Performance Characterization of Hyperscale Applications on NVMe SSDs

Qiumin Xu[§], Mrinmoy Ghosh[ϟ], Manu Awasthi[ϟ], Tameesh Suri[ϟ], Zvika Guz[ϟ], Anahita Shayesteh[ϟ], Vijay Balakrishnan[ϟ]
[§] University of Southern California [ϟ]Samsung Semiconductor Inc

NVMe Analysis

NVMe SSDs build on PCIe to provide:

✦ **Superior interface capabilities**
PCIe x4 supports higher BW (4GB/s) compared to SATA (600MB/s)

✦ **Shorter hardware data path**
NVMe directly connects to a PCIe port; SATA traverses through AHCI HBA

✦ **Simplified software stack**
NVMe driver for SSDs bypasses the traditional OS block layer

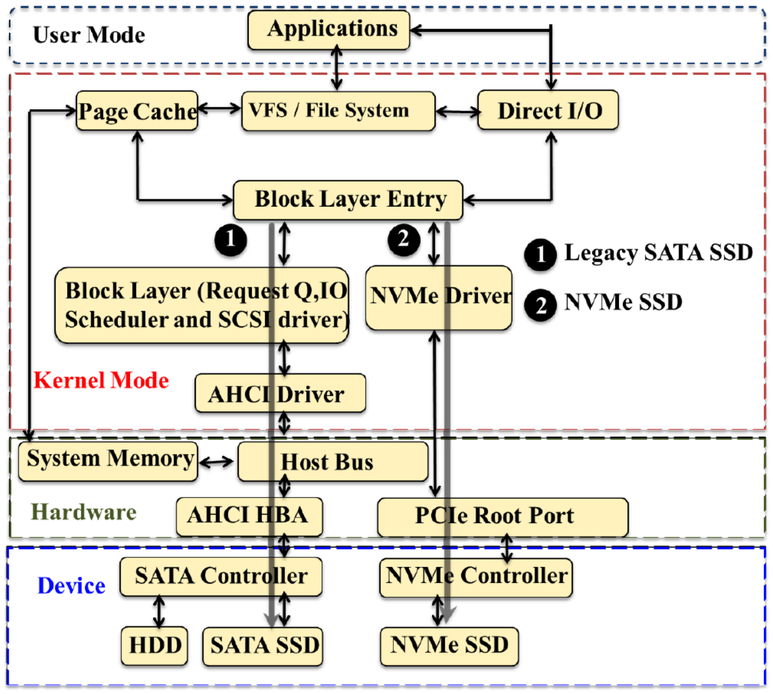

Figure 1: A comparison of the SW/HW architecture of SATA and NVMe SSDs

Samsung XS1715 NVMe SSD

✦ Industry's first NVMe SSD
✦ Densities up to 1.6TB
✦ Random read performance reaches up to 750k IOPS

Comparison of Software Overheads

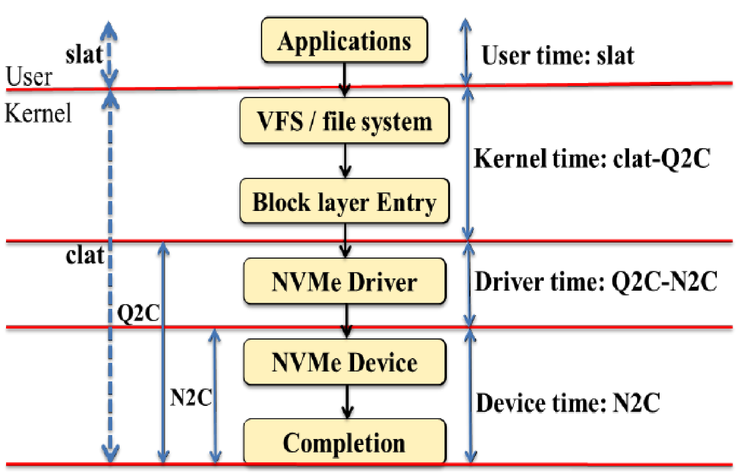

Figure 2: Latency breakdown

✦ Instrumented NVMe driver and *blktrace*

✦ NVMe S/W overhead is much lower than SATA overhead

✦ HDD S/W overhead is negligible because of very large seek and transfer times

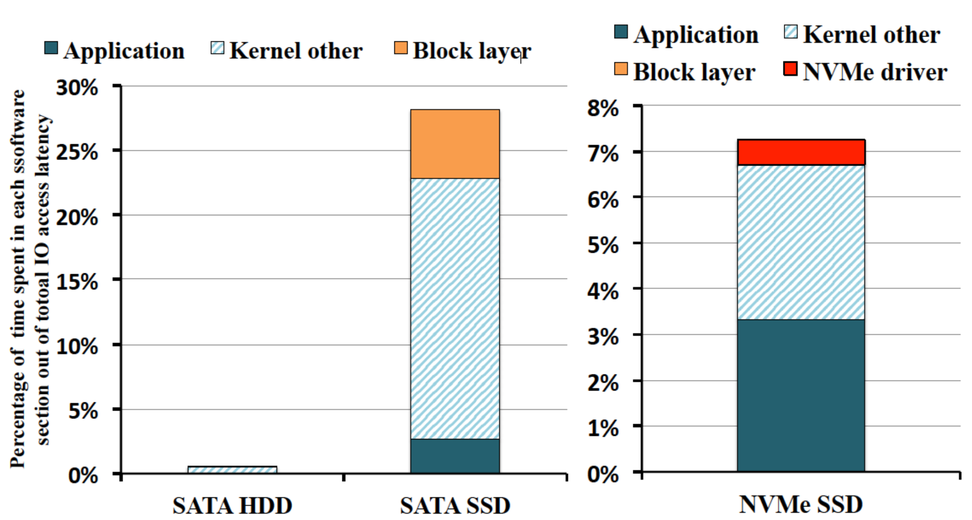

Figure 3: S/W overhead comparison between HDD, SATA and NVMe SSDs

NVMe SSD Performance for Cloud Databases

✦ TPC-C used to drive MySQL and YCSB used to drive Cassandra

✦ NVMe-backed database applications deliver up to 8x superior performance

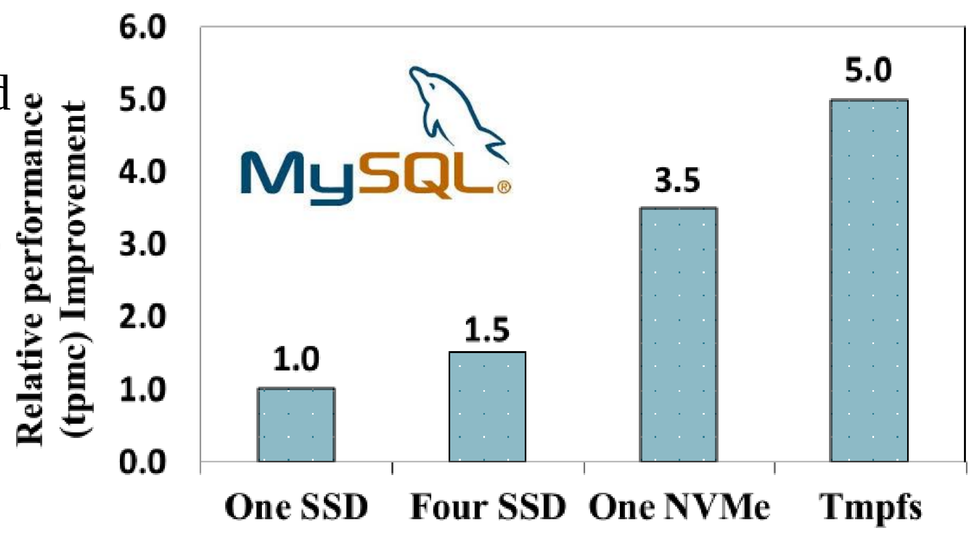

(a)

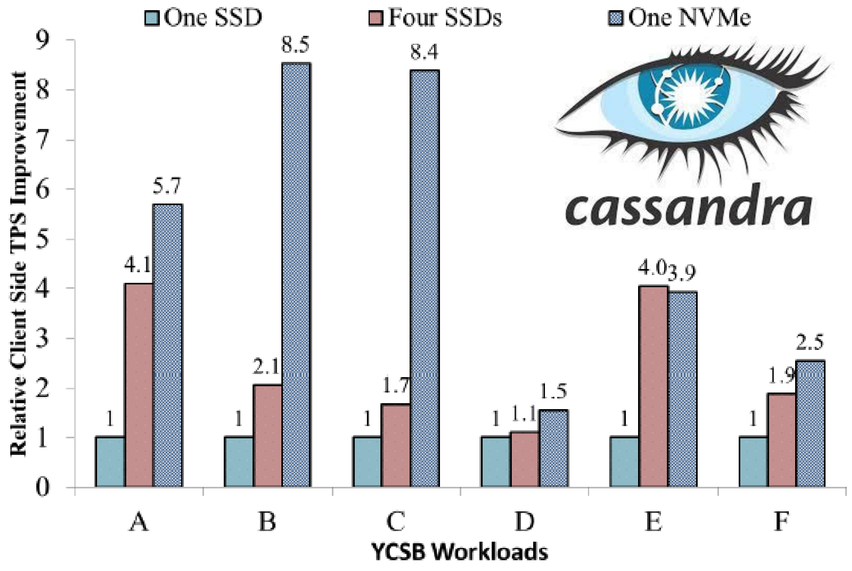

(b)

Figure 4: Performance improvement of (a) TPCC and (b) Cassandra over various storage configurations

Ecosystem MSL Research Development

Memory Solutions Lab

Align with your imagination

SAMSUNG