

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.
 Hardware acceleration with emerging technologies.
 Big data and storage Systems.

EDUCATION

- **University of Illinois at Urbana-Champaign (UIUC)** Urbana, IL
Ph.D. in Computer Science (with Prof. Josep Torrellas) Aug. 2007 – Aug. 2013
 • Dissertation title: *Scalable and Flexible Bulk Architecture*
- **Institute of Computing Technology (ICT), Chinese Academy of Sciences (CAS)** Beijing
M.S. in Computer Engineering Sept. 2004 – Jul. 2007
- **Beihang University (BUAA)** Beijing
B.S. in Computer Engineering (with honor) Sept. 2000 – Jul. 2004

PROFESSIONAL EXPERIENCE

- Assistant Professor, *University of Southern California*. Aug. 2015 – Present
- Postdoctoral Researcher, *UC Berkeley*. Sept. 2013 – June 2015
- Research Intern, *Microsoft Research, Silicon Valley* Summer 2011
- Research Intern, *Microsoft Research, Redmond* Summer 2008

AWARDS AND HONORS

- **ACM Distinguished Speaker** 2017
- **One of 200 selected to participate Heidelberg Laureate Forum Foundation (HLFF)** 2015
- **W.J. Poppelbaum Memorial Award, UIUC** 2013
- **Nominated for an IBM Ph.D. Fellowship, UIUC** 2011
- **Nominated for a Microsoft Research Ph.D. Fellowship, UIUC** 2008
- **Andrew and Shana Laursen Fellowship, UIUC** 2007
- **NOKIA Fellowship, ICT, CAS** 2007
- **QIAN Changzhao Award, Beihang University** 2003
 Awarded to the top computer science student.

PUBLICATIONS

- 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**
- Mengxing Liu, Mingxing Zhang, Kang Chen, Xuehai Qian, Yongwei Wu, Weimin Zheng and Jinglei Ren. *DudeTX: Durable Decoupled Transaction*. To appear in **ACM Transactions on Storage**.
- 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**
- Abdulaziz Tabbakh, Murali Annavaram and Xuehai Qian. *Power Efficient Sharing-Aware GPU Data Management*. **IPDPS'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**
- Hui Wang, Rui Wang, Zhongzhi Luan, Xuehai Qian and Depei Qian. *Improving Multiprocessor Performance with Fine-grain coherence bypass*. **SCIENCE CHINA Information Sciences** 58(1), 2015.
- 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**
- Xuehai Qian, He Huang, Zhenzhong Duan, Junchao Zhang, Nan Yuan, Yongbin Zhou, Hao Zhang, Huimin Cui and Dongrui Fan. *Optimized Register Renaming Scheme for Stack-based x86 Floating Point Operations*. **ARCS'07**
- Xuehai Qian, He Huang, Hao Zhang, Junchao Zhang and Dongrui Fan. *Design and Implementation of Floating Point Stack on General RISC Architecture*. **PDP'07**
- Xuehai Qian, Hao Zhang, Jingang Yang, He Huang, Junchao Zhang and Dongrui Fan. *Circuit Implementation of Floating Point Range Reduction for Trigonometric Functions*. **ISCAS'07**

TALKS

- *Efficient Graph Processing Through Cross-Stack Co-Design*
 - Alibaba Group, Nov. 2017
 - Brown University, Oct. 2017

- *Reducing Data Movements in Graph Processing: from Distributed System to Emerging Technology*
 - University of Texas, Austin, Feb. 2017
 - Rice University, Feb. 2017
 - Cornell University, Dec. 2016
 - University of Illinois, Urbana-Champaign, Dec. 2016
 - University of Pennsylvania, Dec. 2016
 - VMware Research, Dec. 2016
 - Google, Dec. 2016
 - University of California, Santa Barbara, Jan. 2017
- *Breaking the Myth of "Think as a Vertex"*.
 - University of California, Irvine, Oct. 2016
- *Taming Relaxed-Consistency and Non-determinism in Parallel Systems*.
 - CALCM Seminar, Carnegie Mellon University, Jan. 2015.
 - Georgia Tech, Jan. 2015.
 - University of Pennsylvania, Feb. 2015
 - Syracuse University, Feb. 2015
 - Purdue University (CS), Feb. 2015
 - University of California Riverside, Mar. 2015
 - University of Virginia, Mar. 2015
 - Boston University, Mar. 2015
 - University of Southern California, Mar. 2015
 - Purdue University (ECE), Mar. 2015
- *Clearing Chaos: Taming Relaxed-Consistency in Multiprocessors*.
 - Intel Labs, Santa Clara, Jun. 2014.
 - AMD Research, Sunnyvale, Jun. 2014.
- *Scalable and Flexible Bulk Architecture*.
 - CECA, Peking University, Dec. 2012.
 - University at Buffalo, The State University of New York, Feb. 2013.
 - Washington University in St. Louis, Feb. 2013.
 - IIS, Tsinghua University, Feb. 2013.
 - Rutgers University, Mar. 2013.
 - Florida International University, Mar. 2013.
 - Penn State University, Mar. 2013.
 - The George Washington University, Mar. 2013.
 - Oracle, Apr. 2013.
 - Drexel University, Apr. 2013.
 - Microsoft Research, Asia (MSRA), May. 2013.
 - The University of Edinburgh, May. 2013.
 - The University of Manchester, May. 2013.
 - AMD Research Lab, June. 2013.
 - UC Berkeley, Jun. 2013.
 - Intel Labs, Santa Clara, Dec. 2013.

SERVICE

PC member of HPCA 2018, ISCA 2017, SOCC 2017, HiPC 2017, ICPP 2016, IPDPS 2015, HiPC 2016.

External Review Committee of MICRO 2017, ASPLOS 2017, HPCA 2017, MICRO 2016, MICRO 2015, PACT 2016.

Web Co-Chair of ASPLOS 2018.

Sponsorship Chair of HPCA 2018.

Local Arrangement Chair of ISCA 2018.

Publication Chair of HPCA 2017.

Travel Grant Chair of ASPLOS 2017.

Submission Co-chair of ISCA 2012.

Organizer of [Computer Engineering Seminar](#) at USC.

Member of IEEE and ACM.

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.
