A Cycle-level Unified DRAM Cache Controller Model in gem5

Maryam Babaie, Ayaz Akram, and Jason Lowe-Power

DArchR, CS Department, UC Davis

gem5 Users Workshop at ISCA 2022

NYC, NY, USA

UC DAVIS
UNIVERSITY OF CALIFORNIA

Motivation

Hardware-managed DRAM caches: *Intel's Knights Landing, Cascade Lake* and *Sapphire Rapids.*

Previous work has explored DRAM cache design in simulation to improve performance.

The microarchitecture of the cache and the main memory controllers can lead to unexpected performance pathologies.

- Hildebrand et al.[1]: "a dirty miss to the DRAM cache requires up to ***5*** accesses to memory"
  1 DRAM read, 1 NVRAM read (fetch), 1 DRAM write (fill), 1 DRAM write (data), 1 NVRAM write (dirty data)

In this work we try to build a detailed simulation model where a DDR DRAM cache + NVRAM.

UCDAVIS

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.

Unified DRAM Cache Controller (UDCC)

- Modeling Intel's Cascade Lake in 2LM (cache) mode.
  - Direct-mapped.
  - Tag and metadata stored alongside the data.
  - Inserts on miss, write backs the dirty lines.
  - Shared bus.
  - Single Channel.

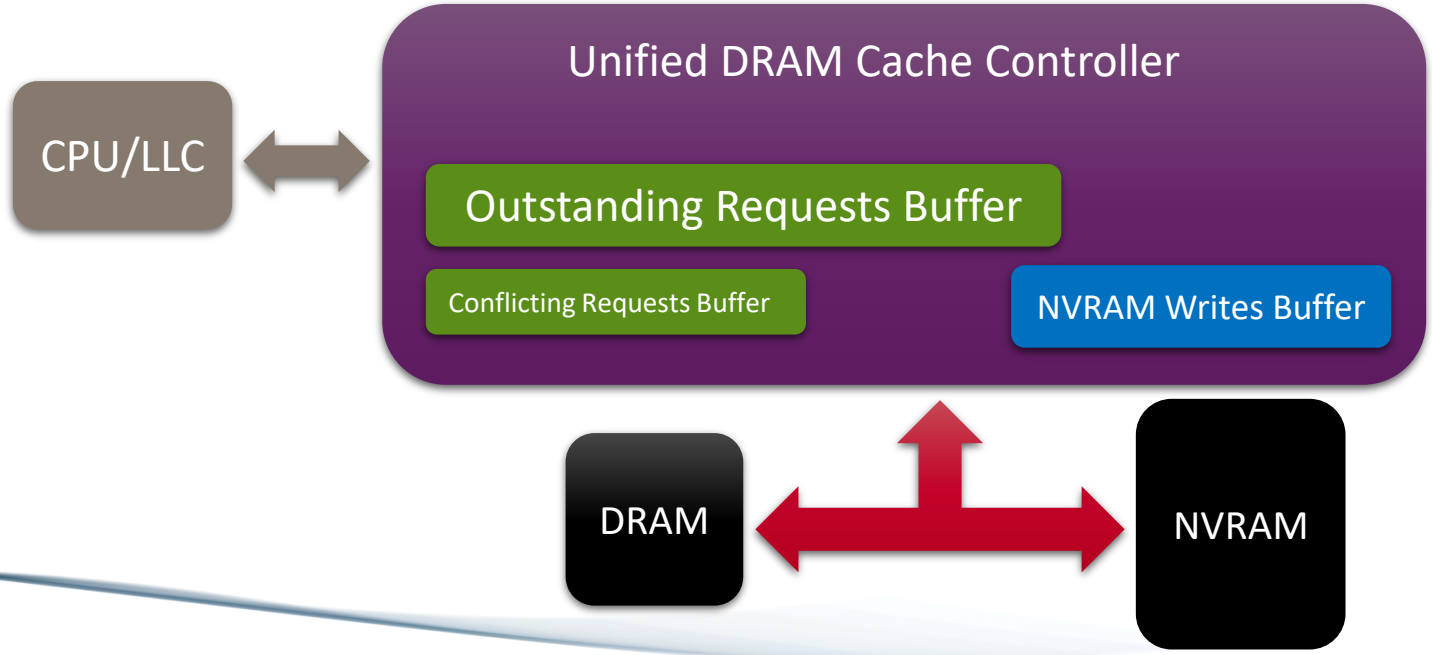

Validation

- Performance validation
  - Read bandwidth compares well with the theoretical peak.
  - Access amplification is same as in real hardware.

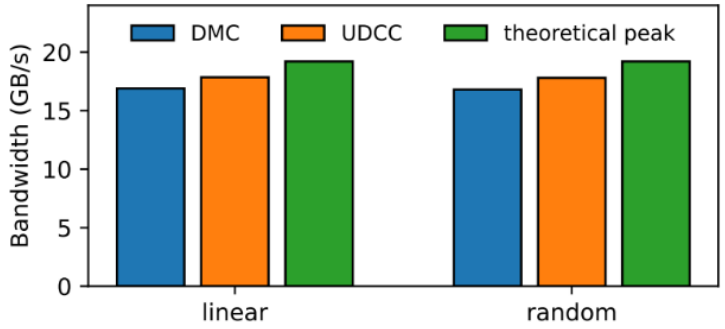

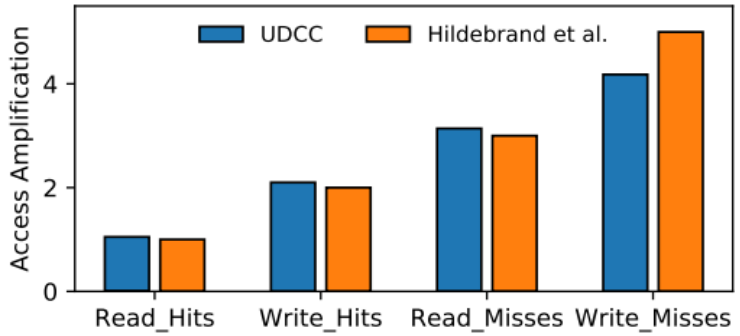

DMC: gem5's default memory controller.
UDCC: DDR4 cache with 256 entries in buffer.

- Functional validation
  - Linux kernel boot test
  - GAPBS and NPB benchmark suites

Case Studies

1. Different **DRAM Technologies for Cache**
2. NVM Wear Leveling Effect
3. Different Main Memory Technologies

Methodology

gem5's traffic generator

- Linear/Random

Read/Write Combination %

- Read-Only (RO)
- Write-Only (WO)

Miss/Hit Ratio

- DRAM cache size (128 MB)
- Range of accessed addresses

Range	6 GB	6 MB
Hit Ratio	0%	100%

gem5's Memory interfaces

- DDR3/ DDR4/ DDR5 + NVM

Other Configurations

- Request Buffer Size = 256 entries
- Block size = 64 B

UCDAVIS

Study#1: Different DRAM Technologies

What buffer size is required for each DRAM technology to fully utilize it as DRAM cache?

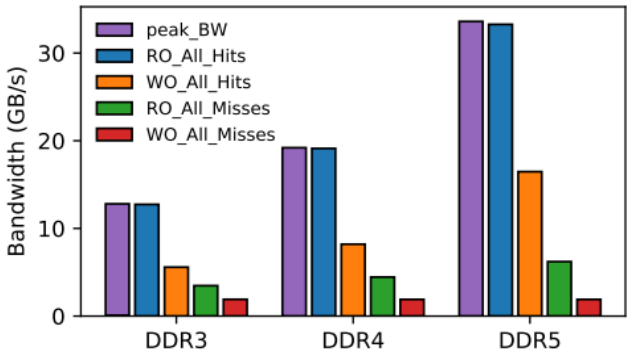

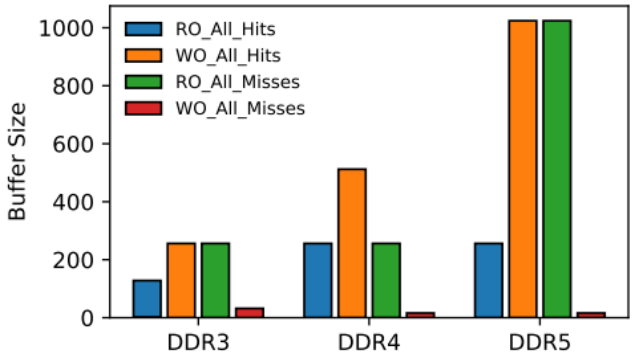

Buffer size could become large (but not impractical) for DDR memories with higher bandwidth.

Buffer size to achieve peak bandwidth depends on memory traffic.

Study #2: Impact of NVRAM Write Wear Leveling

What is the impact of NVRAM wear-leveling on the effective bandwidth?

- NVM wear leveling
  - Every 14000 write accesses the interface adds 60 µs extra latency (based on the observations by Wang et al[1])
- The access pattern
  - 100% writes
  - 100% misses
  - All lines are dirty in DRAM cache.

Parameter	With wear-leveling	Without wear-leveling
Bandwidth	1.77 GB/s	1.92 GB/s
NVM write latency	42.84 μs	39.71 μs

Wear-leveling can cause 8.5% performance degradation.

1 Z. Wang, X. Liu, J. Yang, T. Michailidis, S. Swanson, and J. Zhao, "Characterizing and modeling non-volatile memory systems," in 53rd IEEE/ACM MICRO, 2020, pp. 496–508.

Study #3: Different Main Memory Technologies

How does main memory performance impact the effective bandwidth?

RO_All_Misses & WO_All_Misses traffic patterns make "NVM" be the effective bandwidth determining factor.

Three NVM configurations

Parameter	Slow	Base	Fast
tREAD	300ns	150ns	75ns
tWRITE	1000ns	500ns	250ns
tSEND	28.32ns	14.16ns	7.08ns
tBURST	6.664ns	3.332ns	1.666ns

Access pattern requiring more interaction with NVM (e.g., WO_All_Misses) will have a bigger impact on the observed BW of DRAM cahce, when going from base configuration to slow or fast NVM configurations.

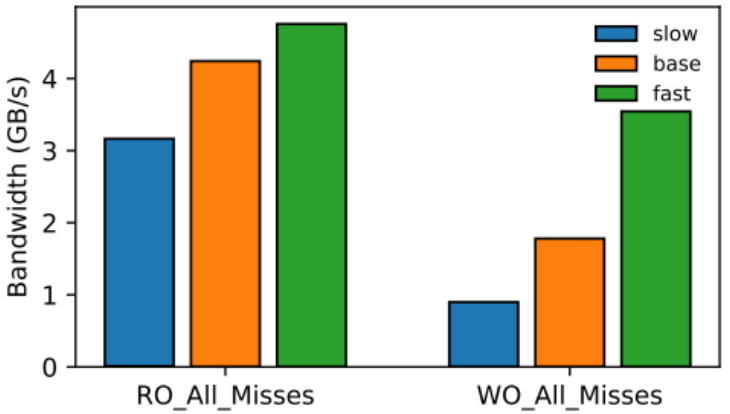

Future Work

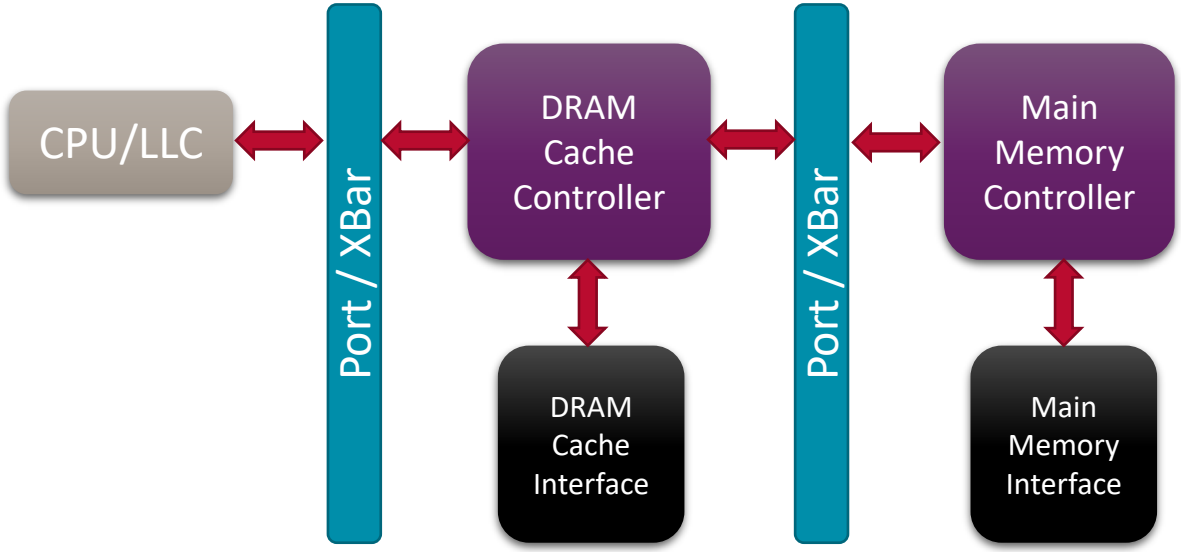

- Pluggable interfaces
  - Modularity
  - Abstracts away the timing and microarchitectural limitations of technologies
  - Enables accurate study of combination of different memory technologies as the cache and/or backing store.
- Controlling links and ports
  - Modeling CXL

Summary

- UDCC is a DRAM cache extension of gem5, modeling Intel's Cascade Lake in 2LM mode.
- UDCC enables a cycle-level analysis of DRAM caches.
- This will facilitate the research on heterogenous memory systems with DRAM caches.
- The code is available on a private repo while we're working on improvements.
  - Let us know if you need to access it before the official release.

Thank You!