

CS162 Operating Systems and Systems Programming Lecture 3

Processes (con't), Fork, Introduction to I/O

September 2nd, 2015
Prof. John Kubitowicz
<http://cs162.eecs.Berkeley.edu>

Recall: Four fundamental OS concepts


- Thread
 - Single unique execution context
 - Program Counter, Registers, Execution Flags, Stack
- Address Space w/ Translation
 - Programs execute in an *address space* that is distinct from the memory space of the physical machine
- Process
 - An instance of an executing program is a *process* consisting of an address space and one or more threads of control
- Dual Mode operation/Protection
 - Only the "system" has the ability to access certain resources
 - The OS and the hardware are protected from user programs and user programs are isolated from one another by *controlling the translation* from program virtual addresses to machine physical addresses

9/2/15

Kubitowicz CS162 ©UCB Fall 2015

Lec 3.2

Recall: give the illusion of multiple processors?



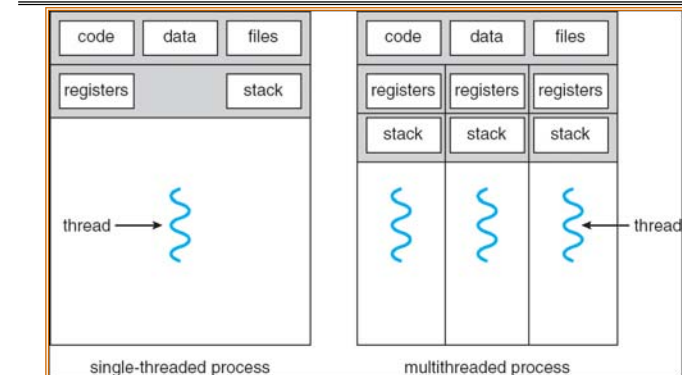
- Assume a single processor. How do we provide the *illusion* of multiple processors?
 - Multiplex in time!
 - Multiple "virtual CPUs"
- Each virtual "CPU" needs a structure to hold:
 - Program Counter (PC), Stack Pointer (SP)
 - Registers (Integer, Floating point, others...?)
- How switch from one virtual CPU to the next?
 - Save PC, SP, and registers in current state block
 - Load PC, SP, and registers from new state block
- What triggers switch?
 - Timer, voluntary yield, I/O, other things

9/2/15

Kubitowicz CS162 ©UCB Fall 2015

Lec 3.3

Single and Multithreaded Processes



- Threads encapsulate concurrency: "Active" component
- Address spaces encapsulate protection: "Passive" part
 - Keeps buggy program from trashing the system
- Why have multiple threads per address space?

9/2/15

Kubitowicz CS162 ©UCB Fall 2015

Lec 3.4

Running Many Programs

- We have the basic mechanism to
 - switch between user processes and the kernel,
 - the kernel can switch among user processes,
 - Protect OS from user processes and processes from each other
- Questions ???
 - How do we represent user processes in the OS?
 - How do we decide which user process to run?
 - How do we pack up the process and set it aside?
 - How do we get a stack and heap for the kernel?
 - Aren't we wasting a lot of memory?
 - ...

9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.5

Process Control Block

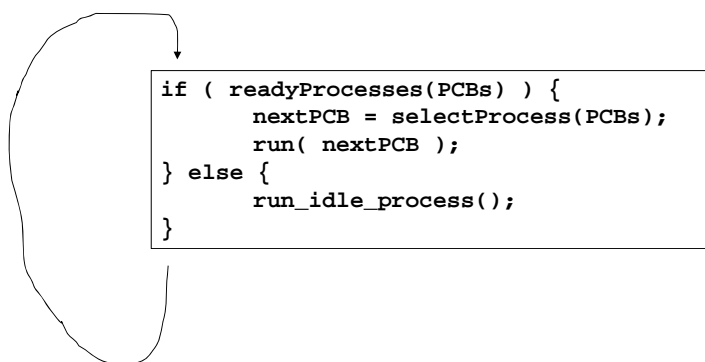
- Kernel represents each process as a process control block (PCB)
 - Status (running, ready, blocked, ...)
 - Register state (when not ready)
 - Process ID (PID), User, Executable, Priority, ...
 - Execution time, ...
 - Memory space, translation, ...
- Kernel Scheduler maintains a data structure containing the PCBs
- Scheduling algorithm selects the next one to run

9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.6

Scheduler



- Scheduling: Mechanism for deciding which processes/threads receive the CPU
- Lots of different scheduling policies provide ...
 - Fairness or
 - Realtime guarantees or
 - Latency optimization or ..

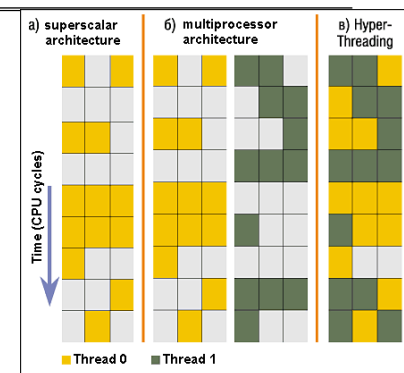
9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.7

Simultaneous MultiThreading/Hyperthreading

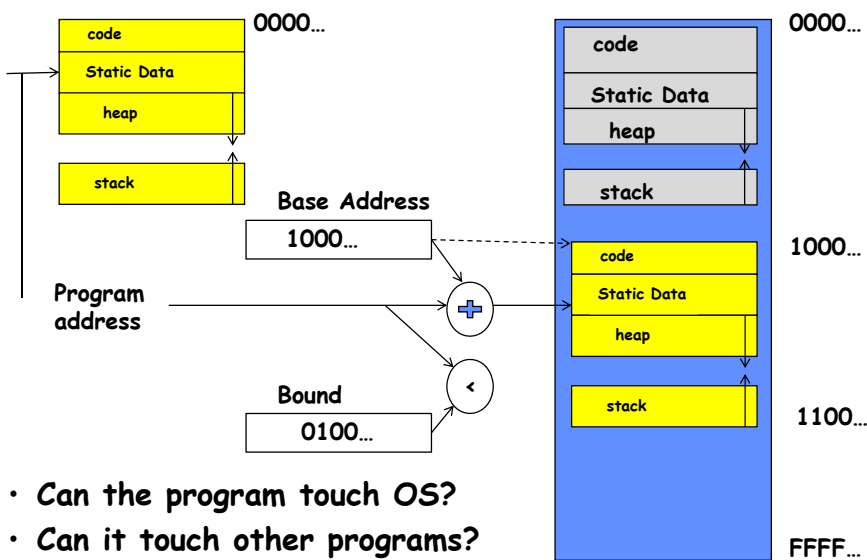
- Hardware technique
 - Superscalar processors can execute multiple instructions that are independent.
 - Hyperthreading duplicates register state to make a second "thread," allowing more instructions to run.
- Can schedule each thread as if were separate CPU
 - But, sub-linear speedup!
- Original technique called "Simultaneous Multithreading"
 - <http://www.cs.washington.edu/research/smt/index.html>
 - SPARC, Pentium 4/Xeon ("Hyperthreading"), Power 5



Colored blocks show instructions executed

Lec 3.8

Recall: A simple address translation (B&B)



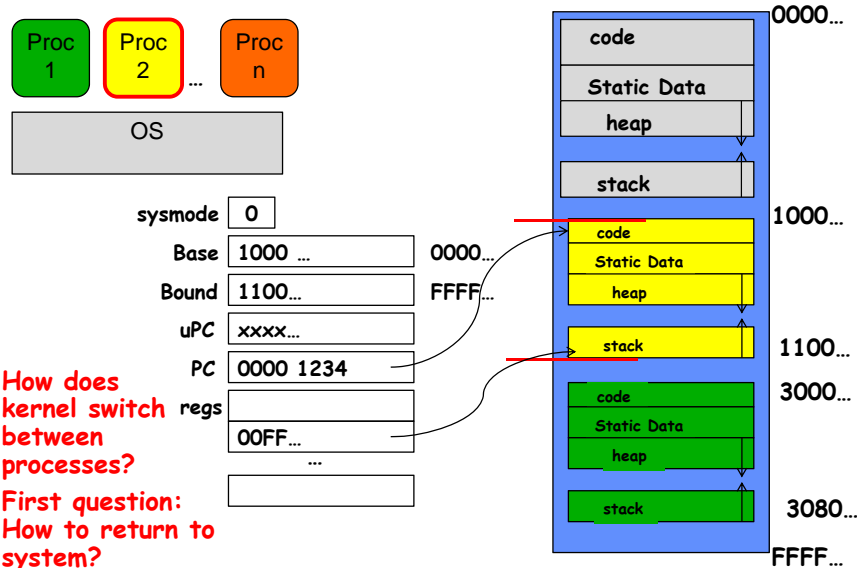
- Can the program touch OS?
- Can it touch other programs?

9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.9

Simple B&B: User code running



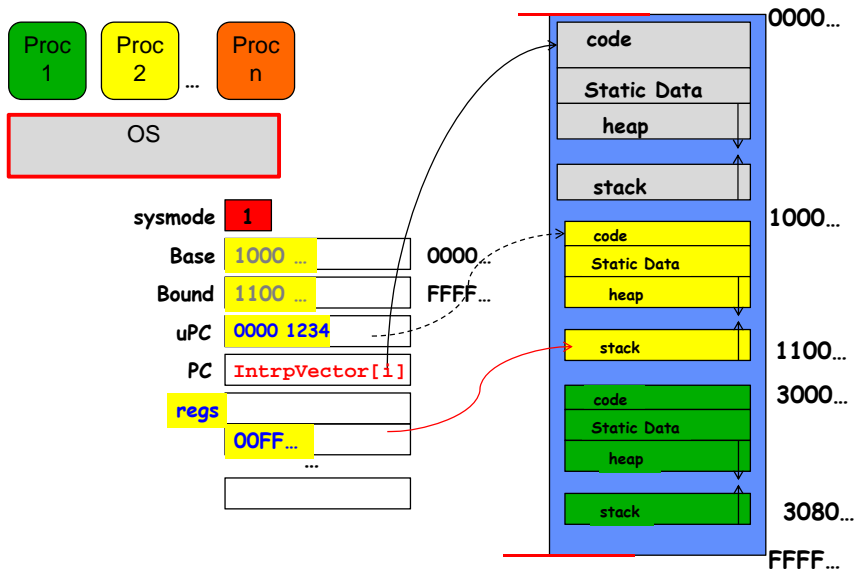
- How does kernel switch between processes?
- First question: How to return to system?

9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.10

Simple B&B: Interrupt

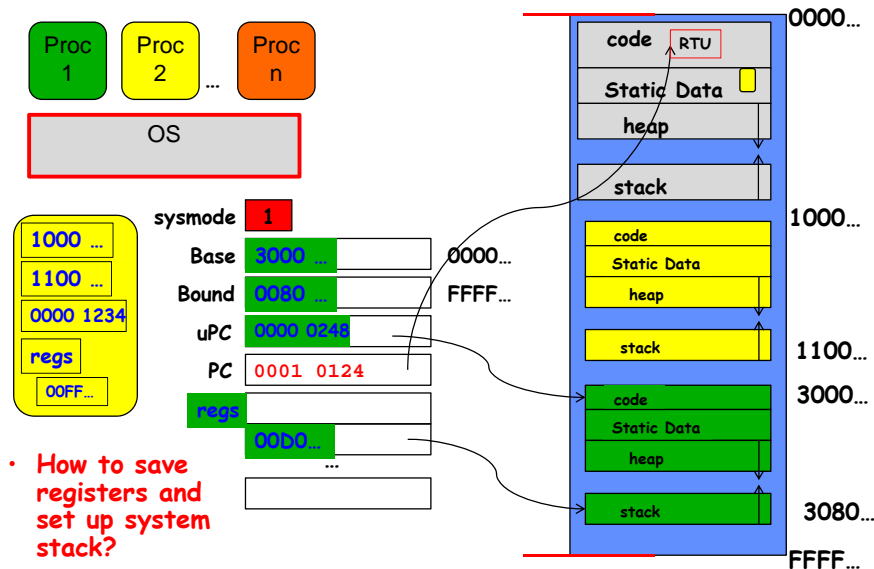


9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.11

Simple B&B: Switch User Process



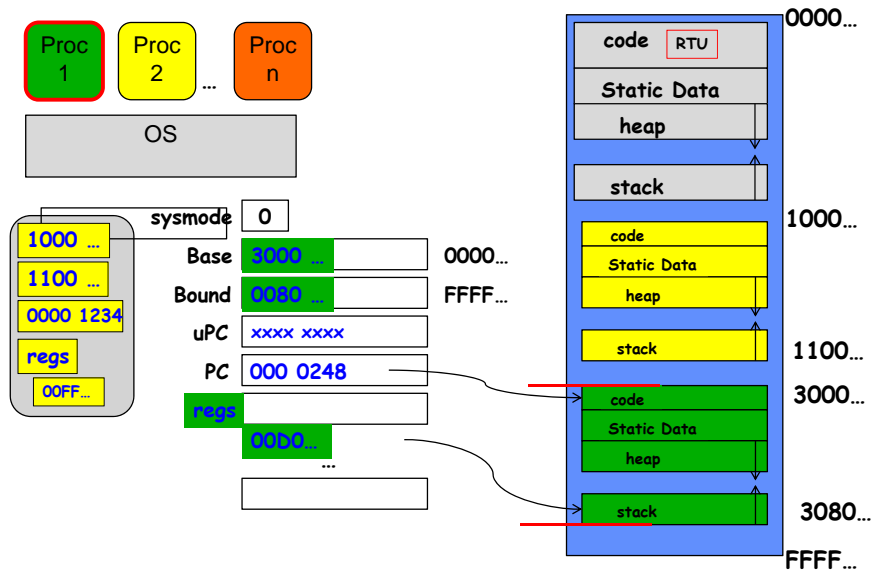
- How to save registers and set up system stack?

9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.12

Simple B&B: "resume"



9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.13

What's wrong with this simplistic address translation mechanism?

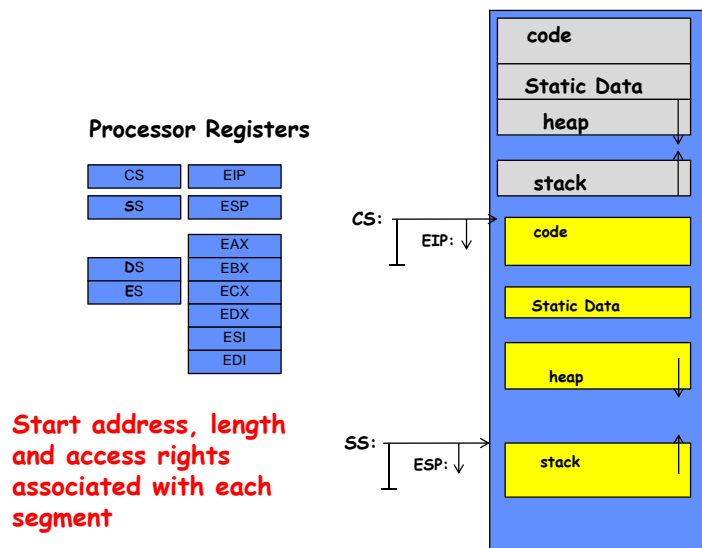
- **Fragmentation:**
 - Kernel has to somehow fit whole processes into contiguous block of memory
 - After a while, memory becomes fragmented!
- **Sharing:**
 - Very hard to share any data between Processes or between Process and Kernel
 - Simple segmentation prevents any memory sharing by its very nature

9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.14

x86 - segments and stacks

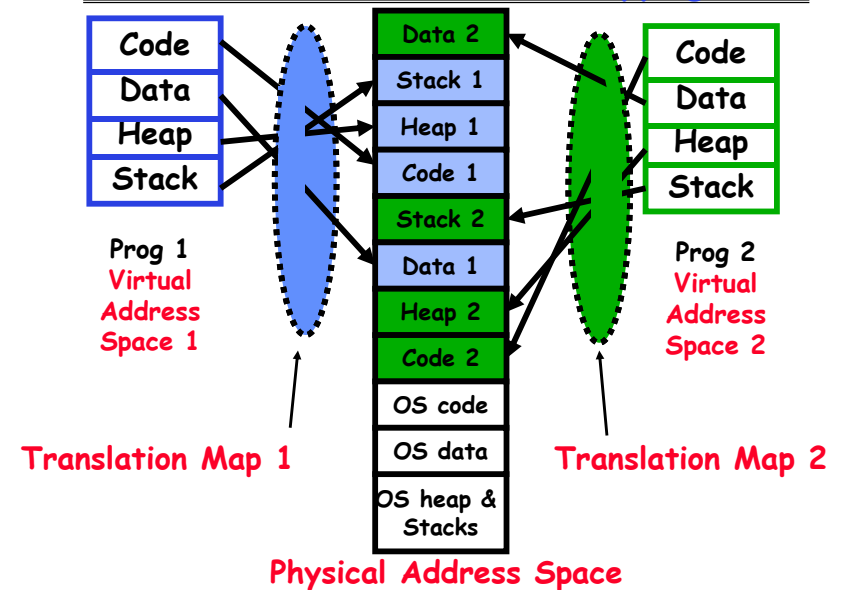


9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.15

Alternative: Address Mapping



9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.16

Administrivia: Getting started

- Kubiawicz Office Hours:
 - 1pm-2pm, Monday/Wednesday
- Homework 0 immediately ⇒ **Due on Monday!**
 - Get familiar with all the cs162 tools
 - Submit to autograder via git
- Should be going to section already!
 - Participation: Get to know your TA!
- **Friday is Drop Deadline!**
- Group sign up form out next week (after drop deadline)
 - Get finding groups ASAP
 - 4 people in a group! Try to keep same section; if cannot make this work, keep same TA

9/2/15

Kubiawicz CS162 ©UCB Fall 2015

Lec 3.17

Administrivia (Con't)

- Conflict between Midterm 2 and EE16A
 - We are thinking of moving Midterm 2 from Wed 11/18
 - Possibilities: Mon 11/23 (my favorite) or Mon 11/16
- Midterm 1 conflicts
 - I know about one problem with Midterm 1 scheduling, and it can be dealt with. Have I missed any others?
- Finals conflicts: We will not be moving the exam or providing makeup finals...
 - I don't know of any current conflicts
 - If you have a significant conflict that you think should cause us to change our policy, let me know now (note that CS186 is not conflicting any more).

9/2/15

Kubiawicz CS162 ©UCB Fall 2015

Lec 3.18

Recall: 3 types of Kernel Mode Transfer

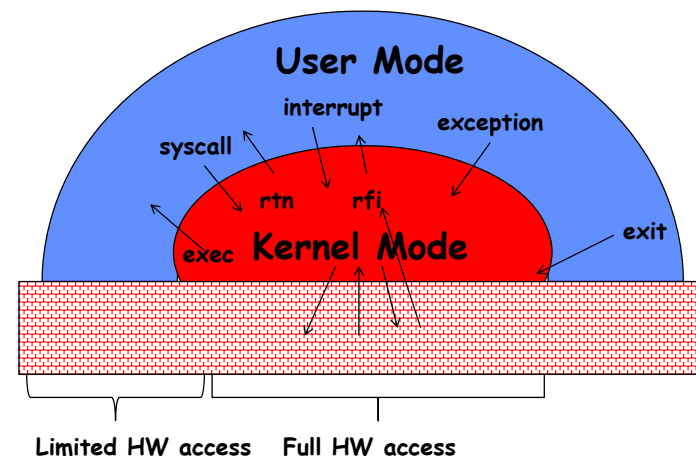
- Syscall
 - Process requests a system service, e.g., exit
 - Like a function call, but "outside" the process
 - Does not have the address of the system function to call
 - Like a Remote Procedure Call (RPC) - for later
 - Marshall the syscall id and args in registers and exec syscall
- Interrupt
 - External asynchronous event triggers context switch
 - eg. Timer, I/O device
 - Independent of user process
- Trap or Exception
 - Internal synchronous event in process triggers context switch
 - e.g., Protection violation (segmentation fault), Divide by zero, ...

9/2/15

Kubiawicz CS162 ©UCB Fall 2015

Lec 3.19

Recall: User/Kernel(Privileged) Mode



9/2/15

Kubiawicz CS162 ©UCB Fall 2015

Lec 3.20

Implementing Safe Kernel Mode Transfers

- **Important aspects:**
 - Separate kernel stack
 - Controlled transfer into kernel (e.g. syscall table)
- Carefully constructed kernel code packs up the user process state and sets it aside.
 - Details depend on the machine architecture
- Should be impossible for buggy or malicious user program to cause the kernel to corrupt itself.

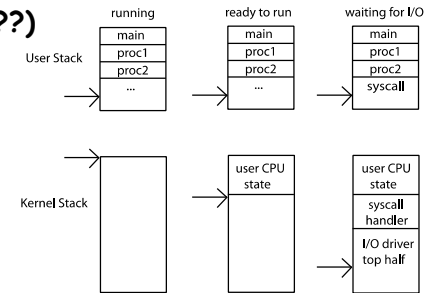
9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.21

Need for Separate Kernel Stacks

- Kernel needs space to work
- Cannot put anything on the user stack (Why?)
- Two-stack model
 - OS thread has interrupt stack (located in kernel memory) plus User stack (located in user memory)
 - Syscall handler copies user args to kernel space before invoking specific function (e.g., open)
 - Interrupts (???)




9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.22

Before

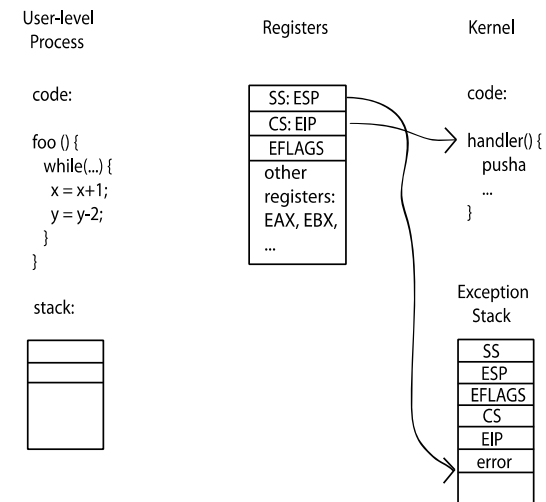


9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.23

During



9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.24

Kernel System Call Handler

- **Vector through well-defined syscall entry points!**
 - Table mapping system call number to handler
- **Locate arguments**
 - In registers or on user(!) stack
- **Copy arguments**
 - From user memory into kernel memory
 - Protect kernel from malicious code evading checks
- **Validate arguments**
 - Protect kernel from errors in user code
- **Copy results back**
 - into user memory

9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.25


Hardware support: Interrupt Control

- **Interrupt processing not be visible to the user process:**
 - Occurs between instructions, restarted transparently
 - No change to process state
 - What can be observed even with perfect interrupt processing?
- **Interrupt Handler invoked with interrupts 'disabled'**
 - Re-enabled upon completion
 - Non-blocking (run to completion, no waits)
 - Pack up in a queue and pass off to an OS thread for hard work
 - » wake up an existing OS thread
- **OS kernel may enable/disable interrupts**
 - On x86: CLI (disable interrupts), STI (enable)
 - Atomic section when select next process/thread to run
 - Atomic return from interrupt or syscall
- **HW may have multiple levels of interrupt**
 - Mask off (disable) certain interrupts, eg., lower priority
 - Certain non-maskable-interrupts (nmi)
 - » e.g., kernel segmentation fault

9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.26



- **Interrupts invoked with interrupt lines from devices**
- **Interrupt controller chooses interrupt request to honor**
 - Mask enables/disables interrupts
 - Priority encoder picks highest enabled interrupt
 - Software Interrupt Set/Cleared by Software
 - Interrupt identity specified with ID line
- **CPU can disable all interrupts with internal flag**
- **Non-maskable interrupt line (NMI) can't be disabled**

9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.27

How do we take interrupts safely?

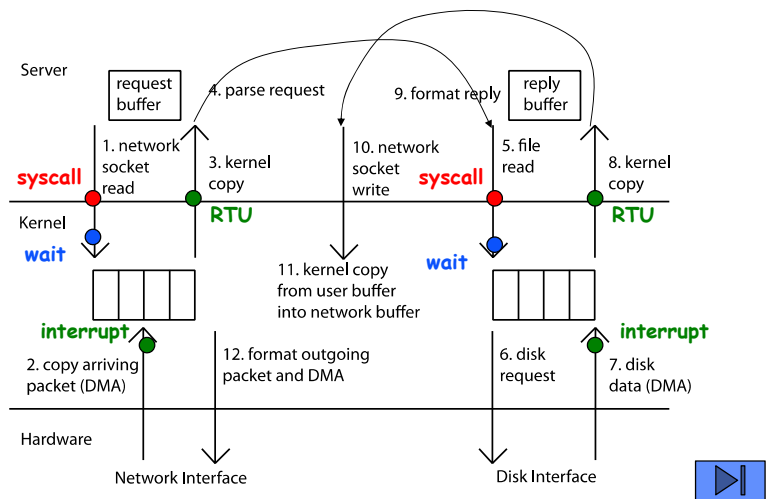
- **Interrupt vector**
 - Limited number of entry points into kernel
- **Kernel interrupt stack**
 - Handler works regardless of state of user code
- **Interrupt masking**
 - Handler is non-blocking
- **Atomic transfer of control**
 - "Single instruction"-like to change:
 - » Program counter
 - » Stack pointer
 - » Memory protection
 - » Kernel/user mode
- **Transparent restartable execution**
 - User program does not know interrupt occurred

9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.28

Putting it together: web server



9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.29

Can a process create a process ?

- **Yes**
 - Unique identity of process is the "process ID" (or pid).
- **Fork() system call creates a *copy* of current process with a new pid**
- **Return value from Fork(): integer**
 - When > 0:
 - » Running in (original) **Parent** process
 - » return value is **pid** of new child
 - When = 0:
 - » Running in new **Child** process
 - When < 0:
 - » Error! Must handle somehow
 - » Running in original process
- **All of the state of original process duplicated in both Parent and Child!**
 - **Memory, File Descriptors (next topic), etc...**

9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.30

fork1.c

```
#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <unistd.h>
#include <sys/types.h>

#define BUFSIZE 1024
int main(int argc, char *argv[])
{
    char buf[BUFSIZE];
    size_t readlen, writelen, slen;
    pid_t cpid, mypid;
    pid_t pid = getpid(); /* get current processes PID */
    printf("Parent pid: %d\n", pid);
    cpid = fork();
    if (cpid > 0) { /* Parent Process */
        mypid = getpid();
        printf("[%d] parent of [%d]\n", mypid, cpid);
    } else if (cpid == 0) { /* Child Process */
        mypid = getpid();
        printf("[%d] child\n", mypid);
    } else {
        perror("Fork failed");
        exit(1);
    }
    exit(0);
}
```

9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.31

UNIX Process Management

- **UNIX fork** - system call to create a copy of the current process, and start it running
 - No arguments!
- **UNIX exec** - system call to *change the program* being run by the current process
- **UNIX wait** - system call to wait for a process to finish
- **UNIX signal** - system call to send a notification to another process

9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.32

fork2.c

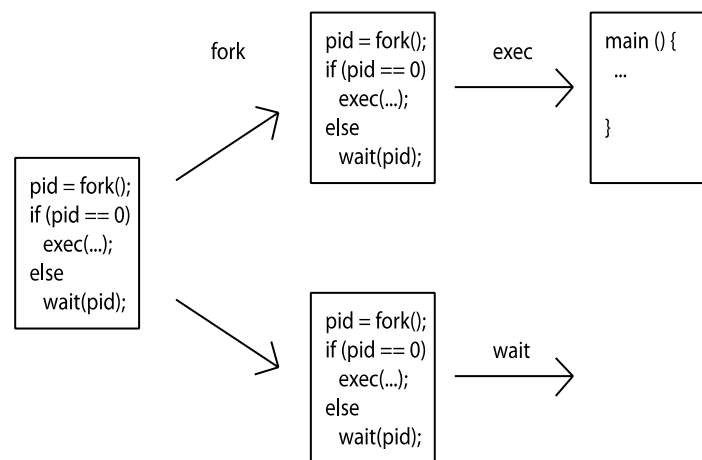
```
int status;
...
cpid = fork();
if (cpid > 0) {
    /* Parent Process */
    mypid = getpid();
    printf("[%d] parent of [%d]\n", mypid, cpid);
    tcpid = wait(&status);
    printf("[%d] bye %d(%d)\n", mypid, tcpid, status);
} else if (cpid == 0) {
    /* Child Process */
    mypid = getpid();
    printf("[%d] child\n", mypid);
}
...
```

9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.33

UNIX Process Management



9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.34

Shell

- A shell is a job control system
 - Allows programmer to create and manage a set of programs to do some task
 - Windows, MacOS, Linux all have shells
- Example: to compile a C program

```
cc -c sourcefile1.c
cc -c sourcefile2.c
ln -o program sourcefile1.o sourcefile2.o
./program
```

HW1

9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.35

Signals - infloop.c

```
#include <stdlib.h>
#include <stdio.h>
#include <sys/types.h>

#include <unistd.h>
#include <signal.h>

void signal_callback_handler(int signum)
{
    printf("Caught signal %d - phew!\n", signum);
    exit(1);
}

int main() {
    signal(SIGINT, signal_callback_handler);

    while (1) {}
}
```

Got top?

9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.36

Process races: fork.c

```

if (cpid > 0) {
    mypid = getpid();
    printf("[%d] parent of [%d]\n", mypid, cpid);
    for (i=0; i<100; i++) {
        printf("[%d] parent: %d\n", mypid, i);
        //      sleep(1);
    }
} else if (cpid == 0) {
    mypid = getpid();
    printf("[%d] child\n", mypid);
    for (i=0; i>-100; i--) {
        printf("[%d] child: %d\n", mypid, i);
        //      sleep(1);
    }
}
    
```

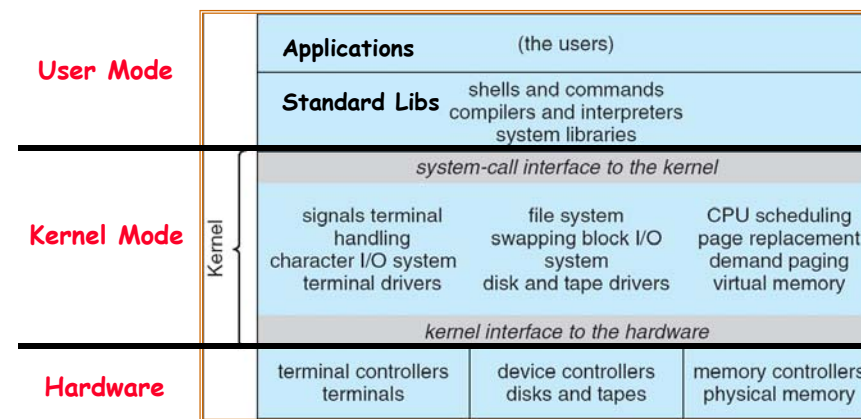
- Question: What does this program print?
- Does it change if you add in one of the sleep() statements?

9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.37

Recall: UNIX System Structure



9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.38

How does the kernel provide services?

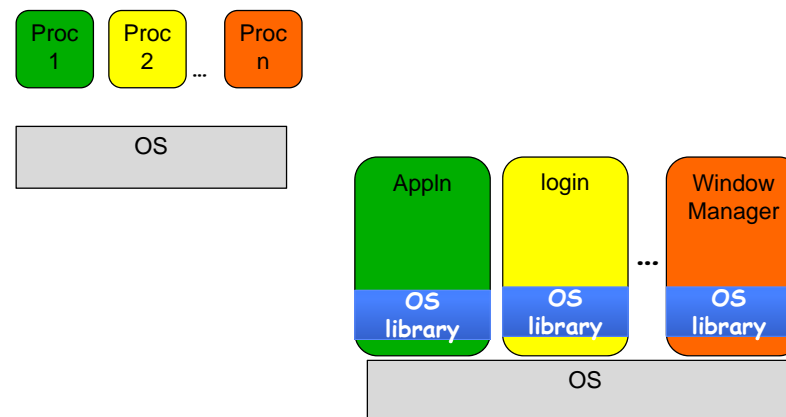
- You said that applications request services from the operating system via *syscall*, but ...
- I've been writing all sort of useful applications and I never ever saw a "syscall" !!!
- That's right.
- It was buried in the programming language runtime library (e.g., libc.a)
- ... Layering

9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.39

OS run-time library

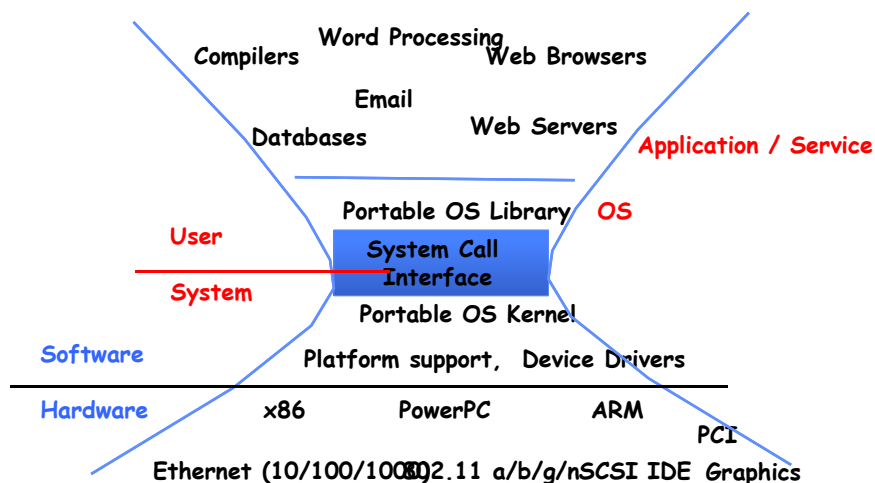


9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.40

A Kind of Narrow Waist



9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.41

Key Unix I/O Design Concepts

- **Uniformity**
 - file operations, device I/O, and interprocess communication through open, read/write, close
 - Allows simple composition of programs
 - » find | grep | wc ...
- **Open before use**
 - Provides opportunity for access control and arbitration
 - Sets up the underlying machinery, i.e., data structures
- **Byte-oriented**
 - Even if blocks are transferred, addressing is in bytes
- **Kernel buffered reads**
 - Streaming and block devices looks the same
 - read blocks process, yielding processor to other task
- **Kernel buffered writes**
 - Completion of out-going transfer decoupled from the application, allowing it to continue
- **Explicit close**

9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.42

Summary

- **Process: execution environment with Restricted Rights**
 - Address Space with One or More Threads
 - Owns memory (address space)
 - Owns file descriptors, file system context, ...
 - Encapsulate one or more threads sharing process resources
- **Interrupts**
 - Hardware mechanism for regaining control from user
 - Notification that events have occurred
 - User-level equivalent: Signals
- **Native control of Process**
 - Fork, Exec, Wait, Signal
- **Basic Support for I/O**
 - Standard interface: open, read, write, seek
 - Device drivers: customized interface to hardware

9/2/15

Kubiatowicz CS162 ©UCB Fall 2015

Lec 3.43