Binary Star: Coordinated Reliability in Heterogeneous Memory Systems for High Performance and Scalability

Xiao Liu¹, David Roberts, Rachata Ausavarungnirun² Onur Mutlu³, Jishen Zhao¹ ¹UC San Diego, ²King Mongkut's University of Technology North Bangkok, ³ETH Zurich

52nd IEEE/ACM International Symposium on Microarchitecture[®] October 12-16, 2019, Columbus, Ohio, USA







Thai-German Graduate School of Engineering



Large memory capacity is in demand



[Ogleari+, HPCA' 2019]



52nd IEEE/ACM International Symposium on Microarchitecture®

Benefit of hybrid memory



Izraelevitz, J., Yang, J., Zhang, L., Kim, J., Liu, X., Memaripour, A., Soh, Y.J., Wang, Z., Xu, Y., Dulloor, S.R., Zhao, J. and Swanson, S. 2019. Basic performance measurements of the intel optane DC persistent memory module. *arXiv preprint arXiv:1903.05714*.



52nd IEEE/ACM International Symposium on Microarchitecture®

Reliability issues with traditional memory hierarchy scaling





Reliability issues with traditional memory hierarchy scaling





Reliability issues with traditional memory hierarchy scaling





Opportunity #1

Inefficiency of decoupled reliability schemes



Duplicated and unaware of each other



Opportunity #2

Most cache lines are clean



Errors on the clean cache line can be corrected by data in NVRAM



Binary Star: design overview



Binary Star: design overview



Binary Star: design overview



Periodic forced writeback

All SRAM caches

On-chip 3D DRAM

LLC

Observation 1: Errors in the LLC can be corrected by consistent data copies in NVRAM main memory.

CRC

Binary Star Daemon: conducts periodic forced writebacks

- Stalls the application
- Saves process state
- Issues cache-line writeback instructions

Periodic forced writeback

(Infrequent, e.g., every 30 minutes)

Consistent NVRAM block



Inconsistent NVRAM block



Off-chip NVRAM

а

Consistent cache writeback

Observation 1: Errors in the LLC can be corrected by consistent data copies in NVRAM main memory.

Observation 2: NVRAM wear leveling naturally redirects and maintains the remapping of data updates to alternative memory locations.





LLC error correction and recovery



- Simulators: McSimA+ (Performance), FaultSim (Reliability)
- Baselines:
 - 3D DRAM cache with DRAM main memory
 - DRAM main memory only
 - 3D DRAM cache with NVRAM main memory (PCM)
 - NVRAM main memory only (PCM)



- Simulators: McSimA+ (Performance), FaultSim (Reliability)
- Baselines:
 - 3D DRAM cache with DRAM main memory
 - DRAM main memory only
 - 3D DRAM cache with NVRAM main memory (PCM)
 - NVRAM main memory only (PCM)



- Resilience schemes:
 - Rank level ECC (RECC)
 - In-DRAM ECC (IECC)
 - Chipkill combined with IECC
 - Binary Star
- Workloads:
 - In-memory/traditional databases workloads (Redis, Memcached, TPCC, YCSB, mysql)
 - Memory intensive workloads from PARSEC



- Resilience schemes:
 - Rank level ECC (RECC)
 - In-DRAM ECC (IECC)
 - Chipkill combined with IECC
 - Binary Star
- Workloads:
 - In-memory/traditional databases workloads (Redis, Memcached, TPCC, YCSB, mysql)
 - Memory intensive workloads from PARSEC



Evaluation: system reliability

System	Better FIT	Storage cost	
		DRAM LLC	Main memory
No-ECC 28nm DRAM	44032-66150	N/A	0%
RECC 28nm DRAM	8806-13230	N/A	12.50%
IECC+RECC sub-20nm 3D DRAM+DRAM	78211-117912	18.75%	18.75%
IECC+Chipkill sub-20nm 3D DRAM+DRAM	37949-59518	18.75%	18.75%
RECC sub-20nm 3D DRAM+PCM	6352-9963	12.50%	12.79%
Binary Star sub-20nm 3D DRAM+PCM	2637-3968	6.25%	12.79%



Evaluation: devices' reliability

- 3D-DRAM device error rate (for sub20 nm)
 - 10¹⁶x reduction compares to RECC
 - 10¹²x reduction compares to Chipkill+IECC

- Number of writes to NVRAM
 - 20% more writes when the periodic force writeback interval is 30 mins



Evaluation: performance





Other results

- NVRAM latency sensitivity study
- Effect on varying the periodic forced writeback interval
 - Performance
 - Rollback rate
 - Number of writes



Summary





Summary





Binary Star: Coordinated Reliability in Heterogeneous Memory Systems for High Performance and Scalability

Xiao Liu¹, David Roberts, Rachata Ausavarungnirun² Onur Mutlu³, Jishen Zhao¹ ¹UC San Diego, ²King Mongkut's University of Technology North Bangkok, ³ETH Zurich

52nd IEEE/ACM International Symposium on Microarchitecture[®] October 12-16, 2019, Columbus, Ohio, USA







Thai-German Graduate School of Engineering



Binary Star triggered Wear-leveling







52nd IEEE/ACM International Symposium on Microarchitecture®

Binary Star Daemon





NVRAM latency sensitivity study





Effect on varying the periodic forced writeback interval





52nd IEEE/ACM International Symposium on Microarchitecture®

Effect on varying the periodic forced writeback interval





Evaluation: 3D-DRAM reliability

-RECC - RECC+IECC - Chipkill+IECC - Binary Star



33

Effect on varying the periodic forced writeback interval





52nd IEEE/ACM International Symposium on Microarchitecture®