

THE UNIVERSITY OF BRITISH COLUMBIA
Curriculum Vitae for Faculty Members

Date: 30 / Jan / 2019

Initials: _____

1. **SURNAME:** Aamodt **FIRST NAME:** Tor
MIDDLE NAME (S): Michael
2. **DEPARTMENT/SCHOOL:** Electrical and Computer Engineering
3. **FACULTY:** Applied Science
4. **PRESENT RANK:** Professor **SINCE:** 1 July 2016
5. **POST-SECONDARY EDUCATION**

University or Institution	Degree	Subject Area	Date of Completion
University of Toronto	Ph.D. ¹	Electrical and Computer Engineering	Jun / 2006
University of Toronto	M.A.Sc.	Electrical and Computer Engineering	Jun / 2001
University of Toronto	B.A.Sc.	Engineering Science	Jun / 1997

- (a) **Special Professional Qualifications**
Association of Professional Engineers and Geoscientists of British Columbia P.Eng., Jun / 2007 - Present
- (b) **Continuing Education / Training (attended)**
TAG Instructional Skills Workshop (Aug / 2006)
Faculty Certificate Program on Teaching and Learning in Higher Education (Sep / 2006 - Apr / 2007)

6. **EMPLOYMENT RECORD**

- (a) *Prior to coming to UBC*

University, Company or Organization	Rank or Title	Dates
NVIDIA Corporation	Sr. Architect	Sep / 2004 – Jan / 2006
Intel Corporation	Consultant	Jun / 2003 – Aug / 2003
Intel Corporation	Intern	Mar / 2002 – May / 2003
AlliedSignal Aerospace Canada	Intern	May / 1995 – Aug / 1996

- (b) *At UBC*

Rank or Title	Dates
Professor	Jul / 2016 – Present
Associate Professor	Jul / 2011 – Jun / 2016
Assistant Professor	Jan / 2006 – Jun / 2011

- (c) *Date of granting of tenure at UBC: July 2011*

7. **LEAVES OF ABSENCE**

University, Company or Organization at which Leave was taken	Type of Leave	Dates
Stanford University (Visiting Associate Professor)	Sabbatical Study Leave	Sep / 2012 – Jun / 2013

¹ Title: Modeling and Optimization of Speculative Threads, Supervisor: Professor Paul Chow

8. **TEACHING**(a) *Briefly describe areas of special interest and accomplishments*

Areas of Interest: Computer Architecture, Digital Design

Accomplishments:

- CPEN 211 (formerly EECE 259) Introduction to Microcomputers: Course redesign for 2014 and 2015 terms. Fall 2014: Introduced VHDL-2008 and ARM instruction set architecture. Basics of pipelining, caches and virtual memory. In addition to developing new lecture material also developed roughly 200 in-class “clicker” questions. Developed 11 new labs introducing students to digital logic through breadboard design (2 labs); VHDL simulation, testing and synthesis including the development of a simple microprocessor (5 labs); ARM assembly programming using a synthesized open source ARM core on the DE2 boards (4 labs). Initiated development of a custom interactive debugging environment suitable for teaching assembly level ARM programming.
Fall 2015: Revised course to use Verilog instead of VHDL for all components; revised all labs due to introduction of more modern hardware platform: DE1-SoC. Created five new online lectures (~1 hr video + 5 to 30 questions) and introduced use of seven flipped lectures into CPEN 211 (two lectures reused videos created for EECE 476). Online lectures and questions uploaded to edge.edx.org. Developed and delivered corresponding class components.
Fall 2016: Minor revisions to labs to improve clarity; Wrote software for and administered a set of Verilog proficiency tests to ensure students who pass course have attain a minimum level of practical competence in digital design.
Fall 2017: Introduced partial auto-grading of lab assignments to reduce pressure on TAs and conflicts over marking. Extended proficiency tests to cover additional labs include ARM assembly programming.
- EECE 476 (Computer Architecture): Fall 2006: Redesigned course to emphasize modern architectures. Incorporated “clickers” into teaching material. Fall 2013: Created 5 flipped lectures.
- EECE 353 (Digital Systems Design): Introduced use of “clickers” questions.
- EECE 527 (Advanced Computer Architecture): Proposed and created graduate course providing students with in depth understanding of microprocessor microarchitecture.

(b) *Courses Taught at UBC*

Year/ Term	Course Number	Scheduled Hours	Class Size	Total Hours Taught			
				Lectures	Labs	Tutorials	Other
2006S	EECE 353	2-3*-1	67	36	0	12	0
2006W T1	EECE 476	3-0-1	72	39	0	0	0
2006W T2	EECE 571B	3-0-0	8	26	0	0	16
2006W T2	EECE 571T	3-0-0	4	24	0	0	0
2007W T1	EECE 476	3-0-1	89	39	0	0	13
2007W T2	EECE 353	2-3*-1	92	26	13	10	13
2007W T2	EECE 527	3-0-0	9	39	0	0	0
2008W T1	EECE 476	3-0-1	69	32	13	0	0
2008W T1	EECE 527	3-0-0	19	27	0	0	0
2008W T2	EECE 353	2-3*-1	115	24	12	15	0
2009W T1	EECE 476	3-0-1	87	32	0	0	0
2009W T1	EECE 527	3-0-0	8	27	0	0	0
2009W T2	EECE 353	2-3*-1	86	24	0	0	0
2010W T1	EECE 476	3-0-1	97	36	0	0	0
2010W T1	EECE 353	2-3*-1	74	24	0	0	0
2010W T2	EECE 527	3-0-0	11	27	0	0	0
2011W T1	EECE 476	3-0-0	76	27	0	0	0
2011W T1	EECE 353	2-3*-1	91	24	0	0	0
2011W T2	EECE 527	3-0-0	9	24	0	0	0
2011W T2	EECE 571M	3-0-0	7	24	0	0	0

2013W T1	EECE 476	3-0-0	55	39	0	0	0
2013W T2	EECE 381	1-8-0	80	11	0	0	0
2013W T2	EECE 527	3-0-0	8	24	0	0	0
2014W T1	EECE 259	4-2-1	320	48	0	0	0
2015W T1	CPEN 211	4-2-1	288	48	0	0	0
2016W T1	CPEN 211	4-2-1	321	48	0	0	0
2017W T1	CPEN 211	4-2-1	291	48	0	0	0

Legend:

CPEN 211: Introduction to Microcomputers
 EECE 259: Introduction to Microcomputers
 EECE 353: Digital Systems Design
 EECE 381: Computer Systems Design Studio
 EECE 476: Computer Architecture
 EECE 527: Advanced Computer Architecture
 EECE 571B: Advanced Computer Microarchitecture
 EECE 571T/M: Optimizing Compilers

(c) Graduate Students Supervised at UBC

Student Name	Program	Year		Principal Supervisor	Co-Supervisor(s)
		Start	Finish		
Ali Bakhoda ²	Ph.D.	2006	2014	T. Aamodt (100%)	
Wilson Fung ³	Ph.D.	2008	2014	T. Aamodt (100%)	
Tim Rogers ⁴	Ph.D.	2010	2015	T. Aamodt (100%)	
Taylor Hetherington ⁵	Ph.D.	2011		T. Aamodt (100%)	
Ayub Gubran ⁶	Ph.D.	2011		T. Aamodt (100%)	
Ahmed El-Shafiey ⁷	Ph.D.	2011	2018	T. Aamodt (100%)	
Dave Evans	Ph.D.	2014		T. Aamodt (100%)	
Wilson Fung	M.A.Sc.	2006	2008	T. Aamodt (100%)	
Henry Wong ⁸	M.A.Sc.	2006	2008	T. Aamodt (100%)	
Xi Chen ⁹	M.A.Sc.	2006	2009	T. Aamodt (100%)	
George Yuan ¹⁰	M.A.Sc.	2007	2009	T. Aamodt (100%)	
Johnny Kuan	M.A.Sc.	2008	2011	T. Aamodt (100%)	
Andrew Turner	M.A.Sc.	2008	2012	T. Aamodt (100%)	
Arun Ramamurthy	M.A.Sc.	2009	2011	T. Aamodt (100%)	
Rimon Tadros	M.A.Sc.	2010	2015 ¹¹	T. Aamodt (100%)	
Jimmy Kwa	M.A.Sc.	2010	2013	T. Aamodt (100%)	
Inderpreet Singh	M.A.Sc.	2011	2013	T. Aamodt (100%)	
Hadi Jooybar	M.A.Sc.	2011	2013	T. Aamodt (100%)	
Dongdong Li	M.A.Sc.	2012	2015	T. Aamodt (100%)	
Shadi Assadi	M.A.Sc.	2014	2017	T. Aamodt (100%)	

² Joined Microsoft Research as a Sr. Hardware Engineer working in Vancouver.

³ Joined Samsung Research America and now a Staff Engineer.

⁴ Joined ECE department at Purdue University as an Assistant Professor.

⁵ Expected completion Fall 2018; has accepted a position at Oracle Research in Vancouver starting Sept 2018.

⁶ Long duration due to multiple internships (Samsung Research America 2015; Google 2017)

⁷ Expected completion Fall 2018; has passed department defense and now working at Huawei in Toronto.

⁸ Joined PhD program at University of Toronto after completing M.A.Sc.

⁹ Joined NVIDIA Corporation after completing M.A.Sc.

¹⁰ Joined NVIDIA Corporation after completing M.A.Sc.

¹¹ Long duration to complete M.A.Sc due to fulltime employment at Microsoft.

Amruth Sandhupatla	M.A.Sc.	2016	2018	T. Aamodt (100%)	
Maria Lubeznov	M.A.Sc.	2017		T. Aamodt (100%)	
Negar Goli	M.A.Sc.	2017		T. Aamodt (100%)	
Aamir Raihan	M.A.Sc.	2017		T. Aamodt (100%)	
François Demoullin	M.A.Sc.	2017		T. Aamodt (100%)	
Andrew Boktor ¹²	M.Eng.	2011	2015	T. Aamodt (100%)	
Ivan Sham	M.Eng.	2006	2008	T. Aamodt (100%)	

(d) *Continuing Education Activities (provided)*

Tutorial on GPGPU-Sim: A Performance Simulator for Massively Multithreaded Processor Research, IEEE/ACM International Symposium on Microarchitecture, New York City, USA, December, 2009

Tutorial on GPGPU-Sim v3.x: A Performance Simulator for Massively Multithreaded Processor Research, ACM International Conference on Parallel Architectures and Compilation Techniques, Minneapolis, MN, USA, September, 2012.

(e) *Visiting Lecturer (indicate university/organization and dates)*(f) *Other*

1. Project supervisor for EECE 496 (Engineering Project) : 31 projects between 2006 and 2013.

2. Staff supervision

Andrew Turner	NSERC USRA	Summer 2007
Aaron Ariel	Undergraduate Research Assistant	Summer 2009
Inderpreet Singh	NSERC USRA	Summer 2010
Roham Sameni	NSERC USRA	Summer 2011
Wandemberg Rodrigues	MITACS Globalink Summer Intern	Summer 2012
Yu Chaowen	MITACS Globalink Summer Intern	Summer 2012
Lotus Fenn	NSERC USRA	Summer 2015
Lotus Fenn	NSERC USRA	Summer 2016
Angus Lin	NSERC USRA	Summer 2016
David Zheng	APSC Work Learn International	Summer 2016
Spencer Spenst	Undergraduate Research Assistant	Summer 2016
Scott Peverelle	Engineering co-op student	Summer 2016
Felix Huang	Visiting CMU undergrad (URA)	Summer 2016
Kahlan Gibson	Undergraduate Research Assistant	Fall 2016
Xiaotong (Kay) Xi	Engineering co-op student	Summer 2017
Jennifer Angelica	APSC Work Learn International	Summer 2017
James Asefa	NSERC USRA	Summer 2017
Kahlan Gibson	NSERC USRA	Summer 2017
Felix Huang	Visiting CMU undergrad (URA)	Summer 2017
Jonathan Lew	NSERC USRA	Summer 2018
Shaylin Cattell	NSERC USRA	Summer 2018

¹² Recently switched to MENG from PHD for personal reasons.

9. SCHOLARLY AND PROFESSIONAL ACTIVITIES

(a) Briefly describe areas of special interest and accomplishments

1. We developed the first simulator, GPGPU-Sim, for studying general-purpose graphics processor unit (GPGPU) architectures [C.9]. This simulator is now the de-facto standard simulator for studying GPGPU architecture. A power model was added [C.23] and GPGPU-Sim has been integrated by others into the Gem5 simulator used by both academia and industry.
2. We introduced dynamic warp formation (DWF) [C.5,J.1]. GPUs replicate control information for many hardware units. This hardware is poorly suited to parallel programs that have significant control flow (e.g., “if-then” statements). DWF exploits the multithreaded nature of GPUs to create new warps that improve hardware utilization. More recently we proposed “thread block compaction” [C.15], which a mechanism that solves the main challenges faced when implementing DWF in practical systems. These works [C.5,J.1,C.15] have been well cited.
3. Hardware Transactional Memory for GPU Architectures: My students (Fung, Singh) and I published paper [C.16] describing a mechanism for supporting the transactional programming model in hardware on modern graphics processor architectures. *This paper was among 12 papers selected by IEEE Micro magazine as “Top Picks” from computer architecture conferences published in 2011.*
4. Cache Conscious Wavefront Scheduling: Our paper [C.18] describes a mechanism for providing feedback on locality from the memory system to the hardware thread scheduler inside a GPU. The purpose was to guide scheduling decisions to improve the locality in the access stream of memory references seen by the cache. *This paper was among 11 papers selected by IEEE Micro magazine as “Top Picks” from computer architecture conferences published in 2012 and also selected to appear as a research highlight by Communications of the ACM Magazine (circulation of roughly 100,000).*
5. Cache Coherence for GPU Architectures: Our paper [C.20] describes the first work showing how to support cache coherence on GPU architectures. *This paper was among 12 papers selected by IEEE Micro magazine as “Top Picks” from computer architecture conferences published in 2013.*

(b) Research or equivalent grants (indicate under COMP whether grants were obtained competitively (C) or non-competitively (NC))

Granting Agency	Subject	COMP	\$ Per Year	Years	Principal Investigator	Co-Investigator(s)
NSERC (Discovery)	Statistical Synthesis of Application Specific Computer Architectures	C	22,000	2006 – 2011	T. Aamodt (100%)	
UBC (Start up)		NC	50,000	2006	T. Aamodt (100%)	
CFI (LOF)	A Computer Cluster and Workstations for Research on Grid Utility Services, Resource Management, and Computer Architecture	C	235,407 ¹³	2007	T. Aamodt (33%)	Ripeanu, Gopalakrishnan
NSERC (RTI)	A cluster for grid utility services and computer architecture research	C	73,716	2007	Ripeanu	T. Aamodt (50%)
NVIDIA (Professor Partnership Program)	A Parallelizing Compiler to Enable GPU Computing	NC	704 (in-kind)	2007	T. Aamodt (100%)	
NVIDIA (Professor)	Exploiting Synchronization on	NC	~500 (in-kind)	2007	T. Aamodt (100%)	

¹³ Includes matching funds from BCKDF and required in-kind CFI special discounts

Partnership Program)	GPUs to Accelerate Algorithms with Fine Grained Inter Thread Communication Patterns						
AMD	Programming Languages and Parallelizing Compilers for GPGPU Software Development	NC	~500 (in-kind)	2007	T. Aamodt (100%)		
NVIDIA (Professor Partnership Program)	Hardware Comparison of the GPGPU-Sim Simulator	NC	4,798 (in-kind)	2009	T. Aamodt (100%)		
NVIDIA (Professor Partnership Program)	Simulating Fermi Using GPGPU-Sim	NC	~500 (in-kind)	2010	T. Aamodt (100%)		
Advanced Micro Devices, Inc.	Support for Complex Data Structures on Accelerator Architectures	C	43,500 USD	2010-2013	T. Aamodt (100%)		
NVIDIA Corporation (Academic Partnership Award)	Simulation Based Performance Tuning of GPU Compute Applications	C	\$25,000 USD	2010	T. Aamodt (100%)		
NSERC (SPG)	Transactional Memory and Language Support for General Purpose Graphics Processors	C	140,500	2010-2013	T. Aamodt (50%)	A. Moshovos (UoFT), T. Abdelrahman (UoFT)	
NSERC (SPG)	Manycore Soft Vector Processors for Single-chip FPGA Applications	C	130,500	2010-2013	G. Lemieux	G. Steffan (UoFT), T. Aamodt (33%)	
NSERC (Discovery)	Heterogeneous manycore accelerator architectures	C	32,000	2011-2016	T. Aamodt (100%)		
NSERC (RTI)	Shared-memory Multiprocessor for Parallel Algorithms and Architectures	C	86,500	2011	G. Lemieux	T. Aamodt (14%) , M. Ripeanu, S. Gopalakrishnan, S. Wilton, A. Hu, K. Pattabiraman	
NSERC (Engage)	Mobile graphics processor unit architecture simulation	NC	25,000	2011	T. Aamodt (100%)		
NSERC (CRD)	Data Supply Architectures for Smartphone and Tablet Devices	NC	94,000	2011 - 2014	A. Moshovos (UoFT)	N. Enright-Jerger (UoFT), T. Aamodt (33%)	
UBC (TLEF-Small Project)	Enhancing Education of Computer Systems with Flipped Lectures and Interactive Laboratories	C	29,562	2015	T. Aamodt (100%)		
NSERC (Discovery)	Energy-Efficient Programmable	C	65,000	2016 - 2021	T. Aamodt (100%)		

	Accelerators					
NSERC (Discovery Accelerator)	Energy-Efficient Programmable Accelerators	C	40,000	2016 - 2019	T. Aamodt (100%)	
NSERC (Strategic Partnership Grants for Networks)	COHESA: Computing Hardware for Emerging Intelligent Sensory Applications	C	1,100,000	2016 ¹⁴ - 2021	Moshovos (Toronto)	Bengio, Fidler, Koudas, Kutulakos, Pal, Urtasun, Aamodt (5%) , Baniasadi, Carusone, Enright Jerger, Hossain, Lis, Shannon, Shriraman, Wilton, Fedorova, Stumm, Verbugge
Google Faculty Research Award	System-on-Chip Architecture Modeling for Mobile Graphics	C	31,153 (USD)	2017 - 2018	T. Aamodt (100%)	
NSERC (Strategic Partnership Grants for Projects)	Error Resilient Machine Learning Systems	C	245,000	2017-2020	T. Aamodt (25%)	Pattabiraman, Moshovos, Urtasun, Fidler
NSERC (Research Tools and Instruments)	Designing Efficient and Resilient Deep Learning Accelerators using an AI Supercomputer	C	150,000	2018 ¹⁵	T. Aamodt (20%)	Wilton, Lis, Federova, Pattabiraman
Activision	Hybrid raytracing	NC	15,000	2018	T. Aamodt (100%)	

(c) *Research or equivalent contracts (indicate under COMP whether grants were obtained competitively (C) or non-competitively (NC)).*

Granting Agency	Subject	COMP	\$ Per Year	Years	Principal Investigator	Co-Investigator(s)
Semiconductor Research Corporation	Combining Formal Analysis, Architectural Features, and Circuit Structures for Post-Silicon Debugging	C	90,000 (USD)	2007-2010	A. Hu	T. Aamodt (22%) , S. Wilton, A. Ivanov
Semiconductor Research Corporation	System Level Post Si Validation Coverage for SoC	C	60,000 (USD)	2010-2012	A. Hu	T. Aamodt (25%) , S. Wilton, A. Ivanov
Semiconductor Research Corporation	CoolCaches: Energy Efficient Cache Coherence for Accelerators	C	70,000 (USD)	2012-2015	T. Aamodt (50%)	A. Shriraman (SFU)

(d) *Invited Presentations*

¹⁴ Officially awarded on 27 February 2017

¹⁵ Funding start date 31 March 2018.

Title	Organization or Event	Location	Date
The GPGPU-Sim Simulator Framework	ModSim 2017: Guiding the Future of Systems and Applications and Enriching the HPC Community	Seattle, WA	Aug / 2017
Simplifying Fine-Grained Synchronization on GPUs	IBM T.J. Watson Research Center	Yorktown Heights, NY	May / 2017
Introduction to Tor's Research	Huawei Canada	Ottawa	Feb / 2017
Easier Programming of Specialized Architectures	MediaTek, Inc.	Taipei, Taiwan	Oct / 2016
Programmer Transparent MIMD Synchronization on GPUs	Georgia Tech, ECE Department	Georgia	Feb / 2016
Reuse Distance Based Probabilistic Cache Replacement	HiPEAC Conference ¹⁶	Prague	Jan / 2016
MemcachedGPU: Scaling-up Scale-out Key-value Stores	Duke University, ECE Department	Durham, NC	Oct / 2015
GPU Computing Architecture	HiPEAC 11 th International Summer School on Advanced Computer Architecture and Compilation for High-Performance and Embedded Systems	Fiuggi, Italy	Jul / 2015
SLIP – Reducing Wire Energy in the Memory Hierarchy	University of Toronto	Toronto, ON	Jun / 2015
Making Accelerators Efficient and Easier to Program	University of Wisconsin-Madison	Madison, WI	Feb / 2014
GPU Computing Architectures	CUSO Winter School on Data-Centric Systems	Veysonnaz, Switzerland	Jan / 2014
Efficient and Easily Programmable Accelerator Architectures	Intel Research	Santa Clara, CA	Jun / 2013
Efficient and Easily Programmable Accelerator Architectures	Stanford CS Department Pervasive Parallelism Lab Retreat	San Francisco, CA	May / 2013
Efficient and Easily Programmable Accelerator Architectures	Qualcomm Research Silicon Valley	Santa Clara, CA	May / 2013
Efficient and Easily Programmable Accelerator Architectures	University of Toronto	Toronto, ON	May / 2013
Efficient and Easily Programmable Accelerator Architectures	UC Berkeley, EECS Department, ASPIRE Lab	Berkeley, CA	May / 2013
Efficient and Easily Programmable Accelerator Architectures	Google, Inc. Platforms Seminar	Mountain View, CA	May / 2013

¹⁶ This talk corresponds to the journal paper [J.13]. From the invitation email, "Beginning in 2011, ACM TACO and the HiPEAC conference have adopted a journal-first publication model. In this model, papers presented at the HiPEAC conference must first be accepted by ACM TACO as a regular paper. Therefore, all manuscripts submitted to the conference are automatically forwarded to ACM TACO and, once accepted, their authors may be invited to present their work in the main track of the conference. This invitation is extended to original-work papers only (excluding extensions of conference papers) that are within the scope of the HiPEAC conference, and that have been accepted by ACM TACO in 2015."

Title	Organization or Event	Location	Date
Efficient and Easily Programmable Accelerator Architectures	University of Texas at Austin	Austin, TX	Apr / 2013
Efficient and Easily Programmable Accelerator Architectures	IBM T.J. Watson Research Center	Yorktown Heights, NY	Apr / 2013
Hardware Acceleration of a Key-Value Store	Google, Inc. HotPar PC Speaker Series	Mountain View, CA	Apr / 2013
Cache Coherence for GPU Architectures	Advanced Micro Devices, Inc	Sunnyvale, CA	Apr / 2013
Cache Coherence for GPU Architectures	NVIDIA	Santa Clara, CA	Mar / 2013
Evolving GPUs into a Substrate for Cloud Computing	Microsoft Research	Redmond, WA	May / 2012
Programmer Friendly Hardware Acceleration	Electronic Arts	Burnaby, BC	Mar / 2012
Hardware Transactional Memory for GPU Architectures	NVIDIA Corp.	Santa Clara, CA	Oct / 2011
Hardware Transactional Memory for GPU Architectures	Intel Corp.	Santa Clara, CA	Oct / 2011
Hardware Transactional Memory for GPU Architectures	Rambus, Inc	Sunnyvale, CA	Oct / 2011
Hardware Transactional Memory for GPU Architectures	Advanced Micro Devices, Inc	Sunnyvale, CA	Oct / 2011
GPU Architecture Challenges for Throughput Computing	University of Toronto, ECE Department Cider Seminar	Toronto, ON	Mar / 2011
GPU Architecture Challenges for Throughput Computing	Qualcomm Ltd.	Markham, ON	Mar / 2011
GPU Architecture Challenges for Throughput Computing	École Polytechnique Fédérale de Lausanne	Lausanne, Switzerland	Feb / 2011
GPU Architecture Challenges for Throughput Computing	University of Saskatchewan	Saskatoon, SK	Jan / 2011
GPU Architecture Challenges for Throughput Computing	IEEE Computer Society, Victoria Chapter	Victoria, BC	Sep / 2010
Leveraging Fine-Grained Multithreading for Efficient SIMD Control Flow	Microsoft Research	Redmond, WA	Feb / 2008
Using Modern Graphics Processors for Non-graphics Applications	IEEE Computer Society, Vancouver Chapter	UBC	Oct / 2007
Modeling and Optimization of Speculative Threads	University of Victoria	Victoria, BC	Apr / 2006

(e) *Other Presentations*

Title	Organization or Event	Location	Date
MemcachedGPU: Scaling-up Scale-out Key-value Stores	Google	Mountain View, CA	Nov / 2015
Enabling MIMD Synchronization on SIMT	NVIDIA	Santa Clara, CA	Nov / 2015

Title	Organization or Event	Location	Date
Architectures			
CoolCaches: Energy Efficient Cache Coherence for Accelerators	SRC GRC ICSS Integrated Systems Contract Review	Hillsboro, OR	May / 2015
CoolCaches: Energy Efficient Cache Coherence for Accelerators	SRC GRC ICSS Integrated Systems Contract Review	Hillsboro, OR	Apr / 2013
Floating-Point to Fixed-Point Compilation and Embedded Architectural Support	CITO Research Review on Software Technology and Distributed Systems Relevant to Communications	Ottawa	Mar / 2001
Embedded ISA Support for Enhanced Floating-Point to Fixed-Point ANSI C Compilation	CITO Knowledge Network Conference	Ottawa	Oct / 2000

(f) *Other*

(g) *Conference Participation (Organizer, Keynote Speaker, Program Committee Member, etc.)*

Conference or Event	Organization	Role(s)	Location	Date(s)
52 nd IEEE/ACM International Symposium on Microarchitecture (MICRO)	IEEE & ACM	Program Co-Chair	Columbus, Ohio	Oct / 2019
46 th International Symposium on Computer Architecture (ISCA)	IEEE & ACM	External Review Committee Member	Phoenix, Arizona	Jun / 2019
45 th International Symposium on Computer Architecture (ISCA)	IEEE & ACM	Program Committee Member	Los Angeles, CA	Jun / 2018
2018 IEEE International Symposium on Performance Analysis of Systems and Software	IEEE	Program Committee Member	Belfast, Northern Ireland	Apr / 2018
44 th International Symposium on Computer Architecture (ISCA)	IEEE & ACM	External Review Committee Member	Toronto, ON	Jun / 2017
49 th IEEE/ACM International Symposium on Microarchitecture (MICRO)	IEEE & ACM	External Review Committee Member	Taipei, Taiwan	Oct / 2016
43 rd International Symposium on Computer Architecture (ISCA)	IEEE & ACM	Technical Program Committee Member	Seoul, Korea	Jun / 2016
22 nd IEEE International Symposium on High Performance Computer Architecture (HPCA)	IEEE	Technical Program Committee Member	Barcelona, Spain	Feb / 2016
48 th IEEE/ACM International Symposium on Microarchitecture (MICRO)	IEEE	Technical Program Committee Member	Hawaii	Dec / 2015
42 nd International Symposium on Computer Architecture (ISCA)	IEEE	External Review Committee Member	Portland, Oregon	Jun / 2015
IEEE International Conference on Embedded Computer Systems: Architectures, MOdeling and Simulation (SAMOS XIV)	IEEE	Keynote Speaker	Samos Island, Greece	Jul / 2014
2014 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)	IEEE	General Chair	Monterey, CA	Mar / 2014
5th USENIX Workshop on Hot Topics in Parallelism (HotPar '13)	USENIX	Technical Program Committee Member	San Jose, CA	Jun / 2013

Conference or Event	Organization	Role(s)	Location	Date(s)
40th International Symposium on Computer Architecture (ISCA)	IEEE & ACM	Technical Program Committee Member	Tel Aviv, Israel	Jun / 2013
27th IEEE International Parallel and Distributed Processing Symposium (IPDPS)	IEEE	Technical Program Committee Member	Boston, MA	May / 2013
2013 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)	IEEE	Program Chair	Austin, TX	Apr / 2013
19th IEEE International Symposium on High Performance Computer Architecture (HPCA)	IEEE	Technical Program Committee Member	Shenzhen, China	Feb / 2013
8th International Conference on High-Performance and Embedded Architectures and Compilers (HiPEAC)	ACM	Board of Distinguished Reviewers (=TPC)	Berlin, Germany	Jan / 2013
45th IEEE/ACM International Symposium on Microarchitecture (MICRO)	IEEE & ACM	Technical Program Committee Member	Vancouver, BC	Dec / 2012
2012 IEEE International Symposium on Workload Characterization (IISWC)	IEEE	Technical Program Committee Member	San Diego, CA	Nov / 2012
26th International Conference on Supercomputing (ICS)	ACM	Technical Program Committee Member	Venice, Italy	Jun / 2012
26th IEEE International Parallel and Distributed Processing Symposium (IPDPS)	IEEE	Technical Program Committee Member	Shanghai, China	May / 2012
2012 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)	IEEE	Technical Program Committee Member	New Brunswick, NJ	Apr / 2012
18th IEEE International Symposium on High Performance Computer Architecture (HPCA)	IEEE	Technical Program Committee Member	New Orleans	Feb / 2012
7th International Conference on High-Performance and Embedded Architectures and Compilers (HiPEAC)	ACM	Board of Distinguished Reviewers (=TPC)	Paris, France	Jan / 2012
44th IEEE/ACM International Symposium on Microarchitecture (MICRO)	IEEE & ACM	Technical Program Committee Member	Porto Alegre, Brazil	Dec / 2011
2011 IEEE International Symposium on Workload Characterization (IISWC)	IEEE	Technical Program Committee Member	Austin, TX	Nov / 2011
20th International Conference on Parallel Architectures and Compilation Techniques (PACT) - Student Research Competition	ACM	Selection Committee Member	Galveston Island, TX	Oct / 2011
8th International Conference on Network and Parallel Computing (NPC 2011)	IFIP	Technical Program Committee Member	Changsha, China	Oct / 2011
25th International Conference on Supercomputing (ICS)	ACM	Technical Program Committee Member	Tucson, Arizona	Jun / 2011
24th Canadian Conference on Electrical and Computer Engineering (CCECE)	IEEE	Technical Program Committee Member	Niagara Fall, ON	May / 2011
21st ACM Great Lakes Symposium on VLSI	ACM	Technical Program Committee Member	Lausanne, Switzerland	May / 2011
2011 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)	IEEE	Technical Program Committee Member	Austin, Texas	Apr / 2011

Conference or Event	Organization	Role(s)	Location	Date(s)
4 th Workshop on General-Purpose Computation on Graphics Processing Units (GPGPU-4)	ACM	Technical Program Committee Member	Newport Beach, California	Mar / 2011
16 th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS 2011)	ACM	External Review Committee	Newport Beach, California	Mar / 2011
43 rd IEEE/ACM International Symposium on Microarchitecture (MICRO)	IEEE & ACM	Technical Program Committee Member	Atlanta, Georgia	Dec / 2010
2010 IEEE International Symposium on Workload Characterization (IISWC)	IEEE	Technical Program Committee Member	Atlanta, Georgia	Dec / 2010
6 th Workshop on Unique Chips and Systems		Co-Organizer	Atlanta, Georgia	Dec / 2010
2 nd Workshop on Ultra Performance and Dependable Acceleration Systems		Technical Program Committee Member	Hiroshima, Japan	Nov / 2010
39 th IEEE International Conference on Parallel Processing (ICPP)	IEEE	Technical Program Committee Member	San Diego, CA	Sep / 2010
1 st International Workshop on Frontier of GPU Computing		Technical Program Committee Member	Bradford, UK	Jun / 2010
2010 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)	IEEE & ACM	Publications Chair	Toronto, ON	Apr / 2010
20 th ACM Great Lakes Symposium on VLSI	ACM	Technical Program Committee Member	Providence, RI	May / 2010
3 rd Workshop on General-Purpose Computation on Graphics Processing Units (GPGPU-3)	ACM	Technical Program Committee Member	Pittsburgh, PA	Mar / 2010
2010 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)	IEEE	Technical Program Committee Member; Publications Chair	White Plains, NY	Mar / 2010
Workshop on Computer Architecture Education		Panel member	New York, NY	Dec / 2009
1st Workshop on Ultra Performance and Dependable Acceleration Systems		Technical Program Committee Member	Hiroshima, Japan	Dec / 2009
2009 CMOS Emerging Technologies Workshop		Session Chair	Vancouver, BC	Sep / 2009
2009 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)	IEEE	Technical Program Committee Member; Session Chair	Boston, NY	Apr / 2009
2009 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)	IEEE & ACM	Registration Chair	Seattle, WA	Mar / 2009
17 th IEEE/ACM International Conference on Parallel Architectures and Compilation Techniques (PACT)	IEEE & ACM	Publicity Chair	Toronto, ON	Oct / 2008
1 st International Conference on Contemporary Computing		Technical Program Committee Member	Noida, India	Aug / 2008
2nd Workshop on Chip Multiprocessor Memory Systems and Interconnects		Technical Program Committee Member	Beijing, China	Jun / 2008
1 st IEEE/ACM International Symposium on Networks-on-Chip (NOCS)	IEEE & ACM	Publications Chair	Princeton, NJ	May / 2007

10. SERVICE TO THE UNIVERSITY

(a) *Memberships in committees, including offices held and dates*

Organizational Unit	Committee Name	Role	Dates	
			Start	End
EECE	Curriculum Committee	Member	Sep / 2007	Apr / 2009
EECE	EECE 301 Digital Systems - Detailed Design Group	Member	Dec / 2010	Jun / 2011
EECE	Scholarship Selection Committee	Member	Feb / 2011	Jun / 2012
EECE	Recruiting Committee	Member	Sep / 2011	Jun / 2012
EECE	Recruiting Committee	Vice Chair, Multicore Computing Subcommittee	Sep / 2013	May/2014
EECE	Recruiting Committee	Member	Sep / 2014	May/2015
UBC VP Research & International	Advanced Research Computing (ARC) Advisory Committee	Member	Oct / 2014	-
EECE	Recruiting Committee	Member	Oct / 2015	May/2016
EECE	Recruiting Committee	Member	Jan / 2018	-
EECE	Undergraduate Scholarship Committee	Chair	Sep / 2017	-

(b) *Other service, including dates*

Organizational Unit	Title / Nature of Duties	Dates	
		Start	End
EECE	APSC "Meet Your Professors" evening: participating faculty member	Sep / 2008	Sep / 2008
EECE	APSC 122: 50 minute presentation on computer and software engineering	Jan / 2009	Jan / 2009
EECE	20 minute presentation on computer architecture research for UBC "Rising Stars of Research"	Aug / 2009	Aug / 2009
EECE	APSC 122: 50 minute presentation on computer and software engineering	Feb / 2010	Feb / 2010
EECE	APSC 122: 50 minute presentation on computer and software engineering	Jan / 2011	Jan / 2011
EECE	APSC 122: 50 minute presentation on computer and software engineering	Jan / 2012	Jan / 2012

Internal Examiner at UBC

Role	Department	Student	Degree	Date
Head's Nominee	ECE	Chris Brouse	M.A.Sc.	May / 2007
Co-Reader	ECE	David K. M. Mak	M.A.Sc.	Sep / 2007

Role	Department	Student	Degree	Date
Heads' Nominee (Qualifying Exam)	ECE	Mandana Sotoodeh	Ph.D.	Jul / 2007
Supervisor	ECE	Wilson W. L. Fung	M.A.Sc.	May / 2008
Co-Reader	ECE	Jason Yu	M.A.Sc.	May / 2008
Head's Nominee (Dept. Defense)	ECE	Brad Quinton	Ph.D.	Jul / 2008
Head's Nominee (Dept. Defense)	ECE	Cristian Grecu	Ph.D.	Jul / 2008
Committee Member (Qualifying Exam)	ECE	David Grant	Ph.D.	Aug / 2008
Supervisor	ECE	Henry Ting-Hei Wong	M.A.Sc.	Sep / 2008
Co-Reader	ECE	Andrew Lam	M.A.Sc.	Jan / 2009
Co-Reader	ECE	Christopher Chou	M.A.Sc.	Apr / 2009
Head's Nominee (Qualifying Exam)	ECE	San-Tsai Sun	Ph.D.	Jul / 2009
Chair and Head's Nominee	ECE	Marcel Gort	M.A.Sc.	Jul / 2009
Committee Member (Qualifying Exam)	ECE	Faizal Karim	Ph.D.	Jul / 2009
Supervisor (Qualifying Exam)	ECE	Ali Bakhoda	Ph.D.	Aug / 2009
Supervisor	ECE	George Yuan	M.A.Sc.	Sep / 2009
Chair and Head's Nominee	ECE	Darius Chiu	M.A.Sc.	Sep / 2009
Committee Member (Qualifying Exam)	ECE	Hanni Bagnordi	Ph.D.	Nov / 2009
Committee Member (Qualifying Exam)	ECE	Eddie Hung	Ph.D.	Nov / 2009
Head's Nominee (Qualifying Exam)	ECE	Samer Al-Kiswany	Ph.D.	Nov / 2009
Committee Member (Qualifying Exam)	ECE	Joydip Das	Ph.D.	Nov / 2009
Committee Member	ECE	Mohammad Jalali	M.A.Sc.	Aug / 2010
Committee Member (Qualifying Exam)	ECE	Assem Bsoul	Ph.D.	Sep / 2010
Committee Member (Dept. Exam)	ECE	David Grant	Ph.D.	Jan / 2011
Committee Member (Final Exam)	ECE	David Grant	Ph.D.	Apr / 2011
Head's Nominee (Dept. Exam)	ECE	Scott Chin	Ph.D.	Apr / 2011
Committee Member (Qualifying Exam)	ECE	Abdullah Gharaibeh	Ph.D.	May / 2011
Committee Member (Qualifying Exam)	ECE	Aaron Severance	Ph.D.	Jun / 2011
Committee Member (Dept. Exam)	ECE	Joydip Das	Ph.D.	Apr / 2012
Supervisor (Qualifying Exam)	ECE	Tim Rogers	Ph.D.	Nov / 2012
Committee Member	ECE	Emalayan Vairavanathan	M.A.Sc.	Nov / 2012
Committee Member	ECE	Samer Al-Kiswany	Ph.D.	Jan / 2012
Supervisor	ECE	Jimmy Kwa	MASc	Apr / 2013
Supervisor (Qualifying Exam)	ECE	Tayler Hetherington	Ph.D.	Apr / 2013
Supervisor	ECE	Inderpreet Singh	MASc	May / 2013
Supervisor	ECE	Hadi Jooybar	MASc	Jul / 2013
Chair (Qualifying Exam)	ECE	Mustafa Fanaswala	Ph.D.	Sep / 2013
Supervisor (Qualifying Exam)	ECE	Ayub Gubran	Ph.D.	Sep / 2013
Supervisor (Department Exam)	ECE	Ali Bakhoda	Ph.D.	Dec / 2013
Supervisor (University Exam)	ECE	Ali Bakhoda	Ph.D.	Mar / 2014
Supervisor (Department Exam)	ECE	Wilson Fung	Ph.D.	May / 2014

Role	Department	Student	Degree	Date
Head's Nominee (Qualifying Exam)	ECE	Jeff Goeders	Ph.D.	Jun / 2014
Head's Nominee	ECE	Bo Fang	M.A.Sc.	Jul / 2014
Committee Member (Dept. Exam)	ECE	Assem Bsoul	Ph.D.	Jul / 2014
Chair (Department Exam)	ECE	Chris Brouse	Ph.D.	Aug / 2014
Committee Member (Qualifying Exam)	ECE	Ameer Abdelhadi	Ph.D.	Aug / 2014
Supervisor (University Exam)	ECE	Wilson Fung	Ph.D.	Oct / 2014
Supervisor (Qualifying Exam)	ECE	Ahmed ElTantawy	Ph.D.	Oct / 2014
Committee Member (Dept. Exam)	ECE	Abdullah Gharaibeh	Ph.D.	Oct / 2014
Head's Nominee and Chair	ECE	Qining Lu	M.A.Sc.	Jan / 2015
Committee Member (Dept. Exam)	ECE	Aaron Severance	Ph.D.	Jan / 2015
Chair (Qualifying Exam)	ECE	Chunsheng Zhu	Ph.D.	Feb / 2015
Committee Member (Univ. Exam)	ECE	Aaron Severance	Ph.D.	Mar / 2015
Supervisor (Department Exam)	ECE	Dongdong Li	M.A.Sc.	Mar / 2015
University Examiner	ECE	Abdullah Gharaibeh	Ph.D.	Apr / 2015
Supervisor (Department Exam)	ECE	Tim Rogers	Ph.D.	Jul / 2015
Supervisor (University Exam)	ECE	Tim Rogers	Ph.D.	Sep / 2015
Head's Nominee (Qualifying Exam)	ECE	Guanpeng Li	Ph.D.	Mar / 2016
Committee Member (Dept. Exam)	ECE	Ameer Abdelhadi	Ph.D.	Apr / 2016
Chair (Qualifying Exam)	ECE	Sujay Bhatt	Ph.D.	Jun / 2016
Head's Nominee (Qualifying Exam)	ECE	Jake Retallick	Ph.D.	Aug / 2016
Head's Nominee (Qualifying Exam)	ECE	Scott Sallinen	Ph.D.	May / 2017
Supervisor (Department Exam)	ECE	Ahmed ElTantawy	Ph.D.	Jun / 2017
Chair (Qualifying Exam)	ECE	Svetozar Miucin	Ph.D.	Aug / 2017
Chair (University Exam)	ECE	Anna Hippmann	Ph.D.	Dec / 2017
Committee Member	ECE	Mohammad Ewais	M.A.Sc.	May / 2018
Committee Member (Qualifying Exam)	ECE	Xiaowei Ren	Ph.D.	May / 2018

11. SERVICE TO THE COMMUNITY

(a) *Memberships in scholarly societies, including offices held and dates*

Scholarly Society	Role	Dates	
		Start	End
Institute of Electrical and Electronics Engineers (IEEE)	Member	2006	Present
Association for Computing Machines (ACM)	Member	2006	Present

(b) *Memberships in other societies, including offices held and dates*

(c) *Memberships in scholarly committees, including offices held and dates*

U.S. Department of Energy, 2012 X-Stack: Programming Challenges, Runtime Systems, and Tools, Panel Member, April 3-5, Washington DC, 2012.

U.S. Department of Energy, 2013 Exascale Operating and Runtime Systems Review, April 11, Washington DC, 2013.

Panel Member, NSERC Electrical and Computer Engineering Evaluation Group (EG1510), Aug/2017 to Jun/2020.

APEGBC Accreditation Committee for University of Manitoba Computer Engineering, November 2018.

NOTE: See section 9(g) for membership in conference technical program committees.

(d) *Memberships in other committees, including offices held and dates*

(e) *Editorships (list journal and dates)*

Associate Editor, Sage Int'l J. High Performance Computing Applications (IJHPCA), May/2012-Aug/2017.
Associate Editor, IEEE Computer Architecture Letters (CAL), Aug/2012-2015.

(f) *Reviewer (journal, agency, etc. including dates)*

Grants	Role	Dates	
		Start	End
NSERC Discovery	Reviewer	2007	2012
U.S. Department of Energy, Computer Science Unsolicited 2011 Mail Review	Reviewer	2011	2011

Journal	Role	Dates	
		Start	End
Communications of the ACM	Reviewer	2017	2017
IEEE Micro Special Issue, Micro's Top Picks from the Computer Architecture Conferences, May/June 2016	Selection Committee	2016	2016
IEEE Micro Special Issue, Micro's Top Picks from the Computer Architecture Conferences, May/June 2014	Selection Committee	2014	2014
IEEE Micro Special Issue, Micro's Top Picks from the Computer Architecture Conferences, May/June 2013	Selection Committee	2013	2013
ACM Transactions on Architecture and Code Optimization (TACO)	Reviewer	2005	2012
Elsevier Journal of Systems Architecture	Reviewer	2006	2009
IEEE Transactions on Very Large Scale Integration Systems (TVLSI)	Reviewer	2007	2009
ACM Transactions on Reconfigurable Technology and Systems (TRETs)	Reviewer	2008	2010
IEEE Transactions on Computers (TC)	Reviewer	2008	2009
Elsevier Journal of Parallel and Distributed Computing (JPDC)	Reviewer	2009	2012
ACM Transactions on Embedded Computing Systems (TECS)	Reviewer	2009	2009
IEEE Transactions on Parallel and Distributed Systems	Editorial Review Board	2009	2009
IEEE Micro Special Issue on Systems for Very Large Scale Computing	Reviewer	2011	2011

Conference*	Role	Dates	
		Start	End
*NOTE: This table does not include reviewing activities for conferences where I am part of the program committee listed in Section 9(g).			
ACM/IEEE International Symposium on Computer Architecture (ISCA)	Reviewer	2003	2012
ACM International Conference on Supercomputing (ICS)	Reviewer	2003	2006
Pacific Conference on Computer Graphics and Applications	Reviewer	2006	2006
IEEE Int'l Symposium on High Performance Computer Architecture (HPCA)	Reviewer	2007	2007
ACM/IEEE Int'l Symposium on Microarchitecture (MICRO)	Reviewer	2007	2009
ACM/IEEE Int'l Conf. on Parallel Architectures and Compilation Techniques (PACT)	Reviewer	2010	2012

(g) *External examiner (indicate universities and dates)*

University	Degree	Student	Date
University of Victoria	Ph.D.	Kaveh Jokar	Aug / 2008
University of Saskatchewan	Ph.D.	Dongdong Chen	Jan / 2011
Stanford University	Ph.D.	Milad Mohammadi	May / 2015
Stanford University	Ph.D.	Subhasis Das	Nov / 2015

(h) *Consultant (indicate organization and dates)*

Organization	Role	Dates	
		Start	End
Intel Corporation	Consultant	Jun / 2003	Aug / 2003
Latham & Watkins LLP	Expert consultant for patent litigation	Mar / 2016	Apr / 2016
McGuireWoods LLP	Expert consultant for patent litigation	May / 2017	Jul / 2017

(i) *Other service to the community*

12. **AWARDS AND DISTINCTIONS**

(a) *Awards for Teaching (indicate name of award, awarding organizations, and date)*

(b) *Awards for Scholarship (indicate name of award, awarding organizations, and date)*

Program Co-Chair MICRO-52, Selected by the Steering Committee of the ACM/IEEE International Symposium on Microarchitecture as Program Chair for MICRO-52.

Google Faculty Research Award The last time someone at UBC won this award was in 2012.

Top Picks “Honorable Mention”: Paper [C.31] was selected by IEEE Micro Magazine as one of 12 papers highlighted as a Top Pick “Honorable Mention” among all computer architecture conference papers published in 2016. The annual “Top Picks” special issue recognizes the best papers published the prior year in terms of novelty and potential for impact. Papers are selected by a rigorous process involving a program committee meeting with industry and academic leaders in the field, which is attended in person. Each paper, which has already been published in a highly selective conference, is again reviewed for this award.

CACM Research Highlight: Paper [C.18] was selected as a Research Highlight in the Communications of the ACM Magazine (CACM) [J.10]. Only 1 to 2 Research Highlight papers are published per month across all fields of computer science and engineering by CACM, a magazine with circulation of over 100,000 readers. An accompanying “technical perspective” article was solicited by CACM from a top industry researcher (Dr. Steve Keckler) to put the impact of our research in context for the broader CACM readership.

MICRO Hall of Fame: Added December 2013. <http://newsletter.sigmicro.org/micro-hof.txt/view>

Top Picks: Paper [C.20] was one of only 12 papers selected by IEEE Micro Magazine as a Top Pick from all computer architecture conference papers published in 2013. The annual “Top Picks” special issue recognizes the best papers published the prior year in terms of novelty and potential for impact.

Top Picks: Paper [C.18] was one of only 11 papers selected by IEEE Micro Magazine as a Top Pick from all computer architecture conference papers published in 2012. The annual “Top Picks” issue acknowledges the best papers in terms of novelty and potential for impact.

Best Paper Runner Up: Paper [C.18] was selected as one of two “best paper runners up” out of 40 papers accepted at MICRO 2012 (in turn out of 228 papers submitted).

Top Picks: Paper [C.16] was one of only 12 papers selected by IEEE Micro Magazine as a Top Pick from all computer architecture conference papers published in 2011. The annual “Top Picks” issue acknowledges the best papers in terms of novelty and potential for impact.

Best Poster, 2nd Place: Paper [C.13] was selected as best poster, 2nd place at the 19th ACM Int’l Conf. on Parallel Architectures and Compilation Techniques, September 2010.

Prior to Final Degree

Best Student Paper, 6th Workshop on Multithreaded Execution, Architecture, and Compilation, 2002
NSERC PGS ‘B’ Postgraduate Scholarship, 2000 to 2002
NSERC PGS ‘A’ Postgraduate Scholarship, 1997 to 1999
United Technologies Scholarship (administered by Pratt & Whitney Canada) 1992-1997

(c) *Awards for Service (indicate name of award, awarding organizations, and date)*

(d) *Other Awards*

Keynote Presentation, Scaling Usable Computing Capability, IEEE International Conference on Embedded Computer Systems: Architectures, Modeling, and Simulation (SAMOS XIV) July 14, 2014.

13. **OTHER RELEVANT INFORMATION (Maximum One Page)**

THE UNIVERSITY OF BRITISH COLUMBIA
Publications Record

SURNAME: Aamodt

FIRST NAME: Tor

Initials: _____

MIDDLE NAME (S): Michael

Date: 30 Jan 2019

Those publications considered to be of primary importance are indicated by an asterisk.

NOTE: For those not familiar with the field of computer architecture, publication in top international conferences is the ***preferred*** means of disseminating research results in this area. Papers in proceedings of top conferences, often with acceptance rates below 20%, are more important than journals. See Chapter 4 in the 1994 National Academy of Sciences report, *Academic Careers for Experimental Computer Scientists and Engineers* <<http://books.nap.edu/html/acesc/>> and the 1999 *Best Practices Memo: Evaluating Computer Scientists and Engineers For Promotion and Tenure* <http://www.cra.org/uploads/documents/resources/bpmemos/tenure_review.pdf>. The top computer architecture conferences are ISCA, MICRO, HPCA, ASPLOS. The review process at these conferences is very rigorous: Program committee (PC) members are internationally recognized experts in the field. Each paper typically receives four or more double-blind reviews (3 or more from PC members) providing detailed feedback. Authors submit responses to questions raised by reviewers *prior* to the PC meeting. The PC meeting for these conferences typically has a policy of mandatory *in-person* attendance by PC members. Accepted papers are revised to reflect reviewer feedback before publication. High quality papers deemed to be on the borderline of accept/reject often undergo an additional round of review by a PC member *after* revision and *before* final acceptance for publication (a process known as “shepherding”). Added together, the four conferences above publish only 200 papers in total per year. Due to their impact, these conferences are very well attended (typically over 200 attendees for a conference with 50 papers and 25% or more of attendees from industry). Given their importance, papers published in the proceedings of these conferences are read and cited by active researchers in the area even if those researchers did not attend the conference. Three papers [C.16, C.18, C.20] have been selected by IEEE Micro magazine as “Top Picks” from computer architecture conferences and one [C.31] as a “Top Pick Honorable Mention”. This means each of these four papers were deemed to be one of the best papers among all papers published in any computer architecture conference in their year. One [C.18] was also selected as a “Research Highlight” by Communications of the ACM Magazine.

Author ordering: In my field the most senior author is usually listed *last*.

Co-authors who are students or other supervised research personnel are in *italics*.

Co-authors who are presenters are underlined.

1. **REFEREED PUBLICATIONS**

(a) *Journals (see note above)*

[J.17] Andreas Moshovos, Jorge Albericio, Patrick Judd, Alberto Delmás Lascorz, Sayeh Sharify, Zissis Poulos, *Taylor Hetherington*, **Tor Aamodt**, Natalie Enright Jerger, Exploiting Typical Values to Accelerate Deep Learning, in IEEE Computer, Volume 51, Issue 5, May 2018.

[J.16] Andreas Moshovos, Jorge Albericio, Patrick Judd, Alberto Delmas Lascorz, Sayeh Sharify, *Taylor Hetherington*, **Tor Aamodt**, Natalie Enright Jerger, Value-Based Deep Learning Hardware Accelerators, in IEEE Micro, Volume 38, Issue 1, pp 45-55, Jan/Feb 2018.

[J.15] Milad Mohammadi, **Tor M. Aamodt**, William J. Dally, CG-OoO: Energy Efficient Coarse-Grain Out-of-Order Execution, in ACM Transactions on Architecture and Code Optimization (TACO), Volume 14, Issue 4, Article 39, 26 pages, December 2017.

[J.14] J. Albericio, P. Judd, T. Hetherington, **T. Aamodt**, N.E. Jerger, R. Urtasun, A. Moshovos, Proteus: Exploiting Precision Variability in Deep Neural Networks, Elsevier Parallel Computing Systems & Applications, available online 24 May, 2017.

- [J.13] *Dongdong Li, Tor M. Aamodt*, Inter-core Locality Aware Memory Scheduling, in IEEE Computer Architecture Letters, Volume 15, Issue 1, pp. 25-28, Jan-Jun 2016.
- [J.12] *Subhasis Das, Tor M. Aamodt*, William J. Dally, Reuse Distance Based Probabilistic Cache Replacement, accepted to appear in ACM Transactions on Architecture and Code Optimization (TACO). Volume 12 Issue 4, January 2016, Article No. 33 (22 pages)
- [J.11] *Milad Mohammadi, Song Han, Tor M. Aamodt*, William J. Dally, On-Demand Dynamic Branch Prediction, IEEE Computer Architecture Letters, 4 pages, published online 13 June 2014.
- [J.10] *Timothy G. Rogers*, Mike O'Connor, **Tor M. Aamodt**, (Research Highlight) Learning Your Limit: Managing Massively Multithreaded Caches Through Scheduling, Communications of the ACM, pp. 91-98, vol. 57, no. 12, December 2014.
- [J.9] *Inderpreet Singh*, Arrvindh Shriraman, *Wilson W. L. Fung*, Mike O'Connor, **Tor M. Aamodt**, Cache Coherence for GPU Architectures, IEEE Micro (Top Pick's Special Issue), pp 69-79, Vol. 34, No. 3, available online 5 March 2014.
- [J.8] *Ali Bakhoda*, John Kim, **Tor M. Aamodt**, Designing On-Chip Networks for Throughput Accelerators, ACM Transactions on Architecture and Code Optimization (TACO), pp 1-35, Vol. 10, No. 3, Article 21, September 2013.
- [J.7] *Timothy G. Rogers*, Mike O'Connor, **Tor M. Aamodt**, Cache-Conscious Thread Scheduling for Massively Multithreaded Processors, in IEEE Micro (Top Pick's Special Issue), Vo. 33, No. 3, pp. 78-85, May/June 2013.
- [J.6] *Marcel Gort, Flavio M. De Paula, Johnny J.W. Kuan, Tor M. Aamodt*, Alan J. Hu, Steven J.E. Wilton, Jin Yang, Formal-Analysis-Based Trace Computation for Post-Silicon Debug, IEEE Transactions on Very Large Scale Integration Systems (TVLSI), pp 1997-2010, Vol. 20, No. 11, Nov 2012.
- [J.5] *Xi E. Chen, Tor M. Aamodt*, *Modeling Cache Contention and Throughput of Multiprogrammed Manycore Processors*, IEEE Transactions on Computers, Vol. 61, No. 7, pp. 913-927, July 2012.
- [J.4] *Wilson Wai Lun Fung, Inderpreet Singh*, Andrew Brownsword, **Tor Aamodt**, Kilo TM: Hardware Transactional Memory for GPU Architectures, IEEE Micro (Top Pick's Special Issue), pp. 7-16, May/June, 2012.
- [J.3] *Xi E. Chen, Tor M. Aamodt*, *Hybrid Analytical Modeling of Pending Cache Hits, Data Prefetching, and MSHRs*, ACM Transactions on Architecture and Code Optimization, ACM Transactions on Architecture and Code Optimization (TACO), Vol. 8, No. 3, Article 10 (October 2011), 28 pages.
- [J.2] *Wilson W. L. Fung, Ivan Sham, George Yuan, and Tor M. Aamodt*, *Dynamic Warp Formation: Efficient MIMD Control Flow on SIMD Graphics Hardware*, In ACM Trans. on Architecture and Code Optimization (TACO), pp. 1-37, Vol. 6, No. 2, June 2009
- [J.1] **Tor M. Aamodt**, Paul Chow, *Compile-Time and Instruction Set Methods for Improving Floating- to Fixed-Point Conversion Accuracy*, ACM Transactions on Embedded Computing Systems (TECS), pp. 1-27, Vol. 7, No. 3, April 2008

(b) *Conference Proceedings (see note above; papers listed here are highly refereed and published in well recognized international conference proceedings)*

- [C.36] *Md Aamir Raihan, Negar Goli, Tor M. Aamodt*, Modeling Deep Learning Accelerator Enabled GPUs, To appear in proceedings of the IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS 2019), 10 pages, March 24-26, 2019 Madison, Wisconsin, USA.

- [C.35] Mahmoud Khairy, Jain Akshay, **Tor M. Aamodt**, Timothy G. Rogers, Exploring Modern GPU Memory System Design Challenges through Accurate Modeling, To appear in proceedings of the IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS 2019), 2 pages, March 24-26, 2019 Madison, Wisconsin, USA. **(poster presentation)**
- [C.34] Ahmed EITantawy, **Tor M. Aamodt**, Warp Scheduling for Fine-Grained Synchronization, In proceedings of the 24th IEEE International Symposium on High-Performance Computer Architecture (HPCA-24), 14 pages, Vienna, Austria, February 24-28, 2018.
- [C.33] Shadi Asadi, Jennifer Ongko, **Tor M. Aamodt**, A State Machine Block for High-Level Synthesis, In proceedings of the IEEE International Conference on Field Programmable Technology (FPT), 8 pages, Melbourne, Australia, December 11-13, 2017.
- [C.32] Ahmed EITantawy, **Tor M. Aamodt**, *MIMD Synchronization on SIMT Architectures*, in proceedings of the ACM/IEEE Int'l Symposium on Microarchitecture (MICRO'16), 14 pages, Taipei, Taiwan, Oct. 15-19, 2016.
- [C.31] Patrick Judd, Jorge Albericio, Taylor Hetherington, **Tor M. Aamodt**, Andreas Moshovos, *Stripes: Bit-Serial Deep Neural Network Computing*, in proceedings of the ACM/IEEE Int'l Symposium on Microarchitecture (MICRO'16), 12 pages, Taipei, Taiwan, Oct. 15-19, 2016.
- [C.30] Jorge Albericio, Patrick Judd, Taylor Hetherington, **Tor Aamodt**, Natalie Enright Jerger, Andreas Moshovos, *Cnvlutin: Ineffectual-Neuron-Free Deep Convolutional Neural Network Computing*, in proceedings of the ACM/IEEE Int'l Symposium on Computer Architecture (ISCA'16), 13 pages, Seoul, Korea, June 18-22 2016 **(acceptance rate: 54/288 \approx 18.8%)**
- [C.29] Patrick Judd, Jorge Albericio, Taylor Hetherington, **Tor Aamodt**, Natalie Enright Jerger, Andreas Moshovos, *Proteus: Exploiting Numerical Precision Variability in Deep Neural Networks*, to appear in proceedings of the ACM International Conference on Supercomputing (ICS 2016), 12 pages, Istanbul, Turkey, June 1-3, 2016. **(acceptance rate: 43/183 \approx 23.5%)**
- [C.28] Taylor H. Hetherington, Mike O'Connor, **Tor M. Aamodt**, *MemcachedGPU: Scaling-up Scale-out Key-value Stores*, In proceedings of the ACM Symposium on Cloud Computing (SoCC'15), pp. 43-57, Kohala Coast, Hawaii, August 27-29, 2015. **(acceptance rate: 34/157 \approx 21.7%)**
- [C.27] Subhasis Das, **Tor M. Aamodt**, William J. Dally, *SLIP: Reducing Wire Energy in the Memory Hierarchy*, In proceedings of the ACM/IEEE International Symposium on Computer Architecture (ISCA 2015), pp. 349-361, Portland, OR, June 13-17, 2015. **(acceptance rate: 58/305 \approx 19.0%)**
- [C.26] Ahmed EITantawy, Jessica Wenjie Ma, Mike O'Connor, **Tor M. Aamodt**, *A Scalable Multi-Path Microarchitecture for Efficient GPU Control Flow*, In proceedings of the 20th IEEE International Symposium on High-Performance Computer Architecture (HPCA-20), pp. 248 - 259, Orlando, FL, February 15-19, 2014. **(acceptance rate: 25%)**
- [C.25] Wilson W. L. Fung, **Tor M. Aamodt**, *Energy Efficient GPU Transactional Memory via Space-Time Optimizations*, In proceedings of the 46th IEEE/ACM International Symposium on Microarchitecture (MICRO-46), pp. 408-420, Davis, CA, December 7-11, 2013. **(acceptance rate: 39/239 \approx 16.3%)**
- [C.24] Timothy G. Rogers, Mike O'Connor, **Tor M. Aamodt**, *Divergence-Aware Warp Scheduling*, In proceedings of the 46th IEEE/ACM International Symposium on Microarchitecture (MICRO-46), pp. 99-110, Davis, CA, December 7-11, 2013. **(acceptance rate: 39/239 \approx 16.3%)**
- [C.23] Jingwen Leng, Syed Gilani, Taylor Hetherington, Ahmed EITantawy, Nam Sung Kim, **Tor M. Aamodt**, Vijay Janapa Reddi, *GPUWatch: Enabling Energy Optimizations in GPGPUs*, In proceedings of the ACM/IEEE International Symposium on Computer Architecture (ISCA 2013), pp. 487-498, Tel-Aviv, Israel, June 23-27, 2013. **(acceptance rate: 56/288 \approx 19.4%)**

- [C.22] *Vitaly Zakharenko, Tor M. Aamodt, Andreas Moshovos, Characterizing the Performance Benefits of Fused CPU/GPU Systems Using FusionSim*, Design, Automation and Test in Europe (DATE), pp. 685-689, Grenoble, France, 18-22 March, 2013. (interactive presentation)
- [C.21] *Hadi Jooybar, Wilson W. L. Fung, Mike O'Connor, Joseph Devietti, Tor M. Aamodt, GPUDet: A Deterministic GPU Architecture*, In proceedings of the Eighteenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS 2013), pp. 1-12 Houston, Texas, March 16-20, 2013. (acceptance rate: 44/191 \approx 23.0%)
- [C.20] *Inderpreet Singh, Arrvinth Shriraman, Wilson W. L. Fung, Mike O'Connor, Tor M. Aamodt, Cache Coherence for GPU Architectures*, In proceedings of the 19th IEEE International Symposium on High-Performance Computer Architecture (HPCA-19), pp. 578-590, February 23-27, 2013, Shenzhen, China. (acceptance rate: 51/249 \approx 20.5%)
- [C.19] *Jimmy Kwa, Tor M. Aamodt, Small Virtual Channel Routers on FPGAs Through Block RAM Sharing*, To appear in proceedings of the IEEE International Conference on Field Programmable Technology (FPT), pp. 71-79, Seoul, Korea, Dec. 10-12, 2012. (acceptance rate: 24/114 \approx 21.1%)
- [C.18] *Timothy G. Rogers, Mike O'Connor, Tor M. Aamodt, Cache Conscious Wavefront Scheduling*, to In proceedings of the 45th IEEE/ACM International Symposium on Microarchitecture (MICRO-45), pp. 72-83, Vancouver, BC, December 1-5, 2012. (acceptance rate: 40/228 \approx 17.5%)
- [C.17] *Taylor Hetherington, Tim Rogers, Lisa Hsu, Michael O'Connor, Tor M. Aamodt, Characterizing and Evaluating A Key-Value Store Application on Heterogeneous CPU-GPU Systems*, In proceedings of the IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS), pp. 88-98, New Brunswick, NJ, April 1-3, 2012. (acceptance rate: 20/65 \approx 30.8%)
- [C.16] *Wilson W. L. Fung, Inderpreet Singh, Andrew Brownsword, Tor M. Aamodt, Hardware Transactional Memory for GPUs*, In proceedings of the 44th IEEE/ACM International Symposium on Microarchitecture (MICRO-44), pp. 296-307, Porto Alegre, Brazil, December 3-7, 2011 (acceptance rate: 44/209 \approx 21.0%)
- [C.15] *Wilson W. L. Fung, Tor M. Aamodt, Thread Block Compaction for Efficient SIMT Control Flow*, In proceedings of the 17th IEEE International Symposium on High-Performance Computer Architecture (HPCA-17), pp. 25-36, February 12-16 2011, San Antonio, TX (acceptance rate: 42/227 \approx 18.5%)
- [C.14] *Ali Bakhoda, John Kim, Tor M. Aamodt, Throughput-Effective On-Chip Networks for Manycore Accelerators*, In proceedings of the 43rd IEEE/ACM International Symposium on Microarchitecture (MICRO-43), pp. 421-432, Atlanta, Georgia, December 4-8, 2010. (acceptance rate: 45/248 \approx 18%)
- [C.13] *Ali Bakhoda, John Kim, Tor M. Aamodt, On-Chip Network Design Considerations for Compute Accelerators*, In proceedings of the 19th ACM International Conference on Parallel Architectures and Compilation Techniques (PACT), pp. 535-536, Vienna, Austria, September 11-15, 2010. (best poster award, 2nd place)
- [C.12] *Aaron Ariel, Wilson W. L. Fung, Andrew Turner, Tor M. Aamodt, Visualizing Complex Dynamics in Many-Core Accelerator Architectures*, In proceedings of the IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS), pp. 164-174, White Plains, NY, March 28-30, 2010. (acceptance rate: 22/64 \approx 34%)
- [C.11] *Johnny Kuan, Steve J. E. Wilton, Tor M. Aamodt, Accelerating Trace Computation in Post-Silicon Debug*, In proceedings of the 11th IEEE International Symposium on Quality Electronic Design (ISQED 2010), pp. 244-249, San Jose, CA, March 22-24, 2010. (poster presentation)

- [C.10] George L. Yuan, Ali Bakhoda, **Tor M. Aamodt**, *Complexity Effective Memory Access Scheduling for Many-Core Accelerator Architectures*, In proceedings of the 42nd IEEE/ACM International Symposium on Microarchitecture (MICRO-42), pp. 34-44, New York, NY, December 12-16, 2009.
(acceptance rate: 52/209 \approx 25%)
- [C.9] Ali Bakhoda, George Yuan, Wilson W. L. Fung, Henry Wong, **Tor M. Aamodt**, *Analyzing CUDA Workloads Using a Detailed GPU Simulator*, In proceedings of the IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS), pp. 163-174, Boston, MA, April 26-28, 2009.
(acceptance rate: 24/86 \approx 28%)
- [C.8] Xi E. Chen and **Tor M. Aamodt**, *A First-Order Fine-Grained Multithreaded Throughput Model*, In proceedings of the 15th IEEE International Symposium on High-Performance Computer Architecture (HPCA-15), pp. 329-340, Raleigh, North Carolina, February 14-18, 2009.
(acceptance rate 35/184 \approx 19%)
- [C.7] Xi E. Chen and **Tor M. Aamodt**, *Hybrid Analytical Modeling of Pending Cache Hits, Data Prefetching, and Limited MSHRs*. In proceedings of the 41st IEEE/ACM Int'l Symp. on Microarchitecture (MICRO-41), pp. 59-70, Lake Como, Italy, November 8-12, 2008.
(acceptance rate 40/210 \approx 19%)
- [C.6] Henry Wong, Anne Bracy, Ethan Schuchman, **Tor M. Aamodt**, Jamison D. Collins, Perry H. Wang, Gautham China, Ankur Khandelwal, Groen, Hong Jiang, and Hong Wang. *Pangaea: A Tightly-Coupled IA32 Heterogeneous Chip Multiprocessor*. In proceedings of the 17th IEEE/ACM Int'l Conf. on Parallel Architectures and Compilation Techniques (PACT 2008), pp. 52-61, Toronto, ON, October 25-29, 2008.
(acceptance rate 30/159 \approx 19%)
- [C.5] Wilson W. L. Fung, Ivan Sham, George Yuan, and **Tor M. Aamodt**. *Dynamic Warp Formation and Scheduling for Efficient GPU Control Flow*. In proceedings of the 40th IEEE/ACM Int'l Symp. On Microarchitecture (MICRO-40), pp. 407-418, Chicago, IL, December 1-5, 2007.
(acceptance rate 35/166 \approx 21%)
- [C.4] **Tor M. Aamodt** and Paul Chow, *Optimization of Data Prefetch Helper Threads with Path-Expression Based Statistical Modeling*, In proceedings of the 21st ACM International Conference on Supercomputing (ICS), pp. 210-221, Seattle, WA, June 16-20, 2007. (acceptance rate: 29/125 \approx 23%)
- [C.3] **Tor M. Aamodt**, Paul Chow, Per Hammarlund, Hong Wang, John P. Shen, *Hardware Support for Prescient Instruction Prefetch*, In proceedings of the 10th International Symposium on High Performance Computer Architecture (HPCA-10), pp. 84-95, Madrid Spain, February 14-18, 2004.
(acceptance rate: 27/153 \approx 18%)
- [C.2] **Tor M. Aamodt**, Pedro Marcuello, Paul Chow, Antonio Gonzalez, Per Hammarlund, Hong Wang, John P. Shen, *A Framework for Modeling and Optimization of Prescient Instruction Prefetch*, In proceedings of the 2003 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems, pp. 13-24, San Diego, CA, June 10-14, 2003. (acceptance rate: 26/222 \approx 12%)
- [C.1] **Tor Aamodt**, and Paul Chow, *Embedded ISA Support for Enhanced Floating-Point to Fixed-Point ANSI C Compilation*, In proceedings of the 3rd International Conference on Compilers, Architectures and Synthesis for Embedded Systems (CASES-2000), pp. 128-137, San Jose, CA. Nov. 17-18, 2000.
(acceptance rate: 25/56 \approx 45%)

(c) *Other (Refereed Workshop Papers, Presentations Refereed by Abstract)*

- [W.12] Patrick Judd, Jorge Albericio, Taylor Hetherington, **Tor Aamodt**, Natalie Enright Jerger and Andreas Moshovos. "Proteus: Exploiting Numerical Precision Variability in Deep Neural Networks", 2nd Workshop On Approximate Computing (WAPCO 2016), In conjunction with HiPEAC 2016, Prague, 18-20 January 2016.

- [W.11] Ayub Gubran, **Tor M. Aamodt**, *Framebuffer Compression Using Dynamic Color Palettes*, Eurographics/ACM SIGGRAPH High Performance Graphics conference, Los Angeles, CA, August 7-9, 2015. (refereed poster "Quick Talks" presentation).
- [W.10] Vitaly Zakharenko, Andreas Moshovos, **Tor Aamodt**, *FusionSim: A Cycle-Accurate CPU + GPU System Simulator*, AMD Fusion Developer Summit, June 11-14, 2012. **(refereed by abstract)**
- [W.9] Taylor Hetherington, Timothy Rogers, Lisa Hsu, Mike O'Connor, **Tor M. Aamodt**, *Characterizing and evaluating the Memcached Key-Value store Application on AMD systems*, AMD Fusion Developer Summit, June 11-14, 2012. **(refereed by abstract)**
- [W.8] Johnny J. W. Kuan, **Tor M. Aamodt**, *Progressive-BackSpace: Efficient Predecessor Computation for Post-Silicon Debug*, 12th IEEE International Workshop on Microprocessor Test and Verification (MTV 2011), 6 pages, Austin, Texas, December 5–7, 2011.
- [W.7] Henry Wong and **Tor M. Aamodt**, *The Performance Potential for Single Application Heterogeneous Systems*, 8th Annual Workshop on Duplicating, Deconstructing, and Debunking (WDDD 2009), held in conjunction with ISCA 2009), 12 pages, Austin, Texas, June 21, 2009.
- [W.6] George L. Yuan and **Tor M. Aamodt**, *An Analytical DRAM Performance Model*, Fifth Annual Workshop on Modeling, Benchmarking and Simulation (MoBS), held in conj. with 36th ACM/IEEE Int'l Symp. on Computer Architecture, (ISCA 2009), 10 pages, Austin, Texas, June 21, 2009.
- [W.5] Xi Chen and **Tor M. Aamodt**. *An Improved Analytical Superscalar Microprocessor Memory Model*. The Fourth Annual Workshop on Modeling, Benchmarking and Simulation [in conj. with 35th ACM/IEEE Int'l Symp. on Computer Architecture, (ISCA 2008)], pp. 7-16, Beijing, China, June 22, 2008.
- [W.4] Ali Bakhoda and **Tor M. Aamodt**, *Extending the Scalability of Single Chip Stream Processors with On-chip Caches*, 2nd Workshop on Chip Multiprocessor Memory Systems and Interconnects [in conj. with 35th ACM/IEEE Int'l Symp. on Computer Architecture, (ISCA 2008)], 9 pages, Beijing, China, June 22, 2008.
- [W.3] **Tor Aamodt**, Pedro Marcuello, Paul Chow, Per Hammarlund, Hong Wang, *Prescient Instruction Prefetch*, 6th Workshop on Multithreaded Execution, Architecture, and Compilation (MTEAC-6) [held in conjunction with the 35th ACM/IEEE International Symposium on Microarchitecture (MICRO-35)], Istanbul Turkey, pp. 3-10, November 2002. **(best student paper award)**
- [W.2] **Tor Aamodt**, Andreas Moshovos, and Paul Chow, *The Predictability of Computations that Produce Unpredictable Outcomes*, 5th Workshop on Multithreaded Execution, Architecture, and Compilation (MTEAC-5) [held in conjunction with the 34th ACM/IEEE International Symposium on Microarchitecture (MICRO-34)], pp. 23-34, Austin TX, December 2001.
- [W.1] **Tor Aamodt**, and Paul Chow, *Numerical Error Minimizing Floating-Point to Fixed-Point ANSI C Compilation*, 1st Workshop on Media Processors and Digital Signal Processing (MPDSP-1) [held in conjunction with the 32nd ACM/IEEE International Symposium on Microarchitecture (MICRO-32)], pp. 3-12, Haifa Israel, November 1999. **(refereed by extended abstract)**

2. NON-REFEREED PUBLICATIONS

- (a) *Journals*
 (b) *Conference Proceedings*

- [NR-C.1] **Tor M. Aamodt**, *Architecting Graphics Processors for Non-Graphics Compute Acceleration*, In proceedings of the 2009 IEEE Pacific Rim Conference on Communications, Computers and Signal Processing, Special Session on Computer Architecture (PACRIM-09), pp. 963-968, Victoria, BC,

August 23-26, 2009.

(lightly refereed, invited paper)**(c) Other**

- [T.11] Md Aamir Raihan, Negar Goli, Tor Aamodt, Modeling Deep Learning Accelerator Enabled GPUs, arXiv preprint arXiv:1811.08309, Nov 2018.
- [T.10] Jonathan Lew, Deval Shah, Suchita Pati, Shaylin Cattell, Mengchi Zhang, Amruth Sandhupatla, Christopher Ng, Negar Goli, Matthew D Sinclair, Timothy .G Rogers, Tor Aamodt, Analyzing Machine Learning Workloads Using a Detailed GPU Simulator, arXiv preprint arXiv:1811.08933, Nov 2018.
- [T.9] Mahmoud Khairy, Jain Akshay, Tor Aamodt, Timothy G. Rogers, Exploring Modern GPU Memory System Design Challenges through Accurate Modeling, arXiv preprint arXiv:1810.07269, Oct 2018.
- [T.8] Yatish Turakhia, Subhasis Das, Tor M. Aamodt, William J. Dally, HoLiSwap: Reducing Wire Energy in L1 Caches, arXiv preprint arXiv:1701.03878, January 2017.
- [T.7] Milad Mohammadi, Tor M Aamodt, William J. Dally, CG-OoO: Energy-Efficient Coarse-Grain Out-of-Order Execution, arXiv preprint arXiv:1606.01607, June 2016
- [T.6] Patrick Judd, Jorge Albericio, *Taylor Hetherington*, **Tor Aamodt**, Natalie Enright Jerger, Raquel Urtasun, Andreas Moshovos, "Reduced-Precision Strategies for Bounded Memory in Deep Neural Nets", arXiv preprint arXiv:1511.05236, 2015
- [T.5] *Owen Kirby*, Shahriar Mirabbasi, **Tor M. Aamodt**, *Mixed-Signal Neural Branch Prediction*, Technical Report, University of British Columbia, 8 June 2007.
- [T.4] **Tor M. Aamodt**, *Modeling and Optimization of Speculative Threads*, Doctoral Thesis, University of Toronto, 2006.
- [T.3] **Tor Aamodt**, Andreas Moshovos, and Paul Chow, *The Predictability of Computations that Produce Unpredictable Outcomes*, Technical Report #TR-01-08-01, EECG, University of Toronto, August 2001.
- [T.2] **Tor M. Aamodt**, *Floating-Point to Fixed-Point Compilation and Embedded Architectural Support*, Masters Thesis, University of Toronto, January 2001.
- [T.1] **Tor M. Aamodt**, *Intelligent Control via Reinforcement Learning: State Transfer and Stabilization of a Rotational Inverted Pendulum*, Bachelors Thesis, University of Toronto, April 1997.

3. BOOKS**(a) Authored**

Tor M. Aamodt, Wilson Wai Lun Fung, Timothy G. Rogers, *General-Purpose Graphics Processor Architectures*, Morgan & Claypool Publishers, Synthesis Lectures on Computer Architecture, 140 pages, May 2018. http://www.morganclaypoolpublishers.com/catalog_Orig/product_info.php?products_id=1245

William J. Dally, R. Curtis Harting, **Tor M. Aamodt**, *Digital Design Using VHDL: A Systems Approach*, Cambridge University Press, 721 pages, January 2016. <http://www.amazon.com/gp/product/1107098866>

(b) Edited**(c) Chapters****4. PATENTS**

[P.5] **Tor M. Aamodt**, Hong Wang, Per Hammarlund, John P. Shen, Steve Shih-wei Liao, Perry H. Wang, Method and apparatus for efficient resource utilization for prescient instruction prefetch, United States Patent #7,818,547, Issued October 19, 2010. Assignee: Intel Corporation.

[P.4] Hong Wang, **Tor M. Aamodt**, Pedro Marcuello, Jared W. Stark, John P. Shen, Antonio Gonzalez, Per Hammarlund, Gerolf F. Hoflehner, Perry H. Wang, Steve Shih-wei Liao, Speculative multi-threading for instruction prefetch and/or trace pre-build, United States Patent #7,814,469, Issued October 12, 2010. Assignee: Intel Corporation.

[P.3] Hong Wang, **Tor Aamodt**, Per Hammarlund, John Shen, Xinmin Tian, Milind Girkar, Perry Wang, Steve Shih-wei Liao, *Safe Store for Speculative Helper Threads*, United States Patent #7,657,880, Issued February 2, 2010. Assignee: Intel Corporation.

[P.2] **Tor M. Aamodt**, Hong Wang, John P. Shen, Per Hammarlund, *Methods and Apparatus for Generating Speculative Helper Thread Spawn-Target Points*, United States Patent #7,523,465, Issued April 21, 2009. Assignee: Intel Corporation.

[P.1] **Tor M. Aamodt**, Hong Wang, Per Hammarlund, John P. Shen, Steve Shih-wei Liao, Perry H. Wang, *Method and Apparatus for Efficient Utilization for Prescient Instruction Prefetch*, United States Patent #7,404,067, Issued July 22, 2008. Assignee: Intel Corporation.

5. **SPECIAL COPYRIGHTS**

6. **ARTISTIC WORKS, PERFORMANCES, DESIGNS**

7. **OTHER WORKS**

GPGPU-Sim Simulator (<http://www.gpgpu-sim.org/>).

8. **WORK SUBMITTED (including publisher and date of submission)**

Several papers under submission.

9. **WORK IN PROGRESS (including degree of completion)**

Many.