



Characterizing a Real PCM Storage Device

University of Pittsburgh & Micron Technology, Inc.‡

Ju-Young Jung, Kelli Ireland, Jiannan Ouyang, Bruce Chiders, Sangyeun Cho, Rami Melhem, Daniel Mosse, Jun Yang, Youtao Zhang, and A. J. Camber‡

1 Motivation

Round 1 is over



HDD

NAND flash-based SSD
(Solid-state Storage Device)

Round 2 already started

SSD versus PSD



Do you keep betting on SSD without knowing a rising star, PSD?

There has been little work evaluating actual PSDs empirically

PCM@PITT gives PSD characterization results

2 Evaluation Approach

Experimental Environment

PSD	Micron's 12 GB PCIe (PCI Express) PSD prototype
SSD	OCZ's 160 GB PCIe Revodrive x2
Host system	Dell OPTIFLEX 980 Desktop computer running Linux 2.6.35 kernel
Benchmark	fiio version 1.41 working with ext4 file system

Preconditioning & Glitch Elimination:

measure performance only after a warm-up period incurring 50% disk fullness, and repeat the measurement process 5 times and average them.

Design parameters identification:

classify parameters that may affect storage system performance into two categories:

Workload-level parameters

- ❖ Block size, file size, I/O depth, read/write mixture ratio
- ❖ Disk fullness, request pattern, multi I/O threads etc.

System-level parameters

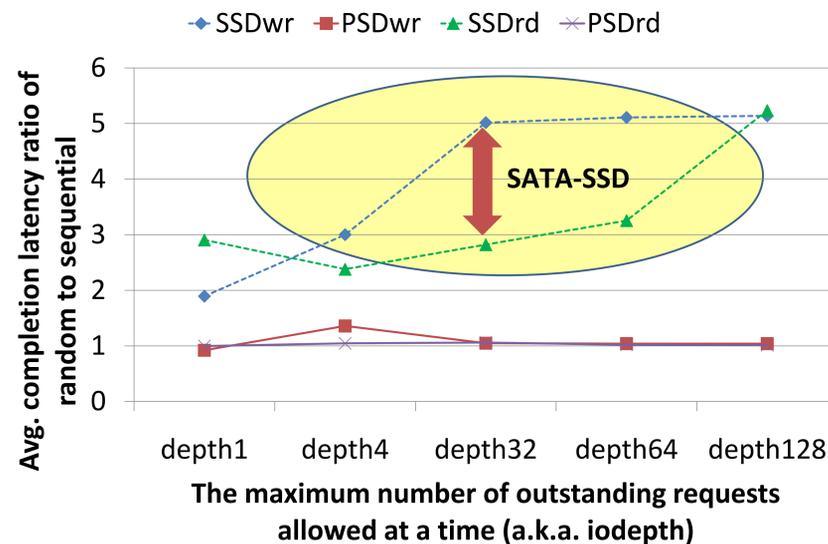
boot time, OS stack latency etc.

3 Preliminary Results

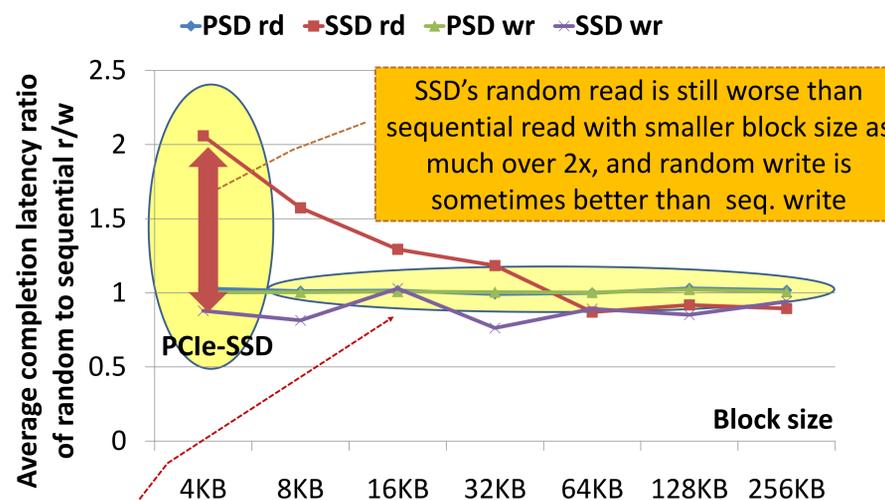
Sensitivity to randomness:

performance of random read and write operations normalized to that of sequential read and write over both PSD and SSD.

Well-known fact: old SATA-interfaced SSD is sensitive to randomness



Did recent PCIe-based SSD overcome this problem?
Not at all! Just get ameliorated a bit

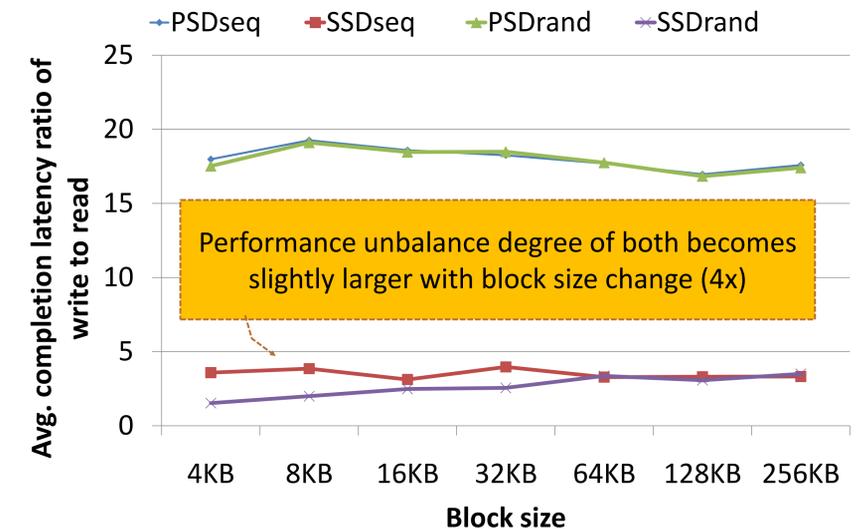
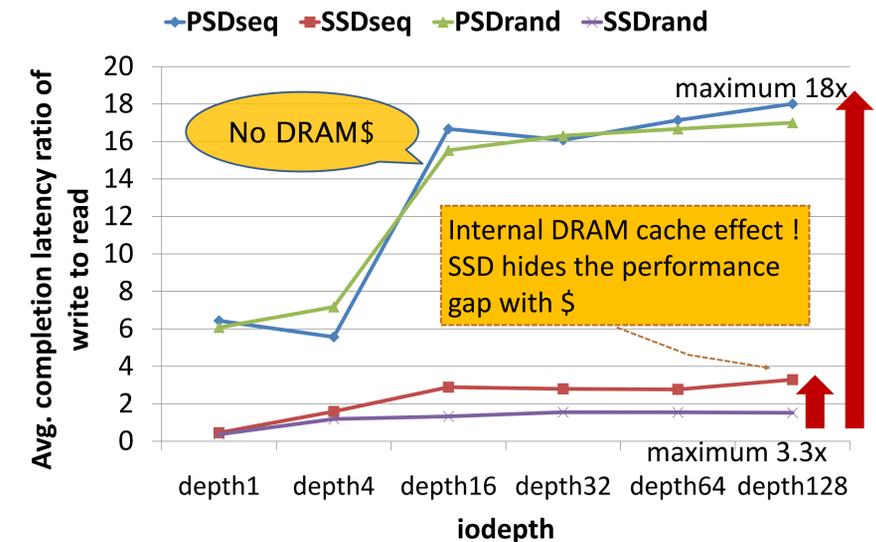


PSD doesn't show any noticeable performance difference between sequential and random read or write

With 20% less writes, SSD improves performance less than 20% except for sequential write, which is more sensitive to mixture ration of read/write. Meanwhile PSD shows nearly constant improvement (40%).

Degree of performance unbalance between read and write:

Both NAND flash memory and PCM suffer from relatively slower write performance than read. See write performance normalized to read performance.



Sensitivity to the mixture ratio of read and write:

This graph explains how the mix ratio of read and writes affect the performance of each operations. The result is R:W = 70%:30% normalized to the mix ratio of R:W = 50%:50%.

