Sunita Chandrasekaran
David L. and Beverly J.C. Mills Career Development Chair
Director, First State AI Institute
Associate Professor, Dept. of Computer & Information Sciences
Computational Research and Programming Lab
417 FinTech Innovation Hub, Newark, DE 19713, USA
University of Delaware
Email: schandra@udel.edu

Professional Experience

- **Director**; First State AI Institute. *July 2025 – present*
- **Vice Chair**; State of Delaware AI Commission. *July 2024 – July 2027*
- **David L. and Beverly J.C. Mills Career Development Chair**; Department of Computer and Information Sciences (CIS), University of Delaware, Newark, USA. *September 2020 – August 2026*
- **Associate Professor**; Department of Computer and Information Sciences (CIS), University of Delaware, Newark, USA. *September 2021 – Present*
- **Directorate Advisory Committee Member (DAC)**; Oak Ridge National Laboratory's Computing and Computational Sciences Directorate (CCSD). *April 2025 – September 2028*
- **Board of Directors**; OpenACC Organization. *Spring 2020 – Present*
- **User Representative Chair**; OpenACC Organization. *June 2017 – Present*
- **ACM Senior Member**. *February 2022 – Present*
- **IEEE Senior Member**. *February 2021 – Present*

Professional Experience (Past)

- **DOE ASCR Advisory Committee Member**; DOE Office of Science Advanced Scientific Computing Advisory Committee (ASCAC). *May 2024 – October 2025*
- **Co-Director**; AI Center of Excellence (AICoE), University of Delaware, Newark, USA. *September 2022 – December 2025*
- **Computational Scientist**; Brookhaven National Laboratory, Upton, NY 11973-5000. *September 2021 – March 2024*
- **Scientific Advisory Board Member**; EuroHPC Center of Excellence Plasma-PEPSC. *August 2023 – August 2026*
- **Global Professorship Program**; Computer Architecture and Parallel Systems, TU Munich, Munich, Germany. *March 2024*

- **Assistant Professor**; Department of Computer and Information Sciences (CIS), University of Delaware, Newark, USA. *September 2015 – August 2021*
- **Postdoctoral Researcher**; Department of Computer Science, University of Houston, Houston, TX, USA. *December 2010 – August 2015*

Education

- Nanyang Technological University (NTU), Singapore, Ph.D., School of Computer Science.
  Title: Tools and Algorithms for High Level Algorithm Mapping to FPGA, 2012
- Anna University, Chennai, India, Bachelor of Engineering, Electrical & Electronics.
  Final Year Thesis Title: Experimental Studies in Statistical Signal Processing, 2005

Awards and Honors

Named Positions

- David L. and Beverly J.C. Mills Career Development Chair, Department of Computer & Information Sciences, University of Delaware. *September 2020 – August 2026*

Society, Technical, and Research Awards

- SCW75 – Recognized by the annual program of Scientific Computing World (SCW) among those driving innovation at the intersection of computing and scientific research, May 2026
- Outstanding Delaware Woman in STEM, Million Women Mentors – Delaware, 2026
- ACM SIGHPC Emerging Woman Leader in Technical Computing, 2025
- Walter O. Roberts Scientific & Technical Advancement Award, for the implementation of MPS/University of Chicago Radiative MHD Models (MURaM) on GPUs, 2023
- ACM Senior Member, February 2022
- SPEC HPG Achievement (Contributor) Award, 2022
- IEEE Senior Member, February 2021
- IEEE TCHPC Award for Excellence for Early Career Researchers in High Performance Computing, 2016
- SPEC HPG Technical Leadership Award, 2016
- SPEC HPG Benchmark Project Leadership Recognition Award, 2014

Editorial Awards

- IEEE Transactions on Parallel and Distributed Systems (TPDS) Award for Editorial Excellence, 2022
- Outstanding Editor Award, Future Generation Computer Systems (FGCS), 2021

Teaching and Educational Recognition

- Google AI for Education Accelerator, Principal Investigator, 2026
- College of Engineering Excellence in Teaching Award, University of Delaware, 2020
- NVIDIA GPU Education Center Award to the University of Delaware, Principal Investigator, 2016

Compute Resources and Allocation

Principal Investigator on compute allocation awards spanning four DOE leadership-class systems at NERSC and OLCF — Summit, Perlmutter, Frontier, and the upcoming Doudna — obtained through, open call, peer-reviewed, highly competitive programs (ECP, CAAR, NESAP, ERCAP). These awards have supported real-world applications, programming-system deployment, and LLM-based software development across successive generations of GPU-accelerated (pre-)exascale architectures.

Publications

Highlights: Publications demonstrate strong collaborations with DOE national labs (ORNL, LLNL, BNL, LBNL, Argonne); medical partners (NCI/NIH, Nemours Children's Hospital, Frederick National Lab of Cancer Research/FNLCR); industry partners (NVIDIA, AMD, Intel, HPE/Cray); and international partners (HZDR, Germany).

Book & Book Chapters

Prior to 2021

- Sunita Chandrasekaran and Guido Juckeland. *OpenACC for Programmers: Concepts and Strategies.* Co-Edited Book published by Pearson Addison-Wesley Professional; 1st edition. ISBN-13: 978-0134694283, November 2017
- Sunita Chandrasekaran, Rengan Xu, and Barbara Chapman. "Using OpenACC for stencil and Feldkamp algorithms." Co-authored chapter in *Parallel Programming with OpenACC* (ed. Rob Farber). Morgan Kaufmann. ISBN-13: 978-0124103979, November 2016
- Barbara Chapman, Deepak Eachempati, and Sunita Chandrasekaran. "Chapter on OpenMP." Co-authored chapter in *Programming Models for Parallel Computing* (ed. Pavan Balaji). MIT Press. ISBN-13: 978-0262528818, 2015

Refereed Journals

- Daria Blinova, †Gayathri Emuru, Rakesh Emuru, †Kushagradheer Shridheer Srivastava, Mina Rulis, Sunita Chandrasekaran (co-author), and Benjamin Bagozzi. "Linked Multi-Model Data on Russian Domestic and Foreign Policy Speeches." *UNDER REVIEW.* Scientific Data, Nature.
  - Daria Blinova, †Gayathri Emuru, Rakesh Emuru, †Kushagradheer Shridheer Srivastava, Mina Rulis, Sunita Chandrasekaran, and Benjamin Bagozzi. 2026. "Linked Multi-Model Data on Russian Domestic and Foreign Policy Speeches." Harvard Dataverse. https://doi.org/10.7910/DVN/SGI0VK. Note: Harvard Dataverse is an open-source data repository for sharing and accessing research data.
- Qile Wang, Abhishek Karwankar, Shukria Muhammadi, †Srikanth K. Perumalla, §Abhin Suresh, Sunita Chandrasekaran (co-author), Kenneth E. Barner, Matthew L. Mauriello. "AI4OSN: A Usability Study of a Human-Centered Agentic AI Platform for Online Social Network Research." *UNDER REVIEW.* ACM Transactions on Social Computing (2026).
- Arthi Jayaraman and Sunita Chandrasekaran (co-author). "Interdisciplinary Polymer and Data Science Graduate Training by Bringing Industry Relevant Problems into the Classroom." *Macromolecules* **Vol. 58**(16), 8569–8571 (2025). *(Impact Factor: 5.1)* DOI: 10.1021/acs.macromol.5c02109

- ‡Christian Munley, *Aaron Jarmusch, and Sunita Chandrasekaran. "LLM4VV: Developing LLM-driven testsuite for compiler validation." *Future Generation Computer Systems*, **Vol. 160**, 1–13 (2024). *(Impact Factor: 6.2)* DOI: 10.1016/j.future.2024.05.034
  **** Spring 2026 Editor's Choice Paper**
- *Vineeth Gutta, Satish Ranganathan Ganakammal, Sara Jones, Matthew Beyers, and Sunita Chandrasekaran. "UNNT: A novel Utility for comparing Neural Net and Tree-based models." *PLOS Computational Biology* **Vol. 20**(4), e1011504 (2024). *(Impact Factor: 3.8)* DOI: 10.1371/journal.pcbi.1011504
- Klaus Steiniger, Rene Widera, Sergei Bastrakov, Michael Bussmann, Sunita Chandrasekaran (co-author), Benjamin Hernandez, ‡Kristina Holsapple, Axel Huebl, Guido Juckeland, Jeffrey Kelling, *Matt Leinhauser, Richard Pausch, David Rogers, Ulrich Schramm, Jeff Young, and Alexander Debus. "EZ: An efficient, charge conserving current deposition algorithm for electromagnetic particle-in-cell simulations." *Computer Physics Communications*, **Vol. 291**, 108849 (2023). *(Impact Factor: 7.2)*
  DOI: 10.1016/j.cpc.2023.108849
- *Mauricio H. Ferrato, Adam G. Marsh, Karl R. Franke, Benjamin J. Huang, E. Anders Kolb, Deborah DeRyckere, Douglas K. Graham, Sunita Chandrasekaran, and Erin L. Crowgey. "Machine learning classifier approaches for predicting response to RTK-Type-III Inhibitors demonstrates high accuracy using transcriptomic signatures and ex vivo data." *Bioinformatics Advances* **Vol. 3**(1), vbad034 (2023). *(Impact Factor: 2.4)* DOI: 10.1093/bioadv/vbad034
- ‡Matt Stack, Paul Macklin, *Robert Searles, and Sunita Chandrasekaran. "OpenACC Acceleration of an Agent-Based Biological Simulation Framework." *IEEE Computing in Science and Engineering* **Vol. 25**(3), 53–63 (2023). *(Impact Factor: 2.1)* DOI: 10.1109/MCSE.2022.3226602
  Note: At the time of this work, Matt Stack was an undergraduate in my lab and Robert Searles was one of my PhD students. Robert earned his doctorate in 2019, followed by Matt's graduation in 2020; both have since joined NVIDIA as solution architects.
- *Matt Leinhauser, Rene Widera, Sergei Bastrakov, Alex Debus, Michael Bussmann, and Sunita Chandrasekaran. "Metrics and Design of an Instruction Roofline Model for AMD GPUs." *ACM Transactions on Parallel Computing* **Vol. 9**(1), 1–14 (2022). *(Impact Factor: 2.0)* DOI: 10.1145/3505285
- *Eric Wright, ‡Mauricio Ferrato, Alex Bryer, *Robert Searles, Juan Perilla, and Sunita Chandrasekaran. "Accelerating Prediction of Chemical Shift of Protein Structures on GPUs." *PLOS Computational Biology* **Vol. 16**(5), e1007877 (2020). *(Impact Factor: 3.8)* DOI: 10.1371/journal.pcbi.1007877
- *Robert Searles, Sunita Chandrasekaran, Wayne Joubert, and Oscar Hernandez. "MPI+OpenACC: Accelerating Radiation Transport Mini-application, Minisweep, on Heterogeneous Systems." *Computer Physics Communications* **Vol. 236**, 176–187 (2019). *(Impact Factor: 7.2)* DOI: 10.1016/j.cpc.2018.10.007
- *Millad Ghane, Sunita Chandrasekaran, and Margaret S. Cheung. "Pointerchain: Tracing pointers to their roots - A case study in molecular dynamics simulations." *Parallel Computing* **Vol. 85**, 190–203 (2019). *(Impact Factor: 2.0)* DOI: 10.1016/j.parco.2019.04.007
- †Jose Manuel Monsalve Diaz, ‡Kyle Friedline, Swaroop Pophale, Oscar Hernandez, David Bernholdt, and Sunita Chandrasekaran. "Analysis of OpenMP 4.5 Offloading in Implementations: Correctness and Overhead." *Parallel Computing* **Vol. 89**, 102546 (2019). *(Impact Factor: 2.0)* DOI: 10.1016/j.parco.2019.102546
- Michael Wolfe, Jungwon Kim, Xiaonan Tian, *Rengan Xu, Barbara Chapman, and Sunita Chandrasekaran. "The OpenACC Data Model: Preliminary Study on Its Major Challenges and Implementations." *Parallel*

Computing **Vol. 78**, 15–27 (2018). *(Impact Factor: 2.0)* DOI: 10.1016/j.parco.2018.07.003

- Sunita Chandrasekaran, Guido Juckeland, Meifeng Lin, et al. "Best Practices in Running Collaborative GPU Hackathons." *IEEE Computing in Science and Engineering* **Vol. 20**(4), 95–106 (2018). *(Impact Factor: 2.1)* DOI: 10.1109/MCSE.2018.042781332
- *Robert Searles, *Stephen Herbein, Travis Johnston, Michela Taufer, and Sunita Chandrasekaran. "Creating a Portable, High-Level Graph Analytics Framework for Compute and Data-Intensive Applications." *International Journal of High Performance Computing and Networking* **Vol. 13**(1), 105–118 (2017). *(Impact Factor: 1.3)* DOI: 10.1504/IJHPCN.2017.10007922

  At the time of publication of the following 2 journal papers, I was a postdoc at the University of Houston mentoring PhD students Wang, Xu, Sun, and Tian. Chapman was their PhD advisor. My role in these projects was to design experiments and to test the compiler transformation for the nested loops in OpenUH compiler.
- *Xiaonan Tian, Rengan Xu, Yonghong Yan, Sunita Chandrasekaran, Deepak Eachempati, and Barbara Chapman. "Compiler Transformation of Nested Loops for General Purpose GPUs." *Concurrency and Computation: Practice and Experience* **Vol. 28**(2), 537–556 (2016). *(Impact Factor: 1.5)* DOI: 10.1002/cpe.3648
- *Rengan Xu, Sunita Chandrasekaran, and Barbara Chapman. "Multi-GPU Support on Shared Memory System using Directive-based Programming Model." *Scientific Programming* **Vol. 2015**, Article ID 621730 (2015). *(Impact Factor: 1.1)* DOI: 10.1155/2015/621730
- Sunita Chandrasekaran, Shilpa Shanbagh, Ramkumar Jayaraman, HuiYan Cheah, and Douglas Maskell. "C2FPGA: A Dependency-Timing Graph Design Methodology." *Journal of Parallel and Distributed Computing* **Vol. 73**(10), 1417–1429 (2013). *(Impact Factor: 3.4)* DOI: 10.1016/j.jpdc.2012.09.001

Refereed Conferences

- *Aaron Jarmusch and Sunita Chandrasekaran. "Microbenchmark-Driven Analytical Performance Modeling Across Modern GPU Architectures." *UNDER REVIEW.* ArXiv: https://arxiv.org/abs/2605.04178
- *Aaron Jarmusch, ‡Connor Vitz, and Sunita Chandrasekaran. "Execution-Centric Characterization of FP8 Matrix Cores, Asynchronous Execution, and Structured Sparsity on AMD MI300A." *UNDER REVIEW.* ArXiv: https://arxiv.org/abs/2602.10262
- *Aaron Jarmusch and Sunita Chandrasekaran. "Microbenchmarking NVIDIA's Blackwell Architecture: An in-depth Architectural Analysis." In *Proceedings of the 2026 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, Milan, Italy, 2026, pp. TBD. DOI: 10.1109/IPDPS2026.XXXXX *(CORE Rank: A)*. ArXiv: https://arxiv.org/abs/2512.02189
- ‡Zach Sollenberger, ‡Rahul Patel, ‡Saieda Ali Zada, and Sunita Chandrasekaran. "LLM4VV: Evaluating Cutting-Edge LLMs for Generation and Evaluation of Directive-Based Parallel Programming Model Compiler Tests." In *Proceedings of the 2025 IEEE 32nd International Conference on High Performance Computing, Data, and Analytics (HiPC)*, Hyderabad, India, 2025, pp. 333–342. DOI: 10.1109/HiPC66333.2025.00041 *(CORE Rank: A)*
- Jeffrey Kelling, Vicente Bolea, Michael Bussmann, Ankush Checkervarty, Alexander Debus, Jan Ebert,

Greg Eisenhauer, *Vineeth Gutta, Stefan Kesselheim, Scott Klasky, Vedhas Pandit, Richard Pausch, Norbert Podhorszki, Franz Pöschel, David Rogers, Jeyhun Rustamov, Steve Schmerler, Ulrich Schramm, Klaus Steiniger, René Widera, Anna Willmann, and Sunita Chandrasekaran. "The Artificial Scientist: In-Transit Machine Learning of Plasma Simulations." In *Proceedings of the 2025 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, Milano, Italy, 2025, pp. 407–418.
DOI: 10.1109/IPDPS64566.2025.00043. **(** Best Paper Nomination)** *(CORE Rank: A)*

- Greg Eisenhauer, Norbert Podhorszki, Ana Gainaru, Scott Klasky, Philip E. Davis, Manish Parashar, Matthew Wolf, Eric Suchtya, Erick Fredj, Vicente Bolea, Franz Pöschel, Klaus Steiniger, Michael Bussmann, Richard Pausch, and Sunita Chandrasekaran. "Streaming Data in HPC Workflows Using ADIOS." In *Proceedings of the 2024 Cray User Group (CUG)*, Perth, Australia, 2024, pp. 31–43. cug.org: pap116.pdf
Note: CUG is a specialized venue to discuss results gathered on Cray/HPE systems, including ORNL's flagship supercomputer Frontier. Hence there is no CORE rating.

- Scott Atchley, Chris Zimmer, John R. Lange, David E. Bernholdt, Verónica G. Melesse Vergara, Thomas Beck, Michael J. Brim, Reuben Budiardja, Sunita Chandrasekaran (co-author), Markus Eisenbach, Thomas Evans, Matthew Ezell, Nicholas Frontiere, Antigoni Georgiadou, Joe Glenski, Philipp Grete, Steven Hamilton, John Holmen, Axel Huebl, Daniel Jacobson, Wayne Joubert, Kim McMahon, Elia Merzari, Stan G. Moore, Andrew Myers, Stephen Nichols, and Sarp Oral. "Frontier: Exploring Exascale." In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC23)*, 2023, pp. 1–16. DOI: 10.1145/3581784.3607089. **(** Best Paper Nomination)** *(CORE Rank: A*)*
Note: I am particularly proud of this paper as we are a set of collaborators from Oak Ridge National Lab, Argonne National Laboratory, Hewlett Packard Enterprise (HPE), Universität Hamburg (Germany), Lawrence Berkeley National Laboratory, Pennsylvania State University, and Sandia National Laboratory who were among the first users and developers of Frontier, ranked 1st in the Top500 supercomputer list in 2023. My particular contribution stems from being PI of 2 critical projects: the Center for Accelerated Application Readiness (CAAR) program at ORNL and the Exascale Computing Project (ECP) Scaling OpenMP with LLVM for Exascale Performance and Portability (SOLLVE).

- Felix Meyer, Benjamin Hernandez, Richard Pausch, René Widera, David Groß, Sergei Bastrakov, Axel Huebl, Guido Juckeland, Jeffrey Kelling, *Matt Leinhauser, David Rogers, Ulrich Schramm, Klaus Steiniger, Stefan Gumhold, Jeff Young, Michael Bussmann, Sunita Chandrasekaran, and Alexander Debus. "Hardware-Agnostic Interactive Exascale In Situ Visualization of Particle-In-Cell Simulations." In *Proceedings of the Platform for Advanced Scientific Computing Conference (PASC)*, 2023, pp. 1–14. DOI: 10.1145/3592979.3593415 *(CORE Rank: B)*

- *Eric Wright, Johannes Doerfert, Shilei Tian, Barbara Chapman, and Sunita Chandrasekaran. "Implementing OpenMP's SIMD Directive in LLVM's GPU Runtime." In *Proceedings of the 52nd International Conference on Parallel Processing (ICPP)*, 2023, pp. 173–182. DOI: 10.1145/3605573.3605613 *(CORE Rank: A)*

- Holger Brunst, Sunita Chandrasekaran, Florina Ciorba, Nick Hagerty, Robert Henschel, Guido Juckeland, Junjie Li, Veronica G. Melesse Vergara, Sandra Wienke, and ‡Miguel Zavala. "First Experiences in Performance Benchmarking with the New SPEChpc 2021 Suites." In *Proceedings of the 2022 22nd IEEE International Symposium on Cluster, Cloud and Internet Computing (CCGrid)*, 2022, pp. 675–684. DOI: 10.1109/CCGrid54584.2022.00077 *(CORE Rank: A)*

- *Eric Wright, Cena Miller, Damien Przybylski, Matthias Rempel, Shiquan Su, Supreeth Suresh, Rich

Loft, and Sunita Chandrasekaran. "Refactoring the MPS/University of Chicago Radiative MHD (MURaM) Model for GPU/CPU Performance Portability Using OpenACC Directives." In *Proceedings of the Platform for Advanced Scientific Computing Conference (PASC)*, 2021, pp. 1–12. DOI: 10.1145/3468267.3470576 *(CORE Rank: B)*

- ‡Joshua Hoke Davis, Tao Gao, Sunita Chandrasekaran (co-author), Heike Jagode, Anthony Danalis, Jack J. Dongarra, Pavan Balaji, and Michela Taufer. "Characterization of Power Usage and Performance in Data-Intensive Applications Using MapReduce over MPI." In *Proceedings of the International Conference on Parallel Computing (ParCO)*, 2019, pp. 287–298. DOI: 10.3233/APC200053 *(CORE Rank: B)*
- *Robert Searles, Sunita Chandrasekaran, Oscar Hernandez, and Wayne Joubert. "Abstractions and Directives for Adapting Wavefront Algorithms to Future Architectures." In *Proceedings of the 5th Platform for Advanced Scientific Computing Conference (PASC)*, 2018, pp. 1–10. DOI: 10.1145/3218176.3218228
- †Jose Monsalve Diaz, Swaroop Pophale, ‡Kyle Friedline, Oscar Hernandez, David E. Bernholdt, and Sunita Chandrasekaran. "Evaluating Support for OpenMP Offload Features." In *Proceedings of the 47th International Conference on Parallel Processing Companion (ICPP)*, 2018, pp. 1–10. DOI: 10.1145/3229710.3229717 *(CORE Rank: A)*
- *Millad Ghane, Sunita Chandrasekaran (co-author), *Robert Searles, Margaret S. Cheung, and Oscar Hernandez. "Path Forward for Softwarization to Tackle Evolving Hardware." In *SPIE*, Vol. 10652, 2018. DOI: 10.1117/12.2304813
  *Millad Ghane, Ph.D. student of Margaret S. Cheung.

  At the time of publication of the following 5 conference papers, I was a postdoc at the University of Houston mentoring PhD students Wang, Xu, Sun, and Tian. Chapman was their PhD advisor. My role was to design experiments and test compiler transformations for nested loops in the OpenUH compiler.
- *Cheng Wang, Sunita Chandrasekaran, and Barbara Chapman. "CusFFT: A High-Performance Sparse Fast Fourier Transform Algorithm on GPUs." In *Proceedings of the 30th IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, Chicago, IL, 2016, pp. 963–972. DOI: 10.1109/IPDPS.2016.95 *(CORE Rank: A)*
- *Rengan Xu, Sunita Chandrasekaran, and Barbara Chapman. "An Analytical Model-based Auto-tuning Framework for Locality-aware Loop Scheduling." In *International Supercomputing Conference (ISC)*, Frankfurt, Germany, 2016, pp. 3–20. DOI: 10.1007/978-3-319-41321-1_1 *(CORE Rank: B)*
- *Cheng Wang, Sunita Chandrasekaran, Barbara Chapman, and Jim Holt. "Portable Mapping of OpenMP to Multicore Embedded Systems Using MCA APIs." In *Proceedings of the 14th ACM SIGPLAN/SIGBED Conference on Languages, Compilers and Tools for Embedded Systems (LCTES)*, Seattle, WA, 2013, pp. 153–162. DOI: 10.1145/2465554.2465573 *(CORE Rank: B)*
- *Peng Sun, Sunita Chandrasekaran, and Barbara Chapman. "Deploying OpenMP Task Parallelism on Multicore Embedded Systems with MCA Task APIs." In *Proceedings of the IEEE High Performance Computing and Communications (HPCC)*, 2015, pp. 843–847. DOI: 10.1109/HPCC-CSS-ICESS.2015.88 *(CORE Rank: B)*
- *Xiaonan Tian, *Rengan Xu, Yonghong Yan, Zhifeng Yun, Sunita Chandrasekaran, and Barbara Chapman. "Compiling A High-Level Directive-based Programming Model for Accelerators." In *Proceedings of the 26th International Workshop on Languages and Compilers for High Performance Computing (LCPC)*, San Jose, CA, 2013, pp. 105–120. DOI: 10.1007/978-3-319-09967-5_6 *(CORE Rank: C)*

- Kevin A. Huck, Oscar Hernandez, Van Bui, Sunita Chandrasekaran, Barbara Chapman, Allen D. Malony, Lois Curfman McInnes, and Boyana Norris. "Capturing Performance Knowledge for Automated Analysis." In *Proceedings of the IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, Austin, TX, 2008, pp. 1–10. DOI: 10.1145/1413370.1413418 *(CORE Rank: A*)*

Refereed Workshops

- ‡Zachary Sollenberger, ‡Jay Patel, ‡Christian Munley, *Aaron Jarmusch, and Sunita Chandrasekaran. "LLM4VV: Exploring LLM-as-a-Judge for Validation and Verification Testsuites." In *Proceedings of the 11th Workshop on Accelerator Programming Using Directives (WACCPD)* co-located with SC24, Atlanta, GA, USA, 2024, pp. 1885–1893. *(Acceptance Rate: 45%)* DOI: 10.1109/SCW63240.2024.00238
- *Aaron Jarmusch, §Felipe Cabarcas, Swaroop Pophale, ‡Andrew Kallai, Johannes Doerfert, Luke Peyralans, Seyong Lee, Joel Denny, and Sunita Chandrasekaran. "CI/CD Efforts for Validation, Verification and Benchmarking OpenMP Implementations." In *Advancing OpenMP for Future Accelerators: 20th International Workshop on OpenMP (IWOMP 2024)*, Perth, WA, Australia, 2024, *Lecture Notes in Computer Science (LNCS)*, Vol. 15195, pp. 111–125. *(Acceptance Rate: 87.5%)* DOI: 10.1007/978-3-031-72567-8_8
- Wael Elwasif, William Godoy, Nick Hagerty, J. Austin Harris, Oscar Hernandez, Balint Joo, Paul Kent, Damien Lebrun-Grandie, Elijah Maccarthy, Veronica Melesse Vergara, Bronson Messer, Ross Miller, Sarp Oral, Sergei Bastrakov, Michael Bussmann, Alexander Debus, Klaus Steiniger, Jan Stephan, René Widera, Spencer Bryngelson, Henry Le Berre, Anand Radhakrishnan, Jeffrey Young, Sunita Chandrasekaran (co-author), Florina Ciorba, Osman Simsek, Kate Clark, Filippo Spiga, Jeff Hammond, John E. Stone, David Hardy, Sebastian Keller, Jean-Guillaume Piccinali, and Christian Trott. "Application Experiences on a GPU-Accelerated Arm-based HPC Testbed." In *Proceedings of the HPC Asia 2023 Workshops (HPC Asia '23 Workshops)*, Singapore, 2023, pp. 35–49. *(CORE Rank: B)* DOI: 10.1145/3578345.3578363
- *Eric Wright, Cena Brown, Damien Przybylski, Matthias Rempel, Supreeth Suresh, and Sunita Chandrasekaran. "Analysis of MURaM, a Solar Physics Application, for Scalability, Performance and Portability." In *Proceedings of the 10th Workshop on Accelerator Programming Using Directives (WACCPD 2023)* co-located with SC23, Denver, CO, USA, 2023, pp. 1929–1938. *(Acceptance Rate: 44%)* DOI: 10.1145/3624062.3624606
- †Thomas Huber, Swaroop Pophale, ‡Nolan Baker, ‡Michael Carr, Nikhil Rao, ‡Jaydon Reap, ‡Kristina Holsapple, ‡Joshua Hoke Davis, Tobias Burnus, Seyong Lee, David E. Bernholdt, and Sunita Chandrasekaran. "ECP SOLLVE: Validation and Verification Testsuite Status Update and Compiler Insight for OpenMP." In *Proceedings of the 2022 IEEE/ACM International Workshop on Performance, Portability and Productivity in HPC (P3HPC)* co-located with SC22, Dallas, TX, USA, 2022, pp. 10–21. *(Acceptance Rate: 41%)* DOI: 10.1109/P3HPC56579.2022.00010
- *Aaron M. Jarmusch, ‡Alana Liu, ‡Christian Munley, ‡David Horta, †Vaishnavi Ravichandran, Joel Denny, and Sunita Chandrasekaran. "Analysis of Validating and Verifying OpenACC Compilers 3.0 and Above." In *Proceedings of the 9th Workshop on Accelerator Programming Using Directives (WACCPD)* co-located with SC22, Dallas, TX, USA, 2022, pp. 1–10. *(Acceptance Rate: 47%)* DOI: 10.1109/WACCPD56842.2022.00006

- Jeffrey Kelling, Sergei Bastrakov, Alexander Debus, Thomas Kluge, *Matt Leinhauser, Richard Pausch, Klaus Steiniger, Jan Stephan, Rene Widera, Jeffrey Young, Michael Bussmann, Sunita Chandrasekaran, and Guido Juckeland. "Challenges Porting a C++ Template-Metaprogramming Abstraction Layer to Directive-based Offloading." In *Proceedings of the 8th Workshop on Accelerator Programming Using Directives (WACCPD 2021)* co-located with SC21, St. Louis, MO, USA, 2021, pp. 46–55. *(Acceptance Rate: 46%)* DOI: 10.1109/WACCPD54688.2021.00010

Prior to 2021

- ‡Joshua Hoke Davis, Christopher Daley, Swaroop Pophale, †Thomas Huber, Sunita Chandrasekaran, and Nicholas Wright. "Performance Assessment of OpenMP Compilers Targeting NVIDIA V100 GPUs." In *Proceedings of the 7th Workshop on Accelerator Programming Using Directives (WACCPD)* co-located with SC20, 2020, pp. 1–10. *(Acceptance Rate: 43%)* DOI: 10.1109/WACCPD51911.2020.00006
- *Millad Ghane, Sunita Chandrasekaran, and Margaret S. Cheung. "Towards a portable hierarchical view of distributed shared memory systems: Challenges and Solutions." In *Proceedings of the 11th International Workshop on Programming Models and Applications for Multicores and Manycores (PMAM)* co-located with PPoPP '20, 2020, pp. 1–10. *(Acceptance Rate: 53%)* DOI: 10.1145/3380536.3380542
  *Millad Ghane, Ph.D. student of Margaret S. Cheung.
- *Millad Ghane, Sunita Chandrasekaran, and Margaret S. Cheung. "Gecko: Hierarchical Distributed View of Heterogeneous Shared Memory Architectures." In *Proceedings of the 10th International Workshop on Programming Models and Applications for Multicores and Manycores (PMAM)* co-located with PPoPP '19, 2019, pp. 21–30. *(Acceptance Rate: 58.8%)* DOI: 10.1145/3303084.3309489
  *Millad Ghane, Ph.D. student of Margaret S. Cheung.
- †Jose Manuel Monsalve Diaz, Swaroop Pophale, Oscar Hernandez, David E. Bernholdt, and Sunita Chandrasekaran. "OpenMP 4.5 Validation and Verification Suite for Device Offload." In *Proceedings of the 14th International Workshop on OpenMP (IWOMP 2018)*, *Lecture Notes in Computer Science (LNCS)*, Vol. 11128, 2018, pp. 82–95. DOI: 10.1007/978-3-319-98521-3_6
- ‡Kyle Friedline, Sunita Chandrasekaran, Graham Lopez, and Oscar Hernandez. "OpenACC 2.5 Validation Testsuite targeting multiple architectures." In *High Performance Computing: ISC 2017 Workshop on Performance Portable Programming Models for Accelerators (P3MA)*, *LNCS*, Vol. 10524, 2017, pp. 557–575. *(Acceptance Rate: 58.8%)* DOI: 10.1007/978-3-319-67630-2_41
- Sergio Pena, Sunita Chandrasekaran, and Lori Pollock. "Exploring translation of OpenMP to OpenACC 2.5: Lessons Learned." In *Proceedings of the 7th International Workshop on Accelerators and Hybrid Exascale Systems (AsHES)* co-located with IPDPS 2017, 2017, pp. 673–682. DOI: 10.1109/IPDPSW.2017.155
- Michael Wolfe, Seyong Lee, Jungwon Kim, *Xiaonan Tian, *Rengan Xu, Sunita Chandrasekaran, and Barbara Chapman. "Implementing the OpenACC Data Model." In *Proceedings of the 7th International Workshop on Accelerators and Hybrid Exascale Systems (AsHES)* co-located with IPDPS 2017, 2017, pp. 662–672. DOI: 10.1109/IPDPSW.2017.156
- *Robert Searles, *Stephen Herbein, and Sunita Chandrasekaran. "A portable, high-level graph analytics framework targeting distributed, heterogeneous systems." In *Proceedings of the 3rd Workshop on Accelerator Programming Using Directives (WACCPD)* co-located with SC16, 2016, pp. 79–88. *(Acceptance Rate: 40%)* DOI: 10.1109/WACCPD.2016.012

Refereed Conference and Workshop Publications before my tenure period (before Fall 2015):

- *Peng Sun, Sunita Chandrasekaran, and Barbara Chapman. "OpenMP-MCA: Leveraging Multiprocessor Embedded Systems using industry standards." In *Proceedings of the 2015 IEEE International Parallel & Distributed Processing Symposium Workshops (PLC)* co-located with IPDPS, Hyderabad, India, 2015, pp. 679–688. DOI: 10.1109/IPDPSW.2015.13
- Guido Juckeland, William Brantley, Sunita Chandrasekaran, et al. "SPEC ACCEL - A Standard Application Suite for Measuring Hardware Accelerator Performance." In *International Workshop on Performance Modeling, Benchmarking and Simulation (PMBS)* co-located with SC14, *LNCS*, Vol. 8966, 2014, pp. 46–67. *(Acceptance Rate: 26%)* DOI: 10.1007/978-3-319-17248-4_3
- *Rengan Xu, Maxime Hugues, Henri Calandra, Sunita Chandrasekaran, and Barbara Chapman. "Accelerating Kirchhoff Migration on GPU using Directives." In *Proceedings of the 1st Workshop on Accelerator Programming using Directives (WACCPD)* co-located with SC14, 2014, pp. 37–46. DOI: 10.1109/WACCPD.2014.10
- *Rengan Xu, *Cheng Wang, Sunita Chandrasekaran, and Barbara Chapman. "An OpenACC 1.0 Validation Suite." In *Proceedings of the 2014 IEEE Workshop on Multi-threaded Architectures and Applications (MTAAP)* co-located with IPDPS, 2014, pp. 1407–1416. DOI: 10.1109/IPDPSW.2014.155
- *Rengan Xu, *Xiaonan Tian, Yonghong Yan, Sunita Chandrasekaran, and Barbara Chapman. "Reduction Operations in Parallel Loops for GPGPUs." In *Proceedings of the 5th Workshop on Programming Models and Applications on Multicores and Manycores (PMAM)* co-located with PPoPP, 2014, pp. 10:10–10:20. DOI: 10.1145/2560683.2560693
- *Rengan Xu, *Xiaonan Tian, Sunita Chandrasekaran, Yonghong Yan, and Barbara Chapman. "NAS Parallel Benchmarks on GPGPUs using a Directive-based Programming Model." In *Proceedings of the 27th International Workshop on Languages and Compilers for Parallel Computing (LCPC)*, *LNCS*, Vol. 8967, 2014, pp. 67–81. DOI: 10.1007/978-3-319-17473-0_5
- *Cheng Wang, Mauricio Araya, Sunita Chandrasekaran, Barbara Chapman, and Detlef Hohl. "Parallel Sparse FFT." In *Proceedings of the 3rd Workshop on Irregular Applications: Architectures and Algorithms (IA3)* co-located with SC13, 2013, pp. 10:1–10:8. DOI: 10.1145/2535753.2535764
- Sayan Ghosh, Sunita Chandrasekaran, and Barbara Chapman. "Statistical Modeling of Power/Energy of Scientific Kernels on a Multi-GPU system." In *Proceedings of the 3rd International Workshop on Power Measurement and Profiling (PMP)* co-located with IGCC, 2013, pp. 1–6. DOI: 10.1109/IGCC.2013.6604513
- *Cheng Wang, Sunita Chandrasekaran, Barbara Chapman, and Jim Holt. "libEOMP: a portable OpenMP runtime library based on MCA APIs for embedded systems." In *Proceedings of the 4th Workshop on Programming Models and Applications for Multicores and Manycores (PMAM)* co-located with PPoPP, 2013, pp. 83–92. DOI: 10.1145/2442992.2443003
- *Cheng Wang, Sunita Chandrasekaran, and Barbara Chapman. "An OpenMP 3.1 Validation testsuite." In *OpenMP in the Petascale Era: 8th International Workshop on OpenMP (IWOMP)*, *LNCS*, Vol. 7312, 2012, pp. 237–249. DOI: 10.1007/978-3-642-30961-8_18
- *Rengan Xu, Sunita Chandrasekaran, Barbara Chapman, and Christoph F. Eick. "Directive-based Programming Models for Scientific Applications - A Comparison." In *Proceedings of the 2nd International Workshop on Domain-Specific Languages and High-Level Frameworks for High Performance Computing*

(WOLFHPC) co-located with SC12, 2012, pp. 1–9. DOI: 10.1109/WOLFHPC.2012.8

- Lei Huang, Eric Stotzer, Hangjun Yi, Barbara Chapman, and Sunita Chandrasekaran. "Parallelizing Ultrasound Image Processing using OpenMP on Multicore Embedded Systems." In *Proceedings of the 2012 IEEE Global High Tech Congress on Electronics (GHTCE)*, 2012, pp. 131–138. DOI: 10.1109/GHTCE.2012.6490139
- Sayan Ghosh, Sunita Chandrasekaran, and Barbara Chapman. "Energy Analysis of Parallel Scientific Kernels on Multiple GPUs." In *Proceedings of the 2012 Symposium on Application Accelerators in High Performance Computing (SAAHPC)*, 2012, pp. 54–63. DOI: 10.1109/SAAHPC.2012.18
- Sunita Chandrasekaran, Shilpa Shanbagh, and Douglas L. Maskell. "A Dependency Graph based Methodology for Parallelizing HLL Applications on FPGA." In *Proceedings of the 18th ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA)*, 2010, p. 283. *(CORE Rank: A)* DOI: 10.1145/1723112.1723169
- Sunita Chandrasekaran, Oscar Hernandez, Douglas Maskell, Barbara Chapman, and Van Bui. "Compilation and Parallelization Techniques with Tool Support to realize Sequence Alignment Algorithm on FPGA and Multicore." In *Proceedings of the 14th IEEE International Conference on High Performance Computing (HiPC)*, Goa, India, 2007, pp. 493–505. *(CORE Rank: A)* DOI: 10.1007/978-3-540-77220-0_46
- Suyang Zhu, Sunita Chandrasekaran, *Peng Sun, Barbara Chapman, Tobias Schuele, and Marcus Winter. "Exploring Task Parallelism for Heterogeneous Systems Using Multicore Task Management API." In *Euro-Par 2016: 4th Workshop on Runtime and Operating Systems for the Many-core Era (ROME)*, *LNCS*, Vol. 10104, 2016, pp. 607–708. DOI: 10.1007/978-3-319-58943-5_49

Invited Technical Reports Not Published Elsewhere

- *Mauricio Ferrato, Erin Crowgey, and Sunita Chandrasekaran. "Developing and Accelerating Predictive Models for Predicting Relapse of Pediatric Oncology patients using Smart Cyberinfrastructure." By Invitation-only NSF Workshop: Developing a Roadmap towards the Next Generation of Smart Cyberinfrastructure, Feb 25–27, Seattle, 2020.
- *Millad Ghane, Sunita Chandrasekaran, and Margaret S. Cheung. "Assessing Performance Implications of Deep Copy Operations via Microbenchmarking." arXiv preprint arXiv:1906.01128, June 2019.
- Sunita Chandrasekaran. "Extreme Heterogeneity for Sn Transport Codes." By Invitation-only Big Data and Extreme-Scale Computing (BDEC) Meeting, Indiana University, November 28–30, 2018.
- Sunita Chandrasekaran. "Development of a parallel algorithm for whole genome alignment for rapid delivery of personalized genomics." By Invitation-only Big Data and Extreme-Scale Computing (BDEC) Meeting, Indiana University, November 28–30, 2018.
- Ritu Arora, Xiaosong Li, Bonnie Hurwitz, Daniel Fay, Dhabaleswar K. Panda, Edward Valeev, Shaowen Wang, Shirley Moore, Sunita Chandrasekaran, and Ting Cao. "Future Directions of the Cyberinfrastructure for Sustained Scientific Innovation (CSSI) Program." Report submitted to NSF, April 2020. https://arxiv.org/pdf/2010.15584.pdf
- Robert Henschel, Junjie Li, Rudolf Eigenmann, and Sunita Chandrasekaran. "Explore True Performance Using Application Benchmark for the Next Generation HPC Systems: First NSF EAGER SPEC HPG Workshop Report." September 2019. https://scholarworks.iu.edu/dspace/handle/2022/25344. DOI: 10.5967/jmkd-

6p64

Software Packages Developed

Designing and developing software is an integral part of research. To that end, this section highlights some of the core software packages created and being created in my group. These packages address the full HPC stack: from understanding hardware capabilities, to validating software correctness, to enabling new science, to establishing best practices for performance-portable, production-ready computational tools.

For some projects, I have listed GitHub metrics as one indicator of community uptake: (a) **Stars** are bookmarks/likes by users; a higher count suggests broader visibility and adoption; (b) **Forks** are independent copies of the repository made for modification or extension, showing how many groups are actively building on the code; and (c) **Watchers/Network** are users monitoring the repository for updates - a more specific signal than stars, representing sustained interest rather than passive bookmarking.

However, GitHub metrics tell only part of the story. For many of the projects below, the more meaningful impact lies in how the software has served as critical infrastructure within a larger scientific or institutional effort, or in how it has measurably advanced research in a domain enabling new science, supporting flagship supercomputer deployments, or shaping decisions at partner labs and agencies. Where applicable, I have flagged this scientific and operational impact alongside the GitHub numbers.

1. GPU & Hardware Characterization

Benchmarking Characterization of Advanced GPUs

This project develops execution-centric microarchitectural characterizations of advanced GPU/accelerator hardware through systematic benchmarking. Papers published so far (a) examine AMD MI300A's FP8 matrix cores, asynchronous execution, and structured sparsity features using microbenchmarks to quantify performance trade-offs in concurrent execution and transformer workloads, and (b) dissect NVIDIA Blackwell's architecture including memory hierarchy, SM pipelines, and 5th-gen tensor cores supporting FP4/FP6 precisions by comparing it against Hopper. Together, these studies provide developers critical insights into modern GPU subsystems, latency characteristics, and optimization strategies for emerging AI workloads. *(Some GitHub repositories are not yet public, due to papers being under double-blind review.)*

Publications & Resources:

- Paper: Microbenchmark-Driven Analytical Performance Modeling | **GitHub:** Analytical Model | *Under Review*
- Paper: Microbenchmarking NVIDIA's Blackwell | **GitHub:** Blackwell B200 | **IPDPS 2026**
- Paper: FP8 Matrix Cores on AMD MI300A | **GitHub:** MI300A Microbenchmarking | *Under Review*

2. Validation & Verification of Programming Models

Section Summary. Three complementary efforts on V&V of parallel programming models. The long-standing **OpenMP_VV** and **OpenACC_VV** testsuites under the Exascale Computing Project (ECP) validate compiler implementations from NVIDIA, Cray, LLVM, Intel, and GNU across DOE flagship systems (Summit, Frontier, Perlmutter, Aurora), and are open-sourced on GitHub with sustained community uptake. The **LLM4VV** framework extends this work by positioning LLMs as autonomous test generators and judges for compiler V&V (3 publications, 70+ citations).

OpenMP Validation & Verification Testsuite (SOLLVE)

This project creates functional, orphan, and cross test codes for the OpenMP offloading programming model (since Version 4.0 through the latest version 6.1). The test codes validate implementations by compilers developed by NVIDIA, Cray, LLVM, Intel, and GNU. The testsuites are tested on platforms such as DOE flagship supercomputers including Summit and Frontier at ORNL, Perlmutter at LBNL, and Aurora at Argonne National Laboratory, among other systems. It has also been integrated into the harness infrastructure of previous systems like Summit at ORNL.

Publications & Resources:

- Paper: ECP SOLLVE Status Update and Compiler Insights (SC22 Workshop)
- **GitHub:** SOLLVE Validation & Verification Testsuite

Metrics

GitHub Metrics: 61 stars · 25 forks · 10 watchers/networks · high commit frequency

OpenACC Validation & Verification Testsuite (SOLLVE)

This project creates functional, orphan, and cross test codes for the OpenACC programming model (Version 1.0 through the latest version 3.4). The test codes validate implementations by compilers developed by NVIDIA, Cray, and GNU. The testsuites are tested on platforms such as DOE flagship supercomputers including Summit and Frontier at ORNL and Perlmutter at LBNL. It has been integrated into the harness infrastructure of previous systems like Summit at ORNL.

Publications & Resources:

- Paper: Analysis of OpenACC Compilers 3.0+ (SC22 Workshop)
- **GitHub:** OpenACC Validation & Verification Testsuite

Metrics

GitHub Metrics: 32 stars · 2 forks · 26 watchers/networks · 12 open issues · 15 pull requests · 2 releases · stable commit frequency

LLM4VV: Using LLMs for Compiler Validation & Verification

This project addresses a critical gap: while LLMs have become powerful tools for code generation, verifying the correctness of LLM-generated code remains a challenge without comprehensive, autonomous solutions.

LLM4VV explores using LLMs-as-judges to validate and verify compiler test suites, enabling more rigorous testing of parallel programming frameworks that have predominantly been a manual process and still continues to be.

Publications & Resources:

- Paper: LLM4VV: Developing LLM-driven testsuite for compiler validation (*Future Generation Computer Systems*); **cited 70+ times**.
- Paper: LLM4VV: Exploring LLM-as-a-Judge for Validation (SC24 Workshop)
- Paper: LLM4VV: Evaluating Cutting-Edge LLMs for Generation and Evaluation of Directive-Based Parallel Programming Model Compiler Tests (IEEE HiPC Conference)
- **GitHub:** LLM4VV Framework *(work-in-progress, not yet fully open-source)*

3. Machine Learning & Biomedical Applications

Drug Response Prediction for Cancer Treatment Optimization

This project, in collaboration with the **Frederick National Lab of Cancer Research (FNLCR)** and the **National Cancer Institute (NCI/NIH)**, produced **UNNT (Utility for comparing Neural Networks and Tree-based models)**, an open-source framework that lets cancer-research scientists drop in their own RNA-Seq gene expression and drug-response tabular data and rigorously benchmark CNN and XGBoost models side by side on the same workflow. Built on the NCI-DOE Pilot 1 (P1B3) project and applied to the NCI60 dataset (RNA-Seq profiles from 60 human cancer cell lines, drug response measurements, and molecular drug descriptors), UNNT closes the "validate this on my own dataset" gap that the original P1B3 pipeline had left open.

Impact: Findings overturned the assumption that GPUs and deep networks always win: on the full NCI60 drug set, XGBoost trained faster than the CNN even on V100 GPUs, and the CNN was faster on CPU than GPU because host-to-device transfers dominated guidance that directly affects how FNLCR and NCI/NIH allocate compute for precision-oncology pipelines. The framework is now reusable infrastructure for the broader drug-response modeling community, enabling biologically interpretable, hardware-aware model selection for cancer treatment optimization. impactbox

Publications & Resources:

- Paper: UNNT: A novel utility for comparing neural net and tree-based models (*PLOS Computational Biology*, 2024)
- **GitHub:** UNNT (Utility for comparing Neural Networks and Tree-based models)

Machine Learning for Rare Disease Drug Response Classification

This project, in collaboration with **Nemours Children's Health System, UC San Francisco, and Emory University School of Medicine**, produced **RCDML (RNA-seq Count Drug-response Machine Learning)**, an open-source, GPU-accelerated machine learning pipeline for drug response classification in rare disease research. The framework predicts patient sensitivity to therapeutic inhibitors directly from RNA-seq transcriptomic data and bundles feature selection, hyperparameter optimization, and a GPU-accelerated classification engine into a

single reusable workflow, applied at scale to the BeatAML clinical dataset.

Impact: The methodological novelty was repositioning **SHAP** from a post-hoc explainability tool into a feature-selection step retaining only the top 30 Shapley-ranked genes from the full 22,844-gene RNA-seq matrix to train the final classifier. This pipeline achieves up to ∼90% AUC in predicting ex vivo RTK-type-III inhibitor response at AML diagnosis, outperforming PCA and differential-gene-expression baselines, while the same gene signatures independently recover known drug-resistance pathways making the model both clinically accurate and biologically interpretable. Released open source as RCDML, it is now reusable infrastructure for pediatric-oncology and rare-disease groups at the intersection of transcriptomics and precision medicine. impactbox

Publications & Resources:

- Paper: Machine learning classifiers for RTK-Type-III inhibitor response (*Bioinformatics Advances*, Oxford University Press, 2023)
- **GitHub:** RNA-seq Count Drug-response Machine Learning (RCDML) Framework

4. Reimagining Algorithms: GPU Refactoring for Faster Science

OpenACC-Accelerated PhysiCell

This project, in collaboration with **Indiana University**, is an OpenACC-accelerated GPU version of PhysiCell, an open-source physics-based cell simulator for 3D multicellular biological systems. The project uses OpenACC (a directive-based programming model) to accelerate the diffusion computations in PhysiCell, which is widely used in cancer and infectious disease modeling.

Impact: The accelerated diffusion equation reduced execution time for 21 days of biological diffusion and data saves from **9.3 hours to 1.7 hours**, a reduction of over **80%** and was subsequently integrated into the main PhysiCell project's development. impactbox

Publications & Resources:

- Paper: OpenACC Acceleration of an Agent-Based Biological Simulation Framework (*IEEE CiSE*, 2022)
- **GitHub:** PhysiCell GPU

Refactoring MURaM for GPU/CPU Performance Portability via OpenACC

This project, in collaboration with the **Max Planck Institute for Solar System Research** and the **National Center for Atmospheric Research (NCAR)**, refactors the MURaM code, a widely used model for simulating the solar atmosphere (from the upper convection zone to the corona) to run efficiently on GPUs while maintaining performance portability. The aim is to keep a single source code that can run on both traditional CPUs and modern GPU accelerators without sacrificing significant speed on either.

Impact: The work enabled (i) longer simulations over extended time periods, higher spatial resolution studies, and better fidelity in multi-physics modeling by reducing time spent on RTS computations; and (ii) more detailed solar-physics studies of quiet and active sun phenomena, eruptive events like solar flares and coronal mass ejections, realistic synthesis of emission across visible to extreme UV and X-ray wavelengths, and more comprehensive comparisons with multi-wavelength observations from space telescopes. impactbox

Publications & Resources:

- Paper: Refactoring the MPS/University of Chicago Radiative MHD (MURaM) model for GPU/CPU performance portability using OpenACC directives (*ACM PASC*, 2021; with **ACM Artifact Availability badge**)
- **GitHub:** MURaM GPU

Metrics

GitHub Metrics: 14 stars · 8 forks · 2 watchers/networks · 1 open issue · emerging contributor base, active development

The Artificial Scientist: In-Transit Machine Learning of Plasma Simulations

This project, in collaboration with **Oak Ridge National Lab, Helmholtz Zentrum Dresden Rossendorf, HPE, and AMD**, addresses a fundamental scaling challenge: large-scale simulations and scientific experiments generate petabytes of data per run, creating massive bottlenecks when attempting to perform analysis manually offline. The project demonstrates a heterogeneous streaming workflow that performs machine learning directly on simulation data as it is being generated, completely bypassing the filesystem bottleneck.

Impact: This work represents a fundamental shift in exascale system architecture from *write-everything-then-analyze* to *stream-and-learn-continuously*. The project paves the way for redesigning HPC workflows, storage systems, and software stacks for the next generation of leadership-class computing. impactbox

Publications & Resources:

- Paper: The Artificial Scientist: In-Transit Machine Learning of Plasma Simulations
- **GitHub:** Plasma Physics in-situ ML

Metrics

GitHub Metrics: 3 stars · 1 fork · 3 watchers/networks · active development

5. NSF SCIPE Democratizing Access to Research Software Engineering (DARSE)

DARSE (Democratizing Access to Research Software Engineering) is funded by the National Science Foundation (grant 2417814). The mission of this project to advance science while we build a pipeline of research software engineering (RSE) professionals and trainees by upskilling people in AI, high-performance computing, and machine learning to support scientific research at UD and across the Mid-Atlantic region.

Impact: DARSE addresses a national gap in the scientific workforce: the chronic shortage of Research Software Engineers (RSEs) capable of operating at the intersection of AI, scientific computing, software sustainability, and domain science. Unlike traditional "research IT" or transactional consulting, DARSE establishes an *embedded RSE model* in which trained engineers and student trainees work alongside domain scientists on long-horizon research problems. In its first two years, the project has led to writing peer-reviewed submissions (e.g., AI4OSN, Russian foreign-policy speech analysis, multicomponent emulsions simulation) - both under review, open com-

munity datasets on Harvard Dataverse, and reproducible software pipelines across nine distinct domains spanning the natural, physical, and social sciences. The project demonstrates that the embedded RSE methodology produces both stronger science and durable computational infrastructure that outlasts individual research projects. impactbox

Publications & Resources:

- **GitHub:** NSF DARSE Software Development; We have approximately 24 on-going projects that includes projects solicited from across campus plus projects mentored via an RSE Class that I taught in Spring 2026.

Mentoring

Key Impact: 27 students mentored across all degree levels, with strong placement in top tech companies (NVIDIA x5, AMD, Google, Meta, JPMorgan, SpaceX) and national laboratories (Lawrence Livermore National Lab).

Postdocs mentees

Student	Year	Status	Current Position
Felipe Cabarcas	Postdoc	2022-2024	Universidad de Antioquia
Pinar Kullu	Postdoc	2022-2024	University of Central Florida

PhD Mentees

Student	Thesis Title	Year	Current Position
Vineeth Gutta	Optimizing and scaling machine learning models for scientific applications on exascale supercomputers	2025	NVIDIA
Fabian Mora	Designing and developing an extensible compiler for heterogeneous computing using MLIR infrastructure	2024	AMD
Eric Wright	Exploring hierarchical parallelism in directive-based models for efficient GPU execution	2023	Lawrence Livermore National Lab
Mauricio Ferrato	Predicting outcomes for rare diseases using machine learning techniques	2023	NVIDIA
Robert Searles	Creating a Portable Programming Abstraction for Wavefront Patterns Targeting HPC Systems	2019	NVIDIA

MS mentees with and without Thesis

Student	Thesis Title	Year	Current Position
Thomas Huber	OpenMP Offloading Validation and Verification Testsuite	2022	AMD
Jose Monsalve Diaz	A Testsuite Design and Implementation for OpenMP 4.5 Offloading Features	2020	AMD
Kshitij Srivatsava	N/A	2017	Uber
Arnov Sinha	High Performance Sparse Fast Fourier Transform Using GPUs	2017	Red Violet
Vaidhyanathan Ravichandran	N/A	2017	Nutrisystem

Undergraduate Thesis Mentees

Student	Thesis Title	Year	Current Status
Andrew Kallai	Statistical Analysis of LLVM IR Compilation	2024	PhD Student (UD)
Talha Mahmood	Scaling MMST-ViT for Nationwide Crop Yield Prediction in the US	2024	PhD Student (UD)
Kristina Holsapple	Is Exascale Enough? Investigation into GPU Performance on AMD MI250x	2023	Recorded Future

Core Undergraduate Research Mentees

Student	Year	Current Position/Status
Zahra Temori	2026	BlackRock
Saieda Ali Zada	2026	PhD Student, UD (Fall 2026)
Rahul Patel	2026	4+1 Degree Program, UD
Christian Munley	2024	NVIDIA
Jay Patel	2024	JP Morgan Chase
Michael Carr	2023	W. L. Gore Associates
Jaydon Reap	2023	Phin Security
Kristina Holsapple	2023	Recorded Future
William Gunter	2022	SpaceX
Miguel Zavala	2021	Capital One
Joshua Davis	2021	PhD completed (UMaryland), Joining NVIDIA
Matt Stack	2020	NVIDIA
Hayden Carter	2020	PhD Student, UMass Amherst
Daniel Liang	2018	Meta

Key Impact: 3 RSE, 4 PhD students, and 6 undergraduate research students are being mentored currently.

Current Research Software Engineering Mentees

Student	Year	Status
Kevin Bhimani	RSE	Current
Abhin Suresh	RSE	Current
Christian Escobar	RSE	Current

Current PhD Mentees

Student	Year	Status
Aaron Jarmusch	3rd year	Successfully completed preliminary examination
Giorgi Gvalia	2nd year	Passed in May 2026
Andrew Kallai	2nd year	Preliminary examination due in Fall 2026
Nikhil Rao	1st year	Preliminary examination due in Spring 2027

Current Undergraduate Mentees

Student	Year	Status
Nathan Graddon	Junior	Class 2027
Zach Sollenberger	Junior	Class 2027
Ryan Padrone	Junior	Class 2027
Connor Vitz	Junior	Class 2028
Aastha Desai	Sophomore	Class 2028
Wyatt Erwin	Sophomore	Class 2028

Tutorials Presented at Conferences

- Tutorial on using OpenACC, a directive-based programming model. *Submitted to SC26, pending review.*
- Tutorial on "Towards Comprehensive System Comparison: Using the SPEC HPG Benchmarks for Better Analysis, Evaluation, and Procurement of Next-Generation HPC Systems," given at:
  - Half-day Tutorial at SC20 (Virtual), November 2020
  - PEARC'19, Chicago, IL, USA
  - Half-day Tutorial at ICS 2019, Phoenix, AZ, USA
  - Half-day Tutorial at ISC 2019, Frankfurt, Germany
  - Half-day Tutorial at SC15, Austin, TX, USA

Invited Technical Talks

2026

- Forthcoming talk, November 2026 - On behalf of the Prime Minister of Estonia, I am invited to join a distinguished circle of global leaders as a speaker at the ninth Tallinn Digital Summit, to be held in Tallinn,

Estonia, on 5–6 November 2026.

- OpenACC User Workshop co-located with SCA2026, Osaka, Japan, January 2026.
- In-Transit Learning at Exascale: A Streaming ML Workflow for Scientific Simulations, SIAM Parallel Processing, Invited Talk, Berlin, Germany, March 2026.
- Delaware's DNA 2026 Life Science Conference, Invited Panelist, Delaware, USA, April 2026.

2025

- The State of AI in Delaware, New Castle Chamber of Commerce, Delaware, USA, July 2025.
- Preparing Graduates for the AI-Driven Future: Multidisciplinary Perspectives from Industry, State of Delaware AI Commission's Subcommittee, Delaware, USA, March 2025.
- The Artificial Scientist - Leveraging In-Transit Machine Learning for Plasma Simulations, Friedrich-Alexander University Erlangen National High-Performance Computing Center (Remote), May 2025.
- The Artificial Scientist - Leveraging In-Transit Machine Learning for Plasma Simulations, Salishan Conference on High Speed Computing, Portland, USA, April 2025.
- The Artificial Scientist - Leveraging In-Transit Machine Learning for Plasma Simulations, MulticoreWorld, Christchurch, New Zealand, February 2025.
- High Performance Computing (HPC) and Artificial Intelligence (AI) Research with Research Software Engineers (RSE), UD's Synthetic Biology Symposium, Delaware, USA, January 2025.
- The Artificial Scientist - Leveraging In-Transit Machine Learning for Plasma Simulations, Princeton Plasma Physics Laboratory, Princeton University, NJ, USA, January 2025.
- Surrogate Models - A Frontier to Explore in Healthcare, AI4Health, Delaware, USA, January 2025.

2024

- Convergence of High Performance Computing and AI/Machine Learning for Real-World Applications, UD's College of Engineering Advisory Council, University of Delaware, USA, October 2024.
- Advancing Interdisciplinary Science using AI & High Performance Computing, Incyte, Delaware, USA, October 2024.
- Software Advancements Meeting Hardware Innovations: Challenges and Solutions, IEEE Cluster Keynote, Kobe, Japan, September 2024.
- Scaling Real-World Applications on Heterogeneous Platforms Including Frontier, TU Munich, Germany, April 2024.
- OpenMP Validation and Verification Test-Suite Insights and Lessons Learnt, SIAM PP, Baltimore, USA, March 2024.

2023

- Advancing Science by Facilitating Easier Access to HPC and Cloud Resources, HPE HPC & AI Cloud Services Teams (Remote), September 2023.
- Plasma Physics on Frontier and Future of LLMs, HPE CTO Labs (Remote), September 2023.

- ECP SOLLVE Test-Suite Infrastructure, HPE Compiler Teams (Remote), September 2023.
- Lecture on ECP and Parallel Programming, KTH Summer School, Stockholm, Sweden, August 2023.
- Leveraging Exascale Computing Resources for Particle-In-Cell on GPU (PIConGPU), SIAM CSE 2023, Amsterdam, The Netherlands, February 2023.
- Leveraging Exascale Computing Resources for Particle-In-Cell on GPU (PIConGPU), Innovative Computing Laboratory, University of Tennessee, TN, USA, 2023.

2022

- Challenges and Success Stories Migrating Software and Applications to Frontier, Keynote at IPDRM Workshop @ SC22, Dallas, TX, USA, November 2022.
- ECP SOLLVE at the DOE Booth, SC22, Dallas, TX, USA, November 2022.
- The Race to Frontier, Stony Brook University, NY, USA, September 2022.
- Preparing Effective Grant Proposals, University Resources, Early Career Program at SC22, Dallas, TX, USA, November 2022.
- The Good, Bad and the Ugly with PIConGPU on Frontier, AMD (Remote), September 2022.
- Preparing PIConGPU for Next-Generation Computing Systems, Brookhaven National Laboratory, NY, USA, January 2022.

2021 (Virtual Talks)

- Experience with Porting and Scaling Codes on AMD GPUs, First PaCER Conference (P'con), December 2021.
- How Ready Are We to Use the First Exascale Supercomputer, Frontier?, DOE ECP Booth, November 2021.
- Best Practices for a Productive Yet Performant Software Development, 2021 Europe ACM Summer School, August 2021.
- An HPC Journey - Porting a Solar Physics Code to Large Scale Systems, The Society of HPC Professionals, July 2021.
- Ten Ways to Build a Productive Yet Performant Software, Supercomputing Frontiers Europe, July 2021.
- Programming Frontier: Challenges and Solutions, OLCF User Group Meeting, June 2021.
- Present and Future of Accelerated Computing Programming Approaches, Panelist for NVIDIA GPU Technology Conference (GTC), March 2021.
- Exascale Simulations for the Next Generation of Plasma Accelerators with PIConGPU, SIAM CSE 2021, February 2021.
- Preparing to Program the World's Fastest Supercomputer, NSF-Funded DARWIN System Symposium, University of Delaware, February 2021.

2020 (Virtual Talks)

- Evolution of a Project, Students at the Supercomputing Conference (SC) Program, November 2020.

- Hierarchical Parallelism for Exascale Computing (HiPar20), SC Panelist, November 2020.
- Developing Software for Today's and Tomorrow's Platforms - Fun or a Nightmare?, University of Tennessee Knoxville Seminar Series, October 2020.
- Research Activities at CRPL, Hierarchical Computations on Manycore Architecture Group at KAUST, Saudi Arabia, August 2020.
- Preparing Software Stack for the Next Generation Systems -An Opportunity or a Nightmare?, P2S2 Workshop at ICPP 2020, August 2020.
- Scientific Software Productivity - Case Studies, Challenges, Opportunities and Potential Solutions, 2020 Collegeville Workshop on Scientific Software (Panel), July 2020.
- ECP SOLLVE: OpenMP Validation and Verification Effort, OpenPOWER Academia and Research Webinar, June 2020.
- GPUs for Science Day 2020, NERSC, Lawrence Berkeley National Laboratory, Berkeley, CA, June 2020.
- Training and Best Practices to Develop Portable Yet Performant Code, SIAM PP, Seattle, USA, February 2020.

2019

- Running PIConGPU on Summit and Preparing PIConGPU for Frontier, 4th OpenPOWER Academic and Research Workshop, SC19, Denver, USA, November 2019.
- OpenMP 4.5 Validation and Verification Test-Suite, 4th OpenPOWER Academic and Research Workshop, SC19, Denver, USA, November 2019.
- Applying Directives to Port MURaM Code to Heterogeneous Systems, ASTRONUM, Paris, France, July 2019.
- Development of a Parallel Algorithm for WGS Alignment for Rapid Delivery of Personalized Genomics, PASC Minisymposium, Zurich, Switzerland, June 2019.
- Impact of Parallel Programming Models on Interdisciplinary Scientific Research, University of Basel, Basel, Switzerland, June 2019.
- OpenACC-Based GPU Acceleration of Chemical Shift Prediction, Invited Featured Speaker, GPU Technology Conference (GTC), CA, USA, March 2019.
- Porting MURaM (Max Planck University of Chicago Radiative MHD) to GPUs Using OpenACC, Invited Speaker, GPU Technology Conference (GTC), CA, USA, March 2019.
- Acceleration of Prediction of Chemical Shift Structures, SIAM CSE, Spokane, USA, February 2019.

2018

- Development of a Parallel Algorithm for Whole Genome Alignment for Rapid Delivery of Personalized Genomics, NSF-Funded International Workshop BDEC Meeting, Indiana University, Bloomington, USA, November 2018.
- 3P to Science Using OpenACC: Performance, Productivity, and Portability, NVIDIA Booth at SC Showfloor, Dallas, November 2018.

- Swiss Army Programming: Performance and Portability from Modern Tools, Invited Panelist, SC18, Dallas, November 2018.
- Using the Parallel Programming Model OpenACC to Do More Science and Less Programming, Bootcamp, Princeton University, NJ, USA, October 2018.
- HPC-as-a-Service to Domain Scientists, PASC Minisymposium, Basel, Switzerland, July 2018.
- Opportunities and Challenges Migrating Scientific Code to Accelerators, National Center for Atmospheric Research (NCAR), Boulder, USA, June 2018.
- Achieving Performance While Preserving Portability for NGS Application, SIAM PP, Tokyo, Japan, March 2018.
- Adapting Minisweep, a Proxy Application, on Heterogeneous Systems Using OpenACC Directives, Featured Speaker, GPU Technology Conference (GTC), CA, USA, March 2018.
- Path Forward for Softwarization to Tackle Evolving Hardware, SPIE, Orlando, USA, April 2018.

2017

- Parallelization and Acceleration of the Nuclear Reactor Mini-App Minisweep on an OpenPOWER Platform, 2nd OpenPOWER Academia Discussion Group Workshop, SC17, Denver, USA, November 2017.
- Building Your Academic Professional Network, Grace Hopper Celebrations, Houston, USA, October 2017.
- Using OpenACC for NGS Techniques to Create a Portable and Easy-to-Use Code Base, GPU Technology Conference (GTC), CA, USA, March 2017.
- Exploring On-Node Programming Models for Irregular Algorithms, SIAM CSE, Atlanta, USA, February 2017.
- Programmer's Perspective on Evolving Hardware: Challenges and Success Stories, Invited Guest Lecture, RWTH Aachen, Germany, February 2017.

2016 and Before

- Hackathons and Best Practices in HPC Training, Workshop Co-located at SC16, Salt Lake City, USA, November 2016.
- OpenACC Status and Feedback, Birds of a Feather (BoF), GPU Technology Conference (GTC), San Jose, USA, May 2016.
- OpenACC API: User Experience, Vendor Reaction, Relevance, and Roadmap, BoF Speaker at SC16, Salt Lake City, November 2016.
- Industry Standards for Programming Multicore Systems: Way to Go!, Multicore Devcon Conference (MDC), Santa Clara, CA, May 2014.
- Exascale Will Soon Be Here - How Prepared Are We?, Argonne National Laboratory, Chicago, USA, April 2014.
- Simplifying Heterogeneous Multicore Programming Using Industry Standards, SIAM PP, Portland, February 2014.
- Open Registry for Accelerated Computing, Many-Core and Reconfigurable Supercomputing Conference

(MRSC), Bristol, UK, 2011.

Professional Activities - External

Special Content Editor

Future Generation Computing Systems (FGCS). 2021–2023.

Journal Guest Co-Editorship

- **Software: Practice and Experience.** Special Issue on New Trends in High Performance Computing: Software Systems and Applications. Co-edited with Min Si, Lena Oden, Jidong Zhai. September 2022.
- **Springer, Communications in Computer and Information Science (CCIS).** Book on Tools and Techniques for High Performance Computing. Selected papers from SC19 Workshop. Co-edited with Dr. Guido Juckeland. ISBN: 978-3-030-44728-1. April 2020.
- **Springer, Lecture Notes in Computer Science (LNCS).** Book on Accelerator Programming Using Directives. Selected papers from SC18 Workshop. Co-edited with Dr. Guido Juckeland. ISBN: 978-3-030-12274-4. February 2019.
- **Springer, LNCS.** Book on Accelerator Programming Using Directives. Selected papers from SC17 Workshop. Co-edited with Dr. Guido Juckeland. ISBN: 978-3-319-74896-2. February 2017.
- **BMC Bioinformatics.** Special Issue: Computational Approaches for Cancer. Co-edited with Dr. Eric Stahlberg. Volume 19, Article 487. December 2018. https://doi.org/10.1186/s12859-018-2502-x
- **Parallel Computing (PARCO).** Special Issue: Applications for the Heterogeneous Computing Era. Volume 77. Co-edited with Dr. Antonio Pena. September 2018. https://doi.org/10.1016/j.parco.2018.06.002
- **Parallel Computing (PARCO).** Special Issue: Topics on Heterogeneous Computing Era. Volume 68. Co-edited with Dr. Antonio Pena. October 2017. https://doi.org/10.1016/j.parco.2017.08.001
- **Inderscience Publishers, International Journal of Embedded and Real-Time Communication Systems.** Special Issue: Novel Strategies for Programming Accelerators. Co-edited with Dr. Guido Juckeland. December 2017.
- **Inderscience Publishers, International Journal of Embedded and Real-Time Communication Systems.** Special Issue: High-level Programming Approaches for Accelerators. Co-edited with Dr. Guido Juckeland. December 2017.
- **Journal of Scientific Programming.** Special Issue: Programming Models, Languages, and Compilers for Manycore and Heterogeneous Architectures. Volume 2015, Article ID 376317. 2015. http://dx.doi.org/10.1155/2015/376

Scientific Advisory Boards and Committees

Current Advisory Boards

- ORNL Computing and Computational Sciences Directorate, Directorate Advisory Committee (DAC), 2025–2028.

Past Advisory Boards and Committees

- NCAR CISL Advisory Panel Member, 2022–2025.
- DOE Advanced Scientific Computing Advisory Committee (ASCAC), 2024–2025.
- Advisory Editorial Board, Software: Practice and Experience (SPE), 2022–2024.
- Organizing Committee, SIAM Parallel Processing Conference, 2022.

  Prior to 2021

- European H2020 Project: EPEEC (European Joint Effort Toward a Highly Productive Programming Environment for Heterogeneous Exascale Computing), Scientific-Industrial Advisory Board (SAIB), Barcelona Supercomputing Center, Spain, 2018–2020.

Proposal Reviewer, US and International

- NSF Directorate for Computer and Information Science and Engineering (CISE), Office of Advanced Cyberinfrastructure (OAC), 2026
- DOE Review Panel for Genesis Mission, 2026
- DOE Laboratory Directed Research and Development (LDRD), 2024
- DOE Presidential Early Career Awards for Scientists and Engineers (PECASE), 2021

  Prior to 2021

- NSF Review Panel, 2016, 2017, 2018, 2019, 2020, 2022
- Natural Sciences and Engineering Research Council of Canada, Review Panel, 2017
- DFG (German Research Foundation), Germany, Review Panel, 2017

Technical Specification and Book Reviewer

- *Computer Systems: An Embedded Approach*, Textbook by Ian McLoughlin, September 2018
- Multicore Association (MCA) Task Management Standard API (MTAPI) and Software-Hardware Interface for Multi-Many-Core (SHIM), 2013–2015

Editorial Affiliations, since 2021

- Special Content Editor, Future Generation Computer Systems, 2021–2023
- IEEE Transactions on Parallel and Distributed Systems, Associate Editor, 2020–2023
- Elsevier's Future Generation Computer Systems, Associate Editor, 2020–2021
- Journal of Parallel Computing (PARCO), Subject Area Editor, 2019–2021
- Journal of Parallel and Distributed Computing (JPDC), Associate Editor, 2019–2021

Steering Committees

- NSF Workshop on Future Directions of the CSSI Program, Steering Committee, 2019
- Women in HPC, SC, Steering Committee, 2016–2018
- Women in HPC, ISC, Steering Committee, 2016–2018

Chair and Co-Chair Roles

Conferences, Workshops, Symposiums, and Scholarships

Current (2026)

- IEEE Technical Computing High Performance Computing (TCHPC) Award Chair, SC26
- Jack Dongarra Early Career Award, Award Committee, ISC 2026
- Birds of a Feather Chair, SC26: The International Conference for High Performance Computing, Networking, Storage, and Analysis, 2026

Recent Past (2020–2025)

- Program Co-Chair, 54th International Conference on Parallel Processing (ICPP), 2025.
- Panels Chair, SC25, 2025.
- Poster Chair, ACM Student Research Competition (SRC), International Symposium on Code Generation and Optimization (CGO), 2025.
- Applications Track Co-Chair, SC24, 2024.
- Technical Program Co-Chair, IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid), 2023.
- IEEE Cluster Student Travel Liaison Chair, 2020–2022.
- Vice Chair, Interdisciplinary Track, International Parallel & Distributed Processing Symposium (IPDPS), 2022.
- Program Co-Chair, 34th International Workshop on Languages and Compilers for Parallel Computing (LCPC), 2021.

  Prior to 2021

- Track Leader, Invited Speaker Program, ISC 2020–2021.
- Technical Program Co-Chair, Platform for Advanced Scientific Computing (PASC), 2020.
- Vice Chair, Early Career Program, SC 2020.

Earlier Roles (2014–2019)

- General Chair, International Workshop on Accelerators and Hybrid Exascale Systems (AsHES), co-located with IPDPS, 2018–2019.
- Research Posters Chair, International Supercomputing Conference (ISC), 2019.

- Workshops Chair, SC19, 2019.
- Technical Program Co-Chair, PASC 2019, 2019.
- Workshop Co-Chair, 1st–4th International Workshop on Performance Portable Programming Models for Accelerators (P^3MA), co-located with ISC, 2016–2019.
- Workshop Co-Chair, 1st–5th Workshop on Accelerator Programming Using Directives (WACCPD), co-located with SC, 2014–2018.
- Research Posters Vice-Chair, ISC, 2018.
- ACM Intel Leadership Fellows, SC 2018.
- General Chair, 7th and 8th International Workshop on Accelerators and Hybrid Exascale Systems (AsHES), co-located with IPDPS, 2017–2018.
- Workshop Co-Chair, 1st–3rd HPC Applications in Precision Medicine, co-located with ISC, 2017–2019.
- Doctoral Showcase Chair, SC 2017.
- Vice-Chair, Software Track, IPDPS, 2017.
- Track Chair, Programming Models and Systems Software, ISC, 2017.
- Workshop Co-Chair, 4th–6th AsHES, co-located with IPDPS, 2014–2016.
- Poster and Research Demo Chair, 15th IEEE/ACM CCGrid, 2015.
- Workshop Chair, Programming Models, Languages and Compilers Workshop for Manycore and Heterogeneous Architectures (PLC), co-located with IPDPS, 2015.
- Workshop Chair, Workshop on Directives and Tools for Accelerators: A Seismic Programming Shift, 2014–2015.

Technical Program Committee - Conferences/Workshops/Symposiums/Scholarships

- International Supercomputing Conference (ISC), 2025.
- LLM4HPC 2025 (Workshop on Large Language Models for High Performance Computing), 2025.
- International Conference on Supercomputing (ICS), 2025.
- SC24, 2024.
- LLVM HPC Workshop (Co-located with SC24), 2024.
- IEEE HiCoMB (High Performance Computational Biology and Bioinformatics), 2022.
- IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid), 2022.
- IEEE Cluster Conference, 2021.
- IEEE International Workshop on OpenMP, 2021.
- International Supercomputing Conference (ISC), 2021.
- International Parallel & Distributed Processing Symposium (IPDPS), 2021.
- SC, 2021.
- CCGrid, 2021.
- SC, 2020.

- International Conference on Supercomputing (ICS), 2020.
- ISC, 2019.
- IPDPS Workshops, 2019.
- SC, 2018.
- DOE Leadership Computing INCITE Program Review, 2018.
- IPDPS, 2018.
- CCGrid (Track: Programming Models and Runtime Systems), 2018.
- International Workshop on FPGAs for Software Programmers (FSP), 2018.
- 3rd Workshop on Open Source Supercomputing, 2018.
- SC (Doctoral Showcase Chair), 2017.
- International Conference on Parallel Processing (ICPP), 2017.
- CCGrid (Track: Programming Models and Runtime Systems), 2017.
- SC, 2016.
- SC (Doctoral Showcase, BoF Algorithms Track), 2016.
- Workshop on Computing and Cancer, 2016.
- Workshops on Energy Efficient Supercomputing (E2SC), 2016.
- FSP, 2016.
- International European Conference on Parallel and Distributed Computing (EuroPar), 2016.
- International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD), 2016.
- International Conference on Network and Parallel Computing, 2016.
- ICPP, 2016.
- DOD Workshop on Mission-Critical Big Data Analytics, 2016.
- First International Workshop on Open POWER for HPC (IWOPH) at ISC, 2016.
- CCGrid, 2014–2016.
- IPDPS - PhD Panel on Research and Career Planning; PhD Poster Judging Committee, 2015.
- SC, 2015.
- SC - Doctoral Showcase Committee; Performance Track, 2016.
- IEEE Cluster Conference, 2014.
- 8th Workshop on General Purpose Processing using GPUs (GPGPU-8), 2014.
- 22nd European Signal Processing Conference (EUSIPCO), 2014.
- International Symposium on Integrated Circuits (ISIC), 2014.
- Workshop on Multicore and GPU Programming Models, Languages and Compilers at IPDPS, 2013–2014.
- International Joint Conference on Neural Networks (IJCNN), 2015.
- SC, 2012–2013.

- ACM Student Research Competition - Poster Committee.
- 50th Design Automation Conference (DAC) - External Reviewer, 2012.

3- to 5-day Programming Hackathons

- UDEL and State of Delaware Hackathon, Upcoming - August 2026.
- UDEL FinTech Hackathon, August 2025. Blog & UDaily article.
- UDEL AI Center of Excellence + Data Science Institute Hackathon, July 2024. Blog & Technical.ly article.
- UDEL AI Center of Excellence + Data Science Institute Hackathon, July 2023. Blog.
- Exascale Computing Project OpenMP Virtual Hackathons, 2021. Blog.

  Prior to 2021
- GPU Programming Hackathon Training in collaboration with Oak Ridge National Laboratory, Brookhaven National Laboratory (BNL), NY, USA, June 2017.
- UDEL GPU Hackathons in collaboration with Oak Ridge National Lab and NVIDIA, May 2016. UDaily article.
- Hands-on training: Introduction to GPGPU Architecture and OpenACC. Center for Advanced Computing and Data Systems (CACDS), University of Houston, TX, April 2014.

Career Development Workshops

- Salishan High Performance Computing Workshop, Portland, OR, USA, 2025
- Dagstuhl Seminar 23132: AI-Augmented Facilities: Bridging Experiment and Simulation with ML, March 2023

  Prior to 2021
- Eastern Nook Promotion and Tenure Workshop, NJ, 2019
- Write Winning Grant Proposals Workshop, VA, 2017
- CRA Career Mentoring Workshop, VA, USA, 2016
- NSF CISE CAREER Workshop, VA, 2016

Professional Activities - UDEL

University

- Research Council, University of Delaware, 2025 – current

- Cost Recovery Working Group, Spring 2020
- Faculty Peer Observation Program (FPOP), Spring 2017 – Spring 2019

College of Engineering

- Dean's Junior Faculty Advisory Council, Spring 2018 – Spring 2019
- Search Committee for Digital Content Specialist, Summer 2016

Department of Computer & Information Sciences

- Department P&T Committee, Spring and Fall 2025
- CIS Bylaws Committee, Fall 2022; Spring 2023; Fall 2024
- CIS Chair Search Committee, Fall 2020 – Spring 2022
- SPICE Committee, Fall 2020 – Spring 2022
- Alumni Committee, Fall 2021 – Spring 2022
- Executive Committee, Fall 2019 – Spring 2020
- Marketing Committee, Fall 2018 – Fall 2019
- Faculty Search Committee - High Performance Computing, Fall 2018 – Spring 2019
- Search Committee - Teaching (Continuing), Fall 2016 – Spring 2017
- Publicity Committee, Fall 2016 – Spring 2017
- Faculty Search Committee - Networking, Fall 2016 – Spring 2017

Media Coverage

- **2026:** Shaping a Human-Centered AI Future: Insights from the First State AI Institute *(UDaily, Feb 27, 2026)*
- **2025:** UDEL & FinTech Hackathon - Hen Street Hacks, August 13–15, 2025
- **2025:** Podcast: Science in Parallel - S6E10: Sunita Chandrasekaran: Computation in Translation
- **2025:** ACM SIGHPC Emerging Women Leader in Technical Computing - HPCwire Coverage
- **2025:** Multicore.world Speaker - Speaking Engagement
- **2024:** Building a Data-Intensive Research Workforce for the Mid-Atlantic (SCIPE Project) *(UDaily, November 2024)*
- **2024:** State of Delaware AI Commission Leadership *(Delaware Government)*
- **2023:** HPC Software Rock Star Sunita Chandrasekaran on Exascale Programming *(insideHPC Podcast)*
- **2023:** Utilizing the Frontier Supercomputer for PIConGPU *(AMD News)*
- **2022:** Interdisciplinary Science: Computing for Pediatric Health *(UDaily)*

- **2021:** Sunita Chandrasekaran Reflects on Teaching Supercomputing and Leading the ECP SOLLVE Project *(Exascale Project)*
- **2021:** ECP Selects Chandrasekaran as PI of SOLLVE Project *(HPCwire)*
- **2021:** Sunita Chandrasekaran to Lead Exascale SOLLVE Project *(Brookhaven National Lab)*
- **2021:** University of Delaware Team Tightens up Code for Exascale Computing on Frontier *(Oak Ridge National Lab)*

  Prior to 2021

- **2020:** Recipient of the Excellence in Teaching Award *(UDaily)*
- **2020:** NVIDIA DLI Releases New Accelerated Computing Teaching Kit *(NVIDIA News)*
- **2016:** IEEE TCHPC Early Career Researcher Award *(IEEE TCHPC News)*