

Mark Christopher Jeffrey (he/him)

CONTACT INFORMATION	The Edward S. Rogers Department of Electrical and Computer Engineering 10 King's College Road University of Toronto Toronto, ON, M5S 3G4, Canada	mcj@ece.utoronto.ca markcjeffrey.com
RESEARCH INTERESTS	Computer architecture, computer systems, parallel computing, parallel programming models, data prefetching, LLM inference, compilers, irregular algorithms, reconfigurable hardware	
EDUCATION	Massachusetts Institute of Technology <i>Doctor of Philosophy, Electrical Engineering and Computer Science</i> Thesis: <i>A hardware and software architecture for pervasive parallelism</i> Advisor: Professor Daniel Sanchez	2019
	University of Toronto <i>Master of Applied Science, Computer Engineering</i> Thesis: <i>Understanding and improving Bloom filter configuration for lazy address-set disambiguation</i> Advisor: Professor J. Gregory Steffan	2011
	<i>Bachelor of Applied Science in Engineering Science with Honours</i>	2009
APPOINTMENT	University of Toronto , Toronto, Canada <i>Assistant Professor, Electrical and Computer Engineering</i> <i>Assistant Professor, Computer Science</i>	August 2020 – present July 2022 – present
INDUSTRY EXPERIENCE	Meta , Cambridge, Massachusetts <i>Research Scientist, Facebook Artificial Intelligence Research</i>	October 2019 – July 2020
	Google , Mountain View, California <i>Software Engineering Intern, Platforms Performance</i>	June 2015 – August 2015
	AeroFS , Palo Alto, California <i>Software Engineer</i>	September 2011 – May 2013
	EPSON , Toronto, Canada <i>Software Development Intern</i>	May 2007 – August 2008
	Neufeld Learning Systems , London, Canada <i>Software Development Intern</i>	Summer 2005, Summer 2006
HONOURS AND AWARDS	2nd Place , 4th Data Prefetching Championship @ HPCA	2026
	Best Paper Nominee, IEEE International conference on Field Programmable Technology	2024
	Connaught Fund New Researcher Award , University of Toronto (\$25,000)	2024
	Best Paper , ACM SIGMICRO International Workshop on Network on Chip Architectures	2023
	2nd Place , MIT EECS George M. Sprowls PhD Thesis Award in Computer Science	2021

HONOURS AND
AWARDS
(CONTINUED)

Best Graduate Poster, Industry-Academia Partnership MIT Cloud Workshop	2018
Facebook PhD Fellowship (\$181,000)	2017
Honourable mention in IEEE Micro “Top Picks from the Computer Architecture Conferences”	2017
Paper selected for IEEE Micro “Top Picks from the Computer Architecture Conferences”	2016
NSERC (NSF-equivalent) Post-Graduate Scholarship (PGS-D3 \$63,000)	2013
MIT Irwin Mark Jacobs and Joan Klein Jacobs Presidential Fellowship (\$69,166)	2013
Best Presentation, Connections Graduate Symposium, University of Toronto	2011
NSERC Alexander Graham Bell Canada Graduate Scholarship (CGS-M \$17,500)	2010
Best Paper, International Symposium on Applied Reconfigurable Computing	2010
L.E. Jones Award of Distinction, Engineering Alumni Association, University of Toronto	2009
Canada Millennium Scholarship Excellence Award (\$4,000)	2006
University of Toronto #2 Canadian Army University Course Award (\$1,300)	2006
University of Toronto Scholar (\$3,000)	2004

PEER-REVIEWED
CONFERENCE
PUBLICATIONS¹

- [C.16] G. Posluns and **M. C. Jeffrey**, “Symbiotic task scheduling and data prefetching,” in *Proc. of the 58th IEEE/ACM International Symposium on Microarchitecture (MICRO-58)*, Oct. 2025, pp. 140–155 (acceptance rate: 21%).
- [C.15] J. Cai, A. Vora, R. Zhang, M. O’Connor, and **M. C. Jeffrey**, “Attention-level speculation,” in *Proc. of the 42nd International Conference on Machine Learning (ICML)*, Jul. 2025 (acceptance rate: 27%).
- [C.14] A. Singer, H. Yan, G. Zhang, **M. C. Jeffrey**, M. Stojilović, and V. Betz, “MultiQueue-based FPGA routing: Relaxed A* priority ordering for improved parallelism,” in *Proc. of the IEEE International Conference on Field-Programmable Technology (FPT)*, Dec. 2024 (acceptance rate: 28%).
(Nominated for Best Paper Award)
- [C.13] G. Zhang, G. Posluns, and **M. C. Jeffrey**, “Multi bucket queues: Efficient concurrent priority scheduling,” in *Proc. of the 36th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, Jun. 2024, pp. 113–124 (acceptance rate: 29%).
- [C.12] J. Abdi, G. Posluns, G. Zhang, B. Wang, and **M. C. Jeffrey**, “When is parallelism fearless and zero-cost with Rust?” In *Proc. of the 36th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, Jun. 2024, pp. 27–40 (acceptance rate: 29%).
- [C.11] J. Zhao, I. Uwizyimana, K. Ganesan, **M. C. Jeffrey**, and N. Enright Jerger, “Altocumulus: Scalable scheduling for nanosecond-scale remote procedure calls,” in *Proc. of the 55th IEEE/ACM International Symposium on Microarchitecture (MICRO-55)*, Oct. 2022, pp. 423–440 (acceptance rate: 24%).
- [C.10] G. Posluns, Y. Zhu, G. Zhang, and **M. C. Jeffrey**, “A scalable architecture for reprioritizing ordered parallelism,” in *Proc. of the 49th ACM/IEEE International Symposium on Computer Architecture (ISCA-49)*, Jun. 2022, pp. 437–453 (acceptance rate: 17%).
- [C.9] K. Maeng, S. Bharuka, I. Gao, **M. C. Jeffrey**, V. Saraph, B.-Y. Su, C. Trippel, J. Yang, M. Rabbat, B. Lucia, and C.-J. Wu, “CPR: Understanding and improving failure tolerant training for deep learning recommendation with partial recovery,” in *Proc. of the 4th Conference on Machine Learning and Systems (MLSys)*, Apr. 2021 (acceptance rate: 24%).
- [C.8] V. A. Ying, **M. C. Jeffrey**, and D. Sanchez, “T4: Compiling sequential code for effective speculative parallelization in hardware,” in *Proc. of the 47th ACM/IEEE International Symposium on Computer Architecture (ISCA-47)*, Jun. 2020, pp. 159–172 (acceptance rate: 18%).

¹Supervised student authors identified with underline.

PEER-REVIEWED
CONFERENCE
PUBLICATIONS
(CONTINUED)

- [C.7] **M. C. Jeffrey**, V. A. Ying, S. Subramanian, H. R. Lee, J. Emer, and D. Sanchez, “[Harmonizing speculative and non-speculative execution in architectures for ordered parallelism](#),” in *Proc. of the 51st IEEE/ACM International Symposium on Microarchitecture (MICRO-51)*, Oct. 2018, pp. 217–230 (acceptance rate: 21%).
- [C.6] M. Abeydeera, S. Subramanian, **M. C. Jeffrey**, J. Emer, and D. Sanchez, “[SAM: Optimizing multi-threaded cores for speculative parallelism](#),” in *Proc. of the 26th International Conference on Parallel Architectures and Compilation Techniques (PACT-26)*, Sep. 2017, pp. 64–78 (acceptance rate: 23%).
- [C.5] S. Subramanian, **M. C. Jeffrey**, M. Abeydeera, H. R. Lee, V. A. Ying, J. Emer, and D. Sanchez, “[Fractal: An execution model for fine-grain nested speculative parallelism](#),” in *Proc. of the 44th ACM/IEEE International Symposium on Computer Architecture (ISCA-44)*, Jun. 2017, pp. 587–599 (acceptance rate: 17%).
- [C.4] **M. C. Jeffrey**, S. Subramanian, M. Abeydeera, J. Emer, and D. Sanchez, “[Data-centric execution of speculative parallel programs](#),” in *Proc. of the 49th IEEE/ACM International Symposium on Microarchitecture (MICRO-49)*, Oct. 2016, 5:1–5:13 (acceptance rate: 21%).
(Honourable mention for IEEE Micro’s Top Picks)
- [C.3] **M. C. Jeffrey**, S. Subramanian, C. Yan, J. Emer, and D. Sanchez, “[A scalable architecture for ordered parallelism](#),” in *Proc. of the 48th IEEE/ACM International Symposium on Microarchitecture (MICRO-48)*, Dec. 2015, pp. 228–241 (acceptance rate: 22%).
(Selected for IEEE Micro’s Top Picks issue of “most significant papers in computer architecture based on novelty and long-term impact”)
- [C.2] **M. C. Jeffrey** and J. G. Steffan, “[Understanding Bloom filter intersection for lazy address-set disambiguation](#),” in *Proc. of the 23rd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, Jun. 2011, pp. 345–354 (acceptance rate: 30%).
- [C.1] M. Labrecque, **M. C. Jeffrey**, and J. G. Steffan, “[Application-specific signatures for transactional memory in soft processors](#),” in *Proc. of the 6th International Symposium on Applied Reconfigurable Computing (ARC)*, Mar. 2010, pp. 42–54 (acceptance rate: 37%).

PEER-REVIEWED
JOURNAL
PUBLICATIONS

- [J.4] P. Golikov, K. Ganesan, G. Pekhimenko, and **M. C. Jeffrey**, “[Fusing adds and shifts for efficient dot products](#),” *IEEE Computer Architecture Letters*, vol. 25, no. 1, pp. 33–36, 2026
- [J.3] J. Li, **M. C. Jeffrey**, and N. Enright Jerger, “[A performance model for disintegrated manycores](#),” *IEEE Computer Architecture Letters*, vol. 24, no. 2, pp. 389–392, 2025
- [J.2] **M. C. Jeffrey**, S. Subramanian, C. Yan, J. Emer, and D. Sanchez, “[Unlocking ordered parallelism with the Swarm architecture](#),” *IEEE Micro’s Top Picks*, vol. 36, no. 3, pp. 105–117, 2016
- [J.1] M. Labrecque, **M. C. Jeffrey**, and J. G. Steffan, “[Application-specific signatures for transactional memory in soft processors](#),” *ACM Transactions on Reconfigurable Technology and Systems (TRETTS)*, vol. 4, no. 3, 21:1–21:14, 2011

PEER-REVIEWED
SHORT
PUBLICATIONS

- [S.4] G. Posluns and **M. C. Jeffrey**, “[Global berti: Simultaneous streaming and spatial prefetching](#),” in *Proc. of the 4th Data Prefetching Championship (DPC)*, Feb. 2026
(2nd place)
- [S.3] A. Plotnik, K. Ganesan, N. Enright Jerger, and **M. C. Jeffrey**, “[Intergenerational embodied carbon](#),” in *Proc. of the 1st Workshop on Hot Topics in Ethical Computer Systems (HotEthics)*, Apr. 2024
- [S.2] I. R. Brkić and **M. C. Jeffrey**, “[Disintegrating manycores: Which applications lose and why?](#)” In *Proc. of the 16th ACM SIGMICRO International Workshop on Network on Chip Architectures (NoCArc)*, Oct. 2023, pp. 3–8 (acceptance rate: 36%).
(Best Paper Award)
- [S.1] **J. Abdi**, G. Zhang, and **M. C. Jeffrey**, “[Brief announcement: Is the problem-based benchmark suite fearless with Rust?](#)” In *Proc. of the 35th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, Jun. 2023, pp. 303–305 (37 regular papers and 10 brief announcements accepted of 104 submissions).

SOFTWARE RELEASES	Attention-level speculation Tenstorrent-based implementation. Accompanies [C.15]	2025
	MultiQueue-based parallel FPGA routing implementation. Accompanies [C.14]	2024
	Multi Bucket Queue implementation and benchmark suite. Accompanies [C.13].	2024
	Rust Parallel Benchmarks suite. Accompanies [C.12][S.1].	2024
	The T4 auto-parallelizing compiler. Accompanies [C.8].	2020
FUNDING	Natural Sciences and Engineering Research Council, USRA (sole PI)	2026
	Total amount: \$6,000	
	Engineering Science Research Opportunities Program (sole PI)	2026
	Total amount: \$3,000	
	University of Toronto Joint EMHSeed and XSeed program	2025-2027
	Total amount: \$120,000	
	Annual amount: \$60,000	
	Annual amount/PI: \$30,000	
	Fujitsu Co-Creation Research Laboratory (sole PI)	2025-2026
	Total amount: \$100,000	
	Annual amount/PI: \$100,000	
	Fujitsu Co-Creation Research Laboratory (sole PI)	2024-2025
	Total amount: \$50,000	
	Annual amount/PI: \$50,000	
	University of Toronto Connaught New Researcher Award (sole PI)	2024-2026
Total amount: \$25,000		
Natural Sciences and Engineering Research Council, USRA (sole PI)	2024	
Total amount: \$6,000		
Natural Sciences and Engineering Research Council, USRA (sole PI)	2023	
Total amount: \$6,000		
Natural Sciences and Engineering Research Council, USRA (sole PI)	2023	
Total amount: \$6,000		
Natural Sciences and Engineering Research Council, USRA (sole PI)	2022	
Total amount: \$6,000		
NSERC Discovery Launch Supplement DGEGR-2022-00117 (sole PI)	2022	
Total amount: \$12,500		
Total amount/PI: \$12,500		
NSERC Discovery Grant RGPIN-2022-05330 (sole PI)	2022-2027	
Total amount: \$145,000		
Annual amount: \$29,000		
Annual amount/PI: \$29,000		
Engineering Science Research Opportunities Program (sole PI)	2021	
Total amount: \$3,000		
TEACHING	University of Toronto , Toronto, Ontario	
	<i>Instructor</i> , ECE1755 Parallel Computer Architecture and Programming	Spring 2021–2025
	<i>Instructor</i> , ECE552 Computer Architecture	Fall 2020–2023, 2025
	<i>Instructor</i> , ECE253 Digital and Computer Systems	Fall 2022, 2023, 2025

TEACHING
(CONTINUED)**Massachusetts Institute of Technology**, Cambridge, Massachusetts*Guest Lecturer*, [6.823](#) Computer System Architecture**Spring 2019***Guest Lecturer*, [6.886](#) Graph Analytics**Spring 2018***Teaching Assistant*, [6.823](#) Computer System Architecture**Spring 2017**STUDENT
SUPERVISION*Current*

Gilead Posluns, Ph.D. student

Awarded a \$20,000 Ontario Bell Graduate Scholarship

Awarded a \$15,000 Ontario Graduate Scholarship

Awarded a \$15,000 Ontario Graduate Scholarship

Aster Plotnik, M.A.Sc. student (co-supervised with Natalie Enright Jerger)

Steven Hill, M.A.Sc. student (co-supervised with Natalie Enright Jerger)

Awarded a \$15,000 Ontario Graduate Scholarship

Parisa Betel Miri, M.A.Sc. student

Jimmy Su, M.A.Sc. student

Athena Cai, B.Sc. research intern and CSC392 project

Kevin Qu, B.A.Sc. thesis

*Alumni (M.A.Sc.)*Guozheng (Ray) Zhang, [M.A.Sc. thesis](#)**2024**

Awarded a \$15,000 Ontario Queen Elizabeth II Graduate Scholarship

First position: Compiler Engineer, Huawei

Mohammad Javad Abdi, [M.A.Sc. thesis](#)**2024**

First position: Member of Technical Staff, Cerebras

Isidor Brkić, [M.A.Sc. thesis](#)**2023**

Awarded a \$15,000 Ontario Queen Elizabeth II Graduate Scholarship

First position: Digital IC Design Engineer, StarIC

Gilead Posluns, [M.A.Sc. thesis](#)**2022**

Awarded a \$15,000 Ontario Queen Elizabeth II Graduate Scholarship

First position: Ph.D. student, University of Toronto

Alumni (M.Eng.)

Hanxiao Wei, M.Eng. research project

2024

Yue Fei, M.Eng. summer research

2024

First position: Design Verification Engineer, Qualcomm

Alumni (B.A.Sc.)

Angela Yu, B.A.Sc. thesis and summer research

2025

First position: M.A.Sc. student, University of Toronto

Rachel Chen, B.A.Sc. summer research

2025

Angus Wu, B.A.Sc. thesis and USRA

2025

First position: M.A.Sc. student, University of Toronto

Jack Cai, B.A.Sc. thesis

2024

First position: Member of Technical Staff, xAI

Abnash Bassi, B.A.Sc. summer research

2024

First position: Hardware System Engineer, Rivian

Edward Wu, B.A.Sc. summer research

2024

Stephen Yang, B.A.Sc. thesis

2024

First position: M.A.Sc. student, University of Toronto

Balaji Venkatesh, B.A.Sc. thesis

2024

First position: Software Developer, GridS2

Leo Han, B.A.Sc. thesis and USRA

2023

First position: Ph.D. student, Cornell Tech

STUDENT SUPERVISION (CONTINUED)	<i>Alumni (B.A.Sc.)</i>	
	Davendra Seunarine Maharaj, B.A.Sc. intern and thesis	2023
	First position: M.Sc. student, Georgia Tech	
	Eugene Lee, B.A.Sc. USRA	2023
	Jerry He, B.A.Sc. thesis	2022
	First position: Software Engineer, Microsoft	
	Larry Wu, B.A.Sc. USRA	2022
	First position: Software Engineer, Qualcomm	
	Yan Zhu, B.A.Sc. ESROP and intern	2021-2022
	First position: Ph.D. student, University of California Berkeley	
Billy Boyle, B.A.Sc. thesis	2021	
First position: Mixed Architecture Specialist, TC Helicon		
INVITED TALKS	Attention-level speculation	
	Fujitsu	May 2025
	Faster priority ordered irregular parallelism through hardware and software	
	Hong Kong University of Science and Technology	October 2024
	Cornell University	August 2024
	Performance for all: simplifying hard parallelism and specialization	
	Fujitsu	February 2024
	Stanford University	May 2023
	Making parallelism pervasive with the Swarm architecture	
	Facebook	September 2019
	Google	May 2019
	University of Pennsylvania	April 2019
	University of Toronto	March 2019
	University of Waterloo	March 2019
	University of Texas at Austin	March 2019
	Simon Fraser University	January 2019
	Facebook	September 2017
Center for Future Architectures Research e-Workshop	July 2017	
University of Toronto Computer Architecture Seminar	March 2017	
BARC: Boston Area Architecture Workshop	January 2017	
Harmonizing speculative and non-speculative execution in architectures for ordered parallelism		
IEEE/ACM International Symposium on Microarchitecture	October 2018	
Data-centric execution of speculative parallel programs		
IEEE/ACM International Symposium on Microarchitecture	October 2016	
A scalable architecture for ordered parallelism		
IEEE/ACM International Symposium on Microarchitecture	December 2015	
Improving Bloom filter configuration for lazy transactional memory		
CASCON, IBM Canada Software Laboratory	November 2011	
Understanding Bloom filter intersection for lazy address-set disambiguation		
ACM Symposium on Parallelism in Algorithms and Architectures	June 2011	
University of Toronto Connections ECE Graduate Symposium	May 2011	
GPU-accelerated software transactional memory		
University of Toronto Connections ECE Graduate Symposium	May 2010	

MENTORING AND OUTREACH	<i>Panelist</i>	
	“Building a Research Program”, University of Toronto Prospective Professors in Training	2024-2025
	“Becoming a Professor”, University of Toronto Division of Engineering Science	2024
	“Applying to Graduate School”, Undergrad Architecture Mentoring (uArch) Workshop @ ISCA	2022
	“Working in Academia”, University of Toronto Division of Engineering Science	2021,2022
	“Former Fellows Panel”, Facebook Fellowship Summit	2020
	<i>Mentor</i>	
	Undergrad Architecture Mentoring (uArch) Workshop	2021–2025
	Meet a Senior Architect Program, ISCA	2020,2021,2023
	Meet a Senior Architect Program, MICRO	2020,2021
	Meet a Senior Architect Program, ASPLOS	2021
PROFESSIONAL SERVICE	<i>Program Committee Area Chair (Conferences)</i>	
	Intl. Symposium on Microarchitecture (MICRO)	2026
	<i>Program Committee Member (Conferences)</i>	
	Intl. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)	2026
	AAAI Conference on Artificial Intelligence (AAAI)	2026
	Intl. Symposium on Computer Architecture (ISCA)	2023,2025
	Intl. Symposium on Microarchitecture (MICRO)	2020,2023
	Intl. Symposium on Workload Characterization (IISWC)	2022
	<i>External Review Committee Member (Conferences)</i>	
	Intl. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)	2021,2022,2023,2025
	Intl. Symposium on Microarchitecture (MICRO)	2021,2022,2024,2025
	Intl. Symposium on Computer Architecture (ISCA)	2022
	<i>Program Committee Member (Workshops/Competitions)</i>	
	Young Architect Workshop	2024
	Student Research Competition @	
	Intl. Conf. on Parallel Architectures and Compilation Techniques (PACT)	2022
	<i>Reviewer</i>	
	IEEE Transactions on Parallel and Distributed Systems (TPDS)	2022
	IEEE Computer Architecture Letters (CAL)	2020
	Symposium on Principles and Practice of Parallel Programming (PPoPP)	2016
<i>External Reviewer</i>		
NSERC Discovery Grants, Electrical and Computer Engineering Committee	2023	
<i>Student Research Competition Co-Chair</i>		
Intl. Symposium on Microarchitecture (MICRO)	2023	
<i>Finance Chair</i>		
Intl. Symposium on High-Performance Computer Architecture (HPCA)	2023	
<i>Web Chair</i>		
Intl. Symposium on Computer Architecture (ISCA)	2022	
<i>Submissions Co-Chair</i>		
Intl. Symposium on Microarchitecture (MICRO)	2017	

PROFESSIONAL
SERVICE
(CONTINUED)

Professional Memberships

Member of IEEE, IEEE Computer Society, Technical Community on Computer Architecture (TCCA),
Technical Community on Microprogramming and Microarchitecture (TCuArch)
Member of ACM, SIGARCH, SIGMICRO

UNIVERSITY
SERVICE

Member

ECE Graduate Awards and Appeals Committee	2025–2026
ECE Graduate Matters Committee	2022–2024
School of Graduate Studies NSERC CGS M Awards Committee	2023
School of Graduate Studies NSERC CGS/PGS D Awards Committee	2022