

NITI MADAN

NSF/CRA Computing Innovation Fellow 2009-2011
Post-doctoral Researcher at IBM T.J. Watson Research Center

Email: nitimadan@gmail.com

Web: <http://www.cs.utah.edu/~niti>

Phone: 801-529-8888

AREAS OF INTEREST

I am interested in all aspects of computer architecture and systems research. My specialized interests include power-efficient, reliable and secure computing systems. I am also interested in emerging 3D die stacking technology, memory system design, and heterogeneous CMPs.

SUMMARY OF EXPERIENCE

- 1.5 years Post-doc research experience at Reliability and Power-aware Micro-architectures group in IBM T.J. Watson Research Center
- 4-months Industry research internship experience at Systems Technology Lab, Intel
- 6-months Post-silicon validation internship experience of Victoria Falls at Sun Microsystems
- 4 years Graduate research experience in Computer Architecture (Performance Analysis/modeling)
- 2 years Masters thesis research experience in Digital VLSI design (Asynchronous circuit design, FPGA based designs)
- 2 years experience as a teaching assistant for various CS/CE courses

EDUCATION

Aug'04-Dec'09

Ph.D. in Computer Science

[University of Utah](#)

Aug'01-Aug'04

M.S. in Computer Science

[University of Utah](#)

Aug'97-May'01

B.E. in Electrical Engineering

[Delhi College of Engineering](#)

University of Delhi, India

HONORS

- [NSF/CRA Computing Innovation Fellowship](#) 2009-2011 ([Overall acceptance rate](#) = 60/526 (11%) and acceptance rate for international students = 15/278 (~5%). Selected by CCC/CRA and funded by NSF a 2 year post-doctoral fellowship to advance the field of computing.
- HPCA-2010 paper nominated to appear in IEEE Micro Jan 2011 edition of Top Picks of Computer Architecture conferences in 2010 (Acceptance rate = 12.6%). This work is one of the eleven papers recognized as "the year's most significant research publications in Computer Architecture based on novelty and industry relevance".

RESEARCH EXPERIENCE

Oct'09-Current: Post-doctoral Research in Power-aware Architectures

Mentor: Pradip Bose, Manager, IBM T.J. Watson Research Center, Yorktown Heights, NY

Research Project: "Guarded Power Gating for Multi-cores"

Power gating or turning idle logic blocks off is becoming an increasingly important actuation knob in chip-level dynamic power management. In a multi-core setting, a key design issue is determining the right balance of gating at the unit-level (within a core) and at the core-level. Another issue is how to architect the predictive control associated with such gating, in order to ensure maximal power savings at minimal performance loss. The goals of this research is to architect efficient mechanisms which we term as “Guard Mechanisms” that can provide performance and power savings guarantees in dynamic power management and to develop new models that can study the effectiveness of these techniques. Following are contributions till date:

- A case for guarded power management in multi-cores (**HPCA-2011** and pending patent)
- Understanding the fundamental tradeoffs between unit-level power gating and per-core power gating algorithms for multi-cores (WEED-2)
- Extended a queuing-model based framework to evaluate various power gating algorithms

Jan’05- Jan’09: Dissertation Research in Computer Architecture

Advisor: [Prof. Rajeev Balasubramonian](#), Assistant Professor, School of Computing, University of Utah

Dissertation: “*Leveraging Mixed Process 3D Technology for Future Cache Hierarchies and Reliability*” – Emerging three-dimensional (3D) integration technology enables vertical stacking of silicon dies with high density and low latency interconnects. This results in increased processor performance as well as reduced power consumption because of smaller on-chip wires. Several research opportunities and challenges in 3D technology are being studied at the fabrication, circuit design and architecture level. This dissertation explores novel applications for 3D die stacking at the micro-architecture level with an emphasis on mixed-process integration. Following are contributions till date:

- Optimizing capacity and communication trade-offs in future reconfigurable cache hierarchies by leveraging 3D stacking and page coloring (**HPCA-2009**)
- Leveraging 3D die stacking technology for improving processor reliability by exploiting heterogeneous process for implementing redundant hardware (**MICRO-40**).
- Employing simple in-order checkers that can be frequency scaled for reducing power overheads of redundancy in chip-level redundant multithreading for single threaded and multi-programmed workloads (**TPDS’07**, SELSE-2).
- Improving register file reliability and reducing performance overheads in a redundantly threaded processor via eager register release (WAR-2).

May’02- Dec’03 : Research Assistant in Asynchronous Circuits and System Design Group

Advisor: [Prof. Erik Brunvand](#) , Associate Professor, School of Computing, University of Utah

- **Masters Thesis:** “[Asynchronous Microengines for Network Processing](#)”
Masters thesis research involved designing an Asynchronous Micro-programmable Network Processing Framework. Asynchronous circuit design was used to prototype this IP router on a Xilinx FPGA board. Extensions such as stateless firewalling etc. were added by modifying the microcode (ANCHOR-2004).
- Explored a design methodology for implementing asynchronous circuits modeled with bundled-delays on Xilinx FPGAs.

INDUSTRY EXPERIENCE

Oct’09-Current: NSF funded Post-doctoral Researcher at Power and Reliability-aware Architectures Group, IBM T.J. Watson Research Center, Yorktown Heights, NY

Manager: Pradip Bose, Research Staff Member and Manager, IBM Research

Co-mentors: Alper Buyuktosunoglu, Research Staff Member, IBM Research and Prof. Murali Annavaram, ECE Dept, USC

Exploring guarded power-gating algorithms.

Apr'08-Aug'08: Graduate Technical Intern at Hardware Architecture Lab, Systems Technology Lab (CTG), Intel, Hillsboro, OR

Mentors: Ravishankar Iyer, Principal Engineer and Srihari Makineni, Senior Researcher

Manager: Donald Newell, Senior Principal Engineer (now CTO at AMD)

Explored DRAM cache design for Intel's next generation processor and reconfigurable SRAM-DRAM 3D stacked cache hierarchies.

Feb'07-Aug'07: Post-Silicon Debug/Validation Intern at Sun Microsystems (now Oracle), Santa Clara, CA

Manager: Ramaswamy Sivaramkrishnan, Principal Engineer

Assisted Root-cause analysis team in system-level post-silicon debug of the next generation Sun's processor "Victoria Falls".

TEACHING EXPERIENCE

- Teaching Assistant for [Engineering Programming](#) in Spring 2005
- Teaching Assistant for [Advanced I.C. Design](#) in Fall 2004
- Teaching Assistant for [Asynchronous VLSI Architecture](#) in Spring 2004
- Teaching Assistant for [Theory of Computation](#) in Spring 2002
- Teaching Assistant for [Computer Architecture](#) in Fall 2001

REFEREED PUBLICATIONS

- [A Case for Guarded Power Gating in Multi-core Processors](#), **Niti Madan**, Alper Buyuktosunoglu, Pradip Bose and Murali Annavaram, *17th International Symposium on High-Performance Computer Architecture 2011 (HPCA-17)*
- [CHOP: Adaptive Filter-based DRAM Caching for CMP Server Platforms](#), Xiaowei Jiang, **Niti Madan**, Li Zhao, Mike Upton, Ravi Iyer, Srihari Makineni, Donald Newell, Yan Solihin, Rajeev Balasubramonian, *IEEE Micro, Special Issue: Micro's Top Picks from 2009 Computer Architecture Conferences (MICRO TOP PICKS)*, January-February 2011
- [Guarded Power Gating in a Multi-core Setting](#), **Niti Madan**, Alper Buyuktosunoglu, Pradip Bose and Murali Annavaram, *2nd Workshop on Energy-Efficient Design (WEED-2) held in conjunction with ISCA-34 2010*
- Power-efficient, Reliable Microprocessor Architectures: Modeling and Design Methods, Several researchers at T.J. Watson Research Center, *ACM Great Lakes Symposium on VLSI 2010 (GLVLSI)* (Invited paper)
- [CHOP: Adaptive Filter-based DRAM Caching for CMP Server Platforms](#), Xiaowei Jiang, **Niti Madan**, Li Zhao, Mike Upton, Ravi Iyer, Srihari Makineni, Donald Newell, Yan Solihin, Rajeev Balasubramonian, *16th International Symposium on High-Performance Computer Architecture 2010 (HPCA-16)* (Acceptance rate 18%)
- [Optimizing Communication and Capacity in a 3D Stacked Reconfigurable Cache Hierarchy](#), **Niti Madan**, Li Zhao (Intel), Naveen Muralimanohar, Aniruddha Udipi, Rajeev Balasubramonian, Ravishankar Iyer (Intel), Srihari Makineni (Intel), Donald Newell (Intel), *15th International Symposium on High-Performance Computer Architecture 2009 (HPCA-15)*. (Acceptance rate 19%)
- ["Scalable and Reliable Communication for Hardware Transactional Memory"](#), Seth Pugsley, Manu Awasthi, **Niti Madan**, Naveen Muralimanohar and Rajeev Balasubramonian, *17th International Conference on Parallel Architectures and Compilation Techniques 2008 (PACT-17)* (Acceptance rate 19%)
- ["Leveraging 3D Technology for Improving Processor Reliability"](#), **Niti Madan** and Rajeev Balasubramonian, *40th International Symposium on Microarchitecture 2007 (MICRO-40)*. (Acceptance Rate 21%)

- [“Power-efficient Approaches to Redundant Multithreading”](#), **Niti Madan** and Rajeev Balasubramonian, *IEEE Transactions on Parallel and Distributed Systems (TPDS), Special Issue on CMPs, August 2007*
- [“Exploiting Eager Register Release in a Redundantly Multi-Threaded Processor”](#), **Niti Madan** and Rajeev Balasubramonian, *2nd Workshop on Architectural Reliability (WAR-2) held in conjunction with MICRO-39 2006*
- [“A First-Order Analysis of Power Overheads of Redundant Multi-Threading”](#), **Niti Madan**, Rajeev Balasubramonian, *2nd Workshop on System Effects of Logic Soft Errors (SELSE-2), April 2006*
- [“A Case for Asynchronous Microengines for Network Processing”](#), **Niti Madan** and Erik Brunvand, *Advanced Networking and Communications Hardware Workshop (ANCHOR 2004) held in conjunction with 31st Annual Symposium in Computer Architecture (ISCA 2004)*

NON-REFEREED PUBLICATIONS

- [Leveraging Mixed-process 3D Die Stacking Technology for Cache Hierarchies and Reliability](#), Niti Madan, *Doctoral Dissertation, December 2009, University of Utah*
- [“Scalable, Reliable, Power-efficient Communication for Hardware Transactional Memory”](#), Seth Pugsley, Manu Awasthi, Niti Madan, Naveen Muralimanohar and Rajeev Balasubramonian, *Technical Report UUCS-08-001, January 2008*
- [“Power-efficient Approaches to Reliability”](#), Niti Madan, Rajeev Balasubramonian, *Technical Report UUCS-05-010, December 2005*
- [“Asynchronous Microengines for Network Processing”](#), Niti Madan, *Masters Thesis, August 2004, University of Utah*

POSTER PRESENTATIONS

- “Guarded Power Gating for Multicore Processors”, Presented at *CRA Career Mentoring Workshop, Washington D.C., December 2010*
- “Power-efficient Approaches to Reliability”, Presented at *CRA-W and CDC Computer Architecture Summer School, Princeton University, July 2006*
- “Power-efficient Approaches to Reliability”, Presented at *Research Day, School of Computing, University of Utah, March 2005*

SKILL SET

- *Architecture Simulators/Models* – SimpleScalar (Modified SimpleScalar to model a dual core SMT simulator), Casper (Intel functional cache model), ManySim (Intel in-house full system simulator), QUTE (Queuing model based simulator at IBM), Thermal and power modeling using Wattch and HotSpot
- *Programming Languages* - C, Perl
- *CAD tools*- Familiar with various Synopsys/Cadence tools and ModelSim

COURSES TAKEN DURING M.S. AND Ph.D.

[Advanced I.C. Design I](#), [Advanced Computer Architecture](#), [Foundations of Computer Science](#), [Advanced I.C. Design II](#), [Programming languages and Semantics](#), [Operating System Design](#), [Async/Architecture Seminar](#), [Special topics in High Performance Architectures](#), [Computer Networks](#), [VLSI Testing and Verification](#), [Parallel Computer Architecture](#), [Research Proposals](#)

TALKS

Presented research work at the following workshops/conferences: ANCHOR-2004, SELSE-2 (2006), WAR-2 (2006), MICRO-2007, WEED-2010, HPCA-2011, Santa Clara University (2011), ISCA 2011 Tutorial.

PROFESSIONAL MEMBERSHIP/SERVICE

- Program committee member
 - IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS) 2011,
 - The Second Exascale Evaluation and Research Techniques Workshop Workshop (EXERT-2) held in conjunction with ASPLOS-2011
 - Workshop on Duplicating, Deconstructing, and Debunking (WDDD 2011) to be held in conjunction with ISCA-2011
- Lead tutorial organizer for tutorial titled “Energy Secure System Architectures” held at ISCA-2011
- External Reviewer for Async’03, HPCA’07, IEEE TCAD’07, ACM TACO’08, HPCA’11, ISCA’11, ACM Computing Frontiers’11, IEEE Micro’11, ACM Computing Surveys’11, MICRO-44
- Member – IEEE, ACM, ACM-SIGARCH

REFERENCES

Available upon request.