

# Youwei Zhuo

(+1)213-255-6108/(+86)188-1136-7397

youweizh@usc.edu

www.alchem.usc.edu/~youwei

## RESEARCH

---

My research interest is system and architecture for accelerating irregular algorithms, including:

- Processing-In-Memory accelerators
- Parallel Finite State Machines
- Distributed system for graph processing
- Distributed system for deep learning on graphs

## EDUCATION

---

2016~Present	Ph.D., Department of Computer Science, <b>University of Southern California</b> Advisor: Prof. <b>Xuehai Qian</b>
2012~2016	B.E., Department of Computer Science, <b>Tsinghua University</b> GPA: 92.3/100 Rank: 5/118
2013~2016	B.A., School of Economics and Management, <b>Tsinghua University</b>

## AWARDS

---

07/2016	ISC'16 Student Supercomputing Competition 2nd place (Team leader)
11/2015	SC'15 Student Supercomputing Competition 1st place (Team leader)
07/2015	ISC'15 Student Supercomputing Competition 1st place (Team member)
05/2015	ASC'15 Student Supercomputing Competition 1st place (Team member)
2012-2016	Outstanding Student Award at Tsinghua University
2012	Outstanding Freshman Fellowship at Tsinghua University
2011	Outstanding Student Award in Jiangxi Province

## EXPERIENCE

---

05/2018~08/2018	<b>Optimize CNN Winograd algorithm on multicore CPU</b> Research Intern, Microsoft Research, Redmond. Advisor: <b>Yuxiong He</b> . Outperform state-of-the-art CNN library (NNPACK, MKLDNN) by 4x.
09/2015~03/2016	<b>Distributed Graph Processing and Machine Learning</b> Research Intern, Microsoft Research Asia, Beijing. Advisor: <b>Ming Wu</b> .
06/2015~09/2015	<b>Optimize BFS on HBM-enabled GPU</b> Research Intern, University of California, Santa Barbara. Advisor: Prof. <b>Yuan Xie</b> . First work on HBM implications for GPU programming. (DATE'17 and TVLSI'17)
2015-2016	<b>Student Supercomputing Competition.</b> Team leader and core member. Tsinghua University. Advisor: Prof. <b>Jidong Zhai</b> . The most significant student competition in HPC community. Built HPC cluster under power constraint and optimized HPC applications. First place in all 3 international competitions in 2015. (First time in history!)
2013~2016	<b>Optimize BFS on Sunway Taihulight Supercomputer.</b> Research Assistant, Tsinghua University. Advisor: Prof. <b>Wenguang Chen</b> . Developed debugging and profiling tools for Sunway architecture. Designed novel on-chip communication for Sunway NOC architecture. (IPDPS'17) Performance ranked 2nd in Graph500.

## PUBLICATION

---

### **Performance Evaluation and Optimization of HBM-Enabled GPU for Data-intensive Applications**

Maohua Zhu, Youwei Zhuo, Chao Wang, Wenguang Chen, Yuan Xie

The 2017 Conference on Design Automation and Test in Europe (DATE'17)

### **Performance Evaluation and Optimization of HBM-Enabled GPU for Data-intensive Applications**

Maohua Zhu, Youwei Zhuo, Chao Wang, Wenguang Chen, Yuan Xie

The IEEE Transactions on VLSI Systems (TVLSI'17)

### **Scalable Graph Traversal on Sunway TaihuLight with Ten Million Cores**

Heng Lin, Xiongchao Tang, Bowen Yu, Youwei Zhuo, Wenguang Chen, Jidong Zhai, Wanwang Yin, Weimin Zhen

The 19th IEEE International Parallel and Distributed Processing Symposium (IPDPS'17)

### **CIRCNN: Accelerating and Compressing Deep Neural Networks Using Block-Circulant Weight Matrices**

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

The 50th International Symposium on Microarchitecture (MICRO'17)

### **GraphR: Accelerating Graph Processing Using ReRAM**

Linghao Song, Youwei Zhuo, Xuehai Qian, Miao Hu, Hai Li, Yiran Chen

The 24th International Symposium on High Performance Computer Architecture (HPCA'18)

### **GraphP: Reducing Communication of PIM-based Graph Processing with Efficient Data Partition**

Mingxing Zhang, Youwei Zhuo (equal contribution), Chao Wang, Mingyu Guo, Yongwei Wu, Kang Chen, Christos Kozyrakis, Xuehai Qian

The 24th International Symposium on High Performance Computer Architecture (HPCA'18)

### **Wonderland: A Novel Abstraction-Based Out-Of-Core Graph Processing System**

Mingxing Zhang, Yongwei Wu, Youwei Zhuo, Xuehai Qian, Chenyin Huan, Kang Chen

The 23rd International Conference on Architectural Support for Programming Languages and Operating Systems Architecture (ASPLOS'18)

### **CSE: Parallel Finite State Machines with Convergence Set Enumeration**

Youwei Zhuo, Jinglei Cheng, Qinyi Luo, Jidong Zhai, Yanzhi Wang, Zhongzhi Luan, Xuehai Qian

The 51st International Symposium on Microarchitecture (MICRO'18)

### **HOP: Heterogeneity-Aware Decentralized Training**

Qinyi Luo, Jinkun Lin, Youwei Zhuo, Xuehai Qian

The 24th International Conference on Architectural Support for Programming Languages and Operating Systems Architecture (ASPLOS'19)

### **E-RNN: Design Optimization for Efficient Recurrent Neural Networks in FPGAs**

Zhe Li, Caiwen Ding, Siyue Wang, Wujie Wen, Youwei Zhuo, Chang Liu, Qinru Qiu, Wenyao Xu, Xue Lin, Xuehai Qian, Yanzhi Wang

The 25th International Symposium on High Performance Computer Architecture (HPCA'19)

### **HyPar: Towards Hybrid Parallelism for Deep Learning Accelerator Array**

Linghao Song, Jiachen Mao, Youwei Zhuo, Xuehai Qian, Hai Li, Yiran Chen

The 25th International Symposium on High Performance Computer Architecture (HPCA'19)

### **GraphQ: Scalable PIM-Based Graph Processing**

Youwei Zhuo, Chao Wang, Mingxing Zhang, Rui Wang, Dimin Niu, Yanzhi Wang, Xuehai Qian

The 52nd International Symposium on Microarchitecture (MICRO'19)

**AccPar: Tensor Partitioning for Heterogeneous Deep Learning Accelerator Arrays**

Linghao Song, Fan Chen, Youwei Zhuo, Xuehai Qian, Hai Li, Yiran Chen

The 26th International Symposium on High Performance Computer Architecture (HPCA'20)

**Prague: High-Performance Heterogeneity-Aware Asynchronous Decentralized Training**

Qinyi Luo, Jiaao He, Youwei Zhuo, Xuehai Qian

The 25th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'20)

**SympleGraph: Distributed Graph Processing with Precise Loop-carried Dependency Guarantee**

Youwei Zhuo, Jingji Chen, Qinyi Luo, Yanzhi Wang, Hailong Yang, Xuehai Qian

The 41st ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'20)

**TEACHING EXPERIENCE**

---

2017 Fall	Teaching Assistant: CSCI 526: Advanced Mobile Devices and Game Consoles
2018/2019/2020 Spring	Teaching Assistant: EE 451/CSCI 452: Parallel and Distributed Computing
2018 Fall	Teaching Assistant: CSCI 570: Analysis of Algorithms
2017 Summer	Mentor: <b>USC-THU Program</b> : Jinglei Chen (Undergraduate student)
2018 Summer	Mentor: <b>USC-THU Program</b> : Jingji Chen (Undergraduate student)

**OTHER EXPERIENCE**

---

06/2020	Presentation at PLDI'20
09/2019	Presentation at MICRO'19
10/2018	Presentation at MICRO'18 [ <a href="#">slides</a> ]
08/2018	Presentation at MSR [ <a href="#">slides</a> ]
05/2018	Student Volunteer at ISCA'18
02/2018	Presentation at HPCA'19 [ <a href="#">slides</a> ]