Comprehensive Benchmark Suite for SSD

Anupam Tambi
Member DISCL

08/06/2011

Introduction

- Solid-State Drives (SSDs) predicted to replace HDDs
  - *SSDs have greater durability and speed over HDDs*
  - *HPC systems start shifting to SSDs*
- Traditional benchmarks inadequate for SSDs
  - *SSDs show best case results on traditional benchmarks*
  - *Do not reflect the actual ability of the SSD well*
  - *Do not offer insights of the measurement*
- We propose a Comprehensive Benchmark Suite for SSDs (CBS-SSD)
  - *Includes SSD specific benchmarks in addition to traditional benchmarks*
  - *A better comparison tool that give realistic values*

Comprehensive Benchmark Suite for SSD

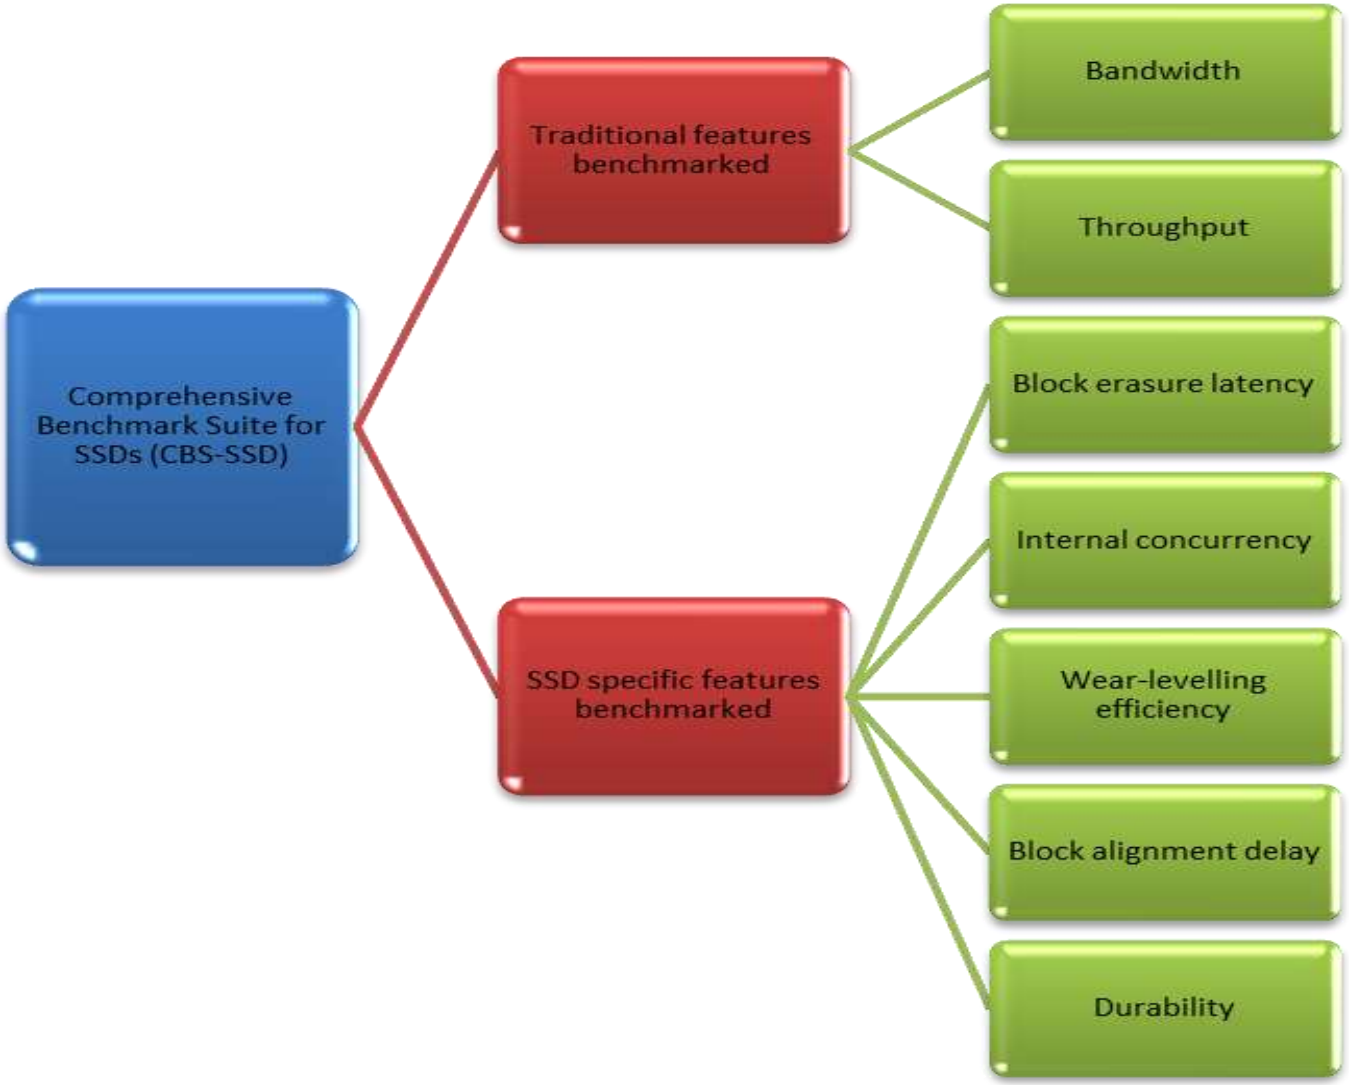

Comprehensive Benchmark Suite for SSD

- CBS-SSD consists of
  - *Traditional features measured*
    - Bandwidth
    - Throughput
  - *SSD Specific features measured*
    - Block Erasure latency
    - Internal Concurrency
    - Wear Leveling efficiency
    - Block Alignment
    - Durability

Block Erasure Latency

- Most SSDs are NAND based
- If a block containing data with a binary value ‘0’ has to be changed to ‘1’, the whole block would have to be erased and the whole block re-written with the new value
- This causes a delay in the modify operation
- Measuring this time delay for modifying one ‘0’ to ‘1’ for various SSD devices would be an effective benchmark
- Implemented by selecting a block that contains a ‘0’ bit and is masked with a ‘1’ bit
- The time taken by the SSD to carry out this operation is measured

Wear Leveling Efficiency Check

- Wear-leveling widely used in SSDs to make sure that no single flash unit is worn out due to repeated usage
- Wear leveling distributes each single new write operation to a new location on the SSD
- Measuring the speed of wear leveling is very important as it can add a lot of over-head
- CBS-SSD injects requests of constantly increasing sizes to SSD and measures the time taken
  - *A black-box approach*
- A perfect wear leveling algorithm would show a balanced pattern

Internal Concurrency Check

- This benchmark measures IOPS when read and write operations are performed concurrently on different planes of the SSD
- This would give a better understanding of the power of the on-chip processor of the SSD in resolving race conditions
- Implemented by a black-box approach as well
- CBS-SSD injects requests of constantly increasing sizes to the SSD and measures the time taken
- This helps find the number of concurrent planes
- Data with sizes that are multiples of the number of planes are written to the SSD and the times are measured
- Comparing these times would show how efficiently the planes are concurrently written.

Block Alignment Delay Check

- NAND flash devices are divided into erasable blocks composed of multiple pages
- A flash block that contains data must be fully erased prior to writing new data to the block
- When the partitions are misaligned, write performance typically suffers a great deal
- The erasure process for a single block can take up to several milliseconds
- Unit block size of the SSD is determined and data sizes that are not multiples of the block size is written
- Time taken to write is measured and the delay is calculated

Block Alignment Delay Check

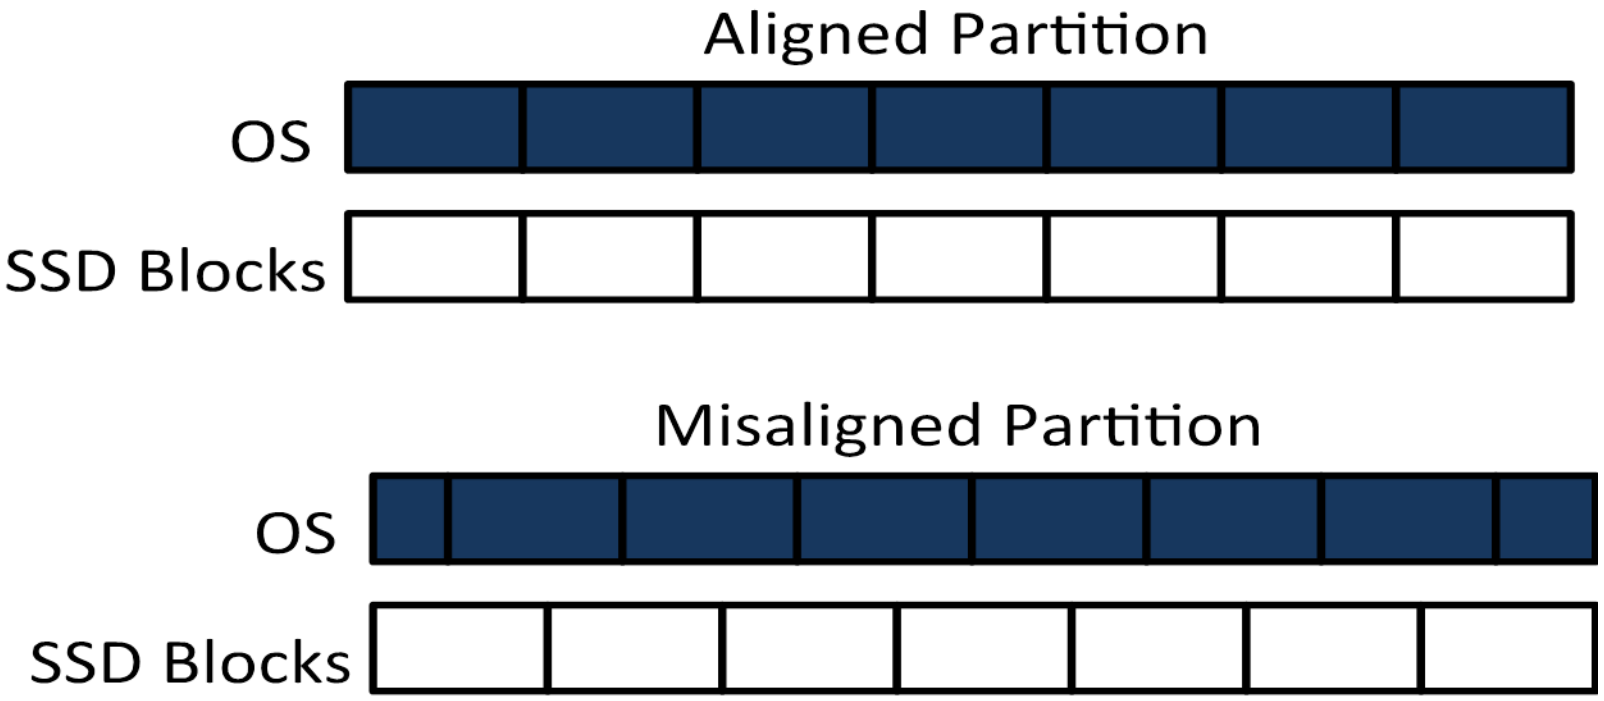

Durability Check

- One drawback of SSDs is that SSDs have a limited number of write operations that can be performed
- The efficiency of the SSD degrades as more and more write operations are performed
- The lifetime of the SSD depends on factors like the wear leveling algorithm used
- A benchmark to measure efficiency of the wear-leveling algorithm would help to check if the every write operation is distributed properly or not

Preliminary Results

- CBS-SSD reports and compares the performance benchmarked for each feature
- The block erasure time and the performance of the internal concurrency check are demonstrated
- Comparisons are based on six SSDs, with block size of 128KB, 256KB, and 512KB and concurrent planes of 2 and 4 respectively

	Block Size	No of Concurrent planes
SSD-A	128 KB	2
SSD-B	256 KB	2
SSD-C	512 KB	2
SSD-D	128 KB	4
SSD-E	256 KB	4
SSD-F	512 KB	4

Preliminary Results (Contd..)

- Performance comparison of block erasure times

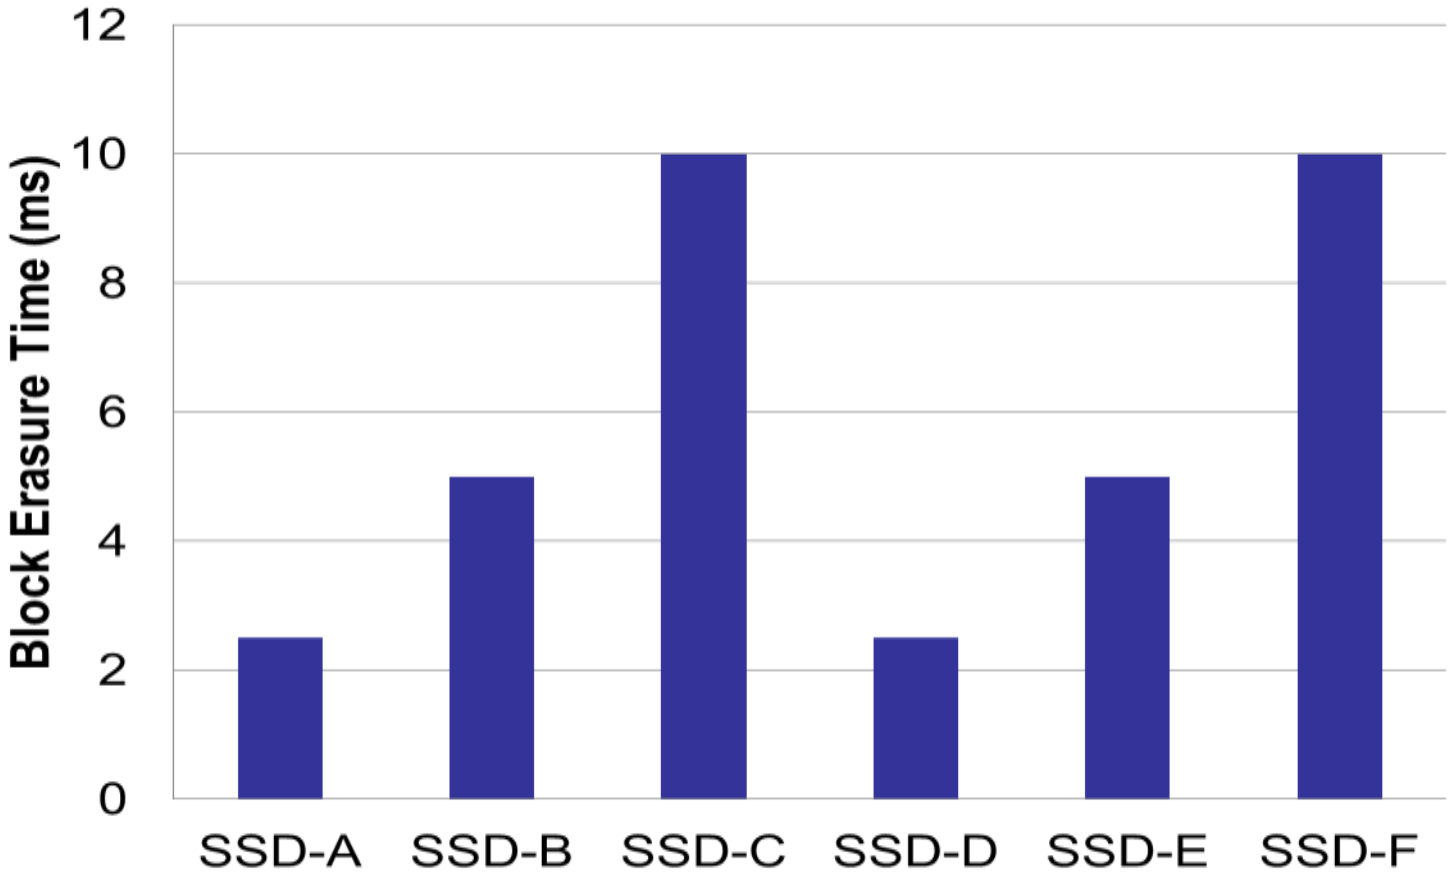

Preliminary Results (Contd..)

- Performance comparison of internal concurrency

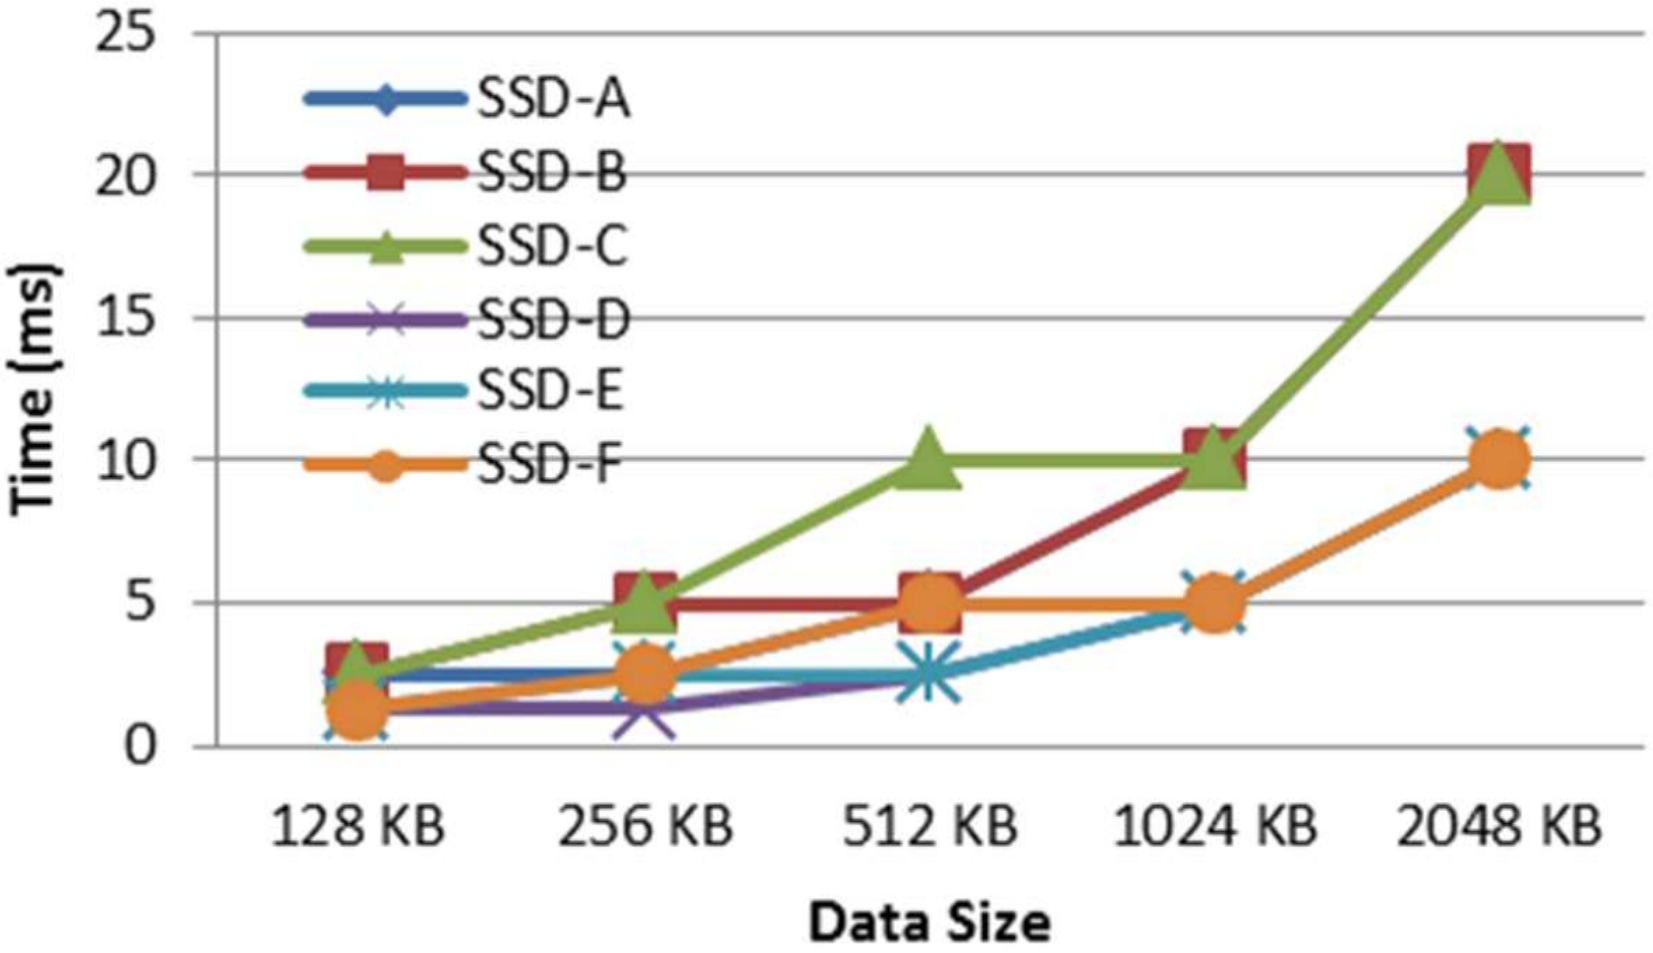

TEXAS TECH UNIVERSITY™