XIAODONG YU

Assistant Computer Scientist
Mathematics and Computer Science Division, Argonne National Laboratory
9700 South Cass Avenue, Building 240, Room 3146, Lemont, IL 60439 USA
xiaodong-yu.github.io | xyu@anl.gov | (630)252-3897
Google scholar: scholar.google.com/citations?user=1sefeCkAAAAJ&hl=en

RESEARCH INTERESTS

My research interests broadly lie in parallel, distributed, and reconfigurable computing, computer systems, and computer architecture. Specifically, I am working on system design and optimization to support large-scale deep learning; performance benchmarking and modeling of next-generation AI hardware; parallel compressor design and optimization for exascale scientific data; accelerator-based (e.g., GPU) application acceleration; and performance optimization of collective communications in GPU-aware MPI.

EDUCATION

Doctor of Philosophy, Computer Science and Applications
Virginia Tech, Blacksburg, VA, USA — August 2019
Dissertation - *Algorithms and Frameworks for Accelerating Security Applications on HPC Platforms*
Thesis committee: Prof. Danfeng (Daphne) Yao (Chair), Prof. Michela Becchi (NCSU), Prof. Ali Butt, Prof. Matthew Hicks, Prof. Xinming (Simon) Ou (USF)

Master of Science, Electrical Engineering
University of Missouri, Columbia, MO, USA — August 2013
Thesis - *Deep Packet Inspection on Large Datasets: Algorithmic and Parallelization Techniques for Accelerating Regular Expression Matching on Many-Core Processors*
Thesis Advisor: Prof. Michela Becchi

Bachelor of Science, Mathematics and Applied Mathematics
China University of Mining and Technology (CUMT), China — June 2008

PROFESSIONAL APPOINTMENTS

Argonne National Laboratory, Lemont, IL
Assistant Computer Scientist — June 2022 - Current
Postdoctoral Appointee — September 2019 - May 2022

- Designing benchmark suite for next-generation AI hardware (LDRD and FAIR-SBI).
- Optimizing data loading for large-scale surrogate training (FAIR-SBI).
- Designing and implementing fast GPU-based lossy compressors for scientific data (VeloC/SZ).
- Designing and implementing MPI collective communications with compression (Exascale MPI).
- Designing and implementing multi-GPU-based ptychographic image reconstruction (RAVEN).

The University of Chicago Consortium for Advanced Science and Engineering, Chicago, IL
Scientist-at-Large — November 2022 - Current

AMD, Austin, TX
Co-op Software Engineer — May 2017 - Aug. 2017

- Developing GPU Deep learning library (MIOpen).

Virginia Tech, Blacksburg, VA
Graduate Research Assistant — August 2013 - August 2019
Graduate Teaching Assistant — August 2013 - May 2017

- Designing and implementing GPU-assisted program analysis for vetting Android apps.
- Studying vulnerability of the systems to cache side-channel attacks.

- Designing and implementing programming framework for Micron's Automata Processor.
- Designing and implementing GPU-based CT image reconstruction.
- Teaching assistant of undergraduate courses: CS2505 and CS2506.

University of Missouri, Columbia, MO
Graduate Research Assistant May 2011 - August 2013

- Designing and implementing GPU-based automata processing.

GRANTS

- **PI**: LDRD Advanced Computing Expedition, "Scalability Study of AI-based Surrogate for Ptychographic Image Reconstruction on Graphcore."
  Award Amount: $50K
  Award Period: 07/01/2022 – 09/30/2022
- **Host and Primary collaborator**: NSF EPSCoR RII Track-4 Program, "Massively Parallel Graph Processing on Next-Generation Multi-GPU Supercomputers." (PI: Prof. Da Yan)
  Award Amount: $275K
  Award Period: 02/01/2023 – 01/31/2025

PROJECT LEADERSHIP

- LDRD Advanced Computing Expedition: *Scalability Study of AI-based Surrogate for Ptychographic Image Reconstruction on Graphcore.*
  My Role: PI
  Period: 07/01/2022 – 09/30/2022
- DOE ASCR: *FAIR-SBI: Surrogate Benchmarks Supporting AI and Simulation Research.*
  My Role: Senior Personnel and Technical Lead
  Period: 03/01/2021 – Current
  Supervision and Mentoring: Baixi Sun (Ph.D. student, August 2021 - Current), Chengming Zhang (Ph.D. student, September 2022 - Current)
- DOE ECP: *VeloC/SZ: VEry-Low Overhead Checkpointing System and Error-bounded Lossy Compression for Scientific HPC Datasets.*
  My Role: Senior Personnel and Technical Lead
  Period: 01/01/2021 – Current
  Supervision and Mentoring: Milan Shah (Ph.D. student, May 2022 - Current)
  Co-Mentoring: Jiannan Tian (Ph.D. student, January 2021 - Current), Cody Rivera (Undergraduate student, May 2021 - August 2021)
- DOE ECP: *Exascale MPI / MPICH.*
  My Role: Senior Personnel
  Period: 10/01/2021 – Current
  Co-Mentoring: Jiajun Huang (Ph.D. student, October 2022 - Current)

PUBLICATIONS

Refereed Journal Articles

1. [**Scientific Reports**] **Xiaodong Yu**, Viktor Nikitin, Daniel J Ching, Selin Aslan, Doga Gursoy, and Tekin Bicer, "Scalable and Accurate Multi-GPU-Based Image Reconstruction of Large-Scale Ptychography Data." *Scientific Reports* 12.1 (2022): 1-16. (2022 impact factor=4.996)
2. [**JSPS**] **Xiaodong Yu**, Hao Wang, Wu-chun Feng, Hao Gong, and Guohua Cao, "GPU-Based Iterative Medical CT Image Reconstructions." *Journal of Signal Processing Systems* 91.3 (2019): 321-338. (2022 impact factor=1.813)
3. [**JSAC**] **Xiaodong Yu**, Bill Lin, and Michela Becchi, "Revisiting State Blow-up: Automatically Building Augmented-FA while Preserving Functional Equivalence." *IEEE Journal on Selected Areas in Communications* 32.10 (2014): 1822-1833. (2022 impact factor=13.081)

Peer-Reviewed Conference Papers (with my student protégés underlined)

1. [**HPDC'22**] **Xiaodong Yu**, Sheng Di, Kai Zhao, Jiannan Tian, Dingwen Tao, Xin Liang, and Franck Cappello, "Ultrafast Error-Bounded Lossy Compression for Scientific Datasets." In *Proceedings of the 31st ACM International Symposium on High-Performance Parallel and Distributed Computing (HPDC 2022)*, Minneapolis, MN, June 27-29, 2022. (acceptance rate=19.4%)

2. [**IPDPS'22**] Cody Rivera, Sheng Di, Jiannan Tian, **Xiaodong Yu**, Dingwen Tao, and Franck Cappello, "Optimizing Huffman Decoding for Error-Bounded Lossy Compression on GPUs." In *Proceedings of IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, Lyon, France, May 30-June 3, 2022. (acceptance rate=27%)

3. [**SSSDU'21**] Guo, Yanfei, Ken Raffenetti, Hui Zhou, Travis Koehring, Sudheer Chunduri, **Xiaodong Yu**, and Rajeev Thakur, "Automated Validation and Verification for Scientific Software." In *Proceedings of the 2022 Workshop on the Science of Scientific-Software Development and Use*, virtual, December 13-15, 2021. (position paper)

4. [**SMC'21**] Tekin Bicer, **Xiaodong Yu**, Daniel Ching, Ryan Chard, Mathew Cherukara, Bogdan Nicolae, Rajkumar Kettimuthu, and Ian T Foster, "High-Performance Ptychographic Reconstruction with Federated Facilities." In *Proceedings of the 2021 Smoky Mountains Computational Sciences and Engineering Conference*, pages 173-189, Oak Ridge, TN, October 18-20, 2021.

5. [**CLUSTER'21**] **Xiaodong Yu**, Sheng Di, Ali Murat Gok, Dingwen Tao, and Franck Cappello, "cuZ-Checker: A GPU-Based Ultra-Fast Assessment System for Lossy Compressions." In *Proceedings of IEEE International Conference on Cluster Computing (CLUSTER)*, pages 307-319, Portland, OR, September 7-10, 2021. (acceptance rate=29%)

6. [**CLUSTER'21**] Jiannan Tian, Sheng Di, **Xiaodong Yu**, Cody Rivera, Kai Zhao, Sian Jin, Yunhe Feng, Xin Liang, Dingwen Tao, and Franck Cappello, "Optimizing Error-Bounded Lossy Compression for Scientific Data on GPUs." In *Proceedings of IEEE International Conference on Cluster Computing (CLUSTER)*, pages 283-293, Portland, OR, September 7-10, 2021. (acceptance rate=29%)

7. [**ICS'21**] **Xiaodong Yu**, Tekin Bicer, Rajkumar Kettimuthu, and Ian Foster, "Topology-Aware Optimizations for Multi-GPU Ptychographic Image Reconstruction." In *Proceedings of the ACM International Conference on Supercomputing (ICS)*, pages 354-366, Worldwide online event, June 14-18, 2021. (acceptance rate=24.2%)

8. [**IPDPS'20**] **Xiaodong Yu**, Fengguo Wei, Xinming Ou, Michela Becchi, Tekin Bicer, and Danfeng Yao, "GPU-Based Static Data-Flow Analysis for Fast and Scalable Android App Vetting." In *Proceedings of IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 274-284, New Orleans, LA, May 18-22, 2020. (acceptance rate=24.7%)

9. [**CSET'19**] **Xiaodong Yu**, Ya Xiao, Kirk Cameron, and Danfeng (Daphne) Yao, "Comparative Measurement of Cache Configurations' Impacts on Cache Timing Side-Channel Attacks." In *Proceedings of the 12th USENIX Conference on Cyber Security Experimentation and Test (CSET 19)*, Santa Clara, CA, August 12, 2019. (acceptance rate=31%)

10. [**SoutheastCon'18**] Thomas C. H. Lux, Layne T. Watson, Tyler H. Chang, Jon Bernard, Bo Li, **Xiaodong Yu**, Li Xu, Godmar Back, Ali R. Butt, Kirk W. Cameron, Yili Hong, Danfeng Yao, "Nonparametric Distribution Models for Predicting and Managing Computational Performance Variability." In *Proceedings of IEEE Southeastcon 2018*, pages 1-7, St. Petersburg, FL, April 19-22, 2018.

11. [**ACMSE'18**] Thomas C. H. Lux, Layne T. Watson, Tyler H. Chang, Jon Bernard, Bo Li, **Xiaodong Yu**, Li Xu, Godmar Back, Ali R. Butt, Kirk W. Cameron, Danfeng Yao, Yili Hong, "Novel Meshes for Multivariate Interpolation and Approximation." In *Proceedings of the ACMSE 2018 Conference*, pages 13:1-13:7, Richmond, KY, March 29-31, 2018.

12. [**BigData'17**] **Xiaodong Yu**, Kaixi Hou, Hao Wang, and Wu-chun Feng, "Robotomata: A Framework for Approximate Pattern Matching of Big Data on an Automata Processor." In *Proceedings of the IEEE International Conference on Big Data (IEEE BigData)*, pages 283–292, Boston, MA, December 11-14, 2017. (acceptance rate=17.9%)

13. [**ICS'17**] Marziyeh Nourian, Xiang Wang, **Xiaodong Yu**, Wu-chun Feng, and Michela Becchi, "De-

mystifying Automata Processing: GPUs, FPGAs or Micron's AP?" In *Proceedings of the International Conference on Supercomputing (ICS)*, pages 1-11, Chicago, IL, June 14-16, 2017. (acceptance rate=15.8%)

14. [**CF'17**] **Xiaodong Yu**, Hao Wang, Wu-chun Feng, Hao Gong, and Guohua Cao, "An Enhanced Image Reconstruction Tool for Computed Tomography on GPUs." In *Proceedings of the ACM International Conference on Computing Frontiers (CF)*, pages 97-106, Siena, Italy, May 15-17, 2017. (acceptance rate=35.5%)

15. [**CCGrid'16**] **Xiaodong Yu**, Hao Wang, Wu-chun Feng, Hao Gong, and Guohua Cao, "cuART: Fine-Grained Algebraic Reconstruction Technique for Computed Tomography Images on GPUs." In *Proceedings of 16th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*, pages 165-168, Cartagena, Colombia, May 16-19, 2016. (short) (acceptance rate=25%)

16. [**ANCS'16**] **Xiaodong Yu**, Wu-chun Feng, Danfeng Yao, and Michela Becchi, "O3FA: A Scalable Finite Automata-based Pattern-Matching Engine for Out-of-Order Deep Packet Inspection." In *Proceedings of the ACM/IEEE Symposium on Architectures for Networking and Communications Systems (ANCS)*, pages 1-11, Santa Clara, CA, March 17-18, 2016. (acceptance rate=20.7%)

17. [**CF'13**] **Xiaodong Yu**, and Michela Becchi, "GPU Acceleration of Regular Expression Matching for Large Datasets: Exploring the Implementation Space." In *Proceedings of the ACM International Conference on Computing Frontiers (CF)*, pages 1-10, Ischia, Italy, May 14-16, 2013.

Pending Conference Papers

1. Baixi Sun, **Xiaodong Yu**, Chengming Zhang, Jiannan Tian, Sian Jin, Kamil A. Iskra, Tao Zhou, Tekin Bicer, Pete Beckman, and Dingwen Tao, "SOLAR: A Highly Optimized Data Loading Framework for Distributed Training of CNN-based Scientific Surrogates." Under review in *49th International Conference on Very Large Data Bases (VLDB)*, Vancouver, Canada, August 28-September 1, 2023.

2. Milan Shah, **Xiaodong Yu**, Sheng Di, Danylo Lykov, Yuri Alexeev, Michela Becchi, and Franck Cappello, "GPU-Accelerated Error-Bounded Compression Framework for Quantum Circuit Simulations." Under review in *IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, St. Petersburg, FL, May 15-19, 2023.

3. Boyuan Zhang, Jiannan Tian, Sheng Di, **Xiaodong Yu**, Dingwen Tao, and Franck Cappello, "Fast GPU Lossy Compressor for Scientific Computing Applications." Under review in *IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, St. Petersburg, FL, May 15-19, 2023.

4. Boyuan Zhang, Jiannan Tian, Sheng Di, **Xiaodong Yu**, Lizhi Xiang, Dingwen Tao, and Franck Cappello, "gpuLZ: Optimizing LZSS for Scientific Lossy Compression on Modern GPUs." Under review in *the 28th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP)*, Montreal, Canada, February 25-March 1, 2023.

Peer-Reviewed Posters/Extended Abstracts

1. [**SC'22**] Baixi Sun, **Xiaodong Yu**, Kamil Iskra, and Dingwen Tao, "SurrogateTrain: Drastically Improving Performance of Data Loading for Training Scientific Surrogate Models." In *Proceedings of the 2022 International Conference for High Performance Computing, Networking, Storage, and Analysis (SC)*, Dallas, TX, November 13-18, 2022. (ACM Student Research Competition track)

2. [**SC'22**] Milan Shah, **Xiaodong Yu**, Sheng Di, Franck Cappello, and Michela Becchi, "Compressing Quantum Circuit Simulation Tensor Data." In *Proceedings of the 2022 International Conference for High Performance Computing, Networking, Storage, and Analysis (SC)*, Dallas, TX, November 13-18, 2022. (ACM Student Research Competition track)

3. [**ECPAM'22**] Yanfei Guo, Kenneth Raffenetti, Rob Latham, Marc Snir, Hui Zhou, Travis Koehring, Sudheer Chunduri, **Xiaodong Yu**, and Rajeev Thakur, "2.3.1.07 Exascale MPI (MPICH)." In *Proceedings of the 2022 Exascale Computing Project Annual Meeting*, virtual, May 2-6, 2022.

4. [**SC'18**] **Xiaodong Yu**, Michela Becchi, and Danfeng (Daphne) Yao, "The Algorithm and Framework Designs and Optimizations for Scalable Automata Processing on HPC Platforms." In *Proceedings of the 2018 International Conference for High Performance Computing, Networking, Storage, and Analysis*

(SC), Dallas, TX, November 11-16, 2018. (Doctoral Showcase track)

5. [**IISWC'17**] **Xiaodong Yu**, Kaixi Hou, Hao Wang, and Wu-chun Feng, "A Framework for Fast and Fair Evaluation of Automata Processing Hardware." In *Proceedings of IEEE International Symposium on Workload Characterization (IISWC)*, pages 120-121, Seattle, WA, October 1-3, 2017.
6. [**PPoPP'13**] **Xiaodong Yu**, and Michela Becchi, "Exploring Different Automata Representations for Efficient Regular Expression Matching on GPUs." In *Proceedings of the 18th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP)*, pages 287-288, Shenzhen China, February 23-27, 2013.

Theses

1. [**Doctoral Dissertation**] **Xiaodong Yu**, "Algorithms and Frameworks for Accelerating Security Applications on HPC Platforms." Virginia Tech, 2019.
2. [**Master Thesis**] **Xiaodong Yu**, "Deep Packet Inspection on Large Datasets: Algorithmic and Parallelization Techniques for Accelerating Regular Expression Matching on Many-Core Processors." University of Missouri-Columbia, 2013.

SELECTED CONFERENCE PRESENTATIONS

- ACM HPDC: *Ultrafast Error-Bounded Lossy Compression for Scientific Datasets*, June 2022.
- IEEE CLUSTER: *cuZ-Checker: A GPU-Based Ultra-Fast Assessment System for Lossy Compressions*, September 2021.
- ACM ICS: *Topology-Aware Optimizations for Multi-GPU Ptychographic Image Reconstruction*, June 2021.
- IEEE IPDPS: *GPU-Based Static Data-Flow Analysis for Fast and Scalable Android App Vetting*, May 2020.
- USENIX CSET: *comparative Measurement of Cache Configurations' Impacts on Cache Timing Side-Channel Attacks*, August 2019.
- IEEE BigData: *Robotomata: A Framework for Approximate Pattern Matching of Big Data on an Automata Processor*, December 2017.
- ACM/IEEE ANCS: *O3FA: A Scalable Finite Automata-based Pattern-Matching Engine for Out-of-Order Deep Packet Inspection*, March 2016.

SUPERVISION/MENTORING EXPERIENCE

Postdoc Mentoring

- Yixuan Sun – Postdoctoral Appointee in ANL MCS, November 2022 - Current.

Student Supervision and Mentoring

- Baixi Sun (Ph.D. student in Intelligent Systems Engineering Department at Indiana University) – Visiting Student–Subcontractor in ANL MCS, August 2021 - Current.
- Milan Shah (Ph.D. student in Electrical and Computer Engineering Department at North Carolina State University) – Visiting Student–Subcontractor in ANL MCS, May 2022 - Current.
- Chengming Zhang (Ph.D. student in Intelligent Systems Engineering Department at Indiana University) – Visiting Student–Subcontractor in ANL MCS, September 2022 - Current.

Student Co-Mentoring

- Jiannan Tian (Ph.D. student in Intelligent Systems Engineering Department at Indiana University) – Visiting Student–Subcontractor in ANL MCS, January 2021 - Current.
- Jiajun Huang (Ph.D. student in Computer Science and Engineering Department at the University of California, Riverside) – Visiting Student–Subcontractor in ANL MCS, October 2022 - Current.
- Cody Rivera (Undergraduate student in the Computer Science Department at the University of Alabama) – Science Undergraduate Laboratory Internship (SULI) Program Intern in ANL MCS, May 2021 - August 2021.

- Boyuan Zhang (Ph.D. student in Intelligent Systems Engineering Department at Indiana University), May 2022 - Current.

PROFESSIONAL SERVICES

Conference Organizing Committee

- **Finance Chair** – IEEE International Symposium on Performance Analysis of Systems and Software (**ISPASS**), 2023.
- **Publicity Co-Chair** – IEEE International Workshop on Signal Processing Systems (**SiPS**), 2022.

Journal Editorial Board

- **Review Board Member** – IEEE Transactions on Parallel and Distributed Systems (**TPDS**), 2022 - Current.

Conference Technical Program Committee (TPC)

- **TPC Member** – The Third International Workshop on Big Data Reduction (**IWBDR-3**) in conjunction with 2022 IEEE International Conference on Big Data (IEEE BigData), 2022.
- **TPC Member** – The 24th IEEE International Conferences on High Performance Computing and Communications (**HPCC**), 2022.
- **TPC Member** – 9th IEEE/ACM International Conference on Big Data Computing, Applications and Technologies (**BDCAT**), 2022.
- **TPC Member** – IEEE International Workshop on Signal Processing Systems (**SiPS**), 2022.
- **Posters Committee Member** – IEEE International Symposium on Performance Analysis of Systems and Software (**ISPASS**), 2022.
- **TPC Member** – The 23rd IEEE International Conferences on High Performance Computing and Communications (**HPCC**), 2021.
- **AD/AE Committee Member** – The International Conference for High Performance Computing, Networking, Storage, and Analysis (**SC**), 2021.
- **TPC Member** – IEEE Hot Interconnects symposium (**HotI**), 2021.
- **TPC Member** – The 22nd IEEE International Conferences on High Performance Computing and Communications (**HPCC**), 2020.

Journal Referee

- ACM Transactions on Storage (**TOS**), 2022.
- Future Generation Computer Systems (**FGCS**) - Elsevier, 2022
- IEEE Transactions on Parallel and Distributed Systems (**TPDS**), 2020 - 2022.
- Journal of Parallel and Distributed Computing (**JPDC**) - Elsevier, 2019 - 2022.
- Journal of Systems Architecture (**JSA**) - Elsevier, 2019, 2022.
- Software: Practice and Experience (**SPE**) - Wiley, 2021.
- Journal of Computer Virology and Hacking Techniques (**JICV**) - Springer Nature, 2021.
- IEEE Access, 2017 - 2021.
- IEEE Transactions on Dependable and Secure Computing (**TDSC**), 2018 - 2020.
- Computer Communications (**COMCOM**) - Elsevier, 2018, 2020.
- Computer Networks (**COMNET**) - Elsevier, 2019.
- Australasian Physical & Engineering Sciences in Medicine (**APES**) - Springer Nature, 2019.
- Journal of Systems and Software (**JSS**) - Elsevier, 2015 - 2017.
- IEEE Journal on Selected Areas in Communications (**JSAC**), 2014.

Conference External Reviewer

- The ACM Symposium on Access Control Models and Technologies (**SACMAT**), 2022.
- 36th IEEE International Parallel & Distributed Processing Symposium (**IPDPS**), 2022.

- 26th International Workshop on High-Level Parallel Programming Models and Supportive Environments (**HIPS**) – held in conjunction with IPDPS, 2021.
- The 27th International Conference on Computer Communications and Networks (**ICCCN**), 2018.
- IEEE Global Communications Conference (**GLOBECOM**), 2018.
- IEEE 88th Vehicular Technology Conference (**VTC**), 2018.

ANL Internal Services

- **Co-Organizer** – CS seminar series hosted by the Mathematics and Computer Science (MCS) Division, 2022 - current.
- **Interviewer** – two MCS postdoc hires, 2022.
- **TPC Member** – Argonne Postdoctoral Research and Career Symposium, 2022.

Other Services

- **Mentor** – SC22 Mentor-Protégé Program, 2022.
- **Judge** – SIGHPC Computational and Data Science Fellowship, 2022.
- **Executive Committee Member** – IEEE Fox Valley Subsection, 2019 - 2020.
- **Student Volunteer** – The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC), 2018.

MAJOR SOFTWARE

- **SZx**: an ultra-fast error-bounded lossy compressors.
  Repository: https://github.com/szcompressor/SZx
  My role: principal designer, developer, and maintainer of the GPU version
- **cuZ-Checker**: a GPU-based library to characterize the data and check the compression results of lossy compressors.
  Repository: https://github.com/CODARcode/cuZ-checker
  My role: principal designer, developer, and maintainer
- **cuSZ**: a CUDA-based error-bounded lossy compressor for scientific data.
  Repository: https://github.com/szcompressor/cuSZ
  My role: core developer and maintainer of the HIP version
- **Tike**: a toolbox for tomographic reconstruction of 3D objects from ptychography data.
  Repository: https://github.com/tomography/tike
  My role: principal designer, developer, and maintainer of the HPC part
- **SOLAR**: a data loading framework for training CNN-based scientific surrogates.
  Repository: https://github.com/ANL-FAIR-SBI/SOLAR
  My role: principal designer and architect
- **MPICH**: a high-performance and widely portable implementation of the MPI-3.1 standard.
  Repository: https://github.com/pmodels/mpich
  My role: developer and maintainer of the GPU-related part

HONORS AND AWARDS

- Second place in ACM Student Research Competition - Graduate Students at ACM/IEEE SC'22 (as the co-advisor of Milan Shah), 2022.
- Selected to present at ACM/IEEE SC'18 Doctoral Showcase (Selection Rate = 16/27), 2018.
- Outstanding Graduate Teaching Assistant Award, Computer Science Department at Virginia Tech, 2016.
- Learning Progress Scholarship, China University of Mining and Technology (CUMT), 2006 - 2007.

REFERENCES

Danfeng (Daphne) Yao
Professor, Elizabeth and James E. Turner Jr. '56 Faculty Fellow, and CACI Faculty Fellow
Virginia Tech, Blacksburg, VA
danfeng@vt.edu

Michela Becchi
Associate Professor
North Carolina State University, Raleigh, NC
mbecchi@ncsu.edu

Dingwen Tao
Associate Professor
Indiana University, Bloomington, IN
ditao@iu.edu

Pete Beckman
Co-Director, Northwestern Argonne Institute of Science and Engineering and Argonne Distinguished Fellow
Argonne National Laboratory, Lemont, IL
beckman@mcs.anl.gov