

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

INTEL CORPORATION,

Petitioner

v.

FG SRC LLC,

Patent Owner

CASE NO.: 2020-01449

PATENT NO. 7,149,867

DECLARATION OF EILEEN D. MCCARRIER

Mail Stop **PATENT BOARD**
Patent Trial and Appeal Board
U.S. Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

I, Eileen D. McCarrier, declare as follows:

1. I am currently Manager of Research Services at Pillsbury Winthrop Shaw Pittman LLP, where I have worked as a reference librarian and manager of research services for 26 years.

2. I make this declaration based on my own personal knowledge, including my knowledge of library science practices and the evidence cited herein.

3. I earned a Master of Arts degree in Library Science from the University of Wisconsin-Madison in 1980, and I have worked as a law librarian for 40 years.

4. I have prepared this declaration in connection with the above-captioned *inter partes* review (“IPR”) proceeding for which I obtained library copies of previously filed Exhibits 1003 and 1004. Each of these articles was published by the Institute of Electrical and Electronics Engineers, Inc. (“IEEE”), a widely recognized publisher of technical papers spanning a broad range of technologies including electronics, electrical engineering, telecommunications, computing, and more. The IEEE publishes thousands of conference papers every year, including by making them publicly available via its Xplore digital library. The IEEE’s collection of publications is recognized within the reference library field as an authoritative source of consolidated published papers in electrical engineering, computer science, and related fields.

5. The primary holders of the original printed versions of IEEE conference proceedings from the 1996–2000 era are generally university libraries across the United States. However, at the time of the drafting of this declaration, safety protocols relating to COVID-19 continue to restrict the ability of members of the public to physically access many of these institutions. To accommodate these safety procedures, the library copies described below were physically located by library staff, digitally scanned, and then emailed directly to me.

Library copies of Zhang, Exhibit 1003

6. X. Zhang et al., Architectural Adaptation for Application-Specific Locality Optimizations, IEEE (1997), published in the Proceedings of the International Conference on Computer Design - VLSI in Computers and Processors (IEEE, October 12–15, 1997), 150–156 (“Zhang”) was filed as Exhibit 1003 with the petition in the above-captioned IPR proceeding.

7. A true and correct copy of Zhang obtained from the Library of the Missouri University of Science and Technology, Rolla, Missouri, part of the University of Missouri system, is attached as Appendix EDM01. I personally received this copy as a pdf from the interlibrary loan department of C. L. Wilson Library.

Library copies of Gupta, Exhibit 1004

8. R. Gupta, Architectural Adaptation in AMRM Machines, Proceedings of the IEEE Computer Society Workshop on VLSI 2000 (IEEE, April 27–28, 2000), 75–79 (“Gupta”) was filed as Exhibit 1004 with the above-captioned IPR petition.

9. A true and correct copy of Gupta obtained from the Georgia Tech Library, Georgia Institute of Technology is attached as Appendix EDM02. I personally received this copy as a pdf from the library’s Interlibrary Loan Office.

I declare under penalty of perjury that the foregoing is true and correct.

Date: March 31, 2021



Eileen D. McCarrier

Appendix EDM01

UNIVERSITY OF MISSOURI-ROLLA



050-102823580

LIBRARY
UNIVERSITY OF MISSOURI-ROLLA
1870 MINER CIRCLE
ROLLA, MISSOURI 65409-0060

Intel Exhibit 1028 - 5

DEMCO

Proceedings

**International Conference on
Computer Design**

VLSI in Computers and Processors

Intel Exhibit 1028 - 6

75.00

TK
7874
.I43
1997

Proceedings

**International Conference on
Computer Design**

VLSI in Computers and Processors

October 12-15, 1997

Austin, Texas

Sponsored by

IEEE Computer Society Technical Committee on Design Automation

IEEE Circuits and Systems Society



Los Alamitos, California

Washington • Brussels • Tokyo

Intel Exhibit 1028 - 7

Copyright © 1997 by The Institute of Electrical and Electronics Engineers, Inc.
All rights reserved

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries may photocopy beyond the limits of US copyright law, for private use of patrons, those articles in this volume that carry a code at the bottom of the first page, provided that the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

Other copying, reprint, or republication requests should be addressed to: IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, P.O. Box 133, Piscataway, NJ 08855-1331.

The papers in this book comprise the proceedings of the meeting mentioned on the cover and title page. They reflect the authors' opinions and, in the interests of timely dissemination, are published as presented and without change. Their inclusion in this publication does not necessarily constitute endorsement by the editors, the IEEE Computer Society, or the Institute of Electrical and Electronics Engineers, Inc.

IEEE Computer Society Order Number PR08026
ISBN 0-8186-8026-X
ISBN 0-8186-8207-8 (case)
ISBN 0-8186-8208-6 (microfiche)
IEEE Order Plan Catalog Number 97CB36149
ISSN 1063-6404

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: + 1-908-981-1393
Fax: + 1-908-981-9667
mis.custserv@computer.org

IEEE Computer Society
13, Avenue de l'Aquilon
B-1200 Brussels
BELGIUM
Tel: + 32-2-770-2198
Fax: + 32-2-770-8505
euro.ofc@computer.org

IEEE Computer Society
Ooshima Building
2-19-1 Minami-Aoyama
Minato-ku, Tokyo 107
JAPAN
Tel: + 81-3-3408-3118
Fax: + 81-3-3408-3553
tokyo.ofc@computer.org

Editorial production by Ian Torwick

Cover design by Joseph Daigle/Studio Productions

Printed in the United States of America by Technical Communication Services

11-25-97
ah

Table of Contents

ICCD '97 Conference Program

1997 Technical Program

Welcome to ICCD'97..... xiv
Program Committee xvi
Additional Referees xix

Session 1.1: Keynote Speech

Intelligent RAM (IRAM): the Industrial Setting, Applications, and Architectures 2
*D. Patterson, K. Asanovic, A. Brown, R. Fromm, J. Golbus, B. Gribstad,
 K. Keeton, C. Kozyrakis, D. Martin, S. Perissakis, R. Thomas, N. Treuhft,
 and K Yelick, University of California at Berkeley, California*

Session 1.2: CAD Plenary

Chair: Andreas Kuehlmann, IBM T. J. Watson Research Center

A Brief History of the Future of Semiconductor Electronic Design Automation 10
Ron Rohrer, TBD Consultants

Concurrent Sessions 1.3

Session 1.3.1: Special Session: Industrial Applications of Formal Verification

Organizer and Chair: Andreas Kuehlmann, IBM T. J. Watson Research Center

**Formal Implementation Verification of the Bus Interface Unit for the Alpha 21264
 Microprocessor** 16
G.P. Bischoff, K.S. Brace, S. Jain, and R. Razdan

Intertwined Development and Formal Verification of a 60x bus Model 25
M. Kaufmann and C. Pixley

**Formally Specifying and Mechanically Verifying Programs for the Motorola Complex
 Arithmetic Processor DSP** 31
B.C. Brock and W.A. Hunt, Jr.

BIST-Based Fault Diagnosis in the Presence of Embedded Memories 37
J. Savir

Built-in Self Test for Content Addressable Memories 48
Y.-S. Kang, J.-C. Lee, and S. Kang

Pseudo-Random Pattern Testing of Bridging Faults 54
N.A. Toubia and E.J. McCluskey

Session 1.3.3: Simulation and Power Estimation

Chair: Teng-Sheng Moh, Silicon Valley Research, Inc.

Novel Simulation of Deep-Submicron MOSFET Circuits 62
S. Bruma and R.H.J.M. Otten

**Time-Stamped Transition Density for the Estimation of Delay Dependent
 Switching Activities** 68
H. Choi and S.H. Hwang

Power Compiler: A Gate-Level Power Optimization and Synthesis System	74
--	----

B. Chen and I. Nedelchev

Session 1.3.4: Branch Prediction

Chair: Jim Bondi, Texas Instruments

Elastic History Buffer: A Low-Cost Method to Improve Branch Prediction Accuracy	82
---	----

M.-D. Tarlescu, K.B. Theobald, and G.R. Gao

Design Optimization for High-Speed Per-address Two-level Branch Predictors	88
--	----

I.-C.K. Chen, C.-C. Lee, M.A. Postiff, and T.N. Mudge

PA-8000: A Case Study of Static and Dynamic Branch Prediction	97
---	----

C. Burch

Concurrent Sessions 1.4

Session 1.4.1: New Techniques for Gate-Sizing and Retiming

Chair: Derek Beatty, Motorola, Inc.

Discrete Drive Selection for Continuous Sizing	110
--	-----

R. Haddad, L.P.P.P. van Ginneken, and N. Shenoy

Continuous Retiming: Algorithms and Applications	116
--	-----

P. Pan

Optimal Clock Period Clustering for Sequential Circuits with Retiming	122
---	-----

A.K. Karandikar, P. Pan, and C.L. Liu

Session 1.4.2: Circuit Modeling

Chair: Sandip Kundu, IBM Corp.

Comparison between nMOS Pass-Transistor logic style vs. CMOS Complementary Cells	130
--	-----

R. Mehrotra, M. Pedram, and X. Wu

Circuit-Based Description and Modeling of Electromagnetic Noise Effects in Packaged Low-Power Electronics	136
---	-----

A.C. Cangellaris, W. Pinello, and A. Ruehli

Transistor-Level Sizing and Timing Verification of Domino Circuits in the Power PC Microprocessor	143
---	-----

A. Dharchoudhury, D. Blaauw, J. Norton, S. Pullela, and J. Dunning

Session 1.4.3: Novel Architectures

Chair: Greg Fisher, Printronix Corporation

Architectural Adaptation for Application-Specific Locality Optimization	150
---	-----

X. Zhang, A. Dasdan, M. Schulz, R.K. Gupta, and A.A. Chien

A New Processor Architecture for Digital Signal Transport Systems	157
---	-----

M. Inamori, K. Ishii, A. Tsutsui, K. Shirakawa H. Nakada, and T. Miyazaki

Short Papers

PROPHID: A Heterogeneous Multi-Processor Architecture for Multimedia	164
--	-----

J.A. Leijten, J.L. van Meerbergen, A.A. Timmer, and J.A.G. Jess

Enhanced Compression Techniques to Simplify Program Decompression and Execution	170
<i>M. Breternitz, Jr. and R. Smith</i>	
Session 1.4.4: Low Power Architectures	
<i>Chair: Tim Brodnax, IBM</i>	
A Low Power Approach to Floating Point Adder Design	178
<i>R.V.K. Pillai, D. Al-Khalili, and A.J. Al-Khalili</i>	
Design and Implementation of Low-Power Digit-Serial Multipliers	186
<i>Y.-N. Chang, J.H. Satyanarayana, and K.K. Parhi</i>	
On Complexity Reduction of FIR Digital Filters Using Constrained Least Squares Solution.....	196
<i>K. Muhammad and K. Roy</i>	
Concurrent Session 1.5	
Session 1.5.1: Timing Optimization for Deep Submicron Technology	
<i>Chair: Masahiro Fujita, Fujitsu Laboratories of America</i>	
An Integrated Placement and Synthesis Approach for Timing Closure of PowerPC™ Microprocessors	206
<i>S. Hojat and P. Villarrubia</i>	
Post-Layout Circuit Speed-up by Event Elimination.....	211
<i>H. Vaishnav, C.-K. Lee, and M. Pedram</i>	
Clustering and Load Balancing for Buffered Clock Tree Synthesis.....	217
<i>A.D. Mehta, Y.-P. Chen, N. Menezes, D.F. Wong, and L.T. Pileggi</i>	
CMOS Gate Delay Models for General RLC Loading	224
<i>R. Arunachalam, F. Dartu, and L.T. Pileggi</i>	
Session 1.5.2: Special Session: The G4 S/390 Microprocessor	
<i>Organizer: Andreas Kuehlmann, IBM T. J. Watson Research Center</i>	
<i>Chair: Sumit Dasgupta, IBM</i>	
Design Methodology for the High-Performance G4 S/390 Microprocessor	232
<i>K.L. Shepard, S. Carey, D.K. Beece, R. Hatch, and G. Northrop</i>	
A High-Frequency Custom CMOS S/390 Microprocessor	241
<i>C.F. Webb and J.S. Liptay</i>	
High Performance CMOS Circuit Techniques for the G-4 S/390 Microprocessor	247
<i>J. Warnock, L. Sigal, B. Curran, and Y. Chan</i>	
A 400 MHz, 144Kb CMOS ROM MACRO for an IBM S/390-Class Microprocessor.....	253
<i>A. Tuminaro</i>	
Session 1.5.3: Multiprocessor Communication	
<i>Chair: Wai-Chi Fang, Jet Propulsion Laboratory</i>	
A Comparative Evaluation of Hierarchical Network Architecture of the HP-Convex Exemplar	258
<i>R. Castaneda, X. Zhang, and J.M. Hoover, Jr.</i>	

Effect of Message Length and Processor Speed on the Performance of the Bidirectional Ring based Multiprocessor	267
<i>H. Oi and N. Ranganathan</i>	
An Approach to Network Caching for Multimedia Objects	273
<i>M.A. Kozuch, W. Wolf, and A. Wolfe</i>	
Development of a High Bandwidth Merged Logic/DRAM Multimedia Chip.....	279
<i>W.K. Luk, Y. Katayama, W. Hwang, M. Wordeman, T. Kirihata, A. Satoh, S. Munetoh, H. Wong, B. El-Kareh, P. Xio, and R. Joshi</i>	

Session 1.5.4: Asynchronous Architectures

Chair: Eric Chou, Hewlett-Packard ULSI Labs

TITAC-2: An asynchronous 32bit microprocessor based on Scalable -Delay-Insensitive model.....	288
<i>A. Takamura, M. Kuwako, M. Ima, T. Fujii, M. Ozawa, I. Fukasaku, Y. Ueno, and T. Nanya</i>	
An Evaluation of Asynchronous and Synchronous Design for Superscalar Architectures.....	295
<i>A. Davey and D. Loyd</i>	
Synthesis of High Speed Delay-Insensitive Combinational Iterative Tree Circuits.....	301
<i>F.-C. Cheng</i>	
Asynchronous Wrapper for Heterogeneous Systems	307
<i>D.S. Bormann and P.Y.K. Cheung</i>	

Concurrent Session 1.6

Session 1.6.1

Panel: The War of the Roses: Designers versus Tool Developers	317
---	-----

Organizer: Andreas Kuehlmann, IBM T. J. Watson Research Center

Moderator: Daniel Beece, IBM T. J. Watson Research Center

Panelists:

- Barbara Chappell, Intel Corp.*
- Robert Damiano, Synopsys, Inc.*
- Charlie Malley, Hewlett-Packard*
- Yiannos Manoli, University of Saarland*
- David F. Reed, Advanced Micro Devices*
- Steven E. Schulz, Texas Instruments*

Session 1.6.2

Panel: If Software is King for Systems-on-Silicon, What's New in Compilers?	322
---	-----

Organizer: Nikil Dutt, University of California at Irvine

Moderator: Sharad Malik, Princeton University

Panelists:

- Lex Augusteijn, Philips Research Laboratory*
- Beatrice Fu, Intel Corporation*
- Alex Nicolau, University of California at Irvine*
- Constantine Polychronopoulos, University of Illinois at Urbana-Champaign*

Session 2.1 Design and Test Plenary

Chair: Magdy Abadir, Motorola, Inc.

Design and Test: The Lost World	328
<i>W. Joyner, IBM T. J. Watson Research Center</i>	

Concurrent Sessions 2.2

Session 2.2.1: Binary Decision Diagrams

Chair: B. Bischoff, Digital Equipment Corporation

Equivalence Checking Using Abstract BDDs	332
<i>S. Jha, Y. Lu, M. Minea, and E.M. Clarke</i>	
Speeding up Variable Reordering of OBDDs	338
<i>C. Meinel and A. Slobodova</i>	
Dynamic Reordering in a Breadth-First Manipulation Based BDD package: Challenges and Solutions	344
<i>R.K. Ranjan, W. Gosti, R.K. Brayton, and A. Sangiovanni-Vincentelli</i>	
Timed Binary Decision Diagrams	352
<i>Z. Li, Y. Zhao, Y. Min, and R.K. Brayton</i>	

Session 2.2.2: Advanced Test Topics

Chair: Magdy Abadir, Motorola

Vector Restoration Based Static Compaction of Test Sequences for Synchronous Sequential Circuits	360
<i>I. Pomeranz and S.M. Reddy</i>	
Nonenumerative Path Delay Fault Coverage Estimation with Optimal Algorithms	366
<i>D. Kagaris, S. Tragoudas, and D. Karayiannis</i>	
Properties of the Input Pattern Fault Model	372
<i>R.D. (Shawn) Blaton and J.P. Hayes</i>	
A new approach for Initialization Sequences Computation for Synchronous Sequential Circuits	381
<i>F. Corno, P. Prinetto, M. Rebaudengo, M.S. Reorda, and G. Squillero</i>	

Session 2.2.3: Embedded Software and Systems

Chair: Rolf Ernst, Technische Universitaet Braunschweig, Germany

Real Time Operating Systems for Embedded Computing	388
<i>Y. Li, M. Potkonjak and W. Wolf</i>	
Allocation and Data Arrival Design of Hard Real Time Systems	393
<i>D.L. Rhodes and W. Wolf</i>	
Improving Design Turnaround Time Via Two-Levels Hw/Sw Co-Simulation	400
<i>A. Allara, S. Filipponi, W. Fornaciari, F. Salice, and D. Sciuto</i>	

Session 2.2.4: Low Power Issues

Chair: Jose L. Cruz-Rivera, University of Puerto Rico-Mayaguez

Power Constrained Design of Multiprocessor Interconnection Networks	408
<i>C.S. Patel, S.M. Chai, S. Yalamanchili, and D.E. Schimmel</i>	

Memory Traffic and Data Cache Behavior of an MPEG-2 Software Decoder	417
<i>P. Soderquist and M. Leeser</i>	
Asynchronous Transpose-Matrix Architectures	423
<i>J.A. Tierno and P. Kudva</i>	
A Low Power Smart Vision System Based on Active Pixel Sensor Integrated with Programmable Neural Processor	429
<i>W.-C. Fang, G. Yang, B. Pain, and B.J. Sheu</i>	

Concurrent Sessions 2.3

Session 2.3.1: Formal Verification Methods

Chair: Warren Hunt, Jr., Computational Logic, Inc.

Formal Verification of the HAL S1 System Cache Coherence Protocol	438
<i>A.J. Hu, M. Fujita, and C. Wilson</i>	
A Survey of Techniques for Formal Verification of Combinational Circuits	445
<i>J. Jain, A. Narayan, M. Fujita, and A. Sangiovanni-Vincentelli</i>	
Checking Formal Specifications under Simulation	455
<i>W. Canfield, E.A. Emerson, and A. Saha</i>	

Session 2.3.2: Mixed Signal Design and Test

Chair: Yervant Zorian, Logic Vision

Built-In Temperature Sensors for On-line Thermal Monitoring of Microelectronic Structures	462
<i>K. Arabi and B. Kaminska</i>	
Develop of Hierarchical Testability Design Methodologies for Analog/ Mixed-Signal Integrated Circuits	468
<i>C.-P. Wang and C.-L. Wey</i>	
A Novel Test Set Design for Parametric Testing of Analog and Mixed-Signal Circuits	474
<i>J. Chen and A. Ramachandran</i>	

Session 2.3.3: FPGA Design

Chair: Paul Franzon, North Carolina State University

On the Construction of Universal Series-Parallel Functions for Logic Module Design.....	482
<i>F.Y. Young and D.F. Wong</i>	
An Universal Pezaris Array Multiplier Generator for SRAM-Based FPGAs	489
<i>J. Stohmann and E. Barke</i>	
Channel Segmentation Design for Symmetrical FPGAs	496
<i>W.-K. Mak and D.F. Wong</i>	

Session 2.3.4: Cache Technology I

Chair: Mauricio Breternitz, Motorola Inc.

Multi-Column Implementations for Cache Associativity	504
<i>C. Zhang, X. Zhang, and Y. Yan</i>	
Design and Performance Evaluation of a Cache Assist to Implement Selective Caching.....	510
<i>L.K. John and A. Subramanian</i>	

On Effective Data Supply for Multi-Issue Processor.....	519
<i>J.A. Rivers, E.S. Tam, and E.S. Davidson</i>	

Concurrent Sessions 2.4

Session 2.4.1: Embedded Tutorial

Practical Issue of Interconnect Analysis in Deep Submicron Integrated Circuits.....	532
<i>Chair: Andreas Kuehlmann, IBM T. J. Watson Research Center</i>	
<i>Presenter: Kenneth L. Shepard, IBM T. J. Watson Research Center</i>	

Session 2.4.2: Fault Diagnosis

<i>Chair: Raj Raina, Motorola, Inc.</i>	
First Results of System Level Fault Tolerant Design Validation Through Laser Fault Injection	544
<i>W. A. Moreno , F.J. Falquez, J.R. Samon Jr., and T. Smith</i>	
Integrated Diagnostics for Embedded Memory Built-in Self Test on PowerPC™ Devices	549
<i>C. Hunter</i>	
A TSC Evaluation Function for Combinatorial Circuits.....	555
<i>C. Bolchini, D. Sciuto, and F. Salice</i>	

Session 2.4.3: Special Session: Low Power Design Issues

<i>Chair: Sarma Vrudhula, University of Arizona</i>	
An Architectural Power Optimization Case Study Using High-level Synthesis	562
<i>C.-T. Chen and K. Kucukcakar</i>	
High-Level Design Synthesis of Low Power, VLIW Processor for the IS-54 VSELP Speech Encoder	571
<i>R. Henning and C. Chakrabarti</i>	

Session 2.4.4: Cache Technology II

<i>Chair: Michael Stumm, University of Toronto</i>	
Fast Cache Access with Full-Map Block Directory	578
<i>J.-K. Peir, W.W. Hsu, H. Young, and S. Ong</i>	
A Data Alignment Technique for Improving Cache Performance.....	587
<i>P.R. Panda, H. Nakamura, N.D. Dutt, and A. Nicolau</i>	
Instruction Prefetching Using Branch Prediction Information.....	593
<i>I.-C.K. Chen, C.-C. Lee, and T.N. Mudge</i>	

Session 3.1: Architecture & Algorithm Plenary

<i>Chair: J. Robert Jump, Rice University</i>	
Is Wireless Data Dead?	604
<i>Randy Katz, University of California at Berkeley, California</i>	

Concurrent Sessions 3.2

Session 3.2.1: Layout Partitioning and Synthesis

Chair: Georg Pelz, Gerhard-Mercator-University-GH, Duisburg, Germany

An Efficient Multi-Way Algorithm for Balanced Partitioning of VLSI Circuits	608
<i>X. Tan, J. Tong, P. Tan, N. Park, and F. Lombardi</i>	
Partitioning Under Timing and Area Constraints	614
<i>G. Tumbush and D. Bhatia</i>	
A Parallel Circuit-Partitioned Algorithm for Timing Driven Standard Cell Placement	621
<i>J.A. Chandy, and P. Banerjee</i>	
Crosstalk-Constrained Maze Routing Based on Lagrangian Relaxation	628
<i>H. Zhou and D.F. Wong</i>	

Session 3.2.2: Design for Testability & Test Synthesis

Chair: Sujit Dey, NEC

High Level Test Synthesis Across the Boundary of Behavioral and Structural Domains	636
<i>K. Lai, C.A. Papachristou, and M. Baklashov</i>	
Power Driven Partial Scan.....	642
<i>J.-Y. Jou and M.-C. Nien</i>	
Synthesis of Delay Verifiable Sequential Circuits using Partial Enhanced Scan.....	648
<i>R.C. Tekumalla and P.R. Menon</i>	
Application of a Testing Framework to VHDL Descriptions at Different Abstraction Levels	654
<i>M. Bacis, G. Buonanno, F. Ferrandi, F. Fummi, L. Gerli, and D. Sciuto</i>	

Session 3.2.3: Embedded Tutorial

Practical Advances in Asynchronous Design	662
---	-----

Chair: Rob Roy, NEC, Inc.

Presentors:

Eric Brunvand of University of Utah

Steve Nowick of Columbia University

Kenneth Yun of University of California at San Diego

Session 3.2.4: Arithmetics

Chair: Joe Cavallaro, Rice University

Benchmarking and Analysis of Architectures for CAD Applications.....	670
<i>A. Mehrotra, S. Qadeer, R.K. Ranjan, and R.H. Katz</i>	
Fast Low-Energy VLSI Binary Addition.....	676
<i>K.K. Parhi</i>	
A Floating-Point Divider using Redundant Binary Circuits and an Asynchronous Clock Scheme	685
<i>H. Suzuki, H. Mankino, K. Mashiko, and H. Hamano</i>	
Parallel-Array Implementations of A Non-Restoring Square Root Algorithm.....	690
<i>Y. Li and W. Chu</i>	

Concurrent Session 3.3

Session 3.3.1: Asynchronous Design

Chair: Daniel Saab, Case Reserve Western University

Optimizing CMOS Implementations of C-element..... 700
M. Shams, J.C. Ebergen, and M.I. Elmasry

A Doubly-Latched Asynchronous Pipeline..... 706
R. Kol and R. Ginosar

A Pulse-To-Static Conversion Latch With a Self-Timed Control Circuit..... 712
W. Hwang, R.V. Joshi, and W.H. Henkels

Session 3.3.2: Special Session: Interconnect Modeling & Repeater Methodologies

Chair: Byron Krauter, IBM Corp.

Fast Generation of Statistically-based Worst-Case Modeling of On-Chip Interconnect..... 720
N. Chang, V. Kanevsky, O.S. Nakagawa, K. Rahmat, and S.-Y. Oh

A Repeater Optimization Methodology for Deep Sub-Micron, High-Performance Processor 726
D. Li, A. Pua, P. Srivastava, and U. Ko

Critical Voltage Transition Logic: An Ultrafast CMOS Logic Family 732
Z. Zhu and B.S. Carlson

Session 3.3.3: Finite-State Machine and High-Level Synthesis

Chair: Ramin Hojati, University of California at Berkeley

Divide and Conquer: A Strategy for Synthesis of Low Power Finite State Machines 740
A. Dasgupta and S. Ganguly

Estimation of Maximum Power for Sequential Circuits Considering Spurious Transitions 746
C.-Y. Wang and K. Roy

Dynamic Bounding of Successor Force Computations in the Force Directed List Scheduling Algorithm 752
S. Govindarajan and R. Vemuri

Author Index 758

Architectural Adaptation for Application-Specific Locality Optimizations

Xingbin Zhang* Ali Dasdan* Martin Schulz† Rajesh K. Gupta‡ Andrew A. Chien*

*Department of Computer Science
University of Illinois at Urbana-Champaign
{zhang,dasdan,achien}@cs.uiuc.edu

†Institut für Informatik
Technische Universität München
schulzm@informatik.tu-muenchen.de

‡Information and Computer Science, University of California at Irvine
rgupta@ics.uci.edu

Abstract

We propose a machine architecture that integrates programmable logic into key components of the system with the goal of customizing architectural mechanisms and policies to match an application. This approach presents an improvement over traditional approach of exploiting programmable logic as a separate co-processor by preserving machine usability through software and over traditional computer architecture by providing application-specific hardware assists. We present two case studies of architectural customization to enhance latency tolerance and efficiently utilize network bisection on multiprocessors for sparse matrix computations. We demonstrate that application-specific hardware assists and policies can provide substantial improvements in performance on a per application basis. Based on these preliminary results, we propose that an application-driven machine customization provides a promising approach to achieve high performance and combat performance fragility.

1 Introduction

Technology projections for the coming decade [1] point out that system performance is going to be increasingly dominated by intra-chip interconnect delay. This presents a unique opportunity for programmable logic as the interconnect dominance reduces the contribution of per stage logic complexity on performance and the marginal costs of adding switching logic in the interconnect. However, the traditional *co-processing* architecture of exploiting programmable logic as a specialized functional unit to deliver a specific application suffers from the problem of machine retargetability. A system generated using this approach typically can not be retargeted to another application

without repartitioning hardware and software functionality and reimplementing the co-processing hardware. This re-targetability problem is an obstacle toward exploiting programmable logic for general purpose computing.

We propose a machine architecture that integrates programmable logic into key components of the system with the goal of customizing architectural mechanisms and policies to match an application. We base our design on the premise that communication is already critical and getting increasingly so [17], and flexible interconnects can be used to replace static wires at competitive performance [6, 9, 20]. Our approach presents an improvement over co-processing by preserving machine usability through software and over traditional computer architecture by providing application-specific hardware assists. The goal of application-specific hardware assists is to overcome the rigid architectural choices in modern computer systems that do not work well across different applications and often cause substantial performance fragility. Because performance fragility is especially apparent on memory performance on systems with deep memory hierarchies, we present two case studies of architectural customization to enhance latency tolerance and efficiently utilize network bisection on multiprocessors. Using sparse matrix computations as examples, our results show that customization for application-specific optimizations can bring significant performance improvement (10X reduction in miss rates, 100X reduction in data traffic), and that an application-driven machine customization provides a promising approach to achieve robust, high performance.

The rest of the paper is organized as follows. Section 2 presents our analyses of the technology trends. Section 3 describes our proposed architecture and the project context. We describe our case studies in Section 4 and discuss related work in Section 5. Finally, we conclude with future directions in Section 6.

2 Background

Technology projections for the coming decade point out a unique opportunity of programmable logic. However, the traditional co-processing approach of exploiting programmable logic suffers from the problem of machine re-targetability, which limits its use for general purpose applications.

2.1 Key Technology Trends

The basis for architectural adaptation is in the key trends in the semiconductor technology. In particular, the difference in scaling of switching logic speed and interconnect delays points out increasing opportunities for programmable logic circuits in the coming decade. Projections by the Semiconductor Industry Association (SIA) [1] show that on-chip system performance is going to be increasingly dominated by interconnect delays. Due to these interconnect delays, the on-chip clock periods will be limited to about 1 nanosecond, which is well above the projections based on channel length scaling [1]. Meanwhile, the unit gate delay (inverter with fanout of two) scales down to 20 pico-seconds. Thus, modern day control logic consisting of 7-8 logic stages per cycle would form less than 20% of the total cycle time. This clearly challenges the fundamental design trade-off today that tries to simplify the amount of logic per stage in the interest of reducing the cycle time [14]. In addition, this points to a sharply reduced marginal cost of per stage logic complexity on the circuit-level performance.

The decreasing delay penalty for (re)programmable logic blocks compared to interconnect delays also makes the incorporation of small programmable logic blocks attractive even in custom data paths. Because the interconnect delays scale down much more slowly than transistor switching delays, in the year 2007, the delay of the average length inter-connect (assuming an average interconnect length of 1000X the pitch) would correspond to approximately three gate delays (see [5] for a detailed analysis). This is in contrast to less than half the gate delay of the average interconnect in current process technology. This implies that due to purely electrical reasons, it would be preferred to include at least one inter-connect buffer in a cycle time. This buffer gate when combined with a weak-feedback device would form the core of a storage element that presents less than 50% switching delay overhead (from 20ps to 30ps), making it performance competitive to replace static wires with flexible interconnect.

In view of these technology trends and advances in circuit modeling using hardware description languages (HDLs) such as Verilog and VHDL, the process of hardware design is increasingly a language-level activity, supported by compilation and synthesis tools [11, 12]. With

these CAD and synthesis capabilities, programmable logic circuit blocks are beginning to be used in improving system performance.

2.2 Co-processing

The most common architecture in embedded computing systems to exploit programmable logic can be characterized as one of *co-processing*, i.e., a processor working in conjunction with dedicated hardware assists to deliver a specific application. The hardware assists are built using programmable circuit blocks for easy interpretation with the predesigned CPU. Figure 1 shows the schematic of a co-processing architecture, where the co-processing hardware may be operated under direct control of the processor which stalls while the dedicated hardware is operational [10], or the co-processing may be done concurrently with software [13]. However, a system generated using this approach typically can not be re-targeted to another application without repartitioning hardware and software functionality and reimplementing the co-processing hardware even if the macro-level organization of the system components remains unaffected. This presents an obstacle of exploiting programmable logic for general-purpose computing even though technology trends make it possible to do so.




Figure 1. A co-processing Architecture

3 Architectural Adaptation

We propose an architecture that integrates small blocks of programmable logic into key elements of a baseline architecture, including processing elements, components of the memory hierarchy, and the scalable interconnect, to provide *architectural adaptation* - the customization of architectural mechanisms and policies to match an application. Figure 2 shows our architecture. Architectural adaptation can be used in the bindings, mechanisms, and policies on the interaction of processing, memory, and communication resources while keeping the macro-level organization the

same and thus preserving the programming model for developing applications. Depending upon the hardware technology used and the support available from the runtime environment, this adaptation can be done statically or at run-time.

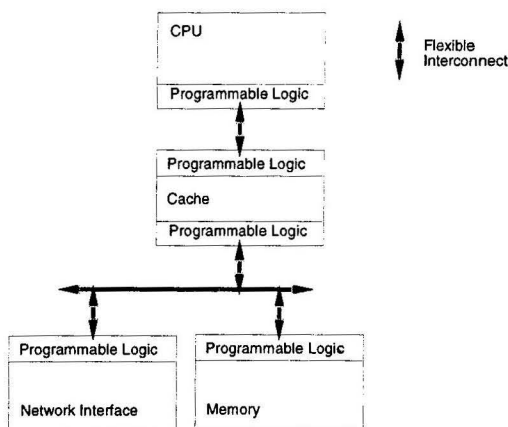


Figure 2. An Architecture for Adaptation

Architectural adaptation provides the mechanisms for application-specific hardware assists to overcome the rigid architectural choices in modern computer systems that do not work well across different applications and often cause substantial performance fragility. In particular, the integration of programmable logic with memory components enables application-specific locality optimizations. These optimizations can be designed to overcome long latency and limited transfer bandwidth in the memory hierarchy. In addition, because the entire application remains in software while the underlying hardware is adapted for system performance, our approach improves over co-processing architectures by preserving machine usability through software. The main disadvantage of our approach is the potential increase on system design and verification time due to the addition of programmable logic. We believe that the advances in design technology will address the increase of logic complexity.

3.1 Project Context

Our study is in the context of the MORPH [5] project, a NSF point design study for Petaflops architectures in the year 2007 technology window. The key elements of the MORPH (MultiprocessOR with Reconfigurable Parallel Hardware) architecture consists of processing and memory elements embedded in a scalable interconnect. With a small amount of programmable logic integrated with key elements of the system, the proposed MORPH architecture aims to exploit architectural customization for a broad range of purposes such as:

- control over computing node granularity (processor-memory association)
- interleaving (address-physical memory element mapping)
- cache policies (consistency model, coherence protocol, object method protocols)
- cache organization (block size or objects)
- behavior monitoring and adaptation

As an example of its flexibility, MORPH could be used to implement either a cache-coherent machine, a non-cache coherent machine, or even clusters of cache coherent machines connected by put/get or message passing. In this paper, we focus on architectural adaptation in the memory system for locality optimizations such as latency tolerance.

4 Case Studies

We present two case studies of architectural adaptation for application-specific locality optimizations. On modern architectures with deep memory hierarchies, data transfer bandwidth and access latency differentials across levels of memory hierarchies can span several orders of magnitude, making locality optimizations critical for performance. Although compiler optimizations can be effective for regular applications such as dense matrix multiply, optimizations for irregular applications can greatly benefit from architectural support. However, numerous studies have shown that no fixed architectural policies or mechanisms, e.g., cache organization, work well for all applications, causing performance fragility across different applications. We present two case studies of architectural adaptation using application-specific knowledge to enhance latency tolerance and efficiently utilize network bisection on multiprocessors.

4.1 Experimental Methodology

As our application examples, we use the sparse matrix library SPARSE developed by Kundert and Sangiovanni-Vincentelli (version 1.3 available from <http://www.netlib.org/sparse/>), and an additional sparse matrix multiply routine that we wrote. This library implements an efficient storage scheme for sparse matrices using row and column linked lists of matrix elements as shown in Figure 3. Only nonzero elements are represented, and elements in each row and column are connected by singly linked lists via the nextRow and nextCol fields. Space for elements, which is 40 bytes per matrix element, are allocated dynamically in blocks of elements for efficiency. There are also several one dimensional arrays for storing the root pointers for row and column lists.





Figure 3. Sparse Library Data Structures

We perform cycle-based system-level simulation using a program-driven simulator based on MINT [22] that interprets program binaries and models configurable logic blocks behaviorally. The details of our simulation environment are presented in [4]. Table 1 shows the simulation parameters used. We report our empirical results for current day computer technologies and then use derivative metrics (such as miss rate) to extrapolate potential benefits for future computer systems which will exhibit much higher clock rates and memory sizes. We also manually translated the C routines modeling the customizable logic blocks into HardwareC [18] to evaluate their hardware cost in terms of space and cycle delays. (However, our recent work is focused on automatic translation of these routines to synthesizable blocks [19].)

	L1 Cache	L2 Cache
Line Size	32B or 64B	32B or 64B
Associativity	1	2
Cache Size	32KB	512KB
Write Policy	Write back + Write allocate	Write back + Write allocate
Replacement Policy	Random	Random
Transfer Rate	(L1-L2) 16B/5 cycles	(L2-Mem) 8B/15 cycles

Table 1. Simulation Parameters

4.2 Architectural Adaptation for Latency Tolerance

Our first case study uses architectural adaptation for prefetching. As the gap between processor and memory speed widens, prefetching is becoming increasingly important to tolerate the memory access latency. However, oblivious prefetching can degrade a program's performance by saturating bandwidth. We show two example prefetching schemes that aggressively exploit application access pattern

information.

Figure 4 shows the prefetcher implementation using programmable logic integrated with the L1 cache. The prefetcher requires two pieces of application-specific information: the address ranges and the memory layout of the target data structures. The address range is needed to indicate memory bounds where prefetching is likely to be useful. This is application dependent, which we determined by inspecting the application program, but can easily be supplied by the compiler. The program sets up the required information and can enable or disable prefetching at any point of the program. Once the prefetcher is enabled, however, it determines what and when to prefetch by checking the virtual addresses of cache lookups to check whether a matrix element is being accessed.




Figure 4. Organizations of Prefetcher Logic

The first prefetching example targets records spanning multiple cache lines and for our example, prefetches all fields of a matrix element structure whenever some field of the element is accessed. The pseudocode of this prefetching scheme for the sparse matrix example is shown below, assuming a cache line size of 32 bytes, a matrix element padded to 64 bytes, and a single matrix storage block aligned at 64-byte boundary. Prefetching is triggered only by read misses. Because each matrix element spans two cache lines, the prefetcher generates an additional L2 cache lookup address from the given physical address (assuming a lock-up free L2 cache) that prefetches the other cache line not yet referenced.

```

/* Prefetch only if vAddr refers to the matrix */
GroupPrefetch(vAddr, pAddr, startBlock, endBlock) {
  if (startBlock <= vAddr && vAddr < endBlock) {
    /* Determine the prefetch address */
    if (pAddr & 0x20) ptrLoc = pAddr - 0x20;
    else ptrLoc = pAddr + 0x20;
    <Initiate transfer of line at ptrLoc to L1 cache>
  }
}

```

The second prefetching example targets pointer fields that are likely to be traversed when their parent structures are accessed. For example, in a sparse matrix-vector mul-

tively, the record pointed to by the nextRow field is accessed close in time with the current matrix element. The pseudocode below shows the prefetcher code for prefetching the row pointer, assuming a cache line size of 64 bytes. Again prefetching is triggered only by read misses, and the prefetcher generates an additional address after the initial cache miss is satisfied using the nextRow pointer value (whose offset is hardwired at setup time) embedded in the data returned by the L2 cache.¹

```

/* Prefetch only if vAddr refers to the matrix */
PointerPrefetch(data, vAddr, startBlock, endBlock) {
  if (startBlock <= vAddr && vAddr < endBlock) {
    /* Get row pointer from returned cache line */
    ptrLoc = data [24]; /* row ptr offset = 24 */
    <Initiate transfer of elt at ptrLoc to L1 cache>
  }
}

```

Our prefetching examples are similar to the prefetching schemes proposed in [23], where they are shown to benefit various irregular applications. However, unlike [23], using architectural customization enables more flexible prefetching policies, e.g., multiple level prefetch, according to the application access pattern.

4.3 Architectural Adaptation for Bandwidth Reduction

Our second case study uses a sparse matrix-matrix multiply routine as an example to show architectural adaptation to improve data reuse and reduce data traffic between the memory unit and the processor. The architectural customization aims to send only used fields of matrix elements during a given computation to reduce bandwidth requirement using dynamic scatter and gather. Our scheme contains two units of logic, an address translation logic and a gather logic, shown in Figure 5.

The two main ideas are prefetching of whole rows or columns using pointer chasing in the memory module and packing/gathering of only the used fields of the matrix element structure. When the root pointer of a column or row is accessed, the gather logic in the main memory module chases the row or column pointer to retrieve different matrix elements and forwards them directly to the cache. The cache, in order to avoid conflict misses, is split into two parts: one small part acting as a standard cache for other requests and one part for the prefetched matrix elements only. The latter part has an application-specific management policy, and can be distinguished by mapping it to a reserved

¹As pointed in [23], the implementation of this prefetching scheme is complicated by the need to translate the virtual pointer address to physical address. We assume that the prefetcher logic can also access the TLB structure. An alternative implementation is to place the prefetcher logic in memory and forward the data of the next record to the upper memory hierarchy. This requires an additional group translation table [23] for address translation.

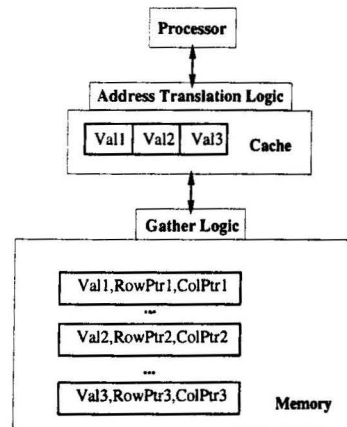


Figure 5. Scatter and Gather Logic

address space. The gather logic in pseudocode is shown below.

```

/* Row gather: pAddr is the start of a row */
Gather(pAddr) {
  chaseAddr = pAddr;
  while(chaseAddr) {
    forward chaseAddr->val
    forward chaseAddr->row
    chaseAddr=virtual-to-physical(chaseAddr->nextRow)
  }
}

```

Because the data gathering changes the storage mapping of matrix elements, in order not to change the program code, a translate logic in the cache is required to present “virtual” linked list structures to the processor. When the processor accesses the start of a row or column linked list, a prefetch for the entire row or column is initiated. Because the target location in the cache for the linked list is known, instead of returning the actual pointer to the first element, the translate logic returns an address in the reserved address space corresponding to the location of the first element in the explicitly managed cache region. In addition, when the processor accesses the next pointer field, the request is also detected by the translate logic, and an address is synthesized dynamically to access the next element in this cache region. The translate logic in pseudocode is shown below.

```

Translate(vAddr, pAddr, newPAddr) {
  /* check if accessing start of a row */
  if(startRowRoot<=vAddr && vAddr<=endRowRoot)
    <Initiate prefetch and return row location>
  /* Similarly for column roots */
  ...
  /* Accessing packed rows */
  if(startPackedRows<=pAddr && pAddr<=endPackedRows) {
    off = pAddr & 63; /* get field offset */
    if( off == 24 ) /* row ptr. at offset 24 */
      return pAddr + 64; /* synthesize next addr. */
    else { /* Two fields at off1 and 2 are packed
           at new_off1 and 2 of new layout */
      if(off < off_2) new_off = off-(off1-new_off1);
      else new_off = off-(off2-new_off2);
    }
  }
}

```

```

return (pAddr>>6) * PACKED_SIZE + new_off; }}
/* Similarly for packed columns */
}

```

As in the case of dense matrix-matrix multiply, blocking of the matrices in the cache is essential to achieve good performance. However, software blocking for sparse matrix-matrix multiply is not as effective as blocking in the dense case because the additional fields in a matrix element structure reduce the number of rows or columns that can fit in the cache, reducing reuse, and because elements are scattered in memory, increasing conflict misses. Figures 6 and 7 compare the miss rates and total data volume transferred between the memory and processor for the straightforward implementation, a pure software blocking scheme, customized scatter and gather, and customized logic with cache bypass for the result matrix (for two 460x460 matrices with 21660 non-zero elements each). The figures show that the hardware assists provided by customization significantly improve the software blocking by reducing row and column footprints to improve reuse. The final results show a 10X in read miss rates and 100X reduction in data volume compared to the straightforward implementation. The implementation costs of these units are less than 1000 in LSI logic 10K family gates and around 1500 in FPGA CLBs, requiring delays of about 3 cycles.




Figure 6. Cache miss rates of different implementations of sparse matrix-matrix multiply.

5 Discussion and Related Work

Traditional computer architectures are designed for best machine performance averaged across applications. Due to the static nature of such architectures, such machines are limited in exploiting an application characteristics unless it is common for a large number of applications. Since no single machine organization fits all applications, the delivered performance is often only a small fraction of the peak ma-




Figure 7. Data traffic volume of different schemes. (Total size of non-zeros, 1.35 MB)

chine performance (frequently less than a tenth [5]). Therefore, we believe that there are significant opportunities for application-specific architectural adaptation. In this paper, we have demonstrated mechanisms for latency hiding and required bandwidth reduction that leverage small hardware support as well as do not change the programming model. Following the same methodology, we can build such assists for other applications as well. Among other examples that applications can benefit from are mechanisms for recognition of working set size for a given application that can be used to alter cache update policies or even use a synthesized small victim cache [16], and mechanisms for monitoring access patterns and conflicts in the caches or memory banks and reconfiguring the assists according to these patterns and conflicts. In summary, we expect an adaptable machine to have a large number of application-specific assists that alter architectural mechanisms and policies in view of application characteristics. The application developer with the help of compilation tools selects appropriate hardware assists to customize the machine to match the application without having to repartition the system functionality or rewrite the application.

There are other approaches for locality optimizations. Researchers have proposed processor-in-memory (PIM) [2, 3, 7, 8, 21] as a solution for solving the latency and bandwidth limitations of the memory hierarchy. We rely on more traditional processor-memory structures with customizable components, which we believe will yield a machine with more accessible performance than an organization in which processors are accessing primarily their local on-chip memory, particularly for irregular applications. In addition, by adding a small amount of programmable logic to the memory units, we can yield some benefits of having computational elements within the memory. Researchers have also proposed various specific architectural mechanisms for locality optimizations, for instance, group prefetching [23]

and informing memory operations [15], with mechanism-specific implementations. As shown in our case studies, we believe that integrating programmable logic into memory components provide a flexible implementation framework for these mechanisms.

6 Conclusions and Future Directions

The increasing dominance of interconnect delays on system performance makes it practical to integrate programmable logic into all key system components, enabling them to be customized to specific applications. As illustrated by two case studies, such architectural adaptation can provide flexible mechanisms for application-specific locality optimizations to combat the increasing gap between processor and memory speed. In addition, system co-design using this approach presents a way to utilize application-specific hardware much more effectively than co-processing architectures, where part of an application is implemented in hardware. Although not explored in this study, architectural adaptations in processor functional units might also be a promising approach to improve instruction throughput. As future directions, we are studying compilation tools in hardware/software co-design for the architectural assists, IC processing issues for implementing customizable system components, such as cache and scalable interconnects, as well as more, larger application studies.

Acknowledgments

The authors gratefully acknowledge support from National Science Foundation award ASC-96-34947. The authors are also supported by an NSF Young Investigator Award CCR-94-57809 and NSF CAREER Award MIP-95-01615.

References

- [1] Semiconductor technology workshop conclusions. Roadmap, Semiconductor Industry Association, 1993.
- [2] BORKAR, S., COHN, R., COX, G., GLEASON, S., GROSS, T., KUNG, H. T., LAM, M., MOORE, B., PETERSON, C., PIEPER, J., RANKIN, L., TSENG, P. S., SUTTON, J., URBANSKI, J., AND WEBB, J. iWarp: An integrated solution to high-speed parallel computing. In *Proc. IEEE Supercomputing '88* (1988), pp. 330–341. Orlando, Florida.
- [3] BORKAR, S., COHN, R., COX, G., GROSS, T., KUNG, H. T., LAM, M., LEVINE, M., MOORE, B., MOORE, W., PETERSON, C., SUSMAN, J., SUTTON, J., URBANSKI, J., AND WEBB, J. Supporting systolic and memory communication in iWarp. In *Proc. 17th IEEE Int. Symp. on Comp. Arch.* (1990), pp. 70–81.
- [4] CHIEN, A., DASDAN, A., GUPTA, R., AND ZHANG, B. Rapid Architectural Design and Validation Using Program-Driven Simulations. In *Symp. on High-level Design, Validation and Test (HLDVT)* (1996).
- [5] CHIEN, A. A., AND GUPTA, R. K. MORPH: A system architecture for robust high performance using customization. In *Frontiers of Massive-Parallelism* (1996). Annapolis, Maryland.
- [6] CHOI, L., AND CHIEN, A. A. The design and performance evaluation of the DI-multicomputer. *J. Parallel and Distributed Comput.* 36:119–143, 1996.
- [7] DALLY, W., CHIEN, A., FISKE, J., FYLER, G., HORWAT, W., KEEN, J., LETHIN, R., NOAKES, M., NUTH, P., AND WILLS, D. The message driven processor: an integrated multicomputer processing element. In *Proc. IEEE Int. Conf. on Comp. Design: VLSI in Computers and Processors* (1992), pp. 416–19.
- [8] DALLY, W. J., FISKE, J. A. S., KEEN, J. S., LETHIN, R. A., NOAKES, M. D., NUTH, P. R., DAVISON, R. E., AND FYLER, G. A. The message-driven processor. *IEEE Micro* (April 1992), pp. 23–39.
- [9] DALLY, W. J., AND SONG, P. Design of a self-timed VLSI multicomputer communication controller. In *Proc. IEEE Int. Conf. on Comp. Design* (1987), pp. 230–4.
- [10] ERNST, R., HENKEL, J., AND BENNER, T. Hardware-Software Cosynthesis for Microcontrollers. *IEEE Design & Test of Computers* (Dec. 1993), pp. 64–75.
- [11] GAJSKI, D., DUTT, N., WU, C. H., AND LIN, Y. L. *High-level Synthesis: Introduction to chip and system design*. Kluwer, 1992.
- [12] GUPTA, R. K. *Co-Synthesis of Hardware and Software for Digital Embedded Systems*. Kluwer, 1995.
- [13] GUPTA, R. K., AND MICHELI, G. D. Hardware-Software Cosynthesis for Digital Systems. *IEEE Design & Test of Computers* (Sept. 1993), pp. 29–41.
- [14] HENNESSY, J. L., AND JOUPPI, N. P. Computer technology and architecture: An evolving interaction. *IEEE Computer* (Sep 1991), pp. 18–29.
- [15] HOROWITZ, M., MARTONOSI, M., MOWRY, T. C., AND SMITH, M. D. Informing memory operations: Providing memory performance feedback in modern processors. In *Proc. 23rd IEEE Int. Symp. on Comp. Arch.* (1996).
- [16] JOUPPI, N. P. Improving direct-mapped cache performance by the addition of a small fully-associative cache and prefetch buffers. In *Proc. the Int. Symp. on Comp. Arch.* (1990), pp. 364–73.
- [17] KOGGE, P. M. Summary of the architecture group findings. In *PetaFlops Architecture Workshop (PAWS)* (Apr. 1996).
- [18] KU, D., AND MICHELI, G. D. HardwareC - A Language for Hardware Design (version 2.0). CSL Tech. Rep. CSL-TR-90-419, Stanford Univ., Apr. 1990.
- [19] LIAO, S. Y., TJIANG, S., AND GUPTA, R. K. An Efficient Implementation of Reactivity in Modeling Hardware in the CSYN Synthesis and Simulation Environment. In *Proc. Design Automation Conf.* (1997).
- [20] SEITZ, C. L. Let's route packets instead of wires. In *Proc. 6th MIT Conf. on Advanced Research in VLSI* (1990), W. J. Dally, Ed., MIT Press, pp. 133–37.
- [21] SPERTUS, E., GOLDSTEIN, S. C., SCHAUSER, K. E., VON EICKEN, T., CULLER, D. E., AND DALLY, W. J. Evaluation of mechanisms for fine-grained parallel programs in the J-Machine and the CM-5. In *Proc. 20th IEEE Int. Symp. on Comp. Arch.* (1993), pp. 302–313.
- [22] VEENSTRA, J. E., AND FOWLER, R. J. MINT: A front end for efficient simulation of shared-memory multiprocessors. In *Proc. 2nd Int. Workshop on Modeling, Analysis, and Simulation of Comp. and Telecommunication Syst. (MASCOTS)* (Jan. 1994), pp. 201–207.
- [23] ZHANG, Z., AND TORRELLAS, J. Speeding up irregular applications in shared-memory multiprocessors: Memory binding and group prefetching. In *Proc. 22th IEEE Int. Symp. on Comp. Arch.* (1995).

NOTICE: WARNING CONCERNING COPYRIGHT RESTRICTIONS

- The copyright law of the United States (Title 17, United States Code) governs the making of photocopies or other reproductions of copyrighted material.
- Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specific conditions is that the photocopy or reproduction is not to be "used for any purpose other than private study, scholarship, or research." If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of "fair use," that user may be liable for copyright infringement.
- This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve violation of copyright law.

Author Index

Otten , R.H.J.M.....	62	Chen, C.-T.....	562
Al-Khalili, A.J.	178	Chen, I.-C.K.....	593
Al-Khalili, D.	178	Chen, I.-C.K.....	88
Allara, A.	400	Chen, J.	474
Arabi, K.	462	Chen, Y.-P.....	217
Arunachalam, R.	224	Cheng , F.-C.....	301
Asanovic, K.....	2	Cheung , P.Y.K.....	307
Augusteijn, L.....	322	Chien, A.A.....	150
Bacis, M.	654	Choi, H.	68
Baklashov, M.....	636	Chu , W.	690
Banerjee, P.	621	Clarke , E.M.	332
Barke , E.....	489	Corno, F.	381
Beece, D.K.	232	Curran, B.....	247
Bhatia , D.	614	Damiano, R.....	317
Bischoff, G.P.....	16	Dartu, F.	224
Blaauw, D.	143	Dasdan, A.	150
Blaton, R.D.	372	Dasgupta, A.	740
Bolchini, C.	555	Davey, A.....	295
Bormann, D.S.....	307	Davidson , E.S.	519
Brace, K.S.....	16	Dharchoudhury, A.....	143
Brayton, R.K.	344	Dunning, J.	143
Brayton, R.K.	352	Dutt, N.D.	587
Breternitz, Jr, M.	170	Ebergen, J.C.	700
Brock, B.C.....	31	El-Kareh, B.....	279
Brown, A.....	2	Elmasry, M.I.....	700
Bruma, S.....	62	Emerson, E.A.....	455
Brunvand, E.	662	Falquez, F.J.	544
Buonanno, G.....	654	Fang, W.-C.....	429
Burch , C.....	97	Ferrandi, F.	654
Canfield, W.	455	Filipponi, S.	400
Cangellaris, A.C.	136	Fornaciari, W.....	400
Carey, S.	232	Fromm, R.....	2
Carlson , B.S.....	732	Fu, B.....	322
Castaneda, R.	258	Fujii, T.	288
Chai, S.M.	408	Fujita, M.	438
Chakrabarti , C.	571	Fujita, M.	445
Chan, Y.....	247	Fukasaku, I.....	288
Chandy, J.A.....	621	Fummi, F.	654
Chang, N.....	720	Ganguly , S.	740
Chang, Y.-N.....	186	Gao, G.R.....	82
Chappell , B.....	317	Gerli, L.	654
Chen, B.	82	Ginosar , R.....	706

Golbus, J.....	2	Kirihata, T.....	279
Gosti, W.....	344	Ko , U.....	726
Govindarajan, S.....	752	Kol, R.....	706
Gribstad, B.....	2	Kozuch, M.A.....	273
Gupta, R.K.....	150	Kozyrakis, C.....	2
Haddad, R.....	110	Kucukcakar , K.....	562
Hamano , H.....	685	Kudva , P.....	423
Hatch, R.....	232	Kuwako, M.....	288
Hayes, J.P.....	372	Lai, K.....	636
Henkels, W.H.....	712	Lee, C.-C.....	593
Henning, R.....	571	Lee, C.-C.....	88
Hojat, S.....	206	Lee, C.-K.....	211
Hoover, Jr., J.M.....	258	Lee, J.-C.....	48
Hsu, W.W.....	578	Leeser, M.....	417
Hu, A.J.....	438	Leijten, J.A.....	164
Hunt, Jr., W.A.....	31	Li, D.....	726
Hunter , C.....	549	Li, Y.....	388
Hwang , S.H.....	68	Li, Y.....	690
Hwang, W.....	279	Li, Z.....	352
Hwang, W.....	712	Liptay , J.S.....	241
Ima, M.....	288	Liu, C.L.....	122
Inamori, M.....	157	Lombardi, F.....	608
Ishii, K.....	157	Loyd, D.....	295
Jain, J.....	445	Lu, Y.....	332
Jain, S.....	16	Luk, W.K.....	279
Jess, J.A.G.....	164	Mak, W.-K.....	496
Jha, S.....	332	Malley, C.....	317
John, L.K.....	510	Mankino, H.....	685
Joshi, R.....	279	Manoli, Y.....	317
Joshi, R.V.....	712	Martin, D.....	2
Jou, J.-Y.....	642	Mashiko, K.....	685
Joyner, W.....	328	McCluskey , E.J.....	54
K Yelick,	2	Mehrotra, A.....	670
Kagaris, D.....	366	Mehrotra, R.....	130
Kaminska , B.....	462	Mehta, A.D.....	217
Kanevsky, V.....	720	Meinel, C.....	338
Kang , S.....	48	Menezes, N.....	217
Kang, Y.-S.....	48	Menon , P.R.....	648
Karandikar, A.K.....	122	Min, Y.....	352
Karayiannis, D.....	366	Minea, M.....	332
Katayama, Y.....	279	Miyazaki, T.....	157
Katz , R.H.....	670	Moreno , WA.....	544
Katz, R.....	604	Mudge , T.N.....	593
Kaufmann, M.....	25	Mudge , T.N.....	88
Keeton, K.....	2	Muhammad, K.....	196
Kenneth, L.S.....	532	Munetoh, S.....	279

Nakada, K.....	157
Nakagawa, O.S.....	720
Nakamura, H.	587
Nanya , T.....	288
Narayan, A.	445
Nedelchev , I.....	82
Nicolau, A.	322
Nicolau, A.	587
Nien , M.-C.....	642
Northrop, G.	232
Norton, J.....	143
Nowick, S.....	662
Oh , S.-Y.	720
Oi, H.....	267
Ong, S.	578
Ozawa, M.....	288
Pain, B.	429
Pan , P.....	116
Pan, P.....	122
Panda, P.R.....	587
Papachristou, C.A.....	636
Parhi , K.K.....	186
Parhi , K.K.....	676
Park, N.	608
Patel, C.S.....	408
Patterson, D.	2
Pedram, M.	130
Pedram, M.	211
Peir, J.-K.....	578
Perissakis, S.	2
Pileggi , L.T.	217
Pileggi , L.T.	224
Pillai, R.V.K.....	178
Pinello, W.	136
Pixley , C.....	25
Polychronopoulus, C.	322
Pomeranz, I.	360
Postiff, M.A.....	88
Potkonjak, M.	388
Prinetto, P.	381
Pua, A.	726
Pullela, S.	143
Qadeer, S.....	670
Rahmat, K.	720
Ramachandran , A.....	474
Ranganathan , N.....	267

Ranjan, R.K.	344
Ranjan, R.K.	670
Razdan , R.	16
Rebaudengo, M.....	381
Reddy , S.M.....	360
Reed, D.F.	317
Reorda, M.S.	381
Rhodes, D.L.	393
Rivers, J.A.	519
Rohrer , R.	10
Roy , K.....	196
Roy , K.....	746
Ruehli, A.	136
Saha , A.....	455
Salice, F.	400
Salice, F.	555
Samon Jr., J.R.....	544
Sangiovanni-Vincentelli, A.....	344
Sangiovanni-Vincentelli, A.....	445
Satoh, A.	279
Satyanarayana, J.H.....	186
Savir , J.....	37
Schimmel , D.E.....	408
Schulz, M.	150
Schulz, S.E.....	317
Sciuto , D.	654
Sciuto, D.	400
Sciuto, D.	555
Shams, M.....	700
Shenoy , N.	110
Shepard, K.L.	232
Sheu , B.J.....	429
Shirakawa H.....	157
Sigal, L.....	247
Slobodova, A.	338
Smith, R.....	170
Smith, T.	544
Soderquist, P.	417
Squillero , G.....	381
Srivastava, P.	726
Stohmann, J.	489
Subramanian , A.	510
Suzuki, H.	685
Takamura, A.....	288
Tam, E.S.	519
Tan, P.....	608

Tan, X.	608
Tarlescu, M.-D.....	82
Tekumalla, R.C.	648
Theobald, K.B.....	82
Thomas, R.....	2
Tierno, J.A.	423
Timmer, A.A.	164
Tong, J.	608
Touba, N.A.....	54
Tragoudas, S.....	366
Treuhaf, N.....	2
Tsutsui, A.	157
Tumbush, G.....	614
Tuminaro , A.	253
Ueno, Y.	288
Vaishnav, H.....	211
van Ginneken, L.P.P.P.	110
van Meerbergen, J.L.....	164
Vemuri , R.	752
Villarrubia , P.....	206
Wang, C.-P.....	468
Wang, C.-Y.	746
Warnock, J.....	247
Webb, C.F.	241
Wey , C.-L.....	468
Wilson, C.	438

Wolf, W.....	273
Wolf, W.....	388
Wolf, W.....	393
Wolfe , A.....	273
Wong , D.F.	482
Wong , D.F.	496
Wong, D.F.	217
Wong, D.F.....	628
Wong, H.	279
Wordeman, M.	279
Wu, X.	130
Xio, P.....	279
Yalamanchili, S.	408
Yan , Y.	504
Yang, G.	429
Young, F.Y.	482
Young, H.	578
Yun, K.	662
Zhang, C.....	504
Zhang, X.....	150
Zhang, X.....	258
Zhang, X.....	504
Zhao, Y.....	352
Zhou, H.	628
Zhu, Z.....	732

NOTES

NOTES

Press Activities Board

Vice President:

I. Mark Haas
Managing Partner
Haas Associates
P.O. Box 451177
Garland, TX 75045-1177
m.haas@computer.org

Editor-in-Chief

Advances in Computer Science and Engineering Board
Pradip Srimani
Colorado State University
Dept. of Computer Science
601 South Hous Lane
Fort Collins, CO 80525
Phone: 970-491-5862 FAX: 970-491-2466
srimani@cs.colostate.edu

Jon T. Butler, Naval Postgraduate School
James J. Farrell III, Motorola
Mohamed E. Fayad, University of Nevada
I. Mark Haas, Haas Associates
Ronald G. Hoelzeman, University of Pittsburgh
Gene F. Hoffnagle, IBM Corporation
John R. Nicol, GTE Laboratories
Yale N. Patt, University of Michigan
Benjamin W. Wah, University of Illinois
Ronald D. Williams, University of Virginia

Editor-in-Chief

Practices for Computer Science and Engineering Board
Mohamed E. Fayad
Computer Science, MS/171
Bldg. LME, Room 308
University of Nevada
Reno, NV 89557
Phone: 702-784-4356 FAX: 702-784-1833
fayad@cs.unr.edu

IEEE Computer Society Executive Staff

T. Michael Elliott, Executive Director
Matthew S. Loeb, Publisher

IEEE Computer Society Publications

The world-renowned Computer Society publishes, promotes, and distributes a wide variety of authoritative computer science and engineering texts. These books are available in two formats: 100 percent original material by authors preeminent in their field who focus on relevant topics and cutting-edge research, and reprint collections consisting of carefully selected groups of previously published papers with accompanying original introductory and explanatory text.

Submission of proposals: For guidelines and information on Computer Society books, send e-mail to cs.books@computer.org or write to the Acquisitions Editor, IEEE Computer Society, P.O. Box 3014, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720-1314. Telephone +1 714-821-8380. FAX +1 714-761-1784.

IEEE Computer Society Proceedings

The Computer Society also produces and actively promotes the proceedings of more than 130 acclaimed international conferences each year in multimedia formats that include hard and softcover books, CD-ROMs, videos, and on-line publications.

For information on Computer Society proceedings, send e-mail to cs.books@computer.org or write to Proceedings, IEEE Computer Society, P.O. Box 3014, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720-1314. Telephone +1 714-821-8380. FAX +1 714-761-1784.

Additional information regarding the Computer Society, conferences and proceedings, CD-ROMs, videos, and books can also be accessed from our web site at <http://computer.org/cspress>

Appendix EDM02

Georgia Institute of Technology ILL

ILLiad TN: 462851

Borrower: Pillsbury Winthrop

Lending String:

Patron: Eileen D. McCarrier

Journal Title: IEEE Computer Society Workshop
on VLSI 2000

Volume: Issue:

Month/Year: 2000 **Pages:** 75-79, title, etc. SEE
NOTES

Article Author: Rajesh Gupta

Article Title: Architectural Adaptation in AMRM
Machines - SEE NOTES IN ILLIAD

Imprint:

ILL Number:



Call #:
50671011813497

Location: LSC

Charge

Maxcost: NOTIFY GB FOR INVOICE

Shipping Address:

Pillsbury Winthrop Shaw Pittman LLP
1650 Tysons Blvd, 14th Floor
McLean, VA 22102-4856

Fax:

Ariel:

Email: eileen.mccarrier@pillsburylaw.com

This article has been provided to you by the
Information Delivery Department at the Georgia Tech Library.
Please contact us at 404-894-3989 or
technicalservices@library.gatech.edu
if you have any questions or concerns.

Notice: Warning Concerning Copyright Restrictions

The Copyright Law of the United States (Title 17, United States Code) governs the making of photocopies or other reproductions of copyrighted material.

Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction.

This material may be protected by copyright law. Any reproduction or distribution of this material, in any format, may be an infringement of the Copyright Law. This reproduction is not to be "used for any purpose other than private study, scholarship, or research." If you make or later uses a photocopy or reproduction for purposes in excess of "Fair Use," that user may be liable for copyright infringement.

Intel Exhibit 1028-34

Architectural Adaptation in AMRM Machines

Rajesh Gupta

*Information and Computer Science
University of California, Irvine
Irvine, CA 92697
rgupta@ics.uci.edu*

Abstract

Application adaptive architectures use architectural mechanisms and policies to achieve system level performance goals. The AMRM project at UC Irvine focuses on adaptation of the memory hierarchy and its role in latency and bandwidth management. This paper describes the architectural principles and first implementation of the AMRM machine proof-of-concept prototype.

1. Introduction

Modern computer system architectures represent design tradeoffs and optimizations involving a large number of variables in a very large design space. Even when successfully implemented for high performance, which is benchmarked against a set of representative applications, the performance optimization is only in an average sense. Indeed, the performance variation across applications and against changing data set even in a given application can easily be by an order of magnitude [1]. In other words, delivered performance can be less than one tenth of the system performance that the underlying hardware is capable of.

A primary reason of this fragility in performance is that rigid architectural choices related to organization of major system blocks (CPU, cache, memory, IO) do not work well across different applications.

Architectural Adaptivity provides an attractive means to ensure robust high performance. Architectural adaptation refers to the capability of a machine to support multiple architectural mechanisms and policies that can be tailored to application and/or data needs [2]. There are a number

of places where architectural adaptivity can be used, for instance, in tailoring the interaction of processing with I/O, customization of CPU elements (e.g., splittable ALU resources) etc.

In view of the microelectronic technology trends that emphasize increasing importance of communication at all levels, from network interfaces to on-chip interconnection fabrics, communication represents the focus of our studies in architectural adaptation. In this context, memory system latency and bandwidth issues are key determining factors in performance of high performance machines because these can provide a constant multiplier on the achievable system performance [1]. Further, this multiplier decreases as the memory latency fails to improve as fast as processor clock speeds.

Consider a hypothetical machine with processing elements running at 2 GHz with eight-way superscalar pipelines. Assuming a typical 1 microsecond round-trip latency for a cache miss, this corresponds to about 16K instructions, with an average 30% or 4800 instructions being load/store. For a single-thread execution a miss rates as low as 0.02% reduces computing efficiency by as much as 50%. This points to a need for very low miss rates to ensure that high-throughput CPUs can be kept busy. A similar analysis of the bisection bandwidth concludes that active bandwidth management is required to reduce the need for communication and to increase the number of operations before a communication is necessary.

2. The AMRM Project

The Adaptive Memory Reconfiguration Management, or the AMRM, project at the University of California, Irvine aims to find ways to improve the memory system performance of a computing system. The basic system architecture reflects the view that communication is already critical and getting increasingly so [3], and flexible

interconnects can be used to replace static wires at competitive performance in interconnect dominated microelectronic technologies [4,5,6]. The AMRM machine uses reconfigurable logic blocks integrated with the system core to control policies, interactions, and interconnections of memory to processing [7]. The basic machine architecture supports application-specific cache organization and policies, hardware-assisted blocking, prefetching and dynamic cache structures (such as stream, victim caches, stride prediction and miss history buffers) that optimize the movement and placement of application data through the memory hierarchy. Depending upon the hardware technology used and the support available from the runtime environment this adaptation can be done statically or at run-time. In the following section we describe a specific mechanism for latency management that is shown to provide significant performance boost for the class of applications characterized by frequent accesses to linked data structures scattered in the physical memory. This includes algorithms that operate on sparse matrices and linked trees.

2.1 Adaptation for Latency Management

Latency management refers to techniques for hiding long latencies of memory accesses by useful computation. Most common technique for latency hiding is by prefetching of data to the CPU. Prefetching is combined with smaller and faster (cache) memory elements that attempt to prefetch application context(s) rather than single data elements. The pointer-based accesses to data items in memory hierarchies typically yield poor results because the indirection introduces main memory and memory hierarchy latencies into the innermost computational loop. Techniques such as software prefetching (loop unrolling and hoisting of loads) do not adequately solve the problem, as prefetching at the processor leaves multi-level memory hierarchy latency in the critical path. Purely hardware prefetching [8] is also often ineffective because the address references generated by an application may contain no particular address structure. We use an application-specific prefetching scheme that resides in dedicated hardware at arbitrary levels of the memory hierarchy, in all of them, or to bypass them completely. This hardware performs application-specific prefetching, based on the address ranges of data structures used. When there is a reference to an address inside in this range, the prefetch hardware will prefetch the “next” element pointed to by the current element. The pointer field for the next element can be changed at runtime.

This prefetch hardware is combined, for some applications, with address translation and compaction hardware in the memory controller that works well with data structures that do not quite fit into a single cache line. The address translation is done transparently from the application using hardware assist *translate* in the cache controller and a corresponding hardware assist *gather* in the memory controller. Simulation results using this prefetch hardware show a 10X reduction in read miss rates and 100X reduction in data volume reduction for sparse matrix multiply operations [9].

3. The AMRM System Prototype

While AMRM simulation results continue to provide valuable insights into the space of architectural mechanisms and their effectiveness [7][9][10], a system implementation is needed to bring together different parts of the AMRM project (including compiler and runtime system algorithms to support adaptivity). The AMRM system prototype is divided into two phases. First phase consists of implementation of a board-level prototype; followed by a second phase single-chip implementation of the cache memory system. At the time of this writing, the first phase of the project prototype implementation has recently completed. The rest of this paper describes the system design and implementation of the Phase I prototype and its relationship to the ongoing second phase ASIC prototype.

The AMRM phase I prototype board is designed to serve two purposes. It can simulate a range of memory hierarchies for applications running on a host processor. The board supports configurability of the cache memory via an on-board FPGA-based memory controller. The board is also designed to be used, in future, as a complete system platform, with on-board memory serving as main memory, via a mezzanine card containing the Phase II AMRM ASIC implementation.

Whereas the goal of the AMRM prototype is to build an adaptive cache memory system, in general, the memory hierarchy performance cannot be decoupled from the processor instruction set architecture, implementation, and compiler implementation. (This is particularly true of the CPU-L1 path that is often pipelined using non-blocking caches.) It would, therefore, be desirable to evaluate any proposed changes to the memory hierarchy for several processor architectures rather than being tied to one specific processor type and its

software. In order to be able to evaluate the effect with different processor architectures as well as to circumvent the implementation difficulties, our Phase I implementation attaches an additional memory hierarchy to a system through a standard peripheral bus. Thus, the board provides a PCI interface that allows a host processor to use the board as a part of its memory hierarchy. Applications running on the host processor are instrumented automatically using the AMRM compiler to use the memory on the AMRM board. Thus direct program execution can proceed on the host processor while the extra memory hierarchy is being exercised.

3.1 AMRM System Goals

One goal of the AMRM prototype system is that it be adaptable to many different memory hierarchy architectures. Another goal of the AMRM system is that it be useful for running real time program execution or even memory simulations. The latter is accomplished by making the AMRM memory available to the user and converting user program to access this memory “directly”. The former is accomplished through the use of a sequence of address/command type requests “run” through various memory system configurations. The AMRM system is to be fast enough to support extensive execution or simulation.

A CPU interfaces to the reconfigurable AMRM memory system through the PCI bus. AMRM accepts CPU PCI requests for memory operations, issues them to the attached memory system, and sends back the data for memory read operations as well as memory access time information.

3.2 AMRM prototype architecture

Figure 1 shows the main components of the AMRM prototype board. It consists of a general 3-level memory hierarchy plus support for the AMRM ASIC chip implementing architectural assists with in the CPU-L1 datapath. The host interface is managed by a Motorola PLX 9080 processor. The FPGAs on the board contain controllers for the SRAM, DRAM and L1 cache. A 1 MB SRAM is used for tag and data store for the L1 cache. A total of 512 MB of DRAM is provided to implement part of the cache hierarchy (and also to serve as main memory by reloading the memory controller into the FPGA.)

The board implementation necessarily hard-wires certain parameters of the memory hierarchy. This includes the board’s clock. In order to perform

detailed and accurate simulation of diverse memory hierarchy configurations at any clock speed, a hardware “virtual” clock has been implemented as part of the performance monitoring hardware. Performance monitoring hardware primarily includes various event counters, which are memory-mapped and readable from the host processor. The “virtual clock” emulates a target system’s clock: the clock rate is determined by the target system’s memory hierarchy design and technology parameters. For example, the delay for an L1 cache hit, miss fetch etc in terms of virtual clocks can be configured by the host to emulate a given target cache design. Thus, the use of virtual clock allows us to simplify the hardware implementation. For instance, the tag and data stores of L1 cache can be a single RAM while the timing may reflect a design with two separate RAM’s.

3.3 Command Interface to the AMRM Board

The memory hierarchy on the AMRM board can be used by an application running on the host processor by writing commands to specific addresses in the PCI address space. Each command consists of a set of four words that specify the operation (e.g., memory read/write, register read/write), the address of the location to access and data in case of a write. For read commands, a read response is generated and data is written into the host’s memory.

For debugging purposes and to enable the cache to be flushed by the host, there are commands to access memory banks directly, i.e., without going through the caches. Commands are also available to read/write the status, configuration registers and performance counters.

The onboard command processors reads a command and launches its execution in the AMRM board. Data is read from the cache and sent back to the processor if it hits in the cache. It takes m Virtual Clock cycles. Otherwise it is requested from the next level in the hierarchy. Writes take n virtual clocks. Upon load command completion the data can be written into the system memory for access by the host processor. Both parameters n and m can be programmed under compiler control.

3.4 Virtual Clock System

The virtual clock system consists of a master clock counter, $Vtime$, the virtual clock signal, $Vclk$,

Ready inputs, and associate virtual clock generation logic. The *Ready* input from each major memory hierarchy module specifies that this unit has completed the current virtual clock cycle activities. A new *Vclk* edge/period is generated when all *Ready* inputs reach 1. The *Vtime* can be read out by the host processor to determine the current virtual time. It can also be automatically supplied to the CPU via host memory (as opposed to the AMRM board memory).

Each major unit in the memory hierarchy is designed to generate *Ready* and wait for the *Vclk*, when appropriate. In most cases the designs actually use *AutoReady* counters local to each unit which can be loaded with a programmable number of cycles. An *AutoReady* counter generates *Ready* using the *Vclk* while its output is non-zero. An idle unit not processing any requests also outputs a *Ready* signal every *Vclk*. For instance, Consider the AMRM Read command addressed to the on-board memory hierarchy. It includes a delay (ΔT) from the previous memory access. This delay is used to advance the virtual clock forward before starting the new access. This is accomplished by loading it into the *AutoReady* counter. This allows other memory hierarchy activity to proceed in parallel with CPU computation. For instance, a prefetch unit may be accessing memory during the ΔT -cycle delay.

3.5 AMRM chip functionality

The AMRM board provides for incorporation of

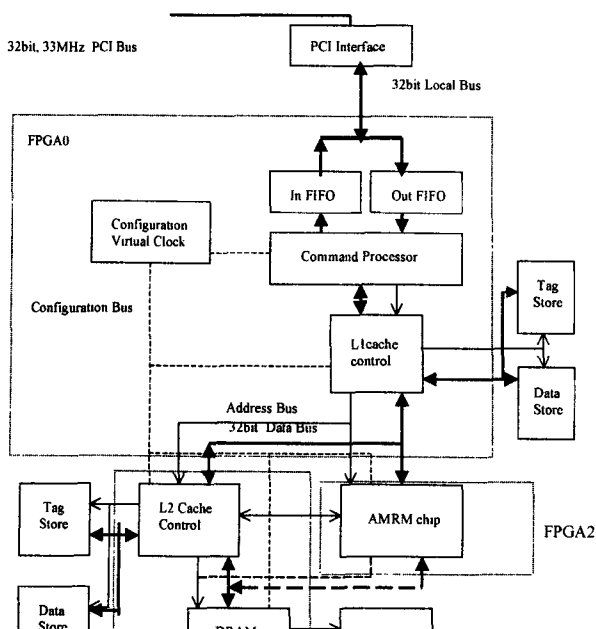


Figure 1: AMRM Phase I Prototype Board

an AMRM chip that uses an ASIC implementation of the AMRM cache assist mechanisms. It is positioned between L1 cache and the rest of the system and can be accessed in parallel with the L2 cache. It can thus accept and supply data coming from or going to the L1 cache. For instance, it may contain a write buffer or a prefetch unit to access L2. It also has access to the memory interface and thus can, for instance, prefetch from memory. The AMRM ASIC design is currently in progress. This chip will include a processor core with adaptive memory hierarchy. When plugged into the AMRM board, the ASIC will use onboard DRAM as main memory by simply reconfiguring the memory controller in FPGA1.

4. Summary

Traditional computer system architectures are designed for best machine performance averaged across applications. Due to the static nature of these architectures, such machines are limited in exploiting application characteristics unless these are common for a large number of applications. Unfortunately, a number of studies have shown that no single machine organization fits all applications, therefore, often the delivered performance is only a small fraction of the peak machine performance. Therefore, we believe that there are significant opportunities for application-specific architectural adaptation. The focus of the AMRM project is on architectural adaptations that close the gap between processor and memory speed by intelligent placement of data through the memory hierarchy. Our current work has demonstrated performance gains due to adaptive cache organizations and cache prefetch assists. In future, we envision adaptive machines that provide a menu of application-specific assists that alter architectural mechanisms and policies in view of the application characteristics. The application developer, with the help of compilation tools, selects appropriate hardware assists to customize the machine to match application needs without having to rewrite the application.

5. Acknowledgement

The AMRM project was conceived through the joint effort of Andrew Chien, Alex Nicolau and Alex Veidenbaum. The team includes Prashant Arora, Chun Chang, Dan Nicolaescu, Rajesh Satapathy, Weiyu Tang, Xiao Mei Ji. The project is sponsored by DARPA under contract DABT63-98-C-0045.

6. References

- [1] P. M. Kogge, "Summary of the architecture group findings," In *PetaFlops Architecture Workshop (PAWS)*, April 1996.
- [2] A. A. Chien, R. K. Gupta, "MORPH: A System Architecture for Robust High Performance Using Customization," In *Proceedings of the Sixth Symposium on the Frontiers of Massively Parallel Computation (Frontiers '96)*, October 1996.
- [3] W. Wulf, S. McKee, "Hitting the memory wall: Implication of the obvious," In *Computer Architecture News*, 1995.
- [4] C. L. Seitz, "Let's route packets instead of wires," In *Proceedings of the 6th MIT Conference on Advanced Research in VLSI*, W. J. Dally, editor, MIT Press, pp. 133-37, 1996.
- [5] W. J. Dally, P. Song, "Design of a self-timed VLSI multicomputer communication controller," In *Proceedings of ICCD*, 1987.
- [6] R. K. Gupta, "Analysis of Technology Trends: A Case of Architectural Adaptation in Custom Datapaths," In *Proceedings of SPIE session on Configurable Computing*, Boston, November 1998.
- [7] AMRM Website: www.ics.uci.edu/~amrm.
- [8] D. F. Zucker, M.J. Flynn, R. B. Lee, "A comparison of hardware prefetching techniques for multimedia benchmarks," *Technical report CSL-TR-95-683*, Computer Systems Laboratory, Stanford University, Stanford University, CA 94305, December 1995.
- [9] X. Zhang, *et al*, "Architectural adaptation for memory latency and bandwidth optimization," In *Proceedings of ICCD*, October, 1997.
- [10] A. Veidenbaum, *et al*, "Adapting Cache Line Size to Application Behavior," In *Proceedings of International Conference on Supercomputing (ICS)*, June 1999.

Proceedings

IEEE Computer Society Workshop on VLSI 2000

System Design for a System-on-Chip Era

27-28 April 2000 Orlando, Florida

Edited by

Asim Smailagic, Robert Brodersen and Hugo De Man

TK
7874.75
.1354
2000

TER
TY

Sponsored by

IEEE Computer Society Technical Committee on VLSI

Intel Exhibit 1028 - 40

Proceedings

**IEEE Computer Society
Workshop on VLSI 2000**

System Design for a System-on-Chip Era



Proceedings

IEEE Computer Society Workshop on VLSI 2000

System Design for a System-on-Chip Era



27–28 April 2000


Orlando, Florida

Edited by

Asim Smailagic, Robert Brodersen, and Hugo De Man

Sponsored by the

IEEE Computer Society Technical Committee on VLSI



Los Alamitos, California

Washington

• **Intel Exhibit 1028-42**

Brussels

Tokyo

15
REC
2005

Copyright © 2000 by The Institute of Electrical and Electronics Engineers, Inc.
All rights reserved

Copyright and Reprint Permissions: Abstracting is permitted with credit to the source. Libraries may photocopy beyond the limits of US copyright law, for private use of patrons, those articles in this volume that carry a code at the bottom of the first page, provided that the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

Other copying, reprint, or republication requests should be addressed to: IEEE Copyrights Manager, IEEE Service Center, 445 Hoes Lane, P.O. Box 133, Piscataway, NJ 08855-1331.

The papers in this book comprise the proceedings of the meeting mentioned on the cover and title page. They reflect the authors' opinions and, in the interests of timely dissemination, are published as presented and without change. Their inclusion in this publication does not necessarily constitute endorsement by the editors, the IEEE Computer Society, or the Institute of Electrical and Electronics Engineers, Inc.

IEEE Computer Society Order Number PR00534
ISBN 0-7695-0534-1
ISBN 0-7695-0536-8(microfiche)
Library of Congress Number 99-069215


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
[http://shop.ieee.org/store/
customer-service@ieee.org](http://shop.ieee.org/store/customer-service@ieee.org)

IEEE Computer Society
Asia/Pacific Office
Watanabe Bldg., 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 Anne Rawlinson
Cover art production by Joe Daigle/Studio Productions
Printed in the United States of America by The Printing House



Intel Exhibit 1028 - 43

Table of Contents



Message from the General Chairs	ix
Message from the Technical Program Chairs	xi
Workshop Committees	xiii
Steering Committee	xv
System Level Design Methods and Examples I	
Mead-Conway VLSI Design Approach and System Design Challenges Ahead	3
<i>Lynn Conway</i>	
Alternative Architectures for Video Signal Processing	5
<i>W. Wolf</i>	
PicoRadio: Ad-hoc Wireless Networking of Ubiquitous Low-energy Sensor/Monitor Nodes	9
<i>J. Rabaey, J. Ammer, J. L. da Silva Jr., and D. Patel</i>	
System Level Design Methods and Examples II	
A System-level Approach to Power/Performance Optimization in Wearable Computers	15
<i>A. Smailagic, D. Reilly, and D. P. Siewiorek</i>	
Emerging Trends in VLSI Test and Diagnosis	21
<i>Y. Zorian</i>	
Multilanguage Design of a Robot Arm Controller: Case Study	29
<i>G. Nicolescu, P. Coste, F. Hessel, P. LeMarrec, and A.A. Jerraya</i>	
Low Power Design	
Instruction Scheduling Based on Energy and Performance Constraints	37
<i>A. Parikh, M. Kandemir, N. Vijaykrishnan, and M.J. Irwin</i>	
Dynamic Voltage Scaling Techniques for Distributed Microsensor Networks	43
<i>R. Min, T. Furrer, and A. Chandrakasan</i>	

Reducing the Power Consumption in FPGAs with Keeping a High Performance Level	47
<i>A. D. Garcia G, W. P. Burlison, and J. L. Danger</i>	
Multiple Access Caches: Energy Implications.....	53
<i>H.S. Kim, N. Vijaykrishnan, M. Kandemir and M.J. Irwin</i>	
System Level Design Examples	
Improved Synchronization Methodologies for High Performance Digital Systems	61
<i>S. Welch and K. Kornegay</i>	
Low Power VLSI Architecture for 2D-Mesh Video Object Motion Tracking	67
<i>W. Badawy and M. Bayoumi</i>	
Timing Issues in System Design	
Architectural Adaptation in AMRM Machines.....	75
<i>R. Gupta</i>	
An Empirical and Analytical Comparison of Delay Elements and a New Delay Element Design	81
<i>N. Mahapatra, S. Garimella and A. Tareen</i>	
Software System Design and Design Environment	
Specification and Validation of Information Processing Systems by Process Encapsulation and Symbolic Execution	89
<i>W. Boßung, T. Geyer, S. A. Huss, and L. Wehmeyer</i>	
Interaction in Language Based System Level Design using an Advanced Compiler Generator Environment	97
<i>I. Poulakis, G. Economakos, and P. Tsanakas</i>	
On Design-for-reusability in Hardware Description Languages.....	103
<i>J. M. Chang and S. K. Agun</i>	
Analysis and Synthesis of Asynchronous Circuits	
Fine-grain Pipelined Asynchronous Adders for High-speed DSP Applications.....	111
<i>M. Singh and S. M. Nowick</i>	
A Low-latency FIFO for Mixed-clock Systems.....	119
<i>T. Chelcea and S. M. Nowick</i>	
Advances in Multiplier Design	
Reconfigurable Low Energy Multiplier for Multimedia System Design.....	129
<i>S. Kim and M. C. Papaefthymiou</i>	
An Algorithmic Approach to Building Datapath Multipliers Using (3,2) Counters.....	135
<i>H. Al-Twaijry and M. Aloqeely</i>	

Issues in System Design

RT-level Interconnect Optimization in DSM Regime 143
S. Katkooi and S. Alupoaei

Validation of Complex Designs through Hardware Prototyping 149
G. Stoler

On-line Error Detection in Multiplexor Based FPGAs 155
A. Jaekel

Author Index 161

Author Index



Agun, S. K.....	103	Kandemi, M.....	37, 53
Aloqeely, M.	135	Katkoori, S.....	143
Al-Twajjry, H.....	135	Kim, H. S.....	53
Alupoaei, S.....	143	Kim, S.	129
Ammer, J.....	9	Kornegay, K.	61
Badawy, W.	67	LeMarrec, P.....	29
Bayoumi, M.	67	Mahapatra, N.	81
Boßung, W.....	89	Min, R.....	43
Burleson, W. P.	47	Nicolescu, G.....	29
Chandrakasan, A.....	43	Nowick, S. M.	111, 119
Chang, J. M.....	103	Papaefthymiou, M. C.	129
Chelcea, T.....	119	Parikh, A.	37
Conway, L.....	3	Patel, D.	9
Coste, P.	29	Poulakis, I.....	97
da Silva Jr., J. L.....	9	Rabaey, J.....	9
Danger, J. L.....	47	Reilly, D.....	15
Economakos, G.....	97	Siewiorek, D. P.....	15
Furrer, T.	43	Singh, M.....	111
Garcia G, A. D.....	47	Smailagic, A.	15
Garimella, S.	81	Stoler, G.....	149
Geyer, T.	89	Tareen, A.	81
Gupta, R.....	75	Tsanakas, P.	97
Hessel, F.....	29	Vijaykrishnan, N.	37, 53
Huss, S. A.	89	Wehmeyer, L.....	89
Irwin, M. J.....	37, 53	Welch, S.....	61
Jaekel, A.	155	Wolf, W.	5
Jerraya, A. A.	29	Zorian , Y.	21



Press Activities Board

Vice President and Chair:

Carl K. Chang
Dept. of EECS (M/C 154)
The University of Illinois at Chicago
851 South Morgan Street
Chicago, IL 60607
ckchang@eecs.uic.edu

Editor-in-Chief

Advances and Practices in Computer Science and Engineering Board
Pradip Srimani
Colorado State University, Dept. of Computer Science
601 South Hows Lane
Fort Collins, CO 80525
Phone: 970-491-7097 FAX: 970-491-2466
srimani@cs.colostate.edu

Board Members:

Mark J. Christensen
Deborah M. Cooper – Deborah M. Cooper Company
William W. Everett – SPRE Software Process and Reliability Engineering
Haruhisa Ichikawa – NTT Software Laboratories
Annie Kuntzmann-Combelles – Objectif Technologie
Chengwen Liu – DePaul University
Joseph E. Urban – Arizona State University

IEEE Computer Society Executive Staff

T. Michael Elliott, Executive Director and Chief Executive Officer
Angela Burgess, Publisher

IEEE Computer Society Publications

The world-renowned IEEE Computer Society publishes, promotes, and distributes a wide variety of authoritative computer science and engineering texts. These books are available from most retail outlets. Visit the Online Catalog, <http://computer.org>, for a list of products.

IEEE Computer Society Proceedings

The IEEE Computer Society also produces and actively promotes the proceedings of more than 141 acclaimed international conferences each year in multimedia formats that include hard and softcover books, CD-ROMs, videos, and on-line publications.

For information on the IEEE Computer Society proceedings, send e-mail to cs.books@computer.org or write to Proceedings, IEEE Computer Society, P.O. Box 3014, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720-1314. Telephone +1 714-821-8380. FAX +1 714-761-1784.

Additional information regarding the Computer Society, conferences and proceedings, CD-ROMs, videos, and books can also be accessed from our web site at <http://computer.org/cspress>

Message from the General Chairs



Welcome to Orlando, to WVLSI '2000. We hope and trust that those of you attending the workshop find it to be both enjoyable and a productive use of your time. WVLSI has become a regular annual forum for researchers to exchange ideas in the area of VLSI and system level design, in particular.

This workshop has been successful over the past decade due to the many members of the VLSI community who have volunteered their efforts. In particular, we would like to thank the Program Co-chairs Asim Smailagic, Robert Broderson and Hugo De Man for having put together a program of high technical excellence.


Srinivas Katkoori and Vamsi Krishna helped in coordinating the publicity for the workshop and earn our thanks for their excellent job. We appreciate the help from the past General Co-Chairs, Nagarajan Ranganathan and Anantha Chandrakhasan, in organizing the conference. Crucial administrative help came from the members of the IEEE Computer Society, in particular, Anne Marie Kelly, Mary-Kate Rada and Maggie Johnson.

Welcome once again and enjoy the workshop program.

Vijaykrishnan Narayanan
The Pennsylvania State University, USA

Mary Jane Irwin
The Pennsylvania State University, USA

Message from the Technical Program Chairs



It is our distinct pleasure to welcome you to the IEEE Computer Society Annual Workshop on VLSI in Orlando, FL.

This Workshop explores emerging trends and novel concepts in the area of VLSI. The theme of the Workshop is System Design for a System-on-Chip Era. System Level Design has been identified as a dominant research theme for the next decade. System Design has been gaining significance and momentum recently due to the emergence of system-on-a-chip designs. New visionary approaches at the system design level are needed to exploit the great opportunities created by the continuous advances in technology and miniaturization of the semiconductor devices.

System design is converging on a paradigm which includes general purpose commodity chips (i.e. processors, memories, DSP) and full custom mixed analogy and digital application specific integrated circuits (ASICs) integrated via programmable gate arrays on custom printed circuit boards or complete silicon boards, System-on-a-Chip. These hardware systems will be driven by custom, real time software that utilizes the latest software design paradigms (i.e. object oriented languages, client-server architecture, browser interfaces) and wireless communications to provide users with unique functionality. To be effective, these systems must be optimized taking into account a variety of constraints including complexity, power consumption, heat dissipation, mechanical packaging, ergonomics, and design effort. Also, future system design methodologies are an important topic at the Workshop.

We are glad to have a number of leading scientists and distinguished speakers on the workshop program, providing an unique opportunity for the attendees to hear the recent research results in this technical area. It is the face to face meetings with each other that attendees will probably value most, which is why we have tried to maintain a schedule permitting such interactions.

We would like to acknowledge the effort and help from the program committee members, and thank the authors and invited speakers for their contributions to an outstanding technical program. We gratefully acknowledge a diligent work of Anne Rawlinson, of the IEEE Computer Society Press, on the workshop proceedings.


It is our sincere hope that each attendee will benefit greatly from participating in this conference, and will find these proceedings to be a valuable source of information for your future work.

Asim Smailagic
Carnegie Mellon University

Robert Brodersen
University of California at Berkeley

Hugo De Man
IMEC, Belgium

Workshop Committees



General Chairs

Vijaykrishnan Narayanan
The Pennsylvania State University, USA

Mary Jane Irwin
The Pennsylvania State University, USA

Technical Program Chairs

Asim Smailagic
Carnegie Mellon University, USA

Robert Brodersen
University of California, Berkeley, USA

Hugo De Man
IMEC, Belgium

Program Committee

Gaetano Borriello, *University of Washington, USA*

Don Bouldin, *University of Tennessee, USA*

Francky Catthoor, *IMEC, Belgium*

Anantha Chandrakasan, *Massachusetts Institute of Technology, USA*

Mike Connors, *Intel Corporation, USA*

Rolf Ernst, *University of Braunschweig, Germany*

Georges Gielen, *Katholieke Universiteit Leuven, Belgium*

Mike Gordon, *Cambridge University, UK*

Rajesh Gupta, *University of California, Irvine, USA*

Sorin Huss, *Darmstadt University of Technology, Germany*

Fritz Prinz, *Stanford University, USA*

Teresa Meng, *Stanford University, USA*

Hidetoshi Onodera, *Kyoto University, Japan*
Jef Van Meerbergen, *Philips, Netherlands*
Pierre Paulin, *SGS-Thomson, Grenoble, France*
Jan Rabaey, *University of California, Berkeley, USA*
Zoran Salcic, *University of Auckland, New Zealand*
Dan Siewiorek, *Carnegie Mellon University, USA*
Mani Srivastava, *UCLA, USA*
Diederik Verkest, *IMEC Belgium*
Vijaykrishnan Narayanan, *The Pennsylvania State University, USA*
Jacob White, *Massachusetts Institute of Technology, USA*
Wayne Wolf, *Princeton University, USA*

Publicity Chairs

Srinivas Katkooi
University of South Florida, Tampa, USA

Vamsi Krishna
HP Labs, USA

Registration Chair

Srinivas Aruru
Intel Corporation, USA

Steering Committee



A. Mukherjee
University of Central Florida, Orlando, USA

D. W. Bouldin
University of Tennessee, Knoxville, USA

N. Ranganathan
University of South Florida, Tampa, USA

P. A. Subramanyam
AT&T Bell Labs, Murray Hill, USA

J. A. B. Fortes
Purdue University, West Lafayette, USA