Hyoukjun Kwon

Research Scientist @ Facebook

1 Hacker Way, Menlo Park, CA 94025

(+1) 404-539-4457 | hyoukjunkwon@fb.com | www.hyoukjunkwon.com | hyoukjun-kwon

Professional Experience

Facebook, Menlo Park, CA
Research Scientist at Facebook Reality Labs — *Oct. 2020 -*
Manager: Dr. Vikas Chandra

Facebook, Menlo Park, CA
Research Intern at AR/VR AI Research — *May. 2019 - July. 2019*
Manager: Dr. Vikas Chandra, Mentor: Dr. Liangzhen Lai

NVIDIA, Westford, MA
Research Intern at Architecture Research Group (ARG) — *May. 2018 - Aug. 2018*
Manager: Dr. Steve Keckler, Mentor: Dr. Michael Pellauer

NVIDIA, Westford, MA
Research Intern at Architecture Research Group (ARG) — *May. 2017 - Aug. 2017*
Manager: Dr. Steve Keckler, Mentor: Dr. Michael Pellauer

Georgia Institute of Technology, Atlanta, GA
Graduate Research Assistant — *Aug. 2015 - Jul. 2020*

Education

Georgia Institute of Technology
PhD in Computer Science — *Aug. 2015 - Aug. 2020*

- Advisor: Prof. Tushar Krishna
- Committee: Dr. Michael Pellauer, Prof. Vivek Sarkar, Prof. Hyesoon Kim, and Prof. Alexey Tumanov
- Thesis Title: Data- and Communication-centric Approaches to Model and Design Flexible Deep Neural Network Accelerators

Seoul National University (SNU)
BS in CSE (Computer Science and Engineering)
BS in EMS (Environmental Material Science) — *Mar. 2007 - Feb. 2015*

- Advisor: Dr. Jihong Kim (CSE) and Dr. Junjae Lee (EMS)

Research Interest

Deep neural network (DNN) accelerators
DNN mapping and dataflow optimization
Network-on-Chips (NoCs)
Reconfigurable architecture
Computer architecture
Machine learning
Design automation

Peer-reviewed Publications

Robert Guirado and **Hyoukjun Kwon (equal contribution)**, Sergi Abadal, Eduard Alarcon, and Tushar Krishna, *"Dataflow-Architecture Co-Design for 2.5D DNN Accelerators using Wireless Network-on-Package"*, *The 26th Asia and South Pacific Design Automation Conference (**ASP-DAC**)*, *2020*

Lei Yang, Zheyu Yan, Meng Li, **Hyoukjun Kwon**, Liangzhen Lai, Tushar Krishna, Vikas Chandra, Weiwen Jiang, Yiyu Shi, *"Co-Exploration of Neural Architectures and Heterogeneous ASIC Accelerator Designs Targeting Multiple Tasks"*, *The 57th Annual Design Automation Conference (**DAC**)*, *2020*

Hyoukjun Kwon, Prasanth Chatarasi, Michael Pellauer, Angshuman Parashar, Vivek Sarkar, and Tushar Krishna, *"MAESTRO: A Data-Centric Approach to Understand Reuse, Performance, and Hardware Cost of DNN Dataflows"*, *IEEE MICRO Top-Picks (**Top-Picks**)*, *2020*

Eric Qin, Ananda Samajdar, **Hyoukjun Kwon**, Vineet Nadella, Sudarshan Srinivasan, Dipankar Das, Bharat Kaul, and Tushar Krishna, *"SIGMA: A Sparse and Irregular GEMM Accelerator with Flexible Interconnects for DNN Training"*, *International Symposium on High-Performance Computer Architecture (**HPCA**)*, *2020*
Received the best paper award

Robert Guirado, **Hyoukjun Kwon**, Sergi Abadal, Eduard Alarcon, and Tushar Krishna, *"Understanding the Impact of On-Chip Communication on DNN Accelerator Performance"*, *International Conference on Electronics Circuits and System (**ICECS**)*, *2019*

Hyoukjun Kwon, Prasanth Chatarasi, Michael Pellauer, Angshuman Parashar, Vivek Sarkar, and Tushar Krishna, *"Understanding Reuse, Performance, and Hardware Cost of DNN Dataflows: A Data-Centric Approach"*, *International Symposium on Microarchitecture (**MICRO**)*, *2019*
Selected as Top Picks in Computer Architecture Conferences in 2019

Jinwoo Kim, Gauthaman Murali, Heechun Park, Eric Qin, **Hyoukjun Kwon**, Venkata Chaitanya Krishna, Nihar Dasari, Arvind Singh, Minah Lee, Hakki Torun, Kallol Roy, Madhavan Swaminathan, Saibal Mukhopadhyay, Tushar Krishna, and Sung Kyu Lim, *"mRNA: Enabling Efficient Mapping Space Exploration on a Reconfigurable Neural Accelerator"*, *Design Automation Conference (**DAC**)*, *2019*

Zhongyuan Zhao, **Hyoukjun Kwon**, Sachit Kuhar, Weiguang Sheng , Zhigang Mao, and Tushar Krishna, *"mRNA: Enabling Efficient Mapping Space Exploration on a Reconfigurable Neural Accelerator"*, *International Symposium on Performance Analysis of Systems and Software (**ISPASS**)*, *2019*

Hyoukjun Kwon, Ananda Smajdar, and Tushar Krishna, *"A Communication-driven Approach for Designing Flexible DNN Accelerators"*, *IEEE Micro Special Issue on Hardware Acceleration (**IEEE Micro**)*, *2018*

Brian Lebiednik, Sergi Abadal, **Hyoukjun Kwon** and Tushar Krishna, *"Architecting a Secure Wireless Network-on-Chip"*, *International Symposium on Network-on-Chips (**NOCS**)*, *2018*

Hyoukjun Kwon, Ananda Samajdar, and Tushar Krishna, *" MAERI: Enabling Flexible Dataflow Mapping over DNN Accelerators via Reconfigurable Interconnects"*, *International Conference on Architectural Support for Programming Languages and Operating Systems (**ASPLOS**)*, *2018*
Honorable mention for Top Picks in Computer Architecture Conferences in 2018

Hyoukjun Kwon, Ananda Samajdar, and Tushar Krishna, *" MAERI: Enabling Flexible Dataflow Mapping over*

DNN Accelerators via Reconfigurable Interconnects", *The Inaugural Sysml Conference (**Sysml**), not archived*, *2018*

Brian Lebiednik, Sergi Abadal, **Hyoukjun Kwon**, and Tushar Krishna, *"Spoofing Prevention via RF Power Profiling in Wireless Network-on-Chip"*, *International Workshop on Advanced Interconnect Solutions and Technologies for Emerging Computing Systems (**AISTECS**)*, *2018*

Hyoukjun Kwon, Ananda Samajdar, and Tushar Krishna, *"Rethinking NoCs for Spatial Neural Network Accelerators"*, *International Symposium on Networks-on-Chips (**NOCS**)*, *2017*

Janardhan Rao Doppa, Ryan Gary Kim, Mihailo Isakov, Michel A. Kinsy, **Hyoukjun Kwon**, and Tushar Krishna, *"Adaptive Manycore Architectures for Big Data Computing"*, *International Symposium on Networks-on-Chips (**NOCS**)*, *2017*

Hyoukjun Kwon, William Harris, and Hadi Esmaeilzadeh, *"Proving Flow Security of Sequential Logic via Automatically Synthesized Relational Invariants"*, *Computer Security Foundations (**CSF**)*, *2017*

Hyoukjun Kwon and Tushar Krishna, *"OpenSMART: Single-Cycle Multi-hop NoC Generator in BSV and Chisel"*, *International Symposium on Performance Analysis of Systems and Software (**ISPASS**)*, *2017*

Hyoukjun Kwon, Dohyun Kim, Jisung Park, and Jihong Kim, *"Improving the Lifetime of NAND Flash-based Storages Using MADE (Minhash-Assisted Delta-compression Engine)"*, *Korean Institute of Information Science and Engineers Annual Conference (KIISE)*, *2014*

Book

Tushar Krishna, **Hyoukjun Kwon**, Angshuman Parashar, Michael Pellauer, and Ananda Samajdar, *"Synthesis lecture on computer architecture: Data Orchestration in DNN Accelerators"*, *Morgan & Claypool*, *2020*

Preprint

Hyoukjun Kwon, Liangzhen Lai, Tushar Krishna, and Vikas Chandra, *"HERALD: Optimizing Heterogeneous DNN Accelerators for Edge Devices"*, *Arxiv preprint*, *2019*

Services

Journal Reviewer
IEEE Computer Architecture Letters *2020*

Journal Reviewer
IEEE MICRO *2019*

Journal Reviewer
ACM Transactions on Architecture and Code Optimization (TACO) *2019, 2020*

Journal Reviewer
IEEE Transactions on Computers (TOC) *2019, 2020*

Journal Reviewer
IEEE Journal on Emerging and Selected Topics in Circuits and Systems (JETCAS) *2020*

Journal Reviewer
IEEE Transactions on Neural Networks and Learning Systems (TNNLS) *2020*

Talks in Industry/National Labs

Understanding Dataflow in DNN Accelerators and Modeling Them with MAESTRO
Sandia National Lab and Pacific Northwest National Lab (Online Talk) *Dec. 2019*

Understanding Dataflows in DNN Accelerators
Facebook, Menlo Park, California *Jun. 2019*

An Open Source Framework for Exploring Dataflow and Generating DNN Accelerators Supporting Flexible Dataflow
IBM Research, Yorktown Heights, New York *Nov. 2018*

Communication-driven Approach to Design DNN Accelerators
Western Digital, Online Talk *Sep. 2018*

Analyzing Dataflows in Accelerators
NVIDIA, Westford, Massachusetts *Aug. 2018*

Optimizing Networks-On-Chip for Deep Learning Accelerators using Micro-switches
NVIDIA, Westford, Massachusetts *Aug. 2017*

Automatic generation of low-latency networks-on-chip
Bluespec Inc., Framingham, Massachusetts *Jul. 2017*

Talks in Academia

Understanding Reuse, Performance, and Hardware Cost of DNN Accelerator Dataflows
Pohang University of Science and Technology (Postech); Online Invited Talk – AI Seminar Series *Aug. 2020*

Understanding Reuse, Performance, and Hardware Cost of DNN Accelerator Dataflows
Seoul National University, Seoul, Korea *Jan. 2020*

Modeling and Analyzing Dataflows in DNN Accelerators
Tokyo City University, 2018, Tokyo, Japan *Dec. 2018*

Enabling Rapid Design Space Exploration and Prototyping of DNN Accelerators
Massachusetts Institute of Technology (MIT), 2018, Cambridge, Massachusetts *Jul. 2018*

A Communication-driven Approach to Designing Flexible DNN Accelerators
CMU, Pittsburgh, Pennsylvania *May. 2018*

[Teaching] Designing CNN Accelerators using Bluespec System Verilog
Seoul National University (SNU), Seoul, Korea *Dec. 2017*
A three-day lecture for undergraduate students
(`https://github.com/hyoukjun/DesignCNNAccelerators`)

Light-weight and High-performance NoC for DNN accelerators
Konkuk University, Seoul, Korea *Oct. 2017*

Talks at Conferences/Workshops

Understanding the Impact of On-Chip Communication on DNN Accelerator Performance
MICRO 2019, Columbus, Ohio *Oct. 2019*

[Tutorial] Enabling Rapid Design Space Exploration and Prototyping of DNN Accelerators
HPCA 2019, Washington D.C. *Feb. 2019*
`http://synergy.ece.gatech.edu/tools/maeri/maeri-tutorial-hpca-2019/`

[Tutorial] MAERI: Enabling Rapid Design Space Exploration and Prototyping of DNN Accelerators
ISCA 2018, Los Angeles, California *Mar. 2018*
`http://synergy.ece.gatech.edu/tools/maeri/maeri_tutorial_isca2018/`

MAERI: Enabling Flexible Dataflow Mapping over DNN Accelerators via Reconfigurable Interconnects
ASPLOS 2018, Williamsburg, Virginia *Mar. 2018*

[Demo] MAESTRO: An Open-source Infrastructure for Modeling Dataflows within Deep Learning Accelerators
Corgarch 2018 (colocated with ASPLOS 2018), Williamsburg, Virginia *Mar. 2018*

[Demo] OpenSMART: An Opensource Single-cycle Multi-hop NoC Generator
SC17 OpenSuCo 2017, Denver, Colorado *Nov. 2017*

Rethinking NoCs for Spatial Neural Network Accelerators
NOCS 2017, Seoul, Korea *Oct. 2017*

Adaptive Manycore Architectures for Big Data Computing
NOCS 2017, Seoul, Korea *Oct. 2017*

Proving Flow Security of Sequential Logic via Automatically-Synthesized Relational Invariants
CSF 2017, Santa Barbara, California *Aug. 2017*

OpenSMART: Single-cycle Multi-hop NoC Generator in BSV and Chisel
ISPASS 2017, Santa Rosa, California *May. 2017*

Skills and Experiences

Software languages	C/C++, JAVA, VBA, OCaml, Scheme, Matlab, and Python
Typesetting System	Latex
Machine Learning Frameworks	PyTorch
Formal verification tools	Coq Proof Assistant and ABC
Hardware languages	Verilog, Chisel, System C, and Bluespec System Verilog
ASIC synthesis tools	Synopsys Design Compiler and Cadence Encounter(Innovus)
FPGA synthesis tools	Xilinx Vivado and Altera Quartus
Parallel programming	OpenMP, OpenCL, and MPI
Image processing	OpenCV and Matlab
Languages	Korean (Native), English (Fluent), and Japanese (Proficient)