Maryam Mehri Dehnavi

1. PERSONAL

Maryam Mehri Dehnavi
Department of Computer Science
University of Toronto
Pratt 398A, 6 King's College Rd.
Toronto, ON M5S 3G4, Canada

Mobile: +1-206-451-0307
mmehride@cs.toronto.edu
www.cs.toronto.edu/ mmehride

3. EMPLOYMENT

University of Toronto, Department of Computer Science — July 2022 to Date
Associate Professor with Tenure
Associate Chair of Research
Canada Research Chair in Parallel and Distributed Computing

Microsoft Research, Redmond, WA — November 2023 to Date
Consultant Researcher

University of Toronto, Department of Computer Science — July 2018 to June 2022
Assistant Professor (Tenure-track)
Canada Research Chair in Parallel and Distributed Computing

Rutgers University, Electrical and Computer Engineering — Sep. 2015 to July 2018
Assistant Professor (Tenure-track)

Massachusetts Institute of Technology, Computer Science — Feb. 2013 to Sep. 2015
Postdoctoral Researcher, Adviser: Professor Charles E. Leiserson

University of California Berkeley, Computer Science — Oct. 2011 Apr. 2012
Visiting Student Scholar, Adviser: Professor James Demmel

University of California Irvine, Computer Engineering — Jan. 2011 to Apr. 2011
Visiting Student Scholar, Adviser: Professor Jean-luc Gaudiot

4. DEGREES

McGill University, Montreal, Canada — 2008 to 2013
Ph.D., Electrical and Computer Engineering
Adviser: Professor Dennis Giannacopoulos
Thesis Topic: *Krylov subspace methods on Graphic Processing Units*

University of Calgary, Calgary, Canada — 2005 to 2007
M.Sc., Computer Engineering
Adviser: Professor Wessam M. Hassanien
Thesis Topic: *Characterizing and enhancing SMT clustered architectures*

Isfahan University of Technology, Isfahan, Iran — 2001 to 2005
B.Sc., Electrical Engineering
Final Project Topic: *Carbon NanoTubes and applications*

Academic History and Awards

5. HONORS

Canada Research Chair in Parallel and Distributed Computing, Tier II.

Ontario Early Researcher Award 2021.

6 RESEARCH AWARDS

Ontario Early Researcher Award, 2021-2026, $200,000.

NSERC New Frontiers in Research Fund - Exploration, Neurocomputing, 2020-2022, $250,000.

Canada Research Chair, Tier II, 2019-2024, $600,000.

XSEDE and SOSCIP advanced computing platform awards, 2020-2022, $40,000.

NSERC Discovery: Automatic Matrix Code Optimization for Performance and Portability, 2019-2024, $165,000.

IBM CAS project: Compilers for sparse algebra, 2020-2023, $100,000.

Connaught New Researcher Award, 2019-2020, $20,000.

Discovery Launch Supplement, 2019-2020, $12,500.

NSF: SMALL: Communication-Efficient Distributed Algorithms for Machine learning, Primary PI, 2018-2021, $610,000.

NSF: CRII: Performance-in-Depth Sparse Solvers for Heterogeneous Parallel Platforms, 2017-2019, $229,000.

Adobe Research Project Funding: Sympiler: Transforming Sparse Matrix Codes. 2019-2020, $8500.

Grand Final Winners of 2017 ACM Student Research Competitions. Acceptance rate: 0.3% (1 out of 330 from all submissions from 27 ACM conferences)

FQRNT postdoctoral fellowship, $70,000, 2013-2015.

NSERC (*Natural Sciences and Engineering Research Council of Canada*) Postdoctoral Fellowship, $80,000, 2013-2015.

7. Awards by Students

First-Place in Student Research Competitions PACT 2022: Student (Laird).

Honourable mention for the 2022 SIGHPC Dissertation Award (Cheshmi).

The 2020 ACM-IEEE CS George Michael Memorial HPC fellowship, $5,000, Student (Cheshmi).

Adobe Research Fellowship Winner 2018, Student (Cheshmi), $13,000.

Wolfond Fellowship, Student (Cheshmi), 2019, $10,000.

MITACS accelerate postdoc and graduate trainee awards, 2019-2022, $80,000.

NSERC USRA undergraduate trainee awards, 2019-2022, $48,000.

First-Place in Student Research Competitions CGO 2017: Student (Cheshmi).

James Leroy Potter Award 2018 for research excellence: Student (Blanco).

SIGPLAN PAC funding, 2020, Student (Cheshmi and Liu), $2,000.

Best Poster Award, CGO 2017: Student (Cheshmi).

Outstanding Poster Paper, HPCS 2018.

NSF Travel Grant for Cluster17 $1,400 and CGO 2017 $1,700 Students: Liu and Cheshmi.

Adobe Research Fellowship Finalist 2017: Student (Cheshmi).

8. PROFESSIONAL AFFILIATIONS AND ACTIVITIES

Associate Editor of JPDC (2019-2021): Journal of Parallel and Distributed Computing

General Chair of PPOPP 2023

Co-chair, Algorithms/Experiments (2021/2022): International Parallel and Distributed Processing Symposium (IPDPS)

Co-chair, Multidisciplinary (2022): International Conference on Parallel Processing (ICPP)

Co-chair, Machine Learning and HPC (2019): ACM/IEEE Intl. Conf. for High Performance Computing, Networking, Storage and Analysis (SC)

Group Chair, SIAM Conference on Parallel Processing for Scientific Computing (SIAM PP2020)

Chair, MLDay and posters at ISC20, ISC21, ISC22

Vice-chair, Algorithms (2018): ACM/IEEE Intl. Conf. for High Performance Computing, Networking, Storage and Analysis (SC)

Organizing Committee, SIAM Conference on Parallel Processing for Scientific Computing (SIAM PP2020)

Co-chair, CHIUW 2019

Scholarly and Professional Work

10. REFEREED CONFERENCE PUBLICATIONS

C-39 M. Mozaffari, S. Li, Z. Zhang, and M. Mehri Dehnavi. MKOR: Momentum-Enabled Kronecker-Factor-Based Optimizer Using Rank-1 Updates. Accepted at *Thirty-seventh Conference on Neural Information Processing Systems, Neurips,* 2023.

C-38 L. Wilkinson, K. Cheshmi, and M. Mehri Dehnavi. Sparse Register Tiling for Sparse Neural Networks. *Programming Language Design and Implementation, PLDI,* pp. 1995-2020, Vol 7, 2023.

C-37 K. Cheshmi, M. Strout, and M. Mehri Dehnavi. Runtime Composition of Iterations for Fusing Loop-carried Sparse Dependence. Accepted at *The International Conference for High Performance Computing, Networking, Storage and Analysis, SC*, 2023.

C-36 A. Jangda, S. Maleki, M. Mehri Dehnavi, M. Musuvathi, O. Saarikivi. A Framework for Fine-Grained Synchronization of Dependent GPU Kernels, conditionally accepted at *IEEE/ACM International Symposium on Code Generation and Optimization, CGO*, 2024.

C-35 Z. Centinic, K. Cheshmi, and M. Mehri Dehnavi. Vectorizing Sparse Matrix Computations with Partially-Strided Codelets. *The International Conference for High Performance Computing, Networking, Storage and Analysis, SC*,pp. 1-15, 2022.

C-34 B. Liu, A. Laird, and M. Mehri Dehnavi. Combining Run-time Checks and Compile-time Analysis to Improve Control Flow Auto-Vectorization. *PACT*, pp. 439-450, 2022.

C-33 S. Soori, B. Mu, B. Ucar, M. Gurbuzbalaban, and M. Mehri Dehnavi. HyLo: A Hybrid Low-Rank Natural Gradient Descent Method. *The International Conference for High Performance Computing, Networking, Storage and Analysis, SC*, pp. 1-16, 2022.

C-32 B. Zare, B. Liu, M. Strout, and M. Mehri Dehnavi. HDAGG: Hybrid Aggregation of Loop-carried Dependence Iterations in Sparse Matrix Computations. Accepted at the *IEEE International Parallel and Distributed Processing Symposium, IPDPS*, pp. 1217-1227, 2022.

C-31 A. Laird, B. Liu, and M. Mehri Dehnavi. Auto-translation of Non-Affine Code. *International Conference on Parallel Architectures and Compilation Techniques, PACT*, 2022. ***PACT Winner of ACM SRC 2022***.

C-30 K. Cheshmi, M. Strout, and M. Mehri Dehnavi. Composing Loop-carried Dependence with Other Loops (Link). *The 25th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPOPP*, 2022.

C-29 B. Ucar, S. Soori, M. Mehri Dehnavi, and M. Gurbuzbalaban. L-DQN: An Asynchronous Limited-Memory Distributed Quasi-Newton Method. *IEEE Conference on Decision and Control, CDC, pp. 2386-2393, 2021*.

C-28 K. Cheshmi, S. Kamil, D. Kaufmann, and M. Mehri Dehnavi. NASOQ: Numerically Accurate Sparsity-Oriented QP Solver. *ACM Transactions on Graphics (TOG), 39(4), 96-1. Presented at SIGGRAPH*, 2020.

C-27 S. Soori, B. Ucar, M. Gurbuzbalaban, and M. Mehri Dehnavi. ASYNC: A Cloud Engine with Asynchrony and History for Distributed Machine Learning. *IEEE International Parallel and Distributed Processing Symposium, pp. 429-439, IPDPS*, 2020.

C-26 S. Soori, K. Mischenko, A. Mokhtari, M. Mehri Dehnavi, and M. Gurbuzbalaban. DAve-QN: A Distributed Averaged Quasi-Newton Method with Local Superlinear Convergence Rate. *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics, 108:1965-1976, AISTATS*, 2020.

C-25 K. Cheshmi, B. Liu, S. Soori, M. Strout, and M. Mehri Dehnavi. MatRox: Modular approach for improving data locality in Hierarchical (Mat)rix App(Rox)imation. *Proceedings of the 25th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, pp. 389-402, PPOPP, 2020*.

C-24 M. Soltan Mohammadi, T. Yuki, K. Cheshmi, E. C. Davis, M. Hall, M. Mehri Dehnavi, P. Nandy, C. Olschanowsky, A. Venkat, M. Strout. Sparse Computation Data Dependence Simplification for Efficient Compiler-Generated Inspectors. *Programming Language Design and Implementation, pp. 594-609, PLDI*, 2019.

C-23 K. Cheshmi, S. Kamil, M. Strout, and M. Mehri Dehnavi. ParSy: Compile-time Inspection and Transformation of Sparse Matrix Computations for Parallelism, *The International Conference for High Performance Computing, Networking, Storage and Analysis, pp. 779-793, SC*, 2018.

C-22 Z. Blanco, B. Liu, and M. Mehri Dehnavi. CSTF: Large-Scale Sparse Tensor Factorizations on Distributed Platforms. *The 47th International Conference on Parallel Processing, pp. 1-10, ICPP*, 2018.

C-21 S. Soori, A. Devarakonda, Z. Blanco, J. Demmel, M. Gurbuzbalaban, and M. Mehri Dehnavi. Reducing Communication in Proximal Newton Methods for Sparse Least Squares Problems. *The 47th International Conference on Parallel Processing, pp. 1-10, ICPP*, 2018.

C-20 K. Cheshmi, S. Kamil, M. Strout, and M. Mehri Dehnavi. Sympiler: Transforming Sparse Matrix Codes by Decoupling Symbolic Analysis. *The International Conference for High Performance Computing, Networking, Storage and Analysis, 13:1–13:13, SC*, 2017. ***Grand Final Winner of ACM SRC 2017***.

C-19 B. Liu, C. Wen, A. Sarwate, and M. Mehri Dehnavi. A Unified Optimization Approach for Sparse Tensor Operations on GPUs. pp. 47-57, *IEEE CLUSTER17*, 2017.

C-18 K. Cheshmi, L. Cheshmi, and M. Mehri Dehnavi. Sparsity-Aware Storage Format Selection. *International Conference on High Performance Computing and Simulation, HPCS*, 2018. ***Outstanding poster paper winner***.

C-17 M. Soltan Mohammadi, K. Cheshmi, M. Venkat, T. Yuki, M. Mehri Dehnavi, and M. Strout. Index-Array Properties for Data Dependence Analysis. *The 30th International Workshop on Languages and Compilers for Parallel Computing, LCPC*, 2018.

C-16 G. Salles-Loustau, L. Garcia, P. Sun, M. Mehri Dehnavi, S. Zonouz. Power grid safety control via fine-grained multi-persona programmable logic controllers, *IEEE International Conference on Smart Grid Communications (SmartGridComm)*, 2017.

C-15 K. Cheshmi, L. Cheshmi, and M. Mehri Dehnavi. Sparsity-Aware Storage Format Selection. *The 2018 International Conference on High Performance Computing & Simulation*, Orleans, France, 2018.

C-14 A. Shukla, Yue Wu, S. Zonouz, and M. Mehri Dehnavi. Fault-tolerant iterative solvers with adaptive reliability. *IEEE Conference on Electromagnetic Field Computation*, Miami, 2016.

C-13 K. Cheshmi, S.Zonouz, and M. Mehri Dehnavi. AXB: A domain-specific compiler for direct solvers. *IEEE Conference on Electromagnetic Field Computation*, Miami, 2016.

C-12 M. Mehri Dehnavi, J. Demmel, and D. Fernández. Communication-avoiding sparse approximate inverse preconditioners. *SIAM Conference on Linear Algebra (SIAM-LA)*, 2015.

C-11 Y. You, D. Bader, and M. Mehri Dehnavi. Designing a heuristic cross-architecture combination for breadth-first search. *The 43th International Conference on Parallel Processing, pp. 70–79, (ICPP)*, 2014.

C-10 Y. You, S. Song, H. Fu, A. Marquez, G. Yang, K. Barker, K. Cameron, M. Mehri Dehnavi, and A. Randles. MIC-SVM: Designing a highly efficient support vector machine for advanced modern multi-core and many-core architectures. *Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS)*, pp. 809-818, 2014.

C-9 M. Mehri Dehnavi, D. Fernández, and D. Giannacopoulos. Finite element sparse matrix vector multiplication on GPUs. *IEEE Conference on Computational Electromagnetics (COMPUMAG)*, 1082–1084, 2009. ***Best paper finalist***.

C-8 M. Mehri Dehnavi, J. Demmel, and D. Giannacopoulos. Communication-avoiding algorithms on GPUs. *IEEE Conference on Electromagnetic Field Computation (CEFC)*, 2012.

C-7 M. Mehri Dehnavi, D. Fernández, and D. Giannacopoulos. Accelerating sparse approximate inverse preconditioners based on matrix entries on GPUs. *IEEE Conference on Computational Electromagnetics (COMPUMAG)*, 2011.

C-6 D. Fernández, J. Zambrano, M. Mehri Dehnavi, Y. El-Kurdi, and D. Giannacopoulos. Accelerating the convergence of the FEMSES method using multi-grid techniques. *XII International Congress on Numerical Methods in Engineering and Applied Sciences*, 2014.

C-5 D. Fernández, M. Mehri Dehnavi, and D. Giannacopoulos. Alternate approach to FEM for parallel processing. *IEEE Conference on Computational Electromagnetics (COMPUMAG)*, 2011.

C-4 M. Mehri Dehnavi, D. Fernández, and D. Giannacopoulos. Enhancing the performance of conjugate gradient solvers on GPUs. *IEEE Conference on Electromagnetic Field Computation (CEFC)*, 2010.

C-3 M. Mehri Dehnavi and D. Giannacopoulos. Enhancing the performance of electromagnetic applications on clustered architectures. *IEEE Conference on Electromagnetic Field Computation (CEFC)*, 2008.

C-2 M. Mehri Dehnavi and W. Hassanein. A thread specific load balancing technique for a clustered SMT architecture. *Proceeding of Canadian Conference on Electrical and Computer Engineering (CCECE)*, 948–951, 2007.

C-1 M. Mehri Dehnavi and W. Hassanein. A clustered SMT architecture for scalable embedded processors. *Practical Real World Technologies for Communications and Embedded Platforms (PRWT)*, 201–203, 2007.

11. Refereed Journal Publications

J-15 M. Pham, A. Angiulli, M. Mehri Dehnavi, R. Chhabra. From Brain Models to Robotic Embodied Cognition: How Does Biological Plausibility Inform Neuromorphic Systems? *Brain Sciences, Vol. 13, Issue 9, 2023.*

J-14 M. Asgari, J. Lindsey, B. Tolson, M. Mehri Dehnavi. Development of a Knowledge-Sharing Parallel Computing Approach for Calibrating Distributed Watershed Hydrologic Models. Accepted at *Environmental modelling & Software Journal, 2023.*

J-13 B. Can, S. Soori, N. Aybat, M. Mehri Dehnavi, M. Gurbuzbalaban, Randomized Gossiping with Effective Resistance Weights: Performance Guarantees and Applications, IEEE Transactions on Control of Network Systems, doi: 10.1109/TCNS.2022.3161201, TCNS, 2022.

J-12 M. Asgari, J. Lindsey, B. Tolson, M. Mehri Dehnavi. A Review of Parallel Computing Applications in Calibration of Watershed Hydrologic Models. *Environmental modelling & Software Journal, vol. 151, issn 1364-8152, 2022.*

J-11 K. Cheshmi, S. Kamil, D. Kaufmann, and M. Mehri Dehnavi. NASOQ: Numerically Accurate Sparsity-Oriented QP Solver. *ACM Transactions on Graphics (TOG) , 39(4), 96-1*, 2020.

J-10 E. Palamadi, M. Mehri Dehnavi, and Charles Leierson. Autotuning divide-and-conquer stencil computations. Concurrency and Computation: Practice and Experience - Decision, 10.1002/cpe.4127, Volume 29, Issue 17:1-16, 2017.

J-9 Y. El-Kurdi*, M. Mehri Dehnavi*, W. Gross, and D. Giannacopoulos. Parallel finite element technique using Gaussian belief propagation. *Computer Physics Communications*, 193: 38-48, 2015. (* equal contribution)

J-8 Y. You, H. Fu, S. Song, M. Mehri Dehnavi, L. Gan, X. Huang, and G. Yang. Evaluating multi-core and many-core architectures through accelerating the three-dimensional Lax Wendroff correction stencil. *International Journal of High Performance Computing Applications (IJHPCA)*, 28(3), 301-318, 2014.

J-7 M.B. Qureshi, M. Mehri Dehnavi, et. al. Survey on grid resource allocation mechanisms. *Journal of Grid Computing (JGC)*, 399–441, 2014.

J-6 M. Mehri Dehnavi, D. Fernández, J.L. Guadiot, and D. Giannacopoulos. Parallel sparse approximate inverse preconditioning on graphic processing units. *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 24(9):1852–1862, 2013.

J-5 M. Mehri Dehnavi, Y. El-Kurdi, J. Demmel, and D. Giannacopoulos. Communication-avoiding Krylov techniques on graphic processing units. *IEEE Transactions on Magnetics (TMAG)*, 49(5):1749–1752, 2013.

J-4 D. Fernández, M. Mehri Dehnavi, W. Gross, and D. Giannacopoulos. Alternate parallel processing approach for FEM. *IEEE Transactions on Magnetics (TMAG)*, 48(2):399–402, 2012.

J-3 M. Mehri Dehnavi, D. Fernández, and D. Giannacopoulos. Enhancing the performance of conjugate gradient solvers on graphic processing units. *IEEE Transactions on Magnetics (TMAG)*, 47(5):1162–1165, 2011.

J-2 M. Mehri Dehnavi, D. Fernández, and D. Giannacopoulos. Finite-element sparse matrix vector multiplication on graphic processing units. *IEEE Transactions on Magnetics (TMAG)*, 46(8):2982–2985, 2010.

J-1 M. Mehri Dehnavi and D. Giannacopoulos. Enhancing the performance of electromagnetic applications on clustered architectures. *IEEE Transactions on Magnetics (TMAG)*, 45(3):1340–1343, 2009.

12. Non-Refereed Publications

P-13 A. Laird and M. Mehri Dehnavi. Auto-translation of non-affine codes. *The 31st International Conference on Parallel Architectures and Compilation Techniques (PACT)-SRC track*, ***First place***, Chicago, 2022.

P-12 B. Liu, M. Mehri Dehnavi. Tensor Factorization on GPUs. *SIAM Conference on Parallel Processing*, 2020.

P-11 K. Cheshmi and M. Mehri Dehnavi. Decoupling Symbolic from Numeric in Sparse Direct Solvers. *The International Symposium on Code Generation and Optimization (CGO)-SRC track*, ***First place***, Austin, 2017.

P-10 S. Soori, R. Shah, M. Mehri Dehnavi, Avoiding Communication in Proximal Methods for Convex Optimization, *SIAM Parallel Processing, 2018*.

P-9 B. Liu, M. Mehri Dehnavi, Sparse Tensor Computations on Graphic Processing Units, *SIAM Parallel Processing, 2018*.

P-8 M. Mehri Dehnavi and D. Giannacopoulos. Fast preconditioning on GPUs. *High Performance Computing Symposium in medical sciences*, 2011.

P-7 M. Mehri Dehnavi and D. Giannacopoulos. Accelerating finite element sparse matrix vector multiplication on GPUs. *Centre de Recherche en Electronique Radiofrequence (CREER)*, 2010.

P-6 M. Mehri Dehnavi and D. Giannacopoulos. Enhancing the performance of clustered architectures. *6th Interdisciplinary Graduate Student Research Symposium (IGTRS)*, 2010.

P-5 M. Mehri Dehnavi. Krylov subspace techniques on graphic processing units. Ph.D. Thesis, McGill University, 2012

P-4 M. Mehri Dehnavi. Characterizing and enhancing SMT clustered architectures. M.Sc Thesis, University of Calgary, 2007

P-3 M. Mehri Dehnavi and W. Hassanein. Characterizing the performance of data base management systems on the Pentium 4 Hyper-Threaded Architecture. Technical Report, University of Calgary, 2006

P-2 M. Mehri Dehnavi and W. Hassanein. CSMT-SIM: A clustered simultaneous multithreaded architecture Simulator. Technical Report, University of Calgary, 2007

P-1 W. Hassanein, L. Rashid, M. Mehri Dehnavi and W. Hassanein. Characterizing the performance of data base management systems on the Pentium 4 hyper-threaded architecture. Technical Report, University of Calgary, 2006

13. Conference and Symposia Presentations

Restructuring computations patterns for performance, SIAM/CSE meeting, 2020 (roundable).

Decoupling Structure in Hierarchical Matrix Approximations, SIAM PP 2020 (roundable).

ParSy, International Conference for High Performance Computing, Networking, Storage and Analysis, SC, 2018 (paper).

13th Biennial IEEE Conference on Electromagnetic Field Computation. Greece, 2008 (paper).

17th Conference on the Computation of Electromagnetic Fields. Brazil, 2009 (paper).

14th Biennial IEEE Conference on Electromagnetic Field Computation. USA, 2010 (paper).

18th Conference on the Computation of Electromagnetic Fields. Australia, 2009 (paper).

6th Interdisciplinary Graduate Student Research Symposium. Canada, 2010 (paper).

High Performance Computing Symposium in Medical Sciences. Canada, 2011 (paper).

SIAM Linear Algebra Symposium, USA, 2015 (roundable).

14. Invited Talks/Lectures

Sparsity: from Algorithms to Systems, **Keynote speaker** at SIAM PP 2024, Baltimore, USA

Sparsity in Machine Learning Pipelines, **Invited Panel Speaker** at SC 2023. Denver, USA

Title: The power of less: Harnessing Sparsity for Performance Optimization, Invited, UT Austin, 2023, USA

Title: The power of less: Harnessing Sparsity for Performance Optimization, *Invited*, Georgia Tech, 2024, Atlanta, USA

Title: The power of less: Harnessing Sparsity for Performance Optimization, *Invited*, UC Berkeley, CA, 2023, USA

Accelerating Sparse Matrix Computations with Code Specialization, *Invited*, Compiler Techniques for Sparse Tensor Algebra (CSTA) 2023. Orlando, USA

HyLo: A Hybrid Low-Rank Second Order Method, virtual, 2022 *Invited*, (Texas Advanced Computing Center, the NSF ICICLE AI Institute).

Scalable Optimization Methods. Seattle, WA, 2023, *Invited*, (SIAM SIAM conference on optimization).

O'Reilly Conference: Invited talk on Efficient Machine Learning methods, 2022, *Invited*, (Virtual, worldwide).

Inspecting Irregular Computation Patterns to Generate Fast Code, 2021, *Invited*, (ETH SPCL).

Fast Code for Sparse Linear Algebra, 2021, *Invited*, (MIT CSAIL Fast Code seminar).

Panelist at the Faculty Panel discussion session on Mitacs programs, Arts & Science, 2019, *Invited*, (UofT Mitacs Panel).

Tensor Computations: Applications and Optimization, 2020, *Invited*, (Dagstuhl Seminar).

A Compiler for Tensor Algebra, ***Invited***, 2019, (Compiler Techniques for Sparse Tensor Algebra at MIT).

Code Generation for Sparse Linear Algebra. Boston, 2019, ***Invited*** (IFIP WG 2.11 on program generation).

Symbolic Decoupling is Sparse linear Algebra, Toronto, 2019, *Invited*, (the CCSE Center at University of Toronto).

Code Generation for Sparse Linear Algebra. McMaster, Canada, 2020, *Invited*, (SIAM Seminar Series).

Transforming Computation and Communication Patterns for High-Performance, Computer Science Department, NYU Courant Institute of Mathematical Sciences, USA, 2018, (department seminar series).

Transforming Computation and Communication Patterns for High-Performance, Computer Science Department, UNC Charlotte, USA, 2018, (department seminar series).

Transforming Computation and Communication Patterns for High-Performance, Computer Engineering Department, George Mason University, USA, 2018, (department seminar series).

Transforming Computation and Communication Patterns for High-Performance, Computer Engineering Department, UMass Amherst, USA, 2018, (department seminar series).

Transforming Computation and Communication Patterns for High-Performance, Computer Science Department, Waterloo University, Canada, 2018, (department seminar series).

Transforming Computation and Communication Patterns for High-Performance, Computer Science Department, University of Toronto, Canada, 2018, (department seminar series).

Transforming Computation and Communication Patterns for High-Performance, Computer Science Department, McMaster University, Canada, 2018.

Transforming Computation and Communication Patterns for High-Performance, Computer Science Department, Concordia University, Canada, 2018, (department seminar series).

Sympiler: Transforming Sparse Matrix codes. Temple University, USA, 2017, (department seminar series).

Sympiler: Transforming Sparse Matrix codes. Lawrence Berkeley National Lab, USA, 2017.

Rutgers University, ECE Department, USA, 2015, (department seminar series).

Algorithms and Architectures for Computational Science and Engineering. Invited lecturer, Italy, 2011, (PhD summer school in Padova University).

Center for Exascale Simulation of Plasma-Coupled Combustion. UIUC, USA, 2014 (department seminar series).

IBM Research Yorktown heights. USA, 2014.

Northeastern University. 2014 (department seminar series).

Qualcomm Canada Inc. Canada, 2012.

Samsung Research America. USA, 2011.

Teaching and Mentoring Experience

15.A. Undergraduate courses taught

University of Toronto, Canada

Instructor

Parallel Computing (CSC367)	Fall 2021
Parallel Computing (CSC367): 4.7/5.0	Winter 2020
Parallel Computing (CSC367): 4.7/5.0	Winter 2019

Rutgers University, USA

Instructor

Adv. Topics in Computer Engineering: Applied Parallel Computing: 4.9/5.0	Spring 2018
Adv. Topics in Computer Engineering: Cloud Computing: 4.8/5.0	Fall 2017
Adv. Topics in Computer Engineering: Applied Parallel Computing: 4.9/5.0	Spring 2017
Adv. Topics in Computer Engineering: Cloud Computing 4.6/5.00	Fall 2016
Adv. Topics in Computer Engineering: Cloud Computing: 4.84/5.0	Fall 2015

15.B. GRADUATE COURSES TAUGHT

University of Toronto, Canada

Instructor

Applications of Parallel and Distributed Computing (CSC2222)	Fall 2021
Applications of Parallel and Distributed Computing (CSC2222): 4.85/5.0	Fall 2019

16. THESIS SUPERVISED

University of Toronto, Canada

PhD students/Postdocs

- Kazem Cheshmi, Transforming Sparse Matrix Code for Performance, 2016-2021. (*now a Tenure track Assistant Professor at McMaster*) [Primary supervisor]
- Bangtian Liu, Optimization Tensor Computations in Parallel Processors, 2016-date. [Primary supervisor]
- Saeed Soori, Communication-Avoiding Machine Learning Algorithms, 2017-2022. (*Now at Huawawei as a senior machine learning research engineer*) [Primary supervisor]
- Amir Akbari (co-advised), 2018-date. [co-supervisor]
- Behrooz Zare, Privacy-Oriented Data Analytics, 2020-date. [Primary supervisor]
- Faraz Shasavan, Machine Learning Algorithms for Applications in Robotics, Fall22. [Primary supervisor]
- Mohammad Mozaffari, Sparse computations in ML, 2021-Date. [Primary supervisor]

MSC/MScAC students

- Zachary Centinic, Decoupling Precedence Constraint to Generate Parallel Specialized Code, 2019-2021. [Primary supervisor]
- Avery Laird, 2021-date. [Primary supervisor]
- Mushegh Shahinyan, 2021-date. [Primary supervisor]
- Adam Adli, Fall22. [Primary supervisor]
- Martin Du Pham, Neurocomputation of normal versus abnormal brain-body interaction, 2021-date. [co-supervisor]
- Yi Xiang, 2021-date. [Primary supervisor]
- Jafeer Uding, 2021-date. [Primary supervisor]
- Laura Walsh, Personalizing Energy Savings for Smart Thermostat Users, 2019. [Primary supervisor]
- Alicia Tang, 2020. [Primary supervisor]
- Jianda Chen, Sparse matrix methods for machine learning 2019-2020. [Primary supervisor]
- Lucas Wilkinson, Fast LDL solvers, 2020-date. [Primary supervisor]

Undergraduate researchers students

- Yongzheng Huang (USRA) (now Quantitative Trading Analyst at RBC Capital Market): 2019-2020
- Jianda Chen: 2019
- Vinit Jogani (now at google): 2019
- Robin Abrahamse: 2020
- Ray Zhang: 2020-2021
- Lawrence Ora (UTEA): 2020
- Junming Zhang (UTEA): 2020
- Ibrahim Hassan: 2020-2021
- Tasbir Rahman: Winter 2021

- Baorun Mu: Summer-2021-Date
- Shujian Qian: Summer 2021
- Tian Yu (UTEA): Summer 2021
- Kaiwen Wu (USRA): Summer 2021
- Ao Li: Winter 2022
- Yuzhe Hua, Winter 2022
- Jim Gao, Sep 2021-March 2022
- Kevin Chen, Sep 2021-March 2022
- Vedang Ashwin, Sep 2022-date
- Haoming Meng, April 2022-date
- Quang Thuy Le, Sep 2022-date
- Ryan Goldenberg, Sep 2022-date
- Smit Lunagaria, Oct 2022-date

Rutgers University, USA

PhD students

- Kazem Cheshmi
- Bangtian Liu
- Saeed Soori
- Amir Akbari (co-advised)

Alumni MSc students

- Thesis-Based: Aadtiya Shukla (now data scientist at IBM Research), Ke Xu
- Project-Based: Yue Wu, Eric Xu, Chaoran Fu, Yuanxi Li (now at Amazon)

Undergraduates

- Zachary Blanco (now: MIT Lincoln Labs and PhD student at UC SanDiego)
- Visiting Students: George Sakkas (now PhD a UCSD PhD student), Nikos Skilikas
- Alex Chan

Administrative Positions

17.A. Positions within the University

- *Associate Chair Research, 2023*
- *Grad Seminar Skills Coordinator, 2021*
- *UG admissions committee, 2021*
- Panelist at the Faculty Panel discussion session on Mitacs programs, Arts & Science, 2019.
- *UofT/RBC PostDoctoral & Graduate Fellowships Adjudication Committee, 2020*
- *Grad Affairs Committee co-chair, University of Toronto, Computer science, 2019-2020*
- *Grad Visit Day Committee Member, University of Toronto, Computer science, 2019-2020*
- *Industrial Relations Committee Member, University of Toronto, Computer science, 2019-2020*
- *Faculty Job Interview Panelist, University of Toronto, Computer science, 2020*
- *University of Toronto MScAC Admissions Committee, 2019*
- *University of Toronto PhD Admissions Committee, 2019-2020*
- *Rutgers ECE IEEE Club Faculty Advisor, 2016-2018*
- *Rutgers Graduate Admissions Committee, 2015-2018*
- *Rutgers IEEE Chapter Head, 2015-2018*
- *The New Jersey Junior Science & Humanities Symposium at Rutgers University–Reviewer and Mentor, 2017*
- *MIT Undergraduate Women's Mentoring Program–Organizer, 2015*
- *McGill Undergraduate Poster Tutorial–Mentor, 2013*

17.B. Program Committee Leadership Outside of University

- *General Chair of PPOPP 2023*
- *Associate Editor of JPDC (2019-2021): Journal of Parallel and Distributed Computing*

- *Co-chair, Multidisciplinary (2022): International Conference on Parallel Processing (ICPP)*
- *Co-chair, Algorithms (2020): IEEE International Parallel and Distributed Processing Symposium (IPDPS)*
- *Co-chair, Machine Learning and HPC (2019): ACM/IEEE Intl. Conf. for High Performance Computing, Networking, Storage and Analysis (SC)*
- *Poster session co-Chair, ISC22*
- *NSERC NFRF, Feb 2022*
- *NSERC Discovery, Feb 2023*
- *INRIA Evaluator, Feb 2022*
- *Group Chair, SIAM Conference on Parallel Processing for Scientific Computing (SIAM PP2020)*
- *Chair, MLDay at ISC20 and ISC21*
- *Vice-chair, Algorithms (2018): ACM/IEEE Intl. Conf. for High Performance Computing, Networking, Storage and Analysis (SC)*
- *Organizing Committee, SIAM Conference on Parallel Processing for Scientific Computing (SIAM PP2020)*
- *Co-chair, CHIUW 2019*

17.C. Program Committee and Reviewer Service outside of University

- *Program Committee Member, SC 2023*
- *Program Committee Member, ASPLOS 2023*
- *Program Committee Member, SPAA 2023, 2024*
- *Program Committee Member, PLDI 2023*
- *Program Committee Member, CHIUW 2023*
- *Program Committee Member, SC 2022*
- *Program Committee Member, IPDPS 2022*
- *Program Committee Member, PPOPP 2022*
- *Inria's evaluation Committee Member, 2021*
- *Committee Member, IBM CASCON 2021*
- *Program Committee Member, ACDA 2021*
- *Program Committee Member, ISC 2021*
- *Program Committee Member, DATE 2020*
- *Program Committee Member, SC 2020, 2021*
- *Program Committee Member, ICPP 2020, 2021*
- *Program Committee Member, IPDPS 2020*
- *Program Committee Member, CHIUW 2020, 2021*
- *Program Committee Member, ICCD 2020*
- *Program Committee Member, GPCE 2019*
- *Program Committee Member, CHIUW 2019*
- *Best Paper Committee, SC 2019*
- *Reviewer for Mitacs Accelerate, 2019*
- *Primary Program Committee Member, IPDPS 19*
- *NSERC CFI Reviewer, 2018, 2019*
- *NSERC CRD Reviewer, 2018, 2019*
- *INTERMAG 2020 Reviewer*
- *NSF Panelist, CISE, March 2019*
- *NSF Panelist, CISE, March 2018*
- *Rutgers ECE IEEE Club Faculty Advisor*
- *Workshops Program Committee Member, SC 17*
- *Program Committee Member, SC 17*
- *NSF Panelist, CISE, March 2017*
- *NSF Panelist, CISE, Feb 2016*
- *NSF Panelist, CISE, March 2016*
- *Program Committee Member, SC 16*
- *Journal of Signal Processing Systems Computing and Visualization in Science, 2017*

- *Canada Foundation for Innovation, 2017*
- *ACM Transactions on Graphics, 2020*
- *SIGGRAPH, 2020*
- *Journal of computation Science, 2020*
- *Parallel Computing, 2020*
- *ACM Transactions on Architecture and Code Optimization*
- *ACM Transactions on Code Optimizations*
- *ACM Transactions on Cloud Computing*
- *Journal of Computational Science, 2020*
- *Journal of Parallel Computing, 2020*
- *Transactions on Re-configurable Technology and Systems, 2017*
- *IEEE Transactions on Magnetics*
- *IEEE Transactions on Parallel and Distributed Systems*
- *International Conference on Parallel Processing*
- *International Conference on Distributed Computing and Networking*
- *Journal of Signal Processing Systems, July 2016*
- *IEEE Conference in Electromagnetic Field Computation, June 2016*
- *Journal of Signal Processing Systems*
- *Workshop on Energy Aware Big Data Computing in Telecommunications*
- *Concurrency and Computation: Practice and Experience, 2019*
- *Transactions on Cloud Computing, 2019*

Industrial and Research Experience

18.A. Industrial Experience

Qualcomm Inc., Canada — Jul. 2012 to Feb. 2013

Senior R&D Engineer

- Supervisor: Alwyn Dos Remedios
- Built CVCL to automatically generate parallel code for computer vision problems.
- Optimized OpenCV and multimedia applications using OpenCL and CUDA.
- Developed an autotuner for CVCL.

18.B. Research Experience

Massachusetts Institute of Technology, USA — Feb. 2013 to Apr. 2015

Postdoctoral researcher

- Developed an autotuner for divide-and-conquer stencil computations.
- Designed and implemented domain-specific compilers for stencil computations.
- Reformulated and re-engineered the finite-element method for better scalability.
- Accelerated machine learning algorithms on heterogeneous hardware platforms.
- Designed a heuristic autotuner to tune the switching point in hybrid breadth-first search algorithms.

University of California Berkeley, USA — Oct. 2011 to Apr. 2012

Visiting student researcher

- Accelerated communication-avoiding (CA) Krylov solvers on GPUs.
- Designed and implemented preconditioning techniques for CA Krylov methods.

University of California Irvine, USA — Jan. 2011 to Apr. 2011

Visiting student researcher

- Accelerated sparse approximate inverse preconditioners on manycore architectures.

McGill University, Canada — 2008 to 2012

Research assistant

- Designed and implemented communication-reducing sparse data structures for sparse matrix computations.
- Developed algorithms for accelerating Krylov solvers on GPUs.
- Designed and implemented a runtime scheduler to improve the performance of electromagnetic simulations on clustered architectures.
- Developed single-element solutions to the finite-element method for better scalability.
- Accelerated preconditioned conjugate gradient methods on manycore hardware.

University of Calgary, Canada — 2005 to 2007

Research assistant

- Designed and implemented a Clustered Simultaneous Multithreaded simulator (CSMT-SIM) to simulate clustering on simultaneous multithreaded processors.