DS-RAID Efficient Parity Update Scheme for SSDs

Jaeho Kim(Presenter), Jongmin Lee, Jongmoo Choi, Donghee Lee, Sam H. Noh

University of Seoul, Dankook University, Hongik University (in Seoul, Korea)

Motivation: Enhanced Reliability for SSDs

- Problem with current SSDs: low reliability
  - ✓ High error rate and limited erase count of flash memory
  - ✓ Multi-level cell (MLC) flash memory aggravates problem
- One solution: provide RAID-5 configuration with chips comprising the SSD device

Problems

- Typical RAID-5 [Logical Block Number (LBN) based striping)]: Parity update burden
  - ✓ Small write problem: read old data, calculate parity, write new data
- Added burden when adopting RAID-5 configuration to SSDs: Out-of-place update
  - ✓ With LBN-based striping ➔ MUST keep (new) updated page in SAME chip as old page

➔ Lower **performance** & **reliability**, higher **cleaning cost**, decreased **lifetime**

- Our solution: Dynamic Striping-RAID (DS-RAID)

RAID-5

- **Logical Block Number (LBN)** based striping
  - ✓ Need Stripe map table in RAID Controller
  - ✓ Must read old data & old parity
  - ✓ May result in particular chips being written to more frequently (skewed writes)
  - ✓ Window of vulnerability for new small writes

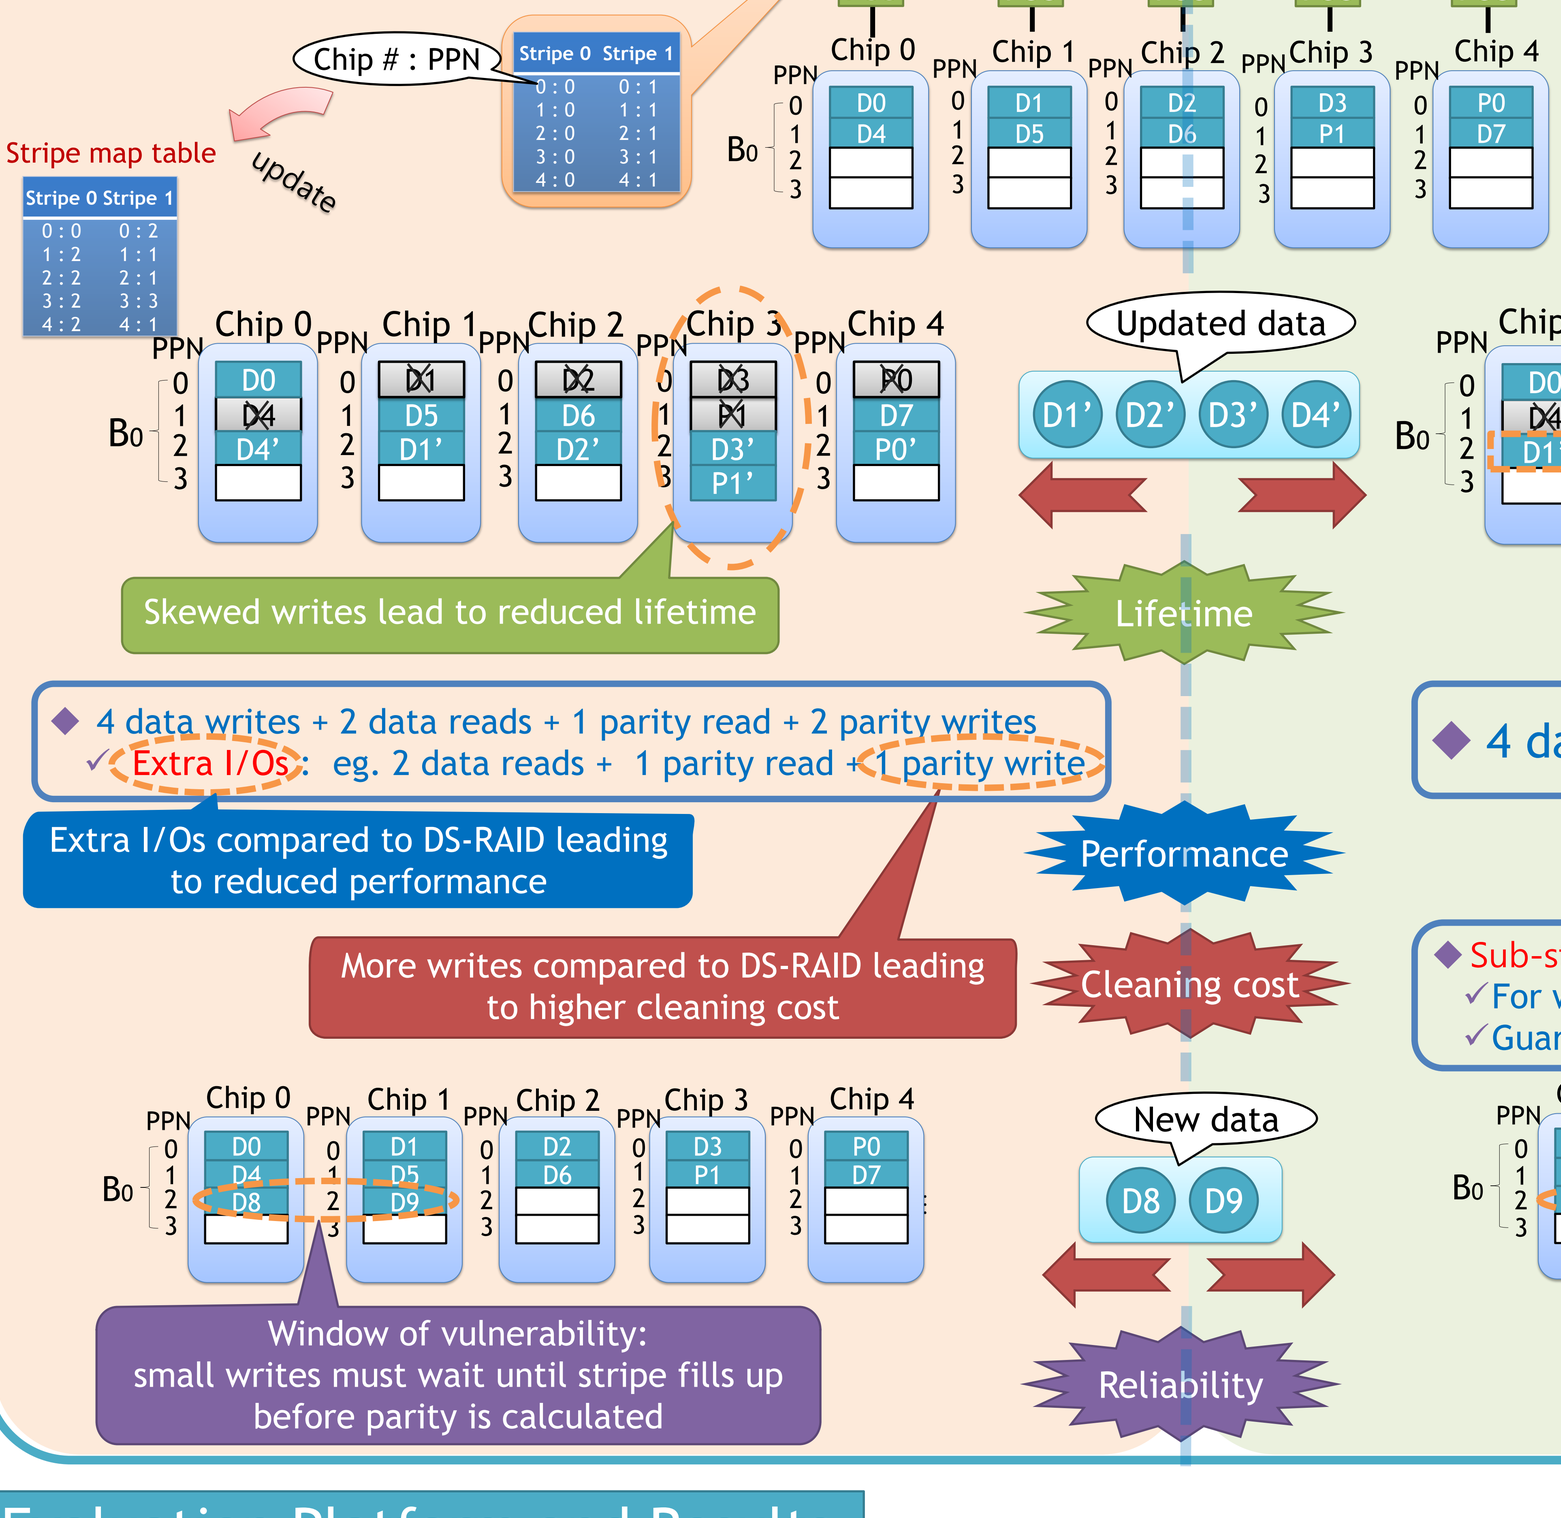

DS-RAID

- **Physical Page Number (PPN)** based striping
  - ✓ No need for Stripe map table
  - ✓ No need to read old data & old parity
  - ✓ All chips are written to evenly
  - ✓ RAID-5 reliability for new small writes (even without non-volatile RAM)

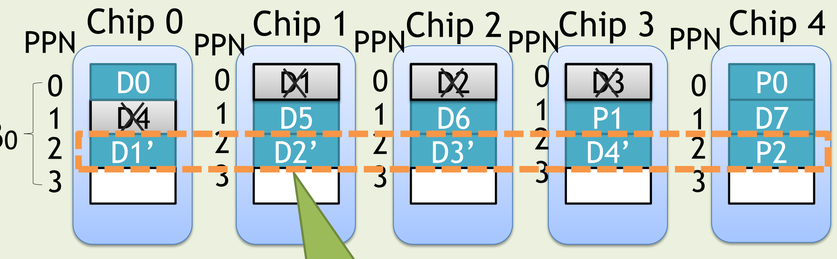

- 4 data writes + 1 parity write
- Sub-stripe parity scheme
  - ✓ For write requests smaller than stripe size
  - ✓ Guarantee RAID-5 reliability via Sub-stripe Parity

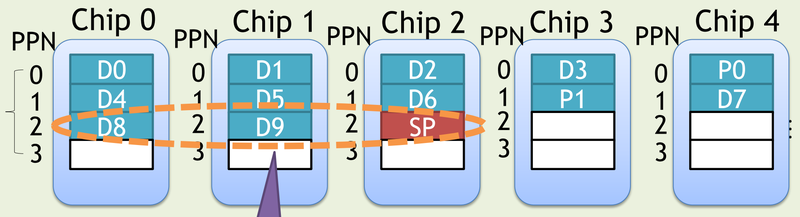

Evaluation Platform and Results

- Evaluation platform
  - ✓ SSD Extension for DiskSim

Parameter	Value	Parameter	Value
Number of chips	8	Planes per chip	8
Blocks per plane	2048	Pages per block	64
Over provision space	5%	Page size	4KB
Stripe size	32KB	Stripe	Consists of 8 pages

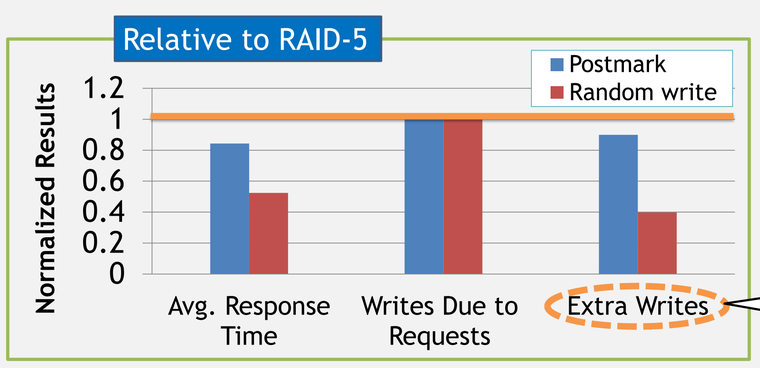

- Workload
  - ✓ Postmark: 1.6MB (Avg. write size)
  - ✓ Random write: 12KB (Avg. write size)

Includes
- Parity writes
- Writes for cleaning
- Sub-stripe parity writes

This work was supported in part by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MEST) (No. R0A-2007-000-20071-0), by the Korea Science and Engineering Foundation (KOSEF) grant funded by the Korea government (MEST) (No. 2009-0085883), and by Basic Science Research Program through the National Research Foundation of Korea(NRF) funded by the Ministry of Education, Science and Technology(2010-0025282).