

JUNGI JEONG

Post Doctoral Research Associate
Department of Computer Science, Purdue University
305 N. University Street, West Lafayette, IN 47907
Office: LWSN 3154C

Email: jungijeong@purdue.edu
Web: <https://jungijeong.github.io>

RESEARCH INTERESTS

My primary area of interest is computer architecture. Throughput my graduate studies, I have explored how computer architecture can support crash consistency and transactions in **persistent memory**. However, my interests are not limited to persistent memory but span across CPU micro-architecture and virtual memory. In particular, I am interested in researching **next-generation processors** with persistent memory and how to run important workloads (e.g., machine learning) on top of them with HW-SW codesign.

Keywords: computer architecture, compilers, persistent memory, CPU micro-architecture

EMPLOYMENT

Purdue University, IN, USA March 2020–Present
Post Doctoral Research Associate
Host: Professor Changhee Jung

EDUCATION

KAIST, Daejeon, South Korea March 2015–February 2020
Ph.D. in School of Computing
Advisor: Professor Seungryoul Maeng and Professor Youngjin Kwon

KAIST, Daejeon, South Korea March 2013–February 2015
M.S. in Computer Science
Advisor: Professor Seungryoul Maeng

Hanyang University, Seoul, South Korea March 2006–February 2013
B.S. in Computer Science and Engineering

PUBLICATIONS

Conference Proceedings

- [1] **Jungi Jeong** and Changhee Jung. **PMEM-Spec: Persistent Memory Speculation (Strict Persistency Can Trump Relaxed Persistency)**. In *Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, April 2021.
- [2] **Jungi Jeong**, Jaewan Hong, Seungryoul Maeng, Changhee Jung, and Youngjin Kwon. **Unbounded Hardware Transactional Memory for a Hybrid DRAM/NVM Memory System**. In *Proceedings of the International Symposium on Microarchitecture (MICRO)*, October 2020.
- [3] Wonsang Kwak, **Jungi Jeong**, Ganguk Lee, Jin-soo Kim, and Seungryoul Maeng. **GPU-delegated Direct I/O Framework Between GPU and NVMe-SSD**. In *Korea Computer Congress (KCC)*, July 2019.

- [4] **Jungi Jeong**, Chang Hyun Park, Jaehyuk Huh, and Seungryoul Maeng. **Efficient Hardware-assisted Logging with Asynchronous and Direct Update for Persistent Memory**. In *Proceedings of the International Symposium on Microarchitecture (MICRO)*, October 2018.
- [5] Chang Hyun Park, Taekyung Heo, **Jungi Jeong**, and Jaehyuk Huh. **Hybrid TLB Coalescing: Improving TLB Translation Coverage under Diverse Fragmented Memory Allocations**. In *Proceedings of the International Symposium on Computer Architecture (ISCA)*, June 2017.
- [6] **Jungi Jeong**, Daewoo Lee, and Seungryoul Maeng. **Application-assisted Writeback for Hadoop Clusters**. In *Proceedings of the International Conference on Cluster Computing (Cluster)*, September 2016.

Journal Articles

- [7] **Jungi Jeong** and Youngjin Kwon. **Supporting Atomic Durability in Persistent Transactions with Hardware Supports**. *KIISE Journal, Invited article*, June 2019.
- [8] Moon Hwan Lee, Yeakyung Row, Oo Sung Son, Uichin Lee, JaeJeong Kim, **Jungi Jeong**, Seungryoul Maeng, and Tek-Jin Nam. **Flower-Pop: Facilitating Casual Group Conversations with Multiple Mobile Devices**. *Wearable and Ubiquitous Technologies (IMWUT)*, December 2017.

RESEARCH PROJECTS

Software-transparent whole-system persistence through architecture/compiler codesign

Proposed *CARP*, providing unmodified programs failure-recovery.

The CARP compiler determines recovery points (e.g., regions), while CARP architecture executes regions in a failure-atomic fashion.

Participating as **the first author** leading the project (**Ongoing work**).

Soft error resilience via selective data duplication

Proposed *iMask*, selectively replicating data based on their influence.

iMask demonstrated that duplicating a carefully chosen minority of stores (14.7%) can recover a majority of soft errors (93.8%).

Participating as **the first author** leading the project (**Ongoing work**).

HW/SW codesign for persist ordering in persistent memory

Proposed *PMEM-Spec*, speculating all PM accesses executed in the right-order and removing all complexities in the core and caches.

Handled the ordering violation (e.g., misspeculation) with software failure-atomicity solutions.

Participated as **the first author** leading the project.

Published and presented at ASPLOS 2021.

Designing unbounded hardware transactional memory for hybrid memory systems

Proposed *UHTM*, unbounded hardware transactional memory for DRAM and NVM that eases the programming and accelerates transactions.

Increased concurrency with the novel hybrid conflict detection scheme and throughput with different logging for DRAM and NVM data.

Participated as **the first author** leading the project.

Published and presented at MICRO 2020.

Designing a HW logging mechanism that supports atomicity and durability of NVM writes

Proposed *ReDU*, redo-based hardware logging which utilized a DRAM cache to reduce latency and saving bandwidth, leading to outperform previous designs.

Participated as **the first author** leading the project.

Published and presented at MICRO 2018.

Providing an adjustable TLB coverage for diverse memory mappings

Non-uniformity and heterogeneity (e.g., 3D-stacked, non-volatile memories) in future memory system require a TLB to perform well diverse memory mappings.

Proposed a HW-SW hybrid TLB coalescing, where the *anchor* TLB that can provide an adjustable coverage using contiguity information encoded in the page table entry by the SW.

Participated as a **collaborator** to help evaluation.

Published in ISCA 2017.

Building a high-speed direct data transfer method between GPU and NVMe-SSD

Developed a library that transfers data between NVMe-SSD and GPU memory without copying to host memory (DRAM) using the NVMe protocol and Nvidia's RDMA APIs.

Proposed new programming model that GPU applications triggers data transfer and removes host intervention.

Participated as a **collaborator** to guide the direction of the project and extended the NVMe device driver in Linux Kernel.

Published in KCC 2019.

TECHNICAL SKILLS

Programming Languages: C, C++, Python, Shell scripts, Java

Software: Linux Kernel, Gem5 simulator, Pin, git, HDFS

PATENTS

KO, 10-2106689-0000, "DATA AVAILABILITY SSD ARCHITECTURE FOR PROVIDING USER DATA PROTECTION" April 24, 2020.

KO, 10-1942663-0000, "METHOD AND SYSTEM TO IMPROVE TLB COVERAGE BY USING CHUNKS OF CONTIGUOUS MEMORY" January 21, 2019.

HONORS AND AWARDS

National Full Scholarship for Graduate Study March 2013 - February 2020

Outstanding Teaching Assistant Award at KAIST Fall, 2013

Academic Scholarship at Hanyang University Spring 2012, Fall 2012

TEACHING EXPERIENCE

Teaching Assistant at KAIST

Operating Systems for SK Hynix Winter 2018

CS211 Digital Systems Spring 2014, 2015, 2016, 2017 (**Head**), 2018

CS310 Embedded Computer Systems Fall 2013, 2015 (**Head**), 2016, 2017, 2018 (**Head**)
SEP561 Embedded ComputingSpring 2014, 2015, 2019

MENTORING EXPERIENCE

Yuchen Zhou (PhD course @ Purdue, Sep 2021 - present)
Jianping Zeng (PhD course @ Purdue, Aug 2020 - present)
Wonsang Kwak (MS course @ KAIST, Mar 2017 - Feb 2019)
Hangyeoul Park (Undergraduate course @ KAIST, Spring 2015)
Jinnapat Indrapiromkul (Undergraduate course @ KAIST, Winter 2014)

INVITED TALKS

PMEM-Spec: Persistent Memory Speculation

Non-Volatile Memory Workshop, Virtual Mar 2021
Uppsala University, Virtual Mar 2021
ASPLOS, Virtual Apr 2021

Architectural Support for Crash Consistency in Persistent Memory

HP Labs, Virtual Jan 2021
IBM Research Almaden, Virtual Sep 2021
Intel, Virtual Sep 2021
Argonne National Lab, Virtual Sep 2021

Unbounded Hardware Transactional Memory for a Hybrid DRAM/NVM Memory System

MICRO, Virtual Oct 2020
Non-Volatile Memory Workshop, Virtual Mar 2021

Efficient Hardware-assisted Logging with Asynchronous and Direct Update for Persistent Memory

MICRO, Fukuoka, Japan Oct 2018
Korea Software Congress, Pyungchang, South Korea Dec 2018
Ajou University, Suwon, South Korea Oct 2019

Application-assisted Writeback for Hadoop Clusters

IEEE Cluster, Taipei, Taiwan Sep 2016

REFERENCES

Available upon request

Last updated: September 23, 2021