Albert Cho

100 6th St NE, Atlanta, GA 30308 | acho44@gatech.edu | U.S. Citizen

Objective

Research opportunities in Computer Architecture

Education

Georgia Institute Of Technology Atlanta GA**| PhD | 2020 ~**

- **Major**: Electrical and Computer Engineering
- **Research Interests**: Computer Architecture, Data Center Server, Memory System
- **Teaching Assistant, High Performance Computer Architecture | Fall 2021**
- **Graduate Research Assistant** (Advisor – Prof. Alexandros Daglis) **| Spring 2022~**

carnegie mellon university pittsburgh pa **| bachelor of science | may 2015**

- **Major**: Electrical and Computer Engineering

Publications

STARNUMA: MITIGATING NUMA CHALLENGES WITH MEMORY POOLING.
Albert Cho, Alexandros Daglis. In Proceedings of the 57th IEEE/ACM International Symposium on Microarchitecture (MICRO), Austin, TX, USA, 2024. **(Accepted)**

> Large multi-socket machines suffer from NUMA effects, where remote and local memory access latency can differ by 4X. StarNUMA introduces a CXL memory device pool directly accessible from all sockets in a single hop, to house shared pages and ameliorate long-latency memory accesses. StarNUMA reduces the average memory access time of a 16-socket system by 35%, yielding performance improvement of 1.13X on average, and up to 1.29X.

CXL-CENTRIC MEMORY SYSTEM FOR SCALABE SERVERS.
Albert Cho, Anish Saxena, Moinuddin Qureshi, Alexandros Daglis. In Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis (SC24), Atlanta, GA, USA, 2024. **(Accepted)**

> Queuing delay from contention at memory controller can become a performance bottleneck for bandwidth-heavy workloads. We propose *Coaxial*, a solution to mitigate the queueing delay by replacing DDR interface with CXL interface, which increases the peak bandwidth by 4X. Coaxial can improve performance of server workloads by 1.39X on average, and by up to 3X.

PATCHING UP NETWORK DATA LEAKS WITH SWEEPER.
Marina Vemmou, **Albert Cho**, Alexandros Daglis. In Proceedings of the 55th IEEE/ACM International Symposium on Microarchitecture (MICRO), Chicago, IL, USA, 2022.

> Modern NICs write/read packets directly into the LLC, bypassing memory. However, dirty packet writebacks still generate a significant amount of memory traffic, even though the evicted packets are never used again. To this end, we propose Sweeper, which marks the used packets to omit their writeback. Sweeper improves the peak sustainable throughput of a server by up to 2.6X.

Honors

CRNCH RESEARCH CENTER FELLOWSHIP | Fall 2022

- Awarded to outstanding Georgia Tech PhD students conducting research related to novel computing paradigms and hierarchies.

Experience

Graphics hardWAre engineer | intel Folsom ca **| June 2015 – Feb 2020**

- Display Microcontroller (DMC) Unit Owner | June 2015 – June 2019
  - DMC manages kernel independent low power states and context save restore, arbitrates display register cycles, and allows advanced debug capabilities
  - RTL Design / Firmware / Feature / Verification ownership across all ongoing projects
  - SoC / Silicon debug support
- Display Register Unit Owner | June 2019 – Feb 2020
  - Instantiating newly defined registers / setting up MMIO bus and register interaction
  - Validation by UVM based CLT
- Championed implementation and verification for new power management features

Skills

- Programming Languages: System Verilog, C/C++, ARM ISA, x86 ISA, Java, Python, Perl
- Application Software: gdb, git, gcc, VCS, DVE, SpyGlass, SST