

鸣谢 :感谢华为、UCloud 和 Tsinghua-Xilinx AI Research Fund
对本次会议的支持

会议地址 : 成都世外桃源酒店 (玫瑰厅 3 六楼-150)

会议网页 : <https://www.acmturc.com/2019/cn/SIGOPS.html>

Day 1: May 17th

Time	Talks
8:00-8:10	Opening Remarks:
8:10-9:10	Keynote 1
	Topic: Accelerating Distributed DNN Training Speaker: Chuanxiong Guo (Director, Bytedance AI Lab) Chair: Wenguang Chen
9:10-10:30	Session 1: Deep learning Chair: TBD
	Large-scale GraphNN System Wei Lin (Alibaba group)
	A Comprehensive Graph Neural Network Overview in Practice Hongxia Yang (Alibaba group), Chang Zhou (Alibaba group)
	SeerNet: Predicting Convolutional Neural Network Feature-Map Sparsity through Low-Bit Quantization Chen Zhang (Microsoft), Yunxin Liu (Microsoft) and Lintao Zhang (Microsoft)
	A Machine Learning Framework with Feature Selection for Floorplan Evaluation in IC Physical Design Shuzheng Zhang (National University of Defense Technology), Zhenyu Zhao (National University of Defense Technology), Chaochao Feng (National University of Defense Technology), Lei Wang (National University of Defense Technology)
10:30-10:50	Tea Break
10:50-12:10	Session 2: Cloud & Virtualization Chair: TBD
	TPShare: A Time-Space Sharing Scheduling abstraction for Shared Cloud via Vertical Labels Yuzhao Wang (Huazhong University of Science & Technology, Shenzhen Institutes of Advanced Technology), Lele Li (Shenzhen Institutes of Advanced Technology), You Wu (University of Southern California), Junqing Yu (Huazhong University of Science & Technology), Zhibin Yu (Shenzhen

	Institutes of Advanced Technology), Xuehai Qian (University of Southern California)
	Phantasy: Low-Latency Virtualization-based Fault Tolerance via Asynchronous Prefetching Zhen Xiao (Peking University)
	Characterization of Co-located Container Workloads with Hardware Performance Counters Wenyan Chen (Shenzhen institutes of advanced technology chinese academy of sciences), Kejiang Ye (Shenzhen institutes of advanced technology chinese academy of sciences) and Cheng-Zhong Xu (Shenzhen institutes of advanced technology chinese academy of sciences)
	A Preliminary Evaluation on Many-to-one Virtualization Jin Zhang (Shanghai Jiao Tong University), Zhuocheng Ding (Shanghai Jiao Tong University), Yubin Chen (Shanghai Jiao Tong University), Zhengwei Qi (Shanghai Jiao Tong University) and Haibing Guan (Shanghai Jiao Tong University)
12:10-14:00	Lunch and Break
14:00-15:20	Session 3: Memory Chair: TBD
	Precise Memory Object Life Cycle Modeling in Memory SSA Shin-Ming Liu (XC5 Limited), Long Li (XC5 Limited) and Sun Chan (XC5 Limited)
	Panthera: Holistic Memory Management for Big Data Processing over Hybrid Memories Chenxi Wang (Institute of Computing Technology, CAS), Huimin Cui (Institute of Computing Technology, CAS), Ting Cao (Microsoft), John Zigman (The University of Sydney), Haris Volos (HP Lab), Onur Mutlu (Carnegie Mellon University), Fang Lv (Institute of Computing Technology, CAS), Xiaobing Feng (Institute of Computing Technology, CAS) and Harry Xu (University of California, Los Angeles)
	Huge Page-Friendly Virtualized Memory Management Sai Sha (Peking University), Jingyuan Hu (Peking University), Yingwei Luo (Peking University), Xiaolin Wang (Peking University) and Zhenlin Wang (Michigan Technolgoical University)
	IMPULP: A Hardware Approach for In-process Memory Protection via User-Level Partitioning Yangyang Zhao (University of Chinese Academy of Sciences; ICT, Chinese Academy of Sciences), Mingyu Chen (Institute of Computing Technology, Chinese Academy of Sciences), Yuhang Liu (Institute of Computing Technology, Chinese Academy of Sciences), Zonghao Yang (University of Chinese Academy of Sciences; ICT, Chinese Academy of Sciences), Xiaojing Zhu (Institute of Computing Technology, Chinese Academy of Sciences), Zonghui Hong (Institute of Computing Technology, Chinese Academy of

	Sciences) and Yunge Guo (Institute of Computing Technology, Chinese Academy of Sciences)
15:20-15:40	Break
15:40-17:00	<p>Session 4: System Optimization Chair: TBD</p> <p>From Static Dataflow to Dynamic, and Pull Back Peng Cao (Ant Financial), Jie Wang (Ant Financial) and Zhenyu Guo (Ant Financial)</p> <p>pLock: A Fast Lock for Architectures with Explicit Inter-core Message Passing Xiongchao Tang (Tsinghua University), Jidong Zhai (Tsinghua University), Xuehai Qian (University of Southern California) and Wenguang Chen (Tsinghua University)</p> <p>MPI-RCDD: A Framework for MPI Runtime Communication Deadlock Detection Wei Hongmei (Shanghai Jiao Tong University) and Gao Jian (Jiangnan Institute of Computing Technology)</p> <p>Optimistic Transaction Processing in Deterministic Database Zhiyuan Dong (Shanghai Jiao Tong University), Chuzhe Tang (Shanghai Jiao Tong University), Jiachen Wang (Shanghai Jiao Tong University), Zhaoguo Wang (Shanghai Jiao Tong University) and Haibo Chen (Shanghai Jiao Tong University)</p>
17:00-18:00	<p>Panel Open Source Chip and Its Ecosystem Chair: Yungang Bao</p>

Day 2: May 18th (afternoon)

Time	Talks
13:30-14:30	<p>Keynote 3 Fast, Safe User-Level Services Speaker: Michael L. Scott (Arthur Gould Yates Professor of Engineering, CS, University of Rochester) Chair: Haibo Chen</p>
14:30-15:50	<p>Session 5: Cloud Chair: TBD</p> <p>Fast Startup for Serverless Computing with Initialization-less Loading Dong Du (Shanghai Jiao Tong University), Tianyi Yu (Shanghai Jiao Tong University), Yubin Xia (Shanghai Jiao Tong University), Haibo Chen (Shanghai Jiao Tong University), Binyu Zang (Shanghai Jiao Tong University), Chenggang Qin (Alibaba group) and Qixuan Wu (Alibaba group)</p> <p>Improving Resource Efficiency with Fine-grained Control over Cloud</p>

	<p>Service Components Laiping Zhao (Tianjin University) and Yanan Yang (Tianjin University)</p> <p>GAugur: Quantifying Performance Interference of Colocated Games for Improving Resource Utilization in Cloud Gaming Yusen Li (Nankai University), Chuxu Shan (Nankai University), Ruobing Chen (Nankai University), Xueyan Tang (Nanyang Technological University), Wentong Cai (Nanyang Technological University), Shanjing Tang (Tianjin University), Gang Wang (Nankai University), Xiaoguang Liu (Nankai University), Xiaoli Gong (Nankai University) and Ying Zhang (Nankai University)</p> <p>Bigflow: a general optimization layer for distributed computing frameworks Yuncong Zhang (Baidu Inc.), Xiaoyang Wang (Peking University), Yao Xu (Baidu Inc.), Jianwei Zhang (Baidu Inc.), Xiaodong Lin (Baidu Inc.), Guangyu Sun (Peking University), Gonglin Zheng (Baidu Inc.), Shanhui Yin (Baidu Inc.), Xianjin Ye (Baidu Inc.), Li Li (Baidu Inc.), Zhan Song (Baidu Inc.), Dongdong Miao (Baidu Inc.) and Xiangwei Liu (Baidu Inc.)</p>
15:50-16:10	Tea Break
16:10-17:30	<p>Session6: Acceleration Chair: TBD</p> <p>LAcc: Exploiting Lookup Table-based Fast and Accurate Vector Multiplication in DRAM-based CNN Accelerator Quan Deng (College of Computer, National University of Defense Technology)</p> <p>SIES: A novel Implementation of the Spiking Convolutional Neural Network Inference Engine on FPGA Shuquan Wang (National University of Defense Technology) and Lei Wang (National University of Defense Technology)</p> <p>A Systolic Accelerator for Real-time Edge Intelligence System Tian Shuo (National University of Defense Technology), Wang Lei (National University of Defense Technology) and Dou Qiang (National University of Defense Technology)</p> <p>Poly: Efficient Heterogeneous System and Application Management for Interactive Applications Shuo Wang (Peking University), Yun Liang (Peking University) and Wei Zhang (Peking University)</p>
17:30-17:40	Concluding Remarks