

Intro Lecture

I. Administrative Matters

Instructors: Eric Brewer and Joe Hellerstien

Eric Brewer

- o PhD MIT, 1994
- o Internet Systems, Mobile computing, Security, Parallel computing
- o Founded Inktomi, Federal Search Foundation
- o Contact: brewer@cs.berkeley.edu

Joe Hellerstien

- o PhD, Wisconsin, 1995
- o Databases, declarative networking
- o Contact: hellerstein@cs.berkeley.edu

Course prerequisite: no **entrance exam** this year, but please review undergrad material

Traditional goals of the course:

- o Introduce a variety of current OS research topics
- o Teach you how to do OS research

New goals:

- o Common foundation for OS and database research
- o Motivation: OS and DB communities historically separate, despite much overlap in goals and ideas. We often use independent vocabulary and have largely separate “classic” papers that the other community typically hasn’t read (especially OS people not reading DB papers, since a DB is “just an application”)
- o Part 1 of a year-long advanced systems course
- o 262A satisfies software breadth requirement by itself

Research = analysis & synthesis

- o Analysis: understanding others’ work - both what’s good AND what’s bad.
- o Systems research isn’t cut-and-dried: few “provably correct” answers.
- o Synthesis: finding new ways to advance the state of the art.

Lecture 1

2 parts to course:

- o literature survey
- o term project

Will not cover basic material.

Analysis: literature survey: read, analyze, criticize papers.

- o All research projects fail in some way.
- o Successful projects: got some interesting results anyway.

In class: lecture format with discussion

Synthesis: do a small piece of real research.

- o Suggested projects handed out in 3-4 weeks
- o Teams of 2-3
- o Poster session and “conference paper” at the end of the semester
- o Usually best papers make it into a real conference (with extra work)

Class preparation:

- o Reading papers is hard, esp. at first.
- o **Read before class.**

Homework: *brief* summary/paper

- o 1/2 page or less
- o 2 most important things in paper
- o 1 major flaw

Course topics: (much of systems design is about sharing)

- o File systems
- o Virtual memory
- o Concurrency, scheduling, synchronization
- o Communication
- o Multiprocessors
- o Distributed Systems
- o Transactions, recovery, & fault-tolerance

Lecture 1

- o Protection & security
- o OS structure
- o “Revealed truth” — overall principles

Grad class \Rightarrow material may be controversial (hopefully!)

Lecture format: cover 1-2 papers plus announcements

Grading: 50% project paper, 15% project demo, 25% midterm, 10% paper summaries

- o can miss 3 summaries without penalty, no reason needed or wanted
- o Summaries: < 1/2 page, at least one criticism, summary should be at most half

Reading list:

- o No textbook
- o Almost everything is online, otherwise I’ll distribute copies in class.
- o “Warm up” paper: “UNIX timesharing”
- o Reading for next time: System R paper (in the Red Book)

II. The UNIX Time-Sharing System

“Warm-up” paper. Material should mostly be a review.

Classic system and paper: described almost entirely in 10 pages.

Key idea: elegant combination of a few concepts that fit together well.

System features:

- o time-sharing system
- o hierarchical file system
- o device-independent I/O
- o shell-based, tty user interface
- o filter-based, record-less processing paradigm

Version 3 Unix:

- o ran on PDP-11’s
- o < 50 KB
- o 2 man-years to write
- o written in C

Lecture 1

File System:

- o ordinary files (uninterpreted)
- o directories (protected ordinary files)
- o special files (I/O)

Directories:

- o root directory
- o path names
- o rooted tree
- o current working directory
- o back link to parent
- o multiple links to ordinary files

Special files:

- o uniform I/O model
- o uniform naming and protection model

Removable file systems:

- o tree-structured
- o *mount*'ed on an ordinary file

Protection:

- o user-world, RWX bits
- o set-user-id bit
- o super user is just special user id

Uniform I/O model:

- o open, close, read, write, seek
- o other system calls
- o bytes, no records

File system implementation:

- o table of i-nodes
- o path name scanning

Lecture 1

- o mount table
- o buffered data
- o write-behind

I-node table:

- o short, unique name that points at file info.
- o allows simple & efficient fsck
- o can't handle accounting issues

Processes and images:

- o text, data & stack segments
- o process swapping
- o pid = fork()
- o pipes
- o exec(file, arg1, ..., argn)
- o pid = wait()
- o exit(status)

The Shell:

- o cmd arg1 ... argn
- o stdio & I/O redirection
- o filters & pipes
- o multi-tasking from a single shell
- o shell is just a program

Traps:

- o hardware interrupts
- o software signals
- o trap to system routine