalable multiprocessor.

The primary benefits from integrating modules onto a single chip arise from more efficient communication interfaces. First, there are fewer signals that cross chip boundaries, leading to lower latency communication. Second, integration allows for substantially higher communication bandwidths by removing constraints imposed by scarce external pin resources. Finally, the tight coupling may in some cases enable novel interactions among modules that would not be feasible across chip boundaries. In the context of high-performance processors, these benefits typically translate into a lower latency and higher bandwidth cache hierarchy and memory system.

Chip-level integration provides other advantages besides higher performance. A key benefit is lower system component counts which leads to more reliable and lower cost systems that are easier to manufacture and maintain. Compute power density (e.g., computes per cubic foot) is also clearly improved through integration. Another possible benefit of integration is reduced design complexity that arises from eliminating generalized transactional interfaces that are typically used when modules communicate across chip boundaries.[1] Finally, designs that provide on-chip cache-coherence support, such as the Alpha 21364, lead to faster turnaround times for incorporating the latest microprocessors in multiprocessor servers.

Given that commercial applications such as databases and Web servers constitute the largest and fastest-growing segment of the market for high-performance servers, it is important to evaluate the impact of chip-level integration in the context of

[1] However, eliminating such interfaces must be done with care to keep verification costs in check; tightly integrated units that lack well structured and observable interfaces must be tested as a single entity and are harder to verify.

0-7695-0550-3/00 $10.00

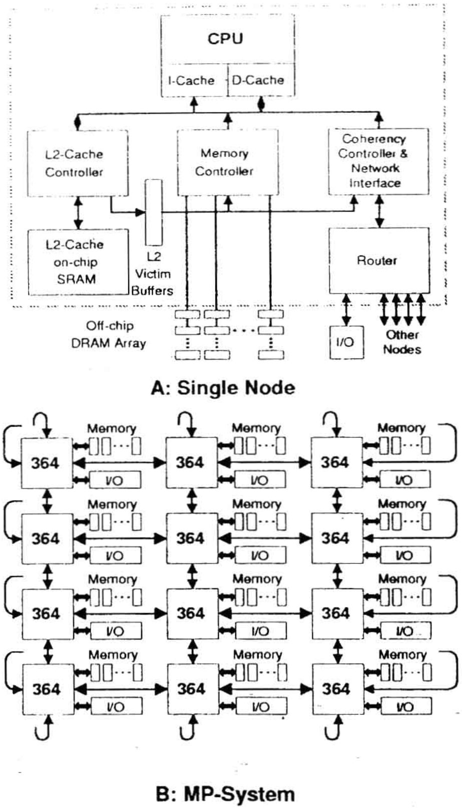

Figure 1: Block diagram of the Alpha 21364.

these workloads. While applications such as decision support (DSS) and Web index search have been shown to be relatively insensitive to memory system performance [1], a number of recent studies have underscored the radically different behavior of online transaction processing (OLTP) workloads [1, 2, 3, 8, 12, 15, 18]. In general, OLTP workloads lead to inefficient executions with a large memory stall component and present a more challenging set of requirements for processor and memory system design. This behavior arises from large instruction and data footprints and high communication miss rates that are characteristic for such workloads [1]. At the same time, the increasing popularity of electronic commerce on the Web further elevates the importance of achieving good performance on OLTP workloads.

This paper explores the general design trade-offs that arise as more system-level modules are incorporated onto the processor chip, and presents a detailed study of the impact of chip-level integration on the performance of OLTP workloads. We consider successively integrating (i) the second-level cache, (ii) memory controllers, and (iii) coherence hardware and network router. In each case, we identify important architectural trade-offs that are influenced by the integration and analyze their performance impact for both uniprocessor and multiprocessor configurations, and in-order and out-of-order processor models. Since previous results on OLTP show the need for extremely large off-chip caches [1, 8], we also consider the viability of larger off-chip caches in a fully integrated design.

Our performance results for OLTP are based on full system simulations, including operating system activity, of the Oracle commercial database engine (version 7.3.2) running under Compaq Tru64 Unix (previously known as Digital Unix). Our simulation parameters are based on an aggressive 0.18um CMOS technology that will be used in next-generation microprocessors such as the Alpha 21364.

Our results show that chip-level integration of the memory system can yield about a 1.4 times improvement in OLTP performance over an aggressive non-integrated design, for both uniprocessor and multiprocessor configurations. Relative to less aggressive off-chip multiprocessor designs, these gains can be as high as 1.5 to 1.6 times. We find that uniprocessor OLTP performance is primarily determined by the size, organization and latency of the L2 cache and is relatively insensitive to the rest of the memory system. Therefore, the primary performance benefit comes from integrating the L2 cache with the processor. Multiprocessor OLTP performance is also dependent on the latency of dirty remote (3-hop) misses. Therefore, the performance benefits in multiprocessor configurations come from both integrating the L2 (lower L2 hit latency) and from integrating the other system components (lower dirty miss latency). We also observe that there is synergy in tightly coupling the coherence and memory controllers, and separating these two system components can lead to sub-optimal designs. Finally, our analysis shows that there is little justification for augmenting a fully integrated design with larger off-chip caches.

One of the most interesting results of this paper is with respect to the effectiveness of smaller on-chip L2 caches in capturing the footprint of OLTP workloads. Previous studies of OLTP have pointed to the need for large off-chip caches. Surprisingly, we find that the limited capacity of an on-chip L2 cache is more than offset by the higher associativity that is made viable by the integration. In fact, a 2MB on-chip cache with 4-way or 8-way associativity exhibits fewer misses than an 8MB direct-mapped off-chip cache. This extreme sensitivity of OLTP workloads to associativity shows that many of the misses eliminated by large direct-mapped off-chip caches are actually conflict misses as opposed to true capacity misses.

The rest of paper is structured as follows. The next section presents our experimental methodology, including a brief description of the OLTP workload and the simulated architectures. Sections 3, 4, and 5 evaluate the trade-offs of successively integrating the second-level cache, memory controller, and coherence controller and network router onto the processor chip. We evaluate the viability of large off-chip third-level caches for integrated designs in Section 6. Section 7 shows that the relative benefits of on-chip integration are virtually identical for in-order and out-of-order processors. Finally, we discuss related work and conclude.