

Junghee Lee⁽¹⁾, Youngjae Kim⁽²⁾, Jongman Kim⁽³⁾, and Galen M. Shipman⁽²⁾

⁽¹⁾ The University of Texas at San Antonio, San Antonio TX

⁽²⁾ Oak Ridge National Laboratory, Oak Ridge, TN

⁽³⁾ Georgia Institute of Technology, Atlanta GA

Introduction

Redundant Array of Inexpensive Disks (RAID)

- ✓ Higher performance and robustness
- ✓ De facto standard for high-performance storage systems

RAID of Solid State Drives (SSDs)

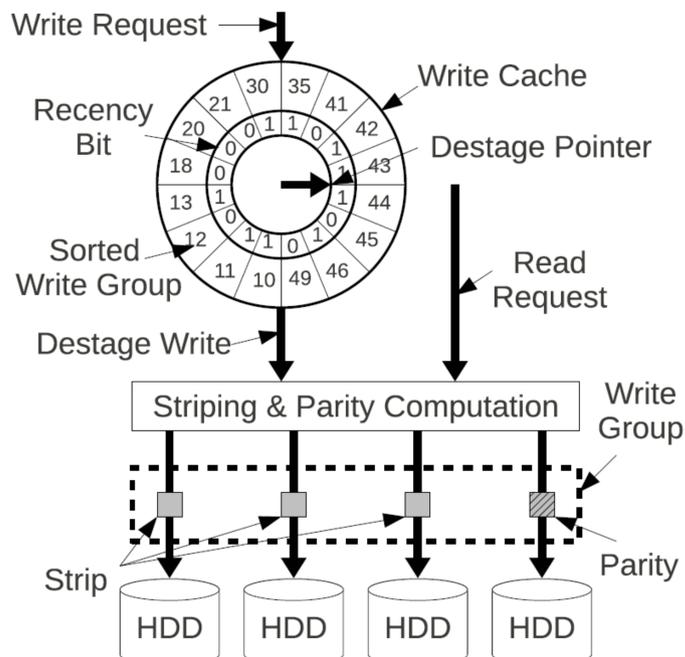
- ✓ Even better performance and robustness
- ✓ The overall performance of RAID of SSDs is limited by the slowest SSD
- ✓ The performance variation of SSDs becomes a major concerns
- ✓ SSD can be delayed by unpredictable internal Garbage Collection (GC) process

Contributions of this paper

- ✓ A new write cache architecture that considers the performance variation of SSDs

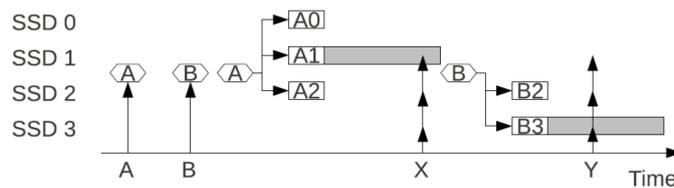
Background

Existing write cache for RAID: WOW

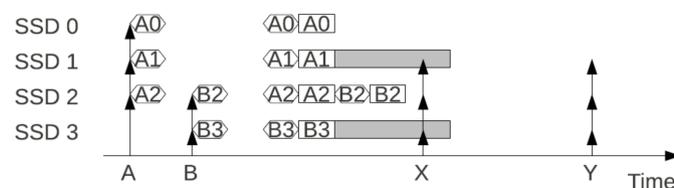


Motivation

◻ Cache Read/Write ◻ Destage ◻ Garbage Collection

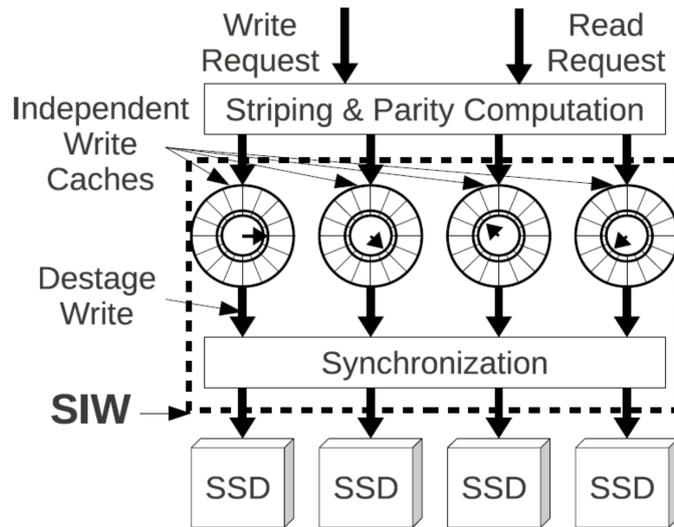


(a) Traditional write cache



(a) Proposed write cache

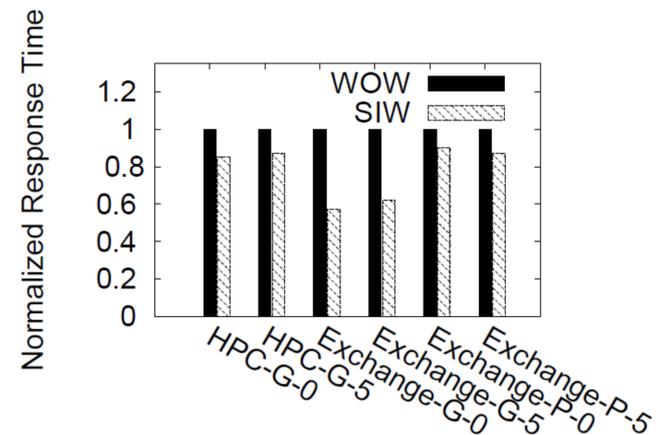
Proposed Write Cache



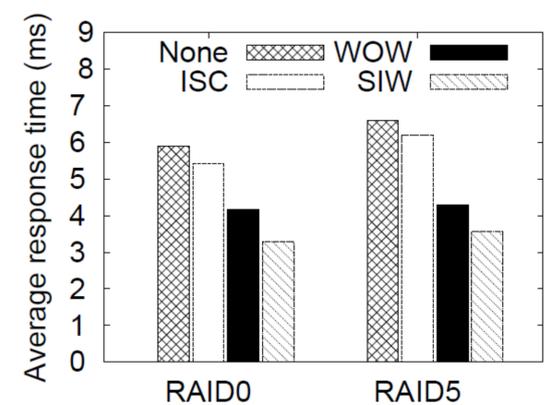
Synchronous Independent Write Cache (SIW)

- ✓ Independent write caches: employing an independent write cache per each SSD
- ✓ Synchronization mechanism: write caches are independent but they are synchronized in that they try to destage together at the same time
- ✓ Pre-parity-computation: compute parity before it is stored in the write cache so that write caches can be scheduled independently

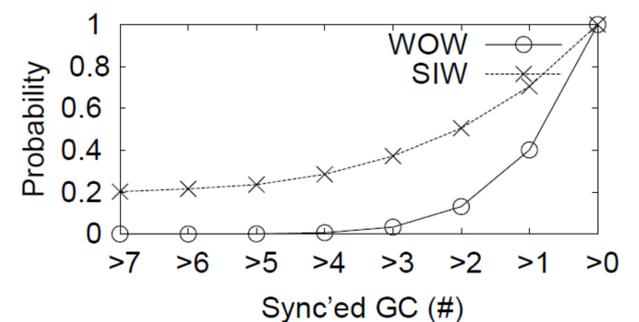
Experiments



(a) Average response time



(b) Placement of write caches



(c) Number of synchronized GCs

Conclusion

A new write cache design is proposed for an array of SSDs

The proposed write cache, Synchronous Independent Write Cache (SIW) improves the response time by up to 50% and 20% on RAID-0 and RAID-5, respectively, compared with the state-of-the-art write cache algorithm