

Bo Fang

Computer Scientist | Pacific Northwest National Laboratory | bo.fang@pnnl.gov

Education

- Ph.D., Electrical and Computer Engineering, University of British Columbia, Canada, Feb 2020. Advisors: Professor Karthik Pat-tabiraman and Professor Matei Ripeanu. Thesis: Approaches for Building Error Resilient Applications.
- MAsc, Electrical and Computer Engineering, University of British Columbia, Canada, May 2014. Advisors: Professor Karthik Pat-tabiraman and Professor Matei Ripeanu. Thesis: Error Resilience Evaluation of GPGPU Applications.
- Master in Software Systems, University of British Columbia, Canada, 2011.
- BE, Information Security, Wuhan University, China, 2006.

Research interest

Distributed Systems, High Performance Computing, Dependable Systems and Fault Tolerance Techniques, GPGPU computing, Quantum Computing.

Research Experience

Computer Scientist

Pacific Northwest National Laboratory

Jan. 2022 to present

Manager: Dr. Kevin Barker

Ongoing projects: explore optimization techniques for quantum circuit transpilation; fault tolerance characterization and improvement for mixed-precision floating point formats, advanced techniques for I/O and network efficiency of large-scale neural network training

Research associate

Pacific Northwest National Laboratory

Apr. 2020 to Dec. 2021

Mentors: Dr. Sriram Krishnamoorthy, Dr. Ang Li

- *Quantum circuit simulation*: designed and implemented a graph-based partitioning approach to achieve efficient quantum circuit simulation on large-scale systems. Designed and developed noise model simulation for the density matrix simulator and evaluated the quantum algorithms' behaviors against various noise models. Key participant in building a quantum circuit simulator that exploits the Partitioned Global Address Space (PGAS) based SHMEM communication model and demonstrated its performance on the Summit supercomputer.
- *Fault tolerance for storage-related errors*: designed and implemented the FUSE-based fault injector to systematically introduce storage-related faults into the application level, and evaluated the error resilience of complex HPC proxy applications.

Research intern

Pacific Northwest National Laboratory

Jun. 2018 to Nov. 2018

Mentor: Dr. Sriram Krishnamoorthy

Explored the spatial data smoothness of scientific application and designed and implemented the BonVoision system that efficiently repairs the errors that cause memory DUEs with the neighboring data at run-time. Proposed a future research direction to employ machine learning techniques to predict the likelihood of a repair to lead to successful execution for an application.

Research intern

Los Alamos National Laboratory

Jan. 2016 to Apr. 2016

Mentor: Dr. Nathan Debardeleben

Investigated how an error propagated through different system layers and proposed that silent data corruptions should be discussed at each layer of the system and are confined within the goal of the approach. Initiated and explored the idea of a roll-forward recovery scheme for checkpointing/restart systems.

Research assistant

University of British Columbia

Sep. 2011 to Jan. 2020

- *Error resilient applications*: proposed comprehensive approaches to improve the effectiveness and efficiency of fault tolerance techniques for HPC applications: ePVF, a methodology that distinguishes crash-causing bits from the architecturally correct execution (ACE) bits and obtains a closer estimate of the SDC rate than PVF analysis; The roll-forward recovery and the roll-back recovery schemes and demonstrates the improvement in the overall efficiency of the C/R with two systems: LetGo (for faults affecting computational components) and BonVoision (for faults affecting DRAM memory).
- *GPU error resilience*: proposed the first instruction-level fault injection methodology and tool, GPU-Qin, to estimate and understand the error resilience of GPGPU applications. Conducted the evaluation and analysis on the error resilience characteristics of various GPU applications.

Major Awards

- Honorable Mention for the 2020 ACM SIGHPC Dissertation Award, Jul. 2020. Awarded by ACM's Special Interest Group on High-performance computing for "making significant progress on protecting large-scale HPC applications against soft errors". This award is given to the best doctoral dissertation completed in high-performance computing in the previous year. One honorable mention is given each year worldwide.
- William C. Carter Ph.D. Dissertation Award in Dependability, May 2020. Awarded jointly by IEEE TC on Dependable Computing and Fault Tolerance (TCFT) and IFIP Working Group 10.4 on Dependable Computing and Fault Tolerance for "addressing the problem of transient hardware faults in high performance computing systems". The William C. Carter Ph.D. Dissertation Award is presented since 1997 "to recognize an individual who has made a significant contribution to the field of dependable and secure computing throughout his or her Ph.D. dissertation". One award is given each year worldwide.

- Canada NSERC Postdoctoral Fellowship, Jan. 2020 - Jan. 2022. \$90,000 for two years, awarded by Natural Sciences and Engineering Research Council of Canada. I am ranked 2nd in the computer science division across Canada.

Teaching Experience

- 2022: mentored Meng Wang from UBC for his PhD internship at PNNL.
- 2015: mentored a Mitacs student from India for his entire internship period.
- 2014 - 2018: teaching assistant for undergraduate and graduate courses: distributed systems, database, Java programming, fault tolerance system design

Other Experience

- Nov. 2011 - May. 2012: Mitacs-Accelerate intern at Singular Software Inc.
- Sep. 2010 - Dec. 2010: Software Engineering Intern at Alcatel-Lucent S.A.

Publications

Journal

- [1] *Fault Injection for TensorFlow Applications*, Niranjhana Narayanan, Zitao Chen, **Bo Fang**, Guanpeng Li, Karthik Pattabiraman, Nathan DeBardeleben, IEEE Transactions on Dependable and Secure Computing (TDSC). accepted May 2022.
- [2] *Improving the Accuracy of IR-level Fault Injection*, Lucas Palazzi, Guanpeng Li, **Bo Fang**, and Karthik Pattabiraman, IEEE Transactions on Dependable and Secure Computing (TDSC). doi: 10.1109/TDSC.2020.2980273.
- [3] *A Systematic Methodology for Evaluating the Error Resilience of GPGPU Applications*, **Bo Fang**, Karthik Pattabiraman, Matei Ripeanu, and Sudhanva Gurumurthi, in the IEEE Transactions on Parallel and Distributed Systems (TPDS). vol. 27, no. 12, pp. 3397-3411, 1 Dec. 2016, doi: 10.1109/TPDS.2016.2517633.

Conference

- [4] *MARS: Malleable Actor-Critic Reinforcement Learning Scheduler*, Betis Baheri, Qiang Guan, Jacob Tronge, **Bo Fang**, Ang Li, Vipin Chaudhary 2022 IEEE International Performance, Computing, and Communications Conference (IPCCC).
- [5] *Efficient Hierarchical State Vector Simulation of Quantum Circuits via Acyclic Graph Partitioning*, **Bo Fang***, M. Yusuf Özkaya*, Ang Li, Ümit V. Çatalyürek, Sriram Krishnamoorthy, IEEE Cluster 2022 (* contributed equally) **Best paper award**.
- [6] *Pinpointing the System Reliability Degradation in NISQ Machines*, Qiang Guan, Betis Baheri, Zixuan Xu, Ying Mao, Vipin Chaudhary, Shuai Xu and **Bo Fang**, 2022 IEEE International Conference on Quantum Computing and Engineering (QCE22).
- [7] *ASAP - Automatic Synthesis of Area-Efficient and Precision-Aware CGRA*, Cheng Tan, Thierry Tambe, Jeff Zhang, **Bo Fang**, Tong Geng, Gu-Yeon Wei, David Brooks, Antonino Tumeo, Ganesh Gopalakrishnan, Ang Li, International Conference on Supercomputing. Jun 27-30, 2022 (Accepted).
- [8] *A Hybrid System for Learning Classical Data in Quantum States*, Samuel A. Stein, Betis Baheri, Ray Marie Tischio, Yiwen Chen, Ying Mao, Qiang Guan, Ang Li, **Bo Fang**, In IEEE 34th International Performance Computing and Communications Conference (IPCCC), 2021 (acceptance rate: 29%).
- [9] *SV-Sim: Scalable PGAS-based State Vector Simulation of Quantum Circuits.*, Ang Li, **Bo Fang**, Christopher Granade, Guen Prawiroatmodjo, Bettina Heim, Martin Roetteler and Sriram Krishnamoorthy, Shuai Xu, The 2021 International Conference for High Performance Computing, Networking, Storage and Analysis (SC), 2021 (acceptance rate: TBA).
- [10] *QuGAN: A generative adversarial network through quantum states*, Samuel A Stein, Betis Baheri, Ray Marie Tischio, Ying Mao, Qiang Guan, Ang Li, **Bo Fang**, Shuai Xu, IEEE International Conference on Quantum Computing and Engineering (QCE), 2021.
- [11] *Characterizing Impacts of Storage Faults on HPC Applications: A Methodology and Insights*, **Bo Fang***, Daoce Wang*, Sian Jin, Quincey Koziol, Zhao Zhang, Qiang Guan, Suren Byna, Sriram Krishnamoorthy, Dingwen Tao, IEEE Cluster, 2021 (acceptance rate: 29%) (* contributed equally).
- [12] *TensorFlowFI: A Flexible Fault Injection Framework for TensorFlow Applications*, Zitao Chen, Niranjhana Narayanan, **Bo Fang**, Guanpeng Li, Karthik Pattabiraman, Nathan DeBardeleben, IEEE International Symposium on Software Reliability Engineering (ISSRE), 2020 (acceptance rate: 25.6%).
- [13] *Chaser: A Enhanced Fault Injection Tool for Tracing Soft Errors in MPI Applications*, Qiang Guan, Xunchao Hu, Terence Grove, **Bo Fang**, Hailong Jiang, Heng Yin, Nathan DeBardeleben, 50th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN) 2020, (acceptance rate: 16.5%).
- [14] *A Tale of Two Injectors: End-to-End Comparison of IR-level and Assembly-Level Fault Injection*, Lucas Palazzi, Guanpeng Li, **Bo Fang**, and Karthik Pattabiraman, IEEE International Symposium on Software Reliability Engineering (ISSRE), 2019 (acceptance rate: 31.4%).
- [15] *BonVoision: Leveraging Spatial Data Smoothness for Recovery from Memory Soft Errors*, **Bo Fang**, Karthik Pattabiraman, Matei Ripeanu, Sriram Krishnamoorthy, ACM International Conference on Supercomputing (ICS), 2019 (acceptance rate: 23.2%).
- [16] *LetGo: A Lightweight Continuous Framework for HPC Applications upon Failures*, **Bo Fang**, Qiang Guan, Nathan Debardeleben, Karthik Pattabiraman, Matei Ripeanu, The ACM International Symposium on High-Performance Parallel and Distributed Computing (HPDC), 2017 (acceptance rate 18%).
- [17] *ePVF: An Enhanced Program Vulnerability Factor Methodology for Cross-Layer Resilience Analysis*, **Bo Fang**, Qining Lu, Karthik Pattabiraman, Matei Ripeanu and Sudhanva Gurumurthi, Proceedings of the IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), 2016, (acceptance rate 21%).
- [18] *GPU-Qin: A Methodology for Evaluating the Error Resilience of GPGPU Applications*, **Bo Fang**, Karthik Pattabiraman, Matei Ripeanu and Sudhanva Gurumurthi, Proceedings of the IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS), 2014, (acceptance rate 31%).
- [19] *GPUS: Combining high-performance with high-reliability*, L. Bautista Gomez, F. Cappello, L. Carro, N. DeBardeleben, **B. Fang**, S.

Gurumurthi, K. Pattabiraman, P. Rech, M. Sonza Reorda, Embedded tutorial paper (invited), Proceedings of the International Symposium on Design Automation and Test in Europe (DATE), 2014 (invited paper).

Workshop, poster and abstract

- [20] *Towards Precision-Aware Fault -olerance Approaches for Mixed-Precision Applications*, **Bo Fang**, Hari Siva, Timonhy Tsai, Xinyi Li, Ganesh Gopalakrishnan, Ignacio Laguna, Kevin Barker, Ang Li, FTXS 2022: Workshop on Fault Tolerance for HPC at eXtreme Scale (co-located at SC22).
- [21] *Visual Analysis on The Resilience of HPC Applications Using Control-Flow Graph*, Hailong Jiang, Shaolun Ruan, **Bo Fang**, Kevin Barker, Ang Li, Yong Wang, Qiang Guan, International Conference for High Performance Computing, Networking, Storage, and Analysis (SC22)
- [22] *Towards Predicting the Impact of Roll-Forward Failure Recovery for HPC Applications*, **Bo Fang**, Jieyang Chen, Karthik Pattabiraman, Matei Ripeanu, Sriram Krishnamoorthy, the 49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), 2019 Fast abstract 2 pages.
- [23] *SDC is in the Eye of the Beholder: A Survey and Preliminary Study*, **Bo Fang**, Panruo Wu, Qiang Guan, Nathan Debardeleben, Laura Monroe, Sean Blanchard, Zhizong Chen, Karthik Pattabiraman, Matei Ripeanu, 3rd IEEE International Workshop on Reliability and Security Data Analysis (RSDA) co-located with DSN 2016, 6 pages.
- [24] *Evaluating the Error Resilience of Parallel Programs*, **Bo Fang**, Karthik Pattabiraman, Matei Ripeanu and Sudhanva Gurumurthi, Workshop on Fault Tolerance for High-Performance at Extreme Scale (FTXS) in conjunction with DSN 2014, 6 pages.
- [25] *Towards Building Error Resilient GPGPU Applications*, **Bo Fang**, Jiesheng Wei, Karthik Pattabiraman, Matei Ripeanu, 3rd IEEE Workshop on Resilient Architecture (WRA) in conjunction with MICRO 2012, 6 pages.
- [26] *Evaluating Error Resiliency of GPGPU Applications*, **Bo Fang**, Jiesheng Wei, Karthik Pattabiraman, Matei Ripeanu, High Performance Computing, Networking, Storage and Analysis (SC), 2012 poster.

Other Awards

University awards

- Nominated for the ACM doctoral dissertation award, Oct. 2020. 1 of 2 nominees, University of British Columbia.
- British Columbia Student Scholarship, Sep. 2019. \$15,000, University of British Columbia.
- Walter C Koerner Fellowship, Aug. 2019. \$5,300, University of British Columbia.
- J K Zee Memorial Fellowship, May. 2018. \$8,800, University of British Columbia.
- Graduate Support Initiative, Nov. 2017. \$12,000, University of British Columbia.
- Graduate Support Initiative, Nov. 2016. \$5,000, University of British Columbia.
- Graduate Support Initiative, Nov. 2014 Nov. \$6,000, University of British Columbia.
- First Place in Coding Competition of UBC Division, Oct. 2014. Microsoft.
- Second Place in Outstanding Annual Undergraduate, May 2005. Wuhan University.

Travel Grants

- Student Travel Grant, Jun. 2019. ACM ICS 2019.
- Student Travel Grant, Jun. 2017. ACM HPDC 2017.
- Student Travel Grant, Jun. 2016. IEEE DSN 2016.

Service

- Program committee member: IEEE/ACM CCGrid 2023, ACM HPDC 2023, IEEE/IFIP International Conference on Dependable Systems and Networks (DSN) 2023, 2022,2021, 2020; IEEE international conference on High-performance computing and Communications (HPCC) 2022, 2021, 2020, The IEEE Workshop on Silicon Errors in Logic – System Effects (SELSE) 2021, 2020; Fault Tolerance for HPC at eXtreme Scales (FTXS) workshop collocated with SC 2022, 2021, 2020.
- Invited reviewers for IEEE Transactions on Parallel and Distributed Systems (TPDS) and Journal of Parallel and Distributed Computing (JPDC).
- External reviewers for conferences such as The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC) 2021, 2020, The ACM International Symposium on High-performance Parallel and Distributed Computing (HPDC) 2020-2014, IEEE/IFIP International Conference on Dependable Systems and Networks (DSN) 2020-2015, IEEE International Symposium on Software Reliability Engineering (ISSRE) 2016, etc.
- Student volunteer for IEEE Pacific Rim International Symposium on Dependable Computing (PRDC) 2012 and IEEE International Conference on Software Quality, Reliability, and Security (QRS) 2015.