

PerMA-Bench: Benchmarking Persistent Memory Access

48th VLDB | September 05 – 09, 2022 | Sydney, Australia

Lawrence Benson, Leon Papke, Tilmann Rabl

Hasso Plattner Institute, Potsdam, Germany

Contribution

We ...

» ... **propose PerMA-Bench**, a benchmark framework to analyze the performance of customizable database-related PMem access.

» ... **evaluate PMem performance** across four PMem servers.

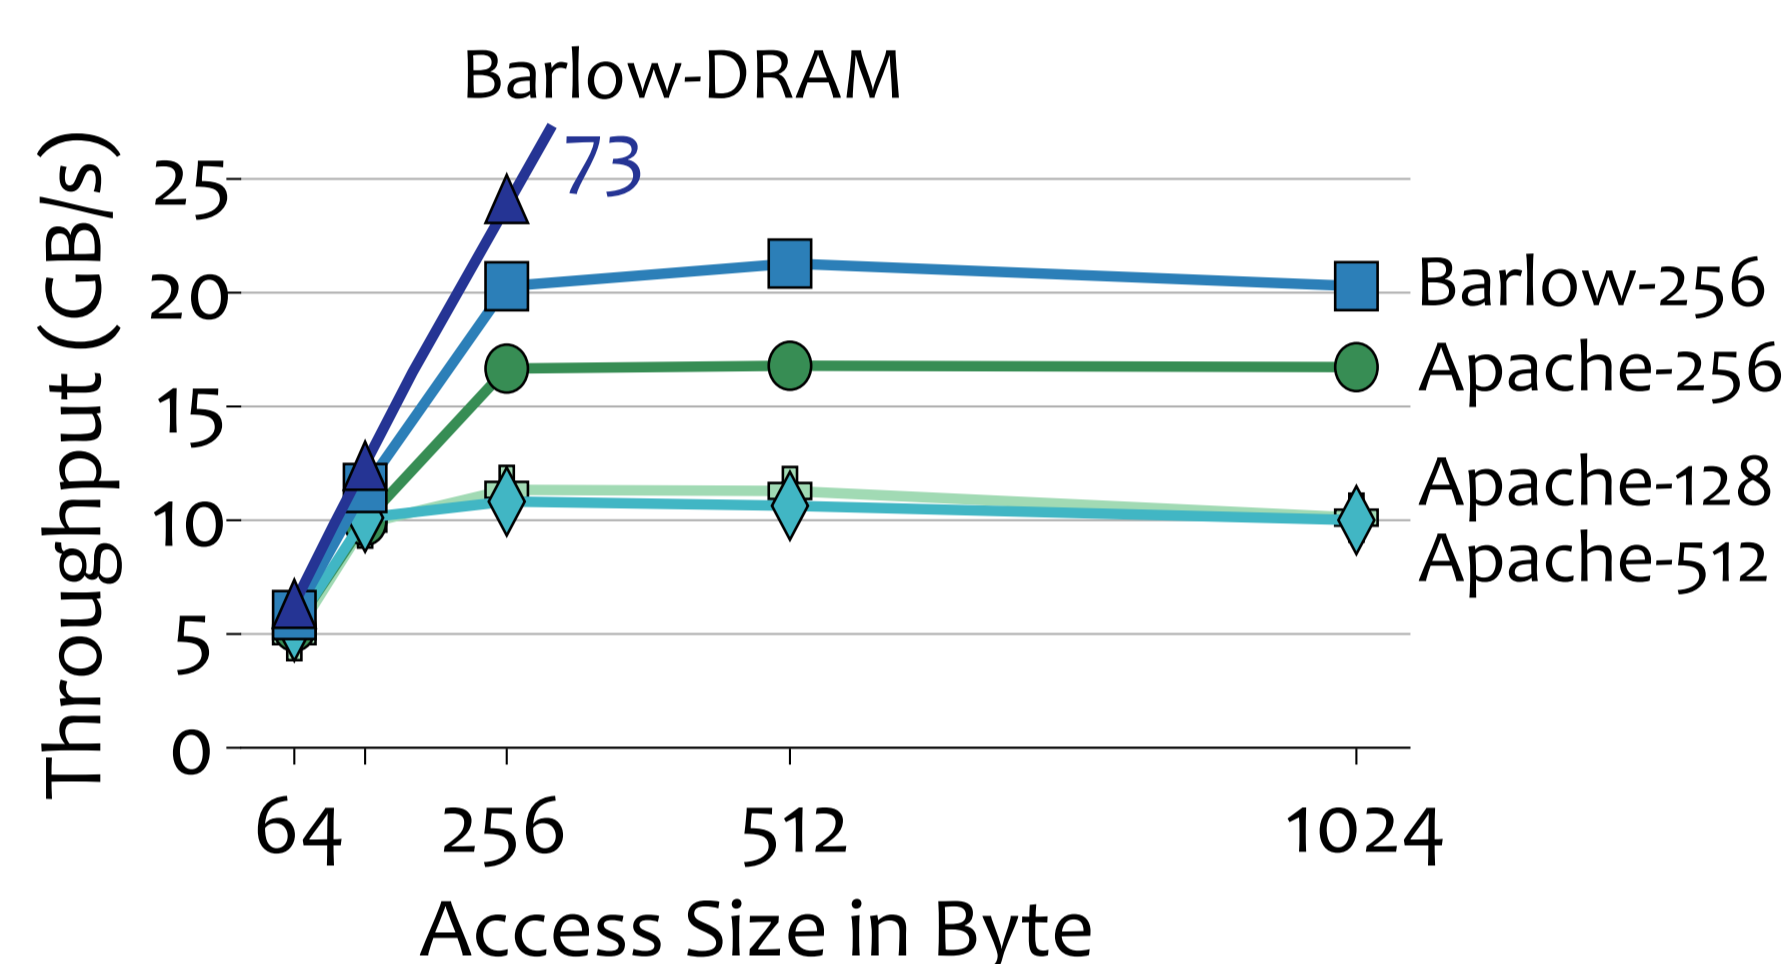
» ... perform a **price-performance analysis**.

» ... **identify eight insights** that influence PMem performance.

Evaluated Servers

System	CPU	PMem	DIMM Power	DRAM
Apache-128 (A-128)	Cascade Lake 18 Cores	6x 128 GB 100 Series	15 Watt	6x 16 GB
Apache-256 (A-256)	Cascade Lake 18 Cores	6x 256 GB 100 Series	18 Watt	6x 16 GB
Apache-512 (A-512)	Cascade Lake 24 Cores	6x 512 GB 100 Series	15 Watt	6x 64 GB
Barlow-256 (B-256/B-D)	Ice Lake 32 Cores	8x 256 GB 200 Series	15 Watt	8x 32 GB


Sequential Writes



- » >20 GB/s in 200 Series
- » Seq. writes >> random writes
- » Higher power == higher bandwidth (A-256 has 40% higher bandwidth than previously reported)

Insight: The DIMM power budget has a large impact, but is not always configurable.

PerMA vs. Real Systems



» Dash^[1] at memory access performance

» F+F^[2] not at memory performance

» Not fully utilizing 200 Series


» Not utilizing 200 Series at all

» 20% cycles for non-memory access

» 30% cycles for non-memory access

Insight: Current designs are too tailored or not tailored enough.

eADR



- » eADR *may* be beneficial to performance
- » Good for <64 Byte (FPTree, LB+Tree)
- » Neutral for ~128 Byte (Dash)
- » Bad for ~200 Byte (Viper)

Legend: ▨ Persist ▨ s-fence only

Insight: Know when to explicitly flush in eADR mode.

Key Insights

PMem Configuration:

1. DIMM size has a performance impact
2. Power budget has a performance impact
3. Number of DIMMs scale linearly (except 1)
4. Memory bus speed has no performance impact

Future PMem Research:

5. Hardware utilization should be improved
6. Persist Instructions should be chosen well
7. eADR is not a performance silver bullet
8. Prefetcher impact should be considered

Price Performance

System	€/GB capacity	€/GB/s seq. read	€/GB/s rnd. read	€/GB/s seq. write	€/GB/s rnd. write
Apache-128	9.21	0.25	0.34	0.78	3.43
Apache-256	10.74	0.25	0.33	0.64	2.78
Apache-512	16.60	0.56	0.61	1.52	6.18
Barlow-256	12.77	0.22	0.33	0.60	2.12
B-DRAM	59.37	0.38	0.46	0.70	0.91

PMem is ...

- » cheaper per GB
- » competitive for reads
- » competitive for sequential writes
- » not competitive for random writes

→ PMem ≈ cheaper, larger DRAM

Code



GitHub:
hpides/perma

Paper



lawben.com/
publication/
perma