15-740/18-740 Computer Architecture, Fall 2012
Schedule for In-Class Discussions

Day 1: *Tuesday, October 16, 2012*

Parallel Programming Models and Languages: What are some recent innovations in how programmers can write parallel software?

- Jeffrey Dean and Sanjay Ghemawat. *"MapReduce: simplified data processing on large clusters,"* in *Commun. ACM 51*, 1 (January 2008), 107-113.
- Michael Isard, Mihai Budiu, Yuan Yu, Andrew Birrell, and Dennis Fetterly. *"Dryad: distributed data-parallel programs from sequential building blocks,"* in *Proceedings of the 2nd ACM SIGOPS/EuroSys European Conference on Computer Systems 2007 (EuroSys '07).*
- Christopher Olston, Benjamin Reed, Utkarsh Srivastava, Ravi Kumar, and Andrew Tomkins. *"Pig latin: a not-so-foreign language for data processing,"* in *Proceedings of the 2008 ACM SIGMOD international conference on Management of data (SIGMOD '08).*
- Sungpack Hong, Hassan Chafi, Edic Sedlar, and Kunle Olukotun. *"Green-Marl: a DSL for easy and efficient graph analysis,"* in *Proceedings of the seventeenth international conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '12).*

Exploiting Parallelism on GPUs: Beyond traditional graphics processing, how can we exploit GPUs as parallel machines?

- John Nickolls, Ian Buck, Michael Garland, and Kevin Skadron. *"Scalable parallel programming with CUDA,"* in *ACM SIGGRAPH 2008 classes (SIGGRAPH '08). ACM, New York, NY, USA, , Article 16 , 14 pages.*
- Shane Ryoo, Christopher I. Rodrigues, Sara S. Baghsorkhi, Sam S. Stone, David B. Kirk, and Wen-mei W. Hwu. *"Optimization principles and application performance evaluation of a multithreaded GPU using CUDA,"* in *Proceedings of the 13th ACM SIGPLAN Symposium on Principles and practice of parallel programming (PPoPP '08).*
- Linchuan Chen and Gagan Agrawal. *"Optimizing MapReduce for GPUs with effective shared memory usage,"* in *Proceedings of the 21st international symposium on High-Performance Parallel and Distributed Computing (HPDC '12).*
- Eddy Z. Zhang, Yunlian Jiang, Ziyu Guo, Kai Tian, and Xipeng Shen. *"On-the-fly elimination of dynamic irregularities for GPU computing,"* in *Proceedings of the sixteenth international conference on Architectural support for programming languages and operating systems (ASPLOS '11).*

Scheduling for Parallelism: Given multicore architectures, etc., what are some new ideas for supporting efficient scheduling of parallel computation?

- Daniel Sanchez, Richard M. Yoo, and Christos Kozyrakis. *"Flexible architectural support for fine-grain scheduling,"* in *Proceedings of the fifteenth edition of ASPLOS on Architectural support for programming languages and operating systems (ASPLOS '10).*

  Background material: Sanjeev Kumar, Christopher J. Hughes, and Anthony Nguyen. *"Carbon: architectural support for fine-grained parallelism on chip multiprocessors,"* in *Proceedings of the 34th annual international symposium on Computer architecture (ISCA '07).*

- Stijn Eyerman and Lieven Eeckhout. *"Probabilistic job symbiosis modeling for SMT processor scheduling,"* in *Proceedings of the fifteenth edition of ASPLOS on Architectural support for programming languages and operating systems (ASPLOS '10).*
- F. Ryan Johnson, Radu Stoica, Anastasia Ailamaki, and Todd C. Mowry. *"Decoupling contention management from scheduling,"* in *Proceedings of the fifteenth edition of ASPLOS on Architectural support for programming languages and operating systems (ASPLOS '10).*
- Sergey Zhuravlev, Sergey Blagodurov, and Alexandra Fedorova. *"Addressing shared resource contention in multicore processors via scheduling,"* in *Proceedings of the fifteenth edition of ASPLOS on Architectural support for programming languages and operating systems (ASPLOS '10).*

Exploiting Heterogeneous Architectures: How can we best take advantage of machines composed of heterogeneous processing units?

- Kenzo Van Craeynest, Aamer Jaleel, Lieven Eeckhout, Paolo Narvaez, and Joel Emer. *"Scheduling heterogeneous multi-cores through Performance Impact Estimation (PIE),"* in *Proceedings of the 39th International Symposium on Computer Architecture (ISCA '12).*
- Ting Cao, Stephen M Blackburn, Tiejun Gao, and Kathryn S McKinley. *"The yin and yang of power and performance for asymmetric hardware and managed software,"* in *Proceedings of the 39th International Symposium on Computer Architecture (ISCA '12).*
- Rachata Ausavarungnirun, Kevin Kai-Wei Chang, Lavanya Subramanian, Gabriel H. Loh, and Onur Mutlu. *"Staged memory scheduling: achieving high performance and scalability in heterogeneous systems,"* in *Proceedings of the 39th International Symposium on Computer Architecture (ISCA '12).*
- Jos A. Joao, M. Aater Suleman, Onur Mutlu, and Yale N. Patt. *"Bottleneck identification and scheduling in multithreaded applications,"* in *Proceedings of the seventeenth international conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '12).*

Table 1: Day 1 Discussion Leaders

Topic	Time Slot	Discussion Leaders
Parallel Prog. Models and Languages	9:00-9:20	Mike Ralph Kun Li Dong Zhou
Exploiting Parallelism on GPUs	9:20-9:40	Matt Mukerjee David Naylor
Scheduling for Parallelism	9:40-10:00	Yixin Luo Rika Nakahara
Exploiting Heterogeneous Architectures	10:00-10:20	Evangelos Papalexakis Alex Beutel John Dickerson

Day 2: *Wednesday, October 17, 2012*

Architectural Support for Security: What can the hardware do to help make systems more secure?

- Mehmet Kayaalp, Meltem Ozsoy, Nael Abu-Ghazaleh, and Dmitry Ponomarev. *"Branch regulation: low-overhead protection from code reuse attacks,"* in *Proceedings of the 39th International Symposium on Computer Architecture (ISCA '12).*
- John Demme, Robert Martin, Adam Waksman, and Simha Sethumadhavan. *"Side-channel vulnerability factor: a metric for measuring information leakage,"* in *Proceedings of the 39th International Symposium on Computer Architecture (ISCA '12).*
- Robert Martin, John Demme, and Simha Sethumadhavan. *"TimeWarp: rethinking time-keeping and performance monitoring mechanisms to mitigate side-channel attacks,"* in *Proceedings of the 39th International Symposium on Computer Architecture (ISCA '12).*
- Jonathan Valamehr, Melissa Chase, Seny Kamara, Andrew Putnam, Dan Shumow, Vinod Vaikuntanathan, and Timothy Sherwood. *"Inspection resistant memory: architectural support for security from physical examination,"* in *Proceedings of the 39th International Symposium on Computer Architecture (ISCA '12).*

Finding and Fixing Software Bugs: What can the system do to help us find and detect bugs in software (especially parallel software)?

- Santosh Nagarakatte, Milo M. K. Martin, and Steve Zdancewic. *"Watchdog: hardware for safe and secure manual memory management and full memory safety,"* in *Proceedings of the 39th International Symposium on Computer Architecture (ISCA '12)".*
- Joseph Devietti, Benjamin P. Wood, Karin Strauss, Luis Ceze, Dan Grossman, and Shaz Qadeer. *"RADISH: always-on sound and complete Race Detection in Software and Hardware,"* in *Proceedings of the 39th International Symposium on Computer Architecture (ISCA '12).*
- Evangelos Vlachos, Michelle L. Goodstein, Michael A. Kozuch, Shimin Chen, Babak Falsafi, Phillip B. Gibbons, and Todd C. Mowry. *"ParaLog: enabling and accelerating online parallel monitoring of multithreaded applications,"* in *Proceedings of the fifteenth edition of ASPLOS on Architectural support for programming languages and operating systems (ASPLOS '10).*
- Wei Zhang, Junghee Lim, Ramya Olichandran, Joel Scherpelz, Guoliang Jin, Shan Lu, and Thomas Reps. *"ConSeq: detecting concurrency bugs through sequential errors,"* in *Proceedings of the sixteenth international conference on Architectural support for programming languages and operating systems (ASPLOS '11).*

Warehouse-Scale Computing: How can we better support computation within large data centers?

- Michael Ferdman, Almutaz Adileh, Onur Kocberber, Stavros Volos, Mohammad Alisafaee, Djordje Jevdjic, Cansu Kaynak, Adrian Daniel Popescu, Anastasia Ailamaki, and Babak Falsafi. *"Clearing the clouds: a study of emerging scale-out workloads on modern hardware,"* in *Proceedings of the seventeenth international conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '12).*
- Lingjia Tang, Jason Mars, Neil Vachharajani, Robert Hundt, and Mary Lou Soffa. *"The impact of memory subsystem resource sharing on datacenter applications,"* in *Proceedings of the 38th annual international symposium on Computer architecture (ISCA '11).*
- Jason Mars, Lingjia Tang, Robert Hundt, Kevin Skadron, and Mary Lou Soffa. *"Bubble-Up: increasing utilization in modern warehouse scale computers via sensible co-locations,"* in *Proceedings of the 44th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO-44 '11).*
- Pejman Lotfi-Kamran, Boris Grot, Michael Ferdman, Stavros Volos, Onur Kocberber, Javier Picorel, Almutaz Adileh, Djordje Jevdjic, Sachin Idgunji, Emre Ozer, and Babak Falsafi. Scale-out processors," in *Proceedings of the 39th International Symposium on Computer Architecture (ISCA '12).*

Optimizing Power and Energy: Power and energy consumption are major constraints on hardware technology these days. What can we do to improve power efficiency?

- Vasileios Kontorinis, Liuyi Eric Zhang, Baris Aksanli, Jack Sampson, Houman Homayoun, Eddie Pettis, Dean M. Tullsen, and Tajana Simunic Rosing. *"Managing distributed UPS energy for effective power capping in data centers,"* in *Proceedings of the 39th International Symposium on Computer Architecture (ISCA '12).*
- Navin Sharma, Sean Barker, David Irwin, and Prashant Shenoy. *"Blink: managing server clusters on intermittent power,"* in *Proceedings of the sixteenth international conference on Architectural support for programming languages and operating systems (ASPLOS '11).*
- Henry Hoffmann, Stelios Sidiroglou, Michael Carbin, Sasa Misailovic, Anant Agarwal, and Martin Rinard. *"Dynamic knobs for responsive power-aware computing,"* in *Proceedings of the sixteenth international conference on Architectural support for programming languages and operating systems (ASPLOS '11).*
- Song Liu, Karthik Pattabiraman, Thomas Moscibroda, and Benjamin G. Zorn. *"Flikker: saving DRAM refresh-power through critical data partitioning,"* in *Proceedings of the sixteenth international conference on Architectural support for programming languages and operating systems (ASPLOS '11).*
- Qingyuan Deng, David Meisner, Luiz Ramos, Thomas F. Wenisch, and Ricardo Bianchini. *"MemScale: active low-power modes for main memory,"* in *Proceedings of the sixteenth international conference on Architectural support for programming languages and operating systems (ASPLOS '11).*

Table 2: Day 2 Discussion Leaders

Topic	Time Slot	Discussion Leaders
Architectural Support for Security	9:00-9:20	Michael Maass Peter Chapman
Finding and Fixing Software Bugs	9:20-9:40	Stefan Muller Andrew Konecki
Warehouse-Scale Computing	9:40-10:00	Kiryong Ha Mu Li
Optimizing Power and Energy	10:00-10:20	Martyn Romanko Lei Fan

Day 3: *Thursday, October 18, 2012*

Caches and Memory Hierarchies: In light of modern multi-core CPUs and memory technologies, how should we design caches and memory hierarchies?

- Zhe Wang, Samira M. Khan, and Daniel A. Jimnez. *"Improving writeback efficiency with decoupled last-write prediction,"* in *Proceedings of the 39th International Symposium on Computer Architecture (ISCA '12).*
- Jaewoong Sim, Jaekyu Lee, Moinuddin K. Qureshi, and Hyesoon Kim. *"FLEXclusion: balancing cache capacity and on-chip bandwidth via flexible exclusion,"* in *Proceedings of the 39th International Symposium on Computer Architecture (ISCA '12).*
- Aamer Jaleel, Kevin B. Theobald, Simon C. Steely, Jr., and Joel Emer. *"High performance cache replacement using re-reference interval prediction (RRIP),"* in *Proceedings of the 37th annual international symposium on Computer architecture (ISCA '10).*
- Gennady Pekhimenko, Vivek Seshadri, Onur Mutlu, Phillip B. Gibbons, Michael A. Kozuch, and Todd C. Mowry. *"Base-delta-immediate compression: practical data compression for on-chip caches,"* in *Proceedings of the 21st international conference on Parallel architectures and compilation techniques (PACT '12).*

Cache Coherence: Can we do a better job of supporting cache coherence, in light of recent architectural trends?

- Ronald G. Dreslinski, Thomas Manville, Korey Sewell, Reetuparna Das, Nathaniel Pinckney, Sudhir Satpathy, David Blaauw, Dennis Sylvester, and Trevor Mudge. *"XPoint cache: scaling existing bus-based coherence protocols for 2D and 3D many-core systems,"* in *Proceedings of the 21st international conference on Parallel architectures and compilation techniques (PACT '12).*
- Alberto Ros and Stefanos Kaxiras. *"Complexity-effective multicore coherence,"* in *Proceedings of the 21st international conference on Parallel architectures and compilation techniques (PACT '12).*
- Blas A. Cuesta, Alberto Ros, Mara E. Gmez, Antonio Robles, and Jos F. Duato. Increasing the effectiveness of directory caches by deactivating coherence for private memory blocks," in *Proceedings of the 38th annual international symposium on Computer architecture (ISCA '11).*

Memory Ordering: Innovations related to memory consistency models and other parallel memory ordering issues.

- Joseph Devietti, Jacob Nelson, Tom Bergan, Luis Ceze, and Dan Grossman. *"RCDC: a relaxed consistency deterministic computer,"* in *Proceedings of the sixteenth international conference on Architectural support for programming languages and operating systems (ASPLOS '11).*
- Abhayendra Singh, Daniel Marino, Satish Narayanasamy, Todd Millstein, and Madan Musuvathi. *"Efficient processor support for DRFx, a memory model with exceptions,"* in *Proceedings of the sixteenth international conference on Architectural support for programming languages and operating systems (ASPLOS '11).*
- Jacob Burnim, George Necula, and Koushik Sen. *"Specifying and checking semantic atomicity for multithreaded programs,"* in *Proceedings of the sixteenth international conference on Architectural support for programming languages and operating systems (ASPLOS '11).*
- Bogdan F. Romanescu, Alvin R. Lebeck, and Daniel J. Sorin. *"Specifying and dynamically verifying address translation-aware memory consistency,"* in *Proceedings of the fifteenth edition of ASPLOS on Architectural support for programming languages and operating systems (ASPLOS '10).*

Transactional Memory: Providing direct hardware support for a transactional programming model would provide some potential advantages. How can we provide this support, and how does it help programmers?

- (*NOTE: the following paper is good background reading for this topic, 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.)
- Jared Casper, Tayo Oguntebi, Sungpack Hong, Nathan G. Bronson, Christos Kozyrakis, and Kunle Olukotun. *"Hardware acceleration of transactional memory on commodity systems,"* in *Proceedings of the sixteenth international conference on Architectural support for programming languages and operating systems (ASPLOS '11).*
- Luke Dalessandro, Franois Carouge, Sean White, Yossi Lev, Mark Moir, Michael L. Scott, and Michael F. Spear. *"Hybrid NOrec: a case study in the effectiveness of best effort hardware transactional memory,"* in *Proceedings of the sixteenth international conference on Architectural support for programming languages and operating systems (ASPLOS '11).*
- Amy Wang, Matthew Gaudet, Peng Wu, Jos Nelson Amaral, Martin Ohmacht, Christopher Barton, Raul Silvera, and Maged Michael. *"Evaluation of blue Gene/Q hardware support for transactional memories,"* in *Proceedings of the 21st international conference on Parallel architectures and compilation techniques (PACT '12).*
- Haris Volos, Andres Jaan Tack, Michael M. Swift, and Shan Lu. *"Applying transactional memory to concurrency bugs,"* in *Proceedings of the seventeenth international conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '12).*

Table 3: Day 3 Discussion Leaders

Topic	Time Slot	Discussion Leaders
Caches and Memory Hierarchies	9:00-9:20	Deby Katz Ameya Ambardekar
Cache Coherence	9:20-9:40	Ross Daly Chan sik Kim
Memory Ordering	9:40-10:00	David Jia Dennis Liang
Transactional Memory	10:00-10:20	Parag Dixit Bruno Vavala