

CSI 62  
Operating Systems and  
Systems Programming  
Lecture 5

Introduction to Networking (Finished),  
Concurrency (Processes and Threads)

February 3<sup>rd</sup>, 2016  
Prof. Anthony D. Joseph  
<http://cs162.eecs.Berkeley.edu>

Recall: Namespaces for communication over IP

- Hostname
  - [www.eecs.berkeley.edu](http://www.eecs.berkeley.edu)
- IP address
  - 128.32.244.172 (IPv4 32-bit)
  - fe80::4ad7:5ff:fecf:2607 (IPv6 128-bit)
- Port Number
  - 0-1023 are “well known” or “system” ports
    - » Superuser privileges to bind to one
  - 1024 – 49151 are “registered” ports ([registry](#))
    - » Assigned by IANA for specific services
  - 49152–65535 ( $2^{15}+2^{14}$  to  $2^{16}-1$ ) are “dynamic” or “private”
    - » Automatically allocated as “ephemeral Ports”

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.2

Recall: Use of Sockets in TCP

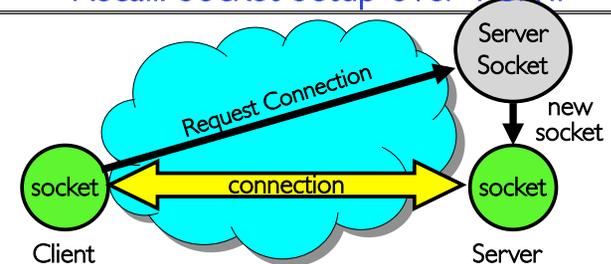
- **Socket**: an abstraction of a network I/O queue
  - Embodies one side of a communication channel
    - » Same interface regardless of location of other end
    - » Local machine (“UNIX socket”) or remote machine (“network socket”)
  - First introduced in 4.2 BSD UNIX: big innovation at time
    - » Now most operating systems provide some notion of socket
- Using Sockets for Client-Server (C/C++ interface):
  - On server: set up “server-socket”
    - » Create socket, Bind to protocol (TCP), local address, port
    - » Call `listen()`: tells server socket to accept incoming requests
    - » Perform multiple `accept()` calls on socket to accept incoming connection request
    - » Each successful `accept()` returns a new socket for a new connection; can pass this off to handler thread
  - On client:
    - » Create socket, Bind to protocol (TCP), remote address, port
    - » Perform `connect()` on socket to make connection
    - » If `connect()` successful, have socket connected to server

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.3

Recall: Socket Setup over TCP/IP



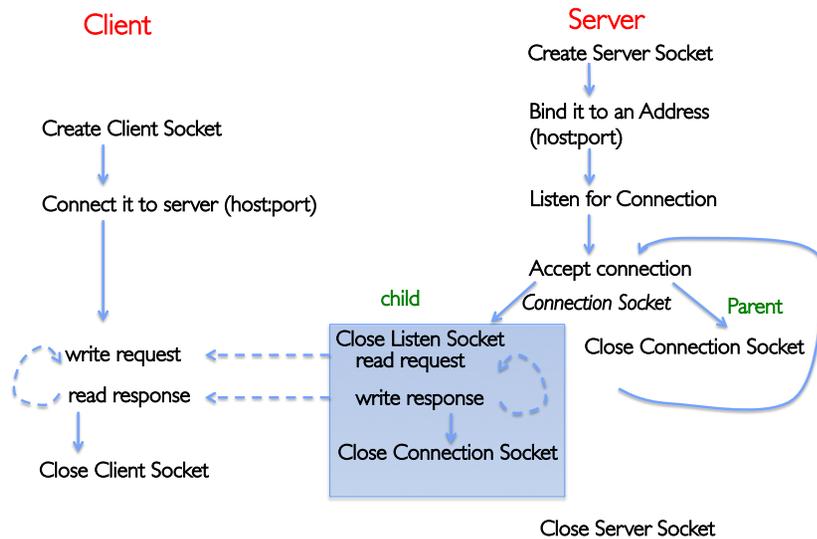
- Server Socket: Listens for new connections
  - Produces new sockets for each unique connection
- Things to remember:
  - Connection involves 5 values:  
[ Client Addr, Client Port, Server Addr, Server Port, Protocol ]
  - Often, Client Port “randomly” assigned
    - » Done by OS during client socket setup
  - Server Port often “well known”
    - » 80 (web), 443 (secure web), 25 (sendmail), etc
    - » Well-known ports from 0–1023

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.4

## Example: Server Protection and Parallelism



2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.5

## Recall: Server Protocol (v3)

```
listen(lstnsckfd, MAXQUEUE);
while (1) {
    consckfd = accept(lstnsckfd, (struct sockaddr *) &cli_addr,
                     &clilen);

    cpid = fork();           /* new process for connection */
    if (cpid > 0) {         /* parent process */
        close(consckfd);
    } else if (cpid == 0) { /* child process */
        close(lstnsckfd);   /* let go of listen socket */

        server(consckfd);

        close(consckfd);
        exit(EXIT_SUCCESS); /* exit child normally */
    }
}
close(lstnsckfd);
```

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.6

## Server Address - Itself

```
struct sockaddr_in {
    short sin_family;
    unsigned short sin_port;
    struct in_addr sin_addr;
    char sin_zero[8];
} serv_addr;

memset((char *) &serv_addr, 0, sizeof(serv_addr));
serv_addr.sin_family = AF_INET;
serv_addr.sin_addr.s_addr = INADDR_ANY;
serv_addr.sin_port = htons(portno);
```

- Simple form
- Internet Protocol
- accepting any connections on the specified port
- In “network byte ordering”

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.7

## Client: Getting the Server Address

```
struct hostent *buildServerAddr(struct sockaddr_in *serv_addr,
                               char *hostname, int portno) {
    struct hostent *server;

    /* Get host entry associated with a hostname or IP address */
    server = gethostbyname(hostname);
    if (server == NULL) {
        fprintf(stderr, "ERROR, no such host\n");
        exit(1);
    }

    /* Construct an address for remote server */
    memset((char *) serv_addr, 0, sizeof(struct sockaddr_in));
    serv_addr->sin_family = AF_INET;
    bcopy((char *) server->h_addr,
          (char *) &(serv_addr->sin_addr.s_addr), server->h_length);
    serv_addr->sin_port = htons(portno);

    return server;
}
```

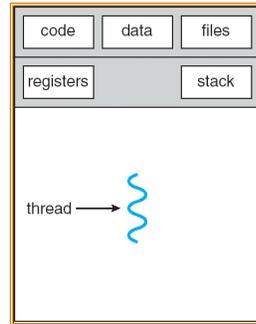
2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.8

## Recall: Traditional UNIX Process

- Process: OS abstraction of what is needed to run a single program
  - Often called a “HeavyWeight Process”
  - No concurrency in a “HeavyWeight Process”
- Two parts:
  - Sequential program execution stream [ACTIVE PART]
    - » Code executed as a sequential stream of execution (i.e., thread)
    - » Includes State of CPU registers
  - Protected resources [PASSIVE PART]:
    - » Main memory state (contents of Address Space)
    - » I/O state (i.e. file descriptors)



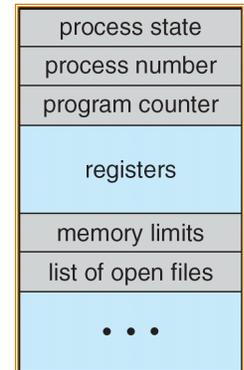
2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.9

## How do we Multiplex Processes?

- The current state of process held in a process control block (PCB):
  - This is a “snapshot” of the execution and protection environment
  - Only one PCB active at a time
- Give out CPU time to different processes (Scheduling):
  - Only one process “running” at a time
  - Give more time to important processes
- Give pieces of resources to different processes (Protection):
  - Controlled access to non-CPU resources
  - Example mechanisms:
    - » Memory Mapping: Give each process their own address space
    - » Kernel/User duality: Arbitrary multiplexing of I/O through system calls



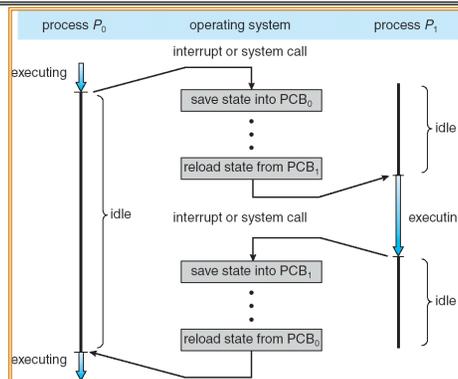
Process Control Block

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.10

## CPU Switch From Process A to Process B



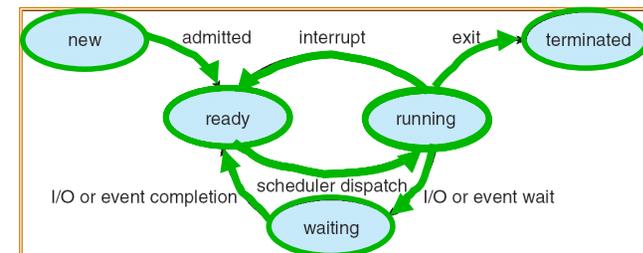
- This is also called a “context switch”
- Code executed in kernel above is overhead
  - Overhead sets minimum practical switching time
  - Less overhead with SMT/hyperthreading, but... contention for resources instead

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.11

## Lifecycle of a Process



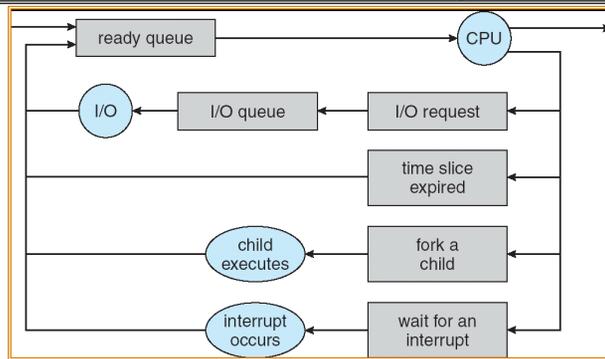
- As a process executes, it changes state:
  - **new**: The process is being created
  - **ready**: The process is waiting to run
  - **running**: Instructions are being executed
  - **waiting**: Process waiting for some event to occur
  - **terminated**: The process has finished execution

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.12

## Process Scheduling



- PCBs move from queue to queue as they change state
  - Decisions about which order to remove from queues are **Scheduling** decisions
  - Many algorithms possible (few weeks from now)

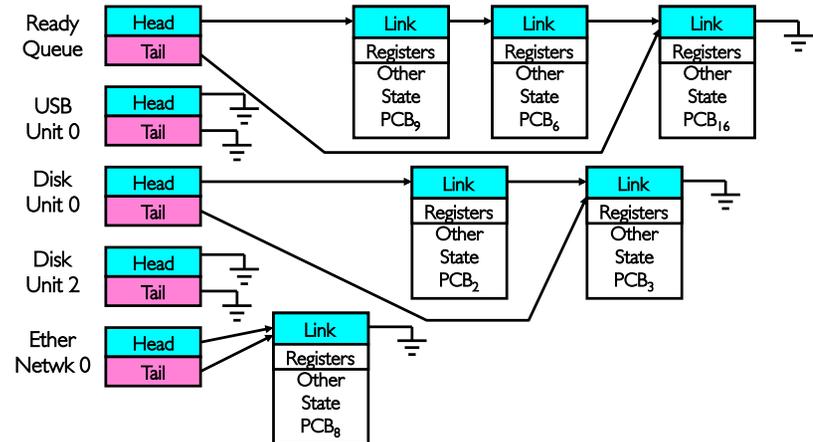
2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.13

## Ready Queue And Various I/O Device Queues

- Process not running  $\Rightarrow$  PCB is in some scheduler queue
  - Separate queue for each device/signal/condition
  - Each queue can have a different scheduler policy



2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.14

## Modern Process with Threads

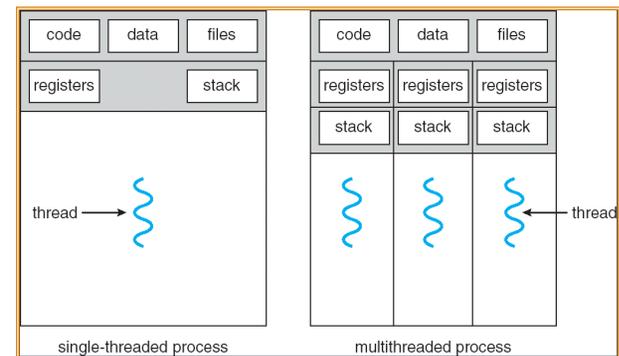
- Thread: *a sequential execution stream within process* (Sometimes called a “Lightweight process”)
  - Process still contains a single Address Space
  - No protection between threads
- Multithreading: *a single program made up of a number of different concurrent activities*
  - Sometimes called multitasking, as in Ada ...
- Why separate the concept of a thread from that of a process?
  - Discuss the “thread” part of a process (concurrency)
  - Separate from the “address space” (protection)
  - Heavyweight Process  $\equiv$  Process with one thread

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.15

## 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?

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.16

## Administrivia

---

- Group signups: 4 members/group
  - Groups need to be finalized today!
  - 36 students without groups
  - We may add a student to groups of 3
  - Sign up with the autograder
- TA Signup form
  - Form asks for 3 section options (ranked) for your group
  - Please sign up by Monday!
  - We will try to accommodate your needs, but may not be able to fill over-popular sections – give us options!
- Need to get to know your TAs
  - Consider moving out of really full sections!

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.17

BREAK

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.18

## Thread State

---

- State shared by all threads in process/address space
  - Content of memory (global variables, heap)
  - I/O state (file descriptors, network connections, etc)
- State “private” to each thread
  - Kept in TCB = Thread Control Block
  - CPU registers (including, program counter)
  - Execution stack – what is this?
- Execution Stack
  - Parameters, temporary variables
  - Return PCs are kept while called procedures are executing

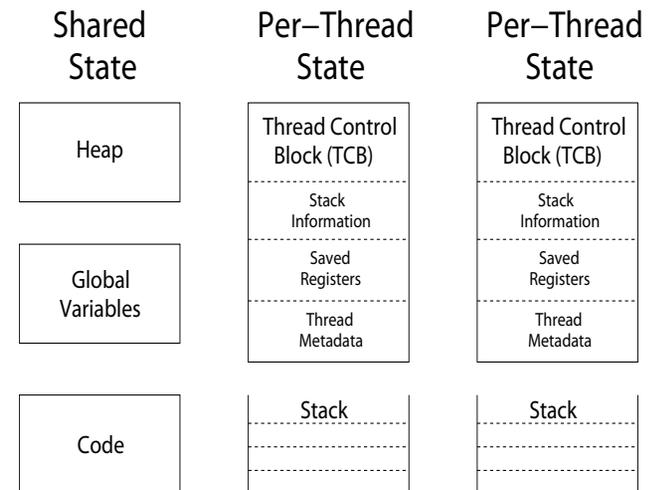
2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.19

## Shared vs. Per-Thread State

---



2/3/16

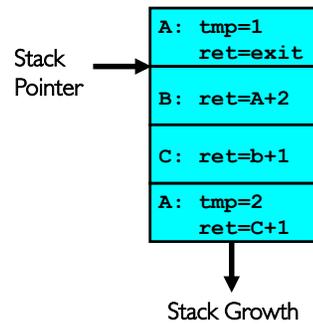
Joseph CSI 62 ©UCB Spring 2016

Lec 5.20

## Execution Stack Example

```

A(int tmp) {
  if (tmp<2)
    B();
  printf(tmp);
}
B() {
  C();
}
C() {
  A(2);
}
A(1);
    
```



- Stack holds temporary results
- Permits recursive execution
- Crucial to modern languages

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.21

## MIPS: Software conventions for Registers

0	zero constant 0	16	s0 callee saves
1	at reserved for assembler	...	(callee must save)
2	v0 expression evaluation &	23	s7
3	v1 function results	24	t8 temporary (cont'd)
4	a0 arguments	25	t9
5	a1	26	k0 reserved for OS kernel
6	a2	27	k1
7	a3	28	gp Pointer to global area
8	t0 temporary: caller saves	29	sp Stack pointer
...	(callee can clobber)	30	fp frame pointer
15	t7	31	ra Return Address (HW)

- Before calling procedure:
  - Save caller-saves regs
  - Save v0, v1
  - Save ra
- After return, assume
  - Callee-saves reg OK
  - gp,sp,fp OK (restored!)
  - Other things trashed

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.22

## Motivational Example for Threads

- Imagine the following C program:

```

main() {
  ComputePI("pi.txt");
  PrintClassList("classlist.txt");
}
    
```

- What is the behavior here?
  - Program would never print out class list
  - Why? ComputePI would never finish

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.23

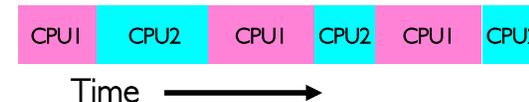
## Use of Threads

- Version of program with Threads (loose syntax):

```

main() {
  ThreadFork(ComputePI("pi.txt"));
  ThreadFork(PrintClassList("classlist.txt"));
}
    
```

- What does "ThreadFork()" do?
  - Start independent thread running given procedure
- What is the behavior here?
  - Now, you would actually see the class list
  - This *should* behave as if there are two separate CPUs



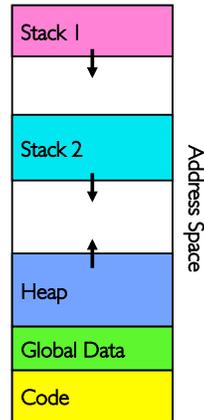
2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.24

## Memory Footprint: Two-Threads

- If we stopped this program and examined it with a debugger, we would see
  - Two sets of CPU registers
  - Two sets of Stacks
- Questions:
  - How do we position stacks relative to each other?
  - What maximum size should we choose for the stacks?
  - What happens if threads violate this?
  - How might you catch violations?



2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.25

## Actual Thread Operations

- **thread\_fork(func, args)**
  - Create a new thread to run func(args)
  - Pintos: **thread\_create**
- **thread\_yield()**
  - Relinquish processor voluntarily
  - Pintos: **thread\_yield**
- **thread\_join(thread)**
  - In parent, wait for forked thread to exit, then return
  - Pintos: **thread\_join**
- **thread\_exit**
  - Quit thread and clean up, wake up joiner if any
  - Pintos: **thread\_exit**
- **pThreads**: POSIX standard for thread programming [POSIX.1c, Threads extensions (IEEE Std 1003.1c-1995)]

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.26

## Dispatch Loop

- Conceptually, the dispatching loop of the operating system looks as follows:

```
Loop {
    RunThread();
    ChooseNextThread();
    SaveStateOfCPU(curTCB);
    LoadStateOfCPU(newTCB);
}
```

- This is an *infinite* loop
  - One could argue that this is all that the OS does
- Should we ever exit this loop???
  - When would that be?

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.27

## Running a thread

Consider first portion: **RunThread()**

- How do I run a thread?
  - Load its state (registers, PC, stack pointer) into CPU
  - Load environment (virtual memory space, etc)
  - Jump to the PC
- How does the dispatcher get control back?
  - Internal events: thread returns control voluntarily
  - External events: thread gets *preempted*

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.28

## Internal Events

- Blocking on I/O
  - The act of requesting I/O implicitly yields the CPU
- Waiting on a “signal” from other thread
  - Thread asks to wait and thus yields the CPU
- Thread executes a `yield()`
  - Thread volunteers to give up CPU

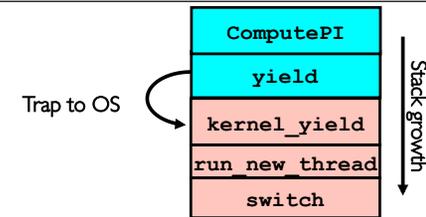
```
computePI() {
    while(TRUE) {
        ComputeNextDigit();
        yield();
    }
}
```

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.29

## Stack for Yielding Thread



- How do we run a new thread?

```
run_new_thread() {
    newThread = PickNewThread();
    switch(curThread, newThread);
    ThreadHouseKeeping(); /* Do any cleanup */
}
```

- How does dispatcher switch to a new thread?
  - Save anything next thread may trash: PC, regs, stack
  - Maintain isolation for each thread

2/3/16

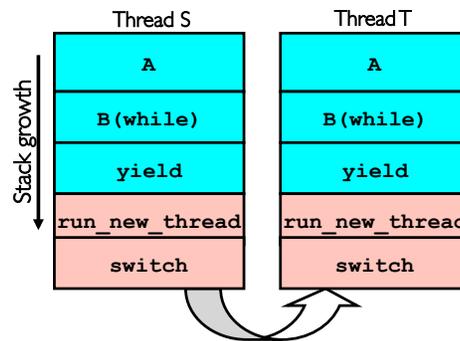
Joseph CSI 62 ©UCB Spring 2016

Lec 5.30

## What Do the Stacks Look Like?

- Consider the following code blocks:

```
proc A() {
    B();
}
proc B() {
    while(TRUE) {
        yield();
    }
}
```



- Suppose we have 2 threads:
  - Threads S and T

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.31

## Saving/Restoring state (often called “Context Switch”)

```
Switch(tCur, tNew) {
    /* Unload old thread */
    TCB[tCur].regs.r7 = CPU.r7;
    ...
    TCB[tCur].regs.r0 = CPU.r0;
    TCB[tCur].regs.sp = CPU.sp;
    TCB[tCur].regs.retpc = CPU.retpc; /*return addr*/

    /* Load and execute new thread */
    CPU.r7 = TCB[tNew].regs.r7;
    ...
    CPU.r0 = TCB[tNew].regs.r0;
    CPU.sp = TCB[tNew].regs.sp;
    CPU.retpc = TCB[tNew].regs.retpc;
    return; /* Return to CPU.retpc */
}
```

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.32

## Switch Details (continued)

- What if you make a mistake in implementing switch?
  - Suppose you forget to save/restore register 32
  - Get intermittent failures depending on when context switch occurred and whether new thread uses register 32
  - System will give wrong result without warning
- Can you devise an exhaustive test to test switch code?
  - No! Too many combinations and inter-leavings
- Cautionary tale:
  - For speed, Topaz kernel saved one instruction in switch()
  - Carefully documented! Only works As long as kernel size < 1MB
  - What happened?
    - » Time passed, People forgot
    - » Later, they added features to kernel (no one removes features!)
    - » Very weird behavior started happening
  - Moral of story: Design for simplicity

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.33

## Some Numbers

- Frequency of performing context switches: 10-100ms
  - Context switch time in Linux: 3-4  $\mu$ secs (Intel i7 & E5)
    - Thread switching faster than process switching (100 ns)
    - But switching across cores  $\sim 2x$  more expensive than within-core
  - Context switch time increases sharply with size of working set\*
    - Can increase 100x or more
- \*The working set is subset of memory used by process in a time window
- Moral: Context switching depends mostly on cache limits and the process or thread's hunger for memory

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.34

## Some Numbers

- Many process are multi-threaded, so thread context switches may be either **within-process** or **across-processes**

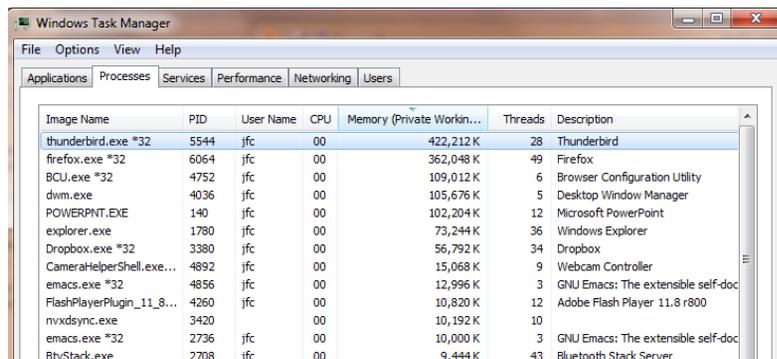


Image Name	PID	User Name	CPU	Memory (Private Workin...	Threads	Description
thunderbird.exe *32	5544	jfc	00	422,212 K	28	Thunderbird
firefox.exe *32	6064	jfc	00	362,048 K	49	Firefox
BCU.exe *32	4752	jfc	00	109,012 K	6	Browser Configuration Utility
dwm.exe	4036	jfc	00	105,676 K	5	Desktop Window Manager
POWERPNT.EXE	140	jfc	00	102,204 K	12	Microsoft PowerPoint
explorer.exe	1780	jfc	00	73,244 K	36	Windows Explorer
Dropbox.exe *32	3380	jfc	00	56,792 K	34	Dropbox
CameraHelperShell.exe...	4892	jfc	00	15,068 K	9	Webcam Controller
emacs.exe *32	4856	jfc	00	12,996 K	3	GNU Emacs: The extensible self-doc
FlashPlayerPlugin_11_8...	4260	jfc	00	10,820 K	12	Adobe Flash Player 11.8 r800
nvxdsync.exe	3420	jfc	00	10,192 K	10	
emacs.exe *32	2736	jfc	00	10,000 K	3	GNU Emacs: The extensible self-doc
BtvStack.exe	2708	jfc	00	9,444 K	43	Bluetooth Stack Server

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.35

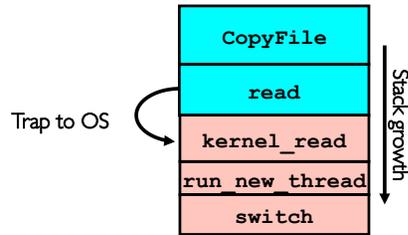
BREAK

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.36

## What happens when thread blocks on I/O?



- What happens when a thread requests a block of data from the file system?
  - User code invokes a system call
  - Read operation is initiated
  - Run new thread/switch
- Thread communication similar
  - Wait for Signal/Join
  - Networking

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.37

## External Events

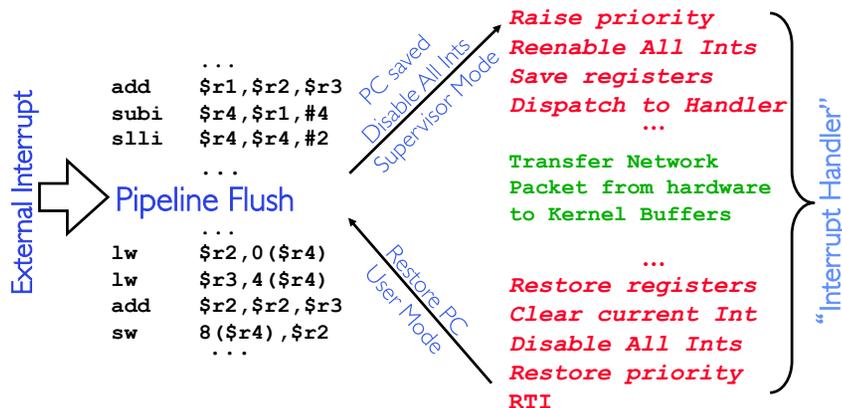
- What happens if thread never does any I/O, never waits, and never yields control?
  - Could the ComputePI program grab all resources and never release the processor?
    - » What if it didn't print to console?
  - Must find way that dispatcher can regain control!
- Answer: Utilize External Events
  - Interrupts: signals from hardware or software that stop the running code and jump to kernel
  - Timer: like an alarm clock that goes off every some many milliseconds
- If we make sure that external events occur frequently enough, can ensure dispatcher runs

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.38

## Example: Network Interrupt



- An interrupt is a hardware-invoked context switch
  - No separate step to choose what to run next
  - Always run the interrupt handler immediately

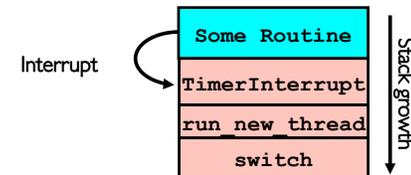
2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.39

## Use of Timer Interrupt to Return Control

- Solution to our dispatcher problem
  - Use the timer interrupt to force scheduling decisions



- Timer Interrupt routine:
 

```

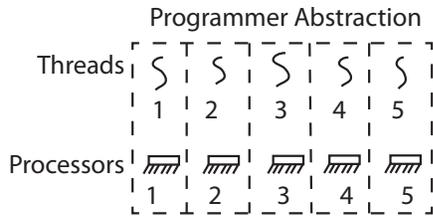
TimerInterrupt () {
    DoPeriodicHouseKeeping ();
    run_new_thread ();
}
            
```

2/3/16

Joseph CSI 62 ©UCB Spring 2016

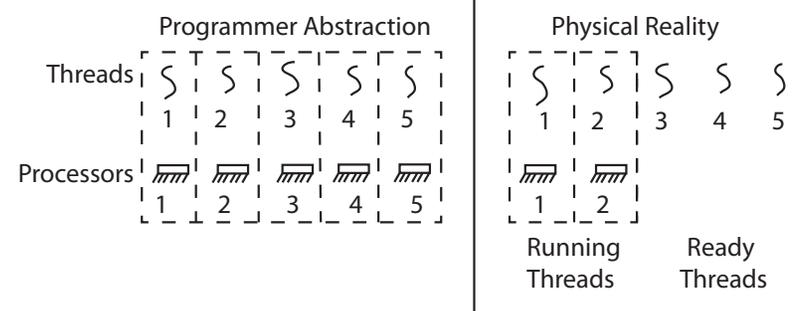
Lec 5.40

## Thread Abstraction



- Illusion: Infinite number of processors

## Thread Abstraction



- Illusion: Infinite number of processors
- Reality: Threads execute with variable speed
  - Programs must be designed to work with any schedule

## Programmer vs. Processor View

Programmer's View	Possible Execution #1
.	.
.	.
.	.
x = x + 1;	x = x + 1;
y = y + x;	y = y + x;
z = x + 5y;	z = x + 5y;
.	.
.	.
.	.

## Programmer vs. Processor View

Programmer's View	Possible Execution #1	Possible Execution #2
.	.	.
.	.	.
.	.	.
x = x + 1;	x = x + 1;	x = x + 1
y = y + x;	y = y + x;	.....
z = x + 5y;	z = x + 5y;	thread is suspended
.	.	other thread(s) run
.	.	thread is resumed
.	.	.....
		y = y + x
		z = x + 5y

## Programmer vs. Processor View

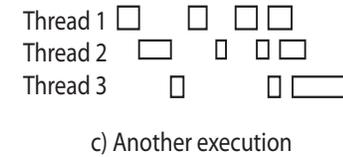
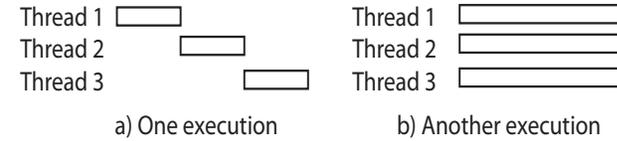
Programmer's View	Possible Execution #1	Possible Execution #2	Possible Execution #3
.	.	.	.
.	.	.	.
.	.	.	.
$x = x + 1;$	$x = x + 1;$	$x = x + 1$	$x = x + 1$
$y = y + x;$	$y = y + x;$	.....	$y = y + x$
$z = x + 5y;$	$z = x + 5y;$	thread is suspended	.....
.	.	other thread(s) run	thread is suspended
.	.	thread is resumed	other thread(s) run
.	.	.....	thread is resumed
		$y = y + x$	.....
		$z = x + 5y$	$z = x + 5y$

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.45

## Possible Executions

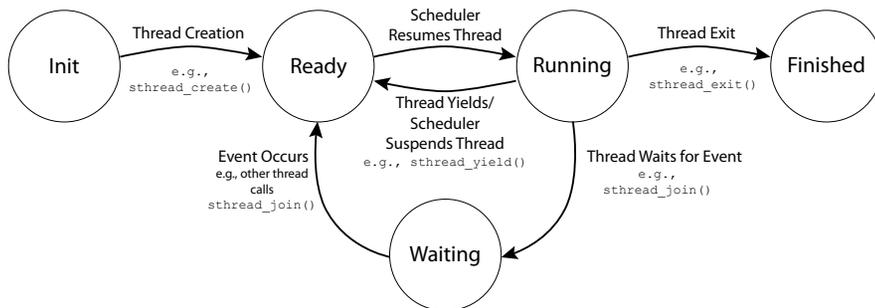


2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.46

## Thread Lifecycle



2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.47

## Summary

- Processes have two parts
  - One or more Threads (Concurrency)
  - Address Spaces (Protection)
- Concurrency accomplished by multiplexing CPU Time:
  - Unloading current thread (PC, registers)
  - Loading new thread (PC, registers)
  - Such context switching may be voluntary (`yield()`, I/O operations) or involuntary (timer, other interrupts)
- Guest lecturers next week

2/3/16

Joseph CSI 62 ©UCB Spring 2016

Lec 5.48