GUANZHOU HU

guanzhou.hu@wisc.edu ⋄ https://josehu.com

EDUCATION

University of Wisconsin—Madison *Aug 2020 - Present*
Ph.D. Student, Computer Science *Madison, WI, USA*

- Advised by Prof. Andrea Arpaci-Dusseau and Prof. Remzi Arpaci-Dusseau
- Research area: Operating systems, Storage systems, Caching, NVM devices

ShanghaiTech University *Sep 2016 - Jul 2020*
B. Eng., Computer Science & Technology *Shanghai, China*

- GPA: 3.9 / 4.0 (rank 2 / 183)
- Honors: Dean's Scholarship (2019), President's Scholarship (2017, 2018)
- Relevant coursework: Operating systems, Computer architecture III, Parallel computing

Massachusetts Institute of Technology *Sep 2019 - Jun 2020*
Special Student, Electrical Engineering & Computer Science *Cambridge, MA, USA*

- GPA: 4.0 / 4.0
- Relevant coursework: Distributed systems, Computer networks, Computer systems security

PUBLICATIONS & PATENTS

BORA: A Bag Optimizer for Robotic Analysis. Jian Zhang, Tao Xie, Yuzhuo Jing, Yanjie Song, Guanzhou Hu, Si Chen, and Shu Yin. 2020. In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC '20). IEEE Press, Article 12, 1–15.

A Storage System Management Policy Based on Data Content Locality. Yin, S. and Hu, G. 2019. CN. Patent application 201910499391.9, filed in June 2019.

RESEARCH EXPERIENCE

Scalable & Affordable GCNs with Serverless Computing *Jul 2019 - Oct 2019*
CSST Research Intern, UCLA, with Prof. Harry Xu *Los Angeles, CA, USA*

- Integrated serverless computing into graph computing to build an affordable, efficient, and scalable graph convolutional networks (GCNs) computation platform without dedicated GPUs.
- Implemented the first workable prototype with AWS Lambdas service, and reached linear scalability and 100% cost-efficiency.

NcTrace: Optimized Trace Data Storage with the netCDF Format *Mar 2019 - Aug 2019*
Leader of Project Team, ShanghaiTech University, with Prof. Shu Yin *Shanghai, China*

- Optimized storage of comma-separated values (CSV) trace data using the netCDF format. Introduced the "dimension packing" model which reduces file size and accelerates analysis tasks.
- Tested with Google cluster traces, and achieved 7:1 size reduction with 2 orders of magnitude acceleration on reading.

Active I/O: Parallel Content-Aware Storage System *Jan 2019 - Aug 2019*
Research Assistant, ShanghaiTech University, with Prof. Shu Yin *Shanghai, China*

- Designed a high-performance, parallel file system which aims to dig out the "content locality" within highly-structured data formats, by clustering data by topics and providing users a better locality when operating on a subset of topics.
- Tested with Robot Operating System bag files, and achieved 6.5x performance improvement on opening and at least 1.4x on reading.

TEACHING EXPERIENCE

Teaching Assistant in Computer Architecture *Feb 2019 - Apr 2019*
School of Information Science & Technology *Shanghai, China*

Teaching Assistant in Operating Systems *Sep 2018 - Jan 2019*
School of Information Science & Technology *Shanghai, China*

- Guided semester-long course projects on the PintOS system kernel from Stanford CS140.

Teaching Assistant in Discrete Mathematics *Mar 2018 - Jul 2018*
School of Information Science & Technology *Shanghai, China*

PRIZES & AWARDS

Outstanding Research Award, CSST Program 2019, UCLA *Sep 2019*

Second Class Prize, ASC Supercomputing Competition 2019 (GeekPie_HPC team leader) *Mar 2019*

Outstanding Teaching Assistant Award, ShanghaiTech University *Jan 2019*

Meritorious Winner, Mathematical Contest in Modelling (MCM) 2018 *Apr 2018*

MISCELLANEOUS

- **Skills**: System programming, C/C++, Rust, Go, Python, Linux server dev/ops, x86, MIPS
- **Languages**: Chinese (native), English (fluent)