Maryam Mehri Dehnavi

CONTACT INFORMATION

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

Mobile: +1-416-948-2008
mmehride@cs.toronto.edu
https://www.cs.toronto.edu/ mmehride

RESEARCH INTERESTS AND VISION

High-Performance Computing, Parallel Algorithms, Compilers, Systems for Machine Learning, Cloud Computing, Computer Architecture, Numerical Analysis and Optimization, Graph Theory

ACADEMIC APPOINTMENTS

University of Toronto, Department of Computer Science — July 2018 to Date
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

EDUCATION

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

Isfahan University of Technology, Isfahan, Iran — 2001 to 2005
B.Sc., Electrical Engineering

HONORS, AWARDS, AND GRANTS

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

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

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

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

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

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

Adobe Research Project Funding: Sympiler: Transforming Sparse Matrix Codes.

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

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

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

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

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

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).

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

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

NSERC Graduate Scholarship-CGSD, $105,000, 2009-2012.

NSERC Michael Smith Foreign Study Scholarship, $6,000, U.C. Berkeley, 2012.

NSERC Industrial Research Fellowship, $60,000, 2012.

Visiting Fellowship in Canadian Government Labs, 2012.

ICPP 2014 and McGill Visiting Researcher (non-conference) 2011 Travel Grants.

FQRNT International Internship Scholarship, $14,470, U.C. Irvine, 2011.

Best paper finalist, *CEFC (IEEE Conference on computational electromagnetics)*, 2009.

ACCEPTED TO-APPEAR AND UNDER-REVIEW PAPERS

A-2 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. Under review.

A-1 S. Soori, B. Ucar, M. Gurbuzbalaban, and M. Mehri Dehnavi. ASYNC: Asynchrous Machine Learning on Distributed Systems. Under review at *The International Conference for High Performance Computing, Networking, Storage and Analysis, SC'19*, 12 pages. Submitted.

PEER REVIEWED CONFERENCE PUBLICATIONS

C-22 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 (PLDI)*, pp. 594-609, 2019.

C-21 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, SC*, pp. 779-793, 2018. Acceptance rate: 19%.

C-20 Z. Blanco, B. Liu, and M. Mehri Dehnavi. CSTF: Large-Scale Sparse Tensor Factorizations on Distributed Platforms. *The 47th International Conference on Parallel Processing, ICPP*, 2018. Acceptance rate: 22%.

C-19 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, ICPP*, 2018. Acceptance rate: 22%.

C-18 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, SC*, 2017: Acceptance rate: 18% (61 out of 327); ***Grand Final Winner of ACM SRC 2017***: Acceptance rate: 0.3% (1 out of 330 from all submissions from 27 ACM conferences)

C-17 B. Liu, C. Wen, A. Sarwate, and M. Mehri Dehnavi. A Unified Optimization Approach for Sparse Tensor Operations on GPUs. *IEEE CLUSTER17*, pp. 47-57, 2017. Acceptance rate: 21% (41 out of 216) .

C-16 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-15 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-14 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)*, pp. 283-288, 2017.

C-13 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-12 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-11 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-10 M. Mehri Dehnavi, J. Demmel, and D. Fernández. Communication-avoiding sparse approximate inverse preconditioners. *SIAM Conference on Linear Algebra (SIAM-LA)*, 2015.

C-9 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 (ICPP)*, 70–79, 2014. Acceptance rate: 21%.

C-8 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. Acceptance rate: 21% (114 out of 541).

C-7 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***. Oral presentation acceptance rate: <10% (24 out of 622).

C-6 M. Mehri Dehnavi, J. Demmel, and D. Giannacopoulos. Communication-avoiding algorithms on GPUs. *IEEE Conference on Electromagnetic Field Computation (CEFC)*, 2012.

C-5 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-4 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-3 D. Fernández, M. Mehri Dehnavi, and D. Giannacopoulos. Alternate approach to FEM for parallel processing. *IEEE Conference on Computational Electromagnetics (COMPUMAG)*, 2011. Oral presentation acceptance rate: <10% (16 out of 756).

C-2 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.Oral presentation acceptance rate: <10% (48 out of 649).

C-1 M. Mehri Dehnavi and D. Giannacopoulos. Enhancing the performance of electromagnetic applications on clustered architectures. *IEEE Conference on Electromagnetic Field Computation (CEFC)*, 2008.

REFEREED JOURNAL PUBLICATIONS

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.

Other Conference publications

P-8 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-7 S. Soori, R. Shah, M. Mehri Dehnavi, Avoiding Communication in Proximal Methods for Convex Optimization, *SIAM Parallel Processing, 2018*.

P-6 B. Liu, M. Mehri Dehnavi, Sparse Tensor Computations on Graphic Processing Units, *SIAM Parallel Processing, 2018*.

P-5 M. Mehri Dehnavi and D. Giannacopoulos. Fast preconditioning on GPUs. *High Performance Computing Symposium in medical sciences*, 2011.

P-4 M. Mehri Dehnavi and D. Giannacopoulos. Accelerating finite element sparse matrix vector multiplication on GPUs. *Centre de Recherche en Electronique Radiofrequence (CREER)*, 2010.

P-3 M. Mehri Dehnavi and D. Giannacopoulos. Enhancing the performance of clustered architectures. *6th Interdisciplinary Graduate Student Research Symposium (IGTRS)*, 2010.

P-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.

P-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.

Thesis and technical reports

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

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

T-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

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

T-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

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.

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.

Teaching Experience

University of Toronto, USA

Instructor

Applications of Parallel and Distributed Computing (CSC2222) — Fall 2019

Parallel Computing (CSC367): 4.6/5.0 — 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

- "The course made me realize which field I would like to pursue my career exactly in. "
- "It covers a lot of important issues for parallel computing, which is really helpful for students to enter this field."

Adv. Topics in Computer Engineering: Cloud Computing 4.6/5.00 — Fall 2016

- "Maryam is very knowledgeable & has a strong grasp on the subject. She teaches really well & is very approachable forming a great learning environment."
- "It was a comprehensive, in-depth, and beneficial coverage of so many important, *cutting edge* topics in cloud computing. I got so much insight out of it. This class was keeping up with new innovations currently coming out in present day, rather than like most classes, where the curriculum is covered in dust because it hasn't been changed since the early 1980's. "

- "It has been extremely insightful, and has provided something that I have seen in almost in no other part of the curriculum so strongly: a direct link to real-world applications, concepts, and implementations. Essential, and there should be more of this in other classes. "
- "I'm glad I took this course as an undergrad. It allowed me to see how graduate students prepare and think. By working with PhD students for the project, it showed me how to write papers and how to analyze papers on a higher level. It made me a stronger undergraduate student by just being around the upperclassmen and learning about how they work."

Adv. Topics in Computer Engineering: Cloud Computing: 4.84/5.0 Fall 2015

- "The instructor has motivated me to keep up with the new technologies and identify them as per need when required, and also to think out of the box and not only focus towards what you are trying to achieve."
- "Best course taken so far at Rutgers, would highly recommend to others."

Massachusetts Institute of Technology, USA

Recitation instructor and teaching assistant Fall 2013

6.172: Performance Engineering of Software Systems

- Designed projects, assignments, and recitation material.
- Held recitation sessions, office hours, and graded exams.

Student mentoring Winter 2013 to present

- Mentored a Ph.D. student at MIT in developing autotuners for stencil code.
- Mentored an M.Sc. student from Tsinghua university in accelerating machine learning and breadth-first search algorithms on heterogeneous architectures.

Padova University, Italy

Invited lecturer Summer 2011

Lecture: Algorithms and Architectures for CSE

- Gave a lecture series (3 days) on computational science and engineering.

McGill University, Canada

Lab. instructor Fall 2009 and Fall 2010

ECSE 291: Electrical measurements lab.

- Supervised lab sessions and graded lab reports.

Recitation instructor and teaching assistant Winter 2008 and Winter 2009

ECSE 425: Computer Organization and Architecture

- Held tutorial sessions and office hours.
- Designed and graded assignments and exams.

Student mentoring 2010 to 2011

- Mentored three undergraduate students on developing a fast MRI imaging software on GPUs.

University of Calgary, Canada

Recitation instructor and teaching assistant Fall 2006 and Winter 2007

ENCM 501: Principles of Computer Architecture

- Held tutorial sessions and office hours.
- Designed and graded assignments and exams.

Lab. instructor Fall 2005 and Winter 2006

ENEL 399: Programming Fundamentals

ENEL 409: Principles of Software Development

- Supervised lab sessions and graded lab reports.

STUDENT SUPERVISION

University of Toronto, Canada

PhD students

- Zachary Centinic
- Amir Masud Zare Bidaki
- YongZheng Huang (USRA)
- Jianda Chen
- Kazem Cheshmi
- Bangtian Liu
- Saeed Soori
- Laura Walsh

Rutgers University, USA

PhD students

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

Alumni MSc students

- Zachary Blanco (now at the MIT Lincoln Lab)
- 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

- Visiting Students: George Sakkas, Nikos Skilikas
- Alex Chan

Selected Invited Lectures and Presentations

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

Decoupling Structure in Hierarchical Matrix Approximations. ***Invited*** SIAM PP 2020.

A Compiler for Tensor Algebra. ***Invited*** Workshop on Compiler Techniques for Sparse Tensor Algebra at MIT, Boston, 2019.

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

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

Transforming Computation and Communication Patterns for High-Performance, Computer Science Department, UNC Charlotte, USA, 2018.

Transforming Computation and Communication Patterns for High-Performance, Computer Engineering Department, George Mason University, USA, 2018

Transforming Computation and Communication Patterns for High-Performance, Computer Engineering Department, UMass Amherst, USA, 2018.

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

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

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.

Sympiler: Transforming Sparse Matrix codes. Temple University, USA, 2017.

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

SIAM Linear Algebra Symposium, USA, 2015.

Rutgers University, Electrical and Engineering Department, USA, 2015.

Arizona State University, Computer Science Department, USA, 2015.

University of Victoria, Computer Science Department, Canada, 2015.

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

Center for Exascale Simulation of Plasma-Coupled Combustion. UIUC, USA, 2014.

IBM Research Yorktown heights. USA, 2014.

Northeastern University. Computer Science Department, USA, 2014.

Qualcomm Canada Inc. Canada, 2012.

Samsung Research America. USA, 2011.

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

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

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

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

6th Interdisciplinary Graduate Student Research Symposium. Canada, 2010.

High Performance Computing Symposium in Medical Sciences. Canada, 2011.

PROFESSIONAL SERVICE AND LEADERSHIP

Program Committee Leadership

- *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)*
- *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*

Program Committee and Other Service

- *Program Committee Member, ICPP 2020*
- *Program Committee Member, GPCE 2019*
- *Program Committee Member, CHIUW 2019*
- *Primary Program Committee Member, IPDPS 19*
- *NSERC CFI Reviewer, 2018, 2019*
- *NSERC CRD Reviewer, 2018, 2019*
- *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*
- *Transactions on Reconfigurable 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*

Other services

- *The New Jersey Junior Science & Humanities Symposium at Rutgers University–Reviewer and Mentor*
- *McGill Undergraduate Poster Tutorial–Mentor*
- *MIT Undergraduate Women's Mentoring Program–Organizer*