



# A Performance Evaluation of Scientific I/O Workloads On Flash –Based SSDs

By Stan Park and Kai Shen

# Review

- Introduction.
- Background and related work.
- Evaluation Setup and Methodology.
- Results and Analysis.
- Conclusion.

# Major Differences

## HDD

- Several seconds for sign up time.
- No Random access.
- Mechanical reliability.
- Not shock resistive.
- Very high power consumption.
- Weight and Size is comparatively high.
- No limits on writes.
- Cost is very low.

## SSD

- Instantaneous.
- Random access.
- No mechanical reliability.
- Shock resistive.
- Very low power consumption.
- Weight and Size is low and hence much portable.
- Limited number of writes.
- Cost is very high.

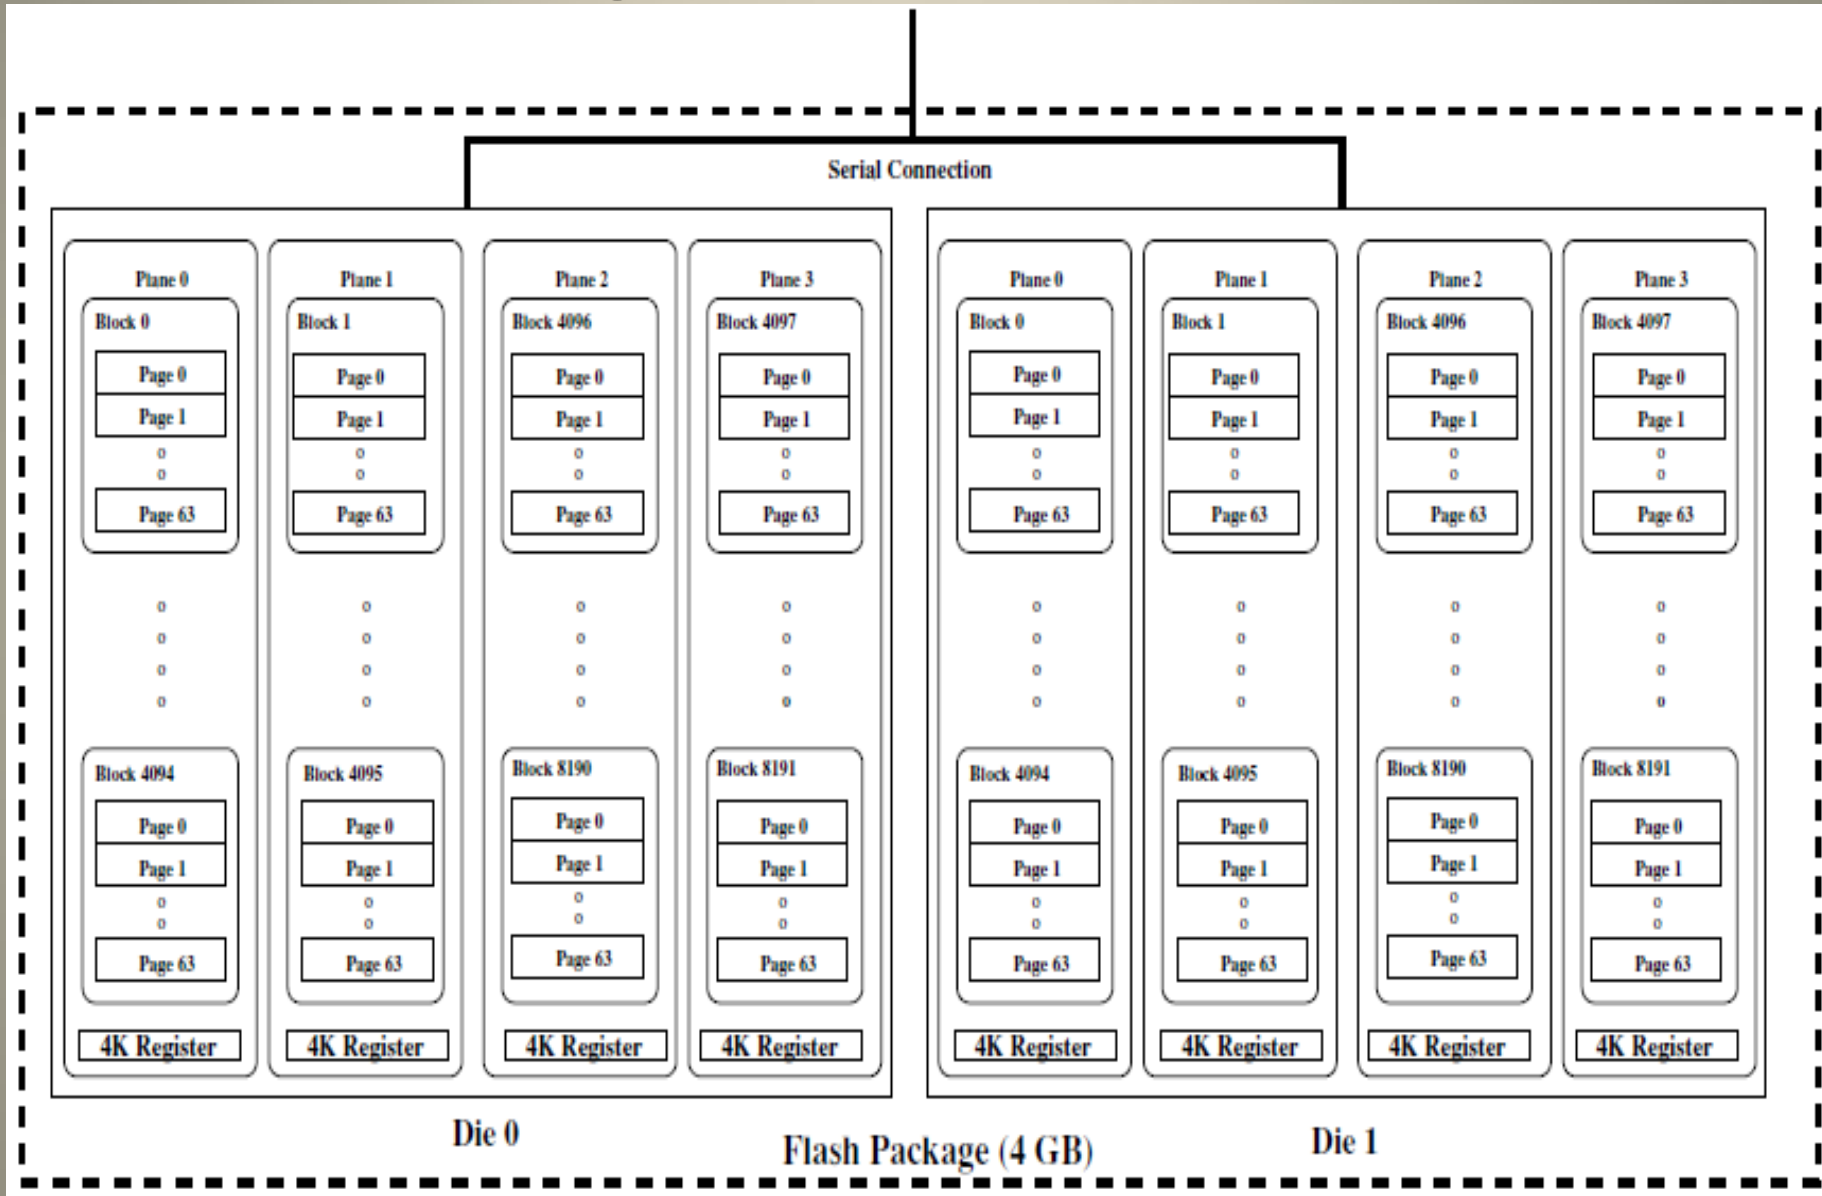
# Our Focus

- Performance evaluation of scientific I/O workloads on SSDs.
- The performance gains of SSDs over mechanical disks under practical circumstances.(Multiple processes in a parallel application access single storage device)
- Performance gaps between different SSDs.

# Review

- Introduction.
- **Background and related work.**
- Evaluation Setup and Methodology.
- Results and Analysis.
- Conclusion.

# Samsung 4GB Flash Internals



# Read and Write Issues

- Read performance is significantly better compared to HDDs.
- Critical limitations with write:
  1. Erase-Before-Write : to overwrite a previously written location must be erased.
  2. Write is performed at page level and erase is performed at block level.



## SLC(Single Level Cell)

- Stores one bit of information.
- Lower storage density.
- Faster write speeds.
- Lower power consumption.
- Higher cell endurance.
- Costlier.

## MLC(Multi Level Cell)

- Stores two or more bits.
- Better storage density.
- Lower write speeds.
- Higher power consumption.
- Lower cell endurance.
- Cheaper.

# Review

- Introduction..
- Background and related work.
- **Evaluation Setup and Methodology.**
- Results and Analysis.
- Conclusion.

# Some Important Definitions

- **Scientific Workloads:** It means, reading in large data sets, transforming (some kind of operation like adding) in some way and writing them back.
- **Trace:** It is simply a logging of a set of data regarding the performance of storage device focusing on I/O requests.

We use: LANL-64KB, LANL-1024 KB, Sandia-alegra, Sandia-cth, Sandia-s3d.

# Experimental System

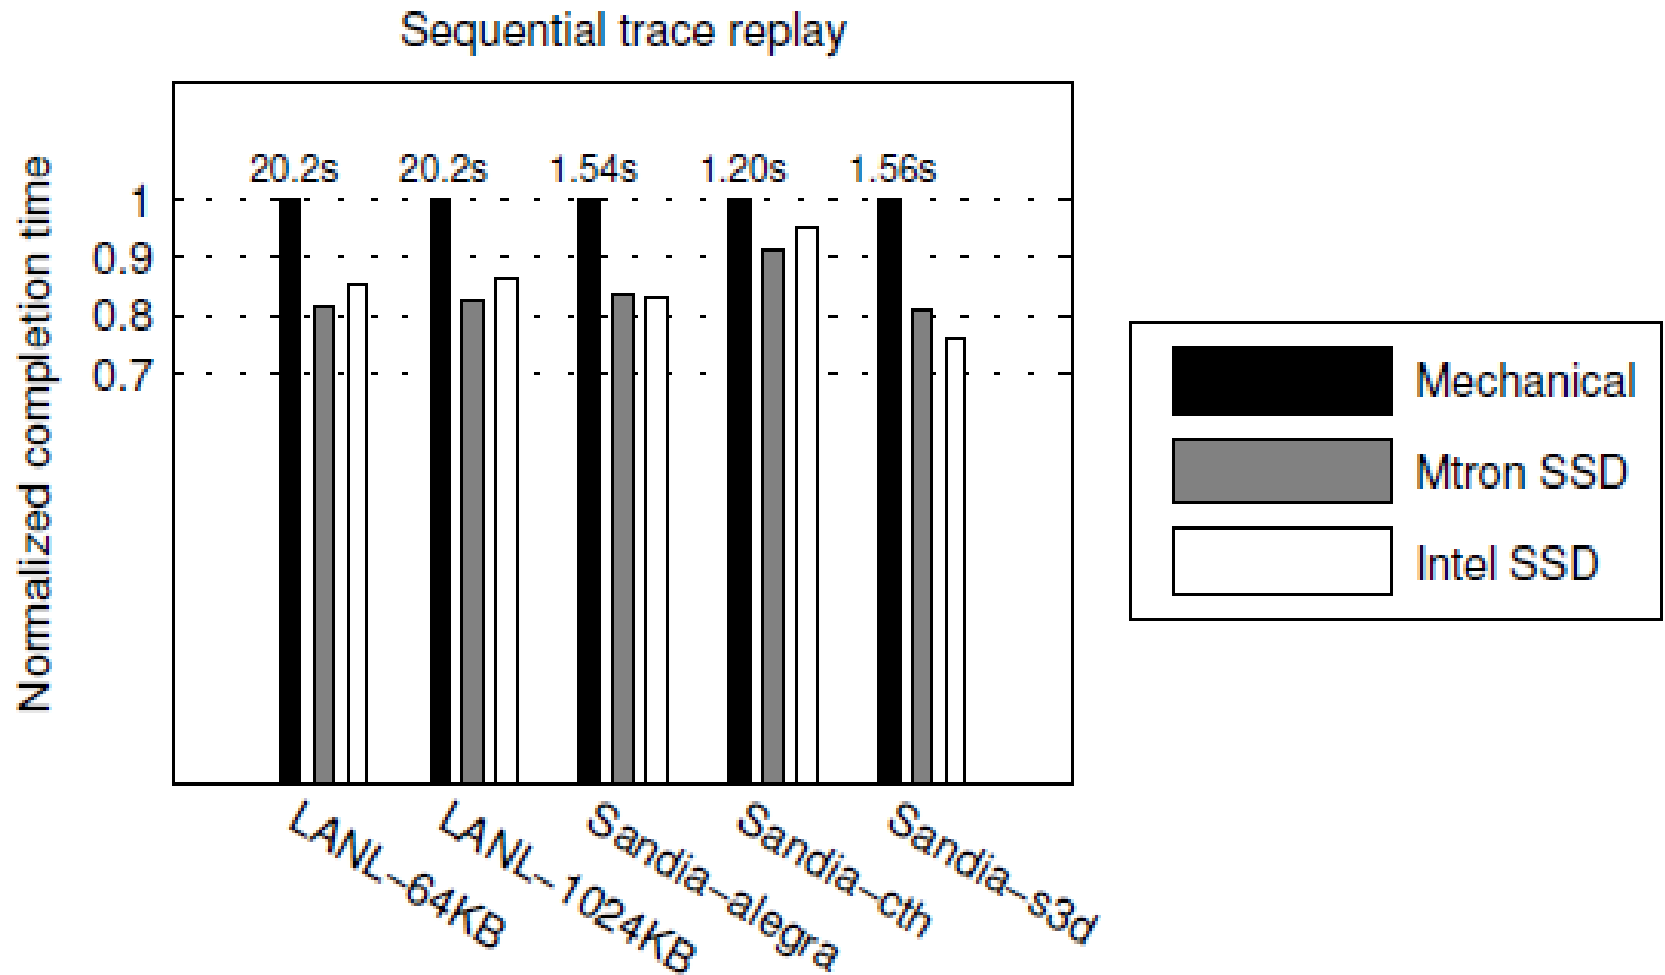
- Two SSDs: Intel X25-M (MLC), Mtron Pro 7500 (SLC).
- One HDD: WD Cavier.
- Processor: Intel core 2 Duo2.0 GHz.
- Memory: 2GB.
- OS: Linux kernel.
- File system: Linux file system ext 3.
- I/O scheduler: cfq scheduler.

# Evaluation Methodology

- Sequential mode: Replay the trace segment of one process on the target storage device w/o considering other process in trace.(each process accesses a dedicated storage device).
- Parallel mode: Multiple processes in a parallel application simultaneously access a single storage device.

In both above cases we don't consider inter-process synchronization.

# Results of Sequential Mode

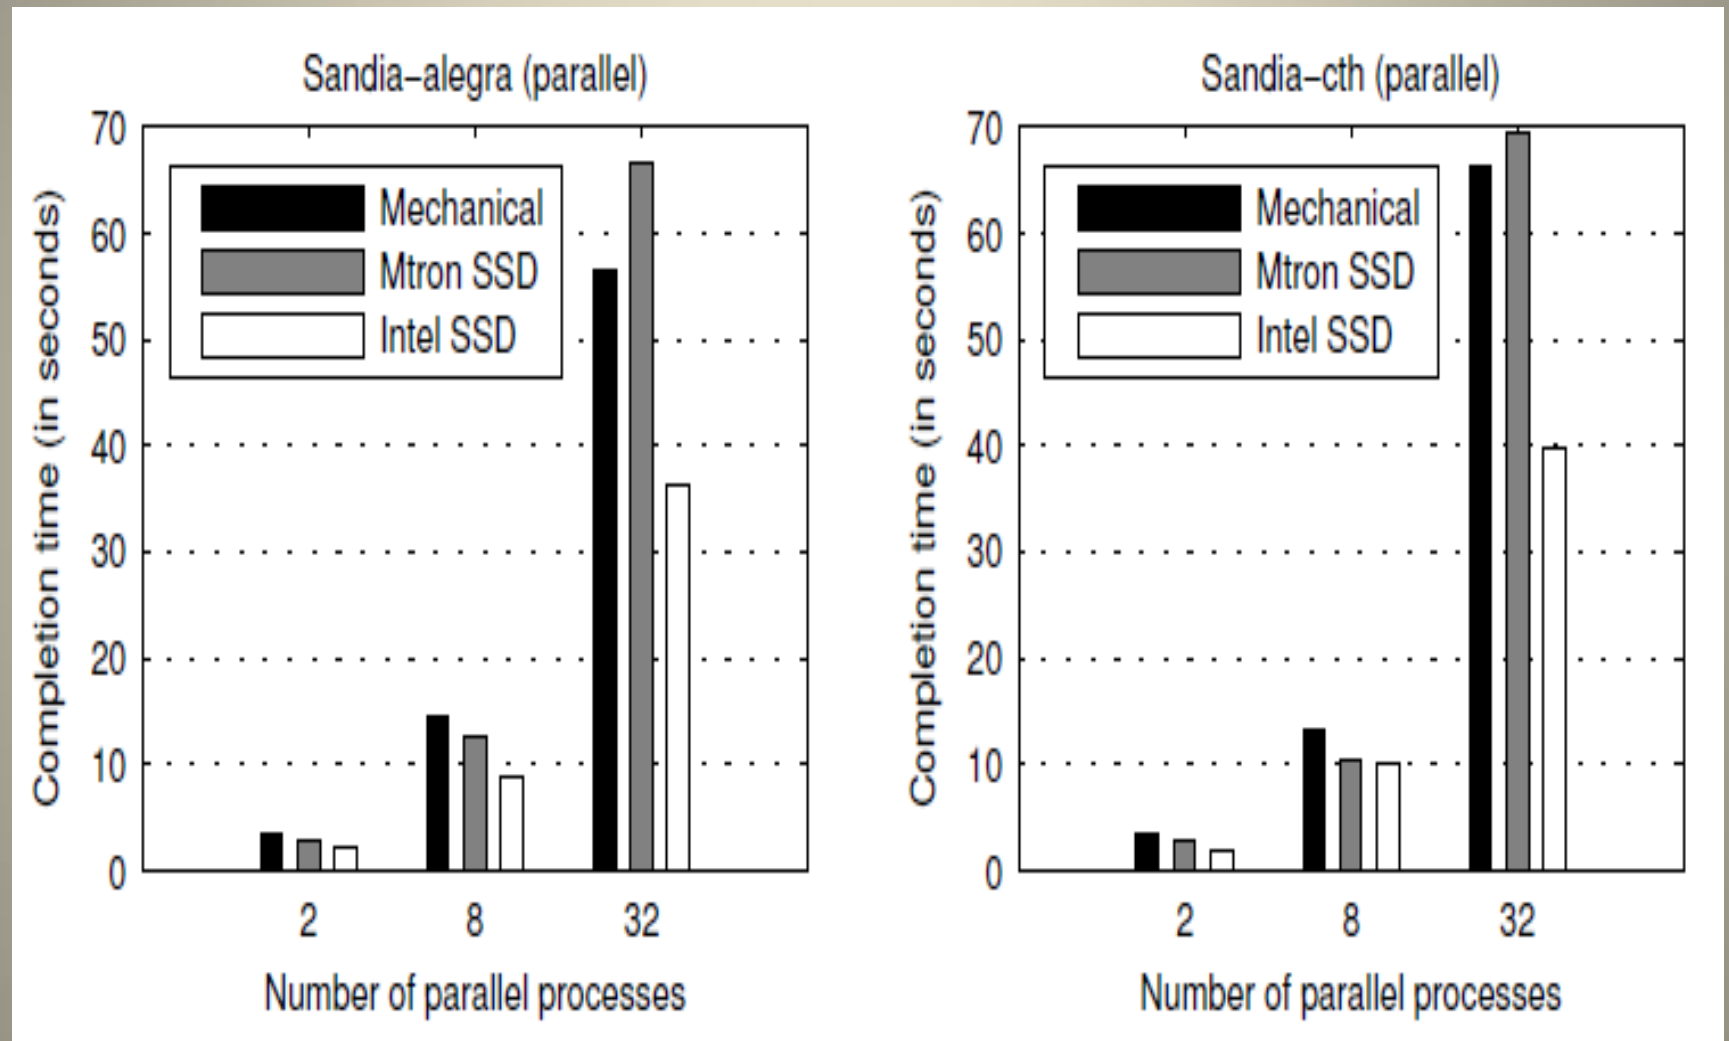


# Analysis

- Sequential Mode:

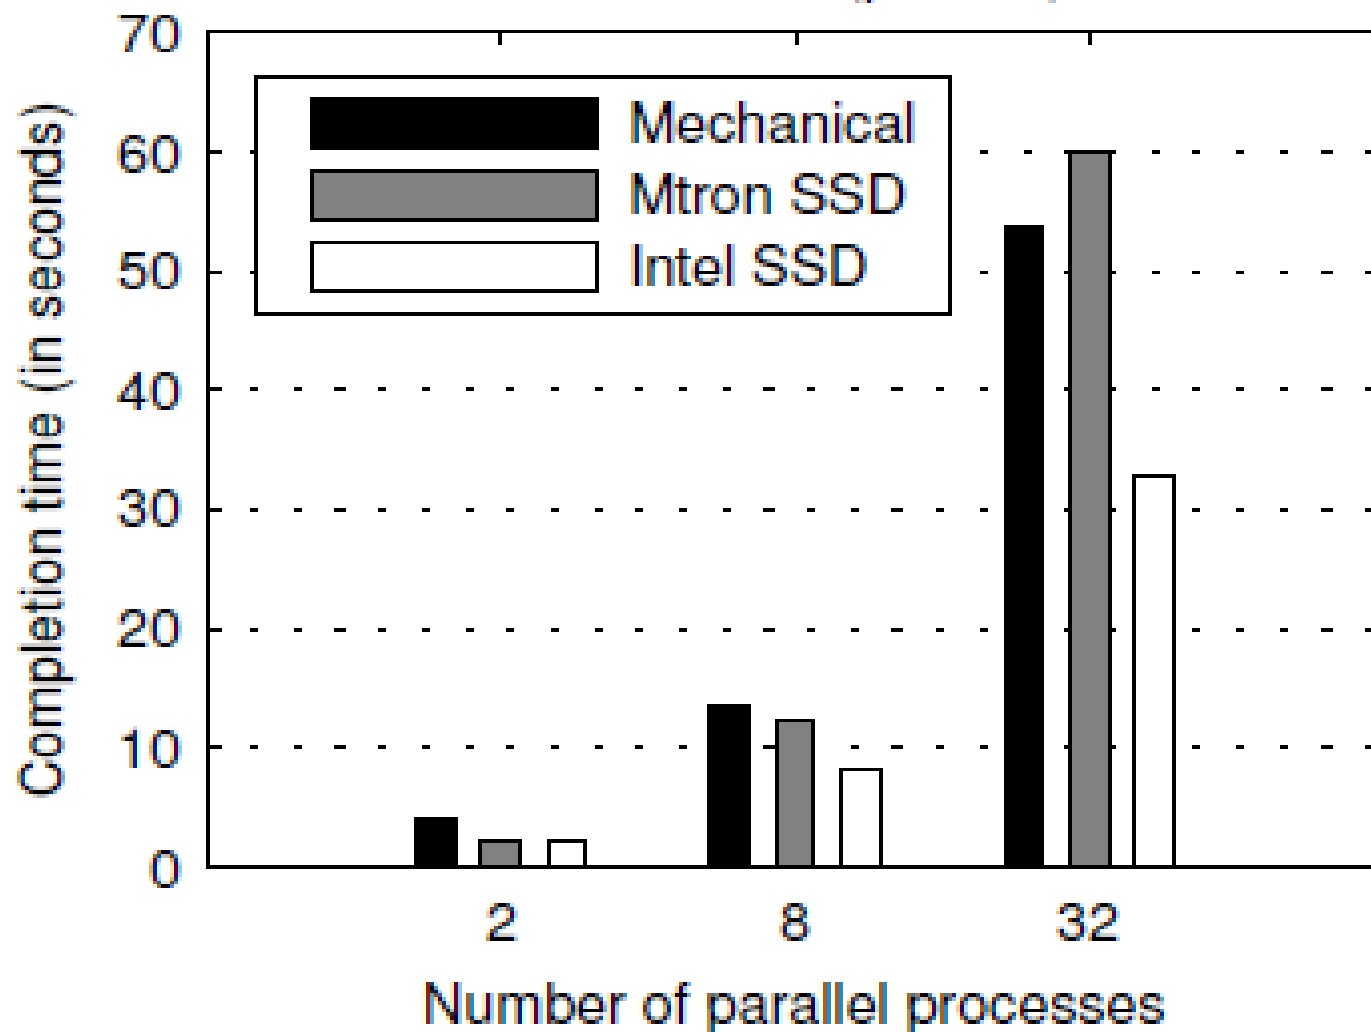
No significant improvement. It is because of the write intensive nature of the workloads.

# Results of Parallel Mode<sub>(continue)</sub>





Sandia-s3d (parallel)



- Parallel Mode:
  1. The ordering of the incoming processes doesn't affect seek time in an SSD. (Absence of mechanical seek latency.)
  2. HDD reorders the concurrent processes' requests to minimize seek distance but still satisfied by a disk head, i.e. sequentially. This is not the case with SSD and hence it performs better. (overlooked)

# Conclusion

- SSDs provide modest performance gains over HDDs when workload is write intensive. For read, it performs superb.
- For parallel process mode, the performance gain is excellent.
- SSD implementation feature is a very important aspect for performance gain.

Thank You.