CS 135: 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/cs135`
- Later papers chosen by you
  - Goal is to have fun, learn lots
  - You will take turns leading discussion

Filesystem Homework

- Tentatively 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
- Possibility of more complex FS later in term

Class Project

- 70% of grade
- Other component: general participation during term
- Written/oral report on some aspect of filesystems research
  - E.g. survey paper on RAID technology
  - But I'm open to ideas & suggestions
    - Including non-paper ideas

Where to Find FS Papers

- Specialized FS conferences
  - File and Storage Technology (FAST) (recent)
  - IEEE Mass Storage Conference (MassStor)
- 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
  - Especially true over 5 years ago
  - Symposium on Operating Systems Principles (SOSP)
  - Operating Systems Design & Implementation (OSDI)
  - Usenix Annual Technical Conference
- Important journals (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)
- Sometimes architecture, networking, applications conferences
- Random other places

First Papers We'll Read

- How ugly 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)

Disk Basics

(To be done on the board)

- Head/platter arrangements
- Motion technology
- Winchester drives
- Sectors and gaps
- Sector alignment
- Encodings and ECC
- General block layout