

Qinghao Hu

S-Lab for Advanced Intelligence
Nanyang Technological University
ABN-02b-11, 61 Nanyang Avenue, Singapore 637335

E-Mail: qinghao.hu@ntu.edu.sg
Tel.: (+65) 8305 3277
Homepage: <https://tonyhao.xyz>

EDUCATION

- Nanyang Technological University, Singapore** 2020 ~ 2023
Ph.D. Student in Computer Science
Supervisor: [Prof. Tianwei Zhang](#) and [Prof. Yonggang Wen](#)
- National University of Singapore, Singapore** 2018 ~ 2020
Master in Electrical Engineering
- Zhejiang University, China** 2014 ~ 2018
Bachelor in Electrical Engineering

RESEARCH INTEREST

- Systems for Large Models
- Datacenter Management and Scheduling
- Machine Learning for Systems

PUBLICATION

Conference & Journal Papers

1. **Characterization of Large Language Model Development in the Datacenter**
[Qinghao Hu*](#), [Zhisheng Ye*](#), [Zerui Wang*](#), [Guoteng Wang](#), [Meng Zhang](#), [Qiaoling Chen](#), [Peng Sun](#), et al.
[\[NSDI '24\]](#) *USENIX Symposium on Networked Systems Design and Implementation*
2. **Hydro: Surrogate-Based Hyperparameter Tuning Service in Datacenters**
[Qinghao Hu](#), [Zhisheng Ye](#), [Meng Zhang](#), [Qiaoling Chen](#), [Peng Sun](#), [Yonggang Wen](#), [Tianwei Zhang](#)
[\[OSDI '23\]](#) *USENIX Symposium on Operating Systems Design and Implementation*
3. **Lucid: A Non-Intrusive, Scalable and Interpretable Scheduler for Deep Learning Training Jobs**
[Qinghao Hu*](#), [Meng Zhang*](#), [Peng Sun](#), [Yonggang Wen](#), [Tianwei Zhang](#)
[\[ASPLOS '23\]](#) *Architectural Support for Programming Languages and Operating Systems*
Distinguished Paper Award
4. **Primo: Practical Learning-Augmented Systems with Interpretable Models**
[Qinghao Hu](#), [Harsha Nori](#), [Peng Sun](#), [Yonggang Wen](#), [Tianwei Zhang](#)
[\[ATC '22\]](#) *USENIX Annual Technical Conference*
5. **Characterization and Prediction of Deep Learning Workloads in Large-Scale GPU Datacenters**
[Qinghao Hu](#), [Peng Sun](#), [Shengen Yan](#), [Yonggang Wen](#), [Tianwei Zhang](#)
[\[SC '21\]](#) *International Conference for High Performance Computing, Networking, Storage, and Analysis*
6. **Deep Learning Workload Scheduling in GPU Datacenters: A Survey**
[Zhisheng Ye*](#), [Wei Gao*](#), [Qinghao Hu*](#), [Peng Sun](#), [Xiaolin Wang](#), [Yingwei Luo](#), [Tianwei Zhang](#), et al.
[\[CSUR\]](#) *ACM Computing Surveys*

Under Review

1. **AMSP: Super-Scaling LLM Training via Advanced Model States Partitioning**
[Qiaoling Chen](#), [Qinghao Hu](#), [Zhisheng Ye](#), [Guoteng Wang](#), [Peng Sun](#), [Yonggang Wen](#), [Tianwei Zhang](#)
[\[Preprint\]](#) *Submitted to a Conference*

2. **Boosting Distributed Full-graph GNN Training with Asynchronous One-bit Communication**
Meng Zhang, Qinghao Hu, Peng Sun, Yonggang Wen, Tianwei Zhang
[\[Preprint\]](#) *Submitted to a Conference*
3. **FedDSE: Distribution-aware Sub-model Extraction for Federated Learning over Resource-constrained Devices**
Haozhao Wang, Yabo Jia, Meng Zhang, Qinghao Hu, Hao Ren, Peng Sun, Yonggang Wen, Tianwei Zhang
[\[Under Review\]](#) *Submitted to a Conference*

PROFESSIONAL SERVICE

[EuroSys '24] EuroSys Conference	Shadow Committee Member
[EuroSys '23] EuroSys Conference	Shadow Committee Member
[OSDI '22] USENIX Symposium on Operating Systems Design and Implementation	AE Committee Member
[ATC '22] USENIX Annual Technical Conference	AE Committee Member
[EuroSys '22] EuroSys Conference	AE Committee Member
[SOSP '21] ACM Symposium on Operating Systems Principles	AE Committee Member

TALK

Hydro: Surrogate-Based Hyperparameter Tuning Service in Datacenters <i>ChinaSys, Wuhan China</i>	<i>Jul. 2023</i>
<i>OSDI, Boston United States</i>	<i>Jul. 2023</i>
Lucid: A Non-Intrusive, Scalable and Interpretable Scheduling System <i>Huawei, Beijing China</i>	<i>May. 2023</i>
<i>MLSys Seminar Singapore</i>	<i>Apr. 2023</i>
<i>ASPLOS, Vancouver Canada</i>	<i>Mar. 2023</i>
Primo: Practical Learning Systems with Interpretable Models <i>ChinaSys, Nanjing China</i>	<i>Dec. 2022</i>
<i>ATC, Carlsbad California United States</i>	<i>Jul. 2022</i>
Scheduling in Large-Scale GPU Datacenters <i>National University of Singapore</i>	<i>Jan. 2022</i>
Characterization and Prediction of DL Workloads in Datacenters <i>SC, St. Louis Missouri United States</i>	<i>Nov. 2021</i>
Cluster Scheduling for Deep Learning <i>S-Lab for Advanced Intelligence, Singapore</i>	<i>Apr. 2021</i>

STUDENT MENTORED

I am fortunate to mentor and co-advise the research of the following students:

- **Ph.D. Students:** Meng Zhang, Zerui Wang (SJTU), Yutong Wu
- **Master Students:** Qiaoling Chen, Tan Ru Phing (NUS)
- **Undergraduates:** Peilin Chen, Amrita Ravishanker

HONOR

Google Ph.D. Fellowship	<i>2023</i>
Distinguished Paper Award of ASPLOS '23	<i>2023</i>
Youth Outstanding Paper Award of WAIC '23	<i>2023</i>
Student Travel Grant of OSDI '23	<i>2023</i>
Best Undergraduate Thesis Award	<i>2018</i>
Outstanding Graduates of Zhejiang University	<i>2018</i>