15-740/18-740 Computer Architecture, Fall 2009
Schedule for In-Class Discussions

Day 1: *Tuesday, October 13, 2009*

Exploiting New Memory Technologies: Phase-change memories (e.g., flash) have unique properties compared with SRAM, DRAM, or disk. How can this new techology be exploited in future memory hierarchies?

- Benjamin C. Lee, Engin Ipek, Doug Burger, and Onur Mutlu. *"Architecting Phase Change Memory as a Scalable DRAM Alternative,"* in *Proceedings of the 36th Annual International Symposium on Computer Architecture (ISCA)*, June 2009.
- P. Zhou, B. Zhao, J. Yang, and Y. Zhang. *"A durable and energy efficient main memory using phase change memory technology,"* in *Proceedings of the 36th Annual International Symposium on Computer Architecture (ISCA)*, June 2009.
- Moinuddin K. Qureshi, Vijayalakshmi Srinivasan, and Jude A. Rivers. *"Scalable High Performance Main Memory System Using Phase-Change Memory Technology,"* in *Proceedings of the 36th Annual International Symposium on Computer Architecture (ISCA)*, June 2009.
- Xiaoxia Wu, Jian Li, Lixin Zhang, Evan Speight, Ram Rajamony, and Yuan Xie. *"Hybrid Cache Architecture with Disparate Memory Technologies,"* in *Proceedings of the 36th Annual International Symposium on Computer Architecture (ISCA)*, June 2009.

Optimizing Shared Caches in Chip Multiprocessors: Rethinking how data should be placed and managed within the on-chip cache hierarchy.

- Jichuan Chang and Gurindar S. Sohi. *"Cooperative Caching for Chip Multiprocessors,"* in *Proceedings of the 33rd Annual International Symposium on Computer Architecture (ISCA)*, June 2006.
- Nikos Hardavellas, Michael Ferdman, Babak Falsafi, and Anastasia Ailamaki. *"Reactive NUCA: Near-Optimal Block Placement and Replication in Distributed Caches",* in *Proceedings of the 36th Annual International Symposium on Computer Architecture (ISCA)*, June 2009.
- Shekhar Srikantaiah, Mahmut Kandemir, and Mary Jane Irwin. *"Adaptive Set-Pinning: Managing Shared Caches in Chip Multiprocessors,"* in *Proceedings of the Thirteenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, October 2008.

On-Chip Interconnects: How should the on-chip interconnect be designed for a chip multiprocessor, and how does this design interact with the on-chip cache hierarchy?

- Martha Mercaldi Kim, John D. Davis, Mark Oskin, and Todd Austin. *"Polymorphic On-Chip Networks,"* in *Proceedings of the 35th Annual International Symposium on Computer Architecture (ISCA)*, June 2008.
- Dongkook Park, Soumya Eachempati, Reetuparna Das, Asit K. Mishra, Yuan Xie, N. Vijaykrishnan, and Chita R. Das. *"MIRA: A Multi-layer On Chip Interconnect Router Architecture,"* in *Proceedings of the 35th Annual International Symposium on Computer Architecture (ISCA)*, June 2008.

Exploiting Asymmetric CMPs: How could we exploit a chip-multiprocessor comprised of a range of heterogeneous cores that differ in their performance/power/area characteristics?

- Saisanthosh Balakrishnan, Ravi Rajwar, Mike Upton, and Konrad Lai. *"The Impact of Performance Asymmetry in Emerging Multicore Architectures,"* in *Proceedings of the 32nd Annual International Symposium on Computer Architecture (ISCA)*, June 2005.
- M. Aater Suleman, Onur Mutlu, Moinuddin K. Qureshi, and Yale N. Patt. *"Accelerating critical section execution with asymmetric multi-core architectures,"* in *Proceedings of the Fourteenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, March 2009.
- Rakesh Kumar, Dean M. Tullsen, and Norman P. Jouppi. *"Core architecture optimization for heterogeneous chip multiprocessors,"* in *Proceedings of the 15th International Conference on Parallel Architectures and Compilation Techniques (PACT)*, September 2006.

Exploiting GPUs and Accelerators: How might we exploit massive amounts of data parallelism via graphics processors and other accelerators.

- David Tarditi, Sidd Puri and Jose Oglesby. *"Accelerator: Using Data-Parallelism to Program GPUs for General-Purpose Uses,"* in *Proceedings of the Twelfth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, October 2006.
- John H. Kelm, Daniel R. Johnson, Matthew R. Johnson, Neal C. Crago, William Tuohy, Aqeel Mahesri, Steven S. Lumetta, Matthew I. Frank, and Sanjay J. Patel. *"Rigel: An Architecture and Scalable Programming Interface for a 1000-core Accelerator,"* in *Proceedings of the 36th Annual International Symposium on Computer Architecture (ISCA)*, June 2009.

Table 1: Day 1 Discussion Leaders

Topic	Time Slot	Discussion Leaders
Exploiting New Memory Technologies	1:30-1:45	Chris Craik Aapo Kyrola Yoshihisa Abe
Optimizing Shared Caches in CMPs	1:45-2:00	Samir Sapra Athula Balachandran Ravishankar Krishnaswamy
On-Chip Interconnects	2:00-2:15	Jiri Simsa Alex Grubb Jenn Tam Harsha Vardhan Simhadri
Exploiting Asymmetric CMPs	2:15-2:30	Sam Ganzfried Aniket Ponkshe Ryan Sukauye
Exploiting GPUs and Accelerators	2:30-2:45	Lawrence Tan Yanlin Li Jonathan Coens

Day 2: *Wednesday, October 14, 2009*

Transactional Memory: Supporting Large Transactions: Conventional designs for hardware-supported transactional memory can only handle transactions of a limited size: how can we extend this support to larger transactions?[1]

- Ravi Rajwar, Maurice Herlihy, and Konrad Lai. *"Virtualizing Transactional Memory,"* in *Proceedings of the 32nd Annual International Symposium on Computer Architecture (ISCA)*, June 2005.
- JaeWoong Chung, Chi Cao Minh, Austen McDonald, Travis Skare, Hassan Chafi, Brian D Carlstrom, Christos Kozyrakis and Kunle Olukotun. *"Tradeoffs in Transactional Memory Virtualizations,"* in *Proceedings of the Twelfth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, October 2006.
- Colin Blundell, Joe Devietti, E. Christopher Lewis, and Milo M. K. Martin. *"Making the Fast Case Common and the Uncommon Case Simple in Unbounded Transactional Memory,"*, in *Proceedings of the 34th Annual International Symposium on Computer Architecture (ISCA)*, June 2007.
- Jayaram Bobba, Neelam Goyal, Mark D. Hill, Michael M. Swift, and David A. Wood. *"TokenTM: Efficient Execution of Large Transactions with Hardware Transactional Memory,"* in *Proceedings of the 35th Annual International Symposium on Computer Architecture (ISCA)*, June 2008.

Transactional Memory: Hybrid Hardware/Software Approaches: How can we combine software- and hardware-based transactional memory to get the best of both worlds?[1]

- Peter Damron, Alexandra Fedorova, Yossi Lev, Victor Luchangco, Mark Moir and Dan Nussbaum. *"Hybrid Transactional Memory,"* in *Proceedings of the Twelfth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, October 2006.
- Chi Cao Minh, Martin Trautmann, JaeWoong Chung, Austen McDonald, Nathan Bronson, Jared Casper, Christos Kozyrakis, and Kunle Olukotun. *"An Effective Hybrid Transactional Memory System with Strong Isolation Guarantees,"* in *Proceedings of the 34th Annual International Symposium on Computer Architecture (ISCA)*, June 2007.
- Arrvindh Shriraman, Michael F Spear, Hemayet Hossain, Virendra J Marathe, Sandhya Dwarkadas, and Michael L Scott. *"An Integrated Hardware-Software Approach to Flexible Transactional Memory,"* in *Proceedings of the 34th Annual International Symposium on Computer Architecture (ISCA)*, June 2007.

Thread-Level Speculation: Recent work on optimistically exploiting thread-level parallelism.

- Christopher B. Colohan, Anastassia Ailamaki, J. Gregory Steffan, and Todd C. Mowry. *"Tolerating Dependences Between Large Speculative Threads Via Sub-Threads,"* in *Proceedings of the 33rd Annual International Symposium on Computer Architecture (ISCA)*, June 2006.
- Luis Ceze, James Tuck, Calin Cascaval, and Josep Torrellas. *"Bulk Disambiguation of Speculative Threads in Multiprocessors,"* in *Proceedings of the 33rd Annual International Symposium on Computer Architecture (ISCA)*, June 2006.

Memory Consistency: Recent innovations in supporting memory consistency models.

- Arvind and Jan Willem Maessen. *"Memory Model = Instruction Reordering + Store Atomicity,"* in *Proceedings of the 33rd Annual International Symposium on Computer Architecture (ISCA)*, June 2006.
- Thomas F Wenisch, Anastasia Ailamaki, Babak Falsafi, and Andreas Moshovos. *"Mechanisms for Store-wait-free Multiprocessors,"* in *Proceedings of the 34th Annual International Symposium on Computer Architecture (ISCA)*, June 2007.

[1] NOTE: the following paper is good background reading for both of the transactional memory discussion topics, but it is not something that you should explicitly cover: Austen McDonald, JaeWoong Chung, Brian D. Carlstrom, Chi Cao Minh, Hassan Chafi, Christos Kozyrakis, and Kunle Olukotun. *"Architectural Semantics for Practical Transactional Memory,"* in *Proceedings of the 33rd Annual International Symposium on Computer Architecture (ISCA)*, June 2006.)

- Luis Ceze, James Tuck, Pablo Montesinos, and Josep Torrellas. *"BulkSC: Bulk Enforcement of Sequential Consistency,"* in *Proceedings of the 34th Annual International Symposium on Computer Architecture (ISCA)*, June 2007.

Deterministic Multiprocessing: What if we wanted a parallel machine to produce the same result each time it executed, while still exploiting parallelism?

- Joseph Devietti, Brandon Lucia, Luis Ceze, Mark Oskin. *"DMP: Deterministic Shared Memory Multiprocessing,"* in *Proceedings of the Fourteenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, March 2009.
- Marek Olszewski, Jason Ansel, and Saman Amarasinghe. *"Kendo: Efficient Determistic Multithreading in Software,"* in *Proceedings of the Fourteenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, March 2009.
- Pablo Montesinos, Matthew Hicks, Samuel King, and Josep Torrellas. *"Capo: A Software-Hardware Interface for Practical Deterministic Multiprocessor Replay,"* in *Proceedings of the Fourteenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, March 2009.

Table 2: Day 2 Discussion Leaders

Topic	Time Slot	Discussion Leaders
Transactional Memory: Large Transactions	1:30-1:45	Kanat Tangwongsan Anvesh Komuravelli Abe Othman
Transactional Memory: Hybrid HW/SW Approaches	1:45-2:00	Vivek Seshadri K. L. Kaushik Mike Rozyczko
Thread-Level Speculation	2:00-2:15	Dafna Shahaf Anthony Gitter Or Sheffet
Memory Consistency	2:15-2:30	Rakesh Iyer Arbob Ahmad Henry DeYoung
Deterministic Multiprocessing	2:30-2:45	Chris Fallin Zongwei Zhou David Lewis

Day 3: *Thursday, October 15, 2009*

Recording Inter-Thread Data Dependencies for Deterministic Replay: To support debugging of parallel programs, we need to be able to reconstruct inter-thread data dependencies. How do we do this efficiently?

- Min Xu, Rastislav Bodik, and Mark D. Hill. *"A 'Flight Data Recorder' for Enabling Full-system Multiprocessor Deterministic Replay,"* in *Proceedings of the 30th Annual International Symposium on Computer Architecture (ISCA)*, June 2003.
- Min Xu, Rastislav Bodik and Mark Hill. *"A Regulated Transitive Reduction for Longer Memory Race Recording,"* in *Proceedings of the Twelfth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, October 2006.
- Satish Narayanasamy, Gilles Pokam, and Brad Calder. *"BugNet: Continuously Recording Program Execution for Deterministic Replay Debugging,"* in *Proceedings of the 32nd Annual International Symposium on Computer Architecture (ISCA)*, June 2005.
- Satish Narayanasamy, Cristiano Pereira and Brad Calder. *"Recording Shared Memory Dependencies Using Strata,"* in *Proceedings of the Twelfth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, October 2006.

Dynamic Checking of Program Invariants: Techniques for monitoring programs as they execute to check for bugs and security problems.

- Michael Dalton, Hari Kannan, and Christos Kozyrakis. *"Raksha: A Flexible Information Flow Architecture for Software Security,"* in *Proceedings of the 34th Annual International Symposium on Computer Architecture (ISCA)*, June 2007.
- Shimin Chen, Michael Kozuch, Theodoros Strigkos, Babak Falsafi, Phillip B. Gibbons, Todd C. Mowry, Vijaya Ramachandran, Olatunji Ruwase, Michael Ryan, and Evangelos Vlachos. *"Flexible Hardware Acceleration for Instruction-Grain Program Monitoring,"* in *Proceedings of the 35th Annual International Symposium on Computer Architecture (ISCA)*, June 2008.

Debugging Parallel Programs: Other techniques for detecting and fixing bugs in parallel programs.

- Shan Lu, Joseph Tucek, Feng Qin and Yuanyuan Zhou. *"AVIO: Detecting Atomicity Violations via Access Interleaving Invariants,"*, in *Proceedings of the Twelfth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, October 2006.
- Brandon Lucia, Joseph Devietti, Karin Strauss, and Luis Ceze. *"Atom-Aid: Detecting and Surviving Atomicity Violations,"* in *Proceedings of the 35th Annual International Symposium on Computer Architecture (ISCA)*, June 2008.

Hardware Faults on CMPs: What can we do if the hardware fails?

- Philip M. Wells, Koushik Chakraborty, and Gurindar S. Sohi. *"Adapting to Intermittent Faults in Multicore Systems,"* in *Proceedings of the Thirteenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, October 2008.
- Philip Wells, Koushik Chakraborty, and Guri Sohi. *"Mixed-Mode Multicore Reliability,"* in *Proceedings of the Fourteenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, March 2009.
- Michael D. Powell, Arijit Biswas, Shantanu Gupta, and Shubhendu S. Mukherjee. *"Architectural Core Salvaging in a Multi-Core Processor for Hard-Error Tolerance,"* in *Proceedings of the 36th Annual International Symposium on Computer Architecture (ISCA)*, June 2009.

Optimizing Power and Heat: Power and thermal issues are major constraints on hardware technology these days. What can we do to improve power efficiency?

- James Donald and Margaret Martonosi. *"Techniques for Multicore Thermal Management: Classification and New Exploration,"* in *Proceedings of the 33rd Annual International Symposium on Computer Architecture (ISCA)*, June 2006.
- Abhishek Bhattacharjee and Margaret Martonosi. *"Thread Criticality Predictors for Dynamic Performance, Power, and Resource Management in Chip Multiprocessors,"* in *Proceedings of the 36th Annual International Symposium on Computer Architecture (ISCA)*, June 2009.
- Krishna K. Rangan , Gu-Yeon Wei, and David Brooks. *"Thread Motion: Fine-Grained Power Management for Multi-Core Systems,"* in *Proceedings of the 36th Annual International Symposium on Computer Architecture (ISCA)*, June 2009.
- Yefu Wang, Kai Ma, and Xiaorui Wang. *"Temperature-Constrained Power Control for Chip Multiprocessors with Online Model Estimation,"* in *Proceedings of the 36th Annual International Symposium on Computer Architecture (ISCA)*, June 2009.

Table 3: Day 3 Discussion Leaders

Topic	Time Slot	Discussion Leaders
Recording Inter-Thread Dependences	1:30-1:45	Arvind Gopalakrishnan Tarun Goyal Kevin Waugh
Dynamic Checking of Program Invariants	1:45-2:00	Gunhee Kim Hyeontaek Lim Aaron Roth Ali Kemal Sinop
Debugging Parallel Programs	2:00-2:15	Chintan Parikh Chris Martens Ivan Jager
Hardware Faults on CMPs	2:15-2:30	Lavanya Subramanian Sven Stork Bodicherla Aditya Prakash
Optimizing Power and Heat	2:30-2:45	Field Cady Bin Fu Kai Ren