

# TCSS 422: OPERATING SYSTEMS

## Processes, Process API, Limited Direct Execution



Wes J. Lloyd  
School of Engineering and Technology,  
University of Washington - Tacoma

October 3, 2018

TCSS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

## FEEDBACK FROM 10/1

- How do processes start threads?
  - Done in code or by compiler?
- VM Survey – results submitted
- Assignment 0 questions
- C Tutorial – to be posted

October 3, 2018

TCSS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L3.2

# OBJECTIVES

- Chapter 4 – Processes
- Chapter 5 – Process API
- Chapter 6 – Limited Direct Execution
  
- Chapter 7 – Introduction to Scheduling
- Chapter 8 – Multi-level Feedback Queue

October 3, 2018	TCSS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma	L3.3
-----------------	---	------

# CHAPTER 4: PROCESSES

```
graph TD
    created((created)) -- admitted --> ready((ready))
    ready -- scheduler dispatch --> running((running))
    running -- interrupt --> ready
    running -- exit --> terminated((terminated))
    running -- I/O or event wait --> waiting((waiting))
    waiting -- I/O or event completion --> ready
```

October 3, 2018	TCSS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma	L3.4
-----------------	---	------

## CPU VIRTUALIZING

- How should the CPU be shared?
- Time Sharing:  
Run one process, pause it, run another
- How do we SWAP processes in and out of the CPU efficiently?
  - Goal is to minimize **overhead** of the swap

October 3, 2018	TCSS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma	L3.5
-----------------	---	------

## PROCESS

A process is a running program.

- Process comprises of:
  - Memory
    - Instructions (“the code”)
    - Data (heap)
  - Registers
    - PC: Program counter
    - Stack pointer

October 3, 2018	TCSS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma	L3.6
-----------------	---	------

## PROCESS API

- Modern OSes provide a Process API for process support
- Create
  - Create a new process
- Destroy
  - Terminate a process (ctrl-c)
- Wait
  - Wait for a process to complete/stop
- Miscellaneous Control
  - Suspend process (ctrl-z)
  - Resume process (fg, bg)
- Status
  - Obtain process statistics: (top)

October 3, 2018	TCSS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma	L3.7
-----------------	---	------

## PROCESS API: CREATE

1. Load program code (and static data) into memory
  - Program executable code (binary): loaded from disk
  - Static data: also loaded/created in address space
  - **Eager loading**: Load entire program before running
  - **Lazy loading**: Only load what is immediately needed
    - Modern OSes: Supports paging & swapping
2. Run-time stack creation
  - Stack: local variables, function params, return address(es)

October 3, 2018	TCSS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma	L3.8
-----------------	---	------

## PROCESS API: CREATE

### 3. Create program's heap memory

- For dynamically allocated data

### 4. Other initialization

- I/O Setup
  - Each process has three open file descriptors:  
Standard Input, Standard Output, Standard Error

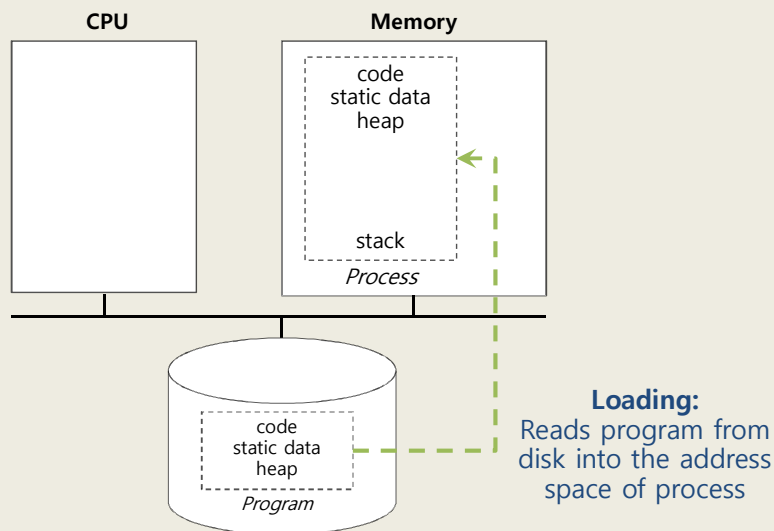
### 5. Start program running at the entry point: `main()`

- OS transfers CPU control to the new process

October 3, 2018

TCSS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L3.9



October 3, 2018

TCSS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L3.10

## PROCESS STATES

- **RUNNING**
  - Currently executing instructions
  
- **READY**
  - Process is ready to run, but has been preempted
  - CPU is presently allocated for other tasks
  
- **BLOCKED**
  - Process is **not** ready to run. It is waiting for another event to complete:
    - Process has already been initialized and run for awhile
    - Is now waiting on I/O from disk(s) or other devices

October 3, 2018	TCSS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma	L3.11
-----------------	---	-------

## PROCESS STATE TRANSITIONS

```
graph TD; Running((Running)) -- "Descheduled" --> Ready((Ready)); Ready -- "Scheduled" --> Running; Running -- "I/O: initiate" --> Blocked((Blocked)); Blocked -- "I/O: done" --> Ready;
```

The diagram illustrates the transitions between three process states: Running, Ready, and Blocked. Running and Ready states are connected by bidirectional arrows labeled 'Descheduled' and 'Scheduled'. Running transitions to Blocked via 'I/O: initiate', and Blocked transitions back to Ready via 'I/O: done'.

October 3, 2018	TCSS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma	L3.12
-----------------	---	-------

## PROCESS DATA STRUCTURES

- OS provides data structures to track process information
  - Process list
    - Process Data
    - State of process: Ready, Blocked, Running
  - Register context
- PCB (Process Control Block)
  - A C-structure that contains information about each process

October 3, 2018

TCSS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L3.13

## XV6 KERNEL DATA STRUCTURES

- Process data structure - textbook: xv6  
Pedagogical implementation of Linux

```
// the information xv6 tracks about each process
// including its register context and state
struct proc {
    char *mem;                // Start of process memory
    uint sz;                  // Size of process memory
    char *kstack;            // Bottom of kernel stack
                             // for this process
    enum proc_state state;   // Process state
    int pid;                 // Process ID
    struct proc *parent;     // Parent process
    void *chan;              // If non-zero, sleeping on chan
    int killed;              // If non-zero, have been killed
    struct file *ofile[NOFILE]; // Open files
    struct inode *cwd;       // Current directory
    struct context context;  // Switch here to run process
    struct trapframe *tf;   // Trap frame for the
                             // current interrupt
};
```

October 3, 2018

TCSS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L3.14

## XV6 KERNEL DATA STRUCTURES - 2

### ■ CPU register context data structure - textbook: xv6

```
// the registers xv6 will save and restore
// to stop and subsequently restart a process
struct context {
    int eip; // Index pointer register
    int esp; // Stack pointer register
    int ebx; // Called the base register
    int ecx; // Called the counter register
    int edx; // Called the data register
    int esi; // Source index register
    int edi; // Destination index register
    int ebp; // Stack base pointer register
};

// the different states a process can be in
enum proc_state { UNUSED, EMBRYO, SLEEPING,
                  RUNNABLE, RUNNING, ZOMBIE };
```

October 3, 2018

TCCS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L3.15

## LINUX: STRUCTURES

- **struct task\_struct**, equivalent to struct proc
  - Provides process description
  - Large: 10,000+ bytes
  - /usr/src/linux-headers-{kernel version}/include/linux/sched.h
    - Starts at 1391
- **struct thread\_info**, provides “context”
  - thread\_info.h is at:  
/usr/src/linux-headers-{kernel version}/arch/x86/include/asm/

October 3, 2018

TCCS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L3.16



## LINUX: THREAD\_INFO

### Linux thread data structure thread\_info

```
struct thread_info {
    struct task_struct      *task;           /* main task structure */
    struct exec_domain     *exec_domain;    /* execution domain */
    __u32                   flags;         /* low level flags */
    __u32                   status;        /* thread synchronous flags */
    __u32                   cpu;           /* current CPU */
    int                     preempt_count; /* 0 => preemptable,
                                           <0 => BUG */

    mm_segment_t           addr_limit;
    struct restart_block   restart_block;
    void __user            *sysenter_return;

#ifdef CONFIG_X86_32
    unsigned long          previous_esp;    /* ESP of the previous stack in
                                           case of nested (IRQ) stacks
                                           */
    __u8                   supervisor_stack[0];
#endif
    int                     uaccess_err;
};
```

October 3, 2018

TCCS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L3.17

## LINUX STRUCTURES - 2

### List of Linux data structures:

<http://www.tldp.org/LDP/tlk/ds/ds.html>

### Description of process data structures:

<http://www.makelinux.net/books/lkd2/ch03lev1sec1>  
2<sup>nd</sup> edition is online (dated from 2005):

**Linux Kernel Development, 2<sup>nd</sup> edition**

Robert Love

Sams Publishing

October 3, 2018

TCCS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L3.18

**W** When a process is in this state, it is advantageous for the Operating System to perform a **CONTEXT SWITCH** to perform other work

RUNNING    READY    BLOCKED    All of the above    None of the above


October 3, 2018    TCSS422: Operating Systems [Fall 2018]    School of Engineering and Technology, University of Washington - Tacoma    Total Re. L3.19

**QUESTION: WHEN TO CONTEXT SWITCH**

- When a process is in this state, it is advantageous for the Operating System to perform a **CONTEXT SWITCH** to perform other work:
  - (a) **RUNNING**
  - (b) **READY**
  - (c) **BLOCKED**
  - (d) **All of the above**
  - (e) **None of the above**

October 3, 2018    TCSS422: Operating Systems [Fall 2018]    School of Engineering and Technology, University of Washington - Tacoma    L3.20


# CHAPTER 5: C PROCESS API



October 3, 2018 TCSS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma L3.21

## fork()

- Creates a new process - think of “a fork in the road”
- “Parent” process is the original
- Creates “child” process of the program from the current execution point
- Book says “pretty odd”
- Creates a **duplicate** program instance (these are processes!)
- **Copy of**
  - Address space (memory)
  - Register
  - Program Counter (PC)
- Fork returns
  - child PID to parent
  - 0 to child



October 3, 2018 TCSS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma L3.22

## FORK EXAMPLE

### ■ p1.c

```
#include <stdio.h>
#include <stdlib.h>
#include <unistd.h>

int main(int argc, char *argv[]){
    printf("hello world (pid:%d)\n", (int) getpid());
    int rc = fork();
    if (rc < 0) { // fork failed; exit
        fprintf(stderr, "fork failed\n");
        exit(1);
    } else if (rc == 0) { // child (new process)
        printf("hello, I am child (pid:%d)\n", (int) getpid());
    } else { // parent goes down this path (main)
        printf("hello, I am parent of %d (pid:%d)\n",
            rc, (int) getpid());
    }
    return 0;
}
```



October 3, 2018

TCCS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L3.23

## FORK EXAMPLE - 2

### ■ Non deterministic ordering of execution

```
prompt> ./p1
hello world (pid:29146)
hello, I am parent of 29147 (pid:29146)
hello, I am child (pid:29147)
prompt>
```

or

```
prompt> ./p1
hello world (pid:29146)
hello, I am child (pid:29147)
hello, I am parent of 29147 (pid:29146)
prompt>
```

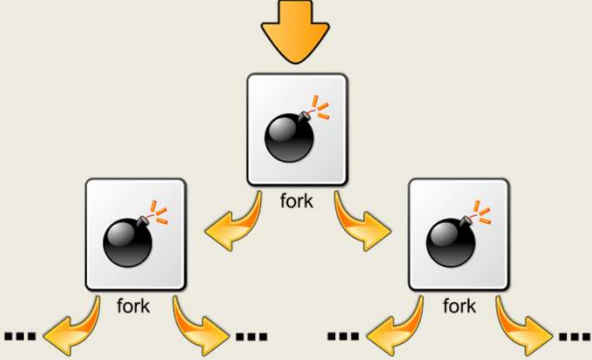
### ■ CPU scheduler determines which to run first

October 3, 2018

TCCS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L3.24

## :(){ :|: & }::




The diagram illustrates a recursive process where a parent process forks into two child processes, which in turn fork into two more child processes each, creating a binary tree structure. Each process is represented by a square box containing a bomb icon. Yellow arrows labeled 'fork' show the flow from parent to child. The process starts with a single parent at the top, which forks into two children. Each of these children forks into two more children, and so on, with ellipses indicating further branching.

October 3, 2018	TCSS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma	L3.25
-----------------	---	-------

## wait()

- `wait()`, `waitpid()`
- Called by parent process
- Waits for a child process to finish executing
- Not a `sleep()` function
- Provides some ordering to multi-process execution



A photograph of a silver alarm clock with a woman's face partially visible behind it, looking at the camera. The clock shows a time around 10:10.

October 3, 2018	TCSS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma	L3.26
-----------------	---	-------

## FORK WITH WAIT

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

int main(int argc, char *argv[]){
    printf("hello world (pid:%d)\n", (int) getpid());
    int rc = fork();
    if (rc < 0) { // fork failed; exit
        fprintf(stderr, "fork failed\n");
        exit(1);
    } else if (rc == 0) { // child (new process)
        printf("hello, I am child (pid:%d)\n", (int) getpid());
    } else { // parent goes down this path (main)
        int wc = wait(NULL);
        printf("hello, I am parent of %d (wc:%d) (pid:%d)\n",
            rc, wc, (int) getpid());
    }
    return 0;
}
```

October 3, 2018

TCSS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L3.27

## FORK WITH WAIT - 2

### ■ Deterministic ordering of execution

```
prompt> ./p2
hello world (pid:29266)
hello, I am child (pid:29267)
hello, I am parent of 29267 (wc:29267) (pid:29266)
prompt>
```

October 3, 2018

TCSS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L3.28

# FORK EXAMPLE

- Linux example

October 3, 2018	TCSS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma	L3.29
-----------------	---	-------

# exec()


- Supports running an external program
- 6 types: `execl()`, `execlp()`, `execle()`, `execv()`, `execvp()`, `execvpe()`
- execl(), execlp(), execle(): const char \*arg**  
Command arguments provided as **LIST of pointers** to strings provided as arguments... (`arg0`, `arg1`, .. `argn`) (terminated by a null pointer)
- execv(), execvp(), execvpe()**  
Command arguments provided as an **ARRAY of pointers** to strings as arguments

Strings are null-terminated  
First argument is name of file being executed

October 3, 2018	TCSS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma	L3.30
-----------------	---	-------

## EXEC() - 2

- Common use case:
  - Write a new program which wraps a legacy one
  - Provide a new interface to an old system: Web services
  - Legacy program thought of as a “black box”
- May not want to know what is inside the black box... 😊
- FORTRAN ???



October 3, 2018

TCSS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L3.31

## EXEC EXAMPLE

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

int main(int argc, char *argv[]){
    printf("hello world (pid:%d)\n", (int) getpid());
    int rc = fork();
    if (rc < 0) {
        // fork failed; exit
        fprintf(stderr, "fork failed\n");
        exit(1);
    } else if (rc == 0) {
        // child (new process)
        printf("hello, I am child (pid:%d)\n", (int) getpid());
        char *myargs[3];
        myargs[0] = strdup("wc");
        myargs[1] = strdup("p3.c");
        myargs[2] = NULL;
        ...
    }
```

October 3, 2018

TCSS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L3.32



## EXEC EXAMPLE - 2

```
...  
➔ execvp(myargs[0], myargs); // runs word count  
printf("this shouldn't print out");  
} else { // parent goes down this path (main)  
int wc = wait(NULL);  
printf("hello, I am parent of %d (wc:%d) (pid:%d)\n",  
rc, wc, (int) getpid());  
}  
return 0;  
}
```

```
prompt> ./p3  
hello world (pid:29383)  
hello, I am child (pid:29384)  
29 107 1030 p3.c  
hello, I am parent of 29384 (wc:29384) (pid:29383)  
prompt>
```

October 3, 2018

TCSS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L3.33

## EXEC WITH FILE REDIRECTION (OUTPUT)

```
#include <stdio.h>  
#include <stdlib.h>  
#include <unistd.h>  
#include <string.h>  
#include <fcntl.h>  
#include <sys/wait.h>  
  
int  
main(int argc, char *argv[]){  
int rc = fork();  
if (rc < 0) { // fork failed; exit  
fprintf(stderr, "fork failed\n");  
exit(1);  
} else if (rc == 0) { // child: redirect standard output to a file  
➔ close(STDOUT_FILENO);  
open("./p4.output", O_CREAT|O_WRONLY|O_TRUNC, S_IRWXU);  
...  
}
```

October 3, 2018

TCSS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L3.34

## FILE MODE BITS

```
S_IRWXU
read, write, execute/search by owner
S_IRUSR
read permission, owner
S_IWUSR
write permission, owner
S_IXUSR
execute/search permission, owner
S_IRWXG
read, write, execute/search by group
S_IRGRP
read permission, group
S_IWGRP
write permission, group
S_IXGRP
execute/search permission, group
S_IRWXO
read, write, execute/search by others
S_IROTH
read permission, others
S_IWOTH
write permission, others
```

October 3, 2018	TCCS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma	L3.35
-----------------	---	-------

## EXEC W/ FILE REDIRECTION (OUTPUT) - 2

```
...
// now exec "wc"...
char *myargs[3];
myargs[0] = strdup("wc");           // program: "wc" (word count)
myargs[1] = strdup("p4.c");         // argument: file to count
myargs[2] = NULL;                   // marks end of array
execvp(myargs[0], myargs);         // runs word count
} else {                             // parent goes down this path (main)
    int wc = wait(NULL);
}
return 0;
}
```

```
prompt> ./p4
prompt> cat p4.output
32 109 846 p4.c
prompt>
```

October 3, 2018	TCCS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma	L3.36
-----------------	---	-------

**W** Which Process API call is used to launch a different program from the current program?

Fork()   Exec()   Wait()   None of the above   All of the above

Start the presentation to see live content. Still no live content? Install the app or get help at [PollEv.com/app](http://PollEv.com/app)   Total Results

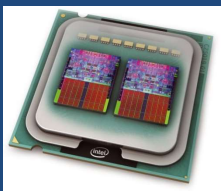
**QUESTION: PROCESS API**

- Which Process API call is used to launch a different program from the current program?
- (a) Fork()
- (b) Exec()
- (c) Wait()
- (d) None of the above
- (e) All of the above

October 3, 2018	TCSS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma	L3.38
-----------------	---	-------

# CH. 6: LIMITED DIRECT EXECUTION

October 3, 2018
TCSS422: Operating Systems [Fall 2018]  
 School of Engineering and Technology, University of Washington - Tacoma
L3.39



# VIRTUALIZING THE CPU

- How does the CPU support running so many jobs simultaneously?
- Time Sharing
- Tradeoffs:
  - Performance
    - Excessive overhead
  - Control
    - Fairness
    - Security
- Both HW and OS support is used

```

top - 18:25:07 up 430 days, 1:03, 3 users, load average: 0.32, 0.28, 0.29
task: 654 total, 1 running, 653 sleeping, 0 stopped, 0 dead
(Cpu): 7.8kusr, 0.5kty, 0.0kni, 91.8kld, 0.0kpa, 0.0ksw, 0.7kai, 0.0kst
Mem: 7423772k total, 5946428k used, 548550k free, 561316k buffers
Swap: 2183572k total, 72252k used, 2176352k free, 5528350k cached

PID USER   PR  NI  VIRT  RES  SHR  S#    CPU MEM%   FILE COMMAND
1526 root    20   0 6036 276 20m  S 152 0.0 4.32 91 postgres
3078 root    20   0 6036 276 20m  S 152 0.0 4.32 91 postgres
3424 root    20   0 6036 228 18m  S 8.6 0.0 0.26 46 postgres
4460 ericson 20   0 15432 1712 208  S 0.7 0.0 0.00 52 top
6280 root    20   0 6036 18m  S 0.7 0.0 0.07 74 postgres
7440 root    20   0 6036 276 20m  S 0.7 0.0 4.21 48 postgres
8520 root    20   0 6036 284 20m  S 0.7 0.0 6.49 42 postgres
10828 root   20   0 6036 20m  S 0.7 0.0 0.02 78 postgres
11912 root   20   0 6036 20m  S 0.7 0.0 0.10 58 postgres
15152 root   20   0 6036 18m  S 0.7 0.0 0.06 78 postgres
17176 root   20   0 6036 20m  S 0.7 0.0 0.00 78 postgres
1904 root    20   0 6036 18m  S 0.7 0.0 0.04 62 postgres
1904 root    20   0 6036 276 20m  S 0.3 0.0 0.28 28 postgres
3504 root    20   0 6036 194 8420  S 0.3 0.0 0.01 16 postgres
4112 root    20   0 6036 276 20m  S 0.3 0.0 4.31 34 postgres
7020 root    20   0 6036 276 20m  S 0.3 0.0 4.31 34 postgres
7080 root    20   0 6036 276 20m  S 0.3 0.0 4.25 40 postgres
8524 root    20   0 6036 276 20m  S 0.3 0.0 4.25 40 postgres
8524 root    20   0 6036 276 20m  S 0.3 0.0 3.52 32 postgres
14280 root   20   0 6036 276 20m  S 0.2 0.0 1.39 20 postgres
15752 root   20   0 6036 276 20m  S 0.3 0.0 1.28 30 postgres
16028 root   20   0 6036 18m  S 0.3 0.0 0.02 46 postgres
16408 root   20   0 6036 18m  S 0.3 0.0 0.02 51 postgres
16528 root   20   0 20916 1404 10m  S 0.2 0.0 1.95 40 postgres
21780 root   20   0 24.4g 5628 10m  S 0.3 0.0 7.39 58 32 java
30704 root   20   0 6036 276 20m  S 0.2 0.0 0.02 46 postgres
31536 root   20   0 6036 248 20m  S 0.3 0.0 5.03 32 postgres
1 root    20   0 0 0 0  S 0 0.0 0.00 02 kthreadd
2 root    20   0 0 0 0  S 0 0.0 0.00 02 kthreadd
3 root    81   0 0 0 0  S 0 0.0 0.22 09 48 migration/0
4 root    20   0 0 0 0  S 0 0.0 0.00 00 kserfsync/0
5 root    81   0 0 0 0  S 0 0.0 0.00 00 stopper/0
6 root    81   0 0 0 0  S 0 0.0 0.00 00 stopper/1
7 root    81   0 0 0 0  S 0 0.0 0.22 31 migration/1
8 root    81   0 0 0 0  S 0 0.0 0.00 00 stopper/1
9 root    20   0 0 0 0  S 0 0.0 0.22 31 migration/1
10 root   81   0 0 0 0  S 0 0.0 0.28 40 migration/2
11 root   81   0 0 0 0  S 0 0.0 0.13 03 04 migration/2
12 root   81   0 0 0 0  S 0 0.0 0.00 00 stopper/2
    
```

October 3, 2018

TCSS422: Operating Systems [Fall 2018]  
 School of Engineering and Technology, University of Washington - Tacoma

L3.40

## COMPUTER BOOT SEQUENCE: OS WITH DIRECT EXECUTION

■ What if programs could directly control the CPU / system?

OS	Program
1. Create entry for process list 2. Allocate memory for program 3. Load program into memory 4. Set up stack with <code>argc / argv</code> 5. Clear registers 6. Execute call <code>main()</code>	7. Run <code>main()</code> 8. Execute <code>return from main()</code>
9. Free memory of process 10. Remove from process list	

October 3, 2018

TCSS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L3.41

## COMPUTER BOOT SEQUENCE: OS WITH DIRECT EXECUTION

■ What if programs could directly control the CPU / system?

OS	Program
1. Create entry for process list 2. Allocate memory for	
Without <i>limits</i> on running programs, the OS wouldn't be in control of anything and would <b>"just be a library"</b>	
5. Clear registers 6. Execute call <code>main()</code>	7. Run <code>main()</code> 8. Execute <code>return from main()</code>
9. Free memory of process 10. Remove from process list	

October 3, 2018

TCSS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L3.42

## DIRECT EXECUTION - 2

- **With direct execution:**

How does the OS stop a program from running, and switch to another to support **time sharing**?

How do programs share disks and perform I/O if they are given direct control? Do they know about each other?

With direct execution, how can dynamic memory structures such as linked lists grow over time?

October 3, 2018

TCSS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L3.43

## CONTROL TRADEOFF

- **Too little control:**

- No security
- No time sharing

- **Too much control:**

- Too much OS overhead
- Poor performance for compute & I/O
- Complex APIs (system calls), difficult to use

October 3, 2018

TCSS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L3.44

## CONTEXT SWITCHING OVERHEAD

### Context Switching

Multitasking

vs. Multitasking with context switching

Sequential

Total cost of context switching

October 3, 2018	TCSS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma	L3.45
-----------------	---	-------

## LIMITED DIRECT EXECUTION

- OS implements LDE to support time/resource sharing
- Limited direct execution means “only limited” processes can execute **DIRECTLY** on the CPU in ***trusted*** mode
- **TRUSTED** means the process is trusted, and it can do anything... (e.g. it is a system / kernel level process)
- Enabled by ***protected (safe) control transfer***
- CPU supported context switch
- Provides data isolation

October 3, 2018	TCSS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma	L3.46
-----------------	---	-------

## CPU MODES

- Utilize CPU Privilege Rings (Intel x86)
  - rings 0 (kernel), 1 (VM kernel), 2 (unused), 3 (user)

access ←———— no access

- **User mode:**  
Application is running, but w/o direct I/O access
- **Kernel mode:**  
OS kernel is running performing restricted operations

October 3, 2018	TCSS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma	L3.47
-----------------	---	-------

## CPU MODES

- **User mode: ring 3 - untrusted**
  - Some instructions and registers are disabled by the CPU
  - Exception registers
  - HALT instruction
  - MMU instructions
  - OS memory access
  - I/O device access
- **Kernel mode: ring 0 – trusted**
  - All instructions and registers enabled

October 3, 2018	TCSS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma	L3.48
-----------------	---	-------



## SYSTEM CALLS

- Implement restricted “OS” operations
- Kernel exposes key functions through an API:
  - Device I/O (e.g. file I/O)
  - Task swapping: context switching between processes
  - Memory management/allocation: malloc()
  - Creating/destroying processes


October 3, 2018

TCSS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L3.49

## TRAPS: SYSTEM CALLS, EXCEPTIONS, INTERRUPTS

- Trap: any transfer to kernel mode
- Three kinds of traps
  - **System call:** (planned) user → kernel
    - SYSCALL for I/O, etc.
  - **Exception:** (error) user → kernel
    - Div by zero, page fault, page protection error
  - **Interrupt:** (event) user → kernel
    - Non-maskable vs. maskable
    - Keyboard event, network packet arrival, timer ticks
    - Memory parity error (ECC), hard drive failure



October 3, 2018

TCSS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L3.50

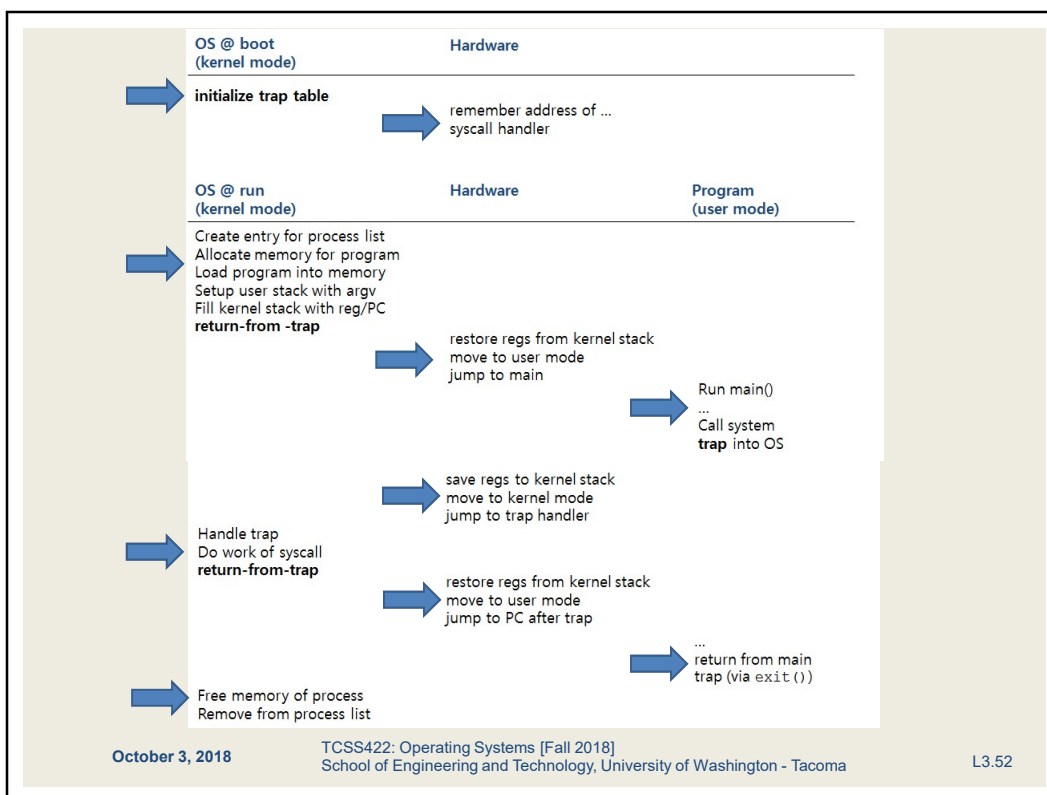
## EXCEPTION TYPES

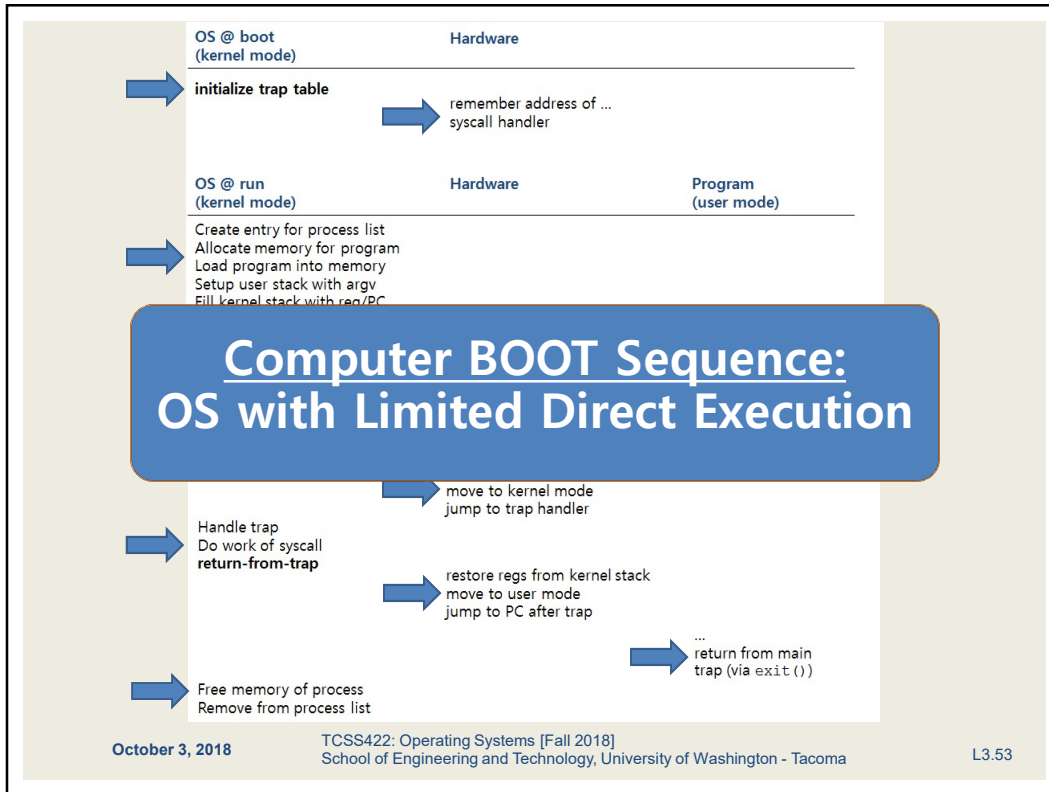
Exception type	Synchronous vs. asynchronous	User request vs. coerced	User maskable vs. nonmaskable	Within vs. between instructions	Resume vs. terminate
I/O device request	Asynchronous	Coerced	Nonmaskable	Between	Resume
Invoke operating system	Synchronous	User request	Nonmaskable	Between	Resume
Tracing instruction execution	Synchronous	User request	User maskable	Between	Resume
Breakpoint	Synchronous	User request	User maskable	Between	Resume
Integer arithmetic overflow	Synchronous	Coerced	User maskable	Within	Resume
Floating-point arithmetic overflow or underflow	Synchronous	Coerced	User maskable	Within	Resume
Page fault	Synchronous	Coerced	Nonmaskable	Within	Resume
Misaligned memory accesses	Synchronous	Coerced	User maskable	Within	Resume
Memory protection violation	Synchronous	Coerced	Nonmaskable	Within	Resume
Using undefined instruction	Synchronous	Coerced	Nonmaskable	Within	Terminate
Hardware malfunction	Asynchronous	Coerced	Nonmaskable	Within	Terminate
Power failure	Asynchronous	Coerced	Nonmaskable	Within	Terminate

October 3, 2018

TCCS422: Operating Systems [Fall 2018]  
School of Engineering and Technology, University of Washington - Tacoma

L3.51





## MULTITASKING

- How/when should the OS regain control of the CPU to switch between processes?
- Cooperative multitasking (mostly pre 32-bit)
  - < Windows 95, Mac OSX
  - Opportunistic: running programs must give up control
    - User programs must call a special **yield** system call
    - When performing I/O
    - Illegal operations
  - (POLLEV)  
 What problems could you for see with this approach?

October 3, 2018	TCCS422: Operating Systems [Fall 2018] School of Engineering and Technology, University of Washington - Tacoma	L3.54
-----------------	---	-------

