Zhenkun Cai

Rm 122A SHB, CUHK
Hong Kong, China
(+86) 18819451537
zekucai@gmail.com

Education

2018–2022 (Expected) **The Chinese University of Hong Kong**,
Ph.D. *in Computer Science and Engineering*,
Supervisor: *Prof. James Cheng*.

2013–2017 **South China University of Technology**,
B.Eng *in Computer Science and Technology*.

Research Interests

Distributed system, machine learning framework, cluster scheduling.

Internships

Oct 2021 - Now **Amazon Web Service, AI Lab**.
- **Applied Scientist Intern** in **DGL** Team
- Distributed sampling on GPUs for GNN training

Nov 2020 - Sept 2021 **Alibaba Group, Apsara Platform**.
- **Research Intern** in **Fuxi** Team
- Resource management for clusters with thousands of GPUs

July 2019 - Sept 2020 **HUAWEI, 2012 Lab**.
- **Research Intern** in **Mindspore** Team
- Auto-parallelism systems for training large-scale DNN models
- Efficient distributed GNN training systems

July 2018 - Aug 2018 **Alibaba Group, Apsara Platform**.
- **Research Intern** in **MaxCompute** Team
- Smart index selection for OLAP on big tables in data warehouses

July 2017 - July 2018 **The Chinese University of Hong Kong**.
- **Research Assistant** supervised by **Prof. James Cheng**
- Distributed Systems for large-scale data analytics

Awards

2018 - 2022 **Postgraduate Studentship, CUHK**.
2014 - 2016 **Undergraduate National Scholarship (three times), SCUT**.
2014 - 2016 **Merit Student, SCUT**.
2017 **ACM-ICPC World Finalist (Rank 56)**.
2016 **ACM-ICPC Regional Contest, Gold Medals**.
2016 **Guangdong Collegiate Programming Contest (GDCPC), Champion**.

Skills

Programming **C/C++, Python, CUDA**.
Systems **MPI, Tensorflow, Pytorch, NCCL**.

Talks

Nov 2021 *Towards Efficient Training for Large-scale Deep Learning Models*, ***AWS AI Lab***.
Apr 2021 *DGCL: An Efficient Communication Library for Distributed GNN Training*, ***Eurosys 2021***.

Projects

Systems for large-scale machine learning

TensorOpt: Training Large-scale DNNs with Auto-parallel.

- Supports distributed training of large models (e.g., Transformer, WideResNet) using limited GPU memory
- Optimally decides the parallelization strategy and automatically generates code for operators in a DNN
- Developed on top of TensorFlow with user-friendly Python APIs

EDL: An Elastic Deep Learning System on GPUs.

- Supports elastic deep learning, i.e., dynamically adjust the number of GPU at runtime
- Stop-free scaling and dynamic data pipeline on Horovod

FlexPS: A Parameter Server with Flexible Parallelism Control.

- A novel multi-stage abstraction to support flexible parallelism control in parameter server
- Optimizations to reduce the overhead of parallelism adjustment

Systems for graph neural networks

DGCL: A Distributed Graph Communication Library for GNN systems.

- A general library to scale single-GPU GNN systems (e.g., DGL and PyG) to the multi-GPU setting
- Efficient communcation kernels optimized for load balancing and bandwidth utilization on NVLinks, PCIe and InfiniBand

Seastar: A Vertex-centric GNN System.

- A vertex-centric programming model for GNNs
- Kernel optimizations such as operator fusion and vertex parallelism to reduce GPU memory consumption and improve training efficiency

Large-scale cluster scheduling

PPS: Fair And Efficient Scheduling for Multi-Tenant GPU Clusters.

- Probabilistic prediction based scheduling for clusters with thousands of GPUs
- Black-box and non-preemptive scheduling for GPU jobs
- Achieves both efficiency and fairness at the same time

Ursa: A Framework conducting both Resource Scheduling and Execution for OLAP Jobs.

- Captures dynamic resource needs at runtime and enables fine-grained, timely scheduling
- Achieves high resource utilization, which translates into shorter makespan and average JCT

Publications

[1] *TensorOpt: Exploring the Tradeoffs in Distributed DNN Training with Auto-Parallelism*, **TPDS 2022**
Z. Cai, K. Ma, X. Yan, Y. Wu, Y. Huang, J. Cheng, T. Su, F. Yu.

[2] *PPS: Fair and Efficient Scheduling for Multi-Tenant GPU Clusters*, **NSDI 2022** (Under Review)
Z. Cai, K. Ma, X. Yan, Y. Zhang, Z. Liu, Y. Feng, C. Li, J. Cheng.

[3] *DGCL: An Efficient Communication Library for Distributed GNN Training*, **Eurosys 2021**
Z. Cai, X. Yan, Y. Wu, K. Ma, J. Cheng, F. Yu.

[4] *Seastar: Vertex-Centric Programming for Graph Neural Networks*, **Eurosys 2021**
Y. Wu, K. Ma, **Z. Cai**, T. Jin, B. Li, C. Zheng, J. Cheng, F. Yu.

[5] *Elastic Deep Learning in Multi-Tenant GPU Clusters*, **TPDS 2021**
Y. Wu, K. Ma, X. Yan, Z. Liu, **Z. Cai**, Y. Huang, J. Cheng, H. Yuan, F. Yu.

[6] *Improving Resource Utilization by Timely Fine-Grained Scheduling*, **Eurosys 2020**
T. Jin, **Z. Cai**, B. Li, C. Zheng, G. Jiang, J. Cheng.

[7] *FlexPS: Flexible Parallelism Control in Parameter Server Architecture*, **VLDB 2018**
Y. Huang, T. Jin, Y. Wu, **Z. Cai**, X. Yan, Y. Guo, F. Yang, J. Li, and J. Cheng.

[8] *Scalable De Novo Genome Assembly Using Pregel*, **ICDE 2018**
D. Yan, H. Chen, **Z. Cai**, J. Cheng, and B. Shao.