Wen-mei William Hwu

PERSONAL INFORMATION

Office:
Coordinated Science Laboratory
1308 West Main Street,
Urbana, Illinois, 61801-2307
(217) 244-8270
(217) 333-5579 (FAX)
Email: w-hwu@illinois.edu

Home:
2709 Bayhill Drive
Champaign, Illinois, 61822-7988
(217) 359-8984

EDUCATION

Ph.D., Computer Science,1987, University of California, Berkeley
B.S., Electrical Engineering, 1983, National Taiwan University, Taiwan

CURRENT POSITION

Professor and Sanders III Advanced Micro Devices, Inc., Endowed Chair, Electrical and Computer Engineering; Research Professor of Coordinated Science Laboratory, University of Illinois, Urbana-Champaign (UIUC).

Chief Technology Officer and Co-Founder, MulticoreWare, Sunnnyvale, California, St. Louis, Missouri, Champaign, Illinois, Chennai, India, Chang-Chun and Beijing, China.

Chief Scientist, Parallel Computing Institute, University of Illinois at Urbana-Champaign

Board Member, Personify, Inc., Champaign, IL

PROFESSIONAL EXPERIENCE

September 2016 to present
Co-Director (with Jinjun Xiong of IBM) of the IBM-Illinois Center for Cognitive Computing Systems Research, funded by IBM at a total of $8M for five years. The center funds a total of 30+ researchers working on hardware, software, and algorithms for building cognitive computing systems for innovative AI applications.

June 2010 to present
Co-Director (with Mateo Valero) of the PUMPS Summer School in Barcelona jointly offered by UIUC and the Universitat Politècnica de Catalunya. The summer school has been attended by about 100 faculty and graduate students worldwide every year to study the advanced parallel algorithm techniques for manycore computing systems.

June 2008 to present
Principle Investigator of the UIUC CUDA Center of Excellence, funded by NVIDIA at over $2.0 M in cash and equipment. The center has been seeded by the UIUC Institute of Advanced Computing Applications and Technology at $400,000. The Center offers a collaborative environment for faculty and staff from Physics, Bioengineering, Chemistry, Astronomy, ECE, CS, Mechanical Engineering, Material Science, Neural Science, Aeronautical Engineering, and NCSA to apply the new inexpensive, massively parallel GPU computing devices to accelerate breakthroughs in critical science and engineering problems.

March 2008 to 2011

Co-Director (with Marc Snir) of the UIUC Universal Parallel Computing Research Center (UPCRC), funded by Intel/Microsoft at $2M/year. The Center is also funded by an $8M match from UIUC. The Center offers a collaborative environment for 16 faculty and 22 graduate students from CS and ECE to work closely with Intel and Microsoft researchers to make parallel programming a ubiquitous practices.

October 2007 to present
Co-PI of the NSF $208M Petascale Leadership Computer Project with $70M matching from the State of Illinois, which will bring the fastest computer for the NSF community to the UIUC campus in 2011. I was responsible for the hardware section of the proposal. As the base hardware specification forms up, I am leading the effort to adapt new accelerator technologies into Blue Waters to extend its capabilities and lifetime.

September 2006 to 2012
Leader of the Concurrent Systems Design Theme of the MARCO FCRP Gigascale Systems Research Center, with 14 faculty from 8 universities (UIUC, UC Berkeley, University of Michigan, Ann Arbor, Princeton, CMU, MIT, U. Penn) and approximate budget of $2M/year for the Theme (www.gigascale.org). The Theme offers a collaborative environment for creating new programming models, frameworks, tools, and architectures to remove the major parallel programming cost for the entire semiconductor and computing industry.

September 2003 to August 2006
Co-lead of the Soft Systems Theme MARCO FCRP Gigascale Systems Research Center, with approximate budget of $1.1M/year for the Theme (www.gigascale.org). The theme offered a collaborative environment for creating programming models and tools for the upcoming acceleration technologies such as many-core chip multiprocessors and FPGA-based reconfigurable computing devices. This Theme was re-organized into the Concurrent Theme in 2006.

August 2003 to present
Sanders III Advanced Micro Devices, Inc., Endowed Chair, ECE
Research Professor of Coordinated Science Laboratory.
University of Illinois, Urbana-Champaign

August 2000 to August 2003
Franklin Woeltge Professor of Electrical and Computer Engineering,
Research Professor of Coordinated Science Laboratory.
University of Illinois, Urbana-Champaign.

August 1996 to July 2000
Professor of Electrical and Computer Engineering,
Research Professor of Coordinated Science Laboratory.
University of Illinois, Urbana-Champaign.

August 1997 to July 1999
Chairman, Computer Engineering Program,
Department of Electrical and Computer Engineering,
University of Illinois, Urbana-Champaign.

August 1992 to July 1996
Associate Professor of Electrical and Computer Engineering,
Research Associate Professor of Coordinated Science Laboratory,
University of Illinois, Urbana-Champaign.

August 1987 to July 1992
Assistant Professor of Electrical and Computer Engineering,
Research Assistant Professor of Coordinated Science Laboratory,
Senior Computer Systems Engineer of Center for Supercomputer Research

and Development. University of Illinois, Urbana-Champaign.

AWARDS AND HONORS

2014 IEEE Computer Society B. R. Rau Award, “for contributions to Instruction Level Parallelism technology, including compiler optimization, program representation, microarchitecture, and applications."

2010 Distinguished Alumni Award, Electrical and Computer Science Department, University of California, Berkeley.

IEEE Computer Society Charles Babbage Award “for visionary leadership & contributions to exploitation of instruction-level parallel processing,” May 28, 2009.

ACM Fellow (elected 2002). “For contributions and technical leadership in the areas of instruction-level parallel processing computer architecture, microarchitecture and compilation.”

Tau Beta Pi Daniel Drucker Eminent Faculty Award (2001). College of Engineering, University of Illinois, Urbana-Champaign.

ACM Grace M. Hopper Award (1999), “For the design and implementation of the IMPACT compiler infrastructure which has been used extensively both by the microprocessor industry as a baseline for product development and by academia as a basis for advanced research and development in computer architecture and compiler design.”

ACM SigArch Maurice Wilkes Award (1998). "For contribution to the creation of the IMPACT compiler and its use to evaluate new architecture features."

IEEE Fellow (elected 1998). "For contributions to high performance compiler and microarchitecture technologies."

1994 University Scholars Award, University of Illinois, "For distinction as a member of the faculty of the University of Illinois."

Teaching:

2014 College of Engineering Collins Award for Innovative Teaching, University of Illinois at Urbana-Champaign. “For the creation of ECE408/CS483, its MOOC version, and its related summer schools worldwide”

2002 ECE Outstanding Teacher Award. “For sustained excellence in the teaching of Compute Engineering courses and in guiding undergraduate students,” University of Illinois, Urbana-Champaign.

1997 Eta Kappa Nu Holmes MacDonald Outstanding Teaching Award.

1997 Pierce Award, College of Engineering, University of Illinois.

Inclusion in the 1992, 1993, 1994, 1996, 1997, 2001 Advisor's List, College of Engineering, University of Illinois.

Inclusion in the Incomplete List of Teachers Ranked as Excellent, University of Illinois,

Spring 2013 (ECE598H), Fall 2012 (ECE408), Fall 2009 (ECE411), Spring 2009 (ECE498AL), Fall 2007(ECE498AL), Spring 2006, Spring 2003, Fall 2002, Spring 2002, Fall 1999, Spring 1999, Fall 1998, Spring 1998, Fall 1997, Fall 1996, Spring 1996, Fall 1995, Fall 1994, Spring 1994, Fall 1993, Spring 1993, Fall 1992, Spring 1992, Fall 1991, Spring 1991, Fall 1990, Spring 1990, Spring 1989, Spring 1988.

Research:

Keynote speaker, "Architecture and Software for Emerging Low-Power Systems," ISLPED, Taipei, July 26 2017

Education/Career Keynote speaker, "What a great time to be a student in computing," SC'2016, Salt Lake City, November 13, 2016.

Keynote speaker, "Innovative Applications and Technology Pivots – A Perfect Storm in Computing" at ICS 2016.

Keynote speaker, "Addressing the Accelerator Programming Challenges in Exascale Systems" at IEEE AsHES 2016.

Distinguished Lecture Series speaker, Electrical and Computer Engineering Department, University of California, Santa Barbara, "What have we learned about programming heterogeneous parallel computing systems?" October 26, 2015.

Distinguished Lecture Series Speaker, Computer Science Department, University of Chicago, "Rethinking Computer Architecture for Energy Limited Computing," January 22, 2015.

2014 MICRO Test-of-Time Award ACM/IEEE International Symposium on Microarchitecture, Yale N. Patt, Wen-mei Hwu, Michael C. Shebanow for their Micro-18 (1985) paper entitled "HPS, a New Microarchitecture: Rationale and Introduction."

2014 MICRO Test-of-Time Award ACM/IEEE International Symposium on Microarchitecture, Yale N. Patt, Wen-mei Hwu, Michael C. Shebanow for their Micro-18 (1985) paper entitled "Critical Issues Regarding HPS, A High Performance Microarchitecture."

Distinguished Lecture Series Speaker, Electrical and Computer Engineering Department, University of California, Riverside, "Scalability, Portability, and Productivity in GPU Computing," March 10, 2014.

NVIDIA CUDA Center of Excellence (CCoE) Achievement Award – Annual Competition among 22 CCoEs worldwide, "For Fighting HIV with CUDA." 2014.

IBM Faculty Award, 2013.

Keynote speaker, "Rethinking Computer Architecture for Throughput Computing," SAMOS 2013 : International Conference on Embedded Computer Systems: Architectures, Modeling, and Simulation (SAMOS XIII), July 15-18, 2013.

NVIDIA CUDA Center of Excellence (CCoE) Achievement Award - Annual Competition among 22 CCoEs worldwide, "For GPU Computing in Blue Waters." 2013

Keynote speaker, "Scalability, Portability, and Numerical Stability in Many-core Parallel Libraries," HiPEAC Conference, Berlin, Germany, January 21-23, 2013.

Keynote speaker, "Heterogeneous Supercomputing in Blue Waters," CyberGIS Champaign, IL 2012.

Keynote speaker, "Language and Compiler Techniques for
Scalable and Portable Applications," International Conference on Compiler Technology for High-Performance Computing (CTHPC), Taipei, 2012

Keynote speaker, “Application Scalability and Portability, key to low power, performance growth, and exascale,” CoolChips Yokohama Japan, April 18, 2012.

Keynote speakers, "Meeting the Multicore Parallel Programming Scalability Challenge", Swidish Multicore Day, 2011

Best Paper Award from FCCM 2011. "Multilevel Granularity Parallelism Synthesis on FPGAs", Papakonstantinou, Alexandros; Liang, Yun; Stratton, John A.; Gururaj, Karthik; Chen, Deming; Hwu, Wen-mei; Cong, Jason. Proceedings of the 2011 International Symposium on Field-Programmable Custom Computing Machines (FCCM).

Keynote speaker, “Raising the Level of Multicore Programming with Compiler Technology, meeting a grand challenge” PACT 2010, Vienna, Austria, September 11, 2010.

MasterWorks speaker, “High-Level Programming Models for Heterogeneous Parallel Computing” at SC2010 (a MasterWorks lecture).

Best Paper Award, “FCUDA – Enabling Efficient Compilation of CUDA Kernels to FPGs,” 2009 IEEE International Symposium on Application-Specific Systems, Architectures, and Processors.

Keynote Speaker, IEEE International Conference on Miltimedia and Expo (ICME), "Parallel Computing Revolution in Video Processing," New York City, July 2, 2009

Plenary Speaker, SIAM Annual Meeting, “The Parallel Revolution in Computational Science and Engineering – applications, education, tools, and impact,” Denver, Colorado, July 7, 2009.

Keynote Speaker, IEEE International Parallel and Distributed Processing Symposium, Rome, Italy, May 2009.

Keynote speaker, HPC Zurich, “Many-core computing, Current Victories and Coming Battles in Application Development,” September 8, 2008. (2008)

Distinguished Lecture Series speaker, Samsung Research, “Many-core GPU Computing – Current Victories and Coming Battles,” Seoul, Korea, June 26, 2008.

Distinguished Lecture Series speaker, Institute of Computing Technology, Chinese Academy of Sciences, “Programming Massively Parallel Processors – The CUDA experience,” June 21-22, 2008.

Keynote speaker, “The Universal Parallelism Vision from Illinois,” Intel Programming Systems Conference, Santa Clara, California, April 22, 2008.

Keynote speaker, “Many-core GPU Computing - Current Victories and Coming Battles in Application Development,” San Jose, California, April 1-2, 2008.

Keynote speaker, “GPU Computing Research at UIUC,” Shanghai Many-Core Workshop, Shanghai, China, March 27-28, 2008.

Distinguished Lecture Series speaker, Ohio Supercomputing Center and the Ohio State University and the Ohio Supercomputing Center, “Accelerating Science and Engineering Applications with GPU Computing,” February, 19, 2008.

Distinguished Lecture Series speaker, Department of Electrical and Computer Engineering, University of Toronto, “GPU Computing – why is it exciting so many application developers, ” November 30, 2007.

Keynote speaker, “A New GPU Computing and Education Environment for Accelerating Science Discoveries and Engineering Breakthroughs,” First GPGPU Workshop, Boston, MA, October 2007.

Keynote speaker, "Top Five Reasons why Sequential Programming Model May Be the Best Way to Program Many-core Microprocessors,” the 39th Annual ACM/IEEE Symposium on Microarchitecture, Orlando, FL, December 2006.

2006 ISCA Influential Paper Award for 1991: “IMPACT: An Architectural Framework for Multiple-Instruction-Issue Processors" by Pohua P. Chang, Scott A. Mahlke, William Y. Chen, Nancy J. Warter, and Wen-mei W. Hwu, in Proceedings of the International Symposium on Computer Architecture, 1991.

IEEE Micro’s Top Picks from the Microarchitecture Conferences in 2005 for “Flea-Flicker Multipass Pipelining: An Alternative to the High-Power Out-of-Order Offense.” By R.D. Barnes, S. Ryoo, and W. W. Hwu.

2002 ComputerWorld Honors Archive Medal, Nominated by Hewlett-Packard, “For Innovation in the Illinois Sandbox e-Services for Campus Life.”

IEEE Fellow (elected 1998). "For contributions to high performance compiler and microarchitecture technologies."

1994 University Scholars Award, University of Illinois, "For distinction as a member of the faculty of the University of Illinois."

1994 Senior Xerox Award for Faculty Research, College of Engineering, University of Illinois, "In recognition of excellence in engineering research during the past five academic years."

Eta Kappa Nu Outstanding Young Electrical Engineer Award for 1993 by the National Jury of Award, "By virtue of his very significant contribution to computer engineering, and for his dedication as an outstanding teacher, advisor, and leader in his profession."
Presidential letter from Bill Clinton, "I am delighted to join your friends, family, and colleagues in congratulating you on winning the 1993 Eta Kappa Nu Outstanding Young Electrical Engineer Award. This award is a fitting treatment to your visionary work in computer engineering. Our nation's continued success requires us to remain at the forefront of innovation in science and technology, and you can take great pride in your contribution endeavor. Your work in this exciting field has already led to new discoveries, and the guidance you have provided to your students has infused the world of computer engineering with the talents of other young achievers..."

Invited Lectures

Distinguished Lecture Series Speaker, Computer Science Department, University of Chicago, “Rethinking Computer Architecture for Energy Limited Computing,” January 22, 2015.

ECE Departmental Colloquium Speaker, University of British Columbia, “Moving Towards Exascale with Lessons Learned from GPU Computing,” October 13, 2014.

ECE Departmental Colloquium Speaker, Columbia University, “Moving Towards Exascale with Lessons Learned from GPU Computing,” October 6, 2014.

ECE Departmental Colloquium, Cornell University, “Scalability, Portability, and Productivity of Heterogeneous Parallel Computing,” April 7, 2014

EECS Departmental Colloquium Speaker, University of Michigan, “Scalability, Portability, and Productivity in GPU Computing,” March 18, 2014

ECE Departmental Colloquium Speaker, Carnegie-Mellon University, “Enabling Unified Code Base for Scalable Heterogeneous Parallel Computing,” October 28, 2013.

EE Departmental Seminar Series, Stanford University, “Why are GPUs so hard to program – or are they?” February 27, 2013

CS Departmental Seminar Series, National Chiao-Tung University, “Application Scalability – Key to Low Power, Performance Growth, and Exascale,” May 21, 2012

W. W. Hwu, “Towards Ultra-efficient Computing Platforms,” ECE Distinguished Lecture Series, University of Texas, Austin, December 5, 2005

W. W. Hwu, “Breaking the Memory Wall for Scalable Microprocessor Platforms,” Keynote, International Conference on Parallel Arhcitectures and Compilers, France, October 1, 2004.

W.W. Hwu, Keynote, “Workload Characterization and Run-time Code Adaptation,” Workshop on Workload Characterization, December 1, 2001.

Keynote speaker, "VLIW: Is it for Real This Time?" the 27th Annual ACM/IEEE Symposium on Microarchitecture, San Jose, CA, 1994.

Keynote speaker, “EPIC Architecture,” All-Spain Parallelism Conference, Murcial, Spain, 1999.

W. W. Hwu, “Ultra-Efficient Computer Architectures, how do we get there?,” Distinguished Lecture Series, Department of Computer Science and Engineering, University of Minnesota, October 20, 2003.

W. W. Hwu, “ROAR: Runtime Optimization Architecture for EPIC Processors,” Distinguished Lecture Series, Department of Electrical Engineering, University of Southern California, April 3, 2002.

W. W. Hwu, “Instruction-Level Parallel Processing – The IMPACT Perspective,” UCLA Jon Postel Lecture Series, Department of Computer Science, March 14, 2002.

Distinguished Lecturer, "EPIC and Beyond, the future of high-end microprocessors," Department of Electrical and Computer Engineering, Purdue University, March 23, 2000.

Courant Institute Series Lecturer, "Predicated Microprocessor Architectures and their Enabling Compiler Technology," NYU, New York City, October 2, 1997.

Intel Associate Professor, Electrical and Computer Engineering, 1992-1993.

Intel Visiting Faculty Award, 1992.

Best paper award for, "Comparing Static and Dynamic Code Scheduling for Multiple-Instruction-Issue Processors," in the Proceedings of the 24th Annual ACM/IEEE International Symposium on Microarchitecture, Albuquerque, New Mexico, November 18-20, 1991.

NSF Research Initiation Award, July 1988.

Best paper award for, "HPSm2: a Refined Single-chip Microengine," presented at the 21st Annual Hawaii International Conference on System Sciences, January 1988.

Best paper award for "An HPS Implementation of VAX; Initial Design and Analysis," presented at The 19th Annual Hawaii International Conference on System Sciences, January 1986.

Service:

1994 ACM Recognition of Service Award, "In appreciation for Contributions to the Association For Computing Machinery - Co-chair ISCA'94." IEEE Computer Society Certificate of Appreciation

1993 IEEE Computer Society Certificate of Appreciation, "For service as Both General and Program Chair for the silver anniversary MICRO conference, making the conference an outstanding success"

IEEE Computer Magazine 1993 list of distinguished reviewers, "for dedication, quality of work, and expertise"

Honor and Professional Societies

Eta Kappa Nu, Institute of Electrical and Electronic Engineers, Association of Computing Machinery

RESEARCH INTERESTS

Architecture, compilation, and microarchitecture of high performance parallel computer systems.

CONSULTING ACTIVITIES

Served as a corporate consultant to Intel, Hewlett-Packard, Advanced Micro Devices, IBM, SUN Microsystems, BOPS, NCR and Hitachi.

Served on the following corporate Boards

Personify, 2009 – present
Founding Director, Series A raised more than $8M
Huawei, 2013-2015
Technical Advisor for the Distributed Computing Lab
ZeroSoft, 2007 – 2010
Technical Advisor, acquired by Synopsis for $24M
Algo-To-Chip, 2009 – 2013
Technical Advisor, Acquired by Nitto Denko for an undisclosed amount
Synfora, 2004-2010
Technical Advisor, Series A-D, raised $22M, acquired by Synopsys
Ageia, 2003-2008
Technical Advisor, Series A raised $3.5M, acquired by NVIDIA for an undisclosed amount
ST Microelectronic, 2006-2010

Served as an expert witness for the following cases:

NAZOMI Communications, Inc. vs. ARM Holdings, PLC, ARM Limited, and ARM Inc., United States District Court, Northern District of California, served as an expert witness retained by ARM (respondent), with declarations, depositions, and testimony, 2002-2006.

EMC vs. Hewlett-Packard, served as an expert witness retained by Hewlett-Packard, case settled before trial, 2003.

Biax Corporation vs. Phillips Semiconductors, United States International Trade Commission, served as an expert witness retained by Phillips (respondent), with declarations, depositions, and testimony, 2006-2007.

ZiiLab vs. Apple Inc., U.S. Patent and Trademark Office, Patent Trial and Appeal Board, served as an expert witness retained by Apple Inc., with a declaration, 2015.

PROFESSIONAL SOCIETIES ACTIVITIES

Program Co-Chair for the 2016 ACM/IEEE International Conference on Parallel Architecture and Compiler Techniques, 2016.

Co-Guest Editor with Sanjay Patel, IEEE MICRO Special Issue on Accelerator Architectures, June/July, 2008.

Program Chair for the 2008 ACM/IEEE International Symposium on Computer Architecture, Beijing, China, June 2008.

Program Chair for the 1st Annual ACM/IEEE International Symposium on Code Generation and optimization, San Francisco, 2003.

Chair, 2002 ACM SigArch Maurice Wilks Award nominations committee. Member 2001-2003.

Program Chair for the 7th Annual IEEE International Symposium on High Performance Computer Architecture (HPCA), Monterrey, Mexico, 2001.

Member of National Jury, 2000 Eta Kappa Nu Holmes MacDonald Outstanding Teaching Award

Chair, 1997 ACM/IEEE Eckert-Mauchly Award committee. Member, 1995-1998.

General Conference Co-Chair, the 21st Annual ACM/IEEE International Symposium on Computer Architecture, Chicago, Illinois, 1994.

General and Program Chair, the 25th Annual ACM/IEEE International Symposium on Microarchitecture, Portland, Oregon, 1992.

Program Chair for the 21st Annual ACM/IEEE International Symposium on Microarchitecture, San Diego, California, 1988.

Served on the program committees for the following years of ACM/IEEE Annual International Symposium on Computer Architecture (ISCA):
42nd, 2015
35th, Beijing, China, 2008
33rd, Boston, 2006
31st, Munich, 2004
28th, Sweden, 2001
25th, Barcelona, Spain, 1998
24th, Denver, Colorado, 1997
22nd, Santa Margherita, Italy, 1995
20th, San Diego, CA, 1993
15th, Hononulu, Hawaii, 1988

Served on the program committees for the following years of ACM/IEEE International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS):

12th, San Jose, 2006

9th, Boston, 2000
8th, Santa Clara, 1998
7th, Boston, 1996.
6th, Santa Clara, 1994.

Served on the program and/or steering committees for the following years of the Annual ACM/IEEE International Symposium on Microarchitecture (MICRO):
48th, Honolulu, HI, 2015
41st, Como, Italy, 2008
40th, Chicago, 2007
39th, Orlando, Florida, 2006
38th, Barcelona, Spain, 2005
37th, Portland, Oregon, 2004
36th, San Diego, 2003
35th, Istanbul, Turkey, 2002
34th, Austin, Texas, 2001
33rd, San Jose, California, 2000
32nd, Haifa, Israel, 1999
31st, Dallas, Texas, 1998
30th, Raleigh, North Carolina, 1997
29th, Paris, France, 1996
28th, Ann Arbor, Michigan, 1995
27th, Santa Clara, California, 1994
26th, Austin, Texas, 1993
25th, Portland, Oregon, 1992
24th, Albuquerque, New Mexico, 1991
21st, San Diego, California, 1988.

Served on the program committee for the following years of the IEEE International Symposium on High Performance Computer Architecture (HPCA), Toulouse, France, 2000, Barcelona, Spain, 2016.

Served on the program committee for the following years of the ACM International Conference on Supercomputing (ICS), Rhode, Greek, 1999, Kos, Greece, 2008

Served on the program committees for the following years of the ACM SIGPLAN Conference on Programming Language Design and Implementation Philadelphia, Pennsylvania, 1996.

Served on the steering committee and/or program committee for the ACM/IEEE Conference on Parallel Architectures and Compiler Technology (PACT): Brosv, Romania, 2007.

Served on the steering committee and/or program committee for the EU Centers of Excellence High-Performance and Embedded Architecture Conference (HiPEAC): Barcelona, Spain 2005, Ghent, Belgiun, 2007, Goteborg, Sweden, 2008.

Guest Editor, IEEE Transactions on Computer Special Issue in Memory of Dr. Bob Rau., October 2005.

Associate Editor, ACM Transactions of Architecture and Code optimization, 2003-present

Associate Editor, IEEE Computer Architecture Letters, 2003-present.

Associate Editor: Journal of Parallel Programming, 1992-2005..

Served as a Minitrack Coordinator for the following years of Hawaii International Conference on System Sciences:
25th, Koloa, Hawaii, 1992,
28th, Maui, Hawaii, 1995.

Served on the following NSF Panels: CAREER, Research Initiation Award, and Small Business Innovation Research Program.

Served as a referee/reviewer in the last two years for papers/books/proposals submitted to:

IEEE Transactions on Computers,
IEEE Transactions on Software Engineering,
IEEE Transactions of Parallel and Distributed Computing,
Journal of Parallel and Distributed Computing, IEEE Computers,
Software - Practice and Experience
International Symposium on Computer Architecture,
Hawaii International Conference on System Sciences,
Annual ACM/IEEE International Symposium on Microarchitectures,
International Conference on Parallel Processing,
Prentice-Hall, McGraw Hill, and National Science Foundation.

Fellow of ACM, Member of EKH. Fellow of IEEE.

UNIVERSITY ADMINISTRATIVE ACTIVITIES

Campus-level

1. MOOC Advisory Committee, 2013-2015.
2. Steering Committee member and Thrust leader, Information Trust Institute, 2006-2008.
3. UIUC Marketing Strategy Committee, 2005-2006
4. University Scholar Selection Committee, University of Illinois, 2002
5. Campus Level Promotion and Tenure Committee, University of Illinois, Urbana-Champaign, 2001-2003
6. Office of Technology Management Advisory Board, University of Illinois, Urbana-Champaign, 2001-present
7. Office of Technology Management Assistant Directory Search Committee, 2002
8. Research Technology Management Office Business Plan Committee, 2000-2001
9. Research Technology Management Office Review Committee, 1999-2000
10. Chair, Research Technology Management Office Technology Transfer Specialist Search Committee, 2000.

Department-level

1. ECE Advisory Committee, 2005-2015
2. CSL Assistant Director Search Committee, 2002
3. ECE Curriculum Committee, 2001-present
4. ECE Faculty Search Committee, 1998-present
5. ECE Long Range Planning Committee, 1999-2001
6. ECE Teaching Evaluation Committee, 1999-present

GRADUATE STUDENTS ADVISED

(a) Ph.D. Degrees Completed: 36 theses (see Ph.D. Theses Supervised Section)

Example of student placements in academia:

Tom Conte: Full Professor, ECE, North Carolina State University
Nancy Warter: Full Professor, California State University, Los Angeles
Scott Mahlke: Assistant Professor, EECS, University of Michigan, Ann Arbor

David August: Assistant Professor, CS, Princeton University
Dan Connors: Assistant professor, ECE, University of Colorado, Boulder
Don Barnes: Assistant Professorm ECE, George Mason University

(b) M.S. Degrees Completed: 70 theses (see M.S. Theses Supervised Section)
(c) Ph.D. Thesis Students Supervised at Present: 7

Sain-Zee Ueng, Chris Rodrigues, Sara Sadeghi (CS), John Stratton, Victor Huang (CS), Deepthi Nandakumar, Xiao-Long Wu

(d) M.S. Thesis Students Supervised at Present: 0

(e) Served on more than 50 other final examination thesis committees.

UNDERGRADUATE RESEARCH STUDENTS ADVISED

John Stratton (with over 25 other previous students)

TEACHING EXPERIENCE

Senior/Junior: Computer Organization and Design, with labs based on hardware description languages
Senior/Junior: Software Engineering Tools for Computer Engineers
Senior/Junior: Microcomputer Laboratory, with labs based on hardware description languages
Senior/Junior: Applied Parallel Programming
Graduate Level: Computer Architecture
Graduate Level: High Performance Computer Architecture and Software
Graduate Level: Parallel Algorithm Techniques for GPU Programming

Courses and Laboratory Developed:

Graduate Level:

Parallel Algorithm Techniques for GPU programming. Created the course and made it a permanent course ECE508. The course covers common algorithm techniques used to regularize data and computation to achieve scalability in parallel applications. The course has also been shared with other institutions such as NCSU, U. Minnesota, U. Tennessee, Georgia Tech, and U. Oklahoma.

High Performance Computer Architecture and Software. Created the course and made it a permanent part of the curriculum. Developed course materials in superscalar microprocessor design, multiprocessor system architecture, multiprocessor interconnect design, multiprocessor memory subsystem design, multiprocessor I/O subsystem design, and compilation for superscalar and multiprocessor systems.

Senior/Junior:

Heterogeneous Parallel Programming. Created the course in partner with Coursera, offered three times in 2013, 2014, and 2015. The course is 8 weeks and covers concepts and techniques programming CPU-GPU heterogeneous parallel computing systems. The main topics include kernel-based parallel programming, hierarchical memory models, performance considerations, important parallel algorithm patterns, and important programming languages. The course is accompanied by a web-based GPU programming environment. It has been taken by more than 70,000 students worldwide.

Programming Massively Parallel Processors. Created the course and made it into a permanent ECE408 course. The course covers concepts and techniques programming CPU-GPU heterogeneous parallel computing systems. The main topics include kernel-based parallel programming, hierarchical memory models, performance considerations, important parallel algorithm patterns, and important programming

languages. Short versions of the course has been offered as VSCSE Summer School, PUMPS Summer School, and other summer schools worldwide.

Created Computer Architecture Laboratory consisting of a network of 50 HP/Apollo, Mentor design and simulation tools, and laboratory projects. Developed laboratory projects on the design of special purpose processors, microprogrammed CPU's, pipelined CPU's, memory management units, and I/O controllers using VHDL. Develop lectures to complement the laboratory projects.

RESEARCH CONTRACTS AND GRANTS

From HP:
(1) Principle Investigator, 2015-2017, $100,000/year, “Programming Models and Systems for Data Centric Computing Architectures.”
(2) Principal Investigator, 2005-2008,$250,000/year, “UIUC Gelato Efforts.”
(3) Co-Principal Investigator with Ravi Iyer and Bill Sanders, 2004-2007, $400,000 “Trustworthy Utility Fabric.”.
(4) Principal Investigator, Equipment Grant, 2002, $350,000, “McKinley Systems for GELATO Software Activities at UIUC/NCSA.”
(5) Principal Investigator, Equipment Grant, 2002, $250,000, “Advanced Operating Structures for the Itanium Processor Family.”
(6) Co-Principal Investigator with Dan Reed (NCSA), 2002-2004, $750,000, “UIUC Gelato Agreement with Hewlett-Packard.”
(7) Principal Investigator, Equipment Grant 2002, $115,260, “Mobile Initiative, Stage II UIUC e-Service
(8) Principal Investigator, Equipment Grant 2001, $500,000, “Mobile Initiative: UIUC Campus Wireless Infrastructure.”
(9) Principal Investigator, 1999-2003, $1,500,000, “Post-Link Optimization Technology for Future Server Architectures.”
(10) Co-Principal Investigator with Thomas Conte (NCSU), 1999-2001, $1,000,000 ($667,000 to UIUC, $333,000 to NCSU), “EPIC Education.”
(11) Principal Investigator, 1998-1999, $260,000, “Real-time Java System Technology.”
(12) Principal Investigator, 1996-1999, $195.000, “Compiler and Debugger Support Debugging Optimized Code."Principal Investigator, 1995, $50,000, “ Instruction-level Parallel Processing Technology.”
(13) Gift, 1992-1998, $450,000, Equipment and software grant through HP’s Research Grants, program in recognition of Hwu's contribution in the areas of compiler optimization and RISC-based architecture.
(14) Principal Investigator, 1992-1995, $300,000, "New Compiler and Architecture Technology for PA and VLIW."

From NVIDIA:
(1) Principal Investigator, 2008-2016, $2,000,000, “UIUC CUDA Center of Excellence”

From Intel:
(1) Co-Principal Investigator (with Marc Snir), 2008-2013, Universal Parallel Computing Research Center (UPCRC), $10M. (Funded by Intel and Microsoft)
(2) Principle Investigator, 2005-present, $100,000/year, “UIUC Gelato Efforts.”
(3) Principal Investigator, 1999-2002, $465,000, “Second-Generation EPIC Compiler Technology.”
(4) Principal Investigator, 1997-2000, $90,000, "Intelligent Cache Hierarchy Management."
(5) Principal Investigator, 1995-1998, $180,000, "Predicated Execution Architectures."
(6) Gift, 1994, $9,928, Equipment gift in recognition of Hwu's contribution in the area of superscalar microprocessors.
(7) Principal Investigator, 1994-1997, $90,000, "Architecture and Compiler Techniques to Improve Performance Through ILP Discovery."
(8) Principal Investigator, 1992-1995, $90,000, "New X86 Architecture Directions and Compilation Technology."

From AMD:
(1) Principal Investigator, 1994-1998 contract, $600,000, "The IMPACT/X86 Compilation and Binary

Reoptimization Technology."
(2) Gift, 1990-2000, $432,000, in recognition of Hwu's contribution in the areas of superscalar microprocessors.

From MARCO/DARPA:
(1) Principle Investigator, 2012-2108, $1,500,000, “UIUC involvement in the C-FAR Research Center.”
(2) Principle Investigator, 2009-2012, $2,600,000, “UIUC involvement in the Gigascale Systems Research Center.”
(3) Principle Investigator, 2006-2009, $800,000, “UIUC involvement in the Center for Circuits, Systems, and Solutions.”
(4) Principle Investigator, 2006-2009, $2,700,000, “UIUC involvement in the Gigascale Systems Research Center.”
(5) Principle Investigator, 2001-2004, $2,500,000, Soft Systems Thrust, Center for Circuits, Systems, and Software.
(6) Principle Investigator, 2006-2009, $600,000, “UIUC involvement in the Center for Circuits, Systems, and Solutions.”

From DARPA:
(1) Co-Principal Investigator (with Andrew Chien, Lautra Caringtn), 2013-2015, $2,000,000, “10x10: Systematic Software-Hardware Heterogeneity for Power-efficient Embedded Computing.”

From NSF:
(1) Co-Principle Investigator (with Bill Kramer and Bill Gropp), 2015-1018, Petascale Application Improvement Discovery (PAID), $2,500,000.
(2) Co-Principal Investigator (with Thom Dunning, Rob Pennington, Marc Snir, Ed Siedel),2008-2016, $208M “Leadership Petascale Computer for the U.S. Science and Engineering Community”
(3) Principle Investigator, 2006-2010, $500,000, NSF Computer Research Infrastructure Grant, “Trusted ILLIAC: A Configurable, Application-Aware, High-Performance Platform for Trustworthy Computing.”
(4) Co-Principal Investigator with S. Adve, D. Padua, L. Kale, S. Patel, Equipment Grant, $120,000, "CISE Research Resources: Programming Environments and Applications for Clusters and Grids."
(5) Co-Principal Investigator, 2000-2003, $1,600,000, “ITR: Experimental Validation of Large-Scale Network Software.”
(6) Principal Investigator, 1998-2000, $158,352, “A New Approach to Accurate and Efficient Pointer Analysis for Large C and Object Oriented Program.”
(7) Principal Investigator, 1996-1998, $109,328, "Data Speculation for Instruction-Level Parallel Processing System.”
(8) Principal Investigator, 1994, $6,000, Student Travel Grant for the 21s International Symposium on Computer Architecture, Chicago, Illinois, 1994.
(9) Principal Investigator, 1993-1995, $163,670, "Predicated and Speculative Execution Support for Instruction Level Parallel Processing."
(10) Co-PI with Yew and Bruner, 1993-1995, $367,599, "Improving the Performance of Scalable Shared-Memory Multiprocessors."
(11) Principal Investigator, 1988-1990, $70,000, "Integrating Compiler Technologies and Parallel Microarchitectures for High Performance Micro System Design."

From Department of Energy
(1) CS Co-Principle Investigator, 2014-2019, $10,000,000 “The Center for Exascale Simulation of Plasma-Coupled Combustion

(2) Co-Principle Investigator (with Jeff Vetter,, Al Malony, Rich Vuduc) , 2010-2013, $2,300,000, “Vancouver: Designing a Next Generation Software Infrastructure for Heterogeneous Exascale Computing.”

From SRC:

(1) Principal Investigator, 2000-2003, $450,000, “Memory Efficient EPIC/VLIW Architecture.”

From AT&T/NCR:
(1) Principal Investigator, 1990-1997, $1,093,624, "The IMPACT Compiler Technology."

From Lucent Technologies:
(1) Principal Investigator, 1999, $25,000, “Computer System Technology for Future High-Performance, Always Available, Versatile Switches.”
(2) Principal Investigator, special purpose grant, 1995, $50,000, "Wireless Information Management."

From SUN Microsystems:
(1) Principal Investigator, 1995, $25,000, "ILP techniques for future SPARC Technology."
(2) Gift, 1993, $150,000, in recognition of Hwu's contribution in the area of architecture research for superscalar , microprocessors.
(3) Principal Investigator, 1993-1994, $60,000, "Design and Analysis of A Superscalar Processor."

Other Grants:
(1) Microsoft, Gift, 1999, $35,000, in recogition of Hwu’s contribution in the area of advanced compiler optimizations for microprocessors
(2) Motorola (Co-PI with Liang Lui, CEE), 1999-, $60,000/year, “Architecture and Compiler Technology for Future Heterogeneous Multiprocessors for Telecommunication”
(3) Motorola (Principal Investigator), 2000-, $60,000/year, “Architecture and Compiler Technology for Future Heterogeneous Multiprocessors for Telecommunication.”
(4) MAZDA (Principal Investigator), 1992-1995, $150,000, "New Computer Architecture Directions for High performance Embedded Applications.”
(5) Matsushita (Principal Investigator), 1991-1993 , $100,000, "Compilation technology for High Performance Parallel Processing Systems."
(6) ONR (Co-Principal Investigator with Prof. K. Fuchs), 1988-1994, $393,168, "Compiler Assisted Recovery for Fault-Tolerant Highly Parallel Multiprocessor Architectures."
(7) JSEP (Faculty Investigator), 1989-1995, $180,000 for research directed by W. W. Hwu, "High-Performance Systems Architecture."
(8) NASA (Faculty Investigator in ICLASS Center), 1987-1997, $4,000,000. ($300,000 for research directed by W. W. Hwu) "Compilation and Architecture for High Performance Parallel Systems."

TECHNOLOGY TRANSFER

OpenIMPACT UIUC/NCSA Open Source License, 2002.

IBM, Research Agreement and Source Code License, the IMPACT Compiler Software, 1998.

Lucent Technologies, Research Agreement and Source Code License, the IMPACT Compiler Software, 1997.

Intel, Research Agreement and Source Code License, the IMPACT Compiler Software, 1995.

AMD, Research Agreement and Source Code License, the IMPACT Compiler Software, 1994.

Hewlett-Packard, Research Agreement and Source License, the IMPACT Compiler Software, 1993.

Mazda, Source License, the IMPACT Compiler Software, 1994.

Matsushita, Source License, the IMPACT Compiler Software, 1994.

AT&T GIS, Source License, the IMPACT Compiler Software, 1993.

SUN Microsystems, Source License, the IMPACT Compiler Software, 1993.

ITRI, ROC, Source License, the IMPACT Compiler Software, 1992.

PRESENTATIONS AND INVITED SEMINARS

(1) W. W. Hwu, invited speaker, "GPU Programming models and tools," NSF/SRC VIA 2020 Workshop, Santa Cruz, July 11, 2008.

(2) W. W. Hwu, Keynote speaker, "The Universal Parallelism Vision from Illinois," Intel Programming Systems Conference, Santa Clara, California, April 22, 2008.

(3) W.W. Hwu, Keynote speaker, "Many-core GPU Computing - Current Victories and Coming Battles in Application Development, San Jose, California, April 1-2, 2008.

(4) W. W. Hwu, Keynote speaker, "GPU Computing Research at UIUC," Shanghai Many-Core Workshop, Shanghai, China, March 27-28, 2008.

(5) W. W. Hwu, "Concurrent Theme Highlights," Plenary Session, The Annual Gigascale Systems Research Center Workshop hosted by IBM T.J. Watson Research Center, NY, March 4, 2008.

(6) W.W. Hwu and Sanjay Patel, Invited Lecture, "Accelerating Science and Engineering Applications with GPU Computing," 2008 Video Analysis and Content Extraction (VACE) Colloquium, Urbana-Champaign, February 26, 2008.

(7) W.W. Hwu, Distinguished Lecture Series speaker, Ohio Supercomputing Center and the Ohio State University and the Ohio Supercomputing Center, "Accelerating Science and Engineering Applications with GPU Computing," February, 19, 2008.

(8) W.W. Hwu, Distinguished Lecture Series speaker, Department of Electrical and Computer Engineering, University of Toronto, "GPU Computing – why is it exciting so many application developers, " November 30, 2007.

(9) W.W. Hwu, Keynote speaker, "A New GPU Computing and Education Environment for Accelerating Science Discoveries and Engineering Breakthroughs," First GPGPU Workshop, Boston, MA, October 2007.

(10) W.W. Hwu, Invited Lecture, "GPU Computing – Programming, Performance, and Scalability," National Research Council Taskforce on Sustaining the Growth of Computer Industry, Spetember 18, 2007.

(11) W.W. Hwu, "Performance Insights of Executing Non-Graphics Applications on the NVIDIA GeForce 8800 and the CUDA Parallel Programming Environment," HotChips 19, Stanford University, August 20, 2007.

(12) W.W. Hwu, "An Implicitly Parallel Compiler Technology Based on Phoenix for Thousand-Core Microprocessors." Microsoft Faculty Summit, Bellevue, Washington, July 17, 2007

(13) W.W. Hwu, "Implicitly Parallel Programming Models for Thousand Core Processors," Design Automation Conference, June 7, 2007.

(14) W.W. Hwu, "Implicitly Parallel Programming Models for Thousand Core Processors," Microsoft Research, Bellevue, WA, May 15, 2007.

(15) W.W. Hwu, "Multi-core Research and the Future of Itanium," Plenary Speech, GELATO ICE Conference, April 16, 2007.

(16) W.W. Hwu, "Concurrency Systems Theme Highlights – Removing Software Development Roadblocks in Massively Parallel Semiconductor Computing Platforms," The Annual Gigascale Systems Research Center Workshop hosted by Intel, Portland, OR, March 14, 2007.

(17) W.W. Hwu, "Top Five Reasons why Sequential Programming Models may be the Best Way to Program Many-core Microprocessors," Keynote, ACM/IEEE 39th Annual Symposium on Microarchitecture, Orland, FL, 2006.

(18) W.W. Hwu, "Top Five Reasons why Sequential Programming Models may be the Best Way to Program Many-core Microprocessors," Invited Lecture, IBM P=AC2 Conference, October 2006.

(19) W.W. Hwu, "Introducing IMPACT Parallelism Discovery and Visualization into Phoenix," Microsoft Faculty Summit, July 17, 2006

(20) W.W. Hwu, Charting the Future of Computing – Trusted ILLIAC and Petascale Computing," EECS College Colloquium, National Taiwan University, Taiwan, June 5, 2006.

(21) W.W. Hwu, Charting the Future of Computing – Trusted ILLIAC and Petascale Computing," Center Colloquium, ITRI, Taiwan, June 2, 2006.

(22) W.W.Hwu, "Hardware Acceleration as a Mainstream Computing Paradigm," EE Departmental Colloquium, National Chao-Tung University, Taiwan, June 1, 2006.

(23) W.W. Hwu, "Non-Traditional Accelerators for Science and Engineering Applications," Institute of Advanced Computing Applications and Technologies, University of Illinois, April 21, 2006.

(24) W.W. Hwu, "Towards Ultra-Efficient Computing Platforms," CSE Departmental Colloquium, Penn State, Januray 12, 2006

(25) W.W. Hwu, "Towards Deep Analysis: Context Sensitive Analysis and Heap Specialization, Intel Santa Clara, December 20 2005

(26) W. W. Hwu, "Towards Deep Program Analysis – Context Sensitive Pointer Analysis and Heap Object Specialization," Intel Multi-core University Research Forum, December 8, 2005

(27) W. W. Hwu, "Towards Ultra-efficient Computing Platforms," ECE Distinguished Lecture Series, University of Texas, Austin, December 5, 2005

(28) W. W. Hwu, "Removing compilation roadblocks for future multiprocessor systems," Wen-mei Hwu, MARCO Live Meeting e-Workshop, September 15, 2005.

(29) W.W. Hwu, "Demolishing Memory Wall for Future Scalable Computing Platforms," IBM T.J. Watson Research Center Seminar on Computer Architecture, March 10, 2005.

(30) W.W. Hwu, "Breaking Down the Memory Wall for Future Scalable Computing Platforms" ACM SigMicro On-line Seminar, http://sigmicro-online.org/seminars.html, January 18, 2005.

(31) W.W. Hwu, "Ultra-efficient Computing Platforms: a Grand Challenge," ECE/CREST Distinguished Lecture, Georgia Institute of Technoloy, October 28, 2004.

(32) W. W. Hwu, "Ultra-efficient Computing Platforms: a Grand Challenge," ECE Departmental Colloquium, North Carolina State University, October 11, 2004.

(33) W. W. Hwu, "Breaking the Memory Wall for Scalable Microprocessor Platforms," Keynote, International Conference on Parallel Arhcitectures and Compilers, France, October 1, 2004.

(34) W. W. Hwu, "Ultra-Efficient Computing Platforms" Infineon, Munich, Germany, June 23, 2004.

(35) W. W. Hwu, “Ultra-Efficient Computer Architectures, how do we get there?” Cornell University School of Electrical and Computer Engineering Colloquium, April 27, 2004.

(36) W. W. Hwu, “Ultra-Efficient Computer Architectures, how do we get there?,” Distinguished Lecture Series, Department of Computer Science and Engineering, University of Minnesota, October 20, 2003.

(37) W.W. Hwu, “OpenIMPACT: Elevating the Performance of IA-64 Linux Platforms,” GELATO Strategy Council Meeting, University of New South Wales, Sydney, Australia, December 6, 2002.

(38) W.W. Hwu, “Open IMPACT,” HP Language Lab, Cupertino, CA, November 22, 2002.

(39) W.W. Hwu, “Addressing Long-Term Software and Architecture Challenges in the MARCO C2S2 Focus Research Center,” IBM T. J. Watson research center, October 14, 2002

(40) W. W. Hwu, “ROAR: Runtime Optimization Architecture for EPIC Processors,” Distinguished Lecture Series, Department of Electrical Engineering, University of Southern California, April 3, 2002.

(41) W. W. Hwu, “Instruction-Level Parallel Processing – The IMPACT Perspective,” UCLA Jon Postel Lecture Series, Department of Computer Science, March 14, 2002.

(42) W.W. Hwu, MARCO C2S2 Focus Research center Annual Review, Carnegie-Mellow University, March 2002.

(43) W.W. Hwu, SRC Annual Research review, “Memory-Efficient EPIC/VLIW Architectures,” Princeton University, February 19, 2002.

(44) W.W. Hwu, “The ROAR Framework for Dynamic Optimization in Future Microprocessors,” Intel, Santa Clara, CA, January 3, 2002.

(45) W.W. Hwu, “Software Customization and Adaptation,” HP Labs, Palo Alto, CA, January 4, 2002.

(46) W.W. Hwu, Keynote, “Workload Characterization and Run-time Code Adaptation,” Workshop on Workload Characterization, December 1, 2001.

(47) W. W. Hwu, "Customization and Adaptation of Future Software Systems," School of EECS Seminar, National Taiwan University, December 24, 2002.

(48) W.W. Hwu, “Composable Memory Reference Analysis for Dynamic Optimization Systems,” HP, Cupertino, CA, December 18, 2001.

(49) W.W. Hwu, “Itanium Performance Insights,” Microprocessor Forum, San Jose, October 16, 2001.

(50) W.W. Hwu, “Itanium Performance Insights from the IMPACT Compiler,” ACM/IEEE HotChips Conference, Palo Alto, August 21, 2001.

(51) W. W. Hwu, “Run-time optimization technology for future microprocessor design,” Transmeta, May, 2000

(52) W. W. Hwu, “Post-Link optimization technology for future applications and architectures,” Hewlett-Packard, January 3, 2000.

(53) W. W. Hwu, “Technology Transfer Practices, Illinois Senate High Technology Task Force Hearings, Beckman Institute, October 14, 1999.

(54) W. W. Hwu, “IMPACT – Past, Present and Future,” IMPACT Partners Meeting, October 8, 1999.

(55) W. W. Hwu, “Microprocessor Architecture and Software Technology for Run-time Optimization: Technical Vision,” Illinois Computer Affiliates Program, October 7, 1999.

(56) W. W. Hwu, “EPIC Architectures and Compiler Technology,” ECE Alumni Board Meeting, University of Illinois, September 12, 1999.

(57) W. W. Hwu, “IMPACT Second Generation EPIC Architecture,” UPC Barcelona, Spain, September 16, 1999.

(58) W. W. Hwu, “EPIC Architectures,” Keynote, All Spain Parallelism Conference, Murcia, Spain, September 12, 1999.

(59) W.W. Hwu, “EPIC Architectures and Enabling Compiler Technology,” IBM T. J. Watsons Research Center, May 14, 1999.

(60) W. W. Hwu, “EPIC Architectures and Enabling Compiler Technology, IBM Toronto Software Laboratory and Center for Advanced Studies, April 22, 1999.

(61) W. W. Hwu, “Beyond EPIC, Semantic Based-Program Optimization,” Computer Science Lecture, Carnegie-Mellon University, April 19, 1999.

(62) W. W. Hwu, “Towards Efficient Exception Detection and Handling in Java,” Illinois Computer Affiliates Program, April 6, 1999.

(63) W. W. Hwu, “EPIC Architectures and Enabling Compiler Technology,” EPFL, Switzerland, January 11, 1999.

(64) W. W. Hwu, “EPIC Architectures and Enabling Compiler Technology,” SUN Microsystems, December 22, 1998.

(65) W. W. Hwu, “EPIC Architecture,” Microprocessor Forum, October 8, 1998.

(66) W. W. Hwu, “EPIC Architectures and Compiler Technology,” September 28, 1998.

(67) W. W. Hwu, "Predicated Microprocessor Architectures and their Enabling Compiler Technology," CASCON Invited Lecture, Toronto, Canada, November 12, 1997.

(68) W. W. Hwu, "On Industry-University Collaborative Research," AMD University Symposium, October 23, 1997.

(69) W. W. Hwu, "Predicated Microprocessor Architectures and their Enabling Compiler Technology,"Courant Institute Lecture Series, NYU, New York City, October 2, 1997.

(70) W. W. Hwu, "Going Beyond 2.0 Effective IPC by Year 1999 with New Compiler and Architecture Techniques," SGI Advanced Technology Lecture Series, Chippewa Falls, WI, August 8, 1997.

(71) W. W. Hwu, "Moving all software into the next level of ILP," Intel Microprocessor Research Forum, Santa Clara, CA, November 1996.

(72) W. W. Hwu, "Some Recent Advances in Static Scheduling," Intel Microprocessor Research Laboratory Santa Clara, CA, August 1996.

(73) W. W. Hwu, "VLIW: Best Road to Even Higher ILP? "Texas Instruments, Dallas, Texas, January, 1996.

(74) W. W. Hwu, "Great Opportunities and Challenges for Microprocessor Architects in the Next Decade," IBM RISC in 95 Symposium, 50 Years of Research: The Science behind the Solutions Yorktown Heights, NY, November 1995.

(75) W. W. Hwu and David August, "Predicated Execution Architectures: New Research Results and Directions," Intel Microprocessor Research Forum, Santa Clara, October 1995.

(76) W.W. Hwu and B. Deitrich, "Data Speculation Architectures: New Research Results and Directions," Intel Microprocessor Research Forum, Santa Clara, October 1995.

(77) W. W. Hwu, "A Comparison of Full and Partial Predicated Execution Support for ILP Processors The 22nd Annual International Symposium on Computer Architecture, Santa Margherita Ligure, Italy, June 1995.

(78) W. W. Hwu, "Instruction-Level Parallel Processing: Sustaining Microprocessor Performance Growth into the Next Millennium, " IBM Yorktown Research Center, NY, March 1995.

(79) W. W. Hwu, "The IMPACT Compiler Technology," Sun Microsystems SPARC Technology Business, Sunnyvale, CA, November 1994.

(80) W. W. Hwu, "VLIW: The next Generation of General Purpose Computers," Microprocessor Forum, October 1994.

(81) W. W. Hwu, "Known Technology vs. Open Research in Instruction-Level Parallel Processing," The First SGI Compiler Summit, San Jose, August 1994.

(82) W. W. Hwu, "Predicated and Speculative Execution," Hewlett-Packard, Palo Alto, California, August, 1994.

(83) W. W. Hwu, "Recent Advances in Predicated Execution," Intel, Santa Clara, California, 1994.

(84) W. W. Hwu, "VLIW Technology," AT&T GIS, Columbia, SC, October 1994.

(85) W. W. Hwu, "Predicated Execution," University of Iowa, Iowa City, IA, October 1994.

(86) W. W. Hwu, "The IMPACT/X86 Superscalar Compiler Technology," Advanced Micro Devices, Austin, Texas, October 1993.

(87) W. W. Hwu, "The IMPACT Project," University of Wisconsin at Madison, October 1993.

(88) W. W. Hwu, "Recent Progress in the IMPACT Compiler Technology," NCR Corporation, Columbia, SC, September 1993.

(89) W. W. Hwu, "Recent Progress in Branch Handling in the IMPACT project," Intel Corporation, Santa Clara, CA, July 1993.

(90) W. W. Hwu, "The IMPACT/SPARC Superscalar Compiler and Architecture Project," SUN Microsystems, Sunnyvale, CA, June 1993.

(91) W. W. Hwu, "Advanced Compiler Technology Lecture Series," A series of eight lectures on the critical compiler technology for high performance computer systems, given at Intel Corporation OR and CA, July - August 1992.

(92) W. W. Hwu, "New Architecture and Compiler Technologies for the HP-PA Architecture," Hewlett-Packard, Palo Alto, CA, June 1992.

(93) W. W. Hwu, "The IMPACT/29K Compiler Technology for Superscalar 29K Implementations," Advanced Micro Devices, Austin, TX, April 1992.

(94) W. W. Hwu, "From IMPACT-I to IMPACT-II, Critical Architecture and Compilation Technology For High Performance Parallel Systems," Intel Corporation, OR, February 1992.

(95) W. W. Hwu, "The IMPACT Compilation Technology for High Performance Parallel Systems," University of Iowa, IA, November 1991.

(96) W. W. Hwu, "The IMPACT Compilation Technology for High Performance Parallel Systems," Carnegie-Mellon University, PA, October 1991.

(97) W. W. Hwu, "IMPACT: An Architectural Framework for Multiple Instruction Issue Processors," Hewlett-Packard Research Laboratories, CA, May 1991.

(98) W. W. Hwu, "IMPACT: An Architectural Framework for Multiple Instruction Issue Processors," Stanford University, Stanford, CA, May 1991.

(99) W. W. Hwu, "IMPACT: An Architectural Framework for Multiple Instruction Issue Processors," Cray Research, Chippewa Falls, WI, March 1991.

(100) W. W. Hwu, "IMPACT C/C++: An Open Foundation for Highly Optimization Compilers," Intel, Santa Clara, CA, January 1991.

(101) W. W. Hwu, "IMPACT C/C++: An Open Foundation for Highly Optimization Compilers," Intel, Portland, OR, November 1990.

(102) W. W. Hwu, "IMPACT C/C++: An Open Foundation for Highly Optimization Compilers," NCR, Columbia, SC, November 1990.

(103) W. W. Hwu, "IMPACT C/C++: An Open Foundation for Highly Optimization Compilers," Hewlett-Packard, Cupertino, CA, May 1990.

(104) W. W. Hwu, "IMPACT C/C++: An Open Foundation for Highly Optimization Compilers," Stanford University, Stanford, CA, May 1990.

(105) W. W. Hwu, "What the coming generation of compiler technology can do for you?" NCR, Dayton, OH, November 1989.

(106) W. W. Hwu, "Compilation Technology and its Impact on the Microarchitecture Tradeoffs," and "The IMPACT Project," IBM T. J. Watson Research Center, Yorktown Heights, NY, August 1989.

(107) W. W. Hwu, "Compilation Technology and its Impact on the Microarchitecture Tradeoffs," and "The IMPACT Project," Digital Equipment Corporation, Marlboro, MA, May 1989.

(108) W. W. Hwu, "Cost-Effective Instruction Caches for Microprocessors Requiring High Instruction Bandwidth," Stanford University, CA, January 1989.

(109) W. W. Hwu, "Cost-Effective Instruction Caches for Microprocessors Requiring High Instruction Bandwidth," NCR, Dayton, OH, December 1988.

(110) W. W. Hwu, "The IMPACT Project," University of Michigan at Ann Arbor, MI, November 1988.

(111) W. W. Hwu, "Exploiting Concurrency to Achieve High Performance in a Single-chip Microarchitecture," Stanford University, CA, April 1987.

(112) W. W. Hwu, "The HPSm Microprocessor Design Project," IBM T. J. Watson Research Center Yorktown Heights, NY, February 1987.

PANELS

Moderator, “Outlook of the Chinese IT Industry,” International Symposium on Computer Architecture, June 15, 2008

Panelist, “Future Course of Computer Microarchitecture,” International Symposium on Microarchitecture, December 3, 2007.

Panelist, “GPU Acceleration of HPC Applications,” Supercomputing 2007, Reno, Nevada, November 16, 2007.

Panelist, “Research in Parallel Programming Tools and Education,” International Workshop on Languages, Compilers for Parallel Computing (LCPC), October 13, 2007

Panelist, “Corezilla: Taming the Multicore Beast,” Design Automation Conference, June 6, 2007.

Panelist, “Computer Technology for Multi-core Architectures,” Intel Multi-core Research Symposium, December 9, 2005.

Panelist, “Future Direction of Computer Architecture Research,” International Symposium on Computer Architecture, June 7, 2005.

Panel moderator, “HPCA Industrial Perspective Panel – new opportunities for computer architecture research,” HPCA, Feb 15, 2005

Panelist, “Where will the Microprocessor Performance Come From in 2010?” ISSCC, February 7, 2000

Panelist, "Microprocessor Architecture for the next Decade," the 28th Annual ACM/IEEE International Symposium on Microarchitecture, Ann Arbor, Michigan, November 1995.

Panelist, "Computer Architecture Research: Blue Sky or Down to the Earth?" the 22nd Annual ACM/IEEE International Symposium on Computer Architecture, Santa Margherita Ligure, June 1995.

INTENSIVE COURSES

W.W. Hwu, “Programming Massively Parallel Processors – the CUDA Experience,” National Center for High-Performance Computing,” Hsingchu, Taiwan, June 30-July 2, 2008.

W.W.Hwu, “Programming Massively Parallel Processors – the CUDA Experience,” Chinese Academy of Sciences, Beijing, China, June 22-23, 2008.

W.W. Hwu, “Compiler Techniques for Multi-core Computing and High-Level Synthesis,” HiPEAC Summer Institute, L’aquila, Italy, July 2006.

W.W. Hwu, “Ultra-efficient Computer Microarchitecture,” Ph.D. Summer Institute, UPC, Barcelon, Spain, June 2004.

W.W. Hwu, “Advanced Topics in Computer Microarchitecture,” Ph.D. Summer Institute, UPC, Barcelon, Spain, June 2002.

W. W. Hwu, “Java Virtual Machine: VM Architecture, Software Architecture, Implementations, and Application Programming Interfaces” NTU Satellite Course from the University of Illinois, March 1999

W. W. Hwu, "Emerging Technology for Microprocessor Architecture, Microarchitecture, and Compilers," Intel Design Center, Haifa, Advanced Engineering Training Course, Israel, January 1996.

Yale Patt and Wen-mei W. Hwu, "Graduate Level Computer Architecture," Digital Equipment Corporation, Advanced Engineering Training Course, November 1985.

TUTORIALS

“Programming Massivley Parallel Processors: the NVIDIA experience,” full-day tutorial with David Kirk and Damir Jamsek, Design Automation Conference (DAC), June 9, 2008.

"Java: VM Architecture, Software Architecture, Implementations, and Applications" the 25 ACM/IEEE International Symposium on Computer Architecture (ISCA), Barcelona, Spain, July 1998

"Compiler and Architecture Support for Java, Conference Tutorial, Seventh International Conference on Architecture Support for Programming Languages and Operating Systems (ASPLOS-VII), Boston, MA, October 1996.

"Advanced Compilation Support for Superscalar and VLIW Processors," Conference Tutorial, 1995 Hawaii International Conference on System Sciences, Maui, HI, January 1995

"New Compiler and Architecture Technologies for Superscalar Processors,"Conference Tutorial, 1992 International Conference on Parallel Processing, St. Charles, IL, August 1992

"New Compiler and Architecture Technologies for Superscalar Processors," Conference Tutorial, 19th International Symposium on Computer Architecture, Queensland, Australia, May 19, 1992

"A New Compiler and Architecture Technology for Superscalar Processors," Advanced Seminars and Tutorials, 25th Hawaii International Conference on System Sciences, Koloa, Hawaii, Jan. 7, 1992.

"Compilation Support for Superscalar Processors," Conference Tutorial, 18th International Symposium on Computer Architecture, Toronto, Canada, May 27, 1991

PUBLICATIONS

BOOKS

(1) D. Kirk and W. Hwu, *Programming Massively Parallel Processors – A Hands-on Approach*, Morgan Kaufmann Publisher, 1st Edition, 2010, ISBN 0123814723.

(2) D. Kirk and W. Hwu, Programming Massively Parallel Processors – A Hands-on Approach, Morgan Kaufmann Publisher, 2nd Edition, 2012, ISBN 0124159923.

(3) D. Kirk and W. Hwu, Programming Massively Parallel Processors – A Hands-on Approach, Morgan Kaufmann Publisher, 3rd Edition, 2016, ISBN 0128119861.

(4) GPU Computing Gems, Emerald Edition, Morgan Kaufmann Publisher, 2011, ISBN 0123849888.

(5) GPU Computing Gems, Jade Edition, Morgan Kaufmann Publisher, 2011, ISBN 0123859638.

(6) Heterogeneous Systems Architecture – A new compute Platform, Morgan Kaufmann Publisher, ISBN

JOURNALS

(1) J. Cabezas, I. Gelado, J. E. Stone, N. Navarro, D. Kirk, W. W. Hwu, "Runtime and Architecture Support for Efficient Data Exchange in Multi-Accelerator Applications", IEEE Transactions on Parallel and Distributed Systems, Issue:99, 2015.

(2) Y. Heo, X-L. Wu, D. Chen, J. Ma, and W.W. Hwu, "BLESS: Bloom-filter-based error correction solution for high-throughput sequencing reads," Bioinformatics, Jan 21, 2014.

(3) C. Atkinson, Geng Liu, N. Obeid, K. R. Thulborn1, W. W. Hwu, "Rapid computation of sodium bioscales using GPU-accelerated image reconstruction," International Journal of Imaging Systems and Technology. 2013; 23: 29-35.

(4) J. Gai, N. Obeid, J.L. Holtrop, X.-L. Wu, F. Lam, M. Fu, J.P. Haldar, W. W. Hwu, Z.-P. Liang, B. P. Sutton, Jiading Gaia, "More IMPATIENT – A Gridding Accelerated Toeplitz-based Strategy for Non-Cartesian, High-Resolution 3D MRI on GPUs," Journal of Parallel and Distributed Computing (2013).

(5) X.-L. Wu, Y. Heo, I. El Hajj, W.W. Hwu, D. Chen, J. Ma, "TIGER: Tiled iterative genome assembler", Journal of BMC Bioinformatics, 2012 .

(6) Huang, Victor; Rodrigues, Christopher I.; Jones, Stephen; Buck, Ian; Hwu, Wen-mei, "Scalable SIMD-parallel memory allocation for many-core machines", The Journal of Supercomputing, 9 Sep 2011.

(7) S. Ryoo, C.I. Ridrigues, S. S. Stone, J. A. Stratton, Z. Ueng, S.S. Baghsorkhi, W. W. Hwu, "Program Optimization Carving for GPU Computing," Journal of Parallel and Distributed Computing (2008), doi:10.1016/j.jpdc.2008.05.011.

(8) S.S. Stone, J.P. Haldar, S. C. Tsao, W.W. Hwu, B.P. Sutton, Z.P. Liang, "Accelerating Advanced MRI Reconstruction on GPUs," Journal of Parallel and Distributed Computing, (2008), doi:10.1016/j.jpdc.2008.05.013.

(9) S. Ryoo, S. Ueng, C. I. Rodrigues, R. E. Kidd, M. I. Frank, W. W. Hwu, "Automatic Discovery of Coarse-Grained Parallelism in Media Applications," Special Issue of the HiPEAC Transactions on High-Performance Embedded Architecture and Compilation, 2007.

(10) R.D. Barnes, J.W. Sias, E.M. Nystrom, S.J.Patel, N. Navarro, W.W. Hwu, "Beating In-order Stalls with Flea-Flicker Two-pass Pipelining," IEEE Transactions on Computers, February 2006.

(11) J. Monks, J.-P. Ebert, W. W. Hwu, and A. Wolisz, "Energy Saving and Capacity Improvement Potential of Power Control in Multi-hop Wireless Networks," Journal of Computer Networks, 2003.

(12) W. W. Hwu, D. I. August and J. W. Sias "Program Decision Logic Optimization Using Predication and Control Speculation," Proceedings of the IEEE, November, 2001, pp. 1660-1675.

(13) M. C. Merten, A. R. Trick, R. D. Barnes, E. M. Nystrom, C. N. George, J. C. Gyllenhaal, and W. W. Hwu, "An Architectural Framework for Run-Time Optimization," IEEE Transactions on Computers, Vol. 50, No. 6, June 2001, pp. 567-589.

(14) T.L. Johnson, D.A. Connors, M.C. Merten, W.W. Hwu, "Run-time Cache Bypassing," IEEE Transactions on Computer, Vol. 48, No. 12, December 1999, pp. 1338-1354.

(15) D.I. August, W.W. Hwu, and S.A. Mahlke, “The Partial Reverse If-Conversion Framework for Balancing Control Flow and Predication,” International Journal of Parallel Programming, Vol. 27, No. 5, October 1999, pp. 381-423.

(16) J. C. Gyllenhaal, W. W. Hwu, and B. R. Rau, "Optimization of Machine Descriptions for Efficient Use," International Journal of Parallel Programming, Special issue on the most significant papers from the 29th ACM/IEEE International Symposium on Microarchitecture, Vol. 26, No. 4, August 1998.

(17) R.E. Hank, W.W. Hwu, and B.R. Rau, "Region-Based Compilation: introduction, motivation, and Initial experience," International Journal of Parallel Programming, Special issue on the most significant papers from 28the ACM/IEEE International Symposium on Microarchitecture, Vol. 25, No. 2, April 1997, pp. 113-146.

(18) T. M. Conte, M. A Hirsch, and W. W. Hwu, "Combining trace sampling with single pass methods for efficient cache simulation," IEEE Transactions on Computers, February 1997. W. W. Hwu, et al, "Compiler Technology for Future Microprocessors," IEEE Proceedings, Vol. 83, No. 12, December 1995.

(19) P. P. Chang, D. M. Lavery, S. A. Mahlke, W. Y. Chen and W. W. Hwu, "The Importance of Prepass Code Scheduling for Superscalar and Superpipelined Processors," IEEE Transactions on Computers, Vol. 44, No. 3, March 1995, pp. 353-370.

(20) P. P. Chang, N. J. Warter, S. A. Mahlke, W. Y. Chen, W. W. Hwu, "Three Architectural Models for Compiler-controlled Speculative Execution," IEEE Transactions on Computers, Vol. 44, No. 4 April 1995, pp. 481-494.

(21) W. W. Hwu and T. M. Conte, “The Susceptibility of Programs to Context Switching Effects,” IEEE Transactions on Computers, Vol. 43, No. 9, September 1994.

(22) S. K. Chen, N. J. Alewine, W. K. Fuchs, W. W. Hwu, "Incremental Compiler Transformations for Multiple Instruction Retry," Software, Practice & Experience, John Wiley & Sons Ltd., Vol. 24(9), pp. 1-20, September 1994. (JSEP/NASA)

(23) C. C. Li, S. K. Chen, W. K. Fuchs, and W. W. Hwu, "Compiler-Assisted Multiple Retry," IEEE Transactions on Computers, Vol. 44, No.1, January 1995.

(24) S. Anik and W. W. Hwu, "Performance Implications of Synchronization Support for Parallel FORTRAN Programs," Journal of Parallel and Distributed Computing, 22, pp. 202-215, 1994.

(25) S. A. Mahlke, W. Y. Chen, R. Bringmann, R. Hank, W. W. Hwu, M. Schlansker and B. Rau "Sentinel Scheduling: A Model for Compiler-Controlled Speculative Execution," ACM Transactions on Computer Systems, Vol. 11, No. 4, November 1993, pp. 376-408.

(26) W. Y. Chen, P. P. Chang, and W. W. Hwu, "The Effect of Code Expanding Optimizations on Instruction Cache Design," IEEE Transactions on Computers, Vol. 42, No. 9, September 1993, pp. 1045-1057.

(27) W. Y. Chen, S. A. Mahlke, N. J. Warter, S. Anik and W. W. Hwu, "Profile-Assisted Instruction Scheduling," International Journal for Parallel Programming, Vol. 22, No. 2, April 1994, pp. 151-181.

(28) N. Alewine, S. K. Chen, W. K. Fuchs, W. W. Hwu, "Compiler-Assisted Multiple Instruction Rollback Recovery Using a Read Buffer," IEEE Transactions on Computers, Vol. 44, No. 9, September 1995.

(29) A. Gupta and W. W. Hwu, "XPROF: An Execution Profiler for Window-Oriented Applications," Software, Practice & Experience, John Wiley & Sons Ltd., Vol. 23(5), pp. 487-510, May 1993.

(30) W. W. Hwu, et al, "The Superblock: An Effective Technique for VLIW and Superscalar Compilation" the Journal of Supercomputing, Kluwer Academic Publishers, 7, 1993, pp. 229-248.

(31) W. W. Hwu and P. P. Chang, "Efficient Instruction Sequencing with Inline Target Insertion," IEEE Transactions on Computers, Vol. 41, No. 12, December 1992, pp. 1537-1551.

(32) P. P. Chang, S. A. Mahlke, W. Y. Chen, and W. W. Hwu, "Profile-guided Automatic Inline Expansion For C Programs," Software Practice and Experience, John Wiley & Sons Ltd., Vol. 22(5), May 1992, pp. 349-369.

(33) G. A. Uvieghara, W. W. Hwu, Y. Nakagome, D-K Jeong, D. D. Lee, D. A. Hodges, Y. N. Patt, "An Experimental Single-chip Data Flow CPU," IEEE Journal of Solid-State Circuits, Vol. 27, No. 1, Jan. 1992, pp. 17-28.

(34) P. P. Chang, S. A. Mahlke, and W. W. Hwu, "Using Profile Information to Assist Classic Compiler Code Optimizations," Software Practice and Experience, John Wiley and Sons Ltd., Vol. 21, No. 12 Dec. 1991, pp. 1301-1321.

(35) W. W. Hwu and Y.N. Patt, "Checkpoint Repair for High Performance Out-of-order Execution Machines", IEEE Transactions on Computers, Vol. C-36, Dec. 1987, pp. 1496-1514.

PATENTS

(1) W.W. Hwu and M. Merten, “Method and Apparatus for Modulo Scheduled Loop Execution in a Processor Architectures,” U.S. Patent #7, 302, 557., November 27, 2007.

(2) T. Kiyohara, W. Y. Chen, and W. W. Hwu, "Memory Conflict Buffer for Achieving Memory Disambiguation," U.S. Patent #5,694,577, issued December 2, 1997.

(3) L.-C. Wu and W.W. Hwu, “Debugging of Optimized Code,” U.S. Patent issued.

(4) W. W. Hwu, et al, “Method and Apparatus for Instruction Execution Hot Spot Detection and Monitoring in a Data Processing Unit,” U.S. Patent #6,681,387, issued January 28, 2004.

(5) W. W. Hwu, , et al “ Method and Apparatus for Enhancing Instruction Level Parallelism,” U.S. Patent #6,640,315, issued October 28, 2003.

MAGAZINE ARTICLES

(1) J.A. Stratton, C. Rodrigues, I. R. Sung, L. Chang, N. Anssari, G. D. Liu, W. W. Hwu, and N. Obeid, “Algorithm and Data Optimization Techniques for Scaling to Massively Threaded Systems,” IEEE Computer, August 2012, pp. 26-32.

(2) Lin, Dennis; Huang, Victor; Nguyen, Quang; Blackburn, Joshua; Rodrigues, Christopher I.; Huang, Thomas; Do, Minh N.; Patel, sanjay J.; Hwu, Wen-mei, "The parallelization of video processing", IEEE Signal Processing Magazine 26(6), 103--112, 2009.

(3) S. Patel and W.W. Hwu, “Accelerator Architectures,” Guest Editors’ Introduction, IEEE Micro, July/August 2008, pp. 4-12.

(4) W. W. Hwu, K. Keutzer, T. Mattson, “The Concurrency Challenge,” IEEE Design and Test of Computers, July/August 2008, pp. 312-320.

(5) R. K. Iyer, Z. Kalbarczyk, K. Pattabiraman, W. Healey, W. W. Hwu, P. Klemerer, R. Farivar, “Toward Application-Aware Security and Reliability,” IEEE Security and Privacy, January/February, 2007.

(6) R. D. Barnes, S. Ryoo, and W. W. Hwu, “Tolerant Cache-Miss Latency with Multipass Pipelines,” Special Issue on Top Picks from Microarchitecture Conferences, IEEE Micro, Volume 26, No. 1,January/February 2006.

(7) W. W. Hwu, "Introduction to Predicated Execution," IEEE Computer, January 1998, pp. 51-52.

(8) C. A. Hsieh, M. T. Conte, T. L. Johnson, J. C. Gyllenhaal, and W. W. Hwu, "Using the NET to Capture Performance in Java-Based Software," IEEE Computer, June 1997, pp. 67-75.

(9) T. M. Conte and W. W. Hwu, "Benchmark Characterization," IEEE Computer, Jan. 1991, pp. 48-56.

RIGOROUSLY REFEREED CONFERENCE PROCEEDINGS

(1) L.-W. Chang, I. El Hajj, C. Rodrigues, J. Gomez-Luna, W.W. Hwu, “Efficient kernel Synthesis for Performance Portable Programming,” Proceedings of the ACM/IEEE 49th International Symposium on Microarchitecture (MICRO’16), October 2016.

(2) I. El Hajj, J. Gomez-Luna, C. Li, L.-W. Chang, D. Milojicic, W.W. Hwu, “KLAP: Kernel Launch Aggregation and Promotion for Optimizing Dynamic Parallelism,” Proceedings of the ACM/IEEE 49th International Symposium on Microarchitecture (MICRO’16), October 2016.

(3) I. El Hajj, A.Merritt, G. Zellweger, D. Milojicic, R. Achermann1, P. Faraboschi, W. Hwu, T. Roscoe, and K. Schwan, “SpaceJMP: Programming with Multiple Virtual Address Spaces,” Proceedings of the ACM/IEEE 21st International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'16), April 2016.

(4) L.-W. Chang, H.-S. Kim, W. W. Hwu, “DySel: Lightweight Dynamic Selection for Kernel-based Data-parallel Programming Model,” Proceedings of the ACM/IEEE 21st International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'16), April 2016.

(5) L. Chang, I. El Hajj, H. S. Kim, J. Gómez-Luna, A. Dakkak, W. W. Hwu, “A Programming System for Future Proofing Performance Critical Libraries,” Proceedings of the 2016 ACM SIGPLAN Conference on Principles and Practice of Parallel Programing (PPoPP), March 2016.

(6) J. Cabezas, L. Vilanova, I. Gelado, T. Jablin, N. Navarro, W. W. Hwu, "Automatic Parallelization of Kernels in Shared-Memory Multi-GPU Nodes", Proceedings of the 29th ACM on International Conference on Supercomputing (ICS '15), June 2015.

(7) H.-S. Kim, I. E. Hajj, J. A. Stratton, S. S Lumetta, W.W. Hwu, "Locality-Centric Thread Scheduling for Bulk-synchronous Programming Models on CPU Architectures", International Symposium on Code Generation and Optimization (CGO), February 2015.

(8) X. Chen, L-W. Chang, C. I. Rodrigues, J. Lv, Z. Wang, W.W. Hwu, "Adaptive Cache Management for Energy-efficient GPU Computing", Proceedings of the 47th Annual IEEE/ACM International Symposium on Microarchitecture, December 2014 .

(9) J. Cabezas, L. Vilanova, I. Gelado, T. Jablin, N. Navarro, W. W. Hwu, "Automatic execution of single-GPU computations across multiple GPUs", Proceedings of the 23rd international conference on Parallel Architectures and Compilation Techniques (PACT), 2014

(10) C. I. Rodrigues,. A. Dakkak, T. Jablin, and W.W. Hwu, "Triolet: A Programming System that Unifies Algorithmic Skeleton Interfaces for High-Performance Cluster Computing", Proceedings of the 2014 ACM SIGPLAN Conference on Principles and Practice of Parallel Programing, February 2014.

(11) I. R. Sung, J. Gómez-Luna, J. M. González-Linares, N. Guil, W. W. Hwu, "In-place transposition of rectangular matrices on accelerators", PPoPP '14 Proceedings of the 19th ACM SIGPLAN symposium on Principles and practice of parallel programming.

(12) A. Papakonstantinou, D. Chen, W.W. Hwu, Wen-mei; Cong, Jason; Liang, Yun, "Throughput-Oriented Kernel Porting onto FPGAs", Proceedings of the 50th Annual Design Automation Conference, May 2013.

(13) L. Chang, J.A. Stratton, H. Kim, and W.W. Hwu, "A Scalable, Numerically Stable Tridiagonal Solver Using GPUs," The International Conference for High-Performance Computing Networking, Storage, and Analysis (SC'12), Salt Lake City, 2012.

(14) I.-J. Sung, G. D. Liu, and W. W. Hwu, "DL: A Data Layout Transformation System for Heterogeneous Computing," The IEEE Innovative Parallel Computing Conference – Foundations and Applications of GPU, Manycore, and Heterogeneous Systems, San Jose, May, 2012.

(15) J. A. Stratton, N. Anssari, C. I. Rodrigues, I. Sung, N. Obeid, L. Chang, G. Liu, and W. Hwu, "Optimization and Architecture Effects on GPU Computing Workload Performance," The IEEE Innovative Parallel Computing – Foundations and Applications, San Jose, May, 2012.

(16) S. S. Baghsorkhi, I. Gelado, M. Delahaye, W. W. Hwu, "Efficient Performance Evaluation of Memory Hierarchy for Highly Multithreaded Graphics Processors," Proceedings of the 17th ACM SIGPLAN symposium on Principles and Practice of Parallel Programming, February, 2012

(17) L. Luo, M. Wong, and W. *Hwu*. An effective gpu implementation of *breath-first* search. ACM/IEEE *Design Automation Conference (DAC)*, 2010.

(18) I. Sung, J. Stratton, W.W. Hwu, "Data Layout Transformation Exploiting Memory-Level Parallelism in Structured Grid Many-core Applications," Proceedings of the 19th international conference on Parallel architectures and compilation techniques, 2010, pp. 513-522.

(19) I. Gelado. J.E. Stone. J. Cabezas, S. Patel, N. Navarro and W.W. Hwu, "An Asymmetric Distributed Shared Memory Model for Heterogeneous Parallel Systems," The ACM/IEEE 15th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'10). March 2010. Pittsburgh, PA.

(20) Stratton, John A.; Grover, Vinod; Marathe, Jaydeep; Aarts, Baastian; Murphy, Mike; Hu, Ziang; Hwu, Wen-mei, "Efficient Compilation of Fine-grained SPMD-threaded Programs for Multicore CPUs", Proceedings of the International Symposium on Code Generation and Optimization, April 2010 .

(21) Baghsorkhi, Sara Sadeghi; Delahaye, Matthieu; Patel, sanjay J.; Gropp, William D.; Hwu, Wen-mei, "An Adaptive Performance Modeling Tool for GPU Architectures", Proceedings of the 15th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP), Jan. 2010.

(22) Pattabiraman, Karthik; Gururaj, Karthik; Stratton, John A.; Chen, DRyooeming; Cong, Jason; Hwu, Wen-mei, (Best Paper Award) Alexander Papakonstantinou, Karthik Gururaj, John A. Stratton, Deming Chen, Jason Cong and Wen-mei W. Hwu, "FCUDA: Enabling Efficient Compilation of CUDA Kernels onto FPGAs", Symposium on Application Specific Processors, July 2009. (Best Paper Award)

(23) Hwu, Wen-mei; Rodrigues, Christopher I.; Ryoo, Shane; Stratton, John A., "Compute Unified Device Architecture Application Suitability", Computing in Science and Engineering Vol. 11 No. 3, May 2009.

(24) Isaac Gelado, John H. Kelm, Shane Ryoo, Nacho Navarro, Steve S. Lumetta, and Wen-mei W. Hwu, "CUBA: An Architecture for Efficient CPU/Co-processor Data Communication," Proceedings of the 22nd ACM International Conference on Supercomputing, June 2008.

(25) C. I. Rodrigues, J. Stone, D. Hardy, W. W. Hwu, "GPU Acceleration of Cutoff-Based Potential Summation," ACM Computing Frontier Conference 2008, Italy, May 2008, pp. 273-282.

(26) S.S. Stone, J.P. Haldar, S. C. Tsao, W.W. Hwu, Z.P. Liang, B.P. Sutton, "Accelerating Advanced MRI Reconstruction using GPUs," ACM Computing Frontier Conference 2008, Italy, May 2008.

(27) S. Ryoo, C. Rodrigues, S. Stone, S. Baghsorkhi, S. Ueng, J. Stratton, W. Hwu, "Program Optimization Space Pruning for a Multithreaded GPU,"Proceedings of the 6th ACM/IEEE International Symposium on Code Generation and Optimization, April 6-9, 2008.

(28) S. Ryoo, C. I. Rodrigues, S. S. Baghsorkhi, S. S. Stone, D. B. Kirk, and W. W. Hwu , "Optimization Principles and Application Performance Evaluation of a Multithreaded GPU Using CUDA," Proceedings of the 13th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, February 2008.

(29) J. Kelm, I. Gelado, M. Murphy, N. Navarro, S. Lumetta, W.W. Hwu, "CIGAR: Application Partitioning for a CPU/Coprocessor Architecture," 9th International Conference on Parallel Architecture and Compiler Techniques (PACT), Brasov, Romania, September 2007.

(30) W.W. Hwu, et al, "Implicitly Parallel Programming Models for Thousand-Core Microprocessors," Proceedings of the 2007 Design Automation Conference, San Diego, California, June, 2007.

(31) Flea-Flicker Multipass Pipelining: An Alternative to the High-Power Out-of-Order Offense, R. Barnes, S. Ryoo, and W.-M. Hwu, Proceedings of the 2005 International Symposium on Microarchitecture (MICRO-05), pp 319-330, November, 2005

(32) Bottom-up and Top-down Context-Sensitive Summary-based Pointer E. M. Nystrom, H. -S. Kim, W. -M. Hwu. The 11th Static Analysis Symposium, Florence, Italy, August 2004.

(33) John W. Sias, Sain-zee Ueng, Geoff A. Kent, Ian M. Steiner, Erik M. Nystrom and Wen-mei W. Hwu, Field-testing IMPACT EPIC Research Results in Itanium 2, Proceedings of the 31st Annual International Symposium on Computer Architecture, Munich, Germany, July 2004.

(34) R. D. Barnes, E. M. Nystrom, J. W. Sias, S. J. Patel, N. Navarro, and W. W. Hwu, "Beating In-order Stalls with Flick-Fleaker Two-Pass Pipelining," Proceedings of the 36th ACM/IEEE Annual International Symposium on Microarchitecture, December 2003.

(35) R. D. Barnes, E.M. Nystrom, M.C. Merten, W.W. Hwu, "Vacuum Packing: Extracting Hardware-Detected Program Phases for Post-Link Optimization," Proceedings of the 35th ACM/IEEE Annual International Symposium on Microarchitecture, November 2002.

(36) J. W. Sias, H. C. Hunter and W. W. Hwu, “Enhancing Loop Buffering of Media and Telecommunications Applications Using Low-overhead Predication,” Proceedings of the 34th International Symposium on Microrchitecture, December, 2001.

(37) M. C. Merten and W. W. Hwu **,** “Modulo Schedule Buffers,” Proceedings of the 34th International Symposium on Microrchitecture, December, 2001

(38) E. M. Nystrom, R. D. Barnes, M. C. Merten, and W. W. Hwu **,** “Code Reordering and Speculation Support for Dynamic Optimization Systems,” Proceedings of the International Conference on Parallel Architectures and Compilation Techniques, September 8-12, 2001.

(39) J. P. Monks, V. Bharghavan, and W. W. Hwu, "Transmission Power Controlled for Multiple Access Wireless Packet Networks," *Proceedings of The 25th Annual IEEE Conference on Local Computer Networks (LCN 2000)*, Tampa, FL, Nov., 2000.

(40) J. P. Monks, V. Bharghavan, and Wen-mei Hwu, "A Power Controlled Multiple Access Protocol for Wireless Packet Networks," *IEEE INFOCOM 2001*, Anchorage, Alaska, April 22-26, 2001.

(41) Sias, J.W., August, D.I., Hwu, W.W., “Accurate and Efficient Predicate Analysis with Binary Decision Diagram,” Proceedings of the 33rd International Symposium on Microarchitecture, December 2000.

(42) Connors, D.A., Hunter, H.C., Cheng, B.C., and Hwu, W.W., “Hardware Support for Dynamic Activation of Compiler-Directed Computation Reuse,” Proceedings of the 9th International Conference on Architecture Support for Programming Languages and Operating Systems, November 2000.

(43) Cheng, B. and Hwu, W. W., “Modular Interprocedural Pointer Analysis Using Access Paths: Design, Implementation, and Evaluation", Proceedings of the ACM SIGPLAN '00 Conference on Programming Language Design and Implementation, June 2000.

(44) M. C. Merten, A. R. Trick, E. M. Nystrom, R. D. Barnes, and W. W. Hwu, "A Hardware Mechanism for Dynamic Extraction and Relayout of Program Hot Spots," *International Symposium on Computer Architecture*, June 2000.

(45) D. Connors and W.W. Hwu, “Compiler-Directed Computation Reuse: Rationale and Initial Results,” Proceedings of the 32nd ACM/IEEE Annual International Symposium on Microarchitecture , Haifa, Israel, November 1999, pp. 158-169 (acceptance rate 21%)

(46) L.-C. Wu, R. Mirani, H. Patil, B. Olsen, W. Hwu, "A New Framework for Debugging Optimized Code," Proceedings of the ACM SIGPLAN '99 Conference on Programming Language Design and Implementation, Atlanta, GA, May 1-4, 1999, pp.181-191. (acceptance rate 20%)

(47) D.I. August, J.W. Sias, J.M. Puiatti, K.M. Crozier, W.W. Hwu, "The Program Decision Logic Approach to Utilizing Predicated Execution," Proceedings of the 26th International Symposium on Computer Architecture, Atlanta, GA, May 2-4, 1999, pp. 208-219. (acceptance rate19%)

(48) M. Merten, A. Trick, C. George, J. Gyllenhaal, W. W. Hwu, “A Hardware-Driven Profiling Scheme for Identifying Program Hot Spots to Support Runtime Optimization,” Proceedings of the 26th International Symposium on Computer Architecture, Atlanta, GA, May 2-4, 1999, pp. 136-147. (acceptance rate 19%)

(49) B. -C. Cheng, D. Connors, W. W. Hwu, "Compiler-Directed Early Load-Address Generation," Proceedings of the 31st ACM/IEEE Annual International Symposium on Microarchitecture, December, 1998, Dallas, TX, pp. 138-147. (acceptance rate 26%)

(50) D.I. August, D. Connors, S. A. Mahlke, J. Sias, K. Crozier, B. Cheng, P. Eaton, Q. Olaniran, W. W. Hwu, "Integrated Predicated and Speculative Execution in the IMPACT EPIC Architecture," Proceedings of the 25th Annual International Symposium on Computer Architecture (ISCA), June 27 - July 1, 1998, Barcelona, Spain, pp. 227-237. (acceptance rate 21%)

(51) T.L. Johnson, M.C. Merten, W.W. Hwu, "Run-time Spatial Locality Detection and Optimization," Proceedings of the 30th Annual International Symposium on Microarchitecture, December 1-3, 1997, Research Triangle, NC, pp. 57-64. (acceptance rate 34%)

(52) D.I. August, W. W. Hwu, S.A. Mahlke, "A Framework for Balancing Control Flow and Predication," Proceedings of the 30th Annual International Symposium on Microarchitecture December 1-3, 1997, Research Triangle, NC, pp. 92-103. (acceptance rate 34%)

(53) T.L. Johnson and W. W. Hwu, "Run-time Adaptive Cache Hierarchy Management via Reference Analysis," Proceedings of the 24th Annual International Symposium on Computer Architecture, (ISCA) June 2-4, 1997, Denver, Colorado, pp. 315-326. (acceptance rate 21%)

(54) D.I. August, D.A. Connors, and W.W. Hwu, "Architectural Support for Compiler-synthesized Dynamic Branch Prediction strategies: Rationale and Initial Results," Proceedings of the Third IEEE Symposium on High-Performance Computer Architecture (HPCA), February 1-5, 1997, San Antonio, TX, pp. 84-93. (acceptance rate 25%)

(55) B.L. Deitrich and W. W. Hwu, "Speculative Hedge: Regulating Compile-Time Speculation Against Profile Variations," Proceedings of the 29th Annual International Symposium on Microarchitecture, December 2-4, 1996, Paris France pp. 70-79. (acceptance rate 25%)

(56) C.-H. A. Hsieh, J. C. Gyllenhaal, W. W. Hwu, "Java Bytecode to Native Code Translation: The Caffeine Prototype and Preliminary Results," Proceedings of the 29th ACM/IEEE Annual Symposium on Microarchitecture, December 2-4, 1996, Paris France pp. 90-97. (acceptance rate 25%)

(57) Daniel M. Lavery and Wen-mei Hwu, "Modulo Scheduling of Loops in Control-Intensive Non-Numeric Programs," Proceedings of the 29th ACM/IEEE Annual International Symposium on Microarchitecture, December 2-4, 1996, Paris France pp. 126-137. (acceptance rate 25%)

(58) J. C. Gyllenhaal, W. W. Hwu, and B. R. Rau, "Optimization of Machine Descriptions for Efficient Use," Proceedings of the 29th ACM/IEEE Annual International Symposium on Microarchitecture December 2-4, 1996, Paris France pp. 349-358. (acceptance rate 25%)

(59) R. Hank, W.W. Hwu, B.R. Rau, "Region-based Compilation: An Introduction and Motivation," Proceedings of the 28th Annual ACM/IEEE International Symposium on Microarchitecture, Ann Arbor, Michigan, November 1995, pp. 158-168. (acceptance rate 24%)

(60) D. Lavery and W. W. Hwu, "Unrolling-Based Optimizations for Modulo Scheduling," Proceedings of the 28th Annual ACM/IEEE International Symposium on Microarchitecture, Ann Arbor, Michigan, November 1995, pp. 327-337. (acceptance rate 24%)

(61) S. A. Mahlke, R. E. Hank, J. E. MCormick, D. I. August, W. W. Hwu, "A Comparison of Full and Partial Predicated Execution Support for ILP Processors," Proceedings of the 22nd Annual International Symposium on Computer Architecture, Santa Margherita Ligure, Italy, June 1995, pp. 138-150. (20% acceptance rate).

(62) S. A. Mahlke, R. E. Hank, R. A. Bringmann, J. C. Gyllenhaal, D. M. Gallagher, W. W. Hwu, "Characterizing the Impact of Predicated Execution on Branch Prediction," Proceedings of the 27th Annual ACM/IEEE International Symposium on Microarchitecture, San Jose, California November 1994, pp. 118-127. (28% acceptance rate).

(63) Y. Yamada, J. Gyllenhaal, G. Haab, W. W. Hwu, "Data Relocation and Prefetching for Large Data Sets," Proceedings of the 27th Annual ACM/IEEE International Symposium on Microarchitecture, San Jose, California, November 1994, pp. 217-227. (28% acceptance rate).

(64) D. M. Gallagher, W. Y. Chen, S. A. Mahlke, J. C. Gyllenhaal, W. W. Hwu, "Dynamic Memory Disambiguation Using the Memory Conflict Buffer," Proceedings of the 6th Conference on Architecture Support for Programming Languages and Operating Systems, San Jose, California, October, 1994, pp. 183-195.

(65) R. E. Hank, S. A. Mahlke, R. A. Bringmann, J. C. Gyllenhaal, W. W. Hwu, "Superblock Formation Using Static Program Analysis," Proceedings of the 26th Annual ACM/IEEE International On Microarchitecture, Austin, Texas, December 1993, pp. 247-256.

(66) R. A. Bringmann, S. A. Mahlke, R. E. Hank, J. C. Gyllenhaal, W. W. Hwu, "Speculative Execution Exception Recovery Using Write-back Suppression," The 26th Annual ACM/IEEE International Symposium on Microarchitecture, Austin, Texas, December 1993, pp. 247-256.

(67) N. J. Warter, S. A. Mahlke, W. W. Hwu, B. Rau, "Reverse If-Conversion," Proceedings of the ACM SIGPLAN '93 Conference on Programming Language Design and Implementation Albuquerque, New Mexico, June 1993, pp. 290-299.

(68) T. Kiyohara, S. A. Mahlke, W. Y. Chen, R. Bringmann, R. E. Hank, S. Anik, W. W. Hwu, "Register Connection: A New Approach to Adding Registers into Instruction Set Architectures," Proceedings of the 20th Annual International Symposium on Computer Architecture, San Diego, CA, May 17-19, 1993, pp. 247-256.

(69) S. A. Mahlke, W. Y. Chen, W. W. Hwu, M. Schlansker, B. Rau, "Sentinel Scheduling for VLIW And Superscalar Processors," Proceedings of the Fifth International Conference on Architecture Support For Programming Languages and Operating Systems (ASPLOS-V), Boston, MA, Oct. 12-15, 1992, pp. 238-247. (16% acceptance rate).

(70) N. J. Alewine, S. K. Chen, C. C. Li, W. K. Fuchs, W. W. Hwu, "Branch Recovery with Compiler-Assisted Multiple Instruction Retry," Conference Proceedings of the 22nd Annual International Symposium on Fault-Tolerant Computing, Boston, MA, July 8-10, 1992, pp. 66-73. (26% acceptance rate)

(71) S. A. Mahlke, N. J. Warter, W. Y. Chen, P. P. Chang, W. W. Hwu, "The Effect of Compiler Optimizations on Available Parallelism in Scalar Programs," Proceedings of the 20th Annual International Conference on Parallel Processing, St. Charles, IL, Aug 12-16, 1991, pp. 142-145. (34% acceptance rate)

(72) P.P. Chang, W.Y. Chen, S.A. Mahlke, N.J. Warter and Wen-mei W. Hwu, "IMPACT: An Architectural Framework for Multiple-Instruction-Issue Processors," Conference Proceedings of the 18th Annual International Symposium on Computer Architecture, Toronto, Canada, May 28, 1991, pp. 266-275. (27% acceptance rate)

(73) N. J. Warter and W. W. Hwu, "A Software Based Approach to Achieving Optimal Performance for Signature Control Flow Checking," Conference Proceedings of the Twentieth Annual Symposium on Fault-Tolerant Computing, Newcastle upon Tyne, UK, June 26-28, 1990, pp. 442 -449. (26% acceptance rate)

(74) W. W. Hwu and P. P. Chang, "Inline Function Expansion for Compiling Realistic C Programs," ACM SIGPLAN '89 Conference on Programming Language Design and Implementation, Portland, Oregon, June 21-23, 1989, pp. 246-257. (18% acceptance rate)

(75) W. W. Hwu and P. P. Chang, "Achieving High Instruction Cache Performance with an Optimizing Compiler," Conference Proceedings of the 16th Annual International Symposium on Computer Architecture, Jerusalem, Israel, May 28 - June 1, 1989, pp. 242-251. (27% acceptance rate)

(76) W. W. Hwu, T. M. Conte, P. P. Chang, "Comparing Software and Hardware Schemes for Reducing the Cost of Branches," Conference Proceedings of the 16th Annual International Symposium on Computer Architecture, Jerusalem, Israel, May 28 - June 1, 1989, pp. 224-233. (27% acceptance rate)

(77) W. W. Hwu and P. P. Chang, "Exploiting Parallel Microarchitectures with a Compiler Code Generator," The 15th Annual International Symposium on Computer Architecture Conference Proceedings, Honolulu, Hawaii, May 1988, pp. 45-53. (19% acceptance rate)

(78) W.W. Hwu and Y.N. Patt, "Checkpoint Repair for Out-of-order Execution Machines," The 14th Annual International Symposium on Computer Architecture Conference Proceedings, Pittsburgh, PA, June 1987, pp. 18-26. (28% acceptance rate)

(79) W.W. Hwu and Y.N. Patt, "HPSm, a High Performance Restricted Data Flow Architecture Having Minimal Functionality," The 13th Annual International Symposium on Computer Architecture Conference Proceedings, Tokyo, Japan, June 1986, pp. 297-306. (29% acceptance rate)

OTHER CONFERENCE PROCEEDINGS

(1) J. Cabezas, M. Jordà, I. Gelado, N. Navarro, W. W. Hwu, “GPU-SM: Shared Memory Multi-GPU Programming,” The 8th GPUPU Workshop, 2015.

(2) Gai, Jiading; Holtrop, Joseph L. ; Wu, Xiao-Long; Lam, Fan; Fu, Maojing; Haldar, Justin P.; Hwu, Wen-mei; Liang, zhi-Pei; Sutton, Bradley P.,"More IMPATIENT : A Gridding - Accelerated Toeplitz - based S trategy for Non - Cartesian High - Resolution 3D MRI on GPU", Proceedings of the International Society for Magnetic Resonance in Medicine (ISMRM), May 2012 .

(3) "Design evaluation of OpenCL compiler framework for Coarse-Grained Reconfigurable Arrays", Kim, Hee-Seok; Ahn, Minwook; Stratton, John A.; Hwu, Wen-mei, Proceedings of the Field-Programmable Technology (FPT) International Conference, Dec 2012 .

(4) H-S. Kim, S. Wu, L. Chang, W. W. Hwu, “A Scalable Tridiagonal Solver for GPUs.” 40th International Conference on Parallel Processing (ICPP2011), Taipei, Taiwan, September 15, 2011.

(5) Lin, Chih-Sheng; Liu, Wei-Lun; Yeh, Wei-Ting; Chang, Li-Wen; Hwu, Wen-mei; Chen, Sao-Jie; Hsiung, Pao-Ann, "A Tiling-Scheme Viterbi Decoder in Software-Defined Radio for GPUs", Proceedings of the 7th International Conference on Wireless Communications, Networking and Mobile Computing (WiCOM), 2011 .

(6) Wu, Xiao-Long; Gai, Jiading; Lam, Fan; Fu, Maojing; Haldar, Justin P.; Zhuo, Yue; Liang, zhi-Pei; Hwu, Wen-mei; Sutton, Bradley P., "IMPATIENT MRI: Illinois Massively Parallel Acceleration Toolkit for Image Reconstruction with ENhanced Throughput in MRI", Proceedings of the International Society for Magnetic Resonance in Medicine (ISMRM), May 2011.

(7) Papakonstantinou, Alexandros; Liang, Yun; Stratton, John A.; Gururaj, Karthik; Chen, Deming; Hwu, Wen-mei; Cong, Jason, "Multilevel Granularity Parallelism Synthesis on FPGAs", Proceedings of the 2011 International Symposium on Field-Programmable Custom Computing Machines (FCCM), May 2011. (Best Paper Award)

(8) Wu, Xiao-Long; Zhuo, Yue; Gai, Jiading; Lam, Fan; Fu, Maojing; Haldar, Justin P.; Hwu, Wen-mei; Liang, zhi-Pei; Sutton, Bradley P., "Advanced MRI Reconstruction Toolbox with Accelerating on GPUs", Proceedings of the IS&T/SPIE Electronic Imaging 2011 Conference on "Parallel Processing for Imaging Applications", January 2011.

(9) Chang, Li-Wen; Lo, Men-Tzung; Anssari, Nasser; Hsu, Liang-Chuan; Huang, Norden E.; Hwu, Wen-mei, "Parallel Implementation of Multi-Dimensional Ensemble Empirical Mode Decomposition", Proceedings of the IEEE Conference on Acoustics, Speech and Signal Processing, May 2011.

(10) Zhuo, Yue; Sutton, Bradley P.; Wu, Xiao-Long; Haldar, Justin P.; Hwu, Wen-mei; Liang, Zhi-Pei,"Sparse Regularization in MRI Iterative Reconstruction using GPUs", Proceedings of the 3rd International Conference on BioMedical Engineering and Informatics (BMEI'10), October 2010.

(11) Wu, Xiao-Long; Obeid, Nady; Hwu, Wen-mei, "Exploiting More Parallelism from Applications Having Generalized Reductions on GPU Architectures", Proceedings of the 10th IEEE International Conference on Computer and Information Technology (CIT 2010), pp.1175-1180, June 2010.

(12) J. A. Stratton, S. S. Stone and W. W. Hwu, "MCUDA: An Efficient Implementation of CUDA Kernels for Multi-Core CPUs," The 21st International Workshop on Languages and Compilers for Parallel Computing, July 30-31, Canada, 2008. Also available as Lecture Notes in Computer Science 2008.

(13) S. Ueng, M. Lathara, S. Baghsorkhi, Hwu "CUDA-lite, Reducing GPU Programming Complexity," The 21st International Workshop on Languages and Compilers for Parallel Computing, July 30-31, Canada, 2008. Also available as Lecture Notes in Computer Science 2008.

(14) S. Stone, H. Yi, J. Haldar, Z.P. Liang, B. Sutton, and W.W. Hwu, "How GPUs can Improve the Quality of Magnetic Resonance Imaging," 1st GPGPU Workshop, Boston, October 2007.

(15) S. Ryoo, C. Rodrigues, S. Stone, S. Baghsorkhi, S. Ueng, W. Hwu, "Program Optimization Study on a 128-Core GPU," 1st GPGPU Workshop, Boston, October 2007.

(16) S. Ryoo, C. Rodrigues, W.W. Hwu, "Iteration Disambiguation for Parallelism Identification in Time-Sliced Applications," the 20th Workshop of Languages and Compilers for Parallel Computing (LCPC), Urbana, Illinois, October 16-19, 2007.

(17) W.W. Hwu, D. Kirk, S. Ryoo, C. Rodrigues, J. Stratton, K. Hwang, "Performance Insights of Executing Non-Graphics Applications on the NVIDIA GeForce 8800 and the CUDA Parallel Programming Environment," HotChips 19, August 2007.

(18) H. Hunter, E. Nystrom, D. Connors, W. Hwu, Hardware-Compiler Co-Design for Adjustable Data Power Savings,", Proceedings of the 7th Workshop on Media and Streaming Processors, Dec. 2005.

(19) R. Barnes and W. Hwu, "Multi-pass pipelining: enhancing in-order microarchitectures to out-of-order performance," EPIC Workshop, San Jose, March 20, 2005.

(20) H. C. Hunter, E. M. Nystrom, S. Ryoo, and W. W. Hwu, "Applying Scalable Interprocedural Pointer Analysis for Power Reduction of Embedded Data Storage," Workshop on Compilers and Tools for Constrained Embedded Systems, September 2004.

(21) E. Nystrom, H. Kim, and W. Hwu , "Importance of Heap Specialization in Pointer Analysis,", in conference proceedings of PASTE'04, June 7-8, 2004, Washington DC.

(22) C. Li, W. W. Hwu “Extracting Data Flow Model from von Neumann Program for Synthesis,”, in conference proceedings of the 13th International Workshop on Logic and Synthesis, June 2-4, 2004, Temecula, CA

(23) H. C. Hunter and W. W. Hwu**,** "Memory Profiling: Expanding the 3G Developer's Bag of Tricks," Workshop on Compilers and Tools for Constrained Embedded Systems, October 2003.

(24) H. C. Hunter, Chien-Wei Li and Wen-mei W. Hwu, "Motivating use of Memory Profiling in the 3G Domain," Proceedings of the SRC TECHCON 2003, August 2003.

(25) R. D. Barnes, E. M. Nystrom, M. T Conte, Wen-mei W, “Phase profiling in a managed code environment,” The 1st Workshop on Managed Run Time Environment Workloads, San Francisco, March 23, 2003.

(26) H. C. Hunter, W.W. Hwu, “Code Coverage and Input Variability: Effects on Architecture and Compiler Research,” Proceedings of the International Conference on Compilers, Architectures, and Synthesis For Embedded Systems, France, October 2002.

(27) W. W. Hwu, J. W. Sias, M. C. Merten, E. M. Nystrom, R. D. Barnes, C. J. Shannon, S. Ryoo, and J. V. Olivier, “Itanium Performance Insights,” Microprocessor Forum, October 2001.

(28) J. W. Sias, M. C. Merten, E. M. Nystrom, R. D. Barnes, C. J. Shannon, J. D. Matarazzo, S. Ryoo, J. V. Olivier, and W. W. Hwu, “Itanium Performance Insights from the IMPACT Compiler,” HotChips 13, August 2001.

(29) E. M. Nystrom, R. D. Ju, and W. W. Hwu, “Characterization of Repeating Data Access Patterns in Integer Benchmarks,” Memory Performance Issues Workshop at the 28th International Symposium on Computer Architecture, July 2001.

(30) D.A. Connors, J.-M. Puiatti, D.I. August, K.M. Crozier, W.W. Hwu, “An Architectural Framework for Introducing Predicated Execution into Embedded Microprocessors,” Proceedings of the 5th International Euro-Par Conference, August, 1999

(31) M. T. Conte, A. Trick, J. Gyllenhaal, W. W. Hwu,"A Study of Code Reuse Characteristics of Java Applications," MICRO-31 Workshop on Workload Characterization, December 1998, Dallas, TX, pp. 3-10.

(32) D.A. Connors, Y. Yamada, W. W. Hwu, "Software Oriented Floating-Point Format for Enhancing Automotive Control Systems," in Proceedings of the 1st Compiler and Architecture Support for Embedded Computing Systems Workshop, December 1998.

(33) C. A. Hsieh, M. T. Conte, T. L. Johnson, J. C. Gyllenhaal, W. W. Hwu, "A Study of the Cache and Branch Performance Issues with Running Java on Current Hardware Platforms," Proceedings IEEE CompCon, February 1997, pp. 211-216.

(34) N. J. Warter, D. M. Lavery, and W. W. Hwu, "The Benefit of Predicated Execution for Pipelining," Proceedings of the 26th Annual Hawaii International Conference on System Sciences, Wailea, Hawaii, Jan. 5-8, 1993, pp. 497-506.

(35) S. A. Mahlke, W. Y. Chen, J. C. Gyllenhaal, W. W. Hwu, P. P. Chang, and T. Kiyohara, "Compiler Code Transformations for Superscalar-Based High-Performance Systems," Proceedings Supercomputing '92, Minneapolis, Minnesota, Nov. 16-20, 1992, pp. 808-817.

(36) W. Y. Chen, S. A. Mahlke, W. W. Hwu, "Tolerating First Level Memory Access Latency in High-Performance Systems," Proceedings of the 21st Annual International Conference on Parallel Processing, St. Charles, IL, Aug., 1992, pp. (I) 36-43.

(37) S. Anik and W. W. Hwu, "Executing Nested Parallel Loops on Shared-Memory Multiprocessors," Proceedings of the 21st Annual International Conference on Parallel Processing, St. Charles, IL, Aug. 1992, pp. (III) 241-244.

(38) W. Y. Chen, S. A. Mahlke, N. J. Warter, R. E. Hank, R. A. Bringmann, S. Anik, D. M. Lavery, J. C. Gyllenhaal, T. Kiyohara, and W. W. Hwu, "Using profile information to assist advanced compiler Optimization and scheduling," Proceedings of the Fifth Workshop on Languages and Compilers for Parallel Computing Yale University, New Haven, CT, August 1992.

(39) W. Y. Chen, S. A. Mahlke, W. W. Hwu, T. Kiyohara, and P. P. Chang, "Tolerating Data Access With Register Preloading," Proceedings of the 1992 International Conference on Supercomputing, Washington D.C., July 1992, pp. 104-113.

(40) T. M. Conte and W. W. Hwu, "Systematic Prototyping of Superscalar Computer Architectures," in Proceedings of the 3rd IEEE International Workshop on Rapid System Prototyping, Research Triangle Park, NC, June 1992.

(41) S. A. Mahlke, W. Y. Chen, P. P. Chang, and W. W. Hwu, "Scalar Program Performance on Multiple-Instruction-Issue Processors with a Limited Number of Registers," Proceedings of the 25th Annual Hawaii International Conference on System Sciences, Jan. 6-9, 1992, pp. 34-44.

(42) P. P. Chang, W. Y. Chen, S. A. Mahlke, and W. W. Hwu, "Comparing Static and Dynamic Code Scheduling for Multiple-Instruction-Issue Processors," Proceedings of the 24th Annual ACM/IEEE International Symposium on Microarchitecture, Albuquerque, New Mexico, November 18-20, 1991, pp. 25-33.

(43) W. Y. Chen, S. A. Mahlke, P. P. Chang, and W. W. Hwu, "Data Access Microarchitectures for Superscalar Processors with Compiler-Assisted Data Prefetching." Proceedings of the 24th Annual ACM/IEEE International Symposium on Microarchitecture, Albuquerque, New Mexico, November 18-20, 1991, pp. 69-73.

(44) T. M. Conte and W. W. Hwu, "Benchmark Characterization," Proceedings of the 24th Annual Hawaii International Conference on System Sciences, Jan. 8-11, 1991, pp. 364-372.

(45) G. A. Uvieghara, W. W. Hwu, Y. Nakagome, D. K. Jeong, D. Lee, D. A. Hodges, and Y. N. Patt, "An Experimental Single-Chip Data Flow CPU," 1990 Symposium on VLSI Circuits Design Digest of Technical Papers, May 1990.

(46) T. M. Conte and W. W. Hwu, "Benchmark Characterization for Experimental System Evaluation," Proceedings of the 23rd Annual Hawaii International Conference on System Sciences, Jan. 2-5 1990, pp. 6-18.

(47) P. P. Chang and W. W. Hwu, "Forward Semantic: A Compiler-Assisted Instruction Fetch Method For Heavily Pipelined Processors," Proceedings of the 22nd International Workshop on Microprogramming and Microarchitecture Dublin, Ireland, August 14-16, 1989, pp. 188-198.

(48) P. P. Chang and W. W. Hwu, "Control Flow Optimization for Supercomputer Scalar Processing," Proceedings of the 1989 International Conference on Supercomputing, Crete, Greece, June 5-9, 1989.

(49) P. P. Chang and W. W. Hwu, "Trace Selection for Compiling Large C Application Programs to Microcode," Proceedings of the 21st International Microprogramming Workshop, San Diego, CA, Nov. 1988, pp. 21-29.

(50) W. W. Hwu and Yale N. Patt, "HPSm2: A Refined Single-chip Microengine," Proceedings of the 21st Annual Hawaii International Conference on System Sciences, Kona, Hawaii, Jan. 1988, pp 30-40.

(51) J. Wilson, S. Melvin, M. Shebanow, W. W. Hwu, and Y. N. Patt, "On tuning the microarchitecture of an HPS implementation of the VAX. Proceedings of the 20th International Microprogramming Workshop, Colorado Springs, CO, Dec. 1987, pp. 162-167.

(52) W. W. Hwu and Y. N. Patt, "Exploiting horizontal and vertical concurrency via the HPSm microprocessor", Proceedings of the 20th International Microprogramming Workshop, Colorado Springs, CO, Dec. 1987, pp. 154-161.

(53) W.W. Hwu and Y.N. Patt, "Design Choices for the HPSm Microprocessor Chip," Proceedings of The 20th Annual Hawaii International Conference on System Sciences, Kona HI, Jan. 1987, pp. 329-335.

(54) Y. N. Patt, S. W. Melvin, W. Hwu, M. C. Shebanow, C. Chen, J. Wei, "Run-Time Generation of HPS Microinstructions From a VAX Instruction Stream," Proceedings of the 19th Annual Workshop Microprogramming, New York, New York, Oct. 1986, pp. 75-81.

(55) Y.N. Patt, W.W. Hwu, S.W. Melvin, M.C. Shebanow, C. Chen, and J.Wei, "Experiments with HPS, a restricted Data Flow Microarchitecture for High Performance Computers," Digest of Papers, Spring Compcon 86, San Francisco, CA, March 1986, pp. 254-258.

(56) W.W. Hwu, S.W. Melvin, M.C. Shebanow, C. Chen, J. Wei, and Y.N. Patt, "An HPS Implementation of VAX; Initial Design and Analysis," Proceedings of the 19th Annual Hawaii International Conference on System Sciences, Honolulu, HI, Jan. 1986, pp. 282-291.

(57) M.C. Shebanow, Y.N. Patt, W.W. Hwu, and S.W. Melvin, "A C Compiler for HPS I, a Highly Parallel Execution Engine," Proceedings of the 19th Annual Hawaii International Conference on System Sciences, Honolulu, HI, Jan. 1986.

(58) Y.N. Patt, W.W. Hwu, and M.C. Shebanow, "HPS, A New Microarchitecture: Rationale and Introduction," Proceedings of the 18th International Microprogramming Workshop, Asilomar, CA Dec. 1985, pp. 103-108.

(59) Y.N. Patt, S.W. Melvin, W.W. Hwu, and M.C. Shebanow, "Critical Issues Regarding HPS, a High Performance Microarchitecture,"Proceedings of the 18th International Microprogramming Workshop, Asilomar, CA, Dec. 1985, pp. 109-116.

(60) P. Chen, Y. Chen, and W. W. Hwu, "On the Duality of Distributed Interprocess Communication," Proceedings of International Computer Symposium 84, Taipei, Taiwan, Dec. 1984.

(61) Y. N. Patt, R. G. Sheldon, M. C. Shebanow, C. Ponder, and W. W. Hwu, "A Comparison of Evolving (University) Supercomputer Architectures," Proceeding of the 4th Jerusalem Conference of Information Technology, IEEE, May 1984.

(62) W. W. Hwu, M. O. Young, and L. Lee, "A Chinese Text-Speech Development System," Proceedings of the National Computer Symposium 83, Taiwan, Dec. 1983.

BOOK CHAPTERS

(1) B. Bode, M. Butler, T. Dunning, W. Gropp, T. Hoefler, W.W. Hwu, and W. Kramer. “The blue waters super-system for super-science.” In J. S. Vetter, editor, *Contemporary High Performance Computing: From Petascale Toward Exascale*, volume 1 of CRC Computational Science Series, page 900. Taylor and Francis, Boca Raton, 1st Edition, 2013.

(2) D.A. Connors and Wen-mei Hwu, “Architecture,” The VLSI Handbook, 2nd Ed, W.-K. Chen (Editor), CRC Press and IEEE Press, 2006, pp 66-1 – 66-23.

(3) D.A. Connors and Wen-mei Hwu, “Architecture,” The VLSI Handbook, W.-K. Chen (Editor), CRC Press and IEEE Press, 1999, pp. (63) 1-22.

(4) M.T. Conte, A.R. Trick, J.C. Gyllenhaal, W.W. Hwu, “A Study of Code Reuse and Sharing Characteristics of Java Applications,” Workload Charaterization: Methodlogy and Case Studies, L.K. John and A.G. Maynard (Editors), IEEE Computer Society Press, 1999, pp.27-35.

(5) W.W. Hwu and Y.N. Patt, “HPSm, a High Performance Restricted Data Flow Architecture Having Minimal Functionality, “ 25 Years of the International Symposium on Computer Architecture – Selected Papers, G. Sohi (Editor), ACM Press, 1998, Retrospective pp. 43-44, main chapter pp. 300-308

(6) P.P. Chang, S.A. Mahlke, W.Y. Chen, N.J. Warter, and W. W. Hwu, “IMPACT: An Architecture Framework for Multiple-Instruction-Issue Processors,” 25 Years of the International Symposium on Computer Architecture – Selected Papers, G. Sohi (Editor), ACM Press, 1998, Retrospective pp. 77-79, main chapter pp. 408-417

(7) N. J. Alewine, W. K. Fuchs, W. W. Hwu, "Application of Compiler-Assisted Rollback Recovery to Speculative Execution Repair," Hardware and Software Architectures for Fault Tolerance, Experiences and Perspectives, M. Banatre and P. A. Lee (Editors), Springer-Verlag, pp. 45-65, 1994.

(8) K. W. Fuchs, N. J. Alewine, W. W. Hwu, "Speculative Execution and Compiler-Assisted Multiple Instruction Recovery" G. Koob (editor), Kluwer Academic Publishers, 1994.

(9) W. Y. Chen, S. A. Mahlke, N. J. Warter, R. E. Hank, R. A. Bringmann, S. Anik, and W. W. Hwu, "Using Profile Information to Assist Advanced Compiler Optimization and Scheduling," Advances in Languages and Compilers for Parallel Processing, U. Banerjee, D. Gelernter, A. Nicolau, and D. Padua (editors), Pitman Publishing: London, 1993.

(10) W. W. Hwu, et al, "The Superblock: An Effective Technique for VLIW and Superscalar Compilation" Instruction-Level Parallelism, B.R. Rau and J.A. Fisher (editors), Kluwer Academic Publishers, ISBN 0-7923-9367-8, 1993, pp. 229-248.

ABSTRACTS

W. W. Hwu and T. M. Conte, "A Simulation Study of Simultaneous Vector Prefetch Performance in Multiprocessor Memory Subsystems (Extended Abstract)," 1989 ACM SIGMETRICS PERFORMANCE '89 International Conference on Measurement and Modeling of Computer Systems, Berkeley, CA, May 23-26, 1989, pp. 227.

NON-REFEREED JOURNAL PAPERS

T. M. Conte and W. W. Hwu, "A Brief Survey of Benchmark Usage in the Architecture Community, Computer Architecture News, Vol. 19, No. 4, June 1991, pp. 37-44.

A. Glew and W. W. Hwu, "Snoopy Cache Test-and-test-and-set Without Excessive Bus Contention," Computer Architecture News, Vol. 18, No. 2, June 1990.

TECHNICAL REPORTS

J. A Stratton, C. Rodrigues, I.-J. Sung, N. Obeid, L.-W. Chang, N. Anssari, G. D. Liu, Wen-mei W Hwu, "Parboil: A revised benchmark suite for scientific and commercial throughput computing," Technical report, IMPACT Group, University of Illinois, 2012. (265 citations according to Google Scholar).

DISSERTATION

W. W. Hwu, Exploiting Concurrency to Achieve High Performance in a Single-chip Microarchitecture, Report No. UCB/CSD 88/398, January 1988, Computer Science Division (EECS), University California, Berkeley.

SOFTWARE

W.W. Hwu, S. Mahlke, R. Hank, J. Gyllenhaal, D. Lavery, R. Bringmann, D. Gallagher, G. Haab, S. Hwu, IMPACT C Compiler, Release 1.1, 220,000 lines of C code, with code generators for HP PA-RISC, Intel Pentium, SPARC, and HP PlaDoh, Jan. 1995.

W. W. Hwu, P. P. Chang, S. A. Mahlke, W. Y. Chen, R. Ouellette, N. J. Warter, R. Bringmann, IMPACT-I C Compiler for Superscalar Systems, Release 0.1, 180,000 lines of C code, with code generators for MIPS R3000, SUN SPARC, AMD 29000, and Intel i860, Feb., 1991.

Ph.D. THESES SUPERVISED

(1) Shane Ryoo, Ph.D. Thesis, "Optimization Space Carving for Massively Parallel Processors," 2008, Advisor W.W. Hwu, NVIDIA.

(2) Chien-Wei Li, Ph. D. Thesis, "On Extracting Course-Grain Function Parallelim in C Programs," 2006, advisor W.W. Hwu, Intel Corporation.

(3) John Sias, Ph.D. Thesis, "A Systematic Approach to Delivering Instruction-Level Parallelism in EPIC System," 2005 advisor W.W. Hwu, Seminary School.

(4) Ron Barnes, Ph.D. Thesis, "Multiple-Pass Pipelining: Enhancing In-Order Microarchitectures to Out-of-Order Performance," 2005, first job Assistant Professor, George Mason University.

(5) Erik Nystrom, Ph.D. Thesis, "Fulcra Pointer Analysis Frameowork," 2005, advisor W.W. Hwu, first job at Universal Network Machines.

(6) Hillery Hunter, Ph.D.. Thesis, "Matching On-Chip Data Storage to Telecommunication and Media Application Properties," 2004, advisor W.W. Hwu, first job at IBM T.J. Watson Research Center, New York.

(7) Hong-Seok Kim, Ph.D. Thesis, "Compaction algorithm for precise modular context-sensitive pointer analysis," 2004, advisor W.W. Hwu, first job at Samsung Research, Korea.

(8) Marie Conte, Ph.D. Thesis, "A Dynamic Application Analysis Framework," 2003, advisor W.W. Hwu, first job at Intel, Portland, OR.

(9) Matthew Mertsn, Ph.D. Thesis, "Run-time Optimization ," 2002, advisor W.W. Hwu, first job at Intel, Portland, OR.

(10) Jeffrey Monks, Ph.D.. Thesis, "Transmission Power Control for Enhancing The Performance of Wireless Packet Data Networks," 2001, Advisor W.W. Hwu and V. Bharghavan, first job at ByteMobile, Sunnyvale, CA.

(11) Daniel Connors, Ph.D. Thesis, “Eliminating Dynamic Computation Redundancy,” 2000, Advisor W. W. Hwu, first job Assistant Professor in Electrical and Computer Engineering Department, University of Colorado, Boulder, CO.

(12) Ben-chung Cheng, Ph.D. Thesis, “Compile-time Memory Disambiguation for C Programs,” 2000, Advisor W. W. Hwu, first job with Transmeta, Santa Clara, CA.

(13) David August, Ph.D. Thesis, “Systematic Compilation for Predicated Execution” 2000, Advisor W. W. Hwu, first job Assistant Professor in Computer Science Department, Princeton University, NJ.

(14) Le-Chun Wu, Ph.D. Thesis, “Interactive Source-Level Debugging of Optimized Code,” 1999, Advisor Wen-mei Hwu, first job with Hewlett-Packard California Language Laboratories, Cupertino, CA.

(15) Teresa Johnson, Ph.D. Thesis, "Run-time Adaptive Cache Management" 1998, Advisor W. W. Hwu, first job with Hewlett-Packard California Language Laboratories, Cupertino, CA.

(16) Brian Deitrich, Ph.D. Thesis, "Static Program Analysis to Enhance Profile Independence in Instruction Level Parallelism," 1998, Advisor W. W. Hwu, first job with Motorola Corporate Research, Schaumburg, IL.

(17) John Gyllenhaal, Ph.D. Thesis, "An Efficient Framework for Performing Execution-Constraint-Sensitive Transformations That Increase Instruction-Level Parallelism," 1997, Advisor W. W. Hwu, first job Research Assistant Professor, Coordinated Science Laboratory, University of Illinois, Urbana-Champaign.

(18) Liang-Chuan Hsu, Ph.D. Thesis, "A Robust Foundation for Binary Translation of X86 Code, 1997, Advisor W. W. Hwu, first job Associate Professor, Department of Computer Science, Chung-Cheng Institute of Technology, Tao-Yuan, Taiwan.

(19) Daniel Lavery, Ph.D. Thesis, "Modulo Scheduling for Control-Intensive General-Purpose Programs," 1997, Advisor W. W. Hwu, first job with Intel Corporation, Santa Clara, CA.

(20) Richard Hank, Ph.D. Thesis, "Region-Based Compilation," 1996, Advisor W. W. Hwu, first job with Hewlett-Packard California Language Laboratories, Cupertino, CA.

(21) Scott A. Mahlke, Ph. D. Thesis, “Exploiting Instruction Level Parallelism in the Presence of Conditional Branches", 1996, Advisor W. W. Hwu, first job with Hewlett-Packard Labs.

(22) David Gallagher, Ph.D. Thesis, "Memory Disambiguation for Instruction-Level Parallelism Compilation," 1995, Advisor W. W. Hwu, first job Assistant Professor, Air Force Institute of Technology, Dayton, OH.

(23) Roger Bringmann, Ph.D. Thesis, "Enhancing Instruction-Level Parallelism through Compiler Controlled Speculation," 1994, Advisor W. W. Hwu, first job Director for Technology Development, QMS Mobile, AL.

(24) Yoji Yamada, Ph.D. Thesis, "Data Relocation and Prefetching for Programs with Large Data Sets, 1994 Advisor W. W. Hwu, first job with Mazda Research Center, Japan.

(25) Sadun Anik, Ph.D. Thesis, "Architecture and Software Support for Executing Numerical Applications on High-Performance Computers," 1993 Advisor W. W. Hwu, first job with HP Labs, Palo Alto, CA.

(26) William Chen, Ph.D. Thesis, "Data Preload for Superscalar and VLIW Processors," 1993, Advisor W. W. Hwu, first job with Intel Corporation, Santa Clara, CA.

(27) Nancy Warter, Ph.D. Thesis, "Modulo Scheduling with Isomorphic Control Transformations," 1993, Advisor W. W. Hwu, first job with California State University, Los Angeles, CA.

(28) A. Gupta, Ph.D. Thesis, "Performance Aspects of Computers with Graphical User Interfaces," 1993, Advisor W. W. Hwu, first job with Hewlett-Packard, Corvallis, OR.

(29) T. Conte, Ph.D. Thesis, "Systematic Computer Architecture Prototyping," 1992, Advisor W. W. Hwu, first job: Assistant Professor, University of South Carolina at Columbia.

(30) P. Chang, Ph.D. Thesis, "Compiler Support for Multiple Instruction Issue Architectures," 1991, Advisor W. W. Hwu, first job with Intel, Beaverton, Oregon.

MS THESES SUPERVISED (needs update from IMPACT Web Site)

(1) Melvin Larthra, "," 2008, Advisor W.W. Hwu, first job with Intel.
(2) Stephanie Tsao, "," 2008, Advisor W.W. Hwu, first job with Microsoft.

(3) James Player, "An Evaluation of Low-Overhead Parial Flow-Sensitivity," 2005, Advisor W.W. Hwu, first job with Universal Network Machines.

(4) Ian Steiner, "Future Compilation Requirements for Emerging Driving General Purpose Applications," 2005, Advisor W.W. Hwu, first job with Intel.

(5) Jeff Cours, "Using Coverage-Based Analysis to Automate the Customization of the Linux Kernel For Embedded Applications," 2004, Advisor W.W. Hwu.

(6) Shane Ryoo, "Partial Code Elimination in the IMPACT Compiler Framework," 2004, Advisor W.W.Hwu, continued for Ph.D.

(7) Sain Ueng, "Template Bundling for EPIC Architectures," 2004, Advisor W.W. Hwu, continued for Ph.D.

(8) Ron Barnes, "Extracting Hardware-Detected Program Phases for Post-Link Optimization," 2002, Advisor W.W. Hwu, continued for Ph.D.

(9) Erik Nystrom, "Architecture Support for Persistent, Dynamic Code Transformations," 2002, Advisor W.W. Hwu, continued for Ph.D.

(10) Chris Shannon, "The IMPACT SC140 Code Generator," 202, Advsor, W.W. Hwu, continued for Ph.D.

(11) Hillery Hunter, "Characterization of Memory Activity in Media and Telecommunication Applications," 2002, Advisor W.W. Hwu, continued for Ph.D.

(12) Kevin Crozier, "Structural and Static Analysis Techniques for Enhancing Compiler Support of Predicated Execution," 1999, Advisor W. W. Hwu, first job with Hewlett-Packard California Language Laboratories, Cupertino, CA.

(13) John Sias, "Condition Awareness Support for Predicate Analysis and Optimization," 1999, Advisor W.W. Hwu, continued for Ph.D.

(14) Matt Merten, "A Framework for Profile-Driven Optimization in the IMPACT Binary Reoptimization System," 1999, Advisor W.W. Hwu, continued for Ph.D.

(15) Chris George, "A Framework for Install-Time Optimization of Binary Dynamic-Link Libraries," 1999, Advisor W. W. Hwu continue to study law.

(16) Qudus Olaniran, “Emulation of the Intermediate Representation in the IMPACT Compiler,” 1998, Advisor W. W. Hwu, first job with Intel Corporation, Santa Clara, CA.

(17) Michael Thiems, “Optimization and Executable Regeneration in the IMPACT Binary Reoptimization Framework,” 1998, Advisor W. W. Hwu, first job with Motorola Corporate Research, Schaumburg, IL.

(18) Marie Conte, “A Characterization of Code Resue Within Java Applets and Applications,” 1998, Advisor W.W. Hwu, continue for Ph.D.

(19) Jaymie Braun, “Dynamic Control of Compile Time Using Vertical Region-Based Compilation,” 1998, Advisor W. W. Hwu, first job with General Dynamics, San Diego.

(20) Mattew Trommer, “”, 1998, Advisor W. W. Hwu, first job with General Motors.

(21) Richard Kutter, “A Run-time Linking Environment to Enable Dynamic Compilation,” 1998, Advisor W. W. Hwu, first job with Veridian Corporation, Dayton, Ohio.

(22) Dan Connors, "Memory Profiling for Directing Data Speculative Optimizations and Scheduling," 1997, Advisor W. W. Hwu, continue for Ph.D.

(23) Kevin Safford, "A Framework for Using the Pentium's Performance Monitoring Hardware," 1997, Advisor W. W. Hwu, first job with Hewlett-Packard, Fort Collins, Co.

(24) Ben-Chung Cheng, "Pinline: A Profile-Driven Automatic Inliner for the IMPACT Compiler," 1997, Advisor W. W. Hwu, continue for Ph.D.

(25) David I. August, "Hyperblock Performance Optimizations for ILP Processors", 1996, Advisor W. W. Hwu, continued for Ph.D.

(26) James McCormick, "Supporting Predicated Execution: Techniques and Tradeoffs," 1996, W. W. Hwu, first job with Hewlett-Packard, Fort Collins, Co.

(27) Derek Cho, "Lanalysis: A Performance Analysis Tool for the IMPACT Compiler," 1996, W. W. Hwu, first job with Hewlett-Packard, Fort Collins, Co.

(28) Dimitri Argyres, "Performance and Cost Analysis of the Execution Stage of Superscalar Microprocessors," 1995, Advisor W. W. Hwu, first job with Hewlett-Packard, Fort Collins, Co.

(29) Benjamin Sander, "Performance Optimization and Evaluation for the IMPACT X86 Compiler," 1995, Advisor W. W. Hwu, first job with Advanced Micro Devices.

(30) Matthew Gavin, "Using VHDL Synthesis and VLSI Layout Tools for Cost Estimation of Superscalar Issue Units," 1995, Advisor W. W. Hwu, first job with Advanced Micro Devices.

(31) Wayne Dugal, "Code Scheduling and Optimization for a Superscalar X86 Microprocessor," 1995, Advisor W. W. Hwu, first job with LCI, OH.

(32) Teresa Johnson, "Automatic Annotation of Instructions with Profiling Information," 1995, Advisor W. W. Hwu, continued for Ph.D.

(33) Grant Haab, "Data Dependence Analysis for Fortran Programs in the IMPACT Compiler," 1995, Advisor W. W. Hwu, first job with Kuck & Associates.

(34) John Gyllenhaal, "A Machine Description Language for Compilation," 1994, Advisor W. W. Hwu continued for Ph.D.

(35) Roland Ouellette, "Compiler Support for SPARC Architecture Processors," 1994, Advisor W. W. Hwu, first job with Digital Equipment Corporation.

(36) Dave McCracken, "Combined Register Allocation and Scheduling," 1993, Advisor W. W. Hwu, first job with Silicon Graphics.

(37) Greg Mendel, "High-Performance Graphics Optimizations," 1993 Advisor W. W. Hwu, first job with Northrup, OH.

(38) Krishna Subramanian, "Restructuring Techniques for Parallelizing Compilers," 1993, Advisor W. W. Hwu, first job with SUN Microsystems.

(39) Stephen Ziegler, "Aggressive Hardware Support for Reduced Execution in Out-of-order Execution Superscalar Processors," 1993, Advisor W. W. Hwu, first job with Northrop, IN.

(40) Richard Hank, "Machine Independent Register Allocation for the IMPACT-I C Compiler," 1993, Advisor W. W. Hwu, continued for Ph.D.

(41) John G. Holm, "Evaluation of Superscalar and VLIW Processor Designs," 1992, Advisor W. W. Hwu, continued for Ph.D.

(42) John, W. Bockhaus, "An Implementation of GURPR*: A Software Pipelining Algorithm," 1992, Advisor W. W. Hwu, first job with Hewlett-Packard, Fort Collins, CO.

(43) James E. Sicolo, "A Multiported Non-blocking Cache for A Superscalar Uniprocessor," 1992, Advisor W. W. Hwu, first job with Intel, Santa Clara, CA.

(44) David C. Lin, "Compiler Support for Predicated Execution in Superscalar Processors," 1992 Advisor W. W. Hwu, first job with Amdahl, Sunnyvale, CA.

(45) Roger A. Bringmann, “A Template for Code Generator Development Using the IMPACT-I C Compiler," 1992, Advisor W. W. Hwu, continued for Ph.D.

(46) Betty. Y. Kikuta, "An Introductory Computer Architecture Laboratory," 1992, Advisor W. W. Hwu, first job with Motorola Austin, TX.

(47) Jack Choquette, "Analysis of Cache Designs for a Multi-Context Processor," 1992, Advisor W. W. Hwu, first job with Silicon Graphics, CA.

(48) Michael. L. Golden, "Issues in Trace Collection through Program Instrumentation," 1991, W. W. Hwu, continued for Ph.D. at U. of Michigan, Ann Arbor.

(49) James Yu, "Performance Evaluation of Shared Memory Multiprocessors with On-chip Caches," 1991, Advisor W. W. Hwu, first job with Silicon Graphics, California.

(50) William Chen, "An Optimizing Compiler Code Generator: A Platform for RISC Performance Analysis," 1991, Advisor W. W. Hwu, continued for Ph.D.

(51) Scot. Mahlke, MS Thesis, "Design and Implementation of a Portable Global Code Optimizer," 1991, Advisor W. W. Hwu, continued for Ph.D.

(52) William Alexander, "Trace Driven Simulation of Multiple Narrow versus Single Wide Bus Shared Memory Multiprocessors," 1991, Advisor W. W. Hwu, first job with Intel Beaverton, Oregon

(53) Andrew Glew, "Synchronization Primitive Implementation Including the Bus Abandonment Lock," 1991, Advisor W. W. Hwu, first job with Intel Beaverton, Oregon.

(54) Brian Upper, "The Application of Inline Target Insertion to an Asynchronous Processor Pipeline," 1991, Advisor W. W. Hwu, first job with Paracom, Chicago, IL.

(55) James Magro, "Sectored Instruction Cache to Take Advantage of Instruction Placement: an Asynchronous Pipelined Approach," 1990, Advisor W. W. Hwu, first job with NCR Clemson, SC.

(56) Terry Tsai, "The Design of an Instruction Cache for the IMPACT RISC Processor," 1989, Advisor W. W. Hwu, first job with Texas Instrument, Dallas, TX.

(57) Chris White, "Design Tradeoffs for A High-Bandwidth, Low-Latency Register File Design," 1989, Advisor, W. W. Hwu, first job with Motorola, Austin, TX.

(58) Sadun Anik, "Compilation and Architecture Issues for Executing Numerical Programs with Super-Scalar Architectures," 1989, Advisor W. W. Hwu, continued for Ph.D.

(59) Pohua Chang, "Aggressive Code Improving Techniques Based on Control Flow Analysis," 1989, Advisor W. W. Hwu, continued for Ph.D.

(60) Tom Conte, "The Simulation and Tuning of the Global Memory Subsystem of a Multiprocessor,"1988, Advisor W. W. Hwu, continued for Ph.D.