CSE 451: Operating Systems
Winter 2017

Module 19
File System Summary

Mark Zbikowski
mzbik@cs.washington.edu
Allen Center 476

UFS

Hardware Device

inodes

data blocks

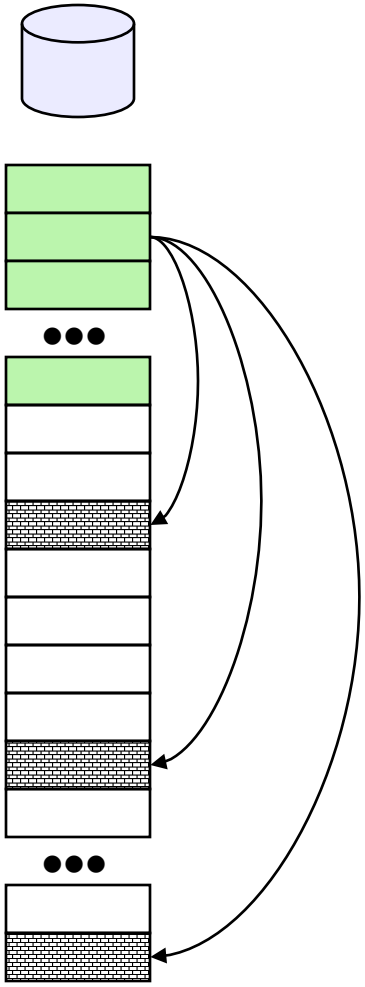

Low throughput addressed by:

- larger blocks
- cylinder groups
- aggressive caching
- hardware awareness

FFS

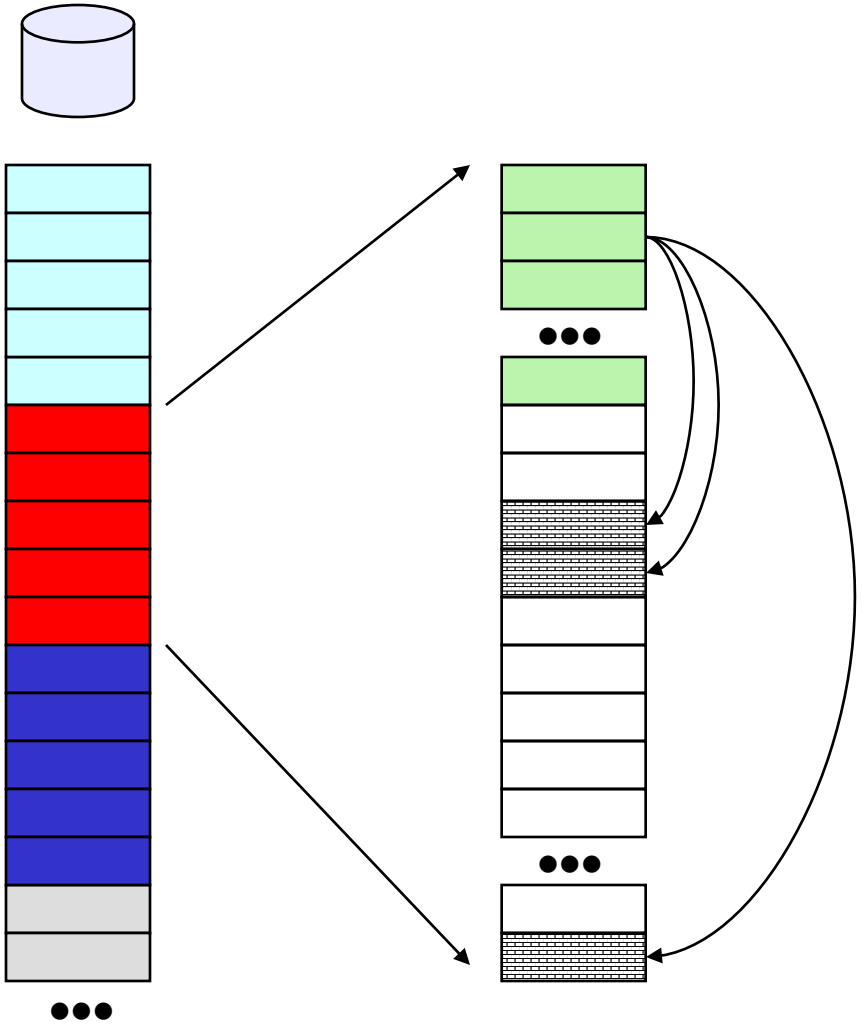

JFS

Long post-crash boot times addressed by:

- transactional journal of changes
- propagated back to "real" file system asynchronously

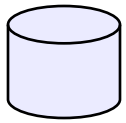

Example: file creation

- …
- Start t
- Alloc inode 1067
- Write inode 1067 w/ [data]
- Write block 22731 w/ [data]
- Commit t
- ….

App requests

Main memory cache

Journal + file system

If data block updates are not journaled, after a crash files may have garbage blocks

(To be clear: FS and FFS have a cache too – I just didn't draw it.)

Write throughput addressed by:

- the file system is a log

LFS

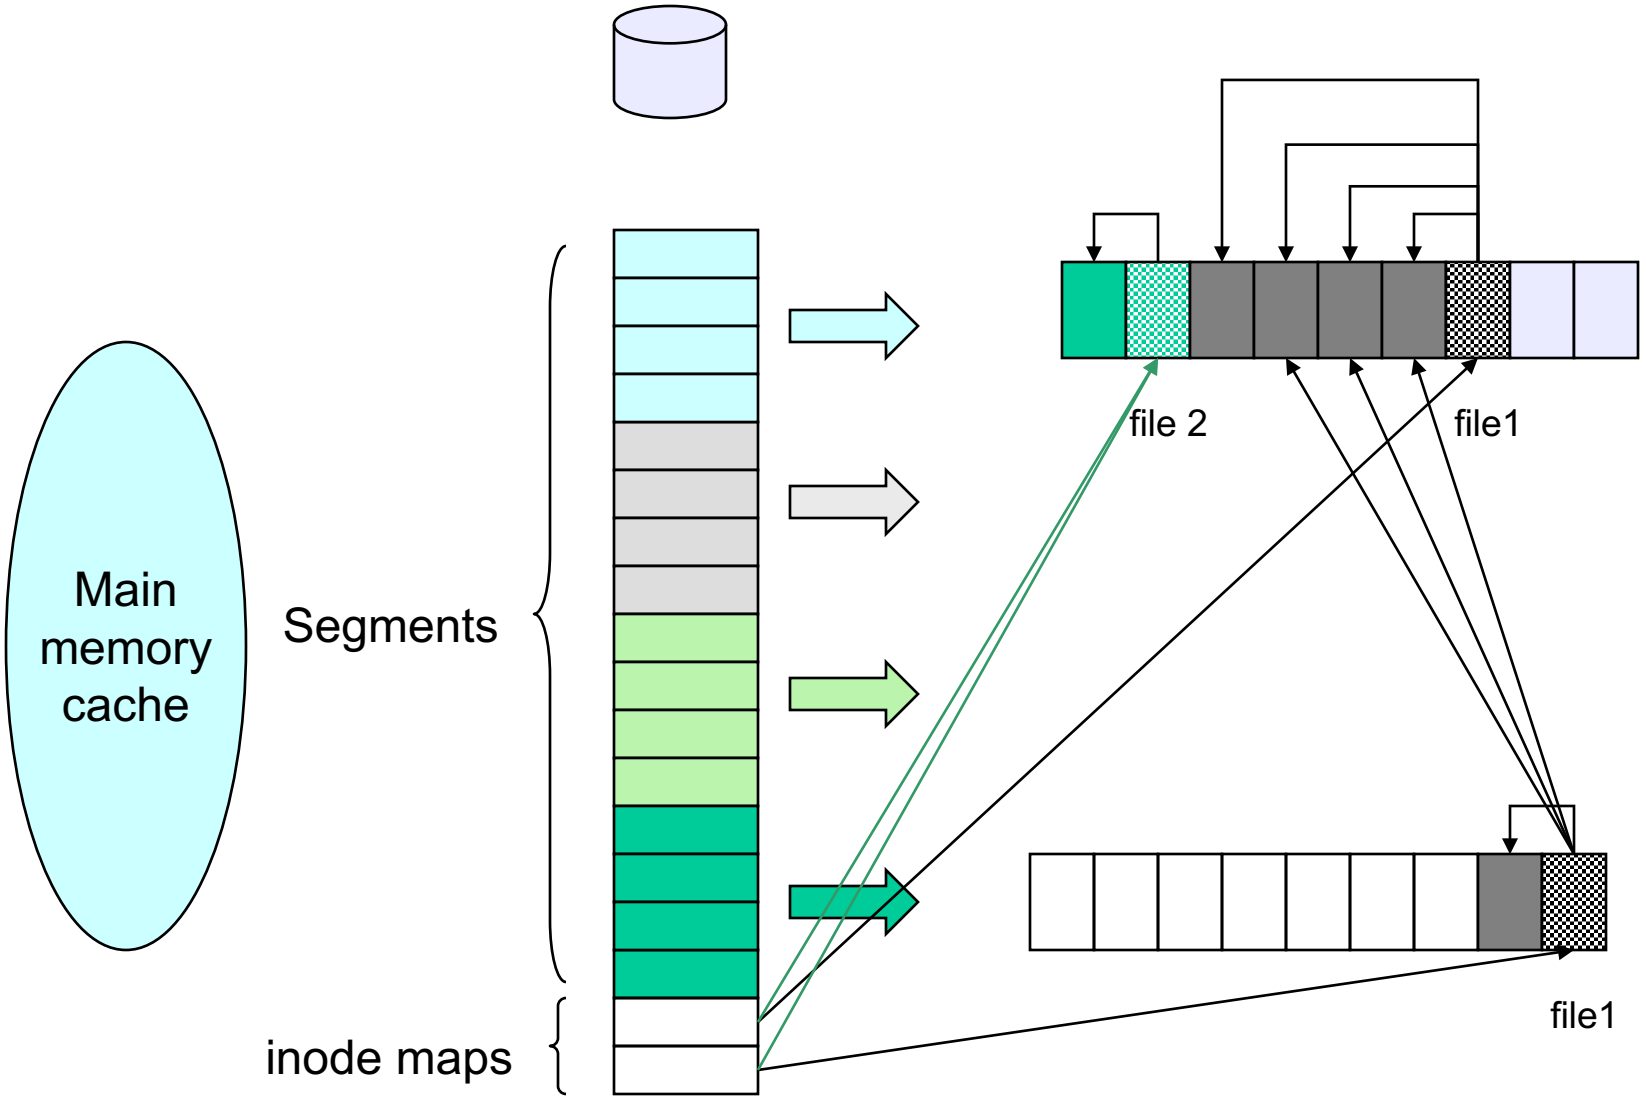

Supporting Multiple File Systems: vfs