

Jaewon Lee

Facebook HQ
1 Hacker Way
Menlo Park, CA 94025

jaewon@fb.com
+1-650-665-0648
hpcs.snu.ac.kr/~jaewonlee

Research Interests

Improving Systems and Workloads with Modeling, Analysis, and Optimization

The scope includes (but is not limited to)

- ♦ Large-Scale Systems and Emerging Workloads (Datacenters / AI)
- ♦ Microarchitecture (Multi-/Many-core Processors)
- ♦ Accelerators (FPGAs/GPUs)
- ♦ Emerging Device Technologies (NVMs/Interconnects)

Work Experiences

Summer 2019 – **Facebook** Menlo Park, CA, US
Current Research Scientist

AI System Software/Hardware Co-design Team

- ♦ Hardware/Software co-optimization of machine learning applications

Spring 2018 – **High Performance Computer System Lab, Seoul National University** Seoul, Korea
Spring 2019 Postdoctoral Researcher

Advisor: Prof. Jangwoo Kim

- ♦ Developing methods to efficiently evaluate large-scale systems such as datacenters (*Under review*)
- ♦ Analyzing the computational behaviors of emerging machine learning models (ISCA '19, MnnFast)
- ♦ Addressing scalability and flexibility issues of device-centric server architectures (*Under review*)
- ♦ Expanded a graph-based single-core processor design space exploration methodology to effectively handle complex multi-core behaviors such as resource sharing and dynamic scheduling (MICRO '18, RpStacks-MT)
- ♦ Evaluated the performance of a device-centric system to effectively support emerging devices in large scales (ISCA '18, DCS-ctrl)

Fall 2012 – **High Performance Computing Lab, POSTECH** Pohang, Korea
Spring 2018 Ph.D. candidate

Advisor: Prof. Jangwoo Kim and Prof. Jong Kim

- ♦ Designed a framework to systematically diagnose the functionality of complex microarchitecture simulators (TACO '18, DiagSim)
- ♦ Proposed a fast and efficient first-order performance model for multi-core processors to facilitate evaluating in-development systems with full-system workloads (ISPASS '17, StressRight)
- ♦ Designed and evaluated a system to effectively utilize future emerging devices such as NVM storages, accelerators, and fast interconnects (MICRO '15, DCS)
- ♦ Proposed a graph-based detailed processor performance model and a novel interpretation method of the model to accelerate processor design space exploration (MICRO '14, RpStacks)
- ♦ Designed and evaluated a new memory architecture for improving GPU performance, utilization, and programmability (HPCA '14 GPUdmm, CAL '13 ScaleGPU)

Fall 2017 **Google** Sunnyvale, CA, US

Software Engineering Intern, Platforms Group

Mentors: Dr. Changkyu Kim, Dr. Liqun Cheng, and Dr. Kun Lin

- ♦ Developed a technique to efficiently capture and reproduce datacenter-scale performance behaviors to facilitate introducing new platforms to Google's fleet

-
- Summer 2016 **Google** Mountain View, CA, US
 Software Engineering Intern, Platforms Group
 Mentors: Dr. Changkyu Kim and Dr. Liqun Cheng
- ♦ Proposed a method to efficiently evaluate the comprehensive performance of next-generation platforms deployed at large scale, using a variety of benchmarks and live testing environments (ASPLOS '18, WSMeter)
- Fall 2013 **Cloud Infrastructure Lab, Korea Telecom** Daejeon, Korea
 Academia-Industry Joint Research Project
- ♦ Designed and evaluated a highly available and cost-effective storage system using OpenStack Swift object storage and cost-effective storage media (e.g., LTO tape) (CCGrid '16, DTStorage)
 - ♦ Domestic patent pending for the proposed online storage management scheme (No.10-2014-00005276)
- Fall 2012 **Microsoft Research – Silicon Valley** Mountain View, CA, US
 Research Intern
 Mentors: Dr. John D. Davis and Dr. Eric S. Chung
- ♦ Contributed to the development of LINQits, a hardware and software accelerator platform for domain specific languages such as C# LINQ, by evaluating the performance of the prototype implemented on Xilinx ZYNQ FPGA (ISCA'13, LINQits)
 - ♦ Implemented benchmarks for evaluating LINQits platform and its counterparts, using C# LINQ and C/C++
- Spring 2011 – **High Performance Computing Lab, POSTECH** Pohang, Korea
 Spring 2012
 Advisor: Prof. Jangwoo Kim
 Undergraduate Researcher
- ♦ Reviewed various FPGA-accelerated full-system and architecture simulators such as Protoflex, RAMP Gold, FAST, and HAsim
 - ♦ Investigated the architectural characteristics and applications of FPGA (e.g., CoRAM and RAMP)

Education

- Spring 2018 – **High Performance Computer System Lab, Seoul National University** Seoul, Korea
 Spring 2019
 Postdoctoral Researcher
 Advisor: Prof. Jangwoo Kim
- Fall 2012 – **High Performance Computing Lab, POSTECH** Pohang, Korea
 Spring 2018
 Ph.D. Candidate (GPA: 4.18 / 4.3)
 Advisor: Prof. Jangwoo Kim and Prof. Jong Kim
- Spring 2008 – **Pohang University of Science and Technology (POSTECH)** Pohang, Korea
 Fall 2012
 Bachelor of Science in Electrical Engineering
 Bachelor of Science in Computer Science and Engineering
 (GPA: 3.75 / 4.3, *Magna cum laude*)
- Fall 2010 **University of Minnesota – Twin Cities** Minneapolis, MN, US
 Exchange Student of College of Science and Engineering (GPA: 4.0 / 4.0)

Publications

- International Conference Hanhwi Jang, Joonsung Kim, Jae-eon Jo, **Jaewon Lee**, Jangwoo Kim, "MnnFast: A Fast and Scalable System Architecture for Memory-Augmented Neural Networks", in *Proceedings of the IEEE/ACM 46th International Symposium on Computer Architecture (ISCA)*, 2019.
- Hanhwi Jang, Jae-eon Jo, **Jaewon Lee**, Jangwoo Kim, "RpStacks-MT: A High-throughput Multi-core Processor Design Evaluation Methodology", in *Proceedings of the IEEE/ACM 51st International Symposium on Microarchitecture (MICRO)*, 2018
- Dongup Kwon, Jaehyung Ahn, Dongju Chae, Mohammadamin Ajdari, **Jaewon Lee**, Suheon Bae, Youngsok Kim, Jangwoo Kim, "DCS-ctrl: A Fast and Flexible Device-control Mechanism for Device-Centric Server Architecture", in *Proceedings of the ACM/IEEE 45th International Symposium on Computer Architecture (ISCA)*, Los Angeles, California, US, 2018
- Jaewon Lee**, Changkyu Kim, Kun Lin, Liqun Cheng, Rama Govindaraju, Jangwoo Kim, "WSMeter: A Performance Evaluation Methodology for Google's Production Warehouse-scale Computers", in *Proceedings of the 23rd ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Williamsburg, Virginia, US, 2018
- Jaewon Lee**, Hanhwi Jang, Jae-eon Jo, Gyu-hyeon Lee, Jangwoo Kim, "StressRight: Finding the Right Stress for Accurate In-development System Evaluation", in *Proceedings of the 2017 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, Santa Rosa, California, US, 2017
- Jaewon Lee**, Jaehyung Ahn, Choongul Park, Jangwoo Kim, "DTStorage: Dynamic Tape-based Storage for Cost-effective and Highly-available Streaming Service", *Proceedings of the 16th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*, Cartagena, Colombia, 2016
- Jaehyung Ahn, Dongup Kwon, Youngsok Kim, Mohammadamin Ajdari, **Jaewon Lee**, and Jangwoo Kim, "DCS: A Fast and Scalable Device-Centric Server Architecture", *Proceedings of the 48th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Honolulu, HI, US, 2015
- Jaewon Lee**, Hanhwi Jang, Jangwoo Kim, "RpStacks: Fast and Accurate Processor Design Space Exploration Using Representative Stall-Event Stacks," *Proceedings of the 47th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Cambridge, UK, 2014
- Youngsok Kim, **Jaewon Lee**, Jae-Eon Jo, Jangwoo Kim, "GPUdmm: A High-Performance and Memory-Oblivious GPU Architecture Using Dynamic Memory Management," *Proceedings of the 20th IEEE International Symposium on High Performance Computer Architecture (HPCA)*, Orlando, FL, US, 2014
- Eric S. Chung, John D. Davis, **Jaewon Lee**, "LINQits: Big data on Little Clients," *Proceedings of the ACM/IEEE 40th International Symposium on Computer Architecture (ISCA)*, Tel-Aviv, Israel, 2013
- International Journal Jae-eon Jo, Gyu-hyeon Lee, Hanhwi Jang, **Jaewon Lee**, Mohammadamin Ajdari, Jangwoo Kim, "DiagSim: Systematically Diagnosing Simulators for Healthy Simulations", in *ACM Transactions on Architecture and Code Optimization (TACO)*, 2018
- Youngsok Kim, **Jaewon Lee**, Donggyu Kim, Jangwoo Kim, "ScaleGPU: GPU Architecture for Memory-Unaware GPU Programming," *IEEE Computer Architecture Letters (CAL)*, 2013

Honors and Awards

- Qualcomm **Qualcomm Innovation Award (2014)**
- ◆ Scholarship awarded to graduate students pursuing innovative future technologies (received USD \$10,000)
 - ◆ Selected as the best presenter

- National Research Foundation of Korea **Global Ph.D. Fellowship (2012 - 2017)**
- ◆ National scholarship awarded to top Ph.D. candidates in Korea (awarded to 38 students (top 9%) in 2012 Fall)
 - ◆ Covers full tuition and stipend

- Korea Student Aid Foundation **National Science & Technology Scholarship (2008 – 2012)**
- ◆ Awarded to top 0.5% of undergraduate students in Korea
 - ◆ Covers full tuition

- POSTECH **Scholarship for the 1st rank in Electrical Engineering Department (2011)**

Dean's List in Electrical Engineering Department
(Fall 2008, Spring 2009, Spring 2010, Spring 2011, Fall 2011)

Teaching Experience

- POSTECH Teaching Assistant for Operating System (2015, received Best TA award)
Teaching Assistant for Computer Architecture (2013)

Technical Skills

Languages Python, Bash, C/C++, Verilog, MATLAB, C#, Labview

Tools Xilinx ISE/EDK (Vivado)

Relevant Coursework

POSTECH Advanced/Parallel Computer Architecture, Advanced Operating Systems, Compiler, Embedded System Architecture, Storage Virtualization, Cloud Computing Systems, Microprocessor Programming, Digital System Design, Advanced Programming, Algorithms, Data Structures, Discrete Optimization, Queueing Theory, Software Defined Networking

University of Minnesota Computer Architecture and Machine Organization, Digital Signal Processing, Linear System