Toward Rethinking the Management Techniques in Emerging Memory Systems

Maryam Babaie, Ayaz Akram, Jason Lowe-Power

DArchR Research Group, CS Department, University of California Davis

{mbabaie, yazakram, jlowepower}@ucdavis.edu

CWIDCA Workshop at MICRO-55
Chicago, IL, USA
Oct. 2022

UC DAVIS
UNIVERSITY OF CALIFORNIA

Motivation

HPC systems rely on heterogeneous memory,
e.g., *Knights Landing, Cascade Lake, Sapphire Rapids.*

We need to rethink memory management.

☹ No accurate model available for community.

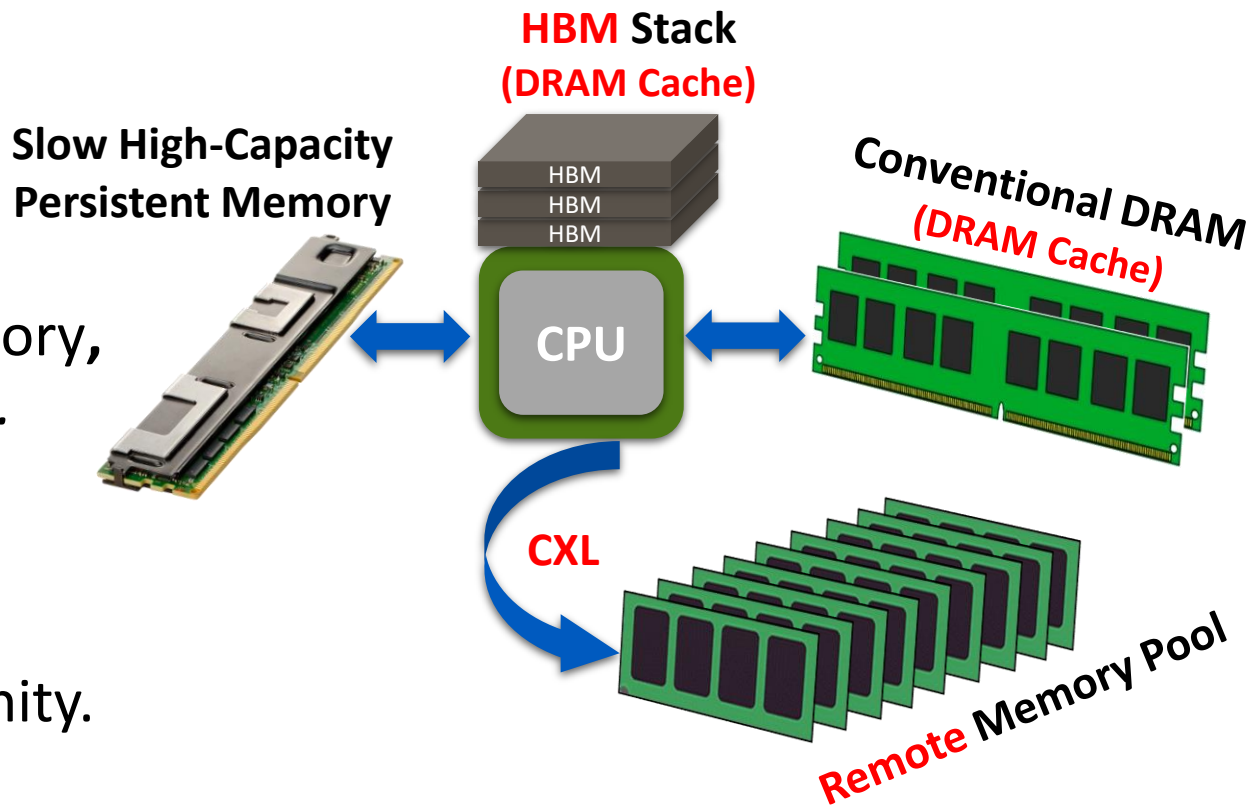

We extend *gem5* [1] with a heterogeneous/disaggregated memory system model for design space exploration of future HPC systems. We add support for:

1. DRAM Cache 2. HBM interface and controller 3. Modular memory controller design

1 J. Lowe-Power et al., "The gem5 simulator: Version 20.0+," 2020

UCDAVIS

gem5's DRAM Cache Support

A. Dedicated DRAM Cache Controller

1. Unified DRAM Cache & Memory Controller
2. Disaggregated DRAM Cache Controller

B. Decoupled DRAM Cache Policy Manager

- Supports multiple cache architectures in parallel

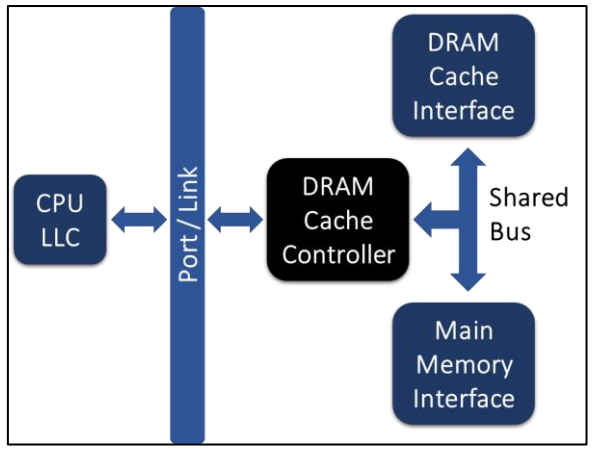

(1)

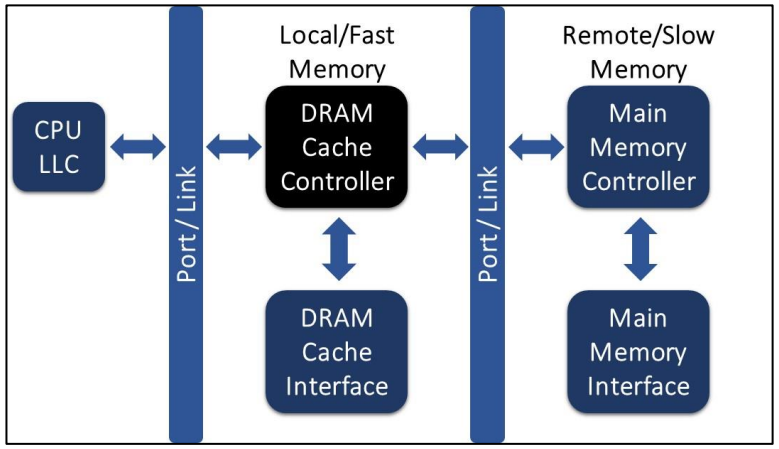

(2)

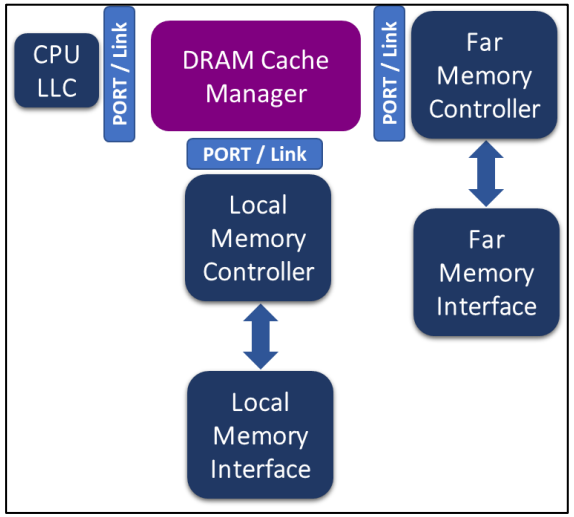

Goal: Enabling cycle-level analysis of DRAM cache designs for **data management** research

gem5's DRAM Cache Validation

Performance Validation

1. Read bandwidth compares well with the theoretical peak.

2. Access amplification is same as in real hardware [1].

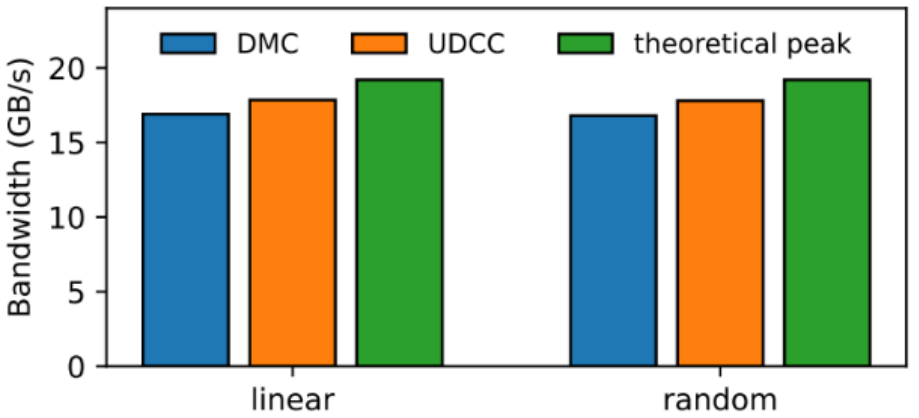

DMC: Default Memory Controller of gem5
UDCC: Unified DRAM Cache/Memory Controller

Functional Validation

- Successfully tested Linux kernel booting
- Tested GAPBS and NPB benchmark suites

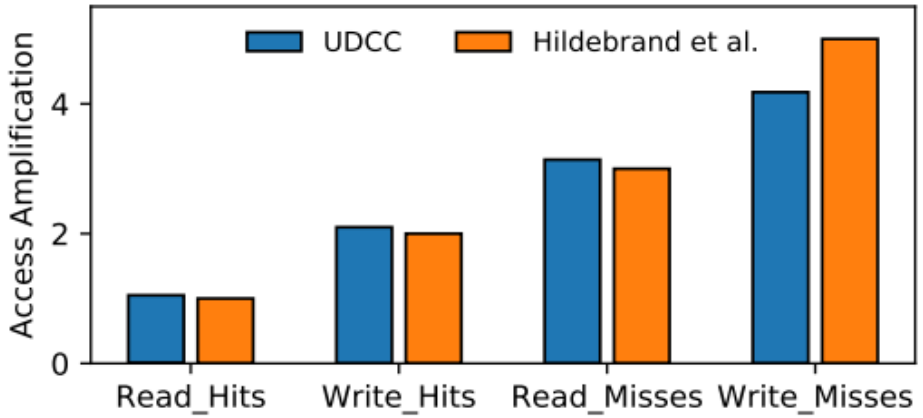

1 M. Hildebrand, J. T. Angeles, J. Lowe-Power, and V. Akella, "A case against hardware managed dram caches for nvram based systems," in 2021 ISPASS.

UCDAVIS

Thank you and see you in the poster session ☺

Questions:
Maryam Babaie
mbabaie@ucdavis.edu