Saturday February 24th, 2018

HPCA		CGO		PPoPP	CC
[08:00 - 18:15] **Registration**					
[08:30 - 10:00] Room: Europa 3 AACBB: Accelerator Architecture in Computational Biology and Bioinformatics **Opening Remarks** **Keynote 1: "Accelerating Genome Analysis: A Primer on an Ongoing Journey"** Onur Mutlu (ETH, CMU) **Exploring Speed/Accuracy Trade-offs** **Accelerating Duplicate Marking In The Cloud**	[08:30 - 10:00] Room: Europa 7 HIPINEB: High-Performance Interconnection Networks in the Exascale and Big-Data Era **Opening** **Keynote: "The three L's in modern high-performance networking: Low latency, Low cost, Low processing load"**	[09:15 - 10:00] Room: Europa 2 LLVM Performance Workshop **How to Evaluate "In-Memory Computing" Performances without Hardware Measurements?**	[09:15 - 10:00] Room: Europa 6 RWDSL'18: 3rd International Workshop on Real World Domain Specific Languages **Welcome** **Industrial Experience with the Migration of Legacy Models using a DSL**	[08:30 - 10:00] Room: Europa 5 WPMVP: Workshop on Programming Models for SIMD/Vector Processing **Keynote TBA** **Vectorization of a spectral finite-element numerical kernel (Application)**	[08:30 - 08:45] Room: Europa 1 CC: International Conference on Compiler Construction Compiler Construction **Opening** [08:45 - 10:00] Room: Europa 1 CC Keynote **Rethinking Compilers in the Rise of Machine Learning and AI** Xipeng Shen (North Carolina State University, USA)
[10:00 - 10:30] **Coffee Break with Snack**					
[10:30 - 12:10] Room: Europa 3 AACBB: Accelerator Architecture in Computational Biology and Bioinformatics **Invited Talk: "Next Generation Sequencing: Big Data meets High Performance Computing Architectures"** Bertil Schmidt (JGU Mainz) **GAME: GPU Acceleration of Metagenomics Clustering** **Exact Alignment with FM-index on the Intel Xeon Phi Knights Landing Processor** **Optimizations of Sequence Alignment on FPGA: A Case Study of Extended Sequence Alignment**	[10:30 - 12:00] Room: Europa 7 HIPINEB Technical Session 1 (research papers) **Analysis and improvement of Valiant routing in low-diameter networks** **Node-type-based load-balancing routing for Parallel Generalized Fat-Trees** **Analyzing topology parameters for achieving energy-efficient k-ary n-cubes**	[10:30 - 12:00] Room: Europa 2 LLVM Performance Workshop **Optimizing LLVM IR for Guided Vectorization** **Efficient use of memory by reducing size of AST dumps in cross file analysis by clang static analyzer**	[10:30 - 11:50] Room: Europa 6 RWDSL'18: 3rd International Workshop on Real World Domain Specific Languages **Saiph: Towards a DSL for High-Performance Computational Fluid Dynamics.** **CFDlang: High-level code generation for high-order methods in fluid dynamics**	[10:30 - 12:00] Room: Europa 5 WPMVP: Workshop on Programming Models for SIMD/Vector Processing **Small SIMD Matrices for CERN High Throughput Computing** **SIMDization of Small Tensor Multiplication Kernels for Wide SIMD Vector Processors** **MIPP: a Portable C++ SIMD Wrapper and its use for Error Correction Coding in 5G Standard**	[10:30 - 12:00] Room: Europa 1 Session 1: Polyhedral Compilation **Modeling the Conflicting Demands of Parallelism and Temporal/Spatial Locality in Affine Scheduling** **A Polyhedral Compilation Framework for Loops with Dynamic Data-Dependent Bounds** **Polyhedral Expression Propagation**
[12:00 - 13:30] **Lunch**					
[13:30 - 15:10] Room: Europa 3 AACBB: Accelerator Architecture in Computational Biology and Bioinformatics **Keynote 2: "Automata Processor and its Applications in Bioinformatics"** Srinivas Aluru (Georgia Tech) **Streaming Gap-Aware Seed Alignment on the Cache Automaton** **Processing-in-Storage Architecture for Large-Scale Biological Sequence Alignment** **The Genomic Benchmark Suite: Characterization and Architecture Implications**	[13:30 - 15:00] Room: Europa 7 HIPINEB Technical Session 2 (research papers) **Evaluating Energy Saving Strategies on Torus, K-Ary N-Tree, and Dragonfly** **VEF3 traces: towards a complete framework for modelling network workloads for exascale systems** **Improving the Efficiency of Future Exascale Systems with rCUDA**	[13:30 - 15:00] Room: Europa 2 LLVM Performance Workshop **Cache-aware Scheduling and Performance Modeling with LLVM-Polly and Kerncraft** **Enabling Automatic Partitioning of Data-Parallel Kernels with Polyhedral Compilation**	[13:30 - 14:50] Room: Europa 6 RWDSL'18: 3rd International Workshop on Real World Domain Specific Languages **dsmodels: A Little Language for Dynamical Systems** **D'Artagnan: An Embedded DSL Framework for Distributed Embedded Systems**	[13:30 - 15:00] Room: Europa 5 WPMVP: Workshop on Programming Models for SIMD/Vector Processing **Ikra-Cpp: A C++/CUDA DSL for Object-Oriented Programming with Structure-of-Arrays Layout** **Usuba, Optimizing & Trustworthy Bitslicing Compiler** **A Data Layout Transformation for Vectorizing Compilers**	[13:30 - 15:00] Room: Europa 1 Session 2: Data-Flow and Pointer/Alias Analysis **Computing Partially Path-Sensitive MFP Solutions in Data Flow Analyses** **An Efficient Data Structure for Must-Alias Analysis** **Parallel Sparse Flow-Sensitive Points-to Analysis**
[15:00 - 15:30] **Coffee Break with Snack**					
[15:30 - 17:50] Room: Europa 3 AACBB: Accelerator Architecture in Computational Biology and Bioinformatics **Invited Talk: "Addressing Computational Burden to Realize Precision Medicine"** Can Alkan (Bilkent University) **Burrows-Wheeler Short Read Aligner on AWS EC2 F1** **Towards BIMAX: Binary Inclusion-MAXimal parallel implementation for gene expression analysis** **Memory: The Dominant Bottleneck in Genomic Workloads** **Gene Sequencing: Where Time Goes** **Are Next-Generation HPC Systems Ready for Population-level Genomics Data Analytics?** **Closing remarks**	[15:30 - 17:00] Room: Europa 7 Panel Session: "Industrial perspective of high-speed communication technology evolution" **Industrial perspective of high-speed communication technology evolution** moderated by Prof. Young Cho (University of Southern California), Panelists: Eitan Zahavi, Mellanox Technologies, Israel, Ola Torudbakken, Skala Norge AS, Norway, Cyriel Minkenberg, Rockley Photonics Inc., Switzrland	[15:30 - 17:00] Room: Europa 2 LLVM Performance Workshop **Tensor Comprehensions** **LLVM Q&A Panel: Questions Welcome**	[15:30 - 17:00] Room: Europa 6 RWDSL'18: 3rd International Workshop on Real World Domain Specific Languages **Q#: Enabling Scalable Quantum Computing and Development with a High-level DSL** **A Task-Based DSL for Microcomputers** **Close**	[15:30 - 17:00] Room: Europa 5 WPMVP: Workshop on Programming Models for SIMD/Vector Processing **Investigating automatic vectorization for real-time 3D scene understanding** **Panel Discussion**	[15:30 - 17:00] Room: Europa 1 Session 3: Code Generation and Optimisation **PAYJIT: Space-Optimal JIT Compilation and Its Practical Implementation** **Finding Missed Compiler Optimizations by Differential Testing** **Fast and Flexible Instruction Selection with Constraints**
[18:15] **Departure of the busses to the Heurigen**					
[18:30] Heurigen: Toni & Birgit Nigl					

Sunday February 25th, 2018

HPCA		CGO	PPoPP					CC
[08:00 - 18:30] **Registration**								
[08:30 - 10:00] Room: Europa 5 WP3: Second Workshop on Pioneering Processor Paradigms **Welcome and Introduction Pradip Bose** **Keynote I: TBD** Mikko H. Lipasti (MICRO 2017 Test of Time Award, University of Wisconsin - Madison) [09:40 - 10:00] Room: Europa 5 WP3: Retrospective Survey I **On the Evaluation of Computer Architectures**	[08:30 - 10:00] Room: Europa 7 Accelerating Big Data Processing with Hadoop, Spark and Memcached on Datacenters with Modern Architectures **Session 1**	[08:30 - 10:00] Room: Pacific 3 Tutorial: Improving security with reversibility and session types **Session 1**	[08:30 - 10:00] Room: Europa 3 PMAM: Workshop on Programming Models and Applications for Multicores and Manycores **Opening Remarks** **Keynote: "Building the next Generation of MapReduce Programming Models over MPI to Fill the Gaps between Data Analytics and Supercomputers"**	[08:30 - 10:00] Room: Europa 2 GPGPU: Workshop on General Purpose Processing Using GPU **Welcome: The Organizers** **Keynote 1: "Initial Steps toward Making GPU a First-Class Computing Resource: Sharing and Resource Management"** Jun Yang (William Kepler Whiteford Professor of Electrical and Computer Engineering, University of Pittsburgh) [09:30 - 10:00] Room: Europa 2 GPGPU Session 1: Persistent Data Structures **A Case For Persist Barriers in GPUs**	[08:30 - 10:00] Room: Pacific 1 An Introduction to Intel® Threading Building Blocks (Intel® TBB) and its Support for Heterogeneous Programming **Session 1**	[08:30 - 10:00] Room: Pacific 2 Productive parallel programming on FPGA with high-level synthesis **Session 1**	[08:30 - 10:00] Room: Europa 6 Debugging and Profiling Task Parallel Programs with TASKPROF **Session 1**	[08:45 - 10:00] Room: Europa 1 CC Keynote **Compiler and Language Design for Quantum Computing** Bettina Heim (Microsoft Research, USA)
[10:00 - 10:30] **Coffee Break with Snack**								
[10:30 - 11:20] Room: Europa 5 WP3: Invited Talk **40 years since dusk: will hardware capabilities finally make our systems more capable?** Lluis Vilanova (Technion) [11:20 - 12:00] Room: Europa 5 WP3: New/Exploratory paradigms **A Multi-component Branch Predictor Design for Low Resource Budget Processors** **FFT implementation using mono-instruction set computer architecture**	[10:30 - 12:00] Room: Europa 7 Accelerating Big Data Processing with Hadoop, Spark and Memcached on Datacenters with Modern Architectures **Session 2**	[10:30 - 12:00] Room: Pacific 3 Tutorial: Improving security with reversibility and session types **Session 2**	[10:30 - 12:00] Room: Europa 3 PMAM Session 1: GPU and Accelerator **Extending ILUPACK with a Task-Parallel Version of BiCG for Dual-GPU Servers** **Reduction to Band Form for the Singular Value Decomposition on Graphics Accelerators** **Combining PREM compilation and ILP scheduling for high-performance and predictable MPSoC execution**	[10:30 - 12:00] Room: Europa 2 GPGPU Session 2: Applications/Frameworks **Overcoming the Difficulty of Large-scale CGH Generation on multi-GPU Cluster** **Transparent Avoidance of Redundant Data Transfer on GPU-enabled Apache Spark** **GPU-based Acceleration of Detailed Tissue-Scale Cardiac Simulations**	[10:30 - 12:00] Room: Pacific 1 An Introduction to Intel® Threading Building Blocks (Intel® TBB) and its Support for Heterogeneous Programming **Session 2**	[10:30 - 12:00] Room: Pacific 2 Productive parallel programming on FPGA with high-level synthesis **Session 2**	[10:30 - 12:00] Room: Europa 6 Debugging and Profiling Task Parallel Programs with TASKPROF **Session 2**	[10:30 - 12:00] Room: Europa 1 Session 4: Compilation for Specialised Domains **Compiling for Concise Code and Efficient I/O** **Termination Checking and Task Decomposition for Task-Based Intermittent Programs** **A Session Type Provider: Compile-Time API Generation of Distributed Protocols with Refinements in F#**
[12:00 - 13:30] **Lunch**								
[13:20 - 14:20] Room: Europa 5 WP3: Second Workshop on Pioneering Processor Paradigms **Keynote II: TBD** TBD [14:20 - 15:00] Room: Europa 5 WP3: Restrospective Survey II **This Architecture Tastes Like Microarchitecture** **Project CrayOn: Back to the future for a more General-Purpose GPU?**	[13:30 - 15:00] Room: Europa 7 PULP: An open hardware platform, the story so far **PULP concept and goals** **State of the art of open source hardware design** **Summary of PULP systems: PULP, PULPino, PULPissimo** **PULP cores: OR10N, RI5CY, Zero-riscy, Ariane**	[13:30 - 15:00] Room: Pacific 2 Turning HPC clusters into High Performance & High Throughput facilities by using remote GPU virtualization **[Session 1.1] Presentation of remote GPU virtualization techniques and rCUDA features (50 minutes)** **[Session 1.2] Practical demonstration about how to install and use rCUDA (40 minutes)**	[13:30 - 15:00] Room: Pacific 3 Tutorial: Improving security with reversibility and session types **Session 3**	[13:30 - 15:00] Room: Europa 3 PMAM Session 2: Fine-grain Parallelism **Fast and Accurate Performance Analysis of Synchronization** **Supporting Fine-grained Dataflow Parallelism in Big Data Systems** **Intra-Task Parallelism in Automotive Real-Time Systems**	[13:30 - 14:30] Room: Europa 2 GPGPU: Workshop on General Purpose Processing Using GPU **Keynote 2: "Generating High Performance GPU Code using Rewrite Rules with Lift"** Christophe Dubach (University of Edinburgh)	[13:30 - 15:00] Room: Pacific 1 An Introduction to Intel® Threading Building Blocks (Intel® TBB) and its Support for Heterogeneous Programming **Session 3**	[13:30 - 15:00] Room: Europa 6 High Performance Distributed Deep Learning: A Beginner's Guide **Session 1**	[13:30 - 15:00] Room: Europa 1 Session 5: Code Translation and Transformation **Tail Call Elimination and Data Representation for Functional Languages on the Java Virtual Machine** **CAnDL: A Domain Specific Language for Compiler Analysis** **Semantic Reasoning about the Sea of Nodes**
[15:00 - 15:30] **Coffee Break with Snack**								

[15:30 - 15:50] Room: Europa 5 WP3: Restrospective Survey III **45-year CPU evolution: one law and two equations** [15:30 - 15:50] Room: Europa 5 WP3: Panel Session **Panel TBD** Invited Pioneers and speakers plus the retrospective paper authors [15:30 - 15:50] Room: Europa 5 WP3: Recap/discussion; clossing remarks, action items **Discussion driven by workshop organizers.**	[15:30 - 17:30] Room: Europa 7 PULP: An open hardware platform, the story so far **Advanced PULP silicon implementations** **Acceleration for PULP systems, examples from cryptography and neural networks** **PULP Programming**	[15:30 - 17:00] Room: Pacific 2 Turning HPC clusters into High Performance & High Throughput facilities by using remote GPU virtualization **[Session 2] Guided exercises so that the audience uses rCUDA in a cluster located at Technical University of Valencia, Spain** **Time for attendees to freely exercise with rCUDA in the remote cluster (a set of exercises is proposed)**	[15:30 - 17:00] Room: Pacific 3 Tutorial: Improving security with reversibility and session types **Session 4**	[15:30 - 17:00] Room: Europa 3 PMAM Session 3: Cache and Pipeline **Understanding Parallelization Tradeoffs for Linear Pipelines** **An Evaluation of Vectorization and Cache Reuse Tradeoffs on Modern CPUs** **VAIL: A Victim-Aware Cache Policy for Improving Lifetime of Hybrid Memory** [17:00 - 17:05] Room: Europa 3 **Closing Remarks**	[15:30 - 16:30] Room: Europa 2 GPGPU Session 3: Concurrent Kernels **MaxPair: Enhance OpenCL Concurrent Kernel Execution by Weighted Maximum Matching**	[15:30 - 17:00] Room: Pacific 1 An Introduction to Intel® Threading Building Blocks (Intel® TBB) and its Support for Heterogeneous Programming **Session 4**	[15:30 - 17:00] Room: Europa 6 High Performance Distributed Deep Learning: A Beginner's Guide **Session 2**	[15:30 - 17:00] Room: Europa 1 Session 6: Compile- and Run-Time Analysis **Towards a Compiler Analysis for Parallel Algorithmic Skeletons** **Generalized Profile-Guided Iterator Recognition** **Efficient Dynamic Analysis for Node.js**
[18:00] **HPCA/CGO/PPoPP Welcome Reception and Poster Session**								
[19:45] (Anthony's Bar) **Women-in-Computer-Architecture (WICARCH) get-together**								

Monday February 26th, 2018

HPCA

[08:00 - 18:00] **Registration**

[08:30 - 08:45] **Opening**

[08:45 - 09:55] (Europa 4) HPCA Keynote: What is the role of Architecture and Software Researchers on the Road to Quantum Supremacy? Margaret Martonosi (Princeton University)

[09:55 - 10:20] **Coffee Break with Snack**

[10:20 - 10:30] Room: Europa 4

Test of Time Award Session

HPCA Test of Time Award

[10:30 - 12:00] Room: Europa 4

Best Paper Session

Session chair: Josep Torrellas (UIUC)

Amdahl's Law in the Datacenter Era: A Market for Fair Processor Allocation

iNPG: Accelerating Critical Section Access with In-Network Packet Generation for NoC based Many-cores

Enabling Efficient Network Service Function Chain Deployment on Heterogeneous Server Platform

Reducing Data Transfer Energy by Exploiting Similarity within a Data Transaction

[11:45 - 13:15] **Lunch**

[13:15 - 14:55] Room: Europa 4

Session 2A: Architecture for Neural Network

Session chair: Rajeev Balasubramonian (University of Utah)

Making Memristive Neural Network Accelerators Reliable

Towards Efficient Microarchitectural Design for Accelerating Unsupervised GAN-based Deep Learning

Compressing DMA Engine: Leveraging Activation Sparsity for Training Deep Neural Networks

In-situ AI: Towards Autonomous and Incremental Deep Learning for IoT Systems

[14:55 - 15:15] **Coffee Break with Snack**

[15:15 - 16:55] Room: Europa 4

Session 3A: Security

Session chair: David R. Kaeli (Northeastern University)

RCoal: Mitigating GPU Timing Attack via Subwarp-based Randomized Coalescing Techniques

Are Coherence Protocol States vulnerable to Information Leakage?

Record-Replay Architecture as a General Security Framework

The DRAM Latency PUF: Quickly Evaluating Physical Unclonable Functions by Exploiting the Latency-Reliability Tradeoff in Modern DRAM Devices

[16:55 - 17:15] **Break**

[17:15 - 18:55] Room: Europa 4

Session 4A: Microarchitecture and Benchmark

Session chair: Benjamin Lee (Duke University)

A Novel Register Renaming Technique for Out-of-Order Processors

Wait of a Decade: Did SPEC CPU 2017 Broaden the Performance Horizon?

Architectural Support for Task Dependence Management with Flexible Software Scheduling

GDP: Using Dataflow Properties to Accurately Estimate Interference-free Performance at Runtime

[19:15 - 20:15] Room: Europa 4

HPCA Business Meeting

CGO

[10:20 - 11:45] Room: Europa 2

Session 1: Managed Runtimes

SIMD Intrinsics on Managed Language Runtimes

CollectionSwitch: A Framework for Efficient and Dynamic Collection Selection

Analyzing and Optimizing Task Granularity on the JVM

[13:15 - 14:55] Room: Europa 2

Session 2: Resilience and Security

Automating Efficient Variable-Grained Resiliency for Low-Power IoT Systems

Resilient Decentralized Android Application Repackaging Detection Using Logic Bombs

nAdroid: Statically Detecting Ordering Violations in Android Applications

SGXElide: Enabling Enclave Code Secrecy via Self-Modification

[15:15 - 15:25] Room: Europa 2

Test of Time Award Session

CGO Test of Time Award

[15:25 - 16:55] Room: Europa 2

Session 3: Best Paper Finalists

Poker: Permutation-based SIMD Execution of Intensive Tree Search by Path Encoding

High Performance Stencil Code Generation with LIFT

Qubit Allocation

Dominance-based Duplication Simulation (DBDS): Code Duplication to Enable Compiler Optimizations

[17:00 - 19:00] Room: Europa 7

Student Research Competition

[17:15 - 17:45] Room: Europa 3

CGO & PPoPP Artifact Evaluation

[18:00 - 19:00] Room: Europa 2

CGO Business Meeting

PPoPP

[10:20 - 11:35] Room: Europa 3

Session 1: Concurrent Data Structures

Session chair: Xipeng Shen (North Carolina State University)

Interval-Based Memory Reclamation

Harnessing Epoch-based Reclamation for Efficient Range Queries

A Persistent Lock-Free Queue for Non-Volatile Memory

[13:15 - 14:55] Room: Europa 3

Session 2: Compilers and runtime systems

Session chair: I-Ting Angelina Lee (Washington University in St. Louis)

Juggler: A Dependency-Aware Task Based Execution Framework for GPUs

HPVM: Heterogeneous Parallel Virtual Machine

Hierarchical Memory Management for Mutable State

SuperNeurons: Dynamic GPU Memory Management for Training Deep Neural Networks

[15:15 - 16:30] Room: Europa 3

Session 3: Performance

Session chair: Milind Chabbi (Baidu Research)

Bridging the Gap between Deep Learning and Sparse Matrix Format Selection

Optimizing N-Dimensional, Winograd-Based Convolution for Manycore CPUs

vSensor: Leveraging Fixed-Workload Snippets of Programs for Performance Variance Detection

[17:15 - 17:45] Room: Europa 3

CGO & PPoPP Artifact Evaluation

[18:00 - 19:00] Room: Europa 3

PPoPP Business Meeting

HPCA / PPoPP

[13:15 - 14:55] Room: Europa 5+6

Session 2B: Cache and Memory

Session chair: Paul V. Gratz (Texas A&M University)

A Hybrid Cache Partitioning-Sharing Technique for Commodity Multicores

SIPT: Speculatively Indexed, Physically Tagged Caches

Domino Temporal Data Prefetcher

ProFess: A Probabilistic Hybrid Main Memory Management Framework for High Performance and Fairness

[15:15 - 16:55] Room: Europa 5+6

Session 3B: GPU Cache and Memory

Session chair: Bradford M. Beckmann (AMD)

Accelerate GPU Concurrent Kernel Execution by Mitigating Memory Pipeline Stalls

LATTE-CC: Latency Tolerance Aware Adaptive Cache Compression Management for Energy Efficient GPUs

GETM: high-performance GPU transactional memory via eager conflict detection

Efficient and Fair Multi-programming in GPUs via Effective Bandwidth Management

[17:15 - 18:55] Room: Europa 5+6

Session 4B: Persistent and NVM memory

Session chair: Hai Li (Duke University)

Crash Consistency in Encrypted Non-Volatile Main Memory Systems

Adaptive Memory Fusion: Towards Transparent, Agile Integration of Persistent Memory

Efficient Hardware-based Undo+Redo Logging for Persistent Memory Systems

Enabling Fine-Grain Restricted Coset Coding Through Word-Level Compression for PCM

Tuesday February 27th, 2018

HPCA		CGO	PPoPP
[08:00 - 17:00] **Registration**			
[08:00 - 09:40] Room: Europa 4 Session 5A: GPU Session chair: Minsoo Rhu (POSTECH) **Perception-Oriented 3D Rendering Approximation for Modern Graphics Processors** **Warp Scheduling for Fine-Grained Synchronization** **WIR: Warp Instruction Reuse to Minimize Repeated Computations in GPUs** **G-TSC: Timestamp Based Coherence for GPUs**	[08:00 - 09:40] Room: Europa 5+6 Session 5B: Secure memory Session chair: Rui Hou (Chinese Academy of Science) **D-ORAM: Path-ORAM Delegation for Low Execution Interference on Cloud Servers with Untrusted Memory** **Secure DIMM: Moving ORAM Primitives Closer to Memory** **Comprehensive VM Protection against Untrusted Hypervisor through Retrofitted AMD Memory Encryption** **SYNERGY: Rethinking Secure-Memory Design for Error-Correcting Memories**	[08:00 - 09:40] Room: Europa 2 Session 4: Linear Algebra and Vectorization **The Generalized Matrix Chain Algorithm** **CVR: Efficient Vectorization of SpMV on X86 Processors** **Look-Ahead SLP: Auto-vectorization in the Presence of Commutative Operations** **Conflict-Free Vectorization of Associative Irregular Applications with Recent SIMD Architectural Advances**	[08:00 - 09:40] Room: Europa 3 Session 4: Best Paper Candidates Session chair: Idit Keidar (Technion) **Cache-Tries: Concurrent Lock-Free Hash Tries with Constant-Time Operations** **Featherlight On-the-fly False-sharing Detection** **Register Optimizations for Stencils on GPUs** **FlashR: Parallelize and Scale R for Machine Learning using SSDs**
[09:40 - 10:05] **Coffee Break with Snack**			
[10:05 - 11:45] Room: Europa 4 Session 6A: Novel Architecture Session chair: Kei Hiraki (University of Tokyo) **A Case for Packageless Processors** **Extending the Power-Efficiency and Performance of Photonic Interconnects for Heterogeneous Multicores** **Routerless Networks-on-Chip** **HeatWatch: Optimizing 3D NAND Read Operations With Self-Recovery and Temperature Awareness**	[10:05 - 11:45] Room: Europa 5+6 Session 6B: In-Memory Computing Session chair: Jishen Zhao (UCSD) **RC-NVM: Enabling Symmetric Row and Column Memory Accesses for In-Memory Databases** **GraphR: Accelerating Graph Processing Using ReRAM** **GraphP: Reducing Communication of PIM-based Graph Processing with Efficient Data Partition** **PM3: Power Modeling and Power Management for Processing-in-Memory**	[10:05 - 11:45] Room: Europa 2 Session 5: Static and Dynamic Analysis **Scalable Concurrency Debugging with Distributed Graph Processing** **Lightweight Detection of Cache Conflicts** **CUDAAdvisor: LLVM-Based Runtime Profiling for Modern GPUs** **May-Happen-in-Parallel Analysis with Static Vector Clocks**	[10:05 - 11:45] Room: Europa 3 Session 5: Concurrency control and fault tolerance Session chair: Walter Binder (USI) **DisCVar: Discovering Critical Variables Using Algorithmic Differentiation for Transient Faults** **Practical Concurrent Traversals in Search Trees** **Communication-Avoiding Parallel Minimum Cuts and Connected Components** **Safe Privatization in Transactional Memory**
[11:45 - 13:15] **Lunch**			
[11:45 - 12:30] (lunch room) Women in Academia and Industry Lunch Session			
[12:35 - 13:10] (Europa 4) Women in Academia and Industry Panel			
[13:15 - 14:25] (Europa 4) CGO Keynote: Biological Computation Sara-Jane Dunn (Microsoft Research Limited)			
[14:25 - 14:50] **Coffee Break with Snack**			
[14:50 - 16:30] Room: Europa 4 Session 7A: Industry Track Session chair: Lieven Eeckhout (Ghent University) **Don't Correct the Tags in a Cache, just Check their Hamming Distance from the Lookup Tag** **Reliability-aware Data Placement for Heterogeneous Memory Architecture** **SmarCo: An Efficient Many-Core Processor for High-Throughput Applications in Datacenters** **Lost in Abstraction: Pitfalls of Analyzing GPUs at the Intermediate Language Level**	[14:50 - 16:30] Room: Europa 5+6 Session 7B: Best of CAL Session chair: Dan Sorin (Duke University) **Resistive Address Decoder** **Transcending Hardware Limits with Software Out-of-order Processing** **Sensing CPU voltage noise through Electromagnetic Emanations**	[14:50 - 16:30] Room: Europa 2 Session 6: Memory usage Optimisation **DeLICM: Scalar Dependence Removal at Zero Memory Cost** **Loop Transformations Leveraging Hardware Prefetching** **Transforming Loop Chains via Macro Dataflow Graphs** **Local Memory-Aware Kernel Perforation**	[14:50 - 16:30] Room: Europa 3 Session 6: Models and Libraries Session chair: Zoltan Majo (Ergon Informatik AG) **Making Pull-Based Graph Processing Performant** **An Effective Fusion and Tile Size Model for Optimizing Image Processing Pipelines** **LazyGraph: Lazy Data Coherency for Replicas in Distributed Graph-Parallel Computation** **PAM: Parallel Augmented Maps**
[17:00] Departure of the busses to Palais Liechtenstein			
[18:00] Banquet at Palais Liechtenstein			

Wednesday February 28th, 2018

HPCA		CGO	PPoPP
[08:00 - 09:00] (Europa 4) **PPoPP Keynote: From confusion to clarity: hardware concurrency programming models 2008-2018** Peter Sewell (University of Cambridge)			
[09:00 - 09:25] **Coffee Break with Snack**			
[09:25 - 11:05] Room: Europa 4 Session 8A: Industry Track (applications) Session chair: Andrew Putnam (Microsoft) **Applied Machine Learning at Facebook: A Datacenter Infrastructure Perspective** **Amdahl's Law in Big Data Analytics: Alive and Kicking in TPCx-BB (BigBench)** **Memory Hierarchy for Web Search** **Characterizing Resource Sensitivity of Database Workloads**	[09:25 - 11:05] Room: Europa 5+6 Session 8B: Memory Session chair: Guangyu Sun (Peking University) **ERUCA: Efficient DRAM Resource Utilization and Resource Conflict Avoidance for Memory System Parallelism** **DUO: Dual Use of On-chip Redundancy for High Reliability** **Memory System Design for Ultra Low Power, Computationally Error Resilient Processor Microarchitectures** **NACHOS : Software-Driven Hardware-Assisted Memory Disambiguation for Accelerators**	[09:25 - 11:05] Room: Europa 2 Session 7: Program Generation and Synthesis **AutoPA: Automatically Generating Active Driver from Original Passive Driver Code** **Synthesizing an Instruction Selection Rule Library from Semantic Specifications** **Synthesizing Programs That Expose Performance Bottlenecks** **Program Generation for Small-Scale Linear Algebra Applications**	[09:25 - 11:05] Room: Europa 3 Session 7: Parallel frameworks and applications Session chair: Bernhard Egger (Seoul National University) **Efficient Shuffle Management with SCache for DAG Computing Frameworks** **High-Performance Genomics Data Analysis Framework with In-Memory Computing** **Griffin: Uniting CPU and GPU in Information Retrieval Systems for Intra-Query Parallelism** **swSpTRSV: a Fast Sparse Triangular Solve with Sparse Level Tile Layout on Sunway Architectures**
[11:05 - 11:20] **Break**			
[11:20 - 12:35] Room: Europa 4 Session 9A: Accelerators Session chair: Xuehai Qian (USC) **OuterSPACE: An Outer product based SPArse matrix multiplication acCElerator** **Searching for Potential gRNA Off-Target Sites for CRISPR/Cas9 using Automata Processing across Different Platforms** **Characterizing and Mitigating Output Reporting Bottlenecks in Spatial-Reconfigurable Automata Processing Architectures**	[11:20 - 12:35] Room: Europa 5+6 Session 9B: Power Session chair: Guru Venkataramani (George Washington University) **Power and Energy Characterization of an Open Source 25-core Manycore Processor** **A Spot Capacity Market to Increase Power Infrastructure Utilization in Multi-Tenant Data Centers** **GPGPU Power Modeling for Multi-Domain Voltage-Frequency Scaling**	[11:20 - 12:35] Room: Europa 2 Session 8: Compilation for Specialised Domains **Optimal DNN Primitive Selection with Partitioned Boolean Quadratic Programming** **Register Allocation for Intel Processor Graphics** **A Compiler for Cyber-Physical Digital Microfluidic Biochips**	[11:20 - 12:10] Room: Europa 3 Session 8: Race Detection Session chair: Jesper Larsson Träff (TU Wien) **VerifiedFT: A Verified, High-Performance Dynamic Race Detector** **Efficient Parallel Determinacy Race Detection for Two-Dimensional Dags**
[12:35] **HPCA Closing**		[12:35 - 12:45] Room: Europa 2 Best Paper Award Session **CGO 2018 Best Paper Award** [12:45] **CGO Closing**	[12:10] **PPoPP Closing**

Venue Floor Plan

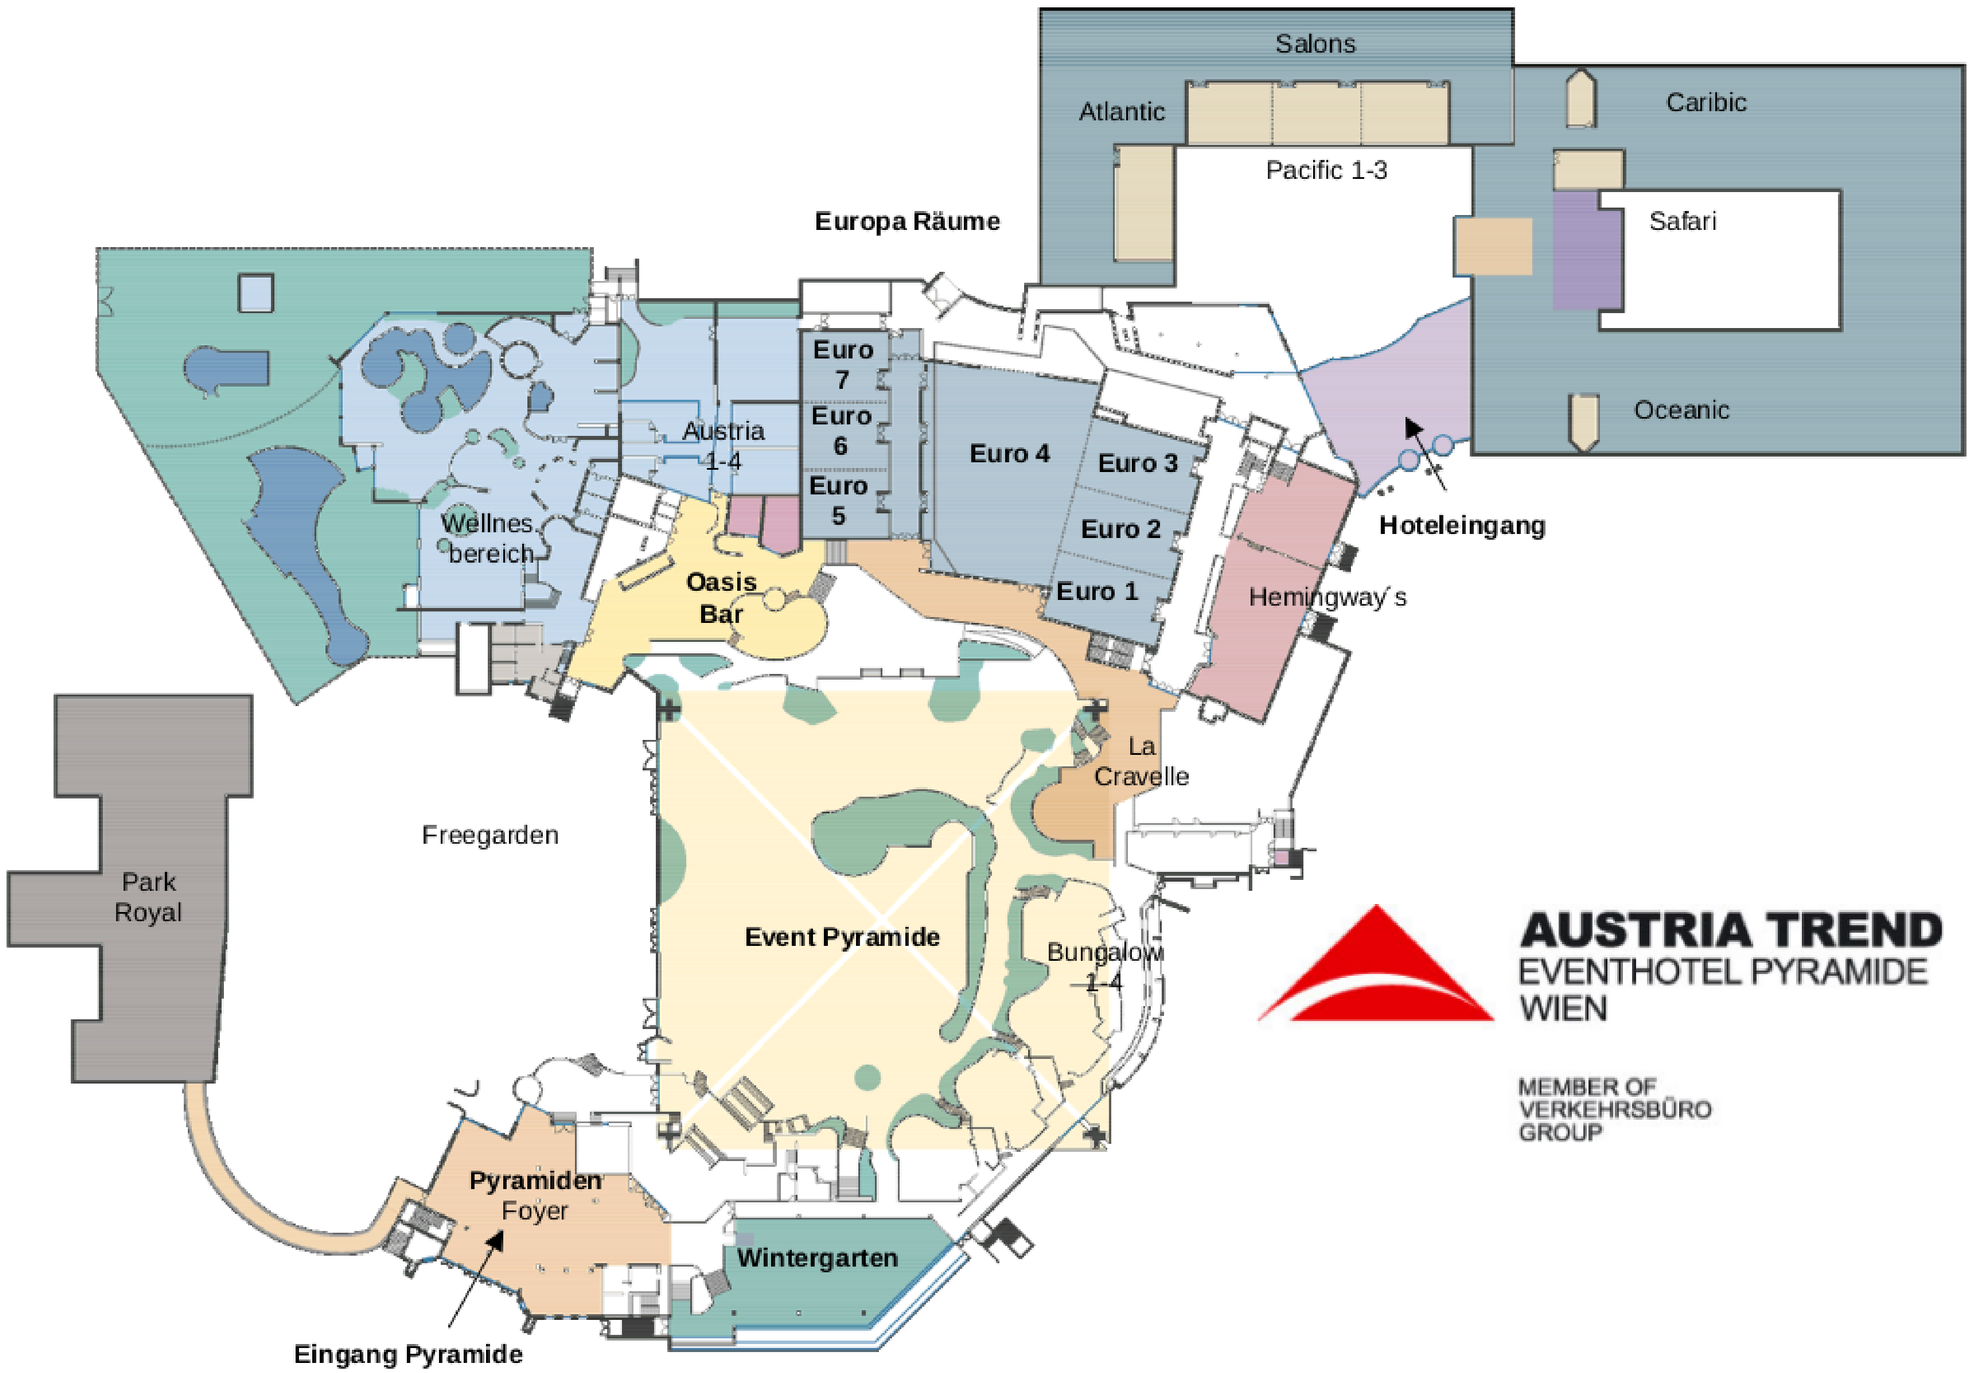

AUSTRIA TREND
EVENTHOTEL PYRAMIDE WIEN

MEMBER OF VERKEHRSBÜRO GROUP