MICHAEL JAEMIN KIM

website-link ◇ google-scholar-link
505 E White Street, Champaign, Illinois
michael.kim.604@gmail.com ◇ michael604@scale.snu.ac.kr

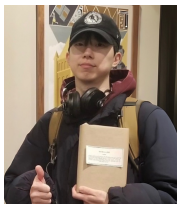

NATIONALITY

United States of America & Republic of Korea (Dual Nationality)

EDUCATION

B.S.E. in Electrical and Computer Engineering, Seoul National University — Aug 2019
Ph.D Candidate, Seoul National University (Professor Jung Ho Ahn) — Aug 2019 - Now
Visiting Research, University of Illinois Urbana-Champaign (Professor Nam Sung Kim) — Mar 2023 - Mar 2024

SUMMARY

Through my research, I have developed a strong understanding of memory systems, including trade-offs in page policies, memory technologies such as DDR/LPDDR/HBM, memory controller (MC) scheduling techniques, tiered memory management, packaging technologies, and detailed DRAM internal microarchitecture. I am also familiar with CPU/GPU/NPU microarchitectures, particularly cache hierarchies, address translation hardware/software, NoC, and interconnection networks.

I have hands-on experience evaluating computer system performance, which involves using and augmenting architectural simulators, writing microbenchmarks, and integrating benchmark suites into simulators with instrumentation and sampling tools. Additionally, I have worked with profiling tools such as Intel VTune, Intel PCM, Linux perf, and studied related theories, including interval analysis.

My research on DRAM security and reliability has also provided me with a deep understanding of streaming algorithms, ECC theory/practice in memory systems, and cryptography. My current work focuses on building cost- and energy-efficient memory systems and DRAM microarchitectures for Large Language Models running on NPU/GPU.

SKILLS

Languages C/C++, Verilog, Python

Tools Simulators (*e.g.*, Ramulator, McSimA+, GPGPU-Sim 4.0/Accel-Sim, SCALE-Sim), Intel Pin, Simpoint, Microbenchmarks (*e.g.*, stream, pointer-chase), Benchmark suites (*e.g.*, SPEC2006/2017, SPLASH-2/3, GAPBS, Redis/YCSB), Intel PCM, PMU/PEBS, VTune, Linux perf, GoogleTest, Pandas, Docker

SELECTED PUBLICATION

- [ASPLOS 2025] "M5: Mastering page migration and memory management for CXL-based tiered memory systems", Yan Sun, Jongyul Kim, Douglas Yu, Jiyuan Zhang, Siyuan Chai, **M. J. Kim**, Hwayong Nam, Jaehyun Park, Eojin Na, Yifan Yuan, Ren Wang, Jung Ho Ahn, Tianyin Xu, Nam Sung Kim
  - Designed hardware-efficient access trackers, based on SpaceSaving and CountMin-Sketch algorithms.
- [ASPLOS 2024] "AttAcc! Unleashing the Power of PIM for Batched Transformer-based Generative Model Inference", J. Park, J. Choi, K. Kyung, **M. J. Kim**, Y. Kwon, N. S. Kim, J. Ahn
  - Architectural simulations and managed overall writing efforts.
- [MICRO 2023] "How to Kill the Second Bird with One ECC: The Pursuit of Rowhammer Resilient DRAM", **M. J. Kim**, J. Park, M. Wi, S. Ko, J. Park, H. Nam, N. S. Kim, E. Lee, J. Ahn,
- [HPCA 2022] "Mithril: Cooperative Row Hammer Protection on Commodity DRAM Leveraging Managed Refresh", **M. J. Kim**, J. Park, Y. Park, W. Doh, N. Kim, T. Ham, J. Lee, and J. Ahn,

FULL PUBLICATION LISTS

- [IEEE CAL 2025] (to appear) "X-PPR: Post package repair for CXL memory",
  Chihun Song, **M. J. Kim**, Yan Sun, Houxiang Ji, Kyungsan Kim, TaeKyeong Ko, Jung Ho Ahn, and Nam Sung Kim, in IEEE Computer Architecture Letters, 2025

- [ASPLOS 2025] (to appear) "Marionette: A RowHammer Attack via Row Coupling",
  Seungmin Baek, Minbok Wi, Seonyong Park, Hwayong Nam, **M. J. Kim**, Nam Sung Kim, Jung Ho Ahn, in ACM International Conference on Architectural Support for Programming Languages and Operating Systems, 2025

  – Identified that Marionette can not only bypass isolation-based but also tracking-based software defenses, and managed the overall writing efforts.

- [ASPLOS 2025] (to appear) "M5: Mastering page migration and memory management for CXL-based tiered memory systems",
  Yan Sun, Jongyul Kim, Douglas Yu, Jiyuan Zhang, Siyuan Chai, **M. J. Kim**, Hwayong Nam, Jaehyun Park, Eojin Na, Yifan Yuan, Ren Wang, Jung Ho Ahn, Tianyin Xu, Nam Sung Kim, in ACM International Conference on Architectural Support for Programming Languages and Operating Systems, 2025

  – Designed hardware-efficient access trackers, based on SpaceSaving and CountMin-Sketch algorithms.

- [IEEE ESL 2024] "Hechi: A Hybrid Approach for Efficient Memory Reclamation Techniques in Mobile Systems",
  Wanju Doh, Seoyoung Ko, **M. J. Kim**, Jung Ho Ahn in IEEE Embedded System Letters, 2024

  – Identified the tradeoff in anon/file page reclamation on application startup time and managed overall writing efforts.

- [ISCA 2024] "DRAMScope: Uncovering DRAM Microarchitecture and Characteristics by Issuing Memory Commands",
  Hwayong Nam, Seungmin Baek, Minbok Wi, **M. J. Kim**, Jaehyun Park, Chihun Song, N. S. Kim, J. Ahn, in Proceedings of the 51st annual international symposium on computer architecture, 2024

- [ASPLOS 2024] "TAROT: A CXL SmartNIC-Based Defense Against Multi-bit Errors by Row-Hammer Attacks",
  C. Park, **M. J. Kim**, Tianchen Wang, Houxiang Ji, Jinghan Huang, Ipoom Jeong, Jaehyun Park, Hwayong Nam, Minbok Wi, J. Ahn, N. S. Kim, in ACM International Conference on Architectural Support for Programming Languages and Operating Systems, 2024

- [ASPLOS 2024] "AttAcc! Unleashing the Power of PIM for Batched Transformer-based Generative Model Inference",
  J. Park, J. Choi, K. Kyung, **M. J. Kim**, Y. Kwon, N. S. Kim, J. Ahn, in ACM International Conference on Architectural Support for Programming Languages and Operating Systems, 2024

  – Architectural simulations and managed overall writing efforts.

- [MICRO 2023] "How to Kill the Second Bird with One ECC: The Pursuit of Rowhammer Resilient DRAM",
  M. J. Kim, J. Park, M. Wi, S. Ko, J. Park, H. Nam, N. S. Kim, E. Lee, J. Ahn, in IEEE/ACM International Symposium on Microarchitecture, 2023, DOI:10.1145/3613424.3623777

- [IEEE CAL 2023] "X-ray: Discovering DRAM internal structure and error characteristics by issuing memory commands.",
  H. Nam, S. Baek, M. Wi, **M. J. Kim**, J. Park, C. Song, N. S. Kim, J. Ahn, in IEEE Computer Architecture Letters, 2023. DOI:10.1109/LCA.2023.3296153

- [HPCA 2023] "SHADOW: Preventing Row Hammer in DRAM with Intra-Subarray Row Shuffling",
  M. Wi, J. Park, **M. J. Kim**, S. Ko, N. S. Kim, E. Lee, J. Ahn, in Proceeding of 29th IEEE International Symposium on High Performance Computer Architecture, 2023. DOI:10.1109/HPCA56546.2023.10070966

- (arXiv) "AESPA: Accuracy Preserving Low-degree Polynomial Activation for Fast Private Inference", J. Park, **M. J. Kim**, W. Jung, J. Ahn, 2022. arXiv
- [IEEE TC 2022] "Future Scaling of Memory Hierarchy for Tensor Cores and Eliminating Redundant Shared Memory Traffic Using Inter-Warp Multicasting", S. Lee, S. Hwang, **M. J. Kim**, J. Choi, and J. Ahn, IEEE Transactions on Computers, Early Access, 2022. DOI:10.1109/TC.2022.3207134
- [ISCA 2022] "BTS: An Accelerator for Bootstrappable Fully Homomorphic Encryption", S. Kim, J. Kim, **M. J. Kim**, W. Jung, M. Rhu, J. Kim, J. Ahn, in Proceedings of the 49th annual international symposium on computer architecture, 2022. DOI:10.1145/3470496.3527415
- [HPCA 2022] "Mithril: Cooperative Row Hammer Protection on Commodity DRAM Leveraging Managed Refresh", **M. J. Kim**, J. Park, Y. Park, W. Doh, N. Kim, T. Ham, J. Lee, and J. Ahn, in Proceeding of 28th IEEE International Symposium on High Performance Computer Architecture, 2022. DOI:10.1109/HPCA53966.2022.00088

PAPER REVIEWS

- IEEE Transactions on Computer 2024

EXPERIENCE

- SNU MMS Lab Internship (Professor Dongsuk Jeon) Dec 2018 - Feb 2019
  - Verilog design project for spike neural network accelerator.
- SK Hynix Internship Dec 2017
  - Verilog design project for toy memory controller.
- Mandatory military service for ROK army, at Korea Military Academy Nov 2014 - Aug 2016
  - Key Resolve ROK-US Combined Exercise at Ministry of National Defense. Mar 2016