

# Tanzima Z. Islam

Associate Professor  
Texas State University

601 University Drive  
San Marcos, TX, 78666  
✉ [tanzima@txstate.edu](mailto:tanzima@txstate.edu)  
🌐 [www.tanzimaislam.com](http://www.tanzimaislam.com)



## Education

- 2013 **Ph.D., Electrical and Computer Engineering**, *Purdue University*.  
*Thesis: Reliable and scalable checkpointing systems for distributed computing environments.*
- 2006 **B.Sc., Computer Science & Engineering**, *Bangladesh University of Engineering & Technology*.

## Professional Experience

- 2025-Present **Associate Professor**, *Department of Computer Science*, Texas State University (TXST).
- 2020-2025 **Assistant Professor**, *Department of Computer Science*, Texas State University (TXST).
- Summer 2024 **Visiting Faculty**, *Oak Ridge National Laboratory*, Computer Science and Mathematics Division (CSMD).
- Summer 2022, 2023 **Visiting Scholar**, *Brookhaven National Laboratory (BNL)*, Computational Science Initiative.
- Summer 2019 **Visiting Scholar**, *Lawrence Berkeley National Laboratory (LBNL)*, Computation Directorate.
- 2017-2019 **Assistant Professor**, *Department of Computer Science*, Western Washington University (WWU).
- 2013-2017 **Postdoctoral Research Staff Member**, *Lawrence Livermore National Laboratory (LLNL)*, Center for Applied Scientific Computing.  
Developed machine-learning techniques for performance analysis.
- 2006-2007 **Member, Research and Development**, *Commlink Info Tech Ltd.*, Bangladesh.  
Developed software for a service-independent telecommunication network (Intelligent Network).

## Awards & Honors & Inventions

- 2025 **NSF CAREER Award**.
- 2025 **Best Paper Award**. IEEE Cloud Summit.
- 2024, 2025 **Research Millionaire Award**, College of Science and Engineering, Texas State University.
- 2023 **Presidential Seminar Award for Excellence in Scholarly/Creative Activities** at Texas State University.
- 2023, 2025 **Presidential Distinction for Excellence in Scholarly/Creative Activities**, College of Science and Engineering (CoSE) at Texas State University.
- 2023 **Valero Award for Excellence in Scholarly/Creative Activities**, Academic Affairs Council at Texas State University.
- 2022 **DOE CAREER Award**.
- 2021 **College of Science and Engineering Achievement Award for Excellence in Scholarly/Creative Activities** at Texas State University.
- 2019 **R&D 100 Award**, in collaboration with Lawrence Livermore National Laboratory for inventing a scalable software solution—the Scalable Checkpoint/Restart Framework 2.0 (SCR)—that makes High Performance Computing Systems reliable and efficient.
- 2014 **Lawrence Livermore National Laboratory (LLNL) Director's Science & Technology Award for Excellence in Publication**.
- 2016, 2015, 2014 Best Poster Award, LLNL Annual Scholars Poster Symposium.
- 2014 2<sup>nd</sup> Place Winner, LLNL Computation Directorate Postdoctoral Poster Symposium.

- 2012, 2009 Best Paper Finalist, International Conference for High Performance Computing, Networking, Storage and Analysis (SC).
- 2010 2<sup>nd</sup> Place Winner, ACM Student Research Competition.

## Research and Other Funding

**Secured \$4.6M+ in external funding as PI/Co-PI over the past five years.**

- 2026 **PI**, Autonomous Performance Optimization, AMD, \$50K.
- 2025 **PI**, CAREER: SPEED: Scalable Performance-aware Engine for Efficient Decision-making on HPC Systems, National Science Foundation (NSF) CAREER program, \$650K.
- 2025 **PI**, AI/ML for Performance Characterization of AMD GPUs, AMD, \$25K.
- 2024 **PI**, Application Fingerprinting, Department of Energy's STEP: Software Tools Ecosystem Project at Oak Ridge National Laboratory, \$16K.
- 2023 **PI**, SRP-HPC: Performance Analysis using Machine Learning: Use Case Chimbuko, DOE SRP Fellowship, \$85K.
- 2023-2026 **Sub-contract**, FRACTAL-SI: Resource Allocation for Center-Wide Throughput Acceleration Leveraging Elasticity, Strategic Initiative from Laboratory Directed Research and Development at Lawrence Livermore National Laboratory, \$3.3M/year (TXST PI: 300K).
- 2022-2027 **PI**, INTELYTICS: An Efficient Data-Driven Decision-Making Engine for Performance In the Era of Heterogeneity, DOE Early Career Research Program (ECRP), \$770K. (10% acceptance rate)
- 2022-2025 **Co-PI**, RECUP: Scalable Metadata And Provenance Services for Reproducible Hybrid Workflows, Department of Energy (DOE), \$2.43M (TXST PI: 300K).
- 2022-2023 **Sub-contract**, ICE4HPC: Toward the Intelligent Center for HPC, Laboratory Directed Research and Development at Lawrence Livermore National Laboratory, \$100K.
- 2022 **PI**, SRP-HPC: Performance Characterization of Workflow Applications, DOE SRP Fellowship, \$68K.
- 2021-Present **PI**, PerfROCm: Cross-platform performance prediction and analysis using deep learning, AMD Research Gift, \$175K.
- 2021 **PI**, Predicting Performance using Few Shot Learning, Research Enhancement Program (REP) at Texas State University, \$8K.
- 2020 **Co-PI**, Equipment grant: 5 Petaflops AMD HPC COVID-19 cluster, \$400K (estimated value).
- Summer 2019 **PI**, Proxy Application Validation for Exascale Co-design, DOE Sustainable Research Pathway Fellowship, \$43K.
- 2018 - Present **PI**, HPC@Scale Education Grant, Time allocation grant from TACC. 5,000+ core-hours.
- 2018 **Co-PI**, Scientific Data Visualization course development, Office of Research and Sponsored Programs at Western Washington University, \$12K.
- 2016 **PI**, VERITAS for Understanding Performance Evolution during Code Development, Linking Exploratory Application Research to Next-gen Development at Lawrence Livermore National Laboratory, \$200K.
- REU & Pathway Development Grants 2021-2024 **Member**, REU Site: Research Experiences for Undergraduates in Edge Computing, NSF, \$389K. **Member**, Cross-Institutional Research Pathways to Graduate Education in the Natural Sciences: A Partnership between Texas State University and the University of Colorado, SLOAN Foundation, \$250K.

## Research Publications

### Papers

- [1] Ashna N. Ahmed, Banooqa Banday, Terry Jones, and **Tanzima Islam**. Attention-Informed Surrogates for Navigating Power-Performance Trade-offs in HPC. In *Workshop on ML for Systems in Conjunction with NeurIPS*, 2025. Accepted.
- [2] Arunavo Dey, Neil Antony, Aakash R. Dhakal, Kowshik Thopalli, Jayaraman J. Thiagarajan, Tapasya

- Patki, Aniruddha Marathe, Tom Scogland, Jae-Seung Yeom, and **Tanzima Islam**. ModelX: A Novel Transfer Learning Approach Across Heterogeneous Datasets. In *The 34th ACM International Symposium on High-Performance Parallel and Distributed Computing (HPDC)*, 2025. Accepted (19% acceptance rate).
- [3] Aakash R. Dhakal, **Tanzima Islam**, Arunavo Dey, Daniel Nichols, Abhinav Bhatele, Tapasya Patki, Tom Scogland, and Jae-Seung Yeom. xamm: "attention" to details improves cross-platform prediction accuracy. In *The 25th IEEE International Symposium on Cluster, Cloud and Internet Computing (CCGRID)*, 2025. Accepted (25% acceptance rate).
- \* [4] Elvis G. Fefey and **Tanzima Islam**. Optimizing deep learning inference on heterogeneous devices on the edge: An ilp-based rate monotonic scheduling approach. In *2025 IEEE Cloud Summit*, pages 23–28, 2025. **Best Paper Award**.
- [5] **Tanzima Islam**, Aniruddha Marathe, Holland Schutte, and Mohammad Zaeed. [Data-driven analysis to understand gpu hardware resource usage of optimizations](#), 2025. Accepted in International Journal of High Performance Computing and Applications (IJHPCA).
- [6] Ankur Lahiry, Banooqa Banday, and **Tanzima Islam**. [Wander: An explainable decision-support framework for hpc](#). In *Arxiv*, 2025.
- [7] Tarek Ramadan, Nathan Pinnow, Chase Phelps, Jayaraman J. Thiagarajan, and **Tanzima Islam**. [Structure-aware representation learning for effective performance prediction](#). *Concurrency and Computation: Practice and Experience*, 37(9-11):e70046, 2025.
- [8] Raül Sirvent, Rocío Carratalá-Saez, Amal Gueroudji, **Tanzima Islam**, Line Pouchard, and Michela Tauber. [Reproducibility for hpc and distributed environments: Committees, nondeterminism, performance and workflows](#). ACM REP '25, New York, NY, USA, 2025. Association for Computing Machinery.
- [9] Mohammad Zaeed, **Tanzima Islam**, and Vladimir Indić. [Opal: A modular framework for optimizing performance using analytics and llms](#). In *Arxiv*, 2025.
- [10] Banooqa Banday, Kowshik Thopalli, **Tanzima Islam**, and Jayaraman J. Thiagarajan. On the role of prompt construction in enhancing efficacy and efficiency of llm-based tabular data generation. In *2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, January 2025. (Accepted).
- [11] Banooqa Banday, **Tanzima Islam**, and Aniruddha Marathe. Perfgen: A synthesis and evaluation framework for performance data using generative ai. In *48th Annual Computers, Software, and Applications Conference (COMPSAC)*. IEEE, 2024. Acceptance rate: 24%.
- [12] Arunavo Dey, Aakash Dhakal, **Tanzima Islam**, Jae-Seung Yeom, Tapasya Patki, Daniel Nichols, Alexander Movsesyan, and Abhinav Bhatele. Relative performance prediction using few-shot learning. In *48th Annual Computers, Software, and Applications Conference (COMPSAC)*. IEEE, 2024. Acceptance rate: 24%.
- [13] Elvis G. Fefey and **Tanzima Islam**. Toward efficient deep learning inference: On-node heterogeneous scheduling in edge-cloud infrastructure. In *2024 IEEE Cloud Summit*, pages 73–78, 2024.
- [14] Christopher Kelly, Wei Xu, Line C. Pouchard, Hubertus Van Dam, **Tanzima Islam**, Shinjae Yoo, and Kerstin Kleese Van Dam. Performance analysis and data reduction for exascale scientific workflows. In *International Journal of High Performance Computing and Applications*. SageJournals, 2024.
- [15] Chase Phelps, Ankur Lahiry, **Tanzima Islam**, and Line Pouchard. Reimagine application performance as a graph: Novel graph-based method for performance anomaly classification in high-performance computing. In *48th Annual Computers, Software, and Applications Conference (COMPSAC)*. IEEE, 2024. Acceptance rate: 24%.
- [16] Mohammad Zaeed, **Tanzima Islam**, and Vladimir Indic. Characterize and compare the performance of deep learning optimizers in recurrent neural network architectures. In *48th Annual Computers, Software, and Applications Conference (COMPSAC)*. IEEE, 2024. Acceptance rate: 24%.
- [17] Amal Gueroudji, Chase Phelps, **Tanzima Islam**, Philip Carns, Shane Snyder, Matthieu Dorier, Robert B. Ross, and Line C. Pouchard. Performance characterization and provenance of distributed task-based workflows on hpc platforms. In *2024 IEEE/ACM Workshop on Workflows in Support of Large-Scale Science (WORKS)*, Nov 2024. (Accepted).
- [18] Arunavo Dey, Chase Phelps, **Tanzima Islam**, and Christopher Kelly. Signal processing based method for real-time anomaly detection in high-performance computing. In *47th Annual Computers, Software, and Applications Conference (COMPSAC)*. IEEE, 2023. (Accepted). Acceptance rate: 27%.
- [19] Bogdan Nicolae, **Tanzima Islam**, Robert Ross, Huub Van Dam, Kevin Assogba, Polina Shpilker, Mikhail Titov, Matteo Turilli, Tianle Wang, Ozgur O. Kilic, Shantenu Jha, and Line C. Pouchard. Building the i

- (interoperability) of fair for performance reproducibility of large-scale composable workflows in recup. In *2023 IEEE 19th International Conference on e-Science (e-Science)*, pages 1–7, 2023.
- [20] Chase Phelps and **Tanzima Islam**. Automatic parallelization of cellular automata for heterogeneous platforms. In *47th Annual Computers, Software, and Applications Conference (COMPSAC)*. IEEE, 2023. (Accepted). Acceptance rate: 27%.
  - [21] Tarek Ramadan, Ankur Lahiry, and **Tanzima Islam**. Novel representation learning technique using graphs for performance analytics. In *2023 International Conference on Machine Learning and Applications (ICMLA)*, pages 1311–1318, 2023.
  - [22] Ishan Abhinit, Emily K. Adams, Khairul Alam, Brian Chase, Ewa Deelman, Lev Gorenstein, Stephen Hudson, **Tanzima Islam**, Jeffrey Larson, Geoffrey Lentner, Anirban Mandal, John-Luke Navarro, Bogdan Nicolae, Line Pouchard, Rob Ross, Banani Roy, Mats Rynge, Alexander Serebrenik, Karan Vahi, Stefan Wild, Yufeng Xin, Rafael Ferreira da Silva, and Rosa Filgueira. Novel proposals for fair, automated, recommendable, and robust workflows. In *2022 IEEE/ACM Workshop on Workflows in Support of Large-Scale Science (WORKS)*, pages 84–92, 2022.
  - [23] Line C. Pouchard, **Tanzima Islam**, and Bogdan Nicolae. [Challenges for implementing fair digital objects with high performance workflows](#). *Research Ideas and Outcomes*, 8:e94835, 2022.
  - [24] Holland Schutte, Chase Phelps, Aniruddha Marathe, and **Tanzima Islam**. LIBNVCD: An extendable and user-friendly multi-gpu performance measurement tool. In *46th Annual Computers, Software, and Applications Conference (COMPSAC)*. IEEE, 2022. Acceptance rate: 20%.
  - [25] **Tanzima Islam**, Philip Wu Liang, Forest Sweeney, Cody Pragner, Jayaraman J Thiagarajan, Moushumi Sharmin, and Shameem Ahmed. College life is hard!-shedding light on stress prediction for autistic college students using data-driven analysis. In *45th Annual Computers, Software, and Applications Conference (COMPSAC)*, pages 428–437. IEEE, 2021. Acceptance rate: 20%.
  - [26] **Tanzima Islam** and Chase Phelps. Hpc@scale: A hands-on approach for training next-gen hpc software architects. In *EduHiPC workshop at the 28th International Conference on High Performance Computing, Data, Analytics (HiPC)*. IEEE, 2021. Invited paper.
  - [27] Quentin Jensen, Filip Jagodzinski, and **Tanzima Islam**. Filcio: Application agnostic i/o aggregation to scale scientific workflows. In *2021 IEEE 45th Annual Computers, Software, and Applications Conference (COMPSAC)*, pages 1587–1592. IEEE, 2021.
  - [28] Nathan Pinnow, Tarek Ramadan, **Tanzima Islam**, Chase Phelps, and Jayaraman J. Thiagarajan. Comparative code structure analysis using deep learning for performance prediction. In *International Symposium on Performance Analysis of Systems and Software (ISPASS)*. IEEE, March 28-30 2021. (Accepted). Acceptance rate: 36%.
  - [29] Jack Stratton, Michael Albert, Quentin Jensen, Max Ismailov, Filip Jagodzinski, and **Tanzima Islam**. Towards aggregation based i/o optimization for scaling bioinformatics applications. In *2020 IEEE 44th Annual Computers, Software, and Applications Conference (COMPSAC)*, pages 1250–1255, 2020.
  - [30] Tapasya Patki, Jayaraman J. Thiagarajan, Alexis Ayala, and **Tanzima Islam**. Performance optimality or reproducibility: that is the question. In *International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, November 17-22 2019. Acceptance rate: 20%.
  - [31] **Tanzima Islam**, Alexis Ayala, Quentin Jensen, and Khaled Ibrahim. Toward A Programmable Analysis and Visualization Framework for Interactive Performance Analytics. In *Workshop on Programming and Performance Visualization Tools held in conjunction with International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, Denver, CO, November 16-23 2019. IEEE.
  - [32] Nicholas Majeske, Filip Jagodzinski, Brian Hutchinson, and **Tanzima Islam**. Low Rank Smoothed Sampling Methods for Identifying Impactful Pairwise Mutations. In *International Conference on Bioinformatics, Computational Biology, and Health Informatics*, pages 681–686. ACM, 2018.
  - [33] Jayaraman J. Thiagarajan, Rushil Anirudh, Bhavya Kailkhura, Nikhil Jain, **Tanzima Islam**, Abhinav Bhatele, Jae-Seung Yeom, and Todd Gamblin. [PADDLE: Performance Analysis using a Data-driven Learning Environment](#). In *IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, May 21-25 2018. Acceptance rate: 24.5%.
  - [34] Teng Wang, Adam Moody, Yeh Zhu, Kathryn Mohror, Kento Sato, **Tanzima Islam**, and Waikuan Yu. [MetaKV: A Key-Value Store for Metadata Management of Distributed Burst Buffers](#). In *IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 1174–1183, May 2017. Acceptance rate: 22%.

- [35] **Tanzima Islam**, Kathryn Mohror, and Martin Schulz. [Exploring the MPI Tool Information Interface: Features and Capabilities](#). In *The International Journal of High Performance Computing Applications (IJHPCA)*, volume 30, pages 212–222, 2016.
- [36] Tania Banerjee, Jason Hackl, Mrugesh Shringarpure, **Tanzima Islam**, S Balachandar, Thomas Jackson, and Sanjay Ranka. [CMT-Bone — A Proxy Application for Compressible Multiphase Turbulent Flows](#). In *IEEE 23rd International Conference on High Performance Computing (HiPC)*, pages 173–182, Dec 2016. Acceptance rate: 23%.
- [37] **Tanzima Islam**, Jayaraman J. Thiagarajan, Abhinav Bhatele, Martin Schulz, and Todd Gamblin. [A Machine-Learning Framework for Performance Coverage Analysis of Proxy Applications](#). In *International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, Salt Lake City, UT, November 13-18 2016. Acceptance rate: 23%.
- [38] Lee Savoie, David K. Lowenthal, Bronis R. de Supinski, **Tanzima Islam**, Kathryn Mohror, Barry Rountree, and Martin Schulz. [I/O Aware Power Shifting](#). In *IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 740–749, May 2016. Acceptance rate: 23%.
- [39] **Tanzima Islam**, Kathryn Mohror, and Martin Schulz. Exploring the Capabilities of the New MPI\_T Interface. In *Proceedings of the 21st European MPI Users' Group Meeting*, page 91. ACM, 2014.
- [40] Anup Mohan, Thomas Hacker, Gregory P. Rodgers, and **Tanzima Islam**. [Batchsubmit: A high-volume Batch Submission System for Earthquake Engineering Simulation](#). In *Concurrency and Computation: Practice and Experience*, volume 26, pages 2240–2252. Wiley Online Library, 2014.
- [41] John Tramm, Andrew Siegel, **Tanzima Islam**, and Martin Schulz. XSBench—the Development and Verification of a Performance Abstraction for Monte Carlo Reactor Analysis. *The Role of Reactor Physics toward a Sustainable Future (PHYSOR)*, 2014.
- [42] **Tanzima Islam**, Saurabh Bagchi, and Rudolf Eigenmann. [Reliable and Efficient Distributed Checkpointing System for Grid Environments](#). In *Journal of Grid Computing (JoGC)*, volume 12, pages 593–613, Dec 2014.
- [43] **Tanzima Islam**, Kathryn Mohror, Saurabh Bagchi, Adam Moody, Bronis R De Supinski, and Rudolf Eigenmann. [McrEngine: A Scalable Checkpointing System Using Data-Aware Aggregation and Compression](#). In *Scientific Programming*, volume 21, pages 149–163. Hindawi, 2013.
- [44] Chris Daw, Brian Barragan Cruz, Nicholas Majeske, Filip Jagodzinski, **Tanzima Islam**, and Brian Hutchinson. Low rank approximation methods for identifying impactful pairwise protein mutations. In *Algorithms and Methods in Structural Bioinformatics*, pages 63–87. Springer, 2012.
- \* [45] **Tanzima Islam**, Kathryn Mohror, Saurabh Bagchi, Adam Moody, Bronis R. de Supinski, and Rudolf Eigenmann. [McrEngine: A Scalable Checkpointing System Using Data-aware Aggregation and Compression](#). In *International Conference on High Performance Computing, Networking, Storage and Analysis (SC)*, pages 17:1–17:11, 2012. Acceptance rate: 20%. **Best student paper finalist**.
- \* [46] **Tanzima Islam**, Saurabh Bagchi, and Rudolf Eigenmann. [FALCON: A System for Reliable Checkpoint Recovery in Shared Grid Environments](#). In *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis (SC)*, pages 1–12, Nov 2009. **Best student paper finalist**.
- Policy Reports, Book Chapters, Articles**
- [47] Lois Curfman McInnes, Dorian Arnold, Prasanna Balaprakash, Mike Bernhardt, Beth Cerny, Anshu Dubey, Roscoe Giles, Denice Ward Hood, Mary Ann Leung, Vanessa Lopez-Marrero, Paul Messina, Olivia B. Newton, Chris Oehmen, Stefan M. Wild, Jim Willenbring, Lou Woodley, Tony Baylis, David E. Bernholdt, Chris Camano, Johannah Cohoon, Charles Ferenbaugh, Stephen M. Fiore, Sandra Gesing, Diego Gomez-Zara, James Howison, **Tanzima Islam**, David Kepczynski, Charles Lively, Harshitha Menon, Bronson Messer, Marieme Ngom, Umesh Paliath, Michael E. Papka, Irene Qualters, Elaine M. Raybourn, Katherine Riley, Paulina Rodriguez, Damian Rouson, Michelle Schwalbe, Sudip K. Seal, Ozge Surer, Valerie Taylor, and Lingfei Wu. [Report of the 2025 workshop on next-generation ecosystems for scientific computing: Harnessing community, software, and ai for cross-disciplinary team science](#), 2025.
- [48] [SIAM Task Force Anticipates Future Directions of Computational Science](#). *Society for Industrial and Applied Mathematics*, June 2024.
- [49] Chris Daw, Brian Barragan Cruz, Nicholas Majeske, Filip Jagodzinski, **Tanzima Islam**, and Brian Hutchinson. Low rank approximation methods for identifying impactful pairwise protein mutations. In Nurit Haspel, Filip Jagodzinski, and Kevin Molloy, editors, *Algorithms and Methods in Structural Bioinformatics*, pages 63–87. Springer Nature, 2022.

- [50] Ritu Arora, Xiaosong Li, Bonnie Hurwitz, Daniel Fay, Dhableswar K. Panda, Edward Valeev, Shaowen Wang, Shirley Moore, Sunita Chandrasekaran, Ting Cao, Holly Bik, Matthew Curry, and **Tanzima Islam**. [Future directions of the cyberinfrastructure for sustained scientific innovation \(cssi\) program](#), 2020. [Ph.D. Dissertation](#)
- [51] **Tanzima Islam**. *Reliable and scalable checkpointing systems for distributed computing environments*. PhD thesis, Purdue University, West Lafayette, IN, May 2013.

---

## Selected Open-Source Software & Dataset Releases

- GPTune** Integrated the interpretable machine learning algorithm for identifying important performance-influencing factors—developed as part of the Dashing framework—into GPTune, an auto-tuner developed at Lawrence Berkeley National Laboratory.
- SCR** Integrated a scalable data compression method—mcrEngine—into Scalable Checkpoint/Restart library (SCR) for MPI applications in collaboration with LLNL. SCR will serve as the reliability backbone for LLNL's upcoming exascale system, El Capitan. <https://github.com/LLNL/scr>.
- libNVCD** Easy-to-use, performance measurement and analysis tool for NVIDIA based GPUs. <https://github.com/tzislam/libnvcd>.
- DASHING** Interpretable machine learning algorithms for performance characterization and comparison. <https://github.com/tzislam/Dashing>.
- Gyan** Performance measurement tool for MPI libraries in collaboration with LLNL.
- PyPerfDump** Hardware performance counter measurement tool for Python applications. <https://github.com/RECUP-DOE/pyperfdump#>.
- Dataset1** On-node scaling data on HPC systems. 2019. [DOI:10.5281/zenodo.4315003](https://doi.org/10.5281/zenodo.4315003).
- Dataset2** Performance characterization data for several DOE Exascale Computing Project (ECP) applications including Particle-In-Cell application WarpX and cosmological simulation code Nyx3D. 2020. [doi:10.5281/zenodo.3403037](https://doi.org/10.5281/zenodo.3403037).

---

## Selected Professional Activities

- Organizer** **Chair** → Student Program@SC'27, Posters@ISC'26, Performance Track@SC'26, Software Track@ICPP'26, Performance Track@IPDPS'25, Applications and Workflows Track@CCGrid'25, Performance Track@SC'22, Posters@HPDC'25; **Organizer** → IEEE International Workshop on Big Data Computation, Analysis & Applications (in conjunction with IEEE COMPSAC'19-'23); **Vice/Co-Chair** → Posters@ISC'25, Performance Track@HiPC'24, Performance Track@IEEE Cluster'23, Performance Track@ICPP'19, Awards@SC'23.
- Subject-Matter Editor** Future Generation Computer Systems (FGCS), Parallel Computing (ParCo).
- Technical Review Board** IEEE Transactions on Parallel and Distributed Systems (TPDS).
- Selected Technical Program Committees** SIGKDD'26; High-Performance Distributed Computing (HPDC)'22–Present; International Parallel & Distributed Processing Symposium (IPDPS)'17–Present; International Conference for High Performance Computing, Networking, Storage and Analysis (SC)'17–Present (Performance track); SC'19–'21 (Posters); SC'18 (HPC for Undergraduates); NeurIPS'22; International Supercomputing Conference (ISC)'19–Present; Women in HPC Technical Conference'20; EuroMPI User Forum'20; The Platform for Advanced Scientific Computing (PASC)'19; International Conference on Parallel Processing (ICPP)'18; International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD)'13–'14; Symposium on Principles and Practice of Parallel Programming (PPoPP)'11; Dependable Systems and Networks (DSN)'10, '22.
- Grant Review Committees** DOE SBIR; DOE SciDAC; NSF CISE/CSSI; NSF HBCU-UP; NSF CISE/OAC; NSF CISE/CCF.
- Selected Journal Reviewing** IEEE Transactions on Parallel and Distributed Systems (TPDS)'19–Present; Parallel Computing (ParCo)'19–'21; Journal of Grid Computing (JoGC)'17–'21; International Journal of High Performance Computing Applications (IJHPCA)'18.
- Memberships** [IEEE](#), [ACM](#).

---

## Invited Talks and Panels

- June and August, 2025 **Multiple Talks**, *From Reactive Debugging to Proactive Detection: AI/ML for Performance-Aware Software Development*, EuroPar (Germany), Monterey Data Conference (Invite-only DOE Workshop in California), PASC (Switzerland).
- September, 2024 **CCDSC**, *Reimagine Performance Modeling: Challenges and Opportunities with Generative AI*, Lyon, France, Personally invited by **Turing Award winner Dr. Jack Dongarra**.
- Aug, 2024 **Coral-2 Meeting at LLNL**, *Augmenting AMD's Omnipar with ML-based In-depth Influence Analysis for Explaining Performance Gap*, Virtual.
- 2024, 2025 **DOE ASCR PI Meeting**, *Intelytics: An Efficient Data-Driven Decision-Making Engine for Performance In the Era of Heterogeneity*, Atlanta and Texas, respectively.
- September, 2022 **Bangladesh University of Engineering and Technology (BUET)**, *How to apply to Graduate School: Things I wish I knew*, Virtual.
- May, 2022 **How to be a Great Mentor**, *Exascale Computing Project (ECP)*, Virtual.
- April, 2022 **Scalability challenges and opportunities for I/O bound applications**, *CHEOPS Workshop at EuroSys*, Virtual, <https://tinyurl.com/2xyb2wnr>.
- March, 2022 **Exascale Computing Project (ECP): Enabling Next-Generation of Hardware-Software Co-design using Data Science**, *LLNL*, Virtual.
- February, 2022 **Sustainable Horizons Institute, SHI**, Virtual.
- Dec, 2021 **Characterizing Performance of Workflow Applications**, *Brookhaven National Laboratory*, Virtual.
- Dec, 2021 **Center of Excellence: Performance Characterization of Deep Learning Workloads**, *AMD*, Virtual.
- Sep, 2021 **Enabling Next-Generation Software and Co-Design**, *AMD at Texas State*, San Marcos, TX.
- Sep, 2021 **HPC Enables Digitalization Panel**, *Digital 360 Summit*, San Marcos, TX.
- Nov, 2020 **Careers in HPC Panel, Taking the Leap: Changing Careers**, *International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, Virtual.
- Oct, 2020 **May the HPC Force be with you!**, *Organized a panel for undergraduate students*, Richard Tapia Celebration of Diversity in Computing.
- Aug, 2020 **Data-Driven Performance Modeling for HPC Co-design**, *Invited scholar's talk*, AMD.
- Aug, 2020 **DASHING: An Interpretable Machine Learning Toolkit for Performance Analysis and Visualization**, *Exascale Computing Project Hackathon*, Oak Ridge National Laboratory.
- May, 2019 **The NIMBioS Workshop on Scientific Collaboration Enabled by High Performance Computing, Scalable I/O Performance for Scientific Applications—Challenges and Potentials**, Knoxville, TN.
- December, 2018 **Lawrence Berkeley National Laboratory**, *Understanding the Performance Portability Challenges and Opportunities using Machine Learning*, Berkeley, CA.
- November, 2018 **HPC for Undergraduates**, *International Conference for High Performance Computing, Networking, Storage and Analysis*, Dallas, TX.
- July, 2018 **The Platform for Advanced Scientific Computing Conference**, *Machine Learning Framework for Performance Coverage Analysis of Proxy Applications*, ACM and the Swiss National Supercomputing Center (CSC), Basel, Switzerland.
- February, 2015 **JOWOG-34**, *Proxy Application Validation using Machine Learning (Veritas)*, Sandia National Laboratories, Albuquerque, NM.
- January, 2014 **Spellman College**, *Opportunities after Graduate School*, Atlanta, GA.

---

## Outreach and Mentoring Students

- 2014-Present **Co-founder**, *Bangladeshi Women in Computer Science and Engineering*, Dhaka, Bangladesh, First research, networking, and mentoring platform for female Computer Science and Engineering students of Bangladesh. Mentored 100s of female students in their undergraduate journey.
- 2017–2019 **Research advisor**, *Undergraduate (7) and graduate (3) students (Western Washington University); Undergraduate (10) and graduate (8) students, Postdoc (1) (Texas State University)*, High school (2); K-5 Robotics coach (4).
- 2017-Present **Faculty advisor**, *ACM-W chapter at Western Washington University*, ACM-ICPC Programming Club, Faculty mentor.

- 2014–2016 **Intern supervisor**, *Undergraduate, M.S. and Ph.D. students from University of Hamburg, University of Illinois Urbana-Champaign, Spelman College, University of California San Diego, NSF REU Advisor, Marshalls College, University of California Berkeley.*
- 2019–Present **Mentor**, *Graduate students (10)*, SC Student Volunteer Program.
- 2011–2013 **Student mentor**, *Undergrad, M.S., Ph. D. students*, Purdue University.