2. Increasing TLB Reach by Exploiting Clustering in Page Translations, *B. Pham, A. Bhattacharjee, Y. Eckert, G. Loh*

3. Supporting x86-64 Address Translation for 100s of GPU Lanes, *J. Power, M. Hill, D. Wood*

9:45-11:00: Session 9B - Power (Plaza F)

1. Scalably Verifiable Dynamic Power Management, *O. Matthews, M. Zhang, D. Sorin*

2. Revolver: Processor Architecture for Power Efficient Loop Execution, *M. Hayenga, M. Lipasti, V. Reddy*

3. Dynamic Management of TurboMode in Modern Multi-core Chips, *D. Lo, C. Kozyrakis*

11:15-12:30: Session 10A - Prefetching and Compression (Plaza E)

1. Spare Register Aware Prefetching for Graph Algorithms on GPUs, *N. Lakshminarayana, H. Kim*

2. Sandbox Prefetching: Safe, Run-Time Evaluation of Aggressive Prefetchers, *S. Pugsley, Z. Chishti, C. Wilkerson, T. Chuang, R. Scott, A. Jaleel, S-L. Lu, K. Chow, R. Balasubramoniam*

3. Memzip: Exploiting Unconventional Benefits from Memory Compression, *A. Shafiee, M. Taassori, R. Balasubramonian, A. Davis*

11:15-12:30: Session 10B - Threading (Plaza F)

1. CDTT: Compiler-generated data-triggered threads, *H-W. Tseng, D. Tullsen*

2. Accelerating Decoupled Look-ahead via Weak Dependence Removal: A Metaheuristic Approach, *R. Parihar, M. Huang*

3. Undersubscribed Threading for High-Performance and Energy-Efficient Many-Core Execution, *W. Heirman, T. Carlson, K. Van Craeynest, I. Hur, A. Jaleel, L. Eeckhout*

12:30-12:45: Conference Closing (Plaza E)

Closing and Best Paper Award

Organizing Committee:

General Chair: *Tao Li (University of Florida)*
Program Chair: *Natalie Enright Jerger (University of Toronto)*
Workshop & Tutorials Chairs: *Satish Narayanasamy (University of Michigan) & Junmin Wu (University of Science and Technology of China)*
Finance Chair: *Xin Fu (University of Kansas)*
Publicity Chairs: *Onur Mutlu (Carnegie Mellon University), Lieven Eeckhout (Ghent University), Rui Wang (Beihang University), Ravishankar Iyer (Intel)*
Local Arrangements Chair: *Jun Wang (University of Central Florida)*
Publications Chair: *Zhibin Yu (Shenzhen Institute of Advanced Technology)*
Registration Chair: *James Poe (Miami Dade College)*
Travel Grant Chairs: *Lu Peng (Louisiana State University) & Xuehai Qian (University of California at Berkeley)*
Industrial Session Chair: *Michael Powell (Intel)*
Submission Chair: *Robert Hesse (University of Toronto)*
Web Chairs: *Chao Li (University of Florida) & Ming Liu (University of Florida)*
Steering Committee: *Lixin Zhang (ICT/Chinese Academy of Sciences), Dean Tullsen (UC, San Diego), David Christie (AMD), David Koppelman (LSU), David Brooks (Harvard), Laxmi Bhuyan (UC Riverside), David Kaeli (Northeastern), Yale Patt (UT Austin), Josep Torrellas (UIUC)*

Tutorials

	Organizers	Date
OC2DS (Blue Spring I)	Jacob Bower, Zivojin Sustran, Oskar Mencer, Veljko Milutinovic	Saturday morning
Manifold (Blue Spring I)	Jun Wang, Sudhakar Yalamanchili, Tom Conte, George Riley	Saturday afternoon
ABDP (Rainbow Spring I)	DK Panda, Xiaoyi Lu	Sunday afternoon

Workshops

	Organizers	Date
REPRODUCE (Blue Spring II)	Alex Jones, Grigori Fursin, Daniel Mosse, Bruce Childers	Saturday morning
EDCS-3 (Blue Spring II)	Rui Hou, Jian Li, Michael C. Huang, Guangdeng Liao	Saturday afternoon
SHAW-5 (Blue Spring I)	Ramesh Illikkal, Ravi Iyer, Raj Yavatkar, Renato Figueiredo	Sunday full-day
HARSH 2014 (Blue Spring II)	Augusto Vega, Alper Buyuktosunoglu, Pradip Bose, Xuan Zhang, David Brooks	Sunday full-day

Sponsored by:

IEEE Computer Society TC on Computer Architecture

Corporate Support:

20th International Symposium on High Performance Computer Architecture

February 15-19, 2014
Orlando, Florida

Sunday February 16, 2014		
6:00-8:00	Welcome Reception (Plaza G)	
Monday February 17, 2014		
8:30-9:50	Welcome and Keynote I: Mark Hill (Florida B+C)	
9:50-10:10	Break	
10:10-10:20	Opening Remarks (Florida B+C)	
10:20-12:00	1A: Caches (Plaza E)	1B: Reliability and Process Variation (Plaza F)
12:00-1:30	Lunch (Plaza G)	
1:30-2:45	2A: Race Detection and Instruction Monitoring (Plaza E)	2B: Data Centers (Plaza F)
2:45-3:15	Break	
3:15-4:30	3A: Coherence & Consistency (Plaza E)	3B: Best of CAL (Plaza F)
4:30-4:45	Break	
4:45-6:25	4A: Security & Cloning (Plaza E)	4B: GPUs (Plaza F)
6:30-8:00	TCCA Business Meeting (Plaza E)	
Tuesday February 18, 2014		
8:05-9:45	5A: Interconnection Networks (Plaza E)	5B: DRAM (Plaza F)
9:45-10:05	Break	
10:05-11:45	Session 6: Best Paper I (Plaza E+F)	
11:45-1:15	Lunch (Plaza G)	
1:15-2:30	Session 7: Best Paper II (Plaza E+F)	
2:30-2:45	Break	
2:45-4:00	8A: Industrial Track (Plaza E)	8B: Non-volatile Mem. (Plaza F)
4:00	Excursion	
Wednesday February 19, 2014		
8:30-9:40	Keynote II: Norm Rubin (Florida B+C)	
9:40-9:45	Room shift	
9:45-11:00	9A: Memory Management (Plaza E)	9B: Power (Plaza F)
11:00-11:15	Break	
11:15-12:30	10A: Prefetching and Compression (Plaza E)	10B: Threading (Plaza F)
12:30-12:45	Closing remarks (Plaza E)	

Monday February 17, 2014

8:30-9:50: Keynote I - Mark Hill, University of Wisconsin (Florida B+C)

21st Century Computer Architecture

10:10-10:20: Opening Remarks (Florida B+C)

10:20-12:00: Session 1A - Caches (Plaza E)

1. Locality-Aware Data Replication in the Last-Level Cache, *G. Kurian, S. Devadas, O. Khan*
2. Adaptive Placement and Migration Policy for an STT-RAM-Based Hybrid Cache, *Z. Wang, D. Jimenez, C. Xu, G. Sun, Y. Xie*
3. DASCA: Dead Write Prediction Assisted STT-RAM Cache Architecture, *J. Ahn, S. Yoo, K. Choi*
4. A Detailed GPU Cache Model Based on Reuse Distance Theory, *C. Nugteren, G.-J. van den Braak, H. Corporaal, H. Bal*

10:20-12:00: Session 1B - Reliability and Process Variation (Plaza F)

1. Precision-Aware Soft Error Protection for GPUs, *D. Palframan, N. Kim, M. Lipasti*
2. Understanding the Impact of Gate-Level Physical Reliability Effects on Whole Program Execution, *R. Balasubramanian, K. Sankaralingam*
3. Accordion: Toward Soft Near-Threshold Voltage Computing, *U. Karpuzcu, I. Akturk, N. Kim*
4. Mosaic: Exploiting the Spatial Locality of Process Variation to Reduce Refresh Energy in On-Chip eDRAM Modules, *A. Agrawal, A. Ansari, J. Torrellas*

1:30-2:45: Session 2A - Race Detection & Instruction Monitoring (Plaza E)

1. Low-Overhead and High Coverage Run-Time Race Detection Through Selective Meta-data Management, *R. Huang, E. Halberg, A. Ferraiuolo, G. E. Suh*
2. FADE: A Programmable Filtering Accelerator for Instruction-Grain Monitoring, *S. Fytraki, E. Vlachos, O. Kocberber, B. Grot, B. Falsafi*
3. Dynamically Detecting and Tolerating IF-Condition Data Races, *S. Qi, A. Muzahid, W. Ahn, J. Torrellas*

1:30-2:45: Session 2B - Data Centers (Plaza F)

1. Exploiting Thermal Energy Storage to Reduce Data Center Capital and Operating Expenses, *W. Zheng, K. Ma, X. Wang*
2. Implications of High Energy Proportional Servers on Cluster-wide Energy Proportionality, *D. Wong, M. Annavaram*
3. Strategies for Anticipating Risk in Heterogeneous System Design, *M. Guervara, B. Lubin, B. Lee*

3:15-4:30: Session 3A - Coherence and Consistency (Plaza E)

1. TSO-CC: Consistency directed cache coherence for TSO, *M. Elver, V. Nagarajan*
2. Stash Directory: A Scalable Directory for Many-Core Coherence, *S. Demetriades, S. Cho*
3. QuickRelease: A Throughput Oriented Approach to Release Consistency on GPUs, *B. Hechtman, B. Beckmann, D. Hower, M. Hill, D. Wood, S. Reinhardt, S. Che, Y. Tian*

3:15-4:30: Session 3B - Best of CAL (Plaza F)

1. The Netflix Challenge: Datacenter Edition, *C. Delimitrou, C. Kozyrakis*
2. High Performance, Energy Efficient Chipkill Correct Memory with Multidimensional Parity, *X. Jian, J. Sartori, H. Duwe, R. Kumar*
3. Shrink-Fit: A Framework for Flexible Accelerator Sizing, *M. Lyons, G-Y. Wei, D. Brooks*
4. Clumsy Flow Control for High-Throughput Bufferless On-Chip Networks, *H. Kim, Y. Kim, J. Kim*

4:45-6:25: Session 4A - Security and Cloning (Plaza E)

1. A Non-Inclusive Memory Permissions Architecture for Protection Against Cross-Layer Attacks, *J. Elwell, R. Riley, N. Abu-Ghazaleh, D. Ponomarev*
2. Suppressing the Oblivious RAM Timing Channel While Making Information Leakage and Program Efficiency Trade-offs, *C. Fletcher, L. Ren, X. Yu, M. Van Dijk, O. Khan, S. Devadas*
3. Timing Channel Protection for Memory Controllers, *Y. Wang, A. Ferraiuolo, G. E. Suh*
4. STM : Cloning the Spatial and Temporal Memory Access Behavior, *A. Awad, Y. Solihin*

4:45-6:25: Session 4B - GPUs (Plaza F)

1. A Scalable Multi-Path Microarchitecture for Efficient GPU Control Flow, *A. ElTantawy, J. Ma, M. O'Connor, T. Aamodt*
2. Improving GPGPU Resource Utilization and Performance Through Alternative Thread Block Scheduling, *M. Lee, S. Song, J. Moon, J. Kim, W. Seo, Y. Cho, S. Ryu*
3. MRPB: Memory Request Prioritization for Massively Parallel Processors, *W. Jia, K. Shaw, M. Martonosi*
4. Warp-Level Divergence in GPUs: Characterization, Impact and Mitigation, *P. Xiang, Y. Yang, H. Zhou*

6:30-8:00: TCCA Business Meeting (Plaza E)

Tuesday February 18, 2014

8:05-9:45: Session 5A - Interconnection Networks (Plaza E)

1. MP3: Minimizing Performance Penalty for Power-gating of Clos Network-on-Chip, *L. Chen, L. Zhao, R. Wang, T. Pinkston*
2. Up By Their Bootstraps: Online Learning in Artificial Neural Networks for CMP Uncore Power Management, *J-Y. Won, X. Chen, P. Gratz, J. Hu, V. Soteriou*
3. QORE: A Fault Tolerant Network-on-Chip Architecture with Power-Efficient Quad-Function Channel (QFC) Buffers, *D. DiTomaso, A. Kodi, A. Louri*
4. Transportation-Network Inspired Network-on-Chip, *H. Kim, G. Kim, H. Yeo, S. Maeng, J. Kim*

8:05-9:45: Session 5B - DRAM (Plaza F)

1. Improving System Throughput and Fairness Simultaneously in CMP Systems via Dynamic Bank Partitioning, *M. Xie, D. Tong, K. Huang, X. Cheng*
2. Improving DRAM Performance by Parallelizing Refreshes with Accesses, *K. Chang, D. Lee, Z. Chishti, C. Wilkerson, A. Alameldeen, Y. Kim, O. Mutlu*
3. CREAM: A Concurrent-Refresh-Aware DRAM Memory System, *T. Zhang, M. Poremba, C. Xu, G. Sun, Y. Xie*
4. DraMon: Predicting Memory Bandwidth Usage of Multi-threaded Programs with High Accuracy and Low Overhead, *W. Wang, T. Dey, J. Davidson, M. L. Soffa*

10:05-11:45: Session 6 - Best Paper I (Plaza E+F)

1. PVCoherence: Designing Flat Coherence Protocols for Scalable Verification, *M. Zhang, J. Bingham, J. Erickson, D. Sorin*
2. Atomic SC for Simple In-order Processors, *D. Gope, M. Lipasti*
3. Concurrent and Consistent Virtual Machine Introspection with Hardware Transactional Memory, *Y. Liu, Y. Xia, H. Guan, B. Zang, H. Chen*
4. Practical Data Value Speculation for Future High-end Processors, *A. Perais, A. Seznec*

1:15-2:30: Session 7 - Best Paper II (Plaza E+F)

1. Tangle: Route-Oriented Dynamic Voltage Minimization for Variation-Afflicted, Energy-Efficient On-Chip Networks, *A. Ansari, A. Mishra, J. Xu, J. Torrellas*
2. Improving Cache Performance by Exploiting Read-Write Disparity, *S. Khan, A. Alameldeen, C. Wilkerson, O. Mutlu, D. Jimenez*
3. NUAT: A Non-Uniform Access Time Memory Controller, *W. Shin, J. Yang, J. Choi, L-S. Kim*

2:45-4:00: Session 8A - Industrial Track (Plaza E)

1. Improving In-Memory Database Index Performance with Intel® Transactional Synchronization Extensions, *T. Karnagel, R. Dementiev, R. Rajwar, K. Lai, T. Legler, B. Schlegel, W. Lehner*
2. BigDataBench: a Big Data Benchmark Suite from Internet Services, *L. Wang, C. Luo, Y. He, J. Zhan, K. Zhan, X. Li, Y. Zhu, S. Zhang, Q. Yang, B. Qiu, Z. Jia*
3. 3D Stacking of High-Performance Processors, *P. Emma, A. Buyuktosunoglu, M. Healy, K. Kailas, V. Puente, R. Yu, A. Hartstein, P. Bose, J. Moreno*

2:45-4:00: Session 8B - Non-volatile memory (Plaza F)

1. Reducing the Cost of Persistence for Nonvolatile Heaps in End User Devices, *S. Kannan, A. Gavrilovska, K. Schwan*
2. Sprinkler: Maximizing Resource Utilization in Many-Chip Solid State Disks, *M. Jung, M. Kandemir*
3. Over-Clocked SSD: Safely Running Beyond Flash Memory Chip I/O Clock Specs, *K. Zhao, K. Venkataraman, X. Zhang, J. Li, N. Zheng, T. Zhang*

4:00-onward: Excursion

Cirque du Soleil (Downtown Disney)

Wednesday February 19, 2014

8:30-9:40: Keynote II - Norm Rubin, NVIDIA (Florida B+C)

Heterogeneous Computing -- what does it mean for compiler research?

9:45-11:00: Session 9A - Memory Management (Plaza E)

1. GPUdmm: A High-Performance and Memory Oblivious GPU Architecture Using Dynamic Memory Management, *Y. Kim, J. Lee, J-E. Jo, J. Kim*