MICRO – 36 Final Program

Monday, December 1, 2003

5th Workshop on Media and Streaming Processors (MSP) Vipin Chaudhary, *Wayne State University*; Alex Dean, *North Carolina State University*; Jason Fritts, *Washington University*
3rd Workshop on Power-Aware Computer Systems (PACS) Babak Falsafi, *Carnegie Mellon University*; T. N. Vijaykumar, *Purdue University*
Morning Tutorial: Challenges in Embedded Computing Wayne Wolf, *Princeton University*; Rajesh Gupta, *University of California - San Diego*.
Afternoon Tutorial: Open Research Compiler for Itanium Processor Family Roy Ju, *Intel*.

Tuesday, December 2, 2003

2nd Workshop on Application Specific Processors (WASP) Alex Orailoglu, *University of California - San Diego*.
Morning Tutorial: Microarchitecture-Level Power-Performance Simulators: Modeling, Validation, and Impact on Design Zhigang Hu, *IBM T. J. Watson Research Center*; David Brooks, *Harvard University*; Pradip Bose, *IBM T. J. Watson Research Center.*
Morning Tutorial: Network Processors David Sonnier, *Agere Systems*.
Afternoon Tutorial: Architectural Exploration with Liberty David August, *Princeton University*.

Wednesday, December 3, 2003

8:00 – 8:10	Welcome
8:10 – 9:10	Keynote Speaker 1
9:10 – 9:30	Break
9:30 – 11:00	Session 1: Voltage Scaling & Transient Faults
11:00 – 11:30	Break
11:30 – 1:00	Session 2: Cache Design
1:00 – 2:30	Lunch
2:30 – 4:30	Session 3: Power and Energy Efficient Architectures
4:30 – 5:00	Break
5:00 – 6:30	Session 4: Application-Specific Optimization and Analysis
6:45	Dinner Cruise

Thursday, December 4, 2003

8:00 – 9:00	Keynote Speaker 2
9:00 – 9:30	Break
9:30 – 11:00	Session 5: Dynamic Optimization Systems
11:00 – 11:30	Break
11:30 – 1:00	Session 6: Dynamic Program Analysis and Optimization
1:00 – 2:30	Lunch
2:30 – 4:30	Session 7: Branch, Value, and Scheduling Optimization
4:30 – 5:00	Break
5:00 – 7:00	Session 8: Dataflow, Data Parallel, and Clustered Architectures
10:00	Business Meeting

Friday, December 5, 2003

8:00 – 10:00	Session 9: Secure and Network Processors
10:00 – 10:30	Break
10:30 – 12:30	Session 10: Scaling Design

Keynote Speaker 1

Chair: Sanjay Patel, *University of Illinois at Urbana-Champaign*

Microarchitecture on the MOSFET Diet
Kerry Bernstein
IBM T. J. Watson Research Center

Back to Top

Session 1: Voltage Scaling & Transient Faults

Chair: Antonio González, *UPC - Barcelona* and *Intel Labs*

Razor: A Low-Power Pipeline Based on Circuit-Level Timing Speculation
Dan Ernst, Nam Sung Kim, Sanjay Pant, Shidhartha Das, Rajeev Rao, Toan Pham, Conrad Ziesler, David Blaauw, Todd Austin, *University of Michigan*; Krisztian Flautner, *ARM Ltd.*; Trevor Mudge, *University of Michigan.*

VSV: L2-Miss-Driven Variable Supply-Voltage Scaling for Low Power
Hai Li, Chen-Yong Cher, T. N. Vijaykumar, Kaushik Roy, *Purdue.*

A Systematic Methodology to Compute the Architectural Vulnerability Factors for a High-

Performance Microprocessor
Shubhendu S. Mukherjee, Christopher T. Weaver, Joel Emer, *Intel Corporation;* Steven K. Reinhardt, *Intel Corporation* and *University of Michigan;* Todd Austin, *University of Michigan.*

Back to Top

Session 2: Cache Design

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

TLC: Transmission Line Caches
Bradford M. Beckmann, David A. Wood, *University of Wisconsin-Madison.*

Distance Associativity for High-Performance Energy-Efficient, Non-Uniform Cache Architectures
Zeshan Chishti, Michael D. Powell, T. N. Vijaykumar, *Purdue.*

Near-Optimal Precharging in High-Performance Nanoscale CMOS Caches
Se-Hyun Yang, Babak Falsafi, *Carnegie Mellon University.*

Back to Top

Session 3: Power and Energy Efficient Architectures

Chair: Todd Austin, *University of Michigan*

Single-ISA Heterogeneous Multi-Core Architectures: The Potential for Processor Power Reduction
Rakesh Kumar, *UCSD;* Keith Farkas, Norman Jouppi, Partha Ranganathan, *HP Labs*; Dean Tullsen, *UCSD.*

Run-time Power Monitoring in High-End Processors: Methodology and Empirical Data
Canturk Isci, Margaret Martonosi, *Princeton University.*

Power-driven Design of Router Microarchitectures in On-chip Networks
Hangsheng Wang, Li-Shiuan Peh, Sharad Malik, *Department of Electrical Engineering - Princeton University.*

Optimum Power/Performance Pipeline Depth
Allan Hartstein, Thomas R. Puzak, *IBM T. J. Watson Research Center.*

Back to Top

Session 4: Application-Specific Optimization and Analysis

Chair: Kemal Ebcioğlu, *IBM T. J. Watson Research Center*

Processor Acceleration through Automated Instruction Set Customization
Nathan Clark, Hongtao Zhong, Scott Mahlke, *University of Michigan.*

The Reconfigurable Streaming Vector Processor (RSVP™)
Silviu Ciricescu, Ray Essick, Brian Lucas, Phil May, Kent Moat, Jim Norris, Michael Schuette, *Motorola Labs;* Ali Saidi, *The Mitre Corporation.*

Scaling and Characterizing Database Workloads: Bridging the Gap between Research and Practice
Richard Hankins, Trung Diep, Murali Annavaram, *Intel;* Brian Hirano, Harald Eri, *Oracle;* Hubert Nueckel, John Shen, *Intel.*

Back to Top

Keynote Speaker 2

Chair: Scott Mahlke, *University of Michigan*

In Memory of Bob Rau
Micheal Schlansker
Hewlett-Packard Laboratories

Back to Top

Session 5: Dynamic Optimization Systems

Chair: Ronny Ronen, *Intel*

Generational Cache Management of Code Traces in Dynamic Optimization Systems
Kim Hazelwood, Michael D. Smith, *Harvard University.*

The Performance of Runtime Data Cache Prefetching in a Dynamic Optimization System
Jiwei Lu, Howard Chen, Rao Fu, Wei-Chung Hsu, Bobbie Othmer, Pen-Chung Yew, *Univ. of Minnesota - Twin Cities;* Dong-Yuan Chen, *Intel.*

IA-32 Execution Layer: a two phase dynamic translator designed to support IA-32 applications on Itanium®-based systems
Leonid Baraz, Tevi Devor, Orna Etzion, Shalom Goldenberg, Alex Skaletsky, Yun Wang, Yigal Zemach, *Intel.*

Back to Top

Session 6: Dynamic Program Analysis and Optimization

Chair: John Shen, *Intel*

LLVA: A Low-level Virtual Instruction Set Architecture
Vikram Adve, Chris Lattner, Michael Brukman, Anand Shukla, Brian Gaeke, *University of Illinois at Urbana-Champaign.*

Comparing Program Phase Detection Techniques
Ashutosh S. Dhodapkar, James E. Smith, *University of Wisconsin-Madison.*

Using Interaction Costs for Microarchitectural Bottleneck Analysis
Brian A. Fields, Rastislav Bodik, *University of California-Berkeley;* Mark D. Hill, *University of Wisconsin-Madison;* Chris J. Newburn, *Intel.*

Back to Top

Session 7: Branch, Value, and Scheduling Optimizations

Chair: Glenn Reinman, *University of California-Los Angeles*

Fast Path-Based Neural Branch Prediction
Daniel A. Jiménez, *Rutgers University.*

Hardware Support for Control Transfers in Code Caches
Ho-Seop Kim, James E. Smith, *University of Wisconsin-Madison.*

Exploiting Value Locality in Physical Register Files
Saisanthosh Balakrishnan, Gurindar S. Sohi, *University of Wisconsin-Madison.*

Macro-op Scheduling: Relaxing Scheduling Loop Constraints
Ilhyun Kim, Mikko H. Lipasti, *University of Wisconsin-Madison.*

Back to Top

Session 8: Dataflow, Data Parallel, and Clustered Architectures

Chair: Matt Farrens, *University of California-Davis*

WaveScalar
Steven Swanson, Ken Michelson, Andrew Schwerin, Mark Oskin, *University of Washington.*

Universal Mechanisms for Data Parallel Architectures
Karthikeyan Sankaralingam, Stephen W. Keckler, William R. Mark, Doug Burger, *University of Texas - Austin.*

Flexible Compiler-Managed L0 Buffers for Clustered VLIW Processors
Enric Gibert, *UPC - Barcelona;* Jesús Sánchez, Antonio González, *UPC - Barcelona* and *Intel Labs - Barcelona.*

Instruction Replication for Clustered Architectures
Alex Aletà, Josep M. Codina, *UPC - Barcelona;* Antonio González, *UPC - Barcelona* and *Intel Labs - Barcelona;* David Kaeli, *Northeastern University.*

Back to Top

Session 9: Secure and Network Processors

Chair: T. N. Vijaykumar, *Purdue*

Efficient Memory Integrity Verification and Encryption for Secure Processor
G. Edward Suh, Dwaine Clarke, Blaise Gassend, *MIT;* Marten van Dijk, *Philips;* Srinivas Devadas, *MIT.*

Fast Secure Processor for Inhibiting Software Piracy and Tampering
Jun Yang, *University of California-Riverside;*Youtao Zhang, *University of Texas-Dallas;* Lan Gao, *University of California-Riverside.*

IPStash: A Power-Efficient Memory Architecture for IP Lookup
Stefanos Kaxiras, *Agere Systems;* Georgios Keramidas, *Department of ECE - Univ. of Patras.*

Design and Implementation of High-Performance Memory Systems for Future Packet Buffers
Jorge Garcia, Jesus Corbal, Llorenç Cerdà, Mateo Valero, *UPC - Barcelona.*

Back to Top

Session 10: Scaling Design

Chair: Mikko Lipasti, *University of Wisconsin-Madison.*

Beating In-Order Stalls with "Flea-Flicker" Two-Pass Pipelining
Ronald D. Barnes, Erik M. Nystrom, John W. Sias, Sanjay J. Patel, Nacho Navarro, Wen-mei W. Hwu, *University of Illinois at Urbana-Champaign.*

Scalable Hardware Memory Disambiguation for High ILP Processors
Simha Sethumadhavan, *Department of Computer Sciences - UT Austin;* Rajagopalan Desikan, *Department of Electrical and Computer Engineering - UT Austin;* Doug Burger, Charles R. Moore, Stephen W. Keckler, *Department of Computer Sciences - UT Austin.*

Reducing Design Complexity of the Load/Store Queue
Il Park, Chong Liang Ooi, T. N. Vijaykumar, *Purdue.*

Checkpoint Processing and Recovery: Towards Scalable Large Instruction Window Processors
Haitham Akkary, Ravi Rajwar, Srikanth T. Srinivasan, *Intel.*