COMPUTE ACCELERATION SYMPOSIUM

Hardware Breakthroughs for the Future of Computing

May 17, 2023 8:50 a.m. - 6:30 p.m.
U of T, St. George Campus
Medical Sciences Building, Rm 3153

DEFY GRAVITY

The Edward S. Rogers Sr. Department of Electrical & Computer Engineering
UNIVERSITY OF TORONTO

8:50 - 9:00	**Welcome Remarks** Stark Draper, Vice-Dean – Research, Faculty of Applied Science and Engineering Deepa Kundur, Chair, Department of Electrical & Computer Engineering
Computer Architecture and Electronics	
9:00 - 9:30	**Boosting Computing Capabilities Beyond Semiconductor Technology Scaling** Andreas Moshovos, University of Toronto
9:30 - 10:00	**Large Language Model Training at Wafer-Scale** Nish Sinnadurai, Cerebras Systems
10:00 - 10:30	**Squeezing the Energy out of Computation** Martin Snelgrove, Untether AI
10:30 - 10:45	**Break**
Datacenters and Co-Packaged Optics	
10:45 - 11:15	**The Role of Co-Packaged Optics in Datacentres** Tony Chan Carusone, Alphawave
11:15 - 11:45	**BSoBOX: Towards Low-Cost Light Sources for Co-Packaged Optics Applications through Heterogeneous Integration** Jason Mak, SCINTIL Photonics
11:45 - 12:15	**Accelerating AI through Photonic Communication and Computing** Darius Bunandar, Lightmatter
12:15 - 13:30	**Lunch and Networking**

Program continued on the next page

COMPUTE ACCELERATION SYMPOSIUM

May 17, 2023 8:50 a.m. - 6:30 p.m.
U of T, St. George Campus, Medical Sciences Building, Rm 3153

Photonic Neural Networks & In-Physics Computing

13:30 - 14:00	**Wavelength Multiplexed Photonic Edge Computing in the Output Stationary Frame** Ryan Hamerly, NTT & MIT
14:00 - 14:30	**Neuromorphic Silicon Photonics and Applications from Classical to Quantum** Bhavin J. Shastri, Queen's University
14:30 - 15:00	**Deep Physical Neural Networks Trained with in-situ Backpropagation** Tatsuhiro Onodera, NTT & Cornell
15:00 - 15:15	**Break**

Quantum Computing

15:15 - 16:15	**Programmable Optics for Quantum Networks and Computing** Dirk Englund, MIT and QuEra
16:15 - 16:45	**Photonic Quantum Computing** Christian Weedbrook, Xanadu Quantum Technologies
16:45 - 17:15	**Break & Networking**

Economics and Societal Impact, Roundtable

17:15 - 17:20	**Remarks** Joyce Poon, University of Toronto & Max Planck Institute of Microstructure Physics
17:20 - 17:40	**The Business Case for Quantum Computing in the NISQ Era and Beyond** Francesco Bova, University of Toronto
17:40 - 18:15	**Roundtable** Moderator: Jackson Hamilton (Business Development Bank of Canada) See participants on p. 8
18:30 - 20:00	**Reception, Massey College (by invitation only)**

COMPUTE ACCELERATION SYMPOSIUM

May 17, 2023 8:50 a.m. - 6:30 p.m.
U of T, St. George Campus, Medical Sciences Building, Rm 3153

Computer Architecture and Electronics

Andreas Moshovos, University of Toronto

Boosting Computing Capabilities Beyond Semiconductor Technology Scaling
9:00 - 9:30

ABSTRACT Computing devices play a vital role in driving innovation across various fields such as science, medicine, and commerce, and have become indispensable in our everyday lives. The reason these devices are so useful lies in their increasing capabilities and the growing interconnectedness of our world. Among the key factors behind the continuous improvement of computing devices is the advancement of semiconductor technology. However, effectively utilizing these advancements is no simple task. It is crucial to harness the full potential of the underlying semiconductor technology in order to deliver more powerful and capable computing devices. This talk reviews the dynamic relationship between semiconductor technology, computing architecture, and applications, highlighting why progress in all is essential.

SPEAKER BIO Andreas Moshovos teaches how to design and optimize computing hardware engines at the University of Toronto where he has the privilege of collaborating with several talented students on techniques to improve execution time, energy efficiency and cost for computing hardware. He has also taught at Northwestern University, USA, the University of Athens, Greece, the Hellenic Open University, Greece, and as an invited professor at the École Polytechnique Fédérale de Lausanne, Switzerland. He has received the ACM SIGARCH Maurice Wilkes award in 2010, an NSF CAREER Award in 2000, two IBM Faculty awards, a Semiconductor Research Corporation Inventor recognition award, a MICRO Hall of Fame award, and the IEEE C.C. Gotlieb Computer Award. He is the Director of the NSERC COHESA Research Network, and a fellow of the ACM and IEEE.

Computer Architecture and Electronics

Nish Sinnadurai, Cerebras Systems

Large Language Model Training at Wafer-Scale
9:30 - 10:00

ABSTRACT Large Language Models (LLMs) have made significant progress over the past several years, unlocking many new capabilities and use cases for AI. However, they are extremely computationally intensive and challenging to train at scale. Recently, Cerebras published a family of open-source LLMs (Cerebras-GPT) trained on its novel Wafer-Scale Engine and the Andromeda supercomputer - the first such models trained on an AI hardware accelerator. This talk will provide an overview of various aspects of the solution that enabled this work: the Wafer-Scale Engine chip architecture, Weight Streaming execution model, and Andromeda Wafer-Scale Cluster.

SPEAKER BIO Nish Sinnadurai is a VP of Engineering and Canada GM at Cerebras Systems, where he leads teams developing graph compilers for ML acceleration. He led the development of the Weight Streaming system for training LLMs with billions-to-trillions of parameters. Nish was previously an Engineering Director at Intel, where he led the Compiler group that develops EDA algorithms for FPGAs.

COMPUTE ACCELERATION SYMPOSIUM

May 17, 2023 8:50 a.m. - 6:30 p.m.
U of T, St. George Campus, Medical Sciences Building, Rm 3153

Computer Architecture and Electronics

Martin Snelgrove, Untether AI

Squeezing the energy out of computation
10:00 - 10:30

ABSTRACT Computing, in particular AI, is on a rapid exponential to use all the power we can generate and all the energy we can store; the one piece of good news being that both the corporate bottom line and the environment want this fixed, fast. The bad news is that there's no "one weird trick" for efficient computing, because we're wasting it in a thousand ways. We'll teach the elements of "whack-a-mole" by introducing the first few moles.

SPEAKER BIO Martin taught electronics at U of T from '82 to '92, then had an NSERC research chair at Carleton until '97 until his entire lab spun out to form an RFIC chip company, Philsar. He's been an entrepreneur since then, and is currently CTO of Untether AI, founded on a side project he did during his PhD.

Datacenter and Co-packaged Optics

Tony Chan Carusone, Alphawave

The Role of Co-Packaged Optics in Datacentres
10:45 - 11:15

ABSTRACT Progress in computation and communication is increasingly bottlenecked by integrated circuit I/O. CMOS technology scaling has enabled complete DSP-based modems operating over 100Gbps. At the same time, 100Gbps optical links have become preferable to copper links in datacentres over any distance beyond a few metres. Co-packaged optics can bring fibre right to the perimeter of CMOS circuits, all but eliminating copper links entirely, but significant technical challenges remain. This talk will elucidate these opportunities and challenges, painting a realistic picture of CPO's role in datacentres.

SPEAKER BIO Dr. Tony Chan Carusone has taught and researched integrated circuits and systems in academia and industry for over 20 years. He has been a faculty member at the University of Toronto since completing his Ph.D. there in 2002. He and his graduate students have received eight best-paper awards at leading conferences for their work on chip-to-chip and optical communication circuits, analog-to-digital conversion, and precise clock generation. He has also been a consultant to industry since 1997, and in 2022 became the Chief Technology Officer of Alphawave Semi in Toronto, Canada.
Dr. Chan Carusone was a Distinguished Lecturer for the IEEE Solid-State Circuits Society 2015-2017 and served on the Technical Program Committee of the International Solid-State Circuits Conference from 2015-2021. He has co-authored the latest editions of the classic textbooks “Analog Integrated Circuit Design” along with D. Johns and K. Martin, and “Microelectronic Circuits” along with A. Sedra and K.C. Smith. He was Editor-in-Chief of the IEEE Transactions on Circuits and Systems II: Express Briefs in 2009, an Associate Editor for the IEEE Journal of Solid-State Circuits 2010-2017 and is now Editor-in-Chief of the IEEE Solid-State Circuits Letters. He is a Fellow of the IEEE.

Datacenter and Co-packaged Optics

Jason Mak, SCINTIL Photonics

BSoBOX: Towards Low-Cost Light Sources for Co-Packaged Optics Applications through Heterogeneous Integration
11:15 - 11:45

ABSTRACT Silicon photonics and adjacent platforms lack intrinsic means to generate light, but progress has been made towards commercializing heterogeneously integrated light sources on silicon platforms. SCINTIL Photonics is aiming to develop low cost light sources addressing co-packaged optics (CPO) applications through its back-side on buried oxide (BSoBOX) platform.

SPEAKER BIO Dr. Jason Mak graduated with a PhD in Electrical Engineering from the University of Toronto in 2018, where he designed passive and tunable optical components on multilayer silicon photonics platforms under the supervision of Professor Joyce Poon. Jason is a founding member of SCINTIL Photonics, which aims to bring the III-V materials necessary for laser integration to standard multilayer silicon photonics platforms. Today, Jason works as an optical engineer in SCINTIL's Toronto offices, where he is responsible for circuit design and component development.

Datacenter and Co-packaged Optics

Darius Bunandar, Lightmatter

Accelerating AI through Photonic Communication and Computing
11:45 - 12:15

ABSTRACT Lightmatter is combining electronics, photonics, and new algorithms to create a next-generation compute platform. At the core of Lightmatter's technology is the use of silicon photonic devices. In this talk, I will introduce two of Lightmatter's photonic computing technologies: Passage and Envise. Passage is a wafer-scale programmable photonic interconnect that uses silicon waveguides to enable heterogeneous chips to communicate with unprecedented bandwidth and energy efficiency. Envise is the world's first major photonic compute processor for accelerating machine learning and digital signal processing workloads.

SPEAKER BIO Darius Bunandar is a Co-Founder and Chief Scientist of Lightmatter. Using a combination of electronics, photonics, and new algorithms, Lightmatter built a next-generation computing platform purpose-built for artificial intelligence. As Chief Scientist, Darius coordinates the company's R&D in new architectures and new algorithms. Darius obtained his PhD in physics at MIT studying quantum communications using compact nanophotonic circuits.

Photonic Neural Networks & In-Physics Computing

Ryan Hamerly, NTT Research & MIT

Wavelength Multiplexed Photonic Edge Computing in the Output Stationary Frame
13:30 - 14:00

ABSTRACT We propose a photonic edge computing architecture based on WDM, broadband modulation, and output-stationary integration. Using this scheme, we demonstrate 98.8%-accurate DNN inference over an 86-km deployed fiber link with 3 THz optical bandwidth.

SPEAKER BIO BS 2010 Caltech, PhD 2016 Stanford, IC Postdoc 2017-18, currently at NTT Research and MIT working on quantum optics, integrated photonics, and machine learning.

Photonic Neural Networks & In-Physics Computing

Bhavin Shastri, Queen's University

Neuromorphic Silicon Photonics and Applications from Classical to Quantum
14:00 - 14:30

ABSTRACT Artificial intelligence (AI), powered by neural networks, has revolutionized various fields. However, its implementation on conventional computers is limited in terms of speed/energy efficiency. Neuromorphic photonics aims to build neural networks that use light and photonic device physics for distributed and parallel processing. This technology offers sub-nanosecond latencies, extending the domain of AI and neuromorphic computing applications. We will discuss photonic neural networks enabled by silicon photonics. We will highlight their applications in scientific computing, online learning, and signal processing (fibre and wireless communications). We will briefly introduce a quantum photonic neural network as near-deterministic quantum gates and operations.

SPEAKER BIO Prof. Bhavin J. Shastri is an Assistant Professor of Engineering Physics at Queen's University and a Faculty Affiliate at the Vector Institute. He was an Associate Research Scholar (2016-2018) and Banting and NSERC Postdoctoral Fellow (2012-2016) at Princeton University. He received a Ph.D. degree in electrical engineering (photonics) from McGill University in 2012.
He is a co-author of the book Neuromorphic Photonics, a term he helped coin. Dr. Shastri is the recipient of a 2022 iCANX Young Scientist Award, the 2022 SPIE Early Career Achievement Award, and the 2020 IUPAP Young Scientist Prize in Optics "for his pioneering contributions to neuromorphic photonics." He is a Senior Member of Optica and IEEE.

COMPUTE ACCELERATION SYMPOSIUM

May 17, 2023 8:50 a.m. - 6:30 p.m.
U of T, St. George Campus, Medical Sciences Building, Rm 3153

Photonic Neural Networks & In-Physics Computing

Tatsuhiro Onodera, NTT & Cornell

Deep Physical Neural Networks Trained with in-situ Backpropagation
14:30 - 15:00

ABSTRACT Deep neural networks are ubiquitous in today's data-driven world, but their increasing energy requirements hinder broader application. In this work, we propose the construction of deep physical neural networks, consisting of layers of controllable physical systems, that learns hierarchical representations of input data. To train these networks efficiently, we introduce physics-aware training, a hybrid in situ-in silico backpropagation-based algorithm. To demonstrate the universality of our approach, we train diverse physical neural networks based on optics, mechanics, and electronics for audio and image classification tasks. This work broadens the possibility of using novel physical systems for energy-efficient machine learning.

SPEAKER BIO Tatsuhiro is a research scientist with NTT Research's PHI Laboratory, and Cornell University, where he works in the laboratory of Peter McMahon. He received his BASc in Engineering Science from the University of Toronto in 2014, and his PhD in Applied Physics from Stanford University in 2019. His research interests include quantum optics, integrated photonics, and the physics of computation.

Quantum Computing

Dirk Englund, MIT and QuEra

Programmable Optics for Quantum Networks and Computing
15:15 - 16:15

ABSTRACT After several decades of intensive theoretical and experimental efforts, the field of quantum information processing is at a critical moment: special-purpose quantum information processors are at or past the "quantum complexity frontier" where classical computers can no longer predict their outputs: we can "program complexity", unable to predict the outcome. Meanwhile, new technologies to connect quantum processors by photons give rise to quantum networks with functions impossible on today's "classical-physics" internet.
But to harness the power of quantum complexity in "noisy intermediate-scale" quantum computers and networks, we need new methods to control and understand them -- and perhaps to manage noise sufficiently to reach fault tolerance. This talk discusses one approach: large-scale programmable photonic integrated circuits (PICs) designed to control photons and atomic or atom-like quantum memories.

SPEAKER BIO Dirk Englund received his BS in Physics from Caltech (2002), MS in Electrical Engineering at Stanford, and PhD in Applied Physics also at Stanford (2008). After a postdoctoral fellowship at Harvard University, he joined Columbia University as Assistant Professor of E.E. and of Applied Physics. He joined the MIT EECS faculty in 2013. Major recognitions include the Presidential Early Career Award in Science and Engineering, the Sloan Fellowship in Physics, the OSA's Adolph Lomb Medal, the Bose Research Fellowship, and the A.v. Humboldt Research Fellowship . He's a fellow of the Optica Society.

COMPUTE ACCELERATION SYMPOSIUM

May 17, 2023 8:50 a.m. - 6:30 p.m.
U of T, St. George Campus, Medical Sciences Building, Rm 3153

Quantum Computing

Christian Weedbrook, Xanadu Quantum Technologies

Photonic Quantum Computing
16:15 - 16:45

ABSTRACT Overview of what we are doing here at Xanadu!

SPEAKER BIO Christian Weedbrook is the Founder and CEO of Xanadu, a photonic quantum computing company based in Toronto. Christian has over a decade of industry, government and research experience in quantum computing. He holds a PhD in Physics from the University of Queensland and was a postdoc at MIT and the University of Toronto.

Economics and Societal Impact

Francesco Bova, Rotman School of Business

The Business Case for Quantum Computing in the NISQ Era and Beyond
17:25 - 17:45

ABSTRACT I will discuss how companies are creating value by working and experimenting with current quantum hardware. I will also discuss how quantum computers may not need to display a quantum advantage in order to be economically useful to a variety of market stakeholders.

SPEAKER BIO Francesco Bova is an Associate Professor at the University of Toronto's Rotman School of Management. Francesco is also the Academic Lead for Creative Destruction Lab's Toronto site, the lab economist for its quantum stream, and a moderator for numerous streams. Francesco has diverse research interests and has published his work in journals from a variety of management disciplines. Most recently, his research has focused on the economics of quantum technologies. This research includes his co-authored work on quantum economic advantage and the commercial applications of quantum computing. Francesco is also an award-winning educator and leads the CDL Partners module on The Simple Economics of Quantum Computing.

Roundtable Discussion

17:45 - 18:15
MODERATOR:
Jackson Hamilton (Business Development Bank of Canada)
ROUNDTABLE PARTICIPANTS:
Chris Ouslis, National Research Council IRAP
Paul Slaby, Canada Semiconductor Council
Elissa Strome, CIFAR / Pan-Canada AI strategy
Rajat Ghosh, Innovation, Science and Economic Development Canada (ISED), Federal S&T
Kofi Kobia, ISED - A/Director, Automotive, Transport & Industry
Andrew Guy, Ontario Ministry of Economic Development, Job Creation, and Trade

COMPUTE ACCELERATION SYMPOSIUM & ACTIVE SILICON PHOTONICS FABRICATION WORKSHOP

Sponsored by:

The Edward S. Rogers Sr. Department of Electrical & Computer Engineering
UNIVERSITY OF TORONTO

PIC
Photonics Innovation Centre
University of Toronto

CMC MICROSYSTEMS

NTT Research

SCINTIL PHOTONICS

OZ Optics LIMITED
shop.ozoptics.com
www.ozoptics.com

LIGHTMATTER

XANADU

bdc

IEEE Advancing Technology for Humanity
TORONTO Section est. 1903

CQC TORONTO

SPIE. STUDENT CHAPTER UNIVERSITY OF TORONTO

OPTICA Advancing Optics and Photonics Worldwide
University of Toronto Student Branch

UNIVERSITY OF TORONTO
FACULTY OF APPLIED SCIENCE & ENGINEERING