

CS 137: File Systems

Class Overview

Today's Topics

- ▶ Purpose of class
- ▶ How class will be run
- ▶ Project
- ▶ Sources of filesystem papers
- ▶ Early reading
- ▶ Introduction to disk technology

Class Purpose

- ▶ Understand how filesystems work
- ▶ Review current research in filesystems
- ▶ Go away with graduate-level understanding

Class Purpose

- ▶ Understand how filesystems work
- ▶ Review current research in filesystems
- ▶ Go away with graduate-level understanding
- ▶ Get me to read good FS papers!

Class Mechanics

- ▶ Begin with general information on disk drives, SSDs, file systems
- ▶ Rest of term will be reading & discussing papers
- ▶ Early papers will be assigned by me
 - ▶ Give you background
 - ▶ Discuss in class
- ▶ See class calendar on Web site:
`http://www.cs.hmc.edu/~geoff/cs137`
- ▶ Later papers chosen by you
 - ▶ Goal is to have fun, learn lots
 - ▶ You will take turns leading discussion

Filesystem Homework

- ▶ 20% of grade
- ▶ We will use FUSE as a development framework
 - ▶ Frees you from kernel development
 - ▶ Otherwise quite similar to “real thing”
- ▶ First assignment: “Hello, world” filesystem
- ▶ Assignments 2 & 3: FAT filesystem

Class Project

- ▶ Written/oral report on some aspect of filesystems research
 - ▶ E.g. survey paper on RAID technology
- ▶ Filesystem development project
- ▶ Filesystem measurement project
- ▶ But I'm open to other ideas & suggestions
- ▶ 70% of grade
 - ▶ Other components: homework (20%), general participation during term (10%)

Where to Find FS Papers

- ▶ Specialized FS conferences
 - ▶ File and Storage Technology (FAST)
 - ▶ IEEE Mass Storage Conference (MassStor)
 - ▶ SIGOPS International Systems and Storage Conference (SYSTOR)
 - ▶ ACM/Usenix HotStorage
- ▶ Supercomputing conferences
 - ▶ IEEE High Performance Distributed Computing
 - ▶ Supercomputing

Where to Find Papers (cont'd)

- ▶ Filesystems are part of operating systems
 - ▶ So big OS conferences have FS papers
 - ▶ Symposium on Operating Systems Principles (SOSP)
 - ▶ Operating Systems Design & Implementation (OSDI)
 - ▶ Usenix Annual Technical Conference
 - ▶ Eurosys
- ▶ Important journals (current & older stuff)
 - ▶ ACM Transactions on Storage
 - ▶ ACM Transactions on Computer Systems
 - ▶ Communications of the ACM
 - ▶ IEEE Computer


Where to Find Papers (cont'd)

- ▶ Database conferences
 - ▶ ACM SIGMOD
 - ▶ Very Large Databases (VLDB)
- ▶ Cloud conferences (e.g., HotCloud)
- ▶ Architectural Support for Programming Languages and Operating Systems (ASPLOS)
- ▶ Sometimes architecture, networking, applications conferences
- ▶ Random other places—list is constantly changing


First Papers We'll Read

- ▶ How nasty disks really are (Ruemmler & Wilkes; Anderson; Patterson et al)
- ▶ Original Unix file system (for elegance)
- ▶ BSD Fast Filesystem (for speed)
- ▶ FAT32 (for ugliness) and NTFS (for breadth)
- ▶ NFS and Coda (for networking)
- ▶ LFS and WAFL (for influence)

Disks In 2 Minutes



Disks In 7 More Minutes



A Running Disk



Slowing That Down

THIS IS A 500GB HARD DRIVE.

Disks From the Beginning

(To be done on the board)

- ▶ Head/platter arrangements
- ▶ Motion technology
- ▶ Winchester drives
- ▶ Sectors and gaps
- ▶ Sector alignment and timing
- ▶ Encodings and ECC
- ▶ General block layout
- ▶ Modern complexities
- ▶ Shingling

Shouting in the Datacenter

