

Yiming Zhang

Dr. Zhang is a professor at School of Informatics, Xiamen University. He is also the director of NICEX Lab (Networked Intelligent Computing at EXascale). His research is focused on operating/networking/storage/AI systems. As the first/corresponding author, he has published many highly-influential papers in top conferences including NSDI, FAST, EuroSys, ATC, VLDB, and famous Journals including TOCS, TON, TOS, TC, JSAC, TPDS. He is an associate editor of IEEE TSC. He has received the National Science and Technology Award (2nd prize), the Natural Science Award of Hunan province (1st prize), and the (top 10) Outstanding Dissertation Award. His research (TianheGraph) ranks No. 1 in the latest Graph500 on Tianhe Supercomputer.



E-mail: zhangyiming@xmu.edu.cn

Interest

My current interests include:

- Storage systems;
- Distributed machine learning systems;
- Secure and encrypted computing;
- Graph computing.

Education

- 2004~2008, **Ph.D. (with honors)** in School of Computer, National University of Defense Technology (NUDT), China.
Advisor: Prof. Xicheng Lu
Dissertation Title: Research on Efficient Overlay Construction in Virtual Computing Environment. The dissertation was awarded the **Outstanding Dissertation Award** by China Computer Federation (CCF) in 2011.
- 2001~2003, **M.S.** in School of Mechanics & Electronics, NUDT, China
Advisor: Prof. Guoxi Li
- 1997~2001, **B.Eng. (with honors)** in School of Mechanics & Electronics, NUDT, China
Ranked first among 67 students in School of Mechanics & Electronics, NUDT

Experiences

- Dec. 2021 ~ Now: **Professor**, *School of Informatics, Xiamen University.*
- Dec. 2013 ~ Dec. 2021: **Associate Professor**, *National Laboratory for Parallel and Distributed Processing (PDL), NUDT.*
- Dec. 2017 ~ Dec. 2020: **Director**, *Big Data and AI Lab, NUDT.*
- July. 2018 ~ Oct. 2018: **Visiting Professor**, *Microsoft Research Asia (MSRA). Studying large scale storage systems*
- Jun. 2012 ~ Jun. 2013: **Visiting Researcher**, *System Research Group, Computer Lab, University of Cambridge.* Implementing the KylinX OS for specialized VM appliances in the cloud.
- Dec. 2008 ~ Dec. 2013: **Assistant Professor**, *National Laboratory for Parallel and Distributed Processing (PDL), NUDT.*
- Apr. 2011 ~ Nov. 2011: **Visiting Researcher**, *Microsoft Research Asia (MSRA).* RAM-based KV.

Selected Publications

Recent Conference Papers:

- Ruiming Lu, Erci Xu, **Yiming Zhang** (Corresponding), Guangtao Xue, et al. NVMe SSD Failures in the Field: the Fail-Stop and the Fail-Slow. **USENIX ATC** 2022.
- Xinbiao Gan, **Yiming Zhang** (Corresponding), Ruigeng Zeng, Jie Liu, Ruibo Wang, Tiejun Li, Li Chen, Kai Lu. XTREE: Traversal-Based Partitioning for Extreme-Scale Graph Processing on Supercomputers. **IEEE ICDE** 2022.
- Wenhao Lv, Youyou Lu, **Yiming Zhang**, Peile Duan, Jiwu Shu. InfiniFS: An Efficient Metadata Service for Large-Scale Distributed Filesystems. **USENIX FAST** 2022.
- Lujia Yin*, **Yiming Zhang*** (Corresponding), Zhaoning Zhang, et al. ParaX: Boosting Deep Learning on Many-Core CPUs. **VLDB** 2021.
*Co-primary authors
- Lujia Yin, Li Wang, **Yiming Zhang*** (Corresponding), Yuxing Peng. MapperX: Adaptive Metadata Maintenance for Fast Crash Recovery of DM-Cache Based Hybrid Storage Devices. **USENIX ATC** 2021.
- Li Wang*, **Yiming Zhang*** (Corresponding), Jiawei Xu, Guangtao Xue. MapX: Controlled Data Migration in the Expansion of Decentralized Object-Based Storage Systems. **USENIX FAST** 2020.
*Co-primary authors
- Huiba Li*, **Yiming Zhang*** (Corresponding), Haonan Wang, Ping Zhong. UrsaL: Ultra-Efficient, Reliable, Scalable, and Available Block Storage at Low Cost. **IEEE INFOCOM** 2020.
*Co-primary authors
- Dongsheng Li*, **Yiming Zhang*** (Corresponding), Jinyan Wang, Kian-Lee Tan. TopoX: Topology Refactorization for Efficient Graph Partitioning and Processing. **VLDB** 2019.
*Co-primary authors
- Huiba Li*, **Yiming Zhang*** (Corresponding), Zhiming Zhang, Shengyun Liu, et al. URSA: Hybrid Block Storage for Cloud-Scale Virtual Disks. **ACM EuroSys** 2019.
*Co-primary authors
- **Yiming Zhang**, Jon Crowcroft, Dongsheng Li, et al. KylinX: A Dynamic Library Operating System for Simplified and Efficient Cloud Virtualization. **USENIX ATC** 2018.
- Jiaxin Li, Yuxi Chen, Haopeng Liu, Shan Lu, **Yiming Zhang**, et al. PCatch: Automatically Detecting Performance Cascading Bugs in Cloud Systems. **ACM EuroSys** 2018.
- Huiba Li, **Yiming Zhang** (Corresponding), Zhiming Zhang, et al. PARIX: Speculative Partial Writes in Erasure-Coded Systems. **USENIX ATC** 2017.
- **Yiming Zhang**, Dongsheng Li, et al. CubeX: Leveraging Glocality of Cube-Based Networks for RAM-Based Key-Value Store. **IEEE INFOCOM** 2017.
- Ziyang Li, **Yiming Zhang**, Dongsheng Li, et al. OPTAS: Decentralized Flow Monitoring and Scheduling for Tiny Tasks. **IEEE INFOCOM** 2016.
- **Yiming Zhang**, Chuanxiong Guo, Dongsheng Li, Rui Chu, Yongqiang Xiong, et al. CubicRing: Enabling One-Hop Failure Detection and Recovery for Distributed In-Memory Storage Systems. **USENIX NSDI** 2015.

Journal Papers:

- **Yiming Zhang**, Lujia Yin, Dongsheng Li, Yuxing Peng, Kai Lu. ParaX: Bandwidth-Efficient Instance Assignment for DL on Multi-NUMA Many-Core CPUs. **IEEE Transactions on Computers (TC)**. Published online, doi: 10.1109/TC.2022.3145164.
- **Yiming Zhang**, Kai Lu, Wenguang Chen. Processing Extreme-Scale Graphs on China's

Supercomputers. **Communications of the ACM (CACM)**, 64(11): 60-63, 2021.

- Xinbiao Gan, **Yiming Zhang** (Corresponding), Ruibo Wang, Chunye Gong, Tiaojie Xiao, Ruigeng Zeng, Jie Liu, Kai Lu. TianheGraph: Customizing Graph Search for Graph500 on Tianhe Supercomputer. **IEEE Transactions on Parallel and Distributed Systems (TPDS)**. Published online, doi: 10.1109/TPDS.2021.3100785.
- **Yiming Zhang**, Chengfei Zhang, Yaozheng Wang, Kai Yu, Jon Crowcroft. KylinX: Simplified Virtualization Architecture for Specialized Virtual Appliances. **ACM Transactions on Computer Systems (TOCS)**, 37(1-4): 1-27, 2021.
- **Yiming Zhang**, Dongsheng Li, et al. GraphA: Efficient Partitioning and Storage for Distributed Graph Computation. **IEEE Transactions on Services Computing (TSC)**, 14(1): 155-166, 2021.
- Zhiyao Hu, Dongsheng Li, Dongxiang Zhang, **Yiming Zhang**, Baoyun Peng. Optimizing Resource Allocation for Data-Parallel Jobs Via GCN-Based Prediction. **IEEE Transactions on Parallel and Distributed Systems (TPDS)**, 32(9): 2188-2201, 2021.
- **Yiming Zhang**, Haonan Wang, Menghan Jia, Jinyan Wang, Dongsheng Li, Guangtao Xue, Kian-Lee Tan. TopoX: Topology Refactorization for Minimizing Network Communication in Graph Computations. **IEEE/ACM Transactions on Networking (ToN)**, 28(6): 2768-2782, 2020.
- **Yiming Zhang**, Huiba Li, Shengyun Liu, et al. PBS: Efficient Erasure-Coded Block Storage System Based on Speculative Partial Write. **ACM Transactions on Storage (ToS)**, 16(1): 1-25, 2020.
- Dongsheng Li, Zhiyao Hu, Zhiquan Lai, **Yiming Zhang**, Kai Lu. Coordinative Scheduling of Computation and Communication in Data-parallel Systems. **IEEE Transactions on Computers (TC)**. Doi: 10.1109/TC.2020.3039238.
- **Yiming Zhang**, Dongsheng Li, Ling Liu. Leveraging Glocality for Fast Failure Recovery in Distributed RAM Storage. **ACM Transactions on Storage (ToS)**, 15(1): 1-24, 2019.
- **Yiming Zhang**, Dongsheng Li, et al. CubicRing: Exploiting Network Proximity for Distributed In-Memory Key-Value Store. **IEEE/ACM Trans. on Networking (ToN)**, 25(4): 2040-2053, 2017.
- **Yiming Zhang**, Dongsheng Li, Zhigang Sun, Feng Zhao, Jinshu Su, Xicheng Lu. CSR: Classified Source Routing in DHT-Based Networks. **IEEE Trans. Cloud Computing (TCC)**, 6(2): 464-477, 2018.
- **Yiming Zhang**, Rong Chang, Paul Townend. Special Section on Virtualization and Services for Cloud-Based Application Systems. **IEEE Transactions on Services Computing (TSC)**, 12(1), 2019.
- Feng Huang, **Yiming Zhang** (Corresponding), Dongsheng Li, et al. meGautz: A High Capacity, Fault-tolerant and Traffic Isolated Modular Datacenter Network. **IEEE Transactions on Services Computing (TSC)**, 11(1): 117-130, 2018.
- Ting Li, **Yiming Zhang** (Corresponding), Hao Liu, et al. Fast Compressive Spectral Clustering for Large-Scale Sparse Graph. **IEEE Transactions on Big Data (TBD)**. Published online. Doi: 10.1109/TBDDATA.2019.2931532.
- **Yiming Zhang**, Ling Liu. Distributed Line Graphs: A Universal Technique for Designing DHTs Based on Arbitrary Regular Graphs. **IEEE Transactions on Knowledge and Data Engineering (TKDE)**, 24(9): 1556-1569, 2012.
- **Yiming Zhang**, Lei Chen, Xicheng Lu, Dongsheng Li. Enabling Routing Control in a DHT. **IEEE Journal on Selected Areas in Communications (J-SAC)**, 28(1): 28-38, 2010.

Awards

- **Science and Technology Award, First Class**, granted by *China Computer Federation (CCF)*, Dec. 2021
- **No. 1 in the Graph500 List (TianheGraph)**, ranked by *Graph500 steering committee*, June 2021
- **Outstanding Dissertation Award**, awarded by *China Computer Federation (CCF)* to the top 10 dissertations chosen nationally, Jan. 2011
- **National Science and Technology Award, Second Class**, granted by the *Chinese Ministry of Science and Technology*, Dec. 2012
- **Natural Science Award, First Class**, granted by the *Hunan Department of Science and Technology*, Sep. 2018
- **Distinguished Chinese Student Scholarship**, granted by the *Chinese Ministry of Education* to the top 12 PhD students chosen nationally, Jul. 2008

Professional Activities

- Member of **2022 EuroSys Roger Needham Award Committee**
- Associate editor of **IEEE Transactions on Computers**
- Associate editor of **IEEE Transactions on Services Computing**
- Chair of **CCF Outstanding Dissertation Award Winners Group (2016-2017)**
- Editor of **TSC special issue** on “Virtualization and Services for Cloud-Based Application Systems”
- PC Chair of **IEEE International Conference on Big Data Processing Systems 2018**
- PC Chair of **IEEE JointCloud Computing Workshop 2017**
- PC Chair of **IEEE Large Scale Clustering Algorithms Workshop 2016**
- PC member of **IEEE AAAI 2020**
- PC member of **IEEE ICDCS 2017, 2011**
- PC member of **IEEE ICWS 2020, 2019, 2018, 2017**
- PC member of **P2P 2010**

Selected Projects (<http://nicexlab.com/>)

- Principal Investigator: **NSFC** (Grant No. 61772541, 61379055, 60903205)
- Executive Principal Investigator: **NSFC** (Grant No. 61872376, 61379053)
- Main Contributor: **National Basic Research Program of China** (973, Grant No. 2016YFB1000100, 2014CB340303, 2011CB302601, 2005CB321801), **National High Technology Research and Development Program of China** (863, Grant No. 2011AA01A202)
- Cloud Operating Systems: KylinX
- Cloud Performance Debugging: PCatch
- Cloud Storage: PariX, Ursa, UrsaL, MapX, MapperX
- Cloud Reliable Memory: MemCube, CubeX
- Cloud Networking: OptaX, DLG
- Cloud VM Fast Booting: VirtMan, C4

- Cloud Graph Computing: TopoX, MST, TianheGraph
- Cloud machine learning: FreeLauncher, ParaX