Write-Rationing Garbage Collection for Hybrid Memories

Shoaib Akram (Ghent), Jennifer B. Sartor (Ghent and VUB),
Kathryn S. Mckinley (Google), and Lieven Eeckhout (Ghent)
Shoaib.Akram@UGent.be

GHENT UNIVERSITY

VUB VRIJE UNIVERSITEIT BRUSSEL

Google

DRAM is facing challenges

Scalability

Cost

Energy

Reliability

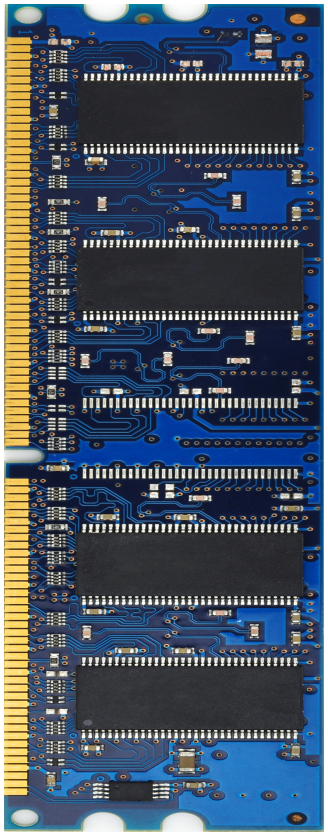

Phase change memory

Persistent

Byte addressable

High latency

Low endurance

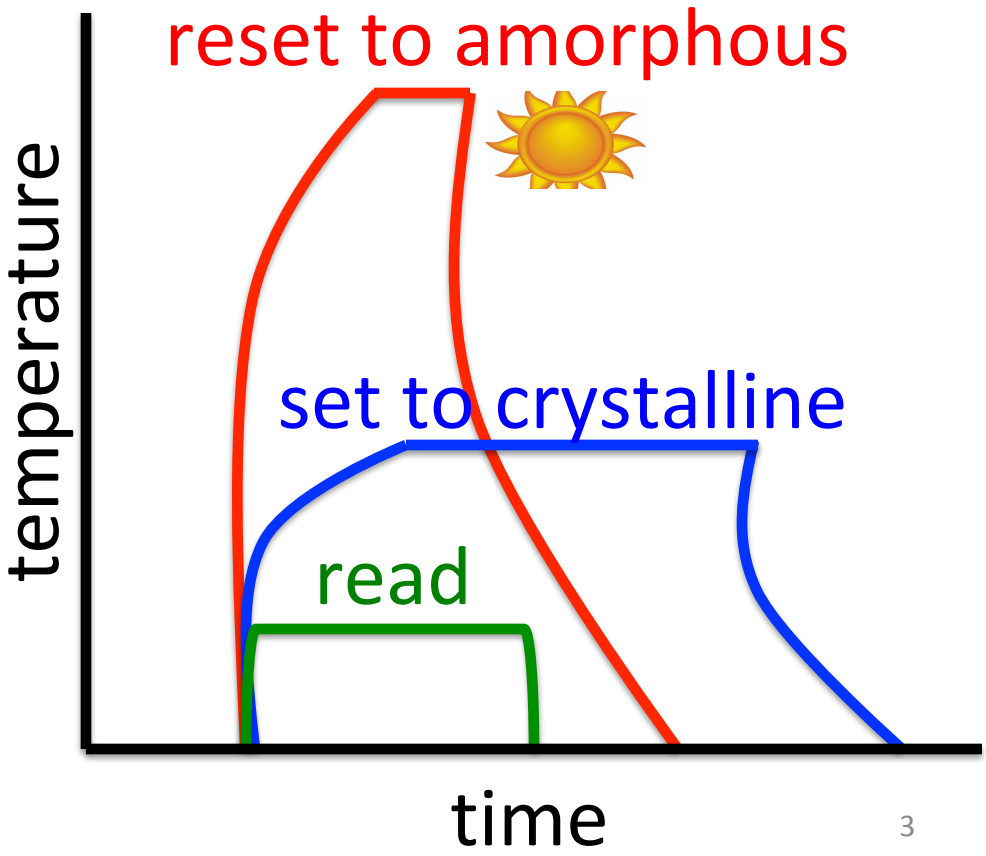

PCM only is not practical

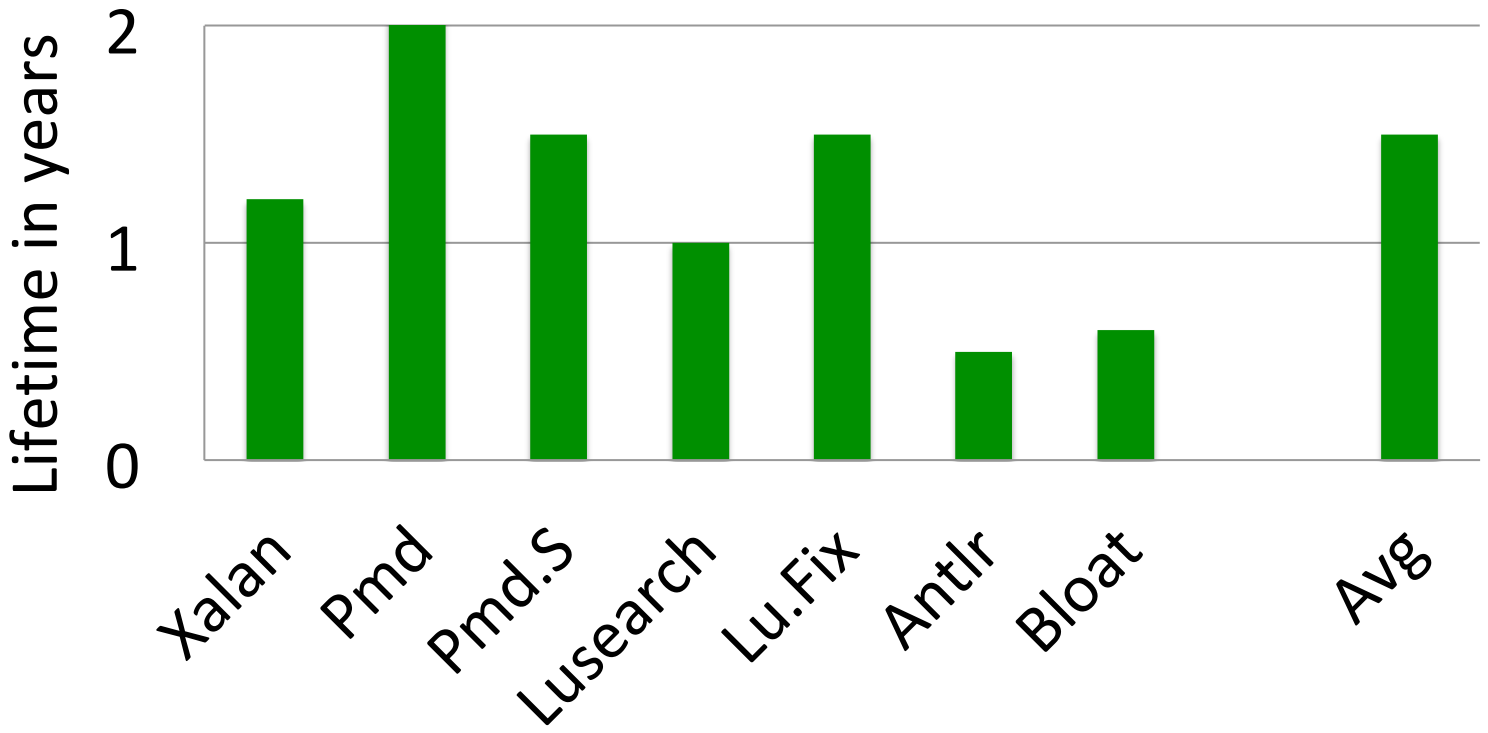

32 GB PCM with hardware wear-levelling

Hybrid DRAM-PCM memory

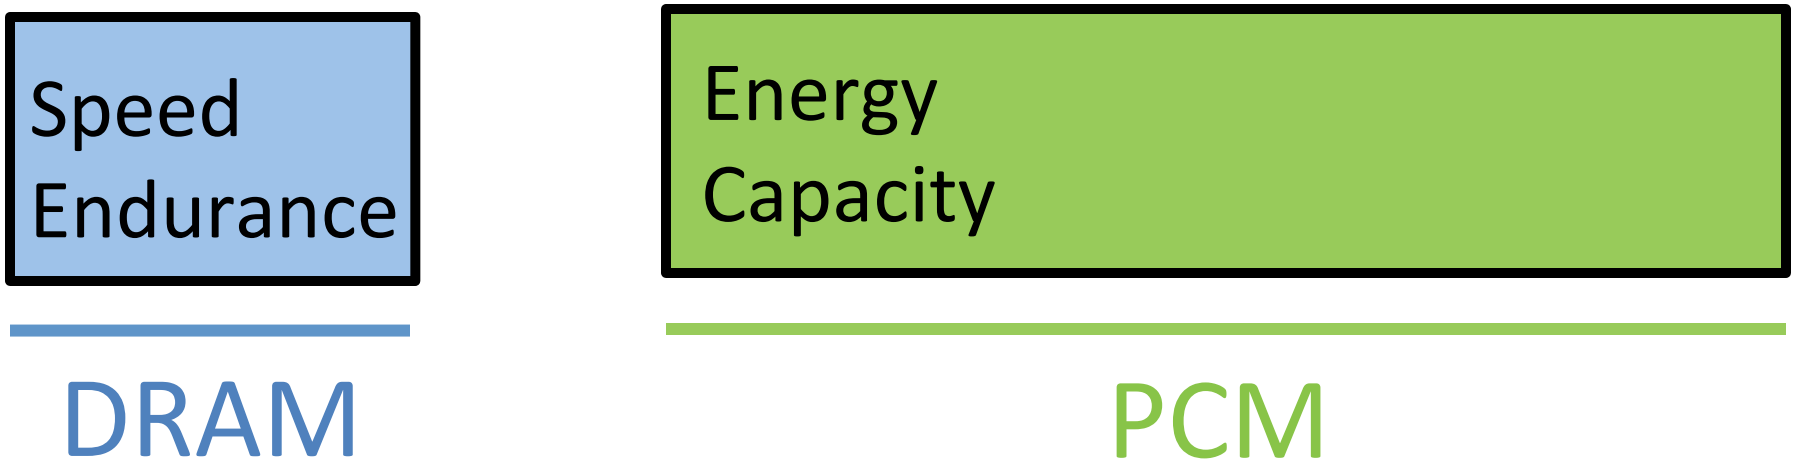

Challenges

Bridging the DRAM-PCM latency gap

Mitigating PCM wear-out

Prior art in mitigating PCM wear-out

Hardware wear-leveling

- Spread writes out across PCM
- 32 GB PCM lasts only two years!

OS write partitioning

- Keep highly written pages in DRAM
- Coarse granularity
- Costly page migrations

Garbage collection for hybrid memory

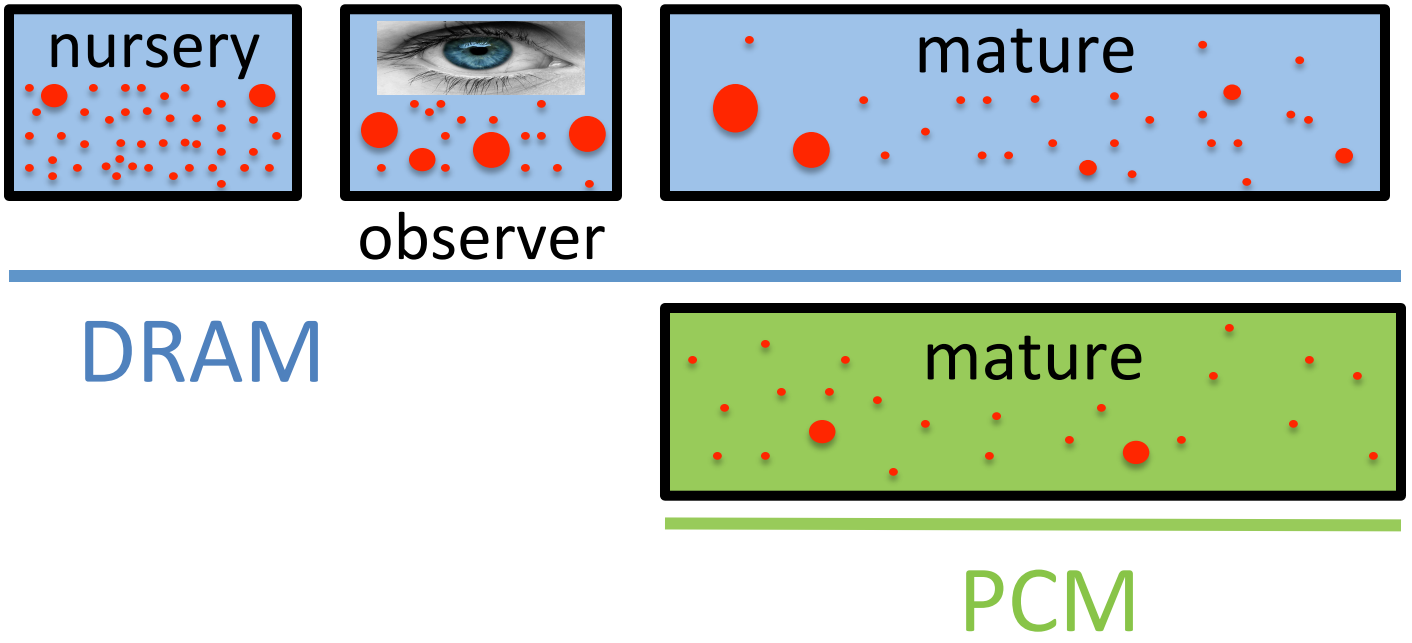

This work uses GC to keep highly written objects in DRAM

Distribution of writes in GC heaps

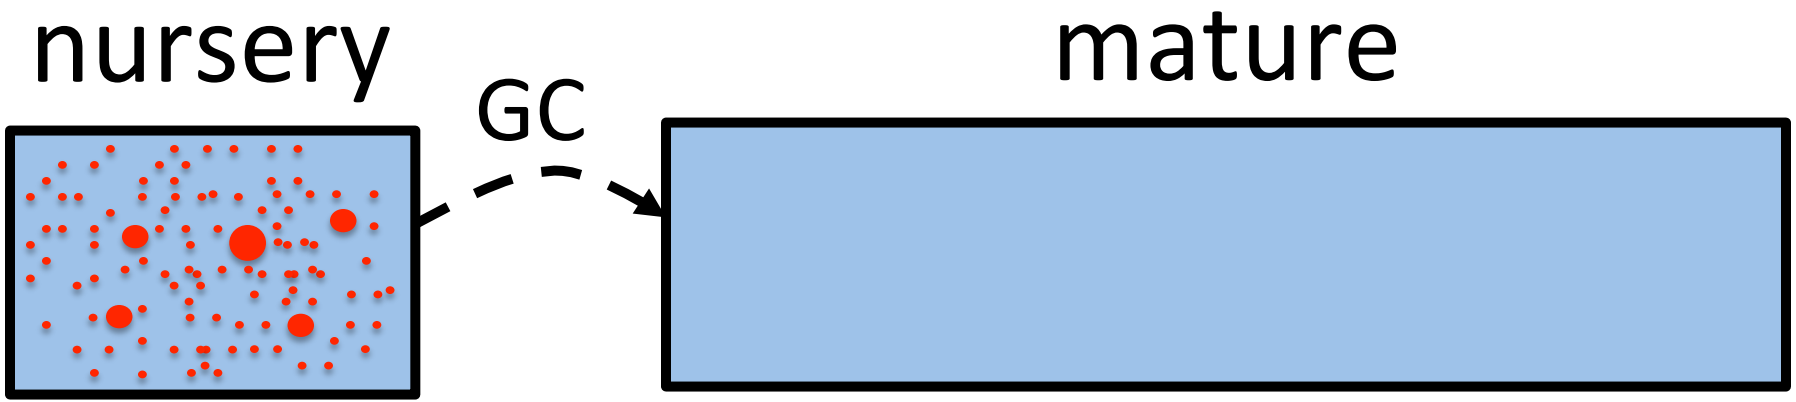

70%

of writes

Distribution of writes in GC heaps

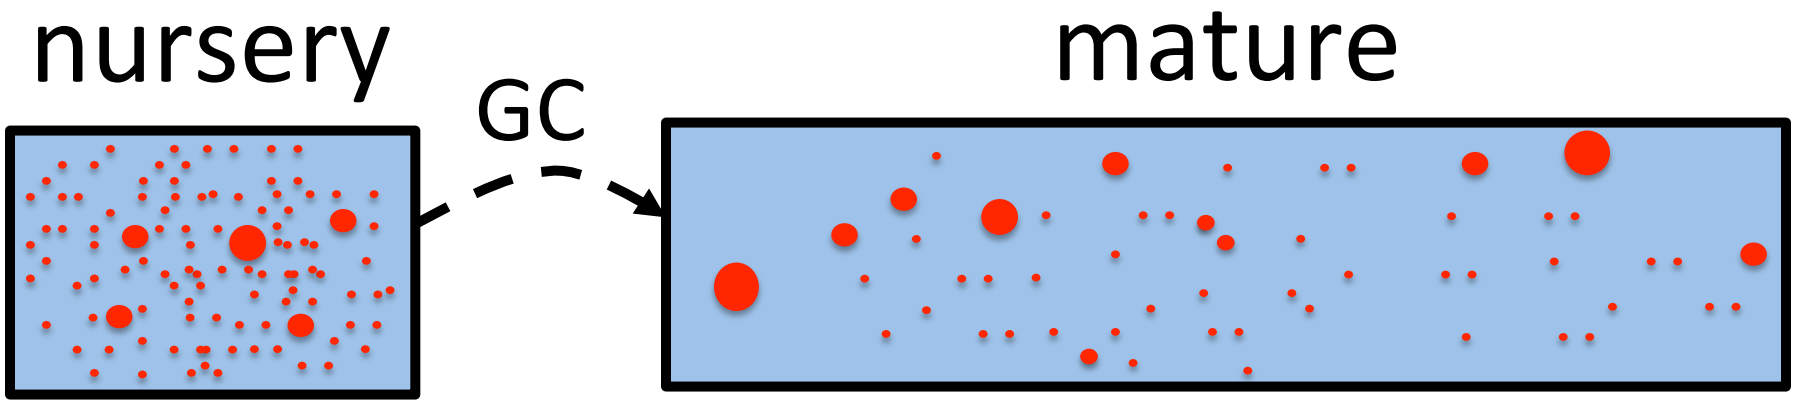

70%

of writes

22%

to 2% of objects

Contribution
Write-Rationing Garbage Collectors

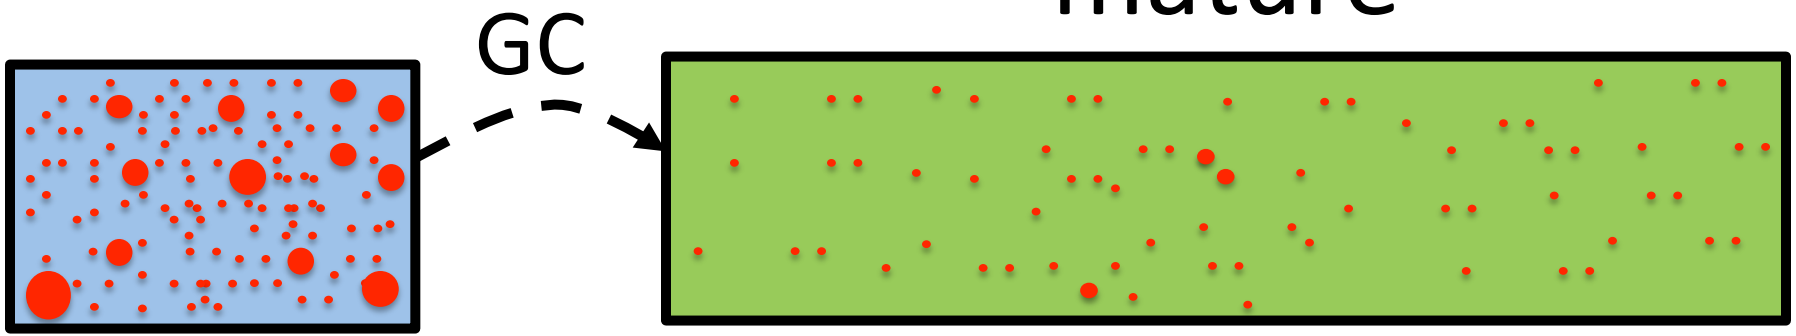

DRAM

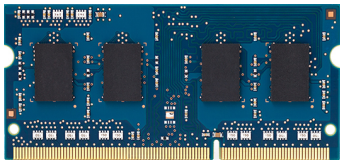

PCM

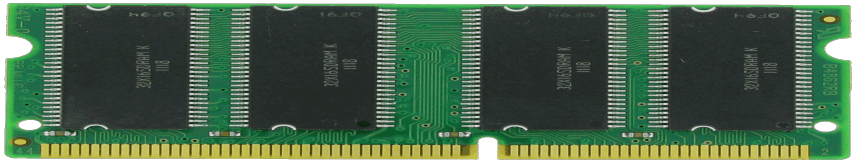

Two write-rationing garbage collectors

Kingsguard-Nursery

Kingsguard-Writers

Heap organization in DRAM

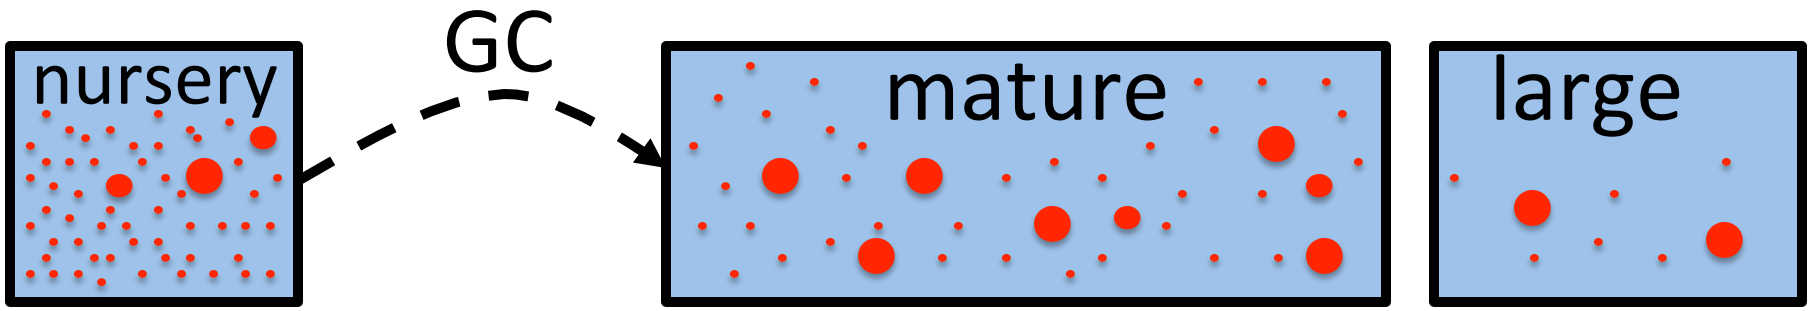

DRAM

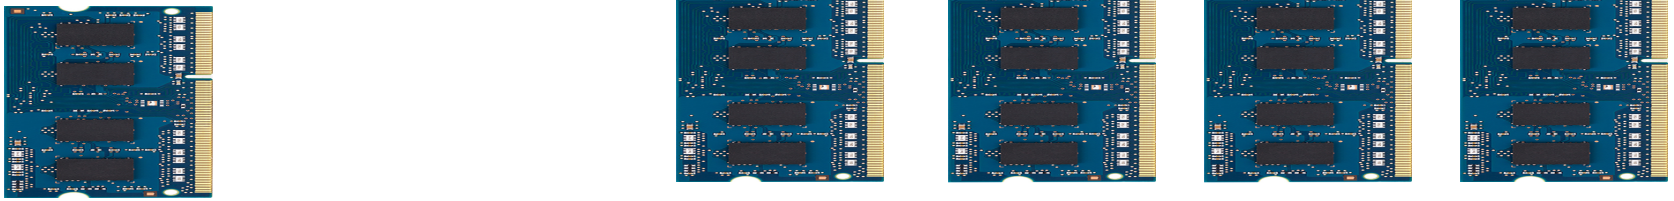

KG-N Kingsguard-Nursery

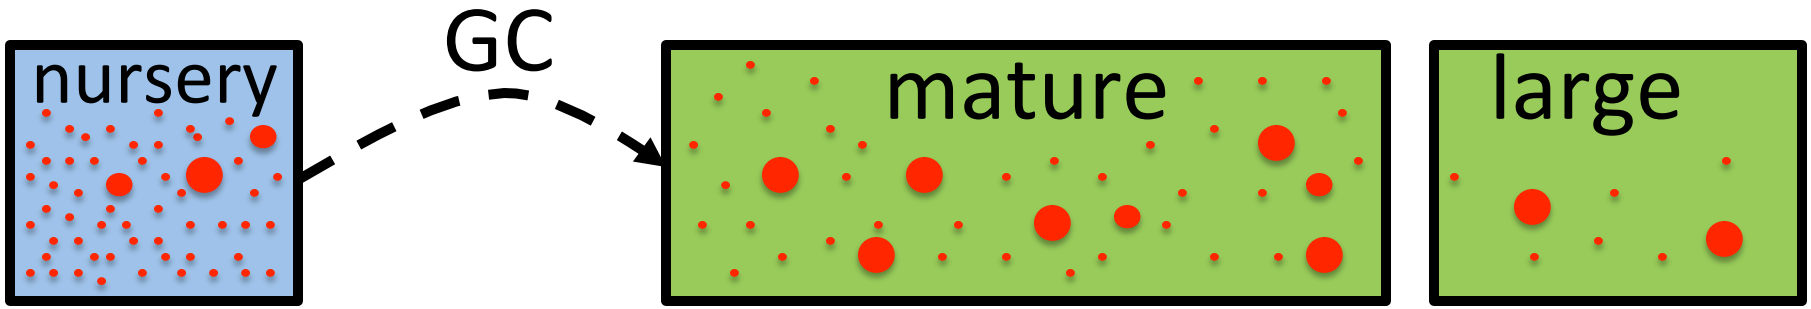

DRAM

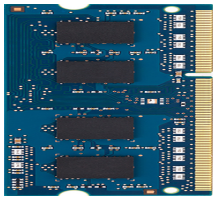

PCM

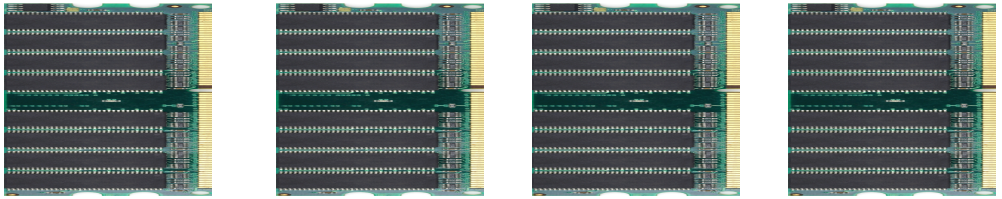

KG-W Kingsguard-Writers

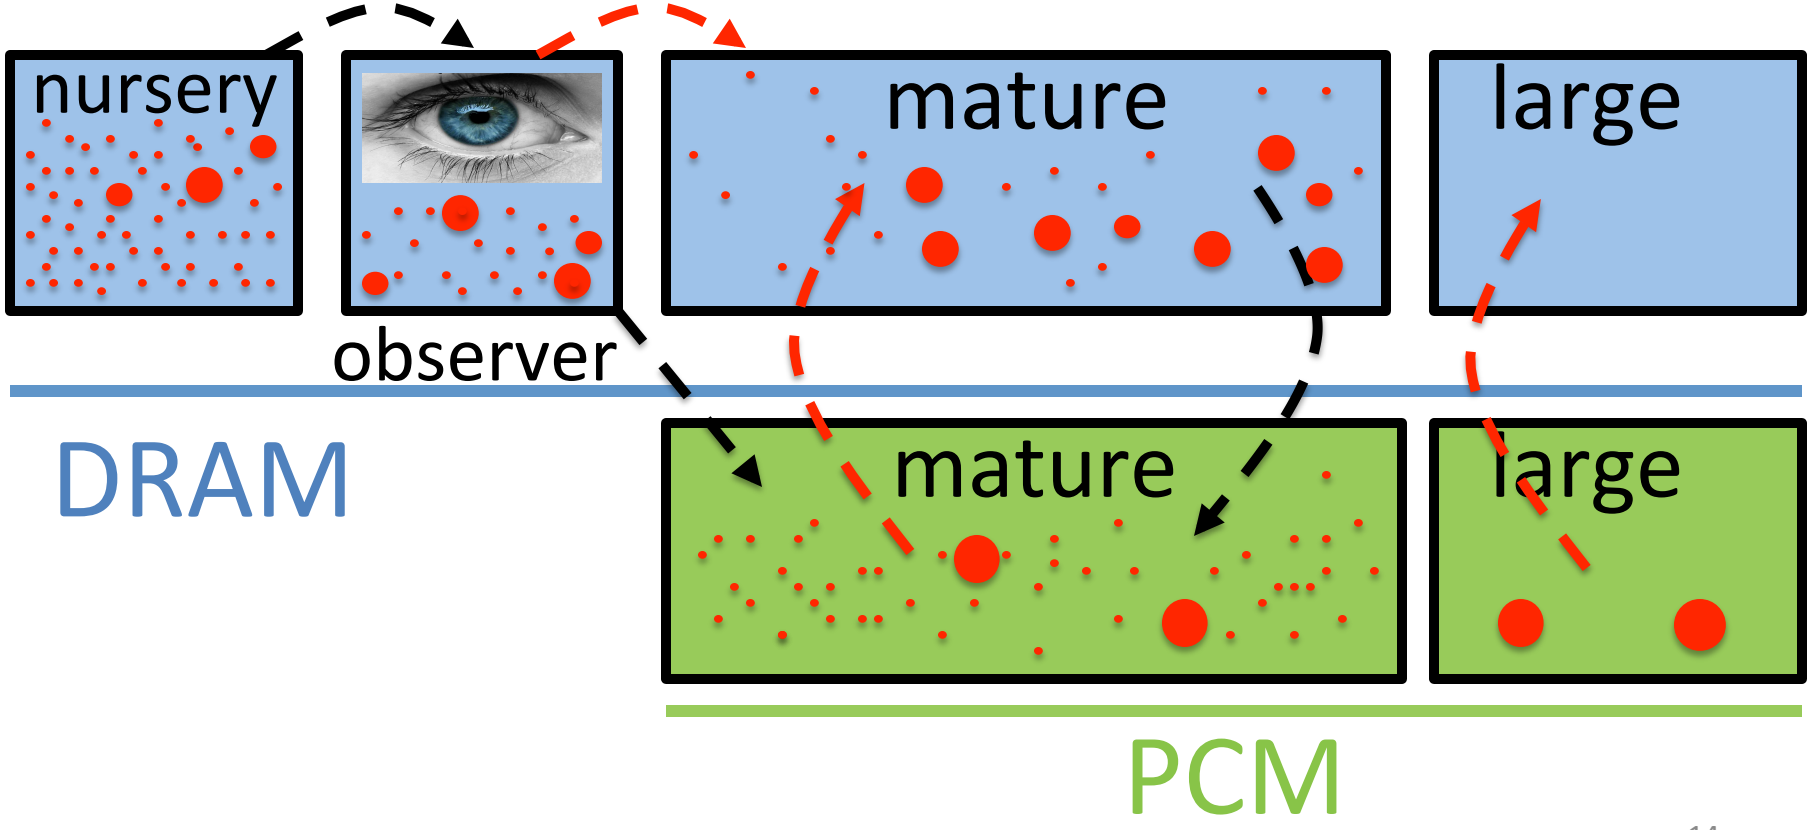

Monitoring writes

On a write to an object

Write barrier sets a bit in header

Write barrier configurations

Monitor references

Monitor references and primitives

Two additional optimizations

Large object optimization

Selectively allocate large objects in DRAM

Metadata optimization

Place mark bits of PCM objects in DRAM

Large object optimization

nursery

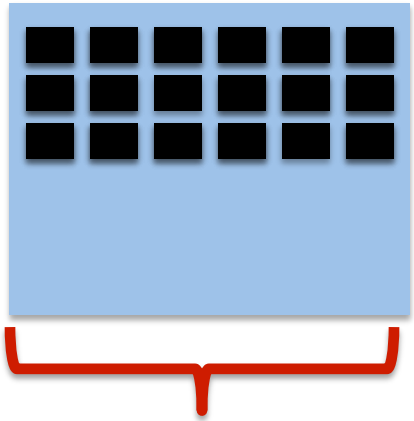

½ of remaining nursery

large

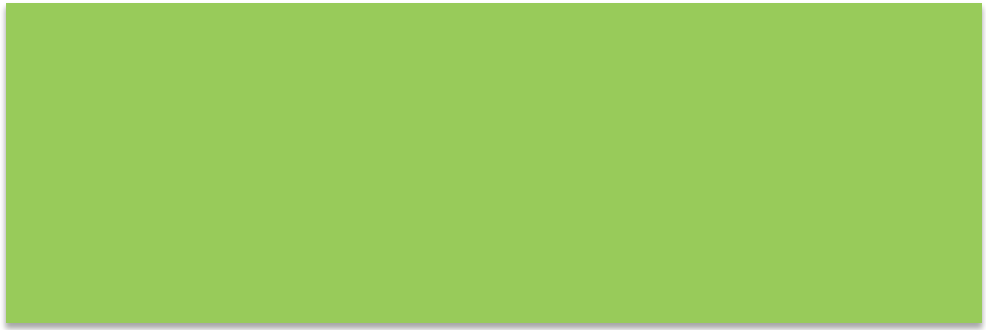

Monitor PCM write rate to tu[redacted]

Results

(1) Measurements on real hardware

(2) Simulation

Jikes research virtual machine

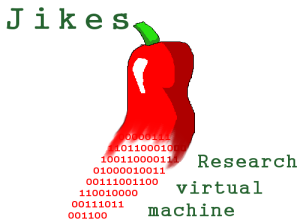

Java applications

Real hardware methodology

Use write barriers to count object writes

Applications: 12 DaCapo, 3 GraphChi, and Pjbb

Configurations

KG-N : 4 MB nursery

KG-W: 4 MB nursery, 8 MB observer

KG-N : 12 MB nursery

Reduction in PCM writes

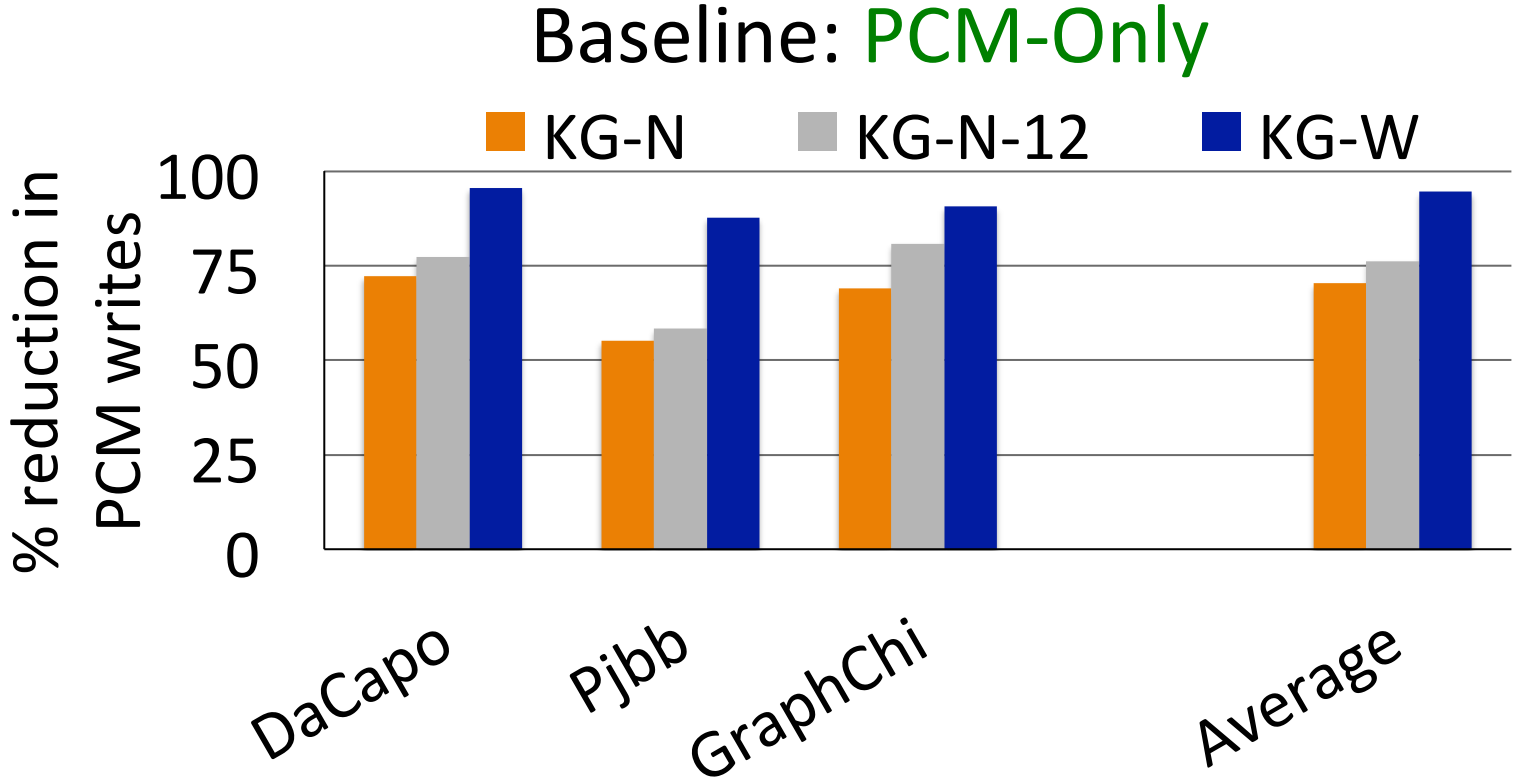

KG-W reduces 95% of writes to PCM

Simulation methodology

7 DaCapo applications

Measure lifetime, energy, and execution time in simulator

Memory systems

Homogeneous

- 32 GB DRAM
- 32 GB PCM

Hybrid

- 1 GB DRAM
- 32 GB PCM

PCM parameters

- 4X read latency
- 4X write energy
- 10 M writes/cell

PCM lifetimes

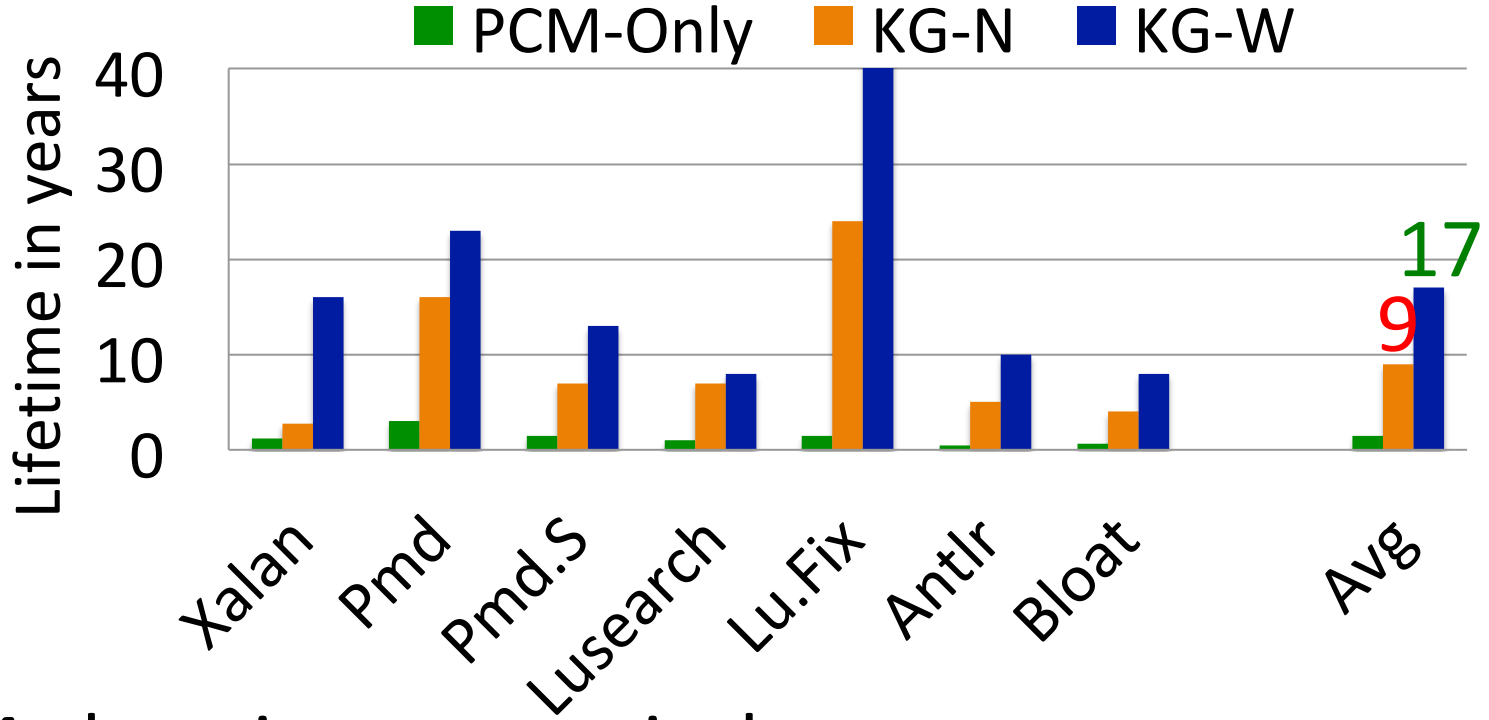

PCM alone is not practical

PCM lasts more than 10 years with KG-W

PCM write rates

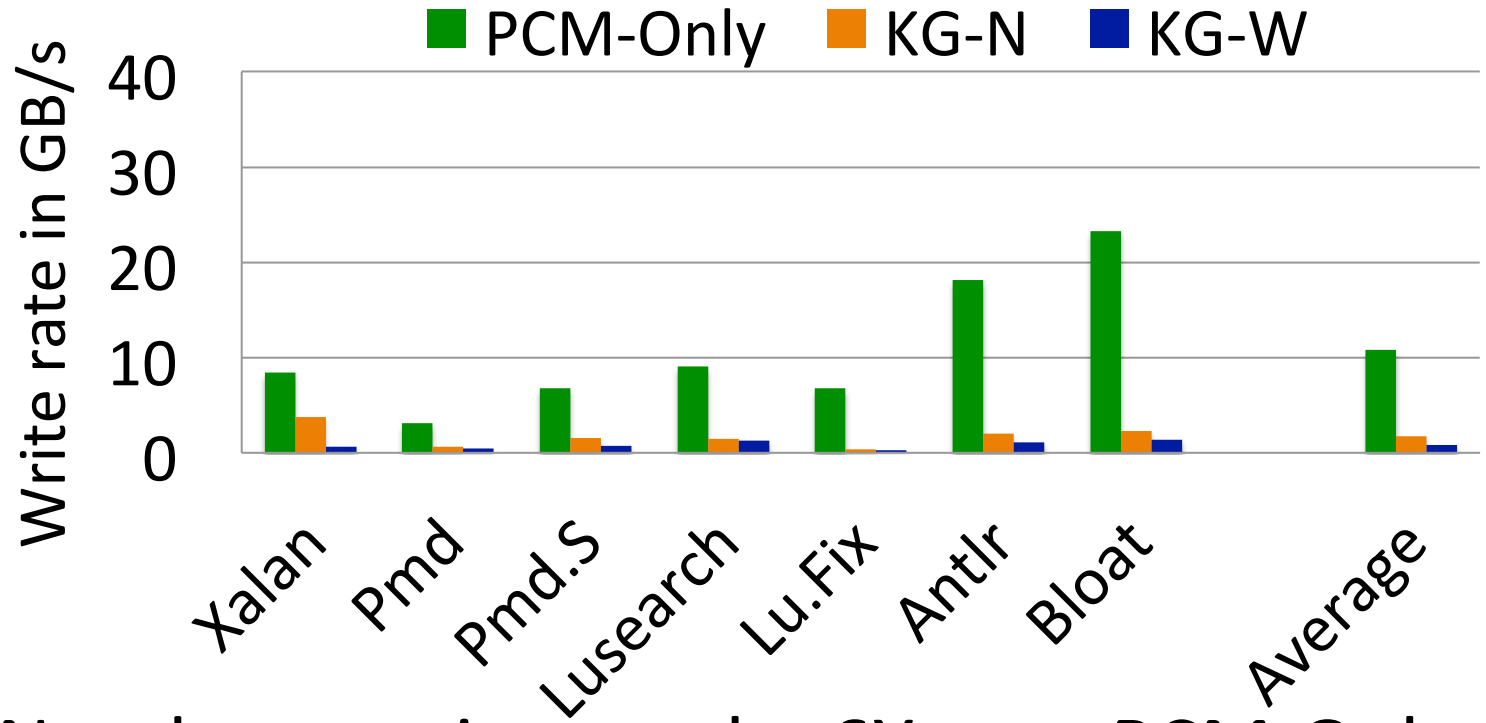

KG-N reduces write rate by 6X over PCM-Only

KG-W reduces write rate by 2X over KG-N

EDP reduction compared to DRAM

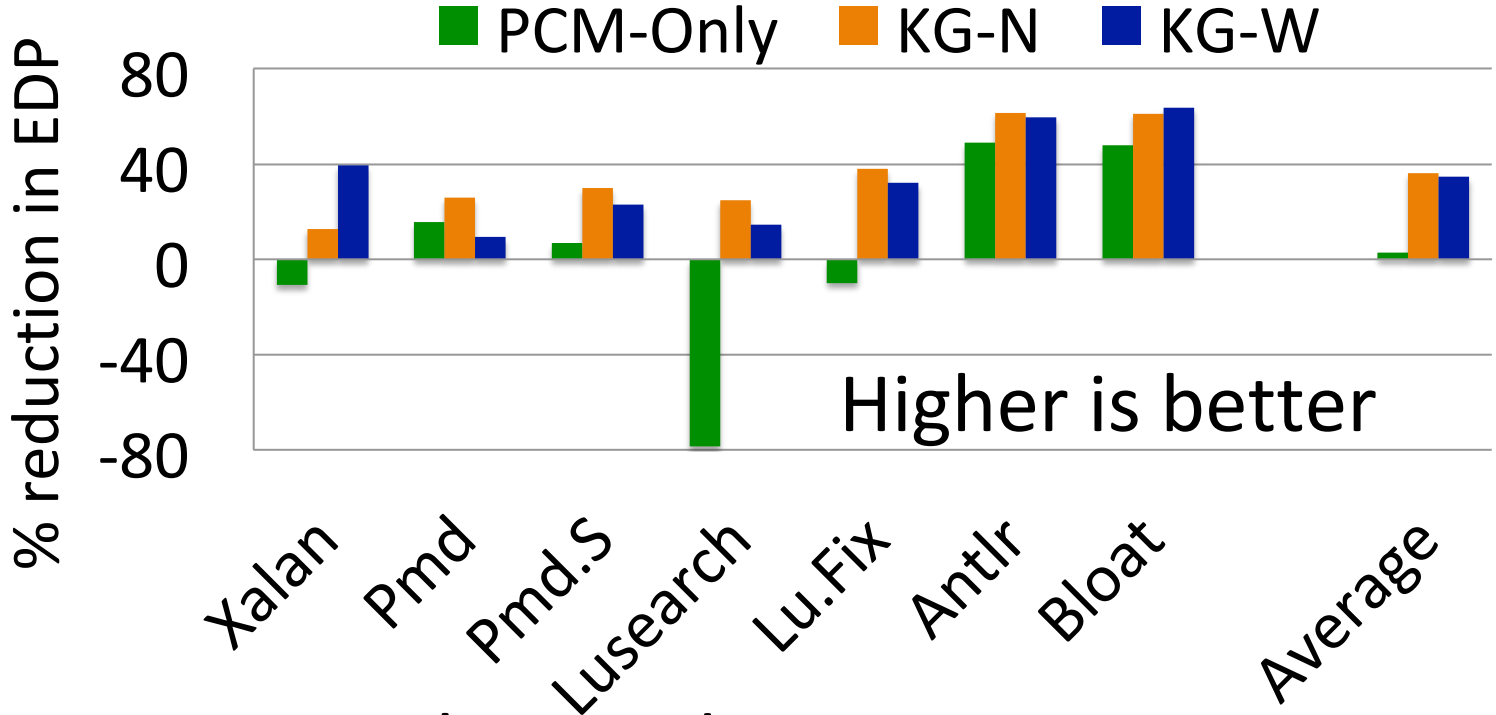

EDP : Energy Delay Product

KG-W has 35% better EDP than DRAM-Only

In the paper

Execution time results

Breakdown of KG-W overheads

Object demographics

Comparison with OS approach

Write rationing garbage collection

Monitor fine grained write behavior of objects

Exploit managed runtimes to organize objects in hybrid memory

Kingsguard collectors improve PCM lifetime

Metadata optimization

GC writes mark bits of PCM objects

KG-W keeps mark bits of PCM objects in DRAM