Tushar **Krishna**

266 Ferst Drive • School of Electrical and Computer Engineering • Atlanta, GA 30332, USA
PHONE (+1) 206 601 6213 • **E-MAIL** tushar@ece.gatech.edu • **WEB** http://tusharkrishna.ece.gatech.edu

RESEARCH INTERESTS

Computer Architecture: Domain-specific accelerators (AI/ML, Sparse Tensors, Graphs)
Computer Systems: Networking, Disaggr Memory, Schedulers for AI/ML and HPC
Interconnection Networks: Networks-on-Chip, Package/Wafer-scale, HPC, Datacenters
Simulation Infrastructures: Multi-fidelity simulators and models, Design-Space Exploration

EDUCATION

Feb 2014 **Massachusetts Institute of Technology**
Ph.D. in Electrical Engineering and Computer Science
- *Advisor:* Li-Shiuan Peh
- *Committee:* Srinivas Devadas and Joel Emer
- *Thesis:* "Enabling Dedicated Single-Cycle Connections Over A Shared Network-on-Chip"

Sep 2009 **Princeton University**
M.S.E. in Electrical Engineering
- *Advisor:* Li-Shiuan Peh
- *Thesis:* "Networks-on-Chip with Hybrid Interconnects"

Aug 2007 **Indian Institute of Technology (IIT), Delhi**
B.Tech. (Honors) in Electrical Engineering

PROFESSIONAL APPOINTMENTS

Aug'23 - present **Massachusetts Institute of Technology,** Cambridge, MA, USA
Visiting Associate Professor
Department of Electrical Engineering and Computer Science and CSAIL

Jan'22 - present **Georgia Institute of Technology,** Atlanta, GA, USA
Associate Director, Center for Research in Novel Computing Hierarchies (CRNCH)
Interdisciplinary Research Center focused on post-Moore computing

Aug'21 – present **Georgia Institute of Technology,** Atlanta, GA, USA
Associate Professor *(with tenure)*
School of Electrical and Computer Engineering.
School of Computer Science (By Courtesy).

Sept'19 – Jul'21 **Georgia Institute of Technology,** Atlanta, GA, USA
ON Semiconductor Professor (Endowed Junior Professorship)
School of Electrical and Computer Engineering

Aug '15 – Jul'21 **Georgia Institute of Technology,** Atlanta, GA, USA
Assistant Professor
School of Electrical and Computer Engineering.
School of Computer Science (By Courtesy).

Jan '15 – Jul '15 **Massachusetts Institute of Technology, SMART Center,** Cambridge, MA, USA
Post-doctoral Researcher.

Nov '13 – Jan '15 **Intel Corporation, VSSAD Group**, Hudson, MA
Research Engineer. Manager: Joel Emer

Jun'10 – Aug '10 **AMD (Advanced Micro Devices) Research**, Bellevue, WA, USA
Co-Op Engineer. Mentors: Bradford Beckmann and Steve Reinhardt

Jun'09 – Aug '09 **AMD Research**, Bellevue, WA, USA
Co-Op Engineer. Mentors: Bradford Beckmann and Steve Reinhardt

Jun'08 – Aug '08 **AMD,** North Bridge Architecture Group, Sunnyvale, CA, USA
Co-op Engineer. Mentor: Pat Conway

May'06 – Jul '06 **NVIDIA**, Digital Hardware Design Group, Bangalore, India
Summer Intern.

HONORS AND AWARDS

International or National Awards

2023 **Best Paper Nominee** at the 17th International Symposium on Networks-on-Chip (**NOCS**)
Paper ranked among top two of presented papers that year.

2022 **Best Paper Award** at 23rd ACM/IFIP International Middleware Conference
Paper ranked top among all accepted papers that year

2022 **HPCA Hall of Fame Award**
Awarded for 8+ papers in IEEE International Symposium on High-Performance Computer Architecture (HPCA)

2022 **Best Paper Nominee** at 36th IEEE International Parallel & Distributed Processing Symposium (**IPDPS**)
Paper ranked among top 5 of all accepted papers that year.

2021 **Qualcomm Faculty Award**
The award supports key professors and their research (without requiring any formal application), with the goal of strengthening Qualcomm's engagement with faculty who also play a key role in Qualcomm's recruiting of top graduate students.

2021 **Honorable Mention at ACM SIGARCH / CS TCCA Outstanding Dissertation Award**
(for Advisee: Hyoukjun Kwon)

2021 **Best Paper Nominee** at International Conference for High Performance Computing, Networking, Storage, and Analysis (**SC**)
Paper ranked among top 4 of 86 accepted papers out of 365 submissions that year

2021 **Best Paper Nominee** at the 15th International Symposium on Networks-on-Chip (**NOCS**)
Paper ranked among top two of presented papers that year.

2021 **Best Paper Nominee** at 26th Asia and South Pacific Design Automation Conference (**ASP-DAC**)
Paper ranked among top 4 of 111 accepted papers out of 327 submissions that year..

2020 **Best Paper Award** at 28th IFIP/IEEE Int. Conference on Very Large Scale Integration (**VLSI-SoC**)
Paper ranked highest among 30 accepted papers out of 78 submissions that year.

2020 **Facebook Faculty Research Award**
9 winners out of 132 worldwide submissions

2020 **Best Paper Award** at 26th IEEE International Symposium on High-Performance Computer Architecture (**HPCA**)
Paper ranked highest among 48 papers accepted out of 248 submissions that year.

2020 **IEEE Micro's Top Picks** from Computer Architecture[1]

2019 **Facebook Faculty Research Award**
8 winners out of 88 worldwide submissions

[1] *Every year, IEEE Micro publishes a special issue called Top Picks from Computer Architecture recognizing the year's top papers (12 as Top Picks and 10-12 more as Honorable Mentions) that have potential for long-term impact. In order for a paper to be considered, it must first have been accepted in a major computer architecture conference the previous year.*

2019 **Google Faculty Research Award**
158 winners out of 910 worldwide submissions
2019 **IEEE Micro's Top Picks** from Computer Architecture
2019 **Honorable Mention in IEEE Micro Top Picks** from Computer Architecture
2019 **Best Paper Nominee** at the IEEE International Symposium on Workload Characterization (**IISWC**)
Paper ranked among top 4 of 27 accepted papers that year.
2018 **NSF CISE Research Initiation Initiative (CRII) Award**
2017 **Best Paper Award** (Design Methods & Tools Track) at Design Automation and Test in Europe (**DATE**)
DATE has 4 tracks (D, A, T and E). Paper was ranked highest in the D track, making it one of the top 4 among 194 accepted papers out of 794 submissions that year.
2014 **Best Paper Award** at the 8th International Symposium on Networks-on-Chip (**NOCS**)
Paper ranked highest among 21 papers accepted out of 83 submissions that year
2014 **IEEE Micro Top Picks** from Computer Architecture
2009 **IEEE Micro Top Picks** from Hot Interconnects Symposium
Special Issue featuring top papers from Hot Interconnects that year
2003 **Gold Medal** at the Indian National Chemistry Olympiad
25 gold medalists out of ~75 invited for the Olympiad after a nationwide exam.

Institute or School Awards

2023 **Richard M. Bass/Eta Kappa Nu Outstanding Junior Teacher Award**
Based on vote by graduating undergraduate class in School of ECE at Georgia Tech
2021 **Roger P. Webb Award for Outstanding Junior Faculty Member** by School of ECE, Georgia Tech
2015-2021 **Certificate of Teaching Excellence** from Georgia Tech's Center of Teaching and Learning
Received 7 times upon recommendation from students in class
2018 **Class of 1940 Course Survey Teaching Effectiveness Award**
Received a teaching evaluation of 5.0 / 5.0
2007-08 **Princeton Graduate Fellowship**
2007 ICIM Stay Ahead Award for the **Best Undergraduate Project in Computer Technology**, IIT Delhi
2004-2006 "National Initiative for Undergraduate Sciences" (NIUS) Fellowship, Homi Bhabha Centre for Science Education (HBCSE), India
2003, 2004 Merit prize for academic excellence, IIT Delhi

RESEARCH FUNDING

My cumulative share: ~$7.3 Million. Current active awards: ~$4.3 Million

NSF (OAC) Collaborative Research: Frameworks: Advancing Computer Hardware and Systems' Research Capability, Reproducibility, and Sustainability with the gem5 Simulator Ecosystem
PI: Jason Lowe Power (UC Davis). Co-PIs: Tushar Krishna (GT), Matthew Sinclair (U Wisc), Chris Batten (Cornell), Mohammad Alian (U Kansas), Elba Garza (U Washington)
Total = $3.5M, My Share = $200,000
Year: 2023-2027

SRC + DARPA (JUMP2.0) "CoCoSys: Center for the Co-Design of Cognitive Systems"
One of seven JUMP2.0 centers from SRC + DARPA
https://www.src.org/newsroom/press-release/2023/1037/
Center Director: Arijit Raychowdhury (GT), Associate Director: Anand Raghunathan (Purdue)
PIs: Kevin Cao (ASU), Michael Carbin (MIT), James DiCarlo (MIT), Anca Dragan (Berkeley), Sumeet Gupta (Purdue), Larry Heck (GT), Tushar Krishna (GT), Yingyan Lin (GT), Azad Naeemi (GT), Bruno Olshausen (Berkeley), Priya Panda (Yale), Jan Rabaey (Berkeley), Vijay Raghunathan (Purdue), Priyanka Raina (Stanford), Tajana Rosing (UCSD), Kaushik Roy (Purdue), Jae-sun Seo (ASU), Naresh Shanbhag (UIUC), Josh Tannenbaum (MIT)
Total: $35M. My-share: $1,250,000.
Year: 2023-2028

SRC + DARPA (JUMP2.0) "ACE Center for Evolvable Computing"
One of seven JUMP2.0 centers from SRC + DARPA
https://www.src.org/newsroom/press-release/2023/1037/
Center Director: Josep Torrelas (UIUC), Associate Director: Minlan Yu (Harvard)
PIs: Tarek Abdelzaher (UIUC), Mohammad Alian (Univ of Kansas), Adam Belay (MIT), Manya Ghobadi (MIT), Rajesh Gupta (UCSD), Tushar Krishna (GT), Arvind Krishnamurthy (UW), Christos Kozyrakis (Stanford), Jose Martinez (Cornell), Charath Mendis (UIUC), Subhashish Mitra (Stanford), Muhammad Shahbaz (Purdue), Edward Suh (Cornell), Steve Swanson (UCSD), Michael Taylor (UW), Radu Teodorescu (OSU), Mohit Tiwari (UT), Zhengya Zhang (UMich), Zhiru Zhang (Cornell).
Total: $39.6M. My-share: $1,250,000.
Year: 2023-2028

SRC (AI HW) "Enabling Simulation and Design of Distributed Heterogeneous AI Accelerator Platforms"
PI: Tushar Krishna
Total = My Share = $270,000
Year: 2023-2025

IARPA (AGILE) "TIGRE -Transactional Integrated Global-memory system with dynamic Routing and End-to-end flow control"
One of six teams selected in IARPA's AGILE Program
https://www.dni.gov/index.php/newsroom/press-releases/press-releases-2022/item/2332-iarpa-and-u-s-army-launch-research-to-engineer-next-generation-of-computers
PIs: Joshua Fryman, Fabrizio Petrini (Intel)
GT PI: Tushar Krishna
My share = $750,000
Year: 2023 - 2028

TSMC A Multi-tier Fine-Grained 3D Architecture for Efficient Inference and Training of Memory-intensive AI Workloads
PIs: Tushar Krishna and Sungkyu Lim (GT)
Total: $300,000. My share = $150,000
Year: 2023 - 2025.

Intel "TRIM: Architecting Systems for Terascale in-Memory Applications"
PI: Moin Qureshi (GT), Co-PIs: Tushar Krishna, Hyesoon Kim, Alex Daglis, Ada Gavrilovska (GT)
Total: $1,000,000, my share: $200,000
Year: 2022 – 2027.

IARPA (MicroE4AI) "Robust Autonomy in UAVs on a Convergent Digital-Analog Ferroelectronics Platform"
PI: Arijit Raychowdhury (GT), Co-PIs: Tushar Krishna, Hyesoon Kim, Vivek Sarkar, Shimeng Yu, Asif Khan (GT), Vijay Reddi (Harvard), Kostas Daniilidis,Vijay Kumar, CJ Taylor (UPenn), Jayakanth Ravichandran (USC), Pradip Bose, Nandhini Chandramoorthy, Karthik Swaminathan (IBM)
Total: $2,000,000, my share: $300,000
Year: 2021-22.

Qualcomm Unrestricted Gift
PI: Tushar Krishna
Total = my share = $75,000.
Year: 2021 – * .

Intel "Network-on-Chip and Memory Systems for Multi-accelerator Systems"

PI: Tushar Krishna
Total = my share = $70,000.
Year: 2021 - *.

Intel "Scalable System Architecture Design for Zeta Scale Computing"
PI: Tushar Krishna
Total = my share = $70,000.
Year: 2021 - *.

Facebook "HW/SW Co-Design of Next-Generation Training Platforms for Recommendation Models"
PI: Tushar Krishna
Total = my share = $100,000.
Year: 2021 - *.

NSF (CNS) "Edge Platform for Enabling Situation Awareness Applications"
PI: Kishore Ramachandran, Co-PI: Tushar Krishna
Total = $500,000, my share = $250,000.
Year: 2020-2023.

NSF (CCF) "PPoSS: Planning: Integrated Scalable Platform for Privacy-aware Collaborative Learning and Inference"
PI: Vivek Sarkar (GT), Co-PIs: Tushar Krishna, Alexey Tumanov (GT), Dawn Song (Berkeley), Jimeng Sun (UIUC)
Total: $250,000, my share = $50,000.
Year: 2020-2021.

Facebook "HW/SW Co-Design of Next-Generation Training Platforms for DLRMs"
PI: Tushar Krishna
Total = my share = $50,000.
Year: 2020 - *.

Intel "System architecture design for large scale AI models"
PI: Tushar Krishna
Total = my share = $75,000.
Year: 2020 - *.

Intel "Tensor Extensions for Accelerating Sparse ML Inference"
PI: Tushar Krishna.
Total = my share = $60,000.
Year: 2020 - *.

DARPA "ARION: A Real-time HPC Platform with In-memory Non-Von-Neumann Processing and Optical Networking"
PI: Saibal Mukhopadhyay (GT), Co-PIs: Tushar Krishna, Madhavan Swaminathan, Santosh Pande, Justin Romberg (GT), Keren Bergman, Luca Carloni, Mingu Seok (Columbia)
Total = $13.4 Million, my share = $500,000.
Year: 2020-2023.

Department of Energy (DoE) "ARtificial Intelligence focused ARchitectures and Algorithms (ARIAA)"
[Department of Energy Center with Pacific Northwest National Lab (PNNL), Sandia National Lab (SNL) and Georgia Tech].
https://www.news.gatech.edu/2019/11/05/national-labs-georgia-tech-collaborate-ai-research
PIs: Roberto Gioiosa (PNNL), Siva Rajamanickam (Sandia National Labs), Tushar Krishna (GT)
Total = 5.5M, my share = $510,000.

Year: 2019-2022.

Intel "Tensor Extensions for Accelerating Sparse ML Inference"
PI: Tushar Krishna.
Total = my share = $60,000.
Year: 2019 - *.

NSF (OAC) "Enabling Rapid Design and Deployment of Deep Learning Models on Hardware Accelerators".
PI: Tushar Krishna.
Total = my share = $500,000.
Year: 2019-2022.

Facebook "ML-Driven HW-SW Co-Design of Efficient Tensor Core Architectures".
PI: Tushar Krishna.
Total = my share = $50,000.
Year: 2019 - *

Google "Using ML to Design Efficient ML Accelerators".
PI: Tushar Krishna.
Total = $66,000.
Year: 2019 - *.

Intel "Design of Scalable Deep Learning Training Accelerators for Sparse and Irregular GEMMS".
PI: Tushar Krishna
Total = $75,000.
Year: 2018 - *.

NSF CRII (SHF) "Enabling Neuroevolution in Hardware".
PI: Tushar Krishna.
Total = $175,000.
Year: 2018-2019.

DARPA "A Vertically-Integrated Design Flow for IP Reuse and Heterogeneous Integration".
PI: Sung-Kyu Lim, Co-PIs: Tushar Krishna, Saibal Mukhopadhyay, Madhavan Swaminathan.
Total = $3,719,691. My share = $400,000.
Year: 2017-2019.

PUBLICATIONS: BOOKS

"Data Orchestration in Deep Learning Accelerators"
Tushar Krishna, Michael Pellauer, Angshuman Parashar, Hyoukjun Kwon, and Ananda Samajdar
Synthesis Lectures on Computer Architecture. Morgan & Claypool Publishers. August 2020

"On-Chip Networks", Second Edition
Natalie Enright Jerger, **Tushar Krishna**, and Li-Shiuan Peh.
Synthesis Lectures on Computer Architecture. Morgan & Claypool Publishers. June 2017

PUBLICATIONS: BOOK CHAPTERS

Interconnect Modeling for Homogeneous and Heterogeneous Multiprocessors
Srikant Bharadwaj and Tushar Krishna
Network-on-Chip Security and Privacy, Springer, May 2021

"Application Benchmarking" chapter
Geoffrey Burr, Tom Conte, Paolo Gargini, Vladimir Getov, Yoshihiro Hayashi, Takeshi Iwashita, Vijay Janapa Reddi, Masaaki Kondo, **Tushar Krishna**, Peter M. Kogge, Anoop Nair, Dam Sunwoo,

Josep Torrelas, Peter Torelli
IEEE International Roadmap for Devices and Systems, 2020 Edition.

PUBLICATIONS: REFEREED CONFERENCES

Total number of citations: 14710, h-index: 37

NOCS 2023 ***Best paper nominee***
"SPOCK: Reverse Packet Traversal for Deadlock Recovery"
Zeyu Chen, Ankur Bindal, Vaidehi Garg and **Tushar Krishna**
In Proc. of the 17th IEEE/ACM International Symposium on Networks-on-Chip, Sep 2023

MLSys 2023
"XRBench: An Extended Reality (XR) Machine Learning Benchmark Suite for the Metaverse", Hyoukjun Kwon, Krishnakumar Nair, Jamin Seo, Jason Yik, Debabrata Mohapatra, Dongyuan Zhan, Jinook Song, Peter Capak, Peizhao Zhang, Peter Vajda, Colby Banbury, Mark Mazumder, Liangzhen Lai, Ashish Sirasao, **Tushar Krishna**, Harshit Khaitan, Vikas Chandra, Vijay Janapa Reddi
In Proc of Sixth Conference on Machine Learning and Systems, Sept 2023

MLSys 2023
"Subgraph-Stationary Hardware-Software Inference Co-Design"
Payman Behnam, Jianming Tong, Alind Khare, Yangyu Chen, Yue Pan, Pranav Gadikar, Abhimanyu Rajeshkumar Bambhaniya, **Tushar Krishna**, Alexey Tumanov
In Proc of Sixth Conference on Machine Learning and Systems, Sept 2023

IMS 2023
"FPGA-Based High-Performance Real-Time Emulation of Radar System using Direct Path Compute Model"
X. Mao, M. Mukherjee, N. M. Rahman, U. Kamal, S. Sharma, P. Behnam, J. Tong, J. Woo, C. DeLude, J. Driscoll, J. Seo, S. Pande, **T. Krishna**, J. Romberg, M. Swaminathan, and S. Mukhopadhyay
In Proc of IEEE MTT-S International Microwave Symposium, Jun 2023

RadarConf'2023
"A High Performance Computing Architecture for Real-Time Digital Emulation of RF Interactions"
M. Mukherjee, N. M. Rahman, C. DeLude, J. Driscoll, U. Kamal, J. Woo, J. Seo, S. Sharma, X. Mao, P. Behnam, S. Khan, D. Kim, J. Tong, P. Sinha, S. Pande, **T. Krishna**, J. Romberg, M. Swaminathan, and S. Mukhopadhyay.
In Prof of IEEE RadarConf 2023

DATE 2023
"Proteus: HLS based NoC Generator and Simulator"
Abhimanyu Rajeshkumar Bambhaniya, Yangyu Chen, Anshuman, Rohan Banerjee and **Tushar Krishna**
In Proc of Design Automation and Test in Europe, Apr 2023

DATE 2023
"AIRCHITECT: Automating Hardware Architecture and Mapping Optimization"
Ananda Samajdar, Jan Moritz Joseph and **Tushar Krishna**
In Proc of Design Automation and Test in Europe, Apr 2023

ASPLOS 2023
"Flexagon: A Multi-Dataflow Sparse-Sparse Matrix Multiplication Accelerator for Efficient DNN Processing"
Francisco Muñoz-Martinez, Raveesh Garg, Michael Pellauer, Jose L. Abellan, Manuel E. Acacio

and **Tushar Krishna**
In Proc of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Mar 2023

ASPLOS 2023 FLAT: An Optimized Dataflow for Mitigating Attention Bottlenecks
Sheng-Chun Kao, Suvinay Subramanian, Gaurav Agrawal, Amir Yazdanbakhsh and **Tushar Krishna**
In Proc of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Mar 2023

HPCA 2023 "VEGETA: Vertically-Integrated Extensions for Sparse/Dense GEMM Tile Acceleration on CPUs"
Geonhwa Jeong, Sana Damani, Abhimanyu Bambhaniya, Eric Qin, Christopher J. Hughes, Sreenivas Subramoney, Hyesoon Kim and **Tushar Krishna**
In Proc of the 29th IEEE International Symposium on High-Performance Computer Architecture, Feb 2023
Acceptance Rate: 25% (91/364)

AAAI 2023 "Efficient Distributed Inference of Deep Neural Networks via Restructuring and Pruning", Afshin Abdi, Saeed Rashidi, **Tushar Krishna**, Faramarz Fekri.
In Proc of the 37'th Association for the Advancement of Artificial Intelligence, Washington, DC, USA, Feb, 2023

IISWC 2022 "Demystifying Map Space Exploration for NPUs"
Sheng-Chun Kao, Angshuman Parashar, Po-An Tsai and **Tushar Krishna**
In Proc of 2022 IEEE International Symposium on Workload Characterization, Nov 2022

Middleware 2022
Best Paper Award "MicroEdge: A Low-Cost Edge Cluster System Architecture for Scalable Camera Processing"
Difei Cao*, Jinsun Yoo*, Zhuangdi Xu, Enrique Saurez, Harshit Gupta, **Tushar Krishna** and Umakishore Ramachandran
In 23rd ACM/IFIP International Middleware Conference, Nov 2022
**equal contribution*

HotI 2022 "Impact of RoCE Congestion Control Policies on Distributed Training of DNNs"
Tarannum Khan, Saeed Rashidi, Pallavi Shurpali, Aditya Akella, **Tushar Krishna** and Srinivas Sridharan
In Proceedings of the 29th International Symposium on High-Performance Interconnects, Aug 2022

DAC 2022 "Self Adaptive Reconfigurable Arrays (SARA): Learning Flexible GEMM Accelerator Configuration and Mapping-space using ML"
Ananda Samajdar, Eric Qin, Michael Pellauer and **Tushar Krishna**
In Proc. of the 59th Annual Design Automation Conference, Jul 2022
Acceptance Rate: 23%

ISCA 2022 "Themis: A Network Bandwidth-Aware Collective Scheduling Policy for Distributed Training of DL Models"
Saeed Rashidi, William Won, Sudarshan Srinivasan, Srinivas Sridharan, and **Tushar Krishna**
In Proc of 49th International Symposium on Computer Architecture, Jun 2022
Acceptance Rate: 16.7% (67/400)

SIGMETRICS 2022 "A Formalism of DNN Accelerator Flexibility"
Sheng-Chun Kao, Hyoukjun Kwon, Michael Pellauer, Angshuman Parashar and **Tushar Krishna**
In Proc of 2022 ACM SIGMETRICS/Performance conference, Jun 2022
Acceptance Rate: 10.6% (13/122)

IPDPS 2022
Best paper nominee "Understanding the Design-Space of Sparse/Dense Multiphase GNN dataflows on Spatial Accelerators"
Raveesh Garg, Eric Qin, Francisco Muñoz-Martínez, Robert Guirado, Akshay Jain, Sergi Abadal, José L. Abellán, Manuel E. Acacio, Eduard Alarcón, Sivasankaran Rajamanickam, and **Tushar Krishna**
In Proc of the 36th IEEE International Parallel & Distributed Processing Symposium, May 2022

DATE 2022 "DiGamma: A Domain-aware Genetic Algorithm for HW-Mapping Co-optimization for DNN Accelerators"
Sheng-Chun Kao, Michael Pellauer, Angshuman Parashar and **Tushar Krishna**
In Proc of Design Automation and Test in Europe, Mar 2022

HPCA 2022 "MAGMA: An Optimization Framework for Mapping Multiple DNNs on Multiple Accelerator Cores"
Sheng-Chun Kao and **Tushar Krishna**
In Proc of the 28th IEEE International Symposium on High-Performance Computer Architecture, Feb 2022
Acceptance Rate: 30% (80/262)

HPCA 2022 "Stay in your Lane: A NoC with Low-overhead Multi-packet Bypassing"
Hossein Farrokhbakht, Paul Gratz, **Tushar Krishna**, Joshua San Miguel and Natalie Enright Jerger
In Proc of the 28th IEEE International Symposium on High-Performance Computer Architecture, Feb 2022
Acceptance Rate: 30% (80/262)

DAC 2021 "RASA: Efficient Register-Aware Systolic Array Matrix Engine for CPU"
Geonhwa Jeong, Eric Qin, Ananda Samajdar, Christopher Hughes, Sreenivas Subramoney, Hyesoon Kim and **Tushar Krishna**
In Proc. of the 58th Annual Design Automation Conference, Dec 2021
Acceptance Rate: 23%

IISWC 2021 "STONNE: Enabling Cycle-Level Microarchitectural Simulation for DNN Inference Accelerators"
Francisco Muñoz-Martinez, Jose L. Abellan, Manuel E. Acacio and **Tushar Krishna**
In Proc. of the IEEE International Symposium on Workload Characterization, Nov 2021

NOCS 2021 "A Novel Network Fabric for Efficient Spatio-Temporal Reduction in Flexible DNN Accelerators"
Best paper nominee
Francisco Muñoz-Martinez, Jose L. Abellan, Manuel E. Acacio and **Tushar Krishna**
In Proc. of the 15th IEEE/ACM International Symposium on Networks-on-Chip), Oct 2021

NOCS 2021 "DUB: Dynamic Underclocking and Bypassing in Network-on-Chip for Heterogeneous GPU Workloads"
Srikant Bharadwaj, Shomit Das, Yasuko Eckert, Mark Oskin and **Tushar Krishna**
Short Paper — In Proc. of the 15th IEEE/ACM International Symposium on Networks-on-Chip, Oct 2021

SC 2021 "SEEC: Stochastic Escape Express Channel"
Best paper nominee
Mayank Parasar, Natalie Enright Jerger, Paul Gratz, Joshua San Miguel and **Tushar Krishna**
In Proc. of International Conference for High Performance Computing, Networking, Storage, and Analysis, Nov 2021
Acceptance Rate: 23% (86/365)

PACT 2021 "Union: A Unified HW-SW Co-Design Ecosystem in MLIR for Evaluating Tensor Operators on Spatial Accelerators"
Geonhwa Jeong, Gokcen Kestor, Prasanth Chatarasi, Angshuman Parashar, Po-An Tsai, Siva Rajamanickam, Roberto Gioiosa and **Tushar Krishna**
In Proc. of 30th International Conference on Parallel Architectures and Compilation Techniques, Sep 2021
Acceptance Rate: 26% (25/96)

IMS 2021 "A Configurable Architecture for Efficient Sparse FIR Computation in Real-time Radio Frequency Systems"
Jamin Seo, Mandovi Mukherjee, Nael Mizanur Rahman*, Jianming Tong, Coleman DeLude, **Tushar Krishna**, Justin Romberg, Saibal Mukhopadhyay,
In Proc of IEEE MTT-S International Microwave Symposium, Jun 2021

ISCA 2021 "Enabling Compute-Communication Overlap in Distributed Training Platforms"
Saeed Rashidi, Srinivas Sridharan, Sudarshan Srinivasan, Matthew Denton, Amoghavarsha Suresh, Jade Nie and **Tushar Krishna**
In Proc of 48th International Symposium on Computer Architecture, Jun 2021
Acceptance Rate: 18.7% (76/406)

IPDPS 2021 "Extending Sparse Tensor Accelerators to Support Multiple Compression Formats"
Eric Qin, Geonhwa Jeong, William Won, Sheng-Chun Kao, Hyoukjun Kwon, Sudarshan Srinivasan, Dipankar Das, Gordon E. Moon, Sivasankaran Rajamanickam and **Tushar Krishna**
In Proc of the 35th IEEE International Parallel & Distributed Processing Symposium, May 2021

ISQED 2021 "Architecture, Dataflow and Physical Design Implications of 3D-ICs for DNN-Accelerators"
Jan Moritz Joseph, Ananda Samajdar, Lingjun Zhu, Rainer Leupers, Sung Kyu Lim, Thilo Pionteck and **Tushar Krishna**
In Proc of the 22nd International Symposium on Quality Electronic Design, Apr 2021

ISPASS 2021 "E3: A HW/SW Co-design Neuroevolution Platform for Autonomous Learning in Edge Device"
Sheng-Chun Kao and **Tushar Krishna**
In Proc of the IEEE International Symposium on Performance Analysis of Systems and Software, Mar 2021

HPCA 2021 "Pitstop: Enabling a Virtual Network Free Network-on-Chip"
Hossein Farrokhbakht, Henry Kao, Kamran Hasan, Paul Gratz, **Tushar Krishna**, Joshua San Miguel and Natalie Enright Jerger
In Proc of the 27th IEEE International Symposium on High-Performance Computer Architecture, Feb 2021
Acceptance Rate: 24% (63/258)

HPCA 2021 "Heterogeneous Dataflow Accelerators for Multi-DNN Workloads"
Hyoukjun Kwon, Liangzhen Lai, Michael Pellauer, Yu-Hsin Chen, **Tushar Krishna**, Vikas Chandra
In Proc of the 27th IEEE International Symposium on High-Performance Computer Architecture, Feb 2021
Acceptance Rate: 24% (63/258)

ASP-DAC 2021 "Dataflow-Architecture Co-Design for 2.5D DNN Accelerators using Wireless Network-on-Package"
Robert Guirado, Hyoukjun Kwon, Sergi Abadal, Eduard Alarcon and **Tushar Krishna**
In Proc of the 26th Asia and South Pacific Design Automation Conference, Jan 2021

ASP-DAC 2021 "Bridging the Frequency Gap in Heterogeneous 3D SoCs through Technology-Specific NoC Router Architecture"
Best paper nominee
Jan Moritz Joseph, Lennart Bamberg, Geonhwa Jeong, Ruei-Ting Chien, Rainer Leupers, Alberto Garcia-Ortiz, **Tushar Krishna** and Thilo Pionteck
In Proc of the 26th Asia and South Pacific Design Automation Conference, Jan 2021
Acceptance Rate: 34% (111/327)

ICCAD 2020 "GAMMA: Automating the HW Mapping of DNN Models on Accelerators via Genetic Algorithm"
Sheng-Chun Kao and **Tushar Krishna**
In Proc of the IEEE/ACM International Conference on Computer-Aided Design, Nov 2020
Acceptance Rate: 27% (127/470)

MICRO 2020 "ConfuciuX: Autonomous Hardware Resource Assignment for DNN Accelerators using Reinforcement Learning"
Sheng-Chun Kao, Geonhwa Jeong and **Tushar Krishna**
In Proc of 53rd Annual IEEE/ACM International Symposium on Microarchitecture, Oct 2020
Acceptance Rate: 19% (82/424)

VLSI-SoC 2020 "Breaking Barriers: Maximizing Array Utilization for Compute In-Memory Fabrics"
Best Paper Award
Brian Crafton, Samuel Spetalnick, Gauthaman Murali, **Tushar Krishna**, Sung-Kyu Lim, and Arijit Raychowdhury
In Proc of the 28th IFIP/IEEE Int.l Conference on Very Large Scale Integration, Oct 2020

HotI 2020 "Scalable Distributed Training of Recommendation Models: An ASTRA-SIM + NS3 case-study with TCP/IP transport"
Saeed Rashidi, Pallavi Shurpali, Srinivas Sridharan, Naader Hassani, Dheevatsa Mudigere, Krishnakumar Nair, Misha Smelyanski, and **Tushar Krishna**
In Proc of the 27th International Symposium on High-Performance Interconnects, Aug 2020

ISPASS 2020 "ASTRA-SIM: Enabling SW/HW Co-Design Exploration for Distributed DL Training Platforms"
Saeed Rashidi, Srinivas Sridharan, Sudarshan Srinivasan, and **Tushar Krishna**
In Proc of the IEEE International Symposium on Performance Analysis of Systems and Software, Aug 2020

ISPASS 2020 "A Systematic Methodology for Characterizing Scalability of DNN Accelerators using SCALE-Sim"
Ananda Samajdar, Jan Moritz Joseph, Yuhao Zhu, Paul Whatmough, Matthew Mattina, and **Tushar Krishna**
In Proc of the IEEE International Symposium on Performance Analysis of Systems and Software, Aug 2020

ISPASS 2020 "CLAN: Continuous Learning using Asynchronous NeuroEvolution on Commodity Edge Devices"
Parth Mannan, Ananda Samajdar, and **Tushar Krishna**
In Proc of the IEEE International Symposium on Performance Analysis of Systems and Software, Aug 2020

DAC 2020 "Kite: A Family of Heterogeneous Interposer Topologies Enabled via Accurate Interconnect Modeling"
Srikant Bharadwaj, Jieming Yin, Bradford M. Beckmann, and **Tushar Krishna**
In Proc. of the 57th Annual Design Automation Conference, Jul 2020
Acceptance Rate: 23% (228/991)

DAC 2020 "Co-Exploration of Neural Architectures and Heterogeneous ASIC Accelerator Designs Targeting Multiple Tasks"
Lei Yang, Zheyu Yan, Meng Li, Hyoukjun Kwon, Liangzhen Lai, **Tushar Krishna**, Vikas Chandra, Weiwen Jiang, Yiyu Shi
In Proc. of the 57th Annual Design Automation Conference, Jul 2020
Acceptance Rate: 23% (228/991)

HPCA 2020
Best Paper Award "SIGMA: A Sparse and Irregular GEMM Accelerator with Flexible Interconnects for DNN Training"
Eric Qin, Ananda Samajdar, Hyoukjun Kwon, Vineet Nadella, Sudarshan Srinivasan, Dipankar Das, Bharat Kaul, and **Tushar Krishna**
In Proc of the 26th IEEE International Symposium on High-Performance Computer Architecture, Feb 2019
Acceptance Rate: 19% (48/248)

HPCA 2020 "DRAIN: Deadlock Removal for Arbitrary Irregular Networks"
Mayank Parasar, Hossein Farrokhbakht, Natalie Enright Jerger, Paul Gratz, **Tushar Krishna**, and Joshua San Miguel

In Proc of the 26th IEEE International Symposium on High-Performance Computer Architecture, Feb 2019
Acceptance Rate: 19% (48/248)

HPCA 2020 "ALRESCHA: A Lightweight Reconfigurable Sparse-Computation Accelerator"
Bahar Asgari, Ramyad Hadidi, **Tushar Krishna**, Hyesoon Kim and Sudhakar Yalamanchili
In Proc of the 26th IEEE International Symposium on High-Performance Computer Architecture, Feb 2019
Acceptance Rate: 19% (48/248)

ICECS 2019 "Understanding the Impact of On-chip Communication on DNN Accelerator Performance"
Robert Guirado, Hyoukjun Kwon, Eduard Alarcon, Sergi Abadal and **Tushar Krishna**
In Proc of 26th IEEE International Conference on Electronics, Circuits and Systems, Nov 2019

IISWC 2019
Best paper nominee
"Characterizing the Deployment of Deep Neural Networks on Commercial Edge Devices"
Ramyad Hadidi, Jiashen Cao, Yilun Xie, Bahar Asgari, **Tushar Krishna**, and Hyesoon Kim
In Proc of IEEE International Symposium on Workload Characterization, Nov 2019

MICRO 2019
Selected for IEEE Micro Top Picks + Finalist at ACM Student Research Competition at MICRO 2018
"Understanding Reuse, Performance, and Hardware Cost of DNN Dataflows: A Data-Centric Approach"
Hyoukjun Kwon, Prasanth Chatarasi, Micheal Pellauer, Angshuman Parashar, Vivek Sarkar, and **Tushar Krishna**
In Proc of 52nd Annual IEEE/ACM International Symposium on Microarchitecture, Oct 2019
Acceptance Rate: 23% (79/344)

MICRO 2019 "SWAP: Synchronized Weaving of Adjacent Packets for Network Deadlock Resolution"
Mayank Parasar, Natalie Enright Jerger, Paul Gratz, Joshua San Miguel, and **Tushar Krishna**
In Proc of 52nd Annual IEEE/ACM International Symposium on Microarchitecture, Oct 2019
Acceptance Rate: 23% (79/344)

NOCS 2019 "BINDU: Deadlock-Freedom with One Bubble in the Network"
Mayank Parasar and **Tushar Krishna**
In Proc of 13th IEEE/ACM International Symposium on Networks-on-Chip, Oct 2019

FPL 2019 "Scaling the Cascades: Interconnect-aware mapping strategies for FPGA implementation of Machine Learning problems"
Ananda Samajdar, Tushar Garg, **Tushar Krishna** and Nachiket Kapre
In Proc of the Field-Programmable Logic and Applications, Sept 2019
Acceptance Rate: 18.5% (28/151)

DAC 2019 "Architecture, Chip, and Package Co-design Flow for 2.5D IC Design Enabling Heterogeneous IP Reuse"
Jinwoo Kim, Gauthaman Murali, Heechun Park, Eric Qin, Hyoukjun Kwon, Venkata Chaitanya Krishna Chekuri, Nihar Dasari, Arvind Singh, Minah Lee, Hakki Mert Torun, Kallol Roy, Madhavan Swaminathan, Saibal Mukhopadhyay, **Tushar Krishna**, and Sung Kyu Lim
In Proc of the 56th ACM/EDAC/IEEE Design Automation Conference, Jun 2019

ISPASS 2019 "mRNA: Enabling Efficient Mapping Space Exploration on a Reconfigurable Neural Accelerator"
Zhongyuan Zhao, Hyoukjun Kwon, Sachit Kuhar, Weiguang Sheng, Zhigang Mao, and **Tushar Krishna**
In Proc of the IEEE International Symposium on Performance Analysis of Systems and Software, Mar 2019
Acceptance Rate: 29.5% (26/88)

AISTECS 2019 "McMahon: Minimum-cycle Maximum-hop network"
Gokul Subramanian Ravi, **Tushar Krishna**, and Mikko Lipasti
In Proc of 4th International Workshop on Advanced Interconnect Solutions and Technologies for Emerging Computing Systems, Jan 2019
[Held in conjunction with the 14th HiPEAC Conference on High-Performance Embedded Architectures and Compilers]

ICRC 2018 "Merge Network for a Non-von Neumann Accumulate Accelerator in a 3D Chip"
Anirudh Jain, Sriseshan Srikanth, Erik Debenedictis, and **Tushar Krishna**
In Proc of 3rd IEEE International Conference on Rebooting Computing, Nov 2018

MICRO 2018
Finalist at ACM Student Research Competition at ASPLOS 2018
"GeneSys: Enabling Continuous Learning through Neural Network Evolution in Hardware"
Ananda Samajdar, Parth Mannan, Kartikay Garg, and **Tushar Krishna**
In Proc of 51st Annual IEEE/ACM International Symposium on Microarchitecture, Oct 2018
Acceptance Rate: 21% (74/351)

MICRO 2018 "Scalable Distributed Last-Level TLBs Using Low-Latency Interconnects"
Srikant Bharadwaj, Guilherme Cox, **Tushar Krishna**, and Abhishek Bhattacharjee
In Proc of 51st Annual IEEE/ACM International Symposium on Microarchitecture, Oct 2018
Acceptance Rate: 21% (74/351)

NOCS 2018 "Brownian Bubble Router: Enabling Deadlock Freedom via Guaranteed Forward Progress"
Mayank Parasar, Ankit Sinha, and **Tushar Krishna**
In Proc of 11th International Symposium on Networks-on-Chip, Oct 2018

ISCA 2018
Selected for IEEE Micro Top Picks
"Synchronized Progress in Interconnection Networks (SPIN) : A New Theory for Deadlock Freedom"
Aniruddh Ramrakhyani, Paul Gratz, and **Tushar Krishna**
In Proc of 45th International Symposium on Computer Architecture, Jun 2018
Acceptance Rate: 16.9% (64/378)

ISCA 2018 "FastTrack: Leveraging Heterogeneous FPGA Wires to Design Low-cost High-performance Soft NoCs"
Nachiket Kapre and **Tushar Krishna**
In Proc of 45th International Symposium on Computer Architecture, Jun 2018

ISCA 2018 "SEESAW: Using Superpages to Improve VIPT Caches"
Mayank Parasar, Abhishek Bhattacharjee, and **Tushar Krishna**
In Proc of 45th International Symposium on Computer Architecture, Jun 2018
Acceptance Rate: 16.9% (64/378)

ISPASS 2018 "Performance Implications of NoCs on 3D-Stacked Memories: Insights from the Hybrid Memory Cube"
Ramyad Hadidi, Bahar Asgari, Jeffrey Young, Burhan Ahmad Mudassar, Kartikay Garg, **Tushar Krishna**, and Hyesoon Kim
In Proc of the IEEE Int. Symp. on Performance Analysis of Systems and Software, Apr 2018
Acceptance Rate: 31.3% (21/67)

ASPLOS 2018 "MAERI: Enabling Flexible Dataflow Mapping over DNN Accelerators via Reconfigurable Interconnects"
Honorable Mention in IEEE Micro Top Picks
Hyoukjun Kwon, Ananda Samajdar, and **Tushar Krishna**
In Proc of the 23rd ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Mar 2018
424 citations
Acceptance Rate: 17.5% (56/319)

ASPLOS 2018 "LATR: Lazy Translation Coherence"
Mohan Kumar, Steffen Maass, Sanidhya Kashyap, Jan Vesely, Zi Yan, Taesoo Kim, Abhishek Bhattacharjee, and **Tushar Krishna**
In Proc of the 23rd ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Mar 2018
Acceptance Rate: 17.5% (56/319)

SysML 2018 "MAERI: Enabling Flexible Dataflow Mapping over DNN Accelerators via Programmable Interconnects"
Hyoukjun Kwon, Ananda Samajdar, and **Tushar Krishna**
In Inaugural SysML Conference, Feb 2018

AISTECS 2018 "Spoofing Prevention via RF Power Profiling in Wireless Network-on-Chip"
Brian Lebiednik, Sergi Abadal, Hyouk Jun Kwon and **Tushar Krishna**
Proc of 3rd International Workshop on Advanced Interconnect Solutions and Technologies for Emerging Computing Systems, Jan 2018
[Held in conjunction with the 13th HiPEAC Conference on High Performance Embedded Architectures and Compilers]

ICCAD 2017 "A Case for Low Frequency Single Cycle Multi Hop NoCs for Energy Efficiency and High Performance"
Monodeep Kar and **Tushar Krishna**
Proc of the IEEE/ACM International Conference on Computer-Aided Design, Nov 2017
Acceptance Rate: 26% (105/399)

NocArc 2017 "Lightweight Emulation of Virtual Channels using Swaps"
Mayank Parasar and **Tushar Krishna**
Proc of 10th International Workshop on Networks-on-Chip Architectures, Oct 2017
[Held in conjunction with the 50th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO-50)]

NOCS 2017 "Rethinking NoCs for Spatial Neural Network Accelerators"
Hyoukjun Kwon, Ananda Samajdar, and **Tushar Krishna**
Proc of 11th International Symposium on Networks-on-Chip, Oct 2017
Acceptance Rate: 32% (14/44)

NOCS 2017 "Adaptive Manycore Architectures for Big Data Computing"
Janardhan Rao Doppa, Ryan Gary Kim, Mihailo Isakov, Michel A. Kinsy, Hyoukjun Kwon, and Tushar Krishna
Proc of 11th International Symposium on Networks-on-Chip, Oct 2017 (Special Session Paper)

ISPASS 2017 "OpenSMART: Single-Cycle Multi-hop NoC Generator in BSV and Chisel"
Hyoukjun Kwon and **Tushar Krishna**
Proc of the IEEE International Symp. on Performance Analysis of Systems and Software, Apr 2017
Acceptance Rate: 29% (24/81)

DATE 2017
Best Paper Award "Automatic Place-and-Route of emerging LED-driven wires within a monolithically-integrated CMOS+III-V process"
Tushar Krishna, Arya Balachandran, Siau Ben Chiah, Li Zhang, Bing Wang, Cong Wang, Kenneth Lee Eng Kian, Jurgen Michel and Li-Shiuan Peh
Proc of Design Automation and Test in Europe, Mar 2017
Acceptance Rate: 24% (194/794)

HPCA 2017 "Static Bubble: A Framework for Deadlock-free Irregular On-chip Topologies"
Aniruddh Ramrakhyani and **Tushar Krishna**
Proc of the 23rd IEEE International Symp. on High-Performance Computer Architecture, Feb 2017
Acceptance Rate: 22% (50/224)

ISSCC 2016
Media Coverage: **IEEE Spectrum**, **Daily Mail**, **The Verge**, **phys.org**, **MIT News** "Eyeriss: An Energy-Efficient Reconfigurable Accelerator for Deep Convolutional Neural Networks"
Yu-Hsin Chen, **Tushar Krishna**, Joel Emer, and Vivienne Sze
Proc. of the IEEE International Solid-State Circuits Conference, Feb 2016

NOCS 2014
Best Paper Award "Single-Cycle Collective Communication Over A Shared Network Fabric"
Tushar Krishna and Li-Shiuan Peh
Proc. of the 8th International Symposium on Networks-on-Chip, Sep 2014
Acceptance Rate: 25% (21/83)

Hot Chips 2014 "SCORPIO: A 36-Core Research Chip Demonstrating Snoopy Coherence on a Scalable Mesh NoC with In-Network Ordering"
Chia-Hsin Owen Chen, Sunghyun Park, Suvinay Subramanian, **Tushar Krishna**, Bhavya K. Daya, Woo-Cheol Kwon, Brett Wilkerson, John Arends, Anantha P. Chandrakasan, and Li-Shiuan Peh
Proc. of Hot Chips 26: A Symposium on High Performance Chips, Aug 2014

ISCA 2014
Media Coverage: ***Wired***, ***PC World***, ***Geek***, ***Phys***, ***Tech***, ***The Registrar***, *etc.* "SCORPIO: A 36-Core Research Chip Demonstrating Snoopy Coherence on a Scalable Mesh NoC with In-Network Ordering"
Bhavya K. Daya, Chia-Hsin Owen Chen, Suvinay Subramanian, Woo-Cheol Kwon, Sunghyun Park, **Tushar Krishna**, Jim Holt, Anantha P. Chandrakasan, and Li-Shiuan Peh
Proc. of the 41st International Symposium on Computer Architecture, Jun 2014
Acceptance Rate: 18% (46/258)

ASPLOS 2014 "Locality-Oblivious Cache Organization leveraging Single-Cycle Multi-Hop NoCs"
Woo-Cheol Kwon, **Tushar Krishna**, and Li-Shiuan Peh
Proc. of the 19th International Conference on Architectural Support for Programming Languages and Operating Systems, Mar 2014
Acceptance Rate: 22% (49/217)

DATE 2013 SMART: A Single-Cycle Reconfigurable NoC for SoC Applications"
Chia-Hsin Owen Chen, Sunghyun Park, **Tushar Krishna**, Suvinay Subramanian, Anantha P. Chandrakasan, and Li-Shiuan Peh
Proc. of Design Automation and Test in Europe, Mar 2013
Acceptance Rate: 25% (206/829)

HPCA 2013 "Breaking the On-Chip Latency Barrier Using SMART"
Selected for IEEE Micro Top Picks **Tushar Krishna**, Chia-Hsin Owen Chen, Woo Cheol Kwon and Li-Shiuan Peh
Proc. of the 19th IEEE International Symp. on High-Performance Computer Architecture, Feb 2013
Acceptance Rate: 20% (51/249)

DAC 2012 "Approaching the Theoretical Limits of a Mesh NoC with a 16-Node Chip Prototype in 45nm SOI"
Media Coverage: EE Times, Slashdot, ACM, IT World, etc. Sunghyun Park, **Tushar Krishna**, Chia-Hsin Chen, Bhavya K. Daya, Anantha Chandrakasan, and Li-Shiuan Peh
Proc. of the 49th Design Automation Conference, Jun 2012
Acceptance Rate: 22% (164/741)

MICRO 2011 "Towards the Ideal On-chip Fabric for 1-to-Many and Many-to-1 Communication"
Tushar Krishna, Li-Shiuan Peh, Bradford M. Beckmann, and Steven K. Reinhardt
Proc. of the 44th IEEE/ACM International Symposium on Microarchitecture, Dec 2011
Acceptance Rate: 21% (44/208)

ICCAD 2011 "A Low-Swing Crossbar and Link Generator for Low-Power Networks-on-Chip"
Chia-Hsin Owen Chen, Sunghyun Park, **Tushar Krishna** and Li-Shiuan Peh
Proc. of the IEEE/ACM International Conference on Computer-Aided Design, Nov 2011
Acceptance Rate: 30% (106/349)

ICCD 2010 "SWIFT: A SWing-reduced Interconnect For a Token-based Network-on-Chip in 90 nm CMOS"
Tushar Krishna, Jacob Postman, Christopher Edmonds, Li-Shiuan Peh and Patrick Chiang,
Proc. of the 28th IEEE International Conference on Computer Design, Oct 2010
Acceptance Rate: 29% (79/267)

NOCS 2010 "Physical vs Virtual Express Topologies with Low-Swing Links for Future Many-core NoCs"
Chia-Hsin Owen Chen, Niket Agarwal, **Tushar Krishna**, Kyung-Hoae Koo, Li-Shiuan Peh and Krishna Saraswat
Proc. of the 4th International Symposium on Networks-on-Chip, May 2010
Acceptance Rate: 27% (32/118)

ISPASS 2009 "GARNET: A Detailed On-Chip Network Model inside a Full-System Simulator"
913 citations Niket Agarwal, **Tushar Krishna**, Li-Shiuan Peh and Niraj K. Jha
Proc. of the International Symp. on Performance Analysis of Systems and Software, April 2009
Acceptance Rate: 28% (24/86)

ICCAD 2008 "Texture Filter Memory – A Power-efficient and Scalable Texture Memory Architecture for Mobile Graphics Processors"
Silpa BVN, Anjul Patney, **Tushar Krishna**, Preeti R. Panda and G.S. Visweswaran

Proc. of the International Conference on Computer-Aided Design, Nov. 2008.
Acceptance Rate: 26% (122/458)

HotI 2008
Selected for IEEE Micro Top Picks

"NoC with Near-Ideal Express Virtual Channels Using Global-Line Communication"
Tushar Krishna, Amit Kumar, Patrick Chiang, Mattan Erez, and Li-Shiuan Peh
Proc. of the 16th International Symposium on High-Performance Interconnects, Aug. 2008.
Acceptance Rate: 30% (19/63)

"Modeling Electron Transport Mechanism in a Molecular Diode through *ab initio* Molecular Energy Calculations"
Tushar Krishna, C Kiran, Dilip K. Maity and Swapan K Ghosh
Proc. of the DAE-BRNS Theme Meeting on Materials Modeling at Different Length Scales, BARC, Mumbai, India, 2006

PUBLICATIONS: REFEREED JOURNALS

TACO 2023

"TNT: A Modular Approach to Traversing Physically Heterogeneous NOCs at Bare-wire Latency"
Gokul Subramanian Ravi, **Tushar Krishna**, and Mikko Lipasti
ACM Transactions on Architecture and Code Optimization, Jul 2023

TACO 2021

"Marvel: A Data-Centric Approach for Mapping Deep Learning Operators on Spatial Accelerators"
Prasanth Chatarasi, Hyoukjun Kwon, Angshuman Parashar, Michael Pellauer, **Tushar Krishna** and Vivek Sarkar
ACM Transactions on Architecture and Code Optimization, 2021

CAL 2021

"STONNE: Enabling Cycle-Level Microarchitectural Simulation for DNN Inference Accelerators"
Francisco Muñoz-Martinez, Jose L. Abellan, Manuel E. Acacio and Tushar Krishna
In IEEE Computer Architecture Letters, Issue 2, Jul-Dec 2021

TPDS 2021

"Evaluating Spatial Accelerator Architectures with Tiled Matrix-Matrix Multiplication"
Gordon E. Moon, Hyoukjun Kwon, Geonhwa Jeong, Prasanth Chatarasi, Sivasankaran Rajamanickam and **Tushar Krishna**
IEEE Transactions on Parallel and Distributed Systems (Special Section on Innovative R&D toward the Exascale Era), 2021

TVLSI 2021

"Clock Delivery Network Design and Analysis for Interposer-Based 2.5-D Heterogeneous Systems"
Gauthaman Murali, Heechun Park, Eric Qin, Hakki Mert Torun, Majid Ahadi Dolatsara, Madhavan Swaminathan, **Tushar Krishna** and Sung Kyu Lim
In IEEE Transactions on Very Large Scale Integration (VLSI) Systems, Apr 2021

CAL 2021

"Flexion: A Quantitative Metric for Flexibility in DNN Accelerators"
Hyoukjun Kwon, Michael Pellauer, Angshuman Parashar and **Tushar Krishna**
IEEE Computer Architecture Letters, Jan 2021

TVLSI 2020

"Architecture, Chip, and Package Co-design Flow for 2.5D IC Design Enabling Heterogeneous IP Reuse"
Jinwoo Kim, Gauthaman Murali, Heechun Park, Eric Qin, Hyoukjun Kwon, Venkata Chaitanya Krishna Chekuri, Nael Mizanur Rahman, Nihar Dasari, Arvind Singh, Minah Lee, Hakki Mert

Torun, Kallol Roy, Madhavan Swaminathan, Saibal Mukhopadhyay, **Tushar Krishna**, and Sung Kyu Lim
IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 2020

TACO 2020
"EcoTLB: Eventually Consistent TLBs"
Steffen Maass, Mohan Kumar, Taesoo Kim, **Tushar Krishna**, and Abhishek Bhattacharjee
ACM Transactions on Architecture and Code Optimization, 2020

IEEE Micro Top Picks 2020
"MAESTRO: A Data-Centric Approach to Understand Reuse, Performance, and Hardware Cost of DNN Mappings"
Hyoukjun Kwon, Prasanth Chatarasi, Michael Pellauer, Angshuman Parashar, Vivek Sarkar, and **Tushar Krishna**
*IEEE Micro (**Special Issue: Top Picks from the Computer Architecture Conferences**), May/Jun 2020*

IEEE Micro Top Picks 2019
"Synchronized Progress in Interconnection Networks (SPIN): A New Theory for Deadlock Freedom"
Aniruddh Ramrakhyani, Paul Gratz, and **Tushar Krishna**
*IEEE Micro (**Special Issue: Top Picks from the Computer Architecture Conferences**), May/Jun 2019*

IEEE Micro Special Issue 2018
"A Communication-Centric Approach for Designing Flexible DNN Accelerators"
Hyoukjun Kwon, Ananda Samajdar, and **Tushar Krishna**
IEEE Micro Special Issue on Hardware Acceleration, Nov/Dec 2018

JSSC 2017
3315 citations
"Eyeriss: An Energy-Efficient Reconfigurable Accelerator for Deep Convolutional Neural Networks"
Yu-Hsin Chen, **Tushar Krishna**, Joel Emer, and Vivienne Sze
IEEE Journal of Solid State Circuits Conference, ISSCC Special Issue, Jan 2017

TOCS 2015
"Efficient Control and Communication Paradigms for Coarse-Grained Spatial Architectures"
M. Pellauer, A. Parashar, M. Adler, B. Ahsan, R. Allmon, N. Crago, K. Fleming, M. Gambhir, A. Jaleel, **T. Krishna**, D. Lustig, S. Maresh, V. Pavlov, R. Rayess, A. Zhai, and J. Emer
ACM Transactions on Computer Systems, Sep 2015

IEEE Micro Top Picks 2014
"SMART: Single-Cycle Multihop Traversals Over A Shared Network-on-Chip"
Tushar Krishna, Chia-Hsin Owen Chen, Woo-Cheol Kwon, and Li-Shiuan Peh
*IEEE Micro (**Special Issue: Top Picks from the Computer Architecture Conferences**), May/Jun 2014*

IEEE Computer 2013
Webex Chat with Guest Editor:
youtu.be/k_I8yc_CjBU
"Single-Cycle Multihop Asynchronous Repeated Traversal: A SMART Future for Reconfigurable On-Chip Networks"
Tushar Krishna, Chia-Hsin Owen Chen, Sunghyun Park, Woo-Cheol Kwon, Suvinay Subramanian, Anantha P. Chandrakasan, and Li-Shiuan Peh
IEEE Computer, 46(10): 48-55, Oct 2013

TVLSI 2012
"SWIFT: A Low-Power Network-On-Chip Implementing the Token Flow Control Router Architecture With Swing-Reduced Interconnects"
Jacob Postman, **Tushar Krishna**, Christopher Edmonds, Li-Shiuan Peh, and Patrick Chiang
IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 21(8): 1432-1446, Aug 2012

CAN 2011 "The gem5 simulator"
5664 citations N. Binkert, B. Beckmann, G. Black, S. K. Reinhardt, A. Saidi, A. Basu, J. Hestness, D. R. Hower, **T. Krishna**, S. Sardashti, R. Sen, K. Sewell, M. Shoaib, N. Vaish, M. D. Hill and D. A. Wood
SIGARCH Computer Architecture News, 39(2): 1-7, May 2011

IEEE Micro Top Picks 2009 "Express Virtual Channels with Capacitively-Driven Global Links"
Tushar Krishna, Amit Kumar, Jacob Postman, Patrick Chiang, Mattan Erez, and Li-Shiuan Peh
*IEEE Micro (**Special Issue: Top Picks from Hot Interconnects 16**), 29 (4): 48-61, Jul/Aug 2009*

PATENTS

"Message Broadcast with Router Bypassing"
Tushar Krishna, Bradford M. Beckmann, Steven K. Reinhardt.
US Patent 2011/0314255 A1, Issued: Dec 22, 2011

OPEN-SOURCE SOFTWARE RELEASES

July 2023 Chakra: Advancing Performance Benchmarking and Co-design using Standardized Execution Traces
https://github.com/mlcommons/chakra

November 2021 STONNE: A Simulation Tool for Neural Networks Engines
https://stonne-simulator.github.io/

August 2020 ASTRA-sim: Distributed Deep Learning Training simulator
https://astra-sim.github.io/

Oct 2018 SCALE-Sim: Systolic CNN Accelerator Simulator
https://github.com/ARM-software/SCALE-Sim

June 2018 MAESTRO: An Analytic Model for Cost-Benefit Analysis of Dataflows in DNN Accelerators
http://maestro.ece.gatech.edu

June 2018 MAERI: An Open Source Framework for Generating Modular DNN Accelerators supporting Flexible Dataflow
https://github.com/maeri-project

May 2017 OpenSMART: A Tool for Automated Integration of Heterogeneous IPs
http://synergy.ece.gatech.edu/tools/opensmart

Oct 2016 Garnet2.0: An On-Chip Network Model for Heterogeneous SoCs
http://synergy.ece.gatech.edu/tools/garnet

INVITED TALKS

Oct 2023 "Chakra and ASTRA-sim: An open- source ecosystem for advancing co- design for future AI system"
At Open Compute Project (OCP) Global Summit, San Jose, CA, USA

Oct 2023 "Chakra and ASTRA-sim: An open- source ecosystem for advancing co- design for future AI system"
At Exascale Simulation of Next-Generation Computing Architectures Workshop @IISWC (Virtual)

Sept 2023 "Data Parallelism and Sharding"
Guest Lecture in CS 243: Advanced Computer Networks, Harvard University, Cambridge, MA

Jun 2023 "Modeling and Mitigating Communication Bottlenecks for Large Model Training at Scale"
Invited Keynote at Cognitive Architectures Workshop @ISCA, Orlando, FL, USA

May 2023 "Communication-centric System Architectures for AI Acceleration"
A*t Hewlett Packard Enterprise (HPE) Labs, Milpitas, CA, USA*

Mar 2023 "The Whole is Greater than the Sum of its Parts – A story of cross-stack collaborations"
Keynote at Young Architect Workshop @ ASPLOS, Vancouver, BC, Canada

Nov 2022 "Communication-centric System Architectures for AI Acceleration"
At Department of Computer Science, Yale University, New Haven, CT, USA

Nov 2022 "Demystifying Map Space Exploration for NPUs"
At IEEE International Symposium on Workload Characterization, Austin, TX, USA

Nov 2022 "A Formalism of DNN Accelerator Flexibility"
At Workshop on Hardware and Algorithms for Learning On-a-chip (HALO), ICCAD (virtual)

Sept 2022 "Formalizing Accelerator Flexibility – Systematic Representations & Search Methods"
At Department of EECS, University of California, Berkeley, CA, USA

Aug 2022 "ML-assisted Communication-aware Co-Design for Next-generation ML Systems"
At Google Workshop on Accelerators, Compute, Reliability, Security Workshop (virtual)

Jul 2022 "Formalizing Design-space Exploration for Next-Generation AI Accelerators"
At IEEE Silicon Integration Initiative (Si2) Low-Power Forum @ Design Automation Conference (DAC), San Jose, CA

Jun 2022 "Formalizing Accelerator Flexibility – Systematic Representations & Search Methods"
at *Department of ECE, CALCM Seminar Series, Carnegie Mellon University, PA, USA*

Apr 2022 "Domain-aware Genetic Algorithms for HW-Mapping Co-optimization for DNN Accelerators"
At Meta Reality Labs, Palo Alto, CA, USA (virtual)

Apr 2022 "Domain-aware Genetic Algorithms for HW-Mapping Co-optimization for DNN Accelerators"
At NVIDIA Research, Westford, MA, USA (virtual)

Feb 2022 " A Communication-centric Approach for Designing Flexible Accelerators"
At Minisymposum on Co-Design of Dataflow Accelerators, SIAM Conference on Parallel Processing for Scientific Computing, Seattle, WA, USA (virtual due to COVID-19)

Dec 2021 "Teaching AI the Game of AI Accelerator Design"
At Rutgers Efficient AI Seminar, Rutgers University, NJ, USA (virtual)

Nov 2021 "SEEC: Stochastic Escape Express Channel"
At International Conference for High Performance Computing, Networking, Storage, and Analysis (SC), Dallas, TX, USA.

Oct 2021 "ASTRA-sim: Enabling SW/HW Co-Design Exploration for Distributed Deep Learning Training Platforms"
At Workshop on Modeling & Simulation of Systems and Applications (Modsim). (virtual)

Sept 2021 "A Communication-centric Approach for Designing Flexible DNN Accelerators"
At Draper Labs, Cambridge, MA, USA (virtual)

Oct 2020 "Teaching AI the Game of AI Accelerator Design"
At Facebook AI Research, Cambridge, MA, USA. (virtual due to COVID-19)

Oct 2020 "Teaching AI the Game of AI Accelerator Design"
At Google Brain, Sunnyvale, CA, USA. (virtual due to COVID-19)

Oct 2020 "A Communication-centric Approach for Designing Flexible DNN Accelerators"
At 3rd China Symposium on Cognitive Computing and Hybrid Intelligence (CCHI'2020) (virtual due to COVID-19)

Sep 2020 "Communication-centric Design of Distributed Deep Learning Training Platforms"
at Computer Engineering and Systems Group, Texas A& M University (virtual due to COVID-19)

Aug 2020 "Benchmarking DNN Models on Future Inference and Training Platforms"
at Workshop on Benchmarking Machine Learning Workloads, co-located with EEE International Symposium on Performance Analysis of Systems and Software (ISPASS) (virtual due to COVID-19)

May 2020 "A Communication-centric Approach for Designing Flexible DNN Accelerators"
at 2nd Workshop on Accelerated Machine Learning (AccML), Co-located with International Symposium on Computer Architecture (ISCA) (virtual due to COVID-19)

May 2020 "A Systematic Methodology for Characterizing Scalability of DNN Accelerators using SCALE-Sim"
Guest Lecture in MIT 6.888 (Hardware Architecture for Deep Learning) *(virtual due to COVID-19)*

May 2020 "Designing Efficient DL Training Platforms by optimizing Communication within and across Accelerators"
at Facebook Research, Menlo Park, CA USA, (virtual due to COVID-19)

Mar 2020 "Enabling Continuous Learning through Synaptic Plasticity in Hardware"
at *Department of CS, Harvard University, Cambridge, MA, USA (virtual due to COVID-19)*

Feb 2020 "Enabling Continuous Learning through Synaptic Plasticity in Hardware"
at *Department of CSE, University of California, San Diego, CA, USA*

Feb 2020 "Enabling Continuous Learning through Synaptic Plasticity in Hardware"
at *Department of ECE, Cornell University, Ithaca, NY, USA*

Jan 2020 "Enabling Continuous Learning through Synaptic Plasticity in Hardware"
at *Department of CSE, University of Michigan, Ann Arbor, MI, USA*

Dec 2019 "Enabling Continuous Learning through Synaptic Plasticity in Hardware"
at *Department of CS, University of Illinois, Urbana Champaign, IL, USA*

Nov 2019 "Enabling Continuous Learning through Synaptic Plasticity in Hardware"
at *Department of ECE, University of Toronto, Toronto, Canada*

Nov 2019 "Enabling Continuous Learning through Synaptic Plasticity in Hardware"
at *Department of Electrical Engineering and Computer Science, MIT, Cambridge, MA, USA*

Nov 2019 "Enabling Continuous Learning through Synaptic Plasticity in Hardware"
at *Department of CS, Columbia University, New York City, NY, USA*

Nov 2019 "Enabling Continuous Learning through Synaptic Plasticity in Hardware"
at *Department of CS, Princeton University, Princeton, NJ, USA*

Nov 2019 "Enabling Continuous Learning through Synaptic Plasticity in Hardware"
at *School of Computer and Communication Sciences, EPFL, Austin, Lausanne, Switzerland*

Sept 2019 "Communication-Aware HW-SW Co-Design for Scalable Deep Learning"
at *Facebook AI Systems Faculty Summit, Menlo Park, CA, USA*

Sept 2019 "Enabling Continuous Learning through Synaptic Plasticity in Hardware"
at *Department of ECE, University of Texas, Austin, TX, USA*

Sept 2019 "Enabling Continuous Learning through Neural Network Evolution in Hardware"
at *ARM Research Summit, Austin, TX, USA*
Winner of Best Presentation Award

July 2019 "Enabling Efficient Dataflow Exploration for Reconfigurable Neural Accelerators"
at *International Conference on Neuromorphic Systems (ICONS), Knoxville, TN, USA*

June 2019 "Using ML to Design ML Accelerators"
at *Google Brain, Sunnyvale, CA, USA*

June 2019 "Enabling Continuous Learning through Synaptic Plasticity in Hardware"
at *4th Workshop on Energy Efficient Machine Learning and Cognitive Computing for Embedded Applications (EMC2), Co-Located with ISCA, Phoenix, AZ, USA*

June 2019 "Enabling Pervasive AI through Dataflow-Hardware Co-Design"
at *Mythic AI, Redwood City, CA, USA*

June 2019 "Dataflow and Communication-Aware Deep Learning Accelerator Design"
at *Google Edge TPU Team, Mountain View, CA USA*

March 2019 "Enabling Continuous Learning through Synaptic Plasticity in Hardware"
at *University of Wisconsin, Madison, WI, USA*

March 2019 "From Dataflows to Communication Flows in Spatial DNN Accelerators"
at *FastPath Workshop, Co-located with ISPASS, Madison, WI, USA*

February 2019 "Dataflow and Communication-Aware Deep Learning Accelerator Design"
at *Xilinx Research, San Jose, CA USA*

February 2019 "DNN-Dataflow-Hardware Co-Design for Enabling Pervasive General-Purpose AI"
at *Facebook Research, Menlo Park, CA USA*

November 2018 "DNN-Dataflow-Hardware Co-Design for Enabling Pervasive General-Purpose AI"
at *Annual CRNCH Summit, Georgia Institute of Technology, Atlanta, GA, USA*

October 2018 "Enabling Pervasive General-Purpose AI"
at *Oakridge National Laboratory (ORNL), Oak Ridge, TN, USA*

August 2018 "MAESTRO: An Analytic Model for Cost-Benefit Analysis of Dataflows in DNN Accelerators"
at *Workshop on Modeling and Simulation on Systems and Applications (Modsim), Seattle, WA, USA*

June 2018 "A Quest to Enable Pervasive General-Purpose AI"
at *IBM T. J. Watson Research Center, Yorktown Heights, NY, USA*

June 2018 "Enabling Rapid Design Space Exploration and Prototyping of DNN Accelerators"
at *CMU/GT Symposium on Machine Learning in Science and Engineering, Pittsburgh, PA, USA*

May 2018 "Enabling Irregular Dataflow Mapping via Reconfigurable Single-Cycle On-Chip Interconnects"
at *Lawrence Berkeley National Laboratory, Berkeley, CA, USA*

May 2018 "A Communication-centric Approach to Designing Flexible DNN Accelerators"
at *TPU Design Team, Google, Sunnyvale, CA, USA*

May 2018 "Communication-centric Computing for Deep Learning"
at *GT Institute for Electronics and Nanotechnology (IEN) Technical Conference, Atlanta, GA, USA*

Sept 2017 "Garnet2.0: A Detailed On-Chip Network Model inside a Full-System Simulator"
at *gem5 workshop, ARM Research Summit, Cambridge, U.K.*

Sept 2017 "Breaking the On-Chip Latency Barrier Using Single-cycle Multi-hop Networks"
at *ARM Research, Cambridge, U.K.*

May 2017 "Scalable and Low-Cost Heterogeneous Design in the Post-Moore's Law Era using Efficient Communication"
at *Georgia Tech ECE Advisory Board Meeting, Cupertino, CA, USA*

Apr 2017 "Breaking the On-Chip Latency Barrier Using Single-cycle Multi-hop Networks"
at *Computer Engineering and Systems Group Seminar, Texas A&M University, TX, USA*

Feb 2017 "Breaking the On-Chip Latency Barrier Using Single-cycle Multi-hop Networks"
at *Baidu Research, Sunnyvale, CA*

Oct 2016 "Breaking the On-Chip Latency Barrier Using Single-cycle Multi-hop Networks"
at *Department of ECE, CALCM Seminar Series, Carnegie Mellon University, PA, USA*

Mar 2016 "Breaking the On-Chip Latency Barrier Using SMART"

at *NVIDIA, Westford, MA, USA*

Mar 2015 "Breaking the On-Chip Latency Barrier Using SMART"
at *Department of CS, University of California at Los Angeles, CA, USA*

Feb 2015 "Breaking the On-Chip Latency Barrier Using SMART"
at *Department of CS, University of Illinois Urbana-Champaign, IL, USA*

Feb 2015 "Breaking the On-Chip Latency Barrier Using SMART"
at *Department of ECE, Georgia Tech, Atlanta, GA, USA*

Jan 2015 "Enabling dedicated single-cycle connections over a shared multi-hop network"
at *Department of ECE, Northeastern University, Boston, MA, USA*

Nov 2014 "Enabling dedicated single-cycle connections over a shared Network-on-Chip"
at *Department of CSE, University of Michigan, Ann Arbor, MI, USA*

Sep 2014 "Single-Cycle Collective Communication Over A Shared Network Fabric"
at *IEEE Intl. Symp. on Networks-on-Chip (NOCS-8), Ferrara, Italy*

Feb 2013 "Breaking the On-Chip Latency Barrier Using SMART"
at *IEEE Intl. Symp. on High-Performance Computer Architecture (HPCA-19), Shenzhen, China*

Jul 2012 "Breaking the On-Chip Latency Barrier Using SMART"
at *VSSAD, Intel Corporation, Hudson, MA, USA*

May 2012 "Reconfigurable on-chip network topologies using SMART links"
at *Industry Affiliates Program, CSAIL, MIT, Cambridge, MA, USA*

Dec 2011 "Towards the Ideal On-chip Fabric for 1-to-Many and Many-to-1 Communication"
at *IEEE/ACM Intl. Symp. on Microarchitecture (MICRO-44), Porte Alegre, Brazil*

Oct 2010 "SWIFT: A SWing-reduced Interconnect For a Token-based Network-on-Chip in 90 nm CMOS"
at *IEEE Intl. Conf. on Computer Design (ICCD-28), Amsterdam, Netherlands*

Feb 2010 "SWing-reduced Interconnect For a Token-based (SWIFT) Network-on-Chip"
at *Student Research Preview, Intl. Solid-State Circuits Conference (ISSCC), San Francisco, CA*

Oct 2008 "NoC with Near-Ideal Express Virtual Channels Using Global-Line Communication"
at *Interconnect Focus Center (IFC) Annual Review, Atlanta, GA*

TEACHING EXPERIENCE

Spring 2023 **Invited Lecturer for HW Accelerated Machine Learning at the Edge, MEAD Education**
Online Class. *Team of Instructors: Marian Verhelst (KU Leaven), Tijmen Blankevoort (Qualcomm), Luca Benini (EPFL, Univ of Bologna), Naveen Verma (Princeton), Vijay Janapa Reddi (Harvard), Tushar Krishna (Georgia Tech), Prasanth Chatarasi (IBM), David Atienza (EPFL), Jan Rabaey (Berkeley), Huichu Liu (Meta), Eduard Alarcon (UPC)*

Fall 2022 **Georgia Tech ECE 3058 (Architecture, Systems, Concurrency and Energy in Computation)**
Undergraduate-level class of 93 students. *Teaching Score: 4.5/5*

Spring 2022 **Georgia Tech ECE 6115 / CS 8803 – ICN (Interconnection Networks)**
Graduate-level class of 52 students. *Teaching Score: 4.9/5*

Spring 2022 **Invited Lecturer for HW Accelerated Machine Learning at the Edge, MEAD Education**
Online Class. *Team of Instructors: Marian Verhelst (KU Leaven), Tijmen Blankevoort (Qualcomm), Luca Benini (EPFL, Univ of Bologna), Naveen Verma (Princeton), Vijay Janapa Reddi (Harvard), Tushar Krishna (Georgia Tech), Prasanth Chatarasi (IBM), David Atienza (EPFL), Jan Rabaey (Berkeley), Huichu Liu (Meta), Eduard Alarcon (UPC)*

Fall 2021 **Georgia Tech ECE 3058 (Architecture, Systems, Concurrency and Energy in Computation)**
Undergraduate-level class of 90 students. *Teaching Score: 4.7/5*

Spring 2021 **Georgia Tech ECE 3058 (Architecture, Systems, Concurrency and Energy in Computation)**
Undergraduate-level class of 78 students. *Teaching Score: 4.7/5*

Spring 2021 **Georgia Tech ECE 6115 / CS 8803 – ICN (Interconnection Networks)**
Graduate-level class of 32 students. *Teaching Score: 4.9/5*

Spring 2020 **Georgia Tech ECE 3057 (Architecture, Systems, Concurrency and Energy in Computation)**
Undergraduate-level class of 59 students. *No teaching score due to COVID-19*

Spring 2020 **Georgia Tech ECE 6115 / CS 8803 – ICN (Interconnection Networks)**
Graduate-level class of 35 students. *No teaching score due to COVID-19*

Fall 2019 **Georgia Tech ECE4100/ECE6100/CS4290/CS6290 (Advanced Computer Architecture)**
Graduate-level class of 54 students. *Teaching Score: 4.85/5*

Spring 2019 **Georgia Tech ECE 8893 B (Hardware Accelerators for Machine Learning)**
NEW Graduate-level class of 60 students. *Teaching Score: 4.5/5*

Spring 2019 **Georgia Tech ECE 8823 A / CS 8803 – ICN (Interconnection Networks)**
Graduate-level class of 48 students. *Teaching Score: 4.9/5*

Fall 2018 **Georgia Tech ECE 3056 (Architecture, Concurrency and Energy in Computation)**
Undergraduate-level class of 86 students. *Teaching Score: 4.2/5*

Spring 2018 **Georgia Tech ECE 3056 (Architecture, Concurrency and Energy in Computation)**
Undergraduate-level class of 90 students. *Teaching Score: 4.3/5*

Spring 2018 **Georgia Tech ECE 8823 A / CS 8803 – ICN (Interconnection Networks)**
Graduate-level class of 26 students. *Teaching Score: 5/5*
+ Won the "Class of 1940 Course Survey Teaching Effectiveness Award"

Summer 2017 **Invited Lecturer for Network-on-Chip course at ACACES Summer School, Fuiggi, Italy**
Graduate-level class with 50+ students
The ACACES Summer School is a one-week summer school for computer architects and tool builders working in the field of high performance computer architecture and compilation for computing systems. It is organized by the High Performance and Embedded Architecture and Compilation (HiPEAC) consortium. The school aims at the dissemination of advanced scientific knowledge and the promotion of international contacts among scientists from academia and industry. It attracts PhD-students, faculty, and industry researchers.

Spring 2017 **Georgia Tech ECE 8823 A / CS 8803 – ICN (Interconnection Networks)**
Graduate-level class of 11 students. *Teaching Score: 4.85/5*

Fall 2016 **Georgia Tech ECE4100/ECE6100/CS4290/CS6290 (Advanced Computer Architecture)**
Graduate-level class of 62 students. *Teaching Score: 4.85/5*

Spring 2016 **Georgia Tech ECE 8823 A / CS 8803 – ICN (Interconnection Networks)**
NEW Graduate-level class of 21 students. *Teaching Score: 4.9/5*

Fall 2015 **Georgia Tech ECE4100/ECE6100/CS4290/CS6290 (Advanced Computer Architecture)**
Graduate-level class of 39 students. *Teaching Score: 4.7/5*

STUDENT ADVISING

PhD Students (Graduated)

Saeed Rashidi (PhD, ECE, January 2023)

- **Thesis:** *HW-SW Methods for Modeling and Optimizing Communication for Scalable Training of Deep Learning Models*
- **First Employment:** HP Labs

Sheng-Chun (Felix) Kao (PhD, ECE, May 2022)

- **Thesis:** *Domain-aware Genetic Algorithms for Hardware and Mapping Co-Optimization for Efficient DNN Acceleration*
- **First Employment:** Waymo

Eric Qin (PhD, ECE, March 2022)

- **Thesis:** *Building Efficient Tensor Accelerators for Sparse and Irregular Workloads*
- **First Employment:** Facebook
- **Select Honors during PhD**
    - Best Paper Candidate at IPDPS 2022
    - Best Paper at HPCA 2020

Ananda Samajdar (PhD, ECE, October 2021)

- **Thesis:** *Methodology and Analysis for Efficient Custom Architecture Design using Machine Learning*
- **First Employment:** IBM Research
- **Select Honors during PhD**
    - Best Paper at HPCA 2020
    - Honorable Mention in IEEE Micro Top Picks 2019
    - Silver Medalist at ACM Student Research Competition, ASPLOS 2019
    - Finalist at ACM Student Research Competition, ASPLOS 2018

Hyoukjun Kwon (PhD, SCS, July 2020)

- **Thesis:** *Data- and Communication-centric Approaches to Model and Design Flexible Deep Neural Network Accelerators*
- **First Employment:** Facebook Reality Labs
    - Now Assistant Professor, Department of EECS, UC Irvine.
- **Select Honors during PhD**
    - Honorable Mention at ACM SIGARCH / IEEE CS TCCA Outstanding Dissertation Award "*For developing mechanisms to quantify the relationship between deep neural network mappings, data reuse and communication flows for system design of flexible deep learning accelerators*
    - Best Paper at HPCA 2020
    - IEEE Micro Top Picks 2020
    - Qualcomm Innovation Fellowship Finalist 2019
    - Honorable Mention in IEEE Micro Top Picks 2019
    - Finalist at ACM Student Research Competition MICRO 2018

Mayank Parasar (PhD, ECE, July 2020)

- **Thesis:** *Subactive Techniques for Guaranteeing Routing and Protocol Deadlock*

Freedom in Interconnection Networks

- **First Employment:** Samsung Austin Research Center
- **Select Honors during PhD**
    - Best Paper candidate at SC 2021

PhD Students (In Progress)

Geonhwa Jeong
William (Jonghoon) Won
Raveesh Garg
Jianming Tong
Difei Cao (co-advised with Kishore Ramachandran)
Jinsun Yoo (co-advised with Kishore Ramachandran)
Abhimanyu Bambhaniya
Canlin Zhang
Jamin Seo
Changhai Man
Hanjiang Wu
Ritik Raj
Hao Kang
Zishen Wan

MS Students (Graduated)

Matthew Denton (MS, ECE, 2021)

- **Thesis:** *Acceleration of Sparse Matrix Multiplication using Bit-Serial Arithmetic*
- **First Employment:** Luminous Computing

Yehowshua Immanuel (MS, ECE, 2020)

- **Thesis:** *Plug-and-Play FOSS ML Accelerator: From Concept to Conception*
- **First Employment:** Founder of ChipEleven (Startup)

Vineet Nadella (MS, ECE, 2020)

- **Thesis:** *Investigating Opportunities and Challenges in Modeling and Designing Scale-Out DNN Accelerators*
- **First Employment:** Amazon

Parth Mannan (MS, ECE, 2018)

- **Thesis:** *Exploring Opportunities and Challenges in Enabling Neuro-Evolutionary Algorithms in Hardware*
- **First Employment:** NVIDIA

Srikant Bharadwaj (MS, ECE, 2017)

- **Thesis:** *Scaling Address Translation in Multi-core Architectures using Low-Latency Interconnects*
- **First Employment:** AMD Research (GPU Micro-architecture)

Aniruddh Ramrakhyani (MS, ECE, 2017)

- **Thesis:** *Deadlock Recovery in On-Chip Interconnection Networks*
- **First Employment:** Apple. Now at Google

PROFESSIONAL SERVICE

Professional Committees:

- Co-Chair of "Chakra Execution Traces and Benchmarks" Working Group at ML Commons (2023 – present)
- International Roadmap for Devices and Systems (IRDS™ 2020 – Present)
    - Applications Benchmarking International Focus Team (AI Team)
- Member of ACM, Senior Member of IEEE

Conference Organization:

- Technical **Program Committee Vice-Chair** for **ISCA 2023**.
- Technical **Program Committee Topic Chair** for *"Design Methodologies for Machine Learning Architectures"* for **DATE 2023**
    - *DATE has one TPC chair, 4 track chairs and multiple topics in each track, each with a chair and a co-chair.*
    - *Role as topic co-chair: pick the TPC for the track (~15-20 members), assign double-blind reviewers, run the TPC meeting for the track, create sessions for the topic.*
- Artifacts Evaluation Committee Co-chair for HPCA 2022.
- **Technical Program Committee Topic Chair** for *"Design Methodologies for Machine Learning Architectures"* for **DATE 2022**
    - *DATE has one TPC chair, 4 track chairs and multiple topics in each track, each with a chair and a co-chair.*
    - *Role as topics co-chair: pick the TPC for the track (~15-20 members), assign double-blind reviewers, run the TPC meeting for the track, create sessions for the topic.*
- **General Co-Chair NOCS 2021**
- Workshops and Tutorials Chair ISPASS 2021
- **Technical Program Committee Topic Co-Chair** for *"Design Methodologies for Machine Learning Architectures"* for **DATE 2021**
    - *DATE has one TPC chair, 4 track chairs and multiple topics in each track, each with a chair and a co-chair.*
    - *Role as track co-chair: pick the TPC for the track (~15-20 members), assign double-blind reviewers, run the TPC meeting for the track, create sessions for the topic.*
- **Technical Program Committee Co-Chair NOCS 2020**
    - *Role: select TPC (~40 members), select organizing committee, assign reviews, create and manage overall technical program, invite and coordinate keynotes.*
- Tutorials and Special Sessions Co-Chair NOCS 2019
- **TPC Track Chair** for track on On-Chip and On-Package Interconnects, **DAC 2019**
    - *DAC has one overall TPC chair and several tracks, each with one chair*
    - *Role as track chair: pick the TPC for the track, assign double-blind reviewers, run the TPC meeting for the track, create session for the track.*
- Publicity Co-Chair NOCS 2018
- Travel Grants Co-Chair PACT 2018
- Travel Grants Co-Chair ISCA 2017
- Workshops and Tutorials Chair ASPLOS 2016

Workshop Organization:

- Co-organized Annual CRNCH Summit at Georgia Tech, February 2023.
    - Invited talks from speakers and panelists from National Labs, Industry and

Academia on Post-Moore Computing

- Co-organized Annual CRNCH Summit at Georgia Tech, February 2022.
  - Invited talks from speakers and panelists from National Labs, Industry and Academia on Post-Moore Computing
- Co-organized a cross-cutting session on "AI Architectures and Co-Design" at the Department of Energy's "AI for Earth System Predictability" Workshop, 2021.
- Started and co-organized a new workshop *"Workshop on Architecture, Compiler, and System Support for Multi-model DNN Workloads"*
  - Offered at MICRO 2021, ISCA 2022.
- Started and co-organized a new workshop *"Systems for Machine Learning for Health (SysML4H)"* at MLSys 2021
- Started and co-organized a new workshop *"Unlocking the power of Edge Computing"* at ASPLOS 2019
- Started a new workshop *"Rising Stars in Computer Architecture"* in 2018 for final-year PhD students across USA interested in academia.
  - Co-Chaired the first (2018) and second iteration (2019) of the workshop. Solicited nominations and invited a select set of candidates to visit and present at Georgia Tech

Tutorial Organization:

- Started a new tutorial series "***ASTRA-sim: Enabling SW/HW Co-Design Exploration for Distributed Deep Learning Training Platforms***"
  - Offered at ASPLOS 2022, ISCA 2022, MLSys 2022, ASPLOS 2023
- Started a new Tutorial series ***"Enabling Rapid Design Space Exploration and Prototyping of DNN Accelerators"*** to aid researchers designing DNN accelerator Hardware-Software using the MAESTRO, MAERI, SCALE-sim, STONNE simulators.
  - MAERI tutorials offered at ISCA 2018, HPCA 2019, ICS 2022
  - MAESTRO tutorial offered a MICRO 2020
  - SCALE-sim tutorials at ISCA 2021 and ASPLOS 2021
  - STONNE tutorials at ASPLOS 2022, ASPLOS 2023

Technical Program Committee Member:

- International Symposium on Computer Architecture (**ISCA**)
  - 2019, 2020, 2021, 2022, 2023, 2024
- IEEE/ACM International Symposium on Microarchitecture (**MICRO**)
  - 2016, 2019, 2020, 2021, 2022, 2023
- IEEE International Symposium on High-Performance Computer Architecture (**HPCA**)
  - 2020, 2021, 2022
- ACM International Conference on Architectural Support for Programming Languages and Operating Systems (**ASPLOS**)
  - 2023
- IEEE Micro **Top Picks** in Computer Architecture
  - 2020, 2022, 2023
- Conference on Machine Learning and Systems (**MLSys**)
  - 2022, 2023, 2024
- IEEE International Symp. on Performance Analysis of Systems and Software (**ISPASS**)
  - 2018
- Design Automation Conference (**DAC**)
  - 2017, 2018, 2019
- Design Automation and Test in Europe (**DATE**)
  - 2017, 2018, 2019, 2020, 2021, 2022, 2023
- IEEE International Symposium on High-Performance Interconnects (**HOTI**)
  - 2016, 2017, 2019

- The International Conference for High Performance Computing, Networking, Storage, and Analysis (**SC**)
  - 2022
- ACM/SIGARCH International Conference on Supercomputing (**ICS**)
  - 2015
- IEEE International Parallel and Distributed Processing (**IPDPS**)
  - 2018
- International Conference on Parallel Architectures and Compilation Techniques (**PACT**)
  - 2019

Conference External Review Committee (ERC) Member:

- IEEE/ACM International Symposium on Microarchitecture (**MICRO**)
  - 2015, 2017
- International Conference on Parallel Architectures and Compilation Techniques (**PACT**)
  - 2016
- IEEE International Symposium on Computer Architecture (**ISCA**)
  - 2017
- IEEE International Symposium on High-Performance Computer Architecture (**HPCA**)
  - 2023

Journal Editorial Team:

- Guest Editor, ACM Journal on Emerging Technologies in Computing Systems, *Special Issue on Next-generation On-Chip and Off-Chip Communication Architectures for Edge, Cloud and HPC, 2022.*
- Guest Editor, *IEEE Transactions on Computing, Special Issue on Machine Learning Accelerators, 2022.*

Journal Reviewer:

- IEEE Computer Architecture Letters (CAL)
  - 2012, 2013, 2014, 2015, 2016, 2017, 2018, 2019, 2020, 2021, 2022, 2023
- IEEE Computer
  - 2013
- IEEE Micro
  - 2020
- ACM Transactions on Architecture and Code Optimization (TACO)
  - 2012, 2014, 2015, 2016, 2017, 2018, 2020
- IEEE Transactions on Very Large-Scale Integration Systems (TVLSI)
  - 2014, 2016, 2019
- IEEE Transactions on Parallel and Distributed Systems (TPDS)
  - 2015
- IEEE Journal on Emerging and Selected Topics in Circuits and Systems (JETCAS)
  - 2010
- IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)
  - 2016, 2017, 2020
- IEEE Transactions on Computers
  - 2016, 2018
- IEEE Transactions on Multi-Scale Computing Systems (TMCS)
  - 2017
- ACM Transactions on Embedded Computing Systems (TECS)
  - 2017

Panels: NSF (2016), DoE (2020)

INTERNAL SERVICE AT GEORGIA TECH

2023-24	**Chair of Computer Software and Systems (CSS) Technical Interest Group** in the School of ECE
2021	**Search Committee for Steve. W. Chaddick Chair** for the School of ECE
2019-22	**Faculty Recruitment Committee** for School of ECE
2017 – 18	**Search Committee for Steve. W. Chaddick Chair** for the School of ECE
2017 and 2018	**Graduate Student Recruitment Committee** for the School of ECE
2017-present	**Faculty Coordinator for CS 8001: Computer Architecture Seminar** *(a weekly seminar series for students and faculty across School of ECE and College of Computing working in Computer Architecture featuring student talks and guest lectures)*