

P A C T 2 0 0 9



The Eighteenth International Conference On Parallel Architectures and Compilation Techniques

Photo by Courtesy of GRCVB/VisitRaleigh.com

C O N F E R E N C E P R O G R A M

Sunday, September 13

6:00-8:00	Opening Reception	Humble Pie Restaurant
-----------	-------------------	-----------------------

Monday, September 14

7:30-8:30	Breakfast	Hobnob/Huddle
8:30-9:00	Welcome	Four Sisters
9:00-10:00	Keynote: Pat Hanrahan. <i>Why are Graphics Systems so Fast?</i>	Four Sisters
10:30-11:30	Software Transactional Memory & Speculation	Four Sisters
	Adaptive Locks: Combining Transactions and Locks for Efficient Concurrency. <i>Takayuki Usui, Yannis Smaragdakis and Reimer Behrends.</i> Anaphase: A Fine-Grain Thread Decomposition Scheme for Speculative Multithreading. <i>Carlos Madriles, Pedro Lopez, Josep M. Codina, Eric Gibert, Fernando Latorre, Alejandro Martinez, Raul Martinez and Antonio Gonzalez.</i>	
11:30-12:00	ACM Student Research Competition	Clambake
12:00-1:30	Lunch	Hobnob/Huddle
1:30-3:00	Best Papers	Four Sisters
	Characterizing the TLB Behavior of Emerging Parallel Workloads on Chip Multiprocessors. <i>Abhishek Bhattacharjee and Margaret Martonosi.</i>	
	Interprocedural Load Elimination for Dynamic Optimization of Parallel Programs. <i>Rajkishore Barik and Vivek Sarkar.</i>	
3:30-4:30	Accelerators	Four Sisters
	Algorithmic Skeletons within an Embedded Domain Specific Language for the CELL Processor. <i>Tarik Saidani, Joel Falcou, Claude Tadonki, Lionel Lacassagne and Daniel Etienneble.</i> A Task-centric Memory Model for Scalable Accelerator Architectures. <i>John Kelm, Daniel Johnson, Steven Lumetta, Matthew Frank and Sanjay Patel.</i>	
5:30-6:30	ACM Student Research Competition Presentations	Four Sisters
6:30-7:30	Poster Reception	Clarion: Top of the Tower Restaurant

Tuesday, September 15 (morning)

7:30-8:30	Breakfast	Hobnob/Huddle
8:30-9:30	Keynote: Frederick H. Streitz. <i>Pushing the Limits of Scientific Computing on Extreme Platforms</i>	Four Sisters
9:30-10:00	Awards	Four Sisters
10:30-12:00	Power and Energy	Four Sisters
	SHIP: Scalable Hierarchical Power Control for Large-Scale Data Centers. <i>Xiaorui Wang, Ming Chen, Charles Lefurgy and Tom Keller.</i>	
	Exploring Phase Change Memory and 3D Die-Stacking for Power/Thermal Friendly, Fast and Durable Memory Architectures. <i>Wangyuan Zhang and Tao Li.</i>	
	Core-Selectability in Chip Multiprocessors. <i>Hashem H. Najaf-abadi, Niket K. Choudhary and Eric Rotenberg.</i>	

W E L C O M E !

It is our pleasure to welcome you to the 18th International Conference on Parallel Architectures and Compilation Techniques (PACT 2009) in beautiful Raleigh, North Carolina. PACT is a multi-disciplinary conference series that brings together researchers from the hardware and software areas to present ground-breaking research related to parallel systems. The technical program includes talks for 35 high quality papers, selected from 188 submissions, covering a wide area of topics. Additionally, the program includes two keynote presentations (see below), four workshops, two tutorials, a separate poster session, as well as an ACM student research competition.

We particularly thank ACM, IEEE, and IFIP for their continued sponsorship of PACT, NSF for providing student travel support, as well as the companies providing financial support for this year's conference: Intel, IBM, Google and Reservoir Labs.


Social events include a Sunday night reception, a poster reception on Monday evening, a banquet on Tuesday evening, as well as plenty of time for socializing and discussing on-going research. In addition, the Raleigh area offers a diversity of tourist activities and attractions. We hope you find the technical program enlightening and enjoy your visit to PACT and the Raleigh area.

K E Y N O T E S P E A K E R S

Pat Hanrahan is the CANON Professor of Computer Science and Electrical Engineering at Stanford University, where he teaches computer graphics. His current research involves visualization, image synthesis, virtual worlds, and graphics systems and architectures. Before joining Stanford he was a faculty member at Princeton. Previously, he developed volume rendering software and was the chief architect of the RenderMan Interface at Pixar. Professor Hanrahan has received two Academy Awards for Science and Technology, the Spirit of America Creativity Award, the SIGGRAPH Computer Graphics Achievement Award, the SIGGRAPH Stephen A. Coons Award, and the IEEE Visualization Career Award. He was recently elected to the National Academy of Engineering, to the American Academy of Arts and Sciences, and is a Fellow of the ACM.

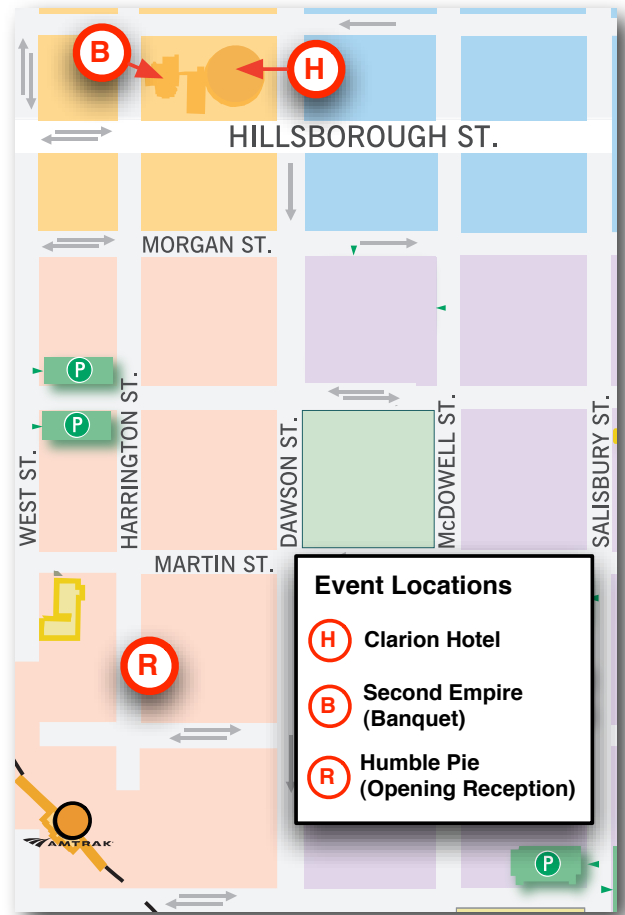


Frederick H. Streitz is Director of the Institute for Scientific Computing Research and Group Leader in the Modeling and Simulations Group at Lawrence Livermore National Laboratory. He joined LLNL's Physical and Life Sciences Directorate as a computational physicist in 1999. He has been active as both an experimentalist and a computational physicist, with recent focus on developing supercomputer applications that push the limits of leadership-class computational capability to address forefront scientific problems. He has twice (2005 and 2007) led multi-institutional teams that were recognized with a Gordon Bell Prize for significant achievement in supercomputing.



Tuesday, September 15 (afternoon)

12:00-1:30	Lunch	Hobnob/Huddle
1:30-3:00	Tools and Testing	Clambake
	Chainsaw: Using Binary Matching for Relative Instruction Mix Comparison. <i>Tipp Moseley, Dirk Grunwald and Ramesh Peri.</i>	
	tm_db: A Generic Debugging Library for Transactional Programs. <i>Yossi Lev and Maurice Herlihy.</i>	
1:30-3:00	StealthTest: Low Overhead Online Software Testing using Transactional Memory. <i>Jayaram Bobba, Weiwei Xiong, Luke Yen, Mark Hill and David Wood.</i>	
	Innovative Hardware	Four Sisters
	CPROB: Checkpoint Processing with Opportunistic Minimal Recovery. <i>Andrew Hilton, Neeraj Eswaran and Amir Roth.</i>	
3:30-5:00	Architecture Support for Improving Bulk Memory Copying and Initialization Performance. <i>Xiaowei Jiang, Yan Solihin, Li Zhao and Ravishanker Iyer.</i>	
	Oblivious Routing on On-Chip Bandwidth-Adaptive Networks. <i>Myong Hyon Cho, Mieszko Lis, Keun Sup Shim, Michel Kinsy, Tina Wen and Srinivas Devadas.</i>	
	Scheduling and Adaptation	Clambake
3:30-5:00	Exploiting Parallelism with Dependence-Aware Scheduling. <i>Xiaotong Zhuang, Alexandre Eichenberger, Yangchun Luo, Kevin O'Brien and Kathryn O'Brien.</i>	
	ITCA: Inter-Thread Conflict-Aware CPU Accounting for CMPs. <i>Carlos Luque, Miquel Moreto, Francisco J. Cazorla, Roberto Gioiosa, Alper Buyuktosunoglu and Mateo Valero.</i>	
	Flexstream: Adaptive Compilation of Streaming Applications for Heterogeneous Architectures. <i>Amir Hormati, Yoonseo Choi, Manjunath Kudlur, Rodric Rabbah, Trevor Mudge and Scott Mahlke.</i>	
6:30-9:30	Novel Cache Systems	Four Sisters
	DDCache: Decoupled and Delegable Cache Data and Metadata. <i>Hemayet Hossain, Sandhya Dwarkadas and Michael C. Huang.</i>	
	Zero-Value Caches: Cancelling Loads that Return Zero. <i>Mafijul Md Islam and Per Stenstrom.</i>	
6:30-9:30	Improving Hardware Cache Performance Through Software-Controlled Object-Level Cache Partitioning. <i>Qingda Lu, Jiang Lin, Xiaoning Ding, Zhao Zhang, Xiaodong Zhang and P. Sadayappan.</i>	
	Banquet	Second Empire Restaurant



Wednesday, September 16

7:30-8:30	Breakfast	Hobnob/Huddle
8:30-10:00	Modeling and Evaluation	Clambake
	Memory Performance and Cache Coherency Effects on an Intel Nehalem Multiprocessor System. <i>Daniel Molka, Daniel Hackenberg, Robert Schöne and Matthias S. Müller.</i>	
	Automatic Tuning of Discrete Fourier Transforms Driven by Analytical Modeling. <i>Basilio Fraguela, Yevgen Voronenko and Markus Puschel.</i>	
8:30-10:00	Analytical Modeling of Pipeline Parallelism. <i>Angeles Navarro, Rafael Asenjo, Siham Tabik and Calin Cascaval.</i>	
	Hardware Transactional Memory	Four Sisters
	FASTM: A Log-based Hardware Transactional Memory with Fast Abort Recovery. <i>Marc Lupon, Grigoris Magklis and Antonio Gonzalez.</i>	
10:30-12:00	Improving Signatures by Locality Exploitation for Transactional Memory. <i>Ricardo Quislan, Eladio Gutierrez and Oscar Plata.</i>	
	Mapping Out a Path from Hardware Transactional Memory to Speculative Multithreading. <i>Leo Porter, Bumyong Choi and Dean Tullsen.</i>	
	Compiler Optimizations	Clambake
10:30-12:00	Polyhedral-Model Guided Loop-Nest Auto-Vectorization. <i>Konrad Trifunovic, Ayal Zaks, Albert Cohen and Dorit Nuzman.</i>	
	Data Layout Transformation for Enhancing Locality on NUMA Chip Multiprocessors. <i>Qingda Lu, Christophe Alias, Uday Bondhugula, Sriram Krishnamoorthy, J. Ramanujam, Atanas Rountev, P. Sadayappan, Yongjian Chen, Haibo Lin and Tin-fook Ngai.</i>	
	Region based Structure Layout Optimization by Selective Data Copying. <i>Sandya Mannarswamy, Ramaswamy Govindarajan and Rishi Surendran.</i>	
10:30-12:00	Cache Management	Four Sisters
	SOS: A Software-Oriented Distributed Shared Cache Management Approach for Chip Multiprocessors. <i>Lei Jin and Sangyeun Cho.</i>	
	Using Aggressor Thread Information to Improve Shared Cache Management for CMPs. <i>Wanli Liu and Donald Yeung.</i>	
10:30-12:00	Cache Sharing Management for Performance Fairness in Chip Multiprocessors. <i>Xing Zhou, Wenguang Chen and Weimin Zheng.</i>	

WORKSHOPS & TUTORIALS

Saturday, September 12

7:30-8:30	Breakfast	Hobnob
8:30-12:00	Tutorial: GPU	Clambake
	Workshop: PMEA	Four Sisters
12:00-1:30	Lunch	Hobnob
1:30-5:00	Tutorial: GPU	Clambake
	Workshop: PMEA	Four Sisters

Sunday, September 13

7:30-8:30	Breakfast	Hobnob
8:30-12:00	Workshop: WPABA	Clambake
	Workshop: MEDEA	Four Sisters
12:00-1:30	Lunch	Hobnob
1:30-5:00	Tutorial: PGAS	Clambake
	Workshop: PACE	Four Sisters