

Ian Lumsden

Curriculum Vitae

☎ (865) 399 1695

✉ ilumsden@vols.utk.edu

🌐 [ian-lumsden-529bb1170](https://www.linkedin.com/in/ian-lumsden-529bb1170)

🌐 [ilumsden](https://www.github.com/ilumsden)

Research Interests

- High-Performance Computing
- Workflow Scheduling, Orchestration, and Data Movement
- Performance Analysis
- Software Engineering
- Data Analysis

Education

2020

PhD Computer Science, *Tickle College of Engineering*,
University of Tennessee, Knoxville, TN
Concentration in High-Performance Computing

2016

2020

B.S. Computer Science, *Tickle College of Engineering*,
University of Tennessee, Knoxville, TN

Major in Computer Science (Honors Concentration) and Part of Chancellor's Honors Program and Cook Grand Challenge Engineering Honors Program

Professional/Research Experience

2019

Graduate (formerly Undergraduate) Research Assistant,
Global Computing Laboratory, University of Tennessee, Knoxville, TN

Advisor Dr. Michela Taufer

- Projects**
- Developing a benchmark to study data movement motifs in scientific workflows and assess the performance of data movement tools in these motifs
 - Developing novel techniques and tools based on Flux to assist in data movement for scientific workflows
 - Developing novel techniques to identify causes of interesting performance phenomena in High-Performance Computing applications using Hatchet
 - Examining the performance and implications of in-situ and in-transit data analysis of molecular dynamics simulations through the Analytics4MD project (<https://analytics4md.org/>)

2024

Summer Graduate Computing Student Intern, Lawrence Livermore National Laboratory, Livermore, CA

Mentors Dr. Olga Pearce, Dr. Jae-Seung Yeom

- Projects**
- Developed a benchmark to study data movement motifs in scientific workflows and assess the performance of data movement tools in these motifs
 - Conducted a performance study of different CPUs and GPUs with different types of memory using the RAJA Performance Suite with other members of the Thicket team
 - Refactored the topdown analysis service in LLNL's Caliper performance monitoring tool and added support for Intel Sapphire Rapids CPUs
 - Developed Python bindings for Caliper and used them to develop performance annotations for LLNL's Hatchet and Thicket tools

2023

Summer Graduate Computing Student Intern, Lawrence Livermore National Laboratory, Livermore, CA

Mentors Dr. Jae-Seung Yeom, Dr. Hariharan Devarajan

- Projects**
- Developed an approach to automatically detect all accessible levels of the storage hierarchy and model it as a bipartite graph
 - Integrated this approach into LLNL's Dynamic and Asynchronous Data Streamliner (DYAD)
 - Improved the structure and performance of DYAD and the test suite developed during the previous summer

2022

Summer Graduate Computing Student Intern, Lawrence Livermore National Laboratory, Livermore, CA

Mentors Dr. Jae-Seung Yeom

- Projects**
- Augmented the workflow framework from the Analytics4MD project to make use of LLNL's Dynamic and Asynchronous Data Streamliner (DYAD)
 - Created a test suite for DYAD based on workflows from the Analytics4MD project
 - Examined the performance of DYAD using our test suite and using the performance data collection tool PerfFlow Aspect

2021

Summer Graduate Computing Student Intern, *Center for Advanced Scientific Computing*, Lawrence Livermore National Laboratory, Livermore, CA

Mentors Dr. Olga Pearce, Dr. Stephanie Brink

- Projects**
- Developed a data analysis workflow that uses profiles of LLNL's MARBL multi-physics simulation tool to predict what compiler MARBL should be built with to get the best performance for a particular simulation workload

2020

Summer Undergraduate Computing Student Intern, Lawrence Livermore National Laboratory, Livermore, CA

Mentors Dr. Stephanie Brink, Dr. Abhinav Bhatele, Dr. Olga Pearce, Dr. Todd Gamblin

- Projects**
- Designed a new graph-based filtering query language for the Hatchet data analysis tool to enable relationship-based analysis of profiling data
 - Implemented the query language and integrated it into Hatchet's data analysis capabilities
 - Used the query language to perform novel analysis of the performance of different MPI calls in HPC benchmark applications
 - Presented the work associated with this internship at the ACM Student Research Competition at the annual Supercomputing (SC) conference, where it won the 1st place award in the Undergraduate category
 - Presented an expanded version of this work at the ACM Student Research Competition Grand Finals

2017
2018

HERE Intern, *Neutron Scattering Division*, Oak Ridge National Laboratory, Oak Ridge, TN

Mentors Dr. Jiao Lin and Dr. Garrett Granroth

- Projects**
- Developing GUIs using Python and JavaScript for user data analysis (<https://github.com/scikit-beam/ipywe>)
 - Updating code to support both Python 2.7 and 3 (<https://github.com/mcvine>)
 - Developing a Python package to convert XML representations of constructive solid geometry into OpenSCAD code for visualization and 3D-printing purposes (<https://github.com/mcvine/SCADGen>)
 - Parallelizing a Monte-Carlo neutron ray-tracing software package with CUDA (<https://github.com/mcvine/McVineGPU>)

2015
2016

Math, Science, and Computer Science Thesis Student, Oak Ridge High School, Oak Ridge, TN

Mentor Dr. Len Gray

- Projects**
- Developing a boundary integral method approach to nonlinear fracture mechanics

Honors and Scholarships

2024

SC24 ACM Student Research Competition: Participant (Graduate Category), *ACM/IEEE International Conference for High Performance Computing, Networking, Storage, and Analysis (SC)*, Atlanta, GA

2024

ACM Travel Grant, *to participate in the ACM Student Research Competition at the ACM/IEEE International Conference for High Performance Computing, Networking, Storage, and Analysis (SC)*, Atlanta, GA

2024

IPDPS PhD Forum, *IEEE International Parallel & Distributed Processing Symposium*, San Francisco, CA

2023

SC23 ACM Student Research Competition: Participant (Graduate Category), *ACM/IEEE International Conference for High Performance Computing, Networking, Storage, and Analysis (SC)*, Denver, CO

2023

ACM Travel Grant, to participate in the ACM Student Research Competition at the *ACM/IEEE International Conference for High Performance Computing, Networking, Storage, and Analysis (SC)*, Denver, CO

2022

NSF Travel Scholarship, to attend the *IEEE International Conference on e-Science*, Salt Lake City, UT

2020

SC20 ACM Student Research Competition: 1st Place (Undergraduate Category), *ACM/IEEE International Conference for High Performance Computing, Networking, Storage, and Analysis (SC)*, Virtual (originally Atlanta, GA)

2020

Tennessee Fellowship for Graduate Excellence, *University of Tennessee*, Knoxville, TN

2016

2020

Member of Chancellor's Honors Program, *University of Tennessee*, Knoxville, TN

2016

2020

Member of Cook Grand Challenge Engineering Honors Program, *University of Tennessee*, Knoxville, TN

2016

2020

Volunteer Scholarship (Highest Level), *University of Tennessee*, Knoxville, TN

2016

2020

Hope Scholarship, *State of Tennessee*

2016

2020

Herbert and Lillian Duggan Engineering Scholarship, *University of Tennessee*, Knoxville, TN

2017

2020

Carol and Malcom Bayless Merit Scholarship, *University of Tennessee*, Knoxville, TN

2018

2020

Charles Weaver Memorial Scholarship, *University of Tennessee*, Knoxville, TN

Papers

- [1] O. Pearce, J. Burmark, R. Hornung, B. Bogale, **I. Lumsden**, M. McKinsey, D. Yokelson, D. Boehme, S. Brink, M. Taufer, and T. Scogland, "RAJA Performance Suite: Performance Portability Analysis with Caliper and Thicket," in *Proceedings of the SC '24 Workshops of The International Conference on High Performance Computing, Network, Storage, and Analysis*, 2024.
- [2] H. Devarajan, **I. Lumsden**, C. Wang, K. Georgouli, T. Scogland, J. S. Yeom, and M. Taufer, "DYAD: Locality-Aware Data Management for Accelerating Deep Learning Training," in *2024 IEEE 36th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD)*, 2024. DOI: 10.1109/SBAC-PAD63648.2024.00010.

- [3] **I. Lumsden**, H. Devarajan, J. Marquez, S. Brink, D. Boehme, O. Pearce, J. S. Yeom, and M. Taufer, "Empirical Study of Molecular Dynamics Workflow Data Movement: DYAD vs. Traditional I/O Systems," in *2024 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, 2024. DOI: 10.1109/IPDPSW63119.2024.00111.
- [4] S. Brink, M. McKinsey, D. Boehme, C. Scully-Allison, **I. Lumsden**, D. Hawkins, T. Burgess, V. Lama, J. Lüttgau, K. E. Isaacs, M. Taufer, and O. Pearce, "Thicket: Seeing the performance experiment forest for the individual run trees," in *Proceedings of the 32nd International Symposium on High-Performance Parallel and Distributed Computing*, 2023. DOI: 10.1145/3588195.3592989.
- [5] **I. Lumsden**, J. Luettgau, V. Lama, C. Scully-Allison, S. Brink, K. Isaacs, O. Pearce, and M. Taufer, "Enabling call path querying in hatchet to identify performance bottlenecks in scientific applications," in *2022 IEEE 18th International Conference on eScience (eScience)*, 2022. DOI: 10.1109/eScience55777.2022.00039.
- [6] S. Brink, **I. Lumsden**, C. Scully-Allison, K. Williams, O. Pearce, T. Gamblin, M. Taufer, K. Isaacs, and A. Bhatele, "Usability and performance improvements in hatchet," in *2020 IEEE/ACM International Workshop on HPC User Support Tools (HUST) and Workshop on Programming and Performance Visualization Tools (ProTools)*, 2020, pp. 49–58. DOI: 10.1109/HUSTProtocols51951.2020.00013.
- [7] J. Bilheux, H. Bilheux, J. Lin, **I. Lumsden**, and Y. Zhang, "Neutron imaging analysis using jupyter python notebook," *Journal of Physics Communications*, vol. 3, no. 8, 2019.
- [8] J. Lin, F. Islam, G. Sala, **I. Lumsden**, H. Smith, M. Doucet, M. B. Stone, D. L. Abernathy, G. Ehlers, J. F. Ankner, and G. E. Granroth, "Recent developments of mcvine and its applications at sns," *Journal of Physics Communications*, vol. 3, no. 8, 2019.
- [9] **I. Lumsden**, L. Gray, and W. Ye, "Grid-based volume integration for elasticity: Traction boundary integral equation," *Engineering Fracture Mechanics*, vol. 176, pp. 74–82, 2017.

Tutorials

- [10] D. Yokelson, S. Brink, D. Boehme, **I. Lumsden**, M. Taufer, and O. Pearce, "Reproducible Performance Measurement and Analysis for High-Performance Computing Applications," presented at the 34th ACM International Symposium on High-Performance Parallel and Distributed Computing, 2025, to appear.
- [11] **I. Lumsden**, N. Tan, D. Milroy, V. Sochat, and M. Taufer, "Introduction to Flux: Next-Generation Resource Management for Exascale Workflows and Job Scheduling," presented at the Joint Laboratory for Extreme-Scale Computing Workshop (JLESC), 2024.

Posters

- [12] **I. Lumsden**, O. Pearce, J. S. Yeom, T. Scogland, and M. Taufer, "Benchmarking and Modeling of Producer-Consumer Data Movement Performance in Scientific Workflows," presented at the International Conference for High Performance Computing, Networking, Storage and Analysis (SC), 2024.
- [13] **I. Lumsden**, H. Devarajan, J. Marquez, S. Brink, D. Boehme, O. Pearce, J. S. Yeom, and M. Taufer, "Empirical Study of Molecular Dynamics Workflow Data Movement: DYAD vs. Traditional I/O Systems," presented at the PhD Forum at 38th IEEE International Parallel & Distributed Processing Symposium, 2024.
- [14] **I. Lumsden** and M. Taufer, "Enabling Transparent, High-Throughput Data Movement for Scientific Workflows on HPC Systems," presented at the Joint Laboratory for Extreme-Scale Computing Workshop (JLESC), 2024.

- [15] **I. Lumsden** and M. Taufer, "Enabling Transparent, High-Throughput Data Movement for Scientific Workflows on HPC Systems," presented at the International Conference for High Performance Computing, Networking, Storage and Analysis (SC), 2023.
- [16] **I. Lumsden**, J. Luettgau, V. Lama, C. Scully-Allison, S. Brink, K. Isaacs, O. Pearce, and M. Taufer, "Identifying Performance Bottlenecks in Scientific Applications with Call Path Querying," presented at the Salishan Conference on High Speed Computing, 2023.
- [17] J. S. Yeom, D. H. Ahn, **I. Lumsden**, J. Luettgau, S. Caino-Lores, and M. Taufer, "Ubique: A new model for untangling inter-task data dependence in complex hpc workflows," presented at the 2022 IEEE 18th International Conference on e-Science (e-Science), 2022. DOI: 10.1109/eScience55777.2022.00068.
- [18] **I. Lumsden** and M. Taufer, "Enabling Graph-Based Profiling Analysis Using Hatchet," presented at the International Conference for High Performance Computing, Networking, Storage and Analysis (SC), 2020.

Professional Services

- 2024

●
Lead Student Volunteer, *ACM/IEEE International Conference for High Performance Computing, Networking, Storage, and Analysis (SC)*, Atlanta, GA
- 2023

●
Student Volunteer, *ACM/IEEE International Conference for High Performance Computing, Networking, Storage, and Analysis (SC)*, Denver, CO
- 2022

●
Lead Student Volunteer, *ACM/IEEE International Conference for High Performance Computing, Networking, Storage, and Analysis (SC)*, Dallas, TX
- 2021

●
Lead Student Volunteer (SCALeR), *ACM/IEEE International Conference for High Performance Computing, Networking, Storage, and Analysis (SC)*, St. Louis, MO
- 2020

●
Student Volunteer, *ACM/IEEE International Conference for High Performance Computing, Networking, Storage, and Analysis (SC)*, Virtual (originally Atlanta, GA)
- 2019

●
Student Volunteer, *ACM/IEEE International Conference for High Performance Computing, Networking, Storage, and Analysis (SC)*, Denver, CO
- 2019

●
Volunteer, *FIRST Robotics Smokey Mountain Regional Competition*, Knoxville, TN

Skills

Programming Languages	C, C++, Python, Rust, JavaScript, TypeScript, Kotlin, Java
Parallelism Libraries	MPI, OpenMP, RAJA, CUDA, ROCm, PThreads
Data Analysis Tools	Thicket, Hatchet, NumPy, SciPy, Matplotlib, Pandas, Jupyter, ipywidgets, Polars, Apache Arrow, Apache Spark, Scikit-Learn, Keras, Tensorflow

Build Systems and Packaging Tools CMake, Spack, Autotools, Cargo, Meson

Testing and Continuous Integration Tools PyTest, GitHub Actions, GoogleTest, Travis CI

Shells Bash, Bourne, Zsh

Other Flux, Caliper, UCX, LaTeX, SQL, Graph Databases, LLVM, HTML, CSS

Languages

English Native

French Good Understanding

6 years of study in Middle and High School

Professional Memberships

IEEE Student Member since April 2017

ACM Student Member since August 2019