

Xuehai Qian

CONTACT INFORMATION	3740 McClintock Avenue, EEB 204 Los Angeles, CA 90089-2562, USA http://alchem.usc.edu/~xuehaiq/ http://alchem.usc.edu	Tel: 213-740-4459 Fax: 213-740-9803 Email: xuehai.qian@usc.edu
RESEARCH INTERESTS	System and Architecture for Graph Processing and Machine Learning. Parallelism and Concurrency. Performance Model. NVM and Cloud System.	
EDUCATION	<ul style="list-style-type: none"> o University of Illinois at Urbana-Champaign (UIUC) Urbana, IL o Ph.D. in Computer Science (with Prof. Josep Torrellas) Aug. 2007 – Aug. 2013 o • Dissertation title: <i>Scalable and Flexible Bulk Architecture</i> 	
PROFESSIONAL EXPERIENCE	<ul style="list-style-type: none"> o Assistant Professor, <i>University of Southern California.</i> Aug. 2015 – Present o Postdoctoral Researcher, <i>UC Berkeley.</i> Sept. 2013 – June 2015 o Research Intern, <i>Microsoft Research, Silicon Valley</i> Summer 2011 o Research Intern, <i>Microsoft Research, Redmond</i> Summer 2008 	
AWARDS AND HONORS	<ul style="list-style-type: none"> o Induce into HPCA Hall of Fame 2019 o Induce into ASPLOS Hall of Fame 2018 o NSF CAREER Award 2018 o ACM Distinguished Speaker 2017 o NSF CRII Award 2017 o One of 200 selected to participate Heidelberg Laureate Forum Foundation (HLFF) 2015 o W.J. Poppelbaum Memorial Award, UIUC 2013 o Andrew and Shana Laursen Fellowship, UIUC 2007 	
SELECTED PUBLICATIONS	<ul style="list-style-type: none"> o Qinyi Luo, Jinkun Lin, Youwei Zhuo Xuehai Qian. <i>HOP: Heterogeneity-Aware Decentralized Training.</i> ASPLOS'19 o Xiongchao Tang, Jidong Zhai, Xuehai Qian, Wenguang Chen. <i>SW-Lock: A Fast Lock for Sunway Taihulight.</i> ASPLOS'19 o Ao Ren, Jiayu Li, Tianyun Zhang, Shaokai Ye, Wenyao Xu, Xuehai Qian, Xue Lin, Yanzhi Wang. <i>ADMM-NN: An Algorithm-Hardware Co-Design Framework of DNNs Using Alternating Direction Methods of Multipliers.</i> ASPLOS'19 o Linghao Song, Jiachen Mao, Youwei Zhuo, Xuehai Qian, Hai Li, Yiran Chen. <i>HyPar: Towards Hybrid Parallelism for Deep Learning Accelerator Array.</i> HPCA'19 o Zhe Li, Caiwen Ding, Siyue Wang, Wujie Wen, Youwei Zhuo, Chang Liu, Qinru Qiu, Wenyao Xu, Xue Lin, Xuehai Qian, Yanzhi Wang. <i>E-RNN: Design Optimization for Efficient Recurrent Neural Networks in FPGAs.</i> HPCA'19 o Xiebing Wang, Kai Huang, Alois Knoll, Xuehai Qian. <i>A Hybrid Framework for Fast and Accurate GPU Performance Estimation through Source-Level Analysis and Trace-Based Simulation.</i> HPCA'19 o Youwei Zhuo, Jinglei Cheng, Qinyi Luo, Jidong Zhai, Yanzhi Wang, Zhongzhi Luan, Xuehai Qian. <i>CSE: Parallel Finite State Machines with Convergence Set Enumeration.</i> MICRO'18 o Yirong Lv, Bin Sun, Qinyi Luo, Jing Wang, Zhibin Yu, Xuehai Qian. <i>CounterMiner: Mining Big Performance Data from Hardware Counters.</i> MICRO'18 	

- Chunhua Deng, Siyu Liao, Yi Xie, Keshab K. Parhi, Xuehai Qian, Bo Yuan. *PermDNN: Efficient Compressed DNN Architecture with Permuted Diagonal Matrices*. **MICRO'18**
- Xiongchao Tang, Jidong Zhai, Xuehai Qian, Bingsheng He, Wei Xue and Wenguang Chen. *vSensor: Leveraging Fixed-Workload Modules of Programs for Performance Variance Detection*. **PPOPP'18**
- Mingxing Zhang, Yongwei Wu, Youwei Zhuo, Xuehai Qian, Chenying Huan, Kang Chen. *Wonderland: A Novel Abstraction-Based Out-Of-Core Graph Processing System*. **ASPLOS'18**
- Ruizhe Cai, Ao Ren, Ning Liu, Caiwen Ding, Luhao Wang, Xuehai Qian, Massoud Pedram, Yanzhi Wang. *VIBNN: Hardware Acceleration of Bayesian Neural Networks*. **ASPLOS'18**
- Zhibin Yu, Zhendong Bei, Xuehai Qian. *DAC: Data-Aware Auto-Tuning High Dimensional Configurations of In-Memory Cluster Computing*. **ASPLOS'18**
- Yanzhi Wang, Caiwen Ding, Geng Yuan, Siyu Liao, Zhe Li, Xiaolong Ma, Bo Yuan, Xuehai Qian, Jian Tang, Qinru Qiu, Xue Lin. *Towards Ultra-High Performance and Energy Efficiency of Deep Learning Systems: An Algorithm-Hardware Co-Optimization Framework*. **AAAI'18**
- Linghao Song, Youwei Zhuo, Xuehai Qian, Miao Hu, Hai Li, Yiran Chen. *GraphR: Accelerating Graph Processing Using ReRAM*. **HPCA'18**
- Mingxing Zhang, Youwei Zhuo, Chao Wang, Mingyu Gao, Yongwei Wu, Kang Chen, Christos Kozyrakis, Xuehai Qian. *GraphP: Reducing Communication of PIM-based Graph Processing with Efficient Data Partition*. **HPCA'18**
- Abdulaziz Tabbakh, Xuehai Qian, Murali Annavaram. *G-TSC: Timestamp Based Coherence for GPUs*. **HPCA'18**
- Caiwen Ding, Yanzhi Wang, Siyu Liao, Zhe Li, Yu Bai, Youwei Zhuo, Chao Wang, Xuehai Qian, Ning Liu, Geng Yuan, Xiaolong Ma, Yipeng Zhang, Xue Lin, Jian Tang, Qinru Qiu, Bo Yuan. *Cir-CNN: Accelerating and Compressing Deep Neural Networks Using Block-Circulant Weight Matrices*. **MICRO'17**
- Zhiyuan Ai, Mingxing Zhang, Yongwei Wu, Xuehai Qian, Kang Chen, Weimin Zheng. *Squeezing out All the Value of Loaded Data: An Out-Of-Core Graph Processing System with Reduced Disk I/O*. **ATC'17**
- Ao Ren, Ji Li, Zhe Li, Caiwen Ding, Xuehai Qian, Qinru Qiu, Bo Yuan and Yanzhi Wang. *SC-DCNN: Highly-Scalable Deep Convolutional Neural Network using Stochastic Computing*. **ASPLOS'17**
- Mengxing Liu, Mingxing Zhang, Kang Chen, Xuehai Qian, Yongwei Wu, Weimin Zheng and Jinglei Ren. *DudeTM: Building Durable Transactions for Persistent Memories with Decoupling*. **ASPLOS'17**
- Linghao Song, Xuehai Qian, Hai Li and Yiran Chen. *PipeLayer: A Pipelined ReRAM-Based Accelerator for Deep Learning*. **HPCA'17**
- Mingxing Zhang, Yongwei Wu, Kang Chen, Xuehai Qian, Xue Li and Weimin Zheng. *Exploring the Hidden Dimension in Graph Processing*. **OSDI'16**
- Xuehai Qian, Koushik Sen, Paul Hargrove and Costin Iancu. *SReplay: Deterministic Group Replay for One-Sided Communication*. **ICS'16**
- Xuehai Qian, Benjamin Sahelices and Depei Qian. *Pacifier: Record and Replay for Relaxed-Consistency Multiprocessors with Distributed Directory Protocol*. **ISCA'14**
- Xuehai Qian, Benjamin Sahelices and Josep Torrellas. *OmniOrder: Directory-Based Conflict Serialization of Transactions*. **ISCA'14**
- Xuehai Qian, Benjamin Sahelices, Josep Torrellas and Depei Qian. *BulkCommit: Scalable and Fast Commit of Atomic Blocks in a Lazy Multiprocessor Environment*. **MICRO'13**
- Xuehai Qian, Benjamin Sahelices, Josep Torrellas and Depei Qian. *Volition: Precise and Scalable Sequential Consistency Violation Detection*. **ASPLOS'13**
- Xuehai Qian, He Huang, Benjamin Sahelices, and Depei Qian. *Rainbow: Efficient Memory Dependence Recording with High Replay Parallelism for Relaxed Memory Model*. **HPCA'13**
- Xuehai Qian, Benjamin Sahelices and Josep Torrellas. *BulkSMT: Designing SMT Processors for Atomic-Block Execution*. **HPCA'12**

- Xuehai Qian, Wonsun Ahn and Josep Torrellas. *ScalableBulk: Scalable Cache Coherence for Atomic Blocks in a Lazy Environment*. **MICRO'10**
-

STUDENTS

- *Ph.D Students*
 - Youwei Zhuo* (<http://alchem.usc.edu/~youwei/>). Ph.D Candidate, Fall 2016 – Present (nominated for **Microsoft Research Ph.D Fellowship**, 2018).
 - Chao Wang* (<http://alchem.usc.edu/~chao/>). Ph.D Candidate, Fall 2016 – Present
 - Qinyi Luo* (<http://alchem.usc.edu/~qinyi/>). Ph.D Candidate, Fall 2017 – Present (nominated for **Microsoft Research Lovelace Fellowship**, 2018)
 - You Wu*. Ph.D Candidate, Fall 2017 – Present
 - Jinglei Cheng*. Ph.D Candidate, Fall 2018 – Present
 - *Undergraduate Internship Students*
 - Jinkun Lin*. Undergraduate Intern from Tsinghua, Summer 2018 (worked on distributed machine learning, paper appears in **ASPLOS'19**)
 - Jingji Chen*. Undergraduate Intern from Tsinghua, Summer 2018 (worked on distributed graph processing, paper in submission)
 - Jinglei Cheng*. Undergraduate Intern from Tsinghua, Summer 2017 (worked on data-parallel finite state machine, paper appears in **MICRO'18**, currently Ph.D student at USC)
 - Fangke Ye*. Undergraduate Intern from Tsinghua, Summer 2016 (currently Ph.D student at Gatech)
 - Andrea Cheng*. Undergraduate Intern from Northeastern University, Summer 2016
 - James Wang*. Undergraduate Intern from Case Western Reserve University, Summer 2016
-

SERVICE

- *External Service*
 - Program Chair of the 8th Workshop on Architecture and Systems for Big Data (ASBD 2018) at ISCA 2018 (<http://acs.ict.ac.cn/asbd2018/>).
 - Program Co-Chair of the 1st International Workshop on Architecture for Graph Processing (AGP-1) Workshop at ISCA 2017. (<https://sites.google.com/view/agp2017/home>).
 - PC member of HPCA 2019, MICRO 2018, HPCA 2018, ISCA 2017, SOCC 2017, HiPC 2017, ICPP 2016, IPDPS 2015, HiPC 2016.
 - External Review Committee of ASPLOS 2018, MICRO 2017, ASPLOS 2017, HPCA 2017, MICRO 2016, MICRO 2015, PACT 2016.
 - Web Co-Chair of ASPLOS 2018, 2019.
 - Sponsorship Chair of HPCA 2018, 2019.
 - Local Arrangement Chair of ISCA 2018.
 - Publication Chair of HPCA 2017.
 - Travel Grant Chair of ASPLOS 2017.
 - Submission Co-chair of ISCA 2012.
 - *Internal Service*
 - Organizer of **Computer Engineering Seminar** at USC.
 - Qualifying Exam committee of Qiumin Xu (2016), Abdulaziz Tabbakh (2017), Aditya Deshpande (2017), Ji Li (2017), Haonan Lu (2017) and Praveen Sharma (2018)
 - Thesis committee of Abdulaziz Tabbakh (2018)
 - Department committee of improving Ph.D student study environment.
-

GRANTS

- CRII: SHF: Improving Programmability of GPGPU/NVRAM Integrated Systems with Holistic Architectural Support (CCF-1657333), \$175,000.
- SHF: Small: Accelerating Graph Processing with Vertically Integrated Programming Model, Runtime and Architecture (CCF-1717754), \$450,000.

CSR: Small: Collaborative Research: GAMBIT: Efficient Graph Processing on a Memristor-based Embedded Computing Platform (CNS-1717984), \$250,000.

CAREER: Algorithm-Centric High Performance Graph Processing (CCF-1750656), \$450,000.
