HPCA/CGO/PPoPP 2017 Program Schedule (Feb 4-8, 2017), Austin, TX

Saturday Feb 4, 2017 (Workshops and Tutorials)

7:30–8:30am: Breakfast (Salon F - 6th Floor) and 10:00–10:30am: Break (Salon F - 6th Floor)		
HPCA	***CGO***	***PPoPP***
PHOTONICS: Photonics-Optics Technology Oriented Networking, Information, and Computing Systems (414) **SCAW 2017: Sensors to Cloud Architectures Workshop (417A)** **W3P: First Workshop on Pioneering Processor Paradigms (415B)**	**LLVM workshop (Salon AB - 4th Floor)** **DCAHPO: Dynamic Compilation for Architectural Heterogeneity and Program Optimization (614)**	**WCIRE 2017 - Workshop on Compiler Infrastructures for Research and Education (417B)** **Heterogeneous CPU+GPU Computing: Models, Tools, and Applications (617)** **CPUs, GPUs, FPGAs: Managing the alphabet soup with Intel Threading Building Blocks (615A)**
12:00–1:30pm: Lunch (Salon F - 6th Floor) and 3:00–3:30pm: Break (Salon GH Foyer - 6th Floor)		
PHOTONICS: Photonics-Optics Technology Oriented Networking, Information, and Computing Systems (414) **SCAW 2017: Sensors to Cloud Architectures Workshop (417A)** **W3P: First Workshop on Pioneering Processor Paradigms (415B)**	**LLVM workshop (Salon AB - 4th Floor)** **RWDSL: Real World Domain Specific Languages (619)**	**WCIRE 2017 - Workshop on Compiler Infrastructures for Research and Education (417B)** **Heterogeneous CPU+GPU Computing: Models, Tools, and Applications (617)**
5:00pm: End		

Sunday Feb 5, 2017 (Workshops and Tutorials)

7:30–8:30am: Breakfast (Salon F - 6th Floor) and 10:00–10:30am: Break (Salon F - 6th Floor)		
HPCA	***CGO***	***PPoPP***
Accelerating Big Data Processing with Hadoop, Spark, and Memcached on Datacenters with Modern	**SAE and BigBench (616A)**	**GPGPU: General-Purpose GPU (614)**

HPCA/CGO/PPoPP 2017 Program Schedule (Feb 4-8, 2017), Austin, TX

Architectures (415A) **Learning gem5 Tutorial and Coding Sprint (615B)** **HiPINEB 2017: The 3rd IEEE International Workshop on High-Performance Interconnection Networks in the Exascale and Big-Data Era (417A)**	**DynamoRIO (616B)** **CCTLib (615A)** **COSMIC: Code Optimizations for Multi- and Many-Cores (619)** **Compiler Construction (404)**	**PMAM: Programming Models and Applications for Multicores and Manycores (415B)** **Ushering OpenMP* Parallelization and Vectorization Forward in LLVM Compilers (414)**
12:00–1:30pm: Lunch (Salon F - 6th Floor) and 3:00–3:30pm: Break (Salon GH Foyer - 6th Floor)		
Learning gem5 Tutorial and Coding Sprint (615B) **HiPINEB 2017: The 3rd IEEE International Workshop on High-Performance Interconnection Networks in the Exascale and Big-Data Era (417A)** **An Introduction to OpenPiton, a Manycore Open Source Processor (617)**	**SAE and BigBench (616A)** **COSMIC: Code Optimizations for Multi- and Many-Cores (619)** **Updates in Heterogeneous Compute (616B)** **Compiler Construction (404)**	**GPGPU: General-Purpose GPU (614)** **PMAM: Programming Models and Applications for Multicores and Manycores (415B)** **TRANSACT: Transactional Computing (417B)** **Advanced MPI (615A)** **Latest Developments of OpenMP 4.0 & 4.5 including OpenMP Offload Model (415A)** **PGAS and Hybrid MPI+PGAS Programming Models on Modern HPC Clusters with Accelerators (414)**
5:00pm: End		
6pm: HPCA/CGO/PPoPP Welcome Reception and Poster Session (Salon H - 6th Floor)		

Monday Feb 6, 2017 (Main Program)

7:30-8:30am: Breakfast (Salon H Foyer)
8:30-8:45am: Opening (Salon H - 6th Floor)
8:45-9:55am (Salon H - 6th Floor) – Keynote: ***Guy Steele (Oracle Labs): It's Time for a New Old Language***

HPCA/CGO/PPoPP 2017 Program Schedule (Feb 4-8, 2017), Austin, TX

9:55-10:20am: Break (Salon H Foyer)		
HPCA	***CGO***	***PPoPP***
10:20-11:45am (Salon FG - 6th Floor) *HPCA Session 1: Lightning Rounds* *Session Chair: Daniel A. Jiménez (Texas A&M)*	*10:20-11:45am (Salon J - 6th Floor)* *CGO Session 1: Shared Memory* *Session Chair: Evelyn Duesterwald (IBM)* **Legato: End-to-End Bounded Region Serializability Using Commodity Hardware Transactional Memory** **Automatic Detection of Extended Data-Race-Free Regions** **FinePar: Irregularity-Aware Fine-Grained Workload Partitioning on Integrated Architectures**	*10:20-11:45am (400/402)* *PPoPP Session 1: GPU I* *Session Chair: Keshav Pingali (UT Austin)* **EffiSha: A Software Framework for Enabling Efficient Preemptive Scheduling of GPU** **Layout Lock: A Scalable Locking Paradigm for Concurrent Data Layout Modifications** **Understanding the GPU Microarchitecture to Achieve Bare-Metal Performance Tuning**
11:45-1:15pm: Lunch (Salon H - 6th Floor)		
1:15-2:55pm (Salon FG - 6th Floor) *HPCA Session 2: Best Paper Nominees* *Session Chair: Yale N. Patt (UT Austin)* **Towards Pervasive and User Satisfactory CNN across GPU Microarchitectures** **Near-Optimal Access Partitioning for Memory Hierarchies with Multiple Heterogeneous Bandwidth Sources** **NCAP: Network-Driven, Packet Context-Aware Power Management for Client-Server Architecture** **Supporting Address Translation for Accelerator-Centric Architectures**	*1:15-2:55pm (Salon J - 6th Floor)* *CGO Session 2: GPU Optimization* *Session Chair: Naveen Kumar (Google)* **TwinKernels: An Execution Model to Improve GPU Hardware Scheduling at Compile Time** **Taming Warp Divergence** **Dynamic Buffer Overflow Detection for GPGPUs** **Lift: A Functional Data-Parallel IR for**	*1:15-2:55pm (400/402)* *PPoPP Session 2: Concurrency* *Session Chair: Michael Scott (Univ. of Rochester)* **Checking Concurrent Data Structures Under the C/C++11 Memory Model** **Hierarchical MCS Locks with Timeout** **Contention in Structured Concurrency: Provably Efficient Dynamic Non-Zero Indicators for Nested Parallelism**

HPCA/CGO/PPoPP 2017 Program Schedule (Feb 4-8, 2017), Austin, TX

		High-Performance GPU Code Generation	**Noise Injection Techniques for Reproducing Subtle and Unintended Message Races**
2:55-3:15pm: Break (Salon H Foyer)			
3:15-4:55pm (Salon F - 6th Floor) *HPCA Session 3A: Industrial Session* *Session Chair: Chris Wilkerson (Intel)* **Vulnerabilities in MLC NAND Flash Memory Programming: Experimental Analysis, Exploits, and Mitigation Techniques** **Defect Analysis and Cost Effective Resilience Architecture for Future DRAM Devices** **Architecting an Energy Efficient DRAM System for GPUs** **Design and Analysis of an APU for Exascale Computing** **BRAVO: Balanced Reliability Aware Voltage Optimization**	*3:15-4:55pm (Salon G - 6th Floor)* *HPCA Session 3B: Cache* *Session Chair: Paul Gratz (Texas A&M)* **Maximizing Cache Performance Under Uncertainty** **SWAP: Effective Fine-Grain Management of Shared Last-Level Caches with Minimum Hardware Support** **A Split Cache Hierarchy for Enabling Data-oriented Optimizations** **Fast and Accurate Exploration of Multi-Level Caches Using Hierarchical Reuse Distance**	*3:15-4:55pm (Salon J - 6th Floor)* *CGO Session 3: Best Paper Nominees* *Session Chair: Aaron Smith (Microsoft)* **Synthesizing Benchmarks for Predictive Modeling** **Formalizing the Concurrency Semantics of an LLVM Fragment** **ThinLTO: Scalable and Incremental LTO** **Automatic Generation of Fast BLAS3-GEMM: A Portable Compiler Approach**	*3:15-4:55pm (400/402)* *PPoPP Session 3: Tools* *Session Chair: Milind Chabbi (HPE)* **Thread Data Sharing in Cache: Theory and Measurement** **Exploiting Vector and Multicore Parallelism for Recursive Data- and Task-Parallel Programs** **Isoefficiency in Practice: Configuring and Understanding the Performance of Task-based Applications** **Processor-Oblivious Record and Replay**
4:55-5:15pm: Break (Salon H Prefunction)			
5:15-6:55pm (Salon F - 6th Floor) *HPCA Session 4A: Power, Energy & Large-Scale Computing* *Session Chair: Benjamin Lee (Duke)* **Enabling Effective Module-oblivious Power Gating for Embedded Processors**	*5:15-6:55pm (Salon G - 6th Floor)* *HPCA Session 4B: Memory* *Session Chair: Mike Ferdman (Stony Brook)* **Tiny Directory: Efficient Shared Memory in Many-core Systems with Ultra-low-overhead Coherence Tracking**	*5:15-6:15pm (Salon J - 6th Floor)* *CGO ACM Student Research Competition (SRC) Presentations* *Session Chair: Ramesh Peri (Intel)*	*5:15-5:45pm (400/402)* *CGO and PPoPP Joint Session: Artifact Evaluation Discussion*

HPCA/CGO/PPoPP 2017 Program Schedule (Feb 4-8, 2017), Austin, TX

Application-Specific Performance-Aware Energy Optimization on Android Mobile Devices **Fast decentralized power capping for Server Clusters** **Random Folded Clos Topologies for Datacenter Networks**	**Partial Row Activation for Low-Power DRAM System** **Understanding and Optimizing Power Consumption in Memory Networks** **SoftMC: A Flexible and Practical Open-Source Infrastructure for Enabling Experimental DRAM Studies**		
7:30-8:30pm (Salon F - 6th Floor): *HPCA Business Meeting*		*6:30-7:30pm (Salon J - 6th Floor): CGO Business Meeting*	*6:30-7:30pm (400/402):* *PPoPP Business Meeting*

Tuesday Feb 7, 2017 (Main Program)

7:30-8:00am: Breakfast (Salon H Prefunction - 6th Floor)			
HPCA		***CGO***	***PPoPP***
8:00-9:40am (Salon F - 6th Floor) *HPCA Session 5A: NOC* *Session Chair: Vijay Nagarajan (University of Edinburgh)* **Static Bubble: A Framework for Deadlock-free Irregular On-chip Topologies** **Designing Low-power, Low-latency Networks-on-Chip by Optimally Combining Electrical and Optical Links** **Near-Ideal Networks-on-Chip for Servers** **Design and Evaluation of AWGR-based**	*8:00-9:40am (Salon G - 6th Floor)* *HPCA Session 5B: Security* *Session Chair: Calvin Lin (UT Austin)* **Secure Dynamic Memory Scheduling Against Timing Channel Attacks** **Cold Boot Attacks are Still Hot: Security Analysis of Memory Scramblers in Modern Processors** **Cooperative Path-ORAM for Effective Memory Bandwidth Sharing in Server Settings** **Camouflage: Memory Traffic Shaping to Mitigate Timing Attacks**	*8:25-9:40am (Salon J - 6th Floor)* *CGO Session 4: Memory Dependencies* *Session Chair: Ayal Zaks (Intel)* **Pointer Disambiguation via Strict Inequalities** **A Collaborative Dependence Analysis Framework** **Characterizing Data Organization Effects on Heterogeneous Memory Architectures**	*8:00-9:40am (400/402)* *PPoPP Session 4: GPU II* *Session Chair: Angelina Lee (Washington Univ. St. Louis)* **Model-based Iterative CT Image Reconstruction on GPUs** **Combining SIMD and Many/Multi-core Parallelism for Finite State Machines with Enumerative Speculation** **S-Caffe: Co-designing MPI Runtimes and Caffe for Scalable Deep Learning on Modern GPU Clusters** **Simple, Accurate, Analytical Time**

HPCA/CGO/PPoPP 2017 Program Schedule (Feb 4-8, 2017), Austin, TX

Photonic NoC Architectures for 2.5D Integrated High Performance Computing Systems			**Modeling and Optimal Tile Size Selection for GPGPU Stencils**
9:40-10:05am: Break (Salon H Foyer)			
10:05-11:45am (Salon F - 6th Floor) *HPCA Session 6A: Emerging Storage* *Session Chair: Samira Khan (University of Virginia)* **SILC-FM: Subblocked InterLeaved Cache-Like Flat Memory Organization** **ATOM: Atomic Durability in Non-volatile Memory through Hardware Logging** **KAML: A Flexible, High-Performance Key-Value SSD** **Balancing Performance and Lifetime of MLC PCM by Using a Region Retention Monitor**	*10:05-11:45am (Salon G - 6th Floor)* *HPCA Session 6B: Scheduling* *Session Chair: Miquel Pericàs (Chalmers)* **Reliability-Aware Scheduling on Heterogeneous Multicore Processors** **Hipster: Hybrid Task Manager for Latency-Critical Cloud Workloads** **Cooper: Task Colocation with Cooperative Games** **MemPod: A Clustered Architecture for Efficient and Scalable Migration in Flat Address Space Multi-Level Memories**	*10:05-11:45am (Salon J - 6th Floor)* *CGO Session 5: Accelerators & Binary Translation* *Session Chair: Milind Chabbi (HP)* **Clairvoyance: Look-Ahead Compile-time Scheduling** **Phase-Aware Optimization in Approximate Computing** **A Space- and Energy-Efficient Code Compression/Decompression Technique for Coarse-Grained Reconfigurable Architectures** **Cross-ISA Machine Emulation for Multicores**	*10:05-11:45am (400/402)* *PPoPP Session 5: Best Paper Nominees* *Session Chair: Lawrence Rauchwerger (Texas A&M Univ.)* **Pagoda: Fine-Grained GPU Resource Virtualization for Narrow Tasks** **Groute: An Asynchronous Multi-GPU Programming Model for Irregular Computations** **Tapir: Embedding Fork-Join Parallelism into LLVM's Intermediate Representation** **A Multicore Path to Connectomics-on-Demand**
11:45am-1:15pm: Lunch (Salon H - 6th Floor)			
1:15-2:25pm (Salon H - 6th Floor) – Keynote: ***Steve Keckler (Nvidia): Everyone Needs High Performance Computing***			
2:25-2:50pm: Break (Salon H Prefunction)			
2:50-4:30pm (Salon F - 6th Floor) *HPCA Session 7A: Novel Architectures* *Session Chair: Carole-Jean Wu (Arizona State University)*	*2:50-4:30pm (Salon G - 6th Floor)* *HPCA Session 7B: Control-Flow and Microarchitecture* *Session Chair: Daniel A. Jiménez (Texas A&M)*	*2:50-4:30pm (Salon J - 6th Floor)* *CGO Session 6: Feedback Directed and Whole Program Optimization* *Session Chair: Alexandra Jimborean (Uppsala)*	*2:50-4:30pm (400/402)* *PPoPP Session 6: Languages & Compilers* *Session Chair: Saday Sadayppan (Ohio State University)*

HPCA/CGO/PPoPP 2017 Program Schedule (Feb 4-8, 2017), Austin, TX

Exploring Hyperdimensional Associative Memory **GraphPIM: Enabling Instruction-Level PIM Offloading in Graph Computing Frameworks** **High-Bandwidth Low-Latency Approximate Interconnection Networks** **Compute Caches**	**Boomerang: A Metadata-Free Architecture for Control Flow Delivery** **PABST: Proportional Allocation of Bandwidth at the Source and Target** **SOUP-N-SALAD: Allocation-oblivious Access Latency Reduction with Asymmetric DRAM Microarchitectures** **Transparent and Efficient CFI Enforcement with Intel Processor Trace**	**Incremental Whole Program Optimization and Compilation** **Optimizing Function Placement for Large-Scale Data-Center Applications** **Minimizing the Cost of Iterative Compilation with Active Learning** **Removing Checks in Dynamically Typed Languages through Efficient Profiling**	**SC-Haskell: Sequential Consistency in Languages that Minimize Mutable Shared Heap** **Synchronized-by-Default Concurrency for Shared Memory Systems** **Function Call Re-Vectorization** **Optimizing the Four-Index Integral Transform Using Data Movement Lower Bounds Analysis**
5:00 – 9:30pm: ***Excursion: Salt Lick BBQ (Vegetarians Welcome!)*** *Buses depart at 5pm and return at 9:30pm*			

Wednesday Feb 8, 2017 (Main Program)

7:30-8:15am: Breakfast (Salon H Foyer - 6th Floor)			
8:15-9:25am (Salon H - 6th Floor) – Keynote: ***Frank Seide (Microsoft): The Computer Science Behind the Microsoft Cognitive Toolkit -- an Open Source Large-Scale Deep Learning Toolkit for Windows and Linux***			
9:25-9:50am: Break (Salon H Foyer - 6th Floor)			
HPCA	***HPCA***	***CGO***	***PPoPP***
9:50-11:30am (Salon F - 6th Floor) *HPCA Session 8A: Accelerators* *Session Chair: Akanksha Jain (UT Austin)* **PipeLayer: A Pipelined ReRAM-Based Accelerator for Deep Learning** **FlexFlow: A Flexible Dataflow**	*9:50-11:30am (Salon G - 6th Floor)* *HPCA Session 8B: GPU Power & Energy* *Session Chair: David Kaeli (Northeastern)* **Pilot Register File: Energy Efficient Register File for GPUs** **G-Scalar: Cost-Effective Generalized**	*9:50-11:30am (Salon J - 6th Floor)* *CGO Session 7: Reductions & Loops* *Session Chair: Michael Laurenzano (Michigan)* **Discovery and Exploitation of General Reductions: A Constraint Based Approach**	*9:50-11:30am (400/402)* *PPoPP Session 7: Data Analytics* *Session Chair: Sam Midkiff (Purdue)* **Using Butterfly-Patterned Partial Sums to Draw from Discrete Distributions** **KiWi: A Key-Value Map for Scalable Real-Time Analytics**

HPCA/CGO/PPoPP 2017 Program Schedule (Feb 4-8, 2017), Austin, TX

Accelerator Architecture for Convolutional Neural Network **Needle : Leveraging Program Analysis to Analyze and Extract Accelerators from Whole Programs** **Radiation-Induced Error Criticality in Modern HPC Parallel Accelerators**	**Scalar Execution Architecture for Power-Efficient GPUs** **Dynamic GPGPU Power Management using Adaptive Model Predictive Control**	**Parallel Associative Reductions in Halide** **Optimistic Loop Optimization** **Software Prefetching for Indirect Memory Accesses**	**Grammar-aware Parallelization for Scalable XPath Querying** **Eunomia: Scaling Concurrent Search Trees under Contention Using HTM**
11:30-11:45am: Break (Salon H Prefunction - 6th Floor)			
11:45am-1:00pm (Salon F - 6th Floor) *HPCA Session 9A: Best of CAL* *Session Chair: Nam Sung Kim (Illinois)* **Hardware Support for Privacy** **Efficient Execution of Bursty Applications** **Non-intrusive Persistence with a Backend NVM Controller**	*11:45am-1:00pm (Salon G - 6th Floor)* *HPCA Session 9B: GPU* *Session Chair: Abdullah Muzahid (UT San Antonio)* **Efficient Sequential Consistency in GPUs with Relativistic Cache Coherence** **Processing-in-Memory Enabled Graphics Processors for 3D Rendering** **Controlled Kernel Launch for Dynamic Parallelism in GPUs**	*11:45-12:30am (Salon J - 6th Floor):* *CGO Closing & Best Paper Award Announcement*	*11:45am-12:35pm (400/402)* *PPoPP Session 8: Fault Tolerance* *Session Chair: E.N. Elnozahy (KAUST)* **Self-Checkpoint: An In-Memory Checkpoint Method Using Less Space and Its Practice on Fault-Tolerant HPL** **Silent Data Corruption Resilient Two-sided Matrix Factorizations**
1:00pm-1:15pm (Salon F - 6th Floor): *HPCA Closing & Best Paper Award Announcement*			*12:35-12:50pm (400/402)* *PPoPP Closing & Best Paper Award Announcement*

HPCA/CGO/PPoPP 2017 Program Schedule (Feb 4-8, 2017), Austin, TX

4th floor layout

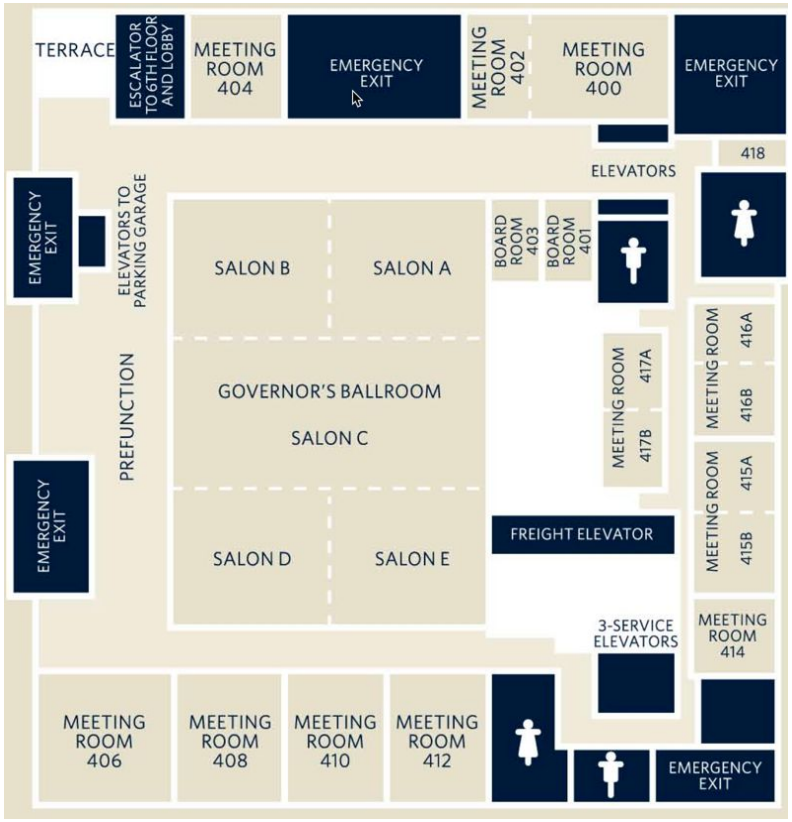

6th floor layout