Budvin Edippuliarachchi

PhD Candidate in Computer Science, Tufts University

budvin.edippuliarachchi@tufts.edu | budvinchathura.github.io | Google Scholar
github.com/budvinchathura | linkedin.com/in/budvin-chathura

Research Interests

I study the performance and accuracy of numerical methods for partial differential equations on modern accelerator architectures and distributed systems, and design methods and software that make them fast and scalable: compressed sparse operators for finite element solvers, scalable mesh partitioning for adaptive mesh refinement, and domain-specific languages that generate efficient solver code through modern compiler infrastructure (MLIR), with formally verified code generation as a foundation.

Education

Tufts University 2023 – expected Dec 2028
Ph.D. in Computer Science (Candidate) Medford, MA

- Dissertation: *Accelerating Numerical Methods on Modern Architectures*.
- Advisor: Prof. Hari Sundar. GPA: 3.96/4.00.
- M.S. in Computer Science, earned en route (2026).
- Coursework: Numerical Analysis (undergraduate and graduate); Numerical Linear Algebra; Numerical Methods for Large-Scale Scientific Computing; Scientific Machine Learning; Compilers; Programming Languages; Software Performance Engineering (MIT).

University of Moratuwa 2018 – 2022
B.Sc. Eng. (Hons), specialized in Computer Science & Engineering Sri Lanka

- CGPA: 4.12/4.20 (3.97/4.00 equivalent); ranked 1st of 128 students; Dean's List in all eight semesters.
- Gold Medal for the best overall GPA in the Computer Science & Engineering specialization.

Publications

Peer-Reviewed Conference Papers

[C2] **Budvin Edippuliarachchi**, David Van Komen, and Hari Sundar. AMRaCut: Scalable Partitioning for Adaptive Mesh Refinement. In *Proc. International Conference for High Performance Computing, Networking, Storage and Analysis (SC '25)*, 2025. doi:10.1145/3712285.3759863
github.com/paralab/AMRaCut

[C1] **Budvin Edippuliarachchi**, Damika Gamlath, Ruchin Amaratunga, Gunavaran Brihadiswaran, and Sanath Jayasena. Quill: A Memory Efficient k-mer Counting and k-mer Querying Tool for Commodity Clusters. In *Proc. 14th International Conference on Bioinformatics and Biomedical Technology (ICBBT '22)*, 2022. doi:10.1145/3543377.3543389

Posters and Student Research Competitions

[P1] **Budvin Edippuliarachchi**, Ponnuswamy Sadayappan, and Hari Sundar. Index Compression for Finite Element Sparse Matrix-Vector Multiplication. *To appear in ACM Student Research Competition at SC '26*, 2026.

Research Experience

Graduate Research Assistant 2023 – present
Tufts University, advised by Prof. Hari Sundar Medford, MA

CLASP: Index Compression for Finite Element SpMV [P1]; full paper under review

- Uses mesh connectivity to split the FEM matrix into small blocks with 8-bit column indices and 16-bit row pointers, about 4× less index storage than CSR.
- Proved the encoding is compressive despite node duplication, with overhead bounded by $1 + O(B^{-1/d})$ for

d-dimensional meshes and block size B.

- Designed two GPU kernels (CSR with asynchronous shared-memory buffering; prefix-sum-based COO): 1.2–1.7× speedup (geometric mean 1.38×) over cuSPARSE across three generations of NVIDIA GPUs, and up to 2× over Intel MKL with a SIMD CPU kernel.
- Co-advised by Prof. Ponnuswamy Sadayappan. Ongoing: a mesh-free variant that derives clusters from the matrix graph alone, extending the method to general sparse matrices.

Finch: A Domain-Specific Language for PDE Solvers ongoing

- Active contributor to Finch, a DSL that generates finite element and finite volume solver code from a high-level problem description.
- Leading development of the next version (not yet public): an end-to-end MLIR compilation pipeline that maps Finch operators to MLIR for portable performance on NVIDIA/AMD GPUs and x86/ARM CPUs via IREE.
- Adds support for coupled multiphysics systems, nonlinear solvers, and flexible time stepping, with a scripting-style API for rapidly prototyping coupling schemes, time integrators, and solvers.

AMRaCut: Scalable Partitioning for Adaptive Mesh Refinement [C2]

- Distributed graph partitioner for AMR combining parallel label propagation and graph diffusion. It reuses the previous partition, so communication stays point-to-point between neighbors without global collectives.
- 5–10× faster than ParMETIS and PT-Scotch on 20,000+ cores, with partition quality comparable to multilevel methods and better than space-filling curves.
- Validated on a downstream finite element SpMV: despite up to 2× load imbalance, performs on par with ParMETIS/PT-Scotch partitions and outperforms spatially ordered ones.

Verified Code Generation for Finite Element Solvers ongoing

- With the Tufts Programming Languages group (TuPL): formalizing, in the Lean theorem prover, the translation of FEM weak forms into quadrature-based element integration code with semantic preservation. Manuscript in preparation.

Student Researcher, ALICE O2 Software Profiling 2022
CERN, in collaboration with the University of Moratuwa

- Profiled ALICE O2 Run 3 workflows with perf and Intel VTune to identify CPU and memory bottlenecks for the ALICE grid.

Undergraduate Researcher, *k*-mer Counting on Commodity Clusters 2021 – 2022
University of Moratuwa Sri Lanka

- Designed Quill, an MPI-based k-mer counter with a pipelined producer–consumer workflow and dynamic load balancing for clusters without high-performance networks or file systems. [C1]

Open-Source Software

Finch `github.com/paralab/Finch`
DSL and code generator for finite element and finite volume PDE solvers

AMRaCut `github.com/paralab/AMRaCut`
Distributed-memory mesh partitioner for adaptive mesh refinement

Quill `github.com/CSE-Optimizers/k-mer_counter`
Distributed k-mer counting for commodity clusters

Talks and Presentations

Highly Scalable Methods for Adaptive Mesh Refinement of Complex Domains Mar 2025
SIAM Conference on Computational Science and Engineering (CSE25) Fort Worth, TX

AMRaCut: Scalable Partitioning for Adaptive Mesh Refinement Nov 2025
SC '25, paper presentation St. Louis, MO

Teaching Experience

Graduate Teaching Assistant, Tufts University 2025
CS 40: Machine Structure & Assembly Language Programming

Teaching Assistant, University of Moratuwa 2021 – 2023
Theory of Computing; Data Structures & Algorithms; Programming Fundamentals

Professional Service

Junior Program Committee, ACM Intl. Conference on Supercomputing (ICS 2027) 2027
Selected for the inaugural Junior PC; review cycle upcoming

Outreach: algorithms workshops for IEEE student branches in Sri Lanka (3 sessions) 2021 – 2022

Honors and Awards

Gold Medal, best overall GPA in Computer Science & Engineering, University of Moratuwa 2022
Selected participant, Cornell, Maryland, Max Planck Pre-doctoral Research School (80 of 400+) 2022
IEEEXtreme global programming competition: 71st (2021) and 92nd (2020) worldwide 2020, 2021
Champions, Aces Coders v8.0 and UoJ Coders, national programming contests, Sri Lanka 2019, 2020
Bronze Medal, Sri Lanka Mathematics Olympiad (Junior) 2008

Industry Experience

Software Engineer, Expert Intelligence — AWS batch pipelines for AI on mass-spectrometry data 2023
Founding Member & DevOps Consultant, IXD Labs — deployment and monitoring 2021 – 2023
Software Engineering Intern, Mitra Innovation — event-driven AWS systems 2020 – 2021

Skills

Numerical Methods	Finite element methods, sparse linear algebra, mesh partitioning and adaptive mesh refinement
Programming	C, C++, Julia, Python, Lean, Java
HPC & Compilers	MPI, CUDA, SIMD vectorization, MLIR/IREE, perf, Intel VTune

References

Prof. Hari Sundar (advisor)
Associate Professor, Computer Science
Tufts University

Prof. Ponnuswamy Sadayappan
Professor, Kahlert School of Computing
University of Utah

Prof. Jeffrey Foster
Professor, Computer Science
Tufts University