

Tanzima Z. Islam

Assistant Professor
Texas State University

601 University Drive
San Marcos, TX, 78666
✉ tanzima@txstate.edu
🌐 www.tanzimaislam.com

Education

- 2013 **Ph.D., 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

- 2020-Present **Assistant Professor**, *Department of Computer Science*, Texas State University (TxState).
- 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

- 2024 **Research Millionaire Award** at Texas State University.
- 2023 **Presidential Seminar Award for Excellence in Scholarly/Creative Activities** at Texas State University.
- 2023 **Presidential Distinction for Excellence in Scholarly/Creative Activities**, College of Science and Engineering (CoSE) at Texas State University. This is the highest honor achievable by a faculty in the CoSE.
- 2023 **Valero Award for Excellence in Scholarly/Creative Activities**, Academic Affairs Council at Texas State University.
- 2021 **Dean's College Achievement Award for Excellence in Scholarly/Creative Activities** at Texas State University.
- 2019-2020 **R&D 100 award**, in collaboration with Lawrence Livermore National Laboratory for the Scalable Checkpoint/Restart Framework 2.0 (SCR).
- 2014 **LLNL Director's Science & Technology Award**.
- 2016, 2015, 2014 Best Poster Award, Lawrence Livermore National Laboratory Annual Scholars Poster Symposium.
- 2014 *2nd* Place Winner, LLNL Computation Directorate Postdoctoral Poster Symposium.
- 2012, 2009 Best Student Paper Nominations, International Conference for High Performance Computing, Networking, Storage and Analysis (SC).
- 2010 *2nd* Place Winner, ACM Student Research Competition, Grace Hopper Celebration of Women in Computing.

Research and Other Funding

Secured a total of \$1.974M in external research grants as Principal Investigator (PI) or Co-Principal Investigator (Co-PI) from agencies including the Department of Energy (DOE), Lawrence Livermore National Laboratory (LLNL), and AMD since 2016, with \$1.722M awarded at Texas State University (TXST) since joining in 2020.

- 2024 **PI**, Application Fingerprinting, Department of Energy's STEP: Software Tools Ecosystem Project, \$16K.
- 2023 **PI**, SRP-HPC: Performance Analysis using Machine Learning: Use Case Chimbuko, DOE SRP Fellowship, \$85K.
- 2021-Present **PI**, PerfROCM: Cross-platform performance prediction and analysis using deep learning, AMD Research Gift, \$175K.
- 2022-2027 **PI**, INTELTYICS: An Efficient Data-Driven Decision-Making Engine for Performance In the Era of Heterogeneity, DOE Early Career Research Program (ECRP), \$770K.
- 2022-2025 **Co-PI**, RECUP: Scalable Metadata And Provenance Services for Reproducible Hybrid Workflows, Department of Energy (DOE), \$2.43M (TxState PI: 300K).
- 2022-2025 **Sub-contract**, ICE4HPC: Towards the Intelligent Center for HPC, Lawrence Livermore National Laboratory, \$750K, (TxState PI: 300K).
- 2022 **PI**, SRP-HPC: Performance Characterization of Workflow Applications, DOE SRP Fellowship, \$68K.
- 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, \$40K.
- 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.

Publications

Peer Reviewed Papers

- [1] Chase Phelps, Tanzima Z. Islam, Tapasya Patki, and Line C. Pouchard. A scalable graph-based influence analysis framework for studying performance reproducibility using machine learning. Journal submission in preparation (Est. submission: Dec. 2024).
- [2] Tarek Ramadan, Nathan Pinnow, Tanzima Z. Islam, Jayaraman J. Thiagarajan, and Chase Phelps. Structure-aware representation learning for effective performance prediction. Nov 2025. (Under minor revision).
- [3] Banooqa H. Bandy, Tanzima Z. 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. (Accepted). Acceptance rate: 24%.
- [4] Arunavo Dey, Aakash Dhakal, Yeom Jae-Seung Islam, Tanzima Z., 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. (Accepted). Acceptance rate: 24%.
- [5] Tanzima Z. Islam, Aniruddha Marathe, Holland Schutte, and Mohammad Zaeed. Data-driven analysis to

understand gpu hardware resource usage of optimizations, 2024. Journal submission in preparation. (Est. submission: Oct. 2024).

- [6] Christopher Kelly, Wei Xu, Line C. Pouchard, Hubertus Van Dam, Tanzima Z. 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. (Under minor revision).
- [7] Chase Phelps, Ankur Lahiry, Tanzima Z. 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. (Accepted). Acceptance rate: 24%.
- [8] Mohammad Zaeed, Tanzima Z. 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. (Accepted). Acceptance rate: 24%.
- [9] Amal Gueroudji, Chase Phelps, Tanzima Z. 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).
- [10] 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%.
- [11] Bogdan Nicolae, Tanzima Z. 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.
- [12] 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%.
- [13] Tarek Ramadan, Ankur Lahiry, and Tanzima Z. Islam. Novel representation learning technique using graphs for performance analytics. In *2023 International Conference on Machine Learning and Applications (ICMLA)*, pages 1311–1318, 2023.
- [14] 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.
- [15] 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%.
- [16] **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%.
- [17] **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.
- [18] 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.
- [19] 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%.
- [20] 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.
- [21] 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%.

- [22] **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.
 - [23] 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.
 - [24] 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%.
 - [25] 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%.
 - [26] **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.
 - [27] 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%.
 - [28] **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%.
 - [29] 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%.
 - [30] **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.
 - [31] 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.
 - [32] 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.
 - [33] **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.
 - [34] **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.
 - [35] 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.
 - * [36] **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.**
 - * [37] **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 nomination.**
- Ph.D. Dissertation**
- [38] **Tanzima Islam**. *Reliable and scalable checkpointing systems for distributed computing environments*. PhD thesis, Purdue University, West Lafayette, IN, May 2013.

Professional Activities

Organizer	Track Chair, IPDPS'25, SC'22, IEEE International Workshop on Big Data Computation, Analysis & Applications in conjunction with IEEE COMPSAC'19-Present; Track Vice/Co-Chair, HiPC'24, IEEE Cluster'23, International Conference on Parallel Processing'19. Award Chair, SC'23.
Editorial Board	Future Generation Computer Systems (FGCS).
Technical Review Board	IEEE Transactions on Parallel and Distributed Systems (TPDS).
Technical Program Committees	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—(Performance track). SC'19-'21 (Posters). SC'18 (HPC for Undergrad), NeurIPS'22, International Conference in Supercomputing (ISC)'19–, Women in HPC Technical Conference'20, Euro/MPI 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 (SBAD-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, NSF CISE/CSSI, NSF HBCU-UP, NSF CISE/OAC, NSF CISE/CCF, DOE SciDAC.
Journal Reviewing	IEEE Transactions on Parallel and Distributed Systems (TPDS)'19–, International Conference on Parallel Computing (ParCo)'19-'21, Journal of Grid Computing (JoGC)'17–'21, International Journal of High Performance Computing (IJHPCA)'18.
Memberships	IEEE, ACM.

Invited Talks and Panels

September, 2024	CCDSC , <i>Reimagine Performance Modeling: Challenges and Opportunities with Generative AI</i> , Lyon, France.
Aug, 2024	Coral-2 Meeting at LLNL , <i>Augmenting AMD's Omniperf with ML-based In-depth Influence Analysis for Explaining Performance Gap</i> , Virtual.
February, 2024	DOE ASCR PI Meeting , <i>Intelytics: An Efficient Data-Driven Decision-Making Engine for Performance In the Era of Heterogeneity</i> , Atlanta.
September, 2022	Bangladesh University of Engineering and Technology (BUET) , <i>How to apply to Graduate School: Things I wish I knew</i> , Virtual.
May, 2022	How to be a Great Mentor , <i>Exascale Computing Project (ECP)</i> , Virtual.
April, 2022	Scalability challenges and opportunities for I/O bound applications , <i>CHEOPS Workshop at EuroSys</i> , Virtual, https://tinyurl.com/2xyb2wnr .
March, 2022	Exascale Computing Project (ECP): Enabling Next-Generation of Hardware-Software Co-design using Data Science , <i>LLNL</i> , Virtual.
February, 2022	Sustainable Horizons Institute , <i>SHI</i> , Virtual.
Dec, 2021	Characterizing Performance of Workflow Applications , <i>Brookhaven National Laboratory</i> , Virtual.
Dec, 2021	Center of Excellence: Performance Characterization of Deep Learning Workloads , <i>AMD</i> , Virtual.
Sep, 2021	Enabling Next-Generation Software and Co-Design , <i>AMD at Texas State</i> , San Marcos, TX.
Sep, 2021	HPC Enables Digitalization Panel , <i>Digital 360 Summit</i> , San Marcos, TX.
Nov, 2020	Careers in HPC Panel, Taking the Leap: Changing Careers , <i>International Conference for High Performance Computing, Networking, Storage and Analysis (SC)</i> , Virtual.
Oct, 2020	May the HPC Force be with you! , <i>Organized a panel for undergraduate students</i> , Richard Tapia Celebration of Diversity in Computing.
Aug, 2020	Data-Driven Performance Modeling for HPC Co-design , <i>Invited scholar's talk</i> , AMD.
Aug, 2020	DASHING: An Interpretable Machine Learning Toolkit for Performance Analysis and Visualization , <i>Exascale Computing Project Hackathon</i> , Oak Ridge National Laboratory.
May, 2019	The NIMBioS Workshop on Scientific Collaboration Enabled by High Performance Computing , <i>Scalable I/O Performance for Scientific Applications—Challenges and Potentials</i> , Knoxville, TN.
December, 2018	Lawrence Berkeley National Laboratory , <i>Understanding the Performance Portability Challenges and Opportunities using Machine Learning</i> , 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.
- 2017–2019 **Research advisor**, *Undergraduate (7) and graduate (3) students (Western Washington University); Undergraduate (6) and graduate (7) students, Postdoc (1), Texas State University, High school (2)*.
- 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.