

## Ding Yuan (Curriculum Vitae)

---

Dept. of Electrical and Computer Engineering    Email: [yuan@ece.toronto.edu](mailto:yuan@ece.toronto.edu)  
University of Toronto    URL: <http://www.eecg.toronto.edu/~yuan/>  
10 King's College Road, 2002E    Phone: (416) 978-5033  
Toronto, ON M5S 3G4 Canada    Fax: (416) 978-5033

### Education

Ph.D. Computer Science    Aug. 2006 - Dec. 2012  
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN, IL

B.E. Computer Science and B.S. in Applied Mathematics    Sept 2002 - Jul 2006  
BEIHANG UNIVERSITY, Beijing, China

### Professional Experience

Professor    2023 - Present  
Associate professor    2018 - 2023  
Assistant professor    2013 - 2018  
University of Toronto, Dept. of Electrical and Computer Engineering

Founder    2021 - Present  
YScope Inc.

### Select Honors and Recognitions

- ◇ Canada Research Chair in Systems Software, 2019 - present
- ◇ ACM SIGOPS Dennis M. Ritchie Thesis Award (my PhD student Yongle Zhang's thesis), 2022
- ◇ Meta (Facebook) Systems Research Award, 2022
- ◇ Gordon Slemmon Award on Teaching Design, 2021
- ◇ Connaught Innovation Award, 2018 and 2021
- ◇ Faculty of Engineering EngSoc Student Choice Award: Best Upper Year Instructors, 2021
- ◇ McCharles Prize for Early Research Distinction, 2018. Highest distinction for pre-tenure faculties in the Faculty of Applied Science and Engineering at University of Toronto, awarded to only one faculty in every two years.)
- ◇ NetApp Faculty Fellowship, 2013 and 2016
- ◇ University of Toronto Connaught New Researcher Award, 2013
- ◇ ACM SIGSOFT Distinguished Paper Award, 2011
- ◇ Saburo Muroga Fellowship, 2007
- ◇ University Golden Medal Prize in Beihang Univ., 2005

### Impact

- ◇ CLP (Compressed Log Processor) [C7], a log compression and search technology, is open-sourced and deployed at a number of Fortune 500 companies. For example, Uber published an engineering blog describing the impact made by CLP:  
<https://www.uber.com/en-US/blog/reducing-logging-cost-by-two-orders-of-magnitude-using-clp>.
- ◇ [C16] and its open-source bug detection tool, Aspirator, have made the following impact:
  - Aspirator is implemented by Google's open-source Error Prone static checker, which is used by hundreds of companies (including Google) and projects, including Hadoop and HBase.

- HBase developers started (and completed) a large project to address the issues discovered in the paper (see HBase-12187: <https://issues.apache.org/jira/browse/HBASE-12187>). It consists of 9 tasks, including using Aspirator as a check-in policy, thorough code review, use other static checkers including Coverity. It triggered a series of changes to their engineering practices. HBase is a widely-used open-source big data analytics system.
- Aspirator has discovered thousands of software defects that could have led to serious service-level outage or data-loss. Developers fixed these bugs, and by our estimation it effectively improve the reliability of these systems by 33%[\[C16\]](#).
- ◇ Hubble[\[C3\]](#), a tracing and performance debugging tool, is shipped on all Android devices manufactured by Huawei.
- ◇ Log20 [\[C11\]](#) is licensed by Netflix.
- ◇ lprof [\[C17\]](#), a performance monitoring system, is licensed by Huawei Technologies. lprof is now used by Huawei developers to monitor a production order processing system that is used by 80 million users and processes up to billions of transactions a day.
- ◇ CODE [\[C22\]](#) is patented and used by Microsoft to detect configuration errors.
- ◇ Contributed over 500 fixes to software defects (i.e. patches) to widely used open-source software projects including Linux Kernel, Hadoop, HBase, Cassandra, Postgres, Apache httpd, etc.

### Media Coverage and Dissemination in Industry Conferences

- ◇ CLP and its deployment results at Uber has been widely reported by media, including InfoQ, Hacker News, and tens of thousands of social media mentions.
  1. Uber Engineering Blog: Reducing Logging Cost by Two Orders of Magnitude using CLP.
  2. InfoQ article: Uber Reduces Logging Costs by 169x Using Compressed Log Processor (CLP). <https://www.infoq.com/news/2022/11/uber-compressed-log-processor/>
  3. Hacker News: Reducing logging cost by two orders of magnitude using CLP. <https://news.ycombinator.com/item?id=33032996>
- ◇ Four papers featured on The Morning Paper. The Morning Paper is an influential blog, with over 20K subscribers, providing in-depth coverage on research papers in computer science and engineering.
  1. The inflection point hypothesis: a principled approach to finding the root cause of a failure. <https://blog.acolyer.org/2019/11/08/the-inflection-point-hypothesis/>
  2. An analysis of performance evolution of Linux’s core operations. <https://shorturl.at/jprL8>
  3. Log20: Fully automated optimal placement of log printing statements under specified overhead threshold. <https://shorturl.at/BIS49>
  4. Simple testing can prevent most critical failures. <https://shorturl.at/afksB>
- ◇ November, 2017, SRE weekly issue 96, a blog on Scalability, Availability, Incident response of computer systems, covered *Log20* [\[C11\]](#).
- ◇ November, 2016, The Next Platform published an in-depth report of the JVM warm-up overhead paper [\[C15\]](#): “JVM Boost Shows Warm Java is Better Than Cold”. <https://www.nextplatform.com/2016/11/10/jvm-boost-shows-warm-java-better-cold/>
- ◇ October, 2016, The Morning Paper considered [\[C16\]](#) as a highlight of the year among all of the papers it has covered.
- ◇ January, 2016, A textbook “Fifty Quick Ideas to Improve Your Tests”, by Gojko Adzic, David Evans, and Tom Roden, ISBN 978-0993088100, Neuri Consulting 2014, reports on the findings from [\[C16\]](#).
- ◇ August, 2015, Dan Luu’s blog wrote an in-depth coverage on [\[C16\]](#). This is a technical blog on computer technologies with over 1 million hits per month.

- ◇ March, 2015, Murat Demirbas’s blog “Metadata” reports on [C16]. This blog has over 1.5 million page views; Demirbas is a professor at University of Buffalo.
- ◇ October, 2014, Patrick Durusau’s blog “Another Word For It” reports on [C16].
- ◇ October, 2014, Lorin Hochstein’s blog “It Will Never Work In Theory” reports on [C16].
- ◇ [C16] is presented in the following software industry conferences:
  1. “The Verification of a Distributed System”, by Caitie McCaffrey, O’Reilly’s Velocity Conference, San Jose, 2017
  2. “The Error of Our Ways”, by Kevin Henney, GOTO conference, the enterprise software development conference, Berlin, 2016
  3. Cassandra Summit Keynote, by Jonathan Ellis, founder of Cassandra, co-founder of DataStax and CTO. Cassandra is the most popular wide-column storage system; DataStax is a database provider with over 500 customers in over 50 countries. 2015.
  4. “Simple Testing Can Prevent Most Critical Failures”, Keynote, by Caitie McCaffrey at InfoQ’s QCon New York, June 2014,
  5. “Outwards from the middle of the maze”, by Peter Alvaro, RICON, the World’s Largest Distributed Systems Conference, San Francisco, 2014

## Research Contributions

### Peer-Reviewed Conference Publications

- [C1]  $\mu$ Slope: High Compression and Fast Search on Semi-Structured Logs. Rui Wang, Devin Gibson, Kirk Rodrigues, Yu Luo, Yun Zhang, Kaibo Wang, Yupeng Fu, Ting Chen, Ding Yuan. *To appear in the Proceedings of the 18th USENIX Symposium on Operating Systems Design and Implementation (OSDI’24)*, July, 2024. Acceptance rate: 15.6% = 44/282
- [C2] Relational Debugging – Pinpointing Root Causes of Performance Problems. Xiang (Jenny) Ren, Sitao Wang, Zhuqi Jin, David Lion, Adrian Chiu, Tianyin Xu, Ding Yuan. *In the Proceedings of the 17th USENIX Symposium on Operating Systems Design and Implementation (OSDI’23)*, pages 65-80, July, 2023. Acceptance rate: 19.6% = 50/255  
**Invited to appear in ACM SIGOPS Blog.**
- [C3] Hubble: Performance Debugging with In-Production, Just-In-Time Method Tracing on Android. Yu Luo, Kirk Rodrigues, Lijin Jiang, Bing Xia, David Lion, Ding Yuan. *In the Proceedings of the 16th USENIX Symposium on Operating Systems Design and Implementation (OSDI’22)*, pages 787-803, July 11-13, 2022. Acceptance rate: 19.5% = 49/251
- [C4] Investigating Managed Language Runtime Performance: Why JavaScript and Python are 8x and 29x slower than C++, yet Java and Go can be Faster? David Lion, Adrian Chiu, Michael Stumm, Ding Yuan. *In the Proceedings of the 2022 USENIX Annual Technical Conference (ATC’22)*, pages 835-851, July 11-13, 2022. Acceptance rate: 16.2% = 64/393  
**Invited to appear in USENIX ;login:**
- [C5] ctFS: Replacing File Indexing with Hardware Memory Translation through Contiguous File Allocation for Persistent Memory. Ruibin Li, Xiang Ren, Xu Zhao, Siwei He, Michael Stumm, Ding Yuan. *In the Proceedings of the 20th USENIX Conference on File and Storage Technologies (FAST’22)*, pages 35-49, February 22-24, 2022. Acceptance rate: 21.5% = 28/130  
**Invited to appear in USENIX ;login: and ACM Transaction on Storage (TOS), Best paper honorable mention**
- [C6] Understanding and Detecting Software Upgrade Failures in Distributed Systems Yongle Zhang, Junwen Yang, Zhuqi Jin, Utsav Sethi, Kirk Rodrigues, Shan Lu, Ding Yuan. *In the Proceedings of The 28th ACM Symposium on Operating Systems Principles (SOSP’21)*, pages 116-131, October 25-28, 2021. Acceptance rate: 15.5% = 54/348

- [C7] CLP: Efficient and Scalable Search on Compressed Text Logs Kirk Rodrigues, Yu Luo, Ding Yuan. *In the 15th USENIX Symposium on Operating Systems Design and Implementation (OSDI'21)*, pages 183-198, July, 2021. Acceptance rate: 18.8% = 31/165  
**Invited to appear in USENIX ;login:**
- [C8] M3: End-to-End Memory Management in Elastic Systems Software Stack. David Lion, Adiran Chiu, Ding Yuan. *In the 16th ACM European Conference on Computer Systems (EuroSys 2021)*, pages 507-522, April 2021. Acceptance rate: 14.1% = 27/191
- [C9] The Inflection Point Hypothesis: A Principled Debugging Approach for Locating the Root Cause of a Failure. Yongle Zhang, Kirk Rodrigues, Yu Luo, Michael Stumm, Ding Yuan. *In the 27th ACM Symposium on Operating Systems Principles (SOSP'19)*, pages 131-146, October 2019. Acceptance rate: 13.8% = 38/276  
**Invited to appear in USENIX ;login:**
- [C10] An Analysis of Performance Evolution of Linux's Core Operations. Xiang (Jenny) Ren, Kirk Rodrigues, Luyuan Chen, Camilo Vega, Michael Stumm, Ding Yuan. *In the 27th ACM Symposium on Operating Systems Principles (SOSP'19)*, pages 554-569, October 2019. Acceptance rate: 13.8% = 38/276
- [C11] Log20: Fully Automated Optimal Placement of Log Printing Statements under Specified Overhead Threshold. Xu Zhao, Kirk Rodrigues, Yu Luo, Michael Stumm, Ding Yuan, and Yuanyuan Zhou *In the 26th ACM Symposium on Operating Systems Principles (SOSP'17)*, pages 565-581, October 2017. Acceptance rate: 16.8% = 39/232
- [C12] Pensieve: Pensieve: Non-Intrusive Failure Reproduction for Distributed Systems using the Event Chaining Approach. Yongle Zhang, Serguei Makarov, Xiang Ren, David Lion, and Ding Yuan. *In the 26th ACM Symposium on Operating Systems Principles (SOSP'17)*, pages 19-33, October, 2017. Acceptance rate: 16.8% = 39/232
- [C13] The Game of Twenty Questions: Do You Know Where to Log? Xu Zhao, Kirk Rodrigues, Yu Luo, Michael Stumm, Ding Yuan, Yuanyuan Zhou. *In the Proceedings of the 16th Workshop on Hot Topics in Operating Systems (HotOS)*, pages 125 - 131, May 2017.
- [C14] Non-intrusive Performance Profiling of Entire Software Stacks based on the Flow Reconstruction Principle. Xu Zhao, Kirk Rodrigues, Yu Luo, Ding Yuan, and Michael Stumm. *In the proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI'16)*, pages 603 - 618, November, 2016. Acceptance rate: 17.6% = 47/267.
- [C15] Don't Get Caught In the Cold, Warm-up Your JVM: Understand and Eliminate JVM Warm-up Overhead In Data-parallel Systems. David Lion, Adrian Chiu, Xin Zhuang, Hailong Sun, Nikola Grcevski, and Ding Yuan. *In the proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI'16)*, pages 383 - 400, November 2016. Acceptance rate: 17.6% = 47/267. **Invited to appear in USENIX ;login:**
- [C16] Simple Testing Can Prevent Most Critical Failures: An Analysis of Production Failures in Distributed Data-intensive Systems. Ding Yuan, Yu Luo, Xin Zhuang, Guilherme Rodrigues, Xu Zhao, Yongle Zhang, Pranay U. Jain, and Michael Stumm. *In the Proceedings of the 11th USENIX Symposium on Operating Systems Design and Implementation (OSDI '14)*, pages 249 - 265, October 2014. Acceptance rate: 18% = 42/228.  
**Invited to appear in USENIX ;login:**
- [C17] lprof: A Non-intrusive Request Flow Profiler for Distributed Systems. Xu Zhao, Yongle Zhang, David Lion, Muhammad Faizan, Yu Luo, Ding Yuan, and Michael Stumm. *In the Proceedings of the 11th USENIX Symposium on Operating Systems Design and Implementation (OSDI '14)*, pages 629 - 644, October 2014. Acceptance rate: 18% = 42/228.
- [C18] Do Not Blame Users for Misconfigurations. Tianyin Xu, Jiaqi Zhang, Peng Huang, Jing Zheng, Tianwei Sheng, Ding Yuan, Yuanyuan Zhou, and Shankar Pasupathy. *In the Proceedings of the 24th ACM Symposium on Operating Systems Principles (SOSP'13)*, pages 244-259, November 2013. Acceptance rate: 19% = 30/160.

- [C19] Be conservative: Enhancing failure diagnosis with proactive logging. Ding Yuan, Soyeon Park, Peng Huang, Yang Liu, Michael M. Lee, Xiaoming Tang, Yuanyuan Zhou, and Stefan Savage. *In the Proceedings of the 10th USENIX Symposium on Operating System Design and Implementation (OSDI'12)*, pages 293-306, October 2012. *Acceptance rate: 11.6% = 25/215.*
- [C20] Characterizing Logging Practices in Open-Source Software. Ding Yuan, Soyeon Park and Yuanyuan Zhou. *In the Proceedings of the 34th International Conference on Software Engineering (ICSE'12)*, pages 102-112, June 2012 *Acceptance rate: 21% = 87/408.*
- [C21] How Do Fixes Become Bugs? – A Comprehensive Characteristic Study on Incorrect Fixes in Commercial and Open Source Operating Systems. Zuoning Yin, Ding Yuan, Yuanyuan Zhou, Shankar Pasupathy and Lakshmi Bairavasundaram. *In the proceedings of the 19th ACM SIGSOFT Symposium on the Foundations of Software Engineering (FSE'11)*, pages 26-36, September 2011. *Acceptance rate: 16.7% = 34/203.*  
**ACM SIGSOFT Distinguished paper award.**
- [C22] Context-based Online Configuration-Error Detection. Ding Yuan, Yinglian Xie, Rina Panigrahy, Junfeng Yang, Chad Verbowski and Arunvijay Kumar. *In the proceedings of the 2011 USENIX Annual Technical Conference (ATC'11)*, pages 313-326, June 2011. *Acceptance rate: 16.7% = 34/203.*
- [C23] Improving Software Diagnosability via Log Enhancement. Ding Yuan, Jing Zheng, Soyeon Park, Yuanyuan Zhou and Stefan Savage. pages 3-14, March 2011. Newport Beach, CA *In the Proceedings of the 16th International Conference on Architecture Support for Programming Language and Operating Systems (ASPLOS'11)*, pages 3-14, March 2011. *Acceptance rate: 21.2% = 32/151.*  
**Nominated for Best Paper Award and fast forwarded to ACM Transaction on Computer Systems (TOCS).**
- [C24] SherLog: Error Diagnosis by Connecting Clues from Run-time Logs. Ding Yuan, Haohui Mai, Weiwei Xiong, Lin Tan, Yuanyuan Zhou and Shankar Pasupathy. *In the Proceedings of the 15th International Conference on Architecture Support for Programming Language and Operating Systems (ASPLOS'10)*, pages 143-154, March 2010. *Acceptance rate: 17.7% = 32/181.*  
**Nominated for Best Paper Award.**
- [C25] CISpan: Comprehensive Incremental Mining Algorithms of Closed Sequential Patterns for Multi-Versional Software Mining. Ding Yuan, Kyuhyung Lee, Hong Cheng, Gopal Krishna, Zhenmin Li, Xiao Ma, Yuanyuan Zhou and Jiawei Han. *In the Proceedings of SIAM: SIAM International Conference on Data Mining (SDM08)*, pages 84-95, April 2008. *Acceptance rate: 14.2% = 40/282.*
- [C26] /\* iComment: Bugs or Bad Comments? \*// Lin Tan, Ding Yuan, Gopal Krishna and Yuanyuan Zhou. *21st ACM Symposium on Operating Systems Principles (SOSP'07)*, pages 145-158, October 2007. *Acceptance rate: 14.2% = 40/282.*
- [C27] HotComments: How to Make Program Comments More Useful? Lin Tan, Ding Yuan and Yuanyuan Zhou. *In the Proceedings of the 11th Workshop on Hot Topics in Operating Systems (HotOS)*, pages 19:1-19:6, May 2007.

### Peer-reviewed Journal Publications

- [J1] ctFS: Replacing File Indexing with Hardware Memory Translation through Contiguous File Allocation for Persistent Memory. Ruibin Li, Xiang Ren, Xu Zhao, Siwei He, Michael Stumm, Ding Yuan. *ACM Transaction on Storage (TOS)*, Volume 18, Issue 4, November 2022, Article No.: 30, pp 1–24
- [J2] Improving Software Diagnosability via Log Enhancement. Ding Yuan, Jing Zheng, Soyeon Park, Yuanyuan Zhou and Stefan Savage. *In ACM Transactions on Computer Systems (TOCS)*, 30(1), pages 1-28, February 2012. (Fast-forwarded from ASPLOS'11.)

## Invited Publications (not peer-reviewed)

- [I1] **Perspect** — Exploiting essential characteristics of performance issues for automatic performance diagnosis. **Xiang (Jenny) Ren, Sitao Wang, Zhuqi Jin, David Lion, Adrian Chiu**, Tianyin Xu, and Ding Yuan. *Published on ACM SIGOPS Blog*, December 2023.
- [I2] **Kairux**: Distributed System Fault Localization based on The Inflection Point Hypothesis Yongle Zhang, Kirk Rodrigues, Yu Luo, Michael Stumm, and Ding Yuan *In USENIX ;login.*, Jan 2023.
- [I3] Investigating Managed Language Runtime Performance: Why JavaScript and Python are 8x and 29x slower than C++, yet Java and Go can be Faster? David Lion, Adrian Chiu, Michael Stumm, Ding Yuan. *In USENIX ;login.*, June 2022.
- [I4] **ctFS**: Converting File Index Traversals to Hardware Memory Translation through Contiguous File Allocation for Persistent Memory Ruibin Li, Xiang Ren, Xu Zhao, Siwei He, Michael Stumm, Ding Yuan. *In USENIX ;login.*, February 2022.
- [I5] Don't Get Caught In the Cold, Warm-up Your JVM: Understand and Eliminate JVM Warm-up Overhead In Data-parallel Systems. David Lion, Adrian Chiu, Xin Zhuang, Hailong Sun, Nikola Grcevski, and Ding Yuan. *In USENIX ;login: 42(1)*, Pages 46-51, March 2017.
- [I6] Simple Testing Can Prevent Most Critical Failures: An Analysis of Production Failures in Distributed Data-intensive Systems. Ding Yuan, Yu Luo, Xin Zhuang, Guilherme Rodrigues, Xu Zhao, Yongle Zhang, Pranay U. Jain, and Michael Stumm. *In USENIX ;login: 40(1)*, Pages 18-23, February 2015.

## Patents

- [P1] **Fingerprinting event logs for system management troubleshooting**. Rina Panigrahy, Chad Verbowski, Yinglian Xie, Junfeng Yang, Ding Yuan. *US patent 8,069,374*, allowed on 2011/11/29. Assigned to Microsoft Corporation.
- [P2] **Systems and processes for computer log analysis**. Muhammad Faizan, David Lion, Yu Luo, Michael Stumm, Ding Yuan, Xu Zhao, Yongle Zhang. *US Patent 9,720,671*. Granted on 2017/8/8.
- [P3] **Systems and processes for computer log analysis**. Muhammad Faizan, David Lion, Yu Luo, Michael Stumm, Ding Yuan, Xu Zhao, Yongle Zhang. *US Patent 10,484,506*. Granted on 2019/11/19. (This is a continuation patent on US Patent 9,720,671.)
- [P4] **Log Processing and Analysis**. Yu Luo, Kirk Rodrigues, Michael Stumm, Ding Yuan, Xu Zhao. *US Patent 1,0642,712 B2*. Granted on 2020-05-05
- [P5] **Compression, Searching, and Decompression of Log Messages**. Yu Luo, Kirk Rodrigues, Ding Yuan. *US patent application No. 17/097550*. Applied in 2020.
- [P6] **Compression, Searching, and Decompression of Log Messages**. Yu Luo, Kirk Rodrigues, Ding Yuan. *Chinese patent application No. 202011282843.7*.

## Advising

### PhD (11):

- Yongle Zhang (2013/9 - 2020/21). First employment: Tenure-track assistant professor, Computer Science Department, Purdue University. Thesis: “Automating Failure Diagnosis for Distributed Systems”. **ACM SIGOPS Dennis M. Ritchie Thesis Award**.
- Xu Zhao (2015/9 - 2020/8). **Facebook PhD Fellow, 2018-2020**. First employment: Research scientist, Facebook. Thesis: “Automating Log-based Software Debugging on Distributed Server Stacks”.
- Kirk Rodrigues (2017/9 - 2023/3). First employment: Co-Founder of YScope. Thesis: CLP: Efficient and Scalable Search on Compressed Text Logs.
- David Lion (2016/9 - 2023/12). First employment: Co-Founder of YScope. Thesis: Investigating and Improving the Performance of Managed Language Runtime Environments.

- Yu Luo (2017/9 - 2024/8 (expected)). First employment: Co-Founder of YScope.
- Xiang Ren (2017/9 - 2024/12 (expected)).
- Adrian Chiu (2021/1 - present).
- Ruibin Li (2021/9 - present).
- Rishi Devsot (2023/9 - present).
- Rui Wang (2023/9 - present).
- Serguei Makarov (2017/9 - 2018), dropped-out.

#### **MASc (18):**

- Xu Zhao (2013/9 - 2015/9). Thesis: “lprof: A Non-intrusive Request Flow Profiler for Distributed Systems”.
- David Lion (2014/9 - 2017/5). Thesis: “Don’t Get Caught In The Cold, Warm-up Your JVM: Understand and Eliminate JVM Warm-up Overhead in Data-parallel Systems”.
- Xiang Ren (2015/9 - 2018). Thesis: “Achieving Scalable and Reliable Non-Intrusive Failure Reproduction in Distributed Systems by Enhancing the Event Chaining Approach”.
- Serguei Makarov (2015/9 - 2017). Thesis: “An Event-based Language for Programmable Debugging”.
- Yu Luo (2017/9 - 2018/1). Thesis: “aprof: An Offline Non-Intrusive Call-tree Profiler for Distributed Systems”.
- Kirk Rodrigues (2017/9 - 2018/1). Thesis: “Non-Intrusive, Automated Log Discovery and Parsing”.
- Adrian Chiu (2018/9 - 2021/1). Thesis: “Profiling the Performance of Programming Language Runtimes and Towards Building a New Python Runtime”.
- Ruibin Li (2019/9 - 2022/1). Thesis: “ctFS: Eliminating File Indexing with Contiguous File System on Persistent Memory”.
- Zhuqi Jin (2020/9 - 2022/1). Thesis: “Understanding and Detecting Software Upgrade Failures in Distributed Systems”. Software engineer at Meta.
- Rui Wang (2021/9 - 2023/9). Thesis: “High Compression and Fast Search on Semi-Structured Logs.” Software engineer at YScope.
- Rishikesh Devsot (2021/9 - 2024/1). Thesis: “CLOUDS: Closed Loop On-site Diagnosis System”. Software engineer at YScope.
- Haiqi Xu (2021/9 - present).
- Devin Gibson (2022/9 - present).
- Sitao Wang (2022/9 - present).
- Xiaochong Wei (2022/9 - present).
- Yifan Yu (2022/9 - present).
- Zhihao Lin (2023/9 - present).
- Chaoyue Gong (2023/9 - present).

#### **Postdoc and Visiting scholar (3):**

- Hailong Sun (2015/8 - 2016/8), visiting scholar, now Professor at Beihang University.
- Huangshi Tian (2022/8 - present), Post-doctoral researcher. PhD from The Hong-Kong University of Science and Technology.
- Ming Zhang (2023/11 - present), visiting PhD student from Huazhong University of Science and Technology.

#### **Undergraduate students (38).**

#### **High school summer students (3).**

## Teaching Experience

ECE 344, “Operating Systems”: 2020-23, 2013-18  
ECE 244, “Programming Fundamentals”: 2023, 2016-17  
ECE 454, “Computer Systems Programming”, 2018, 2013-14  
ECE 1759, “Advanced Operating Systems”: 2020-23, 2014-17

## Service

### Chair and Vice-Chair Positions

2023-present, Vice Chair, ACM SIGOPS - Special Interest Group on Operating Systems  
PC Co-chair, 2023 ACM European Conference on Computer Systems (EuroSys) Poster  
PC Co-chair, 2019 The 17th ACM Workshop on Hot Topics in Operating Systems (HotOS)  
PC Co-chair, 2019 The 10th ACM SIGOPS Asia-Pacific Workshop on Systems (APSys)  
Chair, 2018 OSDI Best Paper Award Selection Committee  
Chair, 2017 Inaugural ACM SOSP Student Research Competition (SRC)  
Chair, 2016 ASPLOS Poster session and Lightning session  
Local Arrangement Co-Chair, 2019 ACM 27th Symposium on Operating Systems Principles (SOSP)

### Program Committee Member

2024 ACM SIGOPS 29th Symposium on Operating Systems Principles (SOSP)  
2024 USENIX Symposium on Operating Systems Design and Implementation (OSDI)  
2023 ACM SIGOPS 28th Symposium on Operating Systems Principles (SOSP)  
2023 USENIX Symposium on Operating Systems Design and Implementation (OSDI)  
2023 ACM European Conference on Computer Systems (EuroSys)  
2022 USENIX Symposium on Operating Systems Design and Implementation (OSDI)  
2021 ACM SIGOPS 28th Symposium on Operating Systems Principles (SOSP)  
2021 USENIX Symposium on Operating Systems Design and Implementation (OSDI)  
2021 International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)  
2021 1st Workshop on High Availability and Observability of Cloud Systems (HAOC)  
2020 USENIX Symposium on Operating Systems Design and Implementation (OSDI)  
2020 USENIX Symposium on Networked Systems Design and Implementation (NSDI)  
2018 USENIX Symposium on Operating Systems Design and Implementation (OSDI)  
2018 European Conference on Computer Systems (EuroSys)  
2018 International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), ERC  
2017 ACM SIGOPS 26th Symposium on Operating Systems Principles (SOSP)  
2016 International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)  
2016 International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), Poster and Lightning session chair  
2015 The 25th ACM Symposium on Operating Systems Principles (SOSP'15) (Poster PC)  
2015 USENIX Annual Technical Conference  
2015 USENIX LISA Conference



2014 USENIX/ACM OSDI (External Review Committee)

2014 USENIX Annual Technical Conference

2014 ACM SIGMETRICS

2014 USENIX ICAC

2014 PACT student competition

2014 IBM CASCON

2012 USENIX Workshop on Managing Systems Automatically and Dynamically (MAD)

### **Other Services**

ACM HotOS Workshop Steering Committee, 2019-present.

Organizer, Papers We Love - Toronto, 2021-present

### **University Services**

ECE Computer Systems User Committee: 2016-Present. (Chair since 2023)

ECE Faculty Search Committee: 2018-19, 2024.

UTEST (University of Toronto Early-Stage Technology) Delegate member to visit The Hebrew University of Jerusalem (HUJI), 2022