Shehbaz Jaffer

647-939-0662 | shehbaz@cs.toronto.edu | LinkedIn | Github | Web

Education

University of Toronto — Canada
PhD in Computer Science — *2015 – Aug 2021 (Expected)*

Indian Institute of Technology, Delhi — India
Master of Technology — *2011 – 2013*

University of Pune — India
Bachelors in Engineering — *2006 – 2010*

Experience

Graduate Researcher — Sept 2015 – Present
University of Toronto — *Toronto, Ontario, Canada*

Rethinking WOM Codes for Increasing QLC Flash lifetime

- Invented a Write-Once memory Coding Scheme that improves QLC Flash reliability by 500%.
- Designed and Implemented WOM Code on a Flash-Emulator and LightNVM Linux Kernel Module.
- Evaluated WOM coding scheme gains on real world traces from Microsoft, Alibaba and academic Institutions.

Evaluating File System Reliability on Solid State Drives

- Evaluated ext4 and BtrFS file systems capability to recover from underlying media errors
- Implemented a device mapper error injection framework to emulate real world SSD based I/O errors.
- Demonstrated 16% File system Operations in real world File Systems can cause data loss due to media errors.

SSD Based Workload Characteristics and Their Performance Implications

- Created a Block tracing framework and traced I/O patterns for Apache Spark, Cassandra and RocksDB.
- Demonstrated the impact of SSD architecture on workload Performance.
- Contributed YCSB and RocksDB traces to Storage Network Industry Association (SNIA)

Research Intern — May 2020 – Aug 2020
VMWare Research — *Palo Alto, USA (Remote Internship)*

- Built a QEMU-based computational storage simulator to offload storage based computational tasks to the simulated in-storage compute drive.
- Worked on QEMU, IPC Optimization, NVMe directives, zlib compression offload, ext4 file system checksum offload, MySQL Project/Scan/Filter Operation offload for OLAP Applications.

Researcher — Aug 2013 – July 2015
Advanced Technology Group, NetApp — *Bengaluru, India*

Kuiper- A Distributed, Reliable and Scalable Peta Byte Scale Object Store for Archival Data

- Designed and developed a distributed Peta Byte scale Key-Value Object Store that stores objects ranging from few MBs to GBs in size reliably over a long period of time.
- Implemented integrity checksum verification, end-to-end Delete work-flow and garbage collection.
- Coordinated work across 3 geographically distributed teams.

Quick Reboot Of Virtual Storage Appliance (VSA)

- Decreased the reboot time of VSA, a Virtual Machine that hosts Data ONTAP Operating System by 22%.
- Created a prototype caching upto 96% file system meta-data of Guest OS on the host RAMDisk.
- Investigated and removed multiple bottlenecks for further decreasing the reboot time of the Virtual Machine.

Graduate Research Assistant Jan 2012 – July 2013

Indian Institute of Technology Delhi *New Delhi, India*

- Implemented a remote desktopping Tool using Virtual machine Record and Replay Framework.
- Decreased Network Traffic by 9X by streaming VM state from server to the client instead of video and replaying the state at client.
- Improved host CPU utlization by upto 40% during passive video and interactive game streaming.

Software Engineer June 2010 – May 2011

Yahoo (Persistent Systems) *Pune, Maharashtra, India*

- Ensured display of text-based ads on Yahoo Partner websites.
- Created reporting and monitoring tools for processing ad-requests received by 80+ ad server nodes in a cluster.

Publications

**Reverse Chronological Order, First author is the primary contributor in each Publication.*

9. SSD Based Workload Characteristics and Their Performance Implications
   Gala Yadgar, Moshe Gabel, **Shehbaz Jaffer**, Bianca Schroeder
   ACM Transactions on Storage '21

8. Rethinking WOM codes to enhance the lifetime in new SSD Generations
   Shehbaz Jaffer, Kaveh Mahdaviani, Bianca Schroeder
   Hot Topics in Storage Systems'20

7. The Reliability of Modern File Systems in the face of SSD Errors
   Shehbaz Jaffer, Stathis Maneas, Andy Hwang, Bianca Schroeder
   ACM Transactions on Storage TOS'20

6. Reclaiming Good Transactions from a Corrupt Journal
   Shehbaz Jaffer
   ACM Student Research Competition@SOSP'19

5. Evaluating File System Reliability on Solid State Drives
   Shehbaz Jaffer, Stathis Maneas, Andy Hwang, Bianca Schroeder
   USENIX Annual Technical Conference'19

4. Dynamic Redundancy based on Media Reliability
   Shehbaz Jaffer, Ashvin Goel, Angela Demke Brown, Bianca Schroeder
   Work In Progress @FileSystem Access and Storage Technology'16

3. Improving Remote Desktopping through Adaptive Record Replay
   Shehbaz Jaffer, Piyus Kedia, Sorav Bansal
   International Conference on Virtual Execution Environments VEE' 15

2. Providing High Availability in Cloud Storage by Decreasing VM Reboot Time
   Shehbaz Jaffer, Mangesh Chitnis, Ameya Usgaonkar
   Hot Topics in Dependability Systems HotDep'2014

1. Offloading File Search Computation for performance improvement in smart phones
   Ashutosh Jain, Vigya Sharma, **Shehbaz Jaffer**, Kolin Paul
   IEEE Connect 2013

Patents

2. **Virtual Machine Reboot Information Persistence Into Host Memory**
   Ameya Usgaonkar, Mangesh Chitnis, **Shehbaz Jaffer**

1. **Risk based rebuild of data objects in an erasure coded storage system**
   David Anthony Slik, **Shehbaz Jaffer**, Sethuraman Subbiah, Keith Arnold Smith, Giridhar Appaji Nag Yasa, Atish Kathpal

Teaching Experience

- TA, ECE 244: Introduction to Data Structures: Fall 2019, Fall 2018
- TA, CSC 209: Software Tools and Systems Programming: Summer 2019
- TA, CSC 22333: Special Topics in Storage Systems: Winter 2019, Winter 2018
- TA, APS105: Introduction to Computer Fundamentals: Winter 2019,Winter 2018
- TA, CSC369: Operating Systems: Winter 2019, Fall 2018, Summer 2018, Winter 2016
- TA, ECE343: Data Structures and Algorithms: Fall 2017
- TA, ECE 344: Operating Systems: Winter 2017, Fall 2016, Fall 2015
- TA, CSC467: Compilers and Interpreters: Fall 2016
- TA, ECE568: System Security: Winter 2016
- TA, ECE243: Computer Organization: Winter 2016

Honours and Awards

- ACM Student Travel Grant for attending SOSP 2019
- DCS PhD Travel Grant for attending SOSP 2019
- Ranked 180/137000 in GATE 2011 (Graduate Aptitude Test in Engineering)
- Scholarship for Graduate Studies awarded by MHRD, Government of India
- Distinguished Alumni, MES College of Engineering, University of Pune.

Presentations

- **HotStorage 2020**: Rethinking WOM Codes to Enhance the lifetime in new SSD Generations (Online Conference)
- **SOSP 2019**: SRC Poster, Reclaiming good transactions from a corrupt journal, Huntsville, ON, Canada
- **ATC 2019**: Poster, Evaluating File System Reliability on Solid State Drives, Washington, USA
- **FAST 2016**: WIP, Dynamic Redundancy Based on Media Reliability, Santa Clara, USA
- **VEE 2015**: Remote Desktopping using adaptive Record-Replay, Istanbul, Turkey
- **HotDep 2014**: Improving VM Reboot for High Availability, Broomfield, Colorado, USA

Mentorship

- **Improving QLC Flash Reliability** *Aug 2020-Present*

  Charles Xu (BASc)
- **SSD Based Workload Characteristics** *May 2018-Jan 2019*

  Brain Chan (MScAc), Tony Xiao (BASc), Jialun Lyu (BASc), Yuhan Shao (BASc), Fengjia Zhang (BASc)

Service

- Undergraduate Student Mentor, Systems Research, Research-A-Thon, UofT, 2020
- Committee Member, DCS Graduate Admissions Triage Committee, UofT, 2020
- PC Member, Artifact Evaluation Committee at ASPLOS 2020
- PC Member, Artifact Evaluation Committee at SOSP 2019
- Proceedings Scribe, Session-8, SOSP 2019

Technical Skills

- Languages: C++, C, Shell Scripting, Python
- Vritualization Tools: QEMU, KVM, Xen, VMWare Server, VMWare ESX, VMWare Player
- System Development and Analysis Tools: sysprof, valgrind, gprof
- Compilers: LLVM/Clang, g++
- Network Analysis: Wireshark, Tshark
- Build Tools: Jenkins, TravisCI
- Testing Tools: GTest, gCov, GMock
- Documentation: Latex, Markdown, gnuplot

Leadership, Diversity and Inclusion

- Co-Founder and CTO, SmartGraid, a ML Based Assignment Grading Startup — May 2017 - April 2018
- Co-Founder and President, Indian Graduate Students Association, University of Toronto — Sept 2016 - May 2017
- Hosted vi-RADIO, a weekly virtual research intern talk attended by 30+ Researchers at VMWare Research with 10+ PhD Research Intern speakers. — Jun 2020 - Aug 2020
- Web-Administrator, Jinbukai Karate Club, University of Toronto — May 2017 - Apr 2018
- Tournament Director, Jinbukai Karate Club, University of Toronto — Sept 2016 - May 2017
- Social Director, ECE Graduate Students Society, University of Toronto — Sept 2015 - Aug 2016
- ECE GSU Representative, University of Toronto — Sept 2015 - Aug 2016