

Yang Wang

CONTACT INFORMATION	Department of Computer Science and Engineering The Ohio State University 2015 Neil Avenue Columbus, Ohio 43210, USA	<i>Office:</i> DL 689 <i>Phone:</i> (614) 292-2577 <i>E-mail:</i> wang.7564@osu.edu http://web.cse.ohio-state.edu/~yangwang
RESEARCH INTERESTS	Large-scale distributed systems, fault tolerance, scalability, performance debugging	
EDUCATION	Ph.D. Computer Science, The University of Texas at Austin, 2014 Advisors: Dr. Lorenzo Alvisi (now at Cornell University) and Dr. Mike Dahlin (now at Google) M.E. Computer Science and Technology, Tsinghua University, 2008 B.E. Computer Science and Technology, Tsinghua University, 2005	
ACADEMIC APPOINTMENTS	Visting Researcher Associate Professor Assistant Professor Research Assistant Teaching Assistant	Meta, 2022 - Present Computer Science, the Ohio State University, 2022 - Present Computer Science, the Ohio State University, 2015 - 2022 Computer Science, the University of Texas at Austin, 2009 - 2014 Computer Science, the University of Texas at Austin, 2008 - 2009
HONORS AND AWARDS	Google PhD Fellowship in Distributed Computing, 2013 - 2014 Best paper award, Systor 2014	
PUBLICATIONS	Yujie Hui, Drew Ripberger, Xiaoyi Lu, and Yang Wang. “Learning Distributed Protocols with Zero Knowledge.” Accepted by ML for Systems Workshop at NeurIPS 2023. Yi Liu, Minmei Wang, Shouqian Shi, Yang Wang, and Chen Qian. “EdgeCut: Fast and Low-overhead Access of User-associated Contents from Edge Servers”. Conditionally accepted by the 8th ACM/IEEE Symposium on Edge Computing (SEC 2023), Wilmington, DE, Dec 2023. Shaobu Wang, Guangyan Zhang, Junyu Wei, Yang Wang, Jiesheng Wu, and Qingchao Luo. “Understanding Silent Data Corruptions in a Large Production CPU Population.” <i>Proceedings of the 29th ACM Symposium on Operating Systems Principles (SOSP 2023)</i> , Koblenz, Germany, Oct 2023. Tianxi Li, Yang Wang, and Xiaoyi Lu. “On the Discontinuation of Persistent Memory: Looking Back to Look Forward”. In the Workshop on Hot Topics in System Infrastructure (HotInfra’23), co-located with ISCA 2023, Orlando, FL, Jun 2023. Boris Grubic, Yang Wang, Tyler Petrochko, Ran Yaniv, Brad Jones, David Callies, Matt Clarke-Lauer, Dan Kelly, Soteris Demetriou, Kenny Yu, and Chunqiang Tang. “Conveyor: One-Tool-Fits-All Continuous Software Deployment at Meta”. <i>Proceedings of the 17th USENIX Symposium on Operating Systems Design and Implementation (OSDI 2023)</i> , Boston, MA, Jul 2023. Chaoyi Cheng*, Mingzhe Han*, Nuo Xu, Spyros Blanas, Michael D. Bond, and Yang Wang (* made equal contributions). “Developer’s Responsibility or Database’s Responsibility? Rethinking Concurrency Control in Databases”. <i>Proceedings of the Conference on Innovative Data Systems Research (CIDR) 2023</i> , Amsterdam, The Netherlands, Jan, 2023.	

Junyu Wei, Guangyan Zhang, Junchao Chen, Yang Wang, Weimin Zheng, Tingtao Sun, Jiesheng Wu, and Jiangwei Jiang. “LogGrep: Fast and Cheap Cloud Log Storage by Exploiting both Static and Runtime Patterns”. *Proceedings of Eurosys 2023*, Rome, Italy, May, 2023.

Drew Ripberger, Yifan Gan, Xueyuan Ren, Spyros Blanas, and Yang Wang. “IsoBugView: Interactively Debugging Isolation Bugs in Database Applications”. *Proceedings of the VLDB Endowment*, Vol. 15, No.12 (VLDB 22 Demo), Sydney, Australia, September, 2022.

Yang Wang, Miao Yu, Yujie Hui, Fang Zhou, Yuyang Huang, Rui Zhu, Xueyuan Ren, Tianxi Li, and Xiaoyi Lu. “A Study of Database Performance Sensitivity to Experiment Settings”. *Proceedings of the VLDB Endowment*, Vol. 15, No.7 (VLDB 22), Sydney, Australia, September, 2022.

Yipei Niu, Panpan Jin, Jian Guo, Yikai Xiao, Rong Shi, Fangming Liu, Chen Qian, and Yang Wang. “PostMan: Rapidly Mitigating Bursty Traffic via On-demand Offloading of Packet Processing”. In *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, June 2021.

Sixiang Ma, Fang Zhou, Mike D. Bond, and Yang Wang. “Finding Heterogeneous-Unsafe Configuration Parameters in Cloud Systems”. *Proceedings of Eurosys 2021*.

Junyu Wei, Guangyan Zhang, Yang Wang, Zhiwei Liu, Zhanyang Zhu, Junchao Chen, Tingtao Sun, and Qi Zhou. “On the Feasibility of Parser-based Log Compression in Large-Scale Cloud Systems”. *Proceedings of the 19th USENIX Conference on File and Storage Technologies (FAST)*, Online, Feb, 2021.

Huazhe Wang, Xin Li, Yang Wang, Yu Zhao, Yu Ye, Hongkun Yang, and Chen Qian. “SICS: Secure and Dynamic Middlebox Outsourcing”. In *IEEE/ACM Transactions on Networking (ToN)*, Vol 6, Issue 28, Sep 2020.

Yifan Gan, Xueyuan Ren, Drew Ripberger, Spyros Blanas, Yang Wang. “IsoDiff: Debugging Anomalies Caused by Weak Isolation”. *Proceedings of the VLDB Endowment*, Vol. 13, No.11 (VLDB 20), Online, August - September, 2020.

Youmin Chen, Youyou Lu, Fan Yang, Qing Wang, Yang Wang, Jiwu Shu. “FlatStore: an Efficient Log-Structured Key-Value Storage Engine for Persistent Memory”. *Proceedings of the 25th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS 20)*, Online, March, 2020.

Zhufan Wang, Guangyan Zhang, Yang Wang, Qinglin Yang, and Jiaji Zhu. “Dayu: Fast and Low-interference Data Recovery in Very-large Storage Systems”. *Proceedings of the 2019 USENIX Annual Technical Conference (USENIX ATC)*, Renton, WA, July, 2019.

Panpan Jin, Jian Guo, Yikai Xiao, Rong Shi, Yipei Niu, Fangming Liu, Chen Qian, and Yang Wang. “PostMan: Rapidly Mitigating Bursty Traffic by Offloading Packet Processing”. *Proceedings of the 2019 USENIX Annual Technical Conference (USENIX ATC)*, Renton, WA, July, 2019.

Donghe Kang, Vadang Patel, Ashwathi Nair, Spyros Blanas, Yang Wang, and Srinivasan Parthasarathy. “Henosis: Workload-driven small array consolidation and placement for HDF5 applications on heterogeneous data stores”. *Proceedings of the 33rd ACM International Conference on Supercomputing (ICS)*, Phoenix, AZ, June 2019.

Chengwen Wu, Guangyan Zhang, Yang Wang, Xinyang Jiang, and Weimin Zheng. “Redio: Accelerating Disk-based Graph Processing by Reducing Disk I/Os”. *IEEE Transactions on Computers*, Vol 68, Issue 3, March 2019.

Fang Zhou, Yifan Gan, Sixiang Ma and Yang Wang. “wPerf: Generic Off-CPU Analysis to Identify Critical Waiting Events”. *Proceedings of the 13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 2018)*, Carlsbad, CA, October, 2018.

Rong Shi, Yifan Gan and Yang Wang. “Evaluating Scalability Bottlenecks by Workload Extrapolation”. *Proceedings of the 26th IEEE International Symposium on the Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2018)*, Milwaukee, WI, September, 2018.

Sixiang Ma and Yang Wang. “Accurate Timeout Detection Despite Arbitrary Processing Delays”. *Proceedings of the 2018 USENIX Annual Technical Conference (USENIX ATC)*, Boston, MA, July, 2018.

Rong Shi and Yang Wang. “Cheap and Available State Machine Replication”. *Proceedings of the 2016 USENIX Annual Technical Conference (USENIX ATC)*, Denver, CO, June, 2016.

Chao Xie, Chunzhi Su, Cody Little, Lorenzo Alvisi, Manos Kapritsos, and Yang Wang. “High-Performance ACID via Modular Concurrency Control”. *Proceedings of the 25th ACM Symposium on Operating Systems Principles (SOSP)*, Monterey, California, October, 2015.

Chao Xie, Chunzhi Su, Manos Kapritsos, Yang Wang, Navid Yaghmazadeh, Lorenzo Alvisi, and Prince Mahajan. “Salt: Combining ACID and BASE in a Distributed Database”. *Proceedings of the 11th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, Broomfield, CO, October, 2014.

Mark Silberstein, Lakshmi Ganesh, Yang Wang, Lorenzo Alvisi, and Mike Dahlin. “Lazy Means Smart: Reducing Repair Bandwidth Costs in Erasure-coded Distributed Storage”. *Proceedings of the 7th ACM International Systems and Storage Conference (Systor)*, Haifa, Israel, June, 2014. Best paper award.

Yang Wang, Manos Kapritsos, Lorenzo Alvisi, and Mike Dahlin. “Exalt: Empowering Researchers to Evaluate Large-Scale Storage Systems”. *Proceedings of the 11th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, Seattle, WA, April, 2014.

Yang Wang, Manos Kapritsos, Zuocheng Ren, Prince Mahajan, Jeevitha Kirubanandam, Lorenzo Alvisi, and Mike Dahlin. “Robustness in the Salus scalable block store”, *Proceedings of the 10th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, Lombard, IL, 2013.

Manos Kapritsos, Yang Wang, Vivien Quema, Allen Clement, Lorenzo Alvisi, and Mike Dahlin. “All about Eve: Execute-Verify Replication for Multi-Core Servers”, *Proceedings of 2012 Symposium on Operating Systems Design and Implementation (OSDI)*, Hollywood, CA, October 2012.

Yang Wang, Lorenzo Alvisi, and Mike Dahlin. “Gnothi: Separating Data and Metadata for Efficient and Available Storage Replication”, *Proceedings of 2012 USENIX Annual Technical Conference (USENIX ATC)*, Boston, MA, June 2012.

Yang Wang, Jiwu Shu, Guangyan Zhang, Wei Xue, and Weimin Zheng. “SOPA: Selecting the Optimal Policy Adaptively”. *ACM transactions on storage (TOS)*, Volume 6 Issue 2, July 2010.

Allen Clement, Manos Kapritsos, Sangmin Lee, Yang Wang, Lorenzo Alvisi, Mike Dahlin, and Taylor Riche. “UpRight Cluster Services”. *Proceedings of the 22nd ACM Symposium on Operating Systems Principles (SOSP)*, Big Sky, MT, October 2009.

PAST STUDENTS Sixiang Ma (Ph.D., 2015 - 2022), “Mitigating Distributed Configuration Errors in Cloud Systems”.

Currently at Databricks.

Yifan Gan (Ph.D., 2015 - 2022), “Exploring Transaction Anomalies under Weak Isolation Levels for General Database Applications”. Currently at Oracle.

Fang Zhou (Ph.D., 2016 - 2021), “Identifying and Understanding Performance Problems in Software Systems”. Currently at Meta.

Rong Shi (Ph.D., 2015 - 2018), “Efficient data and metadata processing in large-scale distributed systems”. Currently at Meta

Chaoyi Cheng (BS 2022). Currently master at University of Pennsylvania.

Nuo Xu (BS 2022). Currently master at University of Southern California.

Mingzhe Han (BS 2022). Currently master at the Ohio State University.

Yuyang Huang (BS 2021). Currently Ph.D. at University of Chicago.

Rui Zhu (BS 2021). Currently master at Purdue.

Zijian Hu (MS 2017). Currently at Meta.

CURRENT STUDENTS

Yujie Hui (Ph.D., 2021 - Present, co-advised with Dr. Xiaoyi Lu at UC Merced)

Tianxi Li (Ph.D., 2021 - Present, co-advised with Dr. Xiaoyi Lu at UC Merced)

Miao Yu (Ph.D., 2020 - Present)

Xueyuan Ren (Ph.D., 2019 - Present)

Drew Ripberger (Undergraduate, 2019 - Present)

TEACHING

CSE6431 Advanced Operating System, Fall 2021.

CSE2421 Computer Organization (System I), Fall 2021.

CSE2431 Operating System (System II), Fall 2021.

CSE2431 Operating System (System II), Spring 2021.

CSE2431 Operating System (System II), Fall 2020.

CSE6431 Advanced Operating System, Spring 2020.

CSE6431 Advanced Operating System, Fall 2019.

CSE5429 System Design for New Hardware, Spring 2019.

CSE6431 Advanced Operating System, Fall 2018.

CSE6431 Advanced Operating System, Fall 2017.

CSE2431 Operating System (System II), Spring 2017.

CSE2431 Operating System (System II), Fall 2016.

CSE5439 Large-scale Data Processing, Spring 2016.

	CSE2431 Operating System (System II), Fall 2015.
DEPARTMENT SERVICE	<p>Co-organizer of the exploreCSR@OSU event (2020 - 2021).</p> <p>Member of the Graduate Admission Committee (2019 - 2021).</p> <p>Co-organizer of the Graduate Student Visit Day (2020).</p> <p>Member of the Graduate Studies Committee (2020 - 2021).</p> <p>Member of the Diversity and Inclusion Committee (2019 - 2020).</p> <p>Member of the Qualification Exam Committee (2016 - 2021).</p> <p>Member of the Curriculum/Undergraduate Study Committee (2015 - 2019).</p>
INVITED TALKS	<p>“Conveyor: One-Tool-Fits-All Continuous Software Deployment at Meta”. Purdue University, Aug 2023.</p> <p>“Conveyor: One-Tool-Fits-All Continuous Software Deployment at Meta”. University of Illinois Urbana-Champaign, Aug 2023.</p> <p>“Continuous and Safe Software Deployment”. Meta Core Systems Faculty Summit, Oct 2022.</p> <p>“Continuous and Safe Software Deployment”. Meta Core Systems Faculty Summit, Oct 2022.</p> <p>“Why is My Program Slow With No Obvious Reason? Understanding Performance Problems Caused by Kernel and Hardware Events”. ByteDance, Jul 2021.</p> <p>“Accurate and Low-overhead Off-CPU Performance Analysis”. University of Central Florida, Nov 2019.</p> <p>“Accurate and Low-overhead Off-CPU Performance Analysis”. Cornell University, Nov 2019.</p> <p>“Accurate and Efficient Off-CPU Performance Analysis”. Uber, Jul 2019.</p> <p>“wPerf: generic Off-CPU analysis to identify bottleneck waiting events”. University of Michigan, Nov 2018.</p> <p>“wPerf: generic Off-CPU analysis to identify bottleneck waiting events”. University of California at Santa Cruz, Aug 2018.</p> <p>“Accurate Timeout Detection Despite Arbitrary Processing Delays”. Tsinghua University, 2018.</p> <p>“Cheap and Available State Machine Replication”, Shanghai Jiaotong University, 2016.</p>
PROFESSIONAL ACTIVITIES	<p>Program committee member: Eurosys 2024, SOSP 2024, OSDI 2023 (Student Research Competition and Poster), Eurosys 2023, Sigmod 2023, Eurosys 2022, APSys 2021, ICDCS 2021, Eurosys 2021, OSDI 2020, USENIX ATC 2019 ERC, ICDCS 2018, SoCC 2016, SysTOR 2016, DCOSS 2016, ICPP 2015</p> <p>Student grant committee member: SOSP 2019, SOSP 2017</p> <p>Reviewer: <i>ACM Transactions on Storage (TOS)</i>, <i>IEEE Transactions on Parallel and Distributed Systems (TPDS)</i>, <i>ACM Transactions on Computer Systems (TOCS)</i>, <i>ACM Transactions on Parallel</i></p>

Computing (TOPC)

GRANTS

07/2022 - 06/2023. Built environmental and infrastructural factors and safe driving among teens. TDAI Pilot Award from the Ohio State University. Co-PI. \$50,000 in total; my share is \$13,000.

10/2021 – 9/2026. Collaborative Research: PPoSS: LARGE: ScaleStuds: Foundations for Correctness Checkability and Performance Predictability of Systems at Scale. National Science Foundation. Site PI at OSU. \$5,000,000 in total; my share is \$624,650.

5/2021-8/2021. Adoption and Implementation of an Evidence-based Safe Driving Program for High-Risk Teen Drivers. Research contract from Nationwide Children’s Hospital. PI. \$31,223 in total.

7/2019 - 12/2019. Intervention to improve driving practices among high-risk teen drivers. Research contract from Nationwide Children’s hospital. PI. \$26,068.00 in total.

7/2019 - Present. CNS Core: SMALL: Clarifying experimenter bias by identifying and visualizing experiment bottlenecks. National Science Foundation. Single PI. \$496,893 in total.

8/2017 - 8/2018. SHF: EAGER: HI-HDFS - Holistic I/O optimizations for the Hadoop distributed filesystem. National Science Foundation. Co-PI. \$150,000 in total; my share is \$50,000.

4/2016 - 4/2019. CRII: CSR: Efficient and Available Replication in Large-scale Datacenters. National Science Foundation. Single PI. \$175,000 in total.