A Report on the FAST 2015 Storage Conference

André Brinkmann

Johannes Gutenberg University Mainz

Some Statistics ...

- Santa Clara has the highest density of (important) storage conferences in the world
- FAST conference takes four days to cover huge set of storage related topics

Some Statistics ...

- 26 Full papers at FAST conference
  - 6 papers on Flash
  - 2 papers on NVRAM
  - 2 papers on Magnetic Disks
  - 1 paper on Tape / Archiving

Some Statistics ...

- 26 Full papers at FAST conference
  - 2 papers including Wide Area Networks
  - 2 papers on storage for mobile devices
  - 0 papers on parallel file systems
  - 0 papers on transactional memory

Selected Papers

- Skylight – A Window on Shingled Disk Operation (Abutalib Aghayev, Peter Desnoyers)
- How Much Can Data Compressibility Help to Improve NAND Flash Memory Lifetime? (Jiangpeng Li, Kai Zhao, Xuebin Zhang, Jun Ma, Ming Zhao, Tong Zhang)
- RAIDShield: Characterizing, Monitoring, and Proactively Protecting Against Disk Failures (Ao Ma, Fred Douglis, Guanlin Lu, Darren Sawyer, Surendar Chandra, Windsor Hsu)

Selected Papers

- SDGen: Mimicking Datasets for Content Generation in Storage Benchmarks (Raúl Gracia-Tinedo, Danny Harnik, Dalit Naor, Dmitry Sotnikov, Sivan Toledo, Aviad Zuck)
- Write Once, Get 50% Free: Saving SSD Erase Costs Using WOM Codes (Gala Yadgar, Eitan Yaakobi, Assaf Schuster)

Skylight

Software part:

- Launch crafted I/O operations using fio
- Disable kernel read-ahead, drive look-ahead, on-board volatile cache
- Use latency to infer drive properties

Hardware part:

- Install a transparent window on the drive
- Track the head movements using a high-speed camera
- Convert movements to head position graphs

Data Compressibility to Improve NAND Flash Memory Lifetime

Data compressibility :α ? Lifetime extends: $1/\alpha$

- Unused space in one NAND Flash page

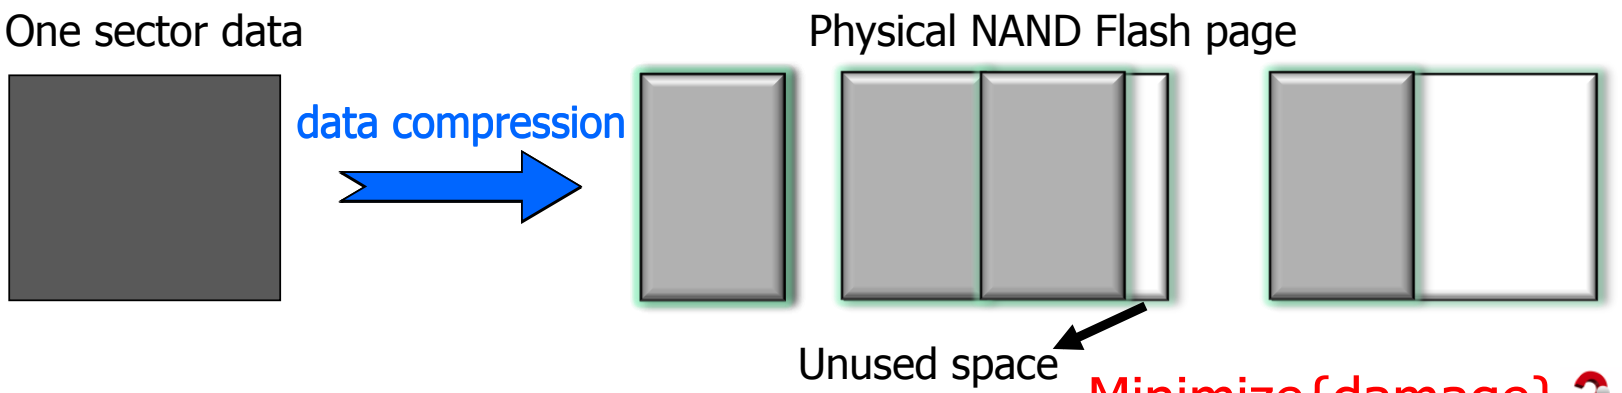

Minimize{damage}

- Impact of compression ratio variance

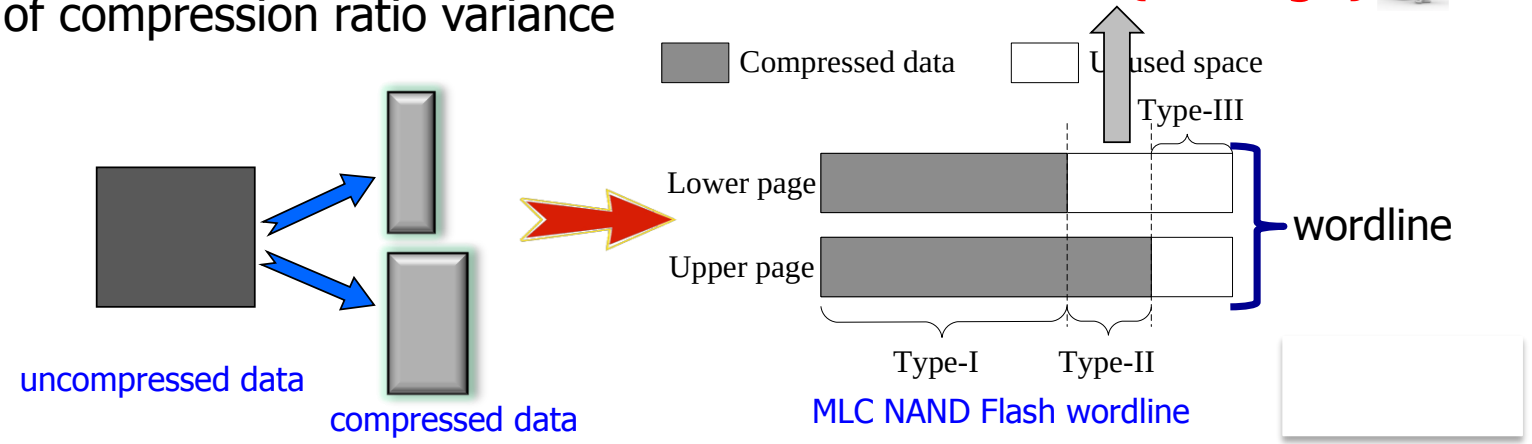

RAID Shield

- Analyzed 1 million SATA disks and revealed
  - Failure modes degrading RAID reliability
  - Reallocated sectors reflect disk reliability deterioration
  - Disk failure is predictable
- Built RAIDSHIELD, an active defense mechanism
  - Reconstruct failing disk before it's too late!
- PLATE: single-disk proactive protection
  - Deployment eliminates 70% of RAID failures
- ARMOR: disk group proactive protection
  - Recognize vulnerable RAID groups

SDGen

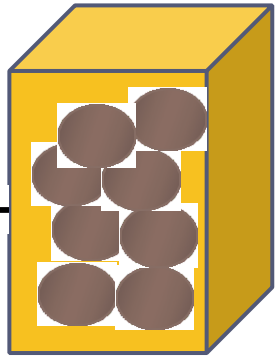

Stones in the backpack!!!

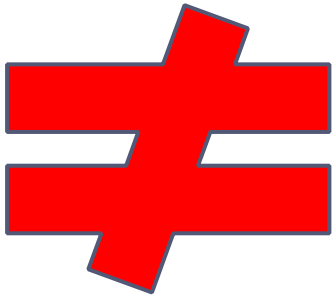

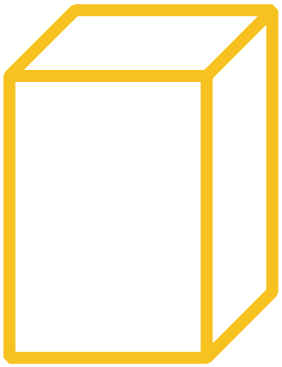

Just thin air ☺

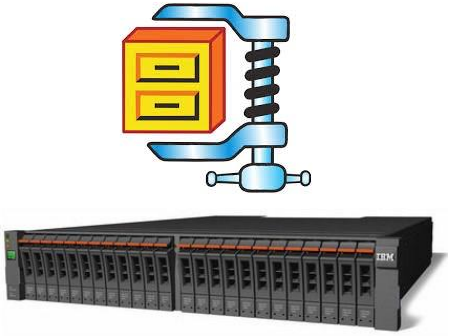

Random Data

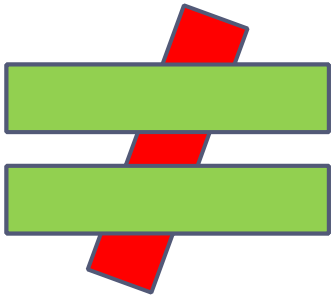

Zero Data

Write Once, Get 50% Free

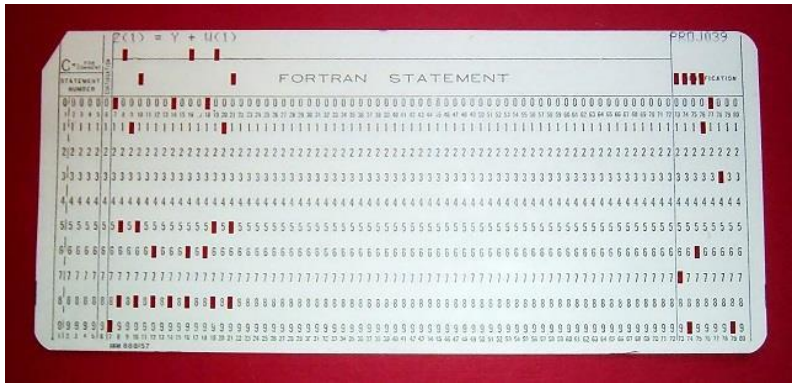

data	1st write	2nd write
00	000	111
10	100	011
01	010	101
11	001	110

(Rivest and Shamir, 1982)

- $0 \longrightarrow 1 \; bu \;\; 1 \nrightarrow 0 \qquad\qquad t$
- WOM Code: write ***n* bits** of information on ***m* cells**, *n>m*
- Example: **write *11*** and then **write *01***
  - Normally: 

  - WOM code: