

Graduate Operating Systems

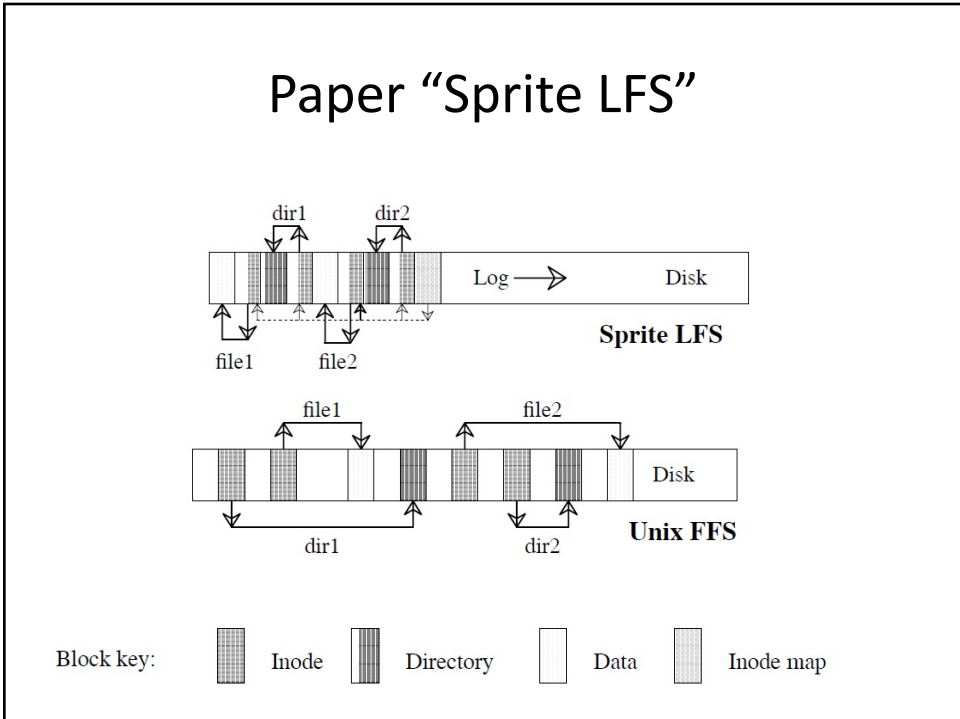
Fall 2021

1


Paper “Sprite LFS”

- Does Sprite LFS improve READ or WRITE performance?
- What is the biggest challenge of a log-structured file system?
- Does Sprite LFS focus on large or small files?
- How does Sprite LFS keep reading performance acceptable?

2



3



4

Paper “Sprite LFS”

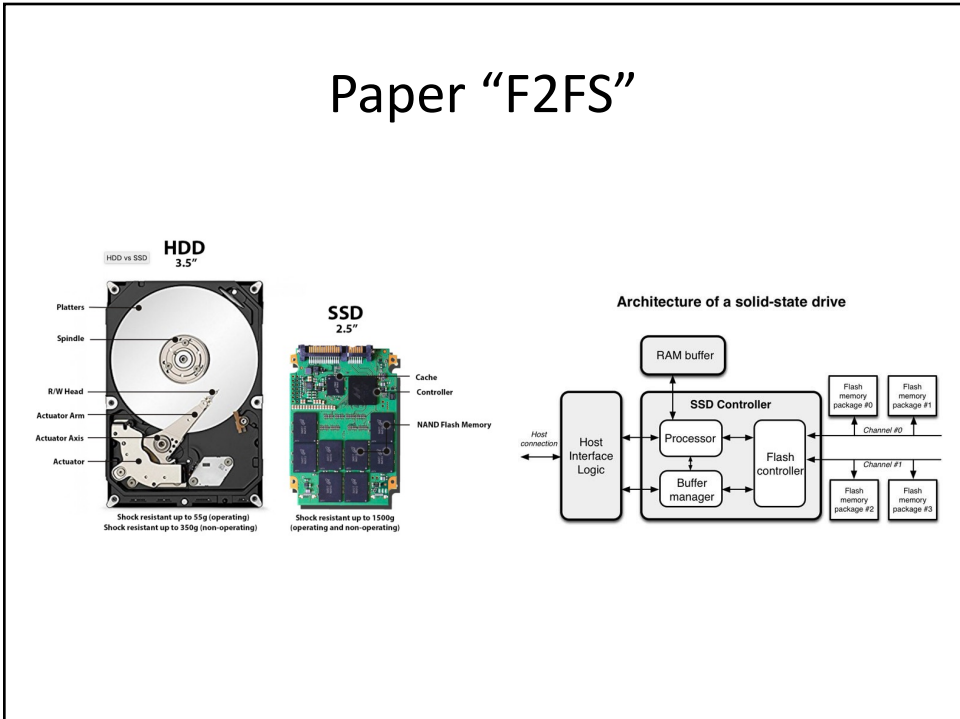
- Segments, segment cleaning, segment summary block
 - When should the segment cleaner execute?
 - How many segments should it clean at a time?
 - Which segments should be cleaned?
 - How should the live blocks be grouped when they are written out?
- Cost-benefit policies
 - Cold segments should be cleaned at high utilization
 - Hot segments should be cleaned at low utilization
- Checkpoints
- Experimentation; metrics (“write cost”); microbenchmarks; pros/cons; overheads

5

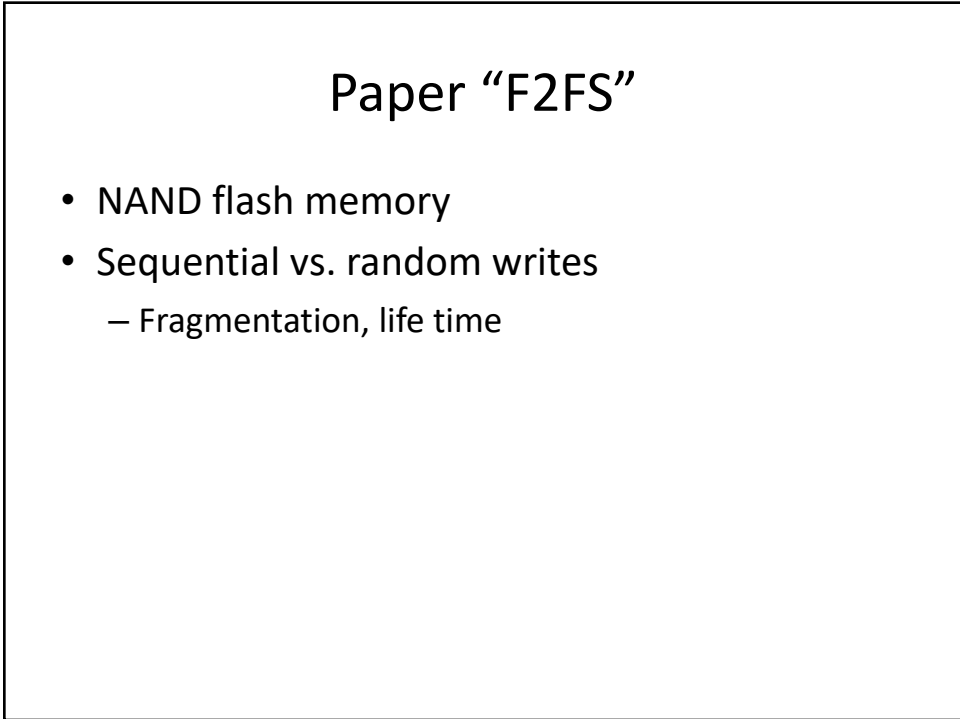
Summary “Sprite LFS”

- Improve write performance
- Crash recovery
- Concept of segments
- Segment cleaning (garbage collection)
- Pros & cons?

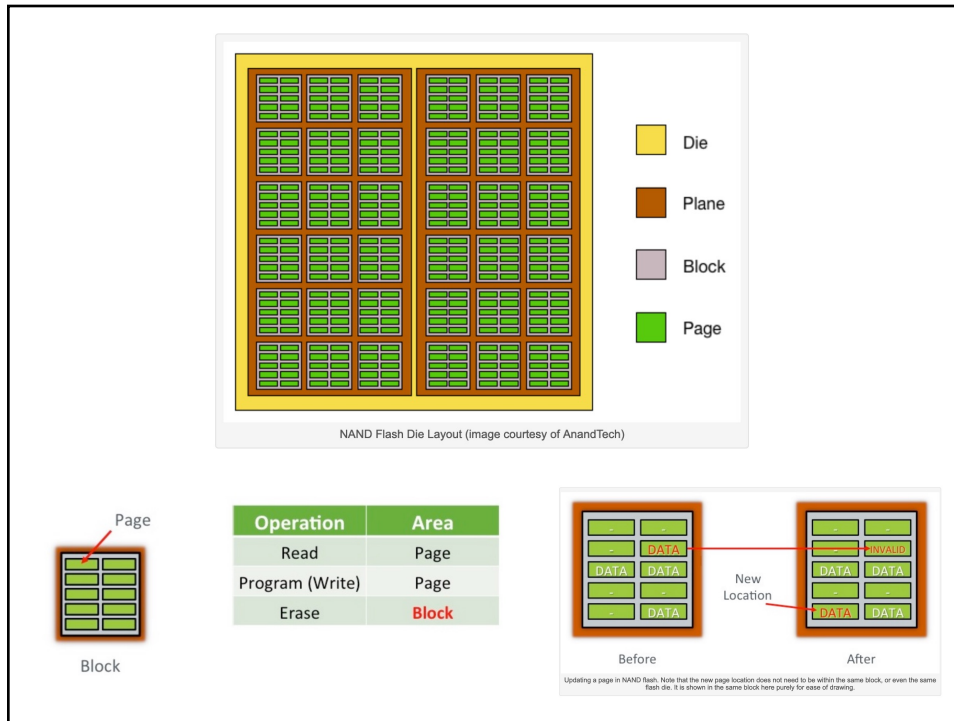
6



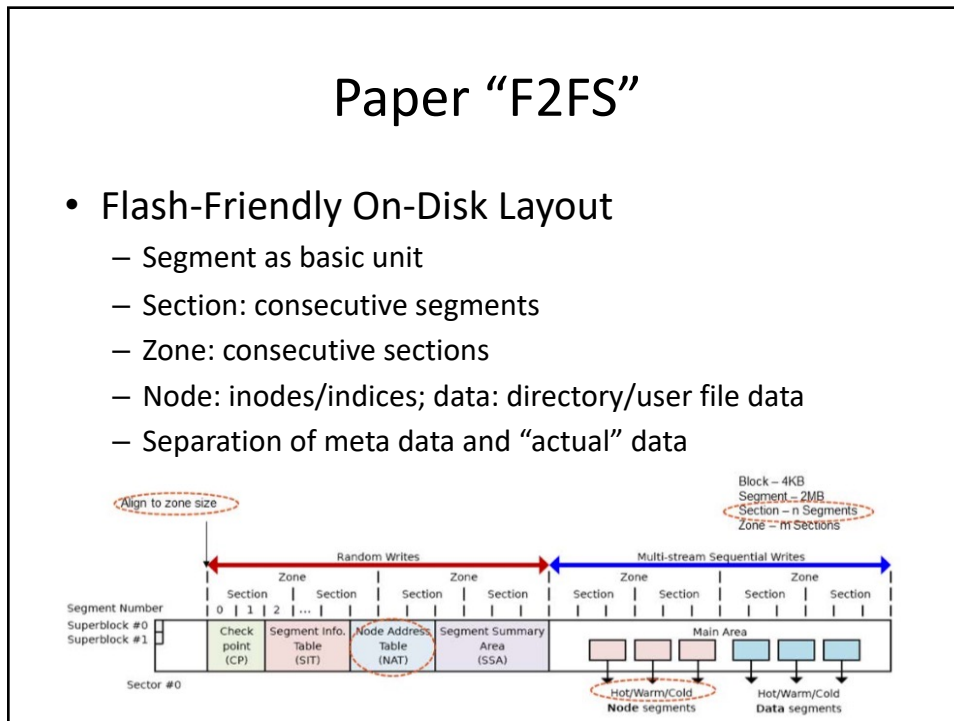
7



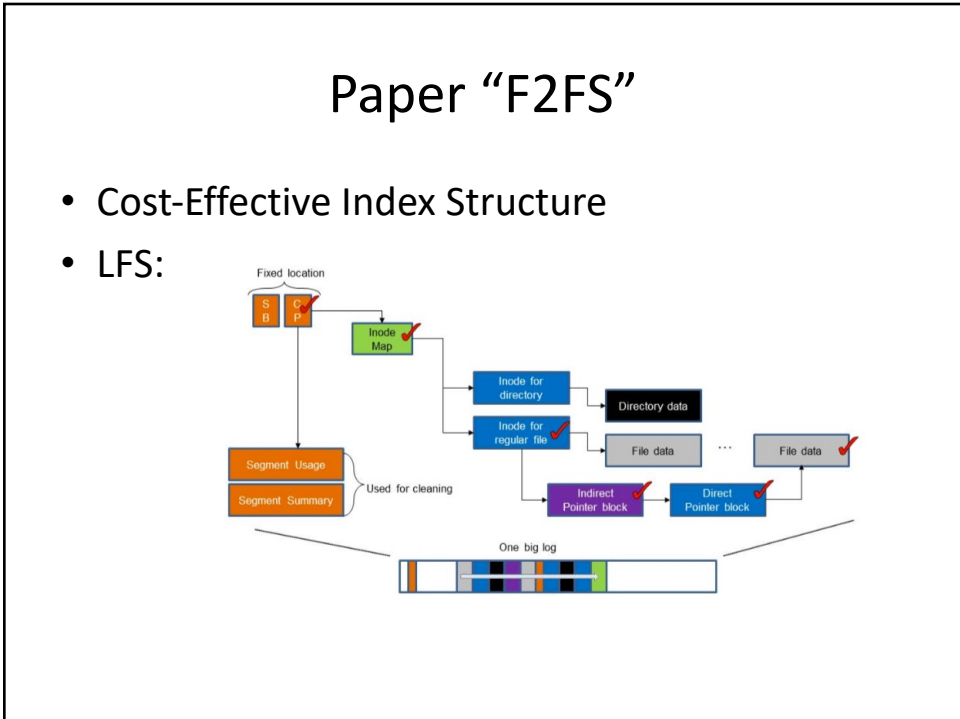
8



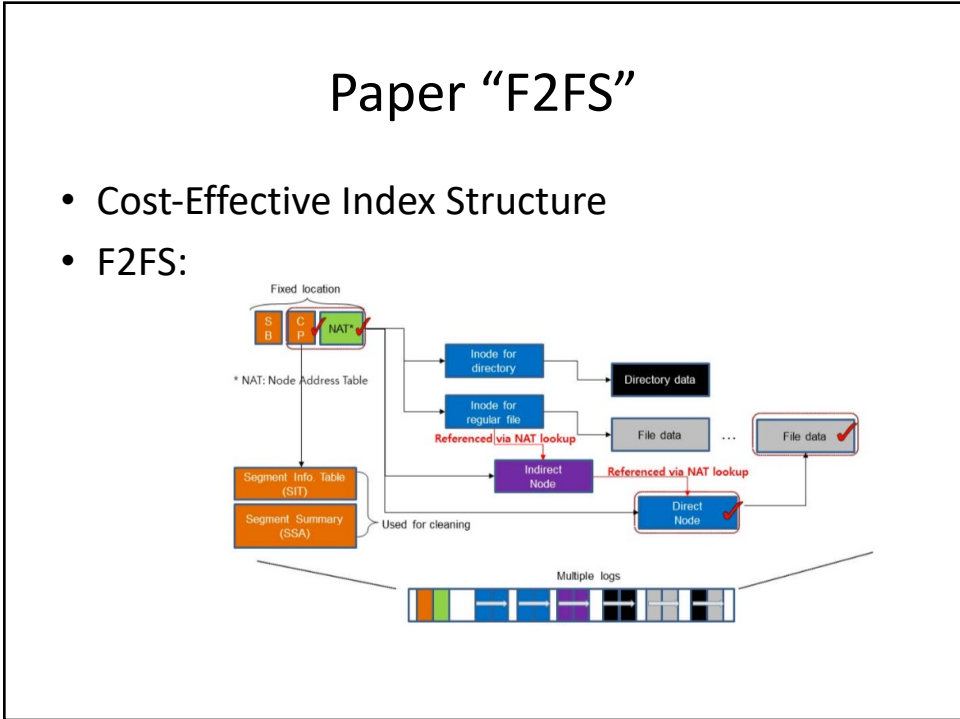
9



10



11



12

Paper “F2FS”

- Wandering tree problem
 - LFS: when a file data is updated and written to the end of log, its **direct** pointer changes, its **indirect pointer block** is also updated, and upper index structures (inode, inode map, checkpoint block) are also changed

13

Paper “F2FS”

- Multi-Head Logging
 - Data temperature classification: hot, warm, cold
 - Six logging segments
- Cleaning
 - Section-level; foreground (need more sections) and background (periodic kernel thread)
 - Greedy for foreground, cost-benefit for background
- Adaptive Logging
 - Switch between normal and threaded logging

14