

# **ECE 598 – Advanced Operating Systems Lecture 13**

Vince Weaver

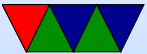
<http://www.eece.maine.edu/~vweaver>

[vincent.weaver@maine.edu](mailto:vincent.weaver@maine.edu)

17 March 2015

# Announcements

- Homework #3 grades were sent out
- Homework #4 posted soon?



# Midterm Review

- OS benefits

We've seen this before. Abstraction, multi-tasking, security, portability, manage resources.

Not really: smoother UI? Hide details from users?

- OS downsides:

overhead, real-time

- Serial: 115200 8N1

- HW flow control: separate wire to indicate DTS/RTS,



not sent over the TX/RX lines. Some mixing up of UART FIFO start/stop and transmission start/stop

- SW flow control: ASCII chars sent down the line. How do you handle if you want to send those chars? Control-S? Story about +++ on modems on old days?
- Abstracting away the Pi2
- Avoiding read/modify/write. Mostly because it's two slow operations plus some somewhat complex bit manipulation.




Can be a race if we have to write existing values back exactly but they change somehow in between (for example, an interrupt is messing with things).

- Re-entering kernel – syscall or interrupt
- Interrupt code, has own stack, so not a problem user messes with it

That's why it does. On some arch other probs, stacks need to be aligned and can crash if IRQ happens and has to use user stack

- IRQ vs FIQ. Faster, as saves more register state so you



don't have to save it out to slow memory. Also higher priority. Only one of them.

- Memory. How much is free? 44kb.

I guess diagram could have been clearer, growing up.

First fit, 0x3

Best fit 0xC

Less fragmentation

malloc() does have realloc(), but the code has to do this explicitly. It's not necessarily easy on C to update every pointer address, especially as you can cast back and forth to int.



- Virtual Mem

Can have more mem than phys, each gets its own view,  
memory protection

A bit unclear. Meant to be easy. Just maps to separate  
phys addresses

Reasons to be shared? Shared memory, shared libraries,  
copy-on-write



# HW3 Review

- Comment your code!
- Enabling blinking, not too bad.
- Command line parsing
- Syscall. Not too bad either. How it works under Linux.
- IRQ/FIQ = see exam review
- Entering user mode, write CSPR. Return to kernel,





interrupt or syscall.

- `int a = 0; stack`  
`static int b = 1; data`  
`static int c = 0; bss`  
`int *d=malloc(sizeof(int))`

```
main=0x400450  
a=0x7fff98145444  
b=0x6009d8  
c=0x6009e0  
d=0x7fff98145448 *d=0x13da010
```



# Filesystems

- What is a file?

Unix is just a stream of bytes. That's not necessarily the only thing. Old systems files were just 80-column punch card images.

- How do you store it?

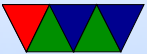
In a “file”?

Everything in root directory? Hiearchical? Database?

- Why not just load everything into memory? Too big?



Share with other processes? Persistent across reboots?



# Filesystems metadata

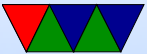
- Name?
  - UNIX – anything that is not / or NUL.  
What about "-rf" or \* or ?
  - Windows, more restrictions. Things like CON and COM and PTR. Can cause problems.
  - DOS famously had 8.3 filenames.
  - Mac used to use : as directory separator.
- Case sensitive? Bob or bob or BoB?



- Special files?
  - Regular files
  - Directories
  - Char devices, block devices
  - Links (hard, soft)
- Attributes
  - Permissions (read, write, execute)
  - Owner (user, group)
  - Access time (atime, ctime, mtime)
  - Current size



- Locks
- Hidden
- Immutable / System
- Extended attributes / capabilities



# File Operations

- Create
- Delete (Unlink)
- Open
- Close
- Read
- Write



- Append
- Seek
- Get attributes
- Set attributes
- Rename/move
- Memory Map





# File Descriptors

- On Linux/Unix "everything is a file"
- When open a file or object, get a number that indexes into a table, each referring to a file.
- Low-level syscalls mostly operate on file descriptors



# Directories/Folders

- Root directory
- Hierarchical
- Path names. /, ., ..
- Operations
  - create
  - delete
  - opendir



- closedir
- readdir
- rename
- link
- unlink



# Filesystems

- Disks divided into partitions
- Often a MBR (master boot record) and partition table
- Then individual filesystems
- Often first is called superblock, containing all master info.
- Some sort of free list, saying what areas are free (bitmap or pointers)



- inodes, an array of data structures containing master info for each file (and if file is small, contents of file)
- root directory entry
- directory layout
- file data

